Journal of Parallel and Distributed Computing 130 (2019) 166–178

Contents lists available at ScienceDirect

J. Parallel Distrib. Comput.

journal homepage: www.elsevier.com/locate/jpdc

Node aware sparse matrix–vector multiplication☆

Amanda Bienz *, William D. Gropp, Luke N. Olson

Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL, 61801, United States

A R T I C L E   I N F O

Article history:
Received 4 November 2017
Received in revised form 13 July 2018
Accepted 23 March 2019
Available online 29 March 2019

Keywords:
Sparse
Matrix–vector multiplication
SpMV
Parallel communication
Node aware

A B S T R A C T

The sparse matrix–vector multiply (SpMV) operation is a key computational kernel in many simulations and linear solvers. The large communication requirements associated with a reference implementation of a parallel SpMV result in poor parallel scalability. The cost of communication depends on the physical locations of the send and receive processes: messages injected into the network are more costly than messages sent between processes on the same node. In this paper, a node aware parallel SpMV (NAPSpMV) is introduced to exploit knowledge of the system topology, specifically the node-processor layout, to reduce costs associated with communication. The values of the input vector are redistributed to minimize both the number and the size of messages that are injected into the network during a SpMV, leading to a reduction in communication costs. A variety of computational experiments that highlight the efficiency of this approach are presented.

1. Introduction

The sparse matrix–vector multiply (SpMV) is a widely used operation in many simulations and the main kernel in iterative solvers. The focus of this paper is on the parallel SpMV, namely

$$w \leftarrow A \cdot v \tag{1}$$

where A is a sparse $N \times N$ matrix and v is a dense N-dimensional vector. In parallel, the sparse system is often distributed across n_p processes such that each process holds a contiguous block of rows from the matrix A, and equivalent rows from the vectors v and w, as shown in Fig. 1. A common approach is to also split the rows of A on a single process into two groups: an on-process block, containing the columns of the matrix that correspond to vector values stored locally, and an off-process block, containing matrix non-zeros that are associated with vector values that are stored on non-local processes. Therefore, non-zeros in the off-process block of the matrix require vector values to be communicated during each SpMV.

The SpMV operation lacks parallel scalability due to large costs associated with communication, specifically in the strong scaling limit of a few rows per process. Increasing the number of processes that a matrix is distributed across increases the number of columns in the off-process blocks, yielding a growth in communication.

Fig. 2 shows the percentage of time spent communicating during a SpMV operation for two large matrices from the SuiteSparse matrix collection at scales varying from 50 000 to 500 000 non-zeros per process [12]. The results show that the communication time dominates the computation as the number of processes is increased, thus decreasing the scalability.

Machine topology plays an important role in the cost of communication [32]. Multicore distributed systems present new challenges in communication as the bandwidth is limited while the number of cores participating in communication increases [14]. Injection limits and network contention are significant roadblocks in the SpMV operation [5], motivating the need for SpMV algorithms that take advantage of the machine topology. The focus of the approach developed in this paper is to use the node-processor hierarchy to more efficiently map communication, leading to notable reductions in SpMV costs on modern HPC systems for a range of sparse matrix patterns. Throughout this paper, the term *node aware* refers to knowledge of the mapping of processes to physical nodes, although other aspects of the topology — e.g. socket information — could be used in a similar fashion. The mapping of virtual ranks to physical processors can be easily determined on many super computers. The flag `MPICH_RANK_REORDER_METHOD` can be set to a predetermined ordering on Cray machines, while modern Blue Gene machines allow the user to specify the ordering among the coordinates

☆ This research is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation, USA (awards OCI-0725070 and ACI-1238993) and the state of Illinois, USA. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications. This material is based in part upon work supported by the National Science Foundation, USA Graduate Research Fellowship Program under Grant Number DGE-1144245. This material is based in part upon work supported by the Department of Energy, USA, National Nuclear Security Administration, USA, under Award Number DE-NA0002374. Portions of this work are funded in part by ExxonMobil Upstream Research Company.

* Corresponding author.
E-mail addresses: bienz2@illinois.edu (A. Bienz), wgropp@illinois.edu (W.D. Gropp), lukeo@illinois.edu (L.N. Olson).

https://doi.org/10.1016/j.jpdc.2019.03.016
0743-7315/

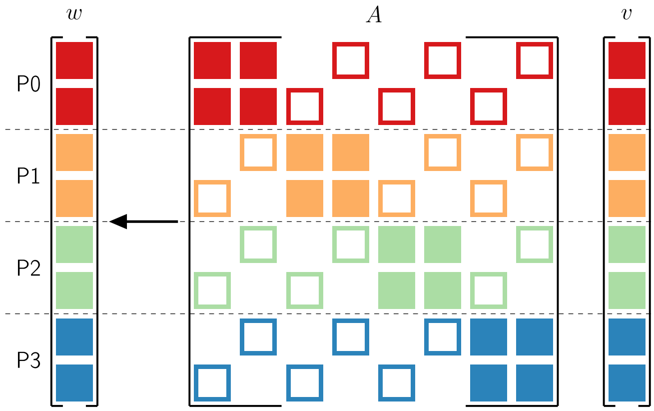

Fig. 1. A matrix partitioned across four processes, where each process stores two rows of the matrix, and the equivalent rows of each vector. The on-process block of each matrix partition is represented by solid squares, while the off-process block is represented by outlined entries.

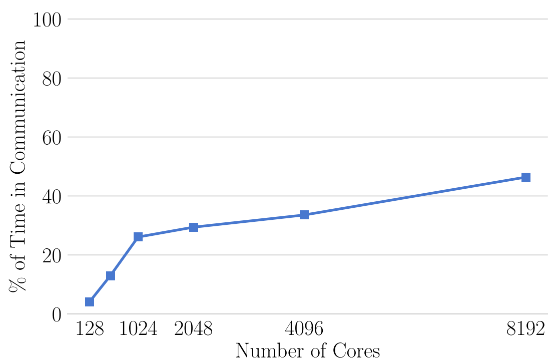

Fig. 2. Percentage of total SpMV time spent during communication for rotated anisotropic diffusion matrix with 10 240 000 rows, with rotation of $\frac{\pi}{4}$ and anisotropy of $\epsilon = 0.001$.

A, B, C, D, E, and T through the variable `RUNJOB_MAPPING` or a runscript option of `--mapping`.

There are a number of existing approaches for reducing communication costs associated with sparse matrix–vector multiplication. Communication volume in particular is a limiting factor and the ordering and parallel partition of a matrix both influence the total data volume. In response, graph partitioning techniques are used to identify more efficient layouts in the data [9,16,22,30]. ParMETIS [19] and PT-Scotch [10], for example, provide parallel partitioning of matrices that often lead to improved system loads and more efficient sparse matrix operations. Communication volume is accurately modeled through the use of a hypergraph [8]. As a result, hypergraph partitioning also leads to a reduction in parallel communication requirements, albeit at a larger one-time setup cost. Similarly, redistribution of data throughout linear solvers such as multigrid yield improved performance of matrix operations [23,33]. Topology-aware task mapping is used to accurately map partitions to the allocated nodes of a supercomputer, reducing the overall cost associated with communication
[1,17,21,27,29]. The approach introduced in this paper complements these efforts by providing an additional level of optimization in handling communication.

Topology-aware methods and aggregation of data are commonly used to reduce communication costs, particularly in collective operations [18,20,24,26]. Aggregation of data is used in point-to-point communication through Tram, a library for streamlining messages in which data is aggregated and communicated only through neighboring processors [31]. The method presented in this paper aggregates messages at the node level and communicates all aggregated data at once, yielding little structural change from standard MPI communication while reducing overall cost.

The performance of matrix operations can also be improved through the use of hybrid architectures and accelerators, such as graphics processing units (GPUs). The throughput of GPUs allows for improved performance when memory access patterns are optimized [3,11,15].

Many preconditioners for iterative methods, such as algebraic multigrid, rely on the SpMV as a dominant operation and therefore lack scalability due to large communication costs. A variety of methods exist for altering the preconditioning algorithms to reduce the communication costs associated with each SpMV [2,4,28].

This paper focuses on increasing the locality of communication during a SpMV to reduce the amount of communication injected into the network. Section 2 describes a reference algorithm for a parallel SpMV, as implemented in RAPtor [6], which resembles the approach commonly used in practice. A performance model is also introduced in Section 3, which considers the cost of intra- and inter-node communication and the impact on performance. A new SpMV algorithm is presented in Section 4, which reduces the number and size of inter-node messages by increasing the significantly cheaper intra-node communication. The code and numerics are presented in Section 5 to verify the performance. Finally, there are a variety of functions defined throughout this paper. These function names and definitions are listed in Table 1 for reference.

2. Background

Modern supercomputers incorporate a large number of nodes through an interconnect to form a multi-dimensional grid or torus network. Standard compute nodes are comprised of one or more multicore processors that share a large memory bank. The algorithm developed in this paper targets a general machine with this layout and the results are highlighted on Blue Waters, a Cray machine at the National Center for Supercomputing Applications. Blue Waters consists of 22 640 Cray XE nodes, each containing two AMD 6276 Interlagos processors for a total of 16 cores per node, and 4228 Cray XK nodes consisting of a single AMD processor along with an NVIDIA Kepler GPU.[1] The nodes are connected through a three-dimensional torus Gemini interconnect, with each Gemini serving two nodes. The remainder of this paper will focus on only the Cray XE nodes within Blue Waters.

Consider a system with n_p processes distributed across n_n nodes, resulting in $\texttt{ppn}$ processes per node. Rank $r \in \left[0, n_p - 1\right]$ is described by the tuple (p, n) where $0 \leq p < \texttt{ppn}$ is the local process number of rank r on node n. Assuming SMP-style ordering, the first $\texttt{ppn}$ ranks are mapped to the first node, the next $\texttt{ppn}$ to the second node, and so on. Therefore, rank r is described by the tuple $\left(r \mod \texttt{ppn}, \lfloor \frac{r}{\texttt{ppn}} \rfloor\right)$. Thus, for the remainder of the paper, the notation of rank r is interchangeable with (p, n).

Parallel matrices and vectors are distributed across all n_p ranks such that each process holds a portion of the linear system. Let $\texttt{Rows}(r)$ be the rows of an $N \times M$ sparse linear system, $w \leftarrow A \cdot v$, stored on rank r. In the case of an even, contiguous partition where the kth partition is placed on the kth rank, $\texttt{Rows}(r)$ is defined as

$$\texttt{Rows}(r) = \left\{ \left\lfloor \frac{N}{n_p} \right\rfloor r, \ldots, \left\lfloor \frac{N}{n_p} \right\rfloor (r+1) - 1 \right\} \tag{2}$$

[1] https://bluewaters.ncsa.illinois.edu/hardware-summary.

Table 1
List of functions defined throughout this paper.

Symbol	Definition
Rows(r)	Rows of linear system stored on rank r
Procs(r)	Processes to which rank r sends
Data(r, t)	Global vector indices that process r sends to process t
Nodes(n)	Set of nodes to which any process on node n sends
Node_Data(n, m)	Global vector indices to be sent from node n to node m
Send_Nodes((p, n))	Nodes to which process (p, n) sends
Recv_Nodes((p, n))	Nodes that sends to process (p, n)
Global_Procs((p, n))	Off-node processes to which (p, n) sends
Global_Data((p, n), (q, m))	Global vector indices that (p, n) sends to (q, m)
locality	Tuple describing locality of both the origin and destination of data
Local_Procs((p, n), locality)	Processes to which process (p, n) sends during communication step described by locality
Local_Data((p, n), (s, n), locality)	Global vector indices send from process (p, n) to (s, n) during communication step described by locality

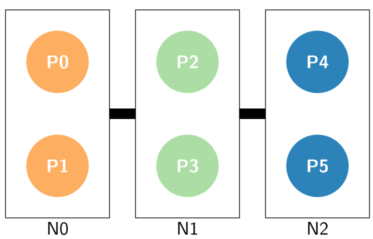

Fig. 3. An example parallel system with six processes distributed across three nodes.

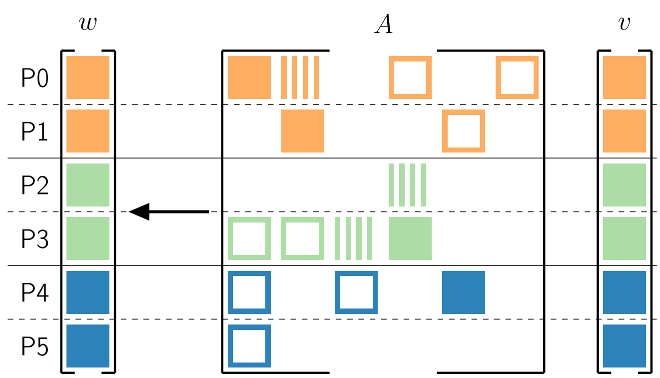

Fig. 4. An example 6 × 6 sparse matrix for the parallel system in Fig. 3. The striped shading denotes blocks that require only on-node communication, while the outlined shading denotes blocks that require communication with distant nodes.

or equivalently as

$$\text{Rows}((p,n)) = \left\{ \left\lfloor \frac{N}{n_p} \right\rfloor (p,n), \ldots, \left\lfloor \frac{N}{n_p} \right\rfloor ((p,n)+1) - 1 \right\}. \quad (3)$$

The rows of a matrix A are partitioned into on-process and off-process blocks, as described in Section 1. Accounting for parallel nodal awareness, the off-process block is further partitioned into on-node and off-node blocks, as described in Example 2.1.

Example 2.1. Suppose the parallel system consists of six processes distributed across three nodes, as displayed in Fig. 3. Let the linear system $w \leftarrow A \cdot v$ displayed in Fig. 4 be partitioned across this processor layout with each process holding a single row of the matrix and associated row of the input vector. In this example, the diagonal entry falls into the on-process block, as the corresponding vector value is stored locally. The off-process block, which requires communication, consists of all off-diagonal non-zeros as the associated vector values are stored on other processes.

For any process (p, n), the on-node columns of A correspond to vector values that are stored on some process (s, n), where $s \neq p$. Similarly, the off-node columns of A correspond to vector values stored on some process (q, m), where $m \neq n$. To make this clear, we define the following

$$\text{on\_process}(A,(p,n)) = \left\{ A_{ij} \neq 0 \mid i, j \in \text{Rows}((p,n)) \right\} \quad (4)$$

$$\text{off\_process}(A,(p,n)) = \left\{ A_{ij} \neq 0 \mid i \in \text{Rows}((p,n)), j \notin \text{Rows}((p,n)) \right\} \quad (5)$$

$$\text{on\_node}(A,(p,n)) = \left\{ A_{ij} \neq 0 \mid \exists q \neq p \text{ with } i \in \text{Rows}((p,n)), j \in \text{Rows}((q,n)) \right\} \quad (6)$$

and

$$\text{off\_node}(A,(p,n)) = \left\{ A_{ij} \neq 0 \mid \exists q, m \neq n \text{ with } i \in \text{Rows}((p,n)), j \in \text{Rows}((q,m)) \right\}. \quad (7)$$

2.1. Standard SpMV

For a sparse matrix–vector multiply, $w \leftarrow A \cdot v$, each process receives all values of v associated with the non-zero entries in the off-process block of A. For example, if rank r contains a non-zero entry of A, A_{ij}, at row i, column j, then rank s with row $j \in \text{Rows}(s)$ sends the jth vector value, v_j, to rank r. Typically, these communication requirements are determined as the sparse matrix is formed [7,13,25].

In the reference SpMV, for each rank r there is a list of processes to which data is sent, as well as the global vector indices to be sent to each. The function Procs(r) defines the list of processes to which a rank r sends. Specifically,

$$\text{Procs}(r) = \left\{ t \mid A_{ij} \neq 0 \text{ with } i \in \text{Rows}(t), j \in \text{Rows}(r), \ r \neq t \right\} \quad (8)$$

For each t in Procs(r), define the function Data(r, t) to return the global vector indices that process r sends to process t. This function is defined as follows.

$$\text{Data}(r,t) = \left\{ i \mid A_{ij} \neq 0 \text{ with } i \in \text{Rows}(t), j \in \text{Rows}(r), \ r \neq t \right\} \quad (9)$$

Consider a standard SpMV for the linear system described in Example 2.1. Table 2 lists the processes to which each rank must send, while Table 3 displays the indices that each rank r sends to any rank t.

With these definitions, the *standard* or reference SpMV is described in Algorithm 1. It is important to note that the parallel communication in Algorithm 1 is executed independent of any locality in the problem. That is, messages sent to another process may be both on-node or off-node depending on the process, however this is not considered in the algorithm.

Table 2
Communication pattern for rank r in Example 2.1, containing the values for $\texttt{Procs}(r)$.

	r					
	0	1	2	3	4	5
$\mathcal{P}(r)$	{3, 4, 5}	{0, 3}	{3, 4}	{0, 2}	{1}	{0}

Table 3
Each column r lists the indices of values sent to each process t in $\texttt{Procs}(r)$, namely $\texttt{Data}(r, s)$.

		r					
		0	1	2	3	4	5
t	0	{}	{1}	{}	{3}	{}	{5}
	1	{}	{}	{}	{}	{4}	{}
	2	{}	{}	{}	{3}	{}	{}
	3	{0}	{1}	{2}	{}	{}	{}
	4	{0}	{}	{2}	{}	{}	{}
	5	{0}	{}	{}	{}	{}	{}

Algorithm 1: `standard_spmv`: Reference Implementation

Input: r
$A|_{\texttt{Rows}(r)}$
$v|_{\texttt{Rows}(r)}$

Output: $w|_{\texttt{Rows}(r)}$

$A_{\text{on\_process}} = \texttt{on\_process}(A|_{\texttt{Rows}(r)})$
$A_{\text{off\_process}} = \texttt{off\_process}(A|_{\texttt{Rows}(r)})$
for $t \in \texttt{Procs}(r)$ **do**
  for $i \in \texttt{Data}(r, t)$ **do**
    $b_{\text{send}} \leftarrow v|_{\texttt{Rows}(r)_i}$
  $\texttt{MPI\_Isend}(b_{\text{send}}, \ldots, t, \ldots)$
$b_{\text{recv}} \leftarrow \emptyset$
for t s.t. $r \in \texttt{Procs}(t)$ **do**
  $\texttt{MPI\_Irecv}(b_{\text{recv}}, \ldots, t, \ldots)$
$\texttt{local\_spmv}(A_{\text{on\_process}}, v|_{\texttt{Rows}(r)})$
`MPI_Waitall`
$\texttt{local\_spmv}(A_{\text{off\_process}}, b_{\text{recv}})$

3. Communication models

The performance of Algorithm 1 is sub-optimal since it does not take advantage of node locality in the communication. To see this, a communication performance model is developed in this section. One approach is that of the *max-rate model* [14], which describes the communication time as

$$T = \alpha + \frac{\text{ppn} \cdot s}{\min(B_{\text{N}}, B_{\max} + (\text{ppn} - 1)B_{\text{inj}})}, \tag{10}$$

where α is the *latency* or start-up cost of a message, which may include preparing a message for transport or determining the network route; s is the number of bytes to be communicated; ppn is again the number of communicating processes per node; B_{inj} is the maximum rate at which messages are injected into the network; $B_{\max}$ is the achievable message rate of each process or *bandwidth*; and B_{N} is the peak rate of the network interface controller (NIC). All bandwidth rates, B_{inj}, $B_{\max}$, and B_{N}, are measured in bytes per second. In the simplest case of $\text{ppn} = 1$, the familiar *postal model* suffices:

$$T = \alpha + \frac{s}{B_{\max}}. \tag{11}$$

Table 4
Measurements for α, B_{inj}, $B_{\min}$, and B_{N} for Blue Waters, as published by Gropp et. al [14].

	α	B_{inj}	$B_{\max}$	B_{N}
Short	$4.0 \cdot 10^{-6}$	$6.3 \cdot 10^{8}$	$-1.8 \cdot 10^{7}$	∞
Eager	$1.1 \cdot 10^{-5}$	$1.7 \cdot 10^{9}$	$6.2 \cdot 10^{7}$	∞
Rend	$2.0 \cdot 10^{-5}$	$3.6 \cdot 10^{9}$	$6.1 \cdot 10^{8}$	$5.5 \cdot 10^{9}$

Table 5
Measurements for intra-node variables, α_ℓ and $B_{\max_\ell}$.

	α_ℓ	$B_{\max_\ell}$
Short	$1.3 \cdot 10^{-6}$	$4.2 \cdot 10^{8}$
Eager	$1.6 \cdot 10^{-6}$	$7.4 \cdot 10^{8}$
Rend	$4.2 \cdot 10^{-6}$	$3.1 \cdot 10^{9}$

MPI contains multiple message passing protocols, including short, eager, and rendezvous. Each message consists of an envelope, including information about the message such as message size and source information, as well as message data. Short messages contain very little data which is sent as part of the envelope. Eager and rendezvous messages, however, send the envelope followed by packets of data. Eager messages are sent under the assumption that the receiving process has buffer space available to store data that is communicated. Therefore, a message is sent without checking buffer space at the receiving process, limiting the associated latency. However, if a message is sufficiently large, rendezvous protocol must be used. This protocol requires the sending process to inform the receiving rank of the message so that buffer space is allocated. The message is sent only once the sending process is informed that this space is available. Therefore, there is a larger overhead with sending a message using rendezvous protocol. Table 4 displays the measurements for α, B_{inj}, $B_{\max}$, and B_{N} for Blue Waters, as determined for the *max-rate* model.

The *max-rate* model can be improved by distinguishing between intra- and inter-node communication. If the sending and receiving processes lie on the same physical node, data is not injected into the network, yielding low start-up and byte transport costs. As intra-node messages are not injected into the network, communication local to a node can be modeled as

$$T_\ell = \alpha_\ell + \frac{s_\ell}{B_{\max_\ell}}, \tag{12}$$

where α_ℓ is the start-up cost for intra-node messages; s_ℓ is the number of bytes to be transported; and $B_{\max_\ell}$ is the achievable intra-node message rate.

Nodecomm,[2] a topology-aware communication program, measures the time required to communicate on various levels of the parallel system, such as between two nodes of varying distances and between processes local to a node. Communication tests between processes local to one node were used to calculate the intra-node model parameters, as displayed in Table 5. These parameters were calculated in python with a least squares data fitting.

Furthermore, Fig. 5 shows the time required to send a single message of varying sizes between two processes. The thin lines display Nodecomm measurements for time required to send a single message, on average, as either inter- or intra-node communication. Furthermore, these timings test various numbers of processes-per-node active in communication, from one to sixteen. Finally, the thick lines represent the time required to send a message of each size, according to the *max-rate* model in (10) and

2 See https://bitbucket.org/william_gropp/baseenv.

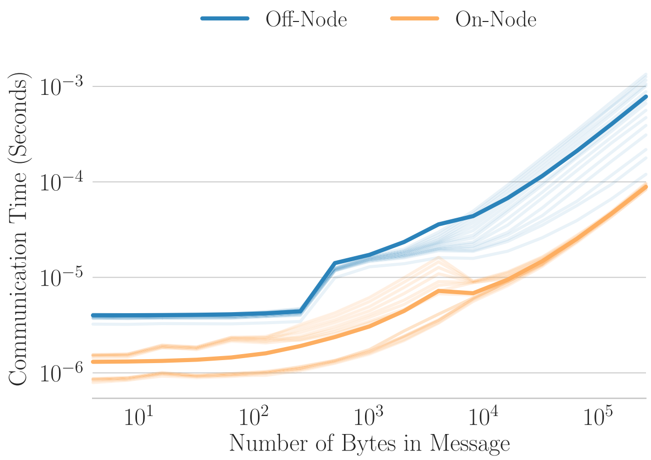

Fig. 5. The time required to send a single message of various sizes, with the thin lines representing timings measured by Nodecomm and the thick lines displaying the *max-rate* and intra-node models in (10) and (12), respectively.

intra-node model in (12). This figure displays a significant difference between the costs of intra- and inter-node communication.

4. Node aware parallel SpMV

To reduce communication costs, the algorithm proposed in this section decreases the number and size of messages being injected into the network by increasing the amount of intra-node communication, which is less-costly than inter-node communication. This trade-off is accomplished through a so-called *node aware parallel* SpMV (NAPSpMV), where values are gathered in processes local to each node before being sent across the network, followed by a distribution of processes on the receiving node. As a result, as the matrix is formed each process (p, n) determines the communicating processes during the various steps of a NAPSpMV, as well as the accompanying data. A high level overview of the process is described in Example 4.1. It is important to note that the communication for each NAPSpMV is load-balanced such that all processes local to node n send and receive both a similar number and size of messages through inter-node communication. Therefore, it is assumed that the nodes n and m in Example 4.1 are only a portion of the parallel system, and n is communicating with other nodes in a similar fashion.

Example 4.1. During each NAPSpMV, off-node data is communicated through a three-step process, as displayed in Fig. 6. This figure displays a portion of a parallel system consisting of 8 processes partitioned across two nodes, labeled n and m. The solid circles on node n represent vector values that must be sent to node m. Therefore, each process on node n must send values to processes on node m. Instead of sending directly to destination processes, each process (s, n) sends to the process labeled (p, n), displayed by the dashed arrows on node n. Process (p, n) then sends all collected values through the network to process (q, m). Finally, process (q, m) distributes received values among the processes local to node m, displayed by the dashed arrows on node m.

On-node data is communicated directly between the process on which the vector values are stored and that which requires the data, as displayed in Fig. 7. In this example, the solid circles represent vector values that are stored on each process (s, n) and needed by (p, n). This data is sent directly between the processes in a single step.

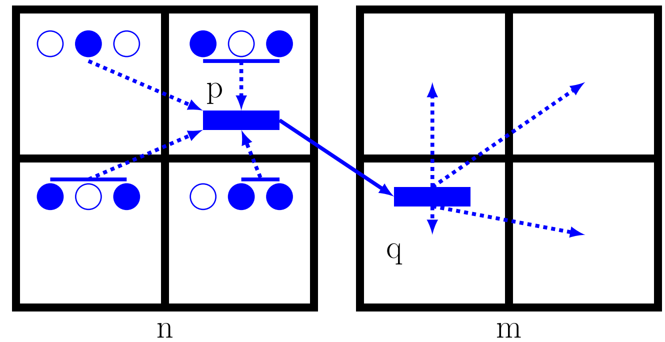

Fig. 6. The various arrows exemplify the process of communicating data from each process on node n to processes on node m through a three-step algorithm. The bold circles on node n represent vector values that must be communicated to node m.

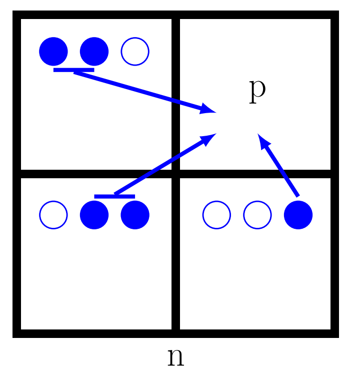

Fig. 7. An example of how vector values corresponding with matrix entries in the on-node block are communicated. All values (p, n) must receive from other processes (q, n) are communicated directly as nothing is injected into the network.

Table 6
Communication requirements for each node n in Example 2.1.

	n		
	0	1	2
Nodes(n)	{1, 2}	{0, 2}	{0}

4.1. Inter-node communication setup

To eliminate the communication of duplicated messages, a list of communicating nodes is formed for each node n along with the accompanying data values. These lists are then distributed across all processes local to n by balancing the number of nodes and volume of data for communication. To facilitate this, the function Nodes(n) defines the set of nodes to which the processes on node n must send,

$$\texttt{Nodes}(n) = \Big\{m \mid \exists\, p, q \text{ s.t. } A_{ij} \neq 0 \quad \text{with } i \in \texttt{Rows}((q, m)), j \in \texttt{Rows}((p, n)), n \neq m\Big\}. \tag{13}$$

Table 6 contains Nodes(n), the list of nodes to which each node n sends. The associated data values are defined for each node $m \in$ Nodes(n) with Node_Data(n, m), which returns the data indices to be sent from node n to node m. That is,

$$\texttt{Node\_Data}(n, m) = \Big\{i \mid \exists\, p, q \text{ s.t. } A_{ij} \neq 0 \quad \text{with } i \in \texttt{Rows}((q, m)), \quad j \in \texttt{Rows}((p, n)), n \neq m\Big\}. \tag{14}$$

Table 7
In Example 2.1, each column n contains the values sent from n to m, as in $\texttt{Node\_Data}(n, m)$.

		n		
		0	1	2
m	0	{}	{3}	{4, 5}
	1	{0, 1}	{}	{}
	2	{0}	{2}	{}

Table 8
Processor mappings for $\texttt{Nodes}(n)$, namely $\texttt{Send\_Nodes}((p, n))$ and $\texttt{Recv\_Nodes}((p, n))$ for Example 2.1.

	(p, n)					
	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
$\texttt{Send\_Nodes}((p, n))$	{1}	{2}	{0}	{2}	{0}	{}
$\texttt{Recv\_Nodes}((p, n))$	{2}	{1}	{}	{0}	{1}	{0}

Table 9
Inter-node communication requirements of each process (p, n) for Example 2.1.

	(p, n)					
	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
$\texttt{Global\_Procs}((p, n))$	{(1, 1)}	{(1, 2)}	{(1, 0)}	{(0, 2)}	{(0, 0)}	{}

Extending Example 2.1, Table 7 displays the global vector indices, $\texttt{Node\_Data}(n, m)$, for each set of nodes n and m.

$\texttt{Send\_Nodes}((p, n))$ defines the nodes to which (p, n) must send, that is the nodes in $\texttt{Nodes}(n)$ that are distributed to process (p, n). Similarly, $\texttt{Recv\_Nodes}((p, n))$ contains the nodes that send to (p, n). Specifically,

$$\texttt{Send\_Nodes}((p, n)) = \{m \in \texttt{Nodes}(n) \mid m \text{ maps to } (p, n)\}\,, \tag{15}$$

$$\texttt{Recv\_Nodes}((p, n)) = \{m \mid n \in \texttt{Nodes}(m),\ n \text{ maps to } (p, n)\}\,. \tag{16}$$

This paper considers a simple distribution in which the node $m \in \texttt{Nodes}(n)$ to which the most data $|\texttt{Data}(n, m)|$ is sent is mapped to process $(0, n)$, the node with the second most data is mapped to process $(1, n)$, and so on. The opposite ordering is used for $\texttt{Recv\_Nodes}((p, n))$, mapping the node $n \in \texttt{Nodes}(m)$ with largest $|\texttt{Data}(m, n)|$ to process $(\text{ppn} - 1, n)$, the second largest to process $(\text{ppn}-2, n)$, etc. If there are fewer nodes in $\texttt{Nodes}(n)$ than there are processes per node, a single node is mapped to multiple local processes so that all processes communicate, such that all processes send a similar number of bytes. There are various other possible mapping strategies, such as mapping a node m to the process (p, n) storing the majority of the data in $\texttt{Data}(n, m)$. However, as this would only affect intra-node communication requirements, these mappings are not explored in this paper.

The processor layout in Example 2.1 is displayed in Table 8, where the columns contain the send and receive nodes that are mapped to each process.

Finally, $\texttt{Global\_Procs}((p, n))$ defines the set of all off-node processes to which process (p, n) sends data during the inter-node communication step of the NAPSpMV. Specifically,

$$\begin{aligned}\texttt{Global\_Procs}((p, n)) = \{(q, m) \mid\ & m \in \texttt{Send\_Nodes}((p, n)),\\ & n \in \texttt{Recv\_Nodes}((q, m))\}\,.\end{aligned} \tag{17}$$

Following Example 2.1, the columns of Table 9 list the indices of the values that each (p, n) sends, $\texttt{Global\_Procs}((p, n))$. Finally, let $\texttt{Global\_Data}((p, n), (q, m))$ define the global data indices corresponding to the values sent from process (p, n) to (q, m):

$$\begin{aligned}&\texttt{Global\_Data}((p, n), (q, m)) =\\ &\quad \{\texttt{Node\_Data}(n, m) \mid m \in \texttt{Send\_Nodes}((p, n)),\\ &\quad n \in \texttt{Recv\_Nodes}((q, m))\}\end{aligned} \tag{18}$$

Table 10
Inter-node communication requirements for each set of processes (p, n) and (q, m). Each column (p, n) contains the indices of values sent from (p, n) to (q, m).

		(p, n)					
		(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
(q, m)	(0, 0)	{}	{}	{}	{}	{4, 5}	{}
	(1, 0)	{}	{}	{3}	{}	{}	{}
	(0, 1)	{}	{}	{}	{}	{}	{}
	(1, 1)	{0}	{}	{}	{}	{}	{}
	(0, 2)	{}	{}	{}	{2}	{}	{}
	(1, 2)	{}	{0, 1}	{}	{}	{}	{}

The global vector indices to which each process (p, n) sends and receives for Example 2.1 are displayed in Table 10.

4.2. Local communication

The function $\texttt{Global\_Procs}_{\text{send}}((p, n))$ for $p = 0, \ldots, \text{ppn} - 1$, describes evenly distributed inter-node communication requirements for all processes local to node n. However, many of the vector indices to be sent to off-node process $(q, m) \in \texttt{Data}((p, n), (q, m))$, are not stored on process (p, n). For instance, in Table 10, process (0, 1) sends global vector indices 0 and 1. However, only row 1 is stored on process (0, 1), requiring vector component 0 to be communicated before inter-node messages are sent.

Similarly, many of the indices that a process (q, m) receives from (p, n) are redistributed to various processes on node n. Table 10 requires process (1, 2) to receive vector data according to indices 0 and 1. Process (0, 2) uses both of these vector values, yielding a requirement for redistribution of data received from inter-node communication. Therefore, local communication requirements must be defined.

Each NAPSpMV consists of multiple steps of intra-node communication. Let a function $\texttt{Local\_Procs}((p, n), \texttt{locality})$ define all processes, local to node n, to which process (p, n) sends messages, where $\texttt{locality}$ is a tuple describing the locality of both the original location of the data as well as its final destination. The locality of each position is described as either $\texttt{on\_node}$, meaning a process local to node n, or $\texttt{off\_node}$, meaning a process local to node $m \neq n$.

There are three possible combinations for $\texttt{locality}$: 1. the data is initialized $\texttt{on\_node}$ with a final destination $\texttt{off\_node}$; 2. the original data is $\texttt{off\_node}$ while the final destination is $\texttt{on\_node}$; or 3. both the original data and the final location are $\texttt{on\_node}$. These three types of intra-node communication are described in more detail in the remainder of Section 4.2.

For each process $(s, n) \in \texttt{Local\_Procs}((p, n), \texttt{locality})$, $\texttt{Local\_Data}((p, n), (s, n), \texttt{locality})$ defines the global vector indices to be sent from process (p, n) to (s, n) through intra-node communication. This notation is used in following sections.

4.2.1. Local redistribution of initial data

During inter-node communication, a process (p, n) sends all vector values corresponding to the global indices in $\texttt{Global\_Data}((p, n), (q, m))$ to each process $(q, m) \in \texttt{Global\_Procs}((p, n))$. The indices in $\texttt{Global\_Data}((p, n), (q, m))$ originate on node n, but not necessarily process (p, n). Therefore, the initial vector values must be redistributed among all processes local to node n.

Let $\texttt{Local\_Procs}((p, n), (\texttt{on\_node}, \texttt{off\_node}))$ represent all processes, local to node n, to which (p, n) sends initial vector

Table 11
Initial intra-node communication requirements for each process (p, n) in Example 2.1. The row of the table describes Local_Procs((p, n), (on_node, off_node)).

	(p, n)					
	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
Local_Procs	{(1, 0)}	{}	{(1, 1)}	{(0, 1)}	{}	{(0, 2)}

Table 12
Global vector indices of initial data that is communicated between processes local to each node n in Example 2.1. Each column contains the indices of values sent from (p, n) to (q, n). Note: dashes (—) throughout the table represent processes on separate nodes, which do not communicate during intra-node communication.

		(p, n)					
		(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
(q, n)	(0, 0)	{}	{}	–	–	–	–
	(1, 0)	{0}	{}	–	–	–	–
	(0, 1)	–	–	{}	{3}	–	–
	(1, 1)	–	–	{2}	{}	–	–
	(0, 2)	–	–	–	–	{}	{5}
	(1, 2)	–	–	–	–	{}	{}

Table 13
Intra-node communication requirements containing processes to which each (p, n) sends received inter-node data, according to Example 2.1. The row of the table describes Local_Procs((p, n), (off_node, on_node)).

	(p, n)					
	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
Local_Procs	{}	{(0, 0)}	{}	{}	{}	{(0, 2)}

values. This function is defined as

$$\texttt{Local\_Procs}((p,n),(\texttt{on\_node},\texttt{off\_node})) = \{(s,n) \mid \exists j \in \texttt{Rows}((p,n)),\ j \in \texttt{Global\_Data}((s,n),(q,m))\}\,. \tag{19}$$

The local processes to which each (p, n) sends initial data in Example 2.1 are displayed in Table 11.

Furthermore, the data global vector indices that must be sent from process (p, n) to each $(s, n) \in$ Local_Procs((p, n), (on_node, off_node)) are defined as

$$\texttt{Local\_Data}((p,n),(s,n),(\texttt{on\_node},\texttt{off\_node})) = \{i \mid i \in \texttt{Rows}((p,n)), \forall i \in \texttt{Global\_Procs}((s,n))\}\,. \tag{20}$$

The global vector indices that each (p, n) must send to other processes on node n in Example 2.1 are displayed in Table 12.

4.2.2. Local redistribution of received off-node data

During inter-node communication, a process (p, n) sends all data with final destination on node m to process $(q, m) \in$ Global_Procs((p, n)). Process (q, m) then distributes these values across the processes local to node m. Let Local_Procs ((q, m), (off_node, on_node)) define all processes local to node m to which process (q, m) sends vector values that have been received through inter-node communication. This function is defined as

$$\texttt{Local\_Procs}((q,m),(\texttt{off\_node},\texttt{on\_node})) = \left\{(s,m) \mid \exists A_{ij} \neq 0 \text{ with } i \in \texttt{Rows}((s,m)),\ j \in \texttt{Global\_Data}((p,n),(q,m))\right\}. \tag{21}$$

This is highlighted, for Example 2.1, in Table 13. Furthermore, the data global vector indices that must be sent from process (q, m) to each $(s, m) \in$ Local_Procs((q, m), (off_node, on_node)) are defined as

$$\texttt{Local\_Data}((q,m),(s,m),(\texttt{off\_node},\texttt{on\_node})) = \left\{j \in \texttt{Global\_Data}((p,n),(q,m)) \mid A_{ij} \neq 0 \text{ with } i \in \texttt{Rows}((s,m))\right\}. \tag{22}$$

The global vector indices, received from the inter-node communication step, which (p, n) must send to each local process (q, n) in Example 2.1 are displayed in Table 14.

Table 14
Global vector indices of received inter-node data that must be communicated between processes local to each node n in Example 2.1. Each column contains the indices of values sent from (p, n) to (q, n). Note: dashes (—) throughout the table represent processes on separate nodes, which cannot communicate during intra-node communication.

		(p, n)					
		(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
(q, n)	(0, 0)	{}	{3}	–	–	–	–
	(1, 0)	{}	{}	–	–	–	–
	(0, 1)	–	–	{}	{}	–	–
	(1, 1)	–	–	{}	{}	–	–
	(0, 2)	–	–	–	–	{}	{1}
	(1, 2)	–	–	–	–	{}	{}

Table 15
Intra-node communication requirements containing processes to which each process (p, n) must send vector data, according to Example 2.1. The row of the table describes Local_Procs((p, n), (on_node, on_node)).

	(p, n)					
	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
Local_Procs	{}	{(0, 0)}	{}	{(0, 1)}	{}	{}

Table 16
Global vector indices that must be communicated between processes local to each node n in Example 2.1. Each column contains the indices of values sent from (p, n) to (q, n). Note: dashes (—) throughout the table represent processes on separate nodes, which cannot communicate during intra-node communication.

		(p, n)					
		(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
(s, n)	(0, 0)	{}	{1}	–	–	–	–
	(1, 0)	{}	{}	–	–	–	–
	(0, 1)	–	–	{}	{3}	–	–
	(1, 1)	–	–	{}	{}	–	–
	(0, 2)	–	–	–	–	{}	{}
	(1, 2)	–	–	–	–	{}	{}

4.2.3. Fully local communication

A subset of the values needed by a process (p, n) are stored on local process (s, n). One advantage is that these values bypass the three-step communication, and are communicated directly. Let Local_Procs((p, n), (on_node, on_node)) define all processes local to node n to which (p, n) sends vector data. This function is defined as

$$\texttt{Local\_Procs}((p,n),(\texttt{on\_node},\texttt{on\_node})) = \left\{(s,n) \mid \exists A_{ij} \neq 0 \text{ with } i \in \texttt{Rows}((s,n)), j \in \texttt{Rows}((p,n))\right\}. \tag{23}$$

The processes local to node n, to which (p, n) must send initial vector data in Example 2.1 are displayed in Table 15. Furthermore, the global vector indices that must be sent from process (p, n) to each $(s, n) \in$ Local_Procs((p, n), (on_node, on_node)) are defined as follows.

$$\texttt{Local\_Data}((p,n),(s,n),(\texttt{on\_node},\texttt{on\_node})) = \left\{j \mid \exists A_{ij} \neq 0 \text{ with } i \in \texttt{Rows}((s,n)), j \in \texttt{Rows}((p,n))\right\}. \tag{24}$$

The global vector indices which (p, n) must send to each local process (s, n) in Example 2.1 are displayed in Table 16.

4.3. Alternative SpMV algorithm

Algorithm 2: `local_comm`

Input: (p, n) : tuple describing local rank and node of process
$v|_{\texttt{Rows}((p,n))}$: rows of input vector v local to process (p, n)
`locality`: locality of input and output data

Output: ℓ_{recv}: values that rank (p, n) receives from other processes

```
// Initialize sends
for (s, n) ∈ Local_Procs((p, n), locality) do
    for i ∈ Local_Data((p, n), (s, n), locality) do
        ℓ_send ← v|Rows((p,n))_i
    MPI_Isend(ℓ_send, ..., (s, n), ...)

// Initialize receives
ℓ_recv ← ∅
for (s, n) s.t. (p, n) ∈ Local_Procs((s, n), locality) do
    MPI_Irecv(ℓ_recv, ..., (s, n), ...)

// Complete sends and receives
MPI_Waitall
```

The method of communicating vector values to on-node processes is described in Algorithm 2. Using the definitions for the various steps of intra- and inter-node communication, the NAPSpMV is described in Algorithm 3, where `local_spmv()` refers to a row-wise, non-distributed SpMV — e.g. with Intel's MKL library or with the Eigen Library. It is important to note that many slight variations to the algorithm are possible. The fully local communication has no dependencies, and can be performed anytime before calling `local_spmv`$(A_{\text{on\_node}}, b_{\ell\to\ell})$. Furthermore, the function `local_spmv`$(A_{\text{on\_process}}, v|\texttt{Rows})$ has no communication requirements and, hence, can be performed at any point in the algorithm.

Algorithm 3: `NAPSpMV`

Input: (p, n): tuple describing local rank and node of process
$A|R$: rows of matrix A local to process (p, n)
$v|R$: rows of input vector v local to process (p, n)

Output: $w|\texttt{Rows}$: rows of output vector $w \leftarrow Av$, local to process (p, n)

```
A_on_process = on_process(A|Rows)
A_on_node = on_node(A|Rows)
A_off_node = off_node(A|Rows)

b_ℓ→ℓ ← local_comm((p, n), v|Rows, (on_node → on_node))
b_ℓ→nℓ ← local_comm((p, n), v|Rows, (on_node → off_node))

// Initialize sends
for (q, m) ∈ Global_Procs((p, n)) do
    for i ∈ Global_Data((p, n), (q, m)) do
        g_send ← b^i_ℓ→nℓ
    MPI_Isend(g_send, ..., (q, m), ...)

// Initialize receives
g_recv ← ∅
for (q, m) s.t. (p, n) ∈ Global_Procs((q, m)) do
    MPI_Irecv(g_recv, ..., (q, m), ...)

// Serial SpMV for local values
local_spmv(A_on_process, v|Rows)

// Serial SpMv for on-node values
local_spmv(A_on_node, b_ℓ→ℓ)

// Complete sends and receives
MPI_Waitall

b_nℓ→ℓ ← local_comm((p, n), v|Rows, (off_node → on_node))

// Serial SpMV for off-node values
local_spmv(A_off_node, b_nℓ→ℓ)
```

5. Results

In this section, the parallel performance and scalability of the NAPSpMV in comparison to the standard SpMV is presented. The matrix–vector multiplication in an algebraic multigrid (AMG) hierarchy is tested for both a structured 2D rotated anisotropic and for unstructured linear elasticity on 32 768 processes in order to expose a variety of communication patterns. In addition, scaling tests are considered for random matrices with a constant number of non-zeros per row to investigate problems with no structure. Lastly, scaling tests on the largest 15 matrices from the SuiteSparse matrix collection are presented. All tests are performed on the Blue Waters parallel computer at University of Illinois at Urbana-Champaign.

AMG hierarchies consist of successively coarser, but denser levels. Therefore, while a standard SpMV performed on the original matrix often requires communication of a small number of large messages, coarse levels require a large number of small messages to be injected into the network. Fig. 8 shows that both the number and size of inter-node messages required on each level of the linear elasticity hierarchy are reduced through use of the NAPSpMV. There is a large reduction in communication requirements for coarse levels of the hierarchy, which includes a high number of small messages. However, as the NAPSpMV requires redistribution of data among processes local to each node, the intra-node communication requirements increase greatly for the NAPSpMV, as shown in Fig. 9. The ratio of inter- to intra-node communication is shown in Fig. 10. This compares both the number and size of messages communicated on-node and through the network.

While there is an increase in intra-node communication requirements, the reduction in more expensive inter-node messages results in a significant reduction in total time for the NAPSpMV algorithm, particularly on coarser levels near the middle of each AMG hierarchy, as shown in Fig. 11.

The reduction in cost associated with communication, particularly on the coarse levels of the hierarchies, yields increased performance and scalability of the AMG solve phase, as displayed in Fig. 12. The NAPSpMV reduces the cost of a single V-cycle of the strongly scaled linear elasticity hierarchy and greatly improves the strong scaling limit.

Random matrices, formed with a constant number of non-zeros per row, lack structure that is found in many finite element discretizations. As these matrices are distributed across an increasingly large number of processes, non-zeros are more likely

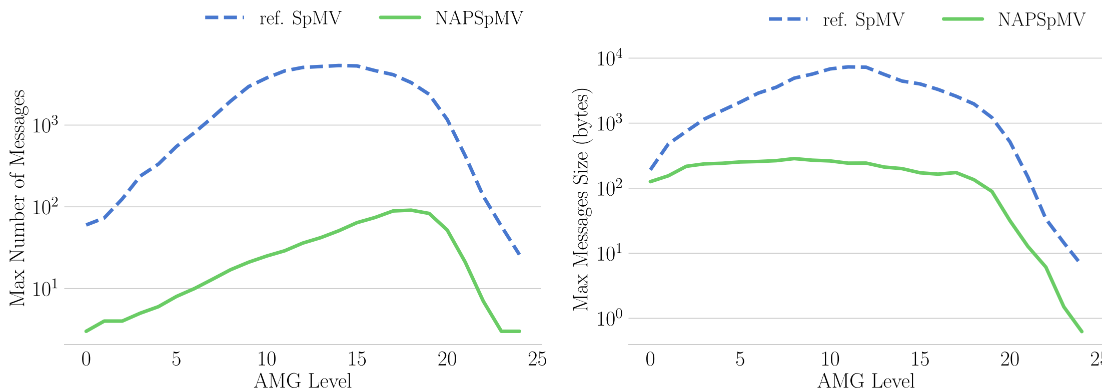

Fig. 8. The maximum number (left) and size (right) of **inter-node** messages communicated by a single process during a standard SpMV and NAPSpMV on each level of the **linear elasticity AMG hierarchy**.

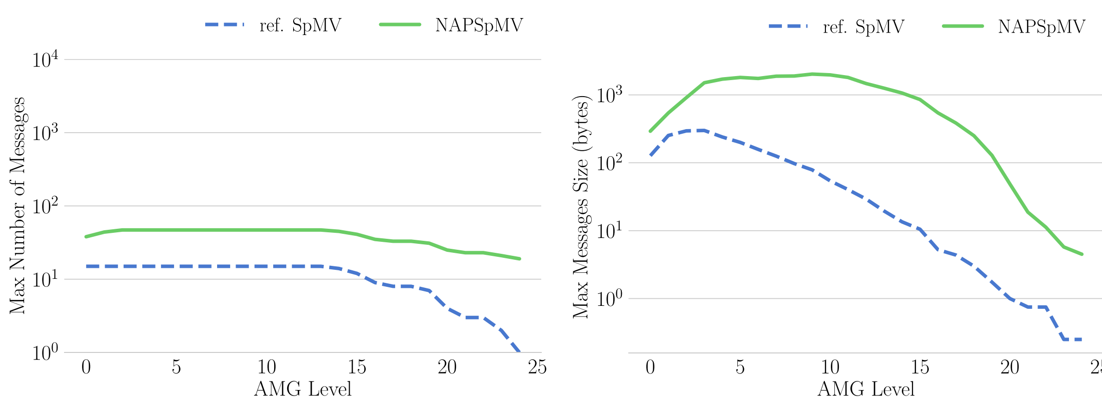

Fig. 9. The maximum number (left) and size (right) of **intra-node** messages communicated by a single process during a standard SpMV and NAPSpMV on each level of the **linear elasticity AMG hierarchy**.

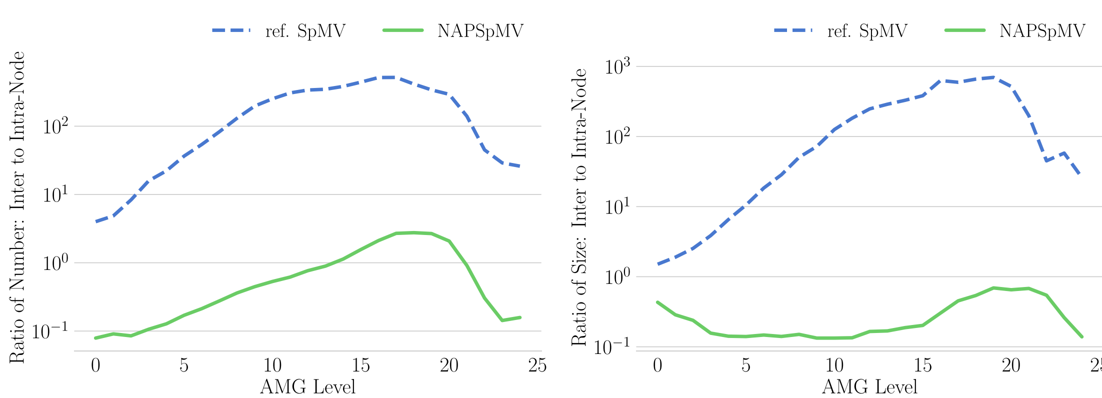

Fig. 10. The ratio of inter- to intra-node messages, with number of messages (left) and size of communicated data (right) during a standard SpMV and NAPSpMV on each level of the **linear elasticity AMG hierarchy**.

to be located in off-process blocks of the matrix. Therefore, both weak and strong scaling studies of random matrices yield increases in communication requirements with scale. The sparsity pattern of random matrices varies with random number generator seeds and are dependent on the number of non-zeros per row. Therefore, the standard SpMV and NAPSpMV were performed on five different random matrices for each tested density of 25, 50, and 100 non-zeros per row, as shown in Fig. 13. The standard and NAPSpMV costs for all random matrices of equivalent density are comparable. Furthermore, there is little difference in costs between each density. Therefore, extended tests are performed on only a single random matrix with 100 non-zeros per row. Fig. 14 displays the time required for a NAPSpMV in comparison to the standard SpMV in both weak and strong scaling studies. For these random matrices, the NAPSpMV exhibits improved performance over the reference implementation by up to two orders of magnitude and also improves scalability.

The time required to perform the various SpMVs on 13 of the 15 largest matrices from the SuiteSparse matrix collection are shown with strided and balanced partitions, in Figs. 15 and 16

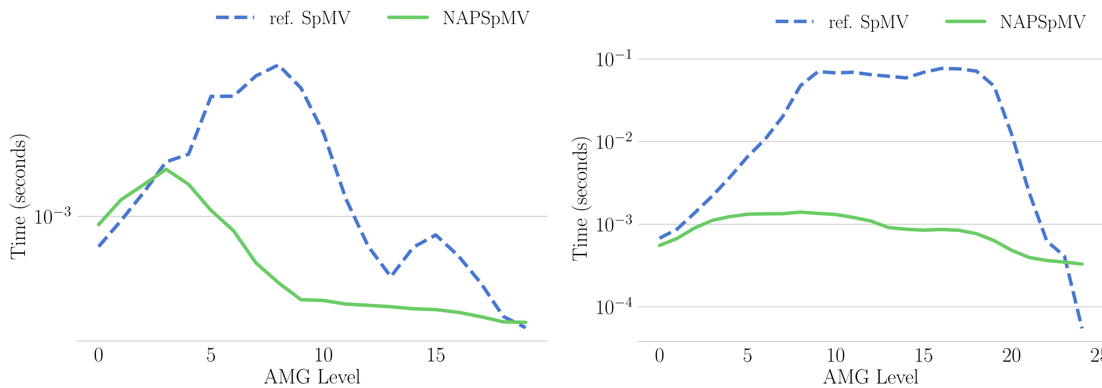

Fig. 11. The time required to perform the various SpMVs on each level of the rotated anisotropic (left) and linear elasticity (right) AMG hierarchies.

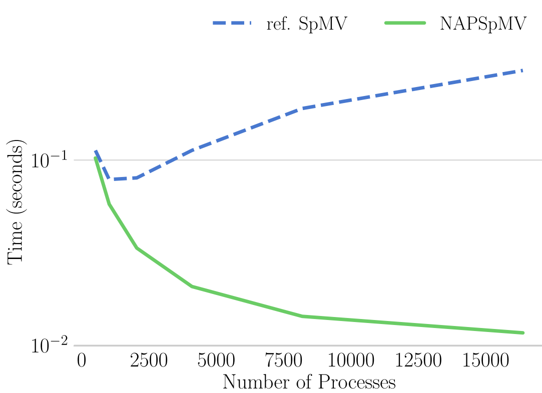

Fig. 12. The time required to perform a single V-cycle of AMG for the linear elasticity hierarchy at various scales, using a standard SpMV and NAPSpMV in every operation.

respectively. The remaining 2 large matrices were not included due to partitioning constraints. For the strided partitions with n_p processes, each row r is local to process $p = r \mod n_p$. Some matrices in this subset have nearly dense blocks of rows, allowing for improved load balancing over each process holding a contiguous block of rows. The balanced partitions were formed with PT Scotch graph partitioning, using the strategy `SCOTCH_STRATBALANCE`.

The NAPSpMV improves upon many of the matrices with strided partitions, as communication patterns are far from optimal, while only minimally improving upon the graph partitioned matrices as expected. However, the cost of partitioning motivates the use of less optimal partitions when a smaller number of SpMVs are anticipated. Fig. 17 shows the time required to perform various numbers of NAPSpMVs on both the strided and balanced partitions at the strongest scale tested, with 50 000 non-zeros per core. In these tests, the balanced partitioned timings include the time required to graph partition and redistribute the matrix. The crossover point for the various SuiteSparse matrices, at which the graph partitioning becomes less costly than performing NAPSpMVs on strided partitions, occurs only after hundreds, or often thousands, of SpMVs have been performed.

6. Conclusion and future work

This paper introduces a method to reduce communication that is injected into the network during a sparse matrix–vector multiply by reorganizing messages on each node. This results in a reduction of the inter-node communication, replaced by less-costly intra-node communication, which reduces both the number and size of messages that are injected into the network. The current implementation could be extended to take various levels of the hierarchy into account, such as splitting intra-node messages into on-socket and off-socket. Fig. 18 shows that on-socket messages are significantly cheaper and could be targeted to further reduce communication costs. Furthermore, node-aware communication could be applied to other instances of irregular point-to-point communication, in both alternative matrix

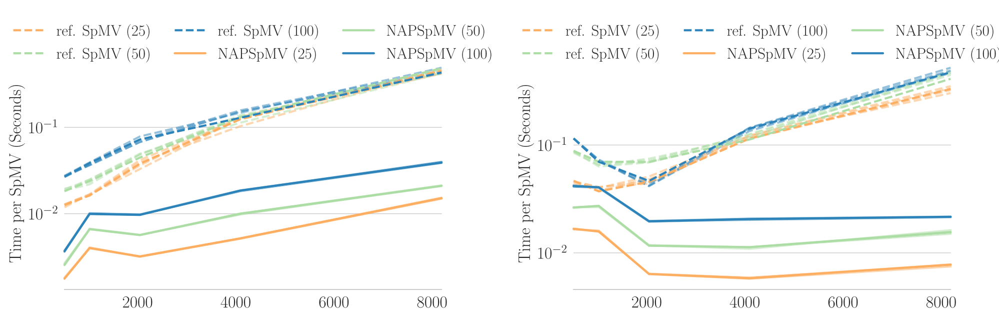

Fig. 13. The time required to perform the various SpMVs on weakly (left) and strongly (right) scaled random matrices. Five different random matrices are tested for each density of 25, 50, and 100 non-zeros per row. The weak-scaling study tests matrices with 1000 rows per process, while the strongly-scaled matrix contains 4 096 000 rows.

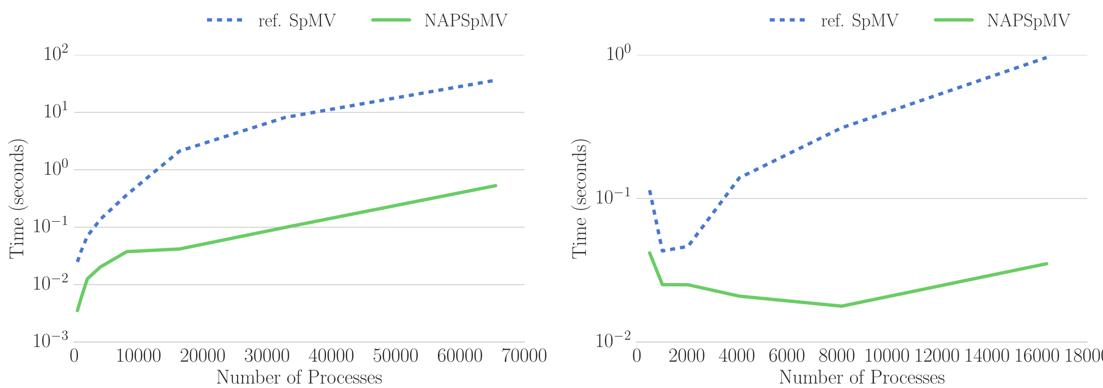

Fig. 14. The time required to perform the various SpMVs on weakly (left) and strongly (right) scaled random matrices, each with 100 non-zeros per row. The weak-scaling study tests matrices with 1000 rows per process, while the strongly-scaled matrix contains 4 096 000 rows.

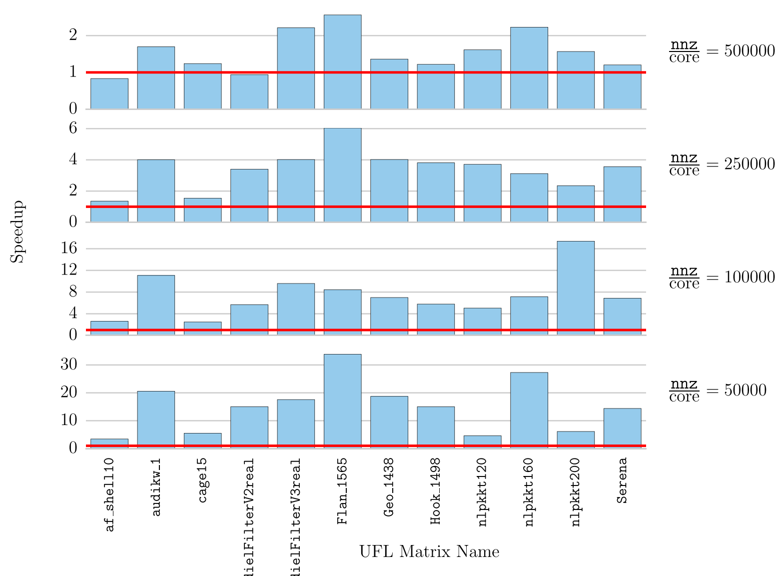

Fig. 15. The speedup of NAPSpMVs over reference SpMVs on a subset of the largest real matrices from the SuiteSparse matrix collection at various scales, where $\frac{nnz}{\text{core}}$ is the average number of non-zeros per core, partitioned so that each row r is stored on process $p = r \mod n_p$, where n_p is the number of processes.

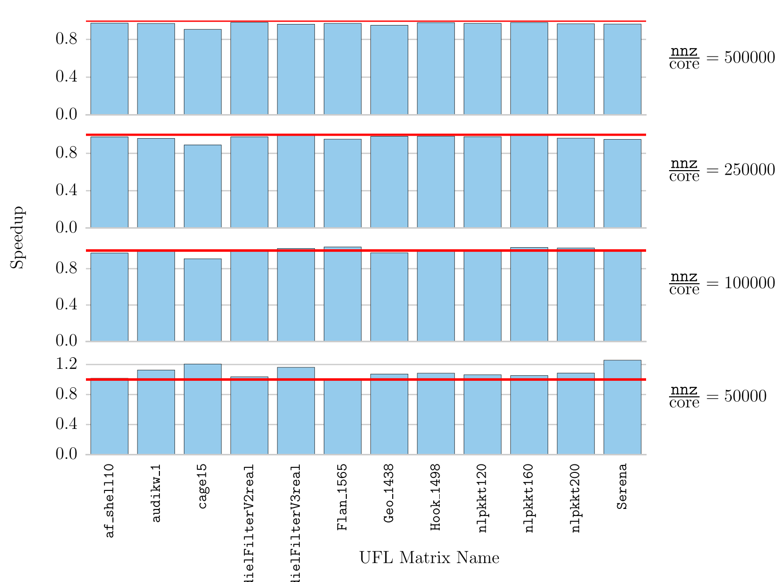

Fig. 16. The speedup of NAPSpMVs over reference SpMVs on a subset of the largest real matrices from the SuiteSparse matrix collection at various scales, where $\frac{nnz}{\text{core}}$ is the average number of non-zeros per core, partitioned with PT Scotch.

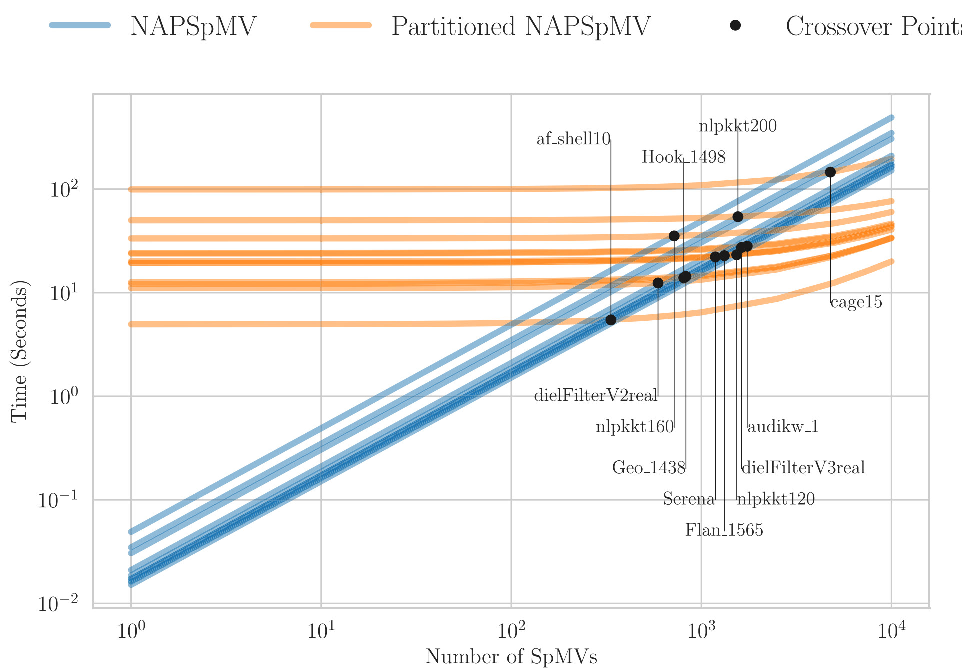

Fig. 17. The time required to perform various numbers of NAPSpMVs on strided and balanced partitions of the largest real SuiteSparse matrices with 50 000 non-zeros per process. The time to perform a NAPSpMV on a balanced partition includes the setup cost of partitioning and redistributing the matrix. The crossover points represent the number of NAPSpMVs required before graph partitioning becomes less costly than performing NAPSpMVs on the strided partition.

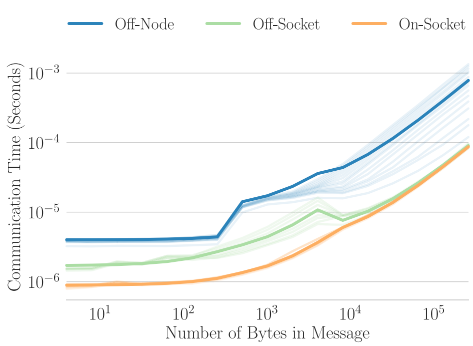

Fig. 18. The time required to send a single message of various sizes, with the thin lines representing timings measured by Nodecomm and the thick lines displaying the *max-rate* and intra-node models in (10) and (12), respectively. The intra-node models are split into two categories, on-socket and off-socket.

operations as well as applications and solvers independent of matrices.

References

[1] T. Agarwal, A. Sharma, L.V. Kalé, Topology-aware task mapping for reducing communication contention on large parallel machines, in: Proceedings of the 20th International Conference on Parallel and Distributed Processing, in: IPDPS'06, IEEE Computer Society, Washington, DC, USA, 2006, 145–145.

[2] N. Bell, S. Dalton, L.N. Olson, Exposing fine-grained parallelism in algebraic multigrid methods, SIAM J. Sci. Comput. 34 (4) (2012) C123–C152, http://dx.doi.org/10.1137/110838844, arXiv:https://doi.org/10.1137/110838844.

[3] N. Bell, M. Garland, Efficient sparse matrix-vector multiplication on CUDA, NVIDIA Technical Report NVR-2008-004, NVIDIA Corporation, 2008.

[4] A. Bienz, R.D. Falgout, W. Gropp, L.N. Olson, J.B. Schroder, Reducing parallel communication in algebraic multigrid through sparsification, SIAM J. Sci. Comput. 38 (5) (2016) S332–S357, http://dx.doi.org/10.1137/15M1026341, arXiv:https://doi.org/10.1137/15M1026341.

[5] A. Bienz, W.D. Gropp, L. Olson, Topology-aware performance modeling of parallel spmvs, in: 17th SIAM Conference on Parallel Processing for Scientific Computing, 2016, URL http://meetings.siam.org/sess/dsp_talk.cfm?p=75934.

[6] A. Bienz, L.N. Olson, RAPtor: parallel algebraic multigrid v0.1, Release 0.1, 2017, URL https://github.com/lukeolson/raptor.

[7] R.H. Bisseling, W. Meesen, Communication balancing in parallel sparse matrix-vector multiplication., ETNA. Electronic Trans. Numer. Anal. [electronic only] 21 (2005) 47–65, URL http://eudml.org/doc/128024.

[8] U.V. Çatalyürek, C. Aykanat, Hypergraph-partitioning-based decomposition for parallel sparse-matrix vector multiplication, IEEE Trans. Parallel Distrib. Syst. 10 (7) (1999) 673–693, http://dx.doi.org/10.1109/71.780863.

[9] U.V. Çatalyürek, C. Aykanat, B. Uçar, On two-dimensional sparse matrix partitioning: models, methods, and a recipe, SIAM J. Sci. Comput. 32 (2) (2010) 656–683, http://dx.doi.org/10.1137/080737770, arXiv:https://doi.org/10.1137/080737770.

[10] C. Chevalier, F. Pellegrini, PT-Scotch: a tool for efficient parallel graph ordering, Parallel Comput. 34 (68) (2008) 318–331, http://dx.doi.org/10.1016/j.parco.2007.12.001, Parallel Matrix Algorithms and Applications. URL http://www.sciencedirect.com/science/article/pii/S0167819107001342.

[11] S. Dalton, L. Olson, N. Bell, Optimizing sparse matrix-matrix multiplication for the gpu, ACM Trans. Math. Software 41 (4) (2015) 25:1–25:20, http://dx.doi.org/10.1145/2699470, URL http://doi.acm.org/10.1145/2699470.

[12] T.A. Davis, Y. Hu, The university of florida sparse matrix collection, ACM Trans. Math. Software 32 (2011) 1:1 – 1:25, URL http://www.cise.ufl.edu/research/sparse/matrices.

[13] R. Geus, S. Röllin, Towards a fast parallel sparse symmetric matrix-vector multiplication, Parallel Comput. 27 (7) (2001) 883–896, http://dx.doi.org/10.1016/S0167-8191(01)00073-4, Linear systems and associated problems. URL http://www.sciencedirect.com/science/article/pii/S0167819101000734.

[14] W. Gropp, L.N. Olson, P. Samfass, Modeling MPI communication performance on SMP nodes: is it time to retire the ping pong test, in: Proceedings of the 23rd European MPI Users' Group Meeting, in: EuroMPI 2016, ACM, New York, NY, USA, 2016, pp. 41–50, http://dx.doi.org/10.1145/2966884.2966919, URL http://doi.acm.org/10.1145/2966884.2966919.

[15] D. Guo, W. Gropp, L.N. Olson, A hybrid format for better performance of sparse matrix-vector multiplication on a gpu, Int. J. High Performance Comput. Appl. 30 (1) (2016) 103–120, http://dx.doi.org/10.1177/1094342015593156, arXiv:https://doi.org/10.1177/1094342015593156.

[16] B. Hendrickson, T.G. Kolda, Graph partitioning models for parallel computing, Parallel Comput. 26 (12) (2000) 1519–1534, http://dx.doi.org/10.1016/S0167-8191(00)00048-X.

[17] E. Jeannot, E. Meneses, G. Mercier, F. Tessier, G. Zheng, Communication and topology-aware load balancing in charm++ with treematch, in: 2013 IEEE International Conference on Cluster Computing (CLUSTER), 2013, pp. 1–8, http://dx.doi.org/10.1109/CLUSTER.2013.6702666.

[18] N.T. Karonis, B.R. de Supinski, I. Foster, W. Gropp, E. Lusk, J. Bresnahan, Exploiting hierarchy in parallel computer networks to optimize collective operation performance, in: Proceedings 14th International Parallel and Distributed Processing Symposium. IPDPS 2000, 2000, pp. 377–384, http://dx.doi.org/10.1109/IPDPS.2000.846009.

[19] G. Karypis, V. Kumar, A parallel algorithm for multilevel graph partitioning and sparse matrix ordering, J. Parallel Distrib. Comput. 48 (1) (1998) 71–95, http://dx.doi.org/10.1006/jpdc.1997.1403.

[20] T. Kielmann, R.F.H. Hofman, H.E. Bal, A. Plaat, R.A.F. Bhoedjang, Magpie: mpi's collective communication operations for clustered wide area systems, SIGPLAN Not. 34 (8) (1999) 131–140, http://dx.doi.org/10.1145/329366.301116, URL http://doi.acm.org/10.1145/329366.301116.
[21] T. Malik, V. Rychkov, A. Lastovetsky, Network-aware optimization of communications for parallel matrix multiplication on hierarchical hpc platforms, Concurr. Comput. : Pract. Exper. 28 (3) (2016) 802–821, http://dx.doi.org/10.1002/cpe.3609.
[22] A. Pinar, C. Aykanat, Fast optimal load balancing algorithms for 1d partitioning, J. Parallel Distrib. Comput. 64 (8) (2004) 974–996, http://dx.doi.org/10.1016/j.jpdc.2004.05.003.
[23] A. Reisner, L.N. Olson, J.D. Moulton, Scaling Structured Multigrid to 500K+ Cores through Coarse-Grid Redistribution, CoRR, abs/1803.02481.
[24] P. Sack, W. Gropp, Faster topology-aware collective algorithms through non-minimal communication, in: Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, in: PPoPP '12, ACM, New York, NY, USA, 2012, pp. 45–54, http://dx.doi.org/10.1145/2145816.2145823, URL http://doi.acm.org/10.1145/2145816.2145823.
[25] G. Schubert, H. Fehske, G. Hager, G. Wellein, Hybrid-parallel sparse matrix-vector multiplication with explicit communication overlap on current multicore-based systems, Parallel Process. Lett. 21 (03) (2011) 339–358, http://dx.doi.org/10.1142/S0129626411000254, arXiv:http://www.worldscientific.com/doi/pdf/10.1142/S0129626411000254. URL http://www.worldscientific.com/doi/abs/10.1142/S0129626411000254.
[26] E. Solomonik, A. Bhatele, J. Demmel, Improving communication performance in dense linear algebra via topology aware collectives, in: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, in: SC '11, ACM, New York, NY, USA, 2011, pp. 77:1–77:11, http://dx.doi.org/10.1145/2063384.2063487, URL http://doi.acm.org/10.1145/2063384.2063487.
[27] S. Sreepathi, E. D'Azevedo, B. Philip, P. Worley, Communication characterization and optimization of applications using topology-aware task mapping on large supercomputers, in: Proceedings of the 7th ACM/SPEC on International Conference on Performance Engineering, in: ICPE '16, ACM, New York, NY, USA, 2016, pp. 225–236, http://dx.doi.org/10.1145/2851553.2851575, URL http://doi.acm.org/10.1145/2851553.2851575.
[28] H.D. Sterck, U.M. Yang, J.J. Heys, Reducing complexity in parallel algebraic multigrid preconditioners, SIAM J. Matrix Anal. Appl. 27 (4) (2006) 1019–1039, http://dx.doi.org/10.1137/040615729, arXiv:https://doi.org/10.1137/040615729.
[29] J.L. Träff, Implementing the mpi process topology mechanism, in: Proceedings of the 2002 ACM/IEEE Conference on Supercomputing, in: SC '02, IEEE Computer Society Press, Los Alamitos, CA, USA, 2002, pp. 1–14, URL http://dl.acm.org/citation.cfm?id=762761.762767.
[30] B. Vastenhouw, R.H. Bisseling, A two-dimensional data distribution method for parallel sparse matrix-vector multiplication, SIAM Rev. 47 (1) (2005) 67–95, http://dx.doi.org/10.1137/S0036144502409019, arXiv:https://doi.org/10.1137/S0036144502409019.
[31] L. Wesolowski, R. Venkataraman, A. Gupta, J.S. Yeom, K. Bisset, Y. Sun, P. Jetley, T.R. Quinn, L.V. Kale, Tram: optimizing fine-grained communication with topological routing and aggregation of messages, in: 2014 43rd International Conference on Parallel Processing, 2014, pp. 211–220, http://dx.doi.org/10.1109/ICPP.2014.30.
[32] S. Williams, L. Oliker, R. Vuduc, J. Shalf, K. Yelick, J. Demmel, Optimization of sparse matrix-vector multiplication on emerging multicore platforms, in: Proceedings of the 2007 ACM/IEEE Conference on Supercomputing, in: SC '07, ACM, New York, NY, USA, 2007, pp. 38:1–38:12, http://dx.doi.org/10.1145/1362622.1362674, URL http://doi.acm.org/10.1145/1362622.1362674.
[33] D.E. WOMBLE, B.C. YOUNG, A model and implementation of multigrid for massively parallel computers, Int. J. High Speed Comput. 02 (03) (1990) 239–255, http://dx.doi.org/10.1142/S0129053390000157, arXiv:https://www.worldscientific.com/doi/pdf/10.1142/S0129053390000157. URL https://www.worldscientific.com/doi/abs/10.1142/S0129053390000157.

Amanda Bienz is a Ph.D. student in the Department of Computer Science at the University of Illinois at Urbana-Champaign. Her research focus is on reducing communication in sparse matrix operations in applications relating to algebraic multigrid.

William Gropp is Director and Chief Scientist of the National Center for Supercomputing Applications and holds the Thomas M. Siebel Chair in the Department of Computer Science at the University of Illinois in Urbana-Champaign. He received his Ph.D. in Computer Science from Stanford University in 1982 and worked at Yale University and Argonne National Laboratory. His research interests are in parallel computing, software for scientific computing, and numerical methods for partial differential equations. He is a Fellow of ACM, IEEE, and SIAM and a member of the National Academy of Engineering.

Luke Olson is a professor in the Department of Computer Science at the University of Illinois at Urbana-Champaign. His research interests include sparse matrix operations, algebraic multigrid, finite elements, and high performance computing. Olson received a Ph.D. in Applied Mathematics from the University of Colorado at Boulder in 2003.