# A componentwise accurate shift technique for M-matrix algebraic Riccati equations 

E. Addis, B. Iannazzo, F. Poloni

ILAS Conference, Madrid, June 2023

## Plan

- What is a fluid queue?
- What is doubling? (ADDA)
- What is componentwise accuracy?
- What is the shift technique?
- How do we combine them all together?


## Fluid queues

Fluid queue: "infinite-size bucket" in which the fluid level $X(t)$ changes with a rate which depends on the state $\varphi(t)$ of a continuous-time Markov chain with generator matrix $Q$.

[Moran '54, Mitra '88, Kulkarni '97, Ahn-Ramaswami '03, Bean-O'Reilly-Taylor '05, etc.] In this talk: rates $\pm 1$; states $S=S_{+} \cup S_{-}$. $\Psi_{i j}=P\left[\right.$ first return to $X(t)=0$ in state $\left.\varphi(t)=j \in S_{-} \mid \varphi(0)=i \in S_{+}\right]$.

## M-matrix algebraic Riccati equation

M-matrix algebraic Riccati equation (MARE):

$$
Q_{+-}+Q_{++} \Psi+\Psi Q_{--}+\Psi Q_{-+} \Psi=0
$$

The matrix

$$
M=-Q=-\left[\begin{array}{ll}
Q_{--} & Q_{-+} \\
Q_{+-} & Q_{++}
\end{array}\right]
$$

is an irreducible, singular M-matrix.
$Q \mathbf{1}=0, \boldsymbol{\pi} Q=0$ for a row vector $\boldsymbol{\pi}>0$.

## Eigenvalues

The most important thing: eigenvalues.
Solving the MARE $\leftrightarrow$ finding an invariant subspace:

$$
\underbrace{\left[\begin{array}{cc}
-Q_{--} & -Q_{-+} \\
Q_{+-} & Q_{++}
\end{array}\right]}_{:=\mathcal{H}}\left[\begin{array}{c}
I_{n_{-}} \\
\Psi
\end{array}\right]=\left[\begin{array}{c}
I_{n_{-}} \\
\Psi
\end{array}\right] \underbrace{\left(-Q_{--}-Q_{-+} \Psi\right)}_{:=-U}
$$

$$
\Lambda(\mathcal{H})=\{\underbrace{\lambda_{1}, \lambda_{2}, \ldots, \lambda_{n_{+}}}_{=\Lambda(V) \subset \text { left half-plane }}, \underbrace{\lambda_{n_{+}+1}, \lambda_{n_{+}+2}, \ldots, \lambda_{n_{+}+n_{-}}}_{=\Lambda(-U) \subset \text { right half-plane }}\} .
$$

For simplicity we assume $\lambda_{n_{+}+1}=0$ : recurrent process.
Critical case: $\lambda_{n_{+}} \approx 0$. Algorithms slow down.

## Spectrum of $\mathcal{H}$ - figure



Figure: $\Lambda(\mathcal{H})$ in the critical case where $\lambda_{n_{+}}=\lambda_{n_{+}+1}=0$

Alternating-direction doubling algorithm [Wang et al. 2012]
(1) Choose rates

$$
\alpha \geq \alpha_{\mathrm{opt}}=\max _{i}\left|\left(Q_{++}\right)_{i i}\right|, \quad \beta \geq \beta_{\mathrm{opt}}=\max _{i}\left|\left(Q_{--}\right)_{i i}\right| .
$$

(2) Compute initial values

$$
\left[\begin{array}{cc}
E_{0} & G_{0} \\
H_{0} & F_{0}
\end{array}\right]=\underbrace{-\left[\begin{array}{cc}
Q_{--}-\alpha I & Q_{-+} \\
Q_{+-} & Q_{++}-\beta I
\end{array}\right]^{-1}}_{\text {M-matrix }} \underbrace{\left[\begin{array}{cc}
Q_{--}+\beta I & Q_{-+} \\
Q_{+-} & Q_{++}+\alpha I
\end{array}\right]}_{\text {non-negative }}
$$

(3) Iterate

$$
\begin{aligned}
& P_{k}=\left[\begin{array}{ll}
E_{k} & G_{k} \\
H_{k} & F_{k}
\end{array}\right] \mapsto P_{k+1}=\left[\begin{array}{ll}
E_{k+1} & G_{k+1} \\
H_{k+1} & F_{k+1}
\end{array}\right], \\
& E_{k+1}=E_{k}\left(I-G_{k} H_{k}\right)^{-1} E_{k}, \\
& F_{k+1}=F\left(I-H_{k} G_{k}\right)^{-1} F_{k}, \\
& G_{k+1}=G_{k}+E_{k}\left(I-G_{k} H_{k}\right)^{-1} G_{k} F_{k} \\
& H_{k+1}=H_{k}+F_{k}\left(I-H_{k} G_{k}\right)^{-1} H_{k} E_{k}
\end{aligned}
$$

## What is ADDA?

The algebra constructing a factorization

$$
f(\mathcal{H})^{2^{k}}=\left[\begin{array}{cc}
l & -G_{k} \\
0 & F_{k}
\end{array}\right]^{-1}\left[\begin{array}{cc}
E_{k} & 0 \\
-H_{k} & I
\end{array}\right], \quad f(z)=\frac{z-\beta}{z+\alpha} .
$$

The interpretation Constructing a "level-crossing" QBD associated to the queue

$$
A_{-1}=\left[\begin{array}{cc}
0 & 0 \\
0 & F_{0}
\end{array}\right], \quad A_{0}=\left[\begin{array}{cc}
0 & G_{0} \\
H_{0} & 0
\end{array}\right], \quad A_{1}=\left[\begin{array}{cc}
E_{0} & 0 \\
0 & 0
\end{array}\right],
$$

and applying Cyclic Reduction to it. [Bean-Nguyen-P. '18]
The practice Fastest iteration in literature for this problem.
Added benefit: componentwise accuracy.

## Componentwise accuracy

An algorithm to compute a matrix/vector quantity $\Psi$ is componentwise accurate if all entries, even tiny ones, are computed with small forward error

$$
\frac{\left|\Psi_{i j}^{\text {computed }}-\Psi_{i j}\right|}{\Psi_{i j}} \approx \mathrm{u}
$$

Very strict requirements:

- High condition numbers prevent forward stable computations.
- Often errors are proportional to $\|\Psi\|$ or $\max _{i j} \Psi_{i j}$. E.g., when computing

$$
\Psi=\left[\begin{array}{lll}
0.5 & 0.499999 & 0.000001
\end{array}\right]
$$

on the last component we expect $\approx 10$ correct digits, not $\approx 16$.

## Keys to componentwise accuracy

## Triplet representations

Given an M-matrix $M$, we can compute $M^{-1}$ with componentwise accuracy if we know (in addition to the entries of $M$ ) vectors $\mathbf{v}>0, \mathbf{w} \geq 0$ such that $M \mathbf{v}=\mathbf{w}$.

GTH-like algorithm Modified Gaussian elimination, using the relation $M \mathbf{v}=\mathbf{w}$ to evaluate pivots more accurately. [Alfa, Xue, Ye '01]

Example In the ADDA initial values, we need $M_{\alpha, \beta}^{-1}$, where

$$
M_{\alpha, \beta}=-\left[\begin{array}{cc}
Q_{--}-\alpha I & Q_{-+} \\
Q_{+-} & Q_{++}-\beta I
\end{array}\right]
$$

Solution Since $Q \mathbf{1}=\mathbf{0}$, we know that the M -matrix satisfies exactly $M_{\alpha, \beta} \mathbf{1}=\left[\begin{array}{c}\alpha 1 \\ \beta 1\end{array}\right]$.

## Keys to componentwise accuracy

More generally:
No Inaccurate Cancellation (NIC) principle [Demmel-Dumitriu-Holtz-Koev '08]
To achieve componentwise accuracy, we must avoid subtractions $a-b$ with $a \approx b$.

Example In the ADDA initial values, we need

$$
Q_{++}+\alpha I, \quad \text { where } \alpha \geq \alpha_{\text {opt }}=\max _{i}\left|\left(Q_{++}\right)_{i i}\right| .
$$

Choosing $\alpha=\alpha_{\text {opt }}$ may lead to trouble if e.g. $\operatorname{diag}\left(Q_{++}\right)=\left[\begin{array}{c}-1 \\ -0.99999\end{array}\right]$.
Solution choose $\alpha$ a bit larger, for instance $\alpha=1.25 \alpha_{\text {opt }}$. Convergence speed is slightly degraded, but we gain in accuracy.

Using these techniques, one can make ADDA componentwise accurate. [Xue-Xu-Li '12, Nguyen-P '15, Xue-Li '17]

## Spectrum of $f(\mathcal{H})$

The most important thing: eigenvalues.
Recall that $f(z)=\frac{z-\beta}{z+\alpha}$. In ADDA we work with $f(\mathcal{H})$, with eigenvalues

- $\Lambda(f(-U))=\left\{f\left(\lambda_{n_{+}+1}\right), f\left(\lambda_{n_{+}+2}\right), \ldots, f\left(\lambda_{n_{+}+n_{-}}\right)\right\} \subset\left\{|z| \leq \frac{\beta}{\alpha}\right\}$ and
- $\Lambda(f(V))=\left\{f\left(\lambda_{1}\right), f\left(\lambda_{2}\right), \ldots, f\left(\lambda_{n_{+}}\right)\right\} \subset\left\{|z| \geq \frac{\beta}{\alpha}\right\}$.

ADDA convergence depends on the parameter

$$
\xi=\rho(f(-U)) \rho\left(f(V)^{-1}\right)=\left|\frac{f\left(\lambda_{n_{+}+1}\right)}{f\left(\lambda_{n_{+}}\right)}\right| \leq 1:
$$

- If $\xi<1,\left\|H_{k}-\Psi\right\| \sim \xi^{2^{k}}$.
- If $\xi=1$ (null recurrent queue), $\left\|H_{k}-\Psi\right\| \sim 2^{-k}$.


## Problem

Slow convergence when $\xi=1$ or $\xi \approx 1$ : how to speed it up?

## Spectrum of $f(\mathcal{H})$ - figure



Figure: $\Lambda(f(\mathcal{H}))$ in the critical case where $\lambda_{n_{+}}=\lambda_{n_{+}+1}=0$

## Shift: from $\mathcal{H}$ to $\hat{\mathcal{H}}$ [Guo-lannazzo-Meini '07]

Shift technique: rank-1 modification to accelerate convergence in (near-)critical cases.
We choose $\eta>0$ and $\mathbf{p} \geq 0$ such that $\mathbf{p}^{\top} \mathbf{1}=1$.

$$
\begin{array}{c|c}
\mathcal{H} & \hat{\mathcal{H}}=\mathcal{H}+\eta \mathbf{1 p}^{T} \\
Q & \hat{Q}=Q-\eta\left[\begin{array}{c}
\mathbf{1}_{n_{-}} \\
-\mathbf{1}_{n_{+}}
\end{array}\right] \mathbf{p}^{T} \\
Q_{+-}+Q_{++} \Psi+ & \hat{Q}_{+-}+\hat{Q}_{++} \Psi+ \\
\Psi Q_{--}+\Psi Q_{-+} \Psi=0 & \Psi \hat{Q}_{--}+\Psi \hat{Q}_{-+} \Psi=0 \\
\Psi & \text { same } \Psi \\
\lambda_{n_{+}+1}=0 & \hat{\lambda}_{n_{+}+1}=\eta \\
\Lambda(-U)=\left\{0, \lambda_{n_{+}+1}, \ldots, \lambda_{n_{+}+n_{-}}\right\} & \Lambda(-\hat{U})=\left\{\eta, \lambda_{n++1}, \ldots, \lambda_{n_{+}+n_{-}}\right\} \\
\xi & \hat{\xi}<\xi \text { (at least when } \eta<\beta) .
\end{array}
$$

The technique may decrease the number of steps dramatically.

## Componentwise accurate construction of $\hat{P}_{0}$

In this talk: Can we combine the two improvements? Shift technique and componentwise accurate computations?

$$
\hat{Q}=Q-\eta\left[\begin{array}{c}
\mathbf{1}_{n_{-}} \\
-\mathbf{1}_{n_{+}}
\end{array}\right] \mathbf{p}^{T}
$$

Problem: The sign properties may be lost, even for tiny values of $\eta$ !

## Delayed shift

Idea: When $\alpha, \beta$ are fixed, $\hat{P}_{0}$ is a rank- 1 modification of $P$ :

$$
\hat{P}_{0}=P_{0}-\eta \underbrace{(\alpha+\beta) \frac{\mathbf{u p}^{T} M_{\alpha, \beta}^{-1}}{1+\eta \mathbf{p}^{T} \mathbf{u}}}_{=: \Sigma_{\eta}}, \quad \mathbf{u}=M_{\alpha, \beta}^{-1}\left[\begin{array}{c}
\mathbf{1}_{n_{-}}  \tag{*}\\
-\mathbf{1}_{n_{+}}
\end{array}\right]
$$

We can first compute $P_{0}>0$, then construct $\hat{P}_{0}$ by subtraction (delayed shift).
$P_{0}-\Sigma_{\eta}$ contains subtractions, but we can compute all the quantities in $\left(^{*}\right)$ and then choose $\eta$ afterwards to satisfy two entrywise conditions:

- $\hat{P}_{0}>0$
- No Inaccurate Cancellation in $P_{0}-\Sigma_{\eta}$.

The range of allowed values for $\eta$ is often much larger than when applying the regular "non-delayed" shift.

## The missing steps

- Triplet representations: can be computed from the relation

$$
\left[\begin{array}{cc}
\tau^{2^{k}} \hat{E}_{k} & \hat{G}_{k}  \tag{T}\\
\hat{H}_{k} & \tau^{-2^{k}} \hat{F}_{k}
\end{array}\right] \mathbf{1}=\mathbf{1}, \quad \text { with } \tau=\frac{\alpha+\eta}{\beta-\eta} .
$$

- Positivity, applicability, convergence: can be proved by mimicking the original proofs in ADDA: the only assumptions they need are $\hat{P}_{0} \geq 0$ and ( T ).
- Forward error bound: can be obtained, though worse than in the non-shifted case:

$$
|\operatorname{computed}(\Sigma)-\Sigma| \leq \underbrace{\mathcal{O}\left(n^{3} \mathbf{u}\right)}_{\text {machine prec. }} \underbrace{\frac{1+\mathbf{p}^{T} M_{\alpha \beta}^{-1} \mathbf{1}}{\left(1+\mathbf{p}^{T} \mathbf{u}\right)^{2}}(\alpha+\beta) M_{\alpha \beta}^{-1} \mathbf{1} \mathbf{p}^{T} M_{\alpha \beta}^{-1}}_{\text {not } \Sigma \text {, possibly larger }}
$$

## Example 1 [Nguyen-P '15, Example 5.1]

An example with $Q$ with imbalanced entries.

$$
\begin{aligned}
\Lambda(\mathcal{H}) & =\{-20.0000,-1.5625,-0.0100,0.0000,2.5575,19.9800\} \\
\Lambda(f(\mathcal{H})) & =\{-6.9970,-1.2314,-1.0003,-0.9990,-0.7078,0.1428\}
\end{aligned}
$$

- without shift: convergence rate $\xi=0.9990$.
- non-accurate shift: $\hat{\xi}=0.7078$.

Optimal $\eta=\beta=14.9850, \mathbf{p}=\frac{1}{n} \mathbf{1}$.

- accurate shift, imposing $\hat{P}_{0} \geq 0: \hat{\xi}=0.8575$.

$$
\eta=1.1429, f(\eta)=-0.8575, \mathbf{p}=\mathbf{e}_{5} .
$$

- accurate shift, imposing NIC in $\hat{P}_{0}=P_{0}-\Sigma_{\eta}: \hat{\xi}=0.9777$.

$$
\eta=0.1614, f(\eta)=-0.9777, \mathbf{p}=\mathbf{e}_{5} .
$$

## Example 1



Figure: Normwise relative error $\left\|H_{k}-\Psi\right\| /\|\Psi\|$ vs. iteration $k$.

## Example 1



Figure: Componentwise relative error $\max _{i, j}\left|\left(H_{k}\right)_{i j}-\Psi_{i j}\right| / \Psi_{i j}$ vs. iteration $k$.

## Example 2



## Example 3



## Conclusions

- You can have your cake (shift) and eat it, too (componentwise accuracy). Se puede estar a la vez en la procesión y repicando las campanas
- In many examples, we can lower the number of iterations to match that of shift, while keeping the original high accuracy.
- Are there benefits in delaying the shift even further?

Reference Elena Addis's thesis at UNIFI, Elementwise accurate algorithms for nonsymmetric algebraic Riccati equations associated with M-matrices, https://hdl.handle.net/2158/1275470. Article version in preparation.

## Conclusions

- You can have your cake (shift) and eat it, too (componentwise accuracy). Se puede estar a la vez en la procesión y repicando las campanas
- In many examples, we can lower the number of iterations to match that of shift, while keeping the original high accuracy.
- Are there benefits in delaying the shift even further?

Reference Elena Addis's thesis at UNIFI, Elementwise accurate algorithms for nonsymmetric algebraic Riccati equations associated with M-matrices, https://hdl.handle.net/2158/1275470. Article version in preparation.

Thanks for your attention!

