Structure preserving low-rank algorithms for plasma simulations Part 3: Physical structure of the Vlasov equation and low-rank projection

Lukas Einkemmer
University of Innsbruck

Structure-Preserving Scientific Computing and Machine Learning Summer School and Hackathon, UW Seattle, 2025

Link to slides: `http://www.einkemmer.net/training.html`

Efficient implementation using Ensign

Software aspects

DLR is mostly formulated in terms of **dense** matrix/vector operations.

- Very well suited for modern computer architectures.

Ensign is a C++ software package that facilitates the implementation of dynamical low-rank algorithms.

- Data structure for DLR (matrix and tensor tree decompositions)
- Functions for computing coefficients, initialization/orthogonalization, ...
- Can be supplemented by user defined space and time discretization routines
- Uses highly optimized linear algebra routines behind the scenes (Intel MKL, cuBLAS, ...)
- Support for **multi-core CPU** and CUDA capable **GPUs**

Ensign is open source under the MIT license.

- `https://github.com/leinkemmer/Ensign`

6D Vlasov–Poisson simulation

A **3+3 dimensional two-stream instability** with $r = 5$, $r = 10$, $r = 15$, $r = 20$.

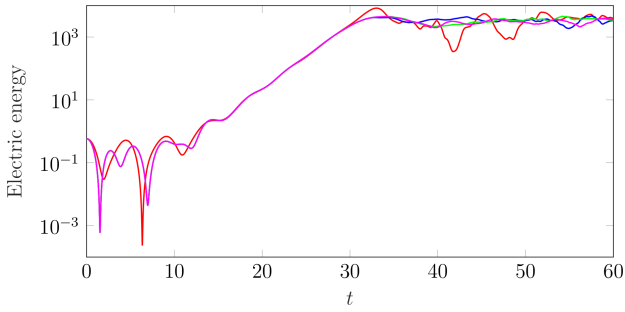

Dense linear algebra is very efficient on GPUs.

- For $128^3 \times 128^3$ grid points and $r = 10$ simulation time on GPU is 20 min.
- Would require at least 70 TB of memory for directly solving Vlasov-Poisson.

F. Cassini, L.E. Comput. Phys. Commun. 280, 2022.

Alfvén waves

Shear Alfvén waves are **electromagnetic** waves in a plasma that **propagate parallel to the magnetic field**.

- Can be destabilized by beams (as in NBI) or energetic particles.

Described by a 4D gyrokinetic equation.

Resolution	rank	Lie (time/step)	speedup	memory usage	memorydown
$64^3 \times 512$	full	2.3e+02 s	–	2.1 GB	–
	2	0.024 s	**9276**	1 MB	**2048**
	5	0.13 s	**1764**	3 MB	**819**
	10	0.8 s	**283**	5 MB	**410**
	15	2.9 s	**78**	8 MB	**273**
$256^3 \times 1024$	full	–	–	**2.7e+02 GB**	–
	2	0.45 s	–	8 MB	32768
	5	4.1 s	–	21 MB	13107
	10	35 s	–	42 MB	6554
	15	1.7e+02 s	–	**63 MB**	4369

L.E. J. Comput. Phys. 501, 2024.

Ensign example

```
#include <lr/lr.hpp>

// vector and matrices (can be either on the CPU or GPU)
using vec  = Ensign::multi_array<double,1>;
using mat  = Ensign::multi_array<double,2>;
using lr   = Ensign::lr2<double>;

// Initialization
lr f(r, {nx, nv});
Ensign::initialize(f, in_x, in_v, hx, hv, blas); // in_x and in_v filled with IC

// K step
mat K({nx,r});
mat c1 = compute_c1(f.V);
mat c2 = compute_c2(f.V);
blas.matmul(f.X, f.S, K);
rk4(deltat, K, [&c1, &c2, &E](const mat& K, mat& Kout) {
  Kout = rhs_K(K, c1, c2, E);
});
```

Ensign example

```
#include <lr/coefficients.hpp>

mat compute_c1(const mat& V) {
  mat Vtmp({nv,r}), c1({r,r});
  Ensign::Matrix::ptw_mult_row(V,vs,Vtmp); // multiply by v
  Ensign::coeff(V, Vtmp, hv, c1, blas);
  return c1;
}

vec compute_rho(const lr& f) {
  mat K({nx, r});
  vec int_V({r}), rho({nx});

  Ensign::integrate(f.V,-hv,int_V,blas);
  blas.matmul(f.X,f.S,K);
  blas.matvec(K,int_V,rho);
  return rho;
}
```

Applications beyond plasma physics

Chemical master equation solver based on Ensign developed by Julian Mangott.

`pip install atropy`

Tree based decompositions (such as tensor trains and hierarchical Tucker)

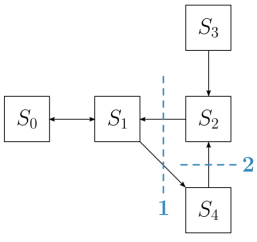

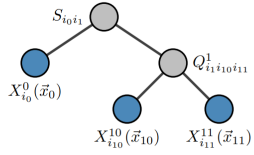

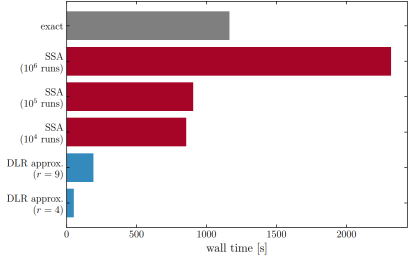

[Video: toogle switch]

Physical structure of the Vlasov equation

Hamiltonian systems

Reminder: In a Hamiltonian system the time evolution of a quantity $F(p, q)$ can be written as

$$\partial_t F = \{F, H\}$$

with Hamiltonian H and the symplectic Poisson bracket

$$\{F, G\} = \nabla_q F \cdot \nabla_p G - \nabla_p F \cdot \nabla_q G.$$

Since $\partial_t H = \{H, H\} = 0$ we follow that the Hamiltonian, i.e. **the energy of the system, is conserved**.

Every quantity C such that $\{C, H\} = 0$ is conserved.

A **non-canonical Hamiltonian** system admits a Poisson bracket $[\cdot, \cdot]$ (not necessarily the symplectic bracket above) that satisfies Anticommutativity, Bilinearity, Leibniz's rule, and the Jacobi identity.

Non-canonical Hamiltonian systems

Example: The generalized Lotka–Volterra model

$$\dot{u} = u(v + w), \qquad \dot{v} = v(u - w + 1), \qquad \dot{w} = w(u + v + 1)$$

is a non-canonical Hamiltonian system with

$$H(u, v) = -u + v + w + \ln v - \ln w \quad \text{and} \quad [F, G] = (\nabla_{uvw} F)^T \underbrace{\begin{bmatrix} 0 & uv & uw \\ -uv & 0 & -vw \\ -uw & vw & 0 \end{bmatrix}}_{B} \nabla_{uvw} G$$

C is a **Casimir invariant** if $\{C, F\} = 0$ holds for all F.

- We call the Poisson bracket degenerate.

Example: For the Lotka–Volterra model $C(u, v) = -\ln u - \ln v + \ln w$ is a Casimir invariant as $(\nabla_{uvw} C)^T B = 0$.

Vlasov–Maxwell equations

Vlasov equation

$$\partial_t f + v \cdot \nabla_x f - (E + v \times B) \cdot \nabla_v f = 0$$

coupled to Maxwell's equations

$$\partial_t E = c^2 \nabla_x \times B - j, \qquad \partial_t B = -\nabla_x \times E,$$

where $j = \int v f \, \mathrm{d}v$.

There are also **constraints** (automatically satisfied in the continuous case)

$$\nabla \cdot B = 0, \qquad \nabla \cdot E = 1 - \rho$$

Hamiltonian structure

The **Vlasov–Maxwell equations** have a **non-canonical Hamiltonian** structure.

$$\begin{aligned} H &= \frac{1}{2}\int |E|^2 \,\mathrm{d}x + \frac{c^2}{2}\int |B|^2 \,\mathrm{d}x + \frac{1}{2}\int v^2 f \,\mathrm{d}(x,v) \\ &=: H_E + H_B + H_f. \end{aligned}$$

Evolution of F

$$\partial_t F = [F, H] = [F, H_E] + [F, H_B] + [F, H_f]$$

with a highly non-canonical Poisson bracket.

We have an **infinite number of Casimir invariants** as any $C(f)$ satisfies $[C, G] = 0$ for arbitrary G.

- $\|f\|_2$ (in fact, any L^p norm)
- Entropy $-\int_\Omega f \log f \,\mathrm{d}(x,v)$.

Accuracy of Vlasov simulation

Performance is often checked by using a work-precision diagram.

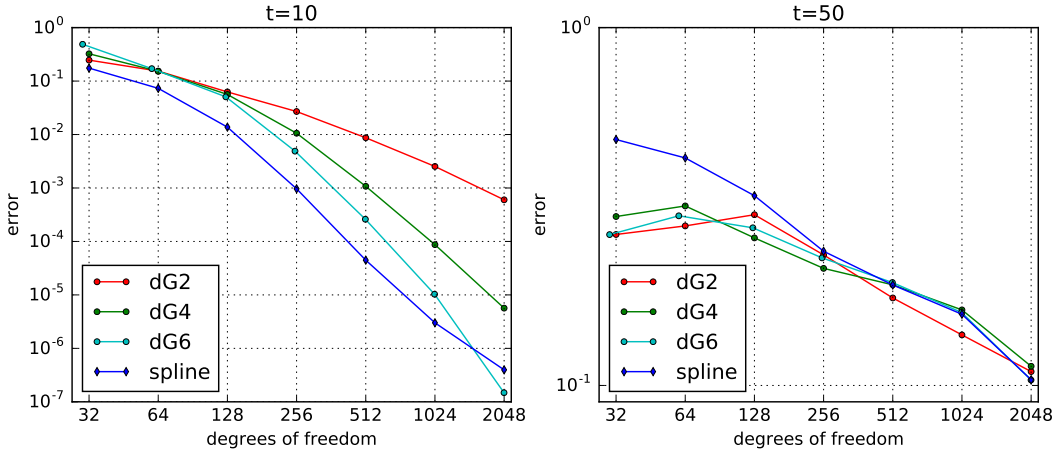

Two-stream instability

Vlasov–Poisson equation with 128^4 degrees of freedom.

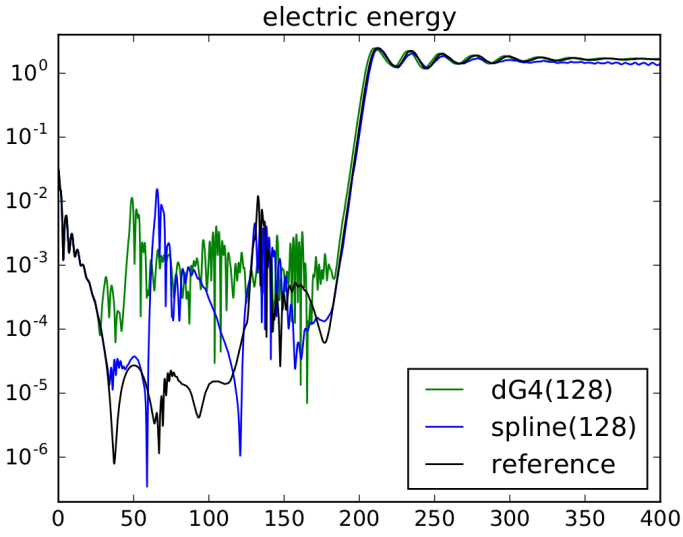

Two-stream instability

Vlasov–Poisson equation with 128^4 degrees of freedom.

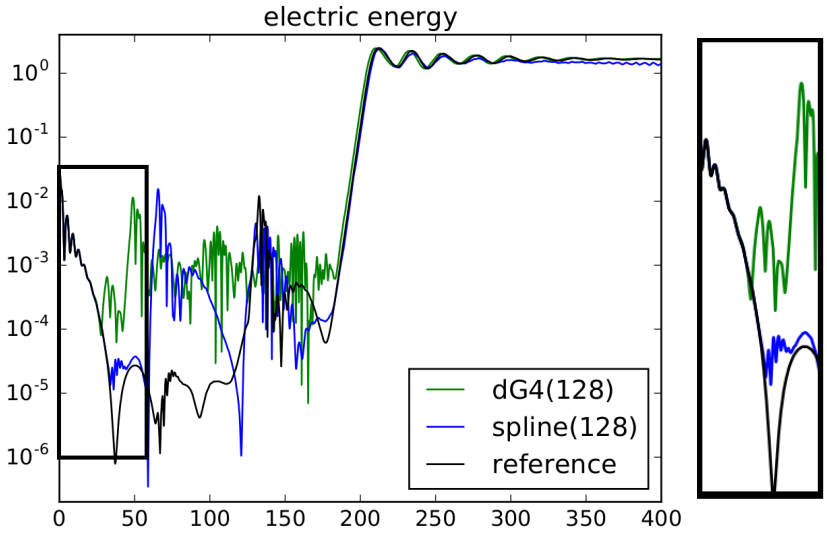

Two-stream instability

Vlasov–Poisson equation with 32^4 degrees of freedom.

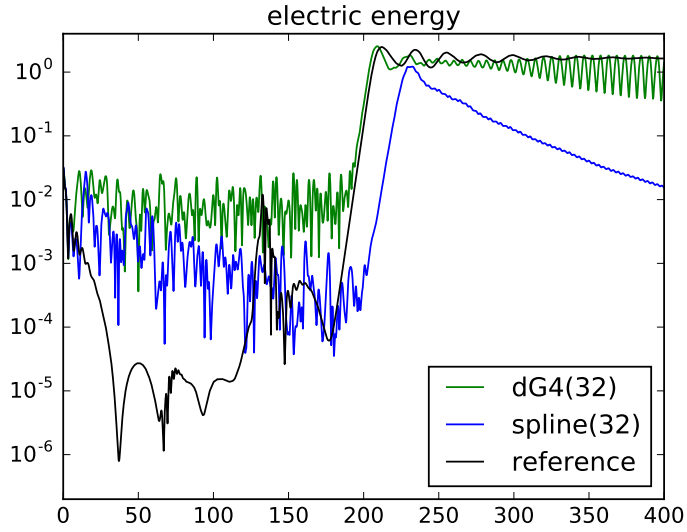

Discussion

We are in the **asymptotic regime** if classic convergence theory applies. That is,

$$\text{error} \leq C\left((\Delta t)^p + (\Delta x)^q\right).$$

gives a tight bound of the error.

Why is this not the case here? Consider

$$\partial_t f(t,x,v) + v\partial_x f(t,x,v) = 0, \qquad f(0,x,v) = \mathrm{e}^{ikx}\mathrm{e}^{-v^2/2}$$

which has the solution

$$f(t,x,v) = \mathrm{e}^{ikt v}\mathrm{e}^{ikx}\mathrm{e}^{-v^2/2}.$$

Small scale structures (e.g. filamentation, turbulence, ...) can not be resolved.

- **All methods are necessarily inaccurate.**
- Often we can still get good physics out of those methods.

Two-stream instability

L^2 **norm** as a measure of **numerical diffusion**.

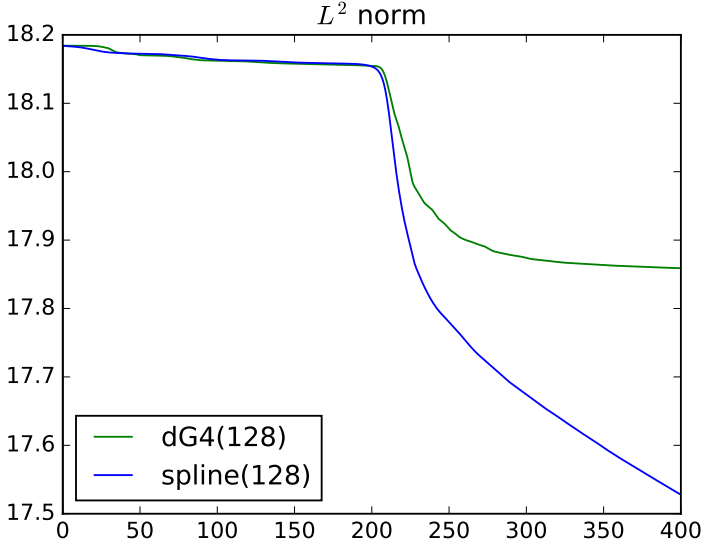

Dynamical low-rank and conservation

Galerkin condition

Orthogonal projection

$$\text{Find} \quad \partial_t f = g \in T_f\mathcal{M} \quad \text{such that} \quad \|g - \text{RHS}\| \text{ is minimal.}$$

That is, $\partial_t f = g = P(f)\text{RHS}$.

Galerkin condition

$$\text{Find} \quad \partial_t f \quad \text{such that} \quad \langle \nu, \partial_t f \rangle = \langle \nu, \text{RHS} \rangle \quad \forall \nu \in T_f\mathcal{M}.$$

For the Schrödinger equation implies symplecticity, energy, and norm conservation.

- But the situation for kinetic equations is very different.

L^2 conservation

Galerkin condition implies **$\mathbf{L^2}$ norm conservation**

$$\partial_t \|f\|^2 = 2\langle f, \partial_t f\rangle_{xv} = 2\langle f, \mathrm{RHS}\rangle_{xv} = 0$$

since $f \in T_f\mathcal{M}$.

But wait, why do we have $\langle f, \mathrm{RHS}\rangle_{xv} = 0$?

This is how we (directly) prove L^2 conservation for the underlying model

$$\begin{aligned} 2\langle f, \mathrm{RHS}\rangle_{xv} &= 2\int -vf \cdot \nabla_x f + Ef \cdot \nabla_v f \,\mathrm{d}(x,v) = 0. \\ &= \int -\nabla_x \cdot (vf^2) + \nabla_v \cdot (Ef^2)\,\mathrm{d}(x,v) = 0. \end{aligned}$$

The analytic argument carries over. **This will be an important technique!**

Mass conservation

From

$$\partial_t f + \nabla_x \cdot (vf) - \nabla_v \cdot (Ef) = 0$$

we follow by integrating in v

$$\partial_t \int f \, \mathrm{d}x + \nabla_x \cdot \int vf \, \mathrm{d}v = 0,$$

which is more commonly written as

$$\partial_t \rho + \nabla \cdot j = 0, \qquad \rho = \int f \, \mathrm{d}v, \qquad j = \int vf \, \mathrm{d}v.$$

Integrating in x we get

$$M = \int f \mathrm{d}(x, v) = \text{const.}$$

That is, **conservation of mass**/conservation of particles/conservation of charge.

Momentum conservation

From

$$\partial_t(vf) + \nabla_x \cdot ((v \otimes v)f) - v\nabla_v \cdot (Ef) = 0$$

we follow by integrating in v

$$\partial_t j + \nabla_x \cdot \sigma = \int v\nabla_v \cdot (Ef)\,\mathrm{d}v = -\int Ef\,\mathrm{d}v = E\rho, \qquad \sigma = \int (v \otimes v)f\,\mathrm{d}v.$$

Since

$$E(1-\rho) = \nabla \cdot (E \otimes E - \tfrac{1}{2}E^2)$$

and $\int E\,\mathrm{d}x = 0$ we obtain

$$P = \int vf\,\mathrm{d}(x,v) = \text{const},$$

That is, **conservation of momentum**/conservation of current.

Energy conservation

We already know that the **energy (i.e. the Hamiltonian) is conserved**

$$H = \frac{1}{2}\int v^2 f \,\mathrm{d}(x,v) + \frac{1}{2}\int E^2 \,\mathrm{d}x.$$

Similar to mass and momentum there is also an associated **local conservation law**

$$\partial_t e + \nabla_x \cdot Q = E \cdot (\partial_t E - j), \qquad e = \frac{1}{2}\int v^2 f \, dv + \frac{1}{2}E^2, \quad Q = \frac{1}{2}\int v v^2 f \, dv.$$

Dynamical low-rank approximation

The dynamical low-rank approximation finds the best L^2 approximation.

- No guarantee that mass, momentum, or energy is conserved.

Linear Landau damping (left) and two-stream instability (right).

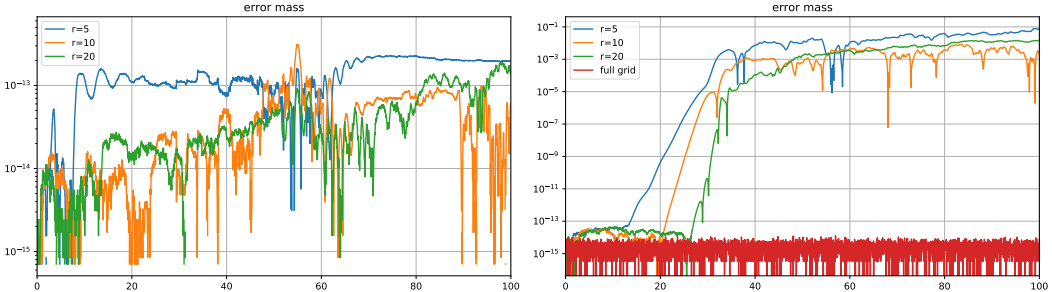

This **failure is in stark contrast to Eulerian and particle methods**.

Literature

[Z. Peng, R. McClarren, M. Frank. J. Comput. Phys., 421 (2020)]

- Rescale solution to obtain mass conservation.
- Global mass conservation only.
- Not extensible to other invariants.

[Z. Peng, and R.G. McClarren. J. Comput. Phys. 447, 2021]

- Couple moments with low-rank approximation of g, where $f = M + g$.
- Needs to enforce $\int g \, \mathrm{d}(x, v) = 0$.
- Global invariants only.

[L. Einkemmer, C. Lubich. SIAM J. Sci. Comput., 40(5) (2018)]

- Add correction $\lambda_{ij} X_i V_j$ to enforce conservation (Lagrange multiplier).
- Conserves either global invariants or (a projected version of) conservation laws.
- Not able to simultaneously conserve both.

Global vs local conservation

We should not forget the **local conservation law!**

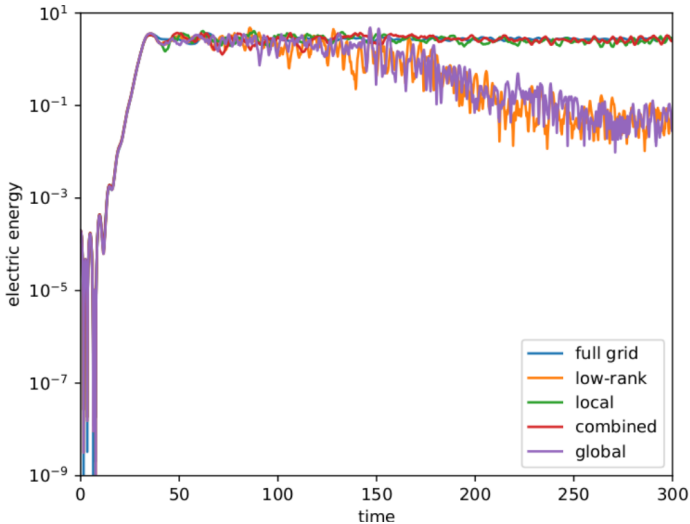

Literature

Literature

[F. Casas, N. Crouseilles, E. Faou. M. Mehrenberger. Numer. Math. 135, 2017]

- Hamiltonian structure for the Vlasov–Poisson equation.
- Used for constructing high-order splitting methods for direct (i.e. not low-rank) solvers.

[P.J. Morrison. Phys. Lett. A 80:4-5, 1980.]
[J.E. Marsden. A. Weinstein. Physica D. 4:3, 1982]

- The rather complicated Hamiltonian structure of Vlasov–Maxwell.

[L.E. J. Comput. Phys. 376, 2019]

- Role of conservative methods for integrating the Vlasov equation.
- Accurate solution despite large phase space error.

[Ensign]

- C++ framework for dynamical low-rank computations.