Collocated Reduced-Order Models for High-Order Schemes

Analysis and Application

M.G. Carlino, A. Del Grosso, A. Iollo, D. Sipp

MONHADE Common Team Inria - Onera
Institut de Mathématiques de Bordeaux, Université de Bordeaux & DAAA, Onera

High-Fidelity Discretization

Full Order Model

We consider a one-step iterative scheme in time arising from the discretization of a PDE:

$$u^{n+1} = \mathcal{N}(u^n; \mu), \quad u^n \in \mathbb{R}^N$$

Computational complexity:

- The dimension N is large ($10^6 - 10^9$ DOFs).
- The operator $\mathcal{N}$ is non-linear and computationally intensive.
- **Objective:** Construct a reduced approximation $u_c \in \mathbb{R}^{N_c}$ with $N_c \ll N$.

The MOR Ansatz: Linear Subspace

Ansatz

MOR assumes the state u evolves in a low-dimensional subspace spanned by a basis Φ:

$$u(t;\mu) \approx u_r(t;\mu) = \sum_{m=1}^{N_r} \phi_m a_m(t;\mu) = \Phi a(t;\mu)$$

- $\Phi \in \mathbb{R}^{N \times N_r}$ is the basis ($N_r \ll N$).
- We assume Φ is W-orthonormal: $\Phi^T W \Phi = I_{N_r}$.
- $a \in \mathbb{R}^{N_r}$ are the reduced (generalized) coordinates.
- W is a symmetric positive-definite weight matrix (e.g., from finite element integration).

pMOR: Projection

Standard (Galerkin) projection enforces orthogonality of the HF residual to Φ:

$$\Phi^T W \left(u_r^{n+1} - \mathcal{N}(u_r^n; \mu)\right) = 0 \implies a^{n+1} = \Phi^T W \cdot \mathcal{N}(\Phi a^n, \mu)$$

Computational complexity

Despite the reduced state dimension (N_r), the evaluation of the non-linear term $\mathcal{N}(\Phi a^n)$ prevents online efficiency.

1. **Reconstruction:** $u^n = \Phi a^n \in \mathbb{R}^N$.
2. **Evaluation:** $\mathcal{O}(N)$ operations to compute $\mathcal{N}(u^n)$.
3. **Projection:** $\mathcal{O}(N \cdot N_r)$ to project back.

pMOR: Projection

Standard (Galerkin) projection enforces orthogonality of the HF residual to Φ:

$$\Phi^T W \left(u_r^{n+1} - \mathcal{N}(u_r^n; \mu)\right) = 0 \implies a^{n+1} = \Phi^T W \cdot \mathcal{N}(\Phi a^n, \mu)$$

Computational complexity

Despite the reduced state dimension (N_r), the evaluation of the non-linear term $\mathcal{N}(\Phi a^n)$ prevents online efficiency.

1. **Reconstruction:** $u^n = \Phi a^n \in \mathbb{R}^N$.
2. **Evaluation:** $\mathcal{O}(N)$ operations to compute $\mathcal{N}(u^n)$.
3. **Projection:** $\mathcal{O}(N \cdot N_r)$ to project back.

Requirement: A hyper-reduction strategy is necessary to decouple the online complexity from the Full Order Model dimension N.

cMOR Collocation: Sparse Evaluation of Differential Operators

We move from reduced coordinates $a \in \mathbb{R}^{N_r}$ to physical nodal values $\bar{u}_c \in \mathbb{R}^{N_c}$. We compute $\mathcal{N}$ only on a subset of indices $\mathcal{J}_c \subset \{1, \ldots, N\}$, with $|\mathcal{J}_c| = N_c \ll N$.

cMOR Collocation: Sparse Evaluation of Differential Operators

We move from reduced coordinates $a \in \mathbb{R}^{N_r}$ to physical nodal values $\bar{u}_c \in \mathbb{R}^{N_c}$. We compute $\mathcal{N}$ only on a subset of indices $\mathcal{J}_c \subset \{1, \ldots, N\}$, with $|\mathcal{J}_c| = N_c \ll N$.

The Stencil Requirement: Differential operators (e.g., $\nabla \cdot F$) are local. For example, evaluating $\mathcal{N}(u)_k$ at node x_k requires the stencil $\{x_{k-1}, x_k, x_{k+1}\}$.

cMOR Collocation: Sparse Evaluation of Differential Operators

We move from reduced coordinates $a \in \mathbb{R}^{N_r}$ to physical nodal values $\bar{u}_c \in \mathbb{R}^{N_c}$. We compute $\mathcal{N}$ only on a subset of indices $\mathcal{J}_c \subset \{1, \ldots, N\}$, with $|\mathcal{J}_c| = N_c \ll N$.

The Stencil Requirement: Differential operators (e.g., $\nabla \cdot F$) are local. For example, evaluating $\mathcal{N}(u)_k$ at node x_k requires the stencil $\{x_{k-1}, x_k, x_{k+1}\}$.

Prolongation: The basis Φ provides a global spectral approximation, allowing implicitly reconstruction of the stencil support required to evaluate the physics at sparse points.

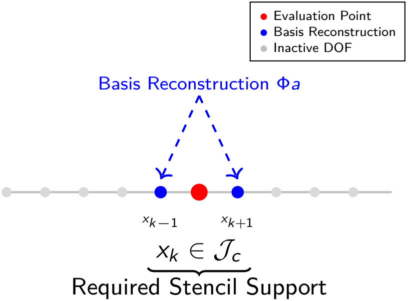

$$\partial_x u \approx \frac{u_{k+1} - u_{k-1}}{2\Delta x}$$

The cMOR Scheme

The scheme operates exclusively on a sparse subset of the grid.

1. **Sparse Reconstruction:** From nodal values $\bar{u}_c$, we reconstruct the field *only* on the active stencils $\mathcal{S}$.
$$u_{\mathcal{S}} = (\mathcal{P}_c)_{\mathcal{S}} \cdot \bar{u}_c^n$$
2. **Local Residual:** We evaluate the nonlinear differential operator locally.
$$r_c = \mathcal{N}(u_{\mathcal{S}})|_{\mathcal{J}_c}$$
3. **Time Integration:**
$$\bar{u}_c^{n+1} = \bar{u}_c^n + \Delta t \cdot r_c$$

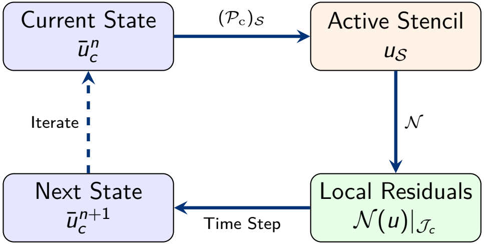

Sparse Quadrature via Non-Negative Least Squares (NNLS)

We define a set of grid indices $\mathcal{J}_c$ and a matrix of corresponding quadrature weights W_ϵ by computing the **projection of snapshots** onto the basis modes.

Input Data

Let $\mathcal{S} = \{u^1, \ldots, u^K\}$ be a collection of solution snapshots. We require that the sparse projection matches the high-fidelity projection for every snapshot and every mode:

$$\langle \phi_m, u^k \rangle_{W_\epsilon} \approx \langle \phi_m, u^k \rangle_W \quad \forall m \in 1..N_r, \forall k \in 1..K$$

Sparse Quadrature via Non-Negative Least Squares (NNLS)

We define a set of grid indices $\mathcal{J}_c$ and a matrix of corresponding quadrature weights W_ϵ by computing the **projection of snapshots** onto the basis modes.

Input Data

Let $\mathcal{S} = \{u^1, \ldots, u^K\}$ be a collection of solution snapshots. We require that the sparse projection matches the high-fidelity projection for every snapshot and every mode:

$$\langle \phi_m, u^k \rangle_{W_\epsilon} \approx \langle \phi_m, u^k \rangle_W \quad \forall m \in 1..N_r, \forall k \in 1..K$$

The NNLS Quadrature Optimization Problem

- **Matrix G:** Rows correspond to the product of mode m and snapshot k at grid point i: $G_{(mk),i} = \phi_m(x_i) \cdot [u^k](x_i)$.
- **Vector b:** high-fidelity projections $b_{(mk)} = \langle \phi_m, u^k \rangle_W$.

The weights w^* are found by solving: $w \geq 0$ s.t. $\|\mathbf{G}w - \mathbf{b}\|_2 < \epsilon \|\mathbf{b}\|_2$. Then:

$$\mathcal{J}_c = \mathrm{supp}(w^*), \quad W_\epsilon = \mathrm{diag}(w^*)$$

Operator Definitions: Inter-Grid Transfer Operators

We define the mappings between the high-dimensional physical space $\mathbb{R}^N$ and the sparse nodal space $\mathbb{R}^{N_c}$ (where $N_c \leq N_r N_K$, Carathéodory theorem for cones) and *vice versa*.

Restricted Basis & Weights

Let $M \in \mathbb{R}^{N_c \times N}$ be the boolean mask for indices $\mathcal{J}_c$. We define:

- **Restricted Basis:** $\overline{\Phi} = M\Phi \in \mathbb{R}^{N_c \times N_r}$.
- **Active Weights:** $\overline{W}_\epsilon = \mathrm{diag}(w|_{\mathcal{J}_c}) \in \mathbb{R}^{N_c \times N_c}$ (Positive definite).

Operator Definitions: Inter-Grid Transfer Operators

We define the mappings between the high-dimensional physical space $\mathbb{R}^N$ and the sparse nodal space $\mathbb{R}^{N_c}$ (where $N_c \leq N_r N_K$, Carathéodory theorem for cones) and *vice versa*.

Restricted Basis & Weights

Let $M \in \mathbb{R}^{N_c \times N}$ be the boolean mask for indices $\mathcal{J}_c$. We define:

- **Restricted Basis:** $\overline{\Phi} = M\Phi \in \mathbb{R}^{N_c \times N_r}$.
- **Active Weights:** $\overline{W}_\epsilon = \mathrm{diag}(w|_{\mathcal{J}_c}) \in \mathbb{R}^{N_c \times N_c}$ (Positive definite).

The Prolongation Operators ($\mathbb{R}^{N_c} \to \mathbb{R}^N$)

- **Least-Squares ($\mathcal{P}_{\mathrm{ls}}$):** Since $\langle \phi_m, u^k \rangle_{W_\epsilon} \approx a_m$, and $u_r = \Phi a$ then:

$$\mathcal{P}_{\mathrm{ls}} = \Phi \overline{\Phi}^T \overline{W}_\epsilon \in \mathbb{R}^{N \times N_c}$$

Operator Definitions: Inter-Grid Transfer Operators

We define the mappings between the high-dimensional physical space $\mathbb{R}^N$ and the sparse nodal space $\mathbb{R}^{N_c}$ (where $N_c \leq N_r N_K$, Carathéodory theorem for cones) and *vice versa*.

Restricted Basis & Weights

Let $M \in \mathbb{R}^{N_c \times N}$ be the boolean mask for indices $\mathcal{J}_c$. We define:

- **Restricted Basis:** $\overline{\Phi} = M\Phi \in \mathbb{R}^{N_c \times N_r}$.
- **Active Weights:** $\overline{W}_\epsilon = \mathrm{diag}(w|_{\mathcal{J}_c}) \in \mathbb{R}^{N_c \times N_c}$ (Positive definite).

The Prolongation Operators ($\mathbb{R}^{N_c} \to \mathbb{R}^N$)

- **Least-Squares ($\mathcal{P}_{\mathrm{ls}}$):** Since $\langle \phi_m, u^k \rangle_{W_\epsilon} \approx a_m$, and $u_r = \Phi a$ then:

$$\mathcal{P}_{\mathrm{ls}} = \Phi \overline{\Phi}^T \overline{W}_\epsilon \in \mathbb{R}^{N \times N_c}$$

- **Interpolatory ($\mathcal{P}_{\mathrm{c}}$):** Corrects $\mathcal{P}_{\mathrm{ls}}$ to ensure $M\mathcal{P}_{\mathrm{c}} = I_{N_c}$ (exactness at nodes).

$$\mathcal{P}_{\mathrm{c}} = \underbrace{(I_N - M^T M)}_{\text{Fill unsampled points}} \mathcal{P}_{\mathrm{ls}} + \underbrace{M^T}_{\text{Exact at } \mathcal{J}_c}$$

The cMOR full formulation

For analysis the cMOR is written

$$u^{n+1} = \mathcal{P}_{\mathrm{c}} M \mathcal{N}(u^n; \mu), \quad u^n \in \mathbb{R}^N$$

In the linear case

$$u^{n+1} = \mathcal{P}_{\mathrm{c}} M A(CFL; \mu) u^n, \quad u^n \in \mathbb{R}^N$$

Property 1: Injectivity and Rank

Theorem

The operator $\mathcal{P}_\mathrm{c}$ has full column rank: rank$(\mathcal{P}_\mathrm{c}) = N_c$.

Proof: It suffices to show $\mathrm{Ker}(\mathcal{P}_\mathrm{c}) = \{0\}$. Let $\overline{v} \in \mathrm{Ker}(\mathcal{P}_\mathrm{c})$.

1. Hypothesis: $\mathcal{P}_\mathrm{c}\overline{v} = 0$.
2. Apply the sampling mask M:
$$M(\mathcal{P}_\mathrm{c}\overline{v}) = 0$$
3. By the definition of $\mathcal{P}_\mathrm{c}$:
$$M[(I - M^T M)\mathcal{P}_\mathrm{ls}\overline{v} + M^T\overline{v}] = 0$$
4. Orthogonality of the mask ($MM^T = I_{N_c}$):
$$\underbrace{M(I - M^T M)}_{=0}\mathcal{P}_\mathrm{ls}\overline{v} + \underbrace{MM^T}_{=I}\overline{v} = 0$$
5. Conclusion: $\overline{v} = 0$. $\square$

Property 2: Idempotency of the Discrete Projector

Theorem

The interpolatory operator $\Pi_c = \mathcal{P}_\mathrm{c} M : \mathbb{R}^N \to \mathbb{R}^N$ is a projector: $(\mathcal{P}_\mathrm{c} M)^2 = \mathcal{P}_\mathrm{c} M$.

Proof: Let $u \in \mathbb{R}^N$. Define the restriction $\overline{u} = Mu$.

$$
\begin{aligned}
(\mathcal{P}_\mathrm{c} M)^2 u &= (\mathcal{P}_\mathrm{c} M)(\mathcal{P}_\mathrm{c} \underbrace{Mu}_{\overline{u}}) \\
&= \mathcal{P}_\mathrm{c}(\underbrace{M\mathcal{P}_\mathrm{c}}_{\text{Identity on } \mathbb{R}^{N_c}} \overline{u}) \\
&= \mathcal{P}_\mathrm{c}(\overline{u}) \\
&= \mathcal{P}_\mathrm{c} M u
\end{aligned}
$$

Property 3: Optimality of the Interpolation Error

Theorem

The approximation error of the cMOR interpolatory projector Π_c is always less than or equal to that of the least square approximation $\Pi_{ls} = \mathcal{P}_{\mathrm{ls}} M$:

$$\|u - \Pi_c u\|_W \leq \|u - \Pi_{ls} u\|_W$$

Proof: Let $D = M^T M$ be the diagonal mask (1 on $\mathcal{J}_c$, 0 elsewhere).

1. **Algebraic Relation:** From the definition $\mathcal{P}_{\mathrm{c}} = (I - D)\mathcal{P}_{\mathrm{ls}} + M^T$, the errors satisfy the identity:

$$e_c = (I - D)e_{ls}$$

2. **Norm Decomposition:** Since D and $(I - D)$ have disjoint supports, the vectors are orthogonal. By the Pythagorean theorem in the weighted norm:

$$\|e_{ls}\|_W^2 = \|(I - D)e_{ls} + De_{ls}\|_W^2 = \|(I - D)e_{ls}\|_W^2 + \|De_{ls}\|_W^2$$

3. **Conclusion:** Substituting e_c into the equation:

$$\|e_{ls}\|_W^2 = \|e_c\|_W^2 + \underbrace{\|De_{ls}\|_W^2}_{\geq 0} \implies \|e_c\|_W \leq \|e_{ls}\|_W$$

$\square$

Linear stability of cROM

The cROM evolves via $\Pi_c = \mathcal{P}_{\mathrm{c}} M$.

Proof of the Operator Norm Bound

Using the property that the interpolatory error is bounded by the spectral error (Property 3):

$$\begin{aligned}
\|\mathcal{P}_{\mathrm{c}} M\mathbf{v} - \mathbf{v}\|_{\mathbf{W}} &\leq \left\|\mathbf{\Phi}\mathbf{\Phi}^T\mathbf{W}_{\varepsilon}\mathbf{v} - \mathbf{v}\right\|_{\mathbf{W}} \\
&\leq \left\|\mathbf{\Phi}\mathbf{\Phi}^T\mathbf{W}_{\varepsilon}\mathbf{v} - \mathbf{\Phi}\mathbf{\Phi}^T\mathbf{W}\mathbf{v}\right\|_{\mathbf{W}} + \left\|\mathbf{\Phi}\mathbf{\Phi}^T\mathbf{W}\mathbf{v} - \mathbf{v}\right\|_{\mathbf{W}} \\
&\leq \epsilon_{\mathrm{NNLS}} \|\mathbf{v}\|_{\mathbf{W}} + \epsilon_{\mathrm{POD}} \|\mathbf{v}\|_{\mathbf{W}}
\end{aligned}$$

Thus, the operator norm is bounded by: $\|\mathcal{P}_{\mathrm{c}} M\|_{\mathbf{W}} \leq 1 + \epsilon_{\mathrm{NNLS}} + \epsilon_{\mathrm{POD}}$.

Theorem (cROM conditional stability)

The cROM is stable only if the time-step provides sufficient dissipation:

$$\beta(CFL) \leq \frac{1}{1 + \epsilon_{\mathrm{NNLS}} + \epsilon_{\mathrm{POD}}}$$

Error Definition

Decomposition of the Error

Consider the error between the high-fidelity solution u^n and the cROM solution u_c^n:

$$e_n = u^n - u_c^n = u^n - \mathcal{P}_c M\mathcal{N}(u_c^{n-1})$$

This can be decomposed as:

$$e_n = e_{c,n} + e_{\text{model},n}$$

where

- $e_{c,n} = u^n - P_c M u^n$: **Collocation error** (NNLS hyper-reduction + POD truncation)
- $e_{\text{model},n} = P_c M\big(u^n - \mathcal{N}(u_c^{n-1})\big)$: **Model error**

In the **linear case** ($\mathcal{N}(u) = Au$):

$$e_{\text{model},n} = P_c MA(u^{n-1} - u_c^{n-1}) = B_c\, e_{n-1}, \quad B_c := P_c MA$$

Error Bound (Linear Case)

Recursive Formulation and Theoretical Bound

By induction, the error admits the following explicit recursive representation:

$$e_n = \sum_{i=1}^{n} B_c^{i-1} \, (e_{\mathrm{NNLS}})_{n-(i-1)}$$

(with $B_c^0 = I$ and initial collocation error $e_1 = 0$, since $u_c^1 = u^1$).

Theorem (Bound of the cROM error). If the stability condition holds, then

$$\|e_n\|_W \leq (\varepsilon_{\mathrm{POD}} + \varepsilon_{\mathrm{NNLS}})\|u^1\|_W \frac{1-\beta^n}{1-\beta}$$

As $n \to \infty$:

$$\|e_n\|_W \leq \frac{(\varepsilon_{\mathrm{POD}} + \varepsilon_{\mathrm{NNLS}})\|u^1\|_W}{1-\beta(\mathrm{CFL})}$$

Example I: 2D SWE Parametric Dam Break

Physical Model: 2D SWE on $\Omega = [0, 6] \times [0, 1]$ with variables $\mathbf{q} = (h, hu, hv)$.

- **Scenario:** Gradual dam opening at $x = L_x$ with parametric bathymetry $b(\mathbf{x}; \boldsymbol{\mu})$ (3 Gaussian bumps).
- **Params:** $\boldsymbol{\mu} \in [1, 5]^2$ controls bump amplitudes.

cMOR Handling of BCs (Lifting): Non-homogeneous Dirichlet BCs are handled via a lifting function q^1_{lift}:

$$h(\mathbf{x}, t) = h_*(\mathbf{x}, t) + q^1_{\text{lift}}(\mathbf{x}, t)$$

This ensures the reduced variable h_* vanishes at the boundary, allowing the basis to respect the changing gate level exactly.

HF Simulation Dynamics

Evolution of the free surface η.

cMOR Hyper-reduction & Numerical Results

Reduced Space Construction:

- **Affine Space:** Solutions approximated as Mean + Variation: $\mathbf{q} \approx \bar{\mathbf{q}} + \mathcal{P}_\epsilon \bar{\mathbf{q}}_c$.
- **Shared Collocation:** NNLS finds a common subset $\mathcal{J}_c$ for all variables (h, hu, hv).
- **Point Distribution:** NNLS clusters points around topographic peaks (regions of high sensitivity to $\boldsymbol{\mu}$).

Discretization:

- Mesh: $60,000$ cells (N).
- Scheme: FV Well-balanced (Hydrostatic reconstruction).
- Snapshots: $N_s = 3775$.

Performance (example):

Metric	Value	Unit
Sampling	$< 0.5\%$	N_c/N
Speed-up	**28**$\times$	vs HF
Error	$\approx 5 \cdot 10^{-4}$	Rel. L^2

Result: cMOR achieves massive speedup (up to 100x for coarser tol) with negligible accuracy loss vs projection.

Example II: High-Fidelity Structure-Preserving Scheme

Continuous periodic wave equation

$$\partial_{tt} u = c^2 \partial_{xx} u, \qquad x \in [0, L].$$

First-order Hamiltonian form

$$y = \begin{pmatrix} u \\ v \end{pmatrix}, \qquad \dot{y} = Ay, \qquad A = \begin{pmatrix} 0 & I \\ c^2 \partial_{xx} & 0 \end{pmatrix}.$$

Spatial discretization (4th-order compact)
Let $h = L/N$, $x_j = jh$. We pose:

$$B(L_h u)_j = \frac{1}{h^2} (u_{j+1} - 2u_j + u_{j-1}) = D_2(u)_j,$$

$$B = \frac{1}{12} \, \mathrm{tridiag}_{circ}(1, 10, 1).$$

Fourth-order accurate periodic Laplacian: $L_h = B^{-1} D_2$ (never explicitly computed).

Discrete Hamiltonian Structure

Semi-discrete system

$$\dot{y} = A_h y, \qquad A_h = \begin{pmatrix} 0 & I \\ c^2 L_h & 0 \end{pmatrix}.$$

Discrete Hamiltonian

Since L_h is symmetric negative semidefinite, define

$$H_h(u, v) = \frac{1}{2} v^T v - \frac{c^2}{2} u^T L_h u.$$

Time integration (exact exponential: Al-Mohy & Higham (2011), SIAM J. Sci. Comput.)

$$y^{n+1} = e^{\Delta t A_h} y^n.$$

The discrete Hamiltonian is exactly conserved

Substituting, we have

$$\frac{d}{dt} H_h(u(t), v(t)) = 0.$$

The fully discrete evolution preserves H_h up to numerical linear algebra accuracy.

Model Reduction and Reduced Hamiltonian

Reduced space and residual

$$u = \mathcal{P}_{\mathrm{c}} \bar{u} \qquad R(\bar{u}) = \mathcal{P}_{\mathrm{c}} \ddot{\bar{u}} - c^2 L_h \mathcal{P}_{\mathrm{c}} \bar{u}$$

Two options: fully collocated or projected reduced system

$$MR(\bar{u}) = 0 \qquad \text{or} \qquad \mathcal{P}_{\mathrm{c}}^T R(\bar{u}) = 0$$

Only the projected residual preserves by construction the Hamiltonian structure. Let

$$G = \mathcal{P}_{\mathrm{c}}^T \mathcal{P}_{\mathrm{c}}, \qquad L_{\mathrm{gal}} = \mathcal{P}_{\mathrm{c}}^T L_h \mathcal{P}_{\mathrm{c}}, \qquad L_{\mathrm{eff}} = G^{-1} L_{\mathrm{gal}}.$$

Reduced Hamiltonian system

$$\dot{y}_r = \begin{pmatrix} 0 & I \\ c^2 L_{\mathrm{eff}} & 0 \end{pmatrix} y_r = A_r y_r.$$

The reduced invariant ...

$$H_r(u_c, v_c) = \frac{1}{2} v_c^T v_c - \frac{c^2}{2} u_c^T L_{\mathrm{eff}} u_c.$$

... is preserved by exact time integration:

$$y_r^{n+1} = e^{\Delta t A_r} y_r^n.$$

Wave solution: training vs extrapolation

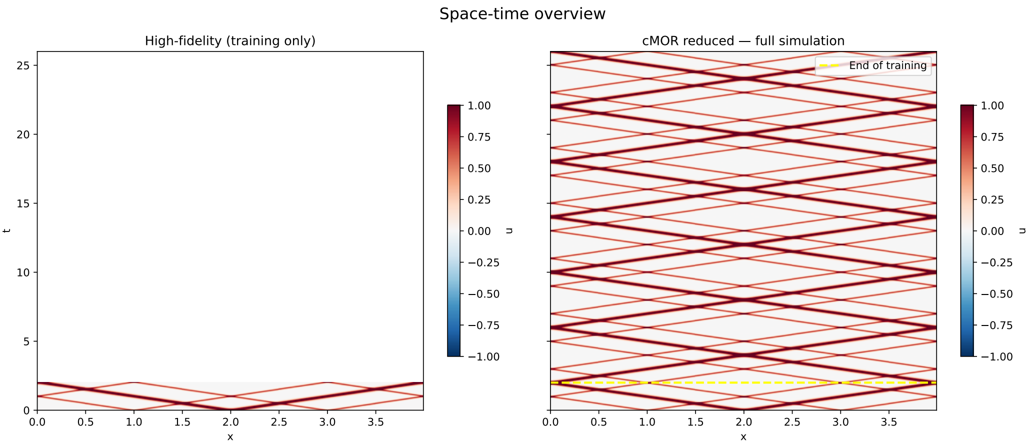

Figure: Comparison between high-fidelity snapshots (training interval) and cMOR reduced reconstruction during extrapolation. The reduced model preserves the wave structure beyond the training horizon.

Example Metrics: $\approx$ 7x speed up

High-Fidelity (HF)

- $N = 512$
- $\Delta t = 0.005$
- Steps: 401
- $t_f = 2.0$
- $\|u\|_\infty = 2.50$

Reduction (SVD)

- $M = 36$
- Energy: 99.9993%

Reduced Model

- $N_c = 71$
- $\Delta t = 0.005$
- Steps: 5201
- $t_f = 26.0$ ($\times 13$)
- $\|u\|_\infty$ norm change: 0.9994 $\times$ training value

Hyper-Reduction (NNLS)

- System: 14436×512
- Iter: 108
- Rel. res: 9.6×10^{-4}

Conclusions

Summary

- The cROM is a rigorously derived, stable, and convergent scheme.

Conclusions

Summary

- The cROM is a rigorously derived, stable, and convergent scheme.
- It inherits the stability of the underlying HF scheme, eventually reduced by controllable POD (ϵ_{POD}) and NNLS (ϵ_{NNLS}) approximation errors.

Conclusions

Summary

- The cROM is a rigorously derived, stable, and convergent scheme.
- It inherits the stability of the underlying HF scheme, eventually reduced by controllable POD (ϵ_{POD}) and NNLS (ϵ_{NNLS}) approximation errors.
- The analysis provides an error bound with respect to the HF solution.

Conclusions

Summary

- The cROM is a rigorously derived, stable, and convergent scheme.
- It inherits the stability of the underlying HF scheme, eventually reduced by controllable POD (ϵ_{POD}) and NNLS (ϵ_{NNLS}) approximation errors.
- The analysis provides an error bound with respect to the HF solution.
- Numerical results confirm the convergence and performance (up to 144x speedup) on linear and (relatively non-trivial) non-linear problems.