Piecewise Omnigenous Magnetohydrodynamic Equilibria as Fusion Reactor Candidates

V. Fernández-Pacheco
Laboratorio Nacional de Fusión, CIEMAT, 28040 Madrid, Spain
and Universidad Carlos III, Madrid, Spain

J. L. Velasco,* E. Sánchez, J. M. García-Regaña, J. A. Alonso, I. Calvo, and D. Carralero
Laboratorio Nacional de Fusión, CIEMAT, 28040 Madrid, Spain

R. Gaur
Princeton University, Princeton, New Jersey, USA

(Received 11 February 2026; accepted 1 June 2026; published 17 November 2026)

In piecewise omnigenous magnetic fields, charged particles remain perfectly confined in the absence of collisions and turbulence. This concept extends the traditional notion of omnigenity, the theoretical principle upon which most of existing magnetic fusion reactor designs, including tokamaks, are based. While piecewise omnigenity broadens the range of potentially viable stellarator reactor candidates, it is achieved by relaxing the requirement of continuity in the magnetic-field strength, which could appear to pose significant challenges for the design of magnetohydrodynamic equilibria. In this work, a stellarator magnetic configuration is presented that satisfies the ideal magnetohydrodynamic equilibrium equation and that achieves unprecedented levels of piecewise omnigenity. As a result, it exhibits favorable transport characteristics, including reduced bulk radial (neoclassical and turbulent) transport, bootstrap current, and fast ion losses. In addition, the configuration displays robust magnetohydrodynamic stability across a range of β values and possesses a rotational transform profile compatible with an island divertor. Collectively, these features satisfy the standard set of physics criteria required for a viable reactor candidate which, until now, were believed to be attainable only by certain types of omnigenous stellarators.

DOI: 10.1103/jcks-2bzg

Subject Areas: Energy Research, Nuclear Physics, Plasma Physics

I. INTRODUCTION

The two leading concepts in magnetic confinement fusion, tokamaks and stellarators [1], confine the charged particles of a fusion plasma [2] by means of magnetic fields in which the field lines form nested toroidal flux surfaces; see, e.g., Fig. 1. Charged particles undergo rapid gyromotion about magnetic-field lines, characterized by a gyrofrequency and an associated gyroradius [3]. To lowest order in an expansion in the gyroperiod (the inverse of the gyrofrequency), the orbit centers, known as guiding centers, follow the magnetic-field lines. On timescales much longer than the gyroperiod, the guiding centers drift perpendicularly to the magnetic field $\mathbf{B}$, in the directions tangential and perpendicular to the flux surfaces. Cross-surface (also termed "radial") transport of energy limits the performance of a fusion reactor, and its minimization is a fundamental design criterion. One of the main channels through which energy and particles are lost is neoclassical transport, associated to the combination of radial drifts and collisions [4]. Another important channel is turbulent transport, which arises when small-scale instabilities in the plasma generate fluctuations in the electric and magnetic fields that modify particle orbits.

According to their lowest-order orbits, particles in a magnetic confinement device can be classified into two groups: passing and trapped particles, depicted in Fig. 1. For passing particles (Fig. 1, black), the component of the velocity that is parallel to $\mathbf{B}$ never vanishes, and they go over the complete flux surface. In the absence of collisions and turbulence, passing particles are always well confined, because their radial drift velocity vanishes when averaged over lowest-order orbits. The situation is, in general, different for trapped particles (Fig. 1, green and red). For trapped particles, the component of the velocity that is parallel to $\mathbf{B}$ vanishes at two points of the trajectory called "bounce points," where it changes sign (at these points, the value of the magnetic-field strength, $B = |\mathbf{B}|$, equals the ratio between the particle energy and the particle

*Contact author: joseluis.velasco@ciemat.es

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

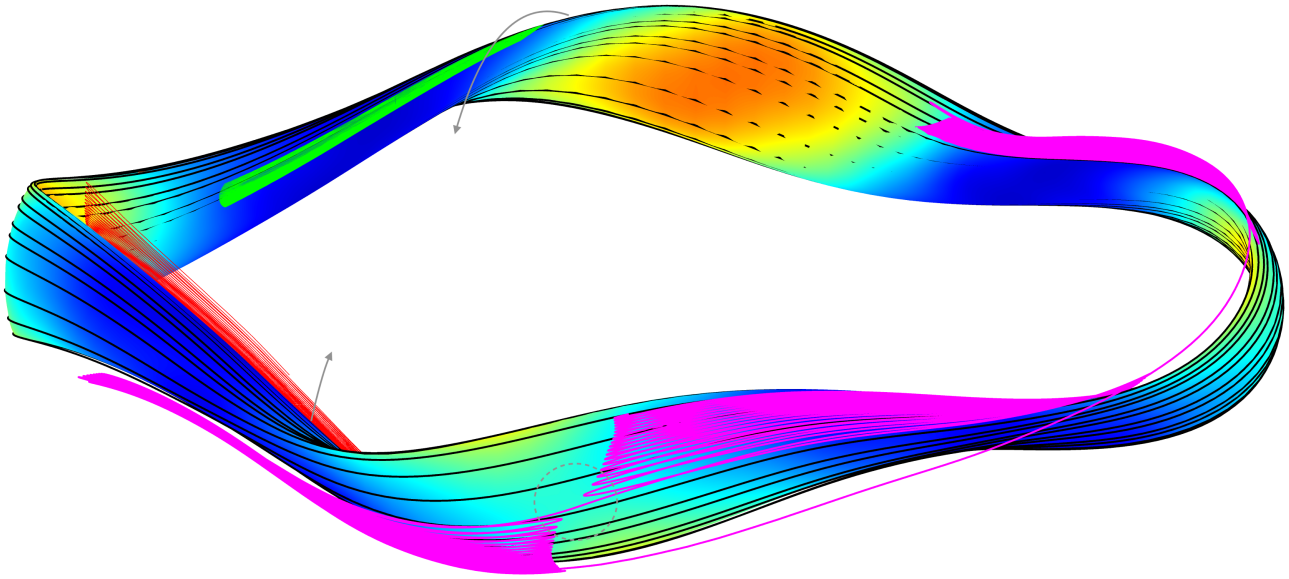


FIG. 1. Flux surface of a stellarator magnetic configuration [red (blue) color corresponds to a larger (smaller) B] with the guiding-center trajectories of a passing particle (black), an unconfined trapped particle (red), a confined trapped particle (green), and a transitioning particle (pink). Gray arrows represent the orbit-averaged drifts, and the dashed circle highlights a transition: at some point of the trajectory of a trapped particle, one of the bounce points lies on a local maximum of B; immediately after this happens, the particle transitions to a different region of the magnetic surface.

magnetic moment $\mathcal{E}/\mu$). Trapped particles move back and forth along the magnetic-field line between the bounce points while slowly drifting perpendicularly to $\mathbf{B}$. Depending on whether their radial drift velocity vanishes when averaged over lowest-order orbits (Fig. 1, green) or remains finite (Fig. 1, red), they are confined (in the absence of collisions) or they quickly leave the device, respectively. Transitioning particles (a subclass of trapped particles [5], Fig. 1, pink; see also below) visit different regions of a flux surface and may, thus, spend part of their trajectory moving on the flux surface and part drifting away from it.

Although the stellarator is considered to have been the first magnetic confinement fusion device [6], the tokamak has been the leading concept since, in the late 1960s, it demonstrated a performance far greater than any other concept [7]. This is a consequence of its axial symmetry, which ensures good confinement of collisionless particles, and, thus, reduced neoclassical transport. Such symmetry does not exist in the stellarator, which produces its magnetic field through a three-dimensional shaping of its magnetic surfaces. While this eliminates the requirement of a large inductive current and, thus, facilitates the operation of the stellarator with respect to that of the tokamak, it generally limits its performance [1]. Specifically, low-collisionality neoclassical transport is intolerably large in a generic stellarator.

It was not until the 1980s [8] that it was discovered that, if the flux surface of the stellarator is mapped with the appropriate set of toroidal and poloidal coordinates, θ and ζ, termed "Boozer coordinates," the guiding-center trajectories are determined only by the spatial variation of B on the surface. Consequently, if B can be made independent of the toroidal angle as in a tokamak (Fig. 2, top left), these orbits are isomorphic to those in a tokamak and, thus, well confined. Such stellarator magnetic fields are termed "quasiaxisymmetric." Good confinement is also the case, more generally, if the spatial dependence of B on the flux surface can be written as a function of $M\theta - N\zeta$, N and M being two integer numbers. The first quasisymmetric (QS) magnetohydrodynamic (MHD) equilibria (of the quasihelically symmetric family, the one with N and M different from zero, Fig. 2, center left) were obtained numerically several years later [9]. In the 1990s, Cary and Shasharina [10] generalized the concept of quasisymmetry by proposing magnetic fields with collisionless confinement of trapped particles without the requirement that all the B contours have to be straight lines in Boozer coordinates. The key point is that, even in the absence of a symmetry direction, trapped particles moving in a magnetic field still display a conserved quantity, the second adiabatic invariant

$$J \equiv 2\int_{l_{b_1}}^{l_{b_2}} \sqrt{2(\mathcal{E}-\mu B)}\,\mathrm{d}l. \tag{1}$$

Here, l is the arclength of field lines, and l_{b_1} and l_{b_2} the bounce-point positions. Cary and Shasharina prescribed the two conditions on the variation of B on the flux surface for a magnetic field to be omnigenous. First, all contours of constant B have to be curves that close in the toroidal,

poloidal, or helical direction (a feature that eliminates transitioning particles; see Fig. 1). Additionally, the distance along the field line between consecutive bounce points must not depend on the field line for a given flux surface. When this happens, the contours of constant J are aligned with the flux surfaces. Since J is a conserved quantity, this implies that the particles must drift, on average, on the flux surface. Around that date, approximately quasi-isodynamic (QI) fields (i.e., omnigenous fields with poloidally closed B contours, Fig. 2, bottom left) were proposed as reactor candidates [11]. QI fields have the additional advantage of displaying zero bootstrap current at low collisionality [12], which makes them compatible with an island divertor [13] for the management of heat and particle exhaust.

Quasisymmetry and quasi-isodynamicity were promptly employed as design targets, but the first examples of QS and QI configurations presented significant deviations from the ideal limit, and so did the first stellarators ever designed or built according to these criteria, at the beginning of this century: the National Compact Stellarator eXperiment (NCSX), the Helically Symmetric eXperiment (HSX), and Wendelstein 7-X (W7-X); see Fig. 2 (right). In all

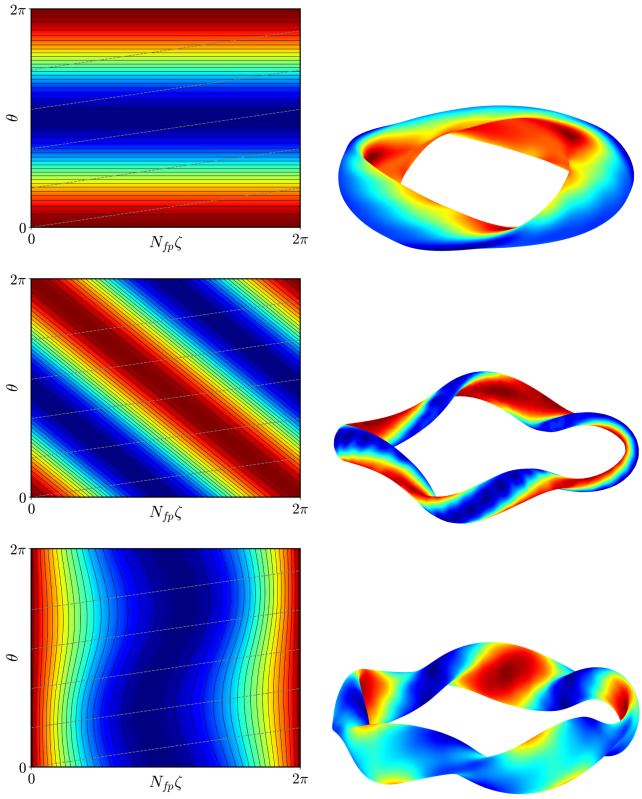


FIG. 2. B on a flux surface of an axisymmetric tokamak or an exactly quasiaxisymmetric stellarator (top left), of an exactly QS field with helical symmetry (center left), and of an exactly QI field (bottom left); B on the boundary of NCSX (top right), HSX (center right), and the high-mirror configuration of W7-X (bottom right). Dashed gray lines represent magnetic-field lines.

cases, violations of the constraints imposed in B by omnigenity and/or quasisymmetry were prominent (note, for instance, the coil ripple in HSX or how the contours of high B do not close poloidally in W7-X). As a consequence of this, the neoclassical transport of these devices (particularly that of energetic particles such as fusion-born alpha particles, which are more affected by these imperfections) is too large. Only recently, in this decade, have theoretical and computational advances allowed for configurations with very precise levels of omnigenity [14–17]. Nevertheless, precise omnigenity should not be expected to be compatible with other reactor relevant criteria. For instance, it is not a guarantee of good energetic particle confinement—see, e.g., Ref. [15]—as this last property is known to depend critically on the radial variation of B [18,19]. Thus, in parallel with these conceptual studies, other, more comprehensive and balanced optimization strategies were being pursued. Finally, almost half a century after the first attempts at stellarator design, this effort has yielded a growing family of candidates for a reactor [20–26]. Most of them are of the QI type, one of the main reasons being compatibility with an island divertor; see, e.g., Ref. [27].

The theoretical development of piecewise omnigenity, on the other hand, is very recent [28–30]. It generalizes the concept of omnigenity by realizing that the latter can be fulfilled *piecewisely* on the flux surface. Rather than J being a flux-surface constant for a given particle velocity, as in an omnigenous field, the flux surface of a piecewise omnigenous (pwO) field can be divided into several regions, within which J can take different constant values. This eliminates the requirement that all the B contours close in a particular direction (i.e., the first of the two conditions prescribed by Cary and Shasharina [10]), thus not necessarily constraining the fields to a limited and discrete set of possible *helicities*, as illustrated by the dashed lines in Fig. 3 (top). Additionally to tokamaklike transport of the bulk species, pwO fields can display zero bootstrap current at low collisionality for arbitrary plasma gradients [30], similarly to QI fields. Two examples of pwO fields are depicted in Fig. 3 (bottom). While Fig. 3 (bottom, left) represents a pwO field that is relatively close to quasi-isodynamicity (deeply trapped particles encounter B contours that close poloidally), Fig. 3 (bottom, right) represents a pwO field completely away from omnigenity. Both cases display wide regions of constant B, but it should be noted that smoother examples exist [29].

Although piecewise omnigenity is a promising concept, there nevertheless exists a tremendous gap between the level of maturity of pwO fields and that of QI or QS fields in the path toward the fusion reactor. Only a very limited set of configurations have been identified that present a spatial variation of B that is reminiscent of that corresponding to piecewise omnigenity [17,28,31–33]. Because they are rough, oftentimes coincidental, approximations to piecewise omnigenity, some of them present a too large

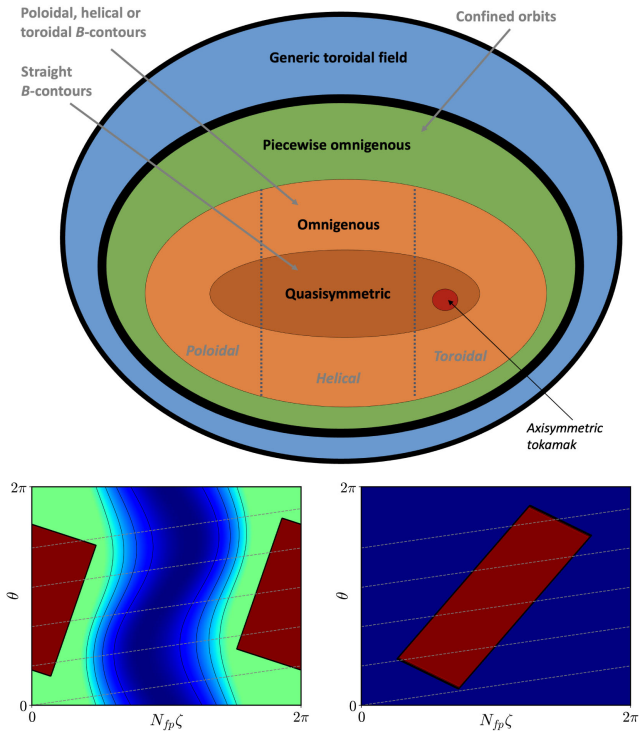


FIG. 3. Piecewise omnigenity in the space of stellarator magnetic fields: diagram depicting the different families of optimized stellarator fields (top) and illustrating that quasisymmetric fields are the subset of omnigenous fields with straight B contours, while piecewise omnigenous fields constitute a broader family obtained by relaxing the constraint of toroidally, helically, or poloidally closed B contours. B on a flux surface of an exactly pwO field close to quasi-isodynamicity (bottom, left) and an exactly pwO field far from omnigenity (bottom, right).

neoclassical transport (in the absence of the notion of piecewise omnigenity, the term quasiomnigenous [31] was occasionally used, despite lacking a precise definition). Moreover, many of them do not fulfill basic reactor criteria such as MHD stability. In this work, we close this gap by presenting a stellarator magnetic configuration explicitly designed to satisfy, thanks to unprecedented levels of piecewise omnigenity, the standard set of physics criteria usually required for a viable reactor candidate [20,21].

II. OPTIMIZATION WITH RESPECT TO PIECEWISE OMNIGENITY

As illustrated by Fig. 3 (bottom, left), piecewise omnigenity is a broader family than the fields explored in Ref. [28] and depicted in Fig. 3 (bottom, right), and examples exist with a smoother spatial variation of B [29]. However, with the goal of assessing the reactor viability of the concept beyond any doubt, in this work we target pwO fields that can be considered as far as possible from omnigenity. As in Ref. [28], we parametrize the variation of B on the flux surface of a pwO field as

$$B_{\mathrm{pwO}}(\theta,\zeta) = B_{\min} + (B_{\max} - B_{\min}) \times e^{-([\zeta-\zeta_c+t_1(\theta-\theta_c)]/w_1)^{2p}-([\theta-\theta_c+t_2(\zeta-\zeta_c)]/w_2)^{2p}}, \tag{2}$$

$$w_1 = \frac{\pi}{N_{\mathrm{fp}}}\frac{1-t_1 t_2}{1+t_2/\iota}, \tag{3}$$

$$w_2 = \pi. \tag{4}$$

Here, $B_{\max}$ and $B_{\min}$ are, respectively, the maximum and minimum values of B on the flux surface, ι is the rotational transform, which determines the direction of the field lines on the flux surface, and N_{fp} is the number of field periods, $B(\theta,\zeta) = B(\theta,\zeta+2\pi/N_{\mathrm{fp}})$. In the limit $p\to\infty$, all the B contours of Eq. (2) collapse into a parallelogram: $B = B_{\max}$ inside it and $B = B_{\min}$ outside it. The rest of the quantities in Eq. (2) parametrize the shape of this parallelogram. First, ζ_c and θ_c determine the position of its center. Then, w_1 and w_2 specify its toroidal and poloidal extension: w_1 is set in Eq. (3) (as a function of other equilibrium quantities, such as ι, that will change during the optimization process) by the requirement of collisionless confinement of trapped particles [28]; w_2 [see Eq. (4)] is set by the requirement of zero bootstrap current [30]. Finally, t_1 and t_2 determine the slope of the two different sides of the parallelogram ($-1/t_1$ and $-t_2$, respectively).

Equation (3) points at a critical dependence of the property of piecewise omnigenity on the value of ι. However, for the family of pwO fields presented in Ref. [28] and targeted in this work, the B contours are composed of straight lines. For this reason, we do not expect a large degradation of the transport properties to be caused by a moderate change in ι due, e.g., to the bootstrap current. In a field like the one in Fig. 3 (bottom, right), two neighboring orbits that originally shared a value of J will, very likely, share a (different) value of J after the change of ι, thus keeping $\partial_\alpha J = 0$ (here, $\alpha = \theta - \iota\zeta$ is employed to label field lines on a flux surface, and, therefore, $\partial_\alpha J$ denotes the variation of J within the flux surface). Only the orbits close to the juncture between regions will start to cause transport, if one of the bounce points *jumps* to a different parallelogram segment. Another way to understand this expected robustness against ι changes is that, within each region of the flux surface, trapped particles behave *as if* they were moving in a QS field (with N/M equal to $1/t_1$ or t_2, depending on the region); the property of zero effective ripple of QS fields, differently to what happens to more general omnigenous fields, does not depend explicitly on ι; see, e.g., Ref. [34].

The MHD equilibrium of a stellarator is determined in fixed-boundary simulations by the plasma pressure profile (in this work, the normalized toroidal flux $s = \psi/\psi_{\mathrm{boundary}}$ is employed as the radial coordinate, and the pressure profile is set to be linear, with normalized volume-averaged

pressure $\beta = 0.5\%$), by the toroidal current through the plasma (which we set to zero), and by the shape of the equilibrium boundary. Stellarator optimization is typically carried out by continuously modifying the shape of the boundary surface in order to explore (a region of) the stellarator configuration space. The optimizer performs such exploration by trying to minimize a cost function in which a set of desired properties have been encoded. Ideally, if a minimum of the cost function is found, the corresponding stellarator equilibrium is close to having the desired properties.

In this work, we have developed specific routines in the code DESC [16] devoted to minimizing the distance of the configuration magnetic-field strength, at a specific flux surface $s = s_0$, to the closest field described by Eqs. (2)–(4). Specifically, the quantity

$$\delta B(\theta,\zeta) = \frac{|B(s=s_0,\zeta,\theta) - B_{\mathrm{pwO}}(\theta,\zeta)|}{B_{\mathrm{pwO}}(\theta,\zeta)} \tag{5}$$

is targeted to be zero. During the optimization, in addition to the shape of the $s = 1$ flux surface, some of the parameters in Eq. (2) are varied. We fix $N_{\mathrm{fp}} = 5$, $\zeta_c = \pi/N_p$, $\theta_c = 0$ (a natural choice given the initial condition; see Appendix A), $s_0 = 0.5$, and $p = 2$. The latter choice implies that, rather than aiming for a precise pwO field, we will actually be aiming for an approximately pwO field that we expect to have (respective to the $p \to \infty$ case) a smoother spatial variation and a nonzero but sufficiently low level of neoclassical transport (both energy flux and bootstrap current) for reactor-relevant collisionalities. The other parameters, $B_{\max}$, $B_{\min}$, t_1, and t_2 (and ι), are left to vary in order not to unnecessarily overconstrain the optimization problem.

In this work, the quantity $\delta B(\theta,\zeta)$ is incorporated into a cost function together with other reactor-relevant physics properties of the MHD equilibrium. Ideally, the result of the optimization process, once the optimizer has minimized the cost function, should be a magnetic configuration that is close to piecewise omnigenity at $s = s_0$ (with the consequent reduction of neoclassical transport) and, at the same time, satisfies the rest of desired reactor-relevant physics criteria. Achievement of small δB is demonstrated in Sec. III A, and the fulfillment of relevant reactor design criteria, such as an appropriate ι profile, is the focus of Sec. III B.

There are, however, some physics criteria whose achievement may be facilitated (although not guaranteed) by optimization with respect to small δB, and they appear in both sections. After all, properties such as MHD stability, turbulence, and fast ion confinement can be straightforwardly connected, to some extent, to the spatial variation of B, both on the flux surface and in the radial direction (in particular, as in the case of omnigenity, good fast ion confinement is significantly facilitated but not guaranteed by closeness to piecewise omnigenity). In the rest of this section, we discuss to what extent optimization with respect to piecewise omnigenity can benefit these aspects of reactor performance.

First, it should be noted that, in this work, piecewise omnigenity has been targeted at a single flux surface, $s_0 = 0.5$. Similarly to what happens in omnigenous fields, closeness to piecewise omnigenity is expected to degrade as one separates from the optimized flux surface [this advises against trying to make $\delta B(\theta,\zeta)$ very small on two different flux surfaces]. In particular, the Fourier spectrum of B_{pwO} is quite broad, and B_{mn} modes with different m number are expected to have a different radial variation since, at least sufficiently close to the axis, $B_{mn} \sim s^{m/2}$ [here, $B_{\mathrm{pwO}}(\theta,\zeta) = \sum_{m,n} B_{m,n}\cos(m\theta - nN_p\zeta)$]. It is shown in Sec. III A that this degradation is weak enough so that the benefits obtained from optimizing at $s = 0.5$ extend to the full plasma volume.

Let us then assume that the magnetic field of a nearly pwO configuration can be approximately written as

$$\begin{aligned} B_{\mathrm{pwO}}(s,\theta,\zeta) &= B_{\min}(s) + [B_{\max}(s) - B_{\min}(s)] \\ &\times \lim_{p\to\infty} e^{-([\zeta-\zeta_c+t_1(\theta-\theta_c)]/w_1)^{2p} - ([\theta-\theta_c+t_2(\zeta-\zeta_c)]/w_2)^{2p}}, \end{aligned} \tag{6}$$

with w_1 and w_2 set by Eqs. (3) and (4) and all the parameters in Eq. (6) (and ι) kept constant throughout the plasma volume except $B_{\max}$ and $B_{\min}$. Then, if $\partial_s B_{\min}$ and $\partial_s B_{\max}$ have the same sign and are both sufficiently large (in absolute value), a enhanced trapped-particle precession on the flux surface, at a stronger rate than in the radial direction, is expected to contribute to good alpha-particle confinement [18,19]. Moreover, if $\partial_s B_{\max}$ and $\partial_s B_{\min}$ are both positive, trapped electrons precess in the direction opposite to the electron diamagnetic direction; this mitigates turbulence originated by trapped-electron modes driven by density gradients and other ion-scale gyrokinetic instabilities involving kinetic electrons [35–38]. This can be understood in terms of the second adiabatic invariant: Straightforward calculations analogous to those of Sec. III. 5 in Ref. [29] show that, for the magnetic field of Eq. (6), if $\partial_s B_{\min}$ and $\partial_s B_{\max}$ are positive,

$$J = J^{(\mathrm{w})}(s,\mathcal{E},\mu), \quad \partial_s J^{(\mathrm{w})} < 0, \quad \mathrm{w} = \mathrm{I},\mathrm{II},\mathrm{III},\ldots, \tag{7}$$

where w labels the different regions of the flux surface within which J is constant. That is, the so-called maximum-J property [35,36] is fulfilled *piecewisely*. Through the connection between orbit-averaged drifts and the spatial variation of J, it is straightforward to see that Eq. (7) guarantees that both trapped ions and trapped electrons are collisionlessly confined and that the latter precess in the direction that mitigates said turbulence. In this work, we try to satisfy Eq. (7) by minimizing $\delta B(\theta,\zeta)$ and

simultaneously achieving $\partial_s B_{\max} > 0$ and $\partial_s B_{\min} > 0$. As we see in Sec. III A, the latter requirement is met by approximately pwO fields at large β thanks to the diamagnetic effect, without the need for a dedicated optimization target. In Sec. III B, the beneficial effect of this approach toward the maximum-J property will be tested with gyrokinetic simulations.

Finally, we note that the maximum-J property has been connected (through the effect of the magnetic well) to MHD stability; see, e.g., Ref. [39].

III. RESULTS

Figure 4 presents two different views of the magnetic-field strength on the $s = 1$ flux surface of an optimized MHD equilibrium that we call CIEMAT-pw1. Its aspect ratio (the ratio between the major and minor radius) is 12.7, and its maximum elongation (at toroidal section $\phi = \pi/N_{\rm fp}$; see Fig. 4, top) is 6.8.

A. CIEMAT-pw1, a magnetic configuration close to piecewise omnigenity

The level of piecewise omnigenity of CIEMAT-pw1 is assessed in Fig. 5: The top row compares $B(\theta, \zeta)$ at the radial position $s = 0.5$ of the optimized equilibrium with $B_{\rm pwO}(\theta, \zeta)$ for $p = 2$ at the end of the optimization process. The relative difference is below 1%, and it is of the order of 0.1% for most of the flux surface. Additionally, the optimized equilibrium is compared with the $p \to \infty$ limit of $B_{\rm pwO}(\theta, \zeta)$ (although this relative difference was not an optimization target, it is strictly speaking the quantity that determines closeness to piecewise omnigenity) and is shown to be larger, around 10%, at specific regions of the flux surface. These values, which are larger than those obtained so far for omnigenous fields [14–17], could likely be made smaller by relaxing some of the other optimization criteria. Nevertheless, in a nonacademic scenario, higher levels of piecewise omnigenity would likely be spoiled by finite β effects or by error fields introduced by the coils (see Sec. II in Ref. [40] for an omnigenous example). As predicted, Fig. 5 (bottom) shows a degradation of the level of piecewise omnigenity with changes in β and, specifically, with the distance to $s = 0.5$. In particular, because of the relative importance of the $m = 0$ modes, the magnetic field at the position closest to the axis [Fig. 5 (bottom left)] resembles the nearly-QI pwO field shown in Fig. 3 (bottom left).

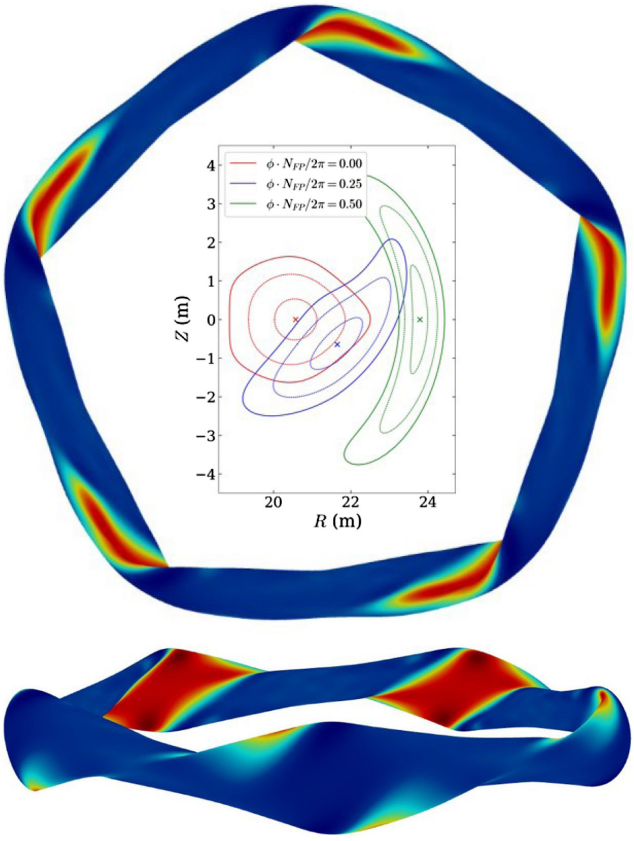


FIG. 4. Top and side views (top and bottom figures, respectively) of the magnetic-field strength on the $s = 1$ flux surface of CIEMAT-pw1 at $\beta = 0.5\%$, with the shape of several toroidal cuts of the flux surfaces (R, Z, and ϕ are cylindrical coordinates).

We next show that this level of piecewise omnigenity is appropriate to ensure robustly good neoclassical transport properties for CIEMAT-pw1. Figure 6 shows the radial profile of the effective ripple $\varepsilon_{\rm eff}$ (top) and the collisionality dependence of the D_{31} bootstrap transport coefficient (bottom). These quantities (the latter computed with MONKES [41]) encapsulate the dependence of neoclassical transport in the radial and parallel directions, respectively, on the details of the magnetic configuration [42]. Comparison with W7-X shows that these quantities are significantly small, as expected [28,30]. The effective ripple is small at $s = 0.5$, the optimized flux surface, for all β values and decreases even further toward the magnetic axis. It increases toward the edge but always below that of the standard configuration of W7-X. The bootstrap transport coefficient is presented for the most adverse situation, zero radial electric field (where larger values are expected [42,43]) and $s = 0.06$ (far from s_0 and, additionally, where the collisionality is lowest for any realistic reactor scenario), and it displays values comparable to the high-mirror configuration of W7-X. The latter has a bootstrap current that, although larger than that of more recent QI reactor candidates (see, e.g., Ref. [26]), is considered small enough for a reactor.

Finally, Fig. 7 (top) shows that, for sufficiently high β, both $\partial_s B_{\min}$ and $\partial_s B_{\max}$ are positive up to $s = s_0$ in CIEMAT-pw1. This, as discussed in Sec. II, will contribute to achieving the maximum-J property (see also Refs. [39,40]). Indeed, the fast-ion proxy Γ_c [18,19], which is, in fact, a phase-space average of $|\partial_\alpha J/\partial_s J|^2$, displays values that decrease with β and that are significantly lower than for W7-X.

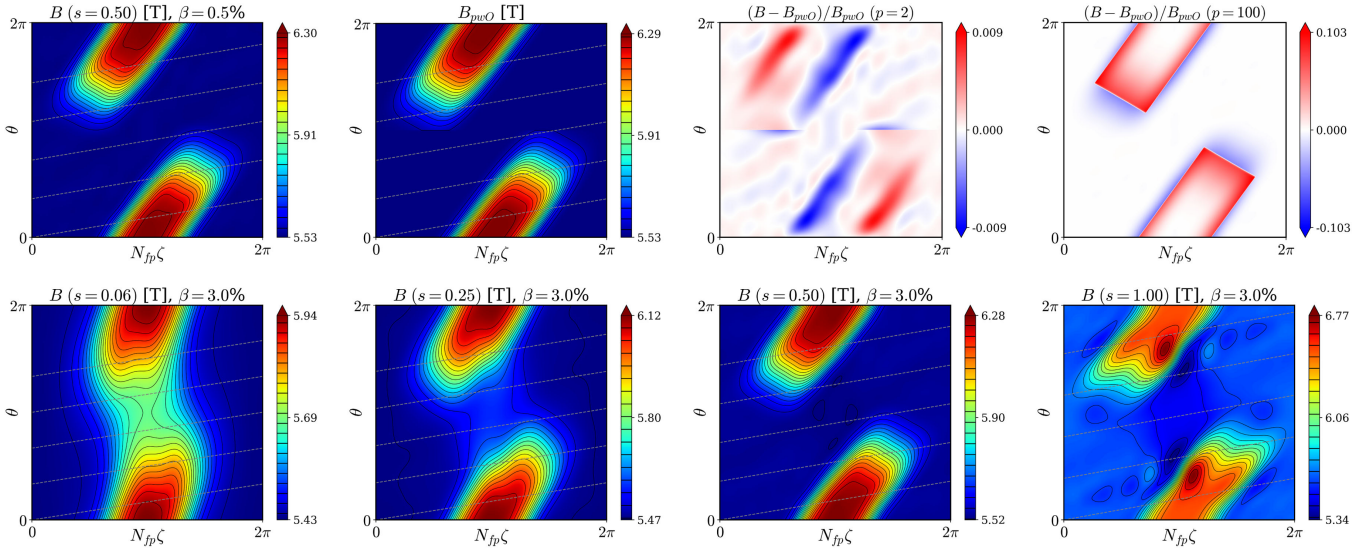


FIG. 5. Closeness of CIEMAT-pw1 to piecewise omnigenity. Top, from left to right: $B(\theta,\zeta)$ for $\beta = 0.5\%$ at the radial position $s = 0.5$ of the optimized equilibrium; $B_{\mathrm{pwO}}(\theta,\zeta)$ at the last iteration of the optimization; relative difference between them, and $p \to \infty$ limit of the relative difference. Bottom, from left to right: $B(\theta,\zeta)$ for $\beta = 3.0\%$ at radial positions $s = 0.06$, $s = 0.25$, $s = 0.50$, and $s = 1.00$.

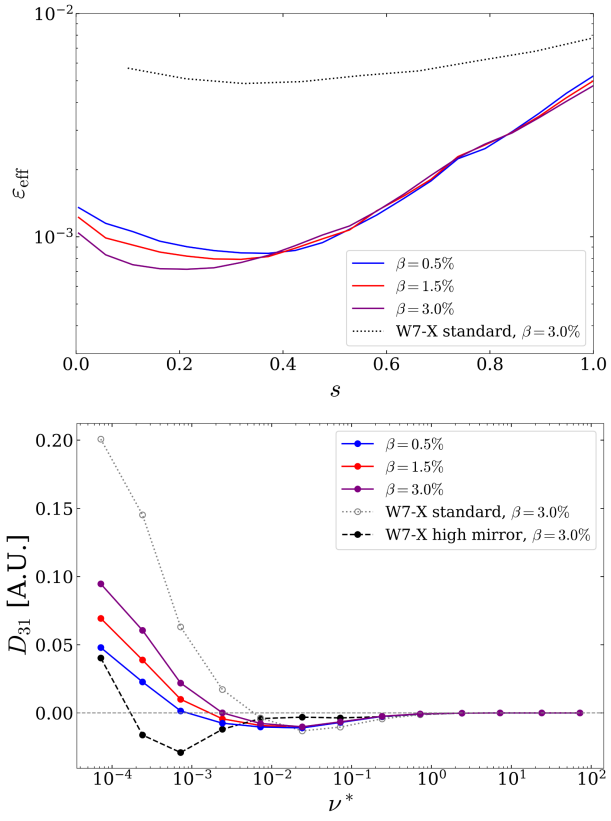


FIG. 6. Neoclassical transport of CIEMAT-pw1: radial profile of the effective ripple (top) and collisionality dependence of the bootstrap transport coefficient at $s = 0.06$ (bottom).

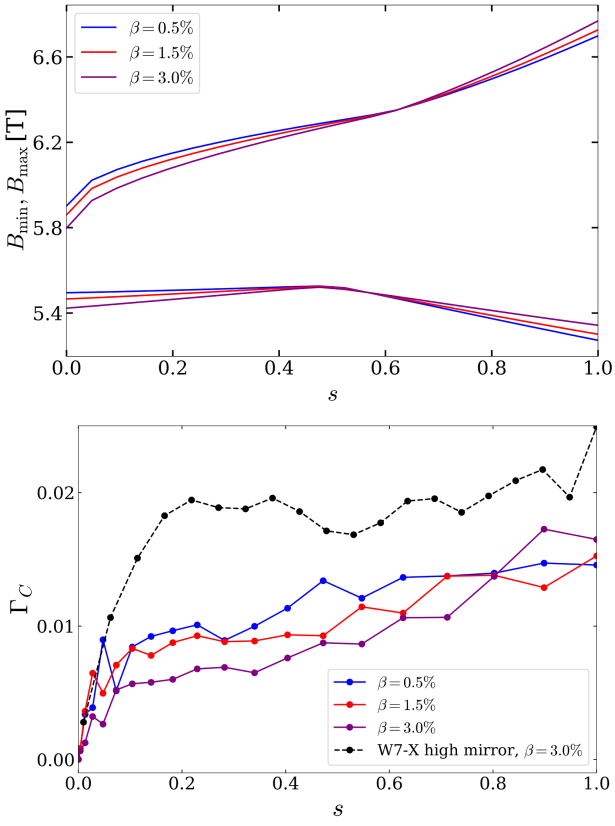


FIG. 7. Optimization of CIEMAT-pw1 with respect to the maximum-J criterion: radial profiles of B_{min} and B_{max} (top) and of the Γ_c figure of merit.

B. Compatibility of CIEMAT-pw1 with reactor physics design criteria

A viable reactor candidate needs to satisfy a set of physics criteria that extend well beyond reduced bulk neoclassical transport, ranging from fundamental requirements to others that are considerably more advanced. Among the former, the magnetic configuration must be magnetohydrodynamically stable and possess good flux surfaces throughout the plasma volume. Among the latter, turbulent transport and alpha-particle losses must remain sufficiently low to reach and sustain the reactor operating point. Finally, the magnetic configuration must be compatible with the chosen divertor concept. These physics requirements were addressed during the optimization of CIEMAT-pw1 (see Appendix A for the specific set of optimization targets). In this section, we analyze, for a range of plasma β, to what extent they were achieved. Where appropriate, this is evaluated by comparison with the properties of W7-X, as it is relatively well established where its stellarator configuration stands in relation to a viable fusion reactor.

The Mercier criterion is employed to assess MHD stability in CIEMAT-pw1, as shown in Fig. 8. The configuration is expected to be stable, since $D_{\mathrm{Mercier}} > 0$ in most of the plasma radius (except at the highest β values very close to the core, where simulations sometimes become unreliable). The increase of D_{Mercier} with β is caused by the magnetic well, the main contributor to stability, which overcomes the destabilizing effect of the geodesic curvature. Ballooning stability has been confirmed by means of calculations with the code COBRA [44], with negative growth rates for all the cases considered.

The rotational transform profile of CIEMAT-pw1 is shown in Fig. 9. It has been designed with the general objective of avoiding low-order rational values in the core, as they could lead to the formation of magnetic islands [45], and to have the $5/5$ rational value just outside the boundary, which could be the basis of an island divertor at $s > 1$ ($s = 1$ then would become the last closed flux surface). In the absence of a plasma current, ι grows approximately linearly with s (with a similar slope to the one that achieved, for the particular case of CIEMAT-QI4X [26], an optimal balance between avoiding island overlap at the regions of $s < 1$, where rational values could be prevented, and permitting a large magnetic island at $s > 1$). In order to assess the effect of the bootstrap current on the rotational transform profile, a reactor physics design point is identified by scaling the magnetic configuration to a major radius of $R = 18.5$ m and a magnetic-field strength on axis of $B_{\mathrm{axis}} = 5.5$ T (keeping $\beta = 3\%$). Using the tools devised in Ref. [46] and similar assumptions (in terms, e.g., of fusion gain and energy confinement time scaling) the plasma profiles in Fig. 9 (bottom) are obtained (quasineutrality between electron, deuterium, and tritium species is imposed by setting $n_e = 2n_D = 2n_T$). Then, neoclassical transport is computed with SFINCS [47], and the resulting bootstrap current is employed to recalculate the $\beta = 3\%$ equilibrium (given the uncertainties in the reactor plasma profiles, which would be reduced only by higher-fidelity modeling—e.g., Ref. [48]—we do not attempt to obtain self-consistency and, rather, limit ourselves to illustrating the size of the change in ι that is to be expected, along with its impact on confinement properties). Compatibility with

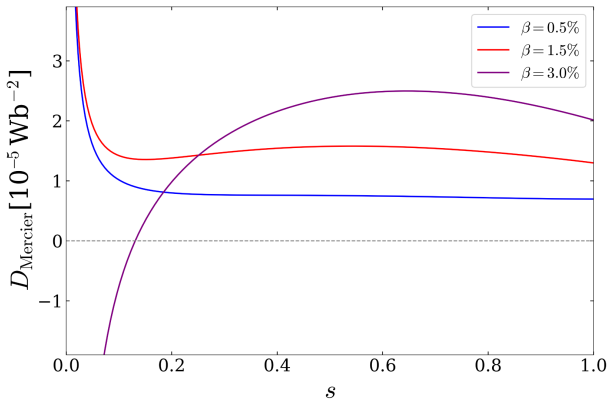


FIG. 8. Mercier stability criterion.

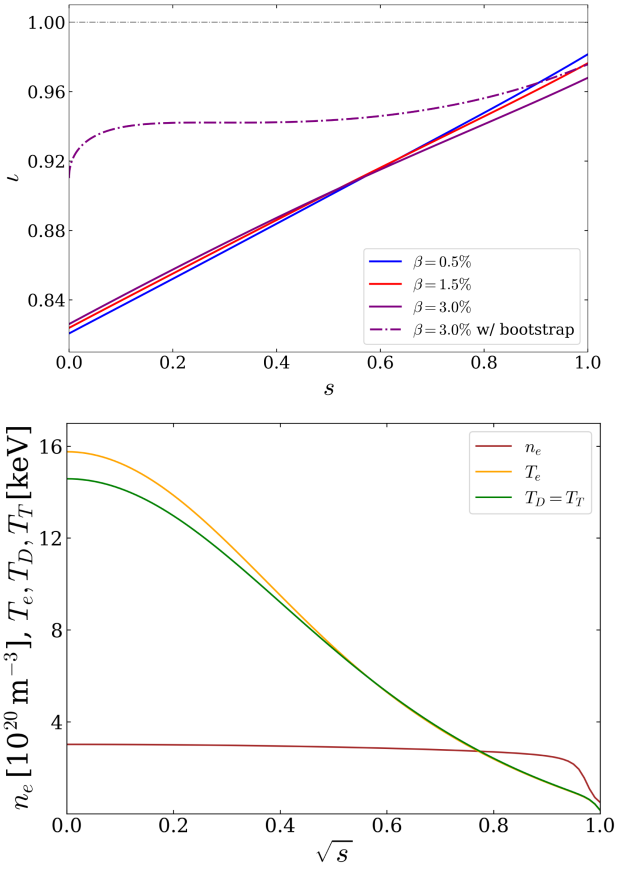


FIG. 9. Rotational transform: radial profile of ι (top) and plasma profiles of the reactor scenario employed for the bootstrap calculation (bottom).

an island divertor is supported by the fact that, in Fig. 9 (top), the change in ι at the plasma edge is small. Close to the axis, where the collisionality is smaller and the bootstrap current level higher, the change is larger and goes in the direction of avoiding the 5/6 rational surface. We confirm below that a change in ι of the size of that in Fig. 9 (top) does not cause a large degradation of the transport properties.

Since the effective ripple is approximately one order of magnitude below that of W7-X (already considered sufficiently small for a reactor), turbulent transport is expected to be the main channel for energy loss of the bulk plasma in CIEMAT-pw1. In this work, it is computed with the gyrokinetic code stella [49], by means of flux-tube simulations with three kinetic species: deuterium, tritium, and electrons. Although electromagnetic effects will likely become important at the β values of a reactor scenario, this type of simulations has been insufficiently validated by experiments in stellarators; we, thus, follow the example of Refs. [21,23,50] and perform electrostatic simulations. Our results should then be considered as a baseline level for turbulent transport, which will then be likely modified (either reduced or increased) by the presence of electromagnetic effects [51], fast particles [52], and/or impurities [53], to name a few effects. For reference, reactor scenarios typically require turbulent heat fluxes $Q/Q_{i,gB}$ between $O(10^{-2})$ (at the core) and $O(10^{-1})$ (toward the edge) [46]. Here, Q is normalized by $Q_{i,gB}$, the gyro-Bohm ion transport level. Obtaining values in this range with our gyrokinetc simulations, as will be the case, should then be considered a necessary, but not sufficient, condition for reactor relevance.

Figure 10 (top) shows the turbulent energy flux at $s = 0.36$ as a function of the normalized temperature gradient a/L_T for several species and values of the normalized density gradient a/L_n (in these simulations, $T_e = T_D = T_T$ and $n_e = 2n_D = 2n_T$). We observe that the deuterium and tritium fluxes are very close, and the electron flux has the same qualitative behavior. By comparing the simulations for $a/L_n = 1$ and $a/L_n = 4$ at large a/L_T, we can confirm the stabilization expected from fulfilling, piecewisely, the maximum-J property. On the other hand, for a small density gradient, a scenario that was not addressed in our optimization strategy, the energy fluxes grow rapidly with the temperature gradient. However, the values remain small (compatible with zero) up to a moderate value of a/L_T: Following the procedure in Ref. [54], a critical gradient $a/L_{T_{\mathrm{crit}}}$ between 1.5 and 2 can be roughly estimated. This value lies between that of W7-X and that of magnetic configurations specifically optimized for this aspect in the referenced work. Along the same line, when comparing the points with $a/L_n = 1$ and $a/L_T = 3$ to those from existing devices (ASDEX Upgrade, TJ-II, LHD, NCSX, and the W7-X standard configuration; see Ref. [55], Figs. 22 and 23), the configuration shows

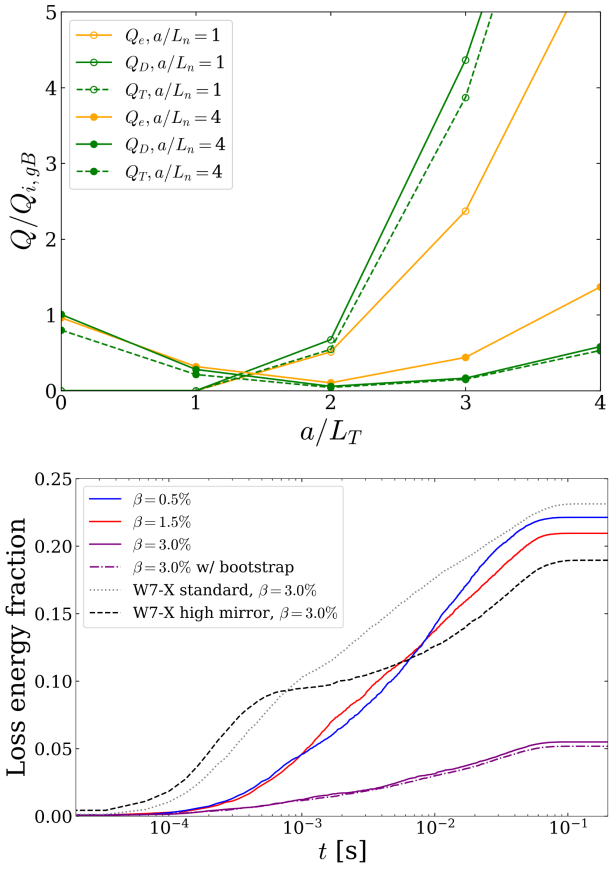


FIG. 10. Consequences of the optimization with respect to the maximum-J criterion: turbulent fluxes of energy for the $\beta = 3\%$ configuration at $s = 0.36$ as a function of the plasma gradients (top) and fraction of alpha particle energy lost through the boundary (bottom).

consistently similar or superior performance, at least for $\beta = 3\%$. Altogether, these preliminary results indicate that pwO configurations are amenable to turbulence optimization.

Finally, the confinement of alpha particles of CIEMAT-pw1 is estimated with the ASCOT code [56] and shown in Fig. 10 (bottom). With an initial distribution that takes into account a realistic fusion reaction rate, the guiding-center trajectories of alpha particles are followed, including collisions with the background, until they have thermalized or have been lost through $s = 1$ (in order to isolate configuration effects and to allow for a comparison with the literature, we employ for the reaction rate and for collisions of all cases the same profiles of density and temperature, taken from Ref. [57]). For $\beta = 3\%$, the optimized equilibrium exhibits an alpha heating efficiency of 95%, compatible with the assumptions of the reactor scenario. Moreover, even at low β, most of the lost alpha particles escape the plasma after they have transferred a fraction of their energy to the plasma (the slowing-down time is approximately 0.1 s), which limits the potential damage to plasma facing components (indeed, the Γ_c proxy

is considered a better predictor of prompt collisionless losses than of total losses [19,58]). Finally, we note that the fast ion confinement of the equilibrium, one of the transport properties most sensitive to details of the magnetic geometry, is not affected by the change in ι caused by the bootstrap current.

IV. DISCUSSION

In summary, in this work, we have radically advanced a new concept, piecewise omnigenity, for the confinement of magnetic fusion plasmas. Starting from a theoretical concept illustrated by a limited set of incomplete examples, we have demonstrated that piecewise omnigenity is a criterion that can actually be employed to design stellarator magnetic configurations that comply with the standard set of physics criteria required for a viable reactor candidate. The magnetic configuration CIEMAT-pw1 is presented in this work as a first example. The exploration of the configuration space of pwO fields has just been initiated, and better candidates, possibly combined with omnigenous fields [29], will likely be found.

A key question, which has to be answered in the new line of investigation that this work opens, is to what extent pwO fields provide additional specific advantages with respect to omnigenous fields. To this respect, the resilience of the transport properties of CIEMAT-pw1 to ι changes presents an interesting example. From the physics points of view, mechanisms such as turbulence stabilization by magnetic shear (see, e.g., Refs. [59,60]) that are, in principle, not viable for QI or QS fields, could be so for pwO fields (in the sense of not being deleterious to neoclassical transport). Additionally, more compact configurations may become accessible by combining piecewise omnigenity with quasi-isodynamicity [61,62] or quasiaxisymmetry [63]. From the technological point of view, robustness against changes in the magnetic configuration is connected to the problem of coil complexity, a major driver of reactor cost. While this work focuses on physics properties and leaves technological details such as coils or compatibility with a breeder blanket for future work, Appendix B shows some preliminary promising results, in line with partial results presented in Ref. [17]. This is a very relevant matter that requires dedicated investigation.

ACKNOWLEDGMENTS

The authors are grateful to EUROfusion's TSVV12 team for its valuable feedback, to the DESC team, and to all the developers of the codes VMEC, COBRA, MONKES, SFINCS, stella, and ASCOT. Calculations for this work made use of computational resources at Xula (CIEMAT), Turgalium (CETA-CIEMAT), and Marenostrum V (Barcelona Supercomputing Center). This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No. 101052200 EUROfusion). This research was supported by Grants No. PID2021-123175NB-I00 and No. PID2024-155558OB-I00, Ministerio de Ciencia, Innovación y Universidades, Spain. This work was conducted as part of the FUSION-EP Master's Programme, and V. F.-P. gratefully acknowledges the program for its support.

Conceptualization: V. F.-P., J. L. V., E. S., and I. C. Data curation: V. F.-P., J. L. V., E. S., and J. M. G.-R. Formal analysis: V. F.-P., J. L. V., E. S., J. M. G.-R., J. A. A., and D. C. Funding acquisition: J. L. V., J. M. G.-R., and I. C. Investigation: V. F.-P., J. L. V., E. S., R. G., J. M. G.-R., J. A. A., and D. C. Methodology: V. F.-P., J. L. V., E. S., R. G., J. M. G.-R., J. A. A., I. C., and D. C. Project administration: V. F.-P., J. L. V., and E. S. Software: V. F.-P. and R. G. Supervision: J. L. V. and E. S. Validation: J. L. V. and E. S. Visualization: V. F.-P., J. L. V., and E. S. Writing—original draft: J. L. V., V. F.-P., and E. S. Writing—review and editing: R. G., J. M. G.-R., J. A. A., I. C., and D. C.

Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

DATA AVAILABILITY

The data supporting the findings of this article are openly available [64].

APPENDIX A: OPTIMIZATION AND SIMULATION DETAILS

Table I provides a list of the different targets employed in the optimization process with DESC, along with some motivation (precise definitions of the targets as implemented in DESC can be found in a Zenodo repository [64]). We note that the general motivation for the ι target is given in Sec. III B, but similar $\iota(s)$ functions would have likely yielded similarly good results. The aspect ratio was monitored but not included as an explicit optimization target. As initial condition, the boundary of configuration A in Ref. [32] was changed to $N_{\mathrm{fp}} = 5$. The Zenodo repository [64] also includes scripts and inputs that allow to reproduce the main calculations of this work with the appropriate resolution.

APPENDIX B: COIL FEASIBILITY

Once the physical properties of the equilibrium have been analyzed, the next step would be to look for a set of magnetic coils that are able to reproduce the obtained magnetic field without compromising its favorable properties (this would, in turn, allow for a detailed design of the edge island structure, so that it can be employed for heat and particle exhaust). This is usually considered an independent

TABLE I. Objectives used for optimization with their target or bounds and purpose.

Objective	Type	Target/bounds	Purpose
Force balance residual	Volumetric	0	Accurate MHD equilibrium
Toroidal current	Profile	0	Stellarator constraint
Major radius	Scalar	Same as initial configuration	No trivial optimization
B on axis	Scalar	Same as initial configuration	No trivial optimization
2nd principle curvature of boundary	Set of scalars	$>(-100, 10)$ (m^{-1})	Simple boundary shape
Maximum elongation	Scalar	7	Simple boundary shape
Rotational transform	Profile	$\iota = 0.82 + 0.16s$ (w/o bootstrap)	Good flux surfaces, compatibility with island divertor
Mercier criterion	Profile	$>10^{-4}$ Wb	MHD stability
Mirror ratio $(B_{\max} - B_{\min})/(B_{\max} + B_{\min})$ at $s = 1$	Scalar	<0.15	Reduced fraction of trapped particles (potential drive of transport)
Deviation from B_{pwO} at $s = 0.5$	Surface	0	Collisionless confinement of orbits

optimization problem (see, e.g., Refs. [65,66]) and is not addressed in this work. However, we use a proxy, the magnetic gradient scale length $L_{\nabla B}$ [67], to obtain a first, albeit rough, indication of coil feasibility. This quantity provides an estimate of the minimum distance at which a set of reasonably simple modular coils should be placed in order to reproduce correctly the magnetic configuration. The criterion of magnetic-field fidelity tends to push the coils to be close to the plasma (otherwise, they may become very complicated if they need to generate high harmonics of B). On the other hand, in a reactor, the plasma-coil distance must exceed a certain threshold in order to have space for the neutron shielding required by the coils and for a breeding blanket able to produce tritium at an adequate rate (these considerations would naturally be less pressing for a non-nuclear experimental device). A compromise needs to be found.

Figure 11 shows $L_{\nabla B}$ at flux surface $s = 1$ of the magnetic configuration. From these results, the minimum plasma-to-coil distance can be estimated to be approximately 1.3 m (details of the model, which should be considered a rough estimate, are given in Ref. [67]). This distance may be somewhat tight to accommodate the required components in a reactor. However, the predicted distance is much larger in other regions, which might compensate for this and provide an adequate tritium breeding ratio. An assessment of this requires the design of coils and neutronics analyses that are left for future work.

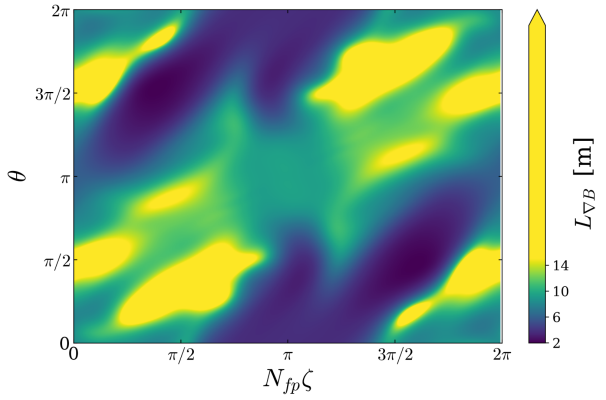


FIG. 11. Estimate of coil complexity.

[1] P. Helander, C. D. Beidler, T. M. Bird, M. Drevlak, Y. Feng, R. Hatzky, F. Jenko, R. Kleiber, J. H. E. Proll, Y. Turkin, and P. Xanthopoulos, *Stellarator and tokamak plasmas: A comparison*, Plasma Phys. Controlled Fusion **54**, 124009 (2012).
[2] J. D. Lawson, *Some criteria for a power producing thermonuclear reactor*, Proc. Phys. Soc. London Sect. B **70**, 6 (1957).
[3] T. G. Northrop, *The guiding center approximation to charged particle motion*, Ann. Phys. (N.Y.) **15**, 79 (1961).
[4] P. Helander and D. J. Sigmar, *Collisional Transport in Magnetized Plasmas* (Cambridge University Press, Cambridge, England, 2002), Vol. 87.
[5] J. R. Cary, D. F. Escande, and J. L. Tennyson, *Adiabatic-invariant change due to separatrix crossing*, Phys. Rev. A **34**, 4256 (1986).
[6] L. Spitzer, *The stellarator concept*, Phys. Fluids **1**, 253 (1958).
[7] L. Arcimovich, G. Bobrovskij, E. Gorbunov, D. Ivanov, V. Kirillov, J. Kuznecov, S. Mirnov, M. Petrov, K. Razumova, V. Strelkov, and D. Tcheglov, *Experiments in Tokamak Devices* (IAEA, Trieste, 1969).
[8] A. H. Boozer, *Transport and isomorphic equilibria*, Phys. Fluids **26**, 496 (1983).
[9] J. Nührenberg and R. Zille, *Quasi-helically symmetric toroidal stellarators*, Phys. Lett. A **129**, 113 (1988).
[10] J. R. Cary and S. G. Shasharina, *Helical plasma confinement devices with good confinement properties*, Phys. Rev. Lett. **78**, 674 (1997).
[11] H. Wobig, *The theoretical basis of a drift-optimized stellarator reactor*, Plasma Phys. Controlled Fusion **35**, 903 (1993).
[12] P. Helander and J. Nührenberg, *Bootstrap current and neoclassical transport in quasi-isodynamic stellarators*, Plasma Phys. Controlled Fusion **51**, 055004 (2009).

[13] T. S. Pedersen *et al.*, *First divertor physics studies in Wendelstein 7-X*, Nucl. Fusion **59**, 096014 (2019).
[14] M. Landreman and E. Paul, *Magnetic fields with precise quasisymmetry for plasma confinement*, Phys. Rev. Lett. **128**, 035001 (2022).
[15] A. G. Goodman, K. Camacho-Mata, S. A. Henneberg, R. Jorge, M. Landreman, G. Plunk, H. M. Smith, R. Mackenbach, C. Beidler, and P. Helander, *Constructing precisely quasi-isodynamic magnetic fields*, J. Plasma Phys. **89**, 905890504 (2023).
[16] D. W. Dudt, A. G. Goodman, R. Conlin, D. Panici, and E. Kolemen, *Magnetic fields with general omnigenity*, J. Plasma Phys. **90**, 905900120 (2024).
[17] H. Liu, G. Yu, C. Zhu, J. L. Velasco, R. Gaur, D. Panici, E. Kolemen, M. Yu, W. Ding, S. Wang, and G. Zhuang, *Optimizing stellarators with hidden symmetry*, arXiv: 2502.09350.
[18] V. V. Nemov, S. V. Kasilov, W. Kernbichler, and G. O. Leitold, *Poloidal motion of trapped particle orbits in real-space coordinates*, Phys. Plasmas **15**, 052501 (2008).
[19] J. L. Velasco, I. Calvo, S. Mulas, E. Sánchez, F. Parra, and Á. Cappa (The W7-X Team Collaboration), *A model for the fast evaluation of prompt losses of energetic ions in stellarators*, Nucl. Fusion **61**, 116059 (2021).
[20] E. Sánchez, J. Velasco, I. Calvo, and S. Mulas, *A quasi-isodynamic configuration with good confinement of fast ions at low plasma β*, Nucl. Fusion **63**, 066037 (2023).
[21] A. G. Goodman, P. Xanthopoulos, G. G. Plunk, H. Smith, C. Nührenberg, C. D. Beidler, S. A. Henneberg, G. Roberg-Clark, M. Drevlak, and P. Helander, *Quasi-isodynamic stellarators with low turbulence as fusion reactor candidates*, PRX Energy **3**, 023010 (2024).
[22] C. Hegna, D. Anderson, E. Andrew, A. Ayilaran, A. Bader, T. Bohm, K. C. Mata, J. Canik, L. Carbajal, A. Cerfon *et al.*, *The infinity two fusion pilot plant baseline plasma physics design*, J. Plasma Phys. **91**, E76 (2025).
[23] J. Lion *et al.*, *Stellaris: A high-field quasi-isodynamic stellarator for a prototypical fusion power plant*, Fusion Eng. Des. **214**, 114868 (2025).
[24] A. G. Goodman, G. G. Plunk, P. Xanthopoulos, M. Drevlak, J. Geiger, R. Davies, H. M. Smith, C. Nührenberg, C. D. Beidler, S. A. Henneberg *et al.*, J. Plasma Phys. **91**, E153 (2025).
[25] C. P. Swanson, *Overview of the helios design: A practical planar coil stellarator fusion power plant*, arXiv:2512.08027.
[26] E. Sánchez, J. Velasco, I. Calvo, J. García-Regaña, C. Salcuni, and J. Alonso, *CIEMAT-QI4X: A reactor-relevant quasi-isodynamic stellarator configuration compatible with an island divertor*, Nucl. Fusion **66**, 056008 (2026),
[27] A. Bader, A. Ayilaran, J. Canik, A. De, W. Guttenfelder, C. Hegna, M. Knilans, A. Malkus, T. Pedersen, P. Sinha *et al.*, *Power and particle exhaust for the infinity two fusion pilot plant*, J. Plasma Phys. **91**, E67 (2025).
[28] J. L. Velasco, I. Calvo, F. J. Escoto, E. Sánchez, H. Thienpondt, and F. I. Parra, *Piecewise omnigenous stellarators*, Phys. Rev. Lett. **133**, 185101 (2024).
[29] J. Velasco, E. Sánchez, and I. Calvo, *Exploration of the parameter space of piecewise omnigenous stellarator magnetic fields*, Nucl. Fusion **65**, 056012 (2025).
[30] I. Calvo, J. L. Velasco, P. Helander, and F. I. Parra, *Piecewise omnigenous stellarators with zero bootstrap current*, Phys. Rev. E **112**, L023201 (2025).
[31] D. A. Spong, S. P. Hirshman, J. C. Whitson, D. B. Batchelor, B. A. Carreras, V. E. Lynch, and J. A. Rome, *J* optimization of small aspect ratio stellarator/tokamak hybrid devices*, Phys. Plasmas **5**, 1752 (1998).
[32] D. Bindel, M. Landreman, and M. Padidar, *Direct optimization of fast-ion confinement in stellarators*, Plasma Phys. Controlled Fusion **65**, 065012 (2023).
[33] R. Gaur, D. Panici, T. Elder, M. Landreman, K. Unalmis, Y. Elmacioglu, D. Dudt, R. Conlin, and E. Kolemen, *Omnigenous umbilic stellarators*, J. Plasma Phys. **91**, E151 (2025).
[34] M. Landreman and P. J. Catto, *Omnigenity as generalized quasisymmetry*, Phys. Plasmas **19**, 056103 (2012).
[35] M. N. Rosenbluth, *Low-frequency limit of interchange instability*, Phys. Fluids **11**, 869 (1968).
[36] P. Helander, J. H. E. Proll, and G. G. Plunk, *Collisionless microinstabilities in stellarators. I. Analytical theory of trapped-particle modes*, Phys. Plasmas **20**, 122505 (2013).
[37] G. G. Plunk, J. W. Connor, and P. Helander, *Collisionless microinstabilities in stellarators. Part 4. The ion-driven trapped-electron mode*, J. Plasma Phys. **83**, 715830404 (2017).
[38] J. Proll, G. Plunk, B. Faber, T. Görler, P. Helander, I. McKinney, M. Pueschel, H. Smith, and P. Xanthopoulos, *Turbulence mitigation in maximum-J stellarators with electron-density gradient*, J. Plasma Phys. **88**, 905880112 (2022).
[39] E. Rodríguez, P. Helander, and A. Goodman, *The maximum-J property in quasi-isodynamic stellarators*, J. Plasma Phys. **90**, 905900212 (2024).
[40] J. L. Velasco, I. Calvo, E. Sánchez, and F. Parra, *Robust stellarator optimization via flat mirror magnetic fields*, Nucl. Fusion **63**, 126038 (2023).
[41] F. J. Escoto, J. L. Velasco, I. Calvo, M. Landreman, and F. I. Parra, *MONKES: A fast neoclassical code for the evaluation of monoenergetic transport coefficients in stellarator plasma*, Nucl. Fusion **64**, 076030 (2024).
[42] C. D. Beidler, K. Allmaier, M. Y. Isaev, S. V. Kasilov, W. Kernbichler, G. O. Leitold, H. Maaßberg, D. R. Mikkelsen, S. Murakami, M. Schmidt, D. A. Spong, V. Tribaldos, and A. Wakasa, *Benchmarking of the mono-energetic transport coefficients. Results from the International Collaboration on Neoclassical Transport in Stellarators (ICNTS)*, Nucl. Fusion **51**, 076001 (2011).
[43] C. G. Albert, C. D. Beidler, G. Kapper, S. V. Kasilov, and W. Kernbichler, *On the convergence of bootstrap current to the Shaing–Callen limit in stellarators*, J. Plasma Phys. **91**, E77 (2025).
[44] R. Sanchez, S. P. Hirshman, A. S. Ware, L. A. Berry, and D. A. Spong, *Ballooning stability optimization of low-aspect-ratio stellarators*, Plasma Phys. Controlled Fusion **42**, 641 (2000).
[45] F. Waelbroeck, *Theory and observations of magnetic islands*, Nucl. Fusion **49**, 104025 (2009).
[46] J. Alonso, I. Calvo, D. Carralero, J. L. Velasco, J. García-Regaña, I. Palermo, and D. Rapisarda, *Physics design point*

of high-field stellarator reactors, Nucl. Fusion **62**, 036024 (2022).

[47] M. Landreman, H. Smith, A. Mollén, and P. Helander, *Comparison of particle trajectories and collision operators for collisional transport in nonaxisymmetric plasmas*, Phys. Plasmas **21**, 042503 (2014).

[48] A. B. Navarro, A. D. Siena, J. Velasco, F. Wilms, G. Merlo, T. Windisch, L. LoDestro, J. Parker, and F. Jenko, *First-principles based plasma profile predictions for optimized stellarators*, Nucl. Fusion **63**, 054003 (2023).

[49] M. Barnes, F. Parra, and M. Landreman, *stella: An operator-split, implicit–explicit delta-f-gyrokinetic code for general magnetic field configurations*, J. Comput. Phys. **391**, 365 (2019).

[50] J. García-Regaña, I. Calvo, E. Sánchez, H. Thienpondt, J. Velasco, and J. Capitán, *Reduced electrostatic turbulence in the quasi-isodynamic stellarator configuration CIEMAT-QI4*, Nucl. Fusion **65**, 016036 (2024).

[51] P. Mulholland, K. Aleynikova, B. J. Faber, M. J. Pueschel, J. H. E. Proll, C. C. Hegna, P. W. Terry, and C. Nührenberg, *Enhanced transport at high plasma pressure and subthreshold kinetic ballooning modes in Wendelstein 7-X*, Phys. Rev. Lett. **131**, 185101 (2023).

[52] A. Di Siena, A. Bañón Navarro, and F. Jenko, *Turbulence suppression by energetic particle effects in modern optimized stellarators*, Phys. Rev. Lett. **125**, 105002 (2020).

[53] J. M. García-Regaña, I. Calvo, F. I. Parra, and H. Thienpondt, *Reduction or enhancement of stellarator turbulence by impurities*, Phys. Rev. Lett. **133**, 105101 (2024).

[54] G. T. Roberg-Clark, *Critical gradient optimisation for quasi-isodynamic stellarators*, J. Plasma Phys. **92**, e21 (2026).

[55] H. Thienpondt, J. García-Regaña, I. Calvo, G. Acton, and M. Barnes, *Influence of the density gradient on turbulent heat transport at ion-scales: An inter-machine study with the gyrokinetic code stella*, Nucl. Fusion **65**, 016062 (2024).

[56] S. Äkäslompolo *et al.*, *Validating the ASCOT modelling of NBI fast ions in Wendelstein 7-X stellarator*, J. Instrum. **14**, C10012 (2019).

[57] M. Landreman, S. Buller, and M. Drevlak, *Optimization of quasi-symmetric stellarators with self-consistent bootstrap current and energetic particle confinement*, Phys. Plasmas **29**, 082501 (2022).

[58] E. Paul, A. Bhattacharjee, M. Landreman, D. Alex, J. Velasco, and R. Nies, *Energetic particle loss mechanisms in reactor-scale equilibria close to quasisymmetry*, Nucl. Fusion **62**, 126054 (2022),

[59] C. Kessel, J. Manickam, G. Rewoldt, and W. M. Tang, *Improved plasma performance in tokamaks with negative magnetic shear*, Phys. Rev. Lett. **72**, 1212 (1994).

[60] M. Nadeem, T. Rafiq, and M. Persson, *Local magnetic shear and drift waves in stellarators*, Phys. Plasmas **8**, 4375 (2001).

[61] J L Velasco *et al.*, *Combination of quasi-isodynamic and piecewise omnigenous magnetic fields*, arXiv:2603.12377.

[62] H. Liu and *et al.*, *Optimization of stellarator configurations combining omnigenity and piecewise omnigenity*, arXiv:2603.12139.

[63] J. L. Velasco *et al.*, *Control of the bootstrap current in approximately quasi-axisymmetric magnetic fields*, arXiv:2603.20125.

[64] V. Fernández-Pacheco *et al.*, Zenodo, 2026, https://zenodo.org/records/20390611.

[65] P F Gil *et al.*, *Augmented Lagrangian methods produce cutting-edge magnetic coils for stellarator fusion reactors*, arXiv:2507.12681.

[66] A Veksler *et al.*, *Stellarator island divertor shape optimization for reduced peak heat fluxes*, arXiv:2602.24049.

[67] J. Kappel, M. Landreman, and D. Malhotra, *The magnetic gradient scale length explains why certain plasmas require close external magnetic coils*, Plasma Phys. Controlled Fusion **66**, 025018 (2024).