Summary of the 2008 ASCAC Balance Panel Report

International Competitiveness: Facilities for Participation Research for Leadership

Presented by: Robert Voigt *(Panel Co-Chair)*

Panel members: F.Ron Bailey, Vincent Chan, Jill Dahlburg *(Chair),* Michael Heath, & Charles McMillan

Report submitted to the ASCAC on 6 February 2008 for consideration as a:
Report of the Advanced Scientific Computing Advisory Committee

Charge

- Assess balance within ASCR between
  - Facilities, particularly LCF
  - Research
    - Near-term, with impact on ongoing applications
    - Fundamental, long term, higher risk
- Assumed background for charge:
  - Accelerate scientific discovery on most important questions
  - Maintain & enhance US leadership in scientific discovery

Balance Panel

- Convened week of January 14 at NRL
  - F. Ronald Bailey, CSC
  - Vincent Chan, General Atomics
  - Jill Dahlburg, NRL
  - Michael Heath, U. of IL
  - Charles McMillan, LANL
  - Robert Voigt, SAIC

Input to Panel

- Numerous DOE and other HPC documents
- Community input via a website
- Presentations
  - Michael Strayer
  - Fred Johnson
  - Jack Dongarra (U. of TN)

In all, over 100 documents

Fundamental Findings

- Leadership in advanced computing must be maintained
- HPC hardware is becoming a world-wide commodity
- Hence, access does not ensure scientific competitiveness
- It's the ability to make effective use of HPC that assures leadership
- Increased HPC complexity requires research in algorithms, tools, software

THE balance issue

Fundamental Recommendation

Concomitant investments in research must be made to ensure that investments in LCF realize their intended goals

In severe funding environments, balance may require delaying acquisition of LCF

The community does not benefit from unusable systems!

From here

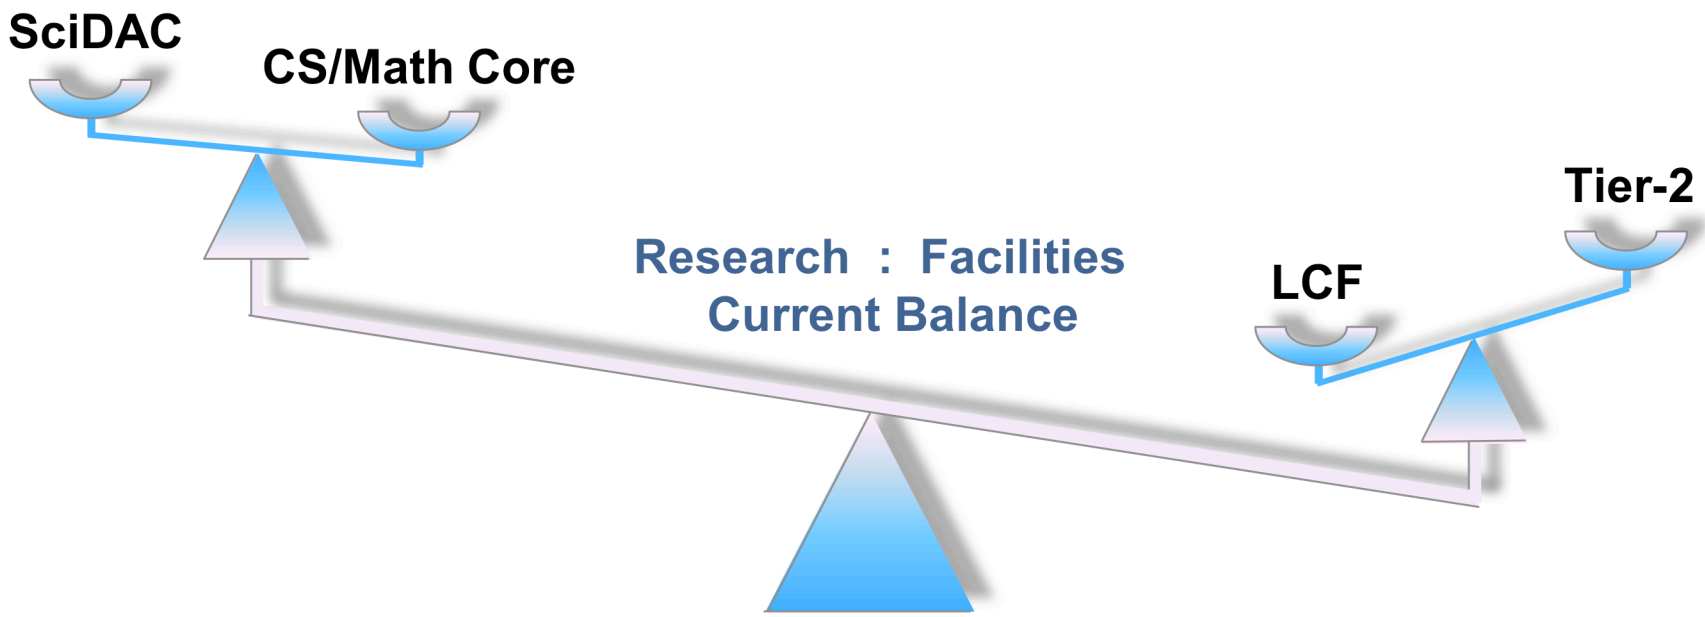

Current balance of programmatic elements in DOE SC ASCR.

To here

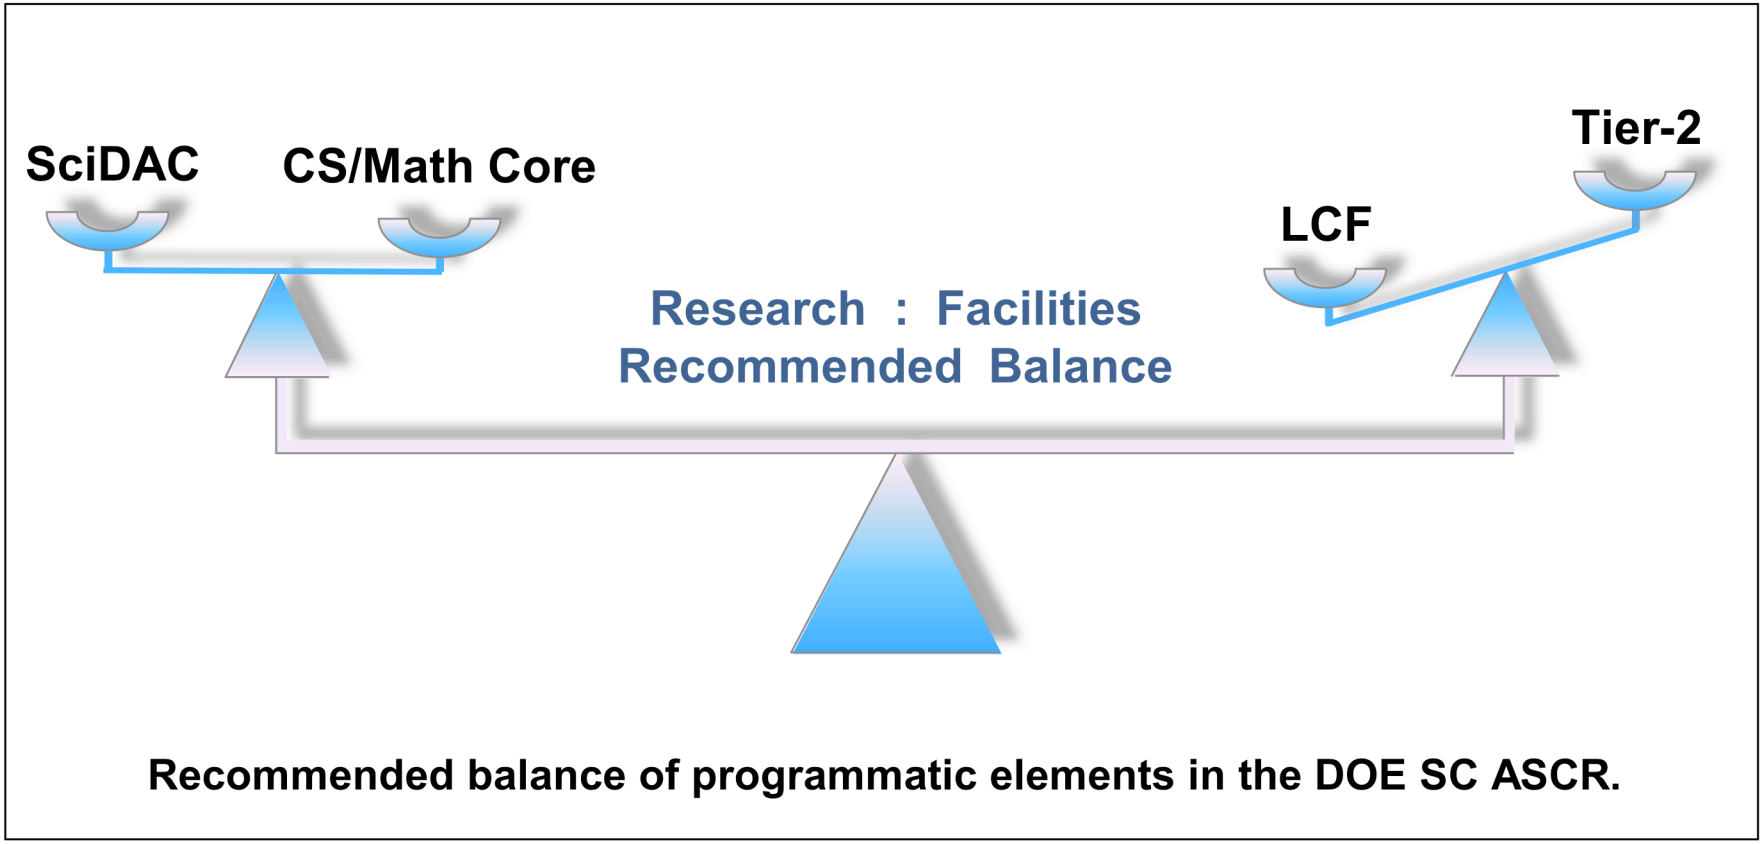

Recommended balance of programmatic elements in the DOE SC ASCR.

I.E., Restore the Balance

Advanced Scientific Computing Research Program Trends (details)

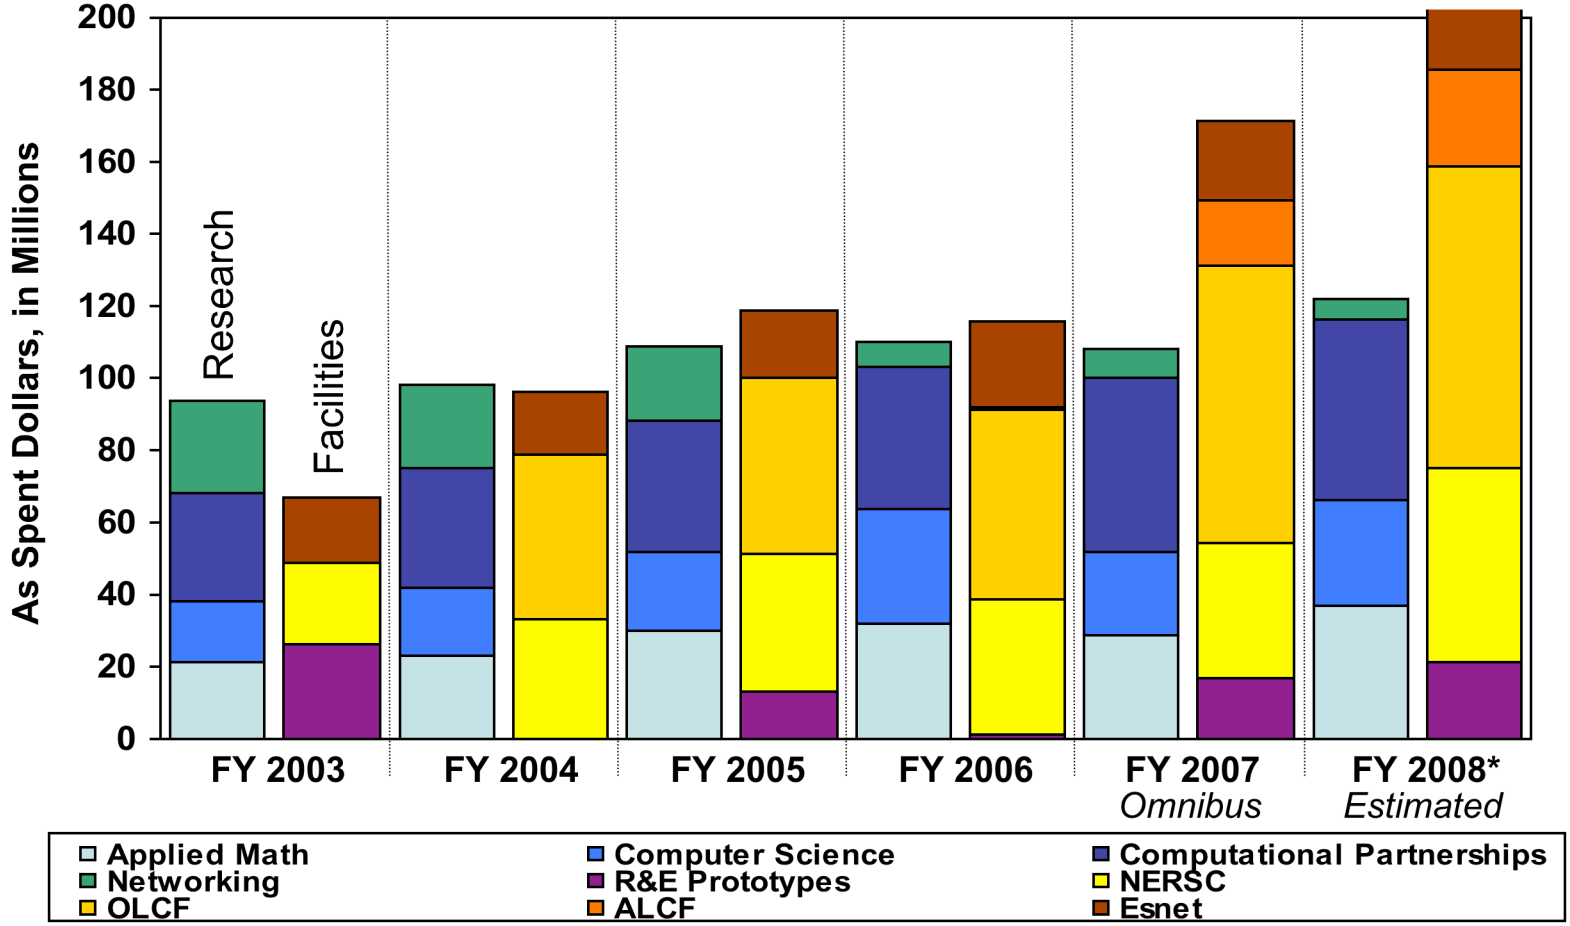

Additional findings and recommendations concerning strategies for a balanced portfolio that

- Supports LCF and capacity computing
- Leverages the SciDAC model
- Supports long-term research with high potential payoff

LCF and Capacity

- It is absolutely critical to US science that ASCR maintain a healthy, effective LCF
- Access to capacity computing is also important
  - LCF architectures change dramatically over short time periods
  - Transition of LCF to capacity as it “ages” maintains software investment and supports broader user community
  - Some important science does not always need LCF
    - E.g., parametric studies involving 1000’s of runs
    - Routine use of verified & validated simulation capability

How to provide capacity and manage transition to Exascale?

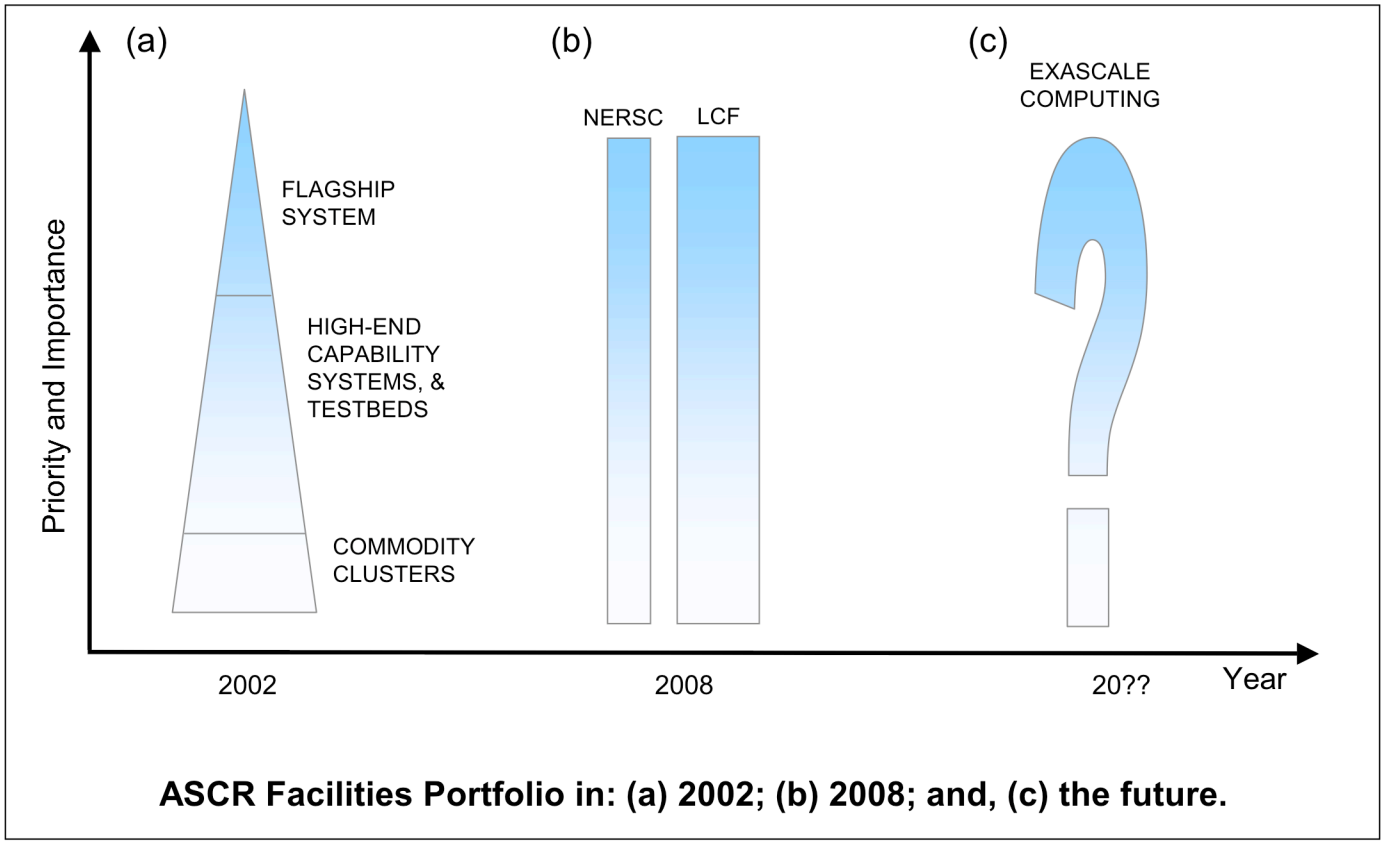

ASCR Facilities Portfolio in: (a) 2002; (b) 2008; and, (c) the future.

Maintaining LCF capability may require using present system as upgrade pathway rather than transitioning to capacity

Recommendation:

- ASCR should plan for re-competing the LCF on a timeframe for exascale computing, in 7-10 years
- Concomitantly, ASCR should develop a plan to
  - Migrate the then-LCF to capacity facilities
  - Support the associated range of different types of core research required

Historically, there have been “knees” in architectural developments, which required revolutionary thinking about the software needed for effective systems.

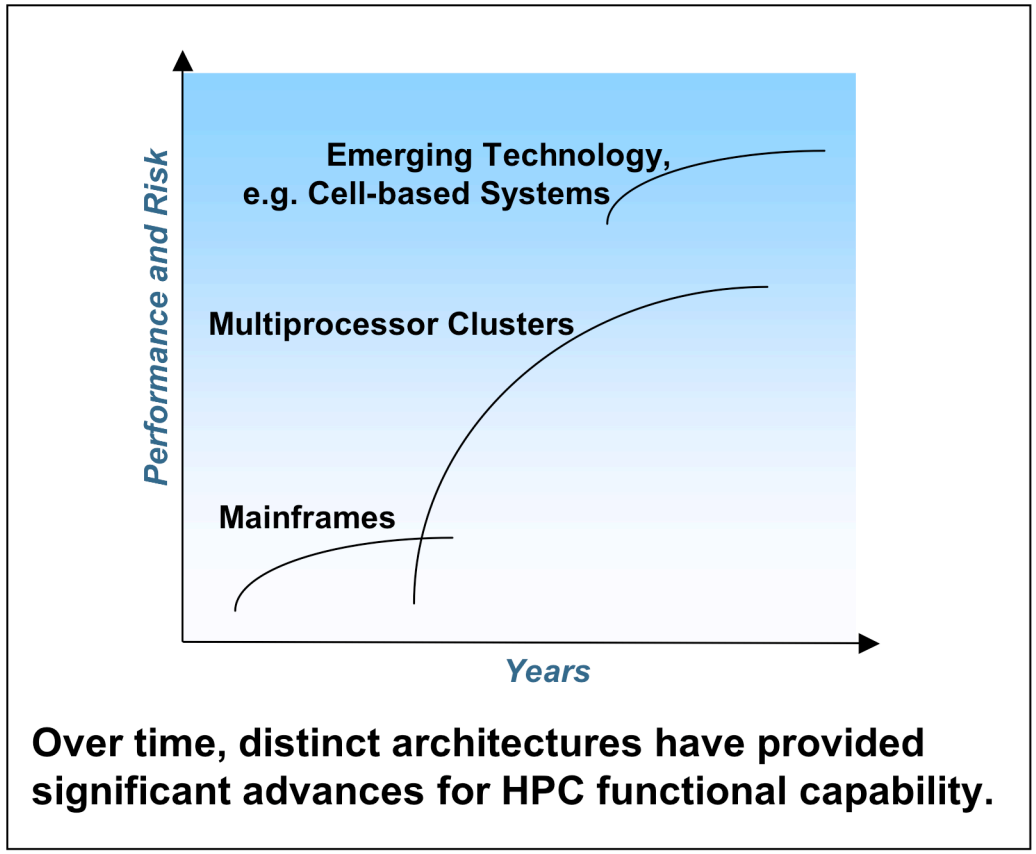

Over time, distinct architectures have provided significant advances for HPC functional capability.

Figure based from communications with Dr. Michel McCoy (LLNL), per: Riding the Waves of Supercomputing Technology [https://www.llnl.gov/str/June03/McCoy.html]

New “Knee”

- Heterogeneous systems with nodes containing hundreds of cores
- Effective utilization requires increased investment in math/cs research

Recommendation:

Support core research in math/cs that can advance scientific discovery

- Leading edge?
- Large scale?
- Agency relevant?

Fraction of research portfolio needs to address new opportunities as they emerge, e.g.,

- Institutionalize multiscale-multiphysics in SC
- Develop programs in data analytics
- Support collaboratories to advance the pace of scientific discovery
- Support research in V&V and Uncertainty Quantification

Recommendation:

Devote a fraction, e.g. 40%, of the research budget in math/cs to support high-risk activities that have the potential to make fundamental changes in how scientific discovery is conducted

SciDAC is highly successful program
SciDAC:Basic Research ~ 40:60

Recommendations:

- Establish 50:50 balance
  - Closer collaboration between math/cs & domain scientists
- Evaluate CET's on basis of productivity of collaboration
- Prune CET's that are not advancing applications

INCITE is important, successful program

- Large amounts of time to small number of users
- Need to demonstrate efficiency of new system may favor applications that can demonstrate high utilization

Recommendation:

Balance in the INCITE program between scientific importance and code efficiency should strongly favor the importance of the science

INCITE selection process favors applications that are "LCF ready"

Assistance from ASCR could broaden base of users and increase scientific discovery and support from other offices

SciDAC awardees have dual review for LCF access

Recommendations:

- Form "SciDAC-like" groups for important applications with specific math/cs expertise to accelerate their readiness for LCF
- Streamline the process for SciDAC access to LCF

The Challenge:
International scientific competitiveness

We must invest in facilities to stay in the game,

but we must invest in research to win.

The ultimate motivation for balance