Long-Term Information Technology Research

Meeting the PITAC Challenge

Ken Kennedy
Center for High Performance Software
Rice University

http://www.cs.rice.edu/~ken/Presentations/SIAMPITAC.pdf

PITAC Charter

- The Committee shall provide an independent assessment of:
  - Progress made in implementing the High-Performance Computing and Communications (HPCC) Program;
  - Progress in designing and implementing the Next Generation Internet initiative;
  - The need to revise the HPCC Program;
  - Balance among components of the HPCC Program;
  - Whether the research and development undertaken pursuant to the HPCC Program is helping to maintain United States leadership in advanced computing and communications technologies and their applications;
  - Other issues as specified by the Director of the Office of Science and Technology.
    - Review of the entire IT investment strategy — is it meeting the nation's needs

PITAC Membership 97-99

- **Co-Chairs:**
  - Bill Joy, Sun Microsystems
  - Ken Kennedy, Rice
- **Members:**
  - Eric Benhamou, 3Com
  - Ching-chih Chen, Simmons
  - Steve Dorfman, Hughes
  - Bob Ewald, SGI
  - Sherri Fuller, U of Washington
  - Susan Graham, UC Berkeley
  - Danny Hillis, Disney, Inc
  - John Miller, Montana State
  - Raj Reddy*, Carnegie Mellon
  - Larry Smarr, UIUC
  - Les Vadasz, Intel
  - Steve Wallach, Centerpoint
  - Vinton Cerf, MCI
  - David Cooper, LLNL
  - David Dorman, PointCast
  - David Farber, Penn
  - Hector Garcia-Molina, Stanford
  - Jim Gray, Microsoft
  - Robert Kahn, CNRI
  - David Nagel, AT&T
  - Ted Shortliffe, Stanford
  - Joe Thompson, Miss. State
  - Andy Viterbi, Qualcomm
  - Irving Wladawsky-Berger*, IBM

* = current co-chair

Methodology

- Evaluation of Federal Research Investment Portfolio
  - —Plans reviewed for each of the major areas:
    - High End Computing and Computation
    - Large Scale Networking
    - Human Centered Computer Systems
    - High Confidence Systems
    - Education, Training, and Human Resources
- Review of Balance in Federal Research Portfolio
  - —Fundamental versus Applied
    - Based on our own definition of these terms
  - —High-Risk versus Low-Risk
  - —Long-Term versus Short-Term

Principal Finding

- Drift Away from Long-Term Fundamental Research

Principal Finding

- Drift Away from Long-Term Fundamental Research
  - Agencies pressed by the growth of IT needs
    - IT R&D budgets have grown steadily but not dramatically
    - IT industry has accounted for over 30 percent of the real GDP growth over the past five years, but gets only 1 out of 75 Federal R&D dollars
    - Problems solved by IT are critical to the nation—engineering design, health and medicine, defense

Principal Finding

- Drift Away from Long-Term Fundamental Research
  - —Agencies pressed by the growth of IT needs
    - IT R&D budgets have grown steadily but not dramatically
    - IT industry has accounted for over 30 percent of the real GDP growth over the past five years, but gets only 1 out of 75 Federal R&D dollars
    - Problems solved by IT are critical to the nation—engineering design, health and medicine, defense
  - —Most IT R&D agencies are mission-oriented
    - Natural and correct to favor the short-term needs of the mission

Principal Finding

- Drift Away from Long-Term Fundamental Research
  - Agencies pressed by the growth of IT needs
    - IT R&D budgets have grown steadily but not dramatically
    - IT industry has accounted for over 30 percent of the real GDP growth over the past five years, but gets only 1 out of 75 Federal R&D dollars
    - Problems solved by IT are critical to the nation—engineering design, health and medicine, defense
  - Most IT R&D agencies are mission-oriented
    - Natural and correct to favor the short-term needs of the mission
- This Trend Must Be Reversed
  - Continue the flow of ideas to fuel the information economy and society

Remedy

Remedy

- Increase the Federal IT R&D Investment by 1.4 billion dollars per year
  - Ramp up over five years
  - Focus on increasing fundamental research

Remedy

- Increase the Federal IT R&D Investment by 1.4 billion dollars per year
  - Ramp up over five years
  - Focus on increasing fundamental research
- Invest in Key Areas Needing Attention
  - Software
  - Scalable Information Infrastructure
  - High-End Computing
  - Social, Economic, and Workforce Issues

Remedy

- Increase the Federal IT R&D Investment by 1.4 billion dollars per year
  - Ramp up over five years
  - Focus on increasing fundamental research
- Invest in Key Areas Needing Attention
  - Software
  - Scalable Information Infrastructure
  - High-End Computing
  - Social, Economic, and Workforce Issues
- Develop a Coherent Management Strategy
  - Establish clear organizational responsibilities
  - Diversify modes of support

Software

- Recommendations
  - Make fundamental software research an absolute priority
  - Invest in key area needing attention
    - Improving programmer productivity
      - Ameliorate the shortage of IT professionals
    - Improving reliability and robustness of software
    - Improving usability through human interface innovations
    - Improving capabilities for information management
  - Make software research a substantive component of every major information technology research initiative.

Scalable Information Infrastructure

- Research Needed:
  - —Understanding the behavior of the global-scale network.
  - —Physics of the network, including optical and wireless technologies such as satellites, and bandwidth issues.
  - —Scalability of the Internet.
  - —Information management, Information and services survivability
  - —Large-scale applications and the scalable services they require.
    - National digital library, Next-generation world-wide web
  - —Fund a balanced set of testbeds that serve the needs of networking research, research in enabling information technologies and advanced applications, and Internet research.

High-End Computing

- Findings:
  - High-end computing is essential for science and engineering research
  - High-end computing is an enabling element of the United States national security program
  - New applications of high-end computing are ripe for exploration
  - Suppliers of high-end systems suffer from difficult market pressures
    - High-end market not large
  - Advances in high-end computing eventually find their way to desktop
- Recommendation
  - Fund high-end computing research (architecture, software, and applications, and testbeds) because it is important to the government and the health, welfare, and security of the population

Social, Economic, Workforce Issues

- Invest in Four Areas:
  - IT-literate population
  - IT workforce
    - More workers, more underrepresented groups
  - Use of IT in education
  - Understanding economic and policy implications of technology

Social, Economic, Workforce Issues

- Invest in Four Areas:
  - —IT-literate population
  - —IT workforce
    - More workers, more underrepresented groups
  - —Use of IT in education
  - —Understanding economic and policy implications of technology
- An Observation on IT Workforce
  - —Research investment in universities is critical
    - Without it, faculty leave
    - Without it, grad students do not go —> no new faculty
    - Without faculty, we cannot produce more BS graduates

Questions

- Can we increase long-term research by rebudgeting?

Questions

- Can we increase long-term research by rebudgeting?
  - No, because the short-term work addresses essential problems

Questions

- Can we increase long-term research by rebudgeting?
  - No, because the short-term work addresses essential problems
- Why doesn't industry fund this?

Questions

- Can we increase long-term research by rebudgeting?
  - No, because the short-term work addresses essential problems
- Why doesn't industry fund this?
  - Industry research focused on product development
    - Enormously expensive

Questions

- Can we increase long-term research by rebudgeting?
  - No, because the short-term work addresses essential problems
- Why doesn't industry fund this?
  - Industry research focused on product development
    - Enormously expensive
  - Thurow:
    - Private rate of return on research — 24%
    - Societal rate of return on research — 66%

Questions

- Can we increase long-term research by rebudgeting?
  - No, because the short-term work addresses essential problems
- Why doesn't industry fund this?
  - Industry research focused on product development
    - Enormously expensive
  - Thurow:
    - Private rate of return on research — 24%
    - Societal rate of return on research — 66%
  - Industry is not good at funding and developing disruptive technologies
  - Federal Government funding creates fuel for the venture capital system

Good News

- Administration Budget
  - Proposed additional $366 million in FY 2000
    - Appropriated: $226 million
  - Proposed $605 million increase for FY 2001
  - Successive years unclear
- Congress
  - Sensenbrenner NITR&D Act from House Science Committee
    - 5 years of funding at PITAC-recommended levels
    - Permanent R&D investment tax credit
    - Passed with near-unanimous support
    - Only partially reflected in the Senate authorization bills
  - Appropriations are year-to-year

IT Grand Challenges

IT Grand Challenges

- Software Reliability
  - Who will pay for bug free, feature-poor software?

IT Grand Challenges

- Software Reliability
  - Who will pay for bug free, feature-poor software?
- Internet Scalability and Security
  - How do we handle 2 billion internet connections at DSL speed?
  - What if those connections are mobile and wireless?

IT Grand Challenges

- Software Reliability
  - Who will pay for bug free, feature-poor software?
- Internet Scalability and Security
  - How do we handle 2 billion internet connections at DSL speed?
  - What if those connections are mobile and wireless?
- Realistic Telepresence
  - Can we put airlines out of business?

IT Grand Challenges

- Software Reliability
  - Who will pay for bug free, feature-poor software?
- Internet Scalability and Security
  - How do we handle 2 billion internet connections at DSL speed?
  - What if those connections are mobile and wireless?
- Realistic Telepresence
  - Can we put airlines out of business?
- The Internet as Problem-Solving Engine*
  - GrADS Project

IT Grand Challenges

- Software Reliability
  - Who will pay for bug free, feature-poor software?
- Internet Scalability and Security
  - How do we handle 2 billion internet connections at DSL speed?
  - What if those connections are mobile and wireless?
- Realistic Telepresence
  - Can we put airlines out of business?
- The Internet as Problem-Solving Engine*
  - GrADS Project
- Software Productivity*
  - Workforce shortage
  - Idea: make it possible for end users to be application developers

National Distributed Computing

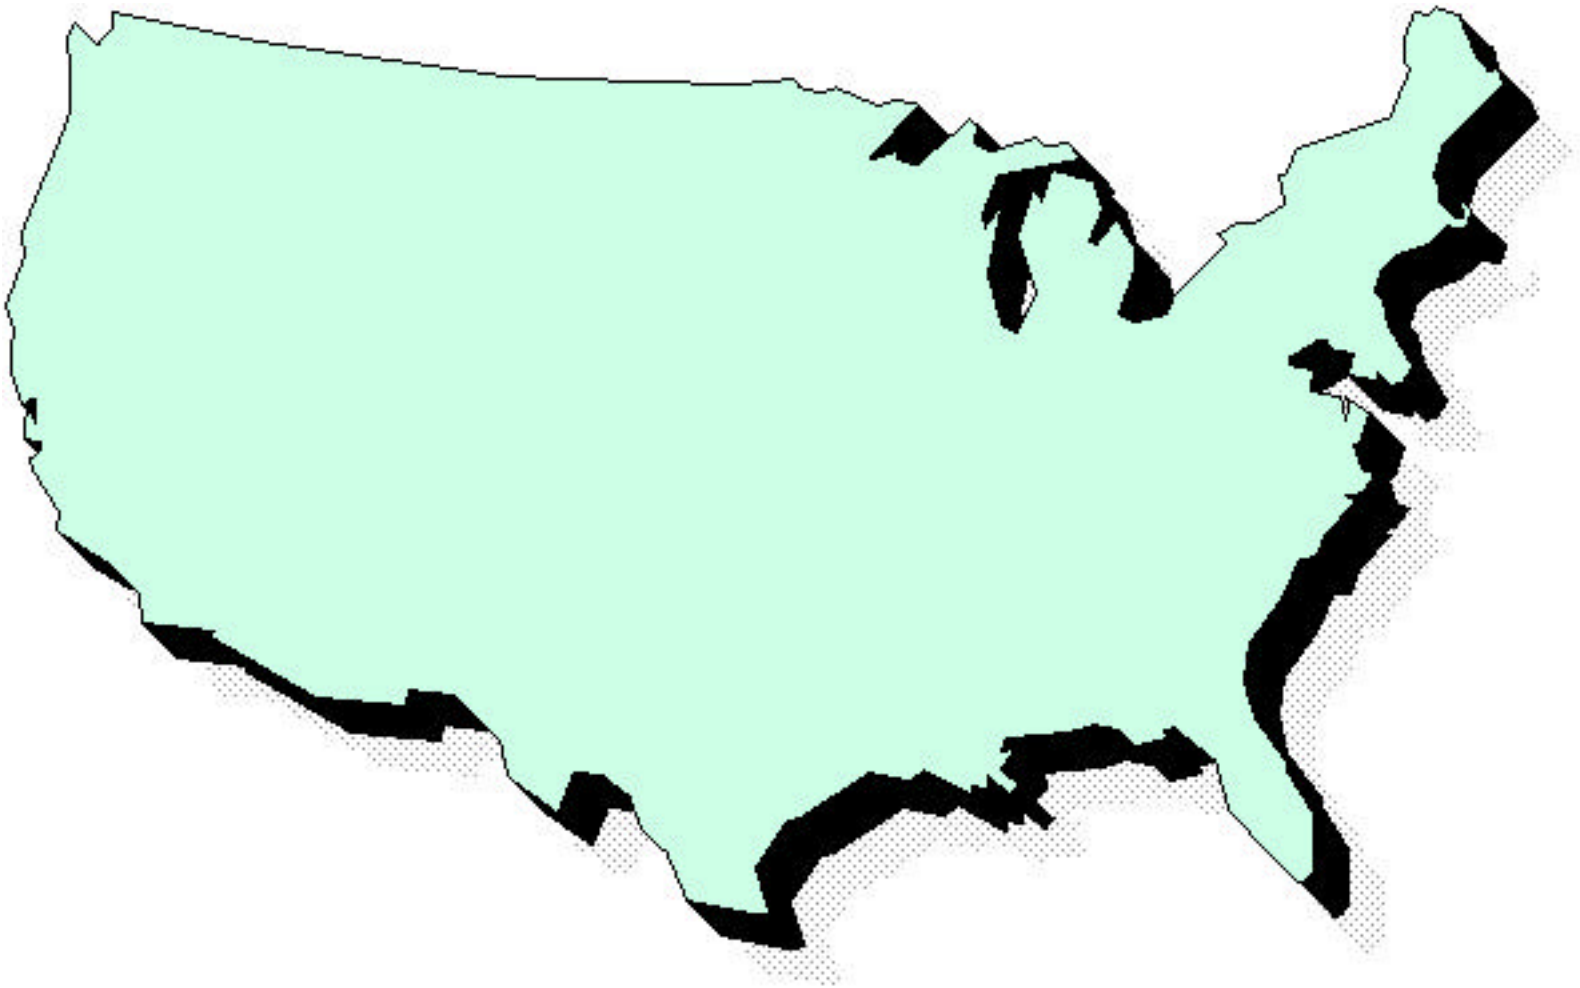

National Distributed Computing

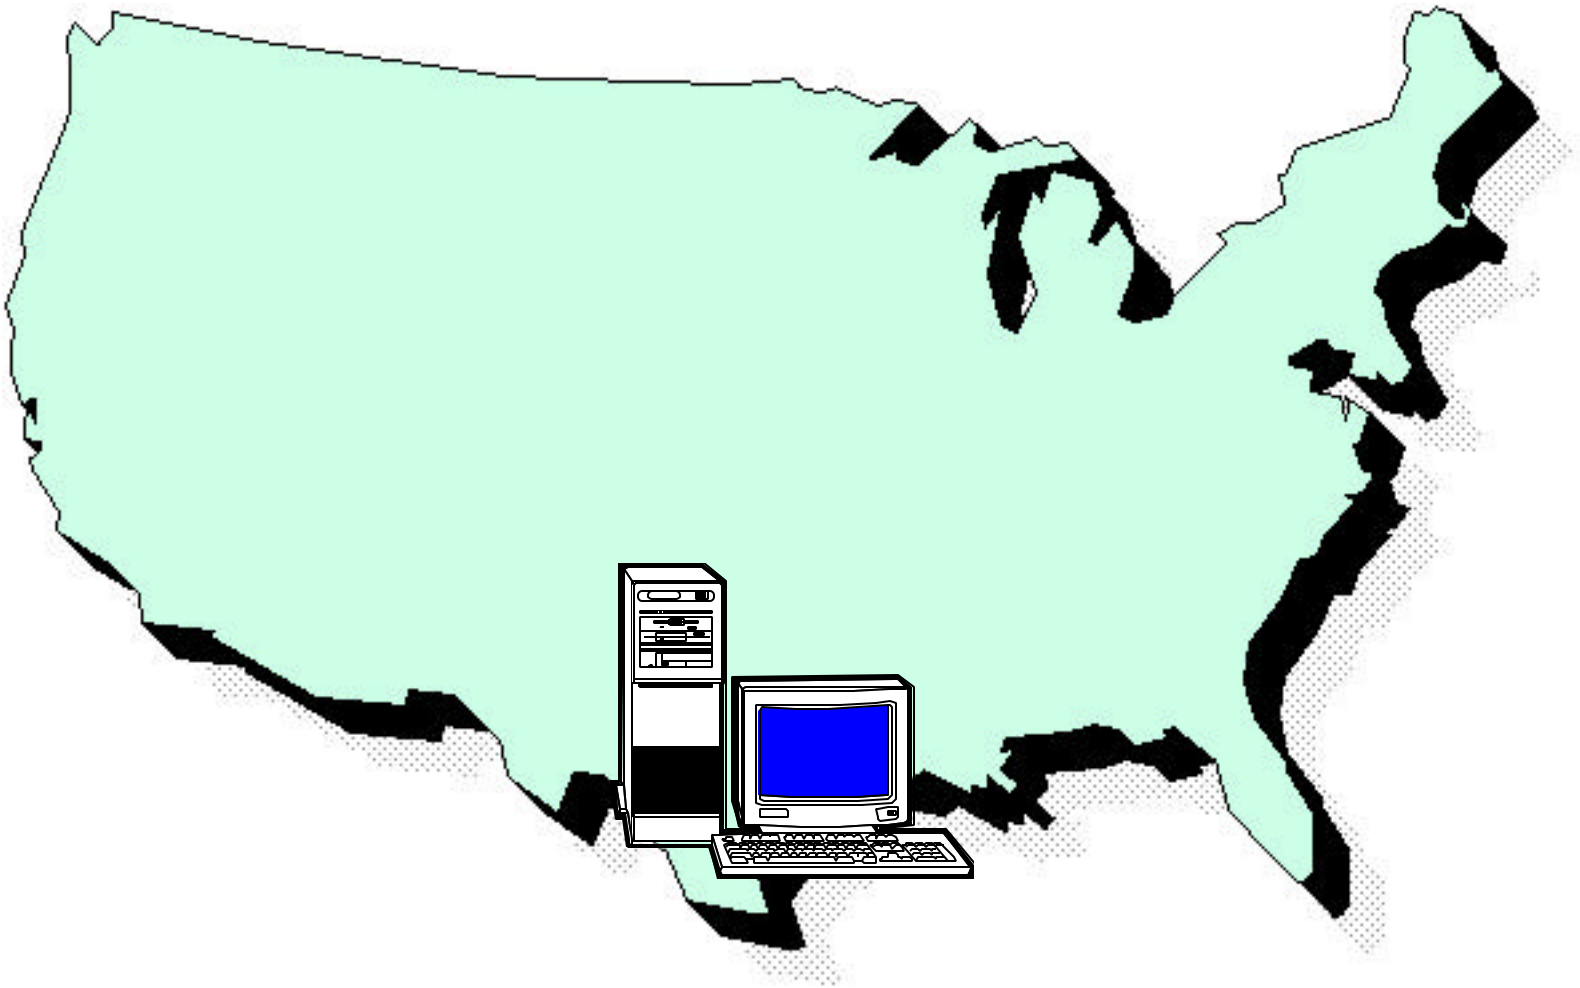

National Distributed Computing

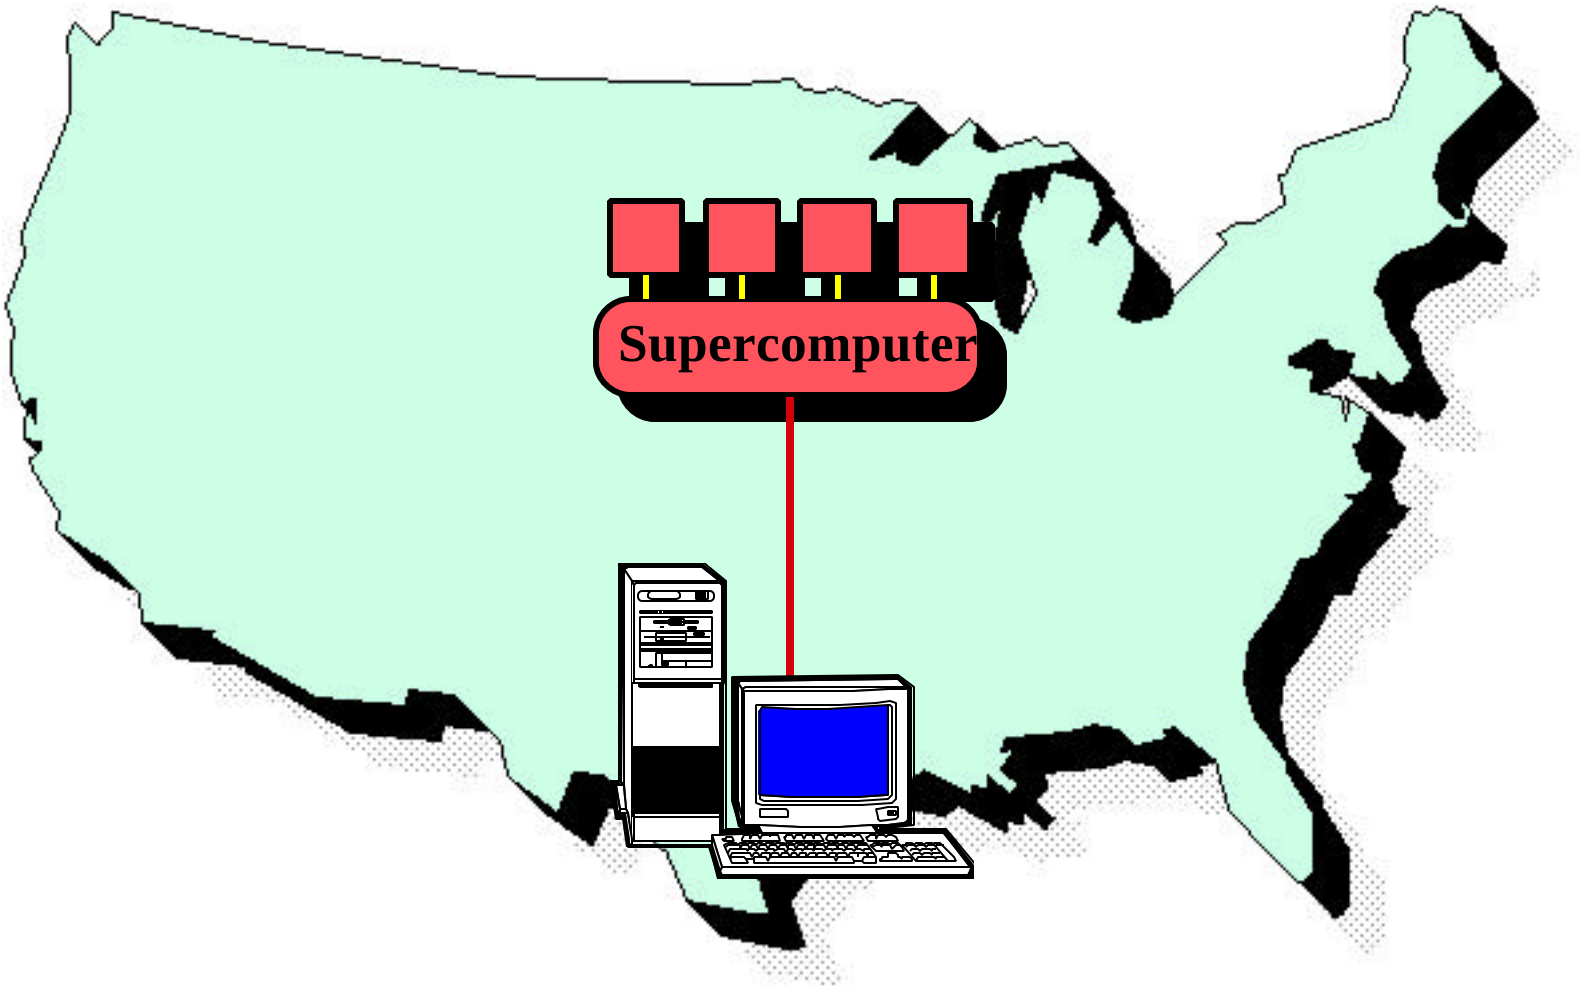

National Distributed Computing

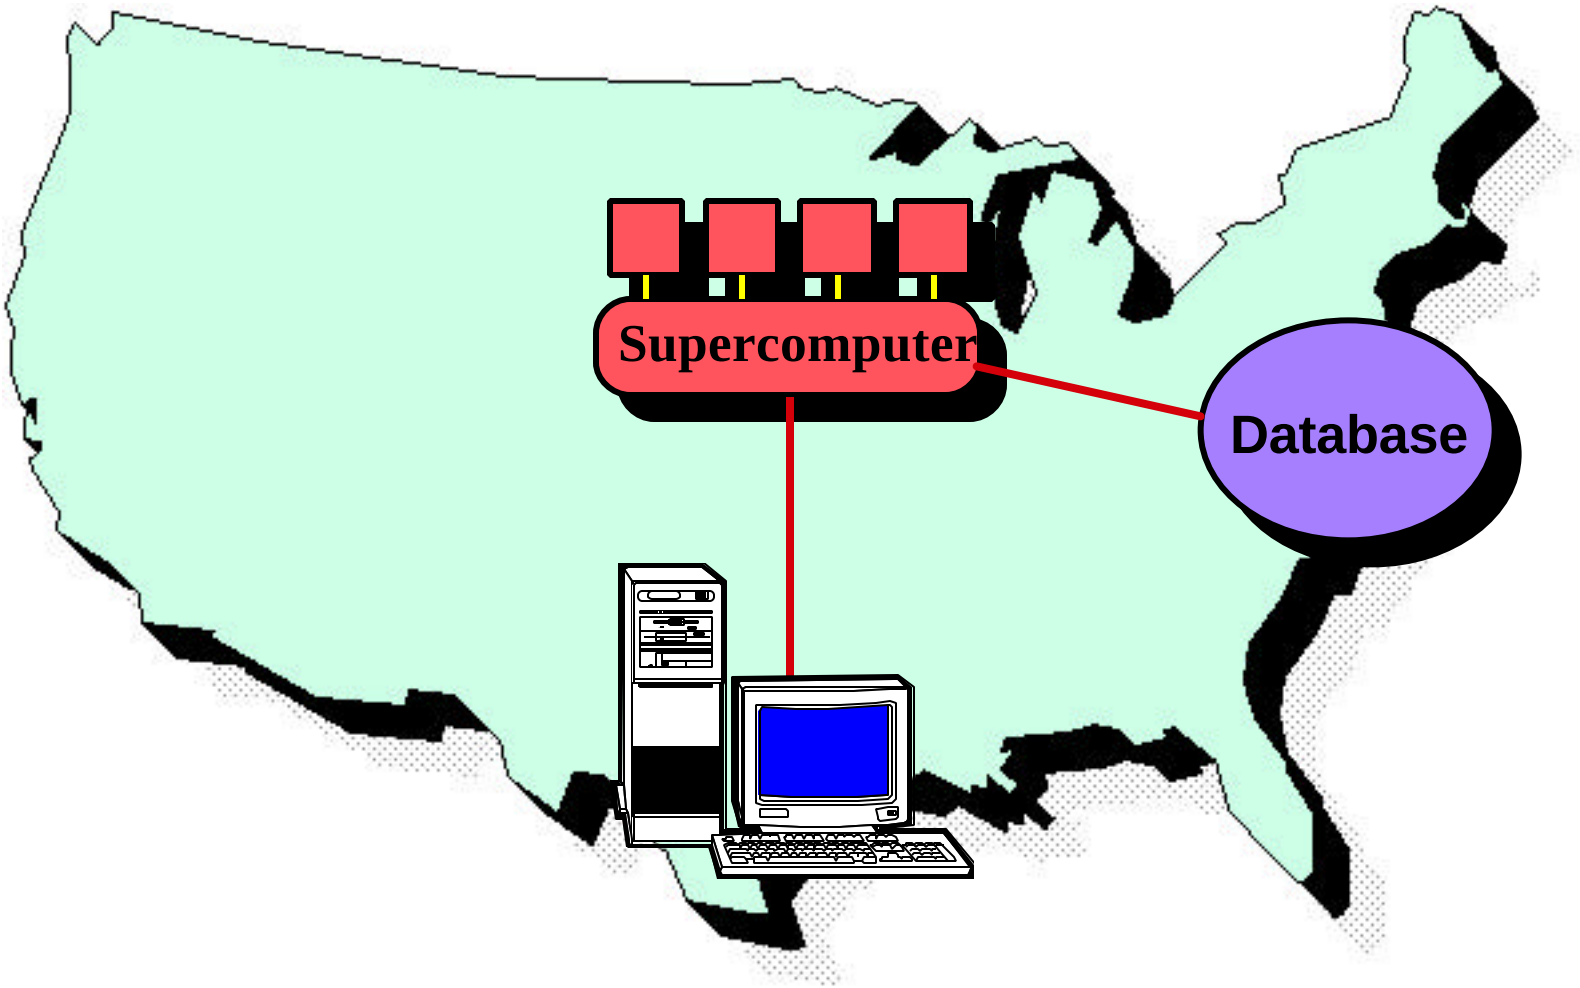

National Distributed Computing

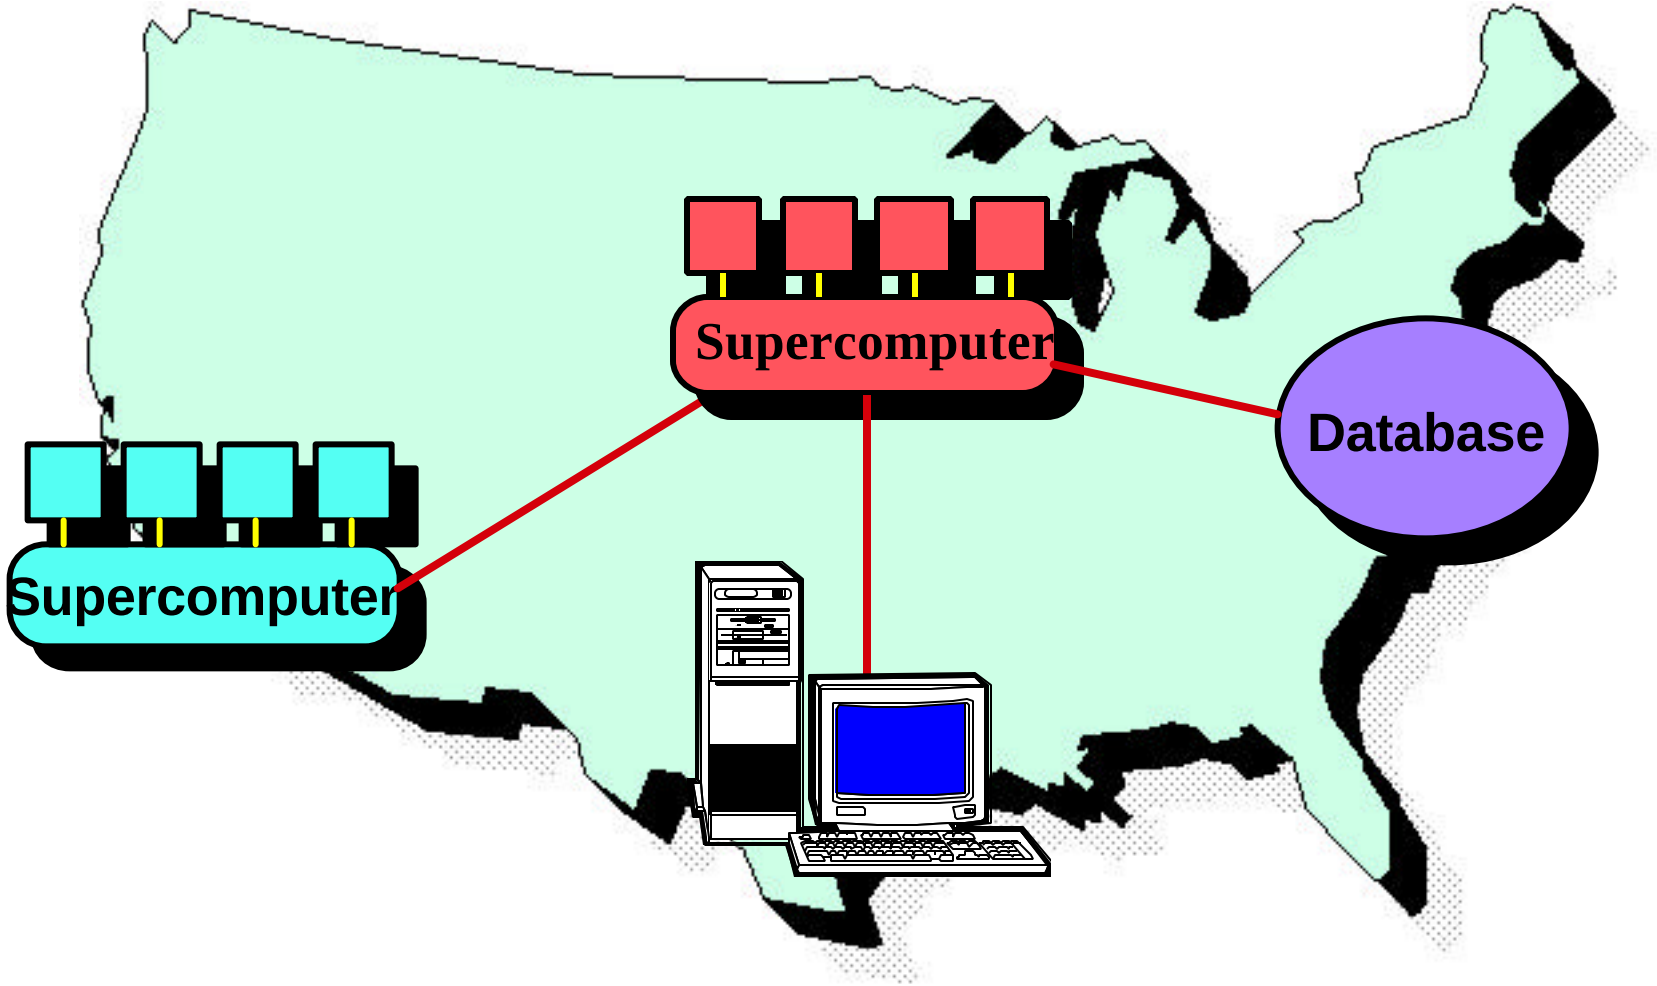

National Distributed Computing

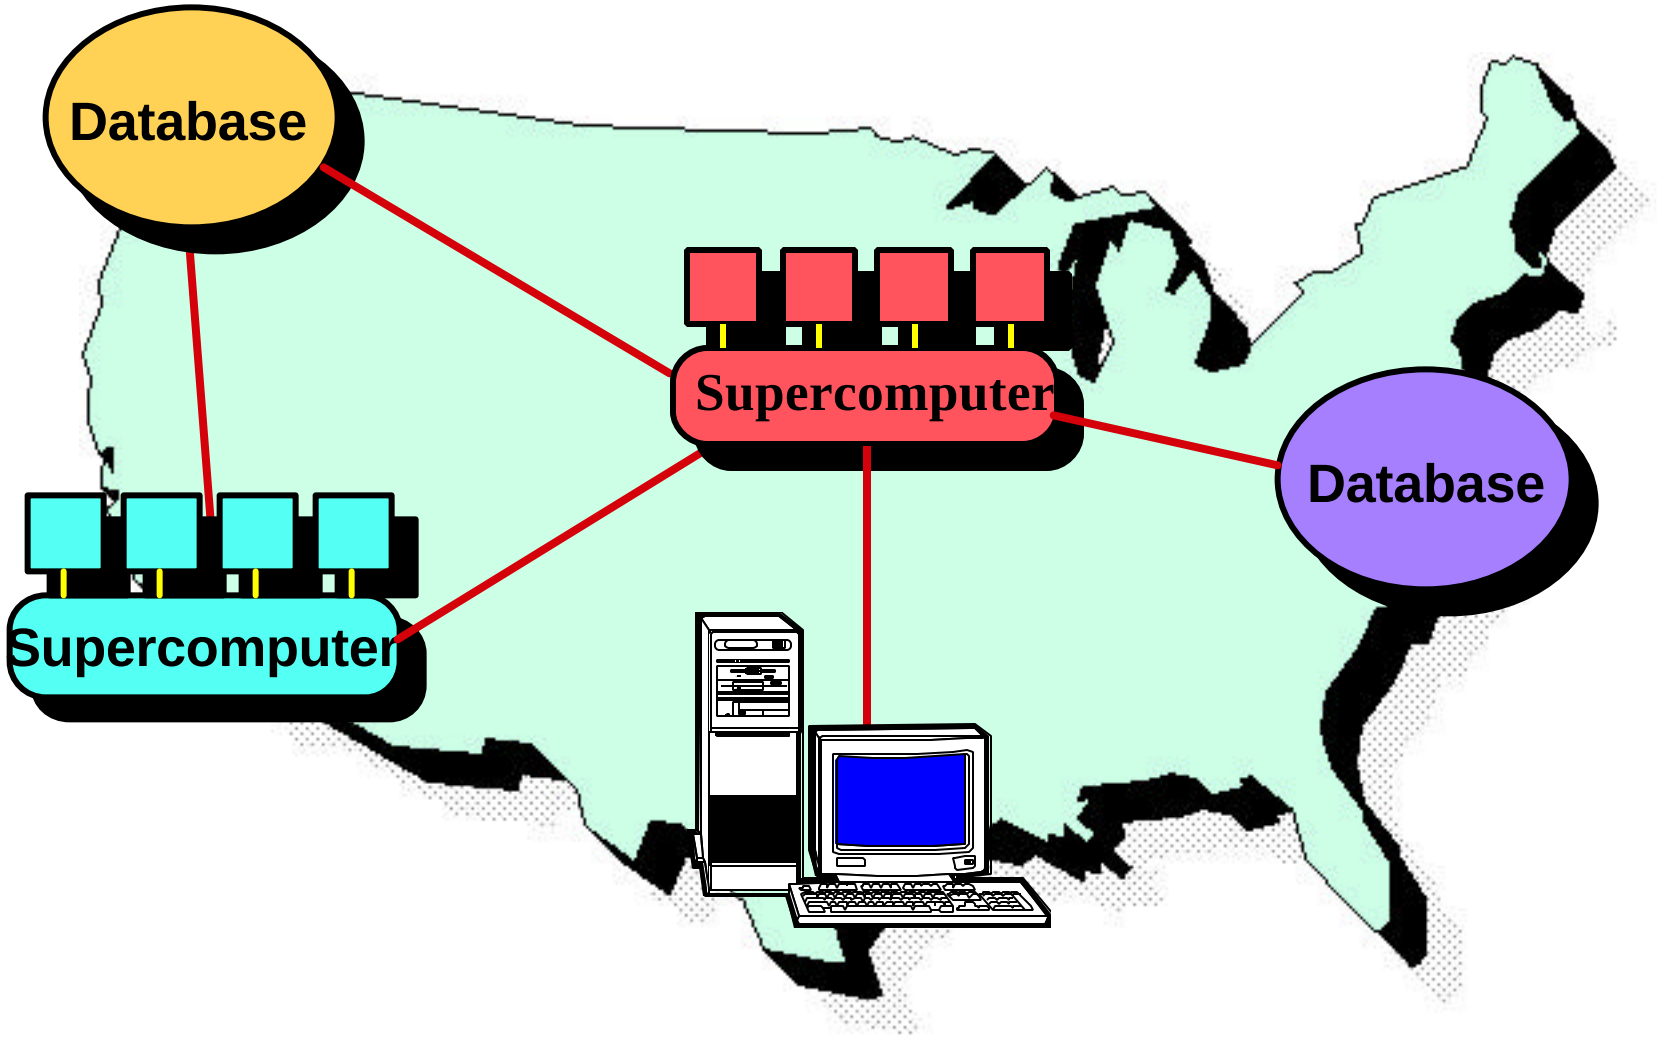

Grid Compilation Architecture

- Goal: reliable performance under varying load

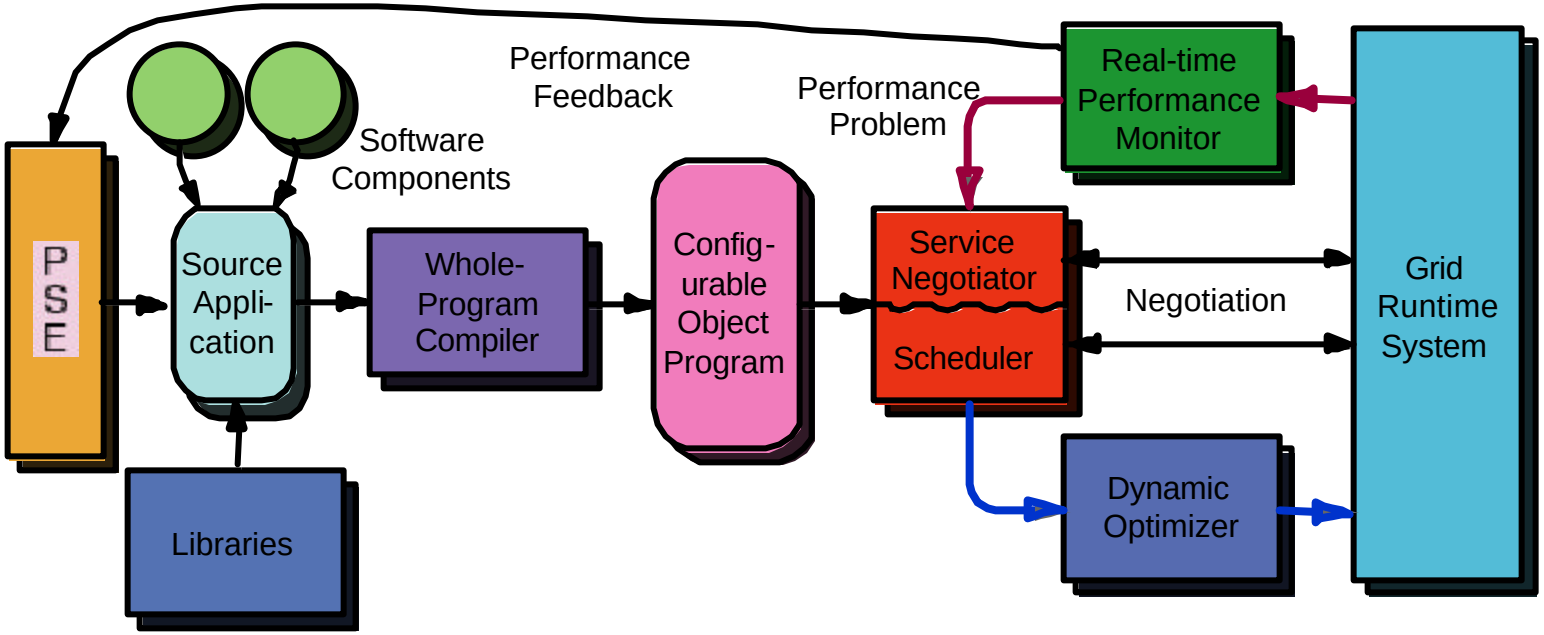

GrADS Project (NSF NGS): Berman, Chien, Cooper, Dongarra, Foster, Gannon, Johnsson, Kennedy, Kesselman, Mellor-Crummey, Reed, Torczon, Wolski

Grid Compilation Architecture

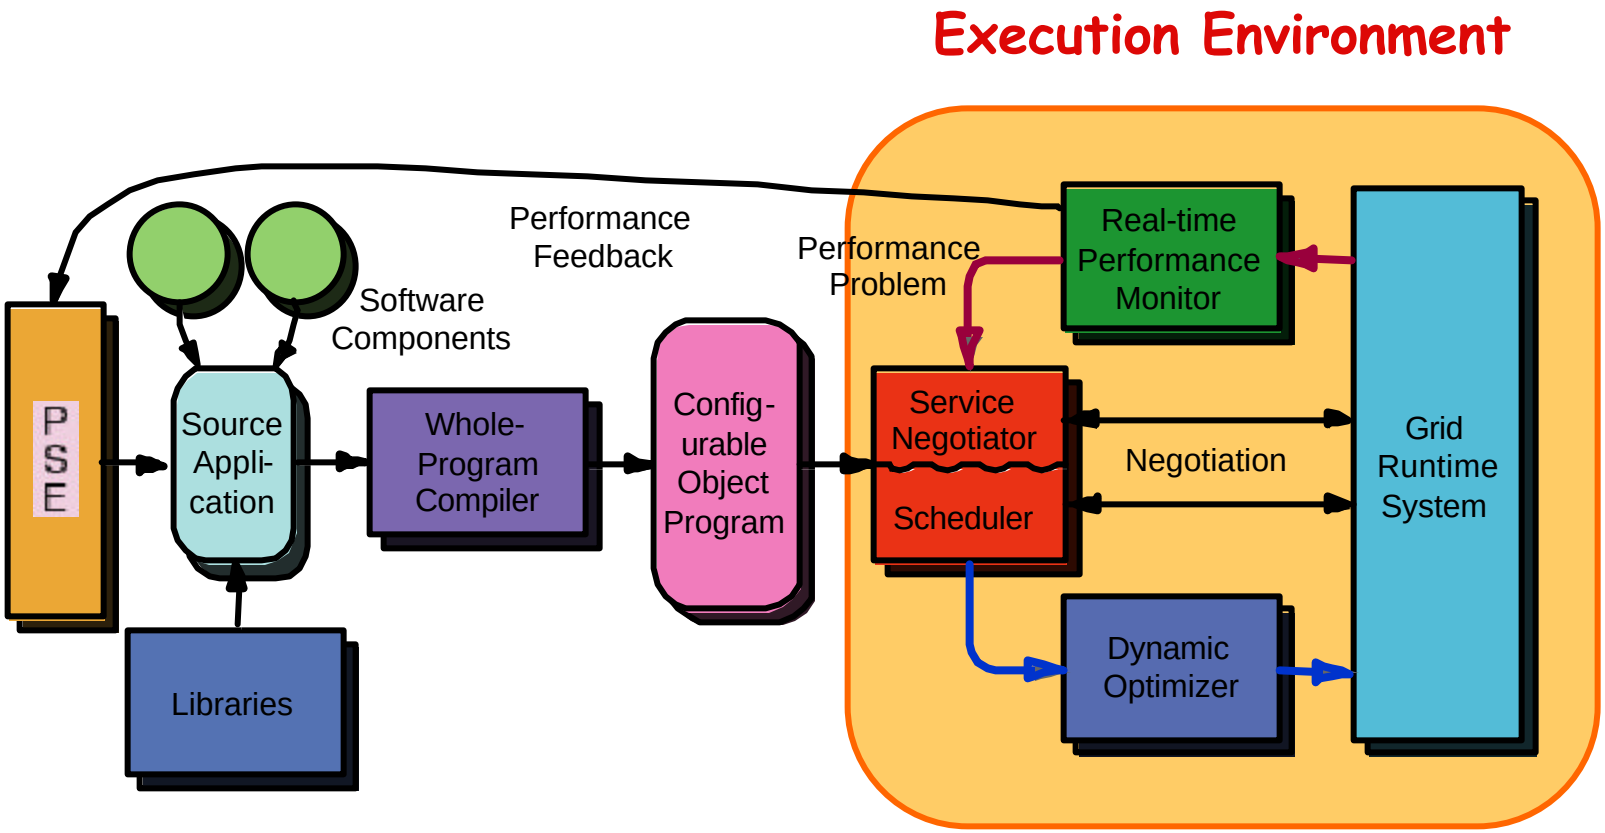

Grid Compilation Architecture

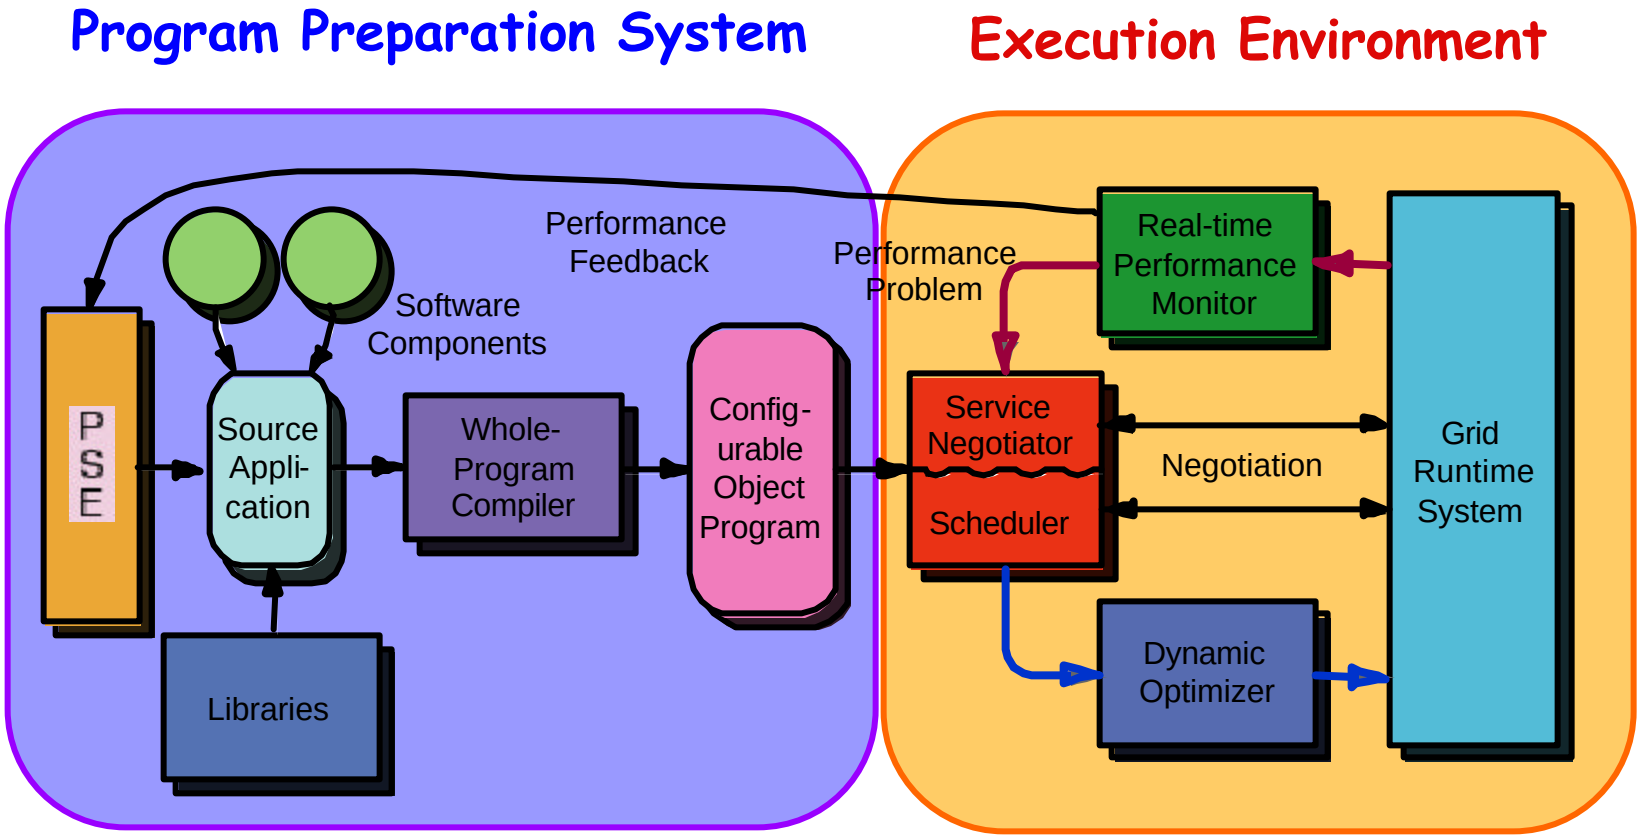

Programming Productivity

- Challenges
  - programming is hard
  - professional programmers are in short supply
  - high performance will continue to be important

Programming Productivity

- Challenges
  - programming is hard
  - professional programmers are in short supply
  - high performance will continue to be important
- One Strategy: Make the End User a Programmer
  - professional programmers develop components
  - users integrate components using:
    - problem-solving environments (PSEs)
    - scripting languages (possibly graphical)

      examples: Visual Basic, Tcl/Tk, AVS, Khoros

Programming Productivity

- Challenges
  - programming is hard
  - professional programmers are in short supply
  - high performance will continue to be important
- One Strategy: Make the End User a Programmer
  - professional programmers develop components
  - users integrate components using:
    - problem-solving environments (PSEs)
    - scripting languages (possibly graphical)

      examples: Visual Basic, Tcl/Tk, AVS, Khoros
- Compilation for High Performance
  - translate scripts and components to common intermediate language
  - optimize the resulting program using interprocedural methods

Telescoping Languages

Telescoping Languages

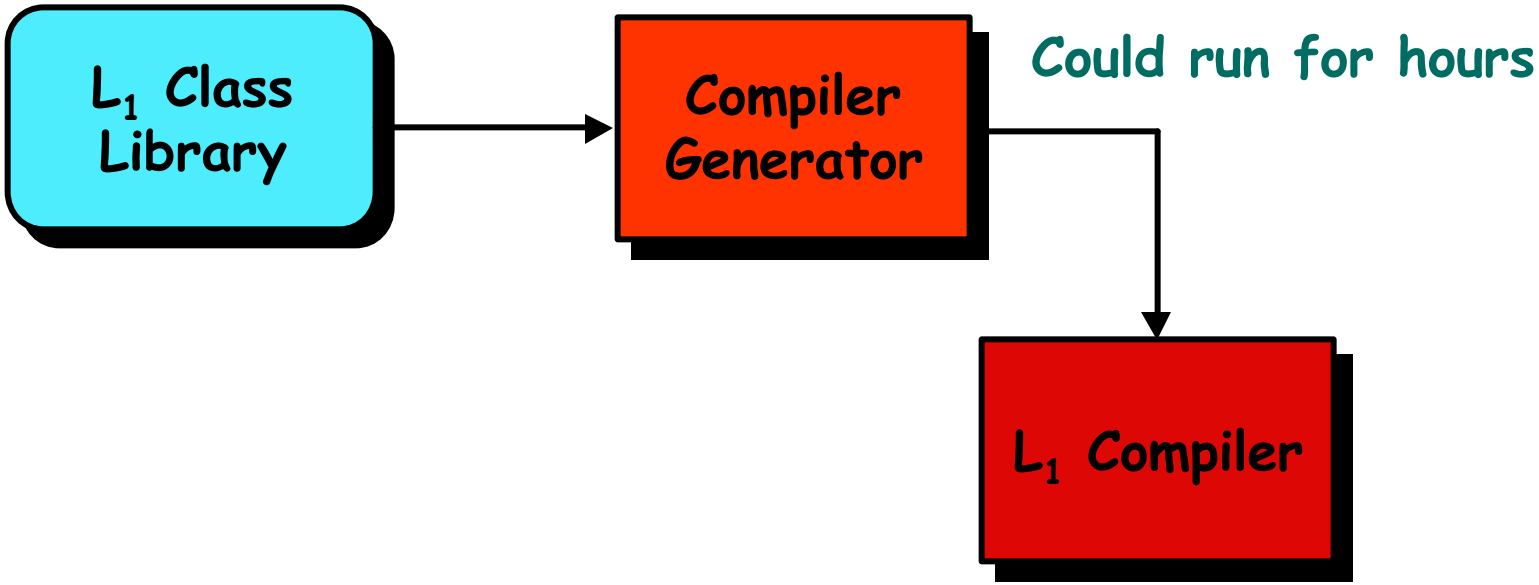

Telescoping Languages

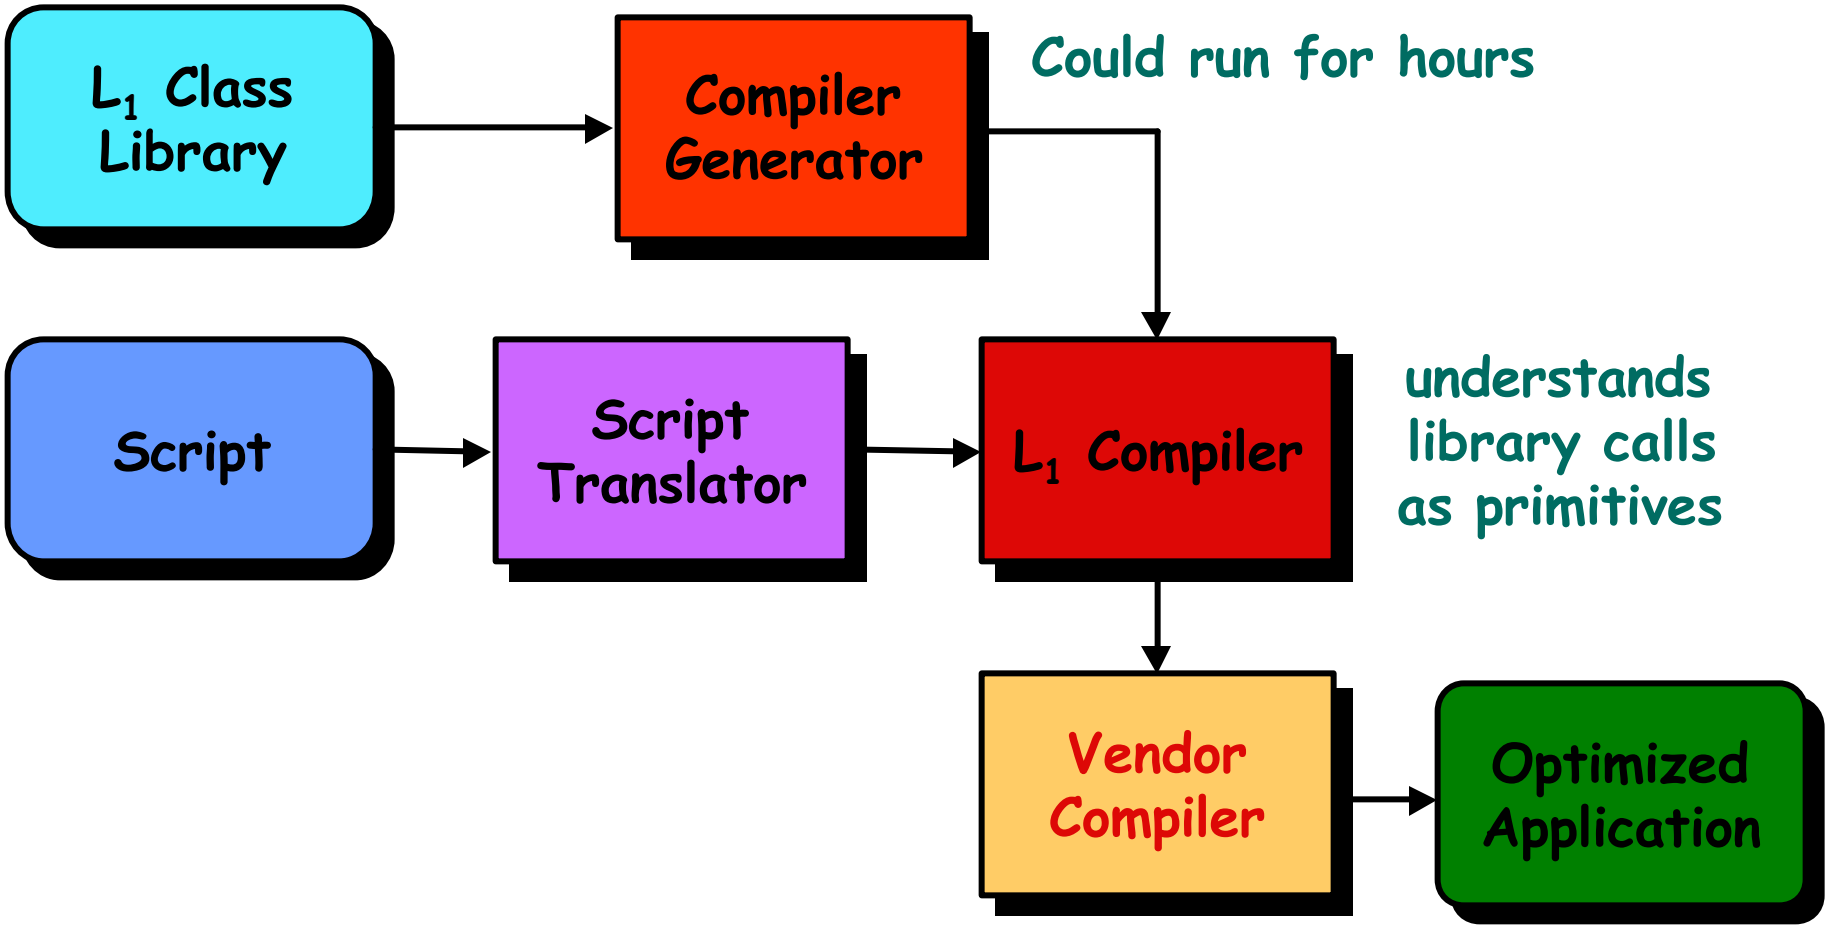

Telescoping Languages: Advantages

- Compile times can be reasonable
  - More compilation time can be spent on libraries
    - Amortized over many uses
  - Script compilations can be fast
    - Components reused from scripts may be included in libraries

Telescoping Languages: Advantages

- Compile times can be reasonable
  - More compilation time can be spent on libraries
    - Amortized over many uses
  - Script compilations can be fast
    - Components reused from scripts may be included in libraries
- High-level optimizations can be included
  - Based on specifications of the library designer
    - Properties often cannot be determined by compilers
    - Properties may be hidden after low-level code generation

Telescoping Languages: Advantages

- Compile times can be reasonable
  - More compilation time can be spent on libraries
    - Amortized over many uses
  - Script compilations can be fast
    - Components reused from scripts may be included in libraries
- High-level optimizations can be included
  - Based on specifications of the library designer
    - Properties often cannot be determined by compilers
    - Properties may be hidden after low-level code generation
- User retains substantive control over language performance
  - Mature code can be built into a library and incorporated into language

Challenges for the SIAM Community

- Theories of security and reliability
  - Most of the current tools are based on mathematics
    - public-key encryption
  - New frontier: Software

Challenges for the SIAM Community

- Theories of security and reliability
  - —Most of the current tools are based on mathematics
    - public-key encryption
  - —New frontier: Software
- Theories of scalable, reconfigurable communication networks
  - —Routing, load balancing
    - Integer and mixed integer programming

Challenges for the SIAM Community

- Theories of security and reliability
  - —Most of the current tools are based on mathematics
    - public-key encryption
  - —New frontier: Software
- Theories of scalable, reconfigurable communication networks
  - —Routing, load balancing
    - Integer and mixed integer programming
- Robust libraries for scientific computation
  - —Components for problem-solving systems
  - —Latency-tolerant algorithms
  - —Management of accuracy in heterogeneous computer configurations

Challenges for the SIAM Community

- Theories of security and reliability
  - —Most of the current tools are based on mathematics
    - public-key encryption
  - —New frontier: Software
- Theories of scalable, reconfigurable communication networks
  - —Routing, load balancing
    - Integer and mixed integer programming
- Robust libraries for scientific computation
  - —Components for problem-solving systems
  - —Latency-tolerant algorithms
  - —Management of accuracy in heterogeneous computer configurations
- Transmission of high-quality video
  - —Compression, compression, compression

Summary

- IT Research Needs Revitalization
  - Focus on the long term

http://www.cs.rice.edu/~ken/Presentations/SIAMPITAC.pdf

Summary

- IT Research Needs Revitalization
  - Focus on the long term
  - University research critical
    - for both ideas and human resources
    - government support role is essential

http://www.cs.rice.edu/~ken/Presentations/SIAMPITAC.pdf

Summary

- IT Research Needs Revitalization
  - Focus on the long term
  - University research critical
    - for both ideas and human resources
    - government support role is essential
  - Areas of focus
    - software, scalable infrastructure, high performance computing, societal and economic impact

http://www.cs.rice.edu/~ken/Presentations/SIAMPITAC.pdf

Summary

- IT Research Needs Revitalization
  - Focus on the long term
  - University research critical
    - for both ideas and human resources
    - government support role is essential
  - Areas of focus
    - software, scalable infrastructure, high performance computing, societal and economic impact
- IT Grand Challenges
  - Software productivity, software reliability, internet security and reliability, telepresence, Grid computing

http://www.cs.rice.edu/~ken/Presentations/SIAMPITAC.pdf

Summary

- IT Research Needs Revitalization
  - Focus on the long term
  - University research critical
    - for both ideas and human resources
    - government support role is essential
  - Areas of focus
    - software, scalable infrastructure, high performance computing, societal and economic impact
- IT Grand Challenges
  - Software productivity, software reliability, internet security and reliability, telepresence, Grid computing
- Critical Role for the SIAM Community

http://www.cs.rice.edu/~ken/Presentations/SIAMPITAC.pdf