

Design and construction of the MUSE permanent magnet stellarator

T.M. Qian^{1,†,‡}, X. Chu¹, C. Pagano¹, D. Patch¹, M.C. Zarnstorff¹,
B. Berlinger¹, D. Bishop¹, A. Chambliss¹, M. Haque¹, D. Seidita¹ and
C. Zhu^{1,‡}

¹Princeton Plasma Physics Laboratory, 100 Stellarator Road, Princeton, NJ 08540, USA

(Received 27 February 2023; revised 25 July 2023; accepted 4 August 2023)

This paper documents the design and construction of MUSE, the world's first permanent magnet (PM) stellarator and the first quasi-axisymmetric experiment. The purpose of MUSE is to develop and assess a new way of building optimised stellarators that uses simple planar coils PMs. Our PM optimisation algorithm consists of initialising a geometry to pack dipoles densely, running the FAMUS code to minimise surface field error subject to PM constraints and applying discrete jumps to reach a physically realisable solution. FAMUS treats the PM system as a set of ideal point dipoles. From there we construct finite-volume magnet towers to be housed in 3D-printed PM holders. We describe the design of the PM holders, which were validated by laser metrology. We analyse the effects of finite permeability, sensitivity to perturbations and magnetostatic forces. An exact analytic formula for the magnetic field from a finite-volume PM tower is presented to compute PM–PM forces and stress on the PM holder. Stellarator construction is complete and experiments are underway.

Key words: plasma confinement, fusion plasma, plasma devices

1. Introduction

Stellarators offer many advantages to magnetic confinement for fusion energy (Boozer 2005; Helander 2014; Imbert-Gerard, Paul & Wright 2020). They are inherently steady state, operate without externally driven plasma current and can be optimised for good magnetohydrodynamic, neoclassical and turbulent plasma confinement (Weller *et al.* 2006; Canik *et al.* 2007; Proll *et al.* 2015; Beidler *et al.* 2021). Compared with other fusion systems, these benefits come at the cost of complexity in engineering design. The magnetic field coils, vacuum vessel (VV) and fusion blanket of a typical stellarator are shaped as three-dimensional (3D) objects. This paper addresses coil complexity by describing in detail how permanent magnets (PMs) can be used in conjunction with simple toroidal field (TF) coils to create a highly optimised stellarator magnetic field.

[†] Email address for correspondence: tqian@pppl.gov

[‡] Present address: University of Science and Technology of China, Hefei 230026, PR China.

PMs were analysed as a method to create optimised stellarator geometries without complex 3D coils by Helander *et al.* (2020). This approach is similar to earlier strategies using arrays of simple window-pane coils (Ku & Boozer 2009) or arrays of superconducting pucks (Bromberg *et al.* 2011). Following the initial paper describing the use of PMs for stellarator shaping, there have been several papers describing methods to determine magnet distributions for optimised stellarator configurations. (Zhu *et al.* 2020b) started from a surface current solution to calculate a single-layer magnetisation in the perpendicular direction and discretised it into a volume distribution. Landreman & Zhu (2021) looked at PMs as if they were a magnetic fluid surrounding the plasma. The thickness varied continuously, as did the direction of magnetisation. This is useful for scoping the thickness of magnetic material needed to realise a certain size stellarator. However, a practical PM array must use a finite set of discrete magnet types. The PM4STELL study achieved a solution that uses only 3 types of cubes, which can be positioned into 56 orientations (Hammond *et al.* 2022; Zhu *et al.* 2022). Lu *et al.* (2021a) used just one cube that is arranged in six orientations. Since those two studies targeted different plasma equilibria, a direct comparison between algorithms is not readily available. Recently, Kaptanoglu introduced a relax-and-split algorithm (Kaptanoglu *et al.* 2022) and a greedy algorithm (Kaptanoglu, Conlin & Landreman 2023) that improved the optimisation process itself. These approaches are agnostic to the specific geometry and can be applied to any target plasma or PM set.

Prior to this work, there were optimised stellarators designed with complex modular coils: HSX (Anderson *et al.* 1995, 2006), W7X (Nührenberg *et al.* 1995; Wolf *et al.* 2017), CFQS (Liu *et al.* 2018; Shimizu *et al.* 2018) and NCSX (Zarnstorff *et al.* 2001; Neilson *et al.* 2010). Stellarators were also built with geometrically simple coils that are not optimised for good plasma confinement: LHD (Iiyoshi *et al.* 1999; Yamada 2011), TJ-II (Alejdre *et al.* 1990; Hidalgo *et al.* 2005), Uragan-2M (Pavlichenko *et al.* 1993; Beletskii *et al.* 2008), H-1NF (Hamberger *et al.* 1990; Kumar, Blackwell & Harris 2009) and Model C (Young 1974).

This paper describes the optimisation, design, analysis and construction of MUSE, the first stellarator to test the PM approach.

- (i) Section 2 describes PM optimisation. We start with initialising an ideal dipole array that geometrically fits outside the plasma and VV but inside circular-bore TF coils.
- (ii) Section 3 describes the design of the PM holder, which is 3D printed to precisely position the PMs and to resolve magnetic forces.
- (iii) Section 4 analyses the effect of finite magnetic permeability, the sensitivity to magnetisation and the near-field forces that arise from finite-volume magnets.
- (iv) Section 5 documents PM assembly methods used in the construction of MUSE.
- (v) Section 6 discusses potential extensions of the PM method.
- (vi) Section 7 summarises the main conclusions.
- (vii) Appendix A presents an exact-analytic formula to calculate the field everywhere in space from a 3D rectangular magnet.

1.1. Stellarator equilibrium

The focus of this paper is to document first-of-a-kind techniques for the design and construction of PM stellarators. The optimisation of the stellarator equilibrium has been described in an earlier publication (Qian *et al.* 2022). For completeness, we briefly outline the method and results.

The equilibrium is computed by VMEC (Hirshman & Whitson 1983), and the optimisation was done using STELLOPT (Lazerson *et al.* 2020). We designed the plasma

to use TF coils that were readily available at PPPL and a VV assembled from 90° glass elbows (described in § 5.3 on Torus Assembly). This constrained the plasma to fit within a circular cross-section torus ($R = 30$ cm, $a = 7.5$ cm). MUSE has two field periods, on-axis magnetic field $B = 1.5$ kG and in-vessel aspect ratio of 7.9. It is a low- β vacuum equilibrium. The rotational transform profile increases from just above 0.18 on the magnetic axis to just below 0.20 at the edge (see the ι profile in figure 12). This range is defined to avoid the 2/11 and 2/10 magnetic islands (where n/m means n toroidal transits for m poloidal transits before the island chain repeats). Within the boundary of the VV, there are good nested flux surfaces (see Poincaré plots in figure 9). The effective ripple $\epsilon_{\text{eff}}^{3/2}$, a confinement metric that is proportional to neoclassical particle flux (Nemov *et al.* 1999), is more than 100 times lower than any existing stellarator experiment. MUSE is also optimised to be quasi-axisymmetric (Garabedian 1996). Figure 1 extends the comparison plot in our earlier paper to include both free boundary and fixed boundary metrics. It includes the ‘precise QA’ case from (Landreman & Paul 2022; Wechsung *et al.* 2022) and a ‘compact QA’ case from adjoint-optimisation (Nies *et al.* 2022; Nies & Paul 2023). These were both published after the MUSE equilibrium was already selected. In principle, they could be targeted for future MUSE-like PM designs.

The shape of the stellarator equilibrium influences the positioning of ports, which constrains the PM design by disallowing magnets from certain locations, as discussed in § 2. The stellarator equilibrium did not influence the choice of heating. A tabletop device such as MUSE has many available heating options, including microwave source electron cyclotron heating (ECH), helicon antenna ion heating, hollow cathode discharge and arc discharge. Since we did not constrain the stellarator equilibrium or PM optimisation to fit requirements for any particular heating scheme, we defer description of heating to the discussion of initial experiments in § 6.

Data files for MUSE are publicly available on GitHub (Qian 2023).

2. PM optimisation

This section describes PM optimisation, assuming that a target stellarator equilibrium has already been found. For this design, we used rare-earth ‘N52’ (neodymium–iron–boron) PMs. The choice and properties are described in our earlier paper (Qian *et al.* 2022). We purchased commercially available magnets and no modifications were needed for construction.

The MUSE PM design procedure consisted of three steps: geometry selection, continuous first pass and discretising refinement.

2.1. Geometry

The first step is to construct a geometry for the PM array. For this section, we treat each PM as an ideal dipole point source. An ideal dipole has six degrees of freedom: three for position, two for spherical orientation and one for magnitude. Our optimisation fixes magnet positions for simplicity. If the position were a free variable, the optimiser would need to constrain neighbouring magnets from intersecting as they move. By fixing the positions, we can initialise an array without intersections and optimise only the magnetisation. To define the initial point cloud of dipoles without intersections, it is necessary to know the 3D build and orientation of the finite-volume magnets.

Some authors choose to align the magnets with a cylindrical grid (Hammond *et al.* 2022). Others use a cartesian coordinate system (Lu *et al.* 2021a). For both cases, the coordinates do not align with a toroidal plasma, and multiple orientations are required

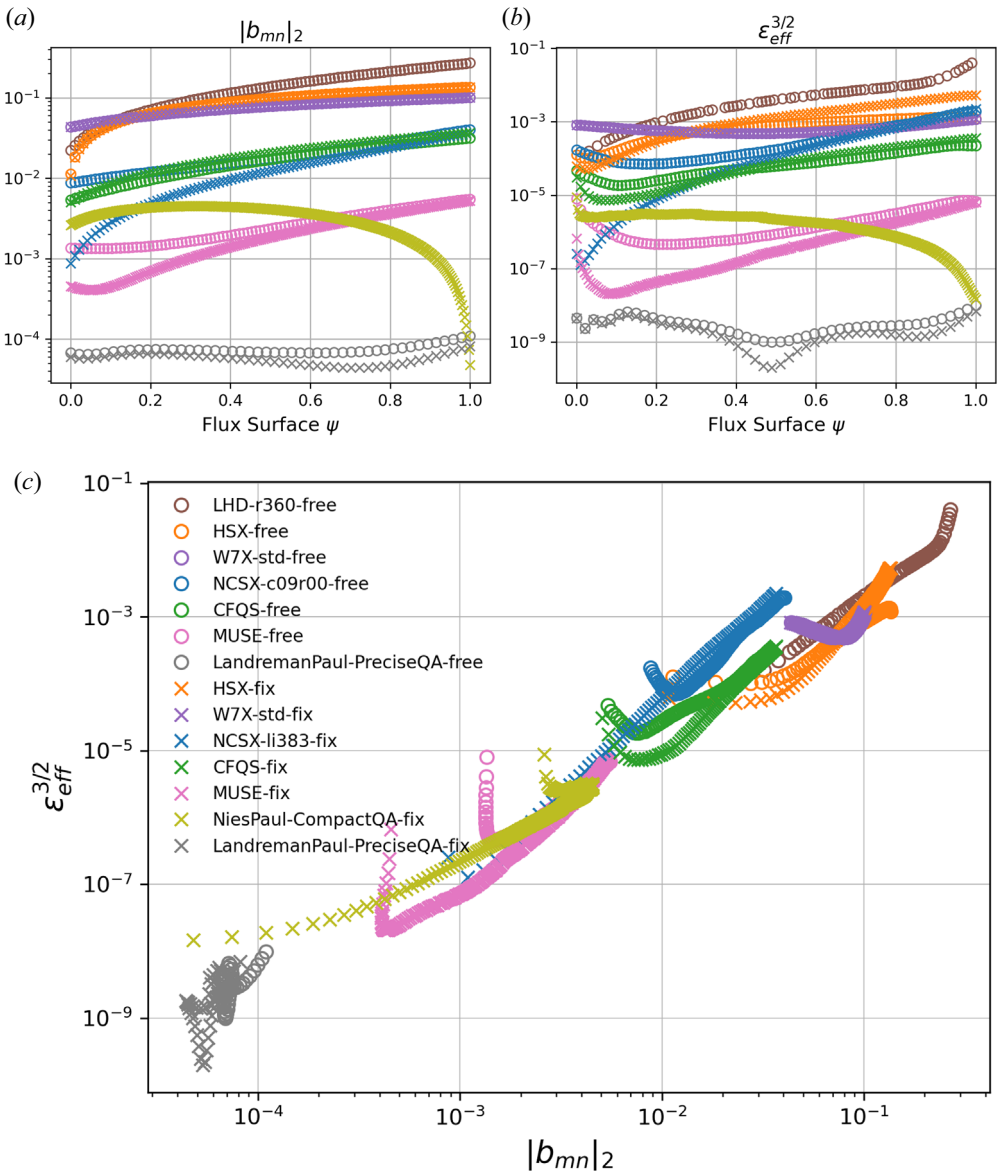


FIGURE 1. (a,b) Quasisymmetry metric $|b_{mn}|_2$ and neoclassical transport metric $\epsilon_{\text{eff}}^{3/2}$, which are computed by STELLOPT. The flux surfaces are indexed by normalised poloidal flux ψ . (c) Parametrisation of both metrics as a function of radius. Free boundary VMEC equilibria are noted by (o) and fixed boundary equilibria by (x). The stellarators are coupled by colour. For most stellarators the fixed boundary is somewhat better than free boundary. LHD, HSX, W7X and MUSE have been completed. NCSX and CFQS are partially constructed. For NCSX ‘c09r00’ is a specific coil design for the ‘li383’ equilibrium.

to facilitate shaping. The MUSE magnets are chosen to align with a toroidal grid. All dipoles are aligned normal to a circular cross-section torus concentric with the VV. The magnetisation vector points either straight in or straight out, and this information is stored

as a signed magnitude ρ such that $\mathbf{m} = \rho \mathbf{m}_0$.¹ Since the orientation and position are fixed by geometry, the optimiser manipulates only a scalar magnitude for each magnet.

It is desirable to pack the magnetisation volume as densely as possible. There are more magnets on the outboard side than on the inboard side due to toroidal curvature. Accounting for this, we divide the magnet volume into N_θ poloidal rings such that each ring has variable $N_\zeta = N_\zeta(\theta)$ magnet positions. At each surface coordinate (ζ, θ) we initialise a magnet tower of N_s radial slices. A ‘PM tower’ refers to the stack of magnets at a single (θ, ζ) position. In the context of optimisation, every tower consists of N_s slices (see figure 4). In the context of assembly, however, the slices are merged into a single physical piece for each tower (see § 5). Note that $N_\zeta(\theta)$ is largest for the outer ring ($\theta = 0$) and smallest for the smallest ring ($\theta = \pi$). Figure 2 shows the non-smooth function of the poloidal angle that arises from packing the magnets as densely as possible. This discrete function takes into account the circumference of a poloidal ring $C(\theta) = 2\pi(R + a \sin \theta)$, the width of a magnet w , and the minimum structural gap which appears between magnets d_m . We must also take into account an assembly gap between support structures d_s . Since MUSE has two field periods, the stellarator is divided into four identical half periods. Let us define the boundaries of these ‘physical half periods’ to be our coordinate system’s x and y axes. However, for the practical purpose of force balance, we will find it useful to let the PM structures straddle the half-period boundary symmetrically. As a result, the ‘PM holder quadrants’ are shifted 45° out of phase, as shown in figure 3. This is also a useful feature for quality checks during assembly because the physical parts exhibit odd symmetry (see figure 19). We emphasise that the information in a two-field-period stellarator is entirely contained in a 90° half period bounded between symmetry planes. We intentionally bisect the half-period into two pairs of distinct PM holders, each centred on one symmetry plane. As a consequence, the appropriate length for toroidal discretisation is actually an eighth of the circumference, $\Delta = \frac{1}{2}C(\theta)/(2N_p) = C/8$. Then the expression for the number of toroidal positions in a poloidal ring is a floor function

$$N_\zeta(\theta) = \left\lfloor \frac{\Delta - d_s}{w + d_m} \right\rfloor. \quad (2.1)$$

This guarantees a gap of at least d_m between every magnet. Due to toroidal curvature, the minimum minor radius a differs for the inner and outer poloidal angles. On the inside the densest ring is determined by the base of the tower; on the outside it is determined by the top of the tower. Figure 3 highlights this feature, and it shows the relative size between the PM grid and the plasma shape. In the final design, $d_m = 1$ mm is the minimum gap between all magnet edges and $d_s = 2d_m$. Because the space d_s is removed from both sides, the physical gap between PM holders is at least 4 mm.

To choose a specific geometry one should consider the total magnetisation thickness and the size of individual magnets. Giving the optimiser a larger region to place magnets expands the solution space, at the cost of larger assembly forces and more magnetic material. Larger footprints for each magnet tower yield fewer total parts, at the cost of fewer degrees of freedom and more non-magnetic gaps in the PM holder. The PM geometry is radially discretised into identical layers by splitting each magnet tower into unit slices. For a fixed (ζ, θ) grid, one can propose multiple radial discretisations based

¹The magnetisation M (A m^{-1}) is an intrinsic material property. When one purchases an N52 grade rare-earth magnet, they are purchasing a high value of M . The dipole moment $m_0 = MV$ is derived from the grid size, where V is the volume of each physical magnet. For the fundamental magnet slice in MUSE, our data files used $m_0 = 0.074625 \text{ A m}^2$, the volume of each slice V is precisely $1/256$ cubic inches (approximately 64 mL) and the modelled magnetisation is $M = 1.165 \times 10^6 \text{ A m}^{-1}$.

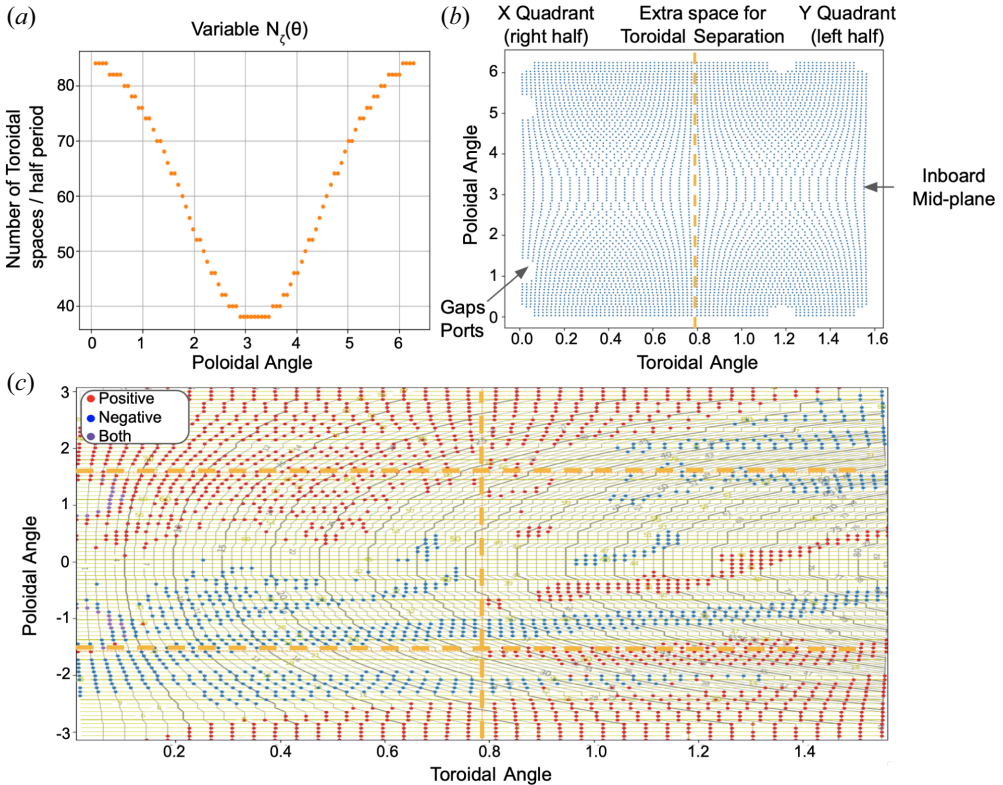


FIGURE 2. Illustrations of the distribution of magnet positions and actual magnets. (a) The number of toroidal magnet positions as a function of poloidal angle $N_z = N_z(\theta)$. (b) The full 2D grid of magnet positions as a function of toroidal and poloidal angles (ζ, θ), including the assembly gap and port holes. (c) The distribution of final MUSE magnets. Red and blue dots mark towers that have only positive or only negatively oriented magnets. The purple dots mark special locations where a single tower has both orientations. These are collected near the vertical port holes at toroidal angle $\zeta = 0$. An orthogonal coordinate grid is indicated in gold and grey contours. It is used to uniquely index each magnet position. Every 5 lines are labelled and bolded to facilitate counting. Dashed orange lines show the $\zeta = \pi/2$ boundary that parts the X and Y quadrants and the $\theta = \pm\pi/2$ boundary that approximately parts the inboard and outboard subsegments. Here, the poloidal axis is shifted to $(-\pi, \pi)$ such that the densest region (the outboard midplane) is centred in the middle. Readers may wish to contrast this representation with alternatives in figures 8 and 10.

on the minimum radius a_0 , the maximum radius a_N and the number of slices N_s . From this the slice thickness is determined, $\Delta a = (a_N - a_0)/N_s$. It is desirable to bring the tower base a_0 as close to the plasma vessel as possible, but a limit is set by the accuracy of the VV's radial extent, the need to leave an assembly gap and the desire to leave space for radiofrequency (RF) antenna straps. The top of the tower a_N is constrained by the inner bore of the TF coils. We considered several geometries using 12–20 layers of $\Delta a \approx 1.6$ mm (to be precise, we specified in customary units $1/16$ inch = 1.5875 mm). We considered cases where the fixed thickness magnetisation layer was shifted closer and farther from the plasma within a region that spans 92–124 mm radially. Ultimately, we selected $a_0 = 100$ mm and $N_s = 14$. In the numeric optimiser, the effect of radial displacement can be achieved by setting a lower magnet to 0 and an upper magnet to 1.

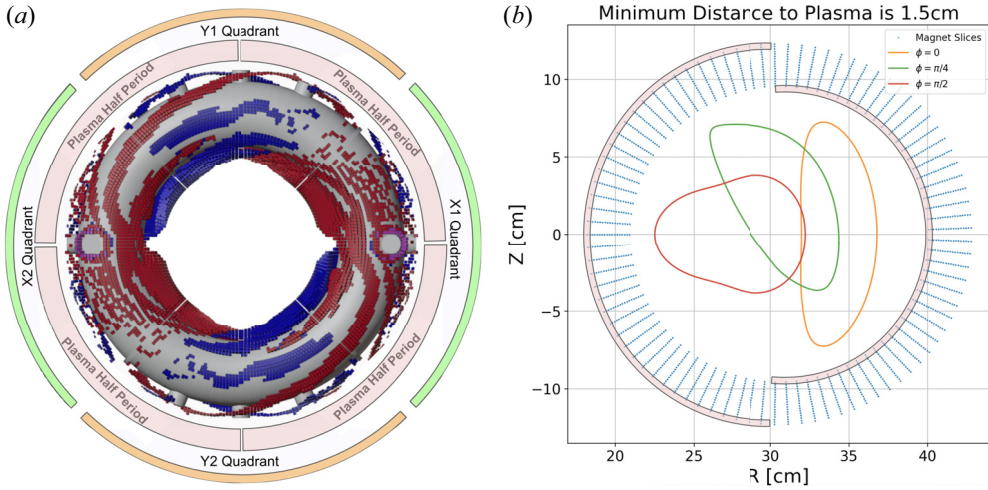


FIGURE 3. (a) Top-down view of MUSE shows the PMs rendered as 3D blocks surrounding the VV. This plot highlights four identical physics half periods, surrounded by two pairs of identical PM holder quadrants, which each mirror a quarter period. (b) A constant toroidal angle cut, visualising the distance between plasma surfaces and the PM geometry. The translucent half-circles indicate the location of the magnet with tightest toroidal curvature, which is used to determine toroidal grid spacing in the magnet geometry.

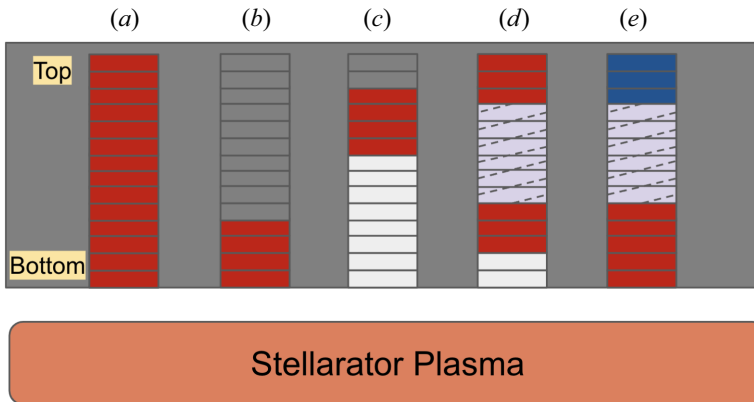


FIGURE 4. Five types of magnet towers. It can be completely filled (a) or partially filled. This gives the effect of magnitude variation. Partially filled magnets can be flush against the PM holder material (b), or they can be recessed with an air gap after assembly (c). The difference between (b) and (c) gives the effect of position variation within one PM tower. At rare locations (about 30 out of 10 000), the optimiser decided that it is best to leave a gap between magnet slices (d) and (e). Undoing the gap was detrimental to our plasma solution. The gap can separate magnets of the same sign (d) or opposite sign (e). This is shown with purple shading and diagonal dashed lines. At construction, these gaps were filled by annular non-magnetic plastic spacers.

Figure 4 illustrates this and other possibilities for a PM tower. Since the minimum distance between PM and plasma surface is 15 mm, and the dipole thickness is 1.6 mm, the plasma gap is approximately 10 times the magnet slice thickness, which justifies the ideal dipole approximation.

In MUSE, the PMs are constrained to fit outside the VV ($r_{\text{outer}} = 85$ mm) and inside the circular-bore coils ($r_{\text{inner}} = 127$ mm). This leaves a radial thickness of 42 mm. We must also take into account the thickness of the PM holder material (see § 3), the imperfect shape of our glass torus and necessary assembly gaps. As a result, our final magnet thickness is approximately 22 mm. There is space left on each side, part of which we used to insert a thin helicon wave antenna between the VV and PM holder. An area for future improvement could be concatenating neighbouring footprints in regions of strong magnetisation (see figure 2). In addition, it would be desirable to define the poloidal rings to include an assembly gap for poloidal subsegments. We did not do this and had to draw an *ad hoc* jagged partition in the helical space between magnetisation bands instead (further discussed in § 3).

2.2. Continuous optimisation

With the geometry fixed, we run the first pass continuous magnet optimisation for a target plasma equilibrium. We use the FAMUS code (Zhu *et al.* 2020a) which takes the following inputs: our magnet grid, a target plasma boundary and a set of TF coils. FAMUS is the PM optimisation branch of the FOCUS code for flexible coil optimisation (Zhu *et al.* 2018). A quasi-Newton solver minimises the objective function

$$f_B = \frac{1}{2} \int_S [(\mathbf{B}_{\text{PM}} + \mathbf{B}_{\text{TF}}) \cdot \mathbf{n}]^2 dA, \quad (2.2)$$

where \mathbf{n} is a unit vector normal to the plasma surface. This is a surface integral of the residual normal field from PM and TF coils on the plasma boundary. Since MUSE is a vacuum equilibrium, the contribution from plasma currents is not considered.

The objective is to find arrangements of dipole magnitudes that best produce the target equilibria, subject to magnet constraints. Because our geometry has just one orientation, we define the dipole moment to be everywhere parallel to the normal vector of a circular-cross section torus. The signed magnetisation is represented by a normalised scalar $\rho \in [-1, 1]$. The optimiser allows ρ to take decimal values in the continuous pass, but in the final configuration only integer values ($-1, 0, 1$) are selected.

The target metric (2.2) is summed over the plasma surface. If there are N dipoles and M plasma surface elements, this makes an $N \times M$ matrix. Because the orientation is fixed, we can greatly improve the speed of iterated calculations by precomputing an induction matrix

$$G_{ij} = \frac{\mu_0}{4\pi} \left[3 \left(\frac{\mathbf{r}_{ij} \cdot \mathbf{n}_i}{|\mathbf{r}_{ij}|^5} \right) \mathbf{r}_{ij} - \frac{\mathbf{n}_i}{|\mathbf{r}_{ij}|^3} \right] \cdot \mathbf{n}_j. \quad (2.3)$$

In this expression i indexes plasma surface elements, while j indexes magnetic dipoles. Thus, \mathbf{n}_i is the normal unit vector on each of N plasma surface elements, \mathbf{n}_j is the normal unit vector from each of M dipole moments and $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$ is the pairwise position vector from magnetic dipole to plasma surface elements. This is particularly useful for MUSE because both sets of unit vectors ($\mathbf{n}_i, \mathbf{n}_j$) are fixed. The distances and dot products can be precomputed once for the entire optimisation. Only the set of scalars $\{\rho_j\}$ are updated such that the local field $(\mathbf{B} \cdot \mathbf{n})_i = G_{ij}\rho_j$. This is analogous to electromagnetic induction in the sense that the matrix is determined entirely by geometry, but there is no actual induction. Compared with the similar quantity defined in Zhu *et al.* (2020a), this induction matrix is dotted into fixed dipole orientations. This simplifies computations in the ‘auto-zot’ algorithm described in Section 2.3.

The space of PM solutions is mathematically ill-posed. Very different input magnet arrays can give the same plasma field output. Figure 5 shows an example of this degeneracy

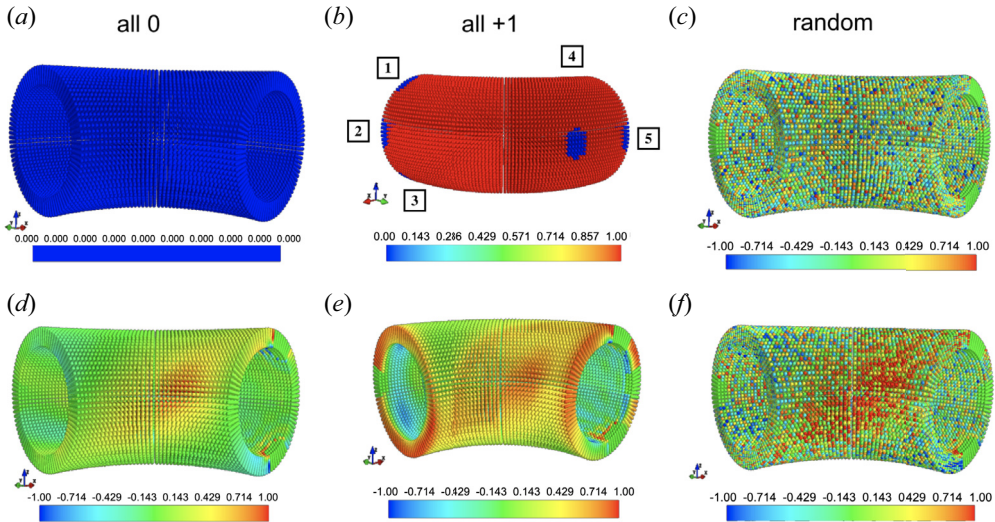


FIGURE 5. Mathematical ill-posedness of PM optimisation. (a,d), (b,e), (c,f) An initial seed and a solution converged to low f_B . Very different solutions all give the similarly low error, therefore the problem is under-constrained. In plot (b), we have highlighted five regions that are excluded from the optimisation for diagnostic ports. Both the initial and final states fix the magnetisation in these regions to zero. When reflected through stellarator symmetry, these five regions become 3 vertical ports and 3 horizontal ports per period, for 12 total port locations (cf. full torus photo in figure 19). Note that zero is blue for the plots (a,b) where the colour bar ranges from [0, 1]. In contrast, zero is green for the plots (c-f) where the colour bar ranges from [-1, 1].

starting from three initial conditions: all 0, all 1 and random. Each subplot shows a half period of magnet geometry in three dimensions. We see that three very different solutions all give ‘optimised’ plasma metrics while some are clearly suboptimal in their use of magnets. This motivates adding target metrics for the total volume (an L-1 norm) of PMs

$$f_V = V \sum |\rho| \tag{2.4}$$

and the elimination of intermediary values, such that only ‘binary’ dipoles $\rho = \{-1, 0, 1\}$ instead of the continuous distribution $(-1, 1)$ remain

$$f_D = \sum |\rho|(1 - |\rho|). \tag{2.5}$$

We use these to find a Pareto frontier between metrics (f_B, f_V, f_D), by comparing differently weighted optimisations. Although the latter two metrics are not differentiable at $\rho = 0$, we did not find convergence issues arising from solutions bouncing at the cusp. Instead, the magnets converged to intermediary solutions between whole numbers to satisfy f_B . Figure 5 also shows the port locations where magnets are forbidden. There are three ports that straddle the X -symmetry plane vertically and three ports that straddle the Y -symmetric plane horizontally. A photograph of the physical ports is given in figure 19.

To find solutions with precisely $-1, 0, 1$ we devised additional discrete refinement methods. Our optimisation used the Fortran code FAMUS. This algorithm has been adapted into the Python code FICUS and the optimisation framework SIMSOPT (Landreman *et al.* 2021). Making a continuous first pass is quick computationally. This stage can be used to scope the trade-off between different magnet array geometries, target

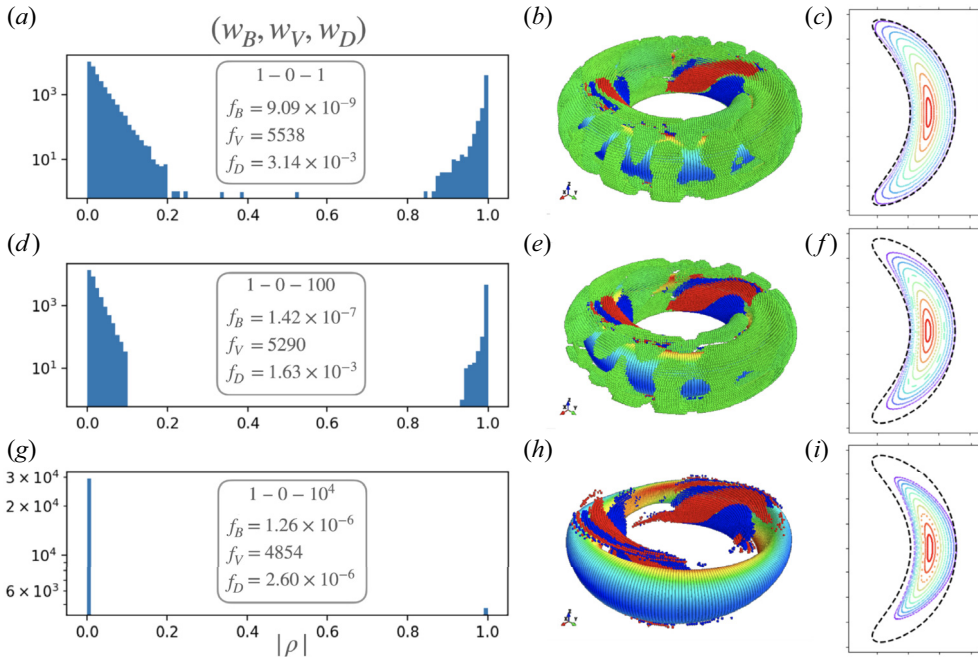


FIGURE 6. (a–c), (d–f), (g–i) A triplet of plots resulting from a specific set of relative weights (w_B, w_V, w_D) . The legend reports the resulting magnetic field error f_B ($\text{T}^2 \text{m}^2$), magnetised volume f_V (cm^3) and the binary dipole metric f_D (dimensionless). While the metrics retain physical units, FAMUS applies relative weights after normalising the metric by their value at the first optimisation step. The histograms show that the distribution of $|\rho|$ tends to collect at 0,1 as $f_D \rightarrow 0$. The 3D plots show the spatial distribution of the dipole set relative to a target plasma surface, with dipoles $|\rho| < 10^{-3}$ suppressed. The colouring is green for 0, blue for -1 and red for $+1$ (see the legend in figure 5). The Poincaré plots trace field lines at the $\zeta = 0$ toroidal cross section. Lowering f_D trades off against maintaining good flux surfaces, which is quantified by low f_B . Let us note that the figure shown is not the final MUSE equilibrium, but an alternative NCSX-based target plasma candidate.

stellarator equilibria, magnetic field accuracy and PM volume. The difficulty is taking discrete steps to set $f_D = 0$ while maintaining a tolerably low f_B .

2.3. Discrete refinement

Discrete refinement uses multiple iterations to force the normalised dipole moment to be exactly -1 , 0 or 1 while preserving the target field optimisation. Achieving $f_D = 0$ is necessary for constructing a set of physical magnets because $|\rho| = 1$ makes all magnet slices identical. We found, however, that it could not be achieved by a continuous optimiser. This is illustrated by the Poincaré plots in figure 6. Even starting from a fairly low value of $f_D \sim 10^{-4}$, naïve jumps to 0 yield intolerable degradation of f_B .

We address this problem by defining two discrete jump operators: (1) a COMPRESS step that sums the magnitudes of all dipoles of each tower and redistributes them into a stack of integer magnetisations plus a single non-integer residual; (2) a ROUND step that forces all dipoles below a threshold $|\rho| < \rho_*$ to 0 while setting all others to either 1 or -1 . Both make a discontinuous jump in solution space. Different choices of rounding threshold ρ_* , each forcing f_D to be exactly 0, lead to different neighbourhoods of local minima.

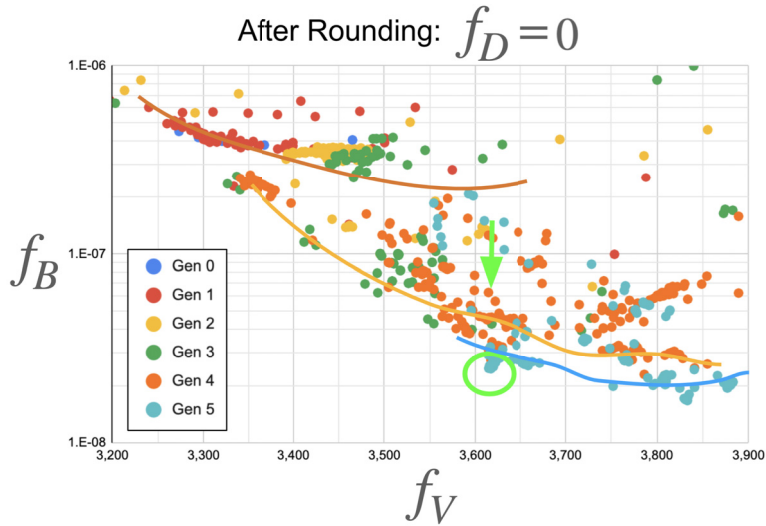


FIGURE 7. Generations of f_B are shown to progressively decrease while maintaining low f_V , subject to the constraint that $f_D = 0$ after each discrete refinement round.

The continuous optimiser is rerun to find the next generation of optimised solutions. A detail of the continuous optimiser is that dipoles get ‘frozen’ when $\rho = 0$ because $\partial_\rho f_B \propto \rho$. Vanishing magnitude implies a vanishing gradient and no change in optimisation. To reset the continuous optimiser, a small residue of 10^{-4} was added to all zero cells after compression, with exceptions for locations where the optimiser is intentionally disabled (e.g. port locations and unused geometry layers). In such locations, $\rho = 0$ is set as a feature to permanently disable magnetisation.

Iterating between cycles of compressing and rounding followed by continuous optimisation improves the minimum f_B when $f_D = 0$ over generations. This is reminiscent of annealing metal over repeated heat cycles. For each generation, we sampled multiple values of rounding threshold ρ_* prior to compressing. Figure 7 shows how the best result improves across generations. The final configuration achieves $f_B = 1.29 \times 10^{-8}$ (Tm)², $f_V = 3.78$ litres total magnetised volume and $f_D = 0$. To reach this, one more type of discrete refinement was used, which we dubbed ‘zotting’.

A ‘zot’ is a discrete change where the magnitude of a single magnet slice ρ_j is set by hand. The surface integral f_B is recomputed as a proxy to evaluate the effect on the plasma surface. We assume that lower f_B corresponds to more faithful reproductions of the target equilibrium. Zotting was used to ‘annihilate’ neighbouring pairs that have opposite signs. In many cases, we found that zotting pairs decreased the integral f_B . This surprising observation can be explained by the fact that ROUND is a majority vote procedure. We chose the threshold that makes most magnets optimal as a group, but there remain local cases where specific magnets are suboptimal as individuals. A ‘wyrms’ is a single tower that houses magnets of both signs (named in anticipation that oppositely signed magnets in close proximity will launch out, like a Norse dragon). Several wyrms were removed completely through annihilation. Where this was not possible, we were able to guarantee

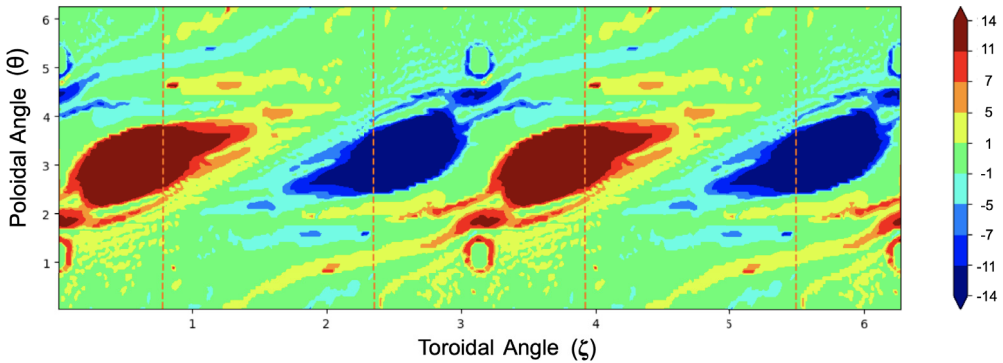


FIGURE 8. Optimised magnet distribution unfolded as a contour plot. For each toroidal and poloidal position, the slices in a magnet tower are summed to a signed integer between $(-14, 14)$. Orange vertical lines indicate the partition between PM holder quadrants. The vertical ports leave a visible ring of magnets at toroidal angles $\zeta = 0, \pi$. Unlike figures 2 and 10, this poloidal axis sets $\theta = \pi$ in the middle to emphasise the region of thickest magnetisation, which appears near the inboard mid-plane. The $\zeta = \pi/4$ cross section from figure 9 falls on the orange dashed line, whereas the $\zeta = 0, \pi/2$ symmetry planes fall in between.

at least three magnet thicknesses between oppositely signed regions for all towers.² Zotting was also used for enlarging the unmagnetised region around ports, cutting a path for the assembly joint between poloidal subsegments of the PM holder (see § 3), and annihilating pairs across neighbouring towers. The final MUSE configuration is called ‘zot-80’ because it is the 80th hand-picked permutation. A ‘gap’ is a wyrm without sign flip (see cases (d,e) in figure 4). There are 24 wyrms and 7 gaps per half period in the zot-80 configuration.

After the design of MUSE, an algorithm called ‘Auto-Zot’ was explored to systematically test every possible magnet after a compress and round. This is a once-through non-accumulating search, where each of the M magnets is perturbed independently relative to a base case. The ensemble of beneficial changes is combined and accepted together. The set of independently beneficial changes does not equal the superposition of those changes. An additional annealing loop is added to relax into a self-consistent solution. This approach is partially inspired by the MASTER algorithm from Lu *et al.* (2021b).

Figures 8 and 9 shows the results of optimisation at three constant toroidal angle cross sections. The final configuration uses 12 674 ideal dipoles per half period. This corresponds to 3.2451 of magnetic material in the full torus. The solution uses about 12.7% of the total volume available to the optimiser. To simplify the design, we concatenate the 50 696 total dipole slices into 9736 magnet towers by merging contiguous slices. While the dipole slices are all one fundamental unit, the magnet towers are distributed into 14 discrete sizes that are integer multiples of the fundamental unit. By volume, size 14 is the most populated bin, containing about a third of the total magnetic volume. By count, size 2 is the most populated bin, containing about a quarter of the magnet pieces (see figure 17).

²As a result, the magnetic pressure between oppositely signed magnets was decreased and the wyrms were pacified. None did jump out during assembly, although some of the smaller inner wyrm halves did tend to flip in the presence of other fields (see § 5.2).

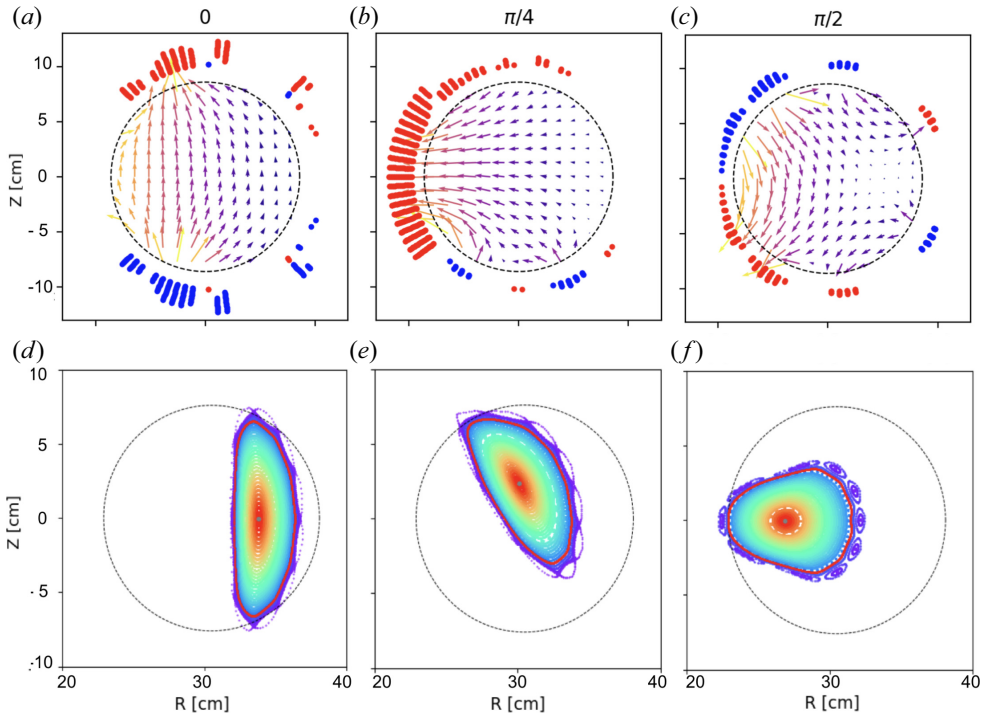


FIGURE 9. (a,d), (b,e), (c,f) Cross sections at constant toroidal angle, ζ . In (a–c), red and blue dots show the signs of slices in a PM tower. Two ‘wyrms’ can be seen at the $\zeta = 0$ plane. The quiver plot depicts the poloidal projection of the vector field \mathbf{B}_{PM} . The magnets are up–down symmetric at symmetry planes $\zeta = 0, \pi/2$. Panels (d–f) show Poincaré plots with good nested flux surfaces. There is a 2/10 magnetic island near the edge, but it lays outside the VV as illustrated by the surface outlined in red. In both (a–c) and (d–f), the dashed black line indicates the inner boundary of the VV at minor radius $a = 7.6$ cm.

3. PM holder

The PM holder accurately positions the PMs in their design location and mechanically supports them to resolve magnetic forces. 3D printing enables complex, precise structures at decreased cost because the process is automated and facilitates rapid iteration. Additive manufacturing was selected due to the complex double-curved geometry of the toroidal PM holders. It would be difficult for machine tools to access the ID surface of the PM holder to perform this through subtractive methods (i.e. CNC machining). The PM cavities also require square holes with tight corner radii. Several subscale and full-scale prototypes were made before selecting the final version for MUSE.

The PM holder design starts with a circular cross-section torus specified by 30.5 cm major radius, 12.5 cm outer minor radius and 10.0 cm inner minor radius. From the ideal point dipoles, a set of rectangular prism blocks is generated using the programmable CAD software OpenSCAD (Kintel 2010). The dipole set is spaced by design in the toroidal and poloidal directions, but they are contiguous in the radial direction. We concatenate neighbouring slices into a solid tower to simplify the system by reducing the total number of parts. In numeric optimisation, ‘magnet’ referred to an ideal dipole that represents a $(4 \times 4 \times 1)$ slice. Going forwards ‘magnet’ refers to a $(4 \times 4 \times H)$ 3D tower, where H is an integer between 1 and 14, the maximum tower height. These dimensions are measured

in units of 1/16 inch. Each point dipole is mapped to the centre of mass of a slice within the set of magnet towers. Then the set of all towers is subtracted from our simple torus to create the PM holder. That leaves a set of square cavities in a toroidal shell.

One must choose whether the square cavities open ‘in’ towards the plasma or ‘out’ away from the plasma. In this section, we use the language ‘out, top and up’ to describe the direction towards larger minor radius and ‘in, bottom and down’ to describe the direction closer to the magnetic axis (see [figure 4](#)). The advantage of opening away from the plasma is that this puts a hard shell of 3D-printed material between the PMs and the glass VV. Should one magnet accidentally break free from its shell (this has never happened) it would move out to larger minor radii and settle on either a neighbouring magnet or a TF coil, rather than impinging the glass vessel. In contrast, there are structural advantages to opening in towards the plasma. First, there is more supporting material between the outer top of the towers than the inner base because the magnet towers in our geometry radiate out from a simple circular cross section. Second, this curvature forms wedge-shaped ‘knife edges’ between the magnet towers. It is easier to 3D print a wedge attached by its wide base than its narrow edge. Third, opening in towards the plasma relaxes the radial build by eliminating the solid layer of material beneath all magnets. The entire magnet shifts closer to the plasma, which improves optimisation efficiency and results in less total volume of magnetised material. Fourth, the optimised magnet slices tend to aggregate towards the base of the PM holder. For most of the toroidal surface, the base is flat while the top is contoured (namely, reflecting the current potential Landreman 2017). This means a PM holder with cavities opening down has less empty space and more structural material than a PM holder opening up towards larger radii. Considering these reasons, we chose to design the MUSE PM holder with cavities that open in towards the plasma.

We chose a multi-jet fusion (MJF) 3D print technology (MJF Design Handbook 2019) to manufacture the PM holders. Key considerations included print accuracy, size, cost and precision. Of the choices considered, MJF had the finest detail for small features ($75\ \mu\text{m}$), the fastest print time and the lowest cost. MJF is a powder-based method in which powder is solidified in specific locations by a printed ink that absorbs heat from an infrared lamp. This is necessary to accommodate the thin cavity walls and overhangs that arise from MUSE’s toroidal curvature. The post-printing cool-down period was extended to 60 h to minimise distortion. Compared with selective laser sintering, MJF printers use a bonding agent rather than a laser introducing less heat into the printed part, leading to less distortion during cooling. We considered nylon-12 and glass-filled-nylon-12 materials. The nylon without glass has higher tensile strength and lower tensile modulus. The tensile strength of nylon without glass is 48 MPa with a tensile modulus of 1800 MPa (HP 2019b), whereas the tensile strength of nylon with glass is 30 MPa but its tensile modulus is 2500 MPa (HP 2019a). The magnetostatic forces between PM can be computed with an exact analytic formula (see § 4.3). Using this, a detailed ANSYS calculation by the PPPL Engineering Department bounded our PM holder stresses below 10 MPa (Bishop 2021). As a result, we opted for the higher-strength, lower-modulus nylon-12 without glass. This analysis validates the existing design for the most part. A minor change was made to add braces to the thinned outboard side to relieve the calculated stress. [Figure 10](#) shows the six unique PM holder parts, how the poloidal subsegments are mapped onto the PM distribution in two dimensions and how they come together in three dimensions.

To complete the design we add the following features.

- (a) Separate the ‘X’ and ‘Y’ quadrants. The physics optimisation is done over a half period (90° for $N_{\text{period}} = 2$). The half period is bisected, with each 45° quarter reflected, to generate the two types of 90° PM holder quadrants (see [figure 3](#)). As

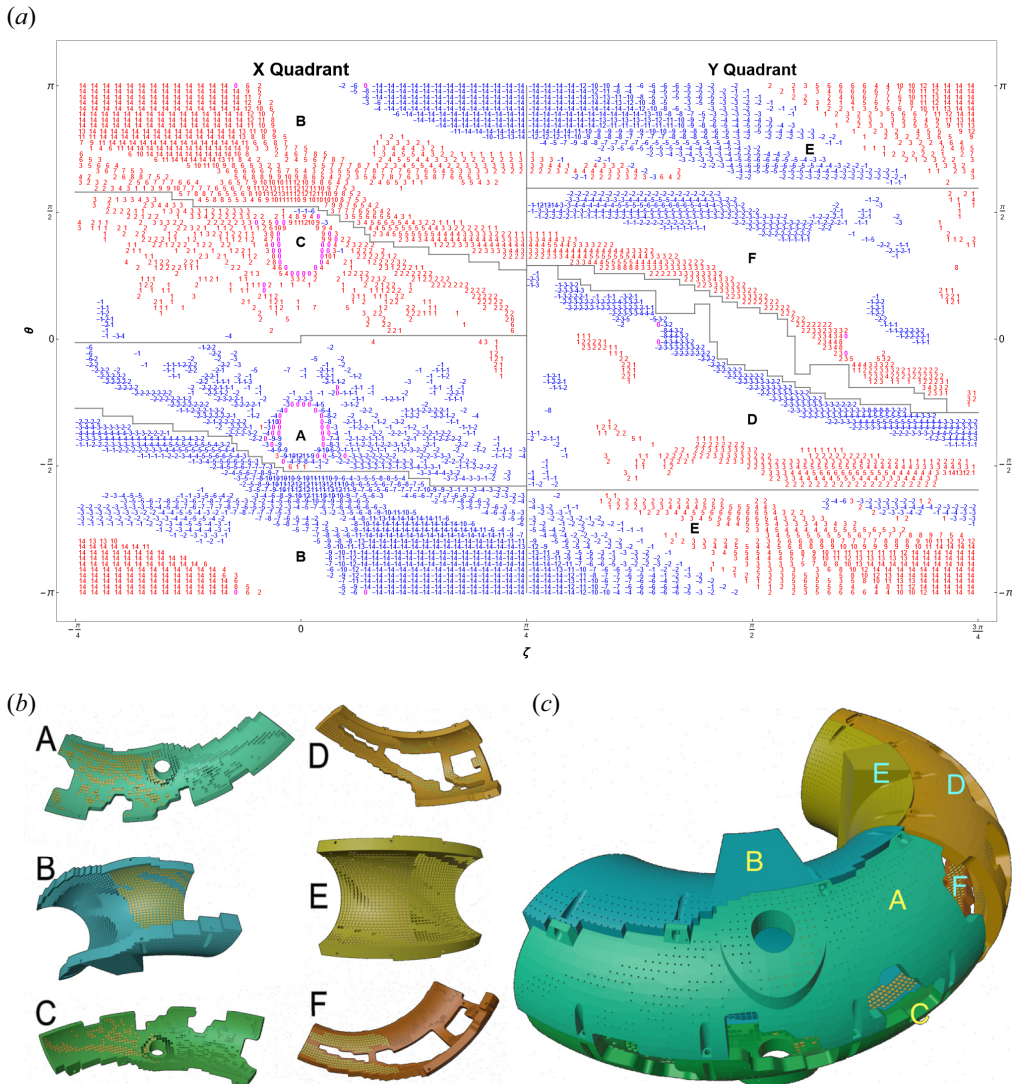


FIGURE 10. (a) This 2D map indicates the exact thickness of magnet towers $[-14, 14]$ where blue and red correspond to signed orientation. The magenta-coloured ‘0’ marks locations where one tower contains two magnets separated by a gap (cases (d,e) in figure 4). In contrast to figure 2, which shows the physics half period $\zeta \in (0, \pi/2)$, this plot shows two engineering half pipes $\zeta \in (-\pi/4, 3\pi/4)$. (b,c) The 2D regions above translate into 3D PM holder parts. Parts A/C and D/F are mechanically identical, while the magnets inside them have opposite signs. Note that the half period of magnets (gold) indicates how the six parts fit together.

described in § 2, a gap was prespaced in the initial dipole geometry to account for this partition.

- (b) Add rectangular support spines for mounting the PM holders to the middle of each engineering quadrant, as shown in figure 11. This straddles the symmetry plane to reduce net forces.
- (c) Introduce thin ‘drain holes’ 1.6 mm in diameter behind every magnet cavity. These holes have dual purpose: (i) they allow for air bubbles and glue to drain out from

beneath the magnets after insertion; (ii) they also allow a long continuous rod to be inserted behind the magnet for ‘emergency ejection’ when a mistake is detected. The length of these drain holes varies from 1 mm (where the magnet cavity is flush against the outside of the PM holder) to greater than 15 cm (where the ‘ray’ traced behind the cavity must pass through the support spine). See § 5.2.

- (d) Cut each quadrant into three poloidal subsegments. This is necessary for assembling around VV ports, for accessing the inside of the toroidal shell and for fitting within available 3D print volumes. The two outboard pieces are identical and each span about 90° of poloidal angle; the inboard piece is larger, contains the square spine and spans a poloidal angle of about 180° . The outboard pieces are further subdivided toroidally into two parts, in order to fit within the print volume (see figure 19). A tongue-and-groove joint is bonded by a cyanoacrylate adhesive-based superglue prior to magnet insertion.
- (e) Thin out the unoccupied thickness of the outboard pieces. The thickest magnet areas are concentrated on the inboard side for MUSE (see figure 8). Since the outboard magnets are comparatively thin, less 3D-printed material is required. We observe that this tends to be the case for QA stellarators, because quasi-axisymmetry keeps high-field regions exclusively on the inboard.
- (f) Remove parts of the outboard pieces for port access and viewing the plasma. Even if the ports themselves are not stellarator symmetric on the VV, the PM holder port holes must be stellarator symmetric when the target plasma equilibrium is stellarator symmetric. This is because a stellarator symmetric surface requires stellarator symmetric sources, but a port is not always desired. There are 12 port holes on the MUSE PM holder and only 7 ports installed on the glass VV.
- (g) Add features to house nuts and bolts to bring the poloidal subsegments together. We placed 4 toroidally spaced junctions on each poloidal cut for a total of 12 bolts.
- (h) Introduce ‘tunnels’ for VV support fingers to reach in through the PM holder to hold the VV.
- (i) Remove mass in hashes (see figure 11) from the inboard pieces in places where the cavities are not full depth. This decreases the total 3D printed material which improves production speed and cost. More critically, it reduces the solid volume and introduces cooling channels, where asymmetric cooling is a known cause of error in 3D print accuracy. This issue and its resolution are further described in § 4.2 on sensitivity.

4. Analysis

In this section, we analyse the MUSE configuration for finite magnet permeability effects, sensitivity to PM perturbations and magnetostatic forces.

4.1. Finite μ

Magnetic permeability adds a nonlinear interaction between PMs. Ideal PMs are linearly independent sources. The magnetisation of real magnets, however, is modified by the field produced by their neighbours. While neodymium PMs are an almost ideal source of magnetisation, they do have a small relative permeability $\mu_{\parallel} = 1.05$ in the direction along the magnetisation. In addition, they have been measured to have a perpendicular permeability of $\mu_{\perp} = 1.15$ (Katter 2005). For optimisation, we assume $\mu = 1$ such that the field from all magnet sources is linearly independent. This section examines the effects of $\mu > 1$ for both isotropic and anisotropic permeability cases. We include only PM–PM interactions. The TF–PM interactions are assumed to be negligible because, at the PM

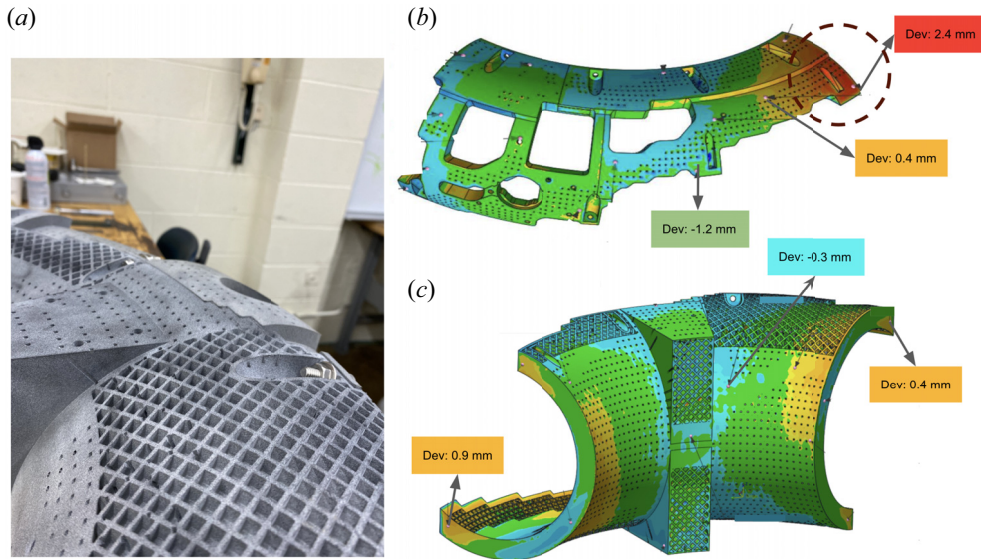


FIGURE 11. (a) Photo of X inboard PM holder showing the drain small drain holes backing each cavity, and the square hashes that are removed for handling heat. (c) Computer rendering of the same part, showing the support spine, drain holes and heat hashes. The colour indicates deviation from the CAD design, as measured by laser metrology on the actual 3D print. (b) Similar rendering and measurement for a Y outboard piece. This is the longest part in our system. The red dashed circle draws attention to an area where deviations were particularly severe.

surface positions, the TF field (0.1–0.2 T) is small compared with the PM–PM fringe fields (≥ 1 T). Moreover, the perpendicular interaction between B_{TF} and μ_{\perp} generates a field aligned dipole moment with characteristic scale ~ 2 cm. The magnitude of this field is small compared with the magnitude of the main TF field in the plasma.

MagTense is an open-source magnetics framework for Python and Matlab that computes finite μ effects (Björk *et al.* 2021). MagTense discretises a volume using tiles and computes the analytic demagnetisation field for finite elements. It assumes that every tile has uniform magnetisation. Since commercial magnets have nominally uniform bulk properties, it is valid to treat finite elements as uniform for our study. The magnetic permeability can be varied in parallel and perpendicular directions. We start with an optimised set of point-dipole slices from FAMUS, expand them into finite-volume magnetisation towers in FICUS, and use MagTense to compute the magnetic field on a spatial grid. The spatial grid is stored in cylindrical coordinates as an MGRID file, which is fed into VMEC (Hirshman & Whitson 1983) to model the free-boundary MHD equilibrium. We report the effects of finite μ_{\parallel} and μ_{\perp} from PM–PM interactions on the rotational transform profile and the neoclassical effective ripple profile which is computed from the STELLOPT code NEO Nemov *et al.* (1999).

Let us start with the isotropic case $\mu_{\perp} = \mu_{\parallel}$. Increasing the relative permeability is found to decrease rotational transform. This can be interpreted as decreasing the poloidal field at a fixed TF. Figure 12 shows that this effect can be compensated by correcting the TF coil current to decrease the TF. A $\delta I/I < 2.5\%$ change in current (10.4 A change over a 433.0 A base current) returns the original target equilibrium. A similar effect was found for the effective ripple. Fluxes from neoclassical transport are proportional to the metric $\epsilon_{\text{eff}}^{3/2}$ in the asymptotically collisionless $1/\nu$ regime (Nemov *et al.* 1999). This proxy

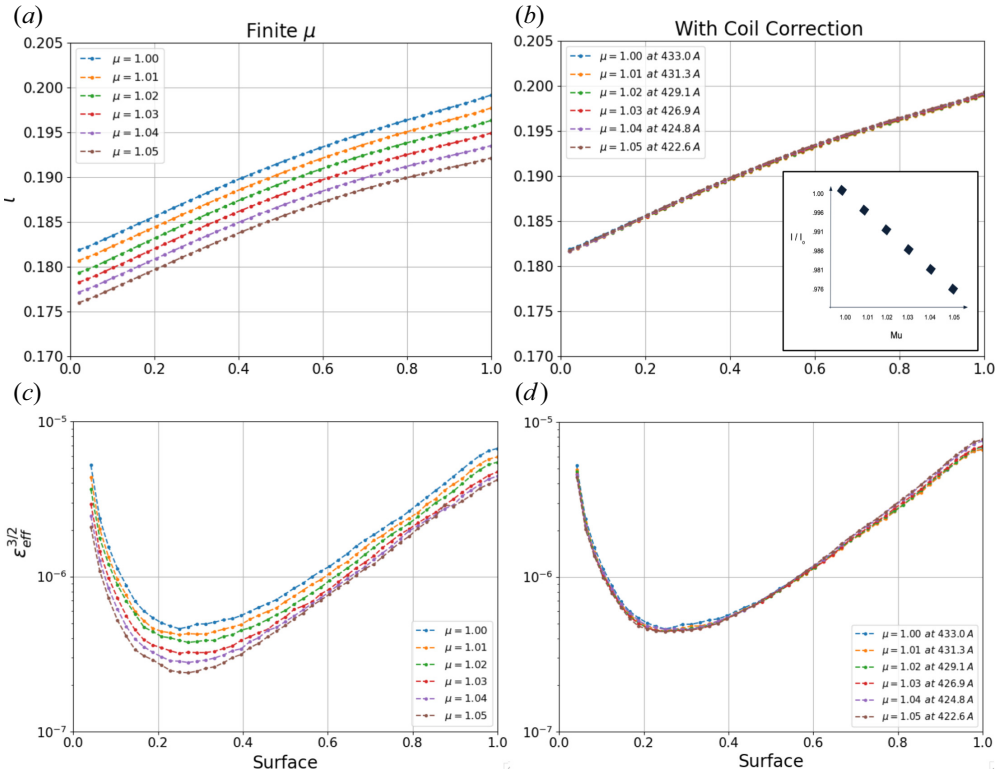


FIGURE 12. (a,c) Effect of finite $\mu > 1$ on rotational transform ι profiles and the neoclassical transport metric $\epsilon_{\text{eff}}^{3/2}$ profiles, using surfaces equally spaced in toroidal flux as the radial coordinate. (b,d) Both sets of profiles can be reduced to the $\mu = 1$ case, provided a small decrement in TF coil current. The inset bounds the magnitude of this change $\delta I/I < 2.5\%$.

decreases in response to magnetic permeability. While having a lower ripple is desirable for particle confinement, we interpret this effect as a decrease in overall shaping, which is undesirable. This is supported by the fact that the iota profile also decreases. The same TF coil adjustment that corrects the iota profile also corrects the epsilon effective profile to that of our base case. Modulating TF current to tune rotational transform was also used, for example, on W7X to correct for elastic deformation under electromagnetic coil loads (Lazerson *et al.* 2019)

Anisotropic permeability effects are shown in figure 13, which varies $(\mu_{\parallel}, \mu_{\perp})$ independently and plots the resulting rotational transform profiles. The effect of $\mu \neq 1$ is dominated by μ_{\parallel} , and μ_{\perp} has a comparatively small effect. The dominance of parallel permeability μ_{\parallel} can be explained by figure 9. Since the radially oriented magnets exist to produce poloidal field, most of the field points parallel to the PM. There is not much field perpendicular to the PM. Thus, the effect of μ_{\perp} is subdominant. A different PM stellarator, one which allows magnets to be orientated orthogonal to each other, may reach another conclusion. These results demonstrate that $\mu > 1$ magnetic permeability effects do not pose a challenge for the MUSE PM stellarator. There is a shift in the ι profile, which can produce islands, but the effect can be compensated by adjusting the TF by 2.5%. This justifies using the $\mu = 1$ model for optimisation.

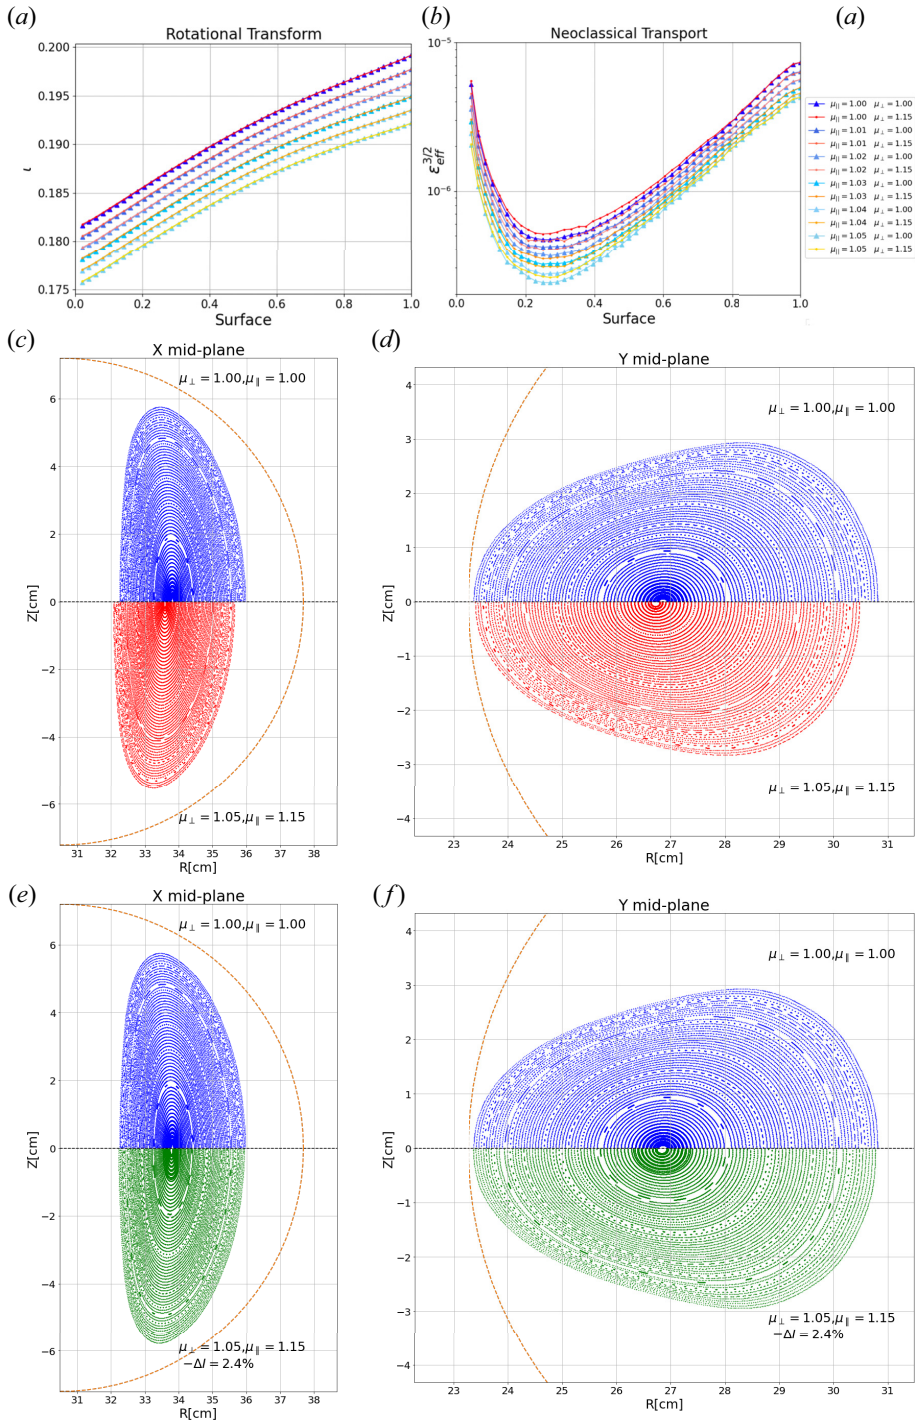


FIGURE 13. Expansion and inclusion of the $(\mu_{\parallel}, \mu_{\perp})$ combinations from figure 12. (a,b) We show 96 profiles, where μ_{\perp} ranges from 1.00 to 1.15, and μ_{\parallel} ranges from 1.00 to 1.05. Observe that the curves varying μ_{\perp} at fixed μ_{\parallel} are clustered together. (c,d) We compare Poincaré plots between the base case (blue) and the most perturbed case (red). Without correction, the axis is shifted, and some flux surfaces are removed from inside the VV volume. (e,f) The most perturbed case with correction (green) matches the base case.

4.2. Sensitivity

The fields in MUSE are designed assuming that every PM has a precise position, magnetisation magnitude and orientation. In reality, the magnetic material itself will have some finite tolerance magnetisation strength and orientation. There will also be imprecisions from assembly that affects position and orientation. In this section, we consider both stochastic and correlated sensitivities, and we set bounds on assembly tolerance.

The effect of uncorrelated perturbations is studied by generating ($N = 600$) magnet sets with stochastic perturbations in orientation, magnitude and position. Then we compute the surface error field metric f_B for each. Figure 14 shows the distributions for four sets of perturbations. Each perturbation is sampled from a uniform distribution. The δM varies magnitude by up to 3%. The $\delta \mathbf{r}$ set perturbs the centre-of-mass position by up to 1 mm. The $\delta \mathbf{n}$ set varies orientation by up to 5° . Each of the above independently varies all dipoles in the system by a random perturbation drawn from a uniform distribution. A fourth set combines the above by perturbing every magnet by δM , $\delta \mathbf{r}$ and $\delta \mathbf{n}$. The threshold for good Poincaré plots and STELLOPT metrics was just above $\tilde{f}_B = 1.5$, a 50% perturbation. Even the outliers of the combined perturbation lie below this threshold. The histograms show that these upper bounds on magnetisation, orientation and position perturbation magnitudes yield acceptable stellarator metrics. Therefore, we used this set of tolerances for magnet purchases.

Laser metrology was performed on the magnet holders (see figure 11). The first round of PM holders found significant deviations up to 2.5 mm. These localised around the edge of the PM holder, and the deviations tended to shrink the printed PM holder relative to the CAD. The distortion is a cumulative error that peaks at the PM holder's edges. It appears to accumulate over the PM holder body and is significantly less over most of the volume. A series of correlated perturbation studies were done to determine what magnitude of distortion is tolerable. Stellarator flux surfaces were found to be relatively insensitive to distortions up to 10 mm. However, we found difficulty inserting test magnets into square cavities where the 3D print distortion exceeds 1.3 mm. Therefore, the PM holder tolerance is not set by the physics of magnetic sensitivity. Instead, the tolerance is set by a mechanical limit on magnet spacing. It was hypothesised that the shrinking distortion results from uneven cooling, as the first print was fan-cooled over a few hours. To address this, the subsequent magnet holders were inertially cooled over multiple days. Cross braces were also added within the torus to distribute thermal stress more evenly. The final 3D prints had bulk deviations less than 1.2 mm, while deviations less than 2.5 mm were tolerated in isolated pockets. For both physical and mechanical limits, we conclude that the final PM holders are satisfactory.

4.3. Magnetostatic forces

To evaluate forces in a PM stellarator, one must use the near-field interaction between neighbouring magnets to determine internal stress and the net force on the support structure. The dipole approximation is insufficient because it diverges as $1/r^3$ near the source. An exact analytic formula that is valid everywhere in space for the magnetic field from a rectangular prism magnet is presented in Appendix A. Given the magnetic field, forces on the PM can be evaluated using the method of magnetostatic charges. Magnetisation is a magnetic surface charge density, $\sigma_m = M$. One can discretise a charge element $\Delta q_m = \sigma_m \Delta A$, where ΔA is an area element such that the net force is $F = \sum B \Delta q_m$.




FIGURE 14. Histograms showing the relative change for the error field metric $\tilde{f}_B = f_{B,\text{perturbed}}/f_{B,\text{nominal}}$. Here $N = 600$ sets of perturbed magnets were generated for each stochastic sample.

This force is analogous to the electrostatic Coulomb force $F = qE$. The magnetic monopole q_m does not physically exist, but it is a valid computational tool in magnetostatics. The magnetic surface charge density, being equal to material magnetisation M , has units of A/m. Then the magnetic charge q_m has units A m, which is consistent with the magnetic dipole moment $m = q_m d$ having units A m^2 . Since all our magnet towers share a common footprint (40 mm^2), they all carry the same total magnetic charge, $q_m \approx 750 \text{ A m}$ or, equivalently, N T^{-1} . Since the total charge scales like surface area R^2 (for photographic enlargement), we find that pressure and stress density is scale-independent for a PM stellarator.

Figure 15 shows the PM–PM forces in our system. As in the finite permeability study, this plot neglects the TF–PM forces because B_{TF} is an order of magnitude smaller than B_{PM} at the surface of the PMs. Let us recall the fact that, while the numerical magnet optimisation decomposed towers into dipole slices, physical construction used single magnets for all towers. We examine the force on each magnet individually by summing the B field from all other magnets and applying the surface charge method. The resulting force peaks at approximately 20 N in the inboard region where magnetisation is thickest. The PM holders straddle the symmetry plane to minimise net force because individual force vectors cancel pairwise. This stores energy inside the PM holder as internal stress instead of external forces or torques. The magnetic pressure inside individual magnet

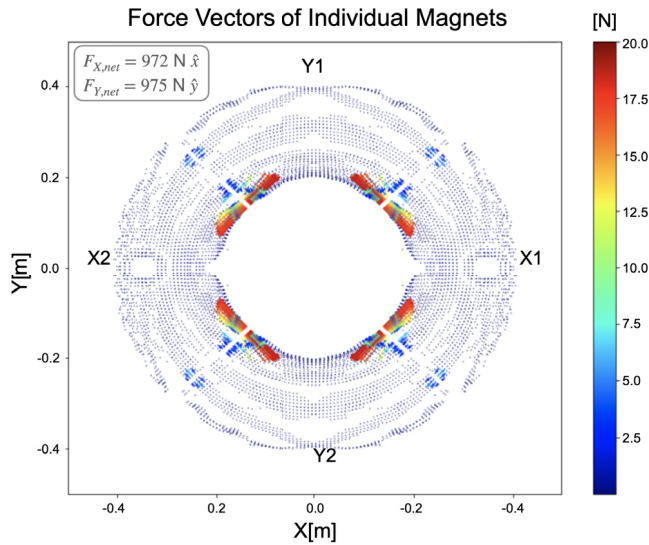


FIGURE 15. Net force (colour bar) on each magnet tower (max ~ 20 N). The legend sums over towers to compute the net total force on each PM holder quadrant (~ 1000 N). There are approximately 2500 towers in each quadrant. Due to symmetry, only the radial component of the net force on each quadrant remains. The peak forces localise to the quadrant boundaries because internal forces cancel (they appear as stress instead) on the interior of magnetisation.

towers (0.5 MPa from 1 T field intensity) is negligible compared with the tensile strength of the magnets (~ 80 MPa). The shear stress from TF field is even less because the TF field is of order 0.1–0.2 T. Therefore, our analysis is only concerned with the force and stresses on the PM holder, and not stresses on the magnets themselves. The remaining radial force is of order 1000 N (the weight of approximately 100 kg), which necessitates special assembly equipment. Section 5.3 describes how a hand-turned crankshaft is deployed to aid construction.

5. Construction

We inserted 9736 magnets towers into the PM holder by hand and fixed them in place with superglue. Tools and procedures were designed to increase accuracy and efficiency. The most important of these are the magnet insertion stencil and the cylindrical guide magnet. This section discusses PM assembly tools and torus assembly.

5.1. Magnet stencil

PM assembly is streamlined when magnets of a specific size and orientation are inserted as a batch. The stencil is designed to identify holes that contain the same magnet size and orientation while covering all others. Magnet size is distinct from cavity depth because some magnets are recessed into their cavities by the optimisation algorithm (see figure 4). We first prototyped 3D-printed plastic covers that conformed to the toroidal shape and later pivoted to 2D paper stencils that can be cut from ordinary card stock. The latter requires an algorithm to map the torus and its magnet holes onto flat 2D sheets, as illustrated in figure 16.

A torus cannot be unwrapped into flat paper because it has non-zero Riemann curvature, so we make a series of approximations. The first step is to partition each toroidal quadrant into straight cylinder wedges. Looking from above, these cylindrical wedges project into

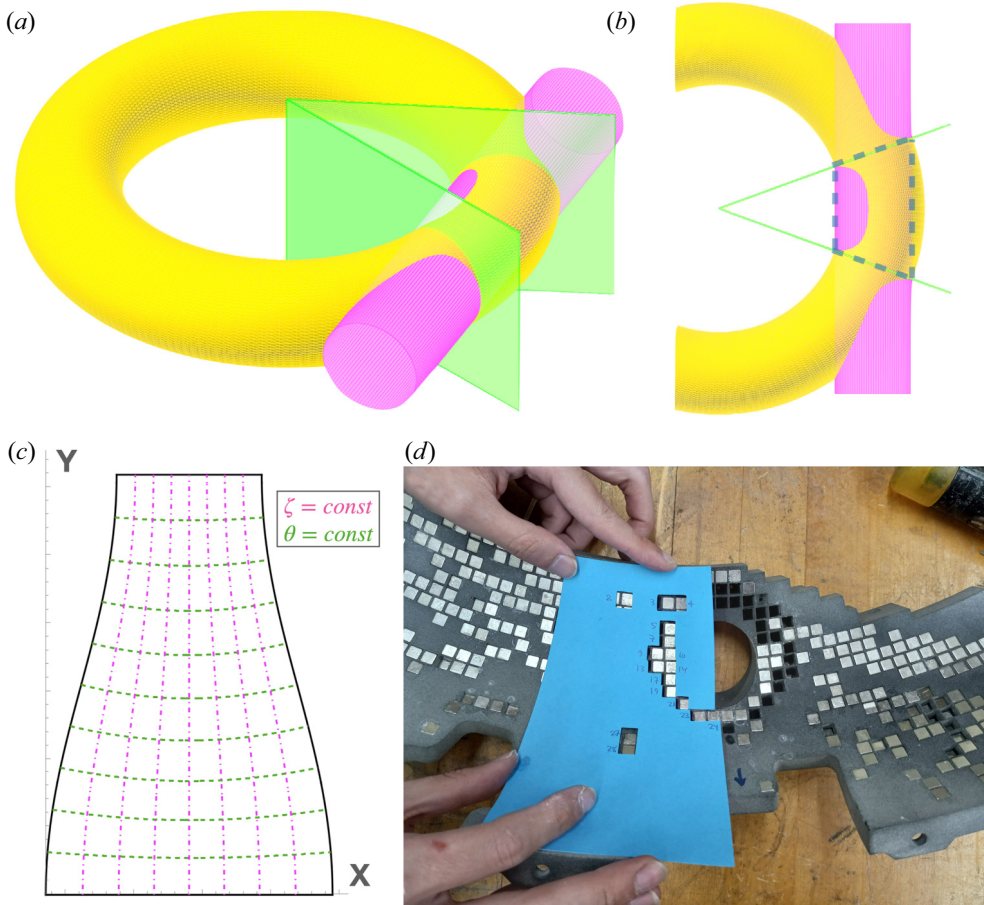


FIGURE 16. The stencils are made by projecting a torus (ζ, θ) onto a cylindrical segment (u, v) , then unrolling the cylinder into a plane (x, y) . (a) 3D perspective. The depicted angle ($\Delta\zeta = 40^\circ$) is exaggerated for illustration (the actual stencils used 15°). (b) 2D projection of the trapezoidal cross section. The cylinder is defined such that its intersection with the $\text{const } \zeta$ planes are circular. This sets $a_{\text{cyl}} = a_{\text{torus}} \cos \phi$ and shifts the centre of the cylinder to $R_{\text{torus}} \cos \phi$, where $\phi = \Delta\zeta/2$. As a result, outboard poloidal angles radially project ‘in’ from the torus to the cylinder, whereas inboard poloidal angles project ‘out’. (c) Curvature of the cylinder unrolled into planar coordinates. All four corners are right angles. The vertical extent covers only 180° of the poloidal extent. The height of the unfolded map as a function of cylindrical–poloidal angle v is an elliptic integral. (d) Photo showing how the stencil selectively reveals holes of a common type on the PM holder.

trapezoids in the x – y plane. This removes curvature from the toroidal direction, and it introduces a shift in the poloidal angle coordinate. Next, we unroll each cylinder into a flat sheet, which maps straight lines into curved boundaries. It is helpful to define a chain of one-to-one maps

$$\begin{pmatrix} \zeta \\ \theta \end{pmatrix} \rightarrow \begin{pmatrix} u \\ v \end{pmatrix} \rightarrow \begin{pmatrix} x \\ y \end{pmatrix} \tag{5.1}$$

that carries points from toroidal coordinates (ζ, θ) , to coordinates on a cylinder (u, v) , to flat stencil coordinates (x, y) . Each step removes a degree of curvature. The cylinder is

parametrised by a sector size $\Delta\zeta$. In our case, each stencil covers a sixth of a quadrant, so $\Delta\zeta = \pi/12$. The (ζ, θ) and (u, v) coordinate systems share the same toroidal angle everywhere, but their poloidal coordinates coincide only at the ends of the cylinder. For this reason, it is convenient to centre the toroidal coordinate $u = 0$ such that it bisects the cylindrical axis of each stencil segment. We define the poloidal coordinate $v = 0$ on the outboard midplane where the major radius is largest. The first map is given by

$$u = \zeta, \quad (5.2)$$

$$v = \theta - \sin(\theta) \left[\cos(\zeta) - \cos\left(\frac{\Delta\zeta}{2}\right) \right] \frac{R + a \cos(\theta)}{a}, \quad (5.3)$$

where R and a are the major and minor radii of the torus. The second term in (5.3) corrects for the shift in poloidal angle between a toroidal wedge and its cylindrical approximation. It vanishes at either end where $\zeta = \pm\Delta\zeta/2$ because the mapping $v = \theta$ is exact on the boundary. The shift in angle is largest at $\theta = \pm\pi/2$, and it is zero on the x - y plane where $\theta = 0, \pi$. The amplitude of this correction vanishes in the limit $\Delta\zeta \rightarrow 0$. Next, the cylinder is mapped to a flat 2D stencil by

$$x = [R + a \cos(v)] \cos\left(\frac{\Delta\zeta}{2}\right) \tan(u), \quad (5.4)$$

$$y = a \cdot E \left[v, \sin\left(\frac{\Delta\zeta}{2}\right) \right], \quad (5.5)$$

where E is an incomplete elliptic integral of the first kind

$$E(\alpha, m) = \int_0^\alpha \sqrt{1 - m^2 \sin^2 x} dx \quad (5.6)$$

that measures the height $y(v)$ for lines of constant u (see [figure 16c](#)).

With this mapping, we project the boundary of the PM holders and the four edges of each square magnet cavity onto a 2D CAD image, which can be cut from light card stock. We used a computer-controlled paper cutter (the Cricut Maker 3), to automate the process. There are 29 magnet types (14 of each sign, plus 1 for holes that have both), and 4 unique magnet holder segments (2 for X/Y times 2 for outboard/inboard). Therefore, there exist at most 116 sets of stencils. In practice, 81 stencil sets were used because not every option was populated. The total stencil area is further reduced by using stellarator symmetry to half the number of printed stencils.

5.2. PM assembly

MUSE uses 14 types of magnets that are composed from a single fundamental unit, as shown in [figure 17](#). The polarity of these magnets can not be distinguished by eye. We define ‘north’ to be the direction that points parallel to a compass aligned with geomagnetic north. Stellarators and tokamaks have a net dipole moment. It is defined by the direction of net plasma current or the equivalent rotational transform. MUSE’s net dipole moment points up out of the table. To match conventions, we arrange the PM such that this net dipole moment agrees with the ‘north’ of the individual PMs. [Figure 9](#) shows the north field at the $\zeta = 0$ symmetry plane. Since stellarators are helical, north points down at the $\zeta = \pi/2$ symmetry plane. Nonetheless, the net dipole moment north

Magnet Count: 9736




FIGURE 17. The distribution of magnet sizes in MUSE is peaked at size 2 and 14. The inset shows a photograph of MUSE magnet types, which are all composed from a single fundamental unit.

points up on average. With magnetic polarity fixed, we use reference magnets to orient the otherwise indistinguishable PM towers. A piece of magnetic material is used to hold and organise magnets during assembly.

To pick up a magnet during assembly, we used a hollow cylinder guide magnet that slides along a non-magnetic wooden stick, as shown in figure 18. The guide magnet selects a polarity that remains fixed when the assembly tool twists and turns in hand. We used superglue (ethyl cyanoacrylate adhesive) to fix the magnets in the PM holder. Glue was preapplied in the square cavity. It takes 5–15 s for the glue to be dry enough to hold a magnet in the presence of background magnetic forces. The bond continues to strengthen over 24 h. Acetone was used to undo glue bonds where necessary. We note that this solvent does not affect the 3D-printed nylon material. Some magnets did not sit squarely into their cavity after insertion. We used a hammer and punch to level the magnet in place. The punches are made from non-magnetic brass. A 3/16-inch diameter is most useful for leveling and 1/16-inch diameter is useful for removing incorrectly placed magnets from behind using the drain hole. We used calipers to measure cavity depth to detect when a magnet was placed incorrectly. We also used calipers to verify when a magnet was positioned correctly in a recessed cavity (see case (c) in figure 4). Non-magnetic titanium pliers are useful for removing incorrect magnets after they have been partially dislodged by applying a thin punch through the drain hole. Dental tools are useful for removing debris such as dried glue from the corners of the square holes. They are also useful for removing residual powder from the powder-based 3D print. For the strongest regions of magnetisation, the cylindrical guide magnet was not sufficient to bring a target magnet into its cavity. The magnets tend to pivot and deflect towards already-glued neighbours. In these cases, we used the square hole of a titanium non-magnetic socket wrench to rigidly guide a magnet into its cavity.




FIGURE 18. A cylindrical guide magnet is used to insert magnets. The guide magnet gives a distinct polarity to the otherwise indistinguishable magnet in hand. It can be detached while normal force is applied to secure the target magnet while glue dries. Also shown are the escape holes, which are used to eject misplaced magnets from the PM holder cavities.

As for the assembly procedure, we found that working in pairs to assemble each poloidal subsegment improves efficiency and decreases errors. Three common errors were identified: inserting a magnet into the wrong hole, inserting a magnet with the wrong polarity and having the glue dry while a partially inserted magnet is tilted. Prior to inserting a set of magnets, we used the stencil to mark all the corresponding holes and then removed the stencils during actual assembly. Since the magnets themselves have unmarked polarity, the guide cylinder is helpful for identifying orientation throughout each stencil set. It is strategic to start with the smallest magnets because these thin slices tend to flip in the presence of larger magnets. Next, we do all the ‘wyrms’ because they contain thin magnets that are susceptible to flipping. Then we go through the remaining magnets in order from smallest to largest. After inserting all the magnets in one stencil set, checks are conducted to verify that the set is indeed complete and that all magnets are oriented correctly and leveled.

Inevitably, some magnets were inserted incorrectly and needed to be removed. This occurred roughly 100 times or on 1 % of the total magnets. A magnet can be driven out by punching a thin rod through the drain hole on the back of the PM holder. When a magnet is removed, the hole is cleaned using acetone and dental tools before reinsertion. On rare occasions, it was necessary to apply acetone carefully to weaken the glue. Also on rare occasions, a rotary tool was required to clean residual glue from a cavity after the magnet was removed. This occurred fewer than 10 times. After all the PM holders were completed, we used an axial Hall probe to check for incorrect magnitudes and orientations. Tilted, displaced and significantly demagnetised magnets were detected by cross-checking the field strength at corresponding locations on identical pairs of PM holders (see figure 19). Approximately 10 magnets were found with the wrong polarity. The resulting error field has potentially large effects, so these magnets were removed and corrected. This check discovered some demagnetised size-1 (thinnest) magnets. Using a Hall probe we found that approximately 30 % of these smallest magnets were demagnetised by a magnitude of approximately 20 %, in spite of a formal quality assurance document from the vendor stating the spread in magnetisation is less than <3 %. A measurement of extra unused

size-1 magnets found the perturbation occurring in similar proportions. This suggests that the ‘demagnetisation’ is actually a quality control issue from the vendor. We assessed the effect of this unexpected error using uncorrelated perturbation calculations similar to those described in § 4.2. Fortunately, it was found that for perturbations applied to the size-1 magnets, even a 50 % change in magnitude was inconsequential. A study of $N = 600$ random draws showed that the nominal $f_B = 1.29 \times 10^{-8}$ was degraded to 1.45×10^{-8} in the worst case, while the threshold for acceptance is $< 2.2 \times 10^{-8}$. Moreover, we found that for the 20 % magnitude perturbations on the size-1 magnets, the error field metric was even improved on average. The surprising observation could be explained by the fact discretising the magnet thickness makes size 1 the smallest non-zero unit. It would appear that the stellarator would gladly take ‘size 0.8’, even if randomly distributed. This fortuitous outcome retires the risk from these partially demagnetised magnets.

Opportunities to improve PM holder designs for future experiments include the following.

- (i) Increasing the thickness of the wall behind the magnets (currently 1 mm). During insertion, we punctured a few cavities when tapping short magnets into deep holes. This was repaired by gluing the punctured fragment back onto the PM holder, and it did not leave any permanent deformation.
- (ii) For parting the poloidal subsegments it is better to make cuts along the minor radius direction, rather than cuts that are horizontal or vertical along the (R, Z) direction because the former mates poloidal subsegments at 90° faces. We unintentionally used radial cuts for the Y quadrants but not for the X quadrants. The slanting face redirects the compression force of the bolt joining two subsegments and introduces about 1 mm of sliding motion, even when the bolts are tightened completely. We plan to resolve this by using Teflon tape to increase friction, but the issue could be avoided altogether by using radial cuts to ensure a 90° mating surface.
- (iii) The X quadrants have jagged toroidal cuts, whereas the Y quadrants have a comparatively simple cut at constant poloidal angle. The latter is more favourable. It is an optimisation coincidence that we were able to find such a cut for Y but not for X . It should be possible to predesign poloidal gaps in the dipole geometry since similar gaps are already included for the toroidal cut that separates X from Y .

5.3. Torus assembly

MUSE torus assembly consists of three systems: the PM holders, a glass VV and the TF coils. A photograph of each system is given in [figure 19](#).

The MUSE VV is made from four borosilicate glass quadrants. Each 90° turn (7.5 cm minor radius, 30 cm major radius) was supplied by Schott-Kimax as off-the-shelf parts designed for use in chemical process plants. The advantage is low cost compared with custom metal fabrication. The transparent glass also allows direct visualisation of the stellarator plasma. Colleagues at the University of Wisconsin–Madison used a similar product to build a VV for the HELICOTOR experiment (O. Schmitz, private communication). Our glass quadrants did not immediately fit into a torus due to the bends being too square, the openings not forming exactly 90° , and the openings tilting out of plane such that two flanges would not meet at a flat surface. We bought 10 pieces and could not find a combination of 4 that fit into a torus. We hired professional glass blowers to align the flanges, smooth the bends and add ports. This process successfully produced a VV that has achieved 7×10^{-8} Torr. The completed VV has seven ports: three vertical ports to match the X pipe, three horizontal ports to match the Y pipe and a single port pipe that can fit in either X or Y positions.

(a)



(b)



(c)




FIGURE 19. (a) The 12 completed PM holder subsegments. Each column is a triplet that makes one quadrant. Y is on the left and X is on the right. (b) The 16 planar, circular TF coils are positioned inside the water-jet cut support structure. (c) The glass VV is joined by 3D-printed low-thickness VV couplers. Glass ports were hot welded to the torus. The PM quadrants accommodate up to three ports each. Seven are installed at present on the north, west and south quadrants.




FIGURE 20. MUSE is assembled in four quadrants of VV, PM holders and TF coils. Only the VV and PM are shown in this photo. The assembly structure brings the four quadrants together using a hand-turned threaded shaft. This is necessary to overcome the mutual repulsion between PM holders.

The TF coil system uses legacy ‘L-2’ coils (Bonanos 1964), which were constructed for Francis Chen’s linear experiments (Chen 1995). They are 33-turn water-cooled copper coils wound in two layers, 16 turns in each layer with 1 turn crossing over. MUSE uses 16 coils. The resistance of each coil is about $23\text{ m}\Omega$. At nominal operation, the TF system produces a 1.5 kG magnetic field on axis, while drawing 70 kW of electrical power (433.0 A, 160 V). We purchased a pair of MagnaPower DC power supplies (500 A, 100 V).

To assemble the torus, we first installed three helicon antennae on two of the VV quadrants for ion heating. Then the three poloidal subsegments were fit around the VV quadrant. Four TF coils were mounted on each quadrant, and the four quadrants were brought together radially using an assembly jig. At the final stage of toroidal mating, conducting grounding bands were inserted inside the glass VV flanges, and VV clamps were installed outside the VV flanges. The VV clamps are custom made to have a low radial profile so that they can fit over the flange and below the PM holders. The torus support structure is water-jet cut from plates of G-10 fiberglass laminate. The TF coils butt against an hourglass-shaped central column. Figure 20 shows the assembly structure, and figure 21 shows the completed device.

5.4. Cost and timeline

The physics design of MUSE’s stellarator equilibrium and PM optimisation was completed in March 2021. The engineering Final Design Review was held in December 2021. Construction was completed in January 2023 with first plasma achieved in February 2023. The entire stellarator was built in approximately 1 year, including procurements. To the best of the authors’ knowledge, this is the fastest build for an optimised stellarator of any size.

In terms of material expense, the PM cost \$16 5000 (2022 USD), which included a 30 % contingency for extra magnets. The VV cost approximately \$5000 for raw material, but




FIGURE 21. The MUSE PM stellarator is mounted on an optical table. Simple planar coils surround the circular-cross-section PM holder and VV. The torus support structure is made from water-jet cut G-10 plates. Lexan panels are installed between TF coils to add torque resistance. The TF power cables and water cooling manifold are above.

\$10 000 for modifications by expert glass blowers because the original parts were out of round. This endeavor was published on the front cover of *Fusion*, a coincidentally named journal of the American Glassblowers Society (Paris & Souza 2022). The TF coils were freely available from PPPL, but the warehouse inventory valued them at \$1000 each. The 3D-printed PM holders cost about \$2500 per quadrant. Therefore, the total material cost of the MUSE PM stellarator is just under \$60 000. This sum includes only the material cost of stellarator components. It does not include the costs of employee salary, administrative overhead or auxiliary equipment such as power supplies.

6. Discussion

The MUSE stellarator experiment will have at least four plasma sources. The two simplest are arc discharge and DC breakdown. The first plasma was achieved by placing a Tesla coil near one of the ports, while the glass vessel was held under vacuum. Since the VV is an insulator, a stainless steel grounding band is inserted under each of four the VV

junctions. The grounding band is a poloidal loop that also mechanically shields the Viton O-rings from plasma interactions. Another breakdown method is to electrically bias one of the grounding bands relative to the others, making a hollow cathode plasma source. As a third method, three helicon antennae of different lengths were installed on the VV. They are made from copper bands, each with a different helical pitch. These sit outside the VV, beneath the PM holders. Fourth, we can use a 4.5 GHz electron cyclotron source for plasma startup and electron heating. Compared with ECH breakdown, the helicon antennae have the advantage of direct ion heating (Granetzny, Schmitz & Zepp 2022).

Initial experiments will start with electron beam mapping to assess the quality of the flux surfaces and validate the magnet construction and alignment. This is typically done on all stellarator experiments, for example (Talmadge *et al.* 2001) on HSX and (Otte *et al.* 2016) on W7-X. In a second set of experiments, the degree of quasi-axisymmetry will be assessed by measuring the flow damping in the quasi-symmetric direction and comparing it with non-symmetric configurations through modulations of the toroidal magnetic field. The reduction of damping in the symmetric mode is a sensitive indication of successful optimisation. This measurement is similar to early experiments on HSX (Gerhardt *et al.* 2005). A third set of experiments will investigate the physics of helicon-wave-heated plasmas in a quasi-symmetric stellarator.

We would like to comment on the future of PM stellarators. With present materials, the poloidal field produced by an individual PM (without the additive effect of Halbach arrays) is less than 1 T. In an experiment such as W7-X, however, the poloidal field is approximately 1/10th of the TF $B_p = \iota \epsilon B_T$. PMs are responsible for the poloidal field, whereas it is the TF that sets magnetic confinement scaling. For W7-X, $B_T = 2.5$ T, the inverse aspect ratio $\epsilon = 1/10$ and the rotational transform $\iota \approx 1$. Therefore, $B_p \approx 0.25$ T is within the available PM field. Nonetheless, PMs have an adverse scaling compared with electromagnetic coils. Assuming that coil field is limited by material current density, $B_{\text{coil}} \propto I/R = J(A/R) \sim R$ scales favourably when dimensions are enlarged photographically (where I is coil current, J is current density, A is conductor cross section and R is coil radius). In contrast, the PM field is limited by material magnetisation M such that $B_{\text{PM}} \propto m/R^3 = M(V/R^3) \sim 1$ (where m is the dipole moment, V is the magnet volume and R is again distance to the plasma). Coils are better at producing large fields when they are larger, but PMs do not. To make a PM stellarator with larger field, one must increase the PM array's relative thickness, in addition to photographically scaling the size of the array. One solution to this problem is optimising PM together with mildly shaped coils to partially alleviate coil complexity. This is a staged retreat: start with tilted circular coils, then go to shaped but still planar coils and, finally, lightly shaped coils with non-zero torsion. An issue often raised is the viability of PM in a neutronic environment. The Curie temperature of neodymium magnets is above 500 K. Superconducting coils must be cooled well below 100 K, which is possible only when the superconductor is kept in a cryostat behind a neutron shield. Provided that neutron shielding and cryostat must exist for the TF coils, one can put the PM behind the blanket and pick up extra magnetisation from the cryogenic temperature. Whether or not the PM stellarator becomes a fusion reactor, the stellarator concept has many paths to choose from, and the PM approach offers a quick and low-cost path to rapid exploration.

As a next step, we propose the design of new 'MUSE-like configurations' defined such that the PM field is produced by single-orientation magnets assembled around a simple circular cross section. We envision a mid-scale experiment CERBERUS ($R = 1$ m, $B = 1$ T) in which three vastly different stellarators (e.g. QA, QH and QI) can be tested in one facility using interchangeable PM sets and a common set of TF coils. Since the axis shape

for QH (helical) differs from QA (circular) and QI (polygonal) it may be necessary for CERBERUS to use multiple sets of VVs, but the circular TF coils can still be shared.

Another natural extension of the PM stellarator concept is the use of ‘window-pane coils’ with adjustable current. These form a surface array that does not link the magnetic axis (Ku & Boozer 2009). They can be made from a high-current-density material such as REBCO high-temperature superconductors (HTS). A PM tower produces a magnetic field externally indistinguishable from that of a square solenoid. In MUSE, the magnetisation $M \sim 1 \times 10^6 \text{ A m}^{-1}$ would be replaced by approximately 30 kA of current around a $0.5 \times 0.5 \times 2.5 \text{ cm}^3$ volume. This would not be feasible mechanically, due to the small scale and large stress. For larger systems, however, the required current can be reached by existing HTS technology. For example, the SPARC Toroidal Field Modular Coil (TFMC) carries 40 kA across a stack of HTS tapes. The key is that the equivalent solenoid for a larger PM requires the same current. Therefore, it could be feasible to make window-pane coils using a larger footprint than that of MUSE to produce of the order of few-tesla fields. This preserves the coil-simplifying properties of the PM stellarator while addressing the two issues of PMs: not being able to be turned off during construction and not being able to be manipulated during experiments. Two independent proposals exploring this concept were presented at the recent APS meeting (Gates 2022; Mauel *et al.* 2022).

7. Conclusion

This paper documents the design and construction of MUSE. A combination of continuous optimisation and discrete refinement realises low magnetic field error using a set of identical magnets in a simple, uniform geometry. Three metrics were targeted in numeric optimisation: low surface error field f_B , low magnetised volume f_V and binary dipoles $f_D = 0$. We initially optimised to reach $f_B < 2.2 \times 10^{-8}$, but the results of discrete refinement and ‘zotting’ in particular improved our final solution to $f_B = 1.29 \times 10^{-8}$. This suggests that our majority vote rounding algorithm can be improved by an algorithm such as ‘auto-zot’ that considers each dipole individually. Reasonable sensitivity limits were found in position, orientation and magnitude. We purchased more than 10 000 magnets at commercially available specifications: 5% tolerance on magnetisation magnitude, 3° tolerance on magnetisation orientation and 1 mm tolerance in physical dimensions. Multi-jet fusion 3D-print technology was selected for fabricating the PM holders using nylon plastic.

Using the magnetic surface charge method, the static assembled force between the 4 PM holders is calculated to be less than 1000 N. Most of the inter-magnet forces are carried as internal stress in the PM holders. Loading these forces into ANSYS, the peak internal stress was found to be less than 7 MPa, which is safely below the 30 MPa tensile strength of our 3D-print material. Finite permeability effects were also modelled, and we found that the small but finite changes to stellarator metrics can be compensated for by slightly adjusting the TF coil current.

Construction of the MUSE stellarator has been completed, and initial physics experiments are underway. We plan to test the PM optimisation using an electron beam flux surface mapping, and we plan to test the quasi-symmetry by measuring viscous flow damping when the symmetry is intentionally broken. MUSE will also test plasma heating using an RF helicon antenna.

Acknowledgements

The authors are grateful to acknowledge useful discussions with A. Brooks, R. Ellis, K. Hammond, J.-K. Park, C. Rana and P. Titus at PPPL. We thank M. Souza, K. Paris and B. Russel for their glassblowing expertise and for providing access to the Salem

Community College Glass Center. We thank D. Fritz and P. DeWys from Forerunner3D for 3D printing and laser metrology services. We also thank J. Abbate, H. Fetsch, E. Jung, D. Liu, R. Nies, J. Parisi, S. Redd, A. Rutkowski, C. Swanson and S. Trinczek for volunteering at MUSE magnet insertion parties.

Editor P. Helander thanks the referees for their advice in evaluating this article.

Funding

This work was supported by the US Department of Energy contract number DE-AC02-09CH11466, the PPPL LDRD program and the NSF Graduate Research Fellowship Program grant number DGE-2039656.

Declaration of interests

The authors report no conflict of interest.

Appendix A. Exact analytic field

An exact analytic field for a rectangular prism magnet can be derived either by treating the two magnetised faces as a magnetostatic capacitor or by treating the four faces parallel to magnetisation as a square solenoid (Thome & Tarrh 1982; Ciftja 2020)

$$\left. \begin{aligned} H_x &= \frac{M}{4\pi} \sum_{i,j,k}^{0,1} (-1)^{i+j+k} \ln(y_j + r_{ijk}), \\ H_y &= \frac{M}{4\pi} \sum_{i,j,k}^{0,1} (-1)^{i+j+k} \ln(x_i + r_{ijk}), \\ H_z &= \frac{M}{4\pi} \sum_{i,j,k}^{0,1} (-1)^{i+j+k} \arctan\left(\frac{z_k r_{ijk}}{x_i y_j}\right), \end{aligned} \right\} \tag{A1}$$

where each sum gives eight terms, the indexed positions $\{x_i, y_j, z_k\}$ measure distances to the eight corners

$$\left. \begin{aligned} x_0 &= x + L/2, & y_0 &= y + W/2, & z_0 &= z + H/2, \\ x_1 &= x - L/2, & y_1 &= y - W/2, & z_1 &= z - H/2, \end{aligned} \right\} \tag{A2}$$

and $r_{ijk} = \sqrt{x_i^2 + y_j^2 + z_k^2}$ is shorthand for the geometric sum of distances to each corner. This expression computes the auxiliary field $\mathbf{H}(\mathbf{r})$ where $\mathbf{r} = (x, y, z)$ is a position in space, and (L, W, H) give the dimensions of the rectangular prism magnet. This expression assumes that magnetisation is perfectly normal to the faces that define the z direction, $\mathbf{M} = M\hat{z}$.

In general, the magnetisation vector \mathbf{M} need not lie parallel to an axis of the prism. This can occur if the magnet tower is cut at an angle relative to magnetisation, e.g. the case when there is uncertainty on orientation $\delta\mathbf{n}$. One can modify (A1) to account for a general orientation $\mathbf{M} = (M_x, M_y, M_z)$ through the principle of superposition. The vector magnetisation \mathbf{M} is decomposed into three magnetic prisms, occupying the same volume in space, each magnetised normal to one pair of faces.

The magnetic field can be computed everywhere in space $\mathbf{B} = \mu_0(\mathbf{H} + \mathbf{M})$. Inside the magnetised volume \mathbf{H} changes sign (like the electric field across a parallel plate capacitor),

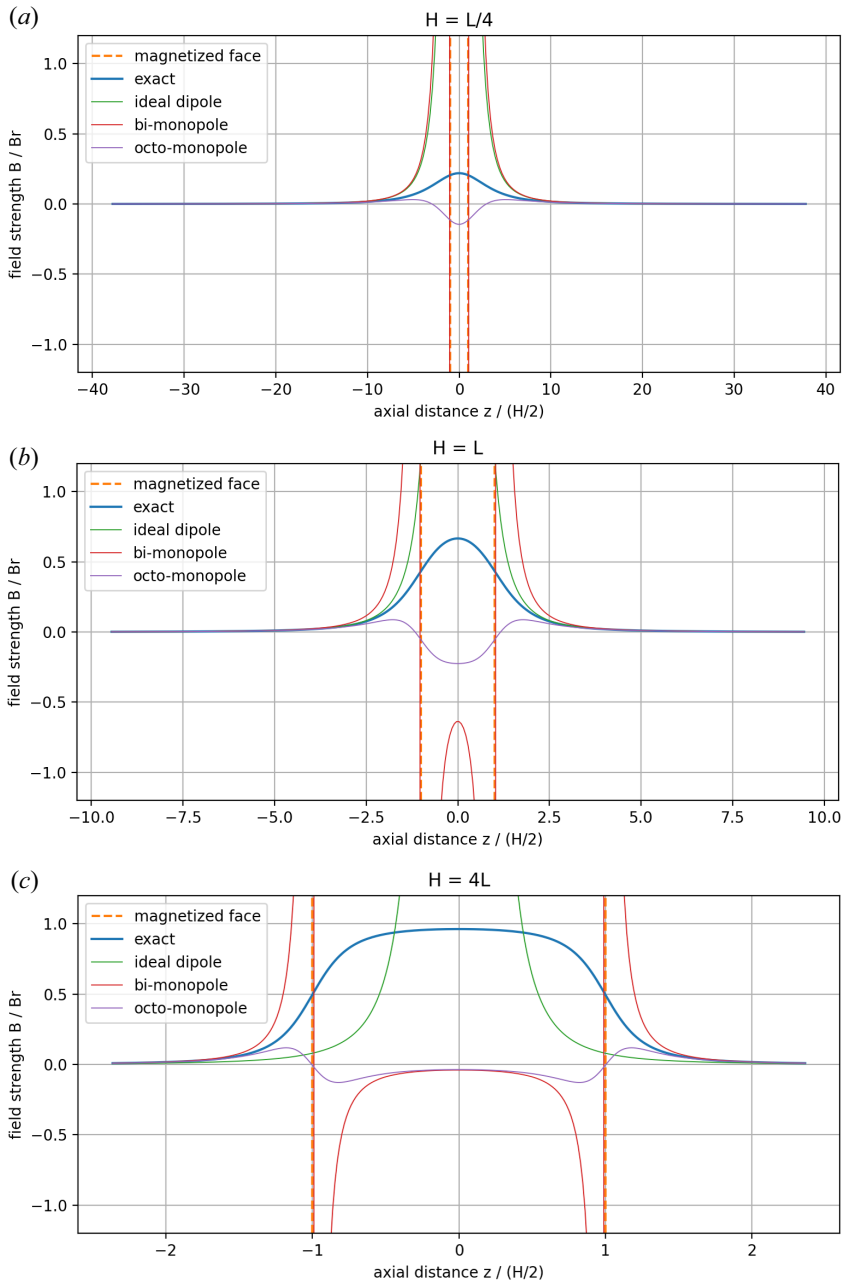


FIGURE 22. Plot of B_z on axis for three rectangular prism geometries: slice ($H < L$), cube ($H = L$) and tower ($H > L$). The axial distance is normalised by half-thickness $H/2$, whereas equal physical lengths are shown across all three plots. The dashed orange vertical line indicates the position of each magnetised face. The exact analytic field is compared with three approximations: an ideal dipole positioned at the centre of mass (green), a pair of monopoles centred on the magnetised faces (red) and a bi-quadrupole (eight monopoles) placed on the magnet corners (purple). The dipole and bi-monopole diverge. The bi-quadrupole is non-singular everywhere, but falsely changes sign near the surface. The magnetic field strength is normalised by remanent field $B_r = \mu_0 M$. The same magnetisation M is used for each case. For asymptotically long magnets $B(0) \rightarrow B_r$ whereas $B(1) \rightarrow B_r/2$ on the surface. For thin magnets $B \ll B_r$.

whereas \mathbf{B} does not. This is analogous to the electric displacement vector $\mathbf{D} = \epsilon_0 \mathbf{E} + \mathbf{P}$ (except that the polarisation field \mathbf{P} includes ϵ_0 by convention). The magnetisation $\mathbf{M} = 0$ is non-zero only inside the magnetic material volume. Figure 22 compares the exact analytic field for slice, cube and tower geometries against various approximations.

REFERENCES

- ALEJALDRE, C., GOZALO, J.J.A., PEREZ, J.B., MAGAÑA, F.C., DIAZ, J.R.C., PEREZ, J.G., LOPEZ-FRAGUAS, A., GARCÍA, L., KRIVENSKI, V.I., MARTÍN, R., *et al.* 1990 TJ-II project: a flexible heliac stellarator. *Fusion Technol.* **17** (1), 131–139.
- ANDERSON, D.T., ABDOU, A., ALMAGRI, A.F., ANDERSON, F.S.B., CANIK, J.M., GUTTENFELDER, W., LECHTE, C., LIKIN, K.M., LU, H., OH, S., *et al.* 2006 Overview of recent results from HSX. *Fusion Sci. Technol.* **50** (2), 171–176.
- ANDERSON, F.S.B., ALMAGRI, A.F., ANDERSON, D.T., MATTHEWS, P.G., TALMADGE, J.N. & SHOHEH, J.L. 1995 The helically symmetric experiment, (HSX) goals, design and status. *Fusion Technol.* **27** (3 T), 273–277.
- BEIDLER, C.D., SMITH, H.M., ALONSO, A., *et al.* 2021 Demonstration of reduced neoclassical energy transport in Wendelstein 7-X. *Nature* **596**, 221–226.
- BELETSKII, A.A., BEREZHNYI, V.L., BURCHENKO, P.Y., CHECHKIN, V.V., CHERNYSHENKO, V.Y., GRIGOR'ÉVA, L.I., GUBAREV, S.P., KONOVALOV, V.G., KULAGA, A.E., KURILO, D.V., *et al.* 2008 First results of the renewed URAGAN-2M torsatron. *Voprosi atomnoi nauki i tekhniki* **6**, 13–15.
- BISHOP, D. 2021 Muse permanent magnet structure structural analysis. *Tech. Rep.* PPPL, mUSE 1-3 CALC 100 Rev. 0.
- BJØRK, R., POULSEN, E.B., NIELSEN, K.K. & INSINGA, A.R. 2021 MagTense: a micromagnetic framework using the analytical demagnetization tensor. *J. Magn. Magn. Mater.* **535**, 168057.
- BONANOS, P. 1964 Coil data compilation. *PPL Tech Memo* (207).
- BOOZER, A.H. 2005 Physics of magnetically confined plasmas. *Rev. Mod. Phys.* **76** (4), 1071–1141.
- BROMBERG, L., ZARNSTORFF, M., MENEGHINI, O., BROWN, T., HEITZENROEDER, P., NEILSON, G.H., MINERVINI, J.V. & BOOZER, A. 2011 Stellarator configuration improvement using high temperature superconducting monoliths. *Fusion Sci. Technol.* **60** (2), 643–647.
- CANIK, J.M., ANDERSON, D.T., ANDERSON, F.S.B., LIKIN, K.M., TALMADGE, J.N. & ZHAI, K. 2007 Experimental demonstration of improved neoclassical transport with quasihelical symmetry. *Phys. Rev. Lett.* **98**, 085002.
- CHEN, F.F. 1995 The “sources” of plasma physics. *IEEE Trans. Plasma Sci.* **23** (1), 20–47.
- CIFTJA, O. 2020 Electrostatic potential of a uniformly charged square plate at an arbitrary point in space. *Phys. Scr.* **95** (9), 095802.
- GARABEDIAN, P.R. 1996 Stellarators with the magnetic symmetry of a tokamak. *Phys. Plasmas* **3** (7), 2483–2485.
- GATES, D.A. 2022 A short introduction to PSI. In *64th Annual Meeting of the Division of Plasma Physics*. American Physics Society conference.
- GERHARDT, S.P., TALMADGE, J.N., CANIK, J.M. & ANDERSON, D.T. 2005 Measurements and modeling of plasma flow damping in the helically symmetric experiment. *Phys. Plasmas* **12** (5), 056116.
- GRANETZNY, M., SCHMITZ, O. & ZEPP, M. 2022 Discharge directionality and dominance of right-handed modes in helicon plasmas due to radial electron density gradients. [arXiv:2212.11401](https://arxiv.org/abs/2212.11401).
- HAMBERGER, S.M., BLACKWELL, B.D., SHARP, L.E. & SHENTON, D.B. 1990 H-1 design and construction. *Fusion Technol.* **17** (1), 123–130.
- HAMMOND, K., ZHU, C., KORRIGAN, K., GATES, D., LOWN, R., MERCURIO, R., QIAN, T. & ZARNSTORFF, M. 2022 Design of an arrangement of cubic magnets for a quasi-axisymmetric stellarator experiment. *Nucl. Fusion* **62**, 126065.
- HELANDER, P. 2014 Theory of plasma confinement in non-axisymmetric magnetic fields. *Rep. Prog. Phys.* **77** (8), 087001.
- HELANDER, P., DREVLAK, M., ZARNSTORFF, M. & COWLEY, S.C. 2020 Stellarators with permanent magnets. *Phys. Rev. Lett.* **124** (9), 095001.

- HIDALGO, C., ALEJALDRE, C., ALONSO, A., ALONSO, J., ALMOGUERA, L., DE ARAGÓN, F., ASCASÍBAR, E., BACIERO, A., BALBÍN, R., BLANCO, E., *et al.* 2005 Overview of TJ-II experiments. *Nucl. Fusion* **45** (10), S266.
- HIRSHMAN, S.P. & WHITSON, J.C. 1983 Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria. *Phys. Fluids* **26** (12), 3553–3568.
- HP, DATASHEET 2019a 3D high reusability PA 12. <https://3dprinting.com/wp-content/uploads/2019/02/4AA6-4895ENA.pdf>, accessed: 2023-06-05.
- HP, DATASHEET 2019b 3D high reusability PA 12 glass beads. <https://static1.sw-cdn.net/files/cms/materials/data-sheets/HP-MJF-PA12GB-datasheet.pdf>, accessed: 2023-06-05.
- IYOSHI, A., KOMORI, A., EJIRI, A., EMOTO, M., FUNABA, H., GOTO, M., IDA, K., IDEI, H., INAGAKI, S., KADO, S., *et al.* 1999 Overview of the large helical device project. *Nucl. Fusion* **39** (9Y), 1245–1256.
- IMBERT-GERARD, L.-M., PAUL, E.J. & WRIGHT, A.M. 2020 An introduction to stellarators: from magnetic fields to symmetries and optimization. [arXiv:1908.05360](https://arxiv.org/abs/1908.05360).
- KAPTANOGLU, A.A., CONLIN, R. & LANDREMAN, M. 2023 Greedy permanent magnet optimization. *Nucl. Fusion* **63** (3), 036016.
- KAPTANOGLU, A.A., QIAN, T., WECHSUNG, F. & LANDREMAN, M. 2022 Permanent-magnet optimization for stellarators as sparse regression. *Phys. Rev. Appl.* **18** (4), 044006.
- KATTER, M. 2005 Angular dependence of the demagnetization stability of sintered Nd-Fe-B magnets. *IEEE Trans. Magn.* **41** (10), 3853–3855.
- KINTEL, M. 2010 OpenSCAD. <https://openscad.org/index.html>, accessed: 2023-02-11.
- KU, L.-P. & BOOZER, A.H. 2009 Nonaxisymmetric shaping of tokamaks preserving quasisymmetry. *Phys. Plasmas* **16** (8), 082506.
- KUMAR, S.T.A., BLACKWELL, B.D. & HARRIS, J.H. 2009 Determination of error field sources by accurate mapping of the magnetic geometry of the H-1 heliac. *Nucl. Fusion* **49** (3), 035001.
- LANDREMAN, M. 2017 An improved current potential method for fast computation of stellarator coil shapes. *Nucl. Fusion* **57**, 046003.
- LANDREMAN, M., MEDASANI, B., WECHSUNG, F., GIULIANI, A., JORGE, R. & ZHU, C. 2021 SIMSOPT: a flexible framework for stellarator optimization. *J. Open Source Softw.* **6** (65), 3525.
- LANDREMAN, M. & PAUL, E. 2022 Magnetic fields with precise quasisymmetry for plasma confinement. *Phys. Rev. Lett.* **128** (3), 035001.
- LANDREMAN, M. & ZHU, C. 2021 Calculation of permanent magnet arrangements for stellarators: a linear least-squares method. *Plasma Phys. Control. Fusion* **63** (3), 035001.
- LAZERSON, S.A., *et al.* 2019 Tuning of the rotational transform in Wendelstein 7-X*. *Nucl. Fusion* **59** (12), 126004.
- LAZERSON, S., SCHMITT, J., ZHU, C., BRESLAU, J., ALL STELLOPT DEVELOPERS & USDOE OFFICE OF SCIENCE 2020 Stellopt, version 2.7.5.
- LIU, H., SHIMIZU, A., ISOBE, M., OKAMURA, S., NISHIMURA, S., SUZUKI, C., XU, Y., ZHANG, X., LIU, B., HUANG, J., *et al.* 2018 Magnetic configuration and modular coil design for the chinese first quasi-axisymmetric stellarator. *Plasma Fusion Res.* **13**, 3405067–3405067.
- LU, Z.Y., XU, G.S., CHEN, D.H., CHEN, L., ZHANG, X.Y., YE, M.Y. & WAN, B.N. 2021b Design of quasi-axisymmetric stellarators with variable-thickness perpendicular permanent magnets based on a two-step magnet design strategy. *Nucl. Fusion* **61** (10), 106028.
- LU, Z.Y., XU, G.S., CHEN, D., ZHANG, X.Y., CHEN, L., YE, M., GUO, H. & WAN, B. 2021a Development of advanced stellarator with identical permanent magnet blocks. *Cell Rep. Phys. Sci.* **3** (1), 100709.
- MAUEL, M.E., LEVESQUE, J.P., NAVRATIL, G.A., PAZ-SOLDAN, C.A. & PAUL, E.J. 2022 Science opportunities for HBT-EP enhanced by REBCO superconducting TF coils. In *64th Annual Meeting of the Division of Plasma Physics*. American Physics Society conference.
- MJF DESIGN HANDBOOK, HP 2019 <https://www.forecast3d.com/wp-content/uploads/2020/08/HP-MJF-Design-Handbook.pdf>, accessed: 2023-02-11.
- NEILSON, G.H., GRUBER, C.O., HARRIS, J.H., REJ, D.J., SIMMONS, R.T. & STRYKOWSKY, R.L. 2010 Lessons learned in risk management on NCSX. *IEEE Trans. Plasma Sci.* **38** (3), 320–327.

- NEMOV, V.V., KASILOV, S.V., KERNBICHLER, W. & HEYN, M.F. 1999 Evaluation of 1/V neoclassical transport in stellarators. *Phys. Plasmas* **6** (12), 4622–4632.
- NIES, R. & PAUL, E.J. 2023 In Preparation.
- NIES, R., PAUL, E.J., HUDSON, S.R. & BHATTACHARJEE, A. 2022 Adjoint methods for quasi-symmetry of vacuum fields on a surface. *J. Plasma Phys.* **88** (1), 905880106.
- NÜHRENBERG, J., LOTZ, W., MERKEL, P., NÜHRENBERG, C., SCHWENN, U., STRUMBERGER, E. & HAYASHI, T. 1995 Overview on Wendelstein 7-X theory. *Fusion Technol.* **27** (3 T), 71–78.
- OTTE, M., *et al.* 2016 Setup and initial results from the magnetic flux surface diagnostics at Wendelstein 7-X. *Plasma Phys. Control. Fusion* **58** (6), 064003.
- PARIS, K. & SOUZA, M.J. 2022 Scientific glassblowers in project MUSE: the world's first permanent magnet stellarator. *Fusion*, LXX.1, 19–24.
- PAVLICHENKO, O.S., *et al.* 1993 First results from the URAGAN-2M torsatron. *Plasma Phys. Control. Fusion* **35** (SB), B223.
- PROLL, J.H.E., MYNICK, H.E., XANTHOPOULOS, P., LAZERSON, S.A. & FABER, B.J. 2015 TEM turbulence optimisation in stellarators. *Plasma Phys. Control. Fusion* **58** (1), 014006.
- QIAN, T. 2023 MUSE PM Stellarator Design Data. <https://github.com/tmqian/Muse-Design-Paper>.
- QIAN, T., ZARNSTORFF, M., BISHOP, D., CHAMBLIS, A., DOMINGUEZ, A., PAGANO, C., PATCH, D. & ZHU, C. 2022 Simpler optimized stellarators using permanent magnets. *Nucl. Fusion* **62** (8), 084001.
- SHIMIZU, A., LIU, H., ISOBE, M., OKAMURA, S., NISHIMURA, S., SUZUKI, C., XU, Y., ZHANG, X., LIU, B., HUANG, J., *et al.* 2018 Configuration property of the chinese first quasi-axisymmetric stellarator. *Plasma Fusion Res.* **13** (0), 3403123–3403123.
- TALMADGE, J.N., SAKAGUCHI, V., ANDERSON, F.S.B., ANDERSON, D.T. & ALMAGRI, A.F. 2001 Experimental determination of the magnetic field spectrum in the Helically Symmetric Experiment using passing particle orbits. *Phys. Plasmas* **8** (12), 5165–5170.
- THOME, R.J. & TARRH, J.M. 1982 *MHD and Fusion Magnets: Field and Force Design Concepts*. Wiley.
- WECHSUNG, F., LANDREMAN, M., GIULIANI, A., CERFON, A. & STADLER, G. 2022 Precise stellarator quasi-symmetry can be achieved with electromagnetic coils. *Proc. Natl Acad. Sci.* **119** (13), e2202084119.
- WELLER, A., SAKAKIBARA, S., WATANABE, K.Y., TOI, K., GEIGER, J., ZARNSTORFF, M.C., HUDSON, S.R., REIMAN, A., WERNER, A., NÜHRENBERG, C., *et al.* 2006 Significance of MHD effects in stellarator confinement. *Fusion Sci. Technol.* **50** (2), 158–170.
- WOLF, R.C., *et al.* 2017 Major results from the first plasma campaign of the Wendelstein 7-X stellarator. *Nucl. Fusion* **57** (10), 102020.
- YAMADA, H. 2011 Overview of results from the Large Helical Device. *Nucl. Fusion* **51** (9), 094021.
- YOUNG, K.M. 1974 The C-stellarator - a review of containment. *Plasma Phys.* **16** (2), 119.
- ZARNSTORFF, M.C., BERRY, L.A., BROOKS, A., FREDRICKSON, E., FU, G.-Y., HIRSHMAN, S., HUDSON, S., KU, L.-P., LAZARUS, E., MIKKELSEN, D., *et al.* 2001 Physics of the compact advanced stellarator NCSX. *Plasma Phys. Control. Fusion* **43** (12A), A237–A249.
- ZHU, C., HAMMOND, K., BROWN, T., GATES, D., ZARNSTORFF, M., CORRIGAN, K., SIBILIA, M. & FEIBUSH, E. 2020a Topology optimization of permanent magnets for stellarators. *Nucl. Fusion* **60** (10), 106002.
- ZHU, C., HAMMOND, K., RUTKOWSKI, A., CORRIGAN, K., BISHOP, D., BROOKS, A., DUGAN, P., ELLIS, R., PERKINS, L., ZHAI, Y., *et al.* 2022 PM4Stell: a prototype permanent magnet stellarator structure. *Phys. Plasmas* **29** (11), 112501.
- ZHU, C., HUDSON, S.R., SONG, Y. & WAN, Y. 2018 Designing stellarator coils by a modified Newton method using focus. *Plasma Phys. Control. Fusion* **60** (6), 065008.
- ZHU, C., ZARNSTORFF, M., GATES, D. & BROOKS, A. 2020b Designing stellarators using perpendicular permanent magnets. *Nucl. Fusion* **60** (7), 076016.