



H2020-FETHPC-2014: GA 671633

NLAFET Working Note 9

Robust Solution of Triangular Linear Systems

Carl Christian Kjelgaard Mikkelsen and Lars Karlsson

April, 2017

Document information

This preprint report is also published as Report UMINF 17.09 at the Department of Computing Science, Umeå University, Sweden.

This version: May 30, 2017.

Acknowledgements

This project has received funding from the *European Union's Horizon 2020 research and innovation programme* under the grant agreement number 671633.

Robust Solution of Triangular Linear Systems *

Carl Christian Kjelgaard Mikkelsen, Lars Karlsson

May 30, 2017

Abstract

Let $T \in \mathbb{C}^{n \times n}$ be a non-singular upper triangular matrix and let $\mathbf{b} \in \mathbb{C}^n$. We show how to compute a scaling $\alpha \in (0, 1]$ and $\mathbf{x} \in \mathbb{C}^n$, such that not only is $T\mathbf{x} = \alpha\mathbf{b}$, but all intermediate results are bounded by a threshold Ω . This problem is relevant when estimating condition numbers or computing eigenvectors. Our scalar algorithm is similar to `xLATRS` from LAPACK. We explain why it is difficult to parallelize these algorithms. We then develop a robust blocked algorithm which can be executed in parallel using a run-time systems such as StarPU.

Contents

1	Introduction	1
2	A scalar algorithm	2
2.1	Robust scalar divisions	2
2.2	Robust linear updates	4
2.3	Linear updates with linear growth	6
2.4	Robust scalar backward substitution	7
3	A blocked algorithm	8
3.1	Robust block solve	9
3.2	Robust block update	9
3.3	Robust blocked backward substitution	10
4	A robust parallel algorithm	11
5	Conclusion	11

1 Introduction

Let $T \in \mathbb{C}^{n \times n}$ be a non-singular upper triangular matrix and let $\mathbf{b} \in \mathbb{C}^n$. In this note we study the problem of computing a scalar $\alpha \in (0, 1]$ such that the solution $\mathbf{x} \in \mathbb{C}^n$ of the scaled linear system

$$T\mathbf{x} = \alpha\mathbf{b}$$

*NLAFFET Working Note 9. Report UMINF 17.09, Dept. Computing Science, Umeå University, SE 901 87 Umeå, Sweden.

can be computed without exceeding a threshold Ω . Useful values of Ω are all less than the largest value which can be represented.

This problem is relevant in the context of estimating condition numbers [1] and computing eigenvectors [3]. In both cases it is necessary to solve triangular linear systems which are potentially very ill-conditioned. When computing eigenvectors we are solving systems of the form $(T - \lambda I)x = b$ where λ . If there are many nearby eigenvalues, then overflow is a distinct possibility, because $T - \lambda I$ has many small diagonal elements.

The algorithms developed by Edward Anderson have been implemented in `xLATRS` and the principles have been integrated into `xTREVC` which computes eigenvectors of triangular matrices. ScaLAPACK contains complex flavors `PCTREVC` and `PZTREVC` which call `PxLATRS` to solve triangular linear systems robustly. These auxiliary routines have been derived from `xLATRS` using PBLAS. This significantly limits the parallel efficiency of `PxTREVC` when numerical scaling is required.

In this note we develop a scalar and a blocked algorithm for solving (1) with respect to (α, \mathbf{x}) . The blocked algorithm can be executed in parallel using a task based run-time system such as StarPU [2]. The solve and update kernels are replaced with robust variants which operate on *augmented* vectors $\langle \alpha, \mathbf{x} \rangle$ rather than regular vectors.

The parallel overhead is increased marginally compared with a task based implementation of non-robust backward substitution. In the long run, this will allow us to remove the parallel bottleneck associated with the robust parallel computation of eigenvectors.

2 A scalar algorithm

In this section we derive a robust scalar algorithm for solving the scaled upper triangular linear system (1). Our algorithm is closely related to the algorithm implemented in `xLATRS` in LAPACK, [1]. Our primary contribution is to simplify and extend the analysis of the underlying problem.

An upper triangular linear system can be solved using scalar backward substitution, i.e., Alg. 1 `ScalarSolve`. It consists of a sequence of scalar divisions

$$x \leftarrow b/t, \quad t \neq 0 \tag{1}$$

and linear updates

$$z \leftarrow y - tx.$$

In Section 2.1 and Section 2.2 we show how to prevent overflow for each of these fundamental operations. We then derive a robust variant of scalar backward substitution in Section 2.4.

2.1 Robust scalar divisions

The scalar division (1) cannot exceed Ω , if $|b| \leq |t|\Omega$. In general, we seek a scaling α such that

$$|\alpha b| \leq |t|\Omega, \tag{2}$$

which implies that the scaled division

$$y \leftarrow \frac{(\alpha b)}{t} \tag{3}$$

cannot exceed Ω . The computation of a suitable scaling α is the subject of the following theorem.

Algorithm 1: $\mathbf{x} = \text{ScalarSolve}(\mathbf{T}, \mathbf{b})$

Data: A non-singular upper triangular matrix $\mathbf{T} \in \mathbb{C}^{n \times n}$ and a vector $\mathbf{b} \in \mathbb{C}^n$.

Result: The solution \mathbf{x} of the linear system $\mathbf{T}\mathbf{x} = \mathbf{b}$.

```

1  $\mathbf{x} \leftarrow \mathbf{b}$ ;
2 for  $j = n, n-1, \dots, 1$  do
3    $x_j \leftarrow x_j/t_{jj}$ ;
4   for  $i = 1, 2, \dots, j-1$  do
5      $x_i \leftarrow x_i - t_{ij}x_j$ ;
6 return  $\mathbf{x}$ ;

```

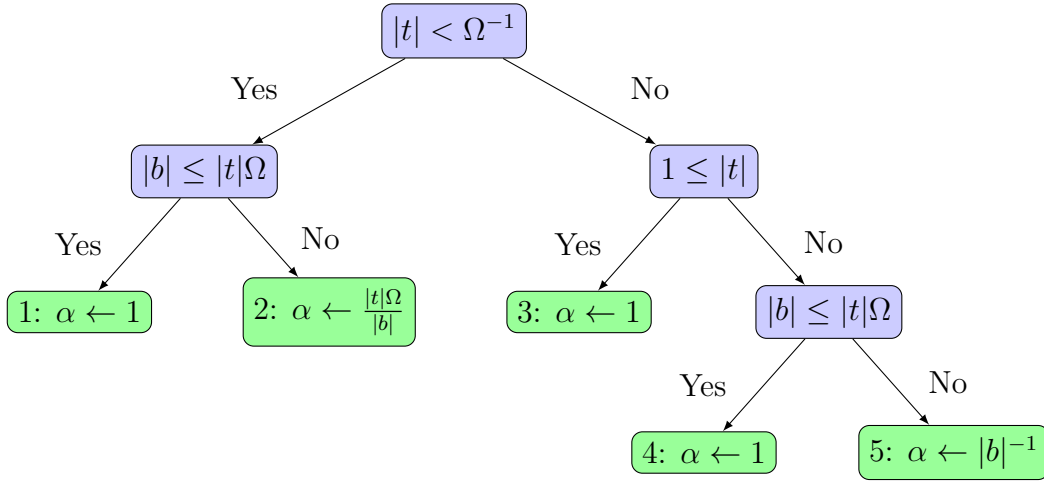


Figure 1: Computing a scaling factor α such that the division (3) cannot overflow.

Theorem 1. Let $b \in \mathbb{C}$ and let $t \neq 0$ such that b/t is defined. If $|b| \leq \Omega$, then the scaling factor α produced by the flowchart in Fig. 1 satisfies $\alpha \in (0, 1]$ and ensures that the result of the scaled division (3) satisfies $|y| \leq \Omega$.

Proof. There are five cases to consider, one for each leaf shown in Fig. 1. In each case, the proof consists of verifying that the central inequality (2) is satisfied for the given value of α .

1. We already have $|b| \leq |t|\Omega$, so no scaling is necessary.
2. We have $|b| > |t|\Omega$, which implies $\alpha = \frac{|t|\Omega}{|b|} < 1$ and $|y| = \frac{\alpha|b|}{|t|} = \Omega$.
3. We have $1 \leq |t|$ and $|b| \leq \Omega$, which implies $|b| \leq |t|\Omega$. Therefore, no scaling is necessary.
4. We already have $|b| \leq |t|\Omega$, so no scaling is necessary.
5. We have $|t| \geq \Omega^{-1}$ and $|b| > |t|\Omega$, so

$$\alpha = |b|^{-1} < |t|^{-1}\Omega^{-1} \leq 1.$$

Moreover,

$$|y| = \frac{\alpha|b|}{|t|} = \frac{1}{|t|} \leq \Omega.$$

This completes the proof. \square

Remark 1. *Thm. 1 requires $|b| \leq \Omega$. If we also have $|t| \leq \Omega$, then the right hand side and the left hand side of each inequality in Fig. 1 can be evaluated without exceeding Ω .*

Alg. 2 `ProtectDivision` summarizes how to safely compute a scaling factor which prevents overflow in the scaled division (3).

Algorithm 2: $\alpha = \text{ProtectDivision}(b, t)$

Data: Numbers b and t such that $|b| \leq \Omega$ and $t \neq 0$.

Result: A scaling $\alpha \in (0, 1]$ such that the scaled division (3) cannot overflow.

```

1  $\alpha \leftarrow 1$ ;
2 if  $|t| < \Omega^{-1}$  then
3   if  $|b| > |t|\Omega$  then
4      $\alpha \leftarrow \frac{|t|\Omega}{|b|}$ ;
5 else
6   if  $|t| < 1$  then
7     if  $|b| > |t|\Omega$  then
8        $\alpha \leftarrow |b|^{-1}$ ;
9 return  $\alpha$ ;
```

2.2 Robust linear updates

We consider more the general problem of computing a scaling $\xi \in (0, 1]$, such that the scaled linear transformation

$$\mathbf{Z} \leftarrow \xi \mathbf{Y} - \mathbf{T}(\xi \mathbf{X})$$

satisfies $\|\mathbf{Z}\|_\infty \leq \Omega$. We begin by stating a condition which ensures that no scaling is necessary.

Theorem 2. *Let \mathbf{Y}, \mathbf{T} and \mathbf{X} be matrices such that $\mathbf{Z} = \mathbf{Y} - \mathbf{T}\mathbf{X}$ is defined. If*

$$\|\mathbf{Y}\|_\infty + \|\mathbf{T}\|_\infty \|\mathbf{X}\|_\infty \leq \Omega,$$

then overflow is impossible in the calculation of \mathbf{Z} regardless of the order of the necessary arithmetic operations.

Proof. By the triangle inequality we have the elementary estimate

$$|z_{ij}| \leq \|\mathbf{Z}\|_\infty \leq \|\mathbf{Y}\|_\infty + \|\mathbf{T}\|_\infty \|\mathbf{X}\|_\infty.$$

It follows that each component of the *final* result \mathbf{Z} is dominated by Ω , i.e., $|z_{ij}| \leq \Omega$. There are many ways to evaluate the the expression

$$z_{ij} = y_{ij} - \sum_k t_{ik} x_{kj},$$

but each *intermediate* result can be written as

$$z_{ij}(\mu, \boldsymbol{\sigma}) = \mu y_{ij} - \sum_k \sigma_k t_{ik} x_{kj}, \quad \mu \in \{0, 1\}, \quad \sigma_k \in \{0, 1\},$$

where the choice of, say, $\sigma_k = 1$ reflects that the corresponding term $t_{ik}x_{kj}$ has been included in the calculation. In all cases, the triangular inequality implies

$$|z_{ij}(\mu, \boldsymbol{\sigma})| \leq |y_{ij}| + \sum_k |t_{ik}| |x_{kj}| = |\mathbf{Y}|_{ij} + (|\mathbf{T}||\mathbf{X}|)_{ij} \leq$$

$$\|\mathbf{Y}\|_\infty + \|\mathbf{T}||\mathbf{X}\|_\infty \leq \|\mathbf{Y}\|_\infty + \|\mathbf{T}\|_\infty \|\mathbf{X}\|_\infty \leq \Omega.$$

It follows, that all components of all intermediate results are bounded by Ω , regardless of the order of the necessary arithmetic operations. \square

In general, we seek a scaling factor ζ such that the scaled linear transformation

$$\mathbf{Z} \leftarrow (\zeta \mathbf{Y}) - \mathbf{T}(\zeta \mathbf{X}) \quad (4)$$

can be computed without exceeding Ω . The computation of a suitable scaling ζ is the subject of the following theorem.

Theorem 3. *Let \mathbf{Y} , \mathbf{T} , and \mathbf{X} be matrices, such that $\mathbf{Z} = \mathbf{Y} - \mathbf{T}\mathbf{X}$ is defined. If*

$$\|\mathbf{Y}\|_\infty \leq \Omega, \quad \|\mathbf{T}\|_\infty \leq \Omega,$$

then the scaling factor ζ produced by the flowchart in Fig. 2 satisfies $\zeta \in (0, 1]$ and prevents overflow in the scaled linear update (4).

Proof. There are four cases to consider, one for each leaf shown in Fig. 2. In each case, the proof consists of verifying that

$$\|\zeta \mathbf{Y}\|_\infty + \|\mathbf{T}\|_\infty \|\zeta \mathbf{X}\|_\infty \leq \Omega \quad (5)$$

for the given value of ζ . By Thm. 2 this ensures that the scaled linear update (4) cannot overflow

1. In this case,

$$\|\mathbf{T}\|_\infty \|\mathbf{X}\|_\infty \leq \Omega - \|\mathbf{Y}\|_\infty$$

and so the central inequality (5) is satisfied for all $|\zeta| \leq 1$.

2. In this case, we have $\zeta = \frac{1}{2}$. We then use

$$\|\mathbf{Y}\|_\infty \leq \Omega, \quad \|\mathbf{T}\|_\infty \leq \Omega, \quad \|\mathbf{X}\|_\infty \leq 1,$$

to estimate

$$\|\mathbf{Z}\|_\infty \leq \|\zeta \mathbf{Y}\|_\infty + \|\mathbf{T}\|_\infty \|\zeta \mathbf{X}\|_\infty = \frac{1}{2} \|\mathbf{Y}\|_\infty + \frac{1}{2} \|\mathbf{T}\|_\infty \|\mathbf{X}\|_\infty \leq \Omega.$$

3. In this case,

$$\|\mathbf{T}\|_\infty \leq \frac{\Omega - \|\mathbf{Y}\|_\infty}{\|\mathbf{X}\|_\infty}$$

and so the central inequality (5) is satisfied for all $|\zeta| \leq 1$.

4. In this case $\|\mathbf{X}\|_\infty > 1$, so $\zeta = \frac{1}{2\|\mathbf{X}\|_\infty} < \frac{1}{2}$. Moreover,

$$\|\zeta \mathbf{Y}\|_\infty + \|\mathbf{T}\|_\infty \|\zeta \mathbf{X}\|_\infty < \frac{1}{2} (\|\mathbf{Y}\|_\infty + \|\mathbf{T}\|_\infty) \leq \Omega.$$

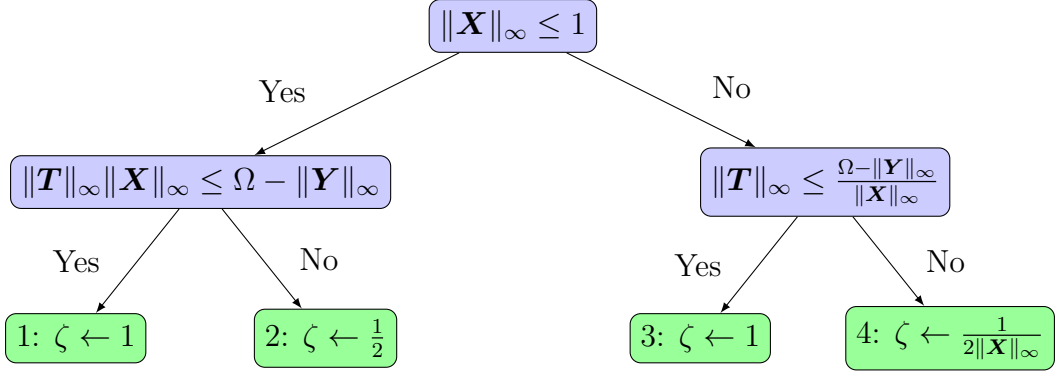


Figure 2: Computing a scaling factor ζ such that the linear update (4) cannot overflow.

This completes the proof. \square

Remark 2. *Thm. 3* requires $\|\mathbf{Y}\|_\infty \leq \Omega$ and $\|\mathbf{T}\|_\infty \leq \Omega$. If we also have $\|\mathbf{X}\|_\infty \leq \Omega$, then the right hand side and the left hand side of each inequality in Fig. 2 can be evaluated without exceeding Ω .

Alg. 3 `ProtectUpdate` summaries how to safely compute a scaling such that the scaled linear update (4) cannot overflow.

Algorithm 3: $\zeta = \text{ProtectUpdate}(y_{\text{norm}}, t_{\text{norm}}, x_{\text{norm}},)$

Data: Non-negative real numbers y_{norm} , t_{norm} and x_{norm} such that

$$\|\mathbf{Y}\|_\infty \leq y_{\text{norm}} \leq \Omega, \quad \|\mathbf{T}\|_\infty \leq t_{\text{norm}} \leq \Omega, \quad \|\mathbf{X}\|_\infty \leq x_{\text{norm}}.$$

Result: A scaling factor ζ such that

$$\zeta (y_{\text{norm}} + t_{\text{norm}} x_{\text{norm}}) \leq \Omega,$$

which implies that the scaled linear update (4) cannot exceed Ω .

```

1  ζ ← 1;
2  if x_norm ≤ 1 then
3  |   if t_norm x_norm > Ω - b_norm then
4  |   |   ζ ← 1/2;
5  else
6  |   if t_norm > (Ω - b_norm) / x_norm then
7  |   |   ζ ← 1 / (2x_norm);
8  return ζ;
  
```

2.3 Linear updates with linear growth

The following theorem is a weaker result which is useful in the context of solving linear systems.¹ It does not require the computation of $\|\mathbf{Y}\|_\infty$. We are interested in the situation where $\Omega' \ll \Omega$, i.e., when there is ample room for growth.

¹This section will be moved and updated with additional information in the future.

Theorem 4. Let ξ be the scaling factor computed using the simplified flowchart in Fig. 3. If $\|\mathbf{T}\|_\infty \leq \Omega'$ and $\|\mathbf{Y}\|_\infty \leq k\Omega'$, then $\mathbf{Z} = \xi\mathbf{Y} - \mathbf{T}(\xi\mathbf{X})$ satisfies $\|\mathbf{Z}\|_\infty \leq (k+1)\Omega'$.

Proof. The proof is straightforward and consists of verifying the conclusion in each of the three cases represented by the leaves in Fig. 3. In general, the triangle inequality implies

$$\|\xi\mathbf{Z}\|_\infty \leq \xi\|\mathbf{Y}\|_\infty + \xi\|\mathbf{T}\|_\infty\|\mathbf{X}\|_\infty. \quad (6)$$

1. In this case, $\xi = 1$, $\|\mathbf{X}\|_\infty \leq 1$ and (6) implies, that

$$\|\mathbf{Z}\|_\infty \leq \|\mathbf{Y}\|_\infty + \|\mathbf{T}\|_\infty \leq k\Omega' + \Omega' = (k+1)\Omega'.$$

2. In this case, $\xi = 1$, $\|\mathbf{T}\|_\infty\|\mathbf{X}\|_\infty \leq \Omega'$ and (6) implies, that

$$\|\mathbf{Z}\|_\infty \leq k\Omega' + \Omega' = (k+1)\Omega'.$$

3. In this case $\xi = \frac{1}{\|\mathbf{X}\|_\infty} < 1$ and

$$\|\mathbf{Z}\|_\infty \leq \|\mathbf{Y}\|_\infty + \|\mathbf{T}\|_\infty \leq (k+1)\Omega'.$$

This completes the proof. □




Figure 3: A simplified computation which allows for linear growth, see Thm. 4.

2.4 Robust scalar backward substitution

In Sections 2.1 and 2.2 we have shown how to protect a division and a linear update against overflow. It is straightforward to combine them into Alg. 4 `RobustScalarSolve`.

The reader should consider the assumptions placed on the input. Assumption (7) ensures that $\|\mathbf{T}_{1:j-1,j}\|_\infty \leq \Omega$, so that $c_j = \|\mathbf{T}_{1:j-1,j}\|_\infty$ is a viable option. Assumption (8) ensures that $x_{\max} \leq \Omega$ initially. This critical property is preserved by virtue of the properties of Alg. 2 `ProtectDivision` and Alg. 3 `ProtectUpdate`.

The reader should also consider the problem of parallelizing `RobustScalarSolve`. Any scaling of \mathbf{x} requires global communication as does the re-calculation of x_{\max} at the end of each iteration. It is clear that alternatives should be explored.

Algorithm 4: $[\alpha, \mathbf{x}] = \text{RobustScalarSolve}(\mathbf{T}, \mathbf{b})$

Data: A non-singular upper triangular matrix $\mathbf{T} \in \mathbb{C}^{n \times n}$, and $\mathbf{b} \in \mathbb{C}^n$, subject to the following constraints

$$|t_{ij}| \leq \Omega \quad \text{for } i < j, \quad (7)$$

and

$$|b_j| \leq \Omega. \quad (8)$$

Result: A scaling $\alpha \in (0, 1]$ and the solution \mathbf{x} of the scaled linear system

$$\mathbf{T}\mathbf{x} = \alpha\mathbf{b}$$

where α ensures that \mathbf{x} can be computed without exceeding Ω .

```
1 Find  $c_j$  such that  $\|T_{1:j-1,j}\|_\infty \leq c_j \leq \Omega$  for  $j = 2, 3, \dots, n$ ;  
2  $\alpha \leftarrow 1, \mathbf{x} \leftarrow \mathbf{b}$ ;  
3  $x_{\max} \leftarrow \|\mathbf{x}\|_\infty$ ;  
4 for  $j \leftarrow n, n-1, \dots, 1$  do  
5    $\beta = \text{ProtectDivision}(x_j, t_{jj})$ ;  
6   if  $\beta \neq 1$  then  
7      $\mathbf{x} \leftarrow \beta\mathbf{x}$ ;  
8      $\alpha \leftarrow \beta\alpha$ ;  
9    $x_j \leftarrow x_j/t_{jj}$ ;  
10  if  $j > 1$  then  
11     $\beta = \text{ProtectUpdate}(x_{\max}, c_j, |x_j|)$ ;  
12    if  $\beta \neq 1$  then  
13       $\mathbf{x} \leftarrow \beta\mathbf{x}$ ;  
14       $\alpha \leftarrow \beta\alpha$ ;  
15     $\mathbf{x}_{1:j-1} \leftarrow \mathbf{x}_{1:j-1} - \mathbf{T}_{1:j-1,j}\mathbf{x}_j$ ;  
16     $x_{\max} \leftarrow \|\mathbf{x}_{1:j-1}\|_\infty$ ;  
17 return  $[\alpha, \mathbf{x}]$ ;
```

3 A blocked algorithm

Let $\mathbf{T} \in \mathbb{C}^{n \times n}$ be a non-singular upper triangular matrix, let $\mathbf{b} \in \mathbb{C}^n$ and partition the linear system $\mathbf{T}\mathbf{x} = \mathbf{b}$ conformally, i.e.,

$$\mathbf{T}\mathbf{x} = \begin{bmatrix} \mathbf{T}_{11} & \mathbf{T}_{12} & \dots & \mathbf{T}_{1N} \\ & \mathbf{T}_{22} & \dots & \mathbf{T}_{2N} \\ & & \ddots & \vdots \\ & & & \mathbf{T}_{NN} \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_N \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \vdots \\ \mathbf{b}_N \end{bmatrix} = \mathbf{b}. \quad (9)$$

This system can be solved using blocked backward substitution as in Alg. 5 `BlockedSolve`.

In order to obtain a robust algorithm we must replace the two kernels

$$f(\mathbf{T}, \mathbf{b}) = \mathbf{T}^{-1}\mathbf{b}, \quad g(\mathbf{y}, \mathbf{T}, \mathbf{x}) = \mathbf{y} - \mathbf{T}\mathbf{x}$$

with robust variants. To this end, we will associate with each block \mathbf{x}_i of the solution \mathbf{x}

Algorithm 5: $\mathbf{x} = \text{BlockedSolve}(\mathbf{T}, \mathbf{b})$

Data: A non-singular upper triangular matrix $\mathbf{T} \in \mathbb{C}^{n \times n}$ and $\mathbf{b} \in \mathbb{C}^n$ partitioned conformally as equation (9).

Result: The solution \mathbf{x} of $\mathbf{T}\mathbf{x} = \mathbf{b}$.

```
1  $\mathbf{x} \leftarrow \mathbf{b}$ ;  
2 for  $j = N, N - 1, \dots, 1$  do  
3    $\mathbf{x}_j \leftarrow f(\mathbf{T}_{jj}, \mathbf{x}_j)$ ;  
4   for  $i = 1, 2, \dots, j - 1$  do  
5      $\mathbf{x}_i \leftarrow g(\mathbf{x}_i, \mathbf{T}_{ij}, \mathbf{x}_j)$ ;  
6 return  $\mathbf{x}$ ;
```

a unique scaling factor $\alpha_i \in (0, 1]$ and all operations will be performed on such pairs. At this point, it is convenient to introduce the concept of an *augmented* vector.

Definition 1. An *augmented vector* $\langle \alpha, \mathbf{x} \rangle$ consists of a scalar $\alpha \in (0, 1]$ and a vector $\mathbf{x} \in \mathbb{C}^n$ and represents the vector $\mathbf{y} = \alpha^{-1}\mathbf{x}$. Two augmented vectors $\langle \alpha, \mathbf{x} \rangle$ and $\langle \beta, \mathbf{y} \rangle$ are said to be *equivalent* if they represent the same vector, i.e.,

$$\langle \alpha, \mathbf{x} \rangle \sim \langle \beta, \mathbf{y} \rangle \Leftrightarrow \alpha^{-1}\mathbf{x} = \beta^{-1}\mathbf{y}.$$

Two augmented vectors $\langle \alpha, \mathbf{x} \rangle$ and $\langle \beta, \mathbf{y} \rangle$ are said to be *consistently scaled* if $\alpha = \beta$.

3.1 Robust block solve

Let $\mathbf{T} \in \mathbb{C}^{n \times n}$ be a non-singular upper triangular matrix such that $|t_{ij}| \leq \Omega$ for $i < j$ and let $\langle \beta, \mathbf{b} \rangle$ be an augmented vector for which $|b_j| \leq \Omega$. Then we can use `RobustScalarSolve` to compute an augmented vector $\langle \gamma, \mathbf{x} \rangle$ such that

$$\mathbf{T}\mathbf{x} = \gamma\mathbf{b},$$

where the scaling factor γ ensures that \mathbf{x} can be computed without exceeding Ω . It follows that

$$\mathbf{T}\mathbf{x} = (\gamma\beta)\beta^{-1}\mathbf{b}$$

or equivalently

$$\mathbf{T}(\alpha^{-1}\mathbf{x}) = \beta^{-1}\mathbf{b}, \quad \alpha = \gamma\beta.$$

This procedure is formalized as Alg. 6.

3.2 Robust block update

Let $\langle \alpha, \mathbf{x} \rangle$ and $\langle \beta, \mathbf{y} \rangle$ be a pair of augmented vectors and let \mathbf{T} be a matrix such that $\mathbf{z} = \mathbf{y} - \mathbf{T}\mathbf{x}$ is defined. In this section we show how to compute an augmented vector $\langle \nu, \mathbf{z} \rangle$ such that

$$\nu^{-1}\mathbf{z} = \beta^{-1}\mathbf{y} - \mathbf{T}(\alpha^{-1}\mathbf{x}).$$

Moreover, the scaling ν will ensure that \mathbf{z} can be computed without overflow. There are two cases to consider, either $\alpha = \beta$ or $\alpha \neq \beta$.

Algorithm 6: $\langle \alpha, \mathbf{x} \rangle = \text{RobustBlockSolve}(\mathbf{T}, \langle \beta, \mathbf{b} \rangle)$

Data: A non-singular upper triangular matrix \mathbf{T} and an augmented vector $\langle \beta, \mathbf{b} \rangle$ such that

$$|t_{ij}| \leq \Omega, \quad \text{for } i < j \quad \text{and} \quad |b_j| \leq \Omega.$$

Result: An augmented vector $\langle \alpha, \mathbf{x} \rangle$ such that

$$\mathbf{T}(\alpha^{-1}\mathbf{x}) = \beta^{-1}\mathbf{b}.$$

Moreover, the computation of \mathbf{x} cannot exceed Ω .

```

1  $[\gamma, \mathbf{x}] \leftarrow \text{RobustScalarSolve}(\mathbf{T}, \mathbf{b});$ 
2  $\alpha \leftarrow \gamma\beta;$ 
3 return  $\langle \alpha, \mathbf{x} \rangle;$ 

```

If $\alpha = \beta$, then we can use Alg. 3 `ProtectUpdate` to compute a scaling ζ such that the scaled linear update

$$\mathbf{z} = \zeta\mathbf{y} - \mathbf{T}(\zeta\mathbf{x})$$

cannot exceed Ω . It follows immediately that

$$\nu^{-1}\mathbf{z} = \alpha^{-1}\mathbf{y} - \mathbf{T}(\alpha^{-1}\mathbf{x}),$$

provided that we choose $\nu = \alpha\zeta$.

If $\alpha \neq \beta$, then we replace $\langle \alpha, \mathbf{x} \rangle$ with $\langle \gamma, \mathbf{x}' \rangle$ and we replace $\langle \beta, \mathbf{y} \rangle$ with $\langle \gamma, \mathbf{y}' \rangle$ given by

$$\gamma = \min\{\alpha, \beta\}, \quad \mathbf{x}' = (\gamma/\alpha)\mathbf{x}, \quad \mathbf{y}' = (\gamma/\beta)\mathbf{y}. \quad (10)$$

It is straightforward to verify that

$$\langle \alpha, \mathbf{x} \rangle \sim \langle \gamma, \mathbf{x}' \rangle, \quad \langle \beta, \mathbf{y} \rangle \sim \langle \gamma, \mathbf{y}' \rangle.$$

Moreover, \mathbf{x}' and \mathbf{y}' can be computed without overflow, because $(\gamma/\alpha) \leq 1$ and $(\gamma/\beta) \leq 1$. Since $\langle \gamma, \mathbf{x}' \rangle$ and $\langle \gamma, \mathbf{y}' \rangle$ are consistently scaled, we can now use Alg. 3 `ProtectUpdate` to compute a scaling ζ such that the scaled linear update

$$\mathbf{z} = \zeta\mathbf{y}' - \mathbf{T}(\zeta\mathbf{x}')$$

cannot exceed Ω . From equation (10) it follows that

$$\mathbf{z} = \zeta(\gamma/\beta)\mathbf{y} - \mathbf{T}(\zeta(\gamma/\alpha)\mathbf{x})$$

or equivalently

$$\nu^{-1}\mathbf{z} = \beta^{-1}\mathbf{y} - \mathbf{T}(\alpha^{-1}\mathbf{x}),$$

provided that we choose $\nu = \zeta\gamma$ (as in the previous case).

The entire procedure is formalized as Alg. 7 `RobustBlockUpdate`.

3.3 Robust blocked backward substitution

Alg. 8 `RobustBlockedSolve` uses `RobustBlockSolve` and `RobustBlockUpdate` to solve an upper triangular system in a robust manner using augmented vectors.

Algorithm 7: $\langle \nu, \mathbf{z} \rangle = \text{RobustBlockUpdate}(\langle \beta, \mathbf{y} \rangle, \mathbf{T}, \langle \alpha, \mathbf{x} \rangle)$

Data: Augmented vectors $\langle \alpha, \mathbf{x} \rangle$, $\langle \beta, \mathbf{y} \rangle$ and a matrix \mathbf{T} such that the linear transformation

$$\mathbf{z} = \mathbf{y} - \mathbf{T}\mathbf{x}$$

is defined and

$$\|\mathbf{y}\|_\infty \leq \Omega, \quad \|\mathbf{T}\|_\infty \leq \Omega, \quad \|\mathbf{x}\|_\infty \leq \Omega.$$

Result: An augmented vector $\langle \nu, \mathbf{z} \rangle$ such that

$$\nu^{-1}\mathbf{z} = \beta^{-1}\mathbf{y} - \mathbf{T}(\alpha^{-1}\mathbf{x})$$

and the scaling ν ensures that the computation of \mathbf{z} cannot exceed Ω .

- 1 $\gamma = \min\{\alpha, \beta\}$;
 - 2 $\mathbf{x} \leftarrow (\gamma/\alpha)\mathbf{x}$;
 - 3 $\mathbf{y} \leftarrow (\gamma/\beta)\mathbf{y}$;
 - 4 $\zeta \leftarrow \text{ProtectUpdate}(\|\mathbf{y}\|_\infty, \|\mathbf{T}\|_\infty, \|\mathbf{x}\|_\infty)$;
 - 5 $\mathbf{z} \leftarrow \zeta\mathbf{y} - \mathbf{T}(\zeta\mathbf{x})$;
 - 6 $\nu = \zeta\gamma$;
 - 7 **return** $\langle \nu, \mathbf{z} \rangle$;
-

The correctness of the algorithm follows from the correctness of the two kernels. It is critical to examine the impact of the assumptions. Assumption (11) ensures that the diagonal blocks $\|\mathbf{T}_{jj}\|_\infty$ satisfy the requirements of Alg. 6 `RobustBlockedSolve`. Assumption (12) ensures that the super diagonal blocks \mathbf{T}_{ij} satisfy the requirements of Alg. 7 `RobustBlockUpdate`. Finally, assumption (13) ensures that the augmented vectors $\langle \alpha_i, x_i \rangle$ are initialized in a manner which satisfy the requirements of both kernels. The critical property, i.e., $\|\mathbf{x}_j\|_\infty \leq \Omega$, is then carried inductively through the execution of the algorithm.

4 A robust parallel algorithm

It is clear that Alg. 8 `RobustBlockedSolve` is virtually identical to Alg. 5 `BlockedSolve`. The robust algorithm has been obtained by replacing the two central kernels with robust variants operating on augmented vectors. Both algorithms can be executed in parallel using a run-time system such as `StarPU`. The overhead for the robust algorithm is marginally larger because each augmented vector contains a single extra element (the scaling) which has to be communicated. Moreover, the robust algorithm requires one final exchange of information in order to obtain a consistent scaling for all the augmented vectors. Otherwise, the number of messages exchanged is the same for the two algorithms. If there are multiple right hand sides, then consistent scalings can be obtained using a single exchange of information.

5 Conclusion

We have carefully derived scalar and blocked algorithms for solving upper triangular linear systems in a robust manner. The robust blocked algorithm can be executed in parallel

Algorithm 8: $\langle \alpha, \mathbf{x} \rangle = \text{RobustBlockedSolve}(\mathbf{T}, \mathbf{b})$

Data: A non-singular upper triangular matrix $\mathbf{T} \in \mathbb{C}^{n \times n}$ and $\mathbf{b} \in \mathbb{C}^n$ partitioned conformally as equation (9) and subject to the following constraints:

1. The matrix: All super diagonal elements t_{ij} must satisfy

$$|t_{ij}| \leq \Omega \quad (11)$$

and all super diagonal blocks must satisfy

$$\|\mathbf{T}_{ij}\|_\infty \leq \Omega. \quad (12)$$

2. The right hand side: All elements must satisfy

$$|b_j| \leq \Omega. \quad (13)$$

Result: An augmented vector $\langle \alpha, \mathbf{x} \rangle$ such that

$$\mathbf{T}\mathbf{x} = \alpha\mathbf{b}.$$

Moreover, the scaling ensures that calculation of \mathbf{x} never exceeds Ω .

```
1 for  $i = 1, \dots, N$  do
2    $\langle \alpha_i, \mathbf{x}_i \rangle \leftarrow \langle 1, \mathbf{b}_i \rangle$ ;
3 for  $j = N, N - 1, \dots, 1$  do
4    $\langle \alpha_j, \mathbf{x}_j \rangle \leftarrow \text{RobustBlockSolve}(\mathbf{T}_{jj}, \langle \alpha_j, \mathbf{x}_j \rangle)$ ;
5   for  $i = 1, 2, \dots, j - 1$  do
6      $\langle \alpha_i, \mathbf{x}_i \rangle \leftarrow \text{RobustBlockUpdate}(\langle \alpha_i, \mathbf{x}_i \rangle, \mathbf{T}_{ij}, \langle \alpha_j, \mathbf{x}_j \rangle)$ 
7  $\alpha = \min\{\alpha_1, \alpha_2, \dots, \alpha_N\}$ ;
8 for  $i = 1, 2, \dots, N$  do
9    $\mathbf{x}_i \leftarrow (\alpha/\alpha_i)\mathbf{x}_i$ ;
10 return  $\langle \alpha, \mathbf{x} \rangle$ ;
```

using a run-time system such as StarPU with a parallel overhead which is marginally larger than the standard algorithm. The cost of obtaining a final consistent scaling can be amortized over multiple right hand sides.

Acknowledgment

This work is part of a project that has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671633. The authors are very grateful to Sven Hammarling (NAG) because he carefully read the manuscript and made several valuable suggestions and comments.

References

- [1] Edward Anderson. Robust Triangular Solves for Use in Condition Estimation. LAWN 36, Cray Research Inc., August 1991.

- [2] Cédric Augonnet, Samuel Thibault, Raymond Namyst, and Pierre-André Wacrenier. StarPU: A Unified Platform for Task Scheduling on Heterogeneous Multicore Architectures. *Concurrency and Computation: Practice and Experience, Special Issue: Euro-Par 2009*, 23:187–198, February 2011.
- [3] Mark R. Fahey. New Complex Parallel Eigenvalue and Eigenvector Routines. LAWN 153, Computational Migration Group, August 2001.