MARITIME FUSION

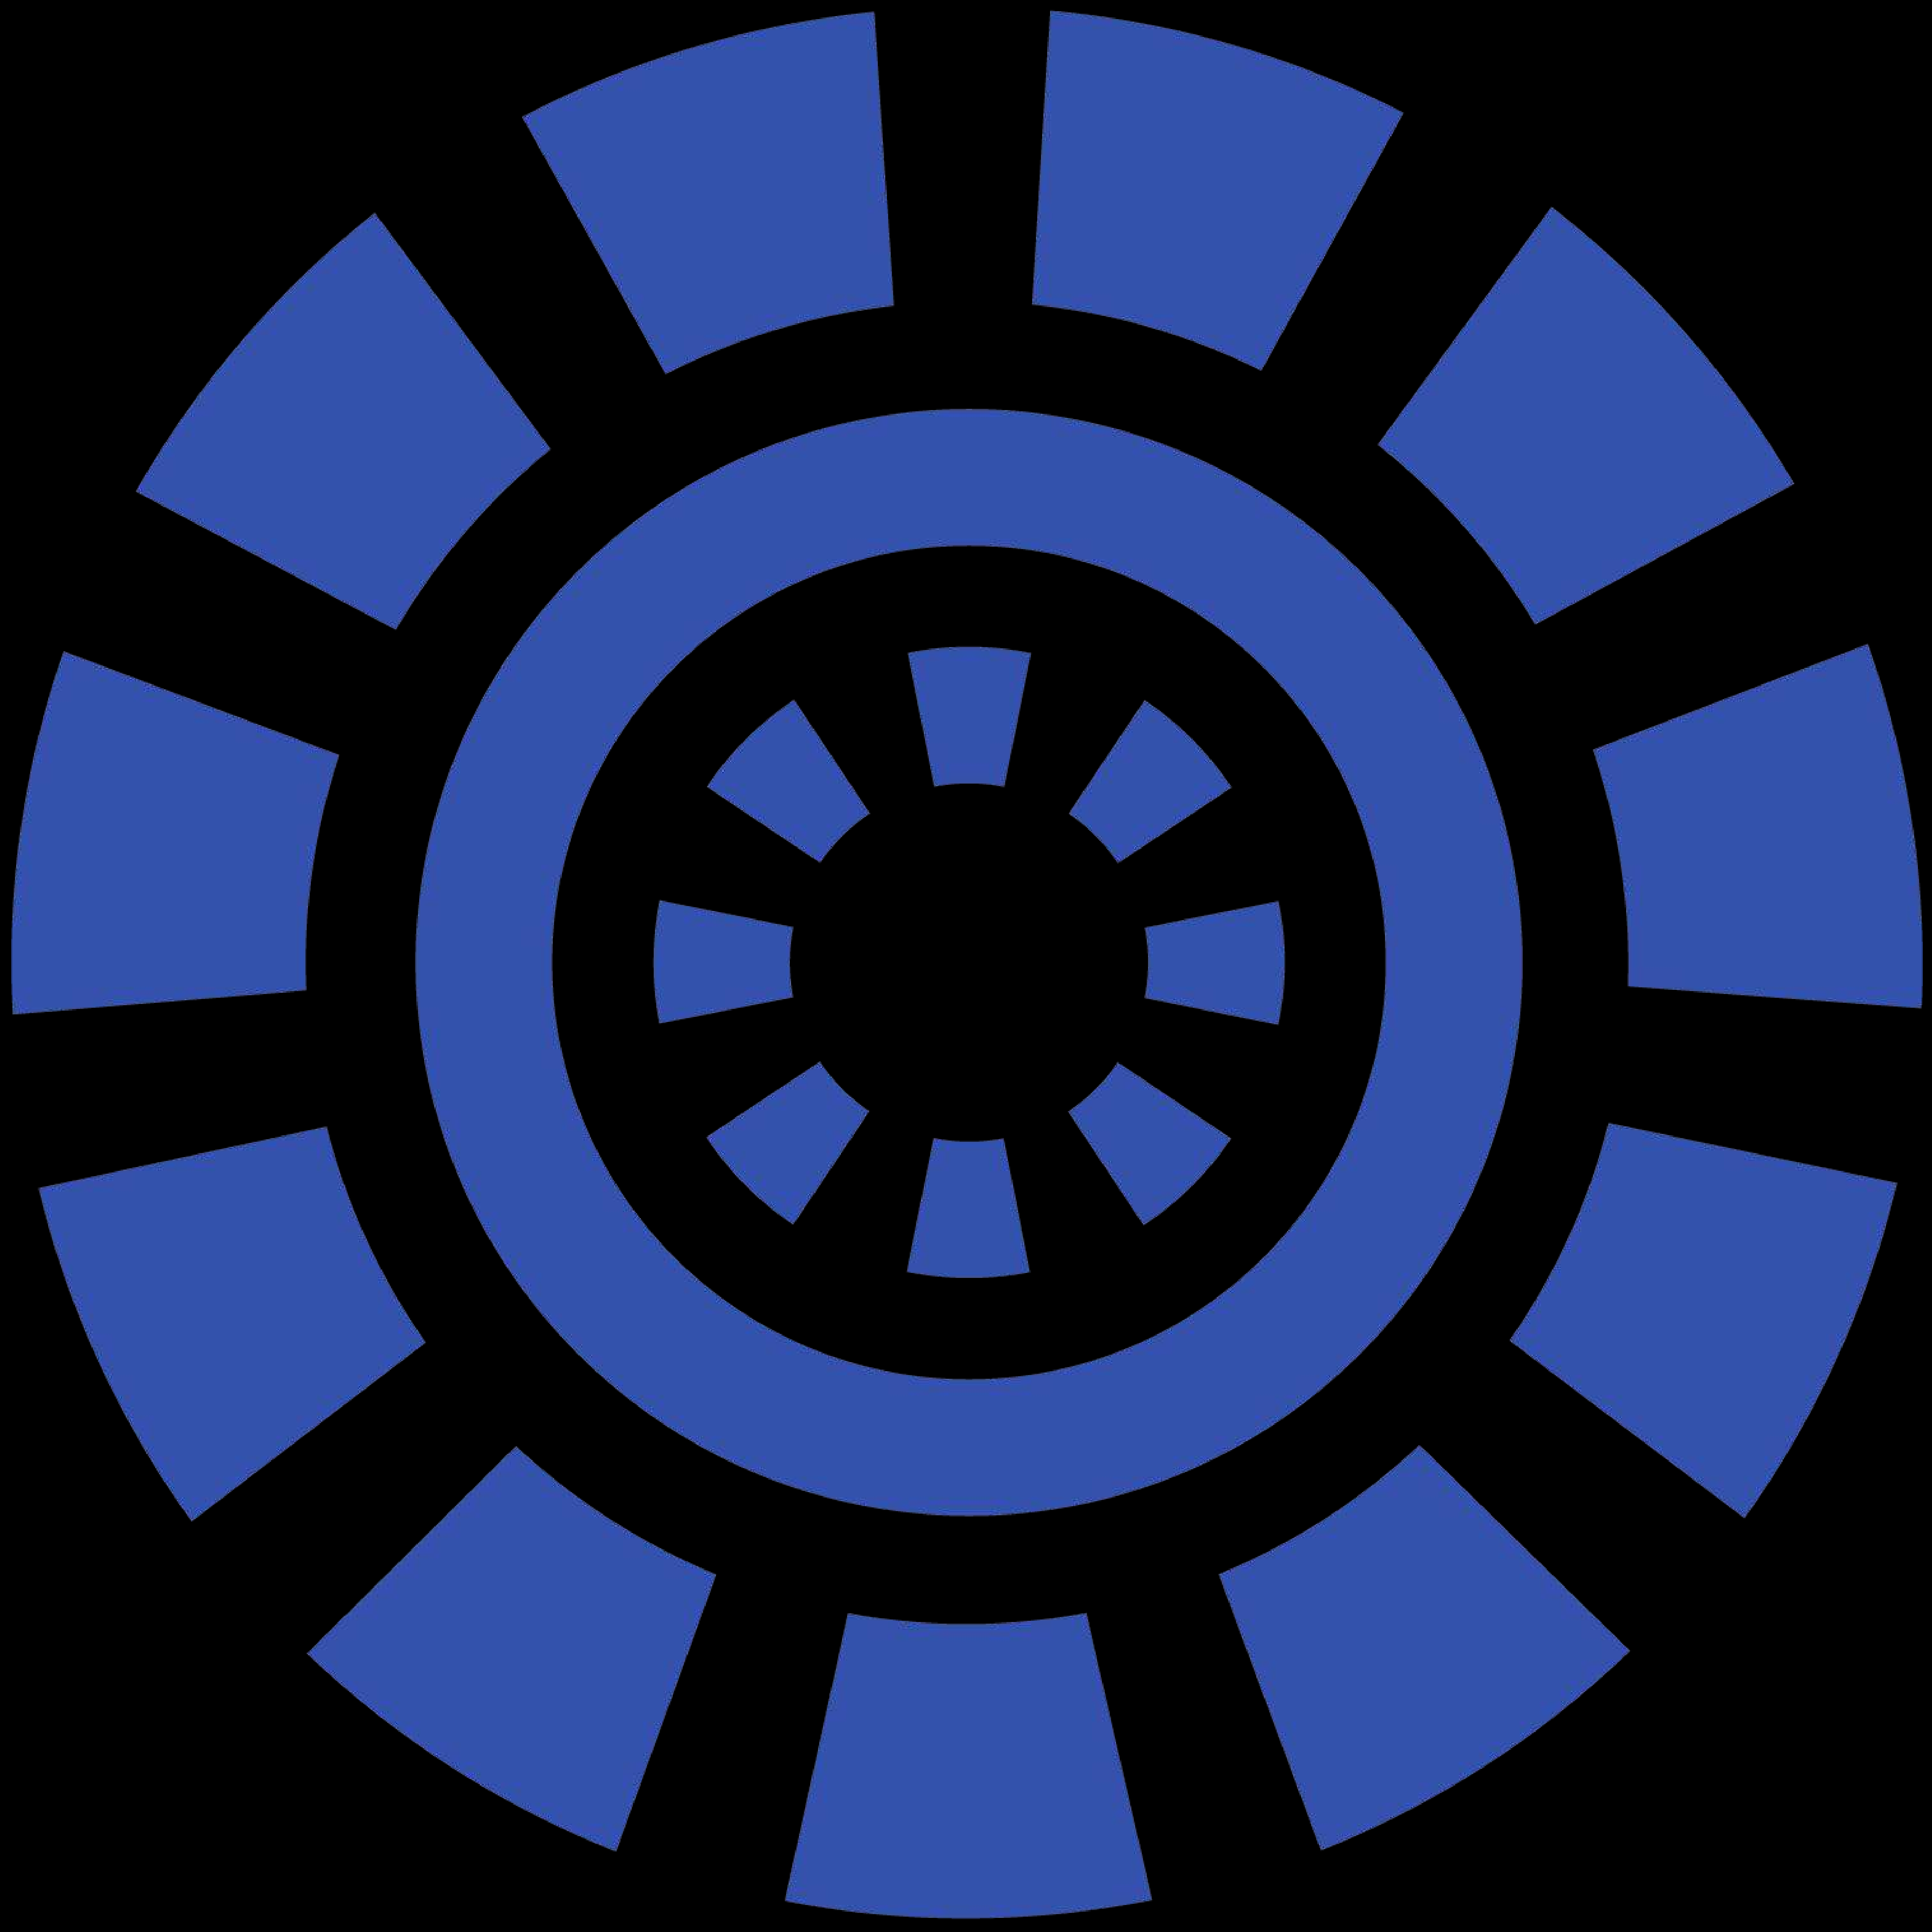

White Paper

v1.0

DRAFT

Justin Cohen + Jason Kaufmann

Table of Contents

MARITIME FUSION

Maritime Fusion makes fusion reactors for ships

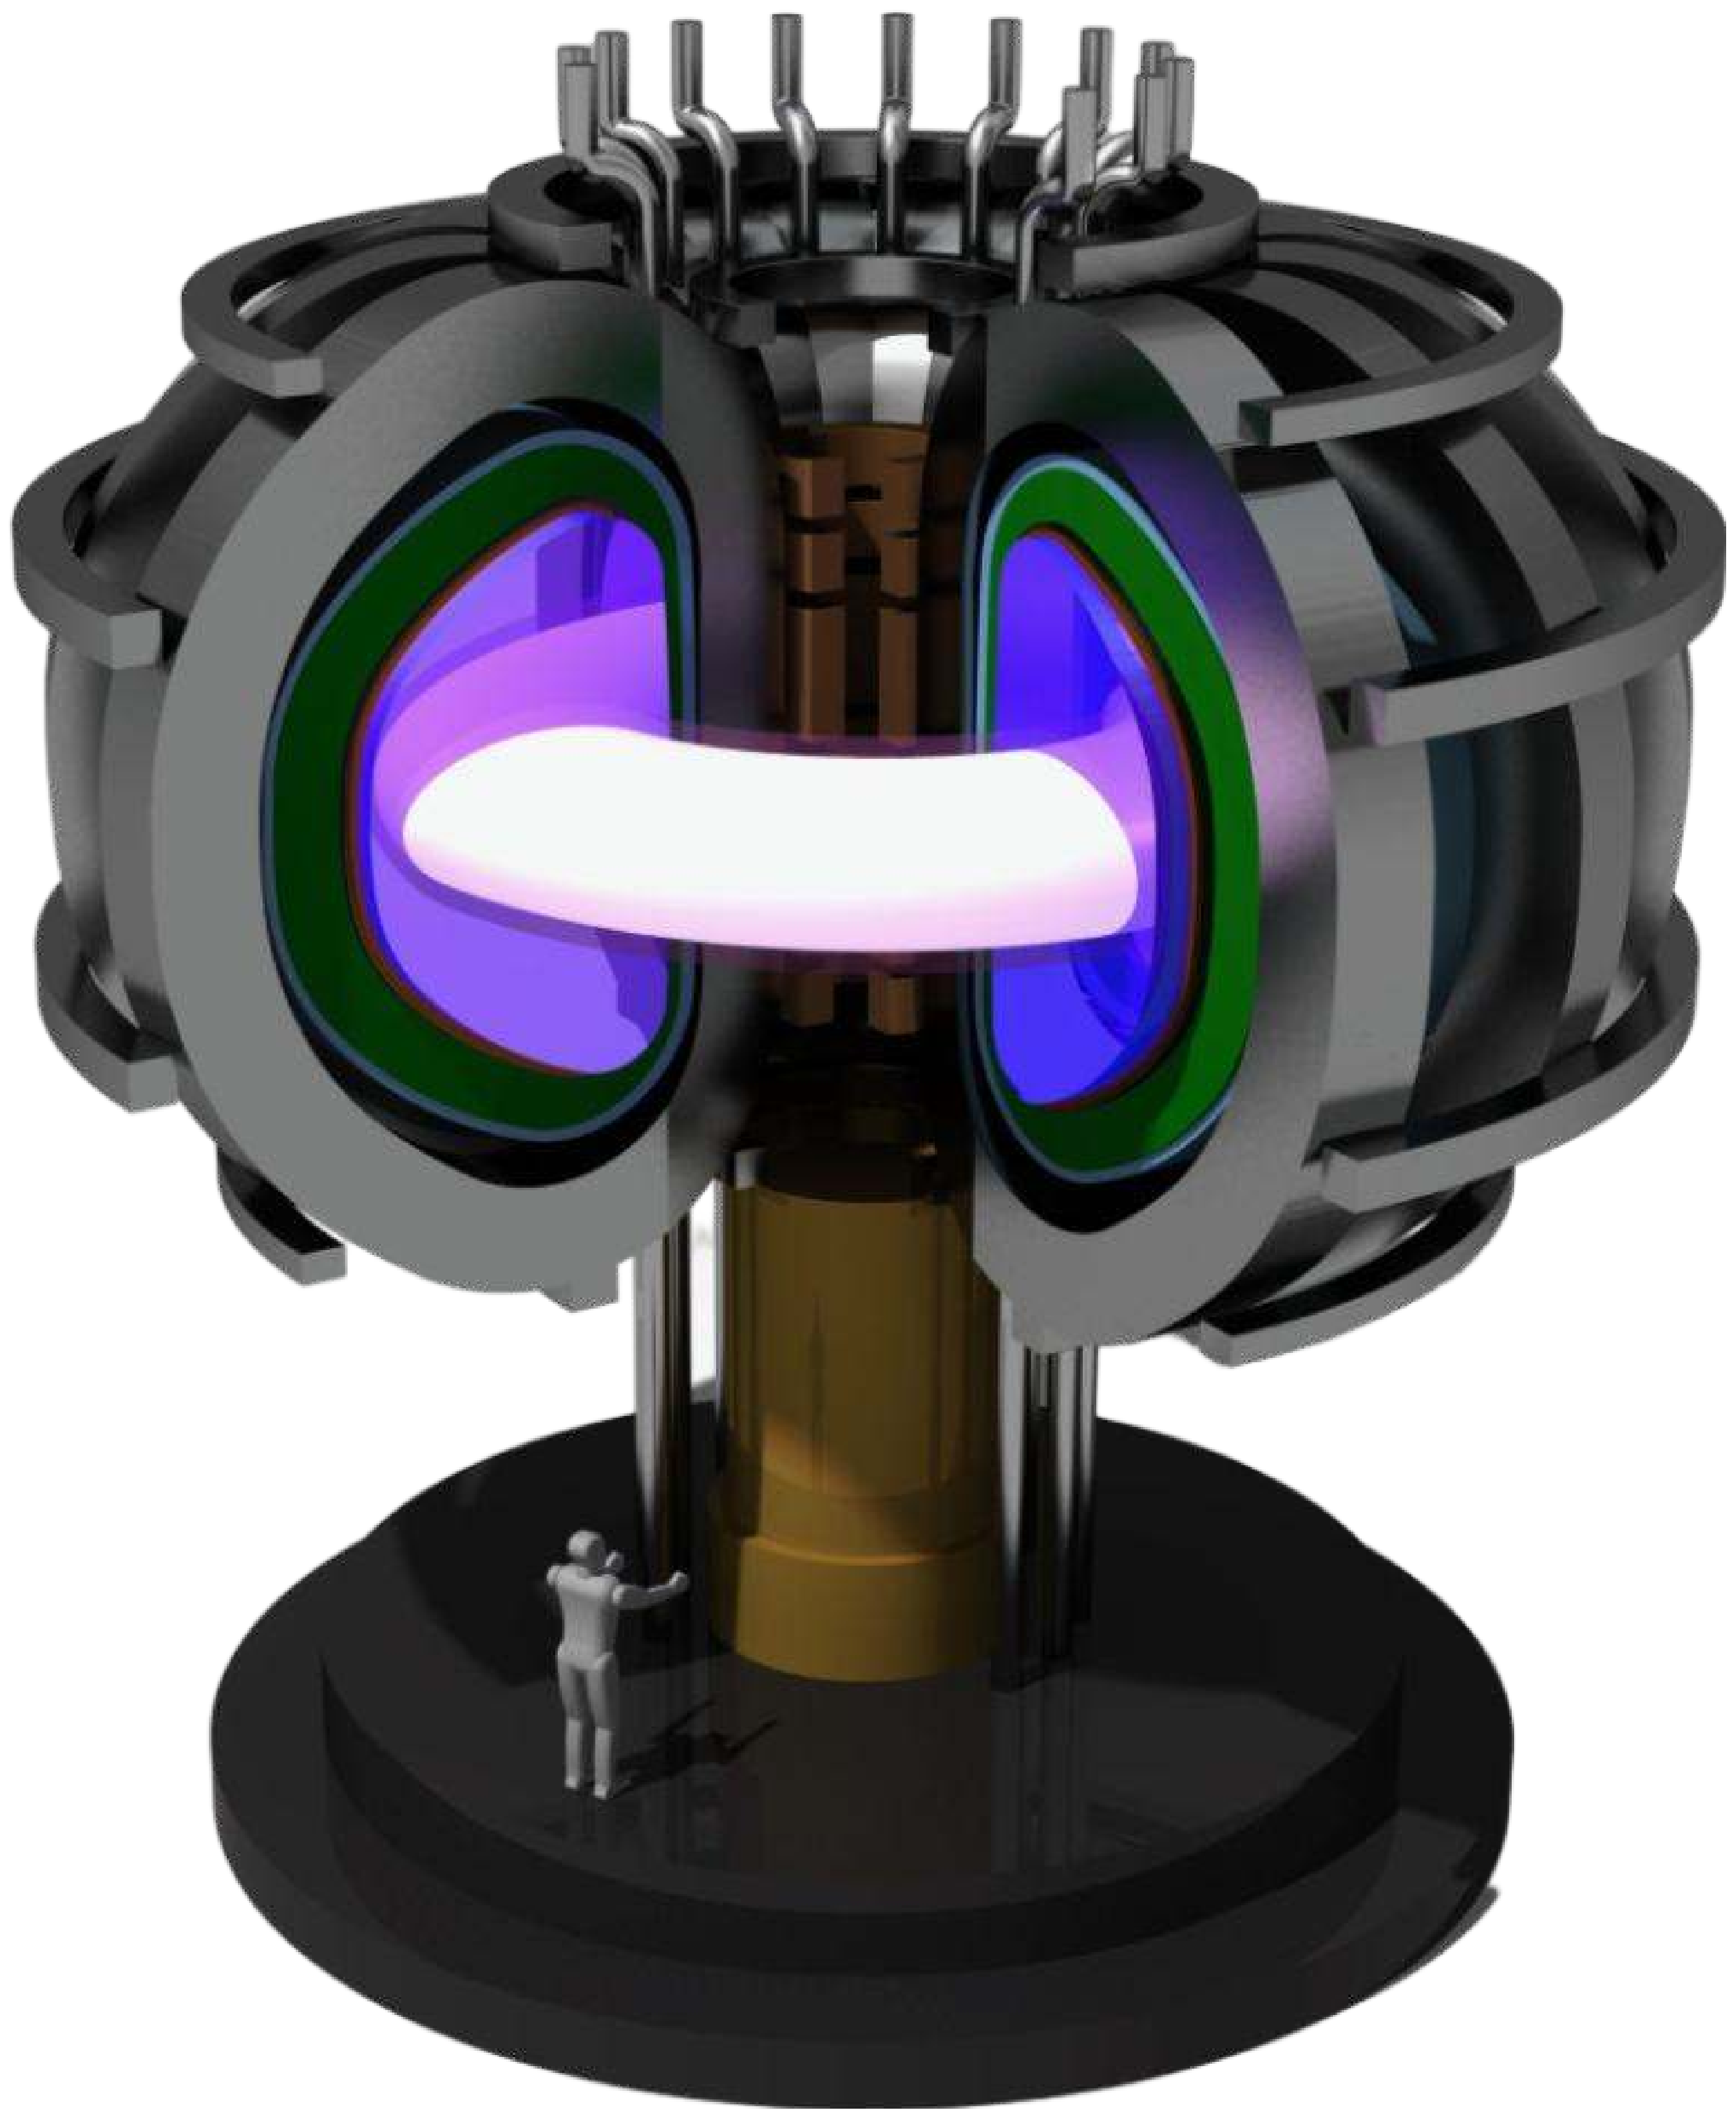

Humanity is on the brink of achieving breakeven fusion, but the first-of-a-kind (FOAK) reactors will be **costly**, **high maintenance**, and have **low capacity factors,** leading to **5-10x higher electricity cost** on the grid. These reactors also face significant physics challenges for steady-state operation - managing first wall heat flux and mitigating nuclear activation of structural components may require novel material breakthroughs to solve.

We're building an HTS (high temperature superconducting) tokamak reactor for **marine applications** that bridges the gap between breakeven fusion and a commercially viable energy source. The market we're targeting **requires 15x less power**, lower up time, and **costs the same** as alternative fuels - **without any emissions**.

Since fusion **does not use highly radioactive fuels** or materials - unlike fission - we sidestep the vast majority of regulatory challenges and safety risks associated with traditional nuclear energy.

We will be the first to deliver a fusion reactor to a paying customer

Top-tier team

Both founders have worked at both SpaceX and Tesla.

- **Justin Cohen** (CEO) - Nuclear Engineering @ NC State and Plasma Physics @ Columbia University, inventor of Cybertruck Range Extender
- **Jason Kaufmann** (CTO) - Electrical Engineering and Physics @ Penn, Intel International Science Fair Winner

Columbia University · Penn University of Pennsylvania · NC State · Tesla · SpaceX

Rapid progress

- Iterated reactor design optimization to yield a 25MWe tokamak ideally suited for our application with **>5x reduction in first wall heat flux**
- Verified magnetic equilibrium and plasma stability at fusion relevant operating conditions
- Iterative neutronics study proving the **reduced radiation fields the structures and magnets are exposed to**
- Completed first pass structural and HTS magnet design
- Demonstrated feasible **divertor survivability** via simulation
- Architected **full scale cryogenic system** applicable for deployment on ships
- Spec'd out remaining auxiliary systems (RF, tritium fueling, power supplies, energy storage) and how they fit on a ship
- Modeled **economic viability** of fusion powered vessels vs alternatives
- Targeting **2032** reactor deployed on ship

Chapter 1

Introduction

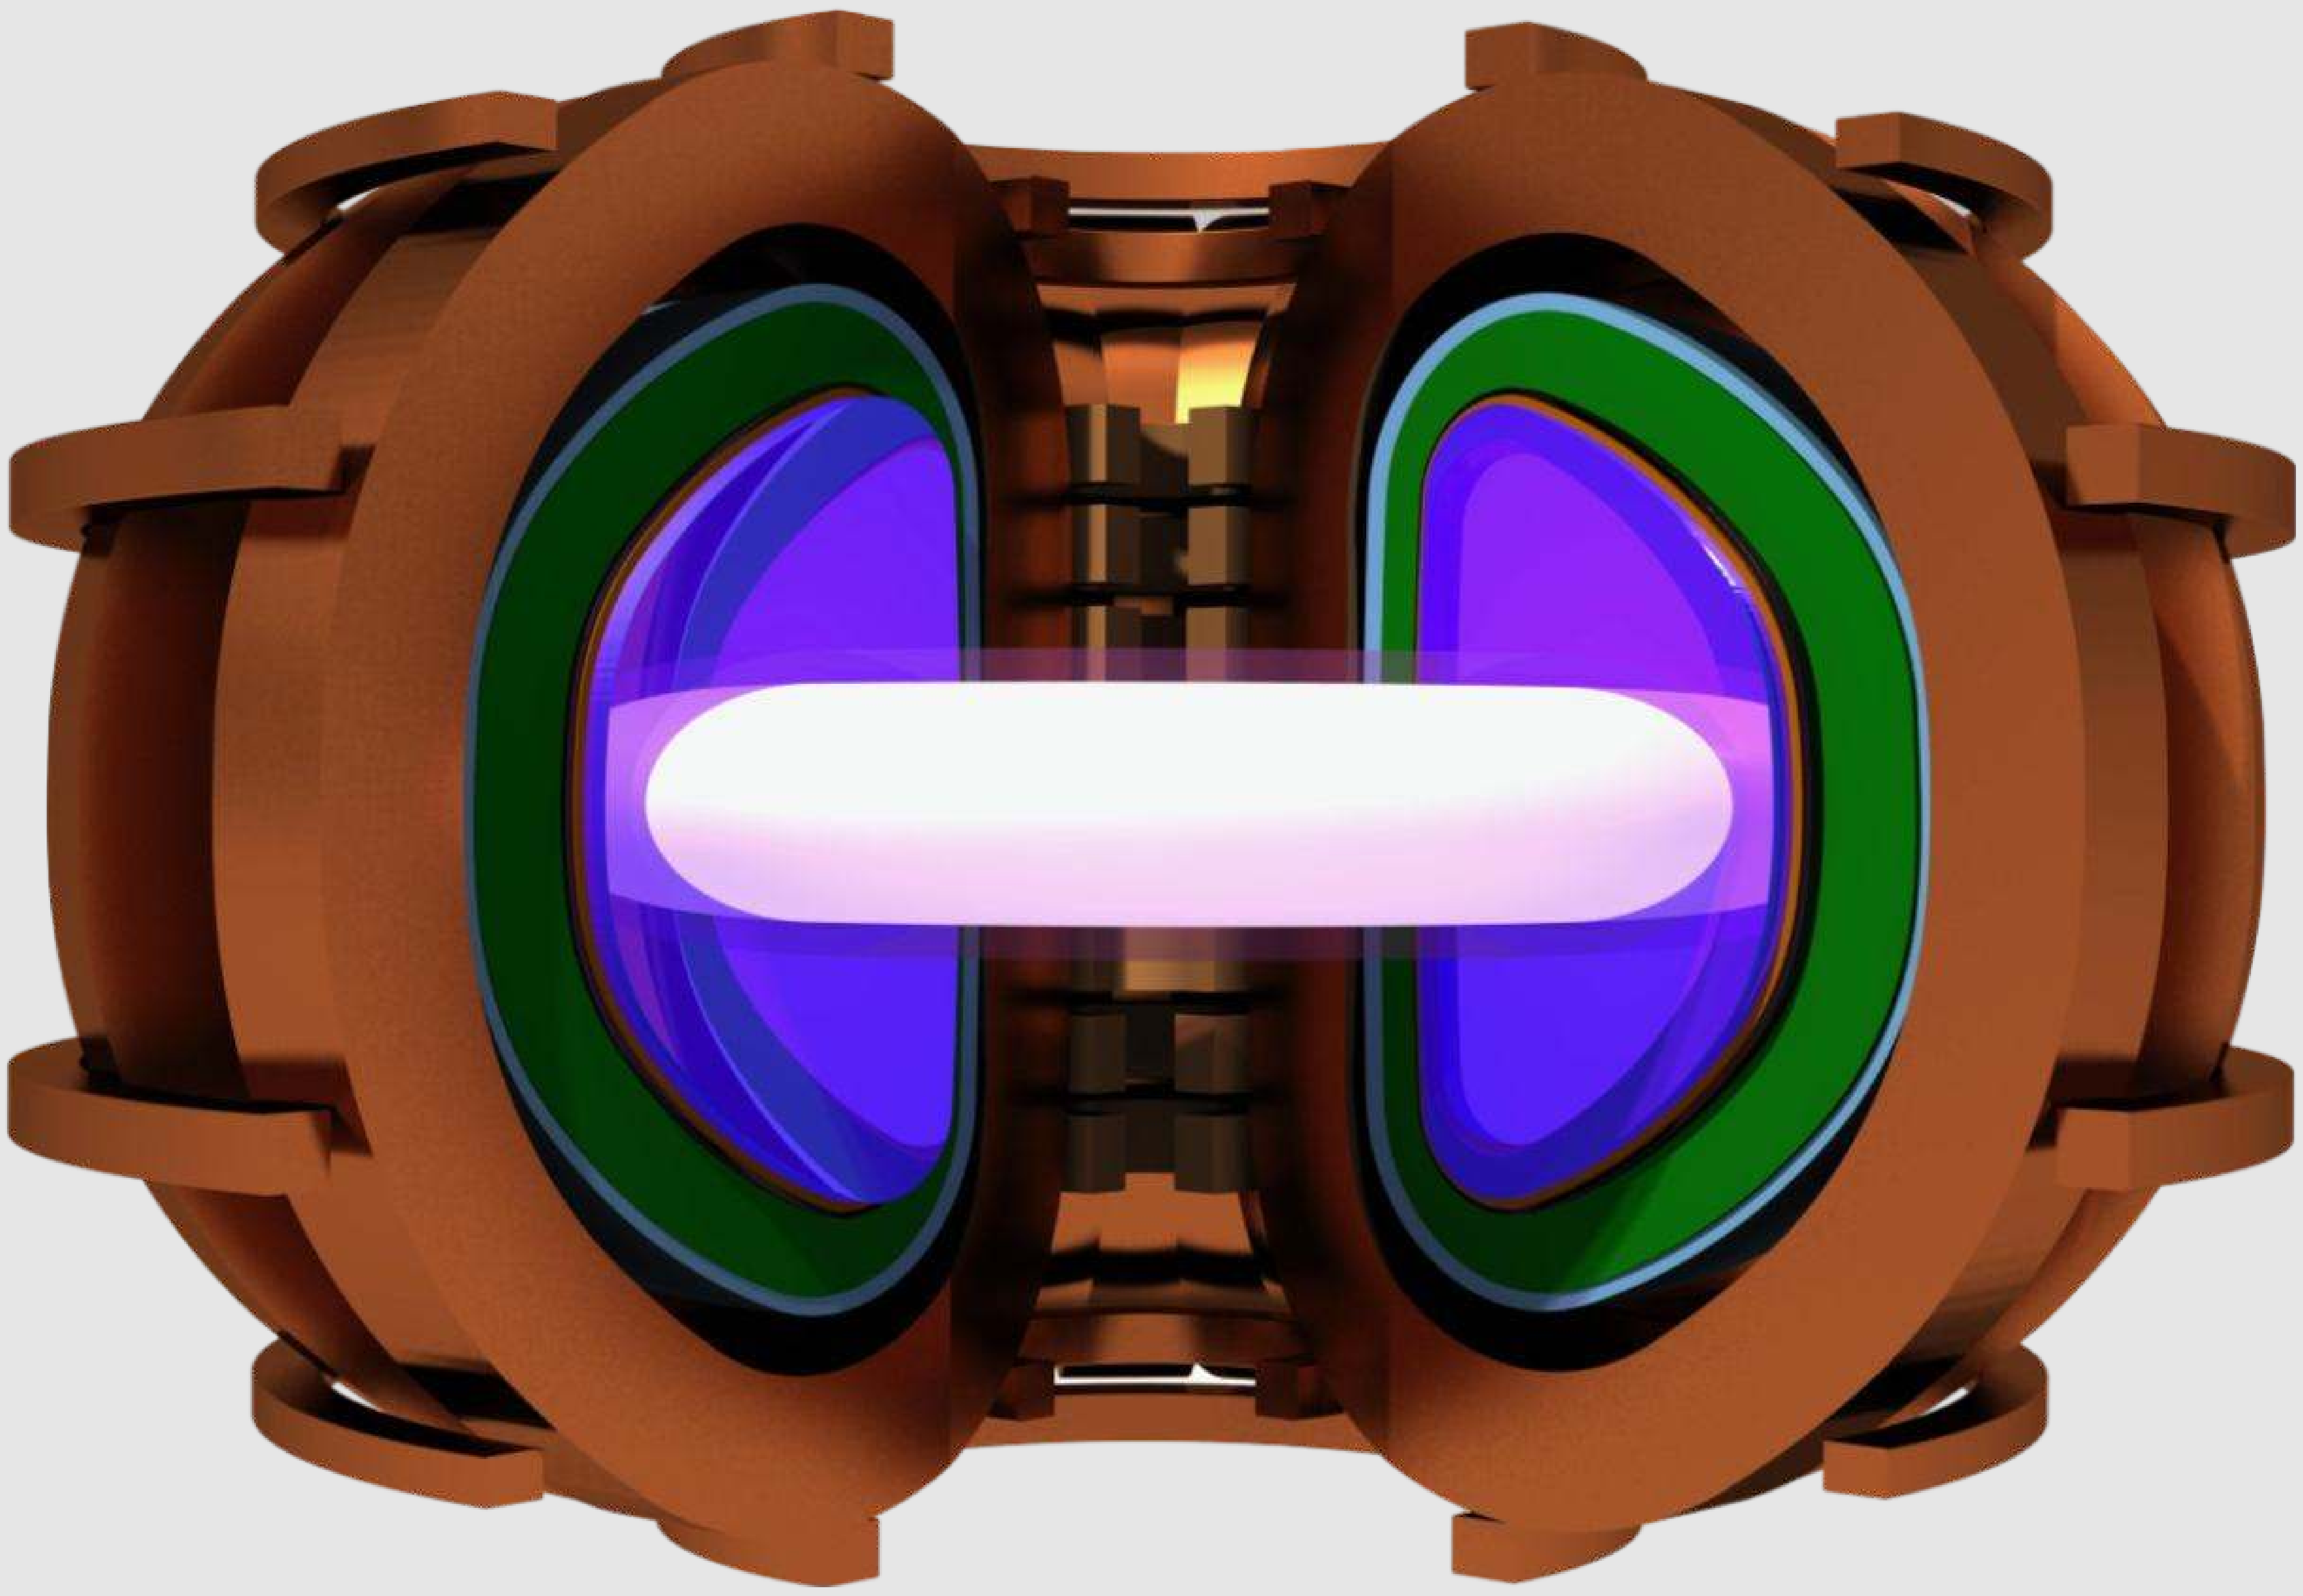

Introduction

Breakeven fusion is coming soon

Fusion has long suffered from being designated as “***the energy source of the future, and it always will be***” … but this is ***no longer true***.

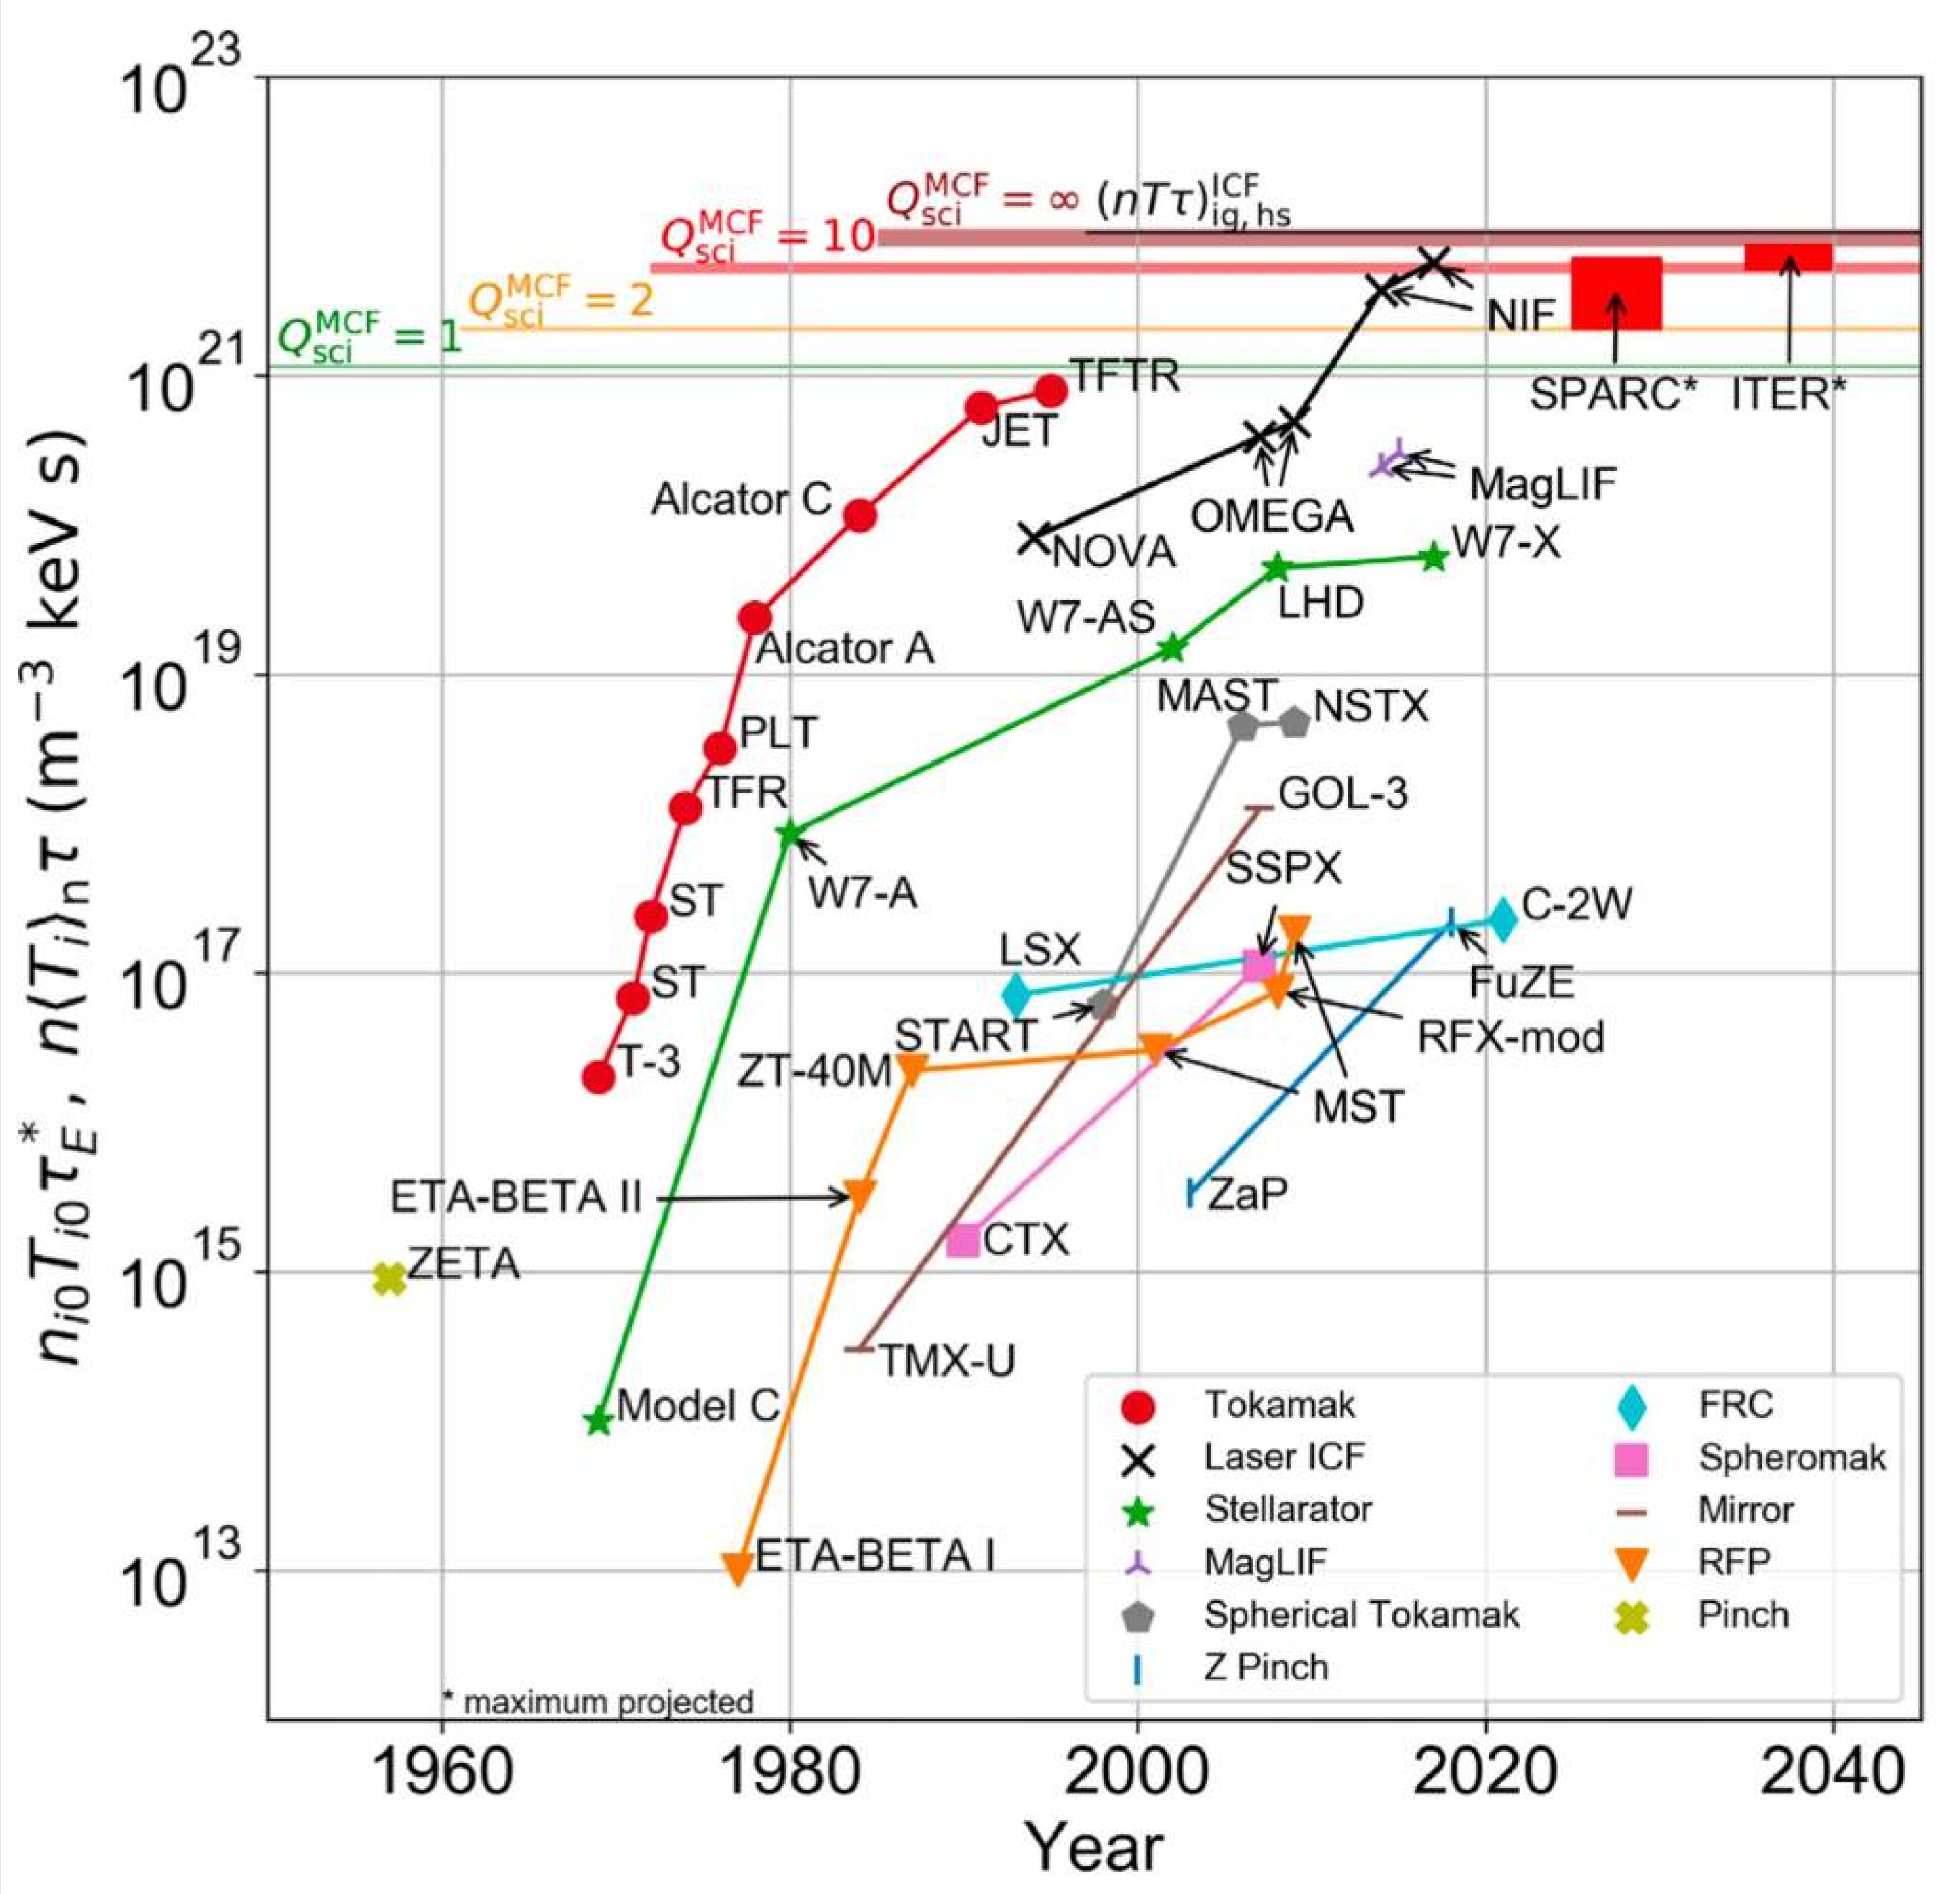

Source: Phys. Plasmas 29, 062103 (2022)

Why now?

For the first time in history, commercially available high-temperature superconductors (HTS) provide the final breakthrough needed for leading magnetic confinement fusion approaches to achieve net energy. This became possible only in the early 2020s with the global ramp-up of HTS tape production.

Since the 1960s, fusion research has advanced primarily in government and academic labs. However, as commercialization nears, **private ventures are now the frontrunners in achieving breakeven energy production.**

Fusion’s lightbulb moment is near. The National Ignition Facility (NIF) offered a preview in 2022 by achieving scientific breakeven, but the real “**moon landing**” will come in 2027 when Commonwealth Fusion Systems (CFS) delivers a **10-second Q>1 pulse on the SPARC tokamak**.

Over $13 billion in private funding, almost all in past 5 years

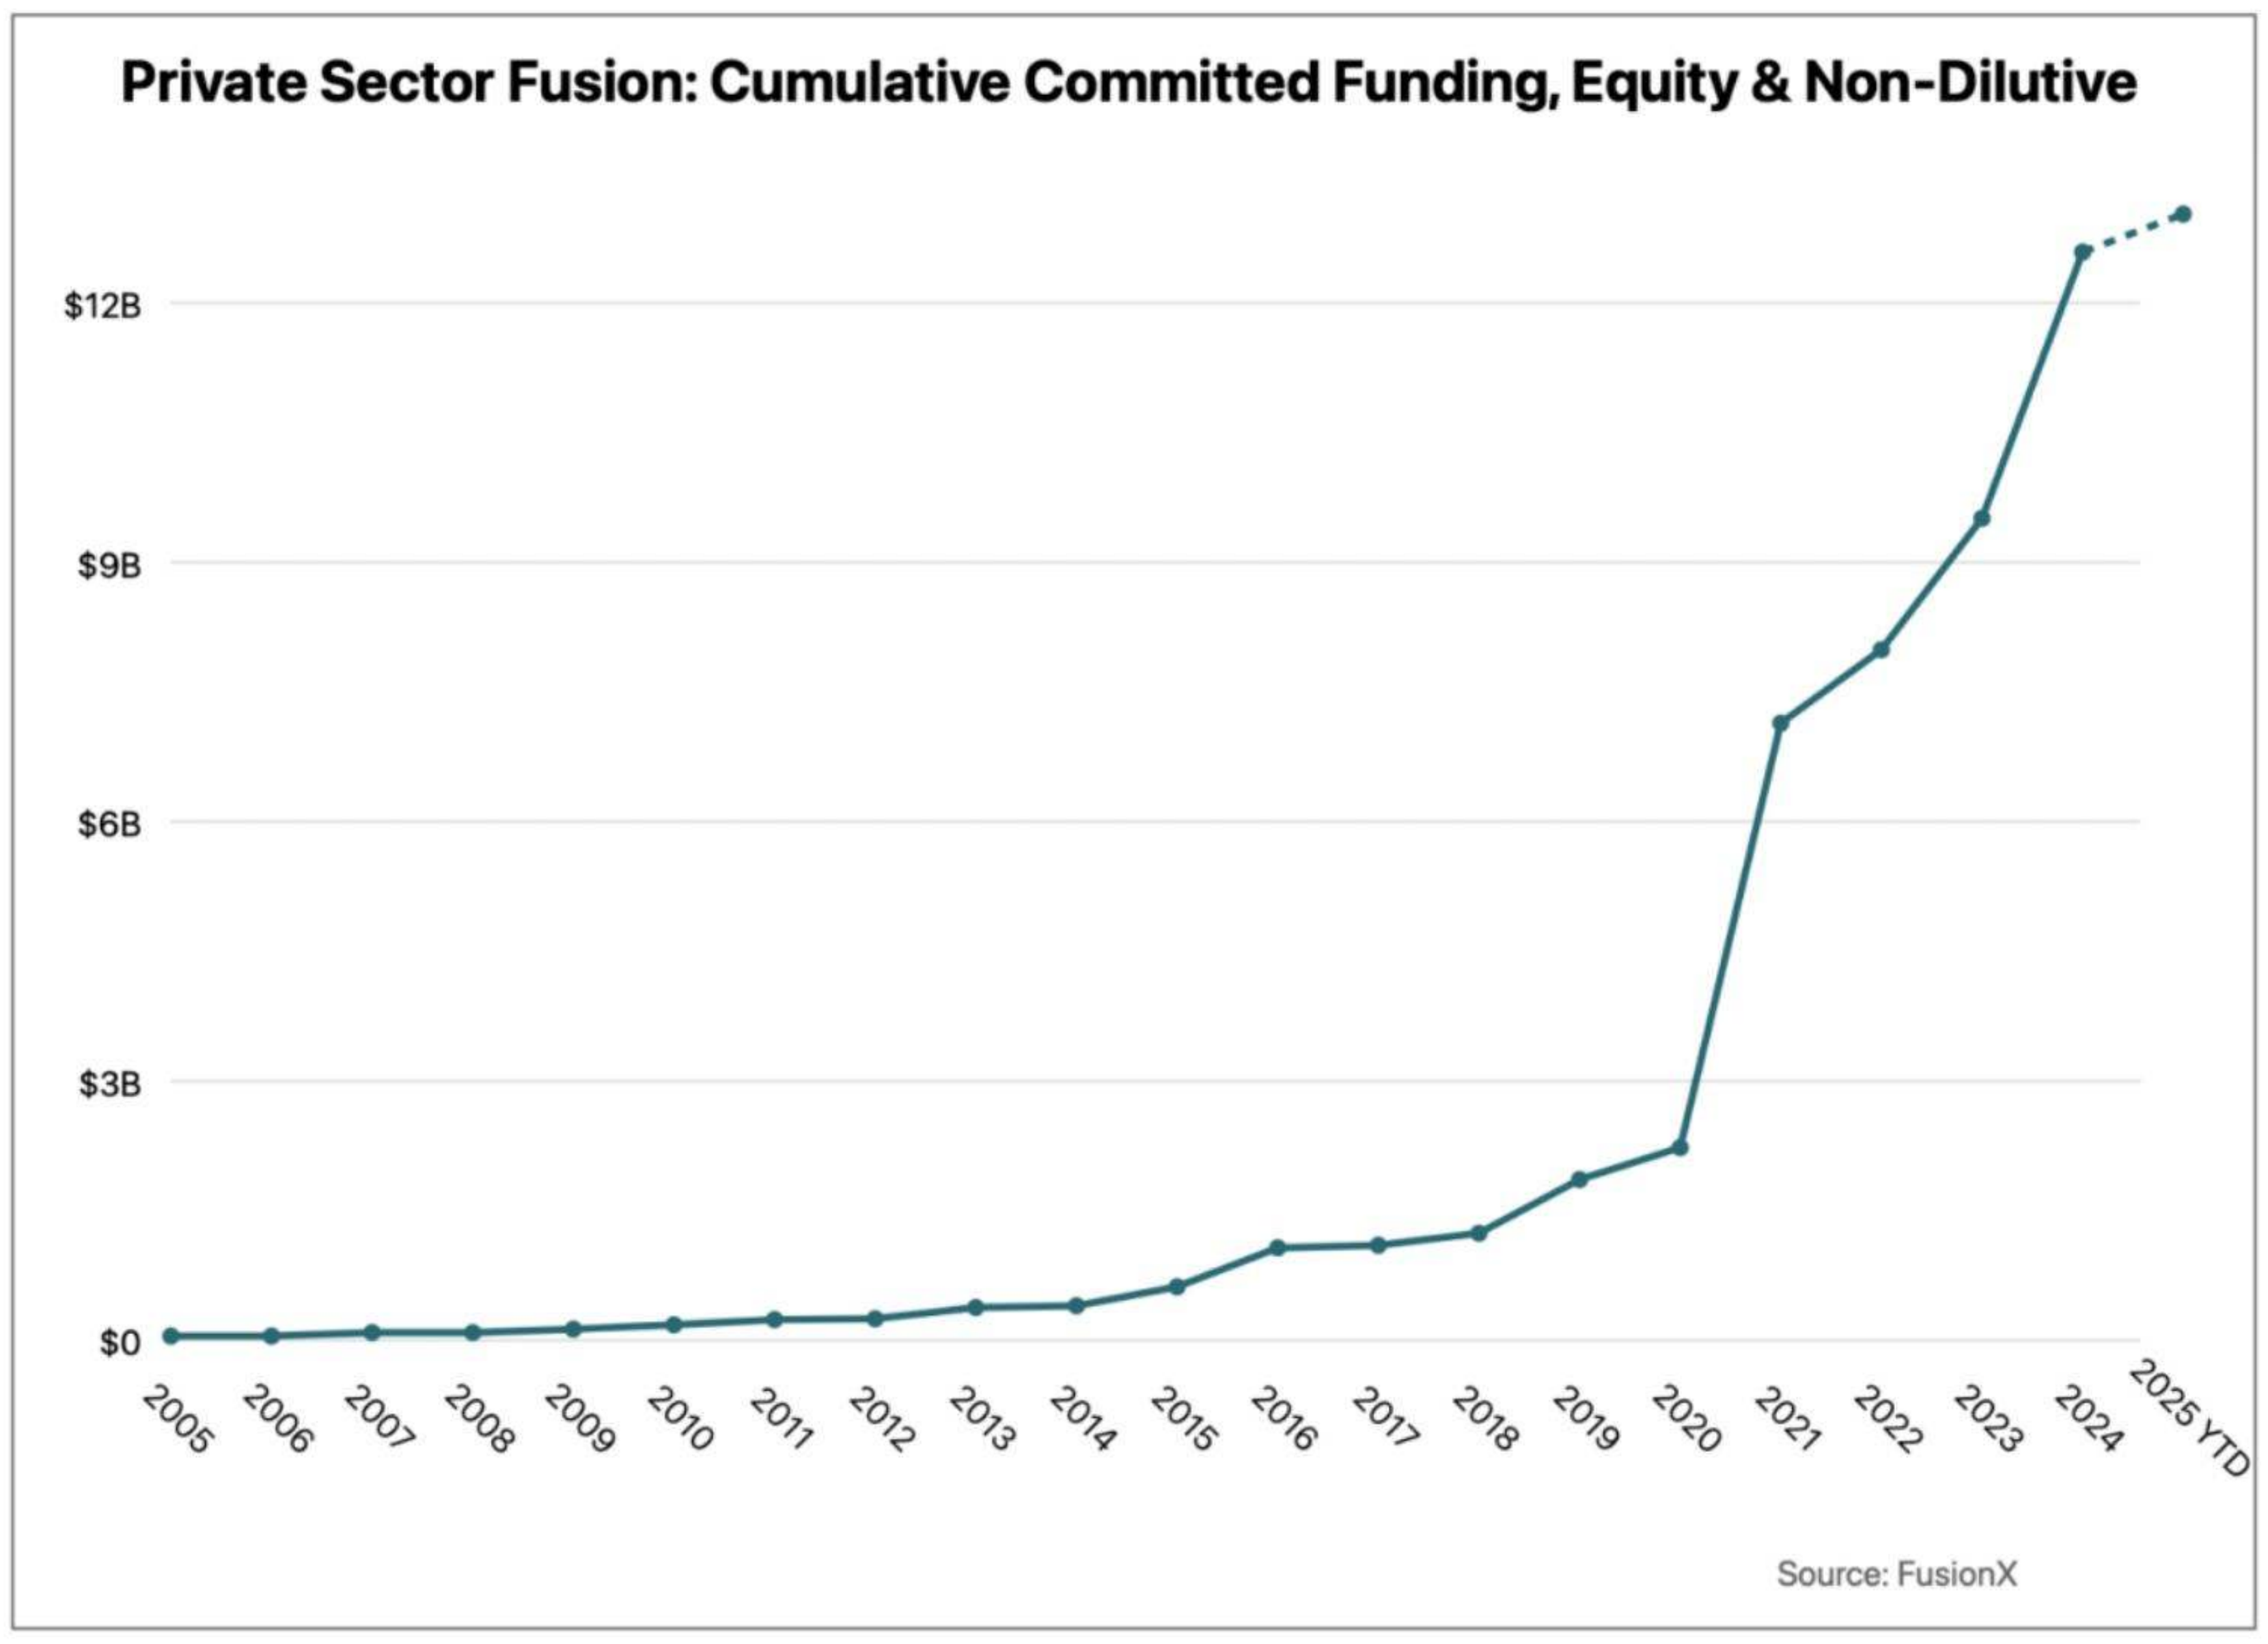

Source: FusionX - Funding Fusion: The State of the Market (2025)

Experts agree: Fusion is within reach.

> *"People see fusion [and think], 'Oh, it's still far away,' or [that] it's been slow progress. It's like when my godmother was born—people had no idea how stars worked... Everybody knew fusion was 40 years away, and **now it's four years away**."*
>
> - Dennis Whyte
>
> *Director of MIT Plasma Science and Fusion Center / “Father” of fusion*

Source: Lex Fridman Podcast #353 (2023)

HTS tokamaks are the most proven path to a power plant

A bunch of fusion approaches have been proposed in the last 60 years:

- Tokamak
- Stellarator
- Z pinch
- Laser ICF
- MagLIF
- FRC
- Mirror
- RFP
- Pinch
- Spheromak

However, tokamaks are the sole fusion approach that have sustained Q=0.67 (JET, 1997) with proven, scalable physics to hit Q>10.

While other methods remain largely experimental, tokamaks stand as the leading contender for achieving true breakeven power within the next decade.

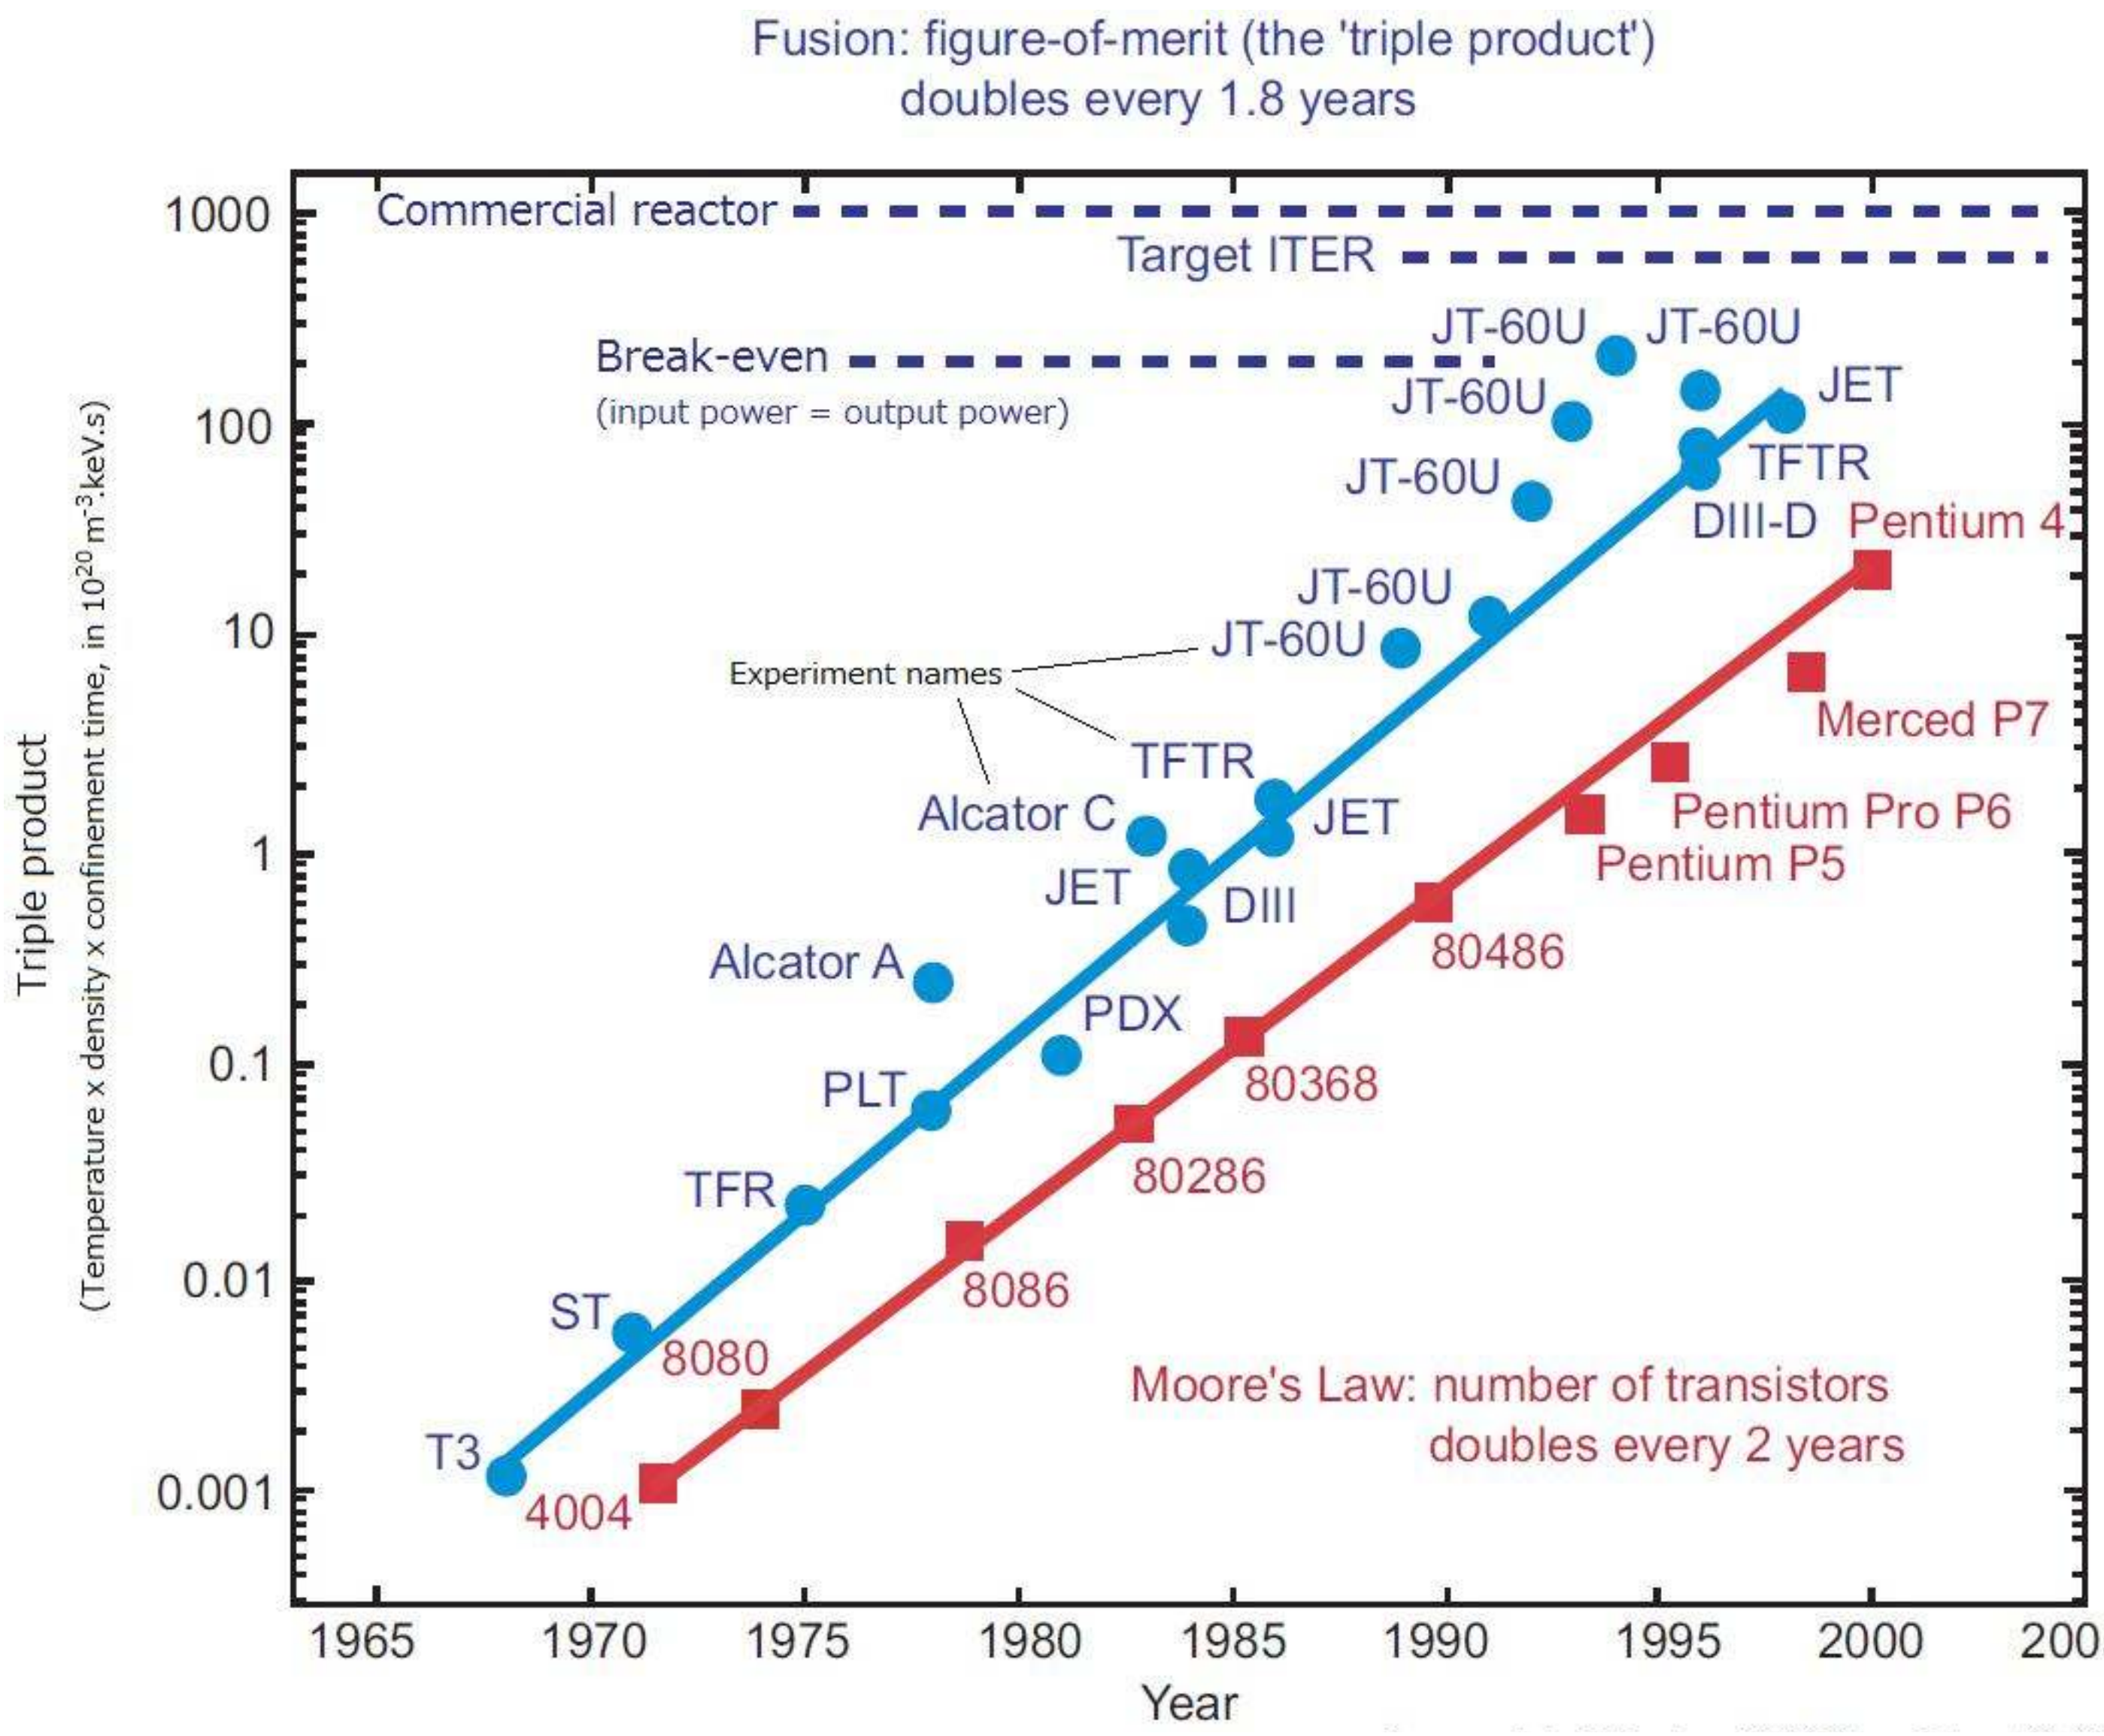

Source: A.J. Webster, 2003 Phys. Educ. 38 135

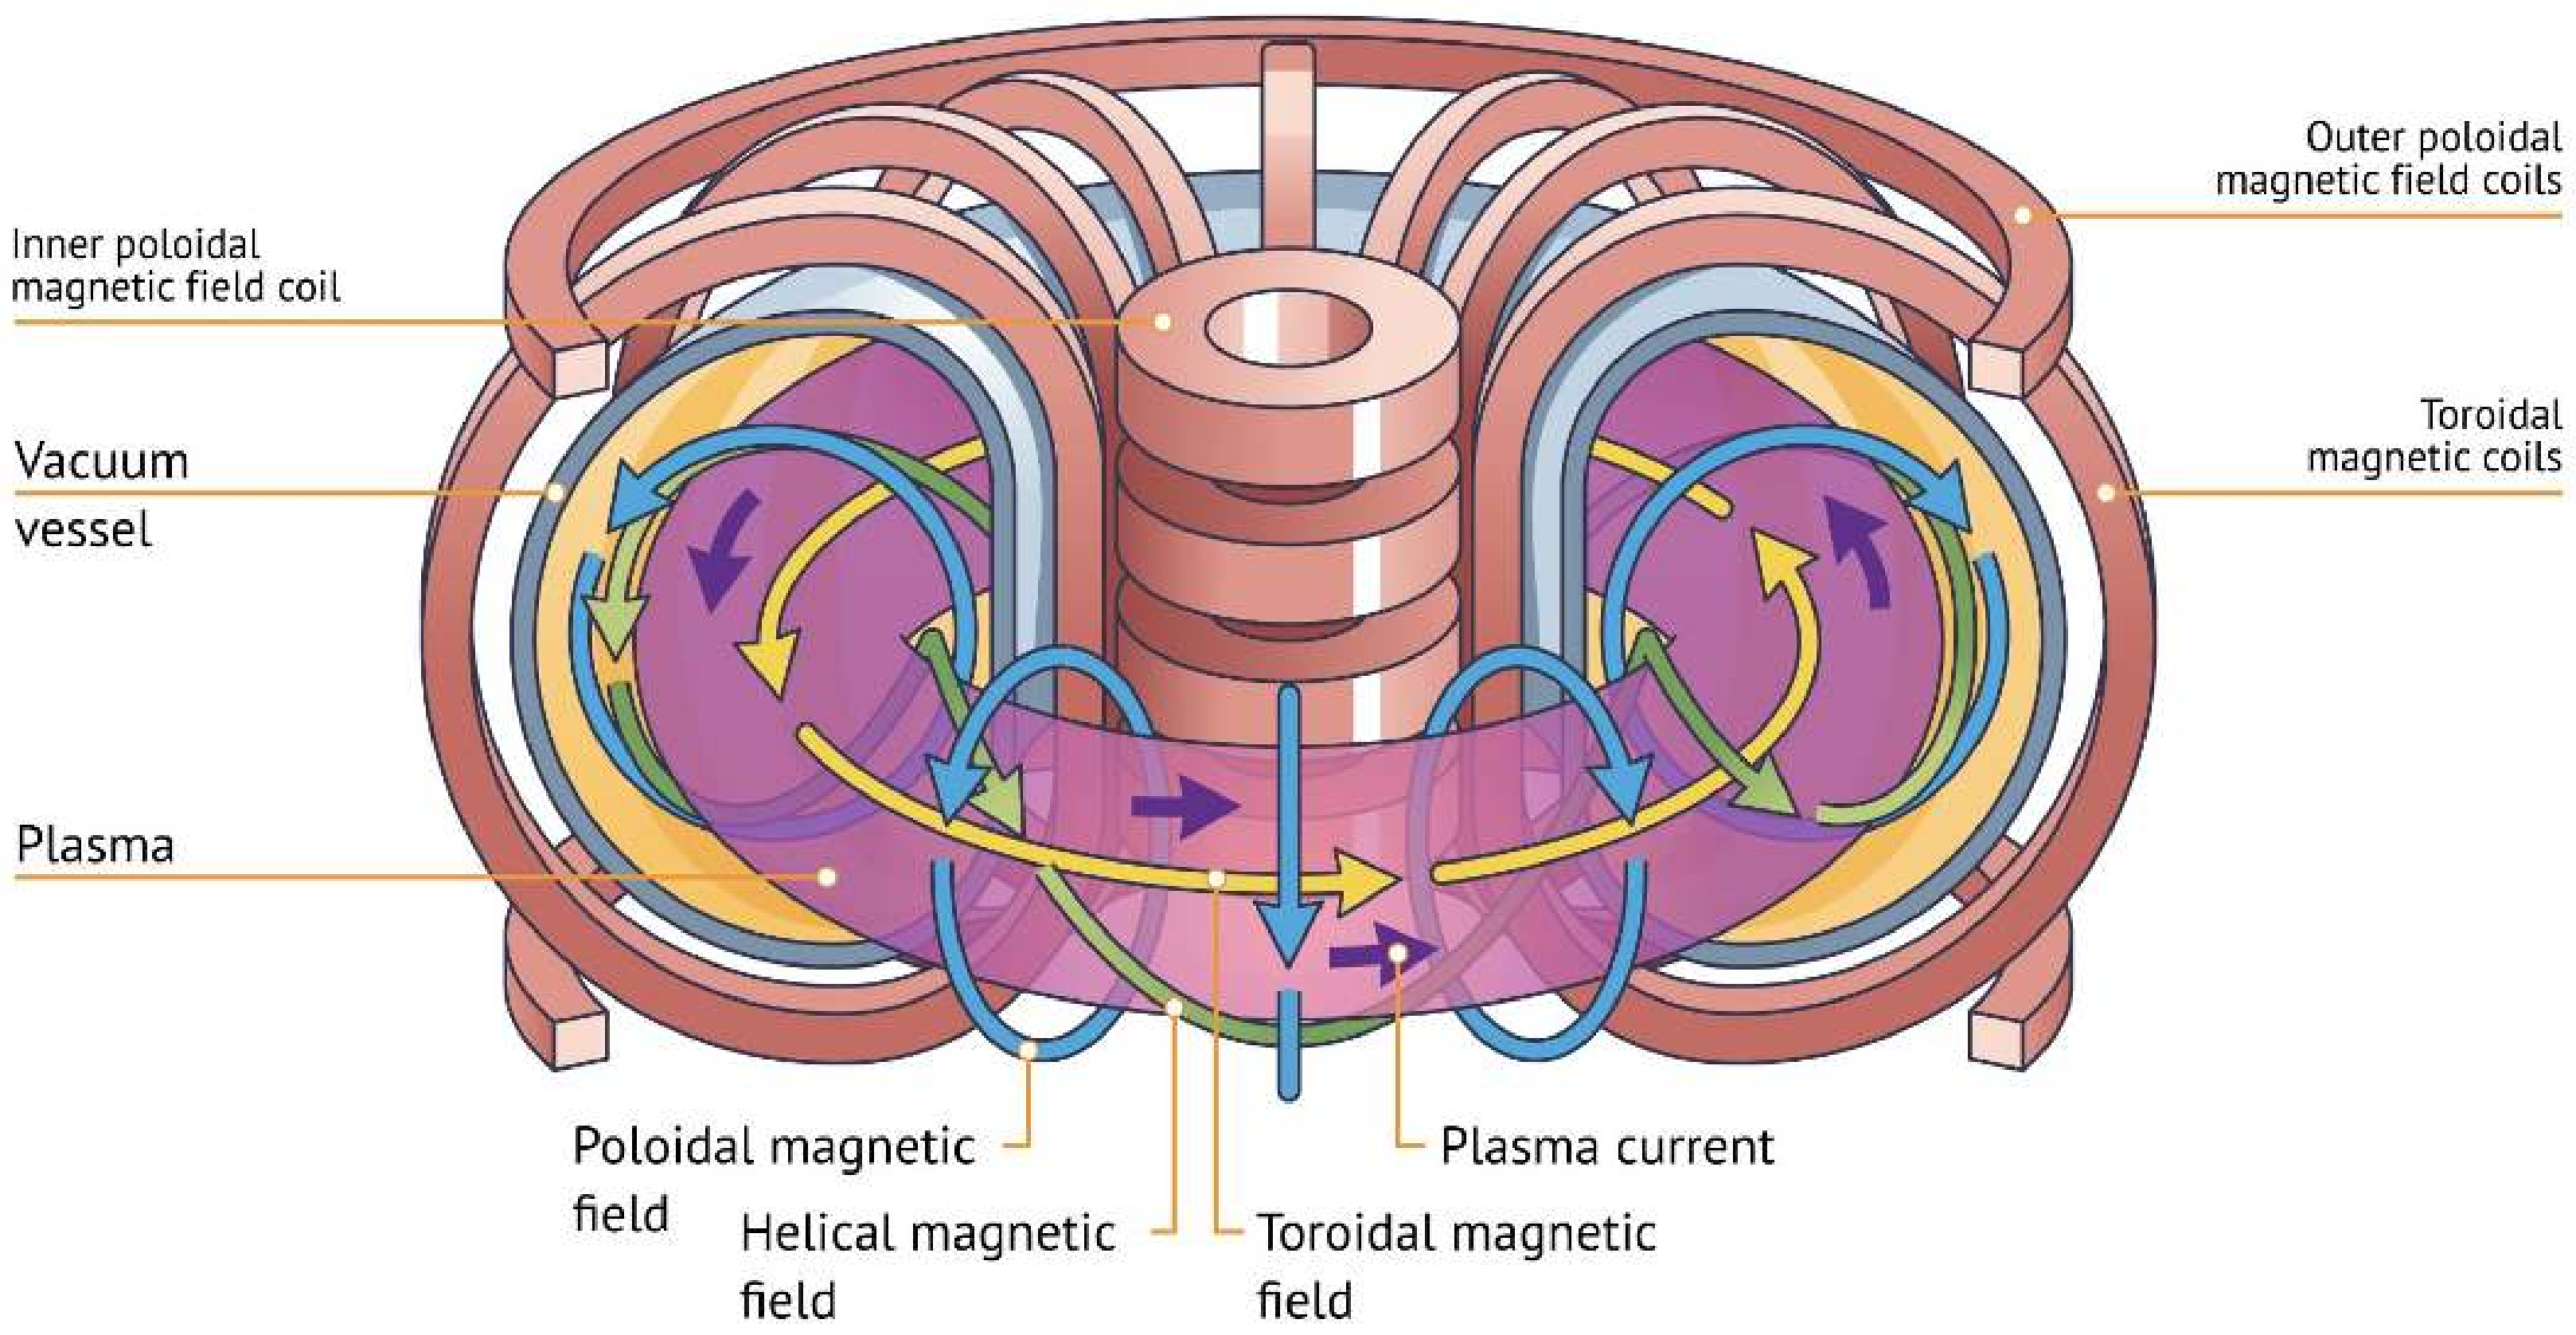

As a good analogy, the triple product for tokamaks was matching Moore's Law until 2000, but hit a wall. Just as transistor scaling may be reaching its physical limits now and a breakthrough is needed, tokamaks hit their peak earlier. However, that breakthrough has come: ***high temperature superconductors***.

$$P_F \propto \beta_t^4 B_{T,0}^4$$

$B_{T,0}$ is TF field at plasma center

Source: https://iopscience.iop.org/article/10.1088/1741-4326/ad2425/pdf

JET (the tokamak with the current highest Q) has a centerline B field of 3.5T and has achieved a Q=0.67. The advent of HTS at scale allows for roughly a ***3-4x increase*** in B field to 12-13T. So, if the scaling law holds where fusion power increases with B field to the fourth power (see equation above), HTS tokamaks should be able to hit Q_{sci}>10, enough to deliver a net-positive-power fusion reactor.

Why?

1. **Demonstrated Physics and Performance**
   Tokamaks have achieved high fusion gains (Q) in experiments like JET, TFTR, and JT-60, consistently demonstrating effective high-temperature plasma confinement and the ability to sustain Q~0.5 conditions even with LTS magnets.
2. **Extensive Research and Global Collaboration**
   Decades of R&D and well-established theoretical models, along with international projects like ITER, have built a mature ecosystem that minimizes risk and paves the way for commercial viability.
3. **Proven Plasma Confinement with HTS Advancements**
   Modern HTS magnets offer higher magnetic fields, improved plasma stability, and more compact reactor designs, reducing operational costs and overcoming previous limitations.
4. **Quasi Steady State Operation**
   Tokamak based power systems with inductively driven plasma currents can achieve sufficiently long discharge pulses, on the order of 10's of minutes. In contrast, ICF and pulsed power systems are limited to discrete shots—NIF, for example, achieved Q_{sci}>1 for only ***mere picoseconds per shot*** and can ***only be fired a few times per day***.
5. **Near-Term Viability**
   While alternative fusion concepts face challenges related to plasma stability and experimental validation, HTS tokamaks are significantly further along, with a focus on power plant design and engineering, making them the most promising near-term option practical fusion power. The cousin of tokamaks, the stellarator, is another viable option but their complex coil geometries and associated manufacturing challenges make it an unlikely candidate for commercial viability in the near-term.

However, this doesn't mean tokamaks have fully solved fusion —key challenges remain...

Chapter 2

Problem

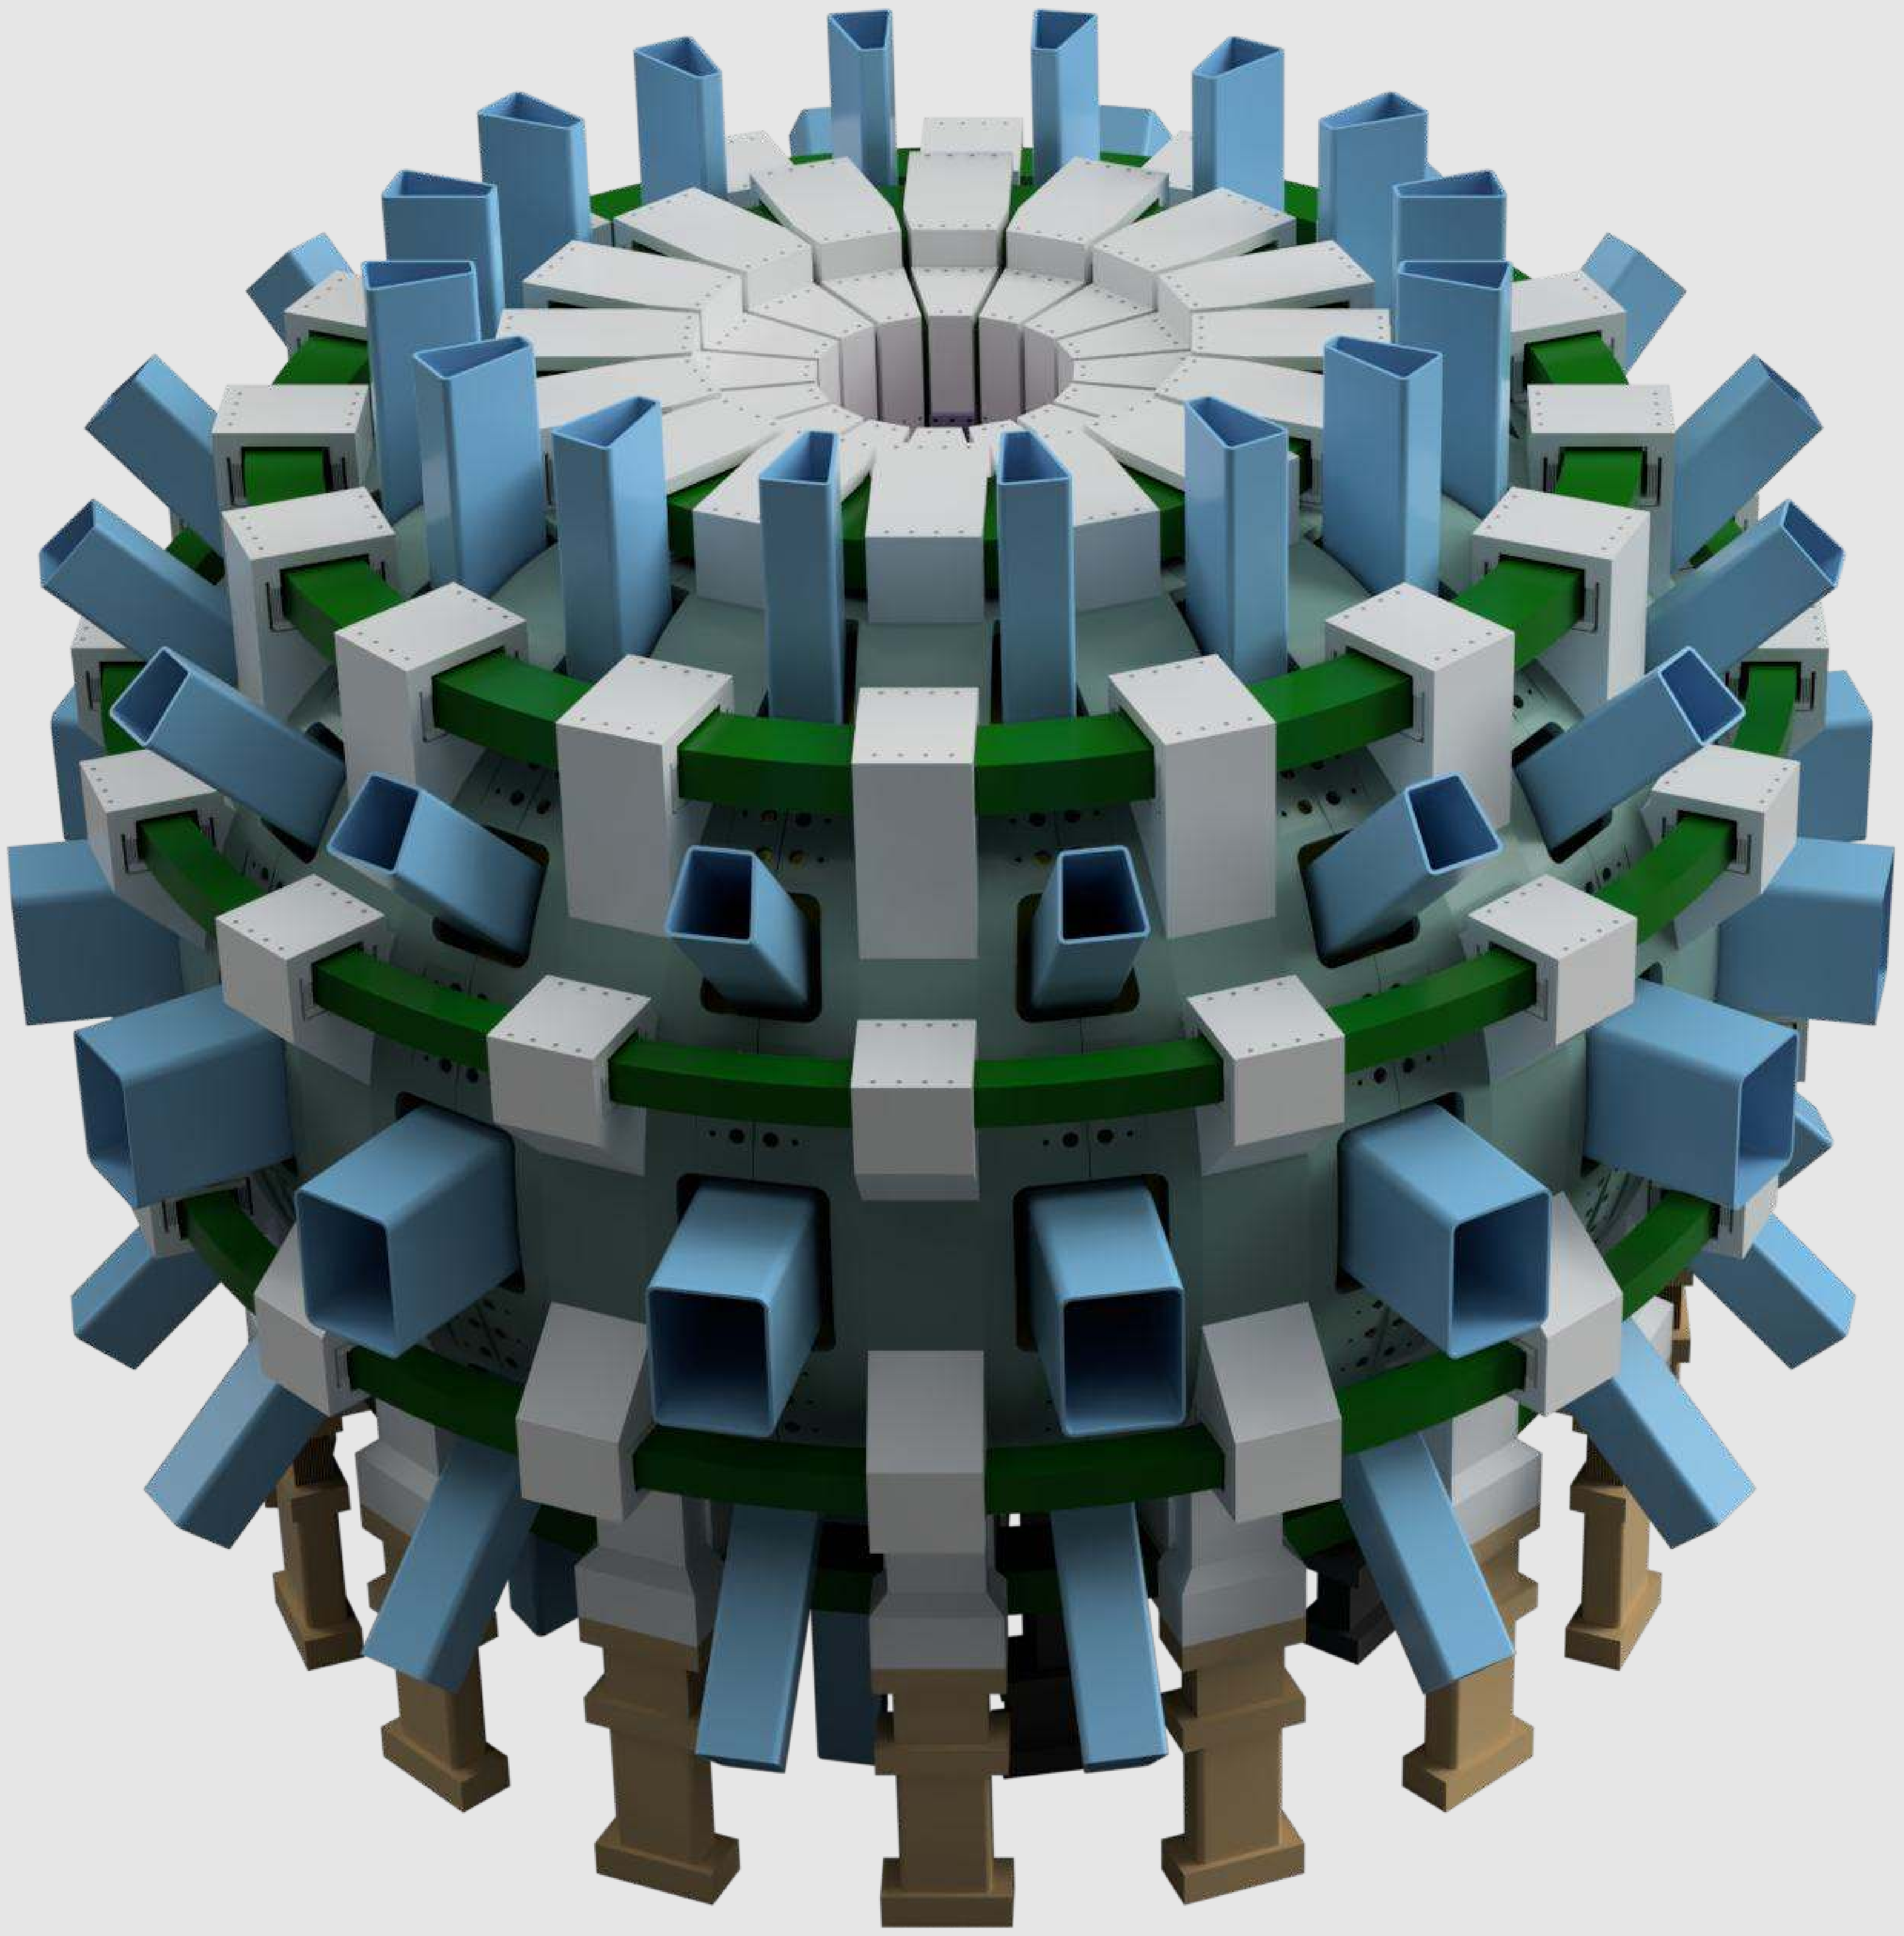

Problem

The grid is a brutal first market to enter

FOAK fusion reactors on the grid are projected to have an **LCOE (Levelized Cost of Electricity) 5–10× higher than current energy sources**, primarily due to steep upfront capital costs—$3,500–$8,500 per kW versus $900–$1,300 for natural gas and $1,000–$1,500 for solar—along with higher operational and maintenance expenses from advanced materials and reactor wear.

Overnight Capital Expenditure
ITER (500MW)

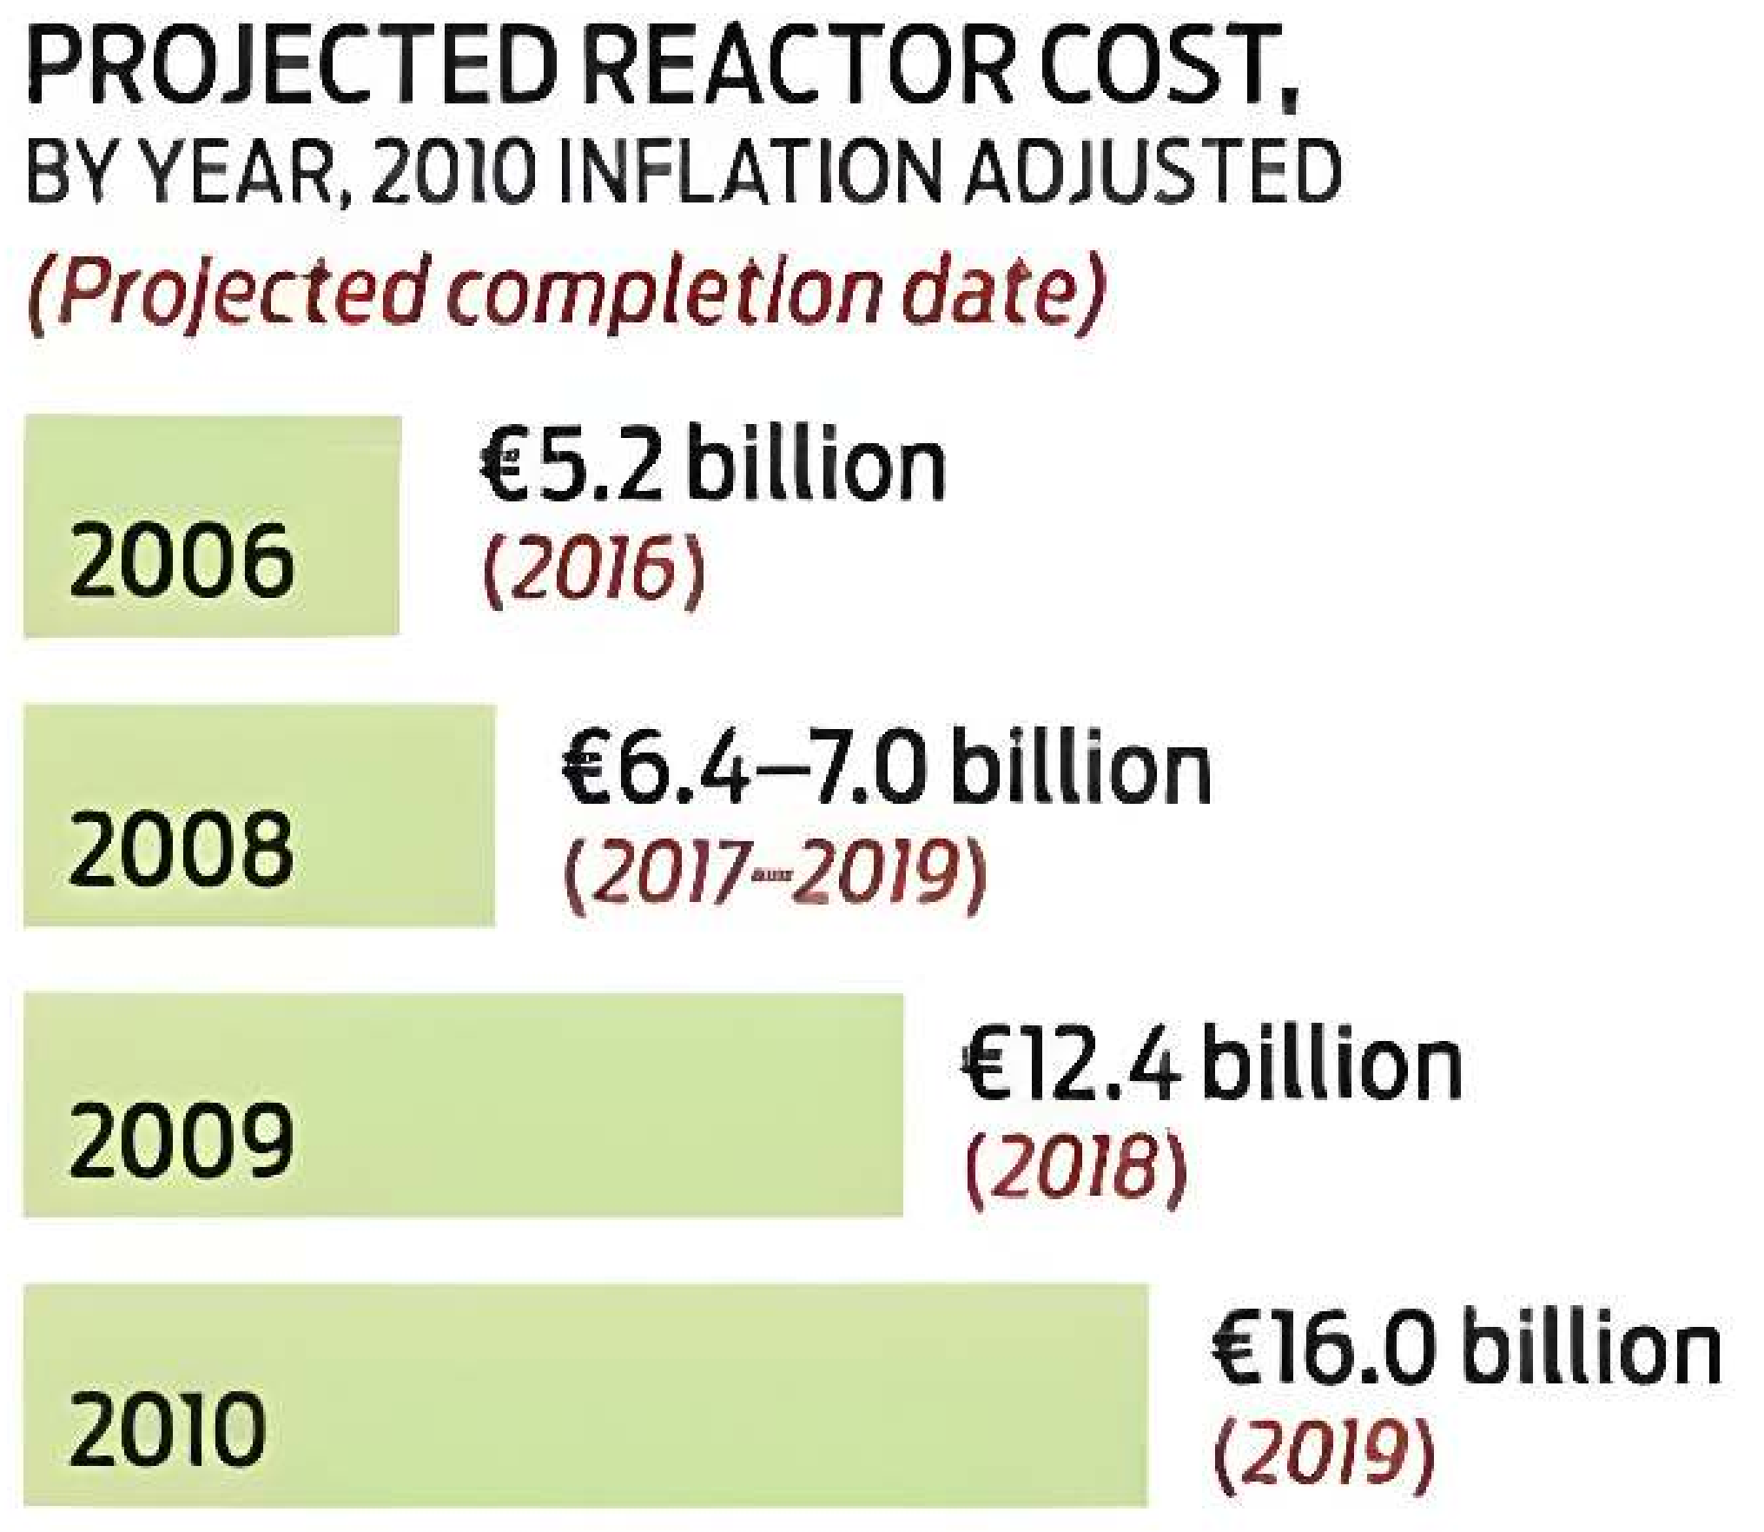

Source: Fusion Is Not Free, Who's paying for ITER's big appetite? Ariel Bleicher, IEEE (2010)

Electricity Cost per Source

*LCOE is the average total cost of producing electricity over the lifetime of a power source, including initial capital, fuel, operation, maintenance, and financing costs.

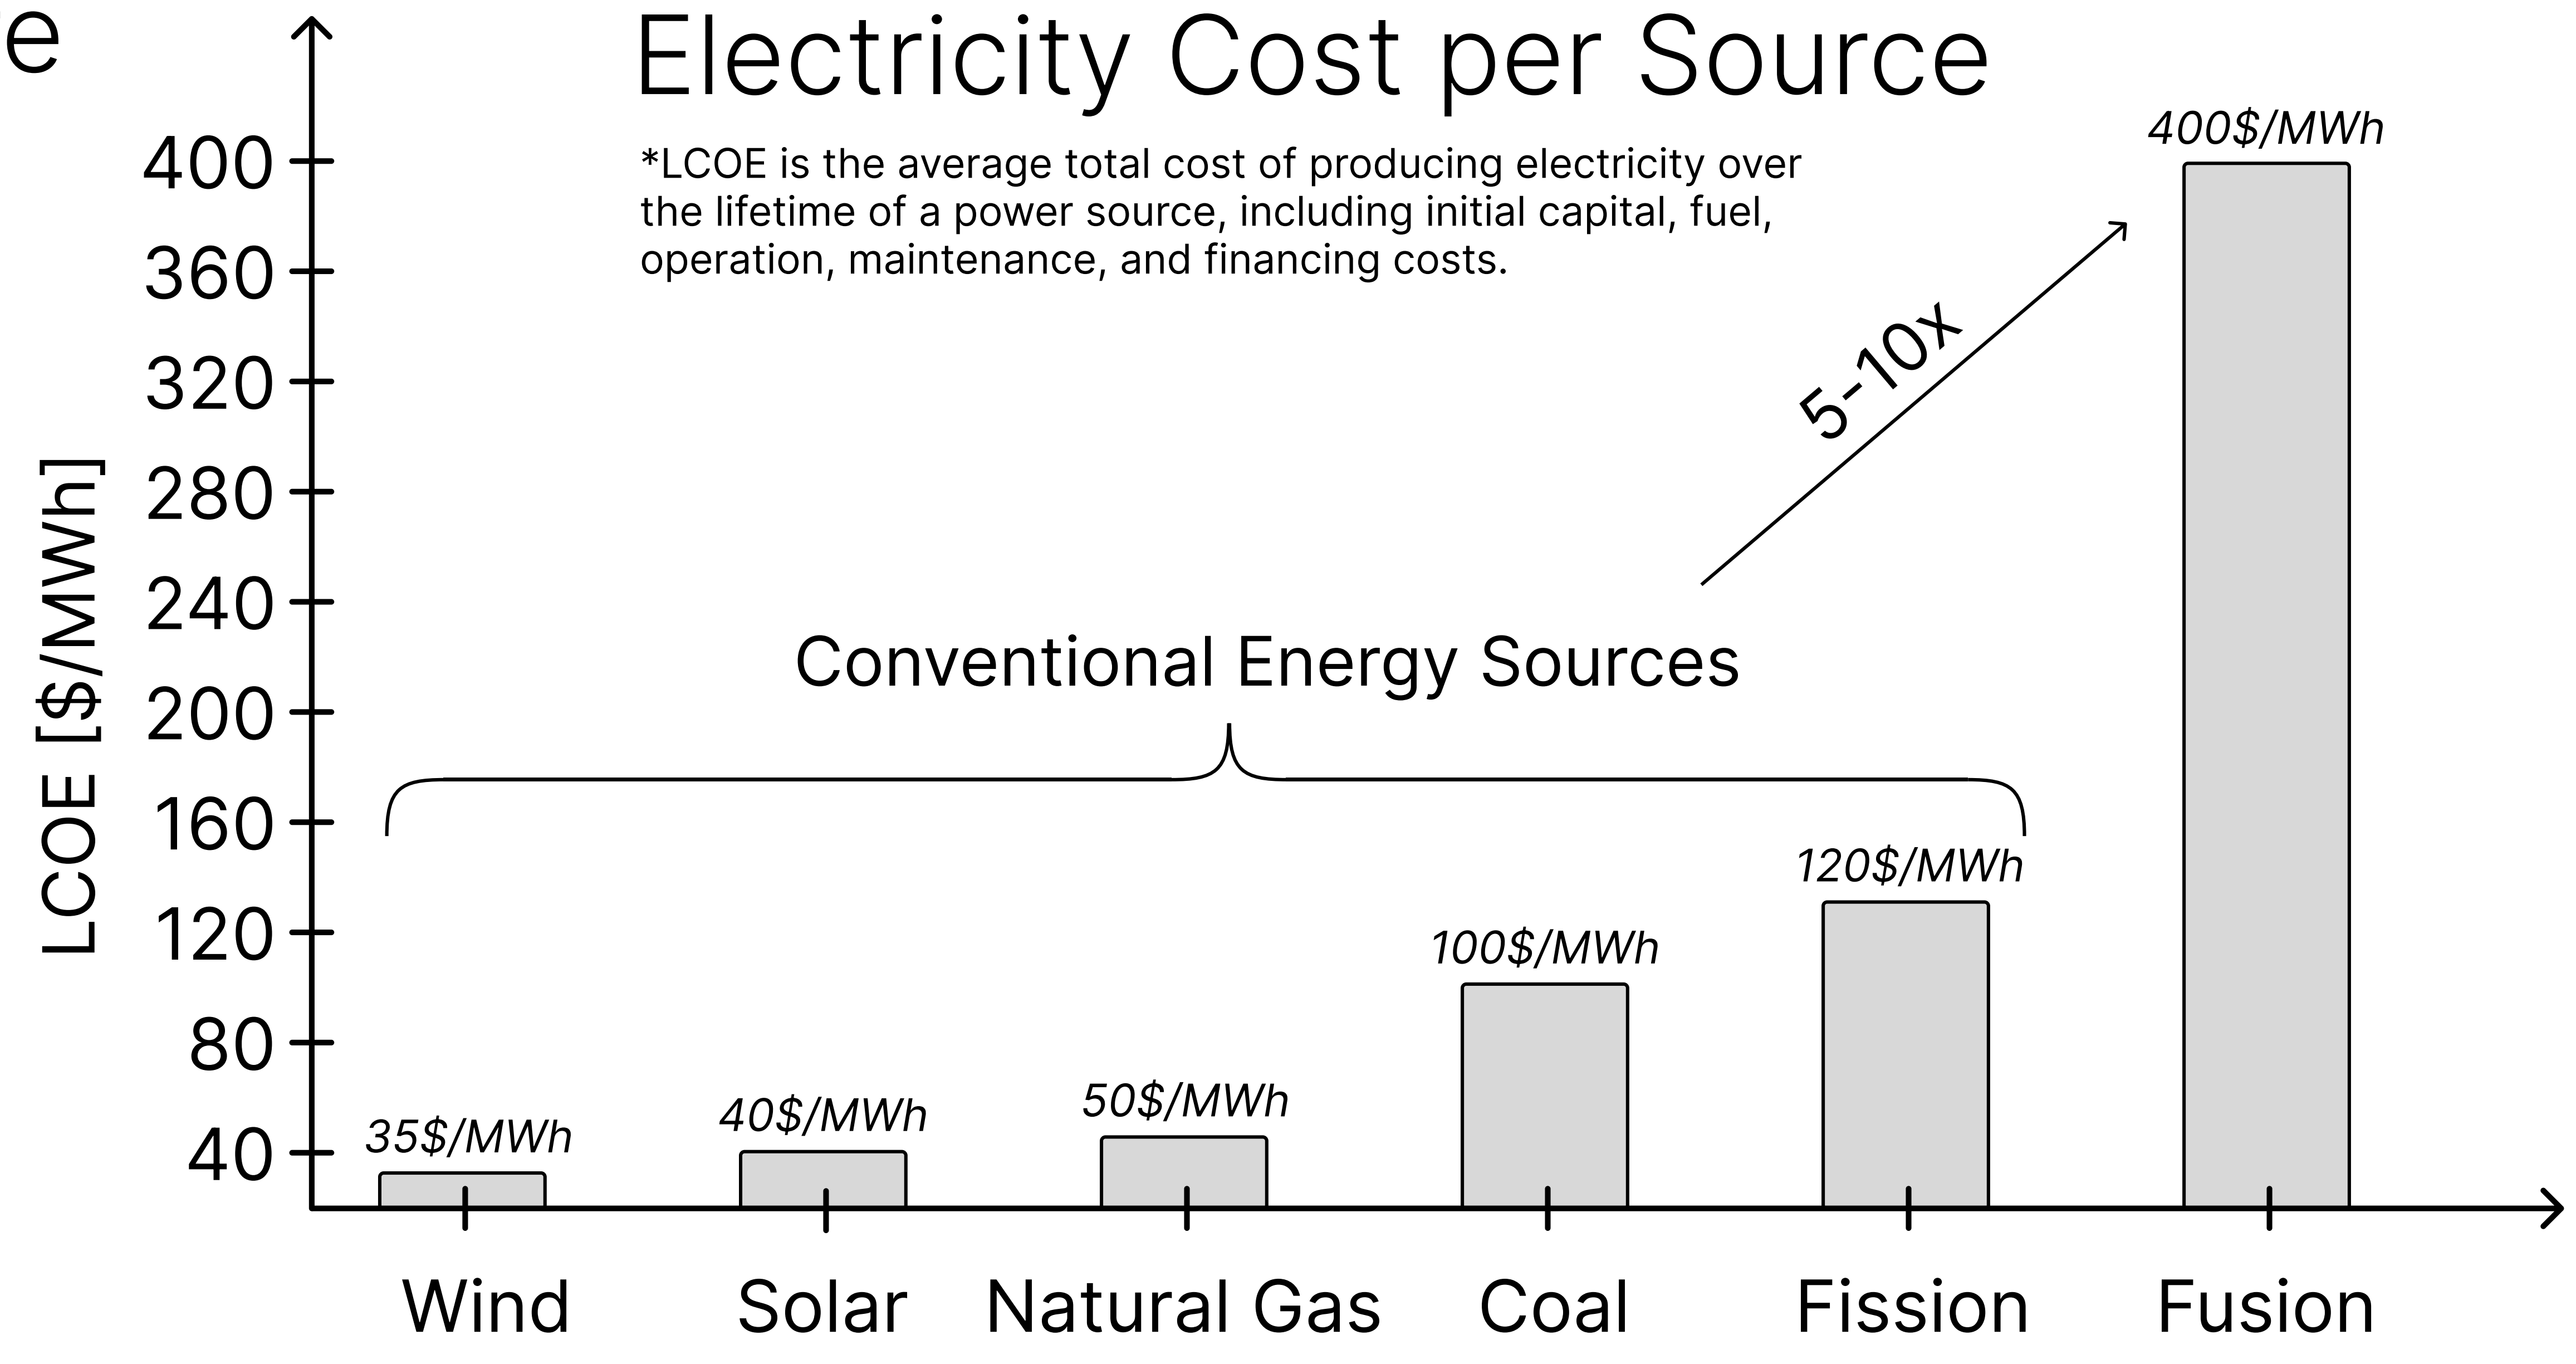

Source: Assessing the Economic Viability of an ARC-Like Fusion Plant: Cost Estimation & LCOE – Dusmet Farina, Festa (2023)
https://world-nuclear.org/information-library/economic-aspects/economics-of-nuclear-power
https://www.eia.gov/outlooks/aeo/pdf/electricity_generation.pdf

Availability Is Crucial

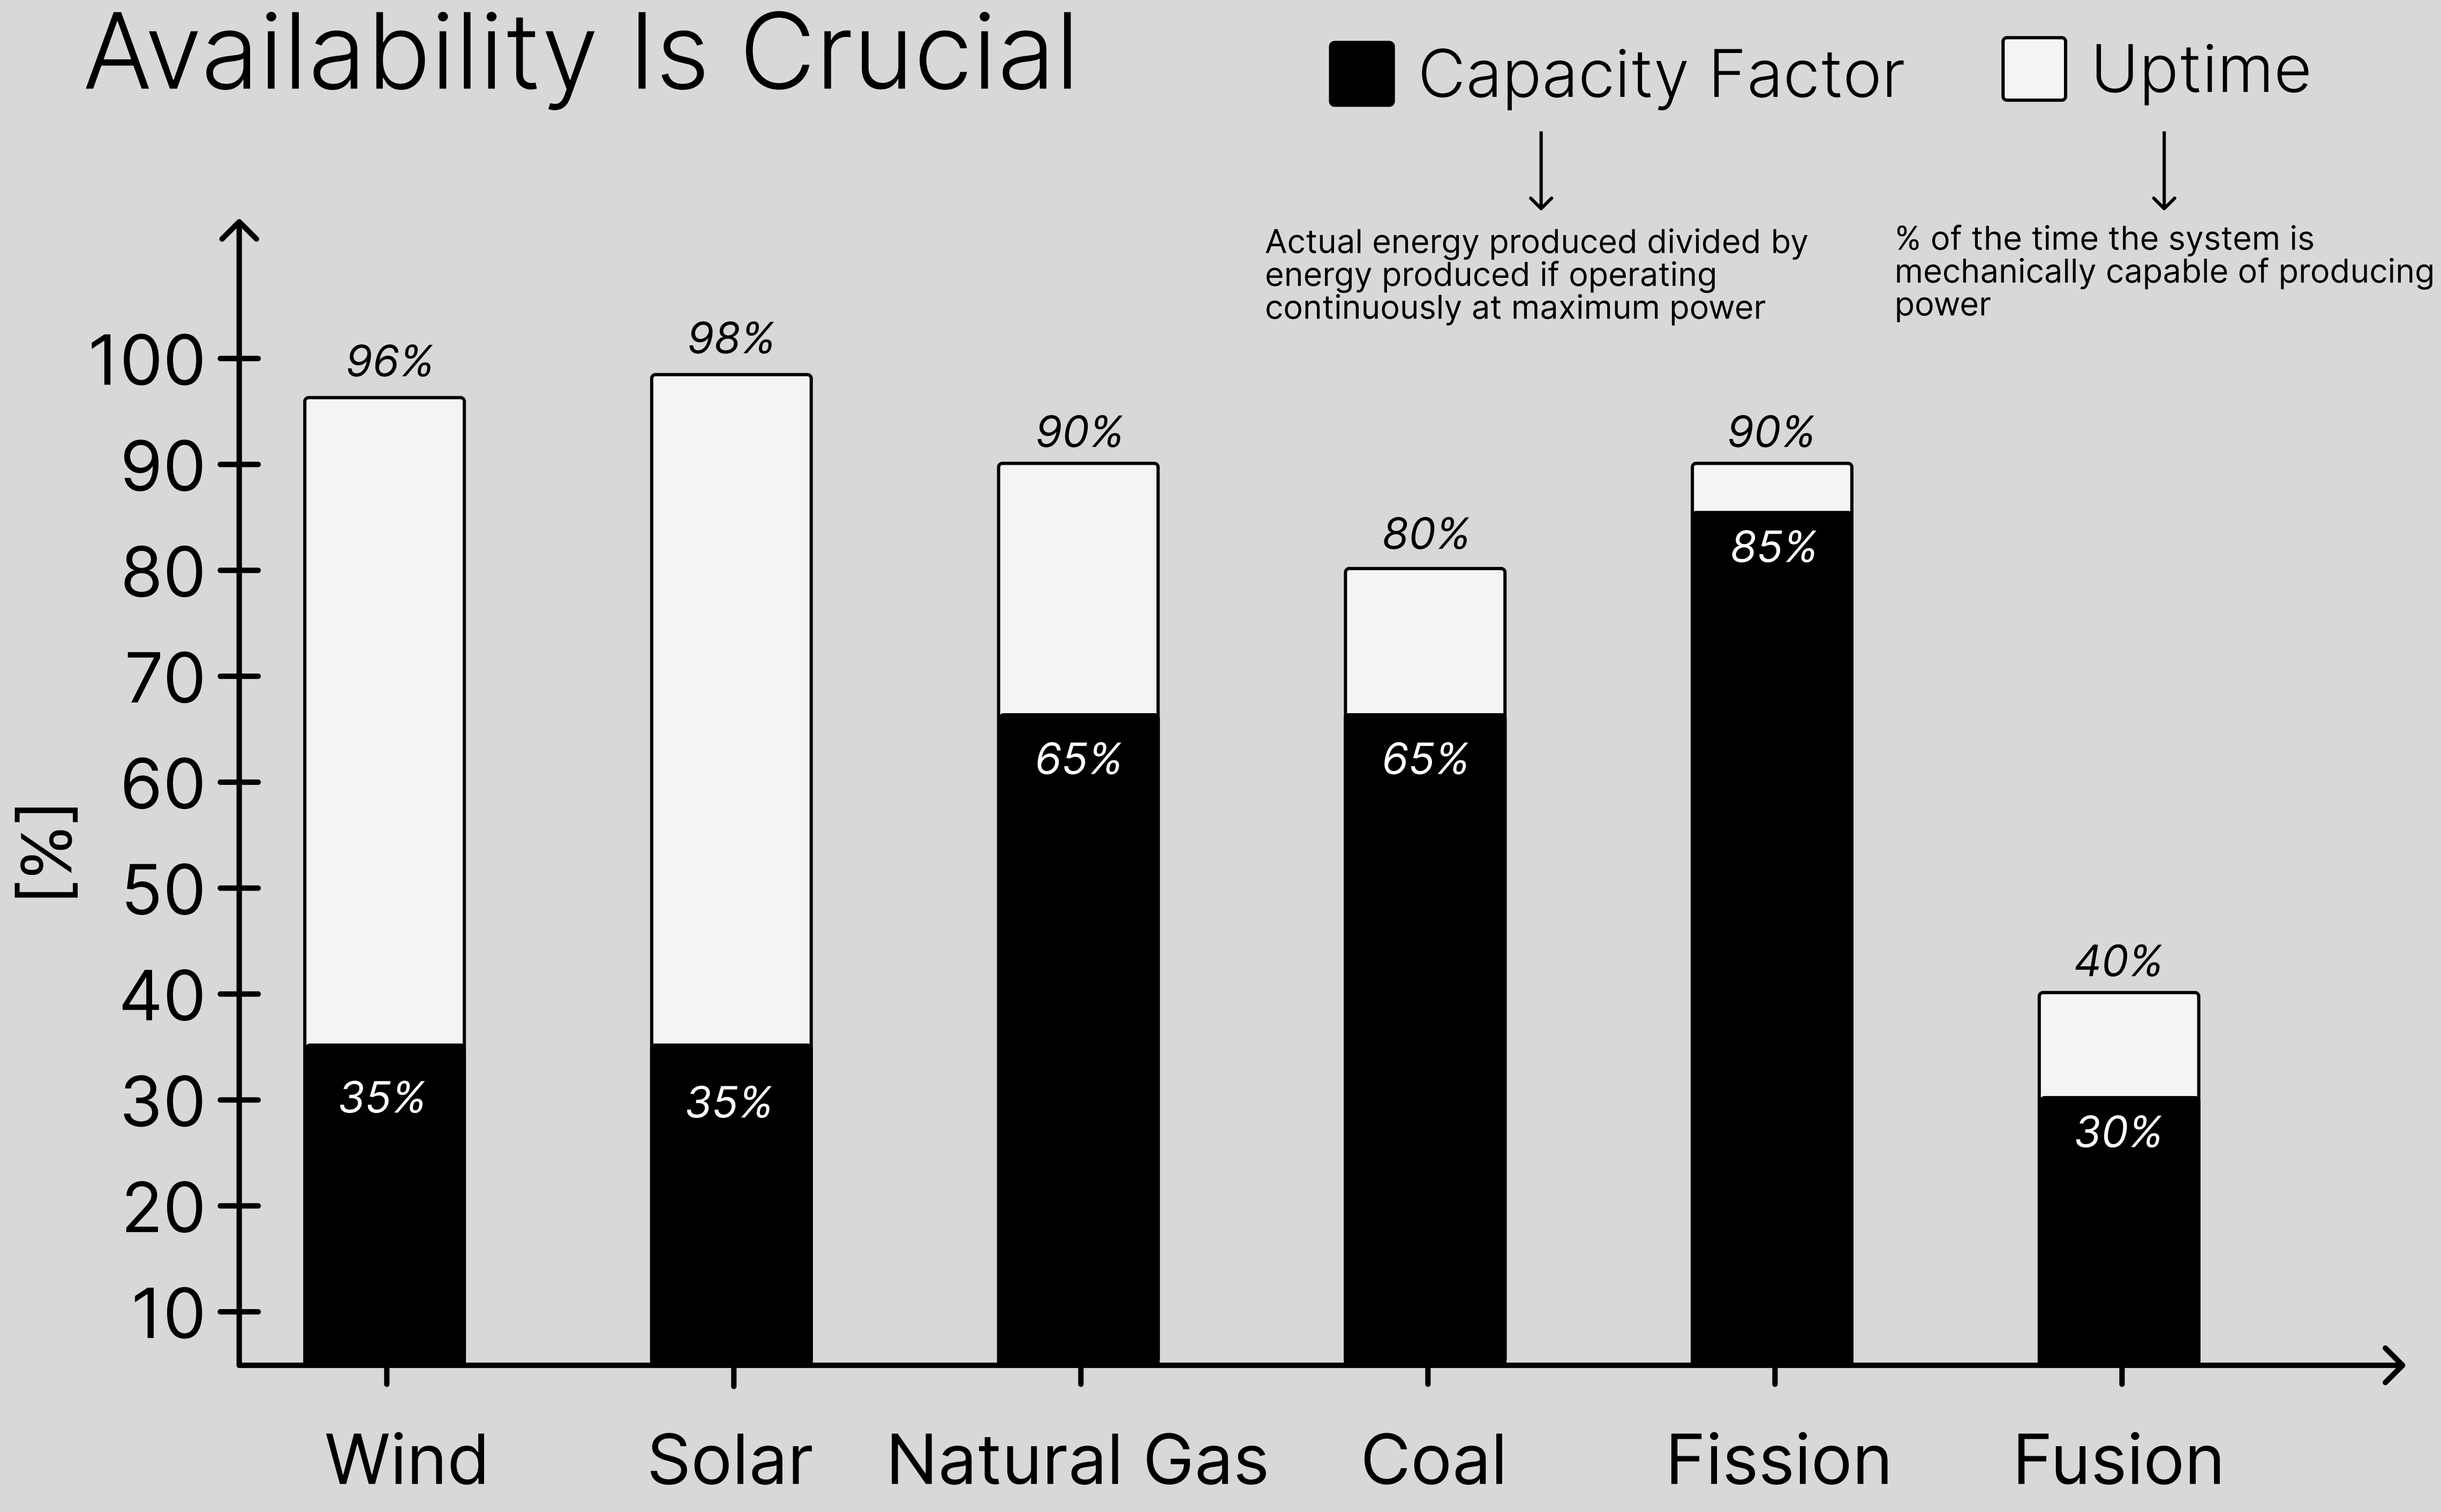

Source: https://world-nuclear.org/information-library/country-profiles/countries-t-z/usa-nuclear-power
https://www.projectfinance.law/publications/2004/december/risk-allocation-in-wind-projects
https://www.climate-and-hope.net/electricity-technologies/maintenance-breakdowns-and-availability
https://www.nrel.gov/docs/fy24osti/88769.pdf

Also, the FOAK reactors will attempt to compete in an energy market where high uptime is standard—80–95% for nuclear and fossil, and 95%+ equipment availability for renewables. Early fusion plants, however, may only operate at maximum power 33% to 50% of the time due to the **challenges of high plasma power density**. Planned downtime will be needed for replacement of in-vessel components. Overcoming disruptions and extending plasma burns will also be a hurdle in achieving baseload reliability.

So, if your fusion reactor doesn't produce electricity at 5¢/kWh—the current market rate—and operate reliably more than 90% of the time, why would someone buy it (on the grid)?

In other words, trying to use any of the first 100 fusion power plants to compete on the grid is equivalent to using a Wright brothers aircraft to operate an airline.

Continuous grid operation requires new materials

First wall material degradation

Tungsten is one of the leading contenders for first wall material due to its high melting point of 3,422°C—the highest among all metals—and its low thermal expansion coefficient.

However, even Tungsten under high heat load will sputter into the plasma and cause it to cool, hindering the fusion power output by increasing impurity concentration in the plasma.

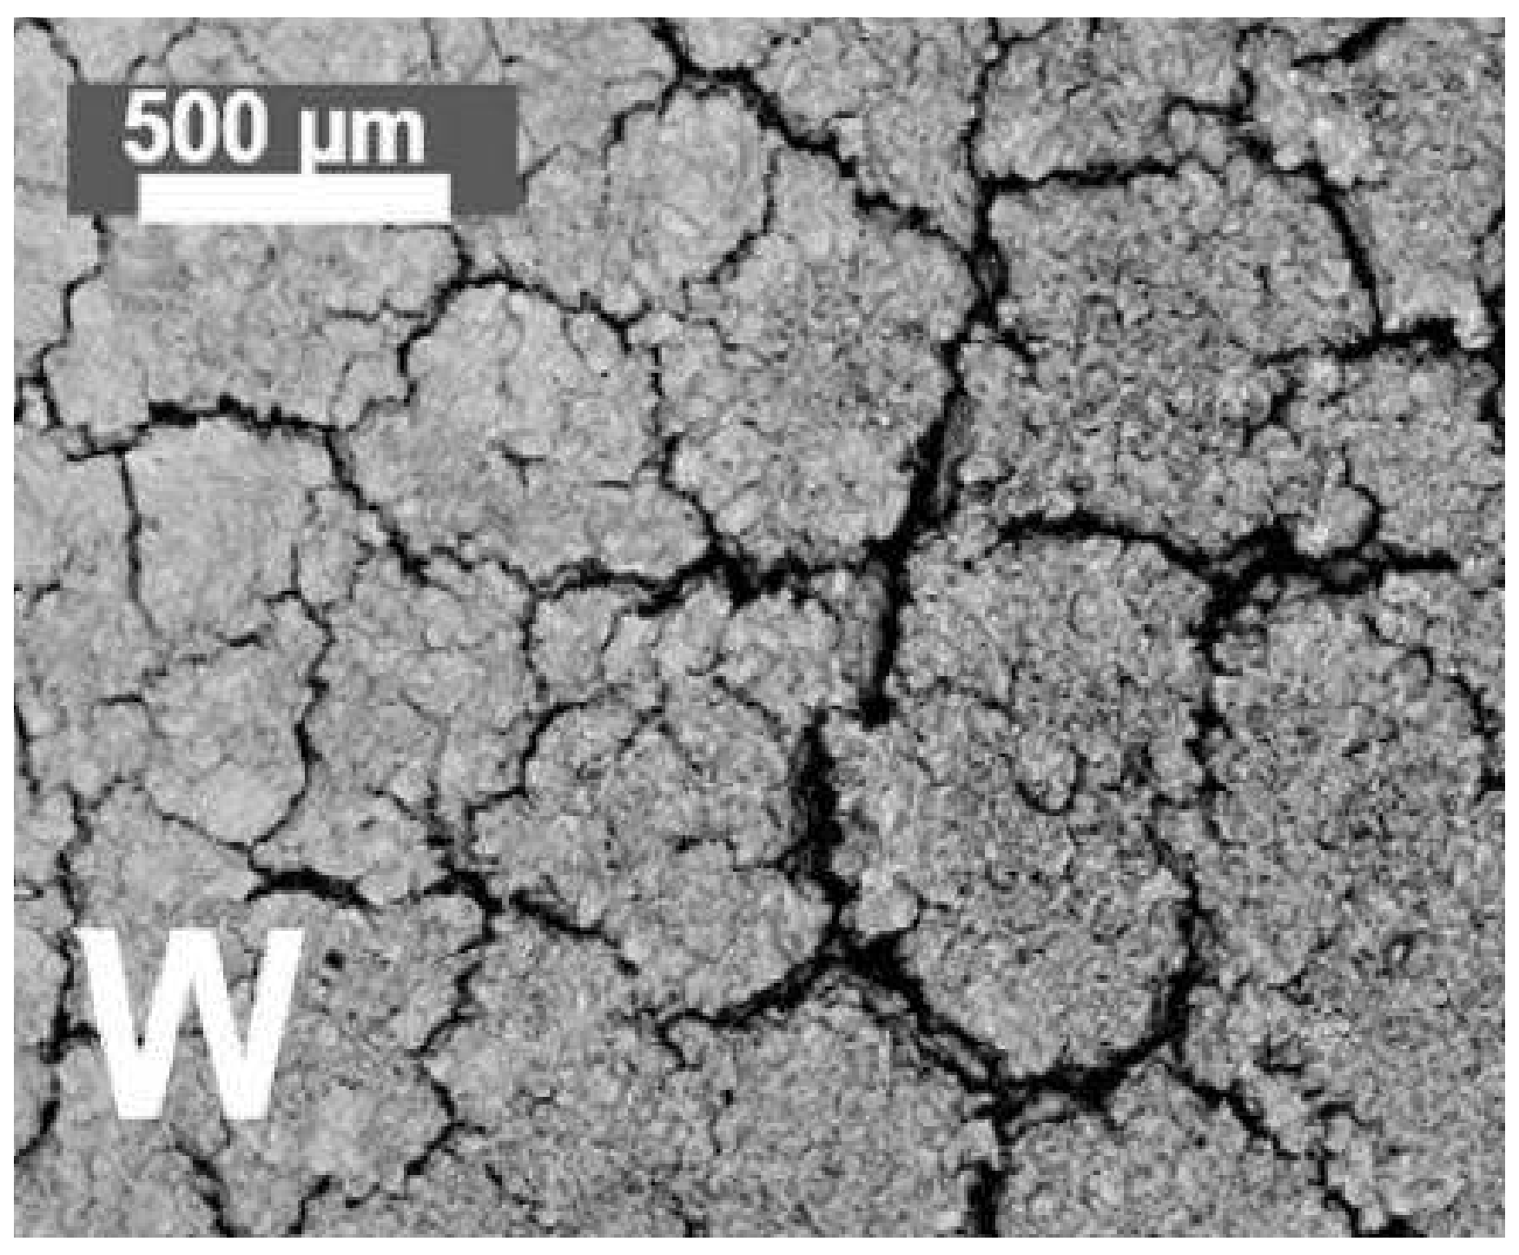

Additionally, the extreme heat loads can lead to microcracking, surface erosion, and embrittlement over repeated cycles. Furthermore, helium ions (alpha particles) from the fusion plasma can accumulate in tungsten and create microscopic bubbles beneath the surface—often leading to “fuzz” formation—which alters the material’s morphology, reduces its thermal conductivity, and accelerates overall degradation.

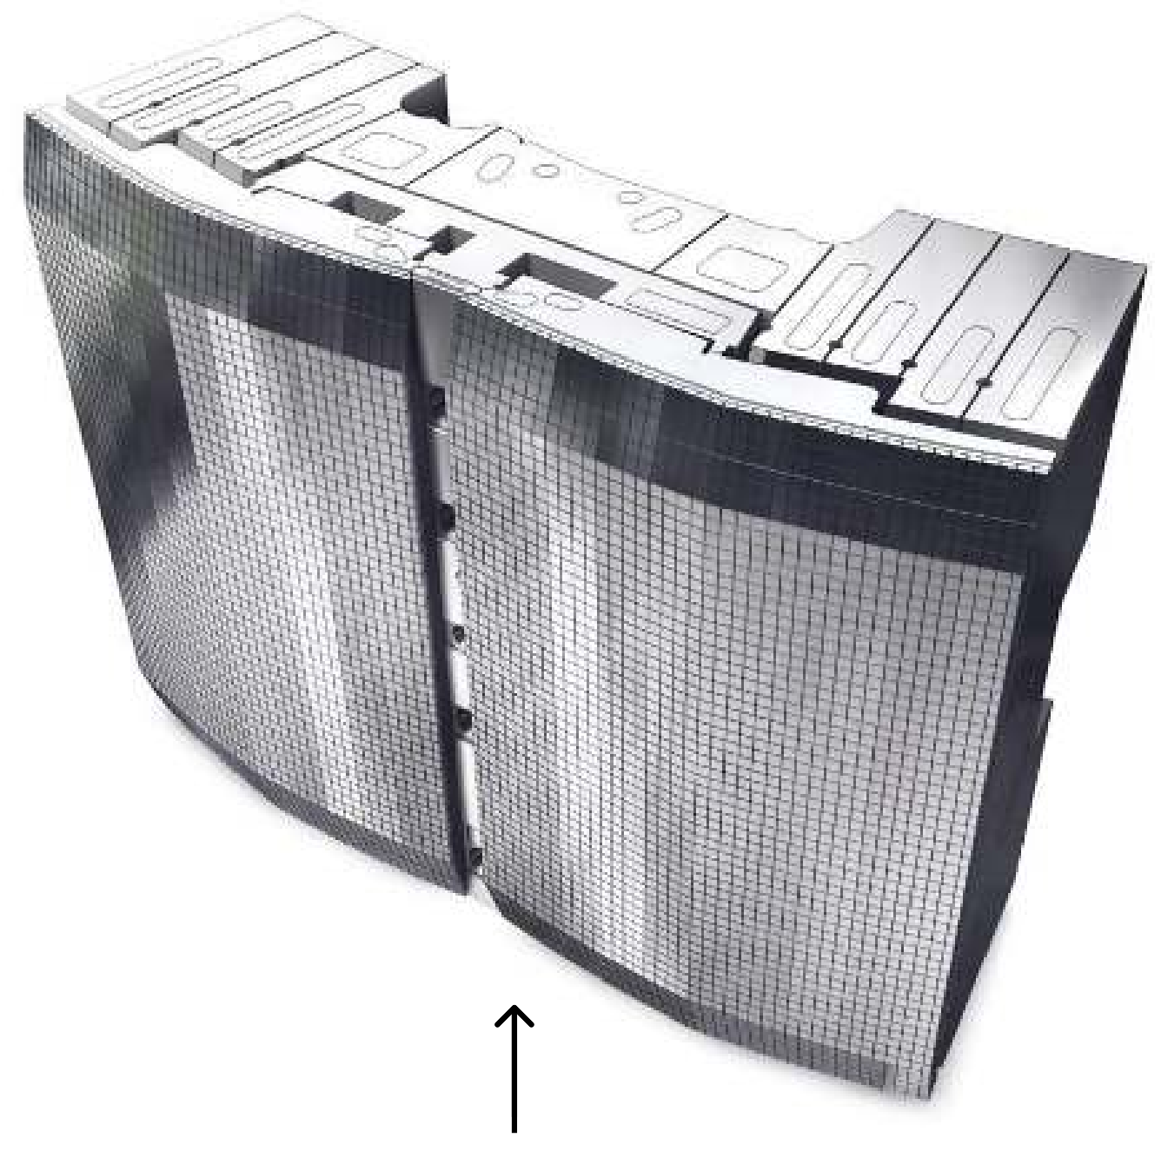

↑
ITER first wall & blanket
↓

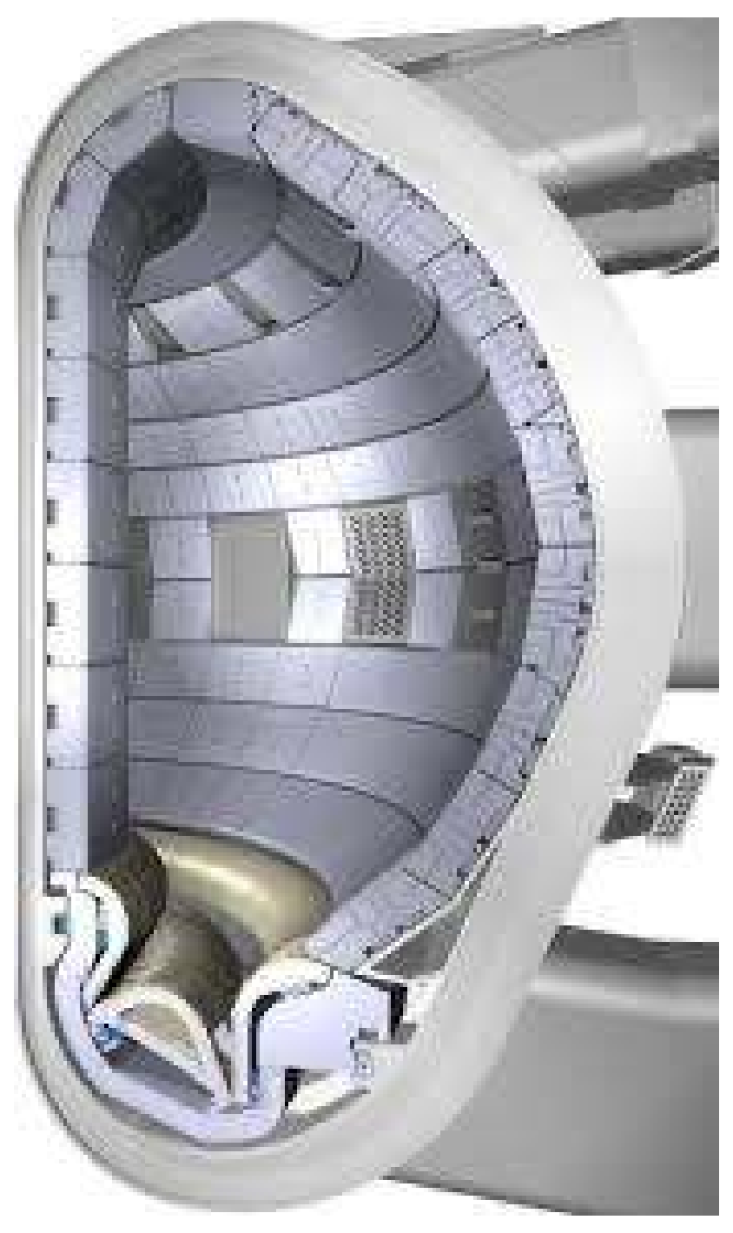

State-of-the-art plasma-facing materials are being validated for ITER relevant operating conditions and lifetimes (**0.3MW*yr/m^2**), but no materials exist that have been shown to last under these power loadings for decades on end.

Source: ITER Research Plan within the Staged Approach (Level III – Final Version) - 2024

	ITER	JT60-SA	JET	SPARC
Auxillary Heating Power	120MW	6.7MW	40MW	25MW
Alpha Heating Power	100MW	41MW	3.2MW	27.8
First Wall Surface Area	600m^2	350m^2	340m^2	100m^2
Radiative Power Loading	0.37MW/m^2	0.14MW/m^2	0.13MW/m^2	0.53MW/m^2
Neutron Power Loading	0.67MW/m^2	0.08MW/m^2	0.04MW/m^2	1.12MW/m^2

Nuclear activation creates radioactive waste

In fusion reactors, deuterium and tritium fuse to release high-energy neutrons (14MeV). These neutrons embed themselves in surrounding reactor components, such as the vacuum vessel, and react with trace impurities, causing activation. This means that although fusion is still no where near fission in-terms of radiological hazard, radioactive waste is still generated that must be handled and stored, adding additional cost to the fusion plant lifecycle operations.

The severity of this problem is directly proportional to the reactor's power density and its operating duration. Thus for proposed high-power-density grid reactors like ARC, the radiation dose rate following shutdown of the reactor can remain **dangerously radioactive for >100 years due to the high power density they operate at.** (see figure to the right)

The goal is to minimize the duration that the structures emit hazardous radiation levels to make maintenance and nuclear waste storage simpler. This can only be achieved by reducing the level of neutron bombardment on the structures or inventing lower activation materials.

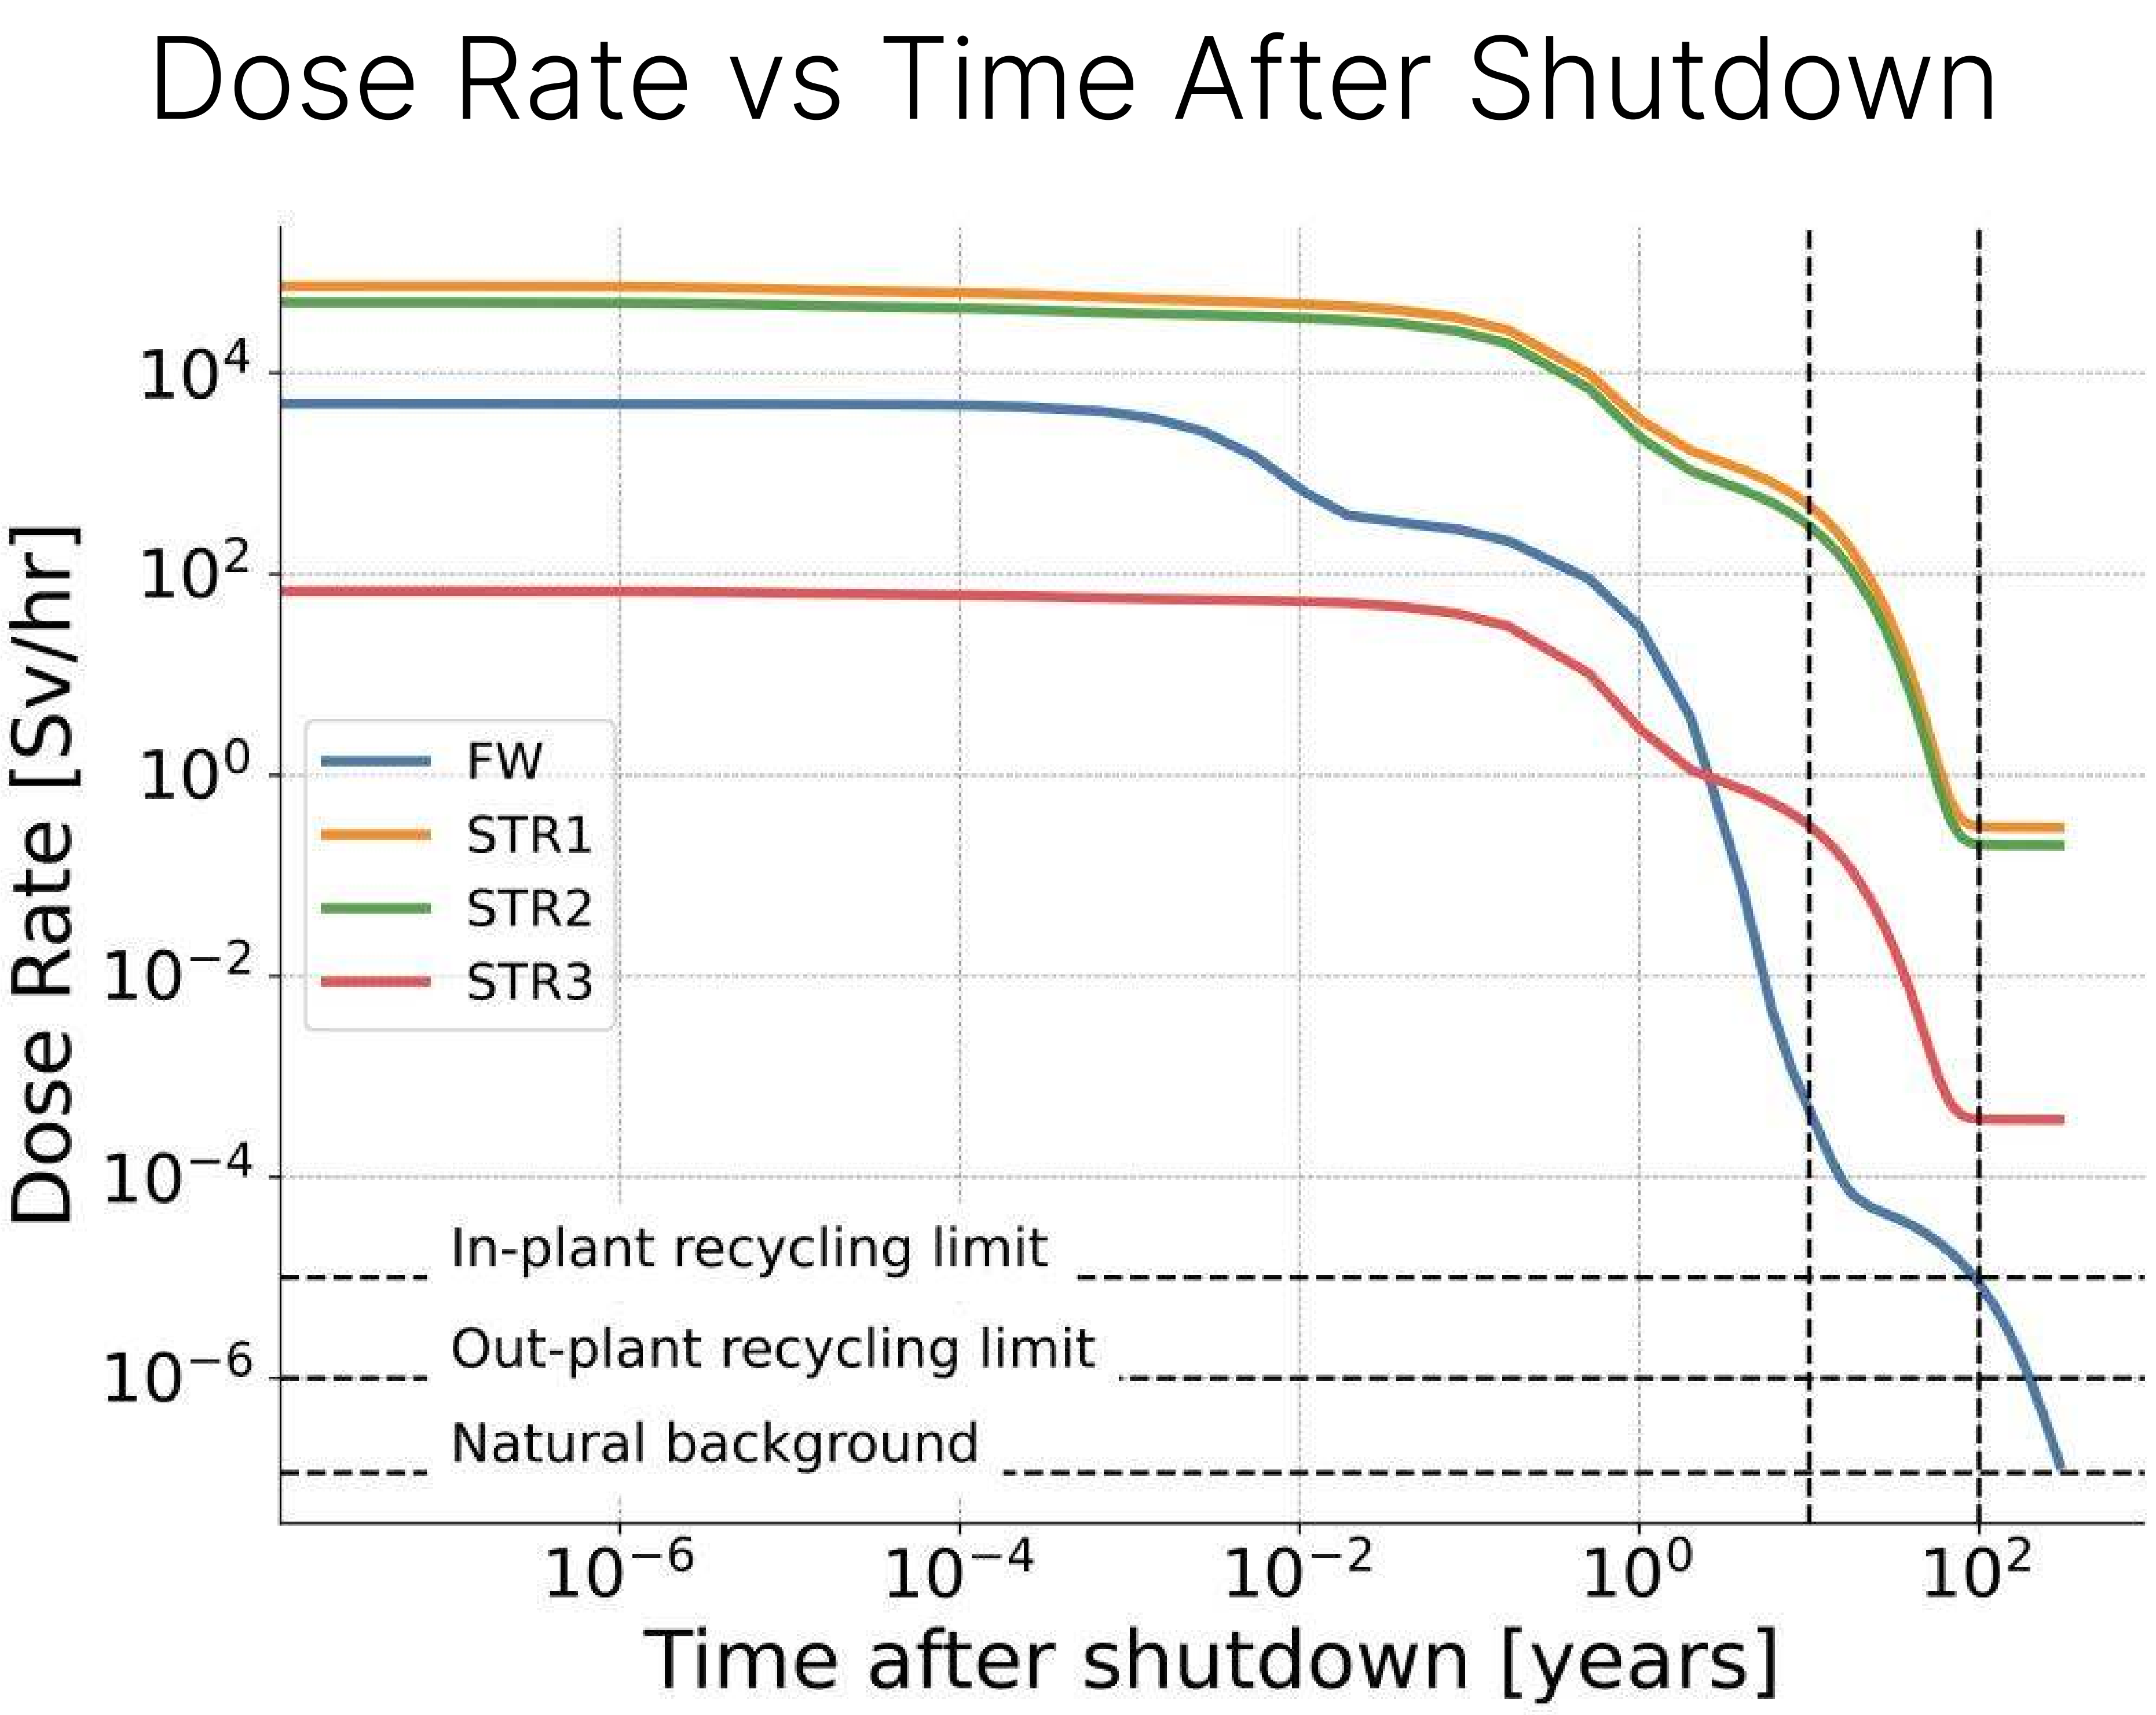

Source: Pettinari, D., Testoni, R., Zucchetti, M., & Parisi, M. (2024). Neutron transport and activation comparison between OpenMC and FISPACT-II in ARC-class reactor. Fusion Engineering and Design, 209, 114713.

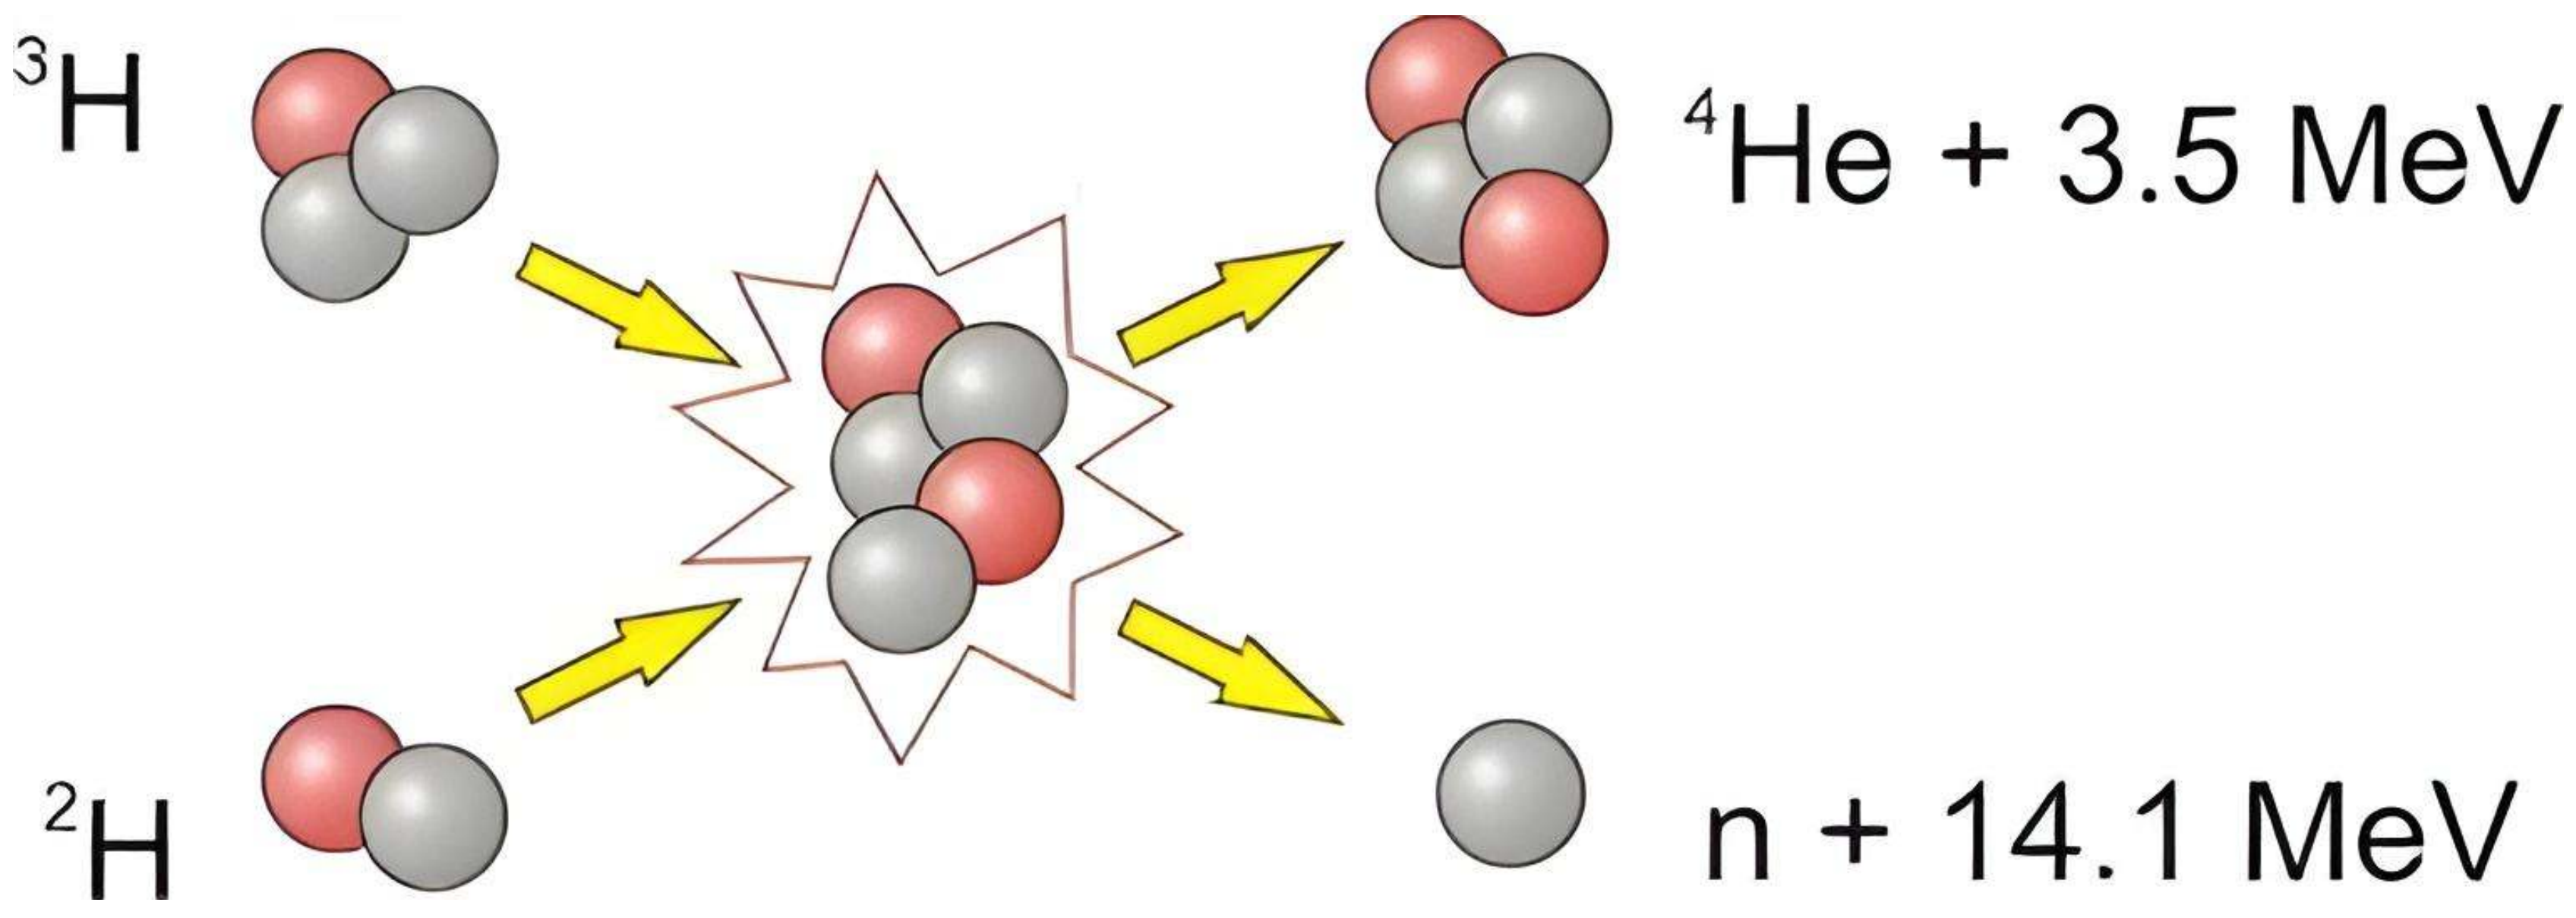

Takeaway

High power density devices make both the **first wall heat flux** and **nuclear activation** problems **exponentially harder**, and we simply do not have the materials to allow for continuous operations in these environments.

Chapter 3

Solution

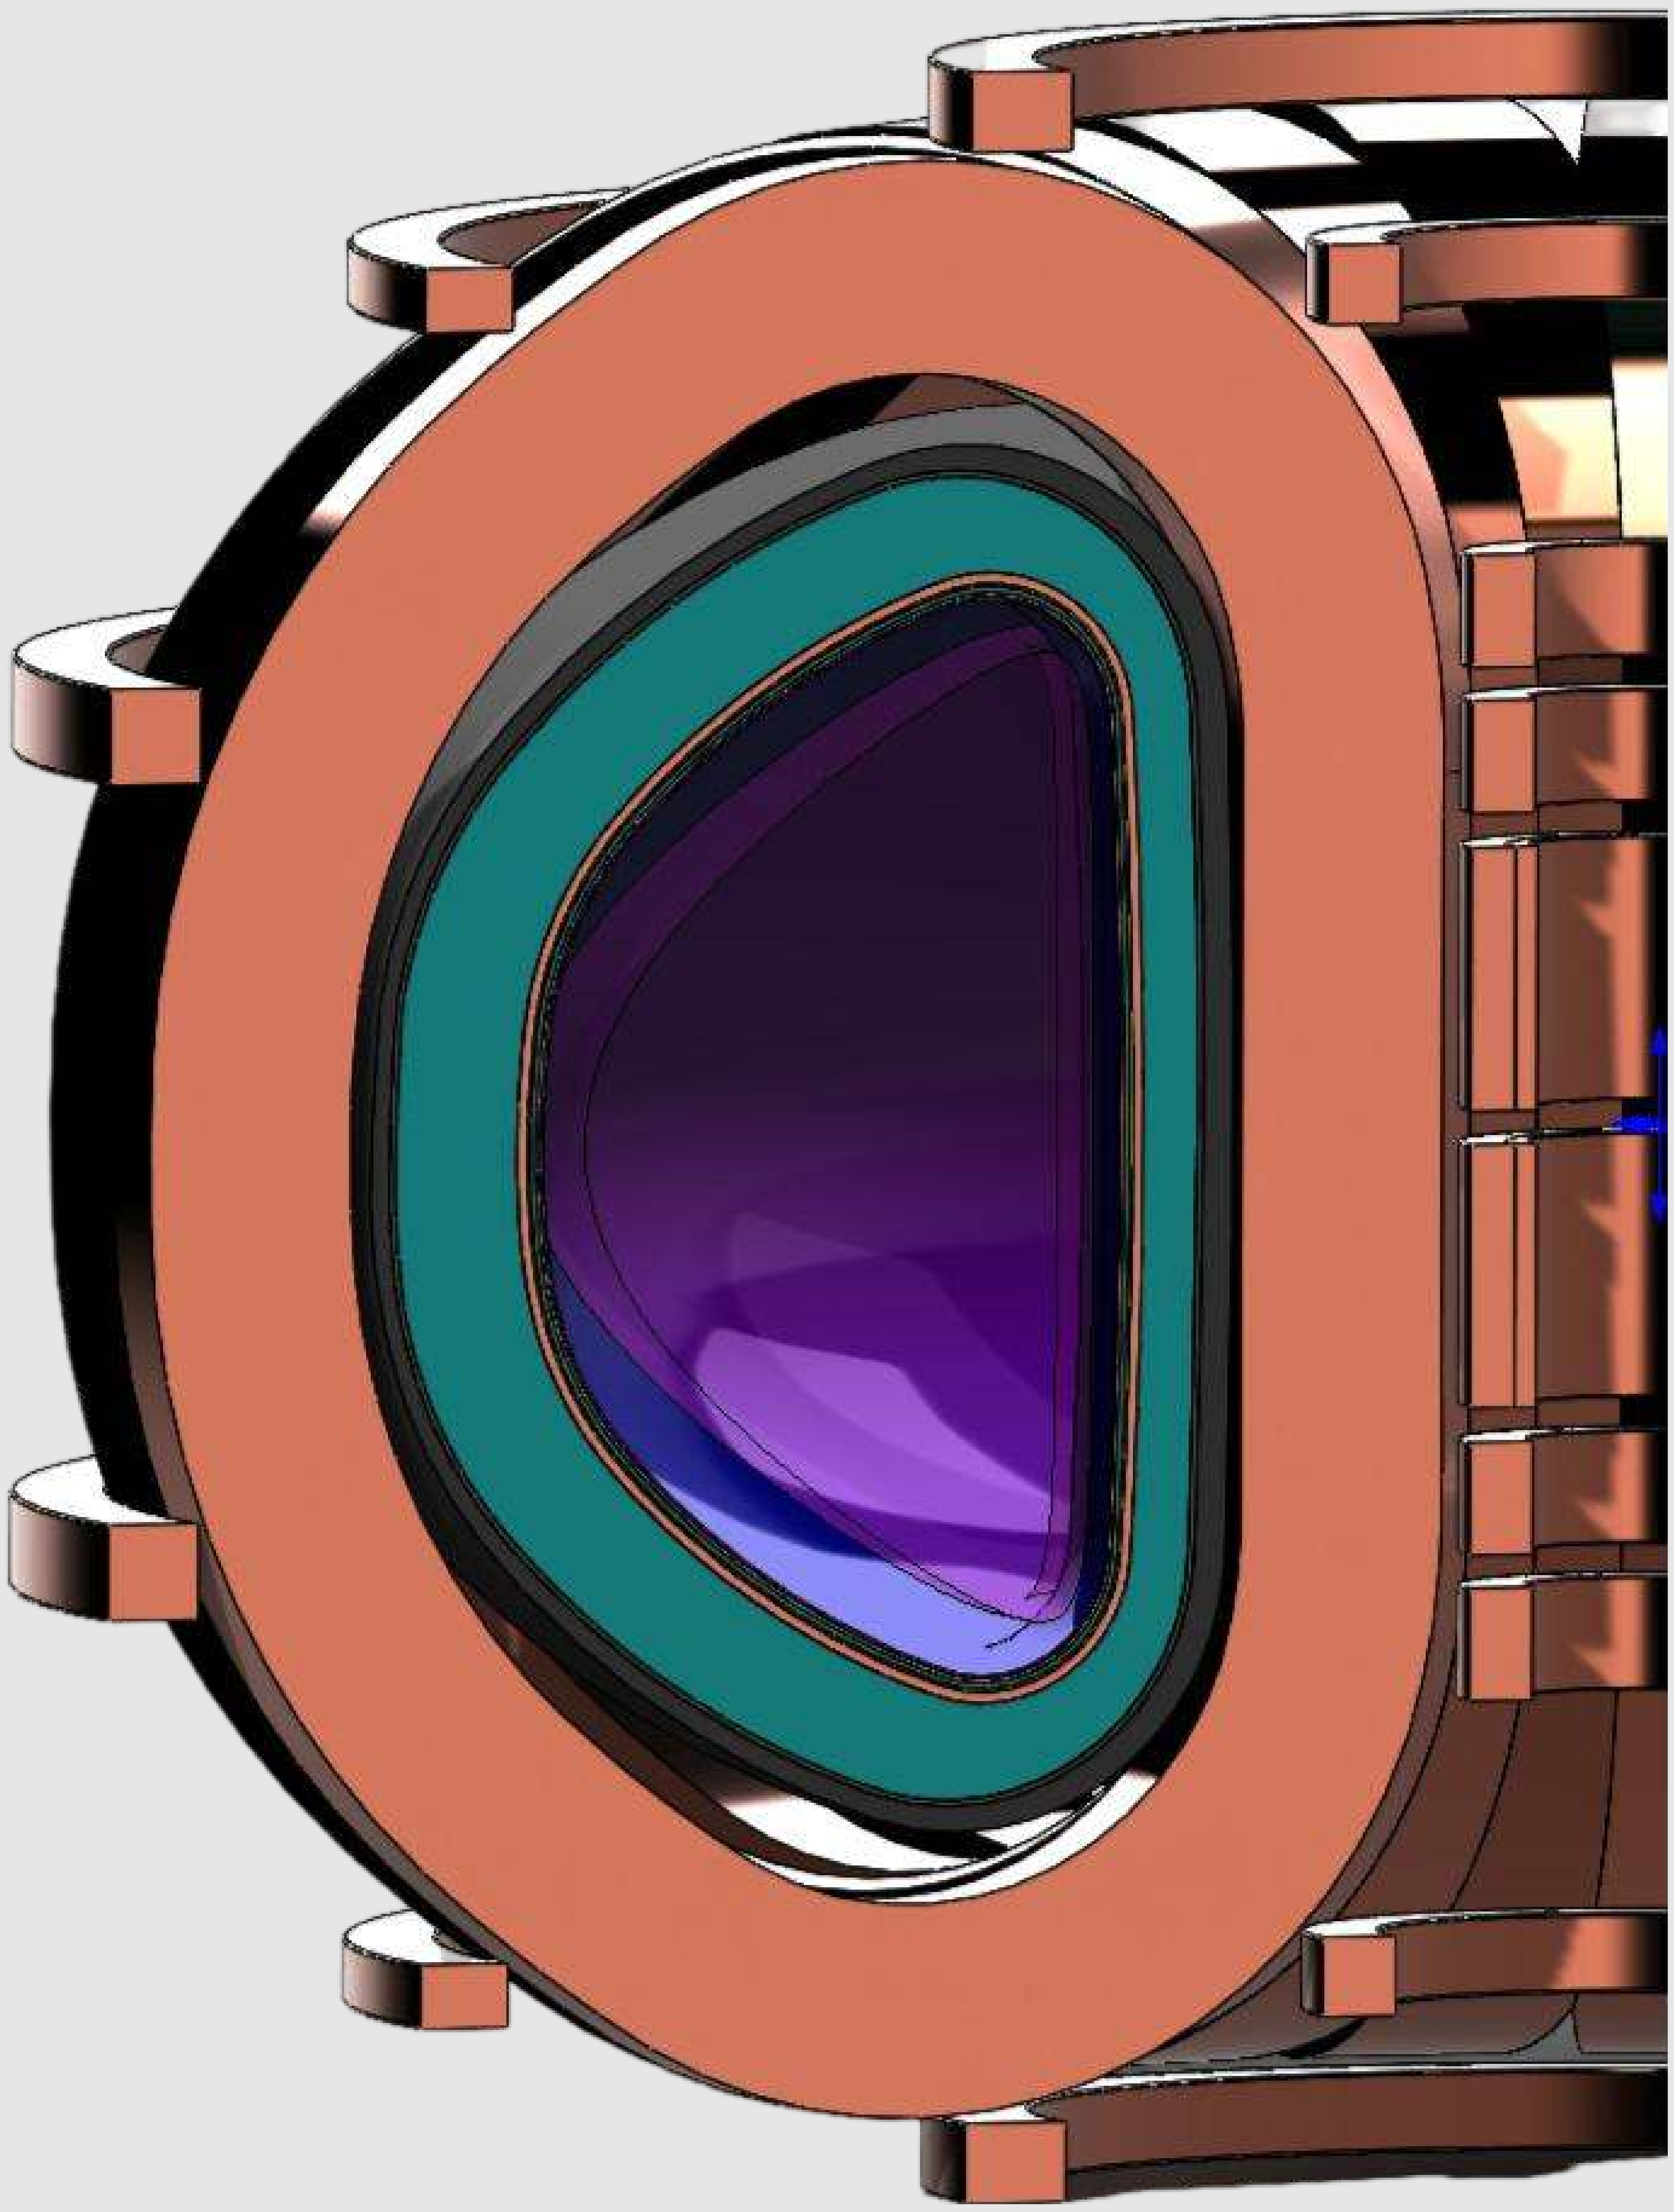

Meet Yinsen

The first of a kind fusion reactor designed for off-grid power

Named after Yinsen, Tony Stark's fellow captive assistant from ***Iron Man*** who helped him assemble the first ***arc reactor***. He sacrificed himself helping Tony escape from their captors. A true unsung hero.

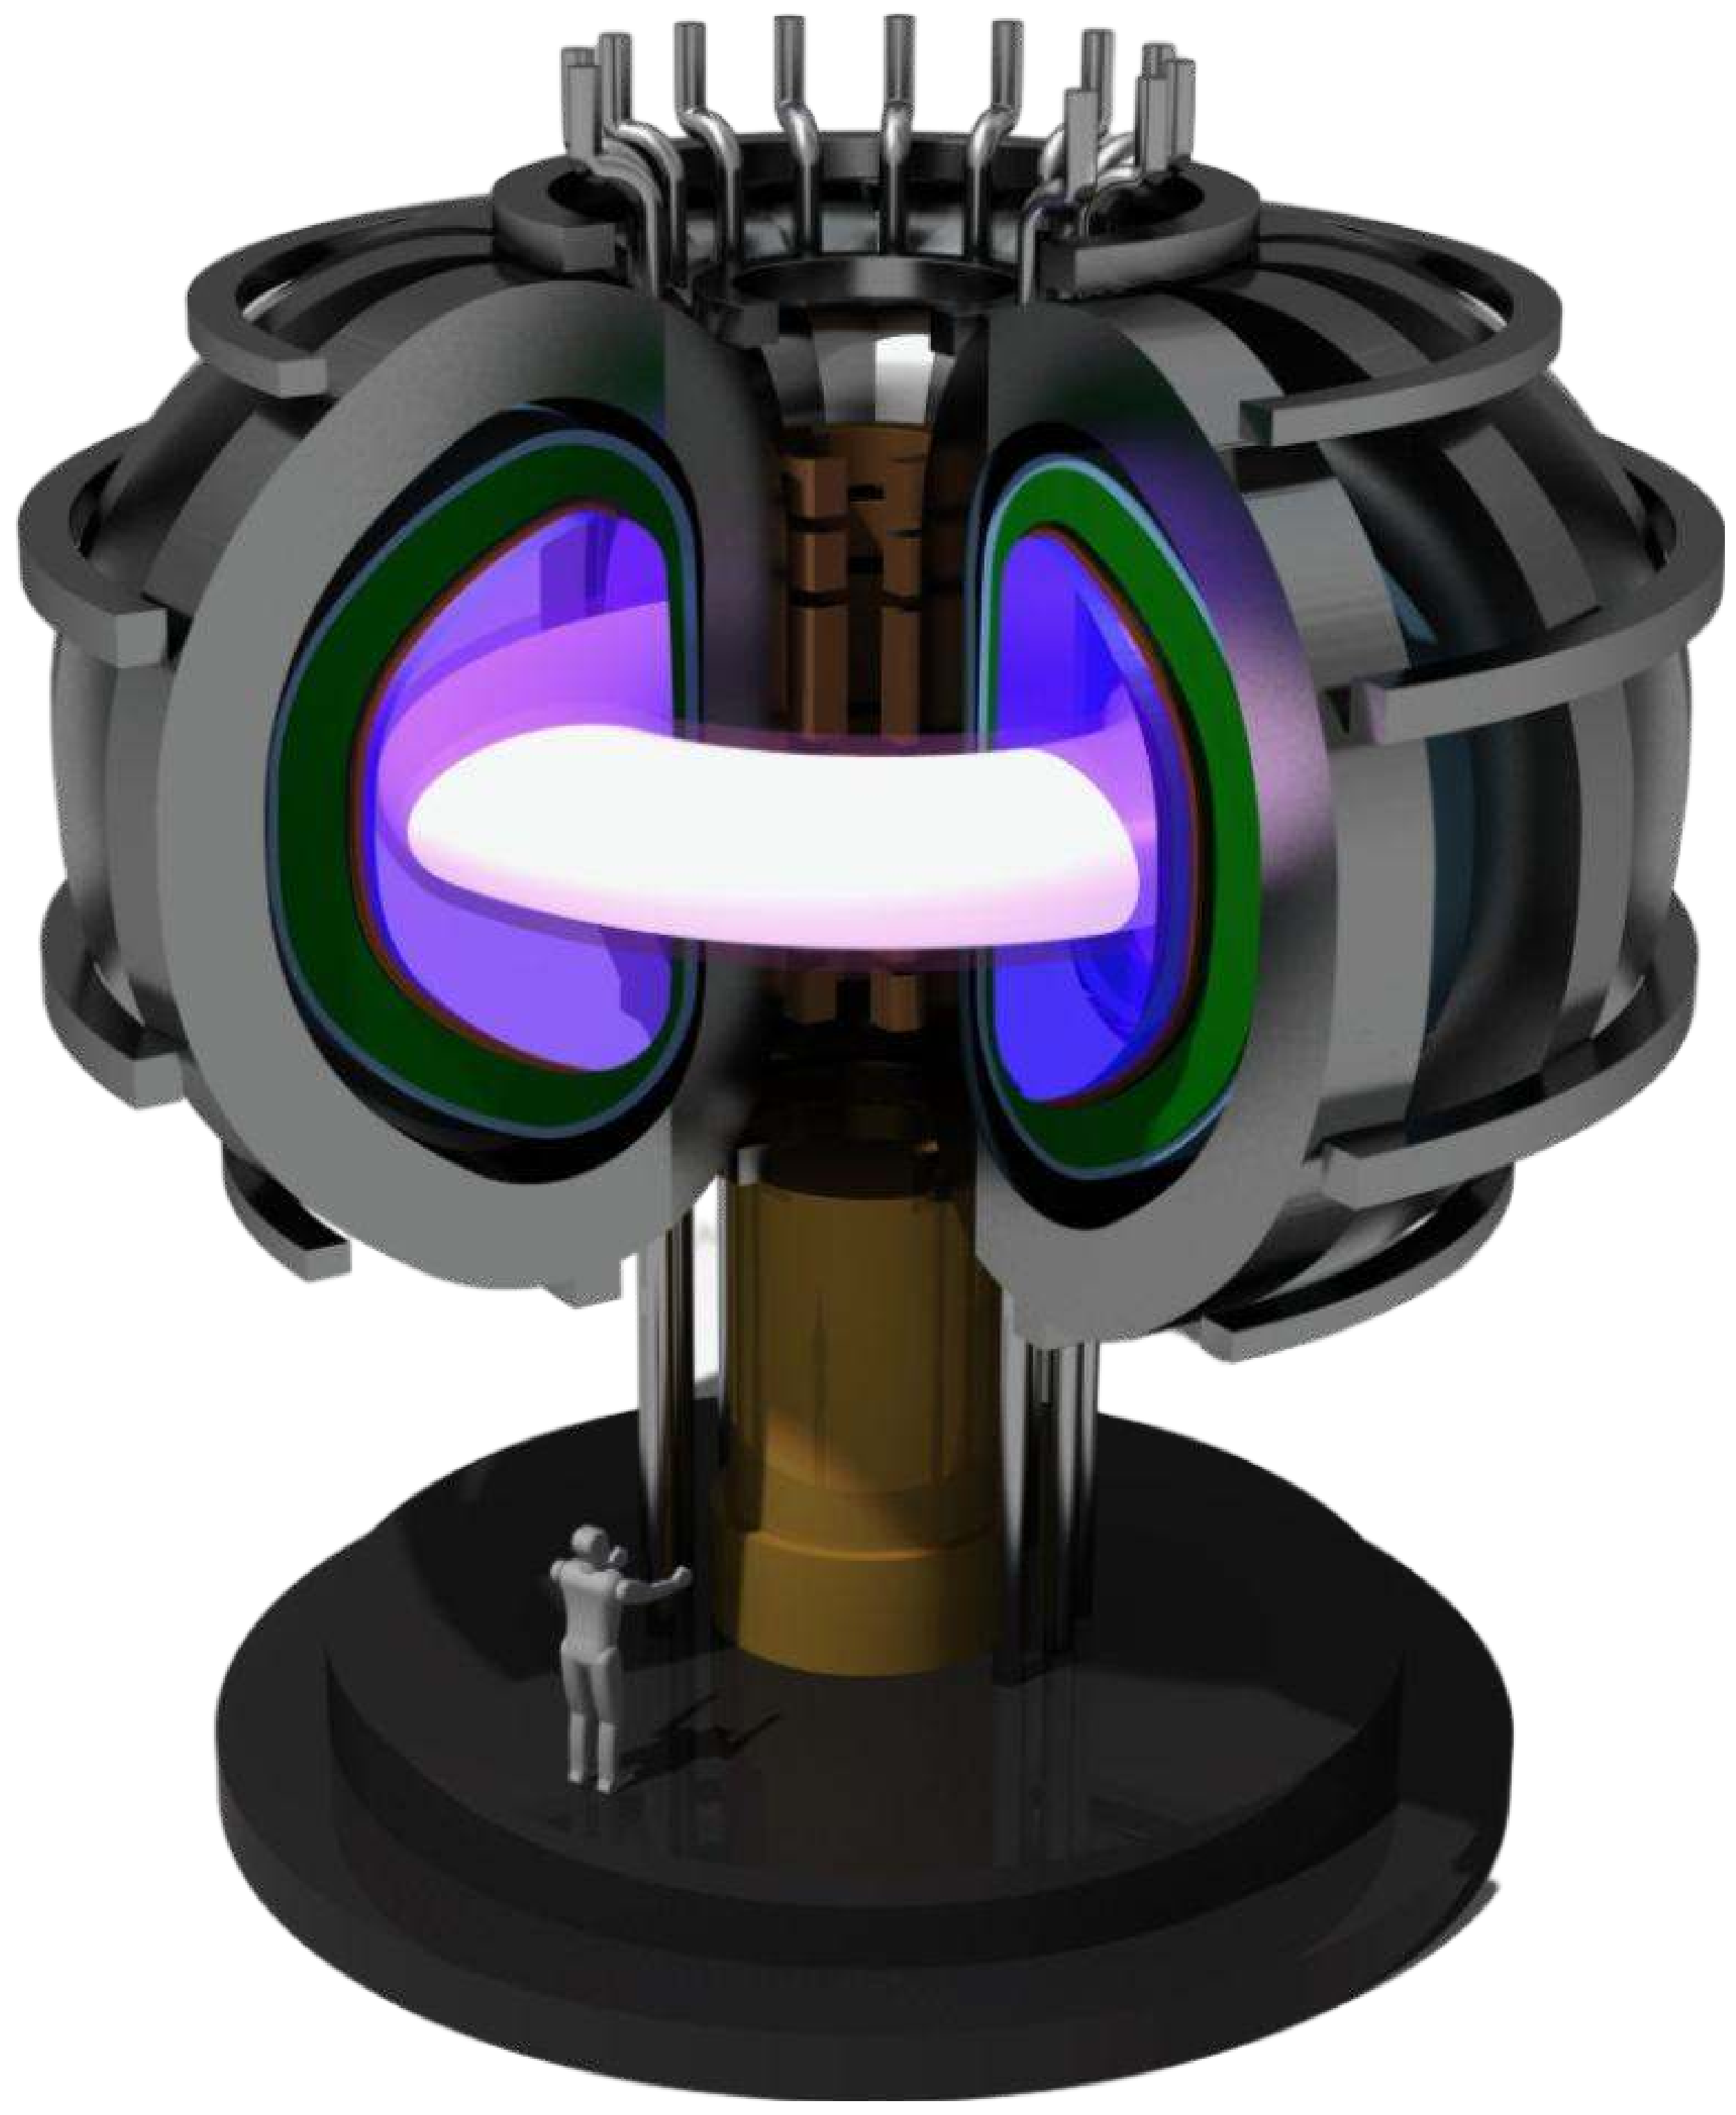

Design Parameter	Symbol	Value
Fusion Power	P_f	115MW
Net Power	P_p	25MW
RF Heating	P_{RF}	7.5MW
Plasma Gain	Q_P	13
Major Radius	R_0	2.85m
Minor Radius	a	0.884m
Elongation	κ	1.8
Triangularity	δ	0.8
Toroidal Field	B_0	8.8T
Plasma Current	I_P	10.3MA
Average Temperature	$\langle T_e \rangle$	~8keV
Average Density	$\langle n_e \rangle$	$\sim 10\times10^{20}\ m^{-3}$
First Wall Power Loading	q_{FW}	$\sim 600kW/m^2$

Yinsen solves both the **physics** and **economic** problems for fusion

Why Yinsen?

Yinsen is designed for **off-grid** applications, prioritizing reliable, zero-emissions operations, in a portable (via ship) form factor. Unlike conventional fusion systems that target large-scale grid energy production, Yinsen focuses on **maritime, industrial, and remote energy needs.**

	Grid Fusion	Yinsen
First Wall Loading	$2.5\text{-}4MW/m^2$	$0.6MW/m^2$ <5x
Net Electric Power	400-1,000MW	25MW <16x
Plasma Volume	$150\text{-}800m^3$	$93m^3$ <2x
Energy Cost with respect to alternative energy sources, LCOE	5-10x	~1x <5x
Capacity Factor	80-90%	65%
Footprint	Large facility	Fits on a ship

Easier Tokamak Engineering

Because the target market requires 15x less power, the **engineering challenges** associated with high power density are greatly reduced as the first wall will see 5× lower power loading - minimizing material degradation, nuclear activation, and extending component lifetimes. Using ITER first wall lifetime scalings, we predict the FOAK Yinsen to require at worse case bi-annual first wall replacement, and best case 10 year first wall lifetime, while never needing to replace the vacuum vessel or HTS magnets over the **30 year lifetime**.

Economic Viability

The **business model** is fundamentally different from grid-scale fusion. Instead of competing with the lowest-cost electricity sources like solar and natural gas, Yinsen is designed for off-grid, maritime applications where renewable energy sources are impractical. By targeting these applications, we avoid the race to ultra-low energy prices that has historically challenged new power technologies. This approach allows Yinsen to become a commercially viable product far earlier than traditional grid-scale fusion concepts.

The perfect market…the ocean

Shipping's Path to Net-Zero: Today's Reality

The shipping industry is **racing to decarbonize** as it seeks alternatives to bunker fuel. The International Maritime Organization (IMO) has adopted a plan targeting **net-zero emissions by 2050**, a commitment embraced by industry giants MSC and Maersk. They have already invested >$1billion into alternatives, yet none fully meets the industry's needs. This creates a unique and timely opportunity for us to introduce a truly viable solution into a market eager for breakthrough innovation.

Problems with Alternative Fuels

Alternative propulsion technologies face **substantial economic and logistical hurdles**. Fuels like ammonia and hydrogen suffer from lower energy density, requiring larger storage tanks and consequently reducing available cargo space. The production, transport, and storage infrastructure required for these fuels introduce additional cost burdens, complicating the operational logistics of shipping fleets. Moreover, traditional fuels such as HFO are increasingly subjected to stringent environmental regulations and emissions penalties, which are expected to escalate significantly over the coming decades, further eroding their economic viability.

Also, the safety and handling complexities associated with ammonia and hydrogen present considerable operational risks, requiring specialized infrastructure and, in the case of ammonia, still produces polluting NOx byproducts.

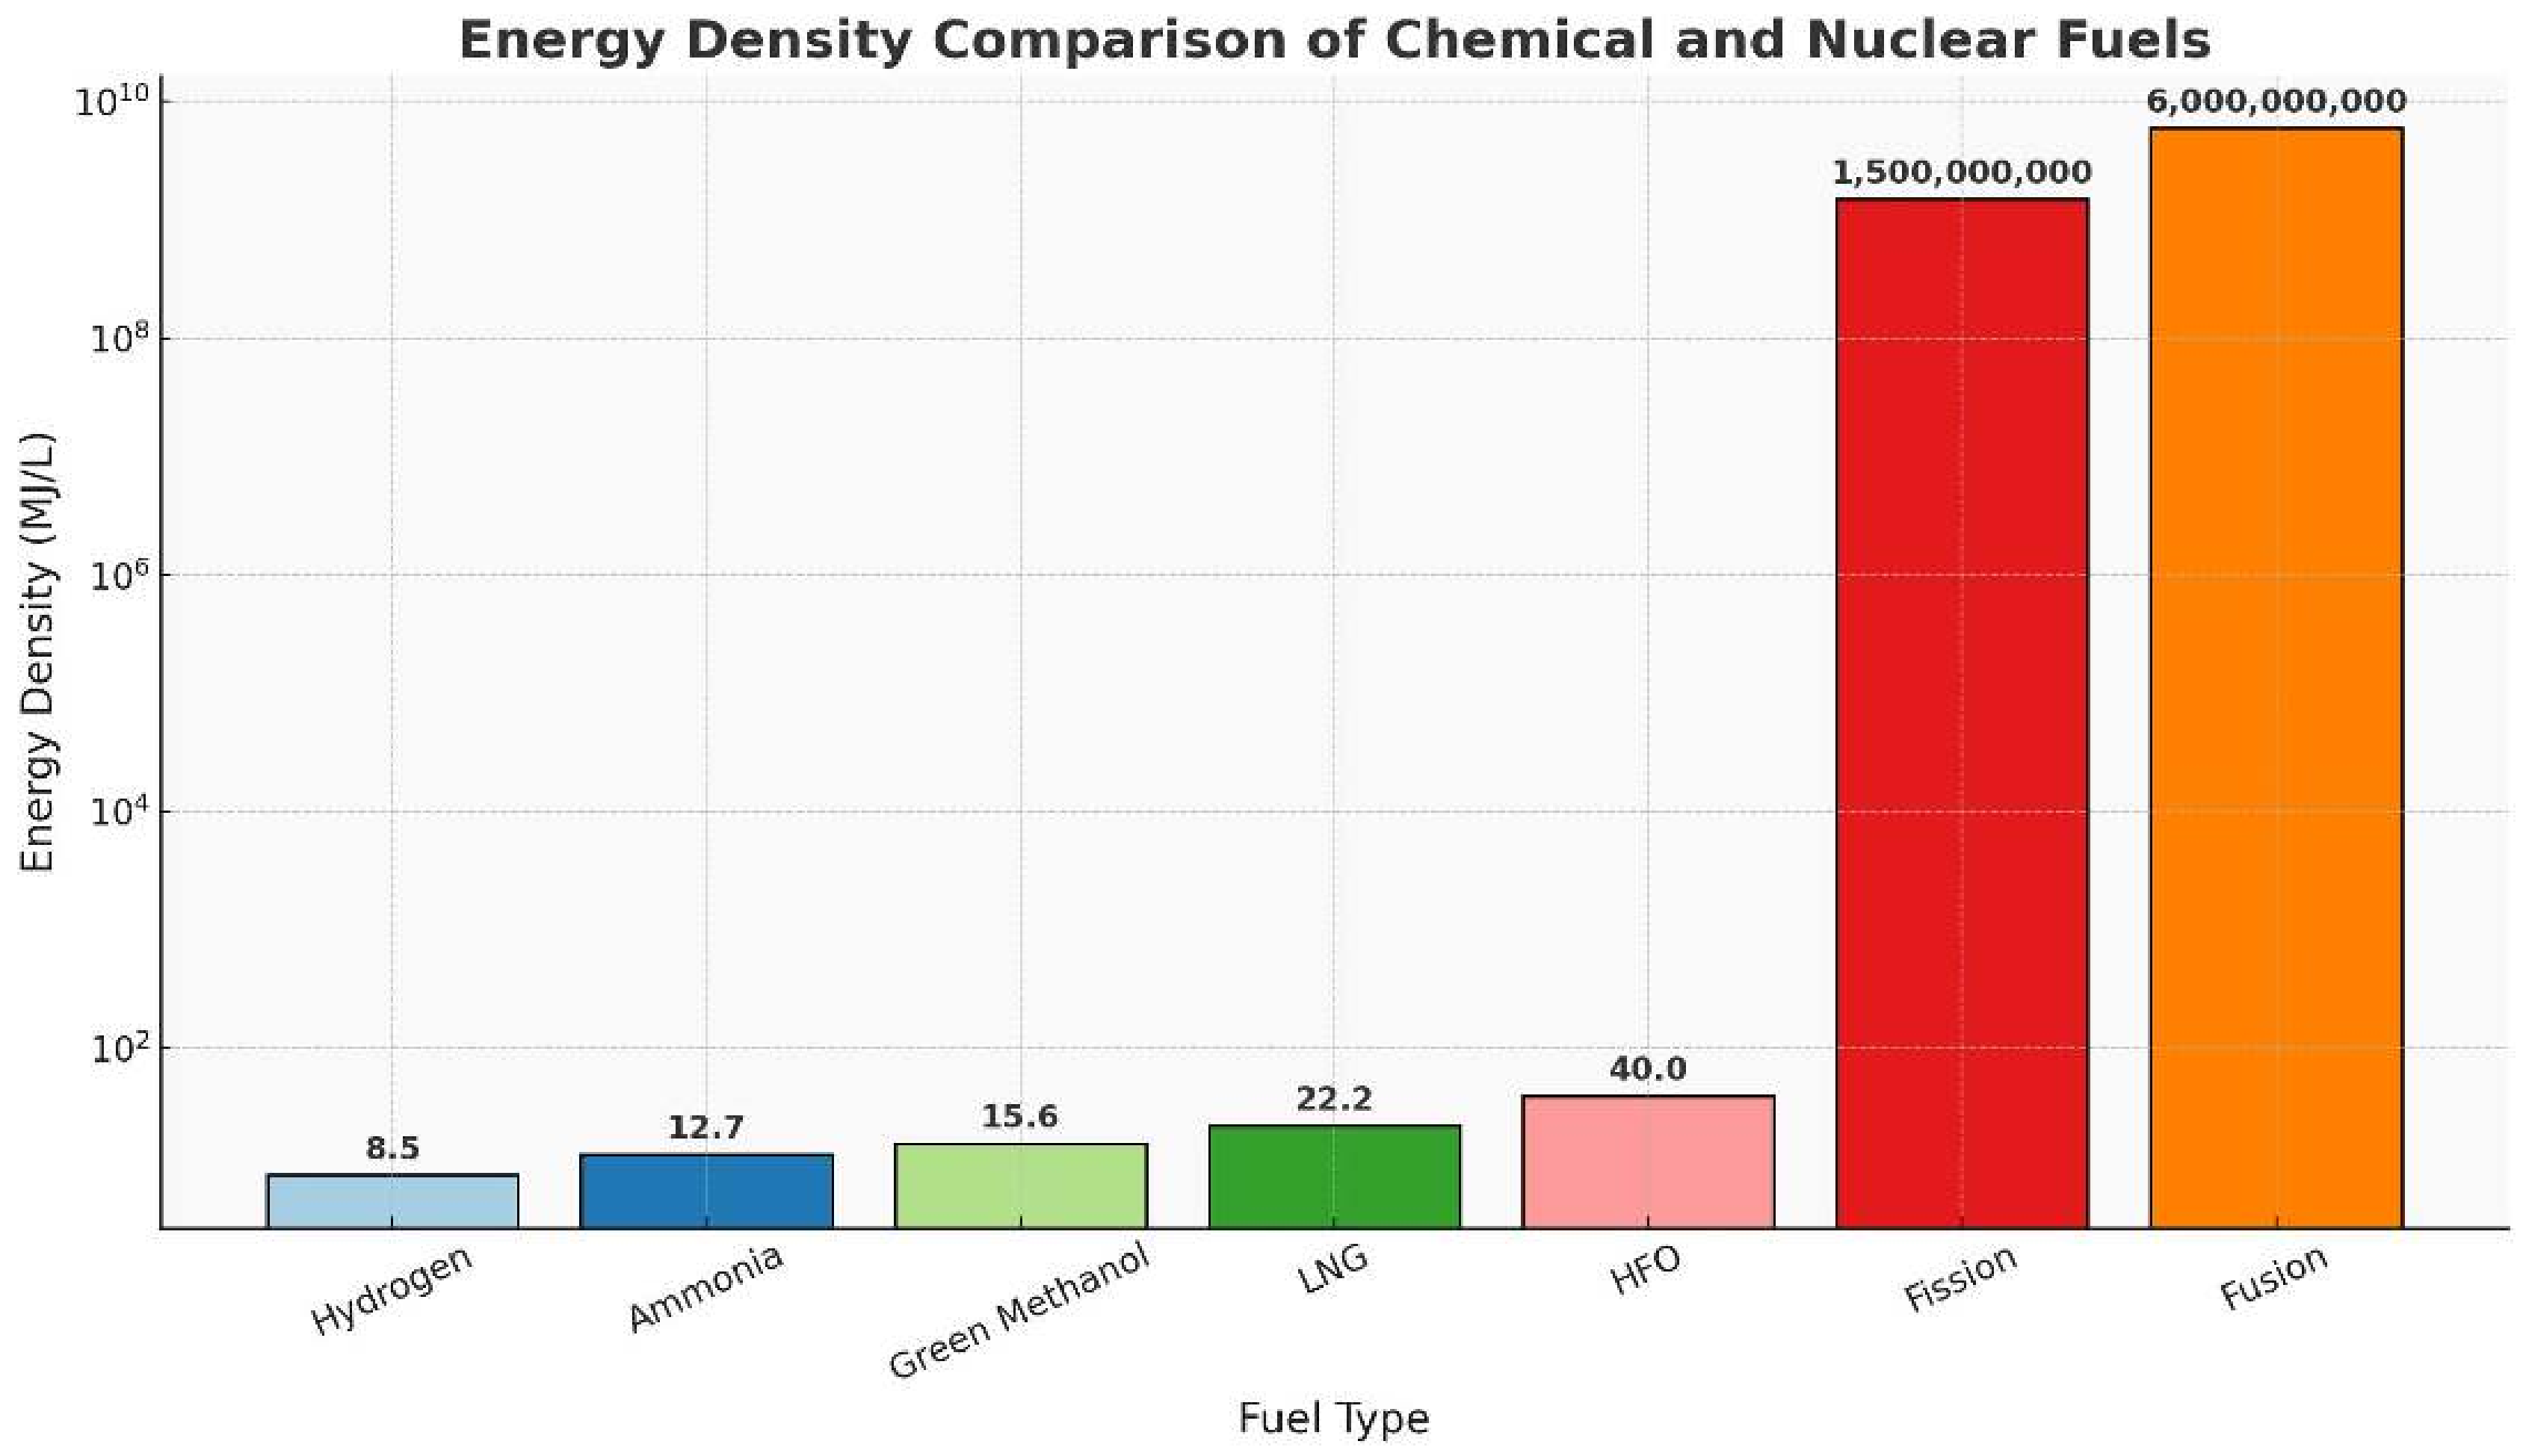

Concern	Fission (SMRs)	Fusion
Meltdown Risk	Possible	Impossible
Radioactive Waste	High	Low
Proliferation Risk	Concerning	Minimal
Attack Vulnerability	Ecological catastrophe	Loss of ship
Public Perception	Skepticism	More positive

Source: https://www.fusionindustryassociation.org/nrc-decision-separates-fusion-energy-regulation-from-nuclear-fission/

What's Wrong with Fission?

While Small Modular Reactors (SMRs) appear very promising from a technological perspective, deploying fission reactors on ships introduces critical challenges. Fission carries inherent meltdown risks - intensified by the harsh maritime environment - and produces significant amounts of highly radioactive waste requiring secure long-term storage. Additionally, enriched uranium used in naval reactors raises proliferation and weaponization concerns. An attack could result in catastrophic environmental contamination as well. Despite recent safety advancements, persistent public apprehension—driven by historical nuclear accidents—has halted many land-based reactors and makes global acceptance of fission-powered ships unlikely. Conversely, fusion reactors present a safer alternative by eliminating meltdown risks, significantly reducing radioactive waste, easing proliferation fears, and benefiting from more favorable public perception. The NRC has recently **separated fusion regulation from fission**, noting they will be regulated like a normal particle accelerator that already exists at labs and universities all over the US.

Benefits of Fusion

Nuclear fusion propulsion offers substantial economic and operational advantages compared to traditional and emerging green fuels:

- **High Fuel Volumetric Energy Density**:
  - Fusion fuel offers energy density far exceeding traditional fuels (HFO, ammonia, hydrogen).
  - Less space needed for fuel storage and bunkering.
  - Enables large ships to transport more goods, boosting voyage profitability.
- **Superior Operational Efficiency**:
  - Fusion propulsion maintains high efficiency even at maximum rated power.
  - Unlike conventional fuels, fusion energy cost doesn't scale significantly with output power.
  - Ships can consistently operate at maximum hull speeds without increased operational costs, improving transit times and competitiveness.
- **Extended Reactor Lifespan**:
  - Fusion reactor cores are expected to last 30+ years vs ~15-20 years for current vessel propulsion systems.
  - Reduces frequency of vessel replacements, lowering long-term capital expenditures enabled by the same fusion powered vessel enduring what would've been 2 procurement cycles of traditional vessels.
- **Economic Feasibility**:
  - Higher initial capital construction and regulatory licensing costs, though notably streamlined compared to fission reactors.
  - Minimal fuel costs significantly offset initial investments and life-cycle expenses.

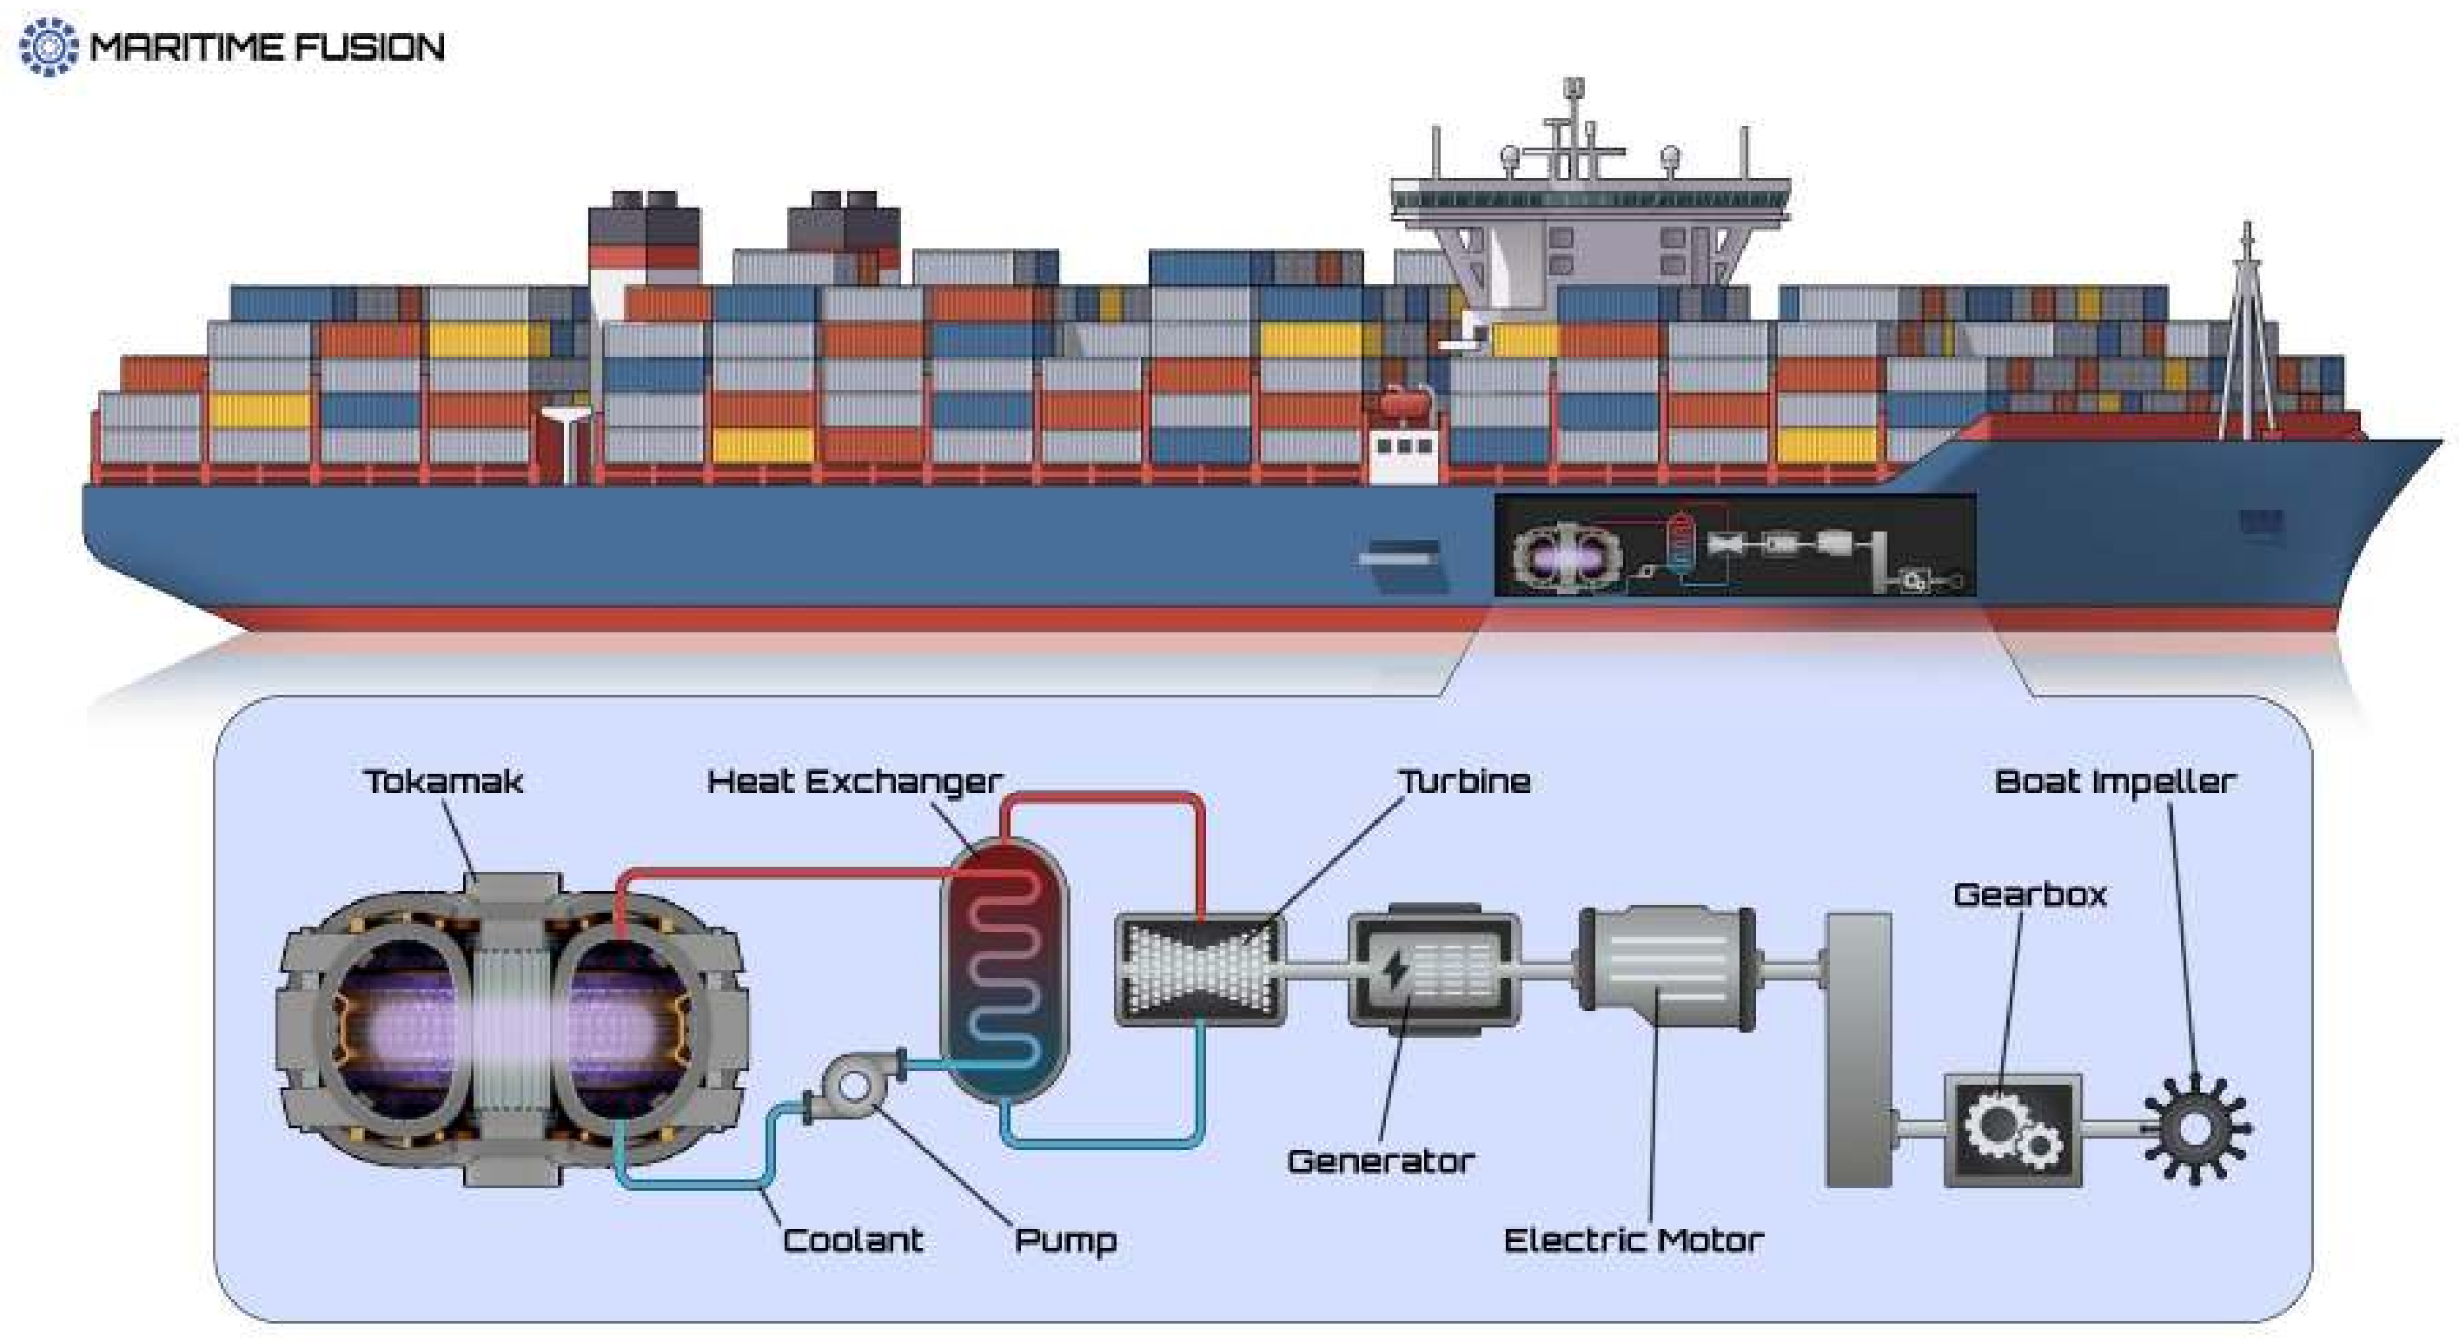

Basic systems layout of our proposed fusion propelled cargo ship

Still from our animated concept video (link here)

How we compete

Projected Costs

The economic feasibility of fusion energy depends on minimizing capital costs while ensuring long-term operational viability. **Yinsen's direct cost is estimated at $1.1 billion**, with the reactor plant accounting for 84.3% of the total. The primary cost drivers include the toroidal field (TF) magnets, shielding and first wall materials, and heating systems. This analysis utilizes the ARIES costing model (UCSD-CER-13-01, Center for Energy Research), where the largest uncertainties stem from fabrication scale factors and material costs. Key assumptions include an HTS cost of $65/kA-m and tungsten priced at $85/kg.

Total Capex Cost:
$1,109,557,403.00

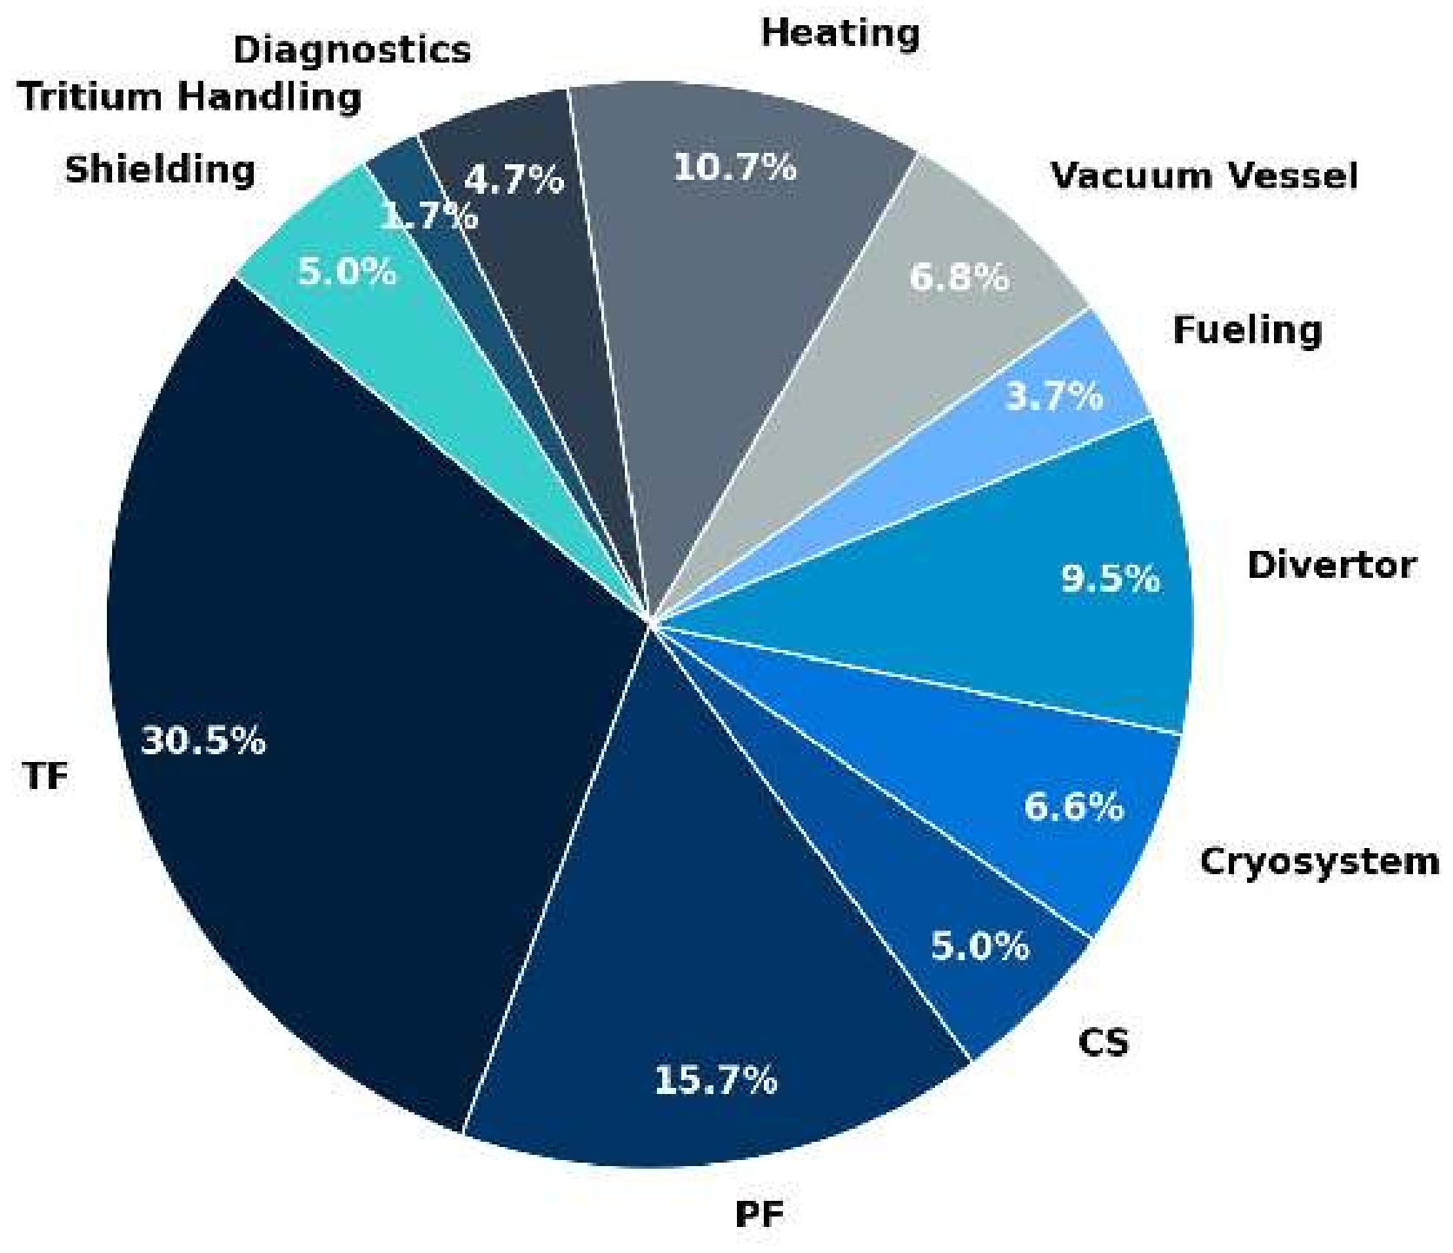

Cost Projection Over Time

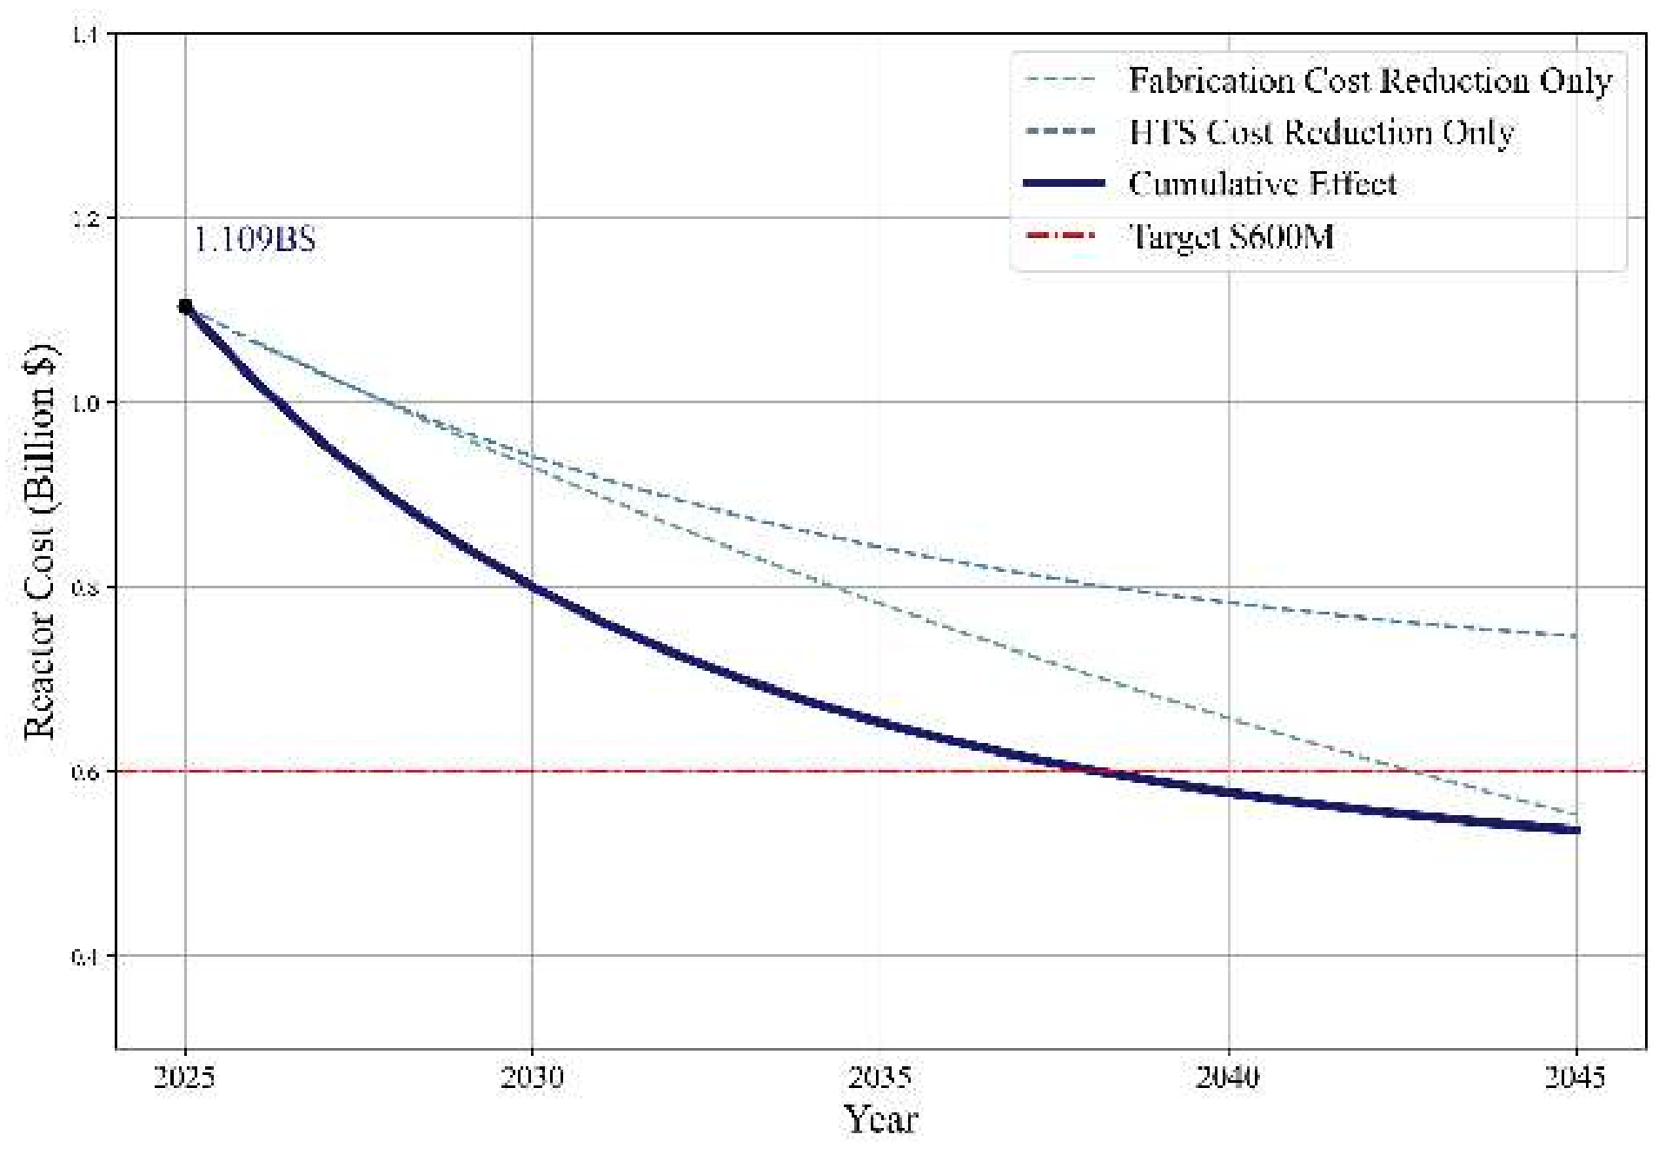

Capex

HTS accounts for ~28% of total costs, or approximately $310 million. Today's HTS pricing is about $100/kA-m, but projections suggest costs will drop to $20/kA-m by the mid 2030s, potentially reducing HTS-related costs by over $200 million and significantly lowering overall reactor costs.

Additionally, as manufacturing scales and supply chains mature, subsequent reactor units will be significantly cheaper than the first-of-a-kind (FOAK) build. With lower HTS costs and improved production efficiencies, Maritime Fusion estimates later units to cost **$600 million per 25MWe reactor**, making fusion a commercially viable energy solution for the maritime sector.

Capex Breakdown by Subsystem

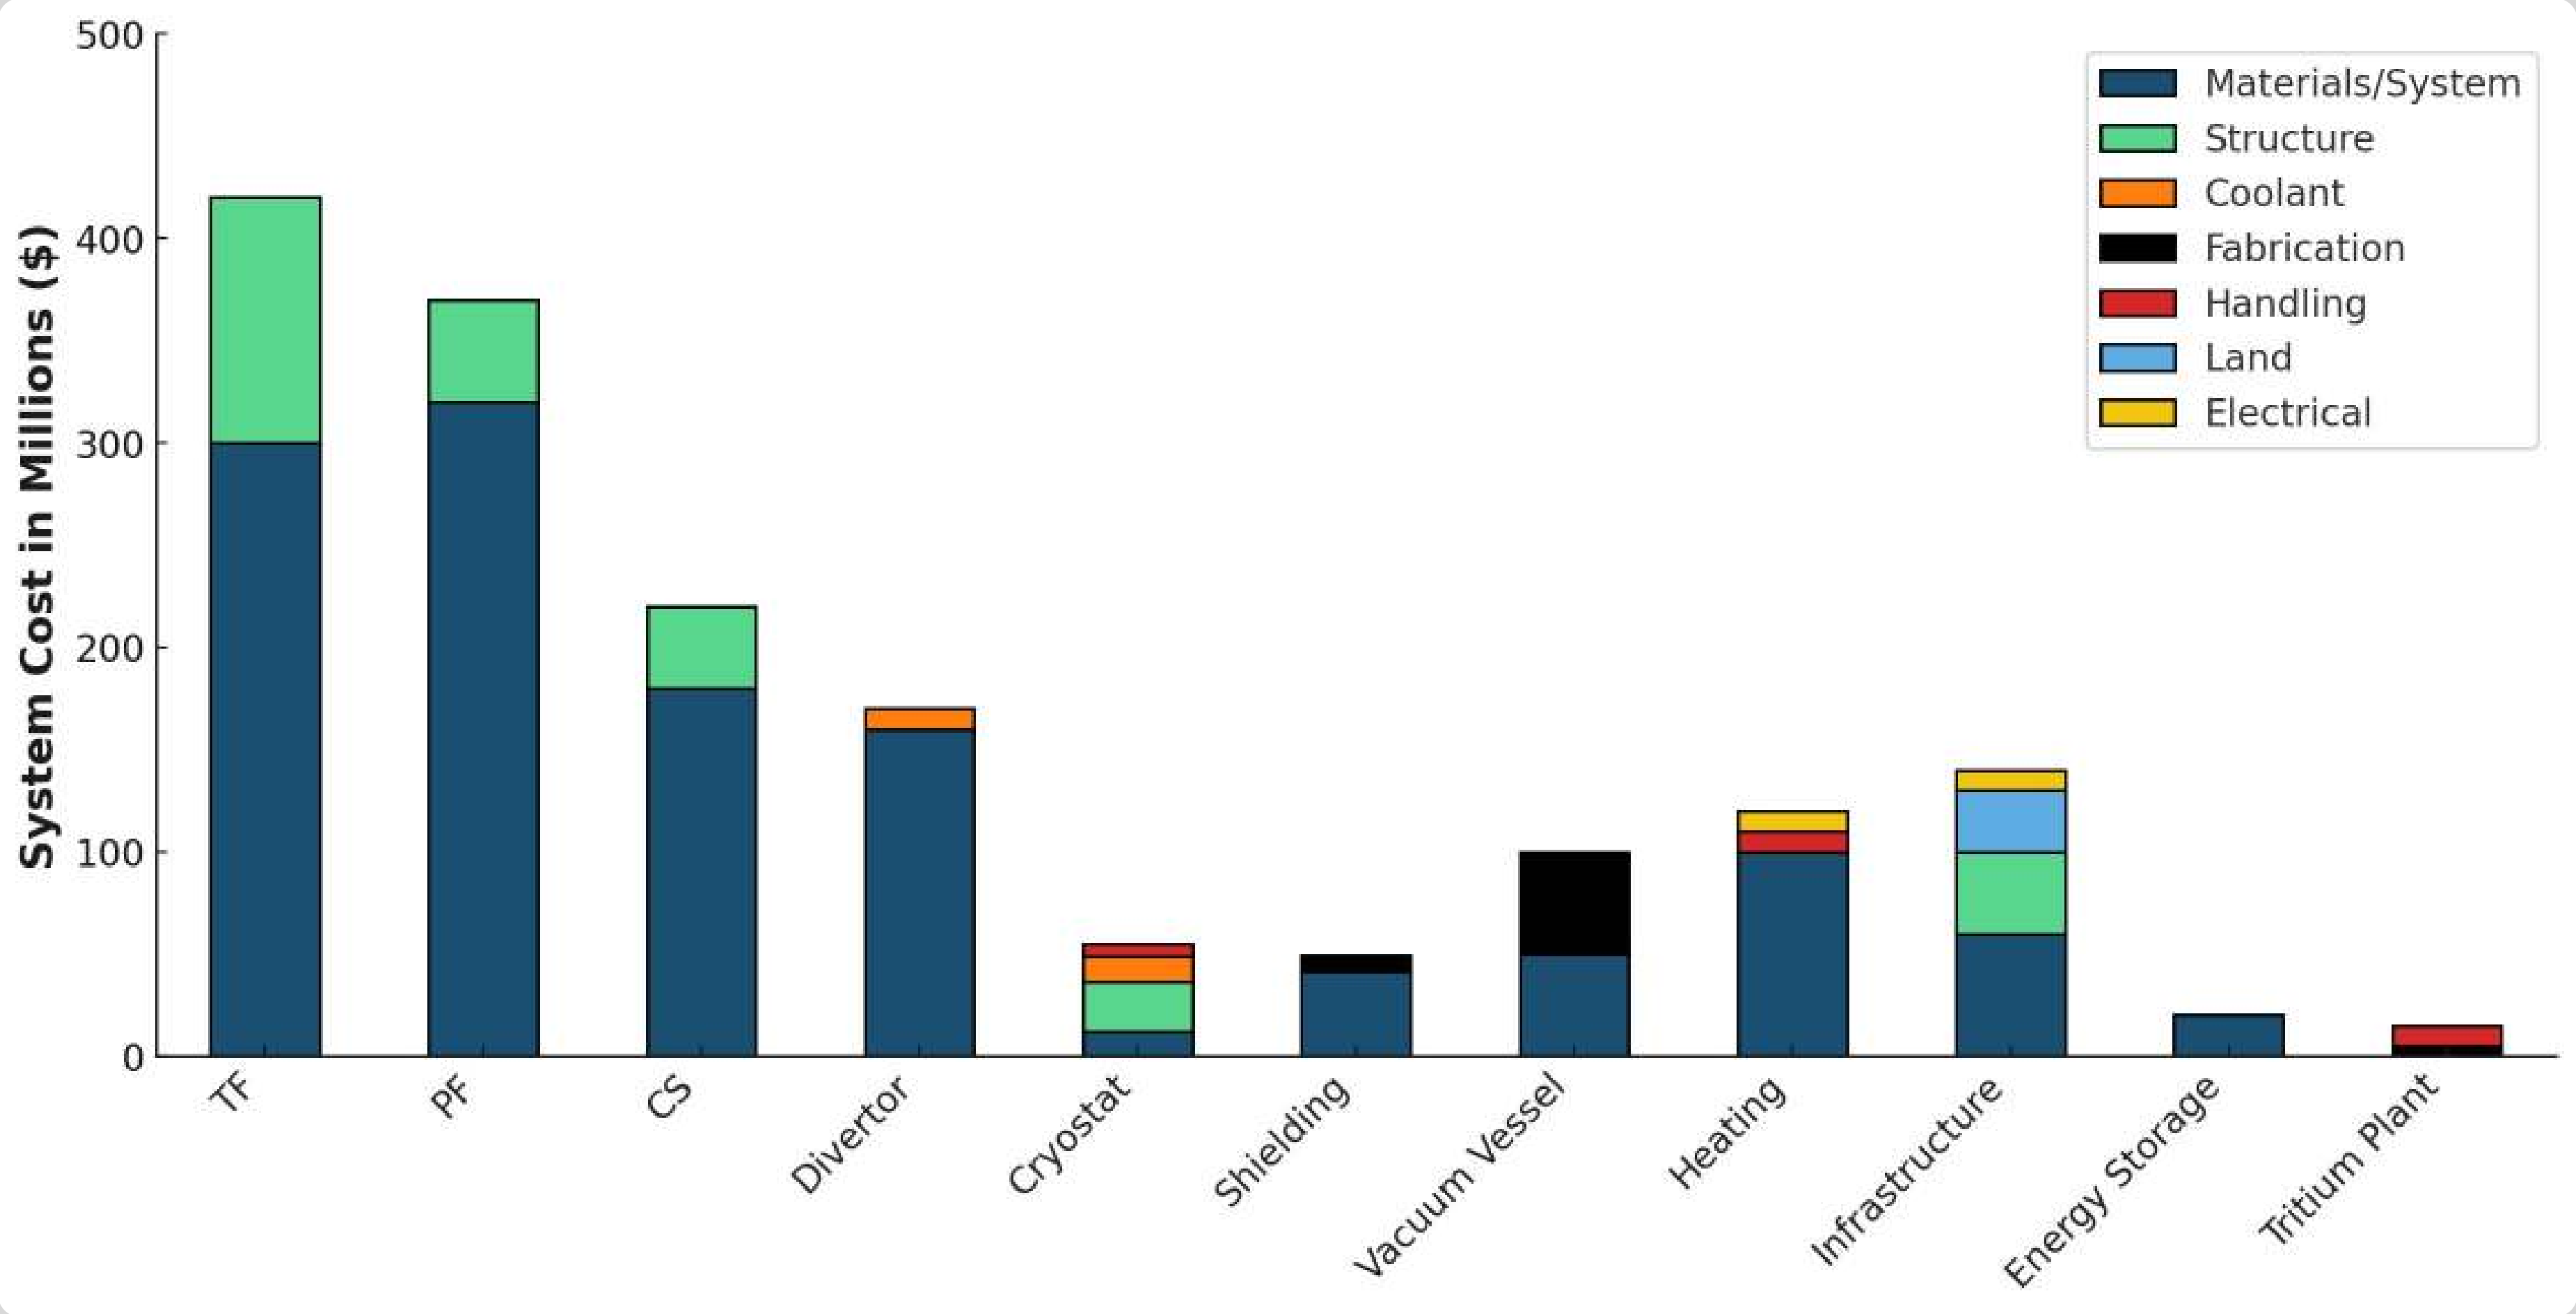

Opex

Based on recent studies, we can estimate the operating expenses to be $5M - $10M per year given a 30yr lifetime for Yinsen. Using ITER scaling, we can estimate a first wall replacement every few years at our power flux levels.

Divertor and blanket/first wall replacements will be much more infrequent compared to grid scale reactors, we can estimate 5% of total capital cost to conduct component replacements, due to the lower power density operation.

Opex Breakdown

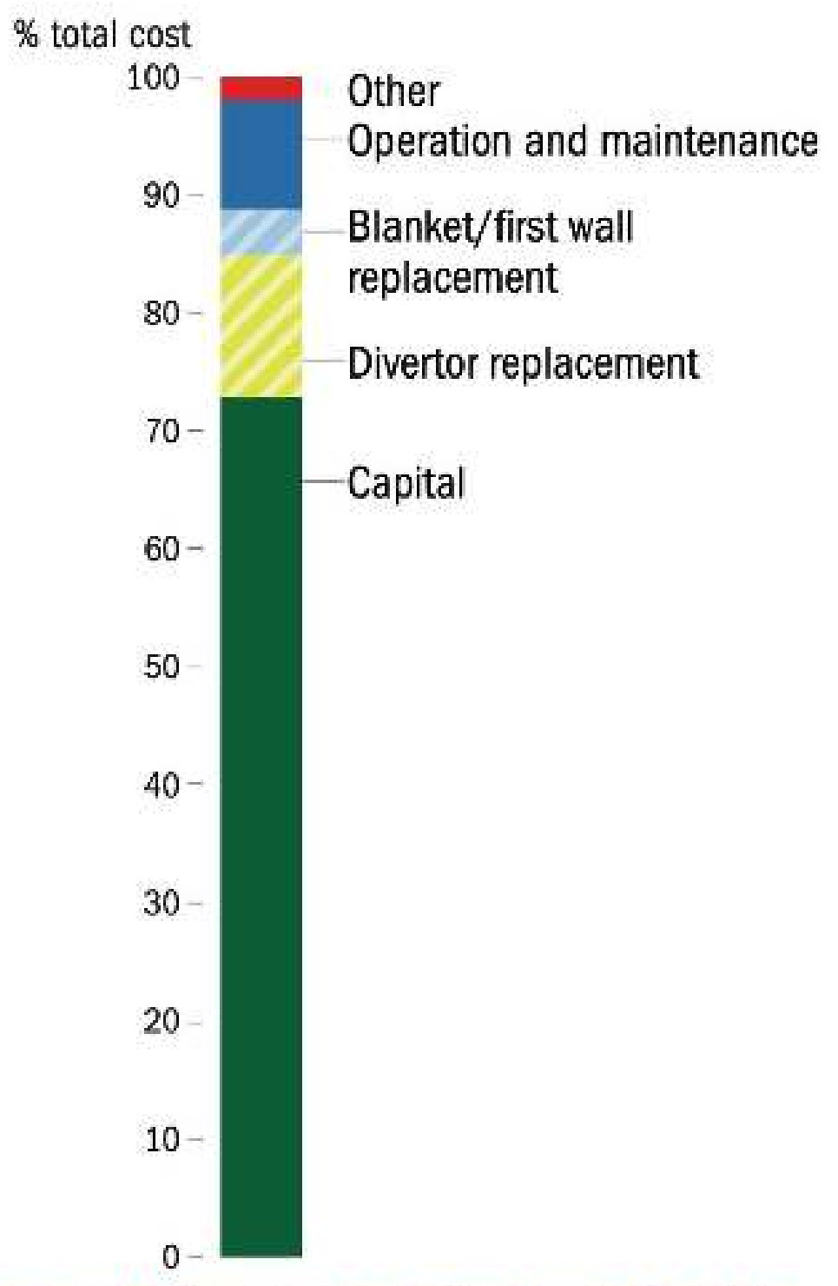

Figure 4.2: Components of the total cost of electricity produced by a fusion power plant based on a tokamak reactor, shown as a percent of total cost *[1]*.

Radiation Absorbed by Blanket Before Replacement

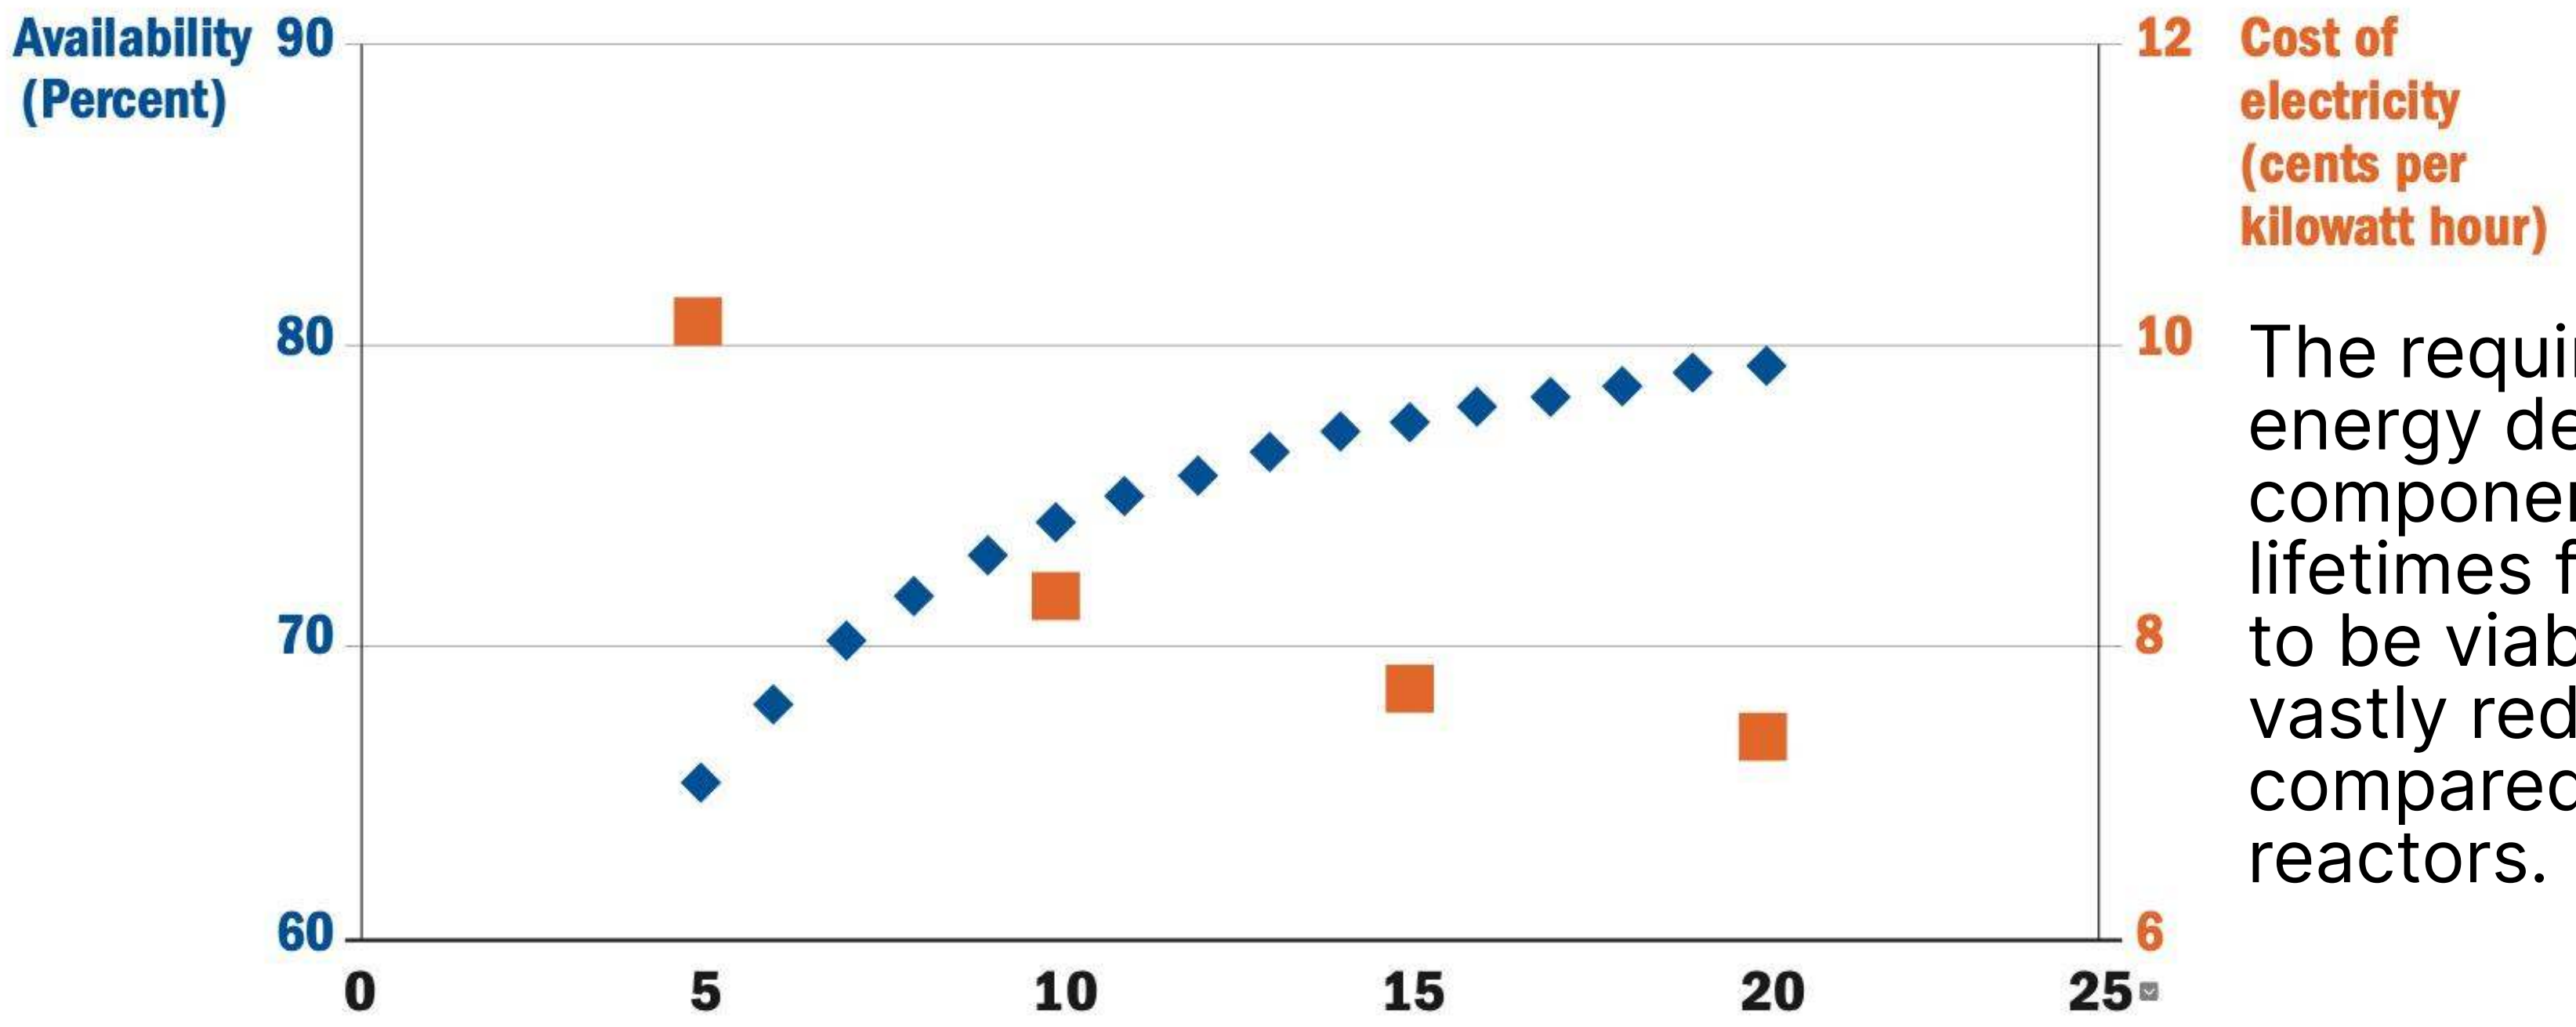

megawatt years per meter²

Source: anlinger center for energy + the environment, Article 4

The requirement on energy density and component lifetimes for Yinsen to be viable is vastly reduced compared to grid reactors.

Cost Comparison to Alternatives

Unlike diesel/alternative fuels, which incurs tens of millions in annual fuel costs - fusion has essentially zero fuel expenses, allowing it to pay for itself over time with a lower operating cost and higher long-term returns. Estimating the future cost of alternative fuels is non-trivial since it involves not only fuel + carbon emission fees + engine + cargo space calculations, but port fueling infrastructure and fuel synthesis/ transportation. Given that, we calculate at 8 **years of operation** for a standard 20,000 TEU cargo ship, the cost of fusion should become less than diesel/ammonia/green methanol/hydrogen, which is **within the operating lifespan** of most commercial shipping vessels. (notable assumptions listed to the right, SOAK used).

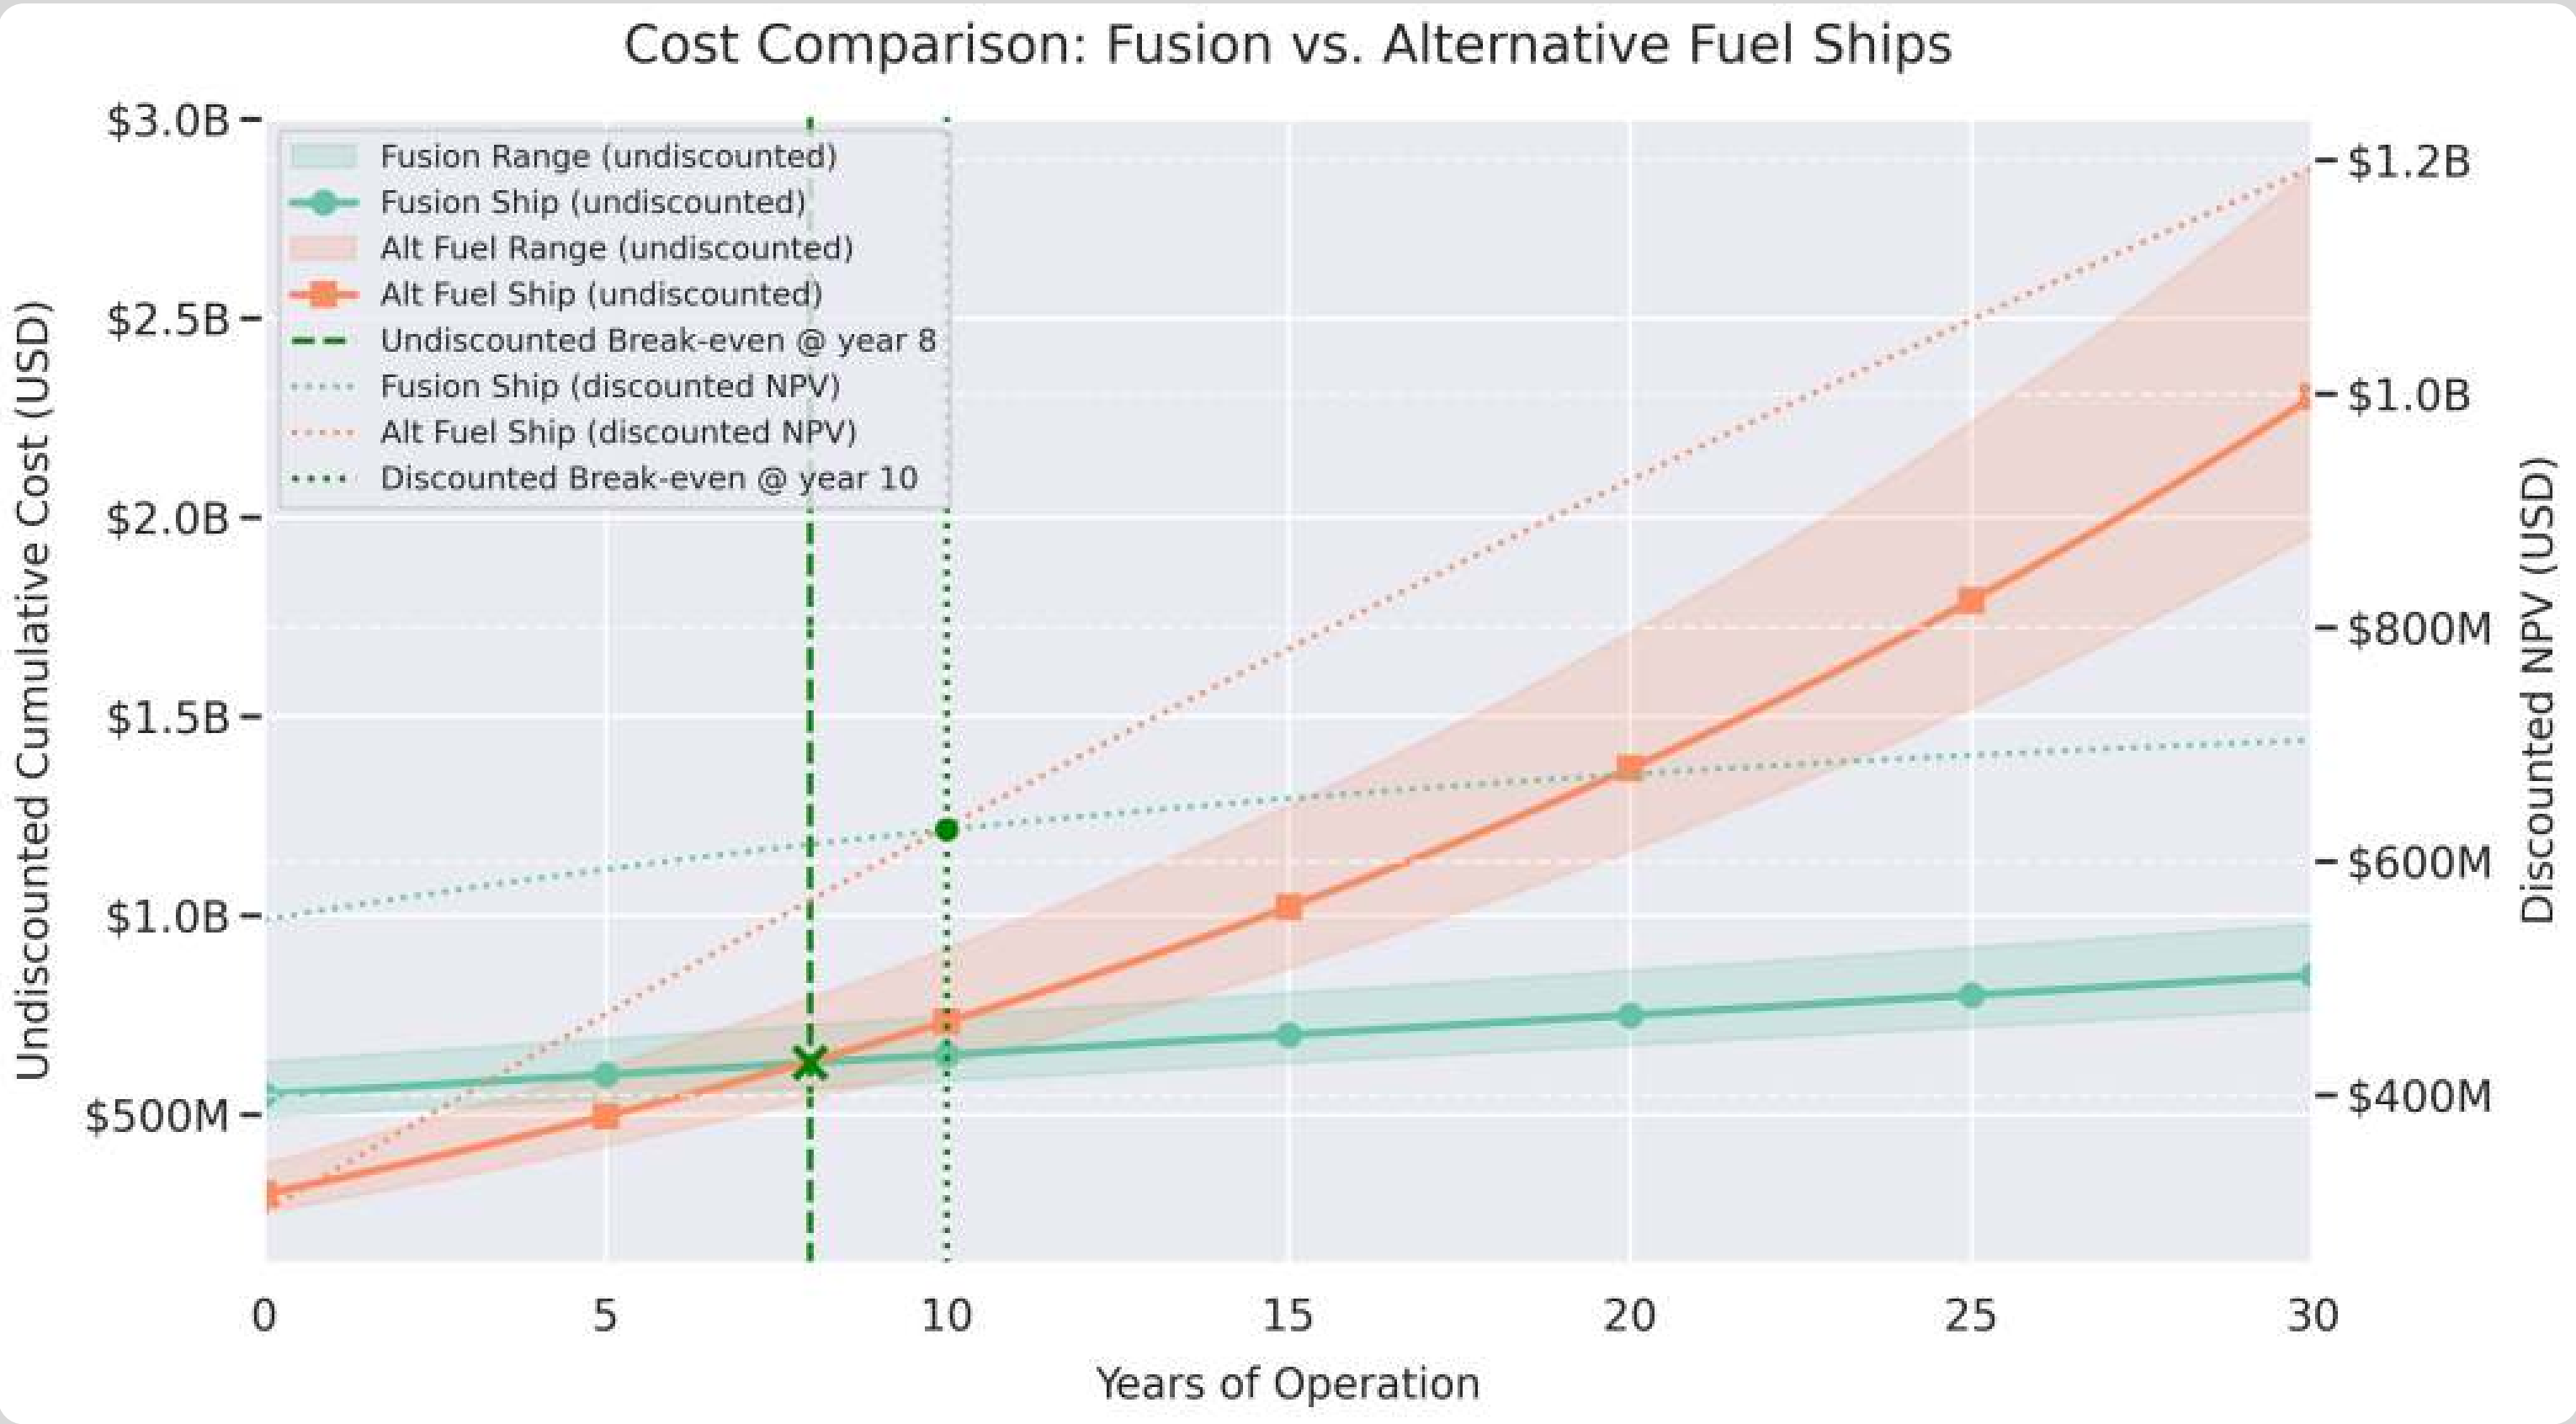

Assumptions:

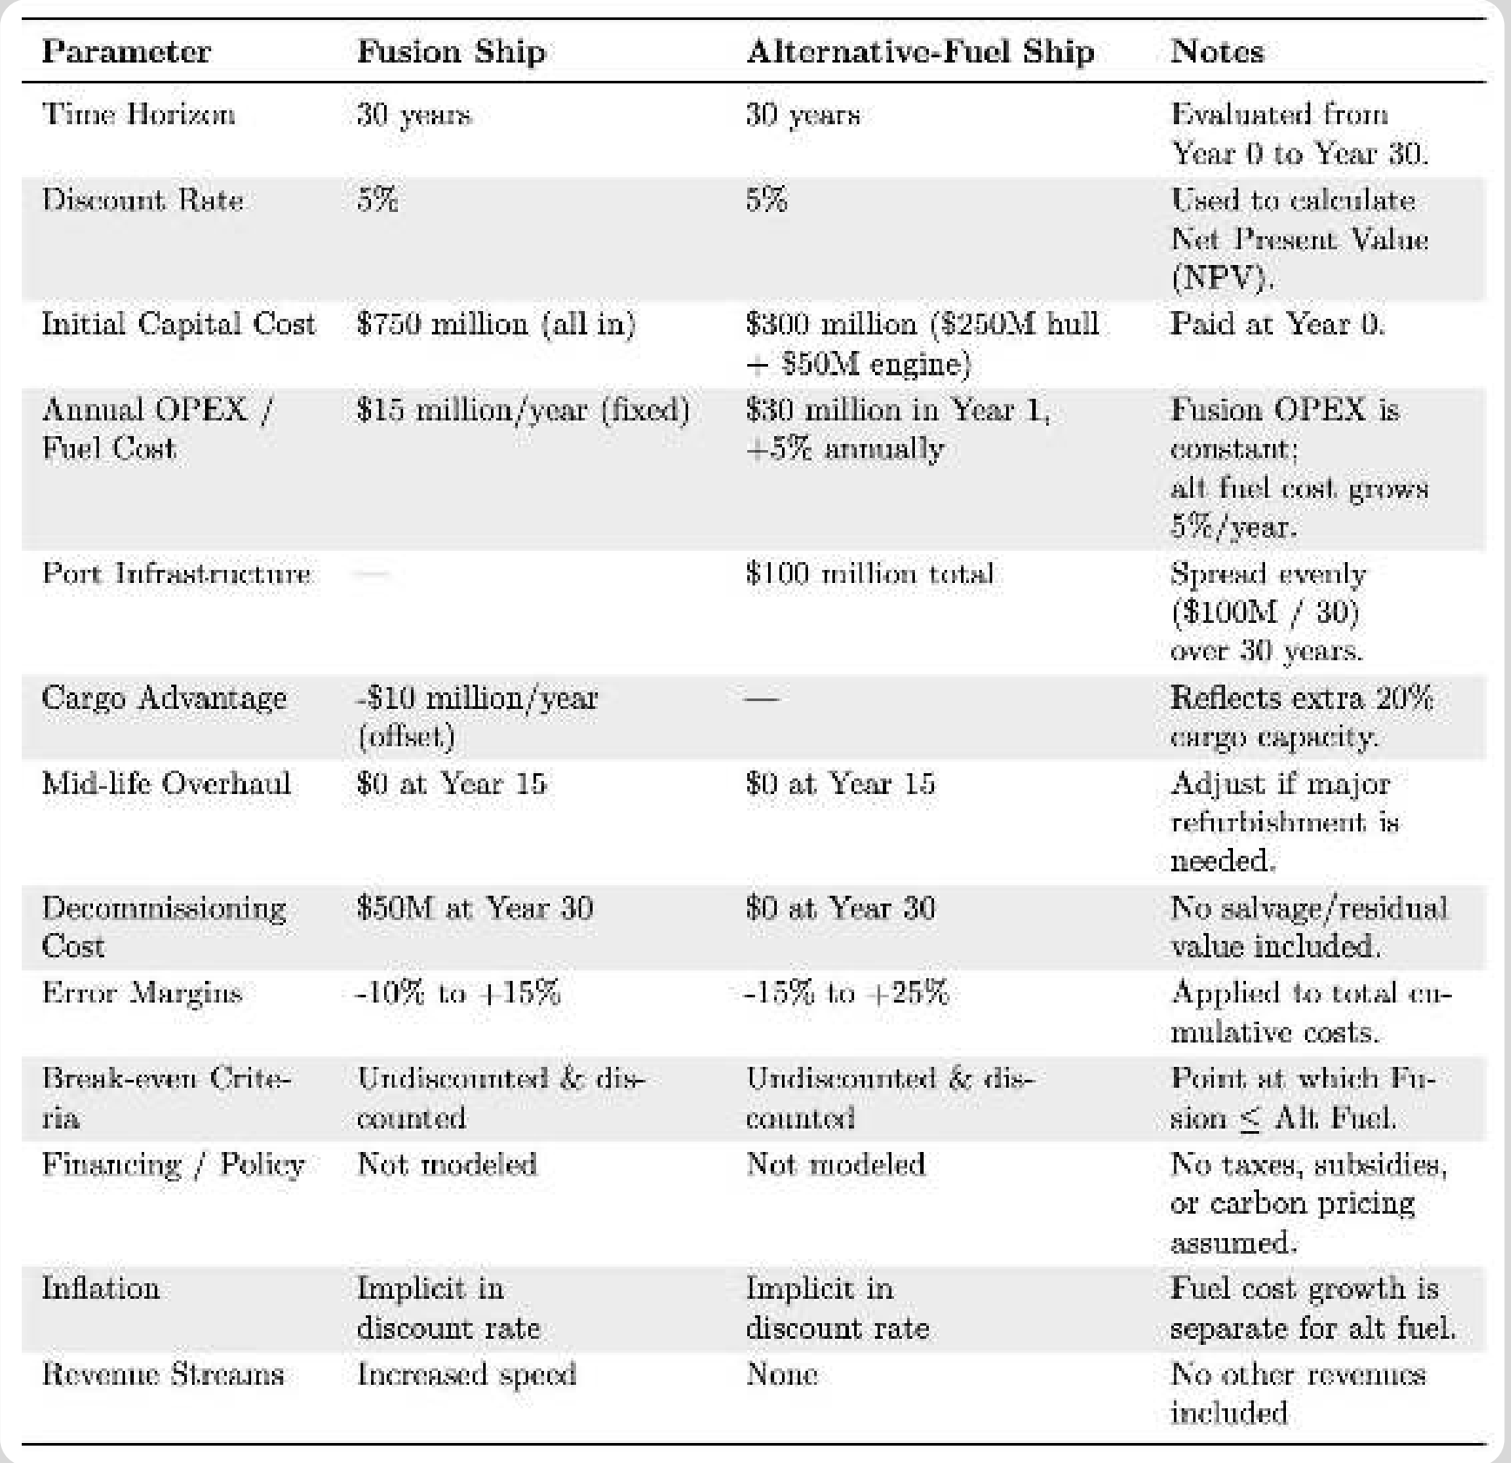

Parameter	Fusion Ship	Alternative-Fuel Ship	Notes
Time Horizon	30 years	30 years	Evaluated from Year 0 to Year 30.
Discount Rate	5%	5%	Used to calculate Net Present Value (NPV).
Initial Capital Cost	$750 million (all in)	$300 million ($250M hull + $50M engine)	Paid at Year 0.
Annual OPEX / Fuel Cost	$15 million/year (fixed)	$30 million in Year 1, +5% annually	Fusion OPEX is constant; alt fuel cost grows 5%/year.
Port Infrastructure	—	$100 million total	Spread evenly ($100M / 30) over 30 years.
Cargo Advantage	-$10 million/year (offset)	—	Reflects extra 20% cargo capacity.
Mid-life Overhaul	$0 at Year 15	$0 at Year 15	Adjust if major refurbishment is needed.
Decommissioning Cost	$50M at Year 30	$0 at Year 30	No salvage/residual value included.
Error Margins	-10% to +15%	-15% to +25%	Applied to total cumulative costs.
Break-even Criteria	Undiscounted & discounted	Undiscounted & discounted	Point at which Fusion ≤ Alt Fuel.
Financing / Policy	Not modeled	Not modeled	No taxes, subsidies, or carbon pricing assumed.
Inflation	Implicit in discount rate	Implicit in discount rate	Fuel cost growth is separate for alt fuel.
Revenue Streams	Increased speed	None	No other revenues included

Chapter 4

Technology

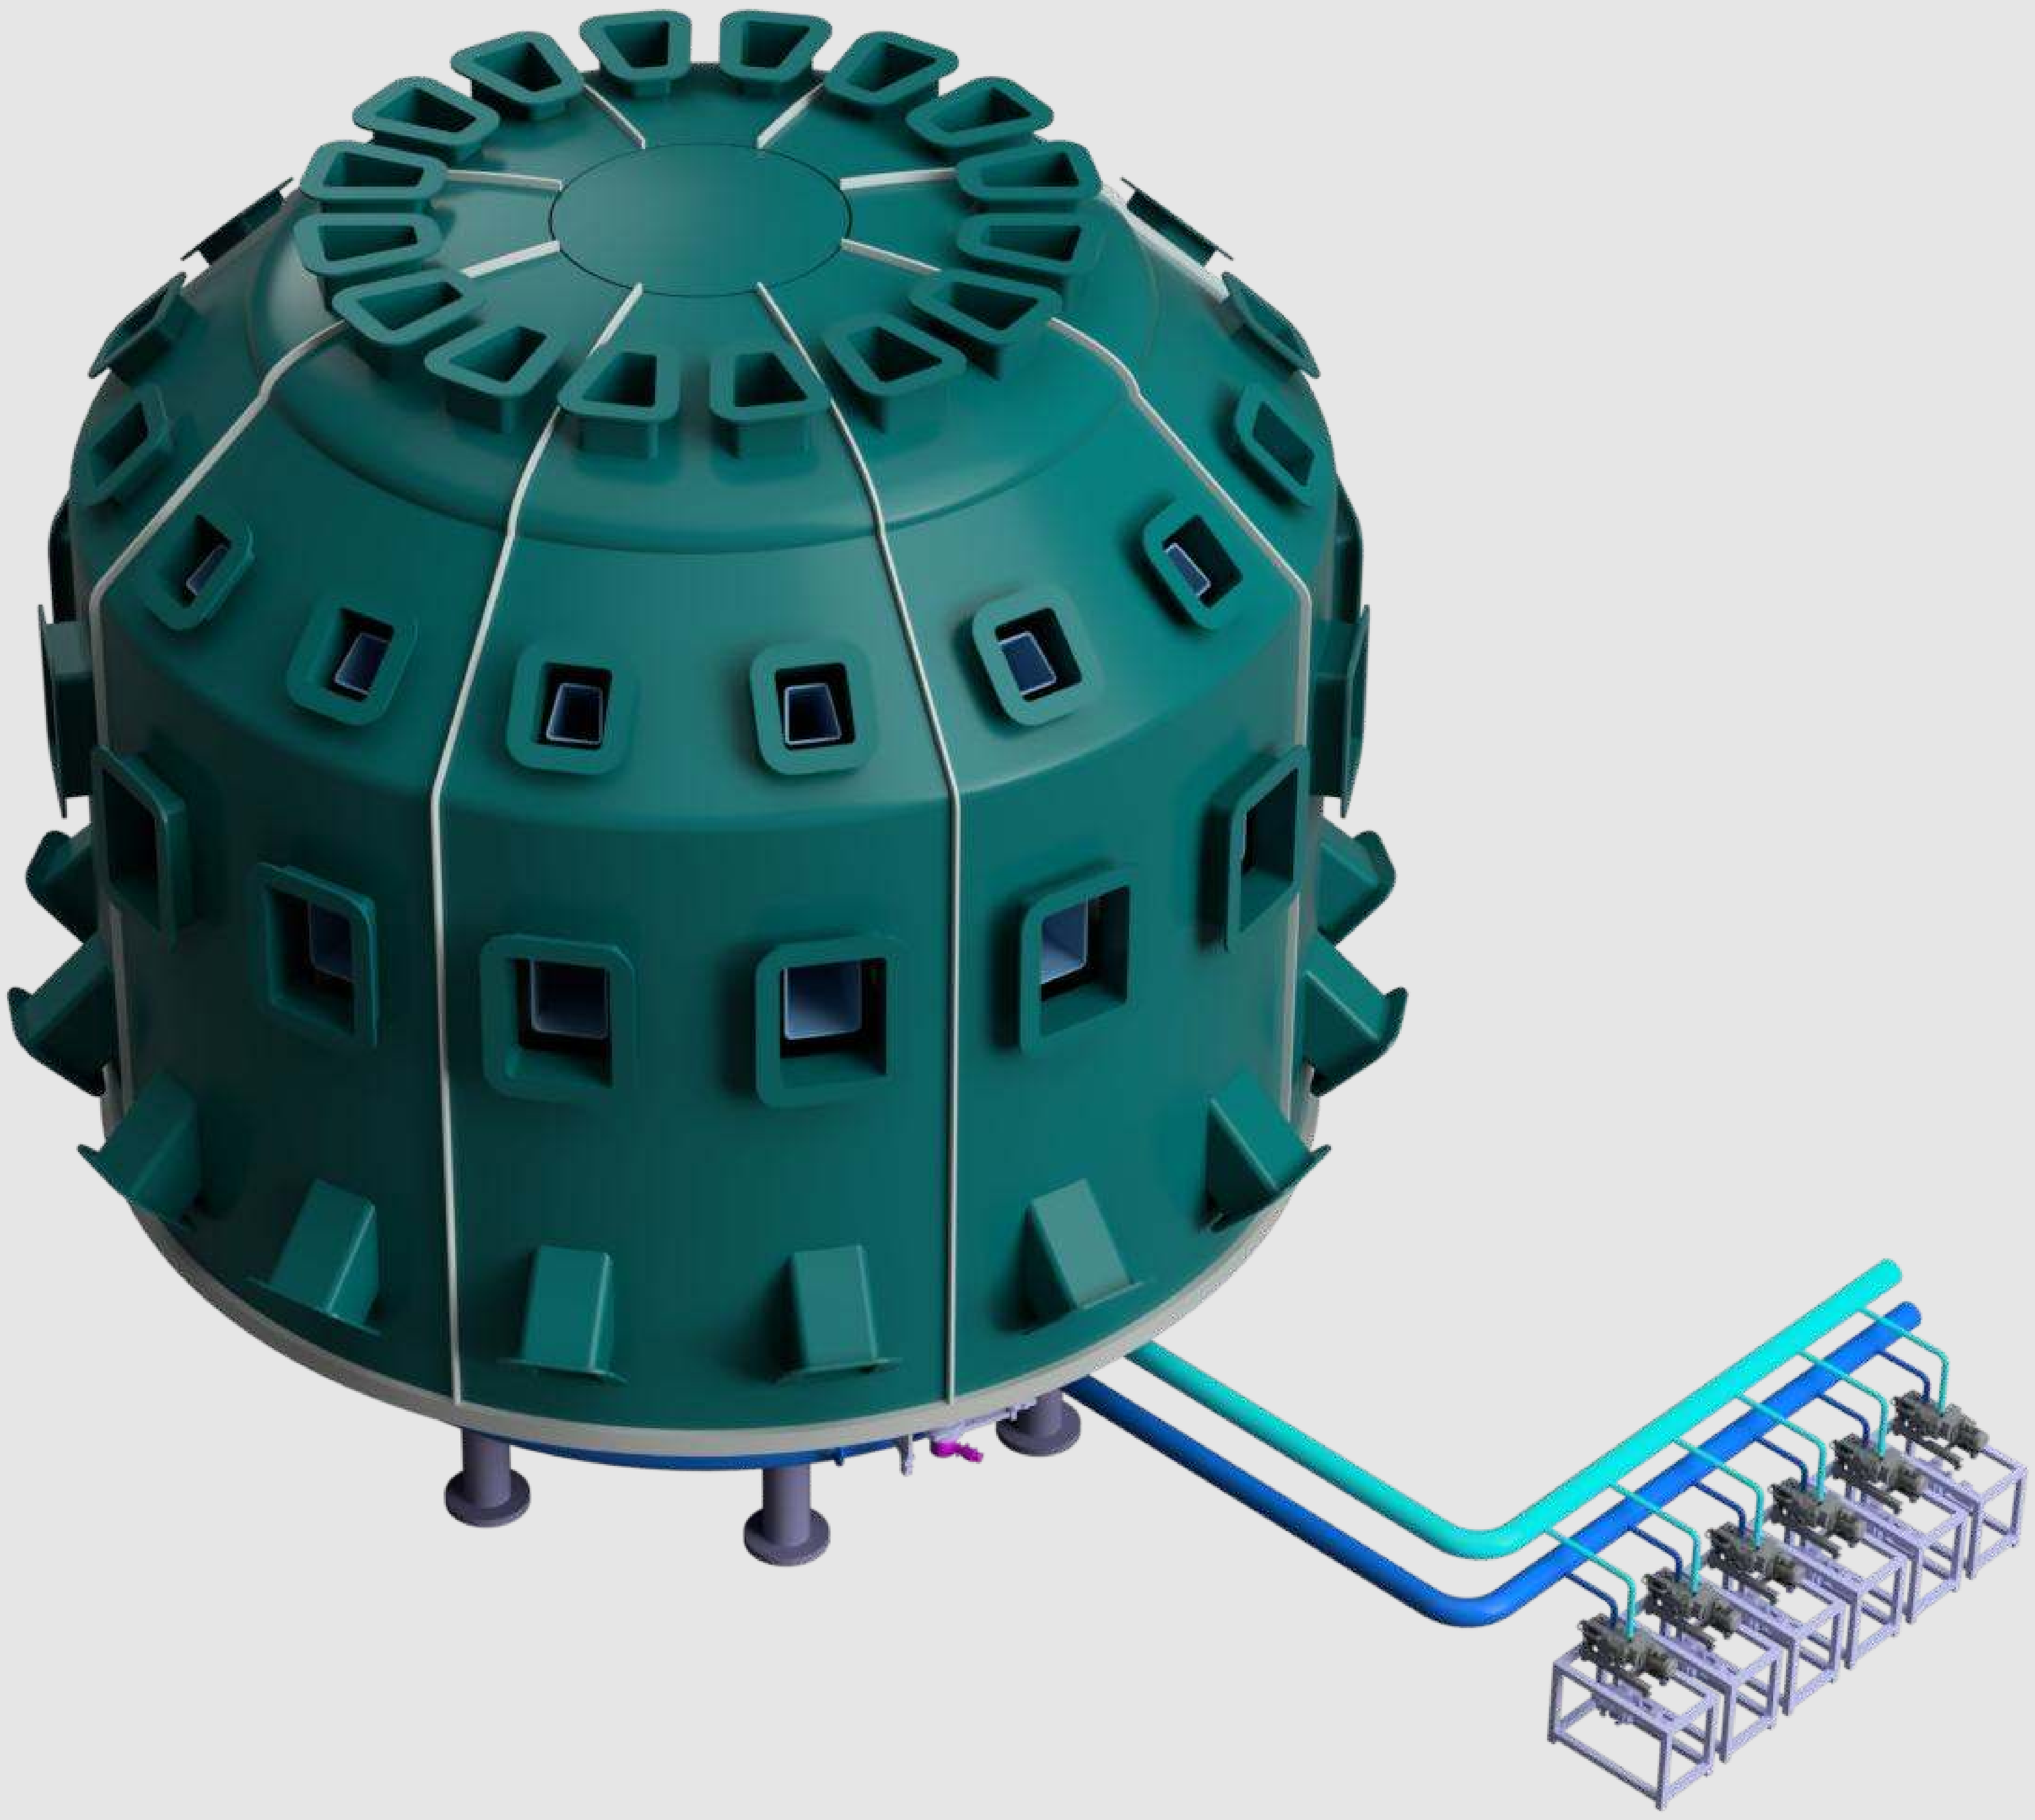

Reactor Design Optimization

0D Optimization

We developed a custom POPCON (Plasma OPerating CONtour) software and optimization algorithm in C++ to determine which set of input parameters (R, a, B_0, I_p) best reduces the first wall power flux - while simultaneously achieving satisfactory plasma stability, pulse duration, and net power.

The software can run over 1 billion simulations per minute, compared to <100 with standard Python alternatives.

This allows us to search the entire operating space with super fine grained steps. Our simulations show that we favor **higher major radius and minor radius to reduce neutron flux**, and its possible to **reduce our centerline magnetic field while still producing enough power** by increasing energy confinement time as we **increase the plasma current**.

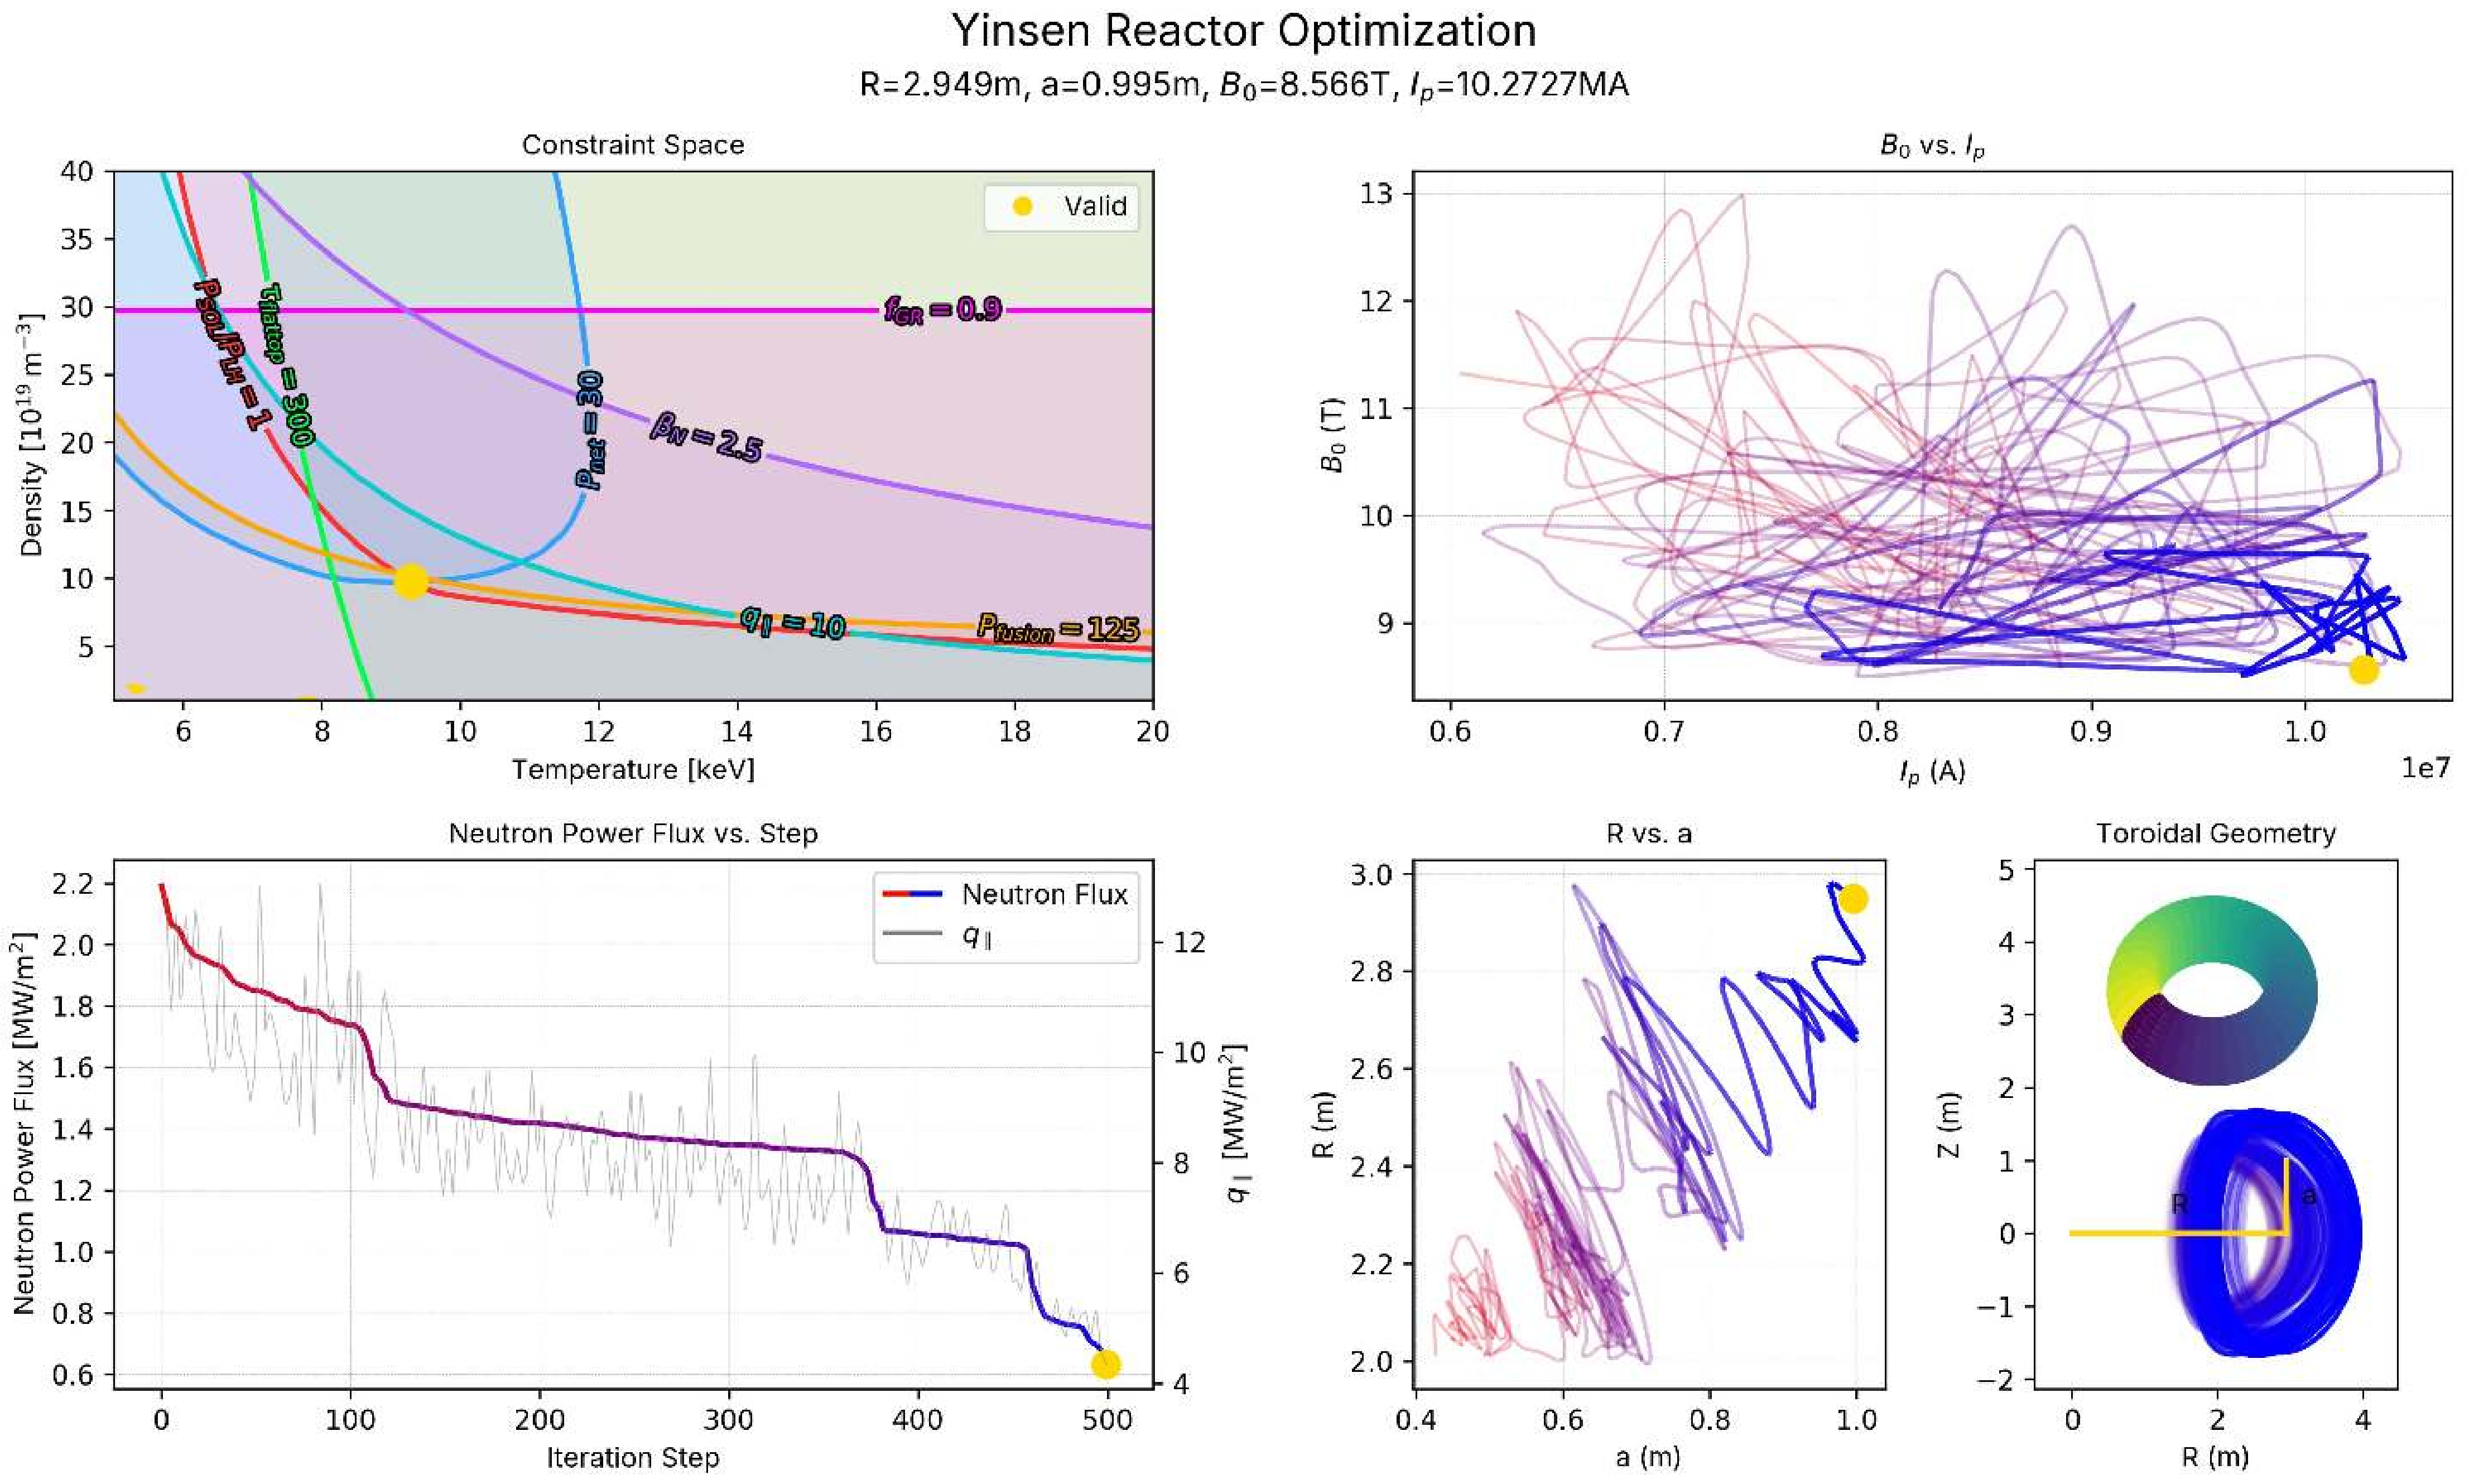

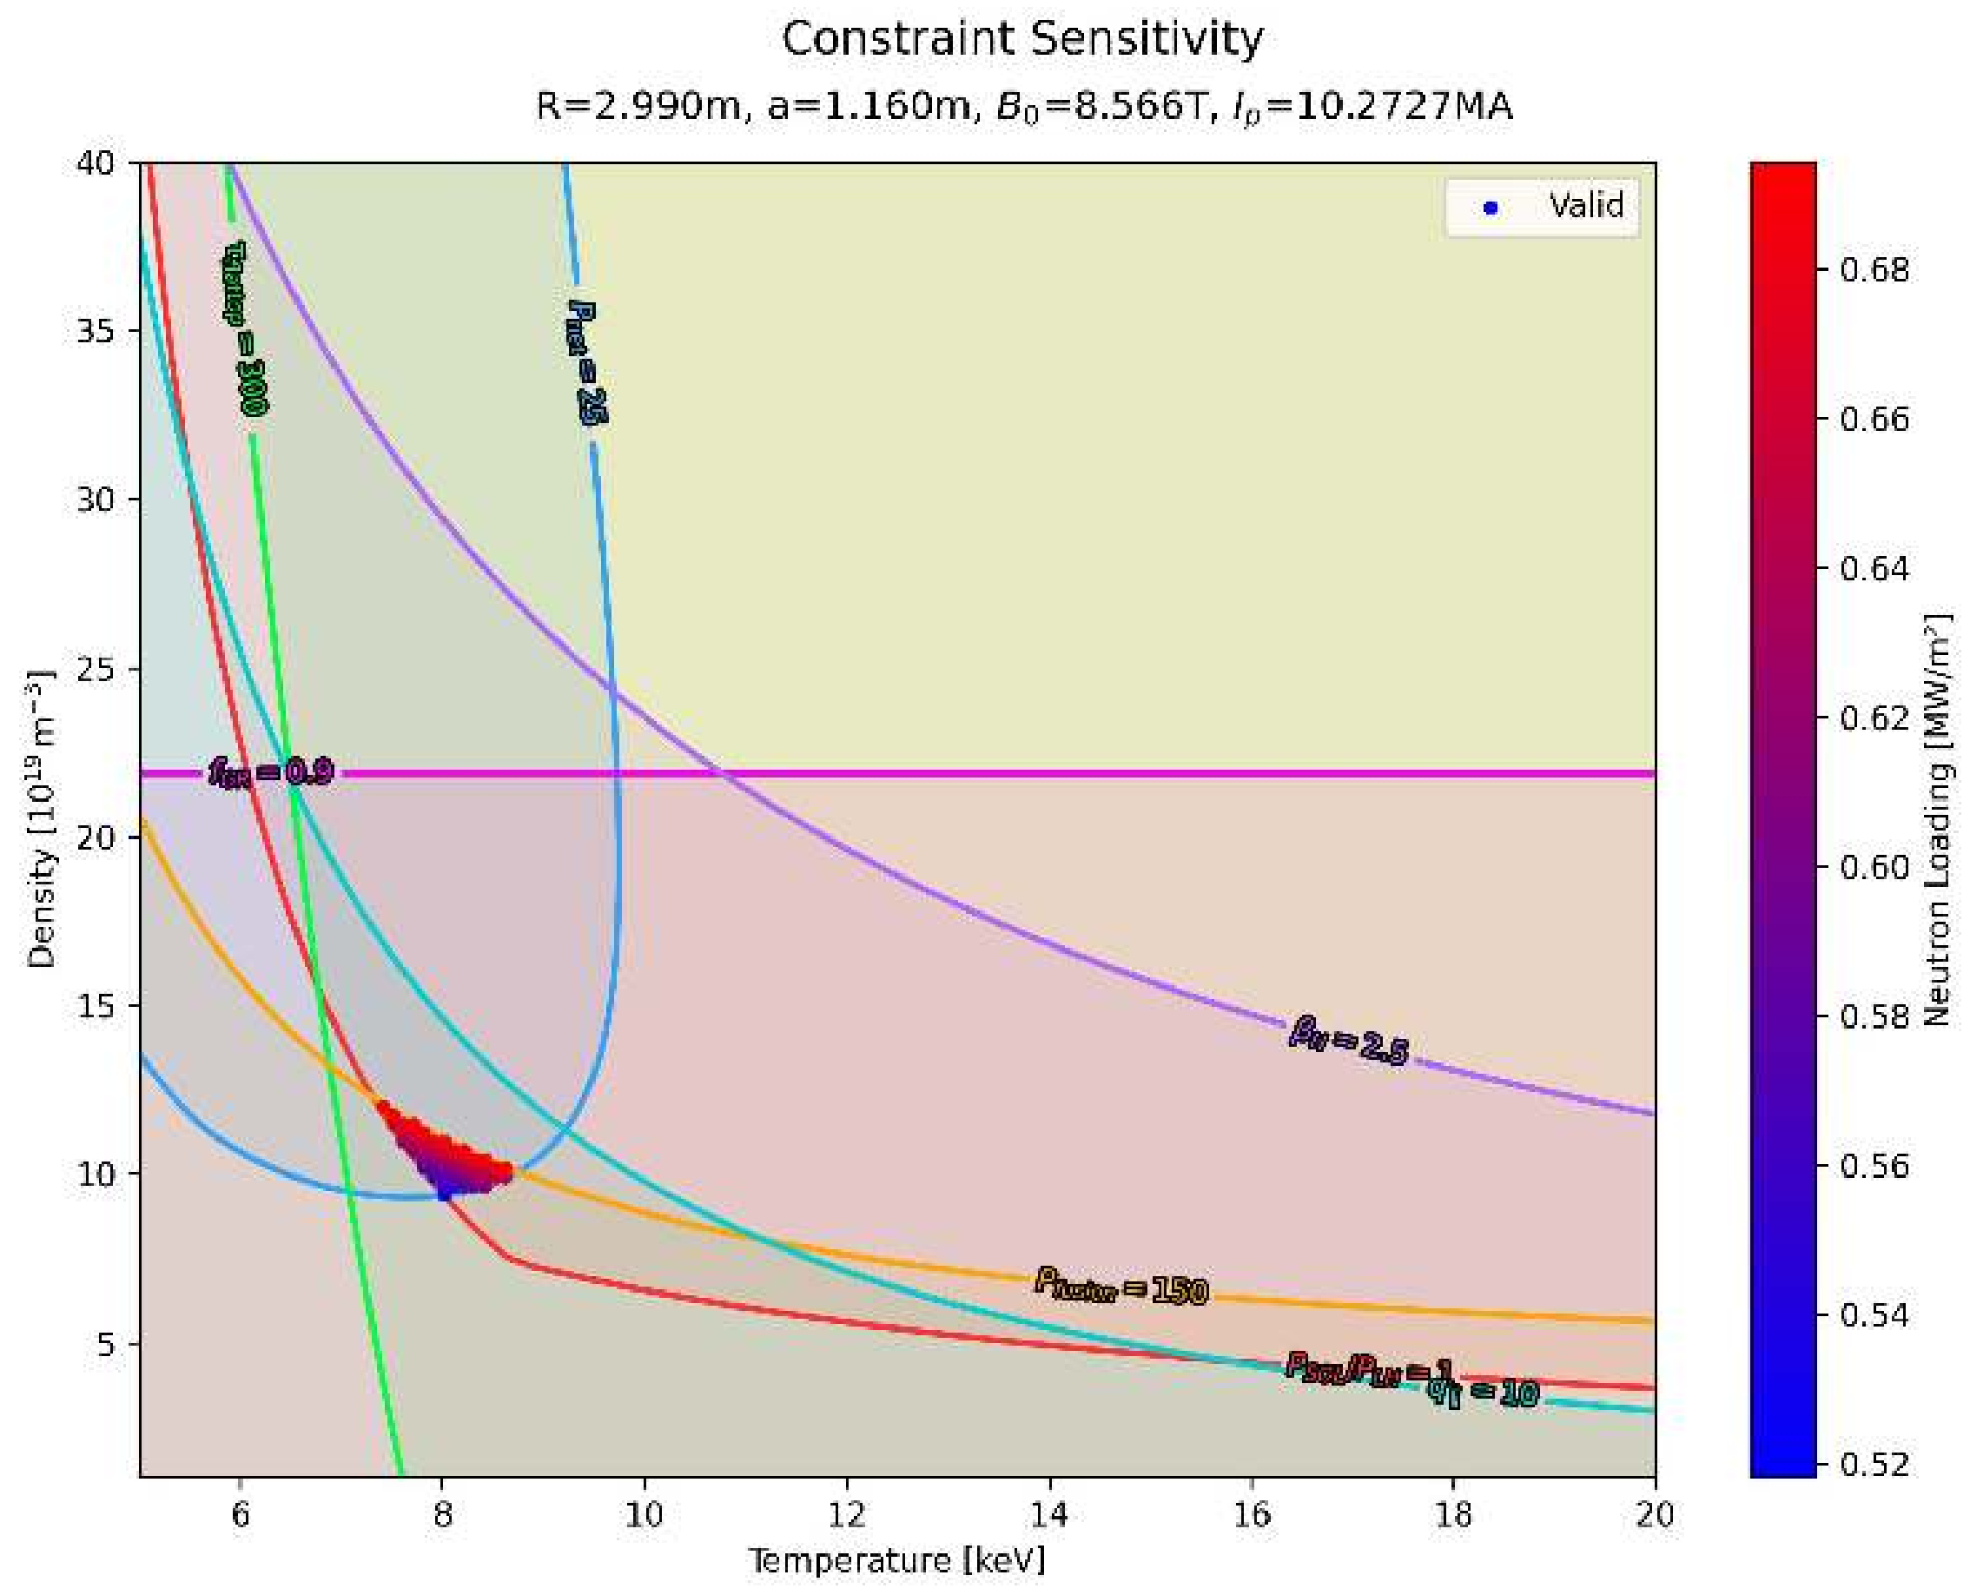

Sensitivity Analysis

If we relax our operating constraints from 30MWe to 25MWe net power and upper constraint thermal limit to 150MW fusion power, our operating region opens significantly, while still maintaining <700kW/m² neutron power flux. This shows that we have some slack at our chosen operating temperature and density, additionally this is with a confinement scalar H set to 1.0, giving further confidence and margin in the plasma performance.

We are also operating well within the Greenwald density limit, the Troyon beta limit, and the constraints on parallel heat flux and maximum flattop duration, so these factors do not directly restrict our operating region. Instead, the primary constraints are net power, the upper thermal plant limit, and the scape-off-layer power required to maintain H-mode confinement. Even though our main operating point is at a lower temperature (8 keV) and density (~8×10^{19}m^{-3}), we can increase these plasma parameters and trade maintenance duration to support higher power operations if needed.

Shaping / Elongation Analysis

Keeping the simulated major parameters from above constant and sweeping over the elongation shows higher elongation is favored for increased net power and Q, at the cost of higher parallel heat flux and lower scrape-off-layer (SOL) width. To stay below the elongation limit for vertical stability, we should not go above an elongation of ~2.1, so we will leave some margin and pick an elongation of ~1.8-1.9 to balance out these competing factors.

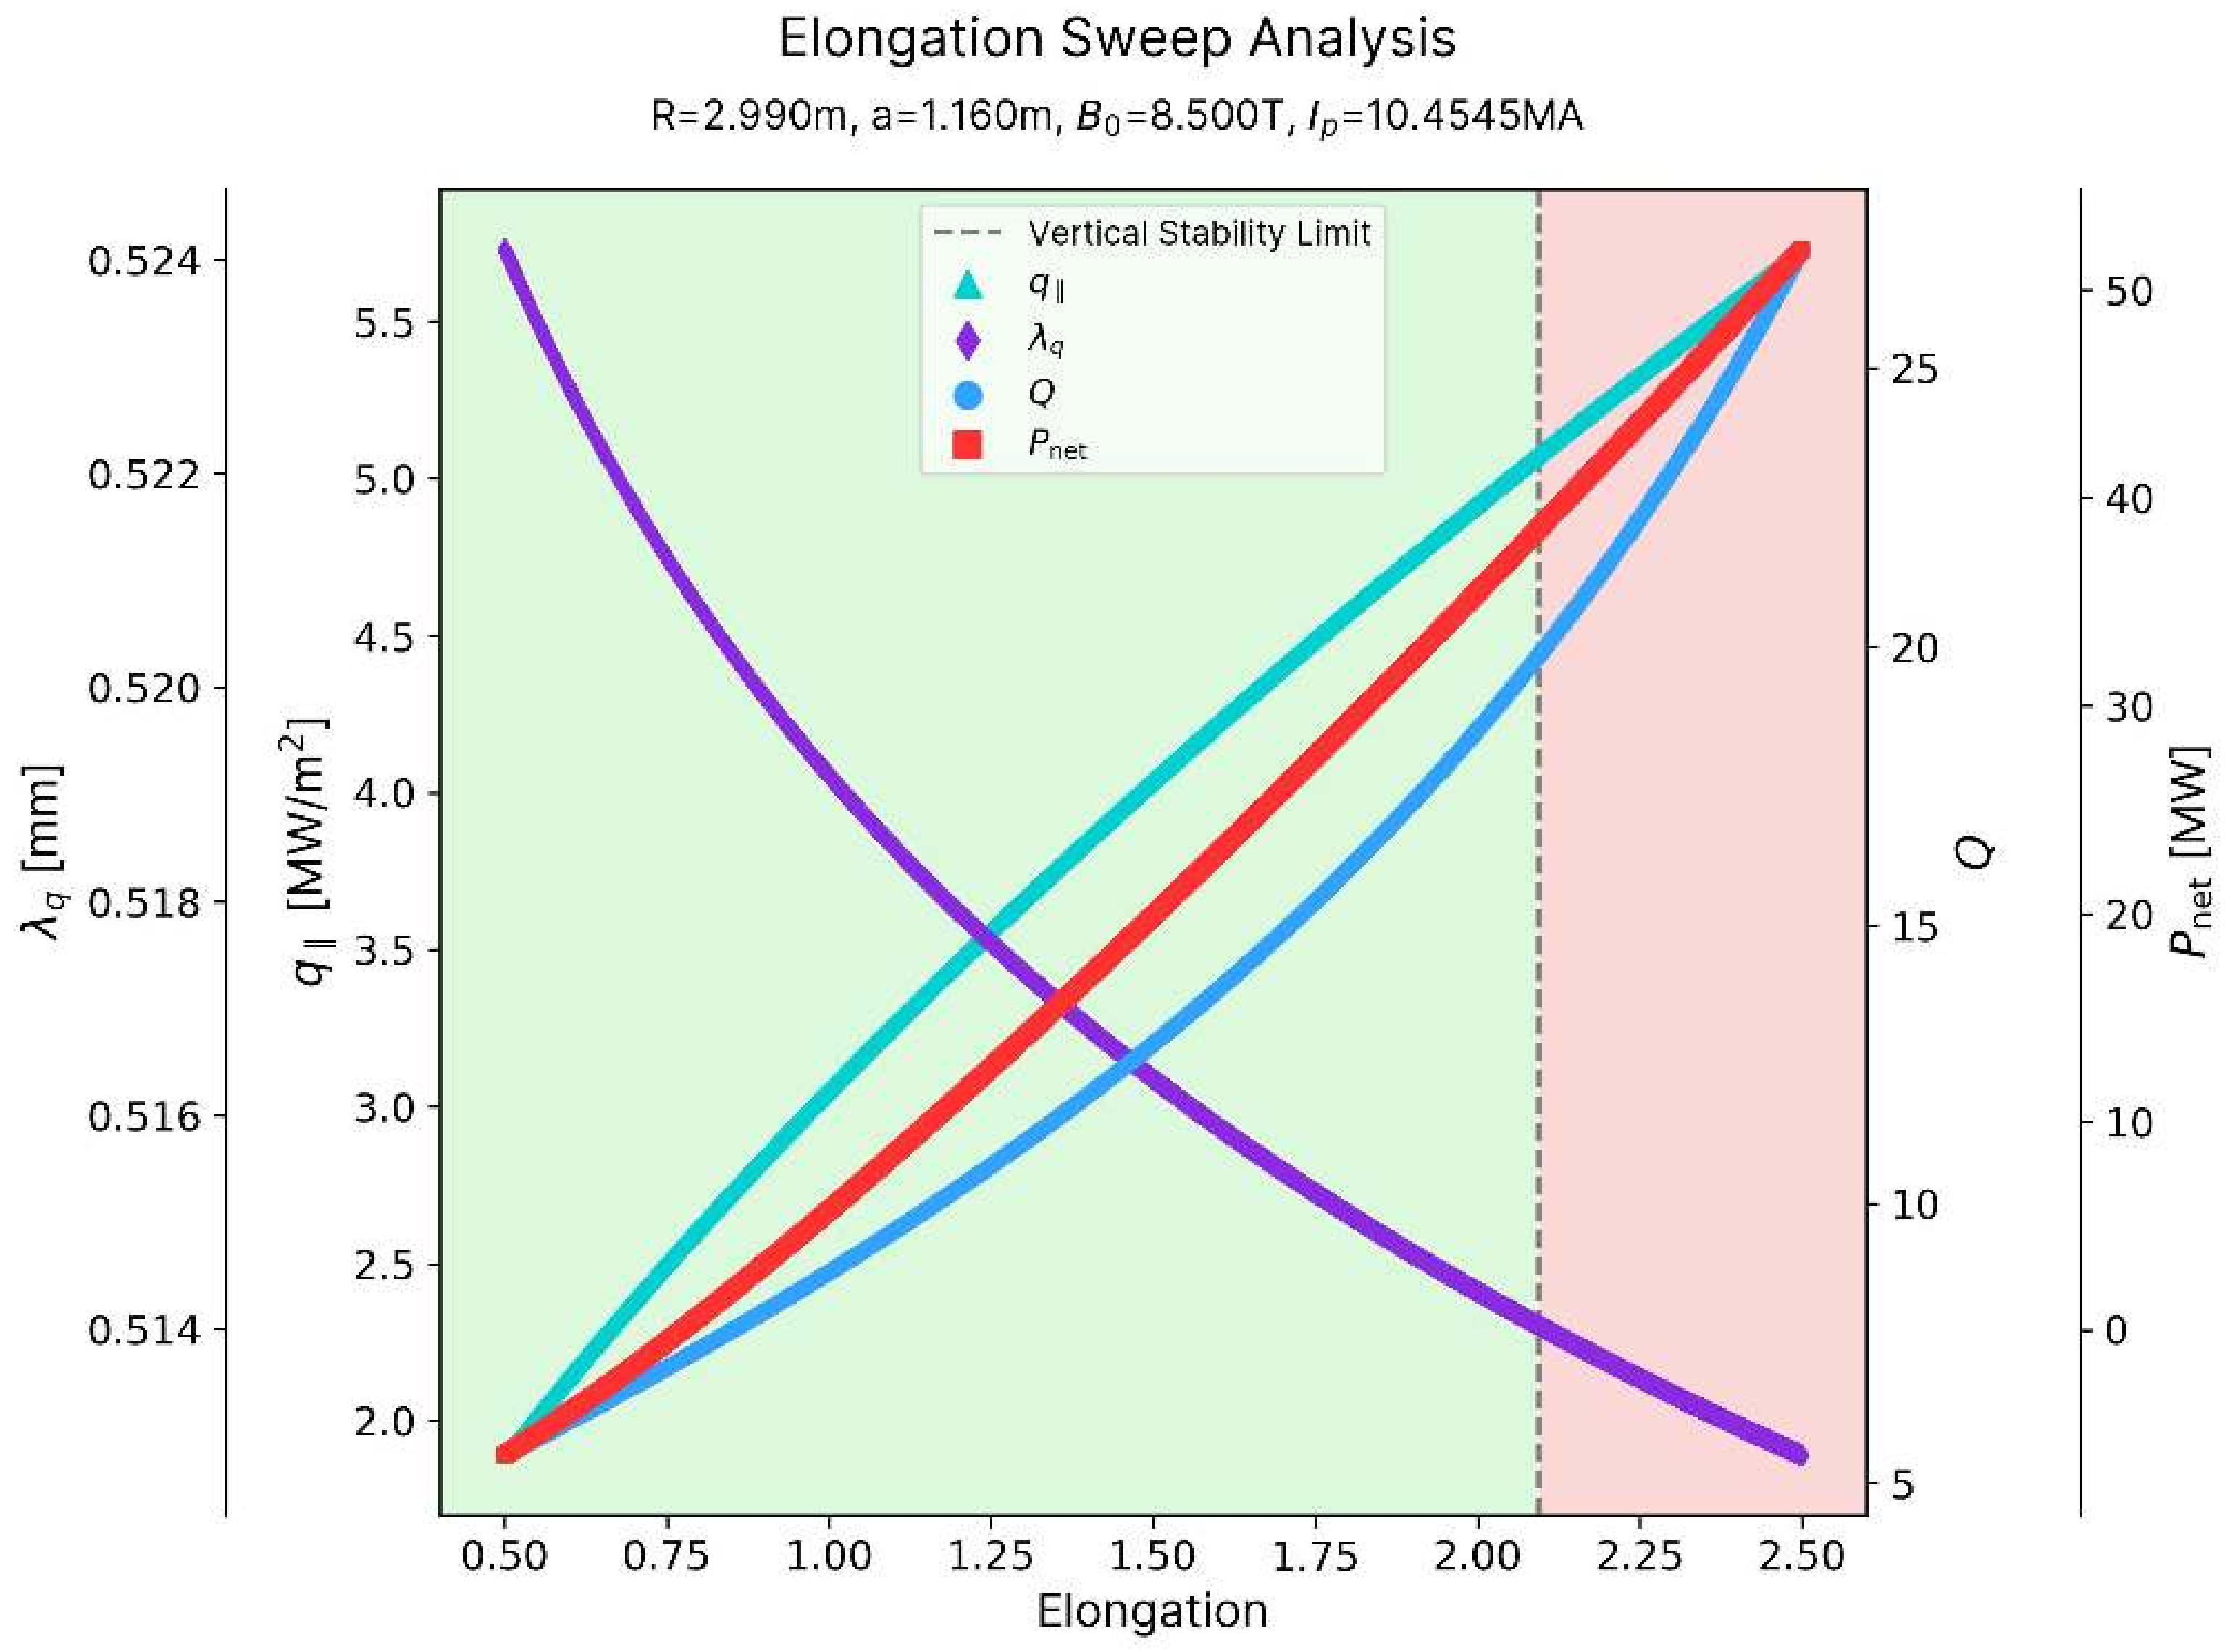

Temp / Density Contours

Based on this rough operating point from our simulations, we can plot the temperature and density profiles for SOL power, parallel heat flux, bootstrap fraction, core radiated power fraction, and heat flux width to further understand the power distribution and identify localized regions of concentrated thermal stress beyond simple uniform power metrics. The valid points are in white and operating region is enclosed with a gold line.

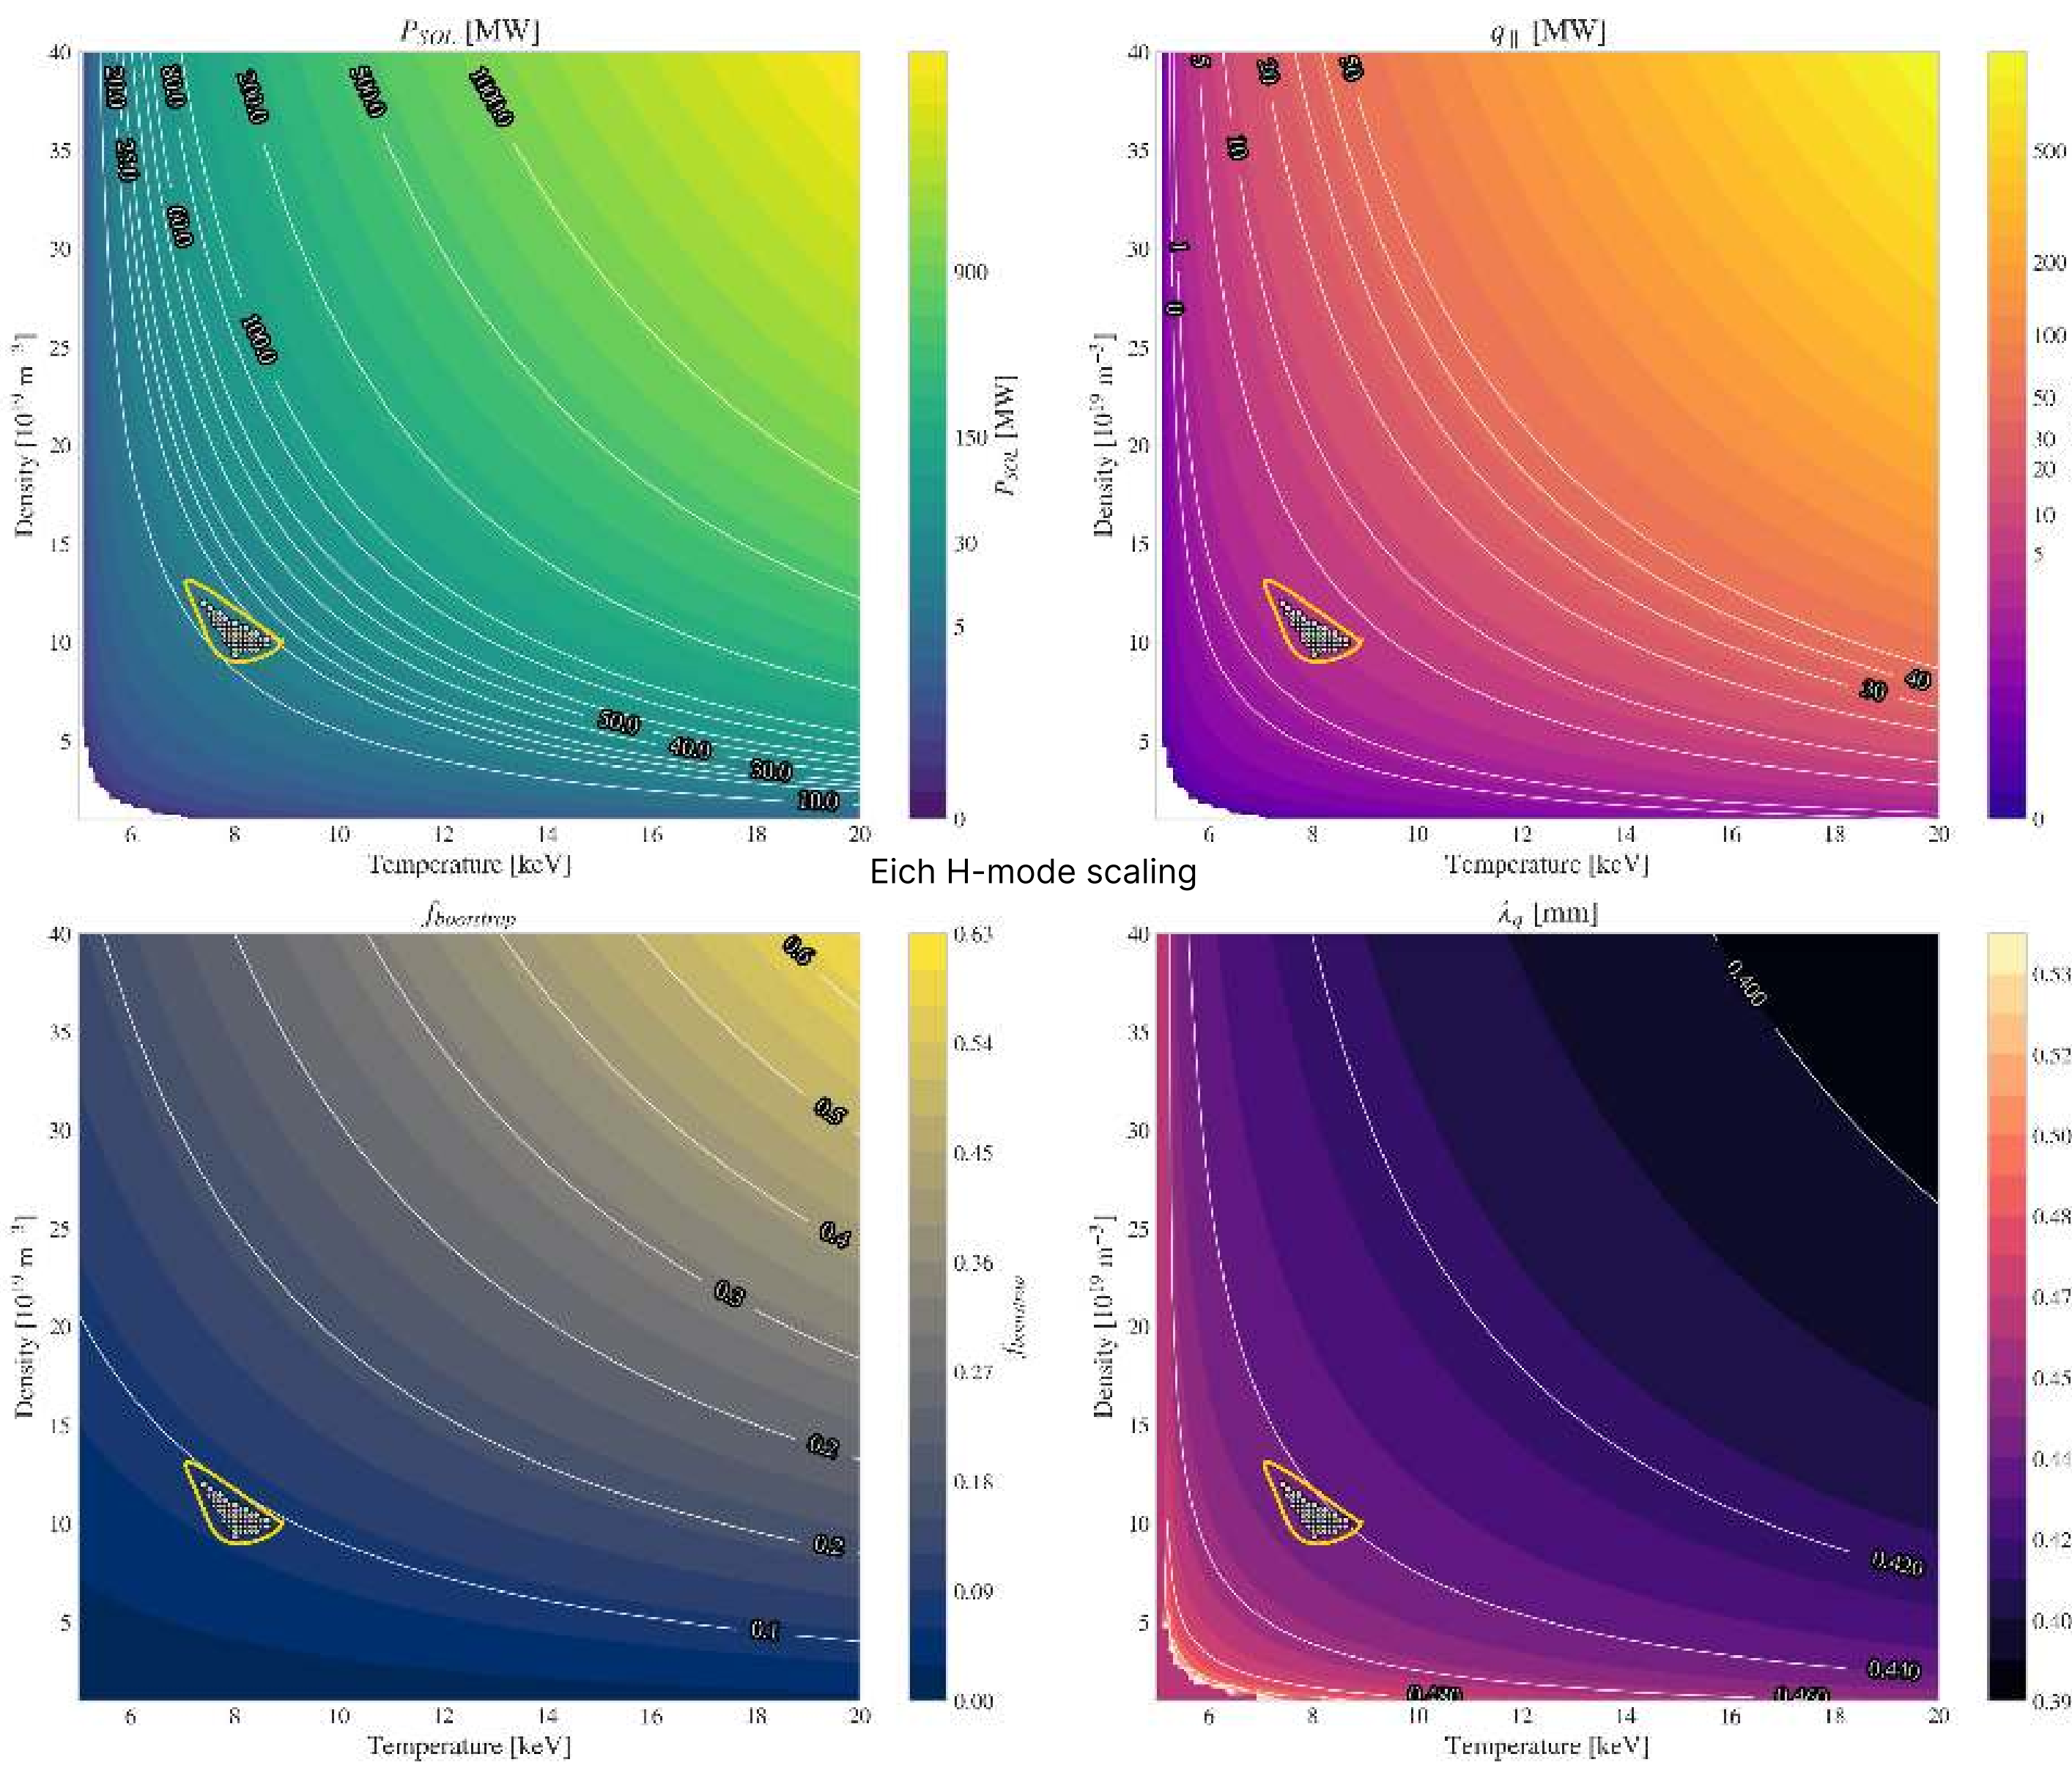

Reactor Design (cont.)

Given the optimization analysis we performed above, we can see if the standard 0D plasma plasma stabilities criteria are satisfied:

Greenwald Fraction

The Greenwald fraction determines the maximum achievable plasma density for stable operation, ensuring the reactor operates safely below density-driven disruptive limits. Operating at or below the Greenwald density maintains sufficient margin against unexpected plasma excursions.

$$n_{20} = \frac{I_p}{\pi a^2} \longrightarrow 24.7\times10^{19}\ \text{m}^3$$

At our target density of $\sim 10\times10^{19}\ \text{m}^3$ we get a Greenwald fraction of ~0.4.

Troyon Beta Limit

The Troyon beta limit defines the maximum stable plasma pressure achievable in a tokamak. Operating near, but below, this limit ensures optimal fusion power density without triggering pressure-driven magnetohydrodynamic (MHD) instabilities. Try to keep below 3.

$$\beta_N \equiv \frac{aB_0}{I_p}\beta_T \longrightarrow 0.94$$

β_T is roughly 1% for our electron density + temp and B field.

Edge Safety Factor

The edge safety factor (q_a) quantifies the stability of the plasma against kink driven instabilities. Maintaining q_a above 2 ensures sufficient safety margin against disruptive kink instabilities, promoting stable, reliable reactor operation.

$$q_a = \frac{5a^2B_0}{R_0 I_p}\left(\frac{1+\kappa^2}{2}\right) \longrightarrow 3.88$$

Elongation is assumed to 1.8 -- κ

This shows we are well within operating limits for plasma stability. Next, we can begin modeling our magnetic equilibrium:

Temp and Density Profiles

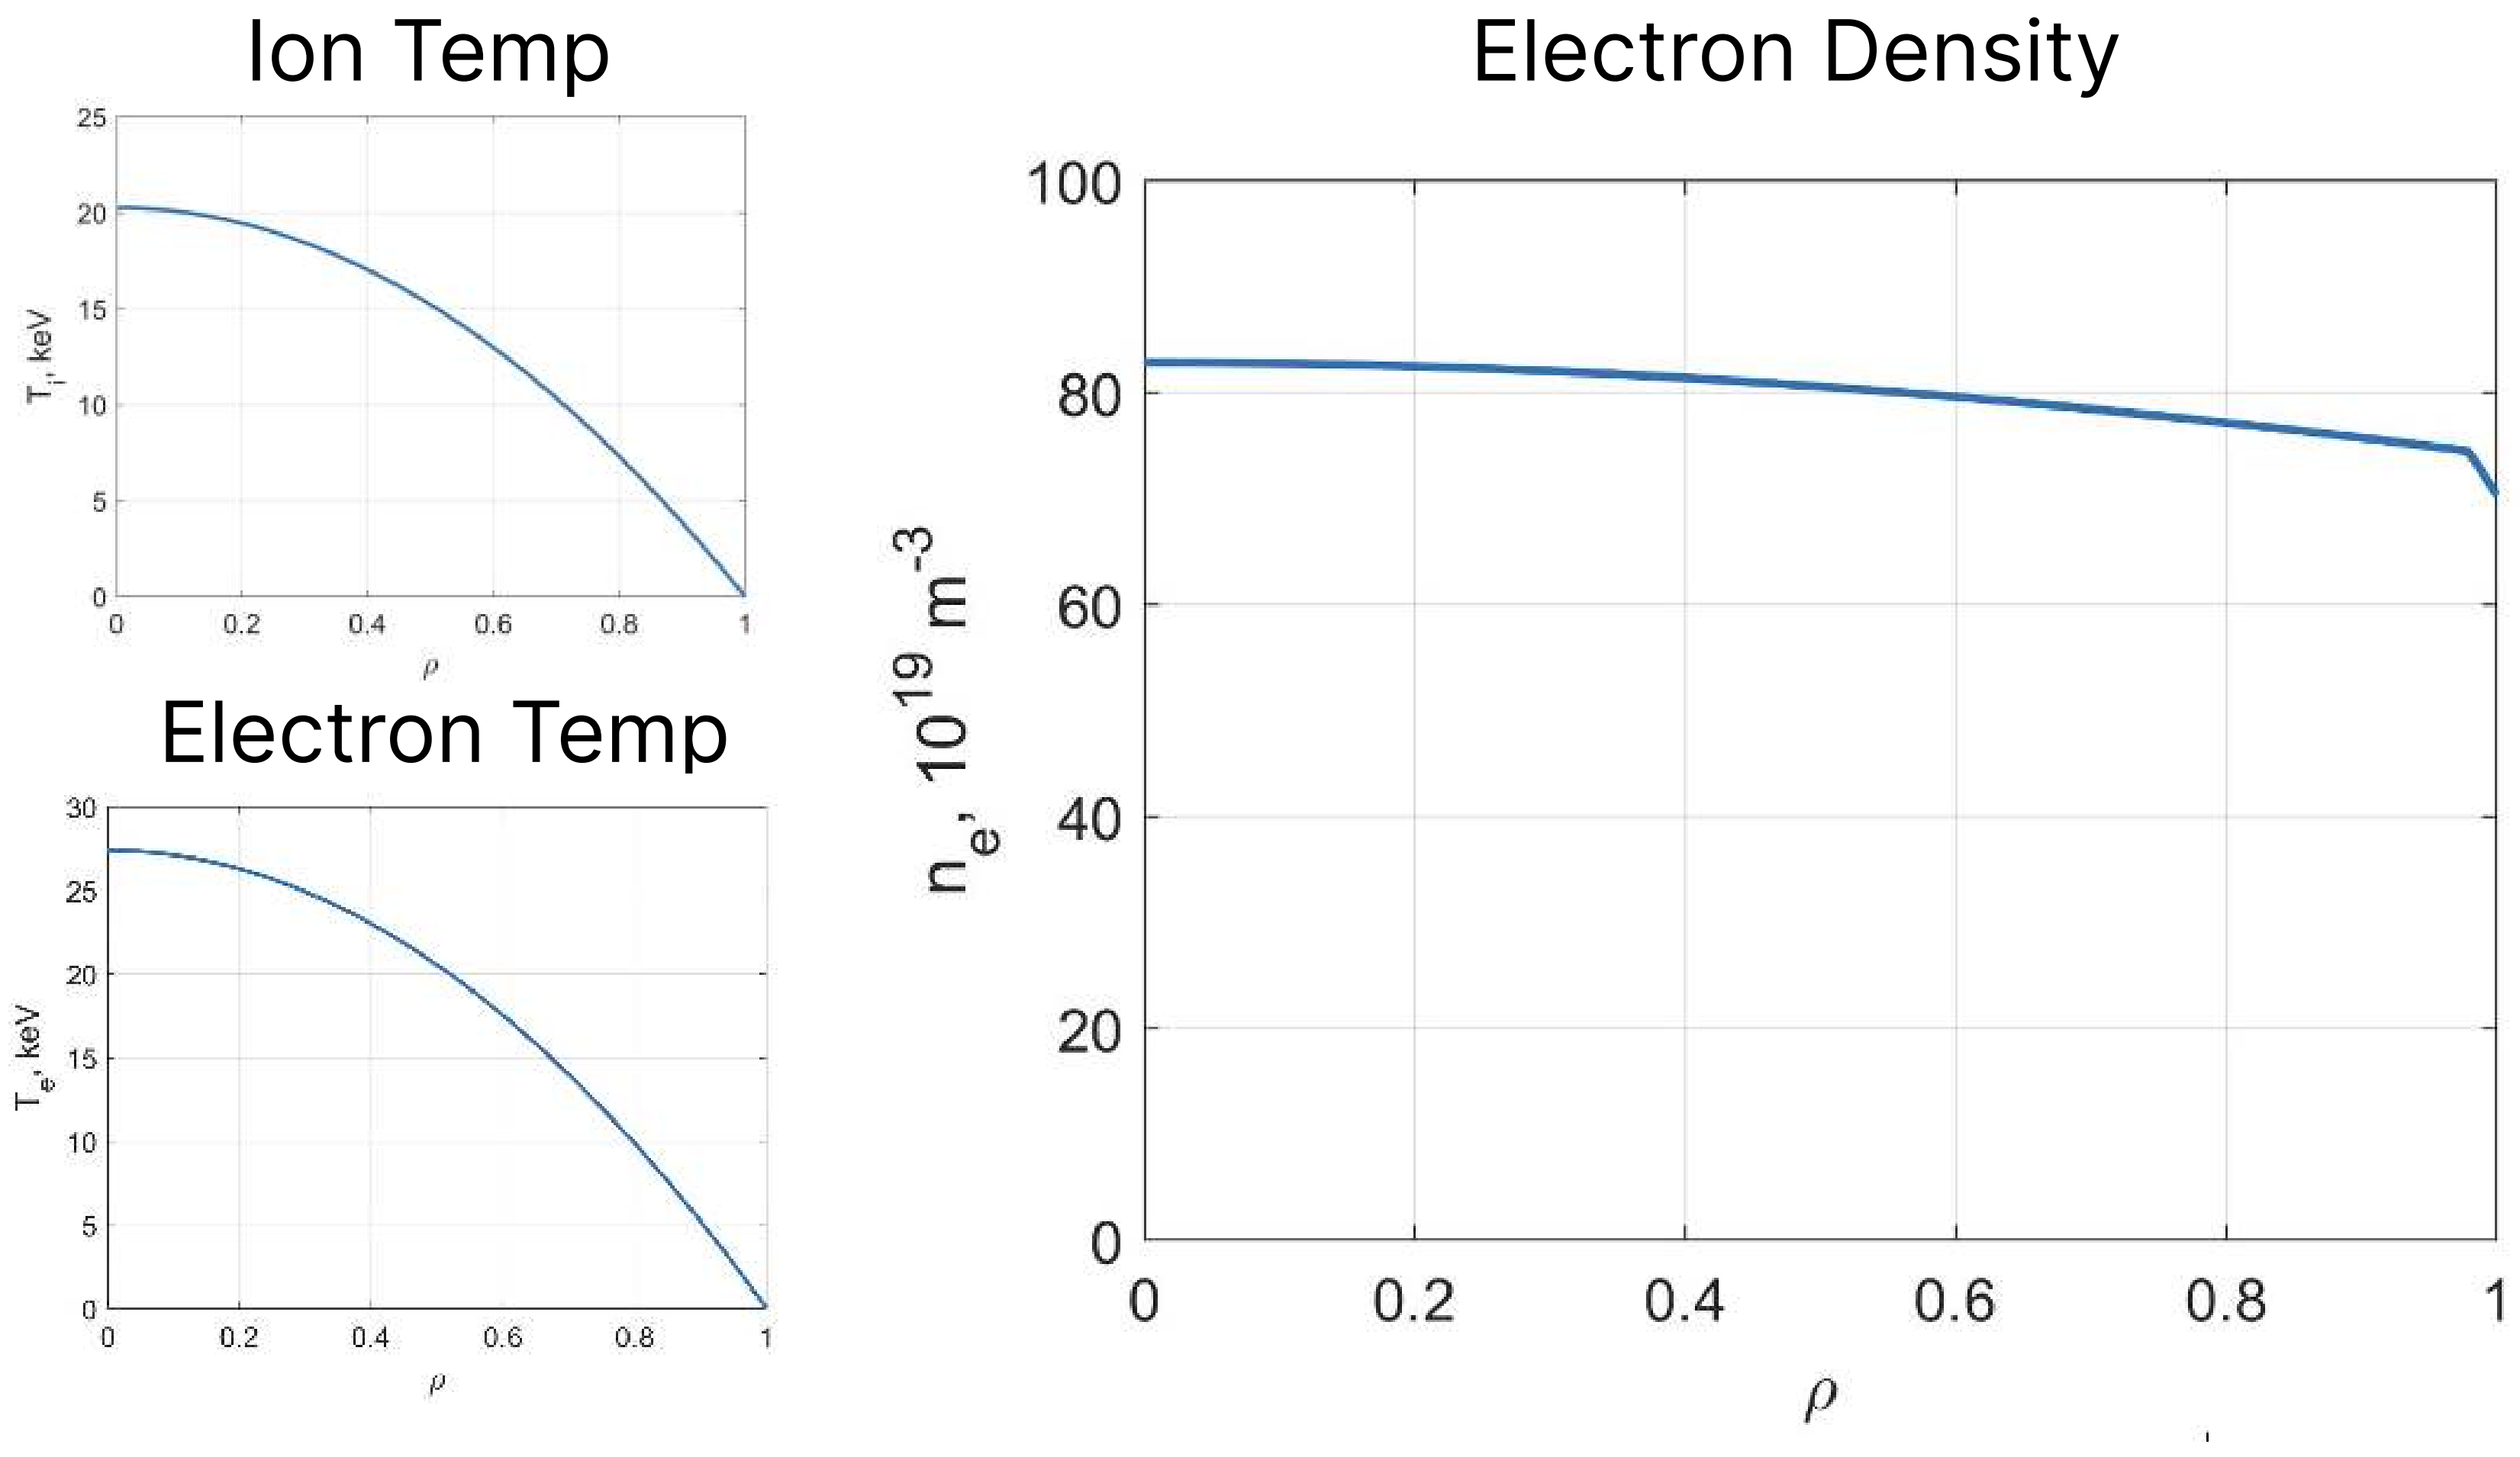
Ion Temp
Electron Temp
Electron Density

These profiles are plotted with ρ on the x-axis, which is a normalized radial coordinate ranging from 0 at the plasma center to 1 at the plasma edge, assuming a parabolic shaping function for the temperature distribution.

Initial Tokamaker Reconstruction

We initially started with Tokamaker to construct a basic coil layout with a simple first-wall and parallel to last closed flux surface vacuum vessel configuration. This preliminary reconstruction allowed us to visualize magnetic flux contours, assess coil placement, and determine approximate coil currents required for desired plasma shaping. The distribution of poloidal field coil currents shown here illustrates the rough balance necessary to achieve vertical stability and proper shaping of the plasma boundary. These initial parameters provided a foundation for iterative refinement using detailed modeling tools such as ACCOME, MCNP, and COMSOL to optimize the reactor design.

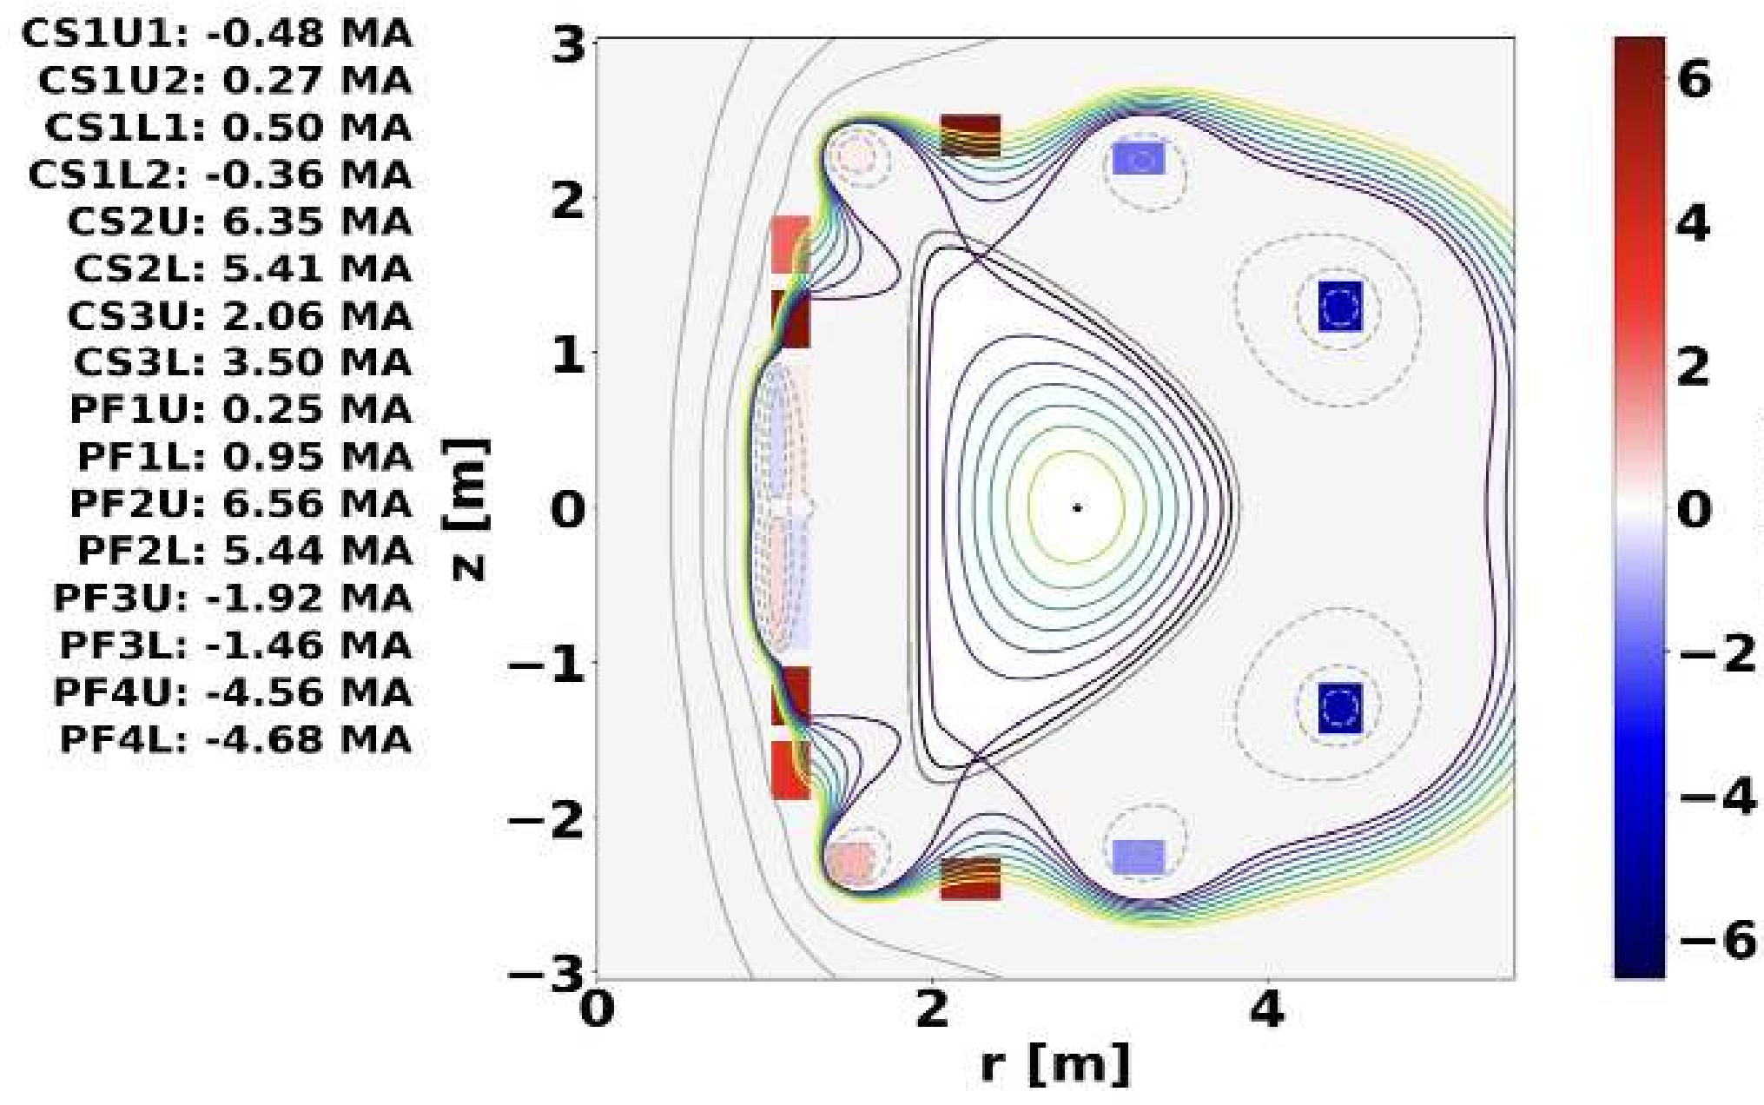

Source: https://github.com/OpenFUSIONToolkit/OpenFUSIONToolkit

Plasma Heating on First Wall

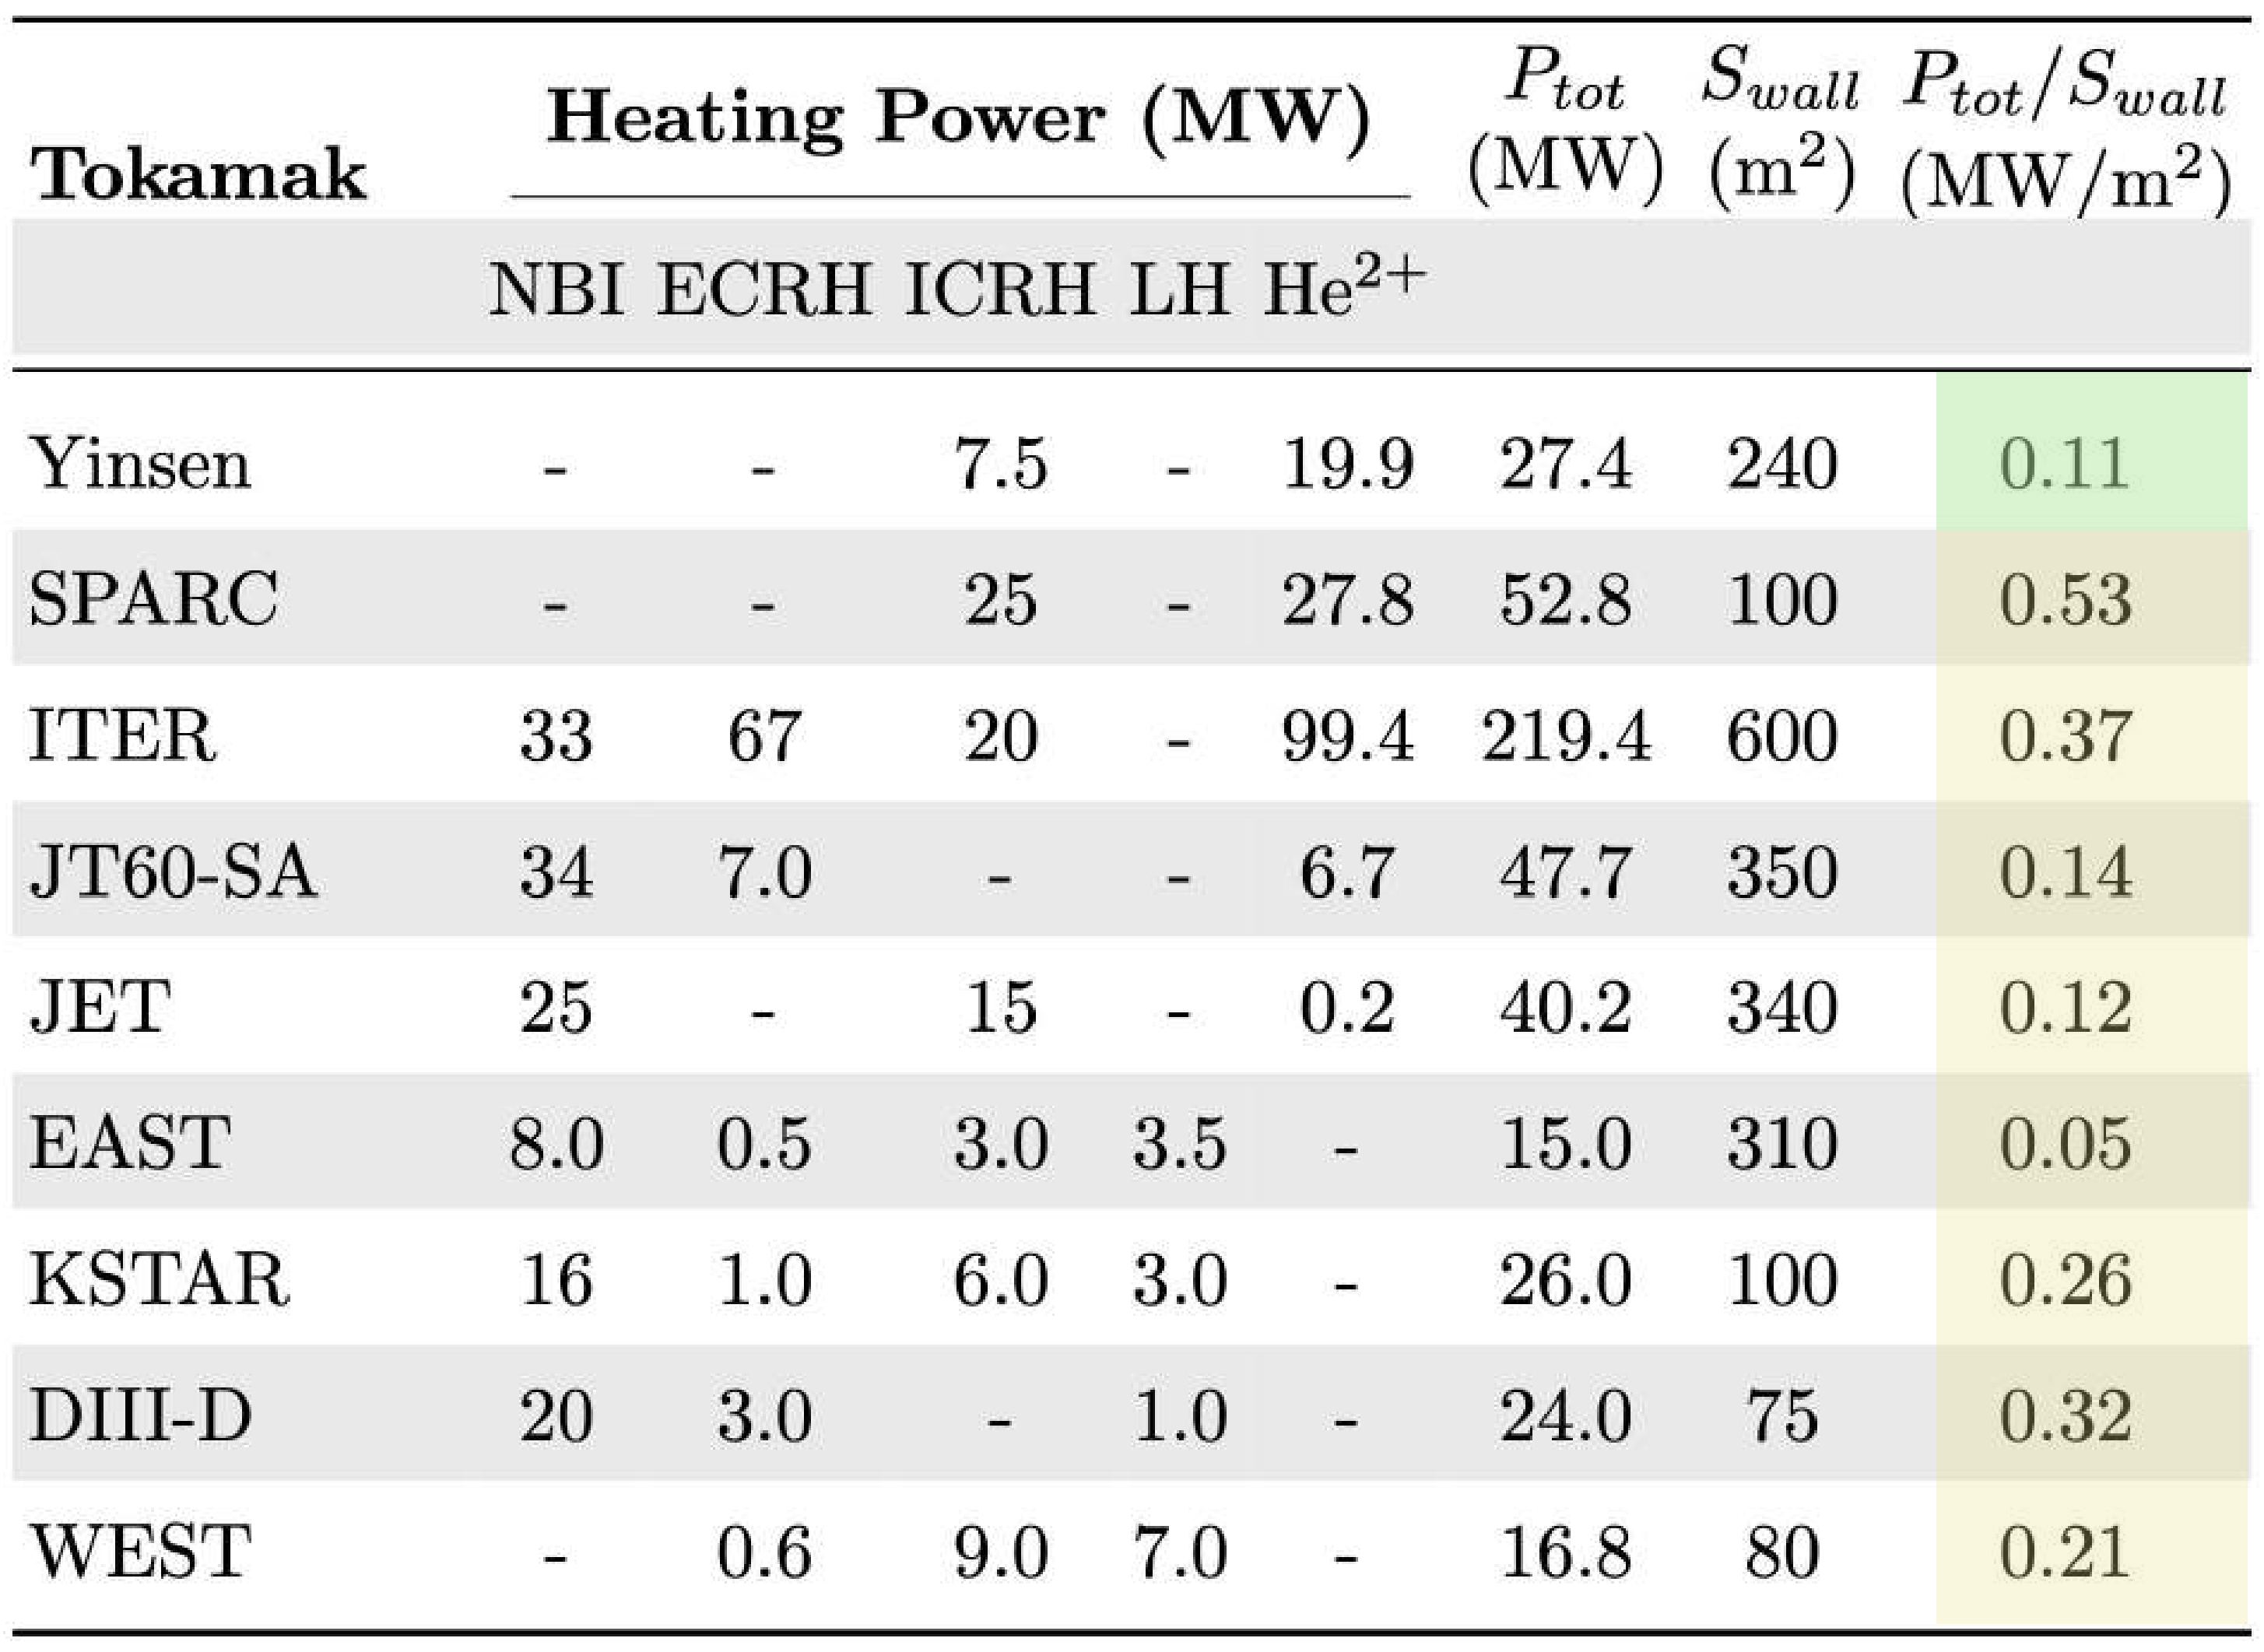

Tokamak	Heating Power (MW)					P_{tot} (MW)	S_{wall} (m^2)	P_{tot}/S_{wall} (MW/m^2)
	NBI	ECRH	ICRH	LH	He^{2+}			
Yinsen	-	-	7.5	-	19.9	27.4	240	0.11
SPARC	-	-	25	-	27.8	52.8	100	0.53
ITER	33	67	20	-	99.4	219.4	600	0.37
JT60-SA	34	7.0	-	-	6.7	47.7	350	0.14
JET	25	-	15	-	0.2	40.2	340	0.12
EAST	8.0	0.5	3.0	3.5	-	15.0	310	0.05
KSTAR	16	1.0	6.0	3.0	-	26.0	100	0.26
DIII-D	20	3.0	-	1.0	-	24.0	75	0.32
WEST	-	0.6	9.0	7.0	-	16.8	80	0.21

Total plasma heating per first wall surface area for reactor grade tokamaks

Triple Product vs Pulse Comparsion

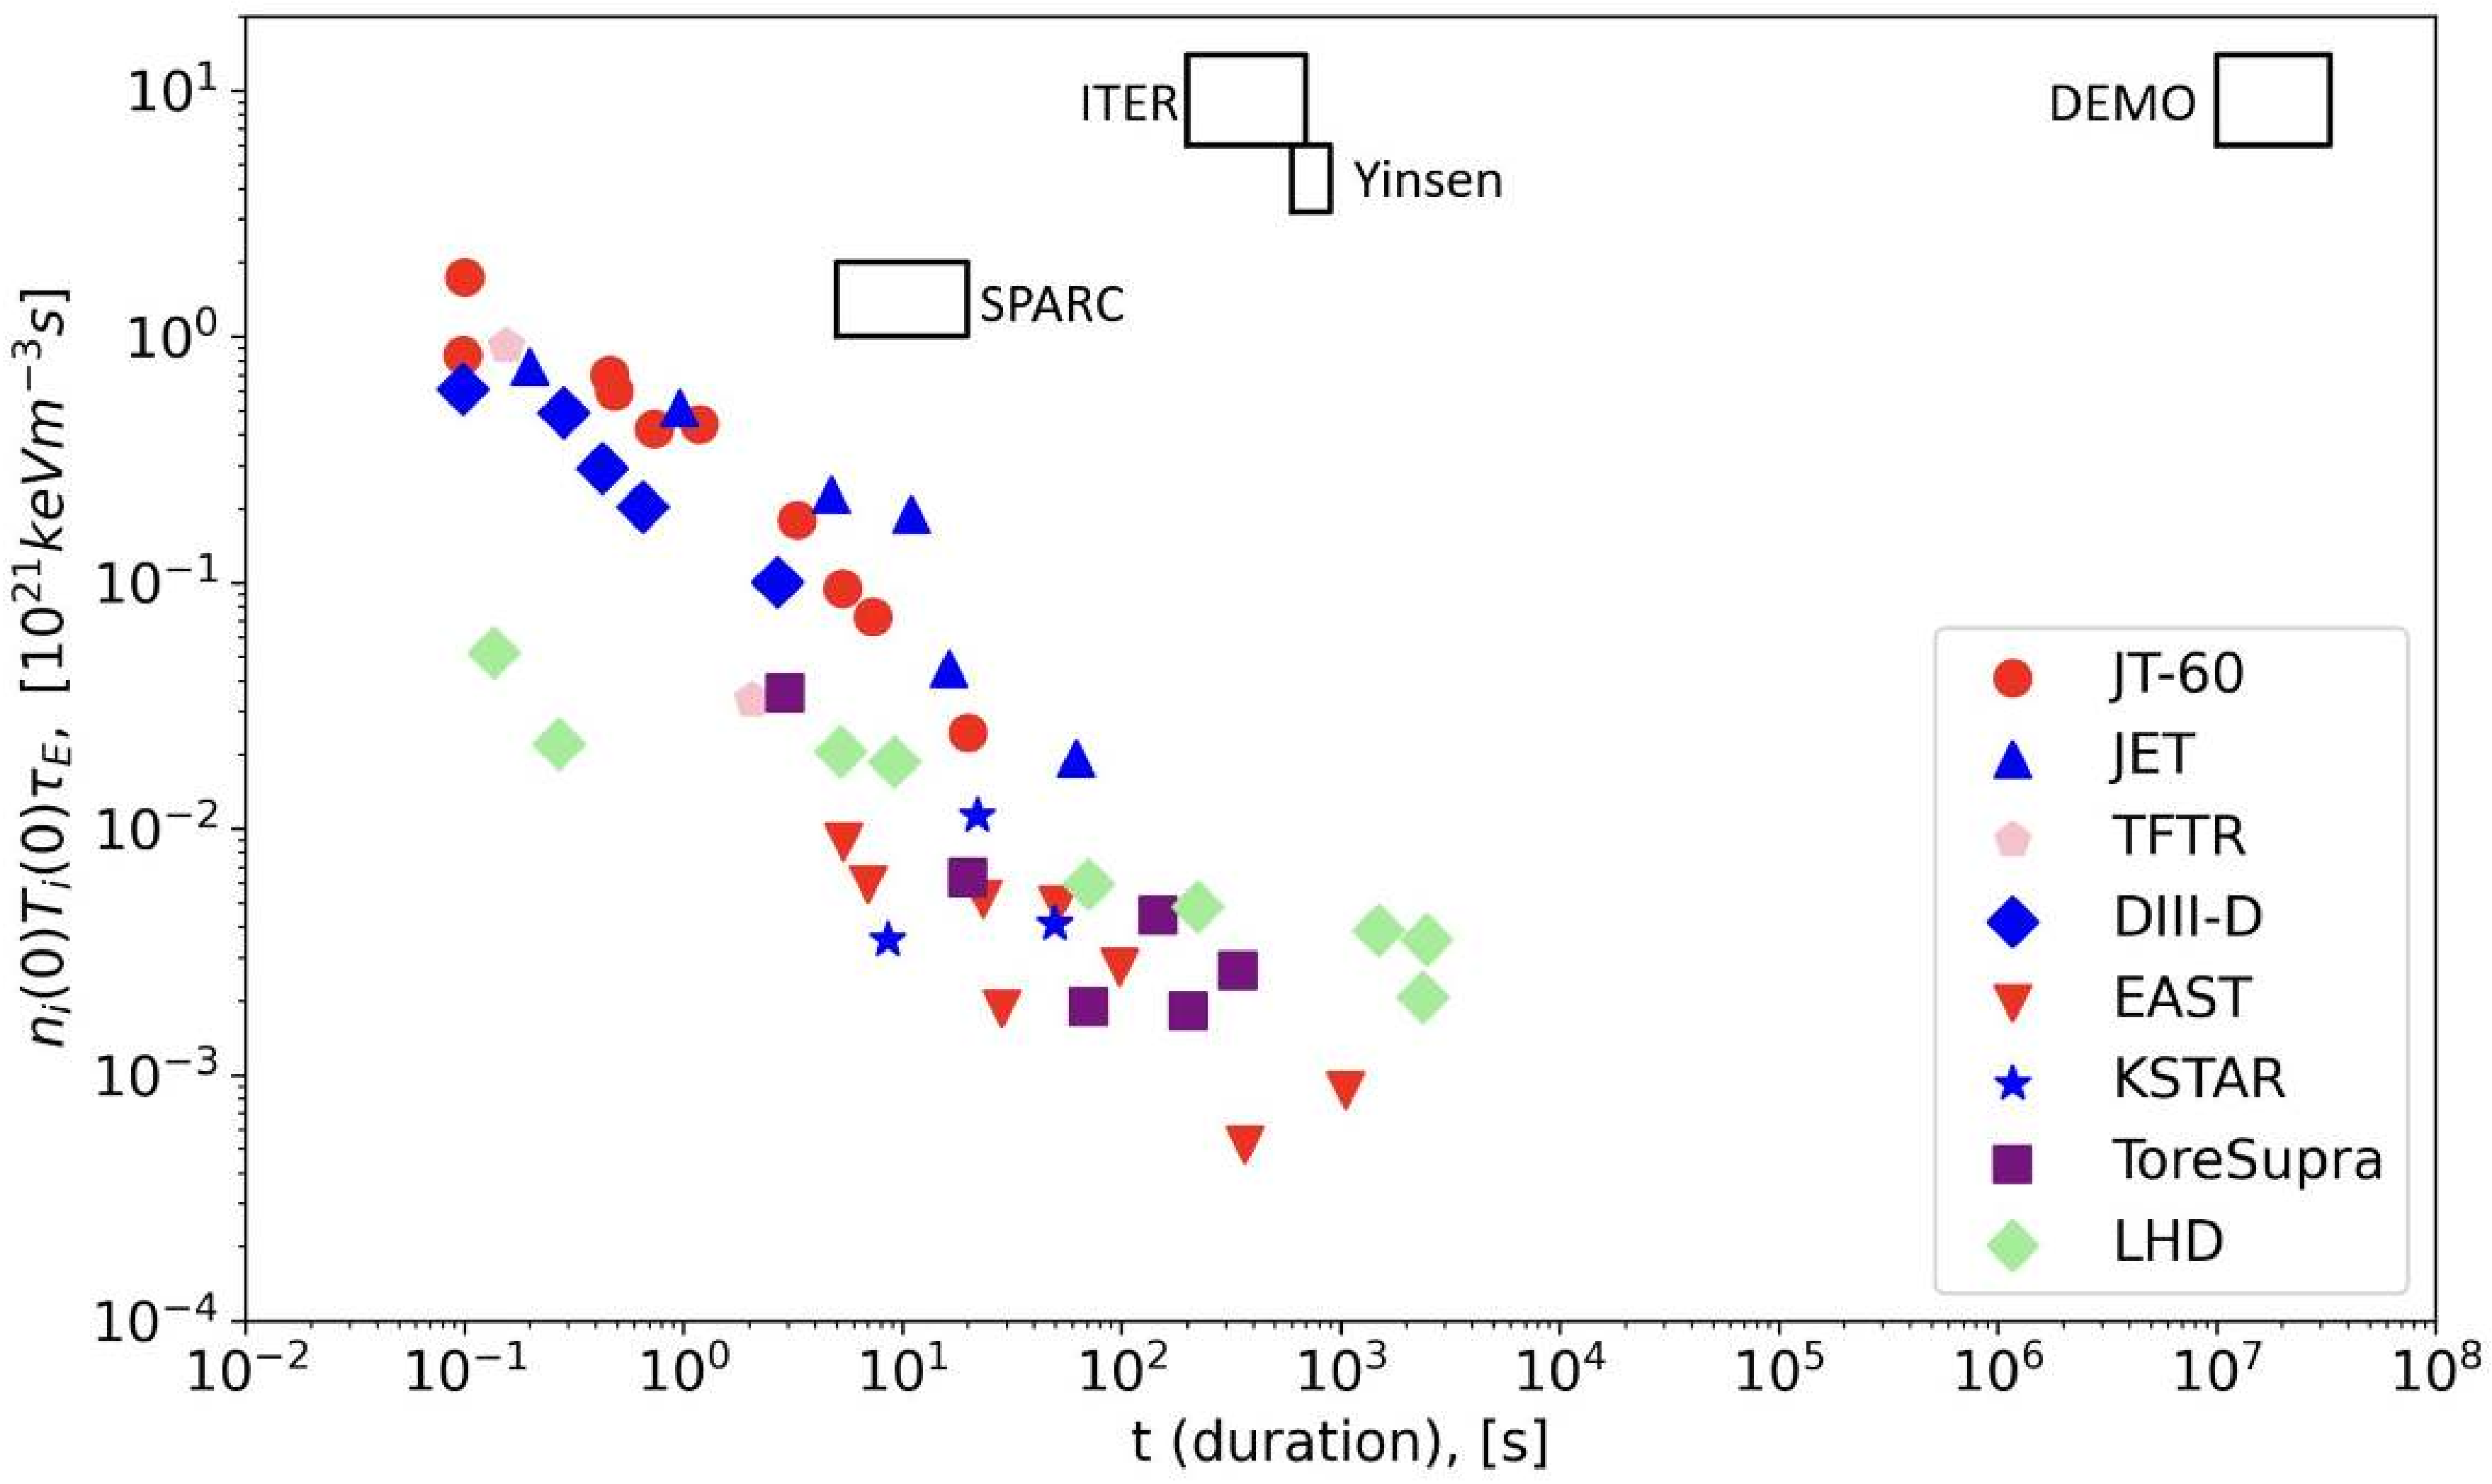

Comparison of triple product vs. pulse duration for existing tokamaks and reactor-grade projects.

Yinsen targets triple product values in between ITER and SPARC tokamaks as shown above (right). Pulse duration is similar to ITER. A distinguishing feature of Yinsen is a lower ratio of total plasma heating (auxiliary heating + alpha-particles) per first wall surface area. Values for Yinsen, ITER, SPARC, and others are given in above left. As one can see, Yinsen has 3 (5) times lower power densities than ITER (SPARC). Therefore, power exhaust management is expected to be solvable with SPARC and ITER-like technology that includes full or partial detachment, closed divertors, and long-leg divertors. These are in line with the 0D approximations we performed in POPCON.

It is important to note that this is a first-pass design of a tokamak fusion reactor for a ship. We do not claim to have every detail fully resolved but aim to demonstrate a clear path toward a viable solution through further analysis and iteration. Our approach ensures that key parameters remain within reasonable margins, with ongoing refinements focused on tuning plasma equilibria as well as optimizing cable current density and conducting further divertor analyses to enhance heat exhaust management and overall reactor performance

2D Plasma Simulations

Advanced MHD Equilibria

To further simulate the plasma confinement and stability, we use a MagnetoHydroDynamic (MHD) simulator based on the Grad-Shafranov equation with a more advanced divertor and better coil placement than the basic simulator in Tokamaker. Since this device will operate at temperatures exceeding 100 million Kelvin, the fuel ionizes and becomes a high-energy plasma that behaves like a magnetic fluid. The ions in this fluid then move via the magnetic fields induced inside the reactor. Understanding the stability and evolution of this plasma is a critical factor in ensuring proper confinement and sustained fusion reactions.

Coil Placement

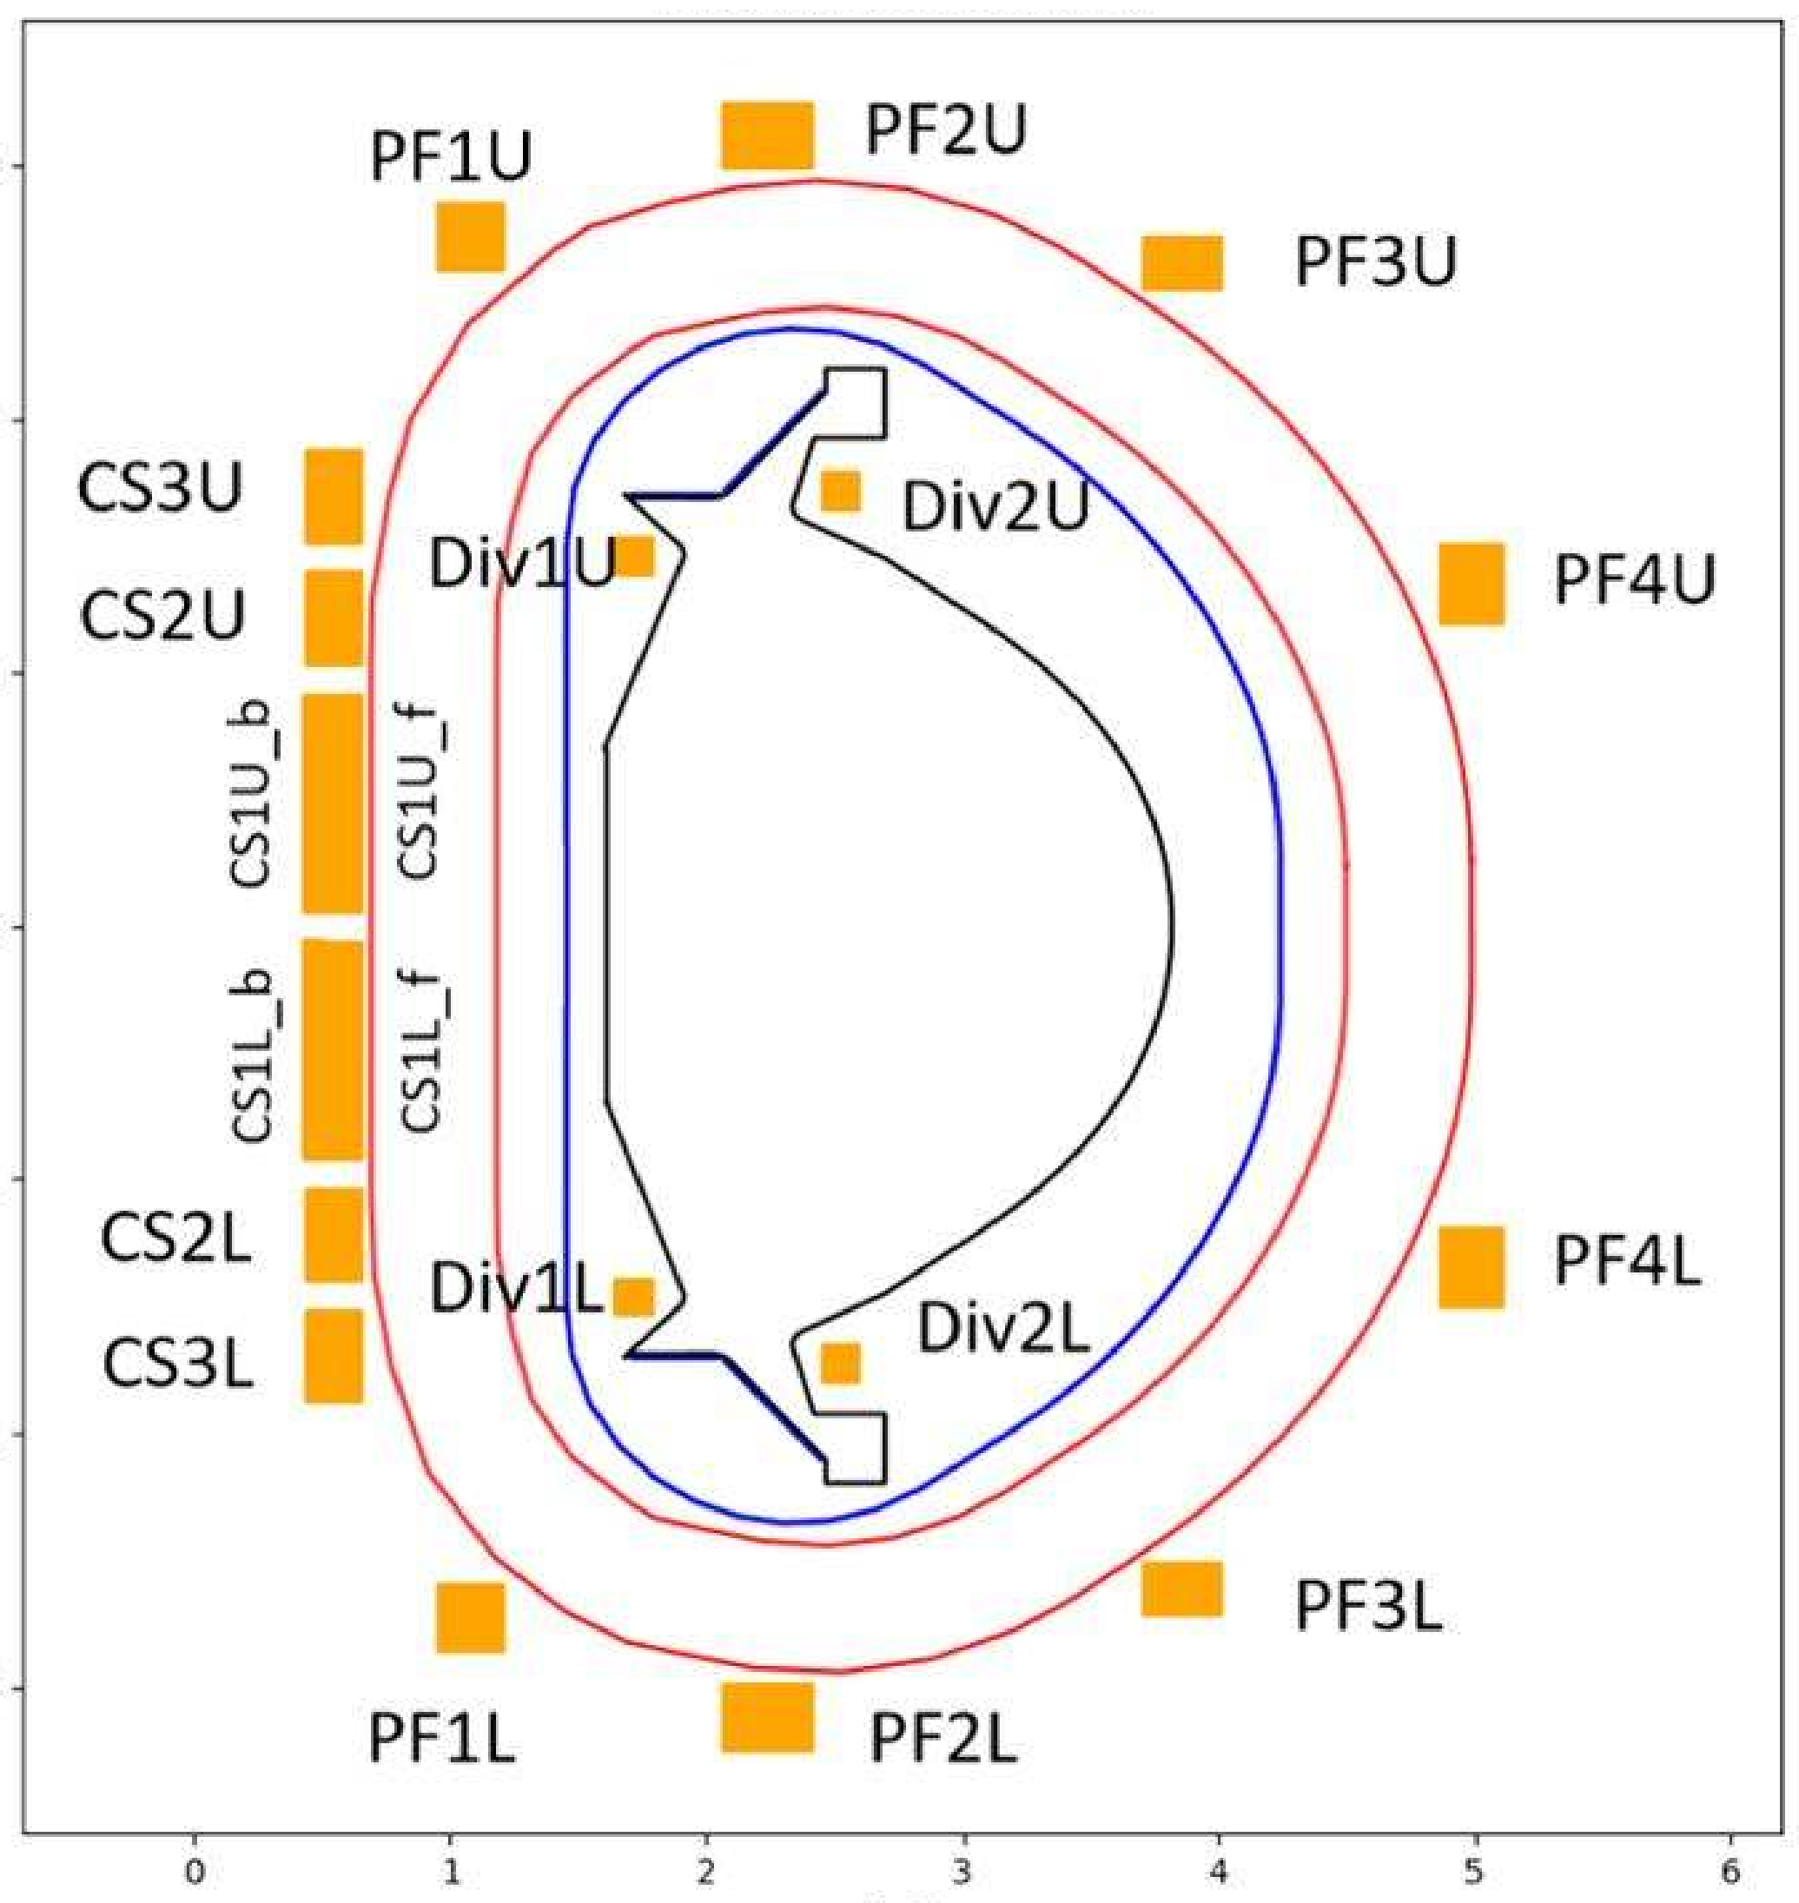

Yinsen's vertical cross-section: toroidal field coil - red line, vacuum vessel - blue line, first wall/limiter – black line, PF1-4U/L – poloidal field coils, CS1-3U/L – sections of central solenoid, Div1-2U/L – divertor coils

The poloidal magnetic system of Yinsen consists of 8 poloidal field coils and a central solenoid divided into 6 sections shown to the left. A stainless steel vacuum vessel that replicates the target plasma shape was used (resistivity of stainless steel is 80×10⁻⁸ Ohm·m). The thickness of the vacuum chamber was set to 1 cm. The discharge pulse to initiate startup in the device is shown to the right. This is achieved by sending varying currents through the CS and PF coils (no active feedback control is shown, this is purely passive stability). The discharge duration is set to be ~10 minutes with ~1 minute for the discharge restart, primarily a function of central solenoid flux.

Startup Scenario

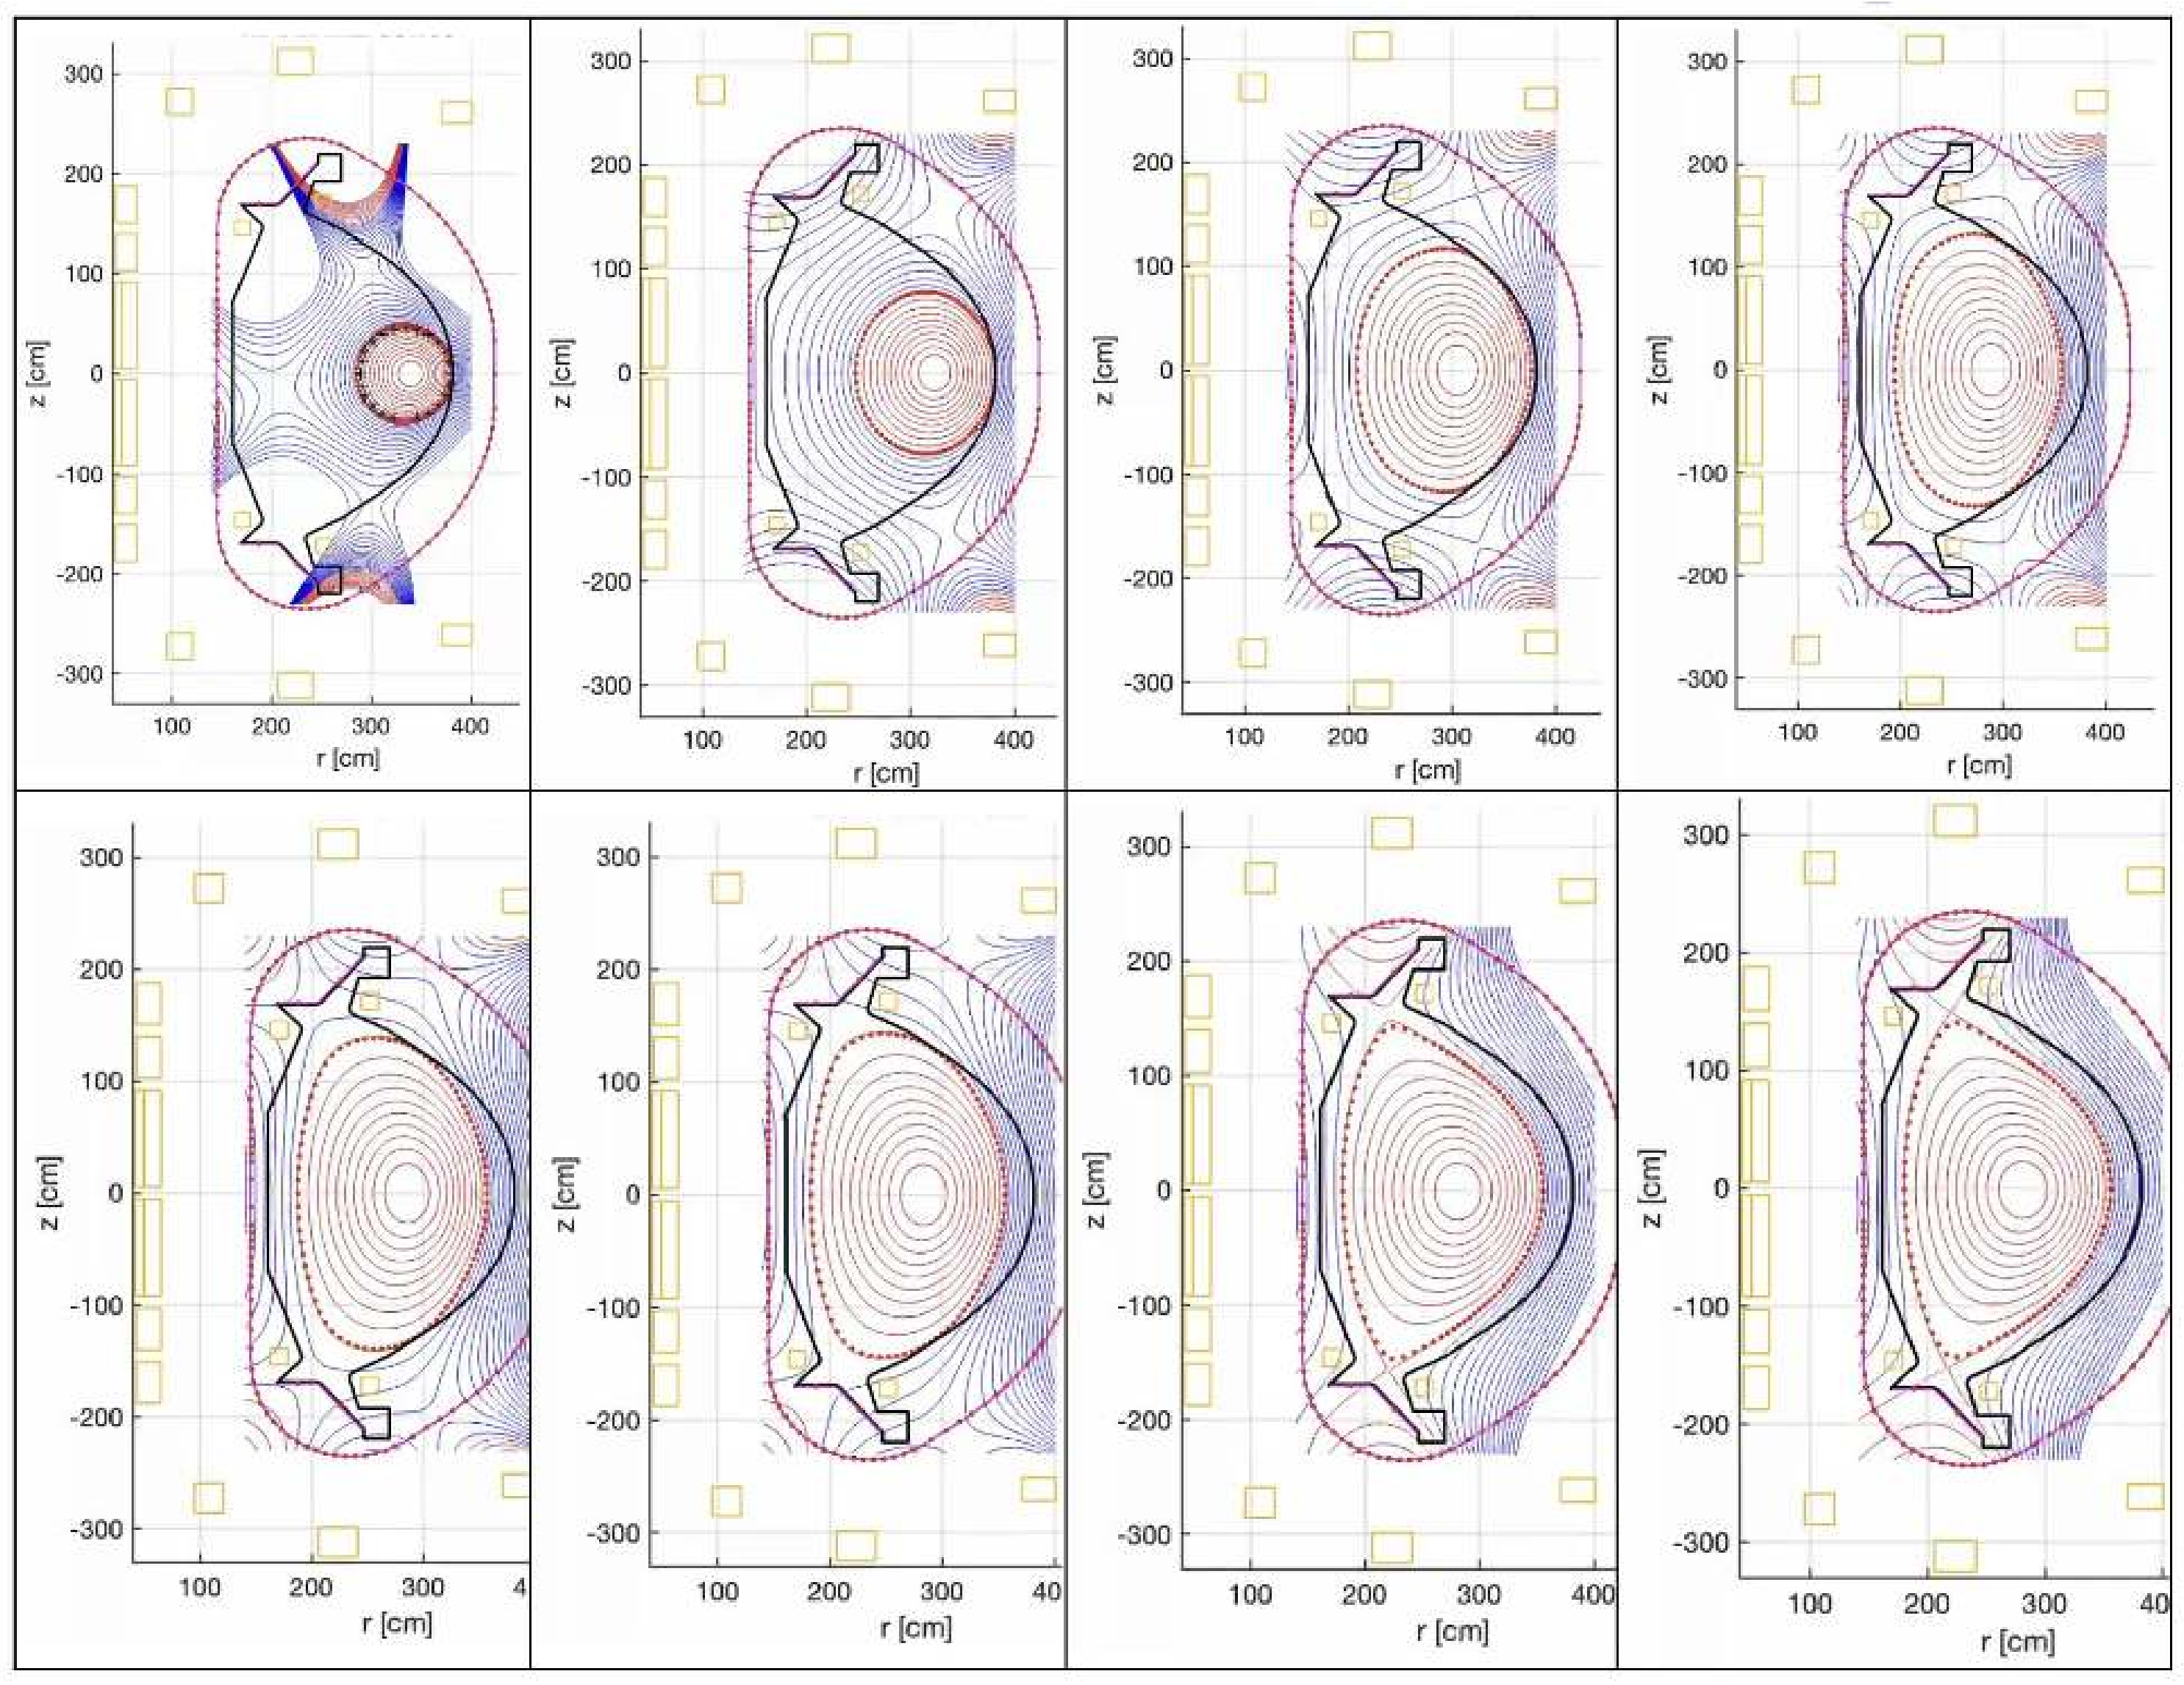

Magnetic equilibrium evolution during discharge: (top left) 50 ms, 199 ms, 599 ms, 1099 ms, 1499 ms, 1799 ms, 2399 ms, 2699 ms (bottom right)

In the startup scenario, plasma discharge starts on the low field side (LFS) with a circular plasma column in a limiter configuration. The plasma column is then extended and moved to the center of the vacuum vessel. At the next stage, strong plasma shaping occurs and a double zero configuration is created (X-point formation). Strike lines of the separatrix are directed to the closed divertor. 2 divertor coils were introduced for sweeping the separatrix legs which will also reduce average heat loads. These coils are placed on the sides of the strike lines which we will use in future simulations for sweeping the strike point to reduce localized hot spots.

Stability Evolution

Using the same shaping parameters values and plasma profiles as in the previous section, a vertical plasma stability evaluation was performed. 3 cases of the plasma-wall gap were considered: 20, 30, and 40 cm. For a more accurate comparison, all 3 cases had the same equilibria at the starting point of the simulation. Vacuum vessel currents were set to zero (stationary condition before the simulation was assumed). Coil currents were equal for each case and fixed for the simulation to evaluate the stabilization of the plasma by the vacuum vessel wall.

The results of the simulations are represented in the figures below. The first case is the most stable with ~215 ms of stable plasma and ~290 ms total VDE time. But this has a drawback, vacuum vessel currents not only stabilize plasma but also make it difficult for the PF coil's magnetic field to diffuse through the vacuum vessel.

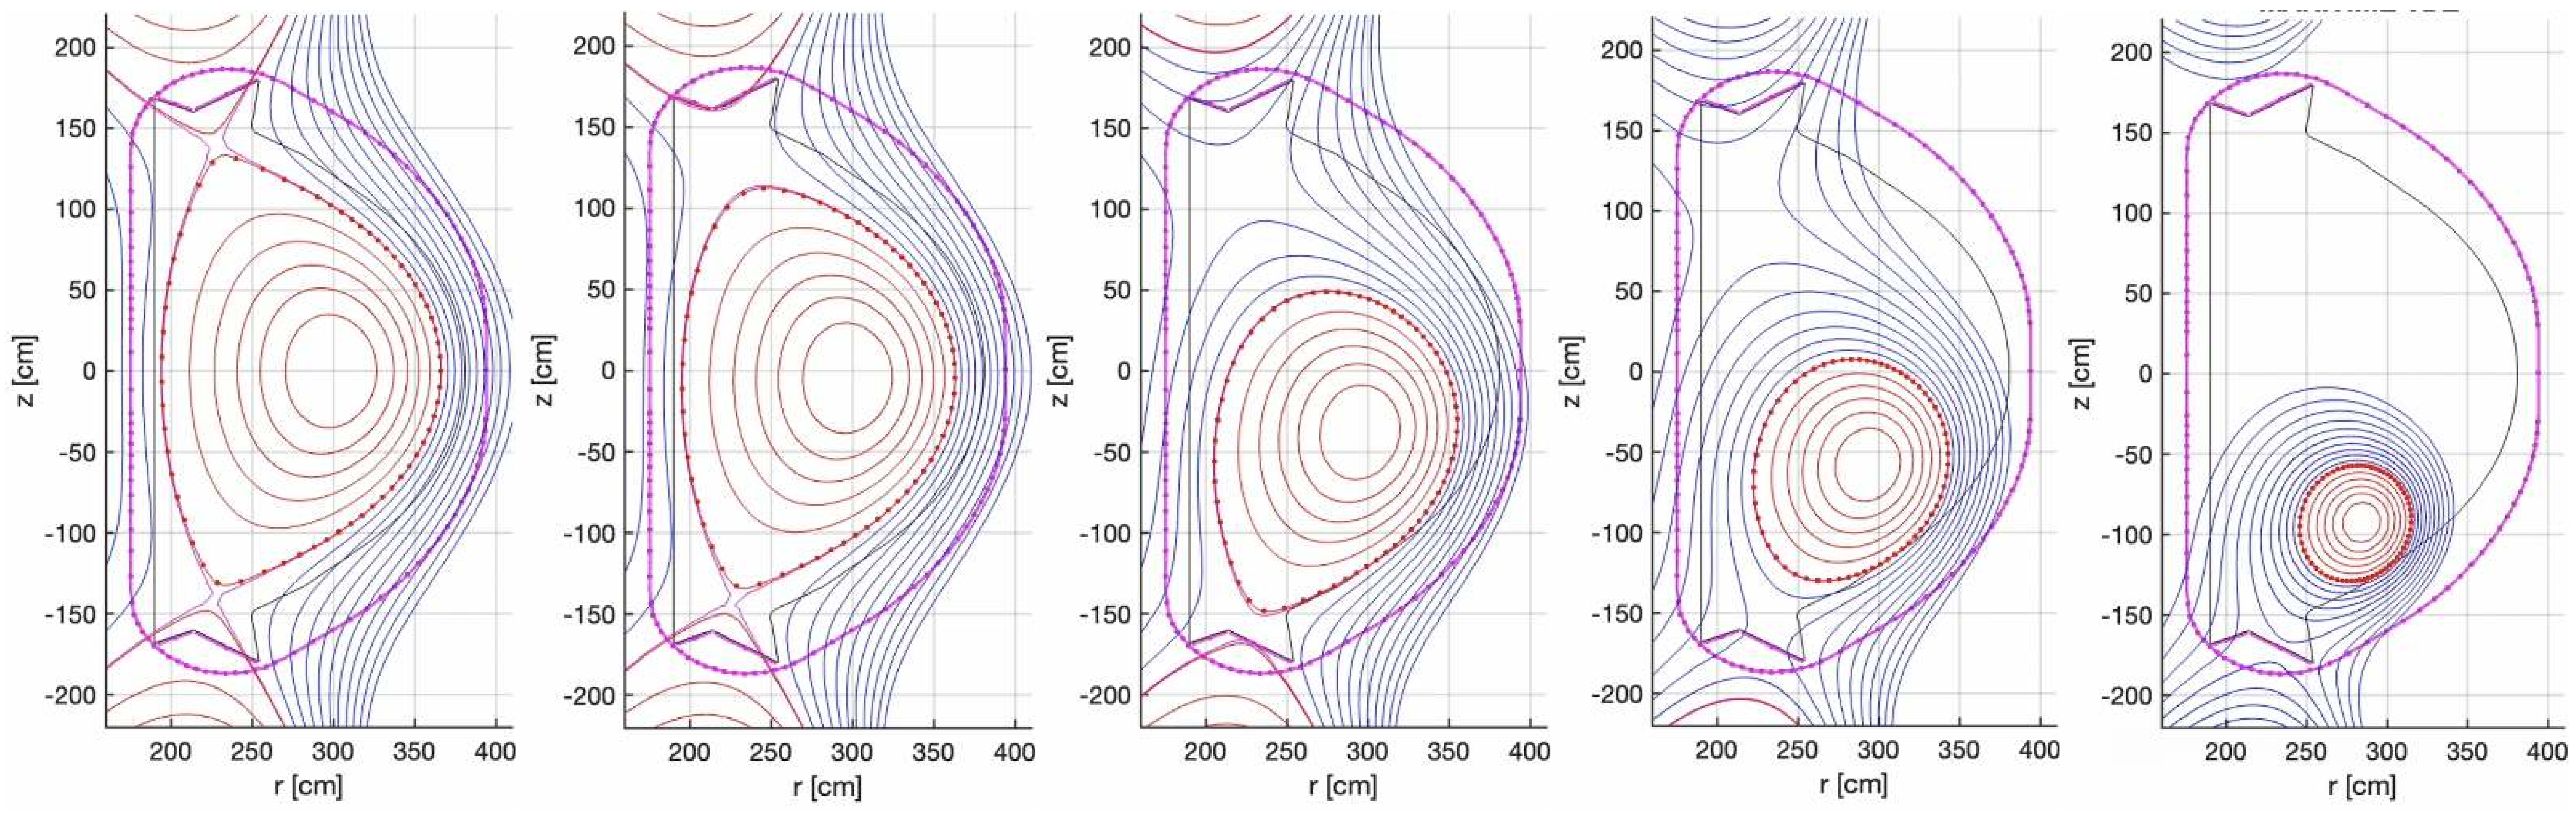

VDE plasma stability for chosen case at (left) 0ms, 120ms, 130ms, 160ms, 180ms (right)

The second and the third cases are less stable but still have quite a long stability phase: ~160 ms of stable plasma and ~215 ms total VDE time for the 30 cm gap and ~130 ms of stable plasma and ~180 ms total VDE time for the 40 cm gap. The last case (40cm gap) was chosen as the optimal one as it provides reasonably good wall stabilization, requires less currents from the PF coils to control the plasma, and offers more space for the first wall integration. Even bigger gaps may provide sufficient wall stabilization.

The main disadvantage of larger gaps between the plasma and the vacuum vessel is an increase in the size of the device and can be enhanced by introducing passive stabilization coils (PSC).

Magnet Power Supplies

Tokamak magnets require substantial currents, necessitating robust power supplies. Superconducting magnets consume the most energy during ramp-up, when currents increase rapidly. Once steady-state (flat-top) operation is reached, coil current stabilizes and di/dt decreases significantly. Thus, power supply requirements primarily depend on the energy needed during initial ramp-up to initiate plasma current, after which flux consumption in the Poloidal Field (PF) system levels off. The current design assumes fully inductive discharges, but adding external current drive technologies could ease PF coil requirements—though it remains uncertain whether they can achieve sufficient efficiencies to meaningfully reduce recirculating power.

These simulations were performed in the COMSOL Multiphysics® software environment. For this purpose, a pre-built complete 3D model (STEP format) of the tokamak magnetic system was imported into the environment. All coils in the model were set as single turns with full current; thus, the specific (per turn) characteristics of the coils can be determined. The coil parameters used for calculations and construction of the 3D model are shown below.

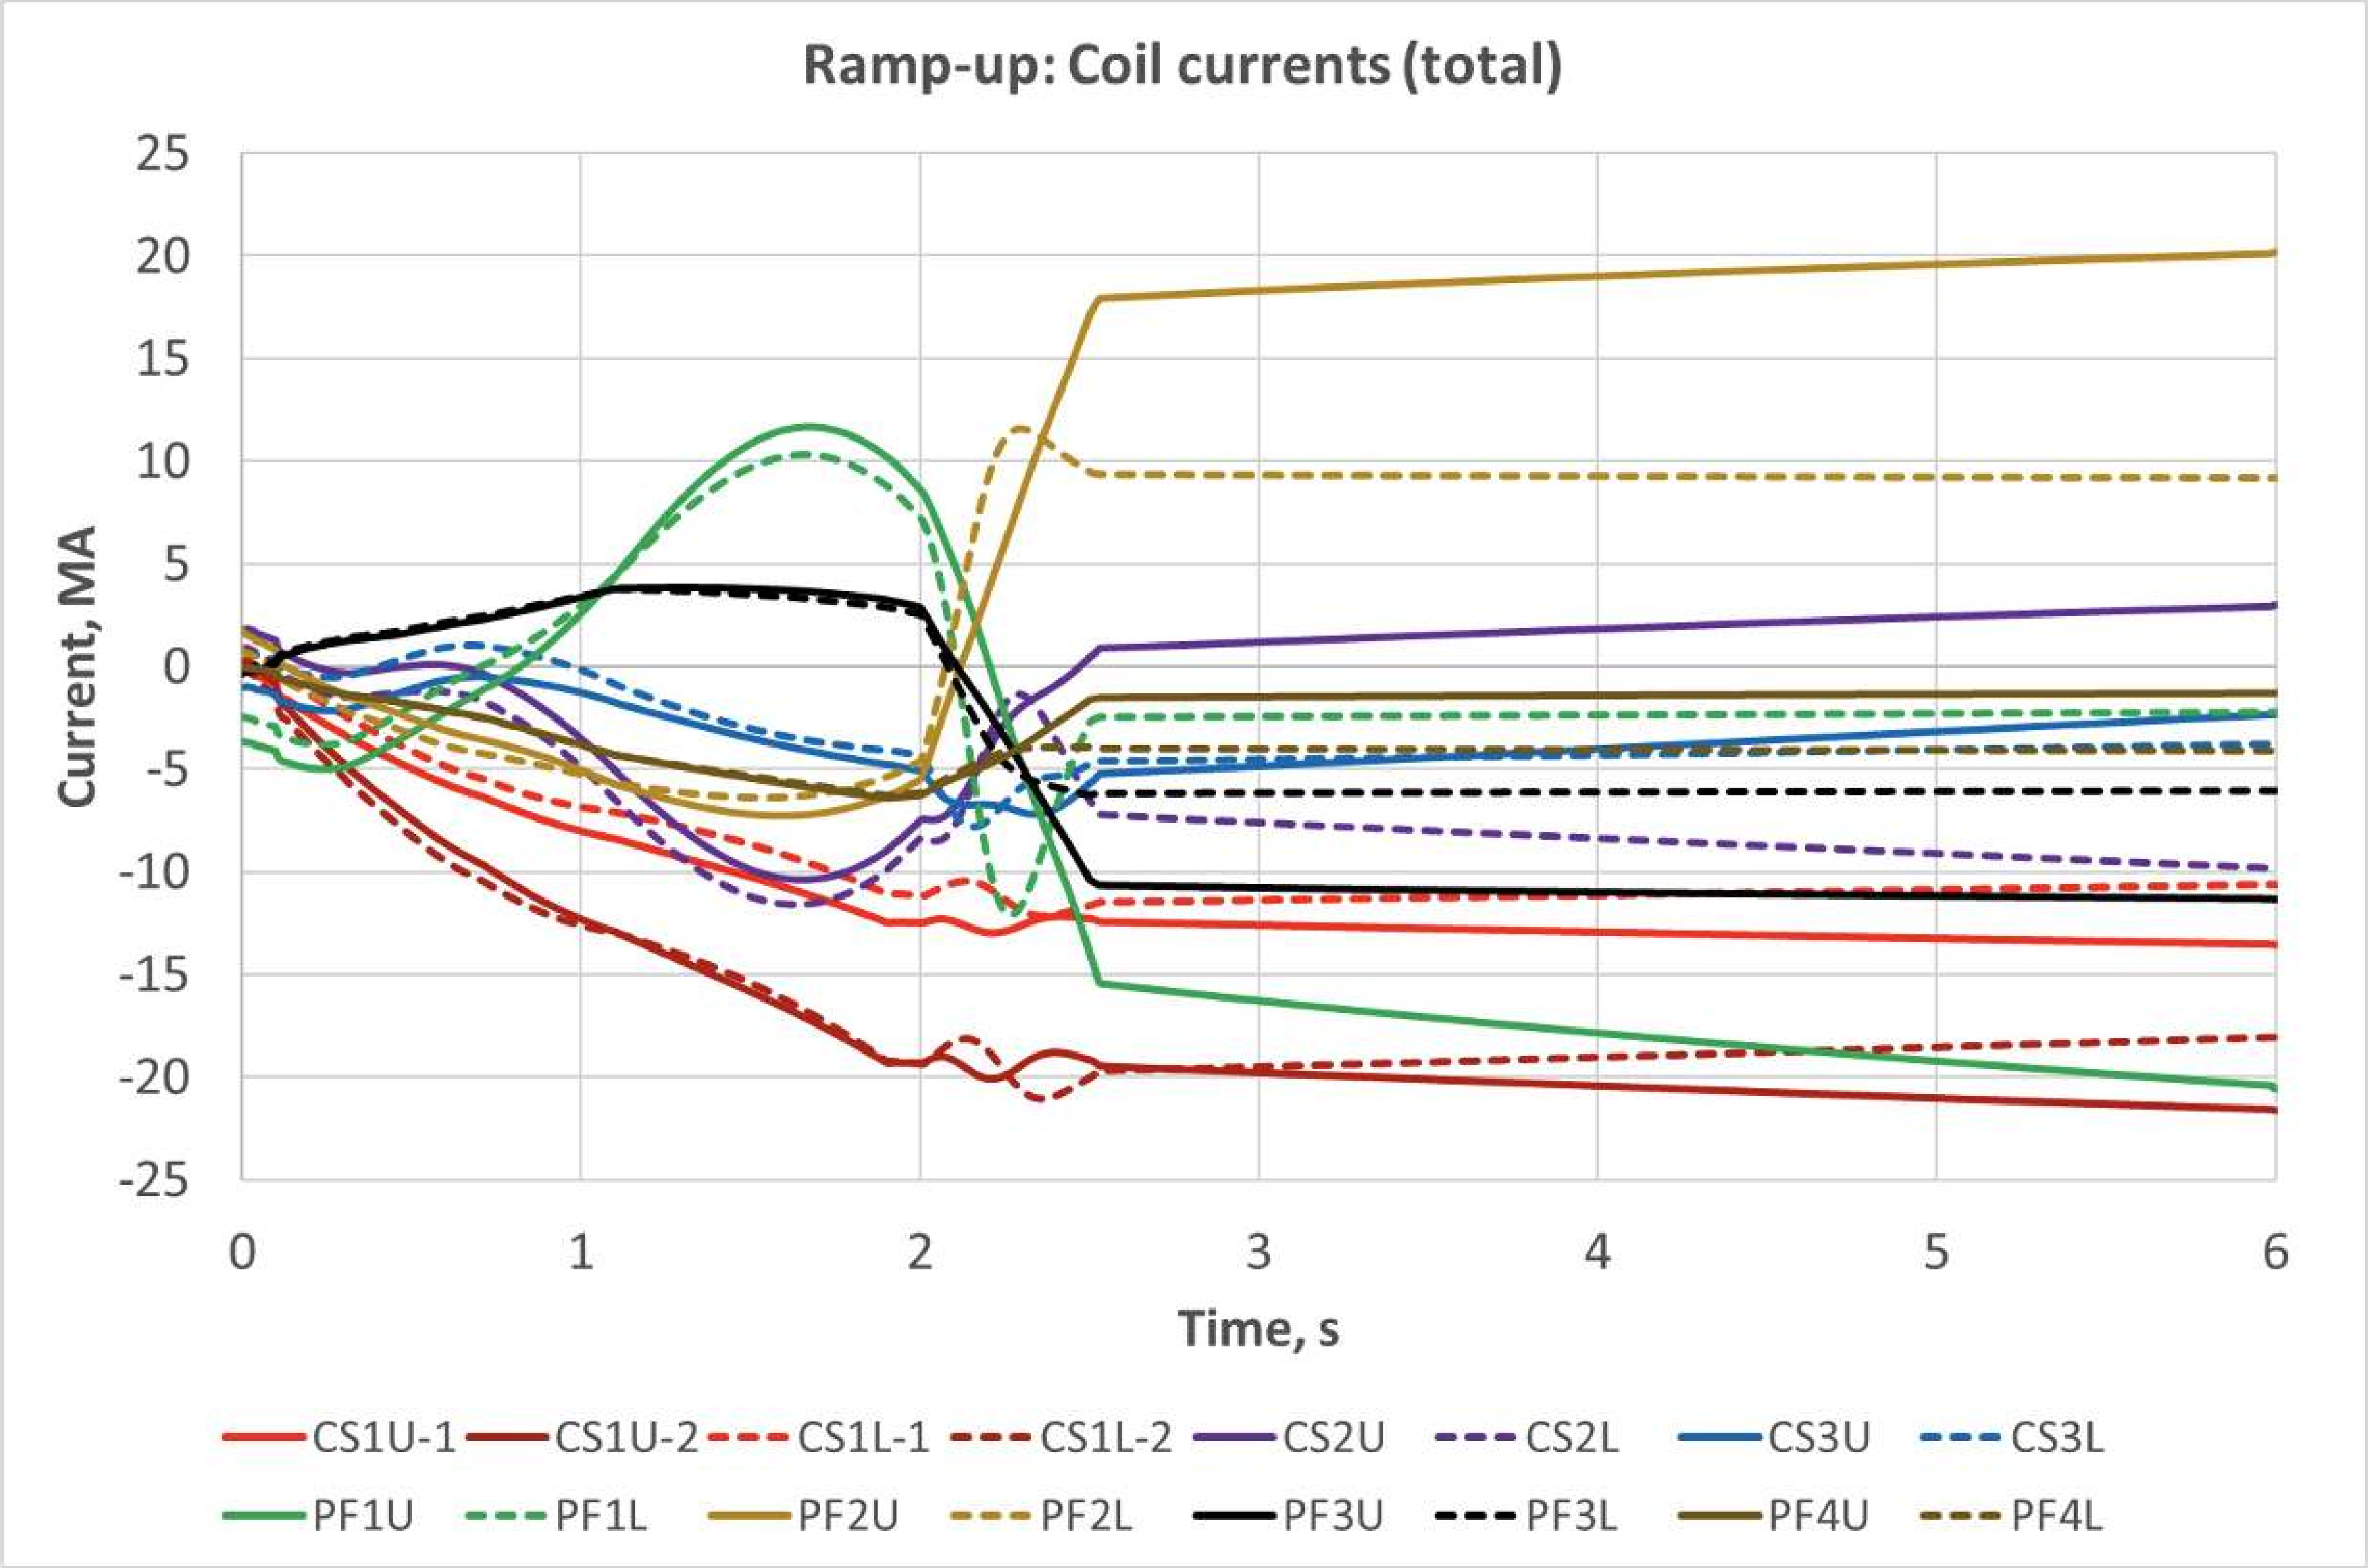

Coil currents evolution during discharge scenario for the inductive current ramp-up phase

Coil	R (Ω)	Center_r (m)	Center_z (m)	Size_r (m)	Size_z (m)	Max mag. field (T)
CS1U-1	1.06E-08	0.4586	0.4865	0.074676	0.866216	23.0
CS1U-2	4.09E-09	0.5751	0.4865	0.153664	0.864076	26.9
CS1L-1	1.06E-08	0.4586	-0.4865	0.074676	0.866216	23.0
CS1L-2	4.09E-09	0.5751	-0.4865	0.153664	0.864076	26.9
CS2U	2.33E-10	0.5415	1.216	0.21952	0.37107	16.0
CS2L	2.33E-10	0.5415	-1.216	0.21952	0.37107	16.0
CS3U	2.33E-10	0.5415	1.694	0.21952	0.37107	6.8
CS3L	2.33E-10	0.5415	-1.694	0.21952	0.37107	6.8
PF1U	4.71E-11	1.0802	2.7243	0.2548	0.262	36.8
PF1L	4.71E-11	1.0802	-2.7243	0.2548	0.262	36.8
PF2U	4.84E-11	2.2345	3.12	0.351753	0.26362	30.7
PF2L	4.84E-11	2.2345	-3.12	0.351753	0.26362	30.7
PF3U	3.04E-11	3.8407	2.605	0.306129	0.2107	21.7
PF3L	3.04E-11	3.8407	-2.605	0.306129	0.2107	21.7
PF4U	1.92E-11	4.9824	1.3479	0.2548	0.315967	5.8
PF4L	1.92E-11	4.9824	-1.3479	0.2548	0.315967	5.8

Table 1: Calculated parameters of the coils and power supplies

Using a series of calculations performed in COMSOL within an integrated 3D model, the total intrinsic magnetic flux and energy of each coil, reduced to a single turn, were determined. From these calculations, the inductance coefficient or the single-turn inductance (k_L) was derived (see Table 2). This coefficient can be used, in conjunction with the formula $k_L = L/N^2$, to calculate the inductance (L) of each coil based on the number of turns (N), or to determine the required number of turns to achieve a required inductance. The voltage per turn (V_{turn}) is calculated by multiplying the single-turn inductance (k_L) by the maximum derivative of the total current (di/dt), $V_{turn} = k_L \cdot di/dt$. It is important to note that the total energy of the coil is independent of the turn number.

Coil	k_L (uH)	I_{max} (MA)	di/dt (MA/s)	U_{turn} (V)	E_{total} (MJ)	Cap. (MJ)	N_{cap}	m_{cap} (t)	V_{cap} (m^3)
CS1U-1	0.60	20	20	12	119.3	155.1	15509	7.8	10.9
CS1U-2	0.82	25	20	16.4	257.2	334.4	33436	16.7	23.4
CS1L-1	0.60	20	20	12	119.3	155.1	15509	7.8	10.9
CS1L-2	0.82	25	20	16.4	257.2	334.4	33436	16.7	23.4
CS2U	1.04	10	45	46.8	52.1	67.7	6773	3.4	4.7
CS2L	1.04	10	45	46.8	52.1	67.7	6773	3.4	4.7
CS3U	1.04	4	30	31.2	8.3	10.8	1079	0.5	0.8
CS3L	1.04	4	30	31.2	8.3	10.8	1079	0.5	0.8
PF1U	2.93	25	100	293	914.6	1189.0	118898	59.4	83.2
PF1L	2.93	25	100	293	914.6	1189.0	118898	59.4	83.2
PF2U	7.96	25	75	597	2486.1	3231.9	323193	161.6	226.2
PF2L	7.96	25	75	597	2486.1	3231.9	323193	161.6	226.2
PF3U	18.75	15	30	562.5	2109.9	2742.9	274287	137.1	192.0
PF3L	18.75	15	30	562.5	2109.9	2742.9	274287	137.1	192.0
PF4U	23.93	5	15	358.95	299.0	388.7	38870	19.4	27.2
PF4L	23.93	5	15	358.95	299.0	388.7	38870	19.4	27.2
TOTAL:					12493	16240.9	1624090	812.0	1136.863

Table 2: Initial parameters of the CS and PF coils

Since grid power is unavailable in remote locations, the only way to supply power to the magnetic system coils is through energy buffering and storage. This can be achieved using either battery storage or capacitive storage. These systems must deliver the peak power required during a ramp-up phase of the discharge. Energy and power density are important parameters for choosing which type of power supply to use. A hybrid system combining supercapacitors and batteries is the optimal choice for an energy buffer.

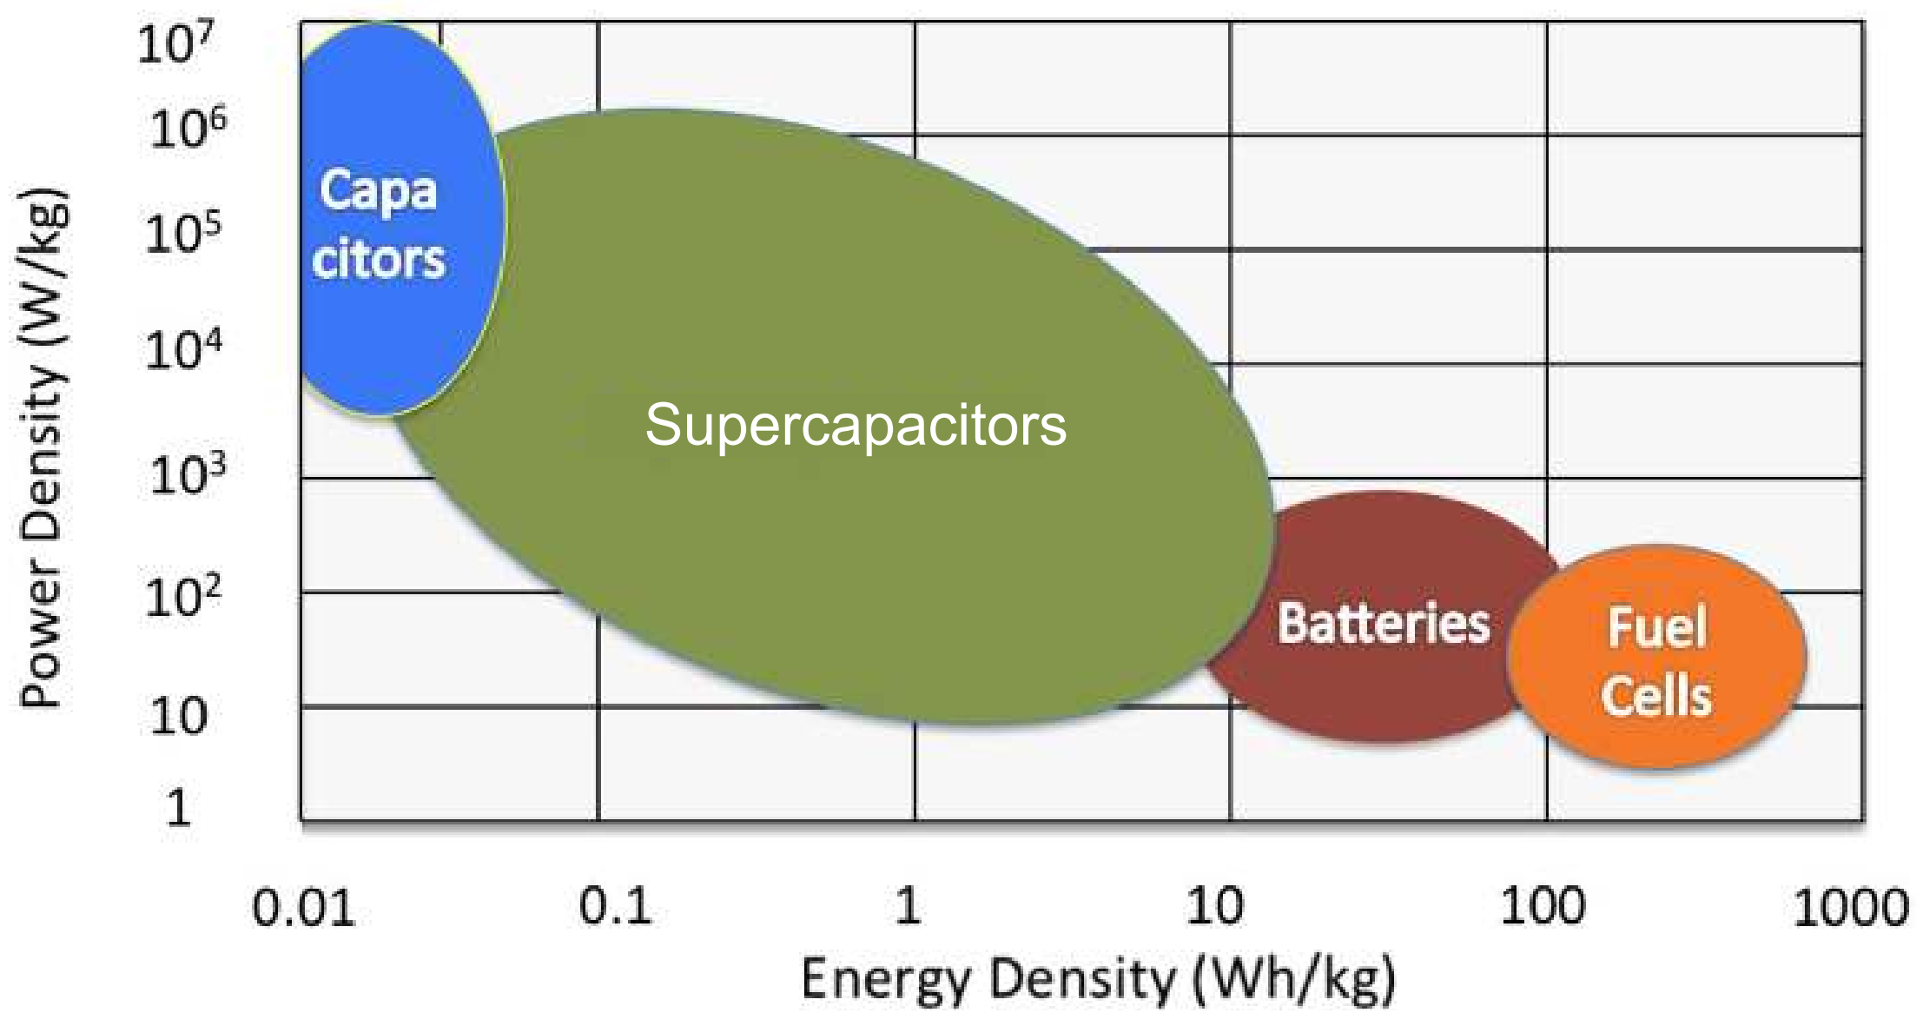

Simplified Ragone plot of the energy storage domains [source]

A rough estimate of the power storage mass and volume:

- TF Battery storage: 30 tons and 250 m^3
- PF Supercapacitor storage: 812 tons and 1140 m^3
- ***TOTAL***: 842 tons and 1390 m^3

A simplified Ragone plot (left) shows that battery storage units offer excellent energy density but can't deliver the high current derivatives required by the poloidal field system. They are most suitable for supplying the TF coils that do not require large values of current derivatives. Traditional capacitor storage provides exceptionally high power density but falls short in energy density. Supercapacitor cells, on the other hand, have characteristics that bridge the gap between batteries and traditional capacitors, offering a balanced combination of energy and power density. This balance makes them well-suited to meet the demands of PF and CS power supply.

Based on the total energy per coil, the required number of supercapacitor cells (3000 F, 3 V) is calculated, assuming an energy capacity of 10 kJ per cell with a 30% safety margin. The gross weight of each cell is estimated at 0.5 kg, accounting for a cell weight range of 0.4–0.45 kg. The average volume per cell is calculated based on its dimensions, assuming efficient dense packaging. All relevant parameters are summarized in Table 2.

HTS Magnets

Our magnet system is composed of:

- ✓ **16 Toroidal Field (TF) Coils** with outer intercoil structures (OIS) in between vacuum vessel ports and gravity supports below;
- ✓ **6 Central Solenoid (CS) Coils** with their pre-compression system and gravity supports on top;
- ✓ **8 Poloidal Field (PF) Coils** supported by the TF Coils with flexible supports, except for PF1U & L which are clamped with sliding interfaces.

Cable Cross Sections

TF Cable (36kA)

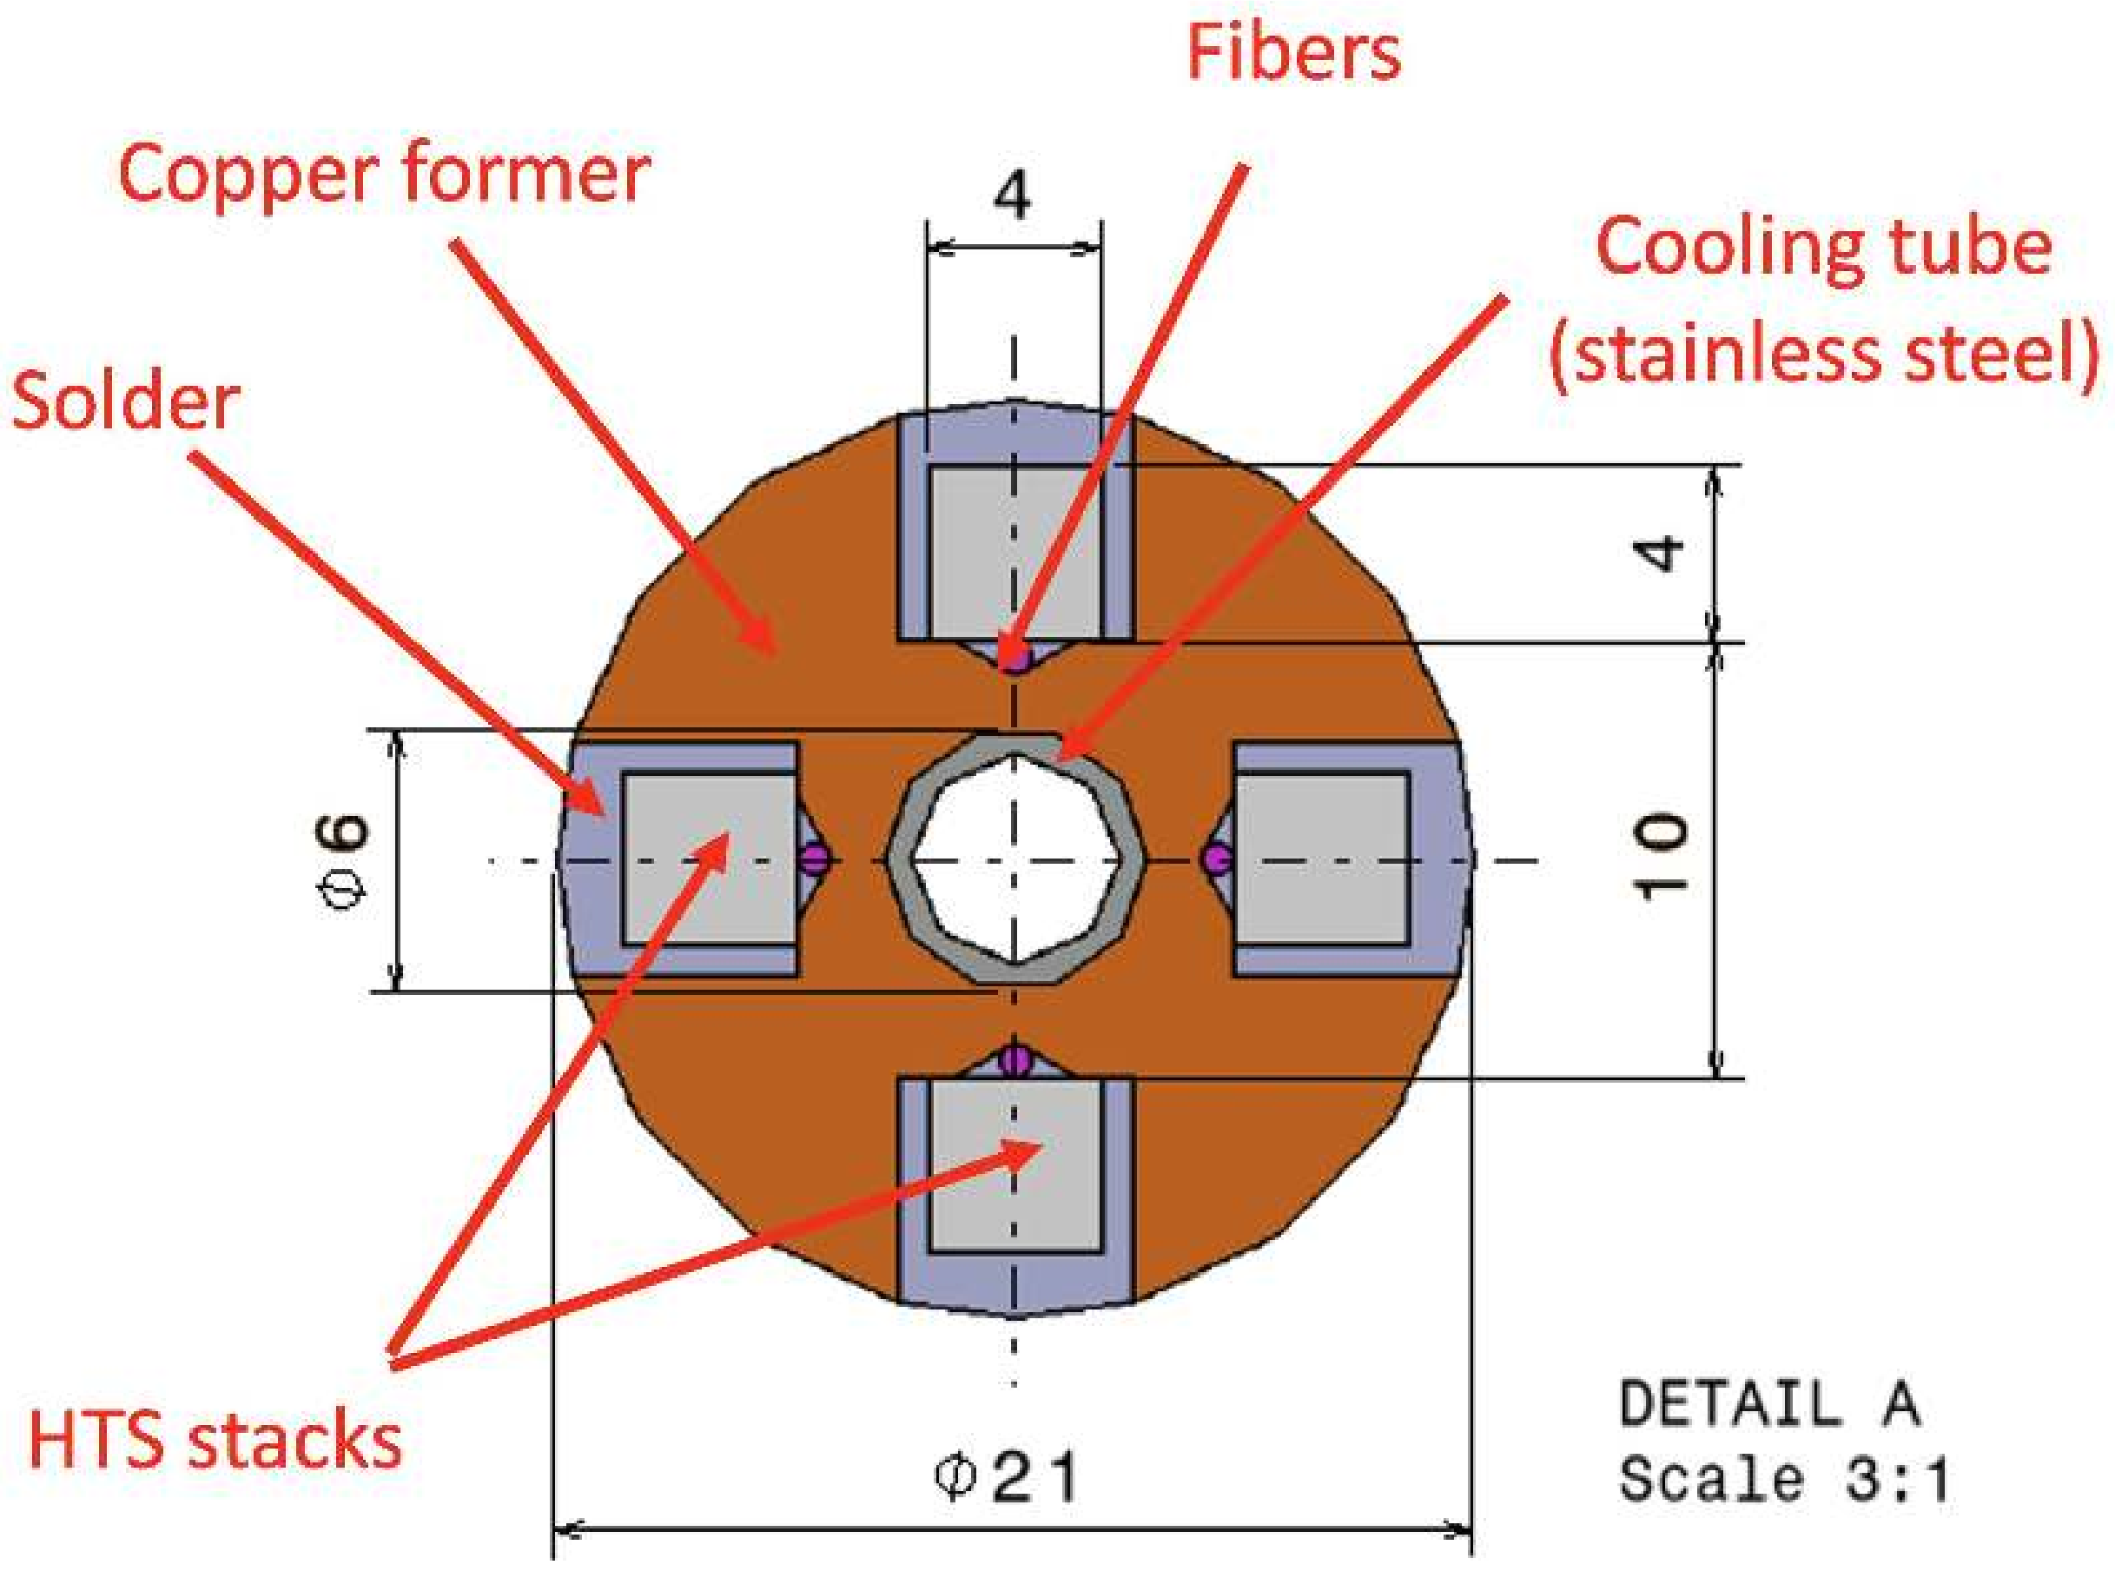

- Uses RF electromagnetics to quickly detect local heat variations and mitigate risk of magnet quenching.
- Resilient to >1000 kilonewtons of force per meter of cable
- Copper "former" which holds 60/40 Pb/Sn electroplated REBCO
- **Total cable area** = 0.073 m2

CS/PF Cables (50kA)

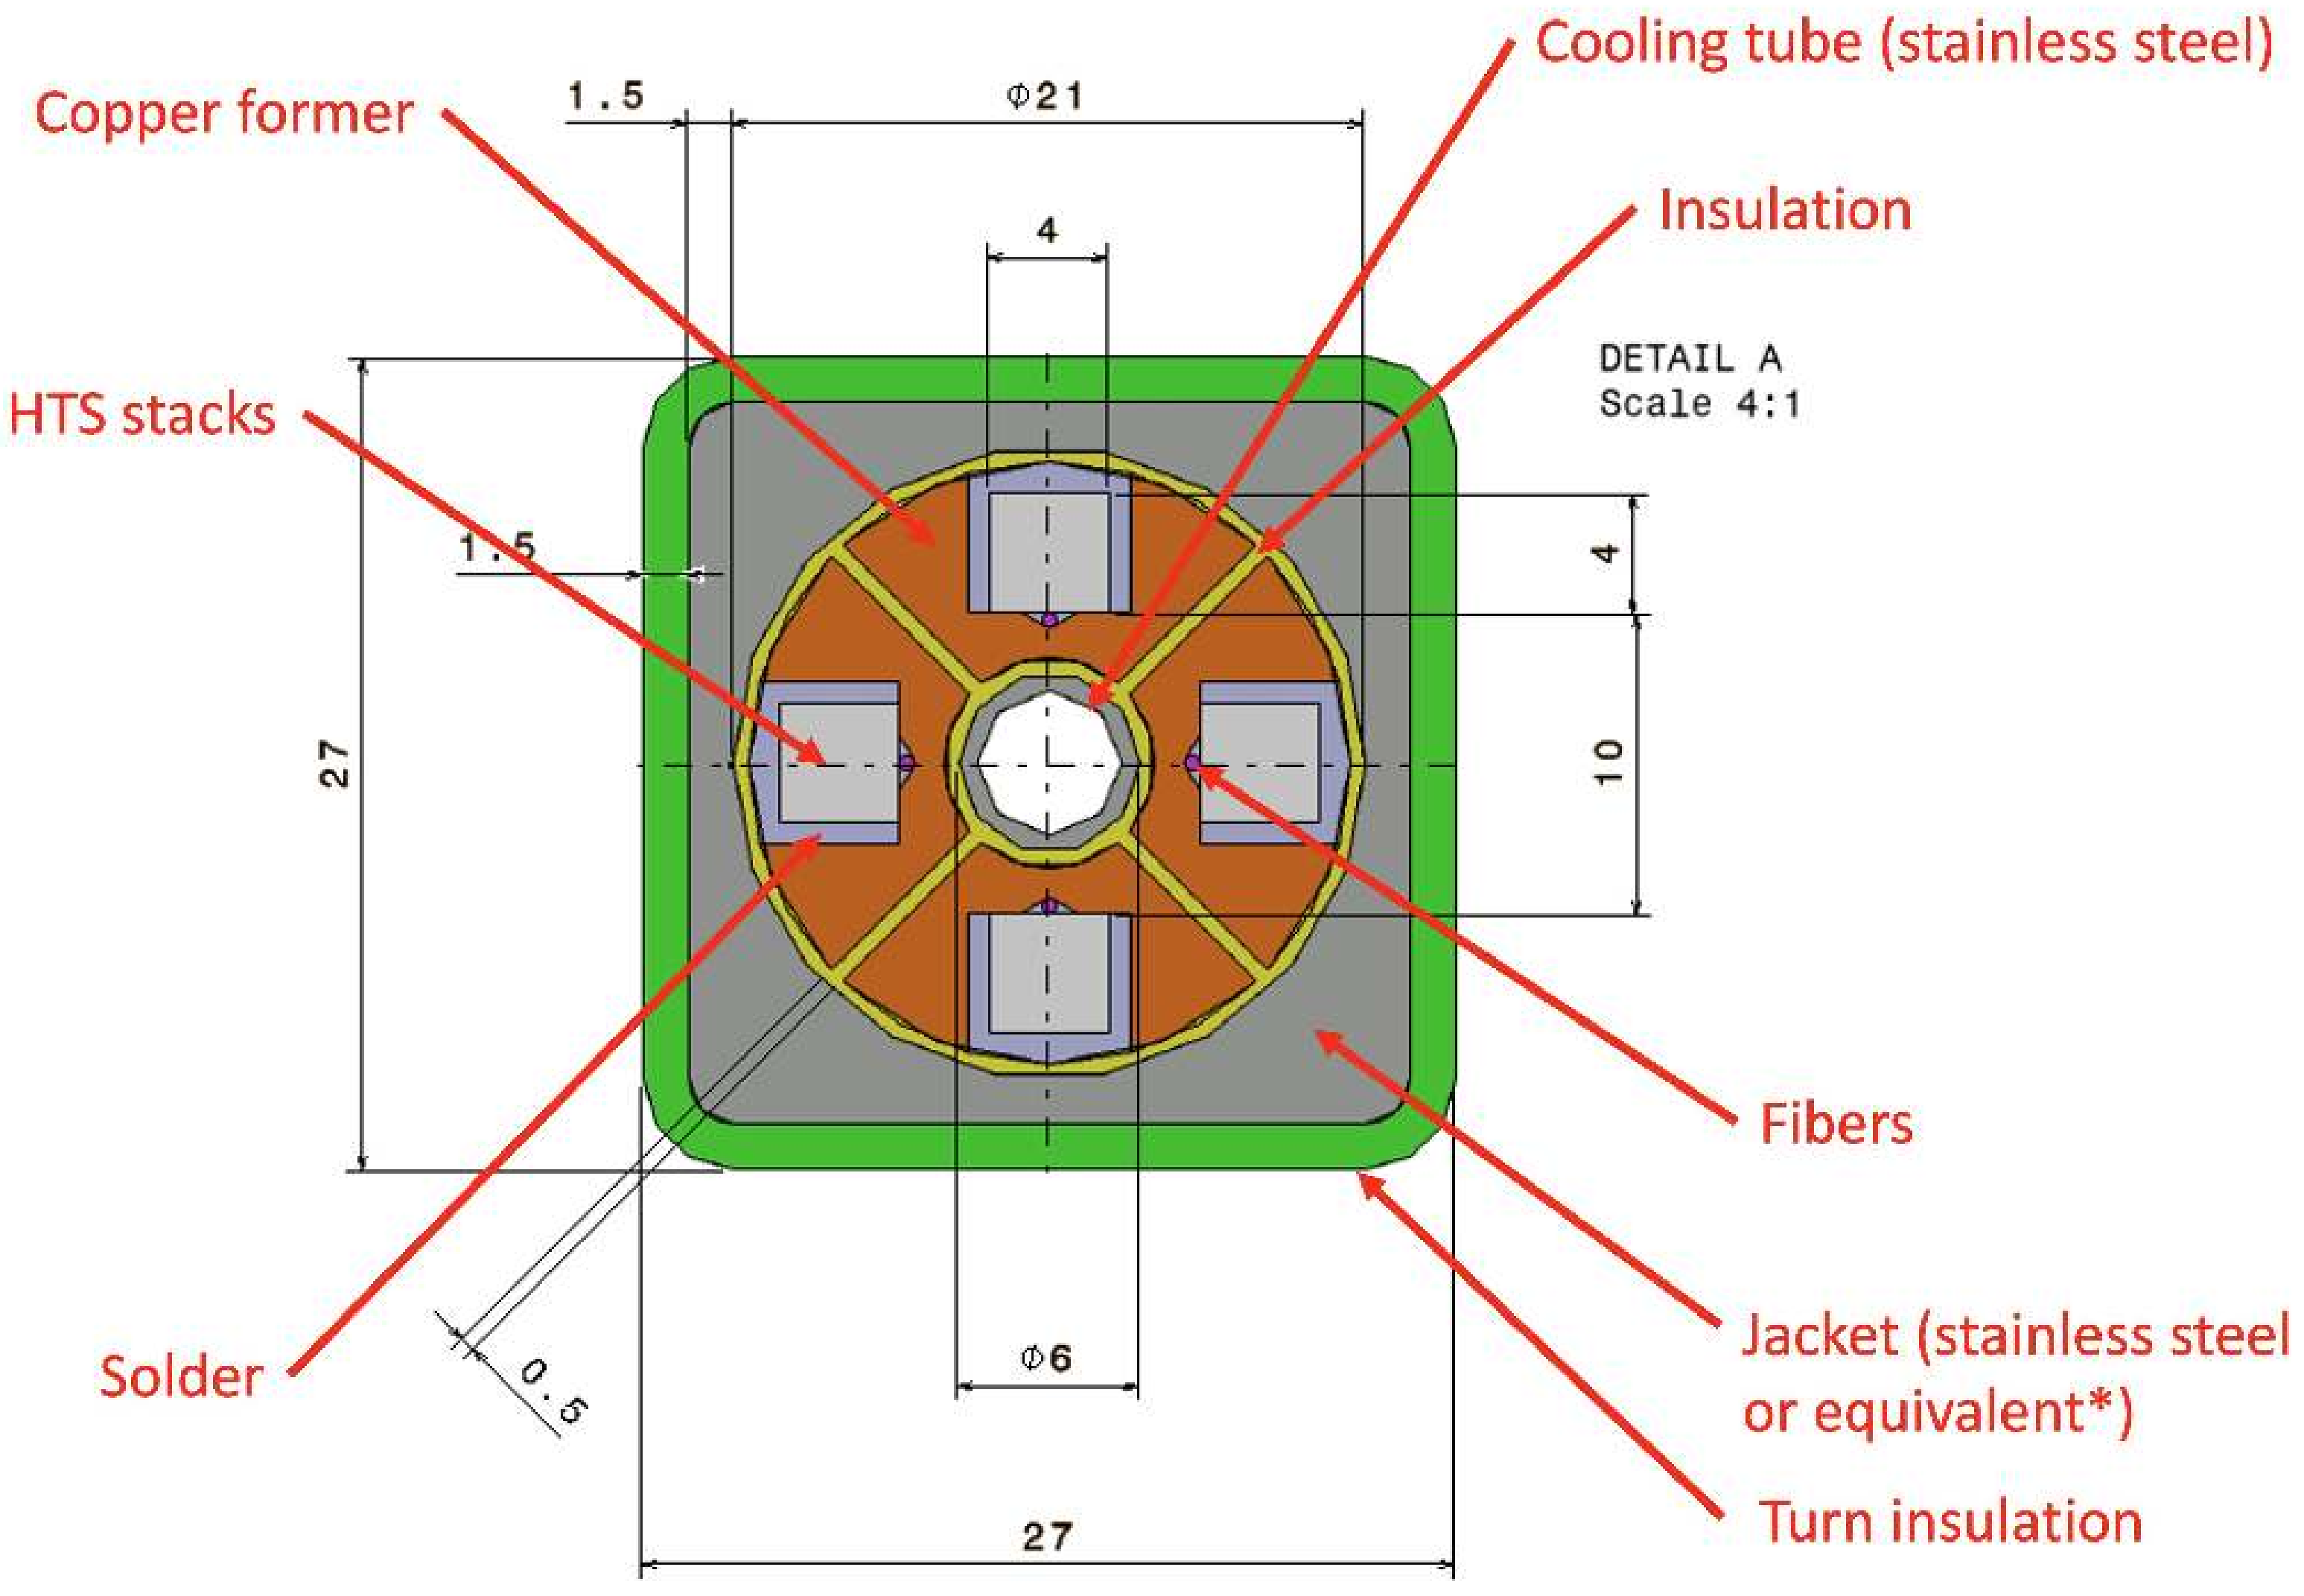

Toroidal Field Coil Assembly

- Winding pack composed of 7 double pancakes (DP) insulated separately
- 216 turns total, and **36 kA/turn** → *this can be increased.*
- **Winding pack area** = 0.17 m2
- Whole winding pack is insulated before insertion into the steel case
- **Total area** (whole TF section with case) = 0.295 m2

Inner Section

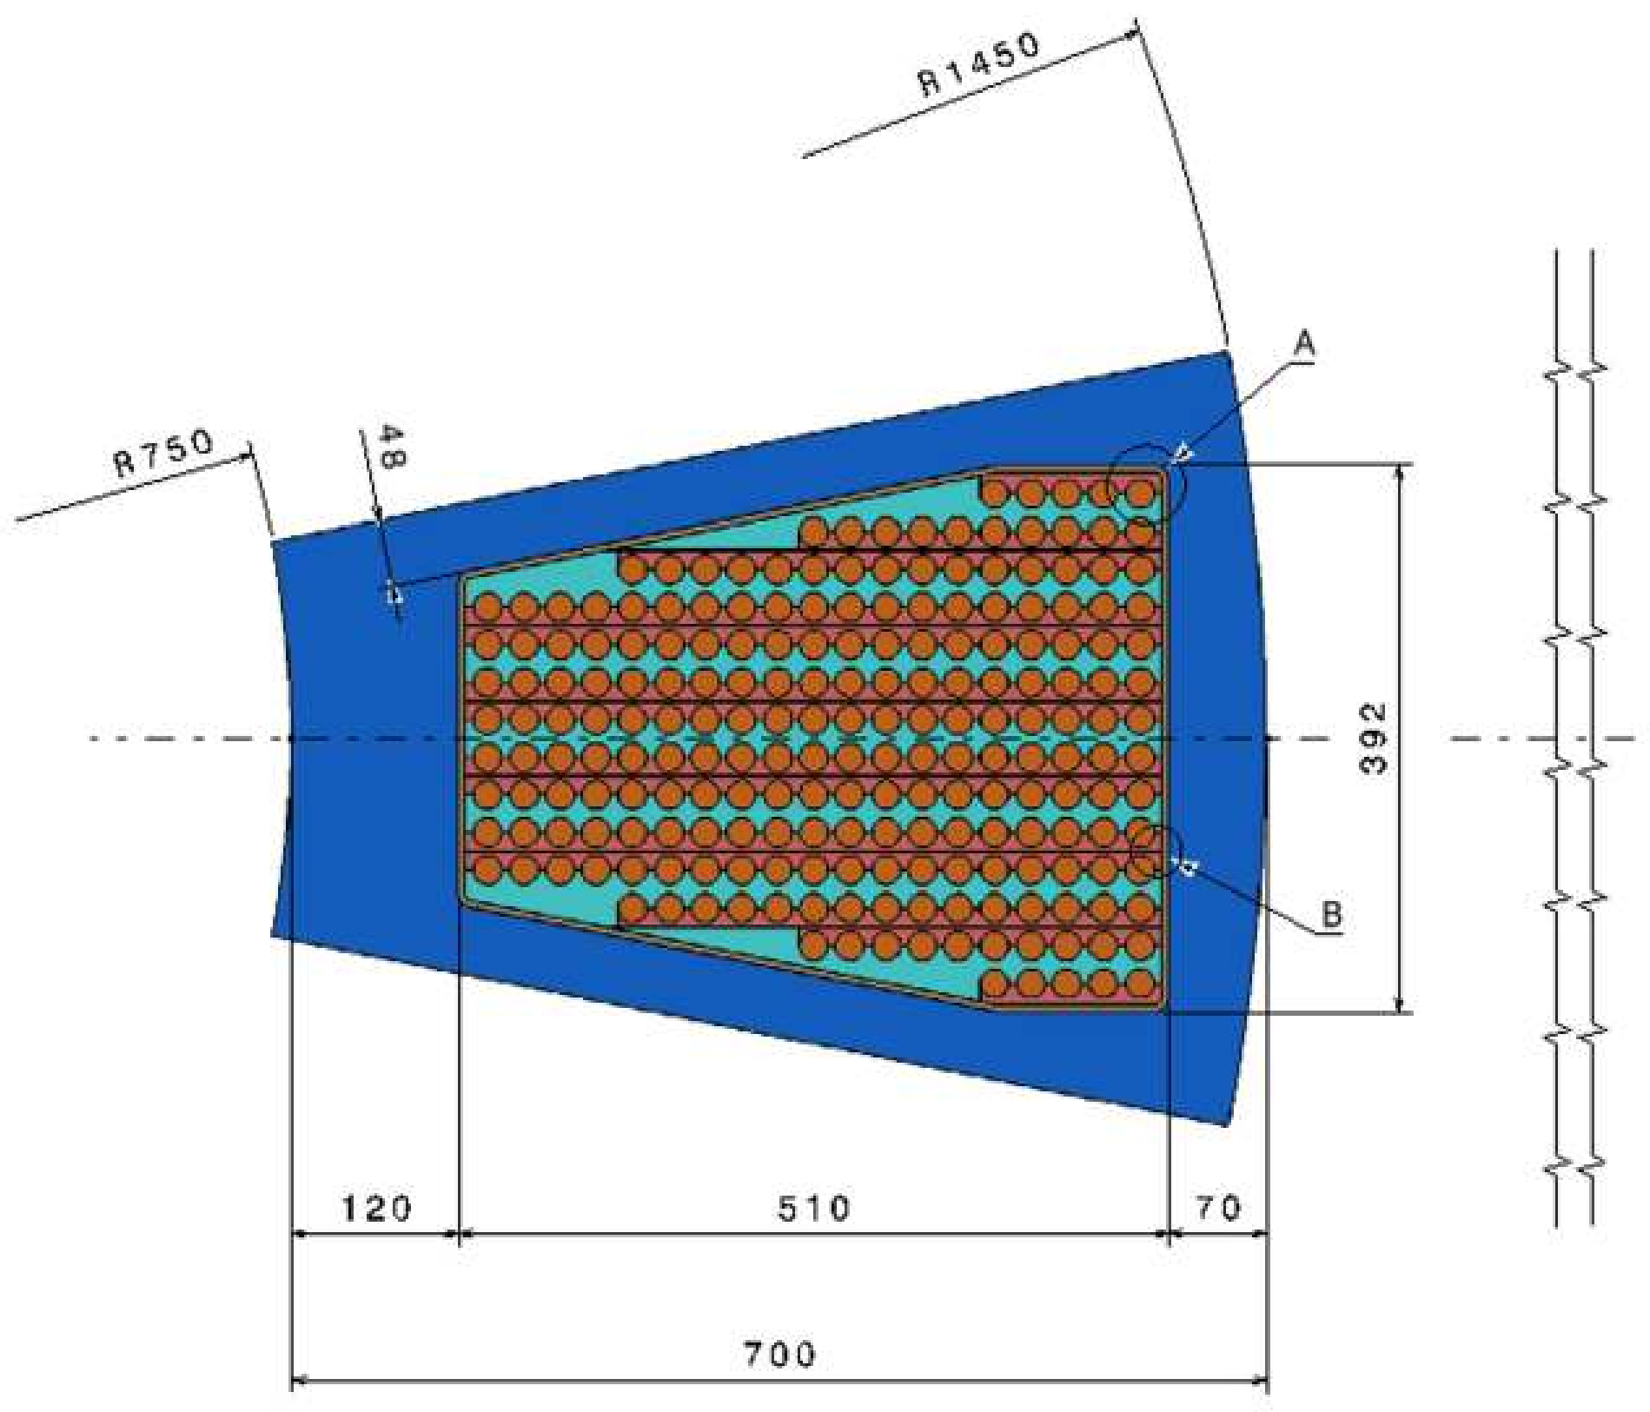

Outer Section

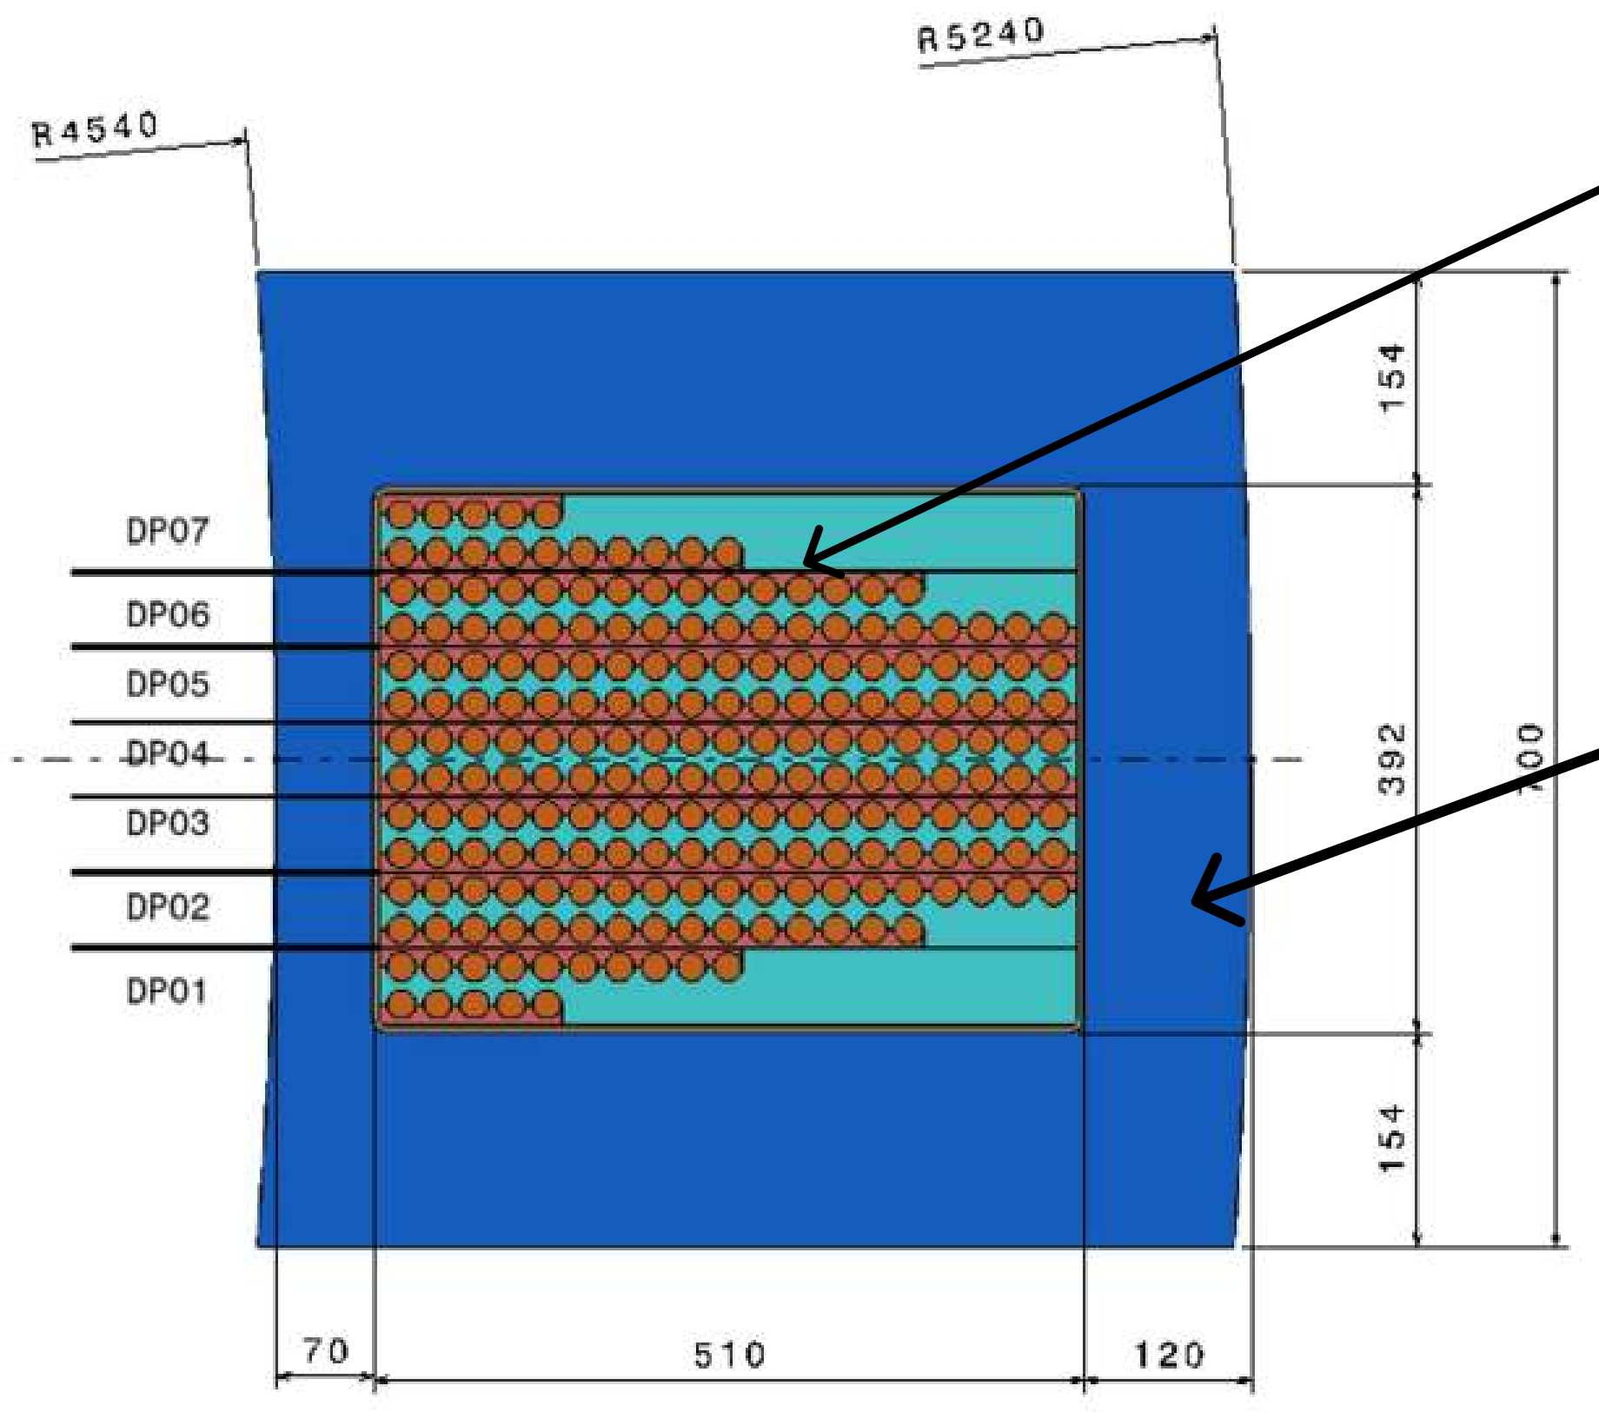

Winding Pack (**current density:** ~ 46 A/mm2)

Coil Case

Wedged Area

TF Coil Supports

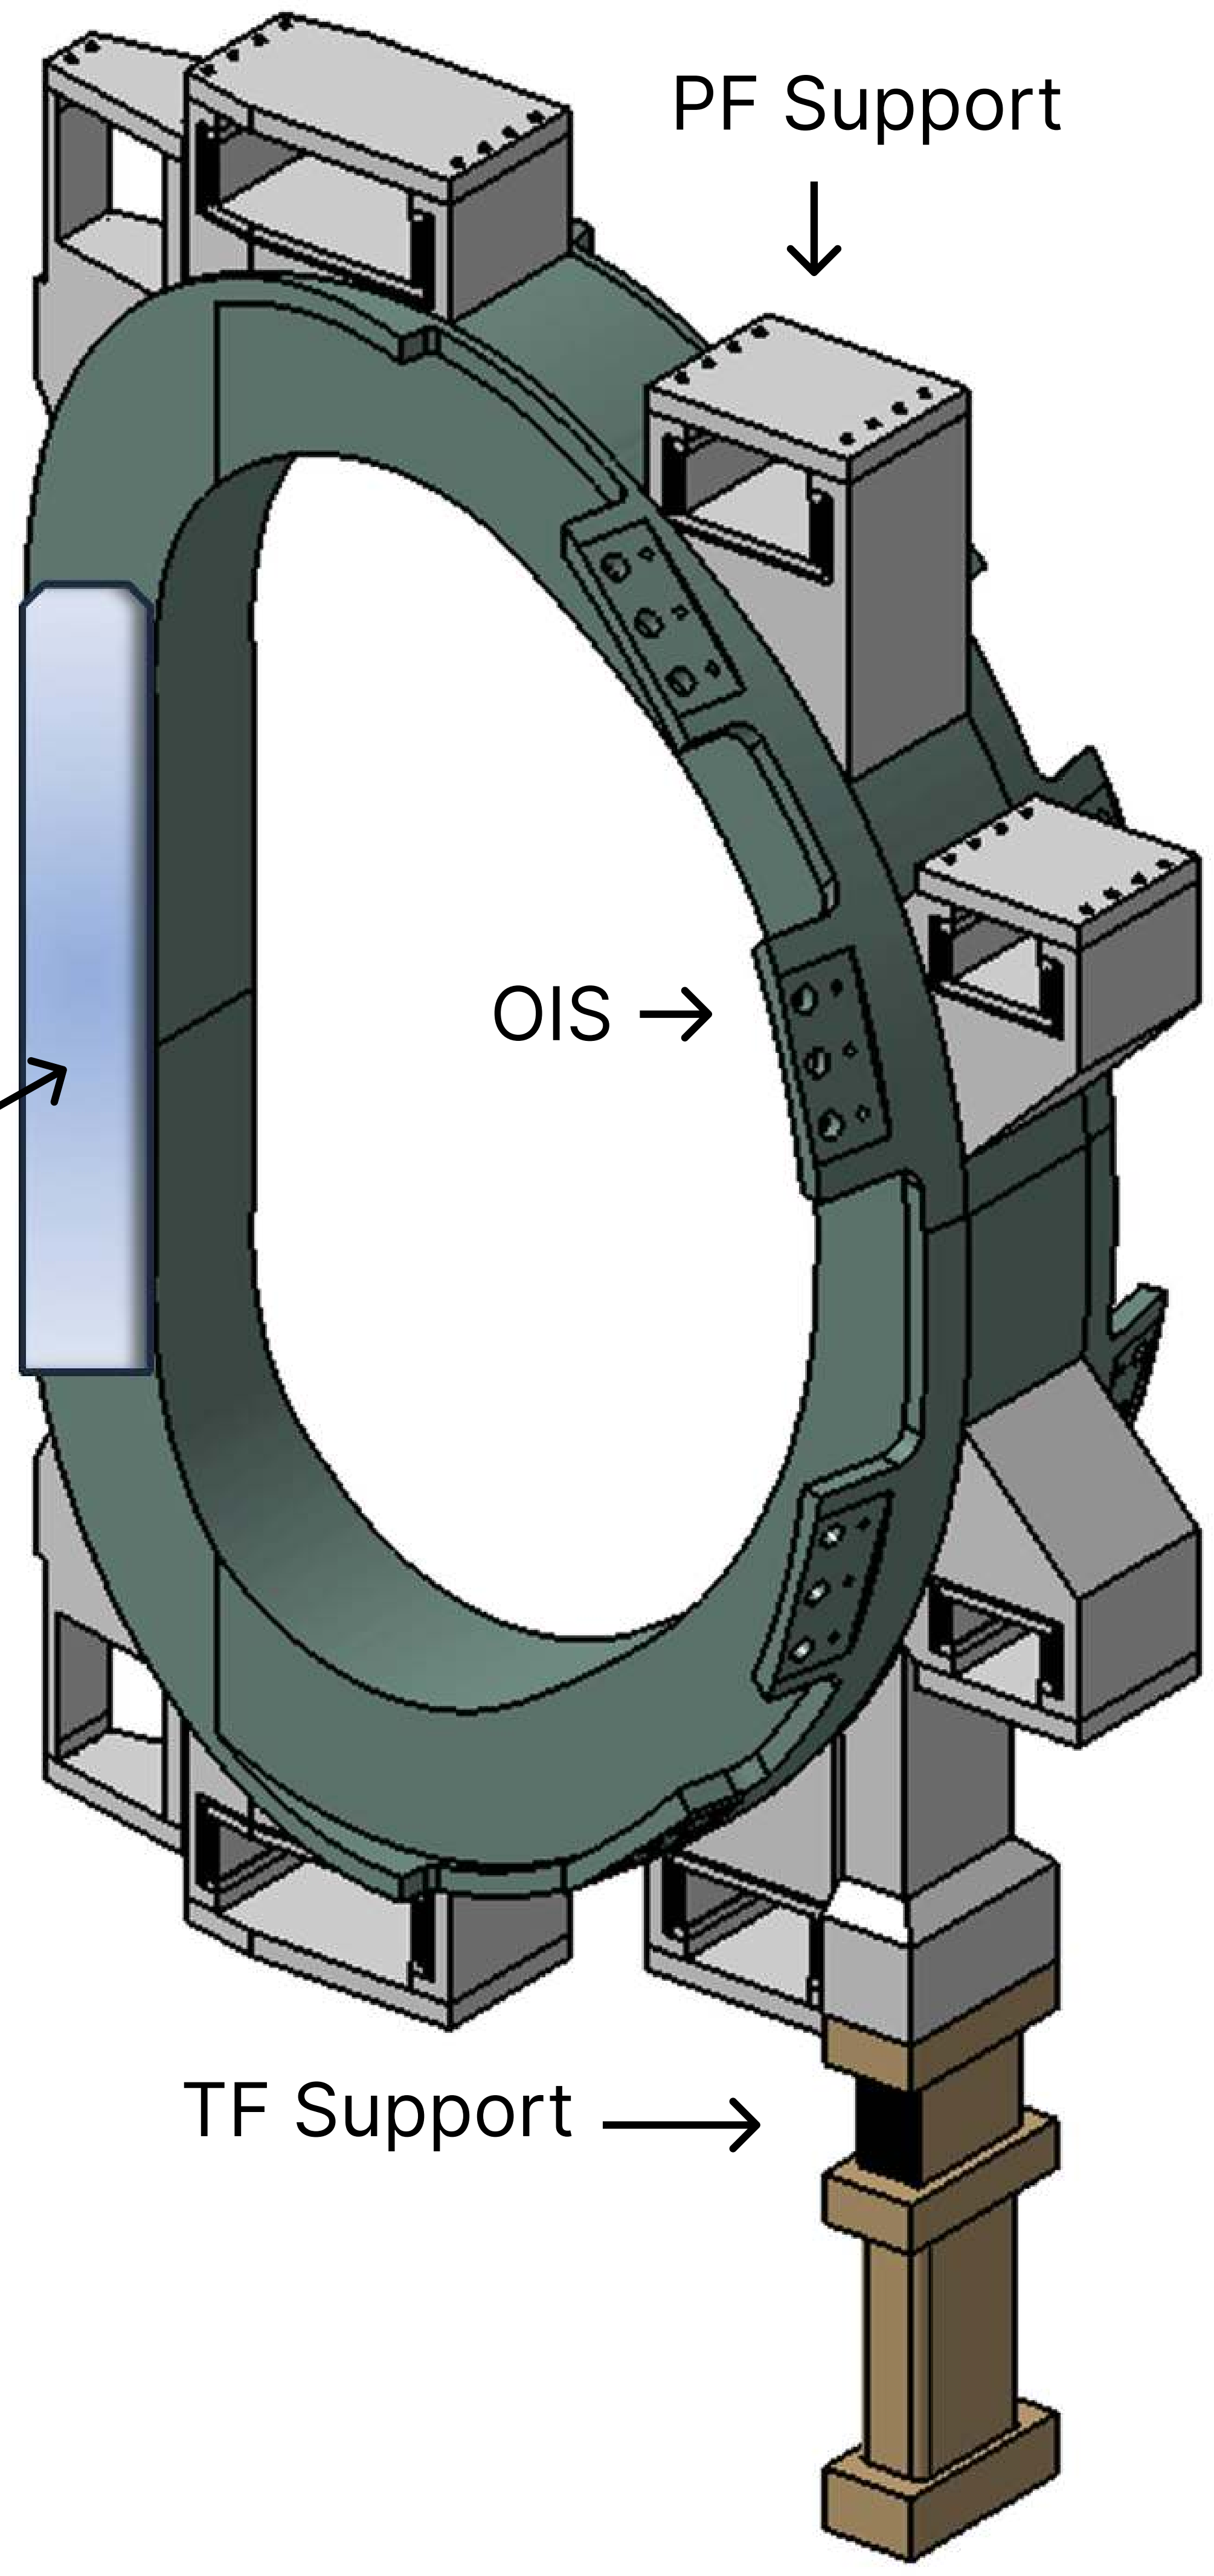

TF Coil Construction

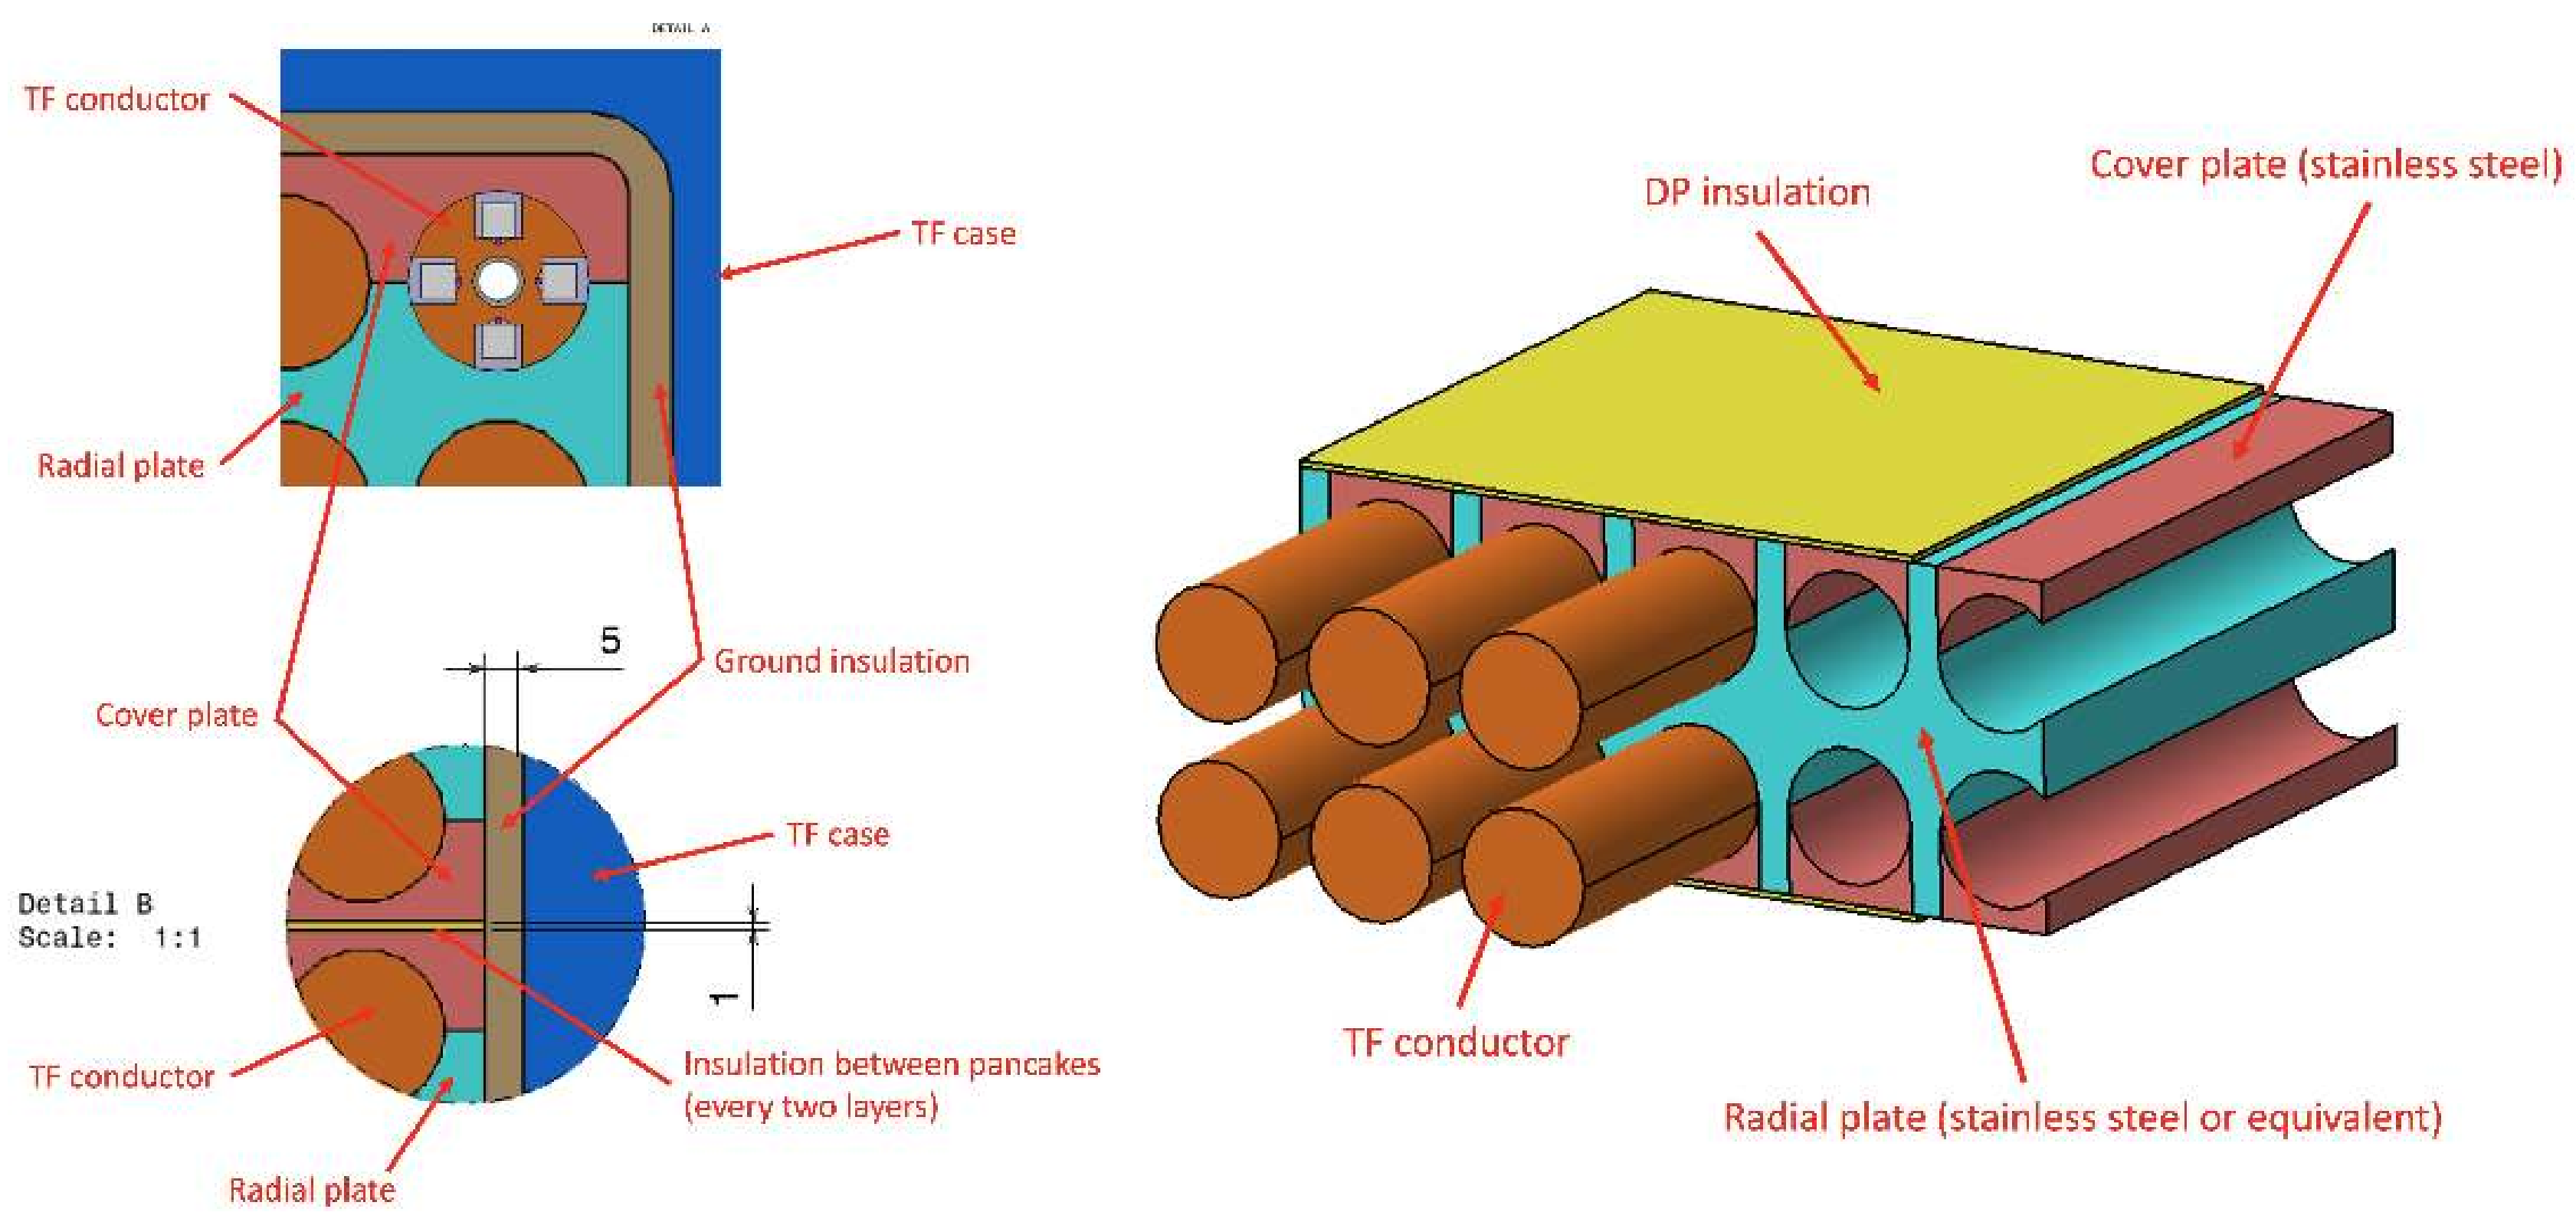

- TF coils are wedged against each at the inboard straight leg.
- A stainless steel (or equivalent material) case encloses the TF winding pack.
- Shear keys (insulated) should be inserted at the inner curved regions to withstand the overturning moment during the pulse.
- Four outer intercoil structures connected to each other by insulated pins.
- PF coil support flanges are integral part of the TF coil case.

HTS Magnets (cont.)

Poloidal Field + Central Solenoid Assembly

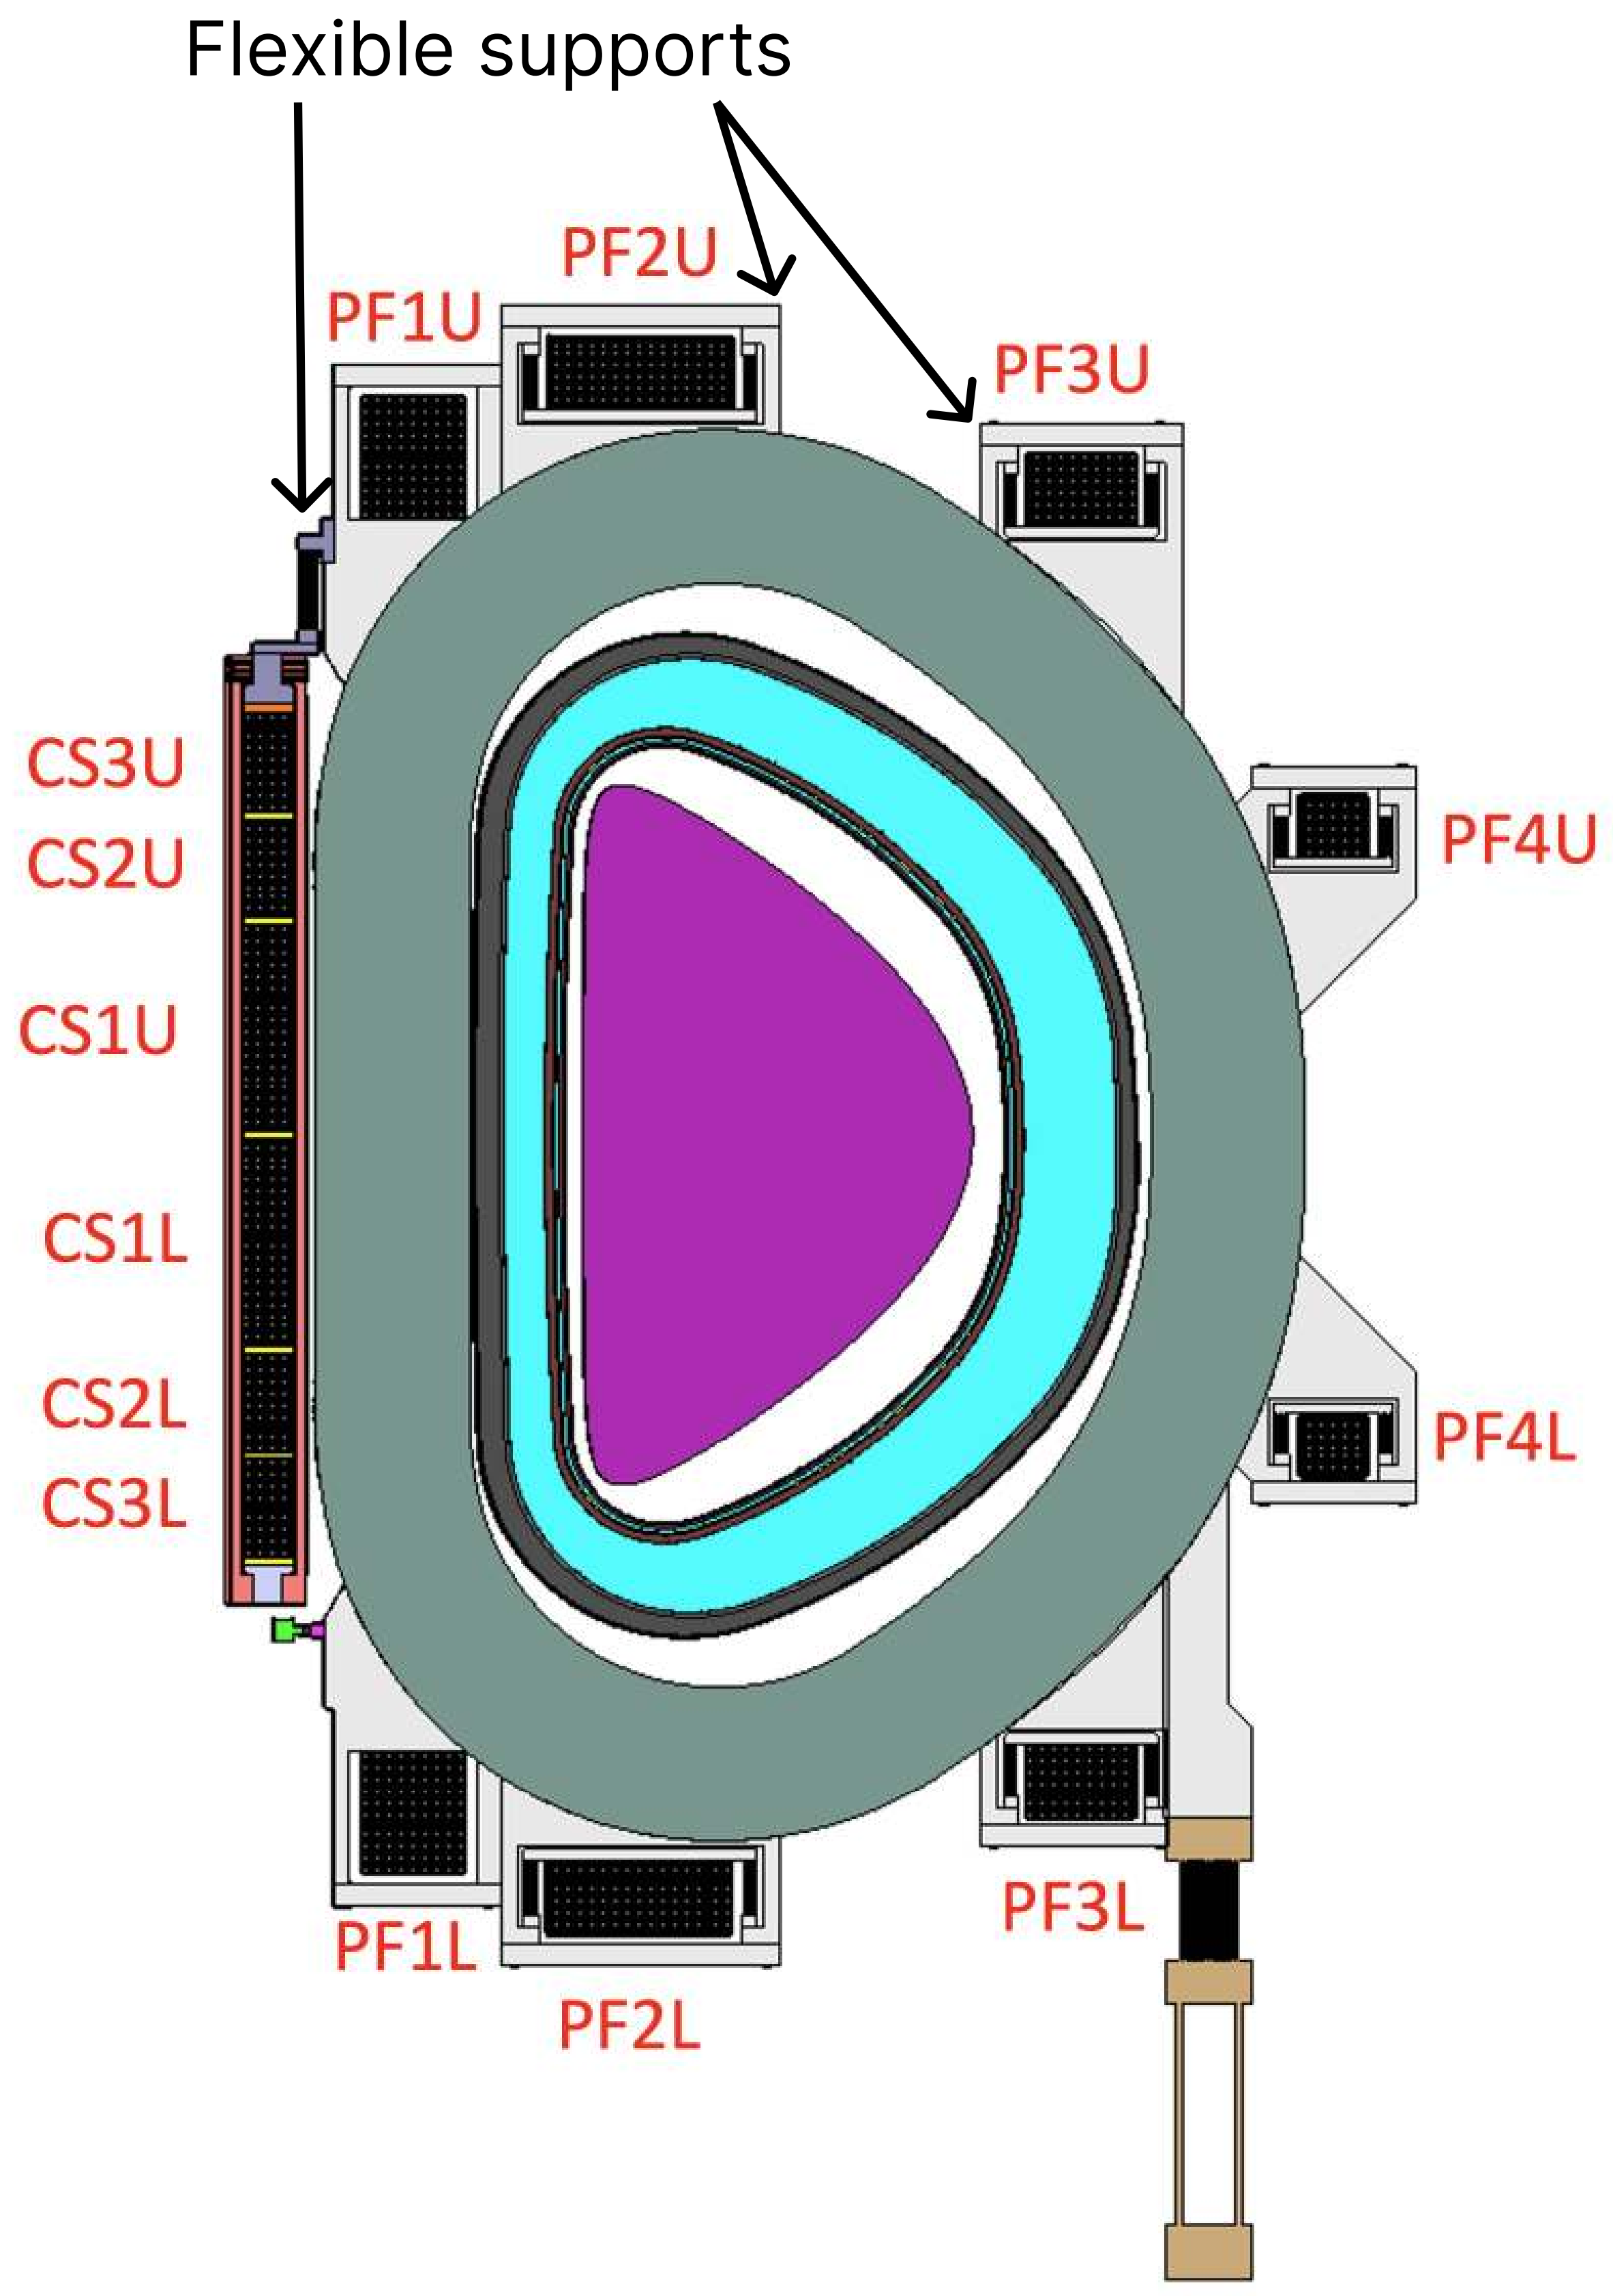

- PF coils are attached to the TF coil cases though welded flanges.
- PF coils are clamped by cover plates after installation on the TF coils.
- At present, it is assumed that PF1 coils can slide in radial direction inside their supports thanks some low-friction interfaces (i.e. Alu-bronze pads).
- The other coils are supported by sets of thin flexible plates which allow for relative radial movements (same as for CS, TF and VV supports).
- For this first pass analysis, a simplified plasma geometry without a divertor was used.

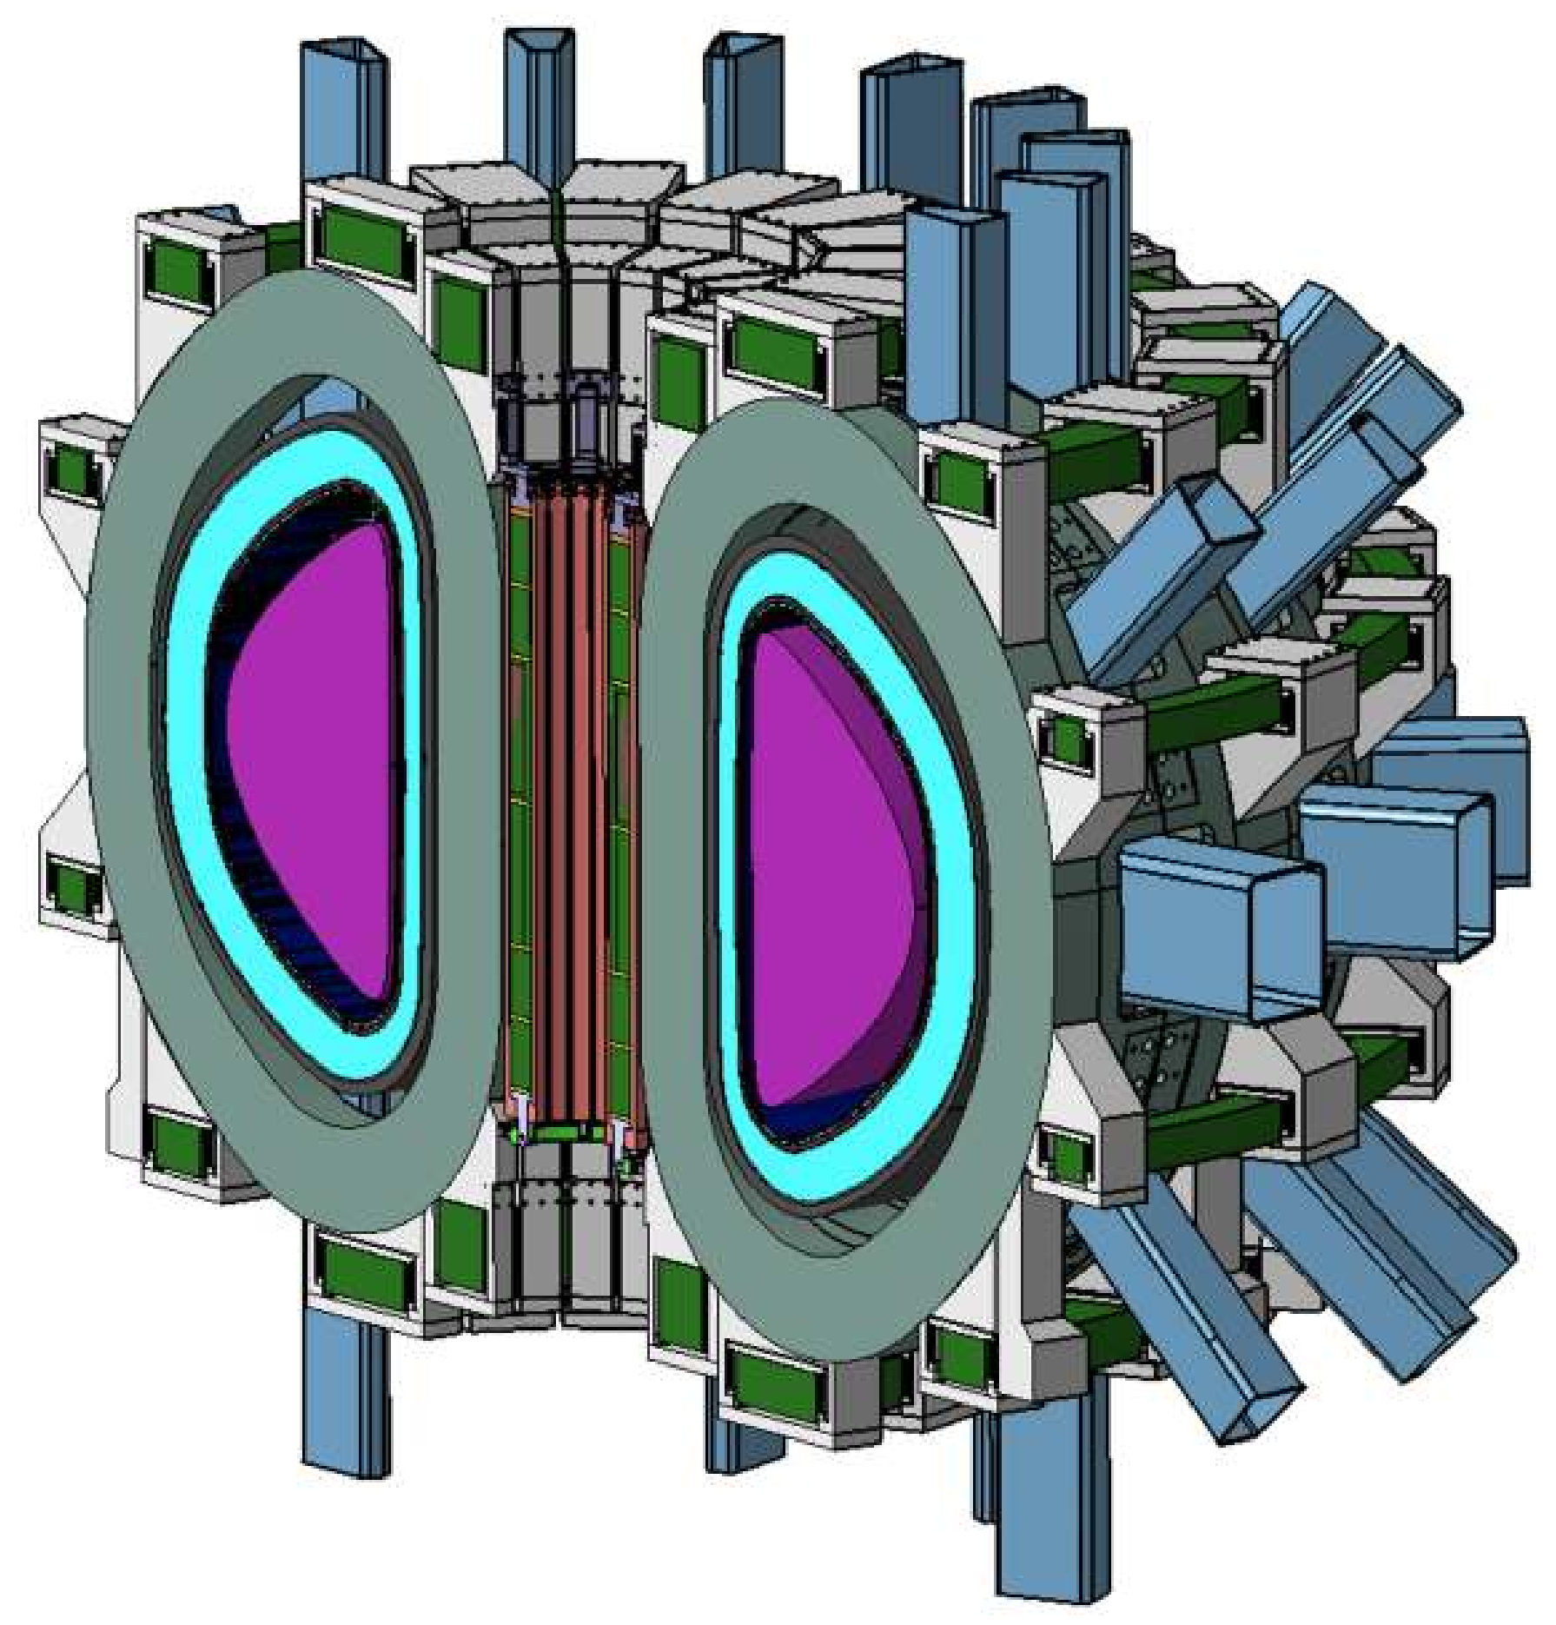

Winding pack current density ~ 68 A/mm²

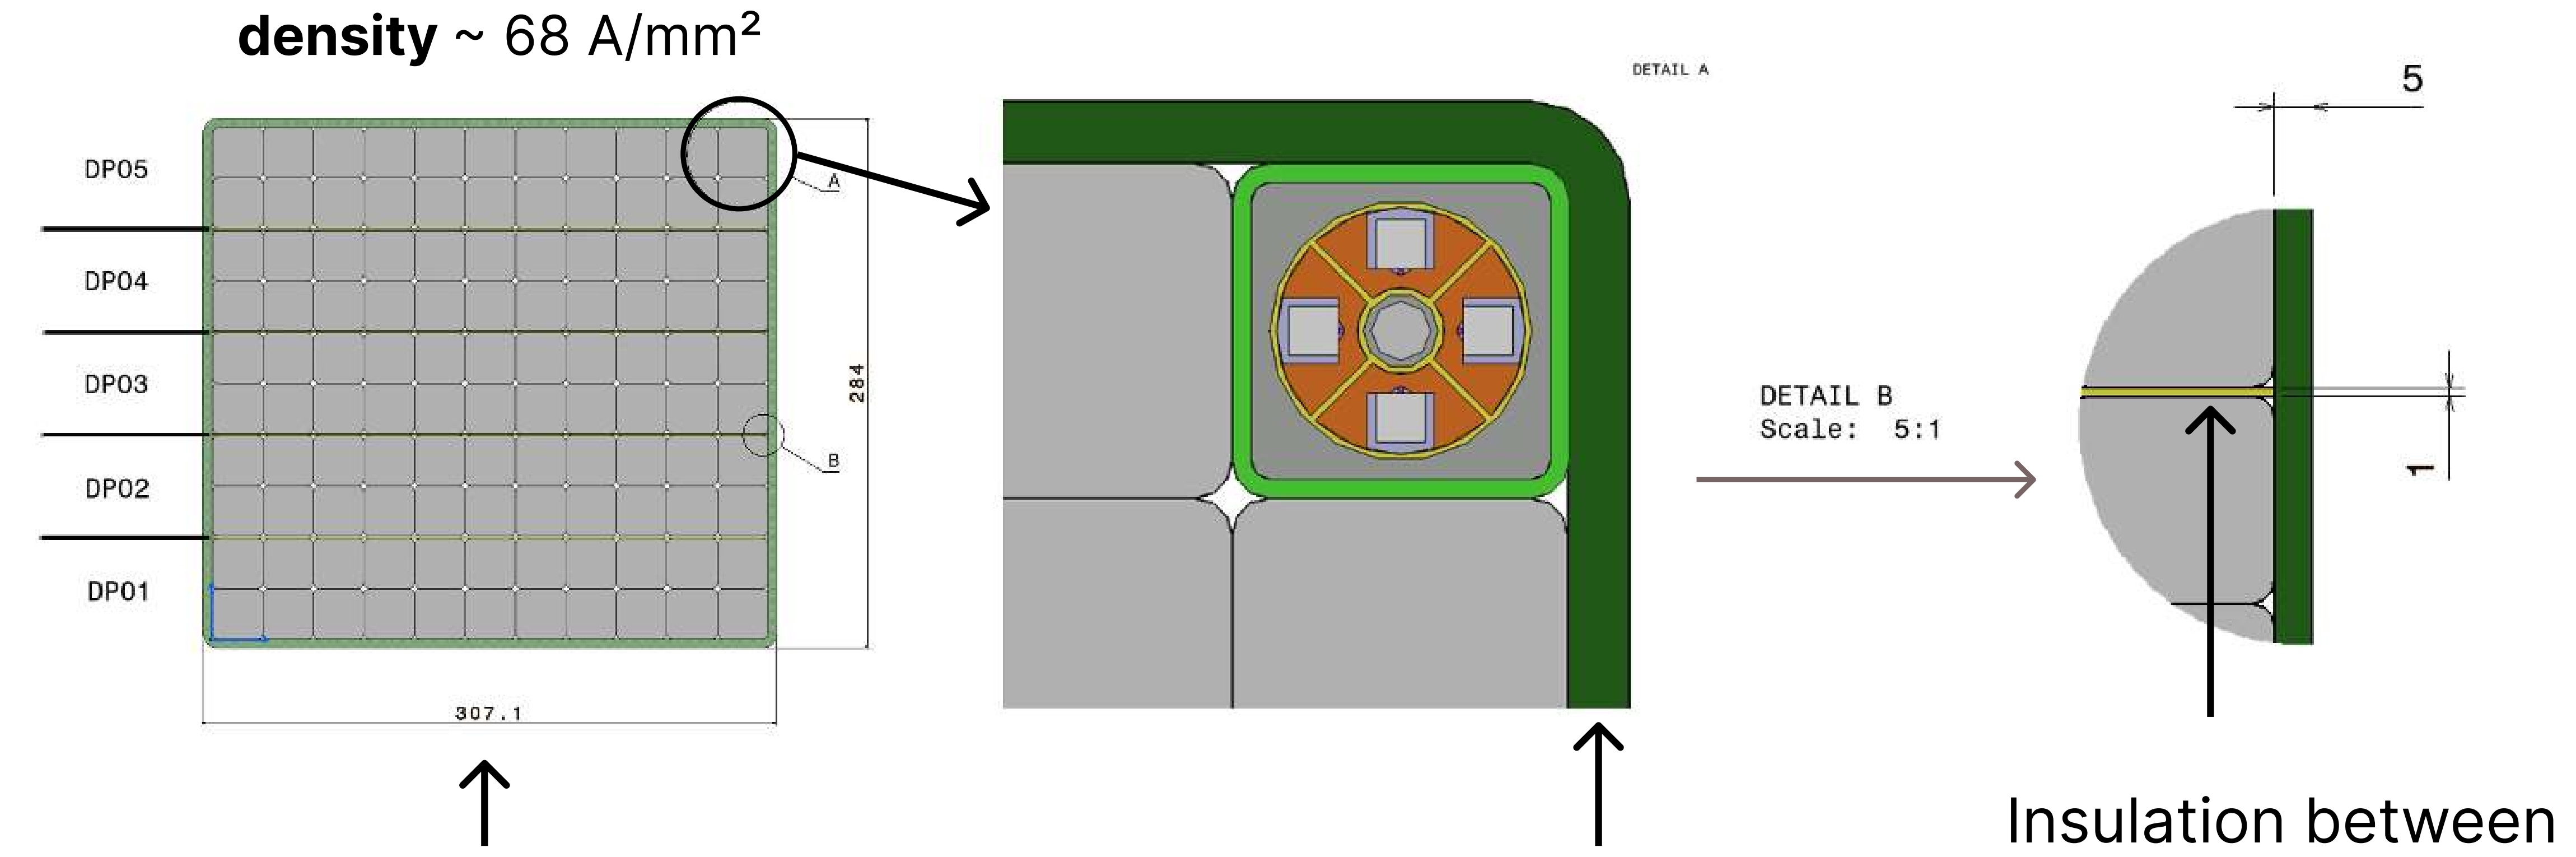

Winding pack with insulated DPs

Ground insulation

Insulation between pancakes

CS + PF Coil Number Conductors / Currents

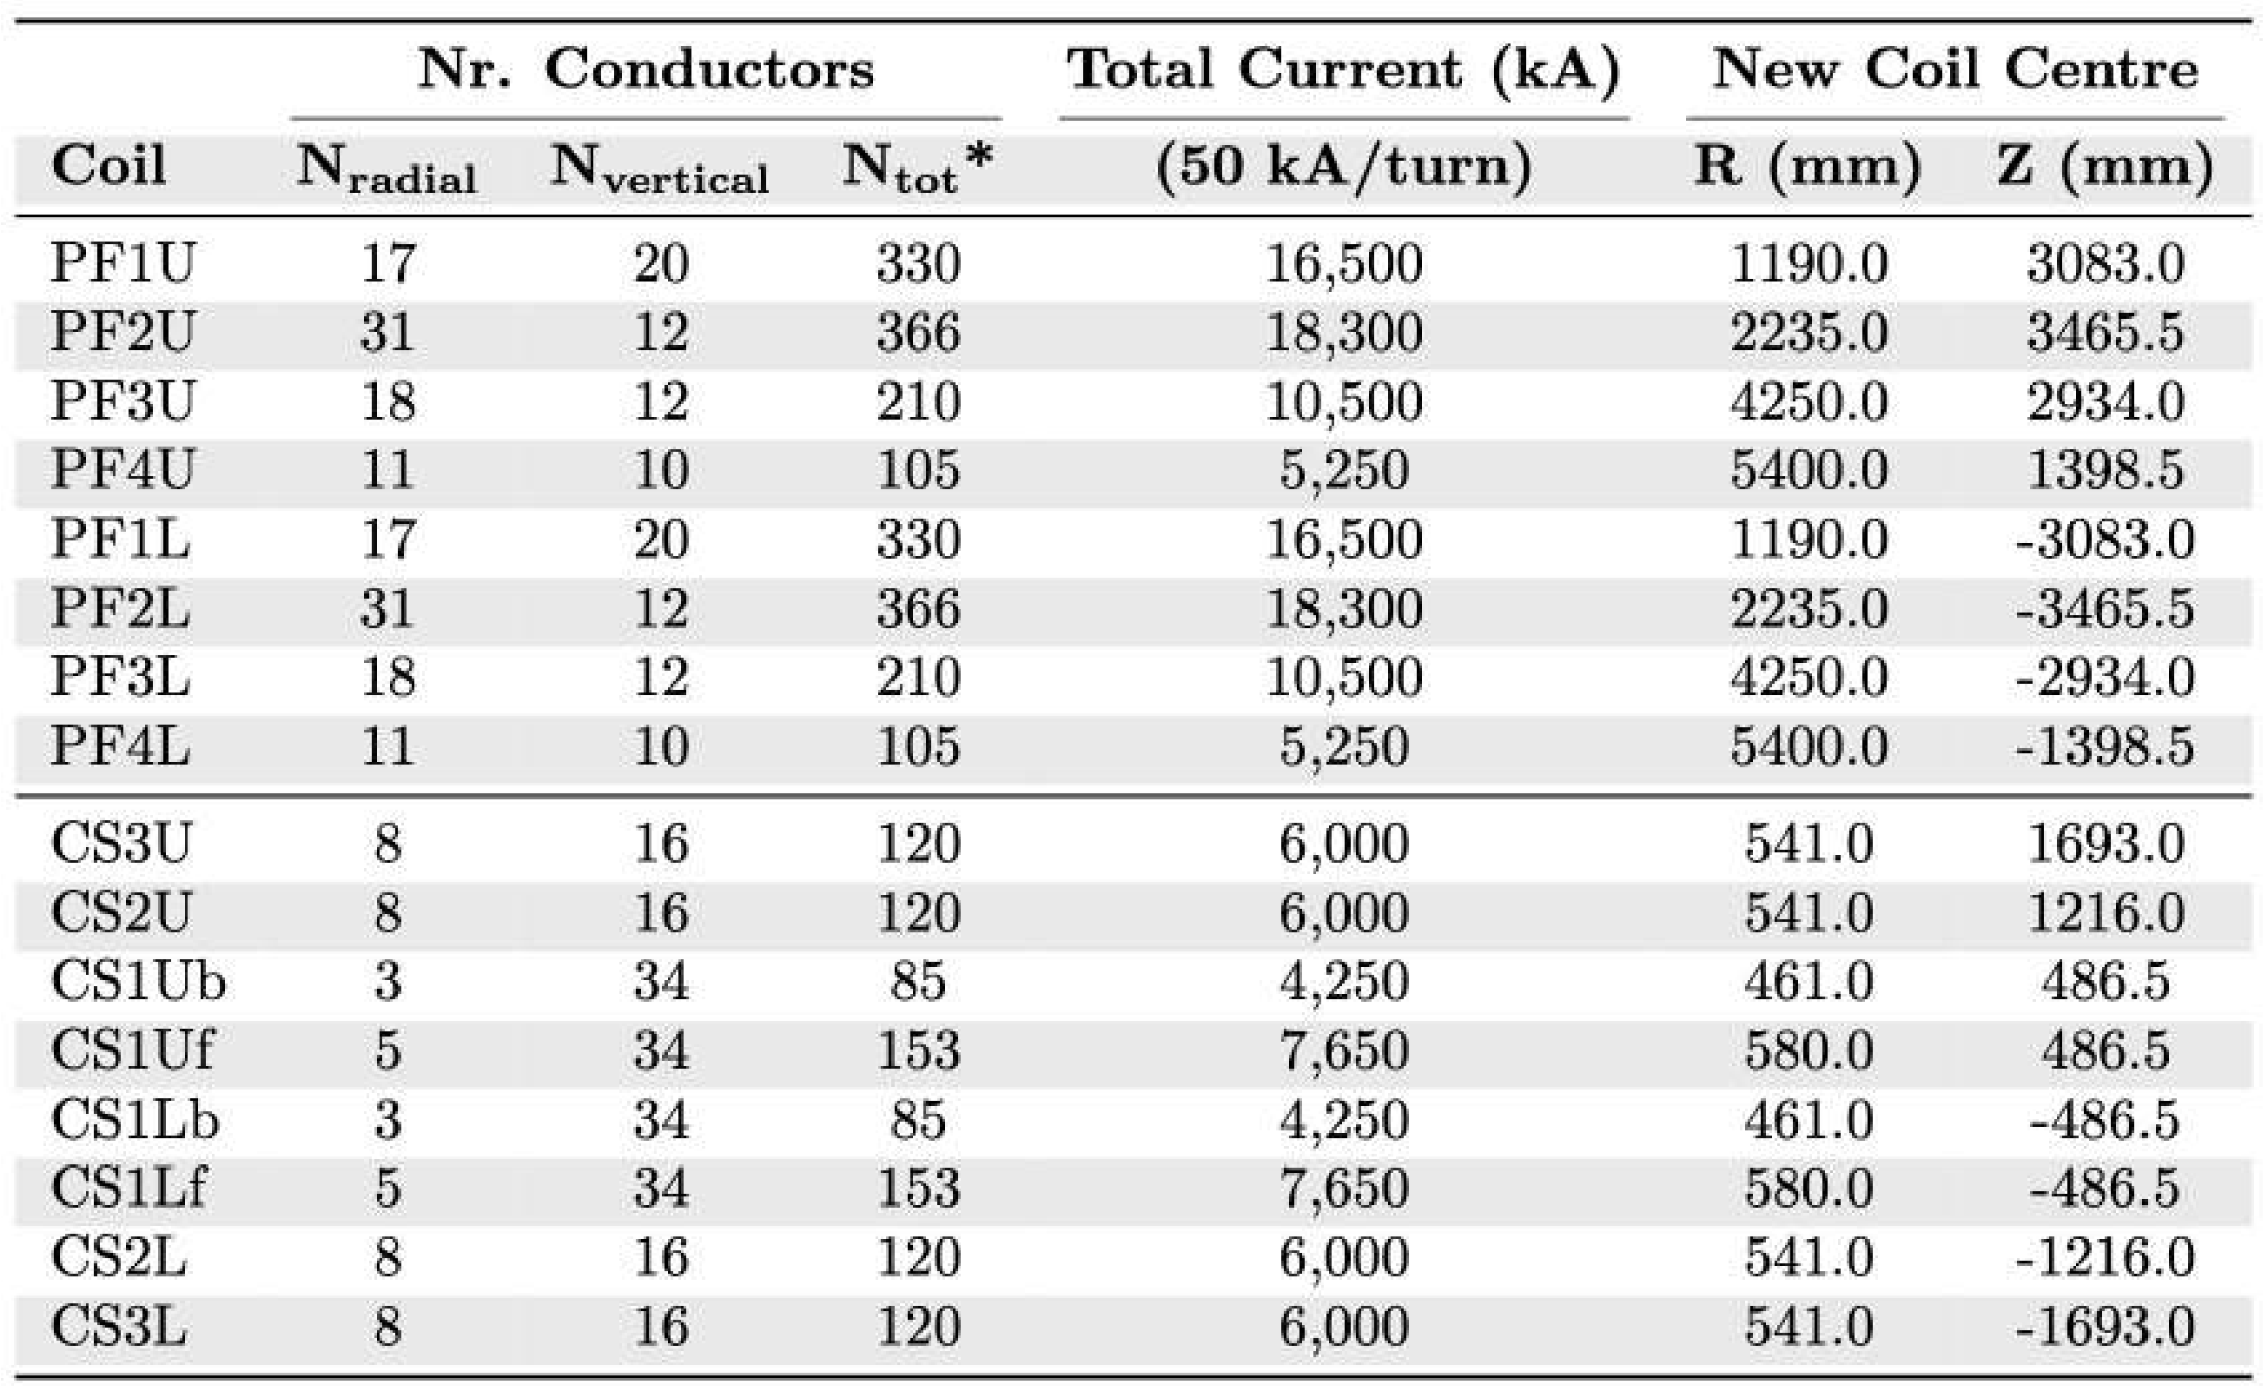

	Nr. Conductors			Total Current (kA)	New Coil Centre	
Coil	N_{radial}	$N_{vertical}$	N_{tot}*	(50 kA/turn)	R (mm)	Z (mm)
PF1U	17	20	330	16,500	1190.0	3083.0
PF2U	31	12	366	18,300	2235.0	3465.5
PF3U	18	12	210	10,500	4250.0	2934.0
PF4U	11	10	105	5,250	5400.0	1398.5
PF1L	17	20	330	16,500	1190.0	-3083.0
PF2L	31	12	366	18,300	2235.0	-3465.5
PF3L	18	12	210	10,500	4250.0	-2934.0
PF4L	11	10	105	5,250	5400.0	-1398.5
CS3U	8	16	120	6,000	541.0	1693.0
CS2U	8	16	120	6,000	541.0	1216.0
CS1Ub	3	34	85	4,250	461.0	486.5
CS1Uf	5	34	153	7,650	580.0	486.5
CS1Lb	3	34	85	4,250	461.0	-486.5
CS1Lf	5	34	153	7,650	580.0	-486.5
CS2L	8	16	120	6,000	541.0	-1216.0
CS3L	8	16	120	6,000	541.0	-1693.0

- **Gravity supports:** The central solenoid is hung on 8 flexible vertical supports attached to the PF1U support flanges: they take the weight of the whole CS stack and its structures and support net vertical or lateral loads due to electromagnetic forces or seismic events, while supporting relative radial movements.
- **Pre-compression:** The central solenoid is pre-compressed vertically by 16 sets of tie-plates located on its inner and outer surface: these tie-plates can be installed on upper and lower flanges by thermal shrink-fitting to provide pre-compression.
- **Radial restrainers** are provided at the bottom of the CS stack to centre the assembly: they will be equipped with spring jacks to allow for relative radial movements.
- CS coils are separated by **interface plates** which can allow for relative radial displacements.
- Separate (electrical-cryogenic) **feeder lines** will feed each coil in the stack.

Central Solenoid Mechanical Structures

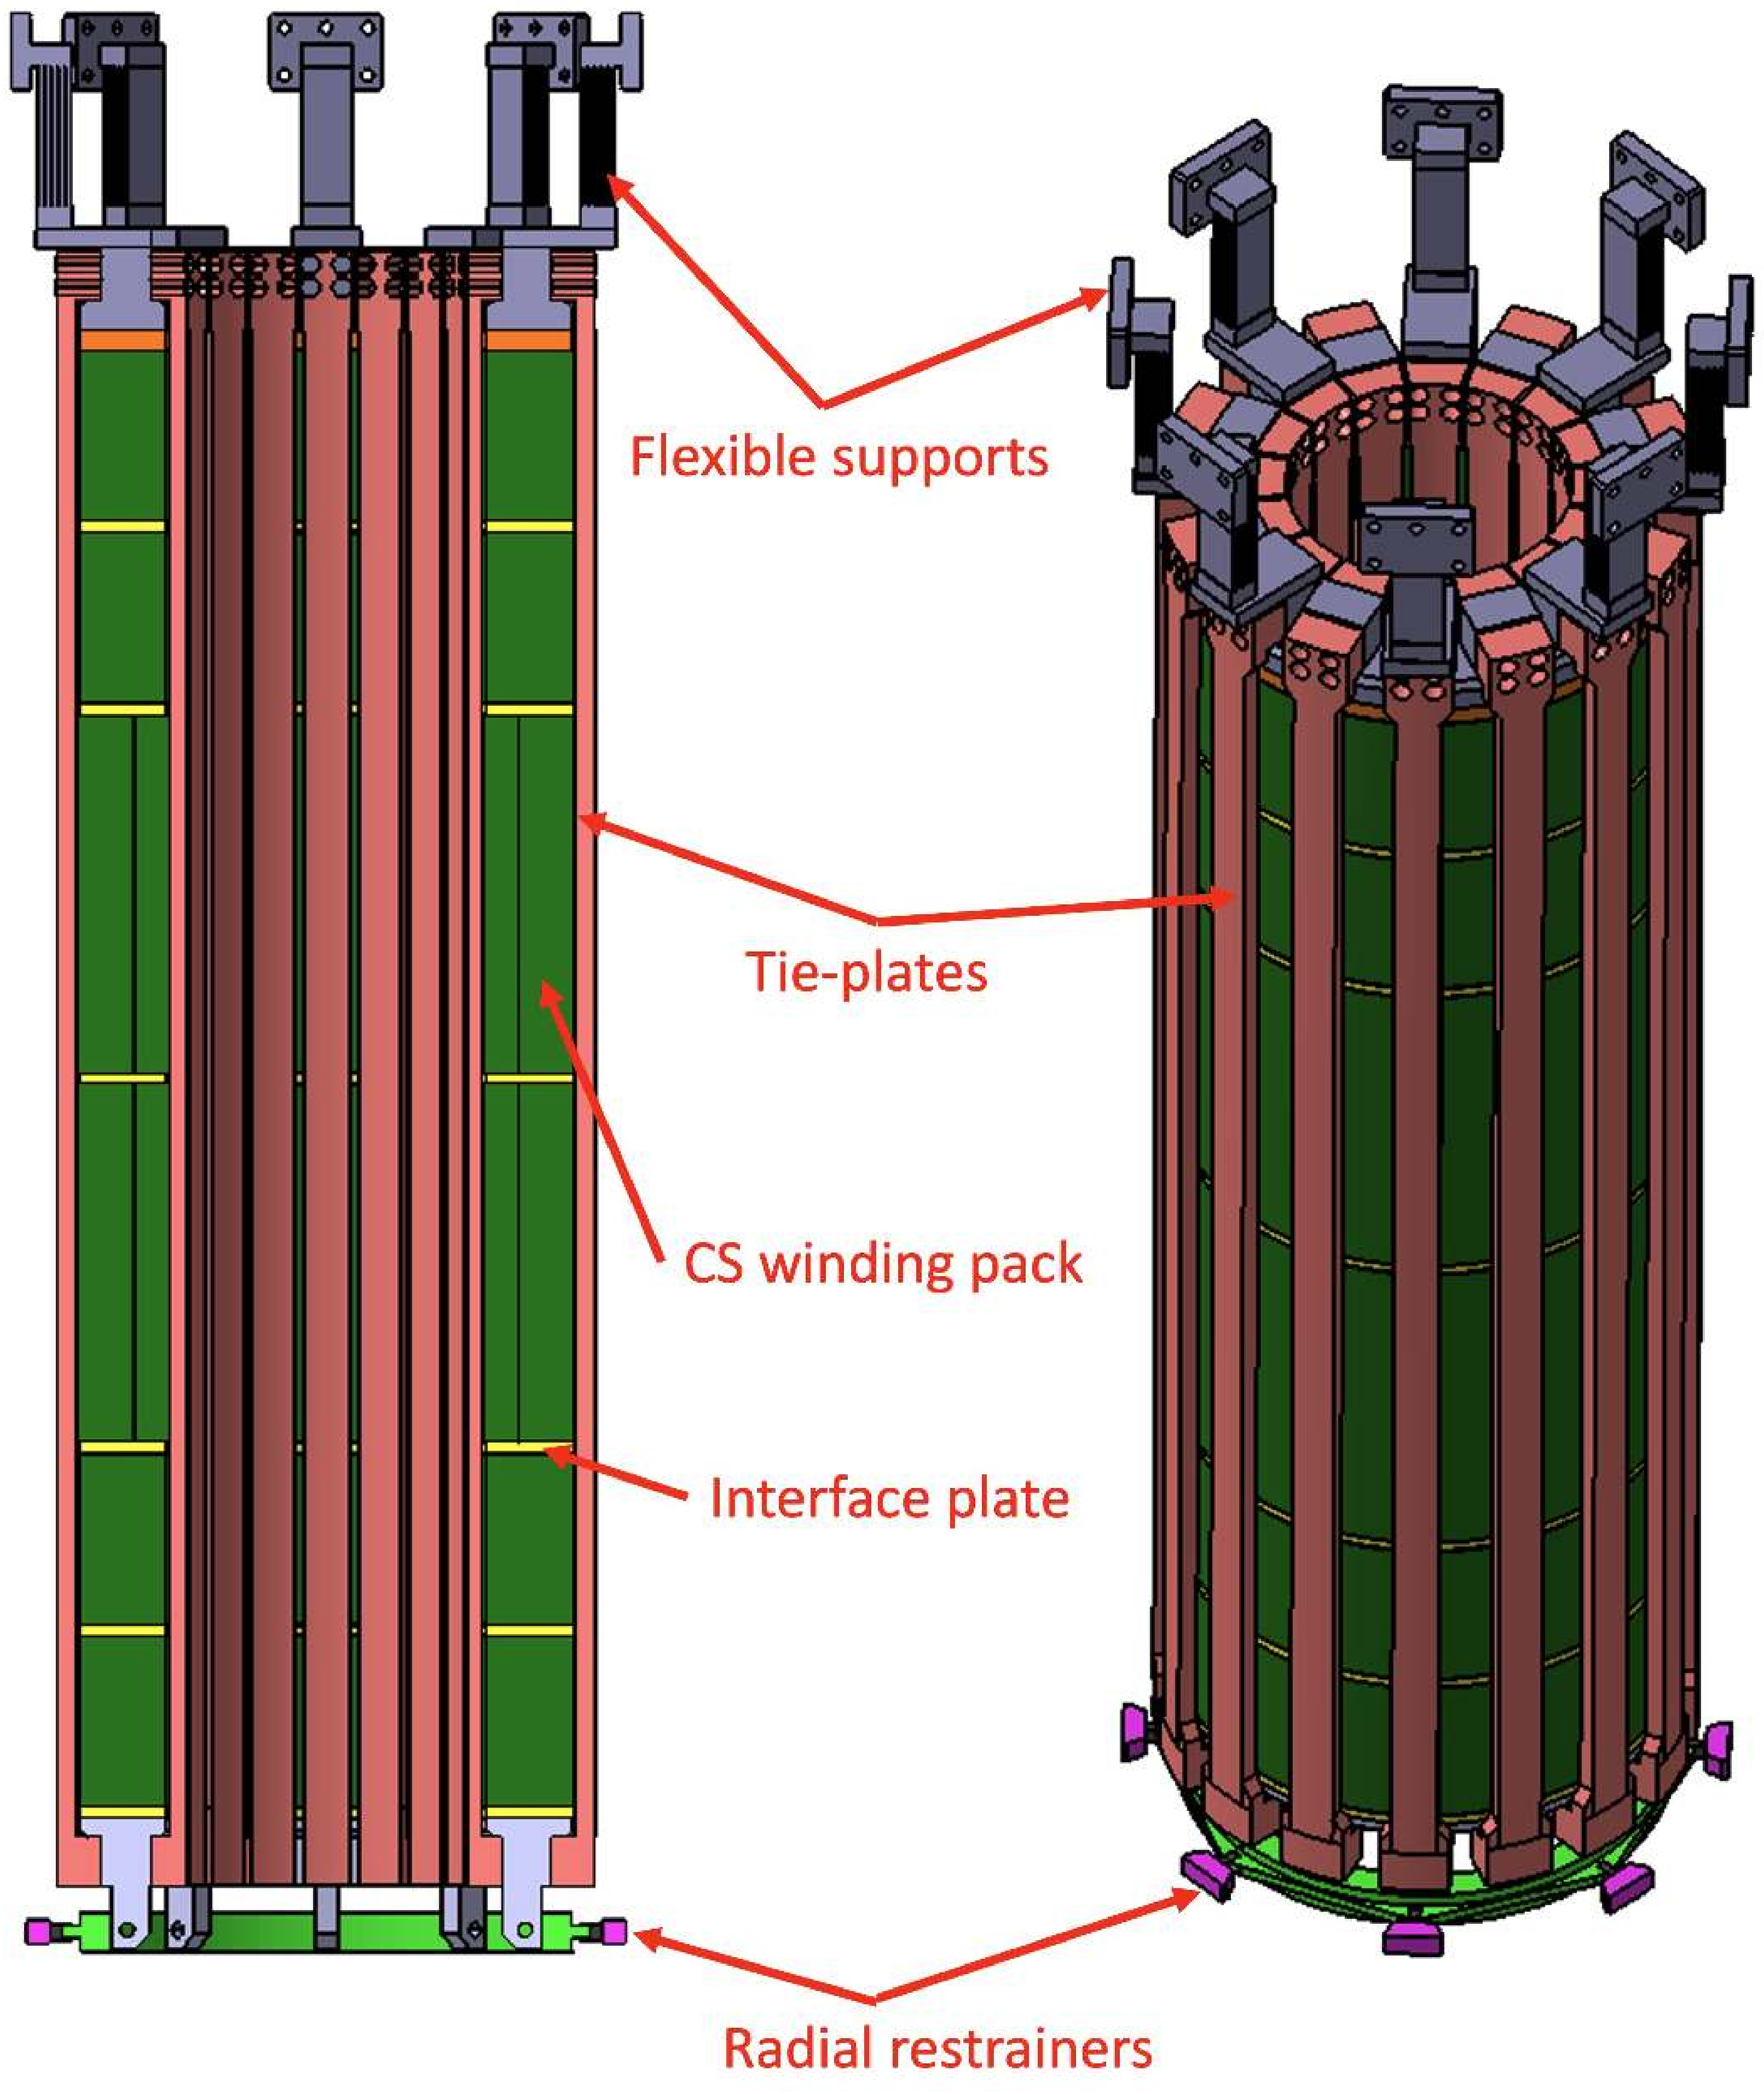

Structural Analysis

Electromagnetic Loads on the Coils

In an HTS tokamak, millions of amps flow through the coils, producing intense magnetic fields. These fields induce significant electromagnetic forces on the coil structures. Accurately determining these structural loads is essential to ensure mechanical integrity and safe reactor operation. By conducting detailed structural analyses, we evaluate stresses, deformation, and fatigue on coil components. The model setup is shown below with the current applied and the current density in all coils + the plasma.

TFC Electromagnetic Model

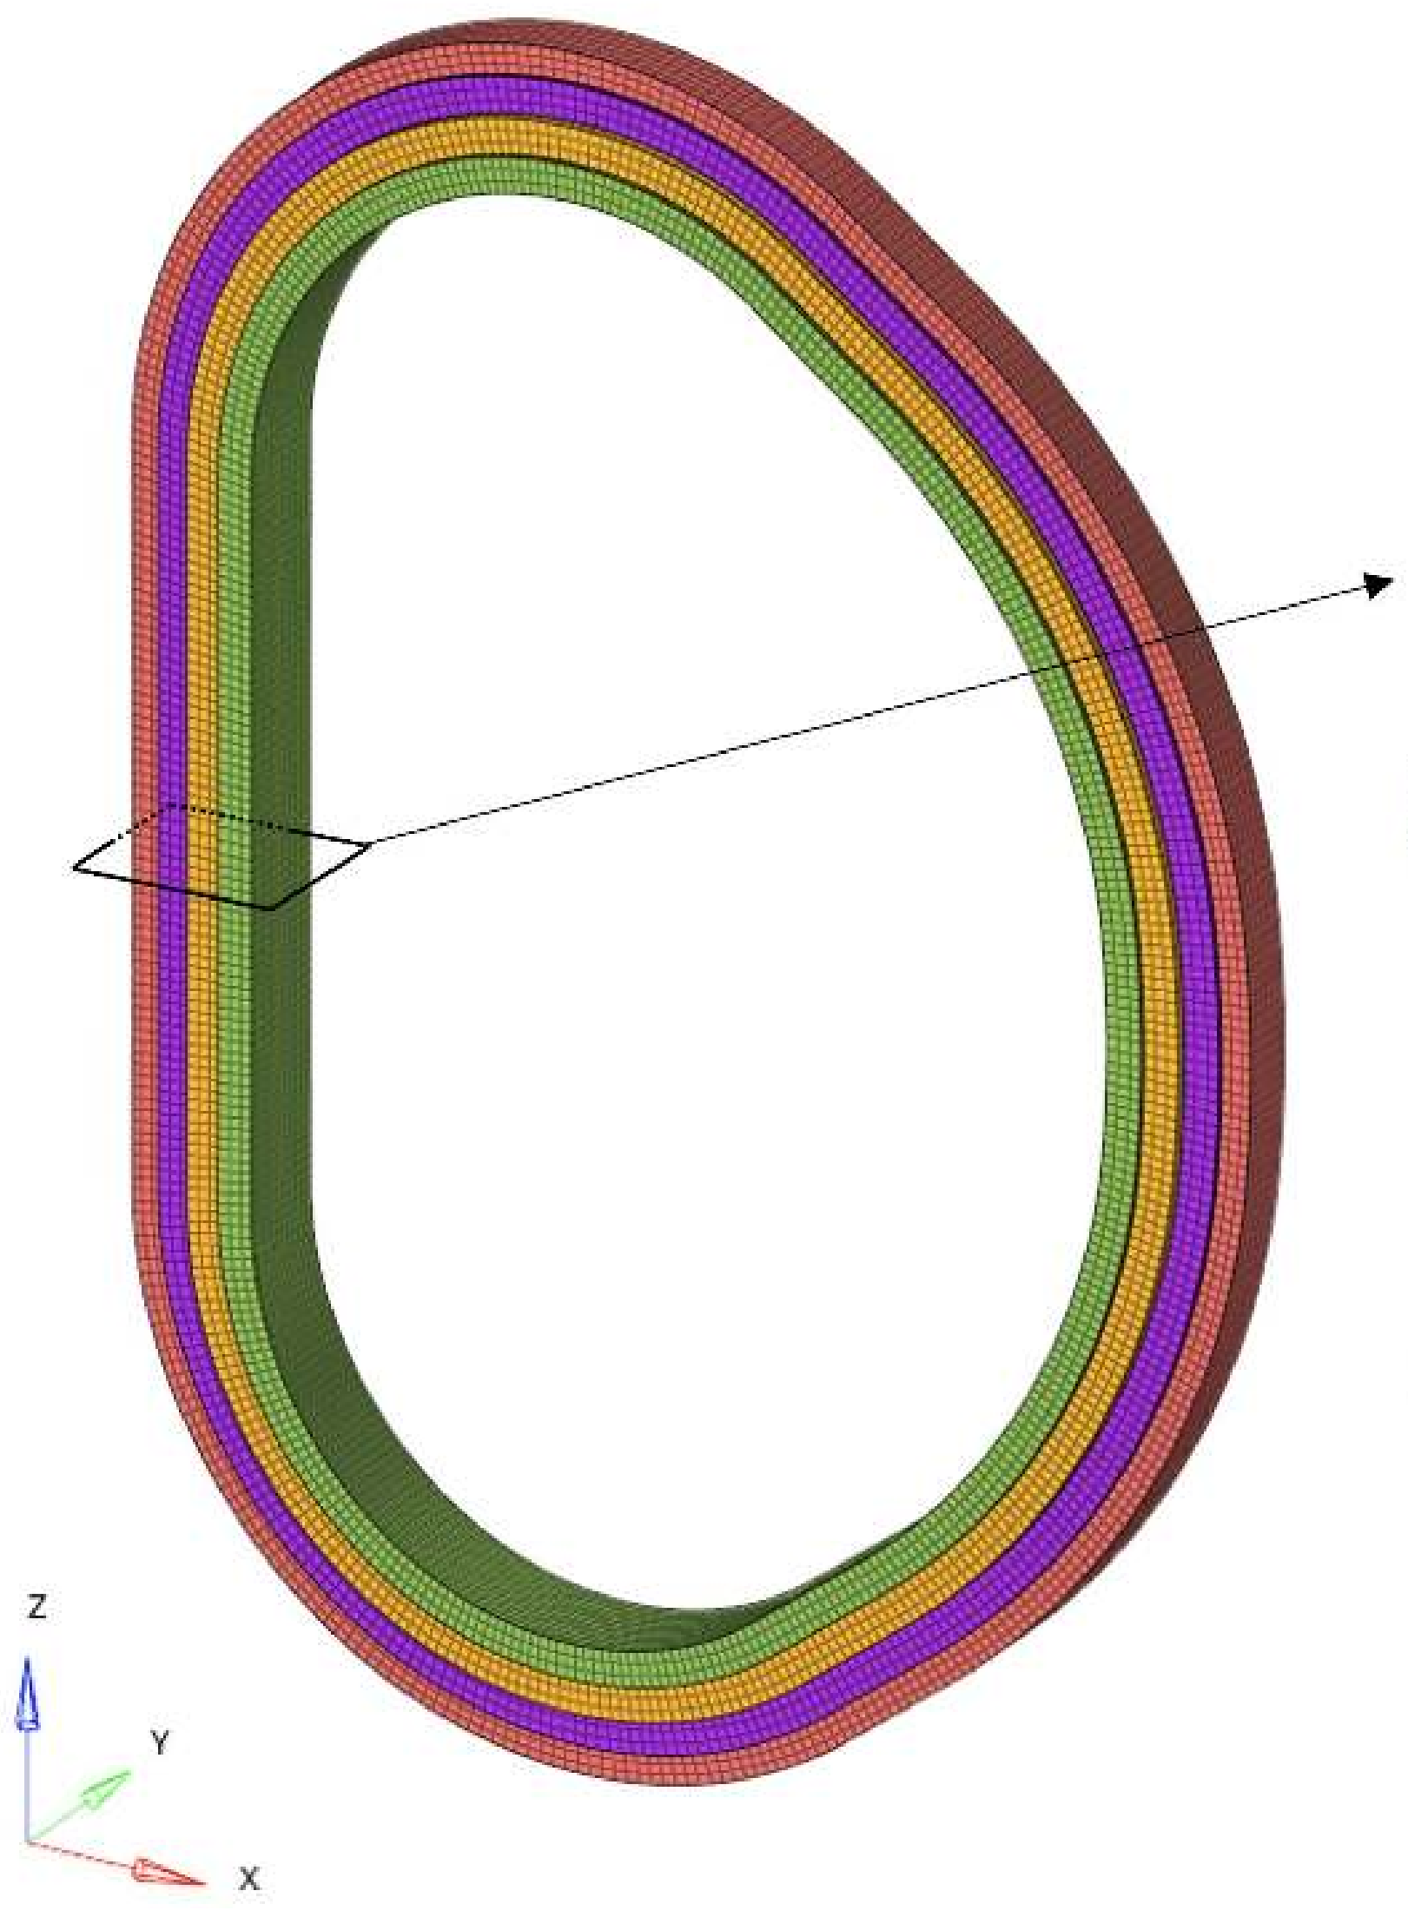

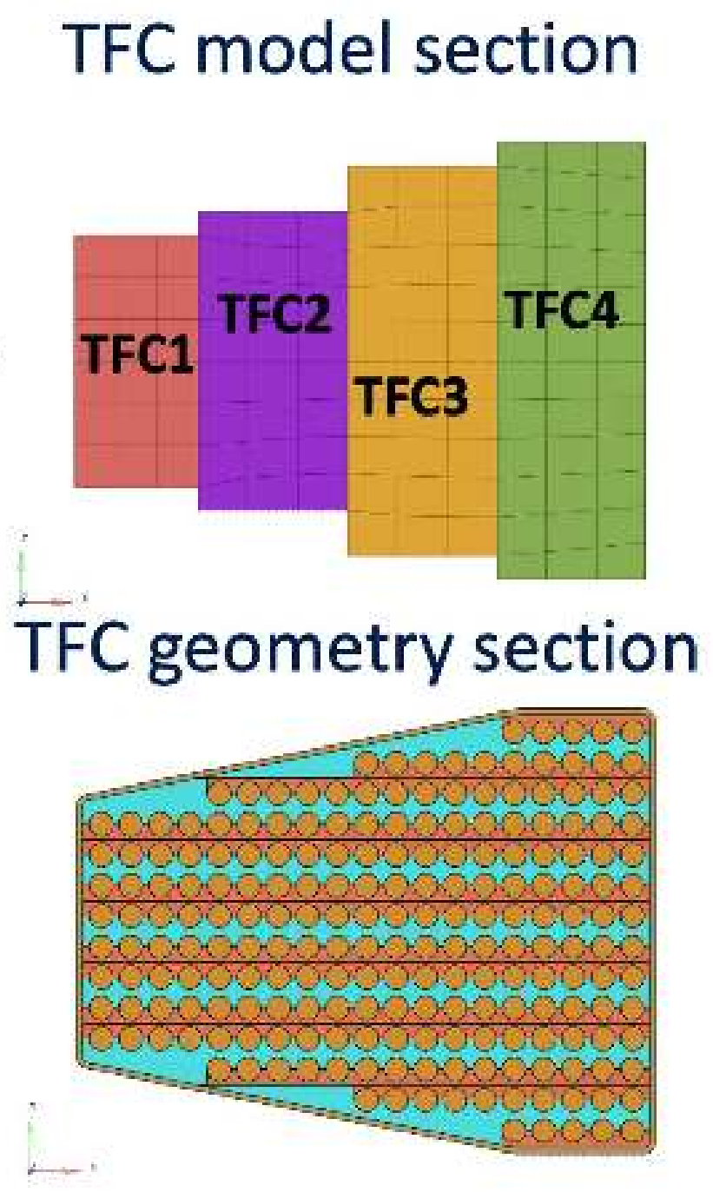

Comp	I (MA)
TFC1	2.553
TFC2	2.294
TFC3	1.754
TFC4	1.237

- One sector model (22.5 degrees)
- Cyclic boundary conditions
- Current applied to TFC to have the uniform current density distribution

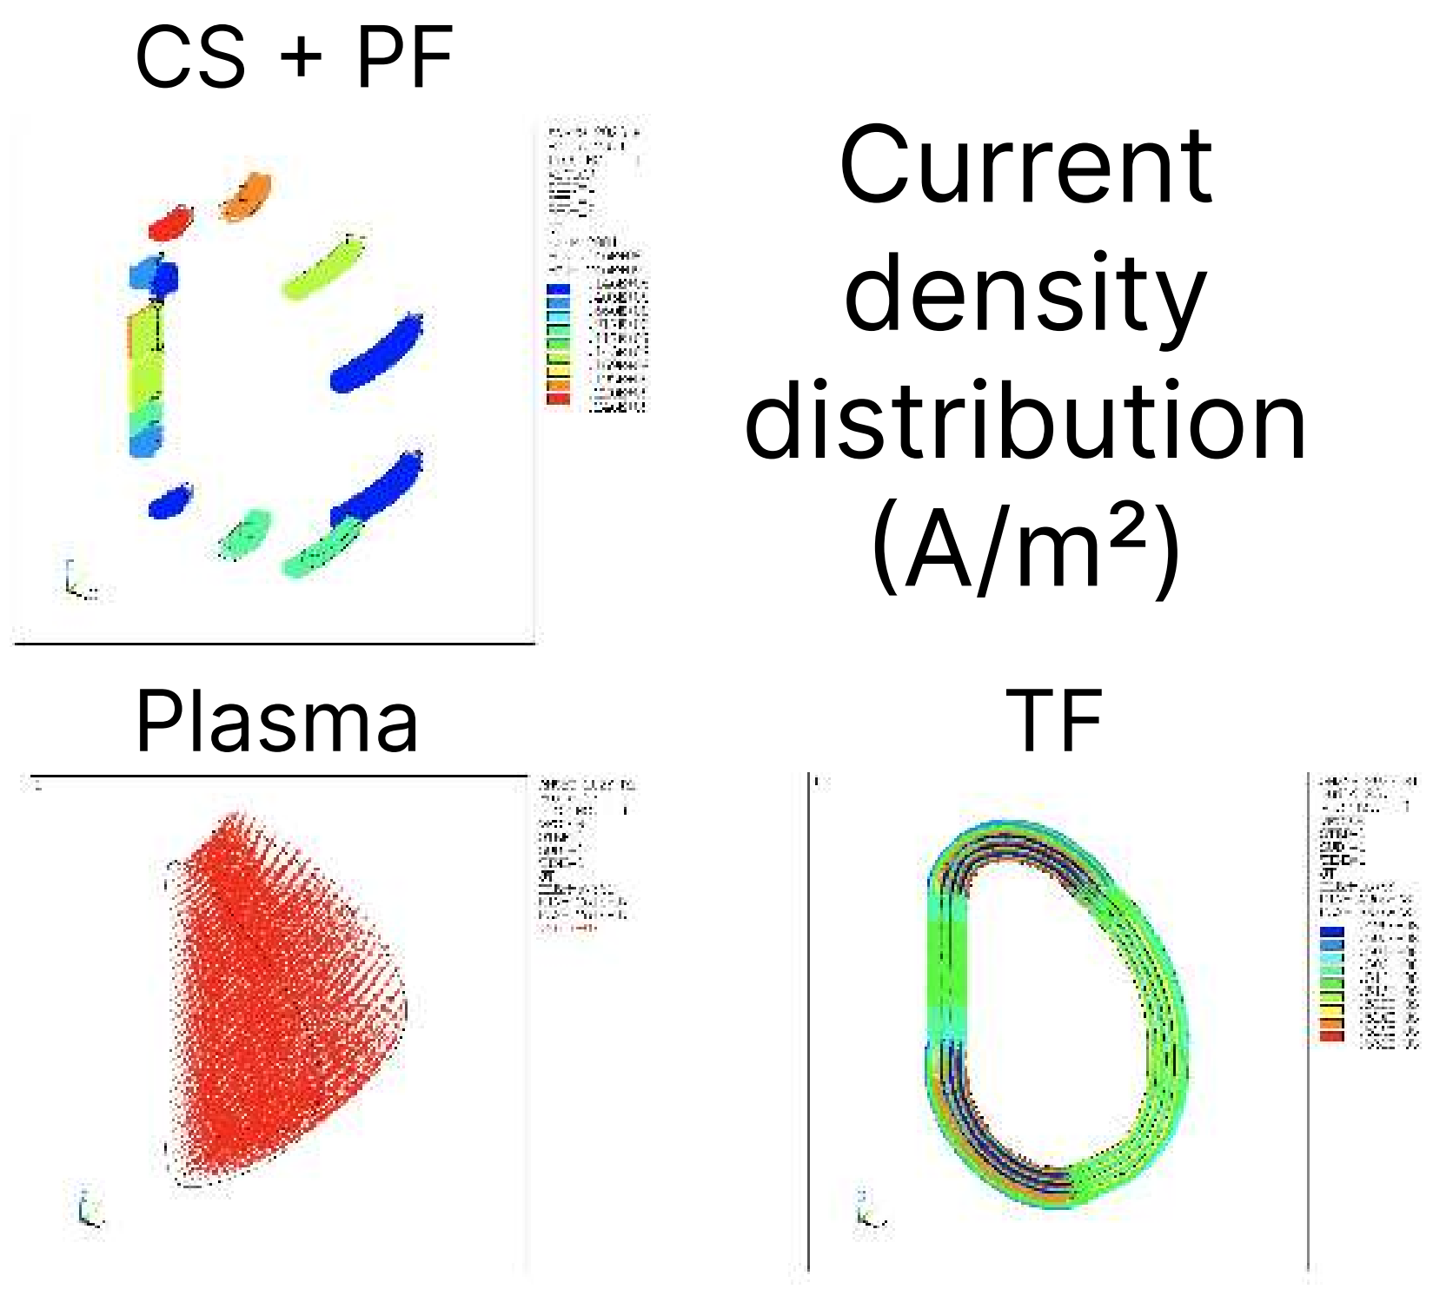

Current density distribution (A/m²)

All Coils + Plasma Model

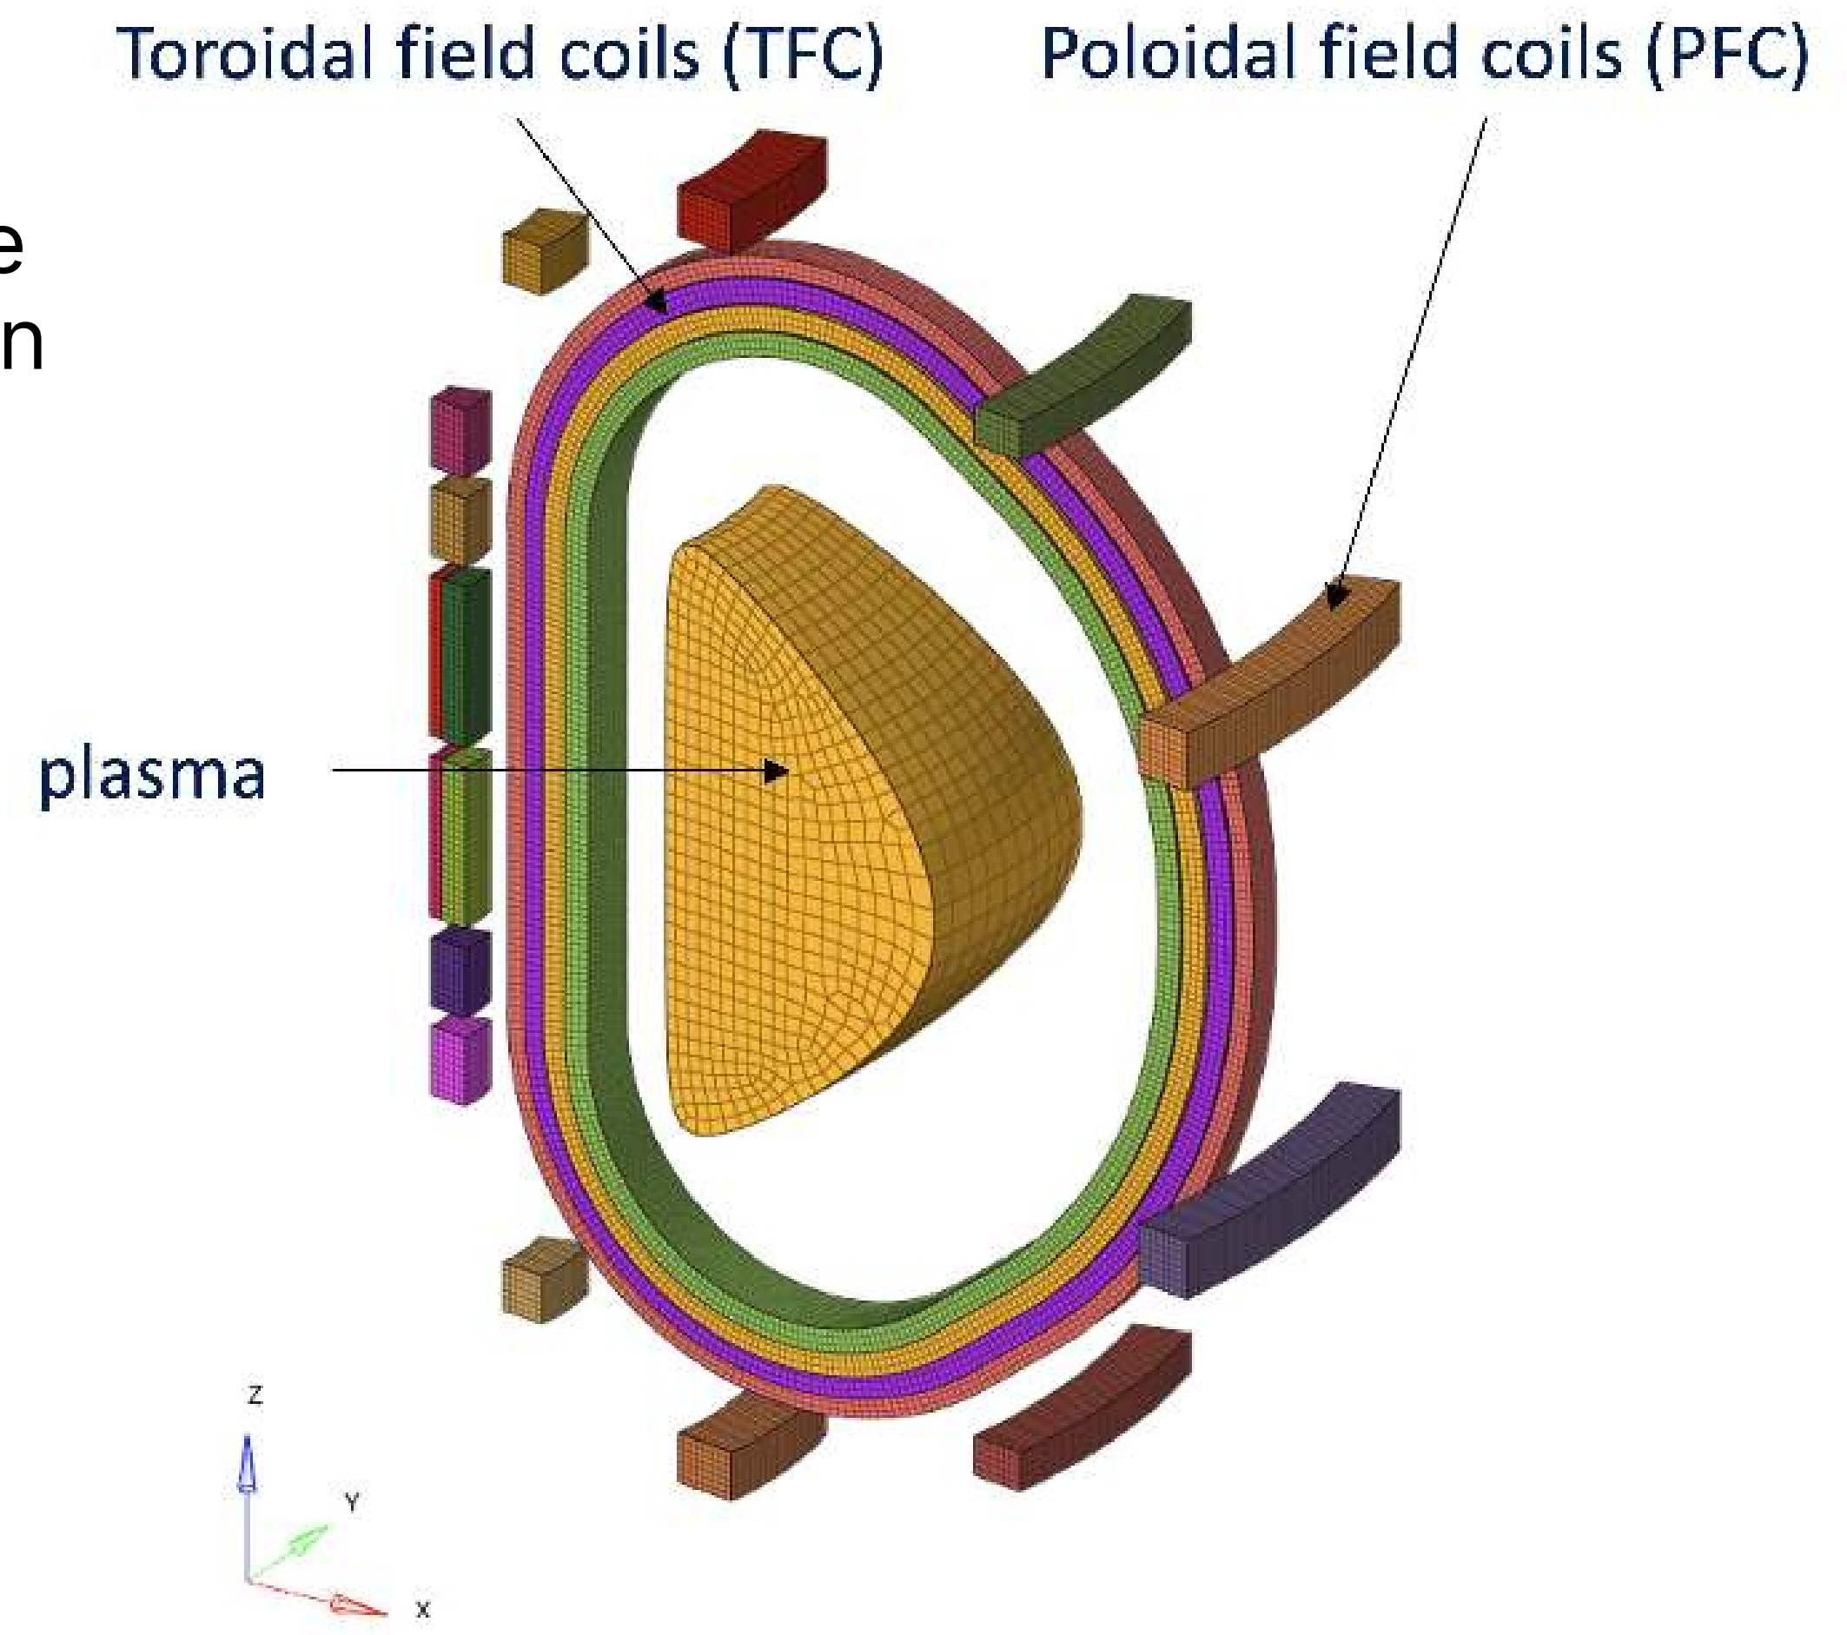

Comp	I (MA)
Plasma	10.315
TFC	7.8375
CS3U	-4.772
CS2U	1.262
CS1UB	-12.636
CS1UF	-19.832
CS1LB	-11.340
CS1LF	-19.433
CS2L	-7.671
CS3L	-4.503
PF1U	-16.435
PF2U	18.400
PF3U	-10.783
PF4U	-1.458
PF1L	-2.412
PF2L	9.293
PF3L	-6.134
PF4L	-4.014

Currents applied to PFC+CS+plasma

After running the simulation in ANSYS, we can determine the magnetic fields at all points in the device. Looking at the images below, we can see that the TFC sees peak ~18.5T on coil and hits our targeted **~8.5T on axis** field. Also, importantly, we can see how quickly the magnetic field drops off outside the reactors, going back to almost 0T within less than <0.5m of the edge of of the device.

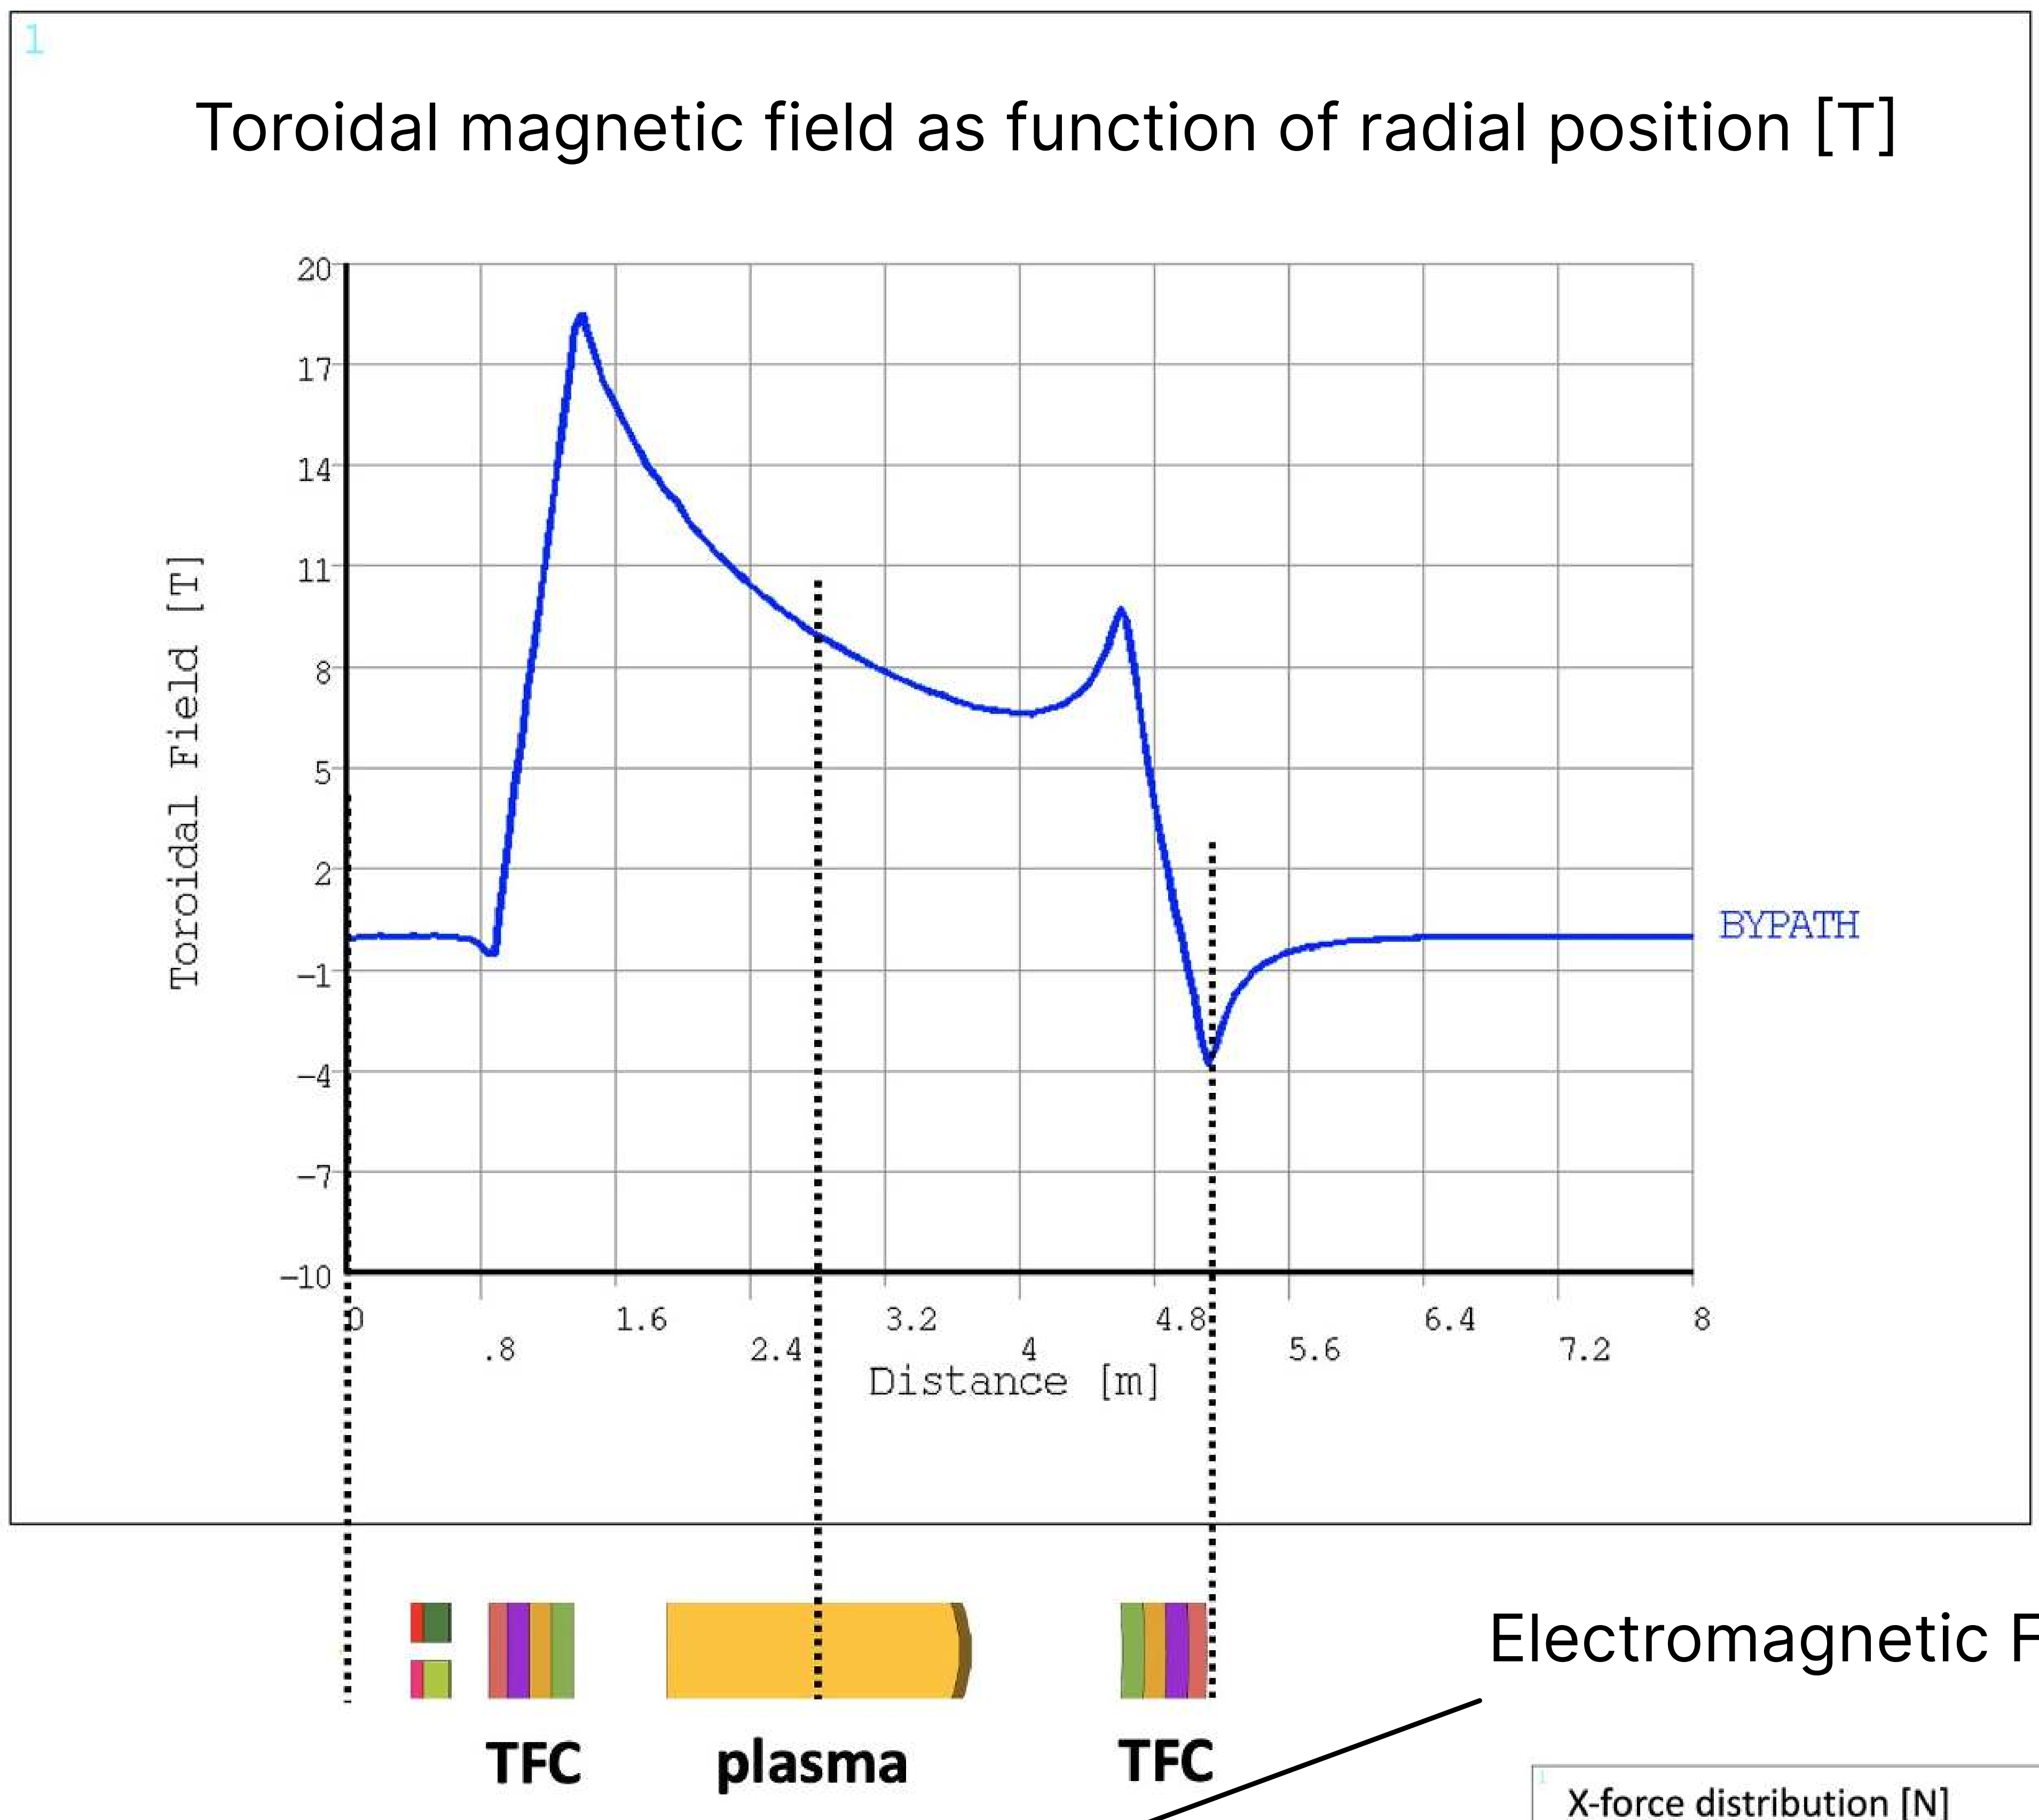

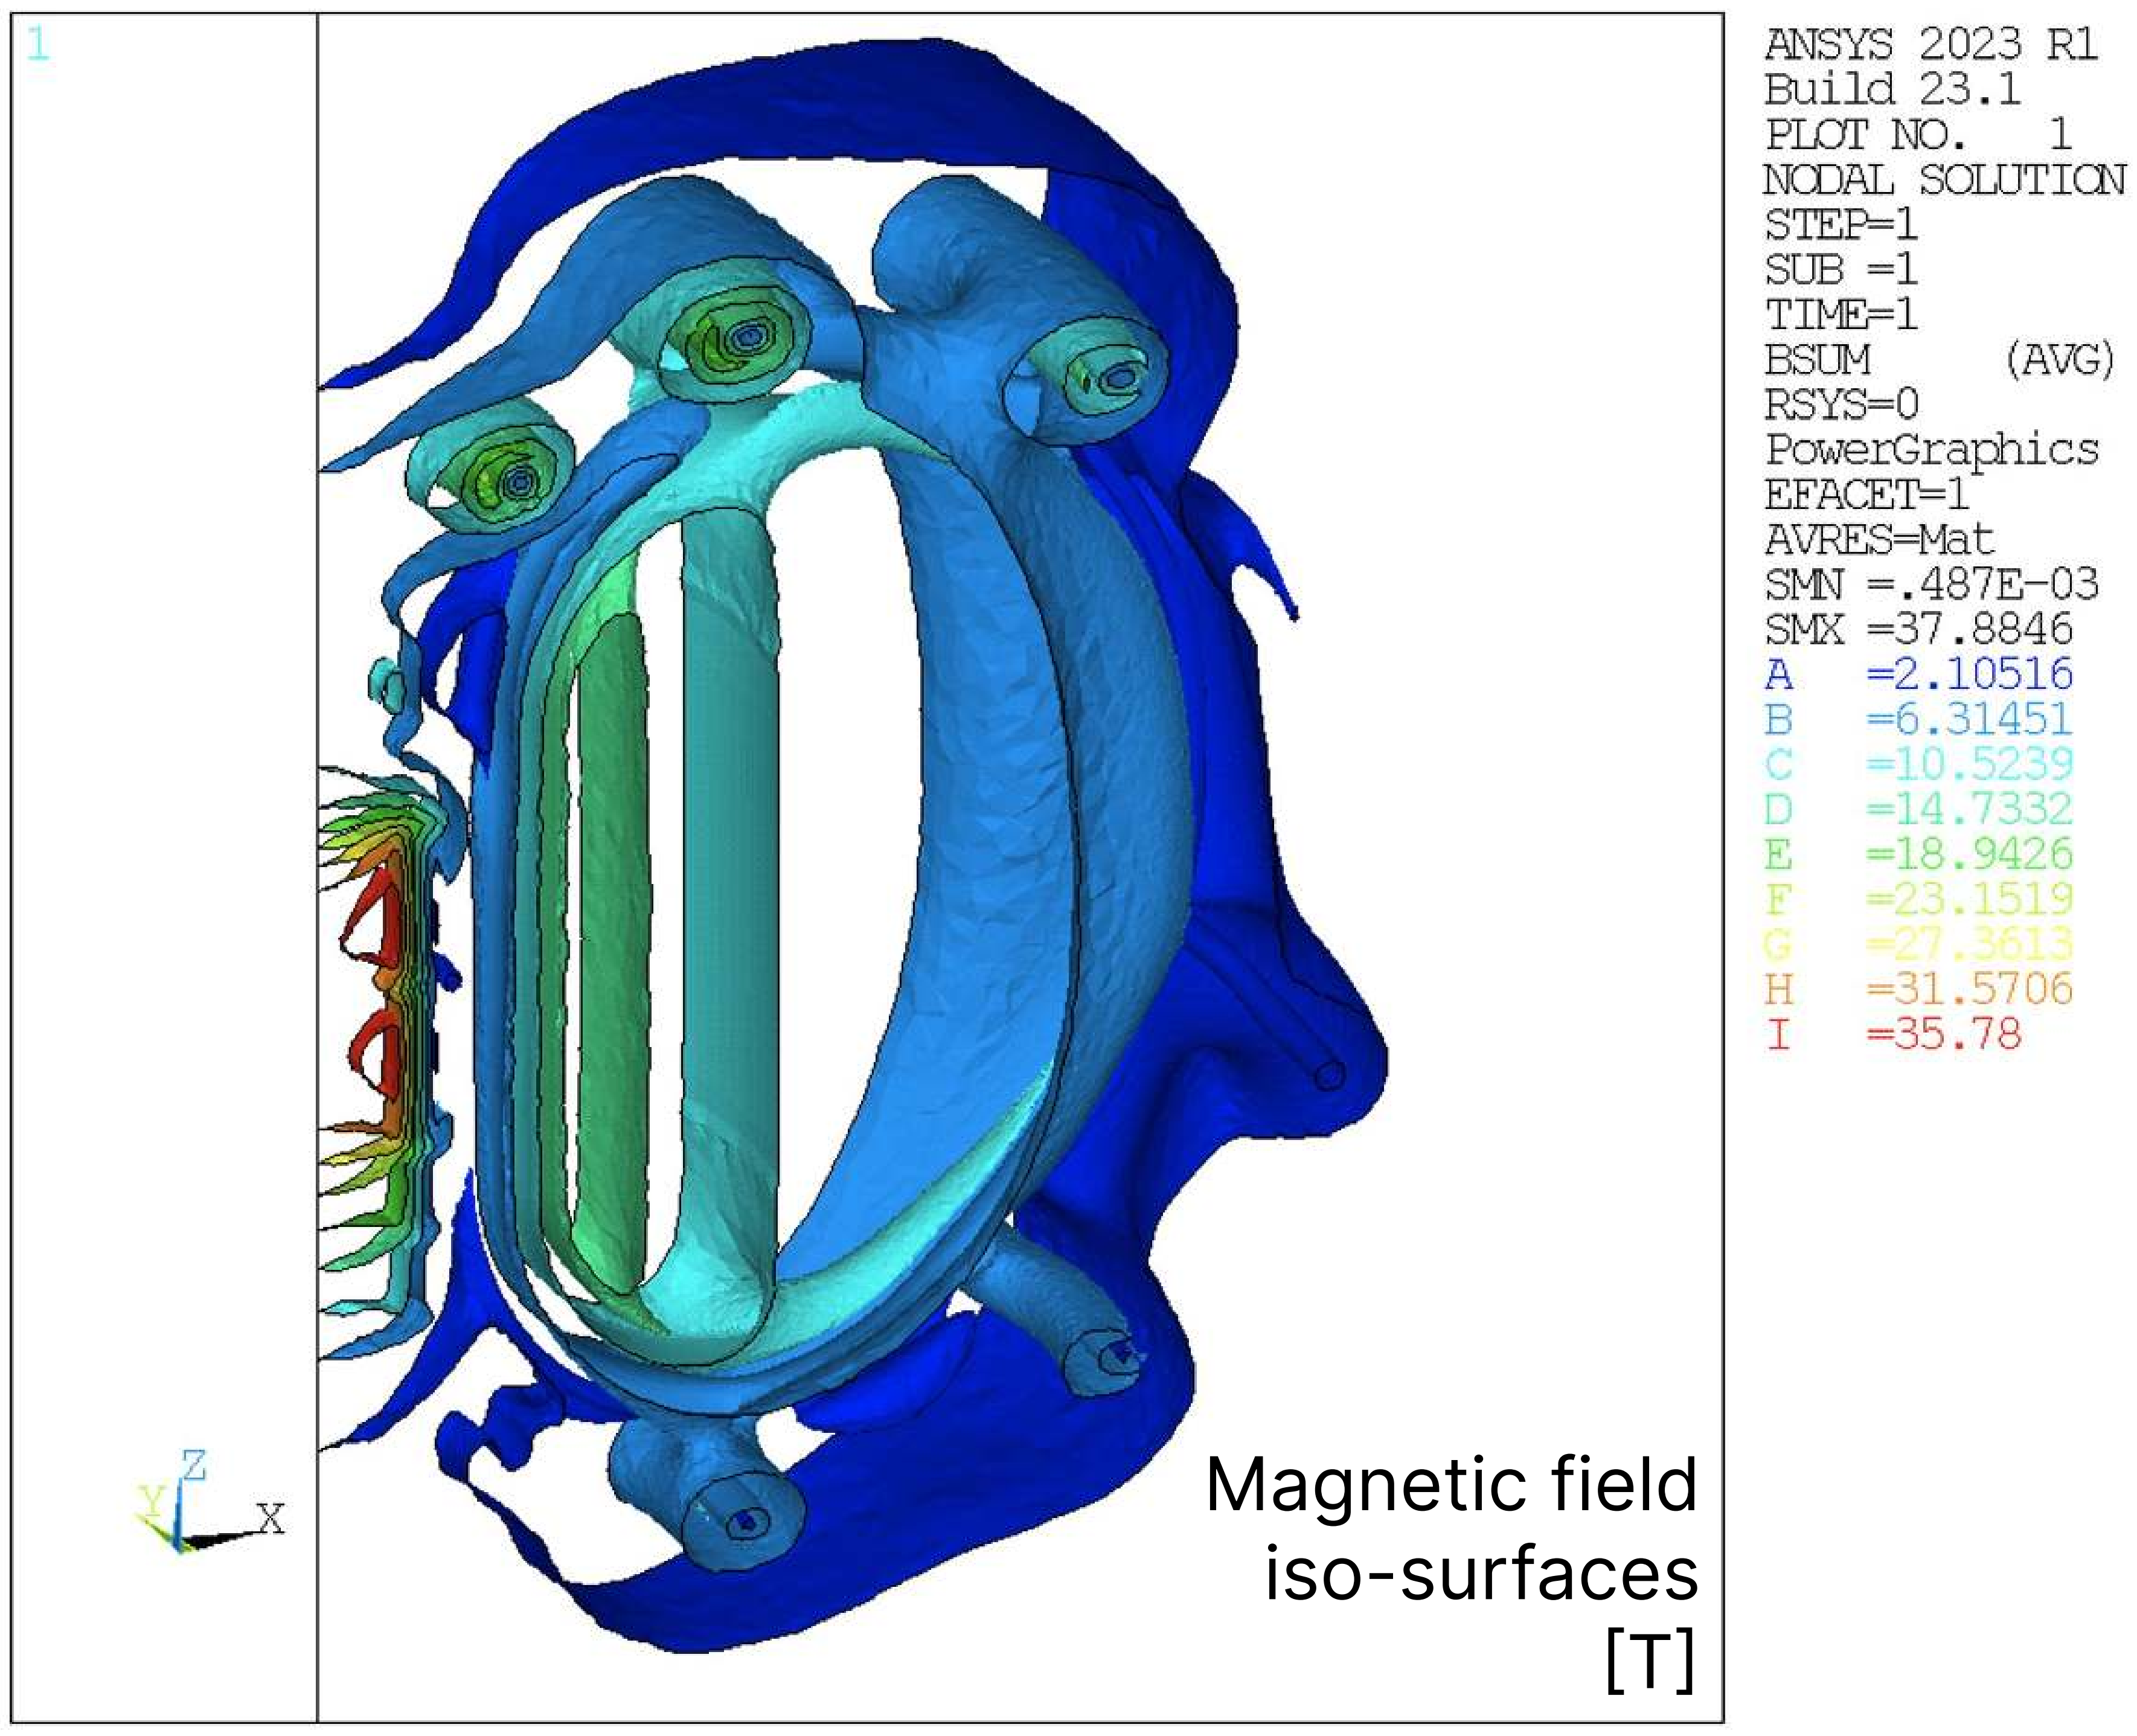

Electromagnetic Forces

on TF (x,y,z)

name	Fx[kN]	Fy[kN]	Fz[kN]	Mx[kNm]	My[kNm]	Mz[kNm]
TFC	-254746	-73	-30	-57442	103	-48

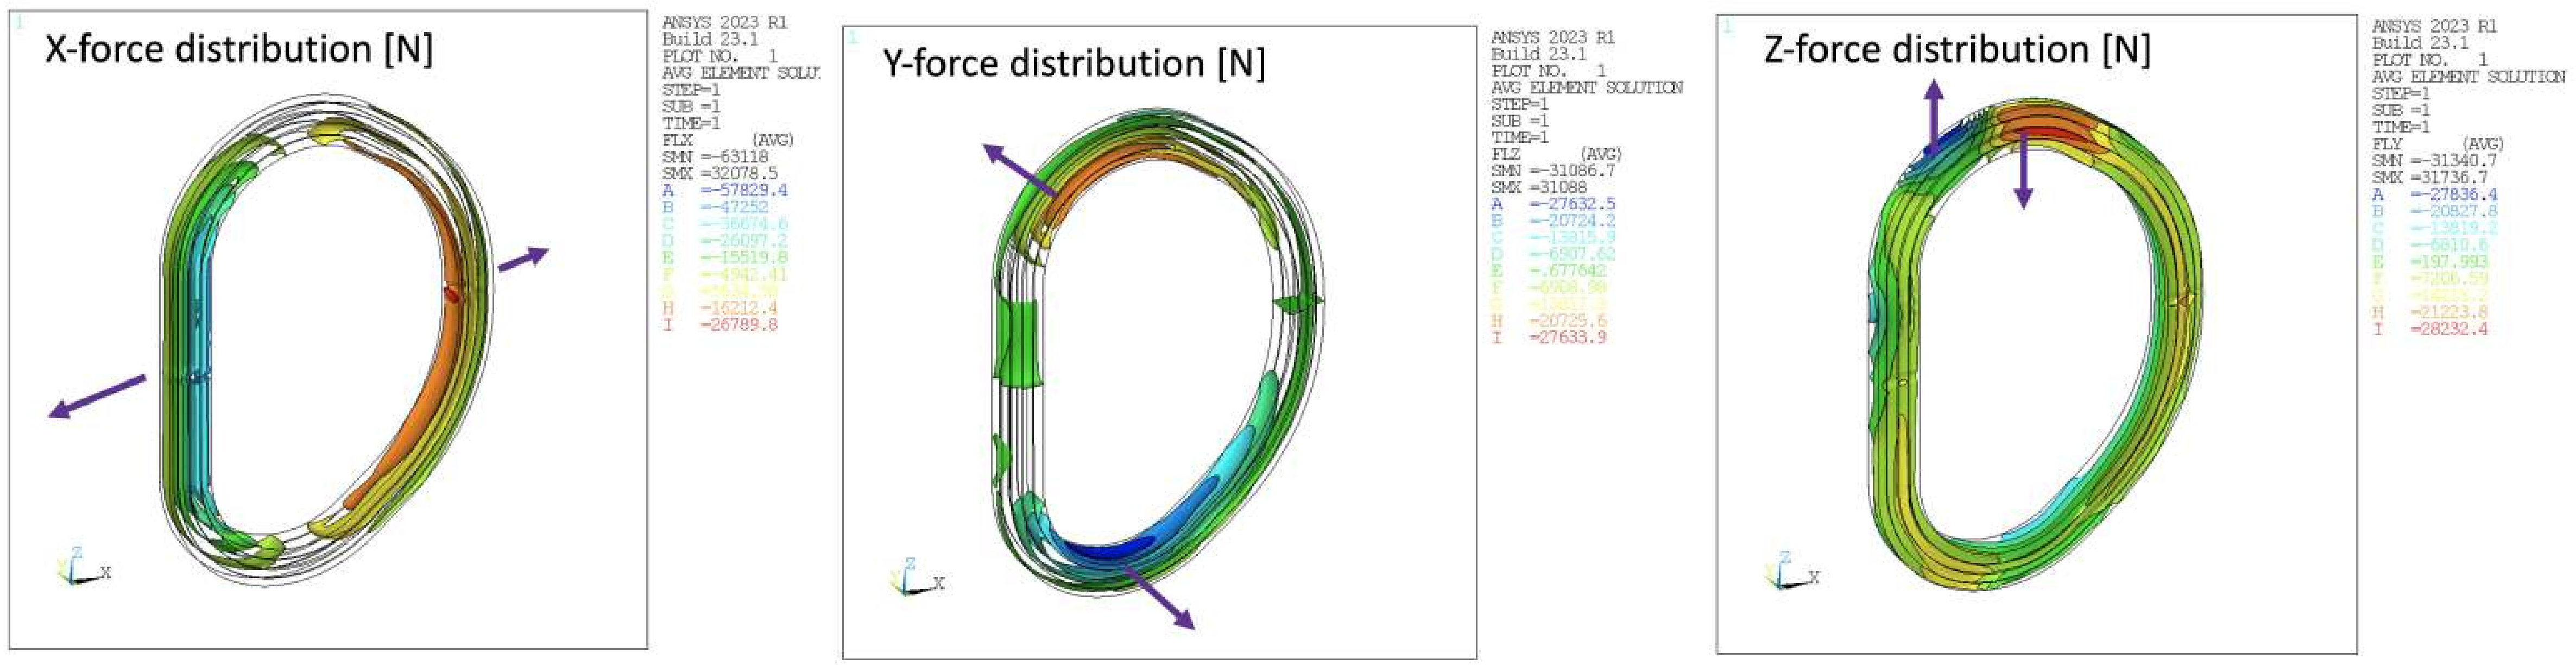

on PF+CS (x,y,z)

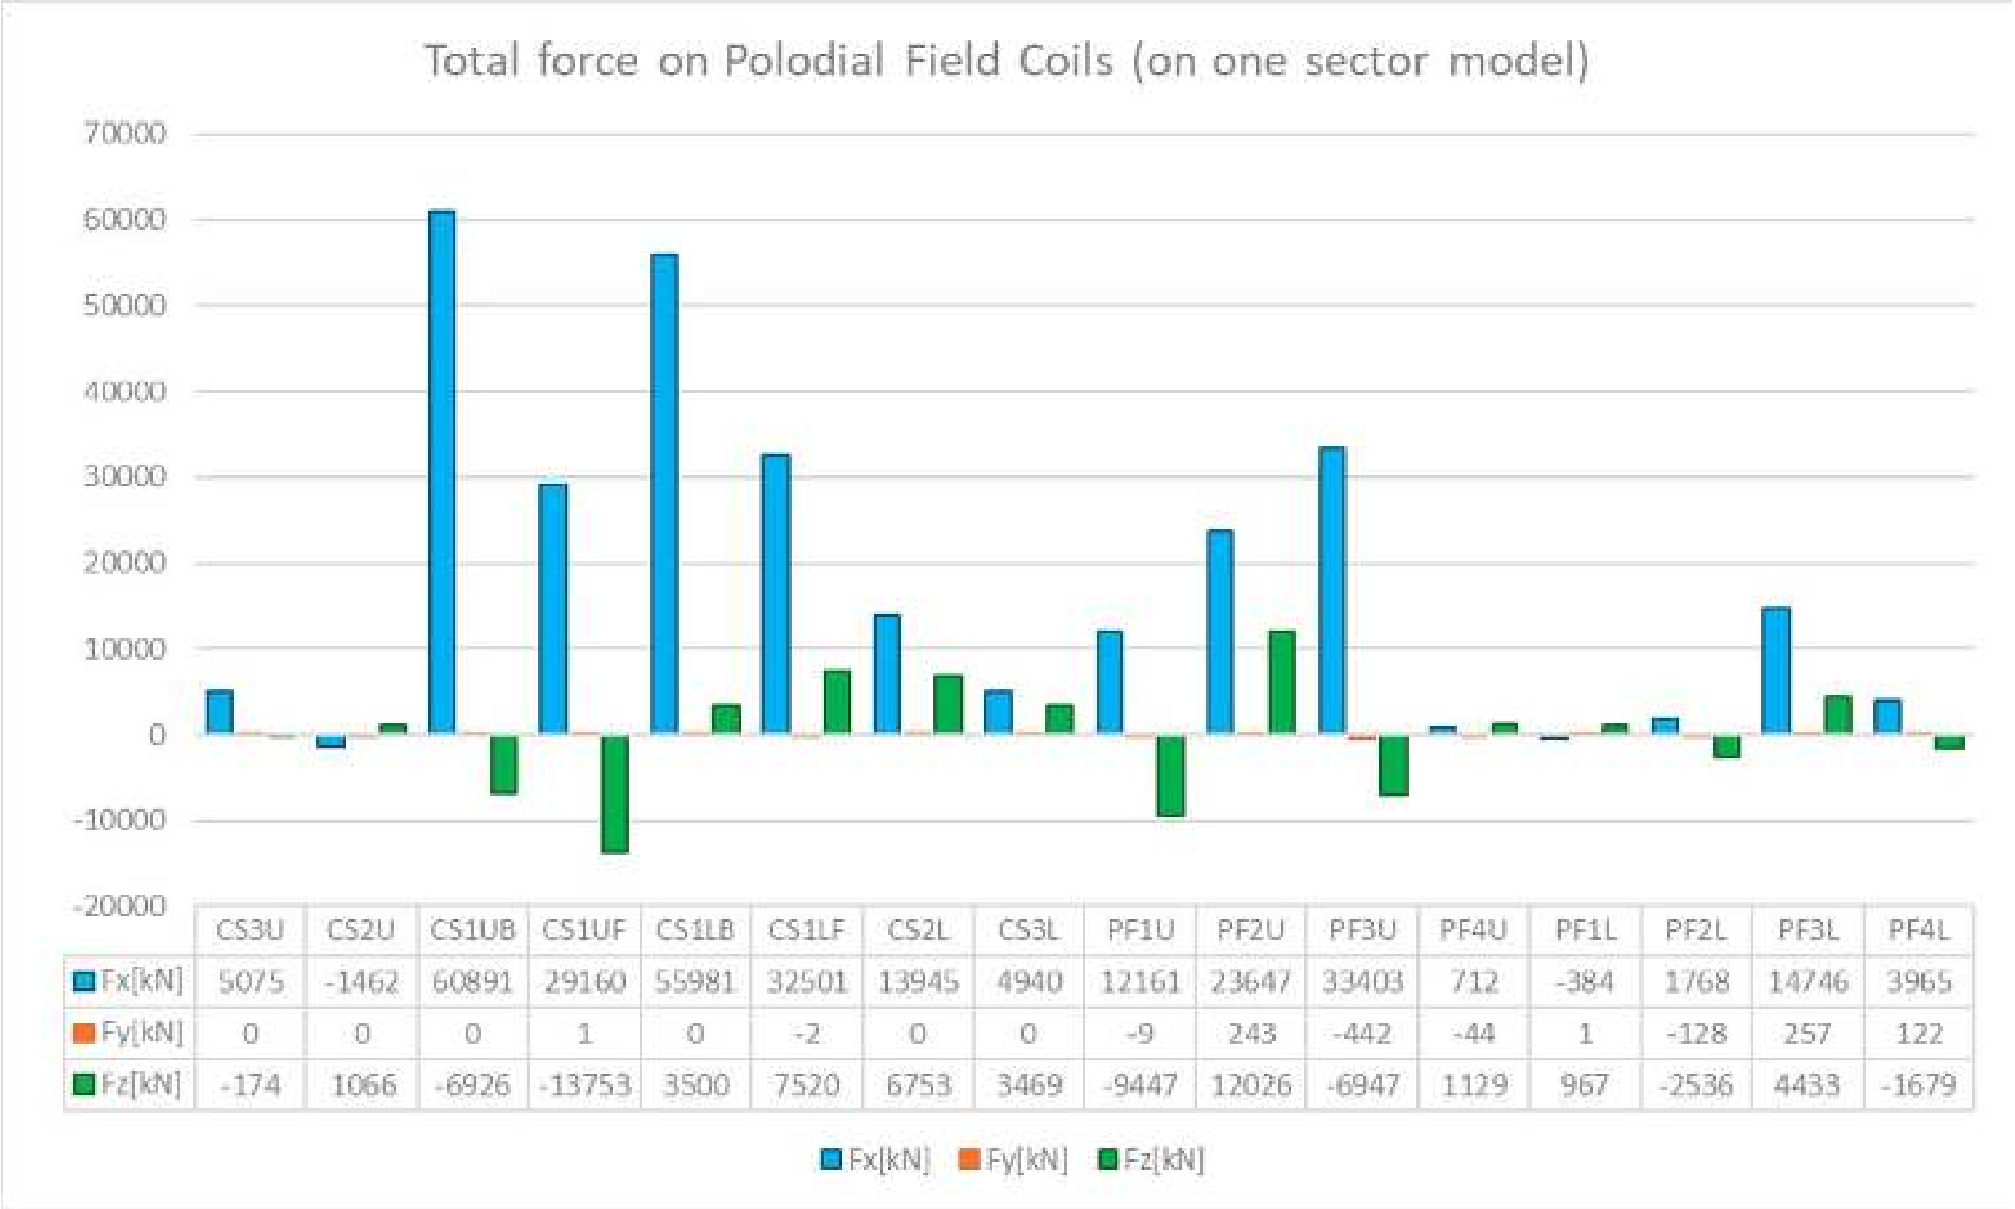

	CS3U	CS2U	CS1UB	CS1UF	CS1LB	CS1LF	CS2L	CS3L	PF1U	PF2U	PF3U	PF4U	PF1L	PF2L	PF3L	PF4L
Fx[kN]	5075	-1462	60891	29160	55981	32501	13945	4940	12161	23647	33403	712	-384	1768	14746	3965
Fy[kN]	0	0	0	1	0	-2	0	0	-9	243	-442	-44	1	-128	257	122
Fz[kN]	-174	1066	-6926	-13753	3500	7520	6753	3469	-9447	12026	-6947	1129	967	-2536	4433	-1679

As we can see, the main forces and torques on all coils are in the radial "X" direction. But these values are the same order of magnitude compared to ITER, proving EM load survivability.

Structural Analysis (cont.)

Vacuum Vessel Loads

Tokamaks operate only under ultra-high vacuum conditions, requiring a robust vacuum vessel to maintain extremely low internal pressures (around 10^{-8} to 10^{-9} atmospheres). The vacuum vessel ensures plasma purity by preventing contamination from external gases. Additionally, it must withstand significant structural loads due to vacuum pressures, thermal cycling, and electromagnetic forces induced by plasma disruptions and coil currents. The simplified model (without divertor) is assumed to be under baking conditions (240°C) and 1bar to see how it displaces under these stresses

- Geometry has been created from CAD model, by extracting the inner/outer skin of the VV
- Internal ribs have been added to connect the skins (4 ribs per typical sector)
- One typical sector FEM model (360-degree 22.5 degrees) considered for structural analysis
- Complete 360 degree model considered for stability analysis against crushing due to vacuum conditions
- Cyclic Symmetry boundary conditions applied to the sector model
- VV material is austenitic stainless steel (e.g. SS316L)
- Internal ribs are connected to the shells to simulate a "perfect" weld

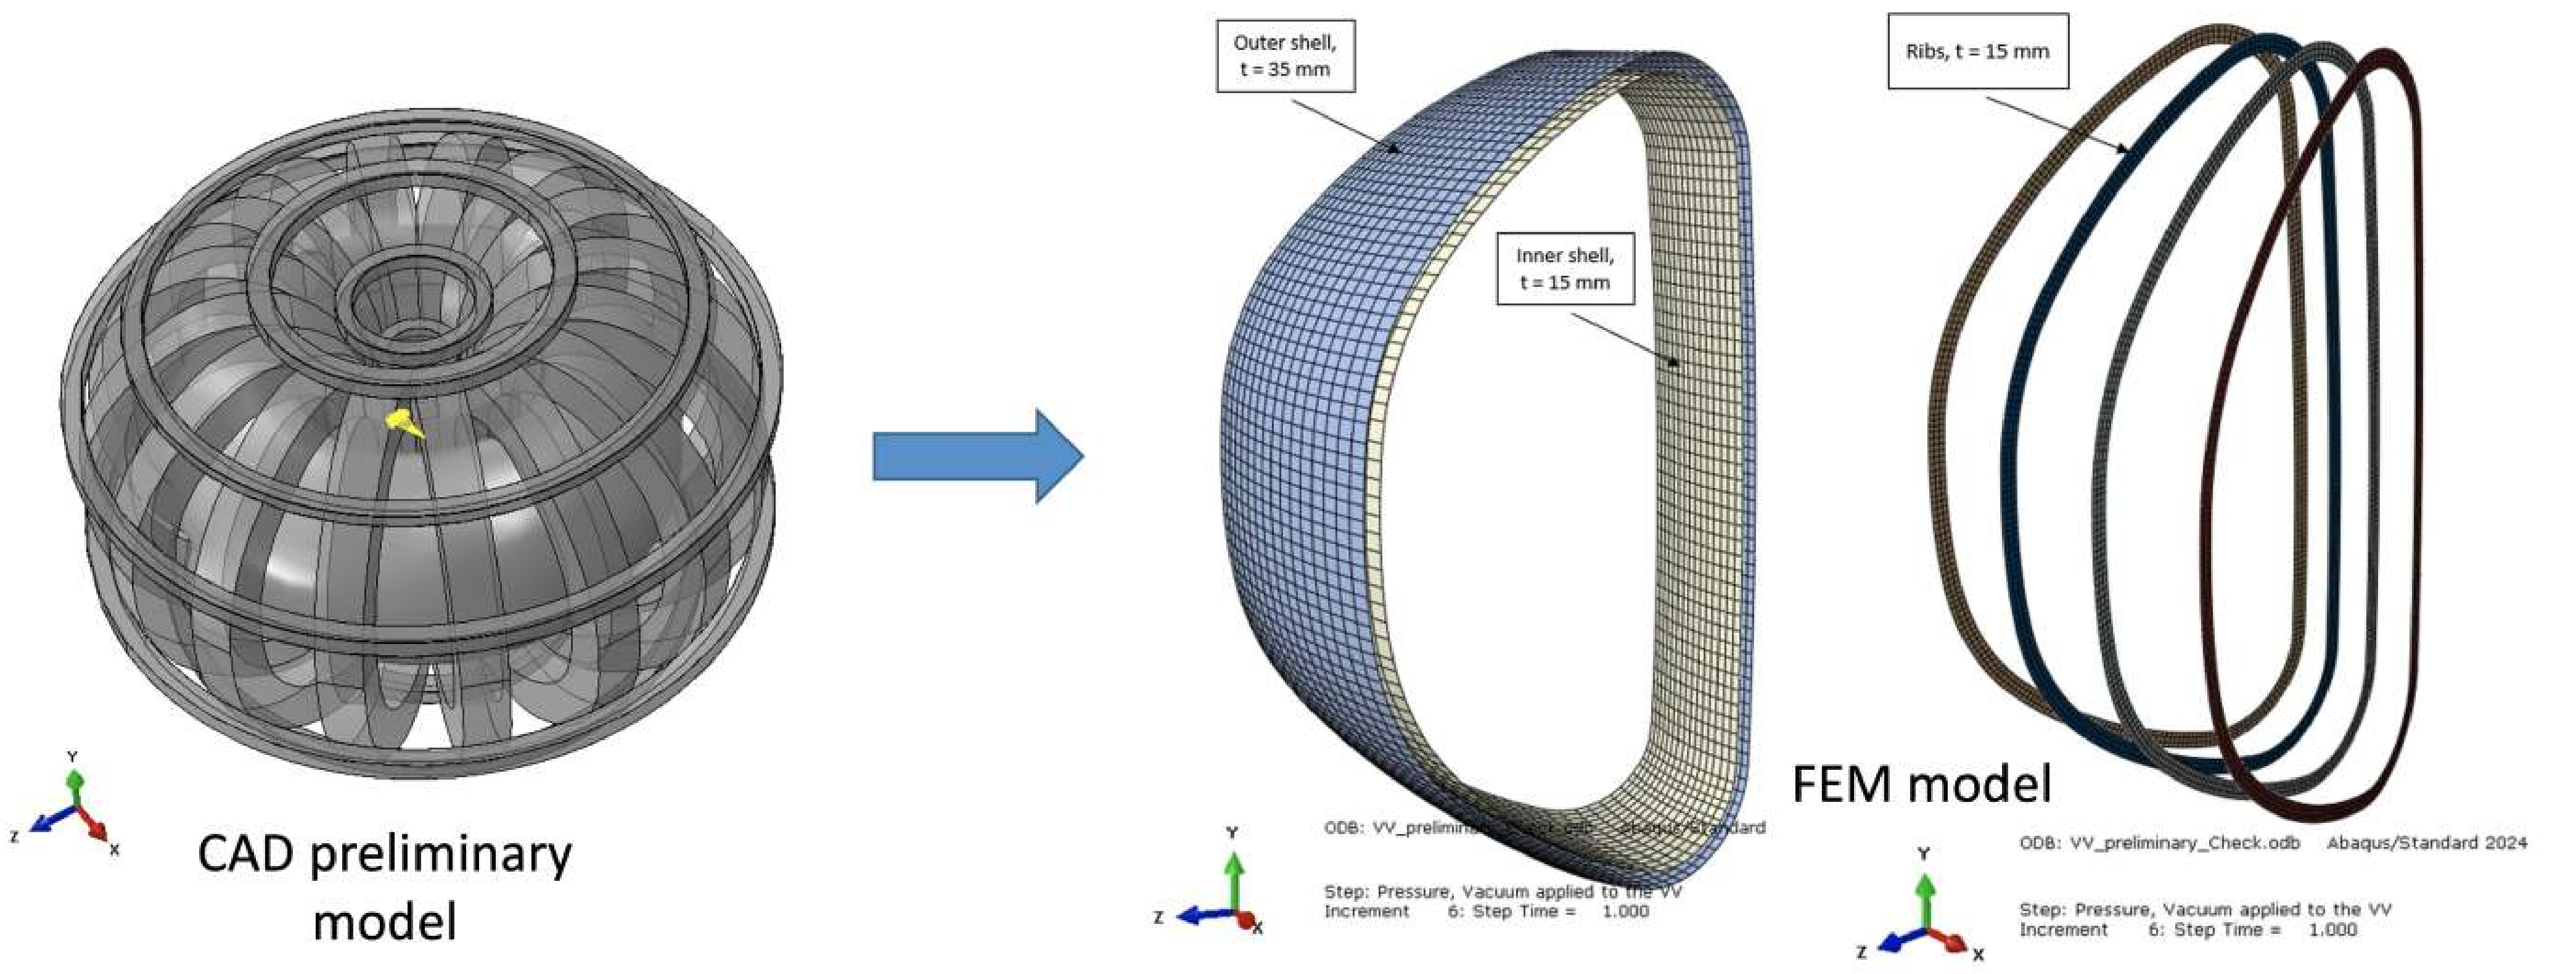

VV Displacement under 240°C + 1bar

Single Sector

Von Mises Stress under 240°C + 1bar

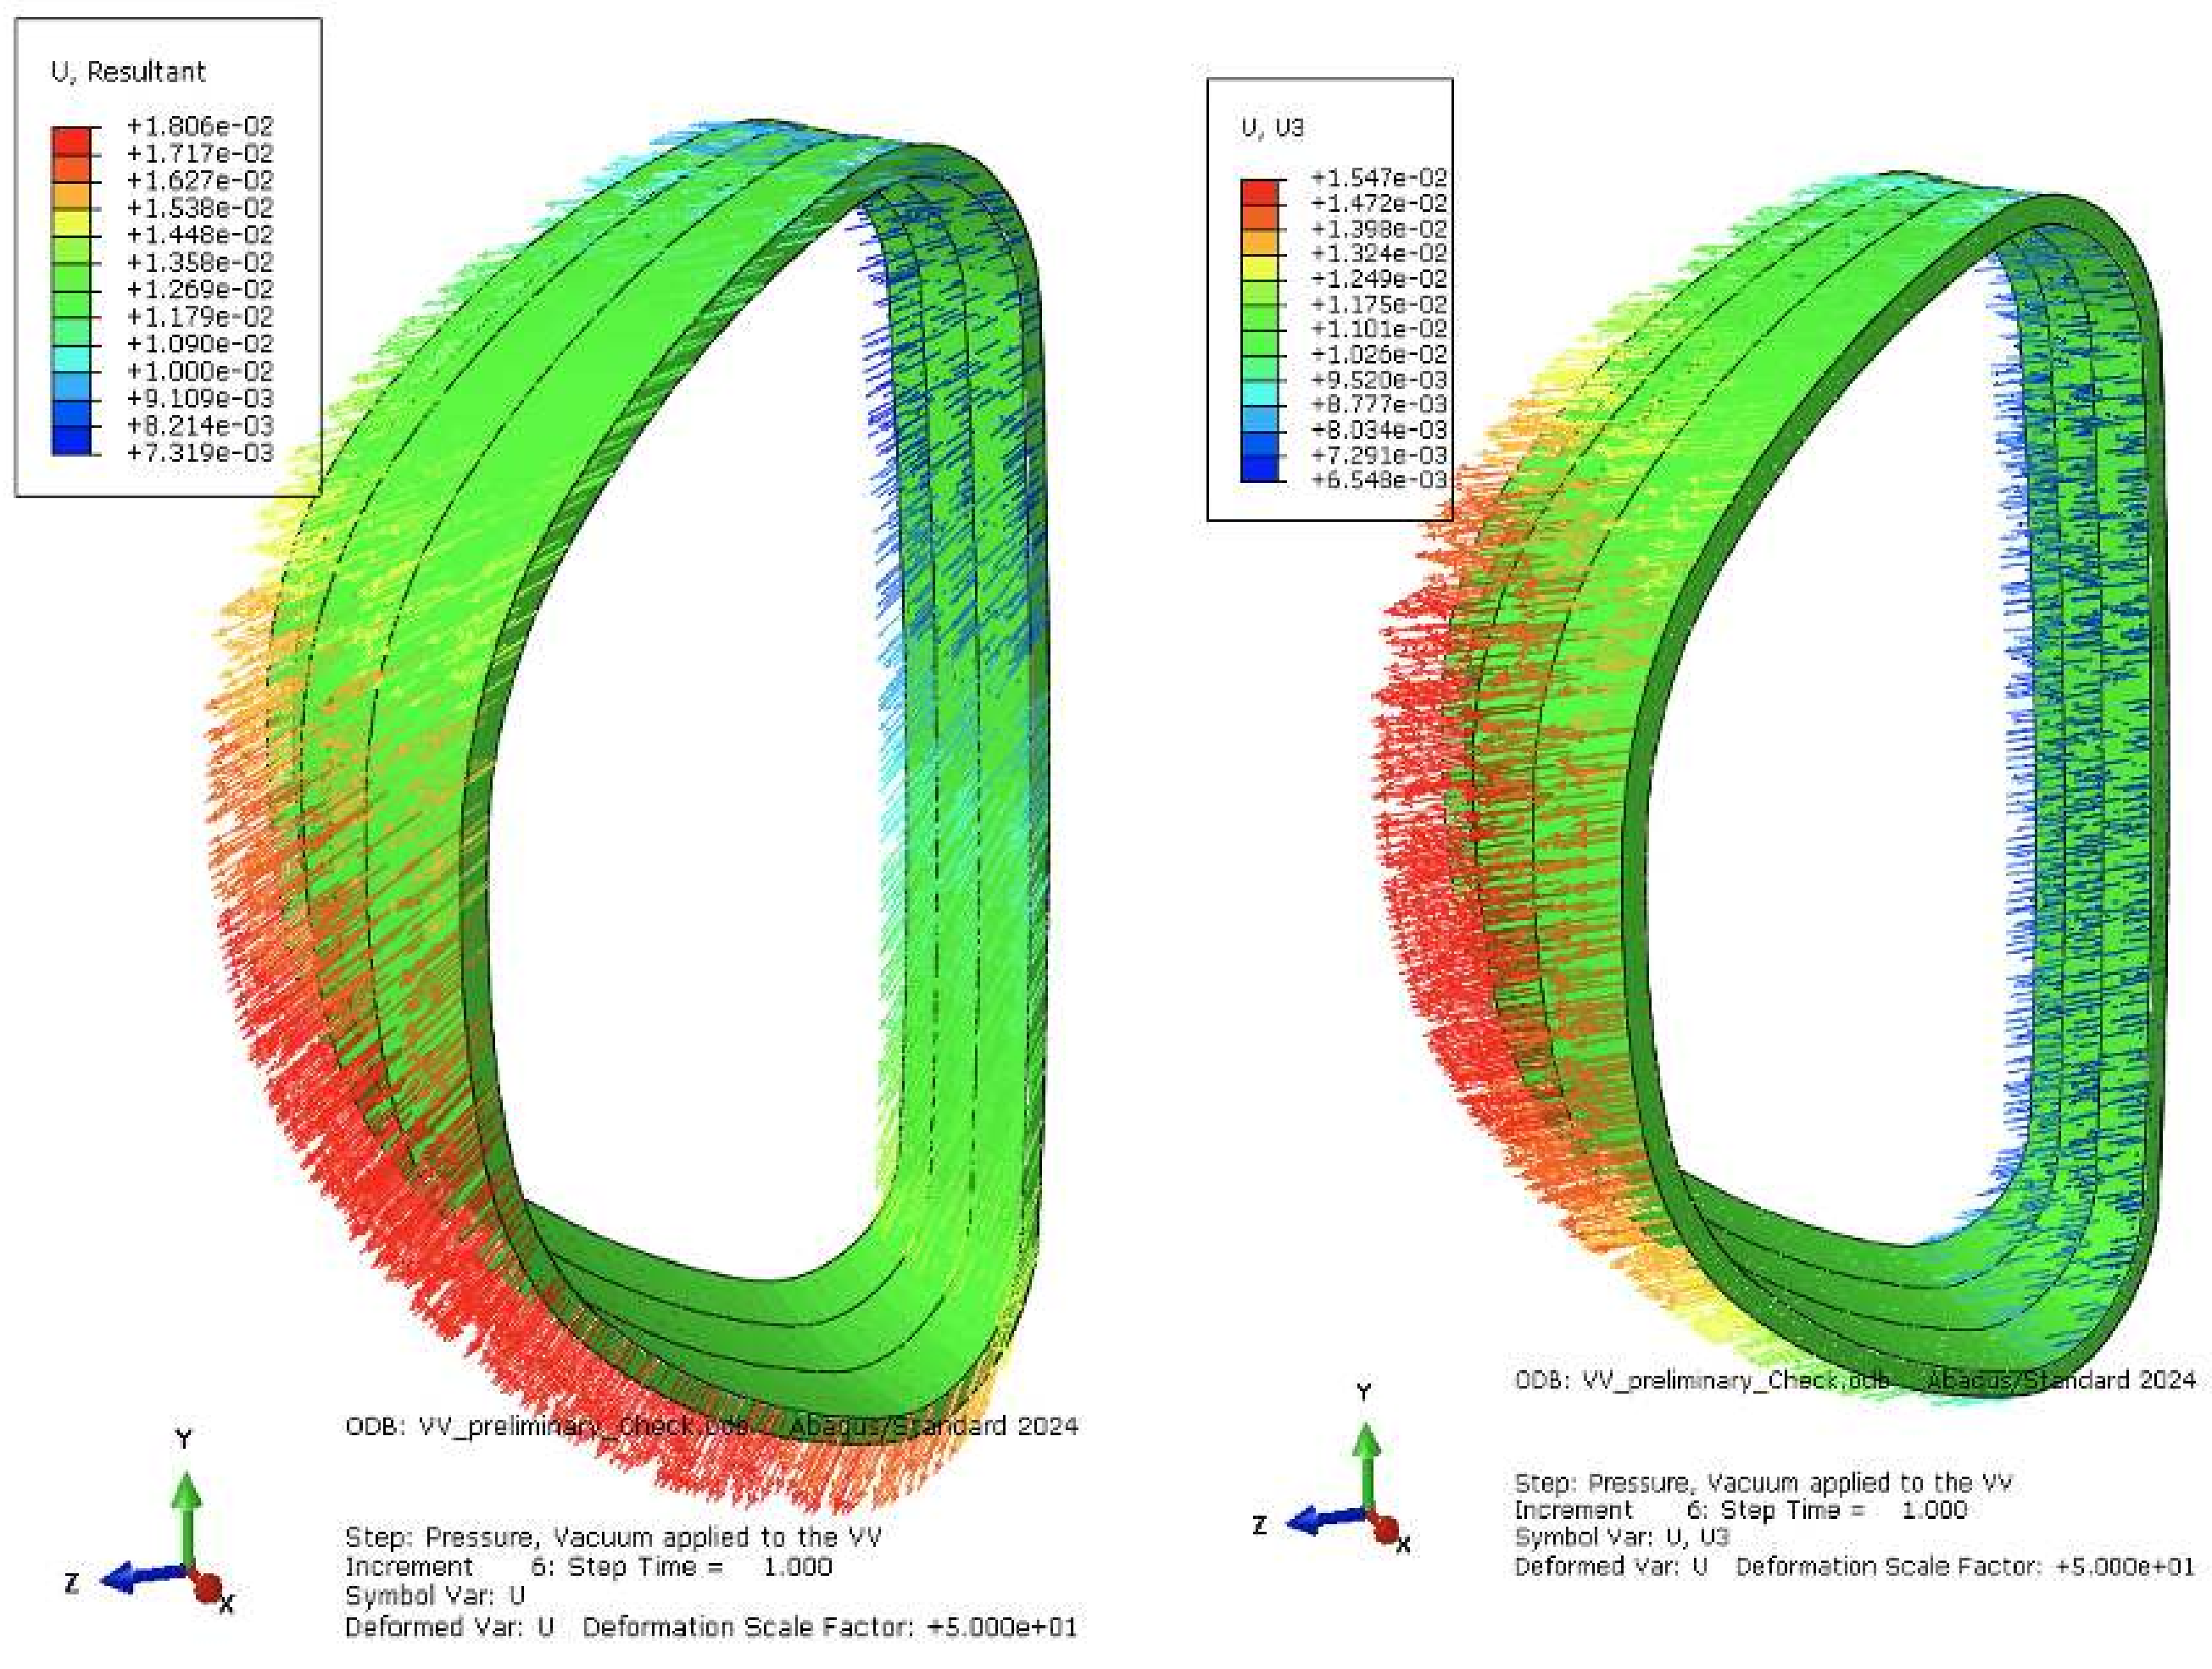

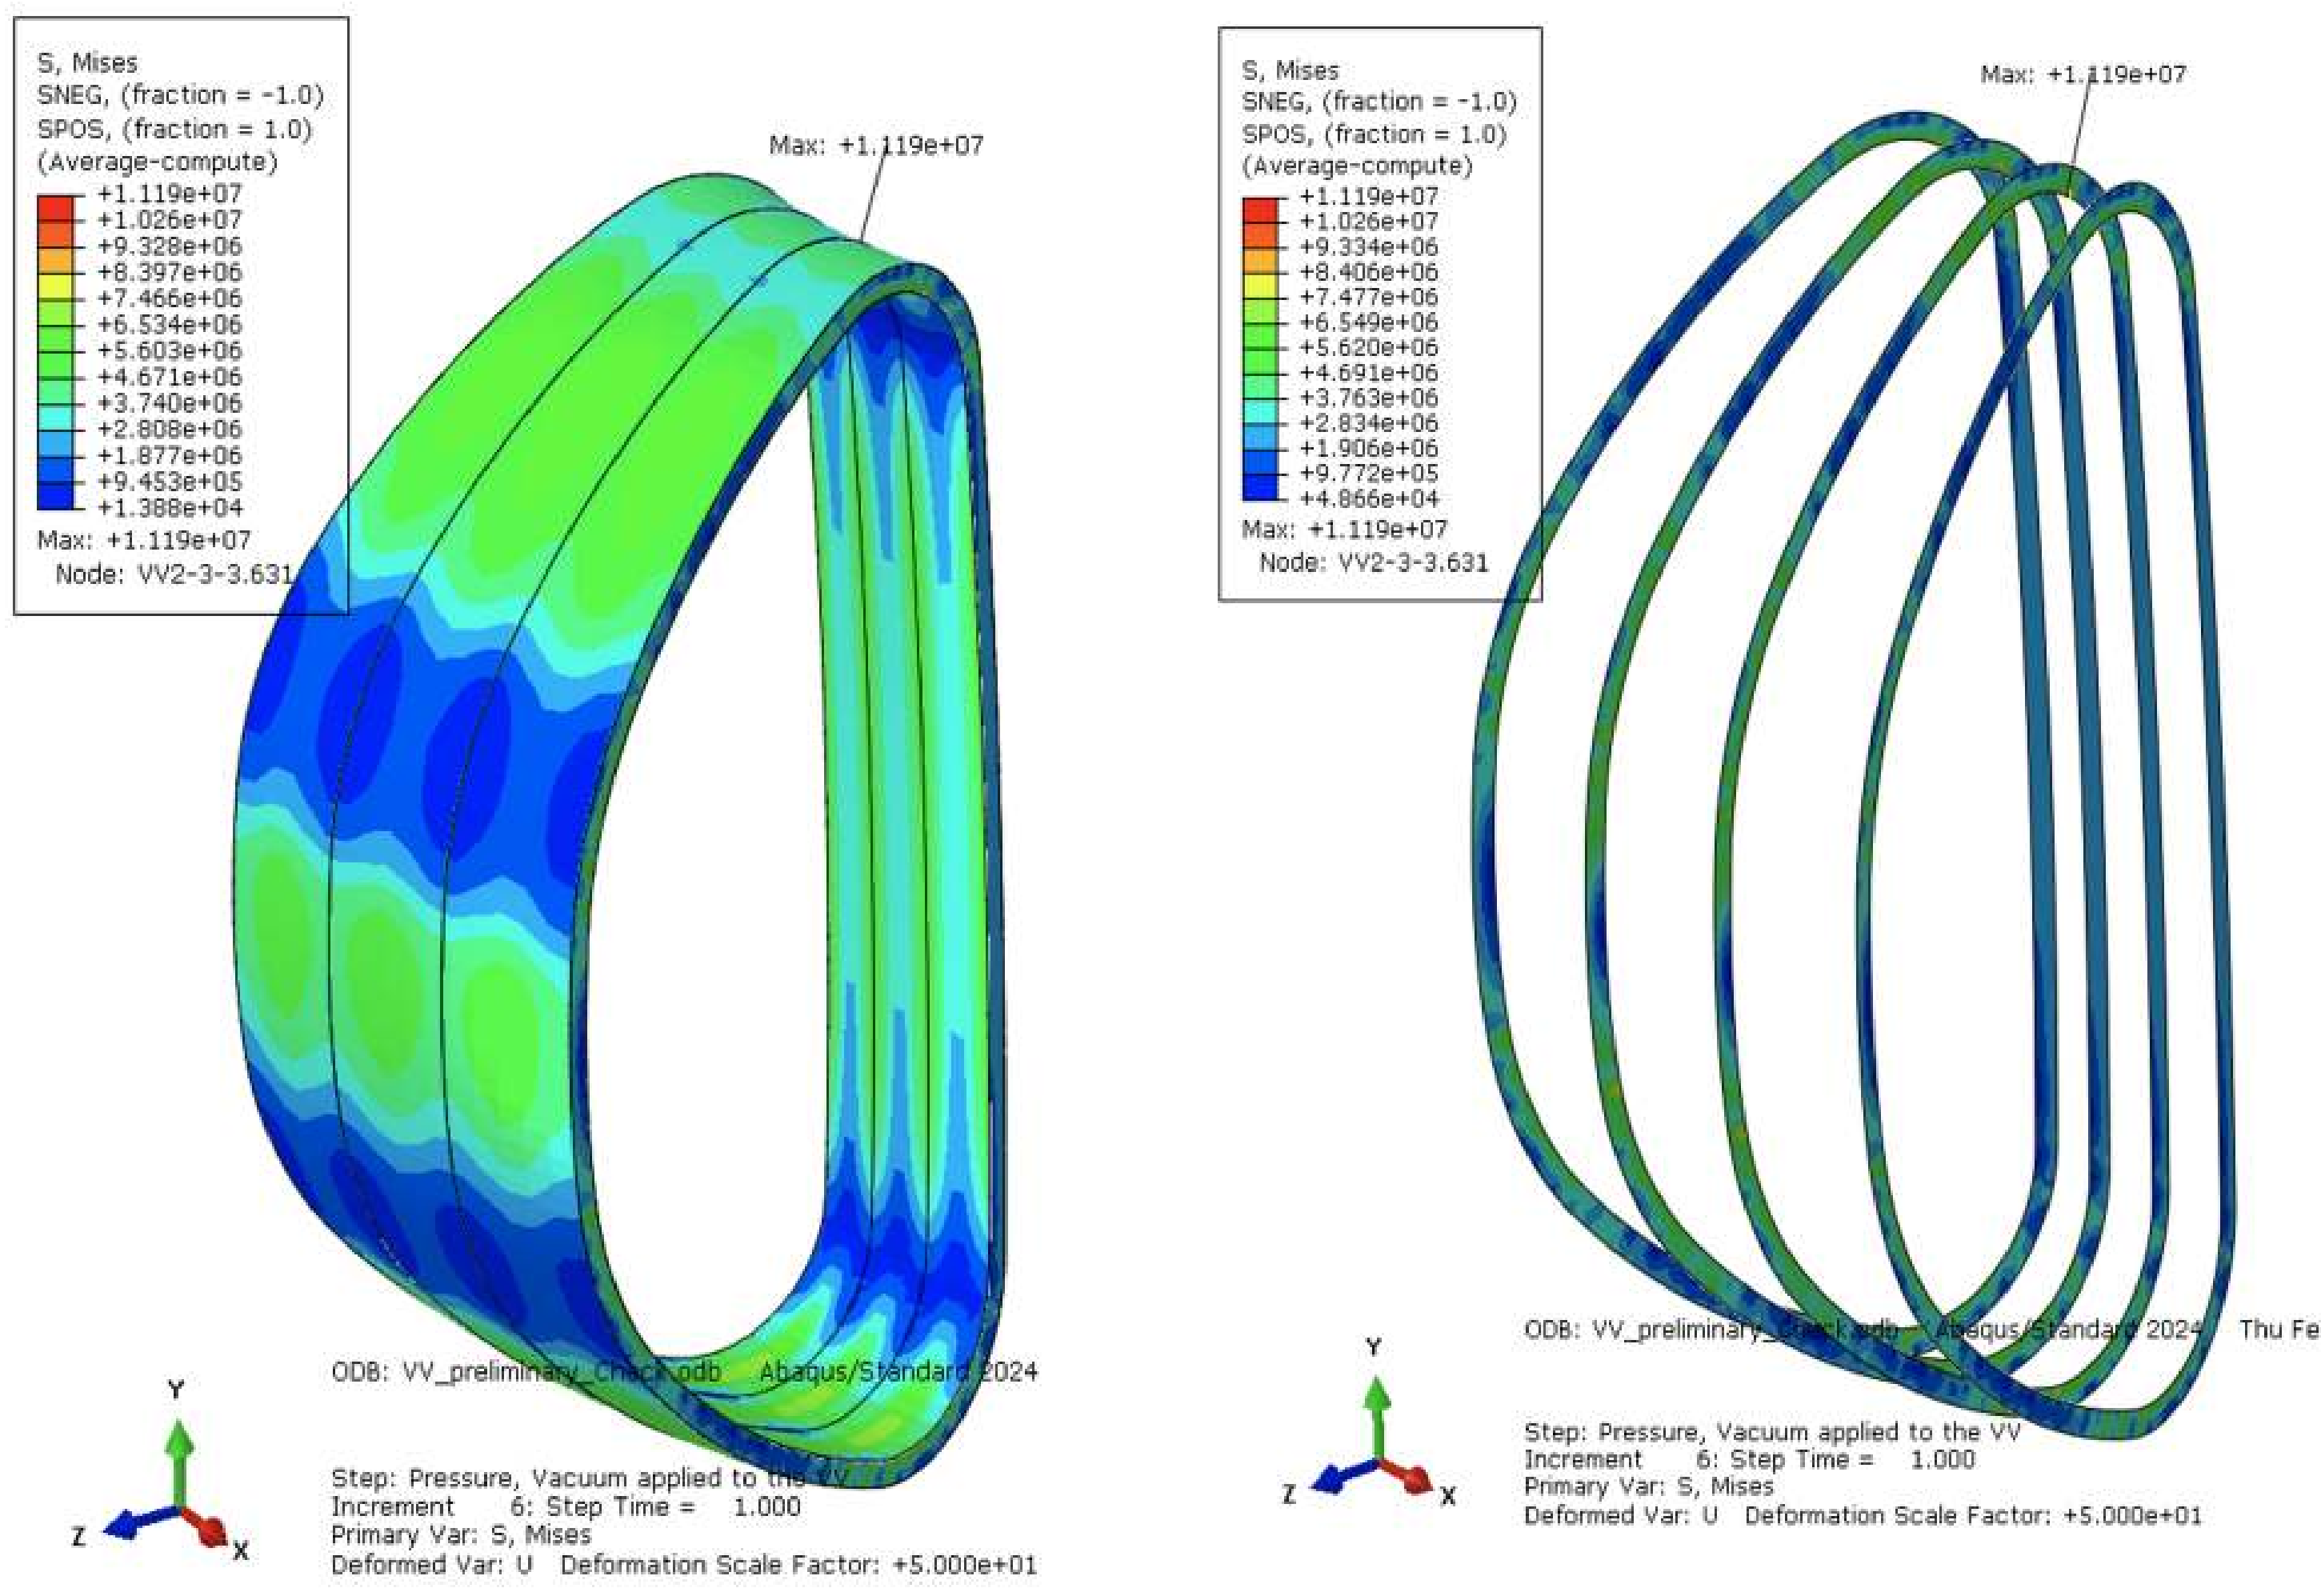

Entire Vessel

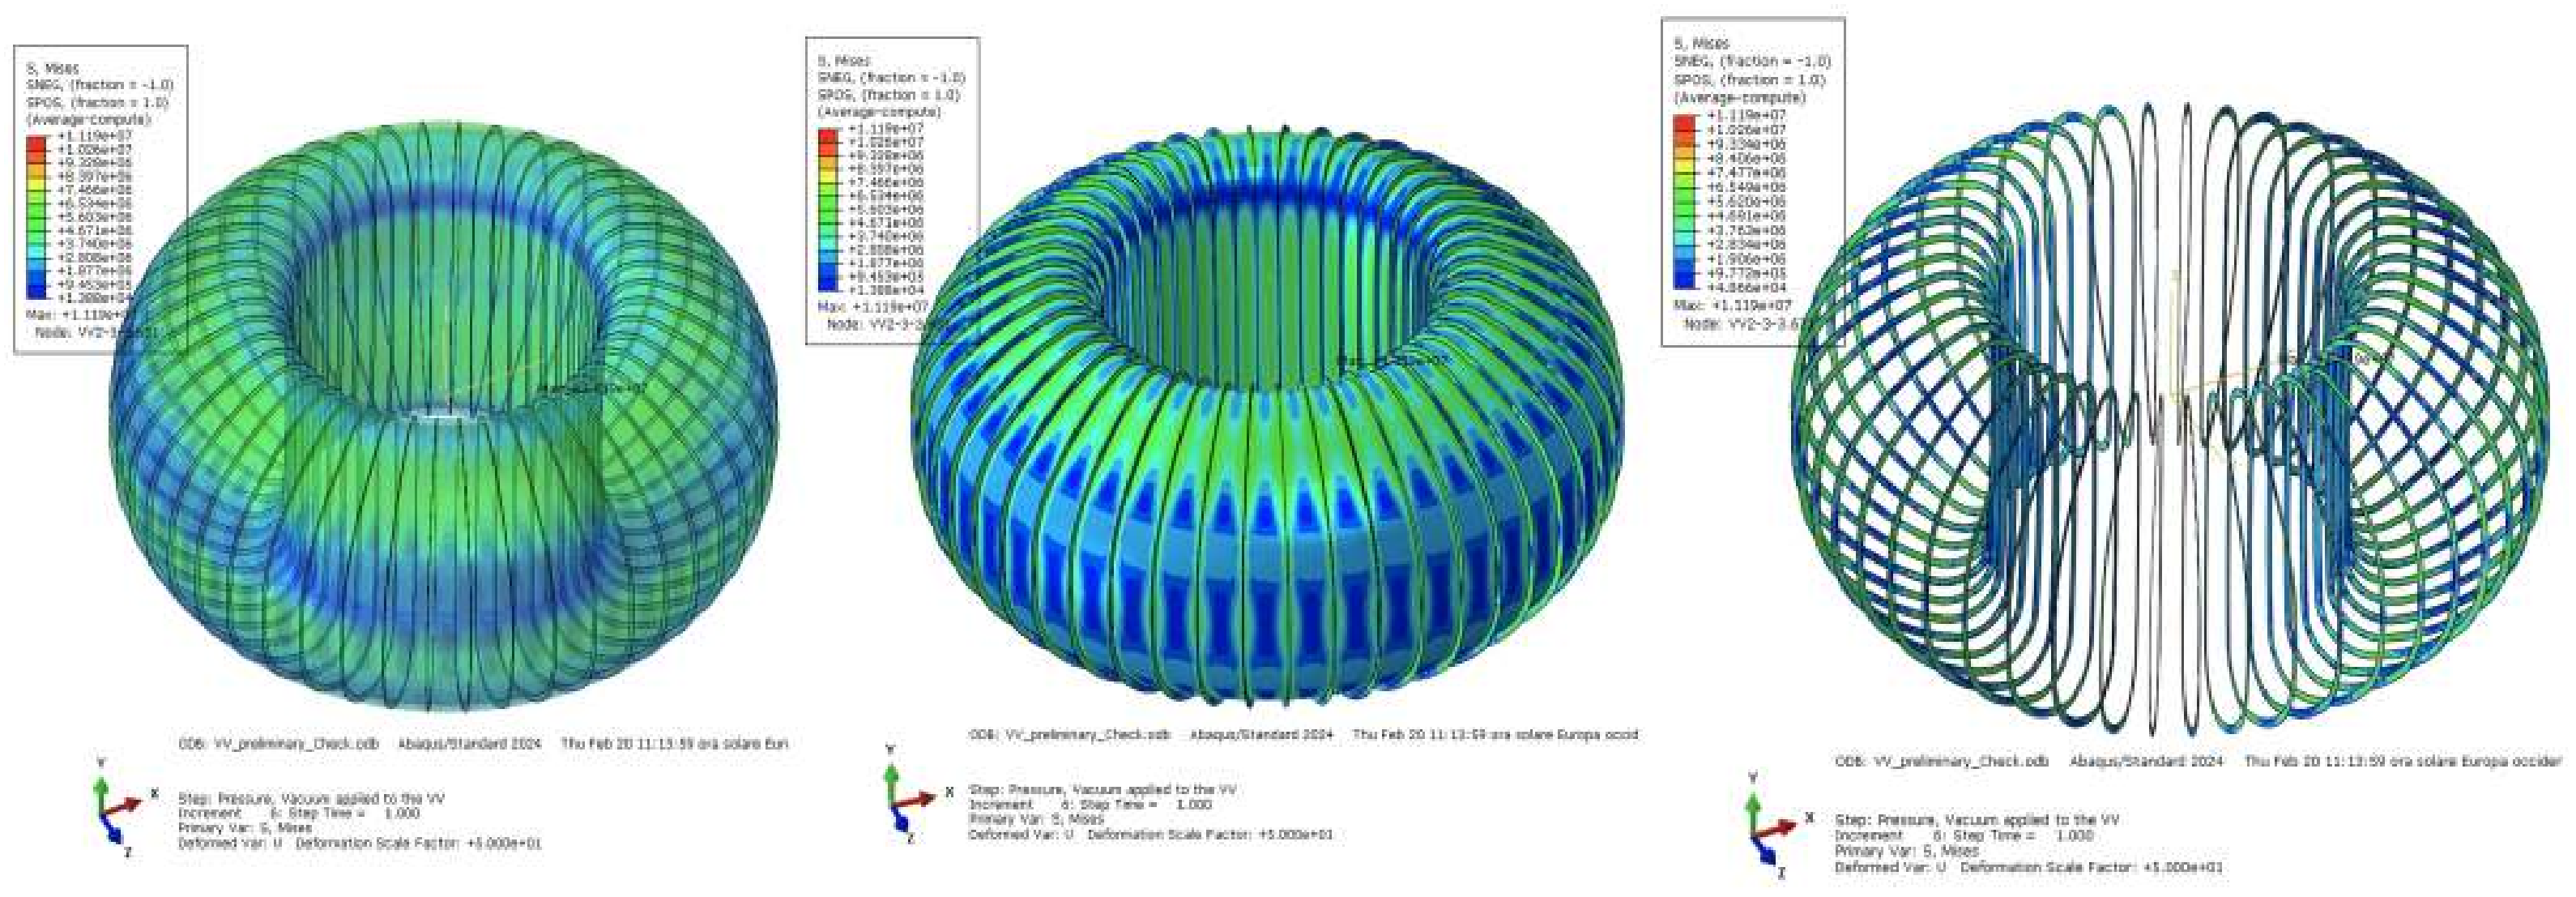

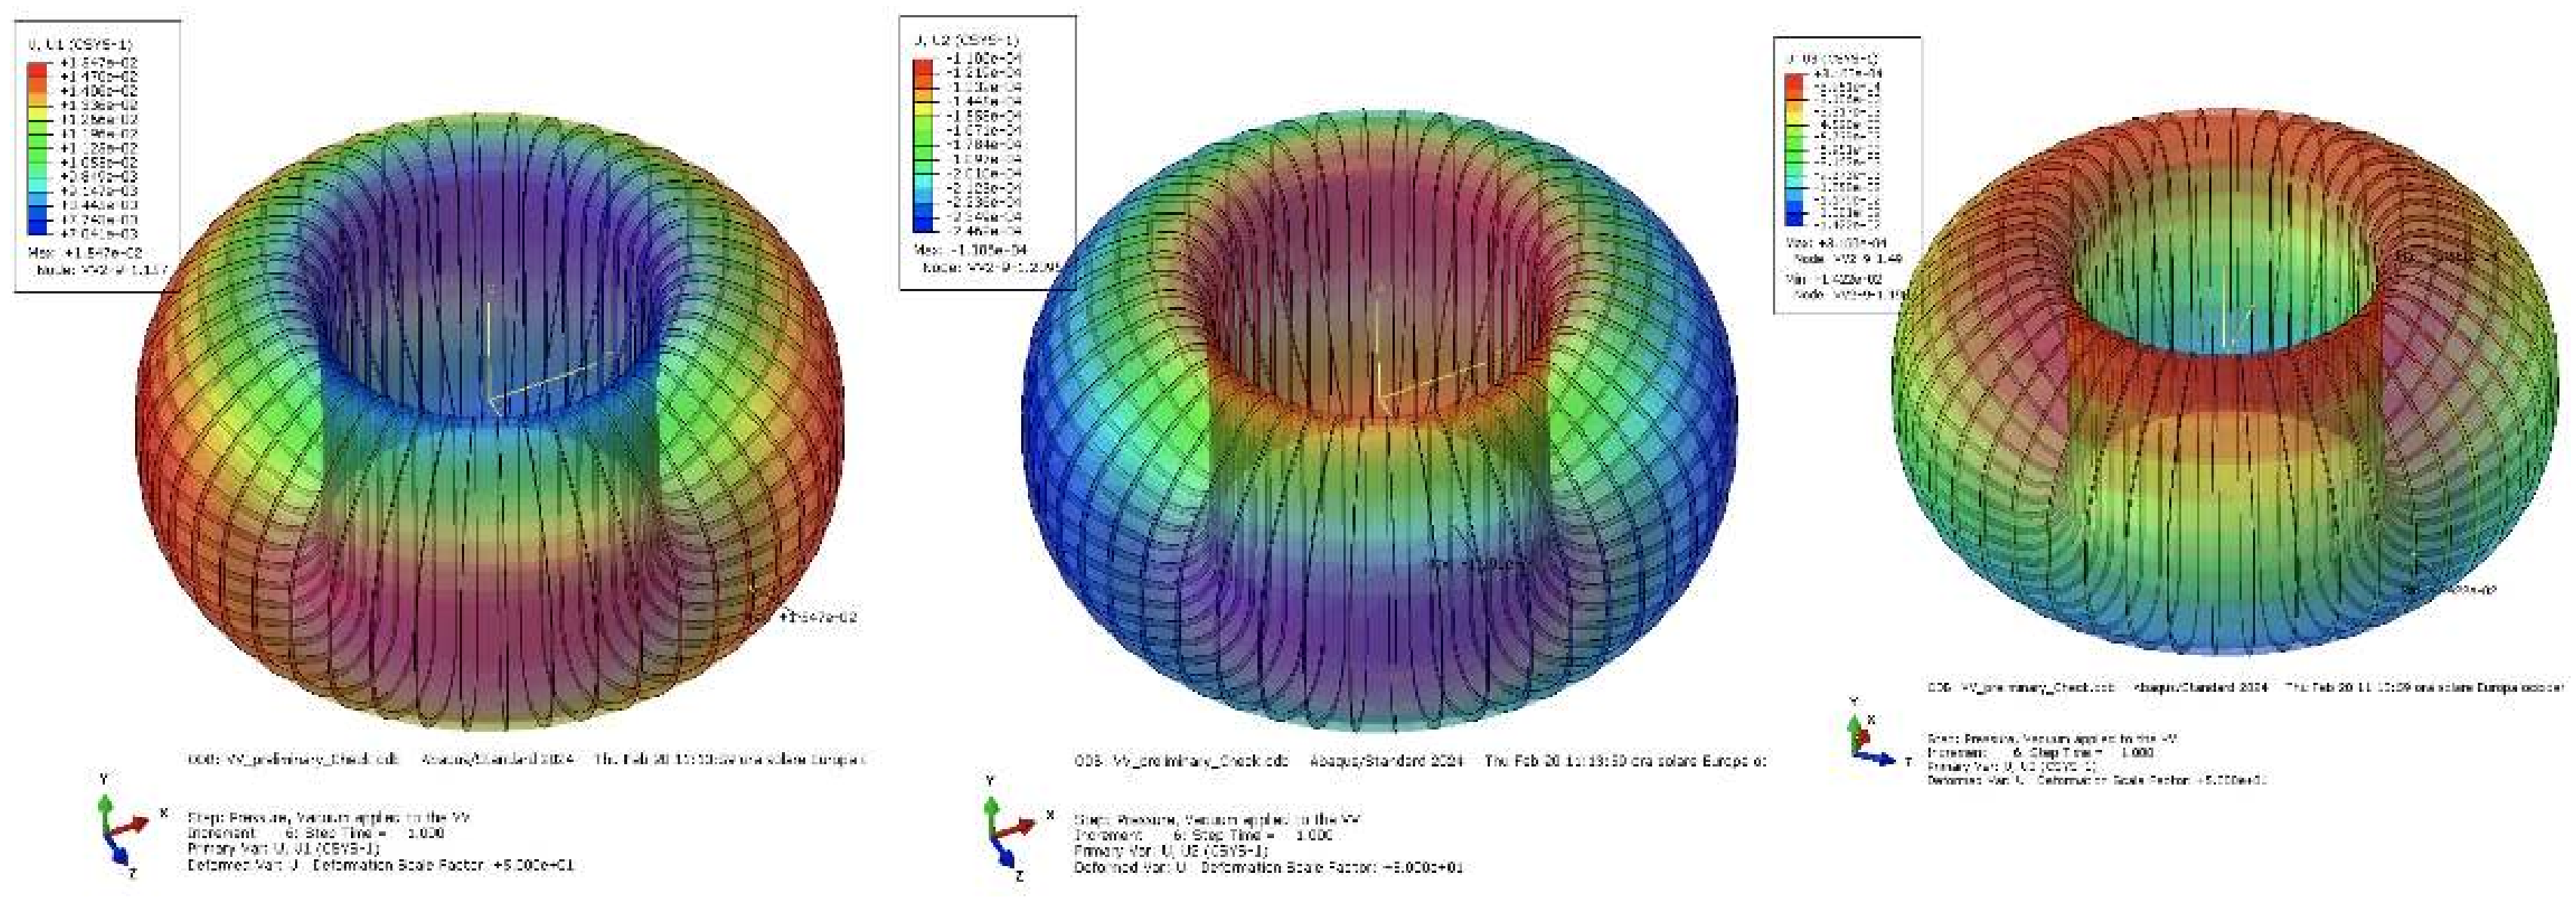

We can see that the vessel displaces no more 2cm at max, which is acceptable. In addition, the Von Mises stress is very low, as expected.

Buckling Analysis

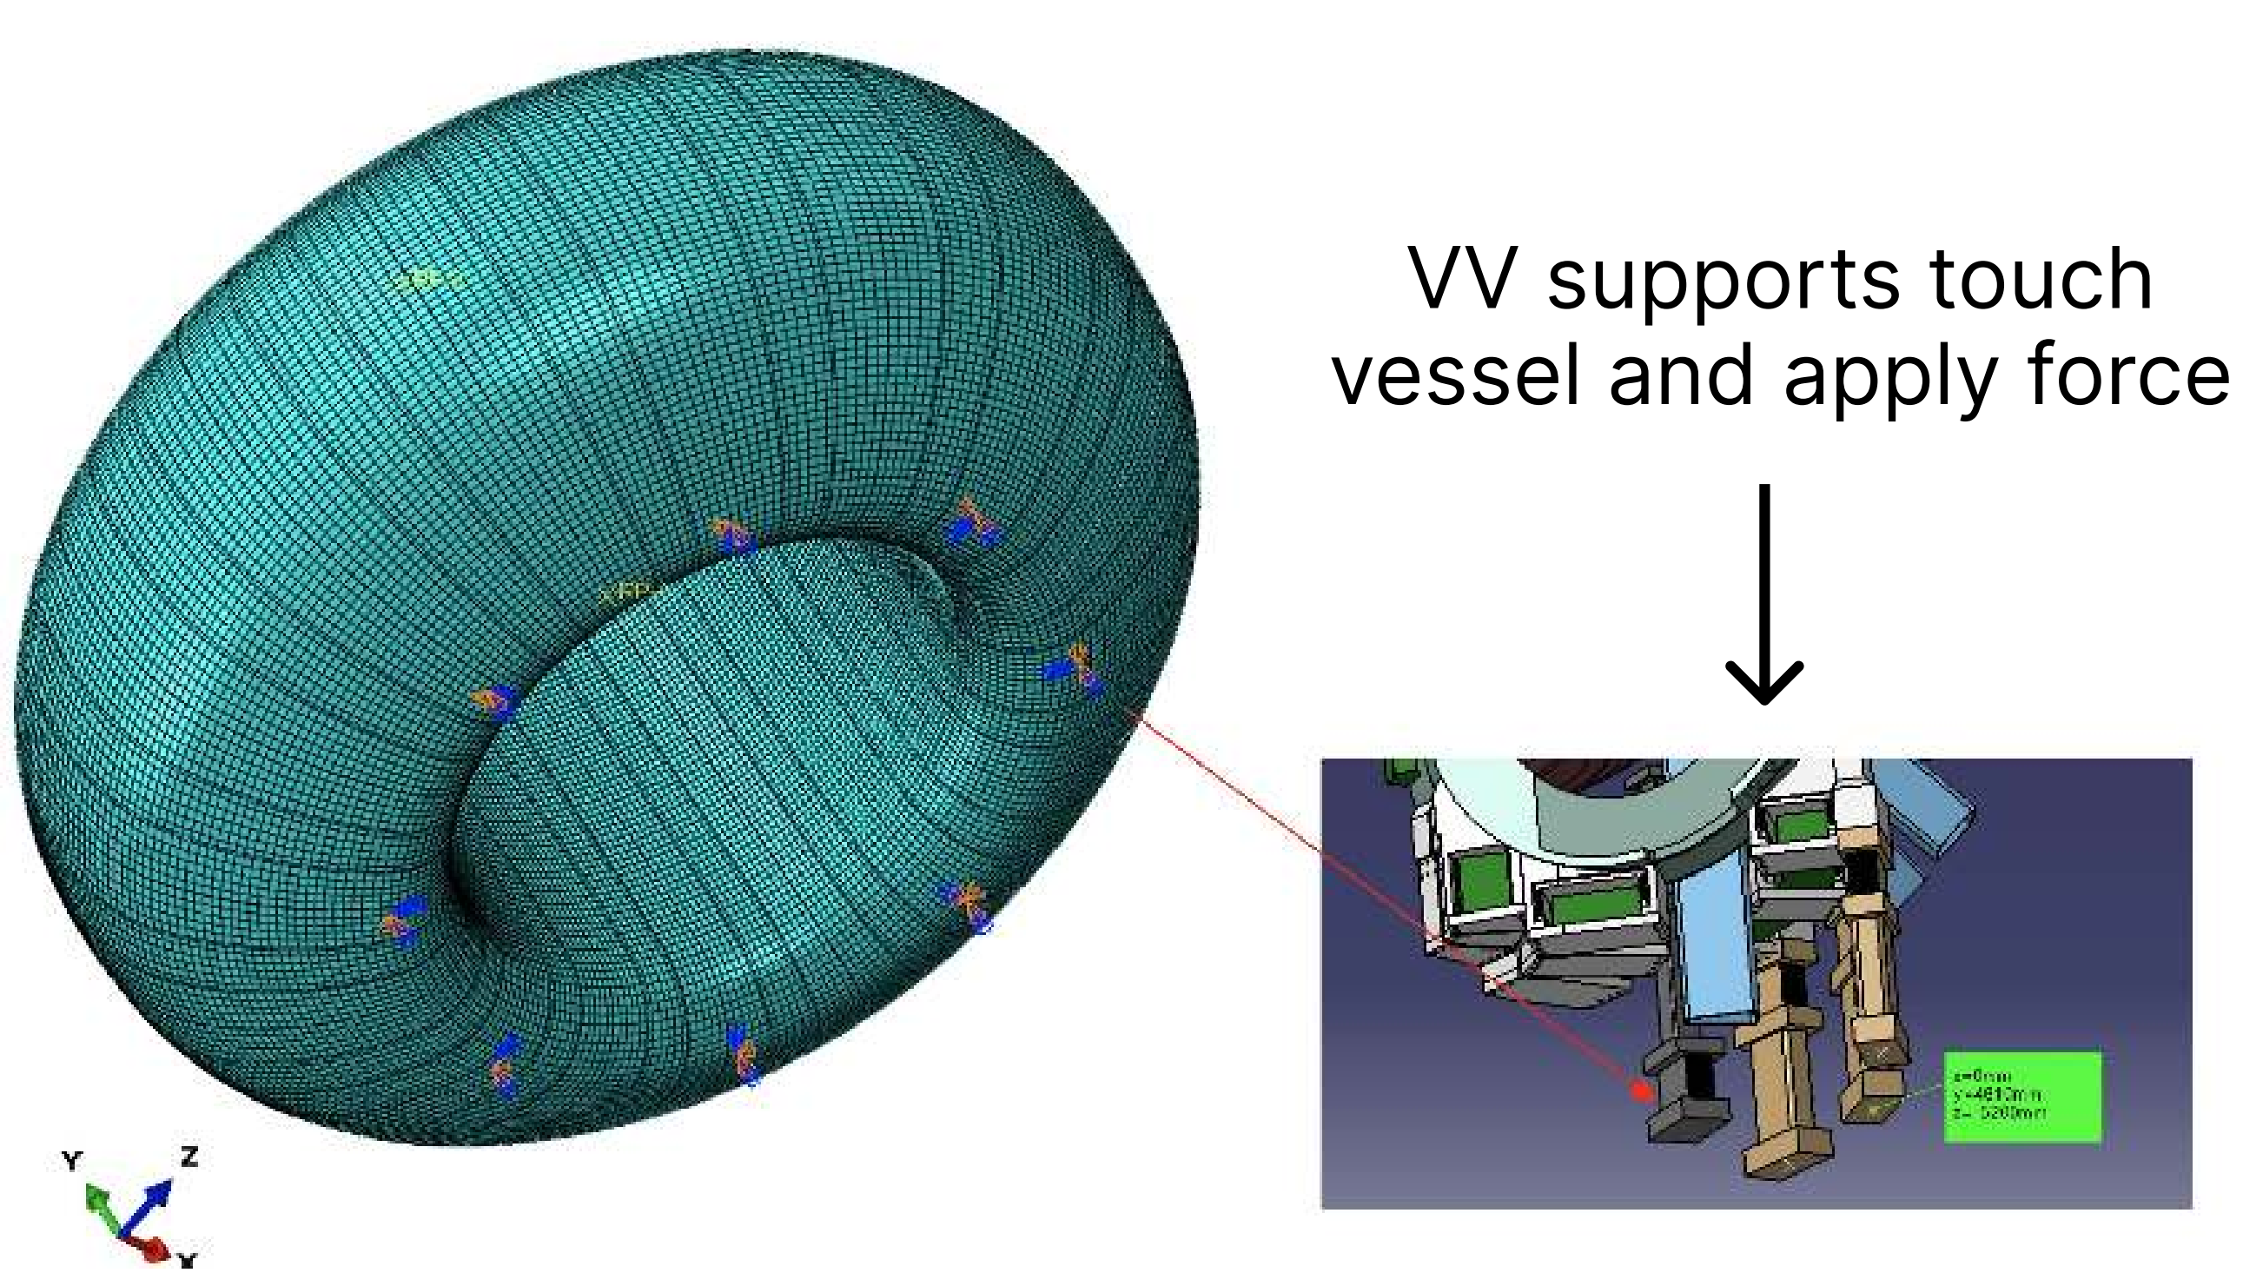

VV supports touch vessel and apply force

- Buckling analysis was performed via FEM to confirm the structural margin against crushing due to vacuum-induced loads. Simple linear buckling procedure with no imperfections has been used: thus, the results obtained shall be considered upper bound values, to be reduced to take into account knock-down factors typically requested by various calculation codes.
- Openings are expected to significantly reduce the buckling load calculated here for a vessel with no opening.
- The buckling analysis model is essentially the sector model copied/pasted 16 times to cover 360 degrees. Additionally, 8 supporting points have been considered at the bottom side, i.e. where the VV supporting columns will be connected.

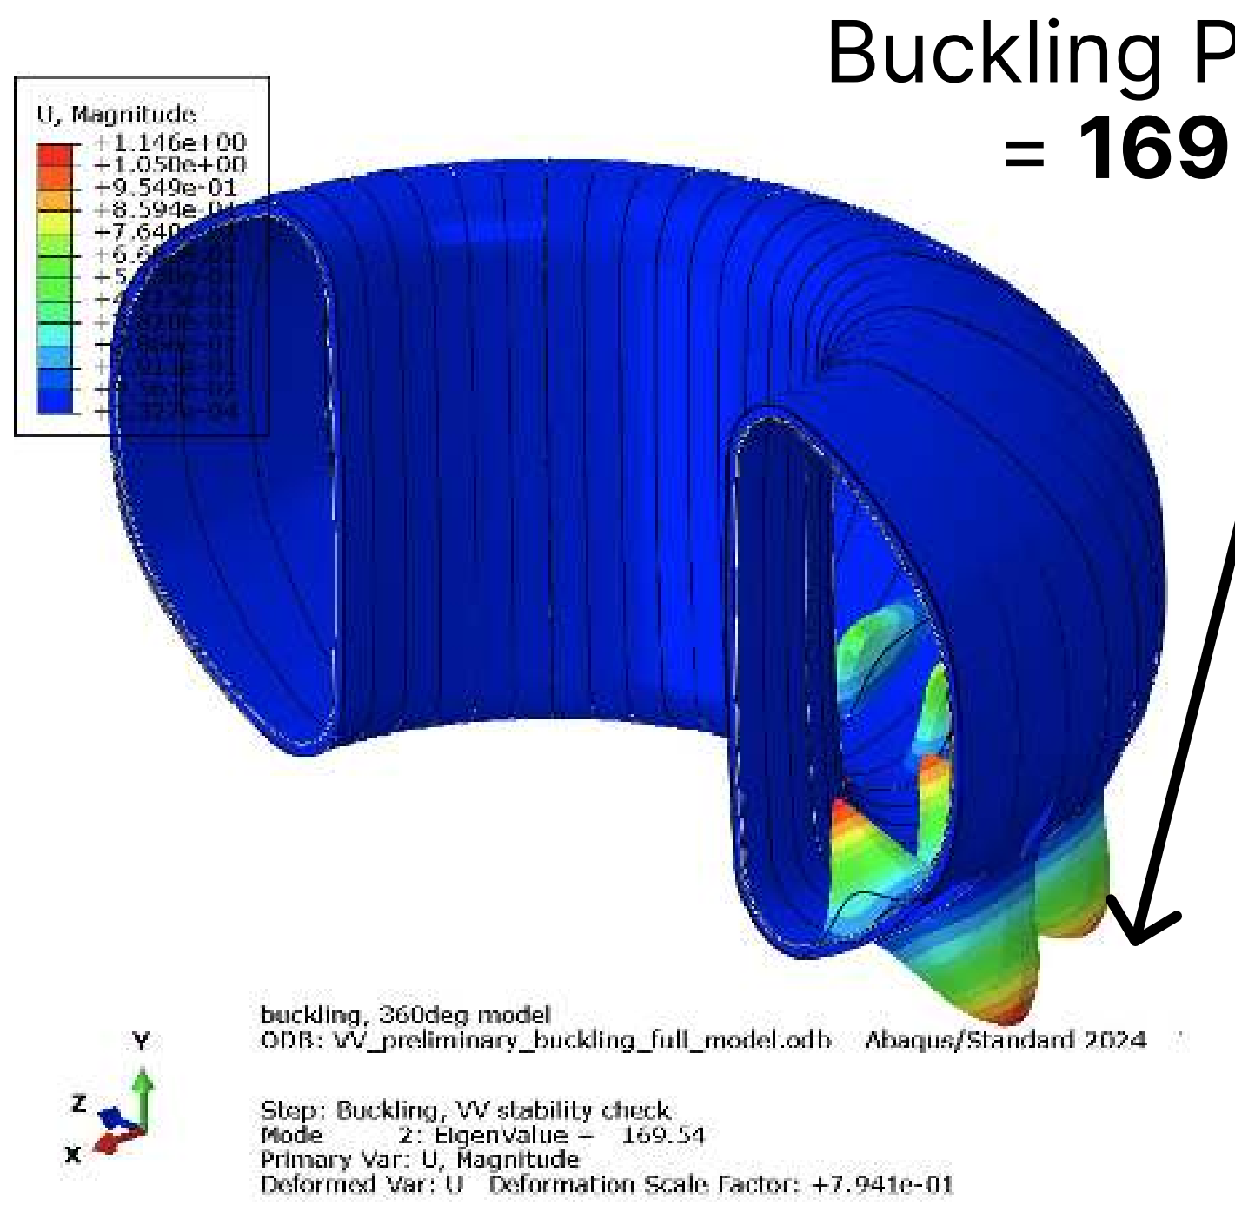

Buckling Pressure = **169bar**

Neutronics

In deuterium-tritium fusion, the primary high-energy particles produced are 14 MeV neutrons. These neutrons are responsible for carrying most of the fusion energy away from the reaction core. When they collide with the reactor's blanket, their kinetic energy is transferred to the working fluid, raising its temperature and ultimately converting the energy into usable power. This energy deposition not only contributes to the reactor's electrical output but also plays a crucial role in processes such as tritium breeding, where neutrons interact with lithium to breed additional fuel, to later be extracted from the coolant.

However, as the high energy neutrons traverse the reactor, not confined by the strong magnetic fields, they deposit energy in structural and ancillary materials, leading to effects such as displacement damage, embrittlement, and nuclear activation. These interactions present significant engineering challenges, as they can compromise the integrity and longevity of reactor components. To study the effects of this, we can run Monte Carlo simulations to determine the heating and nuclear activation in all components of the tokamak.

The **radial cross section of Yinsen** has the following stackup:

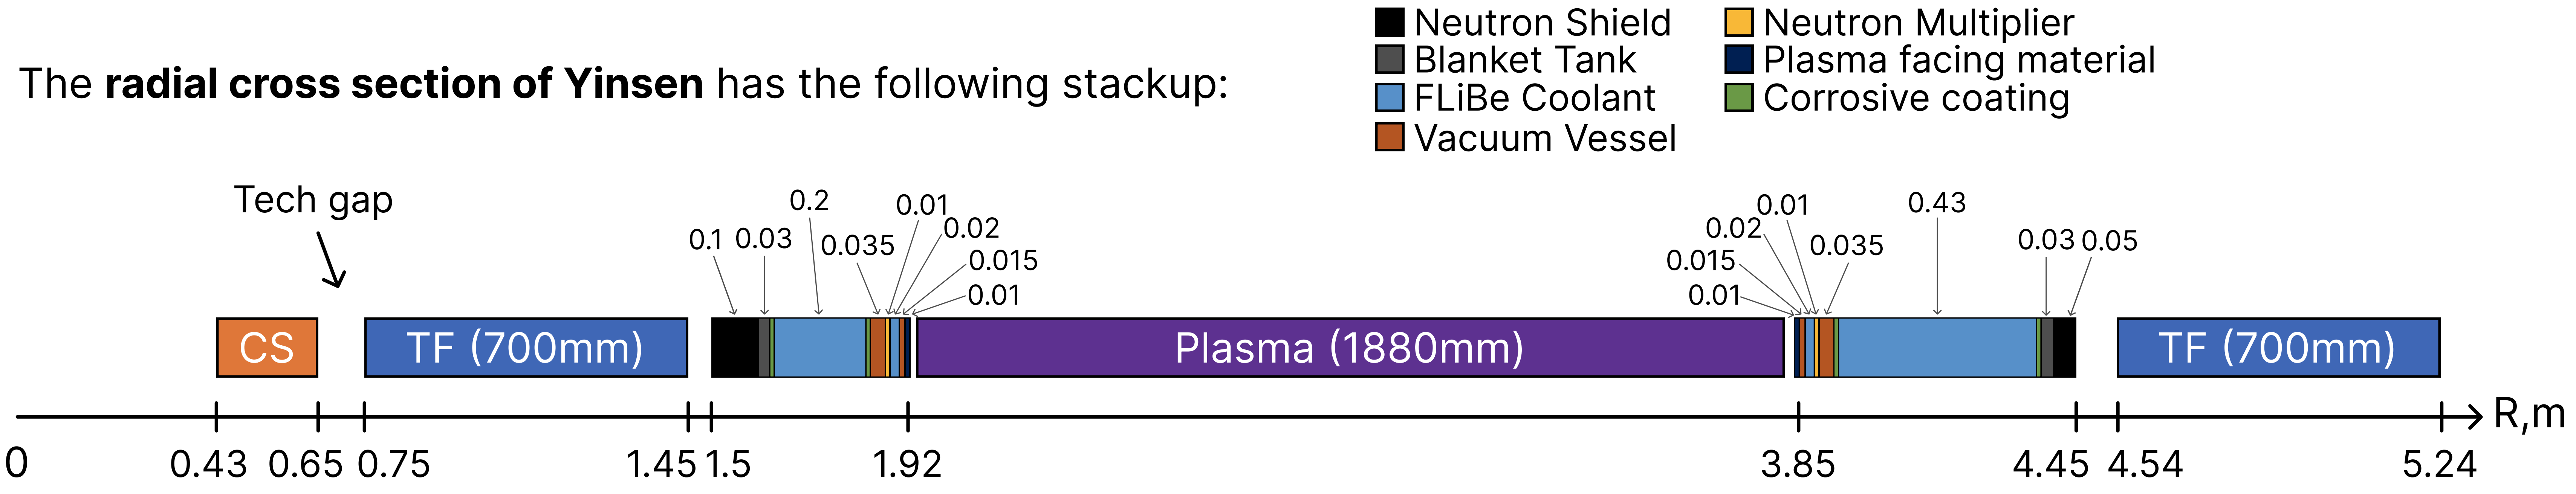

Using OpenMC we can estimate the heating, neutron flux, photon flux, and tritium breeding in all components of the reactor

Neutron Flux

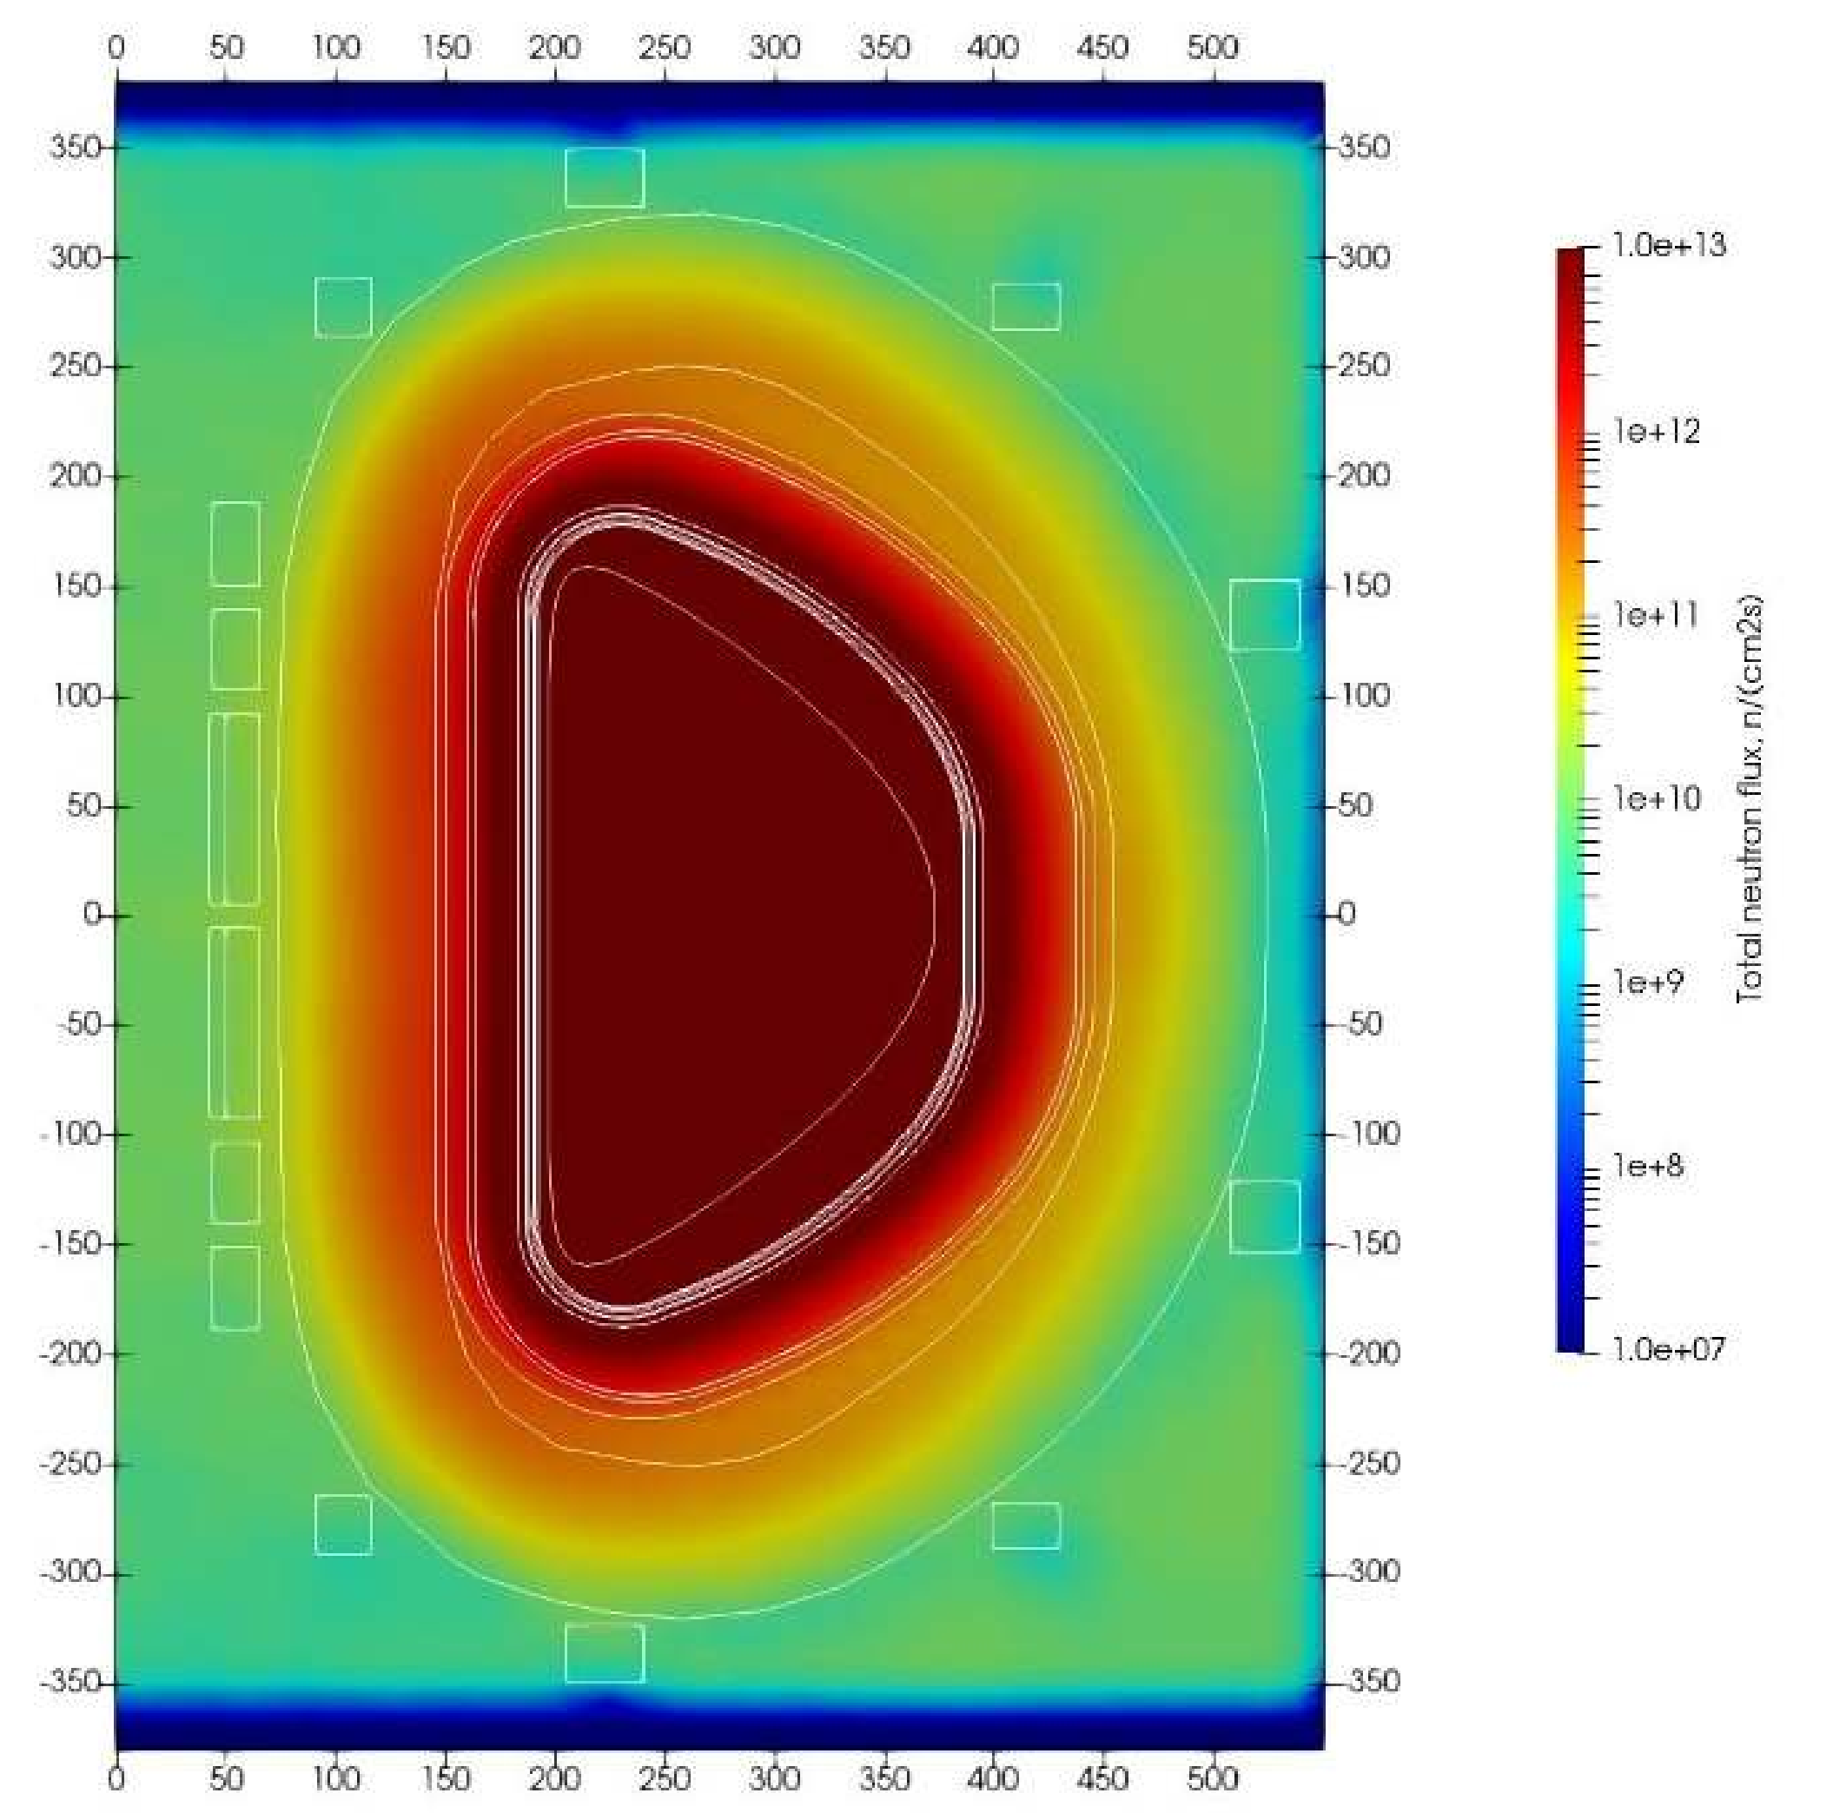

Photon Flux

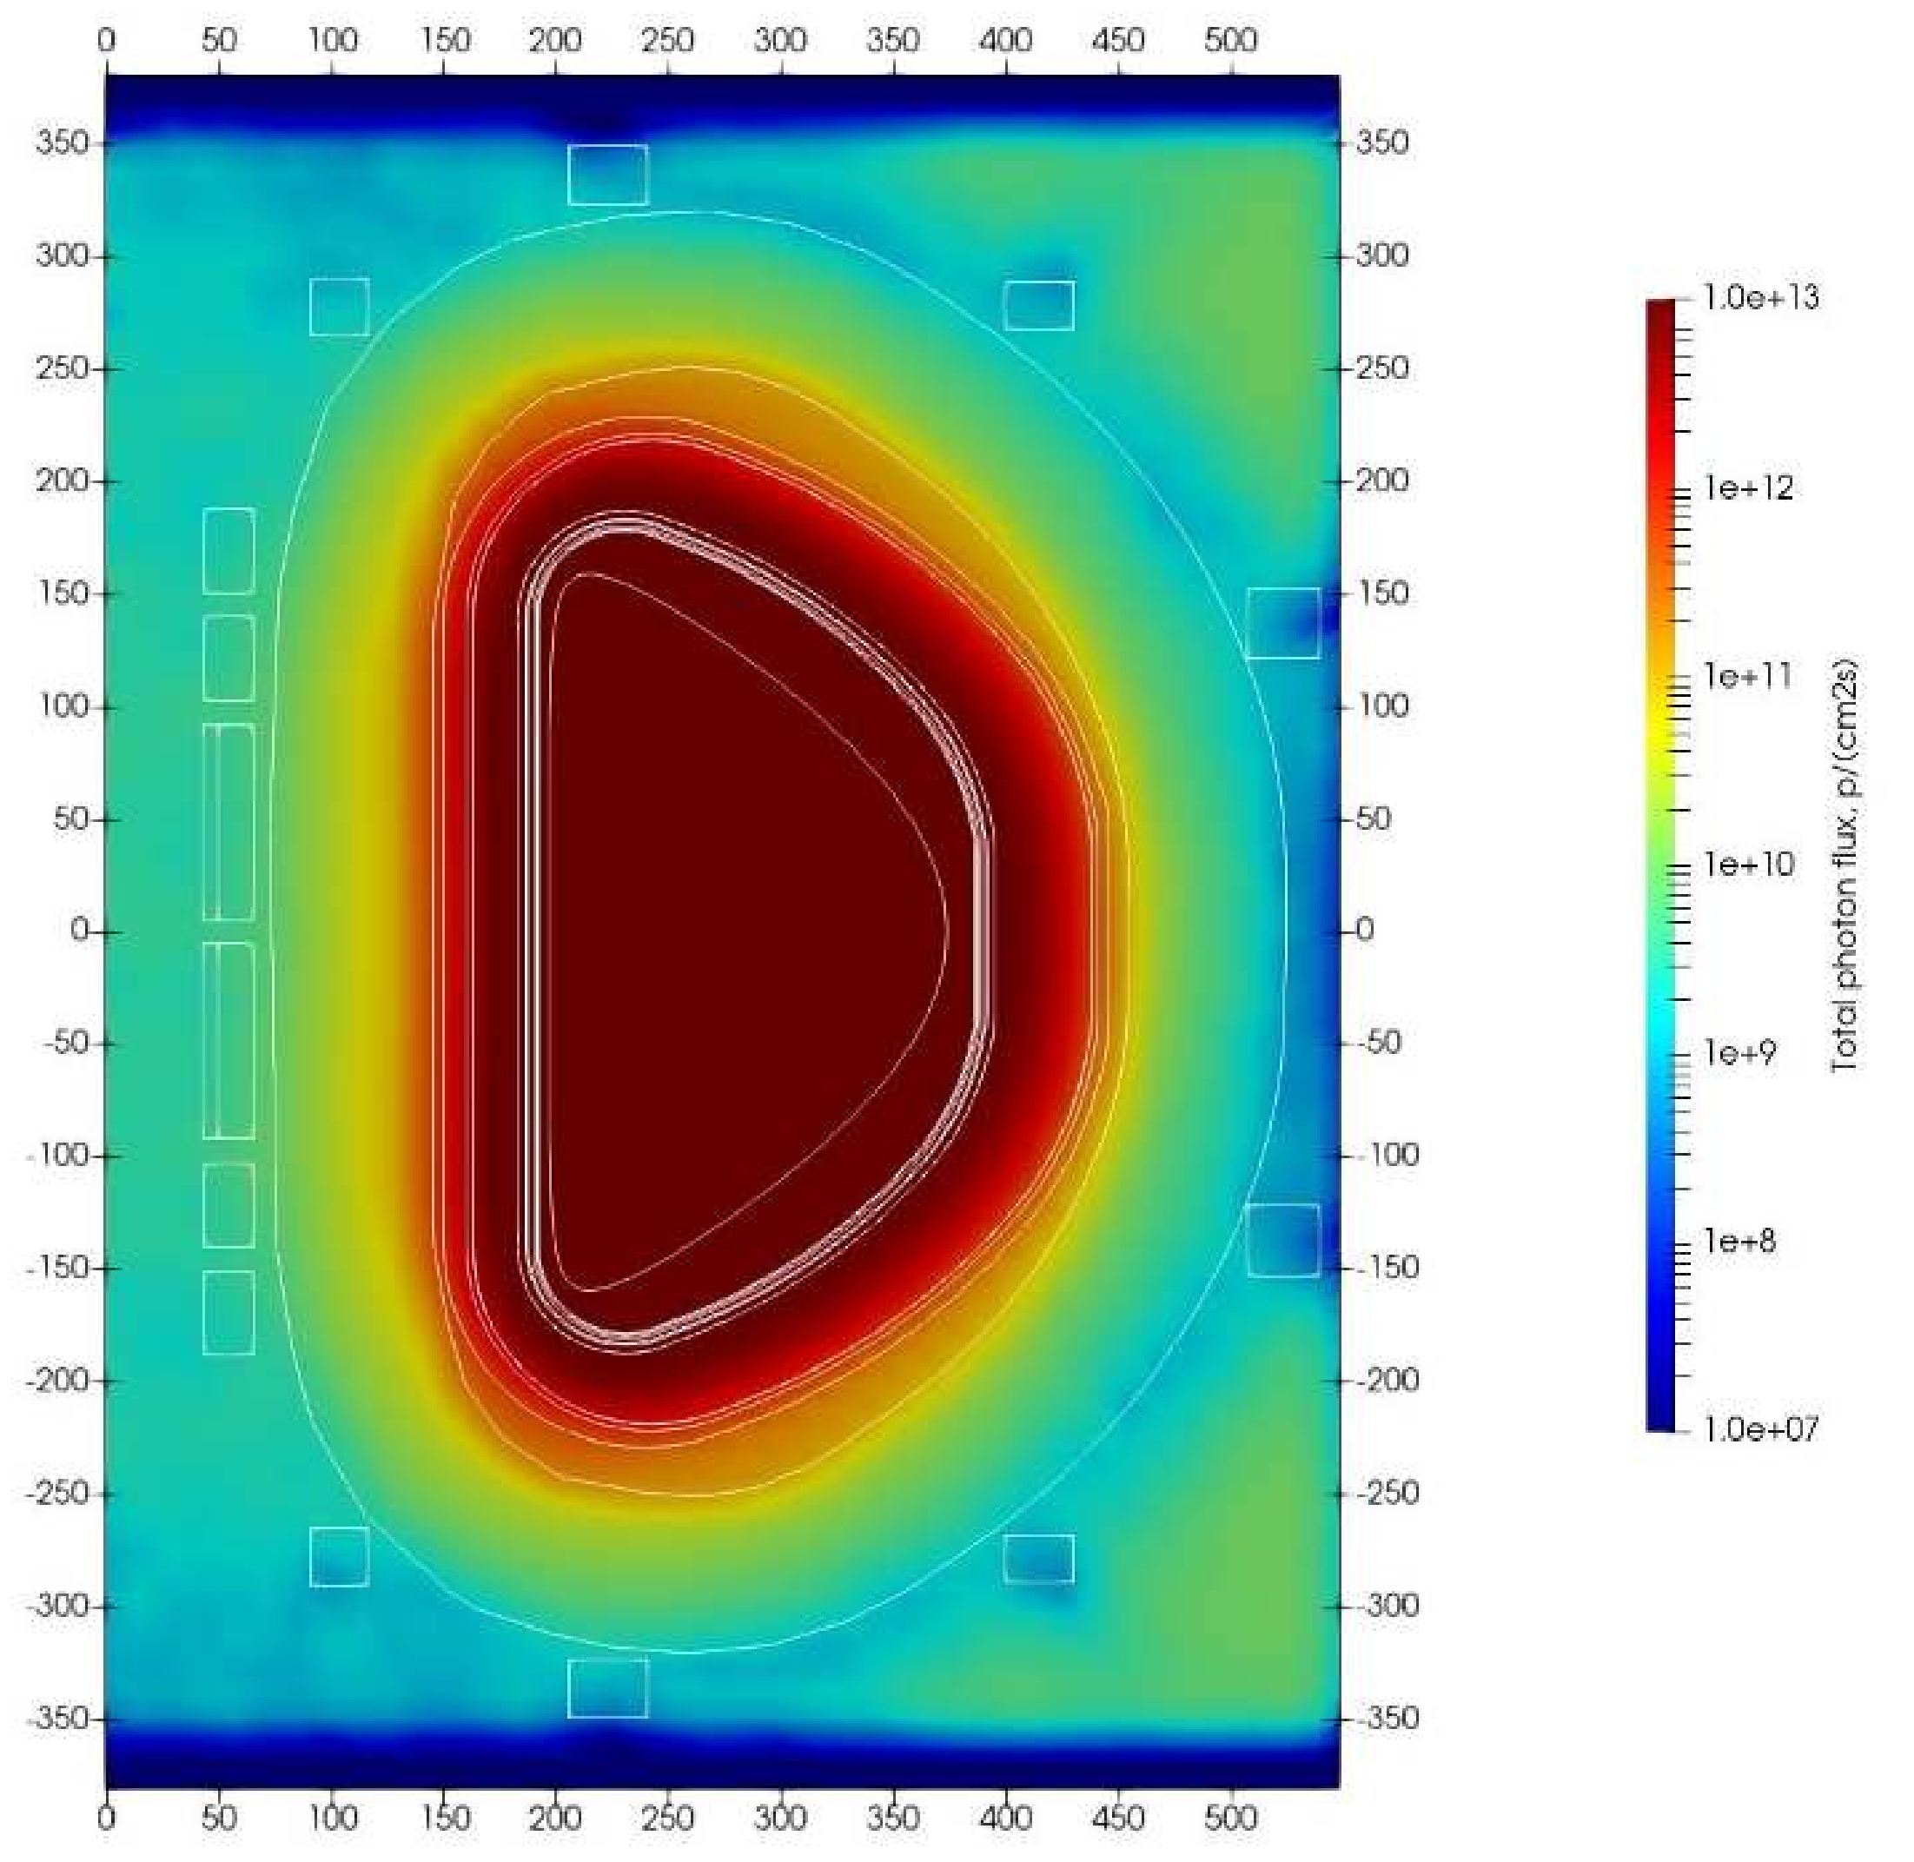

Tritium Breeding Ratio

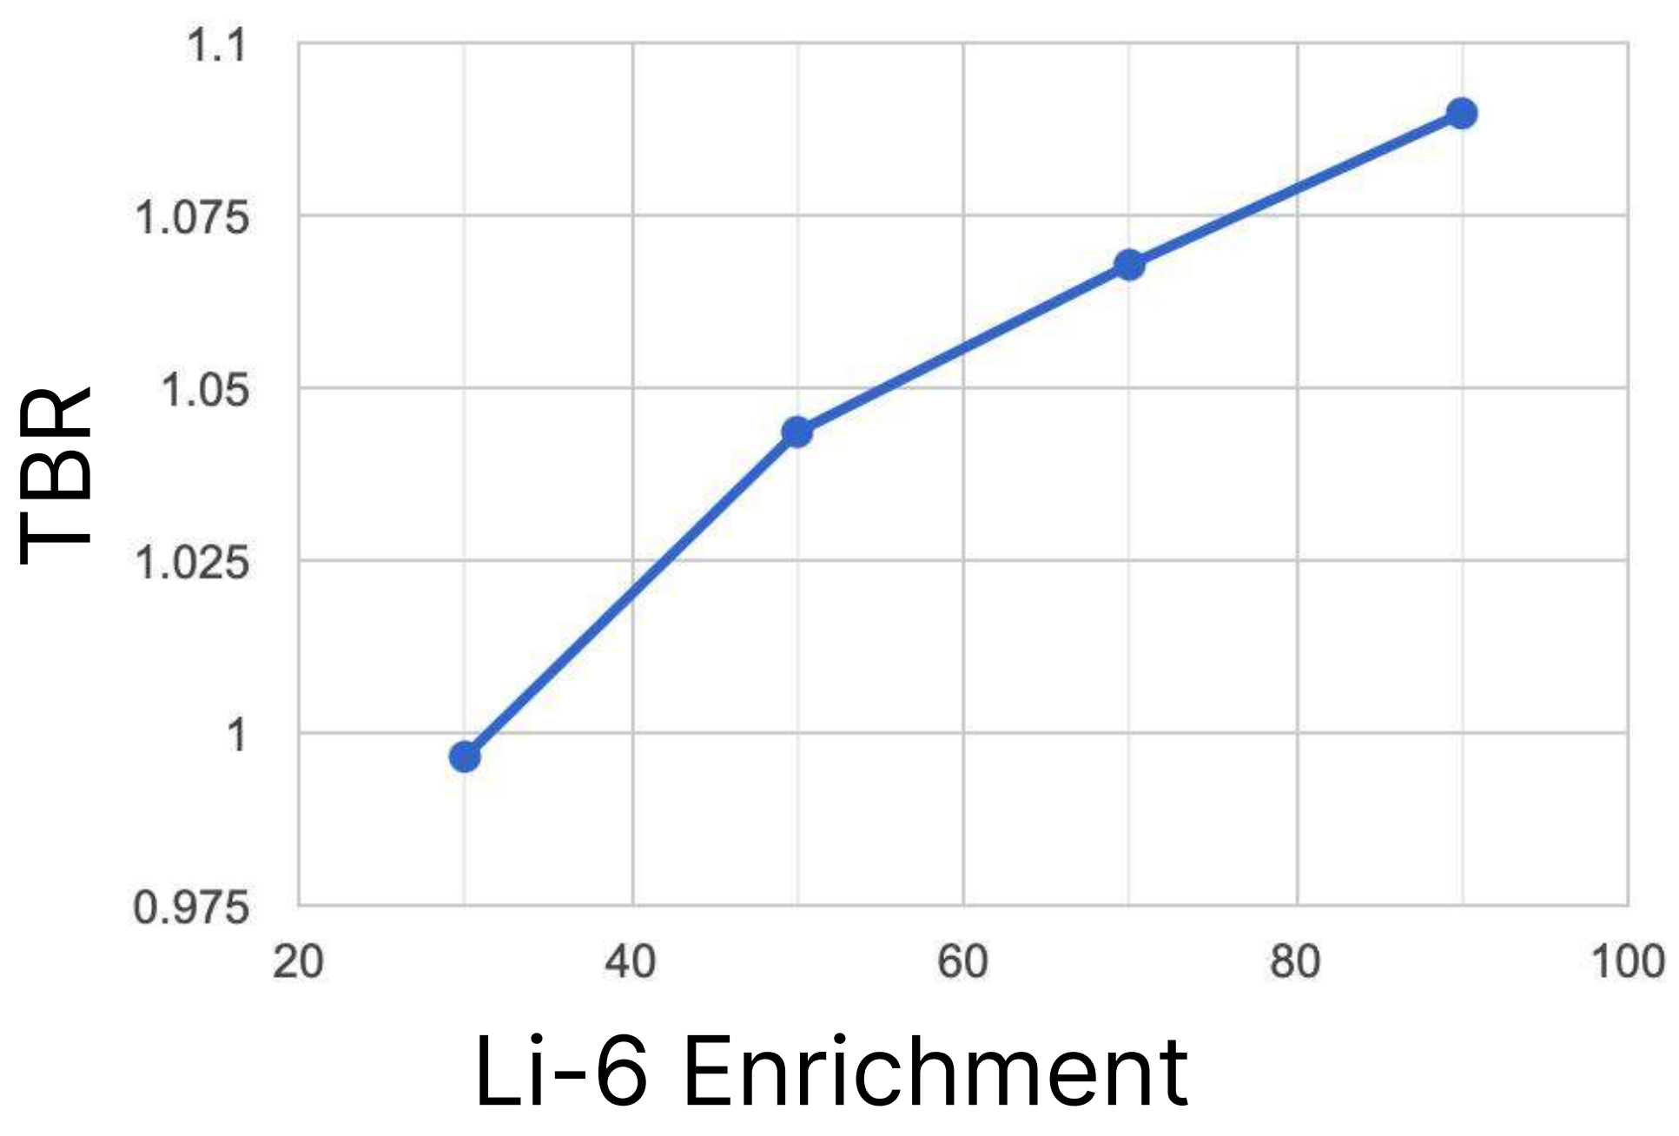

Takeaway

Nuclear heating and flux values are within expected ranges. The total heating to the TF magnets is low, on the order of 10kW with further room for optimization by trading FLiBe thickness for dedicated neutron shielding, particularly on the inner high field side (HFS). In the future, using SS316 with water shielding and refining the model to reduce material density will reduce energy release. Additionally, the Vacuum Vessel is projected to incur 1.2DPA/year, thus giving strong indication replacing the VV and magnets will not be required over the reactor's lifetime.

Component	Material
Neutron Shield	B_4C
Blanket Tank	Inconel
Blanket	FLiBe
VV Inner	Inconel
Multiplier	Be
Blanket	FLiBe
VV Outer	Inconel
First Wall	W
Magnets CC	REBCO

Heating

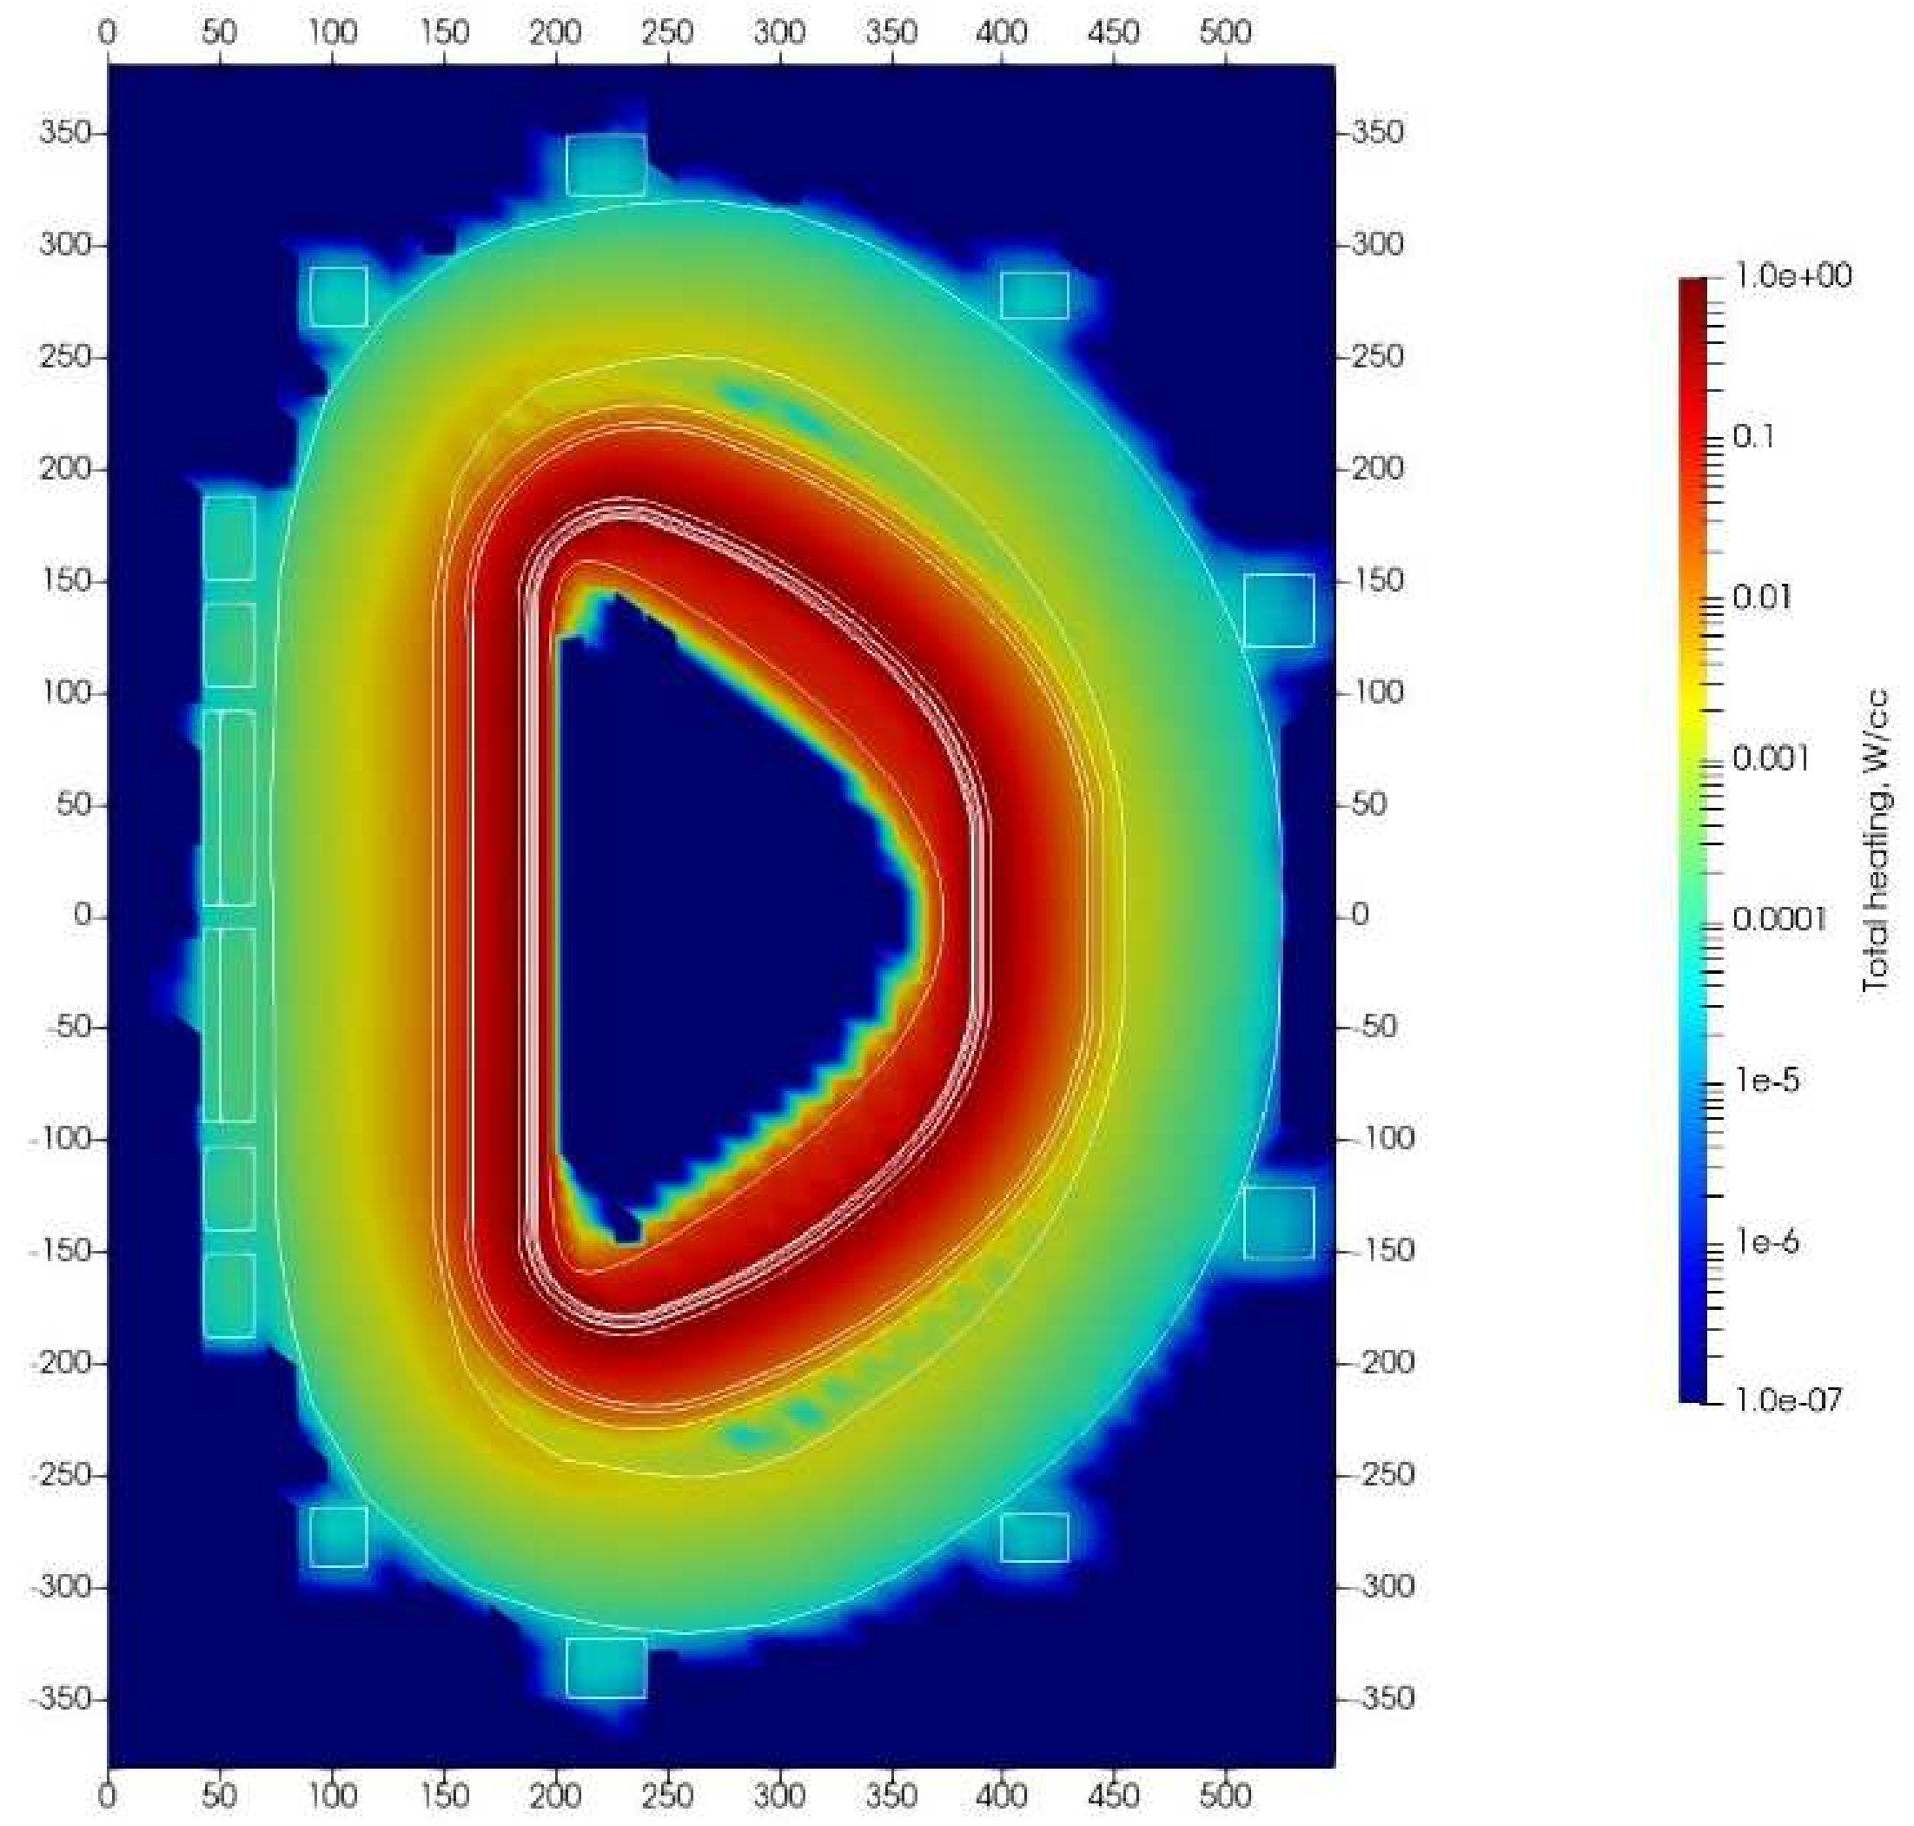

Tritium Generation

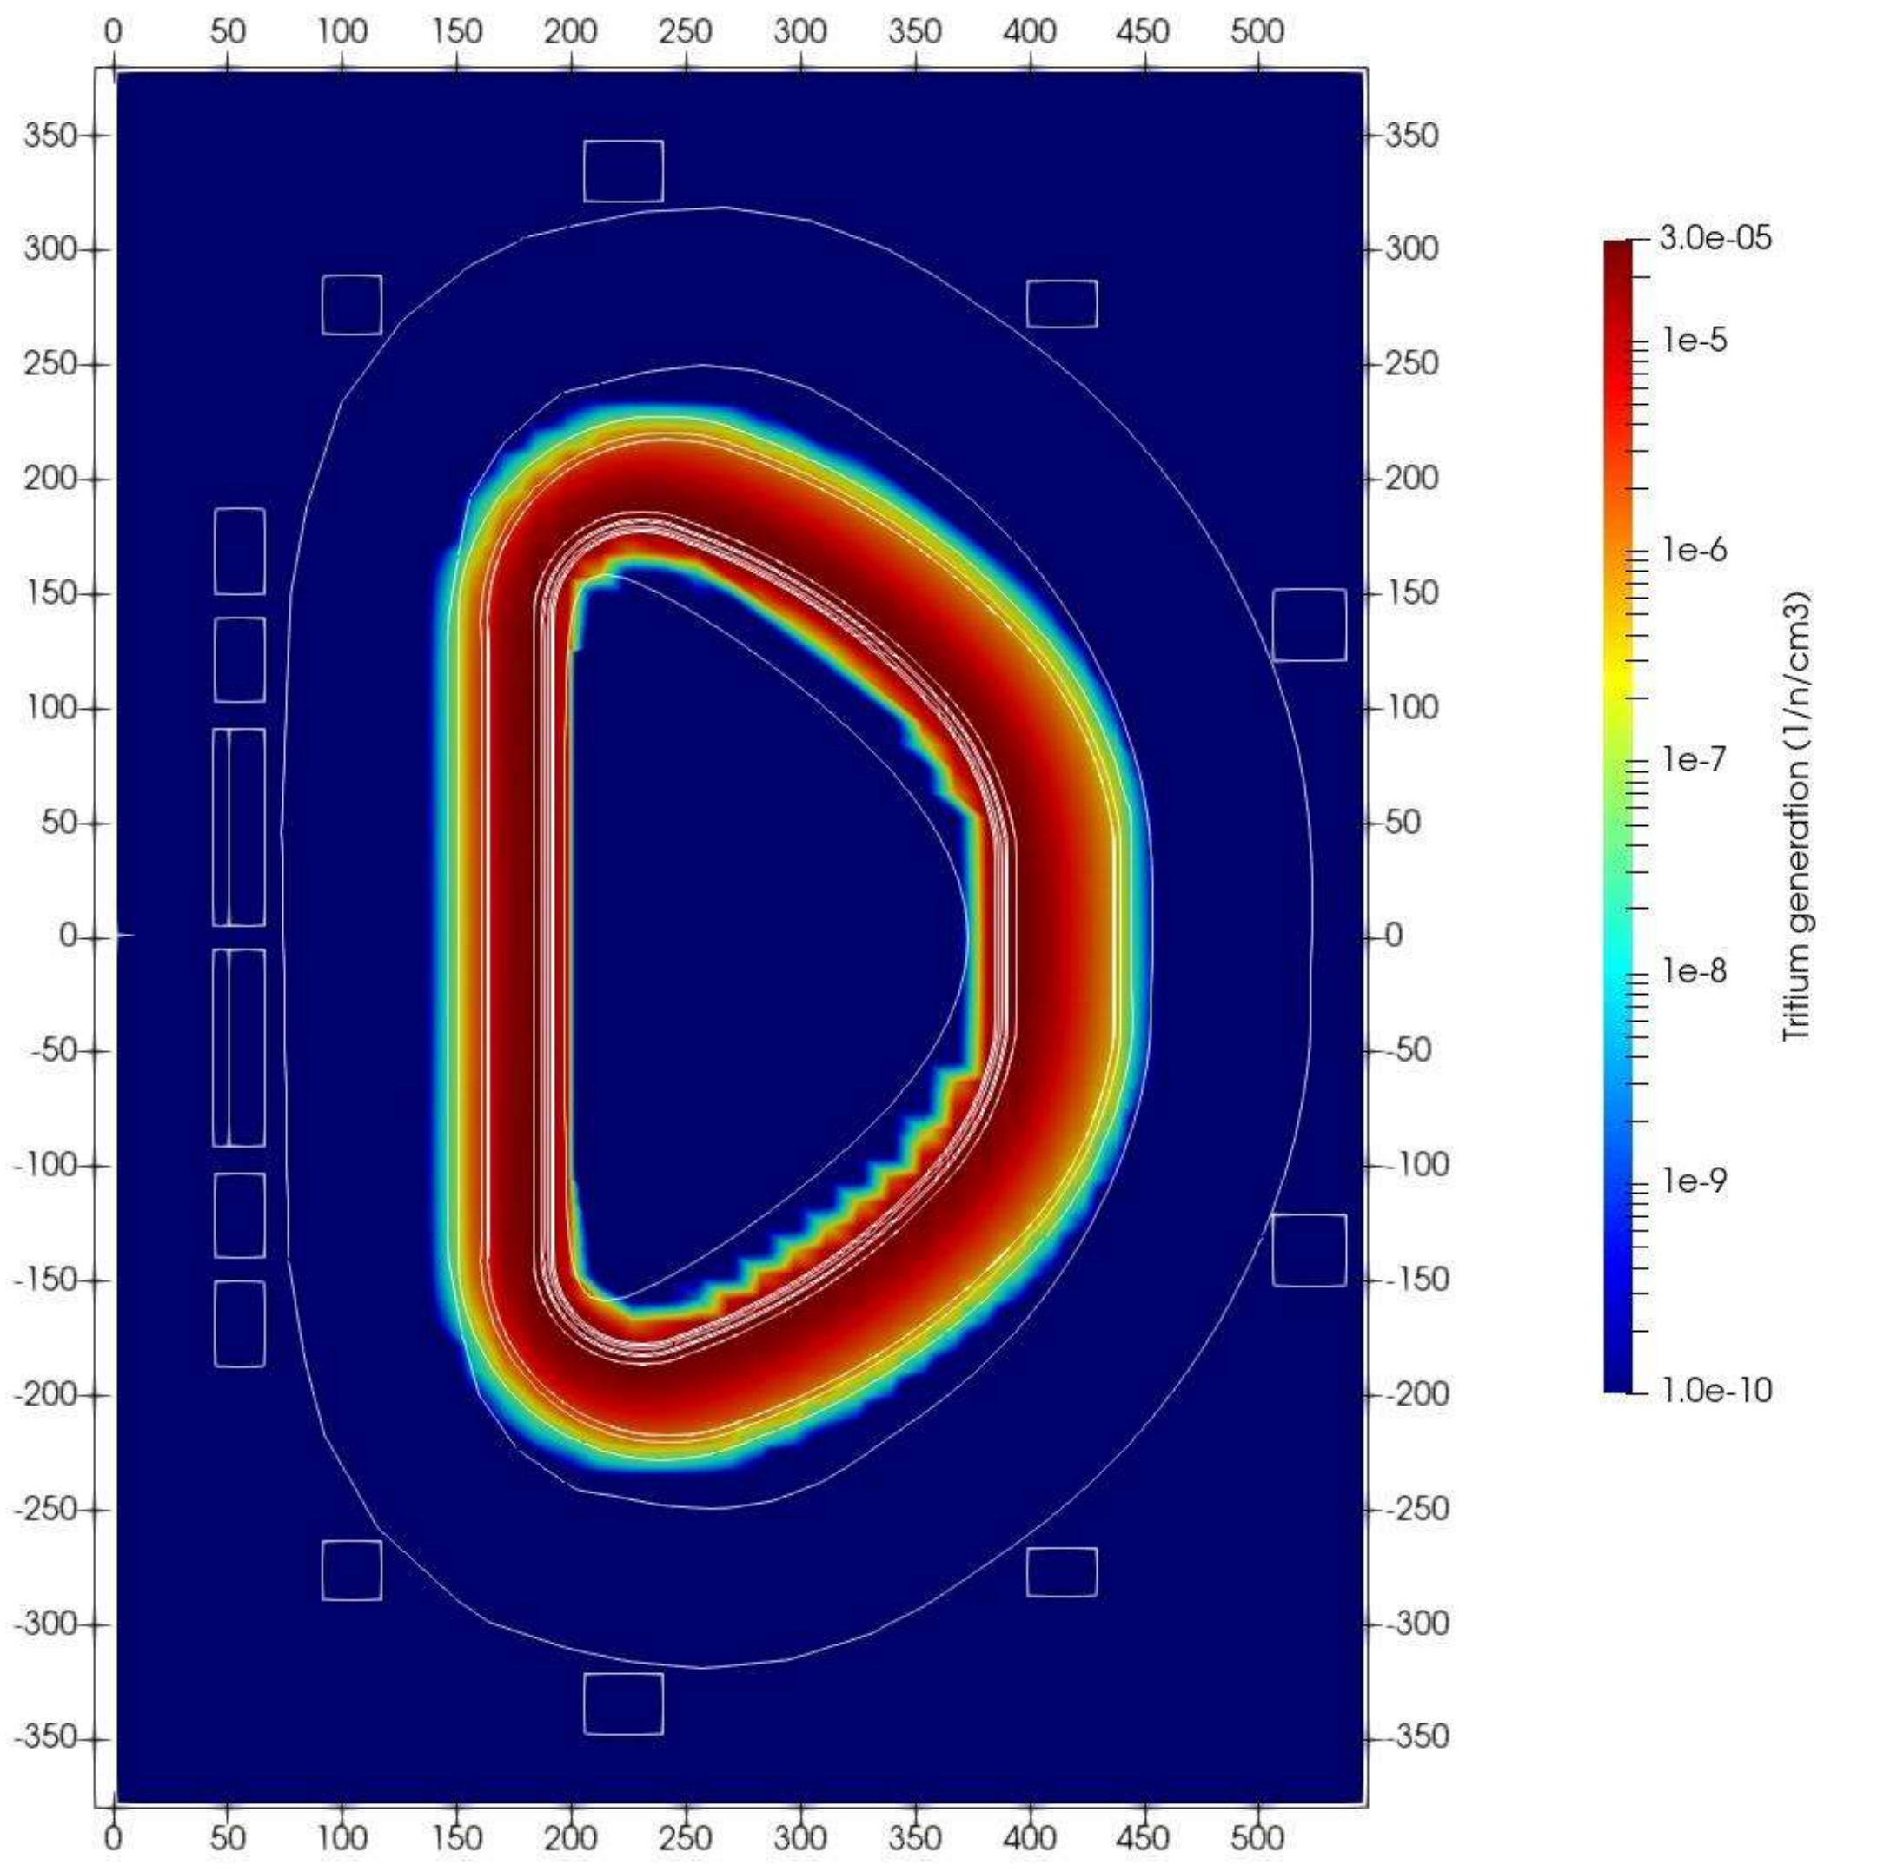

Cryogenics

HTS magnets are only superconducting at below certain temperatures, so cooling systems are needed. The cooling system we will use is based on a **Reverse Turbo Brayton (RTB) cycle**. The principles of operation for this cryogenic system are:

- Compress a fluid (helium, in our case) using centrifugal compressors.
- Cool down the fluid (in the recuperative heat exchangers). The cycle can include a “pre-cooling” loop using Liquid Nitrogen. The use of Liquid Nitrogen as a pre-cooling loop is not mandatory but can increase efficiency.
- Expand the fluid through a high-speed turbine.
- Use this cold helium for the cooling down of the application (through an secondary loop or Application Loop).

To be able to size a cryocooler able to deliver **20 kW at 20K**, we need to determine associated boundary conditions, corresponding to the Application Loop:

- Helium temperature at Cryocooler Inlet (or application outlet, (6) on the sketch below)
- Helium temperature at Cryocooler Outlet (or application inlet, (7) on the sketch below)
- Helium mass flow rate through application

Overview of a RTB cycle

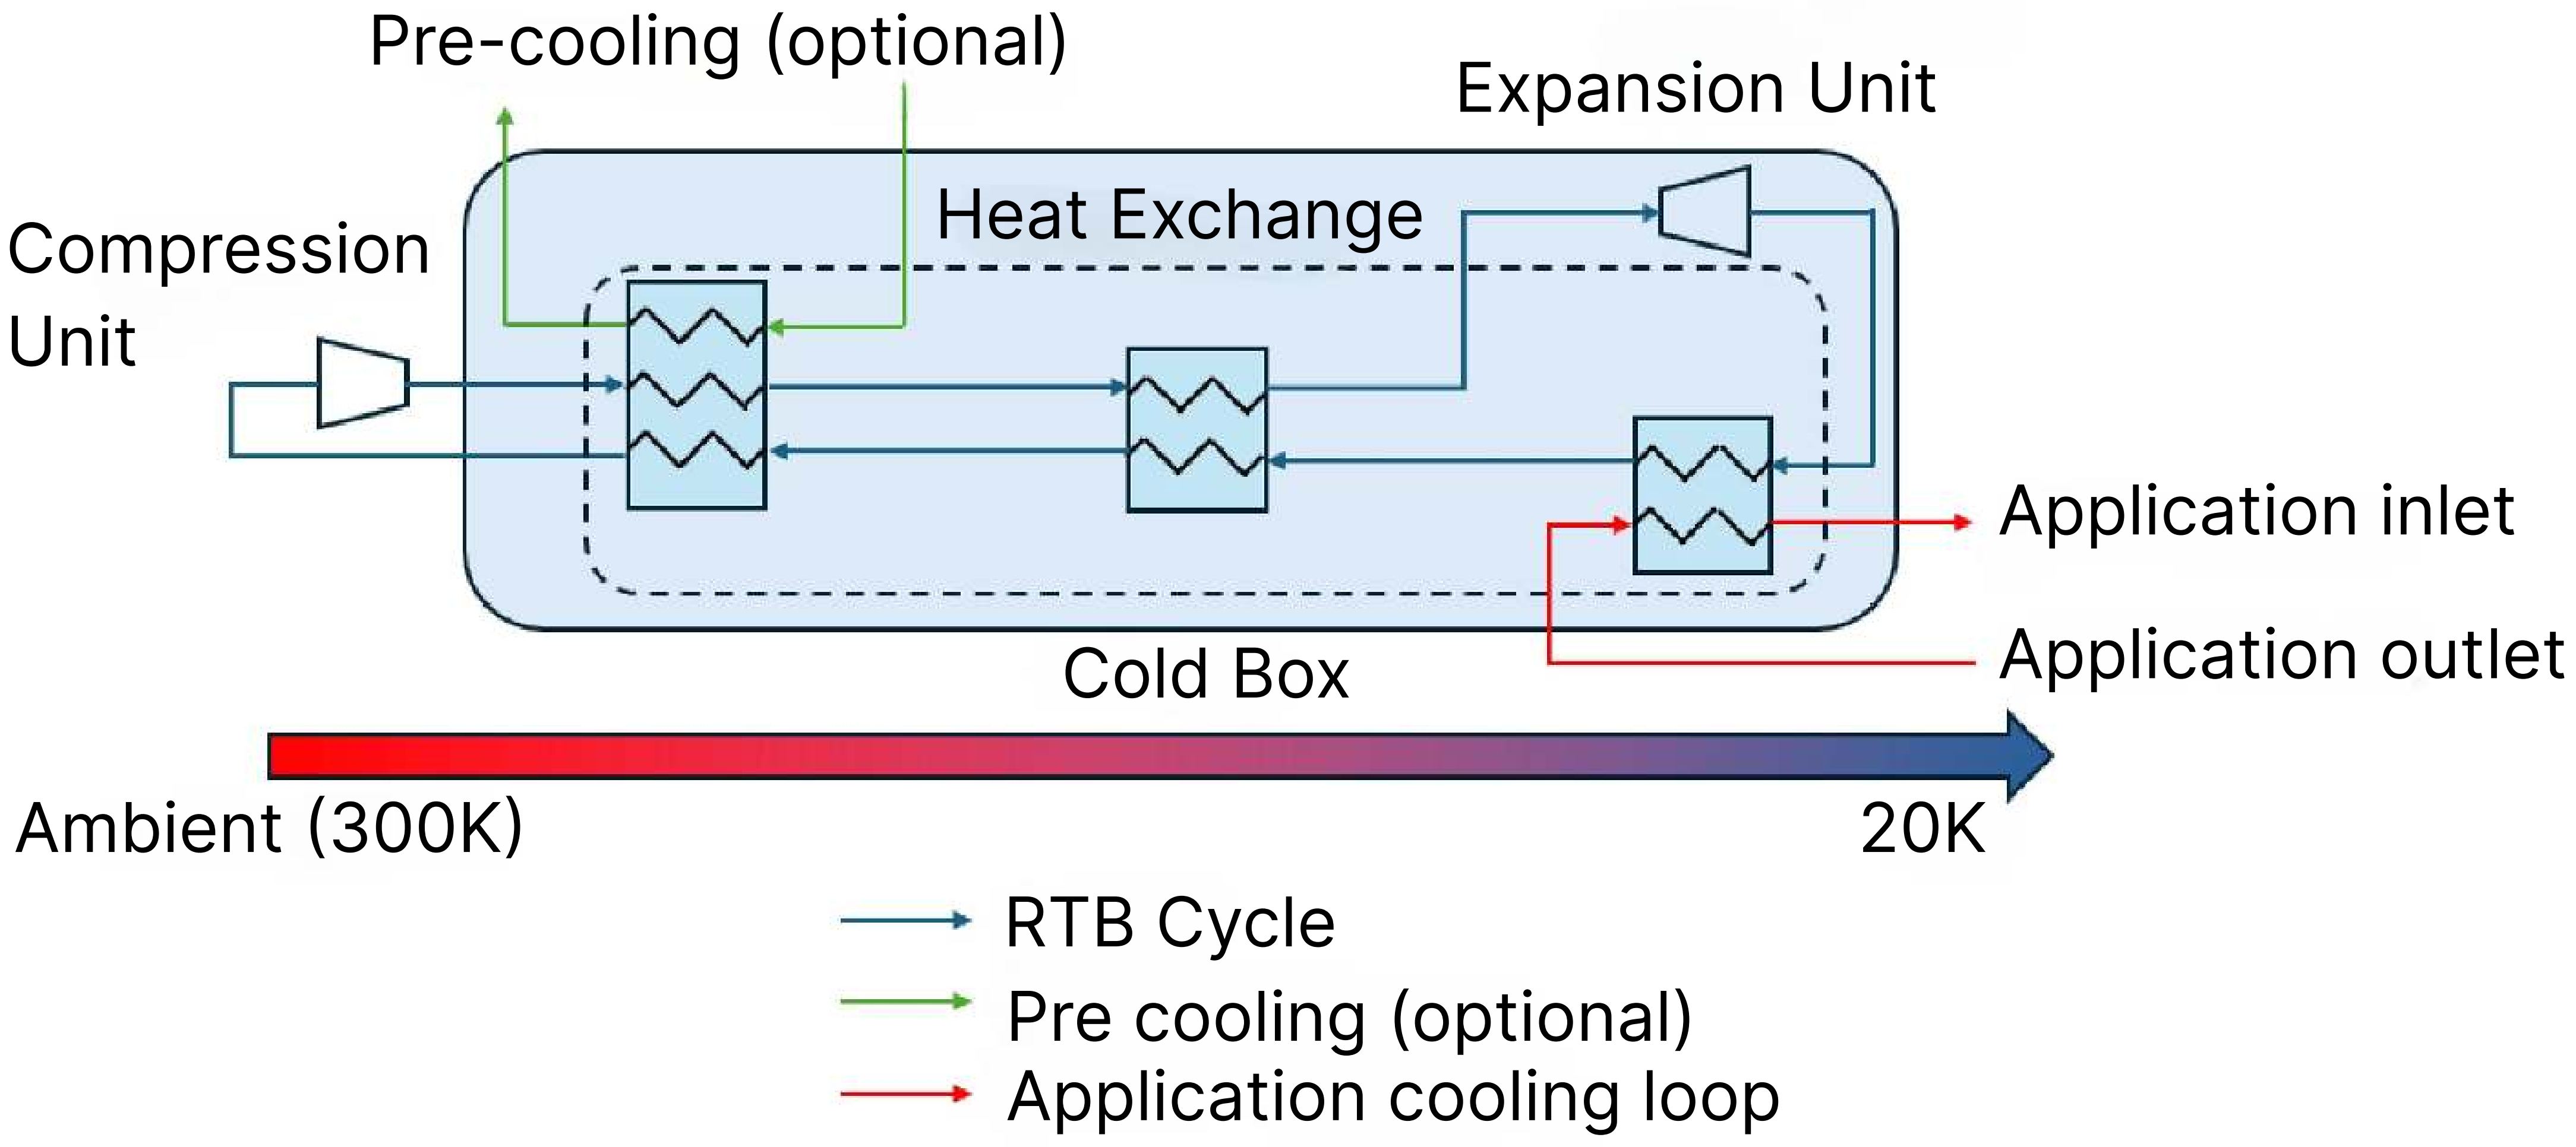

Main Loop (Numbered)

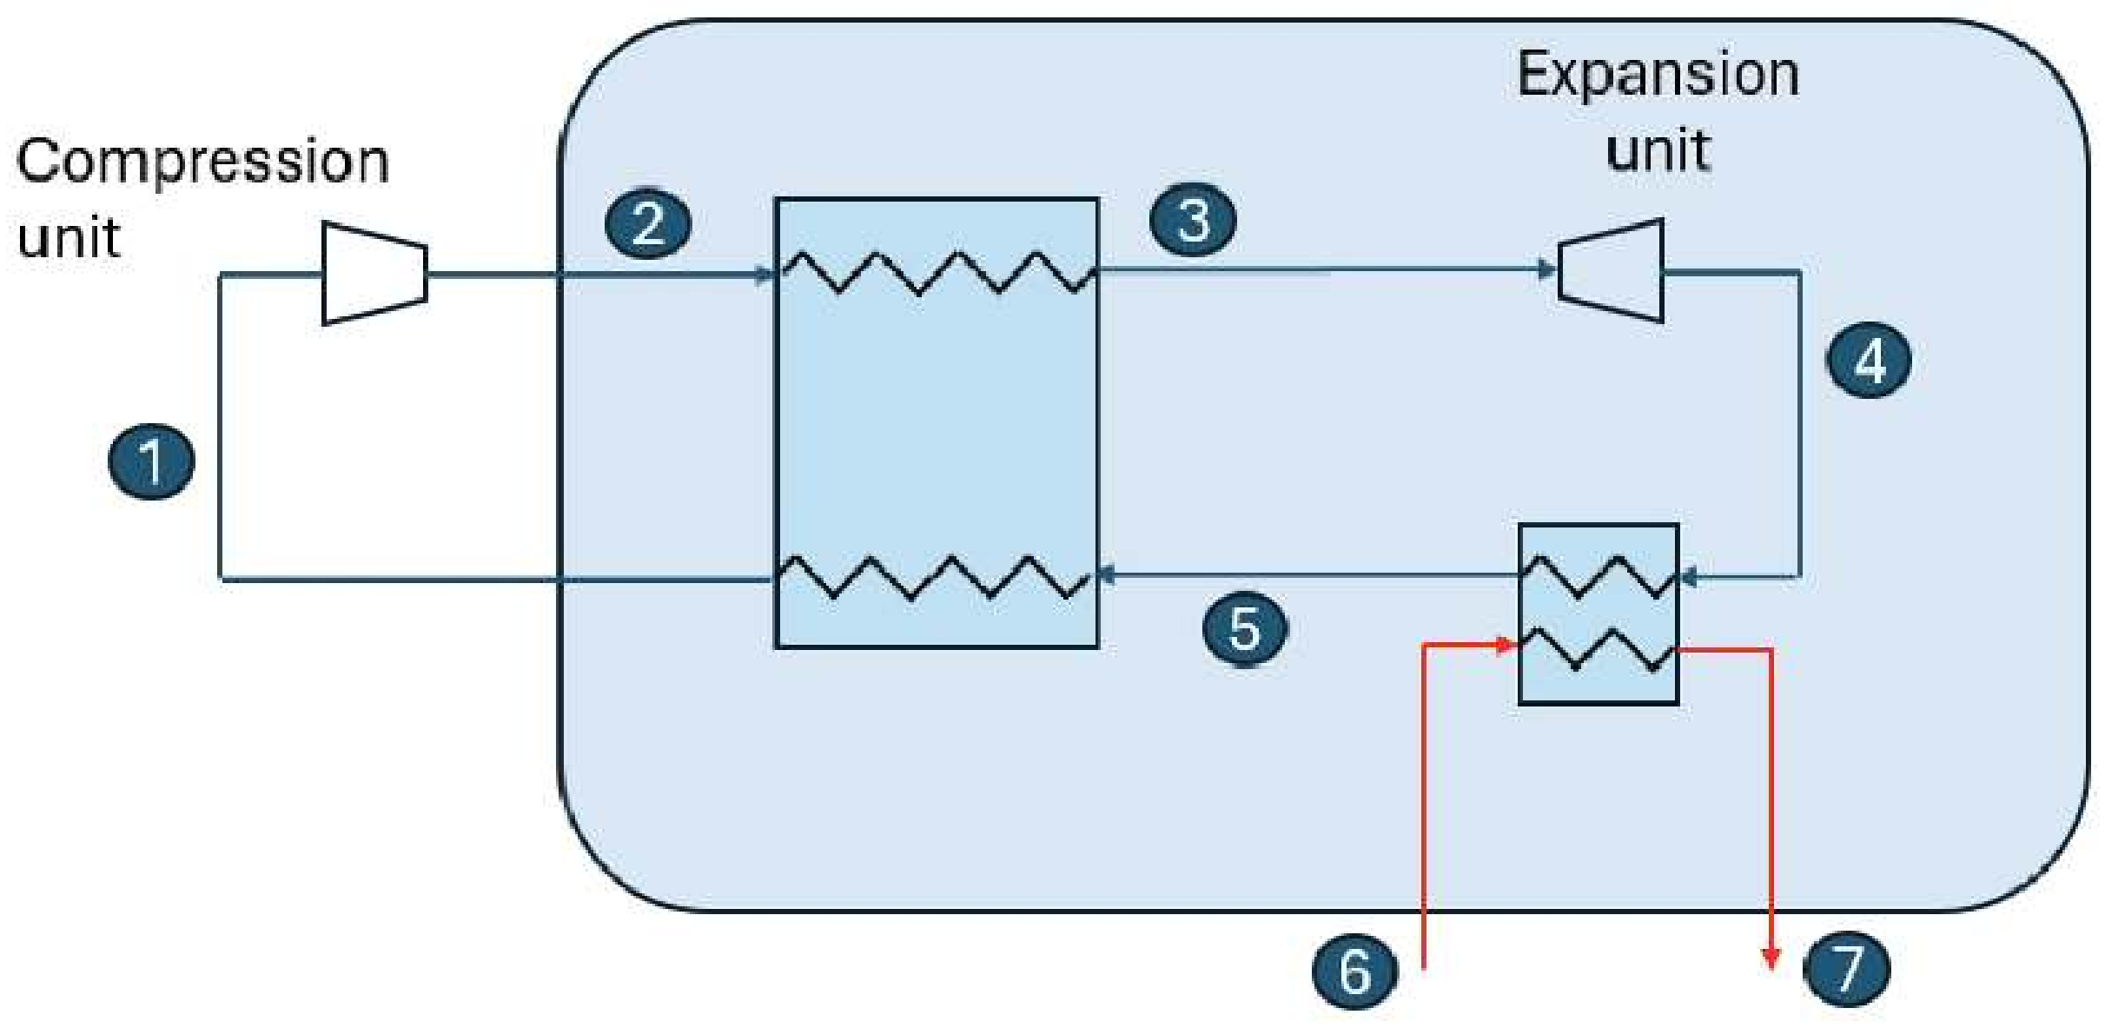

The following cases are considered:

Parameter	Case 1: 23 K	Case 2: 33 K
Helium temperature at Cryocooler Inlet or application outlet	23 K	33 K
Helium temperature at Cryocooler Outlet or application inlet	19 K	29 K
Pressure of helium in application loop	20 bar	20 bar

Based on these cases we can calculate the helium mass flow rate to reach the target cooling power (20kW) :

Parameter	Case 1: 20 K	Case 2: 30 K
Calculated helium mass flow rate through application and cryocooler	860 g/s	910 g/s

An increase of application temperature from 23K to 33K decreases the cryocooler electrical consumption from 1.3MW to 820kW (37% reduction) to provide 20kW of cooling power. Using the maximum electrical consumption in the motor to run the centrifugal compressors can provide either 20kW @ 23K or 32 kW @ 33K. If further footprint reduction is needed, use of Liquid Nitrogen will be considered as well.

Parameter	Case 1	Case 2
Maximum application temperature	23 K	33 K
Electrical consumption to provide 20 kW	1.3 MW	820 kW
Maximum cold power for an electrical power of 1.3 MW	20 kW @ 23 K	32 kW @ 33 K
Delta-T in application loop	4 K	4 K
The footprint of such a system	4 m × 14 m × 2.5 m	4 m × 12 m × 2.5 m

Because the TF magnets in Yinsen are < 9T, the current density required is reduced, allowing us to operate at hotter cryogenic temperatures while still below the superconducting critical current density, thus improving overall efficiency.

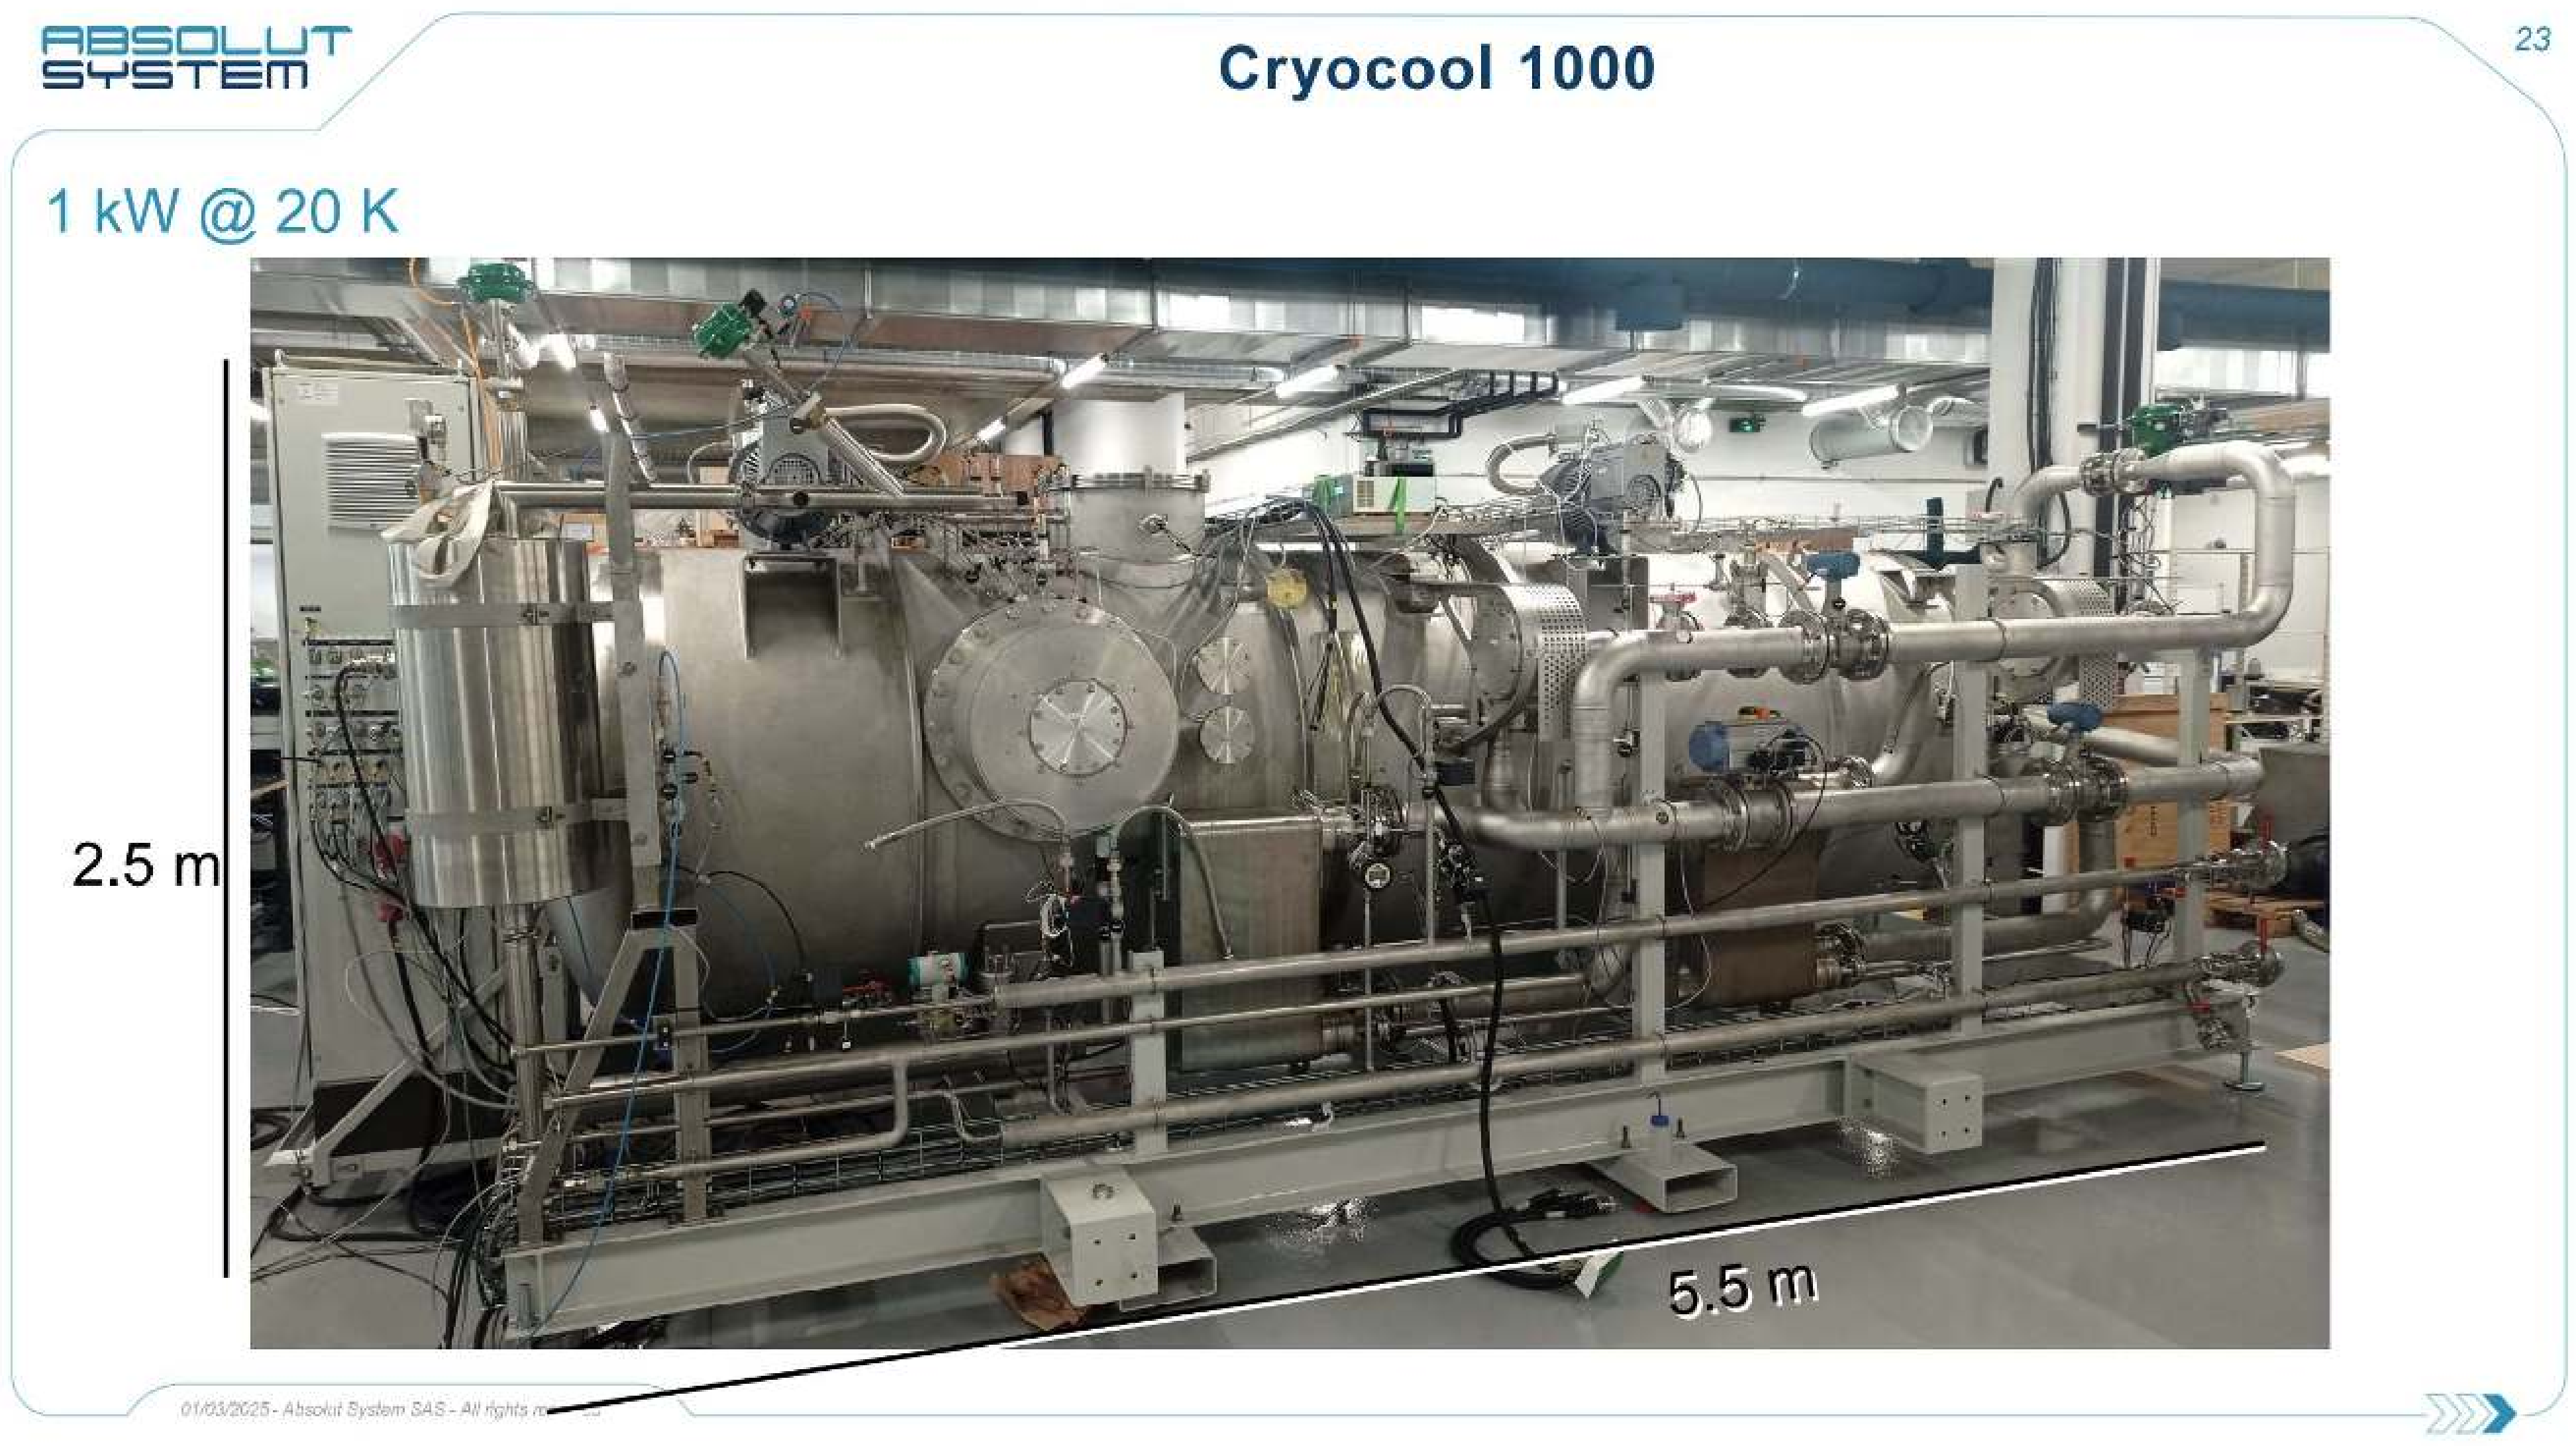

Cryocool 1000 device and size -- *Source*: Absolut Systems

Layout/Floorplanning

A major engineering challenge that we are solving to fit these devices on a ship is the size. Current facilities like ITER take up a full football field. To fit inside a large 20,000TEU vessel we must be below 100,000m³, this is the current maximum size of the fuel tanks being considered for alternative fuel storage.

Below is a rendering of how these systems may fit inside the ship. The exact spacing here is subject to change, but the relative size of the auxiliary systems compared to the tokamak itself is the main point to highlight here. We just want to show that **it is feasible to fit a tokamak fusion system on a large commercial class shipping vessel.**

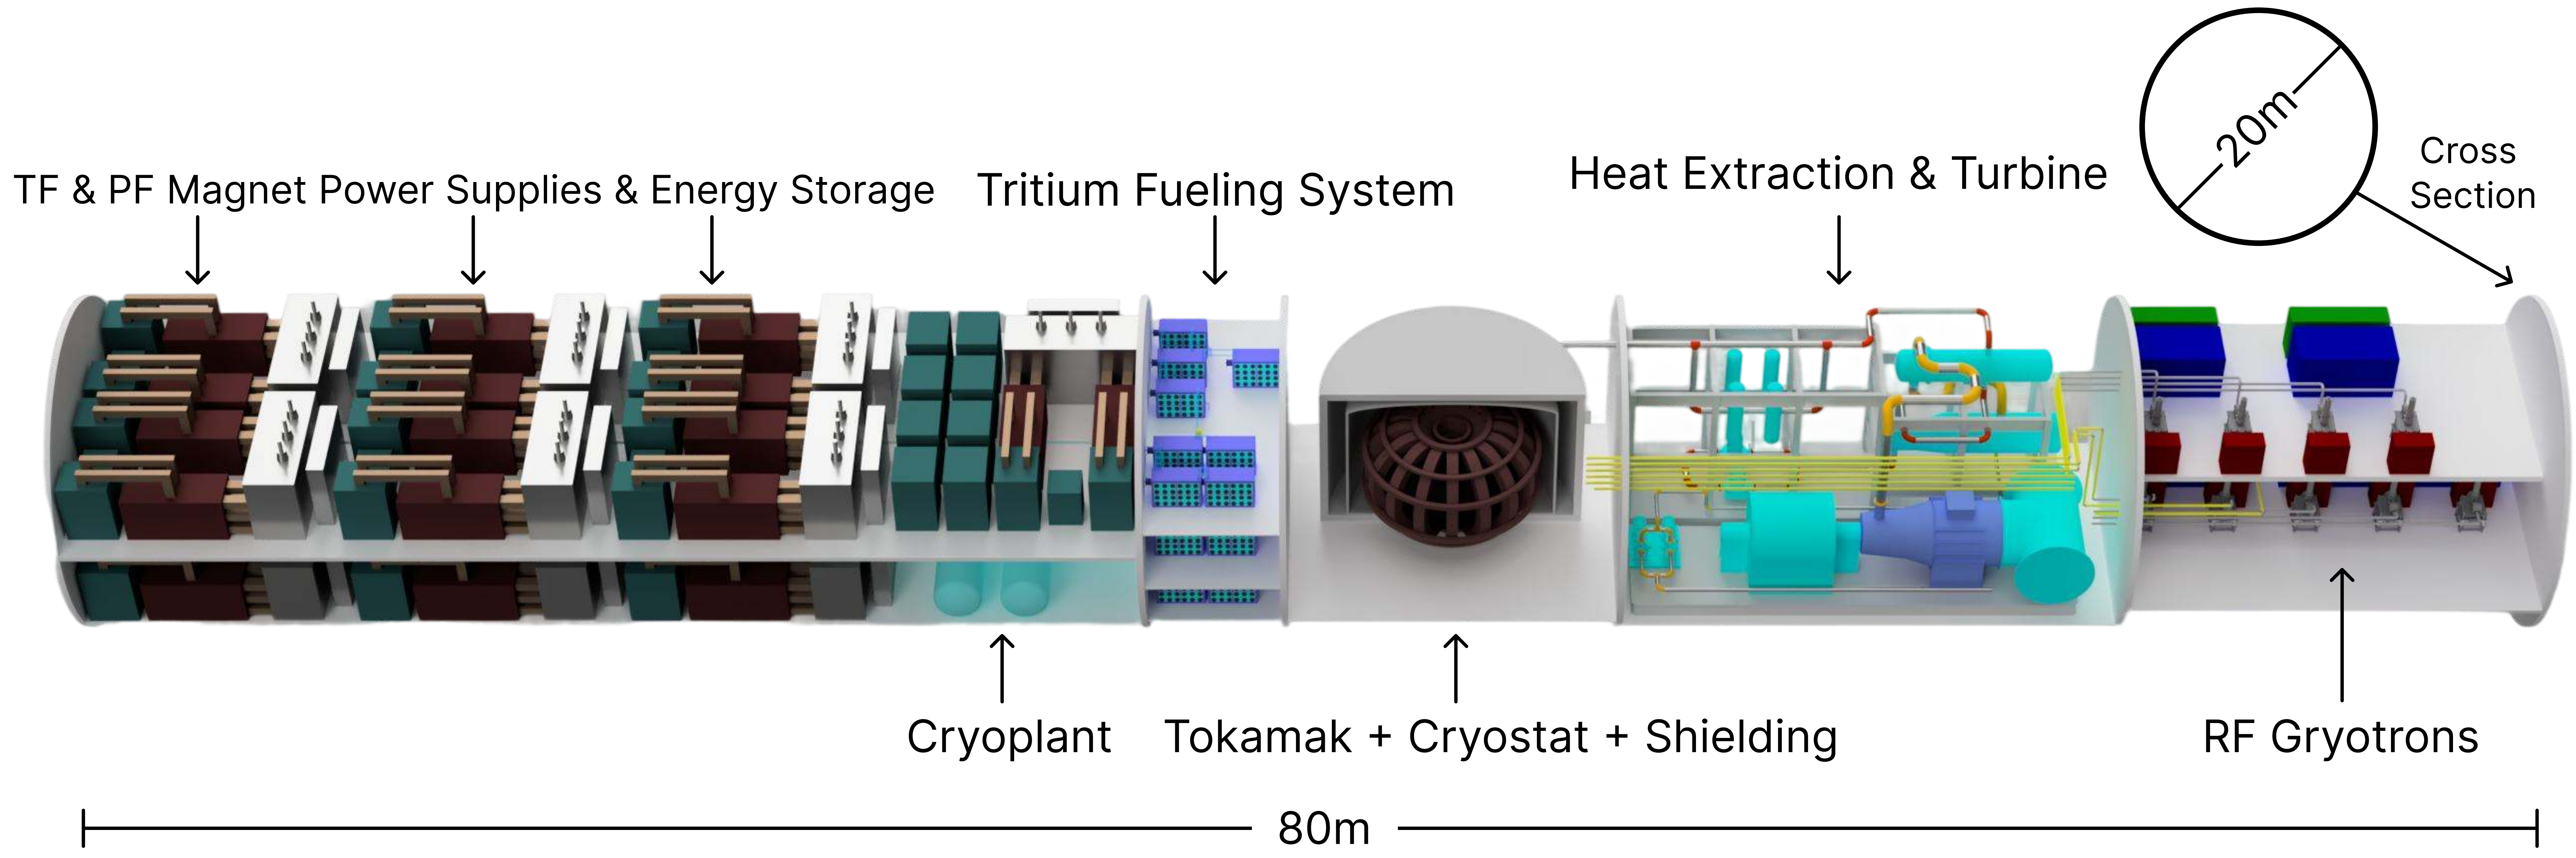

Front View (45° Up)

Based on all this, we estimate the total volume for the entire system to be: **30,000m³.** For reference, the fuel tanks inside of a 20,000TEU ammonia tanker can exceed ~100,000m³ and a similar green methanol tanker ~80,000m³ + ~20,000m³ for the engine room itself. Thus, the reactor and supporting auxiliary system will quite easily fit inside the ship with room to spare. It is slightly larger than the engine room of these tankers, but with the removal of the traditional fuel tanks, a fusion propulsion system can achieve a size reduction of **~3-4x**.

Back View

Reference: Maersk Triple-E Class Ship

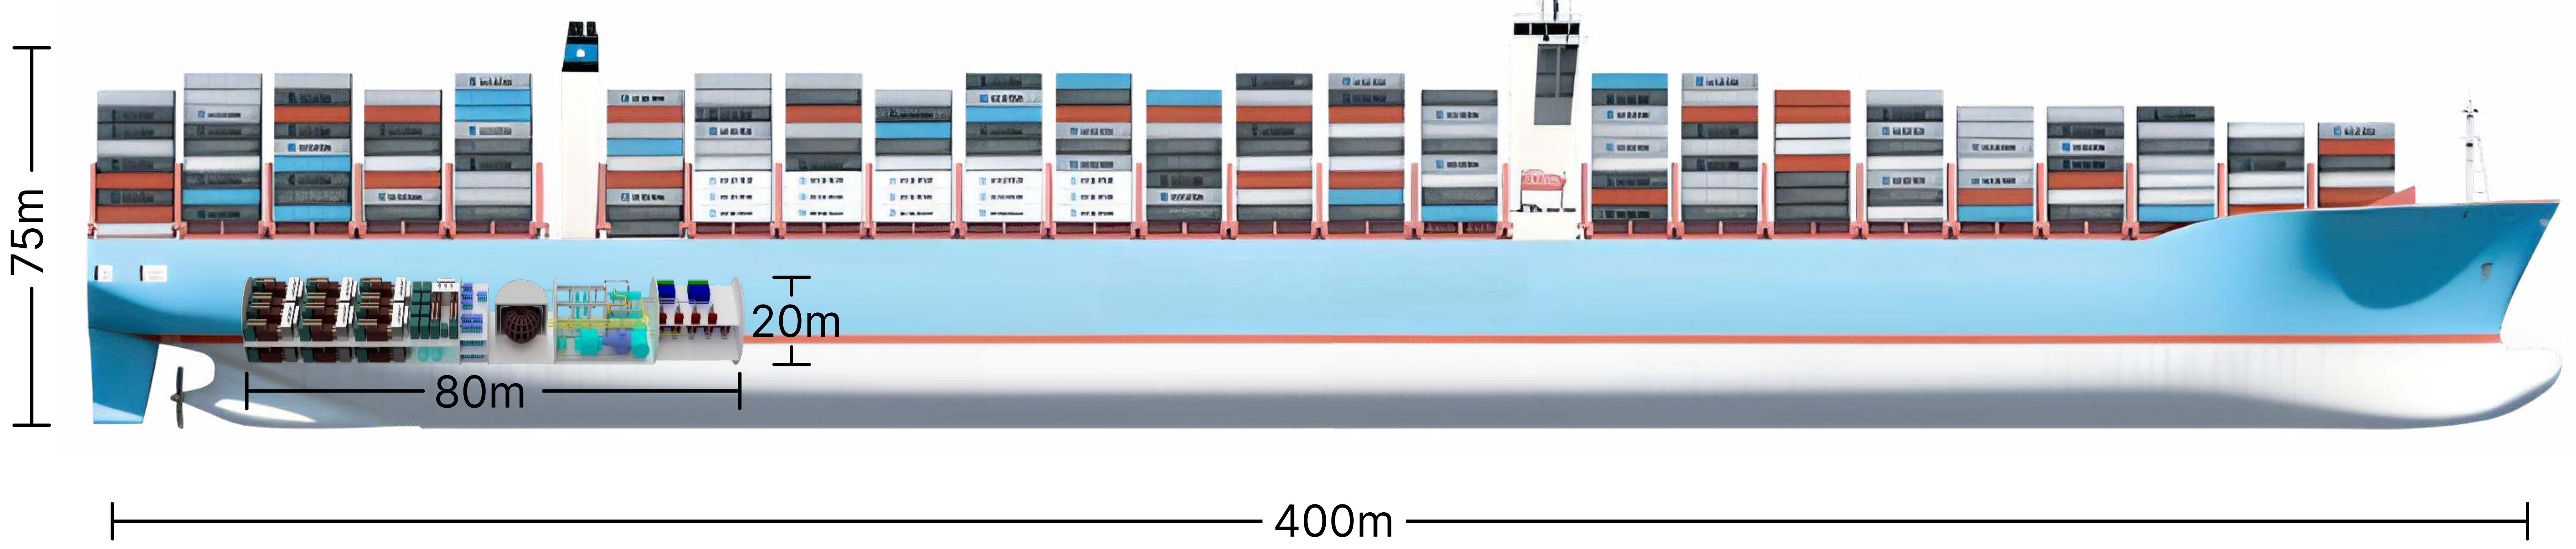

First Wall and Divertor

First Wall Lifetime Estimation

The peak heat flux regions of the reactor are protected by a 1 cm-thick pure tungsten first wall, designed to withstand erosion and thermal stress. First-wall lifetime is primarily determined by power density and particle flux, which drive material degradation. With a lower surface power density (< 0.88MW/m^2), Yinsen's first wall will experience significantly less erosion than grid-scale reactors operating at higher power densities, extending lifespan and reducing maintenance demands.

Based on a tungsten sputtering yield of $Y \approx 0.0048$ at 300 eV [1], the estimated erosion rate for Yinsen's first wall is 21 g/m^2·yr, compared to 7900 g/m^2·yr in grid-scale reactors (>2.5 MW/m^2) [1]. Assuming sputtering is primarily driven by charge-exchange neutrals, this results in a minimal material loss of 5.26 kg/yr·m^2, giving the first wall an erosion-limited, ideal lifespan of ~30 years. However, neutron-induced damage—not sputtering erosion—primarily dictates replacement intervals. The neutronics analysis suggests negligible displacement per atom, 1DPA/yr, however experimental data shows tungsten withstands only up to 10 DPA/year [2]. Based on this data and Yinsen's expected neutron flux, we estimate a 5–10 year replacement schedule, potentially extending to 10–15 years due to the reactor's low capacity factor operation, which reduces integrated neutron exposure.

Divertor Design

The divertor, located at the top and bottom of the tokamak, intercepts heat and particles from the scrape-off layer (SOL) to prevent excessive plasma-wall interactions and provide an energy and particle exhaust pathway. While divertor systems in high-power-density reactors face extreme thermal loads and maintenance challenges, Yinsen's lower-power-density plasma significantly reduces divertor stress, simplifying both its geometry and maintenance requirements.

A double-null tungsten divertor was chosen for its efficient heat and particle exhaust capability, reduced localized thermal stress, and long-term durability - leveraging tungsten's high melting point (3,422°C), low sputtering rate, and neutron resilience under fusion conditions. UEDGE simulations were used to optimize neon (Ne) impurity seeding to maintain safe operating temperatures.By leveraging active liquid cooling, we prevent localized melting and minimize material erosion.

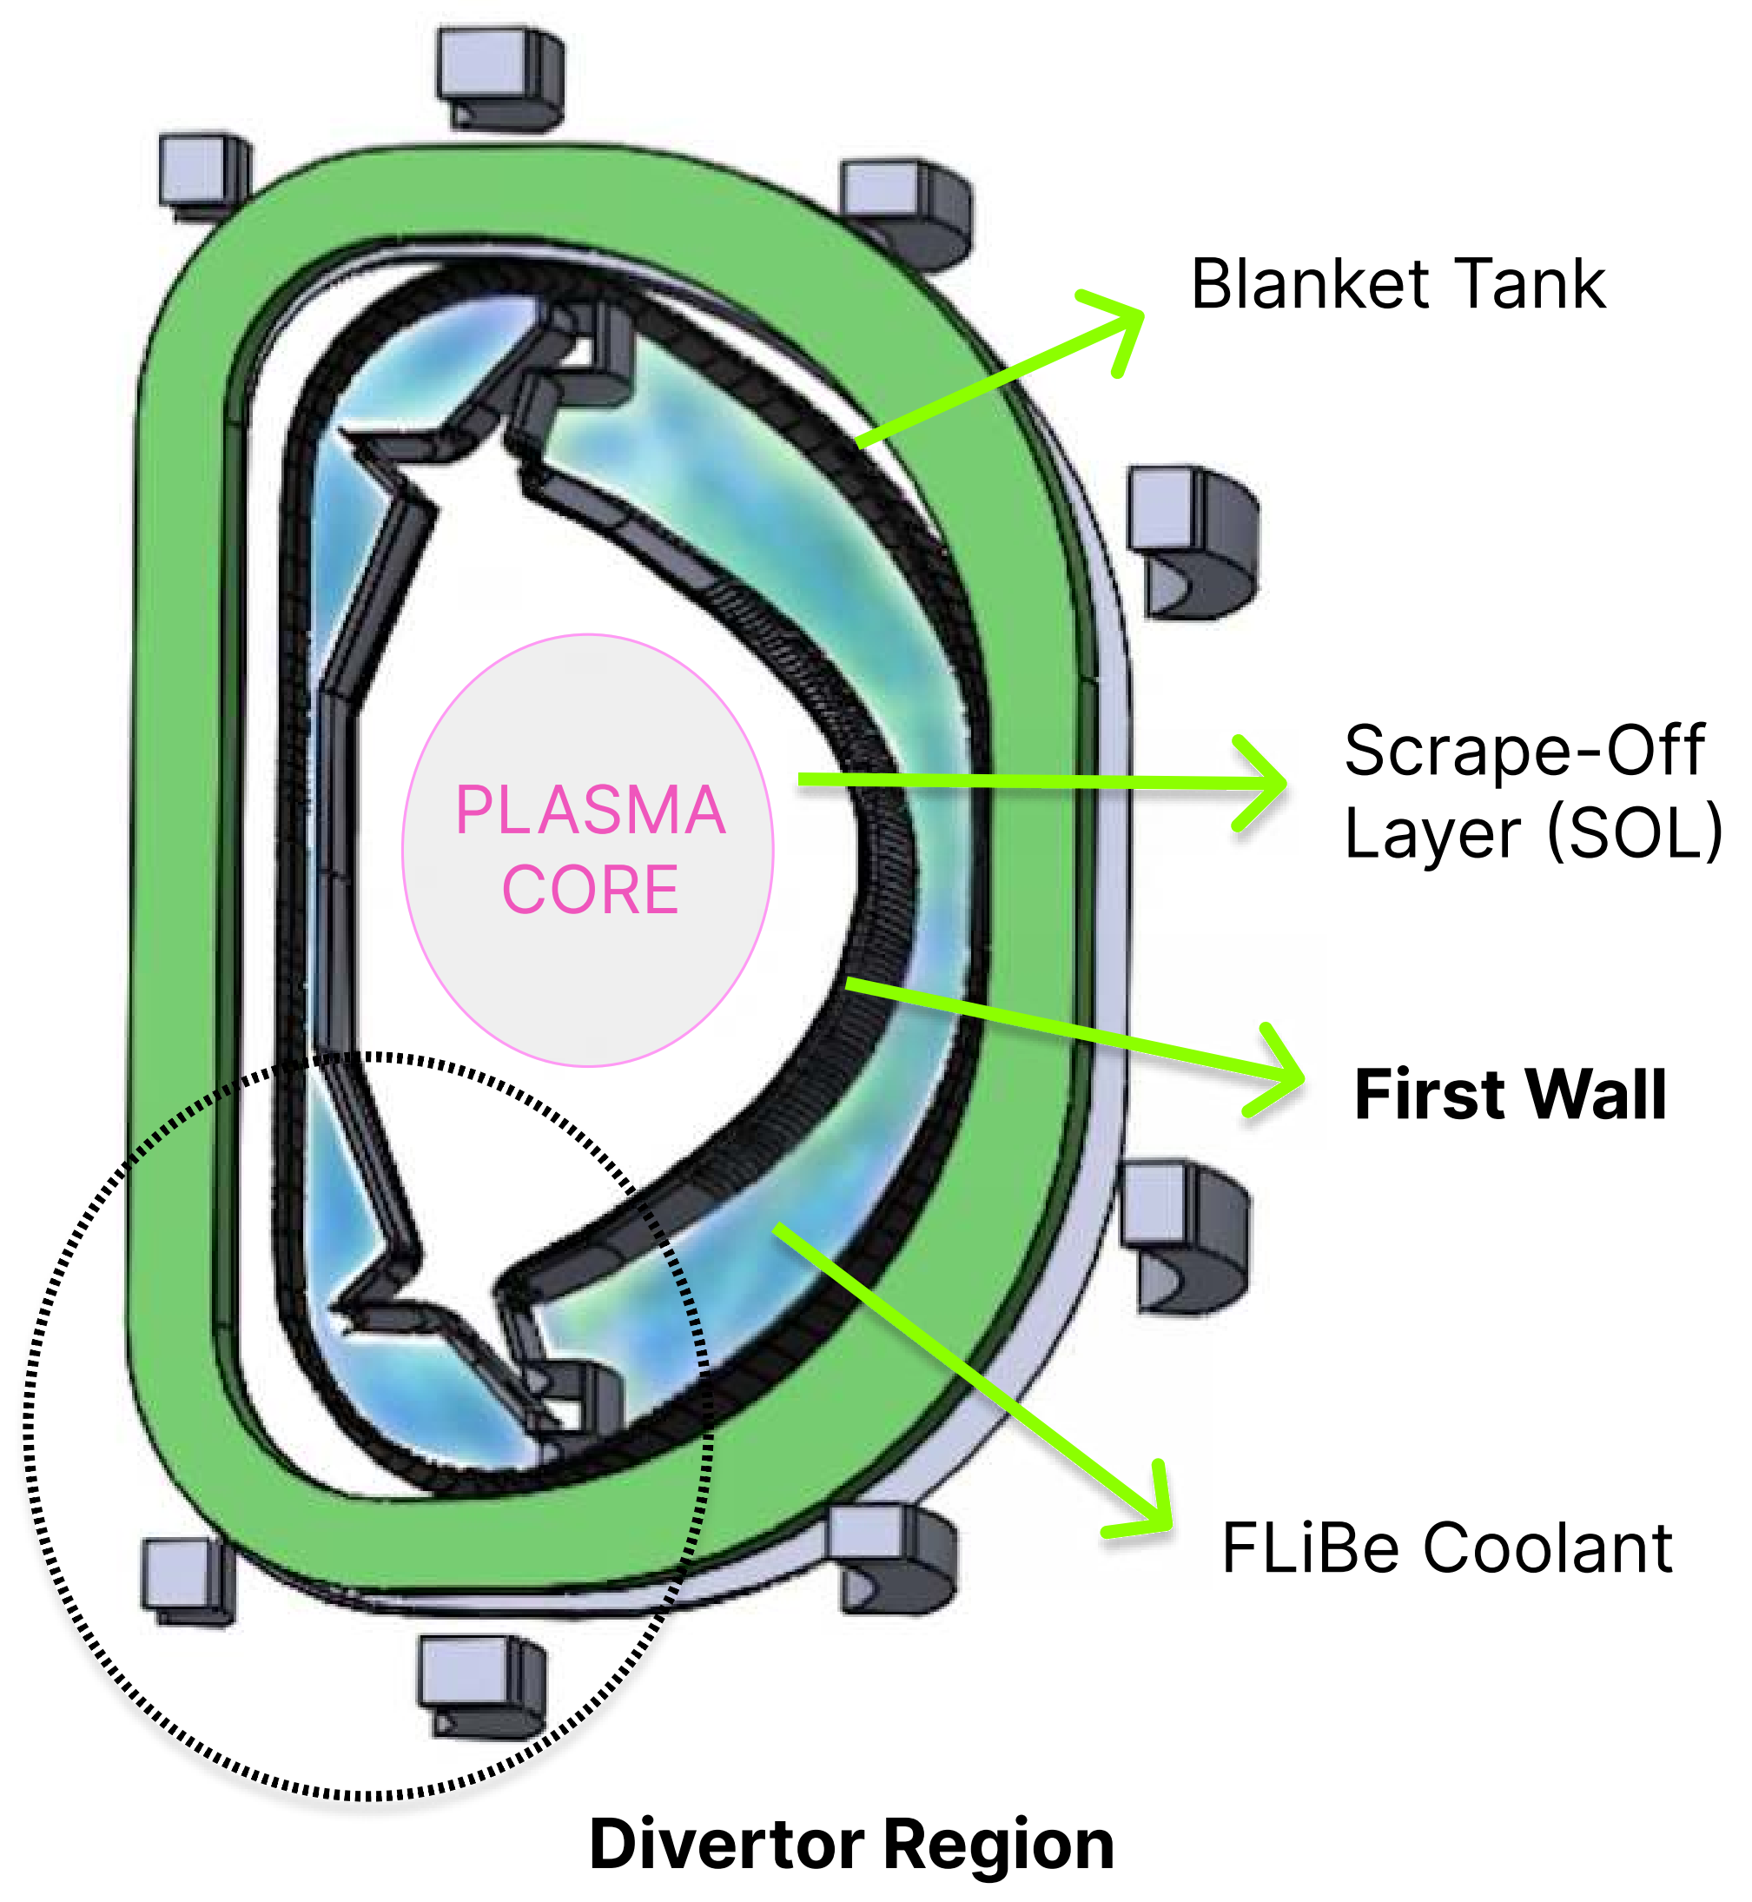

UEDGE Mesh Plots

UEDGE, a 2D fluid transport code developed by Lawrence Livermore National Laboratory (LLNL), was used to model the scrape-off layer (SOL) and divertor plasma, estimating heat fluxes to divertor targets based on our operating parameters. For P_{SOL} = 20 MW and n_separatrix = 2.5 × 10^{20} m^{-3}, the results indicate that only minimal power handling is required for an actively cooled divertor. With a fixed neon impurity seeding fraction of 4%, the peak heat flux at the outer divertor target remains within acceptable limits (max heat flux: 10 MW/m^2), and the tungsten wall temperature stays below the recrystallization threshold of 1,550°C.

Electron Temperature (eV)

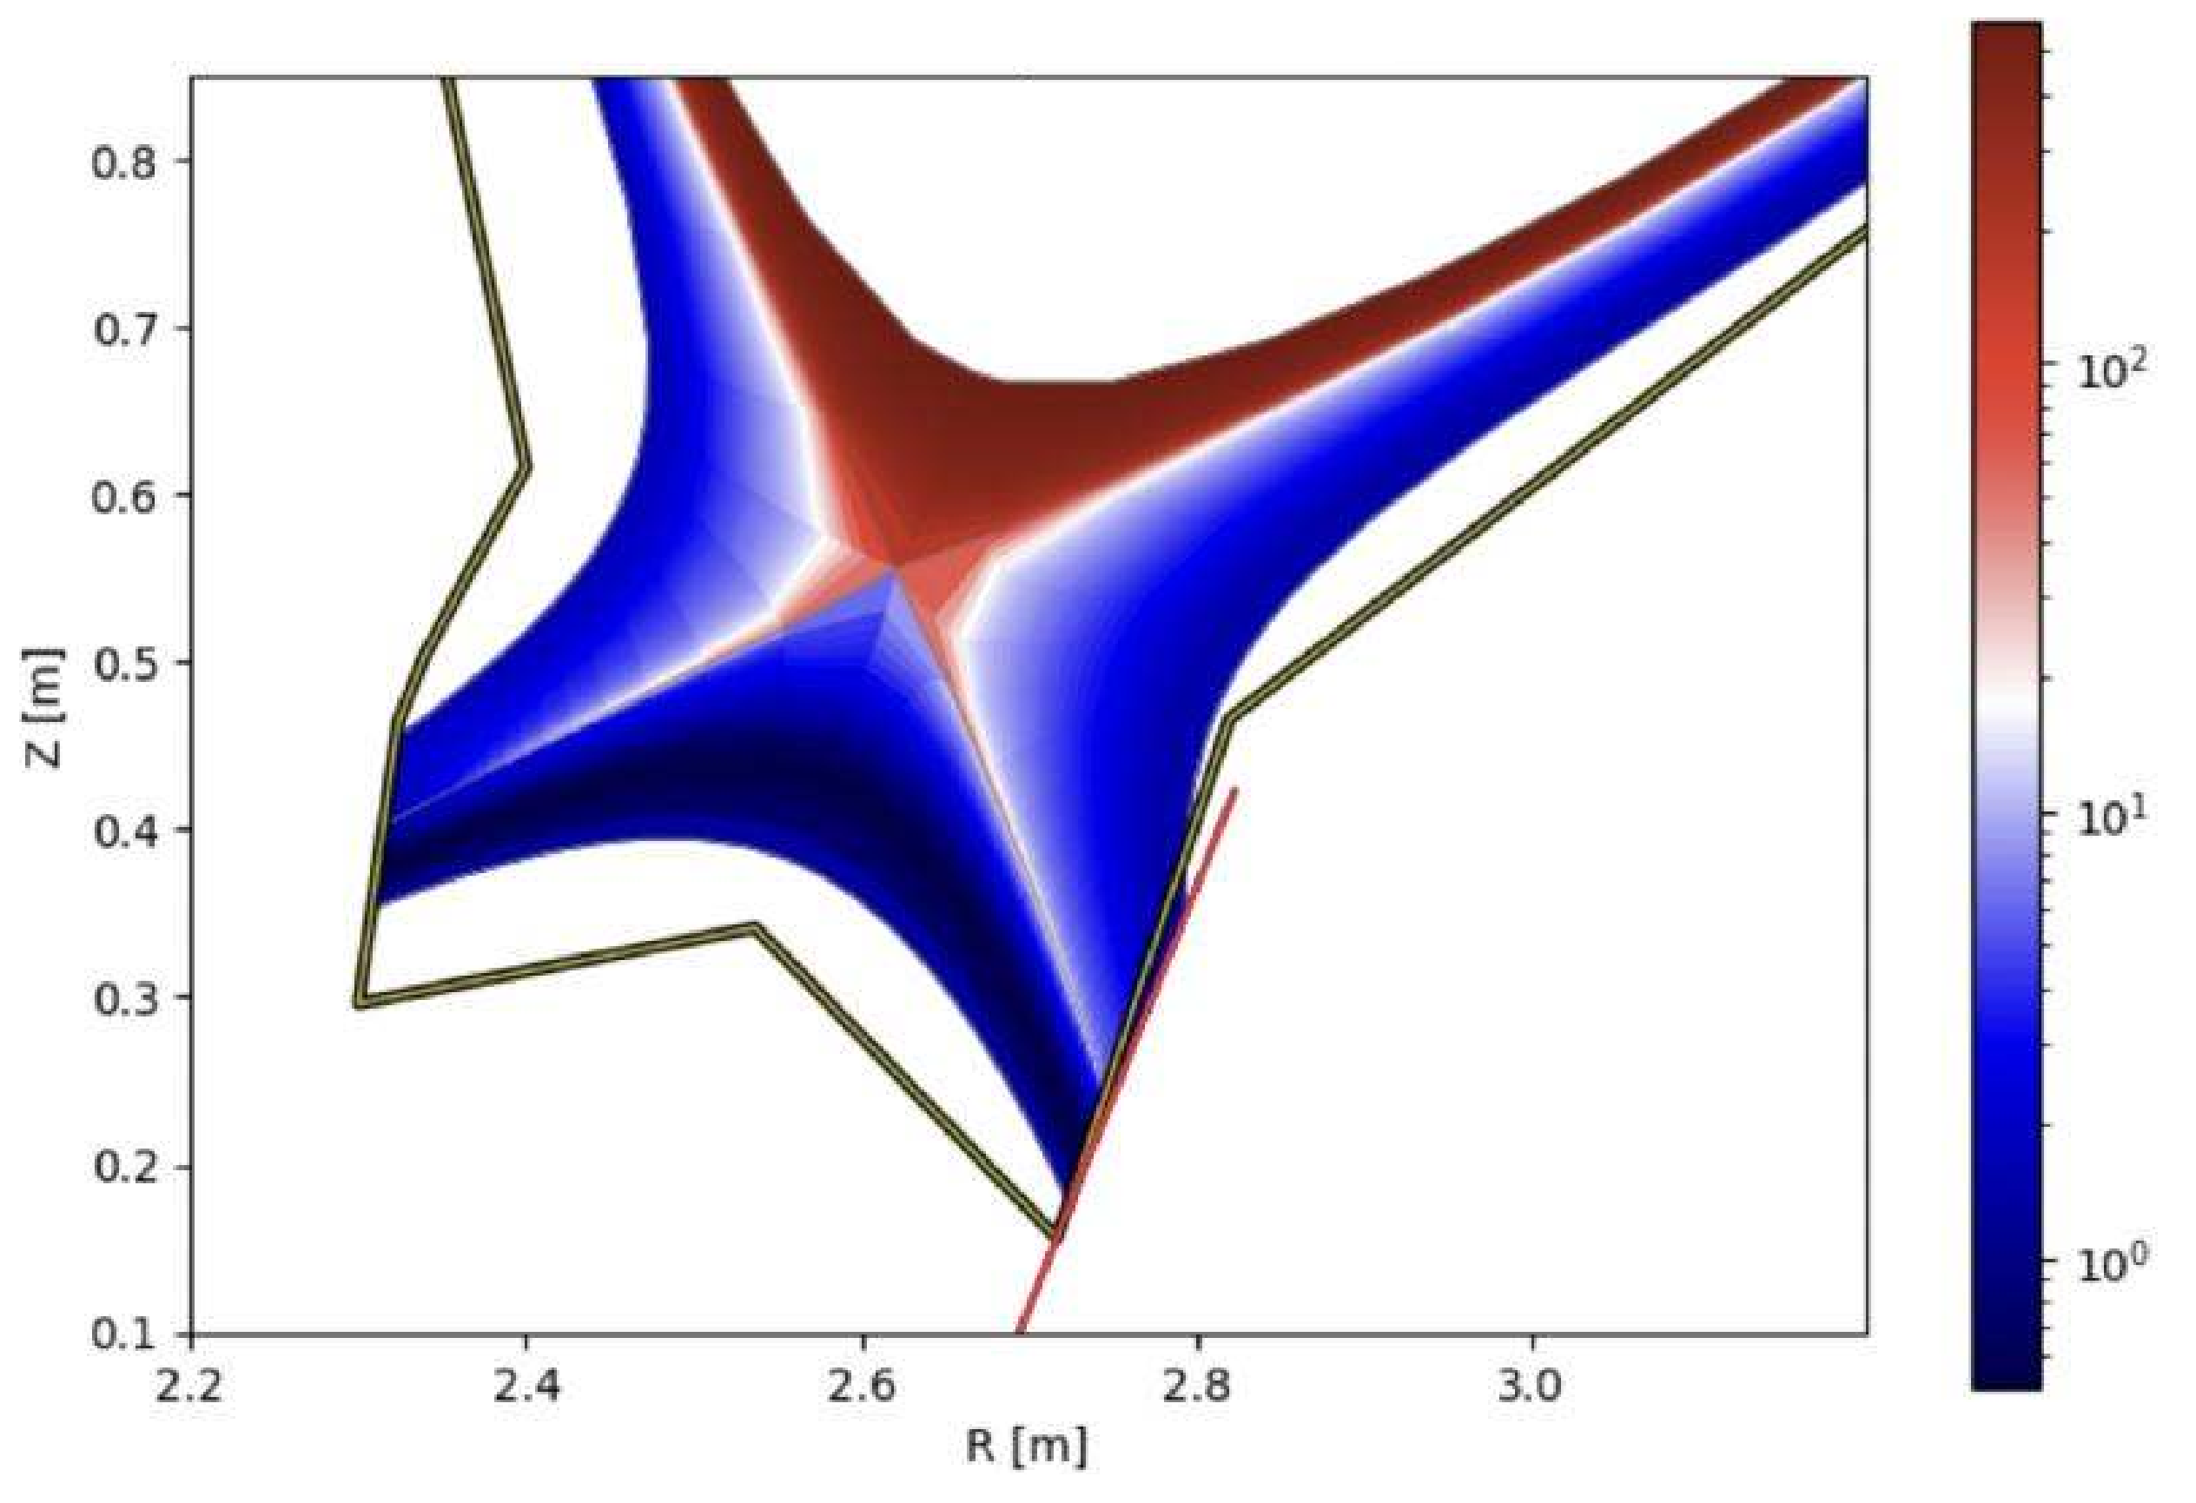

Electron Heat (W)

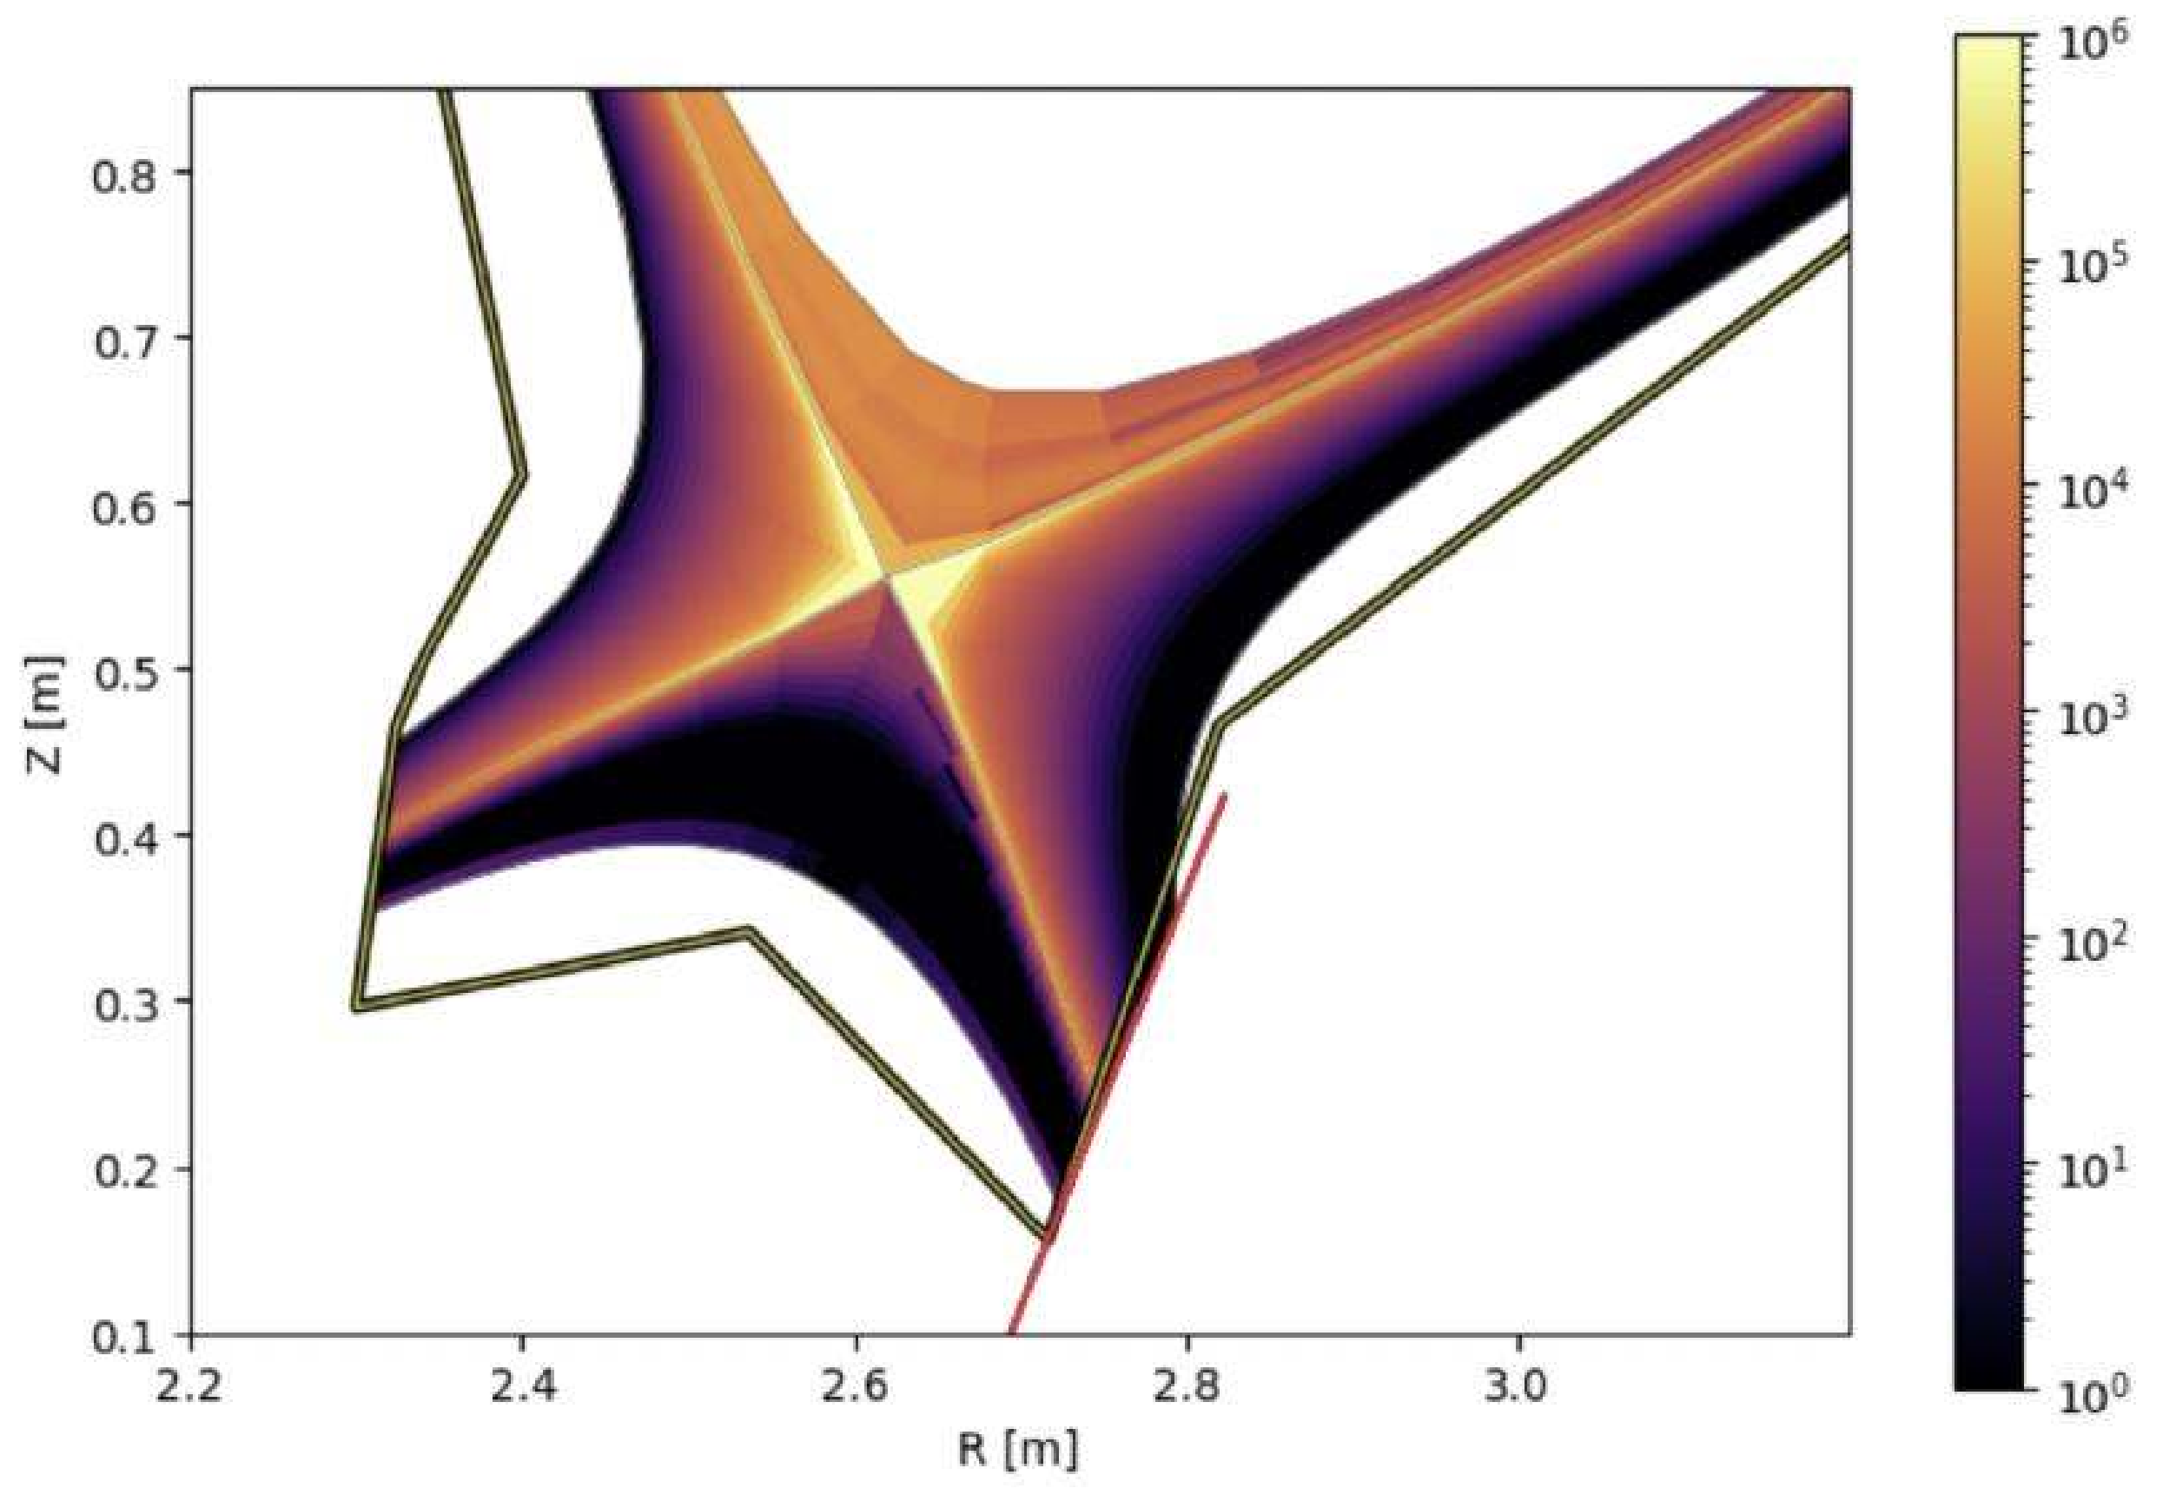

Outer Plate Heat Fluxes

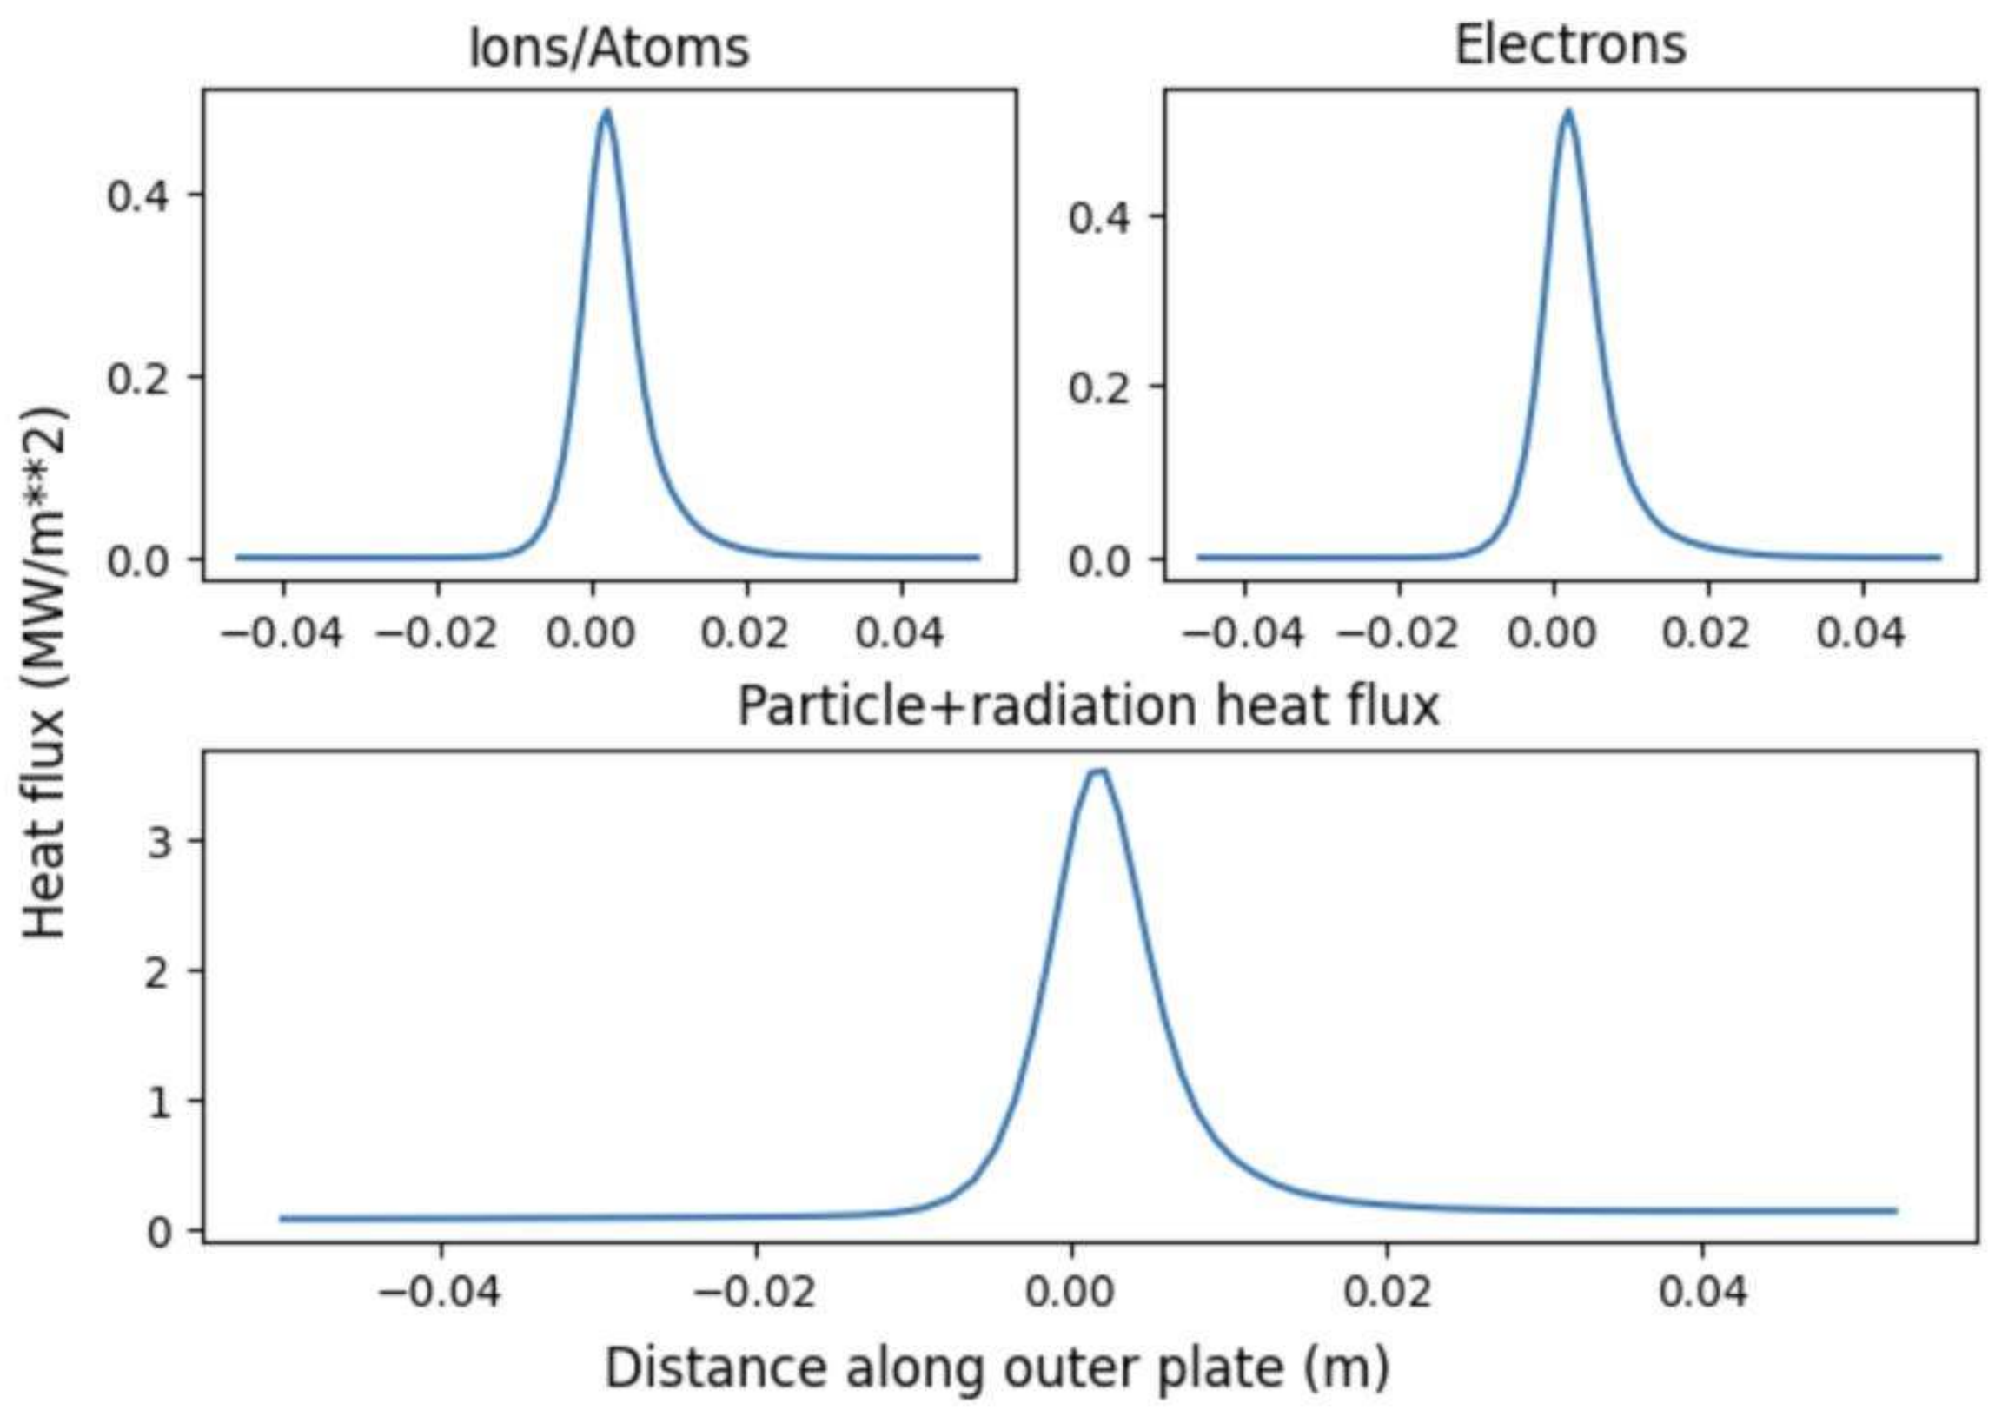

In addition to neon impurity seeding, the first wall will be actively cooled by the blanket behind it, and the divertor will employ strike-point sweeping to distribute heat loads and minimize localized erosion. With these mitigation strategies in place, first wall and divertor erosion are solvable challenges, ensuring extended component lifetimes and reduced maintenance needs.

[1] P.C. Stangeby. Assessing material migration through 13c injection experiments. Journal of Nuclear Materials, 415(1):S1–S8, 2010.
[2] Steven J. Zinkle and Lance L. Snead. Designing radiation resistance in materials for fusion energy. Annual Review of Materials Research, 44:241–267, 2014.

Tritium Fueling

For a 500MW$_{th}$ fusion system, the estimated baseline tritium inventory is approximately 1.6kg, accounting for both the initial startup requirement and the steady-state inventory within the fuel processing systems. The tritium inventory within the gas-phase processing loop is expected to scale proportionally with the reactor's thermal power output. Accordingly, for a 115MW$_{th}$ system, a proportional estimate suggests a tritium inventory of approximately 350g in the gas-phase loop.

The majority of this tritium inventory stays contained within the gas-phase processing loop, however some exists within the blanket itself. For a 500MW$_{th}$ system utilizing a FLiBe blanket, the expected tritium inventory within the blanket is approximately 32g. In contrast, for a 115MW$_{th}$ system, the blanket inventory is anticipated to be significantly lower, as it scales with the blanket's volumetric capacity rather than directly with thermal power. Specifically, for the Yinsen system, a reasonable estimate for the FLiBe blanket **tritium inventory falls within the range of 20–30g.**

Advanced Tokamak Fuel Cycle

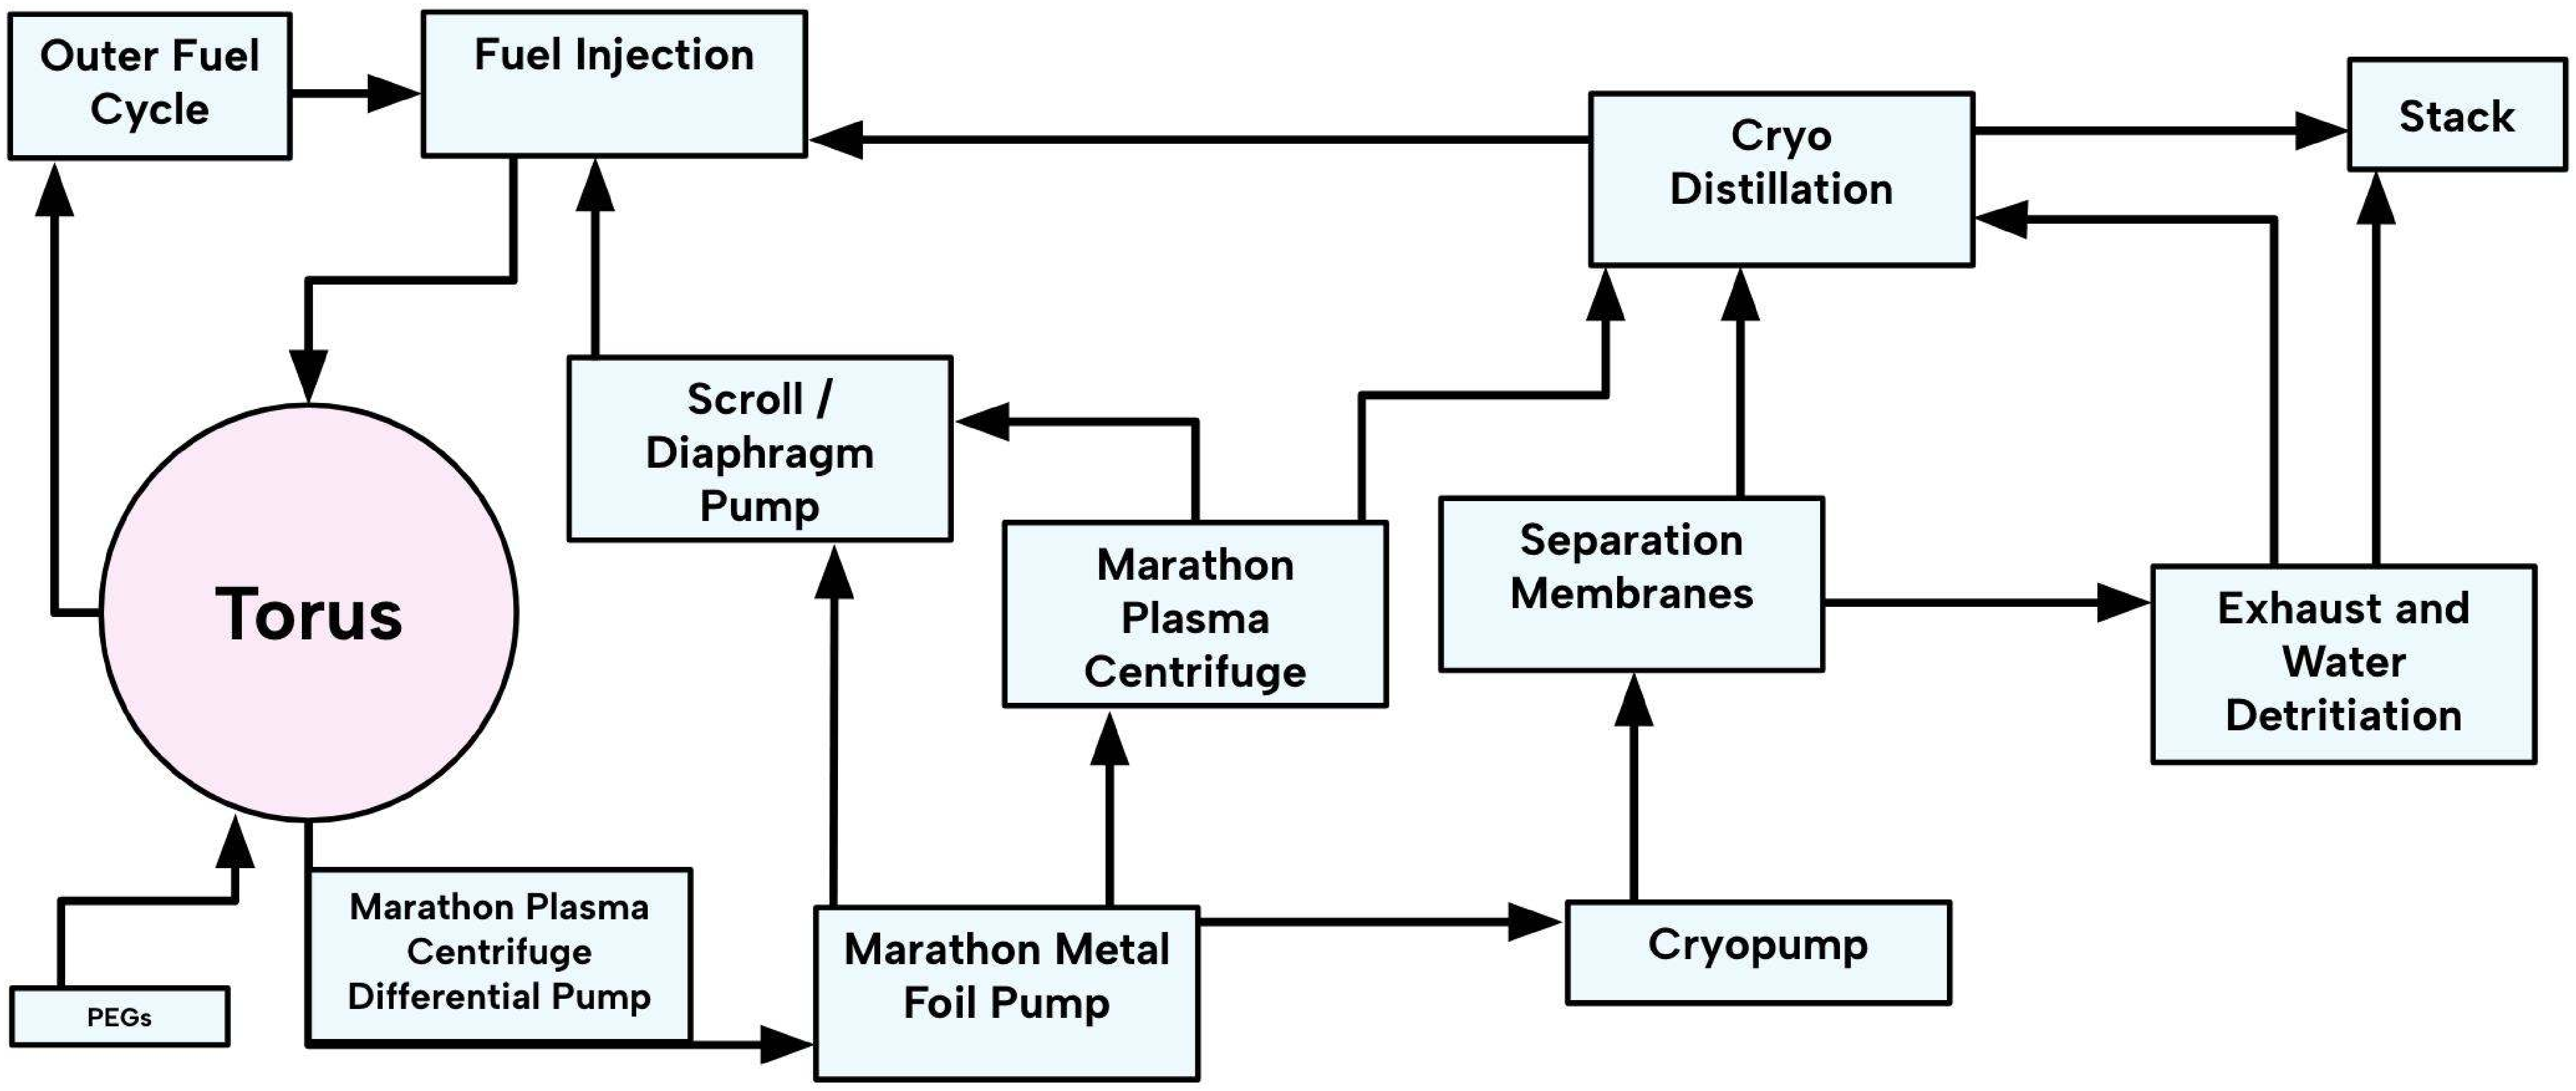

Source: Marathon Fusion

The implementation of a plasma centrifuge enables a significant increase in tritium burn efficiency—by an order of magnitude or more—resulting in a corresponding reduction in the mass flow rate through downstream gas processing systems. This reduction decreases the required tritium inventory from 318g to approximately 31.8g, bringing it in line with the expected inventory within the FLiBe blanket.

Further optimization is achievable through the integration of a superpermeable pump, which enhances the efficiency of gas-phase tritium recovery. This technology is expected to improve processing speed by an additional factor of ~2, further reducing the gas-phase tritium inventory to ~15g, leading to a total system inventory of approximately 47g.

A key advantage of this approach, particularly for Yinsen, is the substantial reduction in system size. By leveraging both the plasma centrifuge and superpermeable pump, the downstream processing systems beyond the superpermeable pump can be scaled down to only a few percent of the size required for conventional grid-scale fusion reactors, based on tritium flow rate. While these fuel cycle technologies have yet to reach full production readiness, significant R&D efforts are underway within the fusion supply chain to bring them to market within a timeframe aligned with the first Yinsen deployment.

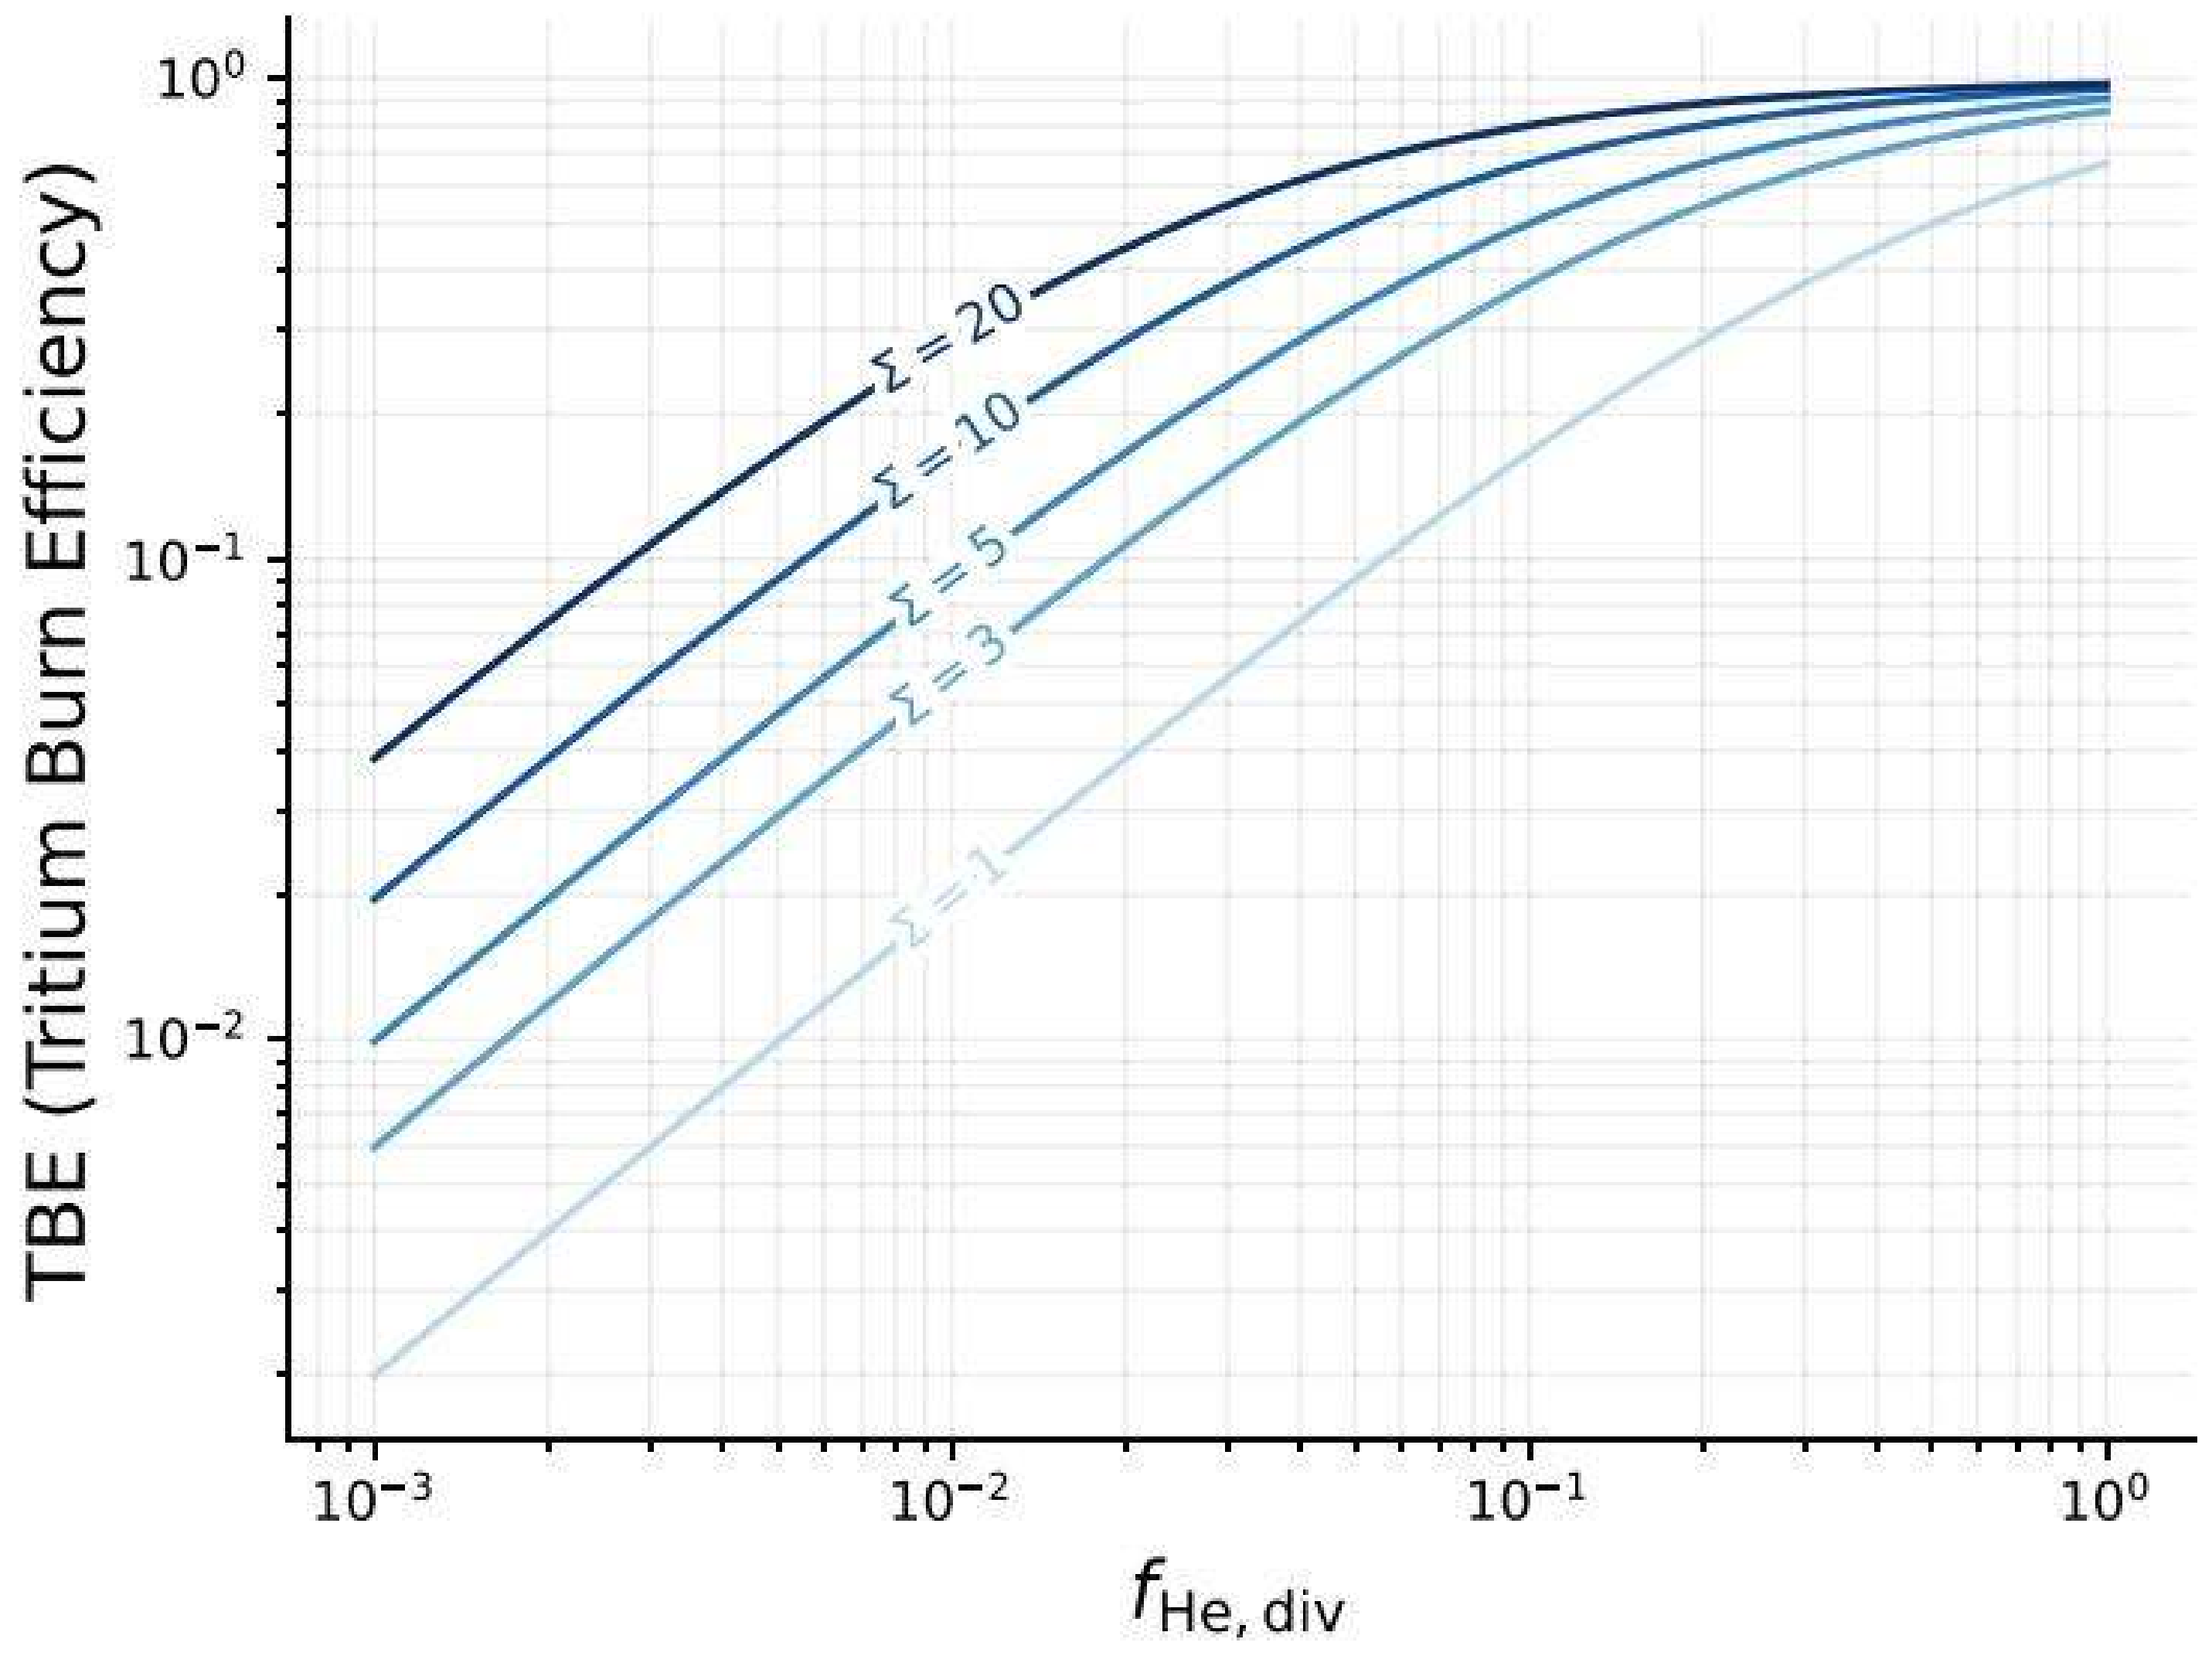

Source: D.G. Whyte et al 2023 Nucl. Fusion 63 126019

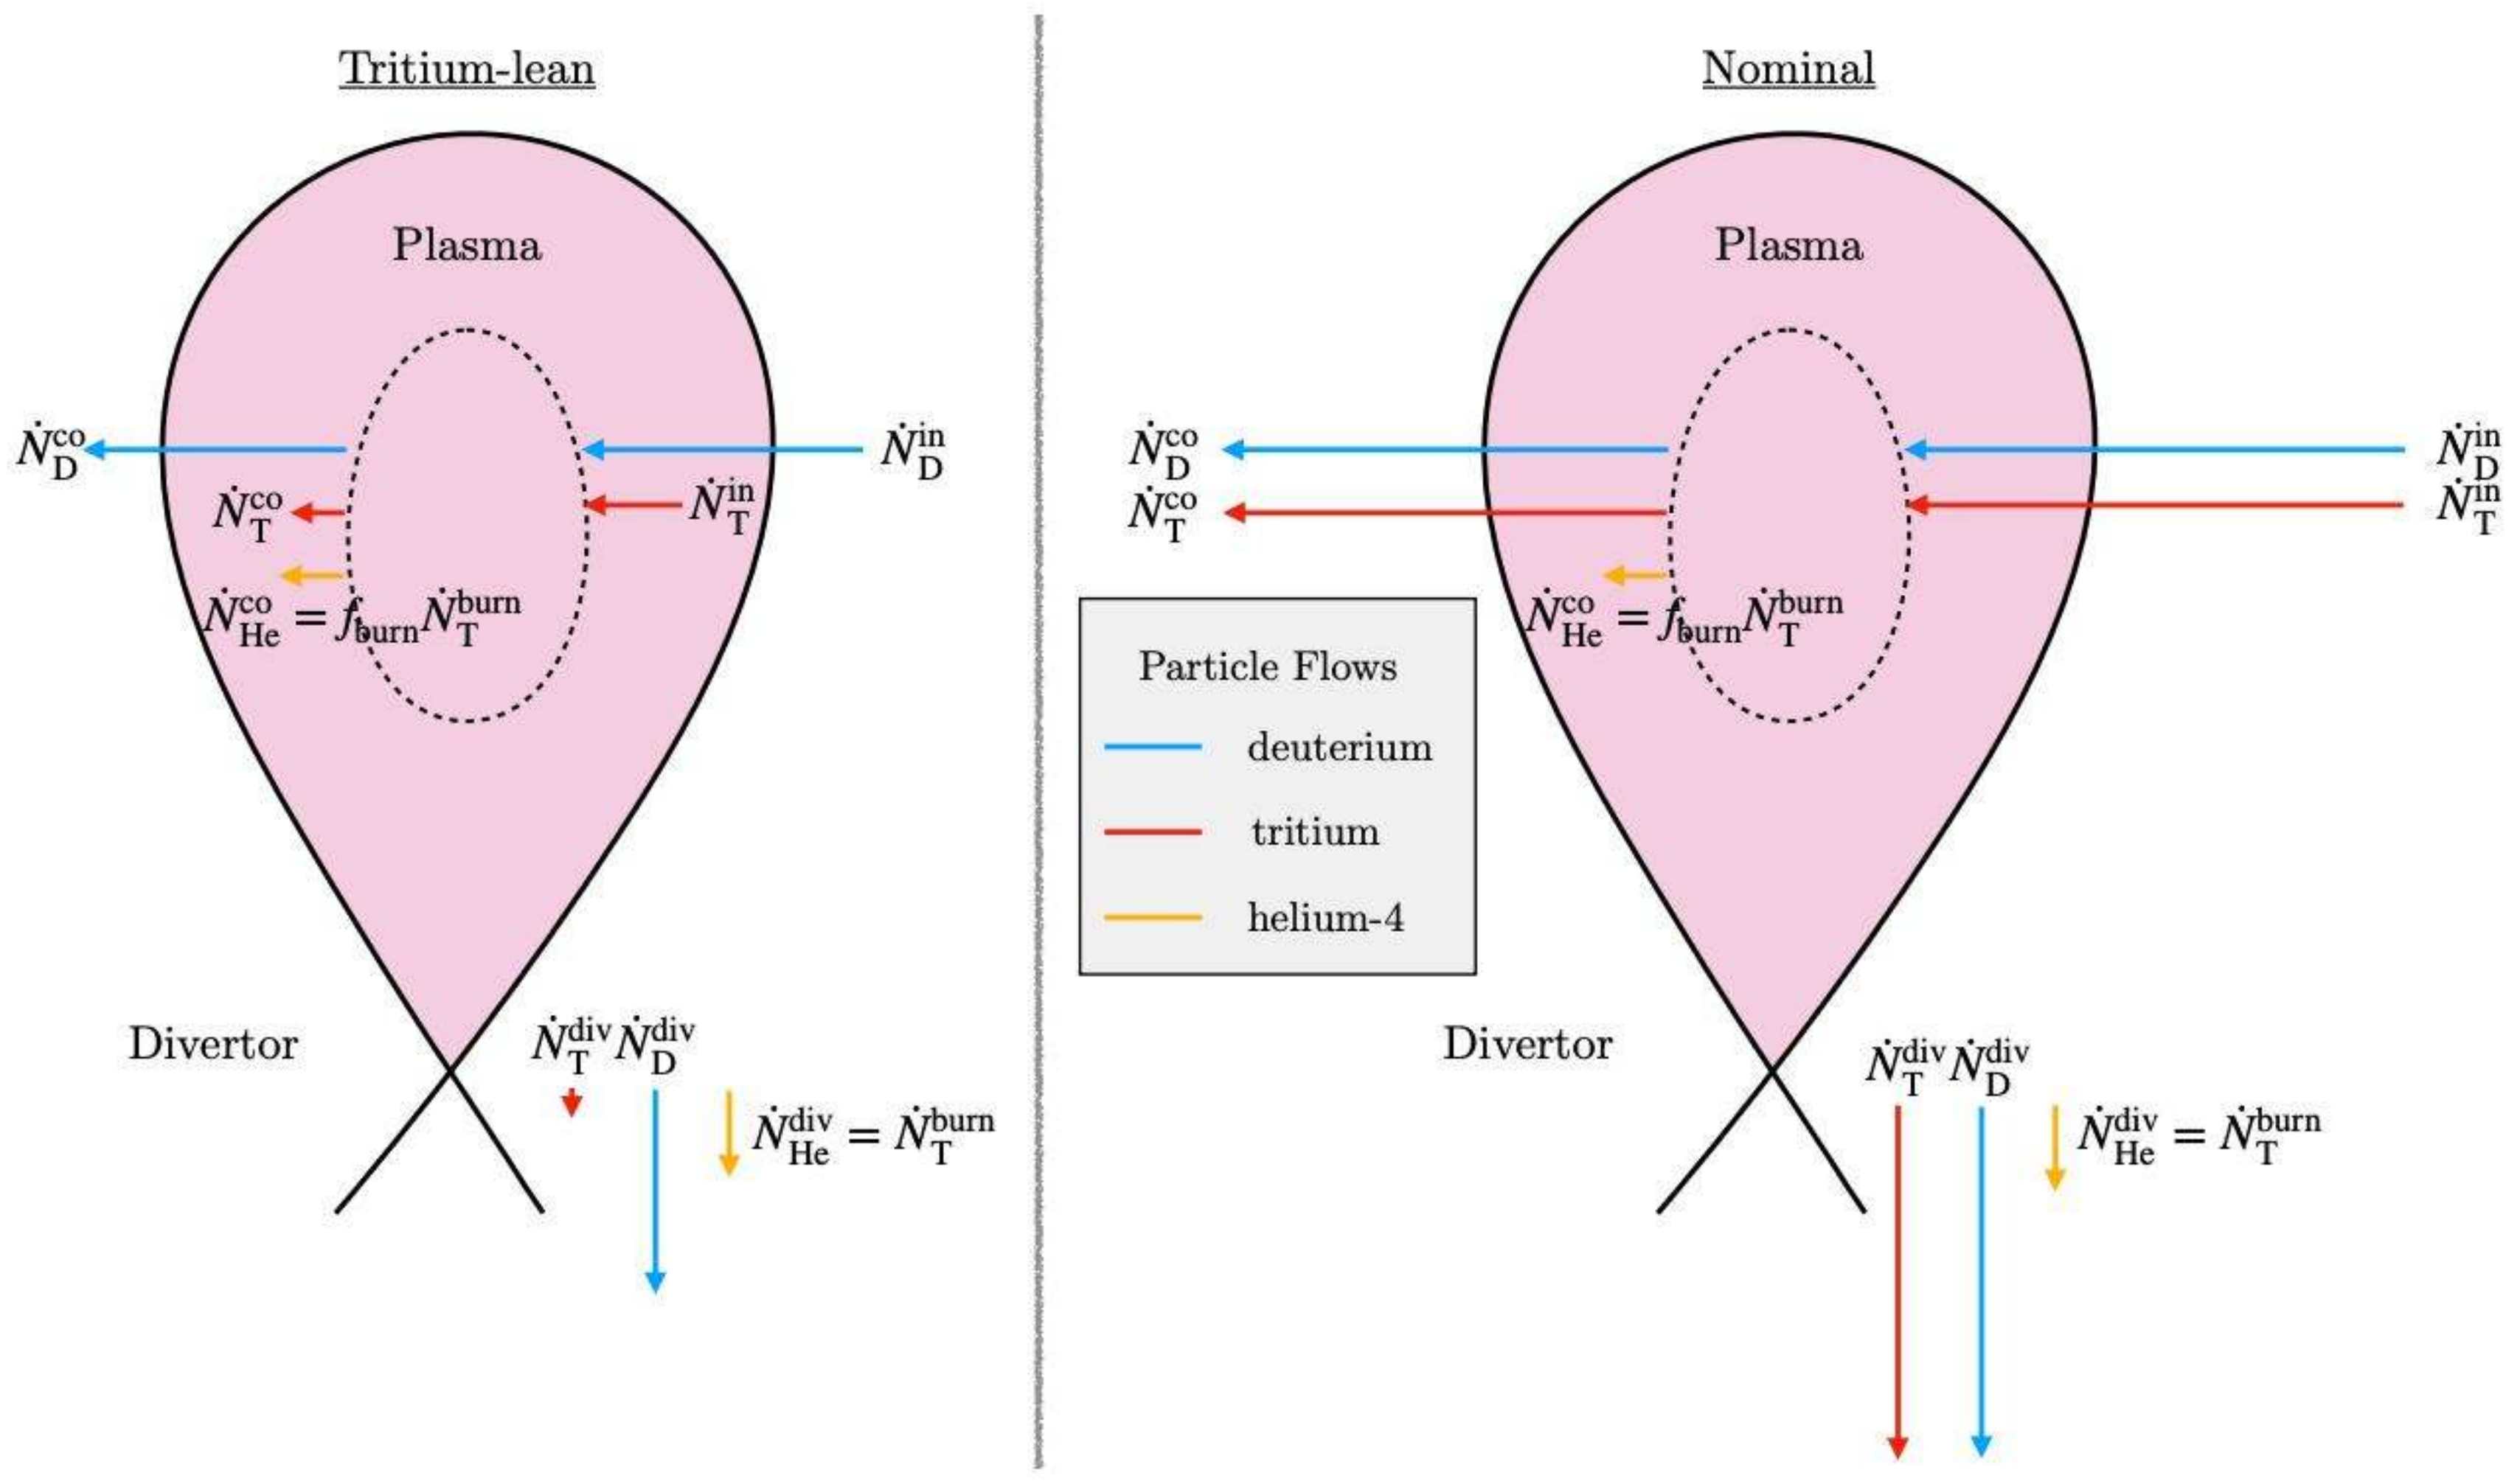

Source: J.F. Parisi et al. 2024

Chapter 5

Go-to-Market

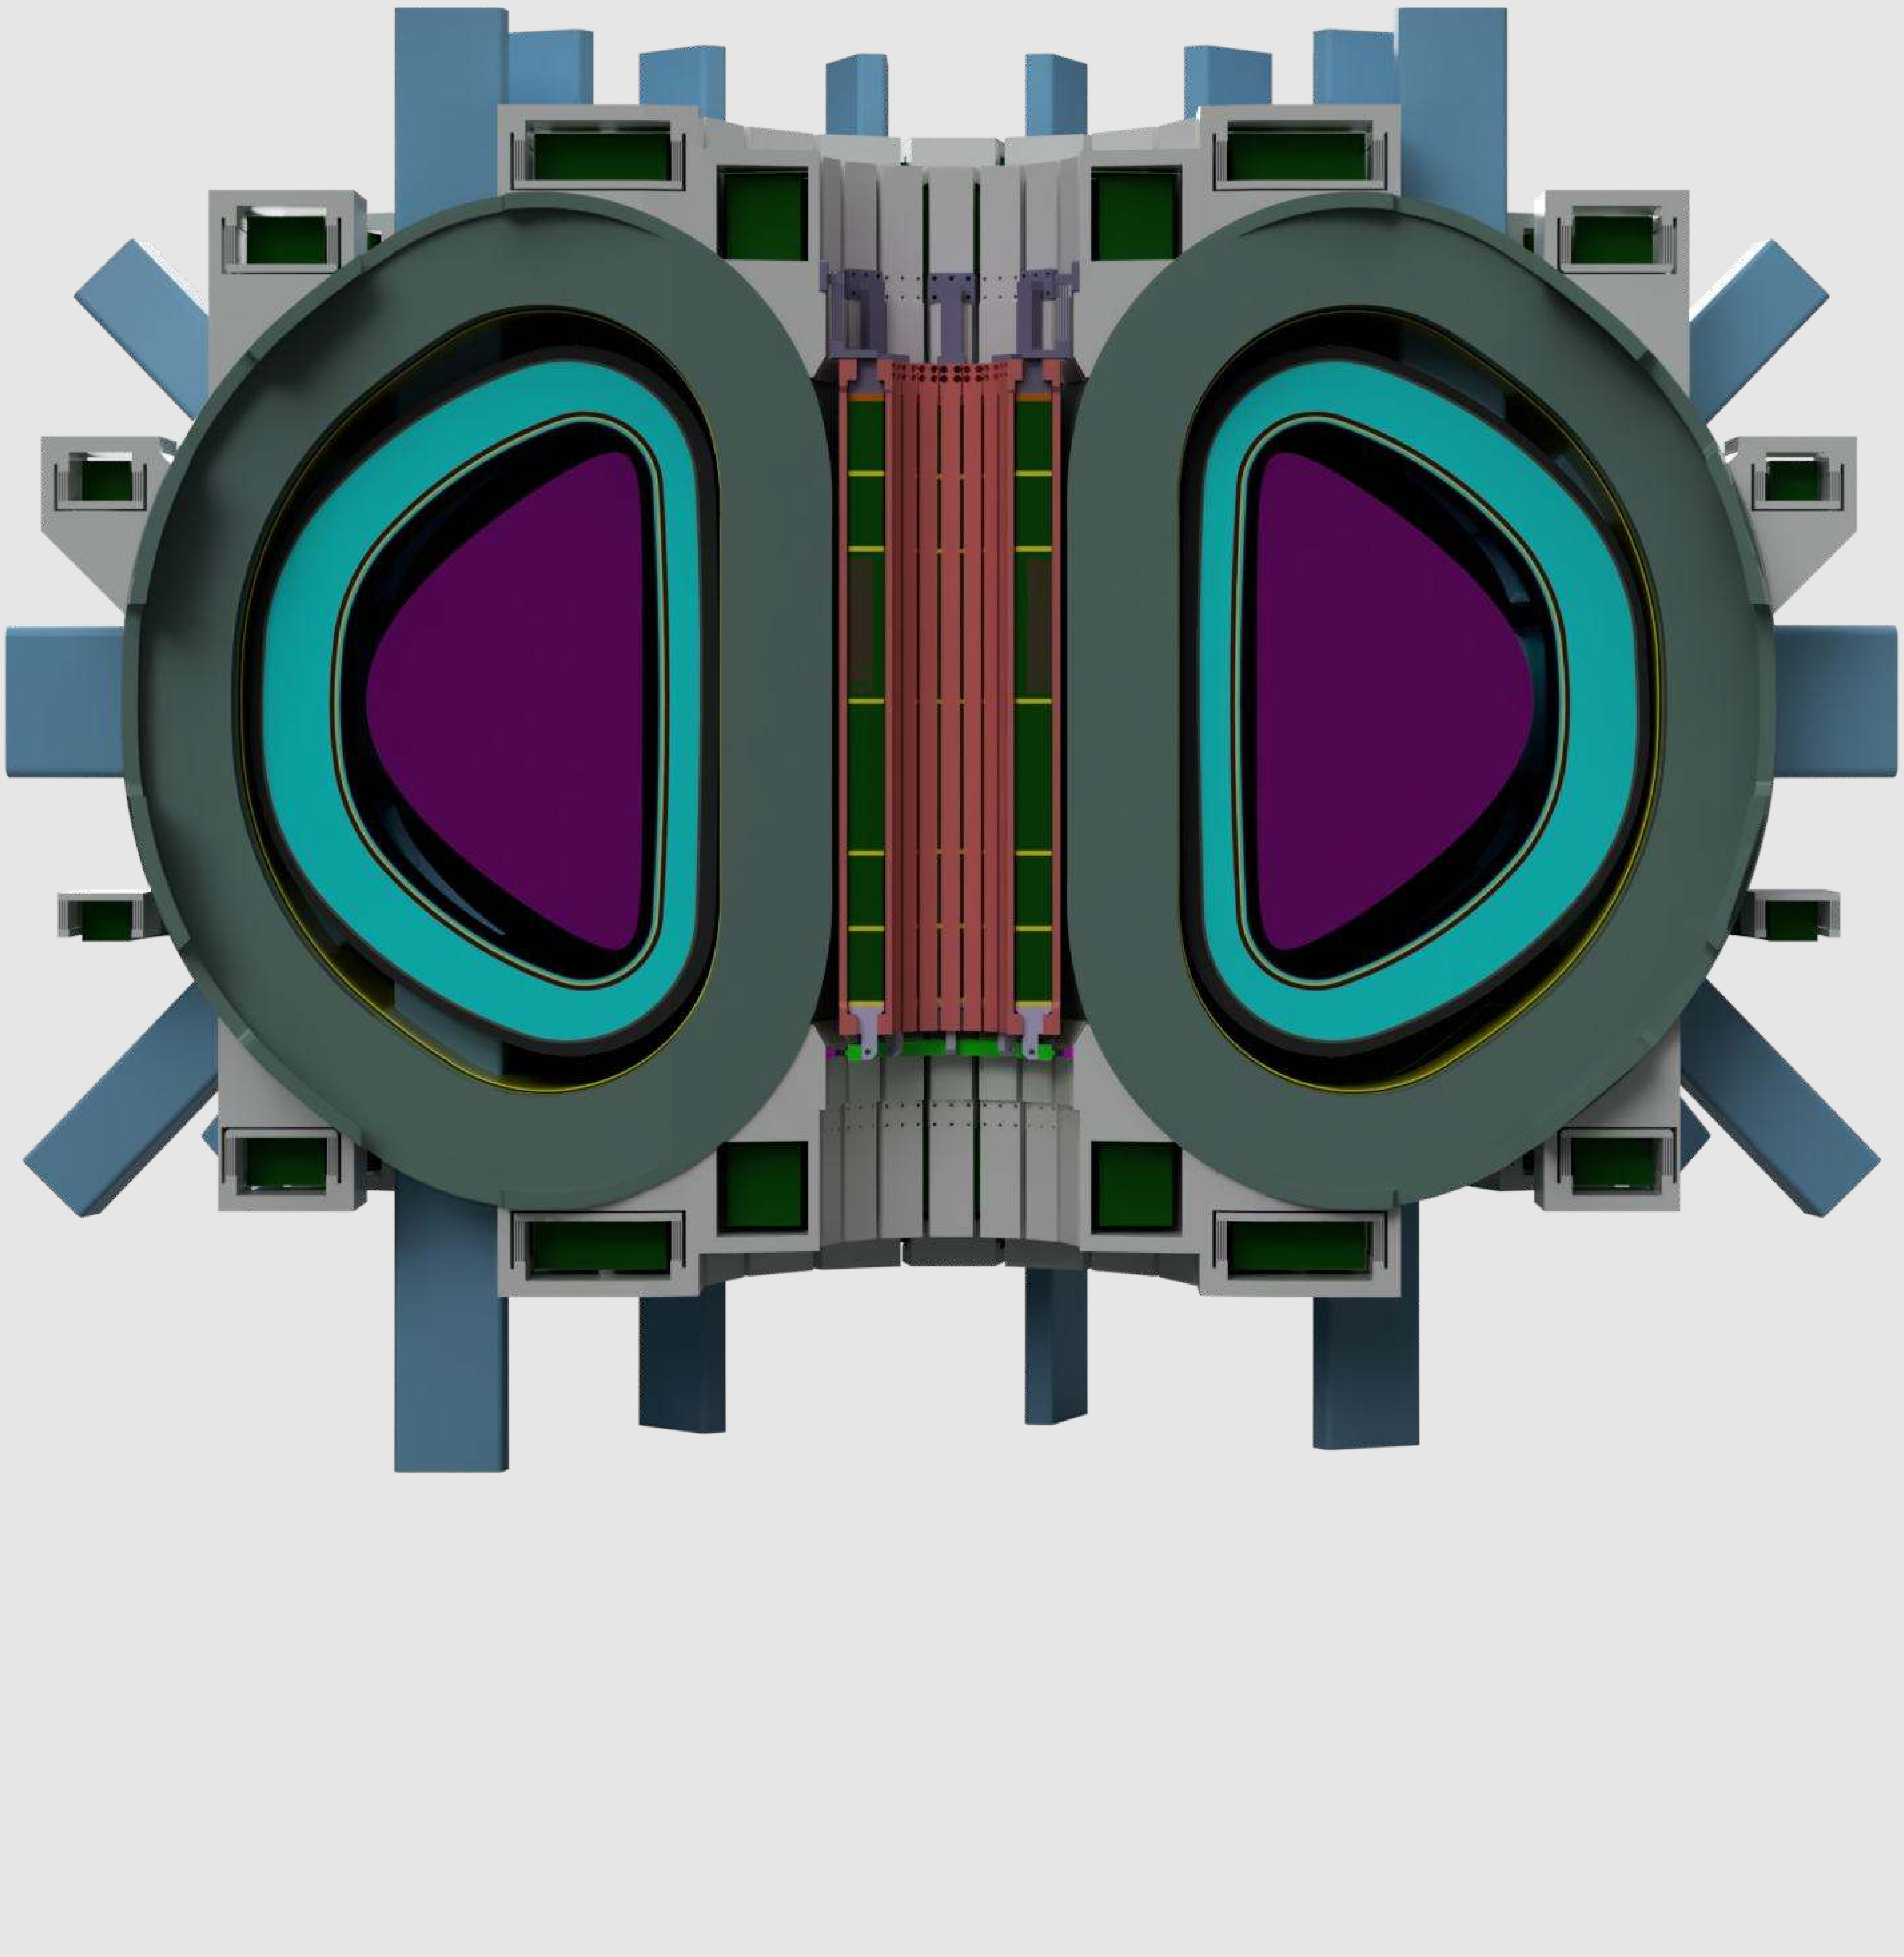

Go-to-Market

Starting with Commercial Shipping and Defense

The commercial shipping industry is actively seeking alternative energy solutions, investing heavily in infrastructure-intensive and costly fuels such as ammonia and hydrogen. Aware of impending carbon taxes, the industry is welcoming of innovative new technologies. It's the ideal gateway market to bring fusion to the world. By starting with ships, fusion gains a critical early market—generating revenue, scaling production, and proving reliability long before achieving competitiveness in traditional energy markets. However, fusion energy's transformative potential **extends far beyond commercial shipping applications.** Here we outline additional applicable markets where FOAK fusion reactors can have a meaningful impact, delivering emission free, safe energy.

Application	Market Opportunity	Value Proposition	Adoption Strategy
Commercial Shipping Primary Application	The maritime shipping industry **moves around 90% of global trade** and faces increasing pressure from the International Maritime Organization (IMO) to rapidly cut emissions. Existing alternative fuels like LNG, hydrogen, and ammonia have significant drawbacks, including storage challenges and infrastructure constraints. With tightening global regulations and rising fuel costs, adopting fusion-based solutions will become a more attractive option.	Fusion-powered ships redefine maritime sustainability by completely eliminating dependence on traditional fuel. They deliver unmatched energy density, extended operational lifespans, and superior efficiency—empowering shipping companies to achieve net-zero emissions goals without the drawbacks and compromises of existing alternative fuels.	Focus on early adopters like **Maersk and MSC**, already investing in alternative fuels. Partner with shipbuilders and regulators for compliance and integration. Launch a pilot vessel to prove viability, then scale production for fleet-wide adoption. 1,000 container ships must decarbonize by 2050. Reactors cost ~$1B initially, dropping to ~$600M at scale. Selling 10–100/year (15% margin) yields **~$1B–$10B annual profit.**
Defense - Ship to Shore Power Primary Application	Military operations increasingly require reliable, large-scale, and mobile power generation, particularly in **contested theaters like the Pacific**. Traditional energy logistics, heavily reliant on liquid fuel, are vulnerable, costly, and impose operational risks. Rapid deployment of resilient **nuclear-powered ships** is critical for both forward military bases and disaster response efforts, underscoring a significant strategic capability need within the Department of Defense (DoD). Fusion systems are also less of a concern for kinetic impact compared to fission, a major concern in the theater of war. More details here.	Fusion energy meets the DoD's critical need for mobile, large-scale power by providing compact, resilient, and fuel-independent solutions. Fusion eliminates reliance on vulnerable fuel supply chains and addresses logistical and security issues associated with diesel and nuclear fission. Minimal refueling requirements reduce operational risks and enhance mission endurance in forward-deployed and disaster-relief scenarios.	Partner with the **DoD, Navy, and allied military forces** to demonstrate fusion-powered mobile units in expeditionary operations, especially in the Pacific theater. Conduct pilot programs with U.S. Pacific Command to validate performance and reliability, positioning fusion as a core element of future military logistics.
Remote Diesel + Desalination Replacement	Remote communities, islands, and industries depend on costly, polluting diesel for power and water desalination. Fusion offers a **sustainable, long-term solution, cutting costs and environmental impact**. Expanding access to clean, limitless energy transforms economic prospects for isolated regions, driving new development without dependence on fossil fuels.	Fusion delivers energy independence by eliminating diesel reliance, enabling **large-scale desalination without constraints**. With zero emissions and no risk of spills, it's a clean, scalable solution for diverse energy needs. The ability to generate massive amounts of electricity in a small form factor makes fusion a game-changer for industrial expansion, rural electrification, and water security.	Partner with **governments and utilities** to bring fusion to remote communities, integrating with existing infrastructure and leveraging renewable energy incentives for deployment. Build partnerships with NGOs and global development organizations to accelerate fusion adoption in regions most affected by energy poverty and climate change.

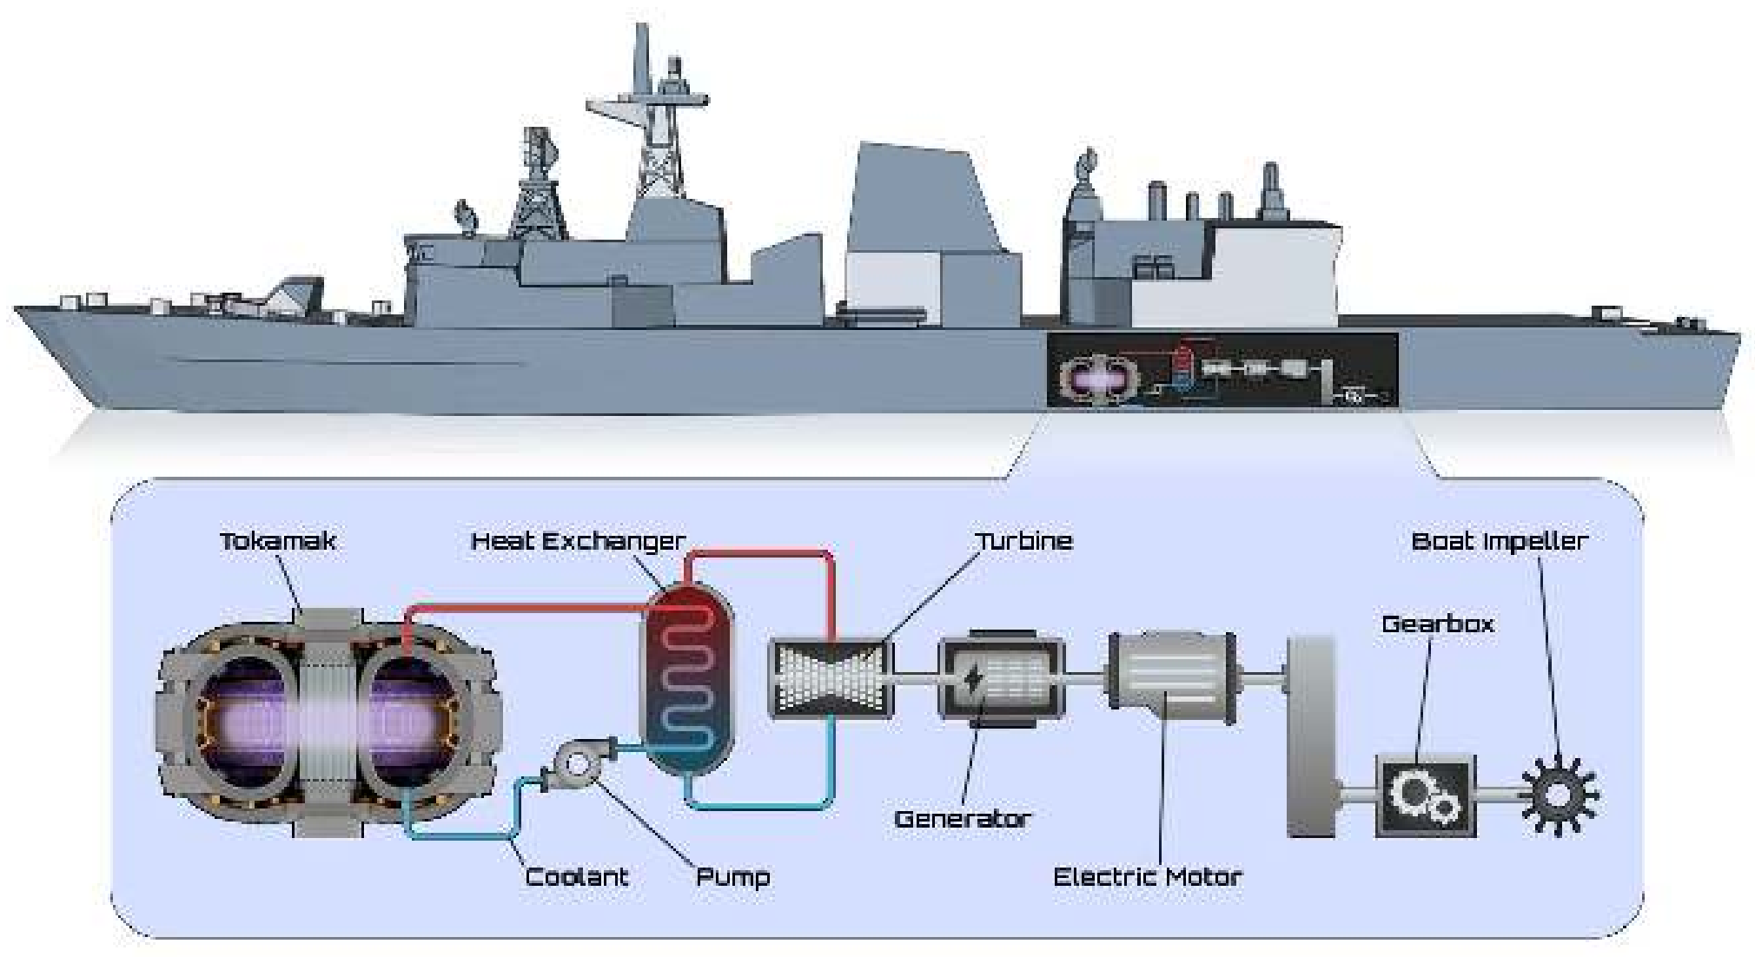

Powering the next frontier of ocean innovation

From Arctic expeditions to deep-sea mining and submerged data infrastructure, the demand for reliable and dependable energy in extreme marine environments is growing. Icebreakers require constant refueling in the harshest conditions, while mining and underwater data centers struggle with the limitations of diesel and battery power. **Fusion energy unlocks new possibilities**. We expect these markets to grow over the next few decades, and they will need a solution capable of providing consistent, long-duration power in challenging environments. The table below explores additional applications where fusion can transform ocean-based industries.

Application	Market Opportunity	Value Proposition	Adoption Strategy
Icebreaker Ships	Arctic operations face logistical hurdles due to **extreme environments and fuel scarcity**. Existing nuclear-powered icebreakers present regulatory and safety concerns, driving demand for cleaner, safer alternatives. Climate change has expanded Arctic shipping routes, increasing the need for reliable vessels with substantial power. Fusion technology addresses these challenges, supporting extended polar missions without typical supply chain issues.	Fusion-powered icebreakers can operate without fuel constraints, eliminating logistical challenges. They produce **no hazardous nuclear waste**, offer high power output for icebreaking, and reduce environmental risks. This will allow icebreakers to support critical research missions, emergency rescue operations, and year-round Arctic navigation.	Partner with **Arctic research initiatives, commercial operators,** and **governments** to integrate fusion technology into icebreaker fleets. Target both retrofits and new builds. Secure strategic collaborations with climate-focused organizations and expedition groups to accelerate adoption and policy support.
Deep Sea Mining	Extracting rare earth elements from the seabed requires self-sustaining energy sources, as diesel-based power is inefficient and polluting. **Battery-powered alternatives lack longevity for continuous operations.** The growing demand for rare earth metals in EVs, semiconductors, and renewable energy makes seabed mining a trillion-dollar industry.	Fusion empowers advanced subsea robotics and automated systems, substantially **improving productivity, operational depth, and environmental stewardship** for mining rare earth metals critical to renewable energy and technology sectors.	Collaborate with **mining companies** to develop pilot projects and advocate for fusion as a sustainable alternative. Work closely with regulatory agencies to position fusion as the standard for responsible deep-sea resource extraction.
Underwater Datacenters	Tech companies are exploring submerged data centers for improved cooling and reduced land usage. **These require stable, long-term power sources, independent of terrestrial grids.** The AI revolution and global data demands are pushing existing infrastructure to its limits. Fusion ensures scalable, emissions-free power, making underwater data centers a viable long-term solution.	Fusion ensures **continuous, off-grid power with minimal environmental impact**. Fusion-powered data hubs enable ultra-secure, resilient computing networks, reducing reliance on unstable grid power and offering redundancy.	Engage with cloud providers like **Microsoft, Google, AWS and other startups** to explore pilot projects. Develop demonstration facilities showcasing fusion's role in sustainable data infrastructure. Work with governments and telecom leaders to establish fusion as the foundation of next-generation global internet infrastructure.


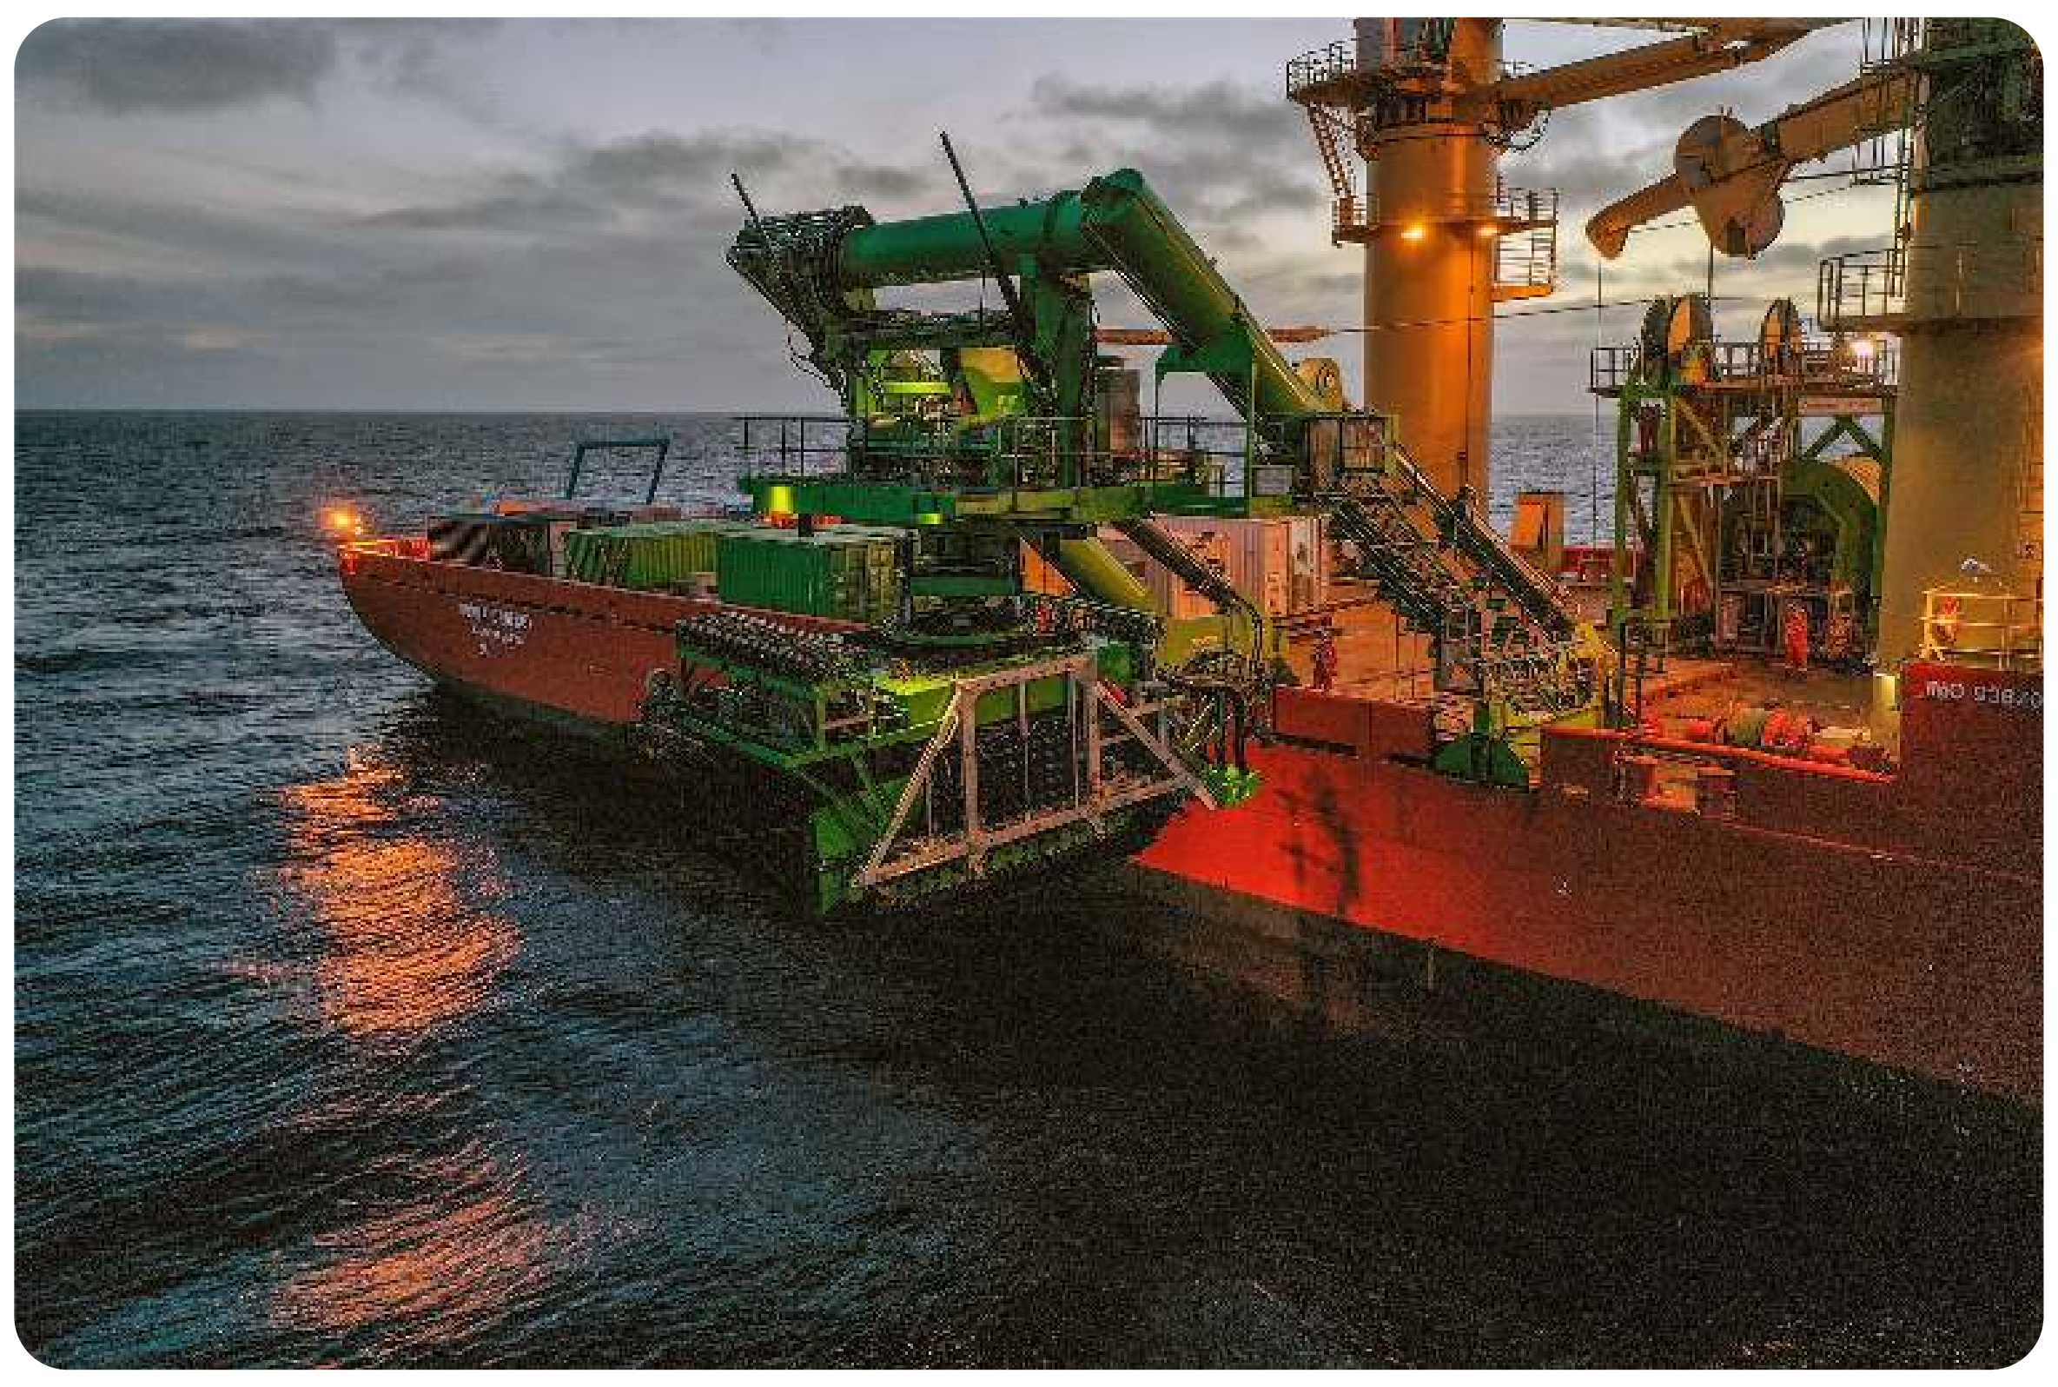


Chapter 6

Company

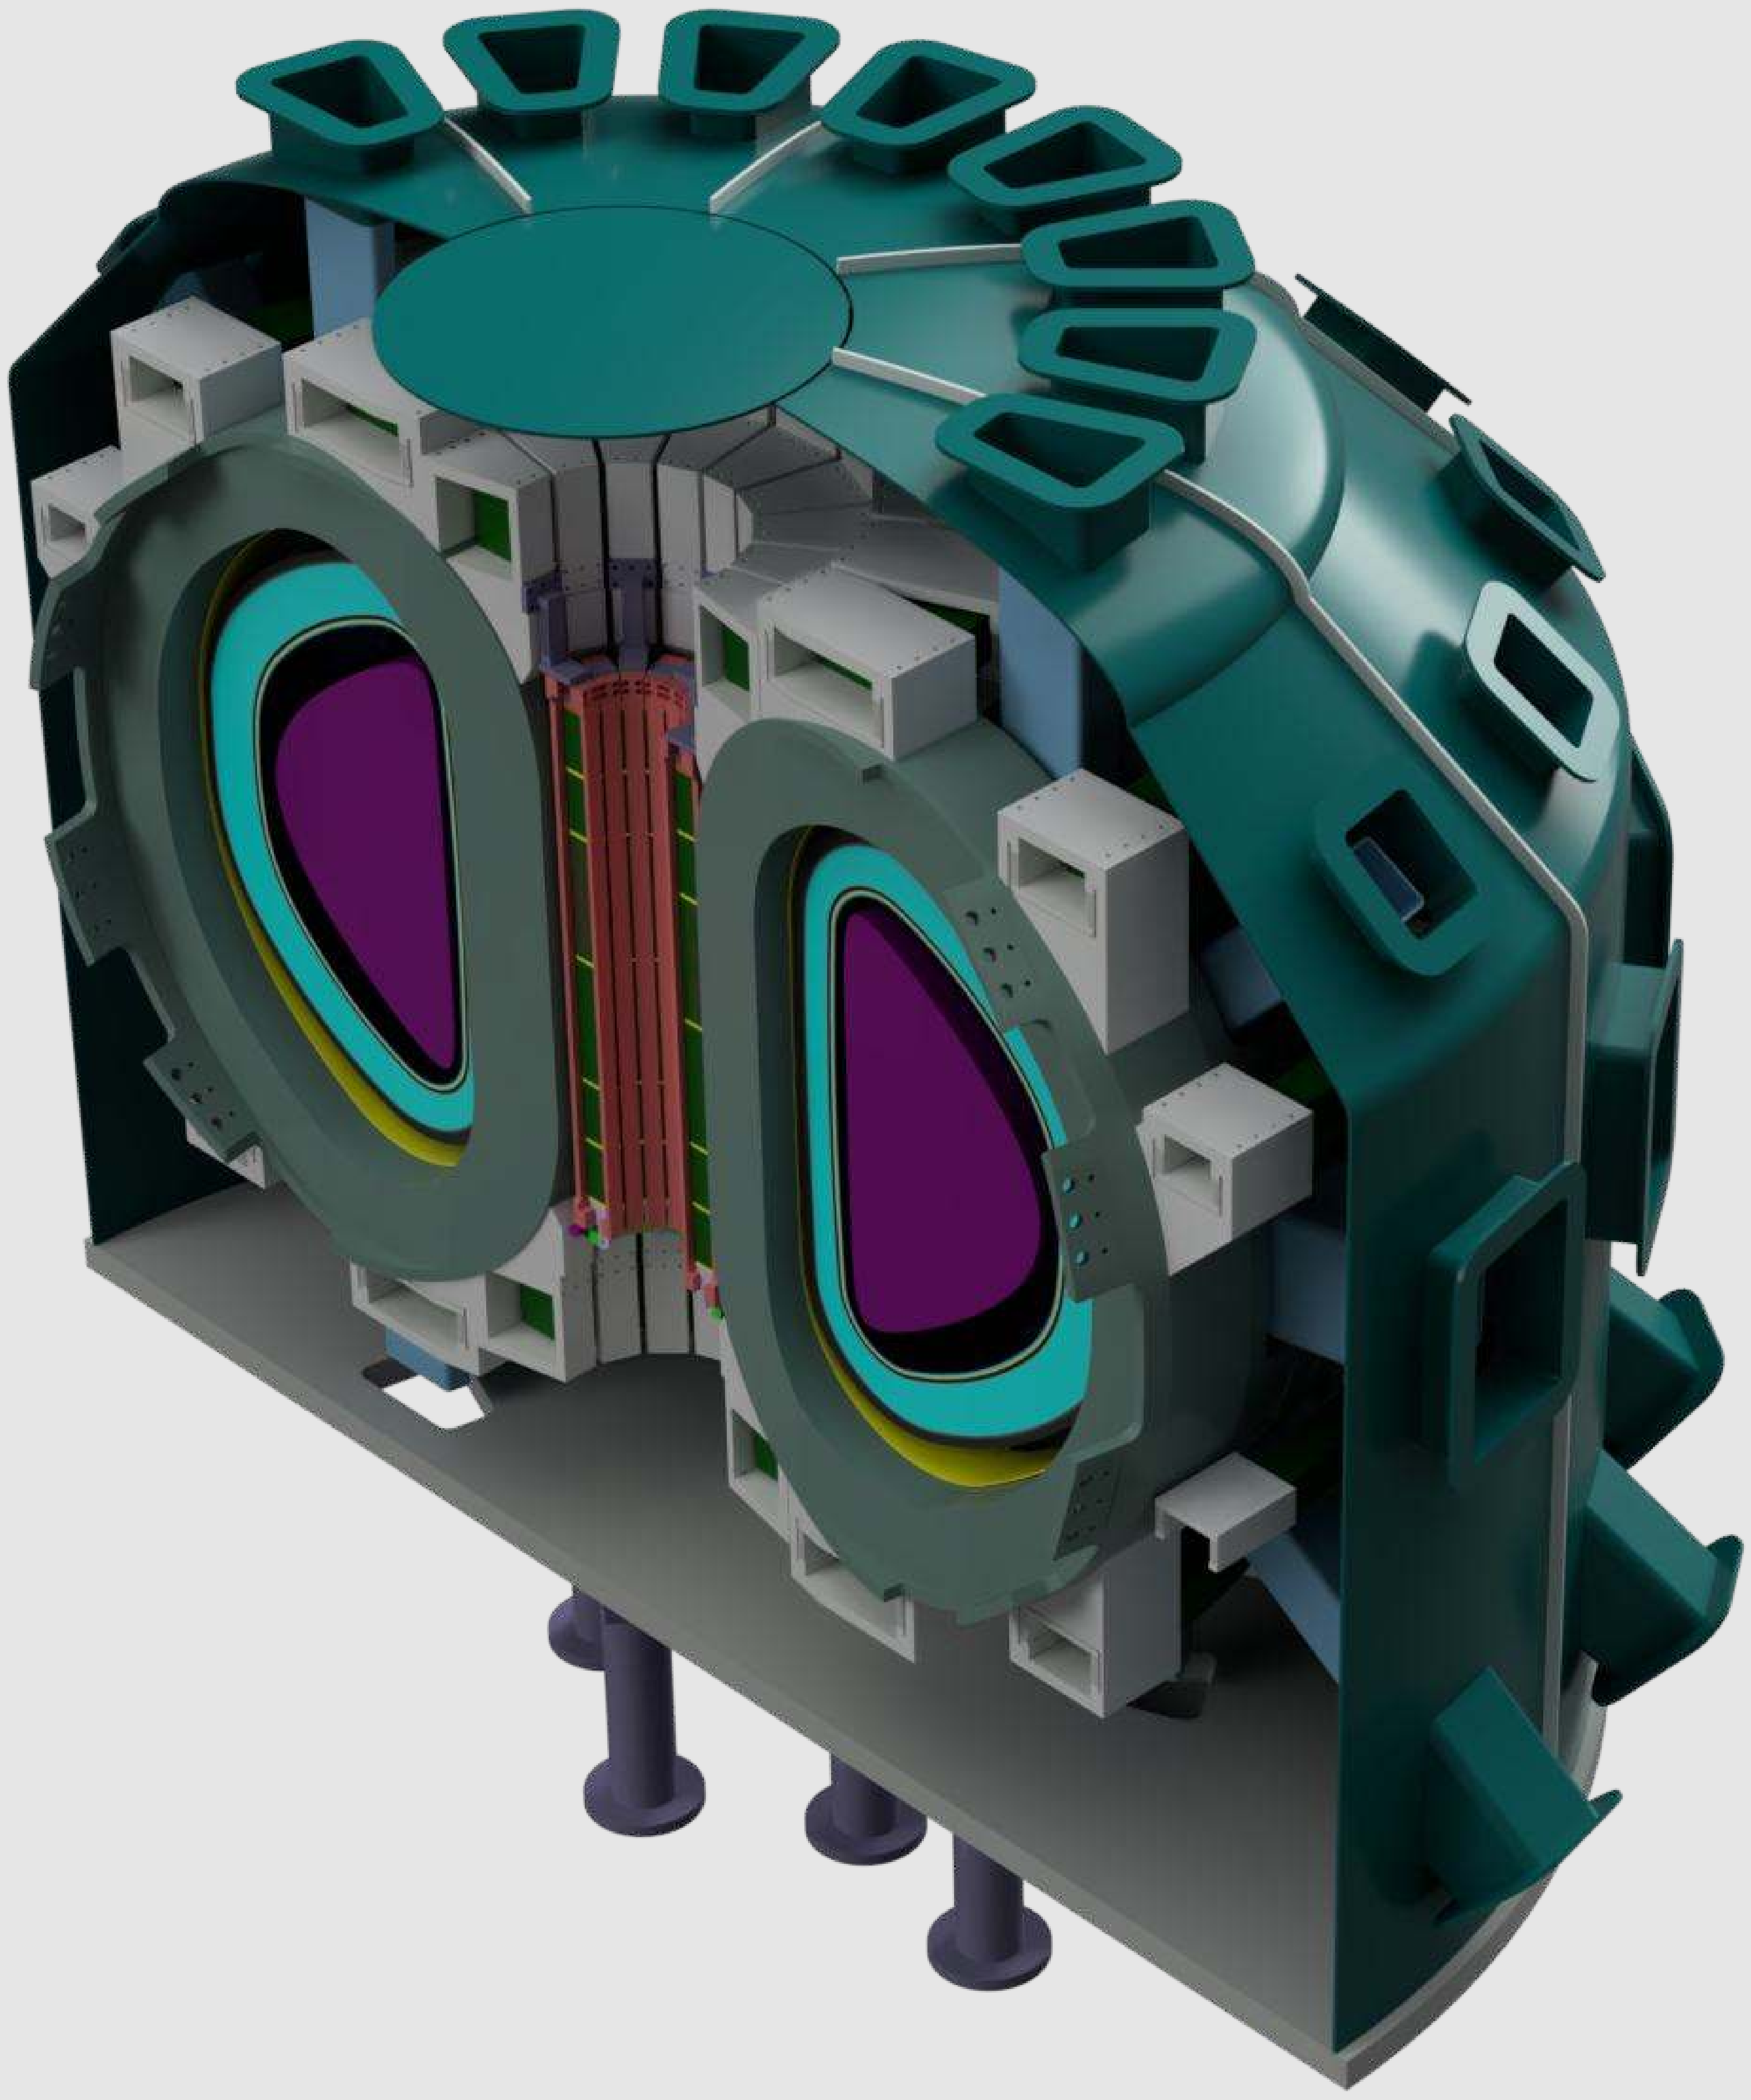

Meet The Team

We are a small but highly skilled team dedicated to bringing fusion power to the world. With deep expertise in nuclear engineering, plasma physics, and hardware systems, we are well suited to grow a company capable of tackling the applied engineering problems that are necessary to build practical fusion reactors for ships. We're here to turn theory into reality.

Justin Cohen, CEO + Co-Founder
Nuclear Engineer & Plasma Physicist

Expert in fusion reactor technology with experience spanning the gauntlet of approaches to fusion. Extensive experience in plasma physics and nuclear engineering, with roles at Tesla, SpaceX, Columbia University, Princeton Plasma Physics Laboratory, and NC State.

Columbia University · NC State · TESLA · SPACEX · PPPL Princeton Plasma Physics Laboratory

Jason Kaufmann
Electrical Engineer & Computational Physics

Physics and computer engineering background from Penn. Pioneered megawatt-scale power electronics for Tesla's Dojo and designed hardware flying on rockets and satellites in orbit right now at SpaceX and Astranis. Has designed a full scale electric race car from scratch.

Penn University of Pennsylvania · TESLA · etched · ASTRANIS · SPACEX

Backed by: Y Combinator

Milestone Plan

Fusion reactors are not cheap. For this to be a viable business, we must devise a plan to warrant raising serious capital on our journey to building to a fusion reactor. To do this, we believe the milestone plan detailed below will allow us to demonstrate meaningful staged progress en route to the full scale reactor and justify our growing intake of capital.

2025 Q1 Pre-seed (YC)	2025Q2 Seed	2026 - 2028 Series A, B	2028 Series C	~2031 Delivery ✓	
Design and analysis, customer engagement	Build and commission full scale auxiliary system, cryogenic plant	Next largest auxiliary system (gyrotrons, magnet power supplies, energy storage, tritium fueling)	Magnets & Tokamak	Full scale application specific fusion reactor	**2030's scale** 10 reactors/yr by 2035, >$1B profit

Pre-seed
Incorporated the company and kicked off detailed design of the reactor. Began forming industry, government, and academic partnerships.

Series A and B will focus on replicating these at larger scale

Seed

- Peer reviewed design with detailed plasma physics basis, operational and maintenance plans
- Experimental campaign on existing facilities to materials for fusion relevant conditions and divertor scenarios.
- Grow the team. +10→100 kick ass engineers (nuclear, mechanical, electrical, thermal), plasma physicists, marine architecture.
- Build a full scale auxiliary system relevant for marine deployment. Demonstrating the de-risking of miniaturizing auxiliary systems.

$100M-$500M for the HTS magnets, vacuum vessel, cryostat.

Fusion powered revenue!

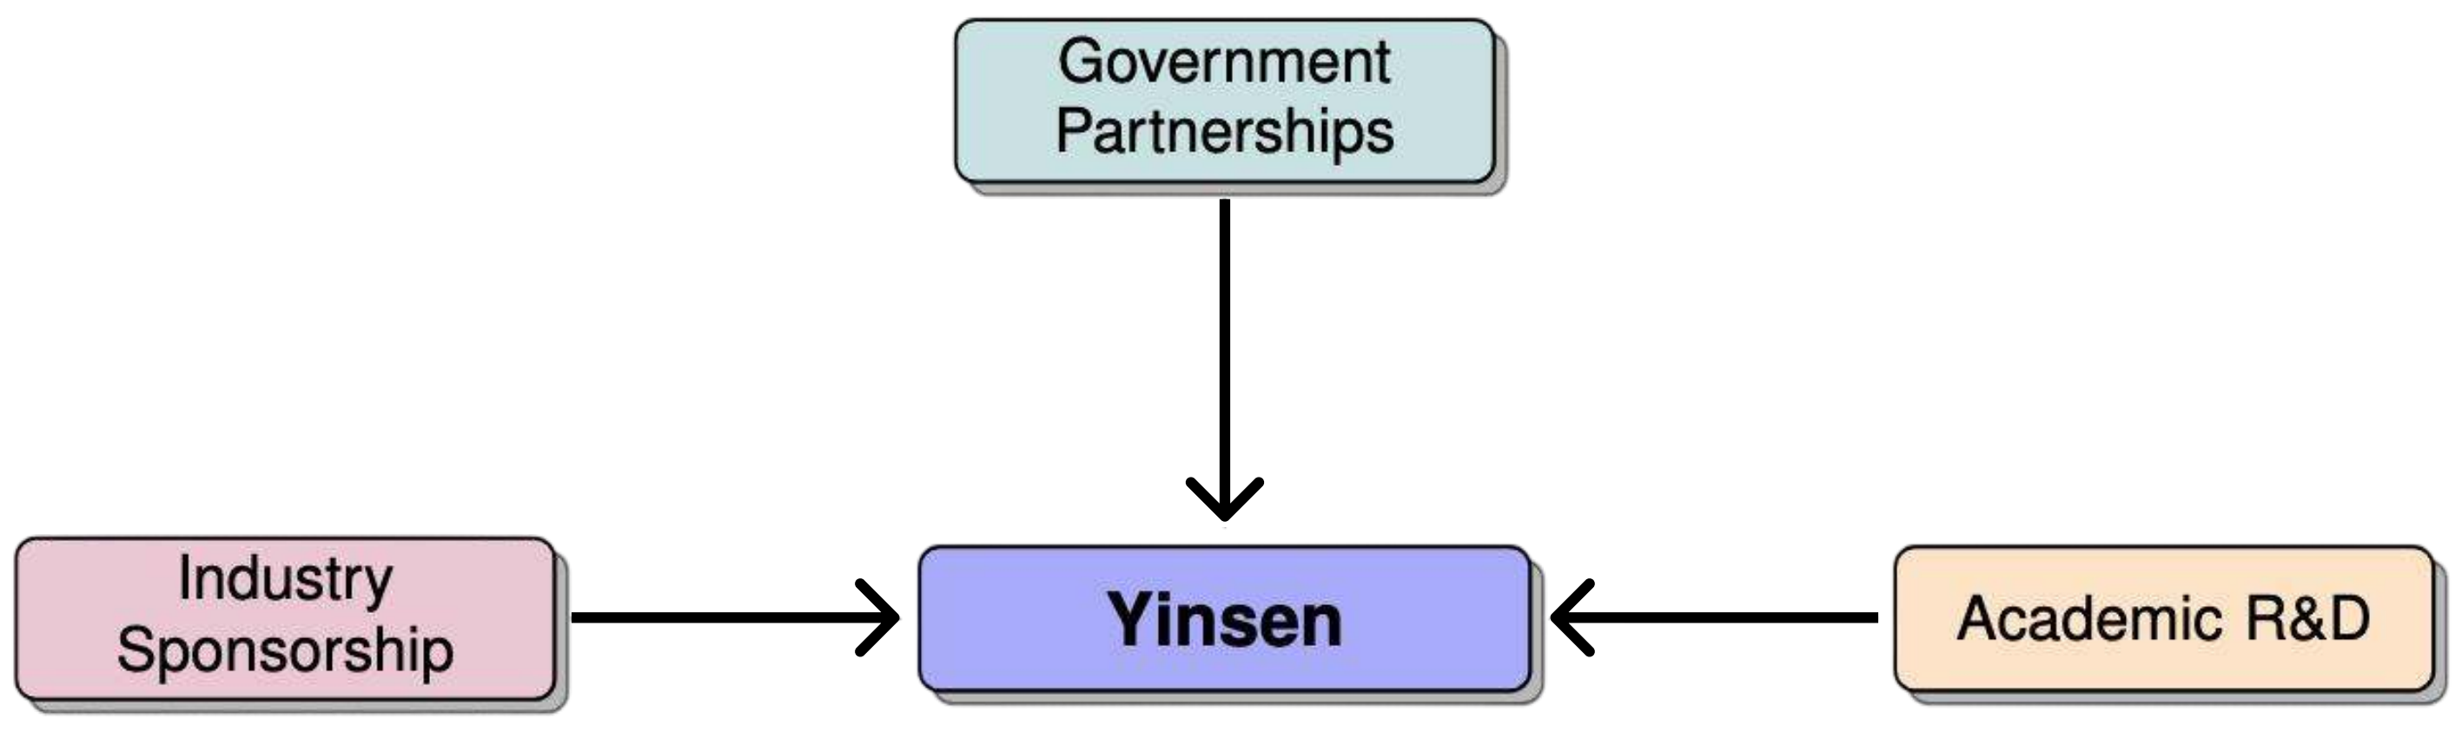

1kW Cryocooler from Absolut System

MW RF Gyrotron from QST & CETD (Japan)

Series C we 'fill in the gap' biting off the HTS magnets, Vacuum Vessel, and heart of the facility; the tokamak core.

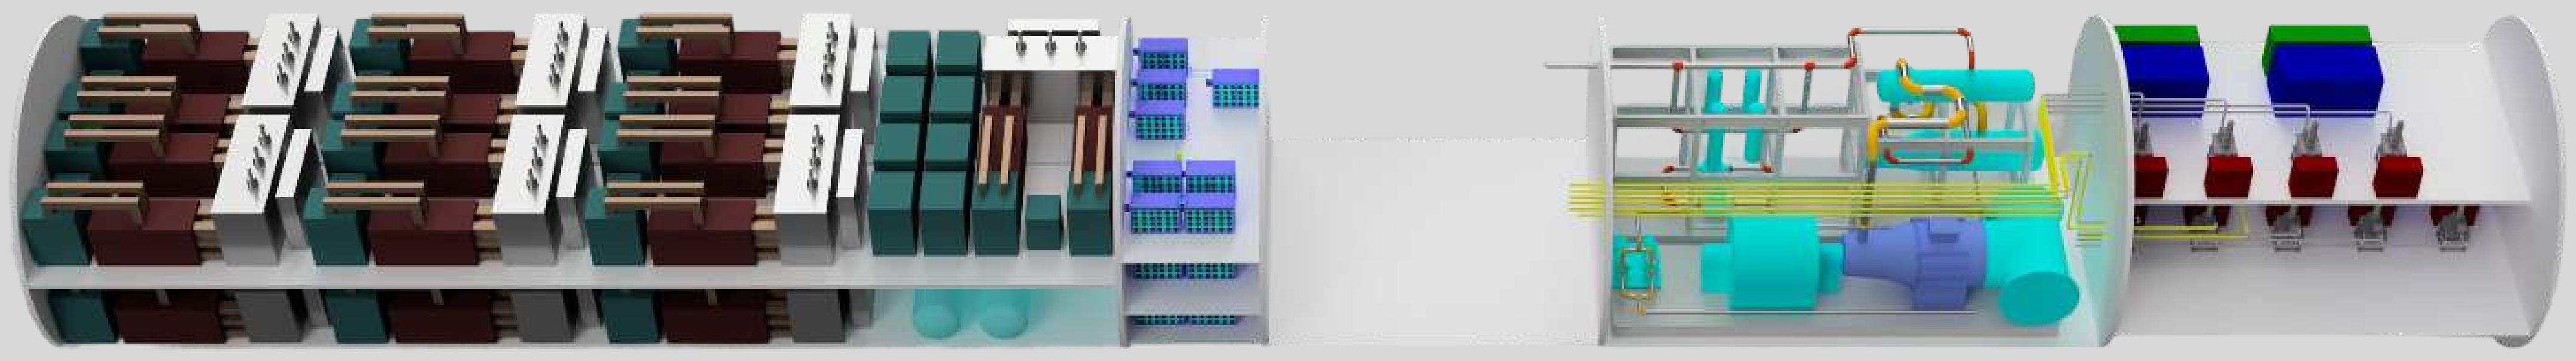

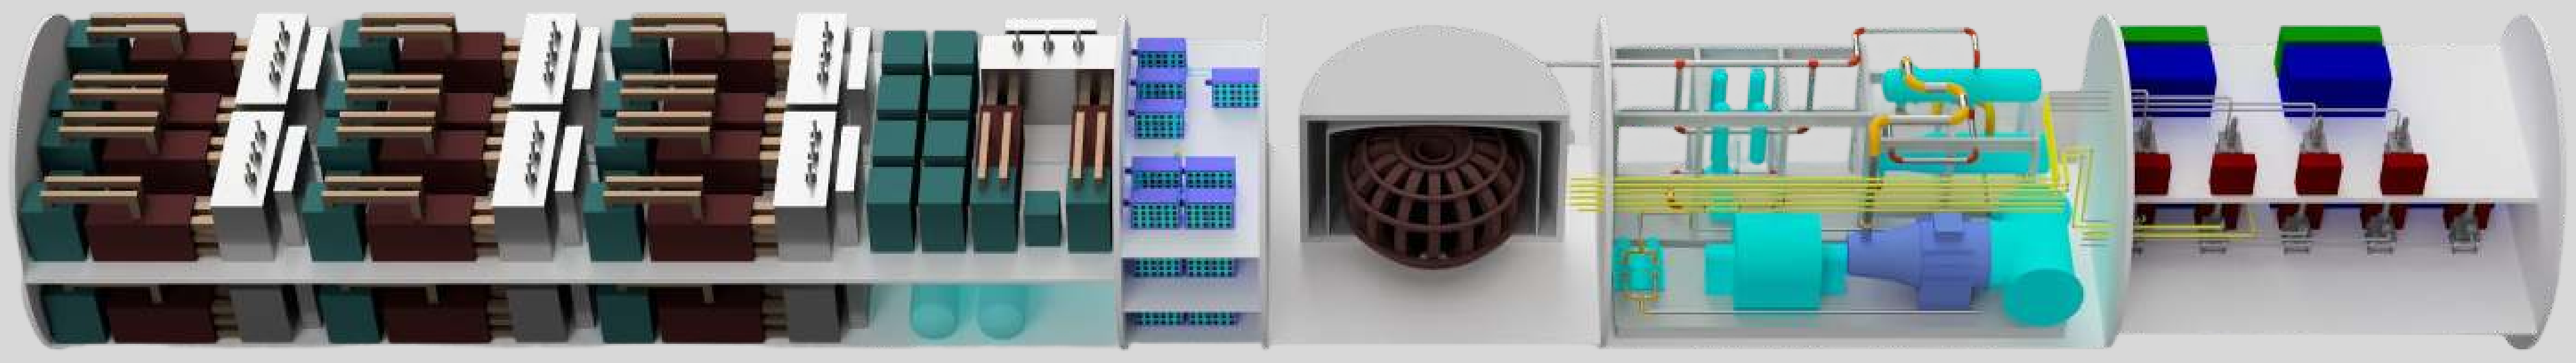

Delivery on a ship

Chapter 7
Vision

Vision

Fusion is the future

The Looming Energy Crisis

The world runs on energy, yet the foundation of our energy systems is crumbling. Fossil fuels—coal, oil, and natural gas—have powered civilization for centuries, but they are finite. **At current consumption rates, known oil reserves could be depleted within 50 years**, while natural gas and coal reserves face similar fates. The harder these fuels are to extract, the more expensive and unsustainable they become. Even as renewables scale, they remain unreliable. Solar and wind fluctuate, requiring storage solutions that are expensive and underdeveloped. **Meanwhile, global energy demand is set to double by 2050, driven by AI computing, urban expansion, and industry**. Without a breakthrough, the world faces an unavoidable energy crisis—one that threatens economic growth, technological progress, and geopolitical stability.

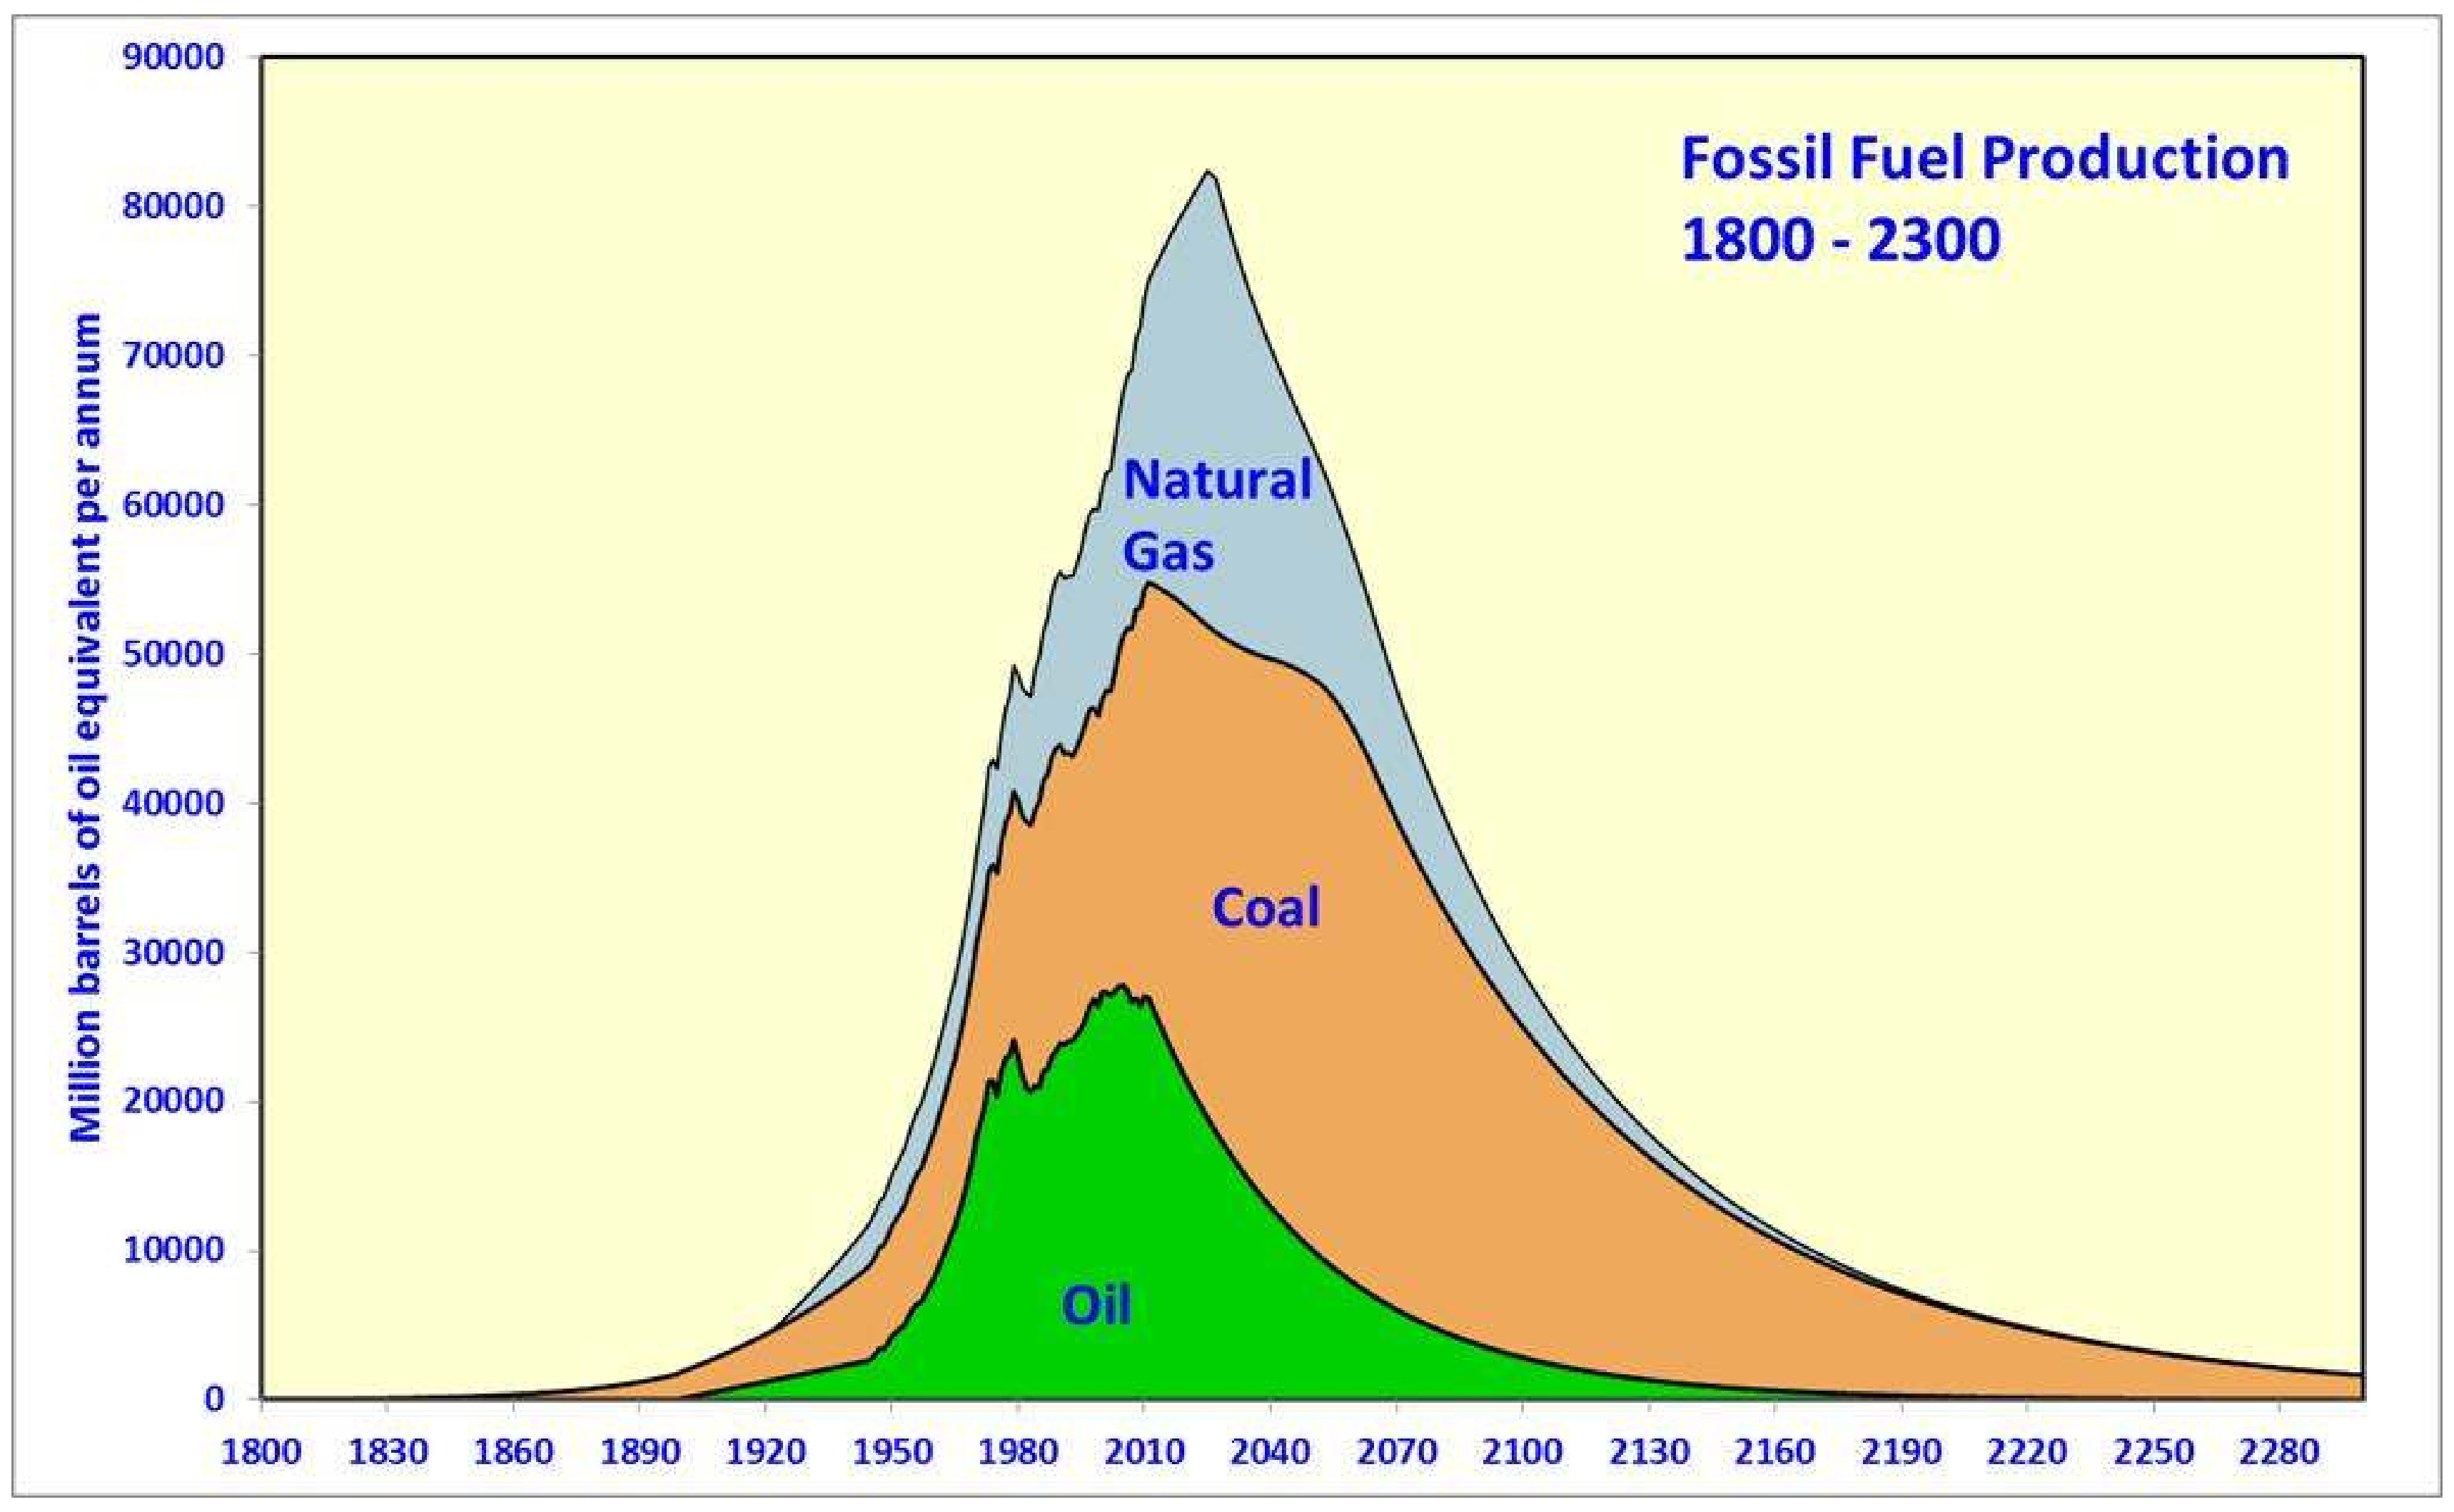

Advanced societies do not use less energy—they use more. Energy consumption is directly correlated with innovation, economic growth, and quality of life. A world constrained by energy shortages is a world of stagnation.

The answer has been in front of us all along

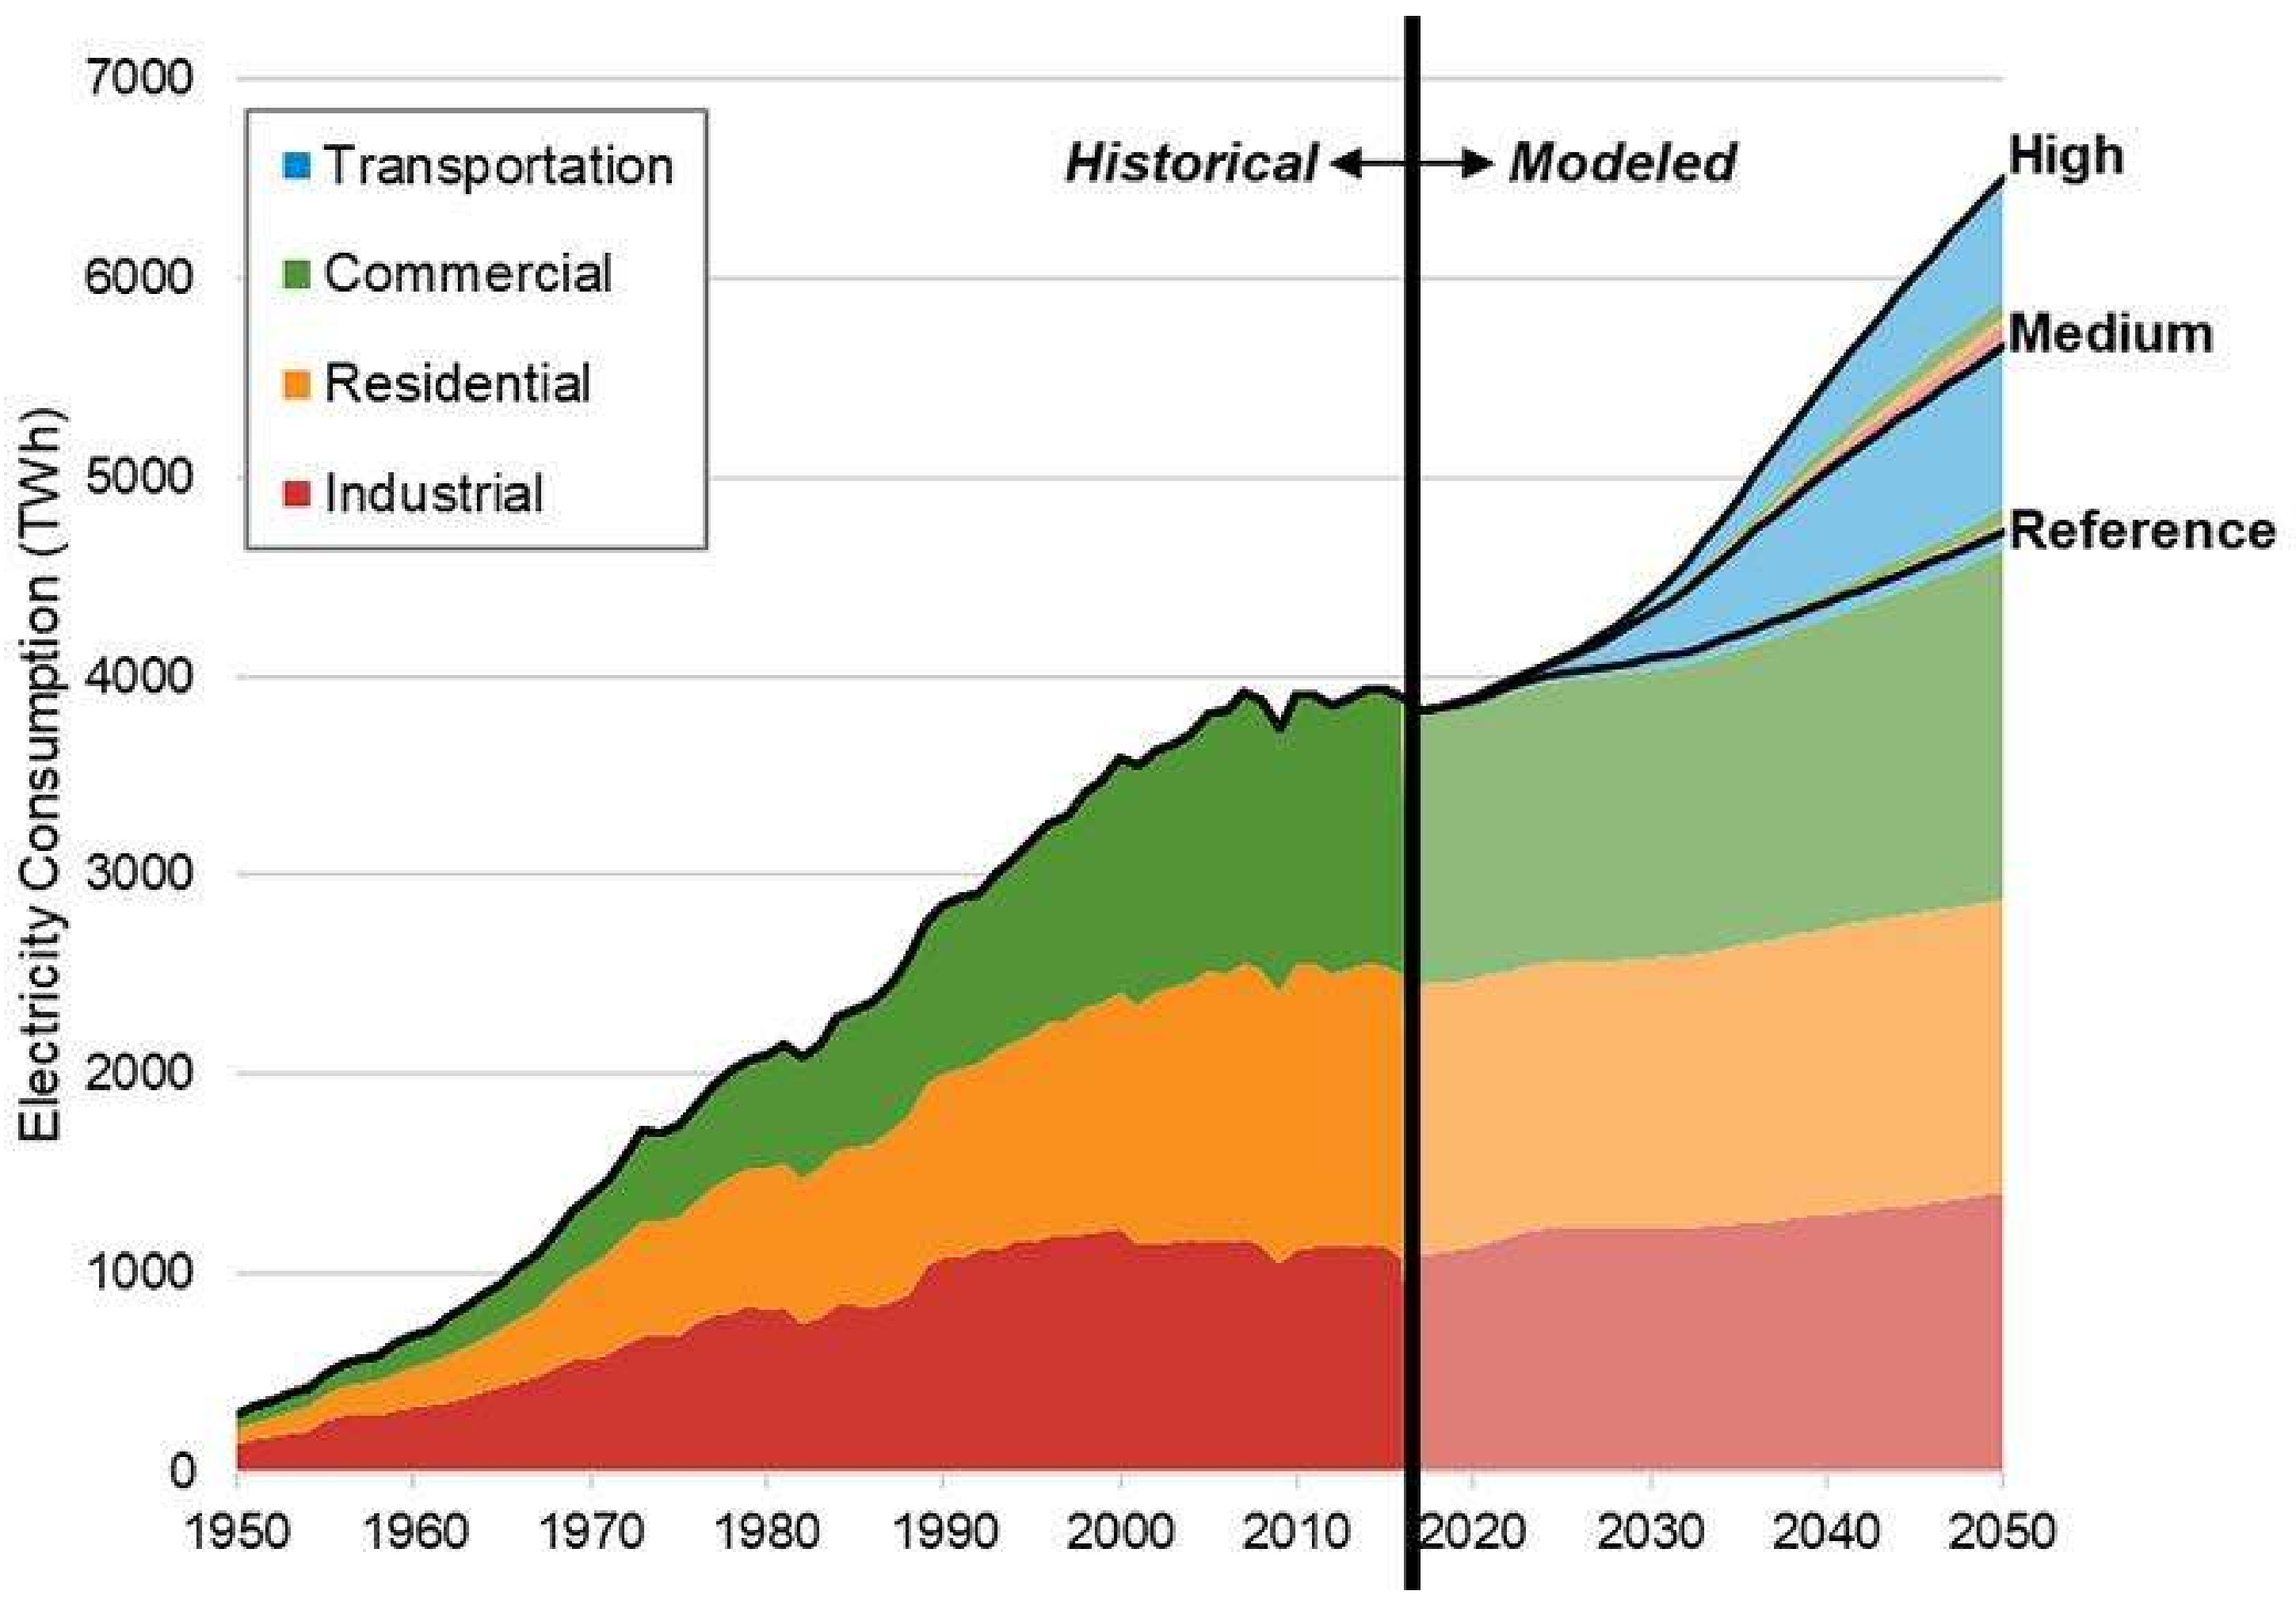

Nearly every watt of energy ever used on Earth—every fire lit, every barrel of oil burned—originates from a single source: **fusion.** The Sun fuses hydrogen atoms, unleashing energy far beyond any chemical reaction. Fossil fuels are simply ancient sunlight, energy stored from past fusion reactions. For the first time in history, we have the technology to replicate the Sun's power here on Earth.

Why burn the leftovers when we can tap into the source itself?

Fusion is not just another energy source—it is a civilizational unlock. It is essentially limitless, emissions-free, and requires fuel sources so abundant that we will never run out.

If energy was no longer scarce, what would that mean for the world?

Water Scarcity

Unlimited clean energy means unlimited de-salination

Food production

Vertical farming and abundant harvests—feeding humanity sustainably.

Computation & AI

No limits on training the most powerful models.

Transportation

Fusion propulsion makes sustainable travel feasible

Fusion is not about consuming less energy—it's about unlocking more.

But we need a thoughtful way to bring the first fusion reactors to market.

Fusion is the future

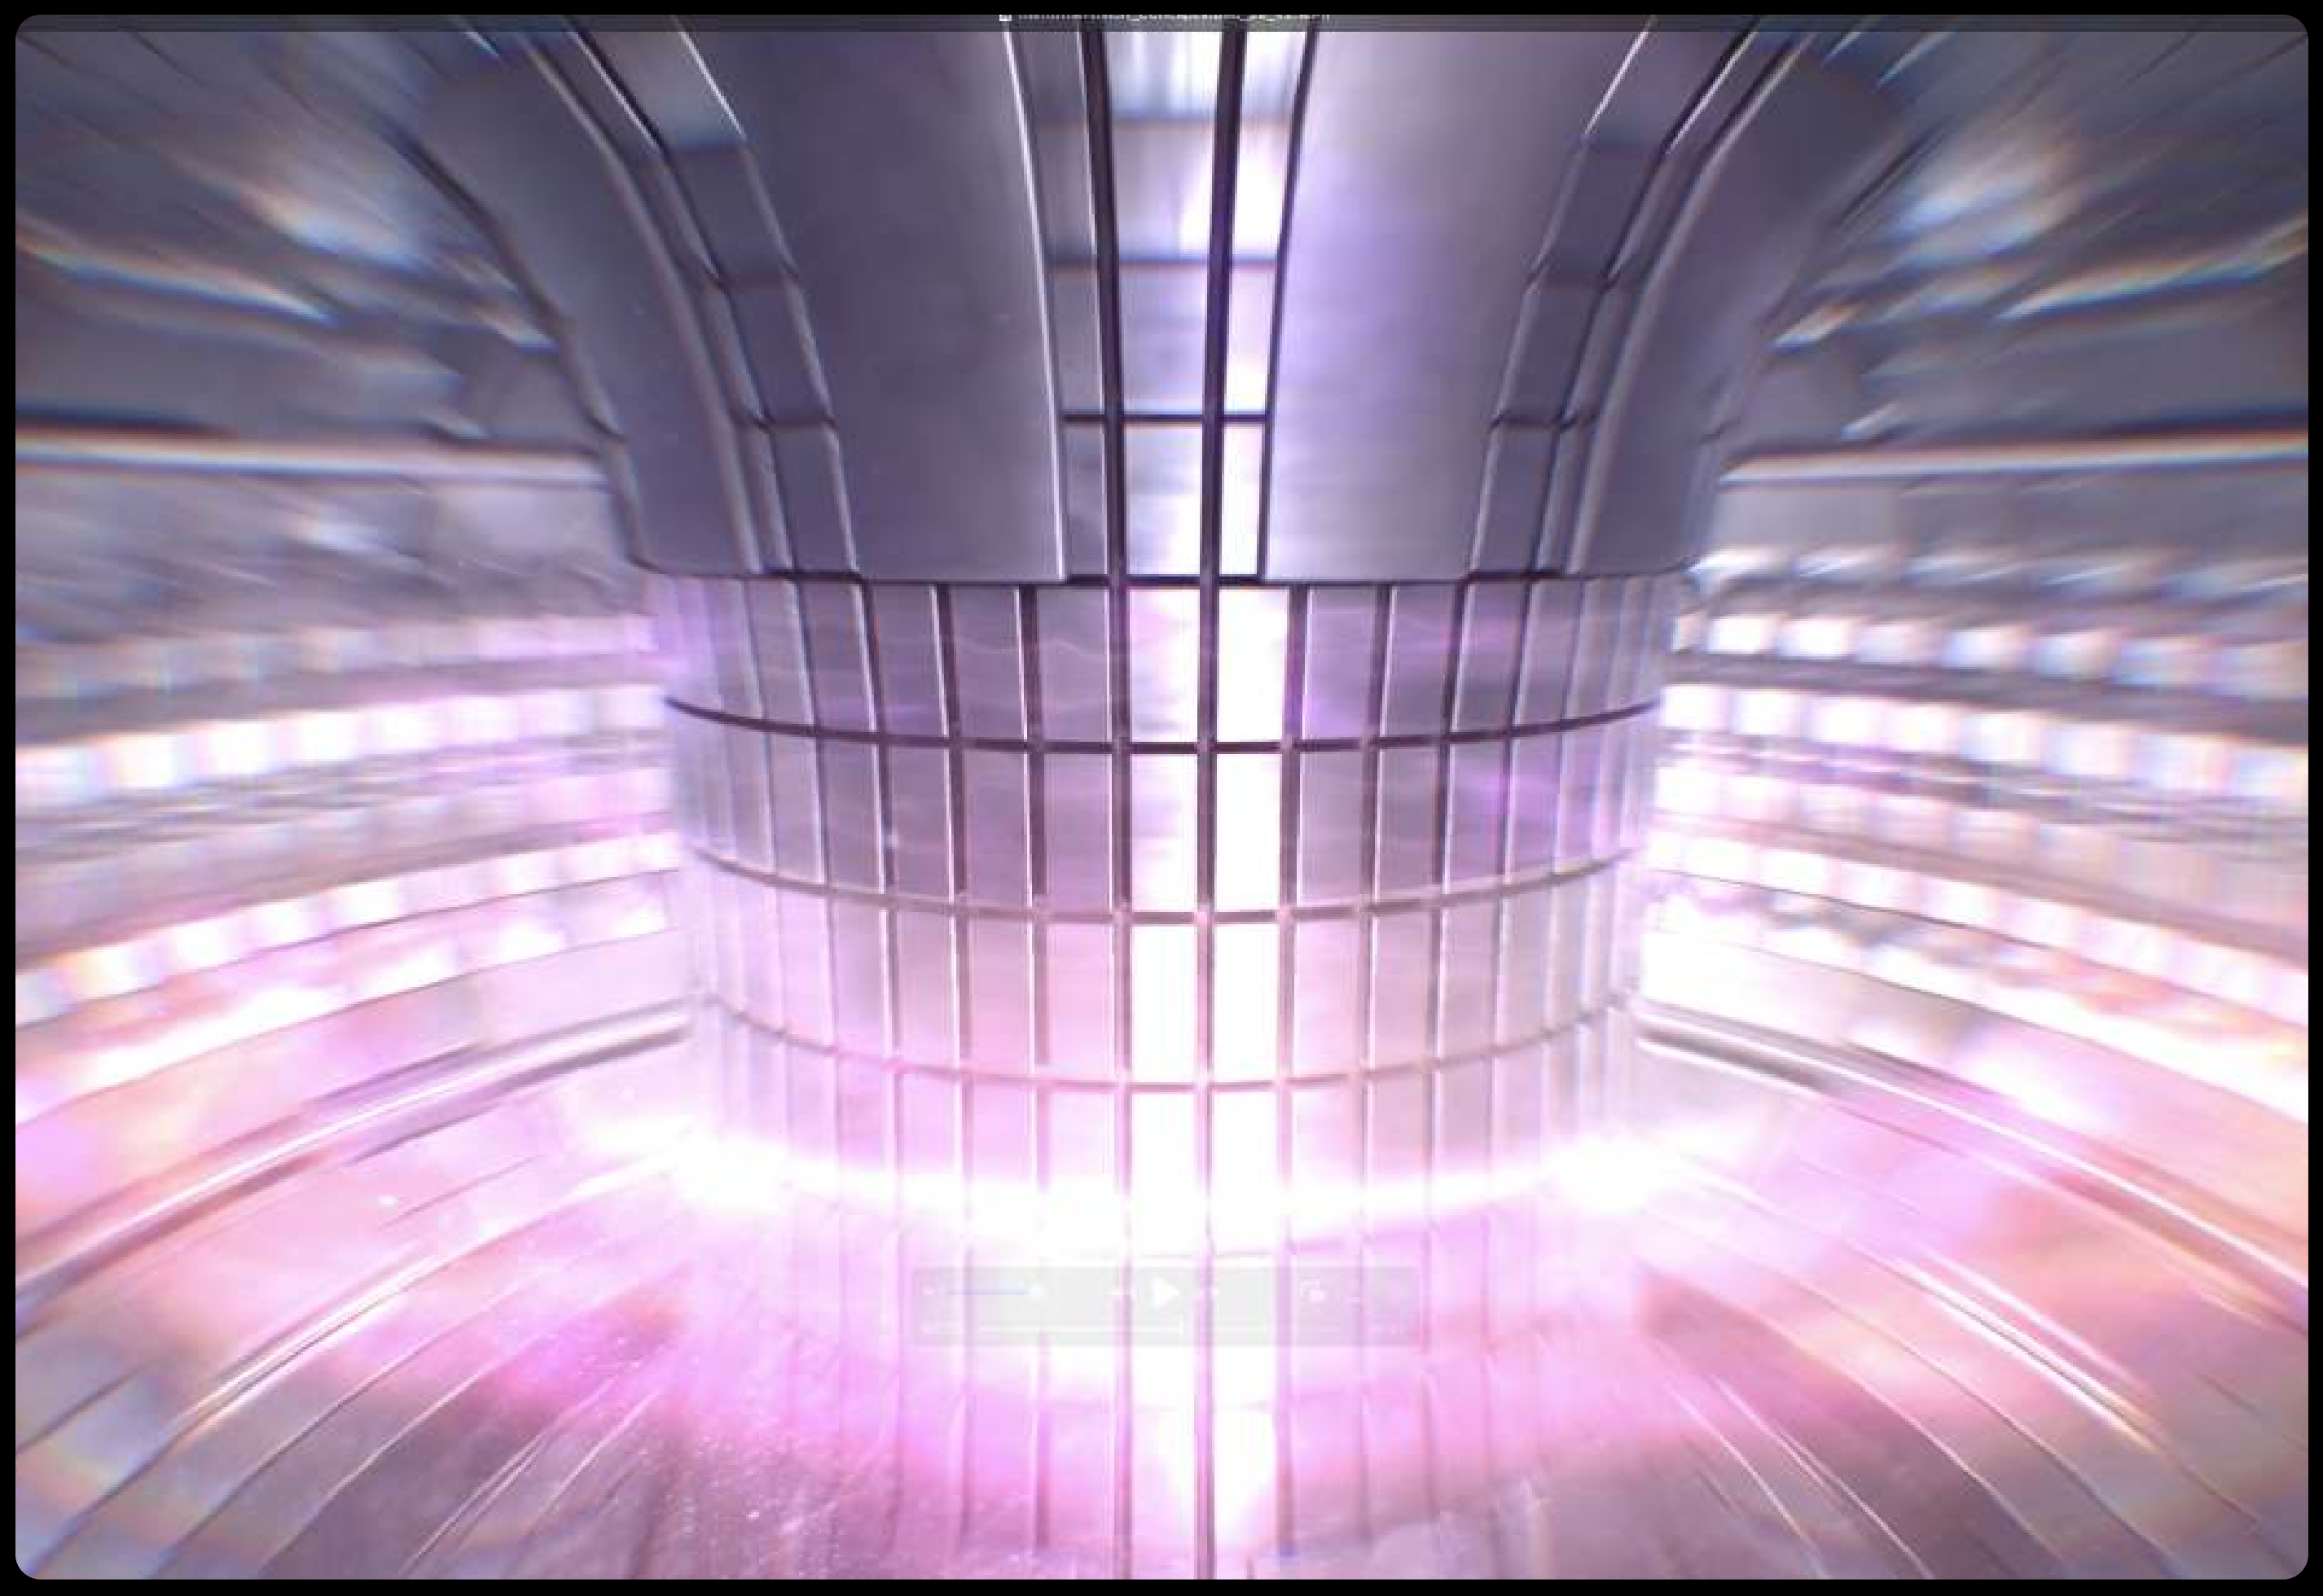

The promise of fusion is undeniable, but realizing it demands more than scientific progress—it requires an industrial revolution the likes of the **Manhattan Project** and **Apollo Program**. Unlike past energy transitions, which repurposed existing infrastructure, fusion needs an entirely new ecosystem. The reactors themselves are just one piece of a far larger puzzle. The true challenge lies in building the supply chains, refining fuel production, and restructuring the energy grid to distribute near-limitless power.

Scaling fusion isn't just about technological breakthroughs—it's about execution. Who will take fusion from lab experiments to real-world deployment? **The transition demands unprecedented collaboration between governments, industries, and innovators.** Policies must evolve, capital must flow, and industries must adapt to a world where energy is no longer a constraint but a catalyst for transformation.

This is no longer a question of if but when—and who will lead. The world's largest economies are already competing to dominate fusion, with billions flowing into research and commercialization. The U.S., China, and the European Union have made bold commitments, but private companies are moving faster than governments. The economic stakes are massive. **The first countries to commercialize fusion will control the energy future—while others risk being left behind in a world they no longer dictate.**

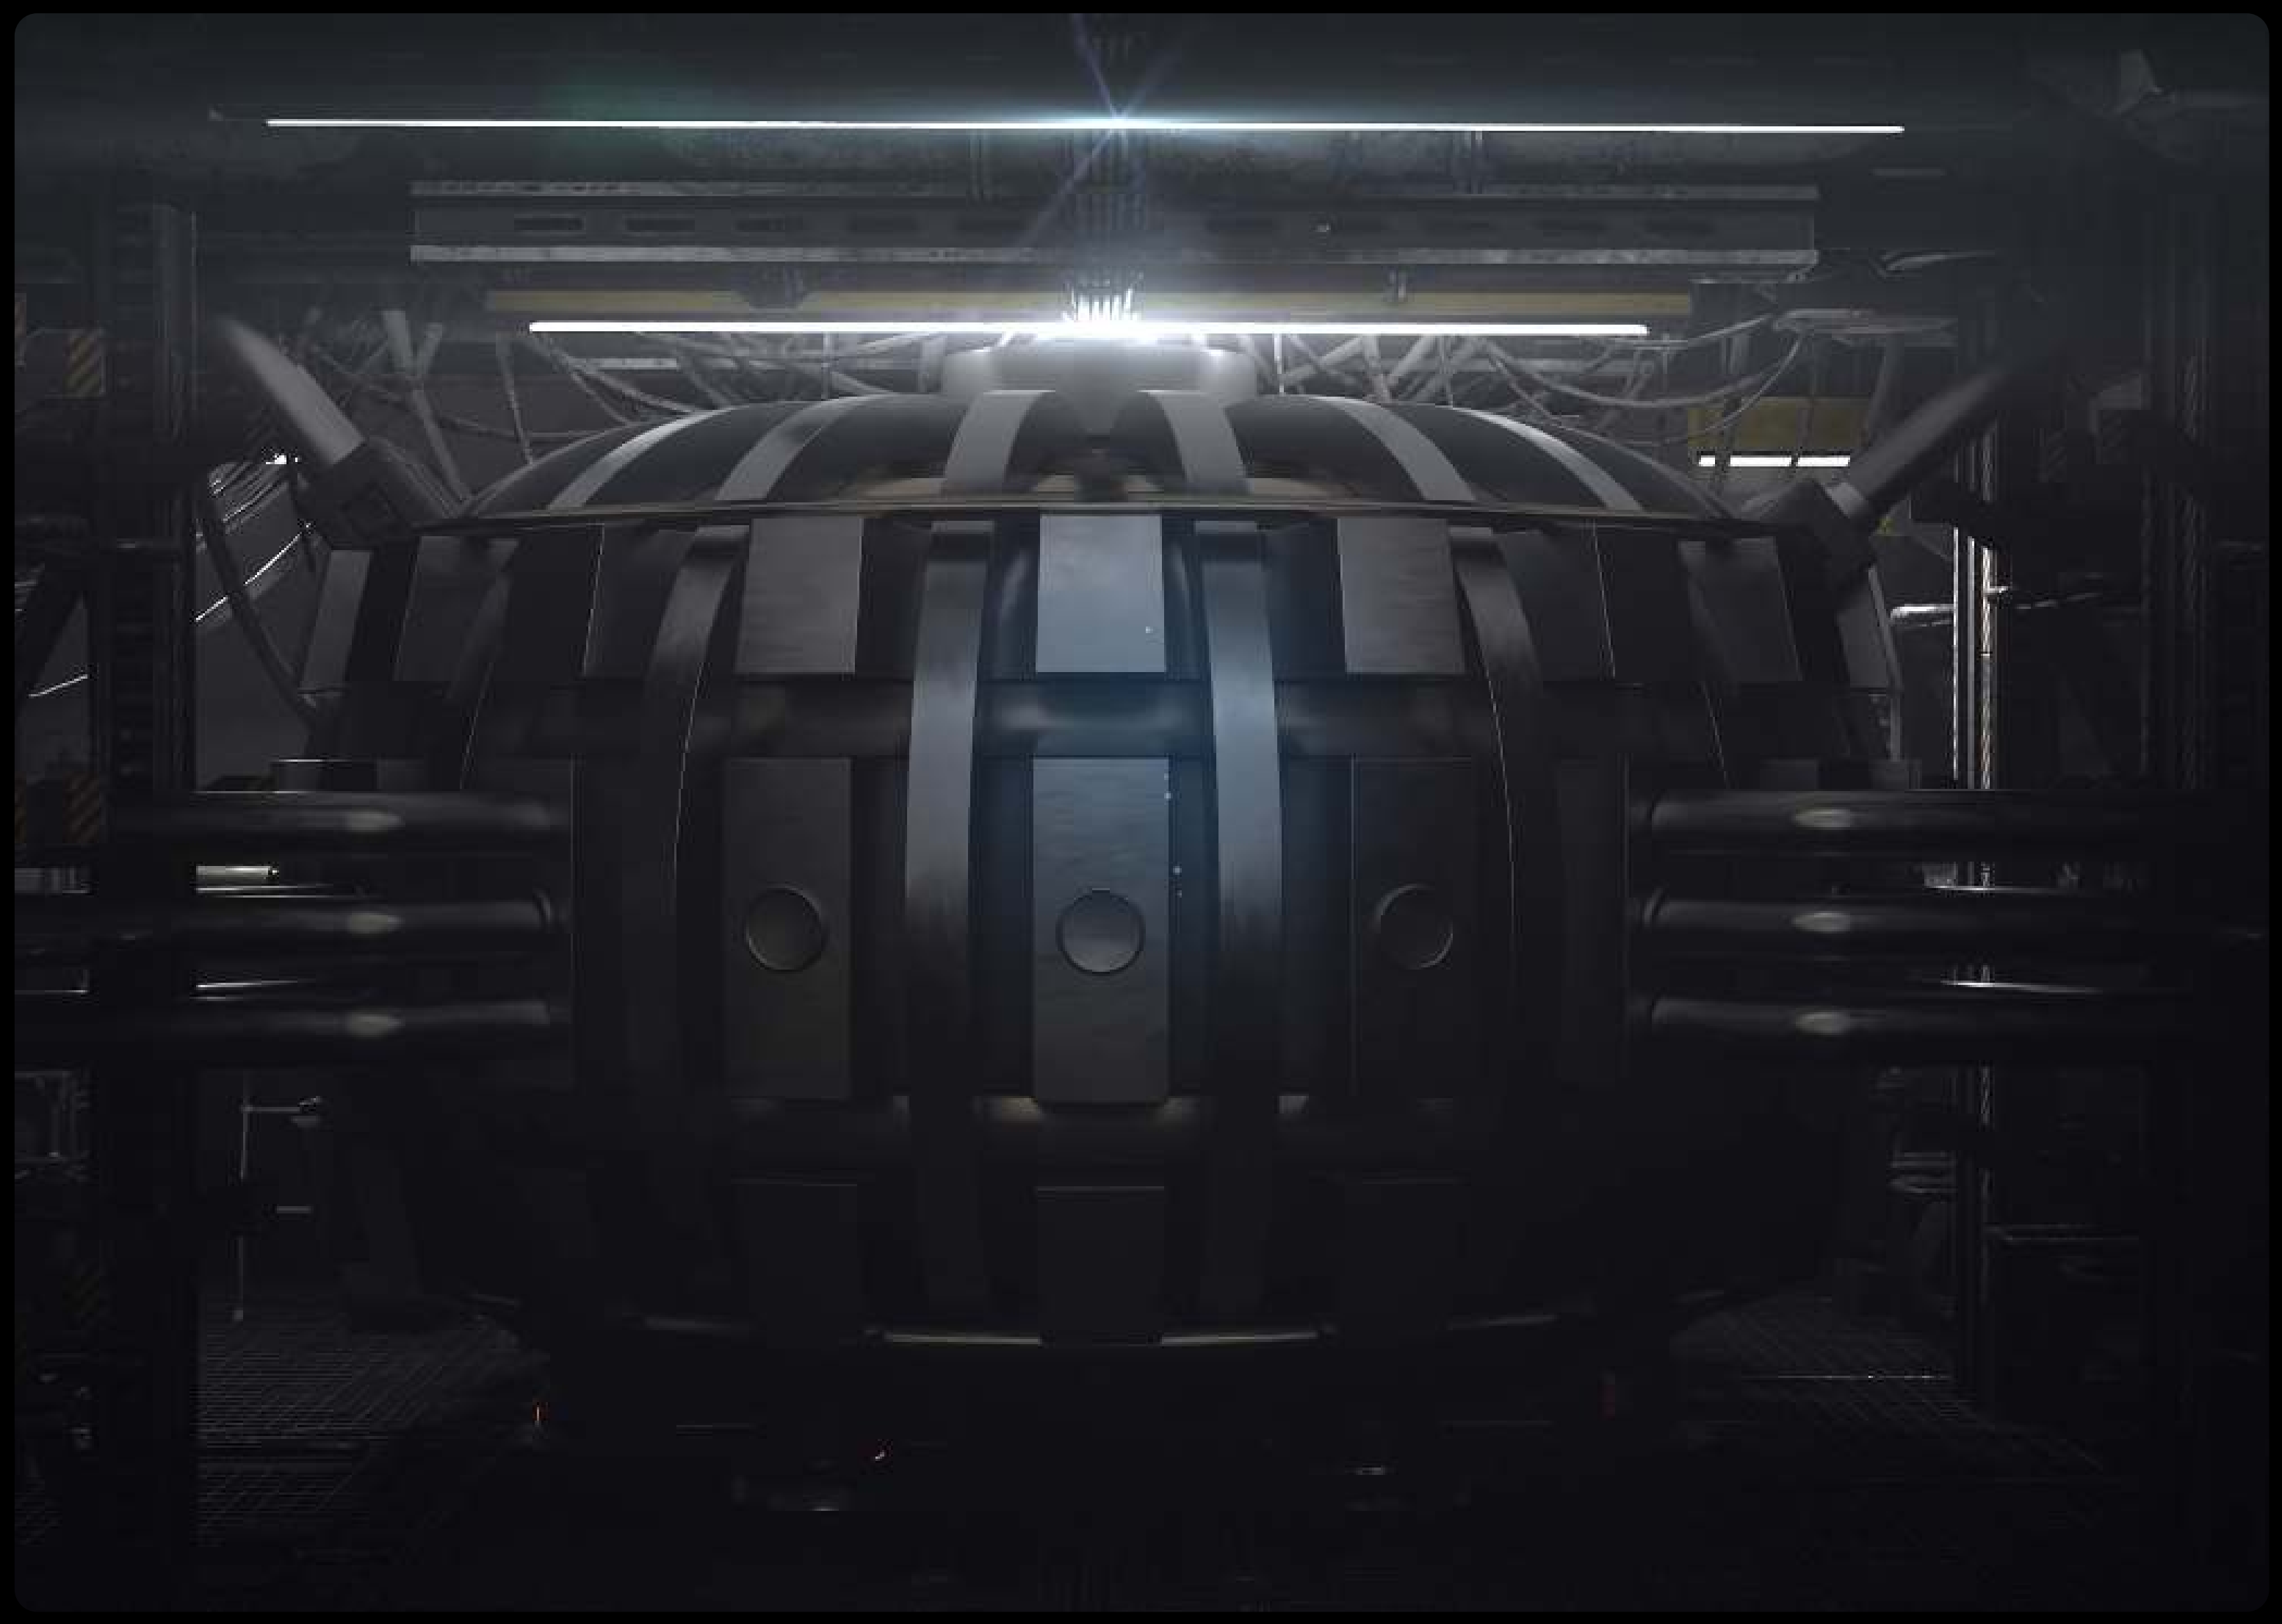

The pioneers of fusion will define the future energy landscape for generations. **Maritime Fusion isn't just watching—we're building it**. This is more than an engineering race; it's the defining economic and technological contest of the 21st century. The moment we crack fusion, energy scarcity soon becomes a relic of the past—fueling infinite innovation, limitless computation, and a world no longer constrained by resources. This isn't just an investment in energy—it's an investment in a new industrial age. **Practicality drives our approach**, ensuring we deliver not just theoretical solutions but tangible results. Fusion ***is*** the future. **With our eyes set firmly on FOAKs today** and SOAKs and NOAKs tomorrow, we're not just preparing to change the maritime industry; we're ready to change the world.

The future belongs to those bold enough to create it.