Fusion Simulation Program Advisory Committee (FSPAC) Report for the September 23-24, 2010, Meeting

Douglass Post (Chair), Allen Boozer, Brian Gross, Greg Hammett, Wayne Houlberg, Earl Marmar, Dan Meiron, Jon Menard, Mike Norman, Carl Sovinec, James van Dam

October 31, 2010

Charge to Panel September 23-24, 2010

(1) FSP Science Goals -- Regarding the current set of science drivers and associated IPT science development road-maps, please comment and advise on:

(a) the current vision for the associated code capabilities that can be expected in nearer-term (within 5 years) and longer term time frames;

(b) an outline of the key prioritized issues faced in carrying out such a program – highlighting integration; and

(c) the cross-disciplinary engagement of communities representing FES theory/computations & experiments and ASCR computer science and applied math.

(2) FSP Management & Governance -- Regarding the FSP Execution Plan, please comment and advise on:

(a) the organizational structure;

(b) the approach for dealing with the distributed project nature of the FSP;

(c) the decision-making process; and

(d) the flow of funds/resources from DoE-SC to the lead institution and to the collaborating research performers at other laboratories, universities, and industries.

Executive Summary

The Fusion Simulation Program Advisory Committee met on September 23 and 24, 2010, for the third of the planned four meetings and reviewed the state of the Fusion Simulation Program Planning Team activities for developing a comprehensive plan for the FSP. The Committee concluded that the Planning Team has made considerable progress on all of the critical issues since the last FSPAC meeting in March, 25-26, 2010. Specifically, focused sets of deliverables were defined for the Science Drivers, and the team was well along in the development of plans for producing codes to address the Science Drivers. The team has made a good start on developing a practical organizational structure that addresses several of the major managerial and organizational issues inherent in a highly challenging development effort. This organizational structure reflects the requirement for multi-disciplinary skills that are essential for a program that is to develop software for use on the present and next generation of supercomputers.

While there has been a lot of progress and the FSP planning team is on a successful path toward developing a solid plan, a great deal remains to be done before the final FSPAC meeting in the Spring of 2011, and before the final plan to DOE is delivered in mid-summer, 2011. These include: 1) resolution and clarification of many organizational issues; 2) development of detailed plans and roadmaps for the Integrated Science Applications and Components; 3) prioritizing the Integrated Science Applications and Components in case funding is inadequate to cover all of the proposed work; and 4) development of detailed plans and roadmaps for software engineering, collaborative development environments, and Validation and Verification.

The FSPAC notes that the FSP planning team has found a highly qualified replacement, Dr. Andrew Siegel from ANL, for Doug Kothe who was pulled away by Oak Ridge to lead the national program to develop a multi-physics US modeling capability for Nuclear Reactors.

For some responses to the charge, the FSPAC has identified particular issues that it deemed useful to express as specific findings and recommendations. In addition, at the end of the report the FSPAC listed a set of action items for reports and other information that should be sent to the FSPAC as soon as it is available. The final FSP meeting will be in the early Spring of 2011. The FSPAC will need to review a large amount of material to develop its final report. This will be facilitated if the could receive these reports and other information as soon as they are developed rather than just before or at the meeting.

FSP Science Goals – (1) Regarding the current set of science drivers and associated IPG science development road-maps, please comment and advise on:

Summary:

The general direction and progress on planning for the Science Drivers are encouraging. The vision and objectives were presented for each, though not yet in sufficient detail for the FSPAC to comment on the code deliverables for the next five years. The FSP team had just completed a community workshop at the time of the previous FSPAC meeting, and a report from the workshop was completed in June. Detailed reports from the Integrated Planning Groups (IPGs) that describe the plans and objectives for each science driver are being completed and will be made available to the FSPAC. Drafts of the Integrated Program Plan and Work Breakdown Structure and other relevant documents will be sent to the FSPAC members before the end of November so that the FSP will be able to provide comments as part of the final review in the Spring of 2011. The FSPAC recommends that "state of the science" reports and the plans for the Science Applications should be made available to the community as soon as possible. An objective of the FSP should be to publish this information in an archival journal during the early part of the FSP. This will provide community access to the FSP basis and plans.

Overall there was a lot of progress on identifying the major issues associated with the individual science drivers, and fleshing out the issues and priorities. The community involvement was extensive. The approach of defining science driver applications that would incorporate components to be shared by the Independent Science Applications where appropriate is sound.

The approach for validation and the level of interaction with the experimental community was also reasonable. The computer science aspects of the FSP plan still need to be more extensively developed. The FSP lost their lead for this area, Doug Kothe, to another national priority, but has acquired another very capable leader, Andrew Siegel from ANL. The outline for the software engineering and computer science approach that Dr. Siegel presented is a good start and follows a reasonable approach. The FSPAC looks forward to a comprehensive plan to be presented at the Spring 2011 review.

(a) The current vision for the associated code capabilities that can be expected in nearer-term (within 5 years) and longer term time frames;

While the FSP has outlined a list of objectives for each science driver, there is insufficient information for the FSPAC to comment in detail on the code capabilities that will be available in 5 years. The achievement of the stated goals would be a major advance in magnetic fusion energy science. The information needed to make a more detailed assessment will be available when a draft Integrated Program Plan and a draft Work Breakdown Structure are issued in October and November, 2010, respectively. The planning will not reach the point at which approximate cost estimates can be placed on the program elements until the end of the calendar year. Each portion of that information should be sent to the FSPAC members as soon as it is developed so that the FSPAC will have time to study and assess it well before the next (and final) FSPAC meeting in Spring, 2011. However, it is clear that the cost of the outlined program would far exceed the expected resources, especially considering that the funding will almost certainly be ramped up over a period of 4-5 years to accommodate the need to grow the program in a controlled and orderly fashion. Priorities will need to be set, and decisions made about the best way to allocate these funds among the various parts of the project, but this funding level should enable a quantum leap in our computational capabilities for fusion energy research. The FSPAC cannot assess the dates when the code capabilities will be delivered until the Spring FSPAC meeting when more detailed information is expected to be available.

The individual science drivers require fundamental advances in understanding to achieve their goals. Before embarking on the development of specific capabilities, the head of each science driver requires knowledge of (a) the present state of the area, (b) the critical issues in the area to progress of fusion energy sciences and which will be addressed, and (c) what code capabilities are required and what choices have been made to achieve these goals.

Finding: The FSP has done a lot of work to develop their plans for the Science Drivers. These plans will have a large impact on the US fusion program and on the OASCR community. The FSP Science Drivers and Components will provide focusing elements for fusion theory and experiments, as well direction for the FSP. It is therefore important that the information developed by the FSP be as widely available to the fusion and OASCR communities as possible.

Recommendation: The FSP should distribute their reports and assessments as widely as possible within the OFES and OASCR communities. A worthwhile goal would be publishing this material in an archival journal during the early phases of the FSP.

(b) an outline of the key prioritized issues faced in carrying out such a program –highlighting integration;

Over the last half year, the FSP planning team has increased its emphasis on the Science Drivers areas. This change addresses the programmatic recommendation of leading with the science. It is also apparent that the Science Drivers focus has given both the planning team and the fusion community, at least the segment involved in the Integrated Planning Groups (IPGs), more identifiable concepts of what the software products will be and how they will contribute to the MFE program. Both are positive outcomes of the shift.

Outlines of key prioritized issues, including components, integration needs, and validation, are being developed by the IPGs. This activity started in earnest with the March planning workshop, and the report from the workshop (June 28, 2010) provides a summary of the issues, existing capabilities, gaps and roadmaps for each of the Science Drivers. Since that workshop report was written, the IPGs have been preparing more complete reports due at the end of September, 2010, that will describe motivation, goals, components, framework requirements, validation, connections to other work (within and outside FSP), schedule and resources, and milestones. There is broad fusion community involvement in the IPGs, including representation from theory, experiment, and technology; however, stronger representation from math and computer science areas may be helpful. The Science Drivers reports are expected to provide the detailed information required to draw specific plans for FSP component and software integration. The FSPPAC agrees that the Science Drivers IPG activity, with its strong community involvement, is a good approach for prioritizing and planning.

Finding: While the FSP has made a strong effort to gather and assess the state of fusion science and identify science drivers from the whole fusion community, it has not succeeded in reaching everyone within the fusion community who is likely to have useful information to contribute. In addition, there needs to be a more comprehensive effort to provide a similar assessment of computational algorithms, software engineering and computer science "drivers".

Recommendation: The FSP should identify those in the fusion community who are likely to have useful information for the FSP, and actively solicit their views and perspectives. In addition, the FSP should mount a similar activity to the fusion science driver assessment for software development challenges including computational algorithms, software engineering, and computer science.

The presentations on components and software integration show that information from the IPGs is being used in the planning of these two 'groups.' Decision-making flowcharts on the acceptability and development needs for components start with needs for the Independent Science Applications (ISAs), where 'application' refers to the integrated software product. Moreover, it was stated that component planning will become more aligned with the Science Drivers by temporarily distributing its activities into the science drivers. This is workable as long as the component planning can be coordinated with the ISAs to ensure that the components can serve multiple Science Drivers (and other integrated studies) to avoid duplication of development effort. The presentation on software integration included specific information on framework needs that have already been communicated from the IPGs. Each component will need to be "interoperable" with multiple applications, which will be challenging given the variety of needs that are already apparent from the Science Drivers and the experience from the proto-FSPs. However, it will be essential if the FSP is to address many application areas.

Overall, the Science Drivers-centric approach to developing priorities for components and integration is on a productive path. The PAC looks forward to seeing the detailed reports and the specific plans for components and software integration that result as soon as they are completed. The next half year will be a critical time for assimilation of the large amount of information from the IPGs and for decisions on the number of Science Drivers activities that can be carried forward into detailed work plans.

(c) The cross-disciplinary engagement of communities representing FES theory/computations & experiments and ASCR computer science and applied math.

The team has considered in detail the role of validation in the FSP project. The validation effort has assembled a validation "best practices" document which lays out the principles for how the FSP will validate the various science drivers. For each science driver the document attempts to identify :

- The scientific issues and the regimes of validity covered by the simulation capability,
- The suite of physics components (codes) to be used or developed,
- Key physics for each component that should be confirmed by validation, especially areas of particular uncertainty in the underlying physical models,
- Important multi-physics predictions that need to be confirmed by validation,
- Key gaps in diagnostics or experimental capabilities,
- The connection between near-term capabilities and outcomes, and
- Areas where iteration with validation would improve fidelity

The FSPAC looks forward to receiving this document. This approach follows closely the validation best practices document that has been developed as part of the Sandia Lab's ASC work to validate their application codes originated by Bill Oberkampf and his colleagues[1]. The document defines a set of practices and a work flow to be followed in performing the validation. In some sense, one might say this is simply the scientific method, but in fact rigorous adherence to this approach is quite rare. Typically, various parameters associated with a complex physics code are simply tuned for best agreement with experiment in which case the code becomes a proxy for the experiment but with questionable predictive capability. Sandia's approach attempts to develop an "evidence file" for the validation activity which tries to take into account all uncertainties (numerical or physical) as well as try to assess the level of maturity of the available models. Most importantly, the approach strives for transparency so that future improvements can be readily identified.

The FSPAC feels that the use of the Sandia validation approach is well-founded, but perhaps the most important thing is that the team recognizes the necessity of a validation discipline. A thorough analysis of validation needs has been assembled through an impressive outreach activity to the experimental and theoretical fusion community. This analysis should be made available for review by the community to make sure that as many of the bases as possible have been covered. The notion of synthetic diagnostics is particularly well conceived and will allow for direct participation and exchange of data by experimental groups.

In all likelihood, the validation plan that has been assembled to date is very comprehensive and it is unlikely that the budget exists within to do all the required experiments and perform the necessary validation process for all the relevant ISA's. Further prioritization will most likely be required and the plan will certainly evolve.

Finding: Validation is an essential part of the FSP. Computational models are only models of nature. Without validation, there is no assurance that they can provide accurate predictions. Experiments are expensive. The FSP funds are very limited.

Recommendation: Coordination with experiments is appropriate, but we support the plans of the FSP team to not fund experiments (i.e. pay for development of diagnostic equipment, experimental operational costs, etc.) However, the FSP should fund the development of synthetic diagnostics, and participation by the FSP team in planning experiments and using the data to validate the FSP codes.

[1] W. Oberkampf and T. Trucano, *Verification and Validation in computational fluid mechanics*, 2002, Progress in Aerospace Studies, **38**, 209-272.

Regardless of the funding source, validation experiments are essential and need to be done as a partnership between the relevant experiments and the FSP team. Validation is not a passive activity with respect to experimental data. As has been the case with the IPGs and their predecessor organizations, the model builders need to be closely tied to the experimental teams, and involved in planning experiments.

Validation campaigns will likely be quite visible within the broader fusion community. In the interest of outreach to the broader community, the FSP should consider the routine inclusion of validation campaign leaders from outside the FSP team, and should attempt to rotate the leadership across institutions.

The connection with the OASCR applied math and computer science activities was addressed by Andrew Siegel who will serve as the deputy director for architecture; he will be responsible for software integration, code development and for the connection with the relevant ASCR research activities including the migration of codes and ISA's to future computing architectures. It is encouraging to note that availability of core OASCR tools will be facilitated through a cross-cutting team approach in the very near future. Each ISA will establish the requirements for computational issues such as code coupling, I/O etc. Where the relevant technology exists, the appropriate OASCR tools will be made available and supported.

However, the FSP will also face algorithmic issues for which further research and development will be required, such as multi-scale hybrid algorithms which will be required to couple aspects of the FSP such as core-edge coupling. This planning for future integration is among the most challenging aspects of the program and will require ongoing interaction with the ASCR community. There is again also the key issue of how such OASCR related activities are funded.

Even less clear at this point is how future computing architectures will be supported. The FSP has as one of its goals the eventual deployment of ISA's on future platforms. Typically, each time a new generation of hardware is developed, it has required disruptive changes to core kernels of codes in order to achieve peak performance. One way to manage this is to identify the core types of kernels required for the FSP and abstract these away so that the ISA developer need only interact with a well-designed interface that isolates the developer from the vicissitudes of low level programming to achieve performance. At this point it is not clear that such a strategy will be successful for a diverse project like the FSP, but some thinking along these lines is called for and this too should be a driver for further integration with OASCR research activities.

Finding: Parts of the fusion community do not feel engaged sufficiently in the FSP planning activities. While the FSP Planning Team has made major efforts to engage the fusion and, to a lesser extent, the OASCR communities in the planning process, it is a very challenging task. it is not surprising that some in the external community do not feel that they have much input to the process. The FSPAC encourages the FSP planning team to develop as many ways as possible for the community to engage with the FSP issues in an ongoing fashion. For instance, the Science Driver teams have recently posted their preliminary reports on the FSP website. Some of the FSPAC observed, that there is, however, no simple way for the community to comment on the reports, except by contacting the authors directly.

Recommendation: The team is urged to explore other methods that can help establish an ongoing dialogue such as review chits, a blog or forum, etc. Given the size and scope of the FSP, continual input from the community will both provide useful information and contribute to establishing community support for the FSP.

(2) FSP Management & Governance-- Regarding the FSP Execution Plan, please comment and advise on the current view of:

Summary

Clearly substantial progress has been made since the last FSPAC meeting in March 2010, but much remains to be done. The FSP team has defined a reasonable organizational structure that aligns responsibility and authority, embodies a workable decision making process, and is a good start on designing a process for distributing resources to the performing organizations with minimal transfer taxes. The latter point is very important, and the team deserves a lot of credit for this. There was less definition of how collaborative code development with non-collocated distributed teams will be supported. Concrete plans for this need to be in place by the Spring 2011 review.

The issue of how to ensure that necessary innovative research and development would be carried out to support the FSP was discussed by the FSP team. It was extensively debated by the FSPAC. Everyone on the FSPAC acknowledges that innovative research and development is needed for areas where existing knowledge and reasonable extrapolations of existing knowledge will not result in achievement of the needed capability. The question is to how to ensure that resources are allocated and managed for this necessary—but high risk and high payoff—innovative research and development.

Finding: It is clear that the R&D needed for a successful FSP ranged from implementation of known methods and algorithms to fundamental research into ways to formulate and solve fusion physics issues that have significant uncertainties. Several options were discussed, but no clear proposal was put forward for how to handle the latter situation. A specific issue that needs to be decided is how such activities will be planned and conducted, including the degree of oversight and control that the FSP director should have over this R&D and to what extent the funds allocated for this R&D should be counted in the FSP budget.

Recommendation: The FSP planning team and DOE need to jointly develop a clear proposal about how to plan and execute research and development of new methods and algorithms needed for the success of the FSP. This R&D will be high payoff but will also be high risk. That proposal should ensure that: 1) the FSP project have adequate influence on and sufficient control over the R&D to ensure that it is focused on and has a good chance for resolving the key issues and problems identified by the FSP and delivering solutions for solving those issues, and 2) that the groups charged with this R&D have the necessary freedom to pursue high risk R&D for which success cannot be assured.

As stated by the OFES theory program manager, the OFES Theory Program for fusion and computational science will also focus some of its work on issues identified by the FSP as crucial. It is also to be expected that the OFES Theory Program will be some of the major users of the FSP tools.

(a) the organizational structure

The proposed organizational structure is reasonable. It describes a clear set of roles and responsibilities and aligns authority and responsibility. Separating out the Independent Science Application (ISA) teams and giving them the primary role for defining requirements for the rest of the project is a sound choice. It gives the primary responsibility for the development of each major application to a single individual. The goal of the rest of the organization is to support the ISA teams. The major potential source of conflict is between the ISA teams and the Physics Components teams. Ensuring that this works well will require close cooperation and supervision by the two ISA and Physics Component leaders, conflict resolution by the Science Deputy Director, and good planning by everyone.

The Deputy Director for Code Architecture will play a key role, especially since many of the functions for that Directorate will be new to much of the existing plasma physics community. However, the FSP cannot succeed in producing high quality, sustainable and useable software without attention to those functions. This will also allow the FSP to actively conduct research and experiments in these issues that will provide valuable lessons for the DOE OS computational science and engineering community.

The planned budget of no more than $25M/year, enough for about 75 FTEs for ~ 18 functions, means that resources will be very constrained, and tasks and deliverables will need to be strongly prioritized. For the final plan, the FSP will need to hone the decision making process to ensure that the project team can define a scope that is consistent with the resources and perform the triage necessary to get a set of tasks that can be supported with the available funds. The plan will need to account for a ramp-up from an initial level to the asymptotic level of ~$25M/year. The FSP team will need to hold further discussions with DOE to ensure that the FSP director has the authority to manage the FSP, i.e. hire, fire, assign tasks, and reallocate resources as necessary to ensure the success of the FSP. The FSP planning team also needs to propose a process for staffing the teams.

(b) the approach for dealing with the distributed project nature of the FSP;

The IPGs are composed of computational personnel and plasma scientists, including representatives of the experimental community. This is a good start on developing collaboration with experimentalists to help define the goals and validation activities for the ISAs. However, the approach for successfully developing software collaboratively with distributed teams will need much more definition for the final report. Dr. Siegel has just joined the FSP team, and even though he has experience with these issues, it is not surprising that detailed plans are not yet available. There is, therefore, a lot to do before the Spring 2011 review.

The FSP team needs to develop strategies and concrete plans for collaborative development by distributed teams. They need to define how they will achieve the envisioned interoperability among the components and science applications. They will need to define how they will handle configuration management; coding and interface standards; software engineering practices; communication among non-collocated teams sufficient for collaborative code development that will alleviate the need for constant travel; document sharing; intellectual property issues; and how they will capture the lessons learned from the proto-FSPs and similar projects. Good approaches for these issues are essential for the success of the FSP, and for that reason, should be a key deliverable of the FSP planning phase. It will also be necessary to lay out a plan for the computer resources required for collaborative code development in a distributed environment.

(c) the decision-making process; and

In a project of this size and complexity, many decisions will need to be made on different timescales. These include staffing, process definition, resource management, component development, framework design, selection of validation experiments, etc. The FSP team has done a good job articulating the roles and responsibilities of 1st, 2nd, and 3rd level managers vertically and to some extent horizontally within the org chart. It is essential for the success of the project that authority and responsibility be aligned in these management positions. The science applications are fundamentally horizontally integrated, and therefore efficient mechanisms for decision making that spans the columns of the org chart also need to be articulated. The PAC recognizes the value of the "cross coupling" between the two deputy directors in this regard. However, more cross coupling may be required on a day-to-day and week-to-week basis than can be achieved in this "top down" approach. The proposed Science Committee and Software Review Board may serve this horizontal integration role. The FSPAC requests a better definition of the function, composition, and meeting frequency of Science Committee and Software Review Board and their role relative to the decision making process. In addition, the FSPAC suggests that the FSP team consider ways to get advice (primarily for internal consumption) on technical and organizational issues from independent external sources in addition to the FSPAC and its successor committee(s). The fusion and computational science community is eager to help with advice, codes, algorithms, etc., and most of this help is available for the asking.

The FSP planning has rightly focused on the integrated science applications (ISAs) to drive the overall program, and has identified the six core applications. Critical decisions are coming up very soon on the leaders of the ISA teams. Until the requirements analysis for the six proposed ISAs is completed, costs assessed, and priorities established, the number of ISAs that will actually go forward is not known. The FSPAC does not endorse any particular mechanism for filling the ISA leadership positions, but rather recommends the FSP management team approach this as a standard recruitment process.

This choice of the ISAs will have a ripple down effect on the physics components to be selected. How the specific needs of the ISAs will be coordinated with respect to the choice and design of physics components needs further clarification. There is a danger that there will be a loss of coherence and support for the ISAs if this is not considered carefully. To provide risk mitigation, multiple approaches for a given physics component should be considered (subject to resource constraints).

Finally, the FSP team has made it clear that science will drive the project, and the PAC fully endorses this. But one can ask what will drive the science? Experiments are a major part of fusion science, and therefore an argument can be made that the validation effort within FSP should be a major driver of the science. The FSPAC feels that FSP has a unique opportunity to bring the experimental and theory/simulation communities together in a way that has not happened before. However this will require not only adequate resources put into the validation effort, but also a visible and identifiable place within the organization chart. This could go as far as designating a V&V group in the ISA column as a separate box. This ISA could provide the interface to experiments and make the validation runs in conjunction with the relevant ISA group members. The ISA Head should be a person with a strong commitment to validation and a commitment to validation should be a key requirement for each ISA leader.

(d) the flow of funds/resources from DOE-OS-OFES to the lead institution and to the collaborating research performers at other laboratories, universities, and industries.

The FSP team as currently envisioned will have performers broadly distributed through the US program. The FSPAC understands that, under the plan, funding will be through a single contract with the lead institution, using flow-through subcontracts to fund work by team members at the other institutions. The FSPAC believes that this is a sound approach. We recommend that the FSP Director have authority to allocate funds as and where they are needed to ensure success of the project. With regard to sub-contracts, it will be critically important to minimize overhead costs associated with managing the flow-through of funds at the lead institution. In this regard, we compliment PPPL management, and the FSP planning team, for their vigorous pursuit of methods to achieve low overhead rates on these transfers.

Given the number and types of institutions involved in the FSP, contracting, conflict of interest, intellectual property, distribution rights, personnel recruitment and retention, succession planning, computer security, property management, and other administrative and legal issues will also need attention.

An Action Item for the FSP

The FSP final meeting will be in the Spring of 2011. To allow adequate time for the FSPAC to assess the FSP report and plans, it will be extremely useful for the FSP to push out the FSPAC members the reports and plans (even in preliminary form) as soon as they are available.

Specific reports that should be sent to the FSPAC before the Spring 2011 review include:

1. Report of the Proto-FSP Workshop (January 11-13, 2010 / ORNL)
2. FSP Planning Workshop Summary Report (March15-18, 2010 /Boulder CO)
3. Report of the June 28, 2010 Science Driver Workshop
4. Draft Work-breakdown Structure—November 2010
5. FSP validation “best practices” document--?
6. Planning Workshop (Feb. 7-11, 2011)