


Parallel robust solution of triangular equations

Carl Christian Kjelgaard Mikkelsen
spock@cs.umu.se

Mirko Myllykoski, Angelika Schwarz, Lars Karlsson

Department of Computing Science
Umeå University

SIAM CSE
Virtual conference
March 1-5, 2021



What is a robust algorithm?

All numbers are in the **representable** range

Example 1: $\mathbf{Ax} = \mathbf{b}$

$$\mathbf{Ax} := \begin{bmatrix} 1 & & & & \\ -1 & 1 & & & \\ -1 & -1 & 1 & & \\ -1 & -1 & -1 & 1 & \\ -1 & -1 & -1 & -1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 4 \\ 8 \\ 16 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} =: \mathbf{b}$$

Example 2: $SX = X\Lambda$

Matrix:

$$S = \begin{bmatrix} 1 & & & & \\ -5 & 2 & & & \\ -5 & -5 & 3 & & \\ -5 & -5 & -5 & 4 & \\ -5 & -5 & -5 & -5 & 5 \end{bmatrix}$$

Eigenvectors:

$$X = \begin{bmatrix} 1 & & & & \\ 5 & 1 & & & \\ 15 & 5 & 1 & & \\ 35 & 15 & 5 & 1 & \\ 70 & 35 & 15 & 5 & 1 \end{bmatrix}$$

Example 3: $\mathbf{AX} + \mathbf{XA}^T = \mathbf{bb}^T$

Matrices:

$$\mathbf{A} = \begin{bmatrix} \frac{1}{2} & & & & \\ -1 & \frac{1}{2} & & & \\ -1 & -1 & \frac{1}{2} & & \\ -1 & -1 & -1 & \frac{1}{2} & \\ -1 & -1 & -1 & -1 & \frac{1}{2} \end{bmatrix}, \quad \mathbf{b} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}$$

Solution:

$$\mathbf{X} = \begin{bmatrix} 1 & 2 & 4 & 8 & 16 \\ 2 & 5 & 12 & 28 & 64 \\ 4 & 12 & 33 & 86 & 216 \\ 8 & 28 & 86 & 245 & 664 \\ 16 & 64 & 216 & 664 & 1921 \end{bmatrix}$$

Robust triangular solvers in LAPACK

| Subroutine | Problem |
|--------------|---|
| <hr/> xlatrs | <hr/> $\mathbf{Ax} = \alpha \mathbf{b}$ <hr/> |
| xtrevc | $\mathbf{AX} = \mathbf{X}\Lambda$ |
| xtgevc | $\mathbf{AX} = \mathbf{BX}\Lambda$ |
| <hr/> xtrsyl | <hr/> $\mathbf{AX} + \mathbf{XB} = \alpha \mathbf{C}$ |

All routines are [scalar](#), [sequential](#) codes descended from xlatrs.

Robust triangular solvers in ScaLAPACK?

From PDLATTRS¹ (parallel dlatrs)

```
***** NO SCALING ***** Call PDTRSV for all cases  
*****
```

From PZLATTRS² (parallel zlatrs)

This is very slow relative to PZTRSV. This should only be used when scaling is necessary to control overflow, or when it is modified to scale better.

¹http://www.netlib.org/scalapack/explore-html/d6/da4/pdlatrs_8f_source.html, line 62. Retrieved March 3rd 2021

²http://www.netlib.org/scalapack/explore-html/df/d88/pzlattrs_8f_source.html, lines 36:38. Retrieved: March 3rd 2021

Regular backward substitution, $\mathbf{T}\mathbf{x} = \mathbf{b}$

Algorithm 1: $[\alpha, \mathbf{x}] = \text{BackSub}(\mathbf{T}, \mathbf{b})$

```
1 for  $j \leftarrow n, n-1, \dots, 1$  do  
2    $x_j \leftarrow x_j / t_{jj};$   
3   if  $j > 1$  then  
4      $\mathbf{x}_{1:j-1} \leftarrow \mathbf{x}_{1:j-1} - \mathbf{T}_{1:j-1,j}x_j;$   
5 return  $\mathbf{x};$ 
```

Robust backward substitution, $\mathbf{T}\mathbf{x} = \alpha\mathbf{b}$

Algorithm 2: $[\alpha, \mathbf{x}] = \text{RobustBackSub}(\mathbf{T}, \mathbf{b})$

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

ProtectDivision

The division $y \leftarrow b/t$ is safe if

$$|b| \leq |t|\Omega, \quad \Omega = \text{largest normal FP number.} \quad (1)$$

ProtectDivision

The division $y \leftarrow b/t$ is safe if

$$|b| \leq |t|\Omega, \quad \Omega = \text{largest normal FP number.} \quad (1)$$

Problem: Find scaling $\alpha \in (0, 1]$ such that

$$|\alpha b| \leq |t|\Omega. \quad (2)$$

ProtectDivision

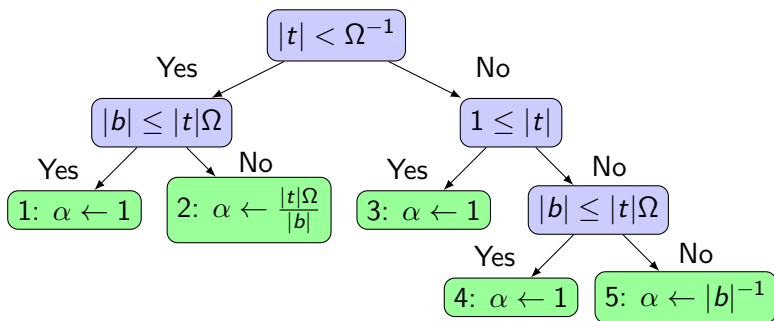
The division $y \leftarrow b/t$ is safe if

$$|b| \leq |t|\Omega, \quad \Omega = \text{largest normal FP number.} \quad (1)$$

Problem: Find scaling $\alpha \in (0, 1]$ such that

$$|\alpha b| \leq |t|\Omega. \quad (2)$$

Solution:



ProtectUpdate

The linear update $\mathbf{y} \leftarrow \mathbf{y} - \mathbf{T}\mathbf{x}$ is safe if

$$\|\mathbf{y}\|_{\infty} + \|\mathbf{T}\|_{\infty}\|\mathbf{x}\|_{\infty} \leq \Omega, \quad \Omega = \text{largest normal FP number.} \quad (3)$$

ProtectUpdate

The linear update $\mathbf{y} \leftarrow \mathbf{y} - \mathbf{T}\mathbf{x}$ is safe if

$$\|\mathbf{y}\|_{\infty} + \|\mathbf{T}\|_{\infty}\|\mathbf{x}\|_{\infty} \leq \Omega, \quad \Omega = \text{largest normal FP number.} \quad (3)$$

Problem: Find $\zeta \in (0, 1]$ such that

$$\|\zeta\mathbf{y}\|_{\infty} + \|\mathbf{T}\|_{\infty}\|\zeta\mathbf{x}\|_{\infty} \leq \Omega. \quad (4)$$

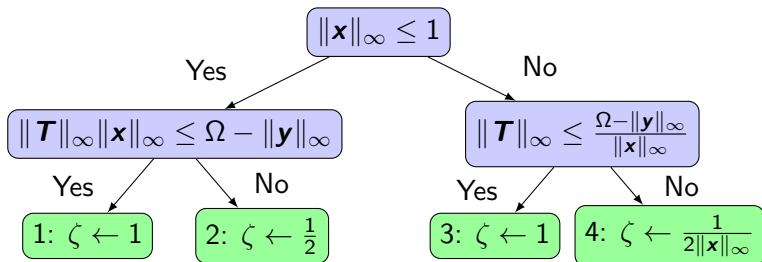
ProtectUpdate

The linear update $\mathbf{y} \leftarrow \mathbf{y} - \mathbf{T}\mathbf{x}$ is safe if

$$\|\mathbf{y}\|_\infty + \|\mathbf{T}\|_\infty \|\mathbf{x}\|_\infty \leq \Omega, \quad \Omega = \text{largest normal FP number.} \quad (3)$$

Problem: Find $\zeta \in (0, 1]$ such that

$$\|\zeta\mathbf{y}\|_\infty + \|\mathbf{T}\|_\infty \|\zeta\mathbf{x}\|_\infty \leq \Omega. \quad (4)$$



How to develop blocked, parallel and robust codes?

How to develop blocked, parallel and robust codes?

Goal: No significant difference between robust and non-robust codes!

How to develop blocked, parallel and robust codes?

Goal: No significant difference between robust and non-robust codes!

Key: Different processing units can use different units of measurement.

Step 1: Partition the problem

$$\begin{aligned} \mathbf{TX} &:= \begin{bmatrix} \mathbf{T}_{11} & \mathbf{T}_{12} & \dots & \mathbf{T}_{1M} \\ & \mathbf{T}_{22} & \dots & \mathbf{T}_{2m} \\ & & \ddots & \vdots \\ & & & \mathbf{T}_{MM} \end{bmatrix} \begin{bmatrix} \mathbf{X}_{11} & \mathbf{X}_{12} & \dots & \mathbf{X}_{1N} \\ \mathbf{X}_{21} & \ddots & \ddots & \mathbf{X}_{2N} \\ \vdots & \ddots & \ddots & \vdots \\ \mathbf{X}_{M1} & \dots & \dots & \mathbf{X}_{MN} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} & \dots & \mathbf{B}_{1N} \\ \mathbf{B}_{21} & \ddots & \ddots & \mathbf{B}_{2N} \\ \vdots & \ddots & \ddots & \vdots \\ \mathbf{B}_{M1} & \dots & \dots & \mathbf{B}_{MN} \end{bmatrix} =: \mathbf{B} \end{aligned}$$

Step 2: Augment each block of data

A separate scaling factor α_{ijk} for each column of every block \mathbf{X}_{ij} .

Step 2: Augment each block of data

A separate scaling factor α_{ijk} for each column of every block \mathbf{X}_{ij} .

An augmented block

$$\langle \boldsymbol{\alpha}, \mathbf{X} \rangle$$

where

$$\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_n), \quad \alpha_j \in (0, 1], \quad \mathbf{X} \in \mathbb{R}^{m \times n}$$

is used to represents the real matrix

$$\bar{\mathbf{X}} \in \mathbb{R}^{m \times n}$$

where

$$\bar{\mathbf{x}}_j = \alpha_j^{-1} \mathbf{x}_j, \quad j = 1, 2, \dots, n.$$

Step 3: Replace kernels with robust kernels ...

Step 3: Replace kernels with robust kernels ...

Problem: Given $\langle \alpha, \mathbf{X} \rangle$ and $\langle \beta, \mathbf{Y} \rangle$ find $\langle \nu, \mathbf{Z} \rangle$ such that

$$\bar{\mathbf{Z}} = \bar{\mathbf{Y}} - \mathbf{T}\bar{\mathbf{X}}$$

or equivalently

$$\nu_j^{-1} \mathbf{z}_j = \beta_j^{-1} \mathbf{y}_j - \mathbf{T} \left(\alpha_j^{-1} \mathbf{x}_j \right), \quad j = 1, 2, \dots, n$$

Step 3: Replace kernels with robust kernels ...

Problem: Given $\langle \alpha, \mathbf{X} \rangle$ and $\langle \beta, \mathbf{Y} \rangle$ find $\langle \nu, \mathbf{Z} \rangle$ such that

$$\bar{\mathbf{Z}} = \bar{\mathbf{Y}} - \mathbf{T}\bar{\mathbf{X}}$$

or equivalently

$$\nu_j^{-1} \mathbf{z}_j = \beta_j^{-1} \mathbf{y}_j - \mathbf{T} \left(\alpha_j^{-1} \mathbf{x}_j \right), \quad j = 1, 2, \dots, n$$

Algorithm 3: $\langle \nu, \mathbf{Z} \rangle \leftarrow \text{RobustGEMM}(\langle \beta, \mathbf{Y} \rangle, \mathbf{T}, \langle \alpha, \mathbf{X} \rangle)$

- 1 **for** $j \leftarrow 1, 2, \dots, n$ **do**
- 2 $\gamma_j \leftarrow \min\{\alpha_j, \beta_j\};$
- 3 $\zeta_j \leftarrow \text{ProtectUpdate}((\gamma_j/\beta_j)\|\mathbf{y}_j\|_\infty, \|\mathbf{T}\|_\infty, (\gamma_j/\alpha_j)\|\mathbf{x}_j\|_\infty);$
- 4 $\nu_j \leftarrow \zeta_j \gamma_j;$
- 5 $\mathbf{Z} \leftarrow \left[\mathbf{Y} \text{diag}(\nu \odot \beta) \right] - \mathbf{T} \left[\mathbf{X} \text{diag}(\nu \odot \alpha) \right];$
- 6 **Return** $\langle \nu, \mathbf{Z} \rangle;$

Cost: $O(n^3)$. Overhead: $O(n^2) + O(n^2)$

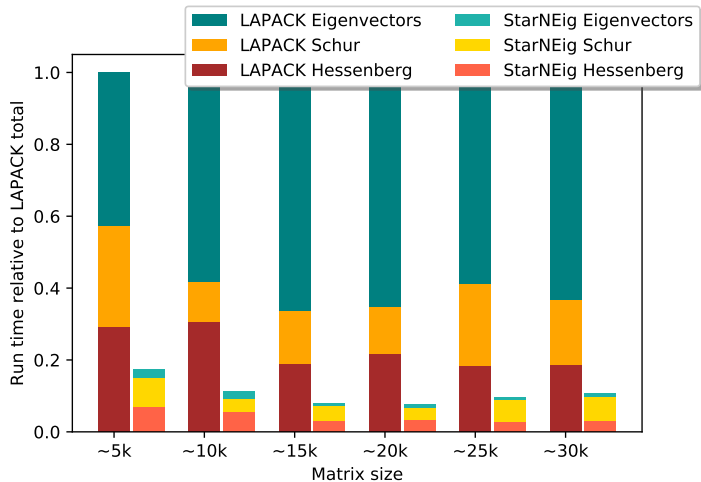
Robust solution of partitioned system $TX = B$.

Robust solution of partitioned system $TX = B$.

Algorithm 4: $\langle \alpha, X \rangle \leftarrow \text{RobustBlockedSolveMRHS}(T, B)$

```
1 for  $(i, j) \in \{1, 2, \dots, M\} \times \{1, 2, \dots, N\}$  do
2    $\langle \alpha_{ij}, X_{ij} \rangle \leftarrow \langle e, B_{ij} \rangle;$ 
3 for  $j \leftarrow M, M - 1, \dots, 1$  do
4   for  $k \leftarrow 1, 2, \dots, N$  do
5      $\langle \alpha_{jk}, X_{jk} \rangle \leftarrow \text{RobustTRSM}(T_{jj}, \langle \alpha_{jk}, X_{jk} \rangle);$ 
6     for  $i \leftarrow 1, 2, \dots, j - 1$  do
7       for  $k \leftarrow 1, 2, \dots, N$  do
8          $\langle \alpha_{ik}, X_{ik} \rangle \leftarrow \text{RobustGEMM}(\langle \alpha_{ik}, X_{ik} \rangle, T_{ij}, \langle \alpha_{jk}, X_{jk} \rangle);$ 
9 for  $k \leftarrow 1, 2, \dots, N$  do
10   $\alpha_k \leftarrow \min_i \{ \alpha_{ik} \};$ 
11  for  $i \leftarrow 1, 2, \dots, M$  do
12     $X_{ik} \leftarrow X_{ik} \text{diag}(\alpha_k \oslash \alpha_{ik});$ 
13  $\alpha \leftarrow [\alpha_1 \quad \alpha_2 \quad \dots \quad \alpha_N];$ 
14 Return  $\langle \alpha, X \rangle;$ 
```

LAPACK versus StarNEig



New Software

StarNEig

1. MS227 talk by Dr. Mirko Myllykoski.
2. A library for solving nonsymmetric eigenvalue problems:

$$\mathbf{Ax} = \lambda \mathbf{x}$$

$$\mathbf{Ax} = \lambda \mathbf{Bx}$$

3. Uses the StarPU runtime system from INRIA.
4. Download

<https://nlafet.github.io/StarNEig/>

Contact: mirkom@cs.umu.se, spock@cs.umu.se

Backup slides

Caution: Underflow *can* still destroy the result!

Ω = largest normal number, ϵ = smallest subnormal number

$$\mathbf{T}_x := \left[\begin{array}{c|cccc} \frac{1}{2} & & & & \\ \hline & 1 & & & \\ & -2 & 1 & & \\ & & \ddots & \ddots & \\ & & & -2 & 1 \end{array} \right] \left[\begin{array}{c} 2\Omega \\ \epsilon \\ 2^1\epsilon \\ \vdots \\ 2^{m-1}\epsilon \end{array} \right] = \left[\begin{array}{c} \Omega \\ \epsilon \\ 0 \\ \vdots \\ 0 \end{array} \right] =: \mathbf{b}.$$

Scaling flushes ϵ to zero and destroys the solution

Performance degradation of robust solvers?




Figure: Triangular linear system $m = 40000$ and $n = 2000$

Bibliography I

- [1] Carl Christian Kjelgaard Mikkelsen, Angelika Beatrix Schwarz, and Lars Karlsson.
Parallel robust solution of triangular linear systems.
Concurrency and Computation: Practice and Experience,
31(19):e5064, 2019.
e5064 cpe.5064.
- [2] Carl Christian Kjelgaard Mikkelsen.
Well-conditioned eigenvalue problems that overflow.
<https://arxiv.org/abs/2005.05356>, 2020.
- [3] Carl Christian Kjelgaard Mikkelsen and Mirko Myllykoski.
Parallel robust computation of generalized eigenvectors of
matrix pencils.
In Roman et. al. Wyrzykowski, editor, *PPAM*, number 12043 in
LNCS, pages 58–69, Cham, 2020. Springer.

Bibliography II

- [4] Mirko Myllykoski and Carl Christian Kjelgaard Mikkelsen. Introduction to StarNEig—A Task-Based Library for Solving Nonsymmetric Eigenvalue Problems. In Roman et. al Wyrzykowski, editor, *PPAM*, number 12043 in LNCS, pages 70–81, Cham, 2020. Springer.
- [5] Mirko Myllykoski and Carl Christian Kjelgaard Mikkelsen. Task-based, GPU-accelerated and robust library for solving dense nonsymmetric eigenvalue problems. *Concurrency and Computation: Practice and Experience*, page e5915, 2020.
- [6] Angelika Schwarz and Carl Christian Kjelgaard Mikkelsen. Robust Task-Parallel Solution of the Triangular Sylvester Equation. In Roman et. al. Wyrzykowski, editor, *PPAM*, number 12043 in LNCS, pages 82–92, Cham, 2020. Springer.

Bibliography III

- [7] Angelika Schwarz, Carl Christian Kjelgaard Mikkelsen, and Lars Karlsson.

Robust parallel eigenvector computation for the non-symmetric eigenvalue problem.

Parallel Computing, 100:102707, 2020.