A 3D Parallel Algorithm for QR Decomposition

Grey Ballard
Wake Forest University
Winston Salem, NC, USA
ballard@wfu.edu

James Demmel
University of California
Berkeley, CA, USA
demmel@berkeley.edu

Laura Grigori
INRIA Paris-Rocquencourt
Paris, France
laura.grigori@inria.fr

Mathias Jacquelin
Lawrence Berkeley Natl. Lab.
Berkeley, CA, USA
mjacquelin@lbl.gov

Nicholas Knight
New York University
New York, NY, USA
nknight@nyu.edu

ABSTRACT

Interprocessor communication often dominates the runtime of large matrix computations. We present a parallel algorithm for computing QR decompositions whose bandwidth cost (communication volume) can be decreased at the cost of increasing its latency cost (number of messages). By varying a parameter to navigate the bandwidth/latency tradeoff, we can tune this algorithm for machines with different communication costs.

CCS CONCEPTS

• Theory of computation → Massively parallel algorithms;

KEYWORDS

communication-avoiding algorithms; bandwidth-latency tradeoffs; communication-memory tradeoffs

ACM Reference Format:

Grey Ballard, James Demmel, Laura Grigori, Mathias Jacquelin, and Nicholas Knight. 2018. A 3D Parallel Algorithm for QR Decomposition. In *SPAA '18: 30th ACM Symposium on Parallelism in Algorithms and Architectures, July 16–18, 2018, Vienna, Austria*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3210377.3210415>

1 INTRODUCTION

A common task in numerical linear algebra, especially when solving least-squares and eigenvalue problems, is *QR-decomposing* a matrix into a unitary Q -factor times an upper trapezoidal R -factor. We present a QR decomposition algorithm, 3D-CAQR-EG, whose bandwidth and latency costs trade off.

We model the cost of a parallel algorithm in terms of the number of arithmetic operations, the number of words moved between processors, and the number of messages in which these words are moved. These three quantities, measured along critical paths in a schedule, characterize the algorithm's arithmetic cost, bandwidth cost, and latency cost, respectively.

ACM acknowledges that this contribution was authored or co-authored by an employee, contractor, or affiliate of the United States government. As such, the United States government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for government purposes only.

SPAA '18, July 16–18, 2018, Vienna, Austria
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5799-9/18/07...\$15.00
<https://doi.org/10.1145/3210377.3210415>

THEOREM 1. *An $m \times n$ matrix, $m \geq n$, can be QR-decomposed on P processors with these asymptotic costs:*

$$\frac{\# \text{operations}}{mn^2/P} \mid \frac{\# \text{words}}{n^2/(nP/m)^\delta} \mid \frac{\# \text{messages}}{(nP/m)^\delta (\log P)^2} \quad (1)$$

where δ can be chosen from $[1/2, 2/3]$, assuming

$$P / (\log P)^4 = \Omega(m/n), \quad \text{and} \quad (2)$$

$$P \cdot (\log P)^2 = O\left(m^{\frac{\delta}{1+\delta}} \cdot n^{\frac{1-\delta}{1+\delta}}\right).$$

This arithmetic cost is optimal [10], assuming conventional algorithms. For the smallest $\delta = 1/2$, the latency cost is optimal, and for the largest $\delta = 2/3$, the bandwidth cost is optimal [2]. (Optimality is precisely defined in References [2, 10].) However, these bandwidth and latency lower bounds are not attained simultaneously: the bandwidth-latency product is $O(n^2(\log P)^2)$. Similar tradeoffs have been observed in LU-decomposition [22], (Householder) tridiagonalization [20], and triangular substitution [27]. Motivated by lower bounds for Cholesky decomposition and triangular substitution [21], we conjecture that the bandwidth-latency product is $\Omega(n^2)$ for QR-decomposition algorithms in our framework, meaning the tradeoff is inevitable.

Our main contribution is the presentation and analysis of 3D-CAQR-EG, which extends Elmroth-Gustavson's recursive algorithm [11] to the distributed-memory setting and uses communication-efficient subroutines. The inductive cases feature *3D matrix multiplication* (3DMM) [1], which incurs a smaller bandwidth cost than conventional (2D) approaches. The base cases feature a new variant of *communication-avoiding QR* (CAQR) [10]. CAQR incurs a smaller latency cost than conventional (Householder) QR. Our variant further improves the bandwidth cost. A similar approach is known for LU-decomposition [22]. The name '3D-CAQR-EG' reflects this lineage.

For sufficiently tall-and-skinny matrices, it's best to directly call 3D-CAQR-EG's base-case subroutine, 1D-CAQR-EG.

THEOREM 2. *An $m \times n$ matrix can be QR-decomposed on $P \geq m/n \geq 1$ processors with these asymptotic costs:*

$$\frac{\# \text{operations}}{mn^2/P} \mid \frac{\# \text{words}}{n^2} \mid \frac{\# \text{messages}}{(\log P)^2} \quad (3)$$

assuming $P \cdot (\log P)^2 = O(n^2)$.

The rest of this work is organized as follows. We start by summarizing relevant mathematical background on computing QR decompositions (Section 2). We then introduce our parallel machine model, formalizing how we quantify the costs of communication and computation (Section 3). Next we review the communication-efficient subroutines mentioned above, 3DMM (Section 4) and CAQR (Section 5). With this background in place, we present and analyze the new algorithms, 1D-CAQR-EG (Section 6) and 3D-CAQR-EG (Section 7), proving Theorems 1 and 2, in reverse order. We conclude by discussing limitations and extensions and comparing with related work (Section 8).

2 QR DECOMPOSITION

Many applications demand the QR-decomposition of large, dense matrices; let's highlight three. First, QR-decompositions are widely used in computational statistics, e.g., performing regressions and principal-component analyses; Reference [9] features a c. 80M×1K image problem. Second, they are widely used within iterative solvers for orthogonalization; Reference [28] solves a c. 1M×100 problem from GMRES. Third, they are the dominant subroutine in (stable) iterative schemes for matrix functions; Reference [15] includes a 40K×20K problem. In Section 2 we summarize the relevant background concerning computing QR decompositions.

After formalizing the problem in Section 2.1, we present a recursive template algorithm, called REC-QR (Algorithm 1 in Section 2.2), which includes many well-known algorithms as special cases. We then specialize REC-QR to utilize compact matrix representations (Section 2.3) and a simpler recursive splitting strategy (Section 2.4). The result of these specializations, called QR-EG (Algorithm 2), serves as a template for our two new algorithms, 1D-CAQR-EG and 3D-CAQR-EG.

2.1 QR Preliminaries

A QR decomposition of a matrix A is a matrix pair (Q, R) such that $A = QR$, the Q -factor Q is unitary, meaning $Q^H Q = I$, and the R -factor R is upper trapezoidal, meaning all entries below its main diagonal equal zero. To specialize for real-valued A , simply substitute $(\cdot)^T$ for $(\cdot)^H$ and 'orthogonal' for 'unitary'.

We will always assume that A has at least as many rows as columns. (This implies that R has the same dimensions as A and is upper triangular.) When A has more columns than rows, we can obtain a QR decomposition by splitting $A = [A_1 \ A_2]$ with square A_1 , decomposing $A_1 = QR_1$, and computing $R = [R_1 \ Q^H A_2]$.

2.2 Recursive QR Decomposition

We consider QR decomposition algorithms based on REC-QR (Algorithm 1), which split A vertically (Line 4), QR-decompose the left panel (Line 5), update the right panel (Line 6), QR-decompose the lower part of the (updated) right panel (Line 7), and then assemble a QR decomposition from the smaller ones (Lines 8 and 9).

Algorithm 1 $(Q, R) = \text{REC-QR}(A)$

```

1: if BASE-CONDITION then
2:    $(Q, R) = \text{BASE-QR}(A)$ .
3: else
4:   SPLIT  $A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}$  with  $A_{11}$  square.
5:    $(Q_L, R_L) = \text{REC-QR}\left(\begin{bmatrix} A_{11} \\ A_{21} \end{bmatrix}\right)$ .
6:    $\begin{bmatrix} B_{12} \\ B_{22} \end{bmatrix} = Q_L^H \cdot \begin{bmatrix} A_{12} \\ A_{22} \end{bmatrix}$ .
7:    $(Q_R, R_R) = \text{REC-QR}(B_{22})$ .
8:    $Q = Q_L \cdot \begin{bmatrix} I & 0 \\ 0 & Q_R \end{bmatrix}$ .
9:    $R = \begin{bmatrix} R_L & B_{12} \\ 0 & R_R \end{bmatrix}$ .
10: end if
```

We call REC-QR a 'template' because it leaves several details unspecified. To instantiate this template and obtain an algorithm, we must pick a base-case condition (BASE-CONDITION, Line 1), a base-case QR-decomposition subroutine (BASE-QR, Line 2), and a splitting strategy (SPLIT, Line 4). Additionally, we must specify how the operations are scheduled and how the data are distributed.

2.3 Compact Representations

In practice, a QR decomposition (Q, R) of an $m \times n$ matrix ($m \geq n$) is typically not represented as a pair of explicit matrices. In Section 2.3 we specialize REC-QR to represent Q and R more compactly.

Since an R -factor is upper triangular, it is identifiable by just its superdiagonal entries. Subsequently, the symbol R may either denote (1) the actual R -factor, an $m \times n$ upper-triangular matrix; (2) the leading n rows of the R -factor, an $n \times n$ upper-triangular matrix; or (3) the upper triangle of the R -factor, a data structure of size $n(n+1)/2$. When presenting algorithms, we will prefer convention (2); to obtain an $n \times n$ upper-triangular R from REC-QR, we agree that BASE-QR (Line 2) returns such an R , and we rewrite the R -factor assembly (Line 9) as

$$R = \begin{bmatrix} R_L & B_{12} \\ 0 & R_R \end{bmatrix}.$$

Any unitary matrix Q can be written as $Q = I - VTV^H$: the matrix pair (V, T) is called a *basis-kernel representation* [23] of Q . If Q is the Q -factor of a QR decomposition of an $m \times n$ matrix ($m \geq n$), then there exists such a representation where the *basis* V is $m \times n$ and the *kernel* T is $n \times n$.

Modifying REC-QR for basis-kernel representations, Line 6 becomes

$$\begin{bmatrix} B_{12} \\ B_{22} \end{bmatrix} = \begin{bmatrix} A_{12} \\ A_{22} \end{bmatrix} - V_L T_L^H V_L^H \begin{bmatrix} A_{12} \\ A_{22} \end{bmatrix}, \quad (4)$$

and Line 8 becomes

$$\mathbf{V} = \begin{bmatrix} \mathbf{V}_L & \begin{bmatrix} \mathbf{0} \\ \mathbf{V}_R \end{bmatrix} \end{bmatrix}$$

$$\mathbf{T} = \begin{bmatrix} \mathbf{T}_L & -\mathbf{T}_L \mathbf{V}_L^H \begin{bmatrix} \mathbf{0} \\ \mathbf{V}_R \end{bmatrix} & \mathbf{T}_R \\ \mathbf{0} & \mathbf{T}_R \end{bmatrix}, \quad (5)$$

where $(\mathbf{V}_L, \mathbf{T}_L)$ and $(\mathbf{V}_R, \mathbf{T}_R)$ represent $\mathbf{Q}_L$ and $\mathbf{Q}_R$.

To simplify the presentation without affecting our asymptotic conclusions, we will not exploit the block-lower-trapezoidal and block-upper-triangular structures of the bases and kernels. With this understanding, it minimizes arithmetic to evaluate the quadruple matrix product in Equation (4) from right to left, and the quadruple product in Equation (5) from inside-out (two possibilities).

When m is close to n , a general basis-kernel representation may require more storage than the explicit $(m \times m)$ $\mathbf{Q}$ -factor. The QR decomposition algorithms in (Sca)LAPACK use a variant [17] of *compact WY representation* [19], which we call *Householder representation* in this work. In Householder representation, $\mathbf{V}$ is unit lower trapezoidal and $\mathbf{T}$ is upper triangular. These properties enable an in-place implementation, where $\mathbf{V}$'s strict lower trapezoid and $\mathbf{R}$'s upper triangle overwrite $\mathbf{A}$ and where $\mathbf{T}$ need not be stored, since in this case

$$\mathbf{T} = \left(\text{triu}(\mathbf{V}^H \mathbf{V}, -1) + \text{diag}(\text{diag}(\mathbf{V}^H \mathbf{V}))/2 \right)^{-1},$$

using the MATLAB operations 'triu' and 'diag'.

Our algorithms will construct, store, and apply $\mathbf{Q}$ -factors in Householder representation. This choice is motivated by our practical goal of integration into the ScaLAPACK library [6]; from a theoretical standpoint, any basis-kernel representation (with $m \times n$ basis) would yield the same asymptotic costs.

2.4 Elmroth-Gustavson's Approach

The recursive framework of REC-QR is quite general. We will obtain our desired algorithmic costs by following an approach of Elmroth-Gustavson [11] (see LAPACK's `_geqrt3`), in which we split $\mathbf{A}$ vertically (roughly) in half, until the number of columns drops below a given threshold $b \geq 1$. The recursive calls define a binary tree whose $\lceil \log_2(n/b) \rceil$ levels are complete except possibly the last. (We always suppose $b \leq n$; when $b \geq n$, the tree has just one node.) We call this specialized template QR-EG (Algorithm 2); QR-EG utilizes the compact representations as explained in Section 2.3.

We mention that [11] actually proposes a hybrid of the stated approach and an iterative approach, switching between the two for a constant-factor improvement in the arithmetic cost. (It still fits in the REC-QR framework.) While our algorithms can also benefit from this optimization, we omit further discussion in the present work since it does not affect our asymptotic conclusions.

3 COMPUTATION MODEL

We model a parallel machine as a set of P interconnected processors, each with unbounded local memory. Processors

Algorithm 2 $(\mathbf{V}, \mathbf{T}, \mathbf{R}) = \text{QR-EG}(\mathbf{A}, b)$

```

1: if  $n \leq b$  then
2:    $(\mathbf{V}, \mathbf{T}, \mathbf{R}) = \text{BASE-QR}(\mathbf{A})$ .
3: else
4:   SPLIT  $\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}$  so  $\mathbf{A}_{11}$  is  $\lfloor n/2 \rfloor \times \lfloor n/2 \rfloor$ .
5:    $(\mathbf{V}_L, \mathbf{T}_L, \mathbf{R}_L) = \text{QR-EG} \left( \begin{bmatrix} \mathbf{A}_{11} \\ \mathbf{A}_{21} \end{bmatrix}, b \right)$ .
6:    $\mathbf{M}_1 = \mathbf{V}_L^H \cdot \begin{bmatrix} \mathbf{A}_{12} \\ \mathbf{A}_{22} \end{bmatrix}$ .
7:    $\mathbf{M}_2 = \mathbf{T}_L^H \cdot \mathbf{M}_1$ .
8:    $\begin{bmatrix} \mathbf{B}_{12} \\ \mathbf{B}_{22} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{12} \\ \mathbf{A}_{22} \end{bmatrix} - \mathbf{V}_L \cdot \mathbf{M}_2$ .
9:    $(\mathbf{V}_R, \mathbf{T}_R, \mathbf{R}_R) = \text{QR-EG}(\mathbf{B}_{22}, b)$ .
10:   $\mathbf{V} = \begin{bmatrix} \mathbf{V}_L & \begin{bmatrix} \mathbf{0} \\ \mathbf{V}_R \end{bmatrix} \end{bmatrix}$ .
11:   $\mathbf{M}_3 = \mathbf{V}_L^H \cdot \begin{bmatrix} \mathbf{0} \\ \mathbf{V}_R \end{bmatrix}$ .
12:   $\mathbf{M}_4 = \mathbf{M}_3 \cdot \mathbf{T}_R$ .
13:   $\mathbf{T} = \begin{bmatrix} \mathbf{T}_L & -\mathbf{T}_L \cdot \mathbf{M}_4 \\ \mathbf{0} & \mathbf{T}_R \end{bmatrix}$ .
14:   $\mathbf{R} = \begin{bmatrix} \mathbf{R}_L & \mathbf{B}_{12} \\ \mathbf{0} & \mathbf{R}_R \end{bmatrix}$ .
15: end if
```

operate on local data and communicate with other processors by sending and receiving messages. A processor can perform at most one task (operation/send/receive) at a time. A (data) word means a complex number; operations are the usual field actions, plus complex conjugation and real square roots. Messages are point-to-point and asynchronous. Each operation takes time γ , while sending or receiving a message of w words takes time $\alpha + w\beta$, α being the latency and β the inverse of the bandwidth.

We model an execution as a DAG whose vertices are tasks and whose edges define P paths, one for each processor's task sequence, plus an inter-path edge for each send/receive pair. Weighting vertices by their tasks' durations, we define runtime as the maximum weight of any path. Therefore, if every path includes at most F operations and at most S messages, containing at most W words in total, the runtime is bounded,

$$\text{runtime} \leq \gamma \cdot F + \beta \cdot W + \alpha \cdot S.$$

When multiple processors send/receive messages simultaneously, it can be more efficient to split the messages into more messages of smaller sizes, to coalesce them into fewer, larger messages, or to route them through intermediate processors. These lower-level implementation details are often irrelevant to our asymptotic analyses, motivating us to express our algorithms' communication patterns abstractly, in terms of collectives.

In the rest of Section 3 we define the eight different collectives appearing in this work, collecting in Table 1 their costs.

	# operations	# words	# messages
SCATTER	0	$(P - 1)B$	$\log P$
GATHER	0	$(P - 1)B$	$\log P$
BROADCAST	0	$\min(B \log P, B + P)$	$\log P$
REDUCE	$\min(B \log P, B + P)$	$\min(B \log P, B + P)$	$\log P$
ALL-GATHER	0	$(P - 1)B$	$\log P$
ALL-REDUCE	$\min(B \log P, B + P)$	$\min(B \log P, B + P)$	$\log P$
ALL-TO-ALL	0	$\min(BP \log P, (B_* + P^2) \log P)$	$\log P$
REDUCE-SCATTER	$(P - 1)B$	$(P - 1)B$	$\log P$

Table 1: Asymptotic costs of the collectives defined in Section 3. P is the number of processors involved, $B = \max_{p,q} B_{pq}$ is the largest block-size, and $B_* = \max(\max_q \sum_p B_{pq}, \max_p \sum_q B_{pq})$ is the maximum number of words any processor holds before/after. Any P can be replaced by $|\{p, q : B_{pq} > 0\}|$ for a possibly smaller (valid) bound.

A general scenario, called an ALL-TO-ALL, is when every processor p initially owns a block of data, containing B_{pq} words, destined for every processor q , including itself. All other collectives we use can be interpreted as special cases of an ALL-TO-ALL. Four of these distinguish a ‘root’ processor r : SCATTER, where only r ’s outgoing blocks are nonempty; GATHER, where only r ’s incoming blocks are nonempty; BROADCAST, a SCATTER where r ’s outgoing blocks are identical; and REDUCE, a GATHER where r ’s incoming blocks have same size and are added entrywise. Four others, including ALL-TO-ALL, can be constructed from the first four: ALL-GATHER, P GATHERS with same outgoing blocks but different roots; ALL-REDUCE, P REDUCES with same outgoing blocks but different roots; ALL-TO-ALL, P GATHERS with different outgoing blocks and different roots; and REDUCE-SCATTER, P REDUCES with different outgoing blocks and different roots. (Since the three ‘REDUCE’ collectives perform arithmetic, they are technically not ALL-TO-ALLS.)

LEMMA 1. *There exist algorithms for the eight collectives satisfying the upper bounds in Table 1.*

PROOF. For all but ALL-TO-ALL we use a *binomial-tree* and possibly a *bidirectional-exchange algorithm*: see, e.g., [8, 24]. In particular, for (ALL-)GATHER and (REDUCE-)SCATTER we use binary tree algorithms, and for BROADCAST and (ALL-)REDUCE we use whichever of the two minimizes all three costs, asymptotically. For ALL-TO-ALL we use the (*radix-2*) *index algorithm* [7], possibly performed twice using the load-balancing approach of [13]. (For a more detailed proof, see [3, Appendix A].) $\square$

Note that when the number of processors is not too large w.r.t. the block size, the bidirectional-exchange algorithm for BROADCAST, built from SCATTER+ALL-GATHER, is asymptotically cheaper than the corresponding binary tree algorithm in terms of bandwidth. Similarly, the bidirectional-exchange algorithm for (ALL-)REDUCE, built from REDUCE-SCATTER+(ALL-)GATHER, improves both arithmetic and bandwidth.

4 3D MATRIX MULTIPLICATION

The key to reducing 3D-CAQR-EG’s bandwidth cost below that of previous approaches is 3D *matrix multiplication* (3DMM) [1].

Here, 3D refers to the parallelization of the operations and distribution of data over a three-dimensional (logical) processor grid. 1D-CAQR-EG will also exploit two special cases, 1DMM, performed on a one-dimensional grid, and MM, performed locally on one processor.

For concreteness, consider multiplying an $I \times K$ matrix $\mathbf{A}$ with a $K \times J$ matrix $\mathbf{B}$ to obtain an $I \times J$ matrix $\mathbf{C}$, via the usual (entrywise) formula:

$$\text{for all } (i, j) \in [I] \times [J], \quad C_{ij} = \sum_{k \in [K]} A_{ik} B_{kj}. \quad (6)$$

We identify each of the IJK (scalar) multiplications with a point (i, j, k) in three-dimensional Euclidean space, so the set of multiplications defines a discrete $I \times J \times K$ brick. We point the reader to [5, 18] for further discussions of this geometric terminology.

LEMMA 2. *Suppose input matrices $\mathbf{A}$ and $\mathbf{B}$ are initially owned by processor p , and output matrix $\mathbf{C}$ is to be finally owned also by processor p . $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$ can be computed with runtime*

$$\gamma \cdot O(IJK). \quad (7)$$

PROOF. Directly evaluating the sums-of-products in Equation (6) takes IJK multiplications and $IJ(K-1)$ additions; no communication is necessary. $\square$

LEMMA 3. *Suppose I, J, K , and P satisfy*

$$P = O\left(\frac{IJK}{\max(I, J, K)}\right) \quad \text{and} \quad P = O(\max(I, J, K)).$$

If $K = \max(I, J, K)$, suppose that matrices $\mathbf{A}^T$ and $\mathbf{B}$ are initially distributed in matching row-wise layouts where each processor owns $O(K/P)$ rows, and that matrix $\mathbf{C}$ is to be finally owned by a single processor r . Alternatively, if $I = \max(I, J, K)$, suppose that matrices $\mathbf{A}$ and $\mathbf{C}$ are initially/finally distributed in matching row-wise layouts where each processor owns $O(I/P)$ rows, and that matrix $\mathbf{B}$ is initially owned by a single processor r . $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$ can be computed with runtime

$$\gamma \cdot O\left(\frac{IJK}{P}\right) + \beta \cdot O\left(\frac{IJK}{\max(I, J, K)}\right) + \alpha \cdot O(\log P), \quad (8)$$

PROOF. In the first case, each processor performs a local MM and then all processors REDUCE to processor r . In the second case, processor r BROADCASTS $\mathbf{B}$ to all processors and then each processor performs a local MM. The hypotheses guarantee that P is not too large for these collectives can leverage the bidirectional-exchange algorithms. (For a more detailed proof, see [3, Appendix B].) $\square$

The bound of Equation (8) also holds in a third case, when $J = \max(I, J, K)$ and the distributions are symmetric to the second case, but we will not need this result.

LEMMA 4. Suppose I, J, K , and P satisfy

$$c \cdot \frac{IJK}{\min(I, J, K)^3} \leq P \leq IJK.$$

There exists a data distribution of $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ such that each processor initially owns at most $O((I + J)K/P)$ entries of $\mathbf{A}$ and $\mathbf{B}$, and finally at most $O(IJ/P)$ entries of $\mathbf{C}$, where $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$ can be computed with runtime

$$\gamma \cdot O\left(\frac{IJK}{P}\right) + \beta \cdot O\left(\left(\frac{IJK}{P}\right)^{2/3}\right) + \alpha \cdot O(\log P). \quad (9)$$

PROOF. Pick $Q = \lfloor I/\rho \rfloor$, $R = \lfloor J/\rho \rfloor$, and $S = \lfloor K/\rho \rfloor$, where $\rho = (IJK/P)^{1/3}$. Under the hypotheses, Q, R, S are positive integers with $QRS \leq P$, thus define a valid $Q \times R \times S$ processor grid. Moreover, $QRS = \Omega(P)$. Pick partitions $\{\mathcal{I}_q\}_q$, $\{\mathcal{J}_r\}_r$, and $\{\mathcal{K}_s\}_s$ of $[I]$, $[J]$, and $[K]$ which are *balanced*, meaning their parts differ in size by at most one. Pick a partition $\{\mathcal{A}_{q,r,s}\}_{q,r,s}$ of $[I] \times [K]$ to be a union of balanced R -way partitions of the sets $\mathcal{I}_q \times \mathcal{K}_s$ ($(q, s) \in [Q] \times [S]$), and similarly for partitions $\{\mathcal{B}_{q,r,s}\}_{q,r,s}$ and $\{\mathcal{C}_{q,r,s}\}_{q,r,s}$ of $[K] \times [J]$ and $[I] \times [J]$. Distribute entries $\mathcal{A}_{q,r,s}$ of $\mathbf{A}$ to each grid processor (q, r, s) , and similarly for $\mathbf{B}$ and $\mathbf{C}$: this distribution satisfies the balance constraint in the theorem statement. The algorithm proceeds with ALL-GATHERS of blocks of $\mathbf{A}$ and $\mathbf{B}$ along processor grid fibers in the Q - and R -directions. then local MMs, then finally REDUCE-SCATTERS of blocks of $\mathbf{C}$ along processor grid fibers in the S -direction. (For a more detailed proof, see [3, Appendix B].) $\square$

We denote by MM, 1DMM, or 3DMM an algorithm that satisfies Lemma 2, Lemma 3, or Lemma 4, resp.

5 COMMUNICATION-AVOIDING QR

Our new algorithms 1D-CAQR-EG and 3D-CAQR-EG are closely related to *communication-avoiding QR (CAQR)* and *tall-skinny QR (TSQR)* [10]: we explore this relationship in Section 8. For now we remark that 3D-CAQR-EG specializes QR-EG to use 1D-CAQR-EG as a base-case, and that 1D-CAQR-EG specializes QR-EG to use TSQR (the variant in [4]) as a base-case.

On input, the $m \times n$ matrix $\mathbf{A}$ is partitioned across the P processors so that each processor p owns $m_p \geq n$ rows, not necessarily contiguous. Thus we require $\mathbf{A}$ be sufficiently tall and skinny: $m/n \geq P$. A single processor r , which owns $\mathbf{A}$'s n leading rows, is designated as the *root processor*.

On output, the Q -factor is stored in Householder representation $(\mathbf{V}, \mathbf{T})$, where $\mathbf{V}$ has the same distribution as $\mathbf{A}$. Both $\mathbf{T}$ and the R -factor are returned only on the root processor.

LEMMA 5. TSQR's runtime is

$$\gamma \cdot O\left(\max_p m_p n^2 + n^3 \log P\right) + \beta \cdot O\left(n^2 \log P\right) + \alpha \cdot O(\log P)$$

PROOF. It is crucial to use the TSQR variant in [4]. (See [3, Appendix C] for detailed proof.) $\square$

Recall from Section 3 that when the block-size is sufficiently large, REDUCE and BROADCAST can be performed more efficiently, by REDUCE-SCATTER+GATHER and SCATTER+ALL-GATHER, resp. Unfortunately, TSQR's REDUCE- and BROADCAST-like collectives preclude these optimizations. Next in Section 6, we will show how similar savings are achievable.

6 1D-CAQR-EG

We now present a new algorithm, 1D-CAQR-EG, an instantiation of the template QR-EG (Algorithm 2). 1D-CAQR-EG effectively reduces TSQR's bandwidth cost by a logarithmic factor, at the expense of increasing its latency cost by a comparable factor.

The input/output data distributions are the same as for TSQR, so we continue notation from Section 5. We specify 1D-CAQR-EG by stepping line-by-line through QR-EG — base case in Section 6.1 and inductive case in Section 6.2 — then prove Theorem 2 in Section 6.3.

6.1 Base Case

1D-CAQR-EG's base-case QR decomposition subroutine (Line 2) is TSQR (Section 5), using the same root processor. Note that $\mathbf{A}$'s distribution satisfies TSQR's requirements, and $\mathbf{V}$, $\mathbf{T}$, and $\mathbf{R}$ are returned distributed as required by 1D-CAQR-EG.

6.2 Inductive Case

Let us walk through the inductive case line-by-line. All algorithmic costs are incurred in the two recursive calls (Lines 5 and 9) and the six matrix multiplications (Lines 6 to 8 and 11 to 13).

(Line 4): the splitting involves no computation nor communication.

(Line 5): the left recursive call is valid since $\begin{bmatrix} \mathbf{A}_{11} \\ \mathbf{A}_{21} \end{bmatrix}$ still satisfies the data distribution requirements (only n decreases).

(Line 6): this is a 3DMM with matrix dimensions $I = \lfloor n/2 \rfloor$, $J = \lfloor n/2 \rfloor$, and $K = m$. We choose a (1D) processor grid with $Q = R = 1$ and $S = P$, thus $T = 0$ and the partitions $\{\mathcal{I}_q\}_q = \{\{I\}\}$ and $\{\mathcal{J}_r\}_r = \{\{J\}\}$ are trivial. We pick the partition $\{\mathcal{K}_s\}_s$ to match the distribution of $\mathbf{A}$'s rows, and pick $\{\mathcal{A}_{q,r,s}\}_{q,r,s} = \{\{I\} \times \mathcal{K}_s\}_s$ and $\{\mathcal{B}_{q,r,s}\}_{q,r,s} = \{\mathcal{K}_s \times [J]\}_s$. Additionally, we set $\mathcal{C}_{q,r,s} = \emptyset$ for all (q, r, s) but the root processor.

(Line 7): this is an MM on the root processor with matrix dimensions $I = K = \lfloor n/2 \rfloor$ and $J = \lceil n/2 \rceil$.

(Line 8): this is a 3DMM with matrix dimensions $I = m$, $J = \lceil n/2 \rceil$, and $K = \lfloor n/2 \rfloor$, followed by a matrix subtraction. We choose a (1D) processor grid with $Q = P$ and $R = S = 1$, thus $T = 0$ and the partitions $\{\mathcal{I}_r\}_r = \{\lfloor J \rfloor\}$ and $\{\mathcal{K}_s\}_s = \{\lfloor K \rfloor\}$ are trivial. We pick the partition $\{\mathcal{I}_r\}_r$ to match the distribution of $\mathbf{A}$'s rows, and pick $\{\mathcal{A}_{q,r,s}\}_{q,r,s} = \{\mathcal{I}_r \times [K]\}_r$ and $\{\mathcal{C}_{q,r,s}\}_{q,r,s} = \{\mathcal{I}_r \times [J]\}_r$. Additionally, we set $\mathcal{B}_{q,r,s} = \emptyset$ for all (q, r, s) but the root processor. The row-wise distribution $\{\mathcal{C}_{q,r,s}\}_{q,r,s}$ enables the subsequent matrix subtraction to be performed without further communication.

(Line 9): the second recursive call is valid since $\mathbf{B}_{22}$ still satisfies the data distribution requirements: the number of rows owned by the root processor decreases by the same amount that n does, while all other processors keep the same number.

(Line 10): each processor assembles its local rows of $\mathbf{V}$: no computation nor communication is required.

(Line 11): this is a 3DMM with matrix dimensions $I = \lfloor n/2 \rfloor$, $J = \lceil n/2 \rceil$, and $K = m - \lfloor n/2 \rfloor$; we choose a processor grid and partitions as in Line 6.

(Line 12): this is an MM on the root processor with the same dimensions as Line 7.

(Line 13): this is an MM on the root processor with the same dimensions as Lines 7 and 12.

(Line 14): each processor assembles its local rows of $\mathbf{R}$: no computation nor communication is required.

Verify that $\mathbf{V}$, $\mathbf{T}$, and $\mathbf{R}$ are distributed as desired.

6.3 Concluding the Analysis

1D-CAQR-EG is valid for any $P, m, n, b \geq 1$ such that $P \leq m/n$, and there is no loss of generality to suppose $b \leq n$. When $b = n$, 1D-CAQR-EG reduces to TSQR. As we will see, picking $b < n$ allows us to reduce 1D-CAQR-EG's arithmetic and bandwidth costs — while increasing its latency cost — to appear as if we had used bidirectional exchange REDUCE and BROADCAST algorithms (Section 3) within TSQR, despite the fact that these algorithms are inapplicable, as we lamented at the end of Section 5. We will navigate the tradeoff with a nonnegative parameter ϵ , taking

$$b = \Theta(n/(\log P)^\epsilon). \quad (10)$$

We will show that taking $\epsilon = 1$ yields Theorem 2.

LEMMA 6. If $P = O(b^2)$, 1D-CAQR-EG has runtime

$$\gamma \cdot \left(\frac{mn^2}{P} + nb^2 \log P \right) + \beta \cdot O\left(n^2 + nb \log P\right) + \alpha \cdot O\left(\frac{n}{b} \log P\right). \quad (11)$$

PROOF. Let us derive an upper bound $T(m, n)$ on the runtime of an 1D-CAQR-EG invocation. (The unchanging parameters P, b are implicit.) We will now assume a balanced data

distribution, meaning the numbers of rows any two processors owns differ by at most one.

When $P = 1$, the runtime of 1D-CAQR-EG for any b is just $\gamma \cdot O(mn^2)$, which satisfies the conclusion, so we may assume $P > 1$ hereafter.

In the base case ($n \leq b$), the algorithmic cost is the 3D-CAQR-EG call, so by Lemma 5 we conclude that

$$T(m, n) = \gamma \cdot O\left(\frac{mn^2}{P} + n^3 \log P\right) + \beta \cdot O\left(n^2 \log P\right) + \alpha \cdot O(\log P).$$

In the inductive case, the algorithmic cost is due to the two recursive calls, the three 1DMM calls (Lines 6, 8 and 11, performed on 1D processor grids, and the three (local) MM calls (Lines 7, 12 and 13), performed by the root.

The local MMs have runtime $\gamma \cdot O(n^3)$.

To apply Lemma 3 to the 3DMM calls, let us now suppose that $P = O(n^2)$ and $P = O(m)$. Actually, only the first assumption is new: we already know that $P \leq m/n$ for the initial m, n , and when $P > 1$ we see that in any recursive call the current m is within a factor of two of the initial m , hence $P = O(m)$ in any recursive call. Confirming that the data distributions chosen in Section 6.2 match those in the proof of Lemma 3, the 1DMMs' runtime is

$$\gamma \cdot O\left(\frac{mn^2}{P}\right) + \beta \cdot O(n^2) + \alpha \cdot O(\log P).$$

Overall, we have found that

$$T(m, n) = T(m, \lfloor n/2 \rfloor) + T(m - \lfloor n/2 \rfloor, \lceil n/2 \rceil) + \gamma \cdot O\left(\frac{mn^2}{P}\right) + \beta \cdot O(n^2) + \alpha \cdot O(\log P).$$

By induction we may take $T(m, n)$ to be nondecreasing in both m and n . The former property justifies replacing $m - \lfloor n/2 \rfloor$ by m in the second recursive call. Hence, we will take m to be its initial value in the analysis of every recursive call.

Supposing $n = b2^L$ for a nonnegative integer L , $T(m, n)$ is bounded by Equation (11).

Now observe that $n = b2^{L+1}$ reproduces the asymptotic bound Equation (11), and since $T(m, n)$ is nondecreasing in n , this bound also holds when $b2^L < n < b2^{L+1}$.

Requiring $P = O(b^2)$ suffices to ensure that $P = O(n^2)$ at every inductive case. $\square$

PROOF OF THEOREM 2. Let us substitute Equation (10) into Equation (11):

$$\gamma \cdot \left(\frac{mn^2}{P} \left(1 + \frac{nP}{m} (\log P)^{1-2\epsilon} \right) + \beta \cdot O\left(n^2 \left(1 + (\log P)^{1-\epsilon} \right) \right) + \alpha \cdot O\left((\log P)^{1+\epsilon}\right) \right),$$

thence the hypothesis is $P(\log P)^{2\epsilon} = O(n^2)$. We conclude by taking $\epsilon = 1$. $\square$

This argument extends to any $\epsilon \geq 0$, assuming $P(\log P)^{2\epsilon} = O(n^2)$, but the asymptotic tradeoff vanishes when $\epsilon > 1$. For

$1/2 \leq \epsilon \leq 1$ the tradeoff is only between bandwidth and latency. A sensible interpretation of the case $\epsilon < 0$ is $b = n$, meaning TSQR is invoked immediately. In this case, the costs are given directly by Lemma 5.

7 3D-CAQR-EG

We now present our second new algorithm, 3D-CAQR-EG, another instantiation of the template QR-EG (Algorithm 2).

On input, the $m \times n$ matrix A ($m \geq n$) is partitioned across the P processors row-cyclically: thus, each processor owns at most $\lceil m/P \rceil$ rows.

On output, the Q -factor is stored in Householder representation (V, T) , where V has the same distribution as A . Both T and the R -factor have the same distribution, matching the top $n \times n$ submatrix of A .

After walking through 3D-CAQR-EG in a similar fashion as we did for 1D-CAQR-EG in Section 6, we collect the results to prove Theorem 1.

In the following, $T_{3D-CAQR-EG}$ denotes an upper bound on the runtime of 3D-CAQR-EG, a function of P, m, n, b .

7.1 Base Case

Recall that b denotes the recursive threshold for 3D-CAQR-EG. 3D-CAQR-EG's base-case QR decomposition subroutine (Line 2) is 1D-CAQR-EG, with a fixed recursive threshold b^* .

To satisfy 1D-CAQR-EG's data distribution requirements, we convert A from row-cyclic to block-row layout, distributed over

$$P^* = \min(P, \lfloor m/n \rfloor)$$

processors, ensuring each owns at least n rows and one owns the top n rows (and perhaps others).

Initially, $P' = \min(m, P)$ processors own rows of A . (Clearly $P' \leq P$ with equality just in case $P \leq m$; note further that $P^* \leq P'$ with equality just in case $P \leq m/n$.) Number these processors from 0 to $P'-1$ according to the cyclic layout of A , so that processor 0 owns the top row of A . Deal these processors among P^* groups, so processor 0 goes into group 0, processor 1 goes into group 1, and so on. Represent each group by its lowest-numbered processor and within each group, GATHER A 's rows to the representative. Since each group contains at most $\lceil P'/P^* \rceil$ processors and each processor initially owns at most $\lceil m/P' \rceil$ rows of A , the largest block-size in any GATHER is at most $\lceil m/P' \rceil n$.

Each of the P^* representatives (including processor 0) now owns at least $\lfloor m/P^* \rfloor \geq n$ rows of A , satisfying the first part of 1D-CAQR-EG's data distribution requirements: it remains to ensure processor 0 owns the top n rows of A . These rows are currently owned by the first $P'' = \min(P^*, n)$ representatives. (Clearly $P'' \leq P^*$ with equality just in case $P \leq n$.) We next perform a GATHER over the representatives of groups 0 through $P''-1$, taking processor 0 to be the root so that afterwards it owns the top n rows of A (and perhaps others). We also perform a SCATTER with the opposite communication pattern so that the overall number of rows per representative

is unchanged. The largest block-size in both the GATHER and the SCATTER is at most $\lceil n/P'' \rceil n$.

We can now invoke 1D-CAQR-EG, with parameters P^*, b^* . After it returns, we redistribute V, T , and R by reversing the preceding GATHERS/SCATTERS, so that V is (resp., T and R are) distributed over all P processors like A was (resp., A 's first n rows were) initially.

7.2 Inductive Case

Let us walk through the inductive case line-by-line, as we did for 1D-CAQR-EG. All algorithmic costs are incurred in the two recursive calls (Lines 5 and 9) and the six matrix multiplications (Lines 6 to 8 and 11 to 13).

(Line 4): the splitting involves no computation nor communication.

(Line 5): the left recursive call is valid since $\begin{bmatrix} A_{11} \\ A_{21} \end{bmatrix}$ still satisfies the data distribution requirements (only n decreases).

(Line 6): this is a 3DMM with matrix dimensions $I = \lfloor n/2 \rfloor$, $J = \lceil n/2 \rceil$, and $K = m$. We do not yet specify the processor grid, but we do suppose that 3D-CAQR-EG uses a balanced parallelization and data distribution as in the proof of Lemma 4: this is possible for any processor grid.

To match this data distribution, we perform an ALL-TO-ALL before and after the 3DMM invocation, each time using the two-phase approach [7]. The first ALL-TO-ALL redistributes the input matrices from column- and row-cyclic to 3DMM layout (the left factor is row-cyclic, transposed); the maximum number of input matrix entries any processor owns before or after this collective is at most

$$\max \left(I \left\lceil \frac{K}{P} \right\rceil + \left\lceil \frac{K}{P} \right\rceil J, \left\lceil \frac{\left\lceil \frac{I}{Q} \right\rceil \left\lceil \frac{K}{S} \right\rceil}{R} \right\rceil + \left\lceil \frac{\left\lceil \frac{J}{R} \right\rceil \left\lceil \frac{K}{S} \right\rceil}{Q} \right\rceil \right),$$

where the processor grid is $Q \times R \times S$. The second ALL-TO-ALL converts the output matrix from 3DMM layout to row-cyclic layout; the maximum number of output matrix entries any processor owns before or after this collective is at most

$$\max \left(\left\lceil \frac{\left\lceil \frac{I}{Q} \right\rceil \left\lceil \frac{J}{R} \right\rceil}{S} \right\rceil, \left\lceil \frac{I}{P} \right\rceil J \right).$$

(Line 7): this is a 3DMM with matrix dimensions $I = K = \lfloor n/2 \rfloor$ and $J = \lceil n/2 \rceil$. We pick a 3D processor grid and partitions satisfying the same constraints as for Line 6, and perform similar ALL-TO-ALLs before and after, so we can reuse the preceding analysis, substituting I, J , and K .

(Line 8): this is a 3DMM with matrix dimensions $I = m$, $J = \lceil n/2 \rceil$, and $K = \lfloor n/2 \rfloor$, followed by a matrix subtraction. We proceed similarly to Lines 6 and 7, except that the left factor is initially in row-cyclic layout, so the first summand in the first term of the maximum becomes $\lceil I/P \rceil K$ (vs. $I \lceil K/P \rceil$).

(Line 9): the second recursive call is valid since B_{22} still satisfies the data distribution requirements: in particular, unlike 1D-CAQR-EG there is no requirement that a fixed processor

owns the first n rows or that every processor owns at least n rows.

(Line 10): each processor assembles its local rows of $\mathbf{V}$: no computation nor communication is required.

(Line 11): this is a 3DMM with matrix dimensions $I = \lfloor n/2 \rfloor$, $J = \lceil n/2 \rceil$, and $K = m - \lfloor n/2 \rfloor$; we choose a processor grid, partitions, and ALL-TO-ALLS as in Line 6.

(Line 12): this is a 3DMM with the same dimensions, processor grid, partitions, and ALL-TO-ALLS as Line 7.

(Line 13): this is a 3DMM with the same dimensions, processor grid, partitions, and ALL-TO-ALLS as Lines 7 and 12.

(Line 14): each processor assembles its local rows of $\mathbf{R}$: no computation nor communication is required.

Verify that $\mathbf{V}$, $\mathbf{T}$, and $\mathbf{R}$ are distributed as desired.

7.3 Concluding the Analysis

3D-CAQR-EG is valid for any $P, m, n, b, b^* \geq 1$, and there is no loss of generality to suppose $b^* \leq b \leq n$. Taking $b = n$ simplifies 3D-CAQR-EG to 1D-CAQR-EG with parameters P^* , b^* and additional data redistributions. As in the case of 1D-CAQR-EG, picking $b < n$ allows us to reduce 3D-CAQR-EG's arithmetic and bandwidth costs, while increasing its latency cost.

We will navigate this tradeoff with two nonnegative parameters δ, ϵ , taking

$$b = \Theta\left(n/(nP/m)^\delta\right), \quad b^* = \Theta\left(b/(\log P)^\epsilon\right). \quad (12)$$

We prove Theorem 1 with $\delta \in [1/2, 2/3]$ and $\epsilon = 1$.

LEMMA 7. *If $P = O(b^2)$ and $P^* = O(b^{*2})$, 3D-CAQR-EG has runtime*

$$\begin{aligned} T(m, n) = & \gamma \cdot \left(\frac{mn^2}{P} + nb^{*2} \log P \right) \\ & + \beta \cdot O\left(\frac{mn}{P} + nb + nb^* \log P + \left(\frac{mn^2}{P} \right)^{2/3} \right. \\ & \left. + \left(\left(\frac{mn}{P} + n \right) \log \frac{n}{b} + \frac{nP^2}{b} \right) \log P \right) \\ & + \alpha \cdot O\left(\frac{n}{b^*} \log P \right). \end{aligned} \quad (13)$$

PROOF. Let us derive an upper bound $T(m, n)$ on the runtime of a 3D-CAQR-EG invocation. (The unchanging parameters P, b, b^* are implicit.)

When $P = 1$, the runtime of 3D-CAQR-EG is just $\gamma \cdot O(mn^2)$, which satisfies the conclusion, so we may assume $P > 1$ hereafter.

In a base case ($n \leq b$), the algorithmic costs are due to the 1D-CAQR-EG invocation and the four communication phases.

The first communication phase, involving P^* independent GATHERS, has runtime bounded by

$$\begin{aligned} & \beta \cdot O\left((\lceil P'/P^* \rceil - 1)\lceil m/P' \rceil n\right) + \alpha \cdot O\left(\log \lceil P'/P^* \rceil\right) \\ & = \beta \cdot O\left(mn/P + n^2\right) + \alpha \cdot O(\log P), \end{aligned}$$

and the same bound applies for the last phase (matching SCATTERS). (Recall that $P^* = \min(P, \lfloor m/n \rfloor)$ and $P' = \min(P, m)$.) The fact that no communication happens when $m \geq nP$ (and thus $P^* = P$) is evident in the first bound but not the second.

The second communication phase, a simultaneous GATHER/SCATTER, has runtime bounded by

$$\begin{aligned} & \beta \cdot O\left((P'' - 1)\lceil n/P'' \rceil n\right) + \alpha \cdot O(\log P'') \\ & = \beta \cdot O\left(n^2\right) + \alpha \cdot O(\log P), \end{aligned}$$

and the same bound applies for the third phase (matching SCATTER/GATHER). (Recall that $P'' = \min(P^*, n)$.)

A runtime bound for the 1D-CAQR-EG invocation is given by Theorem 2, supposing now that $P^* = O(b^{*2})$. Actually we use the more refined bound of Equation (11), substituting P^*, b^* for P, b .

Altogether, in a base case,

$$\begin{aligned} T(m, n) = & \gamma \cdot O\left(\frac{mn^2}{P} + nb^{*2} \log P\right) \\ & + \beta \cdot O\left(\frac{mn}{P} + n^2 + nb^* \log P\right) + \alpha \cdot O\left(\frac{n}{b^*} \log P\right) \end{aligned}$$

In the inductive case ($n > b$), the algorithmic cost is due to the two recursive calls, the six 3DMMs, performed on 3D processor grids, and the twelve ALL-TO-ALLS, performed before and after each 3DMM.

To apply Lemma 4 to the 3DMMs, let us now suppose that $P \geq (3c)^{3/4}$ for some $c > 1$ and $b \geq 2P^{1/3}$: since $m \geq n$ at every recursive call and since $n \geq b + 1$ in the inductive case, $c \cdot IJK/\min(I, J, K)^3 \leq P \leq IJK$ for each 3DMM. Thus the 3DMMs' overall runtime is

$$\gamma \cdot O\left(\frac{mn^2}{P}\right) + \beta \cdot O\left(\left(\frac{mn^2}{P}\right)^{2/3}\right) + \alpha \cdot O(\log P).$$

Moreover, the inequalities derived at the start of Lemma 4's proof also yield the following upper bound on the overall runtime of the ALL-TO-ALLS,

$$\beta \cdot O\left(\left(\frac{mn}{P} + n + P^2\right) \log P\right) + \alpha \cdot O(\log P).$$

Overall, we have found that

$$\begin{aligned} T(m, n) = & T(m, \lfloor n/2 \rfloor) + T(m - \lfloor n/2 \rfloor, \lceil n/2 \rceil) \\ & + \gamma \cdot O\left(\frac{mn^2}{P}\right) + \beta \cdot O\left(\left(\frac{mn^2}{P}\right)^{2/3}\right) \\ & + \left(\frac{mn}{P} + n + P^2\right) \log P + \alpha \cdot O(\log P). \end{aligned}$$

By induction we may take $T(m, n)$ to be nondecreasing in both m and n . The former property justifies replacing $m - \lfloor n/2 \rfloor$ by m in the second recursive call. Hence, we will take m to be its initial value in the analysis of every recursive call.

Supposing $n = b2^L$ for a nonnegative integer L , $T(m, n)$ is bounded by Equation (13). Now observe that $n = b2^{L+1}$ reproduces the asymptotic bound of Equation (13), which thus

holds when $b2^L < n < b2^{L+1}$ since $T(m, n)$ is nondecreasing in n .

In conclusion, 3D-CAQR-EG's runtime $T(m, n)$ satisfies Equation (13) if $P \geq 3$, $b \geq 2P^{1/3}$, $P = O(b^2)$, and $P^* = O(b^{*2})$. $\square$

PROOF OF THEOREM 1. Consider picking b, b^* as in Equation (12). The constraints relating P and b are satisfiable if

$$P = O\left(m^{\frac{2\delta}{1+2\delta}} \cdot n^{\frac{2-2\delta}{1+2\delta}}\right), \quad (14)$$

and the constraint relating P^* and b^* is satisfiable if

$$P \cdot (\log P)^{\frac{2\epsilon}{1+2\delta}} = O\left(m^{\frac{2\delta}{1+2\delta}} \cdot n^{\frac{2-2\delta}{1+2\delta}}\right), \quad (15)$$

a stronger condition than Equation (14).

Substituting Equation (10) into Equation (11),

$$\begin{aligned} & \gamma \cdot \left(\frac{mn^2}{P} \left(1 + \left(\frac{nP}{m} \right)^{1-2\delta} (\log P)^{1-2\epsilon} \right) \right) \\ & + \beta \cdot O\left(\frac{n^2}{\left(\frac{nP}{m} \right)^\delta} \left(1 + (\log P)^{1-\epsilon} \right) + W \right) \\ & + \alpha \cdot O\left((nP/m)^\delta (\log P)^{1+\epsilon} \right), \end{aligned}$$

where W denotes the sum of three terms associated with the ALL-TO-ALLS,

$$\frac{mn}{P} \log \frac{nP}{m} \log P, \quad n \log \frac{nP}{m} \log P, \quad P^2 \left(\frac{nP}{m} \right)^\delta \log P,$$

plus a term $n^2/(nP/m)^{2/3}$ associated with the 3DMMs. We obtain the stated arithmetic and latency costs by taking $\delta \geq 1/2$ and $\epsilon = 1$. To suppress the bandwidth term W , it suffices to require that there exists $\delta' \in (0, 1-\delta)$, hence $\delta < 1$, such that

$$\begin{aligned} P / (\log P)^{\frac{\epsilon}{1-\delta-\delta'}} &= \Omega(m/n), \\ P \cdot (\log P)^{\frac{\epsilon}{\delta+\delta'}} &= O\left(m \cdot n^{\frac{1-\delta-\delta'}{\delta+\delta'}}\right), \\ P \cdot (\log P)^{\frac{\epsilon}{2+\delta}} &= O\left(m^{\frac{\delta}{1+\delta}} \cdot n^{\frac{1-\delta}{1+\delta}}\right). \end{aligned} \quad (16)$$

The 3DMMs' bandwidth cost cannot be reduced, but it is lower-order if $\delta \leq 2/3$.

The hypotheses of Theorem 1, $\delta \in [1/2, 2/3]$ (and tacitly $\epsilon = 1$) and Equation (2), imply Equations (14) to (16). $\square$

This argument extends to a larger range of nonnegative δ, ϵ . Assuming fixed ϵ , for $\delta > 2/3$ the 3DMM invocations dominate the bandwidth cost, whose bound remains as if $\delta = 2/3$, no longer a tradeoff, at least asymptotically. In the case $0 \leq \delta < 1/2$, the additive term in the arithmetic cost, due to the small (mostly) triangular matrix operations on TSQR's critical path, possibly dominates. (A sensible interpretation of the case $\delta \leq 0$ is $b = n$, in which case 1D-CAQR-EG is invoked immediately.) Assuming fixed δ , the tradeoffs due to varying $\epsilon \in [0, 1]$ are just as in the proof of Theorem 2, except now the factor in the arithmetic cost is suppressed by increasing δ .

8 DISCUSSION

We have presented two new algorithms, 1D-CAQR-EG and 3D-CAQR-EG (Sections 6 and 7), for computing QR decompositions on distributed-memory parallel machines. Our analysis (e.g., Equations (11) and (13)) demonstrates tradeoffs between arithmetic, bandwidth, and latency costs, governed by the choice of one (1D-CAQR-EG) or two (3D-CAQR-EG) block sizes. We navigated these tradeoffs in Theorems 1 and 2 by asymptotically minimizing arithmetic, as well as bandwidth in Theorem 2.

8.1 Comparison With Similar Algorithms

Here we compare the two new algorithms with four other instances of the REC-QR framework, deriving Tables 2 and 3. Let us review the other algorithms.

An early and well-known instance of REC-QR (Algorithm 1) was proposed by Householder [14]. It features a splitting strategy (Line 4) where A_{11} is $b \times b$ (*right-looking*) or $n-b \times n-b$ (*left-looking*); a base case threshold (Line 1) asserting $n \leq b$; and a base case subroutine (Line 2) generating a product of b Householder reflectors. Householder's proposal was *right-looking* and *unblocked*, meaning $b = 1$ (vs. *blocked*, $b > 1$).

Let 1D-HOUSE and 2D-HOUSE denote the un/blocked right-looking variants, specialized to use compact representations (Section 2.3); their costs are summarized in the first rows of Tables 2 and 3. 2D-HOUSE invokes 1D-HOUSE as its base case. For 1D-HOUSE we use a 1D processor grid and for 2D-HOUSE we use a 2D processor grid. For 1D-HOUSE we distribute matrices similar to 1D-CAQR-EG and for 2D-HOUSE we distribute matrices (2D-) block-cyclically with $b \times b$ blocks: the distribution block size matches the algorithmic block size. We parallelize 1D-HOUSE and the base case of 2D-HOUSE to match the distribution of A , analogous to 1D-CAQR-EG. We parallelize 2D-HOUSE's inductive-case matrix multiplications to match the output matrix distribution (a 2D parallelization). In the case of 1D-HOUSE we assume $P = O(m)$. In the case of 2D-HOUSE, we choose an $r \times c$ processor grid with $c = \Theta((nP/m)^{1/2})$ and $r = \Theta(P/c)$, and we choose $b = \Theta(1)$. Assuming $P = \Omega(m/n)$ and $P \cdot (\log P)^2 = O(m \cdot n)$, these choices are valid and simultaneously minimize all three costs, asymptotically.

CAQR [10] modifies 2D-HOUSE to invoke TSQR (Section 5) in the base case. Our algorithms make crucial use of the TSQR enhancements in [4]; additionally, we use that paper's improved CAQR in the following comparison. We parallelize and distribute data for TSQR as discussed in Section 5, and for CAQR's inductive case as we did for 2D-HOUSE's. TSQR and CAQR's costs are summarized in the second rows of Tables 2 and 3. In the case of TSQR we assume $P \leq m/n$. In the case of CAQR we use the same $r \times c$ grid as for 2D-HOUSE but now pick $b = \Theta(n/(nP/m)^{1/2})$. Assuming $P/(\log P)^2 = \Omega(m/n)$ and $P \cdot (\log P)^2 = O(m \cdot n)$, these choices are valid and simultaneously minimize all three costs, asymptotically.

The costs of the new algorithms, 1D-CAQR-EG and 3D-CAQR-EG appear in the third rows of Tables 2 and 3. To make the comparison between TSQR and 1D-CAQR-EG more clear, for the latter we use $b = \Theta(n/(\log P)^\epsilon)$ in Theorem 2's proof, allowing

<i>algorithm</i>	<i># operations</i>	<i># words</i>	<i># messages</i>
2D-HOUSE	mn^2/P	$n^2/(nP/m)^{1/2}$	$n \log P$
CAQR	mn^2/P	$n^2/(nP/m)^{1/2}$	$(nP/m)^{1/2}(\log P)^2$
3D-CAQR-EG	mn^2/P	$n^2/(nP/m)^\delta$	$(nP/m)^\delta(\log P)^2$

Table 2: Comparison of approaches for square-ish matrices ($m/n = O(P)$). The algorithms and the assumptions that support these bounds are explained in Section 8.1. (In line 3, δ varies from $1/2$ to $2/3$.)

<i>algorithm</i>	<i># operations</i>	<i># words</i>	<i># messages</i>
1D-HOUSE	mn^2/P	$n^2 \log P$	$n \log P$
TSQR	$mn^2/P + n^3 \log P$	$n^2 \log P$	$\log P$
1D-CAQR-EG	$mn^2/P + n^3(\log P)^{1-2\epsilon}$	$n^2(\log P)^{1-\epsilon}$	$(\log P)^{1+\epsilon}$

Table 3: Comparison of approaches for tall/skinny matrices ($m/n = \Omega(P)$). The algorithms and the assumptions that support these bounds are explained in Section 8.1. (In line 3, ϵ varies from 0 to 1.)

the parameter ϵ to vary over $[0, 1]$, justified by the stronger constraint $P(\log P)^{2\epsilon} = O(n^2)$. For 3D-CAQR-EG we follow Theorem 1's proof and hypotheses.

8.2 Elimination By Blocks

Tiskin [25], working in the BSP model [26], proposed an algorithm outside of the REC-QR framework which demonstrates a similar bandwidth/latency tradeoff as 3D-CAQR-EG. Tiskin's algorithm was designed only for square matrices, but it has been extended to rectangular matrices, using the original algorithm as a black box [20]. This extension achieves BSP bandwidth cost $O(n^2/(nP/m)^\delta)$ and BSP synchronization cost $O((nP/m)^\delta(\log P)^2)$. Despite the fact that, in BSP, 3D-CAQR-EG achieves these same communication costs, we still believe 3D-CAQR-EG is a valuable contribution for multiple reasons. Defining a data distribution is a nontrivial and crucial step in developing a distributed-memory implementation using, e.g., MPI. The aforementioned BSP algorithms do not (and need not) explicitly specify their data distributions. Second, our algorithms are based on Householder's algorithm and use Householder representation. Thus, they are readily assembled from robust, tuned subroutines in standard libraries like (P)BLAS and (Sca)LAPACK. Additionally, all interprocessor communication in our algorithms is expressed in terms of standard MPI collectives. Lastly, we feel that Tiskin's recursive scheme, based on a slope-2 wavefront and 'pseudopanel', is much more demanding from an implementation perspective than Elmroth-Gustavson's (QR-EG). To our knowledge no one has implemented Tiskin's algorithm.

8.3 Lower Bounds

The algorithms discussed in Section 8.1 are all subject to an arithmetic lower bound of $\Omega(mn^2/P)$ [10].

In the tall-skinny case, we have bandwidth and latency bounds $\Omega(n^2)$ and $\Omega(\log P)$ [2, 8]. 1D-HOUSE attains the arithmetic lower bound, but misses the bandwidth and latency

lower bounds by $\Theta(\log P)$ and $\Theta(n)$. TSQR attains the arithmetic lower bound assuming $P \log P = \Omega(m/n)$, but misses the bandwidth and latency lower bounds both by $\Theta(\log P)$. 1D-CAQR-EG attains the latency lower bound when $\epsilon = 0$, the arithmetic lower bound when $\epsilon \leq 1/2$, and the bandwidth lower bound when $\epsilon \geq 1$.

In the (close to) square case, we have bandwidth and latency bounds $\Omega(n^2/(nP/m)^{2/3})$ and $\Omega((nP/m)^{1/2})$ [2]. We restrict parameters so that both CAQR and 3D-CAQR-EG attain the arithmetic lower bound. 2D-HOUSE and CAQR exceed the bandwidth lower bound both by a factor of $\Theta((nP/m)^{1/6})$ and they exceed the latency lower bound by factors of $\Theta(n/(nP/m)^{1/2} \log P)$ and $\Theta((nP/m)^{1/6}(\log P)^2)$, resp. 3D-CAQR-EG attains the bandwidth lower bound when $\delta = 2/3$, and exceeds the latency lower bound by just $\Theta((\log P)^2)$ when $\delta = 1/2$.

We did not prove that 3D-CAQR-EG's bandwidth-latency product is optimal — i.e., that the tradeoff is inevitable — although we conjecture this to be the case. Our intuition is based on bandwidth/latency tradeoffs observed in computations whose dependence graphs have similar structure [16, 21].

8.4 Limitations and Extensions

Our main upper bound Theorem 1 is substantially limited by its restrictions on permissible parallelism: see Equation (2). 3D-CAQR-EG's ALL-TO-ALLs are responsible for these constraints: if we supposed the ALL-TO-ALLs had zero cost, Equation (2) could be weakened to Equation (15). We make three remarks about improving this aspect of our work.

First, the bound used is worst-case; our knowledge of (and control over) data distribution could lead to stronger bounds. Second, it may be that the index algorithm is suboptimal for the data distribution, e.g., many $B_{pq} = 0$ and a specialized algorithm would perform less communication, or at least yield sharper cost bounds. Third and more generally, we should optimize for the data distribution before and after each sub-routine. The constraints are the balance assumptions to invoke Lemma 4 and 1D-CAQR-EG. This is a hard problem.

We have also omitted a number of practical optimizations that do not affect our asymptotic analysis. For example, recall from Section 2.3 that T can be reconstructed from V . If the full T is not desired, by replacing the top level of recursion with a right-looking iterative QR-EG variant, we can avoid ever computing superdiagonal blocks of T ; this does, however, restrict the available parallelism [11].

The constants hidden in our asymptotic analysis are practically important, and parameter choices for particular machines warrant further study. For parallel dense linear algebra, the typical behavior of a strong-scaling regime (increasing P for a fixed problem) is for arithmetic cost to dominate for small P , initially scaling nearly perfectly. However, for larger P , local computation is reduced and the bandwidth cost increasingly dominates runtime. This is the scenario when the new tuning parameter δ will be most useful, reducing the bandwidth cost and extending strong scaling.

Elmroth-Gustavson's recursive approach extends from QR-decomposition to other matrix computations like Cholesky decompositions and solving triangular matrix equations [12]. We see no major obstacles to obtaining the corresponding extensions of Theorems 1 and 2.

ACKNOWLEDGMENTS

G. Ballard was supported by the National Science Foundation (NSF) Grant No. ACI-1642385. The work of L. Grigori was supported by the NLAFFET project as part of the European Union's Horizon 2020 research and innovation program under grant 671633. Support for M. Jacquelin was provided in part through the Scientific Discovery through Advanced Computing (SciDAC) program funded by the US Department of Energy (DOE), Office of Science, Advanced Scientific Computing Research under Contract No. DE-AC02-05CH11231.

REFERENCES

- [1] R. Agarwal, S. Balle, F. Gustavson, M. Joshi, and P. Palkar. 1995. A Three-Dimensional Approach to Parallel Matrix Multiplication. *IBM Journal of Research and Development* 39, 5 (1995), 575–582. <https://doi.org/10.1147/rd.395.0575>
- [2] G. Ballard, E. Carson, J. Demmel, M. Hoemmen, N. Knight, and O. Schwartz. 2014. Communication lower bounds and optimal algorithms for numerical linear algebra. *Acta Numerica* 23 (2014), 1–155. http://journals.cambridge.org/article_S0962492914000038
- [3] G. Ballard, J. Demmel, L. Grigori, M. Jacquelin, and N. Knight. 2018. A 3D Parallel Algorithm for QR Decomposition. Technical Report 1805.05278. arXiv. <https://arxiv.org/abs/1805.05278>
- [4] G. Ballard, J. Demmel, L. Grigori, M. Jacquelin, N. Knight, and H. D. Nguyen. 2015. Reconstructing Householder Vectors from Tall-Skinny QR. *J. Parallel and Distrib. Comput.* 85 (2015), 3–31. <https://doi.org/10.1016/j.jpdc.2015.06.003>
- [5] G. Ballard, A. Drusinsky, N. Knight, and O. Schwartz. 2016. Hypergraph partitioning for sparse matrix-matrix multiplication. *ACM Transactions on Parallel Computing* 3, 3 (2016), 18. <http://doi.acm.org/10.1145/3015144>
- [6] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. 1997. *ScaLAPACK Users' Guide*. SIAM, Philadelphia, PA, USA. Also available from <http://www.netlib.org/scalapack/>.
- [7] J. Bruck, C.-T. Ho, S. Kipnis, E. Upfal, and D. Weathersby. 1997. Efficient algorithms for all-to-all communications in multiport message-passing systems. *IEEE Transactions on Parallel and Distributed Systems* 8, 11 (1997), 1143–1156. <https://doi.org/10.1109/71.642949>
- [8] E. Chan, M. Heimlich, A. Purkayastha, and R. van de Geijn. 2007. Collective communication: theory, practice, and experience. *Concurrency and Computation: Practice and Experience* 19, 13 (2007), 1749–1783. <https://doi.org/10.1002/cpe.1206>
- [9] P. Constantine and D. Gleich. 2011. Tall and Skinny QR Factorizations in MapReduce Architectures. In *Proceedings of the Second International Workshop on MapReduce and Its Applications (MapReduce '11)*. ACM, New York, NY, USA, 43–50. <https://doi.org/10.1145/1996092.1996103>
- [10] J. Demmel, L. Grigori, M. Hoemmen, and J. Langou. 2012. Communication-optimal Parallel and Sequential QR and LU Factorizations. *SIAM Journal on Scientific Computing* 34, 1 (2012), A206–A239. <https://doi.org/10.1137/080731992>
- [11] E. Elmroth and F. Gustavson. 2000. Applying recursion to serial and parallel QR factorization leads to better performance. *IBM Journal of Research and Development* 44, 4 (2000), 605–624. <https://doi.org/10.1147/rd.444.0605>
- [12] E. Elmroth, F. Gustavson, I. Jonsson, and B. Kågström. 2004. Recursive blocked algorithms and hybrid data structures for dense matrix library software. *SIAM Rev.* 46, 1 (2004), 3–45. <http://epubs.siam.org/sam-bin/dbq/article/42869>
- [13] D. Helman, D. Bader, and J. Jájá. 1996. Parallel algorithms for personalized communication and sorting with an experimental study. In *ACM Symposium on Parallel Algorithms and Architectures*. ACM, ACM, New York, NY, USA, 211–222. <http://doi.acm.org/10.1145/237502.237558>
- [14] A. Householder. 1958. Unitary triangularization of a nonsymmetric matrix. *J. ACM* 5, 4 (1958), 339–342. <http://doi.acm.org/10.1145/320941.320947>
- [15] Y. Nakatsukasa and R. Freund. 2016. Computing fundamental matrix decompositions accurately via the matrix sign function in two iterations: The power of Zolotarev's functions. *SIAM Rev.* 58, 3 (2016), 461–493. <https://doi.org/10.1137/140990334>
- [16] C. Papadimitriou and J. Ullman. 1987. A communication-time tradeoff. *SIAM J. Comput.* 16, 4 (1987), 639–646. <https://doi.org/10.1137/0216044>
- [17] C. Puglisi. 1992. Modification of the Householder method based on the compact WY representation. *SIAM J. Sci. Statist. Comput.* 13, 3 (1992), 723–726. <https://doi.org/10.1137/0913042>
- [18] M. Schatz, R. Van de Geijn, and J. Poulson. 2016. Parallel matrix multiplication: A systematic journey. *SIAM Journal on Scientific Computing* 38, 6 (2016), C748–C781. <https://doi.org/10.1137/140993478>
- [19] R. Schreiber and C. Van Loan. 1989. A storage-efficient WY representation for products of Householder transformations. *SIAM J. Sci. Statist. Comput.* 10, 1 (1989), 53–57. <http://epubs.siam.org/doi/abs/10.1137/0910005>
- [20] E. Solomonik, G. Ballard, J. Demmel, and T. Hoeffer. 2017. A Communication-Avoiding Parallel Algorithm for the Symmetric Eigenvalue Problem. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '17)*. ACM, New York, NY, USA, 111–121. <https://doi.org/10.1145/3087556.3087561>
- [21] E. Solomonik, E. Carson, N. Knight, and J. Demmel. 2016. Trade-Offs Between Synchronization, Communication, and Computation in Parallel Linear Algebra Computations. *ACM Transactions on Parallel Computing* 3, 1 (2016), 3. <http://doi.acm.org/10.1145/2897188>
- [22] E. Solomonik and J. Demmel. 2011. Communication-optimal Parallel 2.5D Matrix Multiplication and LU Factorization Algorithms. In *Proceedings of the 17th International Conference on Parallel Processing (Euro-Par '11)*. Springer-Verlag, Berlin, Heidelberg, 90–109. <http://dl.acm.org/citation.cfm?id=2033408.2033420>
- [23] X. Sun and C. Bischof. 1995. A basis-kernel representation of orthogonal matrices. *SIAM journal on matrix analysis and applications* 16, 4 (1995), 1184–1196. <http://epubs.siam.org/doi/abs/10.1137/S0895479894276369>
- [24] R. Thakur, R. Rabenseifner, and W. Gropp. 2005. Optimization of Collective Communication Operations in MPICH. *International Journal of High Performance Computing Applications* 19, 1 (2005), 49–66. <https://doi.org/10.1177/1094342005051521>
- [25] A. Tiskin. 2007. Communication-efficient parallel generic pairwise elimination. *Future Generation Computer Systems* 23, 2 (2007), 179–188. <https://doi.org/10.1016/j.future.2006.04.017>
- [26] Leslie G Valiant. 1990. A bridging model for parallel computation. *Commun. ACM* 33, 8 (1990), 103–111. <http://doi.acm.org/10.1145/79173.79181>
- [27] T. Wicky, E. Solomonik, and T. Hoeffer. 2017. Communication-avoiding parallel algorithms for solving triangular systems of linear equations. In *International Parallel and Distributed Processing Symposium*. IEEE, IEEE Computer Society, Washington, DC, USA, 678–687. <https://doi.org/10.1109/IPDPS.2017.104>
- [28] I. Yamazaki, H. Anzt, S. Tomov, M. Hoemmen, and J. Dongarra. 2014. Improving the Performance of CA-GMRES on Multicores with Multiple GPUs. In *International Parallel and Distributed Processing Symposium*. IEEE Computer Society, Washington, DC, USA, 382–391. <https://doi.org/10.1109/IPDPS.2014.48>