REPORT OF THE REVIEW COMMITTEE ON THE SciDAC ACCELERATOR MODELING PROJECT*

Prepared by

Ronald Davidson, Chair, Princeton Plasma Physics Laboratory
David Brown, Lawrence Livermore National Laboratory
Swapan Chattopadhyay, Jefferson Laboratory
Robert Gluckstern, University of Maryland
Bruce Hendrickson, Sandia National Laboratories
Richard Hubbard, Naval Research Laboratory
Ricky Kendall, Ames Laboratory, U. S. DOE
Jean-Francois Ostiguy, Fermi National Accelerator Laboratory
Jie Wei, Brookhaven National Laboratory

Review Conducted at

Lawrence Berkeley National Laboratory
January 15 - 16, 2003

***Internal review commissioned by the Lawrence Berkeley National Laboratory (LBNL) and the Stanford Linear Accelerator Center (SLAC).**

1. Executive Summary

The Review Committee for the SciDAC Accelerator Modeling Project, commissioned by the Lawrence Berkeley National Laboratory (LBNL) and the Stanford Linear Accelerator Center (SLAC), met in Berkeley on January 15 - 16, 2003, to review the technical progress and plans of the multi-institutional accelerator modeling project funded under the auspices of the Department of Energy's Scientific Discovery through Advanced Computing (SciDAC) initiative. The charge to the Committee is enclosed as Attachment #1 to this report, and the presentation agenda is provided in Attachment #2. The technical presentations to the Committee are available at http://amac.lbl.gov/techreview/. The materials presented at the review were in the following four technical areas:

1. Beam dynamics - Robert Ryne;
2. Advanced accelerators - Warren Mori;
3. Electromagnetics systems simulations - Kwok Ko; and
4. Computer science and applied mathematics - Esmond Ng.

The SciDAC Accelerator Modeling Project began its technical activities about 18 months prior to the review.

As a general remark, the Committee was impressed with the high degree of scientific productivity, technical quality, and effective coordination of this multi-institutional research project. Given the constrained level of resources (about $1.7M, annually), the balance among scientific elements (about 45% - beam dynamics, 45% - electromagnetics, and 10% - advanced accelerators) appears to be about right. While the direct SciDAC support for the advanced accelerator task area is relatively small, this important research activity is highly leveraged through strong coupling to ongoing related activities at UCLA, USC, SLAC and LBNL. The three main scientific modeling efforts are supplemented very effectively by about $650K of direct SciDAC funding in the computer science and applied mathematics areas, augmented by considerable technical support through ISIC. Given the high degree of technical progress and the compelling scientific opportunities facing the project, the Committee believes that all of the research elements would benefit from increased funding.

With regard to technical emphasis, a significant fraction of the computational effort is devoted to high energy physics applications, and less so to nuclear science applications. While some of the advanced computational tools being developed are applicable to both high energy and nuclear science accelerator systems, there is a need for an increased effort in certain high impact areas, such as electron cooling. In addition, while it is important that the project be responsive to near-term programmatic issues (such as beam-beam modeling for the Tevatron), the Committee believes, in the spirit of the original SciDAC charter, that the main focus of the project activities should be on: (a) developing advanced simulation tools for longer term scientific discovery, and (b) developing advanced design tools for next generation accelerator facilities. Of course the relative emphasis on near-term programmatic activities and the development of longer term simulation capabilities should be closely coordinated with Department of Energy priorities. In any case, it is very important to recognize that a sustained focus on near-term technical issues will detract from long-term tool development and opportunities for scientific discovery.

The linkages to the computer science and applied mathematics communities by the SciDAC Accelerator Modeling Project are commendable and merit continued strengthening. The linkages in the electromagnetics area are particularly strong and are having a large impact on this important activity. The linkages in the beam dynamics area are also significant and continue to increase, which is expected to have a similar large benefit to the beam dynamics project. Linkages in the advanced accelerator area should be further strengthened, e.g., through stronger UCLA interactions with ISIC activities in software engineering and language interoperability. Visualization is a common area where all of the technical tasks could benefit from an increased effort. Finally, as a general remark, the SciDAC Accelerator Modeling Project should develop a clear plan for documentation and distribution of codes, including efforts to make the codes modular and portable, and thereby available to the larger technical community.

The SciDAC Accelerator Modeling Project greatly benefits from strong scientific linkages to related accelerator projects at UCLA, SLAC, LBNL, other U.S. institutions, and through targeted international collaborations and expenditure of LDRD funds. This high-leverage situation is highly commendable, and is well documented in the presentations to the Committee. An increased effort to capitalize on potential resources in ISIC would also likely be of further benefit to the SciDAC Accelerator Modeling Project.

Finally, the following are some suggested administrative opportunities to strengthen the SciDAC Accelerator Modeling Project as an integrated national activity:

1. Establish a standing Program Advisory Committee (PAC) that would meet on a regular basis (about once per year) to review technical progress and future plans;
2. Schedule weekly conference calls of the principal investigators and task area leaders to identify and resolve technical and administrative issues affecting the project; and
3. Schedule monthly videoconference seminars/colloquia featuring recent project advances in beam dynamics, computer science, electromagnetics, and advanced accelerators.

As a general remark, Program Advisory Committees often provide valuable technical advice as well as a 'fresh' perspective. They also become part of an informed, knowledgeable constituency.

The following sections provide additional details supporting the findings and recommendations summarized in this Executive Summary.

2. Computer Science and Applied Mathematics

Success of the SciDAC Accelerator Modeling Project will depend on further developments in applied mathematics and computer science. Among the many challenges are the following:

a. The need for faster eigensolvers for complex symmetric matrices;
b. Efficient parallel algorithms for particle-in-cell (PIC) and Poisson solvers on unstructured and adaptive structured grids;
c. Improved numerical discretizations for electromagnetics on unstructured grids;
d. Mesh generation for complex CAD models;
e. Visualization of particles and fields in three dimensions; and
f. Better load-balancing methodologies.

These technology areas will require a mixture of research and development to address the mathematical and computer science needs of the SciDAC Accelerator Modeling Project. Additional resources may be needed to address all of these needs adequately.

For several of these technical challenges, the project is making very good use of the complementary activities in the SciDAC ISIC projects. The task on Electromagnetic Systems Simulations is to be commended for its high level of interaction with and exploitation of SciDAC ISIC activities. This successful leveraging has greatly increased the productivity of the electromagnetics project. The ties between the Beam Dynamics task area and the APDEC ISIC are also good, but don't appear to be as comprehensive. The Advanced Accelerator task area appears to be only beginning to profit from the APDEC ISIC. For all of the task areas, further efforts to take advantage of ISIC activities are strongly encouraged.

The parallel framework for particle-in-cell (PIC) codes that is being developed under the Advanced Accelerator activity is impressive. By developing and deploying general tools for a common computational kernel, this activity supports this specific SciDAC project while also helping the larger community. The plan to make use of more sophisticated load-balancing appears to be a good one. Efforts to extend some of these techniques to support unstructured grids would also be welcome. Interactions with the ISIC activities that are focused on software engineering and interoperability might be of value. Specifically, the common component architecture forum may be able to provide valuable ideas for building libraries that are easier to adapt and use.

The Committee believes that the success of the Accelerator Modeling Project would be greatly enhanced by a clear plan for the documentation, deployment, support, and long-term maintenance of the software modules being developed. The Committee also believes that this would be facilitated by an increase in funding. But even without such an increase, the sophisticated tools being developed in the project need to be made available to the larger technical community. The Committee recognizes that a decision to move forward with an 'open source' has disadvantages as well as advantages, and therefore the Committee is neutral on the precise distribution mechanism to be adopted. However, it is essential to develop detailed plans and begin to make measurable progress.

3. Beam Dynamics

The Beam Dynamics component of the SciDAC Accelerator Modeling Project as reported at the review represents a mature activity whose origin predates even the predecessor Grand Challenge project. This historical background and maturity also explains the current composition and mix of topics and investigators, which are still evolving. The principal investigator in this effort has been a prime mover in the national initiative involving computational beam dynamics and accelerator physics and is to be commended for spearheading this national effort, which has resulted in this important beam dynamics activity that benefits the entire accelerator community. The overall value and quality of the work, the value of the technical connections made with the computer science and large-scale computing communities, and the networking among different theoretical and computational beam dynamics groups at several institutions made possible by this effort --- all deserve to be commended as outstanding. Having made these general comments, the Committee wishes to make more detailed remarks in the following specific areas: (a) technical quality and productivity; (b) utility, programmatic balance and relevance; (c) integration with the applied mathematics and computer science communities; and (c) leveraging the project by other means.

The technical quality of the Beam Dynamics simulations is of the highest order, drawing upon the latest in algorithms from both the physics and the applied mathematics and computer science communities. The split-operator approach, the shifted Green's function algorithm, the parallel particle-based simulations to achieve load balancing, the Lie-algebraic operator methods via 'maps', the particle manager concept and techniques, the parallel Poisson solvers and statistical methods --- are all state-of-the-art simulation and applied mathematics tools that have been very effectively employed to solve large-scale computational problems in beam physics. The beam dynamics group is very well connected with the NERSC computer science group at LBNL, and benefits synergistically from this strong association. Of particular note is the effort made in the area of visualization and display of data, which is of great importance for displaying complex data and patterns in transparent ways. The Committee is very pleased to see a directed effort in this often-neglected area which can have a high impact in the communication of science, both for technical and non-technical purposes. The Committee recognizes the high productivity of the beam dynamics group through presentations at scientific meetings and technical conferences. We encourage the group to be proactive in technical publications not only in specialized scientific journals, but also in high-impact, cross-disciplinary journals such as Science and Nature, to give this field much needed exposure to the broader scientific community.

With regards to utility, programmatic balance and relevance, a few comments are in order. Given the limited time and funding constraints, there is no doubt that the beam dynamics group has been delivering high quality research of great utility to the accelerator community. It is also true that although the project is nominally a joint high energy physics and nuclear physics SciDAC project, it is funded primarily by high energy physics, which is also reflected in the choice of topics - the effort so far has been more on high energy physics applications and less on nuclear physics applications. There are three generic beam dynamics areas that merit advanced simulations - high beam intensity, low temperatures, and high fields. These three areas of beam dynamics match nicely into various areas of high energy physics and nuclear physics projects in the context of HEPAP and NSAC priorities (and BESAC as well). The predominant emphasis on space-charge effects (Poisson solver) and beam-beam issues addresses a large part of the high intensity area for both high energy and nuclear physics. However, aside from collective near-field effects at high intensity, there are also collective far-field effects such as coherent synchrotron radiation which are highly relevant to many accelerator projects in high energy physics, nuclear physics, and basic energy sciences. These also offer an attractive avenue for important scientific discoveries not possible otherwise. The same can be said about the area of low temperatures - the possibility of "frozen" or crystalline beams as a new state of matter can be investigated cost-effectively by high-fidelity simulations before embarking on the expensive construction of specialized machines. Simulations of the physics of low temperature beams by electron cooling, intra-beam Coulomb scattering and various stochastic noise-induced Langevin processes will also open up new vistas in beam physics relevant to other fields and future accelerator facilities. Finally, the physics at extremely high fields will require nonlinear electromagnetic simulations relevant eventually to laser, plasma and atomic processes that can result in new levels of scientific understanding of particle-radiation interactions in the collective regime. We recommend, building on the outstanding performance of the beam dynamics group so far, that the group adopt a more balanced research portfolio aiming at horizons of new knowledge and discovery over the long term, rather than focusing mainly on the details of large-scale computations relevant to large machines. It would seem to be most effective to devote about equal efforts to long-range tasks that are not easily accomplished by individual research groups, and to the more detailed machine-oriented, short-term tasks. Examples of useful applications are Vlasov-Fokker-Planck solvers beyond one dimension that can be used to investigate ultracold states of matter in directed motion. To this end, it would be useful to form a physics steering/advisory committee that would help to maintain a forward-looking, balanced portfolio at the frontier of discovery.

The effective integration of the beam dynamics activities with the applied mathematics and computer science communities is laudable. There is no doubt that the group is thoroughly conversant with the state-of-the art computational algorithms and software issues through the terascale accelerator modeling support provided by ISIC as described in Esmond Ng's presentation to the Committee. The visualization, particle-in-cell, and statistical methods are among the most advanced to date, and these have the promise of evolving at a faster rate in the future. The Committee encourages the beam dynamics group to further strengthen the synergistic relationship with the computer science and applied mathematics communities as much as possible.

The leveraging of the beam dynamics activity through association with other resources is commendable. This group is embedded in one of the nation's premier beam dynamics groups, the Center for Beam Physics, which houses outstanding scientists of high stature who are readily accessible. The scientists throughout the accelerator and fusion research division at LBNL as well as the national accelerator physics community are available as important resources. The proximity of the fusion groups at LBNL and LLNL, and the Heavy Ion Fusion Virtual National Laboratory's beam dynamics group headed by Alex Friedman, makes the connection with plasma physicists direct as well. Such associations add considerable value and status to the project, and the beam dynamics group is to be commended for recognizing and leveraging these resources effectively.

Members of the Committee commented on the issue that the team's products and tools should be made widely available to the general accelerator community beyond the direct collaborators in the SciDAC Accelerator Modeling Project. It is recommended that code documentation and user interface be strengthened. This applies to the beam dynamics packages as well as the electromagnetics and advanced accelerator packages. As the SciDAC activity shifts focus over time, the importance of portability, modularity and documentation of the more mature simulation products cannot be over-stated, as these tools will be of outstanding practical value to a broad segment of the accelerator community.

In summary, the beam dynamics component of the SciDAC Accelerator Modeling Project represents a world class activity. The group's research has been constrained by funding, but the large degree of synergy and leveraging achieved through collaborations with outstanding research groups throughout the community is commendable. The group's achievement in interacting with the applied mathematics and computer science communities, and utilization of advanced visualization tools also deserves special recognition. Overall, the beam dynamics group merits an *outstanding* rating.

4. Advanced Accelerators

The modeling in the Advanced Accelerator area covers plasma-based accelerators, including the plasma wakefield accelerator (PWFA), the laser wakefield accelerator (LWFA), and laser-plasma electron injection schemes. Plasma-based accelerators hold considerable promise for the future of high energy physics because of the extraordinarily high accelerating gradients that can be produced.

Although the Adavanced Accelerator portion of SciDAC Accelerator Modeling Project has a modest level of direct SciDAC funding, the productivity and technical quality of the research is very high. The SciDAC modeling effort has high visibility in the plasma-based accelerator community and impacts most of the major research areas currently being investigated in this field. Substantial recent progress has been made, particularly in the use of the new reduced physics particle-in-cell model QuickPIC. The close connection with major experiments at SLAC and LBNL is a major strength of the program.

The accelerator concepts studied in the SciDAC Advanced Accelerator effort are generally believed to be decades away from having an impact on HENP accelerator systems, but there are several important immediate and near-term applications. The effort on modeling the electron cloud instability could obviously impact a number of current experiments. The PWFA afterburner concept, which involves doubling the energy of the SLAC or NLC in a plasma acceleration stage, involves substantial technical risk but could potentially be fielded in a relatively short time. Compact GeV-class LWFA systems and low-emittance all-optical electron beam injectors could also be available in the near term.

The SciDAC program has promoted a strong collaboration among researchers who have not previously worked together in the advanced accelerator area. This has led to substantial benchmarking efforts between the two primary particle-in-cell (PIC) simulation tools, OSIRIS and OOPIC/VORPAL. The linkage to the computer science and applied mathematics communities is in its early stages. Viktor Decyk (UCLA) has been the primary resource for these issues in the advanced accelerator community, and his development of the parallel framework for PIC codes is very impressive. The Committee believes that his involvement in the Common Component Architecture (CCA) forum would also be useful. Although the parallel particle-pushing algorithms for PIC seem well in hand, there are other computational issues where the involvement of the APDEC ISIC could be useful. APDEC is working on fusion plasmas and hydrodynamics as well as accelerators, and this group may be able to use that experience to contribute to other aspects of plasma-based accelerators. It is also evident that improving visualization tools, particularly for PIC codes, could be an area where collaboration with both the computer science community and other accelerator research areas would prove fruitful. In particular, PIC and laser propagation codes share many of the visualization issues encountered in the beam dynamics area.

The Advanced Accelerator effort is highly leveraged by funds from the DOE Advanced Accelerator program, ongoing related activities at UCLA, USC and SLAC, internal LBNL funding, and DOE SBIR funding. The OOPIC/VORPAL code has applications to a broad range of non-accelerator plasma physics areas, and it has a long history of support from other sources. This leveraging has contributed significantly to the success of this program and will presumably continue at a high level.

5. Electromagnetic Systems Simulations

The Electromagnetics Systems Simulations component of the SciDAC Accelerator Modeling Project is unquestionably very mature. The volume and the quality of the work presented at the review very favorably impressed the Committee. In addition, the general sentiment of the Committee was that the integration and interaction by the electromagnetics group with the computer science component of the program (ISIC) is exemplary. In many ways, the electromagnetics effort is a role model that other components of the program should try to emulate.

This being said, the Committee realizes that the electromagnetics activity represents the continuation of an effort that has been on-going at SLAC for a number of years. Furthermore, mesh generation and the solution of large, sparse eigenvalue problems are very well-established research areas in computer science and numerical analysis. This certainly provides a natural framework for collaboration, which may not be as immediately available for the other task areas.

The Committee applauds the customer-oriented approach and the efforts of the electromagnetics group to apply its expertise to solve specific problems of relevance to accelerator operations and research. Nevertheless, it is felt that more attention should be paid to the documentation and dissemination of the codes. Note that this does not necessarily imply that all software should be made publicly available. One possibility would be to release and document subsets of software with generic usefulness, with a delay of six months to a year. The approach taken by the electromagnetics group of using an SBIR-funded company to write a user interface and provide commercial support is plausible. While this may be an acceptable approach for large organizations, the licensing costs may be high and effectively exclude many universities, graduate students and researchers who cannot justify the extra expenditure to access the software. The Committee believes that regardless of the specifics, it would of great benefit to the SciDAC Accelerator Simulation Project to develop coherent distribution and documentation policies that balance the interests of the code developers with those of the wider technical community.

Research on the computation and characterization of the modes of detuned accelerator structures appears to be in excellent shape. The accurate prediction of the transverse wake in damped, detuned structures requires one to take into account mode losses, especially since these losses can induce enough detuning to affect the effectiveness of the overall wake mitigation scheme. The modeling of mode losses results in large complex symmetric systems (as opposed to Hermitian, which is implied when the field energy is conserved) for which efficient algorithms are not well developed. This is an important problem of generic interest and the Committee believes that advances in this area would impact not only accelerator physics, but other fields as well.

The electromagnetics presentation to the Committee emphasized frequency-domain finite elements codes. Although some impressive results were presented, the status of the time-domain code (Tau3P) was discussed only briefly. There appear to be some issues regarding long-term numerical stability for which a satisfactory solution has not yet been identified and validated. A proposal was also made for the development of a new implicit time-domain finite element code, presumably to overcome geometric limitations inherent to the of the path-integral algorithm used in Tau3P. It is believed that Tau3P and its eventual finite element counterpart would benefit from stronger interactions with the numerical analysis community. These interactions appear to be in progress; it is recommended that the technical status and challenges of the time-domain codes receive more emphasis in future reviews.

Simulation results addressing the important problems of breakdown-induced damage and dark currents were also presented to the Committee. These simulations provide a high level of insight not achievable experimentally. This work constitutes a prime example of the potential of advanced computing to advance the basic understanding of complex physical systems.

In conclusion, Committee believes that the electromagnetics group has carried out and continues to deliver research of the highest quality. Collaborations with the computer science and numerical analysis communities are well established. A coherent policy to document and distribute all or part of this advanced simulation capability to the wider technical community would be very beneficial.

ATTACHMENT #1

Charge to the Internal Review Committee for the SciDAC Accelerator Modeling Project

As part of the initiative, Science Discovery through Advanced Computing (SciDAC), the DOE Office of Science has approved a major multi-institutional project to develop terascale computational tools for the modeling and simulation of accelerator systems. . The host laboratories for the project are LBNL and SLAC under co-PIs Robert Ryne and Kwok Ko respectively. The scientific software to be developed under this project falls into three broad focus areas:

1. Beam Dynamics (BD) covers the dynamics of intense charged particle beams propagating through a variety of accelerator structures that make up an accelerator complex.
2. Electromagnetic Systems Simulation (ES) concentrates on parallel tools for the design, analysis, and optimization of complex electromagnetic components and systems in accelerators.
3. Advanced Accelerators (AA) address the issues of laser-driven and beam-driven plasma-based accelerators.

As the project has now completed half of its three year duration, the Project Steering Committee is commissioning this internal review of the progress of the project team toward meetings its technical objectives. In particular, we would like the reviewers to evaluate the following issues:

4. Technical quality and productivity of the project elements and the project as a whole both from an accelerator physics, and computational science and engineering perspective,
5. Utility, programmatic balance, and application of the project products to high energy and nuclear physics, in the context of HEPAP and NSAC priorities, R&D efforts focused on existing machines, and evolving near-term programmatic issues such as beam-beam modeling of the Tevatron
6. Effectiveness in linking the accelerator modeling efforts to the applied math and computer science communities and integrating the SciDAC components of these communities in collaborative team efforts to accomplish notable accelerator science successes that otherwise would not have been possible.
7. Leveraging of project impact via association with other non-HENP accelerator modeling programs and institutionally supported research at the participating institutions, especially at the host laboratories.

We ask the Committee to summarize their comments in a brief written report to the Steering Committee one month after the completion on the review.

ATTACHMENT #2

SciDAC Accelerator Project Technical Review
LBNL - January 15-16, 2003

Building 71, rm 264 - Albert Ghiorso Conference Room
Break out in 71-180 - Library

Agenda

Wednesday, 15 January

8:20 AM	Continental Breakfast available for Review Committee
8:30- 9:00	Executive Session
9:00- 9:30	Introductions and Overview [Kwok Ko and Rob Ryne]
9:30- 10:45	Bean Dynamics [Rob Ryne]
10:45-11:15	*Break - group refreshments in break-out room*
11:15-12:30	Advanced Accelerators [Warren Mori]
12:30- 2:00	Working Lunch [Committee]
2:00- 3:15	Electromagnetics [Kwok Ko]
3:15	*Refreshments available in break-out room*
3:15- 4:30	Comp Sci, Appl Mathematics, Connections to ICISs [Esmond Ng]
4:30- 5:30	Discussion, Planning for Thursday Session, Executive Session

Thursday, 16 January

8:20	Continental Breakfast available for Review Committee
-	Follow-up Questions and Discussion
10:45	*Break - group refreshments in break-out room*
-	Executive Session
11:00	Close Out
12:00	Adjourn

Review Committee:

Ron Davidson, PPPL - Chair
David Brown, LLNL
Swapan Chattopadhyay, JLAB
Bob Gluckstern, UMD
Bruce Hendrickson, SNLA
Dick Hubbard, NRL
Ricky Kendall, Ames Lab
Jean-François Ostiguy, FNAL
Jie Wei, BNL