

Low-rank variance reduction for uncertain radiative transfer with control variates

Chinmay Patwardhan, Pia Stammer, Emil Løvbak,
Jonas Kusch, Sebastian Krumscheid

CRC Preprint 2025/2, January 2025

KARLSRUHE INSTITUTE OF TECHNOLOGY

CRC 1173



Participating universities



Funded by



Low-rank variance reduction for uncertain radiative transfer with control variates

Chinmay Patwardhan^{1†}, Pia Stammer^{2†}, Emil Løvbak¹, Jonas Kusch³, and Sebastian Krumscheid¹

¹ Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany
{chinmay.patwardhan, emil.loevbak, sebastian.krumscheid}@kit.edu

² Delft University of Technology (TU Delft), Delft, Netherlands
p.k.stammer@tudelft.nl

³ Norwegian University of Life Sciences (NMBU), Ås, Norway
jonas.kusch@nmbu.no

Abstract. The radiative transfer equation models various physical processes ranging from plasma simulations to radiation therapy. In practice, these phenomena are often subject to uncertainties. Modeling and propagating these uncertainties requires accurate and efficient solvers for the radiative transfer equations. Due to the equation’s high-dimensional phase space, fine-grid solutions of the radiative transfer equation are computationally expensive and memory-intensive. In recent years, dynamical low-rank approximation has become a popular method for solving kinetic equations due to the development of computationally inexpensive, memory-efficient and robust algorithms like the augmented basis update & Galerkin integrator. In this work, we propose a low-rank Monte Carlo estimator and combine it with a control variate strategy based on multi-fidelity low-rank approximations for variance reduction. We investigate the error analytically and numerically and find that a joint approach to balance rank and grid size is necessary. Numerical experiments further show that the efficiency of estimators can be improved using dynamical low-rank approximation, especially in the context of control variates.

Keywords: dynamical low-rank approximation, reduced-order modeling, Monte Carlo estimation, control variates, uncertainty quantification

1 Introduction

Kinetic equations play a key role in modeling various natural phenomena from plasma physics [14, 22, 47] to radiation therapy planning [12, 31]. When performing numerical simulations, one often encounters uncertainties due to inaccurate measurements or unknown parameters. In forward uncertainty quantification (UQ), one assumes a probability distribution on these uncertainties and computes quantities of interest, e.g., the expectation or variance, of the corresponding computational results. Despite advances in both computing power and efficient solution methods, numerically solving such equations with uncertainties is still a computationally expensive and memory-intensive

[†]Equal contribution.

task. Hence, a combined approach of efficient sampling methods and efficient solvers is required. To this end, we present a novel combination of hierarchical Monte Carlo sampling methods with the dynamical low-rank approximation [26].

We consider the radiative transfer equation (RTE) which models the transport of particles through a background material

$$\partial_t \psi(t, x, \Omega) + \Omega \cdot \nabla_x \psi(t, x, \Omega) = \sigma_s(x) \left(\int_{\mathbb{S}^2} \psi(t, x, \Omega') d\Omega' - \psi(t, x, \Omega) \right), \quad (1a)$$

$$\psi(t = t_0, x, \Omega) = \psi_0(x, \Omega). \quad (1b)$$

Here, $\psi(t, x, \Omega)$ denotes the particle density at time $t \in \mathbb{R}_{\geq 0}$, spatial position $x \in \mathcal{D}_x$ and travelling in direction $\Omega \in \mathbb{S}^2$. The particle density modeled by (1) is subject to discontinuous changes in direction due to collisions with the background material with scattering rate $\sigma_s(x)$. Moreover, we assume that the particle density does not reach the domain boundary. Due to the 6-dimensional phase space, numerically solving problems involving equations of the form (1) quickly becomes computationally expensive. Hence, a multitude of numerical methods have been developed to tackle such problems, e.g., particle-based Monte Carlo simulation [3] and spectral methods, such as moment models [17]. For an overview of methods, we refer to [11] and references therein.

Recently, dynamical low-rank approximation [26] has been widely used to solve kinetic equations in plasma physics [10, 13], thermal radiative transfer [2, 39] as well as problems related to nuclear power [32] or medical physics [31, 44]. The core idea of dynamical low-rank approximation (DLRA) for radiative transfer is to restrict the underlying radiative transfer equation to low-rank solutions. This is achieved by projecting the dynamics onto the tangent space of the manifold of low-rank functions, leading to novel evolution equations for the low-rank factors of the solution. However, the rank is often overestimated to ensure it is sufficiently high. Overestimated ranks lead to ill-conditioned evolution equations [26], requiring the derivation of novel time integrators tailored to the underlying structure of the manifold of low-rank solutions. The most frequently used robust time integrators for DLRA include projector–splitting [24, 33] and basis-update & Galerkin (BUG) integrators [5–8, 28]. Although numerous works use dynamical low-rank approximation to simulate deterministic kinetic problems efficiently, approaches to include uncertainties as an additional dimension of the phase space have also been explored [29, 38, 42, 43]. Here, most notably, [43] uses a DLRA tensor integrator to include uncertainties into kinetic simulations. However, such approaches require an intrusive modification of a solver’s code. Further, especially in higher dimensions their efficiency depends strongly on the design of the tensor structures and is not yet well explored.

Monte Carlo methods are frequently used for forward uncertainty quantification due to their non-intrusive nature and ease of parallelization. They also tend to be favorable for higher-dimensional uncertainties, as they avoid the construction of parameter-space grids. However, direct Monte Carlo sampling is often prohibitively expensive in practical applications. Hence, variance reduction techniques such as control variates [46] or Multilevel Monte Carlo methods (MLMC) [15, 16, 21] are useful to improve the efficiency of such approaches. While the naive combination of a dynamical low-rank deterministic solver with a Monte Carlo type approach for uncertainty quantification is

straightforward, the analysis and optimal design of a more sophisticated scheme requires knowledge of error and cost convergence rates depending on rank.

Initially, research on using MLMC for stochastic PDE models focused mainly on elliptic problems, notably, Darcy flows with an uncertain diffusion coefficient [9]. However, extensions to an increasing amount of other problems exist, such as parabolic [1] or hyperbolic problems [23, 36]. In particular, MLMC has been applied for the simulation of kinetic equations with particle methods, both at the trajectory level [34, 35, 37] and the ensemble level [20, 41, 45]. Methodologically, work on stochastic PDE problems has focused on solving the PDE with different grid resolutions at different levels in the multilevel hierarchy. In this case, it is known that geometric meshes are, in general, close to optimal [18]. However, recent work has advanced MLMC simulations for elliptic problems by integrating hierarchical solvers. An example being multigrid solvers [27], where it has been shown that one can reduce the total computational cost by recycling work from coarser levels [40]. We also note that multi-index extensions become relevant when considering problems with infinite-dimensional uncertainties, as a hierarchy of approximations is then needed in the random variable [19].

In this work, we combine variance reduction techniques with multi-fidelity dynamical low-rank approximations for the radiation transport equation. We investigate the error of the expectation in terms of both the approximation rank and the spatio-temporal mesh. We then demonstrate how to perform efficient correlated sampling based on the dynamical low-rank method and demonstrate that a control variate can significantly speed up the computation of accurate expectations. The remainder of this paper is structured as follows: In Section 2, we introduce dynamical low-rank approximation [26] as an efficient reduced model-order approach for time-dependent problems and describe the augmented basis update & Galerkin integrator [6]. In Section 3, we extend the well-known robust error bound for the dynamical low-rank approximation [8, 24] to the probabilistic setting based on the proof from [5] and provide insights into the rank-error relation. We then present a low-rank Monte Carlo and control variate strategy based on these insights. In Section 4, we present numerical results for a radiation transport problem that back up our theoretical claims and demonstrates the speedup of our control variate strategy. While our interest lies in higher-dimensional problems, we focus here on a 1D test problem with an uncertain initial condition, to limit computational costs. However, our approach straightforwardly applies to other uncertainties and higher-dimensional cases. Finally, in Section 5, we draw conclusions on our results and discuss the potential for a future extension to multilevel Monte Carlo.

2 Dynamical low-rank approximation

This section summarizes the fundamentals of dynamical low-rank approximation (DLRA) laid out in [26] and presents the augmented basis-update & Galerkin [6] integrator. We start by assuming a discretization of (1) in space and direction of travel and write it as

$$\dot{\Psi}(t) = \mathbf{F}(t, \Psi(t)), \quad \Psi(t_0) = \Psi_0, \quad (2)$$

where $\Psi \in \mathbb{R}^{m \times n}$, $m, n \gg 1$, is the discretized particle density and $\dot{\Psi} = \frac{d}{dt}\Psi$ is the derivative with respect to time. Here, m, n represent the number of spatial and

directional discretization points, respectively. Then DLRA [26] evolves $\Psi(t)$ on the low-rank manifold \mathcal{M}_r , of $m \times n$ rank- r matrices. It does so by projecting the dynamics of the problem onto the tangent space of \mathcal{M}_r .

Any given low-rank approximation $\mathbf{Y}(t) \in \mathcal{M}_r$ to $\Psi(t)$ can be factorized as

$$\mathbf{Y}(t) = \mathbf{X}(t)\mathbf{S}(t)\mathbf{V}(t)^\top,$$

where $\mathbf{X}(t) \in \mathbb{R}^{m \times r}$, $\mathbf{V}(t) \in \mathbb{R}^{n \times r}$ have orthonormal columns and $\mathbf{S}(t) \in \mathbb{R}^{r \times r}$ is invertible. Let $\mathcal{T}_{\mathbf{Y}}\mathcal{M}_r$ denote the tangent space of the low-rank manifold \mathcal{M}_r at \mathbf{Y} . To ensure that the low-rank approximation $\mathbf{Y}(t)$ remains of rank r at any given time, DLRA solves the optimization problem

$$\min_{\dot{\mathbf{Y}}(t) \in \mathcal{T}_{\mathbf{Y}(t)}\mathcal{M}_r} \|\dot{\mathbf{Y}}(t) - \mathbf{F}(t, \mathbf{Y}(t))\|,$$

where $\|\cdot\|$ is the Frobenius norm. This distance is minimized by the orthogonal projection of the right-hand side of (2) onto the tangent space $\mathcal{T}_{\mathbf{Y}(t)}\mathcal{M}_r$, denoted by $\mathbf{P}(\mathbf{Y}(t))$, [26, Lemma 4.1], i.e.

$$\dot{\mathbf{Y}} = \mathbf{P}(\mathbf{Y}(t))\mathbf{F}(t, \mathbf{Y}(t)), \quad \text{where} \quad \mathbf{P}(\mathbf{Y})\mathbf{Z} = \mathbf{X}\mathbf{X}^\top\mathbf{Z} - \mathbf{X}\mathbf{X}^\top\mathbf{Z}\mathbf{V}\mathbf{V}^\top + \mathbf{Z}\mathbf{V}\mathbf{V}^\top.$$

Thus the evolution equations for the factors ensuring $\mathbf{Y}(t) \in \mathcal{M}_r$ for all t , are [26]

$$\begin{aligned} \dot{\mathbf{S}}(t) &= \mathbf{X}(t)^\top \mathbf{F}(t, \mathbf{Y}(t)) \mathbf{V}(t), \\ \dot{\mathbf{X}}(t) &= (\mathbf{I} - \mathbf{X}(t)\mathbf{X}(t)^\top) \mathbf{F}(t, \mathbf{Y}(t)) \mathbf{V}(t) \mathbf{S}(t)^{-1}, \\ \dot{\mathbf{V}}(t) &= (\mathbf{I} - \mathbf{V}(t)\mathbf{V}(t)^\top) \mathbf{F}(t, \mathbf{Y}(t))^\top \mathbf{X}(t) \mathbf{S}(t)^{-\top}. \end{aligned}$$

Since the rank of the solution is not known beforehand, it is often over-approximated in practice. This introduces several near-zero singular values into the approximation, causing conventional time integrators to fail in achieving convergence. Several robust approaches have been developed to counter this problem, such as the projector-splitting [33] or basis-update & Galerkin integrators [6–8] and projection-based methods [25]. In this work, we use the augmented BUG integrator [6] due to its favorable stability properties for radiation transport [30].

Let $\mathbf{Y}_k := \mathbf{Y}(t_k)$, $\mathbf{X}_k := \mathbf{X}(t_k)$, $\mathbf{S}_k := \mathbf{S}(t_k)$, $\mathbf{V}_k := \mathbf{V}(t_k)$, where $t_k = t_0 + kh$ for some starting time $t_0 \in \mathbb{R}_{\geq 0}$ and step size $h > 0$. Then the rank- r approximation at t_0 is $\mathbf{Y}_0 = \mathbf{X}_0\mathbf{S}_0\mathbf{V}_0^\top$. In the following, we outline one update step according to [6], where the factors \mathbf{X}_0 , \mathbf{S}_0 , \mathbf{V}_0 are updated to \mathbf{X}_1 , \mathbf{S}_1 , \mathbf{V}_1 at time t_1 in three sub-steps. Since the goal is to gain a better understanding of uncertainty propagation in the low-rank framework, the rank of the approximation is not adapted according to a truncation tolerance, as proposed in [6], but truncated to a fixed rank r .

1. Basis augmentation

K-step: For $\mathbf{K}(t) = \mathbf{X}(t)\mathbf{S}(t)$, integrate from t_0 to t_1

$$\dot{\mathbf{K}}(t) = \mathbf{F}(t, \mathbf{K}(t)\mathbf{V}_0^\top)\mathbf{V}_0, \quad \mathbf{K}(t_0) = \mathbf{X}_0\mathbf{S}_0,$$

compute $\widehat{\mathbf{X}} \in \mathbb{R}^{m \times 2r}$ as an orthonormal basis of $[\mathbf{K}(t_1), \mathbf{X}_0]$ and store $\mathbf{M} = \widehat{\mathbf{X}}^\top \mathbf{X}_0$.

L-step: For $\mathbf{L}(t) = \mathbf{V}(t)\mathbf{S}(t)^\top$, integrate from t_0 to t_1

$$\dot{\mathbf{L}}(t) = \mathbf{F}(t, \mathbf{X}_0\mathbf{L}(t)^\top)^\top \mathbf{X}_0, \quad \mathbf{L}(t_0) = \mathbf{V}_0\mathbf{S}_0^\top,$$

compute $\widehat{\mathbf{V}} \in \mathbb{R}^{n \times 2r}$ as an orthonormal basis of $[\mathbf{L}(t_1), \mathbf{V}_0]$ and store $\mathbf{N} = \widehat{\mathbf{V}}^\top \mathbf{V}_0$.

2. Galerkin step

S-step: Integrate from t_0 to t_1

$$\hat{\mathbf{S}}(t) = \widehat{\mathbf{X}}^\top \mathbf{F}(t, \widehat{\mathbf{X}}\widehat{\mathbf{S}}(t)\widehat{\mathbf{V}}^\top)\widehat{\mathbf{V}}, \quad \widehat{\mathbf{S}}(t_0) = \mathbf{M}\mathbf{S}_0\mathbf{N}^\top. \quad (3)$$

3. Truncation to rank r

Compute the singular value decomposition of $\widehat{\mathbf{S}}(t_1) = \mathbf{P}\mathbf{\Sigma}\mathbf{Q}^\top$. Let $\mathbf{S}_1 \in \mathbb{R}^{r \times r}$ be the diagonal matrix with the r largest singular values of $\mathbf{\Sigma}$ and set $\mathbf{P}_1 \in \mathbb{R}^{2r \times r}$ and $\mathbf{Q}_1 \in \mathbb{R}^{2r \times r}$ contain the first r columns of \mathbf{P} and \mathbf{Q} . Then, $\mathbf{X}_1 = \widehat{\mathbf{X}}\mathbf{P}_1$ and $\mathbf{V}_1 = \widehat{\mathbf{V}}\mathbf{Q}_1$.

Finally, the approximation at t_1 is set as $\mathbf{Y}_1 = \mathbf{X}_1\mathbf{S}_1\mathbf{V}_1^\top$. Note, if $\widehat{\mathbf{Y}}(t) = \widehat{\mathbf{X}}\widehat{\mathbf{S}}(t)\widehat{\mathbf{V}}^\top$ the augmented solution is obtained by solving the following reformulation of (3):

$$\hat{\mathbf{Y}}(t) = \widehat{\mathbf{X}}\widehat{\mathbf{X}}^\top \mathbf{F}(t, \widehat{\mathbf{Y}}(t))\widehat{\mathbf{V}}\widehat{\mathbf{V}}^\top, \quad \widehat{\mathbf{Y}}(t_0) = \mathbf{Y}_0. \quad (4)$$

3 Low-rank estimators

In this section, we introduce uncertainty to (2) and investigate the propagation of uncertainty through the solution. We do so by constructing Monte Carlo and control variate estimators based on low-rank approximation of the solution. Furthermore, we show that when used together with Monte Carlo sampling, the augmented BUG integrator is robust to the presence of small singular values.

Let ν be a scalar random variable with probability density function $p(\nu)$. Then if $\Psi(t; \nu) \in \mathbb{R}^{m \times n}$ is the discretized particle density subject to the random variable ν , the uncertain matrix differential equation (MDE) reads

$$\dot{\Psi}(t; \nu) = \mathbf{F}(t, \Psi(t; \nu)), \quad \Psi(t_0; \nu) = \Psi_0(\nu). \quad (5)$$

For simplicity, we choose the uncertain initial conditions that, for example, arise from uncertainty in the position or strength of particle sources in (1) [32]. However, the methods developed in this section can be analogously applied to other types of uncertainties. We assume that, for a fixed ν , (5) has a low-rank solution which is a reasonable assumption, for instance for the radiative transfer equation [13, 31, 32].

We are interested in a lower-dimensional function of the particle density, such as the scalar flux, at a given time $t_k > 0$. We therefore define the map $\nu \rightarrow \mathcal{G}(\Psi(t = t_k; \nu)) \in \mathbb{R}^{\tilde{m} \times \tilde{n}}$, where $\tilde{m}, \tilde{n} \in \mathbb{N}$. Then the quantity of interest (QoI), $Q \in \mathbb{R}^{\tilde{m} \times \tilde{n}}$, is the expectation of $\mathcal{G}(\Psi(t_k; \nu))$. The expectation and variance of $\mathcal{G}(\Psi(t_k; \nu))$ are, respectively,

$$Q := \mathbb{E}_p[\mathcal{G}(\Psi(t_k; \nu))] = \int_{-\infty}^{\infty} \mathcal{G}(\Psi(t_k; \nu))p(\nu) \, d\nu \quad \text{and}$$

$$\text{Var}(\mathcal{G}(\Psi(t_k; \nu))) = \int_{-\infty}^{\infty} \|\mathcal{G}(\Psi(t_k; \nu)) - \mathbb{E}_p[\mathcal{G}(\Psi(t_k; \nu))]\|^2 p(\nu) \, d\nu,$$

where $\|\cdot\|$ is a matrix norm induced by a suitable inner product $\langle \cdot, \cdot \rangle$, e.g., the Frobenius norm induced by the Frobenius inner product.

In Section 3.1 we present a low-rank Monte Carlo estimator to Q and provide error bounds on the low-rank approximation under Monte Carlo sampling. In Section 3.2, we introduce a control variate strategy for variance reduction.

3.1 Low-rank Monte Carlo estimator

To estimate Q , we need to know the exact solution, $\Psi(t; \nu)$, to (5) at $t = t_k$. When the exact solution is unknown it is approximated with a standard time-stepping scheme like the explicit Euler or a higher-order Runge-Kutta method. Let $\tilde{\Psi}_k(\nu) := \tilde{\Psi}(t_k; \nu)$ denote the approximation to $\Psi_k(\nu) := \Psi(t = t_k; \nu)$ computed with a time-stepping scheme. Since $\tilde{\Psi}_k(\nu) \approx \Psi_k(\nu)$, we approximate Q by $\mathbb{E}_p[\mathcal{G}(\tilde{\Psi}_k(\nu))] =: Q_{\text{full}}$ and use the Monte Carlo method to construct the estimator, $\hat{Q}_{\text{full};\text{MC}}$. However, this estimator is computationally expensive and memory-intensive since it requires solving and storing the solution of a high-dimensional MDE for each realization of the random variable.

To overcome the computational inefficiency of $\hat{Q}_{\text{full};\text{MC}}$, we make use of the assumption that (5) has a low-rank solution. As shown in Section 2, when the underlying solution to a time-dependent problem has a low-rank structure, DLRA [26] can be used to approximate the solution. Let $\mathbf{Y}_k(\nu) = \mathbf{X}_k(\nu)\mathbf{S}_k(\nu)\mathbf{V}_k(\nu)^\top$ denote the rank- r approximation to $\Psi_k(\nu)$ computed using the augmented BUG integrator described in Section 2. Then instead of Q we compute $Q \approx Q_r := \mathbb{E}_p[\mathcal{G}(\mathbf{Y}_k(\nu))]$ which is computationally efficient and has lower memory requirements.

Let $\nu_1, \dots, \nu_N \sim p(\nu)$ be N independent and identically distributed random realizations of ν . Then we define the low-rank Monte Carlo estimator as

$$\hat{Q}_{r;\text{MC}} := \frac{1}{N} \sum_{i=1}^N \mathcal{G}(\mathbf{Y}_k(\nu_i)).$$

The mean squared error (MSE) of this estimator is

$$\begin{aligned} \text{MSE}(\hat{Q}_{r;\text{MC}}) &= \mathbb{E}_p \left[\left\| \frac{1}{N} \sum_{i=1}^N \mathcal{G}(\mathbf{Y}_k(\nu_i)) - \mathbb{E}_p[\mathcal{G}(\mathbf{Y}_k(\nu))] \right\|^2 \right] \\ &= \mathbb{E}_p \left[\left\| \hat{Q}_{r;\text{MC}} - \mathbb{E}_p[\hat{Q}_{r;\text{MC}}] \right\|^2 \right] + \left\| \mathbb{E}_p[\hat{Q}_{r;\text{MC}}] - Q \right\|^2. \end{aligned} \quad (6)$$

Thus, the MSE of the estimator separates into the variance of the Monte Carlo estimator and a bias term, describing errors caused by the model, numerical discretization, and low-rank approximation. The variance can be written as

$$\mathbb{E}_p \left[\left\| \hat{Q}_{r;\text{MC}} - \mathbb{E}_p[\hat{Q}_{r;\text{MC}}] \right\|^2 \right] = \text{Var}[\hat{Q}_{r;\text{MC}}] = \frac{1}{N} \text{Var}[\mathcal{G}(\mathbf{Y}_k(\nu))].$$

Similarly, the bias can be written as

$$\left\| \mathbb{E}_p[\hat{Q}_{r;\text{MC}}] - Q \right\| = \|Q_r - Q\| = \left\| \mathbb{E}_p[\mathcal{G}(\mathbf{Y}_k(\nu)) - \mathcal{G}(\Psi_k(\nu))] \right\|,$$

where we use $\mathbb{E}_p[\widehat{Q}_{r;\text{MC}}] = Q_r$. The bias further splits into contributions from the low-rank approximation and those from the underlying model and its discretization. We now provide a robust error bound for the low-rank approximation in the presence of small singular values due to over-approximation of rank, i.e. for the error contribution due to low-rank approximation under uncertainty.

In the deterministic setting, the error bound of [6] shows that the augmented BUG integrator is robust to small singular values. However, this result does not hold in the presence of uncertainty. The next theorem shows that the augmented BUG integrator applied to the uncertain matrix-valued differential equation (2) is robust to over-approximation and how the bounding constants depend on the random variable. The proof closely follows the local error bound for the mid-point BUG integrator [5, Theorem 1] and uses the following assumptions, where $\|\cdot\| = \|\cdot\|_F$ is the Frobenius norm:

- A.1 \mathbf{F} , the right-hand side of (2), is Lipschitz-continuous, bounded, and independent of ν : for all $\mathbf{Y}, \widetilde{\mathbf{Y}} \in \mathbb{R}^{m \times n}$ and $t_0 \leq t \leq t_k$,

$$\|\mathbf{F}(t, \mathbf{Y}) - \mathbf{F}(t, \widetilde{\mathbf{Y}})\| \leq L\|\mathbf{Y} - \widetilde{\mathbf{Y}}\|, \quad \|\mathbf{F}(t, \mathbf{Y})\| \leq B.$$

- A.2 For $\mathbf{F}(t, \mathbf{Y}(t; \nu)) = \mathbf{M}(t, \mathbf{Y}(t; \nu)) + \mathbf{R}(t, \mathbf{Y}(t; \nu))$, where $\mathbf{M}(t, \mathbf{Y}(t; \nu)) := \mathbf{P}(\mathbf{Y}(t; \nu))\mathbf{F}(t, \mathbf{Y}(t; \nu))$. The non-tangential part $\mathbf{R}(t, \mathbf{Y}(t; \nu))$ is ε -small:

$$\|\mathbf{R}(t, \mathbf{Y}(t; \nu))\| \leq c_1(\nu)\varepsilon, \quad c_1 \in L^1.$$

- A.3 The error in the initial value is δ -small:

$$\|\mathbf{Y}_0(\nu) - \Psi_0(\nu)\| \leq c_0(\nu)\delta, \quad c_0 \in L^1.$$

Theorem 1 (Local error bound). *Let $\Psi(t; \nu)$ denote the solution of the uncertain matrix differential equation (2) and $\mathbf{Y}_1(\nu)$ denote the rank- r approximation to $\Psi_1(\nu) := \Psi(t_1; \nu)$ at $t_1 = t_0 + h$ obtained after one step of the augmented BUG integrator with step-size $h > 0$ and truncation error $\vartheta_1(\nu)$. Then, if the assumptions (A.1-A.2) are satisfied, we have the following local error bound on the expected value of the solution*

$$\|\mathbb{E}_p[\mathbf{Y}_1(\nu) - \Psi_1(\nu)]\| \leq 2\mathbb{E}_p[c_1(\nu)]\varepsilon h + 5LBh^2 + \mathbb{E}_p[\vartheta_1(\nu)].$$

Remark 1. We fix some notation used in the proof. Let $\mathbf{Y}_0^\nu = \mathbf{X}_0^\nu \mathbf{S}_0^\nu \mathbf{V}_0^{\nu, \top}$ be a realization of $\mathbf{Y}_0(\nu)$, i.e. $\mathbf{Y}_0^\nu \sim \mathbf{Y}_0(\nu)$. Then the exact solution at time $t_1 = t_0 + h$, with the initial condition \mathbf{Y}_0^ν , is denoted by $\Psi_1^\nu \sim \Psi_1(\nu)$. Similarly, $\mathbf{Y}_1^\nu = \mathbf{X}_1^\nu \mathbf{S}_1^\nu \mathbf{V}_1^{\nu, \top} \sim \mathbf{Y}_1(\nu)$ denotes the low-rank approximation at t_1 computed with the augmented BUG integrator and $\vartheta_1^\nu \sim \vartheta_1(\nu)$ denotes the truncation error made at t_1 . Additionally, let $\widehat{\mathbf{Y}}_1^\nu = \widehat{\mathbf{X}}_1^\nu \widehat{\mathbf{S}}_1^\nu \widehat{\mathbf{V}}_1^{\nu, \top} \sim \widehat{\mathbf{Y}}_1^\nu$ denote the augmented solution before truncation.

Proof. We start by using Jensen's and Cauchy-Schwarz (CS) inequality to get,

$$\|\mathbb{E}_p[\mathbf{Y}_1(\nu) - \Psi_1(\nu)]\| \leq \mathbb{E}_p \left[\|\widehat{\mathbf{Y}}_1(\nu) - \Psi_1(\nu)\| \right] + \mathbb{E}_p \left[\|\widehat{\mathbf{Y}}_1(\nu) - \mathbf{Y}_1(\nu)\| \right].$$

For each realization we have

$$\|\widehat{\mathbf{Y}}_1^\nu - \Psi_1^\nu\| \stackrel{\text{CS}}{\leq} \|\Psi_1^\nu - \widehat{\mathbf{X}}_1^\nu \widehat{\mathbf{X}}_1^{\nu, \top} \Psi_1^\nu \widehat{\mathbf{V}}_1^\nu \widehat{\mathbf{V}}_1^{\nu, \top}\| + \|\widehat{\mathbf{X}}_1^\nu \widehat{\mathbf{X}}_1^{\nu, \top} \Psi_1^\nu \widehat{\mathbf{V}}_1^\nu \widehat{\mathbf{V}}_1^{\nu, \top} - \widehat{\mathbf{Y}}_1^\nu\|.$$

Adding and subtracting $\widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top} \boldsymbol{\Psi}_1^\nu$ from the first term on the right-hand side yields

$$\|\boldsymbol{\Psi}_1^\nu - \widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top} \boldsymbol{\Psi}_1^\nu \widehat{\mathbf{V}}^\nu \widehat{\mathbf{V}}^{\nu,\top}\| \leq \|\boldsymbol{\Psi}_1^\nu - \widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top} \boldsymbol{\Psi}_1^\nu\| + \|\boldsymbol{\Psi}_1^\nu - \boldsymbol{\Psi}_1^\nu \widehat{\mathbf{V}}^\nu \widehat{\mathbf{V}}^{\nu,\top}\|. \quad (7)$$

From the definition of $\widehat{\mathbf{X}}^\nu$, $(\mathbf{I} - \widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top}) \mathbf{K}^\nu(t_1) \mathbf{V}_0^{\nu,\top} = 0$, we write the first term of the right-hand side of (7) as

$$\begin{aligned} \|\boldsymbol{\Psi}_1^\nu - \widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top} \boldsymbol{\Psi}_1^\nu\| &= \|(\mathbf{I} - \widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top})(\boldsymbol{\Psi}_1^\nu - \mathbf{K}^\nu(t_1) \mathbf{V}_0^{\nu,\top})\| \\ &\stackrel{\text{FTC}}{\leq} \int_{t_0}^{t_1} \|(\mathbf{I} - \widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top})(\mathbf{F}(t, \boldsymbol{\Psi}^\nu(t)) - \mathbf{F}(t, \mathbf{K}^\nu(t) \mathbf{V}_0^{\nu,\top})) \mathbf{V}_0^{\nu,\top}\| dt, \end{aligned}$$

where FTC stands for the fundamental theorem of calculus. We can write

$$\begin{aligned} \mathbf{F}(t, \boldsymbol{\Psi}^\nu(t)) &= \mathbf{F}(t_0, \mathbf{Y}_0^\nu) + \mathbf{F}(t, \boldsymbol{\Psi}^\nu(t)) - \mathbf{F}(t_0, \mathbf{Y}_0^\nu) \\ &= \mathbf{M}(t_0, \mathbf{Y}_0^\nu) + \mathbf{F}(t, \boldsymbol{\Psi}^\nu(t)) - \mathbf{F}(t_0, \mathbf{Y}_0^\nu) + \mathbf{R}(t_0, \mathbf{Y}_0^\nu), \end{aligned}$$

$$\begin{aligned} \text{and } \mathbf{F}(t, \mathbf{K}^\nu(t) \mathbf{V}_0^{\nu,\top}) \mathbf{V}_0^{\nu,\top} &= \mathbf{F}(t_0, \mathbf{Y}_0^\nu) \mathbf{V}_0^{\nu,\top} \\ &+ \left(\mathbf{F}(t, \mathbf{K}^\nu(t) \mathbf{V}_0^{\nu,\top}) \mathbf{V}_0^{\nu,\top} - \mathbf{F}(t_0, \mathbf{Y}_0^\nu) \mathbf{V}_0^{\nu,\top} \right). \end{aligned}$$

Since $\|\mathbf{I} - \widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top}\| \leq 1$, we have

$$\begin{aligned} \|\boldsymbol{\Psi}_1^\nu - \widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top} \boldsymbol{\Psi}_1^\nu\| &\stackrel{\text{CS}}{\leq} \int_{t_0}^{t_1} \|\mathbf{F}(t, \mathbf{K}^\nu(t) \mathbf{V}_0^{\nu,\top}) - \mathbf{F}(t_0, \mathbf{Y}_0^\nu)\| dt \\ &+ \int_{t_0}^{t_1} \|\mathbf{F}(t, \boldsymbol{\Psi}^\nu(t)) - \mathbf{F}(t_0, \mathbf{Y}_0^\nu)\| dt + \int_{t_0}^{t_1} \|\mathbf{R}(t_0, \mathbf{Y}_0^\nu)\| dt \\ &\stackrel{\text{A.1,A.2}}{\leq} L \int_{t_0}^{t_1} \|\mathbf{K}^\nu(t) \mathbf{V}_0^{\nu,\top} - \mathbf{Y}_0^\nu\| dt + L \int_{t_0}^{t_1} \|\boldsymbol{\Psi}^\nu(t) - \mathbf{Y}_0^\nu\| dt + c_1^\nu h \varepsilon. \end{aligned}$$

Thus using FTC and A.1 we get,

$$\|\boldsymbol{\Psi}_1^\nu - \widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top} \boldsymbol{\Psi}_1^\nu\| \leq 2LBh^2 + c_1^\nu h \varepsilon.$$

Using $(\mathbf{I} - \widehat{\mathbf{V}}^\nu \widehat{\mathbf{V}}^{\nu,\top}) \mathbf{L}^\nu(t_1) \mathbf{X}_0^{\nu,\top} = 0$ we get an analogous expression for $\|\boldsymbol{\Psi}_1^\nu - \boldsymbol{\Psi}_1^\nu \widehat{\mathbf{V}}^\nu \widehat{\mathbf{V}}^{\nu,\top}\|$. Thus, (7) becomes

$$\|\boldsymbol{\Psi}_1^\nu - \widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top} \boldsymbol{\Psi}_1^\nu \widehat{\mathbf{V}}^\nu \widehat{\mathbf{V}}^{\nu,\top}\| \leq 4LBh^2 + 2c_1^\nu h \varepsilon.$$

Now since $\widehat{\mathbf{Y}}^\nu(t) = \widehat{\mathbf{X}}^\nu \widehat{\mathbf{S}}^\nu(t) \widehat{\mathbf{V}}^\nu$ we have $\|\widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top} \boldsymbol{\Psi}_1^\nu \widehat{\mathbf{V}}^\nu \widehat{\mathbf{V}}^{\nu,\top} - \widehat{\mathbf{Y}}_1^\nu\|$

$$\begin{aligned} &\stackrel{\text{FTC}}{\leq} \int_{t_0}^{t_1} \|\mathbf{F}(t, \boldsymbol{\Psi}^\nu(t)) - \mathbf{F}(t, \widehat{\mathbf{Y}}^\nu(t))\| dt \\ &\stackrel{\text{A.1}}{\leq} L \int_{t_0}^{t_1} \|\boldsymbol{\Psi}^\nu(t) - \widehat{\mathbf{Y}}^\nu(t)\| dt \\ &\stackrel{\text{FTC}}{\leq} L \int_{t_0}^{t_1} \int_{t_0}^t \|\mathbf{F}(s, \boldsymbol{\Psi}^\nu(s)) - \widehat{\mathbf{X}}^\nu \widehat{\mathbf{X}}^{\nu,\top} \mathbf{F}(s, \widehat{\mathbf{Y}}^\nu(s)) \widehat{\mathbf{V}}^\nu \widehat{\mathbf{V}}^{\nu,\top}\| ds dt \\ &\stackrel{\text{CS,A.1}}{\leq} LBh^2. \end{aligned}$$

Putting it all together yields

$$\|\widehat{\mathbf{Y}}_1^\nu - \mathbf{\Psi}_1^\nu\| \leq 5LBh^2 + \varepsilon h.$$

We denote the error made by truncating to rank- r by $\vartheta_1(\nu)$, then

$$\|\widehat{\mathbf{Y}}_1^\nu - \mathbf{Y}_1^\nu\| \leq \vartheta_1^\nu$$

Thus we get the following local error bound

$$\|\mathbb{E}_p[\mathbf{Y}_1(\nu) - \mathbf{\Psi}_1(\nu)]\| \leq 2\mathbb{E}_p[c_1(\nu)]\varepsilon h + 5LBh^2 + \mathbb{E}_p[\vartheta_1(\nu)].$$

Theorem 2 (Global error bound). *Let $\mathbf{Y}_k(\nu)$ denote the rank- r approximation to $\mathbf{\Psi}_k(\nu)$ at $t_k = t_0 + kh$ obtained after k steps of the augmented basis-update & Galerkin integrator with step-size $h > 0$ and $k = \frac{k'}{h}$, with k' a fixed integer. Then, if the assumptions (A.1-A.3) are satisfied and $\vartheta_i(\nu)$ denotes the truncation error at t_i , the error satisfies for all k with $t_k = t_0 + kh \leq t_{\text{end}}$*

$$\|\mathbb{E}_p[\mathbf{Y}_k(\nu) - \mathbf{\Psi}_k(\nu)]\| \leq \mathbb{E}_p[c_0(\nu)]\delta + \mathbb{E}_p[\widetilde{c}_1(\nu)]\varepsilon + c_2h + \sum_{i=1}^k \mathbb{E}_p[\vartheta_i(\nu)],$$

where the constants $\widetilde{c}_1(\nu)$ and c_2 only depend on L , B , and t_{end} . In particular, the constants are independent of singular values of the exact or approximate solution.

Thus, the error contribution to the bias due to low-rank approximation remains robust to small singular values in the probabilistic setting, making it a suitable and cheaper alternative to full-rank numerical solvers for uncertainty quantification. In the Section 3.2, we discuss using reduced rank DLRA models as a control variate for the estimator $\widehat{Q}_{r;\text{MC}}$.

3.2 Control variates

The DLRA-based integrator reduces the computational cost of each Monte Carlo sample by using an approximate solution that is ε -close to the low-rank manifold independent of the random variable ν . We thus expect to be able to compute more samples, at the same cost, compared to a full-rank integrator. Given that the variance of the MC estimator scales as $O(N^{-1})$, we expect to produce a lower variance estimate $\widehat{Q}_{r;\text{MC}}$. However, a DLRA-based integrator capturing the true rank of the PDE solution may still be quite expensive in practice. This is further compounded by the fact that one often over-approximates the rank of the solution, out of caution. Hence, for high-dimensional PDEs, like (1), it is infeasible to arbitrarily increase the number of MC samples. We therefore introduce a control-variate strategy to further suppress the variance of $\widehat{Q}_{r;\text{MC}}$.

Simulations based on a reduced rank ($s < r$) can produce good, lower-cost approximations to more accurate higher-rank simulations. Hence, we can use the reduced-rank simulations as a control variate, with the goal of reducing the variance of the estimator $\widehat{Q}_{r;\text{MC}}$. Let $\mathbf{Y}_k^r(\nu)$ and $\mathbf{Y}_k^s(\nu)$ be the low-rank approximations with ranks r and s , respectively. We assume that $\mathbb{E}_p[\mathcal{G}(\mathbf{Y}_k^s(\nu))]$ is known. Then for some $\alpha > 0$, we define

$$\mathcal{G}_{\text{CV}}(\nu) := \mathcal{G}(\mathbf{Y}_k^r(\nu)) - \alpha (\mathcal{G}(\mathbf{Y}_k^s(\nu)) - \mathbb{E}_p[\mathcal{G}(\mathbf{Y}_k^s(\nu))]).$$

With re-arrangement we get the control variate estimator

$$\widehat{Q}_{r;\text{CV}} = \alpha \mathbb{E}_p [\mathcal{G}(\mathbf{Y}_k^s(\nu))] + \frac{1}{N} \sum_{i=1}^N (\mathcal{G}(\mathbf{Y}_k^r(\nu_i)) - \alpha \mathcal{G}(\mathbf{Y}_k^s(\nu_i))). \quad (8)$$

If $\text{Var}_j := \text{Var}(\mathcal{G}(\mathbf{Y}_k^j(\nu)))$, $j \in \{r, s\}$, then the variance of $\mathcal{G}_{\text{CV}}(\nu)$ is

$$\text{Var}(\mathcal{G}_{\text{CV}}(\nu)) = \text{Var}_r + \alpha^2 \text{Var}_s - 2\alpha \text{Cov}_{rs},$$

$$\text{where } \text{Cov}_{rs} = \mathbb{E}_p [\langle \mathcal{G}(\mathbf{Y}_k^r(\nu)) - Q_r, \mathcal{G}(\mathbf{Y}_k^s(\nu)) - Q_s \rangle].$$

The optimal α , that minimizes the variance of the estimator is given by $\alpha^* = \text{Cov}_{rs}/\text{Var}_s$.

Remark 2. In practice, the optimal value of α as well as $\mathbb{E}_p [\mathcal{G}(\mathbf{Y}_k^s(\nu))]$ have to be estimated (on-the-fly or beforehand). As $s < r$, the estimate is cheaper than the equivalent computation with rank- r . Furthermore, the covariance of $\mathcal{G}(\mathbf{Y}_k^r(\nu))$ and $\mathcal{G}(\mathbf{Y}_k^s(\nu))$ increases as s approaches r .

4 Numerical experiments

To demonstrate the efficacy of the proposed low-rank estimators we consider the uncertain 1D radiation transport equation uncertain initial conditions. The source code can be found at github.com/chinsp/publication-Low-rank-for-uncertain-radiative-transfer. Specifically, we consider a projection of (1) onto \mathbb{R} so that $x \in [a, b]$, $\mu \in [-1, 1]$, $t > 0$. We further introduce an uncertain parameter ν with probability density function $p(\nu)$, giving the model equation

$$\begin{aligned} \partial_t \psi(t, x, \mu, \nu) + \mu \partial_x \psi(t, x, \mu, \nu) &= \sigma_s(x) \left(\frac{1}{2} \int_{-1}^1 \psi(t, x, \mu', \nu) d\mu' - \psi(t, x, \mu, \nu) \right), \\ \psi(t_0, x, \mu, \nu) &= \psi_0(x, \mu, \nu). \end{aligned}$$

Note that we assume that the particle density never reaches the boundary. In the numerical experiments the initial condition is set as a cut-off Gaussian with an uncertain amplitude:

$$\psi(0, x, \mu, \nu) = \max \left\{ 10^{-4}, \frac{\nu}{\sqrt{2\pi\sigma}} \exp \left(-\frac{x^2}{2\sigma^2} \right) \right\}, \quad x \in [-1.5, 1.5], \nu \sim \text{U}(0.5, 1.5).$$

The equation is discretized using m grid points in space with an upwind-downwind flux and a spherical harmonics method of order $n-1$ (P_{n-1}) [4] in the direction of travel. The discretization parameters are specified case-wise. This yields a matrix-valued differential equation with an uncertain initial condition.

In this work, we are interested in the expected value of the scalar flux,

$$\phi(x, \nu) = \frac{1}{\sqrt{2}} \int_{-1}^1 \psi(t = 1.0, x, \mu, \nu) d\mu,$$

i.e., $Q = \mathbb{E}_p [\phi(x, \nu)]$. A fine-grid Monte Carlo reference solution is computed for P_{101} approximation with $m = 1601$ spatial grid points and $N = 1.024 \times 10^5$ samples. The CFL number, representing the number of grid cells traveled per time step, is set to 1 in all experiments.

4.1 Low-rank Monte Carlo

While applying the low-rank Monte Carlo estimator is straightforward, it is unclear how the estimator’s parameters affect the MC error and the bias of the estimate. Thus, we conduct a parametric study of the low-rank Monte Carlo method by varying the spatial grids, rank, and the Monte Carlo samples. In preliminary investigations, it was seen that the discretization of μ does not affect the Monte Carlo error or the bias and thus we fix the discretization to the same one as the reference. For the spatial discretization we use nested grids with $m \in \{101, 201, 401\}$, ranks $r \in \{2, 5, 10, 15, 20, 25, 30, 35, 40\}$, and $N \in \{400, 1600, 6400, 25600\}$, where N is the number of samples. The results are plotted in Figure 1.

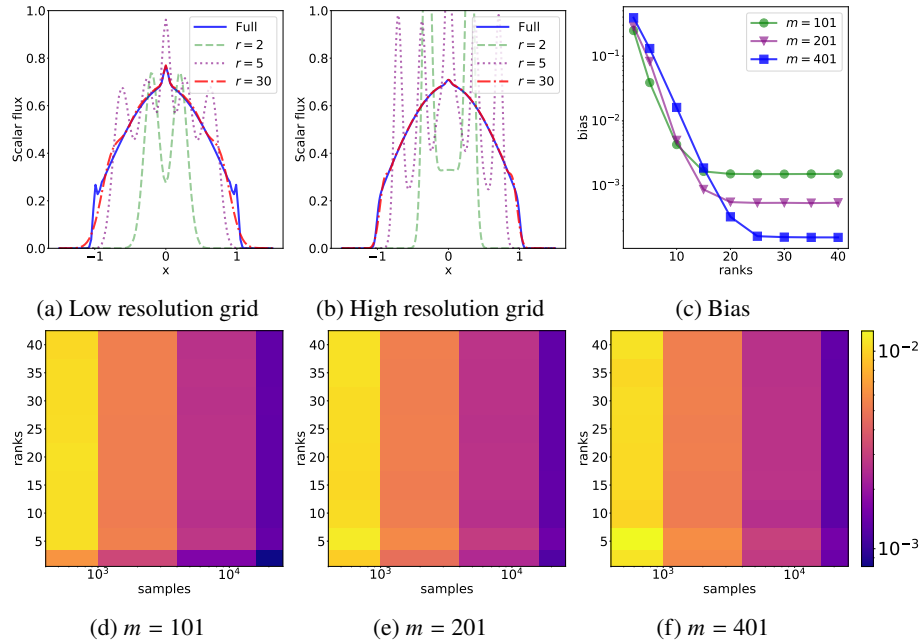


Fig. 1: Top row: Scalar flux computed using the full P_{101} solver and the augment BUG solver at time $t = 1$ for a fixed sample ν and ranks $r \in \{2, 5, 30\}$ with (a) $m = 101$ (b) $m = 401$ spatial grid points. (c) The bias of the low-rank Monte Carlo estimator for $r = 30$ and $N = 25600$ samples. Bottom row: Monte Carlo error of the low-rank Monte Carlo estimator in log-scale for spatial grids of varying widths.

We see from Figure 1a-1c that, independent of the underlying spatial grid, the bias decreases with increasing rank until the spatial discretization error dominates it. Additionally, we see that for smaller ranks ($r \leq 10$), the bias is smaller when the underlying spatial grid is coarse. At the same time, the behavior is flipped for larger ranks ($r \geq 20$), i.e., the error is smaller when the underlying spatial grid is finer. From

Figure 1d-1f we see that as the sample size increases, the Monte Carlo error decreases independent of the rank or the underlying spatial grid. Hence, a good estimate of the QoI requires knowledge of the interplay of rank and the underlying spatial discretization. Here, especially the required rank is difficult to predict in practice, often resulting in over-approximation. Further, even though each individual sample is cheaper to compute using a low-rank approach, reducing the error of the low-rank Monte Carlo estimator can still require a large number of samples. Especially for higher-dimensional problems like (1) it is therefore worthwhile to consider further approaches to variance reduction within our low-rank framework.

4.2 Control variates

To gauge the potential benefits of using lower-rank estimates for variance reduction, we next compare the low-rank Monte Carlo approach to control variate estimators with different levels of coarseness with respect to rank. For this, we perform low-rank Monte Carlo samples as described in Section 4.1, with a fixed discretization of x and μ ($m = 201, n = 101$) but varying rank $r \in \{20, 25, 30, 35, 40\}$. The number of samples is also fixed to $N_{MC} = 2000$. To improve the variance of the low-rank Monte Carlo estimator we use a control variate with a coarser rank $s \in \{2, 5, 10, 15\}$. Since the ultimate goal of variance reduction is typically to achieve a lower error at the same computational costs or vice versa, we fix the error and investigate the effect on the runtime. For this, we use the following heuristic: We compute the same number of samples for the coarse rank s as for the low-rank Monte Carlo estimator, i.e., $N_{CV_{coarse}} = N_{MC} = 2000$. Then we compute the optimal number of samples for the differences such that the total error of the control variate estimator is at most as high as for Monte Carlo

$$N_{CV_{diff}} = \text{Var}_r \cdot \frac{1 - \text{Corr}_{r,s}^2}{\epsilon^2}, \quad \text{Corr}_{r,s} = \frac{\text{Cov}_{r,s}}{\sqrt{\text{Var}_r} \sqrt{\text{Var}_s}},$$

where ϵ is the Monte Carlo error (see 4.1), Var_r is the variance of the rank- r solver, $\text{Corr}_{r,s}$ is correlation and $\text{Cov}_{r,s}$ the covariance between the finer rank- r and coarser rank- s solver. Since the variances and the correlation between coarse and fine levels are not known a priori, these have to be estimated as well. Here, we compare two different approaches: First, we run an extra simulation beforehand to estimate the optimal parameter α as well as the number of samples for the finer level. We then use this as an input in the control variates estimator. Second, we estimate the parameters within the control variate computation, using 200 warm-up samples. Note that this often leads to a higher number of samples than necessary if the computed number is below 200. In Figure 2 the minimal runtime over 20 runs is plotted for both the Monte Carlo estimator at different ranks and the control variate estimators with all combinations of coarse and fine rank. It is apparent that the control variate estimators are consistently faster than low-rank Monte Carlo, both when computing the optimal number of samples beforehand (Figure 2a) and when estimating the parameters during computation using 200 warm-up samples (Figure 2b). While the runtime seems to increase exponentially with the growing rank for Monte Carlo, the runtimes for control variates grow much slower - closer to linearly especially for the higher choices of coarse rank. This implies that it might be worthwhile to consider a multi-level or at least multiple level framework.




Fig. 2: Log-log plot of (minimal) runtime of Monte Carlo vs. different control variates at theoretically constant error. (a) Number of samples and α were estimated based on a separate simulation with 500 samples. (b) Number of samples and α were estimated during runtime based on 200 warm-up samples.

Lastly, in Table 1, the computed values for the parameter α are shown. While the values are almost constant with respect to the rank of the finer level r , they vary between around 0.6 and close to 1 depending on the rank of the coarse level. As expected, the correlation between fine and coarse ranks is higher for larger ranks on the coarse level. Thus, the α values are lower for smaller ranks, especially for $s = 2$ and $s = 4$ and approach 1 for the higher ranks.

$s \setminus r$	20	25	30	35	40
2	0.6060	0.6060	0.6060	0.6060	0.6060
5	0.7834	0.7833	0.7833	0.7833	0.7833
10	0.9937	0.9940	0.9940	0.9940	0.9940
15	0.9979	0.9980	0.9980	0.9980	0.9980

Table 1: Precomputed optimal α values based on a simulation with 500 samples.

5 Conclusion

We have studied the use of the dynamical low-rank approximation in the context of uncertainty quantification, both for cost reduction in a standard Monte Carlo framework and as a lower fidelity control variate. We showed that the robust error bound holds in a probabilistic setting and from the low-rank Monte Carlo experiments we see that, while the variance is only influenced by the sample size, an optimal balance of rank and grid size is necessary to correctly capture the quantity of interest. Computational costs can be further reduced significantly compared to the low-rank Monte Carlo method

using a control variate approach based on a lower rank approximation. Based on our results, future work on the design of a fully multi-level approach balancing both rank and spatial discretization error seems promising for tackling larger scale problems involving the radiative transfer equation, as well as other problems with an identified low-rank structure.

Acknowledgements

The work of Chinmay Patwardhan was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 258734477 – SFB 1173. Pia Stammer received funding from the German National Academy of Sciences Leopoldina for the project underlying this article, under grant number LPDS 2024-03.

References

1. Barth, A., Lang, A., Schwab, C.: Multilevel Monte Carlo method for parabolic stochastic partial differential equations. *BIT Numerical Mathematics* **53**(1), 3–27 (2013). DOI 10.1007/s10543-012-0401-5
2. Baumann, L., Einkemmer, L., Klingenberg, C., Kusch, J.: Energy Stable and Conservative Dynamical Low-Rank Approximation for the Su–Olson Problem. *SIAM Journal on Scientific Computing* **46**(2), B137–B158 (2024). DOI 10.1137/23M1586215
3. Bird, G.A.: Monte Carlo Simulation of Gas Flows. *Annual Review of Fluid Mechanics* **10**, 11–31 (1978). DOI 10.1146/annurev.fl.10.010178.000303
4. Case, K.M., Zweifel, P.F.: *Linear Transport Theory*. Addison-Wesley Publishing Company (1967)
5. Ceruti, G., Einkemmer, L., Kusch, J., Lubich, C.: A robust second-order low-rank BUG integrator based on the midpoint rule. *BIT Numerical Mathematics* **64**(3), 30 (2024). DOI 10.1007/s10543-024-01032-x
6. Ceruti, G., Kusch, J., Lubich, C.: A rank-adaptive robust integrator for dynamical low-rank approximation. *BIT Numerical Mathematics* **62**(4), 1149–1174 (2022). DOI 10.1007/s10543-021-00907-7
7. Ceruti, G., Kusch, J., Lubich, C.: A Parallel Rank-Adaptive Integrator for Dynamical Low-Rank Approximation. *SIAM Journal on Scientific Computing* **46**(3), B205–B228 (2024). DOI 10.1137/23M1565103
8. Ceruti, G., Lubich, C.: An unconventional robust integrator for dynamical low-rank approximation. *BIT Numerical Mathematics* **62**(1), 23–44 (2022). DOI 10.1007/s10543-021-00873-0
9. Cliffe, K.A., Giles, M.B., Scheichl, R., Teckentrup, A.L.: Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Computing and Visualization in Science* **14**(1), 3–15 (2011). DOI 10.1007/s00791-011-0160-x
10. Coughlin, J., Hu, J.: Efficient dynamical low-rank approximation for the Vlasov-Ampère-Fokker-Planck system. *Journal of Computational Physics* **470**, 111590 (2022). DOI 10.1016/j.jcp.2022.111590
11. Dimarco, G., Pareschi, L.: Numerical methods for kinetic equations. *Acta Numerica* **23**, 369–520 (2014). DOI 10.1017/S0962492914000063
12. Ducloux, R., Dubroca, B., Frank, M.: A deterministic partial differential equation model for dose calculation in electron radiotherapy. *Physics in Medicine & Biology* **55**(13), 3843 (2010). DOI 10.1088/0031-9155/55/13/018

13. Einkemmer, L., Lubich, C.: A Low-Rank Projector-Splitting Integrator for the Vlasov–Poisson Equation. *SIAM Journal on Scientific Computing* **40**(5), B1330–B1360 (2018). DOI 10.1137/18M116383X
14. Einkemmer, L., Moriggl, A.: Semi-Lagrangian 4d, 5d, and 6d kinetic plasma simulation on large-scale GPU-equipped supercomputers. *The International Journal of High Performance Computing Applications* **37**(2), 180–196 (2023). DOI 10.1177/10943420221137599
15. Giles, M.B.: Multilevel Monte Carlo Path Simulation. *Operations Research* **56**(3), 607–617 (2008). DOI 10.1287/opre.1070.0496
16. Giles, M.B.: Multilevel Monte Carlo methods. *Acta Numerica* **24**, 259–328 (2015). DOI 10.1017/S096249291500001X
17. Grad, H.: On the kinetic theory of rarefied gases. *Communications on Pure and Applied Mathematics* **2**(4), 331–407 (1949). DOI 10.1002/cpa.3160020403
18. Haji-Ali, A.L., Nobile, F., von Schwerin, E., Tempone, R.: Optimization of mesh hierarchies in multilevel Monte Carlo samplers. *Stochastics and Partial Differential Equations Analysis and Computations* **4**(1), 76–112 (2016). DOI 10.1007/s40072-015-0049-7
19. Haji-Ali, A.L., Nobile, F., Tempone, R.: Multi-index Monte Carlo: when sparsity meets sampling. *Numerische Mathematik* **132**(4), 767–806 (2016). DOI 10.1007/s00211-015-0734-5
20. Haji-Ali, A.L., Tempone, R.: Multilevel and Multi-index Monte Carlo methods for the McKean–Vlasov equation. *Statistics and Computing* **28**(4), 923–935 (2018). DOI 10.1007/s11222-017-9771-5
21. Heinrich, S.: Multilevel Monte Carlo Methods. In: S. Margenov, J. Waśniewski, P. Yalamov (eds.) *Large-Scale Scientific Computing*, vol. 2179, pp. 58–67. Springer Berlin Heidelberg, Berlin, Germany (2001). DOI 10.1007/3-540-45346-6_5. Series Title: Lecture Notes in Computer Science
22. Howes, G., Dorland, W., Cowley, S., Hammett, G., Quataert, E., Schekochihin, A., Tatsuno, T.: Kinetic simulations of magnetized turbulence in astrophysical plasmas. *Physical Review Letters* **100**(6), 065004 (2008). DOI 10.1103/PhysRevLett.100.065004
23. Hu, J., Jin, S., Li, J., Zhang, L.: On multilevel Monte Carlo methods for deterministic and uncertain hyperbolic systems. *Journal of Computational Physics* **475**, 111847 (2023). DOI 10.1016/j.jcp.2022.111847
24. Kieri, E., Lubich, C., Walach, H.: Discretized Dynamical Low-Rank Approximation in the Presence of Small Singular Values. *SIAM Journal on Numerical Analysis* **54**(2), 1020–1038 (2016). DOI 10.1137/15M1026791
25. Kieri, E., Vandereycken, B.: Projection methods for dynamical low-rank approximation of high-dimensional problems. *Computational Methods in Applied Mathematics* **19**(1), 73–92 (2019). DOI 10.1515/cmam-2018-0029
26. Koch, O., Lubich, C.: Dynamical Low-Rank Approximation. *SIAM Journal on Matrix Analysis and Applications* **29**(2), 434–454 (2007). DOI 10.1137/050639703
27. Kumar, P., Oosterlee, C.W., Dwight, R.P.: A Multigrid Multilevel Monte Carlo Method Using High-Order Finite-Volume Scheme for Lognormal Diffusion Problems. *International Journal for Uncertainty Quantification* **7**(1), 57–81 (2017). DOI 10.1615/Int.J.UncertaintyQuantification.2016018677
28. Kusch, J.: Second-order robust parallel integrators for dynamical low-rank approximation (2024). DOI 10.48550/arXiv.2403.02834
29. Kusch, J., Ceruti, G., Einkemmer, L., Frank, M.: Dynamical Low-Rank Approximation for Burgers’ Equation with Uncertainty. *International Journal for Uncertainty Quantification* **12**(5), 1–21 (2022). DOI 10.1615/Int.J.UncertaintyQuantification.2022039345
30. Kusch, J., Einkemmer, L., Ceruti, G.: On the Stability of Robust Dynamical Low-Rank Approximations for Hyperbolic Problems. *SIAM Journal on Scientific Computing* **45**(1), A1–A24 (2023). DOI 10.1137/21M1446289

31. Kusch, J., Stammer, P.: A robust collision source method for rank adaptive dynamical low-rank approximation in radiation therapy. *ESAIM: Mathematical Modelling and Numerical Analysis* **57**(2), 865–891 (2023). DOI 10.1051/m2an/2022090
32. Kusch, J., Whewell, B., McClarren, R., Frank, M.: A low-rank power iteration scheme for neutron transport criticality problems. *Journal of Computational Physics* **470**, 111587 (2022). DOI 10.1016/j.jcp.2022.111587
33. Lubich, C., Oseledets, I.V.: A projector-splitting integrator for dynamical low-rank approximation. *BIT Numerical Mathematics* **54**(1), 171–188 (2014). DOI 10.1007/s10543-013-0454-0
34. Løvbak, E., Samaey, G.: Accelerated simulation of Boltzmann-BGK equations near the diffusive limit with asymptotic-preserving multilevel Monte Carlo. *SIAM Journal on Scientific Computing* **45**(4), A1862–A1889 (2023). DOI 10.1137/22M1498498
35. Løvbak, E., Samaey, G., Vandewalle, S.: A multilevel Monte Carlo method for asymptotic-preserving particle schemes in the diffusive limit. *Numerische Mathematik* **148**(1), 141–186 (2021). DOI 10.1007/s00211-021-01201-y
36. Mishra, S., Schwab, C.: Sparse tensor multi-level Monte Carlo finite volume methods for hyperbolic conservation laws with random initial data. *Mathematics of Computation* **81**(280), 1979–2018 (2012). DOI 10.1090/S0025-5718-2012-02574-9
37. Mortier, B., Robbe, P., Baelmans, M., Samaey, G.: Multilevel asymptotic-preserving Monte Carlo for kinetic-diffusive particle simulations of the Boltzmann-BGK equation. *Journal of Computational Physics* **450**, 110736 (2022). DOI 10.1016/j.jcp.2021.110736
38. Musharbash, E., Nobile, F., Vidličková, E.: Symplectic dynamical low rank approximation of wave equations with random parameters. *BIT Numerical Mathematics* **60**(4), 1153–1201 (2020). DOI 10.1007/s10543-020-00811-6
39. Patwardhan, C., Frank, M., Kusch, J.: Asymptotic-preserving and energy stable dynamical low-rank approximation for thermal radiative transfer equations (2024). DOI 10.48550/arXiv.2402.16746
40. Robbe, P., Nuyens, D., Vandewalle, S.: Recycling Samples in the Multigrid Multilevel (Quasi-)Monte Carlo Method. *SIAM Journal on Scientific Computing* **41**(5), S37–S60 (2019). DOI 10.1137/18M1194031
41. Rosin, M., Ricketson, L., Dimits, A., Caflisch, R., Cohen, B.: Multilevel Monte Carlo simulation of Coulomb collisions. *Journal of Computational Physics* **274**, 140–157 (2014). DOI 10.1016/j.jcp.2014.05.030
42. Sapsis, T.P., Lermusiaux, P.F.: Dynamically orthogonal field equations for continuous stochastic dynamical systems. *Physica D: Nonlinear Phenomena* **238**(23-24), 2347–2360 (2009). DOI 10.1016/j.physd.2009.09.017
43. Stammer, P.: Uncertainty quantification and numerical methods in charged particle radiation therapy. PhD Thesis, Karlsruhe Institut für Technologie (KIT), Karlsruhe, Germany (2023). DOI 10.5445/IR/1000158316
44. Stammer, P., Burlacu, T., Wahl, N., Lathouwers, D., Kusch, J.: A Deterministic Dynamical Low-rank Approach for Charged Particle Transport (2024). DOI 10.48550/arXiv.2412.09484
45. Szpruch, L., Tan, S., Tse, A.: Iterative multilevel particle approximation for McKean–Vlasov SDEs. *The Annals of Applied Probability* **29**(4) (2019). DOI 10.1214/18-AAP1452
46. Wilson, J.R.: Variance Reduction Techniques for Digital Simulation. *American Journal of Mathematical and Management Sciences* **4**(3-4), 277–312 (1984). DOI 10.1080/01966324.1984.10737147
47. Zweibel, E.G., Yamada, M.: Magnetic reconnection in astrophysical and laboratory plasmas. *Annual review of astronomy and astrophysics* **47**, 291–332 (2009). DOI 10.1146/annurev-astro-082708-101726