Stellarator optimization

Where do these shapes come from?

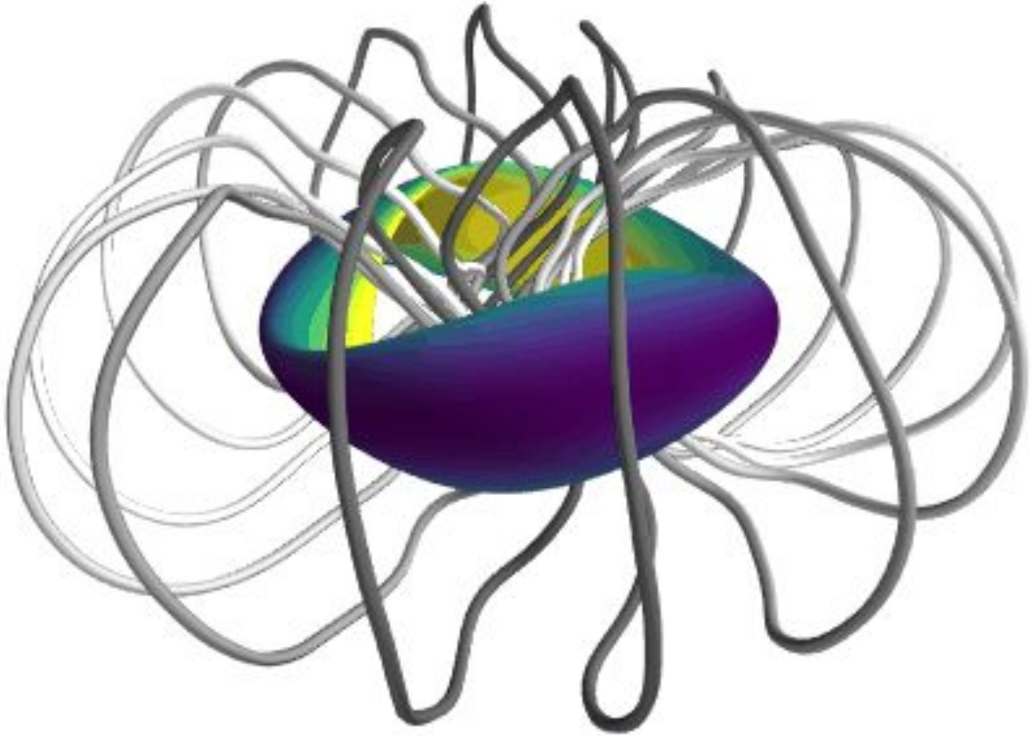

Alan Kaptanoglu – alankaptanoglu@nyu.edu, New York University
Matt Landreman – mattland@umd.edu, University of Maryland

Outline

- Optimization in general
- Optimization for good flux surfaces
- Quasi-axisymmetry, quasi-helical symmetry, and quasi-isodynamic
- ~~Optimization for stability & turbulence~~
- Coil optimization: current potential and filament methods

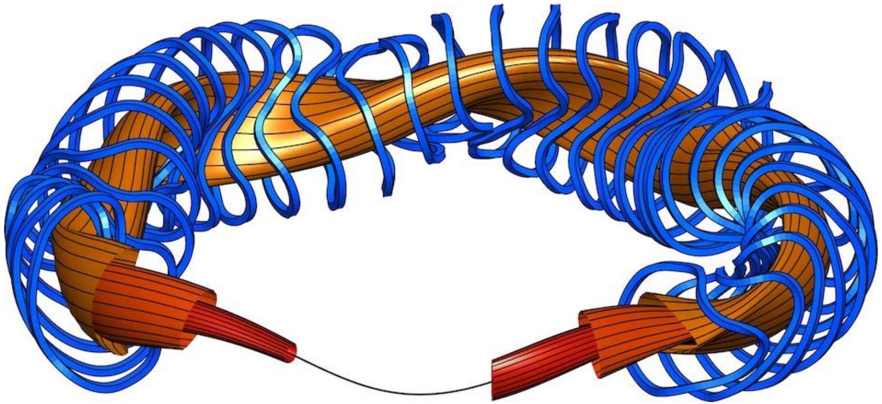

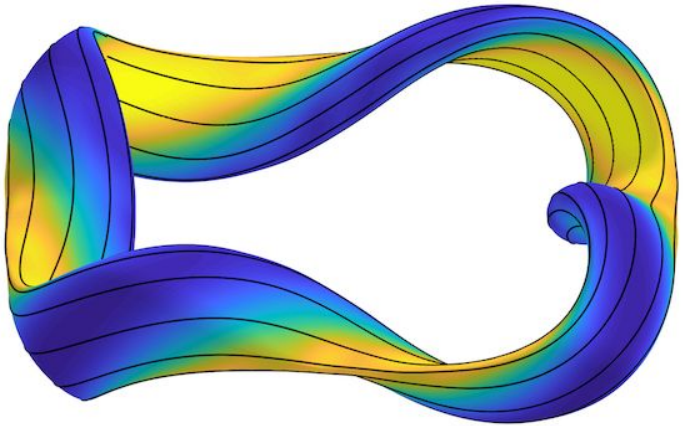

Optimization is a general technique with many applications

Given a "cost function" $f: \mathbb{R}^n \rightarrow \mathbb{R}$, minimize $f(\boldsymbol{x})$

(a.k.a. "objective function", "loss function")

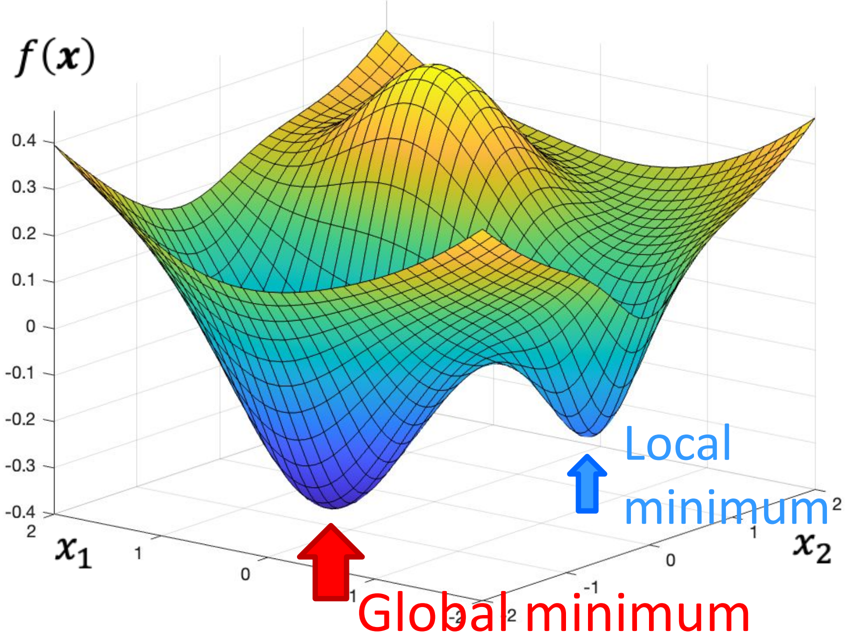

Two important ideas:

1. If f has multiple minima, we say it is "**nonconvex**".
2. Most of optimization is basically "gradient descent" – move in the $-\nabla f$ direction with some step size α.

Optimization is a general technique with many applications

Given a "cost function" $f: \mathbb{R}^n \to \mathbb{R}$,
(a.k.a. "objective function", "loss function")

$$\text{minimize } f(\boldsymbol{x})$$
$$\text{subject to } g_i(\boldsymbol{x}) = 0,$$
$$h_j(\boldsymbol{x}) < 0$$

(Optional) Can add constraints:

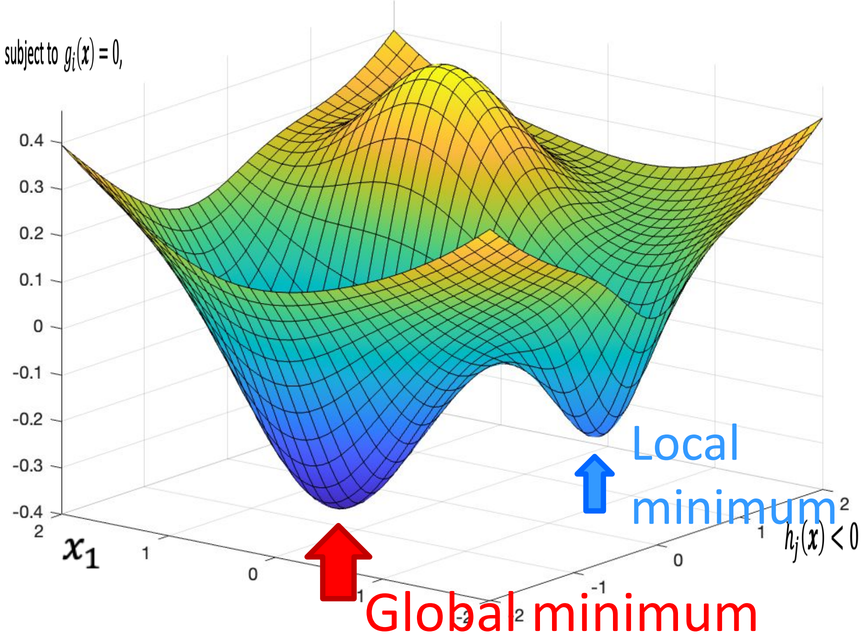

Suppose we want to minimize > 1 quantity, e.g. f_1, f_2, and f_3.

"Scalarization":
minimize $f = w_1 f_1 + w_2 f_2 + w_3 f_3$
where w_j are weights.

Common trick: turn "hard constraint" into "soft constraint": minimize$[f(\mathbf{x}) + w|g(\mathbf{x})|^2]$

Optimization is a general technique with many applications

Given a "cost function" $f: \mathbb{R}^n \to \mathbb{R}$, minimize $f(\boldsymbol{x})$

(a.k.a. "objective function", "loss function")

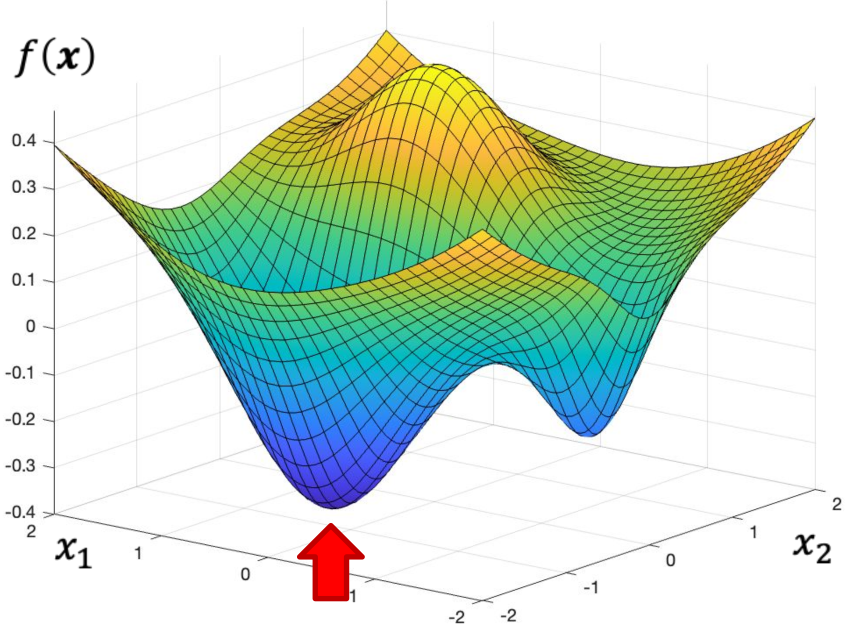

(Optional) stochastic optimization to control for random perturbations:

$$\min_{\mathbf{x}} \mathbb{E}\left[g(\mathbf{x}, \zeta)\right]$$

Why do we need optimization at all?

- Most nuclear fusion devices are planned to be (quasi) steady-state, so to first-order, they are designed as an MHD equilibrium.
- There are lots of physical quantities (nested flux surfaces, rotational transform, quasi-symmetry, linear growth rates, …) we want to prescribe/extremize to achieve a very effective nuclear fusion device, but these quantities are determined by the equilibrium MHD equations: $\mathbf{J} \times \mathbf{B} = \nabla \mathbf{p}$
- However, the solution to these equations depend on boundary conditions. Thus, one of the big advantages of stellarators is that the 3D shape of the plasma can be designed to achieve these goals.

Some stellarator parameters are integers. These are mostly optimized by hand.

- Number of “field periods”.
- Number of coils.

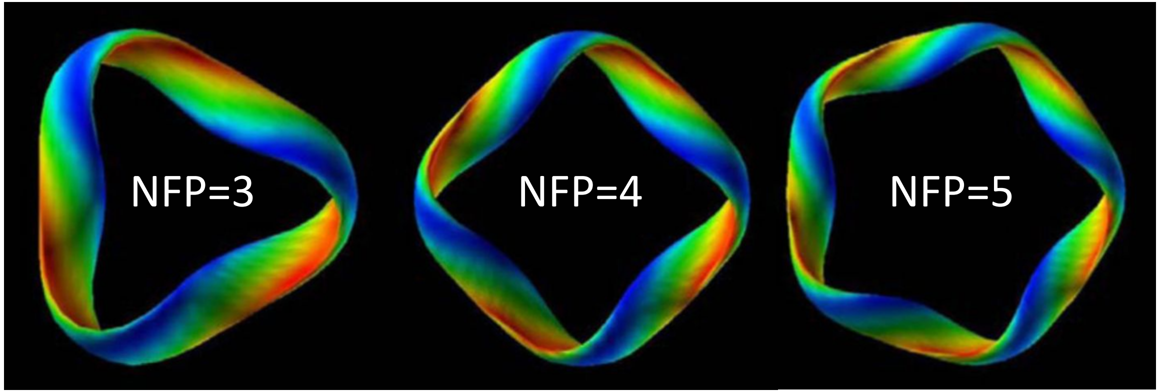

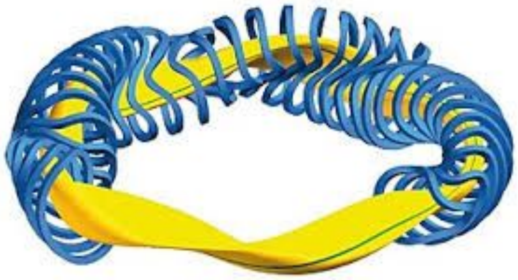

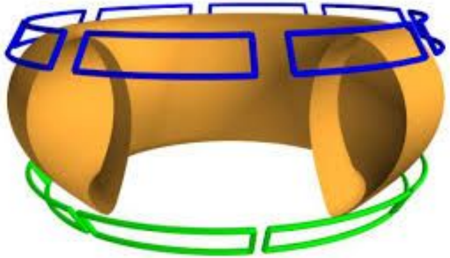

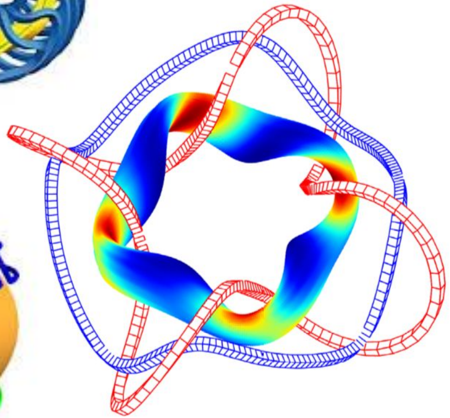

- Do coils link the plasma poloidally, helically, or not at all?
- Do *B* contours link the torus toroidally (QA), helically (QH), or poloidally (QI)?

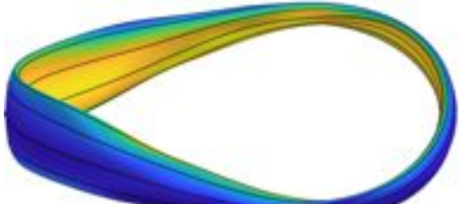

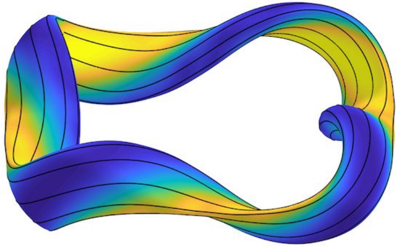

There are several possible choices of parameter space to optimize in

Minimize $f(\boldsymbol{x})$ where $\boldsymbol{x} \in \mathbb{R}^n$. What is $\boldsymbol{x}$?

$\boldsymbol{x}$ = shapes of coils:

- Much of this space is very bad: not good flux surfaces, can't evaluate physics objectives
- "Free boundary equilibrium" calculations can be fragile

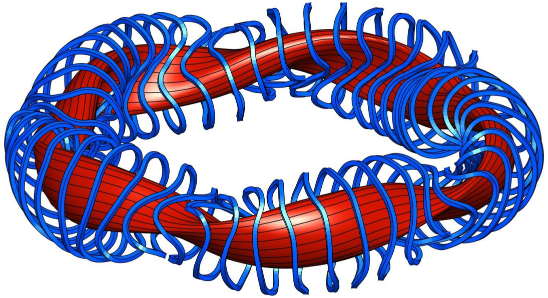

$\boldsymbol{x}$ = shape of plasma:

- Can use "fixed boundary equilibrium", which is very reliable.
- Some of parameter space corresponds to unphysical self-intersecting shapes.
- Otherwise, most of this parameter space is good.
- But the plasma shape is not what you build – you build coils.

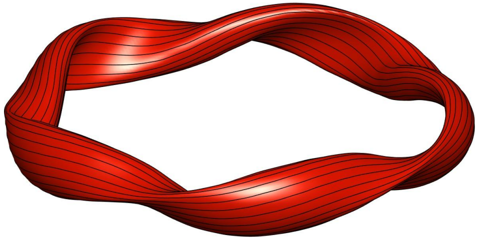

Other options:

$\boldsymbol{x}$ = shape of plasma *and* coils, $\boldsymbol{x}$ = shape of magnetic axis, …

Most transport-optimized stellarators have used 2 sequential optimization stages

1. Parameters = shape of boundary toroidal surface. Objective = physics (confinement, stability, etc.)

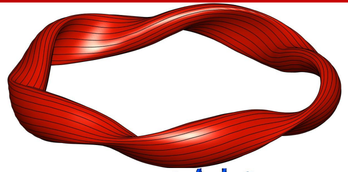

2. Parameters = coil shapes.
Objective = error in **B** on boundary shape from stage 1.

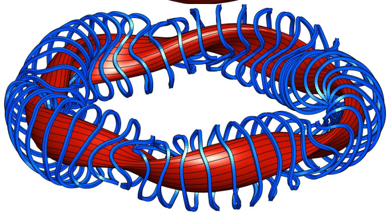

Shape of a toroidal boundary surface (+ pressure & current vs *r* inside, & total **B** flux) determines **B** everywhere inside:

$$\text{Consider a low-pressure plasma so } 0 \approx \mathbf{J} = \nabla \times \mathbf{B} \quad \Rightarrow \quad \mathbf{B} = \nabla \Phi.$$

$$\nabla \cdot \mathbf{B} = 0 \quad \Rightarrow \quad \nabla^2 \Phi = 0.$$

$$\mathbf{B} \cdot \mathbf{n} = 0 \;\text{ on boundary} \quad \Rightarrow \quad \mathbf{n} \cdot \nabla \Phi = 0.$$

$$\Rightarrow \quad \text{Laplace's eq with Neuman condition.}$$

$$\Rightarrow \quad \text{Unique solution up to scale factor + constant.}$$

Most transport-optimized stellarators have used 2 sequential optimization stages

1. Parameters = shape of boundary toroidal surface. Objective = physics (confinement, stability, etc.)

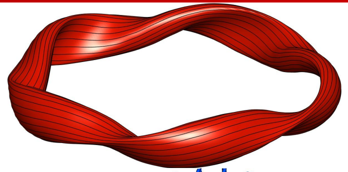

2. Parameters = coil shapes.
   Objective = error in **B** on boundary shape from stage 1.

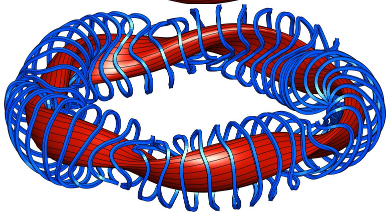

W7-X (Germany)

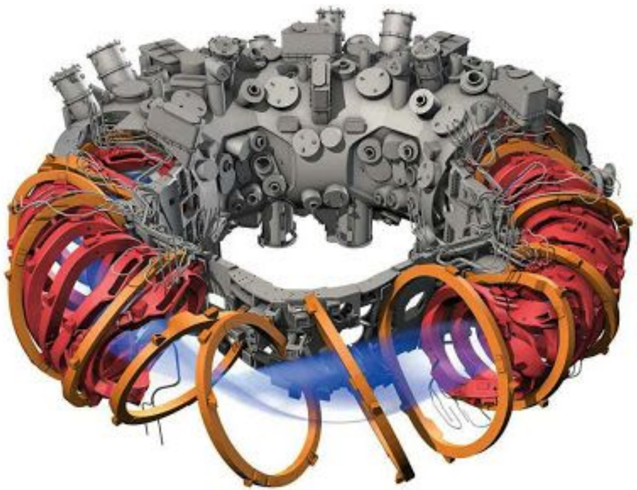

CFQS (China), under construction

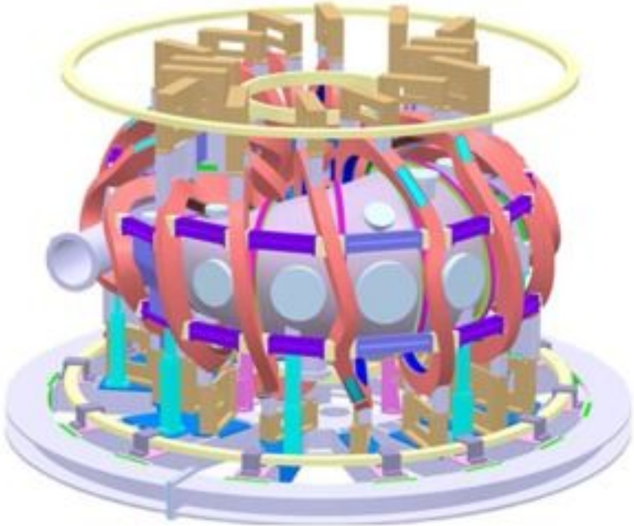

NCSX (Princeton)

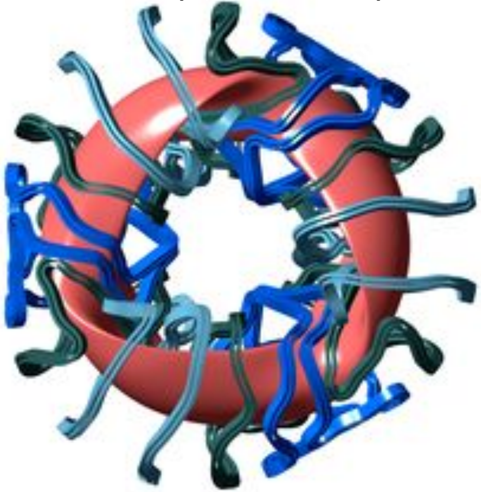

Outline

- Optimization in general
- Optimization for good flux surfaces
- Quasi-axisymmetry, quasi-helical symmetry, and quasi-isodynamic
- Coil optimization: current potential and filament methods

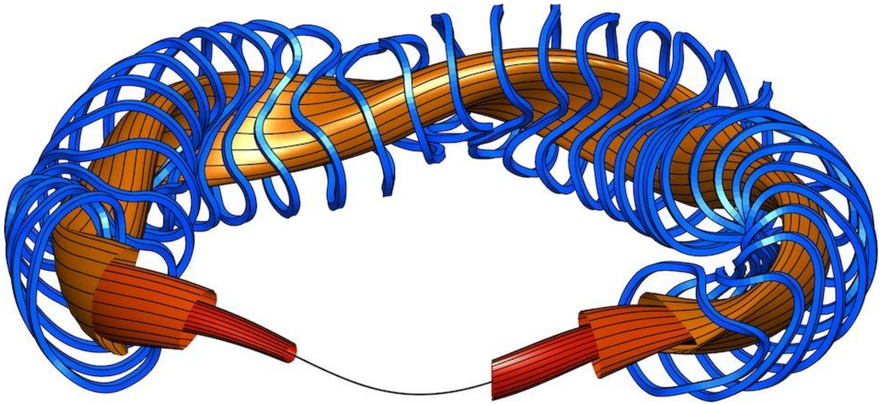

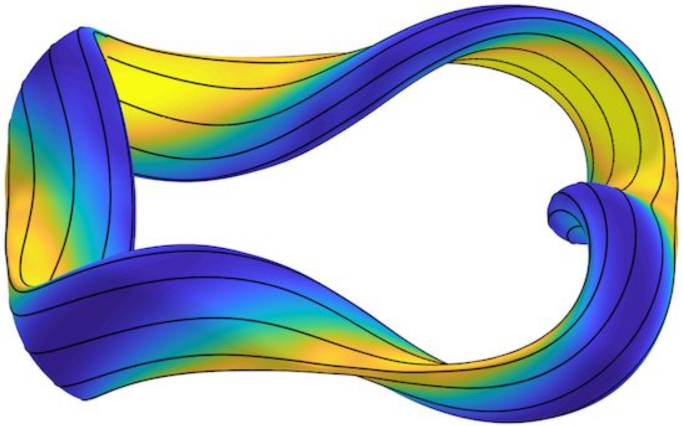

One goal of stellarator optimization is having field lines lie on surfaces.

Chaotic (volume-filling) **B** field lines would allow inside & outside to mix even without cross-**B** drift.

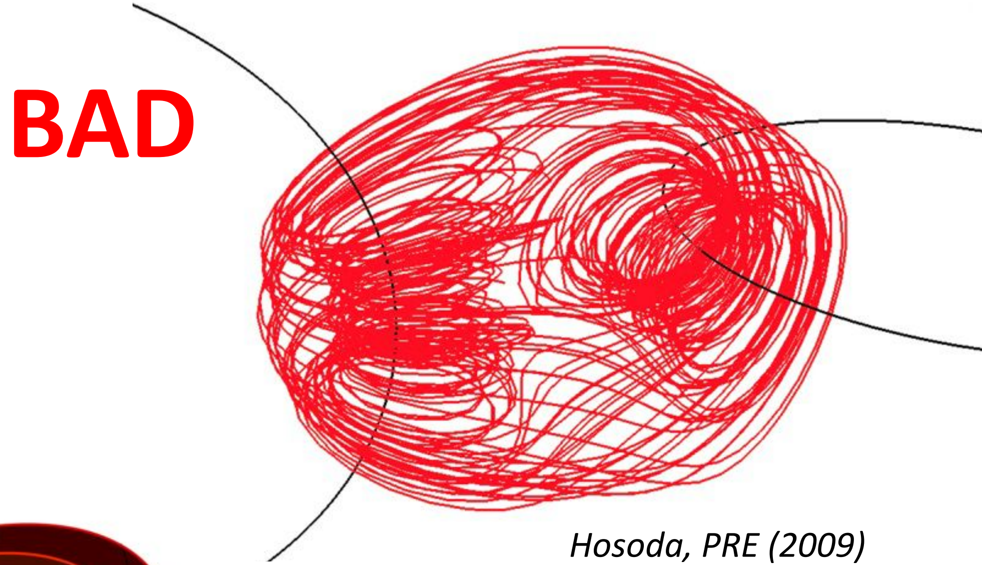

Hosoda, PRE (2009)

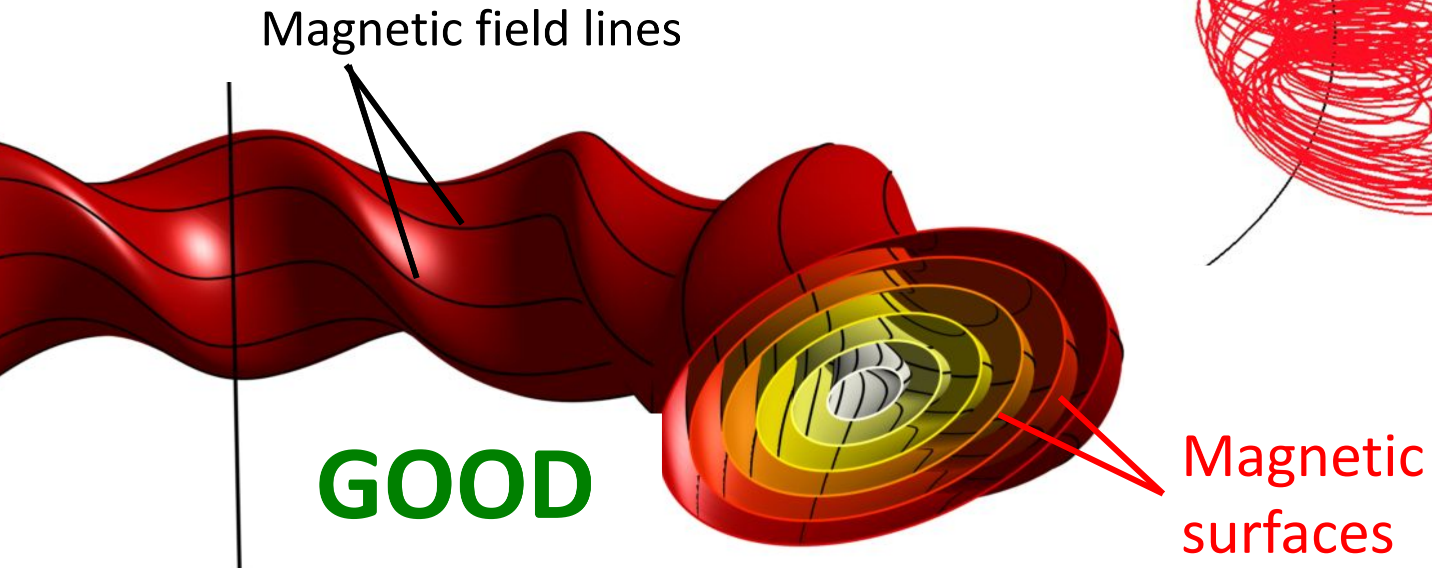

One goal of stellarator optimization is having field lines lie on surfaces.

Magnetic surfaces (a.k.a flux surfaces) can be visualized with a "Poincare plot":

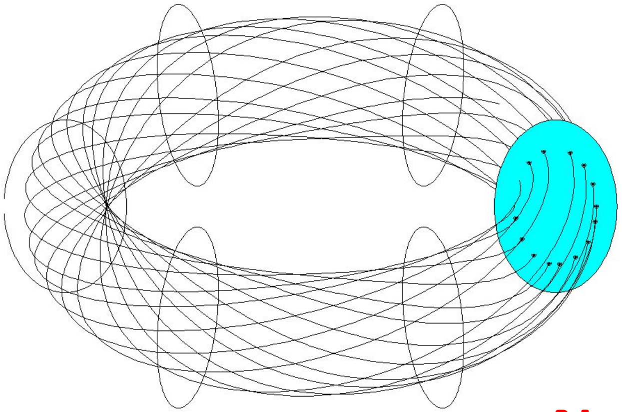

Good

Not so good

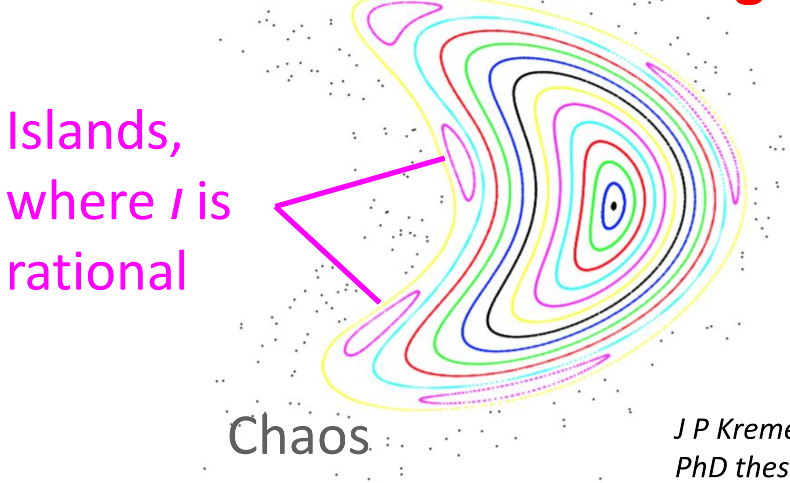

J P Kremer, PhD thesis, Columbia

How much rotational transform do you want?

Avoid rationals like ι = 1 or ½: islands form there.

So, maybe want low “magnetic shear” = $|\nabla\iota|$.

Or, maybe want *high* magnetic shear since it makes islands thin. (width $\propto |\nabla\iota|^{-1/2}$)

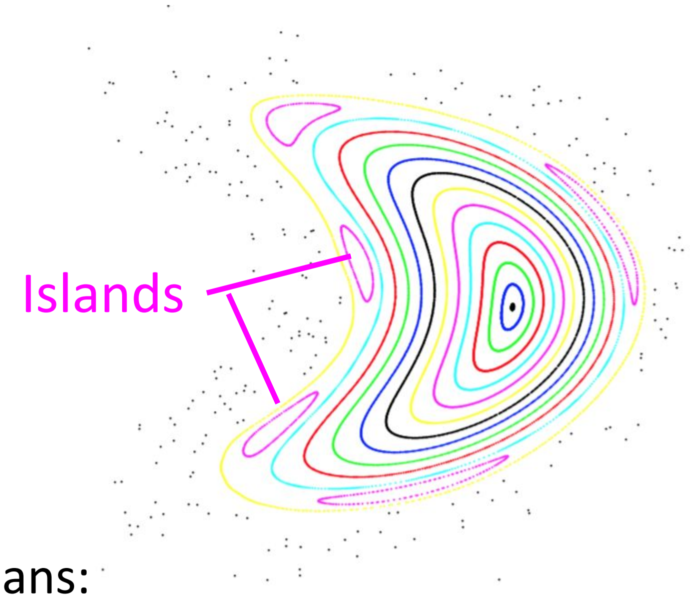

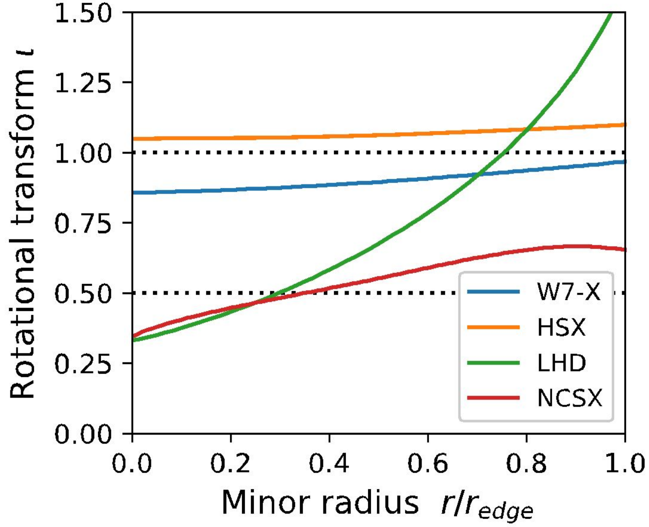

Larger ι means:

- Thinner orbits, so better confinement.
- **B** changes less due to plasma current. (Higher “equilibrium β limit”.)
- But, more complicated coils.

Outline

- Optimization in general
- Optimization for good flux surfaces
- Quasi-axisymmetry, quasi-helical symmetry, and quasi-isodynamic
- Coil optimization: current potential and filament methods

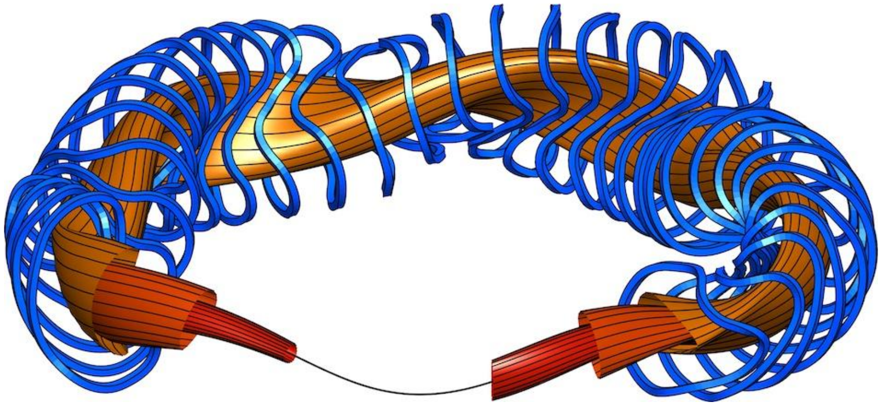

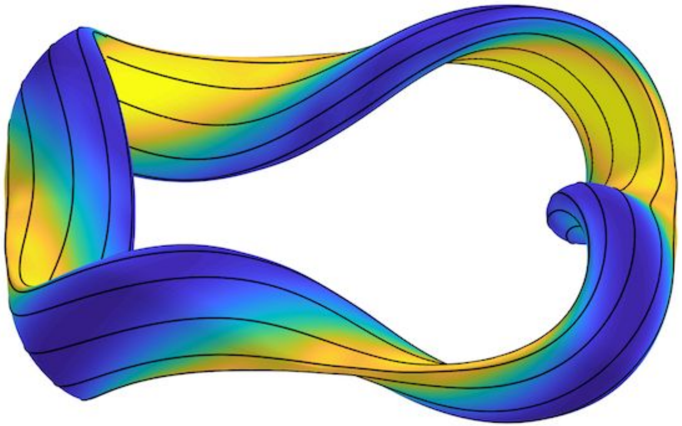

For low neoclassical transport, recent stellarators have come in 3 flavors

- Trapped particles should drift toroidally, helically, or poloidally on a surface.
- *B* contours on a surface have the same topology as these drifts.

Toroidal: “QA” = Quasi-axisymmetric

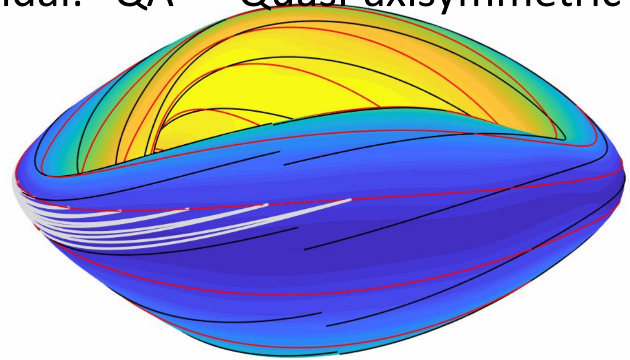

Field lines
|B| contours (slightly idealized)
Trapped particle

Helical: “QH”= Quasi-helically symmetric

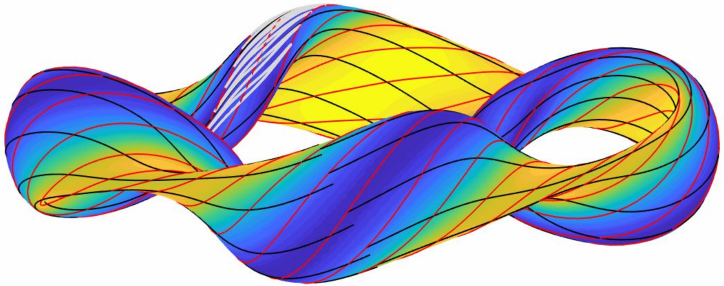

Poloidal: “QI”= Quasi-isodynamic

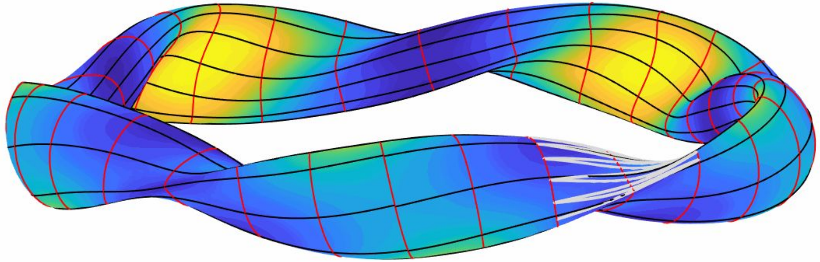

There has been great progress recently in optimizing stellarator neoclassical confinement

Trajectories of fusion-produced distribution of alpha particles followed in many magnetic configurations, all scaled to reactor size and |B|:

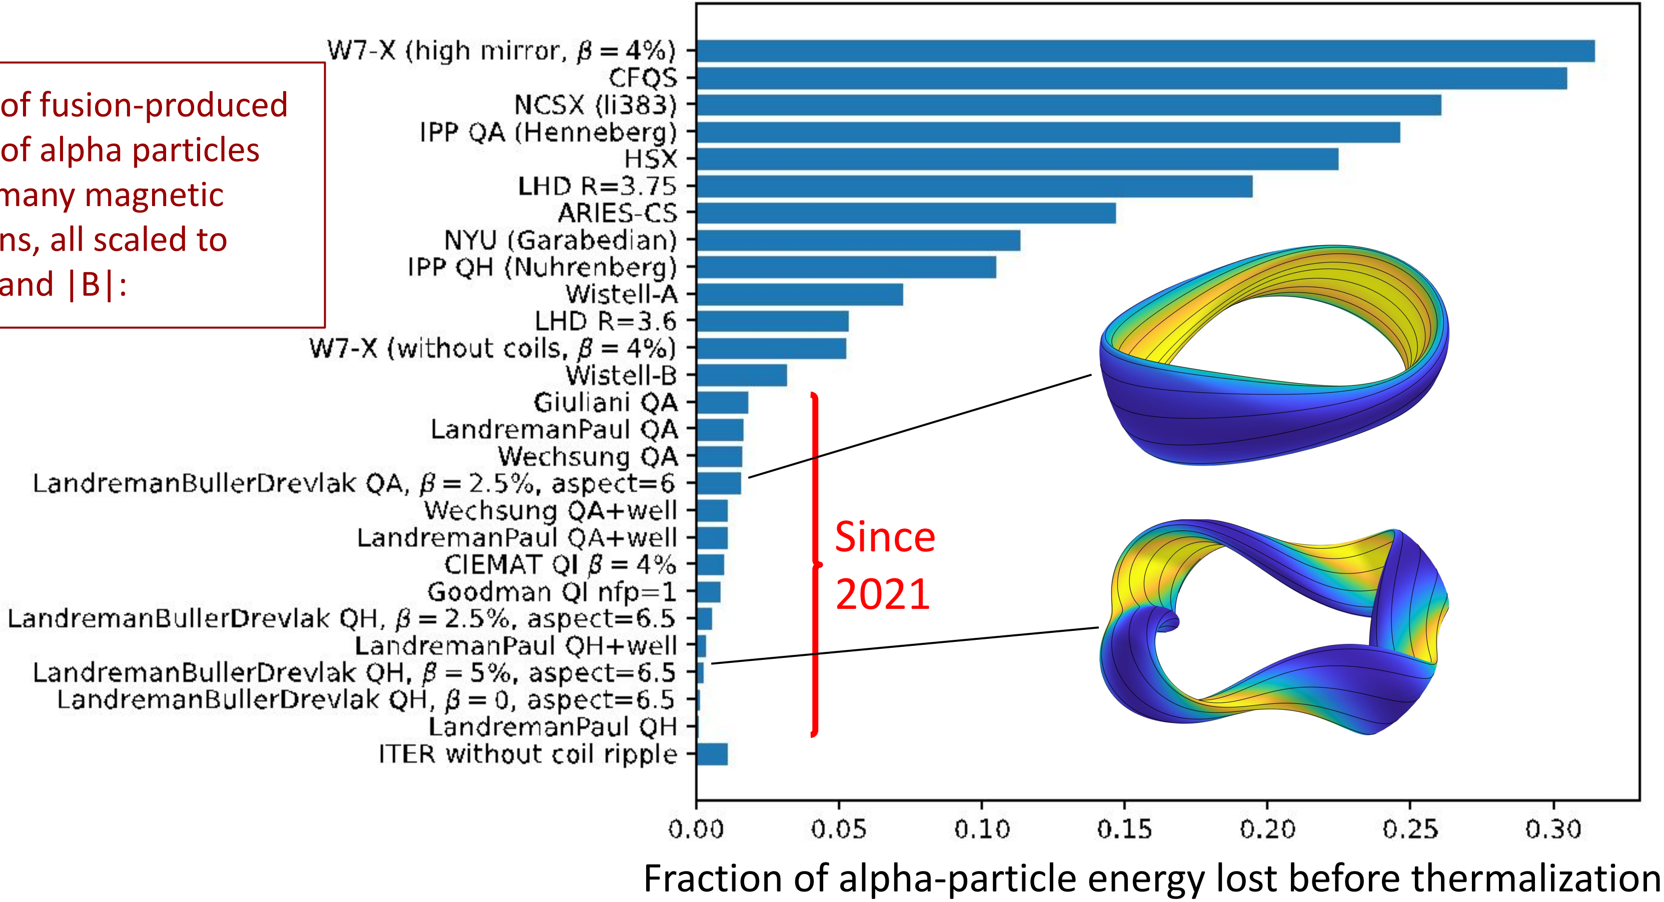

Landreman et al, Phys Plasmas (2022).

The parameter space for stage-1 optimization is typically a set of Fourier amplitudes for the boundary surface in cylindrical coordinates

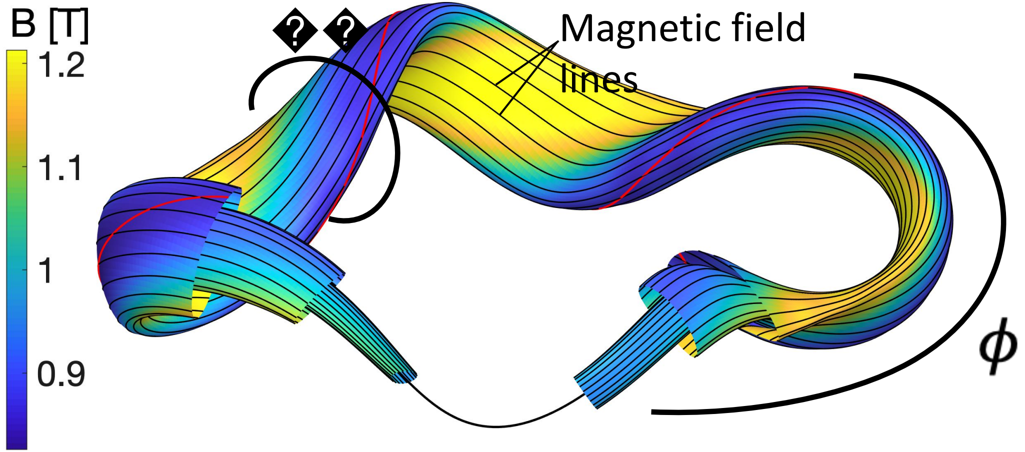

Parameterization of boundary surface:

$$R(\theta,\phi) = \sum_{m,n=-5}^{5} R_{m,n} \cos(m\theta - n_{fp} n\phi) \qquad Z(\theta,\phi) = \sum_{m,n=-5}^{5} Z_{m,n} \sin(m\theta - n_{fp} n\phi)$$

ϕ = standard cylindrical angle, n_{fp} = number of field periods

Parameter space for optimization: $\boldsymbol{x} = [R_{m,n}, Z_{m,n}]$

Notice there are (2M+1)(2N + 1) optimization variables (free parameters)

Quasisymmetry is a sufficient (though not necessary) condition for confinement, & a useful surrogate

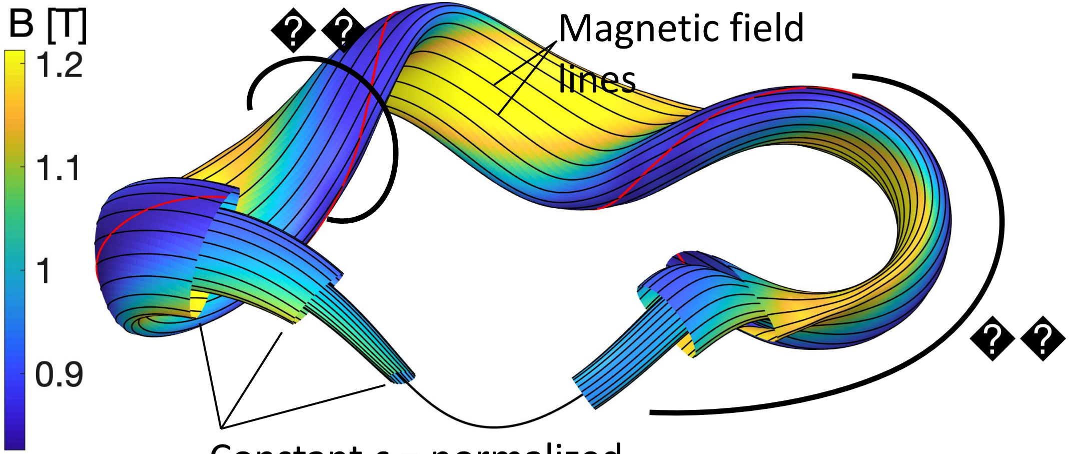

Constant s = normalized toroidal flux

$$B = B(s, \theta - N\varphi)$$

"Boozer angles"

Objective: $$f_{QS} = \int d^3x \left(\frac{1}{B^3} \left[(N - \iota)\boldsymbol{B} \times \nabla B \cdot \nabla \psi - G \boldsymbol{B} \cdot \nabla B \right] \right)^2$$

For quasi-axisymmetry, $N = 0$.

For quasi-helical symmetry, N is the number of field periods,

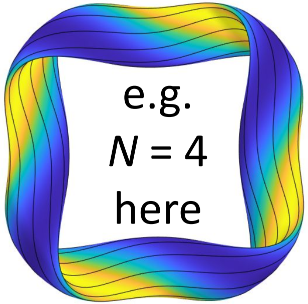

$$f_{QH} = \left(A - A_*\right)^2 + f_{QS}$$

Boundary aspect ratio

So what is the explicit optimization problem?

Inputs:

- two flux functions, $p(\psi)$ and $I_T(\psi)$
- an MHD equilibrium code
- an optimization algorithm,
- values defining the objective function (348) $\{f_i^{\text{target}}, \sigma_i\}_i$,

Problem is very nonconvex with lots of "bad" shapes–need to regularize it, e.g. with "Fourier continuation"

Need a MHD equilibrium code like VMEC or DESC for this step!

for N in {1, 2, …, 5}

for M in {1, 2, …, 5}

- fixed numbers $N + 1 + M(2N + 1)$ of parameters $R_{m,n}, Z_{m,n}$ to describe the boundary,

1. Find an initial boundary shape $\partial\Omega^{\text{init}}$.
   (a) Compute the initial MHD equilibrium magnetic field in Ω from $\partial\Omega^{\text{init}}$, $p(\psi)$, and $I_T(\psi)$.
   (b) Evaluate the set of objectives $\left\{f_i^{\text{equilibrium}}\right\}_i$
2. Until χ^2 satisfies a stopping criterion, repeat the following steps.
   (a) Adjust $\partial\Omega$ according to the optimization algorithm.
   (b) Compute the MHD equilibrium magnetic field in Ω from $\partial\Omega$, $p(\psi)$ and $I_T(\psi)$.
   (c) Evaluate the set of objectives $\left\{f_i^{\text{equilibrium}}\right\}_i$

Example of quasi-axisymmetry optimization

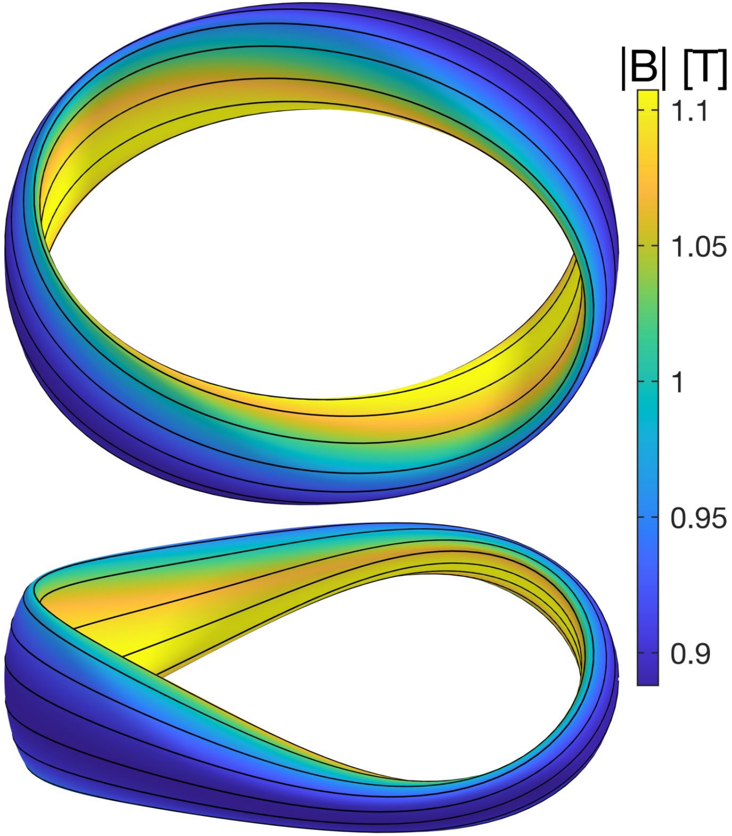

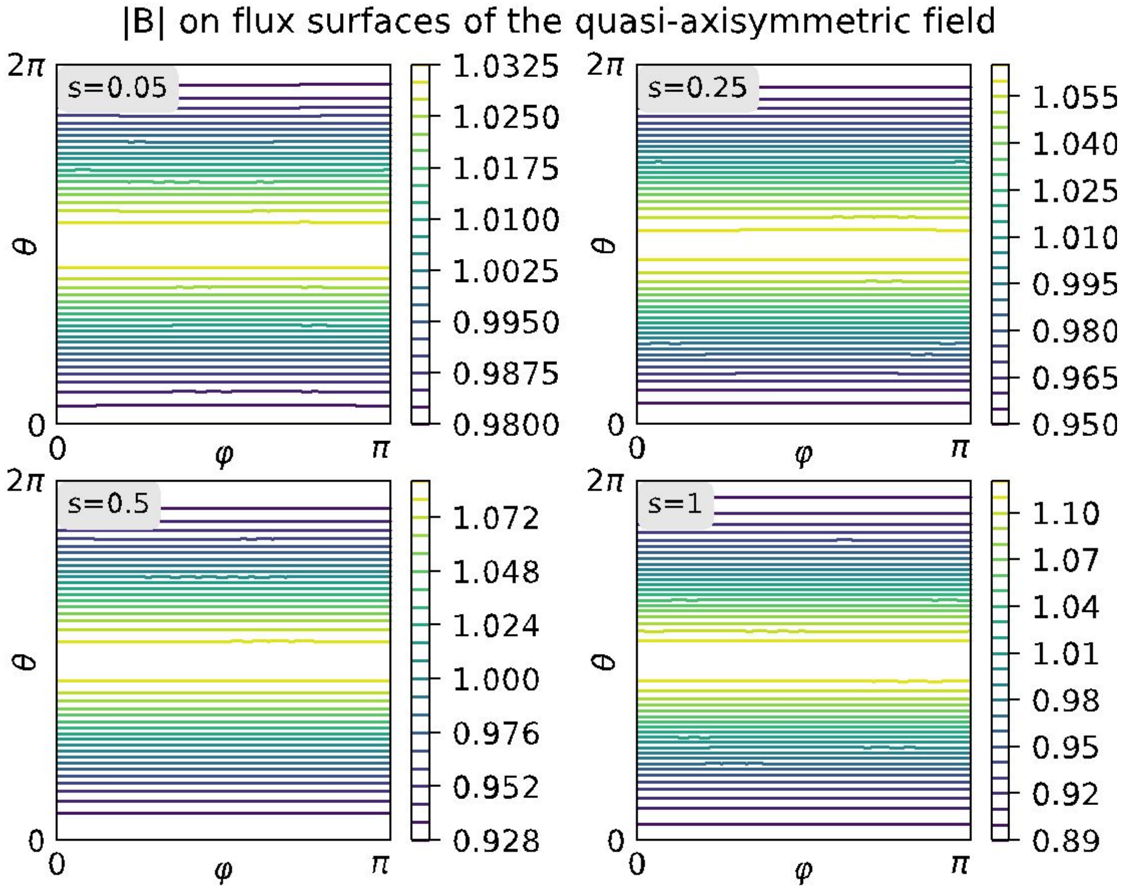

ML & Paul, Phys Rev Lett (2022)

Example of quasi-helical symmetry optimization

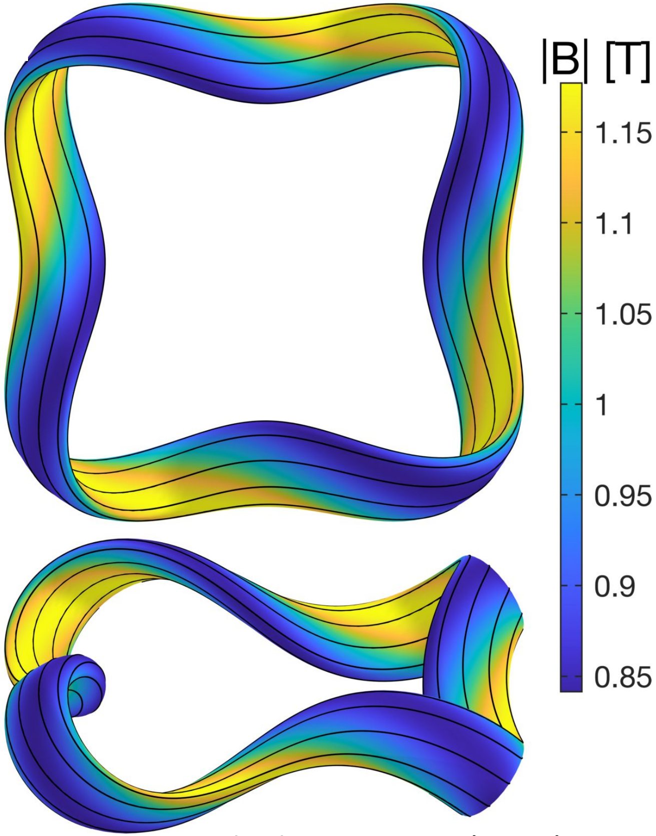

ML & Paul, Phys Rev Lett (2022)

|B| on flux surfaces of the quasi-helically symmetric field

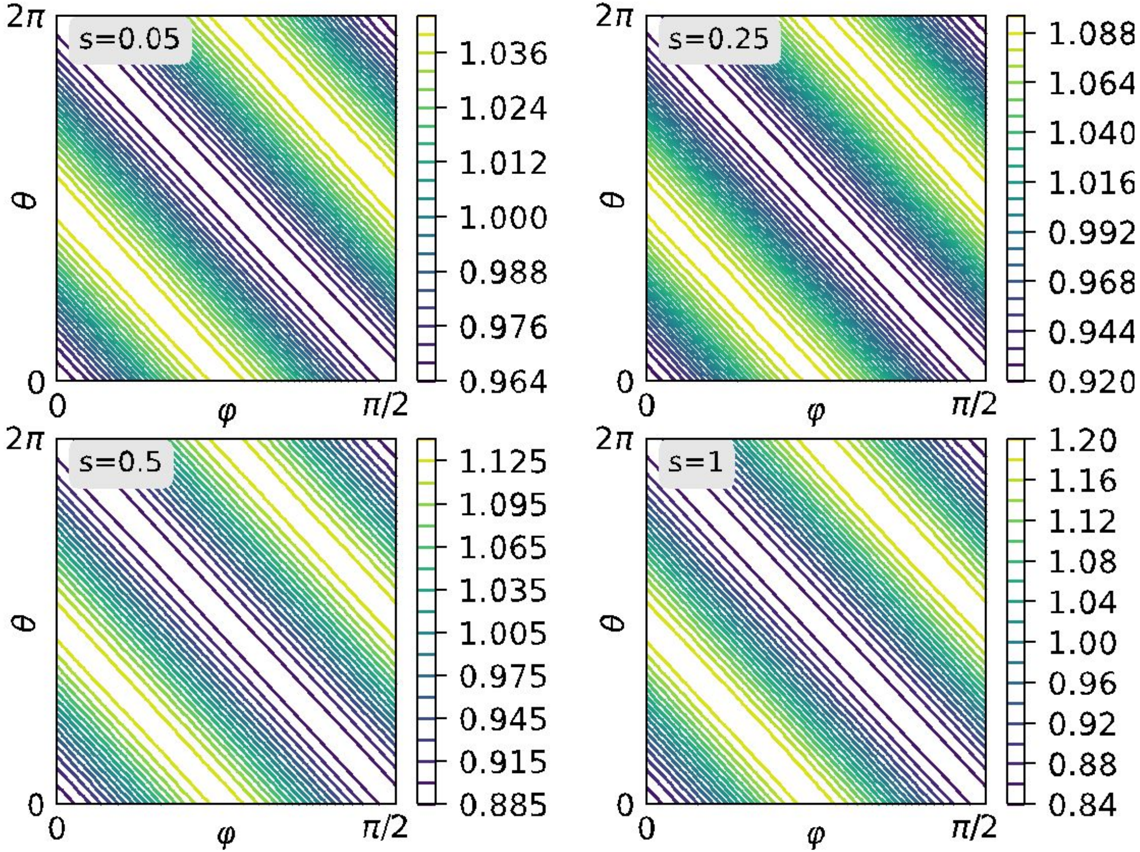

There has been similar recent progress in finding quasi-isodynamic configurations

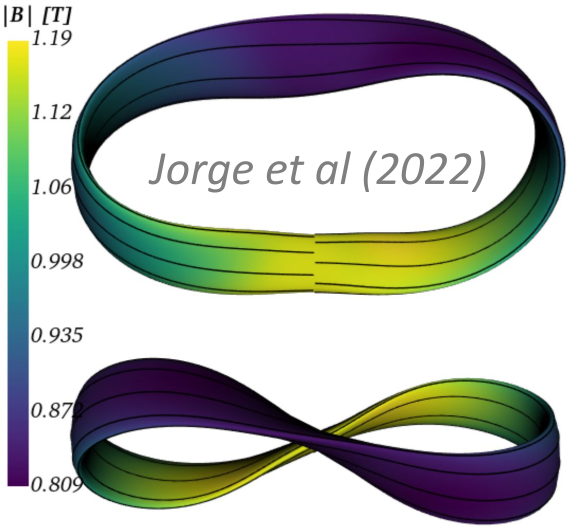

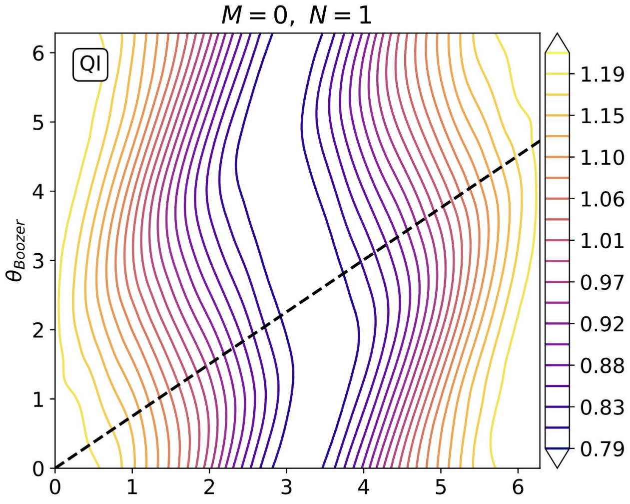

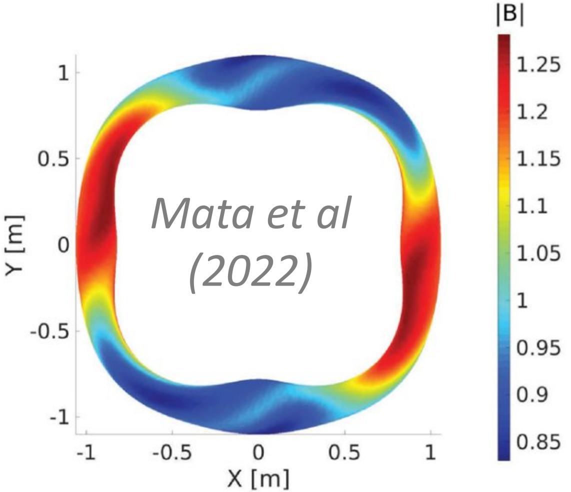

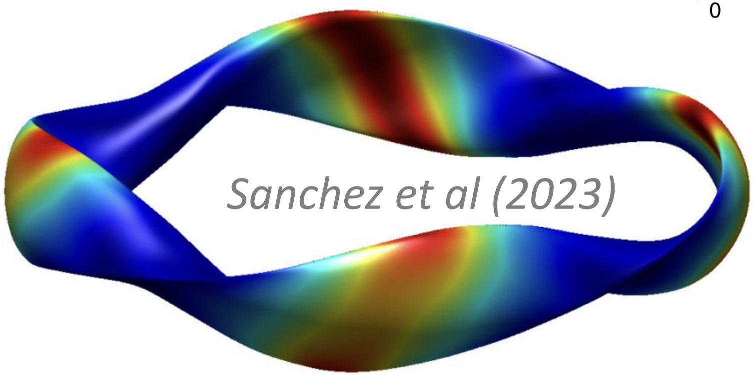

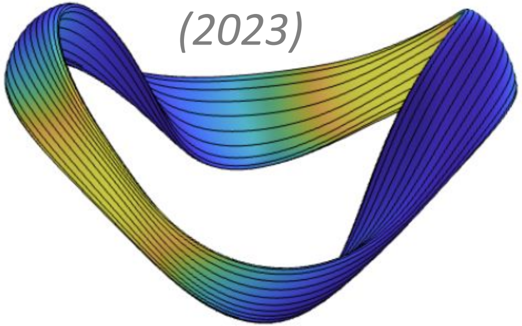

Outline

- Optimization in general
- Optimization for good flux surfaces
- Quasi-axisymmetry, quasi-helical symmetry, and quasi-isodynamic
- Coil optimization: current potential and filament methods

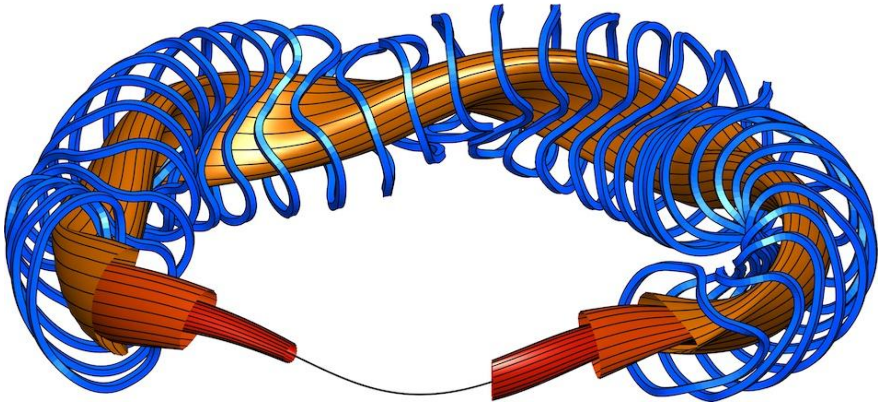

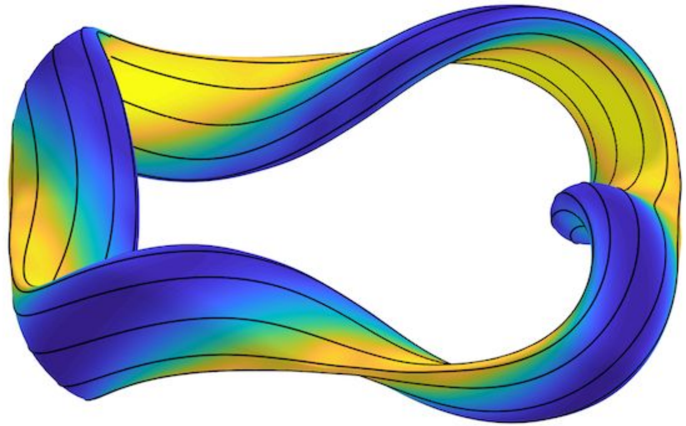

Stellarator designs can use a myriad of different current sources…

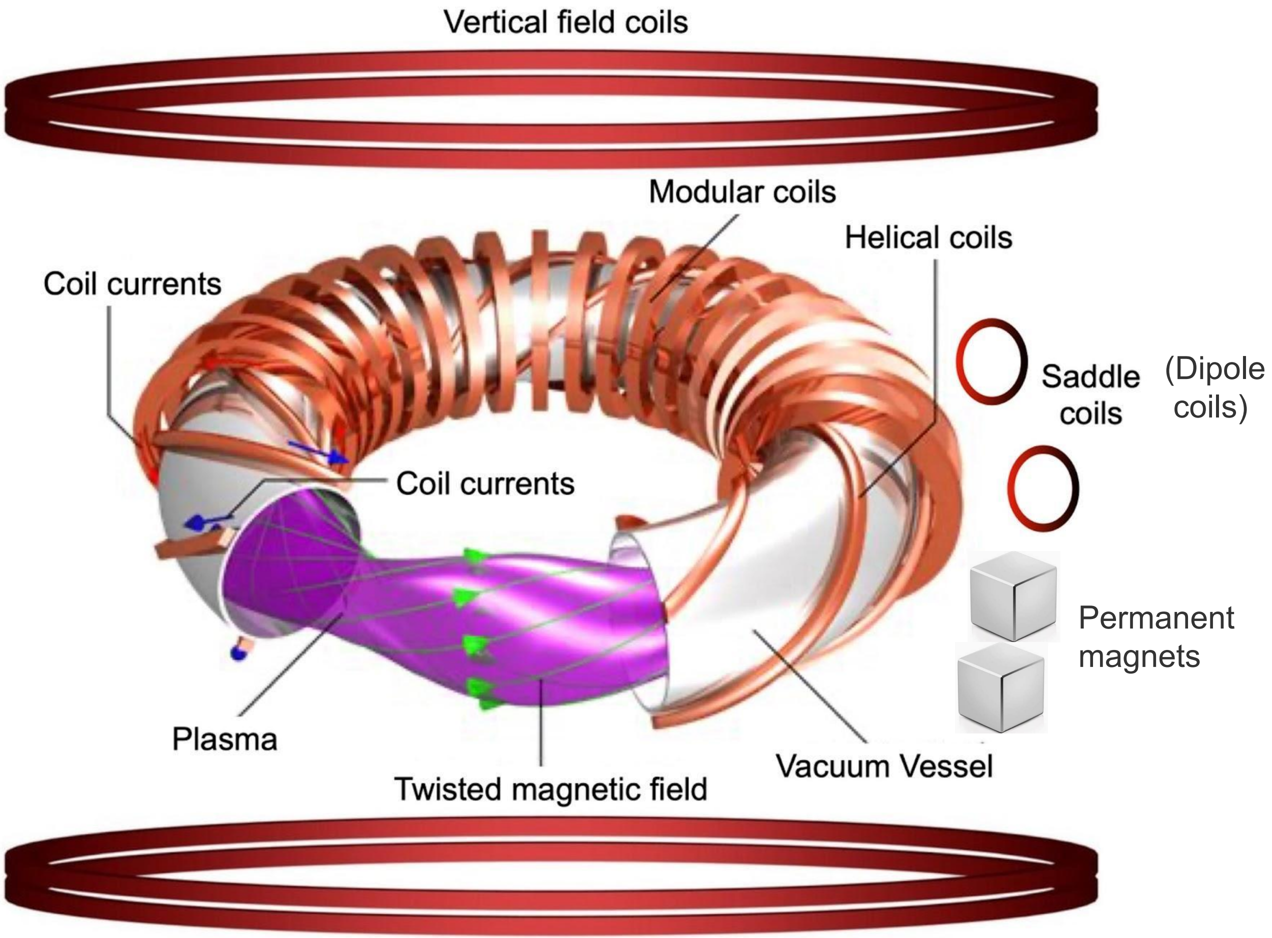

Stellarator coil design is “inverse magnetostatics”

- Inverse magnetostatics: Given a desired magnetic field on a compact surface S or volume V, how to find current sources that produce that field?
- Examples: Design magnetic fields for laboratory experiments (e.g. plasma physics, biology, particle physics), industrial applications (e.g. MRI), etc.
- Degrees of freedom: number, location, shape, strength of current sources.

Example magnetostatic optimization problem

$$\min_{\boldsymbol{\alpha}} \int_S \left(\boldsymbol{B}(\boldsymbol{\alpha}) \cdot \hat{\boldsymbol{n}}\right)^2 d\boldsymbol{r}$$

Biot-Savart Law for current source of volume V’

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi} \int_{V'} \frac{\boldsymbol{J}(\boldsymbol{r}') \times (\boldsymbol{r} - \boldsymbol{r}')}{|\boldsymbol{r} - \boldsymbol{r}'|^3} d\boldsymbol{r}'$$

Finding coils that produce a given B is analogous to fitting data with a polynomial

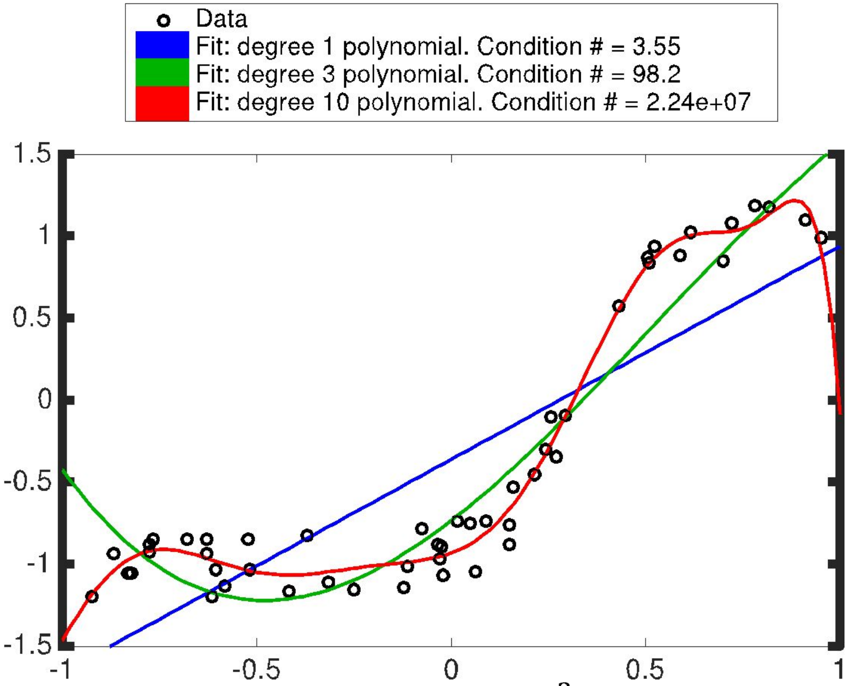

$$\text{Linear least-squares: } \min_{\vec{\mathbf{x}}} \left| \overleftrightarrow{\mathbf{A}}\vec{\mathbf{x}} - \vec{\mathbf{b}} \right|^2 \quad \Rightarrow \quad \vec{\mathbf{x}} = \left(\overleftrightarrow{\mathbf{A}}^T \overleftrightarrow{\mathbf{A}} \right)^{-1} \left(\overleftrightarrow{\mathbf{A}}^T \vec{\mathbf{b}} \right)$$

As polynomial degree increases, fit is closer to data but less ‘regular’.
Polynomial degree is the ‘regularization parameter’.

Calculating the currents that produce a given B is an "ill-posed inverse problem": solution is not unique.

Actually a good thing:

There is a lot of freedom in coil design

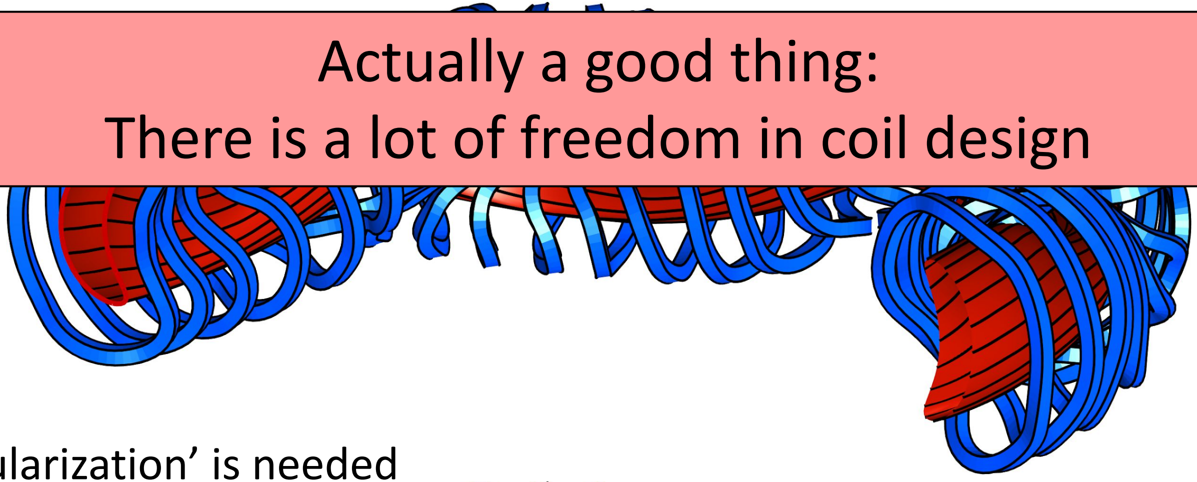

Some kind of 'regularization' is needed to exclude solutions like this:

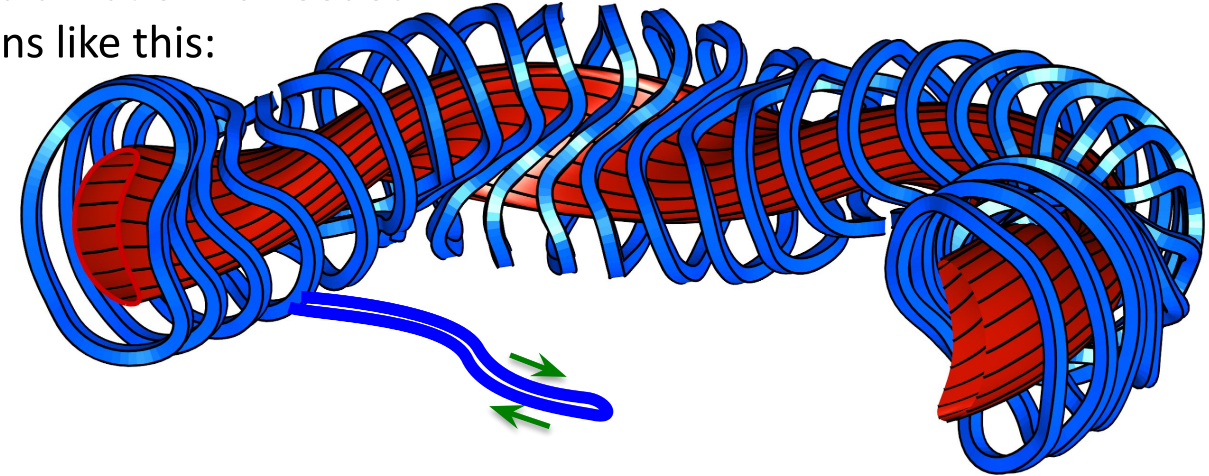

Current potential methods: NESCOIL & REGCOIL

Regcoil: Consider sheet current on a "coil winding surface"

ML, Nuclear Fusion (2017).

$$\boldsymbol{K} = \boldsymbol{n} \times \nabla \phi$$

Surface current

"current potential"

Normal to winding surface

$$\min_{\phi} \left(\int_{\text{Plasma surface}} d^2x \left[(\boldsymbol{B} - \boldsymbol{B}_{\text{target}}) \cdot \boldsymbol{n} \right]^2 + \lambda \int_{\text{Coil surface}} d^2x |\boldsymbol{K}|^2 \right)$$

B field error

Coil complexity

Regularization parameter

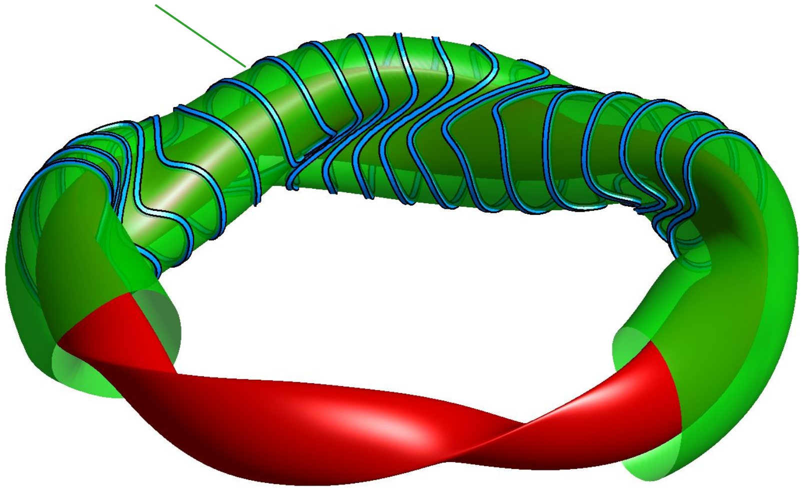

ϕ contours = coils

$K = |\boldsymbol{K}| \propto 1/\text{distance between coils}$

Pros:

- *Linear* least-squares: no local optima besides the global one.
- Only 2 parameters to vary: coil-to-plasma distance and λ.

Cons:

- Neglects ripple from discrete coils.
- Coils can't move in 3rd dimension.

In stage-2 coil optimization, there is a trade-off between field accuracy and coil simplicity

High regularization λ:
Simpler coils
but large field error

Low regularization λ:
Complicated coils
but small field error

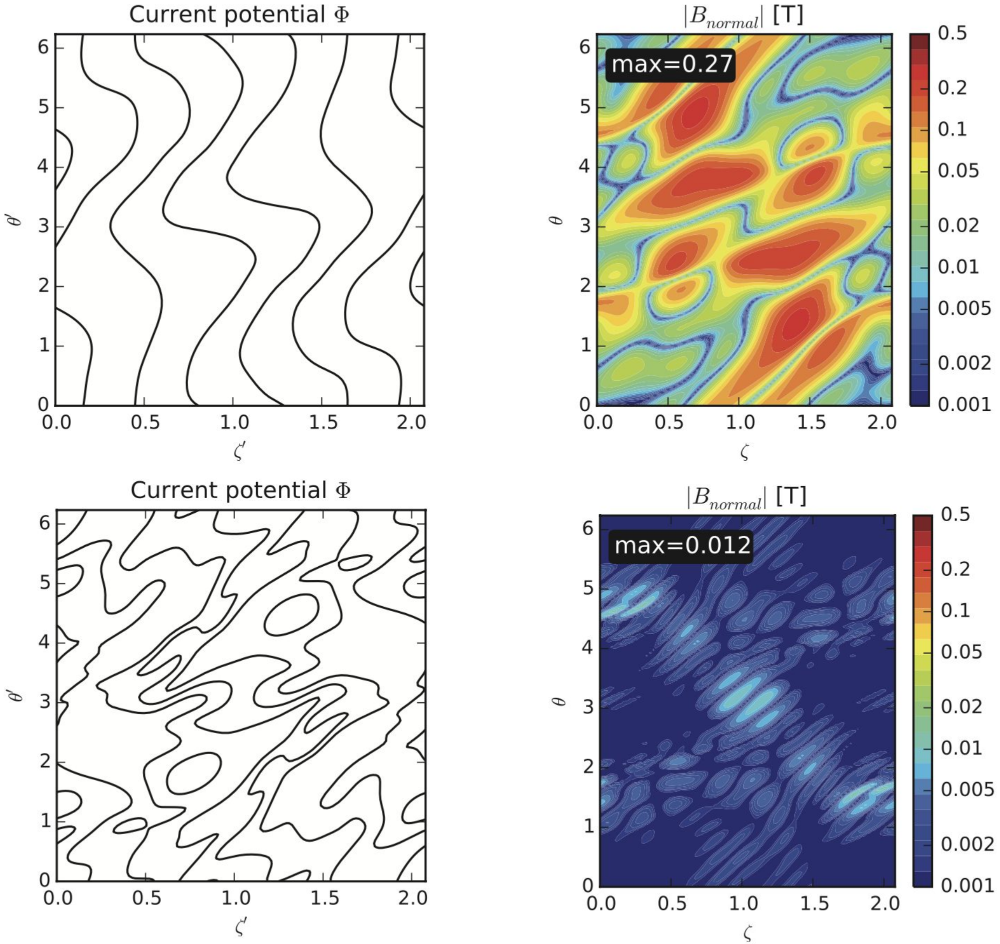

Filament coil optimization

Zhu, Hudson, et al, Nuclear Fusion (2018).

Assume plasma shape has already been optimized, so target **B** field is known.

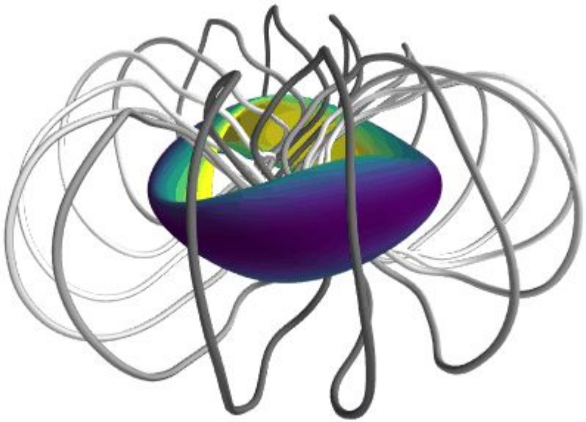

Coils represented as space curves.
Design variables: Fourier modes of Cartesian components.

$$x(t) = x_{c,0} + \sum_{n=1}^{N_F} [x_{c,n} \cos(nt) + x_{s,n} \sin(nt)]$$

Objective:

$$f = \underbrace{\int_{\substack{plasma \\ surf}} [(\boldsymbol{B} - \boldsymbol{B}_{target}) \cdot \boldsymbol{n}]^2}_{\text{Match target B}} + \underbrace{\lambda (length - target)^2 + \ldots}_{\text{Regularization}}$$

- Does account for *B* ripple from discreteness of coils.
- Non-convex, so there are multiple local minima. Need good initial guess.
- Often in practice we use “Fourier continuation” again.

Summary: there is lots of freedom in the shape of a stellarator plasma and coils, which can be used to achieve many objectives

- Large volume of good magnetic surfaces (not islands & chaos)
- Enough rotational transform
- Plasma pressure & current doesn't modify **B** too much, i.e. maximum plasma pressure is not too low.
- Buildable coil shapes: low curvature, large clearances
- Magnetohydrodynamic (MHD) stability
- Good confinement of particle trajectories
- Low neoclassical transport
- Low turbulent transport

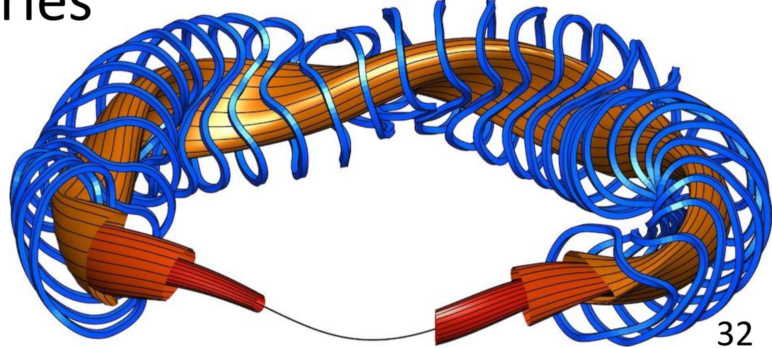

Open questions for stellarator optimization

- How best to combine coil and plasma design?
- How to find designs that tolerate errors in coil shape/position?
- How to avoid getting stuck in little local minima? How to find global optima?
- How to optimize for expensive & noisy objectives (turbulence & fast-particle confinement)?
- How to balance multiple competing objectives?
- How to optimize coil topology?
- How to find configurations that are flexible?
  - Good confinement for different plasma pressures.
  - Ability to tune physics properties by changing coil currents.

More resources

Introductory papers:

Imbert-Gerard, Paul, & Wright, https://arxiv.org/abs/1908.05360
Helander, http://dx.doi.org/10.1088/0034-4885/77/8/087001

Summer schools:

https://hiddensymmetries.princeton.edu/summer-school/summer-school-2020/schedule
https://hiddensymmetries.princeton.edu/summer-school/summer-school-2019/schedule
https://gss.pppl.gov/2021/
https://suli.pppl.gov/2022/course/index.html
https://suli.pppl.gov/2021/course/index.html
https://suli.pppl.gov/2020/course/index.html
https://suli.pppl.gov/2019/course/index.html

Open-source software:

https://github.com/PrincetonUniversity/STELLOPT
https://desc-docs.readthedocs.io/
https://simsopt.readthedocs.io/
https://github.com/landreman/regcoil
https://gitlab.com/wistell/StellaratorOptimization.jl

Extra slides

Optimization for reduced turbulence has mostly used simplified proxies in the cost function

Turbulent heat flux can be simulated, but it is computationally expensive and noisy
⇒ not good for an objective function

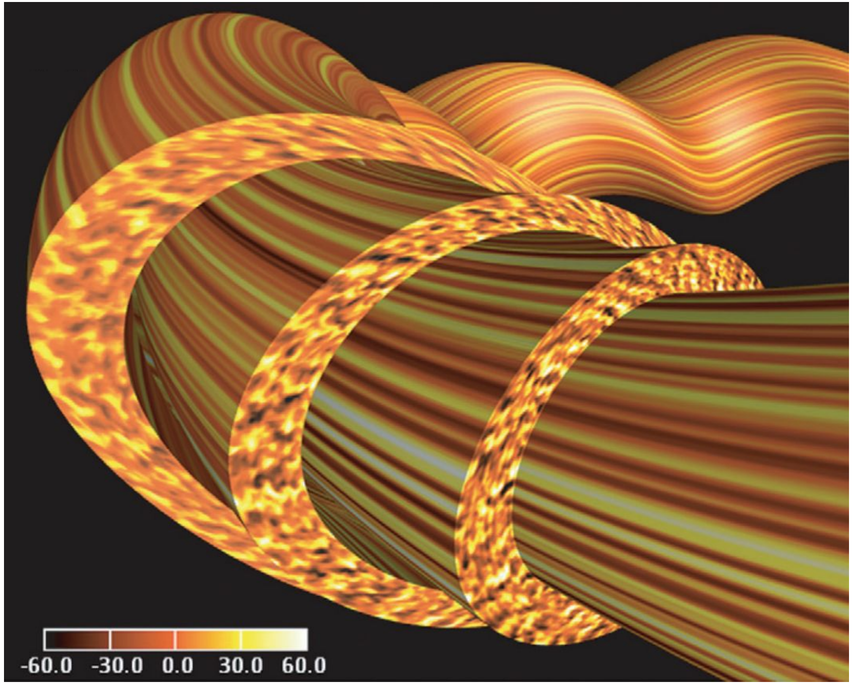

Nunami 2017

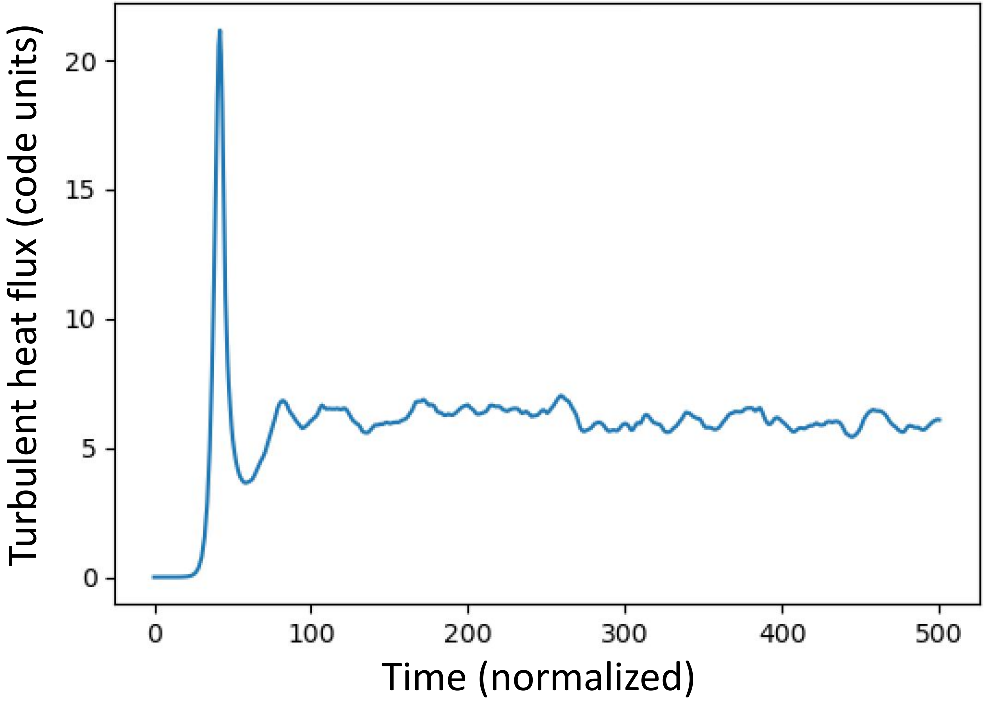

The first optimizations with nonlinear turbulence calculations in the objective are becoming possible

By Patrick Kim (was a Maryland undergraduate, now at Princeton!)

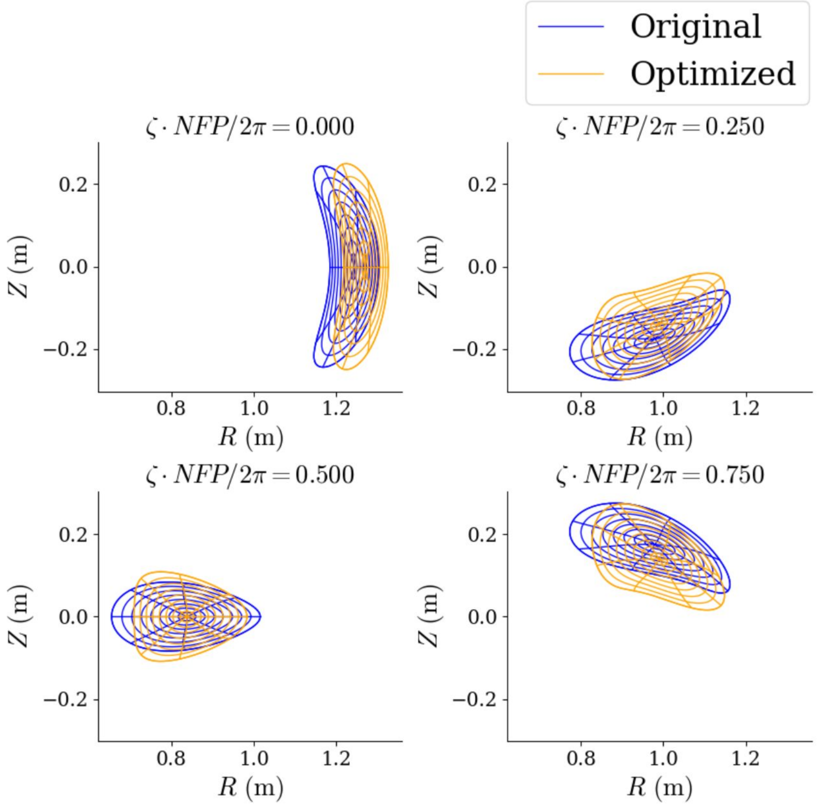

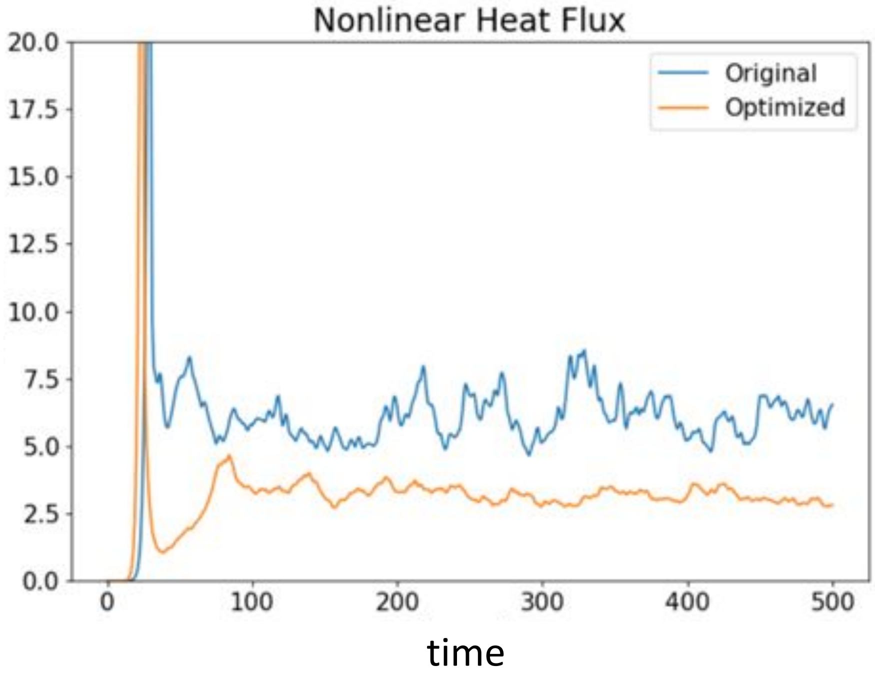

Stellarator geometry can be optimized for MHD stability

Types of MHD stability calculations, in increasing complexity:

- Magnetic well & Mercier's criterion (interchange)
- Ballooning modes (short wavelength ⊥ to **B**)
- Finite wavelength (everything)

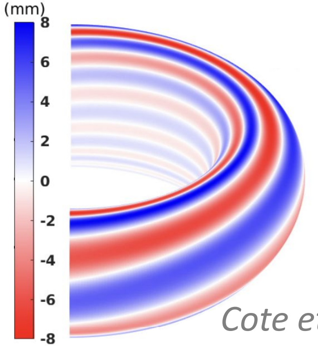

Cote et al, (2019)

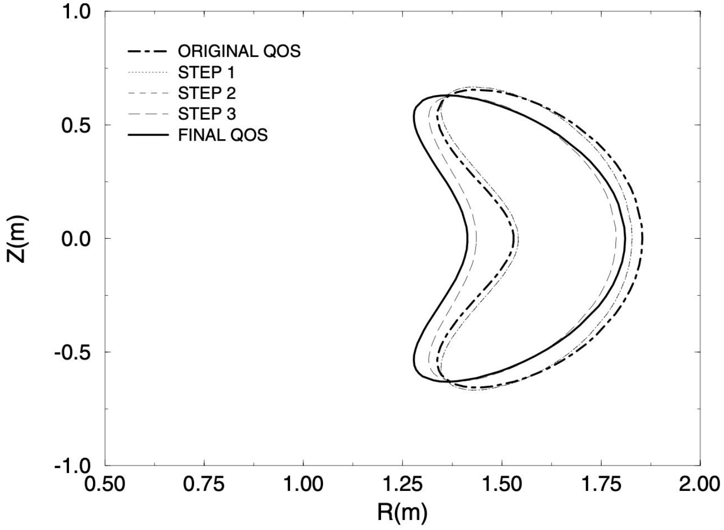

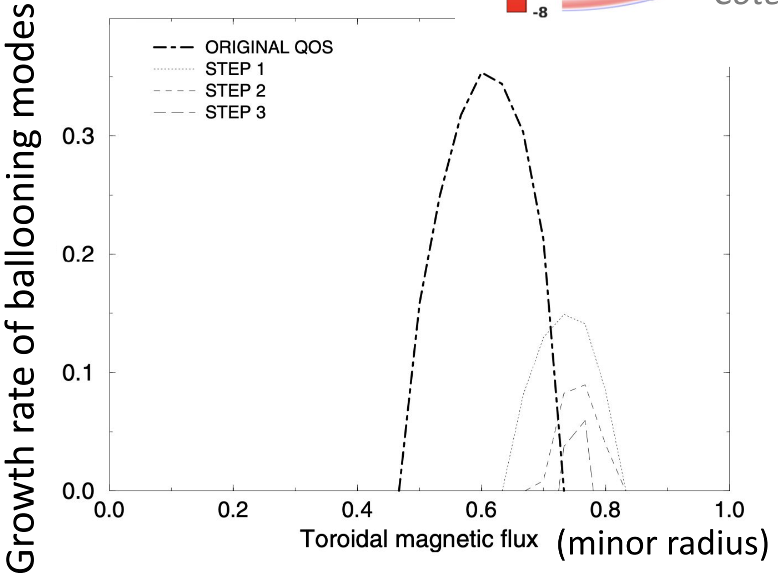

Sanchez et al, Plasma Phys. Control. Fusion (2000)

Pros and cons of using modular superconducting coils

- Magnet fields are all nonlinear in the optimization variables (the coils shapes).
- Machining complex superconducting coils can be very expensive.
- Support structures are complicated and tight tolerances are required.
- Easy to build in holes and diagnostic ports and optimization problem is unchanged.
- Coils operate with power supplies, require significant cooling, generate magnetic ripple
- Coil field strength limited only by large magnetic forces, can be exposed to high fields.
- Can put coils behind neutron shield (or blanket) and can be turned off.
- Can trivially generate a toroidal magnetic flux.

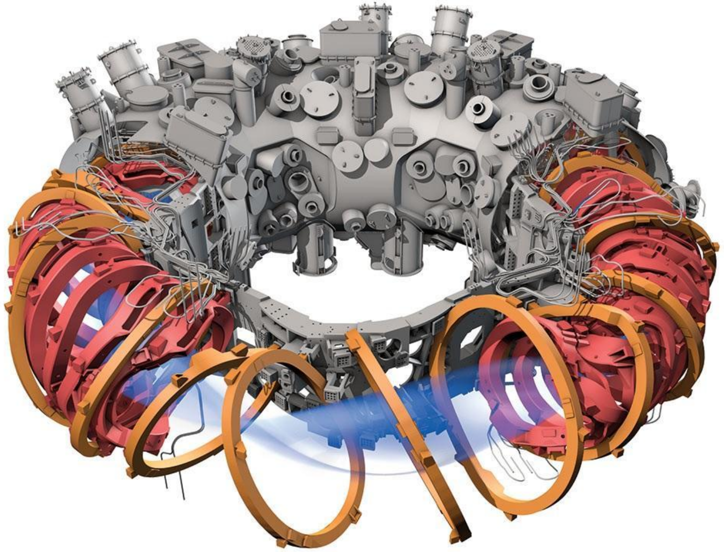

Pros and cons of using permanent magnets

- Magnet fields are all linear in the optimization variables (the dipole vectors).
- Permanent magnets are very cheap compared to machining complex superconducting coils.*
- Support structures can be 3D printed and assembled at remarkably low cost.
- Easy to build in holes and diagnostic ports and optimization problem is unchanged.
- Magnets operate without power supplies, require minimal cooling.
- Magnets limited to ~ 1 Tesla, demagnetize if exposed to high field strengths,
- Cannot withstand neutron bombardment and cannot be turned off.
- Still need some basic coils to generate a toroidal magnetic flux.

$$\boldsymbol{B}_M = \frac{\mu_0}{4\pi}\sum_{i=1}^{D}\left(\frac{3\boldsymbol{m}_i\cdot\boldsymbol{r}_i}{\|\boldsymbol{r}_i\|_2^5}\boldsymbol{r}_i - \frac{\boldsymbol{m}_i}{\|\boldsymbol{r}_i\|_2^3}\right)$$

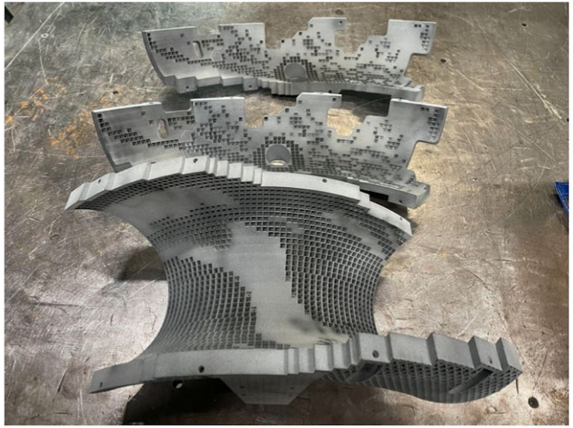

Full sized PM holders 3D printed with nylon. Accuracy measured by laser metrology.

NYU

*~10$ per magnet, so ~10^4 - 10^5 magnet designs are quite inexpensive.

Pros and cons of using a dipole array (DA)

- ○ Magnet fields are all linear in the optimization variables (the current strengths).
- ○ DA coils are very cheap compared to machining complex superconducting coils.
- ○ Support structures are relatively simple and low cost.
- ○ Easy to build in holes and diagnostic ports and optimization problem is unchanged.
- ○ Coils operate with power supplies, require cooling.
- ○ DA field strength limited only by large magnetic forces, can be exposed to high field strengths.
- ○ Can put DA behind neutron shield (or blanket) and can be turned off.
- ○ Still need some basic TF coils to generate a toroidal magnetic flux.

https://thea.energy/

$$\beta^2 \equiv a^2 + r^2 + 2ar\sin\theta,$$

$$k^2 \equiv 1 - \alpha^2/\beta^2, \text{ and } C \equiv \mu_0 I/\pi = \frac{\mu_0 m}{\pi^2 a^2}$$

$$\alpha^2 \equiv a^2 + r^2 - 2ar\sin\theta,$$

$$B_\varphi = 0\,.$$

$$B_r = \frac{Ca^2\cos\theta}{\alpha^2\beta}E(k^2)$$

$$B_\theta = \frac{C}{2\alpha^2\beta\sin\theta}\left[(r^2 + a^2\cos 2\theta)E(k^2) - \alpha^2 K(k^2)\right]$$

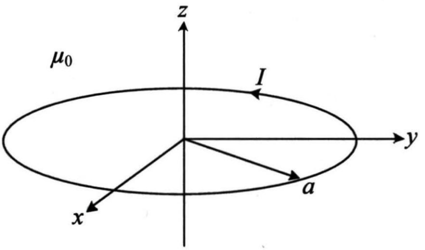

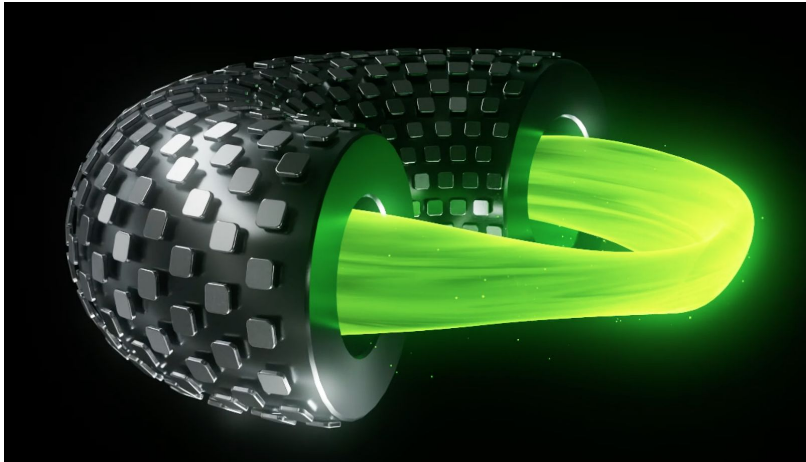

Pros and cons of using superconducting tiles (STs)

- Magnet fields are **not** linear in the optimization variables.
- STs are very cheap compared to machining complex superconducting coils.
- Support structures can be 3D printed and assembled at low cost.
- Easy to build in holes and diagnostic ports and optimization problem is unchanged.
- STs operate without power supplies, require some cooling.
- DA field strength limited only by large magnetic forces, can be exposed to high field strengths.
- Can put DA behind neutron shield (or blanket) and can be turned off (sort of!).
- Still need some basic coils to generate a toroidal magnetic flux + induce ST currents.

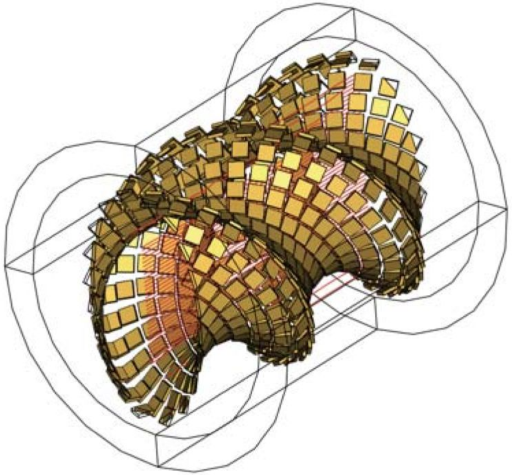

Flux of **B** through ST must vanish, so currents are induced in the ST to cancel any external field.

Bromberg, L., et al. *Fusion Science and Technology* 60.2 (2011): 643-647.

Optimization for reduced turbulence has mostly used simplified proxies in the cost function

$$f = -0.959 n g^{xx} \frac{dT}{dx} \sum_k \frac{\gamma_k}{k_x^2} \qquad \gamma_k = -\frac{c k_y T_i}{eB} \sqrt{\kappa_1 (\kappa_p - \kappa_{cr})} H(\kappa_p - \kappa_{cr}) H(-\kappa_1) \qquad \kappa_{cr} = 0.053$$

$$\frac{1}{k_x^2} = \rho_i^2 + \frac{\rho_i L_p}{1 + \left\langle (1.12 s_l)^2 \right\rangle_\theta} \qquad s_l = \frac{\partial (g^{xy}/g^{xx})}{\partial \theta} \qquad \kappa_1 = (\mathbf{b} \cdot \nabla \mathbf{b}) \cdot \frac{\partial \mathbf{r}}{\partial x} \qquad \kappa_p = \frac{1}{L_p} = -\frac{1}{p}\frac{dp}{dx}$$

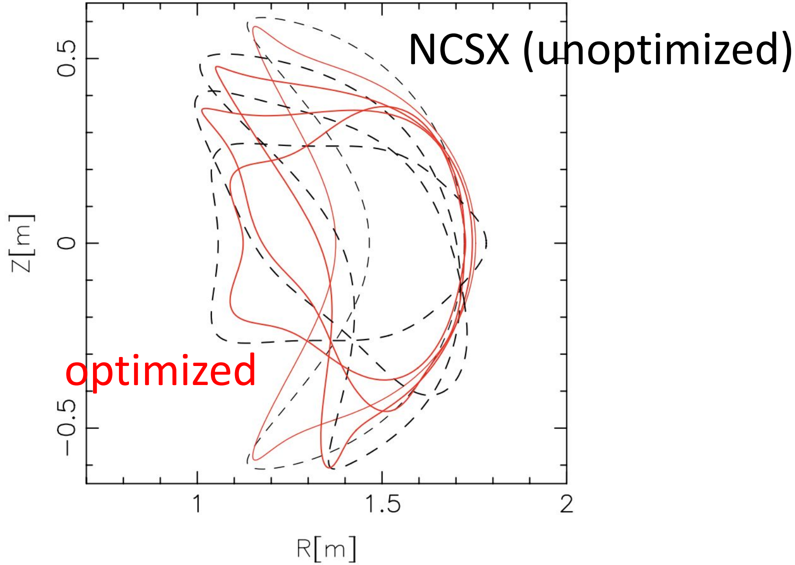

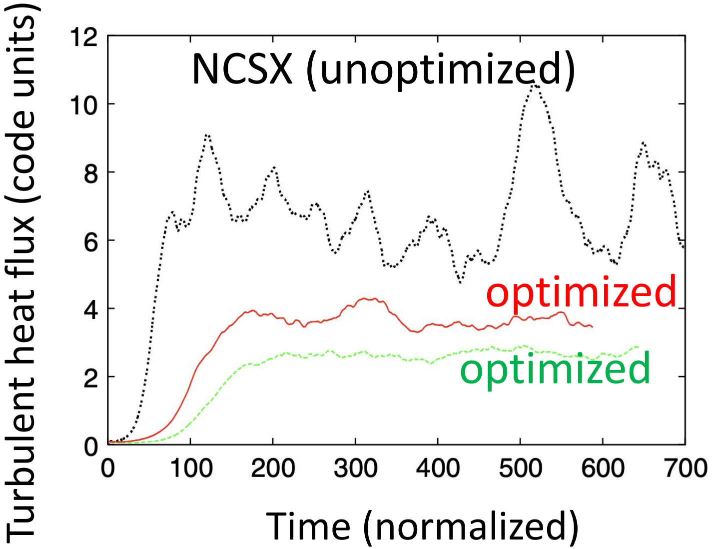

Mynick et al, Physical Review Letters (2010)

Pros and cons of the 3 classes

QA:

+ Lowest aspect ratio
+ Fewest coils, largest clearances
+ Large bootstrap current increases iota
- Wider orbits mean worse confinement
- Large current may contribute to MHD instability

QH:

+ Extremely good confinement
+ Can build on experience with HSX
- Seems to require high aspect ratio and many coils
? Intermediate bootstrap current between QA and QI

QI:

+ Low bootstrap current means high robustness to different pressure profiles
+ Can use island divertor
+ Can build on experience from W7-X
- Seems to require high aspect ratio and many coils
- Optimization is generally trickier

Perhaps the first type of stellarator optimization was to achieve good flux surfaces

Reproduction of Cary & Hanson (1986) by Rogerio Jorge

Unoptimized

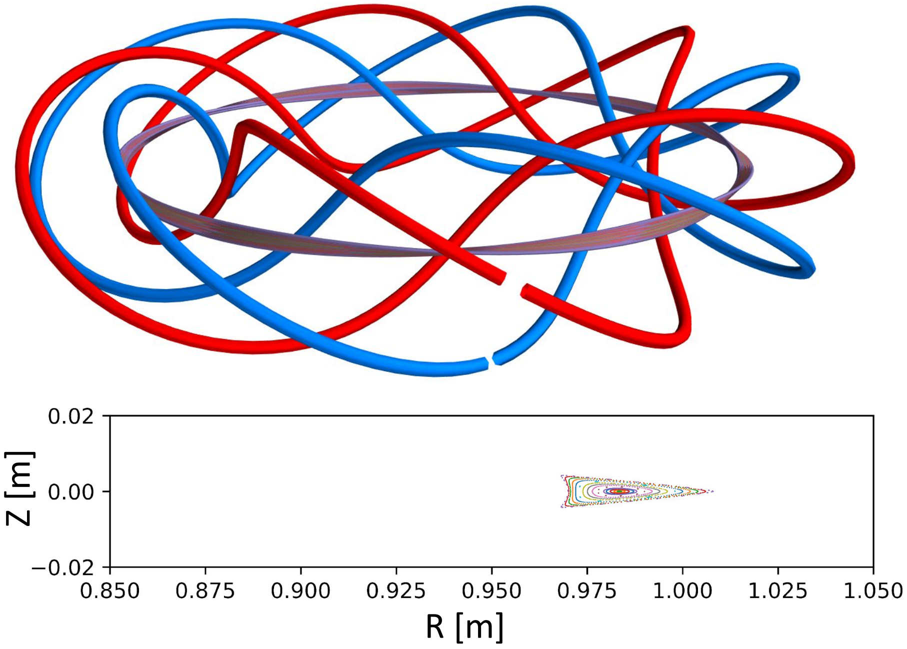

Optimized

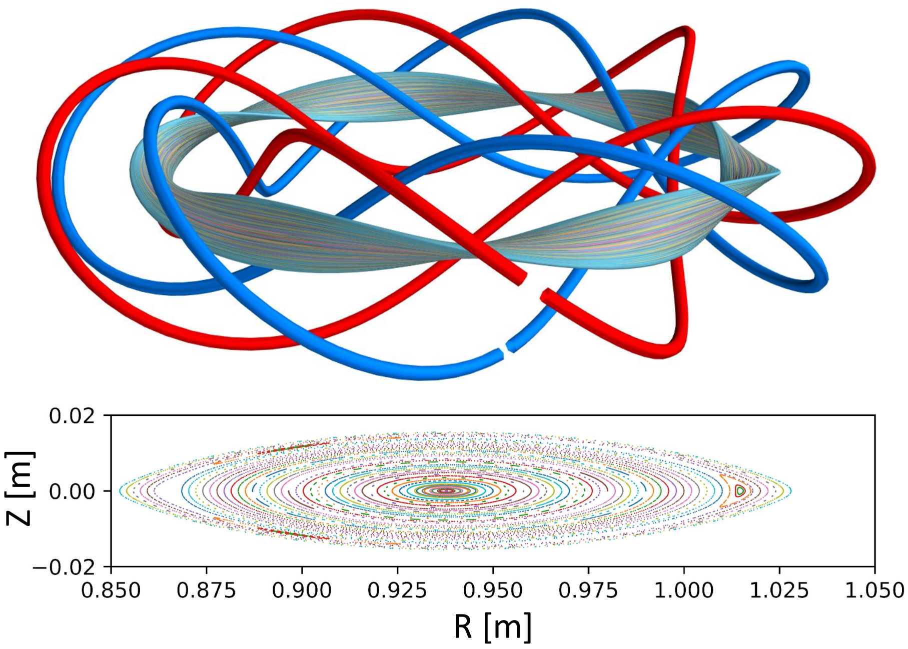

Parameter space *x* = Fourier modes of coil shapes.

Cost function *f* = square of “Greene’s residue”

Optimization for good flux surfaces continues to be a principle behind recent stellarators

CTH:
(Compact Toroidal Hybrid, at Auburn)

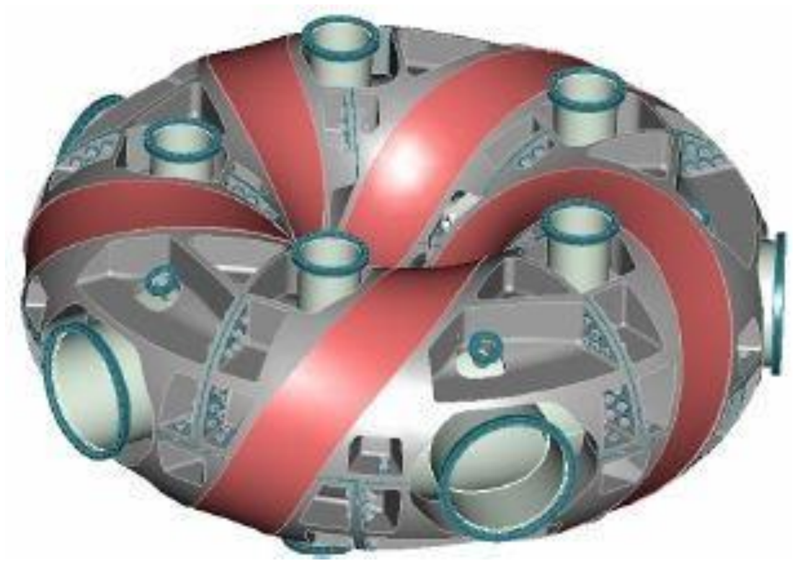

CNT (Columbia Non-neutral Torus):
Optimize *expected* volume over possible coil position errors

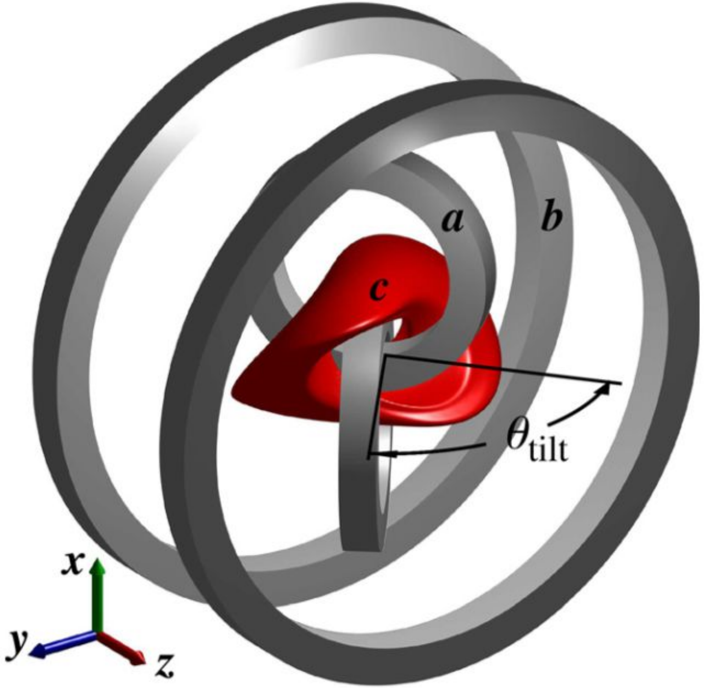

Pedersen (2004), Hammond (2016)

Reminder: the ∇B and curvature drifts make confinement challenging

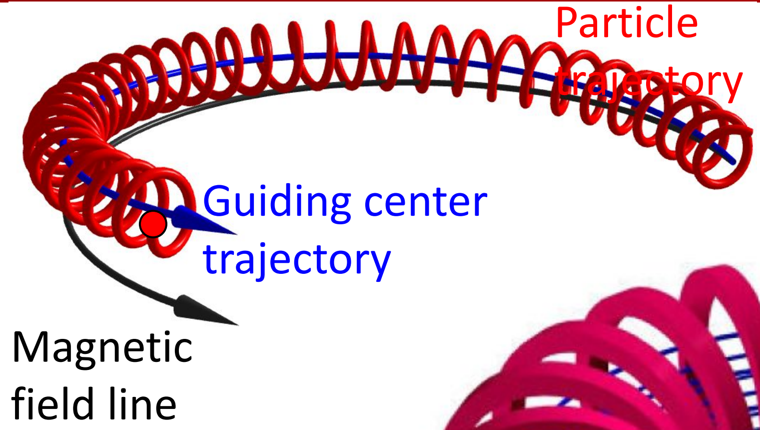

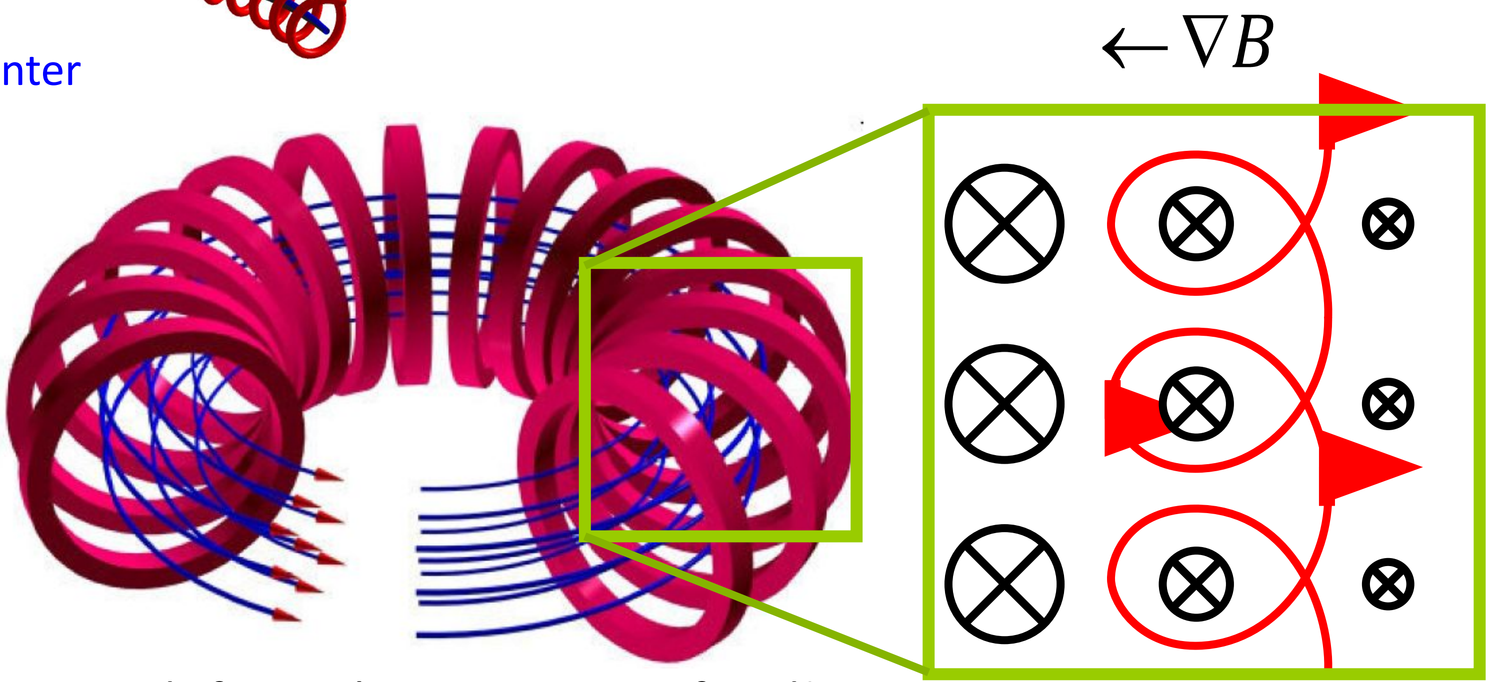

Ions drift up: they are not confined!

Particles drift in the $q\mathbf{B}\times\nabla B$ direction

Mirror force: particles are pushed away from regions of high |B|

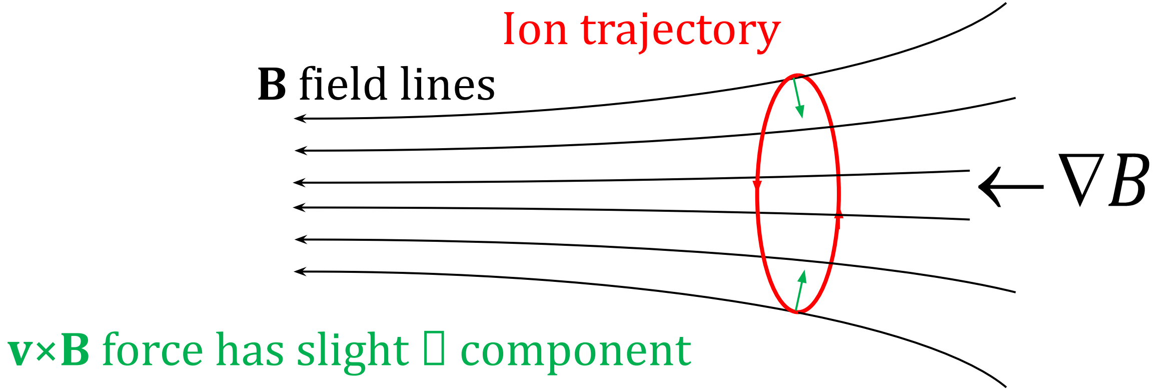

$$\frac{dv_{\parallel}}{dt} = -\frac{v_{\perp}^{2}}{2B}\mathbf{b}\cdot\nabla B$$

A few particles with very small $v_{||}$ = **v** · **B** “bounce” and are “trapped” in low-|B| regions.

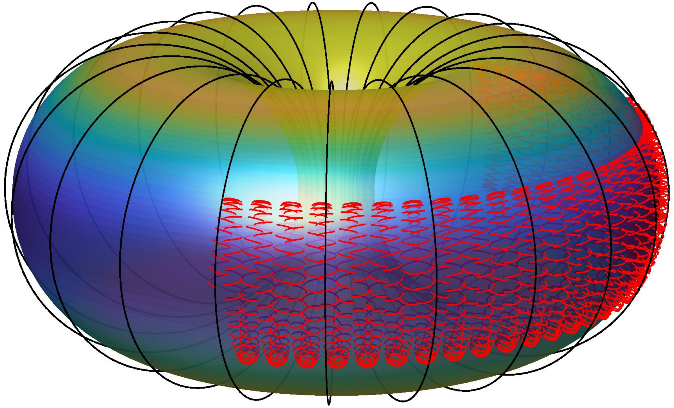

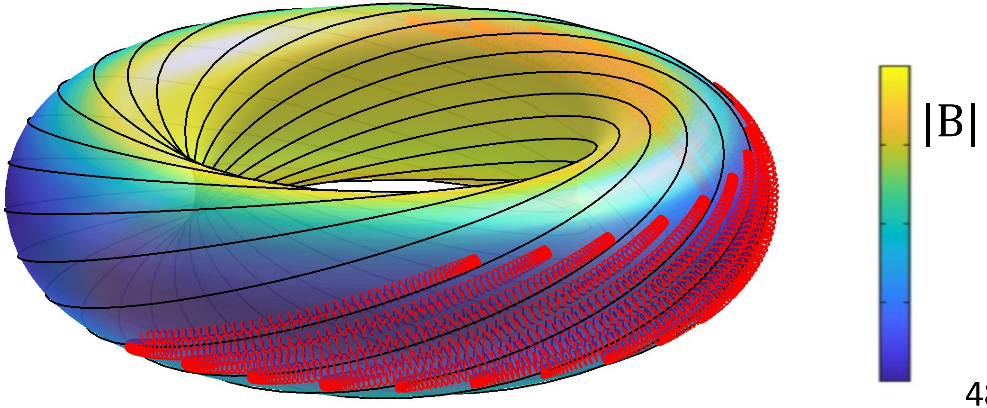

Flux surfaces are not enough: *Trapped* particles are not confined without a further condition like "quasisymmetry" or "omnigenity"

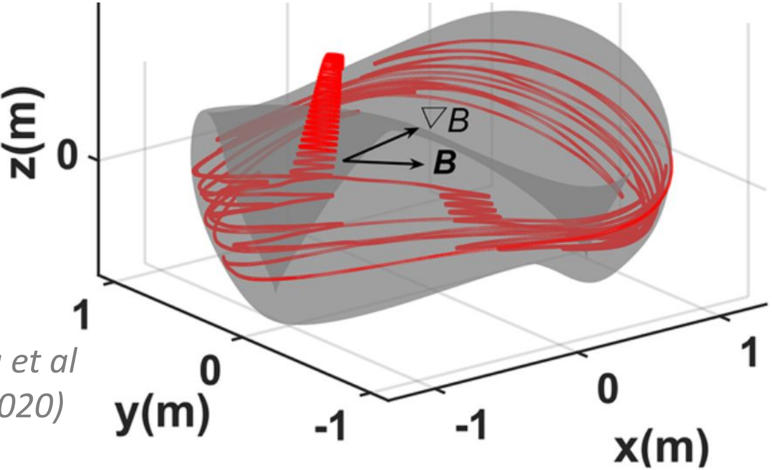

Su et al (2020)

In general: trapped particles do not sample the whole surface, so cross-field drift does not average to 0.

$\Longrightarrow$ Large neoclassical transport.

One solution is quasisymmetry: make $B(r, \theta, \varphi) = B(r, M\theta - N\varphi)$ for special angles θ, φ.

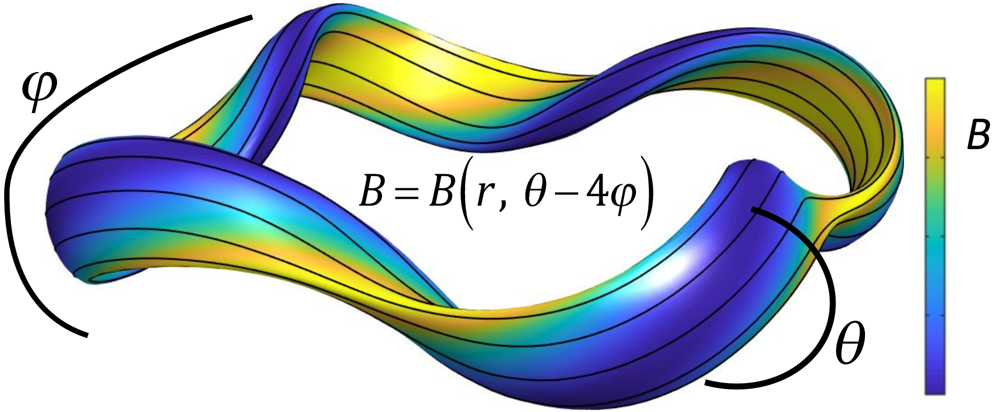

Symmetry direction

$\Longrightarrow$ Conserved quantity.

$\Longrightarrow$ Drift averages to 0.

Lemma: deeply trapped particles move so $|B|$ is constant

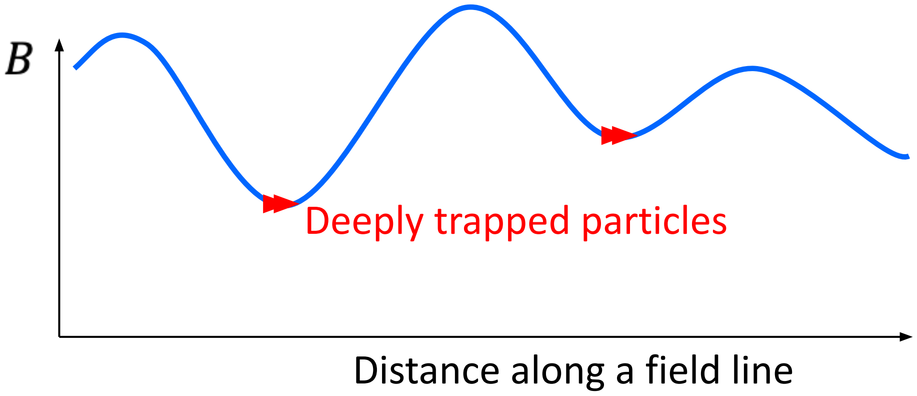

$$\frac{dB}{dt} = \boldsymbol{v} \cdot \nabla B$$

$$v_{||} \approx 0 \ \text{ so } \boldsymbol{v} \approx \text{the } \nabla B \text{ drift} \quad \Longrightarrow \quad \boldsymbol{v} \,||\, \boldsymbol{B} \times \nabla B$$

$$\frac{dB}{dt} \propto (\boldsymbol{B} \times \nabla B) \cdot \nabla B = 0.$$

For low neoclassical transport, recent stellarators have come in 3 flavors

- Trapped particles should drift toroidally, helically, or poloidally on a surface.
- *B* contours on a surface have the same topology as these drifts.

Toroidal:

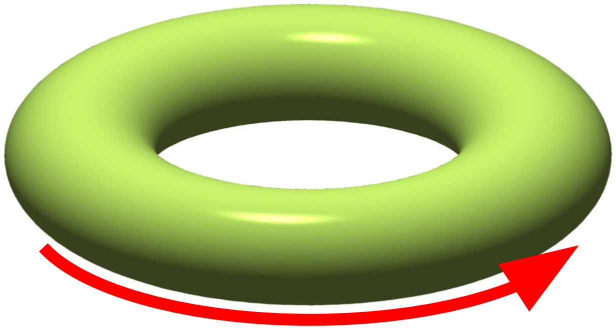

E.g., particles with $v_{||}=0$ move along a constant-*B* contour:

$$\left(\nabla B \text{ drift}\right)\cdot\nabla B \propto \mathbf{B}\times\nabla B\cdot\nabla B = 0$$

Helical:

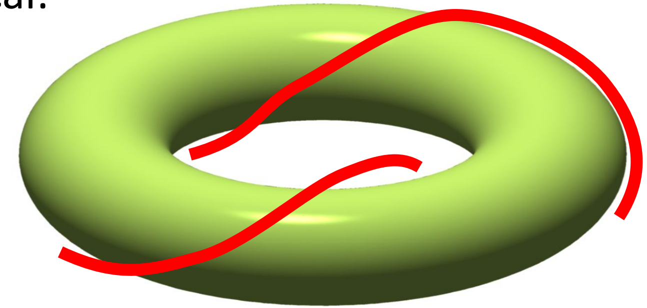

Poloidal: