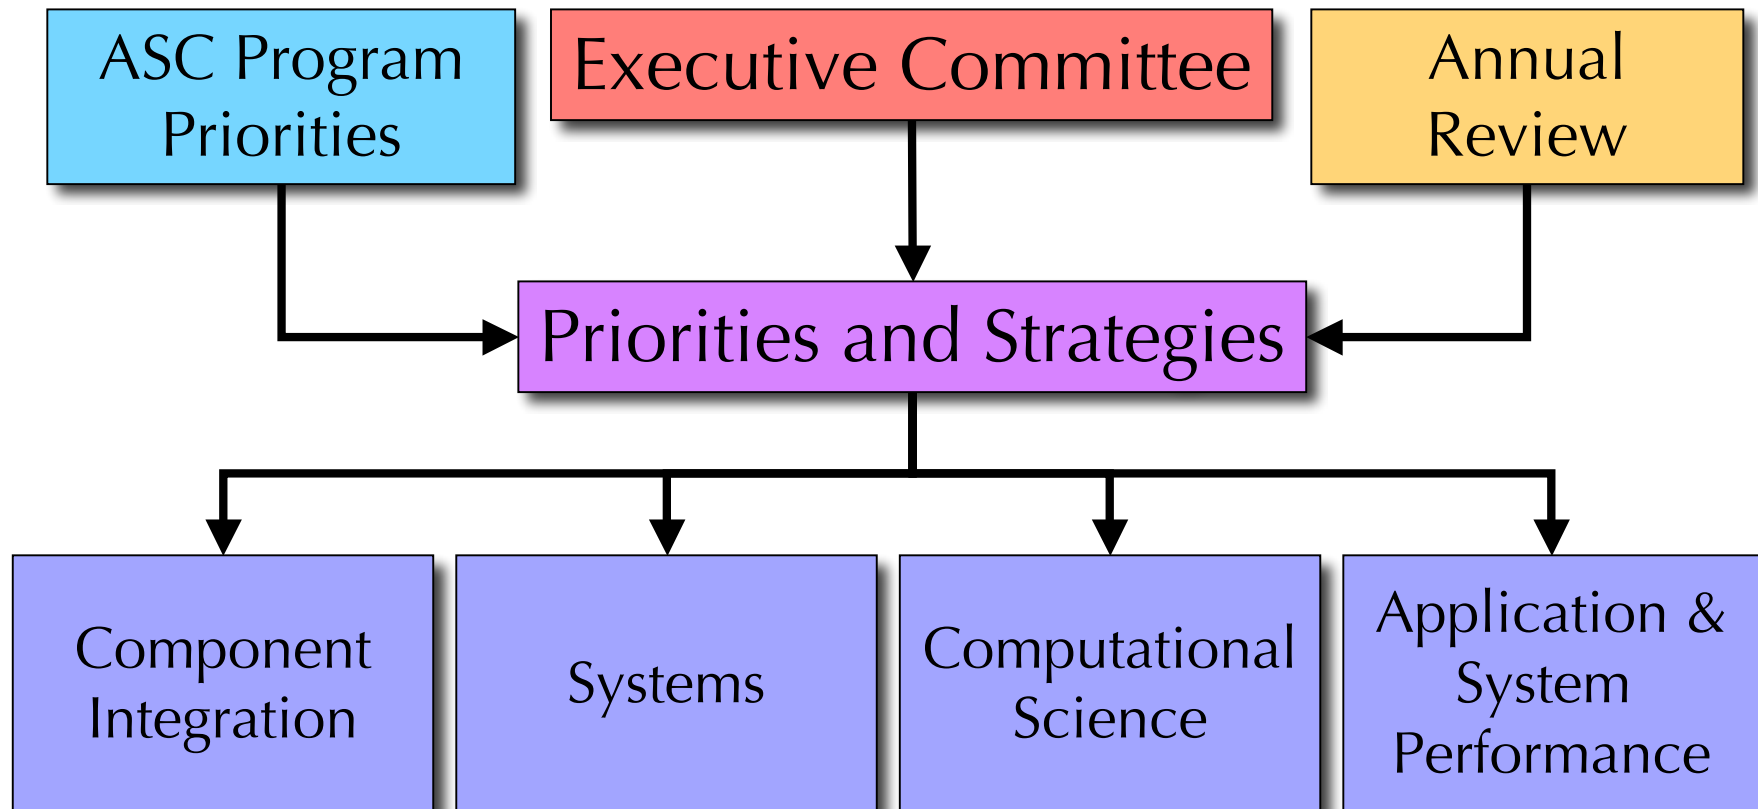

LACSI Management and ASC Impact

Overview

Ken Kennedy
Rice University
LACSI Co-Director

<http://laci.rice.edu/review/2004/slides/management.pdf>

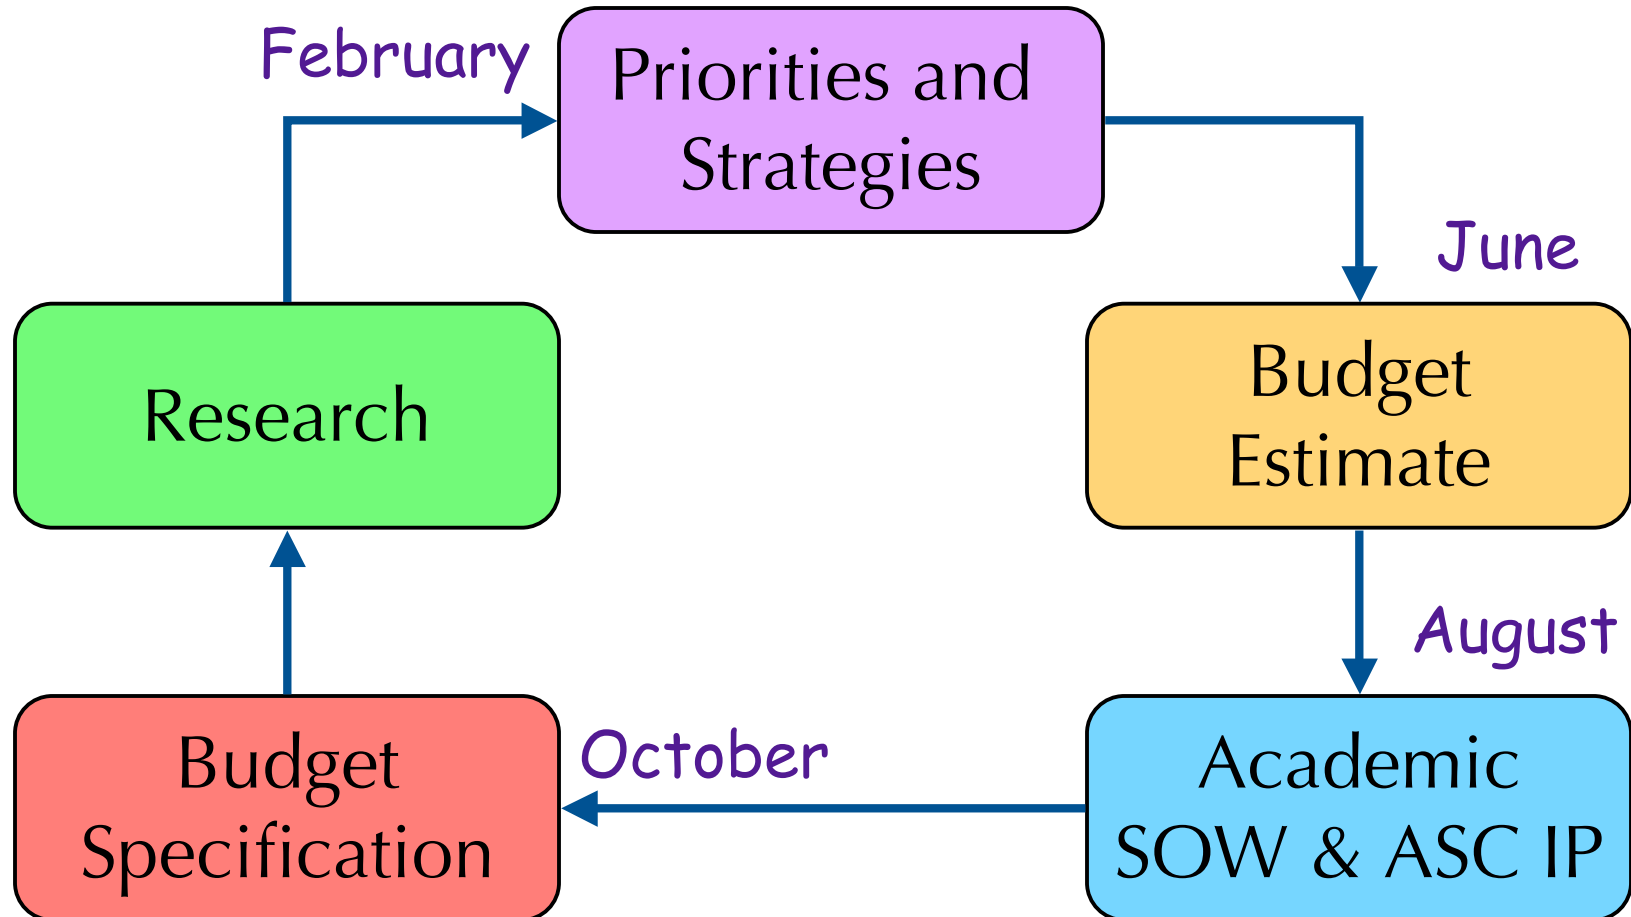
LACSI Structure and Organization



Leadership

- **Co- Directors**
 - Andy White (Los Alamos) & Ken Kennedy (Rice)
- **Executive Committee**
 - **LANL:** Jeff Brown, Bill Feiereisen, Adolfy Hoisie, Doug Kothe, Rod Oldehoeft, John Thorp, Andy White
 - **Rice:** Rob Fowler, Ken Kennedy, John Mellor-Crummey, Linda Torczon
 - **Houston:** Lennart Johnsson, Yuri Kuznetsov
 - **New Mexico:** Deepak Kapur
 - **North Carolina:** Dan Reed
 - **Tennessee:** Jack Dongarra

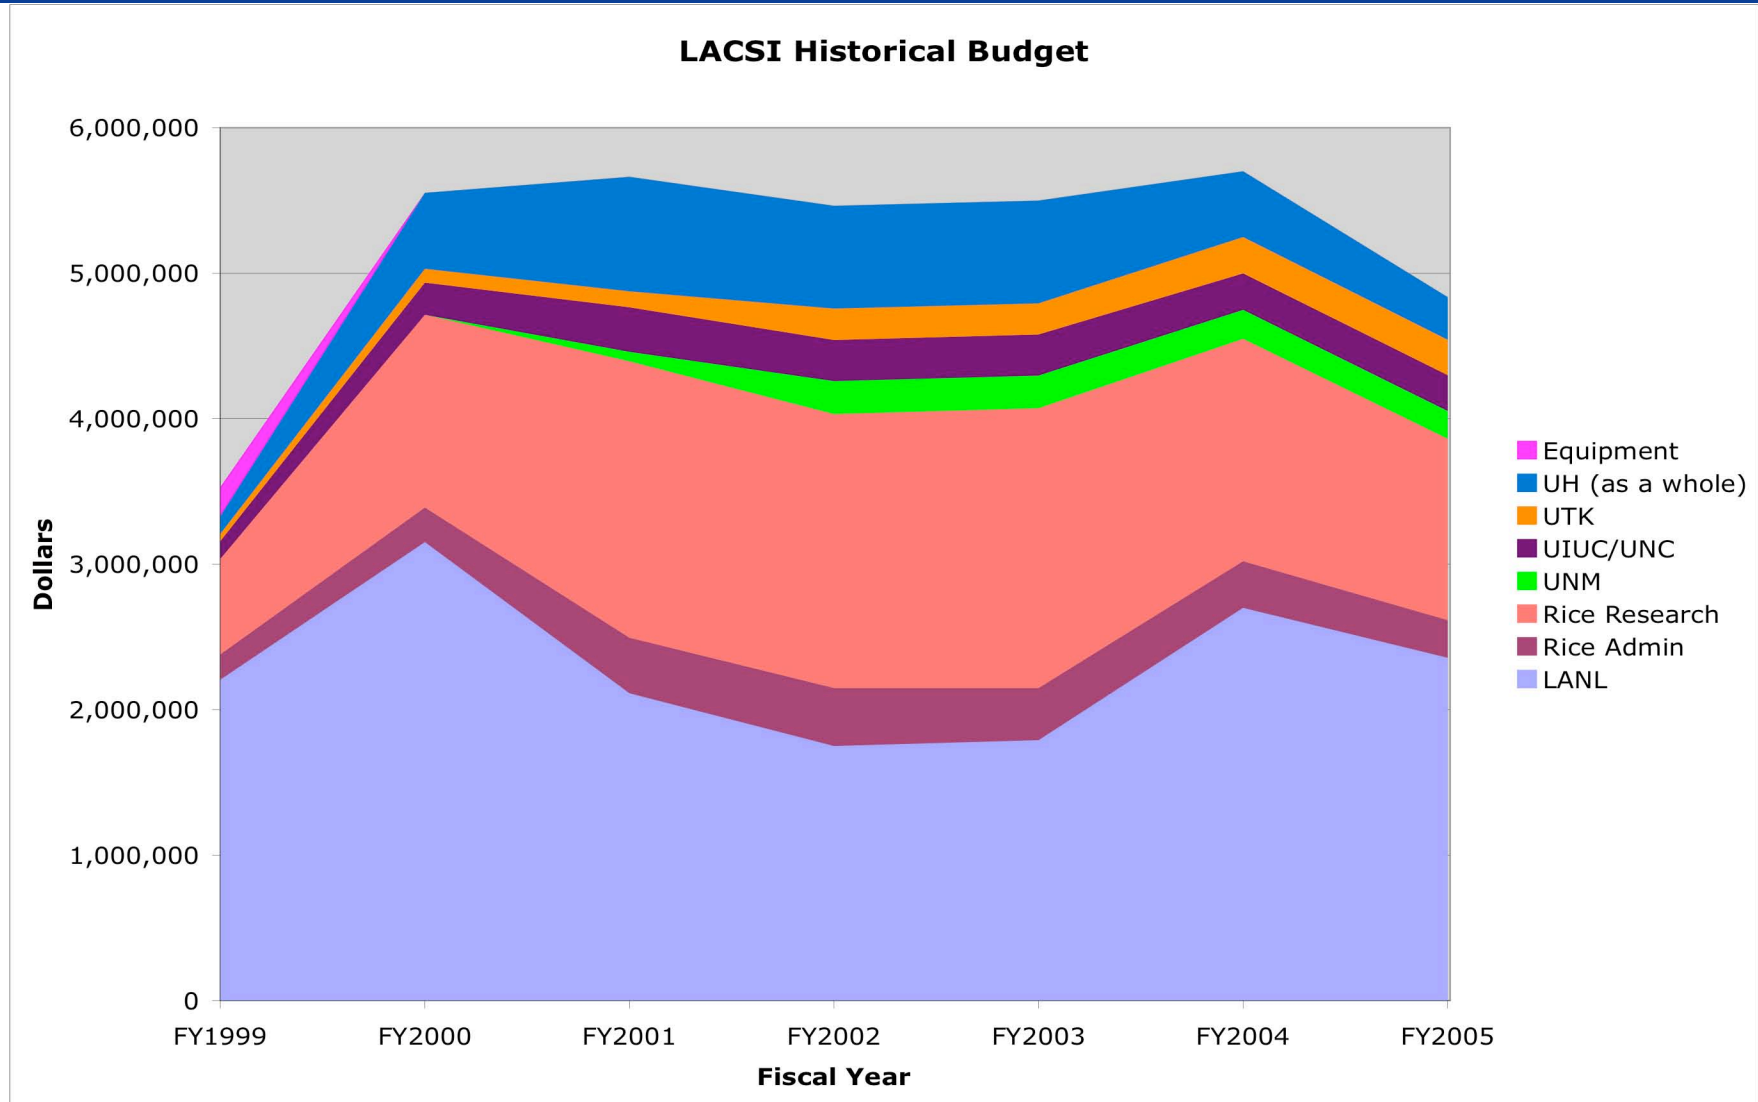
Annual Planning Cycle




Priorities and Strategies Meeting

- **FY05 Attendees:**
 - Bill Archer, Jeff Brown, John Cerutti, Darren Kerbyson, Ken Koch, Stephen Lee, Ron Minnich, Craig Rasmussen, Bill Feiereisen, Adolfy Hoisie, Doug Kothe, Rod Oldehoeft, John Thorp, Andy White, Mike Fagan, Rob Fowler, Ken Kennedy, John Mellor-Crummey, Dan Sorensen, Bill Symes, Linda Torczon, Deepak Kapur, Barney Maccabe, Jack Dongarra, Roland Glowinski, Lennart Johnsson, Yuri Kuznetsov, Dan Reed
- **Typical Agenda**
 - Review of previous year's P&S plan
 - Discussion in plenary session
 - Break into discussion groups for developing plans for the next year
 - **Application and System Performance, Components, Systems, Computational Science**
 - Presentation of revised plans with discussion in plenary session
 - Document developed by email after the meeting


LACSI Budget History



LACSI Budget History by Institution



Budget History by Category



New Review Strategy

- Establish LACSI Review Board (LRB)
 - Include LANL and ASC stakeholders (ASC application developers, computer and computational scientists)
 - External reviewers with multi-year terms
- Combine annual review by LRB with P&S meeting
 - Schedule LRB review on the day before the P&S meeting
 - Outbrief with Executive Committee and formal report
- Use outcome of review as input to P&S planning meeting
 - P&S document becomes proposal for the next year's funding
 - Direct input into academic SOW and LANL ASC IP
- Maintain multi-year stability of projects
 - Phase out unsuccessful projects after 2 or 3 years
 - Intermediate reviews provide constructive criticism

New Annual Planning Cycle



Other Management Challenges

- Enhancing collaboration
 - We understand what works: direct interaction
 - How can we foster more of this?
 - Santa Fe Information Technology lab would be Plan A
 - We are currently engaged in Plan B
 - What incentives are needed on the LANL side?
- Enhancing visibility within LANL
 - Not enough LANL staffers know about LACSI
 - No tangible point of presence
- Applicability: Bringing research (more directly) to bear on ASC problems
 - The problem of classified codes versus sanitized benchmarks
 - More sanitized codes (and data), more cleared academic researchers

One Impact: Tuning of Codes

- **Sage**
 - Elimination of unnecessary copying: 2x on Blue Mountain, 1.5x on Q
 - New sparse matrix representation (not integrated): 2x on Itanium
- **Sweep3D**
 - Memory hierarchy transformations: 1.44x on Alpha, 1.9x on Origin
- **Blanca**
 - Algorithms for reordering vertex lists: $O(n^2)$ -> $O(n \log n)$
 - 26x on the 5-level refinement case
- **CHAD**
 - Reordering of accesses to irregular mesh based on space-filling curves: 2x improvement over random ordering
 - Better scalarization of F90 array accesses: reduced memory traffic by 40%

Some Responses to Questions

LACSI Review Wrap-Up

Andy White

LANL

LACSI Co-Director

Indicators for Success (Last Time)

- **Technology transfer**
 - Joint planning ✓
 - Identify long-term impact of research ✓
 - Communicate with LANL users ✓
- **Planning**
 - Time may be right for reevaluation of foci ✓
 - Each project needs to be aware of its applicability ✓
 - Address relevance regularly ✓
- **Integration**
 - Merge “internal” and “external” LACSI efforts
 - We did increase coordination
 - Review all of LACSI next time ✓

Metrics for Success Next

- All of the above, plus
 - More technologies deployed and used at LANL:
 - E.g., HPCToolkit, distorted polyhedral meshes
- Continued investment in long-term, high-risk research efforts
 - Which, if successful, would have enormous impact on the weapons program

Process for Deployment

- This is difficult, systemic problem that transcends ASC and LANL
 - Not enough resources in research to support deployment
 - Not enough production resources to move all universally accepted research prototypes into production
- There have been successes, as well as frustrations
 - LACSI cannot solve this problem by itself
 - We will continue to try to bring the parties together

Productivity

- Would we object if the report suggested that the unifying theme (end goal) should be human productivity?
 - Short answer: NO
 - Longer answer: There is definitely a way in which everything we are currently doing is having impact on productivity
 - Tools and languages: shorter development time, shorter time to deployment, maintenance costs
 - Improved performance of machines: applications produce answers faster
 - Better (faster, more robust) algorithms: same as above
 - Improved reliability: less time spent in recovering from failures

Other Issues?

- **Work with ASCI ASAP Centers?**
 - All our energy is appropriately focused on the laboratory
 - Plan for SFITL was to serve as focal point for collaborations with ASAPs as well as the LACSI partners
- **Visualization?**
 - LACSI's mission is to look at all problems: cannot solve all of them
 - Need to maintain critical mass in the areas of strength
- **New Efforts? New Partners?**
 - We can fund new research starts, but must cut elsewhere to do so
 - We can add new partners, but our model is to do so only when a long-term commitment makes sense