



A GPU accelerated aggregation algebraic multigrid method



Rajesh Gandham^{a,*}, Kenneth Esler^b, Yongpeng Zhang^b

^a Department of Computational and Applied Mathematics, Rice University, United States

^b Stone Ridge Technology, United States

ARTICLE INFO

Article history:

Received 6 November 2013

Received in revised form 27 July 2014

Accepted 18 August 2014

Available online 15 September 2014

Keywords:

Linear system of equations

Elliptic partial differential equations

Algebraic multigrid

Aggregation

GPGPU

CUDA

ABSTRACT

We present an efficient, robust and fully GPU-accelerated aggregation-based algebraic multigrid preconditioning technique for the solution of large sparse linear systems. These linear systems arise from the discretization of elliptic PDEs. The method involves two stages, setup and solve. In the setup stage, hierarchical coarse grids are constructed through aggregation of the fine grid nodes. These aggregations are obtained using a set of maximal independent nodes from the fine grid nodes. We use a “fine-grain” parallel algorithm for finding a maximal independent set from a graph of strong negative connections. The aggregations are combined with a piece-wise constant (unsmooth) interpolation from the coarse grid solution to the fine grid solution, ensuring low setup and interpolation cost. The grid independent convergence is achieved by using recursive Krylov iterations (K-cycles) in the solve stage. An efficient combination of K-cycles and standard multigrid V-cycles is used as a preconditioner for Krylov iterative solvers such as generalized minimal residual and conjugate gradient. We compare the solver performance with other solvers based on smooth aggregation and classical algebraic multigrid methods.

© 2014 Elsevier Ltd. All rights reserved.

1. Introduction

In the numerical simulation of various physical phenomena, discretizations of partial differential equations result in very large sparse linear system of equations. Computing solutions of these linear systems is a crucial component of the overall simulation. Consequently there has been increasing interest in developing fast linear solvers over the last two decades. Traditionally, these linear solvers are accelerated by using parallel programming techniques on clusters of CPUs. However, hardware architectures such as the one found in the graphics processing unit (GPU) provide efficient, inexpensive alternatives for these computations. The recent developments in general purpose GPUs (GPGPUs) for scientific computations make it feasible to accelerate the numerical simulations by an order of magnitude. This can be achieved by significantly modifying the traditional algorithms to suit the massively parallel hardware architecture of GPUs.

Multigrid methods are among the most efficient and popular solution techniques for solving the linear systems arising from the discretization of elliptic PDEs. Multigrid methods typically fall into one of two classes: geometric multigrid and algebraic multigrid (AMG). Geometric multigrid methods require prior knowledge of the underlying discretization and grid hierarchy, where as algebraic multigrid methods only require the entries of matrix. Furthermore, the algebraic multigrid coarsening process automatically ensures a reduction of the error components that are not reduced by smoothing [1]. Algebraic multigrid requires little or no tuning for various applications and hence make it very easy to use in the form of an efficient “black-box” solver.

* Corresponding author.

E-mail addresses: Rajesh.Gandham@rice.edu, rajeshgandham@gmail.com (R. Gandham), kesler@stoneridgetechnology.com (K. Esler), ypzhang@stoneridgetechnology.com (Y. Zhang).

<http://dx.doi.org/10.1016/j.camwa.2014.08.022>

0898-1221/© 2014 Elsevier Ltd. All rights reserved.




Fig. 1. Example of aggregation. Roman numerals indicate the aggregation index.

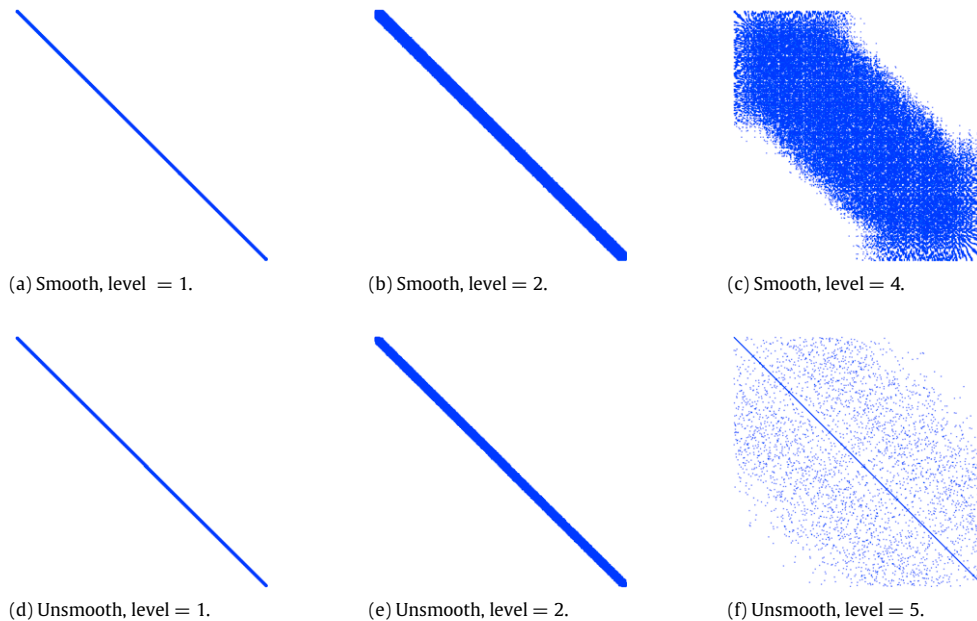


Fig. 2. Comparison of sparsity pattern of CUSP smooth aggregation (top) with GAMPACK unsmooth aggregation (bottom) for 2D anisotropic Poisson problem on a 1000×1000 grid.

Algebraic multigrid methods involve the construction of a hierarchy of matrices (or grids) using the entries of the matrix. At each level in the hierarchy, smoothing is performed to remove the high frequency errors. The low frequency errors are removed by recursively restricting the residual to coarse grids and performing smoothing on the coarse grids. The coarse grid solutions are then interpolated from the coarse grid to a fine grid using an interpolation operator.

Classical algebraic multigrid methods obtain the hierarchical grids by partitioning the nodes into coarse and fine grid nodes. The coarse grid nodes form a coarse level, and an interpolation operator is defined via a weighted sum of the coarse grid nodes. The restriction operator in general is the transpose of the interpolation operator [2].

In contrast, aggregation algebraic multigrid methods obtain the hierarchical grids by aggregating few fine grid nodes to form a coarse grid node (see Fig. 1). An interpolation is defined via piecewise constant interpolation from coarse grid node to a fine grid node. This results in very sparse interpolation and restriction operations compared to classical AMG. The interpolation matrix has exactly one nonzero entry per row, which reduces memory requirements and improves efficiency of the interpolation operation. However, these aggregation schemes are not popular since they do not provide grid independent convergence. The convergence is often improved by smooth interpolation (“smooth aggregation”) [3], but smooth aggregation typically produces a hierarchy with very dense coarse level matrices (see Fig. 2), leading to expensive matrix–vector product computations and increasing the memory requirements. Furthermore, unlike classical algebraic multigrid, smooth aggregation multigrid is not robust for various applications.

The opposite view point to smooth aggregation, presented in [4] uses unsmooth aggregation along with K-cycles as preconditioning for iterative Krylov methods, provide an efficient alternative to smooth aggregation schemes. In this work, we accelerate these unsmooth aggregation methods using GPUs and discuss the algorithmic changes considered for the acceleration, both mathematical and computational. We also discuss the advantages in using these schemes where we can efficiently reuse the sparsity structure of the hierarchical matrices.

Some of the initial works for aggregation AMG focused on accelerating only the solve phase [5]. As a consequence, the setup phase becomes a bottleneck due to the serial nature of algorithms [4] for forming the aggregates. There has been development in accelerating the setup phase for classical AMG [6] and for smooth aggregation [7]. We build on these ideas

and develop fully accelerated unsmooth aggregation AMG as a part of GAMPACK, a GPU accelerated algebraic multigrid package.

For forming the aggregates on GPUs, we follow a fine-grain parallel maximal independent set (MIS) algorithm proposed in [7] with some modifications described in Algorithm 2. With this algorithm on a GPU, we see that the setup is almost three times faster than classical AMG setup.

In Section 2, we describe the algorithms for the setup and solution phase of unsmooth aggregation AMG. In Section 3, we briefly explain GPU acceleration of the algorithms described. In Section 4, we provide numerical results for various classes of problems and compare the performance with GAMPACK classical AMG and CUSP smooth aggregation.

2. Aggregation AMG

Aggregation AMG is used as a preconditioner for iterative Krylov solvers such as generalized minimal residual (GMRES) and conjugate gradient (CG). Consider a linear system of equations,

$$Ax = b, \tag{1}$$

where $A \in \mathbb{R}^{N \times N}$ is a coefficient matrix, $b \in \mathbb{R}^{N \times 1}$ is a right hand side vector, $x \in \mathbb{R}^{N \times 1}$ is the solution vector and N is the number of unknowns. The coarse grid matrices are defined recursively, consider that $A_0 = A$ and $n_{k+1} < n_k$, $n_0 = N$. The coarse grid operators (A_{k+1}) are defined by,

$$A_{k+1} = R_k A_k P_k, \quad k = 0, 1, \dots, L \tag{2}$$

where $P_k \in \mathbb{R}^{n_k \times n_{k+1}}$ is an interpolation operator, $R_k (=P_k^T)$ is a restriction operator, $A_k \in \mathbb{R}^{n_k \times n_k}$ is a fine grid coefficient matrix, while $A_{k+1} \in \mathbb{R}^{n_{k+1} \times n_{k+1}}$ is a coarse grid coefficient matrix. The coarse grid operators are obtained until the size of the coarse grid operator (n_L) is sufficiently small, to obtain the exact solution using a direct solver in a reasonable time.

Aggregation algebraic multigrid uses the prior knowledge of null space or near null space of the linear operator that corresponds to low-energy error. In our work, we consider only one near null space vector of the linear operator and propagate the null space to hierarchical operators during the setup stage. For level k , the near null space vector of matrix A_k is denoted by B_k .

2.1. Setup phase

Due to simple structure of interpolation matrix, unsmooth aggregation schemes have the advantage of low setup and interpolation costs compared to that of smooth aggregation schemes or classical algebraic multigrid schemes. The algorithm for the setup is outlined in Algorithm 1.

Algorithm 1 Aggregation AMG setup

```

1: Input:  $A_0, B_0$ 
2: Output:  $A_k, B_k, P_k, R_k$  for  $k = 1, 2, ..L$ 
3: for  $k = 0, 1, 2, ..$  until  $n_L$  is small do
4:    $C_k \leftarrow \text{strength}(A_k)$ 
5:    $\text{Agg}_k \leftarrow \text{aggregate}(C_k)$ 
6:    $P_k, B_{k+1} \leftarrow \text{prolongate}(\text{Agg}_k, B_k)$ 
7:    $R_k \leftarrow P_k^T$ 
8:    $A_{k+1} \leftarrow R_k A_k P_k$ 
9: end for

```

2.1.1. Strong connections

The aggregation scheme groups a set of fine nodes that are strong negatively coupled with each other into a coarse node, called an aggregate. In general, the strength of the connection between two nodes is defined based on the matrix coefficient corresponding to the connection. Two of the popular strength of measures that are shown to be robust are symmetric strength [3] and classic strength [1]. We observe that classic strength ensures a consistent ratio of the number of aggregates to the fine grid nodes for a large class of problems. A graph of strong connections (C) is constructed based on classic strength of connection given by,

$$C_{ij} = \begin{cases} 1 & \text{if } i = j \text{ or } -s_i A_{ij} > \alpha \max_{-s_i A_{ij} > 0} -s_i A_{ij}, \quad s_i = \text{sign}(A_{ij}) \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

Here α is a threshold parameter for strong and weak connections. The coarsening ratio can be tuned by adjusting α . We choose $\alpha = 0.5$ as default for systems arising from 3D problems, and $\alpha = 0.25$ for systems arising from 2D problems. This selection of α ensures a good ratio of fine grid nodes to coarse grid nodes for a large class of problems that we considered in this paper.

Note that only strong negative connections are considered for the aggregation. This is because the algebraically smooth error changes slowly along the direction of strong negative connections, while oscillating along the strong positive connections. A detailed analysis of the convergence can be found in Stüben [8].

2.1.2. Construction of aggregates

The pairwise aggregation scheme proposed in [9] is simple and efficient on CPUs, but is not readily parallelizable on GPUs. In addition, this scheme requires two passes of pairwise aggregation to obtain a sufficient reduction in the number of unknowns for the coarse grid, increasing the setup cost. The aggregation algorithm based on maximal independent set, proposed in [7] is highly parallelizable on GPUs and constructs the aggregates faster compared to pair-wise algorithms. In the aggregation scheme, a set of nodes that are maximally independent are selected as root nodes, and aggregates are formed by grouping each of these nodes with their neighbors.

An *independent set* is a set of nodes, in which no two of them are adjacent and it is *maximal independent set* (MIS) if it is not a subset of any other independent set. The generalization of MIS is MIS(k), in which the distance between any two independent nodes is greater than k and for every other node there is at least one independent node that is within distance less than or equal to k . For the construction of aggregates, we use MIS(2) nodes as root nodes. For more aggressive coarsening, MIS(k), $k > 2$ can be used, but these may result in poor interpolation leading to slower convergence.

Algorithm 2 Parallel Maximal Independent Set (2)

```

1: Input:  $C, N \times N$  sparse matrix;
2: Output :  $s$ , set of MIS(2) nodes
3:  $s \leftarrow \{0, 0, 0, \dots, 0\}$  {initialize the state as undecided }
4:  $r \leftarrow \text{random}$  { generate random numbers  $\in (0, 1)$  }
5: for all  $i \in I$  do
6:    $v_i = \#\{j : C_{ji} = 1\} + r_i$  { number of strong influencing connections }
7: while  $\{i \in I : s_i = 0\} \neq \emptyset$  do
8:   for all  $i \in I$  do
9:      $T_i \leftarrow (s_i, v_i, i)$  { initialize tuples }
10:  for  $d = 1, 2$  do
11:    for all  $i \in I$  do {for each node in parallel }
12:       $t \leftarrow T_i$ 
13:      for  $j \in N_i$  do
14:         $t \leftarrow \max(t, T_j)$  {compare with the strong neighbors }
15:       $\hat{T}_i \leftarrow t$ 
16:       $T = \hat{T}$ 
17:  for  $i \in I$  do { for each node in parallel }
18:     $(s_{\max}, v_{\max}, i_{\max}) \leftarrow T_i$ 
19:    if  $s_i = 0$  then { if undecided }
20:      if  $i_{\max} = i$  then { if maximal }
21:         $s_i \leftarrow 1$  { mark as MIS }
22:    else
23:      if  $s_{\max} = 1$  then { else.. }
24:         $s_i \leftarrow -1$  { mark as non MIS }

```

In order to aggregate only strongly connected nodes the algorithm utilizes the graph of matrix C instead of A . The parallel MIS Algorithm 2 is similar to the proposed algorithm in [7], additionally it incorporates the number of strongly influencing connections for robustness [4]. Node j is a strongly influencing node of i if A_{ji} is a strong connection, in other words $C_{ji} = 1$. After the MIS(2) nodes are obtained, every other node i , is aggregated with its nearest MIS node.

2.1.3. Interpolation

The interpolation operator for projecting errors from a coarse grid to fine grid, is defined from the aggregations and the near null space vectors (B_k) of the linear operator (A_k). In this paper we consider only one near null space vector for constructing the interpolation operator. This results in exactly one non zero entry per row in the interpolation matrix and is given by,

$$B_k = P_k B_{k+1}, \quad P_k^T P_k = I. \quad (4)$$

Eqs. (4) ensure the near null space vector (B_k) to be in the range space of (P_k) and the interpolation matrix is orthonormal. This is done by copying the entries of B_k to the sparsity pattern of the interpolation matrix and by normalizing the columns. The restriction operator carries the residual from the fine grid to the next coarse grid in the hierarchy. A parallel algorithm for computing the interpolation operator is presented in Algorithm 3.

Algorithm 3 Interpolation

```

1: Input:  $Agg, B_k$ 
2: Output:  $P_k, R_k$ 
3: for all  $1 \leq i \leq n_k$  do                                     { for each fine node in parallel }
4:    $P_k(i, I) \leftarrow B_k(i), \quad I = Agg(i)$                    { copy null space vector entries }
5:  $R_k \leftarrow P_k^T$                                            { transpose }
6: for all  $1 \leq i \leq n_{k+1}$  do                                   { for each coarse node in parallel }
7:    $B_{k+1}(i) \leftarrow \|R_k(i, :)\|_2$                            {  $L_2$  norm of each row }
8:    $R_k(i, :) \leftarrow R_k(i, :)/B_{k+1}(i)$                        { normalize row }
9: for all  $1 \leq i \leq n_k$  do                                     { for each fine node in parallel }
10:   $P_k(i, I) \leftarrow P_k(i, I)/B_{k+1}(I)$                        { normalize interpolator }

```

2.2. Solve phase

We use a multigrid cycle as a preconditioner in a Krylov iteration which is based on either CG or GMRES. Both of these methods can be used for symmetric matrices while only GMRES can be used for non-symmetric matrices.

For unsmooth aggregation, standard V-cycle multigrid given in Algorithm 4 does not provide grid independent convergence. In contrast, the K-cycle Algorithm 5, presented in [4], addresses this issue by recursively applying Krylov iterations on the coarse levels.

Algorithm 4 V-cycle AMG

```

1:  $x_k \leftarrow Vcycle(k, b_k, x_k)$ 
2: Input : level  $k$ , rhs  $b_k$  and initial guess  $x_k$ 
3: Output : updated solution  $x_k$ 
4:  $x_k \leftarrow S_k(b_k, A_k, x_k)$                                { pre-smoothing }
5:  $r_k \leftarrow b_k - A_k x_k$                                    { compute residual }
6:  $r_{k+1} \leftarrow R_k r_k$                                      { restrict the residual to coarse-grid }
7:
8: if  $k + 1 = L$  then
9:    $x_{k+1} \leftarrow A_{k+1}^{-1} r_{k+1}$                              { exact solution of coarse-grid }
10: else
11:   $x_{k+1} \leftarrow Vcycle(k + 1, r_{k+1}, 0)$                    { recursion }
12:
13:  $x_k \leftarrow x_k + P_k x_{k+1}$                                  { prolongation }
14:  $x_k \leftarrow S_k(b_k, A_k, x_k)$                              { post-smoothing }

```

With K-cycles, the number of coarse grid corrections can grow exponentially, leading to a large number of iterations on the coarser grids. Computing these coarse grid corrections on GPUs is inefficient because of the smaller number of nodes. Interestingly, the experimental results suggest that it is sufficient to use K-cycles for only two levels at the top and V-cycles for the remaining levels to achieve grid independence convergence. This improves the overall runtime performance of the aggregation AMG on GPUs even though the number of iterations increases slightly.

2.2.1. Smoothing

In the above algorithms describing V- and K-cycles, S_k represents a smoothing operator that removes high frequency errors corresponding to the matrix A_k . We use one pre- and one post-smoothing step for each level. We consider Jacobi, damped Jacobi and symmetric Gauss–Seidel smoothing, all of which are symmetric. Jacobi and damped Jacobi are easily parallelizable on CPUs and GPUs. However, the Gauss–Seidel algorithm is completely serial in nature and cannot be ported to GPUs efficiently. A multi-color variant of symmetric Gauss–Seidel algorithm has shown to scale well on GPUs for large matrices with fewer numbers of colors. This modified algorithm adds a setup cost of the graph coloring and a reordering of the matrix, and also results in slow down of convergence. Numerical results indicate that damped Jacobi is both scalable and efficient, and requires low setup cost compared to multi-color Gauss–Seidel. In addition, damped Jacobi produces solutions that are independent of number of GPUs/processors unlike multi-color Gauss–Seidel.

Damped Jacobi iteration is given by,

$$x \leftarrow x + \omega D^{-1}(b - Ax), \quad \omega = \frac{4}{3} \frac{1}{\rho(D^{-1}A)}, \tag{5}$$

where $\rho(D^{-1}A)$ is the spectral radius of the matrix $D^{-1}A$ and it is estimated from the eigenvalues of the Hessenberg matrix obtained using m (≤ 5) Arnoldi iterations.

Algorithm 5 K-cycle AMG

```

1:  $x_k \leftarrow \text{Kcycle}(k, b_k, x_k)$ 
2: Input : level  $k$ , rhs  $b_k$  and initial guess  $x_k$ 
3: Output : updated solution  $x_k$ 
4:  $x_k \leftarrow S_k(b_k, A_k, x_k)$  { pre-smoothing }
5:  $r_k \leftarrow b_k - A_k x_k$  { compute residual }
6:  $r_{k+1} \leftarrow R_k r_k$  { restrict the residual to coarse-grid }
7:
8: if  $k + 1 = L$  then
9:    $x_{k+1} \leftarrow A_{k+1}^{-1} r_{k+1}$  { exact solution of coarse-grid }
10: else
11:    $c_{k+1} \leftarrow \text{Kcycle}(k + 1, r_{k+1}, x_{k+1})$  { inner first Krylov iteration }
12:    $v_{k+1} \leftarrow A_{k+1} c_{k+1}$ 
13:    $\rho_1 \leftarrow c_{k+1}^T v_{k+1}, \quad \alpha_1 \leftarrow c_{k+1}^T r_{k+1}$  if CG
14:    $\rho_1 \leftarrow \|v_{k+1}\|^2, \quad \alpha_1 \leftarrow v_{k+1}^T r_{k+1}$  if GMRES
15:
16:    $\tilde{r}_{k+1} \leftarrow r_{k+1} - \frac{\alpha_1}{\rho_1} v_{k+1}$ 
17:   if  $\|\tilde{r}_{k+1}\| \leq t \|r_{k+1}\|$  then
18:      $x_{k+1} \leftarrow \frac{\alpha_1}{\rho_1} c_{k+1}$ 
19:   else
20:      $d_{k+1} \leftarrow \text{Kcycle}(k + 1, \tilde{r}_{k+1}, x_{k+1})$  { inner second Krylov iteration }
21:      $w_{k+1} \leftarrow A_{k+1} d_{k+1}$ 
22:      $\gamma \leftarrow d_{k+1}^T v_{k+1}, \quad \beta \leftarrow d_{k+1}^T w_{k+1}, \quad \alpha_2 \leftarrow d_{k+1}^T \tilde{r}_{k+1}$  if CG
23:      $\gamma \leftarrow w_{k+1}^T v_{k+1}, \quad \beta \leftarrow \|w_{k+1}\|^2, \quad \alpha_2 \leftarrow w_{k+1}^T \tilde{r}_{k+1}$  if GMRES
24:      $\rho_2 \leftarrow \beta - \frac{\gamma^2}{\rho_1}$ 
25:      $x_{k+1} \leftarrow (\frac{\alpha_1}{\rho_1} - \frac{\gamma \alpha_2}{\rho_1 \rho_2}) c_{k+1} + \frac{\alpha_2}{\rho_2} d_{k+1}$ 
26:
27:    $x_k \leftarrow x_k + P_k x_{k+1}$  { prolongation }
28:    $r_k \leftarrow \tilde{r}_k - A_k x_k$  { compute new residual }
29:    $x_k \leftarrow S_k(r_k, A_k, x_k)$  { post-smoothing }

```

3. Implementation

We use the CUDA programming model for our GPU implementations and the OpenMP threading model for CPU counter parts. For the setup phase, the matrices are stored in compressed sparse row (CSR) format for both CPU and GPU implementations. For sparse matrix–matrix products (spmm), we also use the coordinate (COO) format for storing one of the multipliers.

In the solve phase, matrix vector multiplication axpy, is the dominant component. For the CPU implementations we store the matrices in CSR format and for GPU implementations we use a hybrid matrix format [10] that is a combination of the ELL and COO formats. For multi-GPU axpy, the matrix entries corresponding to the columns that are processed by other processors are stored in COO format. This way, the multiplication with the matrix stored in ELL, and the communication of the vector data can be overlapped. We use CUDA peer-to-peer transfers for inter GPU communication of vector entries.

For parallel primitive operations such as reduction, sort by key, scan, and gather required for the aggregation algorithm, we use the libraries thrust [11] and/or GAMPACK.

For the coarsest level, to solve the linear system exactly, we use the CPU based direct solver SuperLU [12], for both GPU and CPU implementations. Only one CPU thread is used for solving the coarse linear system. We ensure that dimension of the coarse linear system is about 600 for GPU implementations to minimize the overall runtime.

3.1. Galerkin product

After the interpolation and restriction matrices are computed at a level k , the coarse grid operator ($A_{k+1} = R_k A_k P_k$) is constructed using two matrix–matrix multiplications, $A_k \times P_k$ and $R_k \times (A_k \times P_k)$. All of these matrices are stored in CSR format. A GPU implementation of sparse matrix–matrix multiplication is described in [6]. We describe another approach that will be useful in the solution of nonlinear systems of equations, where the aggregations are not constructed at every Newton step when the change in matrix entries is not significant. In these cases, sparsity pattern of the matrices in the hierarchy do not change since the pattern depends only on the aggregations and the sparsity pattern of A_0 . This is true only if the sparsity pattern of A_0 does not change at every Newton step.

The Galerkin product can be represented as,

$$(A_{k+1})_{IJ} = \sum_{i \in G_I} \sum_{j \in G_J} (R_k)_{ii} (A_k)_{ij} (P_k)_{jj} = \sum_{i \in G_I} \sum_{j \in G_J} (P_k)_{ii} (A_k)_{ij} (P_k)_{jj} \tag{6}$$

where, G_I is a set of nodes that form the I th aggregate. Each of the fine nodes is grouped into exactly one aggregate. Therefore, for each nonzero of the fine grid matrix $A_{ij} \neq 0$, the index of corresponding coarse grid matrix is IJ , where I and J are aggregation indices corresponding to i and j respectively. These coarse grid nonzero row and column indices are tabulated and sorted based on lexicographic ordering of ordered pair (I, J) . The number of nonzeros in coarse grid matrix is obtained by counting the unique indices (I, J) . All the coefficients corresponding to a same index, will sum to a unique nonzero of the coarse grid. This summation is done using segmented reduction. The pattern of this computation can be repeated by storing the sorting indices and the segmented reduction indices. This way, the entire Galerkin product can be performed in less than 1/10th of the regular spmm compute time. Note that the interpolation operator is not reconstructed, since its entries depend on the near null space vector and not directly on the entries of the level matrix. For the cases in which we do not reuse the hierarchy information; we use GPU sparse matrix–matrix multiplications.

Consider a matrix A_k , with sparsity graph shown in Fig. 1. The aggregates are shown in different colors and enumerated in Roman numerals. The sequence of operations to compute the coarse grid matrix is illustrated in Eq. (7). Initially the nonzero coefficients A_{ij} 's are sorted based on the corresponding aggregation indices (I, J) to obtain unique (I, J) pairs. The Galerkin product is then finished by adding all the duplicate pairs (I, J) to corresponding nonzero of the coarse grid matrix A_{k+1} . Note that each nonzero entry $(A_k)_{ij}$ is multiplied with $(P_k)_{jj}$ and $(R_k)_{ii}$ before sorting the coefficients.

<table style="border-collapse: collapse; width: 100%;"> <tr><th>i</th><th>j</th><th>$(A_k)_{ij}$</th></tr> <tr><td>1</td><td>1</td><td>a_{11}</td></tr> <tr><td>1</td><td>2</td><td>a_{12}</td></tr> <tr><td>1</td><td>6</td><td>a_{16}</td></tr> <tr><td>2</td><td>1</td><td>a_{21}</td></tr> <tr><td>2</td><td>2</td><td>a_{22}</td></tr> <tr><td>2</td><td>4</td><td>a_{24}</td></tr> <tr><td>2</td><td>7</td><td>a_{27}</td></tr> <tr><td>3</td><td>3</td><td>a_{33}</td></tr> <tr><td>3</td><td>5</td><td>a_{35}</td></tr> <tr><td>3</td><td>8</td><td>a_{38}</td></tr> <tr><td>4</td><td>2</td><td>a_{42}</td></tr> <tr><td>4</td><td>4</td><td>a_{44}</td></tr> <tr><td>4</td><td>5</td><td>a_{45}</td></tr> <tr><td>4</td><td>6</td><td>a_{46}</td></tr> <tr><td>4</td><td>7</td><td>a_{47}</td></tr> <tr><td>5</td><td>3</td><td>a_{53}</td></tr> <tr><td>5</td><td>4</td><td>a_{54}</td></tr> <tr><td>5</td><td>5</td><td>a_{55}</td></tr> <tr><td>5</td><td>8</td><td>a_{58}</td></tr> <tr><td>6</td><td>1</td><td>a_{61}</td></tr> <tr><td>6</td><td>4</td><td>a_{64}</td></tr> <tr><td>6</td><td>6</td><td>a_{66}</td></tr> <tr><td>7</td><td>2</td><td>a_{72}</td></tr> <tr><td>7</td><td>4</td><td>a_{74}</td></tr> <tr><td>7</td><td>7</td><td>a_{77}</td></tr> <tr><td>8</td><td>3</td><td>a_{83}</td></tr> <tr><td>8</td><td>5</td><td>a_{85}</td></tr> <tr><td>8</td><td>8</td><td>a_{88}</td></tr> </table>	i	j	$(A_k)_{ij}$	1	1	a_{11}	1	2	a_{12}	1	6	a_{16}	2	1	a_{21}	2	2	a_{22}	2	4	a_{24}	2	7	a_{27}	3	3	a_{33}	3	5	a_{35}	3	8	a_{38}	4	2	a_{42}	4	4	a_{44}	4	5	a_{45}	4	6	a_{46}	4	7	a_{47}	5	3	a_{53}	5	4	a_{54}	5	5	a_{55}	5	8	a_{58}	6	1	a_{61}	6	4	a_{64}	6	6	a_{66}	7	2	a_{72}	7	4	a_{74}	7	7	a_{77}	8	3	a_{83}	8	5	a_{85}	8	8	a_{88}	$\xrightarrow{(i,j) \rightarrow (I,J)}$	<table style="border-collapse: collapse; width: 100%;"> <tr><th>I</th><th>J</th><th>$(A_k)_{ij}$</th></tr> <tr><td>2</td><td>2</td><td>a_{11}</td></tr> <tr><td>2</td><td>1</td><td>a_{12}</td></tr> <tr><td>2</td><td>2</td><td>a_{16}</td></tr> <tr><td>1</td><td>2</td><td>a_{21}</td></tr> <tr><td>1</td><td>1</td><td>a_{22}</td></tr> <tr><td>1</td><td>1</td><td>a_{24}</td></tr> <tr><td>1</td><td>1</td><td>a_{27}</td></tr> <tr><td>3</td><td>3</td><td>a_{33}</td></tr> <tr><td>3</td><td>3</td><td>a_{35}</td></tr> <tr><td>3</td><td>3</td><td>a_{38}</td></tr> <tr><td>1</td><td>1</td><td>a_{42}</td></tr> <tr><td>1</td><td>1</td><td>a_{44}</td></tr> <tr><td>1</td><td>1</td><td>a_{45}</td></tr> <tr><td>1</td><td>2</td><td>a_{46}</td></tr> <tr><td>1</td><td>1</td><td>a_{47}</td></tr> <tr><td>3</td><td>3</td><td>a_{53}</td></tr> <tr><td>3</td><td>1</td><td>a_{54}</td></tr> <tr><td>3</td><td>3</td><td>a_{55}</td></tr> <tr><td>3</td><td>3</td><td>a_{58}</td></tr> <tr><td>2</td><td>2</td><td>a_{61}</td></tr> <tr><td>2</td><td>1</td><td>a_{64}</td></tr> <tr><td>2</td><td>2</td><td>a_{66}</td></tr> <tr><td>1</td><td>1</td><td>a_{72}</td></tr> <tr><td>1</td><td>1</td><td>a_{74}</td></tr> <tr><td>1</td><td>1</td><td>a_{77}</td></tr> <tr><td>3</td><td>3</td><td>a_{83}</td></tr> <tr><td>3</td><td>3</td><td>a_{85}</td></tr> <tr><td>3</td><td>3</td><td>a_{88}</td></tr> </table>	I	J	$(A_k)_{ij}$	2	2	a_{11}	2	1	a_{12}	2	2	a_{16}	1	2	a_{21}	1	1	a_{22}	1	1	a_{24}	1	1	a_{27}	3	3	a_{33}	3	3	a_{35}	3	3	a_{38}	1	1	a_{42}	1	1	a_{44}	1	1	a_{45}	1	2	a_{46}	1	1	a_{47}	3	3	a_{53}	3	1	a_{54}	3	3	a_{55}	3	3	a_{58}	2	2	a_{61}	2	1	a_{64}	2	2	a_{66}	1	1	a_{72}	1	1	a_{74}	1	1	a_{77}	3	3	a_{83}	3	3	a_{85}	3	3	a_{88}	$\xrightarrow{\text{sort}}$	<table style="border-collapse: collapse; width: 100%;"> <tr><th>I</th><th>J</th><th>$(A_k)_{ij}$</th></tr> <tr><td>1</td><td>1</td><td>a_{22}</td></tr> <tr><td>1</td><td>1</td><td>a_{24}</td></tr> <tr><td>1</td><td>1</td><td>a_{27}</td></tr> <tr><td>1</td><td>1</td><td>a_{42}</td></tr> <tr><td>1</td><td>1</td><td>a_{44}</td></tr> <tr><td>1</td><td>1</td><td>a_{47}</td></tr> <tr><td>1</td><td>1</td><td>a_{72}</td></tr> <tr><td>1</td><td>1</td><td>a_{74}</td></tr> <tr><td>1</td><td>1</td><td>a_{77}</td></tr> <tr><td>1</td><td>2</td><td>a_{21}</td></tr> <tr><td>1</td><td>2</td><td>a_{46}</td></tr> <tr><td>1</td><td>3</td><td>a_{45}</td></tr> <tr><td>2</td><td>1</td><td>a_{12}</td></tr> <tr><td>2</td><td>1</td><td>a_{64}</td></tr> <tr><td>2</td><td>2</td><td>a_{11}</td></tr> <tr><td>2</td><td>2</td><td>a_{16}</td></tr> <tr><td>2</td><td>2</td><td>a_{61}</td></tr> <tr><td>2</td><td>2</td><td>a_{66}</td></tr> <tr><td>3</td><td>1</td><td>a_{54}</td></tr> <tr><td>3</td><td>1</td><td>a_{54}</td></tr> <tr><td>3</td><td>3</td><td>a_{33}</td></tr> <tr><td>3</td><td>3</td><td>a_{35}</td></tr> <tr><td>3</td><td>3</td><td>a_{38}</td></tr> <tr><td>3</td><td>3</td><td>a_{53}</td></tr> <tr><td>3</td><td>3</td><td>a_{55}</td></tr> <tr><td>3</td><td>3</td><td>a_{58}</td></tr> <tr><td>3</td><td>3</td><td>a_{83}</td></tr> <tr><td>3</td><td>3</td><td>a_{85}</td></tr> <tr><td>3</td><td>3</td><td>a_{88}</td></tr> </table>	I	J	$(A_k)_{ij}$	1	1	a_{22}	1	1	a_{24}	1	1	a_{27}	1	1	a_{42}	1	1	a_{44}	1	1	a_{47}	1	1	a_{72}	1	1	a_{74}	1	1	a_{77}	1	2	a_{21}	1	2	a_{46}	1	3	a_{45}	2	1	a_{12}	2	1	a_{64}	2	2	a_{11}	2	2	a_{16}	2	2	a_{61}	2	2	a_{66}	3	1	a_{54}	3	1	a_{54}	3	3	a_{33}	3	3	a_{35}	3	3	a_{38}	3	3	a_{53}	3	3	a_{55}	3	3	a_{58}	3	3	a_{83}	3	3	a_{85}	3	3	a_{88}	$\xrightarrow{\text{reduce}}$	<table style="border-collapse: collapse; width: 100%;"> <tr><th>I</th><th>J</th><th>$(A_{k+1})_{IJ}$</th></tr> <tr><td>1</td><td>1</td><td>$a_{22} + a_{24}$ $+ a_{27} + a_{42}$ $+ a_{44} + a_{47}$ $+ a_{72} + a_{74}$ $+ a_{77}$</td></tr> <tr><td>1</td><td>2</td><td>$a_{21} + a_{46}$</td></tr> <tr><td>1</td><td>3</td><td>a_{45}</td></tr> <tr><td>2</td><td>1</td><td>$a_{12} + a_{64}$</td></tr> <tr><td>2</td><td>2</td><td>$a_{11} + a_{16}$ $+ a_{61} + a_{66}$</td></tr> <tr><td>3</td><td>1</td><td>a_{54}</td></tr> <tr><td>3</td><td>3</td><td>$a_{33} + a_{35}$ $+ a_{38} + a_{53}$ $+ a_{55} + a_{58}$ $+ a_{83} + a_{85}$ $+ a_{88}$</td></tr> </table>	I	J	$(A_{k+1})_{IJ}$	1	1	$a_{22} + a_{24}$ $+ a_{27} + a_{42}$ $+ a_{44} + a_{47}$ $+ a_{72} + a_{74}$ $+ a_{77}$	1	2	$a_{21} + a_{46}$	1	3	a_{45}	2	1	$a_{12} + a_{64}$	2	2	$a_{11} + a_{16}$ $+ a_{61} + a_{66}$	3	1	a_{54}	3	3	$a_{33} + a_{35}$ $+ a_{38} + a_{53}$ $+ a_{55} + a_{58}$ $+ a_{83} + a_{85}$ $+ a_{88}$	$\tag{7}$
i	j	$(A_k)_{ij}$																																																																																																																																																																																																																																																																																																					
1	1	a_{11}																																																																																																																																																																																																																																																																																																					
1	2	a_{12}																																																																																																																																																																																																																																																																																																					
1	6	a_{16}																																																																																																																																																																																																																																																																																																					
2	1	a_{21}																																																																																																																																																																																																																																																																																																					
2	2	a_{22}																																																																																																																																																																																																																																																																																																					
2	4	a_{24}																																																																																																																																																																																																																																																																																																					
2	7	a_{27}																																																																																																																																																																																																																																																																																																					
3	3	a_{33}																																																																																																																																																																																																																																																																																																					
3	5	a_{35}																																																																																																																																																																																																																																																																																																					
3	8	a_{38}																																																																																																																																																																																																																																																																																																					
4	2	a_{42}																																																																																																																																																																																																																																																																																																					
4	4	a_{44}																																																																																																																																																																																																																																																																																																					
4	5	a_{45}																																																																																																																																																																																																																																																																																																					
4	6	a_{46}																																																																																																																																																																																																																																																																																																					
4	7	a_{47}																																																																																																																																																																																																																																																																																																					
5	3	a_{53}																																																																																																																																																																																																																																																																																																					
5	4	a_{54}																																																																																																																																																																																																																																																																																																					
5	5	a_{55}																																																																																																																																																																																																																																																																																																					
5	8	a_{58}																																																																																																																																																																																																																																																																																																					
6	1	a_{61}																																																																																																																																																																																																																																																																																																					
6	4	a_{64}																																																																																																																																																																																																																																																																																																					
6	6	a_{66}																																																																																																																																																																																																																																																																																																					
7	2	a_{72}																																																																																																																																																																																																																																																																																																					
7	4	a_{74}																																																																																																																																																																																																																																																																																																					
7	7	a_{77}																																																																																																																																																																																																																																																																																																					
8	3	a_{83}																																																																																																																																																																																																																																																																																																					
8	5	a_{85}																																																																																																																																																																																																																																																																																																					
8	8	a_{88}																																																																																																																																																																																																																																																																																																					
I	J	$(A_k)_{ij}$																																																																																																																																																																																																																																																																																																					
2	2	a_{11}																																																																																																																																																																																																																																																																																																					
2	1	a_{12}																																																																																																																																																																																																																																																																																																					
2	2	a_{16}																																																																																																																																																																																																																																																																																																					
1	2	a_{21}																																																																																																																																																																																																																																																																																																					
1	1	a_{22}																																																																																																																																																																																																																																																																																																					
1	1	a_{24}																																																																																																																																																																																																																																																																																																					
1	1	a_{27}																																																																																																																																																																																																																																																																																																					
3	3	a_{33}																																																																																																																																																																																																																																																																																																					
3	3	a_{35}																																																																																																																																																																																																																																																																																																					
3	3	a_{38}																																																																																																																																																																																																																																																																																																					
1	1	a_{42}																																																																																																																																																																																																																																																																																																					
1	1	a_{44}																																																																																																																																																																																																																																																																																																					
1	1	a_{45}																																																																																																																																																																																																																																																																																																					
1	2	a_{46}																																																																																																																																																																																																																																																																																																					
1	1	a_{47}																																																																																																																																																																																																																																																																																																					
3	3	a_{53}																																																																																																																																																																																																																																																																																																					
3	1	a_{54}																																																																																																																																																																																																																																																																																																					
3	3	a_{55}																																																																																																																																																																																																																																																																																																					
3	3	a_{58}																																																																																																																																																																																																																																																																																																					
2	2	a_{61}																																																																																																																																																																																																																																																																																																					
2	1	a_{64}																																																																																																																																																																																																																																																																																																					
2	2	a_{66}																																																																																																																																																																																																																																																																																																					
1	1	a_{72}																																																																																																																																																																																																																																																																																																					
1	1	a_{74}																																																																																																																																																																																																																																																																																																					
1	1	a_{77}																																																																																																																																																																																																																																																																																																					
3	3	a_{83}																																																																																																																																																																																																																																																																																																					
3	3	a_{85}																																																																																																																																																																																																																																																																																																					
3	3	a_{88}																																																																																																																																																																																																																																																																																																					
I	J	$(A_k)_{ij}$																																																																																																																																																																																																																																																																																																					
1	1	a_{22}																																																																																																																																																																																																																																																																																																					
1	1	a_{24}																																																																																																																																																																																																																																																																																																					
1	1	a_{27}																																																																																																																																																																																																																																																																																																					
1	1	a_{42}																																																																																																																																																																																																																																																																																																					
1	1	a_{44}																																																																																																																																																																																																																																																																																																					
1	1	a_{47}																																																																																																																																																																																																																																																																																																					
1	1	a_{72}																																																																																																																																																																																																																																																																																																					
1	1	a_{74}																																																																																																																																																																																																																																																																																																					
1	1	a_{77}																																																																																																																																																																																																																																																																																																					
1	2	a_{21}																																																																																																																																																																																																																																																																																																					
1	2	a_{46}																																																																																																																																																																																																																																																																																																					
1	3	a_{45}																																																																																																																																																																																																																																																																																																					
2	1	a_{12}																																																																																																																																																																																																																																																																																																					
2	1	a_{64}																																																																																																																																																																																																																																																																																																					
2	2	a_{11}																																																																																																																																																																																																																																																																																																					
2	2	a_{16}																																																																																																																																																																																																																																																																																																					
2	2	a_{61}																																																																																																																																																																																																																																																																																																					
2	2	a_{66}																																																																																																																																																																																																																																																																																																					
3	1	a_{54}																																																																																																																																																																																																																																																																																																					
3	1	a_{54}																																																																																																																																																																																																																																																																																																					
3	3	a_{33}																																																																																																																																																																																																																																																																																																					
3	3	a_{35}																																																																																																																																																																																																																																																																																																					
3	3	a_{38}																																																																																																																																																																																																																																																																																																					
3	3	a_{53}																																																																																																																																																																																																																																																																																																					
3	3	a_{55}																																																																																																																																																																																																																																																																																																					
3	3	a_{58}																																																																																																																																																																																																																																																																																																					
3	3	a_{83}																																																																																																																																																																																																																																																																																																					
3	3	a_{85}																																																																																																																																																																																																																																																																																																					
3	3	a_{88}																																																																																																																																																																																																																																																																																																					
I	J	$(A_{k+1})_{IJ}$																																																																																																																																																																																																																																																																																																					
1	1	$a_{22} + a_{24}$ $+ a_{27} + a_{42}$ $+ a_{44} + a_{47}$ $+ a_{72} + a_{74}$ $+ a_{77}$																																																																																																																																																																																																																																																																																																					
1	2	$a_{21} + a_{46}$																																																																																																																																																																																																																																																																																																					
1	3	a_{45}																																																																																																																																																																																																																																																																																																					
2	1	$a_{12} + a_{64}$																																																																																																																																																																																																																																																																																																					
2	2	$a_{11} + a_{16}$ $+ a_{61} + a_{66}$																																																																																																																																																																																																																																																																																																					
3	1	a_{54}																																																																																																																																																																																																																																																																																																					
3	3	$a_{33} + a_{35}$ $+ a_{38} + a_{53}$ $+ a_{55} + a_{58}$ $+ a_{83} + a_{85}$ $+ a_{88}$																																																																																																																																																																																																																																																																																																					

4. Numerical experiments

In order to determine the efficiency of the aggregation AMG solver, we compare its performance with that of GAMPACK implementation of classical AMG solver. This classical AMG solver was compared with the well known classical AMG solver HYPRE [6]. We compare the best performances of both solvers for each case. We used a flexible GMRES as the outer Krylov solver in all test cases.

In Table 1, we compare the storage complexity of the hierarchy of smooth aggregation from CUSP with unsmooth aggregation for a matrix related to a 2D anisotropic Poisson problem. For smooth aggregation, the density of the coarse grid matrices increase rapidly while it is almost fixed for unsmooth aggregation, ensuring low memory usage by unsmooth aggregation. The dense coarse grids of smooth aggregation increase the communication cost during the smoothing, interpolation,

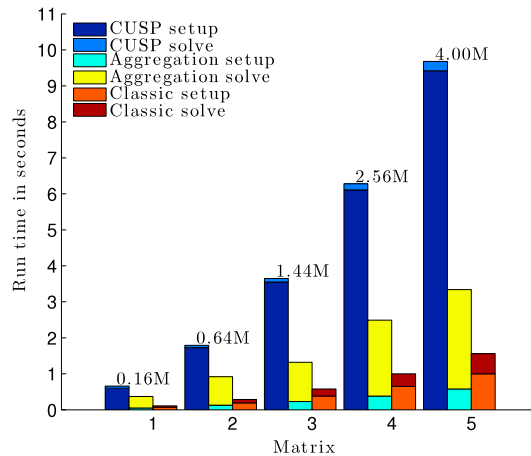


Fig. 3. Comparison of performance of CUSP smooth aggregation, GAMPACK aggregation, and GAMPACK classic AMG for 2D Poisson problem. Experiments ran on a single NVIDIA Tesla M2090 GPU. The numbers on the bars represent the size of the corresponding system in millions.

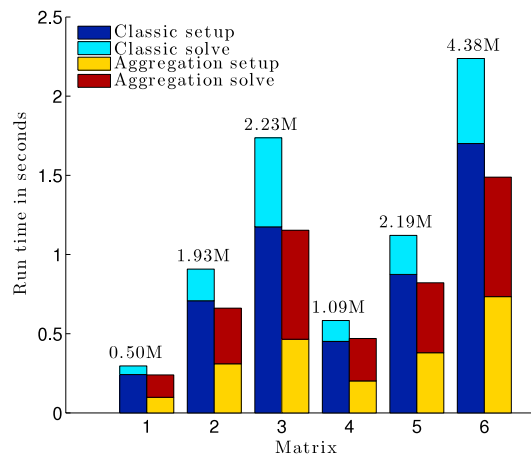


Fig. 4. Comparison of performance of GAMPACK classical AMG and aggregation AMG for RS and pressure matrices. Experiments ran on a single NVIDIA Tesla M2090 GPU. The numbers on the bars represent the size of the corresponding system in millions.

and restriction operations, decreasing the overall parallel efficiency of AMG. In Fig. 2 we compare the growth of density of coarse grid matrices for smooth aggregation with the fixed sparsity pattern of unsmooth aggregation for the same problem.

In Fig. 3, we compare the setup and solve times with those of CUSP smooth aggregation and GAMPACK classical AMG for an anisotropic 2D Poisson problem. The setup times for CUSP are significantly larger while the solve times are significantly smaller compared to that of GAMPACK. Very small solve time of smooth aggregation is due to fewer number of iterations required for the convergence. This is because of very good interpolation between coarse and fine grids. Since the setup is the dominant cost for smooth aggregation, the overall runtime for CUSP smooth aggregation is also very large compared to GAMPACK unsmooth aggregation. The large setup time for smooth aggregation is due to the density of the coarse levels.

In Fig. 4, we compare the setup and solve times with classical AMG. Unlike classical AMG, the setup cost of aggregation is not dominant, and more than half of the time is spent on the solve. The results indicate that the solve stage of V-cycle classical AMG is about two times faster than aggregation AMG while the setup time of classical AMG is three times slower than that of aggregation AMG. Since a major component of computation is the setup for classical AMG, in overall, aggregation AMG is faster than classical AMG for the class of problems we considered. For these test cases, CUSP smooth aggregation could not construct the AMG hierarchy in a reasonable time, hence we do not compare the results with CUSP.

In order to study the scalability of the solver, we consider a series of down-scaled problems corresponding to the SPE10 pressure matrix problem [13]. Since the larger matrices do not fit on one or two GPUs, we include the timings with multiple GPUs for both classical and aggregation AMG in Table 2. As reported in the table, the overall efficiency increases as the problem size increases, for a fixed number of GPUs. However, the efficiency per GPU decreases with the increase in the number of GPUs. This is because of the high ratio of the communication cost to computation cost on GPUs for coarse level matrices. We consider to improve the multi-GPU performance scaling in our future work.

In Fig. 5, we compare the performance of both solvers as the problem size increases. For the smallest problem ($\sim 2M$) we considered, classical AMG out-performs aggregation AMG, while aggregation AMG is about 50% faster to classical AMG

Table 1

Comparison of the hierarchy grid sparsity of CUSP smooth aggregation (left) and GAMPACK unsmooth aggregation (right) for 2D anisotropic Poisson problem on 1000×1000 grid.

CUSP				GAMPACK			
Level	Unknowns	nnz	nnz/row	Level	Unknowns	nnz	nnz/row
0	1,000,000	4,996,000	5.00	0	1,000,000	4,996,000	5.00
1	275,264	4,570,716	16.60	1	274,979	1,770,339	6.44
2	76,002	4,398,710	57.88	2	77,433	525,877	6.79
3	8,601	737,445	85.74	3	21,823	147,803	6.77
4	923	82,045	88.89	4	6,150	39,298	6.39
5	246	28,140	114.39	5	938	5,128	5.47
				6	169	641	3.79

Table 2

Scaling : Timings in seconds for GAMPACK setup and solve for aggregation AMG, and classical AMG on single and multiple GPUs. The first column represents the number of cells in the original SPE10 pressure matrix problem. For each matrix, the first row gives the timings for aggregation AMG while the second row gives the timings for classical AMG. All timings are in seconds and include setup and solve to a relative tolerance of 10^{-6} . The number of GMRES iterations are given in parentheses.

N	1 × M2090	2 × M2090	3 × M2090	4 × M2090
1.1 M	0.469 (15)	0.387 (15)	0.430 (15)	0.409 (15)
	0.583 (10)	0.365 (10)	0.326 (10)	0.312 (11)
2.2 M	0.821 (15)	0.580 (15)	0.561 (15)	0.500 (15)
	1.121 (10)	0.684 (11)	0.526 (11)	0.454 (10)
4.4 M	1.488 (15)	1.079 (15)	0.875 (15)	0.874 (15)
	2.237 (11)	1.302 (11)	0.953 (11)	0.774 (11)
8.9 M	-	1.636 (14)	1.345 (15)	1.177 (15)
	-	2.556 (12)	1.856 (12)	1.564 (15)
13.5 M	-	-	1.813 (15)	1.471 (14)
	-	-	2.769 (12)	2.159 (12)

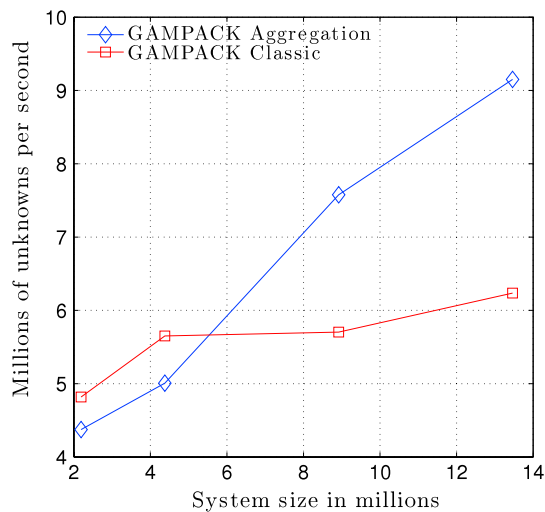


Fig. 5. Scaling of computation rate (in millions of unknowns per second) with system size. The rates include both setup and solve. The computations ran on 4 NVIDIA Tesla M2090 GPUs.

for the largest problem (~14M) we considered. The performance of aggregation AMG improves at a faster rate compared to that of classical AMG.

From Fig. 6, we observe significant decline in the performance of aggregation AMG as the number of GPUs used for the computation increases. This is because K-cycles require more number of coarse grid corrections and GPUs are inefficient for systems of small size. For aggregation AMG to be efficient on multiple GPUs, the system size has to be sufficiently large.

5. Conclusions and future work

We have presented a robust, efficient and scalable aggregation algebraic multigrid solver and have verified the robustness with matrices from various applications. We compared its performance with a GPU accelerated classical AMG solver and a

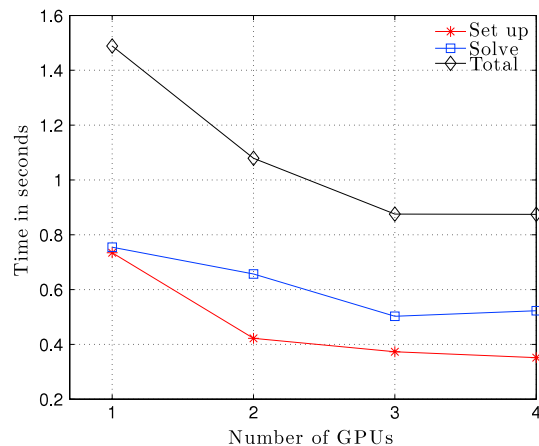


Fig. 6. Strong scaling : Computation time with number of GPUs. Computations ran for a system of four million unknowns on NVIDIA Tesla M2090 GPUs.

GPU accelerated smooth aggregation AMG solver, and observed that the setup cost of aggregation AMG is significantly low and the cost of the solution phase is high. In conclusion, aggregation AMG is extremely efficient for systems of sufficiently large size. Furthermore, for practical applications, that do not require high accuracy of the linear solver, aggregation AMG is extremely efficient due to very low setup costs. Our future work includes using MPI with multi-node computations for very large systems and incorporating polynomial smoothers to improve the convergence of aggregation AMG.

Acknowledgments

The authors gratefully acknowledge the grant from DOE & ANL (ANL Subcontract No. 1F-32301 on DOE grant No. DE-AC02-06CH11357) and fellowship from Ken Kennedy Institute for Information Technology, Rice University to the first author. The authors gratefully acknowledge Tim Warburton, Jesse Chan, and Cynthia Wood for fruitful discussions and suggestions during the preparation of this manuscript.

References

- [1] J. Ruge, K. Stüben, Algebraic multigrid, *Multigrid Methods* 3 (1987) 73–130.
- [2] K. Stüben, A review of algebraic multigrid, *J. Comput. Appl. Math.* 128 (1) (2001) 281–309.
- [3] P. Vaněk, J. Mandel, M. Brezina, Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems, *Computing* 56 (3) (1996) 179–196.
- [4] Y. Notay, An aggregation-based algebraic multigrid method, *Electron. Trans. Numer. Anal.* 37 (6) (2010) 123–146.
- [5] M. Emans, M. Liebmann, B. Basara, Steps towards GPU accelerated aggregation AMG, in: 2012 11th International Symposium on Parallel and Distributed Computing, (ISPDC), IEEE, 2012, pp. 79–86.
- [6] K. Esler, V. Natoli, A. Samardzic, GAMPACK (GPU Accelerated Algebraic Multigrid Package), in: ECMOR XIII-13th European Conference on the Mathematics of Oil Recovery, 2012.
- [7] N. Bell, S. Dalton, L.N. Olson, Exposing fine-grained parallelism in algebraic multigrid methods, *SIAM J. Sci. Comput.* 34 (4) (2012) C123–C152.
- [8] K. Stüben, An introduction to algebraic multigrid, *Multigrid* (2001) 413–532.
- [9] Y. Notay, Aggregation-based algebraic multilevel preconditioning, *SIAM J. Matrix Anal. Appl.* 27 (4) (2006) 998–1018.
- [10] N. Bell, M. Garland, Efficient Sparse Matrix–Vector Multiplication on CUDA, Tech. Rep., NVIDIA Technical Report NVR-2008-004, NVIDIA Corporation, 2008.
- [11] J. Hoberock, N. Bell, Thrust: A parallel template library, version 1.7.0 (2010). URL <http://thrust.github.io/>.
- [12] J.W. Demmel, S.C. Eisenstat, J.R. Gilbert, X.S. Li, J.W.H. Liu, A supernodal approach to sparse partial pivoting, *SIAM J. Matrix Anal. Appl.* 20 (3) (1999) 720–755.
- [13] M. Christie, M. Blunt, Tenth SPE comparative solution project: A comparison of upscaling techniques, *SPE Reservoir Eval. Eng.* 4 (4) (2001) 308–317.