SAND2010-8459P

Some Computational Science Experiences

Richard Barrett

Scalable Architectures Department (1422)

Computer Science Dept Seminar

University of New Mexico

December 3, 2010

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.

Sandia National Laboratories

Iterative methods for solving Ax=b: This was the most fashionable research topic from 1950-1965. Has the point of diminished returns been researched? Here is a quotation from a sophisticated user:

"Progress in Numerical Analysis", Beresford Parlett, SIAM Review, 1978

"Iterative techniques for processing large sparse linear systems were popular in the late 1950's and early 1960's (and their decaying remains still pollute some computational circles). When iterative methods finally departed from the finite element scene in the mid 1960's - having been replaced by direct sparse-matrix methods - the result was a quantum leap in the reliability of linear analysis packages, which contributed significantly to the rapid acceptance of FE analysis at the engineering group level. (This effect, it should be noted, had nothing to do with the relative computational efficiency; in fact iterative methods can run faster on many problems if the user happens to know the optimal acceleration parameters.) Presently, FE analyzers are routinely exercised as black box devices;..."

"Progress in Numerical Analysis", Beresford Parlett, SIAM Review, 1978

Our own view of the situation is different. By their training, the experts in iterative methods *must* collaborate with users. Indeed, the combination of user, numerical analyst, and iterative method can be incredibly effective. Of course, by the same token, inept use can make any iterative method not only slow but prone to failure. Gaussian elimination, in contrast, is a classical black box algorithm demanding no cooperation from the user.

Surely the moral of the story is not that iterative methods are dead, but that too little attention has been paid to the user's current needs?

"Progress in Numerical Analysis", Beresford Parlett, SIAM Review, 1978

Solving Ax = b
Conjugate Gradient Method

```
r = b - A*x;
for iter = 1:max_it                       % begin iteration

  ρ = (r'*r);          dot = r^T_1 r_1+…+r^T_n r_n      MPI_Allreduce ( … )
  if ( iter > 1 ),                        % direction vector
    β = ρ / ρ_1;
    p = r + β*p;       axpy = r_1 + β * p_1+…+r_n + β * p_n
  else
    p = r;
  end
  q = A*p;
  α = ρ / (p'*q );
  x = x + α * p;                          % update approximation vector
  r = r - α *q;                           % compute residual
  error = norm( r ) / bnrm2;              % check convergence
  if ( error <= tol ), break, end
  ρ_1 = ρ ;
end
if ( error > tol ) flag = 1; end
```

$dot = r^T_1 r_1 + \ldots + r^T_n r_n$

$axpy = r_1 + \beta * p_1 + \ldots + r_n + \beta * p_n$

Sandia National Laboratories

Preconditioned Solving $M^{-1}Ax = M^{-1}b$ Conjugate Gradient Method

$M^{-1} \approx A^{-1}$

```
r = b - A*x;
for iter = 1:max_it                        % begin iteration
    z = M / r;                             (If M^-1 A not symmetric.)
    ρ = (r'*z);
    if ( iter > 1 ),                       % direction vector
        β = ρ / ρ_1;
        p = z + β*p;
    else
        p = z;
    end
    q = A*p;
    α = ρ / (p'*q );
    x = x + α * p;                         % update approximation vector
    r = r - α *q;                          % compute residual
    error = f ( A, x, b, N, ε );           % check convergence
    if ( error <= tol ), break, end
    ρ_1 = ρ ;
end
if ( error > tol ) flag = 1; end
```

Sandia National Laboratories

Interaction with “the problem”
q = A * p

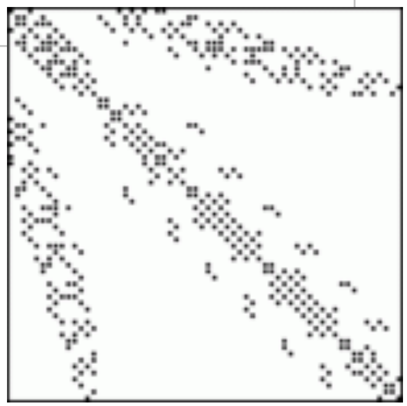

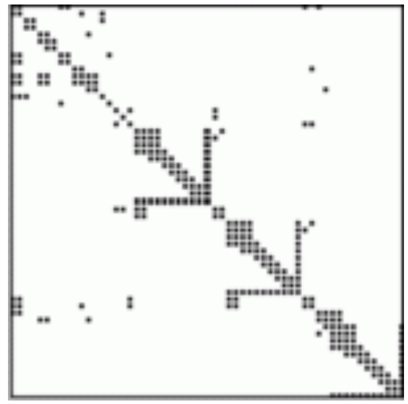

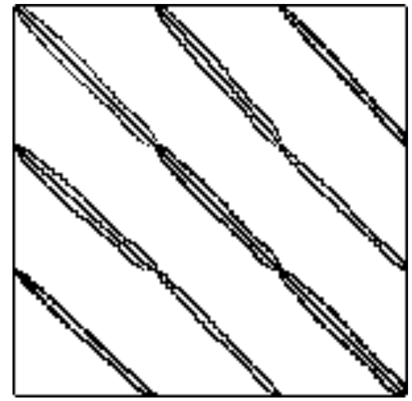

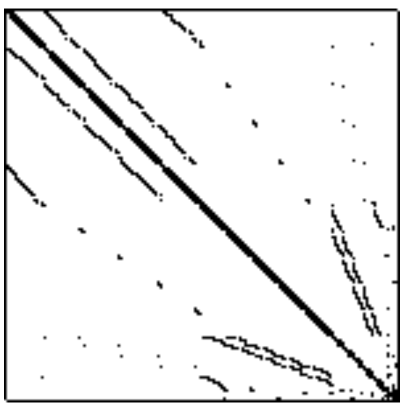

“The Matrix Market”

http://math.nist.gov/MatrixMarket/

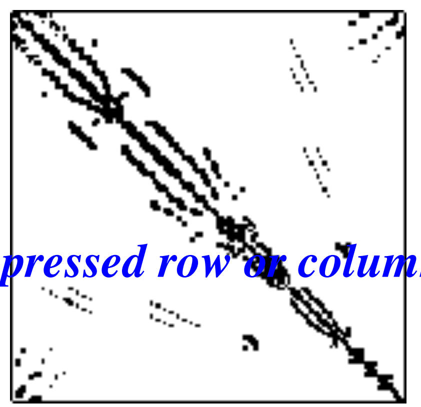

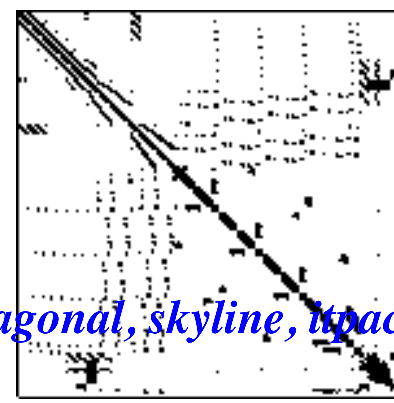

Harwell-Boeing, compressed row or column, blocked, jagged diagonal, skyline, itpack, …

Where's "the matrix"?

Ocean Circulation Model

http://climate.lanl.gov/Models/POP/

2 phases:

1. **Stratification: Baroclinic (3d)**
2. **Surface pressure: Barotropic (2d)**

*A*p: difference stencil "sweep" across mesh:*

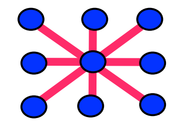

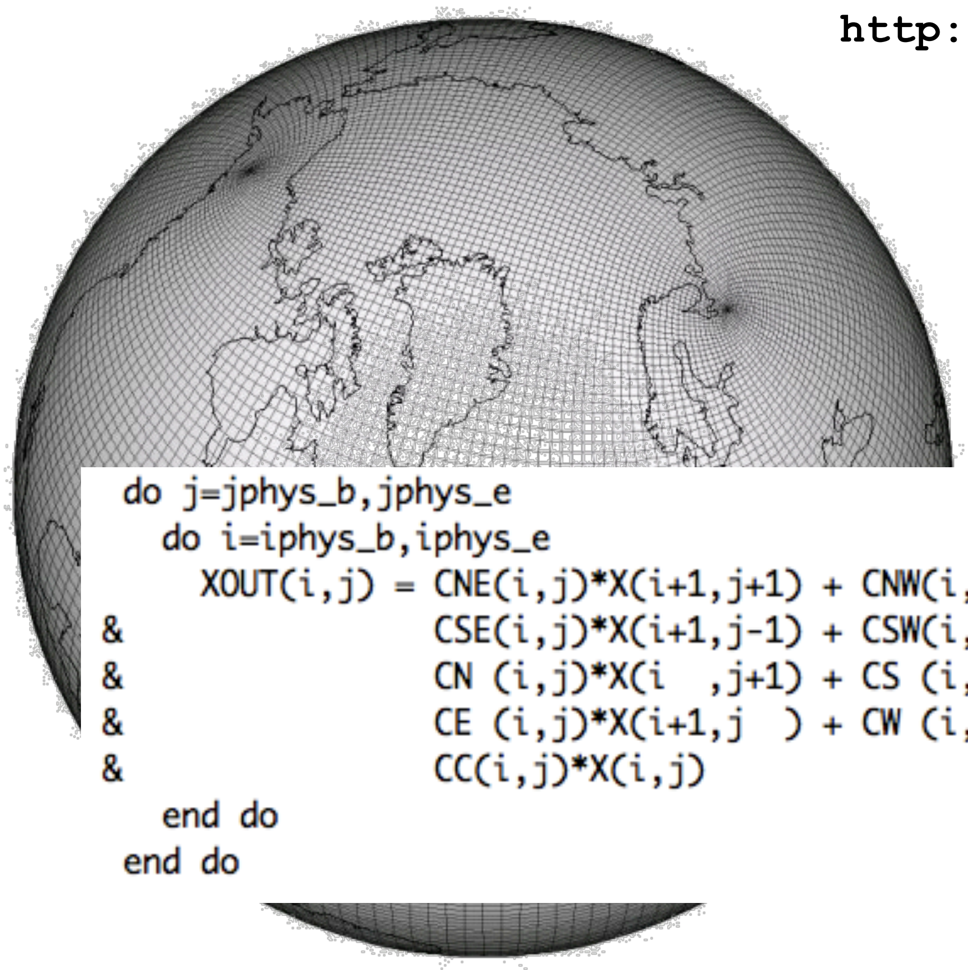

```
 do j=jphys_b,jphys_e
   do i=iphys_b,iphys_e
     XOUT(i,j) = CNE(i,j)*X(i+1,j+1) + CNW(i,j)*X(i-1,j+1) +
&                CSE(i,j)*X(i+1,j-1) + CSW(i,j)*X(i-1,j-1) +
&                CN (i,j)*X(i  ,j+1) + CS (i,j)*X(i  ,j-1) +
&                CE (i,j)*X(i+1,j  ) + CW (i,j)*X(i-1,j  ) +
&                CC(i,j)*X(i,j)
   end do
 end do
```

NATIONAL LEADERSHIP

COMPUTING FACILITY

Preconditioning
$M^{-1}Ax = M^{-1}b$

Want M^{-1} “close to” A^{-1}

i.e. want more favorable system (spectrum)

- Time vs. number of iterations
- If strictly diagonally dominant, M = diag(A)
- A = LU, so M = ILU(d)
- Approximate, Polynomial, Subdomain, …
- for $M=M_1M_2$, $M_1^{-1}AM_2^{-1}(M_2x) = M_1^{-1}bM_2^{-1}$, $M_1 = M_2^T$?
  - Before: $r_0 \leftarrow M_1^{-1}r_0$, after: $x_n \leftarrow M_2^{-1}x_n$

Sandia National Laboratories

Stopping Criteria

Want: *distance of current approximation to exact solution*

$$\mathbf{e}_i = \mathbf{x}_i - \mathbf{x}$$

Have: $\mathbf{x}_i = \mathbf{x}_{i-1} + \alpha * \mathbf{p}_i$ *and* $\mathbf{r}_i = \mathbf{r}_{i-1} - \alpha * \mathbf{q}_i$

Should:

1. identify when the error is small enough to stop,
2. stop if the error is no longer decreasing or is decreasing too slowly,
3. and limit the maximum amount of time spent iterating.

Stopping Criteria

- $\|r_i\|_2 \,/\, \|r_0\|_2$
- $\|r_i\|_2 \,/\, \|b\|_2$

- $\|r_i\|_2 \,/\, \|A\|_\infty$
- $\|r_i\|_\infty / (\|A\|_\infty * \|x\|_1 + \|b\|_\infty)$
- $\|r_i\|_{WRMS}$, *where* $\|\cdot\|_{WRMS} = \sqrt{(1/n)\sum_{i=1}^{n}(r_i/w_i)^2}$
- $\|A x_i - b\|_2 \,/$

A Brief History of Solving Linear Systems of Equations

1959: Conjugate Gradient Method — $A = A^T$ & $x^TAx > 0$ for $x \neq 0$

– Hestenes & Stiefel (& Lanczos)

1976: BiCG, Fletcher — ~~$A = A^T$ & $x^TAx > 0$ for $x \neq 0$~~

1986: GMRes, Saad & Schultz

1989: CGS, Sonneveld.

1991: QMR, Freund & Nachtigal

1992: BiCGStab, van der Vorst

(most cited paper of 1990s (SIAM))

1993: TFQMR, Freund

Luck = Preparation + Opportunity

- 1989: PVM 1.0, Sunderum@Emory, Geist@ORNL
- 1991: PVM 2.0, ORNL + UTK
- 1992: PVM 3.0, ORNL + UTK
- 1994: MPI 1.0

- 1995: T3D@LANL

- 1995: ASCI program@LANL, LLNL, Sandia

Algorithmic Bombardment…

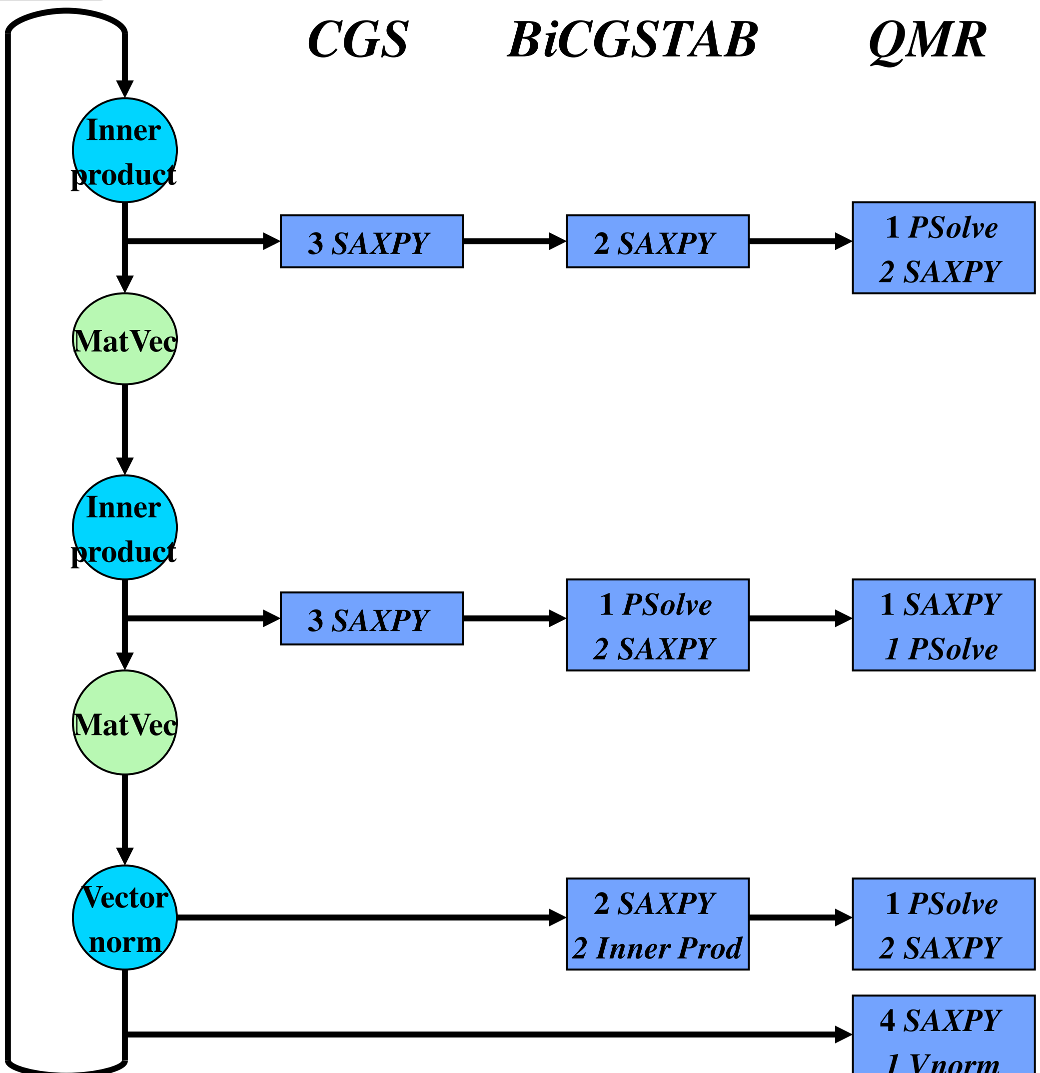

And now you're in luck!
It's the power…

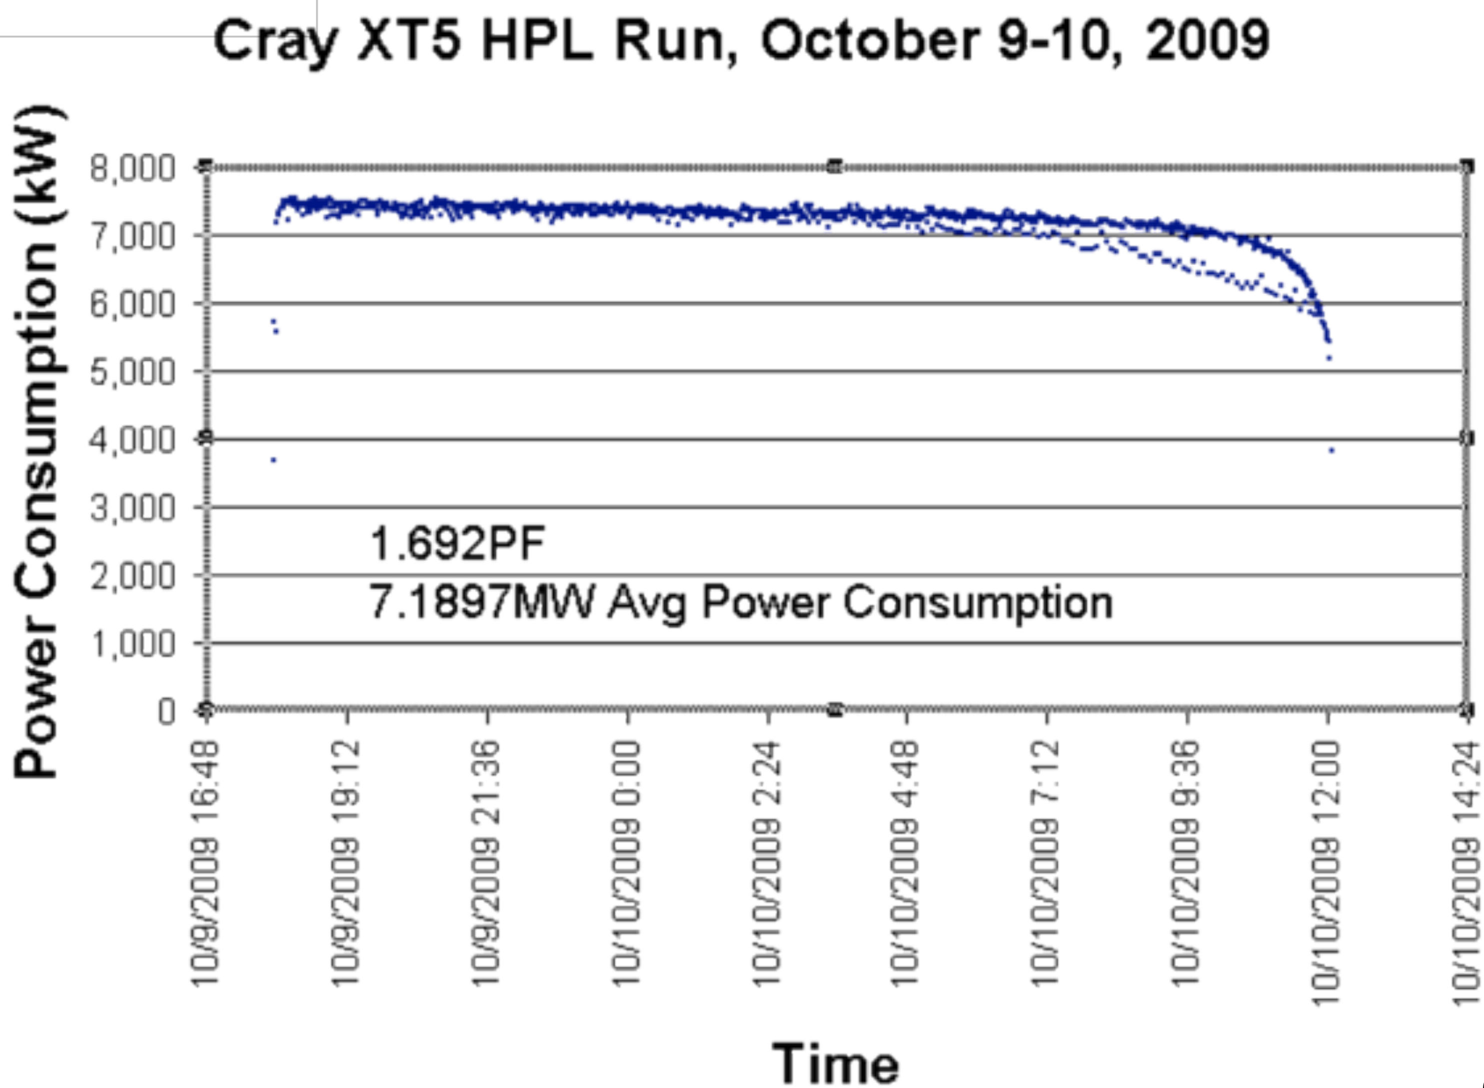

Sandia National Laboratories

AMG2006*

Platform: Jaguar

Architecture: XT4

CPU: AMD Quad

P-states (Frequency States)

- **P0: 2.1 GHz , 1.25V**
- **P1: 2.1 GHz , 1.25V**
- **P2: 1.7 GHz, 1.1625V**
- **P3: 1.4 GHz, 1.125V**
- **P4: 1.1 GHz, 1.1V**

Nodes: 6144

Runtime Increase: 3.2%

Energy Decrease (Savings): 30.6%

Order of magnitude energy savings vs. performance impact!

Two application runs, same physical nodes, statically altering CPU frequency (P-state) allows lowering input voltage to chip resulting in larger energy savings.

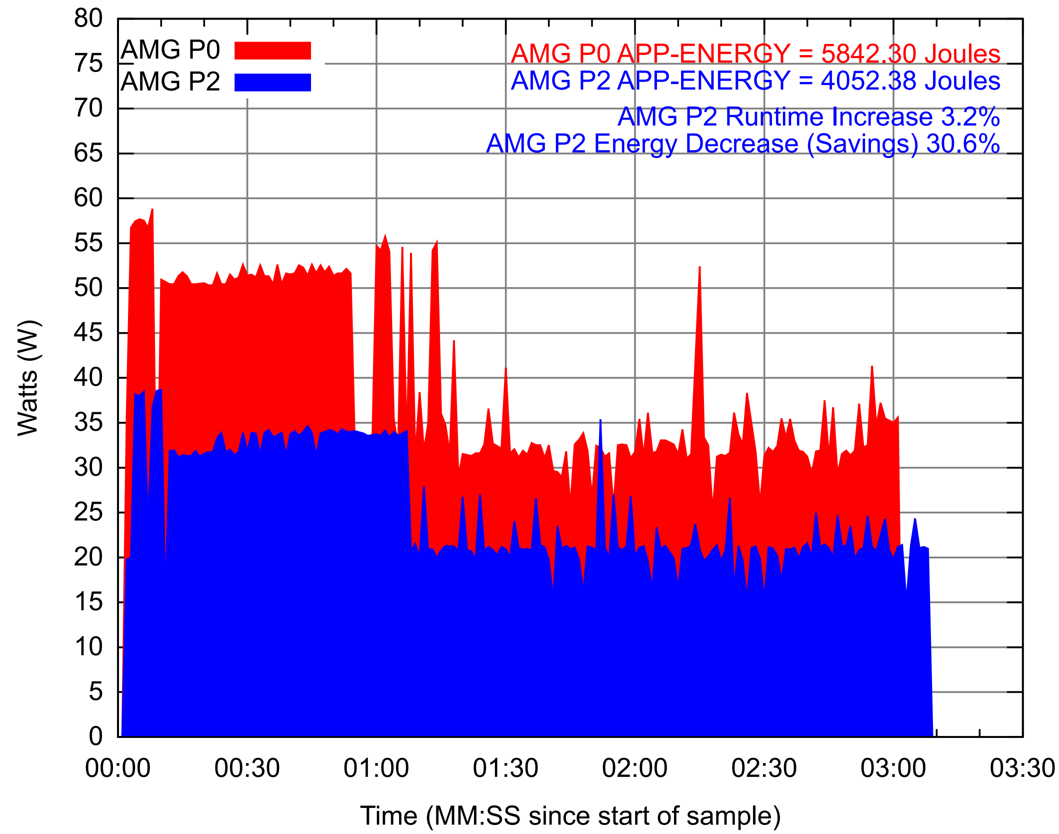

Single node capture of watts over time for each run of AMG2006, varying P-states

** Work of Jim Laros@Sandia*

Sandia National Laboratories

LAMMPS*

Platform: Jaguar
Architecture: XT4
CPU: AMD Quad
P-states (Frequency States)

- **P0: 2.1 GHz , 1.25V**
- **P1: 2.1 GHz , 1.25V**
- **P2: 1.7 GHz, 1.1625V**
- **P3: 1.4 GHz, 1.125V**
- **P4: 1.1 GHz, 1.1V**

Nodes: 4096
Runtime Increase: 16.1%
Energy Decrease (Savings): 21.8%

<u>Compute intensive application, still observe significant energy savings. Illustrates which applications can expect most benefit.</u>

Two application runs, same physical nodes, statically altering CPU frequency (P-state) allows lowering input voltage to chip resulting in larger energy savings.

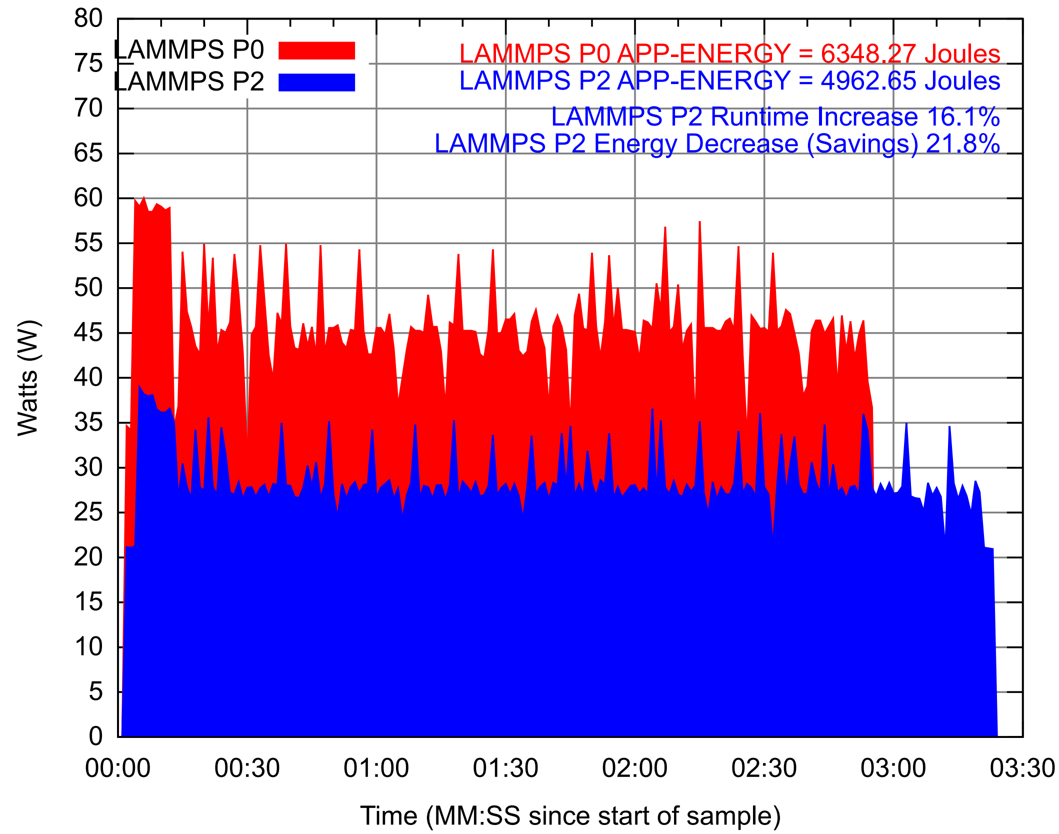

Single node capture of watts over time for each run of LAMMPS, varying P-states

** Work of Jim Laros@Sandia*

Sandia National Laboratories

Self-consistent full-wave and Fokker-Planck calculations for ion cyclotron heating in non-Maxwellian plasmas

Here Maxwell's eqns reduces to Helmholtz wave eqn

$$-\nabla \times \nabla \times \mathbf{E} + \frac{\omega^2}{c^2}\left(\mathbf{E} + \frac{i}{\omega\varepsilon_0}\mathbf{J}_p\right) = -i\omega\mu_0\mathbf{J}_{ant}$$

where

$$\mathbf{J}_p(\mathbf{r},t) = \int_{-\infty}^{t} dt' \sum_s \int d\mathbf{r}' \sigma\left(f_s^0(E), \mathbf{r}, \mathbf{r}', t, t'\right) \cdot \mathbf{E}(\mathbf{r}', t')$$

is a non-local integral operator on the wave electric field.

Which to me comes down to

$$\mathrm{Ax} = \mathrm{b},$$

for $\mathrm{A}^{nxn}, \mathrm{x}^{nx1}, \mathrm{b}^{nx1}$ *in* C.

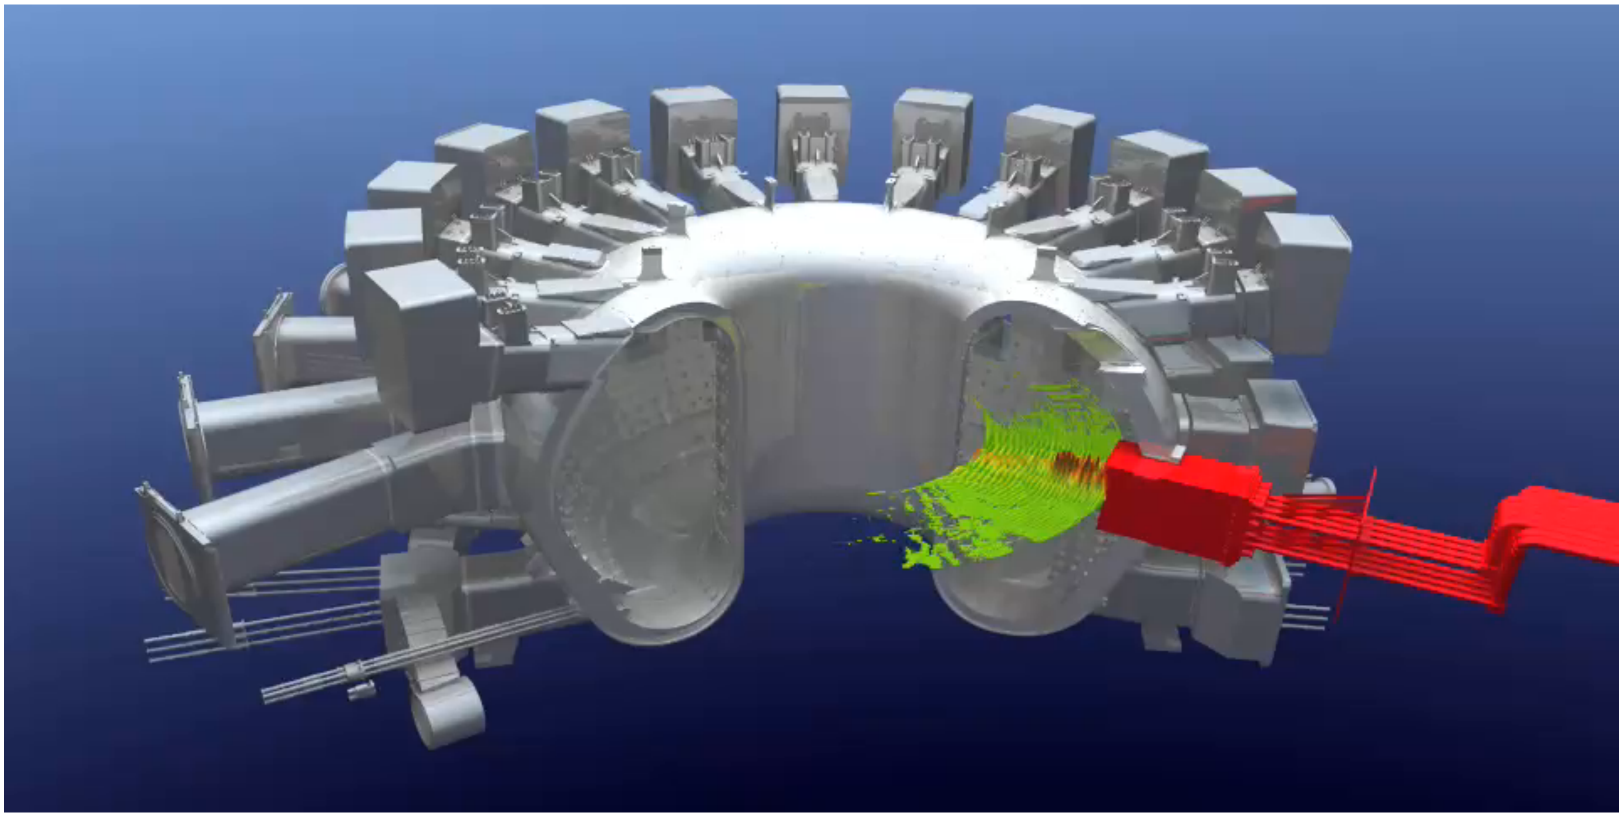

AORSA simulation; movie by Sean Ahern@ORNL

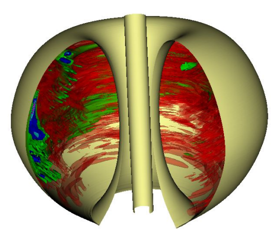

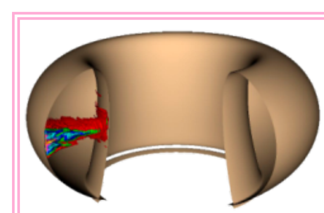

Sandia National Laboratories

- *“Simulation of High Power Electromagnetic Wave Heating in the ITER Burning Plasma”*, E. F. Jaeger, L. A. Berry, E. F. D'Azevedo, R. F. Barrett, S. D. Ahern, D. W. Swain, D. B. Batchelor, R. W. Harvey, J. R. Myra, D. A. D'Ippolito, C. K. Phillips, E. Valeo, D. N. Smithe, P. T. Bonoli, J. C. Wright, and M. Choi, Physics of Plasmas, (15)7, 2008

- *“Global-wave Solutions with Self-Consistent Velocity Distributions in Ion Cyclotron Heated Plasmas”*, E.F. Jaeger, R.W. Harvey, L.A. Berry, J.R. Myra, R.J. Dumont, C.K. Phillips, D.N. Smithe, R.F. Barrett, D.B. Batchelor, P.T. Bonoli, M.D. Carter, E.F. D'Azevedo, D.A. D'ippolito, R.D. Moore and J.C. Wright, Journal of Nuclear Fusion, Volume 46, Number 7, July, 2006.

- *Self-Consistent Full-Wave and Fokker-Planck Calculations for Ion Cyclotron Heating in Non-Maxwellian Plasmas*, E.F. Jaeger, L.A. Berry, S.D. Ahern, R.F. Barrett, D.B. Batchelor, M.D. Carter, E.F. D'Azevedo, R.D. Moore, R.W. Harvey, J.R. Myra, D.A. D'Ippolito, R.J. Dumont, C.K. Phillips, H. Okuda, D.N. Smithe, P.T. Bonoli, J.C. Wright, and M.Choi, Physics of Plasmas, 13, May 2006.

- *Complex version of High Performance Computing LINPACK Benchmark (HPL)*, R. Barrett, T. Chan, E.F. D'Azevedo, E.F. Jaeger, K. Wong, and R. Wong. Concurrency and Computation: Practice and Experience, (22)5, April, 2010 .
- *Performance Analysis and Projections for Petascale Applications on Cray XT Series Systems*, S. R. Alam, R.F. Barrett, J. A. Kuehn, and S. W. Poole, Workshop on Large-Scale Parallel Processing,IEEE International Parallel & Distributed Processing Symposium (IPDPS), Rome, Italy, 2009.
- *Exploring HPCS Languages in Scientific Computing*, R.F. Barrett, S.R. Alam, V. de Almeida, D.E. Bernholdt, W.R. Elwasif, J.A. Kuehn, S.W. Poole, and A.G. Shet, Scientific Discovery Through Advanced Computing (SciDAC 2008), Journal of Physics: Conference Series 125 012034, 2008.
- *An Evaluation of the ORNL Cray XT3*, S.R. Alam, R.F. Barrett, M. R. Fahey, J. A. Kuehn, E.O.B. Messer, R. T. Mills, P.C. Roth, J. S. Vetter, and P. H. Worley; International Journal of High Performance Computing Applications. (22)1, February 2008.
- *Performance Characterization of a Hierarchical MPI Implementation on Large-scale Distributed-memory Platforms*, S.R. Alam, R.F. Barrett, J.A. Kuehn, and S.W. Poole, The 38th International Conference on Parallel Processing (ICPP-2009), Vienna, 2009.
- *Early Evaluation of IBM Blue Gene/P*, S. Alam, R. Barrett, M. Bast, M. Fahey, J. Kuehn, C. McCurdy, J. Rogers, P. Roth, R. Sankaran, J.S. Vetter, P. Worley, W. Yu, Proceedings of the ACM/IEEE Conference on High Performance Networking and Computing (SC08), Austin, TX, 2008.

Co-Design

Application
Methods
Algorithms
Basic Computations
Programming Model
Programming Mechanisms
Libraries
Compiler
Runtime System
Operating System
System Architecture
Node: Processors, Memory
Inter-node: Interconnect
Enabling Technologies
Fundamental Technology Capabilities

Co-Design

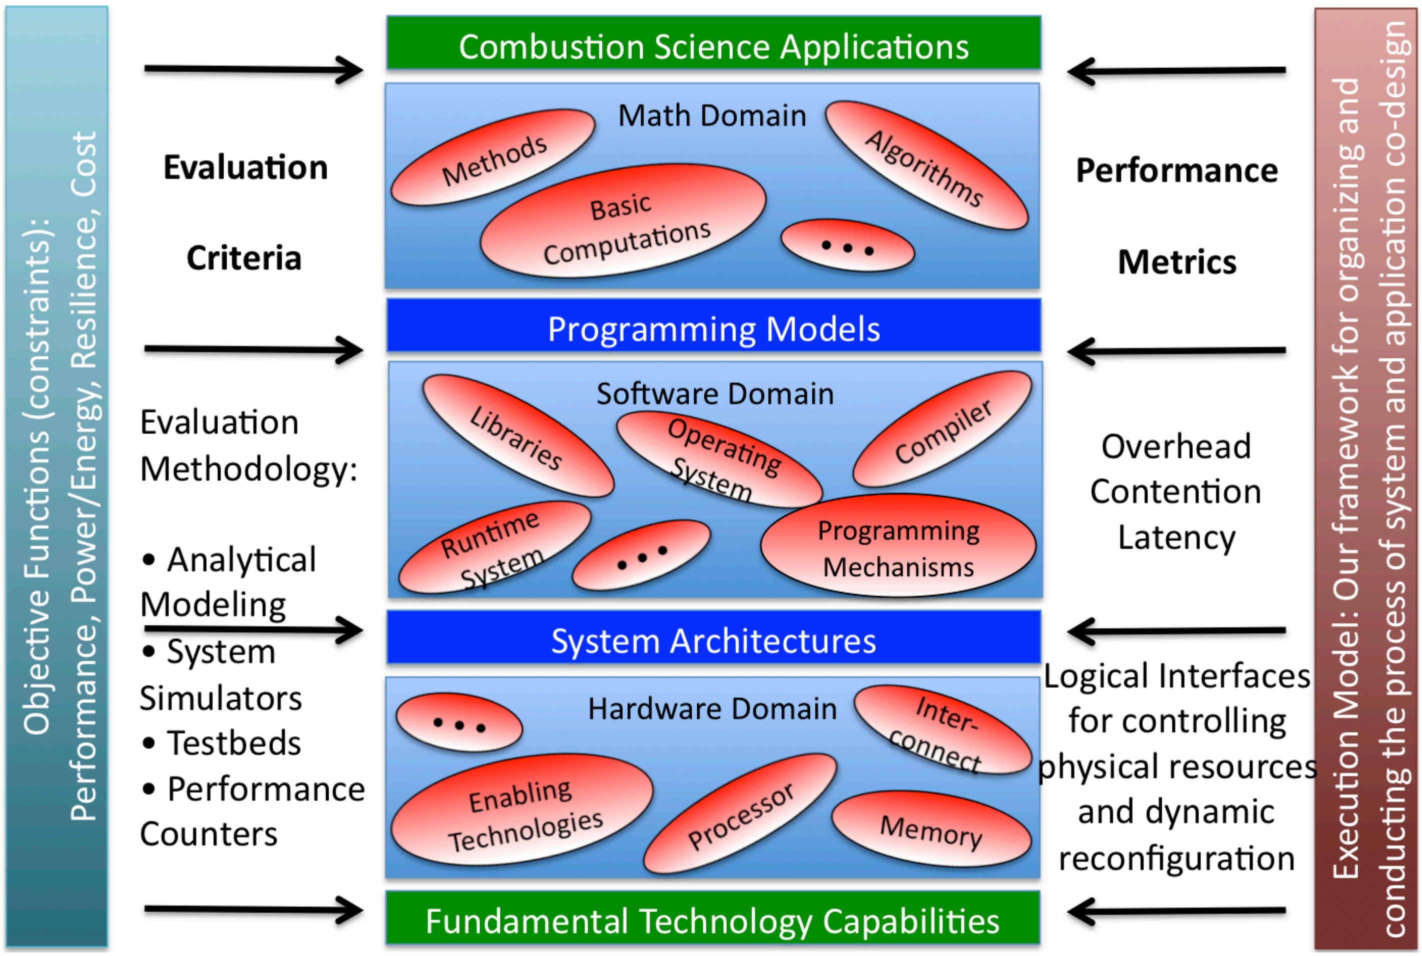
Objective Functions (constraints):
Performance, Power/Energy, Resilience, Cost
Evaluation
Criteria
Evaluation
Methodology:
• Analytical
Modeling
• System
Simulators
• Testbeds
• Performance
Counters
Combustion Science Applications
Math Domain
Methods
Basic
Computations
Algorithms
• • •
Programming Models
Software Domain
Libraries
Operating
System
Compiler
Runtime
System
• • •
Programming
Mechanisms
System Architectures
Hardware Domain
• • •
Inter-
connect
Enabling
Technologies
Processor
Memory
Fundamental Technology Capabilities
Performance
Metrics
Overhead
Contention
Latency
Logical Interfaces
for controlling
physical resources
and dynamic
reconfiguration
Execution Model: Our framework for organizing and
conducting the process of system and application co-design

DOE Exascale Co-design Centers

Exascale Co-Design Consortium (ECDC):

- Exascale Co-Design Center for Materials in Extreme Environments
- Co-design for Exascale Research in Fusion (CERF).
- Chemistry Exascale Co-design Center (CECC)
- High Energy Density Physics
- Center for Exascale Simulation of Advanced Reactors

Co-Design

Performance:

- **HPC: Best case**
- **Embedded systems: Worst case**

Thanks