Please cite as: https://doi.org/10.1007/s11590-024-02095-z
Christopher Lourenco and Erick Moreno-Centeno (2024) Exact QR Factorizations of Rectangular Matrices. *Optimization Letters* 18(1):681-695.

Exact QR Factorizations of Rectangular Matrices

Christopher Lourenco[1*] and Erick Moreno-Centeno[2†]

[1*]Department of Mathematics, US Naval Academy, Annapolis, 21402, MD, USA.
[2]Department of Industrial and Systems Engineering, Texas A&M University, College Station, 77843, TX, USA.

*Corresponding author(s). E-mail(s): lourenco@usna.edu;
Contributing authors: emc@tamu.edu;
†These authors contributed equally to this work.

Abstract

QR factorization is a key tool in mathematics, computer science, operations research, and engineering. This paper presents the roundoff-error-free (REF) QR factorization framework comprising integer-preserving versions of the standard and the thin QR factorizations and associated algorithms to compute them. Specifically, the standard REF QR factorization factors a given matrix $\boldsymbol{A} \in \mathbb{Z}^{m\times n}$ as $\boldsymbol{A} = \boldsymbol{QDR}$, where $\boldsymbol{Q} \in \mathbb{Z}^{m\times m}$ has pairwise orthogonal columns, $\boldsymbol{D}$ is a diagonal matrix, and $\boldsymbol{R} \in \mathbb{Z}^{m\times n}$ is an upper trapezoidal matrix; notably, the entries of $\boldsymbol{Q}$ and $\boldsymbol{R}$ are integral, while the entries of $\boldsymbol{D}$ are reciprocals of integers. In the thin REF QR factorization, $\boldsymbol{Q} \in \mathbb{Z}^{m\times n}$ also has pairwise orthogonal columns, and $\boldsymbol{R} \in \mathbb{Z}^{n\times n}$ is also an upper triangular matrix. In contrast to traditional (i.e., floating-point) QR factorizations, every operation used to compute these factors is integral; thus, REF QR is guaranteed to be an exact orthogonal decomposition. Importantly, the bit-length of every entry in the REF QR factorizations (and within the algorithms to compute them) is bounded polynomially. Notable applications of our REF QR factorizations include finding exact least squares or exact basic solutions, $\mathbf{x} \in \mathbb{Q}^n$, to any given full column rank or rank deficient linear system $\boldsymbol{A}\mathbf{x} = \mathbf{b}$, respectively. In addition, our exact factorizations can be used as a subroutine within exact and/or high-precision quadratic programming. Altogether, REF QR provides a framework to obtain exact orthogonal factorizations of any rational matrix (as any rational/decimal matrix can be easily transformed into an integral matrix).

Keywords: QR Factorization, Exact Solution of Linear Systems, Least squares problems, Quadratic programming.

1 Introduction

QR factorization is widely used throughout mathematics, engineering, and computer science. Specifically, QR is regularly used to solve linear least squares problems [1, 2], reveal the rank of matrices [3–5], solve highly ill-conditioned linear systems [6, 7], compute Eigenvalues [8–10], and as a subroutine in quadratic programming [11–13].

In the context of quadratic programming, QR factorization was shown to be an efficient tool for solving quadratic programs with high (floating-point) accuracy [12]. Likewise, on a parallel effort, [14] developed and implemented an exact quadratic programming algorithm for computational geometry problems; notably, this exact QP solver was based on an extension of an exact linear programming solver, *not* on QR factorization. This paper bridges the gap between these two disjoint approaches by presenting an exact QR factorization, which may be subsequently used to develop an exact QP algorithm. Specifically, we present a general-purpose roundoff-error-free (REF) QR factorization framework that can be used out-of-the-box to exactly and reliably solve any numerically-challenging application.

That said, the benefits of REF QR extend well beyond quadratic programming because general-purpose (floating point) QR factorizations have been shown to suffer from numerical issues in some applications. For example, if A is very ill-conditioned, the Q matrix, when computed with traditional floating-point approaches, tends to lose orthogonality and thus may lead to incorrect solutions [6, 15]. Indeed, in practice, a previous implementation of (floating-point) QR factorization in LAPACK was shown to numerically fail in rank-revealing computations [16]. Though this shortcoming was successfully addressed therein, another far more recent example, in the context of deep neural networks, illustrates that state-of-the-art, out-of-the-box floating-point QR factorization lacks accuracy when solving some least-square problems [17]. Again, though the issue was addressed in that context (via QR code specialized for their application), the authors state, "Our experience seems to suggest that presently with neural engines, matrix factorizations (QR, LU, Cholesky) are best to be co-designed with their applications (linear solver, least square, orthogonalization, SVD, etc.) to achieve high performance and adequate accuracy and reliability." While this suggestion is worthwhile, our REF QR factorization may be directly used to benchmark specialized floating-point matrix factorizations after they are fine-tuned for specific applications.

In pursuit of this, we introduce two new QR factorizations of the form $A = QDR$, where the columns of Q are pairwise orthogonal, D is diagonal, and R is upper trapezoidal. Notably, both Q and R are comprised of exclusively integer entries with polynomially-bounded bit-lengths, while D (whose entries are all reciprocals of integers with also polynomially-bounded bit-lengths) never needs to be explicitly computed or applied when using the REF QR factorization. Given a matrix $A \in \mathbb{Z}^{m \times n}$, with $m \geq n$, we present two variants of the REF QR factorization: (1) the "thin REF QR Factorization" in which $Q \in \mathbb{Z}^{m \times n}$ and $R \in \mathbb{Z}^{n \times n}$ is the REF Cholesky [18, 19] factor of $A^T A$, and (2) the

"standard REF QR Factorization" in which $Q \in \mathbb{Z}^{m \times m}$ and $R \in \mathbb{Z}^{m \times n}$ is the REF Cholesky factor of $A^T A$ appended by $m - n$ rows of zeros. In addition to proving the existence of the REF QR factorizations, we derive several of their properties, which are analogous to the properties of the traditional QR factorizations as well as derive the algorithms to compute them. Moreover, when A has full column rank, we show that solving the linear system $A\mathbf{x} = \mathbf{b}$ using the REF QR factorization provides the exact least squares solution (i.e., the exact minimizer of the two-norm of $A\mathbf{x} - \mathbf{b}$). Conversely, when A is rank deficient, we show that REF QR can be used to compute an exact basic solution to the system $A\mathbf{x} = \mathbf{b}$.

We note that Zhou and Jeffrey [20] were the first to present an exact thin QR factorization called fraction-free QR. The contributions highlighted above significantly expand on their work; as we develop both thin and standard QR factorizations, prove that they have properties analogous to traditional QR factorizations, and show how to use them to solve both full-rank and rank-deficient linear systems. Thus, another contribution of this work is to generalize the ideas of [20] and relate them to the broader linear algebra body of knowledge.

2 Background

This section briefly reviews the theoretical background needed to derive REF QR. Hereafter, it is assumed that $A \in \mathbb{Z}^{m \times n}$ with $m \geq n$. Note that if $n > m$, the derived properties apply to A^T, while if A is decimal or rational, it can be made integral by multiplying by the appropriate power of 10 or the least common multiple, respectively. Lastly, as in several authoritative matrix linear algebra and numerical analysis textbooks (e.g., [6, 15, 21]), we utilize the MATLAB notation, where the (i, j) element of A, the ith row of A, and the ith column of A are denoted as $A(i, j)$, $A(i, :)$, and $A(:, i)$ respectively; while the submatrix at the intersection of (consecutive) rows i_1 to i_2 and (consecutive) columns j_1 to j_2 is denoted as $A(i_1 : i_2, j_1 : j_2)$. Sections 2.1, 2.2, 2.3, and 2.4 review the traditional QR factorizations, the Pursell & Trimble algorithm to compute the traditional QR factorizations via Gaussian elimination, integer-preserving Gaussian elimination, and REF Cholesky, respectively.

2.1 Traditional QR Factorization

Given a matrix $A \in \mathbb{R}^{m \times n}$, a QR factorization factors A into the product $A = QR$, where Q is orthonormal and R is upper trapezoidal. There are two classes of QR factorization: thin QR and standard QR. Thin QR factorization is typically computed via Gram-Schmidt orthogonalization [15, 22] and results in a rectangular $Q \in \mathbb{R}^{m \times n}$ and a square $R \in \mathbb{R}^{n \times n}$. On the other hand, the standard QR factorization, which is typically referred to just as QR factorization, is computed via either Givens rotations [23] or Householder reflections [24] and results in a square $Q \in \mathbb{R}^{m \times m}$ and a rectangular $R \in \mathbb{R}^{m \times n}$. These traditional QR factorizations have several properties of interest to this paper:

1. Q is an orthonormal matrix. In thin QR, Q is left orthonormal; i.e., $Q^TQ = I$. While, in standard QR, Q is fully orthonormal; i.e., $QQ^T = Q^TQ = I$ and thus $Q^{-1} = Q^T$.
2. R contains the Cholesky factor of A^TA. Specifically, in thin QR, R is square and the Cholesky factor of A^TA; i.e., $R^TR = A^TA$. While, in standard QR, R is rectangular, and its first n rows are the Cholesky factor of A^TA (its last $m - n$ rows of R are zeros).

For an in-depth look at QR factorization, we refer the reader to [15].

2.2 QR Factorization via Gaussian Elimination

The three standard approaches to computing a QR factorization (Gram-Schmidt, Givens, and Householder) are orthogonalization algorithms that construct the orthogonal matrix $\hat{Q}$. Pursell and Trimble present a totally different approach to computing a QR factorization [25], which is based on Gaussian elimination. Specifically, they show that performing Gaussian elimination on the matrix $[A^TA \mid A^T]$ to obtain its row echelon form yields $[\hat{R} \mid \hat{Q}^T]$, where $A = \hat{Q}\hat{R}$ is a scaled version (because the columns of $\hat{Q}$ are not normalized) of the thin (floating-point) QR factorization of A. While the computational complexity of this approach is not competitive with that of the other three approaches, Pursell and Trimble's approach is quite useful in proving several of the theorems in this paper.

2.3 IPGE

Integer-preserving Gaussian elimination (IPGE) is an exact variant of Gaussian elimination used for solving a system of linear equations $A\mathbf{x} = \mathbf{b}$. Given a full rank matrix $A \in \mathbb{Z}^{n\times n}$ and right hand side vector $\mathbf{b}$, denote the kth iteration IPGE matrix as $A^{(k)}$ for $k = 0, \ldots, n$ (where $A^{(0)} \triangleq A$)[1]. Then, assuming there are no row permutations, let $a_{i,j}^{(k)}$ and $\rho^k \triangleq a_{k,k}^{(k-1)}$ denote the individual entries of $A^{(k)}$ and the kth pivot element for $1 \le i \le n$, $1 \le j \le n$, and $0 \le k \le n$ (with $\rho^0 \triangleq 1$), respectively. Then, at iteration k, the IPGE algorithm computes the entries $a_{i,j}^{(k)}$ as follows:

$$a_{i,j}^{(k)} = \begin{cases} a_{i,j}^{(k-1)} & \text{if } i = k, \\ \frac{\rho^k a_{i,j}^{(k-1)} - a_{k,j}^{(k-1)} a_{i,k}^{(k-1)}}{\rho^{k-1}} & \text{otherwise} \end{cases} \tag{1}$$

Note that Equation (1) differs from traditional Gaussian elimination only in the denominator (in traditional Gaussian elimination, the division is by the current pivot, ρ^k, instead of the previous pivot). This seemingly minor modification leads to the following two key properties:

Lemma 2.1. *Throughout all of the IPGE algorithm, the divisions in Equation (1) are guaranteed to be integral (have a zero reminder) [26–28].*

[1] Throughout the paper, the symbol $\triangleq$ means "is defined as."

Lemma 2.2. *Given a matrix A, with $\sigma \triangleq \max_{i,j} \mid a_{i,j}^{(0)} \mid$, the maximum bit length to store any IPGE entry, β_{max}, is upper-bounded polynomially as follows [18, 29]:*

$$\begin{cases} \beta_{max} \leq \lceil n \log(\sigma) \rceil & \text{if } A \text{ is SPD,} \\ \beta_{max} \leq \lceil n \log(\sqrt{n}\sigma) \rceil & \text{if } A \text{ is not SPD} \end{cases}$$

This polynomial bound on IPGE entries is generally pessimistic and is tight if and only if A is diagonal (when A is SPD) or A has orthogonal columns (when A is not SPD). [18, 29–31]

2.4 REF Cholesky

The REF Cholesky factorization [18], based on IPGE, factors an SPD matrix A into the product $A = LDL^T$, where L is a lower triangular matrix comprised of integer entries and $D = diag(\rho^0\rho^1, \rho^1\rho^2, \ldots, \rho^{n-1}\rho^n)^{-1}$. Notably, D is never needed to compute the factorization or when using the factorization to solve an SPD linear system. Along with associated REF forward and backward substitution algorithms, REF Cholesky can exactly solve the SPD linear system $A\mathbf{x} = \mathbf{b}$ exclusively in integer arithmetic. Additionally, the left-looking and up-looking sparse REF Cholesky factorization algorithms were derived in [19]. Notably, these algorithms solve a sparse SPD linear system $A\mathbf{x} = \mathbf{b}$ in asymptotically efficient time complexity, meaning that the dominant cost in these algorithms' complexities is that of the arithmetic operations and do not have ancillary operations such as greatest common divisor operations—required by rational-arithmetic algorithms. Accordingly, these exact factorizations outperform competitor rational-arithmetic LDL and unsymmetric exact factorizations [19].

3 Roundoff-Error-Free QR Factorizations

This section derives two REF QR factorizations, with which a given matrix A is factored as $A = QDR$, and proves their key properties. Specifically, Sections 3.1 and 3.2 formally present the thin REF QR factorization (i.e., where $Q \in \mathbb{Z}^{m\times n}$ and $R \in \mathbb{Z}^{n\times n}$) and the standard REF QR factorization (i.e., where $Q \in \mathbb{Z}^{m\times m}$ and $R \in \mathbb{Z}^{m\times n}$), respectively. Lastly, Sections 3.4 establishes the relationship between the REF QR factorizations and the traditional (floating-point) QR factorizations.

3.1 Thin REF QR Factorization

Theorem 3.1 presents the thin REF QR factorization and its key properties.

Theorem 3.1. *Every full column rank integral matrix $A \in \mathbb{Z}^{m\times n}$ has a unique integral factorization $A = QDR$ with the following properties:*

(a) Q and R are both integral matrices; specifically, $Q \in \mathbb{Z}^{m\times n}$ and $R \in \mathbb{Z}^{n\times n}$.

(b) DR is the REF Cholesky factorization of A^TA; that is $R^TDR = A^TA$, where R is an upper triangular integral matrix, and D is a diagonal matrix. Moreover, $D = diag(\rho^0\rho^1, \rho^1\rho^2, \ldots, \rho^{n-1}\rho^n)^{-1}$, where ρ^i is the ith pivot element of the REF Cholesky factorization of A^TA, and thus $D^{-1} \in \mathbb{Z}^{n\times n}$.

(c) R can be computed as $R = Q^TA$ (just like in the traditional QR factorization $A = QR$).

(d) The columns of Q are pairwise orthogonal. Namely, $Q^TQ = D^{-1}$.

Proof We prove this by constructing the matrices Q and R and showing that they satisfy properties (a)-(e) and that $A = QDR$. First, we obtain R by performing the REF Cholesky factorization on the SPD matrix A^TA. Specifically, we factorize $A^TA = R^TDR$ where R is the upper triangular integral REF Cholesky factor of A^TA and $D = diag(\rho^0\rho^1, \rho^1\rho^2, \ldots, \rho^{n-1}\rho^n)^{-1}$. Since this factorization can be computed by performing IPGE operations on A^TA in order to reduce it to an upper triangular matrix, R; these IPGE operations on A^TA are equivalent to left-multiplying a matrix by $(R^TD)^{-1}$. Next, by left multiplying $[A^TA \mid A^T]$ by $(R^TD)^{-1}$ we obtain the integral matrix $[R \mid (R^TD)^{-1}A^T]$. Again, note that this matrix, $[R \mid (R^TD)^{-1}A^T]$, is an integral matrix because multiplying by $(R^TD)^{-1}$ is equivalent to performing the associated IPGE operation on the full matrix $[A^TA \mid A^T]$ which are guaranteed to be integral [26, 27]. Now letting $Q^T = (R^TD)^{-1}A^T$ (i.e., $Q = A(DR)^{-1}$), it follows that $QDR = A(DR)^{-1}(DR) = A$. We next show that the R and Q integral matrices, which satisfy properties (a) and (b), also satisfy properties (c) and (d).

Property (c) is proved by expanding the product Q^TA as $(R^TD)^{-1}A^TA$, then since $A^TA = R^TDR$, we obtain:

$$Q^TA = (R^TD)^{-1}R^TDR = R.$$

In a similar fashion, property (d) is proved by expanding the product Q^TQ as $(R^TD)^{-1}A^TA(DR)^{-1}$, then since $A^TA = R^TDR$, we obtain:

$$Q^TQ = (R^TD)^{-1}R^TDR(DR)^{-1} = D^{-1}.$$

Finally, the uniqueness of the thin REF QR follows from the following observations: (1) DR is unique (because it is the REF Cholesky factorization of A^TA) and (2) $Q = AR^{-1}D^{-1}$ is a unique product (because A has full column rank). □

3.2 Standard REF QR Factorization

Similar to the thin REF QR factorization, the standard REF QR constructs the factorization $A = QDR$ where Q and R are comprised of exclusively integer entries. In contrast to the thin REF QR, the standard REF QR results in a square $Q \in \mathbb{Z}^{m\times m}$ and an upper trapezoidal $R \in \mathbb{Z}^{m\times n}$. Theorem 3.2 formally introduces this version of REF QR and proves some of its properties.

Theorem 3.2. *Given a full column rank integral matrix $A \in \mathbb{Z}^{m\times n}$, there exists an integral factorization $A = QDR$ with the following properties:*

(a) Q and R are both integral matrices. Specifically $Q \in \mathbb{Z}^{m\times m}$ and $R \in \mathbb{Z}^{m\times n}$.

(b) The REF Cholesky factorization of A^TA, $A^TA = \tilde{R}^T\tilde{D}\tilde{R}$, is embedded in DR. Specifically, the first n rows of R and $D(1:n, 1:n)$ are the REF Cholesky factors $\tilde{R}$ and $\tilde{D}$.

(c) R can be computed as $R = Q^TA$ (just like in the traditional QR factorization $A = QR$).

(d) *The columns of Q are pairwise orthogonal. Namely $Q^TQ = D^{-1}$ (note that D^{-1} is a diagonal integral matrix).*

(e) *Q resembles an orthogonal matrix in that $Q^{-1} = DQ^T$ (whereas, in a truly orthogonal matrix, $Q^{-1} = Q^T$).*

Proof To prove this theorem, we define several useful matrices. Without loss of generality, since A is full column rank, assume the first n rows of A are linearly independent. Construct the matrix $\bar{A} \triangleq [A \ \begin{smallmatrix}\mathbf{0}\\ I_{m-n}\end{smallmatrix}] \triangleq [A \ \bar{I}]$ where $\mathbf{0}$ is the zero matrix of the appropriate size. Let the thin REF QR decompositions of A and $\bar{A}$ be given as $A = \tilde{Q}\tilde{D}\tilde{R}$ and $\bar{A} = \bar{Q}\bar{D}\bar{R}$, respectively.

From the sequence of equations

$$\begin{aligned}[A \ \bar{I}] = \bar{A} = \bar{Q}\bar{D}\bar{R} &= [\bar{Q}_1 \ \bar{Q}_2]\begin{bmatrix}\bar{D}_{11} & \mathbf{0}\\ \mathbf{0} & \bar{D}_{22}\end{bmatrix}\begin{bmatrix}\bar{R}_{11} & \bar{R}_{12}\\ \mathbf{0} & \bar{R}_{22}\end{bmatrix}\\ &= [\bar{Q}_1\bar{D}_{11}\bar{R}_{11} \ \ \bar{Q}_1\bar{D}_{11}\bar{R}_{12} + \bar{Q}_2\bar{D}_{22}\bar{R}_{21}]\end{aligned}$$

we make the following observations:

1. Comparing the first and last terms, we obtain $A = \bar{Q}_1\bar{D}_{11}\bar{R}_{11}$. Furthermore, $\bar{Q}_1 = \tilde{Q}$, $\bar{D}_{11} = \tilde{D}$, and $\bar{R}_{11} = \tilde{R}$ because the thin REF QR factorization of $A = \tilde{Q}\tilde{D}\tilde{R}$ is unique.
2. Let $R = \begin{bmatrix}\bar{R}_{11}\\ \mathbf{0}\end{bmatrix}$. Then, $A = \bar{Q}\bar{D}R$ (this can be observed by ignoring the second column in the decomposition of $\bar{R}$, and performing the matrix multiplication yielding $\bar{Q}_1\bar{D}_{11}\bar{R}_{11}$, which in turn equals A).

The second observation establishes that A can always be factorized as $A = \bar{Q}\bar{D}R$. With this factorization in hand, we now proceed to prove properties 3.2(a) to 3.2(e).

Observing how the matrices $\bar{Q}$ and R were obtained (via thin QR factorizations of A and $\bar{A}$), we conclude that the matrices $\bar{Q}$ and R are integral matrices of the sizes specified in Theorem 3.2; i.e., $\bar{Q} \in \mathbb{Z}^{m\times m}$ and $R \in \mathbb{Z}^{m\times n}$. This proves Property 3.2(a).

Since $A = \tilde{Q}\tilde{D}\tilde{R}$ and Property 3.1(b), we have that $\tilde{R}^T\tilde{D}\tilde{R}$ is the Cholesky factorization of A^TA. Thus, Property 3.2(b) follows from Observation 2 above.

Since $\bar{Q}^T\bar{A} = \bar{R}$, the following sequence of equations shows how Property 3.2(c) follows from Property 3.1(c).

$$\bar{Q}^T\bar{A} = \begin{bmatrix}\bar{Q}_1^T\\ \bar{Q}_2^T\end{bmatrix}[A \ \bar{I}] = \begin{bmatrix}\bar{R}_{11} & \bar{R}_{12}\\ \mathbf{0} & \bar{R}_{22}\end{bmatrix} \Rightarrow \bar{Q}^TA = \begin{bmatrix}\bar{R}_{1,1}\\ \mathbf{0}\end{bmatrix} = R$$

Thereby showing that $\bar{Q}^TA = R$.

Next, Property 3.2(d) $\bar{Q}^T\bar{Q} = \bar{D}^{-1}$ follows from Property 3.1(d).

Finally, to show Property 3.2(e), we utilize Property 3.1(d) again. Specifically, since $Q^TQ = D^{-1}$, we can left multiply by D to obtain $DQ^TQ = I$. Finally, since Q is a full rank $m \times m$ matrix, which follows from Property 3.1(d), we can right multiply by Q^{-1} and we obtain $Q^{-1} = DQ^T$ thereby completing the proof.

□

3.3 Computing REF QR and Theoretical Bounds

The proofs above, as well as Pursell & Trimble [25] and Zhou & Jeffrey [20], give a method to compute thin REF QR based on performing IPGE on $[A^TA \mid A^T]$. By exploiting the properties of our REF QR factorizations, this section

presents (1) a slightly modified yet expedited version of the algorithm for thin REF QR, (2) a new algorithm for standard REF QR, (3) proof that the bit-length of each entry is bounded polynomially, and (4) the computational complexity of both algorithms.

Thin REF QR Algorithm: Perform the REF Cholesky factorization of $[A^TA]$, carrying the row operations throughout A^T. By using REF Cholesky, this approach efficiently exploits the symmetry in A^TA and thus requires about $n^3/2$ fewer operations than Zhou and Jeffrey because they compute a full LU factorization on $[A^TA \mid A^T]$. The total number of operations required to compute thin REF QR with our algorithm is $\mathcal{O}(m^2n + n^3)$, as one must compute A^TA ($\mathcal{O}(n^2m)$), perform REF Cholesky on A^TA ($\mathcal{O}(n^3)$) and, concurrently, carry on the row operations on A^T ($\mathcal{O}(n^2m)$). This algorithm returns the unique thin REF QR factorization $A = QDR$. Note that D does not need to be explicitly calculated as its entries are known directly from R, and, furthermore, as the ensuing sections describe, D is not needed to use the QR factorization.

Standard REF QR Algorithm: Convert A to an appropriate full-rank square matrix $\bar{A} = [A\bar{I}]$ where $\bar{I}$ is the last $m-n$ columns of the $m \times m$ identity matrix (note that the augmented matrix needs only be one which guarantees $\bar{A}$ has full rank). Then, perform thin QR on $\bar{A}$ to obtain $A = \bar{Q}\bar{D}\bar{R}$. The standard REF QR is then $A = \bar{Q}\bar{D}R$ where R comprises the first n columns of $\bar{R}$. This approach requires $\mathcal{O}(m^3)$ operations (it simply performs thin QR discussed above with $n = m$). While Standard REF QR is not necessary for the applications discussed in the next section, this approach gives a method to obtain it if desired.

Theorem 3.3 gives the maximum bit-length required to store (and compute) any entry in REF QR.

Theorem 3.3. *Given a matrix $A \in \mathbb{Z}^{m\times n}$, let σ denote the absolute value of the largest entry in A. The maximum bit-length of every entry in the REF QR factorizations of A, denoted β_{max}, is bounded above by:*

$$\begin{cases} \beta_{max} \leq \lceil 2n \log(m\sigma) \rceil & \text{for thin REF QR,} \\ \beta_{max} \leq \lceil 2m \log(m\sigma) \rceil & \text{for standard REF QR.} \end{cases} \tag{2}$$

Proof For thin QR, R and Q can be obtained by applying IPGE on $[A^TA \mid A^T]$. Every entry in $[A^TA \mid A^T]$ is of at most size $m\sigma^2$, which would occur only if A has a fully dense row of entries of magnitude σ. Though this matrix is of size $n \times n+m$, Edmonds [26] showed that every entry throughout IPGE is the determinant of a square submatrix (in this case, at most a $n \times n$ matrix); thus, using Lemma 2.2, we obtain:

$$\beta_{max} \leq n \log(\sqrt{n}m\sigma^2) \leq n \log(m^2\sigma^2) = 2n \log(m\sigma).$$

The proof is complete by noting that the standard QR bound applies by just changing the dimension of the matrix to be $m \times m$ instead of $m \times n$. □

Theorem 3.3 gives a pessimistic bound for two reasons: (1) Hadamard's bound itself is pessimistic, and (2) it assumes A contains a fully dense row comprised of the largest entries in A. Despite these drawbacks, Theorem 3.3 is

valuable as it shows that every entry in REF QR and within their computation is polynomially bounded.

Finally, we present the complexity of computing each factorization.

Corollary 3.1. *Let $A \in \mathbb{Z}^{m\times n}$ and let β_{max} be the maximum bit-length of any entry in the REF QR factorization of A. Then, the worst-case complexity of computing REF QR is:*

$$\begin{cases} \mathcal{O}(n^2 m(\beta_{max} \log \beta_{max} \log\log \beta_{max})) & \text{for thin REF QR,} \\ \mathcal{O}(m^3(\beta_{max} \log \beta_{max} \log\log \beta_{max})) & \text{for standard REF QR.} \end{cases} \tag{3}$$

The above result follows directly from combining Theorem 3.3 with the number of operations within each algorithm discussed above.

3.4 Relationship of REF QR and Traditional QR

This section gives the mathematical relation between the thin REF QR factorization and the traditional thin QR factorization. Through this relation, one can obtain, from the REF QR, the traditional QR factorization in floating-point to any desired level of precision. Note that this section focuses solely on the thin QR factorization as it is the only unique QR factorization.

Theorem 3.4. *Let $A \in \mathbb{Z}^{m\times n}$ be factored via the thin REF QR as $A = QDR$ and via the traditional thin QR factorization as $A = \hat{Q}\hat{R}$. Then, $\hat{Q} = Q\sqrt{D}$ and $\hat{R} = \sqrt{D}R$.*

Proof Recall that, given a matrix $A \in \mathbb{Z}^{m\times n}$ and its traditional QR factorization ($A = \hat{Q}\hat{R}$), $\hat{R}$ is the (unique) Cholesky factor of A^TA; namely, $\hat{R}^T\hat{R} = A^TA$. Likewise, given the REF QR factorization of A, $A = QDR$, and $R^TDR = A^TA$. Thus, $\sqrt{D}R = \hat{R}$. Therefore, $A = QDR = Q\sqrt{D}\hat{R}$. Furthermore, since $\hat{Q}$ and Q are full rank and unique, it must be the case that $\hat{Q} = Q\sqrt{D}$. □

Interestingly, due to the square root operations, Theorem 3.4 implies that it is effectively impossible for a matrix to have a rational QR factorization (at least without some diagonal matrix akin D to "hide" the square roots).

4 Solving Linear Systems with REF QR

The linear system $A\mathbf{x} = \mathbf{b}$ where $A \in \mathbb{Z}^{m\times n}$ and $\mathbf{b} \in \mathbb{Z}^m$ is not guaranteed to have an exact solution because the product $A\mathbf{x}$ lies in the span of the columns of A which is a proper subspace of $\mathbb{Q}^m$ (since $n < m$). In spite of this, several techniques exist to find an acceptable solution to such systems. Specifically, if A has full column rank, the typical strategy is to find the least squares solution to the system $A\mathbf{x} = \mathbf{b}$; i.e., to find the vector $\mathbf{x}$ which uniquely minimizes the two-norm of $A\mathbf{x} - \mathbf{b}$. Alternatively, if A does not have full column rank, the least squares problem has an infinite number of solutions; thus, a common QR factorization-based approach is to find a basic solution. This section describes how to use REF QR to obtain either an exact solution to the least squares problem (full rank A) or a basic solution (rank deficient A).

4.1 Full-Column-Rank Linear Systems

Given a linear system $A\mathbf{x} = \mathbf{b}$ where $A \in \mathbb{Z}^{m\times n}$, $n < m$, $\mathbf{b} \in \mathbb{Z}^m$, and A has full column rank, the typical strategy is to find the unique least squares solution of the given system, $\mathbf{x}$, which is the unique vector minimizing the two-norm of $A\mathbf{x} - \mathbf{b}$. Theorem 4.1 and Corollary 4.1 show how to use the thin and standard REF QR factorizations, respectively, to obtain the unique least squares solution to $A\mathbf{x} = \mathbf{b}$.

Theorem 4.1. *Given a full column rank matrix, $A \in \mathbb{Z}^{m\times n}$, and a vector, $\mathbf{b} \in \mathbb{Z}^m$, the unique least squares solution to $A\mathbf{x} = \mathbf{b}$ can be obtained by solving $R\mathbf{x} = Q^T\mathbf{b}$, where R and Q are the thin REF QR factors of A.*

Proof First, note that the system $R\mathbf{x} = Q^T\mathbf{b}$ has a unique solution because R has a full rank (as it is the REF Cholesky factor of A^TA).

Second, we show that the solution to $R\mathbf{x} = Q^T\mathbf{b}$ is the exact least squares solution to $A\mathbf{x} = \mathbf{b}$. In particular, we give below a series of equivalent systems that show that $R\mathbf{x} = Q^T\mathbf{b}$ is equivalent to the "normal equations" $A^TA\mathbf{x} = A^T\mathbf{b}$, whose solution is well-known to be the unique least squares solution to $A\mathbf{x} = \mathbf{b}$ [15].

Per Theorem 3.1, using the thin REF QR factorization on A yields $A = QDR$ and $A^TA = R^TDR$; moreover since D is diagonal, $A^T = R^TD^TQ^T = R^TDQ^T$. Substituting these expressions for A^TA and A^T into the normal equations $A^TA\mathbf{x} = A^T\mathbf{b}$, we obtain the equivalent system:

$$R^TDR\mathbf{x} = R^TDQ^T\mathbf{b}$$

The proof is completed by left multiplying both sides by $(R^TD)^{-1}$ (note that both R and D have full rank and thus are invertible), thereby yielding:

$$R\mathbf{x} = Q^T\mathbf{b}$$ □

Corollary 4.1. *Given a full column rank matrix $A \in \mathbb{Z}^{m\times n}$ and vector $\mathbf{b} \in \mathbb{Z}^m$, the exact least squares solution to $A\mathbf{x} = \mathbf{b}$ can be obtained by solving $R(1:n, 1:n)\mathbf{x} = [Q^T\mathbf{b}](1:n)$, where R and Q are the standard REF QR factors of A.*

Proof The proof follows directly from Theorems 3.2 and 4.1. Specifically, Theorem 3.2 states that R and Q contain as submatrices the R and Q from the thin REF QR factorization. Using this, this proof follows directly from Theorem 4.1. □

Given the relationship between the QR matrices of the thin and standard QR factorizations used in the above proof and in Theorem 3.2, for simplicity, the remainder of this section uses only the matrices from thin QR factorization.

Theorem 4.1 and Corollary 4.1 establish that the solution to $R\mathbf{x} = Q^T\mathbf{b}$ is the least squares solution to $A\mathbf{x} = \mathbf{b}$. Next, Theorem 4.2 shows how to solve the system $R\mathbf{x} = Q^T\mathbf{b}$ exactly, and thus how to obtain the unique least squares solution to $A\mathbf{x} = \mathbf{b}$ exactly (free of roundoff errors). Specifically, by scaling the right-hand side vector, one can obtain the exact rational solution to $R\mathbf{x} = Q^T\mathbf{b}$ almost entirely in integer arithmetic, using rational numbers only in a final division.

Theorem 4.2. *The exact solution to the linear system* $R\mathbf{x} = Q^T\mathbf{b}$ *can be obtained entirely in integer arithmetic.*

Proof The proof follows from Theorem 4.1, Cramer's rule, and REF backward substitution [18]. From Theorems 4.1 and 4.1, we know that the systems $R\mathbf{x} = Q^T\mathbf{b}$ and $A^TA\mathbf{x} = A^T\mathbf{b}$ are equivalent. By Cramer's rule and since the determinant of A^TA is $R(n,n)$ [18], the system $R\mathbf{x} = R(n,n)Q^T\mathbf{b}$ has an integral solution vector. Therefore, applying REF backward substitution [18] to this scaled system, one can obtain its solution entirely in integer arithmetic. Then, the exact solution of the original system is given as an integral numerator vector $\mathbf{x}$ and integral denominator $R(n,n)$. □

Note that as a consequence of Theorem 4.2, one can obtain the exact two-norm solution either as the above rational vector or to any level of precision. To conclude this subsection, we note that solving the normal equations directly with Cholesky would have an equivalent worst-case computational complexity to QR factorization. Yet, depending on the application and structure of the matrix, one method can be preferred over the other, and it is not always trivial to determine in advance which method would be better (see, for example, the detailed discussion in [15, 32] comparing both approaches in the floating-point case). Examining the same nuances computationally in the exact case would be an interesting problem, but it is beyond the scope of this paper.

4.2 Rank-Deficient Linear Systems

Given a linear system $A\mathbf{x} = \mathbf{b}$ where $A \in \mathbb{Z}^{m\times n}$, $n < m$, $\mathbf{b} \in \mathbb{Z}^m$, and A is rank-deficient, there are an infinite number of solutions to the linear least squares problem $A\mathbf{x} = \mathbf{b}$. Several approaches exist, including finding the minimum norm solution, the truncated SVD solution, or a basic solution [15]. The first approach relies on performing either a complete orthogonal decomposition or finding a pseudoinverse (see [33, 34]), the second approach is SVD based, and the third approach uses QR factorization. This subsection shows how to use the third approach and REF QR to find an exact basic solution of $A\mathbf{x} = \mathbf{b}$.

The QR-based approach is to perform QR with column pivoting; that is, the factorization $QR = AP$ is performed, where P is a permutation matrix chosen during factorization (specifically if a rank-deficient column is found—i.e., a column where one cannot find an eligible non-zero element to pivot—, the rank-deficient column is replaced with a different column; this process is repeated until no more linearly independent columns exist). Given a matrix A with rank r such that $r < n$, without loss of generality, assume that the first r columns of A are linearly independent. Then, the following lemma shows the structure of the thin REF QR factorization of A.

Lemma 4.1. *Given a rank-deficient matrix* $A \in \mathbb{Z}^{m\times n}$ *with rank* $r < n$*, the thin REF QR decomposition of* A *has the following structure:*

$$A = QDR = \begin{bmatrix} Q_{11} & 0 \\ Q_{21} & 0 \end{bmatrix} \begin{bmatrix} D_{11} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} R_{11} & R_{12} \\ 0 & 0 \end{bmatrix}$$

where Q_{11}, D_{11}, R_{11} *are all of dimension* $r \times r$.
Proof Consider performing IPGE on the matrix $[A^TA \mid A^T]$. Since A has column rank r, $A^TA \in \mathbb{Z}^{n\times n}$ has rank r and A^T has row rank r. The result follows trivially from the fact that IPGE on this matrix will lead to a row of $n-r$ zeros on both A^TA and A^T, giving the factorization shown above. □

Thus, a basic solution is found by solving the system $R\mathbf{x} = Q^T\mathbf{b}$, which expands to:

$$\begin{bmatrix} R_{11} & R_{12} \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} = \begin{bmatrix} Q_{11}^T & Q_{21}^T \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{bmatrix}$$

Thus, a unique exact basic solution is found by setting $\mathbf{x}_2 = 0$, and solving the square linear system $R_{11}\mathbf{x}_1 = Q_{11}^T\mathbf{b}_1$. Note that Golub & Pareyra [35] show that such a solution minimizes the two-norm if and only if $R_{12} = 0$. Finding exact two-norm minimizers of rank-deficient systems which have $R_{12} \neq 0$ is an important topic but requires methods outside the scope of this paper.

5 Conclusion

This paper presents and thoroughly analyzes a comprehensive REF QR factorization framework to obtain exact factorizations of any matrix. Specifically, this paper presents the thin and standard REF QR factorizations, which exactly factorize the matrix A into the product $A = QDR$ where Q has pairwise orthogonal columns, D is diagonal, and R is upper trapezoidal. Matrices Q and R comprise exclusively integer entries, while the diagonal matrix D, whose entries are reciprocal of integers, is never explicitly needed (nor computed by the algorithms). Importantly, we derive properties of our REF QR factorizations that are analogous to those of the floating point QR factorizations and illustrate how our exact factorizations are related to the traditional (inexact) floating-point QR factorizations—specifically, we show how to obtain the inexact factorizations from our REF QR factorizations. Furthermore, we present algorithms to compute each factorization and prove that the size of each entry in REF QR (and throughout their computation) is bounded polynomially. Finally, we discuss solving the system $A\mathbf{x} = \mathbf{b}$ for both full-column rank and rank-deficient systems. Specifically, if A has full column rank, REF QR finds the exact least squares solution to $A\mathbf{x} = \mathbf{b}$. Conversely, if A is rank deficient, utilizing column pivoting, REF QR can find an exact basic solution of $A\mathbf{x} = \mathbf{b}$.

Declarations

- The first author was partially supported by the USNA JR NARC. The second author was partially supported by NSF under Grant No OAC-1835499.
- Data sharing does not apply to this article as no datasets were generated or analyzed during this study.

References

[1] Bauer, F.L.: Elimination with weighted row combinations for solving linear equations and least squares problems. Numerische Mathematik **7**(4), 338–352 (1965) https://doi.org/10.1007/BF01436528

[2] Golub, G.H., Wilkinson, J.H.: Note on the iterative refinement of least squares solution. Numerische Mathematik **9**(2), 139–148 (1966)

[3] Chan, T.F.: Rank revealing QR factorizations. Linear algebra and its applications **88**(C), 67–82 (1987) https://doi.org/10.1016/0024-3795(87)90103-0

[4] Chan, T.F., Hansen, P.C.: Some applications of the rank revealing QR factorization. STAT. COMPUT **13**(3), 727–741 (1992)

[5] Higham, N.J.: QR factorization with complete pivoting and accurate computation of the SVD. Linear Algebra and its Applications **309**(1-3), 153–174 (2000)

[6] Higham, N.J.: Accuracy and Stability of Numerical Algorithms, 2nd edn. Siam, Philadelphia, PA, USA (2002)

[7] Fukaya, T., Kannan, R., Nakatsukasa, Y., Yamamoto, Y., Yanagisawa, Y.: Shifted Cholesky QR for computing the QR factorization of ill-conditioned matrices. SIAM Journal on Scientific Computing **42**(1), 477–503 (2020) https://doi.org/10.1137/18M1218212 arXiv:1809.11085

[8] Francis, J.G.F.: The QR Transformation A Unitary Analogue to the LR Transformation–Part 1. The Computer Journal **4**(3), 265–271 (1961) https://doi.org/10.1093/comjnl/4.3.265

[9] Kublanovskaya, V.N.: On some algorithms for the solution of the complete eigenvalue problem. USSR Computational Mathematics and Mathematical Physics **1**(3), 637–657 (1962) https://doi.org/10.1016/0041-5553(63)90168-X

[10] Parlett, B.N., Poole Jr, W.G.: A geometric theory for the QR, LU and power iterations. SIAM Journal on Numerical Analysis **10**(2), 389–412 (1973)

[11] Gill, P.E., Murray, W.: Numerically stable methods for quadratic programming. Mathematical Programming **14**(1), 349–372 (1978) https://doi.org/10.1007/BF01588976

[12] Arioli, M.: The use of QR factorization in sparse quadratic programming

and backward error issues. SIAM Journal on Matrix Analysis and Applications **21**(3), 825–839 (2000) https://doi.org/10.1137/S0895479898338147

[13] Weber, T., Sager, S., Gleixner, A.: Solving quadratic programs to high precision using scaled iterative refinement. Mathematical Programming Computation **11**(3), 421–455 (2019)

[14] Gärtner, B., Schönherr, S.: An efficient, exact, and generic quadratic programming solver for geometric optimization. In: Proceedings of the Sixteenth Annual Symposium on Computational Geometry, pp. 110–118 (2000)

[15] Golub, G.H., Van Loan, C.F.: Matrix Computations, 4th edn. Johns Hopkins Studies in the Mathematical Sciences, vol. 3. Johns Hopkins University Press, Baltimore (2013). http://books.google.com/books?id=mlOa7wPX6OYC

[16] Drmač, Z., Bujanović, Z.: On the failure of rank-revealing QR factorization software - A case study. ACM Transactions on Mathematical Software (TOMS) **35**(2), 1–28 (2008) https://doi.org/10.1145/1377612.1377616

[17] Zhang, S., Baharlouei, E., Wu, P.: High accuracy matrix computations on neural engines: A study of qr factorization and its applications. In: Proceedings of the 29th International Symposium on High-Performance Parallel and Distributed Computing, pp. 17–28. Association for Computing Machinery, Inc, New York, USA (2020). https://doi.org/10.1145/3369583.3392685

[18] Escobedo, A.R., Moreno-Centeno, E.: Roundoff-Error-Free Algorithms for Solving Linear Systems via Cholesky and LU Factorizations. INFORMS Journal on Computing **27**(4), 677–689 (2015) https://doi.org/10.1287/ijoc.2015.0653

[19] Lourenco, C.J., Moreno-Centeno, E.: Exactly Solving Sparse Rational Linear Systems via Roundoff-Error-Free Cholesky Factorizations. SIAM Journal on Matrix Analysis and Applications **43**(1), 439–463 (2022) https://doi.org/10.1137/20M1371592

[20] Zhou, W., Jeffrey, D.J.: Fraction-free matrix factors: new forms for LU and QR factors. Frontiers of Computer Science in China **2**(1), 67–80 (2008) https://doi.org/10.1007/s11704-008-0005-z

[21] Davis, T.A.: Direct Methods for Sparse Linear Systems. SIAM, Philadelphia, PA, USA (2006)

[22] Rice, J.R.: Experiments on gram-schmidt orthogonalization. Mathematics

of Computation **20**(94), 325–328 (1966)

[23] Givens, W.: Computation of plain unitary rotations transforming a general matrix to triangular form. Journal of the Society for Industrial and Applied Mathematics **6**(1), 26–50 (1958)

[24] Householder, A.S.: Unitary Triangularization of a Nonsymmetric Matrix. Journal of the ACM (JACM) **5**(4), 339–342 (1958)

[25] Pursell, L., Trimble, S.Y.: Gram-Schmidt Orthogonalization by Gauss Elimination. The American Mathematical Monthly **98**(6), 544–549 (1991) https://doi.org/10.1080/00029890.1991.11995755

[26] Edmonds, J.: Systems of distinct representatives and linear algebra. J. Res. Nat. Bur. Standards Sect. B **71**(4), 241–245 (1967)

[27] Bareiss, E.H.: Sylvester's identity and multistep integer-preserving Gaussian elimination. Mathematics of computation **22**(103), 565–578 (1968)

[28] Montante-Pardo, R.M., Méndez-Cavazos, M.A.: Un método número para cálculo matricial. Revista Técnico-Científica de Divulgación **2**, 1–24 (1977)

[29] Lourenco, C., Escobedo, A.R., Moreno-Centeno, E., Davis, T.A.: Exact Solution of Sparse Linear Systems via Left-Looking Roundoff-Error-Free LU Factorization in Time Proportional to Arithmetic Work. SIAM Journal on Matrix Analysis and Applications **40**(2), 609–638 (2019) https://doi.org/10.1137/18M1202499

[30] Hadamard, J.: Résolution d'une question relative aux déterminants. Bull. sci. math **17**(1), 240–246 (1893)

[31] Horn, R.A., Johnson, C.R.: Matrix Analysis. Cambridge University Press, New York, NY, USA (2012)

[32] Davis, T.A., Rajamanickam, S., Sid-Lakhdar, W.M.: A survey of direct methods for sparse linear systems. Acta Numerica **25**, 383–566 (2016) https://doi.org/10.1017/S0962492916000076

[33] Foster, L., Kommu, R.: Algorithm 853: An efficient algorithm for solving rank-deficient least squares problems. ACM Transactions on Mathematical Software (TOMS) **32**(1), 157–165 (2006)

[34] Foster, L.V., Davis, T.: Reliable Calculation of Numerical Rank, Null Space Bases, Basic Solutions and Pseudoinverse Solutions using SuiteSparseQR. In: Householder Symposium XVIII on Numerical Linear Algebra, p. 79 (2011)

[35] Golub, G., Pereyra, V.: Differentiation of pseudoinverses, separable nonlinear least squares and other tales. Generalized Inverses and Applications pp. 303-324. Academic Press, New York (1976)