

PARALLEL MATRIX MULTIPLICATION: A SYSTEMATIC JOURNEY*

MARTIN D. SCHATZ[†], ROBERT A. VAN DE GEIJN[†], AND JACK POULSON[‡]

Abstract. We expose a systematic approach for developing distributed-memory parallel matrix-matrix multiplication algorithms. The journey starts with a description of how matrices are distributed to meshes of nodes (e.g., MPI processes), relates these distributions to scalable parallel implementation of matrix-vector multiplication and rank-1 update, continues on to reveal a family of matrix-matrix multiplication algorithms that view the nodes as a two-dimensional (2D) mesh, and finishes with extending these 2D algorithms to so-called three-dimensional (3D) algorithms that view the nodes as a 3D mesh. A cost analysis shows that the 3D algorithms can attain the (order of magnitude) lower bound for the cost of communication. The paper introduces a taxonomy for the resulting family of algorithms and explains how all algorithms have merit depending on parameters such as the sizes of the matrices and architecture parameters. The techniques described in this paper are at the heart of the Elemental distributed-memory linear algebra library. Performance results from implementation within and with this library are given on a representative distributed-memory architecture, the IBM Blue Gene/P supercomputer.

Key words. parallel processing, linear algebra, matrix multiplication, libraries

AMS subject classifications. 65Y05, 65Y20, 65F05

DOI. 10.1137/140993478

1. Introduction. This paper serves a number of purposes:

- Parallel¹ implementation of matrix-matrix multiplication is a standard topic in a course on parallel high-performance computing. However, rarely is the student exposed to the algorithms that are used in practical cutting-edge parallel dense linear algebra (DLA) libraries. This paper exposes a systematic path that leads from parallel algorithms for matrix-vector multiplication and rank-1 update to a practical, scalable family of parallel algorithms for matrix-matrix multiplication, including the classic result in [2] and those implemented in the Elemental parallel DLA library [28].

*Submitted to the journal's Software and High-Performance Computing section October 29, 2014; accepted for publication (in revised form) August 24, 2016; published electronically December 21, 2016. This research used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-06CH11357; early experiments were performed on the Texas Advanced Computing Center's Ranger Supercomputer. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).

<http://www.siam.org/journals/sisc/38-6/99347.html>

Funding: This research was partially sponsored by NSF grants OCI-0850750, CCF-0917167, ACI-1148125/1340293, and CCF-1320112; grants from Microsoft; and an unrestricted grant from Intel. The first author was partially supported by a Sandia Fellowship. The third author was partially supported by a fellowship from the Institute of Computational Engineering and Sciences.

[†]Department of Computer Science, Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX 78712 (martin.schatz@utexas.edu, rvdg@cs.utexas.edu).

[‡]Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305 (poulson@stanford.edu).

¹Parallel in this paper implicitly means *distributed-memory* parallel.

- This paper introduces a set notation for describing the data distributions that underlie the Elemental library. The notation is motivated using parallelization of matrix-vector operations and matrix-matrix multiplication as the driving examples.
- Recently, research on parallel matrix-matrix multiplication algorithms have revisited so-called three-dimensional (3D) algorithms, which view (processing) nodes as a logical 3D mesh. These algorithms are known to attain theoretical (order of magnitude) lower bounds on communication. This paper exposes a systematic path from algorithms for two-dimensional (2D) meshes to their extensions for 3D meshes. Among the resulting algorithms are classic results [1].
- A taxonomy is given for the resulting family of algorithms, all of which are related to what is often called the scalable universal matrix multiplication algorithm (SUMMA) [33].

Thus, the paper simultaneously serves a pedagogical role, explains abstractions that underlie the Elemental library, and advances the state of science for parallel matrix-matrix multiplication by providing a framework to systematically derive known and new algorithms for matrix-matrix multiplication when computing on 2D or 3D meshes. While much of the new innovation presented in this paper concerns the extension of parallel matrix-matrix multiplication algorithms from 2D to 3D meshes, we believe that developing the reader’s intuition for algorithms on 2D meshes renders most of this new innovation a straightforward extension.

2. Background. The parallelization of dense matrix-matrix multiplication is a well-studied subject. Cannon’s algorithm (sometimes called roll-roll-compute) dates back to 1969 [9], and Fox’s algorithm (sometimes called broadcast-roll-compute) dates back to 1988 [15]. Both suffer from a number of shortcomings:

- They assume that p processes are viewed as a $d_0 \times d_1$ grid, with $d_0 = d_1 = \sqrt{p}$. Removing this constraint on d_0 and d_1 is nontrivial for these algorithms.
- They do not deal well with the case where one of the matrix dimensions becomes relatively small. This is the most commonly encountered case in libraries such as LAPACK [3] and `libflame` [35, 36] and their distributed-memory counterparts ScaLAPACK [11], PLAPACK [34], and Elemental [28].

Attempts to generalize [12, 21, 22] led to implementations that were neither simple nor effective.

A practical algorithm, which also results from the systematic approach discussed in this paper, can be described as “allgather-allgather-multiply” [2]. It does not suffer from the shortcomings of Cannon’s and Fox’s algorithms. It did not gain popularity in part because libraries such as ScaLAPACK and PLAPACK used a 2D block-cyclic distribution, rather than the 2D elemental distribution advocated by that paper. The arrival of the Elemental library, together with what we believe is our more systematic and extensible explanation, will, we hope, elevate awareness of this result.

SUMMA [33] is another algorithm that overcomes all of the shortcomings of Cannon’s and Fox’s algorithms. We believe it is a more widely known result, in part because it can already be explained for a matrix that is distributed with a 2D blocked (but not cyclic) distribution, and in part because it is easy to support in the ScaLAPACK and PLAPACK libraries. The original SUMMA paper gives four algorithms as follows:

- For $C := AB + C$, SUMMA casts the computation in terms of multiple rank-k updates. This algorithm is sometimes called the broadcast-broadcast-

multiply algorithm, a label which, we will see, is somewhat limiting. We also call this algorithm “stationary C ” for reasons that will become clear later. By design, this algorithm continues to perform well in the case where the width of A is small relative to the dimensions of C .

- For $C := A^T B + C$, SUMMA casts the computation in terms of multiple panel of rows times matrix multiplications, so performance is not degraded in the case where the height of A is small relative to the dimensions of B . We also call this algorithm “stationary B ” for reasons that will become clear later.
- For $C := AB^T + C$, SUMMA casts the computation in terms of multiple matrix-panel (of columns) multiplications, and so performance does not deteriorate when the width of C is small relative to the dimensions of A . We call this algorithm “stationary A ” for reasons that will become clear later.
- For $C := A^T B^T + C$, the paper sketches an algorithm that is actually not practical.

In [17], it was shown how stationary A , B , and C algorithms can be formulated for each of the four cases of matrix-matrix multiplication, including $C := A^T B^T + C$. This then yielded a general, practical family of 2D matrix-matrix multiplication algorithms, all of which were incorporated into PLAPACK and Elemental, and some of which are supported by ScaLAPACK. Some of the observations about developing 2D algorithms in the current paper can already be found in [2], but our exposition is much more systematic, and we use the matrix distribution that underlies Elemental to illustrate the basic principles. Although the work by Agarwal, Gustavson, and Zubair describes algorithms for the different matrix-matrix multiplication transpose variants, it does not describe how to create stationary A and B variants.

In the 1990s, it was observed that for the case where matrices were relatively small (or, equivalently, a relatively large number of nodes were available), better theoretical and practical performance resulted from viewing the p nodes as a $d_0 \times d_1 \times d_2$ mesh, yielding a 3D algorithm [1]. More recently, a 3D algorithm for computing the LU factorization of a matrix was devised by McColl and Tiskin [27] and Solomonik and Demmel [31]. In addition to the LU factorization algorithm devised in [31], a 3D algorithm for matrix-matrix multiplication was given for nodes arranged as a $d_0 \times d_1 \times d_2$ mesh, with $d_0 = d_1$ and $0 \leq d_2 < \sqrt[3]{p}$. This was labeled a 2.5D algorithm. Although the primary contribution of that work was LU related, the 2.5D algorithm for matrix-matrix multiplication is the portion relevant to this paper. The focus of that study on 3D algorithms was the simplest case of matrix-matrix multiplication, $C := AB$.

In [25], an early attempt was made to combine multiple algorithms for computing $C = AB$ into a poly-algorithm, which refers to “the use of two or more algorithms to solve the same problem with a high level decision-making process determining which of a set of algorithms performs best in a given situation.” That paper was published during the same time when SUMMA algorithms first became popular and when it was not yet completely understood that these SUMMA algorithms are inherently more practical than Cannon’s and Fox’s algorithms. The “stationary A, B, and C” algorithms were already being talked about. In [25], an attempt was made to combine all of these approaches, including SUMMA, targeting general 2D Cartesian data distributions, which was (and still would be) a very ambitious goal. Our paper benefits from decades of experience with the more practical SUMMA algorithms and their variants. It purposely limits the data distribution to simple distributions, namely elemental distributions. This, we hope, allows the reader to gain a deep under-

standing in a simpler setting so that even if elemental distribution is not best for a particular situation, a generalization can be easily derived. The family of presented 2D algorithms is a poly-algorithm implemented in Elemental.

3. Notation. Although the focus of this paper is parallel distributed-memory matrix-matrix multiplication, the notation used is designed to be extensible to computation with higher-dimensional objects (tensors) on higher-dimensional grids. Because of this, the notation used may seem overly complex when restricted to matrix-matrix multiplication. In this section, we describe the notation used and the reasoning behind the choice of notation.

Grid dimension: d_x . Since we focus on algorithms for distributed-memory architectures, we must describe information about the grid on which we are computing. To support arbitrary-dimensional grids, we must express the shape of the grid in an extensible way. For this reason, we have chosen the subscripted letter d to indicate the size of a particular dimension of the grid. Thus, d_x refers to the number of processes comprising the x th dimension of the grid. In this paper, the grid is typically $d_0 \times d_1$.

Process location: s_x . In addition to describing the shape of the grid, it is useful to be able to refer to a particular process's location within the mesh of processes. For this, we use the subscripted s letter to refer to a process's location within some given dimension of the mesh of processes. Thus, s_x refers to a particular process's location within the x th dimension of the mesh of processes. In this paper, a typical process is labeled with (s_0, s_1) .

Distribution: $\mathcal{D}_{(x_0, x_1, \dots, x_{k-1})}$. In subsequent sections, we will introduce notation for describing how data is distributed among processes of the grid. This notation will require a description of which dimensions of the grid are involved in defining the distribution. We use the symbol $\mathcal{D}_{(x_0, x_1, \dots, x_{k-1})}$ to indicate a distribution which involves dimensions x_0, x_1, \dots, x_{k-1} of the mesh.

For example, when describing a distribution which involves the column and row dimension of the grid, we refer to this distribution as $\mathcal{D}_{(0,1)}$. Later, we will explain why the symbol $\mathcal{D}_{(0,1)}$ describes a different distribution from $\mathcal{D}_{(1,0)}$.

4. Of matrix-vector operations and distribution. In this section, we discuss how matrix and vector distributions can be linked to parallel 2D matrix-vector multiplication and rank-1 update operations, which then allows us to eventually describe the stationary C , A , and B 2D algorithms for matrix-matrix multiplication that are part of the Elemental library.

4.1. Collective communication. Collectives are fundamental to the parallelization of dense matrix operations. Thus, the reader must be (or become) familiar with the basics of these communications and is encouraged to read Chan et al. [10], which presents collectives in a systematic way that dovetails with the present paper.

To make this paper self-contained, in Table 4.1 (similar to Figure 1 in [10]) we summarize the collectives. In Table 4.2 we summarize lower bounds on the cost of the collective communications, under basic assumptions explained in [10] (see [8] for an analysis of all-to-all), and the cost expressions that we will use in our analyses.

TABLE 4.1
Collective communications considered in this paper.

Operation	Before				After			
Permute	$\frac{\text{Node 0}}{x_0}$	$\frac{\text{Node 1}}{x_1}$	$\frac{\text{Node 2}}{x_2}$	$\frac{\text{Node 3}}{x_3}$	$\frac{\text{Node 0}}{x_1}$	$\frac{\text{Node 1}}{x_0}$	$\frac{\text{Node 2}}{x_3}$	$\frac{\text{Node 3}}{x_2}$
Broadcast	$\frac{\text{Node 0}}{x}$	$\frac{\text{Node 1}}$	$\frac{\text{Node 2}}$	$\frac{\text{Node 3}}$	$\frac{\text{Node 0}}{x}$	$\frac{\text{Node 1}}{x}$	$\frac{\text{Node 2}}{x}$	$\frac{\text{Node 3}}{x}$
Reduce(-to-one)	$\frac{\text{Node 0}}{x^{(0)}}$	$\frac{\text{Node 1}}{x^{(1)}}$	$\frac{\text{Node 2}}{x^{(2)}}$	$\frac{\text{Node 3}}{x^{(3)}}$	$\frac{\text{Node 0}}{\sum_j x^{(j)}}$	$\frac{\text{Node 1}}$	$\frac{\text{Node 2}}$	$\frac{\text{Node 3}}$
Scatter	$\frac{\text{Node 0}}{x_0}$	$\frac{\text{Node 1}}$	$\frac{\text{Node 2}}$	$\frac{\text{Node 3}}$	$\frac{\text{Node 0}}{x_0}$	$\frac{\text{Node 1}}{x_1}$	$\frac{\text{Node 2}}{x_2}$	$\frac{\text{Node 3}}{x_3}$
Gather	$\frac{\text{Node 0}}{x_0}$	$\frac{\text{Node 1}}{x_1}$	$\frac{\text{Node 2}}{x_2}$	$\frac{\text{Node 3}}{x_3}$	$\frac{\text{Node 0}}{x_0}$	$\frac{\text{Node 1}}{x_1}$	$\frac{\text{Node 2}}{x_2}$	$\frac{\text{Node 3}}{x_3}$
Allgather	$\frac{\text{Node 0}}{x_0}$	$\frac{\text{Node 1}}{x_1}$	$\frac{\text{Node 2}}{x_2}$	$\frac{\text{Node 3}}{x_3}$	$\frac{\text{Node 0}}{x_0}$	$\frac{\text{Node 1}}{x_0}$	$\frac{\text{Node 2}}{x_0}$	$\frac{\text{Node 3}}{x_0}$
Reduce-scatter	$\frac{\text{Node 0}}{x_0^{(0)}}$	$\frac{\text{Node 1}}{x_0^{(1)}}$	$\frac{\text{Node 2}}{x_0^{(2)}}$	$\frac{\text{Node 3}}{x_0^{(3)}}$	$\frac{\text{Node 0}}{\sum_j x_0^{(j)}}$	$\frac{\text{Node 1}}{\sum_j x_1^{(j)}}$	$\frac{\text{Node 2}}{\sum_j x_2^{(j)}}$	$\frac{\text{Node 3}}{\sum_j x_3^{(j)}}$
Allreduce	$\frac{\text{Node 0}}{x^{(0)}}$	$\frac{\text{Node 1}}{x^{(1)}}$	$\frac{\text{Node 2}}{x^{(2)}}$	$\frac{\text{Node 3}}{x^{(3)}}$	$\frac{\text{Node 0}}{\sum_j x^{(j)}}$	$\frac{\text{Node 1}}{\sum_j x^{(j)}}$	$\frac{\text{Node 2}}{\sum_j x^{(j)}}$	$\frac{\text{Node 3}}{\sum_j x^{(j)}}$
All-to-all	$\frac{\text{Node 0}}{x_0^{(0)}}$	$\frac{\text{Node 1}}{x_0^{(1)}}$	$\frac{\text{Node 2}}{x_0^{(2)}}$	$\frac{\text{Node 3}}{x_0^{(3)}}$	$\frac{\text{Node 0}}{x_0^{(0)}}$	$\frac{\text{Node 1}}{x_1^{(0)}}$	$\frac{\text{Node 2}}{x_2^{(0)}}$	$\frac{\text{Node 3}}{x_3^{(0)}}$

4.2. Motivation: Matrix-vector multiplication. Suppose $A \in \mathbb{R}^{m \times n}$, $x \in \mathbb{R}^n$, and $y \in \mathbb{R}^m$, and label their individual elements so that

$$A = \begin{pmatrix} \alpha_{0,0} & \alpha_{0,1} & \cdots & \alpha_{0,n-1} \\ \alpha_{1,0} & \alpha_{1,1} & \cdots & \alpha_{1,n-1} \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_{m-1,0} & \alpha_{m-1,1} & \cdots & \alpha_{m-1,n-1} \end{pmatrix}, \quad x = \begin{pmatrix} \chi_0 \\ \chi_1 \\ \vdots \\ \chi_{n-1} \end{pmatrix}, \quad \text{and } y = \begin{pmatrix} \psi_0 \\ \psi_1 \\ \vdots \\ \psi_{m-1} \end{pmatrix}.$$

	χ_0	\dots		χ_1	\dots		χ_2	\dots
ψ_0	$\alpha_{0,0}$	$\alpha_{0,3}$	$\alpha_{0,6}$	\dots	$\alpha_{0,1}$	$\alpha_{0,4}$	$\alpha_{0,7}$	\dots
	$\alpha_{2,0}$	$\alpha_{2,3}$	$\alpha_{2,6}$	\dots	ψ_2	$\alpha_{2,1}$	$\alpha_{2,4}$	$\alpha_{2,7}$
	$\alpha_{4,0}$	$\alpha_{4,3}$	$\alpha_{4,6}$	\dots		$\alpha_{4,1}$	$\alpha_{4,4}$	$\alpha_{4,7}$
\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\ddots
	χ_3	\dots		χ_4	\dots		χ_5	\dots
ψ_1	$\alpha_{1,0}$	$\alpha_{1,3}$	$\alpha_{1,6}$	\dots	$\alpha_{1,1}$	$\alpha_{1,4}$	$\alpha_{1,7}$	\dots
	$\alpha_{3,0}$	$\alpha_{3,3}$	$\alpha_{3,6}$	\dots	ψ_3	$\alpha_{3,1}$	$\alpha_{3,4}$	$\alpha_{3,7}$
	$\alpha_{5,0}$	$\alpha_{5,3}$	$\alpha_{5,6}$	\dots		$\alpha_{5,1}$	$\alpha_{5,4}$	$\alpha_{5,7}$
\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\ddots

FIG. 4.1. Distribution of A , x , and y within a 2×3 mesh. Redistributing a column of A in the same manner as y requires simultaneous scatters within rows of nodes, while redistributing a row of A consistently with x requires simultaneous scatters within columns of nodes. In the notation of section 5, here the distributions of x and y are given by $x[(1,0),()]$ and $y[(0,1),()]$, respectively, and that of A is given by $A[(0),(1)]$.

node (s_0, s_1) stores submatrix

$$A(s_0:d_0:m-1, s_1:d_1:n-1) = \begin{pmatrix} \alpha_{s_0,s_1} & \alpha_{s_0,s_1+d_1} & \dots \\ \alpha_{s_0+d_0,s_1} & \alpha_{s_0+d_0,s_1+d_1} & \dots \\ \vdots & \vdots & \ddots \end{pmatrix},$$

where the left-hand side of the expression uses the MATLAB convention for expressing submatrices, starting indexing from zero instead of one. This is illustrated in Figure 4.1.

Column-major vector distribution. A *column-major* vector distribution views the $d_0 \times d_1$ mesh of nodes as a linear array of p nodes, numbered in *column-major* order. A vector is distributed with this distribution if it is assigned to this linear array of nodes in a round-robin fashion, one element at a time. In other words, consider vector y . Its element ψ_i is assigned to node $(i \bmod d_0, (i/d_0) \bmod d_1)$, where $/$ denotes integer division. Or, equivalently, in MATLAB-like notation, node (s_0, s_1) stores subvector $y(u(s_0, s_1) : p : m-1)$, where $u(s_0, s_1) = s_0 + s_1 d_0$ equals the rank of node (s_0, s_1) when the nodes are viewed as a one-dimensional (1D) array, indexed in column-major order. This distribution of y is illustrated in Figure 4.1.

Row-major vector distribution. Similarly, a *row-major* vector distribution views the $d_0 \times d_1$ mesh of nodes as a linear array of p nodes, numbered in *row-major* order. In other words, consider vector x . Its element χ_j is assigned to node $(j \bmod d_1, (j/d_1) \bmod d_0)$. Or, equivalently, node (s_0, s_1) stores subvector $x(v(s_0, s_1) : p : n-1)$, where $v(s_0, s_1) = s_0 d_1 + s_1$ equals the rank of node (s_0, s_1) when the nodes are viewed as a 1D array, indexed in row-major order. The distribution of x is illustrated in Figure 4.1.

	χ_0	χ_3	χ_6	\dots		χ_1	χ_4	χ_7	\dots		χ_2	χ_5	χ_8	\dots
ψ_0	$\alpha_{0,0}$	$\alpha_{0,3}$	$\alpha_{0,6}$	\dots	ψ_0	$\alpha_{0,1}$	$\alpha_{0,4}$	$\alpha_{0,7}$	\dots	ψ_0	$\alpha_{0,2}$	$\alpha_{0,5}$	$\alpha_{0,8}$	\dots
ψ_2	$\alpha_{2,0}$	$\alpha_{2,3}$	$\alpha_{2,6}$	\dots	ψ_2	$\alpha_{2,1}$	$\alpha_{2,4}$	$\alpha_{2,7}$	\dots	ψ_2	$\alpha_{2,2}$	$\alpha_{2,5}$	$\alpha_{2,8}$	\dots
ψ_4	$\alpha_{4,0}$	$\alpha_{4,3}$	$\alpha_{4,6}$	\dots	ψ_4	$\alpha_{4,1}$	$\alpha_{4,4}$	$\alpha_{4,7}$	\dots	ψ_4	$\alpha_{4,2}$	$\alpha_{4,5}$	$\alpha_{4,8}$	\dots
\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\vdots	\ddots
	χ_0	χ_3	χ_6	\dots		χ_1	χ_4	χ_7	\dots		χ_2	χ_5	χ_8	\dots
ψ_1	$\alpha_{1,0}$	$\alpha_{1,3}$	$\alpha_{1,6}$	\dots	ψ_1	$\alpha_{1,1}$	$\alpha_{1,4}$	$\alpha_{1,7}$	\dots	ψ_1	$\alpha_{1,2}$	$\alpha_{1,5}$	$\alpha_{1,8}$	\dots
ψ_3	$\alpha_{3,0}$	$\alpha_{3,3}$	$\alpha_{3,6}$	\dots	ψ_3	$\alpha_{3,1}$	$\alpha_{3,4}$	$\alpha_{3,7}$	\dots	ψ_3	$\alpha_{3,2}$	$\alpha_{3,5}$	$\alpha_{3,8}$	\dots
ψ_5	$\alpha_{5,0}$	$\alpha_{5,3}$	$\alpha_{5,6}$	\dots	ψ_5	$\alpha_{5,1}$	$\alpha_{5,4}$	$\alpha_{5,7}$	\dots	ψ_5	$\alpha_{5,2}$	$\alpha_{5,5}$	$\alpha_{5,8}$	\dots
\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\vdots	\ddots

FIG. 4.2. Vectors x and y , respectively, redistributed as row-projected and column-projected vectors. The column-projected vector $y[(0), ()]$ here is to be used to compute local results that will become contributions to a column vector $y[(0, 1), ()]$, which will result from adding these local contributions within rows of nodes. By comparing and contrasting this figure with Figure 4.1, it becomes obvious that redistributing $x[(1, 0), ()]$ to $x[(1), ()]$ requires an allgather within columns of nodes, while $y[(0, 1), ()]$ results from scattering $y[(0), ()]$ within process rows.

4.4. Parallelizing matrix-vector operations. In the following discussion, we assume that A , x , and y are distributed as discussed above.² At this point, we suggest comparing (4.1) with Figure 4.1.

Computing $y := Ax$. The relation between the distributions of a matrix, column-major vector, and row-major vector is illustrated by revisiting the most fundamental of computations in linear algebra, $y := Ax$, already discussed in section 4.2. An examination of Figure 4.1 suggests that the elements of x must be gathered within columns of nodes (allgather within columns) leaving elements of x distributed as illustrated in Figure 4.2. Next, each node computes the partial contribution to vector y with its local matrix and copy of x . Thus, in Figure 4.2, ψ_i in each node becomes a contribution to the final ψ_i . These must be added together, which is accomplished by a summation of contributions to y within rows of nodes. An experienced MPI programmer will recognize this as a reduce-scatter within each row of nodes.

Under our communication cost model, the cost of this parallel algorithm is given

by

$$\begin{aligned}
 T_{y=Ax}(m, n, r, c) = & \underbrace{2 \left\lceil \frac{m}{d_0} \right\rceil \left\lceil \frac{n}{d_1} \right\rceil \gamma}_{\text{local mvmult}} \\
 & + \underbrace{\log_2(d_0)\alpha + \frac{d_0 - 1}{d_0} \left\lceil \frac{n}{d_1} \right\rceil \beta}_{\text{allgather } x} + \underbrace{\log_2(d_1)\alpha + \frac{d_1 - 1}{d_1} \left\lceil \frac{m}{d_0} \right\rceil \beta + \frac{d_1 - 1}{d_1} \left\lceil \frac{m}{d_0} \right\rceil \gamma}_{\text{reduce-scatter } y}
 \end{aligned}$$

²We suggest the reader print copies of Figures 4.1 and 4.2 for easy referral while reading the rest of this section.

$$\approx 2 \frac{mn}{p} \gamma + \underbrace{C_0 \frac{m}{d_0} \gamma + C_1 \frac{n}{d_1} \gamma}_{\text{load imbalance}} + \log_2(p) \alpha + \frac{d_0 - 1}{d_0} \frac{n}{d_1} \beta + \frac{d_1 - 1}{d_1} \frac{m}{d_0} \beta + \frac{d_1 - 1}{d_1} \frac{m}{d_0} \gamma$$

for some constants C_0 and C_1 . We simplify this further to

$$(4.2) \quad 2 \frac{mn}{p} \gamma + \underbrace{\log_2(p) \alpha + \frac{d_0 - 1}{d_0} \frac{n}{d_1} \beta + \frac{d_1 - 1}{d_1} \frac{m}{d_0} \beta + \frac{d_1 - 1}{d_1} \frac{m}{d_0} \gamma}_{T^+_{y:=Ax}(m, n, d_0, d_1)},$$

since the load imbalance contributes a cost similar to that of the communication.³ Here, $T^+_{y:=Ax}(m, n, k/h, d_0, d_1)$ is used to refer to the overhead associated with the above algorithm for the $y = Ax$ operation. In Appendix A we use these estimates to show that this parallel matrix-vector multiplication is, for practical purposes, weakly scalable if d_0/d_1 is kept constant, but it is not if $d_0 \times d_1 = p \times 1$ or $d_0 \times d_1 = 1 \times p$.

Computing $x := A^T y$. Let us next discuss an algorithm for computing $x := A^T y$, where A is an $m \times n$ matrix and x and y are distributed as before (x with a row-major vector distribution and y with a column-major vector distribution).

Recall that $x = A^T y$ (transpose matrix-vector multiplication) means

$$\begin{aligned} \chi_0 &= \alpha_{0,0} \psi_0 + \alpha_{1,0} \psi_1 + \dots + \alpha_{n-1,0} \psi_{n-1} \\ \chi_1 &= \alpha_{0,1} \psi_0 + \alpha_{1,1} \psi_1 + \dots + \alpha_{n-1,1} \psi_{n-1} \\ &\vdots \\ \chi_{m-1} &= \alpha_{0,m-1} \psi_0 + \alpha_{1,m-1} \psi_1 + \dots + \alpha_{n-1,m-1} \psi_{n-1} \end{aligned}$$

or

$$(4.3) \quad \begin{array}{c|c|c|c} \chi_0 = & \chi_1 = & \dots & \chi_{m-1} = \\ \alpha_{0,0} \psi_0 + & \alpha_{0,1} \psi_0 + & \dots & \alpha_{0,n-1} \psi_0 + \\ \alpha_{1,0} \psi_1 + & \alpha_{1,1} \psi_1 + & \dots & \alpha_{1,n-1} \psi_1 + \\ \vdots & \vdots & & \vdots \\ \alpha_{n-1,0} \psi_{n-1} & \alpha_{n-1,1} \psi_{n-1} & \dots & \alpha_{n-1,n-1} \psi_{n-1} \end{array} .$$

An examination of (4.3) and Figure 4.1 suggests that the elements of y must be gathered within rows of nodes (allgather within rows), leaving elements of y distributed as illustrated in Figure 4.2. Next, each node computes the partial contribution to vector x with its local matrix and copy of y . Thus, in Figure 4.2 χ_j in each node becomes a contribution to the final χ_j . These must be added together, which is accomplished by a summation of contributions to x within columns of nodes. We again recognize this as a reduce-scatter but this time within each column of nodes.

The cost for this algorithm, approximating as we did when analyzing the algorithm for $y = Ax$, is

$$2 \frac{mn}{p} \gamma + \underbrace{\log_2(p) \alpha + \frac{d_1 - 1}{d_1} \frac{n}{d_0} \beta + \frac{d_0 - 1}{d_0} \frac{m}{d_1} \beta + \frac{d_0 - 1}{d_0} \frac{m}{d_1} \gamma}_{T^+_{x:=A^T y}(m, n, d_1, d_0)},$$

where, as before, we ignore overhead due to load imbalance since terms of the same order appear in the terms that capture communication overhead.

³It is tempting to approximate $\frac{x-1}{x}$ by 1, but this would yield formulae for the cases where the mesh is $p \times 1$ ($d_1 = 1$) or $1 \times p$ ($d_0 = 1$) that are overly pessimistic.

Computing $y := A^T x$. What if we wish to compute $y := A^T x$, where A is an $m \times n$ matrix and y is distributed with a column-major vector distribution and x with a row-major vector distribution? Now x must first be redistributed to a column-major vector distribution, after which the algorithm that we just discussed can be executed, and finally the result (in row-major vector distribution) must be redistributed to leave it as y in column-major vector distribution. This adds to the cost of $y := A^T x$ both the cost of the permutation that redistributes x and the cost of the permutation that redistributes the result to y .

Other cases. What if when computing $y := Ax$ the vector x is distributed like a row of matrix A ? What if the vector y is distributed like a column of matrix A ? We leave these cases as an exercise for the reader.

Hence, understanding the basic algorithms for multiplying with A and A^T allows one to systematically derive and analyze algorithms when the vectors that are involved are distributed to the nodes in different ways.

Computing $A := yx^T + A$. A second commonly encountered matrix-vector operation is the rank-1 update: $A := \alpha yx^T + A$. We will discuss the case where $\alpha = 1$. Recall that

$$A + yx^T = \begin{pmatrix} \alpha_{0,0} + \psi_0\chi_0 & \alpha_{0,1} + \psi_0\chi_1 & \cdots & \alpha_{0,n-1} + \psi_0\chi_{n-1} \\ \alpha_{1,0} + \psi_1\chi_0 & \alpha_{1,1} + \psi_1\chi_1 & \cdots & \alpha_{1,n-1} + \psi_1\chi_{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_{m-1,0} + \psi_{m-1}\chi_0 & \alpha_{m-1,1} + \psi_{m-1}\chi_1 & \cdots & \alpha_{m-1,n-1} + \psi_{m-1}\chi_{n-1} \end{pmatrix},$$

which, when considering Figures 4.1 and 4.2, suggests the following parallel algorithm: All-gather of y within rows. All-gather of x within columns. Update of the local matrix on each node.

The cost for this algorithm, approximating as we did when analyzing the algorithm for $y = Ax$, yields

$$2 \frac{mn}{p} \gamma + \underbrace{\log_2(p)\alpha + \frac{d_0 - 1}{d_0} \frac{n}{d_1} \beta + \frac{d_1 - 1}{d_1} \frac{m}{d_0} \beta}_{T^+_{A:=yx^T+A}(m, n, d_0, d_1)},$$

where, as before, we ignore overhead due to load imbalance, since terms of the same order appear in the terms that capture communication overhead. Notice that the cost is the same as a parallel matrix-vector multiplication, except for the “ γ ” term that results from the reduction within rows.

As before, one can modify this algorithm for the case when the vectors start with different distributions, building on intuition from matrix-vector multiplication. A pattern is emerging.

5. Generalizing the theme. The reader should now have an understanding of how vector and matrix distributions are related to the parallelization of basic matrix-vector operations. We generalize these insights using sets of indices as “filters” to indicate what parts of a matrix or vector a given process owns.

The results in this section are similar to those that underlie physically based matrix distribution [14], which itself also underlies PLAPACK. However, we formalize the notation beyond that used by PLAPACK. The link between distribution of vectors and matrices was first observed by Bisseling and McColl [6, 7] and, around the same time, by Lewis and Van de Geijn [24].

5.1. Vector distribution. The basic idea is to use two different partitions of the natural numbers as a means of describing the distribution of the row and column indices of a matrix.

DEFINITION 5.1 (subvectors and submatrices). *Let $x \in \mathbb{R}^n$ and $\mathcal{S} \subset \mathbb{N}$. Then $x[\mathcal{S}]$ equals the vector with elements from x , with indices in the set \mathcal{S} , in the order in which they appear in vector x . If $A \in \mathbb{R}^{m \times n}$ and $\mathcal{S}, \mathcal{T} \subset \mathbb{N}$, then $A[\mathcal{S}, \mathcal{T}]$ is the submatrix formed by keeping only the elements of A , whose row indices are in \mathcal{S} and column indices are in \mathcal{T} , in the order in which they appear in matrix A .*

We illustrate this idea with simple examples.

EXAMPLE 1. *Let*

$$x = \begin{pmatrix} \chi_0 \\ \chi_1 \\ \chi_2 \\ \chi_3 \end{pmatrix}$$

and

$$A = \begin{pmatrix} \alpha_{0,0} & \alpha_{0,1} & \alpha_{0,2} & \alpha_{0,3} & \alpha_{0,4} \\ \alpha_{1,0} & \alpha_{1,1} & \alpha_{1,2} & \alpha_{1,3} & \alpha_{1,4} \\ \alpha_{2,0} & \alpha_{2,1} & \alpha_{2,2} & \alpha_{2,3} & \alpha_{2,4} \\ \alpha_{3,0} & \alpha_{3,1} & \alpha_{3,2} & \alpha_{3,3} & \alpha_{3,4} \\ \alpha_{4,0} & \alpha_{4,1} & \alpha_{4,2} & \alpha_{4,3} & \alpha_{4,4} \end{pmatrix}.$$

If $\mathcal{S} = \{0, 2, 4, \dots\}$ and $\mathcal{T} = \{1, 3, 5, \dots\}$, then

$$x[\mathcal{S}] = \begin{pmatrix} \chi_0 \\ \chi_2 \end{pmatrix} \quad \text{and} \quad A[\mathcal{S}, \mathcal{T}] = \begin{pmatrix} \alpha_{0,1} & \alpha_{0,3} \\ \alpha_{2,1} & \alpha_{2,3} \\ \alpha_{4,1} & \alpha_{4,3} \end{pmatrix}.$$

We now introduce two fundamental ways to distribute vectors relative to a logical $d_0 \times d_1$ process grid.

DEFINITION 5.2 (column-major vector distribution). *Suppose that $p \in \mathbb{N}$ processes are available, and define*

$$\mathcal{V}_p^\sigma(q) = \{N \in \mathbb{N} : N \equiv q + \sigma \pmod{p}\}, \quad q \in \{0, 1, \dots, p-1\},$$

where $\sigma \in \{0, 1, \dots, p-1\}$ is an arbitrary alignment parameter. When p is implied from the context and σ is not important in the discussion, we will simply denote the above set by $\mathcal{V}(q)$.

If the p processes have been configured into a logical $d_0 \times d_1$ grid, a vector x is said to be in a column-major vector distribution if process (s_0, s_1) , where $s_0 \in \{0, \dots, d_0 - 1\}$ and $s_1 \in \{0, \dots, d_1 - 1\}$, is assigned the subvector $x(\mathcal{V}_p^\sigma(s_0 + s_1 d_0))$. This distribution is represented via the $d_0 \times d_1$ array of indices

$$\mathcal{D}_{(0,1)}(s_0, s_1) \equiv \mathcal{V}(s_0 + s_1 d_0), \quad (s_0, s_1) \in \{0, \dots, d_0 - 1\} \times \{0, \dots, d_1 - 1\},$$

and the shorthand $x[(0,1)]$ will refer to the vector x distributed such that process (s_0, s_1) stores $x(\mathcal{D}_{(0,1)}(s_0, s_1))$.

DEFINITION 5.3 (row-major vector distribution). *Similarly, the $d_0 \times d_1$ array*

$$\mathcal{D}_{(1,0)} \equiv \mathcal{V}(s_1 + s_0 d_1), \quad (s_0, s_1) \in \{0, \dots, d_0 - 1\} \times \{0, \dots, d_1 - 1\}$$

is said to define a row-major vector distribution. The shorthand $y[(1,0)]$ will refer to the vector y distributed such that process (s_0, s_1) stores $y(\mathcal{D}_{(1,0)}(s_0, s_1))$.

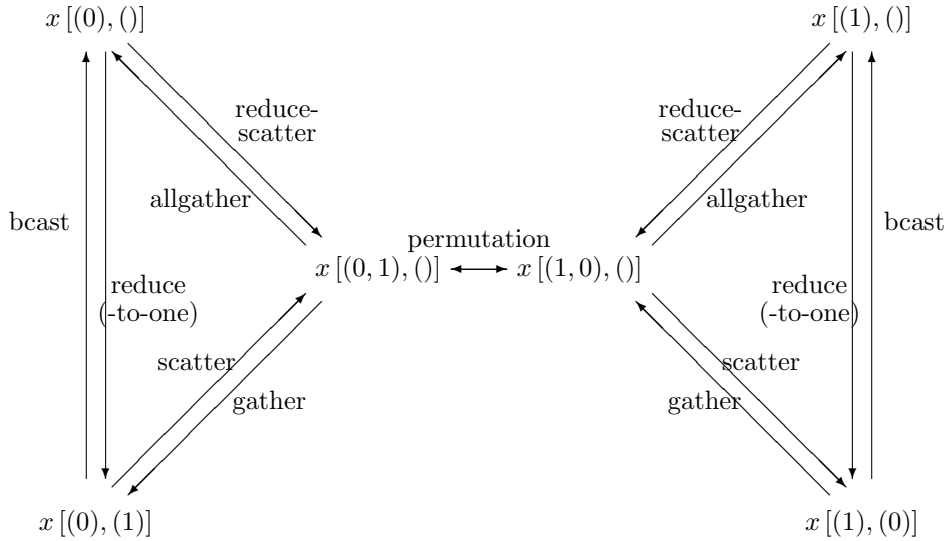


FIG. 5.1. Summary of the communication patterns for redistributing a vector x . For instance, a method for redistributing x from a matrix column to a matrix row is found by tracing from the bottom-left to the bottom-right of the diagram.

The members of any column-major vector distribution, $\mathcal{D}_{(0,1)}$, or row-major vector distribution, $\mathcal{D}_{(1,0)}$, form a partition of \mathbb{N} . The names *column-major vector distribution* and *row-major vector distribution* are derived from the fact that the mappings $(s_0, s_1) \mapsto s_0 + s_1 d_0$ and $(s_0, s_1) \mapsto s_1 + s_0 d_1$, respectively, label the $d_0 \times d_1$ grid with a column-major and row-major ordering.

As row-major and column-major distributions differ only by which dimension of the grid is considered first when assigning an order to the processes in the grid, we can give one general definition for a vector distribution with 2D grids. We give this definition now.

DEFINITION 5.4 (vector distribution). *We call the $d_0 \times d_1$ array $\mathcal{D}_{(i,j)}$ a vector distribution if $i, j \in \{0, 1\}$, $i \neq j$, and there exists some alignment parameter $\sigma \in \{0, \dots, p - 1\}$ such that, for every grid position $(s_0, s_1) \in \{0, \dots, d_0 - 1\} \times \{0, \dots, d_1 - 1\}$,*

$$(5.1) \quad \mathcal{D}_{(i,j)}(s_0, s_1) = \mathcal{V}_p^\sigma(s_i + s_j d_i).$$

The shorthand $y[(i, j)]$ will refer to the vector y distributed such that process (s_0, s_1) stores $y(\mathcal{D}_{(i,j)}(s_0, s_1))$.

Figure 4.1 illustrates that to redistribute $y[(0, 1)]$ to $y[(1, 0)]$, and vice versa, requires a permutation communication (simultaneous point-to-point communication). The effect of this redistribution can be seen in Figure 5.1. Via a permutation communication, the vector y distributed as $y[(0, 1)]$ can be redistributed as $y[(1, 0)]$, which is the same distribution as the vector x .

In the preceding discussions, our definitions of $\mathcal{D}_{(0,1)}$ and $\mathcal{D}_{(1,0)}$ allowed for arbitrary alignment parameters. In the rest of the paper, we will treat only the case where all alignments are zero; i.e., the top-left entry of every (global) matrix and the top entry of every (global) vector are owned by the process in the top-left of the process grid.

TABLE 5.1

The relationships between distribution symbols found in the Elemental library implementation and those introduced here. For instance, the distribution $A[M_C, M_R]$ found in the Elemental library implementation corresponds to the distribution $A[(0), (1)]$.

Elemental symbol	Introduced symbol
M_C	(0)
M_R	(1)
V_C	(0, 1)
V_R	(1, 0)
*	()

5.2. Induced matrix distribution. We are now ready to discuss how matrix distributions are induced by the vector distributions. For this, it pays to again consider Figure 4.1. The element $\alpha_{i,j}$ of matrix A is assigned to the row of processes in which ψ_i exists and to the column of processes in which χ_j exists. This means that in $y = Ax$, elements of x need only be communicated within columns of processes, and local contributions to y need only be summed within rows of processes. This induces a Cartesian matrix distribution: Column j of A is assigned to the same column of processes as χ_j . Row i of A is assigned to the same row of processes as ψ_i . We now answer the following two related questions: (1) What is the set $\mathcal{D}_{(0)}(s_0)$ of matrix row indices assigned to process row s_0 ? (2) What is the set $\mathcal{D}_{(1)}(s_1)$ of matrix column indices assigned to process column s_1 ?

DEFINITION 5.5. *Let*

$$\mathcal{D}_{(0)}(s_0) = \bigcup_{s_1=0}^{d_1-1} \mathcal{D}_{(0,1)}(s_0, s_1) \quad \text{and} \quad \mathcal{D}_{(1)}(s_1) = \bigcup_{s_0=0}^{d_0-1} \mathcal{D}_{(1,0)}(s_0, s_1).$$

Given matrix A , $A[\mathcal{D}_{(0)}(s_0), \mathcal{D}_{(1)}(s_1)]$ denotes the submatrix of A with row indices in the set $\mathcal{D}_{(0)}(s_0)$ and column indices in $\mathcal{D}_{(1)}(s_1)$. Finally, $A[(0), (1)]$ denotes the distribution of A that assigns $A[\mathcal{D}_{(0)}(s_0), \mathcal{D}_{(1)}(s_1)]$ to process (s_0, s_1) .

We say that $\mathcal{D}_{(0)}$ and $\mathcal{D}_{(1)}$ are *induced*, respectively, by $\mathcal{D}_{(0,1)}$ and $\mathcal{D}_{(1,0)}$ because the process to which $\alpha_{i,j}$ is assigned is determined by the row of processes, s_0 , to which y_i is assigned, and the column of processes, s_1 , to which x_j is assigned, so that it is ensured that in the matrix-vector multiplication $y = Ax$ communication need only be within rows and columns of processes. Notice in Figure 5.1 that to redistribute indices of the vector y as the matrix column indices in A requires a communication within rows of processes. Similarly, to redistribute indices of the vector x as matrix row indices requires a communication within columns of processes. The above definition lies at the heart of our communication scheme.

5.3. Vector duplication. Two vector distributions, encountered in section 4.4 and illustrated in Figure 4.2, still need to be specified with our notation. The vector x , duplicated as needed for the matrix-vector multiplication $y = Ax$, can be specified as $x[(0)]$ or, viewing x as an $n \times 1$ matrix, as $x[(0), ()]$. The vector y , duplicated so as to store local contributions for $y = Ax$, can be specified as $y[(1)]$ or, viewing y as an $n \times 1$ matrix, as $y[(1), ()]$. Here the $()$ should be interpreted as “all indices.” In other words, $\mathcal{D}_{()} \equiv \mathbb{N}$.

5.4. Notation in the Elemental library. Readers familiar with the Elemental library will notice that the distribution symbols defined within that library’s

implementation follow a different convention than that used for distribution symbols introduced in the previous subsections. This is due to the fact that the notation used in this paper was devised after the implementation of the Elemental library, and we wanted the notation to be extensible to higher-dimensional objects (tensors). However, for every symbol utilized in the Elemental library implementation, there exists a unique symbol in the notation introduced here. In Table 5.1, the relationships between distribution symbols utilized in the Elemental library implementation and the symbols used in this paper are defined.

5.5. Of vectors, columns, and rows. A matrix-vector multiplication or rank-1 update may take as its input/output vectors (x and y) the rows and/or columns of matrices, as we will see in section 6. This motivates us to briefly discuss the different communications needed to redistribute vectors to and from columns and rows. For our discussion, the reader may find it helpful to refer back to Figures 4.1 and 4.2.

Column to/from column-major vector. Consider Figure 4.1 and let a_j be a typical column in A . It exists within one single process column. Redistributing $a_j [(0), (1)]$ to $y [(0, 1), ()]$ requires simultaneous scatters within process rows. Inversely, redistributing $y [(0, 1), ()]$ to $a_j [(0), (1)]$ requires simultaneous gathers within process rows.

Column to/from row-major vector. Redistributing $a_j [(0), (1)]$ to $x [(1, 0), ()]$ can be accomplished by first redistributing to $y [(0, 1), ()]$ (simultaneous scatters within rows) followed by a redistribution of $y [(0, 1), ()]$ to $x [(1, 0), ()]$ (a permutation). Redistributing $x [(1, 0)]$ to $a_j [(0), (1)]$ reverses these communications.

Column to/from column projected vector. Redistributing $a_j [(0), (1)]$ to $a_j [(0), ()]$ (duplicated y in Figure 4.2) can be accomplished by first redistributing to $y [(0, 1), ()]$ (simultaneous scatters within rows) followed by a redistribution of $y [(0, 1), ()]$ to $y [(0), ()]$ (simultaneous allgathers within rows). However, recognize that a scatter followed by an allgather is equivalent to a broadcast. Thus, redistributing $a_j [(0), (1)]$ to $a_j [(0), ()]$ can be more directly accomplished by broadcasting within rows. Similarly, summing duplicated vectors $y [(0), ()]$ that leaves the result as $a_j [(0), (1)]$ (a column in A) can be accomplished by first summing them into $y [(0, 1), ()]$ (reduce-scatters within rows) followed by a redistribution to $a_j [(0), (1)]$ (gather within rows). But a reduce-scatter followed by a gather is equivalent to a reduce(-to-one) collective communication.

All communication patterns with vectors, rows, and columns. In Figure 5.1 we summarize all the communication patterns that will be encountered when performing various matrix-vector multiplications or rank-1 updates, with vectors, columns, or rows as input.

5.6. Parallelizing matrix-vector operations (revisited). We now show how the notation discussed in the previous subsection pays off when describing algorithms for matrix-vector operations.

Assume that A , x , and y are, respectively, distributed as $A [(0), (1)]$, $x [(1, 0), ()]$, and $y [(0, 1), ()]$. Algorithms for computing $y := Ax$ and $A := A + xy^T$ are given in Tables 5.2 and 5.3.

The discussion in section 5.5 provides insight into generalizing these parallel matrix-vector operations to the cases where the vectors are rows and/or columns of matrices. For example, in Table 5.4 we show how to compute a column of matrix C , \hat{c}_i as the product of a matrix A times the column of a matrix B , b_j . Certain steps in Tables 5.2–5.4 have superscripts associated with outputs of local computations. These superscripts indicate that contributions rather than final results are computed by the operation. Further, the subscript to $\widehat{\sum}$ indicates along which dimension of the

TABLE 5.2
Parallel algorithm for computing $y := Ax$.

Algorithm: $y := Ax$ (GEMV)	Comments
$x[(1), ()] \leftarrow x[(1, 0), ()]$	Redistribute x (allgather in columns)
$y^{(1)}[(0), ()] := A[(0), (1)] x[(1), ()]$	Local matrix-vector multiply
$y[(0, 1), ()] := \widehat{\sum}_1 y^{(1)}[(0), ()]$	Sum contributions (reduce-scatter in rows)

TABLE 5.3
Parallel algorithm for computing $A := A + xy^T$.

Algorithm $A := A + xy^T$ (GER)	Comments
$x[(0, 1), ()] \leftarrow x[(1, 0), ()]$	Redistribute x as a column-major vector (permutation)
$x[(0), ()] \leftarrow x[(0, 1), ()]$	Redistribute x (allgather in rows)
$y[(1, 0), ()] \leftarrow y[(0, 1), ()]$	Redistribute y as a row-major vector (permutation)
$y[(1), ()] \leftarrow y[(1, 0), ()]$	Redistribute y (allgather in columns)
$A[(0), (1)] := x[(0), ()][y[(1), ()]]^T$	Local rank-1 update

TABLE 5.4
Parallel algorithm for computing $\hat{c}_i := Ab_j$, where \hat{c}_i is a row of a matrix C and b_j is a column of a matrix B .

Algorithm: $\hat{c}_i := Ab_j$ (GEMV)	Comments
$x[(1), ()] \leftarrow b_j[(0), (1)]$	Redistribute b_j : $x[(0, 1), ()] \leftarrow b_j[(0), (1)]$ (scatter in rows) $x[(1, 0), ()] \leftarrow x[(0, 1), ()]$ (permutation) $x[(1), ()] \leftarrow x[(1, 0), ()]$ (allgather in columns)
$y^{(1)}[(0), ()] := A[(0), (1)] x[(1), ()]$	Local matrix-vector multiply
$\hat{c}_i[(0), (1)] := \widehat{\sum}_1 y^{(1)}[(0), ()]$	Sum contributions: $y[(0, 1), ()] := \widehat{\sum}_i y^{(1)}[(0), ()]$ (reduce-scatter in rows) $y[(1, 0), ()] \leftarrow y[(0, 1), ()]$ (permutation) $\hat{c}_i[(0), (1)] \leftarrow y[(1, 0), ()]$ (gather in rows)

processing grid a reduction of contributions must occur.

5.7. Similar operations. What we have described is a general method. We leave it as an exercise for the reader to derive parallel algorithms for $x := A^T y$ and $A := yx^T + A$, starting with vectors that are distributed in various ways.

6. Elemental SUMMA: 2D algorithms (eSUMMA2D). We have now arrived at the point where we can discuss parallel matrix-matrix multiplication on a $d_0 \times d_1$ mesh, with $p = d_0 d_1$. In our discussion, we will assume an elemental distribution, but the ideas clearly generalize to other Cartesian distributions.

This section exposes a systematic path from the parallel rank-1 update and matrix-vector multiplication algorithms to highly efficient 2D parallel matrix-matrix multiplication algorithms. The strategy is to first recognize that a matrix-matrix multiplication can be performed by a series of rank-1 updates or matrix-vector multiplications. This gives us parallel algorithms that are inefficient. By then recognizing that the order of operations can be changed so that communication and computation can be separated and consolidated, these inefficient algorithms are transformed into

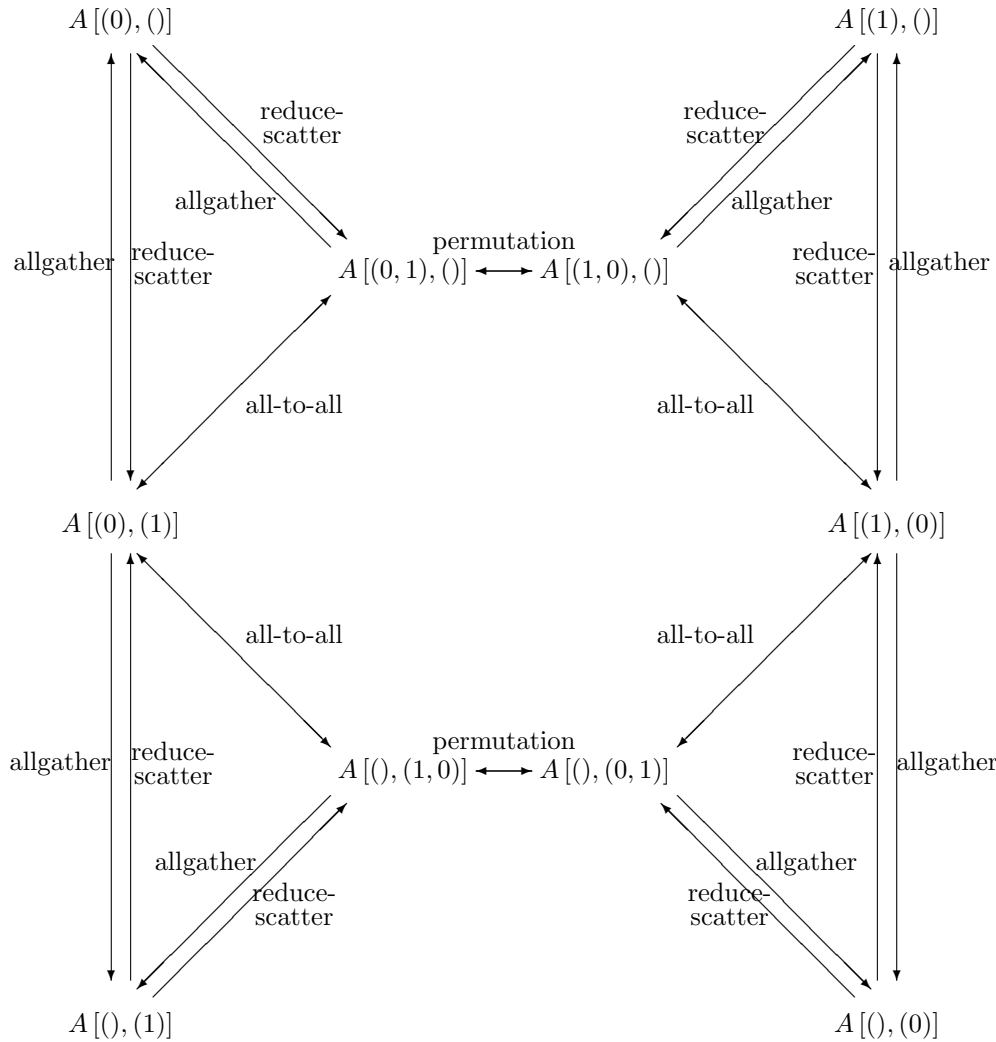


FIG. 6.1. Summary of the communication patterns for redistributing a matrix A .

efficient algorithms. While explained only for some of the cases of matrix multiplication, we believe the exposition is such that the reader can derive algorithms for the remaining cases by applying the ideas in a straightforward manner.

To fully understand how to attain high performance on a single processor, the reader should become familiar with, for example, the techniques in [16].

6.1. Elemental stationary C algorithms (eSUMMA2D-C). We first discuss the case where $C := C + AB$, where A and B have k columns each, with k relatively small.⁴ We call this a rank- k update or panel-panel multiplication [16]. We will assume the distributions $C[(0), (1)]$, $A[(0), (1)]$, and $B[(0), (1)]$. Partition

⁴There is an algorithmic block size, b_{alg} , for which a local rank- k update achieves peak performance [16]. Think of k as being that algorithmic block size for now.

$A = (a_0 \mid a_1 \mid \cdots \mid a_{k-1})$ and

$$B = \begin{pmatrix} \hat{b}_0^T \\ \hat{b}_1^T \\ \vdots \\ \hat{b}_{k-1}^T \end{pmatrix}$$

so that

$$C := ((\cdots ((C + a_0 \hat{b}_0^T) + a_1 \hat{b}_1^T) + \cdots) + a_{k-1} \hat{b}_{k-1}^T).$$

The following loop computes $C := AB + C$:

<pre> for $p = 0, \dots, k - 1$ $a_p [(0), ()] \leftarrow a_p [(0), (1)]$ $b_p^T [(), (1)] \leftarrow \hat{b}_p^T [(0), (1)]$ $C [(0), (1)] := C [(0), (1)] + a_p [(0), ()] \hat{b}_p^T [(), (1)]$ endfor </pre>	<p>(broadcasts within rows)</p> <p>(broadcasts within cols)</p> <p>(local rank-1 updates)</p>
--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------

While section 5.6 gives a parallel algorithm for GER, the problem with this algorithm is that (1) it creates a lot of messages and (2) the local computation is a rank-1 update, which inherently does not achieve high performance since it is memory bandwidth bound. The algorithm can be rewritten as

<pre> for $p = 0, \dots, k - 1$ $a_p [(0), ()] \leftarrow a_p [(0), (1)]$ endfor for $p = 0, \dots, k - 1$ $b_p^T [(), (1)] \leftarrow \hat{b}_p^T [(0), (1)]$ endfor for $p = 0, \dots, k - 1$ $C [(0), (1)] := C [(0), (1)] + a_p [(0), ()] \hat{b}_p^T [(), (1)]$ endfor </pre>	<p>(broadcasts within rows)</p> <p>(broadcasts within cols)</p> <p>(local rank-1 updates)</p>
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------

and finally, equivalently, as

<pre> $A [(0), ()] \leftarrow A [(0), (1)]$ $B [(), (1)] \leftarrow B [(0), (1)]$ $C [(0), (1)] := C [(0), (1)] + A [(0), ()] B [(), (1)]$ </pre>	<p>(allgather within rows)</p> <p>(allgather within cols)</p> <p>(local rank-k update)</p>
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------------------------------------------------------------------------------------------

Now the local computation is cast in terms of a local matrix-matrix multiplication (rank-k update), which can achieve high performance. Here (given that we assume elemental distribution) $A [(0), ()] \leftarrow A [(0), (1)]$ within each row broadcasts k columns of A from different roots: an allgather if elemental distribution is assumed! Similarly, $B [(), (1)] \leftarrow B [(0), (1)]$ within each column broadcasts k rows of B from different roots: another allgather if elemental distribution is assumed!

Based on this observation, the SUMMA-like algorithm can be expressed as a loop around such rank-k updates, as given in Table 6.1 (left).⁵ The purpose of the loop is to reduce workspace required to store duplicated data. Notice that if an

⁵We use FLAME notation to express the algorithm, which we have used in our papers for more than a decade [18].

TABLE 6.1

Algorithms for computing $C := AB + C$. Left: Stationary C . Right: Stationary A .

<p>Algorithm: $C := \text{GEMM_C}(C, A, B)$</p> <p>Partition $A \rightarrow (A_L \mid A_R)$, $B \rightarrow \left(\frac{B_T}{B_B}\right)$ where A_L has 0 columns, B_T has 0 rows while $n(A_L) < n(A)$ do Determine block size b Repartition $(A_L \mid A_R) \rightarrow (A_0 \mid A_1 \mid A_2)$, $\left(\frac{B_T}{B_B}\right) \rightarrow \left(\frac{B_0}{B_1} \mid \frac{B_2}{B_2}\right)$ where A_1 has b columns, B_1 has b rows</p> <hr/> <p>$A_1[(0), (0)] \leftarrow A_1[(0), (1)]$ $B_1[(0), (1)] \leftarrow B_1[(0), (1)]$ $C[(0), (1)] := C[(0), (1)]$ $+ A_1[(0), (0)] B_1[(0), (1)]$</p> <hr/> <p>Continue with $(A_L \mid A_R) \leftarrow (A_0 \mid A_1 \mid A_2)$, $\left(\frac{B_T}{B_B}\right) \leftarrow \left(\frac{B_0}{B_1} \mid \frac{B_2}{B_2}\right)$</p> <p>endwhile</p>	<p>Algorithm: $C := \text{GEMM_A}(C, A, B)$</p> <p>Partition $C \rightarrow (C_L \mid C_R)$, $B \rightarrow (B_L \mid B_R)$ where C_L and B_L have 0 columns while $n(C_L) < n(C)$ do Determine block size b Repartition $(C_L \mid C_R) \rightarrow (C_0 \mid C_1 \mid C_2)$, $(B_L \mid B_R) \rightarrow (B_0 \mid B_1 \mid B_2)$</p> <p style="text-align: center;">where C_1 and B_1 have b columns</p> <hr/> <p>$B_1[(1), (0)] \leftarrow B_1[(0), (1)]$ $C_1^{(1)}[(0), (0)] := A[(0), (1)] B_1[(1), (0)]$ $C_1[(0), (1)] := \widehat{\sum}_1 C_1^{(1)}[(0), (0)]$</p> <hr/> <p>Continue with $(C_L \mid C_R) \leftarrow (C_0 \mid C_1 \mid C_2)$, $(B_L \mid B_R) \leftarrow (B_0 \mid B_1 \mid B_2)$</p> <p>endwhile</p>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

elemental distribution is assumed, the SUMMA-like algorithm should *not* be called a broadcast-broadcast-compute algorithm. Instead, it becomes an allgather-allgather-compute algorithm. We will also call it a stationary C algorithm, since C is not communicated (and hence “owner computes” is determined by what processor owns what element of C). The primary benefit from having a loop around rank- k updates is that it reduces the required local workspace at the expense of an increase only in the α term of the communication cost.

We label this algorithm *eSUMMA2D-C*, an elemental SUMMA-like algorithm targeting a 2D mesh of nodes, stationary C variant. It is not hard to extend the insights to nonelemental distributions (for example, as used by ScaLAPACK or PLAPACK).

An approximate cost for the described algorithm is given by

$$\begin{aligned}
 T_{\text{eSUMMA2D-C}}(m, n, k, d_0, d_1) &= \frac{2mnk}{p} \gamma + \frac{k}{b_{\text{alg}}} \log_2(d_1) \alpha + \frac{d_1 - 1}{d_1} \frac{mk}{d_0} \beta + \frac{k}{b_{\text{alg}}} \log_2(d_0) \alpha + \frac{d_0 - 1}{d_0} \frac{nk}{d_1} \beta \\
 &= \frac{2mnk}{p} \gamma + \underbrace{\frac{k}{b_{\text{alg}}} \log_2(p) \alpha + \frac{(d_1 - 1)mk}{p} \beta + \frac{(d_0 - 1)nk}{p} \beta}_{T^+_{\text{eSUMMA2D-C}}(m, n, k, d_0, d_1)}.
 \end{aligned}$$

This estimate ignores load imbalance (which leads to a γ term of the same order as the β terms) and the fact that the allgathers may be unbalanced if b_{alg} is not an integer multiple of both d_0 and d_1 . As before, and throughout this paper, T^+ refers to the communication overhead of the proposed algorithm (e.g., $T^+_{\text{eSUMMA2D-C}}$ refers to the communication overhead of the *eSUMMA2D-C* algorithm).

It is not hard to see that for practical purposes,⁶ the weak scalability of the eSUMMA2D-C algorithm mirrors that of the parallel matrix-vector multiplication algorithm analyzed in Appendix A: it is weakly scalable when $m = n$ and $d_0 = d_1$, for arbitrary k .

At this point it is important to mention that this resulting algorithm may seem similar to an approach described in prior work [2]. Indeed, this allgather-allgather-compute approach to parallel matrix-matrix multiplication is described in that paper for the matrix-matrix multiplication variants $C = AB$, $C = AB^T$, $C = A^T B$, and $C = A^T B^T$ under the assumption that all matrices are approximately the same size; this is a surmountable limitation. As we have argued previously, the allgather-allgather-compute approach is particularly well suited for situations where we wish not to communicate the matrix C . In the next section, we describe how to systematically derive algorithms for situations where we wish to avoid communicating the matrix A .

6.2. Elemental stationary A algorithms (eSUMMA2D-A). Next, we discuss the case where $C := C + AB$, where C and B have n columns each, with n relatively small. For simplicity, we also call that parameter b_{alg} . We call this a matrix-panel multiplication [16]. We again assume that the matrices are distributed as $C[(0), (1)]$, $A[(0), (1)]$, and $B[(0), (1)]$. Partition

$$C = (c_0 \mid c_1 \mid \cdots \mid c_{n-1}) \quad \text{and} \quad B = (b_0 \mid b_1 \mid \cdots \mid b_{n-1})$$

so that $c_j = Ab_j + c_j$. The following loop will compute $C = AB + C$:

<pre> for $j = 0, \dots, n - 1$ $b_j[(0), (1), ()] \leftarrow b_j[(0), (1)]$ $b_j[(1), (0), ()] \leftarrow b_j[(0), (1), ()]$ $b_j[(1), ()] \leftarrow b_j[(1), (0), ()]$ $c_j[(0), ()] := A[(0), (1)] b_j[(1), ()]$ $c_j[(0), (1)] \leftarrow \widehat{\sum}_1 c_j[(0), ()]$ endfor </pre>	<pre> (scatters within rows) (permutation) (allgathers within cols) (local matvec mult) (reduce-to-one within rows) </pre>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------

While section 5.6 gives a parallel algorithm for GEMV, the problem again is that (1) it creates a lot of messages and (2) the local computation is a matrix-vector multiply, which inherently does not achieve high performance since it is memory bandwidth bound. This can be restructured as

⁶The very slow growing factor $\log_p(p)$ prevents weak scalability unless it is treated as a constant.

for $j = 0, \dots, n - 1$	
$b_j [(0, 1), ()] \leftarrow b_j [(0), (1)]$	(scatters within rows)
endfor	
for $j = 0, \dots, n - 1$	
$b_j [(1, 0), ()] \leftarrow b_j [(0, 1), ()]$	(permutation)
endfor	
for $j = 0, \dots, n - 1$	
$b_j [(1), ()] \leftarrow b_j [(1, 0), ()]$	(allgathers within cols)
endfor	
for $j = 0, \dots, n - 1$	
$c_j [(0), ()] := A [(0), (1)] b_j [(1), ()]$	(local matvec mult)
endfor	
for $j = 0, \dots, n - 1$	
$c_j [(0), (1)] \leftarrow \widehat{\sum}_1 c_j [(0), ()]$	(simultaneous reduce-to-one within rows)
endfor	

and finally, equivalently, as

$B [(1), ()] \leftarrow B [(0), (1)]$	(all-to-all within rows, permutation, allgather within cols)
$C [(0), ()] := A [(0), (1)] B [(1), ()] + C [(0), ()]$	(simultaneous local matrix multiplications)
$C [(0), (1)] \leftarrow \widehat{\sum} C [(0), ()]$	(reduce-scatter within rows)

Now the local computation is cast in terms of a local matrix-matrix multiplication (matrix-panel multiply), which can achieve high performance. A stationary A algorithm for arbitrary n can now be expressed as a loop around such parallel matrix-panel multiplies, given in Table 6.1 (right).

An approximate cost for the described algorithm is given by

$$\begin{aligned}
T_{\text{eSUMMA2D-A}}(m, n, k, d_0, d_1) = & \\
& \frac{n}{b_{\text{alg}}} \log_2(d_1) \alpha + \frac{d_1-1}{d_1} \frac{nk}{d_0} \beta && \text{(all-to-all within rows)} \\
& + \frac{n}{b_{\text{alg}}} \alpha + \frac{n}{d_1} \frac{k}{d_0} \beta && \text{(permutation)} \\
& + \frac{n}{b_{\text{alg}}} \log_2(d_0) \alpha + \frac{d_0-1}{d_0} \frac{nk}{d_1} \beta && \text{(allgather within cols)} \\
& + \frac{2mnk}{p} \gamma && \text{(simultaneous local matrix-panel mult)} \\
& + \frac{n}{b_{\text{alg}}} \log_2(d_1) \alpha + \frac{d_1-1}{d_1} \frac{mn}{d_0} \beta + \frac{d_1-1}{d_1} \frac{mn}{d_0} \gamma && \text{(reduce-scatter within rows)}.
\end{aligned}$$

As we discussed earlier, the cost function for the all-to-all operation is somewhat suspect. Still, if an algorithm that attains the lower bound for the α term is employed, the β term must increase by at most a factor of $\log_2(d_1)$ [8], meaning that it is not the dominant communication cost. The estimate ignores load imbalance (which leads to a γ term of the same order as the β terms) and the fact that various collective communications may be unbalanced if b_{alg} is not an integer multiple of both d_0 and d_1 .

While the overhead is clearly greater than that of the eSUMMA2D-C algorithm when $m = n = k$, the overhead is *comparable* to that of the eSUMMA2D-C algorithm, so the weak scalability results are, asymptotically, the same. Also, it is not hard to see that if m and k are large while n is small, this algorithm achieves better parallelism, since less communication is required: The stationary matrix, A , is then the largest matrix, and not communicating it is beneficial. Similarly, if m and n are large while k

is small, then the eSUMMA2D-C algorithm does not communicate the largest matrix, C , which is beneficial.

6.3. Communicating submatrices. In Figure 6.1 we illustrate the collective communications required to redistribute submatrices from one distribution to another and the collective communications required to implement them.

6.4. Other cases. We leave it as an exercise for the reader to propose and analyze the remaining stationary A, B, and C algorithms for the other cases of matrix-matrix multiplication: $C := A^T B + C$, $C := AB^T + C$, and $C := A^T B^T + C$.

Hence, we have presented a systematic framework for deriving a family of parallel matrix-matrix multiplication algorithms.

7. Elemental SUMMA: 3D algorithms (eSUMMA3D). We now view the p processors as forming a $d_0 \times d_1 \times h$ mesh, which one should visualize as h stacked layers, where each layer consists of a $d_0 \times d_1$ mesh. The extra dimension is used to gain an extra level of parallelism, which reduces the overhead of the 2D SUMMA algorithms within each layer at the expense of communications between the layers.

The approach used to generalize Elemental SUMMA 2D algorithms to Elemental SUMMA 3D algorithms can be easily modified to use Cannon’s or Fox’s algorithm (with the constraints and complications that come from using those algorithms) or any other distribution for which SUMMA can be used (pretty much any Cartesian distribution).

7.1. 3D stationary C algorithms (eSUMMA3D-C). Partition A and B so that $A = (A_0 \mid \cdots \mid A_{h-1})$ and

$$B = \begin{pmatrix} B_0 \\ \vdots \\ B_{h-1} \end{pmatrix},$$

where A_p and B_p have approximately k/h columns and rows, respectively. Then

$$C + AB = \underbrace{(C + A_0 B_0)}_{\text{by layer 0}} + \underbrace{(0 + A_1 B_1)}_{\text{by layer 1}} + \cdots + \underbrace{(0 + A_{h-1} B_{h-1})}_{\text{by layer h-1}}.$$

This suggests the following 3D algorithm:

- Duplicate C to each of the layers, initializing the duplicates assigned to layers 1 through $h - 1$ to zero. This requires no communication. We will ignore the cost of setting the duplicates to zero.
- Scatter A and B so that layer H receives A_H and B_H . This means that all processors $(I, J, 0)$ simultaneously scatter approximately $(m + n)k/(d_0 d_1)$ data to processors $(I, J, 0)$ through $(I, J, h - 1)$. The cost of such a scatter can be approximated by

$$(7.1) \quad \log_2(h)\alpha + \frac{h-1}{h} \frac{(m+n)k}{d_0 d_1} \beta = \log_2(h)\alpha + \frac{(h-1)(m+n)k}{p} \beta.$$

- Compute $C := C + A_K B_K$ simultaneously on all layers. If eSUMMA2D-C is used for this in each layer, the cost is approximated by

$$(7.2) \quad 2 \frac{mnk}{p} \gamma + \underbrace{\frac{k}{hb_{\text{alg}}} (\log_2(p) - \log_2(h)) \alpha + \frac{(d_1 - 1)mk}{p} \beta + \frac{(d_0 - 1)nk}{p} \beta}_{T^+_{\text{eSUMMA2D-C}}(m, n, k/h, d_0, d_1)}.$$

- Perform reduce operations to sum the contributions from the different layers to the copy of C in layer 0. This means that contributions from processors $(I, J, 0)$ through (I, J, K) are reduced to processor $(I, J, 0)$. An estimate for this reduce-to-one is

$$(7.3) \quad \log_2(h)\alpha + \frac{mn}{d_0d_1}\beta + \frac{mn}{d_0d_1}\gamma = \log_2(h)\alpha + \frac{mnh}{p}\beta + \frac{mnh}{p}\gamma.$$

Thus, an estimate for the total cost of this eSUMMA3D-C algorithm for this case of `gemm` results from adding (7.1)–(7.3).

Let us analyze in detail the case where $m = n = k$ and $d_0 = d_1 = \sqrt{p/h}$. The cost becomes

$$\begin{aligned} C_{\text{eSUMMA3D-C}}(n, n, n, d_0, d_0, h) &= 2\frac{n^3}{p}\gamma + \frac{n}{hb_{\text{alg}}}(\log_2(p) - \log_2(h))\alpha + 2\frac{(d_0 - 1)n^2}{p}\beta + \log_2(h)\alpha + 2\frac{(h - 1)n^2}{p}\beta \\ &\quad + \log_2(h)\alpha + \frac{n^2h}{p}\beta + \frac{n^2h}{p}\gamma \\ &= 2\frac{n^3}{p}\gamma + \left[\frac{n}{hb_{\text{alg}}}(\log_2(p) - \log_2(h)) + 2\log_2(h) \right] \alpha + \left[2\left(\frac{\sqrt{p}}{\sqrt{h}} - 1 \right) + 3h - 2 + \frac{\gamma}{\beta}h \right] \frac{n^2}{p}\beta. \end{aligned}$$

Now, let us assume that the α term is inconsequential (which will be true if n is large enough). Then the minimum can be computed by taking the derivative (with respect to h) and setting this to zero: $-\sqrt{p}h^{-3/2} + (3 + K) = 0$ or $h = ((3 + K)/\sqrt{p})^{-2/3} = \sqrt[3]{p}/(3 + K)^{2/3}$, where $K = \gamma/\beta$. Typically $\gamma/\beta \ll 1$, and hence $(3 + K)^{-2/3} \approx 3^{-2/3} \approx 1/2$, meaning that the optimal h is given by $h \approx \sqrt[3]{p}/2$. Of course, details of how the collective communication algorithms are implemented will affect this optimal choice. Moreover, α is typically four to five orders of magnitude greater than β , and hence the α term cannot be ignored for more moderate matrix sizes, greatly affecting the analysis.

While the cost analysis assumes the special case where $m = n = k$ and $d_0 = d_1$, and that the matrices are perfectly balanced among the $d_0 \times d_0$ mesh, the description of the algorithm is general. It is merely the case that the cost analysis for the more general case becomes more complex.

The algorithm and the related insights are similar to those described in Agarwal et al. [1], although we arrive at this algorithm via a different path.

Now, PLAPACK and Elemental both include stationary C algorithms for the other cases of matrix multiplication ($C := \alpha A^T B + \beta C$, $C := \alpha A B^T + \beta C$, and $C := \alpha A^T B^T + \beta C$). Clearly, 3D algorithms that utilize these implementations can be easily proposed. For example, if $C := A^T B^T + C$ is to be computed, one can partition

$$A = \begin{pmatrix} A_0 \\ \vdots \\ A_{h-1} \end{pmatrix} \quad \text{and} \quad B = (B_0 \mid \cdots \mid B_{h-1}),$$

after which

$$C + A^T B^T = \underbrace{(C + A_0^T B_0^T)}_{\text{by layer 0}} + \underbrace{(0 + A_1^T B_1^T)}_{\text{by layer 1}} + \cdots + \underbrace{(0 + A_{h-1}^T B_{h-1}^T)}_{\text{by layer h-1}}.$$

The communication overhead for all four cases is similar, meaning that for all four cases, the resulting stationary C 3D algorithms have similar properties.

7.2. Stationary A algorithms (eSUMMA3D-A). Let us next focus on $C := AB + C$. Algorithms such that A is the stationary matrix are implemented in PLAPACK and Elemental. They have costs similar to that of the eSUMMA2D-C algorithm.

Let us describe a 3D algorithm, with a $d_0 \times d_1 \times h$ mesh, again viewed as h layers. If we partition, conformally, C and B so that $C = (C_0 \mid \cdots \mid C_{h-1})$ and $B = (B_0 \mid \cdots \mid B_{h-1})$, then

$$\underbrace{(C_0 := C_0 + AB_0 \mid C_1 := C_1 + AB_1 \mid \cdots \mid C_{h-1} := C_{h-1} + AB_{h-1})}_{\text{by layer 0}}$$

$$\underbrace{)}_{\text{by layer } h-1}$$

This suggests the following 3D algorithm:

- Duplicate (broadcast) A to each of the layers. If matrix A is perfectly balanced among the processors, the cost of this can be approximated by

$$\log_2(h)\alpha + \frac{mk}{d_0d_1}\beta.$$

- Scatter C and B so that layer K receives C_K and B_K . This means having all processors $(I, J, 0)$ simultaneously scatter approximately $(mn + nk)/(d_0d_1)$ data to processors $(I, J, 0)$ through $(I, J, h - 1)$. The cost of such a scatter can be approximated by

$$\log_2(h)\alpha + \frac{h-1}{h} \frac{(m+k)n}{d_1d_0}\beta = \log_2(h)\alpha + \frac{(h-1)(m+k)n}{p}\beta.$$

- Compute $C_K := C_K + AB_K$ simultaneously on all layers with a 2D stationary A algorithm. The cost of this is approximated by

$$\frac{2mnk}{p}\gamma + T^+_{\text{eSUMMA2D-A}}(m, n/h, k, d_0, d_1).$$

- Gather the C_K submatrices to layer 0. The cost of such a gather can be approximated by

$$\log_2(h)\alpha + \frac{h-1}{h} \frac{mn}{d_1d_0}\beta.$$

Rather than giving the total cost, we merely note that the stationary A 3D algorithms can similarly be stated for general m, n, k, d_0 , and d_1 , and that then the costs are similar.

Now, PLAPACK and Elemental both include stationary A algorithms for the other cases of matrix multiplication. Again, 3D algorithms that utilize these implementations can be easily proposed.

7.3. Stationary B algorithms (eSUMMA3D-B). Finally, let us again focus on $C := AB + C$. Algorithms such that B is the stationary matrix are also implemented in PLAPACK and Elemental. They also have costs similar to that of the SUMMA algorithm for $C := AB + C$.

Let us describe a 3D algorithm, with a $d_0 \times d_1 \times h$ mesh, again viewed as h layers. If we partition, conformally, C and A so that

$$C = \begin{pmatrix} C_0 \\ \vdots \\ C_{h-1} \end{pmatrix}$$

and

$$A = \begin{pmatrix} A_0 \\ \vdots \\ A_{h-1} \end{pmatrix},$$

then

$$\begin{pmatrix} C_0 + A_0B \\ C_1 + A_1B \\ \vdots \\ C_{h-1} := C_{h-1} + A_{h-1}B \end{pmatrix} \begin{array}{l} \text{by layer 0} \\ \text{by layer 1} \\ \vdots \\ \text{by layer } h-1. \end{array}$$

This suggests the following 3D algorithm:

- Duplicate (broadcast) B to each of the layers. If matrix B is perfectly balanced among the processors, the cost can be approximated by

$$\log_2(h)\alpha + \frac{nk}{d_0d_1}\beta.$$

- Scatter C and A so that layer K receives C_K and A_K . This means having all processors $(I, J, 0)$ simultaneously scatter approximately $(mn + mk)/(d_0d_1)$ data to processors $(I, J, 0)$ through $(I, J, h - 1)$. The cost of such a scatter can be approximated by

$$\log_2(h)\alpha + \frac{h-1}{h} \frac{m(n+k)}{d_1d_0}\beta = \log_2(h)\alpha + \frac{(h-1)m(n+k)}{p}\beta.$$

- Compute $C_K := C_K + A_KB$ simultaneously on all of the layers with a 2D stationary B algorithm. The cost of this is approximated by

$$\frac{2mnk}{p}\gamma + T^+_{\text{eSUMMA2D-B}}(m/h, n, k, d_0, d_1).$$

- Gather the C_K submatrices to layer 0. The cost of such a gather can be approximated by

$$\log_2(h)\alpha + \frac{h-1}{h} \frac{mn}{d_1d_0}\beta.$$

Again, a total cost similar to those for stationary C and A algorithms results. Again, PLAPACK and Elemental both include stationary B algorithms for the other cases of matrix multiplication. Again, 3D algorithms that utilize these implementations can be easily proposed.

7.4. Other cases. We leave it as an exercise to the reader to propose and analyze the remaining eSUMMA3D-A, eSUMMA3D-B, and eSUMMA3D-C algorithms for the other cases of matrix-matrix multiplication: $C := A^T B + C$, $C := AB^T + C$, and $C := A^T B^T + C$.

The point is that we have presented a systematic framework for deriving a family of parallel 3D matrix-matrix multiplication algorithms.

7.5. Discussion. This extra level of parallelism gained with 3D SUMMA algorithms allows us to parallelize computation across any one of three dimensions involved in the matrix-matrix multiplication (two dimensions for forming the output C , and one for the reduction of A and B). The particular 3D SUMMA algorithmic variant dictates the dimension along which the extra parallelism occurs. In [13], a

geometric model is developed that views the set of scalar computations associated with a matrix-matrix multiplication as a set of lattice points forming a rectangular prism. This geometric model is based on the Loomis–Whitney inequality [26] that has been used to devise algorithms that achieve the parallel bandwidth cost lower bound for matrix-matrix multiplication [4, 5]. Considering this geometric model, each 3D SUMMA algorithmic variant corresponds to performing computations appearing in different slices⁷ in parallel. The orientation of slices is dictated by the 3D SUMMA algorithmic variant chosen, and the order in which computations are performed within a slice is dictated by the 2D SUMMA algorithm used within each layer of the processing mesh. We now discuss how the communication overhead of Elemental 2D and 3D SUMMA algorithms relates to the lower bounds of both the latency and bandwidth costs associated with parallel matrix-matrix multiplication.

In [23], it was shown that the lower bound on communicated data is $\Omega(n^2/\sqrt{p})$ for a matrix multiplication of two $n \times n$ matrices computed on a processing grid involving p processes arranged as a 2D mesh and $\Omega(n^2/\sqrt[3]{p^2})$ for a matrix multiplication of two $n \times n$ matrices computed on a processing grid involving p processes arranged as a 3D mesh. Examination of the cost functions associated with each eSUMMA2D algorithm and eSUMMA3D algorithm shows that each achieves the lower bound on communication for such an operation.

With regard to latency, the lower bound on the number of messages required is $\Omega(\log(p))$ for a matrix multiplication of two $n \times n$ matrices computed on a processing grid involving p processes arranged as either a 2D or 3D mesh. Examination of the cost functions shows that each achieves the lower bound on latency as well if we assume that the algorithmic block size $b_{alg} = n$. Otherwise, the proposed algorithms do not achieve the lower bound.

8. Performance experiments. In this section, we present performance results that support the insights in the previous sections. Implementations of the eSUMMA-2D algorithms are all part of the Elemental library. The eSUMMA-3D algorithms were implemented with Elemental, building upon its eSUMMA-2D algorithms and implementations. In all of these experiments, it was assumed that the data started and finished distributed within one layer of the 3D mesh of nodes so that all communication necessary to duplicate was included in the performance calculations.

As in [28, 29], performance experiments were carried out on the IBM Blue Gene/P architecture with compute nodes that consist of four 850 MHz PowerPC 450 processors for a combined theoretical peak performance of 13.6 GFlops per node using double-precision arithmetic. Nodes are interconnected by a 3D torus topology and a collective network, each of which supports a per-node bidirectional bandwidth of 2.55 GB/s. In all graphs, the top of the graph represents peak performance for this architecture so that the attained efficiency can be easily judged.

The point of the performance experiments was to demonstrate the merits of 3D algorithms. For this reason, we simply fixed the algorithmic block size, b_{alg} , to 128 for all experiments. The number of nodes, p , was chosen to be various powers of two, as was the number of layers, h . As a result, the $d_0 \times d_1$ mesh for a single layer was chosen so that $d_0 = d_1$ if p/h was a perfect square and $d_0 = d_1/2$ otherwise. The “zig-zagging” observed in some of the curves is attributed to this square vs. nonsquare choice of $d_0 \times d_1$. It would have been tempting to perform exhaustive experiments with various algorithmic block sizes and mesh configurations. However, the performance

⁷As used, the term “slice” refers to a set of “superbricks” in [13].

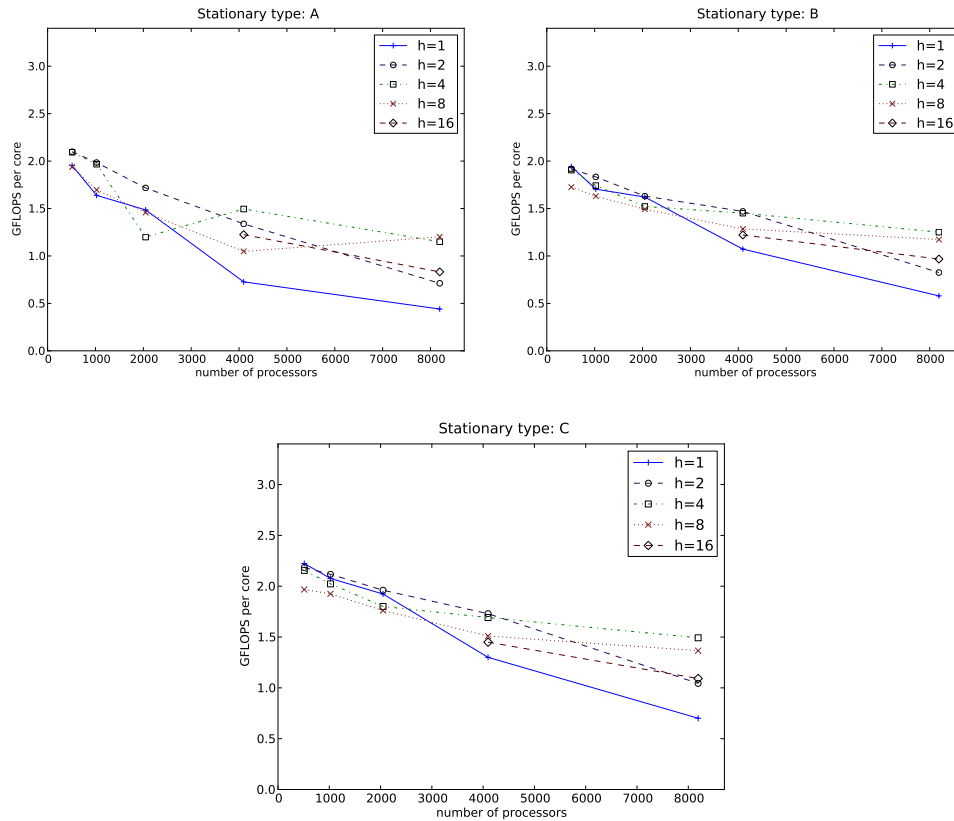


FIG. 8.1. Performance of the different implementations when $m = n = k = 30,000$ and the number of nodes is varied.

results were merely meant to verify that the insights of the previous sections have merit.

In our implementations, the eSUMMA3D-X algorithms utilize eSUMMA2D-X algorithms on each of the layers, where $X \in \{A, B, C\}$. As a result, the curve for eSUMMA3D-X with $h = 1$ is also the curve for the eSUMMA2D-X algorithm.

Figure 8.1 illustrates the benefits of the 3D algorithms. Inherently, efficiency cannot be maintained when the problem size is fixed. In other words, “strong” scaling is unattainable. Still, by increasing the number of layers, h , as the number of nodes, p , is increased, efficiency can be better maintained.

Figure 8.2 illustrates that the eSUMMA2D-C and eSUMMA3D-C algorithms attain high performance already when $m = n$ are relatively large and k is relatively small. This is not surprising: The eSUMMA2D-C algorithm already attains high performance when k is small because the “large” matrix C is not communicated, and the local matrix-matrix multiplication can already attain high performance when the local k is small (if the local m and n are relatively large).

Figure 8.3 similarly illustrates that the eSUMMA2D-A and eSUMMA3D-A algorithms attain high performance already when $m = k$ are relatively large and n is relatively small, and Figure 8.4 illustrates that the eSUMMA2D-B and eSUMMA3D-B algorithms attain high performance already when $n = k$ are relatively large and m is relatively small.

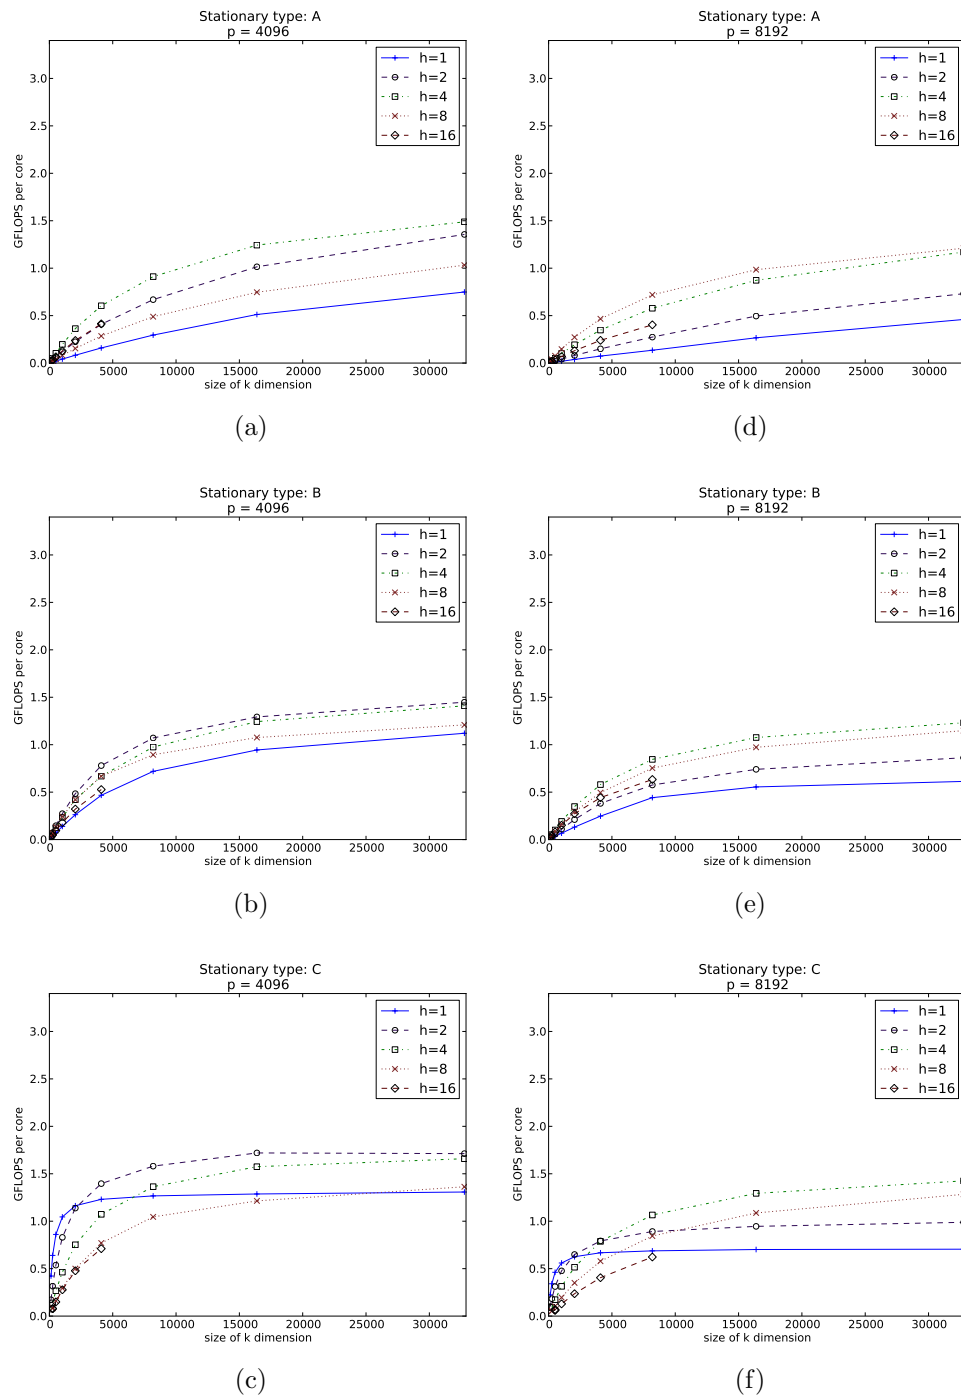


FIG. 8.2. Performance of the different implementations when $m = n = 30,000$ and k is varied. As expected, the stationary C algorithms ramp up to high performance faster than the other algorithms when k is small.

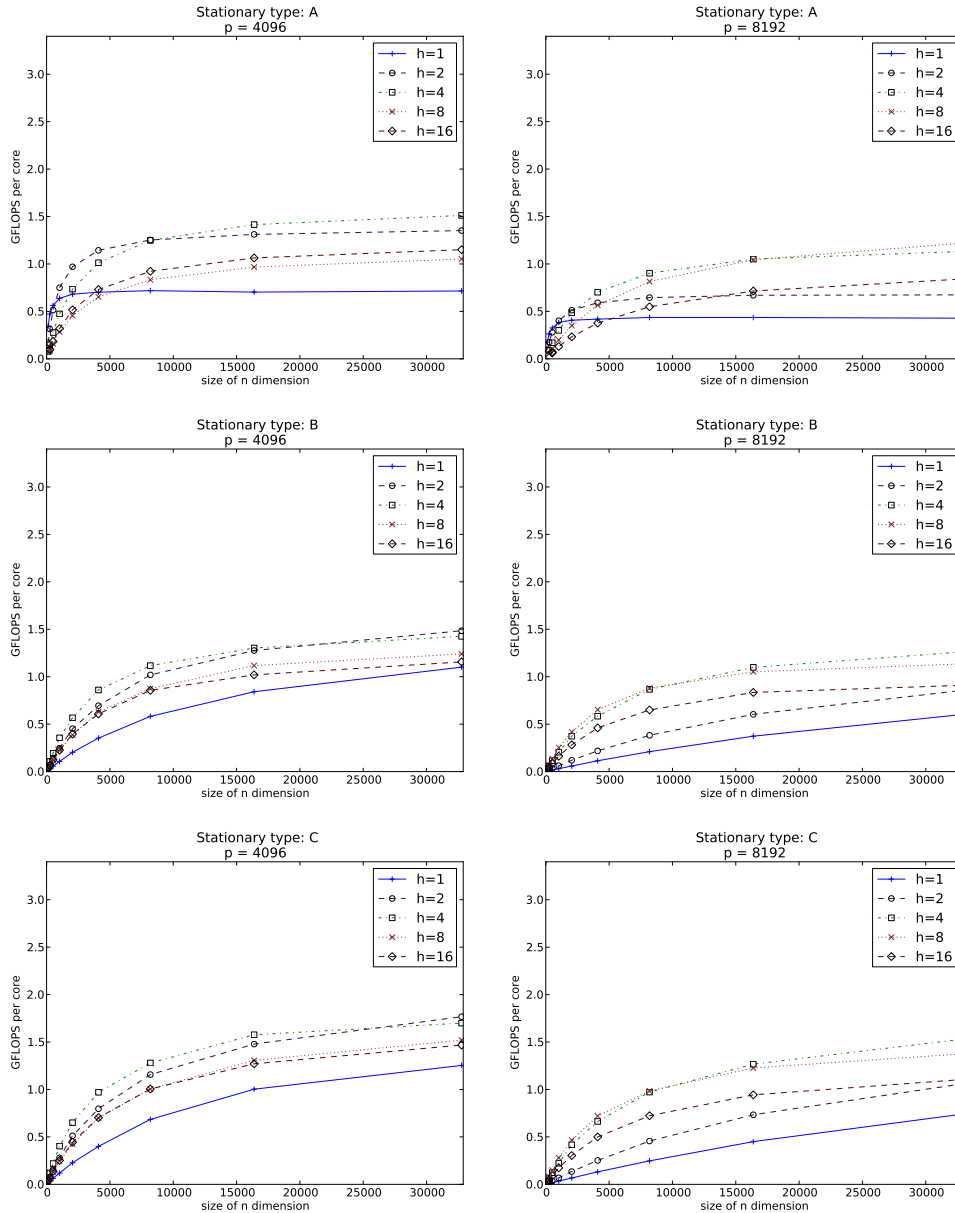


FIG. 8.3. Performance of the different implementations when $m = k = 30,000$ and n is varied. As expected, the stationary A algorithms ramp up to high performance faster than the other algorithms when n is small.

A comparison of Figures 8.2(c) and 8.3(a) shows that the eSUMMA2D-A algorithm (Figure 8.3(a) with $h = 1$) asymptotes sooner than the eSUMMA2D-C algorithm (Figure 8.2(c) with $h = 1$). The primary reason for this is that it incurs more communication overhead. But as a result, increasing h is more beneficial to eSUMMA3D-A in Figure 8.3(a) than is increasing h for eSUMMA3D-C in Figure 8.2(c). A similar observation can be made for eSUMMA2D-B and eSUMMA3D-B in Figure 8.4(b).

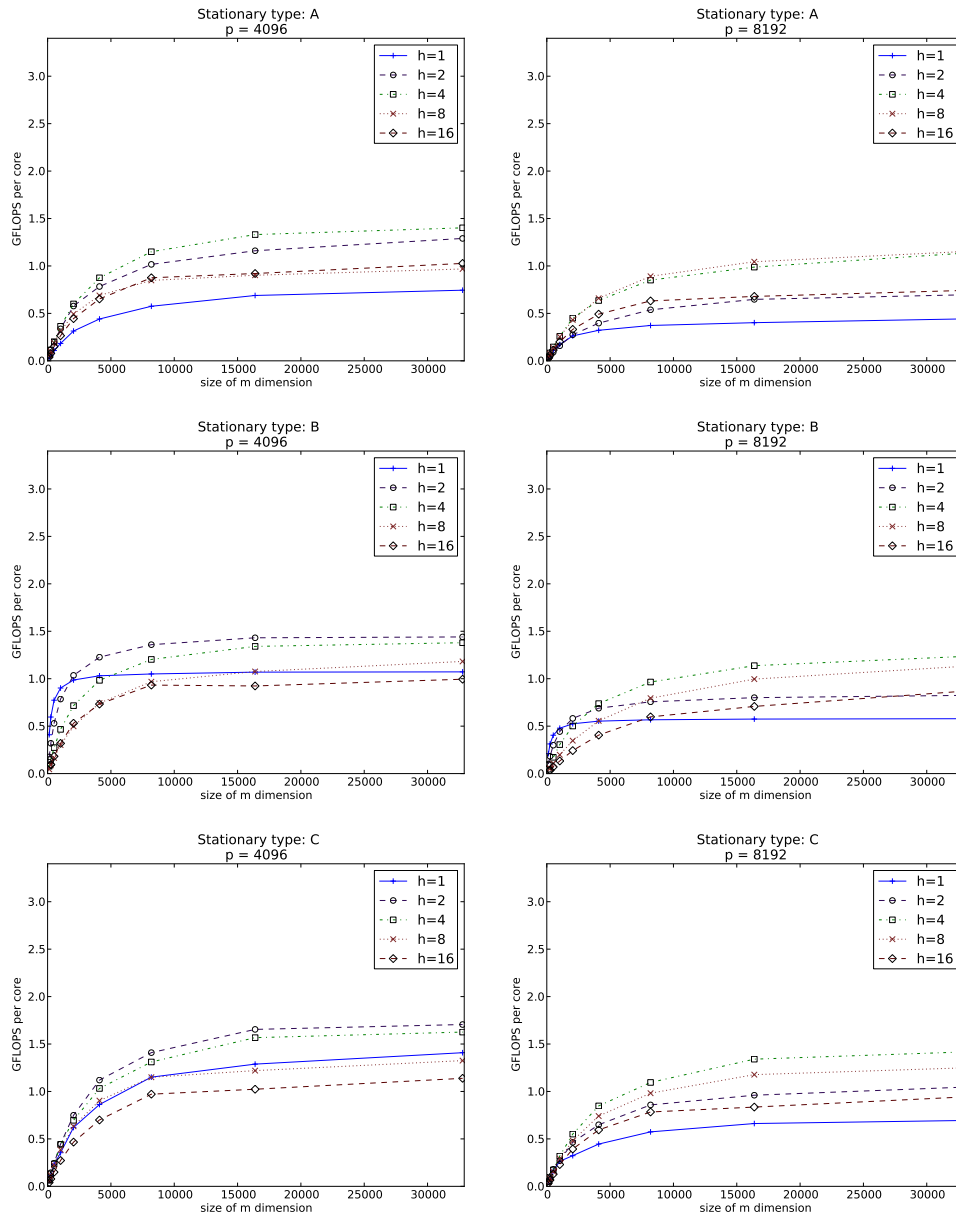


FIG. 8.4. Performance of the different implementations when $n = k = 30,000$ and m is varied. As expected, the stationary B algorithms ramp up to high performance faster than the other algorithms when m is small.

9. Extensions to tensor computations. Matrix computations and linear algebra are useful when the problem being modeled can be naturally described as having up to two dimensions. The number of dimensions that the object (linear or multilinear) describes is often referred to as the *order* of the object. For problems naturally described as higher-order objects, tensor computations and multilinear algebra are utilized.

As an example of tensor computations, the generalization of matrix-matrix multiplication to tensor computations is the tensor contraction. Not only is matrix-multiplication generalized in the respect that the objects represent a greater number of dimensions, but also the number of dimensions involved in the summation or accumulation of the multiplication is generalized (up to all dimensions of a tensor can be involved in the summation of a tensor contraction), and the notion of a transposition of dimensions is generalized to incorporate the higher number of dimensions represented by each tensor.

A significant benefit of the notation introduced in this paper is that generalizing concepts to tensors and multilinear algebra is relatively straightforward. The notation used for an object's distribution is comprised of two pieces of information: how column indices and row indices of the matrix object are distributed. To describe how a higher-order tensor is distributed, the notation need only extend to describing how the additional dimensions are distributed. Further, while this paper focuses predominately on processing grids that are 2D and 3D, modeling higher-order grids is straightforward. By design, we describe the shape of the grid as an array, where each element is the size of the corresponding dimension of the grid. When targeting a higher-order grid, the array need only be reshaped to match the order of the grid.

The challenges of formalizing how the different collective communications relate different distributions of tensors and of systematically deriving algorithms for tensor operations are beyond the scope of this paper but will be a part of future work. Initial results of how the ideas in this paper are extended to the tensor contraction operation are given in the dissertation proposal of the first author [30].

10. Conclusion. We have given a systematic treatment of the parallel implementation of matrix-vector multiplication and rank-1 update. This motivates the vector and matrix distributions that underlie PLAPACK and, more recently, Elemental. Based on this, we have presented a systematic approach for implementing parallel 2D matrix-matrix multiplication algorithms. With that in place, we then extended the observations to 3D algorithms.

The ideas in this paper primarily focus on aspects of distributed-memory architectures that utilize a bulk-synchronous communication model for network communication. The ideas presented do not preclude the use of many-core and/or GPU architectures within each node of such distributed-memory architectures. For distributed-memory architectures that are most appropriately modeled with bulk-synchronous communications, we hope that the ideas presented will allow others to investigate how to effectively utilize various on-node architectures. We recognize that future distributed-memory architectures may be better suited for more asynchronous communication models; however, it is important to understand when the ideas in this paper can be applied to better tune algorithms for given architectures.

We believe that sufficient details have been given so that the reader can now easily extend our approach to alternative data distributions and/or alternative architectures. Throughout this paper, we have hinted at how the ideas can be extended to the realm of tensor computation on higher-dimensional computing grids. A detailed presentation of how these ideas are extended will be given in future work. Another interesting future direction would be to analyze whether it would be worthwhile to use the proposed 3D parallelization, but with a different 2D SUMMA algorithm within each layer. For example, questions such as the following remain: Would it be worthwhile to use the eSUMMA3D-C approach, but with an eSUMMA2D-A algorithm within each layer?

Appendix A. Scalability of matrix-vector operations. Here, we consider the scalability of various algorithms.

A.1. Weak scalability. A parallel algorithm is said to be weakly scalable when it can maintain efficiency as the number of nodes, p , increases.

More formally, let $T(n)$ and $T(n, p)$ be the cost (in time for execution) of a sequential and parallel algorithm (utilizing p nodes), respectively, when computing with a problem with a size parameterized by n . $n_{\max}(p)$ represents the largest problem that can fit in the combined memories of the p nodes, and let the overhead $T^+(n, p)$ be given by

$$T^+(n, p) = T(n, p) - \frac{T(n)}{p}.$$

Then the efficiency of the parallel algorithm is given by

$$E(n, p) = \frac{T(n)}{pT(n, p)} = \frac{T(n)}{T(n) + pT^+(n, p)} = \frac{1}{1 + \frac{T^+(n, p)}{T(n)/p}}.$$

The efficiency attained by the largest problem that can be executed is then given by

$$E(n_{\max}(p), p) = \frac{1}{1 + \frac{T^+(n_{\max}(p), p)}{T(n_{\max}(p))/p}}.$$

As long as

$$\lim_{p \rightarrow \infty} \frac{T^+(n_{\max}(p), p)}{T(n_{\max}(p))/p} \leq R < \infty,$$

then the effective efficiency is bounded below away from 0, meaning that more nodes can be used effectively. In this case, the algorithm is said to be weakly scalable.

A.2. Weak scalability of parallel matrix-vector multiplication. Let us now analyze the weak scalability of some of the algorithms in section 4.4.

Parallel $y := Ax$. In (4.2), the cost of the parallel algorithm was approximated by

$$T_{y:=Ax}(m, n, d_0, d_1) = 2\frac{mn}{p}\gamma + \log_2(p)\alpha + \frac{d_0 - 1}{d_0} \frac{n}{d_1} \beta + \frac{d_1 - 1}{d_1} \frac{m}{d_0} \beta + \frac{d_1 - 1}{d_1} \frac{m}{d_0} \gamma.$$

Let us simplify the problem by assuming that $m = n$ so that $T_{y:=Ax}(n) = \frac{2n^2\gamma}{p}$ and

$$T_{y:=Ax}(n, d_0, d_1) = 2\frac{n^2}{p}\gamma + \underbrace{\log_2(p)\alpha + \frac{d_0 - 1}{d_0} \frac{n}{d_1} \beta + \frac{d_1 - 1}{d_1} \frac{n}{d_0} \beta + \frac{d_1 - 1}{d_1} \frac{n}{d_0} \gamma}_{T^+_{y:=Ax}(n, d_0, d_1)}.$$

Now, let us assume that each node has memory to store a matrix with M entries. We will ignore memory needed for vectors, workspace, etc. in this analysis. Then

$n_{\max}(p) = \sqrt{p}\sqrt{M}$ and

$$\begin{aligned} \frac{T^+(n_{\max}(p), p)}{T(n_{\max}(p))/p} &= \frac{\log_2(p)\alpha + \frac{d_0-1}{d_0} \frac{n_{\max}}{d_1} \beta + \frac{d_1-1}{d_1} \frac{n_{\max}}{d_0} \beta + \frac{d_1-1}{d_1} \frac{n_{\max}}{d_0} \gamma}{2n_{\max}^2/p\gamma} \\ &= \frac{\log_2(p)\alpha + \frac{d_0-1}{p} n_{\max}\beta + \frac{d_1-1}{p} n_{\max}\beta + \frac{d_1-1}{p} n_{\max}\gamma}{2n_{\max}^2/p\gamma} \\ &= \frac{\log_2(p)\alpha + \frac{d_0-1}{\sqrt{p}} \sqrt{M}\beta + \frac{d_1-1}{\sqrt{p}} \sqrt{M}\beta + \frac{d_1-1}{\sqrt{p}} \sqrt{M}\gamma}{2M\gamma} \\ &= \log_2(p) \frac{1}{2M} \frac{\alpha}{\gamma} + \frac{d_0-1}{\sqrt{p}} \frac{1}{2\sqrt{M}} \frac{\beta}{\gamma} + \frac{d_1-1}{\sqrt{p}} \frac{1}{2\sqrt{M}} \frac{\beta}{\gamma} + \frac{d_1-1}{\sqrt{p}} \frac{1}{2\sqrt{M}}. \end{aligned}$$

We will use this formula to now analyze scalability.

Case 1. $d_0 \times d_1 = p \times 1$. Then

$$\begin{aligned} \frac{T^+(n_{\max}(p), p)}{T(n_{\max}(p))/p} &= \log_2(p) \frac{1}{2M} \frac{\alpha}{\gamma} + \frac{p-1}{\sqrt{p}} \frac{1}{2\sqrt{M}} \frac{\beta}{\gamma} \\ &\approx \log_2(p) \frac{1}{2M} \frac{\alpha}{\gamma} + \sqrt{p} \frac{1}{2\sqrt{M}} \frac{\beta}{\gamma}. \end{aligned}$$

Now, $\log_2(p)$ is generally regarded as a function that grows slowly enough that it can be treated almost like a constant. This is not so for \sqrt{p} . Thus, even if $\log_2(p)$ is treated as a constant, $\lim_{p \rightarrow \infty} \frac{T^+(n_{\max}(p), p)}{T(n_{\max}(p))/p} \rightarrow \infty$, and eventually efficiency cannot be maintained. When $d_0 \times d_1 = p \times 1$, the proposed parallel matrix-vector multiply is not weakly scalable.

Case 2. $d_0 \times d_1 = 1 \times p$. We leave it as an exercise for the reader that the algorithm is not scalable in this case.

The case where $d_0 \times d_1 = 1 \times p$ or $d_0 \times d_1 = p \times 1$ can be viewed as partitioning the matrix by columns or rows, respectively, and assigning these in a round-robin fashion to the 1D array of processors.

Case 3. $d_0 \times d_1 = \sqrt{p} \times \sqrt{p}$. Then

$$\begin{aligned} \frac{T^+(n_{\max}(p), p)}{T(n_{\max}(p))/p} &= \log_2(p) \frac{1}{2M} \frac{\alpha}{\gamma} + \frac{\sqrt{p}-1}{\sqrt{p}} \frac{1}{2\sqrt{M}} \frac{\beta}{\gamma} + \frac{\sqrt{p}-1}{\sqrt{p}} \frac{1}{2\sqrt{M}} \frac{\beta}{\gamma} + \frac{\sqrt{p}-1}{\sqrt{p}} \frac{1}{2\sqrt{M}} \\ &\approx \log_2(p) \frac{1}{2M} \frac{\alpha}{\gamma} + \frac{1}{2\sqrt{M}} \frac{\beta}{\gamma} + \frac{1}{2\sqrt{M}} \frac{\beta}{\gamma} + \frac{1}{2\sqrt{M}}. \end{aligned}$$

Now, if $\log_2(p)$ is treated like a constant, then $R(n_{\max}, \sqrt{p}, \sqrt{p}) \frac{T^+(n_{\max}(p), p)}{T(n_{\max}(p))/p}$ is a constant. Thus, the algorithm is considered weakly scalable for practical purposes.

A.3. Other algorithms. All other algorithms discussed, or derivable with the methods, in this paper can be analyzed similarly.

REFERENCES

- [1] R. C. AGARWAL, S. M. BALLE, F. G. GUSTAVSON, M. JOSHI, AND P. PALKAR, *A three-dimensional approach to parallel matrix multiplication*, IBM J. Res. Develop., 39 (1995), pp. 575–582.
- [2] R. C. AGARWAL, F. GUSTAVSON, AND M. ZUBAIR, *A high-performance matrix multiplication algorithm on a distributed memory parallel computer using overlapped communication*, IBM J. Res. Develop., 38 (1994), pp. 673–681.
- [3] E. ANDERSON, Z. BAI, C. BISCHOF, L. S. BLACKFORD, J. DEMMEL, J. DONGARRA, J. DU CROZ, A. GREENBAUM, S. HAMMARLING, A. MCKENNEY, AND D. SORENSEN, *LAPACK Users' Guide*, 3rd ed., Software Environments Tools 9, SIAM, Philadelphia, 1999, <https://doi.org/10.1137/1.9780898719604>.
- [4] G. BALLARD, *Avoiding Communication in Dense Linear Algebra*, Ph.D. thesis, EECS Department, University of California, Berkeley, 2013.
- [5] G. BALLARD, E. CARSON, J. DEMMEL, M. HOEMMEN, N. KNIGHT, AND O. SCHWARTZ, *Communication lower bounds and optimal algorithms for numerical linear algebra*, Acta Numer., 23 (2014), pp. 1–155.
- [6] R. H. BISSELING, *Parallel iterative solution of sparse linear systems on a transputer network*, in Parallel Computation, The Institute of Mathematics and Its Applications Conference 46, A. E. Fincham and B. Ford, eds., Oxford University Press, Oxford, UK, 1993, pp. 253–271.
- [7] R. H. BISSELING AND W. F. MCCOLL, *Scientific computing on bulk synchronous parallel architectures*, in Technology and Foundations: Information Processing '94, Vol. I, IFIP Transactions A 51, B. Pehrson and I. Simon, eds., Elsevier Science Publishers, Amsterdam, 1994, pp. 509–514.
- [8] J. BRUCK, C.-T. HO, S. KIPNIS, E. UPFAL, AND D. WEATHERSBY, *Efficient algorithms for all-to-all communications in multi-port message-passing systems*, IEEE Trans. Parallel Distrib. Syst., 8 (1997), pp. 1143–1156.
- [9] L. E. CANNON, *A Cellular Computer to Implement the Kalman Filter Algorithm*, Ph.D. thesis, Montana State University, Bozeman, MT, 1969.
- [10] E. CHAN, M. HEIMLICH, A. PURKAYASTHA, AND R. VAN DE GEIJN, *Collective communication: Theory, practice, and experience*, Concurrency and Computation: Practice and Experience, 19 (2007), pp. 1749–1783.
- [11] J. CHOI, J. J. DONGARRA, R. POZO, AND D. W. WALKER, *ScaLAPACK: A scalable linear algebra library for distributed memory concurrent computers*, in Proceedings of the Fourth Symposium on the Frontiers of Massively Parallel Computation, IEEE Press, Piscataway, NJ, 1992, pp. 120–127.
- [12] J. CHOI, D. W. WALKER, AND J. J. DONGARRA, *PUMMA: Parallel universal matrix multiplication algorithms on distributed memory concurrent computers*, Concurrency and Computation: Practice and Experience, 6 (1994), pp. 543–570.
- [13] M. CHRIST, J. DEMMEL, N. KNIGHT, T. SCANLON, AND K. YELICK, *Communication Lower Bounds and Optimal Algorithms for Programs That Reference Arrays—Part 1*, preprint, 2013, <https://arxiv.org/abs/1308.0068>.
- [14] C. EDWARDS, P. GENG, A. PATRA, AND R. VAN DE GEIJN, *Parallel Matrix Distributions: Have We Been Doing It All Wrong?*, Technical Report TR-95-40, Department of Computer Sciences, The University of Texas at Austin, Austin, TX, 1995.
- [15] G. FOX, M. JOHNSON, G. LYZENGA, S. OTTO, J. SALMON, AND D. WALKER, *Solving Problems on Concurrent Processors. Vol. I*, Prentice Hall, Englewood Cliffs, NJ, 1988.
- [16] K. GOTO AND R. A. VAN DE GEIJN, *Anatomy of high-performance matrix multiplication*, ACM Trans. Math. Softw., 34 (2008), 12.
- [17] J. GUNNELS, C. LIN, G. MORROW, AND R. VAN DE GEIJN, *A flexible class of parallel matrix multiplication algorithms*, in Proceedings of the First Merged International Parallel Processing Symposium and Symposium on Parallel and Distributed Processing (IPPS/SPDP 1998), IEEE Press, Piscataway, NJ, 1998, pp. 110–116.
- [18] J. A. GUNNELS, F. G. GUSTAVSON, G. M. HENRY, AND R. A. VAN DE GEIJN, *FLAME: Formal linear algebra methods environment*, ACM Trans. Math. Softw., 27 (2001), pp. 422–455.
- [19] B. HENDRICKSON, R. LELAND, AND S. PLIMPTON, *An efficient parallel algorithm for matrix-vector multiplication*, Int. J. High Speed Comput., 7 (1995), pp. 73–88.
- [20] B. A. HENDRICKSON AND D. E. WOMBLE, *The torus-wrap mapping for dense matrix calculations on massively parallel computers*, SIAM J. Sci. Comput., 15 (1994), pp. 1201–1226, <https://doi.org/10.1137/0915074>.

- [21] S. HUSS-LEDERMAN, E. JACOBSON, AND A. TSAO, *Comparison of scalable parallel matrix multiplication libraries*, in Proceedings of the Scalable Parallel Libraries Conference, IEEE Press, Piscataway, NJ, 1993, pp. 142–149.
- [22] S. HUSS-LEDERMAN, E. JACOBSON, A. TSAO, AND G. ZHANG, *Matrix multiplication on the Intel Touchstone DELTA*, Concurrency and Computation: Practice and Experience, 6 (1994), pp. 571–594.
- [23] D. IRONY, S. TOLEDO, AND A. TISKIN, *Communication lower bounds for distributed-memory matrix multiplication*, J. Parallel Distrib. Comput., 64 (2004), pp. 1017–1026.
- [24] J. G. LEWIS AND R. A. VAN DE GEIJN, *Implementing matrix-vector multiplication and conjugate gradient algorithms on distributed memory multicomputers*, in Proceedings of the Scalable High-Performance Computing Conference, IEEE Press, Piscataway, NJ, 1994, pp. 542–550.
- [25] J. LI, A. SKJELLUM, AND R. D. FALGOUT, *A poly-algorithm for parallel dense matrix multiplication on two-dimensional process grid topologies*, Concurrency and Computation: Practice and Experience, 9 (1997), pp. 345–389.
- [26] L. H. LOOMIS AND H. WHITNEY, *An inequality related to the isoperimetric inequality*, Bull. Amer. Math. Soc., 55 (1949), pp. 961–962.
- [27] W. F. MCCOLL AND A. TISKIN, *Memory-efficient matrix multiplication in the BSP model*, Algorithmica, 24 (1999), pp. 287–297.
- [28] J. POULSON, B. MARKER, R. A. VAN DE GEIJN, J. R. HAMMOND, AND N. A. ROMERO, *Elemental: A new framework for distributed memory dense matrix computations*, ACM Trans. Math. Softw., 39 (2013), 13.
- [29] J. POULSON, R. VAN DE GEIJN, AND J. BENNIGHOF, *Parallel Algorithms for Reducing the Generalized Hermitian-Definite Eigenvalue Problem*, FLAME Working Note 56, Technical Report TR-11-05, Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX, 2011.
- [30] M. D. SCHATZ, *Anatomy of Parallel Computation with Tensors*, Technical Report TR-13-21, Department of Computer Science, The University of Texas at Austin, Austin, TX, 2013.
- [31] E. SOLOMONIK AND J. DEMMEL, *Communication-optimal parallel 2.5d matrix multiplication and lu factorization algorithms*, in Proceedings of the 17th International Conference on Parallel Processing (Euro-Par 2011), Part II, Lecture Notes in Comput. Sci. 6853, Springer, Berlin, 2011, pp. 90–109.
- [32] G. W. STEWART, *Communication and matrix computations on large message passing systems*, Parallel Comput., 16 (1990), pp. 27–40.
- [33] R. VAN DE GEIJN AND J. WATTS, *SUMMA: Scalable universal matrix multiplication algorithm*, Concurrency and Computation: Practice and Experience, 9 (1997), pp. 255–274.
- [34] R. A. VAN DE GEIJN, *Using PLAPACK: Parallel Linear Algebra Package*, The MIT Press, Cambridge, MA, 1997.
- [35] F. G. VAN ZEE, *libflame: The Complete Reference*, Lulu Press, Raleigh, NC, 2009.
- [36] F. G. VAN ZEE, E. CHAN, R. A. VAN DE GEIJN, E. S. QUINTANA-ORTÍ, AND G. QUINTANA-ORTÍ, *The libflame library for dense matrix computations*, IEEE Comput. Sci. Engrg., 11 (2009), pp. 56–63.