
Magnetic Confinement via Dual Counter-Rotating Tungsten-Lined Plate Discs

Geometry, governing equations and design inputs for the dual rotating plate disc configuration, with three of its published quantities recomputed from the formulas that define them.

Thomas Garren

SOPHIA XT LLC

Source document 2025-11-14. Rebuilt and re-audited 2026-09-19. · Design study, revision 2

sophiaxt.com

CC BY 4.0. Code referenced herein is Apache 2.0.

Abstract

This document specifies a magnetic confinement configuration in which two counter-rotating tungsten-lined conducting discs generate opposing fields at slightly different angular velocities, compressing flux surfaces between them and stabilising resistive modes through the resulting sheared poloidal flow. It gives the geometry, the governing extended-magnetohydrodynamic system, the intended discretisation and the design inputs in full: a plasma of major radius 0.85 m and minor radius 0.40 m between 1.2 m tungsten-lined discs at an axial separation of 0.25 m, driven at up to $2\pi \times 1500$ rad/s. Three of the quantities the source document publishes are defined by an equation whose inputs the same document supplies, so each one can be recomputed, and section 6 carries those recomputations. The skin depth reproduces: 2.35 mm against a quoted 2.3 mm, which puts the field-penetration depth alongside the 3 mm tungsten armour, as the design intends. The volume-averaged beta does not: the stated definition with the stated pressure and field gives 0.40% against a printed 8.3%. The magnetic compression ratio does not: the stated formula with the stated rotation rates gives 1.005 against a printed 3.2, and the quoted field amplification from 3.5 T to 11.2 T restates that same factor rather than confirming it independently. Two of the document's headline numbers therefore fail a check against the document's own equations, and this rebuild says which and by how much. No substitute value is printed for either, because which input is wrong is not determinable from the source. The work is a design specification and a simulation plan. Every quantity in it is calculated from a stated formula, assumed as a design input, or attributed by the source document to its own solver; section 8 classifies each one by provenance and section 9 states the scope of what the study supports.

1. Context

Conventional tokamaks face divertor erosion, tritium retention and capital cost at gigawatt scale. Stellarators such as Wendelstein 7-X avoid the disruption problem but require large numbers of non-planar superconducting coils held to roughly millimetre geometric precision. Compact alternatives, including field-reversed configurations and spheromaks, reduce engineering complexity but have historically been limited in achievable plasma pressure and confinement time by edge instabilities and limited flux availability. The Princeton Field-Reversed Configuration experiments demonstrated current drive and heating by rotating magnetic fields in metre-scale devices.

The configuration specified here, referred to below as the dual rotating plate disc or DRPD configuration, generates the dynamic field by mechanically rotating conductive plates rather than by a static coil set, and obtains flow shear from the difference between the two rotation rates rather than from external momentum injection. The tungsten lining follows published work on tungsten monoblocks for the ITER divertor, which have been tested to heat loads of order 20 MW/m².

Everything in the preceding two paragraphs about other devices is published work by other groups and is cited in the references. Everything about the DRPD configuration is a proposal.

2. Theoretical framework

2.1 Field generation by rotating conductors

Each disc is specified as a copper-chromium-zirconium substrate, CuCrZr, with an assumed conductivity $\sigma \approx 46$ MS/m at 300°C, lined with 3 mm of tungsten armour, rotating about a common axis at angular velocities ω_1 and ω_2 . The induced field at a point r from the disc surface is written from the Biot-Savart law applied to the rotating conductor:

$$B(r, t) = \text{integral_disc} \left(\mu_0 / 4 \pi \right) \left[\sigma \left(\mathbf{v} \times \mathbf{E} \right) \times \mathbf{rhat} / r^2 \right] dV$$

$\mathbf{v}$ = $\omega \times \mathbf{r}'$ local velocity of the conducting material
 $\mathbf{E}$ induced field from the time-varying flux
 μ_0 = $4 \pi \times 10^{-7}$ H/m

The tungsten lining modifies the boundary condition for field penetration through the skin depth

$$\delta = \sqrt{2 / (\mu_0 \sigma \omega)}$$

At the specified $\omega \approx 2\pi \times 1000$ rad/s and $\sigma = 46$ MS/m this evaluates to 2.35 mm, which the source document quotes as 2.3 mm. This is the one quantity in the document that reproduces exactly from its own stated inputs. The design intent is that the field is confined within the copper substrate while the tungsten protects against plasma erosion, which

requires the armour thickness and the skin depth to be comparable; at 3 mm of armour against a 2.35 mm skin depth they are.

2.2 Differential-velocity compression

The mechanism maintains a small fractional difference between the two rotation rates, creating a sheared rotational transform that compresses flux surfaces radially inward. The source defines a compression ratio

$$CR = \left(1 + (d\omega / \omega_0)^2 * (r/a)^4 \right) ^ \gamma$$

ω_0 = $(\omega_1 + \omega_2) / 2$ mean rotation rate
 a = 0.40 m minor radius (assumed)
 γ = 5/3 adiabatic index, deuterium

with design inputs $\omega_1 = 2\pi \times 1050$ rad/s and $\omega_2 = 2\pi \times 950$ rad/s, from which it claims $CR \approx 3.2$ at $r = 0.3$ m and a field amplification from an assumed $B_0 = 3.5$ T to $B_{\max} = 11.2$ T in the compression zone. Section 6 shows that this formula and these inputs give 1.005.

$$(2r / Bz^2)(dP/dr) + (r/q)(dq/dr) + \alpha_{\text{flow}} (r / v_{\text{theta}})(d v_{\text{theta}} / dr) > 0$$

$q(r)$ safety factor profile
 α_{flow} 0.85, taken from published NSTX work
 v_{theta} poloidal flow velocity induced by magnetic drag

The source reports that the differential rotation term contributes +0.42 to the left-hand side at $r/a = 0.75$ with a pressure gradient $dp/dr = -2.1 \times 10^5$ Pa/m. The structure of the criterion is standard and the flow-shear coefficient is taken from published work on another device. The +0.42 itself comes out of the solver rather than from a closed form, and section 8 records it under that provenance.

$$\begin{aligned} d(\rho)/dt &+ \text{div}(\rho \mathbf{v}) = 0 \\ d(\rho \mathbf{v})/dt &+ \text{div}[\rho \mathbf{v} \mathbf{v} + P \mathbf{I} + (B^2 / 2 \mu_0) \mathbf{I} - B \mathbf{B} / \mu_0] = \mathbf{F}_{\text{ext}} \\ dB/dt &= \text{curl}(\mathbf{v} \times \mathbf{B}) + \eta \text{lap } B - (1/n_e) \text{curl}[\mathbf{J} \times \mathbf{B}] \\ dU/dt &+ \text{div}[(U + P) \mathbf{v}] = Q_{\text{fusion}} - Q_{\text{rad}} - \text{div } \mathbf{q} \end{aligned}$$

The stated discretisation is a toroidal grid of $256 \times 128 \times 64$ points in the radial, poloidal and toroidal directions, with adaptive refinement near the disc surfaces and the compression zone reaching a minimum spacing of 0.8 mm. The domain runs from the magnetic axis at $R = 0.85$ m to a boundary defined by the separatrix intersecting the tungsten surface. Resistivity follows the Spitzer model with neoclassical corrections. Radiation losses include line radiation from tungsten impurities at concentrations up to $n_W/n_e = 10^{-5}$.

The specification stops short of five things a reader would need in order to reproduce any output of it: the time integration

$$Y(E_{\text{ion}}) = 0.6 * S_n(\text{eps}) / [1 + 6 \text{eps}^{0.5} + 9 \text{eps}^{1.2}]$$

$\text{eps} = E_{\text{ion}} / E_{\text{TF}}$
 $E_{\text{TF}} = 281$ eV Thomas-Fermi energy, deuterium on tungsten
 S_n nuclear stopping power

From this the source derives a tungsten influx of order $10^{18} \text{ m}^{-2} \text{ s}^{-1}$ under assumed edge conditions of $T_e \approx 50$ eV and $n_e \approx 10^{20} \text{ m}^{-3}$. The formula is published and the edge conditions are assumed design inputs, so section 8 classifies the influx as calculated from assumptions.

2.3 Stability from sheared flow

The stabilising claim is that differential rotation suppresses resistive tearing modes. The criterion used is the Suydam criterion with an added flow-shear term:

3. Numerical method, as specified

The source specifies a three-dimensional extended magnetohydrodynamic model with two-fluid effects and anisotropic thermal conduction, solving the following normalised system:

scheme and its order, the timestep or the CFL condition governing it, the divergence-cleaning method for B , the numerical treatment of the rotating boundaries, and a grid-convergence study. Section 10 lists them as the first things to publish.

3.1 Material and surface modelling

The tungsten-copper interface is modelled as a multi-layer thermal resistance network with an assumed contact resistance of $2.5 \times 10^{-5} \text{ m}^2 \text{ K/W}$, a value typical of plasma-sprayed coatings. Impurity sputtering follows the Bohdansky yield formula:

4. Proposed apparatus

The table below is the machine the configuration would require: what would have to exist for it to be tested, with every number a design input rather than an as-built dimension.

Table 1. Design inputs for the proposed apparatus. Every row is a design choice, and the basis column says which choices follow from another row and which are free.

Design input	Value	Basis
Major radius R_0	0.85 m	assumed
Minor radius a	0.40 m	assumed
Aspect ratio A		2.1 follows from R_0 and a
Elongation κ		1.8 assumed
Disc diameter	1.2 m	assumed
Axial gap between discs	0.25 m, discs at $Z = \pm 0.125$ m	assumed
Disc mass	847 kg each	assumed
Maximum rotation	$2\pi \times 1500$ rad/s, 14,300 rpm	assumed
Drive power	1.2 MW per disc, steady state	assumed
Substrate	CuCrZr, $\sigma \approx 46$ MS/m at 300°C	published material property
Armour	3 mm tungsten	assumed
Heating	3.5 MW at 84 MHz, plus 200 kW neutral beam at 30 keV	assumed, patterned on published RMF drive systems




Figure 1. Schematic cross-section of the proposed configuration: two tungsten-lined discs counter-rotating about a common axis inside a vacuum vessel, with the intended compression region between them and indicative diagnostic port positions. This is a labelled drawing of the design, not output from any solver. The rotation rates shown on it are the design inputs of Table 1. Redrawn from the source document's vector artwork by scripts/papers/make_mc_fig1.py; every label and coordinate is read out of that artwork rather than off a screenshot of it.

5. Configuration reported by the specified simulation

The values in this section are what the source document reports as output of the magnetohydrodynamic model of

section 3. The source publishes the model and none of the artifacts that would let a reader recompute them: there is no log, no input deck, no output file, no mesh and no post-processing script, and the code it names is not available for inspection. They are carried here under a single provenance class, source simulation, which section 8 defines.

Table 2. Quantities the source attributes to its own solver. The status column is the point of the table: every entry carries its provenance, and two of them fail a recomputation from the document's own equations.

Reported quantity	Value	Status
Central electron density n_{e0}	$2.4 \times 10^{21} \text{ m}^{-3}$	source simulation
Peak electron temperature T_{e0}	12.3 keV	source simulation
Ion temperature T_i	11.1 keV	source simulation
Plasma current I_p	1.25 MA	source simulation
Energy confinement time τ_E	2.34 s	source simulation
Volume-averaged β	8.3%	does not reproduce, see section 6
Magnetic compression ratio CR	3.2	does not reproduce, see section 6
Tungsten impurity fraction n_W/n_e	3.2×10^{-6}	source simulation
Heat flux to disc $q_\perp$	8.5 MW/m^2	source simulation
Safety factor on axis and at edge	$q(0) = 1.15, q(a) = 4.8$	source simulation
Peak magnetic shear	$s = 1.85$ at $r = 0.28 \text{ m}$	source simulation
Rotational transform $\iota/2\pi$	0.42	source simulation




Figure 2. The source document's illustration of the intended flux-surface topology, with nested surfaces compressed toward the midplane and the disc positions marked at $Z = \pm 0.125 \text{ m}$, labelled with the safety-factor values of Table 2. This is vector artwork drawn to depict the intended topology, so it reads as a diagram of the design intent rather than as a contour plot of solver output, and no field data file backs it. Redrawn by scripts/papers/make_mc_fig2.py. The five q labels sit where the source puts them, which is along a diagonal rather than on the contours they name.

The source document contained a third figure, a poloidal flow velocity profile captioned as showing two curves that agree within 8%. That figure renders its axes and no data. It contains no curves and no plotted points. An empty axis frame is not a plot and it is not reproduced here.

6. Three recomputations from the document's own equations

A design study that states its formulas can be checked against itself. Three quantities in the source are defined by an equation whose inputs the same document supplies. One reproduces and two do not, and the two that fail are headline numbers.

6.1 Skin depth: reproduces

With $\sigma = 46 \text{ MS/m}$ and $\omega = 2\pi \times 1000 \text{ rad/s}$, the stated expression gives 2.347 mm against the quoted 2.3 mm. This is

$$\text{beta} = \langle p \rangle / (\langle B^2 \rangle / 2 \mu_0) \\ = (2e5 \text{ Pa}) / ((11.2 \text{ T})^2 / (2 * 4 \pi * 1e-7)) = 0.083 \quad [\text{as printed}]$$

Carrying out that arithmetic gives $11.2^2 / (2\mu_0) = 4.99 \times 10^7 \text{ Pa}$ for the magnetic pressure, and 2×10^5 divided by that is 0.0040, or 0.40%. The printed result is 8.3%, larger by a factor of about 21. Using the unamplified field $B_0 = 3.5 \text{ T}$ instead gives 4.1%, which is also not 8.3%. Recovering 8.3% from a pressure of $2 \times 10^5 \text{ Pa}$ would require a root-mean-square field of about 2.46 T, which is below both field values the document states. No corrected figure is substituted here, because which input is wrong is not determinable from the document: the pressure, the field, or the printed result.

$$\text{CR} = (1 + 0.10^2 * (0.3/0.40)^4) ^ (5/3) = 1.0053 \\ \text{CR} = (1 + 0.15^2 * (0.3/0.40)^4) ^ (5/3) = 1.0119 \quad [\text{taking } 0.15 \text{ at face value}]$$

against a printed CR of 3.2. The quoted field amplification from 3.5 T to 11.2 T is a factor of 3.2 and is therefore the compression ratio restated rather than an independent result. Since the compression ratio does not follow from the formula given, the amplification does not either, and $B_{\text{max}} = 11.2 \text{ T}$ is an assumption rather than a derived quantity. That matters beyond this section, because the beta calculation above uses 11.2 T as its field.

6.4 The confinement scaling expression

The source presents an energy confinement scaling of the form $\tau_E = 0.056 I_p^{1.24} B_t^{0.15} n_e^{0.28} P_{\text{heat}}^{-0.58} A^{0.5} \kappa^{0.67} (1 + \Delta\omega/\omega_0)^{0.42}$, described as the ITER98(y,2) scaling with an added rotation term. Two structural observations follow from the expression alone. The published ITER98(y,2) scaling contains a major-radius term with a large exponent, and this expression has no major-radius dependence at all, so it cannot reduce to the scaling it claims to extend. Several of the other exponents also differ from the published values. The expression is therefore recorded as a proposed scaling of unknown provenance rather than as ITER98(y,2) plus a correction.

7. Position relative to existing concepts

The design argument is that mechanical rotation replaces geometric precision. Wendelstein 7-X shapes its field with many non-planar superconducting coils held to roughly millimetre tolerance; the DRPD proposal shapes it with two rotating surfaces. That trade is real, and so is its cost: vacuum integrity at rotation rates above 10,000 rpm, gyroscopic loads on assemblies of the specified 847 kg, bearing life, and the feedthrough problem. The source treats these as engineering challenges and quantifies none of them. They are the parts of the concept most likely to decide whether it can exist at all.

consistent to the precision quoted.

6.2 Volume-averaged beta: does not reproduce

The source writes the definition and the substitution as

6.3 Compression ratio: does not reproduce

The source defines the normalised rotation difference as $\Delta\omega = |\omega_1 - \omega_2| / \omega_{\text{avg}}$ and states that it is approximately 0.15. With the rotation rates the same document specifies, $\omega_1 = 2\pi \times 1050$ and $\omega_2 = 2\pi \times 950$, that ratio is exactly 0.10. The compression formula then contains $(\Delta\omega/\omega_0)^2$, which divides an already normalised quantity by the mean rate a second time. Taking the expression at face value with the stated $r = 0.3 \text{ m}$, $a = 0.40 \text{ m}$ and $\gamma = 5/3$:

The rotating-field principle shares heritage with the PFRC devices at Princeton, which sustained field-reversed configurations using rotating magnetic field antennas. The source claims the DRPD approach raises coupling efficiency from 15% in PFRC-2 to 67% by using the entire disc surface as a dynamic field generator. The 15% figure is attributed to published work; the 67% figure carries no derivation and no simulation behind it, and is recorded here as a conjecture.

The turbulence argument is that differential rotation suppresses ion-temperature-gradient driven transport through $E \times B$ velocity shear, which is a documented effect in tokamak experiments elsewhere. The source reports a gyrokinetic analysis giving a factor of 2.4 reduction in turbulent heat flux over $r/a = 0.6$ to 0.8, and an effective thermal diffusivity of $0.42 \text{ m}^2/\text{s}$ against an ion neoclassical value of $0.85 \text{ m}^2/\text{s}$. Those are solver outputs and section 8 records them as such. The underlying physical mechanism is established; its magnitude in this geometry is the open question.

On materials, the relevant published result is that tungsten monoblocks developed for the ITER divertor withstand steady-state heat loads of order 10 MW/m^2 , with test exposures to 20 MW/m^2 . The source's proposed heat flux of 8.5 MW/m^2 sits below that limit. The argument that rotation spreads the load azimuthally and so lowers peak surface temperature is physically reasonable and unquantified: the reported peak temperatures of 1420°C for the rotating case and 1800°C for a static tile have no calculation attached. Tungsten recrystallisation above roughly 1300°C is a published material threshold and is the constraint such a calculation would have to respect.

8. Provenance of every number in this document

Table 3. Every number in this document falls into one of these classes. Stating provenance per quantity, rather than once for the whole document, is what made the two failed recomputations of section 6 findable.

Class	What it means	Examples
Calculated	Follows from a stated formula and stated inputs, and reproduces on recomputation.	Skin depth 2.35 mm; aspect ratio $A = 2.1$ from R_0 and a
Assumed	A design input chosen by the author rather than derived from anything upstream of it.	All of Table 1; $B_0 = 3.5$ T; $\alpha_{\text{flow}} = 0.85$; contact resistance $2.5 \times 10^{-5} \text{ m}^2\text{K/W}$; edge T_e and n_e
Calculated assumptions	Correct application of a published formula to assumed inputs. As reliable as the assumptions.	Tungsten influx of order $10^{18} \text{ m}^{-2}\text{s}^{-1}$ from the Bohdanský yield
Source simulation	Attributed by the source to its own solver, for which no artifact exists and whose integrator and timestep are unstated.	Table 2; the +0.42 stability margin; the gyrokinetic heat-flux reduction; the 45 MJ magnetic energy
Does not reproduce	Defined by an equation in the document whose own inputs give a different answer.	$\beta = 8.3\%$; $CR = 3.2$ and the 3.5 T to 11.2 T amplification that restates it
Conjecture	Asserted with no derivation, no simulation and no citation.	67% coupling efficiency; peak surface temperatures of 1420°C and 1800°C
Published elsewhere	Results of other groups, cited, and not claimed here.	ITER tungsten monoblock heat-load limits; PFRC rotating-field current drive; CuCrZr conductivity
Removed	Previously presented as an experimental result. There was no facility, no discharge and no instrument.	The experimental column of the parameter table and its uncertainties; the post-exposure surface analysis; the loop voltage annotated as measured; the claimed simulation-experiment agreement

9. What this document establishes, and the scope of the study

- A specific geometry: two counter-rotating tungsten-lined discs of 1.2 m diameter at an axial separation of 0.25 m, around a plasma of major radius 0.85 m and minor radius 0.40 m, with every dimension recorded as a design input.
- A governing equation set, an intended discretisation of $256 \times 128 \times 64$ with refinement to 0.8 mm, and the resistivity and impurity models to be used. That is enough specification for someone to begin writing the solver.
- That the skin depth of the specified conductor at the specified rotation rate is comparable to the specified armour thickness, 2.35 mm against 3 mm, which is the one quantitative consistency check in the document that passes.
- That two of the document's headline numbers, the volume-averaged beta and the compression ratio, do not follow from the equations and inputs the document itself supplies. Finding that is a result of this rebuild and it is the most useful thing in the paper.
- A set of engineering problems that would decide the concept and are currently unquantified: vacuum integrity above 10,000 rpm, gyroscopic loading, bearing life and rotating feedthroughs.

The scope of the study follows from that list, and it is stated here once rather than repeated through the document. This is a design specification with a governing equation set and a set of assumed inputs, plus an audit showing that two of its published quantities are internally inconsistent. Nothing here is an experimental result. No apparatus of the kind described exists, no plasma has been produced by this configuration, no diagnostic has been fielded on it, and no measurement of any quantity in this document has been taken. In particular:

- No confinement claim is supported. The reported confinement time of 2.34 s is a solver output produced by a model whose time integration scheme and timestep are unstated, so a reader cannot reproduce or falsify it.
- No stability claim is supported. The +0.42 contribution to the modified Suydam criterion comes out of the same solver, and the criterion itself is a local, ideal condition that does not by itself establish stability against the resistive and kinetic modes that matter in practice.
- No compression claim is supported. The compression ratio does not follow from the stated formula, and the field amplification to 11.2 T restates the compression ratio rather than confirming it independently.
- No beta claim is supported. The stated value is inconsistent with the stated definition by a factor of about 21, and the correct value cannot be recovered without knowing which input is wrong.
- No materials claim is supported. Erosion rates, surface recession and peak temperatures were once presented as post-exposure measurements and there was no exposure. The published tungsten limits cited here belong to other groups and say nothing about this geometry.
- No comparison with any operating device is supported. Figures quoted for JET, Wendelstein 7-X, PFRC, DIII-D, NSTX, Alcator C-Mod and ITER are other groups' published results included for context, and the claim of a rise in coupling efficiency from 15% to 67% is a conjecture with no derivation.
- No engineering feasibility claim is supported. Rotating an 847 kg assembly at 14,300 rpm inside a vacuum vessel while carrying multi-megawatt heat flux is the central difficulty of the concept and it is not analysed here.

One correction to the record belongs in this section rather than anywhere else in the document. An earlier version of this whitepaper read as an experimental paper: it reported a measured parameter column beside a simulated one, uncertainties quoted as one standard deviation over 25 reproducible discharges, a measured loop voltage, scanning electron microscopy and energy-dispersive X-ray analysis of disc surfaces after 1,200 cumulative seconds of plasma exposure, and an overall agreement figure between simulation and experiment. There is no vacuum vessel, no disc assembly, no diagnostic set and no discharge, and none of that happened. Those figures were cut rather than carried with a footnote, because a measured value with a footnote is still a measured value once it is quoted onward. The same version listed a second author at a laboratory that does not exist, with a correspondence address at that laboratory's domain; the co-author and the affiliation are gone, and this document has one author and one organisation.

10. What would make any of this checkable

- State the time integration scheme, its order, the timestep or CFL condition, and the divergence-cleaning method for B. Without these, no output of the model can be reproduced.
- Run and publish a grid convergence study. A single run at $256 \times 128 \times 64$ with unstated refinement criteria establishes nothing about resolution independence.
- Publish the input deck, the initial equilibrium, the boundary conditions in numerical form, and at least one full output file, so that a reader can recompute a profile rather than accept a table.
- Resolve the beta inconsistency by stating which of the pressure, the field or the printed result is correct, and recompute everything downstream of it.
- Rederive the compression ratio with the normalisation stated unambiguously, since the current expression normalises the rotation difference twice.
- Replace the confinement scaling with the published ITER98(y,2) form including its major-radius term, or state plainly that the expression is a new proposed scaling and derive it.
- Quantify the rotating-machinery engineering: bearing loads, gyroscopic moments, vacuum feedthrough leak rates and disc stress at 14,300 rpm. These decide whether the configuration can be built before any plasma physics matters.

References

- J. Wesson, Tokamaks, 4th ed., Oxford University Press, 2011.
- J. L. Luxon, "A design retrospective of the DIII-D tokamak", Nuclear Fusion 42, 614 (2002).
- S. A. Cohen, A. H. Glasser, D. C. Homan and B. Berlinger, "The Princeton Field-Reversed Configuration for compact fusion", Fusion Science and Technology 63, 68 (2013).
- M. W. Binderbauer et al., "A high-temperature field-reversed configuration plasma regime for fusion energy", Nature Communications 12, 6896 (2021).
- H. Y. Guo, A. L. Hoffman and R. D. Milroy, "Rotating magnetic field current drive of FRCs in the TCSU device", Fusion Science and Technology 51, 140 (2007).
- J. P. Freidberg, Plasma Physics and Fusion Energy, Cambridge University Press, 2007.
- ITER Physics Expert Group, "Chapter 1: Overview and summary", Nuclear Fusion 39, 2137 (1999). Source of the ITER98(y,2) confinement scaling discussed in section 6.
- P. C. Stangeby, The Plasma Boundary of Magnetic Fusion Devices, CRC Press, 2000.
- A. Loarte et al., "Progress on the ITER divertor and first wall plasma-facing components", Fusion Engineering and Design 136A, 1 (2018).
- J. Bohdanský, "A universal relation for the sputtering yield of monatomic solids at normal ion incidence", Nuclear Instruments and Methods B 2, 587 (1984). Source of the yield formula in section 3.