

Sensitivity Studies of Power Deposition Computed on Tokamak First Wall

Wayne Arter¹, Senior Member, IEEE

Abstract—Profiles of power deposition on the first wall of the Joint European Torus (JET) tokamak experiment are fit by tracing an analytic representation for the distribution of power from midplane along lines of magnetic field. The technique is used to help design plasma-facing components (PFCs) in reactor-scale magnetic confinement devices; hence, understanding how to employ it both efficiently and accurately is important. Focusing on JET divertor geometry, the work examines the sensitivity of the integrated power and maximum power per PFC tile to the representation of the magnetic field and to the discretization of the tile's own geometry and that of other PFCs. For design, it is helpful to reduce computation costs per realization to a few seconds of elapsed time, and the work concludes with recommendations and guidelines for minimizing cost while retaining adequate accuracy.

Index Terms—Joint European Torus (JET), numerical sensitivity, plasma-facing components (PFCs), power deposition, tokamak, uncertainty quantification.

I. INTRODUCTION

DETAILED motivation has been presented in the recently published [1] for the accurate simulation of the power deposited on the first wall of a tokamak due to leakage from the main plasma confinement region. There are three aspects to be considered to provide confidence in any simulation, namely, verification, validation, and uncertainty quantification (VVUQ). The recent literature [1] and other literature [2], [3], [4] give evidence for the adequacy of the SMARDDA modules in respect of the first two (VV) aspects. Specifically, for the Joint European Torus (JET) tokamak, unpublished reports from a multiyear campaign to speed the execution of the software provide further VV. Extracted from this report material, this work seeks to quantify uncertainty arising due to the discretization of tokamak magnetic field and plasma-facing component (PFC) surface geometry employed by the SMARDDA-PFC code. Such understanding ultimately enables accurate and efficient UQ of different field and geometrical designs.

The presence of plasma is key to the effect that the energy flux on first wall is expected to be directed parallel to the

magnetic field; hence, the power Q deposited per unit area can be reduced dramatically by arranging for grazing incidence of magnetic fieldlines on PFCs, with obvious benefits for reducing thermal stress. Observing that, for a PFC tile of indicative extent 200 mm, even a single 1-mm corrugation of the surface can equate to half a degree of normal deflection, it is a challenge to discretize field and geometry, so that power deposition at angles of as little as 2° or less can be economically calculated. This challenge forms the primary focus of this work.

The challenge is greater, because the present strategy for triangulating geometry imported from a computer-aided design (CAD) package does not allow for adaptive mesh refinement; for example, it is not possible to insert more triangles where gradients of power density Q are large. However, the surface mesher used, provided by the CADFIX package, allows for a wide range of factors to influence triangulation, and locally written software enables systematic refinement of an initial mesh [3, Fig. 3].

The representation of the magnetic field $\mathbf{B}$ is by the de facto eqdsk standard for an axisymmetric field in terms of poloidal magnetic flux function $\Psi(R, Z)$ and toroidal magnetic field component B_T via the plasma current expressed as a function of flux $I(\Psi)$, where (R, Z) are position coordinates in a vertical plane. The discretization of $\mathbf{B}$ is all the more restricted, since Ψ is specified as point values on a uniform rectangular grid. Previously published work on convergence has concerned relatively small areas of “target” PFC, as little as one tile, where calculations of shadowing by immediately adjacent tiles can be made in a few seconds on a desktop for a detailed target with 50 000 triangles [3]. However, the JET divertor consists of approximately 600 tiles, and reactor PFCs will have surface areas that are at least an order of magnitude larger still; thus, the novelty and importance of this work lie in enabling fast design job turnaround by use of coarser triangulations in configurations that allow more complex shadowing of one tile by another.

Section II describes the JET test cases in more detail; then, Section III begins the presentation of results in a “best case” situation to set the context and largely disposes the issue of magnetic field discretization. The main study of sensitivity to surface meshing appears in Section IV, where many anomalies in the Q results become apparent. Their causes are identified, enabling the straightforward removal of most of them. Section V presents analysis that explains how the most awkward anomaly arises; then, meshings, which confirm the explanation and show how the anomalies may be

Received 28 June 2024; revised 31 August 2024; accepted 18 October 2024. Date of publication 4 February 2025; date of current version 20 March 2025. This work was supported in part by the EUROfusion Consortium funded by the European Union via the Euratom Research and Training Program under Grant 101052200—EUROfusion and in part by the U.K. Engineering and Physical Sciences Research Council (EPSRC) under Grant EP/W006839/1. The review of this article was arranged by Senior Editor M. Kovari.

The author is with U.K. Atomic Energy Authority, Culham Science Centre, OX14 3DB Abingdon, U.K. (e-mail: wayne.arter@iee.org).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TPS.2024.3485530>.

Digital Object Identifier 10.1109/TPS.2024.3485530

eliminated, are presented in Section VI. The implications of the results are discussed in Section VII.

II. TEST CASE CONSTRUCTION

Calculations are given a runid as follows:

$$\text{runid} = \text{eqid-gshad-gres}$$

where eqid identifies the magnetic equilibrium, usually in terms of the last three digits of the JET shot number, and gshad and gres are, respectively, the roots of the .vtk filenames of the shadow and target or “results” geometry triangulations, respectively. Shadowing of one surface by another PFC is key to their mutual protection, and it is helpful to be able to distinguish the two geometries for reasons explained in [3]; see also Section VI. Specification of the physics of the simulation is completed by giving the leakage power profile at midplane. This power is assumed to spread both upward and downward, with equal total amounts directed toward the inboard and outboard of the divertor. Other input parameters for SMARDDA-PFC, other than those that specify output diagnostics, serve to define the numerical algorithm. Sensitivities to most of the latter have been studied previously in [3], enabling the present work to focus on the discretization of the magnetic field and of the geometry in the simulations.

A. Equilibria

For shot #89297 at $t = 45.011$ s, the equilibrium was specified by two .eqdsk files each with a different spacing for the samples of magnetic flux Ψ . These two solutions were produced with, respectively, a coarse 32×32 Ψ -mesh and a much finer 256×256 Ψ -mesh. The former is given the simple eqid = 297, and the latter is referred to as fi297. The deposition profile for the shot is exponential decay [1, Sec. III] with $\lambda_q = 0.0066$ m and $P_{\text{sep}} = 10.5$ MW (P_{sep} was formerly referred to as P_{loss} [3], but either this new notation or the alternative P_{SOL} is now preferred for the total particle energy escaping across the separatrix).

Since strike-point sweeping to spread the power deposition more widely over the PFCs might well be used in high power density discharges, a second shot is considered. The periodic sweeping is represented by a sequence of five equilibria equally spaced in time with the first and the last in the sequence at approximately the same phase; i.e., sampling is at an interval $\tau_P/4$, where τ_P is sweep period. The flux-sweeping shot considered is JET #90271, for which $\tau_P = 0.25$ s and where the strike point oscillates by 3 cm across the surface of the T6 tile. A 128×128 Ψ -mesh is used throughout to represent the magnetic field. Since more than one equilibrium is taken from this shot, they are numbered in time order after a decimal point, so that eqid = 271.1 corresponds to the first equilibrium in the strike-point sweeping sequence at $t = 49.0$ s, 271.2 to $t = 49.0624$ s, 271.3 to $t = 49.1368$ s, 271.4 to $t = 49.198$ s, and 271.5 to $t = 49.2788$ s. The formula for power deposition profile is due to Eich, which is mathematically defined in [1, Sec. A] and [3, Sec. II-B] following [5], with parameters $\lambda_q = 0.0170$ m $S = 0.0011$ m

and $P_{\text{sep}} = 10.2$ MW. Each parameter is assumed constant throughout the cycle, because the sweep takes a time relatively short compared with overall discharge evolution. These shots are in the JET “hybrid” ELMy H-Mode scenario, for fuller experimental details of which see [6].

B. Geometry

Since the original design of JET predates the universal employment of CAD systems, the acquisition and assembly of geometrical representations suitable for meshing required a deal of assistance as noted in the acknowledgment section at end. Fortunately, the divertor design belongs to a later period where CAD files including construction details, such as holes and fastenings, were available. It followed that a workflow of modern relevance was possible for the calculation of Q on the divertor, which is anyway a region of high interest to tokamak designers. The divertor consists of 24 repeats of a set of 24 tiles that spans 15° in toroidal angle, together with 48 repeats of a set of vertical plates joined (T5) to present their edges to the plasma.

The workflow begins by defeaturing the CAD description of the 15° toroidal segment, removing irrelevant construction details, using the commercially provided CADFIX software. Attention focuses on the surface interfaces with vacuum, particularly the PFCs, which as the name implies those directly facing the hot, centrally confined plasma, and especially on the adjacent narrow fillets illustrated in Fig. 1. The retention of the latter demands the extra expense of locally fine mesh; thus, two attempts were made to remove all such fillets “by hand,” as the automatic defeaturing facility was defeated by some of the three-way joins. Surface meshing of the defeatured CAD was also carried out using CADFIX, which has a GUI that allows for a wide range of different meshing strategies. Given that four defeatured representations (two with and two without fillets) of the CAD were available and two levels of automatic mesh refinement possible, this led to some 20-plus different meshes for the divertor repeat geometry, with a range of sizes of approximately 100 in terms of the total number of triangles N_Δ in the target. All the meshes were given labels traceable to a .vtk file, keeping each one for future reference. The labels begin with “c” if fillets have been removed and with “d” or “f” if retained, but as explained in the Appendix, the remainder of the label is an arbitrary string, excepting that a terminating “2” or “4” denotes, respectively, a refined or doubly refined mesh (although not all automatically refined meshes are labeled in this way). Each separate meshing of a 15° segment is referred to as the “a” model. The combination of three adjacent copies of the “a” model is called the “b” model, which, therefore, represents one octant (eighth) of the divertor. Replicating, translating, and joining the “a” model 24 times in toroidal angle give the “c” model, and combining “c” with meshes for other PFC geometry gives the full 360° “d” model (see Fig. 2), corresponding to the initial 15° “a” mesh. Used as part of the geometry label, “a–d” are prefixed by “fs” and “ft” to denote whether the file represents a shadow or target, respectively. To save storage, instructions only are saved for regenerating the “c” and “d” models by SMARDDA code vtktfm.




Fig. 1. JET divertor tile T6 design, illustrating fillet boundaries as pale green lines, closeup at bottom left.




Fig. 2. 360° model for shadow viewed from outside looking down, with superimposed results for a 15° geometry providing an indicative fieldline, drawn in white.

III. PRELIMINARY STUDIES

A. Results for One Octant

Initial studies that included the JET limiters showed that little power was intercepted by them. Since the outer limiters are likely to be the most important PFC lacking discrete rotational symmetry in their siting around the major axis, it is reasonable to exploit the inherent symmetries of the divertor, when power deposition upon it is the main focus.

First, exploiting an eightfold symmetry, in order to verify the use of the software on a single JET octant, a series of calculations was conducted, as indicated in Table I. As Table I shows, these were used to test the adequacy of the meshing. It immediately makes clear that the total power deposited on the 15° or 1/24th segment of the divertor is approximately 400 kW and, therefore, significantly lower than the value of $P_{\text{sep}}/24 = 425$ kW expected, given that total power in the midplane profile was $P_{\text{sep}} = 10.2$ MW. It was then realized that $\lambda_q = 17$ mm was of order of one-third of the width of the scrape-off layer (SOL) at midplane, so that only the first three e-folds of the approximately exponential distribution of power could contribute to the total. Elementary integration of the decaying exponential reveals that this amounts to 95% of the total, i.e., $0.95 \times 425 \approx 404$ kW.

TABLE I
SCAN OVER STRIKE-POINT SWEEPING OF 45° TARGET GEOMETRIES FOR GEOMETRY “CLE.” “REFINED” DENOTES THAT BOTH TARGET AND SHADOW MESH REFINED $\times 4$; THUS, COLUMNS 2 AND 3 ILLUSTRATE THE SENSITIVITY TO MESH REFINEMENT

Equil -	Q_{max} MW/m ²	$\int Q/3$ kW
271.1	6.884	407
Refined	6.735	405
271.2	5.380	399
Refined	5.412	403
271.3	5.002	400
Refined	5.046	400
271.4	5.326	399
Refined	5.380	400
271.5	6.947	407
Refined	6.735	405




Fig. 3. Time history throughout strike-point sweep of power on different parts of the geometry. Labels finishing r denote results from refined calculations in Table I.

Following this renormalization, Table I shows that the integrated power is accurately computed on the coarser mesh, and this level of accuracy extends even to total power on subsets of the geometry (see Fig. 3). However, it should be noted that in the first version of Table I, there was a maximum power for $\text{eqid} = 271.3$ on the coarse mesh given as 7.261 MW, which was identified as spurious due to poor grid alignment (see Section V), and as there was only one offending triangle, the entry in Table I was easily correctable. Compute times for a single processor desktop were of order a minute for the coarser grid and scaling approximately as N_{Δ} (multicore parallelism has been shown to bring wall-clock times down to 30 s for $N_{\Delta} = 10^6$).

The segregation of the geometry indicated by Fig. 3 is about the best that can be achieved using the SMARDDA software *smanal*, which implements a clustering algorithm in poloidal angle, because seen from any point inside the PFCs, there is significant overlapping of at least two tile rows. Given that the variation between the 15° segments is negligible, from the simple practical standpoint of reducing by a factor of three, the work needed to segregate tile rows, and further exploration was conducted using a 15° toroidal segment as target, although using a full “d” model as the shadow.

A compact example of the plot of the power deposited on the divertor is provided by Fig. 5, where it will be seen from T6 that shadowing leads to sharp cutoffs both in the toroidal




Fig. 4. Flux surfaces $\Psi(R, Z)$ for $\text{eqid} = 297$, drawn in purple, superimposed on a profile or “silhouette” of the JET first wall plotted in green. Distances are in meters, flux contour spacing 0.1 Wb, except for the contour at the plasma boundary drawn in orange and its extensions down from the X-point.




Fig. 5. Power deposition Q in MWm^{-2} for a 15° toroidal segment of JET divertor, labeled with tile row numbers, for case fi297-fip-fiar. The plasma is sited above the divertor geometry, as indicated by Fig. 4.

direction (decreasing y in Fig. 5) and the radial direction corresponding to increasing x . Despite this, good accuracy is achievable using relatively coarse meshes with sidelengths of order λ_q for recently explained reasons [1]. Since the coarse mesh of Table I has a side of mean length one-third of λ_q , a tenfold increase in execution speed is conceivable.

B. Effects of Flux Mesh Size

Section II-A describes the equilibria labeled $\text{eqid} = 297$ and $\text{eqid} = \text{fi297}$ as having size 32×32 and 256×256 meshes,




Fig. 6. Close-up of contour plots of flux $\Psi(R, Z)$ is superimposed on JET silhouette (green), flux contour spacing 0.02 Wb, with plasma boundary in orange. The equilibrium $\text{eqid} = \text{fi297}$ has contours drawn in purple, whereas those of the coarse mesh $\text{eqid} = 297$ are shown as blue dots.

respectively. Fig. 6 shows that their flux contour plots overlay very closely in the region of most interest for divertor power deposition, excepting what is confirmed to be a plotting issue at the X-point by inspection of the gradients of Ψ . At close to grazing incidence, Q depends sensitively on the angle between the field and the surface normal, but it does not depend greatly on the *poloidal component* of field, since this is much smaller than the toroidal component of 2 T. To be specific, the poloidal contribution needed to deflect $\mathbf{B}$ by one degree is approximately 0.04 T, whereas the entire poloidal field at the wall is approximately 0.2 T; thus, Fig. 6 suggests a negligible change in field direction from the different samplings.

In contrast, changes in the point value of Ψ feed into a formula involving an exponential. Error estimation begins by using the fact that the software outputs and uses in the power deposition formula, a quantity somewhat confusingly referred to as *psista*, which is the value of Ψ at the midplane, i.e., the computed end of the fieldline given a starting point on the PFC. Complete information needed to evaluate Q is available at the latter point, including a value for Ψ (which is, however, ignored by default for consistency with cases with field ripple, wherein flux varies along a fieldline). By accessing the data for the flux at the PFC, it is, thus, possible to plot the discrepancy in Ψ over the wetted area as in Fig. 7, since fieldlines must to lie on flux surfaces when the field is toroidally axisymmetric (as is normally the case). Since $\Psi = \mathcal{O}(1)$, the discrepancy even on the coarse mesh seems, at first sight, negligible, and indeed, the shadow patterns are visually indistinguishable, but sadly the error of $|\Delta\Psi| \approx 2 \times 10^{-3}$ implies that Q is in error by approximately 10%. However, the resulting error in Q on the fine mesh will be 0.1% and easily tolerable.

The standard JET.eqdsk file contains a somewhat coarser 128×128 Ψ -mesh, so it is useful to understand the mechanism underlying the failure in flux conservation. For a randomly selected long fieldline, Fig. 8 shows that it is not a consequence of inaccuracy in the embedded Runge–Kutta (ERK) solver that computes the fieldline




Fig. 7. Discrepant flux for `powcal` with 32×32 Ψ -mesh (top), $\max |\Delta\Psi| = 2 \times 10^{-3}$ and 256×256 Ψ -mesh (bottom), smaller $\max |\Delta\Psi| = 2 \times 10^{-5}$. `eqid` = (fi)297-fsdfip-fiir.

trajectories. Further investigation confirms the suggestion from Fig. 6 that the error correlates with flux gradients. The system of fieldline equations solved for JET divertor models advances positions in cylindrical polar coordinates by the corresponding components of axisymmetric equilibrium $\mathbf{B}$. The finite-size mesh used to represent Ψ implies that the fieldline equations are solved subject to subgrid spatial uncertainty in the field. Generically, this deficiency leads to a failure of flux conservation, even when (pseudo-)time step $\rightarrow 0$. This explanation is supported by the facts that the error is small and accumulates only over large spatial scales, but is worse the coarser the Ψ -mesh, and also greater where the flux gradients are larger.

The above arguments and simulations imply an empirical formula for the relative error in $Q = 100\Delta\Psi_{\text{num}}$, where Ψ_{num} is the numerical value of Ψ in SI units, which is confirmed by Fig. 10, that indicates 2% or less error. Hence, the error is bearable for Q calculations using a 128×128 Ψ -mesh.

From the plot of flux contours in Fig. 6, it is possible to estimate graphically that the spatial errors (in m) at the endpoint of the fieldline are $0.2\Delta\Psi_{\text{num}}$, equating to at most 0.04 mm for the 128×128 Ψ -mesh. Hence, there is no visible effect on the extent of the intertile shadowing, since the triangle dimensions are typically at least several mm. The difference




Fig. 8. `eqid` = 297-fsdfip-fiir. Failure of magnetic flux conservation along fieldline 753 with 32×32 Ψ -mesh, diagnostic activated by uncommenting source code. ϵ is the tolerance used in the ERK solver. The fieldline computed with greater accuracy was allowed to continue beyond the inner midplane.




Fig. 9. Failure of flux conservation over divertor using the 128×128 Ψ -mesh, $\max |\Delta\Psi| = 2 \times 10^{-4}$. Equilibrium 271.1, #90 271 at $t = 49.0$ s.

between flux values found at the two ends of a numerically computed fieldline is noted to be a powerful diagnostic of its accuracy in the typical, toroidally axisymmetric field.

IV. SENSITIVITY STUDIES

The 25 different `gshad-gres` combinations of meshes were considered for a given `eqid`, using the notation of Section II. Mesh triangle counts are listed in Table II. There is a significant reduction in execution time through pairing a relatively coarse `fsd/360°` shadow with a finer `fta/15°` target, because this considerably reduces the number of fieldline-surface intersection tests required at each step, as well as the storage needed for the calculation. Frequently, a refined mesh, typically with a name containing a terminating “2” or “4,” as mentioned in Section II-B, was paired with the original mesh as its shadow.

The, otherwise, trivial matter of labeling the 15° geometry segment by tile row and location enables completion of a workflow that is almost completely automatic for processing and analyzing power deposition on the set of mesh combinations. The script that coordinates the execution of the different modules `geoq`, `hdsgen`, and `powcal` needed to



Fig. 10. Absolute (top) and relative error (below) in power deposition due to failure of flux conservation along a fieldline. runid = 271.1-fsdfia-fiar.

perform a SMARDDA-PFC simulation was, in turn, invoked by a higher level script covering the 25 cases. A second script was then written to invoke the `smanal` module for analyzing `powcal` output to give maximum and integrated values of Q on each tile row and to combine the results both in tabular and graphical form. Both equilibria jet271.1 and jet271.3 (subsequently abbreviated to 271.1 and 271.3, respectively) were analyzed in this way. The combined plots were initially ordered by number of triangles in the target; later, when it was realized that there was negligible power deposition on T5, plots were based on the total number of triangles representing the other tiles.

The results obtained by the automated process showed a number of anomalies. First, although the results for 271.1 and 271.3 tended to share the same features, those for the *gshad-gres* combinations of *challmr-challmr2* and *challmr2-challmr2* differed. This was traced to the fact that the sorting algorithm, since the simulations shared the same size target, had reversed the order of analysis of their output and simply corrected. Simulations will, henceforth, unless stated otherwise, be for 271.1 and so may be identified by the labels of their shadows and targets only (further omitting *fsd* and *fta* label prefixes). The second set of anomalies was confined to the T1 results and found to occur, because a smaller area of the T1 surface had been meshed in simulations made before it was realized that

notably in the flux sweeping calculations, power was deposited over much of T1 and indeed spread to T0. Depending on which T1 mesh was used, the total power deposited could change by a factor of approximately two. The simple fix made here was to omit T1 from the plotting and focus on the other tile rows.

Attention now transferred to the integrated and maximum powers on T6, which contains the region where the largest values of Q are expected in 271.1. Unexpected variability in their values was traced to inappropriate shadowing of the concave region on T6, as illustrated by Fig. 11 (top). Such a problem had been anticipated, and a control parameter ℓ_{crit} introduced so as to ignore fieldline intersections with surfaces when the length of fieldline was below this value. However, the critical length parameter was specified in units of the underlying uniform grid, partly for convenience of implementation, but also because it gives rise to a useful anisotropy in the metric. Spurious collisions may also occur with neighboring triangles on the surface, in which case the length along the fieldline to intersection may be on the order of the triangle size, rather than the considerably smaller distance between two triangulations at the same point. Specifically, all calculations in this section use a “uniform” grid of size $16\,384 \times 65\,536 \times 16\,384$ to represent a volume of approximately $2 \times 20 \times 2$ m, so that, e.g., $\ell_{\text{crit}} = 3$ corresponds to 0.4 mm in a radial direction, but over 0.9 mm in the toroidal (Y) direction.

Unfortunately, it seems that consequent on this choice of metric, premature fieldline termination took place over a significant area of T6, as indicated by Fig. 11, leading to greatly reduced maximum and integrated Q on refined target meshes. Increasing ℓ_{crit} from 3 to 20 removed the effect for six of the visibly anomalous cases, and the remaining two, *fias-fiar2* and *challmr-challmr4*, were fixed by use of $\ell_{\text{crit}} = 50$. However, in the latter case, fieldline intersection with objects over 10 mm toroidally from the starting location will be ignored, which since the intertile separations range from approximately 5 to 14 mm means that other triangles on tile edges start to be incorrectly illuminated.

The remaining obvious anomalies are confined to the highest values of Q on different tiles, some of which are capable of affecting the global maximum. Inspection of the results using ParaView [7] indicated that the anomalous global maxima were confined typically to a small number, often four or less of isolated triangles, which appeared to be spurious as a result of distortions to the T6 tile resulting from the defeaturing of the fillets (see Fig. 12). However, all T6 results for the defilleted tiles seemed to be affected to an extent, which prevents detailed comparison between filleted and defilleted cases. In fact, in all the cases considered, the fillets receive no power. The difficulties encountered in their removal, and the likelihood that if fillets are not shadowed, large angles of fieldline incidence will lead to high values of power deposition, are, nonetheless, two strong arguments for retaining fillets except in special circumstances.

Anomalies in the T3 maximum are attributable to the field misaligned meshing described in Section V. Anomalous T1 maxima are linked to the effect discussed in Section VI, but, in any event, do not appear in the corrected plots that




Fig. 11. Top: schematic of the spurious shadowing of a concave surface of the blue volume. The surface is represented by two elements in the finer meshing (red) and one surface in the coarser meshing (green). Fieldlines starting on the finer mesh soon intersect the coarse mesh. Bottom: overlay of meshes of the divertor geometry, `clas` and `clar2`. The first mesh is shown in green, and the second, its refinement by a factor of four, in red.




Fig. 12. Power deposition on T6, case *fsdchasm-chasm*. Highlighted triangles have anomalously high values of Q .

follow as Figs. 13 and 14, where it will be seen that the results for each tile are, otherwise, generally not too sensitive to triangulation size.

V. SURFACE ELEMENT PROPERTIES

For SMARDDA-PFC calculation of point Q , the key quantities are the barycenter of each triangle and its normal $\mathbf{n}$ outward from the solid, which will here be assumed cylindrical with radius R . In the analysis, it is convenient to work in polar coordinates centered on the cylinder axis (see Figs. 15 or 17),




Fig. 13. Effects of mesh, broken down by tile row (15° segment). Total power (integrated Q) per row for each meshing, for $\text{eqid} = 271.1$.

so that

$$2R \sin \frac{\gamma}{2} = h_\Delta, \quad \text{so } h_\Delta \approx \gamma R \quad (1)$$

in the limit of small h_Δ/R . Computation starts by writing down the Cartesian coordinates of each point $ABCDE$ relative to an origin at the cylinder axis in the normal plane containing the midpoint of AB ; thus,

$$\begin{aligned} A, B: & (0, R, \pm \ell) \\ C: & (R \sin \gamma, R \cos \gamma, 0) \\ D, E: & (R \sin 2\gamma, R \cos 2\gamma, \pm \ell). \end{aligned} \quad (2)$$

Triangles BCE and ACD share an unsatisfactory property, now discussed in detail for BCE in Section V-A. Triangles ABC and CDE have very similar and overall satisfactory properties, shared with all the triangles considered in Section V-B thereafter. There follows a study in Section V-C as to how small vertex displacements can lead to unsatisfactory triangle properties, and the importance of triangle alignment is summarized in Section V-D.




Fig. 14. Effects of mesh, broken down by tile row (15° segment). Maximum Q per row for each meshing, for $I_{eqid} = 271.1$. Meshes for additional sensitivity studies described in Section VI.

A. Problematic Surface Element

Elementary calculation gives the barycenter of BCE as follows:

$$\frac{1}{3}(R[\sin \gamma + \sin 2\gamma], R[1 + \cos \gamma + \cos 2\gamma], 2\ell) \quad (3)$$

which corresponds to a polar angle β of

$$\tan \beta = \frac{\sin \gamma + \sin 2\gamma}{1 + \cos \gamma + \cos 2\gamma}, \quad \text{implying } \beta \approx \gamma \quad (4)$$

in the limit of small h_Δ/R . The normal to the triangle BCE lies in the direction given by the vector cross product

$$(R \sin \gamma, R[\cos \gamma - 1], -\ell) \times (R \sin 2\gamma, R[\cos 2\gamma - 1], 0) \\ = \ell R(1 - \cos 2\gamma, \sin 2\gamma, R/\ell \cdot 2 \sin \gamma[\cos \gamma - 1]) \quad (5)$$




Fig. 15. Schematic of triangulation of cylinder of radius R . The triangle ABC has base 2ℓ and height h_Δ , where “height” corresponds to an angle γ subtended at the cylinder axis. Conversely, triangle BCE has base $2h_\Delta$ and height ℓ .




Fig. 16. MVT in 1-D. At a value of x between the ends of the secant (shown in red), the curve $y(x)$ has the same gradient dy/dx as the secant.

using the trigonometric identity $\sin \gamma(\cos 2\gamma - 1) - \sin 2\gamma(\cos \gamma - 1) = 2 \sin \gamma(\cos \gamma - 1)$, so that the normal is directed in polar angle η_θ , where

$$\tan \eta_\theta = \frac{1 - \cos 2\gamma}{\sin 2\gamma}, \quad \text{so } \eta_\theta \approx \gamma \quad (6)$$

which accurately corresponds to the normal to the cylinder surface at the barycenter angle β . However, the normal has a component in the z -direction, so is tilted out of the vertical by an angle η_z given by

$$\tan \eta_z = \frac{R}{\ell} \cdot \frac{\sin \gamma(\cos \gamma - 1)}{\sqrt{2 - 2 \cos 2\gamma}}, \quad \text{so } |\eta_z| \approx \frac{R \gamma^3/2}{\ell \cdot 2\gamma} = \frac{R \gamma^2}{\ell \cdot 4} \quad (7)$$

in the limit of small h_Δ/R .

1) *Mean Value Theorem*: It is worth noting that triangles BCE and ACD are counter examples to the possibility of a fully 2-D application of the mean value theorem (MVT). In one dimension, meaning application to a scalar function y




Fig. 18. Power deposition on T6, case *fsdfitt3-fitt3*, showing “speckling” in some of the regions near the boundary of high- Q deposition.

obscure genuine “hot spots” from the designer’s attention, and it should be noted that there is a possible situation where field error produces a similar effect and will need eliminating earlier in the workflow, viz., where there is a checkerboard distribution of residuals in the equilibrium solution.

Provided at least one edge of a triangle lies approximately in the toroidal direction (y in Fig. 18), then the Q point values are more accurate. The effect of misalignment on the quadrature is also further reduced, since the spurious contributions tend to cancel, e.g., in Fig. 15. *ACD* and *BCE* have precisely opposite normals, so any spurious contribution will be limited by the change in $\mathbf{B}$ over the separation of their barycenters.

VI. ADDITIONAL SENSITIVITY STUDIES

Drawing on the above experience, the CAD description of T4 and T6 was defeatured for a third time, retaining the fillets. Meshings of the new geometry were combined with meshings of T0, T1, T3, T7, and T8 to produce.vtk files for the divertor “repeat,” and another 11 cases were constructed.

Although, for most SMARDDA-PFC fieldline-tracing calculations, the default size of ℓ_{crit} is adequate, a more reliable way to avoid problems might be to shadow the target with an identical mesh. This was not initially considered, because previously studied tokamak geometries had all had discrete toroidal symmetries, whence SMARDDA-PFC could impose periodic boundary conditions to reduce the problem size to the geometry “repeat,” i.e., by $1/12$ in MAST-U application [3]. The shadow geometry had thus to have the same meshing in every repeat. However, the JET results of Section IV confirm SMARDDA-PFC capability to treat multimillion triangulations, hence inhomogeneous shadowing to be feasible. The preceding experience indicates that the shadowing geometry needs to be as fine as the target geometry “locally.” The meaning of locally in this context was explored by using 360° divertor shadow meshes generated, as shown in Fig. 19. Replacing only the target geometry with a fine mesh is shown to be vulnerable to spurious deposition when edges of convex tiling are exposed, as shown in Fig. 20, indicating that adjacent “local” repeats should also be meshed finely for robustness.




Fig. 19. Schematics of toroidally 360° shadowing geometry, imagined viewed from above. “= T” here denotes the location of the target segment. At the top is the original periodic arrangement of repeats with the coarse (C) meshing. Inhomogeneous shadows are sketched underneath: the left one shows a fine mesh (F) coinciding with the target, whereas the right shows the target bracketed by fine meshes (note that a more accurate schematic of the JET meshes would have 24 segments rather than the 12 drawn above).




Fig. 20. Exploratory case *filp4-dint4* with inhomogeneous shadow, showing fieldline incorrectly missing adjacent coarsely meshed convex tile geometry.

The meaning of locally in the context of accurate shadowing will depend on precise details of geometry, but a common situation will be a configuration involving rows of rectangular tiles, with approximate sizes $L_\phi \times L_\theta$, where L_ϕ and L_θ are the dimension in the toroidal and poloidal directions, respectively. Supposing shadowing by tiles on other rows to be less critical than by tiles of the same row, then coarser meshing might be allowed once a path following a fieldline has moved a poloidal distance of L_θ , equating to a poloidal angle of $\Delta\theta = L_\theta/r$, where r is the minor radius. Assuming the local fieldline safety factor is q , the corresponding toroidal angle is $q\Delta\theta$, equating to a distance traveled of approximately $qR\Delta\theta$.




Fig. 21. Effects of mesh, broken down by tile row (15° segment). Total power (integrated Q) per row for each meshing, for eqid = 271.1. Meshes for additional sensitivity studies described in Section VI.

Supposing each tile to subtend a toroidal angle of approximately L_ϕ/R , the number of additional fine copies of a tile needed by the shadow geometry at each side of the target is, thus,

$$N_{f\pm} = \frac{qL_\theta R}{L_\phi r}. \quad (16)$$

Evidently, if the target consists of n_T copies of the tile in a row of interest, then the total number of fine copies needed by the shadow is

$$N_f = 1 + 2 \left[\frac{qL_\theta R}{n_T L_\phi r} \right]^+ \quad (17)$$

where $[\cdot]^+$ denotes that the bracketed quantity should be rounded up to the nearest whole number.

The new 11 cases, all with fillets, were run semiautomatically, as described in Section IV for 271.1. The results are shown in Figs. 21 and 22, which now show an absence of major anomalies, with again the coarser grids producing results very little different from those obtained with up to




Fig. 22. Effects of mesh, broken down by tile row (15° segment). Maximum Q per row for each meshing, for eqid = 271.1.

a hundred times more triangles (and, therefore, given that computational cost is approximately proportional to N_Δ , up to a hundred times more expensive). The most significant variation is in $Q_{\max}$ on T6, but inspection shows that the higher values are all associated with a curvature sensitive meshing of T6 and its successive refinement. The *fit-ditt* meshes responsible, although arguably a more accurate representation of the surfaces than the others, are misaligned, accounting for the anomaly in $Q_{\max}$.

VII. CONCLUSION

The work presented here in Sections IV–VI has shown that it is possible to compute power deposition in the divertor to an accuracy better than 10% and probably exceeding that

of the physical model, with a remarkably small number of target triangles and a relatively coarse mesh for the flux function. The achievement is assisted by the relatively simple pattern of power deposition in the divertor. Unfortunately, Sections IV–VI have also shown that it is quite possible to have comparable or worse errors with much finer geometry and flux function discretizations.

The following guidelines are proposed based on the current work and the companion paper [1].

- 1) Great care should be taken in defeaturing CAD, especially when the plasma-facing surfaces might be affected. Neither user-supervised nor automatic routine removal of adjacent fillets is recommended. It may be safer, albeit slower, to rebuild surfaces using specially selected curves and face geometry extracted from the CAD database.
- 2) Make an initial coarse mesh of the main plasma-facing surfaces with triangles of side no greater than the typical length scale of input power profile, viz., λ_q in the case of a simple exponentially decaying profile. Check results using a $4\times$ refined meshing. Certain critical areas, normally tile or limiter edges adjacent to the main surfaces, may require much finer meshing, down to assembly tolerance of typically a millimeter, and in the absence of special indication; all such areas should be gridded to a finer length scale.
- 3) Meshing should be aligned with the main direction of the magnetic field, viz., one side of each triangle should be aligned with or close to the toroidal direction in a tokamak.
- 4) Where shadowing and target geometry coincide, they should share a common triangulation. In cases where results are needed only over a fraction of the toroidal angle, coarser triangulation of more distant shadow features is acceptable, based on the formula of Section VI.

There is an important practical point that execution speed may be significantly increased by inserting into the shadow geometry “cutouts” and “beancans,” nonphysical surfaces that intercept fieldlines certain *not* to connect with the midplane, and “skylights,” nonphysical surfaces certain to connect to the midplane. These additions may greatly reduce the total length of fieldline that is followed, easily outweighing the costs of the relatively minor increase in size of the shadow geometry.

APPENDIX MESHES

For the meshing of the JET divertor geometries, use was made exclusively of the CADFIX mesher. This allows meshing strategies that may insert points into surfaces, to satisfy the Delaunay (empty triangle circumcircle) property, or try to account for surface curvature, the latter giving the option to increase element size on flatter surfaces by setting suitable control parameters. Other parameters set the maximum and minimum acceptable lengths for mesh edges, the extreme values of acceptable angles in the triangulation, and the maximum and minimum number of divisions allowed along an edge &c., together with further parameters that set the relative

TABLE II
MESHED GEOMETRIES, ORDERED IN TERMS OF TOTAL NUMBER OF TRIANGLES IN THE MESH, FOLLOWING THE SUBTRACTION OF THE NUMBER OF TRIANGLES IN T5

Label	Triangles	Label	Triangles
challmr	2 851	clet204	51 635
dint	3 631	clar	52 963
challma	6 792	dint4	58 096
challmr2	11 404	filt204	58 206
dint2	14 524	fiar	63 375
ditt	16 236	ditt3	64 944
cfallmm	19 883	challnm	69 469
challn	22 159	challn2	88 636
ditt2	22 248	chasm2	92 868
chasm	23 217	clet2042	208 492
ftavm	36 486	clar2	211 852
fsazm	41 405	filt2042	232 824
chatm	42 084	ditt4	235 728
challmr4	45 616	fiar2	253 500
ftaum	47 871		

importance of satisfying the preceding constraint parameters. This parameter space was freely explored for five different CAD descriptions over several years, in such a way as to defeat a systematic naming convention, hence the arbitrary choice of mesh label, beyond the choice of beginning with “c” if fillets have been removed and with “d” or “f” if retained in the CAD prior to meshing.

The problem with describing any meshing is not only the sheer number of parameters that may affect the result, but also the difficulty that even small changes to an algorithm, as minor as changing the order of processing a CAD part, may make significant changes to the result. The decision whether to join point $P1$ to $P2$ or instead to $P3$ has to be taken on the basis of comparison of real numbers, which may be very close to coincidence if there is symmetry in the local geometry, and thus, the decision effectively be made at random. Since the first marginal choice may thereafter affect all subsequent mesh connectivity, there is no expectation that meshing will be reproducible from one software release to another, or even from one computer architecture to another. Hence, the only reliable way to achieve reproducibility is to retain the mesh files, as described in the main text files, which may then be inspected using 3-D viewing software, such as ParaView should it be necessary to check details of the triangulations. The number of triangles N_Δ in each description of the geometry is, however, recorded herein, as it gives an approximate value for the average mesh length $h_\Delta = \sqrt{A_\Delta/(2N_\Delta)}$. Since the total area A_Δ is here approximately a square meter, $N_\Delta = 20\,000$ corresponds to $h_\Delta = 5$ mm. Triangle numbers are presented in Table II. Note that, in some cases (calculations), the shadow mesh of the 15° cell differs from the target by the insertion of cutouts, skylights, and so on and so has a label differing by one character, typically an “r” or “t” is replaced by “s.”

ACKNOWLEDGMENT

The author thank students Dominic Calleja and Ted Stokes for their contributions to the UQ and meshing aspects of this work, respectively. Paul Carman and Martin Kocan provided important information, including CAD files, about the JET PFCs. Gabor Szepesi and Nick Hawkes provided eqdsk files.

Special thanks are due to Scott Silburn for the profile data and much advice besides. To obtain further information on the data and models underlying this article, please contact PublicationsManager@ukaea.uk. Views and opinions expressed are, however, those of the author 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.

REFERENCES

- [1] W. Arter, "Error estimation for plasma power deposition in tokamak first wall designs," *IEEE Trans. Plasma Sci.*, vol. 52, no. 5, pp. 1909–1918, May 2024, doi: [10.1109/TPS.2024.3399929](https://doi.org/10.1109/TPS.2024.3399929).
- [2] W. Arter, E. Surrey, and D. B. King, "The SMARDDA approach to ray tracing and particle tracking," *IEEE Trans. Plasma Sci.*, vol. 43, no. 9, pp. 3323–3331, Sep. 2015, doi: [10.1109/TPS.2015.2458897](https://doi.org/10.1109/TPS.2015.2458897).
- [3] W. Arter, V. Riccardo, and G. Fishpool, "A CAD-based tool for calculating power deposition on tokamak plasma-facing components," *IEEE Trans. Plasma Sci.*, vol. 42, no. 7, pp. 1932–1942, Jul. 2014, doi: [10.1109/TPS.2014.2320904](https://doi.org/10.1109/TPS.2014.2320904).
- [4] W. Arter, G. Fishpool, D. King, H. Leggate, V. Riccardo, and E. Surrey, "Modelling of tokamak plasma interaction with engineered surfaces," in *Proc. 14th World Congr. Comput. Mech. (WCCM) ECCOMAS Congr., Online Video*, 2021. [Online]. Available: <https://slideslive.com/38943983>
- [5] T. Eich, B. Sieglin, A. Scarabosio, W. Fundamenski, R. J. Goldston, and A. Herrmann, "Inter-ELM power decay length for JET and ASDEX upgrade: Measurement and comparison with heuristic drift-based model," *Phys. Rev. Lett.*, vol. 107, no. 21, Nov. 2011, Art. no. 215001.
- [6] S. A. Silburn et al., "Mitigation of divertor heat loads by strike point sweeping in high power JET discharges," *Phys. Scripta*, vol. T170, Dec. 2017, Art. no. 014040.
- [7] J. Ahrens, B. Geveci, and C. Law, "ParaView: An end-user tool for large data visualization," in *Visualization Handbook*. Amsterdam, The Netherlands: Elsevier, 2005.