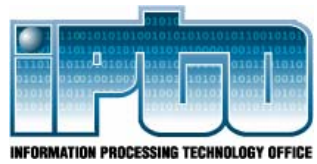

Real Productivity
The Next HPC Frontier

Salishan High-Speed Computing Conference

Robert Graybill
DARPA IPTO
April 22, 2004

What is in: HPC Productivity "Value"
What is out: HPC Peak FLOPS

Parallel Vector Systems

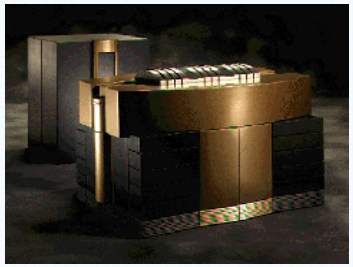

1980's Technology

Vector

Commodity HPCs

Tightly Coupled Parallel Systems

HPCS

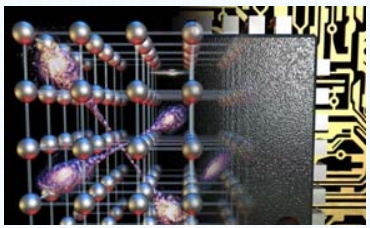

2010 High-End Computing Solutions

Moore's Law Double Raw Performance every 18 Months

New Goal: Double Value Every 18 Months

Industry Accepted Metrics Drive End Products - Time to Augment TOP 500 Criteria

Reason # 1: Technology Near End of Life

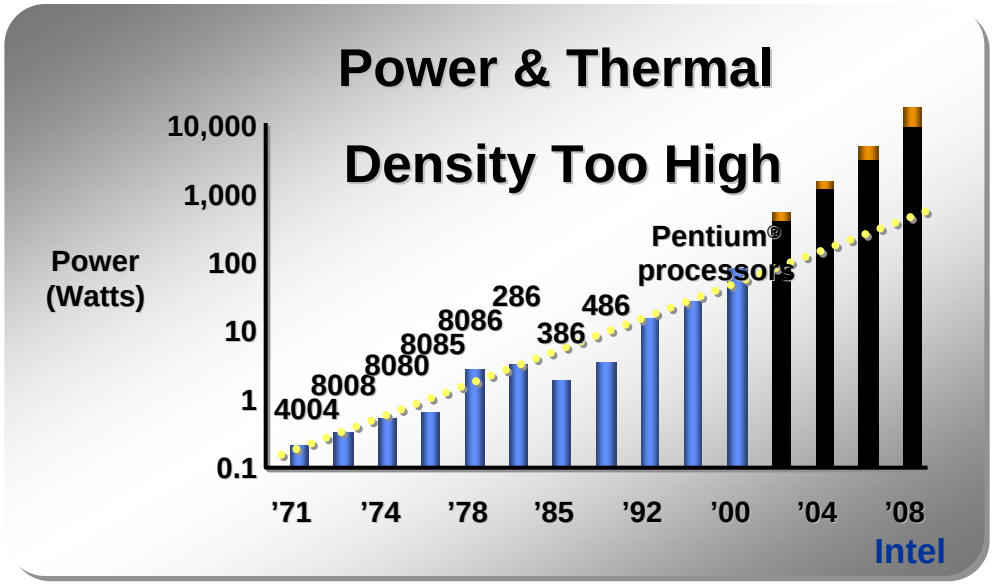

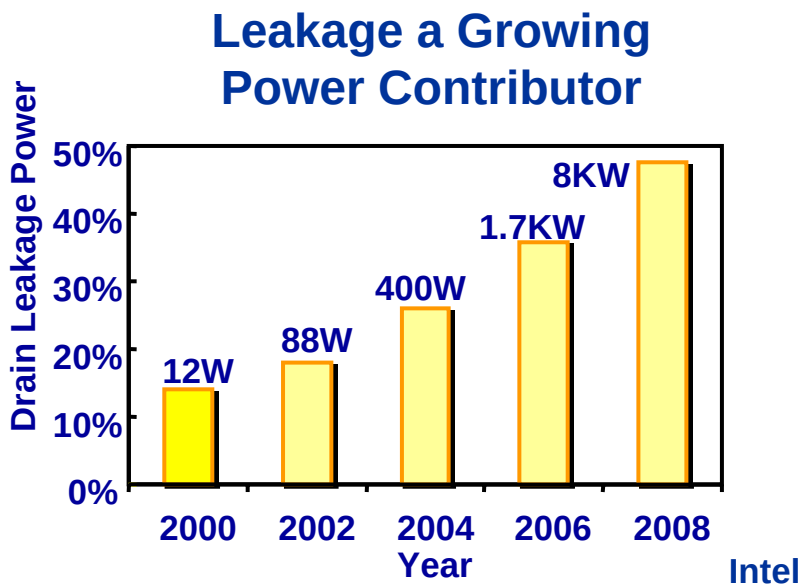

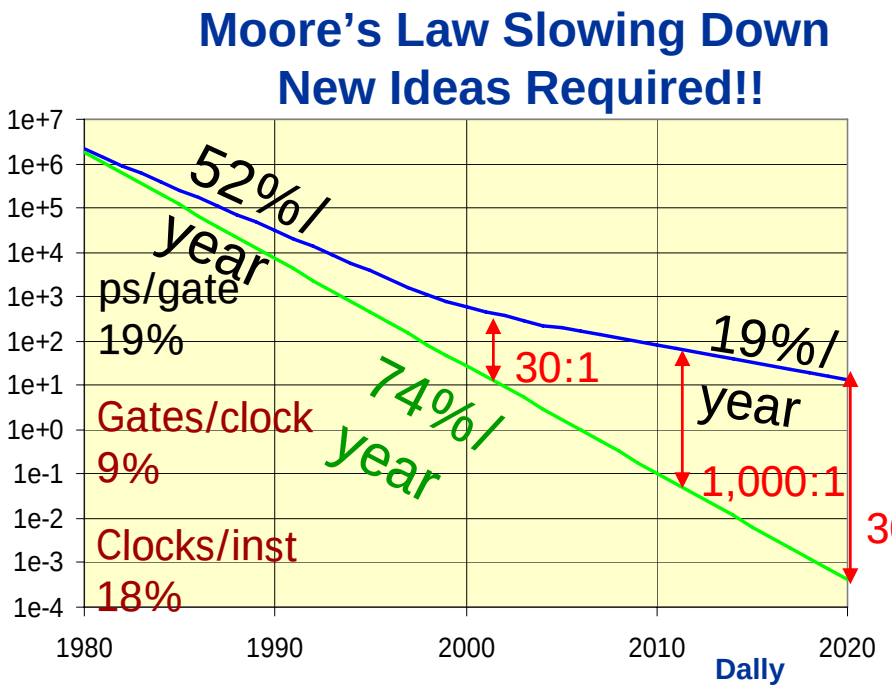

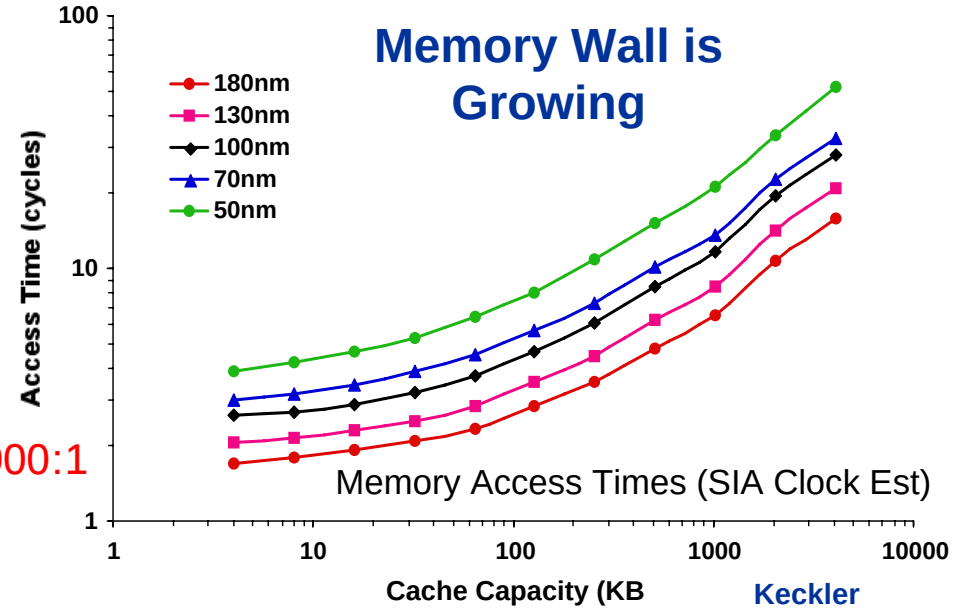

Reason # 2: Loss of a Primary Technology Performance Driver

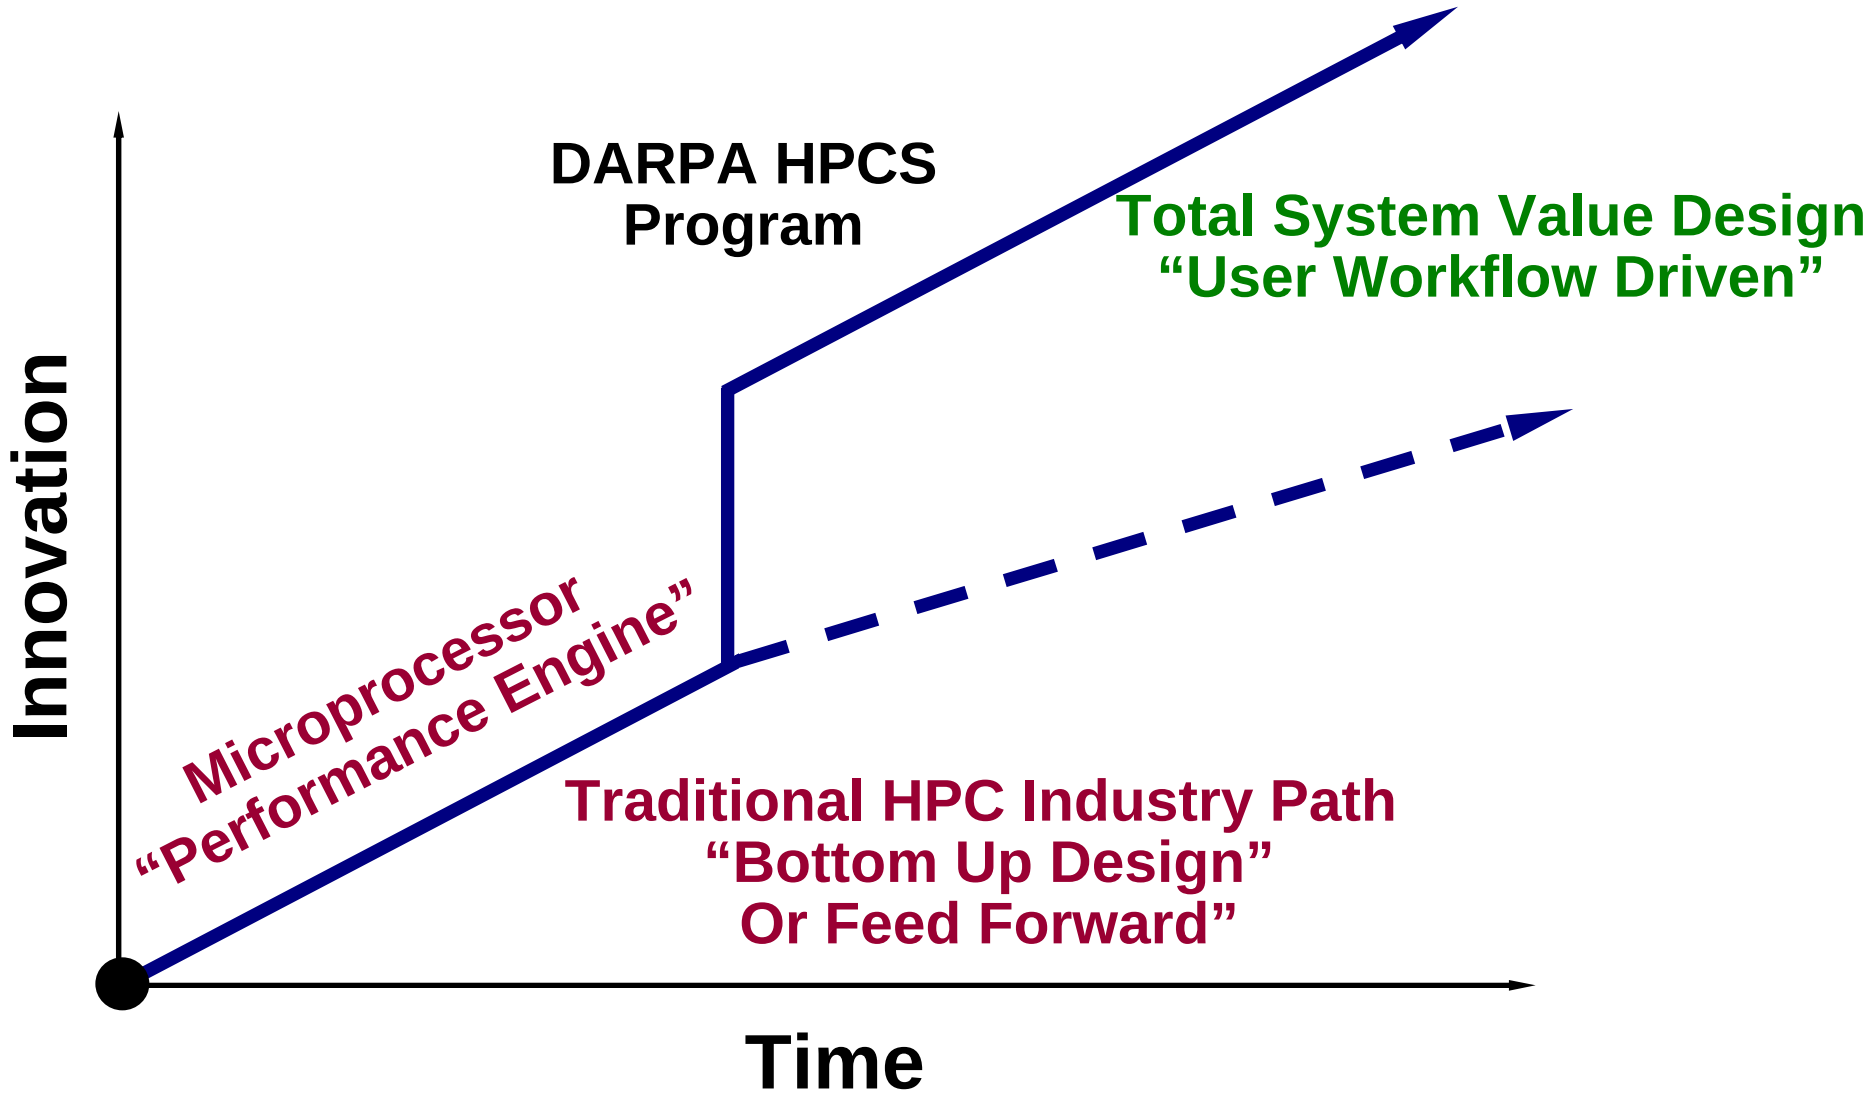

Revolutionary - - System Engineering is Required!

Reason # 3: Growing HPC " Time to Solution Challenge"

- Today: Application development roles
  - Domain Expert
  - Programmers
  - Correctness Expert
  - HPC System Expert
- Future: Peer reviewed functional components to optimized parallel code to final solution
- Compute in the presence of failures and achieve .9999 in reliability

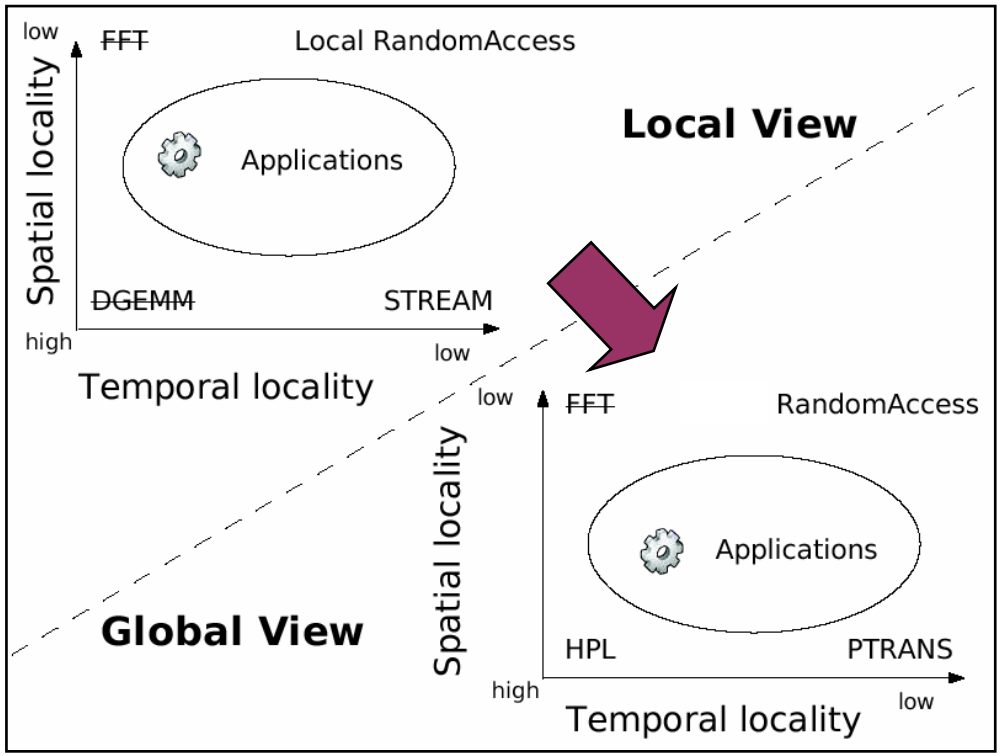

Movement from tens to thousands today to hundred thousand processors in the future

Reason # 4: Growing Application to HPC Architecture Mismatch

DDR&E Study Applications

- Operational weather and ocean forecasting
- Planning activities for dispersion of airborne/waterborne contaminants
- Cryptanalysis
- Intelligence, surveillance, reconnaissance
- Improved armor design
- Engineering design of large aircraft, ship and structures
- National missile defense
- Test and evaluation
- Weapon (warheads and penetrators)
- Survivability/stealth design

IHEC Study Applications

- Comprehensive Aerospace Vehicle Design
- Signals Intelligence (Crypt)
- Signals Intelligence (Graph)
- Operational Weather/Ocean Forecasting
- Stealthy Ship Design
- Nuclear Weapons Stockpile Stewardship
- Signal and Image Processing
- Army Future Combat Systems
- Electromagnetic Weapons Development
- Geospatial Intelligence
- Threat Weapon Systems Characterization

High Productivity Computing Systems

Goal:

- **Provide a new generation of economically viable high productivity computing systems for the national security and industrial user community (2009 – 2010)**

Impact:

- **Performance** (time-to-solution): speedup critical national security applications by a factor of 10X to 40X
- **Programmability** (idea-to-first-solution): reduce cost and time of developing application solutions
- **Portability** (transparency): insulate research and operational application software from system
- **Robustness** (reliability): apply all known techniques to **protect against outside attacks**, hardware faults, & programming errors

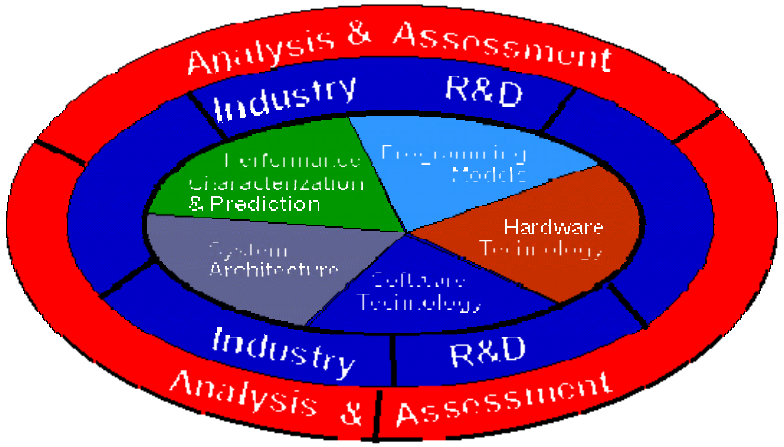

HPCS Program Focus Areas

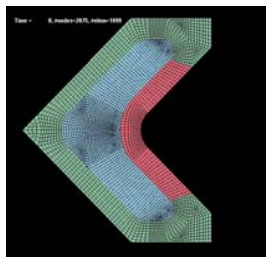

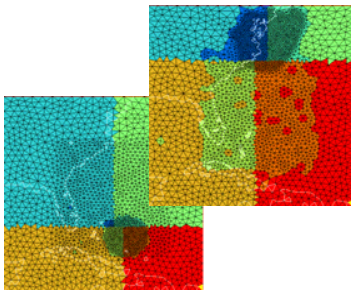

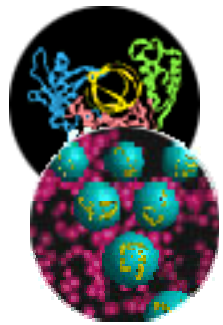

Applications:

- Intelligence/surveillance, reconnaissance, cryptanalysis, weapons analysis, airborne contaminant modeling and biotechnology

Fill the Critical Technology and Capability Gap

Today (late 80's HPC technology).....*to*.....Future (Quantum/Bio Computing)

HPCS Technical Considerations

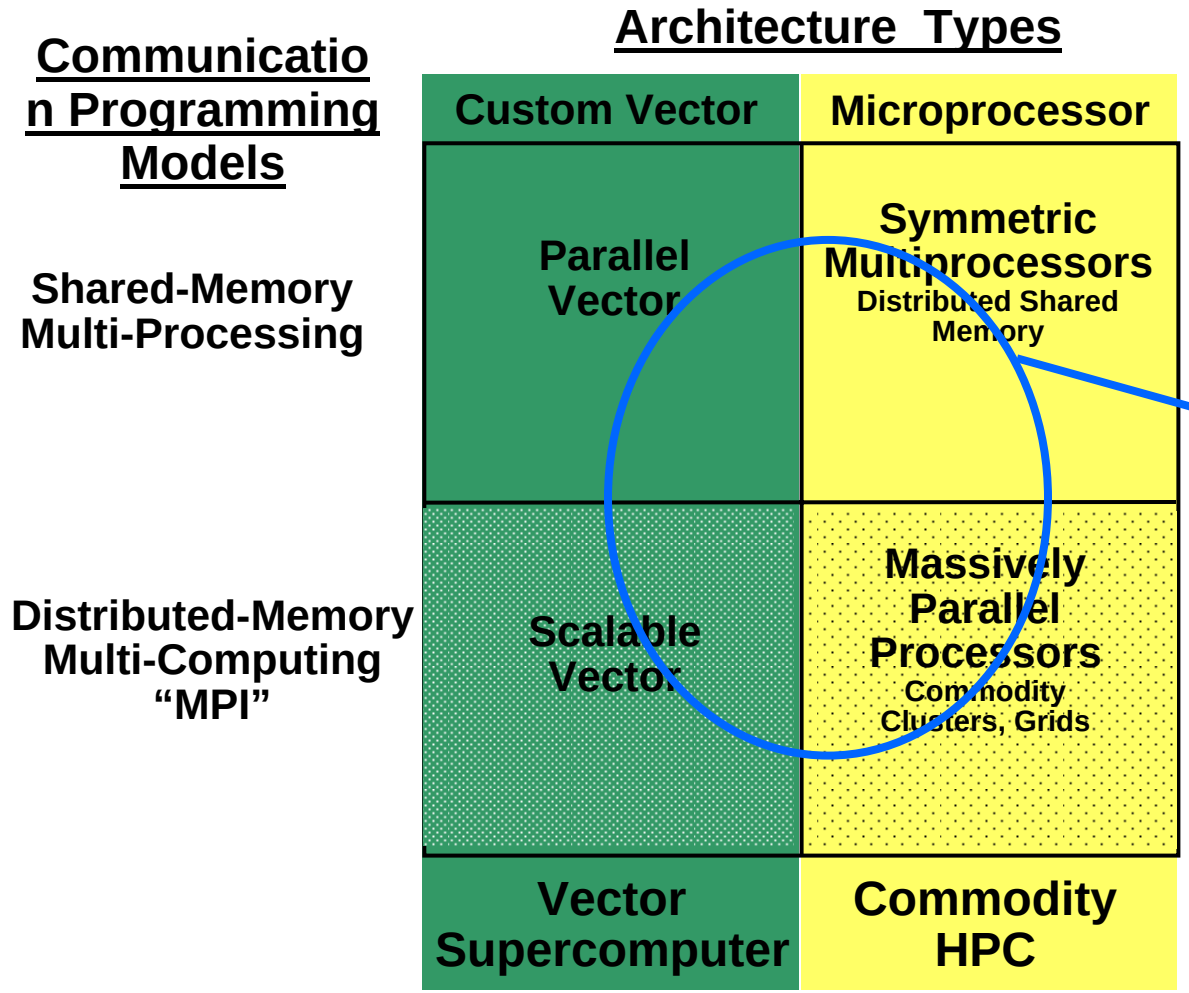

HPCS Focus
Tailorable Balanced Solutions

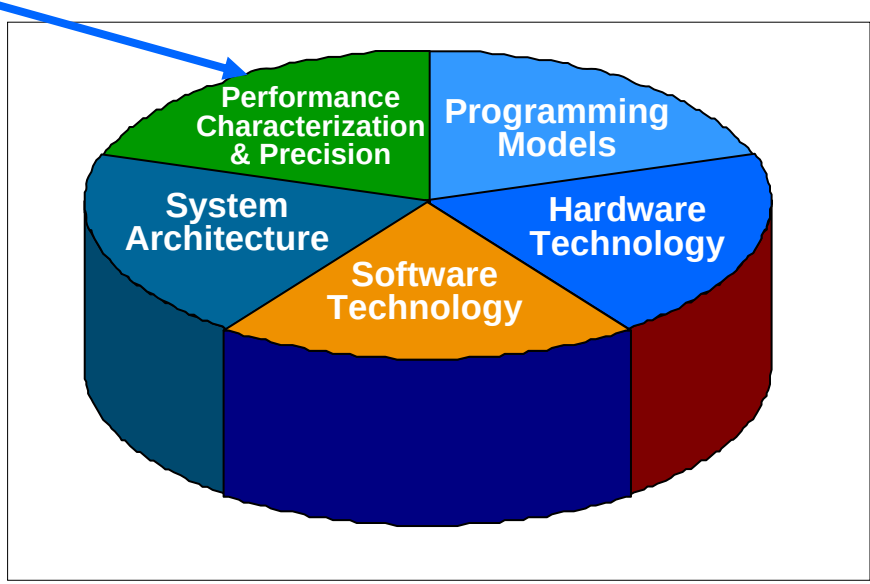

High-End Capability Systems
Need to Adapt to Applications and End Users

HPCS Program Phases I - III

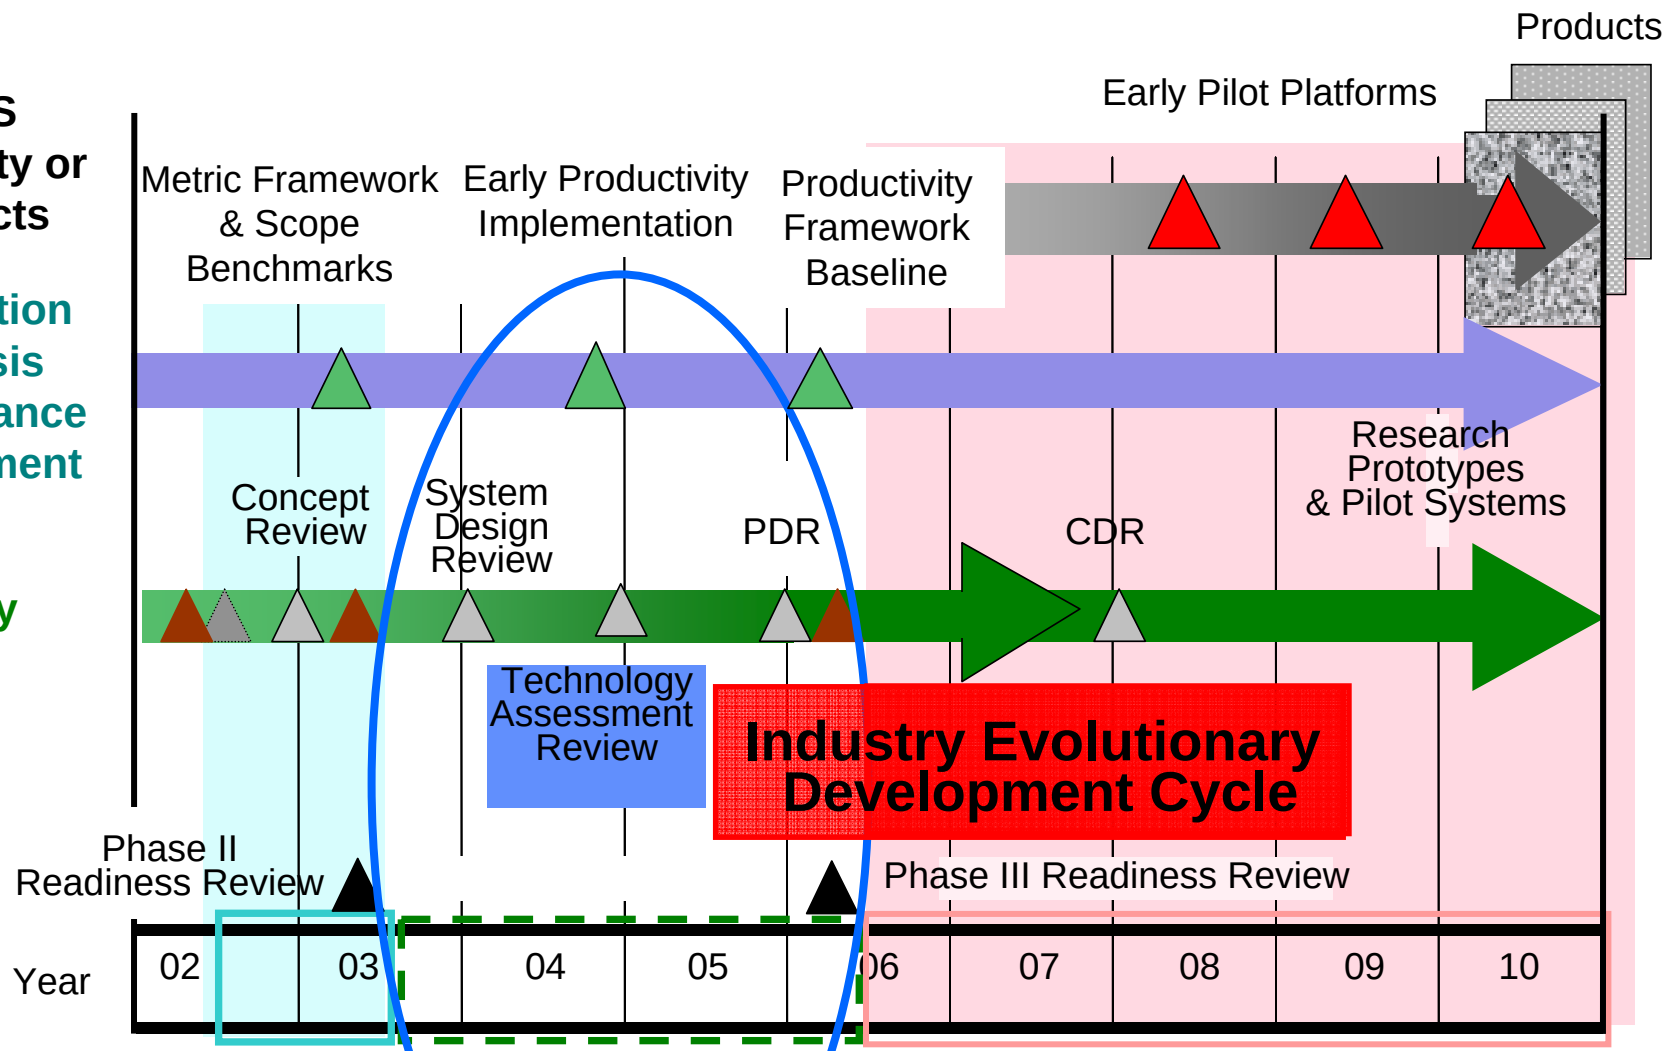

HPCS Phase II Teams

Industry Teams Sponsored by DARPA, NSA, and NNSA

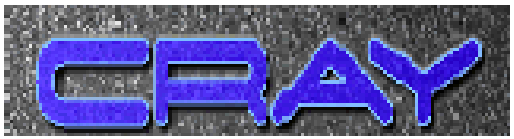

Cray, Inc. (Burton Smith)

International Business Machines Corporation
(Mootaz Elnozahy)

Sun Microsystems, Inc. (Jim Mitchell)

Productivity Team Sponsored by DARPA, DOE, NASA, and NSF

Team Lead: MIT Lincoln Laboratory

Remember our Customers!!

DOD/HPC Modernization Office

DOE

NSA

NNSA

NASA

NRO

NSF

NIH

DHS

NIST

Commercial Sector

Unknown Market Sectors

HPCS Phase II Program Goals

- **Phase II overall productivity goals are as follows:**
  - **Execution** (sustained performance) – 1 Petaflop/sec (scalable to greater than 4 Petaflop/sec). Reference: Workflow 3
  - **Development –** 10X over today's systems. Reference: Workflows 1 and 2
- **Productivity Framework -** Productivity framework that has been base lined for today's systems, successfully used to evaluate the vendors emerging productivity techniques, and provide a solid reference for evaluation of vendor's proposed Phase III designs.
- **Subsystem Performance Indicators –** 3.2 PB/sec bisection bandwidth; 64,000 GUPS; 6.5 PB/sec data streams bandwidth; 2+ PF/s LINPACK

(New!!! HPC Challenge Benchmarks)

Phase II Industry Team Message

Phase II Is Not for "Wimps"

Innovation Now – Technology Freeze in 2006

We have an unique opportunity to impact high performance computing – Let's not drop the ball!!

Productivity Team

Industry:

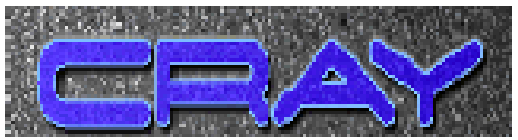

Non-HPCS Industry Members

Mission Partners:

Productivity Team (Lincoln Lead)

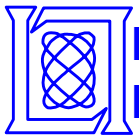

PI: Kepner

PI: Lucas

PI: Basili

PI: Benson & Snavely

PI: Dongarra

PI: Koester

PIs: Vetter, Lusk, Post, Bailey

PIs: Gilbert, Edelman, Ahalt, Mitchell

Productivity = Utility/Cost

$$\Psi \equiv \frac{U}{C} = \frac{U(T)}{C_S + C_O + C_M}$$

Ψ = productivity [utility/$]
U = utility [user specified]
T = time to solution [time]
C = total cost [$]

C_S = software cost [$]
C_O = operation cost [$]
C_M = machine cost [$]
$C_S+C_O+C_M = (c_S+c_O+c_M) \times T$

- **Utility is value user places on getting a result at time T**

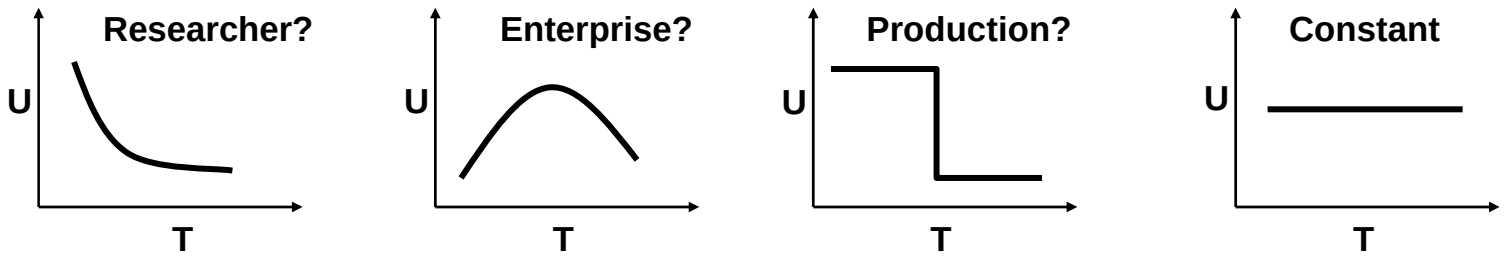

- **T = T(P,Q) and C = C(P,Q) are functions system parameters P and application characteristics Q**

Productivity Is User Specific

Productivity Approach

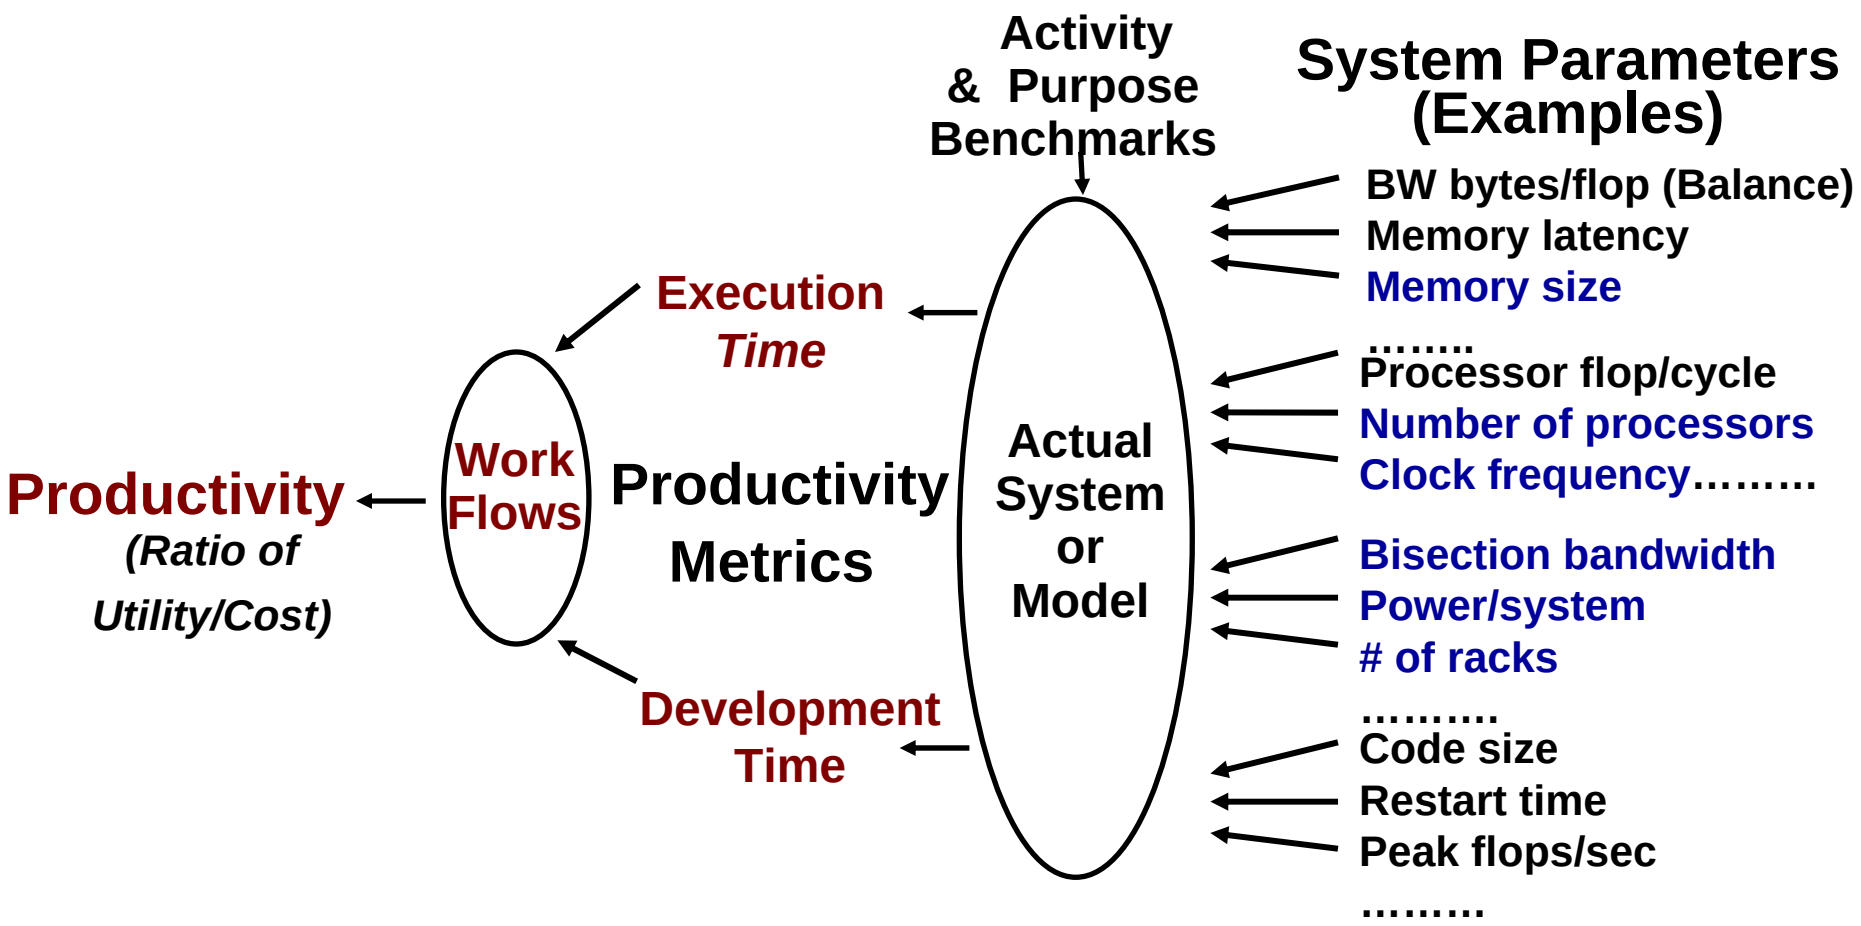

HPCS Mission Work Flows

INFORMATION PROCESSING TECHNOLOGY OFFICE

DARPA

Overall Cycle

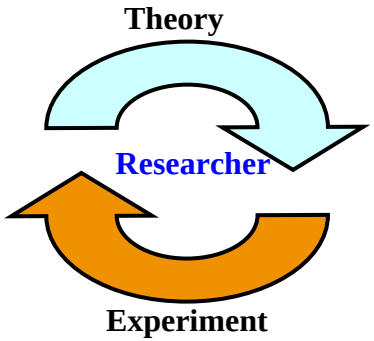

Days to hours

Development

Execution

Development Cycle

Researcher

Hours to minutes

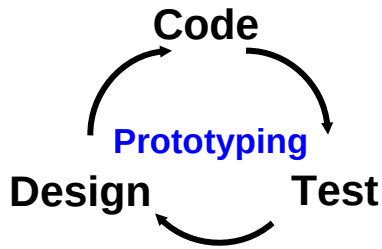

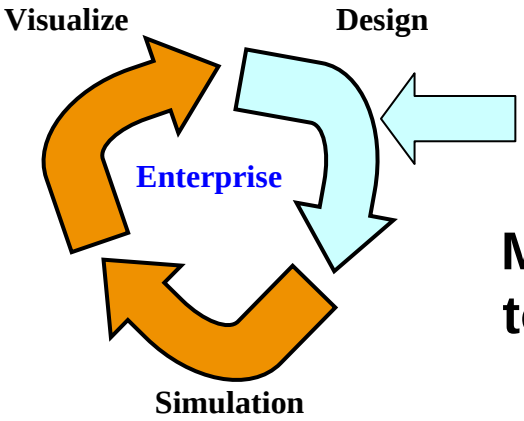

Port Legacy Software

Months to days

Enterprise

Months to days

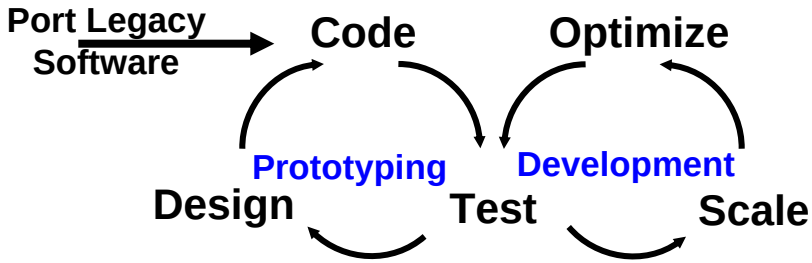

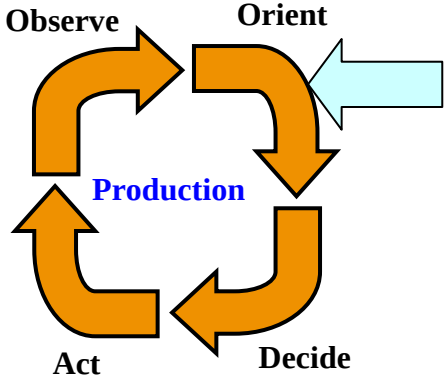

Initial Product Development

Hours to Minutes (Response Time)

Production

Years to months

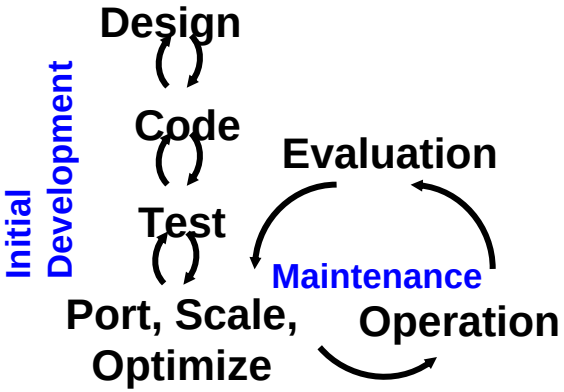

Fourth Work Flow Under Construction
"Administration"

Productivity Team Tasks and Schedule

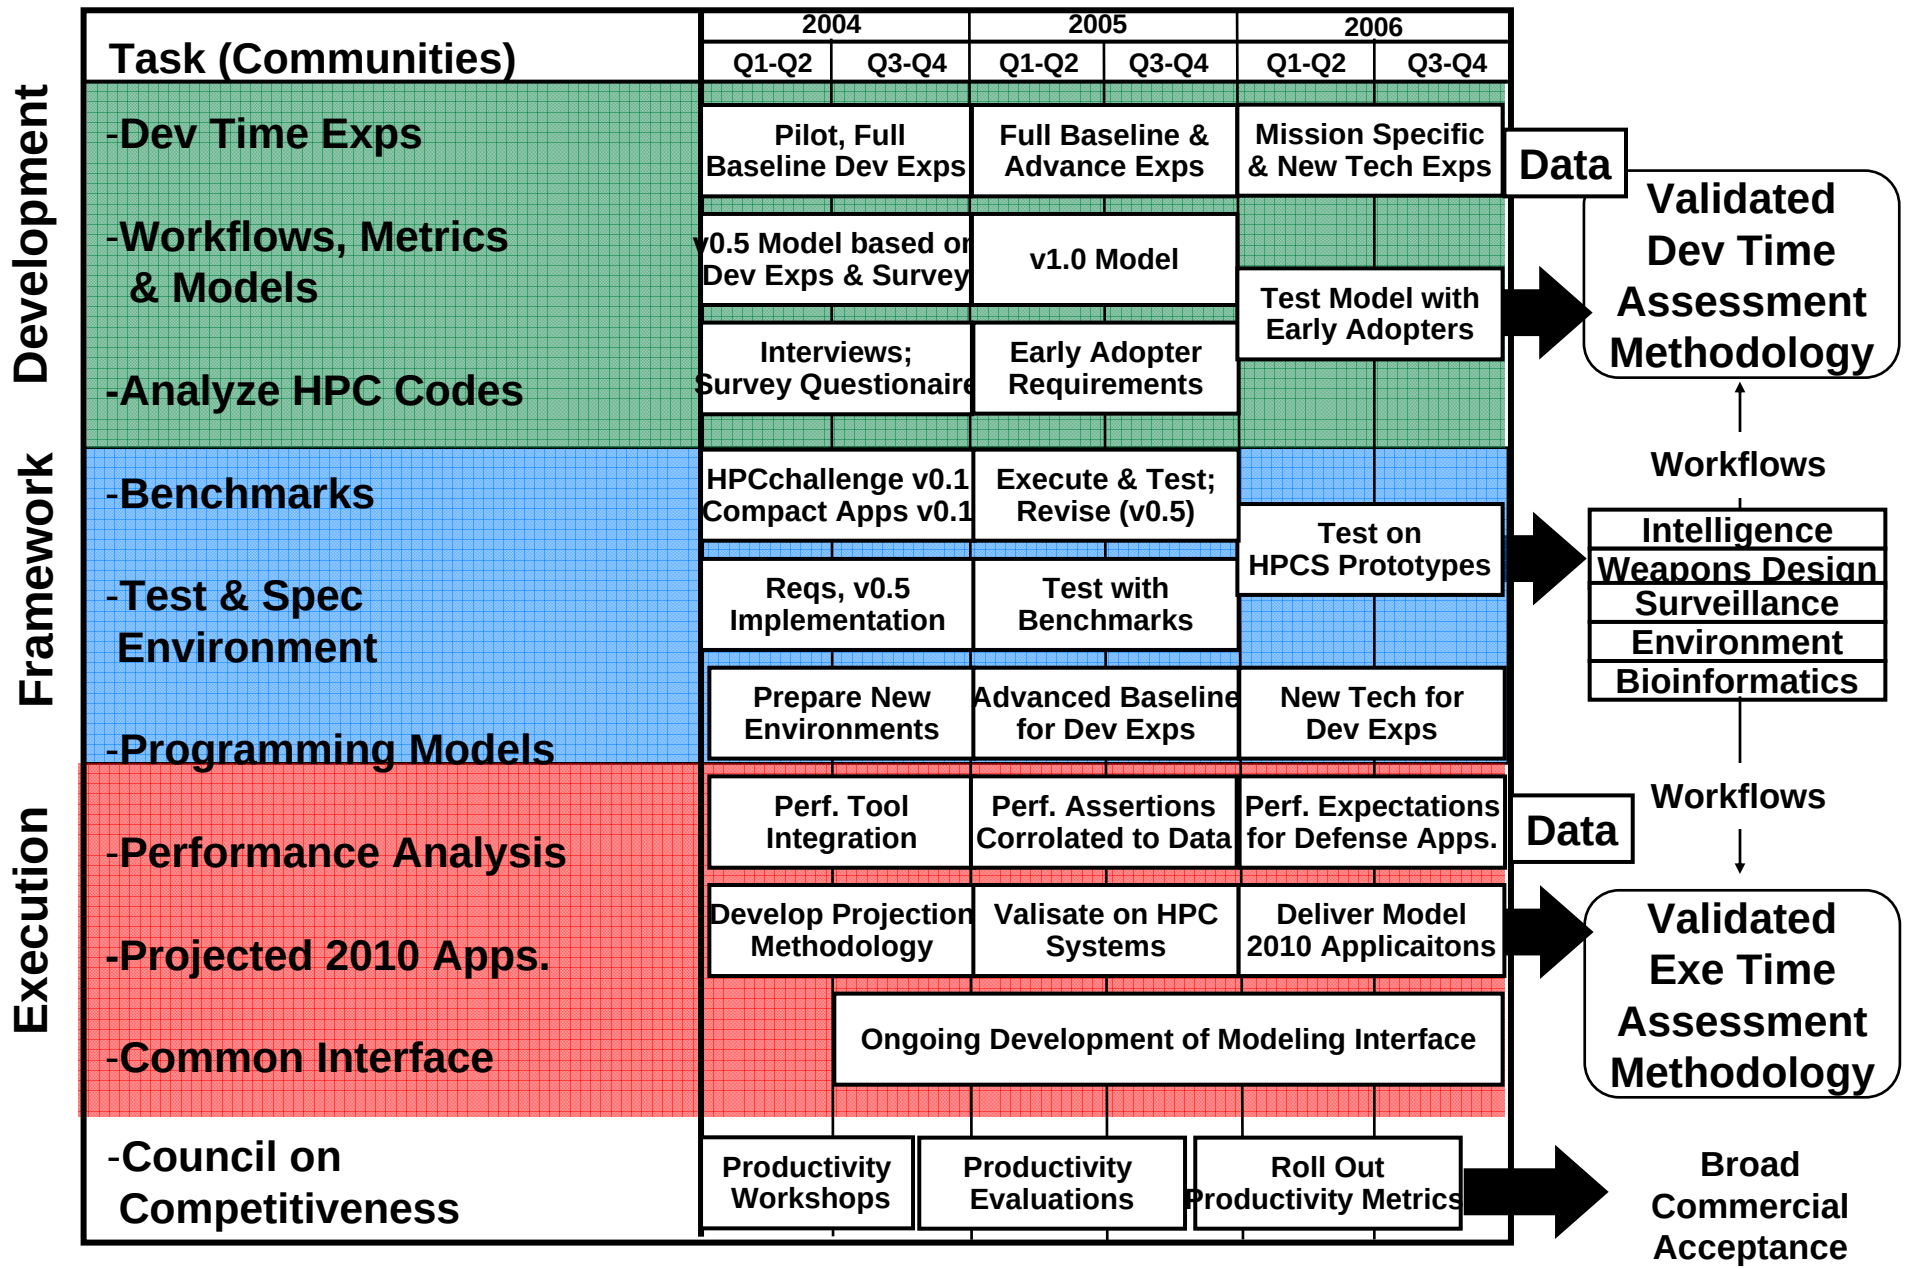

HPCS Benchmark Relationships Preliminary

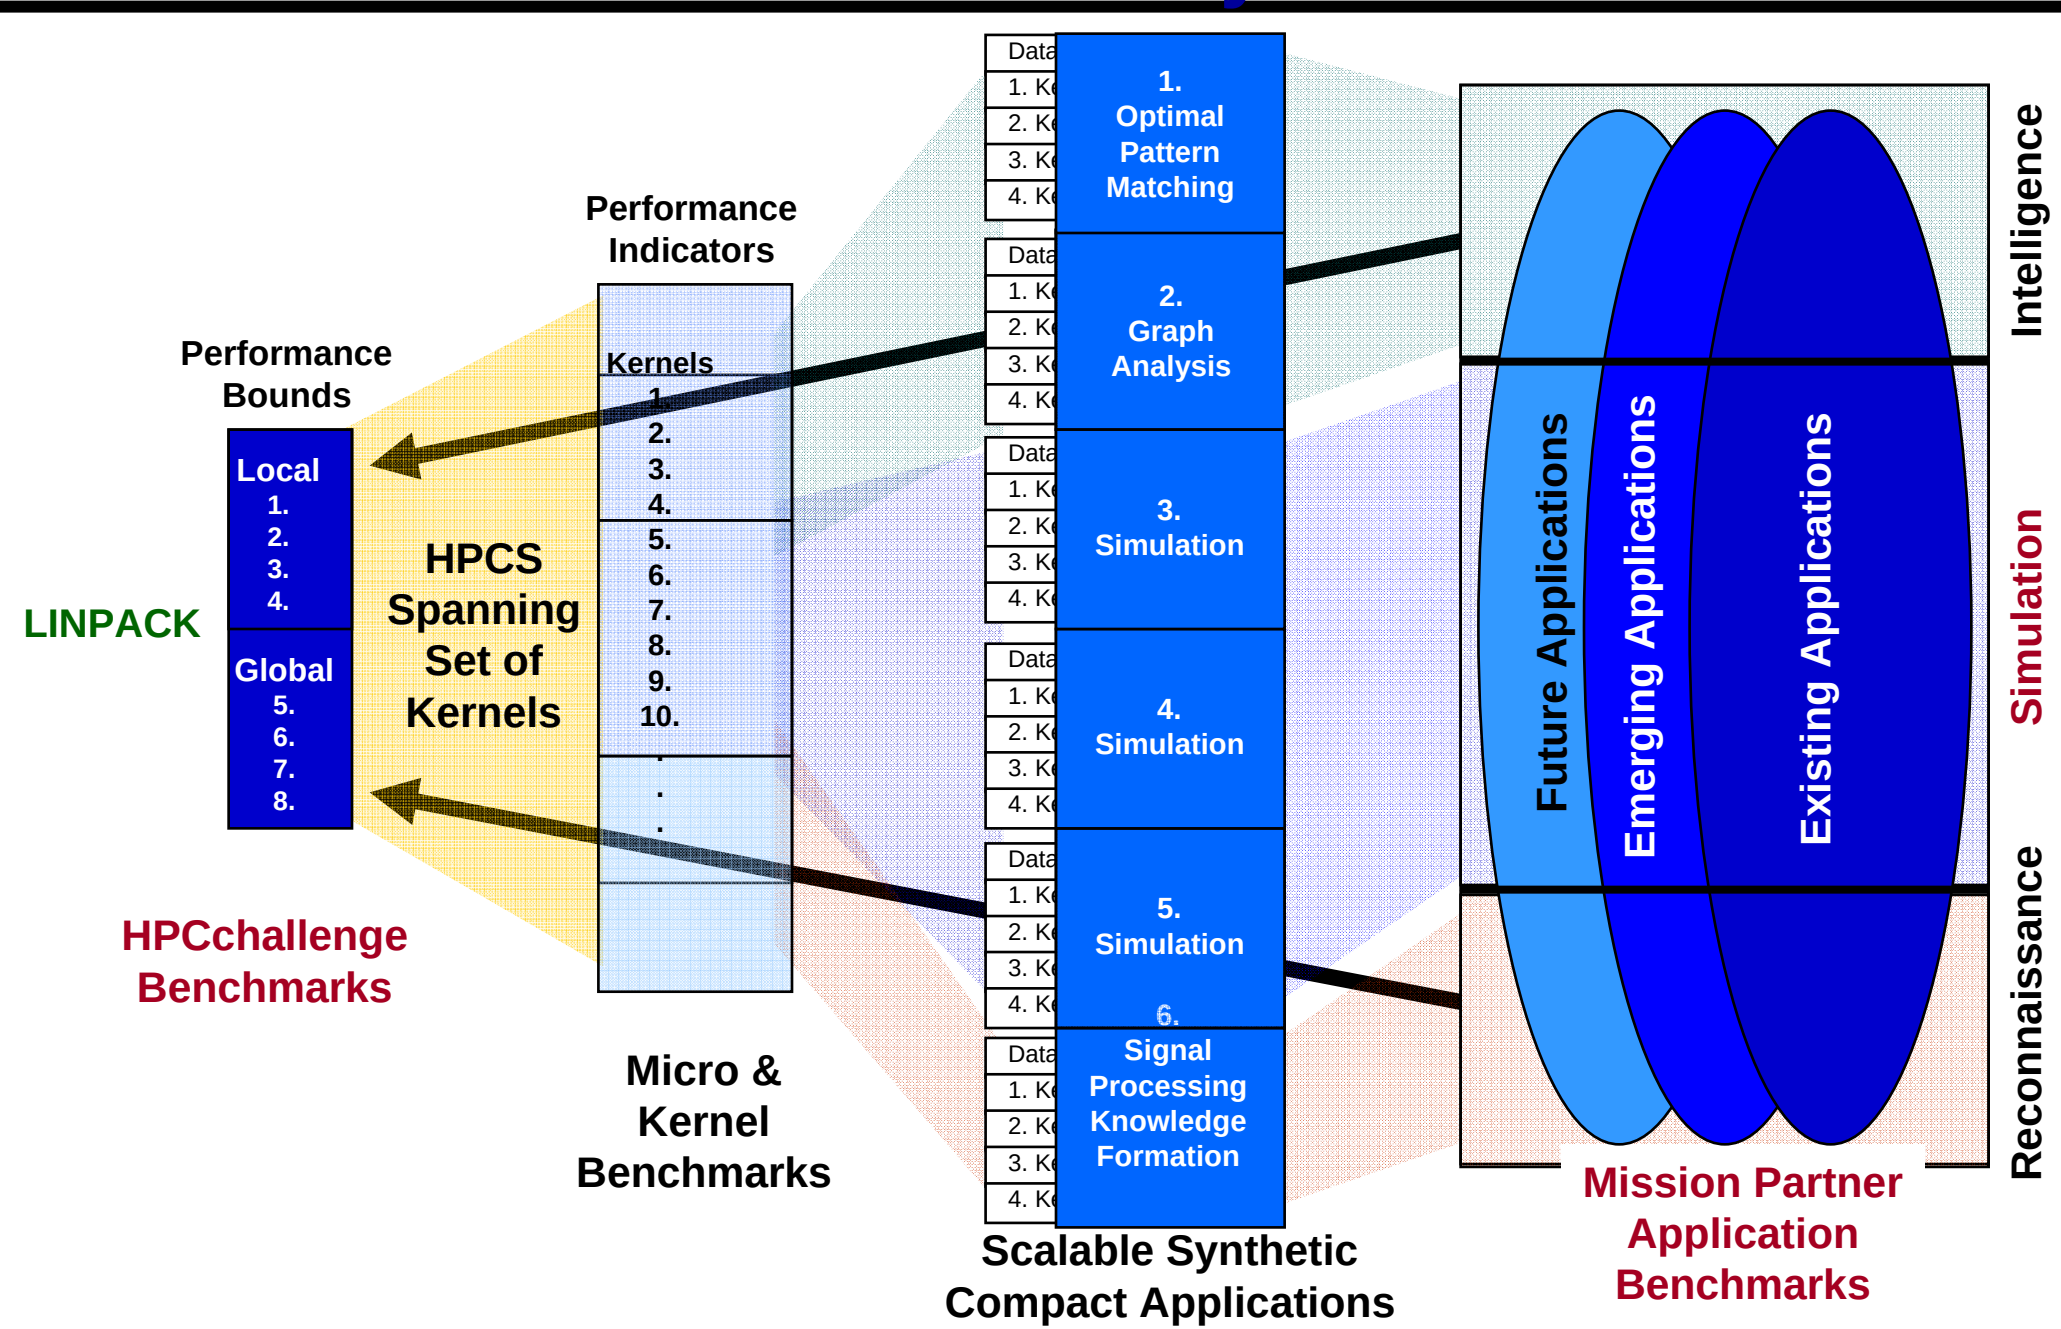

HPCchallenge

Local

- DGEMM (matrix x matrix multiply)
- STREAM
  - COPY
  - SCALE
  - ADD
  - TRIADD
- RandomAccess
- FFT (under development)

Global

- High Performance LINPACK
- PTRANS — parallel matrix transpose
- Random Access
- FFT (under development)

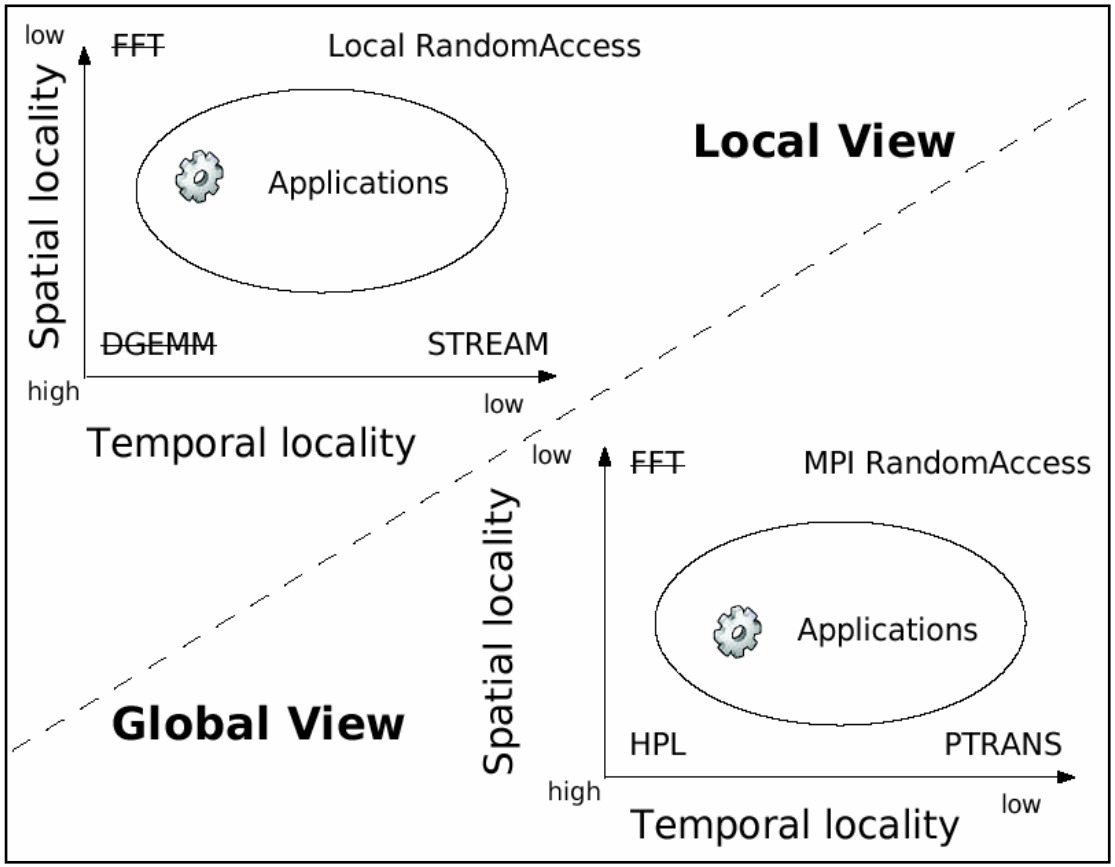

Development Time Pilot Studies

- **Goals:**
  - Give other HPC educators experience in conducting experiments
  - Test the design on a broader range of technologies and contexts
  - Look at hypotheses generated in Pre-Pilot study
  - Formulate some real hypotheses for the full study
- **Locations:**
  - Spring 2004 classes being taught by:
    - Alan Snavely [UCSD]
    - John Gilbert [UCSB]
    - Mary Hall [USC]
    - Alan Edelman [MIT]
    - Uzi Vishkin, Alan Sussman [UMD]
- **Experiment Team:**
  - The studies will be conducted through the joint efforts of the professors and the UMD empirical research team (Vic Basili – UMD)
- **Pre-Pilot Class Assignments (UMD) – Already Completed !**
  1. Develop "Game of Life" from scratch in C using MPI
  2. Parallelize a serial version of a weather prediction code in Fortran using OpenMP

Pre-Pilot Study: Potential Hypothesis

MPI Speedup has more variation than OpenMP Speedup

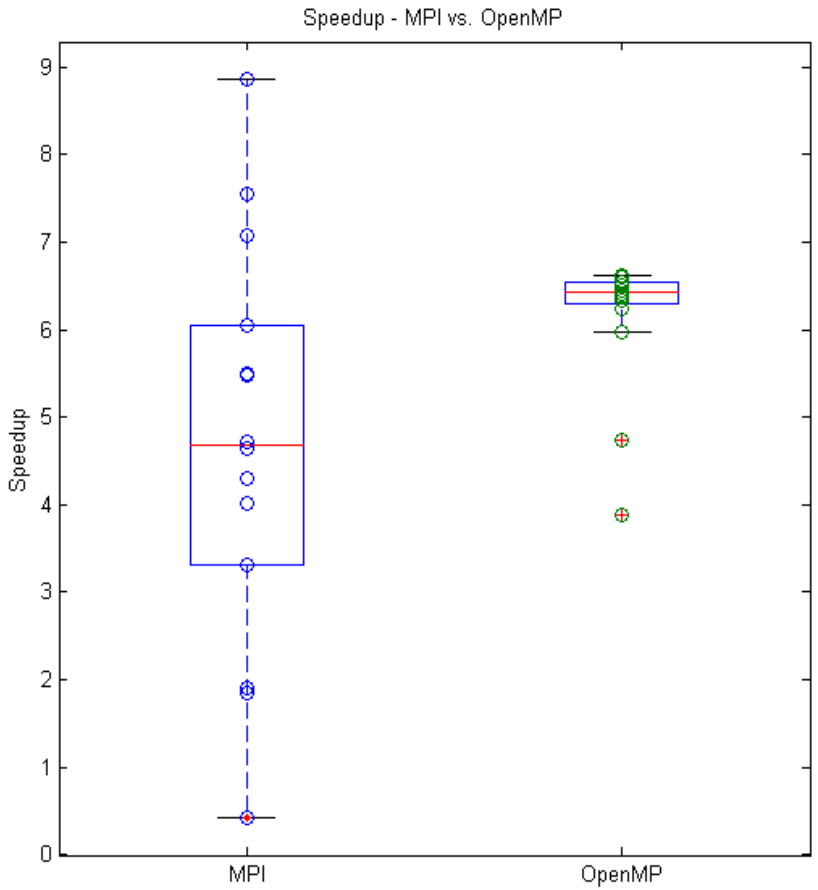

Pre-Pilot Study: Potential Hypothesis

Novices can achieve speedup

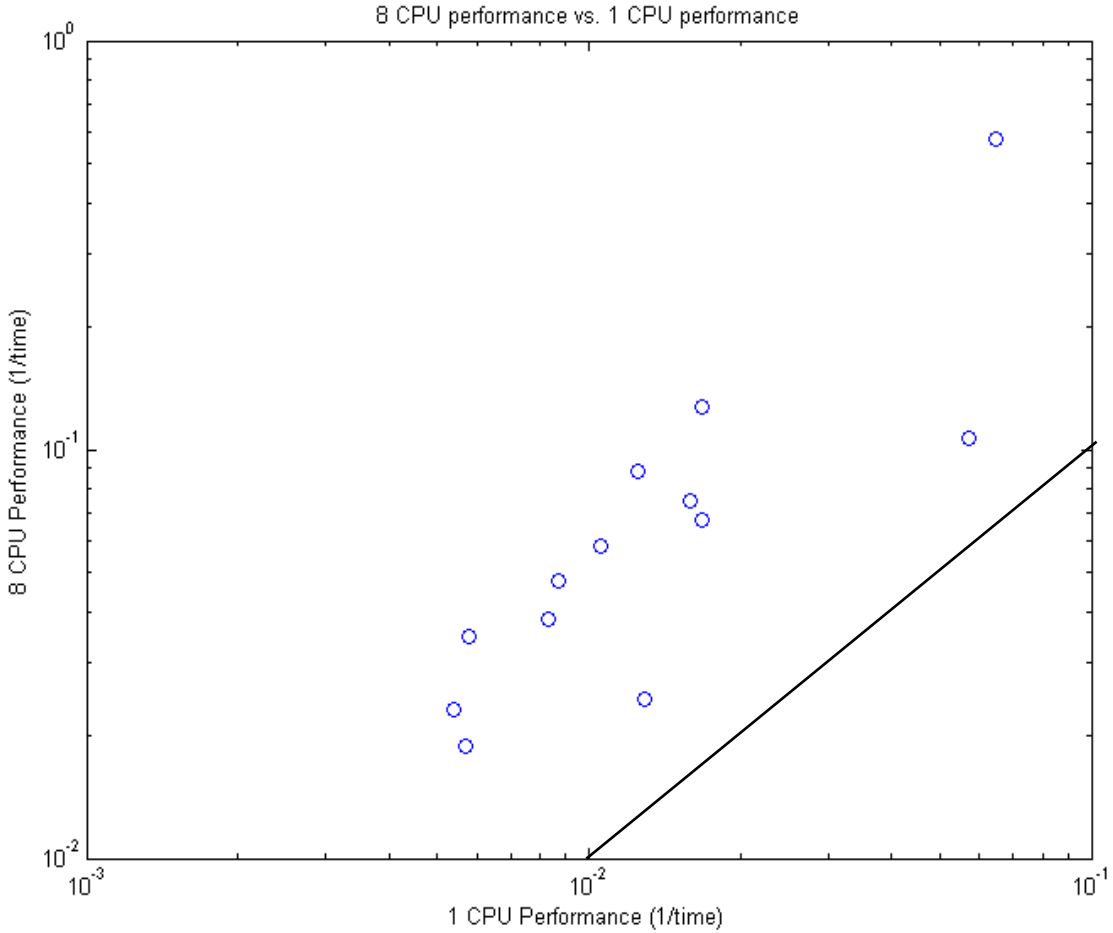

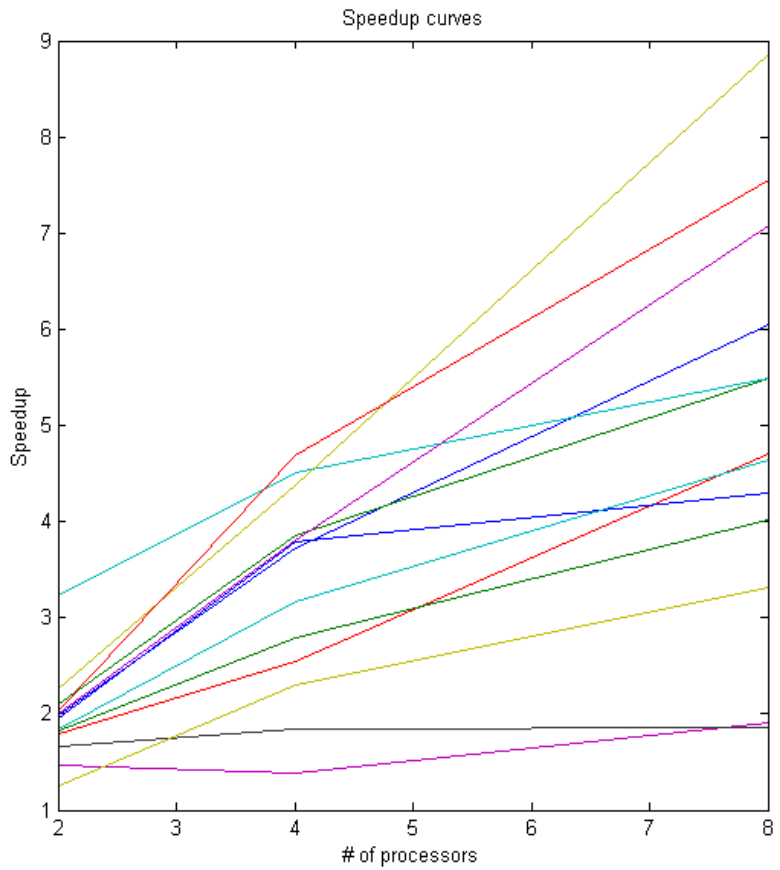

Pre-Pilot Study: Potential Hypothesis

Non CS Majors achieve better speedup than CS Majors

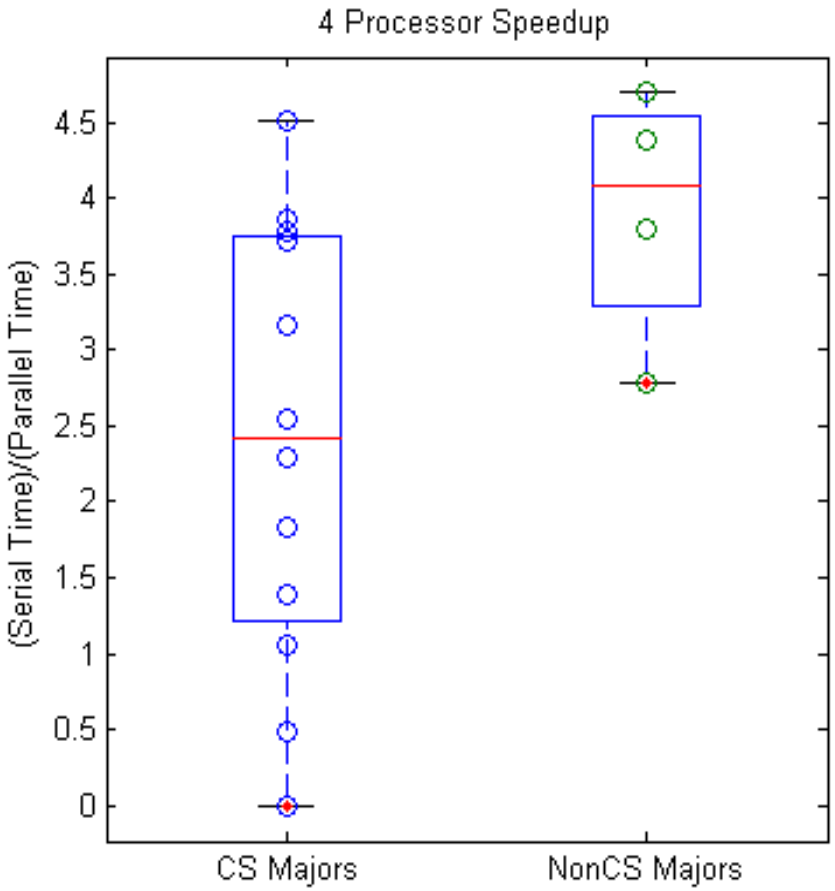

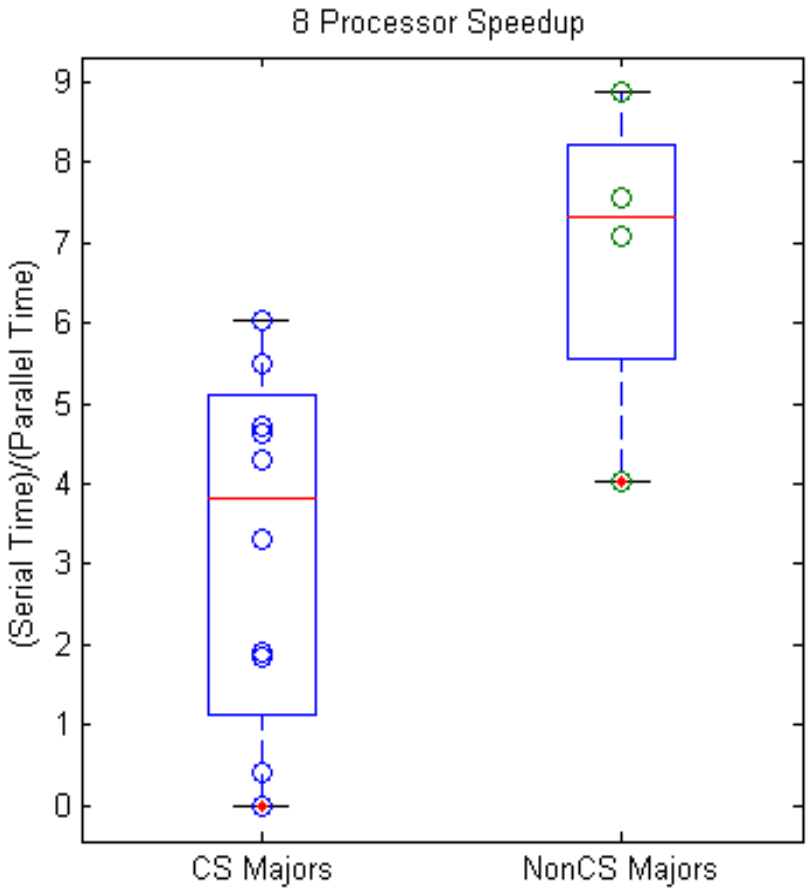

Council on Competitiveness HPC Initiative

Project at a Glance: Funders/Deliverables

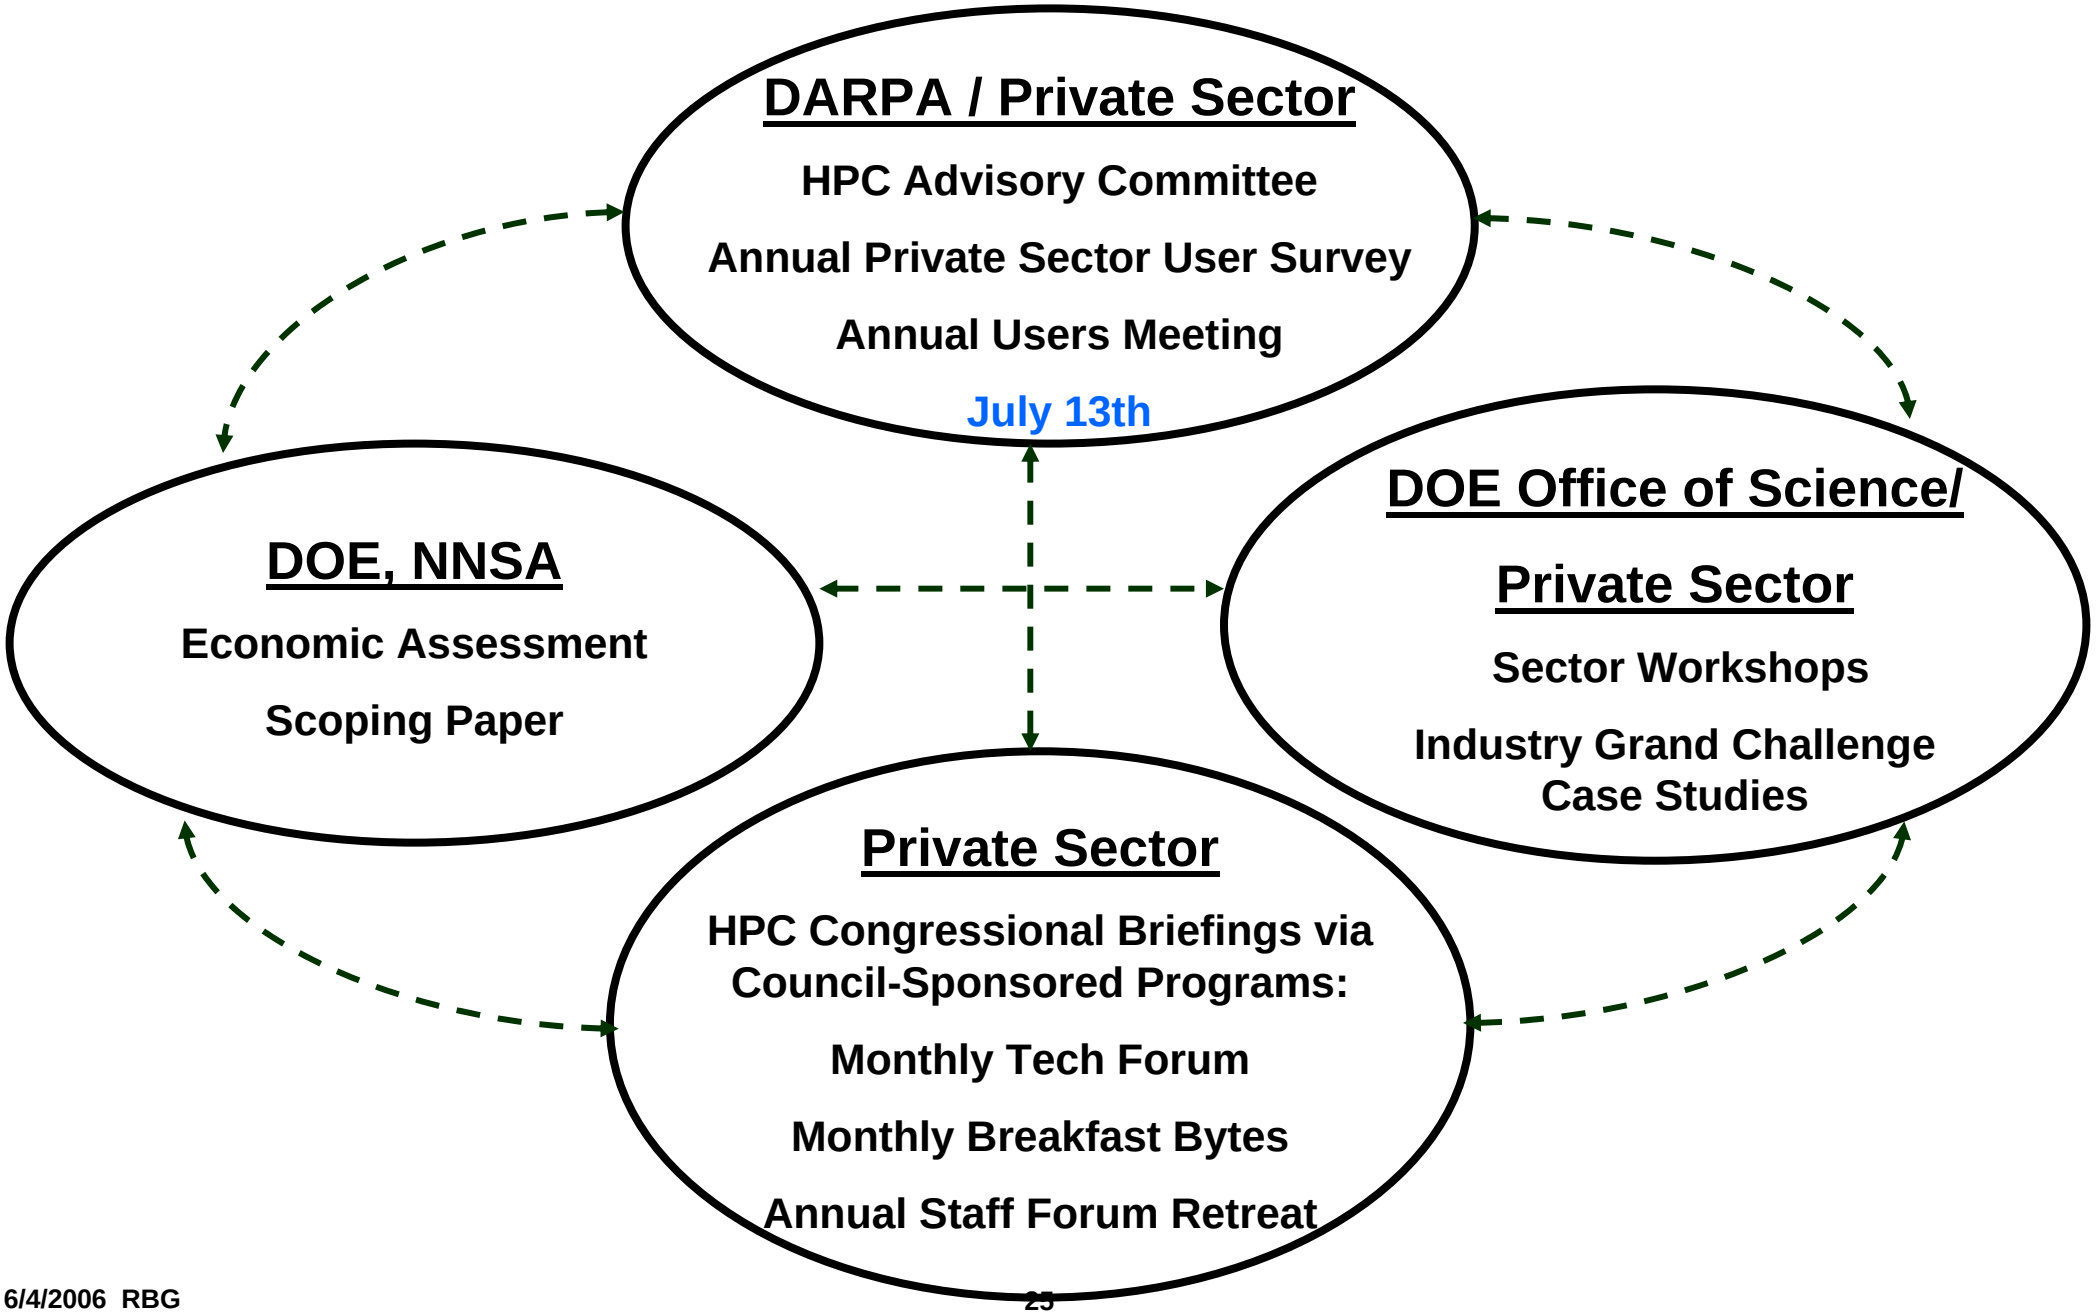

HPCS Activities of Interest

- Next HPCS Productivity Team/Task Group meeting (Target expanded technical audience)
  - June 29-30 2004 (Hyatt – Fair Lakes VA)

- First Annual Council on Competitiveness Conference - HPC: Supercharging U.S. Innovation and Competitiveness (Senior application users, HPC Center Directors, industry executives, and policy/funding decision representatives)
  - July 13, 2004 (Capital Hilton – Washington DC)

???????

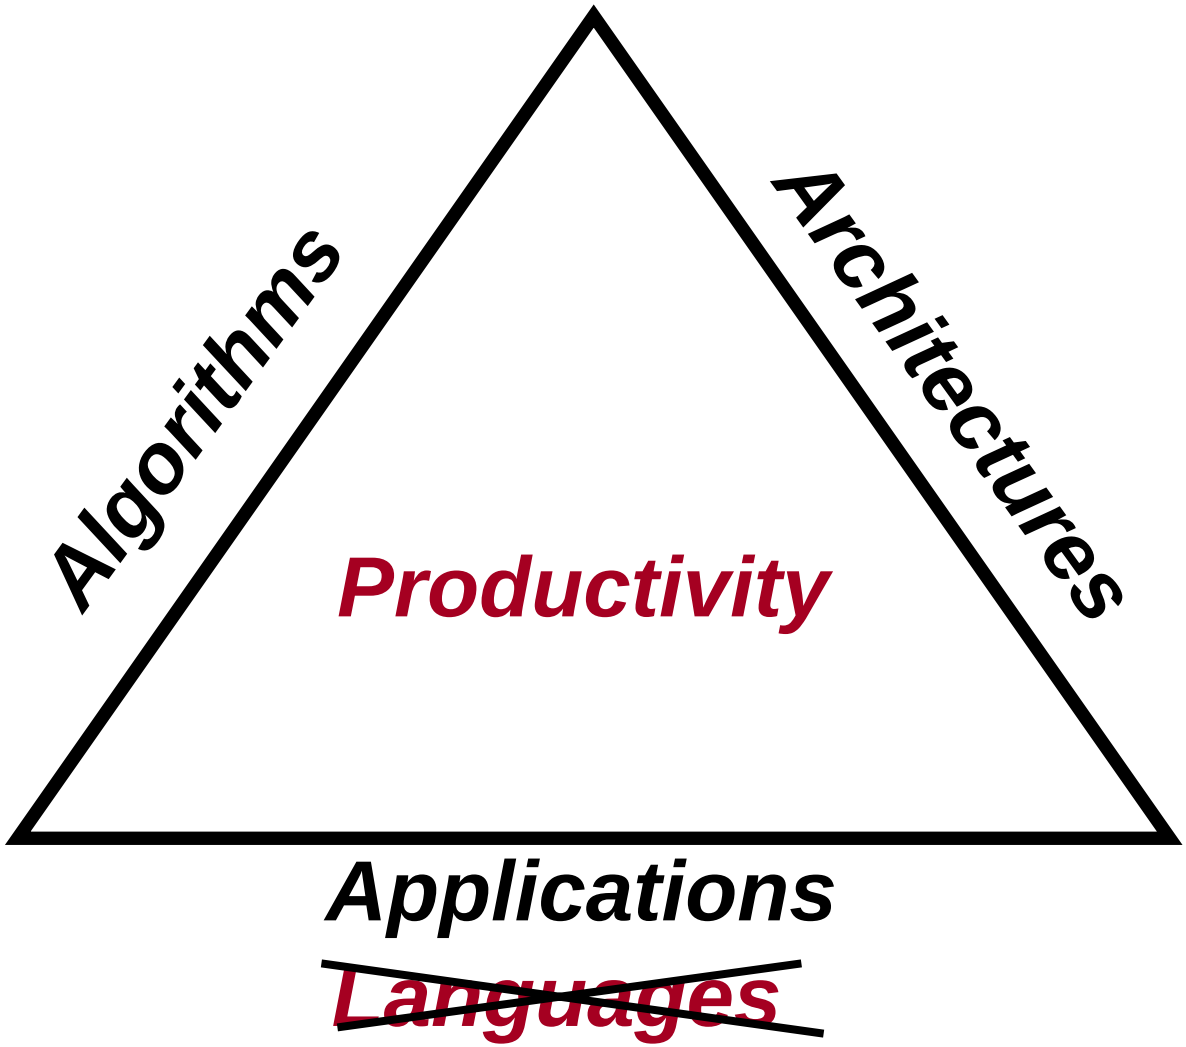
Algorithms
Architectures
Productivity
Applications
Languages

Four Tiers of Agile Processing

Systems That Know What They're Doing

- Intelligent Systems — ACIP
  - Architectures for Cognitive Information Processing (ACIP)

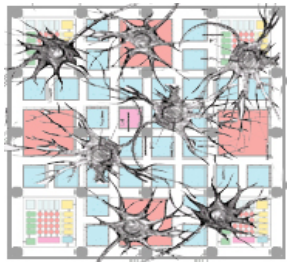

- High-End Application Responsive Computing — HPCS + HECURA (NEW!!!)
  - High Productivity Computing Systems Program (HPCS)

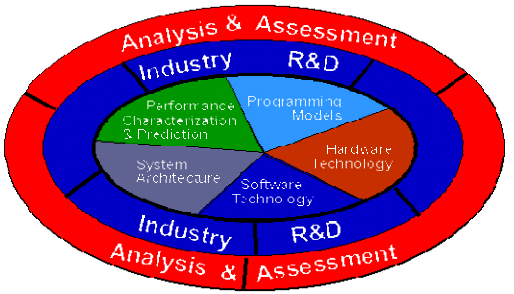

- Mission Responsive Architectures — PCA + OneSAF
  - Polymorphous Computing Architectures Program (PCA)

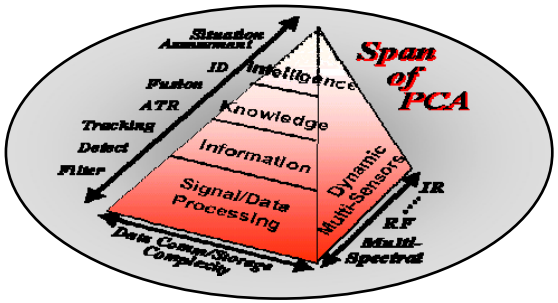

- Power Management — PACC
  - Power Aware Computing and Communications Program (PAC/C)