

A Predictor-Corrector Strategy for Adaptivity in Dynamical Low-Rank Approximations*

Cory D. Hauck[†] Stefan R. Schnake[‡]

September 9, 2022

Abstract

In this paper, we present a predictor-corrector strategy for constructing rank-adaptive dynamical low-rank approximations (DLRAs) of matrix-valued ODE systems. The strategy is a compromise between (i) low-rank step-truncation approaches that alternately evolve and compress solutions and (ii) strict DLRA approaches that augment the low-rank manifold using subspaces generated locally in time by the DLRA integrator. The strategy is based on an analysis of the error between a forward temporal update into the ambient full-rank space, which is typically computed in a step-truncation approach before re-compressing, and the standard DLRA update, which is forced to live in a low-rank manifold. We use this error, without requiring its full-rank representation, to correct the DLRA solution. A key ingredient for maintaining a low-rank representation of the error is a randomized singular value decomposition (SVD), which introduces some degree of stochastic variability into the implementation. The strategy is formulated and implemented in the context of discontinuous Galerkin spatial discretizations of partial differential equations and applied to several versions of DLRA methods found in the literature, as well as a new variant. Numerical experiments comparing the predictor-corrector strategy to other methods demonstrate robustness to overcome short-comings of step truncation or strict DLRA approaches: the former may require more memory than is strictly needed while the latter may miss transient solution features that cannot be recovered. The effect of randomization, tolerances, and other implementation parameters is also explored.

1 Introduction

Low-rank methods are a class of numerical tools that are used to represent high-dimensional data with low-dimensional objects. In many scenarios, they offer substantial memory savings in the storage and transmission of data which can be leveraged to construct efficient models of high-dimensional systems [12]. Two common low-rank methods for dynamical systems are step-truncation methods and dynamical low-rank approximation (DLRA).

*Notice: This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a non-exclusive, paid up, irrevocable, world-wide license to publish or reproduce the published form of the manuscript, or allow others to do so, for U.S. Government purposes. The DOE will provide public access to these results in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

[†]Mathematics in Computation Section, Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA and Mathematics Department, University of Tennessee, Knoxville, TN 37996, USA (hauckc@ornl.gov).

[‡]Mathematics in Computation Section, Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA (schnakesr@ornl.gov).

Step truncation methods are a generic tool for sparse approximation of dynamical systems [31, 30, 14, 4]. These methods update a sparse initial condition in the ambient space and then compress the solution back to a sparse representation through a truncation procedure that removes elements considered unnecessary for maintaining a certain level of accuracy. Low-rank representations are an effective way to enforce a sparse structure.

Dynamical low-rank approximations solve an approximate model that, in the absence of adaptivity, maintains a fixed rank for all time. First introduced as the Dirac-Freknel-McLachlan variational principle in the 1930s [6, 11], DLRA methods were revived in the 2000s for the simulation of high-dimensional quantum systems and biological cellular systems [18, 1, 24]. Even more recently, DLRA methods have been applied to kinetic equations [10, 28, 29, 8, 7, 5], hyperbolic problems with uncertainty [23], and even neural network training [32]. DLRA methods evolve a tensor-valued ODE on a rank- r manifold \mathcal{M}_r by projecting the time derivative onto the tangent space of \mathcal{M}_r . This forces the DLRA to be rank- r for all time and can make certain high-dimensional problems computationally tractable [18].

In [21], DLRA methods were rigorously analyzed showing that they are robust and accurate. Additionally, in [21] the authors show the low-rank factors of the DLRA solution solve a coupled system, but this system is often numerically unstable to march forward in time unless an unreasonably small timestep is taken [25]. Recent work has produced several numerically stable and accurate DLRA temporal integrators [25, 19, 20, 3]. Extensions of these methods to asymptotically preserving or conservative schemes have also been of interest [5, 9, 27].

The current work focuses on adaptivity in the DLRA framework. In practice, the rank required to accurately approximate a dynamical system may change in time, due to non-linear phenomena such as filamentation and turbulence, initial and boundary layers, contributions from sources, and coordinate-based numerical artifacts. Numerical methods that allow the rank of the solution to change in time are called *rank-adaptive* methods and are vital to efficiently resolve these dynamics. Several rank-adaptive methods have recently been developed [2, 14, 4, 17], and, in general, these methods take one of two strategies. Some stay completely in the DLRA regime, where they are prone to modelling error by the low-rank projection [2]. Others construct a full-rank object that is accurate and compress back to a low-rank object [14, 13], but this can be inefficient if the rank of the solution does not change too much.

In the current paper, we develop a rank-adaptive predictor-corrector (RAPC) strategy for existing DLRA methods [3, 20] that strikes a balance between the advantages of step truncation methods and a strict DLRA-based approach. The DLRA serves as the predictor and its error to the full-rank forward Euler update serves as the corrector. We use a randomized SVD [15] to efficiently approximate the corrector by a low-rank object, and then modify the rank of the predictor utilizing low-rank information from the corrector. This yields a strategy that both allows for information outside of the low-rank space to enter the solution and does not explicitly construct a high-rank object. The rank-adaptive tensor scheme of [4] follows a similar strategy as the RAPC strategy, but the estimator is only used to determine the possible rank at the next time step; the predictor-corrector strategy uses the corrector information to fill in the rank deficiency of the DLRA update if need be. Numerical results illustrated in this paper show that the predictor-corrector strategy nearly matches the rank of the rank-adaptive DLRA method of [2]. While the RAPC method uses more memory than rank-adaptive DLRA methods, we present an example where the DLRA regime is not rich enough to accurately approximate a solution while the RAPC approach is still accurate.

The paper is organized as follows. In Section 2, we introduce preliminary notation, an illustrative two-dimensional solid-body rotation problem, and a standard discontinuous Galerkin discretization that serves as a “full-rank” numerical reference. In Section 3, we introduce separability of a bilinear form, show how low-rank methods can take advantage of separability, and finally show that the 2D

solid-body problem is a sum of separable forms. In Section 4, we summarize the DLRA method, motivate several numerical integrators recently developed in the DLRA framework, and construct a new integrator that is accurate and has favorable properties with respect to the predictor-corrector strategy. In Section 5, we list several rank-adaptive integrators and give a detailed motivation and construction of the predictor-corrector strategy. In Section 6, we compare the predictor-corrector strategy against other rank-adaptive integrators on several related test problems. Additionally, we test the effects of the randomized SVD on the accuracy of the rank-adaptive DLRA solution.

2 Notation and Preliminaries

2.1 Notation

For $z \in \{x, y\}$, let $\Omega_z = (-L, L) \subset \mathbb{R}$ for some $L > 0$ be an open bounded domain with boundary $\partial\Omega_z = \{-L, L\}$. Denote the computational domain $\Omega = \Omega_x \times \Omega_y$ and let $(\cdot, \cdot)_D$ denote the standard L^2 inner product on $L^2(D)$ for any domain D .

Given a scalar mesh parameter $h_z > 0$, let $\mathcal{T}_{z,h} := \mathcal{T}_{z,h_z}$ be a mesh on Ω_z with interior skeleton $\mathcal{E}_{z,h}^I := \mathcal{E}_{z,h_z}^I$. Given an edge $e \in \mathcal{E}_{z,h}^I$ with $e = T^+ \cap T^-$ for $T^+, T^- \in \mathcal{T}_{z,h}$, let n^\pm be the outward normal vector on T^\pm . Given a scalar function q defined on $\partial T^+ \cap e$ and $\partial T^- \cap e$ define the average and jump of q , respectively, by

$$\{\{q\}\} = \frac{1}{2}(q|_{T^+} + q|_{T^-}) \quad \text{and} \quad \llbracket q \rrbracket = q|_{T^+} n^+ + q|_{T^-} n^-. \quad (2.1)$$

Let $V_{z,h} := V_{z,h_z}^{k_z}$ be standard the discontinuous Galerkin Lagrange finite element space:

$$V_{z,h_z}^{k_z} = \{q \in L^2(\Omega_z) : q|_T \in \mathbb{Q}_{k_z}(T) \quad \forall T \in \mathcal{T}_{z,h}\}, \quad (2.2)$$

where $\mathbb{Q}_{k_z}(T)$ is the set of all polynomials on T with degree less than or equal to k_z in each coordinate. Let $V_h := V_{x,h} \otimes V_{z,h}$. Additionally let $\langle \cdot, \cdot \rangle_{\mathcal{E}_{z,h}^I} := \sum_{e \in \mathcal{E}_{z,h}^I} \langle \cdot, \cdot \rangle_e$.

Given two $m \times n$ matrices A, B , let

$$(A, B)_F = \text{tr}(B^T A) = \sum_{i=1}^m \sum_{j=1}^n a_{ij} b_{ij}, \quad (2.3)$$

be the Frobenius inner product with induced norm $\|A\|_F = \sqrt{(A, A)_F}$. Denote by $\{\sigma_i(A)\}_{i=1}^{\min\{m,n\}}$ the singular values of A , arranged in a non-increasing order with respect to i .

We utilize MATLAB-style slicing notation for matrices. If A is an $m \times n$ matrix, $I = \{i_1, \dots, i_k\}$ is a subset of $\{1, \dots, m\}$ and $J = \{j_1, \dots, j_l\}$ is a subset of $\{1, \dots, n\}$, then $\tilde{A} := A(I, J)$ is a $k \times l$ matrix defined by $\tilde{A}_{kl} = A_{i_k, j_l}$. The notation $i : k \subseteq \{1, \dots, n\}$ is a shorthand for the subset $\{i, i+1, \dots, k\}$, $:$ is taken to mean $1:n$, and $\text{diag}(A_1, A_2, \dots, A_N)$ is a block diagonal matrix.

2.2 The PDE and its Discretization

The techniques presented in this paper can be applied to many advective or diffusive PDEs; however, we will focus on a solid-body rotation problem as a model for the topics covered in Section 3. Other operators of physical interest are discussed in Section 6.

The 2D solid-body rotation is defined by the following system:

$$\frac{\partial u}{\partial t} - y \cdot \nabla_x u + x \cdot \nabla_y u = s, \quad (x, y) \in \Omega, \quad t > 0; \quad (2.4a)$$

$$u(x, y, t) = g_x(x, y, t), \quad (x, y) \in \partial\Omega_x^-, \quad t > 0; \quad (2.4b)$$

$$u(x, y, t) = g_y(x, y, t), \quad (x, y) \in \partial\Omega_y^-, \quad t > 0; \quad (2.4c)$$

$$u(x, y, 0) = u_0(x, y), \quad (x, y) \in \Omega, \quad (2.4d)$$

where $s(\cdot, \cdot, t) \in L^2(\Omega)$ is a source term and g_z is the inflow data. Here $\partial\Omega_z^\pm$ is the outflow (+) and inflow (-) boundaries on $\partial\Omega_z$, given by

$$\partial\Omega_x^\pm = \{(x, y) \in \partial\Omega_x \times \Omega_y : \pm y \cdot n(x) < 0\} \quad \text{and} \quad \partial\Omega_y^\pm = \{(x, y) \in \Omega_x \times \partial\Omega_y : \mp x \cdot n(y) < 0\}. \quad (2.5)$$

We discretize (2.4a) by the standard discontinuous Galerkin method with upwind fluxes. The discretized problem is as follows: Find $u_h(\cdot, \cdot, t) \in V_h$ such that

$$\left(\frac{\partial u_h}{\partial t}, q_h \right)_\Omega = \mathcal{A}(u_h, q_h) + \mathcal{G}(q_h, t) + (s, q_h)_\Omega \quad \forall q_h \in V_h, \quad (2.6)$$

where $\mathcal{A} : V_h \times V_h \rightarrow \mathbb{R}$ and $\mathcal{G} : V_h \times \mathbb{R}^+ \rightarrow \mathbb{R}$ are given by

$$\begin{aligned} \mathcal{A}(w_h, q_h) &= -(-yw_h, \nabla_x q_h)_\Omega + \langle -y\{\{w_h\}\} + \frac{1}{2}|y|[[w_h]], [[q_h]] \rangle_{\mathcal{E}_{x,h}^I \times \Omega_y} \\ &\quad + \langle -yw_h, nq_h \rangle_{\partial\Omega_x^+} \\ &\quad - (xw_h, \nabla_y q_h)_\Omega + \langle x\{\{w_h\}\} + \frac{1}{2}|x|[[w_h]], [[q_h]] \rangle_{\Omega_x \times \mathcal{E}_{y,h}^I} \\ &\quad + \langle xw_h, nq_h \rangle_{\partial\Omega_y^+}, \end{aligned} \quad (2.7)$$

$$\mathcal{G}(q_h, t) = \langle yg_x(\cdot, \cdot, t), nq_h \rangle_{\partial\Omega_x^-} + \langle -xg_y(\cdot, \cdot, t), nq_h \rangle_{\partial\Omega_y^-}. \quad (2.8)$$

3 Conversion and Evaluation of Matrix Forms

3.1 Low-rank savings

To extract low-rank features of u_h , we convert the coefficients of u_h to a matrix form. Let $\{\varphi_i(x)\}_{i=1}^m$ and $\{\psi_j(y)\}_{j=1}^n$ be a basis for $V_{x,h}$ and $V_{y,h}$, respectively, that are orthonormal in $L^2(\Omega_x)$ and $L^2(\Omega_y)$, respectively. Then u_h has an expansion

$$u_h(x, y, t) = \sum_{i=1}^m \sum_{j=1}^n u_h^{ij}(t) \varphi_i(x) \psi_j(y) \quad (3.1)$$

where the coefficients u_h^{ij} are given by

$$u_h^{ij} = (u_h, \varphi_i \psi_j)_\Omega. \quad (3.2)$$

These coefficients may be assembled into the $m \times n$ matrix $U_h(t)$ where $(U_h)_{ij} = u_h^{ij}$, in which case the expansion for u_h in (3.1) can be written as

$$u_h(x, y, t) = \Phi(x) U_h(t) \Psi(y)^T \quad (3.3)$$

where $\Phi(x) = [\varphi_1(x), \dots, \varphi_m(x)]$ and $\Psi(y) = [\psi_1(y), \dots, \psi_n(y)]$. For any function $q_h \in V_h$, we use the capitalized version Q_h denote its matrix valued coordinate representation derived as in (3.3).

This paper is largely devoted the case when U_h is a low-rank matrix.

Definition 3.1. A matrix $U \in \mathbb{R}^{m \times n}$ is said to be rank- r if the column space

$$\text{Col}(U) = \{Uz : z \in \mathbb{R}^n\}. \quad (3.4)$$

has dimension r . The manifold of all such matrices is denoted \mathcal{M}_r .

The following is a standard result from linear algebra

Proposition 3.1. If a matrix $U \in \mathbb{R}^{m \times n}$ is rank- r , then $r \leq \min\{m, n\}$. Moreover, U has a decomposition of the form $U = CSD^T$ where $C \in \mathbb{R}^{m \times r}$ and $D \in \mathbb{R}^{n \times r}$ are orthogonal matrices and $S \in \mathbb{R}^{r \times r}$ is non-singular.

If $U_h(t)$ is rank- r with decomposition $U_h(t) = C(t)S(t)D(t)^T$, then according to (3.3),

$$u_h(x, y, t) = [C(t)\Phi(x)]S(t)[D(t)\Psi(y)]^T. \quad (3.5)$$

The matrices C and D map the time-independent bases functions in $\Phi(x)$ and $\Psi(y)$ to the time-dependent *low-rank basis functions* $\xi_k(x, t) := C(:, k)(t)\Phi(x)$ and $v_k(y, t) := D(:, k)(t)\Psi(y)$. Thus if $X(x, t) = [\xi_1(x, t), \dots, \xi_r(x, t)] = C(t)\Phi(x)$ and $Y(y, t) = [v_1(y, t), \dots, v_r(y, t)] = D(t)\Psi(y)$, then

$$u_h(x, y, t) = \sum_{k,l=1}^r S_{kl}(t)\xi_k(x, t)v_l(y, t) = X(x, t)S(t)Y(y, t)^T. \quad (3.6)$$

The matrix S glues the low-rank basis functions together to reconstruct u_h , and for $r \ll \min\{m, n\}$, it is much cheaper to store the low-rank factors C , D , and S , which together require $mr + r^2 + nr$ floating point values, than U_h directly, which requires mn floating point values.

Dynamic low-rank approximation (DLRA) [21] methods update X , S , and Y in tandem, but since Φ and Ψ are independent of t , tracking the evolution of X and Y in time is equivalent to tracking the evolution of C and D . While DLRA methods on paper can be applied to a generic matrix-valued ODE

$$\frac{\partial U}{\partial t} = F(U), \quad U \in \mathbb{R}^{m \times n}, \quad F : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}, \quad (3.7)$$

the memory and computational savings provided by DLRA can be realized only if the operator F respects the low-rank framework—that is, only if $F(CSD^T)$ can be evaluated without reconstructing the coefficients of U_h by direct multiplication of the low-rank factors C , S , and D . Bilinear forms that can separate their actions on the x and y domains tend to respect the low-rank structure. This notion of separability is codified in Definition 3.2.

Definition 3.2. Let \mathcal{B} be a bilinear form on $V_h \times V_h$. Given $N \in \mathbb{N}$ we say that \mathcal{B} is N -separable with respect to the basis $\{\varphi_i\psi_j\}$ if for every $\kappa \in \mathbb{N}$ with $1 \leq \kappa \leq N$ there exists bilinear forms $\mathcal{B}_\kappa : V_h \times V_h \rightarrow \mathbb{R}$, $\mathcal{B}_{\kappa,x} : V_{x,h} \times V_{x,h} \rightarrow \mathbb{R}$, $\mathcal{B}_{\kappa,y} : V_{y,h} \times V_{y,h} \rightarrow \mathbb{R}$ such that

$$\mathcal{B}(\varphi_i\psi_j, \varphi_k\psi_l) = \sum_{\kappa=1}^N \mathcal{B}_\kappa(\varphi_i\psi_j, \varphi_k\psi_l) = \sum_{\kappa=1}^N \mathcal{B}_{\kappa,x}(\varphi_i, \varphi_k)\mathcal{B}_{\kappa,y}(\psi_j, \psi_l). \quad (3.8)$$

We call \mathcal{B}_κ the terms of \mathcal{B} , and call $\mathcal{B}_{\kappa,x}$ and $\mathcal{B}_{\kappa,y}$ the factors of \mathcal{B}_κ .

We can exploit separability to evaluate N -separable bilinear forms by the following proposition.

Proposition 3.2. *Given $N \in \mathbb{N}$, let \mathcal{B}_κ be an N -separable bilinear form with expansion given in Definition 3.2. Let $A_\kappa \in \mathbb{R}^{n \times n}$ and $B_\kappa \in \mathbb{R}^{m \times m}$ be matrices with elements $(A_\kappa)_{ik} = \mathcal{B}_{\kappa,x}(\varphi_k, \varphi_i)$ and $(B_\kappa)_{jl} = \mathcal{B}_{\kappa,y}(\psi_l, \psi_j)$ for all $1 \leq \kappa \leq N$. Then for any $w_h, q_h \in V_h$,*

$$\mathcal{B}(w_h, q_h) = \sum_{\kappa=1}^N (A_\kappa W_h B_\kappa^T, Q_h)_F, \quad (3.9)$$

where W_h and Q_h are the matrix coefficient representations of w_h and q_h respectively.

Proof. Let $w_h = \sum_{i=1}^m \sum_{j=1}^n (W_h)_{ij} \varphi_i \psi_j$ and $q_h = \sum_{k=1}^m \sum_{l=1}^n (Q_h)_{kl} \varphi_k \psi_l$. Then for any $1 \leq \kappa \leq N$,

$$\begin{aligned} \mathcal{B}_\kappa(w_h, q_h) &= \sum_{ijkl} (W_h)_{ij} (Q_h)_{kl} \mathcal{B}_{\kappa,x}(\varphi_i, \varphi_k) \mathcal{B}_{\kappa,y}(\psi_j, \psi_l) = \sum_{ijkl} (W_h)_{ij} (Q_h)_{kl} (A_\kappa)_{ki} (B_\kappa)_{lj} \\ &= \sum_{kl} (A_\kappa W_h B_\kappa^T)_{kl} (Q_h)_{kl} = (A_\kappa W_h B_\kappa^T, Q_h)_F. \end{aligned} \quad (3.10)$$

Summing (3.10) over κ yields (3.9). The proof is complete. \square

Definition 3.3. *A linear operator $\mathcal{L} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ is called N -separable provided it is of the form*

$$\mathcal{L}W = \sum_{\kappa=1}^N A_\kappa W B_\kappa^T \quad \forall W \in \mathbb{R}^{m \times n}, \quad (3.11)$$

where $A_\kappa \in \mathbb{R}^{m \times m}$ and $B_\kappa \in \mathbb{R}^{n \times n}$ for all $1 \leq \kappa \leq N$.

From Proposition 3.2, Definition 3.3, and the Riesz Representation theorem, the existence of the linear operator $\mathcal{L} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ by

$$(\mathcal{L}W_h, Q_h)_F = \mathcal{B}(w_h, q_h) \quad \forall w_h, q_h \in V_h \quad (3.12)$$

is immediate.

Proposition 3.3 shows that the low-rank structure of W_h allows for memory efficient evaluations of $\mathcal{L}W_h$ and the action $Q \rightarrow (\mathcal{L}W_h)Q$ – both of which do not require assembly of U_h by multiplication of the low-rank factors.

Proposition 3.3. *Suppose W_h has a rank- r decomposition $W_h = CSD^T$ where $C \in \mathbb{R}^{m \times r}$, $S \in \mathbb{R}^{r \times r}$, and $D \in \mathbb{R}^{n \times r}$. Furthermore for $1 \leq \kappa \leq N$, suppose the matrices A_κ and B_κ in (3.11) are sparse with $\mathcal{O}(m)$ and $\mathcal{O}(n)$ elements respectively. Then $\mathcal{L}W_h$ can be stored using $Nr(m+n+r)$ floating point values and computed in $\mathcal{O}(Nr(m+n))$ floating point operations (flops). Moreover, suppose Q is an $n \times q$ matrix. Set $\rho = \max\{r, q\}$. Then the action $Q \rightarrow (\mathcal{L}W_h)Q \in \mathbb{R}^{m \times q}$ requires $\mathcal{O}(Nr\rho(n+m))$ flops. A similar operation count holds for the action $Q \rightarrow Q^T(\mathcal{L}W_h)$ for $Q \in \mathbb{R}^{m \times q}$.*

Proof. By (3.11),

$$\mathcal{L}W_h = \sum_{\kappa=1}^N A_\kappa CSD^T B_\kappa^T = \sum_{\kappa=1}^N (A_\kappa C)S(B_\kappa D)^T. \quad (3.13)$$

Since $A_\kappa C \in \mathbb{R}^{m \times r}$ and $B_\kappa D \in \mathbb{R}^{n \times r}$ for all κ , we can store a copy of $A_\kappa C$, S , and $B_\kappa D$ which requires $r(m+n+r)$ floating point values for each κ and $Nr(m+n+r)$ in total. The sparsity of A_κ and B_κ imply that the operations $A_\kappa C$ and $B_\kappa D$ require $\mathcal{O}(mr)$ and $\mathcal{O}(nr)$ flops respectively.

Thus storage of $\mathcal{L}W_h$ in this way requires $\mathcal{O}(Nr(m+n))$ flops. To evaluate $(\mathcal{L}W_h)Q$, we use (3.13):

$$(\mathcal{L}W_h)Q = \sum_{\kappa=1}^N (A_\kappa CS)(D^T B_\kappa^T Q). \quad (3.14)$$

The matrix $A_\kappa CS \in \mathbb{R}^{m \times r}$ and $D^T B_\kappa^T Q \in \mathbb{R}^{r \times q}$ require $\mathcal{O}(mr + mr^2)$ and $\mathcal{O}(nr + nrq)$ flops to evaluate while their product requires $\mathcal{O}(mrq)$ flops to multiply. Dropping the lower order mr and nr terms, the evaluation of $(A_\kappa CS)(D^T B_\kappa^T Q) \in \mathbb{R}^{m \times q}$ requires $\mathcal{O}r(mr + mq + nq) \leq \mathcal{O}r\rho(m+n)$ flops. Therefore (3.14) implies that the evaluation of $(\mathcal{L}W_h)Q$ has a cost of $\mathcal{O}(Nr\rho(m+n))$. The proof is complete. \square

Remark 3.1.

1. (3.13) is used in [14] in order to store $\mathcal{L}W_h$ in a low memory fashion as compared to the direct storage of the $m \times n$ matrix. Memory savings of this storage technique rely on Nr being small. If either N or r becomes large, then storage by (3.13) becomes impractical.
2. (3.14) shows the memory footprint of $(\mathcal{L}W_h)Q$ is not dependent on N .
3. DLRA methods heavily use (3.14) with Q being one of the basis matrices C or D . In addition, the predictor-corrector based DLRA method introduced in this paper also uses (3.14) with the columns of Q small.

3.2 The low-rank structure of the PDE

It turns out that the bilinear form \mathcal{A} defined in (2.7) is a N -separable. The proof of the following result is given in the Appendix.

Proposition 3.4. *Given any basis $\{\varphi_i \psi_j\}$ on V_h , the bilinear form \mathcal{A} in (2.7) is N -separable with $N = 4$. Moreover, the terms in the decomposition (3.8) are given by*

$$\mathcal{B}_{1,x}(\varphi_i, \varphi_j) = -(\varphi_i, \nabla_x \varphi_k)_{\Omega_x} + \langle \{\{\varphi_i\}\} - \frac{1}{2} \llbracket \varphi_i \rrbracket, \llbracket \varphi_k \rrbracket \rangle_{\mathcal{E}_{x,h}^I} + \langle \varphi_i n, \varphi_k \rangle_{\{x=-L\}} \quad (3.15a)$$

$$\mathcal{B}_{1,y}(\psi_j, \psi_l) = (-y\psi_j, \psi_l)_{\{y>0\}} \quad (3.15b)$$

$$\mathcal{B}_{2,x}(\varphi_i, \varphi_j) = -(\varphi_i, \nabla_x \varphi_k)_{\Omega_x} + \langle \{\{\varphi_i\}\} + \frac{1}{2} \llbracket \varphi_i \rrbracket, \llbracket \varphi_k \rrbracket \rangle_{\mathcal{E}_{x,h}^I} + \langle \varphi_i n, \varphi_k \rangle_{\{x=L\}} \quad (3.15c)$$

$$\mathcal{B}_{2,y}(\psi_j, \psi_l) = (-y\psi_j, \psi_l)_{\{y<0\}} \quad (3.15d)$$

$$\mathcal{B}_{3,x}(\varphi_i, \varphi_k) = (x\varphi_i, \psi_k)_{\{x>0\}} \quad (3.15e)$$

$$\mathcal{B}_{3,y}(\psi_j, \psi_l) = -(\psi_j, \nabla_y \psi_l)_{\Omega_y} + \langle \{\{\psi_j\}\} + \frac{1}{2} \llbracket \psi_j \rrbracket, \llbracket \psi_l \rrbracket \rangle_{\mathcal{E}_{y,h}^I} + \langle \psi_j n, \psi_l \rangle_{\{y=L\}} \quad (3.15f)$$

$$\mathcal{B}_{4,x}(\varphi_i, \varphi_k) = (x\varphi_i, \psi_k)_{\{x<0\}} \quad (3.15g)$$

$$\mathcal{B}_{4,y}(\psi_j, \psi_l) = -(\psi_j, \nabla_y \psi_l)_{\Omega_y} + \langle \{\{\psi_j\}\} - \frac{1}{2} \llbracket \psi_j \rrbracket, \llbracket \psi_l \rrbracket \rangle_{\mathcal{E}_{y,h}^I} + \langle \psi_j n, \psi_l \rangle_{\{y=-L\}} \quad (3.15h)$$

Remark 3.2. *If central fluxes and periodic boundary conditions were used in the DG discretization, then only two terms would be needed: one for $-y \cdot \nabla_x u$ and one for $x \cdot \nabla_y u$. However, more terms are needed to handle the upwind numerical flux and the outflow boundary condition.*

Remark 3.3. $\mathcal{B}_{1,x}$ and $\mathcal{B}_{2,x}$ are a downwind and upwind flux DG discretizations of the operator ∂_x respectively. While $\mathcal{B}_{1,x}$ is a negative semi-definite form due to the downwind numerical flux, $\mathcal{B}_{1,y}$ is also negative semi-definite. Therefore their product is a positive semi-definite form on $V_h \times V_h$. Indeed, using the notation in (3.8), \mathcal{B}_κ is positive semi-definite for $\kappa = 1, 2, 3, 4$.

To handle the other terms in (2.6), let $s_h \in V_h$ be the L^2 projection of the source s in (2.4a) onto V_h , defined by the relation

$$(s_h, q_h)_\Omega = (s, q_h) \quad \forall q_h \in V_h, \quad (3.16)$$

and let S_h be the matrix representation of s_h with coefficients derived using (3.2). We assume the operations $Q \rightarrow S_h Q$ and $Q^T \rightarrow Q^T S_h$ are both efficiently computed. This can be done if the rank of S_h is small. Additionally, let $g_h(t) \in V_h$ represent the action of \mathcal{G} ; that is,

$$(g_h(t), q_h)_\Omega = \mathcal{G}(q_h, t) \quad \forall q_h \in V_h. \quad (3.17)$$

The existence and uniqueness of g_h is given by the Riesz representation theorem. Let G_h be the matrix representation of g_h with coefficients derived using (3.2). We can show that the rank of G_h is at most four by the following proposition, whose proof is in the Appendix.

Proposition 3.5. $G_h(t) = C(t)S(t)D(t)^T$ where $C(t) \in \mathbb{R}^{m \times 4}$, $S(t) \in \mathbb{R}^{r \times r}$ and $D(t) \in \mathbb{R}^{n \times 4}$ for any $t \geq 0$. Additionally, the rank of G_h is at most 4.

In matrix form, (2.6) becomes

$$\frac{\partial U_h}{\partial t} = F(U_h, t) := \mathcal{L}U_h + G_h(t) + S_h(t). \quad (3.18)$$

4 Dynamic Low-Rank Approximation

In this section we give a brief introduction to Dynamic Low-Rank Approximation (DLRA) and several associated temporal integrators. As the content in this section does not depend on the DG discretization, we consider the general ODE

$$\begin{aligned} \frac{\partial U}{\partial t}(t) &= F(U(t), t), \quad t > 0; \\ U(0) &= U_0. \end{aligned} \quad (4.1)$$

Abusing notation, we will suppress the t argument of F and write $F(U, t) = F(U)$.

4.1 A Short Introduction

DLRA was developed in [21] as efficient way to produce a low-rank approximation to $U(t)$ that doesn't require compression at each step of a time integration scheme. This is achieved by evolving the continuous ODE on the manifold \mathcal{M}_r . Indeed if $U_{\text{DLRA}} \in \mathcal{M}_r$ is a rank- r approximation at time t , then one can guarantee $U_{\text{DLRA}} \in \mathcal{M}_r$ for all future time by forcing its time derivative to live in $T_{U_{\text{DLRA}}} \mathcal{M}_r$ — the tangent space of \mathcal{M}_r at U_{DLRA} . The rank- r solution U_{DLRA} is defined by the following evolution equation:

$$\frac{\partial U_{\text{DLRA}}}{\partial t} = \underset{Z \in T_{U_{\text{DLRA}}} \mathcal{M}_r}{\operatorname{argmin}} \|F(U_{\text{DLRA}}) - Z\|_F, \quad (4.2)$$

which is equivalent to the following Galerkin condition:

$$\left(\frac{\partial U_{\text{DLRA}}}{\partial t}, Z \right)_F = (F(U_{\text{DLRA}}), Z)_F \quad \forall Z \in T_{U_{\text{DLRA}}} \mathcal{M}_r; \quad (4.3)$$

which is again equivalent to

$$\frac{\partial U_{\text{DLRA}}}{\partial t} = P_{T_{U_{\text{DLRA}}}\mathcal{M}_r} F(U_{\text{DLRA}}), \quad (4.4)$$

where $P_{T_{U_{\text{DLRA}}}\mathcal{M}_r}$ is the orthogonal projection onto $T_{U_{\text{DLRA}}}\mathcal{M}_r$.

Any implementation of (4.2)-(4.4) should leverage the low-rank structure of U_{DLRA} in order to be memory efficient. If $U_{\text{DLRA}} = C(t)S(t)D(t)^T$, then $T_{U_{\text{DLRA}}}\mathcal{M}_r$ is given by [21]:

$$T_{U_{\text{DLRA}}}\mathcal{M}_r = \{\delta CSD^T + C\delta SD^T + CS\delta D^T : \delta C \in \mathbb{R}^{m \times r} \text{ with } \delta C^T C = 0, \\ \delta S \in \mathbb{R}^{r \times r}, \delta D \in \mathbb{R}^{n \times r} \text{ with } \delta D^T D = 0\}, \quad (4.5)$$

where the gauge conditions $\delta C^T C = 0$ and $\delta D^T D = 0$ guarantee a unique representation. Moreover, C, S, D satisfy the ODE system [21]:

$$C'(t) = (I - CC^T)F(CSD)DS^{-1}, \quad (4.6a)$$

$$S'(t) = C^T F(CSD)D, \quad (4.6b)$$

$$D'(t) = (I - DD^T)F(CSD)^T C(S^{-1})^T. \quad (4.6c)$$

It was shown in [21] that $\|U_{\text{DLRA}} - U\|_F$ where U solves (4.1) is of $\mathcal{O}(\varepsilon)$ for small time where ε is the distance from U to \mathcal{M}_r . Therefore U_{DLRA} is a quasi-optimal rank- r approximation of U . Additionally, in the formulation of (4.6), $F(U_{\text{DLRA}})$ is evaluated only via products of the form $F(U_{\text{DLRA}})D$ or $C^T F(U_{\text{DLRA}})$. If these products can be evaluated efficiently, like we have shown in Section 3, then the solution U_{DLRA} need not be explicitly constructed; and in such cases, one may expect computational and memory savings over the full-rank system.

While (4.6) has a lot of promising features, it is often numerically unstable to advance in time. This is because if the effective rank of U_{DLRA} is less than the prescribed r , then numerically S will contain singular values near machine epsilon. Thus timestepping (4.6) will be unstable unless Δt is taken to be impracticably small (see [25]). Because of this, other integrators have been developed to bypass this obstacle while keeping the main advantage of (4.6) – the evolution of the low-rank factors. The remainder of Section 4 is devoted to such integrators.

4.2 The Unconventional Integrator

The *unconventional integrator* for (4.6) was introduced in [3]. For completeness, we list a motivation and derivation of the algorithm here.

The goal is to create an integrator that does not invert S . This is achieved by first isolating the combined variable $K = CS$. Right multiplying (4.6a) by S , left multiplying (4.6b) by C , then summing yields a modified system of (4.6):

$$C'(t)S + CS'(t) = F(CSD)D, \quad (4.7a)$$

$$D'(t) = (I - DD^T)F(CSD)^T C(S^{-1})^T. \quad (4.7b)$$

The right hand side of (4.7a) is equal $K'(t)$. Additionally, the left hand side of (4.7a) can be written in terms of K rather than C . Applying such modifications gives

$$K'(t) = F(KD)D, \quad (4.8a)$$

$$D'(t) = (I - DD^T)F(CSD)^T C(S^{-1})^T. \quad (4.8b)$$

However, (4.8b) is explicitly dependent on C , not K . Given a $\Delta t > 0$, we remove the C coupling in (4.8b) by instead solving the approximate system

$$K'(t) = F(KD^T)D, \quad (4.9a)$$

$$D'(t) = 0. \quad (4.9b)$$

on $[t_0, t_0 + \Delta t]$. Since D remains constant in time we can replace D with $D = D_0$ where $C_0 S_0 D_0$ is the approximation to U at t . Additionally we set the initial condition for $K(t_0) = C_0 S_0$. While we cannot recover C using (4.9), the purpose of C is to track the evolution of the basis functions. Therefore, any orthogonal matrix that spans the same column space of $K_1 := K(t_0 + \Delta t)$ is sufficient. Such an updated matrix, denoted as C_1 , can be obtained by a QR or SVD decomposition of K_1 .

A similar system holds for evolving the D basis. Setting $L = DS^T$, then we solve the approximate ODE system

$$L'(t) = F(CL^T)^T C, \quad (4.10a)$$

$$C'(t) = 0. \quad (4.10b)$$

on $[t_0, t_0 + \Delta t_0]$. Again $C = C_0$, $L(t_0) = D_0 S_0^T$, and D_1 is an orthogonal matrix with equal column space of $L(t_0 + \Delta t)$. Once the new bases C_1 and D_1 are known. S is updated in the new space spanned by C_1 and D_1 through a Galerkin projection of (4.1). It was shown in [3] that the global timestepping error of the unconventional integrator is $\mathcal{O}(\varepsilon + \Delta t)$ where ε is the distance from \mathcal{M}_r to the solution U of (4.1).

Algorithm 4.1 details a forward Euler timestepping method applied to the unconventional integrator. Due to the line 5 of the algorithm, C_0 and C_1 must be kept in memory at the same time. Thus the memory requirement for Algorithm 4.1 is near double the storage cost of storing U via a low-rank factorization.

Algorithm 4.1: Forward Euler Timestepping with Unconventional Integrator

Input : $C_0 \in \mathbb{R}^{m \times r}$, $S_0 \in \mathbb{R}^{r \times r}$, $D_0 \in \mathbb{R}^{n \times r}$ // $U^n = C_0 S_0 D_0^T$

Output : $C_1 \in \mathbb{R}^{m \times r}$, $S_1 \in \mathbb{R}^{r \times r}$, $D_1 \in \mathbb{R}^{n \times r}$ // $U^{n+1} = C_1 S_1 D_1^T$

- 1 $K_1 = C_0 S_0 + \Delta t F(C_0 S_0 D_0^T) D_0$
 - 2 $[C_1, \sim] = \text{qr}(K_1)$
 - 3 $L_1 = D_0 S_0^T + \Delta t F(C_0 S_0 D_0^T)^T C_0$
 - 4 $[D_1, \sim] = \text{qr}(L_1)$
 - 5 $\tilde{S} = (C_1^T C_0) S_0 (D_0^T D_1)$
 - 6 $S_1 = \tilde{S} + \Delta t C_1^T F(C_1 \tilde{S} D_1^T) D_1$
-

4.3 Tangent Projection DLRA

Another integrator comes from a direct temporal discretization of the low-rank ODE (4.4). Given the decomposition $U = CSD^T$, an explicit representation of the projection operator $P_{T_U} := P_{T_U \mathcal{M}_r}$ is given by [25]:

$$P_{T_U} Z = CC^T Z - CC^T ZDD^T + ZDD^T. \quad (4.11)$$

Due to (4.11), we write the forward Euler update of $U^n = C_0 S_0 D_0^T$ as

$$\begin{aligned} U^{n+1} &= U^n + \Delta t P_{T_{U^n}} F(U^n) \\ &= U^n + \Delta t (C_0 C_0^T F(U^n) - C_0 C_0^T F(U^n) D_0 D_0^T + F(U^n) D_0 D_0^T) \end{aligned} \quad (4.12)$$

Since $U^n \in T_{U^n} \mathcal{M}_r$, then U^{n+1} is the projection of the full-rank forward Euler update onto the tangent space. Rewriting U^{n+1} in a low-rank fashion yields Algorithm 4.2. The algorithm shows the construction of U^{n+1} requires only two evaluations of F – unlike the three required for the unconventional integrator. However, one downside comes from lines 3 and 5 that show U^{n+1} must be stored in a rank $2r$ factorization. Thus the rank will exponentially grow without some culling mechanism (see Definition 4.1). In [20], the authors discuss such culling methods as well as extensions to higher order timestepping schemes.

Definition 4.1. Suppose $U \in \mathbb{R}^{m \times n}$ is a rank- r matrix U with SVD decomposition $U = CSD^T$, where $C \in \mathbb{R}^{m \times r}$, $S \in \mathbb{R}^{r \times r}$, and $D \in \mathbb{R}^{n \times r}$. Given $r_1 < r$, the cull U to rank r_1 to replace U by \tilde{U} where

$$\tilde{U} = \sum_{i=1}^{r_1} \sigma_i(U) c_i d_i^T = \tilde{C} \tilde{S} \tilde{D}^T. \quad (4.13)$$

Here \tilde{U} is rank r_1 , and $\tilde{C} \in \mathbb{R}^{m \times r_1}$, $\tilde{S} \in \mathbb{R}^{r_1 \times r_1}$, and $\tilde{D} \in \mathbb{R}^{n \times r_1}$ are defined by

$$\tilde{C} = C(:, 1:r_1), \quad \tilde{S} = S(1:r_1, 1:r_1), \quad \tilde{D} = D(:, 1:r_1). \quad (4.14)$$

Algorithm 4.2: Forward Euler Timestepping with Tangent Integrator

Input : $C_0 \in \mathbb{R}^{m \times r}, S_0 \in \mathbb{R}^{r \times r}, D_0 \in \mathbb{R}^{n \times r}$	// $U^n = C_0 S_0 D_0^T$
Output : $C_1 \in \mathbb{R}^{m \times 2r}, S_1 \in \mathbb{R}^{2r \times 2r}, D_1 \in \mathbb{R}^{n \times 2r}$	// $U^{n+1} = C_1 S_1 D_1^T$
1 $K_1 = F(C_0 S_0 D_0^T) D_0$	
2 $\tilde{S} = C_0^T K_1$	
3 $[C_1, R_C] = \text{qr}([C_0 \quad K_1])$	
4 $L_1 = F(C_0 S_0 D_0^T)^T C_0$	
5 $[D_1, R_D] = \text{qr}([D_0 \quad L_1])$	
6 $S_1 = R_C \begin{bmatrix} S_0 - \Delta t \tilde{S} & \Delta t I \\ \Delta t I & 0 \end{bmatrix} R_D^T$	

4.4 Projected Unconventional Integrator

Finally, we present a perturbation of the unconventional integrator from Section 4.2 that is beneficial to our predictor-corrector based discussed in Section 5.3. This integrator evolves the basis functions C and D in the exact same way as the unconventional integrator. Given the evolved basis C_1, D_1 from lines 3 and 5 of Algorithm 4.1, let H be the subspace of $\mathbb{R}^{m \times n}$ defined by

$$H = \{C_1 W D_1^T : W \in \mathbb{R}^{r \times r}\} \quad (4.15)$$

The update of the coefficient matrix S in the unconventional integrator (lines 6 and 7 of Algorithm 4.1) can be written using the projector P_H as

$$U^{n+1} = P_H(U^n + \Delta t F(P_H U^n)) \quad (4.16)$$

where

$$P_H Z = C_1 C_1^T Z D_1 D_1^T. \quad (4.17)$$

Again, this is applying a Galerkin projection to the full-rank ODE (4.1) and then discretizing in time using forward Euler. However, we can reverse the order of these operations. We instead project the full-rank forward Euler update onto H , that is,

$$U^{n+1} = P_H(U^n + \Delta t F(U^n)). \quad (4.18)$$

The projected unconventional integrator is given for a forward Euler timestep in Algorithm 4.3. One advantage of the projected unconventional integrator is that the error between the full-rank forward Euler update U_{FE}^{n+1} and the low-rank approximation is the orthogonal complement of $P_H U_{\text{FE}}^{n+1}$. This property has a benefit when our rank-adaptive algorithm in Section 5.3 is applied to the projected unconventional integrator.

Algorithm 4.3: Forward Euler Timestepping with Projected Unconventional Integrator

Input : $C_0 \in \mathbb{R}^{m \times r}, S_0 \in \mathbb{R}^{r \times r}, D_0 \in \mathbb{R}^{n \times r}$ // $U^n = C_0 S_0 D_0^T$
Output : $C_1 \in \mathbb{R}^{m \times r}, S_1 \in \mathbb{R}^{r \times r}, D_1 \in \mathbb{R}^{n \times r}$ // $U^{n+1} = C_1 S_1 D_1^T$

- 1 $K_1 = C_0 S_0 + \Delta t F(C_0 S_0 D_0^T) D_0$
- 2 $[C_1, \sim] = \text{qr}(K_1)$
- 3 $L_1 = D_0 S_0^T + \Delta t F(C_0 S_0 D_0^T)^T C_0$
- 4 $[D_1, \sim] = \text{qr}(L_1)$
- 5 $S_1 = (C_1^T C_0) S_0 (D_0^T D_1) + \Delta t C_1^T F(C_0 S_0 D_0^T) D_1$

We now show that this method is first order accurate in the following theorem which borrows heavily from Theorem 4 of [3].

Theorem 4.1. *Suppose $U(t)$ is the solution to (4.1) with initial condition U_0 on the interval $[0, T]$ for some final time $T > 0$. Suppose following assumptions are held:*

1. *F is Lipschitz continuous in Y , uniformly in t , and bounded uniformly in (Y, t) ; that is, there exist $L, B > 0$ such that*

$$\|F(Y, t) - F(Z, t)\|_F \leq L \|Y - Z\|_F \text{ and } \|F(Y, t)\|_F \leq B \quad (4.19)$$

for all $Y, Z \in \mathbb{R}^{m \times n}$ and $0 \leq t \leq T$

2. *There is an $\varepsilon > 0$ such that*

$$\|F(Y, t) - P_{T_Y} F(Y, t)\|_F \leq \varepsilon \quad (4.20)$$

for all $Y \in \mathcal{M}_r$ in a neighbourhood of $U(t)$ and all $0 \leq t \leq T$.

3. *There is a $\delta > 0$ such that*

$$\|U^0 - U(t^0)\|_F \leq \delta \quad (4.21)$$

Suppose Δt is sufficiently small such that the the standard forward Euler timestepping iteration is stable w.r.t $\|\cdot\|_F$, and furthermore suppose U is C^2 on $(0, T)$. Then setting $t_n = n\Delta t$ and letting U^n be the result of n steps of the projected unconventional integrator, we have

$$\|U^n - U(t^n)\|_F \leq c_0 \delta + c_1 \varepsilon + c_2 \Delta t \quad (4.22)$$

provided $t_n \leq T$. The constants c_i depend only on L, B , and T .

Before showing the proof of Theorem 4.1, we first derive a local truncation error estimate.

Lemma 4.1. *Suppose $U(t^n) = U^n$ where each are defined in Theorem 4.1, then assuming all suppositions in Theorem 4.1, we have*

$$\|U(t^{n+1}) - U^{n+1}\|_F \leq \Delta t (\widehat{c}_1 \varepsilon + \widehat{c}_2 \Delta t) \quad (4.23)$$

where \widehat{c}_1 and \widehat{c}_2 are positive constants that depend of L, B , and T .

Proof. We introduce ϑ from Equation 6 of [3]:

$$\vartheta = (4e^{L\delta t}BC + 9BL)\Delta t^2 + (3e^{L\Delta t} + 4)\varepsilon\Delta t + e^{Lh}\delta. \quad (4.24)$$

Since $U(t^n) = U^n$, then $\delta = 0$ for this estimate. Let $U_{\text{FE}}^{n+1} = U^n + \Delta t F(U^n)$. Since $U^{n+1} = P_H U_{\text{FE}}^{n+1}$ and by a triangle inequality we have

$$\begin{aligned} \|U(t^{n+1}) - U^{n+1}\|_F &\leq \|U(t^{n+1}) - P_H U(t^{n+1})\|_F + \|P_H U(t^{n+1}) - P_H U_{\text{FE}}^{n+1}\|_F \\ &\leq \|(I - P_H)U(t^{n+1})\|_F + \|U(t^{n+1}) - U_{\text{FE}}^{n+1}\|_F =: I_1 + I_2. \end{aligned} \quad (4.25)$$

By Lemma 3 of [3], $\|I_1\|_F \leq \vartheta$. Moreover, since U is sufficiently smooth, the full-rank forward Euler update satisfies

$$\|I_2\|_F \leq C(\Delta t)^2 \quad (4.26)$$

for C dependent on L and B . Therefore we have (4.23). The proof is complete. \square

Borrowing the stability of the full-rank forward Euler update, we can now prove Theorem 4.1

Proof of Theorem 4.1. Since the full-rank forward Euler iteration is stable, then so must the projection unconventional integrator since it is a projected version of the full-rank update. Thus we can extend the local error estimate from Lemma 4.1 to the global estimate (4.22). The proof is complete. \square

5 Adaptive Algorithms

The algorithms listed in Section 4 all operate on a fixed rank- r manifold \mathcal{M}_r . Choosing r too small risks polluting the solution with error from the low-rank approximation, while choosing r too large wastes memory and computational resources. To illustrate this point, we show in Figure 5.1 the results for the discretized 2D solid body rotation problem (2.6) with a box initial condition. For small values of r accuracy of the solution is severely degraded. However, for r sufficiently large, but still much less than the full-rank, the accuracy of the solution is reasonable and further increasing r does not significantly improve the solution.

Unfortunately, an efficient choice for r is rarely known a priori, and adaptive rank algorithms are needed. In the remainder of this section, we will present two recent adaptive algorithms from the literature and introduce a new predictor-corrector approach.

5.1 A Rank-Adaptive Unconventional Integrator

A rank-adaptive modification of the unconventional integrator was recently developed in [2]. From Section 4.2, the basis at t^n , C_0 and D_0 , and updated basis, C_1 and D_1 , must both be kept in memory at the same time in order to project the solution at t^n into the subspace spanned by the updated basis. The rank-adaptive unconventional integrator takes advantage of this by modifying lines 5-6 of Algorithm 4.1 and defining the basis at t^{n+1} to be composed of $[C_0 \ C_1]$ and $[D_0 \ D_1]$. Just like the tangent integrator, since $[C_0 \ C_1]$ contains $2r$ basis vectors, then the rank would exponentially




Figure 5.1: Discrete DLRA solutions for (2.4a) with fixed rank 4, 9, 17, 256 (left to right then top to bottom) at $T = \pi$. See Section 6 for spacial discretization details. Here the unconventional integrator with forward Euler timestepping was used (see Algorithm 4.1) with $\Delta t = 1/4096$. The initial condition (6.1) is rank 1 and the rank grows as the box steps through its first 90° rotation. It appears that a rank of 17, much lower than the full-rank of 256, is sufficient to accurately approximate the solution.

grow in time; thus culling the rank is required. The rank-adaptive unconventional integrator with forward Euler timestepping is shown in Algorithm 5.1.

Algorithm 5.1: Forward Euler Timestepping with rank-adaptive Unconventional Integrator

Input : $\tau > 0$ // tolerance
Input : $C_0 \in \mathbb{R}^{m \times r_0}, S_0 \in \mathbb{R}^{r_0 \times r_0}, D_0 \in \mathbb{R}^{n \times r_0}$ // $U^n = C_0 S_0 D_0^T$
Output : $C_1 \in \mathbb{R}^{m \times r_1}, S_1 \in \mathbb{R}^{r \times r_1}, D_1 \in \mathbb{R}^{n \times r_1}$ // $U^{n+1} = C_1 S_1 D_1^T$

- 1 $K_1 = C_0 S_0 + \Delta t F(C_0 S_0 D_0^T) D_0$
- 2 $C_1 = \text{qr}([C_0 \ K_1])$
- 3 $L_1 = D_0 S_0^T + \Delta t F(C_0 S_0 D_0^T)^T C_0$
- 4 $D_1 = \text{qr}([D_0 \ L_1])$
- 5 $\tilde{S} = (C_1^T C_0) S_0 (D_0^T D_1)$
- 6 $\tilde{S}_1 = \tilde{S} + \Delta t C_1^T F(C_1 \tilde{S} D_1^T) D_1$
- 7 $[S_C, S_1, S_D] = \text{svd}(\tilde{S}_1)$
- 8 $r_1 = \min\{r \in \mathbb{N} : 1 \leq r \leq 2r_0 \text{ and } \sum_{i=r+1}^{2r_0} \sigma_i(S_1)^2 < \tau^2\}$
- 9 $C_1 = C_1 S_C(:, 1:r_1); D_1 = D_1 S_D(:, 1:r_1); S_1 = S_1(1:r_1, 1:r_1)$

5.2 Adaptivity Via Step Truncation

Methods that compute the full-rank update, that is storing the evaluation $F(U)$, and then cull in post-processing are often labeled as step truncation methods [4]. Here we present one such method from [14] that does not rely on the DLRA framework. For simplicity, assume $F(U) = AUB^T$ for some $m \times m$ matrix A and $n \times n$ matrix B . If U^n is rank- r with decomposition $U^n = CSD^T$, then

$$U_{\text{FE}}^{n+1} = CSD^T + \Delta t ACS D^T B^T = [C \ AC] \begin{bmatrix} S & 0 \\ 0 & \Delta t S \end{bmatrix} [D \ BD]^T. \quad (5.1)$$

In terms of the QR factorizations $[C \ AC] = C_1 R_C$ and $[D \ BD] = D_1 R_D$, where C_1 and D_1 are orthogonal, (5.1) can be rewritten in the low-rank format

$$U_{\text{FE}}^{n+1} = C_1 \tilde{S} D_1^T, \quad \text{where } \tilde{S} = R_C \begin{bmatrix} S & 0 \\ 0 & \Delta t S \end{bmatrix} R_D^T. \quad (5.2)$$

The factorization in (5.2) shows that U_{FE}^{n+1} can have rank at most $2r$.

As with the tangent projection method in Section 4.3, the rank of the forward Euler update will grow exponentially with the number of timesteps, unless culling is applied. However, unlike the tangent projection method, (5.2) contains no modeling error; thus the only error from the low-rank approximation is made in the culling. To maintain an $\mathcal{O}(\Delta t)$ method, the tail of the singular values is culled until the portion removed is $\mathcal{O}(\Delta t^2)$. Extensions to higher order methods can be found in [14]. Algorithm 5.2 shows the rank-adaptive step truncation with forward Euler time-stepping

applied to an N -separable operator.

Algorithm 5.2: Forward Euler Timestepping with rank-adaptive Step Truncation

Input : $\tau > 0$ // tolerance
Input : $C_0 \in \mathbb{R}^{m \times r_0}, S_0 \in \mathbb{R}^{r_0 \times r_0}, D_0 \in \mathbb{R}^{n \times r_0}$ // $U^n = C_0 S_0 D_0^T$
Input : $A_\kappa \in \mathbb{R}^{m \times m}, B_\kappa \in \mathbb{R}^{n \times n}, \kappa = 1, \dots, N$ // $F(U) = \sum_{\kappa=1}^N A_\kappa U B_\kappa^T$
Output : $C_1 \in \mathbb{R}^{m \times r_1}, S_1 \in \mathbb{R}^{r_1 \times r_1}, D_1 \in \mathbb{R}^{n \times r_1}$ // $U^{n+1} = C_1 S_1 D_1^T$

- 1 $K_1 = [C_0 \ A_1 C_0 \ A_2 C_0 \ \dots \ A_N C_0] \in \mathbb{R}^{m \times (N+1)r_0}$
- 2 $[C_1, R_C] = \text{qr}(K_1)$.
- 3 $L_1 = [D_0 \ B_1 D_0 \ B_2 D_0 \ \dots \ B_N D_0] \in \mathbb{R}^{n \times (N+1)r_0}$
- 4 $[D_1, R_D] = \text{qr}(L_1)$
- 5 $\tilde{S} = R_C \text{diag}(S, \Delta t S, \Delta t S, \dots, \Delta t S) R_D^T$
- 6 $[S_C, S_1, S_D] = \text{svd}(\tilde{S})$
- 7 $r_1 = \min\{r \in \mathbb{N} : 1 \leq r \leq (N+1)r_0 \text{ and } \sum_{i=r+1}^{(N+1)r_0} \sigma_i(S_1)^2 < \tau^2\}$
- 8 $C_1 = C_1 S_C(:, 1:r_1); D_1 = D_1 S_D(:, 1:r_1); S_1 = S_1(1:r_1, 1:r_1)$

An advantage of step truncation is that in a single timestep consistency error created in order to obtain a low-rank solution is only created via a post-processing step that can be easily controlled. A disadvantage is that the memory footprint for storing the full-rank update is directly proportional to the number of terms in F . For a general N -separable bilinear form, a rank $(N+1)r$ decomposition needs to be stored even if the solution $U(t^{n+1})$ is close to a rank- r matrix.

5.3 Predictor-Corrector Rank Adaptivity

We now present a new rank-adaptive algorithm that involves two steps. The predictor step produces a low-rank solution U_{P}^{n+1} by any of the numerical DLRA integrators from Section 4. The corrector step builds a low-rank approximation of the modeling error of the DLRA integrator which we also refer to as the residual R :

$$R = U_{\text{FE}}^{n+1} - U_{\text{P}}^{n+1}. \quad (5.3)$$

where

$$U_{\text{FE}}^{n+1} := U^n + \Delta t F(U^n). \quad (5.4)$$

We summarize the two steps below; details are provided in Algorithm 5.4. For the remainder of this subsection we assume $U^n = C_0 S_0 D_0^T$ is a rank- r decomposition of U^n and U_{P}^{n+1} has a rank r_1 decomposition $U_{\text{P}}^{n+1} = C_1 S_1 D_1^T$ where $r_1 = r$ if Algorithm 4.1 or Algorithm 4.3 is used to compute U_{P}^{n+1} and $r_1 = 2r$ if Algorithm 4.2 is used to compute U_{P}^{n+1} .

5.3.1 An randomized algorithm-based error estimator

The goal is to approximate $\|R\|_F$ and use this approximation to determine whether to add or decrease the rank of the predictor. Rank can then be added as needed using the principal singular vectors of the singular value decomposition (SVD) for R . Rather than compute the mn individual entries of R , we instead seek a low-rank approximation, which we denote \tilde{R} , that can be computed efficiently in terms of operation count and memory requirements. Matrix-free SVD algorithms are memory efficient, but numerical tests of the svds algorithm in MATLAB reveal that the time to compute \tilde{R} is often several orders of magnitude longer than the computation time of the U_{P}^{n+1} .

To efficiently compute \tilde{R} , we turn to a randomized SVD (R-SVD) approach. The randomized SVD algorithm shown in Algorithm 5.3 is sourced from Algorithms 4.1 and 5.1 of [15]. To summarize, a random sketching matrix $Y \in \mathbb{R}^{n \times (l+p)}$ is created, and then an approximation to the $\text{Col}(R)$ is formed using RY . The parameter $l \in \mathbb{N}$ specifies that \tilde{R} should be close to a rank l approximation of R . The *oversampling parameter* $p \in \mathbb{N}$ is the number of additional samples used in order to increase the probability that $\|R - \tilde{R}\|_F$ is small.

To extract an orthogonal basis for $\text{Col}(RY)$, we use the QR factorization $RY = QP$ where $Q \in \mathbb{R}^{m \times (l+p)}$ is orthogonal and $P \in \mathbb{R}^{(l+p) \times (l+p)}$ is not used. Then we set $\tilde{R} = QQ^T R$. Because Q is orthogonal, we only need the SVD of $B^T := Q^T R \in \mathbb{R}^{(l+p) \times n}$ to form an SVD of \tilde{R} . Specifically if $B = \tilde{D}\tilde{S}X^T$ is the SVD of B , then

$$\tilde{R} = QQ^T R = QB^T = QX\tilde{S}\tilde{D}^T = \tilde{C}\tilde{S}\tilde{D}^T \quad (5.5)$$

where $\tilde{C} = QX$. Moreover, since B has only $l+p$ columns, a direct SVD can be computed quite cheaply. This yields a low-rank decomposition of \tilde{R} .

Algorithm 5.3: Randomized SVD

Input: $l \in \mathbb{N}$ // Approximation rank
Input: $p \in \mathbb{N}$ // Oversampling parameter
Input: R_{LR} // $R_{\text{LR}} \sim R$ (see Definition 5.1)
Output: $\tilde{C} \in \mathbb{R}^{m \times l+p}$, $\tilde{S} \in \mathbb{R}^{l+p \times l+p}$, $\tilde{D}^{n \times l+p}$ // $\tilde{R} = \tilde{C}\tilde{S}\tilde{D}^T$

- 1 Construct $Y \in \mathbb{R}^{n \times l+p}$ whose entries are identically and independently distributed Gaussian random variables
- 2 $[Q, \sim] = \text{qr}(R_{\text{LR}}Y)$
- 3 $B = R_{\text{LR}}^T Q$
- 4 $[\tilde{D}, \tilde{S}, X] = \text{svd}(B)$
- 5 $\tilde{C} = QX$

Since the residual R from (5.3) contains the term $\Delta t F(U^n)$, the most expensive operation in Algorithm 5.3 is the evaluation of the products $F(U^n)Y$ and $F(U^n)^T Q$. If $l+p \approx r$, then the computational cost of building \tilde{R} will be similar to any of the DLRA updates in Section 4 as all of these methods require the products $F(U^n)C_0$ and $F(U^n)^T D_0$.

The statistical accuracy of this approximation is provided by [15, Theorem 10.7]:

Proposition 5.1. *Set $d = \min\{m, n\}$. Suppose $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_d$ are the singular values of R . Suppose $l \geq 2$ and $p \geq 4$; then \tilde{R} created by Algorithm 5.3 satisfies the following error bound:*

$$\|R - \tilde{R}\|_F \leq \left(1 + 10\sqrt{\frac{l}{p}}\right) \left(\sum_{j>l} \sigma_j^2\right)^{1/2} + 12\frac{\sqrt{p(l+p)}}{p+1}\sigma_{l+1} \quad (5.6)$$

with failure probability at most $7e^{-p}$.

Proof. This result follows from setting $t = e$ and $u = \sqrt{2p}$ in [15, Theorem 10.7] and bounding the resulting constants by the nearest whole number. \square

Using the R-SVD, an error indicator for the predictor and a subsequent corrector can be constructed. While the rank of R is not known a priori, it can be bounded when F is an N -separable linear operator.

Proposition 5.2. *Let F be an N -separable linear operator; let U^n have rank- r ; and let R be given in (5.3). Then the maximum rank of R is either (i) $(N+2)r$ if Algorithm 4.1 or Algorithm 4.3 are used to compute U_P^{n+1} or (ii) Nr if Algorithm 4.2 is used to compute U_P^{n+1} .*

Proof. Let $U^n = C_0 S_0 D_0^T$ be a rank- r decomposition of U^n , Since F is N -separable, it has a decomposition (see (3.11)) of the form

$$F(U^n) = \sum_{\kappa=1}^N A_\kappa U^n B_\kappa^T = \sum_{\kappa=1}^N (A_\kappa C_0) S_0 (B_\kappa D_0)^T \quad (5.7)$$

If Algorithm 4.1 or Algorithm 4.3 are used to compute U_P^{n+1} , then $U_P^{n+1} = C_1 S_1 D_1^T$ also has rank- r , and a direct calculation shows that

$$\begin{aligned} R &= U_{FE}^{n+1} - U_P^{n+1} = C_0 S_0 D_0^T + \Delta t \sum_{\kappa=1}^N (A_\kappa C_0) S_0 (B_\kappa D_0)^T - C_1 S_1 D_1^T \\ &= \begin{bmatrix} C_0 & A_1 C_0 & \cdots & A_N C_0 & C_1 \end{bmatrix} \text{diag}(S_0, \Delta t S_0, \dots, \Delta t S_0, -S_1) \begin{bmatrix} D_0 & B_1 D_0 & \cdots & B_N D_0 & D_1 \end{bmatrix}^T \end{aligned} \quad (5.8)$$

has a decomposition that is at most rank- $(N+2)r$. If Algorithm 4.2 is used to compute U_P^{n+1} , then

$$\begin{aligned} R &= U_{FE}^{n+1} - U_P^{n+1} \\ &= U^n + \Delta t F(U^n) - (U^n + \Delta t (C_0 C_0^T F(U^n) - C_0 C_0^T F(U^n) D_0 D_0^T + F(U^n) D_0 D_0^T)) \\ &= \Delta t (I - C_0 C_0^T) F(U^n) (I - D_0 D_0^T) \\ &= \Delta t \sum_{\kappa=1}^N ((I - C_0 C_0^T) A_\kappa C_0) S_0 ((I - D_0 D_0^T) B_\kappa D_0)^T \\ &= \left[\widetilde{A_1 C_0} \quad \cdots \quad \widetilde{A_N C_0} \right] \text{diag}(\Delta t S_0, \dots, \Delta t S_0) \left[\widetilde{B_1 D_0} \quad \cdots \quad \widetilde{B_N D_0} \right]^T, \end{aligned} \quad (5.9)$$

where $\widetilde{A_\kappa C_0} = (I - C_0 C_0^T) A_\kappa C_0$ and $\widetilde{B_\kappa D_0} = (I - D_0 D_0^T) B_\kappa D_0$. (5.9) implies that R has decomposition that is at most rank Nr . The proof is complete. \square

Definition 5.1. *We denote by R_{LR} the low-rank representation of R given by the decomposition in (5.8) or (5.9) and write $R_{LR} \sim R$.*

The rank bound from Proposition 5.2 is used in the estimator ν of $\|R\|_F$. We use the $l+p$ singular values of \tilde{R} to approximate the first $l+p$ singular values of R and the smallest singular value of \tilde{R} to approximate the remaining non-trivial singular values of R . This gives

$$\|R\|_F^2 \approx \|\tilde{R}\|_F^2 + \sum_{l+p+1}^{\theta(r)} \sigma_i^2(R) \approx \|\tilde{R}\|_F + [\theta(r) - (l+p)] (\sigma_{l+p}^2 \tilde{R}) =: \nu^2 \quad (5.10)$$

where $(\ell+p) \leq \theta(r)$ and

$$\theta(r) = \begin{cases} (N+2)r & \text{if Algorithm 4.1 or Algorithm 4.3 is used to compute } U_P^{n+1}, \\ Nr & \text{if Algorithm 4.2 is used to compute } U_P^{n+1}. \end{cases} \quad (5.11)$$

Accounting for the additional singular values in (5.10) gives a conservative estimate that favors accuracy over memory savings; a less conservative estimate would be to simply use $\|\tilde{R}\|_F$ to approximate $\|R\|_F$.

5.3.2 Adding and Removing Rank

Let $\tilde{R} = \tilde{C}\tilde{S}\tilde{D}^T$ be the low-rank approximation that is constructed by Algorithm 5.3. Rank is added or removed to the predictor U_P^{n+1} based on the size of the estimator ν , defined in (5.10), relative to a prescribed tolerance $\tau > 0$. There are three cases:

Case 1: $\nu \geq \tau$. In this case, rank is added to the predictor by moving rank one components from \tilde{R} to U_P^{n+1} and updating the corrector:

$$U_{(j)}^{n+1} = U_{(j-1)}^{n+1} + \tilde{C}(:, j)\sigma_j(\tilde{R})\tilde{D}(:, j)^T, \quad U_{(0)}^{n+1} = U_P^{n+1}, \quad (5.12)$$

$$\tilde{R}_{(j)} = \tilde{R}_{(j-1)} - \tilde{C}(:, j)\sigma_j(\tilde{R})\tilde{D}(:, j)^T, \quad \tilde{R}_{(0)} = \tilde{R}, \quad (5.13)$$

$$\nu_{(j)}^2 = \nu_{(j-1)}^2 - \sigma_j^2(\tilde{R}), \quad \nu_{(0)}^2 = \nu^2, \quad (5.14)$$

for $j \in \{1, \dots, l\}$. These rank one updates are added until $\nu_{(j)} < \tau$ or until $j = l$ and $\nu_{(l)} \geq \tau$. In the latter case the predictor is updated and a new residual is constructed:

$$R \leftarrow R - \tilde{C}(:, 1:l)\tilde{S}(1:l, 1:l)\tilde{D}(:, 1:l)^T, \quad U_{n+1}^P \leftarrow U_{(j)}^{n+1}, \quad \theta(r) \leftarrow \theta(r) - l, \quad (5.15)$$

and the process repeats until the tolerance criterion is satisfied. The final low-rank update is then called to remove possibly redundant information introduced by the correction.

Case 2: $\nu \leq \delta\tau$ for some prescribed constant $\delta \in (0, 1)$. In this case, rank is removed by culling the predictor. Let $S_1 = C_S \hat{S} D_S^T$ be the SVD decomposition of $S_1 \in \mathbb{R}^{r_1 \times r_1}$. Thus $U_P^{n+1} = \hat{C} \hat{S} \hat{D}^T$ where $\hat{C} = C_1 C_S$ and $\hat{D} = D_1 D_S$. The algorithm removes components corresponding to the smallest singular values of S from U_P^{n+1} and appends them to R . Specifically let

$$\nu_{(-j)}^2 := \nu^2 + \sum_{\kappa=j+1}^r \hat{S}(\kappa, \kappa)^2 \quad \text{and} \quad j_* = \min\{j \in \mathbb{N} : 1 \leq j \leq r_1 \text{ and } \nu_{(-j)} \geq \delta\tau\} \quad (5.16)$$

Then the updated approximation

$$U^{n+1} = \hat{C}(:, 1:j_*) \hat{S}(1:j_*, 1:j_*) \hat{D}(:, 1:j_*)^T. \quad (5.17)$$

has rank j_* and induces an error

$$\|U_{\text{FE}}^{n+1} - U^{n+1}\| \leq \|R\|_F + \|U_P^{n+1} - U^{n+1}\|_F \approx \nu^2 + \sum_{\kappa=j+1}^r \hat{S}(\kappa, \kappa)^2 = \nu_{(-j_*)}. \quad (5.18)$$

In general, $\nu_{(-j)}$ is usually a slight overestimation of the true error $\|R + (U_P^{n+1} - U^{n+1})\|_F$ because of the conservative estimate of ν^2 by the R-SVD algorithm and the fact that R and $(U_P^{n+1} - U^{n+1})$ are not orthogonal in the Frobenius norm. However, if the projected unconventional integrator from Section 4.4 is used, then R and $(U_P^{n+1} - U^{n+1})$ are orthogonal and $\nu_{(-j)}$ is a much better representation of the error.

Case 3: $\delta\tau < \nu < \tau$. In this case, the predictor is considered a sufficiently accurate and efficient

approximation. Thus $U^{n+1} = U_p^{n+1}$.

Algorithm 5.4: predictor-corrector Rank Adaptivity

Input: $l \in \mathbb{N}$ // Approximation rank for R
Input: $p \in \mathbb{N}$ // Oversampling parameter
Input: $\tau > 0$ // Tolerance
Input: $\tau_{\text{cull}} > 0$ // Redundant Culling tolerance
Input: $0 < \delta < 1$ // Rank removal factor
Input: $C_0 \in \mathbb{R}^{m \times r}, S_0 \in \mathbb{R}^{r \times r}, D_0 \in \mathbb{R}^{n \times r}$ // $U^n = C_0 S_0 D_0^T$
Output: $C_1 \in \mathbb{R}^{m \times r_1}, S_1 \in \mathbb{R}^{r_1 \times r_1}, D_1 \in \mathbb{R}^{n \times r_1}$ // $U^{n+1} = C_1 S_1 D_1^T$

- 1 Create $C_1 \in \mathbb{R}^{m \times r_1}, S_1 \in \mathbb{R}^{r_1 \times r_1}, D_1 \in \mathbb{R}^{n \times r_1}$ // $U_p^{n+1} = C_1 S_1 D_1^T$
using Algorithm 4.1, 4.2, or 4.3
- 2 Construct $R_{\text{LR}} \sim R = U_{\text{FE}}^{n+1} - C_1 S_1 D_1^T$ // see Definition 5.1
- 3 $[\tilde{C}, \tilde{S}, \tilde{D}] = \text{r-svd}(l, p, R_{\text{LR}})$ using Algorithm 5.3 // $\tilde{R} = \tilde{C} \tilde{S} \tilde{D}^T$
- 4 $\nu = \sqrt{\sum_{i=1}^{l+p} \tilde{S}_{i,i}^2 + (\theta(r) - (l+p)) \tilde{S}_{l+p,l+p}^2}$ // $\theta(r)$ from (5.11)
- 5 **if** $\nu \geq \tau$ **then**
 - 6 $c = 0$
 - 7 **while** $\nu \geq \tau$ **do**
 - 8 Find minimum $1 \leq j \leq l$ such that $\nu_{(j)} < \tau$ where

$$\nu_{(j)} = \sqrt{\sum_{i=1}^{l+p} \tilde{S}_{i,i}^2 + (\theta(r) - (l+p+c+j)) \tilde{S}_{l+p,l+p}^2}$$
 - 9 **if** no such j exists **then**
 - 10 $C_1 = [C_1 \ \tilde{C}(:, 1:l)], D_1 = [D_1 \ \tilde{D}(:, 1:l)]; S_1 = \text{diag}(S_1, \tilde{S}(1:l, 1:l))$
 - 11 Construct $R_{\text{LR}} \sim R = U_{\text{FE}}^{n+1} - C_1 S_1 D_1^T$
 - 12 $[\tilde{C}, \tilde{S}, \tilde{D}] = \text{r-svd}(l, p, R_{\text{LR}})$ using Algorithm 5.3
 - 13 $c = c + l$
 - 14 $\nu = \sqrt{\sum_{i=1}^{l+p} \tilde{S}_{i,i}^2 + (\theta(r) - (l+p+c)) \tilde{S}_{l+p,l+p}^2}$
 - 15 **else**
 - 16 $C_1 = [C_1 \ \tilde{C}(:, 1:j)], D_1 = [D_1 \ \tilde{D}(:, 1:j)]; S_1 = \text{diag}(S_1, \tilde{S}(1:j, 1:j))$
 - 17 **end**
 - 18 **end**
 - 19 $[C_1, R_C] = \text{qr}(C_1); [D_1, R_D] = \text{qr}(D_1)$
 - 20 $S_1 = R_C S_1 R_D^T$
 - 21 $[S_C, S_1, S_D] = \text{svd}(S_1); r_1 = \text{size}(S_1)$ // $\text{size}(S_1)$ is # of columns of S_1
 - 22 $r_1 = \min\{j \in \mathbb{N} : 1 \leq j \leq r_1 \text{ and } \sigma_j(S_1) < \tau_{\text{cull}}\}$
 - 23 $C_1 = C_1 C_S(:, 1:r_1), D_1 = D_1 D_S(:, 1:r_1), S_1 = S_1(1:r_1, 1:r_1)$
 - 24 **else if** $\nu \leq \delta \tau$ **then**
 - 25 $[C_S, S_1, D_S] = \text{svd}(S_1)$
 - 26 Find maximum $1 \leq j \leq r$ such that $\nu^2 + \sum_{\kappa=j+1}^r \sigma_{\kappa}(S_1)^2 > (\delta \tau)^2$
 - 27 $C_1 = C_1 C_S(:, 1:j), D_1 = D_1 D_S(:, 1:j)$
 - 28 $S_1 = S_1(1:j, 1:j)$
 - 29 **else**
 - 30 **continue**
 - 31 **end**

5.4 Memory Comparisons

In this subsection, we give a comparison of the memory requirements to run each of the adaptive algorithms presented above: the *rank-adaptive unconventional integrator* (RAUC) algorithm, the *rank-adaptive step truncation* (RAST) algorithm and the new (RAPC) *rank-adaptive predictor-corrector* algorithm. Each of these algorithms primarily involve function evaluations which have been made memory efficient by Proposition 3.3 or modern QR and SVD decompositions that are also memory efficient. Thus we gauge the memory requirements of each adaptive algorithm by the max of the sum of ranks of all objects held in memory at one time. As an example, for the unconventional integrator (Algorithm 4.1), the rank- r matrix U^n must be stored at the same time as the rank- r matrix $C_1 D_1^T$; thus we list the memory requirements as $r + r = 2r$. Using this metric, Table 5.1 shows the results below.

Method	Section	Algorithm	Memory
RAUC [2]	Section 5.1	Algorithm 5.1	$2r$
RAST [14]	Section 5.2	Algorithm 5.2	$(N + 1)r$
RAPC	Section 5.3	Algorithm 5.4	$r + \max\{r_1, r\} + l + p$

Table 5.1: Memory comparisons of the rank-adaptive algorithms list in Section 5. Here r is the rank of U^n , F is N -separable, r_1 is the final rank of the U^{n+1} , and the parameters l and p are given in Algorithm 5.3.

The RAUC algorithm is the most memory efficient; this is expected, as it reuses information obtained through the DLRA process. However, because of this convergence is entirely dependent on the modeling error introduced from the DLRA approximation. The RAST algorithm uses the most memory per run due to storing the full-rank update in a low memory fashion; this allows for accurate approximations to be guaranteed, but use excessive memory in cases were DLRA methods work well. The RAPC algorithm sits in between the these two. The maximum memory is not known a-priori in that it is dependent on the tolerance given and the oversampling parameter p used in the R-SVD, but by using the SVD to obtain new basis vectors for U^{n+1} the algorithm seeks to prioritize memory as much as possible.

6 Numerical Results

In this section we compare the three adaptive strategies listed in Table 5.1: the rank-adaptive unconventional integrator (RAUC) algorithm, the rank-adaptive step truncation (RAST) algorithm and the new (RAPC) rank-adaptive predictor corrector algorithm. The RAPC algorithm is applied to all the three DLRA integrators in Section 4: the unconventional integrator (RAPC-UC) in Section 4.2, the tangent projector integrator (RAPC-Tan) in Section 4.3, and the projected unconventional integrator (RAPC-Proj) in Section 4.4, for a total of five distinct methods.

We apply the five methods to three test problems. The first problem is the 2D solid-body rotation problem from (2.6); these results are presented in Section 6.1. A large fraction of the results are dedicated this problem, which assess the ability of the adaptive methods to match the rank of the true solution as it oscillates in time. The second and third test are modifications of solid-body rotation problem. The problem in Section 6.2 introduces a relaxation problem common to radiation transport and gauges how the rank-adaptive methods perform when the rank does not change after a short time. The problem in Section 6.3 also contains a relaxation along with localized sources.

The RAUC and RAST algorithms require only a tolerance τ to run while all of the RAPC algorithms require a tolerance τ , a culling fraction δ , and a redundant culling parameter τ_{cull} . In our tests RAPC-Tan is much more apt to decrease rank than the other two RAPC algorithms, so we choose $\delta = 0.5$ for RAPC-Tan and $\delta = 0.7$ for RAPC-UC and RAPC-Proj to make up for this difference. In all tests we set $\tau_{\text{cull}} = 10^{-14}$. For the randomized SVD in Algorithm 5.4, we use $l = 3$ in all cases, but the oversampling parameter p is problem-specific. All of the algorithms are implemented and run using MATLAB [26]. Random vectors are drawn using MATLAB’s built-in `randn` function. Finally $\Omega = [-1, 1]^2$ is used for all problems and we set the basis $\{\varphi_j(x)\}$ for $V_{x,h}$ to be the shifted localized Legendre polynomials on every interval $T \in \mathcal{T}_{x,h}$ that are scaled to be orthonormal in $L^2(\Omega_x)$. The same basis is used for $V_{y,h}$.

6.1 Test Problem 1 - 2D solid-body Rotation

The equation for solid-body is given in (2.4a) with initial condition

$$u_0(x, y) = \chi_{\{|x|<0.5\}}(x)\chi_{\{|y|<0.5\}}(y), \quad (6.1)$$


where χ_E is the indicator function on a set E . The DG discretization and low-rank structure for this problem is discussed extensively in Section 2.2 and Section 3.2. We choose a uniform discretization with $h_x = h_y = 1/128$ and $k_x = k_y = 0$ so that the degrees of freedom in each variable is 256. The oversampling parameter p is set to 7.

Because u_0 is separable, the L^2 projection of u_0 onto V_h has a rank of one. However, as the solution evolves via a counter-clockwise rotation, the numerical rank of the true solution oscillates with its numerical rank, i.e., the total number of singular values above some tolerance, increasing and decreasing back to rank 1 at $t = \pi/2$. Figure 6.1a shows the numerical rank of the discrete full-rank forward Euler update as a function of time. From $t = \pi/2$ to $t = \pi$, the rank does not increase as much as the first 90 degree rotation. This is because the numerical dissipation in the system smooths out the solution and forces the singular values to decay faster. In order to measure rates of convergence in time, we set $\omega = 1/256$ to be our “base timestep” and take Δt to be some fraction of ω .

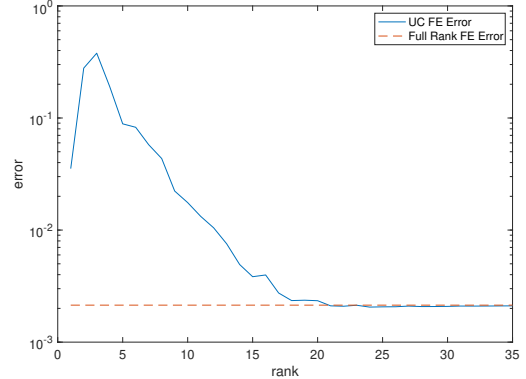
To gauge how much rank is expected to resolve the discrete solution up to timestepping error, we plot the error of a rank- r unconventional integrator at time $t = \pi$ against the full-rank SSP-RK3 integrator. The initial condition in this rank is a rank- r truncation of the SVD of the projected initial condition. One important detail to observe is that that all of the DLRA integrators in Section 4 are dependent on the low-rank components C, S, D rather than the product CSD^T ; because of this, even though the second through r -th singular values are near machine epsilon, all of the singular vectors are equally weighted in creating the updated basis vectors in time (see Algorithm 4.1). Thus the modelling error is dependent on the initial basis vectors C_0 and D_0 of U^n – even if that information is essentially not used to construct U^n when the product $U^n = C_0 S_0 D_0^T$ is formed. This allows DLRA integrators to perform better with a larger starting rank and can be seen in Figure 6.1b. As the initial rank grows, the errors decrease until the error is saturated by the forward Euler timestepping error. The dotted line in Figure 6.1b is the full-rank forward Euler error – 2.138×10^{-3} whose error is only dependent on the temporal discretization. An initial rank of at least 21 is needed to remove the modelling error and produce a DLRA solution that is as good as a full rank solution.

6.1.1 Isolating Modelling Error

For the first test we wish to compare the rank of the five methods when modelling error is the dominant source of error in the problem. In this test each of the low-rank solutions are less



(a) Numerical rank of the forward Euler full-rank iteration of (2.6) with $\Delta t = 1/4096$ as a function of time. Given a tolerance tol , we define the numerical rank of the discrete function by the number of singular values larger than tol .



(b) Errors of the unconventional integrator with forward Euler timestepping (see Algorithm 4.1) applied to (2.6) at time $T = \pi$ as a function of the initial rank- r . The errors are created by comparison against the full-rank SSP-RK3 approximation with $\Delta t = 1/4096$ for all runs.

Figure 6.1

Method	RAUC	RAST	RAPC-UC	RAPC-Tan	RAPC-Proj
τ	$25\Delta t^2$	$19\Delta t^2$	$880\Delta t^2$	$11255\Delta t^2$	$857\Delta t^2$
Error	$1.4901e-6$	$1.1325e-6$	$5.2452e-5$	$6.7085e-4$	$5.1081e-5$
τ	Δt^2	$5\Delta t^2$	$400\Delta t^2$	$500\Delta t^2$	$200\Delta t^2$
Error	$5.9605e-8$	$2.9802e-7$	$2.3842e-5$	$2.9802e-5$	$1.1921e-5$
Error	$2.592e-3$	$2.108e-3$	$2.1348e-3$	$2.1414e-3$	$2.1642e-3$

Table 6.1: Tolerances and Errors for all five rank-adaptive methods in Section 6.1.1. The errors are computed at $T = \pi$ against the full-rank SSP-RK3 discretization. $\Delta t = \omega/16$ where $\omega = 1/256$.

accurate than the full rank solution. We set $\Delta t = \omega/16$ and run each method with a tolerance τ to achieve an error as close to 3×10^{-3} when compared against the SSP-RK3 approximation. The tolerances are listed in Table 6.1 and were chosen by trial and error. The error of the full-rank forward Euler iteration against the SSP-RK3 iteration is 2.138×10^{-3} so our choice of 3×10^{-3} is to judge the modelling error by the low-rank approximation rather than the temporal discretization error. The errors for each method are given in Table 6.1, and Figure 6.2 shows the rank of each of the five methods as a function of time. The maximum rank of each of the methods far less than the full-rank forward Euler update (Figure 6.1a). However, each of the RAPC methods do much better than other two methods. Both the RAPC-UC and RAPC-Proj methods need about half of the rank of the RAUC implementation in order to achieve the same error. RAPC-UC and RAPC-Proj have a maximum rank of $r = 17$ and an error of 3.039×10^{-3} (see Table 6.1). From Figure 6.1b, the error of 3.039×10^{-3} fits nicely between errors of the fixed-rank unconventional integrator for the rank-16 and rank-17 runs. The RAPC-UC and RAPC-Proj solutions capture the same accuracy as a fixed rank-17 unconventional integrator and these methods only keep a rank 17 solution for less than half of the simulation. For a majority of the time, the RAPC-UC and RAPC-Proj solutions are below a rank of 17.

The RAST and RAPC-Tan methods best capture the temporal oscillations in rank while the




Figure 6.2: Rank (left) and memory requirements (right) of each of the five methods over time for the 2D solid-body rotation problem (Section 6.1) with $\Delta t = \omega/16$ where $\omega = 1/256$. The tolerance τ for each method is set so that the error at time $T = \pi$ against the full-rank SSP-RK3 method is near 3×10^{-3} and can be found in Table 6.1, and the memory footprint is calculated using Table 5.1. For readability, the data is plotted at every 200 timesteps. The legend labels provide the error of the method, normalized with respect to the error of the RAST algorithm.

rank of the methods based on the unconventional integrator (RAUC, RAPC-UC, and RAPC-Proj) appear to flatten out over the second 90° rotation.

6.1.2 Resolving Modelling Error

We now run the same test but tune to the largest tolerances, again by trial and error, so that the error of each run is as close to the temporal discretization error, 2.138×10^{-3} , as possible. Thereby the modelling error is resolved meaning that each of the low-rank solutions is a good approximation as the full rank solution. The rank of the runs are shown in Figure 6.3. Similar to the previous run, the RAPC-UC and RAPC-Proj perform the best with the RAPC-UC needing at maximum a rank of 23 for approximately a quarter of the run. Figure 6.1b shows that this is close to the rank of 21 that was required to capture this error in the non rank-adaptive case. Also after $t = 3\pi/8$, a rank of 23 is not required to resolve the modelling error.

The RAPC-Tan and RAST integrators perform similarly in rank, but the RAPC-Tan integrator requires storage of only $2r + 10$ vectors in each direction (see Table 5.1) while the RAST requires $5r$ vectors in each direction. When $r = 38$, this difference is quite large as evidenced in the memory plot of Figure 6.3; thus the RAPC-Tan integrator is much more memory efficient than the RAST integrator. RAST is practical when the rank of the solution is guaranteed to be small; indeed, the authors in [14] propose a different discretization of (2.4a) where the characteristics are instead approximated which leaves the numerical solution low-rank for all time.

Finally, the rank-adaptive unconventional integrator [2] (RAUC) described in Section 5.1 was not able to capture the required error for any reasonable rank – any tighter tolerances did not produce a more accurate result. This is because the method can only gain information from the current and updated basis created by the unconventional integrator. If this space is not rich enough, then the method cannot resolve the modelling error. After the first timestep, the numerical rank of the full-rank forward Euler update is 5 and the other four methods correctly have a rank of 5. However, because the initial condition is only rank 1, the RAUC integrator can only have a




Figure 6.3: Rank (left) and memory requirements (right) of each of the five methods over time for the 2D solid-body rotation problem (Section 6.1) with $\Delta t = \omega/16$ where $\omega = 1/256$. The tolerance τ for each method was set so that the error at time $T = \pi$ against the full-rank SSP-RK3 method was near 3×10^{-3} and can be found in Table 6.1, and the memory footprint was calculated using Table 5.1. For readability, the data is plotted at every 200 timesteps. The legend labels provide the error of the method, normalized with respect to the error of the RAST algorithm.

maximum of two basis vectors in the updated components C and D . Therefore there is no way to resolve the modelling error. A simple remedy is to increase the rank of the initial condition to five (using four additional singular values are near machine epsilon, with corresponding singular vectors that result from MATLAB’s `svd` algorithm). In this case, the algorithm, denoted as RAUC-5 in Figure 6.3, with $\tau = 50\Delta t^2$ resolves the modelling error with an error of 2.138×10^{-3} . While the rank of RAUC-5 ($r = 32$) is larger than the RAPC-Proj run ($r = 27$), the memory requirements are the same because of the additional memory overhead of the R-SVD corrector estimation call in the RAPC algorithm. We note that this fix to the unconventional integrator is not guaranteed to work and shows that the performance of the DLRA integrators is highly dependent on the initial basis chosen and not on the initial condition itself.

Similar to the last test, RAPC-UC algorithm is the most memory efficient in the sense that the maximum memory used by RAPC-UC is the least of the five runs.

6.1.3 Rates of Convergence

Here we test that the RAPC methods and the other adaptive low-rank methods are able to recover an $\mathcal{O}(\Delta t)$ method. This is achieved by setting $\tau = \mathcal{O}(\Delta t^2)$ in order to guarantee that the update is second-order accurate locally and first-order accurate globally. The results are given in Figure 6.4 and they show setting $\tau = \mathcal{O}(\Delta t^2)$ is sufficient to achieve a second order method. Tests not included show that a tolerance of order Δt does not yield a globally first order accurate method.

6.1.4 Statistical Analysis

The purpose of this test is to provide statistical information about how variations in the randomized SVD (Algorithm 5.3) affect the error of the RAPC algorithm. Fixing a tolerance, we run Algorithm 5.4 with $\Delta t = \omega/16$ from $t = 0$ to $T = \pi$ for 50 runs and measure the error between the RAPC-UC, RAPC-Tan, and RAPC-Proj solution and the full-rank SSP-RK3 approximation. We




Figure 6.4: Rate of convergence of each low-rank method the 2D solid-body rotation problem (Section 6.1). $\Delta t = \omega/2^k$ where $k = 1, \dots, 4$ and the tolerances are shown in Row 3 of table 6.1. For RAUC-5, we set $\tau = 50\Delta t^2$.

then compute the sample mean and variance by the standard formulae:

$$\mu = \frac{1}{N} \sum_{i=1}^N X_i \quad \text{and} \quad \sigma^2 = \frac{1}{N-1} \sum_{i=1}^N (X_i - \mu)^2, \quad (6.2)$$

where $\{X_i\}_{i=1}^N$ is the collection of random samples and $N = 50$.

Method	RAPC-UC	RAPC-Tan	RAPC-Proj
τ	$880\Delta t^2$	$11255\Delta t^2$	$857\Delta t^2$
	$5.2452\text{e-}5$	$6.7085\text{e-}4$	$5.1081\text{e-}5$
μ	$2.9862\text{e-}3$	$2.9644\text{e-}3$	$3.052\text{e-}3$
σ	$6.449\text{e-}5$	$1.594\text{e-}5$	$9.214\text{e-}5$
τ	$400\Delta t^2$	$500\Delta t^2$	$200\Delta t^2$
	$2.3842\text{e-}5$	$2.9802\text{e-}5$	$1.1921\text{e-}5$
μ	$2.1328\text{e-}3$	$2.1423\text{e-}3$	$2.1967\text{e-}3$
σ	$5.563\text{-}6$	$9.191\text{e-}7$	$2.148\text{e-}5$

Table 6.2: Mean and standard deviations of 50 samples of the error of the Residual based rank-adaptive methods Section 5.3 with all three DLRA integrators in Section 4 measured against the standard full-rank SSP-RK3 approximation. The final time is $T = \pi$ and $\Delta t = \omega/16$.

Table 6.2 shows the average error and the standard deviation of the runs. The standard deviation of the runs is at least an order of magnitude from the mean-error. Thus the randomized algorithm SVD has little impact on the approximation of the solution.

6.2 Test Problem 2 - Advection with Relaxation

We consider an equation with advection and relaxation that is often used to model time-dependent radiation transport. Following, [16], the relaxation coefficient is allowed to be spatially dependent

and discontinuous. Let $\Omega_x = \Omega_y = (-1, 1)$, and consider the following initial value problem:

$$\frac{\partial u}{\partial t} + \frac{1}{\varepsilon} y \cdot \nabla_x u + \frac{\sigma(x)}{\varepsilon^2} (u - Pu) = 0 \quad (6.3a)$$

$$u(x, y, 0) = u_0(x, y) = \rho_0(x). \quad (6.3b)$$

where $\varepsilon > 0$, $Pu = \frac{1}{2} \int_{\Omega_y} u(x, y) dy$, and

$$\sigma(x) = \begin{cases} 0.02 & \text{if } x \in (-0.75, -0.25) \cup (0.25, 0.75) \\ 1 & \text{if } x \in [1, -0.75] \cup [-0.25, 0.25] \cup [0.75, 1] \end{cases} \quad (6.4)$$

We equip (6.3) with zero-inflow boundary conditions and initial data $\rho_0(x) = \chi_E(x)$ where $E = (-0.2, 0.2)$, and set $\varepsilon = 1/5$.

6.2.1 Matrix Representation

With upwind numerical fluxes, the DG discretization of (6.3) gives the following ODE for $u_h(t) \in V_h$:

$$\begin{aligned} \left(\frac{\partial u_h}{\partial t}, q_h \right)_\Omega - \frac{1}{\varepsilon} (y u_h, \nabla_x q_h)_\Omega + \left\langle y \{ \{ u_h \} \} + \frac{|y \cdot n|}{2} \llbracket u_h \rrbracket, \llbracket q_h \rrbracket \right\rangle_{\mathcal{E}_x^1 \times \Omega_y} + \langle y u_h, n q_h \rangle_{\{y \cdot n(x) > 0\}} \\ + \frac{1}{\varepsilon^2} (\sigma u_h, q_h)_\Omega - \frac{1}{\varepsilon^2} \left(\frac{\sigma}{2} (u_h, 1)_{\Omega_y}, q_h \right)_\Omega = 0 \end{aligned} \quad (6.5)$$

for all $q_h \in V_h$. Using the notation of Section 3, we may rewrite (6.5) into a matrix ODE

$$\frac{\partial U_h}{\partial t} = F(U_h) \quad (6.6)$$

where F has three terms: two for the advection operator and one for the relaxation operator $u - Pu$. Rewriting the advection operator is similar to the process in Section 3.2, while the relaxation operator can be expressed with one term. Given the bases $\{\varphi_i(x)\}_{i=1}^m$ and $\{\psi_j(y)\}_{j=1}^n$ for $V_{x,h}$ and $V_{y,h}$, let $\bar{\zeta} \in \mathbb{R}^{n \times 1}$ be the coefficient representation of the function $\zeta(y) = 1$ with respect to $\{\psi_j(y)\}_{j=1}^n$. Then

$$(\sigma u_h, q_h)_\Omega - \left(\frac{\sigma}{2} (u_h, 1)_{\Omega_y}, q_h \right)_\Omega = (AU_h(M - \frac{1}{2}M\bar{\zeta}\bar{\zeta}^T), Q_h)_F, \quad (6.7)$$

where $M_{ij} = (\psi_j, \psi_i)_{\Omega_y}$ and $A_{ij} = (\sigma \varphi_j, \varphi_i)_{\Omega_x}$. Depending on the basis, the matrix $\bar{\zeta}\bar{\zeta}^T$ may be dense, but the action $Z \rightarrow Z\bar{\zeta}\bar{\zeta}^T$ can be computed without construction of $\bar{\zeta}\bar{\zeta}^T$ thereby avoiding the storage of an $\mathcal{O}(m^2)$ object.

6.2.2 Numerical Test

We set $h_x = h_y = 1/64$ and $k_x = k_y = 1$ so that the number of degrees of freedom in each variable is $m = n = 256$. The timestep is $\Delta t = \frac{1}{12800}$, the final time is $T = 3$, and the discrete initial condition is given as a true rank-one matrix $U_h|_{t=0} = \bar{\rho}_0 \bar{\zeta}^T$ where $\bar{\rho}_0 \in \mathbb{R}^{m \times 1}$ is the coefficient representation of ρ_0 with respect to $\{\varphi_i(x)\}_{i=1}^m$. The oversampling parameter p is set to 10.

At time $t = T$, the error between the full-rank forward Euler and SSP-RK3 method is 8.95×10^{-6} . We choose the tolerance in each algorithm, by trial and error, such that the error between the low-rank approximations and the SSP-RK3 reference solution is roughly 2.5×10^{-3} . Every method is able to recover the specified accuracy, and the tolerances used to do so are given in Table 6.3.

The solution rank as a function of time for each method is plotted in Figure 6.5. The rank of every method except PCRA-Tan is similar over time. The RAUC method uses the least memory; the RAPC-Tan method and the RAST method use the most memory; and the memory usage of the other RAPC methods lie somewhere in between. The rank disparity between the RAPC-Tan integrator and the other two RAPC algorithms is larger than in the Section 6.1; this is likely because the rank of the tangent integrator (Section 4.3) can double every timestep. If the residual error is still not small enough, then the RAPC algorithm will then add more vectors onto to discrete approximation yielding a large rank solution.

Method	RAUC	RAST	RAPC-UC	RAPC-Tan	RAPC-Proj
τ	$47700\Delta t^2$	$47100\Delta t^2$	$56000\Delta t^2$	$165000\Delta t^2$	$58000\Delta t^2$
	$2.91e-4$	$2.87e-4$	$3.4e-4$	$1.00e-3$	$3.54e-4$
Error	$2.51e-3$	$2.71e-3$	$2.36e-3$	$2.22e-3$	$2.39e-3$

Table 6.3: Tolerances and errors for all five rank-adaptive methods in Section 6.2. The errors are computed at $T = \pi$ against the full-rank SSP-RK3 discretization. $\Delta t = 1/128000$.




Figure 6.5: Rank (left) and memory requirements (right) of each of the five methods over time for the advection-relaxation problem in Section 6.2. The tolerances for each adaptive algorithm are given in Table 6.3 and are chosen to produce errors near 6.5×10^{-3} when measured against the SSP-RK3 full-rank solution. The memory footprint is calculated using Table 5.1. For readability, the data is plotted at every 160 time steps. The legend labels provide the error of the method, normalized with respect to the error of the RAST algorithm.

6.2.3 Statistical Tests

It turns out that the advection-relaxation problem is much more sensitive to the errors created by the randomized SVD than the solid-body rotation problem from Section 6.1. To demonstrate this behavior, we consider a simulation with $h_x = h_y = 1/64$, $k_x = k_y = 1$, $T = 2$, and $\Delta t = 1/6400$. The tolerance for each of the RAPC methods is set to $40000\Delta t^2 \approx 9.77 \times 10^{-4}$ which produces errors that are roughly between 6×10^{-3} and 7×10^{-3} . For a given RAPC algorithm and oversampling parameter p , we run the algorithm 20 times and save the error with respect the SSP-RK3 full-rank reference solution. We repeat this procedure for all three RAPC algorithms and $p \in \{5, 7, 10, 20, 30\}$.

A statistical summary of the results is plotted in Figure 6.6. The RAPC-UC method gives the expected result: increasing the oversampling parameter increases the accuracy in the randomized SVD (see Proposition 5.1) and reduces the statistical variation. The standard deviation for the RAPC-Proj method also correlates with p , but not as strongly. Meanwhile, the RAPC-Tan results do not appear to correlate with p at all.

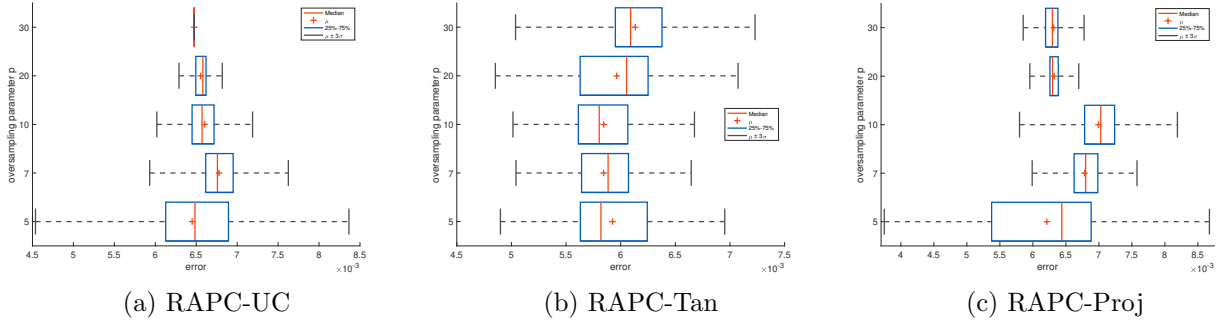


Figure 6.6: Statistical data for the advection-relaxation problem given in Section 6.2. For each oversampling parameter listed on the y -axis, 20 runs are computed. The errors are created via comparison to the SSP-RK3 full-rank reference solution. Parameters for the test are listed in Section 6.2.3.

While the randomness of the randomized SVD may increase the sensitivity of the algorithm, the process of culling process is also very sensitive to the tolerance chosen. To demonstrate, using the parameters in Section 6.2.2, we only consider the methods which are deterministic: the RAUC and RAST methods. In Figure 6.7 results are displayed for the both of these algorithms and the tolerances used are within 1% of the tolerances given in Table 6.3. The maximum rank of the RAST test varies significantly as the tolerance is changed by only a slight amount. As the tolerance increases, the algorithms should be willing to cut more and more rank which should yield lower rank solutions – this trend does happen on the longer time scales. However, during the initial rise this trend is reversed: the higher tolerance methods actually require more rank in the beginning. The errors in both methods are not monotonic with respect to the tolerance as well; the RAUC integrator gives the best error with the largest tolerance of the sample. This test shows that with all of these algorithms – random or not – choosing the tolerance is a delicate process and can greatly affect the rank and accuracy of the corresponding solution.




Figure 6.7: Rank for the radiation transport problem given in Section 6.2 for the RAUC (left) and RAST (right) algorithms with varied tolerances given in the legend. The errors are created via comparison to the SSP-RK3 full-rank reference solution.

6.3 Test Problem 3 - Rotation, Relaxation, and Localized Sources

We consider the following PDE with sources:

$$\frac{\partial u}{\partial t} - y \cdot \nabla_x u + x \cdot \nabla_y u + \frac{1}{\varepsilon}(u - Pu) = S, \quad (x, y) \in \Omega, \quad t > 0; \quad (6.8a)$$

$$u(x, y, t) = 0, \quad (x, y) \in \partial\Omega_x^-, \quad t > 0; \quad (6.8b)$$

$$u(x, y, t) = 0, \quad (x, y) \in \partial\Omega_y^-, \quad t > 0; \quad (6.8c)$$

$$u(x, y, 0) = u_0(x, y), \quad (x, y) \in \Omega \quad (6.8d)$$

where $\varepsilon > 0$, $Pu = \frac{1}{2} \int_{\Omega_y} u(x, y) dy$, u_0 is given in (6.1), and S is given by

$$S(x, y) = -(\chi_{E_1}(x) + \chi_{E_2}(x) + \chi_{E_3}(x) + \chi_{E_4}(x) + \chi_{E_5}(x)) \cdot (\chi_{E_1}(y) + \chi_{E_2}(y) + \chi_{E_3}(y) + \chi_{E_4}(y) + \chi_{E_5}(y)) \quad (6.9)$$

where $E_i = \left[\frac{20(i-3)-3}{60}, \frac{20(i-3)+3}{60} \right]$ for $i = 1, \dots, 5$. The discretization and matrix evaluations of all terms on the left-hand side of (6.8a) have been treated in Section 3 and Section 6.2. Here S_h , the matrix representation of the L^2 projection of S onto V_h is a rank-one matrix and thus can be stored in a low-memory format.

For this test we set $h_x = h_y = 1/128$ and $k_x = k_y = 0$ which yield 256 degrees of freedom in each direction. Additionally, we set $\varepsilon = 1/5$, $T = \pi$, and $\Delta t = 1/1024$. The goal of this test is to see how well the adaptive integrators can resolve the high frequency source S . The oversampling parameter p is set to 15.

At $t = \pi$, the error between the full-rank forward Euler and SSP-RK3 approximations is 3.94×10^{-5} . The tolerance for each of the five methods is given in Table 6.4 and is chosen to produce an error of approximately 1.50×10^{-3} . Figure 6.8 provides the rank of the methods over time. In this case the rank of each method quickly rises and then plateaus as the projection operator forces the solution to a steady state. All methods but PCRA-Tan arrive at the same final rank. Similar to Section 6.2, the RAUC method is able to capture the equilibrium with the smallest memory footprint.

Method	RAUC	RAST	RAPC-UC	RAPC-Tan	RAPC-Proj
τ	$90\Delta t^2$	$100\Delta t^2$	$125\Delta t^2$	$1300\Delta t^2$	$120\Delta t^2$
Error	$1.5e-3$	$1.50e-3$	$1.50e-3$	$1.72e-3$	$1.50e-3$

Table 6.4: Tolerances and Errors for all five rank-adaptive methods in Section 6.3. The errors are computed at $T = \pi$ against the full-rank SSP-RK3 discretization. $\Delta t = 1/1024$.




Figure 6.8: Rank (left) and memory requirements (right) of each of the five methods over time for the problem given in Section 6.3 with $\Delta t = 1/1024$. The tolerances for the adaptive algorithms are given in Table 6.4 and is chosen to produce errors near 1.50×10^{-3} when measured against an SSP-RK3 reference solution. The memory footprint is calculated using Table 5.1. For readability, the data is plotted at every 65 timesteps. The legend labels each algorithm as well as the ratio of the method’s error over the RAST algorithm error which is used as a reference solution.

7 Conclusion

In this paper, we introduced a predictor-corrector strategy that modifies the rank of the DLRA solution by building a low-rank approximation of the error made against the full rank forward Euler method. We presented several numerical results showing our method is accurate in resolving the modelling error of the low-rank projection and that the methods rank is comparable with other rank-adaptive DLRA methods. The paper also included a discussion on how a variety of PDEs discretized by the discontinuous Galerkin method can be rewritten to fit the DLRA framework. Future topics include:

- Building higher order methods in time using the predictor-corrector strategy with forward Euler timestepping as a building block. This approach is similar to the work in [20].
- Using more advanced randomized techniques, such as power iterations [15], to diminish the effects of randomness on the solution.
- Extensions to dynamic low-rank approximation on higher order tensors [22].




Figure 6.9: Plots of the discrete approximations at $t = \pi$ for the problem given in Section 6.3 with the RAPC-Proj Algorithm (left) with a rank of 26 and the full-rank SSP-RK3 method (right). The tolerance for the adaptive algorithm is given in Table 6.4.

8 Appendix

Proof of Proposition 3.4. We focus on (3.15e)-(3.15h). Let $w_h = \varphi_i \psi_j$ and $q_h = \varphi_k \psi_l$. The portion of \mathcal{A} that (3.15e)-(3.15h) separates, namely,

$$-(xw_h, \nabla_x q_h)_\Omega + \langle x \{ \{ w_h \} \} \rangle + \frac{1}{2} |x| \llbracket [w_h], [q_h] \rrbracket_{\Omega_x \times \mathcal{E}_{y,h}^I} + \langle xw_h, nq_h \rangle_{\partial\Omega_y^+} \quad (8.1)$$

is not separable because the flux depends on the sign of x . To create the minimal number of separable terms, we split the mass integral $(x\varphi_i, \psi_k)_{\Omega_x}$ about $x = 0$ to obtain $(x\varphi_i, \psi_k)_{\Omega_x} = (x\varphi_i, \psi_k)_{\{x>0\}} + (x\varphi_i, \psi_k)_{\{x<0\}}$. Splitting (8.1) yields

$$-(xw_h, \nabla_y q_h)_\Omega + \langle x \{ \{ w_h \} \} \rangle + \frac{1}{2} |x| \llbracket [w_h], [q_h] \rrbracket_{\Omega_x \times \mathcal{E}_{y,h}^I} + \langle xw_h, nq_h \rangle_{\partial\Omega_y^+} = I_1 + I_2 \quad (8.2)$$

where

$$I_1 = -(xw_h, \nabla_y q_h)_{\{x>0\} \times \Omega_y} + \langle x \{ \{ w_h \} \} \rangle + \frac{1}{2} |x| \llbracket [w_h], [q_h] \rrbracket_{\{x>0\} \times \mathcal{E}_{y,h}^I} + \langle xw_h, nq_h \rangle_{\partial\Omega_y^+ \cap \{x>0\} \times \partial\Omega_y} \quad (8.3)$$

$$I_2 = -(xw_h, \nabla_y q_h)_{\{x<0\} \times \Omega_y} + \langle x \{ \{ w_h \} \} \rangle + \frac{1}{2} |x| \llbracket [w_h], [q_h] \rrbracket_{\{x<0\} \times \mathcal{E}_{y,h}^I} + \langle xw_h, nq_h \rangle_{\partial\Omega_y^+ \cap \{x<0\} \times \partial\Omega_y} \quad (8.4)$$

We focus on I_1 . For $x > 0$, $|x| = x$ and thus the flow direction is independent of x . Moreover, (2.5) and the sign of x implies the outflow boundary is only on the top of Ω_y . Thus

$$\partial\Omega_y^+ \cap \{x > 0\} \times \partial\Omega_y = \{x > 0\} \times \{y = L\} \quad (8.5)$$

Substituting (8.5) into (8.3) yields

$$\begin{aligned}
I_1 &= -(x\varphi_i\psi_j, \nabla_y(\varphi_k\psi_l))_{\{x>0\}\times\Omega_y} + \langle x\{\{\varphi_i\psi_j\}\} + \frac{1}{2}|x|[\varphi_i\psi_j], [\varphi_k\psi_l]\rangle_{\{x>0\}\times\mathcal{E}_{y,h}^I} \\
&\quad + \langle x\varphi_i\psi_j, n\varphi_k\psi_l\rangle_{\{x>0\}\times\{y=L\}} \\
&= -(x\varphi_i, \psi_k)_{\{x>0\}}(\psi_j, \nabla_y\psi_l)_{\Omega_y} \\
&\quad + (x\varphi_i, \psi_k)_{\{x>0\}} \langle \{\{\psi_j\}\} + \frac{1}{2}[\psi_j], [\psi_l]\rangle_{\mathcal{E}_{y,h}^I} \\
&\quad + (x\varphi_i, \psi_k)_{\{x>0\}} \langle \psi_j n, \psi_l \rangle_{\{y=L\}} \\
&= (x\varphi_i, \psi_k)_{\{x>0\}} \left(-(\psi_j, \nabla_y\psi_l)_{\Omega_y} + \langle \{\{\psi_j\}\} + \frac{1}{2}[\psi_j], [\psi_l]\rangle_{\mathcal{E}_{y,h}^I} + \langle \psi_j n, \psi_l \rangle_{\{y=L\}} \right) \\
&= \mathcal{B}_{3,x}(\varphi_i, \varphi_k)\mathcal{B}_{3,y}(\psi_j, \psi_l).
\end{aligned} \tag{8.6}$$

We now focus on I_2 . Since $x < 0$, $|x| = -x$ and

$$\partial\Omega_y^+ \cap \{x < 0\} \times \partial\Omega_y = \{x < 0\} \times \{y = -L\} \tag{8.7}$$

Therefore similar to the derivation of (8.6) we have

$$\begin{aligned}
I_2 &= -(x\varphi_i\psi_j, \nabla_y(\varphi_k\psi_l))_{\{x<0\}\times\Omega_y} + \langle x\{\{\varphi_i\psi_j\}\} + \frac{1}{2}|x|[\varphi_i\psi_j], [\varphi_k\psi_l]\rangle_{\{x<0\}\times\mathcal{E}_{y,h}^I} \\
&\quad + \langle x\varphi_i\psi_j, n\varphi_k\psi_l\rangle_{\{x<0\}\times\{y=-L\}} \\
&= -(x\varphi_i, \psi_k)_{\{x<0\}}(\psi_j, \nabla_y\psi_l)_{\Omega_y} \\
&\quad + (x\varphi_i, \psi_k)_{\{x<0\}} \langle \{\{\psi_j\}\} - \frac{1}{2}[\psi_j], [\psi_l]\rangle_{\mathcal{E}_{y,h}^I} \\
&\quad + (x\varphi_i, \psi_k)_{\{x<0\}} \langle \psi_j n, \psi_l \rangle_{\{y=-L\}} \\
&= (x\varphi_i, \psi_k)_{\{x<0\}} \left(-(\psi_j, \nabla_y\psi_l)_{\Omega_y} + \langle \{\{\psi_j\}\} - \frac{1}{2}[\psi_j], [\psi_l]\rangle_{\mathcal{E}_{y,h}^I} + \langle \psi_j n, \psi_l \rangle_{\{y=-L\}} \right) \\
&= \mathcal{B}_{4,x}(\varphi_i, \varphi_k)\mathcal{B}_{4,y}(\psi_j, \psi_l).
\end{aligned} \tag{8.8}$$

(3.15a)-(3.15d) can be similarly shown. The proof is complete. \square

Proof of Proposition 3.5. By (3.2), (3.1), (3.17), and (2.8), we have

$$\begin{aligned}
(G_h(t))_{ij} &= g_h(t)^{ij} = (g_h(t), \varphi_i\psi_j)_\Omega = \mathcal{G}(\varphi_i\psi_j, t) \\
&= \langle yg_x(\cdot, t), n\varphi_i\psi_j \rangle_{\partial\Omega_x^-} + \langle -xg_y(\cdot, t), n\varphi_i\psi_j \rangle_{\partial\Omega_y^-}.
\end{aligned} \tag{8.9}$$

Spitting the terms in (8.9) about $x = 0$ and $y = 0$ yields

$$\begin{aligned}
(G_h(t))_{ij} &= \langle yg_x(\cdot, \cdot, t), n\varphi_i\psi_j \rangle_{\partial\Omega_x^- \cap \partial\Omega_x \times \{y>0\}} + \langle yg_x(\cdot, \cdot, t), n\varphi_i\psi_j \rangle_{\partial\Omega_x^- \cap \partial\Omega_x \times \{y<0\}} \\
&\quad + \langle -xg_y(\cdot, \cdot, t), n\varphi_i\psi_j \rangle_{\partial\Omega_y^- \cap \{x>0\} \times \partial\Omega_y} + \langle -xg_y(\cdot, \cdot, t), n\varphi_i\psi_j \rangle_{\partial\Omega_y^- \cap \{x<0\} \times \partial\Omega_y} \\
&= \langle yg_x(\cdot, \cdot, t), n\varphi_i\psi_j \rangle_{\{x=L\} \times \{y>0\}} + \langle yg_x(\cdot, \cdot, t), n\varphi_i\psi_j \rangle_{\{x=-L\} \times \{y<0\}} \\
&\quad + \langle -xg_y(\cdot, \cdot, t), n\varphi_i\psi_j \rangle_{\{x>0\} \times \{y=-L\}} + \langle -xg_y(\cdot, \cdot, t), n\varphi_i\psi_j \rangle_{\{x<0\} \times \{y=L\}} \\
&= \langle 1, n\varphi_i \rangle_{\{x=L\}} (yg_x(L, \cdot, t), \psi_j)_{\{y>0\}} + \langle 1, n\varphi_i \rangle_{\{x=-L\}} (yg_x(-L, \cdot, t), \psi_j)_{\{y<0\}} \\
&\quad + \langle -xg_y(\cdot, -L, t), \varphi_i \rangle_{\{x>0\}} \langle 1, n\psi_j \rangle_{\{y=-L\}} + \langle -xg_y(\cdot, L, t), \varphi_i \rangle_{\{x<0\}} \langle 1, n\psi_j \rangle_{\{y=L\}} \\
&:= c_i^1 d_j^1(t) + c_i^2 d_j^2(t) + c_i^3(t) d_j^3 + c_i^4(t) d_j^4.
\end{aligned} \tag{8.10}$$

where $c^1, c^2, c^3, c^4 \in \mathbb{R}^m$ and $d^1, d^2, d^3, d^4 \in \mathbb{R}^n$. (8.10) yields

$$G_h(t) = c^1(t)(d^1)^T + c^2(t)(d^2)^T + c^3(t)(d^3)^T + c^4(t)(d^4)^T = C(t)S(t)D(t)^T \quad (8.11)$$

where

$$C(t) = [c^1(t), c^2(t), c^3, c^4], \quad S(t) = I_{4 \times 4}, \quad D(t) = [d^1, d^2, d^3(t), d^4(t)]. \quad (8.12)$$

Since C and D are not guaranteed to be linearly independent, then it is readily seen that the rank of G_h is at most 4. The proof is complete. \square

References

- [1] Michael H Beck, Andreas Jäckle, Graham A Worth, and H-D Meyer. The multiconfiguration time-dependent hartree (mctdh) method: a highly efficient algorithm for propagating wavepackets. *Physics reports*, 324(1):1–105, 2000.
- [2] Gianluca Ceruti, Jonas Kusch, and Christian Lubich. A rank-adaptive robust integrator for dynamical low-rank approximation. *BIT Numerical Mathematics*, pages 1–26, 2022.
- [3] Gianluca Ceruti and Christian Lubich. An unconventional robust integrator for dynamical low-rank approximation. *BIT Numerical Mathematics*, pages 1–22, 2021.
- [4] Alec Dektor, Abram Rodgers, and Daniele Venturi. Rank-adaptive tensor methods for high-dimensional nonlinear pdes. *Journal of Scientific Computing*, 88(2):1–27, 2021.
- [5] Zhiyan Ding, Lukas Einkemmer, and Qin Li. Dynamical low-rank integrator for the linear boltzmann equation: error analysis in the diffusion limit. *SIAM Journal on Numerical Analysis*, 59(4):2254–2285, 2021.
- [6] Paul AM Dirac. Note on exchange phenomena in the thomas atom. In *Mathematical proceedings of the Cambridge philosophical society*, volume 26, pages 376–385. Cambridge University Press, 1930.
- [7] Lukas Einkemmer, Jingwei Hu, and Yubo Wang. An asymptotic-preserving dynamical low-rank method for the multi-scale multi-dimensional linear transport equation. *Journal of Computational Physics*, 439:110353, 2021.
- [8] Lukas Einkemmer, Jingwei Hu, and Lexing Ying. An efficient dynamical low-rank algorithm for the boltzmann-bgk equation close to the compressible viscous flow regime. *SIAM Journal on Scientific Computing*, 43(5):B1057–B1080, 2021.
- [9] Lukas Einkemmer and Ilon Joseph. A mass, momentum, and energy conservative dynamical low-rank scheme for the vlasov equation. *Journal of Computational Physics*, 443:110495, 2021.
- [10] Lukas Einkemmer and Christian Lubich. A low-rank projector-splitting integrator for the vlasov–poisson equation. *SIAM Journal on Scientific Computing*, 40(5):B1330–B1360, 2018.
- [11] J Frenkel. *Wave mechanics*, clarendon, 1934.
- [12] Lars Grasedyck, Daniel Kressner, and Christine Tobler. A literature survey of low-rank tensor approximation techniques. *GAMM-Mitteilungen*, 36(1):53–78, 2013.

- [13] Wei Guo and Jing-Mei Qiu. A conservative low rank tensor method for the vlasov dynamics. *arXiv preprint arXiv:2201.10397*, 2022.
- [14] Wei Guo and Jing-Mei Qiu. A low rank tensor representation of linear transport and nonlinear vlasov solutions and their associated flow maps. *Journal of Computational Physics*, page 111089, 2022.
- [15] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [16] Cory D Hauck and Robert B Lowrie. Temporal regularization of the p_n equations. *Multiscale Modeling & Simulation*, 7(4):1497–1524, 2009.
- [17] Jingwei Hu and Yubo Wang. An adaptive dynamical low rank method for the nonlinear boltzmann equation. *Journal of Scientific Computing*, 92(2):1–24, 2022.
- [18] Tobias Jahnke and Wilhelm Huisinga. A dynamical low-rank approach to the chemical master equation. *Bulletin of mathematical biology*, 70(8):2283–2302, 2008.
- [19] Emil Kieri, Christian Lubich, and Hanna Walach. Discretized dynamical low-rank approximation in the presence of small singular values. *SIAM Journal on Numerical Analysis*, 54(2):1020–1038, 2016.
- [20] Emil Kieri and Bart Vandereycken. Projection methods for dynamical low-rank approximation of high-dimensional problems. *Computational Methods in Applied Mathematics*, 19(1):73–92, 2019.
- [21] Othmar Koch and Christian Lubich. Dynamical low-rank approximation. *SIAM Journal on Matrix Analysis and Applications*, 29(2):434–454, 2007.
- [22] Othmar Koch and Christian Lubich. Dynamical tensor approximation. *SIAM Journal on Matrix Analysis and Applications*, 31(5):2360–2375, 2010.
- [23] Jonas Kusch, Gianluca Ceruti, Lukas Einkemmer, and Martin Frank. Dynamical low-rank approximation for burgers’equation with uncertainty. *International Journal for Uncertainty Quantification*, 12(5), 2022.
- [24] Christian Lubich. On variational approximations in quantum molecular dynamics. *Mathematics of computation*, 74(250):765–779, 2005.
- [25] Christian Lubich and Ivan V Oseledets. A projector-splitting integrator for dynamical low-rank approximation. *BIT Numerical Mathematics*, 54(1):171–188, 2014.
- [26] MATLAB. *version 9.8.0 (R2020b)*. The MathWorks Inc., Natick, Massachusetts, 2020b.
- [27] Zhuogang Peng and Ryan G McClarren. A high-order/low-order (holo) algorithm for preserving conservation in time-dependent low-rank transport calculations. *Journal of Computational Physics*, 447:110672, 2021.
- [28] Zhuogang Peng and Ryan G McClarren. A sweep-based low-rank method for the discrete ordinate transport equation. *arXiv preprint arXiv:2206.14404*, 2022.

- [29] Zhuogang Peng, Ryan G McClarren, and Martin Frank. A low-rank method for two-dimensional time-dependent radiation transport calculations. *Journal of Computational Physics*, 421:109735, 2020.
- [30] Hayden Schaeffer, Russel Caflisch, Cory D Hauck, and Stanley Osher. Sparse dynamics for partial differential equations. *Proceedings of the National Academy of Sciences*, 110(17):6634–6639, 2013.
- [31] Kai Schneider, NK-R Kevlahan, and Marie Farge. Comparison of an adaptive wavelet method and nonlinearly filtered pseudospectral methods for two-dimensional turbulence. *Theoretical and computational fluid dynamics*, 9(3):191–206, 1997.
- [32] Steffen Schotthöfer, Emanuele Zangrando, Jonas Kusch, Gianluca Ceruti, and Francesco Tudisco. Low-rank lottery tickets: finding efficient low-rank neural networks via matrix differential equations. *arXiv preprint arXiv:2205.13571*, 2022.