

INCITE COV Report

Committee of Visitors Report

on

Innovative and Novel Computational Impact on

Theory and Experiment

Marsha Berger (chair)

Courant Institute, New York University

Outline

- COV timeline
- INCITE overview
- INCITE proposal review and allocation process
- Issues and Findings
- Recommendations

COV Charge

Excerpt from Dr. Orbach's charge:

*“To help the research communities tap into the capabilities of current and future supercomputers, the Office of Science launched the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program in 2003. I am very pleased with the response from the community and the numbers and the quality of the proposals we have been receiving. **Given the high visibility of the INCITE program and the importance of this program to national competitiveness, I now ask the ASCAC to form a COV to review the process used to manage the INCITE program. A report from ASCAC is expected by the Fall 2008 ASCAC meeting.**”*

COV Committee

Marsha Berger
Courant Institute, NYU

Roscoe Giles
Boston University

Cray Henry
High Performance Computing
Modernization Program, DOD

James Kinter
COLA, George Mason University

Dimitri Mavriplis
University of Wyoming

Gopal Shenoy
Argonne National Laboratory

COV Timeline

- Solicited input from
 - users
 - computer center Directors (ORNL, ANL, NERSC)
 - other DOE program offices (BES, BER, OFES)
- Read successful and unsuccessful proposals, progress reports, and reviews
- COV met April 23-24 at DOE HQ
- Draft report circulated, follow-up deliberations by phone and email

INCITE Overview

- Announced by the Director of the Office of Science in 2003.
- Outgrowth of HEC-RTF report, which defined capability and capacity computing, and recommended establishing Leadership Class computing systems.
- INCITE mission is to advance a small number of computationally intensive large-scale projects needing Leadership Class computing.

The INCITE program was conceived specifically to seek out computationally intensive, large-scale research projects with the potential to significantly advance key areas in science and engineering.

from INCITE web page

INCITE Overview

- Originally based at NERSC: 3 projects used 10% of the cycles.
- In 2004, the Office of Science received funding to create the LCFs at Oak Ridge (OLCF) and Argonne (ALCF).
- Starting in 2006, 10% of OLCF and ALCF added to INCITE. Multiple year proposals allowed.
- In 2007, 80% of OLCF and ALCF made available to INCITE (in response to OMB request). 10% of NERSC remained in INCITE. 5% of PNNL added in 2006.
- 10% of all cycles reserved for Dr. Orbach's discretionary use. 10% for Center Director's use.

INCITE Overview

- OLCF - CRAY XT4 (31K cores)
- ALCF - Blue Gene/P (32K cores), partners with IBM (#3 on Top 500 list; 537 TFLOPS/s peak)
- NERSC - Upgraded from Seaborg (IBM) to Cray XT4 (19K cores)
- OLCF and ALCF provide staff support to INCITE projects (one person for two projects). User feedback highly appreciative.

INCITE Overview

Tremendous growth in INCITE submissions and allocations. Off to a good start.

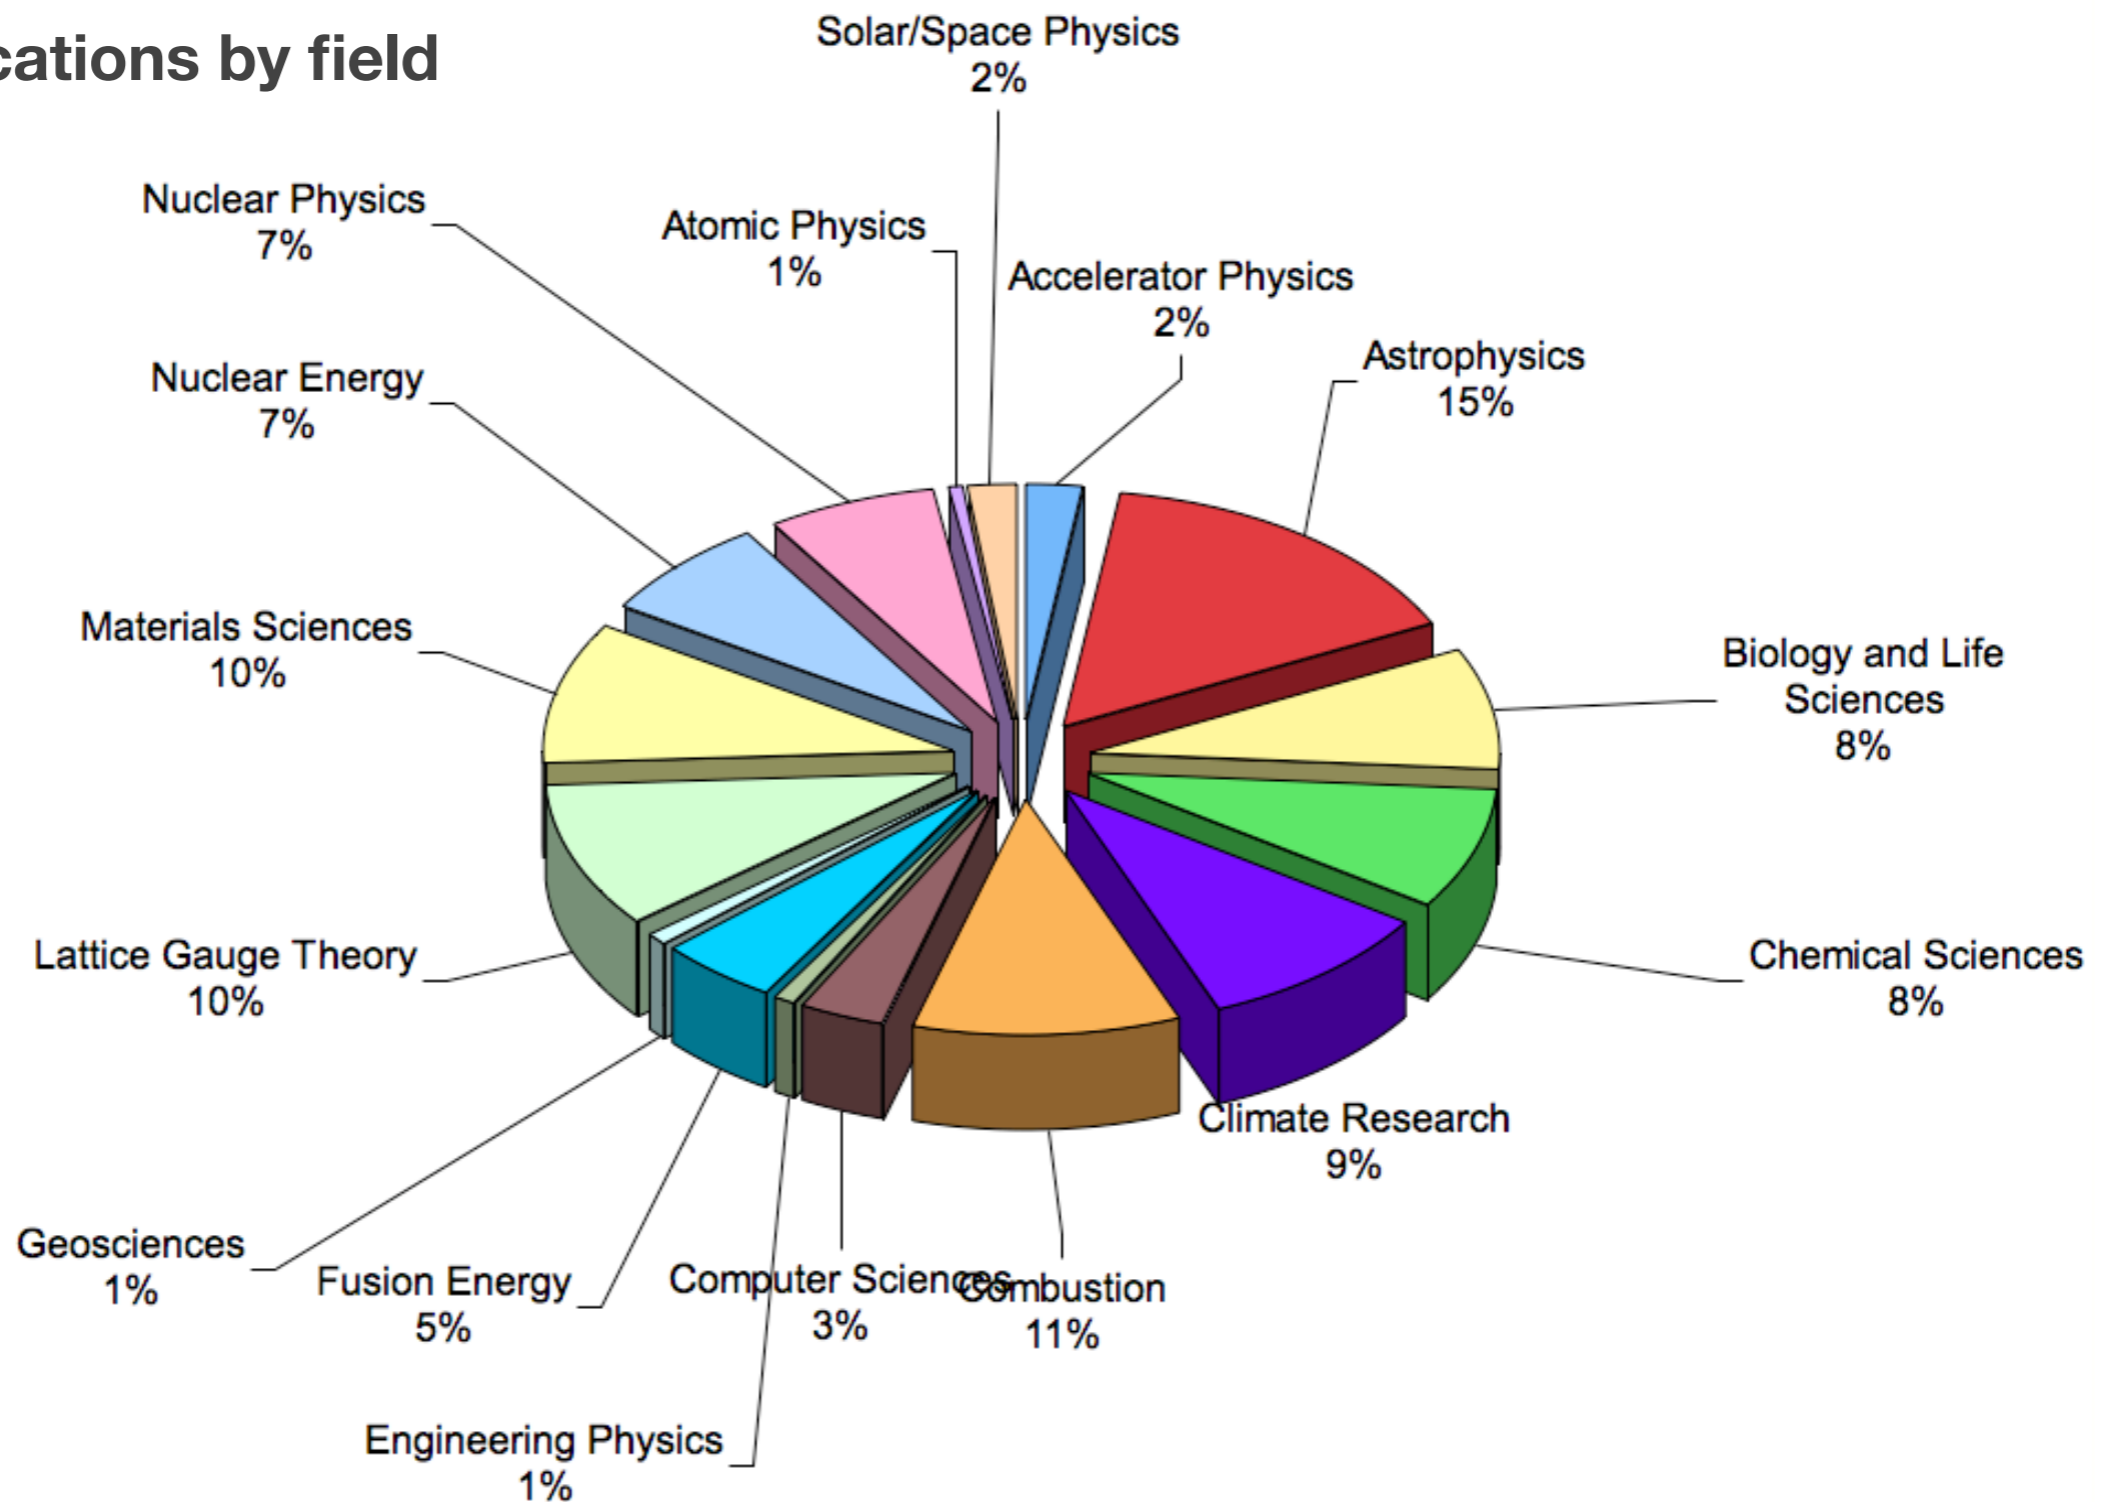
Year	Proposals Submitted	Proposals Approved	Processor-Hours Requested	Processor-Hours Allocated
2004	52	3	130M	5M
2005	23	3	28M	6M
2006	43	15	95M	18M
2007	107 (20 renewals)*	45	250M	95M
2008	112 (24 renewals)	55	600M	265M

*9 were INCITE renewals, 11 were previously allocated DOE renewals.

INCITE provides the largest national leadership class capability to the national science community and also high-end capability class support for the DOE community. DOE is to be commended for taking the lead and accepting the risk inherent in pursuing big computational science projects. (Finding 1)

INCITE Overview

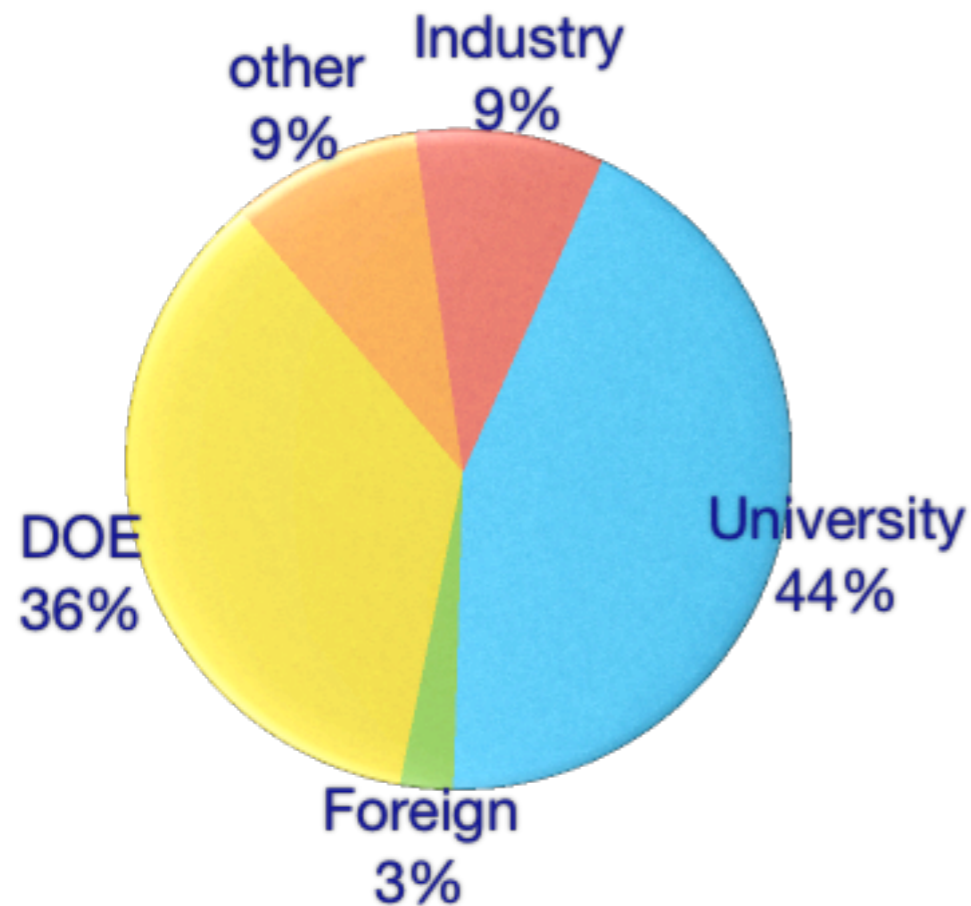
2008 allocations by field



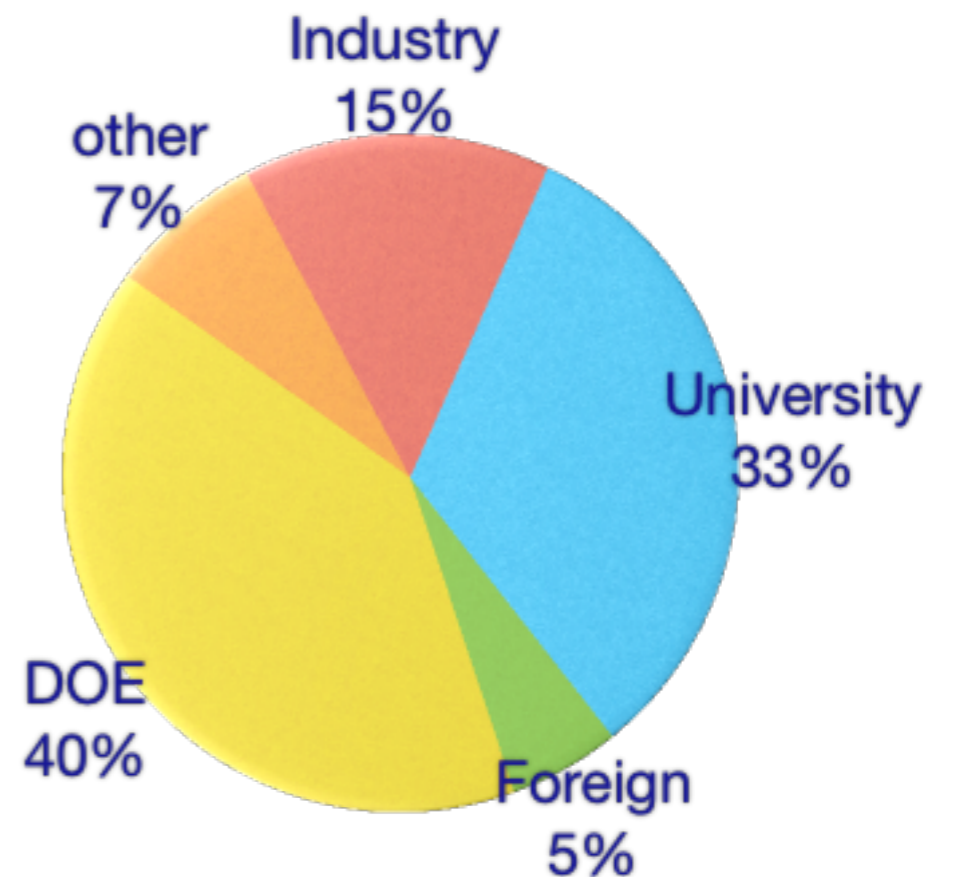
INCITE Overview

2008 distribution by type

Proposed



Allocated



- 16 of 20 INCITE SciDAC proposals received allocations
- 14 of 18 University recipients had DOE funding (DOE leads in high-end computing)

INCITE Review and Allocation Process

- INCITE open to national and international researchers, including industry.
- No requirement for DOE funding (or any other).
- Proposals receive scientific review and computational readiness review. Renewals (2nd and 3rd year of multiyear proposal) undergo only readiness review.
- SC program offices prioritize INCITE proposals
- INCITE program manager prepares recommendations for SC Director, who makes final decision.

INCITE Review and Allocation Process

- **Scientific Review Criteria: (scale 0-10)**

- Scientific and/or technical merit of project
- Appropriateness of proposed method or approach
- Competency of the personnel
- Reasonableness and appropriateness of proposed request for resources

For proposals rated 7 or above, provide justification of importance of the research to the scientific field.

- **Computational Readiness Review Criteria: (scale 0/1)**

- Reasonableness of request for computational resources
- Appropriateness of computational approach
- Technical readiness: performance and scalability

Renewals only receive Readiness Review

INCITE Review and Allocation Process

- **Scientific Review Process: (scale 0-10)**
 - Peer reviews: some mail reviews, some panel reviews
- **Computational Readiness Review Process: (scale 0/1, results averaged)**
 - Proposals should include performance and scalability data
 - Each LCF evaluates every proposal; 1 or 2 reviewers per proposal
 - Evaluators interact with proposers (new this year)
- **Allocation Process:**
 - SC program directors prioritize proposals
 - Renewals usually awarded full requested amount
 - INCITE Program Manager prepares initial allocations; final decision by SC Director

FINDINGS and ISSUES

- Tension between user communities: DOE and at large
- Process Issues
 - Non-uniform and non-transparent process
 - Perfunctory readiness reviews
- Number of Projects Supported
- Timing of RFPs
- Renewals (annual and re-competes)

Tension between DOE and Broader Community

FINDING 2: INCITE has become the umbrella program that encompasses both national leadership class computing and high-end capability computing that directly supports DOE science. There is tension between meeting the DOE science mission objectives for high-end capability class computing and the national objective of providing leadership class computing for the broader scientific community.

- LCFs are DOE's most advanced computing resource.
- Much DOE research calls for most advanced computing available
- “Double jeopardy” for SciDAC projects
- Program managers “wasted” resources by funding SciDACs without sufficient computing resources
- Projects distorted to “package” themselves for INCITE
- DOE computing has different characteristics

Process Issues

Reviewing Process:

- Different disciplines reviewed in different ways (hard to find reviewers)
- Little programmatic prioritization for projects outside DOE (hard to get outside input)

Computational Readiness:

- Not completely effectively in identifying projects ready to go on day one
- Binary decisions + averaging neglects relevant review content
- Some projects deemed not computationally ready got allocations
- LCFs review their own proposals, creating the appearance of conflict of interest.
- Different LCFS use different procedures.

Number of Projects Supported

FINDING 3: The INCITE portfolio has evolved so that there are a large number of projects taking up a smaller fraction of the facilities.

- 55 projects stretches the highly-skilled support staff
- Some projects receiving allocations are outside Leadership Class
- INCITE projects no longer the top users at NERSC (was 3, now 11 projects)
- Only 11 of 55 projects requested more than 10M CPU hours

FINDING 4: The INCITE program nicely balances the computing and support resources needed to effectively make use of complex computing environments inherent in leadership class supercomputers. The emphasis is appropriately split between operating the very large supercomputers and in providing direct project level support to the individual science team to improve their ability to make the best use of the supercomputer resources.

- If a project is highly rated scientifically but not ready, it may be referred to SciDAC center to help it get ready, or LCF's can work with it to get ready using discretionary time.

Proposal Timing

- INCITE proposals solicited annually
- LCFs felt more frequent solicitations would generate more even flow of new projects
- Users felt one year too long to wait to reapply if proposal declined, adversely impacts scientific progress
- Concern that increasing frequency doubles the workload for the program office

Renewals

- Review process almost always resulted in continuing project allocations, in spite of minimal previous year accomplishments.
- Second and third years of multi-year proposals receive requested allocation in almost all cases. In some cases justified since projects are more mature and their requests are presumably more reliable given their experience from past years.

FINDING 5: Because of the mixed objectives, the INCITE selection process attempts to balancing conflicting requirements and lacks a transparent set of selection criteria. While scientific merit and computational readiness are the primary criteria, the decision also considers unevenly collected project information and specific supercomputer system requests. The scientific reviews are accomplished differently for different disciplines making the comparison of scores problematic, and creating the opportunity for proposers to question the outcome. We found no evidence that the outcome was unfair or biased but user comments highlighted a concern.

RECOMMENDATIONS

1. The selection processes for leadership class and high-end DOE capability class computing should be separated. A significant portion, but less than half, of INCITE computational resources should be allocated to high-end DOE capability-class computing.
2. INCITE awards should be fewer in number and larger in size with the expectation of demonstrated concurrency across a very large number of cores. For projects deemed important but not ready, some resources could be reserved for development. Renewals should meet an achievement threshold below which projects are rejected, or referred to additional technical support on smaller platforms to make way for more promising new projects.
3. INCITE should continue to provide robust expert assistance to the science teams performing leadership class computing.

RECOMMENDATIONS (cont.)

4. Review Process:

- a. The selection process should be made as transparent and uniform across disciplines as practical. Selection criteria should be formulated and published. These should include scientific promise and importance, appropriateness of the computational technique, and potential impact on overall technological capability.
- b. When this has been accomplished, consideration should be given to increasing the frequency of INCITE calls for proposals, or at least staggering the annual call with other relevant calls.
- c. The computational readiness review process should adopt a more descriptive outcome, for example an overall grade (0-5) could be used. The readiness review would also benefit from a more systematic process such as a panel review performed by a group of computational experts from all the leadership class facilities as well as high-end computing experts from outside DOE.
- d. An appeals process for allocation decisions should be implemented.

RECOMMENDATIONS (cont.)

5. The COV recommends that in approximately 5 years a formal review panel be convened to assess the impact of the INCITE program. This would be done through examination of project final reports, publication records, and the assessment of other types of impact, and by collecting feedback on what worked and what didn't from past users. Such information needs to be collected more systematically to be able to measure more precisely the scientific impact of INCITE.