

Adaptive symplectic model order reduction of parametric particle-based Vlasov–Poisson equation

Jan S. Hesthaven*, Cecilia Pagliantini[†], and Nicolò Ripamonti[‡]

Published on Math. Comp. 93 (2024), 1153-1202
DOI: <https://doi.org/10.1090/mcom/3885>

Abstract

High-resolution simulations of particle-based kinetic plasma models typically require a high number of particles and thus often become computationally intractable. This is exacerbated in multi-query simulations, where the problem depends on a set of parameters. In this work, we derive reduced order models for the semi-discrete Hamiltonian system resulting from a geometric particle-in-cell approximation of the parametric Vlasov–Poisson equations. Since the problem’s non-dissipative and highly nonlinear nature makes it reducible only locally in time, we adopt a nonlinear reduced basis approach where the reduced phase space evolves in time. This strategy allows a significant reduction in the number of simulated particles, but the evaluation of the nonlinear operators associated with the Vlasov–Poisson coupling remains computationally expensive. We propose a novel reduction of the nonlinear terms that combines adaptive parameter sampling and hyper-reduction techniques to address this. The proposed approach allows decoupling the operations having a cost dependent on the number of particles from those that depend on the instances of the required parameters. In particular, in each time step, the electric potential is approximated via dynamic mode decomposition (DMD) and the particle-to-grid map via a discrete empirical interpolation method (DEIM). These approximations are constructed from data obtained from a past temporal window at a few selected values of the parameters to guarantee a computationally efficient adaptation. The resulting DMD-DEIM reduced dynamical system retains the Hamiltonian structure of the full model, provides good approximations of the solution, and can be solved at a reduced computational cost.

2020 Mathematics Subject Classification. 65M99, 37N30, 65P10.

1 Introduction

The kinetic modeling of collisionless magnetized plasmas is based on the Vlasov–Maxwell equations, which describe the evolution of the distribution function of a collection of charged particles under the action of self-consistent electromagnetic fields. Because of the high dimensionality of the phase space, the large separation of scales, the inherent nonlinearity, and the infinitely many conserved quantities, the numerical treatment of the Vlasov–Maxwell equations, and of its electrostatic limit of Vlasov–Poisson equations, is a challenging task.

Arguably, the most widely used family of numerical methods for the solution of kinetic plasma models are Particle-In-Cell (PIC) methods [35]. The idea of PIC schemes is to sample the distribution function in velocity space using a finite number of macro-particles that are evolved along their characteristics. The electromagnetic fields are discretized on a grid in the computational domain, and the macro-particles move through the grid according to the Lorentz force. To preserve key physical properties of the problem, such

*Chair of Computational Mathematics and Simulation Science (MCSS), École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland. Email: Jan.Hesthaven@epfl.ch

[†]Department of mathematics, University of Pisa, Italy. Email: cecilia.pagliantini@unipi.it, c.pagliantini@tue.nl

[‡]Chair of Computational Mathematics and Simulation Science (MCSS), École Polytechnique Fédérale de Lausanne (EPFL), CH-1015 Lausanne, Switzerland. Email: nicolo.ripamonti@epfl.ch

as conservation of total energy, PIC schemes have evolved into variational algorithms based on least action principles [37, 36, 60] or on the Hamiltonian formulation of the Vlasov–Maxwell and Vlasov–Poisson equations [40, 17]. In parallel, several numerical methods [60, 17] for kinetic plasma models have leveraged discrete differential forms and de Rham complexes for the geometric approximation of the electromagnetic fields through Maxwell’s equations. Combining these ideas of structure-preserving approximations has led to finite element PIC methods able to exactly satisfy physical constraints, like the Gauss laws, and to guarantee the preservation of the Hamiltonian structure of the problem. Examples include the canonical [51] and non-canonical [66] symplectic particle-in-cell algorithms, the Hamiltonian particle-in-cell method of [22], and its generalization, the Geometric Electromagnetic PIC (GEMPIC) method [32].

The multiscale nature of plasmas implies that PIC codes require a significant amount of computational resources to resolve the shortest length scale and the fastest plasma frequency, and, thus, to yield stable and accurate numerical approximations. Moreover, the slow convergence rate of particle-based methods necessitates the use of many particles to achieve sufficient accuracy, to capture, for example, time-dependent physical phenomena such as plasma instabilities. As a result, PIC methods can be prohibitively expensive in terms of computational cost. This computational burden can become intractable in the parametric case, when simulations for many input parameters are of interest. This problem has been tackled from an algorithmic standpoint via the improvement of algorithms’ structure and the use of suitable computational hardware. In this work, we propose to address this computational issue through model order reduction. Starting from the high-resolution geometric PIC approximation of the Vlasov–Poisson problem, the idea is to derive a low-dimensional surrogate model that can be solved at a reduced computational cost and still provides accurate approximate solutions.

In recent years, several works have been devoted to the development of reduced representations of kinetic plasma dynamics based on model order reduction techniques, sparse grids and low-rank approximation [30, 31, 13, 15, 16, 14]. A dynamical low-rank approximation of the infinite-dimensional Vlasov–Poisson problem has been proposed in [15, 16, 14] with the intent to optimize the number of degrees of freedom needed for a sufficiently accurate and stable approximation of the solution. The continuous distribution function of the Vlasov–Poisson problem is expanded into a finite sum of low-rank factors, for which a new dynamical system is derived. Discretization in space and velocity of the resulting problem yields a low-rank approximation of the original dynamics. Another class of methods has considered the model reduction of the ordinary differential equation ensuing from the semi-discretization of the Vlasov–Poisson problem. A semi-Lagrangian method to solve the problem on a sparse grid is proposed in [31]. In [30], time splitting is combined with a spline-based semi-Lagrangian scheme where each step consists of linear combinations of low-rank approximations. In [13], the solution of the full-Eulerian time-dependent Vlasov–Poisson system is approximated using a tensor decomposition whose rank is adapted at each time step. To address the problem of the fluid closure for the collisionless linear Vlasov system, an interpolatory order reduction is proposed in [18]. In the context of a particle-based discretization of kinetic plasma models, a dynamic mode decomposition (DMD) strategy has been proposed in [44] to reconstruct the electric field within an Electromagnetic particle-in-cell (EMPIC) algorithm. Although the proposed approach can effectively capture and extrapolate the electric field behavior around equilibria, the computational burden associated with the high number of particles is not overcome. Model order reduction — in the number of particles — of *parametric* plasma models is, to the best of our knowledge, an open problem. Moreover, the conservation of the Hamiltonian structure has not been explored in reduced representations of plasma dynamics. This work aims at addressing these issues.

We focus on the parametric 1D-1V Vlasov–Poisson problem and consider its discretization in space and velocity via a geometric PIC method. The resulting dynamical system has a Hamiltonian form and, hence, corresponds to a symplectic flow. Model order reduction of the Vlasov–Poisson problem poses some major challenges, and standard reduced basis techniques are prone to fail in terms of numerical stability, computational efficiency, and accuracy of the simulations. The application [65] of the symplectic reduced basis methods proposed in [48, 1] to the Vlasov equation, with a fixed external electrostatic field, shows the importance of retaining the symplectic structure of Vlasov’s equation in the reduced model. However, the multi-scale nature of the problem makes it difficult for a reduced order model to characterize, with sufficient accuracy, the plasma behavior using a small number of degrees of freedom. This implies that accurate reduced representations of the solution may require large approximation spaces that eliminates the benefits of model order reduction. To overcome this limitation, we focus on methods

that adapt the approximate reduced space in time while preserving the geometric properties of the original problem. This approach allows to accurately describe the plasma evolution with a considerably reduced number of particles without compromising the quality of the simulation. The bulk of the computational effort to solve the reduced dynamics is due to the nonlinearity of the particles-to-grid mapping, and thus the Hamiltonian, whose evaluation needs to be performed in the original high-dimensional space. To alleviate these computational inefficiencies, we propose a strategy that approximates the reduced Hamiltonian gradient via a combination of hyper-reduction techniques and parameter sampling procedures. A reduction in the computational runtimes of the algorithm is achieved by decoupling the operations that depend on the number of particles from those that depend on the number of parameters while retaining an accurate representation of the plasma dynamics. The resulting dynamical system preserves the symplectic structure of the problem, ensures the stability of the approximation, and exploits the local-in-time low-rank nature of the solution.

The remainder of this work is organized as follows. In Section 2, the parametric Vlasov–Poisson problem is introduced both in its classical Eulerian formulation and in the Hamiltonian formulation. Moreover, the semi-discrete approximation of the problem via a particle method coupled with a finite element discretization of the Poisson problem is described. In Section 3, the model order reduction of the parametric dynamical system originating from the semi-discretization of the Vlasov–Poisson equation is considered. First, we describe a global-in-time symplectic reduced basis method, and then we consider a nonlinear approach, based on dynamical low-rank approximation, still preserving the symplectic structure of the flow field. After discussing the computational complexity of the dynamical reduced basis algorithm, Section 4 is devoted to the DMD-DEIM structure-preserving approximation of the nonlinear Hamiltonian gradient and the particles-to-grid mapping. The numerical temporal discretization of the resulting reduced model is discussed in Section 5. Numerical experiments in Section 6 on benchmark tests show that the proposed method can accurately reproduce the dynamics of particle-based kinetic plasma models with significant speedups compared to solving the original system. Concluding remarks are presented in Section 7.

2 The physical model

We consider the parametric Vlasov–Poisson problem with parameters that describe physical properties of the system. In particular, we focus on the study of the effect of parametrized initial distributions on the plasma dynamics. Let us assume that the parameters range in a compact set $\Gamma \subset \mathbb{R}^q$, with $q \geq 1$. The plasma, at any time $t \in \mathcal{T} \subset \mathbb{R}$, is described in terms of the distribution function $f(t, x, v; \eta)$ in the Cartesian phase space domain $(x, v) \in \Omega := \Omega_x \times \Omega_v \subset \mathbb{R}^2$. We focus on single species plasmas but the proposed method can be trivially extended to the multi-species case. Assume that $\Omega_x := \mathbb{T}^d = \mathbb{R}^d / (2\pi\mathbb{Z}^d)$ is the d -dimensional torus and $\Omega_v := \mathbb{R}^d$. For $\eta \in \Gamma$ fixed, we introduce the space

$$V_\eta := \{f(t, \cdot, \cdot; \eta) \in L^2(\Omega) : f(t, x, v; \eta) > 0 \text{ for all } (x, v) \in \Omega, f(t, \cdot, v; \eta) \sim e^{-v^2} \text{ as } |v| \rightarrow \infty\}.$$

The 1D-1V Vlasov–Poisson problem ($d = 1$) reads: For each $\eta \in \Gamma$ and $f_0(\eta) \in V_\eta|_{t=0}$, find $f(\cdot, \cdot, \cdot; \eta) \in C^1(\mathcal{T}; L^2(\Omega)) \cap C^0(\mathcal{T}; V_\eta)$, and the electric field $E(\cdot, \cdot; \eta) \in C^0(\mathcal{T}; L^2(\Omega_x))$ such that

$$\begin{cases} \partial_t f(t, x, v; \eta) + v \partial_x f(t, x, v; \eta) + \frac{q}{m} E(t, x; \eta) \partial_v f(t, x, v; \eta) = 0, & \text{in } \Omega \times \mathcal{T}, \\ \partial_x E(t, x; \eta) = q \int_{\Omega_v} f(t, x, v; \eta) dv, & \text{in } \Omega_x \times \mathcal{T}, \\ f(0, x, v; \eta) = f_0(\eta), & \text{in } \Omega. \end{cases}$$

Here q is the charge and m is the particle mass. The boundary conditions for f are periodic in space and prescribed via the space V_η in velocity.

Since the electric field can be written as the spatial derivative of the electric potential ϕ , namely

$E(t, x; \eta) = -\partial_x \phi(t, x; \eta)$, the Vlasov–Poisson problem can be recast, for each $\eta \in \Gamma$, as

$$\begin{aligned} \partial_t f(t, x, v; \eta) + v \partial_x f(t, x, v; \eta) - \frac{q}{m} \partial_x \phi(t, x; \eta) \partial_v f(t, x, v; \eta) &= 0, & \text{in } \Omega \times \mathcal{T}, \\ -\partial_{xx} \phi(t, x; \eta) = \rho(t, x; \eta) := q \int_{\Omega_v} f(t, x, v; \eta) dv, & & \text{in } \Omega_x \times \mathcal{T}, \end{aligned} \quad (2.1)$$

where ρ is the electric charge density. In the presence of a background charge, as in the numerical experiments, the right hand side of the Poisson problem reads $1 - \int_{\Omega_v} f(t, x, v; \eta) dv$ and the formulation of the method presented in the following sections can be extended straightforwardly.

The Lagrangian and Hamiltonian formulation of the Vlasov–Poisson and Vlasov–Maxwell equations reveals a set of mathematical and geometric features that encode the physical properties of these systems. The Vlasov–Poisson problem admits a Hamiltonian formulation, *cf.* [8, Sections 1 and 2], with a Lie–Poisson bracket, and Hamiltonian given by the sum of the kinetic and electric energy as

$$H(f, \eta) = \frac{m}{2} \int_{\Omega} v^2 f(t, x, v; \eta) dx dv + \frac{1}{2} \int_{\Omega_x} |E(t, x; \eta)|^2 dx. \quad (2.2)$$

Eulerian-based discretizations of kinetic plasma models in Hamiltonian form, with general noncanonical Poisson brackets, do not appear to inherit the phase space structure of the continuous problem, as has been observed in [42, 58]. On the contrary, particle-in-cell methods have led to the geometric approximation of these models when coupled to the discretization of the electromagnetic fields via discrete differential forms [66, 51, 22, 32]. For the structure-preserving approximation and reduction of the Vlasov–Poisson system (2.1) we rely on the Hamiltonian structure of its semi-discrete formulation obtained via particle-based methods as derived in the following.

2.1 Geometric particle-based discretization

We consider a particle method for the approximation of the Vlasov equation, coupled with a H^1 -conforming discretization of the Poisson problem for the electric potential. In detail, the distribution function f^s is approximated by the superposition of $N \in \mathbb{N}$ computational macro-particles as

$$f(t, x, v; \eta) \approx f_h(t, x, v; \eta) = \sum_{\ell=1}^N \omega_{\ell} S(x - X_{\ell}(t, \eta)) \delta(v - V_{\ell}(t, \eta)), \quad (2.3)$$

where $\omega_{\ell} \in \mathbb{R}$ is the weight of the ℓ -th particle, δ is the Dirac delta, S is a compactly supported shape function, and, for each $\eta \in \Gamma$ and $t \in \mathcal{T}$, $X(t, \eta) \in \mathbb{R}^N$ and $V(t, \eta) \in \mathbb{R}^N$ denote the vectors of the position and velocity of the macro-particles, respectively. In this work, we assume homogeneous macro-particles' weights so that $m_{\ell} = m$ and $\omega_{\ell} = \omega$, for all $\ell = 1, \dots, N$. The proposed method can be extended to the general case *mutatis mutandis*.

The idea of particle methods is to derive the time evolution of the approximate distribution function f_h by advancing the macro-particles along the characteristics of the Vlasov equation, i.e. the particles' positions and velocities satisfy the following set of ordinary differential equations

$$\begin{aligned} \dot{X}(t, \eta) &= V(t, \eta), & \text{in } \mathcal{T}, \\ \dot{V}(t, \eta) &= \frac{q}{m} E(t, X(t, \eta); \eta), & \text{in } \mathcal{T}, \end{aligned} \quad (2.4)$$

under suitable initial conditions. Here and throughout, the dot denotes the derivative with respect to time. The macro-particles move through a computational grid under the influence of electromagnetic fields. The latter are self-consistently calculated from the positions of the particles on the grid via the Poisson equation (2.1). On a partition of the spatial domain Ω_x , we consider a finite element discretization of the Poisson equation in the space $\mathcal{P}_k \Lambda^0(\Omega_x) \subset H^1(\Omega_x)$ of continuous piecewise polynomial functions of degree at most $k \geq 1$. The semi-discrete variational problem reads: For each $\eta \in \Gamma$, find $\phi_h(\cdot; \eta) \in C^1(\mathcal{T}; \mathcal{P}_k \Lambda^0(\Omega_x))$ such that

$$a_h(\phi_h(t, \cdot; \eta), \psi) = g(\psi), \quad \forall \psi \in \mathcal{P}_k \Lambda^0(\Omega_x), \quad (2.5)$$

where the bilinear form a_h corresponds to the Laplace operator, while the linear function g is associated with the density ρ ; thereby

$$a_h(\varphi, \psi) := \int_{\Omega_x} d_x \varphi(x) d_x \psi(x) dx, \quad g(\psi) := q \int_{\Omega} f_h(t, x, v; \eta) \psi(x) dv dx, \quad \forall \psi, \varphi \in \mathcal{P}_k \Lambda^0(\Omega_x).$$

Let $\{\lambda_i^0(x)\}_{i=1}^{N_x}$ be a basis of the space $\mathcal{P}_k \Lambda^0(\Omega_x)$, so that the semi-discrete potential can be written as $\phi_h(t, \cdot; \eta) = \sum_{i=1}^{N_x} \Phi_i(t, \eta) \lambda_i^0(x)$, for any $t \in \mathcal{T}$ and $\eta \in \Gamma$. We introduce the matrices $L \in \mathbb{R}^{N_x \times N_x}$ and $\Lambda^0(X) \in \mathbb{R}^{N \times N_x}$, defined as

$$L_{i,j} := (d_x \lambda_j^0, d_x \lambda_i^0)_{L^2(\Omega_x)}, \quad (\Lambda^0(X))_{\ell,i} := \lambda_i^0(X_\ell), \quad i, j = 1, \dots, N_x, \ell = 1, \dots, N. \quad (2.6)$$

Then the time-dependent algebraic system ensuing from (2.5) reads

$$L\Phi(t, \eta) = \Lambda^0(X(t, \eta))^\top M_q =: \rho_h(X(t, \eta); \eta), \quad (2.7)$$

where $M_q \in \mathbb{R}^N$ is the vector of entries $(M_q)_\ell = q\omega_\ell$, for $\ell = 1, \dots, N$. The proposed discretization of the electromagnetic field allows to recast the characteristic equations (2.4) as a Hamiltonian system. The phase space of Hamiltonian systems is characterized by a symplectic geometric structure. We denote with $\mathcal{V}_{2N} \subset \mathbb{R}^{2N}$ the phase space of (2.4) and we assume it is a $2N$ -dimensional symplectic vector space.

Let $W(t, \eta) = [X(t, \eta); V(t, \eta)] \in \mathcal{V}_{2N}$ denote the vector of all particle positions and velocities at a given time $t \in \mathcal{T}$ and parameter value $\eta \in \Gamma$, obtained by concatenating the vectors $X(t, \eta)$ and $V(t, \eta)$. The latter are also known as generalized position and momentum in the symplectic formalism. We denote with $W \in \mathcal{V}_{2N} \subset \mathbb{R}^{2N}$ the state vector collecting the positions and velocities of all particles of all species. To simplify the notation, we define the Hamiltonian $\mathcal{H}$ resulting from the proposed discretization of (2.2) up to the constant $m_p := m\omega$. First, we introduce the (nonlinear) electric energy $\mathcal{E} : \mathbb{R}^{2N} \times \Gamma \rightarrow \mathbb{R}$ defined as

$$\begin{aligned} \mathcal{E}(X(t, \eta); \eta) &:= \frac{m_p^{-1}}{2} \int_{\Omega_x} |\partial_x \phi_h(t, x; \eta)|^2 dx = \frac{m_p^{-1}}{2} \Phi(t, \eta)^\top L \Phi(t, \eta) \\ &= \frac{m_p^{-1}}{2} M_q^\top \Lambda^0(X(t, \eta)) L^{-1} \Lambda^0(X(t, \eta))^\top M_q. \end{aligned} \quad (2.8)$$

The Hamiltonian is the sum of the kinetic and electric energy and it is given, for any $W(t, \eta) \in \mathcal{V}_{2N}$, by

$$\mathcal{H}(W(t, \eta)) = \frac{1}{2} V(t, \eta)^\top V(t, \eta) + \mathcal{E}(X(t, \eta); \eta). \quad (2.9)$$

Differentiating the discrete Hamiltonian in (2.9) with respect to the vector X of particles' positions and the vector V of particles' velocities, results in the semi-discrete system in Hamiltonian form

$$\begin{pmatrix} \dot{X}(t, \eta) \\ \dot{V}(t, \eta) \end{pmatrix} = J_{2N} \begin{pmatrix} m_p^{-1} \text{diag}(M_q) \nabla \Lambda^0(X(t, \eta)) L^{-1} \Lambda^0(X(t, \eta))^\top M_q \\ V(t, \eta) \end{pmatrix}. \quad (2.10)$$

Here $\nabla \Lambda^0(X) \in \mathbb{R}^{N \times N_x}$ is defined as $(\nabla \Lambda^0(X))_{\ell,i} := (d_x \lambda_i^0)(X_\ell)$, for $i = 1, \dots, N_x$ and $\ell = 1, \dots, N$, and $J_{2N} \in \mathbb{R}^{2N \times 2N}$, called Poisson tensor, is the block matrix

$$J_{2N} = \begin{pmatrix} 0_N & I_N \\ -I_N & 0_N \end{pmatrix}, \quad (2.11)$$

with 0_N and I_N being the null and identity matrices of size N , respectively. A generalization of this discretization to the case of Vlasov–Maxwell's equations leads to the GEMPIC method introduced in [32]. Note that the proposed semi-discretization preserves the Hamiltonian, which corresponds to the discrete energy of the system, but not the momentum. In PIC it does not appear possible to simultaneously conserve momentum and energy [35].

3 Model order reduction of the Vlasov–Poisson problem

The parametric Hamiltonian system (2.10) reads: For each $\eta \in \Gamma$ and for $W_0(\eta) = [X(0, \eta); V(0, \eta)] \in \mathcal{V}_{2N}$, find $W(\cdot, \eta) \in C^1(\mathcal{T}, \mathcal{V}_{2N})$ such that

$$\begin{cases} \dot{W}(t, \eta) = J_{2N} \nabla_W \mathcal{H}(W(t, \eta)), & \text{in } \mathcal{T}, \\ W(0, \eta) = W_0(\eta), \end{cases} \quad (3.1)$$

where the initial condition $W_0(\eta) \in \mathcal{V}_{2N}$ is prescribed by the initial distribution $f_0(\eta)$.

3.1 Projection-based symplectic reduced basis methods

The parametric differential problem (3.1) is said to be reducible if the solution set $\mathcal{M} := \{W(t, \eta) \in \mathcal{V}_{2N} \subset \mathbb{R}^{2N} : t \in \mathcal{T}, \eta \in \Gamma\}$, which collects the solution of (3.1) under variation of time and parameter, can be well approximated by a $2n$ -dimensional subspace $\mathcal{M}_{2n}$, with $n \ll N$. The reducibility of a problem is measured via the so-called Kolmogorov n -width [49]: roughly speaking, the faster its decay as n increases, the better $\mathcal{M}$ can be approximated by a low dimensional space.

Within model order reduction techniques, projection-based reduced basis methods construct a modal approximation of the solution $W(t, \eta)$, for each $(t, \eta) \in \mathcal{T} \times \Gamma$, of the form $UZ(t, \eta)$, $U \in \mathbb{R}^{2N \times 2n}$ is a basis spanning the reduced approximation space $\mathcal{M}_{2n}$, and $Z(t, \eta) \in \mathcal{M}_{2n} \subset \mathbb{R}^{2n}$ are the expansion coefficients of the approximate reduced solution in the reduced basis. The reduced solution is then obtained by solving a lower dimensional dynamical systems obtained by the Galerkin projection of the full model (3.1) onto the reduced space. If the basis U is orthogonal, the reduced problem reads: For each $\eta \in \Gamma$ and $W_0(\eta)$, find $Z(t, \eta) \in \mathcal{M}_{2n}$ such that

$$\begin{cases} \dot{Z}(t, \eta) = U^\top J_{2N} \nabla_W \mathcal{H}(UZ(t, \eta)), & \text{in } \mathcal{T}, \\ Z(0, \eta) = U^\top W_0(\eta). \end{cases} \quad (3.2)$$

The reduced basis U is constructed from a collection of full order solutions at sampled values of time and parameters, via SVD-type methods, like POD, or greedy algorithms.

Note that (3.2) does not possess the Hamiltonian structure of the full model (3.1). In the setting of conservative systems, it is vital not only to provide accurate reduced representations of the solution but also to preserve the geometric structure underlying the physical properties and conservation laws of the dynamics. Moreover, violating intrinsic properties of the dynamics raises concerns about the reduced model's validity and reliability as a surrogate model. To preserve the Hamiltonian structure of the problem after reduction, symplectic reduced basis techniques have been proposed in [48, 1, 25]. Despite symplectic reduced basis methods providing stable and structure-preserving surrogate models, these might not be considerably more efficient to solve than the original model. Indeed, the non-dissipative nature of Hamiltonian problems is associated with poor reducibility and a slowly decaying Kolmogorov n -width. As a consequence, accurate reduced representations of the solution of (3.1) may require large approximation spaces that jeopardize the benefits of model order reduction. To overcome this limitation, we focus on adaptive methods where the reduced space changes in time while preserving the geometric properties of the original problem.

We remark that, in this work, we aim at tackling the challenges posed by problems that are not globally reducible *in time*. We implicitly assume that, for a fixed time, the solution under variation of parameter can be accurately represented by a low-rank approximation. The more general case of local reducibility with respect to both time and parameter is out of the scope of the present work and a short discussion on possible research in this direction is presented in the conclusions.

3.2 Dynamical structure-preserving model order reduction

For the model order reduction of the semi-discrete Vlasov–Poisson problem (3.1), we adopt the structure-preserving dynamical reduced basis method for Hamiltonian systems introduced in [45]. This approach can be related to both classical reduced basis methods and dynamical low-rank approximation [28].

Differently from a traditional reduced basis approach, the method does not have a separation into an offline phase and an online phase and it does not rely on snapshots. On the other hand, when the reduced basis is fixed, i.e. the approach becomes global, we recover a symplectic reduced basis method as in [48, 1]. As described in this section, the dynamical reduced models derived in [45] are obtained via a structure-preserving projection of the original evolution problem onto the tangent space of the reduced manifold, rather than on the reduced manifold itself. This is typical of dynamical low-rank approximation techniques and dynamically orthogonal (DO) field equations [54, 43]. We refer to [45, Introduction] for a detailed discussion on the relationship among these approaches.

Assume we seek to solve the semi-discrete Vlasov–Poisson problem (3.1) for p parameters $\Gamma_h := \{\eta_j\}_{j=1}^p \subset \Gamma$. The idea of the dynamical approach [45] is to find a low-rank approximation of the solution of problem (3.1) at each time instant. For this reason it is of interest to look at the properties of the solution, for each fixed time, under variation of the p parameters of interest. We therefore incorporate the parameters of interest in the problem formulation by recasting (3.1) as an evolution equation in the matrix-valued unknown $\mathcal{R}(t) := [W(t, \eta_1) | \dots | W(t, \eta_p)] \in \mathcal{V}_{2N}^p \subset \mathbb{R}^{2N \times p}$, with $\mathcal{V}_{2N}^p := \mathcal{V}_{2N} \times \dots \times \mathcal{V}_{2N}$. Given $\mathcal{R}_0 = [W_0(\eta_1) | \dots | W_0(\eta_p)]$, we look for $\mathcal{R} \in C^1(\mathcal{T}, \mathcal{V}_{2N}^p)$, such that

$$\begin{cases} \dot{\mathcal{R}}(t) = J_{2N} \nabla_{\mathcal{R}} \mathcal{H}^p(\mathcal{R}(t)), & \text{in } \mathcal{T}, \\ \mathcal{R}(0) = \mathcal{R}_0. \end{cases} \quad (3.3)$$

The Hamiltonian of (3.3) is a vector-valued quantity $\mathcal{H}^p : \mathcal{V}_{2N}^p \rightarrow \mathbb{R}^p$ that collects the values of the Hamiltonian of (3.1) for each parameter in Γ_h , namely $(\mathcal{H}^p(\mathcal{R}(\cdot)))_j = \mathcal{H}(W(\cdot, \eta_j))$, for $j = 1, \dots, p$. Moreover, the gradient $\nabla_{\mathcal{R}} \mathcal{H}^p(\mathcal{R}) \in \mathcal{V}_{2N}^p$ is defined as

$$(\nabla_{\mathcal{R}} \mathcal{H}^p(\mathcal{R}))_{\ell, j} = \frac{\partial \mathcal{H}(W(\cdot, \eta_j))}{\partial W_{\ell}(\cdot, \eta_j)}, \quad \ell = 1, \dots, 2N, \quad j = 1, \dots, p.$$

The idea of dynamical low-rank approximations is to expand the full order solution in a truncated modal decomposition where both the basis and the expansion coefficients are time-dependent. For all $t \in \mathcal{T}$, we approximate $\mathcal{R}(t)$, in (3.3), as $R(t) = U(t)Z(t)$, where $U(t) \in \mathbb{R}^{2N \times 2n}$ is the time-dependent basis and $Z(t) \in \mathbb{R}^{2n \times p}$ are the associated expansion coefficients, with $Z(t) := [Z_1(t) | \dots | Z_p(t)]$, and $Z_i(t) := Z(t, \eta_i)$. The reduced space is then defined as

$$\mathcal{M}_{2n}^p := \{R \in \mathbb{R}^{2N \times p} : R = UZ \text{ with } U \in \mathcal{U}(2n, \mathbb{R}^{2N}), Z \in \mathcal{Z}\}, \quad (3.4)$$

where

$$\mathcal{U}(2n, \mathbb{R}^{2N}) := \text{St}(2n, \mathbb{R}^{2N}) \cap \text{Sp}(2n, \mathbb{R}^{2N}),$$

is the manifold of orthosymplectic matrices, being $\text{St}(2n, \mathbb{R}^{2N}) := \{U \in \mathbb{R}^{2N \times 2n} : U^T U = I_{2n}\}$ the Stiefel manifold and $\text{Sp}(2n, \mathbb{R}^{2N}) := \{U \in \mathbb{R}^{2N \times 2n} : U^T J_{2N} U = J_{2n}\}$ the manifold of symplectic matrices. The factor Z in (3.4) is assumed to belong to the space

$$\mathcal{Z} := \{Z \in \mathbb{R}^{2n \times p} : \text{rank}(ZZ^T + J_{2n} Z Z^T J_{2n}) = 2n\}, \quad (3.5)$$

with $n \ll N$ and $2n < p$. The full-rank condition, (3.5), guarantees the uniqueness of U in the decomposition in (3.4), for all $t \in \mathcal{T}$, once Z is fixed. Let $X_r^i(t) := U_X(t)Z_i(t) \in \mathbb{R}^N$ and $V_r^i(t) := U_V(t)Z_i(t) \in \mathbb{R}^N$ denote the reduced velocity and position vectors, respectively, associated with the parameter η_i , for $i = 1, \dots, p$. The submatrices $U_X, U_V \in \mathbb{R}^{N \times 2n}$ are defined as $(U_X)_{i,j} = (U)_{i,j}$, $(U_V)_{i,j} = (U)_{i+N,j}$ for any $i = 1, \dots, N$ and $j = 1, \dots, 2n$. For fixed U and for each parameter $\eta_i \in \Gamma_h$, the flow map of the coefficient equation is a canonical symplectic map with a Hamiltonian having the i -th entry equal to

$$\mathcal{H}_{U,i}(Z_i(t)) := \mathcal{H}(U(t)Z_i(t)) = \frac{1}{2} V_r^i(t)^T V_r^i(t) + \mathcal{E}_{U,i}(Z_i(t)), \quad (3.6)$$

where the first part is quadratic in the coefficients Z_i , and $\mathcal{E}_{U,i} : \mathbb{R}^{2N} \rightarrow \mathbb{R}$ is the nonlinear electric energy component (2.8) of the Hamiltonian function, i.e., for all $i \in S_p$,

$$\mathcal{E}_{U,i}(Z_i(t)) = \mathcal{E}(X_r^i; \eta_i) = \frac{m_p^{-1}}{2} M_q^T \Lambda^0(X_r^i(t)) L^{-1} \Lambda^0(X_r^i(t))^T M_q. \quad (3.7)$$

We introduce the matrices

$$G_{\mathcal{H}}^p(U, Z) := [G_{\mathcal{H}}^p(U, Z_1) | \dots | G_{\mathcal{H}}^p(U, Z_p)] \in \mathbb{R}^{2N \times p},$$

and

$$g_{\mathcal{H}}^p(U, Z) := [g_{\mathcal{H}}^p(U, Z_1) | \dots | g_{\mathcal{H}}^p(U, Z_p)] \in \mathbb{R}^{2n \times p}$$

having as columns the p instances of the gradient of the Hamiltonian and of the reduced Hamiltonian, respectively,

$$G_{\mathcal{H}}^p(U, Z_i) := \nabla_{UZ_i} \mathcal{H}_{U,i}(Z_i) \in \mathbb{R}^{2N}, \quad g_{\mathcal{H}}^p(U, Z_i) := U^\top G_{\mathcal{H}}^p(U, Z_i) = \nabla_{Z_i} \mathcal{H}_{U,i}(Z_i) \in \mathbb{R}^{2n}. \quad (3.8)$$

where

$$\nabla_{UZ_i} \mathcal{H}_{U,i}(Z_i(t)) = \begin{pmatrix} m_p^{-1} \text{diag}(M_q) \nabla \Lambda^0(X_r^i(t)) L^{-1} \Lambda^0(X_r^i(t))^\top M_q \\ V_r^i(t) \end{pmatrix}. \quad (3.9)$$

Under these assumptions, a dynamical system for the reduced solution is characterized via the symplectic projection of the velocity field of the full dynamical system (3.3) onto the tangent space of $\mathcal{M}_{2n}^p$ at each state [45, 43]. The resulting reduced dynamics is given in terms of evolution equations for the reduced basis and for the expansion coefficients as

$$\begin{cases} \dot{Z}(t) = J_{2n} g_{\mathcal{H}}^p(U, Z) = J_{2n} \nabla_Z \mathcal{H}_U(Z), & \text{in } \mathcal{T}, & (3.10a) \\ \dot{U}(t) = (I_{2N} - UU^\top)(J_{2N} G_{\mathcal{H}}^p(U, Z) Z^\top - G_{\mathcal{H}}^p(U, Z) Z^\top J_{2n}^\top) S(Z)^{-1}, & \text{in } \mathcal{T}, & (3.10b) \\ U(t_0) Z(t_0) = U^0 Z^0, & & (3.10c) \end{cases}$$

where $S(Z) = ZZ^\top + J_{2n}^\top ZZ^\top J_{2n}$, and the initial condition $U^0 Z^0 \in \mathcal{M}_{2n}^p$ is computed via a truncated complex SVD of $\mathcal{R}_0 \in \mathbb{R}^{2N \times p}$. More details on the symplectic projection and on the gauge conditions enforced on the tangent space of $\mathcal{M}_{2n}^p$ can be found in [45].

Equation (3.10a) describes the evolution of the coefficients $Z(t)$ and is a system of p independent equations, each in $2n$ unknowns. It corresponds to the Galerkin projection of the full order Hamiltonian systems onto the space spanned by the columns of $U(t)$, as obtained with the global symplectic reduced basis method. Here, however, the basis U changes in time, and the evolution problem (3.10b), for the basis U , is a matrix-valued problem in $2N \times 2n$ unknowns on the manifold of ortho-symplectic rectangular matrices. Observe that the reduced basis depends on the parameters, but it is the same for all parameters in the set Γ_h .

4 Efficient treatment of nonlinear terms

In this section, we discuss the computational cost of the numerical solution of the reduced problem (3.10), and propose a novel algorithm for the efficient and structure-preserving treatment of the nonlinear operators.

The proposed reduction of the nonlinear terms is independent of the numerical time integrators used to solve the reduced dynamical system (3.10). However, the algorithm can be optimized depending on the time integrator of choice. We consider the structure-preserving partitioned Runge–Kutta temporal integrators proposed in [45, 26]. In particular, the evolution of the basis U is approximated with an explicit method, while a symplectic temporal integrator is employed for the evolution of the coefficients Z , and this latter will generally be an implicit scheme. Observe that we do not require that the stages of the RK integrators for the basis U and coefficients Z coincide. We will discuss the details and implementation of such schemes in Section 6. Although not strictly necessary, here we also assume that the first step of the partitioned RK method involves the evolution of the reduced basis; this assumption implies that we have the information on the Hamiltonian gradient at the beginning of the temporal interval (at least for some parameter values).

Let us split the temporal domain $\mathcal{T}$ into sub-intervals $\mathcal{T}_\tau := (t_{\tau-1}, t_\tau]$, for any $\tau = 1, \dots, N_\tau$, where $t_0 = 0$ and $\Delta t = t_\tau - t_{\tau-1}$ is the uniform time step. For each temporal interval $\mathcal{T}_\tau$, the dynamical reduced basis method involves the following operations.

- The evolution of the basis U requires $O(Nnp) + O(N_x^2 p) + O(Npc) + O(Nn^2) + O(n^2 p)$ flops, where $c \in \mathbb{N}$ is the number of finite element basis functions whose support is contained in a given mesh element, and we recall that N is the number of particles, n the size of the reduced basis and p denotes the number of parameter values. Note that c is a mild constant, and it is equal to 2 for piecewise polynomial functions in 1D, as in the discretization discussed in Section 2.1. The computational costs of this step are distributed as follows

Arithmetic complexity	Operation
$O(Nnp)$	computation of $V_r = U_V Z$ and $X_r = U_X Z$
$O(Npc)$	assembly of $\nabla \Lambda^0(X_r^i)$ and of $\Lambda^0(X_r^i)$, for all $i \in S_p$
$O(Npc) + O(N_x^2 p)$	computation of $\text{diag}(M_q) \nabla \Lambda^0(X_r^i(t)) L^{-1} \Lambda^0(X_r^i(t))^\top M_q$
$O(pn^2) + O(n^3)$	construction and inversion of the matrix $S(Z)$
$O(Nnp) + O(Nn^2)$	matrix-matrix multiplications in the r.h.s. of (3.10b)

The first three rows of the table correspond to the assembly and evaluation of $G_{\mathcal{H}}^p(U, Z)$ for all p parameters in Γ_h .

- The integration, using an implicit time scheme, of the evolution equation (3.10a) for the p vector-valued coefficients requires $O(Nnp) + O(N_x^2 p) + O(Npc)$ flops:

Arithmetic complexity	Operation
$O(Nn)$	computation of $U_V^\top V_r^i = U_V^\top U_V Z_i$
$O(Nn) + O(N_x^2) + O(Nc)$	assembly and evaluation of $G_{\mathcal{H}}^p(U, Z_i)$
$O(Nn)$	computation of $U_X^\top G_{\mathcal{H}}^p(U, Z_i)$

Each of the operation listed in the table needs to be performed for each parameter $\eta_i \in \Gamma_h$, at each stage of the RK scheme, and at each iteration of the nonlinear solver.

The leading computational cost in both steps depends on the product of the number of particles N and of the number of parameter p , both potentially large in multi-query simulations of high-dimensional problems. This cost is associated with the remapping of the particles to the full dimensional space, in each temporal interval and for each parameter, and with the evaluation of the velocity field of the reduced flow. Indeed, the sole knowledge of the expansion coefficients with respect to the reduced basis is not enough to compute the particles-to-grid mapping needed to evaluate the electric field and, hence, the Hamiltonian. Even in the reduced model (3.10), these operations require the reconstruction of the approximate particle positions, at a cost proportional to the size of the full model. This lifting to the high-dimensional space needs to be performed for each instance of the parameter, at each stage of the RK time integrator, and at every iteration of the nonlinear solver. Analogous computational problems are common in model order reduction and emerge whenever non-affine and nonlinear operators are involved, *cf.* e.g. [52, Chapters 10 and 11]. In the presence of low-order polynomial nonlinearities, tensorial techniques [61] can be used to separate terms that depend on full spatial variables and on reduced coefficients to allow efficient computations of the nonlinear terms. The non-polynomial nature of the nonlinearity in the gradient of (3.6) prevents us from using the aforementioned tensorial approach to accelerate the computation. For general nonlinear operators, several so-called *hyper-reduction* techniques have been proposed to mitigate or overcome the computational bottleneck associated with the lifting. The empirical interpolation method (EIM) [3] and its discrete counterpart (DEIM) [9] are interpolatory techniques used to approximate the nonlinearity in the projection-based ROM, requiring the computations of only a few components of the original nonlinearity. While effective in the case where each component of the nonlinearity depends only on a few components of the input, they are not suited for the treatment of not component-wise nonlinear terms. Using a sparsity argument [9] or the introduction of auxiliary variables [46], DEIM has been

adapted to deal with the approximation of the nonlinear terms at interpolation points that require the evaluation of the reduced solution on a limited number of neighboring mesh points, as it happens for high-order spatial discretization schemes with large stencils. However, the same strategy does not work for the treatment of the gradient of (3.6), as the inverse of the discrete Laplacian operator is generally dense, and hence each of its entries requires the computation of X_r^i for p sampled parameters, making traditional approaches computationally impractical. Moreover, traditional hyper-reduction techniques, like EIM and DEIM, applied to a gradient vector field do not result in a gradient field, which means that the geometric structure of the Hamiltonian dynamics is compromised in the hyper-reduction process.

To achieve computational efficiency in the simulation of (3.10) without compromising its geometric structure, we propose a strategy that approximates the reduced Hamiltonian gradient via a combined hyper-reduction technique and sampling procedure. A reduction in the computational runtimes of the algorithm is achieved by decoupling the operations that depend on N from those that depend on p , while retaining an accurate representation of the plasma dynamics. There are several challenges that we need to face in the development of such techniques: (i) the preservation of the Hamiltonian structure of the dynamics; (ii) the lack of information on the full model solution and nonlinear operators, traditionally collected in an offline phase via snapshots; and (iii) the lack of a sparsity pattern in the nonlinear Hamiltonian gradient, i.e., the fact that each entry of the electric energy vector (3.7) depends on all N computational particles.

4.1 Parameter sampling

The reduced dynamics (3.10) involve p evolution equations for the expansion coefficients, one per parameter value, and one evolution problem for the matrix-valued reduced basis. Since, at each time, the reduced basis is the same for all parameters, one can reduce the computational cost required for its evolution by sampling over the parameter space and constructing a reduced basis for only a subsample of parameters, but which remains accurate for all other parameters in Γ_h . This corresponds to a reduction in parameter space. Let us denote by p the cardinality of the set Γ_h and assume that the parameters in Γ_h are indexed from the set $S_p := \{1, \dots, p\}$. Let us consider a subset Γ_h^* of Γ_h of size $p^* \ll p$. Define $S_{p^*} \subset S_p$ to be the set of indices corresponding to the parameters in the selected subset so that $\Gamma_h^* = \{\eta_i \in \Gamma_h \mid i \in S_{p^*}\} \subset \Gamma_h$. The idea of the proposed sampling approach is to replace, in the evolution of the basis (3.10b), the matrix of the expansion coefficients $Z(t) \in \mathbb{R}^{2n \times p}$ by the matrix obtained via the concatenation of the columns of Z with indices in S_{p^*} . Following the discussion at the beginning of the section, this approximation leads to a computational complexity for the basis evolution of the order of $O(Nnp^*) + O(N_x^2 p^*) + O(Np^*c) + O(Nn^2) + O(n^2 p^*)$. To preserve the accuracy of the method, we must ensure that the chosen subset Γ_h^* is representative of the entire parameter set Γ_h . For the sake of simplicity, in this work, we set it at $t = 0$, and we keep it fixed over time. Starting from $\Gamma_h^* = \emptyset$, the set Γ_h^* is constructed using a greedy algorithm that, at each iteration, adds to the index subset S_{p^*} the index i that satisfies

$$\max_{i \in \Gamma_h \setminus \Gamma_h^*} \min_{j \in \Gamma_h^*} \|Z_i^0 - Z_j^0\|_2, \quad (4.1)$$

until a user-defined threshold value of the cost function or a maximum number of iterations. Possible research directions to improve the selection strategy would be to adapt in time the set Γ_h^* to capture significant changes in the behavior of the solution relative to the parameters or to modify the cost function to incorporate errors in the evaluation of physical quantities, such as the electric field. However, as (4.1) is an NP-complete problem [59], it is not currently known if it is possible to find an optimal parameter selection strategy with a polynomial cost in p . Thus, in the current form, while affordable if performed only once at $t = 0$, the selection strategy may become computationally expensive for large p and could compromise the efficiency of the proposed dynamical RB method if repeated in each temporal sub-interval. Suboptimal, yet more efficient, algorithms could be adopted by framing the problem into the more general column subset selection problem (CSSP) [5], that consists in finding an optimal subset of columns of a given matrix that minimizes the residual of the projection of the given matrix onto the selected column subset. Other parameter reduction strategies, like active subspaces [10], might also be envisioned.

Concerning the expansion coefficients, for which one differential equation per parameter needs to be solved, subsampling is not an option.

4.2 DMD-DEIM approximation of the Hamiltonian gradient

In this section, we develop a reduction algorithm where, in each temporal interval $\mathcal{T}_\tau$, DMD is used for the hyper-reduction of the electric potential $\Phi(X_r^i(t)) = L^{-1}\Lambda^0(X_r^i(t))^\top M_q$ in (3.7), while a DEIM strategy is developed to approximate the component $\Lambda^0(X_r^i(t))$ of the particles-to-grid mapping. Note that in (3.6), the quadratic term involving the particles' velocity represents a linear contribution in the gradient of the reduced Hamiltonian and, hence, does not require any hyper-reduction.

4.2.1 Dynamic Mode Decomposition of the electric potential

Dynamic mode decomposition is an equation-free data-driven approach, proposed in [56, 57], that uses only data measurements of a given dynamical system to approximate the dynamics and predict future states. The idea is to decompose the problem into a set of coherent spatial structures, known as DMD modes, and associate correlated data to specific Fourier modes that capture temporal variations. DMD was initially employed as a spectral decomposition method for complex fluid flows [53]. More recently, it has proved successful in a wide range of settings such as background/foreground separation in real-time video [19], characterization of dynamic stall [39], and analysis of the propagation of infectious diseases [50]. DMD hinges on the theory of Koopman operators [29], which allows to represent the flow of a nonlinear dynamical system via an infinite-dimensional linear operator on the space of measurement functions. Without explicit knowledge of the operator describing the dynamics, DMD computes a least-squares regression of data measurements to an optimal finite-dimensional linear dynamical system that approximates the infinite-dimensional Koopman operator. In this subsection, we first describe the classical DMD algorithm following [34]. Next, we introduce a sliding-window based DMD formulation for the hyper-reduction of the electric potential in the dynamical reduced model (3.10) of the Vlasov–Poisson problem.

Consider a general nonlinear dynamical system: Find $\mathbf{y} : \mathcal{T} \rightarrow \mathbb{R}^\ell$, for $\ell \in \mathbb{N}$ such that

$$\begin{cases} \dot{\mathbf{y}}(t) = F(t, \mathbf{y}(t)), & t \in \mathcal{T}, \\ \mathbf{y}(t_0) = \mathbf{y}_0. \end{cases} \quad (4.2)$$

Assume that we have as data measurements exact values or approximations of the state at different time instants, namely

$$\mathbf{Y} = [\mathbf{y}_0 \quad \mathbf{y}_1 \quad \dots \quad \mathbf{y}_{\tau-1}] \in \mathbb{R}^{\ell \times \tau}, \quad \mathbf{Y}' = [\mathbf{y}_1 \quad \mathbf{y}_2 \quad \dots \quad \mathbf{y}_\tau] \in \mathbb{R}^{\ell \times \tau}, \quad (4.3)$$

where $\mathbf{y}_k = \mathbf{y}(k\Delta t)$ and Δt is the uniform time step. In the DMD method, data measurements are used to approximate the nonlinear dynamics (4.2) by a locally linear system $\dot{\mathbf{y}} = A\mathbf{y}$, where $A \in \mathbb{R}^{\ell \times \ell}$ is the matrix that best fits the measurements in a least-square sense, i.e., $A = \arg\min_B \|\mathbf{Y}' - B\mathbf{Y}\|_F$, with $\|\cdot\|_F$ being the Frobenius norm. Then, A is given by $A = \mathbf{Y}'\mathbf{Y}^\dagger$, where $\dagger$ denotes the Moore–Penrose pseudoinverse.

From the linear approximation of the dynamics, the DMD algorithm computes a low-rank eigendecomposition of the matrix A by extracting its r_τ largest eigenvalues Λ^A and corresponding eigenvectors $\Theta^A = [\theta_1^A \dots \theta_{r_\tau}^A] \in \mathbb{R}^{\ell \times r_\tau}$. The resulting DMD approximation of the state $\mathbf{y}(t)$, for $t > \tau\Delta t$, reads

$$\mathbf{y}(t) \approx \mathbf{y}_{\text{DMD}}(t) = \Theta^A \left(\Pi \odot e^{\Omega(t-\tau\Delta t)} \right) = \sum_{j=1}^{r_\tau} \theta_j^A \pi_j e^{\omega_j(t-\tau\Delta t)}, \quad (4.4)$$

where, for any $j = 1, \dots, r_\tau$, $\omega_j := \ln(\Lambda_j^A)/\Delta t$ is the j -th entry of the vector $\Omega \in \mathbb{R}^{r_\tau}$, while π_j is the j -th entry of the vector $\Pi = (\Theta^A)^\dagger \mathbf{y}_0 \in \mathbb{R}^{r_\tau}$ containing the coordinates of the initial condition $\mathbf{y}_0$ with respect to the DMD modes.

If the size of the matrix A is large, A might be severely ill-conditioned and not directly tractable. In this situation, a different version of the DMD algorithm, proposed in [64], projects the data into a

Algorithm 1: DMD algorithm

Input : $\mathbf{Y}, \mathbf{Y}', \text{tol}$.

Output: Λ, Θ .

- 1 Compute the truncated SVD of $\mathbf{Y}, \mathbf{Y} = U\Sigma V^\top$, using tol as tolerance for singular values selection;
 - 2 Define $A_{\text{tol}} = U^\top \mathbf{Y}' V \Sigma^{-1}$;
 - 3 Compute the eigendecomposition of A_{tol} : $A_{\text{tol}} W = W \Lambda$;
 - 4 Reconstruct the eigendecomposition of A by defining its eigenvectors as $\Theta = \mathbf{Y}' V \Sigma^{-1} W$.
-

low-rank subspace instead of deriving A directly from the data, as described in Algorithm 1. Moreover, given the sensitivity of the DMD algorithm to the duration and sampling of the series $\mathbf{Y}$ and $\mathbf{Y}'$, [12] proposes a sliding-window approach where the measurement data are not taken in the whole temporal interval but only in the sampling window $[t_{\tau-T}, t_\tau]$ of length $T\Delta t$, with $T \in \mathbb{N}$. The rationale is that if the system is time-varying and the incoming data is harvested in a streaming fashion, it may be beneficial to accuracy and memory storage to consider only the most recent data. The only computational overhead is the computation of the DMD modes and weights in the DMD approximation (4.4) as new data are collected. This cost may be mitigated by efficient online updates of the eigenvalues and eigenvectors of A [23] or by means of incremental SVD algorithms [41].

In the context of kinetic plasma PIC simulations, a DMD strategy has been used in [44] to detect and track equilibrium states. The aforementioned method relies on snapshots of the high-fidelity simulation until an equilibrium is detected and, after this time, the solution is extrapolated via the DMD modes. Here, we propose to employ a DMD strategy in a different way, namely to hyper-reduce the self-consistent electric potential $\Phi(X_r^i(t)) = L^{-1} \Lambda^0 (X_r^i(t))^\top M_q \in \mathbb{R}^{N_x}$ that enters the reduced Hamiltonian (3.6) for each parameter $\eta_i \in \Gamma_h$. The idea is to extract low-dimensional dynamical features from a time-series of the electric potential and use them, as part of the DMD algorithm in (4.4), to extrapolate the value of $\Phi(X_r^i(t))$ needed for the computation of the reduced Hamiltonian (3.6) in each temporal interval $\mathcal{T}_\tau$. In details, let Φ_τ^i , for a fixed parameter with index $i \in S_{p^*}$ and $\tau = 1, \dots, N_\tau$, be the approximation of $\Phi(X_r^i(t))$ at $t = t_\tau$ for each parameter in the subset Γ_h^* . Since the first step of the temporal integrator involves the evolution of the reduced basis, these quantities are computed while assembling the right hand side of the basis evolution equation (3.10b). For each $\mathcal{T}_\tau$, we collect the time-discrete approximations of the electric potential obtained in a past time window of length $(T+1)\Delta t$;

$$\mathbf{Y}_i = \begin{bmatrix} \Phi_{\tau-T-1}^i & \Phi_{\tau-T}^i & \cdots & \Phi_{\tau-2}^i \end{bmatrix}, \quad \mathbf{Y}'_i = \begin{bmatrix} \Phi_{\tau-T}^i & \Phi_{\tau-T+1}^i & \cdots & \Phi_{\tau-1}^i \end{bmatrix}, \quad (4.5)$$

where $\mathbf{Y}_i, \mathbf{Y}'_i \in \mathbb{R}^{N_x \times T}$ for $i = 1, \dots, p^*$. Extracting the dominant modes from each realization of the electric potential associated with a fixed parameter is a cumbersome task. To the best of our knowledge, DMD-based methods for the model order reduction of parametric problems have not been developed. In our setting, the dependence on the parameter comes from the state, and it is propagated via the parametric initial distribution f_0 . This suggests that, instead of extracting the DMD modes for each fixed parameter $\eta_i \in \Gamma_h^*$, we can incorporate the parameter in the DMD procedure to approximate the dynamics of the electric potential for all parameters. A similar approach can be found in [55] when dealing with bifurcation parameters in thermo-acoustic systems. For each parameter index $i \in S_{p^*}$, the datasets $\mathbf{Y}_i$ and $\mathbf{Y}'_i$ are concatenated column-wise to form two global datasets $\mathbf{Y}$ and $\mathbf{Y}'$, i.e.

$$\mathbf{Y} = \begin{bmatrix} \mathbf{Y}_1 & \mathbf{Y}_2 & \cdots & \mathbf{Y}_{p^*} \end{bmatrix}, \quad \mathbf{Y}' = \begin{bmatrix} \mathbf{Y}'_1 & \mathbf{Y}'_2 & \cdots & \mathbf{Y}'_{p^*} \end{bmatrix}, \quad (4.6)$$

with $\mathbf{Y}, \mathbf{Y}' \in \mathbb{R}^{N_x \times p^* T}$. This procedure is justified by the absence of an explicit dependence of the electric potential on the parameter. Following Algorithm 1, we generate the DMD eigenvectors $\Theta \in \mathbb{R}^{N_x \times r_\tau}$ and eigenvalues $\Lambda \in \mathbb{R}^{r_\tau \times r_\tau}$ of the linear approximation of the dynamics for the problem of interest. The resulting DMD approximation of the self-consistent electric potential reads

$$\Phi(X_r^i(t)) \approx \Phi_{\text{DMD}}^i(t) = \Theta \left(\Pi_i \odot e^{W(t-(\tau-1)\Delta t)} \right), \quad \forall i \in S_{p^*}, \forall t \in \mathcal{T}_\tau, \quad (4.7)$$

with $W \in \mathbb{R}^{r_\tau}$ is the vector of entries $\omega_j = \ln(\Lambda_j)/\Delta t$, for any $j = 1, \dots, r_\tau$, and $\Pi_i := \Theta^\dagger \Phi_{\tau-1}^i \in \mathbb{R}^{r_\tau}$ for any $i \in S_{p^*}$.

Assuming a smooth dependence of the DMD coordinates Π_i on the parameter, interpolation techniques can be used to recover the DMD coordinates for parameters not included in Γ_h^* , similarly to the POD with interpolation (PODI) [7]. In this work, we adopt the radial basis interpolation [6], with a Gaussian kernel, as interpolation algorithm to reconstruct the DMD coordinates Π_i for $i \in \Gamma_h \setminus \Gamma_h^*$. The computational cost of the interpolation step is negligible as compared to the cost of Algorithm 1, as we comment on at the end of the section. For ease of the notation, we use the same symbol Π_i to represent the interpolated DMD coefficients for all $i \in S_p$. Knowing Π_i for all $i \in S_p$, the electric potential is reconstructed using (4.7). The resulting sampling error can be controlled by enriching the subset of parameters Γ_h^* and by optimal placement of the location of the parameters with indices in S_{p^*} in the parameter space.

Using the DMD estimate of the potential Φ , the Hamiltonian function (3.6) is approximated as

$$\mathcal{H}_{U,i}^{\text{DMD}}(Z_i) = \frac{1}{2} V_r^i(t)^\top V_r^i(t) + \frac{m_p^{-1}}{2} M_q^\top \Lambda^0(X_r^i(t)) \Phi_{\text{DMD}}^i(t), \quad \forall i \in S_p, \forall t \in \mathcal{T}_\tau. \quad (4.8)$$

Remark 4.1. If the Hamiltonian function depends explicitly on the parameter, the approach outlined above is not legitimized because a non-parametric operator would be used to approximate the parametric potential. An alternative strategy would require an approximation of the form (4.7) for each parameter realization η_i , with $i \in S_{p^*}$, using different W_i and Θ_i for each i . The resulting DMD approximations of the potential $\Phi_{\text{DMD}}^i(t)$ could then be directly interpolated on S_p or, as suggested in [11], interpolated based on physical concepts as in PODI.

The DMD approximation based on the matrix A , instead of its projection A_{tol} defined in Algorithm 1, would result in a computational complexity $O(N_x^3)$. Although this cost might still be tractable in one dimension, it becomes prohibitive when considering the Vlasov–Poisson problem in higher dimension. The method described in Algorithm 1 is, therefore, the preferred choice.

The computational cost of the proposed DMD strategy, applied to the electric potential, reduces to the cost needed to perform Algorithm 1 from the datasets $\mathbf{Y}, \mathbf{Y}' \in \mathbb{R}^{N_x \times p^* T}$ in (4.6). The truncated SVD decomposition of $\mathbf{Y}$ in Line 1 has arithmetic complexity $O(N_x p^* T r_\tau)$, where r_τ is the number of retained modes [21]. Observe that, if the number r_τ of truncated modes is chosen based on a tolerance to control the magnitude of the neglected singular values, Algorithm 1 computes the full SVD of $\mathbf{Y}$ and then performs the truncation. This variant of Line 1 has computational complexity $O(N_x (p^* T)^2)$, under the assumption that the number T of chosen DMD samples and the number of parameter subsamples p^* satisfy $N_x > p^* T$. The eigendecomposition of $A_{\text{tol}} \in \mathbb{R}^{r_\tau \times r_\tau}$ in Line 3 costs $O(r_\tau^3)$. Finally, the matrix-matrix multiplications to compute A_{tol} and Θ in Line 2 and Line 4, respectively, require $O(N_x r_\tau^2) + O(N_x p^* T r_\tau)$ operations. The computational cost to compute Φ_{DMD}^i for every parameter $\eta_i \in \Gamma_h$ – including sampling parameters and reconstructed parameters – is $O(N_x r_\tau p)$. The leading cost is, therefore, $O(N_x r_\tau p) + O(N_x p^* T r_\tau)$, with the last term replaced by $O(N_x (p^*)^2 T^2)$ for a naive implementation of the truncated SVD. This cost is linear in N_x , does not depend on the number N of particles, and only the computation of the DMD coordinates Π_i depends on the number of parameters p .

4.2.2 Discrete Empirical Interpolation Method for reduction in the number of particles

The DMD approach described in the previous section allows to derive an approximate electric potential that can be evaluated independently on the number of particles. However, the evaluation of the electric energy component of the approximate reduced Hamiltonian (4.8) still requires the particles-to-grid mapping for $\Lambda^0(X_r^i(t))$, for each value of the parameter, at each stage of the temporal solver, and at each iteration of the nonlinear solver. The computational cost of this step is a major bottleneck of the algorithm. To overcome this computational burden, we propose hyper-reduction of the approximate reduced Hamiltonian (4.8) with a DEIM-based strategy.

The DEIM approach is a discrete variant of the empirical interpolation method (EIM) introduced in [3] to approximate nonlinear functions via a combination of projection and interpolation. DEIM constructs carefully selected interpolation indices to specify an interpolation-based projection so that the complexity of evaluating the nonlinear term becomes proportional to the (small) number of selected spatial indices.

The gist of the DEIM approximation of a given nonlinear operator $F : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is to replace it with its oblique projection into a lower-dimensional space obtained from the most relevant components of F . Starting from snapshots of the nonlinear function F , DEIM constructs a lower-dimensional space as the span of the left singular vectors associated with the $d \ll N$ largest singular values of the snapshot matrix. If $\Psi = [\psi_1 \dots \psi_d] \in \mathbb{R}^{N \times d}$ denotes the DEIM basis, the DEIM approximation of F reads

$$F(\mathbf{x}) \approx \Psi(P^\top \Psi)^{-1} P^\top F(\mathbf{x}), \quad \forall \mathbf{x} \in \mathbb{R}^N, \quad (4.9)$$

where $P \in \mathbb{R}^{N \times d}$ is defined as $P := [\mathbf{e}_{\ell_1} \dots \mathbf{e}_{\ell_d}]$, with $\mathbf{e}_k = \{0, 1\}^N$ being the k -th canonical vector. We denote by I_{DEIM} the set of indices corresponding to the DEIM sampling points so that $I_{\text{DEIM}} = \{\ell_k\}_{k=1}^d$. The DEIM sampling indices are iteratively chosen according to the greedy procedure proposed in [9, Algorithm 1]. The k -th step, for $1 \leq k \leq d$, of the algorithm consists in choosing the index of the k -th DEIM vector ψ_k that maximizes the DEIM projection error of ψ_k into the current DEIM basis $[\psi_1 \dots \psi_{k-1}] \in \mathbb{R}^{N \times (k-1)}$. Under the assumption that each entry of the nonlinear vector-valued function F depends only on a few entries of the input vector $\mathbf{x}$, the computation of $P^\top F(\mathbf{x})$ requires only d evaluations of the nonlinear function leading to a computational cost independent of the original dimension N .

The application of the classical DEIM procedure for the hyper-reduction of the nonlinear Hamiltonian gradient (3.9) is challenged by several factors. As stated at the beginning of the section, applying the DEIM interpolation directly to the right-hand side of the coefficients evolution equations (3.10a) arising from the dynamic reduced basis approach, would not result in a structure-preserving approximation. Moreover, the classical DEIM algorithm hinges on the availability of snapshots of the full model nonlinear operator of interest collected in the offline phase. In our dynamical model order reduction approach, there is no offline phase and, therefore, snapshots are not available.

We consider the Hamiltonian splitting in (3.6), and the approximation of the reduced electric energy (3.7) resulting from DMD, namely

$$\mathcal{E}_{U,i}^{\text{DMD}}(Z_i(t), t) := \frac{m_p^{-1}}{2} M_q^\top \Lambda^0(X_r^i(t)) \Phi_{\text{DMD}}^i(t), \quad \forall i \in S_p, \forall t \in \mathcal{T}_\tau,$$

where Φ_{DMD}^i is defined in (4.7). Approximating directly the vector $\nabla \mathcal{E}_{U,i}^{\text{DMD}}$ by a DEIM interpolation, as in (4.9), would not preserve the geometric structure of the problem because it is not possible to define explicitly an Hamiltonian gradient from the interpolated vector field. We propose a DEIM approximation of the reduced electric energy via hyper-reduction of the term $\Lambda^0(X_r^i(t)) \in \mathbb{R}^{N \times N_x}$, which otherwise would require the evaluation of the finite element basis functions at each particle position.

Let us introduce the function $\mathcal{N}_i(X_r^i(t), t) := \Lambda^0(X_r^i(t)) \Phi_{\text{DMD}}^i(t) \in \mathbb{R}^N$; approximated using a DEIM approach in each temporal interval $\mathcal{T}_\tau$ as follows. First, we consider samples of the nonlinear term associated with the electric potential Φ (2.7) at p^* instances of the parameter and over a temporal window of length $(T+1)\Delta t$. These are obtained from simulations of the reduced dynamical system in the past window of interest. The data matrix $\mathbf{Y} \in \mathbb{R}^{N \times p^*(T+1)}$ is defined as

$$\mathbf{Y} = \begin{bmatrix} \mathbf{Y}_1 & \mathbf{Y}_2 & \dots & \mathbf{Y}_{p^*} \end{bmatrix}, \quad \mathbf{Y}_i := \begin{bmatrix} \Lambda^0(X_r^i(t_{\tau-T-1})) \Phi_{\tau-T-1}^i & \dots & \Lambda^0(X_r^i(t_{\tau-1})) \Phi_{\tau-1}^i \end{bmatrix}. \quad (4.10)$$

Note that the terms $\Lambda^0(X_r^i(t_{\tau-j-1}))$ and $\Phi_{\tau-j-1}^i$, for $i \in \Gamma_h^*$ and $j = 0, \dots, T$, are available from the evolution equation (3.10b) for the reduced basis solved at previous time steps. The DEIM basis matrix $\Psi_\tau \in \mathbb{R}^{N \times d}$ is obtained by taking the first d left singular vectors of the matrix $\mathbf{Y}$, where the value d is fixed at the beginning of the simulation and might differ for different problems. We will comment on this in the numerical experiments in Section 6. Denoting with $P_\tau \in \mathbb{R}^{N \times d}$ the matrix corresponding to the DEIM indices obtained as describe above, the nonlinear term $\mathcal{N}_i(X_r^i(t), t)$ is approximated by

$$\Psi_\tau^\top (P_\tau^\top \Psi_\tau)^{-1} P_\tau^\top \mathcal{N}_i(X_r^i(t), t), \quad \forall i \in S_p, \forall t \in \mathcal{T}_\tau.$$

Observe that, although the basis Ψ_τ is constructed from the parameter subsample Γ_h^* , the nonlinear term $\mathcal{N}_i$ is approximated by its DEIM projection onto the DEIM space for all instances of the parameter, i.e., for all $i \in S_p$.

To reduce the computational burden associated with the computation of the DEIM sampling points P_τ in each temporal interval $\mathcal{T}_\tau$, we follow an update strategy similar to the one proposed in [47]. All the

interpolation indices in the set I_{DEIM} are computed using the standard DEIM greedy method; not at all time steps but only every $k_{\text{DEIM}} > 1$ time steps. In other temporal intervals, we proceed as follows. Assume we have computed the set of DEIM indices $I_{\text{DEIM}}^{\tau-1}$ in the temporal interval $\mathcal{T}_{\tau-1}$, then, in the following interval $\mathcal{T}_\tau$, we update only the indices in the subset $I^* \subset I_{\text{DEIM}}^{\tau-1}$ of cardinality n_{DEIM} given by

$$I^* = \underset{\substack{I \subset I_{\text{DEIM}}^{\tau-1}, \\ \dim(I) = n_{\text{DEIM}}}}{\operatorname{argmax}} \sum_{k \in I} (\psi_k^\tau)^\top \psi_k^{\tau-1}, \quad (4.11)$$

where ψ_k^τ denotes the k -th vector of the DEIM basis Ψ_τ at time t_τ . The remaining $d - n_{\text{DEIM}}$ indices in $I_{\text{DEIM}}^{\tau-1} \setminus I^*$ are inherited by I_{DEIM}^τ . The rationale for the choice of I^* is to only update the indices associated with the DEIM basis vectors at $t_{\tau-1}$ that have undergone the largest rotations in the DEIM basis update from $\Psi_{\tau-1}$ to Ψ_τ .

The resulting approximate reduced Hamiltonian, associated with the parameter $\eta_i \in \Gamma_h$, reads

$$\mathcal{H}_{U,i}^{\text{DD}}(Z_i(t), t) = \frac{1}{2} V_r^i(t)^\top V_r^i(t) + \frac{m_p^{-1}}{2} M_q^\top \Psi_\tau (P_\tau^\top \Psi_\tau)^{-1} P_\tau^\top \Lambda^0(X_r^i(t)) \Phi_{\text{DMD}}^i(t), \quad \forall i \in S_p, \forall t \in \mathcal{T}_\tau. \quad (4.12)$$

Observe that the multiplication of the matrix $\Lambda^0(X_r^i)$, of the finite element basis functions evaluated at the particles' position, by the DEIM sampling matrix $P_\tau^\top$, corresponds to evaluating the finite element basis functions only on a subset of $d \ll N$ particles. Hence, this operation represents a substantial reduction in the number of particles.

The computational cost of the DEIM algorithm can be summarized as follows. The computation of the data matrix in (4.10) only involves the multiplications of the terms $\Lambda^0(X_r^i(t_{\tau-j-1}))$ and $\Phi_{\tau-j-1}^i$, for $i \in \Gamma_h^*$ and $j = 0, \dots, T$. Indeed, since these terms are available from the solution of the reduced basis evolution at previous time steps, there is no cost associated with their assembly, at least at this stage of the proposed model order reduction algorithm. The matrix-matrix multiplications require $O(Np^*Tc)$, where p^* is the dimension of the subset of sampling parameters, T is the number of samples in the temporal window, and c is a mild constant that depends only on the support of the finite element basis functions. The truncated SVD decomposition of the matrix $\mathbf{Y} \in \mathbb{R}^{N \times p^*(T+1)}$ has arithmetic complexity $O(Np^*Td)$, where d is the number of DEIM modes. The computational cost required to assemble the interpolation matrix $P_\tau \in \mathbb{R}^{N \times d}$ using [9, Algorithm 1] only depends on d . This cost is further reduced by updating the indices according to the strategy described above and inspired by the adaptive sampling of [47]. Hence, the leading computational cost of the DEIM algorithm is $O(Np^*Td)$.

4.3 DMD-DEIM reduced dynamics and computational complexity

From (4.12), the Hamiltonian gradient $G_{\mathcal{H}}^p(U, Z_i)$ in (3.8) is approximated as

$$G_{\mathcal{H}}^{\text{DD}}(U, Z_i, t) := \nabla_{U, Z_i} \mathcal{H}_{U,i}^{\text{DD}}(Z_i(t), t) = \begin{pmatrix} m_p^{-1} \operatorname{diag}(\nabla \Lambda^0(X_r^i(t)) \Phi_{\text{DMD}}^i(t)) P_\tau (P_\tau^\top \Psi_\tau)^{-\top} \Psi_\tau^\top M_q \\ V_r^i(t) \end{pmatrix}, \quad (4.13)$$

for all $i \in S_p$, and $t \in \mathcal{T}_\tau$. Similarly, the approximation of the gradient of the reduced Hamiltonian $g_{\mathcal{H}}^p(U, Z_i)$ reads

$$\begin{aligned} g_{\mathcal{H}}^{\text{DD}}(U, Z_i, t) &= U(t)^\top G_{\mathcal{H}}^{\text{DD}}(U, Z_i, t) = \nabla_{Z_i} \mathcal{H}_{U,i}^{\text{DD}}(Z_i(t), t) \\ &= U_V(t)^\top V_r^i(t) + m_p^{-1} U_X(t)^\top \operatorname{diag}(\nabla \Lambda^0(X_r^i(t)) \Phi_{\text{DMD}}^i(t)) P_\tau (P_\tau^\top \Psi_\tau)^{-\top} \Psi_\tau^\top M_q, \end{aligned} \quad (4.14)$$

for all $i \in S_p$, and $t \in \mathcal{T}_\tau$.

The reduced dynamical system (3.10) is approximated by replacing the gradient of the reduced Hamiltonian $g_{\mathcal{H}}^p(U, Z_i) \in \mathbb{R}^{2N}$ with its DMD-DEIM approximation $g_{\mathcal{H}}^{\text{DD}}(U, Z_i, t)$ from (4.14) in the evolution equations of the expansion coefficients. The DMD-DEIM reduced dynamics reads

$$\begin{cases} \dot{Z}(t) = J_{2n} g_{\mathcal{H}}^{\text{DD}}(U, Z, t), & \text{in } \mathcal{T}, \end{cases} \quad (4.15a)$$

$$\begin{cases} \dot{U}(t) = (I_{2N} - UU^\top)(J_{2N} G_{\mathcal{H}}^p(U, Z) Z^\top - G_{\mathcal{H}}^p(U, Z) Z^\top J_{2n}^\top) S(Z)^{-1}, & \text{in } \mathcal{T}, \end{cases} \quad (4.15b)$$

$$\begin{cases} U(t_0) Z(t_0) = U^0 Z^0, \end{cases} \quad (4.15c)$$

Note that this approximate reduced model retains the geometric structure of the full model.

We analyze the computational cost to assemble and evaluate the right hand side of the DMD-DEIM reduced model (4.15) at each time instance. We then compare the results with the ones at the beginning of Section 4 corresponding to the reduced model. The evolution of the reduced basis requires $O(Nnp^*) + O(N_x^2p^*) + O(Np^*c) + O(Nn^2) + O(n^2p^*)$ flops owing to the parameter sampling discussed in Section 4.1. This cost also includes the assembly of the quantities $\Lambda^0(X_r^i)$ and Φ^i needed in the DMD and DEIM algorithms. The computational cost required to assemble and evaluate the velocity field of the flow characterizing the evolution of the coefficients reduces to the cost of the evaluation of the gradient (4.14) of the DMD-DEIM Hamiltonian at each time instant t . This includes:

- (1) The cost to compute the linear part $U_V^\top V_r$ of (4.14), for all instances of the parameter in Γ_h , is $O(Nn^2) + O(pn^2)$. Note that the cost $O(Nn^2)$ required to assemble the matrix $U_V^\top U_V$ is performed once per stage of the RK time integrator, while the matrix-matrix product $(U_V^\top U_V)Z$, at cost $O(pn^2)$, has to be performed, at every RK stage, and for each iteration of the nonlinear solver.
- (2) The cost to compute the DMD approximation of the electric potential $\Phi_{\text{DMD}}^i \in \mathbb{R}^{N_x}$, for any $i \in S_p$, is $O(N_x p^* Tr_\tau) + O(N_x r_\tau p)$, as shown in Section 4.2.1. This cost is linear in N_x , and does not depend on the number N of particles. The evaluation of Φ_{DMD}^i , that requires $O(N_x r_\tau p)$ flops, needs to be performed at each stage of the RK scheme and at each iteration of the nonlinear solver. The other cost $O(N_x p^* Tr_\tau)$ is accounted for once per time step.
- (3) The cost to run the DEIM algorithm is $O(Np^*Td) + O(Np^*Tc)$, as described in Section 4.2.2. The matrix-matrix multiplication $(P_\tau^\top \Psi_\tau)^{-\top} \Psi_\tau^\top M_q$ costs $O(Nd) + O(d^3)$. These operations are performed once per time step.
- (4) The computation of the nonlinear time-dependent part of (4.14) for all parameters $\eta_i \in \Gamma_h$, namely $U_X^\top \text{diag}(\nabla \Lambda^0(X_r^i(t)) \Phi_{\text{DMD}}^i(t)) P_\tau$, requires $O(pdn) + O(pdc)$ operations. This includes the cost of the matrix-matrix multiplications and the cost $O(pdc)$ to assemble $\nabla \Lambda^0(X_r^i)$ for d particles.

To summarize, we report, in the following table, the leading arithmetic complexity of the various steps required to solve the DMD-DEIM reduced system (4.15) in a fixed time interval $\mathcal{T}_\tau$. Here n_s denotes the number of stages of the RK time integrator and i_s the number of iteration of the nonlinear solver in the implicit timestepping, see also Section 5 for further details.

Leading arithmetic complexity	Operation
$O(Nnp^*n_s) + O(N_x^2p^*n_s) + O(Nn^2n_s)$	evolution of the reduced basis
$O(Nn^2n_s)$	computation of $U_V^\top U_V$ in (4.14)
$O(N_x r_\tau p^* T)$	construction of the DMD approximation as in Section 4.2.1
$O(Np^*Td) + O(Np^*Tc)$	construction of the DEIM approximation as in Section 4.2.2
$O(pn^2n_s i_s) + O(N_x r_\tau p n_s i_s)$	evolution of (4.14) for all parameters $\eta_i \in \Gamma_h$

The first two rows of the table report the operations required once per stage of the RK time integrator, while the third and forth rows refer to the operations required once per time interval. Note that these operations are shared by all parameters, resulting in a computational cost independent of the size p of the parameter set Γ_h , but only dependent on the number of parameter subsamples p^* . The last row reports the arithmetic complexity of the parameter-dependent computations. These need to be performed at each RK stage and iteration of the nonlinear solver, but their computational cost is independent of the number of particles N . Therefore, the DMD-DEIM approximation allows a complete separation of the costs involving the number N of particles and the number p of parameters, both potentially large.

5 Numerical temporal integration of the DMD-DEIM system

For the numerical time integration of the DMD-DEIM reduced dynamics (4.15), we adopt the partitioned RK method of order 2 proposed in [26]. The idea is to combine a symplectic temporal integrator for the

evolution (4.15a) of the expansion coefficient, with a time discretization of the basis evolution (4.15b) able to preserve the ortho-symplectic constraint. For the latter, we adopt the tangent method proposed in [45], summarized next.

In each temporal sub-interval $\mathcal{T}_\tau = (t_{\tau-1}, t_\tau]$, given the approximate reduced basis $U_{\tau-1}$, the method constructs a local retraction $\mathcal{R}_{\tau-1}$ from the tangent space $T_{U_{\tau-1}}\mathcal{U}$ into $\mathcal{U}$ so that $U(t) = \mathcal{R}_{\tau-1}(\xi(t))$ for some ξ in the tangent space at $U_{\tau-1}$. For the computation of U_τ , the idea is to evolve $\xi(t)$ in the tangent space and then recover U via the retraction. The reduced problem (4.15) in each temporal sub-interval $\mathcal{T}_\tau$ is re-written in terms of the variable $Z(t)$ and $\xi(t)$ as: Given $Z_{\tau-1}$ and $\xi_{\tau-1} = 0$, find $Z(t)$ and $\xi(t)$ such that

$$\begin{cases} \dot{Z}(t) = J_{2n} g_{\mathcal{H}}^{\text{DD}}(\mathcal{R}_{\tau-1}(\xi(t)), Z(t), t), & \text{in } \mathcal{T}_\tau, \\ \dot{\xi}(t) = \mathcal{Y}(\xi(t), Z(t)), & \text{in } \mathcal{T}_\tau. \end{cases} \quad (5.1a)$$

$$\quad (5.1b)$$

The velocity field $\mathcal{Y} : T_{U_{\tau-1}}\mathcal{U} \times \mathcal{Z} \rightarrow \mathbb{R}^{2N \times 2n}$, describing the local flow on the tangent space at $U_{\tau-1}$, is

$$\mathcal{Y}(\xi, Z) := -U_{\tau-1}(\mathcal{R}_{\tau-1}^\top(\xi)U_{\tau-1} + I_{2n})^{-1}(\mathcal{R}_{\tau-1}(\xi) + U_{\tau-1})^\top \Upsilon(\xi, Z) + \Upsilon(\xi, Z) - U_{\tau-1}\Upsilon^\top(\xi, Z)U_{\tau-1},$$

where $\Upsilon(\xi, Z)$ is given by

$$\Upsilon(\xi, Z) := (2\mathcal{X}(\mathcal{R}_{\tau-1}(\xi), Z) - (WU_{\tau-1}^\top - U_{\tau-1}W^\top)\mathcal{X}(\mathcal{R}_{\tau-1}(\xi), Z))(U_{\tau-1}^\top \mathcal{R}_{\tau-1}(\xi) + I_{2n})^{-1},$$

with $2W := (2I_{2N} - U_{\tau-1}U_{\tau-1}^\top)\xi$ and $\mathcal{X} : \mathbb{R}^{2N \times 2n} \times \mathcal{Z} \rightarrow \mathbb{R}^{2N \times 2n}$ being the velocity field of the approximate basis evolution in (4.15b), i.e.

$$\mathcal{X}(U, Z) := (I_{2N} - UU^\top)(J_{2N}G_{\mathcal{H}}^{p*}(U, Z)Z^\top - G_{\mathcal{H}}^{p*}(U, Z)Z^\top J_{2n}^\top)(ZZ^\top + J_{2n}^\top ZZ^\top J_{2n})^{-1}.$$

The retraction is defined according to [45, Proposition 5.6] as $\mathcal{R}_{\tau-1}(\xi) = \text{cay}(WU_{\tau-1}^\top - U_{\tau-1}W^\top)U_{\tau-1}$ where cay denotes the Cayley transform. We refer the reader to [45, 26] for further details regarding the formal derivation of (5.1). Note that, with the algorithm proposed in [45, Section 5.3.1], the computation of the retraction $\mathcal{R}$ and the assembly of the operator $\mathcal{Y}$ have arithmetic complexity $O(Nn^2)$.

The partitioned Runge–Kutta scheme applied to (5.1) reads

$$Z_\tau = Z_{\tau-1} + \Delta t \sum_{\ell=1}^{n_s} b_\ell k_\ell, \quad (5.2a)$$

$$\xi_\tau = \Delta t \sum_{\ell=1}^{n_s} \hat{b}_\ell \hat{k}_\ell, \quad U_\tau = \mathcal{R}_{U_{\tau-1}}(\xi_\tau), \quad (5.2b)$$

$$k_1 = J_{2n} g_{\mathcal{H}}^{\text{DD}}\left(U_{\tau-1}, Z_{\tau-1} + \Delta t \sum_{j=1}^{n_s} a_{1,j} k_j, t_{\tau-1}\right), \quad \hat{k}_1 = \mathcal{X}\left(U_{\tau-1}, Z_{\tau-1} + \Delta t \sum_{j=1}^{n_s} a_{1,j} k_j\right), \quad (5.2c)$$

$$k_\ell = J_{2n} g_{\mathcal{H}}^{\text{DD}}\left(\mathcal{R}_{U_{\tau-1}}\left(\Delta t \sum_{j=1}^{\ell-1} \hat{a}_{\ell,j} \hat{k}_j\right), Z_{\tau-1} + \Delta t \sum_{j=1}^{n_s} a_{\ell,j} k_j, t_{\tau-1} + c_\ell \Delta t\right), \quad \ell = 2, \dots, n_s, \quad (5.2d)$$

$$\hat{k}_\ell = \mathcal{Y}\left(\Delta t \sum_{j=1}^{\ell-1} \hat{a}_{\ell,j} \hat{k}_j, Z_{\tau-1} + \Delta t \sum_{j=1}^{n_s} a_{\ell,j} k_j\right), \quad \ell = 2, \dots, n_s, \quad (5.2e)$$

where $\{a_{\ell,j}, b_j, c_j\}$ and $\{\hat{a}_{\ell,j}, \hat{b}_j\}$ are the set of coefficients corresponding to the implicit midpoint rule and the explicit midpoint method, respectively, cf. [26, Appendix A]. Note that the Hamiltonian (2.9) is separable, i.e. the gradient of the electric energy determines the evolution of the state variable and the gradient of the kinetic energy describes the dynamics of the momentum. The separability of the full order Hamiltonian is, however, not inherited by the reduced model, as seen in (3.6). This precludes the explicit integration of (5.2a). Even though explicit numerical integrators for non-separable Hamiltonian, based on Hamiltonian extensions, have been recently proposed [63], further investigations are required to assess their accuracy in the framework of partitioned Runge–Kutta schemes.

For comparison purposes, in the numerical experiments, we will solve the full order model (2.10). The Störmer-Verlet scheme [20, Section I.1.4] is the most popular symplectic integrator for separable Hamiltonian systems and yields the following system of equations

$$X_\tau^i = X_{\tau-1}^i + \Delta t \left(V_{\tau-1}^i + \frac{\Delta t}{2} E_h(X_{\tau-1}^i; \eta_i) \right), \quad (5.3a)$$

$$V_\tau^i = V_{\tau-1}^i + \frac{\Delta t}{2} (E_h(X_{\tau-1}^i; \eta_i) + E_h(X_\tau^i; \eta_i)), \quad (5.3b)$$

to be solved for each of the p parameters $\eta_i \in \Gamma_h$. Here E_h denotes the approximate electric field (up to constants), i.e. $E_h(X_\tau^i; \eta_i) = -m_p^{-1} \text{diag}(M_q) \nabla \Lambda^0(X_\tau^i) L^{-1} \Lambda^0(X_\tau^i) M_q$.

We summarize the main steps of the proposed method in Algorithm 2.

Algorithm 2: DMD-DEIM dynamical model order reduction

Input : $f(0, x, v; \eta)$, $\{\eta_i\}_{i=1}^p$, Γ_h , p^* , T , k_{DEIM} .

- 1 Initialize particles' positions $\{X_0^i\}_{i=1}^p$ and velocities $\{V_0^i\}_{i=1}^p$ via inverse transform sampling from $f(0, x, v; \eta_i)$. Define $W_0(\eta_i) = [X_0^i; V_0^i]$;
 - 2 Compute U^0 and Z^0 via truncated complex SVD of $\mathcal{R}_0 = [W_0(\eta_1) | \dots | W_0(\eta_p)]$;
 - 3 Define the parameter subset $\Gamma_h^* \subset \Gamma_h$, of size $p^* \ll p$, and the indices subset S_{p^*} as in Section 4.1;
 - 4 Compute the potential Φ_0^i , for any $i \in S_{p^*}$, by solving the Poisson problem (2.7).;
 - 5 Initialize $\mathbf{Y}_{\text{DMD}}^i := [\Phi_0^i]$ and $\mathbf{Y}_{\text{DMD}}^{\prime, i} := []$. Initialize $\mathbf{Y}_{\text{DEIM}}^i := [\Lambda^0(X_r^i(t_0)) \Phi_0^i]$;
 - 6 **for** $\tau = 1, \dots, N_\tau$ **do**
 - 7 **if** $\tau \leq T$
 - 8 Solve the system (3.10a)-(3.10b) in $\mathcal{T}_\tau$ with the partitioned RK described in Section 5;
 - 9 For any $i \in S_{p^*}$, update $\mathbf{Y}_{\text{DMD}}^i = [\mathbf{Y}_{\text{DMD}}^i \quad \Phi_\tau^i]$ and $\mathbf{Y}_{\text{DMD}}^{\prime, i} = [\mathbf{Y}_{\text{DMD}}^{\prime, i} \quad \Phi_\tau^i]$;
 - 10 For any $i \in S_{p^*}$, update $\mathbf{Y}_{\text{DEIM}}^i := [\mathbf{Y}_{\text{DEIM}}^i \quad \Lambda^0(X_r^i(t_\tau)) \Phi_\tau^i]$.
 - 11 **elseif** $\tau > T$
 - 12 At each stage of the RK time integrator:
 - 13 Solve the evolution equation (4.15b) for the reduced basis as in (5.2b);
 - 14 For any $i \in S_{p^*}$, update $\mathbf{Y}_{\text{DMD}}^i = [\mathbf{Y}_{\text{DMD}}^i \quad \Phi_\tau^i]$ and $\mathbf{Y}_{\text{DMD}}^{\prime, i} = [\mathbf{Y}_{\text{DMD}}^{\prime, i} \quad \Phi_\tau^i]$;
 - 15 Remove the outdated terms at $\tau - T - 1$ from $\mathbf{Y}_{\text{DMD}}^i$ and at $\tau - T$ from $\mathbf{Y}_{\text{DMD}}^{\prime, i}$;
 - 16 For any $i \in S_{p^*}$, update $\mathbf{Y}_{\text{DEIM}}^i := [\mathbf{Y}_{\text{DEIM}}^i \quad \Lambda^0(X_r^i(t_\tau)) \Phi_\tau^i]$;
 - 17 Remove the outdated terms at $\tau - T - 1$;
 - 18 Compute Φ_{DMD} using (4.7) as described in Section 4.2.1;
 - 19 Approximate $\Lambda^0(X_r^i) \Phi_{\text{DMD}}^i$, for any $i \in S_p$, as described in Section 4.2.2;
 - 20 Build $g_{\mathcal{H}}^{\text{DD}}$ as in (4.14);
 - 21 Solve the evolution equation (4.15a) for the coefficients as in (5.2a).
 - 22 **end if**
 - 23 **end for**
-

6 Numerical experiments

6.1 Implementation and numerical study

In this section, we apply the proposed structure-preserving dynamical model order reduction approach to several periodic electrostatic benchmark problems. In all the examples, computational macro-particles are loaded from a perturbed initial distribution given by

$$f(0, x, v; \eta) = f_v(v; \eta) f_x(x; \eta), \quad (6.1)$$

where $f_v(v; \eta)$ is the initial velocity distribution, and $f_x(x; \eta) := 1 + \alpha \cos(kx)$ is the initial perturbation, with k as the wavenumber and α as the amplitude of the perturbation. We have chosen physical units such that the particle mass and particle charge are normalized to one for electrons, i.e., $q = -1$ and $m = 1$, and the weight w of the computational macro-particles is set to N^{-1} . To reduce the statistical

noise [27, 33] introduced by the particle discretization (2.3) of the initial condition (6.1), particles are loaded following a quiet start procedure based on a quasirandom sequence of samples. In detail, particles' positions and velocities are initialized by evaluating the inverse cumulative distribution function of $f(0, x, v; \eta)$ at the points defined by the Hammersley sequence [62] of length N . The distribution is defined over $\Omega := \Omega_x \times \Omega_v$, with $\Omega_v = [-10, 10]$ for all numerical experiments and Ω_x specified for each example. The quasirandom Hammersley sequence is characterized by a discrepancy value proportional to N^{-1} , whereas for a random distribution, the discrepancy is proportional to $N^{-1/2}$. Since the discrepancy measures the highest and lowest densities of points in a sequence, the Hammersley sequence guarantees that the particles are almost evenly distributed, and a significant noise reduction in the electrostatic field is therefore achieved. The p test parameters are sampled from the set Γ also using a quasirandom Hammersley sequence.

The DMD-DEIM reduced order model (4.15) is numerically integrated in time, in the interval $\mathcal{T} = (0, t_f]$, according to the scheme described in Section 5, resulting in the system of equations in (5.2). The full order model is solved using the Störmer-Verlet scheme (5.3). The same uniform time step Δt and number of time steps $N_\tau := t_f/\Delta t + 1$ are considered for the numerical integration of the two models. In the following, we adopt the notation $W_\tau^i := [X_\tau^i; V_\tau^i] \in \mathbb{R}^{2N}$ and $R_\tau^i := [X_{r,\tau}^i; V_{r,\tau}^i] = U_\tau Z_{i,\tau} \in \mathbb{R}^{2N}$ to denote the numerical solutions of the discrete full order model (5.3) and the discrete reduced order model (5.2) for the parameter η_i at time t_τ , respectively.

The reducibility of the considered benchmark tests is studied in terms of the decay of the singular values of the snapshots matrices

$$S_X = \begin{bmatrix} X_0^1 & \cdots & X_0^p & \cdots & X_{N_\tau}^1 & \cdots & X_{N_\tau}^p \end{bmatrix} \quad \text{and} \quad S_V = \begin{bmatrix} V_0^1 & \cdots & V_0^p & \cdots & V_{N_\tau}^1 & \cdots & V_{N_\tau}^p \end{bmatrix}, \quad (6.2)$$

collecting the position and velocity components of W_τ^i . As in traditional reduced basis methods, the space spanned by the selected snapshots is assumed to be representative of the solution set. The behavior of the singular values of S_X and S_V is also compared to the decay of the singular values of the *local* snapshots matrices

$$S_X^\tau = \begin{bmatrix} X_\tau^1 & \cdots & X_\tau^p \end{bmatrix} \quad \text{and} \quad S_V^\tau = \begin{bmatrix} V_\tau^1 & \cdots & V_\tau^p \end{bmatrix}, \quad \forall \tau = 1, \dots, N_\tau, \quad (6.3)$$

to assess the applicability and the benefits of the dynamical approach over standard global reduction methods. For the local snapshots matrices, we compute the ordered singular values $\{\sigma_{X,j}^\tau\}_j$ of S_X^τ , for each $\tau = 1, \dots, N_\tau$, and normalize them with respect to the maximum singular value $\sigma_{X,1}^\tau$. Then, for each j , we consider the average over time, i.e., $\sum_\tau \sigma_{X,j}^\tau$, and the maximum over time, i.e., $\max_\tau \sigma_{X,j}^\tau$. The same study is carried out for the matrix S_V^τ . A further indicator of the reducibility properties of the problem is given by the numerical rank of S_X^τ and S_V^τ , defined as the number of singular values larger than a user-defined threshold tolerance. In the following, different tolerances are considered.

The accuracy of the DMD-DEIM-ROM is evaluated by computing, for each $\tau = 1, \dots, N_\tau$, the relative errors

$$\varepsilon_{\text{rel},X}(t_\tau) = \frac{\|S_X^\tau - X_{r,\tau}\|_F}{\|S_X^\tau\|_F}, \quad \text{and} \quad \varepsilon_{\text{rel},V}(t_\tau) = \frac{\|S_V^\tau - V_{r,\tau}\|_F}{\|S_V^\tau\|_F}, \quad (6.4)$$

where $X_{r,\tau}, V_{r,\tau} \in \mathbb{R}^{N \times p}$ are the position and velocity components of the discrete reduced order solution $R_\tau = U_\tau Z_\tau \in \mathbb{R}^{2N \times p}$, respectively. We study the error in the position and velocity of the particles separately because they are characterized by different scales of absolute error. The relative errors (6.4) are compared to the target values given by the projection errors

$$\varepsilon_{\text{rel},X}^{\text{Target}}(t_\tau) = \frac{\|S_X^\tau - S_{X,\text{cSVD}}^\tau\|_F}{\|S_X^\tau\|_F}, \quad \text{and} \quad \varepsilon_{\text{rel},V}^{\text{Target}}(t_\tau) = \frac{\|S_V^\tau - S_{V,\text{cSVD}}^\tau\|_F}{\|S_V^\tau\|_F}, \quad (6.5)$$

where $S_{X,\text{cSVD}}^\tau, S_{V,\text{cSVD}}^\tau \in \mathbb{R}^{N \times p}$ are the position and velocity components of the projection of the snapshots onto the space spanned by the ortho-symplectic basis of size $2N \times 2n$ obtained from the Complex SVD [48] of the matrix $S_X^\tau + \iota S_V^\tau$. Since the Complex SVD provides the ortho-symplectic basis that minimizes the projection error in the snapshots [48, Theorem 4.6], comparing (6.4) and (6.5) allows to test the approximability properties of the reduced basis constructed in the dynamical approach.

Moreover, we analyze the evolution of the electric field energy (2.8) for the reduced order approximation, i.e., $\mathcal{E}(X_{r,\tau}^i; \eta_i)$, and for the full order solution, i.e., $\mathcal{E}(X_{\tau}^i; \eta_i)$, for each instance η_i of the parameter. This term gives information of the macroscopic behavior of the plasma, and it is also the one affected by more levels of approximation.

Finally, the efficiency of the proposed approach is investigated by comparing the running times required for the integration over a single time step of the fully-discrete DMD-DEIM reduced model (5.2) and of the discrete full order model (5.3). The running time for the full model is obtained by summing the times required for each instance of the parameter $\eta_i \in \Gamma_h$. The comparison focuses on the scalability in the approximation of parametric problems as the size p of the parameter set increases, a typical scenario in a multi-query context. To compare the efficiency of the different methods, we analyze the runtime required for integration over a single time interval for all parameter values considered. The values reported were obtained as the average of the runtimes obtained in the first 25 time intervals. For the dynamical reduced basis method, we also analyze separately the contributions to the computational cost due to the basis evolution (5.2b)-(5.2e) and to the coefficients evolution (5.2a)-(5.2c)-(5.2d), in line with the theoretical findings of Section 4.3.

In the construction of the DMD-DEIM reduced model, we consider a tolerance equal to 10^{-5} in the computation of the DMD modes in Algorithm 1. Moreover, the nonlinear system, (5.2c) and (5.2d), describing the evolution of the increments k_ℓ , $\ell = 1, \dots, n_s$, is solved using the fixed point iteration method. As a stopping criterion for the nonlinear solver, we check when the relative norm of the update to k_ℓ is smaller than the threshold value 10^{-9} . All numerical simulations are performed using Matlab on computer nodes with Intel Xeon E5-2643 (3.40GHz). The code and the data supporting the findings of this study are available upon request.

6.2 Weak Landau damping of 1D Langmuir waves

The first application we consider is the study of the damped propagation of small amplitude plasma waves, also known as Landau damping (LD). The resonance between physical particles and the propagating wave generates damping of the electric field energy, without particle collisions. This process is used in particle accelerators to prevent coherent beams oscillations that could cause potential instabilities [24]. The initial condition is given by (6.1) with the velocity distribution function

$$f_v(v; \eta) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{v^2}{2\sigma^2}\right), \quad (6.6)$$

where the amplitude of the perturbation α and the standard deviation σ of the velocity Maxwellian are the study parameters $\eta = (\alpha, \sigma)$, with $\eta \in \Gamma = [0.03, 0.06] \times [0.8, 1]$, and the perturbation wavenumber k is fixed to 0.5. Following the sampling procedure described in Section 6.1, we solve the Landau damping problem for $p = 300$ different realizations of the parameter η . In Figure 1(a) and 1(b) the initial position and velocity distributions are shown for several of the selected parameter values.




Figure 1: LD: (a) – (b) Initial position and velocity distributions for selected values of the parameter in Γ_h . (c) Exponential time decay of the electrostatic energy $\mathcal{E}(X_{\tau}^i; \eta_i)$ obtained from the full model solution, for selected values of η_i in Γ_h . Since not all parameters in Γ_h are reported, the black lines in each subplot are used to mark the region where the plotted quantity is contained, for any value of the parameter in Γ_h .

We consider periodic boundary conditions on the physical space domain $\Omega_x := (0, \frac{2\pi}{k})$ with a uniform neutralizing background charge. For the numerical solution of the full order model, we use $N_x = 32$ piecewise linear basis for the Poisson solver, and $N = 5 \times 10^4$ macro-particles for the approximation of the solution density. A uniform time step $\Delta t = 0.0025$ has been adopted for the evolution of particles' positions and velocities over the time interval $\mathcal{T} = (0, 20]$.

In Figure 2, the decay of the singular values of the global snapshots matrices S_X and S_V , normalized with respect to the corresponding largest singular value, are compared to the maximum and averages over τ of the normalized singular values of the local counterparts S_X^τ and S_V^τ , computed as described in Section 6.1. Concerning the particles position, although a plateau of the singular values can be seen for both global and local matrix, the initial decay is sharper in the local case with singular values that are two orders of magnitude smaller than in the global case, suggesting a more efficient representation using a local low-rank model. This gap increases when considering the particles velocity, suggesting that a *global* reduced basis approach would not be effective in reducing the computational cost of the Landau damping simulation.




Figure 2: LD: Singular values of the global snapshots matrices S_X and S_V compared to the maximum and time average (in τ) of the singular values of the local matrices S_X^τ and S_V^τ .

In Figure 3, we report the numerical rank of the matrices S_X^τ and S_V^τ as a function of τ and for different values of the threshold. The numerical rank remains constant and below 4 for tolerances larger than 10^{-3} , grows to 8 for a tolerance of 10^{-4} , and reaches a maximum of 13, for positions and 32 for velocities, when the tolerance is set to 10^{-5} . The increase in the solution complexity over time is partially due to the accumulation of statistical noise associated with the discretization of $f(t, x, v; \eta)$ by macro-particles. In Figure 4, in support of this conclusion, we note that as the average number of particles per cell increases during the initial particle loading phase, the numerical ranks of S_X^τ and S_V^τ , at fixed tolerance, decrease.




Figure 3: LD: Numerical rank of S_X^τ in (a) and S_V^τ in (b), as a function of τ . Different colors are associated with different values of the threshold, according to the legend.




Figure 4: LD: Evolution of the numerical ranks of S_X^τ (a) – (c) and of S_V^τ (d) – (f) for different threshold indicated by Tol. In each subfigure, the rank behavior for different values of N_x and p , in the discretization of the full order model, is compared.

This behavior of the numerical rank suggests that evolving the basis, but keeping the rank of the approximation constant, is sufficient to accurately approximate the solution of the full order model, at least in this test case.

Concerning the reduced dynamical model, we consider $2n = 4$ as the reduced manifold dimension. For the DEIM reduction described in Section 4.2.2, $d = 32$ interpolation points have been used to reduce the approximation error, and $n_{\text{DEIM}} = 12$ DEIM indices are updated at each time step, for the sake of efficiency, according to (4.11). All DEIM indices are recomputed every $k_{\text{DEIM}} = 3$ time steps using [9, Algorithm 1]. We refer to Section 4.2.2 for further details.

For this test case, we include a numerical study of the evolution of the approximation errors $\varepsilon_{\text{rel},X}$ and $\varepsilon_{\text{rel},V}$ in (6.4) under variations of the size p^* of the subset of the parameters used to evolve the basis effectively, according to Section 4.1, and the length $(T+1)\Delta t$ of the time window adopted to harvest the self-consistent electric potential $\Phi(X_r^i)$ for the DMD extrapolation step, as described in Section 4.2.1. In particular, we consider $p^* \in \{8, 12, 16\}$ and $T \in \{3, 5\}$ and the results are shown in Figure 5. In all tested combinations of T and p^* , the error is proportional to the best approximation error, both in position and velocity.




Figure 5: LD: Evolution of the position (a) and velocity (b) relative errors, as defined in (6.4), for different choices of p^* and T . These errors are compared to the target values given by the position component $\varepsilon_{\text{rel},X}^{\text{Target}}$ and the velocity component $\varepsilon_{\text{rel},V}^{\text{Target}}$ of the relative projection errors defined in (6.5). The target reduced basis has dimension 4 and is computed, for each time step, using the Complex SVD algorithm, as described in Section 6.1.

As p^* increases, both errors decrease: this is expected since a more refined sampling of the subset Γ_h^* results in a more accurate representation of the dynamics of evolution of the bases in (4.15b). We also note that, for $p^* = 16$, the error $\varepsilon_{\text{rel},V}$ is, at several time instances, smaller than the target value. This performance can be explained by the fact that the optimality of the Complex SVD algorithm concerns the projection of the entire state $[S_X^T S_V^T]$ and not of its components individually. We observe that the DMD number of samples, $T + 1$, has no impact on the error when p^* is large, and a small accuracy degradation is even registered for $p^* = 8$ when $T = 5$ is chosen over $T = 3$. The optimal choice of T remains an open problem: as pointed out in [12], it should capture slow and fast scales of the local dynamics, but a rigorous optimization strategy would require a study of the multi-scale properties of the solution to the Vlasov–Poisson equation for each of the parameter realization considered. However, we stress that the results are relatively robust concerning this parameter.

We remark that in this test case the time step Δt is rather small, $\Delta t = 0.0025$. The rationale is that a small time step allows us to gauge the error introduced by the reduction and hyper-reduction without pollution from other sources of errors, such as the temporal discretization. However, larger time steps can be chosen without a significant decrease of accuracy, as shown in Figure 6. There we report the evolution of the relative errors $\varepsilon_{\text{rel},X}$ and $\varepsilon_{\text{rel},V}$ for values of the time step $\Delta t \in \{0.0025, 0.01, 0.04, 0.08, 0.16, 0.26\}$ in the case $T = 3$, $p^* = 12$, and $k_{\text{DEIM}} = 1$. It can be observed that the timestepping is not the dominant source of error up to the value $\Delta t = 0.04$, and that, for $\Delta t > 0.04$, increasing the time step does not significantly affect accuracy.




Figure 6: LD: Evolution of the position (a) and velocity (b) relative errors, as defined in (6.4), for different values of the time step Δt . The hyperparameters are $p^* = 12$, $T = 3$, and $n_{\text{DEIM}} = 1$.

Landau theory [4] establishes that, for small perturbations of the initial analytical data of the form (6.6), the electric energy $\mathcal{E}(X^i; \eta_i)$ decays (in time) exponentially with a damping factor that depends on the standard deviation σ of the Maxwellian distribution f_v but is independent of the amplitude of the perturbation α . In Figure 7 we report the damping rate of the electric energy for each value of the parameter $\eta = (\alpha, \sigma) \in \Gamma$: for the p test parameters the damping rate is computed as the slope of the peaks of the electric energy (excluding the first one); all other values are generated via linear interpolation/extrapolation on a uniform grid of 90000 points. As can be seen from Figure 7, the theoretical dependence of the damping rate on the considered parameters is captured by the reduced model numerical solution.




Figure 7: LD: Damping rates of the exponential time decay of the electric energy $\mathcal{E}(X^i; \eta)$, defined in (2.8), as a function of the two-dimensional parameter $\eta = (\alpha, \sigma)$. The plots refer to (a) the full order model; (b) the dynamical reduced model with $T = 3$ and $p^* = 16$; and (c) the dynamical reduced model with $T = 5$ and $p^* = 8$.

In Figure 8(a), we report the evolution of the relative error of the Hamiltonian (2.9) computed in the reduced and full model solutions, i.e.

$$\frac{\|\mathcal{H}([X_\tau, V_\tau]) - \mathcal{H}(U_\tau Z_\tau)\|_2}{\|\mathcal{H}([X_\tau, V_\tau])\|_2}, \quad \forall \tau = 1, \dots, N_\tau. \quad (6.7)$$

It is observed that the error is bounded and grows only slowly over time. The reason why the Hamiltonian is not exactly preserved is twofold: the numerical temporal integrator is symplectic but not Hamiltonian-preserving, and the reduced model possesses a Hamiltonian structure but with an approximate Hamiltonian function. As pointed out in Section 4, applying DEIM directly to the nonlinear Hamiltonian vector field yields an approximation that, in general, does not possess a gradient structure and hence fails to provide a Hamiltonian dynamics. To better understand the aforementioned two sources of errors, we consider the error in the Hamiltonian at two consecutive time instances of the solution. For the DMD-DEIM reduced model discretized with the partitioned Runge-Kutta method described in Section 5, it holds

$$\Delta \mathcal{H}_{\tau-1 \rightarrow \tau} := \|\mathcal{H}(U_\tau Z_\tau) - \mathcal{H}(U_{\tau-1} Z_{\tau-1})\|_2 \leq \|\mathcal{H}(U_\tau Z_\tau) - \mathcal{H}(U_{\tau-\frac{1}{2}} Z_\tau)\|_2 \quad (6.8a)$$

$$+ \|\mathcal{H}(U_{\tau-\frac{1}{2}} Z_{\tau-1}) - \mathcal{H}(U_{\tau-1} Z_{\tau-1})\|_2 \quad (6.8b)$$

$$+ \|\mathcal{H}(U_{\tau-\frac{1}{2}} Z_\tau) - \mathcal{H}(U_{\tau-\frac{1}{2}} Z_{\tau-1})\|_2. \quad (6.8c)$$

The first two terms (6.8a) and (6.8b) depend on the numerical time integration of the basis equation (5.2b), while the last term (6.8c) also depends on the DMD-DEIM approximation of the Hamiltonian. In particular, it holds

$$\Delta \mathcal{H}_{\tau-1 \rightarrow \tau}^Z := \|\mathcal{H}(U_{\tau-\frac{1}{2}} Z_\tau) - \mathcal{H}(U_{\tau-\frac{1}{2}} Z_{\tau-1})\|_2 \leq \|\mathcal{H}(U_{\tau-\frac{1}{2}} Z_\tau) - \mathcal{H}_{U_{\tau-\frac{1}{2}}}^{\text{DD}}(Z_\tau, t_\tau)\|_2 \quad (6.9a)$$

$$+ \|\mathcal{H}_{U_{\tau-\frac{1}{2}}}^{\text{DD}}(Z_{\tau-1}, t_{\tau-1}) - \mathcal{H}(U_{\tau-\frac{1}{2}} Z_{\tau-1})\|_2, \quad (6.9b)$$

$$+ \|\mathcal{H}_{U_{\tau-\frac{1}{2}}}^{\text{DD}}(Z_\tau, t_\tau) - \mathcal{H}_{U_{\tau-\frac{1}{2}}}^{\text{DD}}(Z_{\tau-1}, t_{\tau-1})\|_2, \quad (6.9c)$$

where $\mathcal{H}_U^{\text{DD}}$ is defined in (4.12). The first two terms (6.9a) and (6.9b) depend on the approximation of the Hamiltonian introduced by the DMD-DEIM method, while the last term (6.9c), that we dub $\Delta\mathcal{H}_{\tau-1\rightarrow\tau}^{\text{Z,DD}}$, is only associated with the numerical time integrator of the coefficient equation. In Figure 8(b) we report the time evolution of $\Delta\mathcal{H}_{\tau-1\rightarrow\tau}$, $\Delta\mathcal{H}_{\tau-1\rightarrow\tau}^{\text{Z}}$ and $\Delta\mathcal{H}_{\tau-1\rightarrow\tau}^{\text{Z,DD}}$, for the hyper-parameters $T = 3$ and $p^* = 12$. We can observe that the DMD-DEIM method provides a good approximation of the Hamiltonian (dashed line).




Figure 8: LD: (a) Evolution of the relative error (6.7) of the Hamiltonian. (b) Evolution of the components $\Delta\mathcal{H}_{\tau-1\rightarrow\tau}$, $\Delta\mathcal{H}_{\tau-1\rightarrow\tau}^{\text{Z}}$ and $\Delta\mathcal{H}_{\tau-1\rightarrow\tau}^{\text{Z,DD}}$ of the error bound (6.8), (6.9) in the local conservation of the reduced Hamiltonian. The values of the hyper-parameters are set to $p^* = 12$ and $T = 3$, respectively.

To study the algorithm efficiency, we investigate the runtime as a function of the number p of tested parameters. The proposed approach outperforms the full order solver, as shown in Figure 9(a), and the gap widens as the value of p increases. Depending on the choice of hyper-parameters of the reduced model, the algorithm speed-up varies between 1.9 and 3.3 when $p = 30$ and between 46 and 71 when $p = 1000$. For $p \geq 2000$, the evolution of the expansion coefficients (4.15a) becomes computationally more demanding than the evolution of the reduced basis (4.15b), as shown in Figure 9(b), and the overall computational cost of the reduced model begins to grow approximately linearly as a function of p . Thus, for values of p larger than 2000, the ratio between the time required to integrate the full model and the time to integrate the reduced model remains constant, with speed-ups ranging between 141 and 183, depending on the values of the hyper-parameters. We also remark that the computational cost to evolve the reduced basis (dashed lines in Figure 9(b)) is independent of p because it only depends on the number p^* of subsamples.




Figure 9: LD: (a) Comparison of the runtime per time step (in seconds) between the full order solver and the dynamical reduced basis approach for different hyper-parameter configurations, as function of the parameter sample size p . (b) Separation of contributions to the runtime per time step of the reduced model due to basis evolution (5.2b) (dashed lines) and coefficients evolution (5.2a) (continuous line).

The choice and the number p^* of subsamples has also an effect on the accuracy of the reduced order approximation. In the following test, we numerically study the convergence of the algorithm for different numbers of sampling points and for different parameter domains. In particular, we compare the performances of the algorithm when the parameter $\eta = (\alpha, \sigma)$ belongs to the domain $\Gamma \in \{\Gamma_1, \Gamma_2\}$, where $\Gamma_2 = [0.02, 0.07] \times [0.7, 1.1]$ and $\Gamma_1 = [0.03, 0.06] \times [0.8, 1]$ is the domain considered above. Although the length of the 1D parameter intervals used to define Γ_1 and Γ_2 might not appear very different, what matters for the numerical approximation is the reducibility of the solution as a function of the parameter, that is the approximability of the parameter-to-solution map. To gauge the reducibility of the solution under variation of the parameter we perform a qualitative (Figure 10) and a quantitative assessment (Figure 11). More in details, analogously to Section 6.2, we consider $p = 300$ parameter samples in the domain Γ . In Figure 10 we report the initial distribution of position (plot (a)) and velocity (plot (b)) and the time evolution of the electrostatic energy (plot (c)) for different samples of the parameters in Γ_2 . Comparing with Figure 2, where $\Gamma = \Gamma_1$, we observe a greater variability of all quantities, with a variation of the maximum and minimum values of the electrostatic energy of one order of magnitude at time $t = 20$.




Figure 10: LD: (a) – (b) Initial position and velocity distributions for selected values of the parameter η_i in $\Gamma_{2,h}$. (c) Exponential time decay of the electrostatic energy $\mathcal{E}(X_\tau^i; \eta_i)$ obtained from the full model solution, for selected values of η_i in $\Gamma_{2,h}$.

The increased complexity is also certified by the decay of the singular values of the global and local snapshots matrix, as shown in Figure 11: the singular values decay more slowly for $\Gamma = \Gamma_2$ compared to the case $\Gamma = \Gamma_1$, see Figure 2. In particular, looking at Figure 11(b) (zoomed box), we can see that, for the local snapshots matrix, the fourth singular value reaches the value of around 10^{-4} for $\Gamma = \Gamma_1$, while it is equal to 10^{-2} for $\Gamma = \Gamma_2$.




Figure 11: LD: Singular values of the global snapshots matrices S_X and S_V compared to the singular values of the local matrices S_X^τ and S_V^τ for $\Gamma = \Gamma_2$.

We study the evolution of the error for the two intervals of parameters as a function of the number of sampling points p^* . We select the size $2n$ of the reduced space so that the errors when $p = p^*$ are comparable in magnitude in the two cases: This entails taking $n = 4$ for $\Gamma = \Gamma_1$ and $n = 8$ for $\Gamma = \Gamma_2$.

The fact that Γ_2 requires a larger reduced space to achieve an accuracy comparable to the one obtained for Γ_1 is expected, given the lower reducibility of the problem (see Figure 11). All hyper-parameters (i.e. N , d , n_{DEIM} , k_{DEIM} , etc.) are set as above, $T = 5$ and $\Delta t = 0.0025$.

We report the errors in the particles' position and velocity in Figure 12 for the case $\Gamma = \Gamma_1$ and in Figure 13 for the case $\Gamma = \Gamma_2$. The black dashed line represents the target error: since for this study storing all required quantities for $p = p^*$ is demanding on computer memory we consider $p^* = 56$ as reference value. No difference in the error was recorded for values of p^* greater than 56. Notice also that $p^* \geq n$ is necessary condition for the invertibility of $S(Z) = ZZ^\top + J_{2n}^\top ZZ^\top J_{2n}$.

It can be observed that, as expected, the error improves as the number of sampling points p^* increases. However, for values larger than a certain p^* , approximately $p^* = 20$ for $\Gamma = \Gamma_1$ and $p^* = 28$ for $\Gamma = \Gamma_2$, the error does not improve any longer. The reason for this behavior is that, because of the local low-rank structure of the solution, only a subset of the parameters is significant to describe the solution-to-parameter map at a given time and the proposed algorithm is able to select those parameters. Moreover, for values larger than a certain p^* , the error introduced by the other approximations is dominating.

Finally, by comparing the behavior of the errors for the two parameter ranges Γ_1 and Γ_2 , it can be observed that the number p^* of sampling points required to achieve a certain accuracy is, as expected, larger for larger parameter domains (more precisely for less reducible solutions). It can also be observed that not all parameters p are required to achieve the target error, suggesting that the algorithm reaches the fixed accuracy at a reduced computational cost compared to evolving the reduced basis for $p^* = p$.




Figure 12: LD: Evolution of the relative errors in position (a) and velocity (b) for different choices of p^* and $n = 4$ for the parameter interval $\Gamma_1 = [0.03, 0.06] \times [0.8, 1]$.




Figure 13: LD: Evolution of the relative errors in position (a) and velocity (b) for different choices of p^* and $n = 8$ for the parameter interval $\Gamma_2 = [0.02, 0.07] \times [0.7, 1.1]$.

6.3 Nonlinear Landau damping of 1D Langmuir waves

For larger initial perturbation amplitudes, the linear theory does not hold and, after an initial shearing in phase space, leading to Landau damping, the damping is halted, and strong particle-trapping vortices are formed, leading to a growth of the potential energy of the system [38]. To simulate this scenario, starting from the same initial condition (6.6) and periodic domain $\Omega_x := (0, \frac{2\pi}{k})$ of the previous test, we take the parameter $\eta = (\alpha, \sigma)$ in the domain $\Gamma = [0.46, 0.5] \times [0.96, 1]$ and consider $p = 300$ different realizations. In Figure 14, we report the behavior of the initial velocity and position distributions along with the evolution in time of the electric field energy.

The full order simulations are conducted using $N_x = 64$ degrees of freedom for the discretization of the Laplacian operator, and $N = 10^5$ particles for the approximation of the distribution function. We consider the time interval $\mathcal{T} = (0, 40]$, with $\Delta t = 0.002$ and the numerical time integrators described in Section 5.




Figure 14: NLD: (a) – (b) Initial position and velocity distributions for selected values of the parameter in Γ_h . (c) Exponential time decay of the electrostatic energy $\mathcal{E}(X_\tau^i; \eta_i)$ obtained from the full model solution, for selected values of η_i in Γ_h . Since not all parameters in Γ_h are reported, the black lines in each subplot are used to mark the region where the plotted quantity is contained, for any value of the parameter in Γ_h .

The decay of the singular values of the global and local snapshot matrices, defined in (6.3) and (6.2), is shown in Figure 15). Compared to the linear Landau damping, the nonlinear test case is unsuitable for reduction with a global reduced basis approach both in terms of particles position and velocity. Regarding reducibility via a local basis in time, we note that although the problem is more challenging than the weak Landau damping, the normalized singular values of S_X^τ and S_V^τ reach $3.9 \cdot 10^{-4}$ and 2.2×10^{-3} , respectively, at the sixth singular value, making the problem amenable to local reduction.




Figure 15: NLD: Singular values of the global snapshots matrices S_X and S_V compared to the maximum and time average (in τ) of the singular values of the local matrices S_X^τ and S_V^τ .

Similar conclusions are drawn from the behavior of the numerical rank, shown in Figure (16) as a function of time, from which we also note that the problem becomes significantly more complex in the final part of the time interval considered, corresponding to the formation of the particle attractor vortices and as the nonlinear contribution to the dynamics of the particles becomes dominant.




Figure 16: NLD: Numerical rank of S_X^τ in (a) and S_V^τ in (b), as a function of τ . Different colors are associated with different values of the threshold, according to the legend.

To reduce this test problem, we consider a symplectic dynamical basis of dimension $2n = 6$ and the same number $d = 32$ of DEIM interpolatory indices as used for the weak Landau damping. In addition, a subset of $p^* = 8$ parameters, taken according to Section (4.1), is considered for the efficient evolution of the basis. The relative errors for the different choices of the $T + 1$ DMD samples and frequency k_{DEIM} are shown in Figure 17: the error does not deteriorate over time for any of the chosen hyper-parameters, and the increase of k_{DEIM} only marginally impacts the performances of the reduced model.




Figure 17: NLD: Evolution of the position (a) and velocity (b) relative errors, as defined in (6.4), for different choices of p^* and T . These errors are compared to the target values given by the position component $\varepsilon_{\text{rel},X}^{\text{Target}}$ and the velocity component $\varepsilon_{\text{rel},V}^{\text{Target}}$ of the relative projection errors defined in (6.5). The target reduced basis has dimension 6 and is computed, for each time step, using the Complex SVD algorithm, as discussed in Section 6.1.

In Figure 18, we plot the distribution function $f_h(t, x, v; \eta)$ reconstructed from the macro-particles for the parameter $\eta = (0.4912, 0.9889)$. The numerical solution of the approximate reduced model is in good agreement with the full model solution, and the various dynamical stages, from the initial shearing to the development of the two particle-trapping vortices, are correctly captured. Furthermore, although tiny artifacts in the vortex structure can be observed in the case of hyper-parameters $T = 3$ and $k_{\text{DEIM}} = 3$ at $t = 40$, this is not the case for the choice $T = 5$ and $k_{\text{DEIM}} = 1$.




Figure 18: NLD: Numerical distribution function for $\eta = (0.4912, 0.9889)$ at different times obtained from (a) the full order model; (b) the dynamical reduced model with $T = 3$ and $k_{\text{DEIM}} = 3$; and (c) the dynamical reduced model with $T = 5$ and $k_{\text{DEIM}} = 1$. Starting from the perturbed Maxwellian distribution, particles with different energies oscillate with different frequencies leading to the typical filamentation that starts developing at $t = 13.33$. Two trapping vortices, centered at opposite phase velocities, form at $t = 26.66$ and fully develop at $t = 40$.

To better understand the macroscopic effects of the order reduction on the numerical solution, we consider, in Figure 19, the evolution of the electric field energy for different realizations of the parameter. The reduced model solution gives accurate results, both in terms of the amplitude and frequency of the peaks. As a further analysis, in Figure 21, we report the exponential damping rate of $\mathcal{E}(X_{r,\tau}^i; \eta_i)$, which is

obtained during the initial phase of Landau damping, and the exponential growth rate of $\mathcal{E}(X_{r,\tau}^i; \eta_i)$ that characterizes the subsequent formation of particle-trapping vortices in phase space. For $\eta_i = (0.5, 1)$, the values obtained are around -0.287 and 0.078 , which is in agreement with the literature, *cf.* for example [32, Table 3].

As a general remark, from the analysis of Figures 19 and 21, we can see that the reduced model solution slightly underestimates the growth rate of the electric field energy for values of the parameter σ greater than 0.98.

We show in Figure 20 the peaks of $\mathcal{E}(X_{r,\tau}^i)$ that have been fitted for the calculations of the damping and growth rates.




Figure 19: NLD: Evolution of the electric field energy $\mathcal{E}(\cdot; \eta_i)$. The energy is evaluated at the positions X_{τ}^i computed using the high-fidelity solver and at the positions $X_{r,\tau}^i$ computed using the reduced model, for different values η_i of the parameter.




Figure 20: NLD: Peaks of the electric field energy $\mathcal{E}(X_{r,\tau}^i; \eta_i)$ selected for the computation of the exponential damping and growth rates.




Figure 21: NLD: Contour plots of the damping rate ((a), (c), (e)) and growth rate ((b), (d), (f)) of the electric field energy $\mathcal{E}(X_{r,\tau}^i; \eta_i)$ for the solution of the reduced model, with different values of k_{DEIM} and T , and for the solution obtained with the high-fidelity solver

In Figure 22, we compare the running times of the full order solver and the reduced order solver. For this numerical simulation, the choice of the hyper-parameters T and k_{DEIM} has a mild impact on the computational cost required to advance the reduced state of a single time step. Once the cost to integrate the evolution of the coefficients has exceeded the cost to integrate the basis equation, the most computationally expensive choice of hyper-parameters (i.e., $T = 5$ and $k_{DEIM} = 1$) is only around 1.35 times more demanding than the computationally cheapest choice (i.e., $T = 3$ and $k_{DEIM} = 3$). This result is in agreement with the analysis of the arithmetic complexity of the reduction algorithm presented in Section 4.3: for large p , the dominant cost has order $O(pn^2) + O(N_x r_\tau p) + O(pdn) + O(pdc)$, which depends on the hyper-parameters only via the number r_τ of retained DMD eigenvalues. Although the value r_τ might be different for different choices of the the number T of DMD samples, this does not significantly affect the computational cost of the algorithm.




Figure 22: NLD: (a) Comparison of the runtime per time step (in seconds) between the full order solver and the dynamical reduced basis approach for different hyper-parameter configurations, as a function of the parameter sample size p . (b) Separation of contributions to the runtime per time step of the reduced model due to basis evolution (5.2b) (dashed lines) and coefficient evolution (5.2a) (continuous line).

6.4 Two-stream instability

The two-stream instability is a well-known instability in plasma physics generated by two counterstreaming beams, where the kinetic energy of particles excites a plasma wave and, consequently, transfers to electric potential energy [2]. In this study, we focus on the temporal interval that includes the first two stages in which the evolution of electric field energy is distinct, namely the initial, short transient stage, and the subsequent growth stage. For the latter stage, the dynamics is defined by the interplay between harmonics characterized by different growth rates. We consider the spatial domain $\Omega_x := (0, \frac{2\pi}{k})$ with periodic boundary conditions. The initial velocity distribution is given by

$$f_v(v; \eta) = \frac{1}{2\sqrt{2\pi}\sigma} \exp\left(-\frac{(v-v_0)^2}{2\sigma^2}\right) + \frac{1}{2\sqrt{2\pi}\sigma} \exp\left(-\frac{(v+v_0)^2}{2\sigma^2}\right), \quad (6.10)$$

where $v_0 = 3$ is the initial velocity displacement in phase space. The wavenumber k of the perturbation is set to 0.2, and the parameter $\eta = (\alpha, \sigma)$ varies in the domain $\Gamma = [0.009, 0.011] \times [0.98, 1.02]$ discretized using $p = 300$ samples. Figures 23(a)-(b) show the initial parametric distributions of position and velocity. The evolution of the electric energy is shown in Figure 23(c). We note that the ratio between the maximum and the minimum of $\mathcal{E}(X_{r,\tau}^i; \eta_i)$ under variations of the parameter η_i , is slightly larger than 2, indicating a certain variability of the solution in the range of parameters considered.




Figure 23: TSI: (a) – (b) Initial position and velocity distributions for selected values of the parameter in Γ_h . (c) Exponential time decay of the electrostatic energy $\mathcal{E}(X_{r,\tau}^i; \eta_i)$ obtained from the full model solution, for selected values of η_i in Γ_h . Since not all parameters in Γ_h are reported, the black lines in each subplot are used to mark the region where the plotted quantity is contained, for any value of the parameter in Γ_h .

The distribution function $f(t, x, v; \eta)$ is approximated with $N = 1.5 \cdot 10^5$ computational macro-particles, and $N_x = 64$ piecewise linear functions have been adopted to discretize the Poisson equation. We solve the discrete systems (5.2) and (5.3) with a time step $\Delta t = 0.0025$ over the temporal domain $\mathcal{T} = (0, 20]$.

Compared to previous tests, there is a dissimilarity between the decays of the singular values of the snapshots matrices for positions and velocities, as shown in Figure 24. The decay of the singular values of the S_X and S_X^τ is rather fast, and the singular values of the global snapshot matrix become smaller than 10^{-3} after the fourteen-th singular value. On the contrary, the decay of the singular values of the snapshots matrices associated with the particles velocity suggests that a local basis might be more effective in approximating the evolution of the particles velocity.




Figure 24: TSI: Singular values of the global snapshots matrices S_X and S_V compared to the maximum and time average of the singular values of the local trajectories matrices S_X^τ and S_V^τ . We study position (a) and velocity (b) variables separately. The singular values are normalized using the largest singular value for each case.

A similar conclusion can be drawn from the evolution of the numerical ranks of S_X^τ and S_V^τ , shown in Figures 25(a)-(b). In Figure 25(c), we also report the numerical rank of the self-consistent electric potential $\Phi(X_i^\tau)$ obtained from the full model at different time instants t_τ . It can be observed that the electric potential is low-rank throughout the simulation, which justifies the use of hyper-reduction strategies, in our case provided by the DMD-DEIM approach, to accelerate the computation of the nonlinearity in the Vlasov–Poisson equation. This speed up is ensured on the entire time interval since the numerical rank remains, on average, constant over time. We observe that, in principle, the rank of the hyper-reduced approximation provided by DMD and DEIM can change over time. As a general consideration, although the electric potential depends on the particles’ positions, there seems to be no straightforward connection between the reducibility properties of the sets $\{X_i^\tau\}_\tau$ and $\{\Phi(X_i^\tau)\}_\tau$.




Figure 25: TSI: Numerical rank of S_X^τ in (a) and S_V^τ in (b), as a function of τ . Different colors are associated with different values of the threshold, according to the legend. In (c) is reported the evolution of the numerical rank of the electric potential obtained from the full model.

In Figure 26, the evolution of the errors in the positions and velocities of the particles are reported: the approximability properties of the dynamical approach is not affected by the choice of the number

T of the DMD samples and the frequency k_{DEIM} of full updates of the DEIM indices. The dominant component of the error is the projection error. The growth in the error can be explained by the increase in time of the rank of the full model solution. This conclusion is supported by the same growth rate and the small difference between the relative error for the proposed approach and the relative projection error committed using an optimal ortho-symplectic basis of dimension $2n$ for both positions and velocities. We also stress that the error scales for the two components are different, as to be expected from the trend of singular values, and the greater accuracy in approximating the position is not affected by the velocity error.




Figure 26: TSI: Evolution of the position (a) and velocity (b) relative errors, as defined in (6.4), for different choices of k_{DEIM} and T . These errors are compared to the target values given by the position component $\varepsilon_{\text{rel},X}^{\text{Target}}$ and the velocity component $\varepsilon_{\text{rel},V}^{\text{Target}}$ of the relative projection errors defined in (6.5). The target reduced basis has dimension 4 and is computed, for each time step, using the Complex SVD algorithm, as described in Section 6.1.

To study the convergence properties of the proposed scheme in terms of the reduced basis size $2n$, we consider the same test case but in the parametric domain $\Gamma = [0.0075, 0.0125] \times [0.98, 1.02]$ and with a larger number $N = 5 \cdot 10^5$ of macro-particles. Increasing the size of the parametric domain produces an increase in the rank of the initial datum, while increasing the number of macro-particles reduces the statistical noise in the numerical rank that plagues particle simulations. As expected, Figure 27 shows a decrease of the error between the reduced and the full order solution as the size n of the dynamical reduced basis is increased. The numerical rank of the full model solution, shown in Figure 25, and the error evolution in Figures 26 and 27 suggest that enlarging the reduced basis U over time to increase the rank of the reduced model solution may improve the accuracy. Although a structure-preserving rank-adaptive algorithm has been proposed in [26], its direct application to the Vlasov–Poisson DMD–DEIM reduced model (4.15) would require the solution of a linear system of dimension proportional to the number of particles, to determine a candidate vector for the expansion of U . This cost would limit the computational speed-up obtained in the reduced model when compared to the full order model. The exploration of rank-adaptive algorithms for this case provides a possible direction of future investigation.




Figure 27: TSI: Evolution of the position (a) and velocity (b) relative errors, as defined in (6.4), for different choices of the reduced basis dimension $2n$. The values of the hyper-parameters k_{DEIM} and T are both set to 3.

The evolution of the electric energy (2.8) is shown in Figure 28: the behavior of the electric energy obtained from the approximate reduced model almost coincides with the full model except for slight mismatches in the amplitude of the oscillations during the transient phase.




Figure 28: TSI: Evolution of the electric field energy $\mathcal{E}(\cdot; \eta_i)$. The energy is evaluated at the positions X_τ^i computed using the high-fidelity solver and at the positions $X_{r,\tau}^i$ computed using the reduced model, for different values η_i of the parameter.

Similarly to the two numerical tests on the Landau damping, the proposed dynamical model order reduction method outperforms, in terms of efficiency, the full order solver, with speed-ups reaching 340 for $p = 10^5$ parameters, as shown in Figure 29.




Figure 29: TSI: (a) Comparison of the runtime per time step (in seconds) between the full order solver and the dynamical reduced basis approach for different hyper-parameter configurations, as function of the parameter sample size p . (b) Separation of contributions to the runtime per time step of the reduced model due to basis evolution (5.2b) (dashed lines) and coefficients evolution (5.2a) (continuous line).

To numerically show the effectiveness of our proposed hyper-reduction algorithm in highly nonlinear regimes, we have tested the method for times larger than $t = 20$ as follows. Until time $t = 30$ the error of the reduced solution behaves as in Figure 27, with no significant increase in the interval $[20, 30]$. After time $t = 30$ a reduced basis of size $2n \leq 8$ is no longer large enough to capture the dynamics with the required level of accuracy. We have considered as initial data for the algorithm (U_0, Z_0) obtained from complex SVD of the matrix $\mathcal{R}(30) := [W(30, \eta_1) | \dots | W(30, \eta_p)] \in \mathcal{V}_{2N}^p \subset \mathbb{R}^{2N \times p}$, that collects the full order solution at $t = 30$ for all parameters in Γ_h . These initial data allow reduced spaces of higher dimensions. In Figure 30 we report the results for $2n \in \{12, 18, 24\}$ in the case $T = 3$, $p^* = 8$, and $k_{DEIM} = 3$. The change of the rank of the approximation due to the re-initialization of the problem at time $t = 30$ ensures the required accuracy. A small deterioration over time can be ascribed to the fact that the complexity of the solution in terms of numerical rank increases.




Figure 30: TSI: Evolution of the position (a) and velocity (b) relative errors for different choices of the reduced basis dimension $2n$, with initial condition defined in the regime of highly nonlinear dynamics.

This test shows that the efficient treatment of the nonlinear operators proposed in this work has significant advantages over solving the full order model and the proposed method is able to capture the low-dimensional feature of the dynamics whenever the reducibility properties of the solution space are not considerably changing over time. As mentioned above, a combination with rank-adaptivity could further improve the proposed method whenever the rank of the solution undergoes significant variation during the evolution.

7 Concluding remarks

We proposed reduced order models of parametric particle-based kinetic plasma problems. High resolution simulations of such problems may require a high number of particles and thus can become computationally intractable in multi-query scenarios. Moreover, the lack of *global-in-time* reducibility of kinetic plasma models makes it hard for traditional reduced order models to provide accurate approximations at a competitive computational cost. We developed a dynamical approach to tackle this issue for problems that possess *local* low-rank structures. The proposed method combines dynamical low-rank approximation where the reduced space changes in time with adaptive hyper-reduction techniques to efficiently deal with the nonlinear operators. The resulting DMD-DEIM reduced dynamical system retains the Hamiltonian structure of the full model and, whenever the problem solution is locally low-rank, is able to provide accurate approximations that can be evaluated efficiently thanks to low-dimensional approximation spaces. Moreover, a considerable reduction in the computational runtime of the algorithm is achieved by decoupling the operations having a cost dependent on the number of particles from those that depend on the parameters. Several benchmark plasma models have been used to numerically assess the performances of the proposed method in dynamical regimes where the full order solution is locally low-rank.

The study of the more general case of solution sets with poor reducibility properties with respect to both time and parameter is left for future investigation. This might include the development of efficient error estimators to dynamically adapt the rank of the reduced order solution, and the combination of the proposed dynamical approach with localized MOR techniques. Future work might also include the derivation of partitioned Runge–Kutta schemes that can ensure the preservation of the Hamiltonian (at least when this is a lower degree polynomial), and the study of parameter sampling to speedup the computation of the reduced basis.

8 Acknowledgement

This work was partially supported by AFOSR under grant FA9550-17-1-9241.

References

- [1] B. M. Afkham and J. S. Hesthaven. “Structure preserving model reduction of parametric Hamiltonian systems.” *SIAM Journal on Scientific Computing* 39.6 (2017), A2616–A2644.
- [2] D. Anderson, R. Fedele, and M. Lisak. “A tutorial presentation of the two stream instability and Landau damping.” *American Journal of Physics* 69.12 (2001), pp. 1262–1266.
- [3] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera. “An ‘empirical interpolation’ method: application to efficient reduced-basis discretization of partial differential equations.” *Comptes Rendus Mathématique* 339.9 (2004), pp. 667–672.
- [4] G. Berge. “Landau damping in a plasma.” (1969).
- [5] C. Boutsidis, M. W. Mahoney, and P. Drineas. “An improved approximation algorithm for the column subset selection problem.” *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*. SIAM. 2009, pp. 968–977.
- [6] D. S. Broomhead and D. Lowe. *Radial basis functions, multi-variable functional interpolation and adaptive networks*. Tech. rep. Royal Signals and Radar Establishment Malvern (United Kingdom), 1988.
- [7] T. Bui-Thanh, M. Damodaran, and K. Willcox. “Proper orthogonal decomposition extensions for parametric applications in compressible aerodynamics.” *21st AIAA Applied Aerodynamics Conference*. 2003, p. 4213.
- [8] F. Casas, N. Crouseilles, E. Faou, and M. Mehrenberger. “High-order Hamiltonian splitting for the Vlasov–Poisson equations.” *Numer. Math.* 135.3 (2017), pp. 769–801.
- [9] S. Chaturantabut and D. C. Sorensen. “Nonlinear model reduction via discrete empirical interpolation.” *SIAM Journal on Scientific Computing* 32.5 (2010), pp. 2737–2764.
- [10] P. G. Constantine. *Active Subspaces*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2015.
- [11] P. Druault. “Développement d’interfaces expérience/simulation. Application à l’écoulement de couche de mélange plane turbulente.” PhD thesis. Université de Poitiers (France), 1999.
- [12] D. Dylewsky, M. Tao, and J. N. Kutz. “Dynamic mode decomposition for multiscale nonlinear physics.” *Physical Review E* 99.6 (2019), p. 063311.

- [13] V. Ehrlacher and D. Lombardi. “A dynamical adaptive tensor method for the Vlasov-Poisson system.” *J. Comput. Phys.* 339 (2017), pp. 285–306.
- [14] L. Einkemmer and I. Joseph. “A mass, momentum, and energy conservative dynamical low-rank scheme for the Vlasov equation.” *J. Comput. Phys.* 443 (2021), Paper No. 110495, 16.
- [15] L. Einkemmer and C. Lubich. “A Low-Rank Projector-Splitting Integrator for the Vlasov–Poisson Equation.” *SIAM Journal on Scientific Computing* 40.5 (2018), B1330–B1360.
- [16] L. Einkemmer and C. Lubich. “A Quasi-Conservative Dynamical Low-Rank Algorithm for the Vlasov Equation.” *SIAM Journal on Scientific Computing* 41.5 (2019), B1061–B1081.
- [17] E. G. Evstatiev and B. A. Shadwick. “Variational formulation of particle algorithms for kinetic plasma simulations.” *Journal of Computational Physics* 245 (2013), pp. 376–398.
- [18] C. Gillot, G. Dif-Pradalier, X. Garbet, P. Ghendrih, V. Grandgirard, and Y. Sarazin. “Model order reduction approach to the one-dimensional collisionless closure problem.” *Physics of Plasmas* 28.2 (2021), p. 022111.
- [19] J. Grosek and J. N. Kutz. “Dynamic mode decomposition for real-time background/foreground separation in video.” *arXiv preprint arXiv:1404.7592* (2014).
- [20] E. Hairer, C. Lubich, and G. Wanner. *Geometric numerical integration*. Second. Vol. 31. Springer Series in Computational Mathematics. Springer-Verlag, Berlin, 2006.
- [21] N. Halko, P. G. Martinsson, and J. A. Tropp. “Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions.” *SIAM Rev.* 53.2 (2011), pp. 217–288.
- [22] Y. He, Y. Sun, H. Qin, and J. Liu. “Hamiltonian particle-in-cell methods for Vlasov-Maxwell equations.” *Physics of Plasmas* 23.9 (2016), p. 092108.
- [23] M. S. Hemati, M. O. Williams, and C. W. Rowley. “Dynamic mode decomposition for large and streaming datasets.” *Physics of Fluids* 26.11 (2014), p. 111701.
- [24] H. Hereward. *Landau damping*. Tech. rep. 1977.
- [25] J. S. Hesthaven and C. Pagliantini. “Structure-preserving reduced basis methods for Poisson systems.” *Math. Comp.* 90.330 (2021), pp. 1701–1740.
- [26] J. S. Hesthaven, C. Pagliantini, and N. Ripamonti. “Rank-adaptive structure-preserving model order reduction of Hamiltonian systems.” *ESAIM Math. Model. Numer. Anal.* 56.2 (2022), pp. 617–650.
- [27] R. Hockney. “Characteristics of Noise in a Two-Dimensional Computer Plasma.” *The Physics of Fluids* 11.6 (1968), pp. 1381–1383.
- [28] O. Koch and C. Lubich. “Dynamical low-rank approximation.” *SIAM J. Matrix Anal. Appl.* 29.2 (2007), pp. 434–454.
- [29] B. O. Koopman. “Hamiltonian systems and transformation in Hilbert space.” *Proceedings of the national academy of sciences of the united states of america* 17.5 (1931), p. 315.
- [30] K. Kormann. “A semi-Lagrangian Vlasov solver in tensor train format.” *SIAM J. Sci. Comput.* 37.4 (2015), B613–B632.
- [31] K. Kormann and E. Sonnendrücker. “Sparse grids for the Vlasov-Poisson equation.” *Sparse grids and applications—Stuttgart 2014*. Vol. 109. Lect. Notes Comput. Sci. Eng. Springer, Cham, 2016, pp. 163–190.
- [32] M. Kraus, K. Kormann, P. J. Morrison, and E. Sonnendrücker. “GEMPIC: geometric electromagnetic particle-in-cell methods.” *Journal of Plasma Physics* 83.4 (2017), p. 905830401.
- [33] J. A. Krommes. “Nonequilibrium gyrokinetic fluctuation theory and sampling noise in gyrokinetic particle-in-cell simulations.” *Physics of Plasmas* 14.9 (2007), p. 090501.
- [34] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor. *Dynamic Mode Decomposition*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2016.
- [35] A. Langdon and C. Birdsall. *Plasma Physics via Computer Simulation*. First. Springer Series in Computational Mathematics. CRC Press, 1991.
- [36] H. R. Lewis. “Energy-conserving numerical approximations for Vlasov plasmas.” *Journal of Computational Physics* 6.1 (1970), pp. 136–141.
- [37] F. E. Low. “A Lagrangian formulation of the Boltzmann-Vlasov equation for plasmas.” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 248.1253 (1958), pp. 282–287.
- [38] G. Manfredi. “Long-time behavior of nonlinear Landau damping.” *Physical review letters* 79.15 (1997), p. 2815.
- [39] S. Mariappan, A. D. Gardner, K. Richter, and M. Raffel. “Analysis of dynamic stall using dynamic mode decomposition technique.” *AIAA journal* 52.11 (2014), pp. 2427–2439.
- [40] J. E. Marsden and A. Weinstein. “The Hamiltonian structure of the Maxwell-Vlasov equations.” *Physica D: Nonlinear Phenomena* 4.3 (1982), pp. 394–406.
- [41] D. Matsumoto and T. Indinger. “On-the-fly algorithm for dynamic mode decomposition using incremental singular value decomposition and total least squares.” *arXiv preprint arXiv:1703.11004* (2017).

- [42] P. J. Morrison. “Hamiltonian field description of the one-dimensional Poisson-Vlasov equations” (July 1981).
- [43] E. Musharbash, F. Nobile, and E. Vidličková. “Symplectic dynamical low rank approximation of wave equations with random parameters.” *BIT Numerical Mathematics* (2017), pp. 1–49.
- [44] I. Nayak, M. Kumar, and F. L. Teixeira. “Detection and prediction of equilibrium states in kinetic plasma simulations via mode tracking using reduced-order dynamic mode decomposition.” *Journal of Computational Physics* 447 (2021), p. 110671.
- [45] C. Pagliantini. “Dynamical reduced basis methods for Hamiltonian systems.” *Numer. Math.* 148.2 (2021), pp. 409–448.
- [46] M. F. de Pando, P. J. Schmid, and D. Sipp. “Nonlinear model-order reduction for compressible flow solvers using the Discrete Empirical Interpolation Method.” *Journal of Computational Physics* 324 (2016), pp. 194–209.
- [47] B. Peherstorfer and K. Willcox. “Online adaptive model reduction for nonlinear systems via low-rank updates.” *SIAM Journal on Scientific Computing* 37.4 (2015), A2123–A2150.
- [48] L. Peng and K. Mohseni. “Symplectic model reduction of Hamiltonian systems.” *SIAM Journal on Scientific Computing* 38.1 (2016), A1–A27.
- [49] A. Pinkus. *n-widths in approximation theory*. Vol. 7. Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]. Springer-Verlag, Berlin, 1985.
- [50] J. L. Proctor and P. A. Eckhoff. “Discovering dynamic patterns from infectious disease data using dynamic mode decomposition.” *International health* 7.2 (2015), pp. 139–145.
- [51] H. Qin, J. Liu, J. Xiao, R. Zhang, Y. He, Y. Wang, Y. Sun, J. W. Burby, L. Ellison, and Y. Zhou. “Canonical symplectic particle-in-cell method for long-term large-scale simulations of the Vlasov–Maxwell equations.” 56.1 (2015), p. 014001.
- [52] A. Quarteroni, A. Manzoni, and F. Negri. *Reduced basis methods for partial differential equations*. Vol. 92. Unitext. Springer, Cham, 2016.
- [53] C. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. Henningson. “Spectral analysis of nonlinear flows.” *Journal of fluid mechanics* 641.1 (2009), pp. 115–127.
- [54] T. P. Sapsis and P. F. J. Lermusiaux. “Dynamically orthogonal field equations for continuous stochastic dynamical systems.” *Phys. D* 238.23–24 (2009), pp. 2347–2360.
- [55] T. Sayadi, P. J. Schmid, F. Riechecoeur, and D. Durox. “Parametrized data-driven decomposition for bifurcation analysis, with application to thermo-acoustically unstable systems.” *Physics of Fluids* 27.3 (2015), p. 037102.
- [56] P. J. Schmid. “Dynamic mode decomposition of numerical and experimental data.” *Journal of fluid mechanics* 656 (2010), pp. 5–28.
- [57] P. J. Schmid. “Application of the dynamic mode decomposition to experimental data.” *Experiments in fluids* 50.4 (2011), pp. 1123–1130.
- [58] C. Scovel and A. Weinstein. “Finite-dimensional Lie-Poisson approximations to Vlasov-Poisson equations.” *Comm. Pure Appl. Math.* 47.5 (1994), pp. 683–709.
- [59] Y. Shitov. “Column subset selection is NP-complete.” *Linear Algebra and its Applications* 610 (2021), pp. 52–58.
- [60] J. Squire, H. Qin, and W. M. Tang. “Geometric integration of the Vlasov-Maxwell system with a variational particle-in-cell scheme.” *Physics of Plasmas* 19.8 (2012), p. 084501.
- [61] R. Ștefănescu, A. Sandu, and I. M. Navon. “Comparison of POD reduced order strategies for the nonlinear 2D shallow water equations.” *International Journal for Numerical Methods in Fluids* 76.8 (2014), pp. 497–521.
- [62] R. Sydora. “Low-noise electromagnetic and relativistic particle-in-cell plasma simulation models.” *Journal of computational and applied mathematics* 109.1-2 (1999), pp. 243–259.
- [63] M. Tao. “Explicit symplectic approximation of nonseparable Hamiltonians: Algorithm and long time performance.” *Physical Review E* 94.4 (2016), p. 043303.
- [64] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz. “On dynamic mode decomposition: Theory and applications.” *Journal of Computational Dynamics* 1.2 (2014), pp. 391–421.
- [65] T. M. Tyranowski and M. Kraus. “Symplectic model reduction methods for the Vlasov equation.” *arXiv preprint arXiv:1910.06026* (2019).
- [66] J. Xiao, H. Qin, J. Liu, Y. He, R. Zhang, and Y. Sun. “Explicit high-order non-canonical symplectic particle-in-cell algorithms for Vlasov-Maxwell systems.” *Physics of Plasmas* 22.11 (2015), p. 112504.