SAND2020-3112C

Discretization of the Multi-fluid Plasma Model using IMEX and Mixed Continuous/Discontinuous FEM

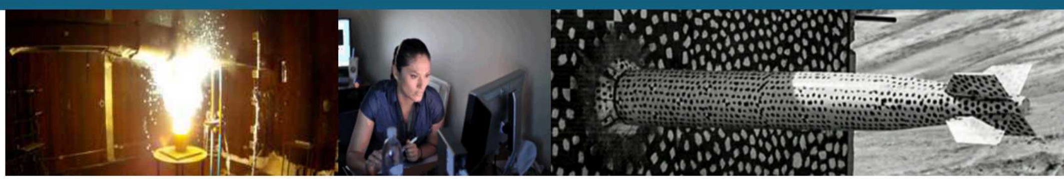

The 44th SIAM Southeastern Atlantic Section Conference
March 14-15, 2020, Auburn University

Sean T. Miller, Eric C. Cyr, Edward G. Phillips, Roger P. Pawlowski, Kris Beckwith, Matt Bettencourt, Keith Cartwright, Sidney Shields

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Motivation

Sandia's applications span the plasma regime from rarefied to HEDP plasma physics.

Multi-fluid plasma models simulate each species independently and couple the species together through collisional and electromagnetic operators.

- Primarily designed for systems with fast time scales where electron inertial effects and charge separation are resolved.

Main drawback: Model is computational expensive due to the necessity of resolving fast plasma scales associated with electron dynamics and Maxwell's equations.

Objective: Explore efficient treatments of the multi-fluid plasma model using implicit-explicit (IMEX) time integration and mixed discontinuous and continuous finite element methods.

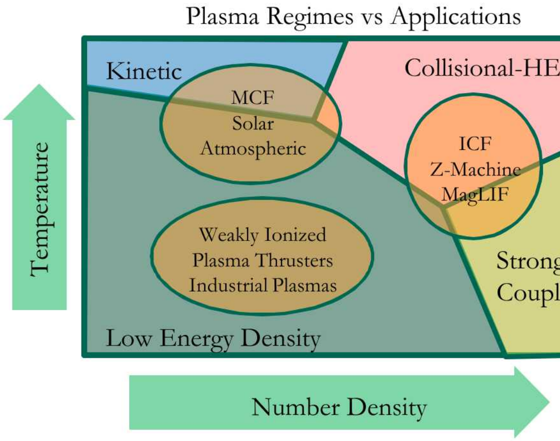

EMPIRE: A hierarchy of capabilities

EMPIRE: A high-performance, multi-scale plasma physics modeling code developed at Sandia National Laboratories for accurately simulating systems across a wide parameter regime on next-generation exascale computing platforms.

Goal: Expand the capability for modeling electromagnetic pulse and Z-power flow applications with high confidence and fidelity.

Features include:

- Relativistic particle-in-cell (PIC).
- Fluid-based neutral/plasma models (Euler, MHD, multi-fluid, relativistic, etc).
- Hybrid capabilities (multi-species and delta-f).
- DSMC based collision/reaction models.

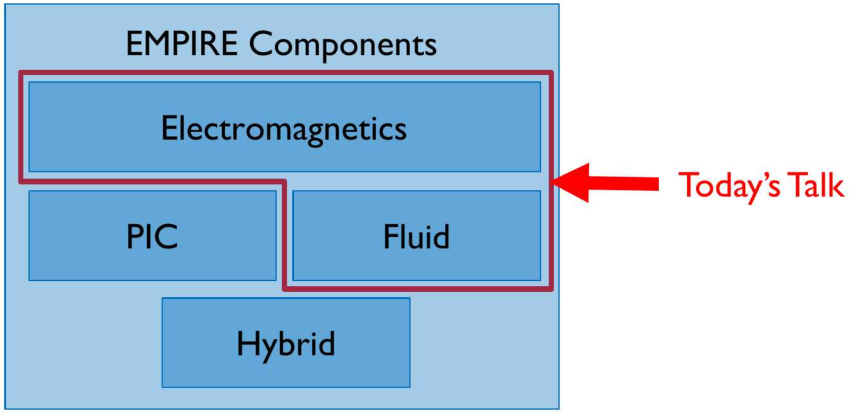

Multi-fluid plasma model

Continuity Equation:

$$\partial_t \rho_\alpha + \nabla \cdot (\rho_\alpha \boldsymbol{u}_\alpha) = S_\alpha$$

Each species α is represented by a separate density ρ, momentum $\rho\boldsymbol{u}$, and isotropic energy ϵ.

Momentum Equation:

$$\partial_t (\rho_\alpha \boldsymbol{u}_\alpha) + \nabla \cdot (\rho_\alpha \boldsymbol{u}_\alpha \otimes \boldsymbol{u}_\alpha + \boldsymbol{P}_\alpha) = \frac{q_\alpha}{m_\alpha} \rho_\alpha (\boldsymbol{E} + \boldsymbol{u}_\alpha \times \boldsymbol{B}) + \boldsymbol{R}_\alpha + \boldsymbol{u}_\alpha S_\alpha$$

Fluid

Electromagnetic

Inter-fluid

Energy Equation:

$$\partial_t \epsilon_\alpha + \nabla \cdot (\boldsymbol{u}_\alpha \cdot (\epsilon_\alpha \boldsymbol{I} + \boldsymbol{P}_\alpha) + \boldsymbol{q}_\alpha) = \frac{q_\alpha}{m_\alpha} \rho_\alpha \boldsymbol{u}_\alpha \cdot \boldsymbol{E} + Q_\alpha + \boldsymbol{u}_\alpha \cdot \boldsymbol{R}_\alpha + \frac{1}{2} \boldsymbol{u}_\alpha^2 S_\alpha$$

Ampere's Law:

$$\partial_t \boldsymbol{E} - c^2 \nabla \times \boldsymbol{B} = -\frac{1}{\epsilon_0} \sum_\alpha \frac{q_\alpha}{m_\alpha} \rho_\alpha \boldsymbol{u}_\alpha$$

Spatial operators are discretized using a discontinuous Galerkin finite element method.

Faraday's Law:

$$\partial_t \boldsymbol{B} + \nabla \times \boldsymbol{E} = 0$$

Cacophony of plasma time scales

Multi-fluid plasma models are littered with multi-scale physics.

- Strongly dependent on species mass, density, and temperature.
- Speed of light, plasma and cyclotron frequency are often stiff!
- Can be grouped into frequency, velocity, and diffusion scales:

Plasma frequency

$$\omega_{p\alpha} = \sqrt{\frac{q_\alpha^2 n_\alpha}{m_\alpha \epsilon_0}}$$

Cyclotron frequency

$$\omega_{c\alpha} = \frac{q_\alpha B}{m_\alpha}$$

Collision frequency

$$\nu_{\alpha\beta} \sim \frac{n_\beta}{\sqrt{m_\alpha} T_\alpha^{\frac{3}{2}}} \frac{1 + \frac{m_\alpha}{m_\beta}}{\left(1 + \frac{m_\alpha}{m_\beta}\frac{T_\beta}{T_\alpha}\right)^{\frac{3}{2}}}$$

Flow velocity

$$u_\alpha$$

Speed of sound

$$v_{s\alpha} = \sqrt{\frac{\gamma P_\alpha}{\rho_\alpha}}$$

Speed of light

$$c \gg u_\alpha, v_{s\alpha}$$

Momentum diffusivity

$$\upsilon_\alpha = \frac{\mu_\alpha}{\rho_\alpha}$$

Thermal diffusivity

$$\kappa_\alpha \sim \frac{k_\alpha}{\rho_\alpha}$$

IMEX time integration

IMEX gives a framework for splitting the model up into implicit and explicit terms:

- Explicit for slow, non-stiff terms
- Implicit for fast, stiff terms

$$\partial_t u = f(u,t) + g(u,t)$$

$$u^{(i)} = u^n + \Delta t \sum_{j=0}^{j<i} \hat{A}_{ij} f\left(u^{(j)}, t_n + \hat{c}_j \Delta t\right) + \Delta t \sum_{j=0}^{j\leq i} A_{ij} g\left(u^{(j)}, t_n + c_j \Delta t\right)$$

$$u^{n+1} = u^n + \Delta t \sum_{i=0}^{i<s} \hat{b}_i f\left(u^{(i)}, t_n + \hat{c}_i \Delta t\right) + \Delta t \sum_{i=0}^{i\leq s} b_i g\left(u^{(i)}, t_n + c_i \Delta t\right)$$

Implicit tableau

$$\begin{array}{c|c} c & A \\ \hline & b^t \end{array}$$

Explicit tableau

$$\begin{array}{c|c} \hat{c} & \hat{A} \\ \hline & \hat{b}^t \end{array}$$

3 Stage IMEX-RK Algorithm

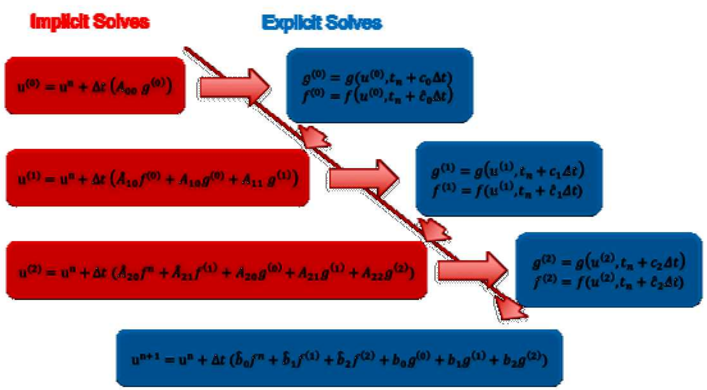

Objective: Take advantage of expensive implicit solver to overstep fast scales, and cheap explicit solver to resolve slow scales.

Breaking up plasma model for IMEX

Each operator is associated with one or more plasma scales.

- Here we group them by color representing their explicit stability limits:

Equations

$$\partial_t \rho_\alpha + \nabla \cdot (\rho_\alpha \boldsymbol{u}_\alpha) = S_\alpha$$

$$\partial_t (\rho_\alpha \boldsymbol{u}_\alpha) + \nabla \cdot (\rho_\alpha \boldsymbol{u}_\alpha \otimes \boldsymbol{u}_\alpha + P_\alpha) = \frac{q_\alpha}{m_\alpha}\rho_\alpha \boldsymbol{E} + \frac{q_\alpha}{m_\alpha}\rho_\alpha \boldsymbol{u}_\alpha \times \boldsymbol{B} + \boldsymbol{R}_\alpha + \boldsymbol{u}_\alpha S_\alpha$$

$$\partial_t \epsilon_\alpha + \nabla \cdot \left(\boldsymbol{u}_\alpha \cdot (\epsilon_\alpha + P_\alpha)\right) = \frac{q_\alpha}{m_\alpha}\rho_\alpha \boldsymbol{u}_\alpha \cdot \boldsymbol{E} + Q_\alpha + \boldsymbol{u}_\alpha \cdot \boldsymbol{R}_\alpha + \frac{1}{2}\boldsymbol{u}_\alpha^2 S_\alpha$$

$$\partial_t \boldsymbol{E} - c^2 \nabla \times \boldsymbol{B} = -\frac{1}{\epsilon_0}\sum_\alpha \frac{q_\alpha}{m_\alpha}\rho_\alpha \boldsymbol{u}_\alpha$$

$$\partial_t \boldsymbol{B} + \nabla \times \boldsymbol{E} = 0$$

Approx. Explicit Stability Limits

$$|u_\alpha| + |v_{s\alpha}| < \frac{\Delta x}{\Delta t}$$

$$\nu_\alpha \Delta t < 1$$

$$\omega_{p\alpha} \Delta t < 1$$

$$\omega_{c\alpha} \Delta t < 1$$

$$c < \frac{\Delta x}{\Delta t}$$

Move operators between explicit/implicit treatment depending on their "stiffness".

Discontinuous Galerkin (DG) has a local stencil

DG breaks discretization into "local" interior and "non-local" interface operations.

- Interior operations: Internal flux, source terms
- Interface operations: Numerical flux

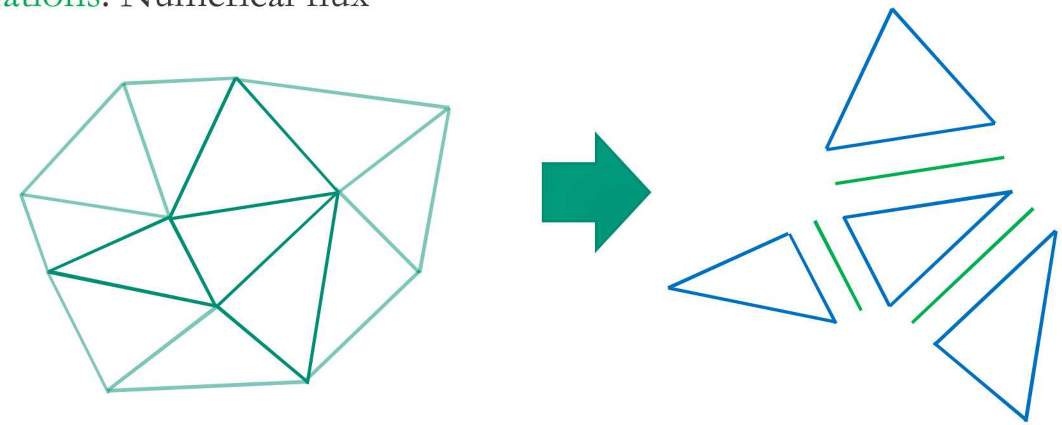

For IMEX-DG this means:

- **Explicit flux:** Take advantage of nearest-neighbor communication stencil.
  - Fast to compute, and scales as well as any explicit code.
  - Limited by explicit CFL.
- **Implicit source:** Block-diagonal Jacobian only couples degrees of freedom within a cell.
  - Much faster than a globally connected solve and supports perfect scaling.
  - No preconditioning required for small dense matrices on diagonal.

Early Success: Example two-fluid, unmagnetized plasma shock

Goal: Test shock capturing techniques for DG in the presence of strong implicit source terms.

- The electrostatic shock is based on the Sod-Shock problem for neutral Euler model.
- Stiffness is dominated by the plasma frequency.
- Converges to neutral/MHD shock in the limit of large plasma frequency and mass ratio.

Plasma Scales			
		Electrons	Ions
$\epsilon_0 = 10^{-4}$	$\omega_p \Delta t$	0.095	0.003
$\epsilon_0 = 10^{-8}$	$\omega_p \Delta t$	9.5	0.3
$\epsilon_0 = 10^{-16}$	$\omega_p \Delta t$	95000	3000
	$v_{se} \frac{\Delta t}{\Delta x}$	0.22	0.007

Plasma parameters

$m_i = 1$

$m_e = 10^{-3}$

$q_i = 1$

$q_e = -1$

$\gamma = 5/3$

Initial conditions

$n_L = 1$	$n_R = 0.125$
$u_L^e = u_L^i = 0$	$u_R^e = u_R^i = 0$
$P_L^e = P_L^i = 0.5$	$P_R^e = P_R^i = 0.05$
$E_L = B_L = 0$	$E_R = B_R = 0$

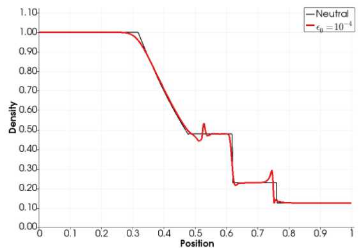

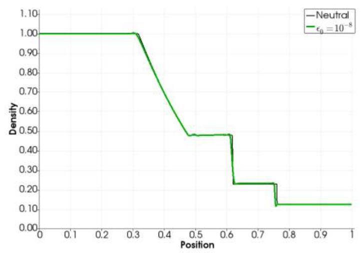

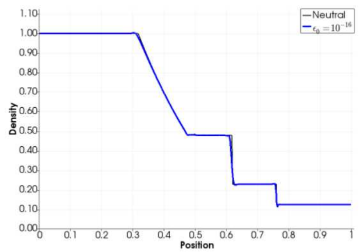

Result: Able to step over plasma frequency by arbitrary amount, and capture shock physics.

Problem: EMPIRE's PIC feature does not like a DG-based Maxwell solve.

DG-multi-fluid with exact-sequence-Maxwell's equations

EMPIRE uses PIC for kinetic physics, which requires Gauss' Laws to be satisfied.

- No magnetic mono-poles: $\nabla \cdot \boldsymbol{B} = 0$.
- Charge conservation: $\epsilon_0 \nabla \cdot \boldsymbol{E} = \rho_c$.

These involutions can be satisfied by using the proper discretization.

- Continuous FEM support exact-sequence formulations that preserve these properties[1]:

Discretization:
Fields continuous by construction (CG)

$$\begin{bmatrix} \rho_\alpha \\ \rho_\alpha u_\alpha \\ \mathcal{E}_\alpha \end{bmatrix} \in V_\nabla^h \qquad \boldsymbol{E} \in \boldsymbol{V}_{\nabla\times}^h \qquad \boldsymbol{B} \in \boldsymbol{V}_{\nabla\cdot}^h$$

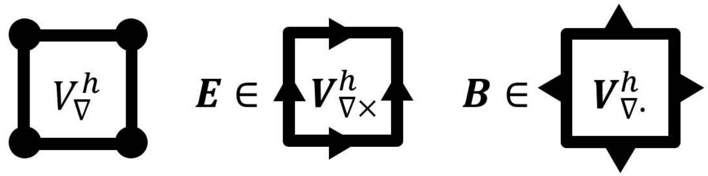

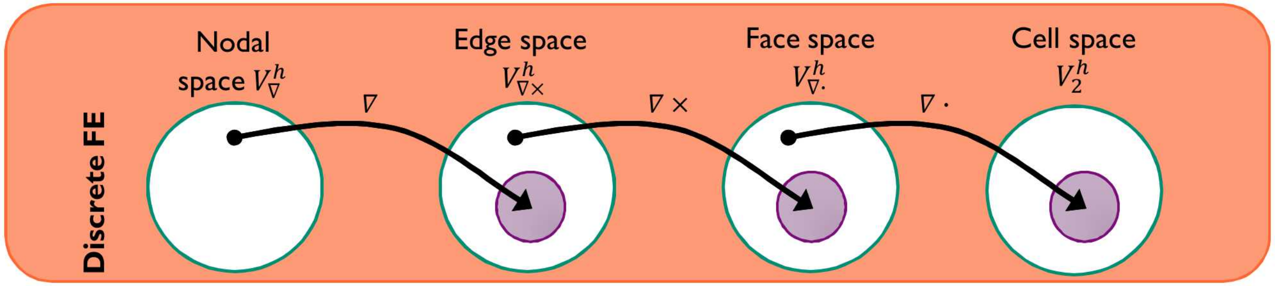

This approach has been shown to work for multi-fluid plasma models in a continuous setting (in the Drekar code[2]), however the extension to discontinuous systems is challenging.

[1] P. Bochev, H.C. Edwards, R.C. Kirby, K. Peterson, D. Ridzal. Solving PDEs with intrepid. *Scientific Programming* (2012).
[2] S.T. Miller, E.C. Cyr, J.N. Shadid, et-al., IMEX and exact sequence discretization of the multi-fluid plasma model. *JCP* (2019).

Divergence involutions when mixing exact-sequence and DG

A compatible finite element discretization is used to enforce the divergence constraints for the electric and magnetic fields.

- Fluids are represented by an **HGrad** (node) basis $\rho \in V_{\nabla}$.
- The electric field is represented by an **HCurl** (edge) vector basis $\mathbf{E} \in \boldsymbol{V}_{\nabla\times}$.
- The magnetic field is represented by an **HDiv** (face) vector basis $\mathbf{B} \in \boldsymbol{V}_{\nabla\cdot}$.
- Compatibility is defined by the discrete preservation of the **De Rham Complex**:

$$\nabla\phi_{\nabla} \in \boldsymbol{V}_{\nabla\times} \longrightarrow \nabla\times\boldsymbol{\phi}_{\nabla\times} \in \boldsymbol{V}_{\nabla\cdot} \longrightarrow \nabla\cdot\boldsymbol{\phi}_{\nabla\cdot} \in V_{\mathrm{L}_2}$$

For Faraday's law, we choose a basis for the electric field such that its curl is in the null space of the divergence operator.

$$\partial_t \boldsymbol{B} + \nabla\times\boldsymbol{E} = 0 \qquad \nabla\cdot\boldsymbol{B} = 0$$

Since the curl of the electric field is supported by the exact-sequence property, it's divergence must be zero in the strong form:

$$\nabla\cdot(\partial_t\boldsymbol{B} + \nabla\times\boldsymbol{E}) = \partial_t(\nabla\cdot\boldsymbol{B}) + \nabla\cdot\nabla\times\boldsymbol{E} = \partial_t(\nabla\cdot\boldsymbol{B}) + \sum_i E_i \underbrace{\nabla\cdot\nabla\times\boldsymbol{\phi}^i_{\nabla\times}}_{0} = \partial_t(\nabla\cdot\boldsymbol{B})$$

Result: The magnetic field remains divergence free after applying $\nabla\times\boldsymbol{E}$ flux.

- Charge conservation, where DG plays a role, is still under development.

Examining the exact-sequence-DG IMEX scheme

Fluid solve is block-diagonal.

- Small dense solves for each block, with no off-process communication.

Maxwell solver is efficient (and should remain unperturbed).

- Handles speed of light coupling.
- Comprised of many linear operators that can be computed once and reused.
- Algebraic-multigrid preconditioning is only updated on time step size changes.

Challenge: Coupling is characterized by plasma/cyclotron frequencies.

- Usually handled by preconditioning.
- These are local (ODE-like) coupling terms.

Goal: Attempt to construct a scheme that…

- Takes advantage of the local coupling in fluid operators.
- Handle plasma/cyclotron frequency coupling efficiently.
- Minimizes the number of re-computations required per nonlinear step.

Nonlinear system (broad strokes)

For IMEX we solve a nonlinear system:

- **Nonlinear terms:** Fluid Jacobian, Lorentz force
- **Linear terms:** Maxwell's equations, Current operator

The nonlinear system can be formulated from the linearized fluid model:

Newton Iteration

$$J_k \Delta x_k = -f(x_k)$$
$$x_{k+1} = x_k + \Delta x_k$$

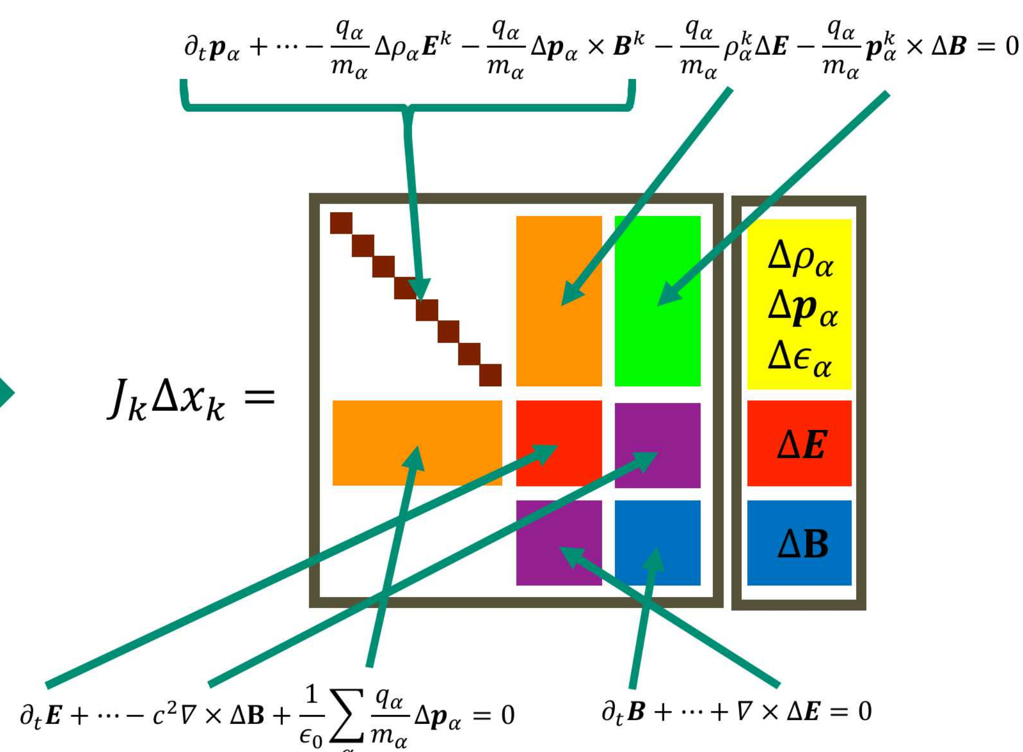

Nonlinear system simplification

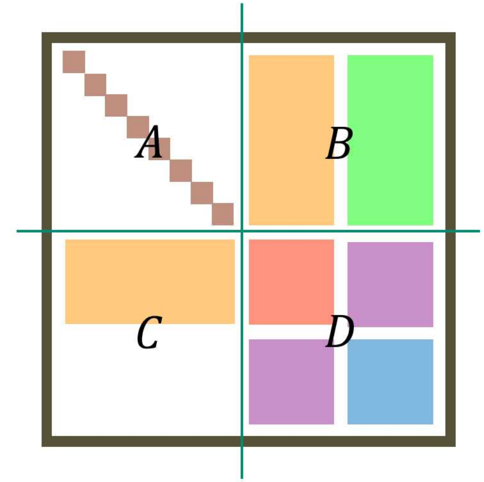

Instead of solving a full system we break it into sub-blocks:

Instead of Full-Newton, we try Quasi-Newton iteration.

$$J_k = \begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} I & BD^{-1} \\ 0 & I \end{bmatrix} \cdot \underbrace{\begin{bmatrix} A - BD^{-1}C & 0 \\ C & D \end{bmatrix}}_{\tilde{J}_k}$$

Looking at the operators, we can approximate the effect of the Shur complement term:

$$\left.\begin{aligned} \partial_t(\rho_\alpha \boldsymbol{u}_\alpha) &= \frac{q_\alpha}{m_\alpha}\rho_\alpha \boldsymbol{E} \\ \partial_t \boldsymbol{E} &= -\frac{1}{\epsilon_0}\frac{q_\alpha}{m_\alpha}\rho_\alpha \boldsymbol{u}_\alpha \end{aligned}\right\} \partial_t^2(\rho_\alpha \boldsymbol{u}_\alpha) + \omega_{p\alpha}^2(\rho_\alpha \boldsymbol{u}_\alpha) = 0$$

We apply the local Schur complement approximation to represent the plasma frequency:

$$\partial_t^2(\rho_\alpha \boldsymbol{u}_\alpha) + \omega_{p\alpha}^2(\rho_\alpha \boldsymbol{u}_\alpha) = 0$$

$$\underbrace{\frac{\rho_\alpha \boldsymbol{u}_\alpha}{\Delta t}}_{A} \Rightarrow \underbrace{\frac{\rho_\alpha \boldsymbol{u}_\alpha}{\Delta t} + \frac{q_\alpha}{m_\alpha \epsilon_0}\Delta t \sum_\beta \frac{q_\beta}{m_\beta}\rho_\beta \boldsymbol{u}_\beta}_{\tilde{A} \approx A - BD^{-1}C}$$

(Unlike the "analysis" above, we use the full current in the Schur complement correction)

Striking a balance with Newton method

Quasi-Newton method will converge slower than Full-Newton method.

- Lower triangular block solve:

Block Lower Triangular

$$\begin{bmatrix} \tilde{A} & 0 \\ C & D \end{bmatrix} \cdot \begin{bmatrix} f \\ m \end{bmatrix} = \begin{bmatrix} b_f \\ b_m \end{bmatrix}$$

→ Local Fluid Solve

$$\tilde{A} \cdot f = b_f$$

→ Global Maxwell Solve

$$D \cdot m = b_m - C \cdot f$$

Maxwell solve uses AMG-preconditioned solve using Trilinos::Belos.

- Step over speed of light by arbitrary factor.
- Preconditioner is generated using Trilinos::MueLu (refMaxwell) whenever time step changes.
  - Constant time step → only precondition once (useful optimization for PIC).

Fluid solve uses dense LU solve using Kokkos::Kernels.

- Step over collision/reaction time scales by arbitrary factor.
- Plasma/cyclotron frequency supported by Schur complement approximation.
- No preconditioning needed!

Result: More iterations than a Newton method, but faster iterations.

Building a fast electromagnetic solve

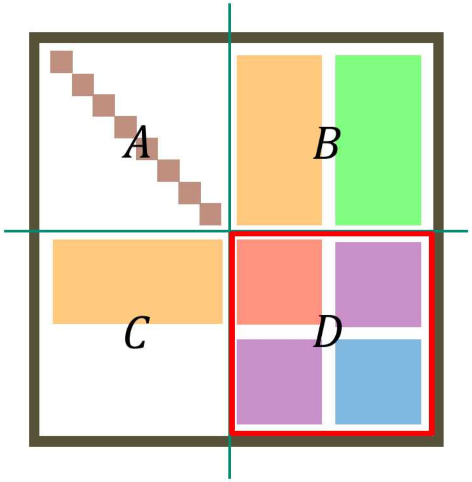

The electromagnetic solve also has a trick to it.

- Here we are solving another block LU factorized system:

$$D = \begin{bmatrix} M_E & C_E \\ M_B^{-1}C_B & I \end{bmatrix} = \begin{bmatrix} I & C_E \\ 0 & I \end{bmatrix} \cdot \begin{bmatrix} M_E - C_E M_B^{-1} C_B & 0 \\ M_B^{-1} C_B & I \end{bmatrix}$$

- Running through the triangular solves leaves us with a two-stage process:

$$\begin{bmatrix} M_E & C_E \\ M_B^{-1}C_B & I \end{bmatrix} \cdot \begin{bmatrix} \Delta\boldsymbol{E} \\ \Delta\boldsymbol{B} \end{bmatrix} = \begin{bmatrix} b_E \\ b_B \end{bmatrix} \quad \Rightarrow \quad \begin{aligned} (M_E + C_E M_B^{-1} C_B) \cdot \Delta\boldsymbol{E} &= b_E - C_E \cdot b_B \\ \Delta\boldsymbol{B} &= b_B - M_B^{-1} C_B \cdot \Delta\boldsymbol{E} \end{aligned}$$

- Note that the exact-sequence discretization has a couple of interesting consequences:
  - $M_B^{-1}C_B$ is sparse, making $M_E + C_E M_B^{-1} C_B$ and $M_B^{-1}C_B$ objects we can assemble and store (linear).
  - Curl operator $C_E = C_B^T$, making $M_E + C_E M_B^{-1} C_B$ symmetric → CG solve!
- Finally, we use algebraic multigrid[1] to precondition the solve for the electric field.
  - For small speed-of-light CFLs we can get away with Jacobi-preconditioning as an alternative.

[1]Bochev, Hu, Siefert, and Tuminaro. An Algebraic Multigrid Approach Based on a Compatible Gauge Reformulation of Maxwell's Equations. *SIAM* (2008).

Impact of Schur complement approximation

Test the Schur complement approximation using a linear plasma wave test.

- Still has strong growth in iteration count with increasing time steps
- Cost/benefit tradeoff study against Newton-Krylov with similar preconditioner required.

Example: O-wave

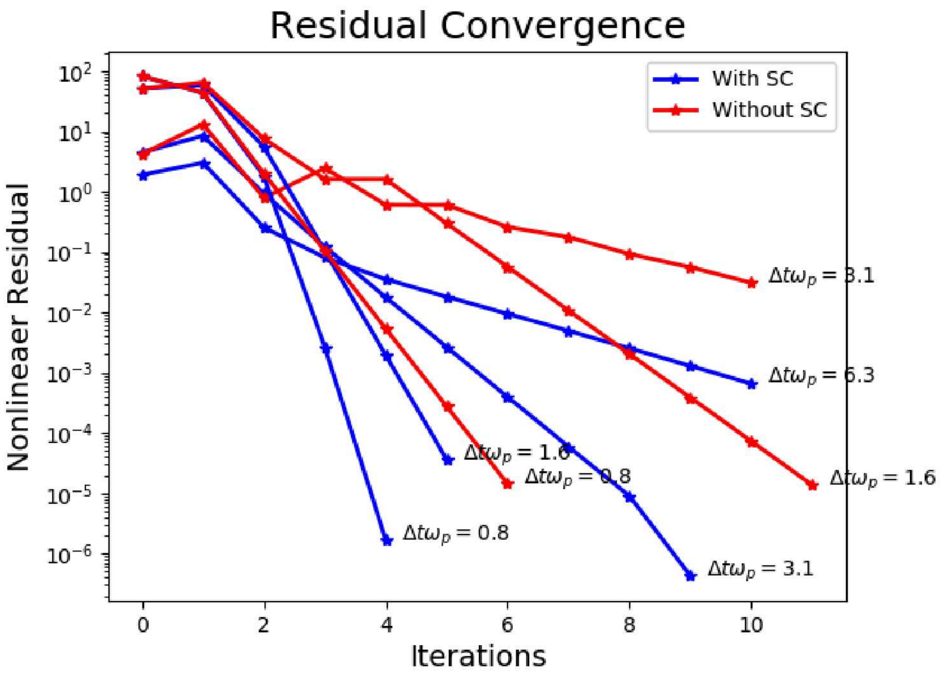

Result: Schur complement approximation is seen to improve nonlinear convergence by appreciable factor when stepping over plasma frequency.

One more trick: Anderson acceleration

Anderson Acceleration though Trilinos::NOX.

- Attempts to improve convergence by combining multiple nonlinear steps.
- Similar computation footprint to Quasi-Newton method.
- Fixed point around $x = g(x)$.
- Typically less complex to implement than a full Newton method.
- Walker and Ni, *SINUM* (2011): "Essentially equivalent" to GMRES

Algorithm AA: Anderson Acceleration

GIVEN x_0 AND $m \geq 1$.
SET $x_1 = g(x_0)$.
FOR $k = 1, 2, \ldots$ (UNTIL CONVERGED) DO:
  SET $m_k = min\{m, k\}$.
  DETERMINE $\gamma^{(k)} = (\gamma_0^{(k)}, \ldots, \gamma_{m_k-1})^T$ THAT SOLVES
  $min_{\gamma^{(k)}=(\gamma_0^{(k)},\ldots,\gamma_{m_k-1})^T} \|f_k - \mathcal{F}_k\|_2$.
  SET $x_{k+1} = g(x_k) - \mathcal{G}_k \gamma^{(k)}$.

$f_i = g(x_i) - x_i$
$\mathcal{F}_k = (\Delta f_{k-m_k}, \ldots, \Delta f_{k-1})$ with $\Delta f_i = f(x_{i+1}) - f(x_i)$.
$\mathcal{G}_k = (\Delta g_{k-m_k}, \ldots, \Delta g_{k-1})$ with $\Delta g_i = g(x_{i+1}) - g(x_i)$.

Anderson, *ACM* (1965)

Impact of Anderson acceleration: Two-fluid plasma vortex

Two fluid plasma vortex in MHD limit.

- Plasma vortex supported by combination of magnetic field (z-pinch) and pressure gradient.
- Using Schur complement approximation.

Convergence study:

- $N_x \times N_y \times N_z = [8 \times 8 \times 8, 16 \times 16 \times 16, 32 \times 32 \times 32]$
- $N_t = [10,20,40]$
- Speed of light: $\frac{c\Delta x}{\Delta t} = 8$

Nonlinear iteration study on $8 \times 8 \times 8$ grid:

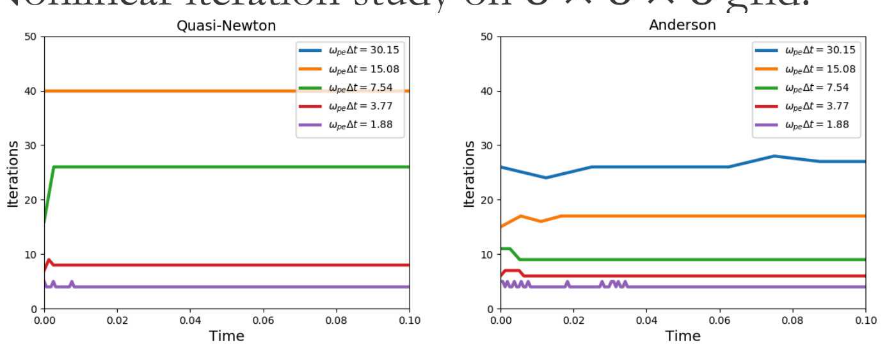

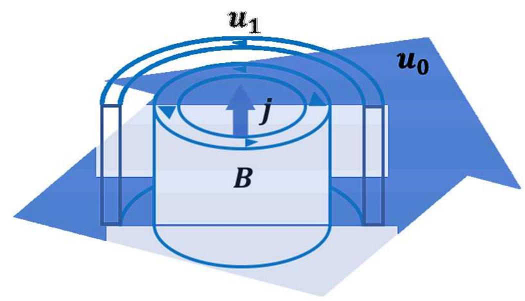

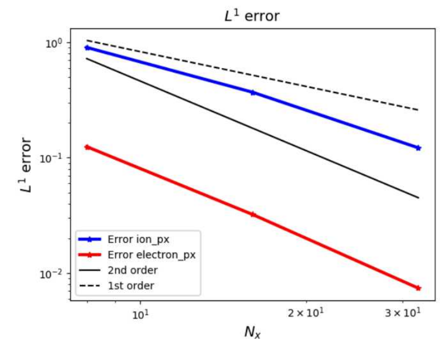

Result: Anderson acceleration shows improvement over Quasi-Newton method for iteration count (and runtime).

20

Two-fluid plasma wave convergence

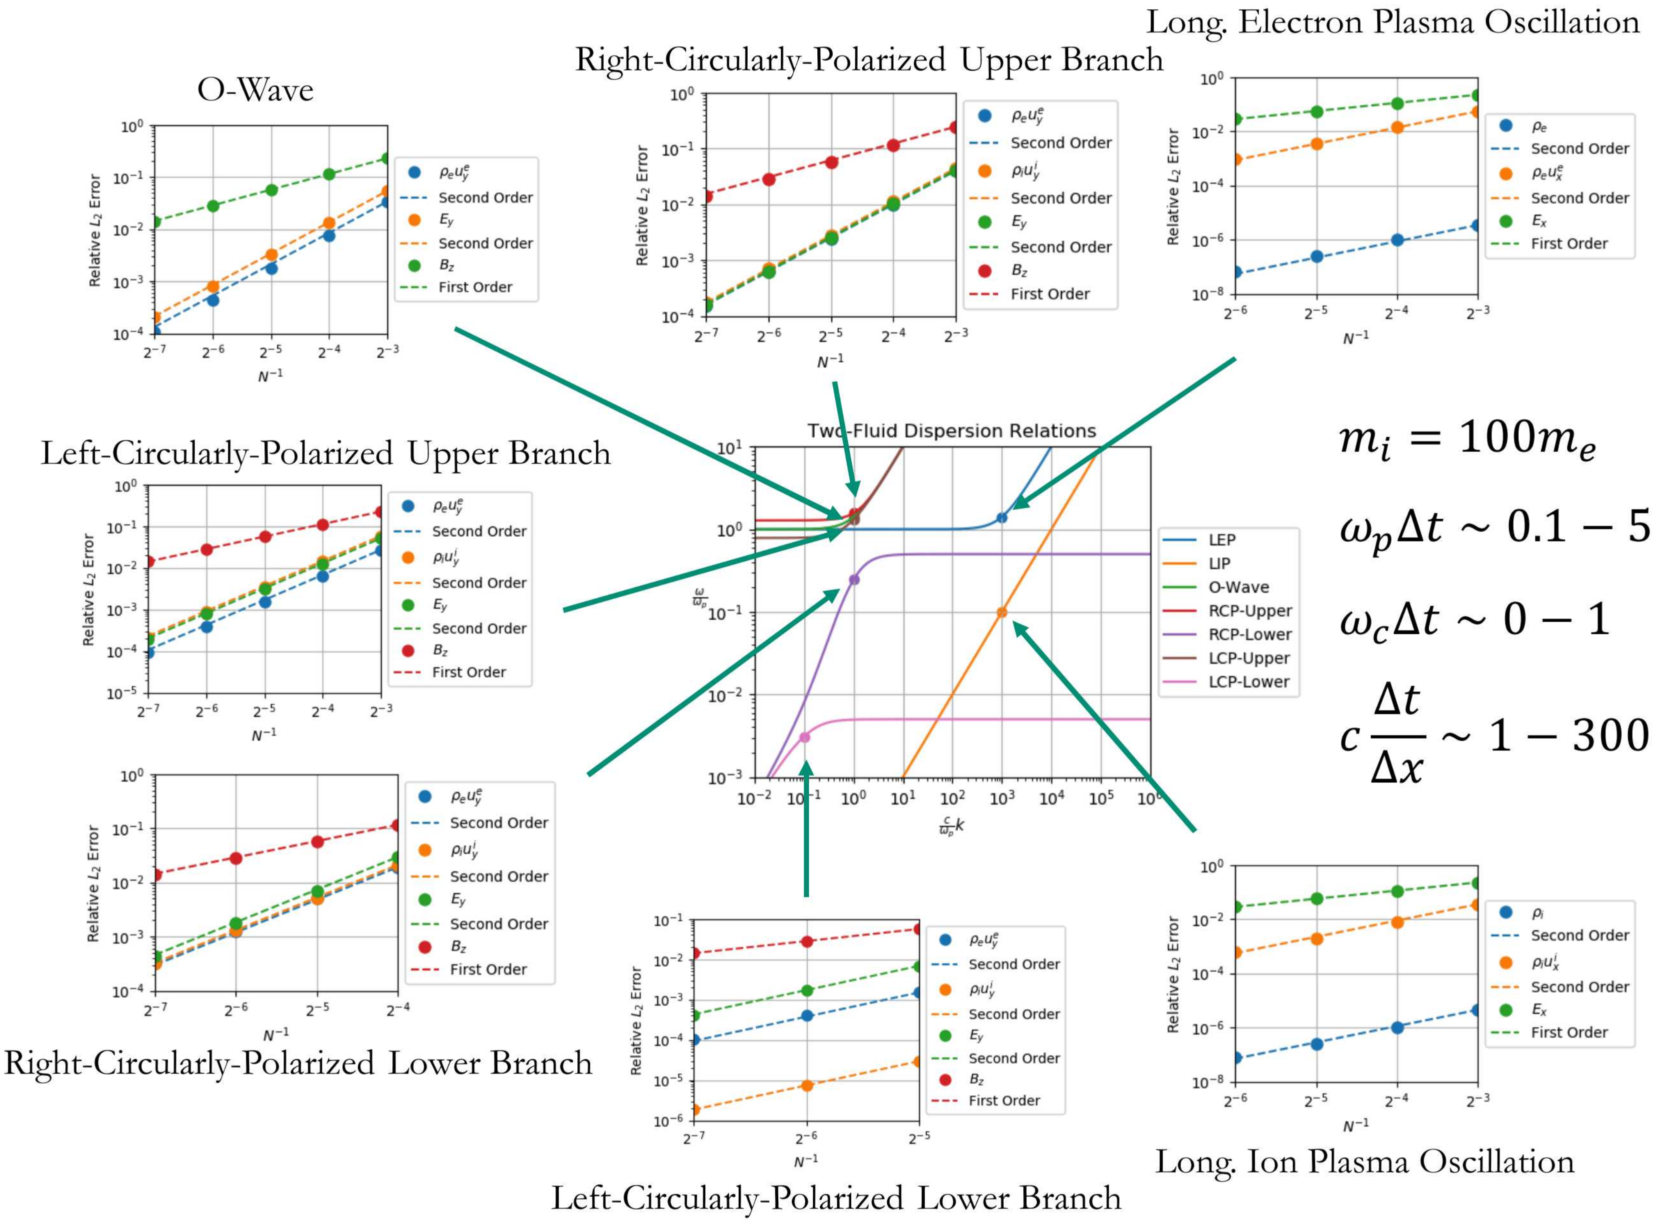

$$m_i = 100 m_e$$

$$\omega_p \Delta t \sim 0.1 - 5$$

$$\omega_c \Delta t \sim 0 - 1$$

$$c \frac{\Delta t}{\Delta x} \sim 1 - 300$$

Summary

We have given a brief overview of how the EMPIRE code approaches the multi-fluid plasma model:

- Fluid components discretized by Discontinuous-Galerkin method.
- Maxwell components discretized by exact-sequence discretization.
- Model is grouped into explicit and implicit components based on time scales.

Showed how we approach the nonlinear solver components:

- Nonlinear system is represented by Quasi-Newton method.
- Fluid solve is block-diagonal and local.
- Maxwell solve is a global, two-stage CG solve with AMG-based preconditioning.
- Schur complement approximation added to the fluid solve to represent plasma oscillations.
- Anderson acceleration is used to reduce the number of nonlinear iterations.

Improvement in nonlinear solver convergence was shown using Schur complement and Anderson acceleration.

Convergence for linear waves and two-fluid plasma vortex was shown when stepping over various time scales.