

An Advanced Computing Initiative To Study Methods of Improving Fusion¹

G.W. Hammett, hammett@pppl.gov (PPPL), Aug. 10, 2014

with input from

C.S. Chang, S. Kaye, A. H. Hakim (PPPL), A. Pletzer, J. Cary (Tech-X)

2-Page Executive Summary:

The exponential growth of computer power (a factor of a million in the past 30 years, with more to come), combined with major advances in our ability to simulate key aspects of plasmas in fusion devices, has led many to conclude that the time is ripe for an initiative to develop integrated computer simulations of fusion devices. There are now detailed 5D gyrokinetic simulations in the main core region of tokamaks ($r/a < \sim 0.9$) that can predict fluctuation spectra and turbulent transport fairly well in many regimes. The general feasibility of coupling short time scale gyrokinetic turbulence codes to long-time-scale transport codes has been demonstrated, using massively parallel computers and implicit projective integration algorithms. However, these simulations need to use measured boundary conditions (or semi-empirical models) at $r/a \sim 0.9$ and so are not yet predictive. One of the biggest remaining challenges is to develop gyrokinetic simulations that can handle the additional complexities of the pedestal and scrape-off-layer edge region ($r/a > \sim 0.9$). (These computational challenges include the need to handle the separatrix and open and closed field lines, large amplitude fluctuations, and magnetic fluctuations near the beta limit.) These are hard problems, but progress is being made, and there are recent advances in algorithms that may significantly help, such as discontinuous Galerkin algorithms. The success of core gyrokinetic simulations gives us encouragement that a significant initiative should be able to develop codes that fully simulate the edge as well, and thus enable integrated simulations for prediction of the whole device.

While we believe this is feasible, much work remains to realize it. Besides developing codes that are complete enough to successfully handle the edge region, more work is




Fig. 1: Example of the kind of comprehensive gyrokinetic turbulence simulations that are fairly successful in the core region of tokamaks. A computational initiative could develop gyrokinetic codes capable of handling the additional complexities of the edge region, and thus enable predictive integrated simulations of the whole device. (Candy & Waltz, GA.)

¹ This white paper covers both the Priorities and Initiatives topics. The slides for the related talk at http://fire.pppl.gov/FESAC_2014_IC_Initiative_Hammett.pdf contain some additional material.

needed to understand different types of core and edge turbulence, to test these codes against experiments in a wider range of parameters, and to improve some of the source and sink modules. Because of the importance of this area, the Panel heard multiple presentations on different aspects of a computational initiative, focusing on projects that can both undertake standalone goals (like MHD disruption studies) and serve as modules in an integrated simulation initiative. For example, a transport framework could evolve profiles with transport from small-scale turbulence calculated by gyrokinetic modules running on extreme scale computers, and extended MHD modules would be periodically called to check large-scale stability and calculate the spreading of energetic particles. While different levels of integration can be pursued, all will involve some level of modularity, and a modular framework is essential so that different options can be used for different purposes. Independent modules are needed to cross-check each other and find efficient algorithms for these very hard problems (kinetic plasma dynamics in 5D is a grand challenge). Reduced transport models would be used in interpretive and predictive runs for fast turn-around for discharge analysis, transport model validation, and shot scenario development and prediction. Such tools would be used for an extensive validation campaign comparing with present experiments, and would be used to make projections to future devices. Every shot on ITER will first have to be simulated with these types of codes, to predict if they will avoid disruption limits.

Relation To FESAC Reports and DOE Goals: The high priority of an integrated simulation initiative (including edge / boundary layer simulations) has been recognized in several recent FESAC/DOE reports. Such an initiative would directly address the top 4 ReNeW Thrusts as prioritized in the 2013 “Report of the FESAC Subcommittee on the Priorities of the Magnetic Fusion Energy Science Program”, the top 2 being “Thrust 2: Control Transient Events in Burning Plasmas”, and “Thrust 6: Develop Predictive Models for Fusion Plasmas, Supported by Theory and Challenged with Experimental Measurement”. A computational initiative is obviously well aligned with the DOE Office of Fusion Energy’s goals, including the first goal to “Advance the fundamental science ... to develop the predictive capability needed for ... fusion energy”. An integrated initiative addresses the challenges posed by Synakowski in his charge to the panel, such as the challenge to “understand the fundamentals of transport, macro-stability, wave-particle physics, plasma-wall interactions”.

Leadership and Funding: This initiative will build on US leadership in general computing and will make effective use of extreme-scale computing. The US has been a historical leader in theory and advanced simulations for fusion. The US could play a leadership role in the world fusion program through advanced simulations if a significant initiative was begun soon, otherwise overseas programs may take the lead. An investment of order \$10-15M/y initially could have a significant impact, eventually ramping up to the ~\$25M/y level of the original FSP (Fusion Simulation Project) proposal.

Strategic Impact: The FESAC Strategy panel has been hearing about a number of innovative ideas for improving fusion (including liquid metals, advanced tokamak and ST regimes, new divertor and stellarator concepts, and high field superconductors). The goal of a major computational initiative would be to develop comprehensive simulations within five years that are being tested against experiments and that can be used to help evaluate and optimize these innovative concepts. If begun soon, these comprehensive simulations can help make a strong case to policy makers for the next steps in fusion research.

The rest of this document gets into a number of technical issues, some of which are hopefully useful but some of which are more detailed than is needed by the FESAC Strategic Planning Panel. Hopefully this does not obscure the main points that are made in the above 2-page summary.

Motivation: The need for predictive integrated simulations for fusion energy.

The overall advantages of developing better computer simulations for fusion are well recognized in DOE's goals, in various FESAC reports, and in past community-based reports for the Fusion Simulation Project (FSP), as documented in sections below. The main factor that determines the minimum acceptable size (and thus the construction cost) of a fusion reactor is the achievable fusion triple product $nT\tau_E$ (the product of the density, temperature, and confinement time), which controls the fusion gain Q . Thus finding ways to improve the confinement time can significantly improve the attractiveness of fusion reactors (as illustrated in one way in slide #3 of the presentation corresponding to this paper). While ITER's main pulsed operating scenario is conservative and assumes a confinement factor (a multiplier on a standard empirical confinement time) of $H=1$, the steady-state scenario assumes 50% improvements can be achieved, as does the ARIES-AT reactor design study and the MIT ARC design. While these ~50% levels of improvement have been achieved in some experiments, comprehensive simulations would help understand how these improved regimes work and how they scale from present experiments to reactors, increasing our confidence in projections for reactors and giving us tools that we can use to try to optimize the design of future reactors. Confinement has other benefits as well (particularly if it can be combined with improvements in the normalized beta limit). Higher confinement allows one to reduce the current needed for a tokamak, and thus reduces the current drive and recirculating power requirements.

The need for an initiative in simulating the edge region.

As illustrated in slides number 4-6 of the presentation, a lot of progress has been made in understanding turbulence in the main core region of tokamaks, and we now have quite comprehensive gyrokinetic simulations that include all of the key effects thought to be important in the core region and that can predict the measured turbulence levels and fluctuation spectra fairly well in the core region of many experiments. However, these codes are highly optimized for the core region of tokamaks, and can't handle the additional complexities of the edge region, such as the change in magnetic topology at the separatrix, the large amplitude fluctuations in the edge (the main algorithms used in core simulations would have difficulties with negative density overshoots in the edge region), and the complex plasma-wall interactions near the boundary.

Transport models (such as the TGLF and GLF23 models) based on 5-D gyrokinetic turbulence simulations have been used to predict temperature and density profiles in the core region of ITER, and thus to predict the performance of ITER, as illustrated in Fig. 2, but the results are strongly dependent on the assumed pedestal temperature. [This strong dependence of the core on the edge is due to stiff transport with marginal stability effects, where the turbulence would increase strongly if the normalized temperature gradient (or in some regimes the density gradient) $R/L_T = -R d\log(T)/dr$ exceeds a certain threshold. The-

se marginal stability effects explain how many experiments, like supershots and H-modes, show a strong sensitivity to the edge.] Comprehensive gyrokinetic simulations that can handle the edge region are needed to predict not only the height of the H-mode pedestal temperature but other important questions too, such as: What are the threshold conditions to achieve an H-mode transport barrier in the first place? Can RMP coils suppress ELMs without degrading performance too much? Why is the SOL width observed to be so narrow, and are there ways it can be broadened? Can advanced divertor concepts handle the extreme power densities at reactor scales? Can liquid metals, perhaps with continuous vapor shielding (as discussed in the presentations by Jaworski and Goldston) help with the divertor power handling problem at reactor scales?

As illustrated in Fig. 3, there is a SciDAC project working on developing a PIC-based edge gyrokinetic turbulence code, XGC1. It is producing encouraging initial electrostatic results, giving similar blobby edge turbulence features and a similar SOL width as in the experiment. This is an example of extreme-scale computing, solving for 40B particles on 131k processors and 8k GPUs on the Cray XK-7 at ORNL for a total of 6M processor hours (about 5% of the yearly allocation for this work). [Continuum codes can also use extreme-scale computing efficiently, as one can parallelize over grid points in phase space just as one parallelizes over particles in a PIC code, though continuum codes tend to be more complicated and will take more human effort to optimize.] This correspond to about 1 ms of physical time in the tokamak, short compared to the global energy confinement time but comparable to the edge/SOL dynamical times. XGC1 is the only gyrokinetic code at present that can handle a separatrix and open and close field lines simultaneously, which is important for accurate treatment of the edge region. While these initial electrostatic results are encouraging, work is ongoing to include more physics in this code, and in particular to include magnetic fluctuations, which has been a challenge in gyrokinetic PIC codes because of a cancellation problem (two large terms that need to nearly offset each other with high accuracy) that becomes more difficult at longer wavelengths and higher beta. Two algorithms for magnetic fluctuations with kinetic electrons have been tried in XGC1 so far. One has a general formulation but has numerical difficulties at long wavelengths and high beta. The other works at long wavelengths but doesn't allow magnetic field lines to break. There is ongoing work to investigate some new algorithms and see how much they can improve the handling of magnetic fluctuations. Electromagnetic effects are particularly important in the edge region because pressure gradients there are often near the beta limit.

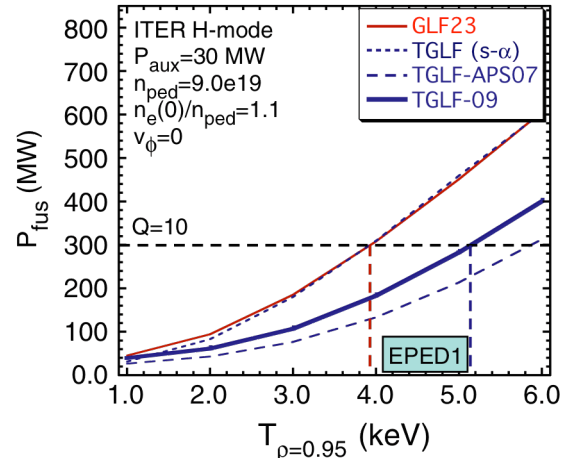


Fig. 2. Gyrokinetic-based predictions of fusion power in ITER as a function of the assumed pedestal temperature (the curves correspond to different assumptions about the gyrokinetic model, but all show a strong dependence on the pedestal temperature). Significant work is needed to develop gyrokinetic codes that can handle the edge region to predict the pedestal temperature and answer other important questions. (Kinsey et al. Nucl. Fus. 2011 <http://stacks.iop.org/NF/51/083001>)

Continuum gyrokinetic codes (i.e., codes that use Eulerian algorithms instead of PIC algorithms) are being widely used to study the core region of tokamaks, and they have been fairly successful in including magnetic fluctuations (perhaps cancellations are easier to preserve because the particle distribution function is on the same spatial grid as the fields, unlike randomly positioned particles in a PIC code). This is motivation for developing continuum codes for the edge region as well. There are some smaller efforts in this regard, including the Edge Simulation Laboratory project using a finite-volume algorithm at LLNL, and a project at PPPL exploring Discontinuous Galerkin (DG) algorithms for edge gyrokinetics (at present funded by LDRD). DG algorithms have been a hot topic in the applied math and computational fluid dynamics communities in the past 15 years, and combine certain attractive features of finite-volume and finite-element methods. Certain versions of DG can conserve energy exactly for kinetic Hamiltonian problems even with upwind fluxes. [Energy conservation is not exact for standard finite volume methods, because energy conservation is different for kinetic problems than fluid problems, where one of the equations is explicitly an energy conservation law.] DG may have several other advantages as well, including the flexibility to consider Maxwellian-weighted basis functions. Ultimately one could consider a combination of algorithms, such as a PIC algorithm for ions and a continuum DG algorithm for electrons to handle magnetic fluctuations.

Some of these issues are too detailed for the FESAC Strategic Planning Panel to concern themselves with, but they are mentioned to point out that these problems are computationally and algorithmically very challenging, and part of any computational initiative in fusion needs to invest in multiple approaches on such problems. It is essential to have multiple codes, particularly for very difficult problems like plasma turbulence in high-dimensional kinetic phase space, as they provide valuable cross-checks against each other (along with verification and validation tests). The excellent progress made in core gyrokinetic turbulence over the past 15 years has been in part because of the existence of multiple independent groups exploring different approaches and different physical effects.

While the edge region is challenging, the excellent progress that has been made in understanding turbulence in the core region of tokamaks gives us encouragement that the edge region should be tractable as well, with a major initiative to develop codes with ad-




Fig. 3: Example of a PIC gyrokinetic simulation (XGC1) of electrostatic turbulence in the outer regions of a DIII-D plasma, giving similar blobby edge features and SOL width as observed in the experiment. At present this is the only gyrokinetic edge turbulence code that can simultaneously handle the separatrix, closed, and open field lines. Work is ongoing to include more physics, including magnetic fluctuations.

http://tff2014.ucsd.edu/TTF_2014/Presentations_files/Ku.pdf

vanced algorithms and capabilities that can handle the additional complexities of the edge region.

Before moving on, we briefly mention that although core gyrokinetic codes agree well with many experiments in many regimes, some codes are underpredicting the turbulence level in the outer regions ($r/a > 0.6$) of some cold L-mode plasmas (the “shortfall problem”). The reasons for this are still under investigation. Some codes do not see a significant shortfall for these cases, even though in another parameter regimes they agree with the codes that do see a shortfall. This might be because different algorithms are able to resolve different features of the solution more easily, indicating that higher resolution than usual is needed for these cases, though more work is needed to draw a firm conclusion. In any case, it points to the importance of having multiple independent codes, and the value of validation tests with experimental data over a wide range of parameters.

Modular approaches to integrated simulations for whole-device modeling.

Fusion devices are complex devices involving physics on many different scales, and computational models of them will inevitably involve multi-physics/multi-scale coupling of codes that handle the physics of different parts of the problem. For example, the TRANSP code is an interpretive/predictive transport code used to analyze many experiments. Even though it uses reduced transport models (not full gyrokinetic calculations) of the core region and does not model the scrape-off-layer region in detail, it still has over 2 million lines of code, in part because of the many different options it has for modules with different types of assumptions and different levels of complexity and expense. A modular approach with modern software engineering is essential to managing code of this complexity.

Fig. 4 illustrates some of the standard modules included in present transport codes, with several major extensions needed for a fully predictive, integrated simulation of a whole device. Fig. 5 shows how several of the computational talks to the FESAC panel are discussing different aspects of what can fit together in a modular way for an Integrated Fusion Simulation Initiative.

One major extension is that it would have the option for calculating transport either with reduced models (as used in present transport codes), or directly from fully nonlinear gyrokinetic turbulence codes running on the latest massively parallel computers. The general feasibility of this kind of direct coupling between a long time scale transport code and a short time scale turbulence code has been demonstrated in recent papers by Barnes et al (Ph.D. Thesis 2008, Phys. Plasmas 2010, <http://dx.doi.org/10.1063/1.3323082>) and Candy et al., (Phys. Plasmas 2009, <http://dx.doi.org/10.1063/1.3167820>). (These demonstrations were done with simplified transport codes that neglected the edge and took sources and sinks from a previous full transport code calculation, but they can be generalized.) This is made possible both by the major advances in computational power and by the use of an implicit form of a kind of projective integration technique. This implicit coupling involves carrying out multiple turbulence simulations with slightly different parameters to calculate the Jacobian, as illustrated roughly by the sketch in slide 12 of the presentation corresponding to this white paper. This is faster than a brute force approach (of trying to directly calculate ion-scale or electron-scale turbulence for equilibrium time scales) by a factor of 100's to 1000's. The present coupling is done with multiple flux-tube simulations in the

core, which works well in most cases. In principle one could extend these to full-torus simulations to include non-local turbulence spreading effects, which would be particularly important for the edge region. These initial coupled turbulence-transport simulations took $\sim 100,000$ CPU hours (and will require more CPU time if the edge region and electron-scale turbulence is included), which can be done more routinely over the next 5 years as computer power continues to grow. This kind of implicit projective integration to couple transport and turbulence works well in many cases, but there can be challenges for this kind of multiscale technique, such as events like sawteeth or ELMS which are highly localized in time. This could be an area for applied math / physics research, but it does not seem insurmountable and one can envision various special techniques, such as averaging over these events if the transport code is trying to take time steps that are large compared to the period of those events. In any case, the present implicit projective integration methods appear to work well for MHD quiescent cases or cases where those events are resolved by the transport time step and then directly simulated in the turbulence or MHD modules.

Another major extension in Fig. 4 required for a fully predictive ability is the coupling to gyrokinetic codes that can handle the edge region (the outer 10% of the plasma, from near the top of the pedestal out through the scrape-off layer). As described above, this is a key capability that does not yet fully exist and needs new gyrokinetic codes to be developed.

One could consider evolving to different levels of integration as the capabilities of some of the modules mature. The early version of an integrated simulation would involve initially separate modules for core turbulence (this would use existing gyrokinetic codes) and edge turbulence (this would require a new code), both coupled to a top-level transport code framework for long time scales (and which also couples in heating and fueling sources, radiation sinks, etc.). This framework would periodically call extended MHD codes to check stability as profiles evolve. As edge gyrokinetic codes mature, one could consider extending them to handle core turbulence as well. If such a global gyrokinetic code could also handle magnetic fluctuations, one could imagine it simultaneously also simulating kinetic effects on low- n instabilities that are currently studied by extended MHD codes. However, many years of effort and expertise have gone into the physics and computational efficiency of the existing extended MHD codes so that they can tackle challenging problems, and a lot of work would need to be done before a gyrokinetic code is demonstrated to be able to do similar calculations including kinetic effects. In any case, it should be noted that even approaches that try to handle more of the physics in a single code will still inevitably be relying on a modular framework of some sort, because of the wide range of physical processes that would be treated by different codes (from RF heating to atomic physics to plasma-wall interactions to wall evolution), and because of the need to carry out different types of simulations for different applications. In some cases we will want the ability to run with reduced transport models for fast parameter scans, while in other cases we will want to be able to run with direct calculation of transport by full core and edge gyrokinetic codes at the highest resolution (which may require over 10M CPU hours in some cases) to check the accuracy of the results. Independent modules need to be developed to cross-check each other and find efficient algorithms for different problems. Reduced transport models would be used in interpretive and predictive runs for fast turnaround for discharge analysis (and maybe even real-time feedback control eventually), transport model validation, and shot scenario development and prediction.

There are of course similarities between the kind of Integrated Fusion Simulation Initiative proposed here and the previously proposed Fusion Simulation Project. One difference is that we would start at a lower funding level and initially focus on key missing modules that are needed for a fully predictive simulation (such as edge gyrokinetic codes), and focus on initial applications of MHD modules to the high priority problem of disruptions. I would also advocate trying to develop a more decentralized structure, with Silicon-valley style innovation groups developing key modules, trying different approaches internally, but maintaining standardized external interfaces to facilitate easier testing of different modules in the integrated framework. But these are details that the Strategy Panel does not need to get into.

The Strategy Panel also does not need to get into the details of exactly what would and would not be funded in an Integrated Fusion Simulation Initiative, but here I will outline some of what it could include and estimated budget levels. At the top level, I think it needs at least 2 modular framework teams, maybe TRANSP and something like the FACETS/MEMFIS project. TRANSP is a legacy workhorse code that is widely used by experiments now, and could be extended to directly call gyrokinetic codes to calculate core turbulence, though a more extensive modification would be needed to handle the edge. The FACETS/MEMFIS project (which was supported as a proto-FSP project but not at present) includes the edge region and provides a more modern and more flexible framework to build on. The funding level for TRANSP at present is relatively modest (about \$1M/y including support from PPPL experiments) and is only enough to maintain it (including user support) with few changes. An additional \$1.5M/y would enable a healthier development path for it. A similar ~\$1.5M/y of additional funding would be needed a more modern framework like FACETS/MEMFIS.

Fig. 5 lists estimated costs for initiatives in various other parts of an Integrated Fusion Simulation Initiative (including the need for additional work in edge turbulence simulations), for an estimated total in the range of \$10-\$15M/y, part of which could be from ASCR instead of FES. It could later ramp up to a range comparable to the ~\$25M/y level envisioned for the previously proposed FSP.

Modular Whole-Device Simulations

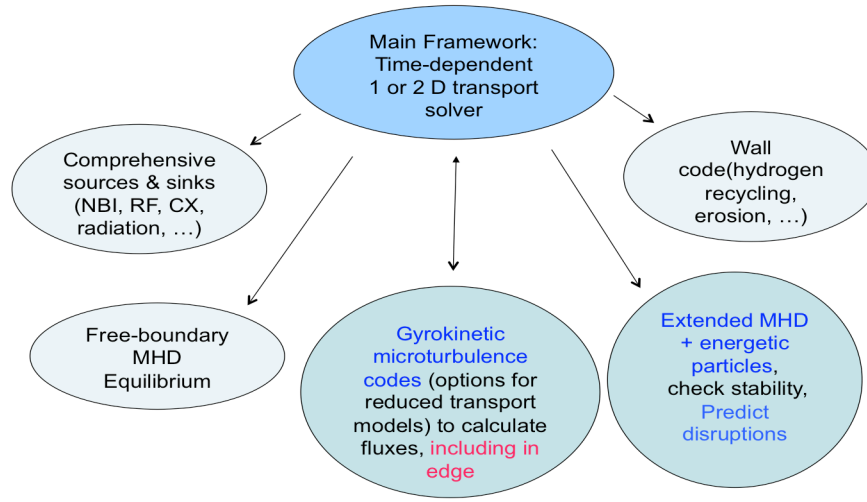


Fig. 4. An integrated simulation initiative would use modular frameworks to call different modules optimized for different purposes and provide a fully-predictive whole-device capability. While various pieces of this exist at present, a major initial focus would be on the lowest two circles. One involves the need to develop gyrokinetic simulations that are comprehensive enough to handle the edge region. The other involves extended MHD simulations, because of the important short-term standalone goals involving disruptions.

Relation to previous FESAC/DOE reports

“Priorities, Gaps and Opportunities: Towards A Long-Range Strategic Plan For Magnetic Fusion Energy”, 2007: This report was at about the same time as the early planning stages for the Fusion Simulation Project (FSP). It mentions how an FSP could help the fusion program in significant ways. There have been several FSP reports since then. As described previously, one of the significant advances since 2007 is the demonstration of the feasibility of transport/turbulence code couplings.

“Report of the FESAC Subcommittee on the Priorities of the Magnetic Fusion Energy Science Program” (Feb. 10, 2013): It is clear that an Integrated Simulation Initiative as proposed here would directly address the top 4 ReNeW Thrusts as prioritized in this 2013 report:

- * Thrust 2: Control Transient Events in Burning Plasmas
- * Thrust 6: Develop Predictive Models for Fusion Plasmas, Supported by Theory and Challenged with Experimental Measurement
- * Thrust 9: Unfold the Physics of Boundary Layer Plasmas
- * Thrust 10: Decode and Advance the Science and Technology of Plasma-Surface Interactions

An Integrated Simulation Initiative would of course directly address Thrust 6. It could be used as a “flight simulator” to simulate all discharges for ITER ahead of time, to predict if they will stay far enough away from known disruption boundaries (and thus aids Thrust 2). The key missing piece in developing a completely predictive capability is the ability to predict the edge region (Thrust 9), which is why we propose this as an early high priority. Eventually codes like this will be needed to explore plasma-surface interactions and ad-

vanced divertor concepts (including the use of liquid metals), and so is part of Thrust 10 as well.

There were several detailed reports written as part of the proposed Fusion Simulation Project planning process, including:

2007 Fusion Simulation Project (FSP) Panel Report (Tang et al):
http://science.energy.gov/~media/fes/fesac/pdf/2007/Fesac_fsp_report.pdf

2007 Fusion Simulation Project Workshop Report (Kritz and Keyes):
http://science.energy.gov/~media/fes/pdf/about/Fusion_simulation_report_may_2007.pdf

2009 Scientific Grand Challenges: Fusion Energy Sciences and The Role of Computing at the Extreme Scale (Tang and Keyes):
<http://science.energy.gov/ascr/news-and-resources/workshops-and-conferences/grand-challenges/>

2011 Final FSP Detailed Plans: <http://w3.pppl.gov/fsp/Overview.html>

Relation to DOE Goals and Mission

This initiative is very well aligned with DOE-FES's goals (<http://science.energy.gov/fes/>) and mission. It is obviously central to the first goal "Advance the fundamental science ... to *develop the predictive capability* needed for ... fusion energy". The second goal is on materials for a burning plasma, and the edge codes proposed here are essential to understanding the plasma input to plasma-material interactions, and to studying innovative ideas (including liquid metals and continuous vapor shielding) to handle the extreme power densities.

This also contributes to the fourth goal because it will develop "fundamental understanding of basic plasma science, including both burning plasma and low temperature plasma science and engineering, to enhance economic competitiveness...". Some of the codes developed for plasma-material interactions and the edge region of fusion devices can also be applied to low temperature plasmas. The edge region of fusion devices is computationally very challenging, and would benefit from the development of kinetic codes using advanced algorithms (such as recent advances in discontinuous Galerkin methods) that could be useful to a wide range of kinetic problems where the mean free path is long, including certain astrophysical plasmas, hypersonics, MicroElectroMechanical Systems (MEMS), and semiconductor nanodevices.

Furthermore, an Integrated Simulation Initiative would clearly be addressing the first challenge on p.9 of Synakowski's presentation to the FESAC panel on "The charge for advice on strategic planning", both in developing individual modules that predict in a fundamental way "transport, macro-stability, wave-particle physics, and plasma-wall interactions", and in integrating them together to study and predict their interactions. Such a whole-device simulation can be used to study high-beta, high-bootstrap regimes needed for steady state, and would incorporate plasma-material-interaction modules and so would also play a role in his second challenge involving materials and steady-state.

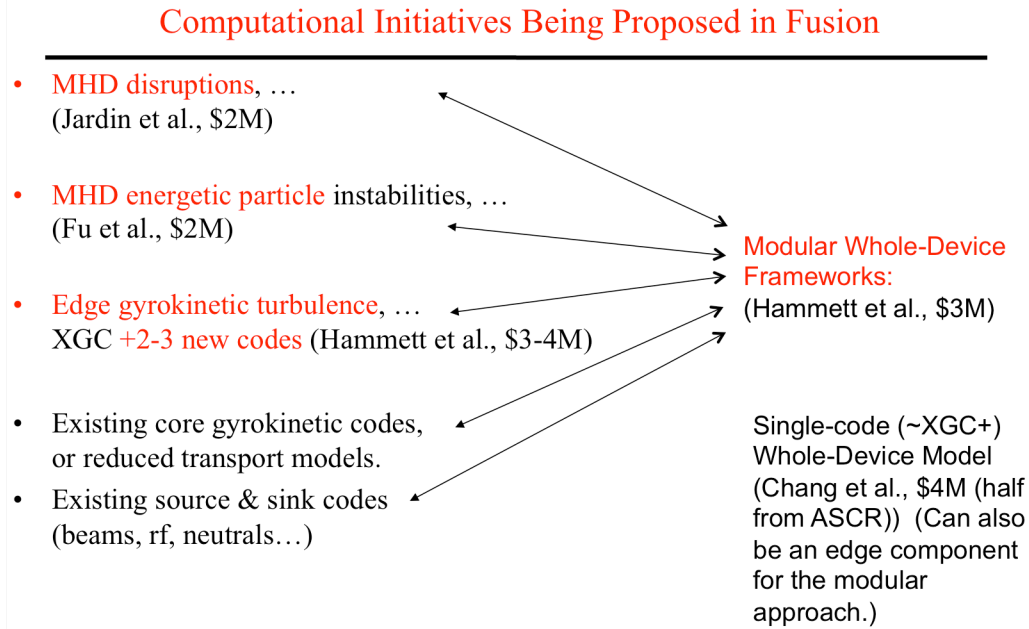


Fig. 5. Illustration of how the various computational presentations to the panel fit together into a single integrated fusion simulation initiative. The presentations addressed stand-alone research topics but were also about codes that would be components in an integrated whole-device framework. (The names above indicate the speaker on that topic.) Note that a range of degree of integration/modularity may be possible eventually, including a more monolithic “single-code” that handles turbulent transport and kinetic effects on MHD simultaneously (as discussed in Chang’s talk), once this capability is demonstrated. But it would still involve some degree of modularity (using other modules for things like RF heating and wall interactions). A top-level modular transport framework is needed both to allow switching to faster reduced transport models for some applications, and to parameterize the “slow manifold” (slowly varying profiles) used to make long-time-scale implicit projective integration practical. This concept of an Integrated Fusion Simulation Initiative has similarities to the earlier Fusion Simulation Project (FSP) proposal that Tang talked about (but with some differences and starting at a lower funding rate), and to Snyder’s talk on Validated, Highly Integrated Models and Simulations (VHIMS).