TexMex: A Coupled Transport-Equilibrium Solver Written in JAX for Gradient-Based Trajectory Optimization

Greta Hibbard*[1], Aaron Ho[1], Garud Snoep[1], Pablo Rodriguez-Fernandez[1], Nathan Howard[1]

[1]*Massachusetts Institute of Technology, Department of Nuclear Science and Engineering - 77 Massachusetts Avenue, Cambridge, MA 02139, USA,*

Questions? hibbard@mit.edu

Plasma Science & Fusion Center | Commonwealth Fusion Systems | Google DeepMind

Motivation

Trajectory optimization is important in reactor scenario design.

- Trajectory optimization necessitates auto-differentiable codes in order to use gradient-based optimization methods.

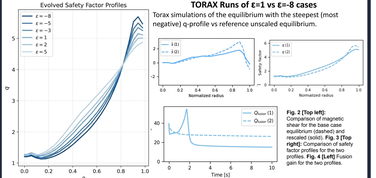


Fig. 2 [Top left]: Comparison of magnetic shear for the base case equilibrium (dashed) and rescaled (solid). **Fig. 3 [Top right]:** Comparison of safety factor profiles for the two profiles. **Fig. 4 [Left]:** Fusion gain for the two profiles.

Fig. 1 [Above]: Scaled example safety factor profiles at the end of ramp up. These profiles are evolved into steady state to compare performance.

Our previous work found that the safety factor profile exhibits path-dependent behavior (different trajectories result in different steady-state profiles). **The safety factor profile is therefore a strong candidate for trajectory optimization studies.**

Methods

Auto-Differentiation (AD) computes accurate and fast derivatives by applying the chain rule to each operation.



Fig. 5 example of AD method of computing derivatives via the chain rule [1].

- Fully AD-coupled systems allow derivatives to flow through the entire simulation.

TORAX[2] is a differentiable tokamak core transport model . This work utilizes TGLFNN.

- TORAX is a 1D transport solver within a 2D axisymmetric geometry used to evolve equilibria to steady state.

$$(1) \quad \left(\frac{\partial}{\partial t}\right)[n_e V'] = \frac{\partial}{\partial \hat{\rho}}\left[D_e n_e \frac{g_1}{V'}\frac{\partial n_e}{\partial \hat{\rho}} - g_0 V_e n_e\right] + V' S_n$$

MEXCAL[3] is a differentiable Grad-Shafranov equilibrium solver formulated in inverse flux coords.

$$(2) \quad \frac{\partial^2 \psi}{\partial r^2} - \frac{1}{r}\frac{\partial \psi}{\partial r} + \frac{\partial^2 \psi}{\partial z^2} = -\mu_0 r^2 \frac{dp}{d\psi} - \frac{1}{2}\frac{dF^2}{d\psi}$$

- MEXCAL was built to sit directly inside a transport solver's computational graph to compute exact gradients of any output with respect to any input.

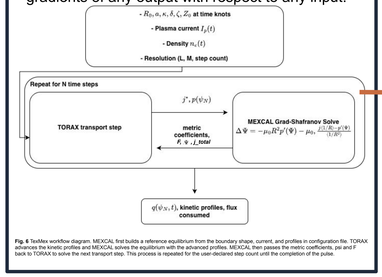


Fig. 6 TexMex workflow diagram. MEXCAL first builds a reference equilibrium from the boundary shape, current, and profiles in configuration file. TORAX advances the kinetic profiles and MEXCAL solves the equilibrium with the advanced profiles. MEXCAL then passes the metric coefficients, psi and F back to TORAX to solve the next transport step. This process is repeated for the user-declared step count until the completion of the pulse.

TexMex Trajectory Optimization Demonstration on SPARC and ARC

SPARC

Optimization Results

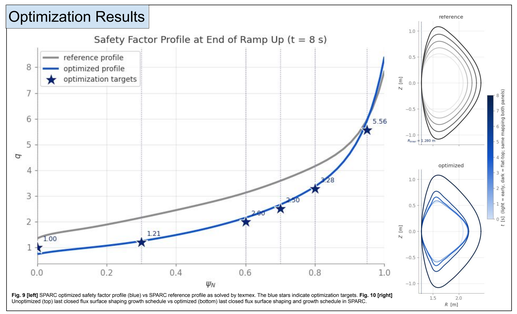


Fig. 9 [left] SPARC optimized safety factor profile (blue) vs SPARC reference profile as solved by texmex. The blue stars indicate optimization targets. **Fig. 10 [right]** Unoptimized (top) last closed flux surface shaping growth schedule vs optimized (bottom) last closed flux surface shaping and growth schedule in SPARC.

Constraints

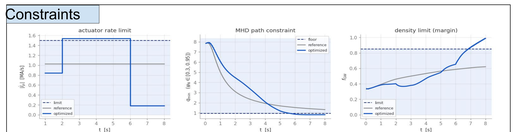


Fig. 11 SPARC optimized constraints vs reference constraints. The actuator rate limit stays below 1.2. MA/s, the q(0) path constraint is lower than the reference across the whole trajectory, and the Greenwald density limit is marginally violated at the end of the trajectory.

Actuators

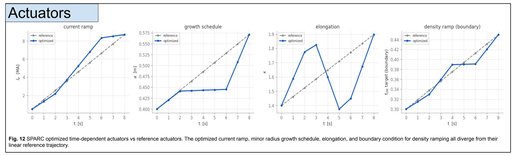


Fig. 12 SPARC optimized time-dependent actuators vs reference actuators. The optimized current ramp, minor radius growth schedule, elongation, and boundary condition for density ramping all diverge from their linear reference trajectory.

Kinetic Profiles

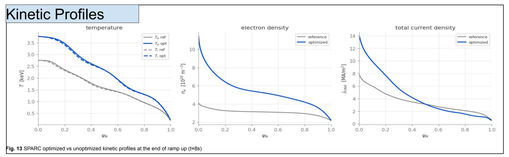


Fig. 13 SPARC optimized vs unoptimized kinetic profiles at the end of ramp up (H98).

Time Evolution

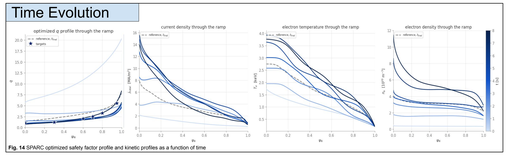


Fig. 14 SPARC optimized safety factor profile and kinetic profiles as a function of time

ARC

Optimization Results

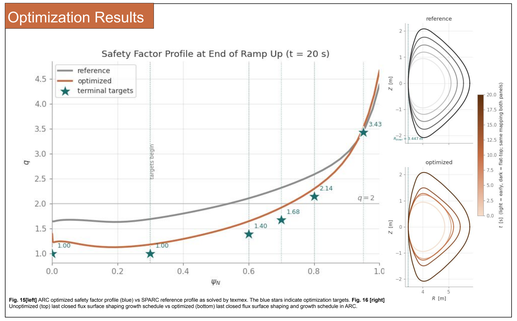


Fig. 15[left] ARC optimized safety factor profile (blue) vs SPARC reference profile as solved by texmex. The blue stars indicate optimization targets. **Fig. 16 [right]** Unoptimized (top) last closed flux surface shaping growth schedule vs optimized (bottom) last closed flux surface shaping and growth schedule in ARC.

Constraints

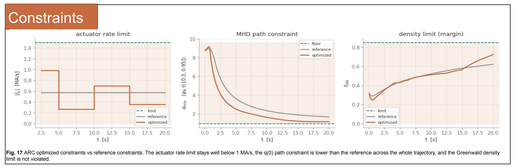


Fig. 17 ARC optimized constraints vs reference constraints. The actuator rate limit stays well below 1 MA/s, the q(0) path constraint is lower than the reference across the whole trajectory, and the Greenwald density limit is not violated.

Actuators

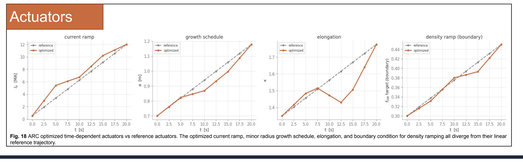


Fig. 18 ARC optimized time-dependent actuators vs reference actuators. The optimized current ramp, minor radius growth schedule, elongation, and boundary condition for density ramping all diverge from their linear reference trajectory.

Kinetic Profiles

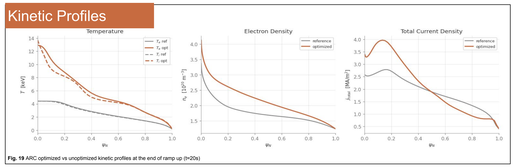


Fig. 19 ARC optimized vs unoptimized kinetic profiles at the end of ramp up (H=20s)

Time Evolution

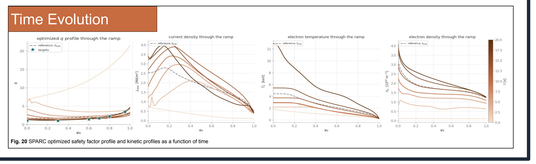
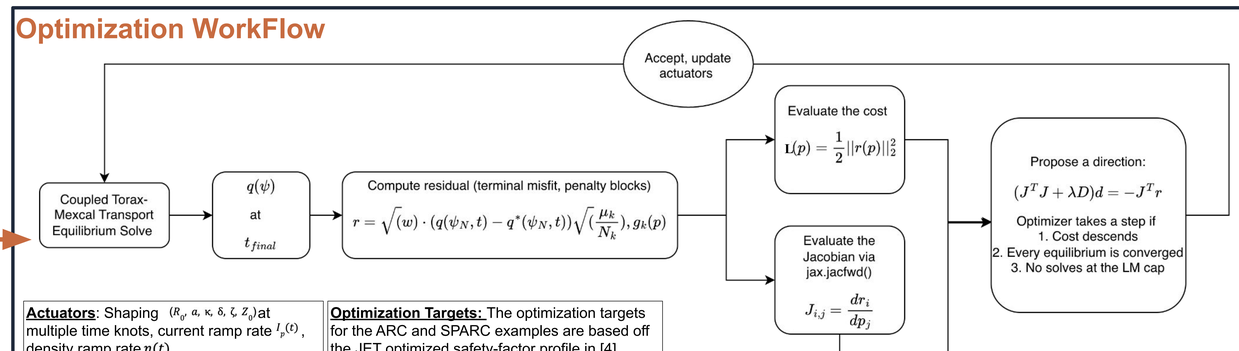


Fig. 20 SPARC optimized safety factor profile and kinetic profiles as a function of time

Optimization WorkFlow



Actuators: Shaping (R_0, a, κ, δ, ζ, Z_0)at multiple time knots, current ramp rate $I_p(t)$, density ramp rate $n(t)$.
Constraints: Actuator ramp limit, q(0) > 1 path constraint, density limit F_GW < 0.85.

Optimization Targets: The optimization targets for the ARC and SPARC examples are based off the JET optimized safety-factor profile in [4] as

$$(3) \quad q_{target,i} = q_{0,ref,i} + q_{JET} * q_{0,ref,i}$$

Fig 7: The outer optimization loop wrapped around the coupled TEXMEX solve. After each TexMex simulation the solved safety factor profile at t_final is used to compute the residual against the optimization targets, weighted with penalty blocks. The residual is used to evaluate the cost as well a the jacobian. The cost and Jacobian are used by the optimizer to propose the new set of input actuators based on gradient descent. This process is repeated for 20 iterations or until the cost converges.

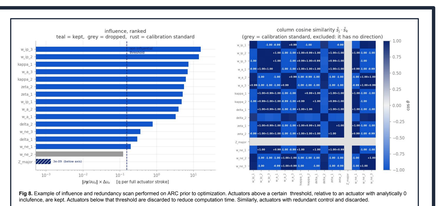


Fig 8. Example of influence and redundancy scan performed on ARC prior to optimization. Actuators above a certain threshold, relative to an actuator with analytically 0 incidence, are kept. Actuators below that threshold are discarded to reduce computation time. Similarly, actuators with redundant control are discarded.

- Before optimization, it is important to eliminate non-influential and redundant actuators to reduce computational cost.
- A non-influential actuator is defined as having influence less than a threshold relative to that of an actuator with analytically zero influence: $J_{j,k} = \frac{\partial y_j}{\partial u_k}$
- A redundant parameter is defined as having the dot product of the actuators be greater than 0.99:

$$\hat{s}_j \cdot \hat{s}_k > 0.99$$

Results

Actuator Movement in ARC vs SPARC

Open circle = reference, filled = optimized, bold rows = moved most in both cases, diamond = boundary

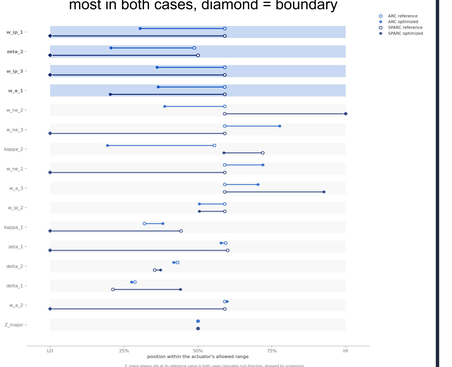


Fig. 21 Summary of optimized actuator changes in ARC and SPARC relative to the reference case.

- Both optimizers find changes in the current and density ramping from the linear reference case to be favorable.

Time-dependent shaping and ramping parameters have successfully been optimized towards target safety factor profiles, resulting in higher temperature and density profiles compared to the reference case at the end of ramp-up in ARC and SPARC.

Future Directions

- With the **x-point solver extension** now enabled on MEXCAL, this example will be redone with the full resolution LCFS including x-points.
- At L = M = 16 running locally on a 32 GB RAM Macbook, one forward pass of TexMex takes 104s to simulate 8s of physics time. At this speed, TexMex optimization and simulation could be used in the control room. There is room for **additional speed optimization and parallelization.**
- **Improved engineering constraints** can be imposed on the optimization routine, eg feasibility of coils creating intermediary ramp up geometries.
- Perform **benchmarking** with other AD equilibrium codes.
- **This routine enables investigation of applying trajectory optimization for MHD mode avoidance.**

Acknowledgements

This work is supported by Google Deepmind and Commonwealth Fusion Systems (RPP020). The authors extend their thanks to A. O. Nelson and F. Sheehan at Columbia University for their help using TMTX in the early stages of this project.

Sources

[1] A Gentle Introduction to Automatic Differentiation, Pablo Monteagudo Lago. Accessed 19 Oct. 2026. [2] Citrin et al., TORAX: A Fast and Differentiable Tokamak Transport Simulator in JAX, arXiv:2406.06718 (2024). [3] Snoep, G., MEXCAL, GitHub, https://github.com/gsnoep/mexcal [4] H. J. Hobirk et al., "Improved confinement in JET hybrid discharges,"