

December 14, 2004

Dr. N. Anne Davies  
Associate Director  
Office of Fusion Energy Sciences  
SC-J204/50 Germantown Building  
U.S. Department of Energy  
1000 Independence Avenue, SW  
Washington, D.C. 20585-1290

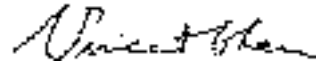
Dear Dr. Davies:

The attached White Paper on Fusion Computing was recently presented to the Theory Coordinating Committee (TCC). The TCC endorses the following main points, which we hope OFES will consider:

- Cost-effective computational research needs a mix of computing hardware resources, ranging from high-end, centrally located supercomputers (“capability” or cycles/sec) to low- and mid-range clusters (“capacity” or cycles/year). Progress in Fusion computing is presently most limited by computing capacity. Productive use of human resources requires minimal job wait times, particularly while complex codes are being developed, which requires many small jobs. Properly allocated clusters could provide a very cost-effective increase in computing capacity and a significant reduction in wait times.
- Largely for historical and funding reasons OFES fusion computing, unlike other Energy Research divisions, relies almost exclusively on supercomputers for nearly all its parallel computing needs. The two main such supercomputers are Seaborg at NERSC, where Fusion is the largest user, and Cheetah at CCS. Of the total usage on Seaborg, 40% is done on fewer than 256 of the 6656 Seaborg processors-work that could be done on large clusters if they were available.
- Seaborg (despite some relief from Cheetah) became so over allocated last summer that the loading efficiency [run-time/(run-time+wait-time)] dropped to 10%, well below the nominal value of 50%. The long wait-times were a great drain on scientific productivity and made the facility unusable for most Fusion researchers.
- The situation last summer resulted from two factors. The first was the NERSC policy of providing “free” or reduced charges to the largest jobs in order to meet the NERSC milestone of having a certain fraction of its jobs use more than 1/8 of the machine. This effectively over-allocated Seaborg. The second factor was the decision of making large awards to a few “INCITE” projects that were not ready early in the year, and then providing these projects with front-of-the-line priority late in the year to use their large allocations over a few months. These policies essentially froze-out non-INCITE users.
- In the short term, we encourage OFES to urge NERSC to look for ways to reduce the wait-times and to increase the loading efficiency above 50%. Methods for doing so could include allocating less than full capacity by holding back a reserve, giving out allocations in quarterly amounts with small penalties for non-use, and requiring the INCITE awards to start running earlier in the year so that they don’t perturb the system as much, thus minimizing the need for special “priority” that distorts the system.

- We applaud the fact that NERSC now has plans to provide a mix of capability and capacity computing by installing a new 600+ processor cluster, Jaquard. We urge MICS to monitor the success of this approach and, for long-term planning, to consider the idea of a “mixed size cluster farm” center preserving the present economy of scale.
- We encourage OFES to provide funding for acquisition and maintenance of “local clusters” to theory groups that have staff and infrastructure to manage them.

Sincerely,



Vincent S. Chan, Chair  
For the Theory Coordinating Committee  
<http://web.gat.com/theory/tcc/>

cc: Dr. John Willis, OFES/DOE