Yearly Progress Report (May 2018):
Plasma Surface Interactions: Predicting the Performance and Impact of Dynamic PFC Surfaces

Lead PI: Brian Wirth, UTK-ORNL Governor's Chair Professor of Computational Nuclear Engineering

Institution – Agreement/FWP #	*Principal Investigator*	*Senior Personnel*
ORNL - ERAT953	Brian Wirth	David Bernholdt, John Canik, David Green, James Kress, David Pugmire, Phil Roth, Clayton Webster
ANL –	Barry Smith	
LLNL - SCW1617 / SCW1620	Ilon Joseph	Mikhail Dorf, Milo Dorr, Maxim Umansky
LANL -	Blas Uberuaga	Sham Bhat, Danny Perez, Arthur Voter
PNNL – FWPs 70444/70230	Rick Kurtz	Giridhar Nandipati, Kenny Roche
SNL – FWPs 18-020352/18-020349	Habib Najm	Aidan Thompson, Khachik Sargsyan, Mitch Wood
GA – DE-SC0018423	Phil Snyder	Rui Ding, Jerome Guterl, Orso Meneghini
UCSD – DE-SC0018302	Sergei Krasheninnikov	Russ Doerner, Roman Smirnov
UIUC – DE-SC0018141	Davide Curreli	Jon Drobny, Rinat Khaziev
UMass Amherst – DE-SC0018421	Dimitrios Maroudas	Chao-Shou Chen, Asanka Weerasinghe
UMissouri	Karl Hammond	
University of Tennessee, Knoxville	Brian Wirth	Sophie Blondel, Dwaipayan Dasgupta, Ane Lasa, David Martin, Li Yang, Tim Younkin
RPI -	Mark Shephard	Onkar Sahni, Gopan Perumpilly

The objective of this project is to develop, and integrate, high-performance simulation tools capable of predicting plasma facing component (PFC) operating lifetime and the impact of the evolving surface morphology of tungsten-based PFCs on plasma contamination, including the dynamic recycling of fuel species and tritium retention, in future magnetic fusion devices. Establishing a fundamental physical understanding and developing predictive capabilities of plasma-surface interactions (PSI) requires simultaneously addressing complex and diverse physics occurring over a wide range of length (Angstroms to meters) and time (femtoseconds to years) scales, as well as integrating extensive physical processes across the plasma – surface interface. This requires development of not only detailed physics models and computational strategies at each scale, but also algorithms and methods to couple them effectively in a way that can be robustly validated. Deploying these tools requires the continued development and coupling of leadership-scale computational codes to describe the boundary plasma and the evolving PFC surface, as well as a host of simulations that bridge disparate scales to address complex physical and computational issues at the plasma – surface interface in multi-component materials systems for magnetic fusion energy development beyond ITER.

This project will enable discovery of the key physical phenomena controlling critical PFC performance issues, and the quantitative prediction of their impact on PFC performance during both steady-state and transient plasma conditions. Such phenomena include: (i) surface evolution in regions of either net erosion or net deposition; (ii) the impact of the evolving surface composition and roughness on the

retention and recycling of hydrogenic fuel isotopes; (iii) the impact of dilute impurities on surface morphological evolution and plasma contamination; and (iv) the effects of high-energy neutron damage on surface properties that could influence helium/hydrogenic species retention and recycling. The research activities within this proposed project will focus on three broadly defined research thrusts:

- Bridging the scales between atomistic/microstructural modeling and the continuum-based PFC simulations through the adaptation of multiscale modeling techniques. Specifically, *ab initio*/atomistic/microstructural studies will be used to develop the requisite knowledge to 'coarse grain' disparate, heterogeneous material structures and 'scale bridge' across diverse spatial/temporal scales to provide detailed input to, and calibrate, the Xolotl PFC simulator;
- Integrating boundary plasma and surface evolution models, specifically investigating effects of the plasma sheath and evolving surfaces, in order to predict PFC performance under both steady-state plasma operation and in the presence of transient events. This includes the development of a new code to model the fate of eroded impurities, including their migration and re-deposition, specifically designed to take advantage of leadership-class computing facilities; and
- Studying the dynamic response of the surface to transient plasma events and exploring synergistic phenomena between the near-surface plasma and wall response, with emphasis on dynamic recycling processes that couple the divertor plasma to the material surface layer. The coupled models will explore the impact of wall response on boundary turbulence and divertor detachment dynamics.

The outcome of this project will be a suite of coupled plasma and materials modeling tools, and a leadership class PFC simulator to predict PFC evolution and feedback to the boundary plasma both during steady-state plasma operation and transient events. Towards this end, we plan to have strong coordination of our efforts with the other successful fusion SciDAC projects proposals in the Boundary Physics and Whole Device Modeling topical areas. Success in the proposed research tasks will enable the prediction of both plasma fueling and the sources of impurity contamination that impact core plasma performance, and will lay the foundation for understanding, designing and developing the materials required to meet the performance objectives of future fusion reactors.

The project has established a web presence at: https://collab.cels.anl.gov/display/PSIscidac2/ leveraging R&D infrastructure available in the ANL Mathematics and Computer Sciences Division. This wiki site has both public-facing and project-private spaces. Currently, the public space provides basic information about project goals, participants, sponsors, key collaborators, and a catalog of project presentations and publications. The private space includes records of project and working group meetings and other information. Additionally, we use a Subversion repository, also hosted at ANL, as a repository for internal and under-development documents, and the project mailing list hosted at ORNL.

This document describes the initial progress of the integrated project during the first year from the initial kick-off meeting held in early October 2017 at Argonne National Laboratory, through the development of specific tasks on both the applied mathematics/computer science and domain physics aspects of the program, and finally culminating with the current status of activities in May 2018.

STATUS AND INITIAL PROGRESS ON PROJECT RESEARCH ACTIVITIES

Computer Science/Applied Math

Computer science and applied math (CS/AM) are strongly driven by the science goals of the project and are well integrated through the three major Tasks described blow. In this section, we call out and describe certain of these activities in greater detail in order to illustrate the engagement and interactions within the project. We would also like to note our engagement with researchers at Rensselaer Polytechnic Institute (RPI) as part of Task 2, where the current focus is on infrastructure development in support of parallel, unstructured meshing for particle in cell (PIC) calculations. They are providing a separate progress report, which encompasses their work on multiple SciDAC projects.

Computational readiness for large-scale MD simulations: This work, carried out by Ken Roche (PNNL), was intended to support effective MD analysis of project problems that combine complex analytic, empirically fit, and computationally approximated gas-gas, gas-lattice, and lattice-lattice potentials to study pressure-, temperature-, and time- dependent evolution of helium-hydrogen gas bubbles in tungsten for Task 1.

LAMMPS is our tool of choice for the MD component of these studies. The library was configured to support these project problems specifically, and to exploit architectural details on the IBM BG/Q at the ALCF with GCC and IBM compilers, and on the Cray XK7 at the OLCF with combined GCC + Nvidia compilers. Weak and strong scaling performance of these builds were studied for supported cases, and compared for a range of atom counts per process on single oxide and mixed oxide equilibration test problems, with and without fission gas bubbles, for a range of temperatures to identify 'optimal' execution modes for MPI, MPI + OpenMP, and MPI + CUDA use cases. The IBM BG/Q was selected for a production study as our builds passed fidelity tests and were a factor of two, or more, faster than the software modules provided by the center on the same representative atomic systems with complex potentials. The LAMMPS builds were made accessible to project researchers for production use.

Current and future activities will examine the fidelity and performance of MPI+Kokkos enabled multi-force kernels within LAMMPS, i.e. the use of the Spectral Neighbor Analysis Potential (SNAP), on leadership class computers. Additionally, we are interested to study LAMMPS OpenMP 4.5 kernels designed to exploit low-level parallelism for many-core and GPU platforms.

Conversion of Xolotl to use VTK-m: Xolotl originally used the EAVL library [1] to provide visualization capabilities. EAVL is no longer supported, but its key capabilities, along with those of several similar visualization libraries, have been incorporated in a new community effort, VTK-m [2], which is also designed to take advantage of modern HPC architectures. Recent work by James Kress and Dave Pugmire (ORNL) has focused on replacing EAVL with VTK-m in Xolotl, as well as "upstreaming" changed needed in VTK-m to support features required by Xolotl (one of the first codes to make the transition), including: a software renderer for making 1D and 2D plots, the development of better text annotations to annotate the new plots, and updates to the color table infrastructure in VTK-m. One significant development and improvement over the EAVL implementation, is that the software renderers in VTK-m do not rely on a 3rd party-rendering library to be linked to Xolotl. This makes the visualization and Xolotl build process easier and reduces software dependencies when Xolotl is run on leadership class machines.

At this point, all EAVL visualization components have been replaced with their VTK-m counterparts. One outstanding requirement involves adding support for logarithmic scale color table transforms to VTK-m. These are used in the surface-plotting infrastructure in Xolotl. This work is in progress and has required multiple refactoring of the original VTK-m color table infrastructure. After this work is

completed, we will update the VTK-m annotation infrastructure to meet the Xolotl requirements for being able to annotate output images.

This work impacts all uses of Xolotl, including both Tasks 1 and 2, and in the future within Task 3.

Performance and scalability improvements in Xolotl: In order to reach experimental time scales, Xolotl needs to be able to model larger and larger networks of clusters. However, the Xolotl version used for the FY2018 Q2 PMI Theory Milestone experiments (Task 2) suffered from a memory usage problem, namely that simulation runs would fail due to out-of-memory errors at smaller network sizes than expected on many high-performance computing (HPC) systems.

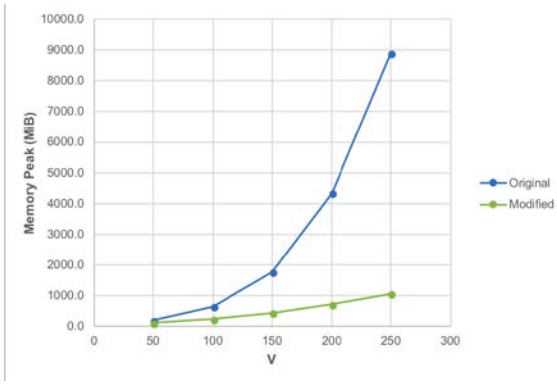

Figure 1. Peak memory usage for original Xolotl implementation and our modified version when running 1D problem on 16 processes in 1 node of the OLCF Eos Cray XC30.

Phil Roth (ORNL), working with Sophie Blondel (UTK), explored two different approaches to address this problem: the use of a grouping scheme to gather clusters together and reduce the number of degrees of freedom but still model the same physics, and profiling Xolotl to purely and simply reduce the memory usage. We identified that the high memory consumption of Xolotl resulted from the need to create two very large two-dimensional matrices to specify the fill pattern of the sparse matrix to the PETSc solver. We determined that the number of non-zeros in these matrices was small, and that with a minor addition to the PETSc application programming interface (API), we could specify the same information using a sparse matrix representation. As shown in Figure 1, our modification reduced the peak memory requirement by 88% for the largest problem that the original version could run on the Oak Ridge Leadership Computing Facility (OLCF) Eos Cray XC30, namely {He,V}={1000,250}. The modification also resulted in a performance benefit, running approximately 1.2× faster than the unmodified code for the V=200 problem size. Our addition to the PETSc API has been integrated into the PETSc development repository by Barry Smith (ANL) for inclusion in future releases. Similarly, at each time step Xolotl was using a non-sparse matrix to compute and transfer the Jacobian to the PETSc solver, which is then using a sparse matrix in the solver. This Xolotl matrix has now been reduced to a sparse matrix which lead to a reduction of memory used from 1.4 GB per core with V=200 to 230 MB during the time steps.

As part of our ongoing effort to improve Xolotl performance, we have also been investigating the use of threads on systems with no graphics processing units (GPUs), since our target system for the Q3 PMI Theory Milestone does not contain GPUs. We recently conducted a study comparing runtime configurations that vary the ratio of threads to processes used in a Xolotl run. Figure 2 shows the time required to complete the most expensive parts of program initialization and time stepping on the OLCF Eos system. We determined that a high threads-to-processes ratio

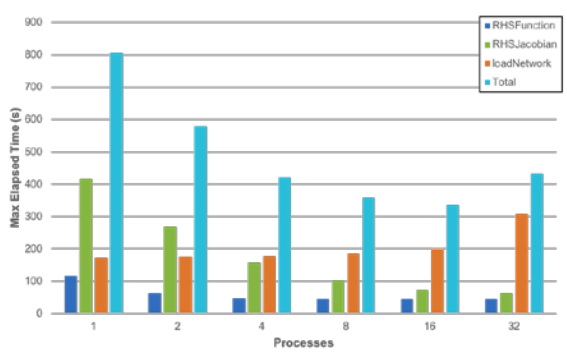

Figure 2. Time required to run the three most expensive activities of a Xolotl simulation run when trading off processes for threads. Timings shown for 1D problem on 32 total threads in 1 OLCF Eos compute node.

benefits the initialization, but time-stepping does best with a single thread per process in the current implementation. We are currently working on an implementation that changes the number of threads used as the simulation runs.

General co-design approach to GITR and overall coupling: The development of the global impurity transport code (GITR) has focused on modularity to interface with plasma sheath and boundary simulations and capability for coupling with materials surface codes, and performance to simulate impurity transport in large sections of fusion machines in high geometric detail. Interactions across the plasma boundary and material surface are of utmost importance to fusion plasma performance and material degradation. Therefore inclusion of phenomena on different time and spatial grid sizes must be included. Data and reduced models derived from the science codes Xolotl-PSI, SOL-PS, F-TRIDYN and hPIC are made available to GITR (and vice versa) via interfacing through the integrated plasma simulator (IPS) python framework. In addition to the data processing/passing, the IPS manages the compute resources on which these codes run. GITR performance has focused on simulating sufficient numbers of Monte Carlo particles in a large high fidelity geometry to derive good statistics from the simulation output. This has been done using a CUDA Thrust library implementation which iterates the integration scheme and physics operators on particle objects to evolve the physical system. This implementation can be run taking advantage of all GPUs on a compute node or CPUs by specifying the CUDA or OpenMP device backend within the Thrust library at compilation. The MPI implementation allows GITR to take advantage of many nodes which is efficient for GITR because of minimal intra-node communication during the computation. Presently we are investigating improving both the weak and strong scaling via alternate implementations of the global reduction operations required to compute the volumetric impurity densities. In addition to these advancements in the GITR code development, a simultaneous development of advanced meshing software on accelerated architectures is being advanced at the Rensselaer Polytechnic Institute. The needs for unstructured meshing of near surface background plasma and material boundary data are motivated both by the input of physics data to GITR as well as performance. The developers of the parallel unstructured mesh infrastructure (PUMI) are porting the advanced meshing algorithms for capability for more broad use on accelerated architectures.

Computational Aspects of the Modeling of Linear Plasma Device Experiments: The deliverables for the first two quarters of the FY2018 Fusion PMI Theory milestone involved the development and use of coupled plasma-material interaction simulations to model experiments conducted at the PISCES-A linear plasma device. This extensive effort has included numerous PSI team members on both the physics and CS/AM sides from ORNL, UTK, UIUC, and UCSD, as well as Wael Elwasif (ORNL) of the AToM project. The IPS framework, developed primarily by the AToM project and its precursors, was used to couple F-TRIDYN, GITR, and Xolotl in order to track the migration and re-deposition of tungsten impurities sputtered off of a tungsten target under irradiation by a mixed He/D plasma. As part of the Q2 deliverable, an initial performance characterization of the coupled simulation was performed to set the baseline for future improvements.

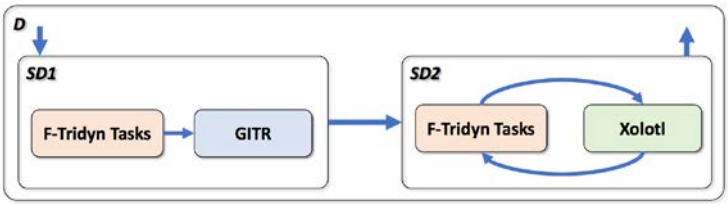

Figure 3. IPS workflow for coupled F-TRIDYN/GITR/Xolotl simulation of He-D plasma in PISCES linear device.

The simulation workflow followed several steps, as shown in Figure 3. First, F-TRIDYN is used to generate tungsten sputtering yields, and the sputtering products are tracked with GITR. F-TRIDYN is a sequential code, and execution is parallelized over the spaces of interest in projectile species, energy, and angle with many instances of F-TRIDYN run concurrently in an IPS task pool. These results are collected and fed into GITR, a parallel code, which carries out the impurity transport simulation (in this case, 10^5 particles for

10^5 time steps). The results of this phase of the simulation are then fed into a time-stepped simulation, which uses Xolotl to model the evolution of the tungsten surface under irradiation by the plasma and re-

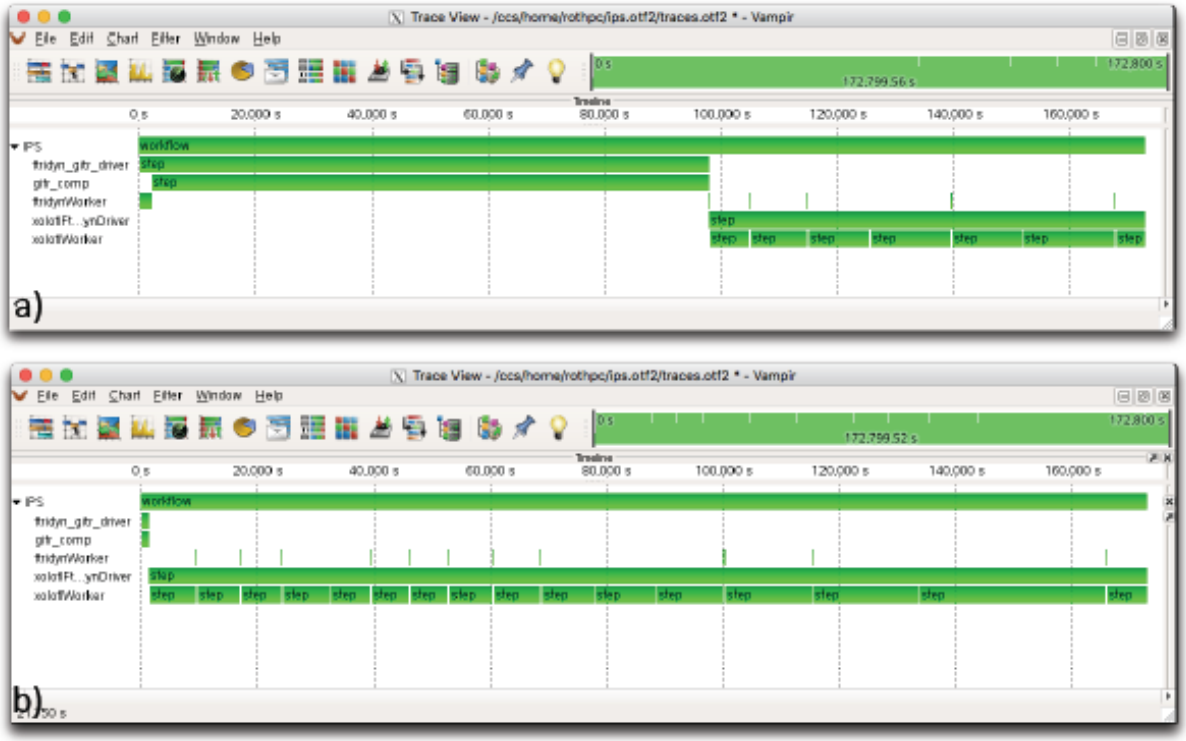

Figure 4. Timeline visualization of a) original 48-hour coupled IPS/F-TRIDYN/GITR/Xolotl simulation of He-D plasma in the PISCES linear device, and b) after correctly setting the OpenMP environment variables for GITR.

deposition of the tungsten impurities, and F-TRIDYN to model the implantation and sputtering. Once again, F-TRIDYN is parallelized via a task pool, while Xolotl is parallel.

Simulation runs carried out on the Cray XC30 (Edison) at NERSC initially produced 0.28 s of simulated time in a 48-hour (wall-clock time) run using 16 nodes (384 cores), which was much less than desired. Initial performance analysis, led by Phil Roth (ORNL), showed that most of the time in the simulation was spent in GITR and Xolotl, which was expected, but the amount of time in GITR was unexpectedly high. Further analysis showed that the configuration parameters specifying the OpenMP thread deployment were suboptimal, as shown in Figure 4. When the configuration was corrected, the same 48-hour, 16-node simulation produced 7.95 s of simulated time, an increase of 28x.

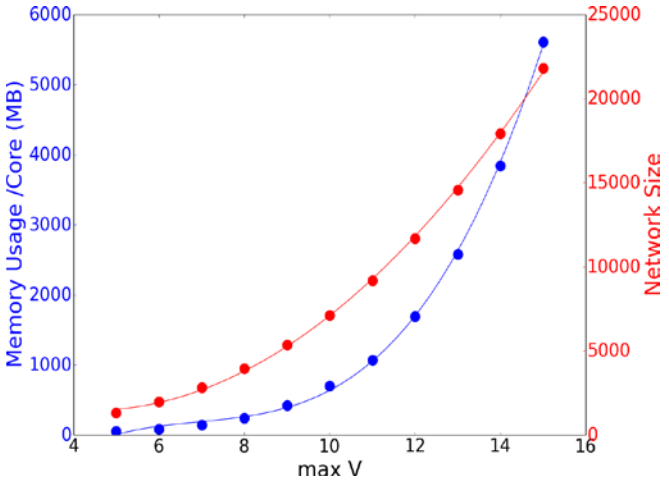

Figure 5. The maximum physical memory usage per core, and reaction network size as a function of the maximum vacancy size of the He-D-vacancy mixed clusters in Xolotl. The simulations were performed with a simulation setup similar to the PISCES He-D simulation.

We also found that even though these simulations used a relatively small reaction network, Xolotl tended to be memory-limited, and had to be deployed with less than one MPI rank per CPU core in order to ensure sufficient memory. A more detailed examination of the scaling of memory usage with network size in Xolotl is shown in Figure 5, and provided some of the motivation for the memory usage improvements discussed above.

Uncertainty Quantification Strategy and Activities: Given a computational forward model, with uncertain inputs, our UQ strategy involves a general workflow, relying on probabilistic forward and inverse UQ methods, as follows:

1. Given experimental measurements, use Bayesian inference to estimate a joint density on uncertain model parameters
2. Employ global sensitivity analysis (GSA), coupled with polynomial chaos smoothing and sparsity constraints, to identify the subset of important parameters for UQ
3. Using adaptive sparse quadrature, propagate uncertainty forward through the computational forward model, providing prediction of uncertain model outputs.

Where needed, if the measurement model used in step (1) above is computationally expensive, step (3) can be used to build a polynomial chaos surrogate to be used in the Bayesian inference process. Also, and again where needed, we rely on Bayesian model selection to identify the best model for fitting the data in case there are alternatives. Finally, this strategy accommodates coupled models as needed, incorporating the full set of parameters in multiple models as needed. This work is being carried out by Habib Najm, Khachick Sargsyan, and Tiernan Casey (SNL). We also expect to bring Clayton Webster (ORNL) into these efforts going forward.

UQ in Xolotl: We are working on employing elements of this workflow as needed on different parts of the project. In the case of Xolotl, we have been collecting expert-defined uncertainty specifications on each of the uncertain input parameters to proceed directly to GSA. This collection of relevant inputs is near completion, and will be followed by generation of samples for GSA in Xolotl.

Langmuir Probe Processed Data Fitting for GITR: Plasma profiles were fitted using *processed* data obtained from Langmuir probe measurements. More specifically, we have employed a Bayesian approach for fitting ion flux and temperature profiles as functions of radial distance. The form of the parametric fit is a third-order polynomial expansion with two features: a) scaled Legendre polynomials were chosen as expansion bases, since their orthogonality makes the coefficient inference better conditioned, and b) the right end of the profile is constrained to have vanishing derivative, manifested as a simple linear dependence of the highest-order coefficient on the rest of the expansion coefficients.

The initial Bayesian inference formulation via error-in-variable models revealed the major drawback in interpreting the processed data error bars -- both horizontal (in radial distance), and vertical (in quantity of interest, Temperature or Ion Flux) -- as data noise. In fact, we subsequently understood that neither of these two error bars were representing data noise. We have reformulated the fitting problem where the vertical error bars are interpreted as modeling noise and parameterized as standard deviation together with the mean values, while ignoring the horizontal error bars given our improved understanding of the data gathering and processing scheme. The approximate likelihood for fitting is derived from a Kullback-Leibler distance between uncertain model prediction and the uncertain representation of processed data. Figure 6 illustrates the posterior predictive distribution of the Ion Flux profile. Figure 7

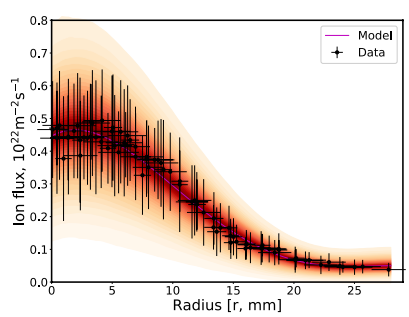

Figure 6. Posterior predictive distribution as compared to the ion flux profile.

demonstrates a set of samples drawn from the posterior predictive distribution. The uncertain representation of both Ion Flux and Temperature (not shown) profiles will augment the input list of GITR, feeding into forward uncertainty propagation.

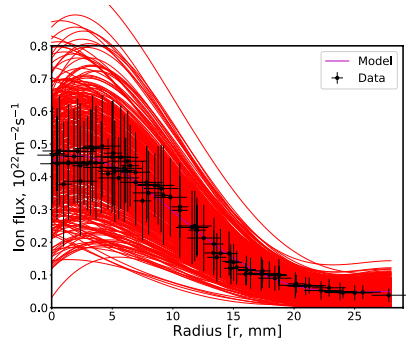

Figure 7. Samples drawn from the posterior predictive distribution for forward uncertainty propagation in GITR.

Langmuir Probe Raw Data Fitting for GITR: Raw experimental data from Langmuir probe measurements collected from a translating sensor are available as probe voltage [Volts], probe current [Amps], and radial position [mm] at $1x10^5$ measurement points polluted with sensor noise. In order to extract meaningful physical quantities from these measurements the data must be fit using concepts from Langmuir probe theory.

The probe voltage arises from a voltage signal driven at ~140 Hz during the probe flight, and is shown in Figure 8, which we segment into an ensemble of i-V curves at the associated radial probe positions as shown in Figure 9. We are implementing a Bayesian statistical inference procedure to estimate the physical parameters (electron temperature, number density) with uncertainty from these noisy i-V profiles. Figure 10 shows where the data has been fit using a multiplicative Gaussian noise model. Ultimately we aim to employ a piecewise fitting approach across V-space, such that we can adequately identify the regions of the profiles associated with ion current saturation (V<<0), Maxwellian electron current, and electron saturation (V>>0) directly from the data. This approach will be embedded into a hierarchical modeling framework to represent electron temperature and number density as uncertain spatial fields and will be compared to the original domain expert fitting performed by the experimentalists to assess consistency.

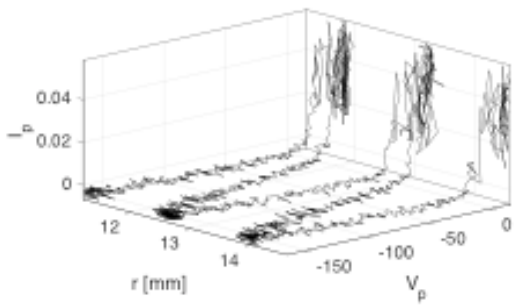

Figure 8. Experimental raw data for the Langmuir probe current versus voltage traces with time.

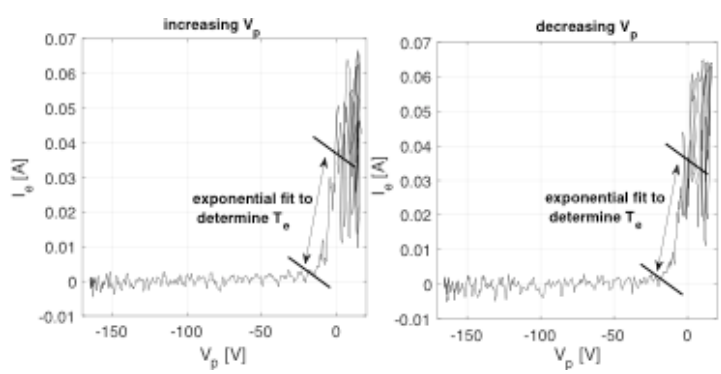

Figure 9. Extracted Langmuir i-V curves with either increasing or decreasing Vp.

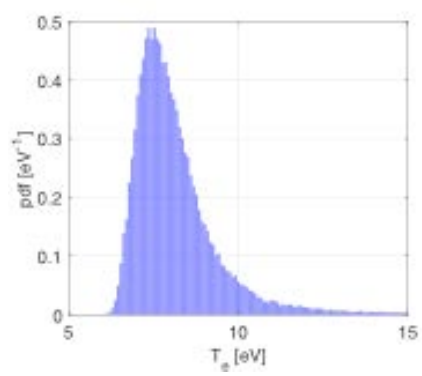

Figure 10. Resulting distribution of electron temperature to be used in the UQ analysis.

Improved Solvers for COGENT: The preliminary ELM heat pulse investigation described in the section on dynamic wall response and feedback to the divertor plasma (Task 3) assumed adiabatic electrons. The extension of these results to include kinetic electrons requires the ability to treat the associated fast time

scales implicitly in COGENT. Milo Dorr (LLNL) has therefore been working with the COGENT team to extend the implicit-explicit (IMEX) capability in COGENT to achieve this goal. In an IMEX approach, a nonlinear equation must be solved at each implicit stage, yielding a self-consistent, coupled update of all fast processes (e.g., collisions, Alfven waves). We employ a Jacobian-free, Newton-Krylov algorithm for the solution of this nonlinear equation, in which the construction of an efficient pre-conditioner is critical for the performance of the method. Because our initial IMEX implementation addressed fast collisional scales only, we were able to utilize a pre-conditioner with a relatively homogeneous banded structure. To incorporate more general fast scales, including those arising from fast electron dynamics involving both configuration and phase space couplings, a more general blocked pre-conditioner is needed. We have therefore been refactoring parts of COGENT to achieve the required flexibility. This includes the definition of interfaces enabling individual plasma species to provide their contribution to the pre-conditioner when called from a high-level driver. We are currently performing tests including an electron vorticity model combined with collisions and anomalous transport.

Task 1. PFC Surface Response (Scale-bridging to extend capability of Xolotl PFC Simulator)
Development of Be-W potentials: Simulation methods such as Molecular Dynamics (MD) are powerful tools for examining behaviors of materials that originate on the atomic scale. On modern supercomputing platforms it is possible to simulate many millions of atoms on time scales approaching 1μs. This is possible because an explicit treatment of the electronic degrees of freedom can be avoided and instead a classical interatomic potential (IAP) can be constructed to represent the Born-Oppenheimer potential energy surface for the nuclear degrees of freedom. The IAP is the leading approximation made in MD simulation and special care must be taken to construct and test it prior to deploying it in a research program. In other words, the IAP is the critical multiscaling link between quantum electronic structure methods, such as density functional theory (DFT), and the materials behaviors that emerge in multi-million atom MD simulations.

Our approach to constructing this multiscaling link employs machine learning in the form of the Spectral Neighborhood Analysis Potential (SNAP) to 'learn' the features of a DFT database in order to fit the IAP. This automated learning is done by translating the DFT data into a basis set of descriptors of the atomic environment. SNAP uses the 4D bispectrum [3] as a descriptor. The optimal values of SNAP coefficients are determined using the FitSNAP.py software, while SNAP hyperparameters are determined via a genetic algorithm search (GA) using the Dakota package [4]. A schematic view of the optimization procedure is shown in Figure 11. The GA operates on the outer loop connecting Dakota and FitSNAP.py. Since the GA requires only infrequent evaluation of global information, the main optimization loop can be efficiently parallelized across a cluster of connected compute nodes.

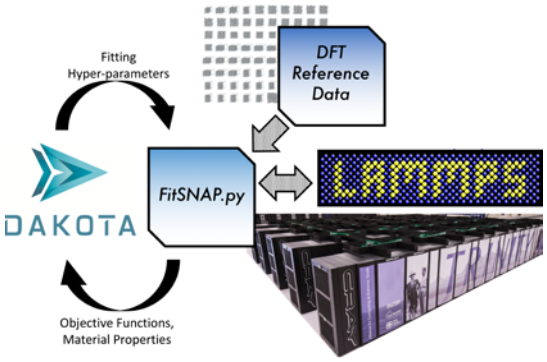

Figure 11. Schematic of the fitting process used for SNAP interatomic potentials. Material properties are learned through an optimization of the weights applied to the input DFT training data.

The first step in this process is to build up an extensive training set of DFT calculations that represents the material system of interest. Here we will report progress on a SNAP IAP for the W-Be system. Over the past few months we have assembled a very large database of DFT calculations (~25k individual DFT calculations, ~1.8M training points) for both single-element and mixed configurations of tungsten and beryllium. This database includes elastic deformations, defects, liquids, free surfaces, and many other configurations that are representative of those encountered during an MD simulation of plasma surface

interactions. To make a suitable W-Be IAP we first fit a potential for both pure component using just a subset of the total training set, then optimize the fit on the training data of the mixed components.

Objectives for the GA optimization must characterize the accuracy of candidate IAPs with respect to the training data, as well as suitability for MD simulations of plasma-surface interactions. This led us to a combination of material properties including elastic constants, relative phase stability, and point defects in otherwise perfect crystals. Figure 12 compares properties calculated with our converged W-Be SNAP potential (shown in green) to those of an existing IAPs for tungsten [5] and beryllium [6] (shown in red), as well as single-element SNAP potentials (shown in blue). For each of the self-interstitial (SIA) defects the SNAP potentials predict extremely errors with respect to DFT calculations. Most notably are the di-vacancy binding energy in tungsten and single vacancy formation energy in beryllium. These were not well represented by the existing IAPs. The elastic constant errors are averaged over all the principle elastic moduli of the most stable phase (W: BCC, Be: HCP). With the exception of the Be elastic constants, the full W-Be SNAP potential preserves or improves on the accuracy of the single-element SNAP potentials.

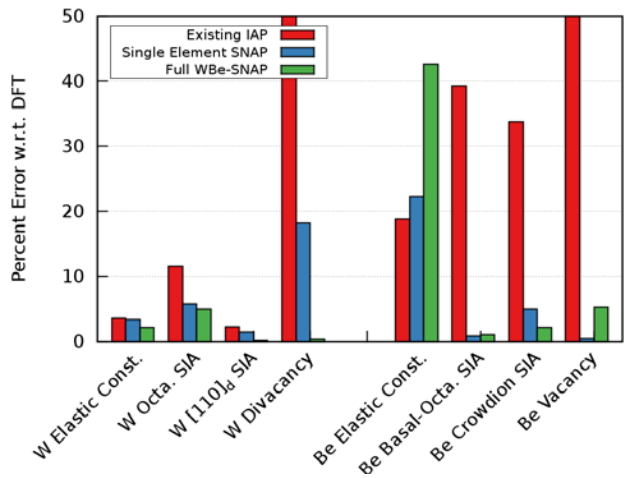

Figure 12. Demonstration of the accuracy gained by using a machine learning SNAP interatomic potential for the tungsten – beryllium system. Potentials taken from the literature (red, [5,6]) are compared to the single-element SNAP potential for pure component (blue) or the full W-Be SNAP potential (green).

While the DFT training set used here is extensive, it is inherently limited to describing material properties representable by a few dozen atoms. The next step in this IAP development work is to test the extrapolation capability of the full W-Be SNAP model to situations that are much closer to the large scale simulations of plasma-facing materials. While this is ongoing work, we have gathered evidence that the new IAP predicts correct melting temperatures for the WBe_2 intermetallic phase, as well as good phonon band structures. Neither of these features were included in the training data or the GA objective functions. This gives us some confidence that this potential is suitable for broad use within this material system.

Large-Scale MD Simulations of H-He behavior in tungsten: Large-scale MD simulations have begun in earnest on hydrogen–tungsten systems at both NERSC and ALCF. We have exhausted our allocation at ALCF (Mira) and are now working either in the backfill queue or are using NERSC's Cori system, in anticipation of moving some calculations to Theta once it comes online in the 2018-2019 ALCC allocation.

The first task within this part of the project is the continued development of an atomistic database of hydrogen and helium aggregation kinetics and retention in sub-surface gas bubbles. We have accumulated a database of helium–tungsten bubble aggregation simulations (primarily from the prior SciDAC) consisting of more than 15 simulations. Each of these requires between 250 and 33000 cores (depending on system size, which due to limitations of run time between insertions is roughly inversely proportional to flux in our database) run continually for a year or more. These simulations span several decades of flux (10^{25} $m^{-2}s^{-1}$ to 10^{27} $m^{-2}s^{-1}$) and four different surface orientations [(100), (110), (111), and (211)]. A comparison of these simulations across different surface orientations was recently published [7]. A

manuscript telling the story of the effect of flux and fluence—including one simulation carried out as far as 2.5 μs, well beyond the onset of frequent bubble-bursting—is anticipated for submission in the very near future. These calculations have been an invaluable benchmark for Xolotl development, and have spawned several "child" manuscripts on the mechanisms of bubble formation and other tangential issues.

The second activity within this part of the project involves studying hydrogen diffusion in the presence of helium bubbles, there are two large-scale MD simulations involving a 50 nm × 50 nm × 25 nm slab of tungsten with a free surface in the *z*-direction and periodic boundaries in the *x*- and *y*- directions. The total number of tungsten atoms in the simulation cell is about 4 million. In one case, approximately 1000 hydrogen atoms were inserted in a 1 nm layer below the helium bubbles, as formed from a 1 μs simulation of helium plasma exposure. The diffusion of hydrogen appears to be relatively independent of the presence of helium, and is quite slow compared to the rate of helium motion. An example is shown in Figure 13.

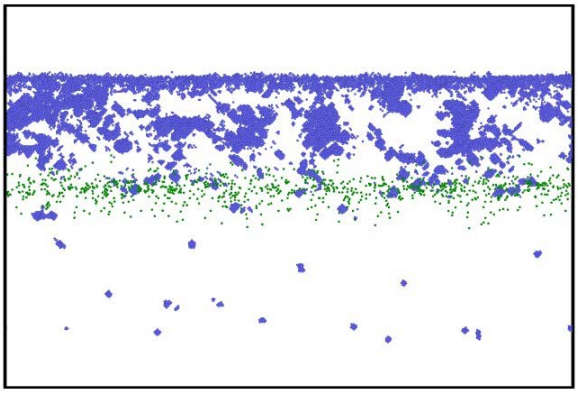

Figure 13. Side view of the MD simulation at 3 ns following introduction of hydrogen (green), in which the hydrogen was introduced below a helium (blue) bubble microstructure that developed over 1 μs of 100 eV helium implantation at a flux of 4.0×10^{25} $m^{-2}s^{-1}$.

This research into hydrogen diffusion behavior was recently aided by the adoption of a new interatomic potential, which decreases computer time dramatically. Preliminary results for hydrogen plasma exposure in tungsten indicate that hydrogen behavior in tungsten (in the absence of helium) is *fundamentally different* from helium behavior. Not only are retention patterns and surface behavior significantly different, which was anticipated, but we have observed that hydrogen is generally immobile in tungsten on short time scales unless it encounters another hydrogen atom nearby. Since two hydrogen atoms in proximity tend to repel each other, this leads to much deeper hydrogen penetration, a general lack of clustering, and very different retention patterns. This causes problems with the insertion of hydrogen, as the location of an existing hydrogen may well determine the behavior of a freshly inserted one (since hydrogen atoms do not cluster in the absence of impurities or defects). Specifically, the approximation we have used to great effect during the previous SciDAC—that helium can be directly inserted into tungsten stochastically, without the need to dynamically observe the energy dissipation of each bombarding ion explicitly, thus saving on 90% or more of the computer's time—is not a good approximation for hydrogen. Paradoxically, direct implantation of thermalized hydrogen actually results in *less* retention than direct hydrogen bombardment of the surface at the same actual flux!

We are continuing to explore hydrogen clustering and transport behavior, with emphasis on both long- and short-term effects of hydrogen bombardment on hydrogen transport.

Explaining mechanisms of hydrogen trapping in tungsten: UCSD has used molecular dynamics simulations to investigate the possibility of formation of self-trapped hydrogen platelets in tungsten. Such platelets can form either spontaneously, when the interstitial hydrogen concentration is larger than about 10 at.% (Fig. 14a), or assisted by dislocations and external non-hydrostatic stresses at lower hydrogen concentrations on the order of about 1 at.% (as shown in Fig. 14b and 14c). The MD studies of the hydrogen platelet structure reveals that hydrogen self-trapping occurs without formation of tungsten lattice defects. Instead, the tungsten lattice within the platelets changes from bcc to fcc, and a significant part of the trapped hydrogen atoms occupy octahedral interstitial positions, as compared to the usual

tetrahedral interstitial hydrogen positions in bcc tungsten. The hydrogen-to-tungsten atomic ratio within the platelet is close to 1:1, which corresponds to the stoichiometry of the high-pressure tungsten hydride phase WH. We also found that the hydrogen platelets induced by dislocations can contain thousands of hydrogen atoms per one tungsten atom on a dislocation line. Thus, the trapping capacity of the dislocation

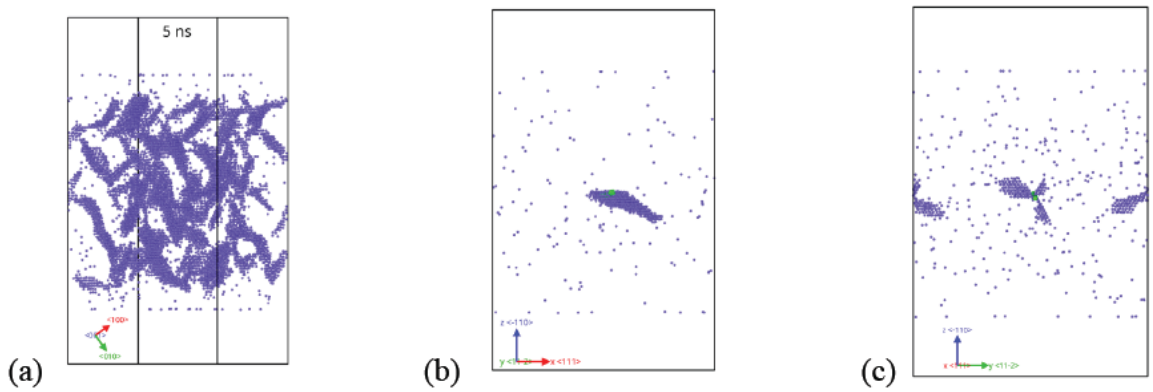

Figure 14. MD simulation snapshots showing hydrogen atoms (blue spheres) forming self-trapped platelets in tungsten, a) spontaneously at 10 at.% hydrogen concentration, b) near a tungsten edge dislocation, or c) near a tungsten screw dislocation. The images in b) and c) are obtained with a hydrogen concentration of 1 at. %.

significantly exceeds the trapping capacity of other non-cavity kinds of defects in tungsten.

Our analysis of stresses in the tungsten lattice during the platelet formation reveals that the mechanism of the hydrogen agglomeration is related to the interaction of interstitial hydrogen atoms with the non-hydrostatic stress components, which allows the hydrogen to overcome the repulsion from interstitial hydrogen in tungsten to form hydrogen clusters. The hydrogen clusters induce local transformation of the tungsten lattice from bcc to fcc and generate stresses, which cause further growth of the platelets by hydrogen self-trapping. This mechanism is different from formation of Cottrell atmospheres near dislocations in metals. We have also evaluated the de-trapping energy of hydrogen atoms from the platelets. It is found to have tendency to increase when the size of the platelets or the temperature increase and in a steady-state at the high temperatures ~1000 K is around ~0.9 eV, which is in the range of the de-trapping energies of hydrogen in plasma exposed tungsten samples experimentally measured using (thermal desorption spectroscopy) TDS technique. These findings may explain the experimentally observed widening of the hydrogen TDS peaks for tungsten, as well as the experimental observations showing very large near surface hydrogen concentrations in low-energy deuterium plasma exposed tungsten. Such a mechanism may also impact the efficiency of isotope exchange process for tritium extraction, which otherwise is limited by two-stage mechanisms of hydrogen de-trapping from vacancies and bubbles.

We have also developed a hydrogen retention model to evaluate the quantity of hydrogen trapped in tungsten by the dislocation-induced platelets under steady-state conditions. The model shows that for the typical temperatures of plasma exposed tungsten samples and the measured dislocation densities in ITER grade tungsten, the platelets can accommodate a concentration of retained hydrogen corresponding to the high ~0.1 at.% average density found experimentally in plasma exposed tungsten samples. Moreover, we found that presence of the hydrogen platelets facilitates generation of dislocations under stress conditions, which can further explain the even higher ~1 at.% density of retained hydrogen seen in the first ~0.1 μm below the surface of plasma exposed tungsten.

Experimental investigations of tungsten fuzz: We have experimentally studied the mechanism of tungsten fuzz formation at the UCSD PISCES laboratory. Measurements were conducted to identify the motion of tungsten and helium atoms during the formation of tungsten fuzz. The experiments employed helium and tungsten isotopes, to address the following conceptual issues: (a) does the fuzz become opaque for helium ions impinging on the sample? (b) are helium nano-bubbles in both the fuzz tendrils, and at their base, 'frozen' after they are formed and no longer undergo changes? (c) does the fuzz growth, and nano-bubble layer formation, result in the mixing of tungsten atoms and thereby impact the distribution of atoms in the sample?

In a first series of experiments the mobility of helium within the growing fuzz was measured by adding He^3 to the different stages of plasma exposure under conditions that promoted tungsten fuzz growth. Ion beam analysis was used to quantify the amount of He^3 remaining in the samples following the plasma exposure. Specifically, He^3 was included in the initial phase of the plasma discharge in one sample and toward the end of the plasma discharge in other samples. No evidence of He^3 was observed retained in the tungsten when the He^3 was only present in the plasma initially. On the other hand, He^3 was detected in the samples exposed to He^3 containing plasma at the end of the discharge. It was also found that the amount of He^3 in the tungsten sample scales directly with the fluence to the samples, rather than with a square root of time, which would be expected for diffusive-like process. These results demonstrate that the retention of helium in bubbles within tungsten is best described as a dynamic process with direct helium implantation rather than diffusion into the bubbles.

In the second experiment, an isotopically enriched layer of tungsten (~92.99% of W^{182}) was deposited on the surface of a bulk tungsten sample with the natural abundance of the isotopes. This sample was then exposed to helium plasma at the conditions necessary to support the formation of tungsten 'fuzz'. Depth profiles of the concentration of each of the tungsten isotopes were obtained using secondary ion mass spectrometry (SIMS) before and after the plasma exposure. The depth profiles clearly showed mixing of tungsten atoms from the bulk sample toward the surface of the fuzz, as shown in Figure 15. These results demonstrate that: (i) fuzz is not completely opaque for helium ions impinging on the sample and, at least, some incident ions penetrate through to the fuzz base; (ii) helium nano-bubbles exhibit a dynamic behavior which is manifested in nano-bubble bursting and reformation; (iii) as a result of this dynamic process the motion of tungsten atoms in the region of helium filled bubbles demonstrate enhanced mobility which results in their strong mixing. Thus, our experimental findings support a physical picture of the dynamic behavior of helium bubbles that also causes an enhanced mixing of tungsten atoms.

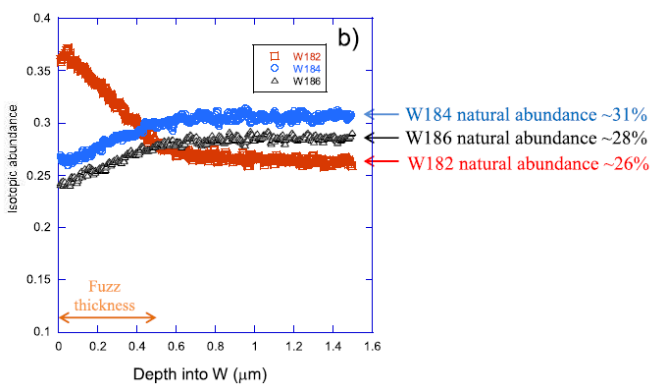

Figure 15. Depth profiles of tungsten isotopes from a sample coated with a 15 nm thick layer of W182 isotope after 1 hour of helium plasma exposure that produced fuzz growth.

Accelerated Molecular Dynamics simulations: Continuing the theme of previous studies, we have continued to examine the evolution of He in tungsten at He arrival rates representative of fusion conditions using accelerated molecular dynamics (AMD) methods. Two recent results, both submitted for review, are highlighted here. First, we have examined how He bubble growth differs at grain boundaries than in the bulk. In the bulk, as a bubble grows, it emits interstitials which form dislocation lines on the

bubble surface and, if conditions are right, detach from the bubble and escape from the bubble vicinity. At least for a Σ5 grain boundary, this does not happen. Instead, as highlighted in Figure 16, the interstitials remain trapped in the vicinity of the bubble, forming an interstitial halo around the bubble. This is important as the transport rate of He is much slower in this halo region than in the rest of the boundary, and this limits the supply of He to the bubble. In essence, the formation of the interstitial rich halo region reduces the growth of the bubble. This is a different bubble growth mode than in the bulk and would lead to a very different bubble distribution than if the halo did not form. This work is under review at *Materials Research Letters*.

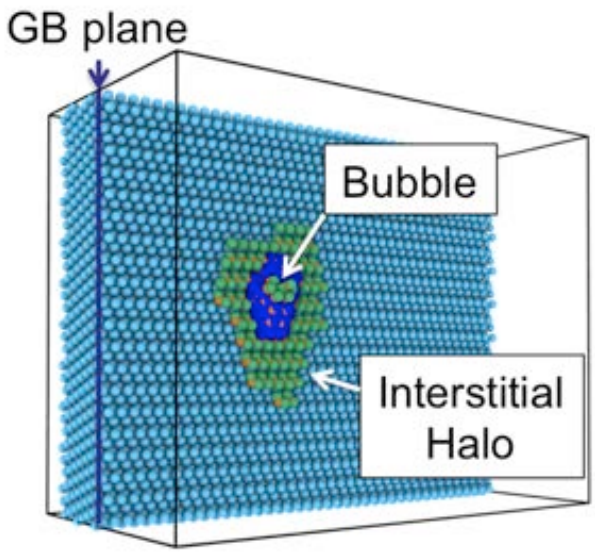

Figure 16. AMD simulation snapshot of a growing helium bubble (dark blue) at a tungsten (light blue) Σ5 grain boundary, which is surrounded by a 'halo' of interstitial atoms (green).

In parallel, we have been examining the major assumption of our previous bubble growth work, the assumption about how He arrives at the bubble in the first place. Before, we simply inserted He into the center of the bubble. However, work by Smirnov et al. using extremely high helium introduction rates of 10^{12} helium atoms/s [8] has shown that He will interact with the dislocation loops that are generated on the surface of growing bubbles. To analyze this more systematically, we again used AMD methods to examine the growth of bubbles, allowing the He to diffuse in from the tungsten matrix. We also see that He is trapped at evolving dislocation lines, as illustrated in Figure 17. This leads to a distributed network of bubbles rather than a single, growing bubble. The nature of this network depends on both the temperature and the implantation flux of helium. For higher temperature or lower flux, there is a greater tendency to form one larger bubble, a consequence of the faster rate of He migration compared to implantation which allows for more time between implantation events for the He to find the central bubble. Our results indicate that the kinetic interaction of He with generated dislocations is a key factor dictating the evolution of bubble distributions in plasma-exposed tungsten. We speculate that the formation of this network is correlated with the formation of fuzz. This work is under review in *Acta Materialia*.

Finally, we have been involved in studies of the effects of grain boundaries on damage evolution in tungsten. In an effort in collaboration with experimentalists at LANL, we have been using kinetic Monte Carlo models to examine how grain boundaries impact defect evolution to potentially mitigate the overall accumulation of damage in materials. Our results indicate that the ability of a grain boundary to mitigate damage depends on whether the grain boundary saturates or not. That is, if a grain boundary fills with defects, such that a back-flux back into the grain is established, the ability of the grain boundary to enhance the annihilation of defects is reduced. Several papers exploring this behavior are under preparation.

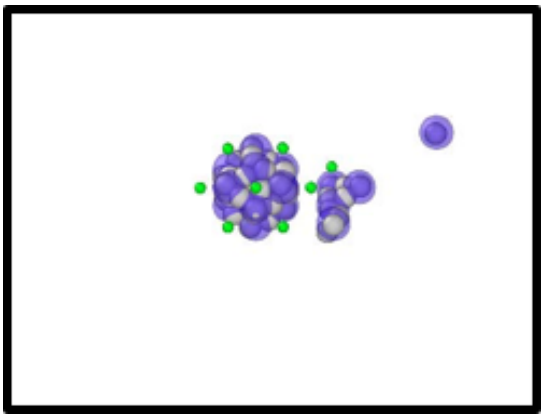

Figure 17. AMD simulation snapshot of small growing helium bubble (dark blue spheres) interacting with tungsten interstitial defects (green spheres), which limits adsorption of additional helium atoms.

Analysis of Helium Segregation on Surfaces of Plasma-exposed Tungsten at Different Levels of Helium Implantation: Using molecular-dynamics (MD) simulations and molecular-statics (MS) computations based on a reliable many-body interatomic potential [5], we have analyzed helium segregation on tungsten surfaces for tungsten that has been exposed to different levels of helium implantation. With increasing levels of helium implantation at higher fluence, mobile He_n clusters ($1 \leq n \leq 7$) are subjected to cluster-cluster and cluster-bubble interactions in addition to cluster-surface interactions, which complicate cluster dynamics beyond the dilute limit of helium content in the PFC material. We have characterized in detail configurations generated by large-scale MD simulations of implanted helium evolution in plasma-exposed tungsten with W(100), W(110), W(111), and W(211) surfaces facing the plasma. Specifically, we have examined the effects of varying helium fluence due to increased plasma exposure of the tungsten PFC by performing systematic MS computations of small He_n cluster energetics near the above low-Miller-index W surfaces as a function of distance (depth) of the cluster center from the surface on a grid of lateral locations on the surface. We have analyzed the defect interactions that mediate the energetics of small He_n clusters migrating to the surface, taking into account that the migrating cluster also is subjected to the stress fields generated by larger He bubbles, as well as other He_n clusters, and quantified the strength of these interactions for different levels of He implantation. The outcome of this analysis is the systematic parameterization of mobile helium cluster energetics at varying levels of He implantation through functional forms that include contributions from cluster-cluster and cluster-bubble interactions in addition to the cluster-surface interactions. Such parameterizations are important for developing atomistically-informed, hierarchical multiscale models of helium cluster dynamics in PFCs. This work is currently in preparation to be submitted for publication in the *Journal of Applied Physics*.

KMC-MD studies to predict structure evolution and properties: This task aims to predict microstructure and property evolution of plasma-facing materials exposed to particle bombardment from the plasma, including neutrons, under isothermal and variable temperature conditions. We have completed the first stage of computations of the thermal expansion coefficient of plasma-exposed tungsten, $\alpha(T, p)$, consisting of reference-state computations on defect-free crystalline tungsten over a range of temperature, T ($0 \leq T \leq 2000$ K), and pressure, p ($0 \leq p \leq 10$ GPa). A systematic computational protocol has been implemented to determine the lattice parameter as a function of temperature, $a(T)$, and its temperature derivative $\partial a/\partial T$ at various pressures based on isothermal-isobaric, i.e., (NpT), molecular-dynamics (MD) simulations according to the Juslin-Wirth interatomic potential for W [5] for well-converged supercell sizes; $\alpha \equiv [(\partial a/\partial T)/a]_{N,p}$. The results are shown in Figure 18a, and demonstrate a ~40% reduction in the value of α at 1000 K when the pressure increases from 0 to 10 GPa (pressure scale characteristic of the helium nanobubble region of plasma-exposed tungsten). This substantial reduction in α has far-reaching implications for thermal stress generation under plasma exposure conditions. The generated thermal expansion coefficient database, augmented by forthcoming results for damaged tungsten due to plasma exposure, will be utilized as input to our coarse-grained models and continuum-scale codes to calculate thermal stresses and their contribution to the total stress in the material; these models are coupled with our PFC surface morphological evolution models.

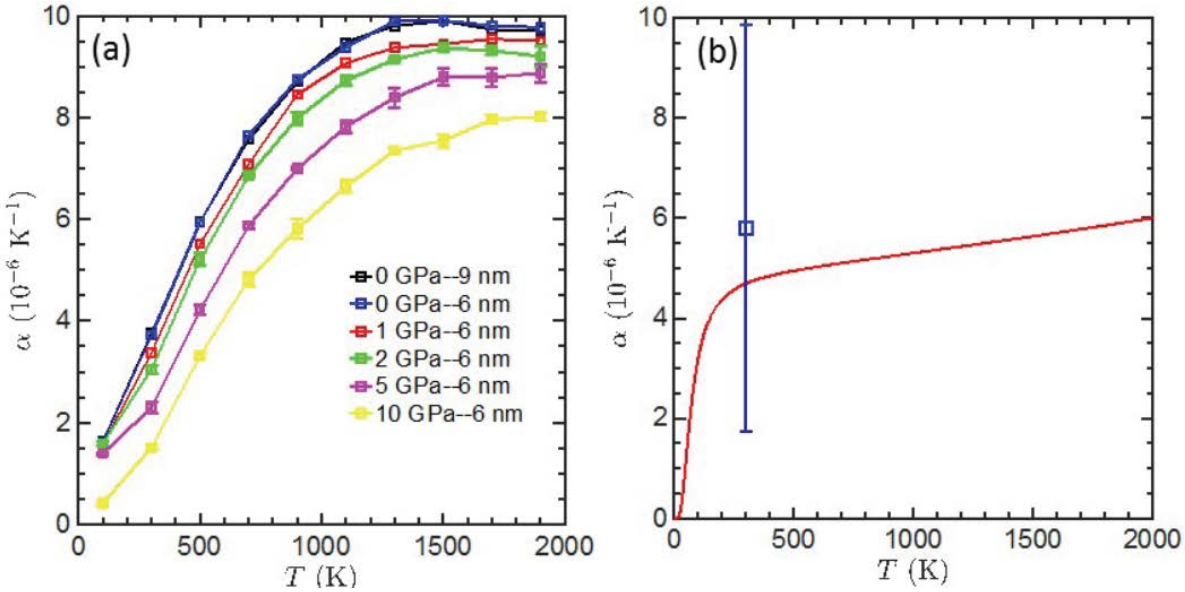

Figure 18. Predictions of thermal expansion coefficient, α, of defect-free W at temperatures up to 2000 K based on (a) (*NpT*) MD simulations according to the Juslin-Wirth potential for pressures up to 10 GPa and (b) DFPT-QHA calculations at zero pressure (red solid line). A preliminary AIMD prediction (blue open square) is included in (b).

To assess our classical MD predictions, we have also carried out first-principles computations of the thermal expansion coefficient of tungsten within density functional theory (DFT) as implemented in VASP. These computations involve free-energy minimization at given T with respect to the lattice parameter and lattice-dynamics calculations of phonon spectra within density functional perturbation theory (DFPT) for determining vibrational free energies; the free energy computations have been carried out within the quasi-harmonic approximation (QHA). Our DFPT-QHA results for $\alpha(T, p = 0)$ are shown in Fig. 18b, and are in excellent agreement with the latest published results for W [9]. Such results can be used to both evaluate our current classical MD predictions and contribute to our property database toward classical potential improvement and development. Our comparisons indicate that the Juslin-Wirth potential under-predicts α at T = 300 K by ~20%. Finally, we have implemented first-principles (*NpT*) MD (AIMD) computations of $\alpha(T,p)$ for direct *ab initio* predictions of the thermal expansion coefficient as the ultimate means for such thermodynamic property determination; a preliminary AIMD result at T = 300 K (longer MD trajectories for better sampling are required for higher accuracy and statistical significance) is included in Fig. 18b.

To investigate PFC structure-property relations DFT calculations and atomistic simulations using reliable interatomic potentials for W-based materials, along with object kinetic Monte Carlo (KMC) simulations to evolve spatially significant microstructures to experimentally relevant time scales are being performed.

To study the formation and evolution of the damage microstructure, a lattice-based object kinetic Monte Carlo tool, KSOME [10,11] (developed under SciDAC-3) is being enhanced to incorporate the ability to simulate diffusion of point defects and their clusters under the influence of long-range interactions with extended defects, particularly with free surfaces and grain boundaries. This enhancement is essential for improving the fidelity of microstructural evolution simulations of the near-surface region. In a recent study of near-surface helium bubble growth, higher helium retention was observed in KMC simulations using KSOME when compared to the retention observed in molecular dynamics simulations, as shown in Fig. 19 [12]. This difference was attributed to a decrease in the activation barrier for helium atom diffusion near the surface, as illustrated in Fig. 20 [13], which resulted in drift of helium bubbles towards

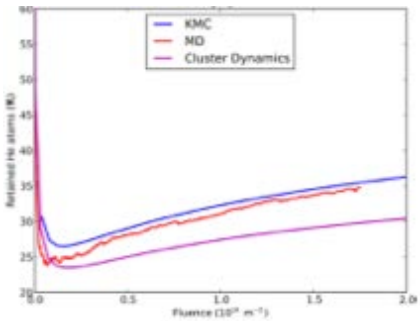

Figure 19. Comparison of helium retention obtained using KSOME (KMC), MD and Xolotl (Cluster Dynamics) [12].

the surface, thus lowering helium retention. The version KSOME used for that study had no provision to vary the helium atom diffusion activation barrier as a function of distance from the surface. In addition to an increase in the rate of drift as helium bubbles approach the surface, the trap mutation rate also increases. More importantly, the strength of the interaction of helium bubbles with free surfaces increases with bubble size as depicted in Fig. 20. Not including these interactions will significantly affect the results obtained from long-time simulations of near-surface microstructure evolution under both isothermal and variable temperature conditions.

To treat depth-dependent activation barriers additional subroutines were added to KSOME to give it the ability to simultaneously track distances of point defects from multiple extended microstructural features. With this improvement KSOME can now perform simulations with depth-dependent activation barriers, including the long-range interaction effects of grain boundaries and dislocations on point defect diffusion. For simplicity, these extended microstructural features are assumed to have regular geometric shapes. Furthermore, additional subroutines were added to parse mathematical functions related to point defect capture radii from a text-based input file. In the previous version of KSOME, capture radii were hardwired into the code, which required mathematical expressions to recalculate capture radii as the number of defect types increased. In addition, information regarding the characteristics of planar and line defects are now inputted via text-based files. In the future, these enhancements will make incorporation of long-range interactions between point and extended defects feasible with minimal code modification (e.g. elastic trapping of large dislocations loops). All necessary modifications to KSOME are complete and the new version is being benchmarked using the helium retention data obtained from MD simulations shown in Fig. 19 [12].

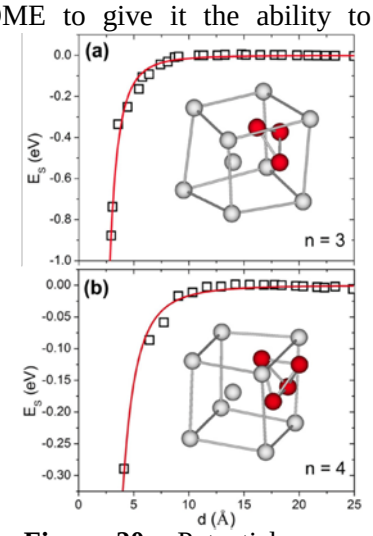

Figure 20. Potential energy profiles for the interaction of (a) He_3 and (b) He_4 clusters with W(211) surface [13].

Xolotl developments, including bubble bursting model/simulations and incorporation of He-H energetics for PSI: Xolotl solves the spatially dependent DDR equations in multiple dimensions to predict the evolution of the concentration fields of the various cluster species for a plasma exposed tungsten material. We use a finite difference implementation in which the grid along the depth direction has a variable step size, with a finer grid near the surface, in order to better encompass the profile of the incoming helium flux. In the interest of brevity, the model governing master equation and solver will not be discussed here. These details have been described in our earlier publication [13]. The clusters originally under consideration are tungsten self-interstitial (I), vacancies (V), helium (He), as well as vacancy-trapped helium (HeV) that we will also call a bubble. Single vacancy, interstitial clusters, and small helium clusters are mobile and diffuse isotropically in bulk tungsten. The surface is able to move as a result of tracking the number of interstitial clusters diffusing to the surface and adding a grid point when this number exceeds the tungsten atomic density. Additionally, a biased drift term toward the surface and trap mutation for small He clusters near the surface were modeled after MD simulations as described in Ref. [14].

In order to continue towards a more realistic model for Xolotl, we recently added a mechanism for bubble bursting. As observed in MD simulations, a bursting event occurs when a cluster of helium atoms is

trapped in vacancies of the tungsten material, grows close to the surface and ruptures, leading to the release of the helium atoms. The mechanism is implanted as follows in Xolotl: at the end of each time step, the total quantity of helium atoms is computed (both mobile and trapped) on each grid point, and the radius of the helium bubble is calculated, the ligament thickness is then computed from the distance to the grid point (depth) minus the radius of the bubble; the bubble is burst if the ligament thickness is zero or negative; otherwise, the bursting is considered a stochastic phenomenon, in which a probability proportional to the ligament thickness (L) and depth of the bubble (d), and including an adjustable factor (τ_d) determines the probability of bubble bursting, as:

$$P_{bursting} \propto \left(1 - \frac{L}{d}\right) \times f \times \min(1, e^{-(d-\tau_d)/2\tau_d}) \qquad (1).$$

The last term in the probability is introduced to take into account the unlikeliness of deep bubbles to burst, and is here chosen to be 10 nm. When a bubble bursts, the concentrations of all helium clusters are set to zero at this location (grid point), and the concentrations of HeV clusters are transferred to V clusters of the same size. This bursting behavior mimics a pinhole bursting, when a cavity releases the helium but gets sealed and is able to refill with helium atoms. The factor f was then selected based on benchmark comparisons to MD simulations. In order to ensure that the MD simulation database included sufficient number of bubble bursting events, we have compared the newly implemented bubble bursting model in Xolotl to an MD simulation at a high flux of 5.0×10^{27} He m^{-2} s^{-1}, which has accumulated an implantation fluence of 1.5×10^{21} He m^{-2}.

Figure 21 shows a comparison of the helium retention as a function of implantation fluence as predicted by Xolotl versus that observed in the MD simulation. Both the MD and Xolotl simulations use a high helium flux of 5.0×10^{27} He m^{-2} s^{-1} incident on a (100)-oriented tungsten surface at 933 K. The sharp drops in retention in the MD simulations (black line) correspond to bubble bursting events that release helium. Notably, the Xolotl simulation without bubble bursting grossly over-predicts helium retention, while the predictions with the incorporation of our bubble bursting model reproduce the sharp drops in retention. As expected, a higher value of f leads to a lower quantity of helium retained as well as more frequent drops in the distribution (bursting events). Based on the overall agreement, we conclude that a value of $f = 0.1$ best reproduces the MD results.

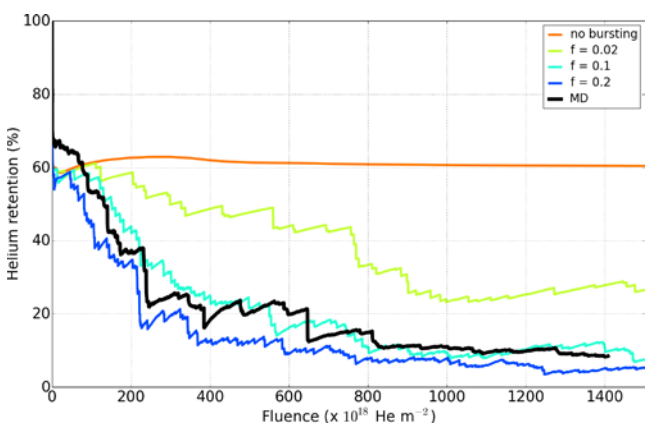

Figure 21. Percentage of helium retention as a function of helium fluence according to MD simulations and Xolotl predictions. Xolotl results were obtained with four different bursting probability factors. The incoming helium flux is 5.0×10^{27} He m^{-2} s^{-1} for a (100) oriented tungsten surface at 933 K.

Another recent addition to Xolotl is the extension of the cluster network to hydrogen. Atomistic modeling, based on a combination of density functional theory and molecular statics calculations using interatomic potentials have been used to populate a database describing the interactions of helium and hydrogen with vacancy clusters in tungsten [15-18]. The activation energy and pre-exponential factor for hydrogen diffusion in tungsten has been set to experimentally known values [19], and an inverse mass square root dependence will be used to scale the appropriate diffusivities of deuterium and tritium, respectively. The binding energy of hydrogen to helium-vacancy clusters will be extrapolated to

larger vacancy cluster sizes, and we plan to perform a select set of atomistic modeling calculations to assess the uncertainty of these extrapolations.

This addition of deuterium to our previously helium-only model greatly increased the total number of clusters we have to consider for the same maximum size of mixed cluster. Indeed, mixed clusters now have a 3-dimensional phase space where they had a 2-dimensional one before, and in order to model mixed clusters up to 10 vacancies (reaction radius of 0.317 nm) 7090 clusters are needed where 312 were sufficient for helium-only. As discussed in the CS/AM section describing the improvements to Xolotl performance, more clusters mean more memory usage so we are limited in the total number of clusters we can model. In parallel to our effort to reduce the memory usage purely and simply, we decided to implement an extension of the moment-based grouping method [20] in which we gather mixed clusters that are close in size and properties into a "super" cluster, and then only track the evolution of the super cluster. This method was already present in Xolotl for helium-vacancy clusters where each super cluster was represented by three equations (one for the zeroth moment, one for the first moment in the helium and vacancy directions) instead of variables for every individual cluster that each super cluster represents. In the helium-deuterium case, each super cluster is now represented by four equations (one additional for the first moment in the deuterium direction), which still reduces memory footprint and improves performance even for small groups of 2x2x2 mixed clusters compared to the non-grouped implementation.

Task 2. Integrating boundary plasma and PFC Surface models to predict PFC Performance

The vast majority of activities within Task 2 this year have focused on supporting the DOE Office of Fusion Energy Sciences FY2018 Performance Measure Management (PMM) reportable milestone in Fusion Theory and Simulation, which involves modeling Plasma Materials Interaction (PMI). 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, this milestone effort will involve performing 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.

Modeling geometry for ITER PMI simulations: The ITER geometry that we have identified for the modeling effort within the FY2018 milestone has been chosen to meet two principal 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 currently anticipate that 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.,

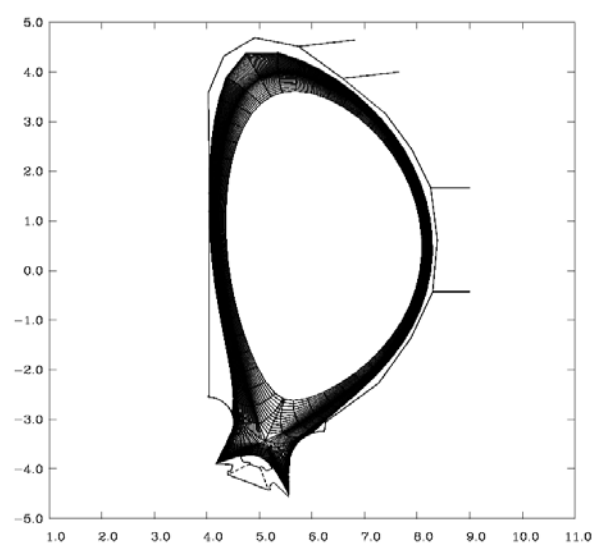

Figure 22. The poloidal cross section of ITER that is modeled in SOLPS [6], including the wall for EIRENE and the grid for B2.5 [7,8].

surface growth, erosion, variation in gas retention, etc).

The equilibrium plasma characteristics are calculated using the SOLPS package, which includes the neutral transport calculated by the kinetic Monte Carlo code EIRENE and plasma transport simulated by the 2-D fluid model B2.5 [6-8]. Given the toroidal symmetry assumed in the fluid description, only a poloidal cross section of the tokamak is modeled in SOLPS. The physics of neutral particles dominates the areas of low plasma temperature and thus the modeled geometry in EIRENE will encompass the entire first wall, including the main wall, all divertor surfaces and some level of its sub-structure, such as the entrance to the pumping ducts located under the dome in the private flux region. Figure 22 documents the ITER geometry in B2.5 [21], which involves a grid encompassing most of the edge plasma and aligned with the plasma flux surface, with 90 points in the poloidal direction and 36 flux surfaces in the radial direction. The flux surface shape does not conform to that of the wall, such that areas near the wall, which are dominated by neutral transport, are excluded from the fluid model geometry. The grid extends 10-20 cm into the core (at the mid-plane) and includes several scrape-off layer heat flux widths (λ_q) in the scrape-off layer (SOL). The last grid surface is located as close as 2 cm off the inner wall, and at a distance of 8 cm off the inner mid-plane, 9 cm off the outer mid-plane and 15-25 cm off the top of the machine. The B2.5 geometry also resolves the divertor in greater detail, extending the grid all the way to the wall in the inner and outer vertical targets, and 2 cm away from the dome.

The ITER divertor is divided into 54 essentially identical cassettes. Therefore, any given toroidal 'slice', which encompasses the entire width of a cassette should be representative of the physics throughout the entire toroidal symmetry. Thus, we plan to model the impurity migration and re-deposition in a single divertor cassette, assuming periodic boundary conditions, in anticipation that this approach will provide the most efficient, yet complete representation of the entire ITER divertor geometry.

The divertor cassette is represented in our global impurity transport code, GITR, by a 1/54th scale representation of a fully cylindrically symmetric model of ITER. The divertor tiles within a cassette incorporate 2816 units, with tile gap widths and depths specified as 0.5 mm and 2 mm, respectively, as shown in Figure 23. The mono block tile faces range in size from 24 to 44 mm in either of the “face” directions. These sizes were chosen to “fit” with the given gap width and toroidal extrusion. At a minimum, each tile is resolved in the geometry grid, resulting in 66,000 triangle faces for the geometry, which is an easily manageable geometry for GITR. The fidelity can be increased to isolate areas of special interest within our modeling. For instance, the area around the outer strike point, where the highest plasma ion fluxes, erosion and W prompt re-deposition are expected, can be modeled with higher spatial resolution to better capture variations in surface conditions, necessary to accurately model the evolution of the divertor tile target.

The fully integrated model, including plasma and impurity transport calculations, in addition to surface modeling of the ion implantation and gaseous species diffusional transport and clustering behavior, will be applied to resolve the surface and sub-surface evolution of a 3x3 divertor tile region, including the gaps between tiles as shown in Figure 24. The geometry will be centered (poloidally) in the outer strike point (OSP), where we expect net erosion. We expect the two tile locations adjacent to the OSP to be regions of net deposition. If our modeling predicts that these two outer arrays also experience net erosion, then we will extend the tile geometry in the poloidal direction (where the largest plasma variations take place) in order to include a region of net deposition in our PMI modeling effort. Alternatively, we could also consider the option to model an additional, similar tile domain (e.g., 2x2 or 3x3) in a separate location where net deposition occurs within the outer divertor. Initial impurity transport calculations in GITR will provide an estimate on areas of the outer divertor with net erosion and net deposition, and thus help us determine the need for, size and precise location of the possible additional area. We will perform at least one simulation of surface and sub-surface evolution per ITER divertor tile within the FY2018 milestone. Additional simulations of the tungsten tile divertor material evolution are possible if conditions

change significantly across a given tile (e.g., between center and near a gap, or between areas very close to the strike point and elsewhere poloidally in the same tile). This need will be determined by the sensitivity of the sub-surface evolution to local plasma conditions, studied by applying the integrated model to linear machine experiments.

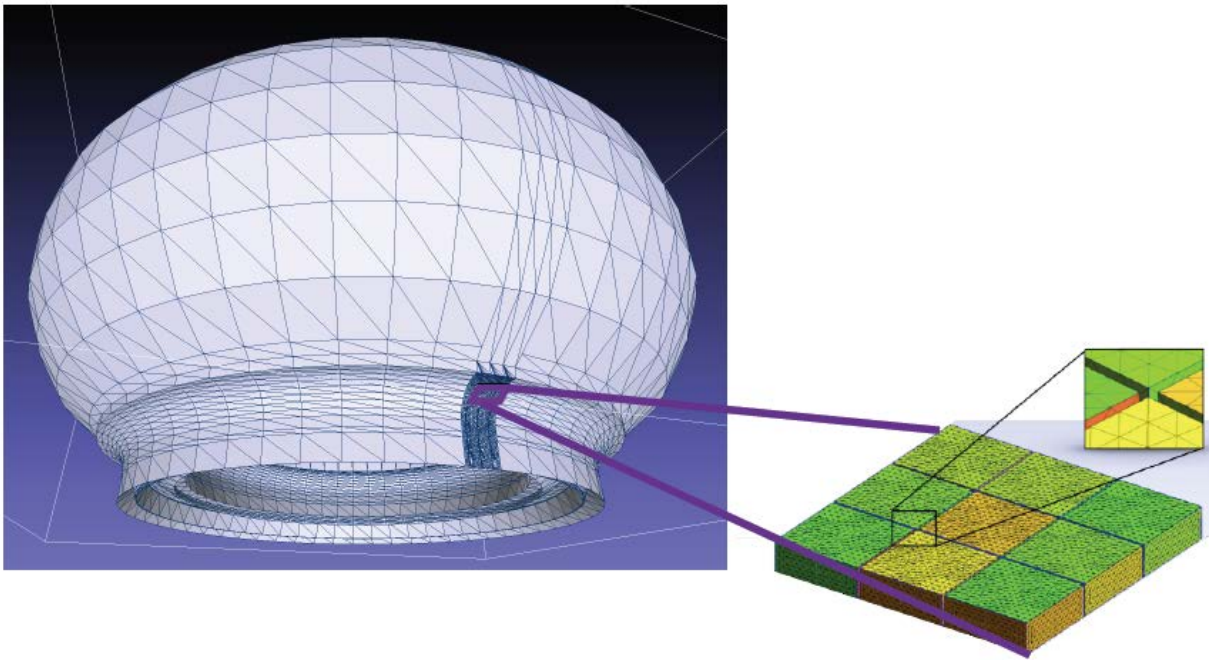

Figure 24. Representation of the complete ITER geometry to be modeled, with a magnified view of a three by three divertor tile region geometry including tile gaps.

Computational codes and Code coupling strategy: PMI 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 envision a code coupling workflow as illustrated in Figure 25, 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, 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 plan to 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 Integrated Plasma Simulator (IPS) [22] provides the framework and serves to integrate our models. In the remainder of this Section, we first describe the IPS framework for code coupling, and then describe the models necessary to model the complex, multiscale nature of the material surface response and the boundary plasma, respectively.

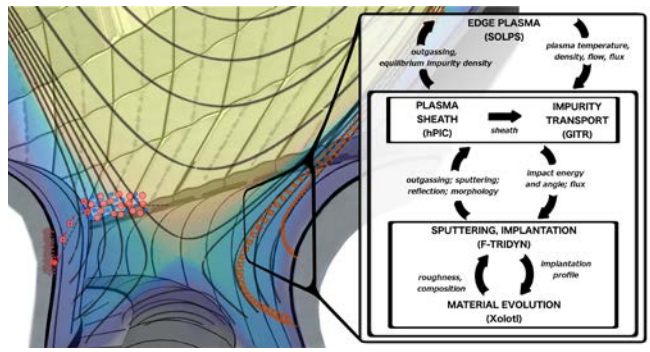

Figure 25. Schematic illustration of a tokamak divertor outlining the magnetic field line structure, plasma density (gradient), sputtered impurity trajectories (orange) and potential drop at the sheath for ions (red) and electrons (blue); next to the code-coupling workflow to model PFC surface evolution.

<u>Description of Coupled Code Modeling Predictions compared to Experimental Measurements – PISCES-A linear device with mixed He-D plasma</u>

Experimental setup and modeling parameters

Within the last year, a number of dedicated linear machine experiments have been performed in PISCES-A [23] to benchmark the integrated plasma surface interaction model. In the PISCES experiments, there is an array of experimental observables, which vary from experiment to experiment. In summation of all of the experiments, the 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 have all been experimentally measured. Common to all of the plasma exposed W experiments are the plasma column strikes the center of a circular W base plate which is biased to obtain 250 eV incident plasma ions and the temperatures is held between 400 and 800 degrees C in order to avoid the fuzz forming regime.

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. The variations in the experiments include W target radius and thickness, plasma species mixtures, background plasma temperature, density, and flux, exposure times, and experimental observations available for comparison to modeling. An overview of the measurements performed to date is provided in Table 1.

The largest difference between the He only and He/D mixed plasma experiments is the reduced He fraction (10%) in the plasma. The plasma parameters of electron temperature, electron density, and the applied bias voltage of the target (-250 V) are reasonably consistent between the He only and He/D mixed experiments, and thus result in similar energetic interactions of the background plasma ions with the W target (e.g., ion impact energies). However, in the mixed He-D plasmas, the D ion energy is effectively below the sputtering threshold (e.g., 250 eV He on W has sputtering yield ~3.9×10^{-3}, whereas 250 eV D on W has sputtering yield ~7×10^{-6}), thus only the He will cause any significant sputtering and correspondingly much less sputtering will occur in the mixed plasma experiments due to the lower He implantation flux.

Table 1. Summary of experimental conditions in dedicated PISCES-A experiments on tungsten targets for benchmarking the coupled code simulations.

	20170313	20170314	20180205	20180213
W Target Radius [cm]	4.875	4.875	1.27	1.27
Plasma Species	He	He	He	10% He / 90% D
Exposure Time [s]	10,000	5,000	100	1,000
Peak Electron Temperature [eV]	6	9	7	10
Peak Electron Density [m^{-3}]	0.8×10^{18}	5×10^{18}	6.4×10^{18}	1.1×10^{18}
Peak Ion Flux [$m^{-2}s^{-1}$]	0.5×10^{22}	4×10^{22}	5.4×10^{22}	1.5×10^{22}
Surface temperature [K]	673 - 873	673 - 873	1093	808
Experimental Observations	Target mass loss, titania tower mass gain, axial profile of W neutral emission spectroscopy		Ex-situ analysis using laser-based spectroscopy and surface roughness measurement	

The more recent 2018 measurements used one inch diameter W discs, in order to provide samples for ex-situ measurements of the surface roughness and the sub-surface gas distributions. Measurements in the present experiments, such as surface temperature and plasma gas composition, focus on ensuring an accurate analysis of the samples and modeling of the sub-surface evolution. The experimental set-up is similar to that in the impurity-migration experiments performed in 2017. However, the 1 inch disc is mounted onto a target holder, which does not provide the tower to measure tungsten re-deposition that was used in the earlier experiments. Plasma characteristics (density, temperature, etc.) in the 2018 experiments were only measured in the center of the plasma beam. Thus, we have assumed a similar

radial profile of the plasma parameters for these measurements as in the previous experiments, and we have scaled the plasma parameters by the peak values.

These experiments have been modeled by integrating the impurity transport code GITR, the binary collision approximation code F-TRIDYN and the continuum based cluster dynamics code Xolotl. The code coupling was performed within the IPS framework, as noted in the previous section and described in the quarter one milestone report. Langmuir probe data was used to obtain the radial profiles of n_e, T_e and particle flux for the initial 2017 experimental measurements, and the background plasma profiles have been modified for the 2018 experiments based on the electron densities.

F-TRIDYN generates a reduced model (in the phase-space of impact energy, angle and surface roughness) for W sputtering yields by D and He, which is input to GITR. We assume an initially smooth surface (no roughness) and perform 10^5 computational ion impacts for each combination of projectile species, energy and angle. The angles vary from 0 to 89.9, every 10 degrees, and the impact energies range from 1 to 10^3 eV, with 50 logarithmically spaced energy binds.

GITR then evaluates the W impurity transport. We use a sampling of 10^5 computational particles, which are followed for 10^5 time steps, with a time step of 10 ns. Subsequently, as noted in Figure 1, F-TRIDYN and Xolotl are coupled, initially every 0.02 s. Given that the implantation profiles and surface composition approach a saturation, the coupling time-step is doubled every two coupled iterations. F-TRIDYN calculates sputtering yields and implantation profiles for D, He and the sputtered and re-implanted W. In these simulations, we assume that the gas species (He or D) impact with a kinetic energy of 250 eV with normal incidence, while the description of the re-depositing W ions is provided by GITR at the center of the target.

Xolotl is used to calculate the evolution of implanted species and surface morphology. The ion flux is provided by GITR, including the fraction of each species (D, He, W) in the plasma (at the center of the 1inch W target), for scaling implantation profiles and sputtering yields (calculated by F-TRIDYN in the second sub-component). We include a wide range of reactions that may take place between the network species (clusters, vacancies, etc), such as reaction, advection, modified trap-mutation, diffusion, moving surface, attenuation and bursting. The surface temperature is set to the measured value of 808 K in the 2018 experiment with the mixed 90%D-10%He plasma. The Xolotl spatial domain is a slab with a thickness of 5 μm, with a non-uniform mesh that has a much finer grid spacing near the surface. For the 10 nm closest to the surface, the bursting of high-pressure bubbles is modeled with a proportional to the depth of the bubble below the surface. The initial vacancy concentration was set to 10^{-9} m^{-3}.

<u>Benchmarking the tungsten erosion and re-deposition behavior (gross versus net erosion) in PISCES-A experiments</u>

He only exposure

A He only exposure of the one inch diameter W disc was performed at PISCES-A in experiment 20180205, as noted in Table 1. This experiment has background plasma conditions similar to a prior experiment, number 20170314, with comparable background temperatures, densities, and ion flux. In the 20170314 experiment, the experimentally measured mass loss value of the W target was 79.533 mg, and that mass loss occurred over a 5,000 second exposure. This leads to an average mass loss rate of approximately .0159 mg/sec. GITR predicted both the net erosion and the deposition of sputtered tungsten or this experiment, and the predicted values were in good agreement with the experimental measurements.

GITR simulations of the experiment 20180205 predict gross tungsten erosion of 2.704 mg with 1.201 mg being re-deposited on the surface. Thus, the GITR predictions are for a net erosion and mass loss of 1.503 mg. Since this experiment only involved a 100 second exposure, it is relevant to compare the average rate of mass loss. GITR predicts a tungsten mass loss rate of 0.015 mg/sec per second shows consistent results with the previously measured and modeled experimental results. Simulation parameters which can affect the gross erosion and deposition of W include the ion flux to the target as a scaling factor for the erosion from the background plasma, and background electron temperatures and densities which affect ionization rates which affect prompt re-deposition rates. Comparing the percentage of re-deposited impurities, the GITR simulation of the 2017 experiments indicated that 65%-81% of the eroded impurities were re-deposited on the target. However, the current GITR simulations of the more recent experiments indicate a significantly lower re-deposition percentage of 44%, which likely results from using an updated surface model in addition to the geometrical effects of the smaller specimen diameter.

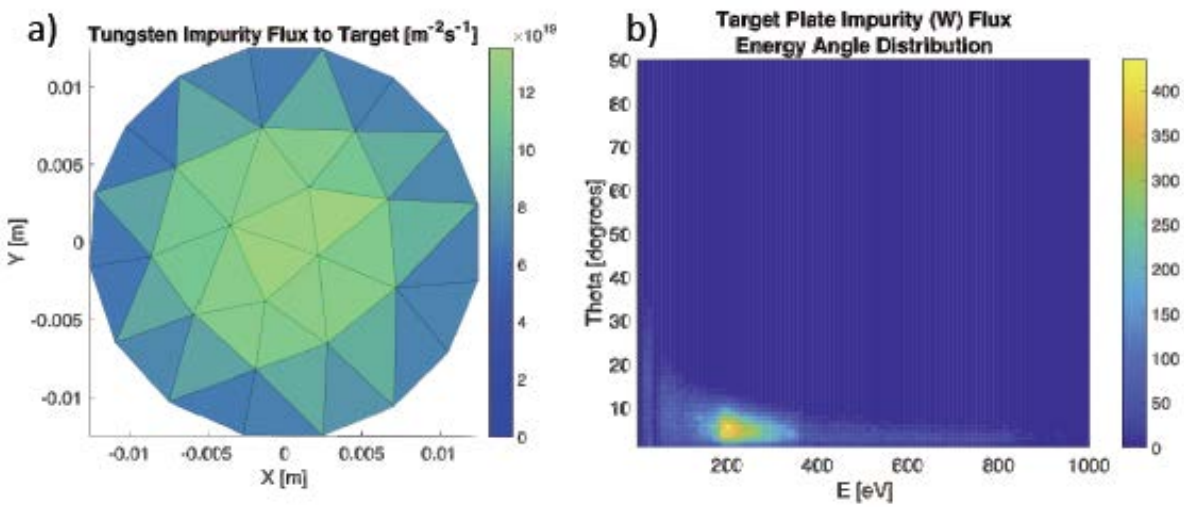

Figure 26. a) Characterization of W ion flux returning to the tungsten surface for the pure He plasma, and b) the energy and angular distribution of the re-deposited tungsten ions.

Figure 26 shows the GITR results of the tungsten re-deposition flux for the helium only plasma experiment that are passed to the surface evolution iteration simulation performed by F-TRIDYN and Xolotl. The W impurity flux at the target surface is shown in Fig. 26a, and shows that the flux is approximately 1.4×10^{19} $m^{-2}s^{-1}$ at the center of the target and then decreases with radial position. Examining the kinetics of the W impurity flux, it is clear that there is a peak in energy at 250 eV (singly charged W) near normal incidence, as shown in Fig. 26b. A tail of higher energy particles is also predicted, which indicates ionization of W to higher states (W^{+2} and W^{+3}) because of the relatively high electron temperatures and densities in this experiment. The continuous nature of this energy tail indicates that many of these atoms undergo ionization while still within the sheath electric field region near the target. Finally the peak in angle is slightly away from normal incidence, which can be attributed to the lower energy impacts that are a result of prompt ionization and deposition. Low energy ions that impact the surface are the main contributor to the distribution that does not have normal incidence.

Mixed D-He exposure

Another recent experiment, denoted as number 20180213 experiment involved mixed deuterium – helium plasma conditions for which there is a decrease in the helium ion flux to the tungsten surface, which significantly reduces the tungsten erosion. As well, the ion flux and electron densities are lower and are

similar to values in the 20170313 experiment. The low flux exposure 20170313 had a 10,000 second duration and the experimentally measured mass loss was 18.83 mg, resulting in a mass loss rate of 1.883 μg/sec. The GITR simulations of the D-He mixed plasma experiment predict a tungsten gross erosion of 0.757 mg of W with the re-deposition of 0.237 mg, resulting in a net mass loss of 0.520 mg and a corresponding mass loss rate of 0.520 μg/sec. The ratio of the He/D exposure target mass loss rate (0.520 μg/sec) to the low flux He 20170313 exposure (1.883 μg/sec) is 0.276 which is approximately the same as the ratio of He flux to the target in each experiment (1.5×10^{21} versus 5×10^{21} $m^{-2}s^{-1}$). The fraction of eroded impurities that GITR predicts to re-deposit on the target surface for the D-He mixed plasma is 31.3% which is less than the approximately 32-67% predicted for the He only exposure due to the surface model and geometric considerations mentioned previously.

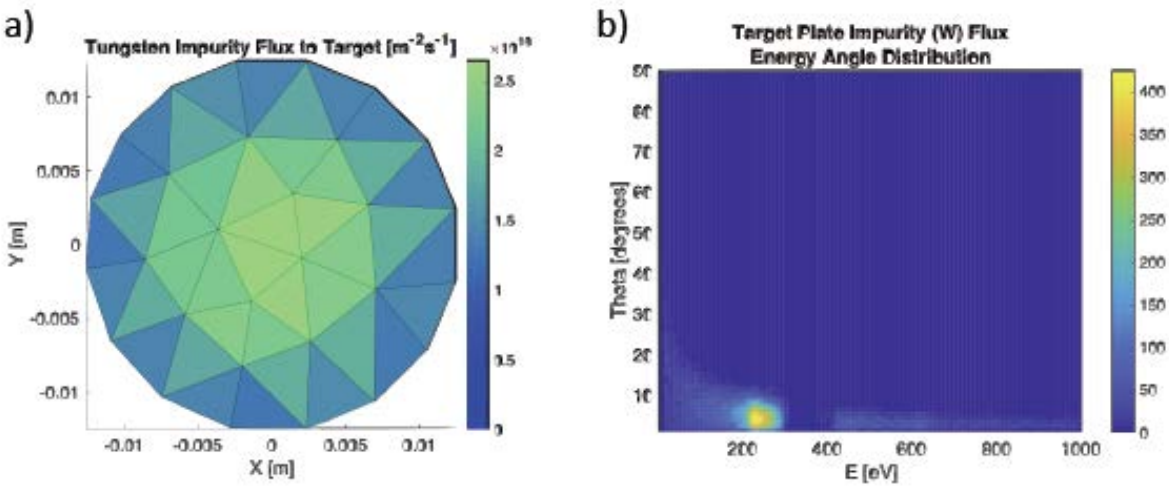

Figure 27. a) Characterization of W ion flux returning to the tungsten surface for the mixed D-He plasma, and b) the energy and angular distribution of the re-deposited tungsten ions.

Figure 27 plots the GITR output of the tungsten re-deposition flux to the tungsten surface in the mixed D-He plasma experiment that is input to the coupled F-TRIDYN Xolotl simulation. Fig. 27a indicates that the peak tungsten flux is 2.6×10^{18} $m^{-2}s^{-1}$, which is consistent with the reduced erosion yield and almost an order of magnitude below the helium only plasma conditions shown in Fig. 26a. Fig. 27b shows the angular and energy distribution of the re-deposited tungsten ion flux, which again shows a peak at an impact energy of 250 eV, as well as a smaller peak around 500eV. A gap in deposition energy is shown between 300 and 400 eV, which is attributed to reduced ionization rates of singly charged W near the surface because of lower electron temperatures and densities.

Predictions of deuterium (and He) recycling during the PISCES-A experiments

Figure 28 shows the deuterium and helium implantation profiles predicted by F-TRIDYN in our coupled simulations for the mixed D-He plasma experiments. Comparison of Fig. 28a and 28b indicates that the D is implanted much deeper than He as expected based on the mass difference. The D implantation profile peaks at ~5nm and extends to ~20 nm, as compared to the peak implantation depth of helium at ~3 nm with a tail of the implanted He distribution profile to a depth of about 12 nm. However, it is important to point out that this subtle difference in implantation depth does not correlate to an expectation of higher deuterium retention (relative to He). This is because D does not self-trap in W in the manner that He will self-cluster to limit the repulsive He – metal interaction potential. Thus, D implanted beyond the He range will diffuse either into the bulk or back towards the surface, where it can trap with the developing helium

cluster population or desorb. Notably, the implantation profiles shown in Figure 28 do not indicate any significant variation in the implantation depths for the relatively short simulated times of 20 seconds. We attribute this to the relatively low implantation flux, which results in predictions of small values of retained gas (discussed in the following section), as well as the relatively short accumulated simulation time to date.

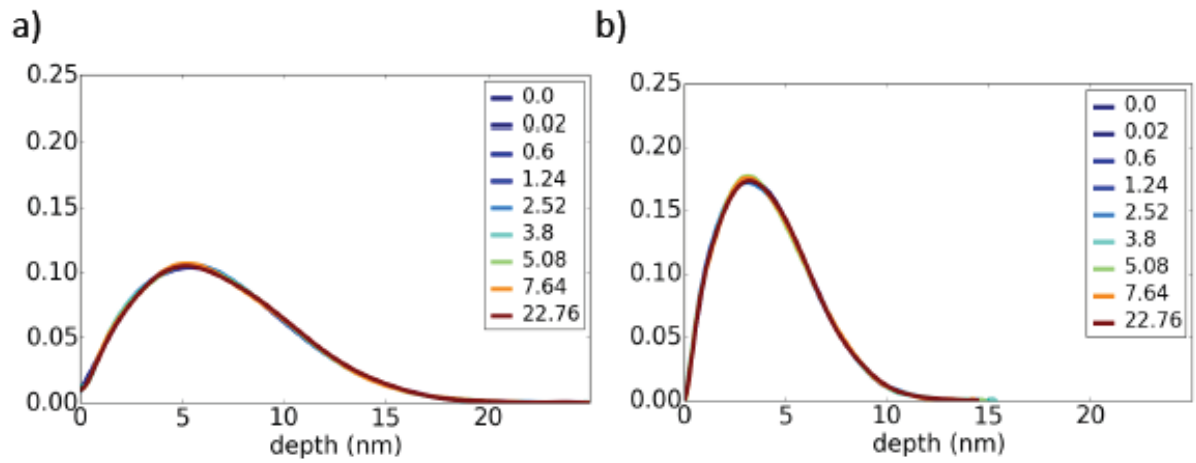

Figure 28. Evolution with time (values shown in legend have units of seconds) of the a) deuterium and b) helium implantation profile.

Gas recycling from a plasma-exposed material can be separated into two main processes, namely, prompt recycling, also referred to as reflection, and longer-term recycling which results from the diffusion back to the surface of the implanted gas and its subsequent desorption, which is also referred to simply as recycling. The former represents the fraction of ions that are not implanted, but rather are reflected from the surface. In our calculations, the prompt recycling (reflection) of deuterium has a value of 55%, while helium reflection has a value of 53%.

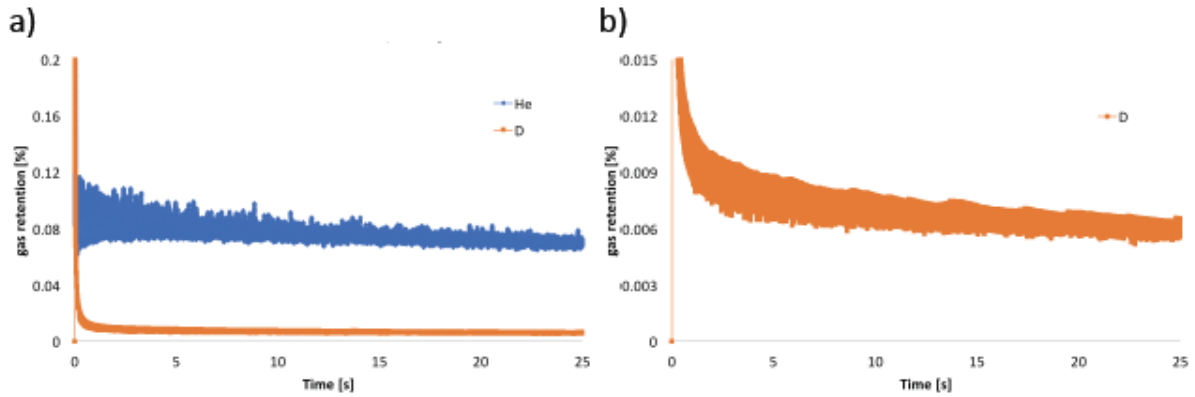

Figure 29. a) Percentage of gas retention of helium (blue) and deuterium (orange) as a function of time predicted for the mixed D-He plasma exposure experiment 20180213. **b)** Magnified view of the deuterium retention.

The longer term recycling can result from different processes, such as gas release that results in an instantaneous fashion as an over-pressurized gas bubble bursts and releases gas, as well as the diffusion of

implanted gas to the surface and the subsequent recombination of deuterium atoms to deuterium molecules that subsequently desorb (e.g., surface D+D → D_2). Often, the quantity of retained gas in the solid is analyzed, which is essentially the reverse metric to retention. Figure 29 plots the gas retention predicted by our coupled PSI simulation of the mixed D-He plasma exposure in the 20180213 experiment. Retention is defined simply as the fraction of gas atoms that are predicted to be within the sample at a given time relative to the total number of implanted ions. By this definition, the retention is initially 100% (at the time of the first gas implantation), but then decreases until clusters are produced that trap gas. After this nucleation event, the retained gas fraction generally trends towards a saturation value with fluctuating values that result from the repeated gas bubble bursting and re-filling with gas atoms.

Fig. 29a indicates that the saturation retention value for helium is about 0.08%, while it is much lower for the deuterium (as expected). Fig. 29b shows a magnified view of the deuterium retention, and reveals both a fairly significant fluctuation about the saturation value due to the bubble bursting events, as well the lower level of retention around 0.007%. The implication of this calculation of the deuterium retention is that the recycling fraction of deuterium is predicted to be 99.993%.

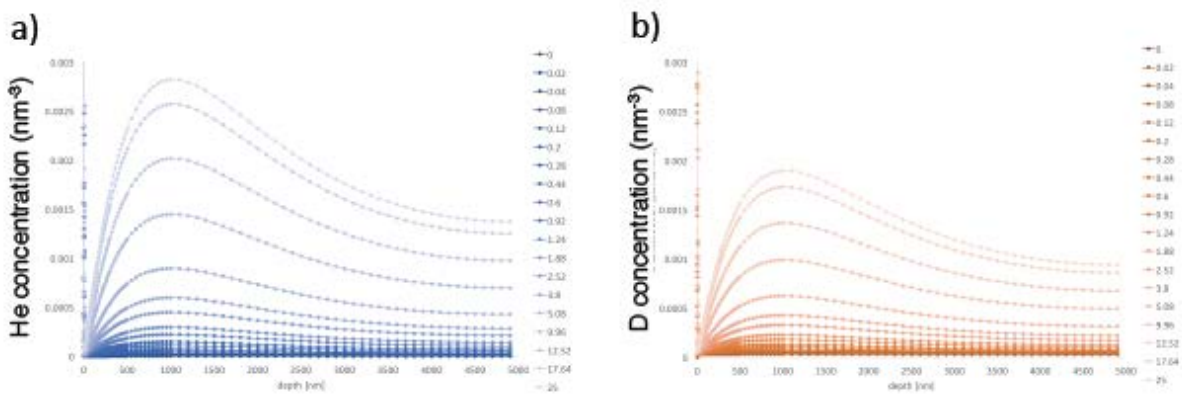

Figure 30. a) Concentration of helium as a function of depth and time during exposure to a 250 eV mixed D-He plasma in PISCES-A, **b)** concentration of deuterium as a function of depth and time (legend shows the time in seconds).

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 [24]. 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. We do not expect the deuterium to self-cluster, but do expect that it will partition to and bind with both vacancies and helium-vacancy clusters. Figure 30 shows the predicted concentration profiles of helium and deuterium as a function of depth below the surface and time. A comparison of Fig. 30a and 30b indicates that the helium and deuterium concentration profiles are similar, presumably because the trapping sites for deuterium, like helium, are associated with the helium self-clustering process. However, it is also clear that the helium concentrations are slightly larger than the deuterium concentrations, which is consistent with the lower deuterium retention shown in Fig. 29. At the experimental exposure temperature of approximately 800 K, the helium has a slightly higher diffusion coefficient than deuterium (4.4×10^{-9} m^2/s versus 1.2×10^{-9} m^2/s for He versus D, respectively), which are consistent with the higher helium concentrations at deeper depths below the surface.

Figure 31 plots the integrated concentration data of helium and deuterium within 60 nm layers. To date, the results presented for the coupled simulations (Figures 29-31) have not yet reached sufficiently long

times for comparison to experimental techniques, although they clearly indicate a difference in deuterium behavior for the mixed D-He plasma exposure relative to pure D. However, because of the limited simulation time to date, these results should be considered as preliminary predictions. These results do, however, highlight the opportunity presented by coupling the plasma physics and surface response to better understand plasma material interactions and to validate these predictions against experimental observations.

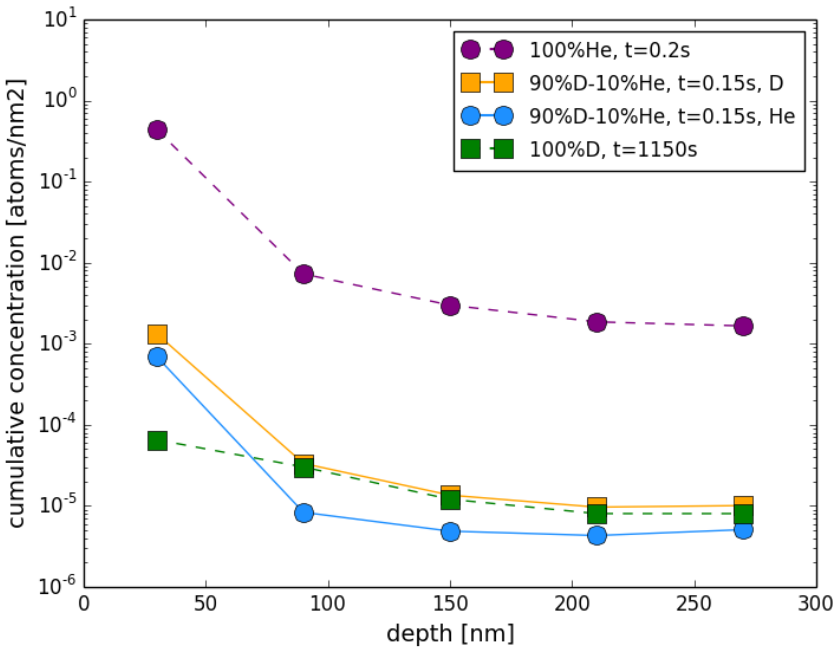

Figure 31. a) Xolotl predictions of the depth integrated concentration of helium (purple or blue circles) and deuterium (orange or green squares) versus depth, comparing the behavior in pure He (purple) or D (green) plasma exposure to that in a mixed D-He plasma exposure (orange and blue). Although the simulation times are different, the model predicts that the deuterium concentration near the surface is clearly increased in the mixed He-D exposure condition relative to pure D.

Documenting initial background plasma conditions for ITER with helium plasma operation

The background scrape-off layer and divertor plasma has been calculated for ITER using the SOLPS code, 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 disperse 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.

The first SOLPS case to serve as the basis for the PMI simulations is the helium plasma scenario. In the SOLPS runs described here, 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 $\chi_e=\chi_i=1.0$ m^2/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 ~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~1.5×10^{19} m^{-3} as shown in Figure 32.

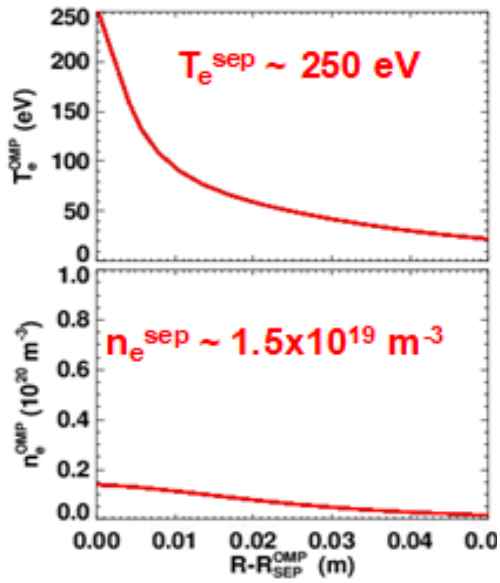

Figure 32. Radial profiles of electron temperature (top) and density (bottom) as calculated by SOLPS for a helium plasma operation in ITER.

Strong radiation is predicted for the helium scenarios under these conditions. SOLPS calculates that 25 MW out of the 40 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). As a result the parallel heat flux near the strike point is significantly reduced at the target compared to the midplane (Figure 33). The resulting heat fluxes on the target are quite manageable, with ~2 MW/m^2 or less carried by the plasma onto the target plates.

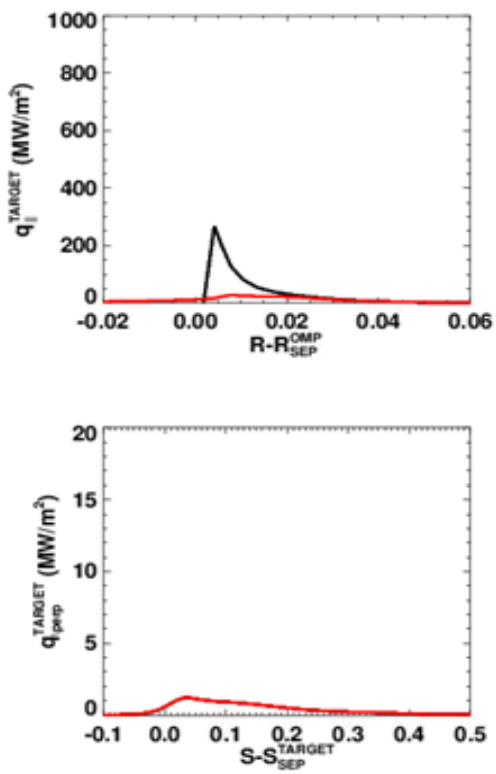

Figure 33. Radial profiles of parallel (top) and perpendicular (bottom) heat flux for the ITER helium plasma conditions.

The radial profiles of the divertor electron density and temperature and ion flux are shown in Figure 34. These show the profile shapes characteristics of the partially detached divertor state, with very low temperature (~eV) near the strike point indicated a local detached plasma, and radially increasing temperatures into the ~20 eV range farther into the SOL where the plasma remains attached. The electron density is highest near the strike point, but reaches relatively modest values of ~4×10^{20} m^{-3} in this low power scenario. Likewise the ion flux is high near the strike point, with a peak value of ~2×10^{23} $m^{-2}s^{-1}$. The density and ion flux both decrease significantly farther into the SOL. The presence of the sharp changes in electron temperature and density in this radial region is expected to lead to strong changes in the erosion and re-deposition behavior along the divertor target.

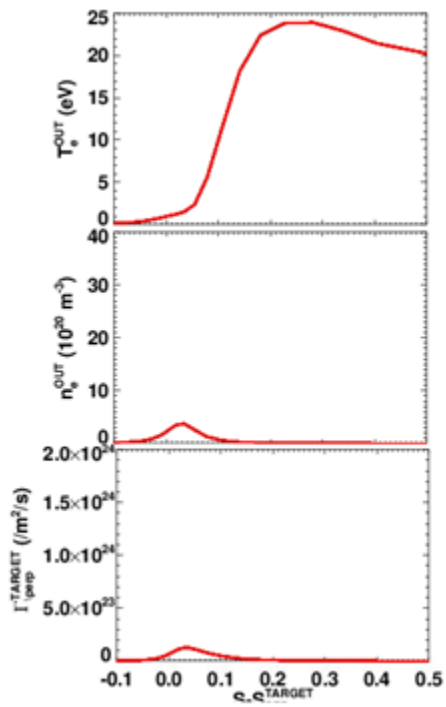

Figure 34. Radial profile of SOLPS-calculated electron temperature (top), electron density (middle), and c) helium ion flux (bottom) for representative operating conditions in ITER with a helium plasma.

These SOLPS simulations provide the background plasma conditions that will be used in the coupled PMI code modeling of the tungsten divertor response to helium plasma operation. Those results, along with an update on the experimental benchmarking of the linear plasma device predictions will be presented in the next quarterly milestone report, along with the background plasma conditions for burning plasma operation in ITER.

Assessment of the role of collisional drag in the magnetic pre-sheath

A first year activity at UIUC has been on the assessment of the role of collisional drag in the magnetic pre-sheath, by means of both detailed Particle-in-Cell simulations using the hPIC code, and validation of the simulation results against 3D tomographic Laser-Induced-Fluorescence measurements performed at the HELIX facility at WVU. Laser-Induced-Fluorescence (LIF) measurements can resolve the velocity distribution functions, temperatures, and drift velocities of plasma particles directly via fluorescence and Doppler effects and thus offer a methodology for direct experimental validation of kinetic plasma codes such as hPIC. In the experiments performed at the HELIX facility at West Virginia University, tomographic LIF techniques were applied to measure the ion velocity distribution functions of an Argon plasma along three independent spatial coordinates (x,y,z coordinates with respect to the target surface), and reconstruct the three-dimensional features of the ion flow during the acceleration toward a target plate and the impact of the plasma on the surface. The experiments have produced a large database of Ion Velocity Distribution Functions, which can be used for direct comparison against PIC data.

Particle-in-cell simulations of the HELIX experiment were carried out using hPIC in 1D3V mode (1 spatial dimension and 3 velocity coordinates). Plasma properties were resolved along the x-axis in a 2.5 cm domain perpendicular to the target plate. Simulations were performed to evaluate both collisionless and collisional conditions. Collision physics were handled using Monte-Carlo collisions and argon cross-sections. Using models of the collector optics and the finite laser beam width, simulated profiles were created for direct comparison to the LIF measurements. A fluid 1D3V plasma model was also run under analogous conditions. Figure 35 shows the comparison of particle-in-cell and fluid model simulation results to the LIF measurements at various conditions. Collisionless simulations overestimate the ion drift velocity along all three directions. Once collisional effects are taken into account, the ion drift velocity is in excellent agreement with experimental measurements. Near the Debye sheath, LIF measurements are affected by spurious reflections, resulting in subsonic drift velocities measured at the Debye sheath entrance that are 20% slower than those predicted by the fluid and kinetic models. Furthermore, Particle-in-Cell simulations reproduce the experimental data better than the fluid model. Quantitative agreement has been met for the electrostatic potential, particle fluxes, and for the x and z components of ion drift velocity. Disagreement in the y directions has been observed but may be attributed to the presence of either a neutral flow or a small longitudinal (field aligned) electric field. Careful analysis of these results suggests that the presence of a longitudinal electric field in the range of 150 to 200 V/m is the most likely explanation. Subsequent measurements using two-photon laser-absorption spectroscopy have revealed a

lack of neutral flow along the longitudinal direction. This effort offers not only quantitative validation of plasma codes, but also provides a comprehensive strategy for doing so.

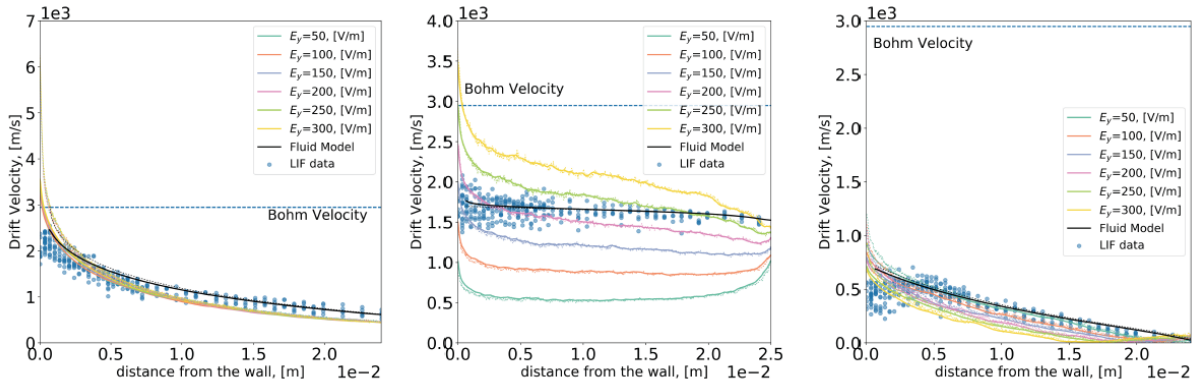

Figure 35. Ion drift velocities along three independent spatial coordinates within a magnetic presheath, obtained from: the hPIC particle-in-cell code, a simplified fluid model, and Laser-Induced-Fluorescence measurements at the HELIX facility. These plots show the effect of a presheath electric field along the longitudinal (B-parallel) direction on the 3D sheath structure.

Task 3. Dynamic wall response & feedback to divertor plasma

The goal of this thrust is to develop fundamental understanding of the temporal dynamics of the interaction between plasma and material surfaces at the edge of magnetically confined fusion energy devices. Do these interactions cause new types of coupled plasma-wall oscillations and instabilities? Will plasma-wall interactions change the character of turbulence near material surfaces? The objectives of the research program, co-led by Ilon Joseph (LLNL) and Sergei Krashenninikov (UCSD), are to develop predictive computational models of plasma-material interactions (PMI) for two of the most important outstanding issues:

1. material surface erosion due to transient power and particle loads, and,
2. dynamic recycling of main ions and impurity species between the plasma and wall.

In order to develop predictive capability, high-fidelity models for both the edge plasma and material PFCs must be coupled together. The goal is to perform the first studies of dynamic recycling and material erosion caused by transient events that include kinetic effects from the SOL plasma to the sheath to the material surface.

In the first year of the project, the main tasks are to develop divertor-relevant kinetic and fluid plasma computational models and utilize them to dynamically predict the flux of particles and energy impinging on plasma facing components (PFCs). We are simultaneously working to improve the physics modeling and computational capabilities of the wall codes in order to become technically ready for plasma-wall coupling. Finally, we are developing strategies for implementing the coupling of plasma and wall models by exploring the available numerical methods and by analyzing of the physics of the coupled models. At first we are focusing on explicit coupling strategies and an initial attempt at coupling has demonstrated the feasibility of this technique.

Summary of Task 3 Accomplishments (Sept. 2017-May 2018):

- Began COGENT studies of ELMs and sheath physics; implemented benchmark ELM case; implemented sheath boundary conditions for the gyro-Poisson equation

- Began development of BOUT++ divertor-relevant turbulence model including resistive-ballooning physics and linear sheath boundary conditions in mapped slab geometry; performed verification study of linear stability
- Formulated initial strategies for explicit coupling of plasma and wall codes; performed first wall-plasma coupling with UEDGE and FACE; developed strategy for coupling SOLPS and FACE

Roadmap and plans for dynamic recycling plasma-wall model: For plasma turbulence in the divertor region there are multiple instability drives. The two most important modes are the drift-resistive-ballooning mode (DRBM) and the sheath-driven conducting wall mode (CWM) instability. A flute-reduced turbulence model that captures the physics of these modes has been implemented within BOUT++. The model includes dynamic evolution of vorticity, ion density, electron temperature and includes linear sheath boundary conditions. The model has been verified by comparing the computed growth rates for these modes to an analytic model of the linear dispersion relations. The remainder of the year will focus on completing linear and nonlinear verification tests of these models. We will also add a model for neutral species that includes the ionization source so that the plasma model will respond dynamically to the recycling in the wall model.

The drift-resistive-ballooning mode (DRBM) is destabilized both by resistive dissipation and by regions of bad magnetic curvature. The simplest DRBM model that captures this physics evolves the plasma density, N_i, and the potential vorticity, ϖ:

$$\frac{\partial N_i}{\partial t} = -V_E \bullet \nabla N_{i0} + \nabla_{\parallel}(j_{\parallel}/e)$$

$$\frac{\partial \varpi}{\partial t} = 2\omega_{ci} b_0 \times \kappa \bullet \nabla P + N_{i0} Z_i e \frac{4\pi V_A^2}{c^2} \nabla_{\parallel} j_{\parallel}$$

$$j_{\parallel} = \frac{eN_{i0}}{0.51\nu_{ei}}\left(-\frac{e}{m_e}\partial_{\parallel}\phi + \frac{T_{e0}}{N_{i0}m_e}\partial_{\parallel}N_i\right)$$

$$\varpi = N_{i0}\nabla_{\perp}^2\phi$$

$$P = T_{e0}N_i$$

The plasma is assumed to be quasi-neutral so that the ion and electron densities are approximately equal. The equations are closed through Ohm's law, which determines the electric current, $j_{\parallel}$, parallel to the magnetic field lines. The magnetic field direction is denoted b_0 and the field line curvature is denoted κ. The ion cyclotron frequency is ω_{ci}, and the ions are assumed to be cold so that the total pressure, P, is the given by the density times the electron temperature T_{e0}. Figure 36 shows a verification of the linear dispersion relation against the analytic solution.

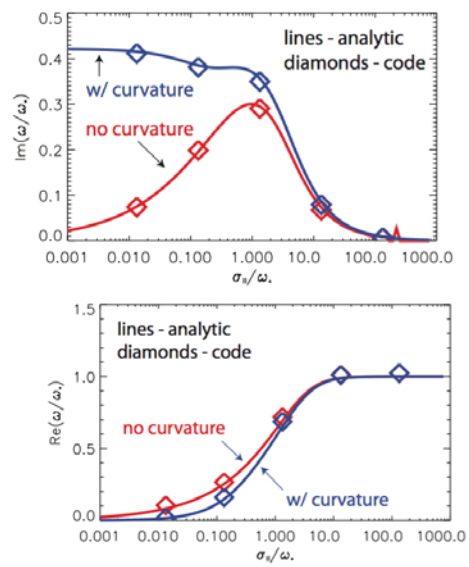

Figure 36. The computed growth rate and real frequency (diamonds) of the linear drift-resistive ballooning model (DRBM) implemented within BOUT++ verified against analytic theory (solid lines). Results are shown with (blue) and without (red) the curvature.

The conducting wall mode (CWM) is one of the most important instabilities for the divertor plasma because it is in contact with the electric sheath above the target plates and because the growth rate of this mode can be quite large. The simplest model that captures the physics of the CWM evolves the electron temperature, T_e, and the potential vorticity, ϖ:

$$\frac{\partial T_e}{\partial t} = -V_E \bullet \nabla T_{e0}$$

$$\frac{\partial \varpi}{\partial t} = -V_E \bullet \nabla \varpi + N_{i0} Z_i e \frac{4\pi V_A^2}{c^2} \nabla_{\parallel} j_{\parallel}$$

$$j_{\parallel} = \frac{eN_{i0}}{0.51\nu_{ei}} \left(-\frac{e}{m_e} \partial_{\parallel} \phi \right)$$

$$\varpi = N_{i0} \nabla_{\perp}^2 \phi$$

$$D \nabla_{\parallel} \phi = \pm \left(\Lambda_1 \phi - \Lambda_1 T_e \right)$$

$$D = \sigma \parallel T_0^2 / N_0$$

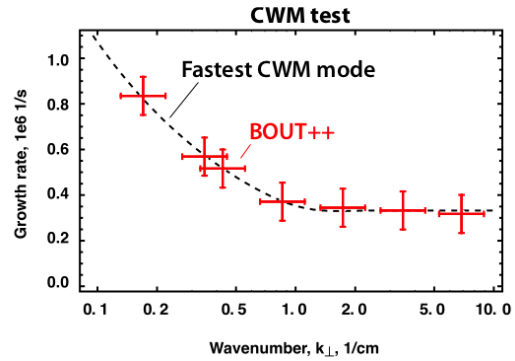

Figure 37. Computed growth rate for the BOUT++ conducting wall mode model compared to analytic theory.

In this case, the radial gradient of T_e drives the instability through the sheath boundary conditions represented by the constant Λ_1 and through the background electric field $eE_0=\Lambda_1 dT_{e,0}/dr$. While the temperature equation is passively advected by the ExB flow, it provides feedback to the instability through the sheath boundary conditions. This is because the electric current that passes through the target plate depends on sensitively on the difference between the electric potential and the temperature on the plate.

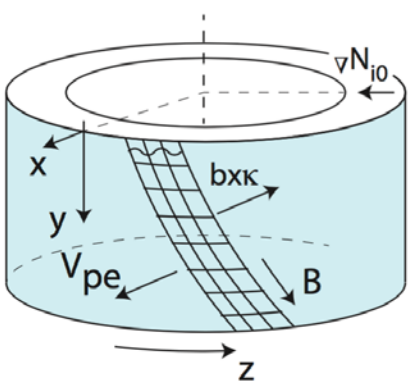

Figure 38. The slab geometry is topologically equivalent to a cylinder (torus). The directions of the magnetic field B, the density gradient N_{i0}', the curvature drift bxκ, and the electron diamagnetic flow V_{pe} are shown.

The initial stages of the project will focus on slab geometry, which is the geometry that is most appropriate for coupling to the wall models. This simplifies the potential numerical issues that can arise while capturing the majority of the relevant physics. As shown in Fig. 38, the slab geometry retains magnetic curvature and field line pitch.

We have developed an efficient conformal mapping approach to creating grids that add realistic shaping to a slab-like domain. A wide range of geometric shapes can be produced with an appropriate choice of an analytic function w(z) of complex coordinates z=x+iy. In principle, an arbitrarily detailed geometry can be generated using the Schwartz-Christoffel mapping of the upper half complex plane to an arbitrary polygon. An example of using the conformal mapping technique to generate a "divertor leg" is shown in Figure 39.

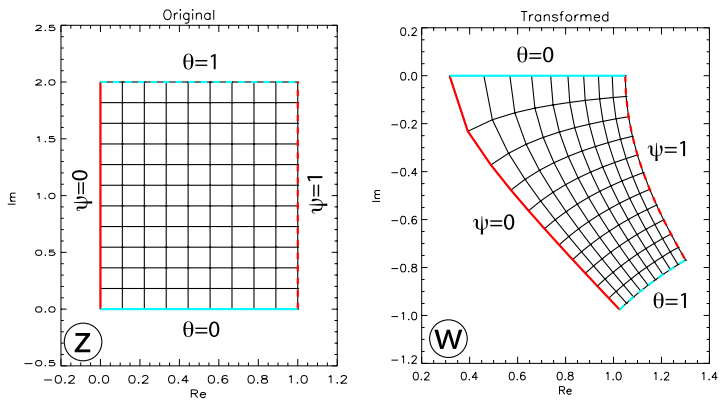

Figure 39. An example "divertor leg" geometry that was generated by applying the conformal mapping $w(z)=(z+0.1)^{1/2}$ to the slab.

Development of FACE wall model: a simple test bed for plasma-wall coupling algorithms

The FACE wall model and its numerical solver have been improved to facilitate coupling with plasma codes SOLPS and UEDGE. The FACE upgrades included:

- Improving precision of simulated volumetric sources of hydrogen, produced by plasma implantation into wall material, to ensure particle conservation in the wall. The sources are now numerically integrated for a given hydrogen implantation depth profile and rescaled in order to correspond exactly to the plasma flux to the wall.
- Implementing automatically adjustable variable time step in FACE to guarantee stability of the numerical solver. The time step is now selected based on the calculated instantaneous rates of physical processes in the wall material to satisfy CFL stability condition for backward Euler numerical scheme.
- Implementing globalization method for iterative solution search, which prevents the implicit Newton iterations from entering a non-converging cycle or deviating too far from current physical state when equations Jacobian becomes small. This is achieved by automatic reduction of the Newton iteration step vector with preservation of its direction or by reduction of the solver time step, when the changes in physical parameters or residual norm become too large.
- FACE has now been fully parallelized using MPI, so that several instances of this wall code can be run simultaneously for each plasma time step.

Coupling strategies for dynamic recycling of plasma-wall models: An approximate method has been developed to study coupled plasma-wall instabilities numerically without directly coupling the plasma and wall codes together. This method is based on a numerical study of the linear response of the spatiotemporal variation of the plasma and wall parameters on the plasma and neutral particle and heat fluxes. The technique is based on the determination of the linear response functions of the plasma and wall codes: $j^s_{wp}(\omega)$, $j^c_{wp}(\omega)$, $j^s_{pw}(\omega)$, $j^c_{pw}(\omega)$, $j^s_{wR}(\omega)$, and $j^c_{wR}(\omega)$ as a function of frequency ω and toroidal mode number n. These functions can be deduced from the numerical solutions obtained by running the plasma and wall codes separately and determining the response to a sinusoidal time-dependent perturbation. After analytically continuing the response functions to the complex plane, we obtain the following dispersion equation for the coupled plasma-wall response:

$$\pm i\left\{j_{wp}^{c}(\omega)j_{pw}^{s}(\omega)+j_{wp}^{s}(\omega)j_{pw}^{c}(\omega)+j_{wR}^{c}(\omega)\right\}$$
$$=1-j_{wp}^{s}(\omega)j_{pw}^{s}(\omega)+j_{wp}^{c}(\omega)j_{pw}^{c}(\omega)-j_{wR}^{s}(\omega)$$

One of the important advantages of this approach is that plasma and wall models can be easily be replace. Hence, the results of different plasma and wall models can be used in the above equation without performing any additional simulations. In addition, we note that the response functions needed for dispersion within the equation could be deduced from laboratory experimental data in an experimental set up that is designed to oscillate the particle and heat fluxes exchanged between the plasma and wall.

Technical progress on SOLPS-FACE coupling

An interface to couple a generic wall code to SOLPS-ITER has been developed to run instances of a wall code on various elements of the divertor and wall components. This interface uses the formalism of the EIRENE neutral code, which computes the particles sources due to wall recycling in SOLPS-ITER (Fig. 40). A separate instance of the wall code can be called for each "stratum," which may correspond to one or more elements of the wall. A recycling coefficient is calculated by the wall code for each stratum, making the use of wall code thus fully compatible with the various options of neutral sourcing in EIRENE (e.g., pumping, puffing, …). To take advantage of the parallelization of the calls to EIRENE in SOLPS-ITER, the FACE wall model will be used as a simplified test bed for developing plasma-wall algorithms. FACE has been fully parallelized using MPI, so that several instances of this wall code can be run simultaneously for each SOLPS-ITER time step. The package SOLPS-ITER/FACE will be applied to the investigation and the modeling of fuel recycling after transient wall loading due to ELMs in collaboration with the Boundary Group at DIII-D.

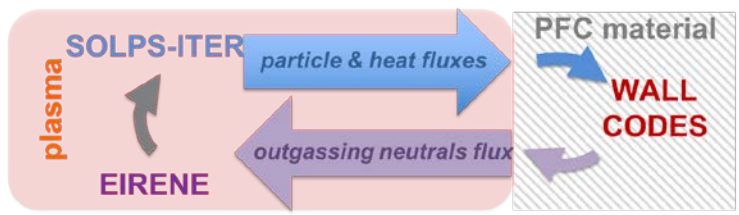

Figure 40. Information flow between plasma and wall codes.

Explicit plasma-wall coupling between UEDGE and FACE

In order to begin to prototype plasma-wall coupling algorithms, we have coupled the UEDGE plasma code and the FACE wall code. We consider a 1D simulation where the UEDGE code solves for plasma and neutral dynamics with fixed temperatures, T_e=T_i=10 eV, and the FACE code simulates the dynamics of hydrogen in a tungsten wall. On the left end, we use Dirichlet boundary conditions (N_i=5e18 m-3, U_{pi}=0.0, N_g=1e12 m-3, U_{pg}=0.0) for the plasma. On the right end, UEDGE calculates the particle and heat flux impinging on the wall. The FACE simulation determines the flux of neutrals back into the plasma. For this model, the only information supplied by FACE is the "effective recycling coefficient", *recycp*=-Γ_n/Γ_i. Note, however, that prompt reflection, which is usually calculated with the TRIM code is not included in this model.

The tungsten wall has a width of 1cm and initially has no hydrogen in it. The hydrogen ions incoming from plasma are implanted in the wall material at depth 2 nm below the surface with a Gaussian profile having a characteristic width of 1 nm. The wall material initially has a high concentration (equivalent to 1% of tungsten atoms) of uniformly distributed empty hydrogen traps with a de-trapping energy equal to 0.95eV. Transport of the implanted hydrogen in the wall is simulated with the FACE code at each time step including hydrogen diffusion, trapping, de-trapping, and desorption from the wall surface. The simulated transport determines the recycled flux of hydrogen, Γ_n.

The temperature profile in the wall material is also simulated with FACE by solving the heat transport equation. The boundary conditions for the heat equation use the plasma heat flux provided by UEDGE to the wall surface and a constant room temperature at the far end of the wall. As the hydrogen diffusion and trapping/de-trapping rates in the material are dependent on the temperature, the evolution of temperature is important for determination of the hydrogen recycling.

For the results shown below, we use an initial steady-state UEDGE solution, obtained with unity recycling coefficient. Then, using our iterative procedure with a time step, $dt=10^{-6}$ s, we produce a sequence of 1D plasma profiles that after 10 steps reach what looks like a self-consistent state (Figure 41). At this stage, the implanted hydrogen is quickly captured by the empty traps in the material and this suppresses hydrogen diffusion back to the wall surface. Thus, the recycling coefficient becomes small significantly affecting the plasma density profiles.

However, this self-consistent state is not a steady state, it is actually slowly evolving. For a long run, using $dt=10^{-3}$ s, and running to 1.0 s, we see that the ion density (N_i) profile returns to the original state, corresponding to recycp=1.0 (see Fig. 41). Indeed, the recycling coefficient becomes nearly unity toward the end of the run (see Fig. 41), as the density of empty traps and hydrogen trapping rate drop near the wall surface. By the end of the run, a wall surface layer of width ~1μm has traps nearly saturated with hydrogen, allowing the implanted hydrogen to freely diffuse in this region. As the hydrogen implantation depth of 2 nm is much smaller than the saturated layer width, most of the hydrogen diffuses back to the surface producing the recycling coefficient close to 1 at this stage. This test shows that the plasma response time to the highly non-equilibrium recycling changes is of order of 10 μs that is much shorter than the wall response time ~1s.

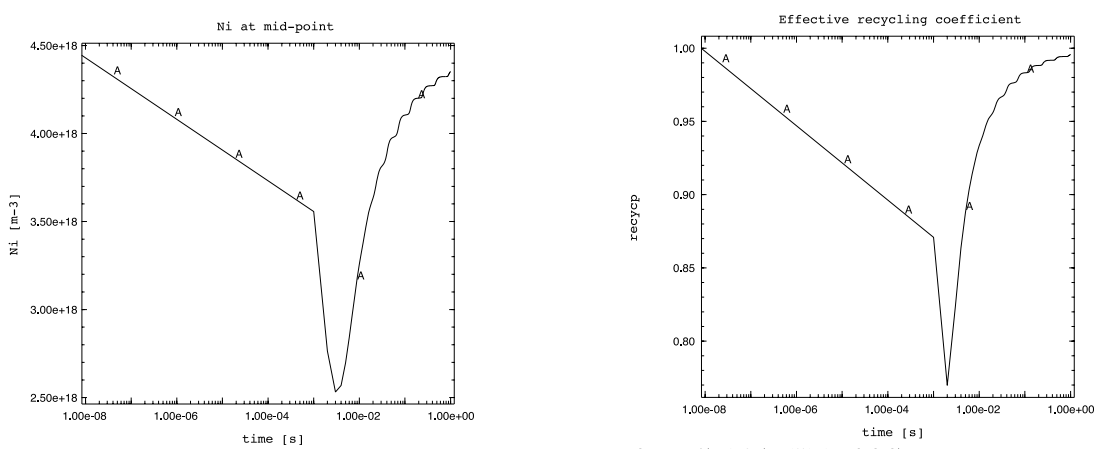

Figure 41. Dynamic simulation of an initially pure tungsten wall that absorbs hydrogen until the recycling coefficient becomes nearly unity. The evolution of the ion density Ni (left) and recycling coefficient (right) are shown for an explicitly coupled simulation using the UEDGE and FACE codes.

<u>Roadmap and plans for kinetic simulation of ELM pulse</u>

Development of COGENT ELM model

In the first year of this project, the main goal is to develop a kinetic model for an ELM heat pulse within the COGENT code and to perform model verification studies. The first objective is to perform a benchmark of the model against a well-known ELM test problem. There are multiple steps that need to be performed including: (i) setting up the benchmark problem, (ii) studying the ELM problem with simplified electron models, (iii) developing sheath boundary conditions that are appropriate for kinetic electrons, (iv) studying the ELM problem with kinetic electrons, and (v) exploring ways of mathematically accelerating the two-species kinetic problem. To this end, we have begun studies of ELM

heat pulses with kinetic ions and the simple adiabatic electron model. We have also implemented sheath boundary conditions for the gyro-Poisson equation within COGENT and performed initial verification studies of kinetic electron and kinetic ion species sheath simulations. We will present our results at the upcoming PSI conference and will publish our work in the associated conference proceedings.

We have set up an initial ELM simulation similar to the benchmark case but have not yet attempted to reproduce the exact parameters. The geometry corresponds to the scrape-off layer of the JET tokamak with a length of $2L_{\|}$ = 100 m, major radius, R= 3 m, nominal SOL thickness 0.1 m, a magnetic field strength B = 3T, and the ratio of poloidal field to toroidal field is B_p/B_t = 0.1; i.e. a field line pitch of approximately 6°. The initial hydrogen plasma is maintained by a Maxwellian source of particles that delivers ions with temperature T_i = 260 eV at a fixed rate S_0 = 6.24x10^{22} m^{-3}s^{-1}. In order to provide meaningful initial condition for the ELM, the profiles are initialized with an initial guess at the equilibrium and the simulation is run until an approximate steady-state is achieved. The ELM is turned on for τ_{ELM} = 200 μs by increasing the source temperature to T_i = 1.0 keV and by increasing the source magnitude by a factor of 10. If equal energy was delivered to ions and electrons, this would correspond to an ELM size of W_{ELM}=75 kJ. Finally, the source is returned to its initial state, and the simulation follows the dynamics of the particle and heat fluxes through the SOL.

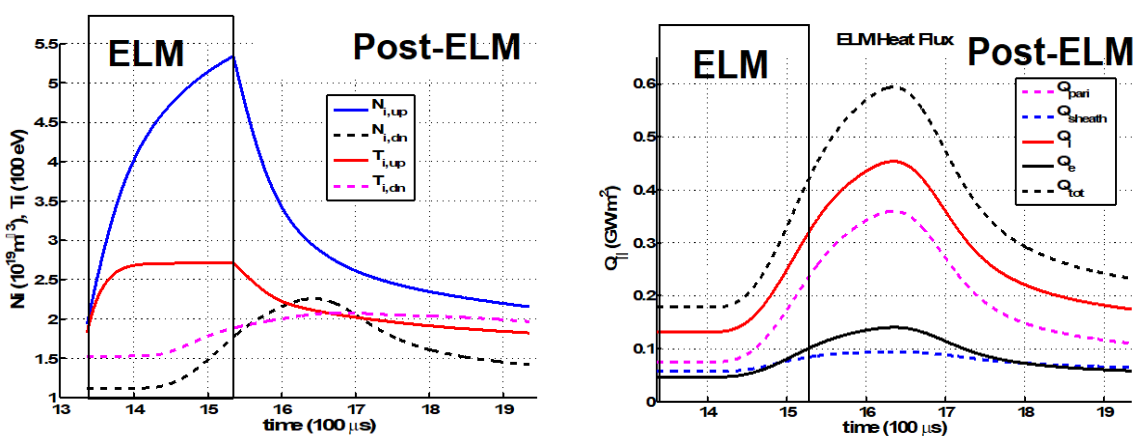

(a) (b)

Figure 42. Simulation results for an ELM heat pulse impinging on the divertor of a JET-like tokamak scrape-off layer using the COGENT kinetic ion and adiabatic electron model. (a) The ion density (N_i) and temperature (T_i) at the midplane (up) and at the target plate (dn): $N_{i,up}$ (blue solid), $N_{i,dn}$ (black dashed), $T_{i,up}$ (red solid) $T_{i,dn}$ (magneta dashed). (b) The heat flux impinging on the divertor target plate (Q_{tot} black dashed) and the individual components: electron heat flux (Q_e black solid), ion heat flux (Q_i red solid), ion heat flux on the plasma side of the sheath (Q_{pari}, magenta dashed), and ion heat flux due to sheath acceleration ($Q_{i,sheath}$ blue dashed).

The ELM model reported on here uses kinetic ions and an adiabatic model of electrons at a fixed temperature T_e = 100 eV; collisions are also neglected. The adiabatic assumption allows the electric potential to be solved for via

$$\phi = \phi_s + (T_e/e)\log(n_i/n_{i,s})$$
$$\phi_s = (T_e/e)\log(\Gamma_{i,s}/n_{i,s}(T_e/2\pi m_e)^{1/2})$$

where $n_{i,s}$ is the ion density, $\Gamma_{i,s}$ is the ion flux at the sheath entrance, and ϕ_s is the potential difference across the sheath for insulating logical sheath boundary conditions. The source is assumed to have a Gaussian profile along the poloidal direction with a length of 3.125 m; i.e. corresponding to a length of L_s

= 31 m along the field line. With these conditions, a steady-state is achieved after running the simulation time for 13 ms, at which point the upstream density is $n_i = 2x10^{19}$ m^{-3} and the upstream temperature is T_i = 175 eV. The ELM and post-ELM be seen in Figure IV.C.1(a). After the ELM is turned on as described above, the density rises above $5x10^{19}$ m^{-3} and the temperature rises to 275 eV. The particles and heat reach the divertor targets roughly 100 μs after the ELM is turned off, which is characteristic of the ion sound transit time. The heat flux delivered to the target plate is shown in Fig. 42b. The steady-state heat flux Q_{tot} = 0.2 GW/m^2 is increased by a factor of 3 during the ELM to a maximum of Q_{tot} = 0.6 GW/m^2. The total heat flux (Q_{tot}, dashed black) is the sum of the ion heat flux (Q_i) and the electron heat flux (Q_e = 2 T_e Γ_{ion}, solid black). The ion heat flux (Q_i, solid red) is calculated as the sum of the conductive and convective components (Q_{pari}, dashed magenta) as well as the contribution due acceleration through the plasma sheath (Q_{sheath} = $e\phi$ Γ_{ion}). This latter term exchanges electron heat flux on the plasma side of the sheath with ion heat flux on the wall side of the sheath.

We are actively working on extending the results to two kinetic species: both ions and electrons. In order to extend to the case of kinetic electrons, we have implemented a version of the gyrokinetic sheath boundary condition similar to that in Shi [25]. We have begun the study of sheath simulations small domain with the full Poisson equation and compared the results to using the gyro-Poisson equation. We have also exercised the sheath BCs to verify their performance. Using the gyrokinetic Poisson equation instead of the full Poisson operator has the beneficial effect of eliminating the fast time scale due to the electron plasma wave. Instead, the plasma wave dispersion relation turns into the electrostatic Alfven wave dispersion relation:

$$\omega^2 = V_A^2 k_{||}^2 / k_\perp^2 d_e^2$$

where V_A is the Alfven speed and $d_e = c/\omega_{pe}$ is the electron skin depth. However, this dispersion relation is incorrect and unphysical if the wavelength is longer than the electron skin depth k_{perp} d_e<1. In the limit $k_{||}/k_{perp}$<<1, the frequency grows without bound, leading to a stringent CFL condition. In a presentation at the Transport Task Force Workshop [6], it was shown that this problem can be eliminated by solving for a separate Helmholtz equation for $E_{||}$ in addition to the gyro-Poisson equation for ϕ. This yields the physically correct shear Alfven wave dispersion relation

$$\omega^2 = V_A^2 k_{||}^2 / (1 + k_\perp^2 d_e^2)$$

and eliminates the singularity.

Discussion of Status of Year 1 milestones:
We now discuss the status of the year 1 milestones, which were documented in the proposal. Four main milestones were proposed for the initial year of research activities. Table 2 provides a list of those milestones, along with a description of the status of the milestones.

Table 2. Summary of the proposed year 1 milestones, along with a description of the status and plans for successful completion.

Milestone descriptions	**Status**
Coupled boundary plasma-surface evolution simulations of ITER tungsten divertor tiles, demonstrating predicted fuel recycling and tritium retention, and evaluating impurity transport to determine net erosion characteristics	This milestone is making good progress, and is on track for completion by September 30, 2018. The results of this milestone are being documented as the DOE Fusion Theory PMM milestone.
Initial parameterization of W-Be-He IAP, targeted MD/AMD simulations to evaluate modifications in He transport and clustering behavior; KMC/MD isothermal simulations of sub-surface gas bubble & defect populations	As documented earlier in the report, the initial SNAP Be-W potential has been fit and is currently being tested. Atomistic simulations have been performed to evaluate the effect of helium bubbles on tungsten thermal expansion, and identified a possible new mechanism of enhanced hydrogen trapping in tungsten. Ongoing atomistic and mesoscale simulations are making good progress , and will provide additional benchmarking data for the Xolotl code.
Initiate investigation of magnetic pre-sheath turbulence; begin extending hPIC capability to incorporate non-planar surfaces and introduce unstructured meshing into GITR and hPIC	As noted in Figure 35, efforts are underway to validate the hPIC magnetic pre-sheath models and are making good progress. The incorporation of the non-planar surfaces has been completed within GITR and hPIC as part of the FY2018 Fusion Theory PMM milestone. Introduction of the unstructured meshing capability will begin in the last quarter of FY2018, and extend into year 2 of the project.
Implement PSI-relevant turbulence model within BOUT++ and ELM-relevant heat pulse model within COGENT incorporating sheath boundary conditions and neutral re-cycling	Development of the divertor-relevant BOUT++ turbulence model and the ELM-relevant heat pulse model within COGENT are in progress and the initial models will be completed by the end of FY2018. The BOUT++ turbulence physics model now includes drift-resistive ballooning physics and linear sheath boundary conditions that drive preseath turbulence modes such as the conducting wall mode. A neutral physics module will be added and additional linear and nonlinear verification studies will be performed over the remainder of the year. ELM heat pulse simulations have been performed with kinetic ions and adiabatic electrons and we have initiated the study of sheath simulations with two kinetic species: both ions and electrons using gyrokinetic boundary conditions. Over the remainder of the year we plan to perform our first ELM heat pulse simulations with two kinetic species.

Plans for Year 2 activities:
Overall, the project has had a very successful first year. Plans are underway to host a mini-conference on plasma material interactions at the 2018 APS-DPP meeting in November, and we will use that venue to have an annual team meeting to review progress and finalize the plans for upcoming activities. We expect to continue to make progress at the atomic scale on discovering the behavior of helium and hydrogen below mixed material surfaces as a result of beryllium implantation, and to continue to refine the code coupling strategy for simulating plasma surface interactions, including the dynamic feedback of the surface to the edge plasma. In the next year, we will also perform a systematic uncertainty quantification of the coupled plasma surface interaction simulations of the ITER divertor.

PROJECT BIBLIOGRAPHY

Peer-reviewed articles acknowledging SciDAC project funding:

- K. D. Hammond, "Helium, Hydrogen, and Fuzz in Plasma-Facing Materials," Materials Research Express 4, 104002 (2017).
- R. P. Doerner, D. Nishijima, S. I. Krasheninnikov, T. Schwarz-Selinger and M. Zach "Motion of W and He atoms during formation of W fuzz", Nucl. Fusion 58 (2018) 066005.
- K. D. Hammond, S. Blondel, L. Hu, D. Maroudas, and B. D. Wirth, "Large-scale atomistic simulations of low-energy helium implantation into tungsten single crystals," *Acta Materialia* **144**, 561-578 (2018).
- J. Drobny, D. Curreli, "F-TRIDYN Simulations of Tungsten Self-Sputtering and Applications to Coupling Plasma and Material Codes", Computational Materials Science, Volume 149, 15, Pages 301–306 (2018)
- S. I. Krasheninnikov "On numerical study of coupled plasma-wall instability," accepted by Phys. Plasmas, 2018.
- L. Yang, Z.J. Bergstrom and B.D. Wirth, "First-principles study of helium-vacancy complexes below tungsten surfaces", Journal of Applied Physics, **123 (**2018) 205108.
- L. Yang and B.D. Wirth, "Energetics of hydrogen and helium-vacancy complexes in bulk and near surfaces of tungsten: First-principles study", Journal of Applied Physics, **123** (2018) 215104.
- Z. Yang, L. Hu, D. Maroudas, and K.D. Hammond, "Helium Segregation and Transport Behavior Near Symmetric Tilt Grain Boundaries in Tungsten", accepted by Journal of Applied Physics, 2018.

Presentations:

- (Invited) Thompson, A. P., "SNAP: Automated Generation of High-accuracy Interatomic Potentials Using Quantum Data," CECAM Workshop "Machine Learning at Interfaces," June 4-8, 2018, Lausanne, Switzerland
- (Invited) Thompson, A. P. "It's a SNAP: Automated Generation of High-accuracy Interatomic Potentials Using Quantum Data" The Minerals, Metals and Materials Society Annual Meeting. Phoenix, AZ. March 11 – 15, 2018.
- Wood, M. A. and A. Thompson. "Linear Scaling, Quantum-Accurate Interatomic Potentials with SNAP; Reaching those Hard-to-Reach Places in Classical Molecular Dynamics" The Minerals, Metals and Materials Society Annual Meeting. Phoenix, AZ. March 11 – 15, 2018.
- Wood, M. A. and A. Thompson. "Linear Scaling, Quantum-Accurate Interatomic Potentials with SNAP" American Physical Society March Meeting. Los Angeles, CA. March 4 – 10, 2018.
- (invited) A. P. Thompson, "Molecular Dynamics Simulation: Engine of Discovery or Bridge to Nowhere?," Department of Chemistry, University of Missouri, Columbia, Missouri, February 23, 2018.
- (Review) S. I. Krasheninnikov "Divertor Plasma Detachment: Past and Future", 23rd International Conference on Plasma Surface Interactions in Controlled Fusion Devices, Princeton, NJ, USA, June 18, 2018.

- R. D. Smirnov and S. I. Krasheninnikov “Model of hydrogen retention in tungsten with self-induced trap formation”, 23rd International Conference on Plasma Surface Interactions in Controlled Fusion Devices, Princeton, NJ, USA, June 19, 2018.
- I. Joseph, M. Dorf, and M. A. Dorr, “Kinetic simulation of heat pulse propagation through the tokamak scrape-off layer,” Plasma Surface Interactions Workshop, June 22, 2018.
- I. Joseph, “Taming the Electrostatic Alfven Mode and the Ampere Cancellation Problem,” Transport Task Force Workshop, May 9, 2018.
- Tim Younkin, “Impurity Migration Simulation of PISCES-A He Exposed Tungsten With GITR”, APS-DPP 2017, Milwaukee, Wisconsin.
- Ane Lasa, “Multi-physics modeling of the long-term evolution of surfaces exposed to steady-state plasmas”, APS-DPP 2017, Milwaukee, Wisconsin.
- Dwaipayan Dasgupta, “Modeling of Fuzz Formation on Helium-Ion-Irradiated Tungsten Surfaces”, APS-DPP 2017, Milwaukee, Wisconsin.
- Sophie Blondel, “Calculations of Helium Bubble Evolution in the PISCES Experiments with Cluster Dynamics”, APS-DPP 2017, Milwaukee, Wisconsin.
- Jon Drobny, “Fast, Statistical Model of Surface Roughness for Ion-Solid Interaction Simulations and Efficient Code Coupling”, APS-DPP 2017, Milwaukee, Wisconsin.
- D. Dasgupta, K. D. Hammond, L. Hu, D. Maroudas, and B. D. Wirth. American Institute of Chemical Engineers (AIChE) Annual Meeting, Minneapolis, Minnesota, November 1, 2017. *Modeling of Nano-Fuzz Formation in Helium-Ion-Irradiated Tungsten.*
- K. D. Hammond, S. Blondel, L. Hu, D. Maroudas, and B. D. Wirth. American Institute of Chemical Engineers (AIChE) Annual Meeting, Minneapolis, Minnesota, November 1, 2017. *Large-Scale Atomisic Simulations of Low-Energy Helium Implantation into Tungsten Single Crystals.*
- Brian Wirth, (Invited) plenary presentation, “Overview of modeling plasma surface interactions in tungsten with a focus on connecting computational predictions to experimental observations”, 18th International Conference on Fusion Reactor Materials, Aomori, Japan, 10 November 2017

REFERENCES CITED

[1] Sisneros, Robert & S. Meredith, Jeremy & Ahern, Sean & Pugmire, Dave. (2012). EAVL: The extreme-scale analysis and visualization library. Eurographics Symposium on Parallel Graphics and Visualization (EGPGV).

[2] Moreland, Kenneth, et al. "Vtk-m: Accelerating the visualization toolkit for massively threaded architectures." IEEE computer graphics and applications 36.3 (2016): 48-58.

[3] Bartok, A. P. *et. al.*, "On representing chemical environments", PRB 87 (2013) 184115.

[4] Eldred, M. et. al. Technical Report SAND 2006–6337, Sandia National Laboratories, 2006.

[5] N. Juslin and B. D. Wirth, "Interatomic potentials for simulation of He bubble formation in W," *Journal of Nuclear Materials* **432**, 61-66 (2013).

[6] Björkas, C., *et al.*, "A Be-W Interatomic Potential", J. Phys.: Condens. Matter, 22 (2010) 352206.

[7] K. D. Hammond, S. Blondel, L. Hu, D. Maroudas, and B. D. Wirth, "Large-scale atomistic simulations of low-energy helium implantation into tungsten single crystals," *Acta Materialia* **144**, 561-578 (2018).

[8] R. D. Smirnov, S. I. Krasheninnikov, J. Guterl, Atomistic modeling of growth and coalescence of helium nano-bubbles in tungsten, J. Nucl. Mater. 463 (2015) 359.

[9] T. Dengg, V. Razumovskiy, L. Romaner, G. Kresse, P. Puschnig, and J. Spitaler, "Thermal expansion coefficient of WRe alloys from first principles," *Physical Review B* **96**, 035148 (2017).

[10] G. Nandipati, W. Setyawan, H. L. Heinisch, K. J. Roche, R. J. Kurtz, B. D. Wirth, "Displacement cascades and defect annealing in tungsten, Part II: Object kinetic Monte Carlo simulation of tungsten cascade aging," *Journal of Nuclear Materials* 462 (2015) 338.

[11] G. Nandipati, W. Setyawan, H. L. Heinisch, K. J. Roche, R. J. Kurtz, B. D. Wirth, "Object Kinetic Monte Carlo Simulations of Microstructure Evolution," *Semiannual Progress Report DOE/ER-0313/54* (2013) 179.

[12] Z. Yang, Y. S. Blondel, K. D. Hammond, B. D. Wirth, "Kinetic Monte Carlo Simulations of Helium Cluster Nucleation in Tungsten with Preexisting Vacancies" *Fusion Science & Technology* 71 (2017) 60.

[13] D. Maroudas, S. Blondel, L. Hu, K. D. Hammond, B. D. Wirth, "Helium segregation on surfaces of plasma-exposed tungsten," *Journal of Physics: Condensed. Matter* 28 (2016) 064004.

[14] S. Blondel, K. D. Hammond, L. Hu, D. Maroudas and B. D. Wirth, "Modeling Helium Segregation to the Surfaces of Plasma-Exposed Tungsten as a Function of Temperature and Surface Orientation," Fusion Sci. Technol., 71 (2017) 22; http://dx.doi.org/10.13182/FST16-112

[15] L. Yang and B.D. Wirth, "Energetics of hydrogen and helium-vacancy complexes in bulk and near surfaces of tungsten: First-principles study", accepted by Journal of Applied Physics, 2018.

[16] L. Yang, Z.J. Bergstrom and B.D. Wirth, "First-principles study of helium-vacancy complexes below tungsten surfaces", accepted by Journal of Applied Physics, 2018.

[17] N. Juslin and B.D. Wirth, Molecular dynamics simulation of the effect of sub-surface helium bubbles on hydrogen retention in tungsten, Journal of Nuclear Materials 438, S1221 (2013).

[18] L. Wang, X. Shu, G. Lu, and F. Gao, Journal of Physics: Condensed Matter 29 (2017) 435401.

[19] R. Frauenfelder, J. Vacuum Science Technology 6 (1969) 388.

[20] A. A. Kohnert and B. D. Wirth, "Grouping techniques for large-scale cluster dynamics simulations of reaction diffusion processes," Model. Simul. Mater. Sci. Eng., 25, 015008 (2017); http://dx.doi.org/10.1088/1361-651X/25/1/015008

[21] B. Braams, Radiative Divertor Modelling for ITER and TPX, *Contributions to Plasma Physics* **36**, 276 (1996).

[22] W. R. Elwasif *et al*., "The Design and Implementation of the SWIM Integrated Plasma Simulator," *2010 18th Euromicro Conference on Parallel, Distributed and Network-based Processing*, Pisa, 2010, pp. 419-427.

[23] G.R. Tynan, M. Baldwin, R. Doerner, E. Hollman, D. Nishijima, K. Umstadter, and J. Yu, *AIP Conference Proceedings* **1237** (2010) 78; https://doi.org/10.1063/1.3447995

[24] S. Blondel, D.E. Bernholdt, K.D. Hammond, L. Hu, D. Maroudas, and B.D. Wirth, “Benchmarks and Tests of a Multidimensional Cluster Dynamics Model of Helium Implantation in Tungsten,” Fusion Sci. Tech., 71 84-92 (2017).
[25]. E.L. Shi, A.H. Hakim, G.W. Hammett, Phys. Plasmas, 22 (2015) 022504.