Parallel Set-Valued Rational Approximation for PDEs

Simon Dirckx (Oden institute, Austin TX)
Karl Meerbergen (KU Leuven, Leuven BE)
Daan Huybrechs (KU Leuven, Leuven BE)

SIAM Conference on Computational Science and Engineering 2025 (CSE25)

The University of Texas at Austin
Oden Institute for Computational Engineering and Sciences

Outline

Rational approximation for PDEs

Some applications

Discretized parametized PDE

$$\mathbf{A}(s)\mathbf{u}(s) = \mathbf{f}(s)$$

Often, but not necessarily,

$$\mathbf{A}(s) = \sum_i g_i(s)\mathbf{A}_i$$

e.g. nonlinear damping, where $\mathbf{A}(s) = \mathbf{K} + g(s)\mathbf{D} + s^2\mathbf{M}$.

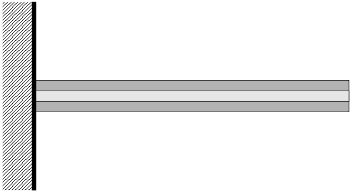

Figure: A clamped sandwich beam

Nonlinearities often introduce poles close to the domain of of approximation.

Rational approximation for PDEs

Some applications

Compactification of BEM operators (Dirckx, S., Meerbergen, K., Huybrechs, D., 2024):
After moving to a BEM formulation, say for the Helmholtz equation

$$\Delta u + \kappa^2 u = 0 \text{ in } \Omega$$
$$\mathcal{B}u = g \text{ on } \partial\Omega$$

we obtain

$$\mathbf{T}_1(\kappa)\mathbf{u}(\kappa) = \mathbf{T}_2(\kappa)\mathbf{g}(\kappa)$$

with $\mathbf{T}_1, \mathbf{T}_2$ discretized non-local compact operators on a boundary mesh $\Gamma \approx \partial\Omega$. The dependence $\kappa \mapsto \mathbf{T}_i(\kappa)$ is extremely expensive, especially in the near-field regime.

Rational approximation for PDEs

Some applications

Hedgehog method with rational approximation (Bagge, J., Thornberg, A.-K.,2023):

1. Rational functions as natural choice for DLP approximation
2. Large collection of (similar) rational functions

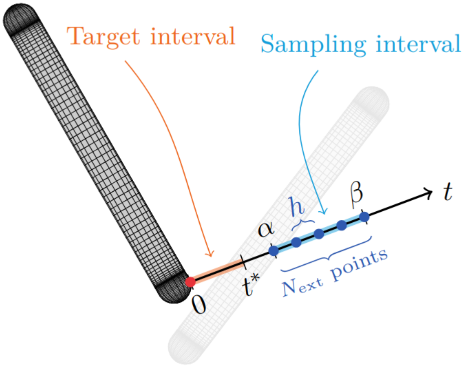

Figure: Sampling and target interval for one DLP function to be approximated and extrapolated.

Outline

Set-Valued AAA

Greedy Rational Approximation with Shared Poles

The SV-AAA Algorithm, analogous to scalar AAA (Adaptive Antoulas-Anderson):

Set-Valued AAA (SV-AAA)

Given a vector-valued function $\mathbf{f} : \mathbb{C} \mapsto \mathbb{C}^N$, SV-AAA satisfies after iteration m:

$$\mathbf{f}(z) \approx \mathbf{r}_m(z) = \mathbf{n}_m(z) \Big/ \mathbf{d}_m(z) = \left(\sum_{\nu=1}^{m} \frac{w_\nu \mathbf{f}(z_\nu)}{z - z_\nu} \Big/ \sum_{\nu=1}^{m} \frac{w_\nu}{z - z_\nu} \right)$$

i.e. a degree $(m-1, m-1)$ vector-valued rational interpolant in barycentric form, subject to $\sum_\nu |w_\nu|^2 = 1$ and

- $z_m := \arg\max_{Z \setminus Z_{m-1}} \|\mathbf{f}(z) - \mathbf{r}_{m-1}(z)\|_\infty$,
- $\sum_i \|\mathbf{d}_{m,i} \mathbf{f}_i - \mathbf{n}_{m,i}\|^2$ is minimal over $Z \setminus Z_m$,

i.e., greedy optimization.

Matrix interpretation

The $\mathbf{F}$-matrix

F-matrix (w.r.t. Z)

We collect the component fibers of $\mathbf{f} = [f_1, \ldots, f_N]$ discretized on Z:

$$\underbrace{\begin{bmatrix} f_1(z_1) & f_2(z_1) & \cdots & f_N(z_1) \\ f_1(z_2) & f_2(z_2) & \cdots & f_N(z_2) \\ \vdots & \vdots & \ddots & \vdots \\ f_1(z_{|Z|}) & f_2(z_{|Z|}) & \cdots & f_N(z_{|Z|}) \end{bmatrix}}_{:=F\in\mathbb{C}^{Z\times N}}$$

- column j of $F \leftrightarrow$ evaluating f_j on Z
- row i of $F \leftrightarrow$ evaluating $\mathbf{f}(z_i), z_i \in Z$
- $\mathbf{r}_m \approx \mathbf{f} \iff \widetilde{F}_m \approx F$

Row Interpolative Decomposition (RID)

Approximate *row interpolative decomposition* (RID) for $A \in \mathbb{C}^{n_1\times n_2}$ w.r.t. $\|\cdot\|$:

$$A \approx A_m = H_m \cdot A(I_m, :)$$

with $H_m \in \mathbb{C}^{n_1\times m}$ and $I_m \subseteq \{1, \ldots, n_1\}$, such that $\|A - A_m\|$ is sufficiently small.

Two important norms:

$$\|A\|_{p,\infty} := \max_{i\in|Z|} \|A(i,:)\|_p$$

$$\|A\|_{\max} := \max_{ij} |A(i,j)|$$

Interpolative decompositions

SV-AAA as an RID

We can make our intuition more precise:

Theorem (SV-AAA as an RID)

Let $\mathbf{f} \approx \mathbf{r}_m$ *on* $Z \subset \mathbb{C}$ *with support points* $Z_m \subset Z$ *such that for some* $\epsilon > 0$

$$\textbf{res}_m := \sup_{z \in Z} \|\mathbf{f}(z) - \mathbf{r}_m(z)\|_p < \epsilon. \tag{1}$$

Let F *and* $\widetilde{F}_m$ *as before. Then* $\widetilde{F}_m$ *is of rank* m*, and can be written as*

$$\widetilde{F}_m = H_m \cdot F(Z_m, :),$$

which constitutes an approximate RID with respect to the $\| \cdot \|_{p,\infty}$*-norm, of the same tolerance* ϵ*.*

We have $H_m(\nu, j) = \frac{w_\nu}{z_i - z_\nu} \Big/ \left(\sum_{j=1}^{m} \frac{w_j}{z_i - z_j} \right)$ with $\nu \in \{1, \dots, m\}$ and $z_j \in Z$.

Theorem (transitive property of RIDs)

Let $A \in \mathbb{C}^{n_1 \times n_2}$ *be a matrix that can be factorized as* $A = XY$, $X \in \mathbb{C}^{n_1 \times k}$ *and* $Y \in \mathbb{C}^{k \times n_2}$ *and let* $X \approx X_m = H_m \cdot X(I_m, :)$ *be an approximate RID for* X *w.r.t. the norm* $\| \cdot \|_{p,\infty}$ *such that*

$$\|X - X_m\|_{p,\infty} < \epsilon.$$

Then $A \approx A_m = H_m \cdot A(I_m, :)$ *is an approximate RID for* A *w.r.t.* $\| \cdot \|_{\max}$ *such that*

$$\|A - A_m\|_{\max} < \frac{k\epsilon}{\sqrt[p]{k}} \|Y\|_{\max}$$

More generally, if $\| \cdot \|$, $\| \cdot \|_\alpha$ and $\| \cdot \|_\beta$ s.t. $\forall M_1 \in \mathbb{C}^{n_1 \times k}$, $M_2 \in \mathbb{C}^{k \times n_2}$:

$$\|M_1 M_2\| \leq c_{\alpha,\beta}(k) \|M_1\|_\alpha \|M_2\|_\beta.$$

Then, with $\|X - X_m\|_\alpha < \epsilon$,

$$\|A - A_m\| < \epsilon c_{\alpha,\beta}(k) \|Y\|_\beta$$

Outline

QR-based SV-AAA

The basic idea

Simple QR-AAA factorization

After RRQR(ϵ), by the transitive property:

$$\begin{aligned} F &\approx QR \\ &\approx \tilde{Q}_m R \\ &= H_m Q(Z_m, :)R \\ &= H_m F(Z_m, :) \end{aligned}$$

QR-AAA:

1. Assemble F
2. $F \approx QR$ (pivoted RRQR)
3. $Q\Gamma \approx \widetilde{Q}_m \Gamma$ (SV-AAA)
4. Store (Z_m, W_m)

Weighted Approximation

For stability, we approximate $Q\Gamma$, $\Gamma := \text{diag}(R)$, since $\|\Gamma^{-1}R\|_{\max} = 1$, so

$$\begin{aligned} \text{err}_{\max} &\leq \epsilon + \sqrt{k}\|Q\Gamma - \widetilde{Q}_m\Gamma\|_{2,\infty}\|\Gamma^{-1}R\|_{\max} \\ &\leq \epsilon(1 + \sqrt{k}). \end{aligned}$$

Intuitively: simple smoothing/noise reduction, by weighting basis vectors in Q.

The Algorithm

Schematic overview

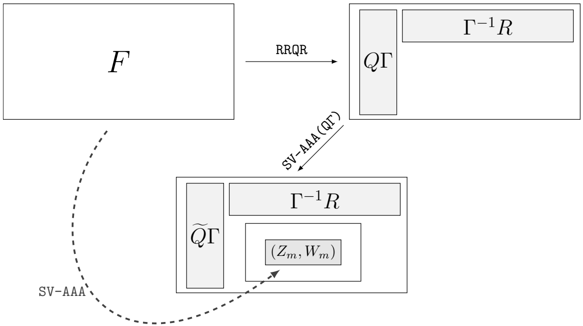

Figure: Diagram showing the principle of QR-AAA. Here $\Gamma = \text{diag}(R)$.

Custom Bases

Generalized QR-AAA

In many applications we can attempt to exploit additional knowledge about our problem:

Generalized QR-AAA

Let $\Phi \in \mathbb{C}^{Z \times N}$ correspond to a set of functions $\{\varphi_j\}_{j=1}^{k}$ i.e. $\Phi(i,j) = \varphi_j(z_i)$. Then, suppose $F \approx \Phi C$, with $C = \Phi^\dagger F$. QR-AAA then still works, with Q replaced by Φ and Γ the diagonal matrix defined by $\Gamma_{ii} := \|C(i,:)\|_\infty$.

The matrix Φ can either be given by analytical knowledge, or computed using some other scheme. The assumption that $F \approx \Phi\Phi^\dagger F$ just means that Φ constitutes an approximate generating set for the columns of F.

Performance and Stability

Numerical results

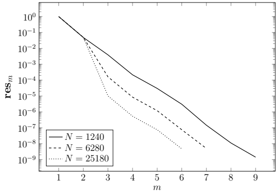

(a) `sandwich_beam`: SV-AAA

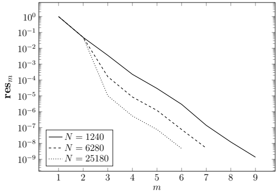

(b) `sandwich_beam`: QR-AAA

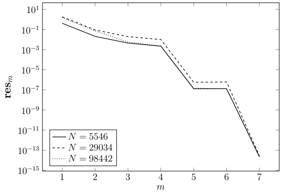

(c) `photonic_crystal`: SV-AAA

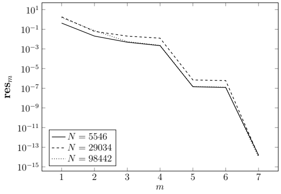

(d) `photonic_crystal`: QR-AAA

Figure: Residue $\mathbf{res}_m$ for SV-AAA (left) and QR-AAA (right) over the degree m, for the first two selected problems.

Performance and Stability

Numerical results

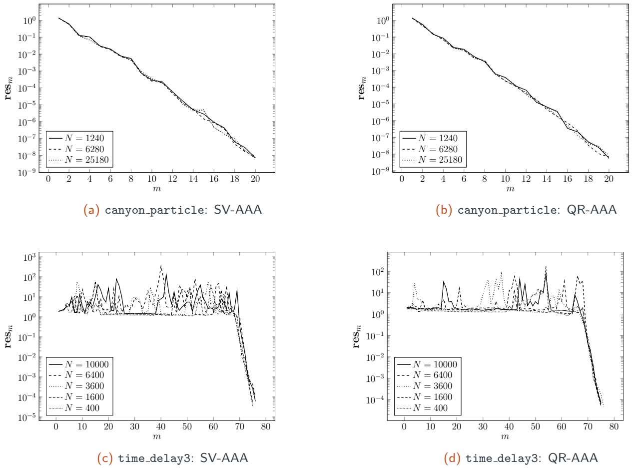

Figure: Residue $\mathbf{res}_m$ for SV-AAA (left) and QR-AAA (right) over the degree m, for the second two selected problems.

Performance and Stability

Numerical results

N	t_QR	t_AAA_Q	t_AAA_F
1240	0.025	4.02e−3	1.93
6280	0.124	2.24e−3	6.04
25180	0.470	1.63e−3	20.41

(a) sandwich_beam

N	t_QR	t_AAA_Q	t_AAA_F
5546	0.108	5.56e−3	2.61122
29034	0.527	1.26e−3	13.7263
98442	1.756	1.14e−3	54.847

(b) photonic_crystal

N	t_QR	t_AAA_Q	t_AAA_F
5971	0.278	8.351e−2	70.4021
7432	0.381	8.39e−2	86.9885
10157	0.447	8.32e−2	119.217
15121	0.683	8.38e−2	177.209

(c) canyon_particle

N	t_QR	t_AAA_Q	t_AAA_F
400	0.010	9.21	232.38
1600	0.091	9.41	1033.47
3600	0.241	9.51	2285.06
6400	0.440	9.67	4431.65
10000	0.701	9.41	6370.27

(d) time_delay3

Figure: Tables of absolute timings for SV-AAA and QR-AAA, for the four NLEVP problems. All timings are reported in seconds.

Outline

Parallelization

Motivation

1. Large-scale problems
2. Trivial parallelization of the construction of F (dominant cost!)
3. Expected: local and global correlation

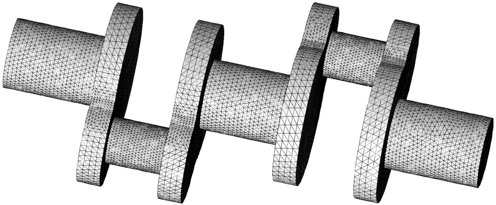

Figure: Mesh of a Crankshaft

Parallelization

Overview

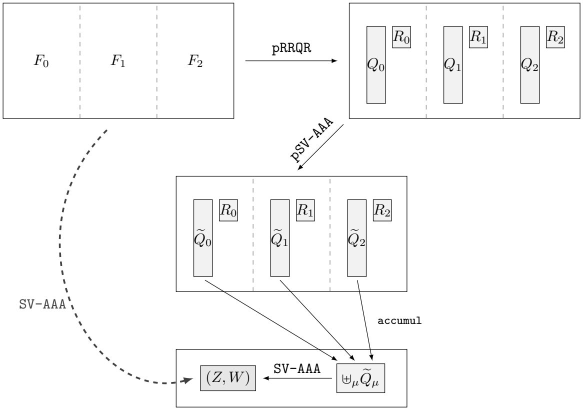

Figure: Overview of the parallel QR based set-valued AAA approach.

Local Approximations

The accumulate support problem

Accumulate support problem

Given a set of set-valued rational functions $\{\widetilde{Q}_\mu\}_\mu$ with support nodes and weights $\{Z_\mu, W_\mu\}$ such that $Z_\mu \subset Z$, find a global set-valued rational approximation $\widetilde{Q}$ using only samples in $Z^+ \subset Z$, where Z^+ is as small as possible.

Gluing Local Approximations

Some theoretical insights

Naively we might think of setting $Z^+ := \cup_\mu Z_\mu$. However:

Minimal size

Rational interpolation of degree $m-1$ requires at least $2m-1$ sample points, but $|\cup_\mu Z_\mu| \geq 2m-1$ often does not hold! We must extend by some $Z^{\mathrm{e}} \subset Z$.

Ill-conditioning

Only requiring $|Z^+| = 2m-1$ may still not be sufficient. Approximating a degree $m-1$ rational function on $\cup_\mu Z_\mu \cup Z^e$ can be very ill-conditioned; given some rational function f on Z, it is possible that its rational approximant r on $Z^+ := (\cup_\mu Z_\mu \cup Z^{\mathrm{e}})$ satisfies

$$\|f - r\|_{Z^+,\infty} < \epsilon$$

while the error $\|f - r\|_{Z,\infty}$ on the full grid grows exponentially large in m.

Gluing Local Approximations

Some theoretical insights

Definition (Reciprocal Christoffel)

For an ordered set $Z^e = \{z_i\}_i$ in $[-1, 1]$:

$$\zeta(Z^{\mathsf{e}}) := \max_i \int_{z_i}^{z_{i+1}} \frac{1}{\sqrt{1-x^2}}\,\mathrm{d}x.$$

Definition (Good extension)

Given QR-AAA approximations $\{\widetilde{Q}_\mu\}_{\mu=1}^{l}$ with support points and weights $\{(Z_\mu, W_\mu)\}_{\mu=1}^{l}$, we set

$$m^+ := 2\left|\cup_{\mu=1}^{l} Z_\mu\right| - 2 \text{ and } Z^+ := (\cup_\mu Z_\mu) \cup Z^{\mathsf{e}}.$$

$Z^{\mathsf{e}} \subset Z$ is a *good extension set for degree* m^+ if $m^+\zeta(Z^{\mathsf{e}}) = \alpha$ with $0 < \alpha < 1$ sufficiently small. Define

$$\widetilde{Q}_1 \uplus \cdots \uplus \widetilde{Q}_l := [\widetilde{Q}_1(Z^+, :) \cdots \widetilde{Q}_l(Z^+, :)]$$

Gluing Local Approximations

Some theoretical insights

As shown in (Adcock, B. et al., 2018), if $m^+\zeta(Z^e) = \alpha < 1$ we have that

$$B(Z^{\mathsf{e}}, m^+) := \sup_{p \in \mathbb{P}_{m^+}} \{\|p\|_{\infty,[-1,1]} \mid \|p\|_{\infty,Z^{\mathsf{e}}} \leq 1\} < \frac{1}{1-\alpha} \tag{2}$$

Theorem

Suppose $f := n/d = p/q$ is a rational function on $[-1,1]$ of degree $m-1$, defined by support points and weights (Z_m, W_m). Then if $Z_m^+ = Z_m \cup Z^{\mathsf{e}}$, we have that any AAA approximant $\hat{f} = \hat{n}/\hat{d} = \hat{p}/\hat{q}$ of degree at most $m-1$ such that

$$\|\hat{d}f - \hat{n}\|_{Z_m^+,\infty} < \epsilon \tag{3}$$

satisfies

$$\frac{\|p\hat{q} - \hat{p}q\|_\infty}{\|q\|_\infty} < \|\hat{\ell}\|_{\infty,Z_m^+} B(Z^{\mathsf{e}}, m^+)\epsilon$$

where $\hat{\ell} := \prod_{\nu=1}^{m^+}(z - z_\nu)$ is the node polynomial for $\hat{f}$.

Parallel QR-based SV-AAA (PQR-AAA)

The algorithm

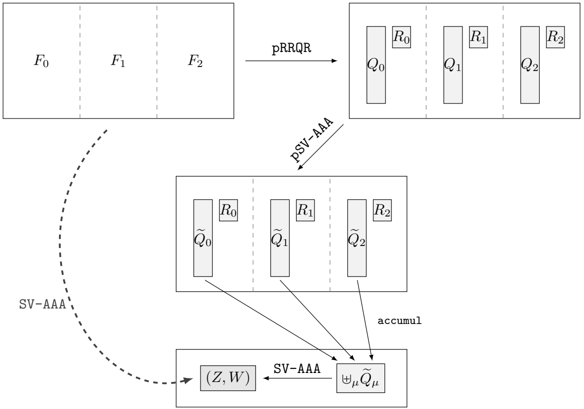

Figure: Overview of the parallel QR based set-valued AAA approach.

Performance

Near-field BEM compression

Single layer boundary operator:

$$\mathcal{S}(\kappa) : H^{-1/2}(\partial\Omega) \to H^{1/2}(\partial\Omega) : u \mapsto \mathcal{S}(\kappa)[u]$$

where

$$(\mathcal{S}(\kappa)[u])(\mathbf{x}) = \lim_{\mathbf{x}\to\partial\Omega} \int_{\partial\Omega} G(\mathbf{x}, \mathbf{y}; \kappa) u(\mathbf{y}) dS_{\mathbf{y}}$$

with $G(\mathbf{x}, \mathbf{y}; \kappa)$ the Green's kernel at wave number κ. The operator $\mathcal{S}(\kappa)$ is discretized to $S(\kappa)$ using Galerkin discretization on a triangular surface mesh. Approximation over 28 cores of $N = 1000000$ near-field BEM fibers. The required tolerance was set to 10^{-6}. The (dimensionless) wavenumber range was set to $[1., 80.]$, and discretized into 500 equispaced points.

Performance

Near-field BEM compression

1. `t_F`: the maximal time needed to assemble the matrix F_μ, over the processors $\mu \in \{1, \dots, 28\}$,
2. `t_QR`: maximal time needed for RRQR over the processors $\mu \in \{1, \dots, 28\}$,
3. `t_AAA_Q`: the maximal time needed to compute $\widetilde{Q}_\mu$ over the processors $\mu \in \{1, \dots, 28\}$.
4. `t_fin`: the time needed to compute the final SV-AAA for $\uplus_\mu Q_\mu$

t_F	t_QR	t_AAA_Q	t_fin	**Err**$_\infty(F)$
343.213	.926	.019	.259	$6.64897 \cdot 10^{-8}$

Table: Error and components of the PQR-AAA timings (in seconds) for the the near-field of $S(\kappa)$ defined on the sphere S^2.

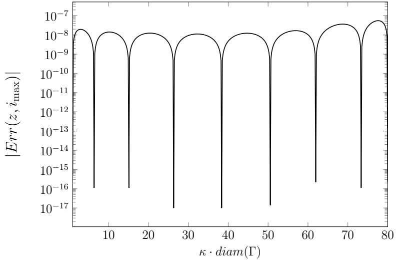

Figure: $\|\cdot\|_\infty$ error of the PQR-AAA approximation over $\kappa \cdot \text{diam}(\Gamma)$

Additional comments

1. This talk is based on (Dirckx, S. et al., 2024)
2. Alternative is ‘sketchAAA’ (Güttel, S. et al., 2024).
3. General Z still not understood.
4. More refined Christoffel function analysis is welcome.
5. Many additional applications.