

Final Report for the FY2018 Fusion Theory and Simulation Milestone on Plasma Materials Interaction (PMI)

David Bernholdt¹, Sophie Blondel², John Canik¹, Mark Cianciosa¹, Davide Curreli³, Russ Doerner⁴, Jon Drobny³, Wael Elwasif¹, David Green¹, Ane Lasa¹, David Martin², Larry Owen¹, Philip C. Roth¹, Guin Shaw², Li Yang², Tim Younkin², and Brian D. Wirth^{1,2}

¹ Oak Ridge National Laboratory

² University of Tennessee, Knoxville

³ University of Illinois, Urbana-Champaign

⁴ University of California, San Diego

The FY-2018 FES Theory and Simulation Performance Measure Management Target

“The interaction of the boundary plasma with the material surfaces in magnetically confined plasmas is among the most critical problems in fusion energy science. In FY2018, perform high performance computing simulations with coupled boundary plasma physics and materials surface models to predict the fuel recycling and tritium retention of the ITER divertor for D-T burning plasma conditions, accounting for erosion, re-deposition and impurity transport in the plasma boundary, and an initial evaluation of the influence of material deposition on the recycling and retention.”

Executive Summary

This report documents the activities performed to complete the DOE Office of Fusion Energy Sciences FY2018 Performance Measure Management (PMM) reportable milestone in Fusion Theory and Simulation, which involves modeling the plasma material interaction (PMI) in the ITER tungsten divertor during a deuterium (D) – tritium (T) burning plasma discharge. The interaction of the boundary plasma with the divertor and plasma facing component (PFC) surfaces in magnetically confined plasmas is among the most critical problems in fusion energy science. In FY2018, this milestone effort involved performing high performance computing simulations with integrated boundary plasma physics and materials surface models to predict the fuel recycling and tritium retention of the ITER divertor for D-T burning plasma conditions, accounting for erosion, re-deposition and impurity transport in the plasma boundary, and also included an initial evaluation of the influence of material deposition on the recycling and retention.

PMI produce long time scale evolution of plasma exposed surfaces and involve processes occurring at a wide range of time and spatial scales [1-8]. Thus, simulation of the plasma material interface by its nature requires applying and integrating multiple material and plasma models. Our PMI modeling and code coupling strategy has involved using the Integrated Plasma Simulator (IPS) framework [9] developed within the AToM SciDAC project. Further, we take advantage of the separation of time scales that govern important physical phenomena involved in PMI and the trace impurity approximation for modeling the transport and fate of eroded surface tungsten atoms, which enables us to perform these aspects of the PMI modeling sequentially.

The file based code integration provides the background plasma, impurity generation through sputtering yields, transport and re-deposition, and the ion implantation profiles. We then couple a binary collision approximation code with a continuum, spatially dependent, drift-diffusion-reaction cluster dynamics code to predict the tungsten surface height and composition evolution, along with the sub-surface gas diffusion and clustering evolution that controls retention and permeation of hydrogenic species.

The integrated modeling of tungsten (impurity) erosion, transport and re-deposition, and the D, T and helium (He) implantation and retention in the tungsten (W) divertor of ITER begin with SOLPS [10-12] simulations of the equilibrated background plasma, which define the density, temperature, flow, gradients, etc., in the divertor. The background plasma conditions inform hPIC [13,14], a particle-in-cell code to assess the effect of the electric and magnetic sheath on the ion energy-angle distributions (IEAD) to the tungsten divertor. Given the shallow magnetic field angle expected for the ITER divertor, a magnetic pre-sheath will form in front of the surface. Furthermore, at the most shallow inclination angle locations, the Debye sheath will vanish and the magnetic pre-sheath will dominate. In order to capture the plasma sheath physics, including both the magnetic pre-sheath and the full-orbit effects, we have performed hPIC simulations using the plasma profiles provided by SOLPS plasmas to obtain accurate and highly spatially resolved IEADs.

The background plasma and IEADs are provided to both F-TRIDYN [15] and GITR [16]. The erosion rates and energy / angular distribution of sputtered particles (used by GITR), as well as ion implantation profiles needed to compute the sub-surface gas and surface evolution (by Xolotl) are calculated by the binary collision approximation code Fractal (F)-TRIDYN, which accounts for key parameters in sputtering, such as ion impact energy, angle and surface roughness. F-TRIDYN also accounts for substrate composition when characterizing ion implantation profiles, and surface damage to the material. GITR models the gross and net erosion of plasma facing components as well as mapping the kinetic re-distribution and re-deposition of impurities to evaluate the gas implantation conditions which are passed to Xolotl [17-19], as well as to track the sputtered tungsten (impurities) and evaluate the re-deposition. The drift-diffusion-reaction cluster dynamics model Xolotl is then used to predict the surface evolution, and sub-surface gas dynamics controlling recycling, retention and permeation. F-TRIDYN and Xolotl are interactively coupled together using IPS, and a more detailed description of both IPS and the codes used in this analysis is provided in Appendix A.

The following physics goals were accomplished as part of this Theory and Simulation PMM Target:

- Experimentally validated PMI modeling predictions that helium implantation during mixed hydrogen-helium plasma exposure of tungsten leads to increased near surface concentrations of hydrogen isotopes, but decreased permeation of hydrogenic species into the tungsten.

- Integrated PMI modeling predicts that a 10-second, 40 MW helium plasma discharge in ITER (containing 5% hydrogen) will produce net deposition in the tungsten divertor across a poloidal region spanning from the strike point ($R-R_{sep} = 0$ cm) to about $R-R_{sep} \sim 20$ cm, with a net erosion region that extends along the poloidal direction from $20 \text{ cm} < R-R_{sep} < 120$ cm. The implanted helium concentrations are predicted to be in excess of 100 atomic parts per million (appm) from $10 \text{ cm} < R-R_{sep} < 40$ cm, resulting from variations in ion flux and temperature although the helium recycling coefficient remains nearly 1.0 for all locations.

- Integrated PMI modeling predictions of a 100 MW, D-T burning plasma discharge in ITER indicate that neon is the dominant contributor to the tungsten sputtering outside the private flux region, and that the size of the net deposition and erosion regions decrease (relative to the lower power, helium-only discharge) to about $2 \text{ cm} < R-R_{\text{sep}} < 10 \text{ cm}$ and $10 \text{ cm} < R-R_{\text{sep}} < 45 \text{ cm}$, respectively. Within this region, the highest hydrogenic species and helium concentrations are predicted to occur. The recycling coefficients of hydrogenic species are essentially 1.0, with some deviation predicted to occur from about $2 < R-R_{\text{sep}} < 45 \text{ cm}$. However, it is important to point out that these simulations assumed initially pure tungsten divertor tiles without prior exposure to helium plasma discharges.

- Integrated PMI modeling predictions indicate a strong effect of the peak heat flux within the region from $0 \text{ cm} < R-R_{\text{sep}} < 20 \text{ cm}$ in increasing tungsten divertor surface temperature to values from 450 to 525 K, and that these higher surface temperatures lead to a decrease in the peak hydrogenic species concentrations. Again, it is important to point out that these simulations assumed initially pure tungsten divertor tiles without prior exposure to helium plasma discharges.

- PMI predictions indicate that much higher T concentrations (by nearly 10X) develop as a result of the initial sub-surface helium concentrations implanted during prior helium plasma discharge. As well, and consistent with experimental observations and PMI modeling predictions of linear plasma devices, this near surface helium cluster microstructure is predicted to reduce the deeper permeation of hydrogenic species, at least for the approximately 1 second interval reported here. This indicates that the pre-implantation of helium into the tungsten divertor tiles could produce complex hydrogenic species retention behavior during the burning plasma operation phase of ITER, in which the very near surface regions exhibit increased tritium retention at potentially strongly bound trapping sites while the net permeation through the divertor is decreased.

This Theory and Simulation PMM target has demonstrated the existence of a new, high-fidelity PMI modeling capability, involving multiple boundary plasma and materials surface evolution codes, which can predict the fuel recycling and tritium retention of the ITER divertor for D-T burning plasma conditions. Our integrated simulations fully account for erosion, re-deposition and impurity transport in the plasma boundary, and the influence of material deposition and evolving divertor surface composition on the hydrogenic species recycling and retention. While this demonstration has initially simulated relatively short time plasma discharges, we have nonetheless documented several key physical insights, and we anticipate future effort to extend our ability to model longer timescales, in addition to incorporating additional physics related to modeling the effect of radiating impurities such as neon, incorporating the influence of beryllium erosion and re-distribution from the main PFC walls, and developing modeling capability to assess the feedback from PMI on the boundary plasma physics.

1. Completion of the First Quarterly Milestone

Document the ITER geometry to be modeled, describing the divertor geometry spanning multiple divertor tiles encompassing regions of expected net erosion and net deposition. Specify the codes involved, the code coupling strategy and information flow. Define the initial thermodynamic and kinetic database of W-He-D-T interactions. Describe linear device (e.g.,

PISCES) experimental measurement conditions, and document the geometry to be modeled, that provide opportunity for benchmarking modeling predictions of *W* gross erosion and re-deposition, and *D* recycling and retention.

1a. Summary

The first quarter milestone defined the code coupling strategy for ITER simulations, utilizing IPS [9]. Additionally, the first quarter milestone report fully defined all of the simulation codes in our workflow, and we documented the ITER geometry to be modeled during the FY2018 FES theory and simulation PMM target, which spans multiple divertor tiles to encompass regions of net erosion and net deposition. Within this quarter, we also documented the thermodynamic and kinetic database to be used to model the behavior of helium and hydrogen in tungsten. Further, we provided an initial experimental validation of our impurity transport modeling of sputtered tungsten in the linear plasma device by benchmarking the modeling predictions of tungsten erosion, re-deposition and transport of tungsten under helium plasma exposure to experimental measurements of the WI light emission and the mass gain of tungsten deposited on titania beads mounted on a rod designed to collect eroded species in PISCES-A.

1b. Details of First Quarter Progress

The ITER geometry that we have identified for the modeling effort within the FY2018 milestone has been chosen to meet two principle criteria: i) to ensure that all sub-systems include most relevant areas (for each model) and are resolved at sufficient resolution; ii) to optimize our use of computational resources. The latter is achieved by taking advantage of the toroidal symmetry of the system, given that we model the steady state solution of a tokamak plasma, and by limiting the scope of each model to the geometry essential to that sub-system. That is, we will only model a 2D poloidal cross section to calculate the equilibrium plasma background, a single divertor cassette for modeling impurity migration and re-deposition, and select representative divertor tiles with most complete representation of the surface physics (e.g., surface growth, erosion, variation in gas retention, etc).

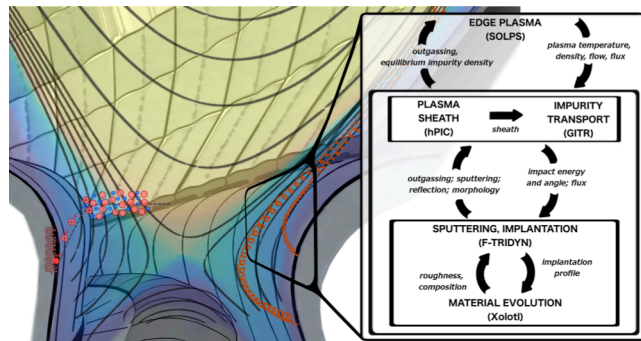


Figure 1. Schematic illustration of a tokamak divertor outlining the magnetic field line structure, plasma density (gradient), sputtered impurity trajectories (orange) and potential drop at the sheath for ions (red) and electrons (blue); next to the code-coupling workflow to model PFC surface evolution.

PMIs produce long time scale evolution of plasma exposed surfaces and involve processes occurring at a wide range of time and spatial scales. Thus, simulation of the plasma material interface by its nature requires applying and coupling multiple material and plasma models. In order to predictively model the plasma surface interface, and the natural coupling across that interface, we have used a code coupling workflow as illustrated in Figure 1, along with a

schematic illustration of the divertor region of a tokamak. As we have identified the PMI modeling and code coupling strategy, we have focused on taking advantage of the separation of time scales that govern important physical phenomena and the trace impurity approximation for modeling the transport and fate of eroded surface tungsten atoms that enable us to perform the PMI models sequentially that provides the ability to use a file transfer based code integration.

For steady-state plasmas, only the equilibrated plasma conditions are of interest. Therefore, the plasma characteristics (n_e , T_e , T_i , etc.) are calculated only once for a steady-state plasma and are assumed to remain constant throughout the entire simulated time of impurity migration, re-deposition and surface evolution. Further, impurity transport calculations assume the trace impurity approach/approximation, namely that the concentration of species we track (W) is negligible compared to that of main plasma ions (D , T , He). Therefore, the impurity transport model neglects feedback to the plasma background and impurity-impurity interactions. Given that the plasma is in steady-state, and that the surface evolves over much longer time scales ($O(s)$) than impurity transport ($O(ms)$), migration (e.g., fraction and pattern of re-deposited particles) calculated using these steady-state plasma parameters as input is also representative of steady-state conditions. However, changes in surface characteristics such as roughness or composition may involve characteristic length and time scales of impurity migration ($O(mm$ and $ms)$) thus requiring an update of the transport calculation. In fact, we anticipate that the main non-equilibrium sub-system of our model is the description of a substrate exposed to continuous gas implantation. Atomistic modeling techniques, such as the binary collision approximation provide the initial profiles of ion implantation and damage to the substrate (sputtering yields, primary knock-on atom rates, etc.). The fractal TRIDYN code for the eroded particles and plasma ion implantation is run for steady-state conditions that is essentially time independent, until statistics are sufficient for smooth, stable profiles. The evolution of the implanted species below the tungsten divertor tile surface, in addition to any surface roughness or patterning, is then calculated by Xolotl [17-19], our continuum based cluster dynamics code. The model includes changes in substrate composition and surface morphology, which greatly impact the interaction of energetic ions with the substrate. Thus, ion implantation and substrate damage data input to the cluster dynamics code is updated to ensure a complete and correct description of the (plasma exposed substrate) system. That is, the time-dependent and time-independent models are run in a loop (two-way coupling). However, these processes are loosely coupled (modest information exchange rate) as also confirmed by sensitivity analysis on the coupling frequency. A file-based system is therefore sufficient to exchange the time-evolving data. Thus, we take advantage of this time-scale separation (in steady-state conditions), to perform the plasma and material simulations sequentially, and exchange information a file-based system. The IPS [10] provides the framework and serves to integrate our models. A more detailed description of IPS and the codes we have used within the PMM target is provided in Appendix A.

Dedicated linear machine experiments have been performed in PISCES-A [20] to benchmark the integrated PSI model. In the PISCES experiments, we measure W surface erosion, transport, re-deposition and evolution of the sub-surface composition caused by the exposure to pure He and mixed He-D plasmas. The plasma column strikes the center of the W base plate, which has a radius of 48.75 mm, is biased to obtain 250 eV incident plasma ions and the temperatures is held between 400 and 600 °C in order to avoid a regime where nanoscale fuzz is formed on the tungsten surface. To measure W that is transported along the Z -axis away from the target, a titania tower is mounted in one of the holes that sit on a radius of 44.6 mm, outside the plasma column to avoid any W re-erosion. The first bead on the tower is a 12.75 mm high ceramic

(TiO₂) break, acting to isolate the tower beads from the target bias voltage. 13 titania beads sit on top of the ceramic bead and each TiO₂ bead is 10 mm tall and has an outer radius of 5 mm. The linear magnetic field strength in the target region is 800 G, perpendicular to the target to minimize any sheath effects that are commonly seen at more oblique angles. The background neutral pressure is 2.5 mTorr.

Initial experiments involving He plasma exposure were performed for time durations from 5000 to 10000 s, using high and low plasma fluxes ($\Gamma=4$ and $0.5 \cdot 10^{18} \text{ cm}^{-2}\text{s}^{-1}$ for high versus low flux, respectively). After the plasma exposure, the mass loss was measured from the W base plate along with the mass gain of each of the 13 TiO₂ beads, and the axial profiles of W I (429.4 nm) and He I (447.1 nm) light emission near the W base plate. Langmuir probe data was used to obtain the radial profiles of n_e , T_e and particle flux. The mass difference of the target plate is an indication of the erosion source strength as well as re-deposition, and thus net erosion. The mass difference of the titania tower collector beads is an axial measurement of long range material migration and deposition. The axial array of W I emission spectroscopy sight lines is an indication of the rate of transport and ionization near the target plate. To compare with the experimental values, the input to GITR used to model this initial PISCES experiment includes fits to the experimental profiles measured by Langmuir probes that are 30 cm upstream from the target.

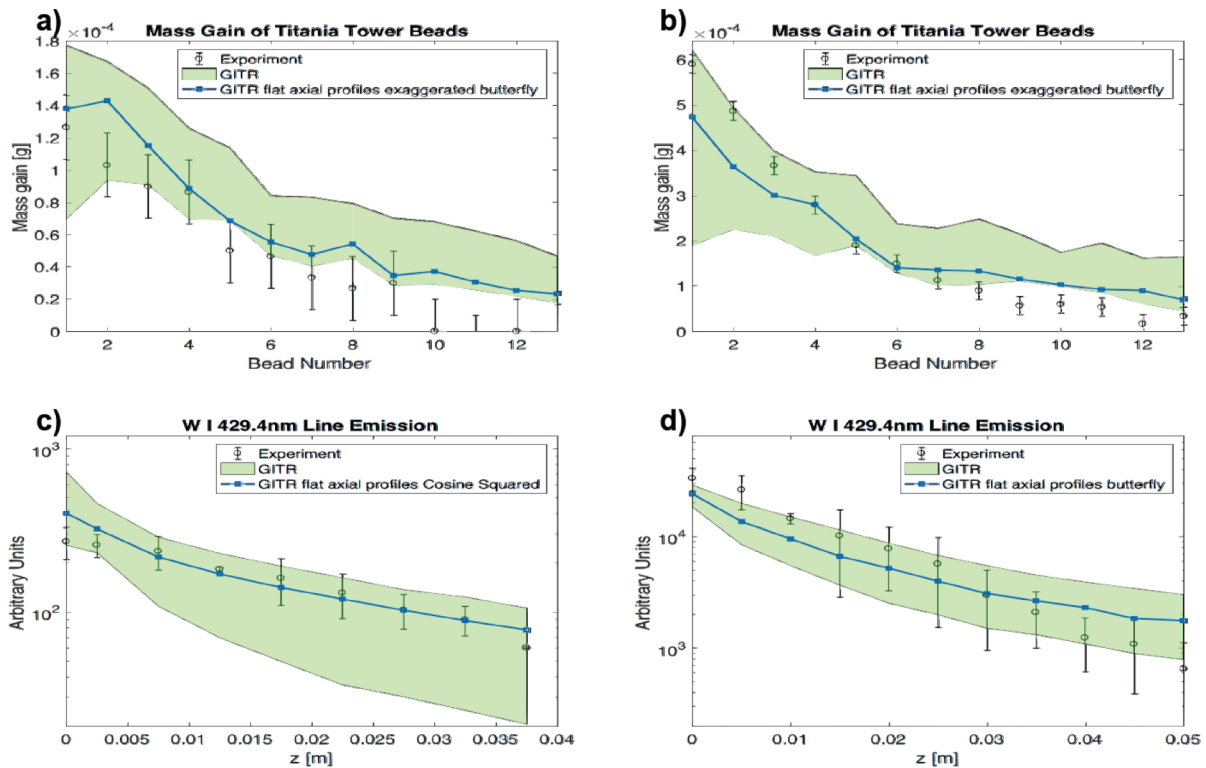


Figure 2. Comparison of GITR modeling predictions (blue data points/line plus light green error band) to PISCES experimental measurements (black circles) for a 250 eV helium plasma exposure of tungsten. a) and b) W mass gain on the titania tower beads as a function of position at low versus flux, respectively. c) and d) W I emission (429.4 nm) measured as a function of height above the W surface for low versus high flux, respectively.

Figure 2 shows a comparison between the GITR predictions and the experimentally measured PISCES data. The GITR predictions include an error band, which has been calculated based on performing 20 independent simulations that vary the plasma temperature and profile within the experimentally measured error bars in addition to utilizing different angular distributions of sputtered tungsten. For the low flux exposure (Fig. 2a), experimental measurement of mass loss of the W target plate is well reproduced by GITR, within the uncertainty caused by input parameters mentioned above. The mass gain of the titania beads also shows good qualitative and quantitative agreement between experiments and the model (Fig. 2a). There is a slight over-prediction of the mass gain of the beads far from the erosion source. The variance in the “bead 1” mass gain comes predominately from the variation in sputtered W angular distribution. The synthetic diagnostic of WI 429.4 nm line emission decreases at a function of distance moving away from the target plate, at a similar fall-off rate both in experiments and in GITR (Fig. 2c). Although the trend has a slightly different shape, dipping down in the model but rolling over in experimental measurements, the variations in the simulations from this initial sensitivity analysis are on the order of the experimental error bars.

The high flux exposure shows similar qualitative trends to the low flux case (Fig. 2b and 2d), with the magnitude of the observable features in experiment and simulation increased relative to the low flux exposure conditions (Fig. 2a and 2c). The experimentally measured mass loss of 79.5 mg is well reproduced utilizing the original fits to the background plasma profiles. This high flux case shows an increased percent of eroded material returning to the target (with predicted values ranging from 65-81%), which is expected compared to the low flux case (with predicted values from 32-67%). The variability in the mass gain of the titania tower beads at the base is once again present as is the slight over-prediction of mass gain of the beads at the top of the collector tower (Fig. 2b). Good qualitative agreement is shown for the spectroscopy although the slight difference in fall-off shape remains (Fig. 2d).

These initial comparisons for a helium only plasma exposure provide confidence in the predictions of GITR in advance of modeling the more complex helium – deuterium plasma conditions, before adding tritium. Although not included in the first quarter milestone report, we subsequently characterized the surface response and sub-surface gas concentrations using laser-based characterization techniques to compare to Xolotl predictions. That comparison is discussed in the second milestone report.

2. Completion of the Second Quarterly Milestone

Document the computational performance and coupled plasma physics - materials surface predictions of coupled code simulations of linear device experiments with mixed He-D plasma conditions. Describe results of modeling comparisons of W gross versus net erosion and D recycling/retention using coupled code PMI strategy for these linear device experiments. Document initial background plasma conditions calculated for the ITER geometry with He (only) plasma operation.

2a. Summary

During the second quarter, we successfully completed the task of assessing the computational performance of coupled plasma physics and materials surface simulation codes in modeling an

experiment in which tungsten was exposed to mixed D-He, versus D-only plasma exposure in the linear plasma device PISCES-A. We then successfully compared the initial experimental measurements to the coupled code modeling predictions, with a focus on the net versus gross erosion of W, the recycling and retention behavior of deuterium and the depth profiles of helium and deuterium. Finally, we described the initial background plasma conditions in ITER for He plasma operation, which provided the basis of the coupled code plasma surface interaction modeling effort that was performed in the third quarter.

2b. Details of Second Quarter Progress

As described previously in Section 1b and in Appendix A, we used the IPS [9] framework to integrate F-TRIDYN, GITR, and Xolotl. Our primary focus in this quarter is on the performance of the coupled code workflow because this performance measurement establishes a baseline against which we will compare performance and scalability as we make optimizations to the individual components, and the coupled workflow during the remainder of FY2018. It also defines a procedure for collecting, analyzing, and visualizing performance data about our workflow from the IPS framework. IPS calls two sub-driver components, defined as *SD1* and *SD2*, that simulate tungsten impurity transport and plasma surface interactions, respectively.

For this milestone, we simulated a 10% helium – 90% deuterium plasma environment in the PISCES-A linear device in which our *SD1* component performed 10^5 GITR time steps involving 10^5 particles and F-TRIDYN modeled 10^5 particle impacts. The *SD2* component performed Xolotl simulations with 256 grid points and a reaction network with maximum cluster sizes of 8 He interstitials, 50 tungsten vacancies, and 6 tungsten interstitials. These simulations were performed in a 48-hour batch job on 16 compute nodes of the Edison Cray XC30 at the US Department of Energy National Energy Research Scientific Computing Center (NERSC) [21]. Edison has 5586 compute nodes, each with 64GB memory and two Intel Xeon E5-2695 “Ivy Bridge” 12-core processors running at 2.4Ghz. The Edison compute nodes are connected using a Cray Aries interconnect with Dragonfly topology.

We configured our workflow so that the F-TRIDYN task pools and GITR run used all available processor cores, but ran the memory-bound Xolotl with four MPI processes per node. Figure 3 shows a high-level timeline visualization of our original simulation. The events used to generate this visualization were taken from the IPS framework event log and each line in this visualization reflects an activity within the workflow (rather than a specific process or thread). Because the simulation as configured did not reach its target end time before the 48 hour time limit of the batch job, we added synthetic activity end events so that the visualization would more clearly show the activities that were in progress when the job ended. In addition to giving us timings for each activity in our workflow, the event log and our visualization exposes the nesting, ordering, and dependencies within our workflow, including the iterative nature of *SD2* between F-TRIDYN and Xolotl. It also clearly shows that the vast majority of the workflow time was spent in GITR and Xolotl during this original simulation (Fig. 3a), rather than the F-TRIDYN tasks or IPS itself. Because the amount of time spent in GITR was so much larger than what we expected, we hypothesized that our IPS workflow configuration was not correctly specifying the OpenMP environment variables needed for GITR to use all available processor cores. We modified our workflow configuration to specify those environment variables, and re-ran the 48-hour job. Fig. 3b documents that we observed much better GITR performance.

Whereas our run with the original configuration allowed Xolotl to achieve only 0.28 s of simulated time, the improved GITR performance resulting from the modified workflow configuration allowed Xolotl to reach 7.95s of simulated time, an improvement of over 28X.

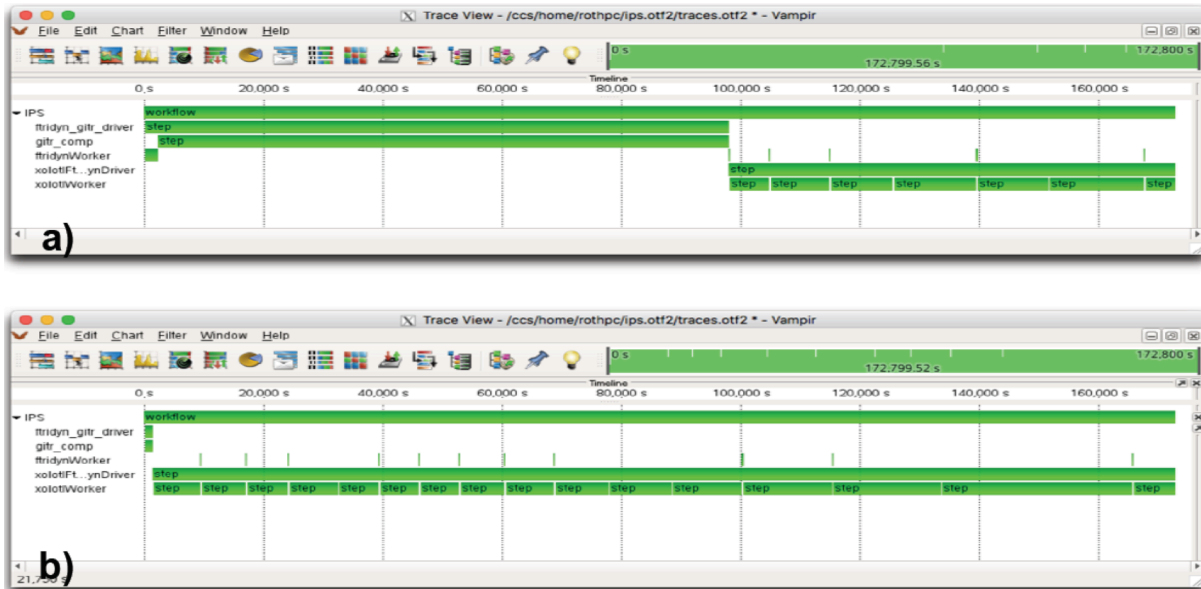


Figure 3. Timeline visualization of a) original 48-hour coupled IPS/F-TRIDYN/GITR/Xolotl simulation of He-D plasma in the PISCES linear device, and b) after correctly setting the OpenMP environment variables for GITR.

Xolotl predicts the sub-surface diffusion and clustering behavior of the retained gas species, as we have extensively reported for the case of helium implantation below tungsten surfaces and computationally benchmarked against very high flux molecular dynamics simulations [17]. In this FY2018 fusion theory modeling effort, we have extended Xolotl within the coupled code simulation framework to predict the sub-surface helium and deuterium evolution. Deuterium will not self-cluster, but is predicted to partition to and bind with both vacancies and helium-vacancy clusters. Figure 4 shows a comparison of the Xolotl predictions and laser-induced breakdown spectroscopy (LIBS) and laser ablation mass spectrometry (LAMS) experimental measurements of the He and D concentration as a function of depth below the surface. Our laser ablation measurements (LIBS and LAMS) were performed in ultra-high vacuum (3×10^{-8} Torr), and use a Nd:YAG laser with a 532 nm wavelength coupled to an energy modulator to produce a 5 ns pulse width and an energy of 1.8 mJ to produce an ablation crater with depth less than 100 nm. We have previously performed an experimental assessment to indicate the reproducibility of the ablation crater, and to demonstrate our ability to detect sub-surface helium [22]. The ablated material plume is first analyzed by filterscopes to provide the LIBS measurements, and then pumped with a cryogenic vacuum pump into a quadrupole mass spectrometer, which can provide sufficient mass to charge resolution to distinguish between atomic deuterium and helium. We refer to this technique as laser ablation mass spectroscopy (LAMS).

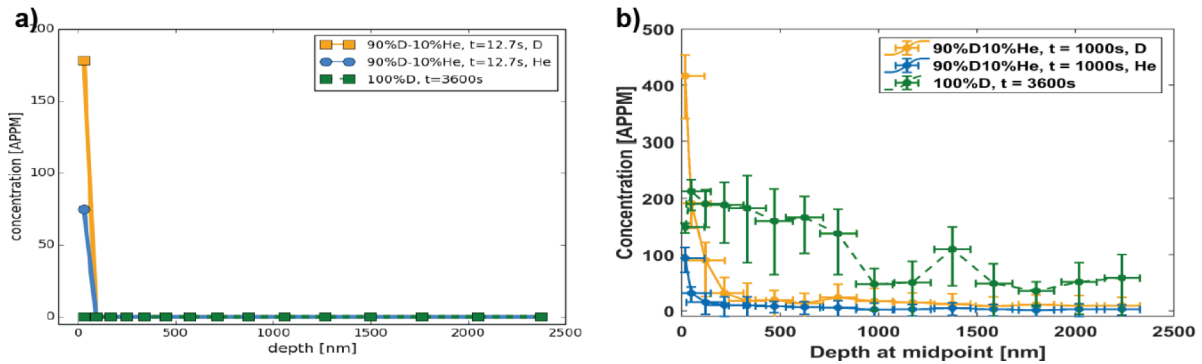


Figure 4. a) Xolotl predictions of the depth dependent D (green or yellow) and He (blue) concentrations following exposure to 90%D-10%He, or 100%D, plasmas in PISCES-A with a bias voltage of 250 eV. The accumulated simulation time is shown in the legend. **b)** LIBS and LAMS measurements of the depth dependent D and He concentration. The experimental exposure times in PISCES-A are shown in the legend.

In comparing Fig. 4a and 4b, we see very similar trends in the D and He depth dependent concentrations for the mixed 90%D-10%He plasma exposure, although the simulated time of ~ 13 seconds is significantly less than the experimental time of 1000 s. These simulations are ongoing, and we believe that the similar depth dependent profiles do indicate successful validation of the model, with the anticipation that continued simulation of the 90%D-10%He plasma exposure of tungsten will bring the predicted concentration up to the levels experimentally obtained. The most striking difference with the predictions relates to the pure D exposure. In our modeling, we have assumed elementally pure, single crystal tungsten containing a thermal equilibrium concentration of vacancies ($\sim 10^9 \text{ m}^{-3}$), without any impurities or pre-existing dislocations or grain boundaries that are known to trap hydrogen. Since helium naturally self-clusters, and is a trap for H isotopes, the increase in D concentration predicted by the Xolotl simulations relates entirely to the trapping interaction of D at small He clusters that nucleate below the tungsten surface. In the case of D-only plasma exposure, there are no atomic-scale or microstructural features included in the model (beyond the extremely low, thermodynamic equilibrium vacancy concentration), and the D is predicted to permeate through the sample without any noticeable buildup in concentration. However, the experimental measurements in Fig. 4b clearly indicate that the presence of He in the plasma greatly reduces the total amount of D retained in the tungsten, although with a much larger near surface concentration of D that is co-located with the helium. Future simulations will incorporate vacancies as a surrogate for the approximately 20-100 appm of impurities that are known to exist in tungsten, and for which we believe will improve the comparison with the experimental results for the case of the D-only plasma exposure.

Finally, the background scrape-off layer and divertor plasma has been calculated for a helium plasma discharge in ITER using the SOLPS code [10-12], which couples a 2D fluid transport model with a Monte Carlo treatment of neutral transport in the plasma edge. Two primary scenarios for ITER are of interest for the plasma-surface interaction simulations targeted as part of this Milestone: full-power operations, and a representative scenario during the helium campaign. The full power case represents burning plasma conditions, and is characterized by high power levels, strong impurity presence as needed to radiate the heat flux, and a mixed plasma species made up predominantly of the main-ion deuterium and tritium ions but with a significant helium component due to helium ash. The helium campaign is planned for early in ITER operations, with helium plasma being used to enable H-mode access in the initial reduced magnetic field scenarios. While not presenting the same level of power load challenge, these plasmas are of special interest for the PMI simulations due to the intense interactions between helium and the tungsten surface anticipated.

In the SOLPS simulations, the input power was set to 40 MW, representative of these early ITER operations. The transport was fixed as spatially constant, with values for the particle and electron/ion thermal diffusivities of $D = 0.3$ and $\chi_e = \chi_i = 1.0 \text{ m}^2/\text{s}$. These represent the standard values for ITER SOLPS simulations, and produce SOL solutions that have been well documented. Helium and hydrogen species are included in the simulation; the plasma is predominantly helium, with $\sim 5\%$ hydrogen content included to model residual hydrogen present in the system (potentially due to pellet fueling). Realistic pumping in the divertor is included, and the plasma is fueled by edge gas puffing set to a level that produces a mid-plane separatrix density of $n_e \sim 1.5 \times 10^{19} \text{ m}^{-3}$ as shown in Figure 5. These SOLPS simulations provide the background plasma conditions at the divertor that were used in the coupled PMI code modeling of the tungsten divertor response to helium plasma operation in the Third Quarterly Milestone.




Figure 5. Radial profiles of electron temperature (top) and density (bottom) as calculated by SOLPS for a helium plasma operation in ITER.

3. Completion of the Third Quarterly Milestone

Document the computational performance and coupled plasma physics - materials surface simulations of W erosion, re-deposition and impurity transport in ITER (multiple divertor tile geometry) and He recycling/retention with a He plasma operation. Models will also predict the evolving divertor surface composition (W-He). Describe opportunities to improve computational performance for modeling W-D-T-He burning plasma conditions. Document initial background plasma conditions in ITER geometry for He-D/T burning plasma operation.

3a. Summary

During the third quarter, we successfully completed a coupled code simulation of a 10 second long helium plasma discharge in ITER using 160 nodes of the National Energy Research Scientific Computing Center (NERSC) Edison Cray XC30 system [21]. We documented the computational performance of the coupled plasma physics and materials surface simulation codes in modeling the He plasma operation of ITER, and discuss opportunities for improving performance when we simulate the PMI behavior during a burning plasma discharge. The physics predictions from these coupled PMI models show a variation in the tungsten gross versus net erosion, the tungsten surface height evolution and the sub-surface helium cluster concentration as a function of varying position along the outboard divertor. This report also describes the background plasma conditions in ITER expected for burning plasma operation with a 100 MW input power, which will be the basis of the FY2018 Theory and Modeling PMM target to couple codes in order to predict plasma surface interactions in the ITER tungsten divertor during a burning plasma discharge.

3b. Details of Third Quarter Progress

Figure 6 shows a timeline visualization of the entire workflow running for 48 hours on 160 nodes of the National Energy Research Scientific Computing Center (NERSC) Edison Cray XC30 system. In the workflow visualizations, each line within the timeline represents an activity (as opposed to a process or thread as shown in a traditional parallel program visualization), and time increases moving left to right. As shown in the visualization, the workflow begins with an F-TRIDYN+GITR component followed by an F-TRIDYN+Xolotl component that alternates between F-TRIDYN and Xolotl simulations until the end of the job. Because we are now simulating a 3D problem domain with toroidal symmetry, the F-TRIDYN+GITR component takes a slightly larger fraction of the 48-hour run time than it did in our previous simulations of the 1D PISCES experiments.

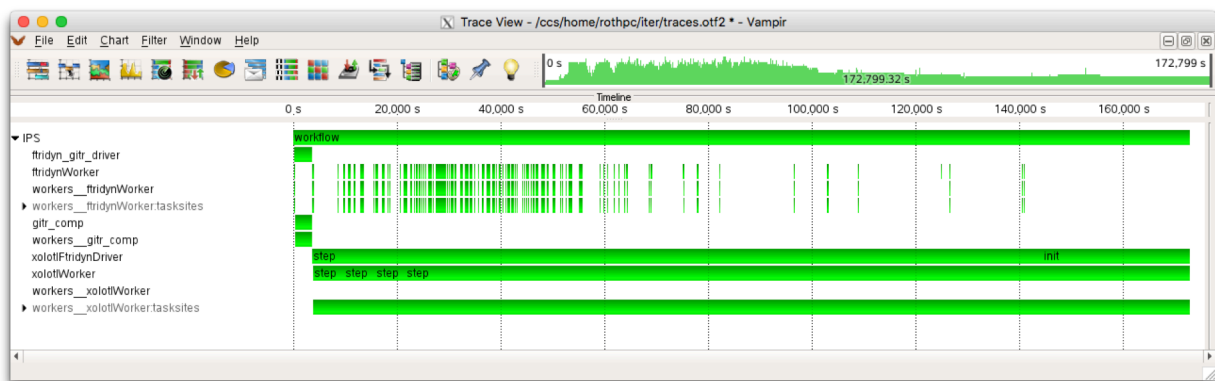


Figure 6. Timeline visualization for the ITER simulation workflow running on NERSC Edison.

The workflow retains the same high-level structure as described last quarter and the workflow is still implemented using IPS. However, the complexity of the coupled simulations has increased dramatically due to the 2D problem domain and to take better advantage of the available computational resources. In the Xolotl part of the workflow, we use multiple 1D Xolotl instances, each simulating a “pencil-like geometry” extending below surface. We used the

IPS task pool concept (i.e., performing multiple instances of the same component with different initial conditions) to run these Xolotl simulations concurrently. Because of the large number of F-TRIDYN instances needed to generate inputs for both Xolotl and GITR, we also use the task pool concept for running these single-process F-TRIDYN instances.

The tiled divertor surfaces in ITER consist of 2816 total tiles with gap widths and depths specified to 0.5mm and 2mm, respectively, as shown in Figure 7. The mono block tile faces range from 24mm to 44mm in either of the face directions. Each of the tile faces is composed of at least 2 triangular elements where the area of increased surface resolution each tile face is made up of 16 triangular elements. Including the tile faces, tile gaps, and boundaries the simulated divertor is made up of 79281 total elements. The main wall is composed of 405 elements to give a total of 79686.

Given the 3D nature of the problem, it is useful to reduce this space to a single variable for purposes of plotting and visualization. Taking advantage of the toroidal symmetry of ITER, we will commonly use the coordinate system such that we define “ $R-R_{sep}$ ” in units of distance (m). This coordinate system describes the length along the outer divertor target with respect to where the separatrix crosses the target. Thus the magnetic strike point is the reference location $R-R_{sep}=0$. Going vertically down the target is thus represented as a negative direction, while moving vertically (and poloidally) up the target and radially outward is represented as a positive direction.

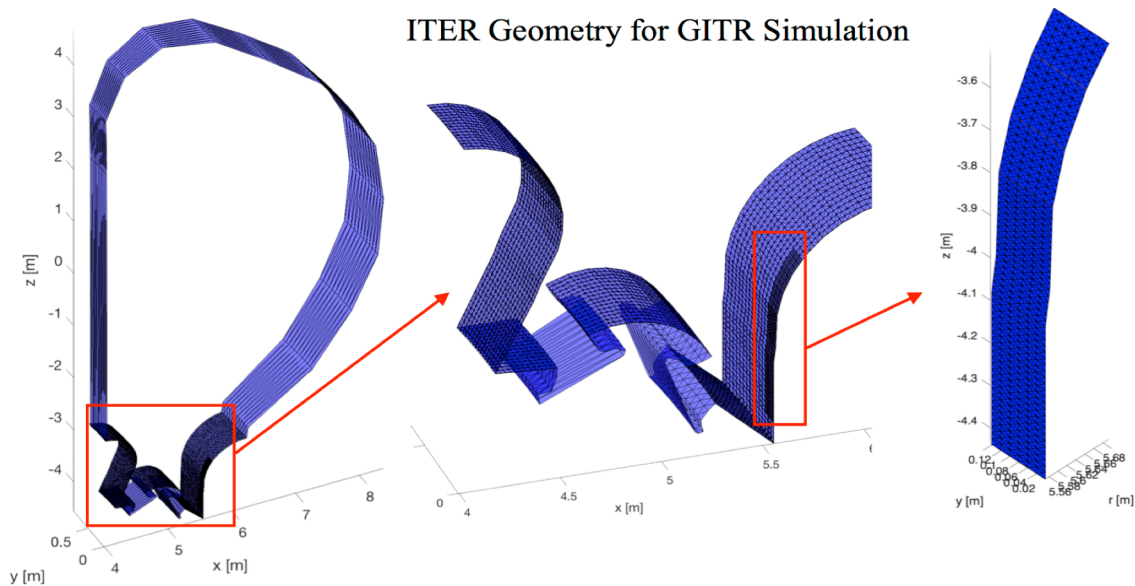


Figure 7. 3D surface geometry modeled in GITR, comprised of 1/54th toroidal symmetry (e.g., 1 divertor cassette). The divertor tiles and refined GITR mesh along the tile faces of the target are also shown.

Each of the coupled F-TRIDYN and Xolotl simulations is setup to evenly distribute the 160 nodes requested for GITR to allocate 6 nodes, that is 60 processes to each of the 25 coupled simulations. With the current substrate depth (5 μ m, i.e., 256 grid points), the optimal number of processes (in resource to speed-up ratio) is ~64. So, even without hyper-threading, 6 nodes per Xolotl simulation are sufficient. We model the surface evolution over 10 seconds of plasma

exposure. Initially we update the implantation profiles every 0.05 sec, and regularly increase this interval to 3.2 sec. GITR and hPIC provide the input for modeling the evolution of the W substrate at each location, and this input includes the concentration of different species reaching the surface as well as the impact angle-energy distributions for re-deposited W and for the background plasma species (He in the present case), respectively.




Figure 8. a) W sputtering yield as a function of incident angle and energy of He ions, as calculated by F-TRIDYN, b) effective sputtering yield and eroded W flux along the target, and c) 3D representation of the helium-induced tungsten erosion along the outboard divertor target.

Based on the background plasma data (SOLPS background plasma profiles and ion energy-angle distributions of He plasma species at the outer divertor target) F-TRIDYN produces the initial conditions for the He sputtered W source in GITR. These sputtering yields which are highly dependent on incident ion energy and angle can then be integrated with the incoming ion energy-angle distribution to produce an effective sputtering yield, as shown in Fig. 8a. The spatially resolved yields are multiplied by the local flux to produce the impurity source term to GITR, as shown in Fig. 8b, and the sputtered W energy and angle distributions are formulated using F-TRIDYN output of the kinetics. The He-eroded W source on the outer ITER divertor target is shown in Fig. 8c.

The results of the impurity transport simulation show net erosion along the ITER outer divertor target over a distance of approximately 90 cm in the positive $R-R_{sep}$ coordinate, as shown in more detail in Fig. 9a. The shape of this erosion profile mimics the erosion profile from the background He^+ and He^{++} plasma flux. The gross deposition flux is slightly less than the gross erosion flux resulting in the net erosion. In these simulations approximately 80% of the eroded tungsten is re-deposited while 20% of the eroded impurities are lost to the private flux region, scrape-off layer, or other surface locations away from the outer divertor target. Fig. 9b shows the gross erosion




Figure 9. a) Tungsten gross erosion, gross deposition and net erosion flux, as a function of position along the outboard divertor, and b) the contributions to the gross tungsten erosion flux, showing the ion responsible for sputtering, as a function of position.

flux based on the species causing the sputtering, and reveals that He^{++} is the dominant erosion source followed by He^+ and then the self-erosion caused by sputtered, ionized and re-deposited W.

Figure 10a shows the variation in tungsten surface height that is predicted to result from the combination of net erosion and the sub-surface gas evolution, while Fig. 10b shows the retained helium concentration as a function of surface location. Taking a closer look at the He concentrations predicted by Xolotl, the He retention (relative to the implanted flux) is the largest for the plasma conditions with peak plasma temperature, although the total accumulated helium inventory in the tungsten divertor is dominated by the regions of higher implantation flux. The Xolotl predictions indicate higher helium concentrations at the locations of higher implantation flux, with a value of about 75 ppm at the peak flux location and increasing to a maximum value of about 390 ppm at $R-R_{\text{sep}} = 0.2$ m, and then decreasing sharply as the implantation flux decreases and plasma temperature increases. Notably, the regions with retained helium concentrations above 50 ppm correspond to the spatial locations exhibiting the largest surface height changes associated with either net deposition or net erosion.

Finally, the background scrape-off layer and divertor plasma for a burning D-T plasma discharge of ITER was calculated using the SOLPS code, which couples a 2D fluid transport model with a Monte Carlo treatment of neutral transport in the plasma edge. This case is now documented representative of Fusion Power Operations (FPO) for use in the final milestone activities. This case is largely similar to the previous in its setup, with the main differences in input power and species included, as will be discussed in the Section below on the Details of Fourth Quarter Progress.

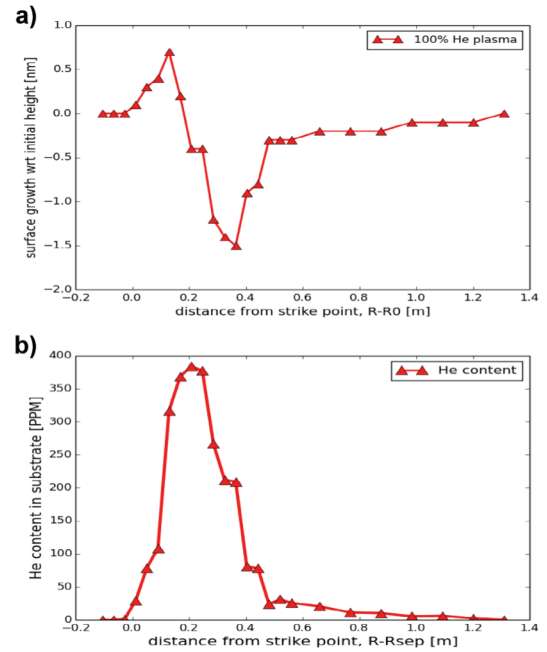


Figure 10. a) net surface height change (positive referring to growth/deposition) and b) retained He gas content as predicted by Xolotl, which accounts for surface sputtering, re-deposited W and the surface variations induced by the sub-surface gas dynamics.

4. Completion of the Fourth Quarterly Milestone

Perform coupled plasma physics - materials surface simulations of gas (He, D, T) implantation, recycling and retention in the ITER divertor (multiple divertor tile geometry) with burning plasma conditions. Document the surface erosion, transport through the plasma boundary and re-deposition. Predict T recycling and retention as a function of spatial location in the divertor, as well as the surface concentration (W-He-D/T) as a function of shot duration and spatial location. Document code performance and quantify high performance computing resources used for coupled simulations.

4a. Summary

During the fourth quarter, we successfully completed a coupled code simulation for a 100 MW, deuterium (D) – tritium (T) burning plasma discharge in ITER using 240 nodes of the National Energy Research Scientific Computing Center (NERSC) Edison Cray XC30 system, for a plasma duration of about one second. The geometry is the same as previously documented in Quarter 3, and shown in Figure 7. The toroidal symmetry of ITER allows for the use of a coordinate system such that we define “R-R_{sep}” in units of distance (m). This coordinate system describes the length along the outer divertor target with respect to where the separatrix crosses the target. Thus the magnetic strike point is the reference location R-R_{sep} = 0. Going vertically down the target is thus represented as a negative direction, while moving vertically (and poloidally) up the target and radially outward is represented as a positive direction.

The integrated simulation workflow is also similar to that previously described in the quarter 3 results for a helium (He) plasma discharge, with three notable exceptions. First, we have performed detailed F-TRIDYN and Xolotl calculations at 40 spatial locations (rather than 25 locations in Q3). Second, we are now explicitly utilizing the plasma heat flux to the tungsten divertor and calculating the resulting temperature distribution as a function of depth in Xolotl by solving the one-dimensional thermal diffusion equation. And, third, the integrated simulations now incorporate a much larger set of species, with the addition of D, T and neon (Ne). Neon has been incorporated into the boundary plasma calculations, the sheath model and used in F-TRIDYN for calculating the sputtering yield of tungsten. This information is passed to Xolotl, but at this time, Xolotl does not account for implanted neon below the tungsten surface. It is also important to note that while the SOLPS simulations include beryllium, the other simulations within our workflow do not.

4b. Details of Fourth Quarter Progress

The starting point for the integrated model of impurity erosion, transport and re-deposition, and He implantation and retention in the ITER divertor, is a SOLPS simulation performed to characterize the equilibrated background plasma (density, temperature, flow, gradients, etc.). This gives background profiles for the scrape-of-layer (SOL) plasma and divertor plasma including values at the divertor targets.

The SOLPS code couples a 2D fluid transport model with a Monte Carlo treatment of neutral transport in the plasma edge. This is the standard edge modeling tool used for ITER divertor predictions, and a large database of SOLPS runs has been generated [23]. Our prior quarterly reports documented a first SOLPS case representing early helium operations and representative of Fusion Power Operations (FPO). Both cases are largely similar in their setup, with the main differences being in input power and species included.

In the SOLPS simulations of FPO described here, the input power was set to 100 MW. The plasma is predominantly hydrogenic, with deuterium, helium, beryllium, and neon included in the simulation; and strong radiation is provided by neon. SOLPS predicts that 73 MW out of the 100 MW input power is radiated, with nearly the entirety being localized in the divertor region (with only modest differences between the inner and outer leg), and over 80% of the radiation is due to neon, which is seeded into the boundary plasma for that purpose. As a result, the parallel heat flux near the strike point is significantly reduced at the target compared to the midplane, as shown in Figure 11. The resulting heat fluxes on the target are reduced to ~ 7 MW/m² carried by the plasma onto the target plates.

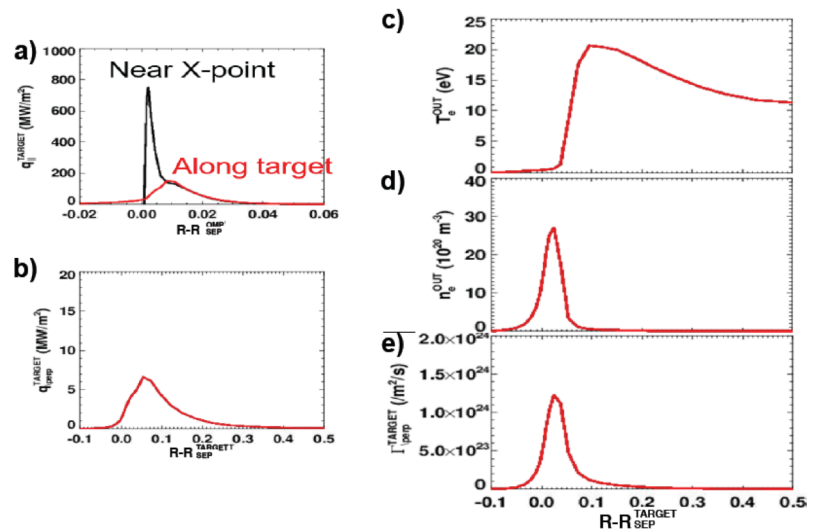


Figure 11. Radial profiles of a) parallel and b) perpendicular heat flux, c) electron temperature, d) electron density, and e) ion flux for the ITER, 100 MW D-T burning plasma discharge.

The radial profiles of the divertor electron density and temperature and ion flux are shown in Figure 11. The profile shapes are similar to the previously described helium cases, with very low temperature (\sim eV) near the strike point and increasing farther into the SOL, but now with much higher densities and ion fluxes. The electron density is very high near the strike point, reaching values of $\sim 3 \times 10^{21}$ m⁻³. Likewise the ion flux is high near the strike point, with a peak value of $\sim 1 \times 10^{24}$ m⁻²s⁻¹. The density and ion flux both decrease significantly farther into the SOL.

The results of the SOLPS simulation and the re-gridding/processing scheme were documented in the Q3 milestone report (and also described in Appendix A), and provide a well-characterized background plasma solution as well as conditions at the divertor targets. The plasma conditions at the divertor target along with the magnetic field equilibrium and the input geometry are our basis for the ion energy-angle distributions (IEADs) simulated by hPIC. It is these IEADs, in combination with the surface model from F-TRIDYN and the surface fluxes from SOLPS, that drive the background plasma erosion of W for which the migration and self-erosion is simulated by GITR as well as the hydrogen-helium implantation and resulting surface evolution simulated by Xolotl.

The 2D background plasma profiles obtained from SOLPS include temperature, density, electric potential, and flow velocities. Figure 12 shows example output for the ion density and average background plasma ion charge, showing the area covered by the SOLPS grid. This shows a gap between the simulated SOL grid and the first wall, but a fully simulated plasma in contact with the divertor targets. As a reference for the physical locations of the Xolotl-F-TRIDYN simulations as well as an understanding of the R-R_{sep} coordinate at the target used

throughout this report, the electron temperature in the divertor with a scatter of the locations along with R-R_{sep} coordinate, r, and z coordinates is shown in Figure 13.

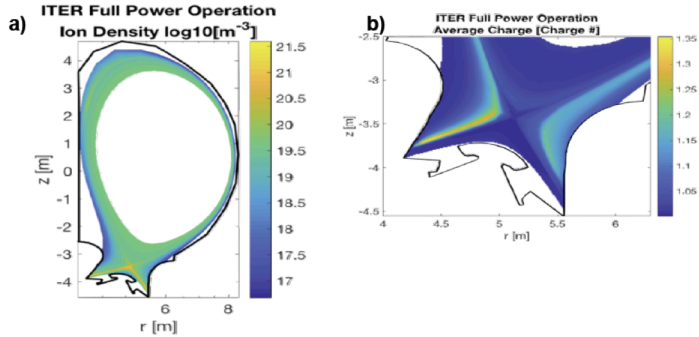


Figure 12. a) Re-gridded SOLPS output for ITER full power operation showing ion density on a log scale and the domain covered, and b) weighted average of SOLPS output (density and charge number) to provide the average charge number in the ITER divertor region.

the relative concentrations of the background ion species are also changing. By comparing Fig. 14a and 14b, it can be seen that the heat flux to the target is proportional to both the ion flux as well as the ion temperature. While the ion heat flux is the dominant contributor near the divertor strike point, the radiated heat flux (defined as the difference between the P_{tot} and P_{ion}) is a significant contributor at points moving vertically upward along the outer divertor target.

The SOLPS data of the ITER full-power D-T operation scenario have been sampled at 36 locations along the poloidal coordinate ‘R - R_{sep}’ at the outer divertor target. The data provide the plasma conditions in proximity to the surface, for a total of 18 charged species including D, T, He, Ne (each with multiple charged states). The SOLPS data are passed as an input to the hPIC Particle-in-Cell (PIC) code, a full-f, full-orbit PIC resolving the plasma sheath physics, which can now be interfaced with an external plasma fluid code, such as SOLPS or similar solver. The hPIC code has been modified to input an arbitrary number of plasma species, each species represented by a density and temperature. The hPIC code can now be interfaced with an external plasma fluid code, such as SOLPS or similar solver. The hPIC code requires information on the local magnetic field (magnitude of the field and angle with respect to the wall) at

In addition to the 2D plasma profiles, a full set of the target values are given by SOLPS. The target profiles are characterized by a peak in ion flux near the divertor strike point accompanied by very low electron and ion temperatures shown in Figure 14. Moving away from the strike point in the positive R-R_{sep} direction we see an increase in temperature and a decrease in flux until the peak temperature is reached at which point we see a monotonic decrease in plasma temperature. While the total ion flux to the target decreases with increasing distance,

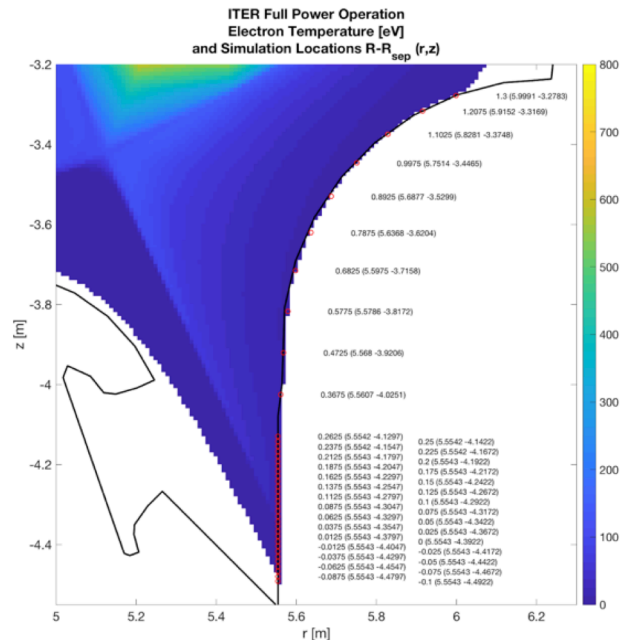


Figure 13. Electron temperature at the outer divertor target with reference points for the locations of the simulated hPIC and F-TRIDYN-Xolotl simulations given with R-R_{sep} value as well as coordinates in (r,z).

the point of interest on the target plate, plus a number of necessary discretization parameters (number of grid points per Debye length, number of time steps per gyro-period, number of ion transit times, numbers of particle per cell). hPIC provides as an output the energy-angle distribution of the ions impacting on the surface, and then passes this information to the material codes handling surface erosion and ion implantation (F-TRIDYN and GITER). One hPIC simulation per each SOLPS point along the $R-R_{\text{sep}}$ coordinate has been run in 1D3V electrostatic mode (one spatial dimension and 3 velocity dimensions).

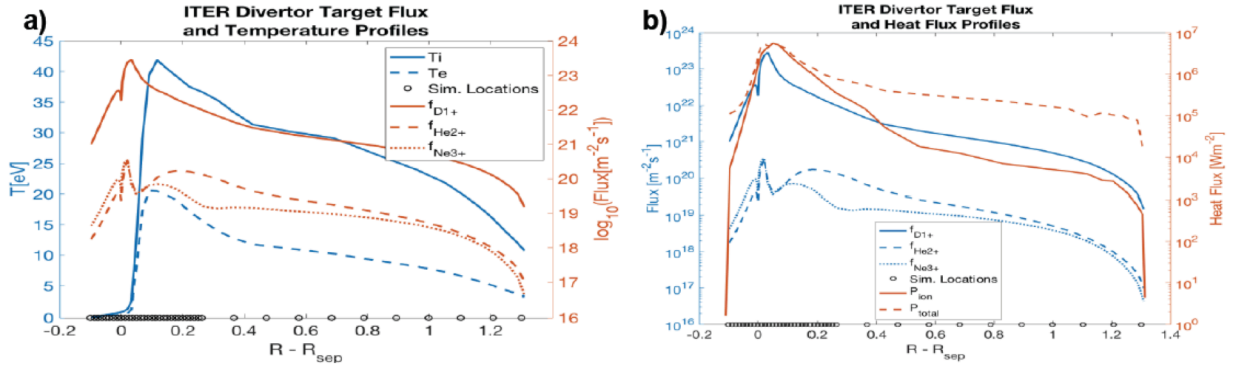


Figure 14. a) The SOLPS outer divertor target grid points for electron and ion temperature, as well as representative D, He and Ne fluxes as a function of radial position. b) Particle fluxes to the outer ITER divertor target, and the parallel heat flux contribution due to ions (P_{ion}). The total heat flux (P_{tot}) includes a contribution from radiation. For reference, the locations of coupled F-TRIDYN-Xolotl simulations are provided by circles in the $R-R_{\text{sep}}$ dimension.

Each simulation resolves a region encompassing a few ion Larmor radii (of a thermal tritium ion) with the spatial dimension along the direction perpendicular to the surface, using an average of 60k particles-per-cell, 6 grid points per Debye length, and 30 time steps per ion gyro-period. The IEADs of each of the 18 plasma species have been recorded at each spatial location. The distributions are saved on a grid spanning over 24 times the local electron temperature ($24T_e$), with energy bins of amplitude $T_e/10$ uniformly-spaced along the energy coordinate. The factor of 24 was empirically chosen as to properly sample the vast majority of the ion energy distribution for the ITER case. The minor fraction of particle population arriving at energies larger than $24T_e$ was collected on the bin of highest energy of the energy grid, and treated as an equivalent mono-energetic beam. We estimate that this approximation only introduces very minor uncertainties, as the sputtering yield for neon (the main contributor to erosion, as discussed below) plateaus for energies in this range ($E_{\text{in}} > 300\text{eV}$). Future improvements of hPIC will include options of re-binning the IEADs on the fly, for a better representation of the high-energy tails. Figure 15 shows several examples of IEADs predicted by hPIC at three selected points along the $R-R_{\text{sep}}$ coordinate (below, at, and above the strike point), of the species (D , T , Ne^{2+}).

The individual plots within Fig. 15 report the distribution of the ions at the wall on a two-dimensional energy-angle plane, where the angle is measured with respect to the normal to the surface (0° means that the ions arrive at normal incidence on the surface, 90° means parallel to the surface). The two most abundant plasma species, deuterium and tritium, are thermalized, and have very similar behavior in the energy-angle phase-space, except for small differences due to their masses. The top red line on each plot marks the most probable energy (eV), as expected from the classical sheath theory modified to account for the relative concentrations of a multi-

species plasma. Most of the D-T ions hit the surface at angles between 65° and 89.9° ; such an angular interval is as expected for ITER's magnetic angles comprised between 86.8° and 89.6° [24]. Indeed, the ion angular distributions at the wall are primarily affected by the local value of the magnetic inclination with respect to the surface. In ITER's DT case, the most probable impact angle of the DT ions at the surface is comprised within the interval 70-80 degrees for all cases under analysis.




Figure 15. hPIC ion energy-angle distribution (IEAD) functions for three simulated SOLPS input locations (number 1, 7 and 24) for species number 0 (D^+), 1 (T^+) and 9 (Ne^{++}). The IEADs show a strong dependence on electron temperature, with the T_e/T_i temperature ratio affecting the spread along the energy axis. For the locations 1, 7 and 24, the electron temperature at these positions varies from 0.1 eV (pos. 1) to 0.45 eV (pos. 7) to 11.8 eV (pos. 24). Angular trends are affected by the local magnetic field inclination with respect to the surface.

As part of the integrated simulation, a reduced model of the surface is generated by performing many simulations with the F-TRIDYN ion-surface interaction code. The number of simulations is set by the energy and angle range, in addition to the fidelity of the impacting background plasma species, and input on the number of species comes from the SOLPS simulations. Of most relevance to this part of the integrated simulation are the reflection yield, the sputtering yield as well as the sputtered energy and angle distributions. The combination of

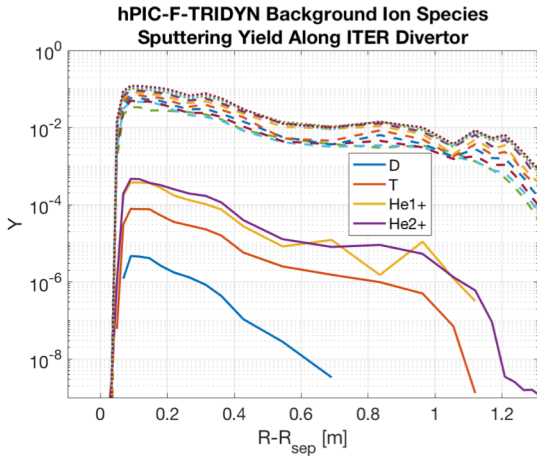


Figure 16. Sputtering yield Y_s along the ITER divertor. The Ne charge states are provided in the family of dashed lines.

the F-TRIDYN sputtering model along with the hPIC IEADs and the SOLPS surface fluxes for each ion species gives the erosion flux of the W divertor by the background plasma ions. As a first step we can compute the sputtering yield for each element and charge state. Taking $Y_s(E, \theta)$ as the F-TRIDYN surface models given before, the subscript indicating different elemental species and $f_s(E, \theta)$ as the IEADs given by hPIC for each charge state species, the sputtering yield by a given ion species is given by:

$$Y_s = \iint f_s(E, \theta) Y_s(E, \theta) dE d\theta \quad (1).$$

Upon doing this integration, Figure 16 shows that the neon species have much larger sputtering yield than the background hydrogen and helium species. The neon sputtering yields reach as high as 0.12, and the sputtering yield is generally higher for the higher charge state neon species. This trend comes about because of the increased impact energies of higher charge states accelerated through the sheath electric field.

When the relative flux magnitudes of each ion species at the target are taken into account we can produce an effective yield. This operation is done by dividing each sputtering yield Y_s by the summed integral of each species distribution function and is defined by:

$$\bar{Y}_s = \frac{\iint f_s(E, \theta) Y_s(E, \theta) dE d\theta}{\sum_s \iint f_s(E, \theta) dE d\theta} \quad (2).$$

Upon acquiring these effective yields \bar{Y}_s , Figure 17 plots which background plasma ion species are responsible for most of the sputtering of W divertor material and the sputtering trend along the target. It can be seen that T is a significant sputtering contributor while D erodes much less W material and that the contribution of He to W sputtering is nearly negligible. By normalizing to the relative background flux, the effective yields can then be added together to give a total yield for each element. As shown in Fig. 17b, the sputtering contribution along the outer ITER divertor target is dominated by Ne, but with some non-negligible contribution from T.

One notable feature in the neon effective yield is the double peak with a reduced effective sputtering yield between the peaks at $R-R_{sep} = 0.15$ m. This feature can be explained by examining the Ne energy distributions along the ITER outer divertor target as well as the relative flux fractions of Ne. By plotting the weighted average of the neon energy distributions (the sum of all Ne charge species energy distributions) in Fig. 18a, we can see that in general, the impact energies are proportional to electron temperature and fall between $3 \cdot T_e$ and $24 \cdot T_e$. Between these two bounds, we can see that the relative populations of the energy distribution shift with respect to the upper and lower bounds as well as the average energy. In particular between $R-R_{sep}$

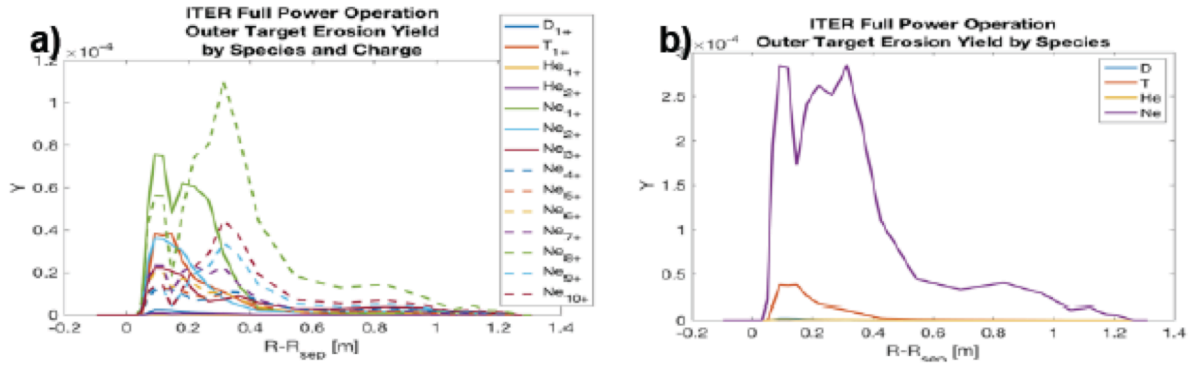


Figure 17. Effective sputtering yield, \bar{Y}_s , along the ITER divertor. a) plots the sputter yield resulting from individual ion charge states, while b) plots the total elemental sputtering yields.

of 0.3 m to 1 m there is an increase in the energy distribution that arises because of an increase in the higher charge (Ne^{8+} and Ne^3) fluxes to the target. This becomes more apparent when multiplying the energy distributions by the sputtering yield as shown in Fig. 18b, which shows the relative contribution of each energy to the sputtering of the tungsten target by neon. This highlights the contribution of the high energy IEADs of neon and their impact on erosion of the W target. By summing over these relative contributions a total effective yield for neon is acquired (white line using the right axis), and the double peaked behavior of the effective sputtering yield of the background plasma (dominated by neon) is reproduced.

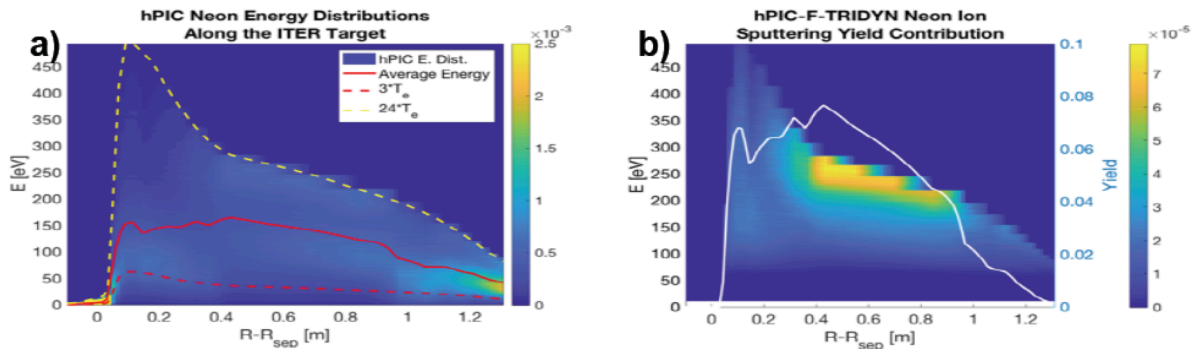


Figure 18. a) Neon ion energy distribution along the ITER divertor showing the relative concentration of the ions at each energy. b) Neon ion relative sputtering contribution as a function of energy along the ITER divertor.

The final step to reach the spatially resolved absolute erosion of the W outer divertor by the background plasma is multiplying the effective sputtering yield by the total background ion flux. Following with the previous definition of effective sputtering yield, we multiply by the total ion flux:

$$\Gamma_W = \bar{Y}_s * \Gamma_{ionTotal} \quad (3).$$

The general flux trend in the area with temperatures high enough to create erosion by background plasma species is to decrease away from the strike point. This results in a significant change in the shape from the effective sputtering yield to the eroded W flux. Figure 19 plots the W sputtering calculated for the 100 MW D-T burning plasma discharge in ITER, as a function of the elemental contribution, clearly showing that neon is the species responsible for the majority of the tungsten erosion.

Applying the background plasma eroded W profile along the divertor target plotted in Fig. 19, we acquire our initial conditions for the spatial distribution of the eroded W impurity flux in the 3D ITER geometry. In addition to the spatial distribution of eroded W impurities, the integration of the F-TRIDYN-hPIC simulation data provides the sputtered W energy and angle distributions. The sputtered W energy distribution involves an average energy of 9 eV for the sputtered W along the majority of the eroded area. The median energy, however, is only on the order of 4 eV. For these simulations, we make use of the conservative assumption in F-TRIDYN of zero bulk binding energy for the target and this results in a conservative (over-prediction) of the erosion yields, as well as the sputtered energies. The additional initial conditions include the sputtered angle distributions of W by the background plasma. The average sputtered W azimuthal angle is between 50 and 60 degrees where 0 degrees is normal incidence, with the peak of the distribution between 60 and 70 degrees. This indicates that most W atoms sputtered by the background plasma are sputtered more along the direction of the surface (greater than 45 degrees) than in the direction perpendicular to the surface. With the additional information from the poloidal angle of the sputtered W averaging 90 degrees, this shows that the W atoms are sputtered in the “forward” direction with respect to the magnetic field angle and the impact of the background plasma ions.

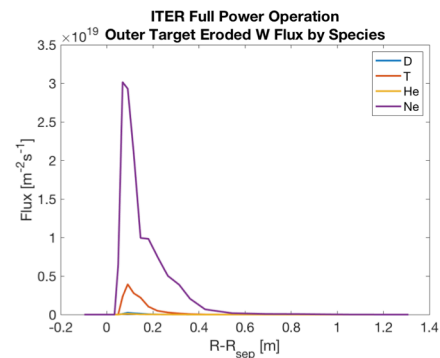


Figure 19. Total W sputtered flux from D, T, He and Ne.

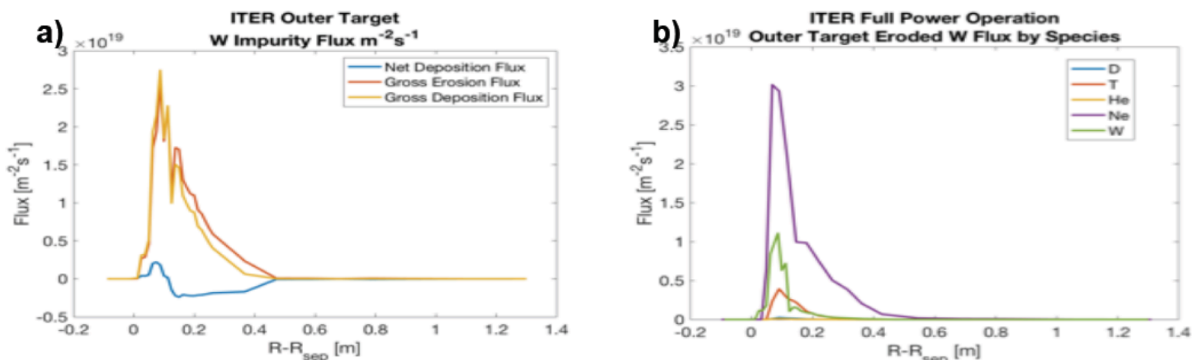


Figure 20. a) Gross erosion, gross deposition and net deposition flux to the ITER target as simulated by GTR. Positive values indicate net deposition of W, while negative values indicate net erosion of the W divertor. b) Gross erosion flux contributions by element along the ITER divertor target.

The results of the W impurity transport simulation performed with GTR predict a gross W erosion profile along the target, shown by the orange line in Figure 20. Much of the eroded W

impurity flux is promptly re-deposited, as shown by the yellow line in Fig. 20a. However, a fraction of the eroded impurities are transported in the direction of the strike point creating a net deposition region of W ranging from $R-R_{\text{sep}} = 0$ to $R-R_{\text{sep}} = 0.1$ m. In total for the impurity transport simulation, 93% of the eroded W is deposited in the divertor region with approximately 7% of the eroded W being trapped in the scrape-off-layer or divertor plasma (including the private flux and open magnetic field regions). Taking the gross erosion flux contribution from W self-erosion in GITR and comparing it to the magnitudes of the erosion due to the background plasma ions, Fig. 20b shows that Ne is still the dominant eroding species, but W becomes a significant contributor, especially at $R-R_{\text{sep}} < 0.05$ m.

Following the integrated workflow, we next perform coupled F-TRIDYN and Xolotl simulations (FTX) to model gas (D, T, He) implantation, retention and recycling, as well as variations in surface height during the 1st second of plasma exposure of the ITER outer divertor target during a burning plasma discharge. We performed an independent, coupled FTX simulation for 40 locations across the W target. To highly resolve the region with highest variations in ion temperature and particle and heat fluxes, 30 of these points are located near the strike point (within $R-R_{\text{sep}} = -0.1 - 0.25$ m). 10 more locations are modeled further along the target ($R-R_{\text{sep}} = 0.25 - 1.3$ m). Each FTX task is performed on Edison at NERSC on 6 nodes and using 144 processes.

In the coupled FTX workflow, F-TRIDYN calculates the implantation profile and sputtering rate for each species in the plasma taking as input IEADs provided by hPIC (for D, T, He, Ne) and GITR (for W). Xolotl then uses these implantation profiles (except that of Ne, which is not tracked in the code), as well as the effective sputtering yield as input. With this information, Xolotl evolves of the sub-surface gas (D, T, He), W interstitial and vacancy concentrations. The code also models changes in surface height driven by sputtering and by gas dynamics (nucleation, bubble bursting, etc.), as described below. Given that the implantation profiles and sputtering rates can vary with surface composition, the workflow updates the profiles and yields regularly by making a new call to F-TRIDYN (also referred to as ‘a loop’ in what follows). Xolotl solves the spatially dependent DDR equations in one spatial dimension (depth) to predict the evolution of the concentration fields for the plasma exposed tungsten divertor. We use a finite difference implementation in which the grid along the depth direction has a variable grid size, with a finer grid near the surface, in order to better describe the profile of the incoming ion fluxes. In the interest of brevity, the model governing the master equations and solver will not be discussed here, but are presented in detail in Ref. [19].

The clusters under consideration in this work are tungsten self-interstitials (I), vacancies (V), helium (He), deuterium (D), tritium (T), as well as vacancy-trapped impurities (He-D-T-V), which we also refer to as a bubble since the incorporation of one or more vacancies makes the cluster immobile. Single vacancies, interstitial clusters, helium clusters, single deuterium, and single tritium are mobile and diffuse isotropically in bulk tungsten. The reaction network includes maximum cluster sizes of 8 He interstitials, 50 W vacancies and 6 W interstitials. The surface is able to move upward in opposition of the erosion flux as a result of tracking the number of tungsten self-interstitial atoms diffusing to the surface and adding a grid point when this number exceeds the tungsten atomic density. As described in Ref. [25-27] our model includes a biased drift term and modified trap mutation near the surface for mobile helium clusters. Bursting of bubbles reaching the surface is also included, freeing all the impurities (He, D, T) and leaving an empty cavity (V), which can subsequently refill with gas. Isotope exchange

is not yet included in Xolotl, and will be a focus of future research activities, along with extending the absolute time that the simulations reach.

At each point along the divertor, a 6 mm deep substrate is simulated in Xolotl, that is, the thickness of the W divertor tiles up to the copper cooling pipe [28]. This is modeled using 367 grid points, which include to a 3.6 nm void at the top to allow for surface growth (as described in the prior paragraph). The temperature is initially uniform across the substrate, at 343K. This value is fixed at the deepest grid-point (6 mm deep) throughout the simulation, and all other points evolve based on solving the thermal diffusion equation based on the local value of the incoming heat flux.

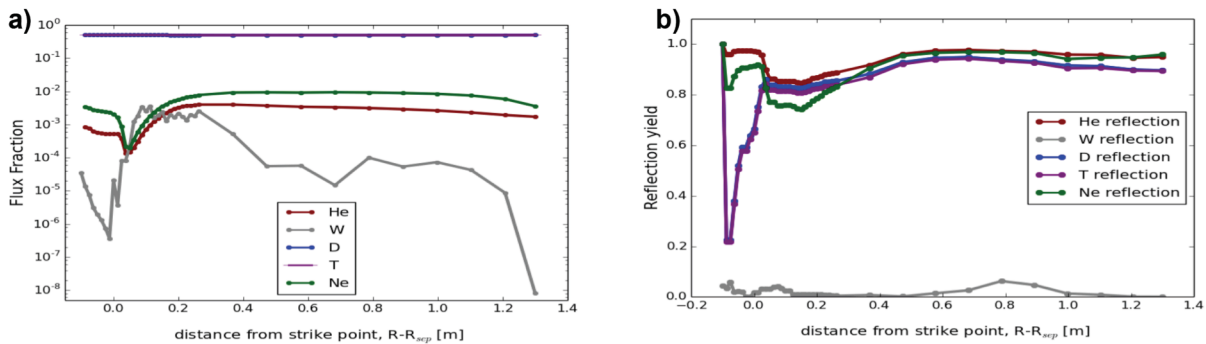


Figure 21. a) Fraction of each species in the plasma, with respect to the surface, which is input to FTX from hPIC (for D, T, He and Ne) and GTR (for W). b) Reflection (or prompt recycling) integrated over the IEADs and averaged over all F-TRIDYN loops.

Figure 21 shows the input data to the FTX sub-workflow. Within Fig. 21a, it is possible to define three ‘populations’ that exist in the fraction of each species (relative to the total flux) that reach the wall. These are the main plasma components (D, T), the extrinsic impurities (i.e., particles generated or seeded outside the simulated plasma-wall volume, namely He and Ne) and the eroded and re-deposited W. The main species (D,T) dominate the flux composition and their fraction is nearly constant across the target. In contrast, the fraction of extrinsic impurities decreases near the strike point, due to the very high recycling of hydrogenic species [29]. Finally, the W concentration (relative to the total flux) varies strongly across the target, due to the complex interplay between erosion and re-deposition. Note that reflection is automatically set to 1.0 for a species in areas with no flux of it, so that the implanted fraction (1-reflection) is zero. This is the case for background plasma species at around $R-R_{sep} = -0.1$ m. Overall, the hydrogenic species recycling coefficient is effectively 1.0 across the divertor, although there are slight variations that lead to increasing implanted hydrogenic species concentrations, as discussed below.

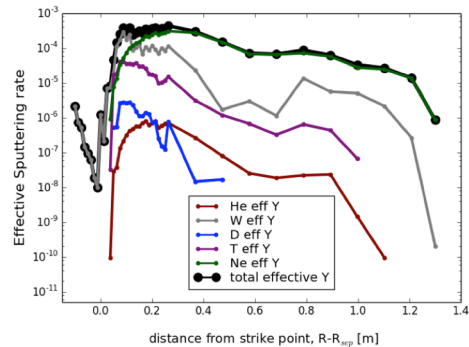


Figure 22. Effective sputtering rates for each species, and the total sputtering rate, as a function of distance to the strike point.

The effective sputtering rates (calculated from the sputtering yield of a species times its concentration) are plotted in Figure 22 and show the

relative contributions of each species to the total sputtering in that location. Near the strike point and in the private flux region ($R-R_{\text{sep}} < 0.1$ m), the ion temperature, which most strongly correlates with the impact energy of background plasma species, is low. Thus, in this region, He, D, T, as well as Ne (the main sputtering contributor for locations $R-R_{\text{sep}} > 0.1$ m) impact the surface at low energies and their contribution to erosion is negligible. This leads to quite low sputtering values in this region, which result entirely from W self-sputtering. For locations further up the W target ($R-R_{\text{sep}} > 0.1$ m), Ne dominates the total effective sputtering yield, due to its high sputtering yield at mid-to-high ion temperatures ($T_i > \sim 10$ eV).

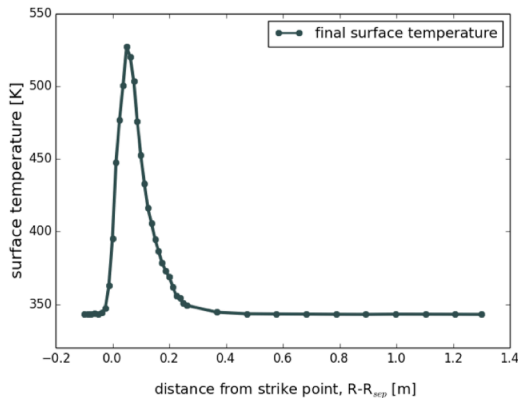


Figure 23. Surface temperature along the ITER divertor at a time of approximately 1 second of 100 MW, D-T burning plasma operation.

Figure 23 shows the surface temperatures at a time of approximately one second after the initiation of the 100 MW burning D-T plasma discharge, and reveal an increase on surface temperature of up to nearly 200 K. Although this temperature increase will not pose a threat to the mechanical properties of the W, nor melting, these calculations are for steady state plasma operation and do not incorporate edge localized modes, nor plasma disruptions. However, these temperature increases are sufficient to impact the sub-surface gas dynamics, as discussed below for D and He depth profiles at the location of peak heat flux.

The gas content as a function of position along the ITER divertor is shown in Figure 24. Both the H isotopes (D, T) and the He concentrations peak close to the strike point (Fig. 24a). Fig. 24b shows the growth rate of the sub-surface gas concentration (defined as concentration increase per unit time), and shows that the rate of D accumulation in the tungsten is nearly constant in the region with significant ion flux ($R-R_{\text{sep}} \leq 0.5$ m). Further along the target, the D concentration decreases, consistent with the decrease in implantation flux and ion temperature. The T content is slightly but consistently higher than that of D, possibly due to a minimal difference in reflection yields that are slightly lower for T than D, thus lead to a slightly larger implanted T flux. Slower T diffusion compared to D (due to the mass difference) may also contribute to the difference in the sub-surface hydrogenic species retention and spatial concentration profiles. Although D is expected to implant deeper than T due to the mass difference, the difference in our calculated implantation profiles for D versus T is negligible. Across most of the divertor target ($R-R_{\text{sep}} > 0.2$ m), the He content shows a similar trend as the hydrogenic species. However, the He fraction in the plasma dips slightly around the strike point leading to a lower He concentration, possibly due to the low flux in the private flux region, which produces reduced He implantation fluxes and thus, a lower He content. In addition, the flux is low in the private flux region, thus less He is implanted, leading to a lower content. The latter may also be responsible for the significant fluctuations in D and T content at locations around $R-R_{\text{sep}} = -0.1 - 0.05$ m.

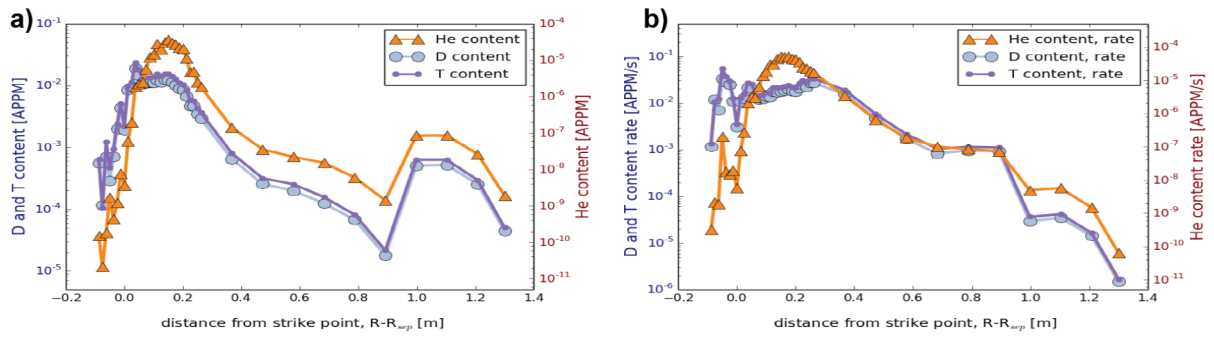


Figure 24. a) D, T and He concentration as a function of position along the ITER divertor, where the concentration values have been integrated in depth, b) the growth rate (concentration/second) of the D, T and He gas concentrations.

These gas concentrations result from the implantation, and build up, due to the gases that either cluster in the near surface region or begin to diffuse deeper below the surface of the divertor tile. In locations of low ion temperature ($R-R_{sep} \leq 0.1$ m), the calculated ion implantation profiles are very shallow and, in fact peak at or close to the surface (<1.5 nm), as shown in Figure 25. Indeed, at the location of the peak particle flux location, shown in Fig. 25a, the highest implanted D and T fluxes are at the surface, while the He peaks is approximately 1 nm below the surface. Regions in the divertor exposed to higher ion temperatures experience deeper ion implantation, which is more representative of the profiles implanted in linear plasma devices of normal incident, low-energy ion implantation with implantation peaks located a few nm deep in the substrate. This is the case for the location of the peak heat flux, as shown in Fig. 25b (D, T implantation profiles) and Fig. 25c (He and Ne implantation profiles), where the implantation of lighter ions peaks at ~ 2.5 nm, whereas the heavier Ne ions are implanted at a depth of ~ 1 nm. It is important to note that while we have included neon within our workflow, our sub-surface gas dynamics model in Xolotl does not currently include the effect of neon clustering or diffusion, beyond the impact of neon on near surface erosion.




Figure 25. Gas implantation profiles at a) the location of peak particle flux showing D, T and He, and b) and c) the location of peak heat flux, with D and T implantation profiles shown in b) and He and Ne implantation profiles shown in c).

After implantation, the spatial distribution of He is predicted by Xolotl to initially increase in concentration near the surface with limited diffusion into the bulk, within the first second. This trend, plotted in Figure 26 as a function of time shows that the He concentrations are approximately two orders of magnitude lower than the hydrogenic species, consistent with the relative concentration of helium in the boundary plasma. Note that the D and T concentration

profiles are very similar, with the T concentration slightly larger than that of D, and correspondingly, only the T concentration profile is shown in Fig. 26.

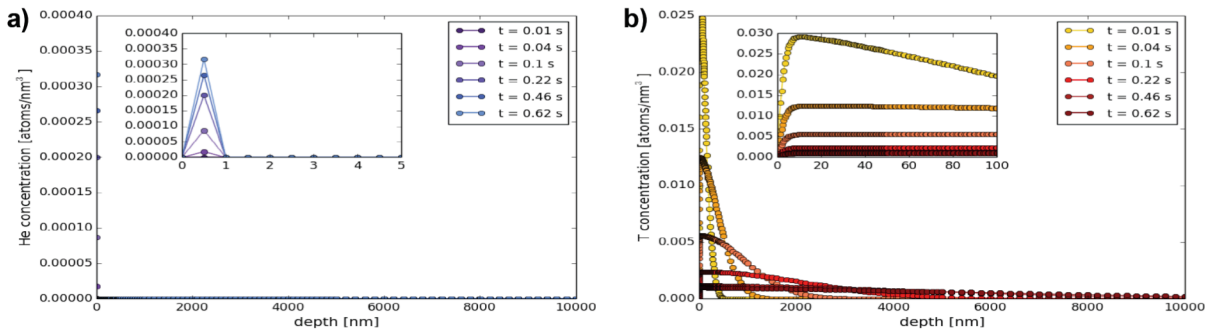


Figure 26. Depth profiles of the cluster-size integrated gas concentrations (in units of gas atoms nm^{-3}) as a function of time at the spatial location of peak heat flux for a) helium and b) tritium. The insets highlight the concentration profiles very near the surface in the first 5 or 100 nm, respectively for He and T, while the larger plots show the profiles to a depth of 10 microns.

The He concentration (integrated over all cluster sizes) does increase with increasing time, as shown in Fig. 26a, and indicates significant He cluster nucleation occurs near the surface, likely due to the modified helium trap mutation processes with increased reaction rate near tungsten surfaces. In contrast, some locations (i.e., those of peak particle flux, peak heat flux and peak ion temperature) show an interesting evolution of the D and T concentrations, as shown in Fig. 26b (Note that the D and T concentration profiles are very similar, so only the T concentration profile is shown in Fig. 26, which has been integrated over all cluster sizes at each depth). While initially unexpected, the time-dependence of the T peak concentration is explained by accounting for the significant change in surface temperature (~ 150 K), which rapidly increases due to the local surface heat flux and reaches an approximately steady state temperature distribution after about the first 0.5 seconds.

Thus, the initial D and T concentration profiles are the values expected from an initial temperature of 343K, with slower hydrogenic species diffusion. As the surface temperature increases, such as in the location of peak heat flux shown in Fig. 26b, the T (& D) diffusivity increases. Effectively, the peak implantation decreases as a result of diffusion to the surface leading to desorption, as well as deeper below the surface of the tungsten divertor tile. The peak of the cluster size integrated T concentration profile decreases to a value consistent with what is expected at the final temperature. The final two time-frames ($t=0.46$ and $t=0.62$ s, respectively) nearly overlap, reflecting that the temperature distribution reaches a steady-state at around 0.5 s. Separate Xolotl simulations of gas implantation with constant temperature values between 343 and 525 K have been used to confirm this explanation of the time dependence of the hydrogenic species concentrations, but are not included here in the interest of brevity.

At locations further from the strike point, the heat flux is much lower and the surface temperature does not increase significantly, as was shown in Fig. 23. Figure 27 shows the cluster-size integrated tritium concentration profile as a function of time at a location of approximately $R - R_{\text{sep}} = 1.0$ m, for which the surface temperature increase is minimal, and the simulation has accumulated a time slightly more than 15 seconds. At this location, the T concentration profile quickly builds up to a value of about 0.00011 nm^{-3} , and as time increases,




Figure 27. Depth profile of the cluster-size integrated tritium concentration as a function of time at $R - R_{sep} = 1.0$ m, where there is low particle and heat flux.

and therefore the ions are implanted on or near the top-most surface and easily escape the substrate. The shallow implantation range and the high gas implantation fluxes contribute to a rapid saturation of the surface hydrogenic species concentration, leading to essentially zero T (and also D) retention, such that the recycling coefficient of D and T at the peak particle flux is nearly 1.0. Similar behavior is seen for the He in Fig. 28a. Although the peak particle flux and peak heat flux are located only 2.5 cm apart, the ion temperature increases in this region from $T_i \sim 2.5$ eV to 10 eV. Thus ions are implanted deeper beneath the surface, and T retention (highly dominated by diffusion in cases with low impurity content, as at present) is higher. The same reasoning follows for the higher tritium retention observed in the case of the peak ion temperature and the 2nd He peak, as ion temperatures are even higher for these two locations ($T_i \sim 42$ eV and 39 eV, respectively).

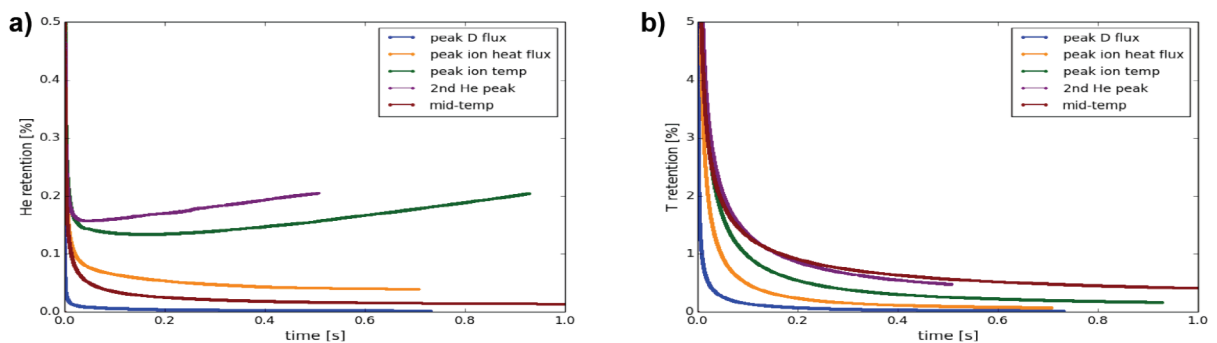


Figure 28. The time evolution of the retained gas within the tungsten divertor for a) helium and b) tritium. Retention as defined as depth integrated concentration divided by the implanted fluence. The retention values vary with spatial position from the peak particle flux, peak heat flux, peak ion temperature, secondary helium implantation flux and the mid-temperature location.

In contrast, He retention shows a more complex pattern of interplay between ion implantation depths, He diffusion and nucleation. The two locations with highest He implantation rate (peak ion temperature and 2nd He peak, with $\Gamma_{impl, He} \sim 5.5$ and 8 atoms/nm² s, respectively) show an

the characteristic deeper diffusion and permeation of the T is observed. The diffusional behavior shown in Fig. 27 is much more representative of implantation, diffusion and permeation at a constant temperature, as in our prior modeling of the PISCES-A linear device experiments. At this location there is also a very low implanted helium flux, so the helium concentration has not been plotted.

Figure 28 shows a comparison of the helium (Fig. 28a) and tritium (Fig. 28b) retention (relative to the implanted fluence) across the 5 spatial locations that we have analyzed in detail. Retention of both He and T is low (and in fact lowest) at the location of peak flux, as the impact energy is lowest

increasing trend in He retention, expected when trap mutation takes over, as more and more He can be absorbed by bubbles. The three locations with lowest He implanted fluxes (peak particle flux, peak heat flux and the mid-temperature area, with $\Gamma_{\text{impl, He}} \sim 4, 3$ and 0.05 atoms/nm² s, respectively) show a retention trend similar to that of T and seems dominated by diffusion. Note that although the He implantation rate for the peak D flux is close to that at the peak for ion temperature, He is implanted very near the top-most surface in the former case, thus desorbs quickly and nucleation of He clusters is less likely. The very low particle flux in the mid-temperature location may require longer time scales to reach the equilibrium retention values, as shown by the still significant decay of T and He retention at $t=1.0$ s.

Finally, we have performed additional simulations that have an initial helium concentration below the tungsten divertor surface predicted to result from the He-only plasma discharge of ITER operation that was simulated in our third quarter activity (as shown previously in Fig. 10b). These simulations were performed to meet our milestone objective of evaluating the influence of material deposition and evolving near surface concentration on the recycling and retention of the hydrogenic fuel species. The substrate composition for the location at $R-R_{\text{sep}} \sim 0.05$ m, near the peak heat flux, was set to the exact He-cluster distribution present in the W substrate at the end 10 s He-plasma exposure reported in the 3rd quarter. This pre-implanted divertor tile was then exposed to the same conditions as the initially pure tungsten divertor tile at $R-R_{\text{sep}} \sim 0.05$ m, for which the results were previously shown in Fig. 26. This Xolotl simulation has reached a somewhat shorter accumulated time of 0.51 seconds, compared to the 0.8 s accumulated for the pure W divertor tile that did not contain any pre-existing impurities. However, comparing Fig. 29a to Fig. 29a clearly indicates that the He concentration is significantly higher (by about 1000X) in the substrate with pre-implanted He (than in the simulations with an initially clean W substrate). Again, this higher helium concentration was the initial state of the divertor tile resulting from the exposure to a large He fluence during the 10 second He-only plasma discharge, compared with the small He fraction in the D-T-He exposure. This statement applies to both to near-surface concentrations, where distinctive peaks due to He clustering are visible, as well as deeper locations that have increased helium concentration from helium diffusion well below the surface. During exposure to the burning D-T-He plasma conditions, the near-surface He concentration decreases as the pre-existing high pressure helium bubbles are able to burst, leaving empty vacancy cavities in the divertor substrate. These vacancy clusters are quickly filled with the implanted species, in this case, mainly D and T.

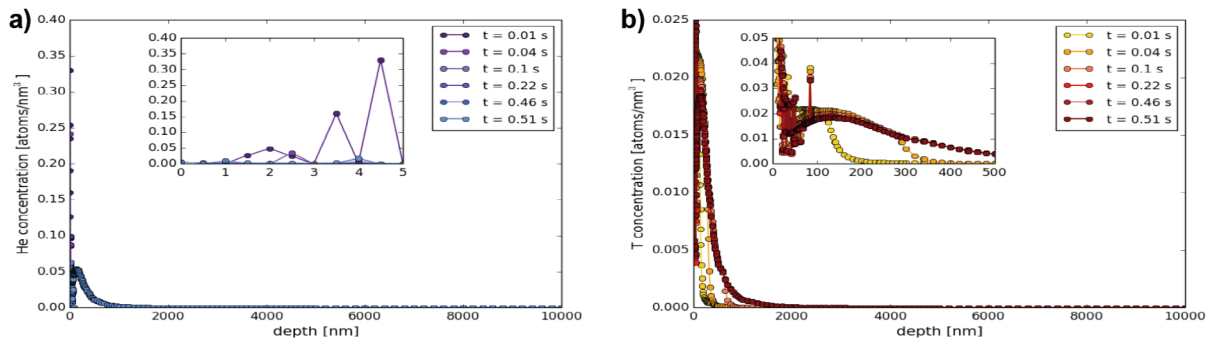


Figure 29. Depth profiles of the cluster-size integrated gas concentrations (in units of gas atoms nm⁻³) as a function of time at the spatial location of peak heat flux for a) helium and b) tritium, for an initial substrate composition that includes pre-implanted He. The insets highlight the concentration profiles very near the surface in the first 5 or 100 nm, respectively for He and T, while the larger plots show the profiles to a depth of 10 microns.

Correspondingly, comparing Fig. 29b to Fig. 26b, indicates a much higher T concentration (by nearly 10X) develops as a result of the pre-existing sub-surface gas cluster microstructure. Likewise, the near-surface concentration peak does not decrease with increasing time, since these hydrogenic species are more strongly trapped at helium-vacancy and vacancy clusters with de-trapping energies that range from about 1.1 to in excess of 1.5 eV, as predicted by our first principles density functional theory calculations [30] reported in the quarter 1 milestone report. Correspondingly, this near-surface concentration peak is expected to remain nearly constant without decreasing caused by the increasing near surface temperature, although we can not preclude further increases in near surface concentration at later time with continued D, T and He implantation. Another important observation can be noted by comparing Fig. 28b to Fig. 26b, namely that this near surface helium cluster microstructure is predicted to reduce the deeper permeation of T (and also D, although not shown here), at least for the approximately 1 second interval reported here. This indicates that the pre-implantation of helium into the tungsten divertor tiles could produce complex hydrogenic species retention behavior during the burning plasma operation phase of ITER, in which the very near surface regions exhibit increased tritium retention at potentially strongly bound trapping sites while the net permeation through the divertor is decreased. Further research will be required to confirm this hypothesis.

Computational Performance

In order to obtain the results documented in this report, we have used both coupled-code and standalone simulations. During the fourth quarter, we improved our ability to monitor the performance of the coupled-code simulations, and improved the performance of two of the standalone simulation codes. As in previous quarters, our Q4 coupled-code simulations integrated the F-TRIDYN ion-solid interactions simulation [15], the GITER Tungsten impurity transport simulation [16], and the Xolotl cluster dynamics simulation [17-19] using the IPS [9,31] framework. During Q4, we augmented our workflow-level performance monitoring approach to capture more detail regarding when individual workflow activities occurred and on which compute nodes they ran. Like our Q2 and Q3 approach, we still post-process the IPS event log into an Open Trace Format [32] event trace file that can be visualized using the Vampir performance tool [33]. We now also post-process some of our workflow scripts' output files to collect information about which compute nodes were allocated to each F-TRIDYN and Xolotl instance, and correlate that information with the events from the IPS event log. This additional information allowed us to correct a mistaken assumption regarding how IPS allocates compute nodes to each task, and will form the basis for a future investigation of the performance impact of being allocated compute nodes that are not "close" in terms of their location within the physical system running our simulations. Figure 30 shows an example visualization of a restart job involving F-TRIDYN and Xolotl from our Q4 ITER simulation with He-D-T burning plasma conditions. The visualization reflects Xolotl's increased computational demands for this simulation compared to Q2 and Q3 simulations, and shows our simulations exhibited good resource utilization during the job.

Toward the end of the reporting period, we developed an alternative performance data collection library that can provide both workflow- and application-level performance information in integrated event trace files. The library provides C, C++, Fortran, and Python interfaces and so can be used to instrument all of the components of our coupled-code workflow. Due to the high-priority need to focus on Xolotl performance improvements to support our Q4 ITER He-D-

T plasma simulations, we have not yet incorporated this performance data collection library into our workflow but plan to do so in FY2019.

As revealed by our workflow-level performance visualizations, Xolotl performance is the primary determinant of our overall workflow’s performance. As our simulations became more complex during FY2018, we identified and addressed several performance and memory usage problems with the Xolotl version used in our workflow. During Q4, our Xolotl performance engineering activities mainly involved improvements to Xolotl’s I/O operations.

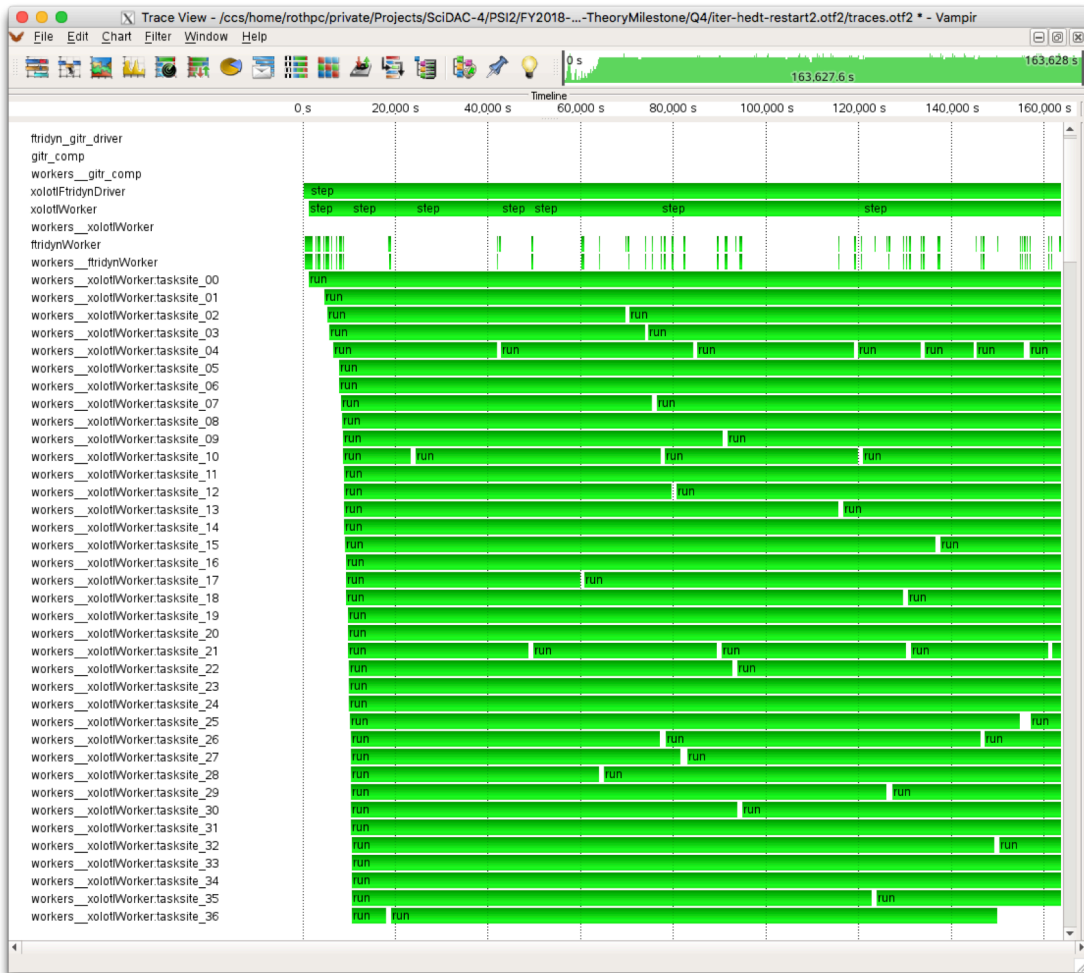


Figure 30. Timeline visualization from a restart run of the ITER D-T (+ He) burning plasma simulation showing the performance of multiple F-TRIDYN and Xolotl instances running concurrently on the available resources.

As we increased the scope and complexity of our simulation for Q4, Xolotl’s built-in performance data collection infrastructure indicated poor performance and high performance variability of two Xolotl functions called by PETSc on every time step. Subsequent fine-grained performance diagnosis using the HPCToolkit tool [34] pointed to problems with the way that Xolotl was writing its state to its checkpoint file and to the files to be used as input by subsequent F-TRIDYN instances. By modifying Xolotl to use parallel I/O and to avoid redundant I/O, we improved the checkpoint performance by 57.5X, as shown in Fig. 31a on four nodes of the Oak Ridge Leadership Computing Facility’s Eos Cray XC30. Because of the

similarity between the Eos system and the NERSC Edison Cray XC30 used for our Q4 simulations, this modification provided a positive performance impact to every Xolotl instance run during our final simulation. A similar modification greatly reduced Xolotl's performance variability when writing files to be used by subsequent F-TRIDYN instances. Using one of the Xolotl runs from our early Q4 simulations as a benchmark, the modification reduce the total time spent writing F-TRIDYN files from over 700s to approximately 30 s.

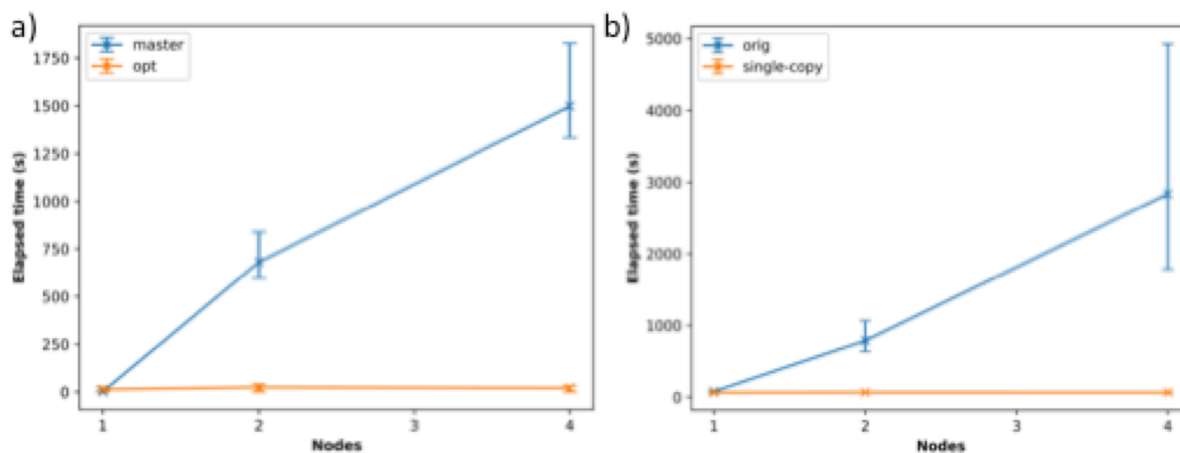


Figure 31. a) Time required to write Xolotl checkpoint file, as reported by built-in timers and running on the OLCF Eos Cray XC30 with 32 processes per node. Data shown is the max of reported time values across all processes, averaged over at least three simulations. b) Time required to copy Xolotl reaction network data from input HDF5 file to checkpoint file from simulations on the Eos Cray XC30 with 16 processes per node.

During Q4, we also experienced Xolotl runs that exhibited poor performance during initialization. We determined that approach used in Xolotl for copying its reaction network data from its input HDF5 file to its checkpoint file (needed so that it is available for a restart), was using an HDF5 function that is inherently sequential but Xolotl was assuming it was a collective operation that needed to be done by all Xolotl processes. We modified Xolotl so that only one of its processes copies the reaction network data, thus eliminating redundant I/O and contention to the file system. The modified version did Xolotl initialization (including the reaction network copy I/O) approximately 43X faster than the original version on four nodes of the OLCF Eos system, as shown in Fig. 31b.

In addition to the improvements highlighted here, our performance engineering work with Xolotl for the PMI milestone experiments uncovered several directions for future Xolotl performance improvements that will require more substantial code modifications than we could accomplish during FY2018, such as major data structure reorganization to better support multithreading and GPU acceleration of the most computationally expensive functions in Xolotl.

GITR replicates its background mesh and static fields on each node. It also initializes all particles on each node. This initialization uses data read from input files, and during Q4 we determined that I/O contention caused poor performance for GITR initialization. We addressed

this problem by changing GITER so that one process reads the input data and then broadcasts to all other processes, thus avoiding the contention for the file system.

We also improved GITER performance by improving its load balance. We realized that the original approach within GITER of simulating the same number of time steps for each particle was inefficient, because particles that were experiencing fewer simulated events can be simulated using fewer time steps than those experiencing many events. To leverage this, we modified the code to detect particles with larger computational demands and then adjust how many particles each GITER thread “owns” and how many time steps each must simulate so that the overall computation time of each thread is roughly the same. We do this in a way that retains GITER’s good performance running on graphics processing units (GPUs), by avoiding branching that causes performance-robbing divergence across GPU threads. This led to an approximately 40% decrease in GITER run time.

REFERENCES

1. D.E. Post and R. Behrisch, editors, *Physics of plasma-wall interactions in controlled fusion*, Basic Books, 1986.
2. G. Federici, C.H. Skinner, J.N. Brooks, J.P. Coad, C. Grisolia, A.A. Haasz, A. Hassanein, V. Philipps, C.S. Pitcher, J. Roth, W.R. Wampler, and D.G. Whyte, *Nuclear Fusion* **41** (2001) 1967.
3. J. Roth, E. Tsitrone, T. Loarer, V. Philipps, S. Brezinsek, A. Loarte, G. Counsell, R. Doerner, K. Schmid, and O. Ogorodnikova, *Plasma Physics and Controlled Fusion* **50** (2008) 103001.
4. M. Greenwald, R. Callis, D.W. Dorland, D. Gates, J. Harris, R. Linford, M. Mael, K. McCarthy, D. Meade, F. Najmabadi, W. Nevins, J. Sarff, M. Ulrickson, M. Zarnstorff, and S. Zinkle, Technical Report DOE/SC-102, U.S. Department of Energy Fusion Energy Sciences Advisory Committee, 2007.
5. R. Maingi, S. Zinkle, M. Foster, H. Guo, B. LaBombard, J.P. Allain, R. Doerner, C. Kessel, D. Youchison, A. Hubbard, and A. Leonard, https://science.energy.gov/~media/fes/pdf/workshop-reports/2016/PMI_fullreport_21Aug2015.pdf, 2015.
6. B. D. Wirth, K. D. Hammond, S. I. Krasheninnikov, and D. Maroudas, *Journal of Nuclear Materials* **463** (2015) 30.
7. B.D. Wirth, K. Nordlund, D.G. Whyte, and D. Xu, *MRS Bulletin* **36** (2011) 216.
8. P. Bonoli, L. McInnes, C. Sovinec, D. Brennan, T. Rognien, P. Snyder, J. Candy, C. Kessel, J. Hittinger, L. Chacón, D. Estep, T. Munson, W. Bethel, M. Greenwald, D. Bernholdt, and B. Lucas, https://science.energy.gov/~media/fes/pdf/workshop-reports/2016/ISFusionWorkshopReport_11-12-2015.pdf, 2015.
9. W.R. Elwasif, D.E. Bernholdt, A.G. Shet, S.S. Foley, R. Bramley, D.B. Batchelor, and L.A. Berry, 18th Euromicro Conference on Parallel, Distributed and Network-based Processing (PDP 2010), Pisa, Italy, Feb. 2010.
10. R. Schneider, X. Bonnin, K. Borrass, D.P. Coster, H. Kastelewicz, D. Reiter, V. A. Rozhansky, and B. J. Braams, *Contributions to Plasma Physics, Plasma Edge Physics with B2-Eirene* **46** (2006) 3.
11. B. Braams, Radiative Divertor Modelling for ITER and TPX, *Contributions to Plasma Physics* **36** (1996) 276.

12. D. Reiter, M. Baelmans, and P. Börner, *Fusion Science and Technology*. **47** (2005) 172.
13. R. Khaziev and D. Curreli, *Physics of Plasmas* **22** (2015) 043503 (2015).
14. D. Curreli and R. Khaziev, *Towards large scale kinetic simulations of the plasma-material interface*. Blue Waters annual report book, NCSA, 2016.
<https://bluewaters.ncsa.illinois.edu/annual-report>
15. J. Drobny, A. Hayes, D. Curreli, D.N. Ruzic, *J. Nucl. Mat.* **494** (2017) 278-283.
16. GITR Github repository, <https://github.com/ORNL-Fusion/GITR>
17. S. Blondel, D.E. Bernholdt, K.D. Hammond, L. Hu, D. Maroudas, and B.D. Wirth, *Fusion Sci. Tech.* **71** (2017) 84.
18. Xolotl Github repository, <https://github.com/ORNL-Fusion/Xolotl>
19. D. Maroudas, S. Blondel, L. Hu, K.D. Hammond, and B.D. Wirth, *Journal of Physics: Condensed Matter* **28** (2016) 064004
20. R.P. Doerner, D. Nishijima and T. Schwarz-Selinger, *Nuclear Fusion* **52** (2012) 103003.
21. <http://www.nersec.gov/users/computational-systems/edison/>
22. G. Shaw, M. Bannister, T.M. Biewer, M.Z. Martin, F. Meyer, and B.D. Wirth, *Applied Surface Science* **477** (2018) 695-703
23. A.S. Kukushin, H.D. Pacher, G. Janeschitz, A. Loarte, D.P. Coster, G. Matthews, D. Reiter, R. Schneider, and V. Zhogolev, *Nuclear Fusion* **42** (2002) 187.
24. R. Khaziev, D. Curreli, *Physics of Plasmas* **22** (2015) 043503.
25. L. Hu, K. D. Hammond, B. D. Wirth and D. Maroudas, *J. Appl. Phys.* **115** (2014) 173512.
26. L. Hu, K. D. Hammond, B. D. Wirth and D. Maroudas, *Surf. Sci.* **626** (2014) 21.
27. L. Hu, K. D. Hammond, B. D. Wirth and D. Maroudas, *J. Appl. Phys.* **118** (2015) 163301.
28. RA Pitts, S Bardin, B Bazylev, MA van den Berg, P Bunting, S Carpentier-Chouchana, JW Coenen, Y Corre, R Dejarnac, F Escourbiac, J Gaspar, JP Gunn, T Hirai, SH Hong, J Horacek, D Iglesias, M Komm, K Krieger, C Lasnier, GF Matthews, TW Morgan, S Panayotis, S Pestchanyi, A Podolnik, RE Nygren, DL Rudakov, G De Temmerman, P Vondracek, JG Watkins, *Nuclear Materials & Energy* **12** (2017) 60.
29. S. Krasheninnikov, A.S. Kukushkin and A.A. Pshenov, *Phys. of Plasmas* **23** (2016) 055602.
30. L. Yang and B.D. Wirth, *Journal of Applied Physics* **123** (2018) 215104.
31. S.S. Foley, W.R. Elwasif, and D.E. Bernholdt, "The Integrated Plasma Physics Simulator: A Flexible Python Framework for Coupled Multiphysics Simulation," ORNL Technical Report ORNL/TM-2012/57.
32. D. Eschweiler, M. Wagner, M. Geimer, A. Knüpfer, W.E. Nagel, F. Wolf, *Advances in Parallel Computing* **22** (2013) 481.
33. A. Knüpfer, H. Brunst, J. Doleschal, M.J. Jurenz, M. Lieber, H. Mickler, M. S. Müller, W.E. Nagel, "The Vampir Performance Analysis Tool-Set." In: Resch M., Keller R., Himmler V., Krammer B., Schulz A. (eds) *Tools for High Performance Computing*. Springer, Berlin, Heidelberg, (2008) pp. 139-155.
34. L. Adianto, S. Banerjee, M. Fagan, M. Krentel, G. Marin, J. Mellor-Crummey, and N.R. Tallent, *Concurrency and Computation: Practice and Experience* **22** (2010) 685.
35. S. Wiesen, D. Reiter, V. Kotov, M. Baelmans, W. Dekeyser, A.S. Kukushkin, S.W. Lisgo, R.A. Pitts, V. Rozhansky, G. Saibene, I. Veselova, and S. Voskoboinikov, *Journal of Nuclear Materials* **463** (2015) 480-484.

36. A. Kirschner, D. Tskhakaya, G. Kawamura, D. Borodin, S. Brezinsek, R. Ding, Ch. Linsmeier, and J. Romazanov, *Contributions to Plasma Physics* **56** (2016) 622–627
37. P.C. Stangeby and J.D. Elder. “Calculation of observable quantities using a divertor impurity interpretive code, DIVIMP”, *J. Nucl. Mater.* **196-198** (1992), 258–263.
38. J. Hogan, C.C. Klepper and J. Harris (1997). “Integrated impurity model for actively cooled plasma facing components”. International Atomic Energy Agency (IAEA): IAEA.
39. J. N. Brooks, *Physics of Fluids B: Plasma Physics* **2** (1990) 1858.
40. J. Nickolls, I. Buck, M. Garland and K. Skadron. Scalable Parallel Programming with CUDA. *ACM Queue*, vol. 6 no. 2, March/April 2008, pp. 40-53
41. J. Hoberock and N. Bell. "Thrust: A parallel template library." (2010).
42. S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, K. Rupp, B. F. Smith, S. Zampini, and Hong Zhang, PETSc web page, <http://www.mcs.anl.gov/petsc/>, 2016.

Appendix A:

Description of the Codes used in our Integrated ITER burning plasma simulations

Our ITER workflow couples three distinct applications - F-TRIDYN, GTR, and Xolotl - using the IPS framework, and receives one-way input from SOLPS for the background plasma conditions and hPIC for the ion energy angle distributions. In this appendix, we provide a detailed descriptions of these codes, the integration between simulation codes performed within the IPS framework, and our re-gridding strategy for translating SOLPS results to the data format required for hPIC and GTR.

Integrated Plasma Simulator (IPS) for code coupling

IPS [9,30] is a high-performance computing framework developed with focus on flexibility for loosely-coupled, component-based simulations and is maintained by the AToM Fusion SciDAC project. IPS provides services to manage resources and data, and to execute, coordinate and communicate between components. Using these tools, IPS provides a platform for improving use of resources by concurrent “multi-tasking” execution model. IPS is built following a component-based software engineering design, for targeting systems with relatively loose coupling between physics components (modest data exchange, both in volume and frequency). Within this approach, each component acts as a unit of software development and composition, with well-defined boundaries. Components interact with each other through well-defined interfaces. The IPS framework provides the environment for composing the components, as well as a set of services for the interaction between them. In the present case, each physics model maps quite naturally onto the concept of a component.

Further, a lightweight framework approach is taken in the design of IPS to provide a modest set of services to manage resource allocation, configuration, data and task execution. The framework delegates more functionality to components, allowing for flexible – yet robust – environment for fusion simulations. The IPS structure is greatly defined by the ‘driver component’ and ‘plasma state’. The ‘driver’ is a component that controls the simulation workflow. Thus, its inclusion allows the user to control the simulation workflow. The ‘plasma state’ is a repository for all collective information or data that needs to be shared between components. This structure facilitates data management and ensures consistent state of the simulation and parameterization across models.

SOLPS for modeling the ITER plasma conditions

The SOLPS code package [10-12] models the steady-state conditions in the scrape-off layer, including the plasma and neutral transport. The B2.5 code [11] models the plasma, using a 2-D fluid description, allowing calculations over long time and length scales. Neutral particle transport is calculated using either an internal fluid neutral module within B2.5, or via coupling to the kinetic Monte Carlo code EIRENE [13]. A classical plasma transport model is applied in the direction parallel to magnetic field lines (with corrections for kinetic effects), and empirical transport coefficients are used in the cross-field direction. Models for PMI processes are included (sputtering, recycling) via simple models, and multiple particle species and charge states can be included, along with comprehensive atomic physics (ionization, radiation, etc). By modeling the scrape-off layer, SOLPS acts as a bridge between the core plasma and the plasma-materials interface. For a given set of core plasma conditions (e.g., input power and density at the

separatrix), SOLPS calculates the 2-D distribution of density and temperature of all charge states, including the divertor plasma in contact with the sheath at material surfaces and the heat and particle fluxes onto those surfaces. SOLPS has already been used for predictive modeling of the ITER divertor characteristics. The SOLPS-ITER version [35] is used in the present exercise, for containing the most advanced models of the plasma and neutral transport modules.

hPIC Particle-in-Cell (PIC) code for predicting sheath effects on the ion energy angle distributions

The plasma conditions provided by SOLPS are passed as an input to the hPIC Particle-in-Cell (PIC) code. The code is a full-f, full-orbit PIC resolving the plasma sheath physics. In the present work it provides as an output the ion energy-angle distribution (IEAD) of the particles striking the wall. The code has been modified to accept as an input the local plasma properties from SOLPS at the point of interest on the target plate, plus information on the local magnetic field (magnitude of the field and angle with respect to the wall). The code also requires a number of discretization parameters necessary to setup the PIC simulation in ITER-relevant conditions (number of grid points per Debye length, number of time steps per gyro-period, number of ion transit times, numbers of particle per cell). hPIC provides as an output the energy-angle distribution of the ions impacting on the surface, and then passes this information to the material codes handling surface erosion and implantation (F-TRIDYN and GTR). A single hPIC simulation per SOLPS point has been performed with a 1D3V electrostatic mode (one spatial dimension and 3 velocity dimensions).

GTR for modeling W impurity transport in boundary plasma

In order to efficiently model the gross and net erosion of plasma facing components as well as the mapping the kinetic re-distribution and re-deposition of impurities, a leadership scale impurity transport code has been developed. The global impurity transport code (GTR) is designed with modular physics components that accept data for the magnetic field, background plasma profiles, sheath characteristics, and surface characterizations, which are then used to simulate particle tracks. While the physics basis for GTR follows other similar Monte Carlo codes (e.g., ERO [36], DIVIMP [37], BBQ [38], or WBC [39]), for the problems to be addressed within this milestone, we require that both prompt re-deposition (i.e., resolve the full Lorentz motion, as opposed to a gyro- or drift- averaged representation), and global deposition physics are included. In addition, we require the model to interact with realistically detailed PFC surface representation in three dimensions (3D) over large regions of the SOL and far-SOL to capture the impurity sources. Other available codes tend to focus on either small or large domains with commensurate physics (e.g., using guiding center orbits for a large domain, which precludes prompt losses). GTR fills this need within our project by targeting implementation on advanced computing architectures in order to drastically improve efficiency. The physical model is described by the statistics over a set of Lorentz particle trajectories, which yield the $E \times B$ and ∇B particle drifts. Monte Carlo operators describe atomic physics (ionization, recombination), interaction with the background plasma (Coulomb collision effects, including the thermal force), anomalous perpendicular diffusion, and interactions with material surfaces (via sputtering/reflection simulation data or experimental fit), and a modular sheath model for fields near the PFC surface. This approach relies on a trace approximation for the impurities, with the intended use being tracking the tungsten eroded from the surface. From a computer science perspective, this means solving many data parallel ODEs (trace impurity particle trajectories)

with a set of Monte Carlo operators on time steps of ~ 1 ns, out to times of $O(\text{ms})$ and even $O(\text{s})$ for long range transport. Resolving both erosion and volumetric quantities to statistical significance, and at the PFC discretization level we require, takes $O(10^8)$ particles. The parallelization is well suited to GPUs, utilizing the CUDA [40] THRUST library [41]. GITR has already achieved tracking the trajectories of $O(10^8)$ particles in modeling of linear devices experiments.

The erosion rates and energy / angular distribution of sputtered particles (used by GITR), as well as ion implantation profiles needed to compute the material evolution (by Xolotl) are calculated by the binary collision approximation code Fractal (F)-TRIDYN [15], which accounts for key parameters in sputtering, such as ion impact energy, angle and surface roughness. F-TRIDYN also accounts for substrate composition when characterizing ion implantation profiles, and surface damage to the material.

F-TRIDYN for modeling erosion and plasma ion implantation

Fractal-TRIDYN (F-TRIDYN) is a stable and robust version of the binary collision approximation (BCA) code TRIDYN [15], which includes the effect of surface morphology on ion-solid interactions. The code resolves the physics of surface erosion and particle implantation, producing relevant quantities such as sputtering yields, energy-angle distribution of the emitted particles (both material impurities and pre-implanted gas), reflection and backscattering coefficients, and Frenkel pair production.

F-TRIDYN includes two descriptions of surface roughness: an explicit, fractal surface model, and an implicit, statistical surface model based on an arbitrary distribution of surface heights. Fractal models offer an explicit description of the surface morphology, by describing the natural corrugation of the surface as a polygonal fractal. The polygonal fractal surface is resolved directly by exploiting an efficient algorithm of $O(n)$ complexity based on the Jordan Curve Theorem. However, a direct measurement of the fractal dimension of a given surface requires techniques of molecular absorption, which are not readily available for a wide range of material surfaces. Statistical models of surface morphology, on the other hand, are more computationally efficient than explicit models and directly comparable to experiments using standard methods of surface roughness measurement (AFM, SEM, etc.). Complex statistical surfaces can be modeled using a Kernel Density Estimation technique to represent an arbitrary distribution of surface heights in F-TRIDYN.

The sub-surface evolution of the gas species implanted into the tungsten divertor tiles, along with any surface patterning or height changes, will be modeled from a continuum perspective using the Xolotl code [9], which is informed by a “bottom-up” atomistic-based approach.

Xolotl for modeling tungsten divertor tile surface evolution

Xolotl is a continuum based cluster dynamics code solving the spatially dependent Drift-Diffusion-Reaction (DDR) equations in multi-dimension to predict the evolution of the concentration fields for a mixed helium-hydrogen implanted tungsten material. It aims at reaching time scales of 100-1000s and length scales of nm- μm . We use a finite difference implementation in which the grid along the depth direction has a variable step size, with a finer mesh near the surface, in order to better encompass the profile of the incoming fluxes. The code is developed in C++, uses PETSc [42] as its solver, and relies on MPI for high performance parallelization.

The clusters under consideration consist of six types: tungsten self-interstitials (I), tungsten self-vacancies (V), helium (He), deuterium (D), tritium (T), and trapped defect-vacancy clusters (He-D-T-V); the size of each type can vary in a domain suggested by MD simulations. Only small helium clusters, single deuterium and tritium atoms, interstitial clusters, and the single tungsten vacancy mobile in this model. The drift of mobile helium clusters caused by surfaces has been studied by Hu et al. [25-27] and is implemented here as a drift term toward the surface. The set of possible reactions, including formation and dissolution of various clusters is defined to be the same at every spatial grid point. The parameters needed in the previous equations are obtained from other methods: MD simulations can provide the range for the incoming helium flux, the diffusion parameters, and the formation energies; the capture radii of a cluster are computed from geometric considerations.

Re-gridding of SOLPS calculation results of divertor plasma conditions

Extracting quantities from SOLPS requires determining in which cell of the unstructured grid a spatial point resides inside. In order to make this search efficient, the SOLPS cells are loaded into an R-tree data structure. Each node contains a bounding box sufficiently large to encapsulate two sub-trees. Leaf nodes contain a bounding box large enough to contain a single SOLPS cell. When adding a new leaf to the tree, the bounds of the current node are expanded to include the new leaf. A heuristic approach based on minimizing the resulting bounding area of the new sub-trees is used to determine where the leaf node is added. Thus, the new leaf is added to one of the two sub-trees or the current sub-trees are combined into a single branch and the leaf is added as a new branch. Using this data structure, a SOLPS cell containing an arbitrary point can be located with an average $O(\log_m n)$ time complexity where M is the depth of the tree and n is the number of SOLPS cells. At each leaf, a point is determined to exist within a cell, by dividing the cell into two triangles and computing the barycentric coordinate weights.

Once a cell containing the point is located the current cell is divided into 4 non-overlapping quadrants sharing the cell center as a common vertex. Then the values from the neighboring cells are averaged to provide quantities at the corners of the enclosing quadrant. The final quantity is linearly interpolated using barycentric weights of corners and the current cell center. In its current C++ implementation, a million point grid can be interpolated in roughly half a minute using 4 processors. The approach is shown in Figure A-1.




Figure A-1. The SOLPS grid cells (left) are loaded into a R-tree data structure (center). Limiting searches to regions bounding a point, allows rapid interpolation of SOLPS electron density (right) and other quantities.

The 2D background plasma profiles taken from SOLPS include temperature, density, electric potential, and flow velocities. Figure A-2 shows example output for the ion density (Fig.

A-2a) and average background plasma ion charge (Fig. A-2b), which show the area covered by the SOLPS grid. This shows a gap between the simulated SOL grid and the first wall, but a fully simulated plasma in contact with the divertor targets.

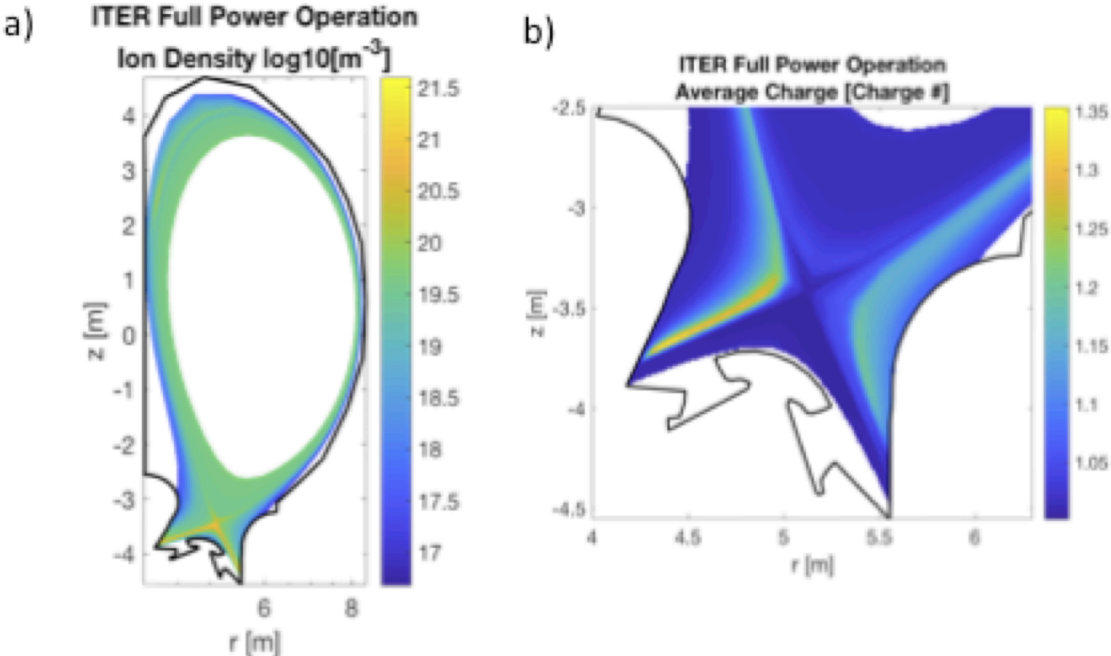


Figure A-2. a) Re-gridded SOLPS output for ITER full power operation showing Ion density on a log scale and the domain covered, and b) a weighted average of SOLPS output (density and charge number) to give average charge number in the ITER divertor region.