



---

# Panel: Performance Metrics for HPCS: Holy Grail or Fata Morgana

David Koester / MITRE

Workshop on Performance Characterization, Modeling  
and Benchmarking for HPC Systems

Emeryville, CA

5-7 May 2003



# Indiana Jones and the Last Crusade The Grail



- A Knight of The Order of the Grail bemusedly calls upon seekers of the Holy Grail to 'choose wisely.'



# Fata Morgana



- **The fata morgana is a mirage produced by reflections of light from strong, low level temperature inversions in the atmosphere**
- **Also called a superior mirage, named after King Arthur's wicked stepsister, the sorceress Morgan le Fay, who, according to some legends, lived in a crystal palace under the waves and manifested her magical powers by creating mirages**

<http://jackstephensimages.com/Merchant/photographicgallery/fatamorgana/fatamorganapage.html>  
<http://virtual.finland.fi/finfo/english/mirage.html#shif>



# Fata Morgana Images



## Dalrymple Rock Thule AFB, Greenland



[http://jackstephensimages.com/  
Merchant/photographicgallery/  
fatamorgana/fatamorganapage.html](http://jackstephensimages.com/Merchant/photographicgallery/fatamorgana/fatamorganapage.html)



# Metrics



- The panel topic is *Performance Metrics*
- But what about
  - Programmability?
  - Portability?
  - Robustness?



# The Questions



- **(1) Why don't we have a good metric for HPC performance?  
Do we have one?**
- **(2) Is there any chance to define a single metric for HPC performance?  
Does everybody needs his/her own?**
- **(3) What are the requirements (theoretical, practical, and political) for such metrics?**
- **(4) What needs to be done to get a new metrics accepted?  
If it's possible at all?**
- **(5) "Sustained to Peak" metric and its repercussions**
- **(6) Implications of performance metrics on the political level**



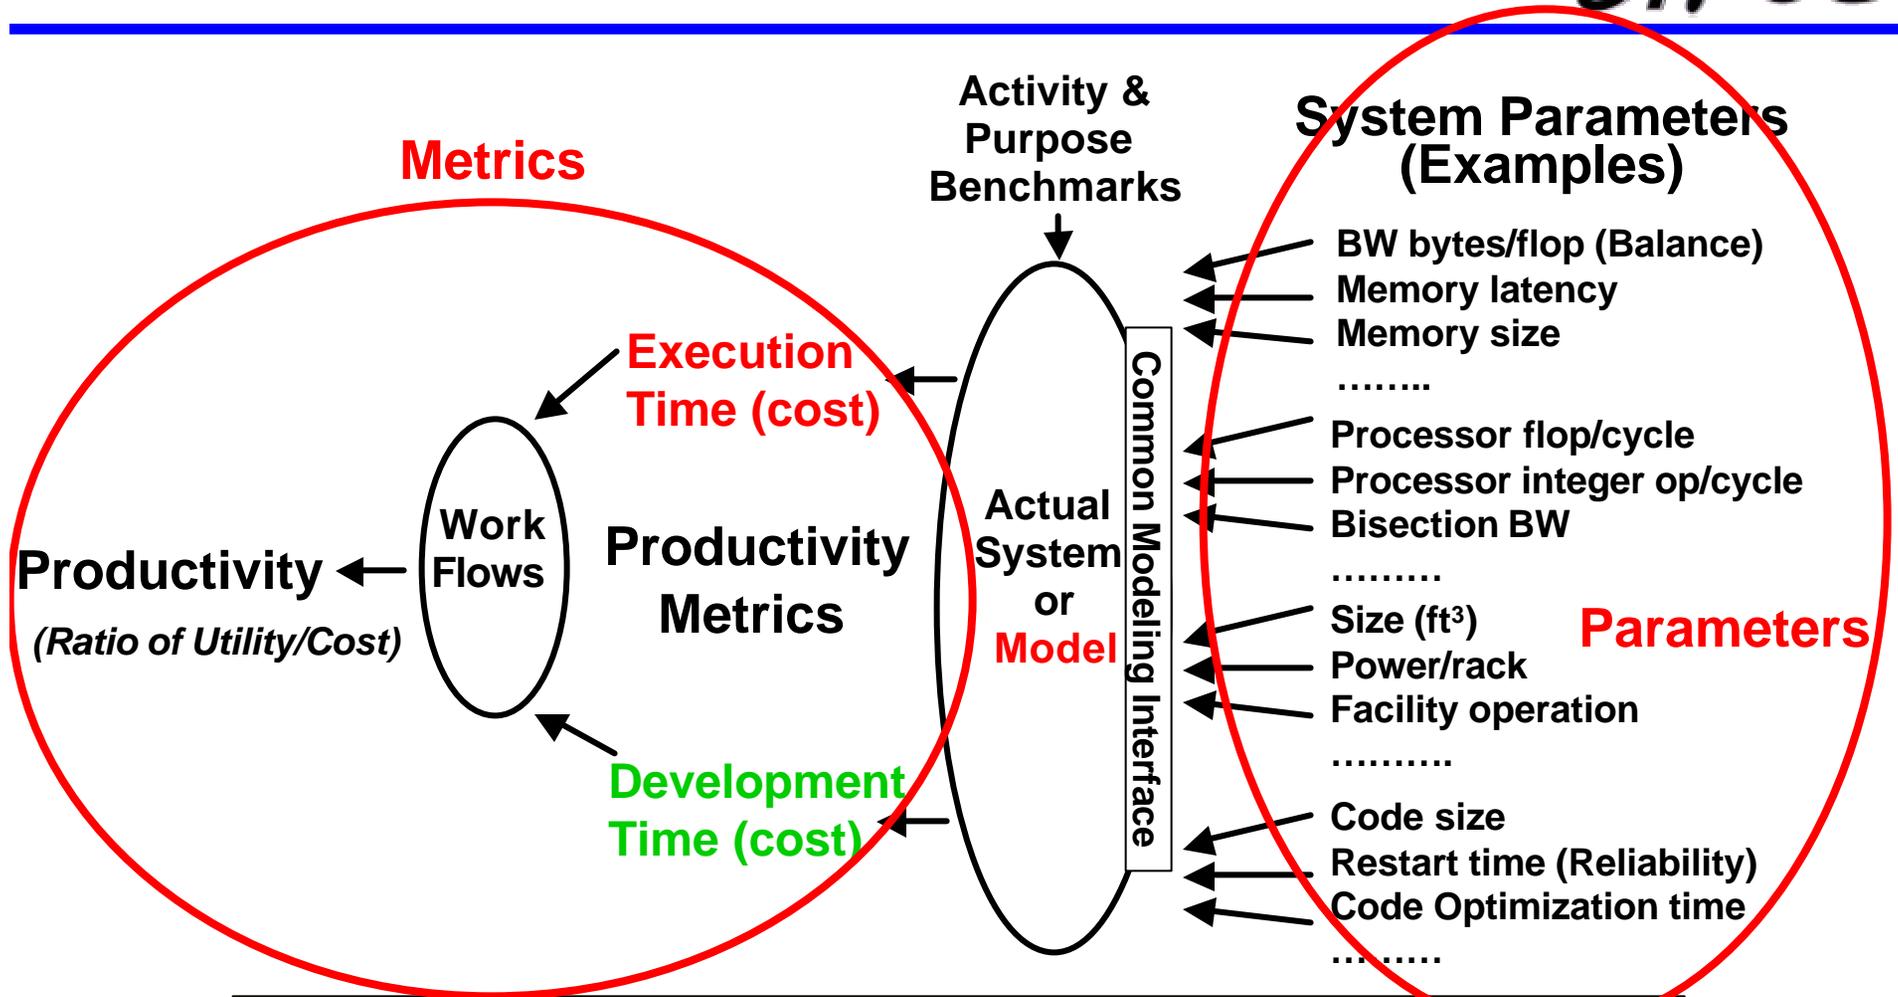
# A Single “Good” Metric? (1 & 2) Digression



- ***Metric*** versus
  - ***Model***
  - ***Parameter***



# HPCS Assessment Framework



**HPCS Productivity Factors:  
Performance, Programmability, Portability, and Robustness**



# A Single “Good” Metric? (1 & 2)



- **Why don't we have a good metric for HPC performance?**
  - We do have good metrics — well sort of...
    - Wall clock time to solution for *your* application
    - Sustained Performance for *your* application
    - Productivity for *your* application
  - Because not really a metric, can impose limitations with comparisons
- **Is there any chance to define a single metric for HPC performance?**
  - No, everybody needs his/her own — applications and workflows vary too greatly for a single performance metric
    - Scientific calculations — sustained Petaflop/s
    - Table Toy — (sustained) GUP/S
    - etc.
  - Yes, we *must* define a single metric
    - The CNN sound bite metric
    - The metric for the Gordon Bell award



# HPCS Phase 1 Kernel and Application Scope Benchmarks



Mission Area	Kernels	Application	Source
Stockpile Stewardship	Random Memory Access	UMT2000	ASCI Purple Benchmarks
	Unstructured Grids		
	Eulerian Hydrocode	SAGE3D	ASCI Purple Benchmarks
	Adaptive Mesh		
	Unstructured Finite Element Model	ALEGRA	Sandia National Labs
	Adaptive Mesh Refinement		
Operational Weather and Ocean Forecasting	Finite Difference Model	NLOM	DoD HPCMP TI-03
Army Future Combat Weapons Systems	Finite Difference Model	CTH	DoD HPCMP TI-03
	Adaptive Mesh Refinement		
Crashworthiness Simulations	Multiphysics Nonlinear Finite Element	LS-DYNA	Available to Vendors

Other Kernels	Kernels	Application	Source
	Lower / Upper Triangular Matrix Decomposition	LINPACK	Available on Web
	Conjugate Gradient Solver		DoD HPCMP TI-03
	QR Decomposition		Paper & Pencil for Kernels
	1D FFT		Paper & Pencil for Kernels
	2D FFT		Paper & Pencil for Kernels
	Table Toy (GUP/s)		Paper & Pencil for Kernels
	Multiple Precision Mathematics		Paper & Pencil for Kernels
	Dynamic Programming		Paper & Pencil for Kernels
	Matrix Transpose [Binary manipulation]		Paper & Pencil for Kernels
	Integer Sort [With large multiword key]		Paper & Pencil for Kernels
	Binary Equation Solution		Paper & Pencil for Kernels
	Graph Extraction (Breadth First) Search		Paper & Pencil for Kernels
	Sort a large set		Paper & Pencil for Kernels
	Construct a relationship graph based on proximity		Paper & Pencil for Kernels
	Various Convolutions		Paper & Pencil for Kernels
	Various Coordinate Transforms		Paper & Pencil for Kernels
	Various Block Data Transfers		Paper & Pencil for Kernels

Bio-Application	Kernels	Application	Source
Quantum and Molecular Mechanics	Macromolecular Dynamics	CHARMM	<a href="http://yuri.harvard.edu/">http://yuri.harvard.edu/</a>
	Energy Minimization		
	MonteCarlo Simulation		
Whole Genome Analysis	Sequence Comparison	Needleman-Wunsch	<a href="http://www.med.nyu.edu/rcr/rcr/course/sim-sw.html">http://www.med.nyu.edu/rcr/rcr/course/sim-sw.html</a>
		BLAST	<a href="http://www.ncbi.nlm.nih.gov/BLAST/">http://www.ncbi.nlm.nih.gov/BLAST/</a>
		FASTA	<a href="http://www.ebi.ac.uk/fasta33/">http://www.ebi.ac.uk/fasta33/</a>
		HMMR	<a href="http://hmmer.wustl.edu/">http://hmmer.wustl.edu/</a>
Systems Biology	Functional Genomics	BioSpice (Arkin, 2001)	<a href="http://genomics.lbl.gov/~aparkin/Group/Codebase.html">http://genomics.lbl.gov/~aparkin/Group/Codebase.html</a>
		Biological Pathway Analysis	
Bio-Application	Kernels	Application	Source
Quantum and Molecular Mechanics	Macromolecular Dynamics	CHARMM	<a href="http://yuri.harvard.edu/">http://yuri.harvard.edu/</a>
	Energy Minimization		
	MonteCarlo Simulation		
Whole Genome Analysis	Sequence Comparison	BLAST	<a href="http://www.ncbi.nlm.nih.gov/BLAST/">http://www.ncbi.nlm.nih.gov/BLAST/</a>
Systems Biology	Functional Genomics	BioSpice (Arkin, 2001)	<a href="http://genomics.lbl.gov/~aparkin/Group/Codebase.html">http://genomics.lbl.gov/~aparkin/Group/Codebase.html</a>
		Biological Pathway Analysis	



# Interrelationships



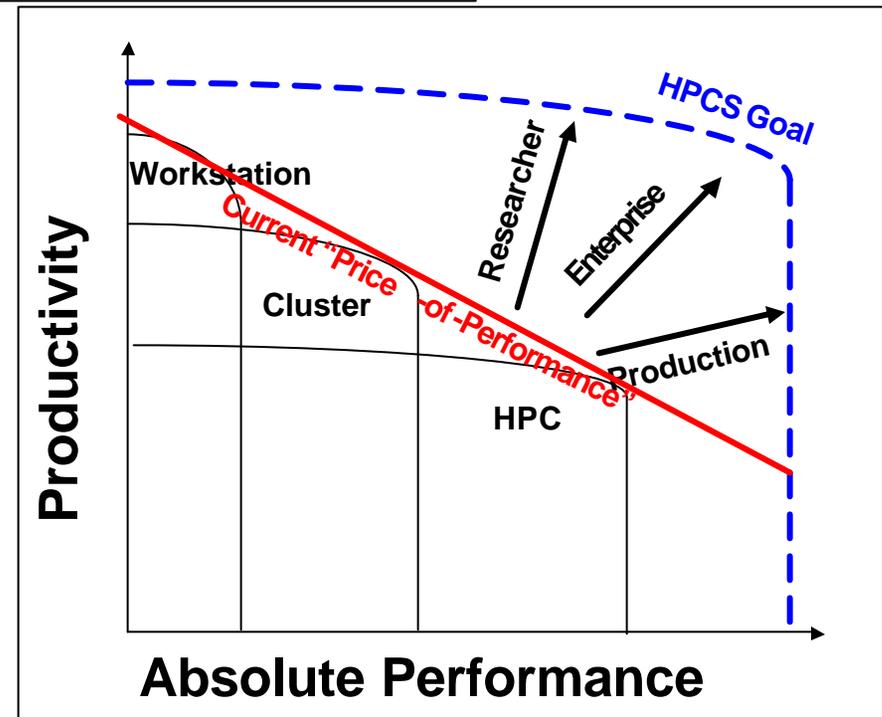
<u>Workflow</u>	<u>Productivity Factors</u>			
	<u>Perf.</u>	<u>Prog.</u>	<u>Port.</u>	<u>Robust.</u>
Lone Researcher		high		
Domain Researcher	high			high
Enterprise	high	high	high	high
Production	high			high

Mission Needs



System Requirements

- Workflows define scope of customer priorities
- Activity and Purpose benchmarks will be used to measure Productivity
- HPCS Goal is to add value to each workflow
  - Increase productivity while increasing performance





## Requirements for Metrics (3)



- **“Demonstrate (not claim) benefits across all mission areas” (Kuck)**
- **“You get what you measure”**
  
- **Theoretical**
  - Stable
  - Quantifiable
  - Reasonable
- **Practical**
  - K.I.S.S.
- **Political**
  - Understandable by someone outside the community with limited time, limited knowledge of HPC, ... , (who controls your funding)



# Accepting New Metrics (4)



- **What needs to be done?**
  - **View as a science**
    - **Perform research and publish!**
    - **Fund work in this area!**
  - **Work toward consensus**
- **Is it possible?**
  - **Most definitely!**



# Sustained or Peak? Repercussions (5)



**↑ Sustained Performance**

- ? Honest
- ? Application dependent

**↓ Peak Performance**

- ? Dishonest
- ? Misleading
- ? Processing in HPC will be *free!*





# Metrics — The Layer 9 Protocol Implications (6)



- National level politics and the media
- Metric(s) need to be understandable
  - By individuals outside the community with limited time, limited knowledge of HPC, ... , (who controls your funding)
  - The New York Times level of in-depth media coverage
  - The CNN Headline News soundbite
- ? Need main stream acceptance and understanding of HPC metrics
- Business
- Metric(s) need to be useful
  - For procurement planning and decision-making by various organizations and business entities



# HPC Performance Metrics Conclusions



- **Grail?**  
In the movie “Indiana Jones and the Last Crusade” the true Grail was a simple wooden cup that stood out among the false golden, jewel encrusted grails — remember, choose wisely!
- **Fata Morgana?**  
At least a superior mirage!