AMS-UMI International
Joint Meeting 2024
Palermo, July 23–26, 2024

Exploiting Low-rank Structures for the Solution of PDEs. Special Session B1

Alessandro Alla
Università Ca' Foscari Venezia, Italy
Boris Kramer
University of California San Diego, USA

Exploiting low-rank structures for the solution of partial differential equations (PDEs) has emerged as a powerful technique in the field of numerical simulation and scientific computing. Many real-world systems exhibit inherent low-dimensional patterns among solution variables, allowing for more efficient computation. Techniques like reduced-order modeling, tensor decomposition, matrix factorizations and certain machine learning methods help uncover and exploit these structures, enabling faster simulations and reduced memory usage. This approach not only accelerates the analysis of physical phenomena but also finds applications in fields such as optimal control, uncertainty quantification and computational design. By efficiently approximating complex PDE solutions with fewer degrees of freedom, low-rank methods facilitate scalable algorithms and offer enhanced interpretability, making them invaluable tools for tackling high-dimensional problems in various scientific and engineering domains.

Schedule and Abstracts

July 25, 2024

11:30–12:00 Low-rank Matrix Solvers for Evolutionary PDEs
Davide Palitta (Università di Bologna, Italy)

Abstract. We consider time-dependent partial differential equations (PDEs) of the form

$$\text{(1)} \qquad \begin{cases} u_t &= \mathcal{L}(u) + f, & \text{in } \Omega \times (0,T], \\ u &= g, & \text{on } \partial\Omega, \\ u(x,0) &= u_0(x), \end{cases}$$

where $\Omega \subset \mathbb{R}^d$, $d = 1, 2, 3$, and $\mathcal{L}$ is a linear differential operator involving only spatial derivatives.

Many efficient and diverse numerical methods have been developed in the literature for the solution of (1). Roughly speaking, these can be divided into two main categories: *time-marching* schemes and *all-at-once* methods.

In the former class, we first semi-discretized (1), namely we discretize only the space component of the differential problem by applying our favorite scheme. Then, a time integrator is applied to deal with the time component of the problem. The latter step leads to a sequence of linear systems, one at each time step, of the form

$$\text{(2)} \qquad A_i u_i = f_i, \qquad A_i \in \mathbb{R}^{\bar{n}\times\bar{n}},\ i = 1, \dots, \ell,$$

where $\bar{n}$ is the number of spatial degrees of freedom and ℓ is the number of time steps.

All-at-once methods significantly differ from time-marching schemes. Indeed, in place of a two-step strategy where we first discretize in space and then in time, the space and time components of (1) are discretized at once whenever an all-at-once approach is adopted. This different point of view in the discretization phase impacts the linear algebra step as well. In particular, rather than the sequence of linear systems in (2), we need to solve a single, yet much larger, linear system of the form

$$\text{(3)} \qquad \mathcal{A}\mathbf{u} = \mathbf{f},$$

where $\mathcal{A} \in \mathbb{R}^{\bar{n}\ell\times\bar{n}\ell}$, $\mathbf{u} = [u_1^T, \dots, u_\ell^T]^T, \mathbf{f} = [f_1^T, \dots, f_\ell^T]^T \in \mathbb{R}^{\bar{n}\ell}$. The dimensions of $\mathcal{A}$ may make look the all-at-once formulation (3) computationally prohibitive. Nevertheless, the coefficient matrix $\mathcal{A}$ is often extremely structured. This feature, along with other characteristics of the

problem at hand, can be fully exploited to design a very efficient numerical linear algebra phase to the point that the all-at-once scheme oftentimes outperforms the time-marching approach (2).

We leverage this structure to reformulate (3) in terms of a single Sylvester matrix equation of the form

$$K\mathbf{U} + M\mathbf{U}B = \mathbf{F}, \tag{4}$$

where $K, M \in \mathbb{R}^{\bar{n}\times\bar{n}}$ are the spatial stiffness and mass matrices, respectively, $B \in \mathbb{R}^{\ell}$ is the discrete time operator stemming from the adopted time integrator, and $\mathbf{U} = [u_1, \ldots, u_\ell], \mathbf{F} = [f_1, \ldots, f_\ell] \in \mathbb{R}^{\bar{n}\times\ell}$. The matrix equation formulation (4) of the discrete problem naturally encodes the separability of the spatial and time derivatives of the underlying differential operator and the tensorized nature of the cylinder $\Omega \times (0, T]$ in (1). We show that this lets us employ different strategies to deal with the spatial and time components of the algebraic problem and combine them in a very efficient solution procedure. In particular, state-of-the-art projection techniques are proposed to tackle the spatial operator while the circulant-plus-low-rank structure of the time discrete operator is exploited to derive effective solution schemes [1]. The resulting algorithm is able to efficiently solve problems with a tremendous number of degrees of freedom while maintaining a low storage demand as we will illustrate in several numerical results.

References

[1] D. Palitta. Matrix Equation Techniques for Certain Evolutionary Partial Differential Equations. *J Sci Comput*, 87 (99), 2021

12:00–12:30 Matrix-oriented discretizations for evolutionary PDEs and applications
Ivonne Sgura (Università del Salento, Italy)

Abstract. Among evolutionary PDEs we focus on reaction-diffusion systems where the so-called *Turing instability* can give rise to a rich scenario of spatial patterns attained as stationary solutions. To capture the morphological features of Turing patterns is a challenging task because both very fine space discretization and longtime integration are needed. We show that, after semi-discretization in space, the Kronecker structure of the diffusion matrix can be exploited to build matrix-oriented (MO) versions of some classical time integrators. In particular, we consider finite differences on square domains and classical Lagrangian FEM on more general domains and special surfaces. In the first case, the fully discrete problem is reformulated as a sequence of *Sylvester matrix equations*, that we solve by the *reduced approach* in the associated spectral space [1]. In the second case, at each time step we solve a *multiterm Sylvester matrix equation* by the matrix form of the Preconditioned Conjugate Gradient (MO-PCG) [2].
As an application, we consider Turing pattern approximation in the DIB morphochemical PDE system for battery modeling [3] and we present encouraging results wrt the classical vector approach (solving large sparse linear systems) in terms of execution times and memory storage. Moreover, we show that the MO approach based on the *reduced method* can be extended to high order diffusion PDEs like the Cahn-Hilliard equation (4th order) among phase-field models.

References

[1] M.C. D'Autilia, I. Sgura, V. Simoncini, *Matrix-oriented discretization methods for reaction-diffusion PDEs: comparisons and applications*, CAMWA 79, (2020) 2067–2085, doi.org/10.1016/j.camwa.2019.10.020

[2] M. Frittelli, I. Sgura - *Matrix-oriented FEM formulation for reaction-diffusion PDEs on a large class of 2D domains*, Applied Numerical Mathematics (APNUM) 200 (2024) 286–308, doi.org/10.1016/j.apnum.2023.07.010

[3] D. Lacitignola, B. Bozzini, M. Frittelli, I. Sgura, *Turing pattern formation on the sphere for a morphochemical reaction-diffusion model for electrodeposition*, Comm Nonlin Sci Numer Simul, 48 (2017) 484–508, doi.org/10.1016/j.cnsns.2017.01.008

12:30–13:00 Low rank methods for Turing pattern approximation in reaction-diffusion PDE systems

Alessandro Alla (Università Ca' Foscari Venezia, Italy)

Abstract. In this talk we will discuss two low rank methods for the numerical approximation of Turing patterns, that are stationary solutions of reaction-diffusion PDE (RD-PDE) systems by means of Proper Orthogonal Decomposition (POD) and Dynamic Mode Decomposition (DMD). Both techniques present inaccurate approximations, therefore we will introduce two novel algorithms that aim at stabilizing the studied problem. In the first part of the talk we focus on the stabilization of the POD-DEIM technique. We show that solutions of surrogate models built by classical POD-DEIM exhibit an unstable error behaviour over the dimension of the reduced space. To overcome this drawback, we add a correction term that provides missing information to the reduced model and we apply the POD-DEIM technique to the corrected model. To further improve the computational efficiency, we propose an adaptive version of this algorithm in time that accounts for the peculiar dynamics of the RD-PDE in presence of Turing instability.

In the second part we show that DMD does not provide accurate approximation for datasets describing oscillatory dynamics, like spiral waves, relaxation oscillations and spatio-temporal Turing instability. Inspired by the classical "divide and conquer" approach, we propose a piecewise version of DMD (pDMD) to overcome this problem. The main idea is to split the original dataset in N submatrices and then apply the exact (randomized) DMD method in each subset of the obtained partition. We describe the pDMD algorithm in detail and we introduce some error indicators to evaluate its performance when N is increased.

Throughout the talk, we will show the effectiveness of the proposed methods in terms of accuracy and computational cost for a selection of RD systems.

References

[1] A. Alla, A. Monti, I. Sgura, *Adaptive POD-DEIM correction for Turing pattern approximation in reaction-diffusion PDE systems*, Journal of Numerical Mathematics, 2023, **31**, 205–229.

[2] A. Alla, A. Monti, I. Sgura, *Piecewise DMD for oscillatory and Turing spatio-temporal dynamics*, Computers and Mathematics with Applications, 2024, **160**, 108–124

14:30–15:00 Polynomial Approximation for Nonlinear Model Reduction by Moment Matching

Carlos Doebeli (Imperial College London, UK)

Abstract. Many physical applications involve modeling the state of a complex system as a high dimensional system of ODEs with algebraic constraints. The high dimensionality of these models introduces significant computational challenges in trying to analyse the behaviour of the system, which is often known as the "curse of dimensionality". One way to alleviate the computational complexity of such problems is by using reduced order modeling to create a simpler system that preserves important characteristics.

This talk focuses on nonlinear model reduction using the technique of moment matching for nonlinear dynamical systems whose input is generated by a signal generator system. Under certain assumptions, there exists an invariant manifold associated with the solution of a "Sylvester-like" PDE which can then be related to the steady-state response of the system. This talk introduces a procedure to numerically approximate the solutions to the invariance equations that arise in moment matching. The Galerkin residual method is employed to find an approximate solution to the invariance equation using a Newton iteration on the coefficients of a monomial basis expansion of the solution. These approximate solutions to the invariance equation can then be used to construct reduced order models.

15:00–15:30 Residual Data-Driven Variational Multiscale Reduced Order Models for Convection-Dominated Flows

Birgul Koc (University of Seville, Spain)

Abstract. We investigate the modeling of sub-scale components of proper orthogonal decomposition reduced order models (POD-ROMs) of convection-dominated flows. We propose ROM closure models that depend on the ROM residual. We illustrate the new residual-based data-driven ROM closure within the variational multi-scale (VMS) framework and investigate it in the numerical simulation of a one-dimensional parameter-dependent convection-dominated problem and two-dimensional time-dependent advection-diffusion-reaction problem and flow past a cylinder. Our numerical investigation shows that the new residual-based data-driven VMS-ROM is more accurate than both the coefficient-based data-driven ROMs and the standard Galerkin ROM.

References

[1] Mou, C., Koc, B., San, O., Rebholz, L. G., & Iliescu, T. (2021). *Data-driven variational multiscale reduced order models.* Computer Methods in Applied Mechanics and Engineering, 373, 113470.

[2] Koc, B., Chacón, T., & Iliescu, T. (2022). *Residual Data-Driven Variational Multiscale Reduced Order Models for Parameter Dependent Problems.* arXiv preprint arXiv :2208.00059. https://arxiv.org/abs/2208.00059.

15:30–16:00 Structure-Preserving Learning of High-Dimensional Lagrangian Systems

Boris Kramer (University of California San Diego, USA)

Abstract. Lagrangian mechanics is a foundational modeling approach in diverse areas such as structural mechanics, aerospace engineering, wave propagation, and soft robotics. The state equations of these systems are derived from by applying the Euler-Lagrange equations to the Lagrangian energy function, which is specified by the difference of kinetic and potential energy. This results in a highly structured state space and model equations, which are, however, too high-dimensional for fast simulation and control. This talk gives an overview of a few recently developed approaches for learning structure-preserving reduced-order models.

First, we discuss Lagrangian Operator Inference (LOPINF), a data-driven reduced-order model learning framework that respects the Lagrangian structure and learns the associated linear operators for the kinetic and potential energies. We demonstrate the wide applicability and effectiveness of the approach an Euler-Bernoulli beam model from structural mechanics, the sine-Gordon (nonlinear) wave equation, and two large-scale discretization of a soft robot fishtail with 251,000 and 779,232 degrees of freedom, respectively. With much fewer training data than non-structured learning methods require, the Lagrangian learned reduced models provides stability and robustness, and have long-term predictive accuracy.

Second, we present an extension of LOPINF that adds structure-preserving machine learning to learn the nonlinear terms in the Lagrangian models. Through a two-step approach, we first learn the best-fit linear Lagrangian ROM via LOPINF and then use a structure-preserving machine learning method to learn nonlinearities in the reduced space. The proposed approach can learn a structure-preserving nonlinear ROM purely from data, unlike the previously introduced Lagrangian Operator Inference approach that required knowledge about the mathematical form of nonlinear terms (if used on nonlinear models). From a machine learning perspective, it accelerates the training of the structure-preserving neural network by providing an informed prior (i.e., the linear Lagrangian ROM structure), and it reduces the computational cost of the network training by operating on the reduced space. The method is first demonstrated on two simulated examples: a conservative nonlinear rod model and a two-dimensional nonlinear membrane with nonlinear internal damping. Finally, the method is demonstrated on an *experimental dataset* consisting of digital image correlation measurements taken from a lap-joint beam structure from which a predictive model is learned that captures amplitude-dependent frequency and damping characteristics accurately. The numerical results demonstrate that the proposed approach yields

generalizable nonlinear ROMs that exhibit bounded energy error, capture the nonlinear characteristics reliably, and provide accurate long-time predictions outside the training data regime.

We highlight that the proposed approaches exploit knowledge of the structure of the governing equations (but not their discretization) to define the form and parametrization of a Lagrangian ROM which can then be learned from projected snapshot data. The method does not require access to FOM operators or computer code.

References

[1] H. Sharma, B. Kramer, *Preserving Lagrangian structure in data-driven reduced-order modeling of large-scale mechanical systems*, Physica D: Nonl. Phenomena, 462, p. 134128, 2024.

[2] H. Sharma, D. Najera-Flores, M. Todd, B. Kramer, *Lagrangian operator inference enhanced with structure-preserving machine learning for nonintrusive model reduction of mechanical systems*, Computer Methods in Applied Mechanics and Engineering, Vol. 423, 116865, 2024.

16:00–16:30 Flexible domain decomposition-based couplings of conventional and data-driven models via the Schwarz alternating method

Irina Tezaur (Sandia National Laboratories, Livermore, CA, USA.)

Abstract. This talk will describe some recent advancements in developing a rigorous mathematical framework for the domain decomposition-based coupling of arbitrary combinations of first-principles numerical methods (i.e., full order models or FOMs) with data-driven models (i.e., projection-based reduced order models or ROMs) under the flexible Heterogeneous Numerical Methods (fHNM) project at Sandia National Laboratories. Specifically, I will present a recent extension of the Schwarz alternating method [1,2] that enables the creation of FOM-ROM and ROM-ROM couplings from nonlinear monolithic problems [3]. This method works by performing an overlapping or non-overlapping domain decomposition (DD) of the physical domain, and solving a sequence of problems on these subdomains, with information propagating through carefully-constructed transmission conditions on the subdomain boundaries. The solution in each subdomain may be characterized by much simpler, localized dynamics, and hence more easily modeled and solved. We will showcase recent results obtained by implementing this method in the open-source Pressio demo-apps library, which demonstrate that the Schwarz alternating method is capable of delivering stable and accurate hybrid models when applied to advection-dominated fluid flow problems with moving shocks, for which monolithic approaches are prohibitively expensive or grossly inaccurate. We will additionally demonstrate that online CPU-time gains are achievable through an implementation of the additive variant of the Schwarz alternating method, which admits more parallelism by solving all subdomain-local problems simultaneously on different processes/threads. Time-permitting, we will describe some ongoing work aimed at developing an automated learning algorithm leveraged to select "optimal" (i.e, accuracy and efficiency maximizing) DDs and ROM vs. FOM placements/assignments.

References

[1] A. Mota, I. Tezaur, C. Alleman. The Schwarz alternating method in solid mechanics", Comput. Meth. Appl. Mech. Engng. 319 19-51, 2017.

[2] A. Mota, I. Tezaur, G. Phlipot. The Schwarz alternating method for dynamic solid mechanics, Int. J. Numer. Meth. Engng. 123 (21) 5036-5071, 2022.

[3] J. Barnett, I. Tezaur, A. Mota. The Schwarz alternating method for the seamless coupling of nonlinear reduced order models and full order models, ArXiV pre-print, 2022. https://arxiv.org/abs/2210.12551

17:00–17:30 Dynamical Low-Rank Approximations for parametric infinite horizon optimal control problems

Maria Strazzullo (Politecnico di Torino, Italy)

Abstract. Optimal control is a strategy to enhance the performance of complex dynamical systems by minimizing a predefined cost functional, with the overarching goal of steering the system dynamics towards a desired stable configuration. However, its applicability is still limited due to

high-dimensional state spaces, especially in parametric settings. To address these challenges, we exploit the Dynamical Low-Rank Approximation (DLRA) methodology for the efficient and precise resolution of high-dimensional feedback control problems, relying on State Dependent Riccati Equations (SDREs) for infinite horizon optimal control. DLRA not only furnishes a compact, low-dimensional representation of the system dynamics but also dynamically evolves alongside the problem, ensuring increased accuracy and enabling effective real-time control settings. In this work, we propose two novel algorithms for tackling feedback control problems:

- Monolithic DLRA (mDLRA-SDREs) for SDRE control, where the SDRE is solved for all parametric instances simultaneously.
- Cascade DLRA (cDLRA-SDREs) for SDRE control, where each SDRE Riccati solution is obtained by leveraging the solution of the previous parameter, enhancing convergence and resulting in a faster evolution of the controlled low-rank system.

We validate the robustness and effectiveness of the algorithms through nonlinear test cases, such as optimal control problems governed by Burgers' equation and the Allen–Cahn equation. Moreover, we compare the DLRA-based control strategies with standard global Proper Orthogonal Decomposition model order reduction, showing advantages in terms of speed-up and accuracy with respect to the ground truth solution.

17:30–18:00 Optimal bounds for POD approximations of infinite horizon control problems based on time derivatives

Julia Novo (University Autónoma of Madrid, Spain)

Abstract. In this talk we consider the numerical approximation of infinite horizon problems via the dynamic programming approach. The value function of the problem solves a Hamilton-Jacobi-Bellman (HJB) equation that is approximated by a fully discrete method. It is known that the numerical problem is difficult to handle by the so called curse of dimensionality. To mitigate this issue we apply a reduction of the order by means of a new proper orthogonal decomposition (POD) method based on time derivatives. We carry out the error analysis of the method using recently proved optimal bounds for the fully discrete approximations. Moreover, the use of snapshots based on time derivatives allow us to bound some terms of the error that could not be bounded in a standard POD approach. Some numerical experiments show the good performance of the method in practice.

18:00–18:30 Using Empirical Tensor Train Approximation for Solving High-Dimensional Optimal Control Problems

Mathias Oster (RWTH Achen, Germany)

Abstract. We display two approaches to solve finite horizon optimal control problems. First we solve the Bellman equation numerically by employing the Policy Iteration algorithm. Second, we introduce use open loop methods to learn the value function. To overcome computational infeasability we use tensor trains and multi-polynomials, together with high-dimensional quadrature, e.g. Monte-Carlo. Furthermore, we compare the tensor methods to neural networks and kernel approaches. By controlling a destabilized version of viscous Burgers and a diffusion equation with unstable reaction term numerical evidence is given.

References

[1] Oster, Mathias and Sallandt, Leon and Schneider, Reinhold, *Approximating Optimal feedback Controllers of Finite Horizon Control Problems Using Hierarchical Tensor Formats*, SIAM Journal on Scientific Computing, 40, 2022, B746-B770.

[2] Mathias Oster and Luca Saluzzi and Tizian Wenzel, *A comparison study of supervised learning techniques for the approximation of high dimensional functions and feedback control*, arXiv (2024).

E–mail: alessandro.alla@unive.it.