Optimization of M3D for the Cray X1 and Neoclassical Closures

Don Spong , Steve Hirshman, Diego Castillo-Negrete, Ed D'Azevedo, Mark Fahey, Richard Mills

Oak Ridge National Laboratory

Introduction

- ORNL LDRD awarded in October under the Terascale Computing and Simulation Science Initiative for "Terascale Computations of Multiscale Magnetohydrodynamics"
  - Team members: Don Spong (PI), Ed D'Azevedo (Co-PI), Steve Hirshman, Diego Castillo-Negrete, D. Batchelor, M. Fahey, R. Mills, Steve Jardin, W. Park, G. Y. Fu

Introduction

- ORNL LDRD awarded in October under the Terascale Computing and Simulation Science Initiative for "Terascale Computations of Multiscale Magnetohydrodynamics"
  - Team members: Don Spong (PI), Ed D'Azevedo (Co-PI), Steve Hirshman, Diego Castillo-Negrete, D. Batchelor, M. Fahey, R. Mills, Steve Jardin, W. Park, G. Y. Fu
- Motivated by ORNL's selection for the National Leadership Computing Facility (NLCF)

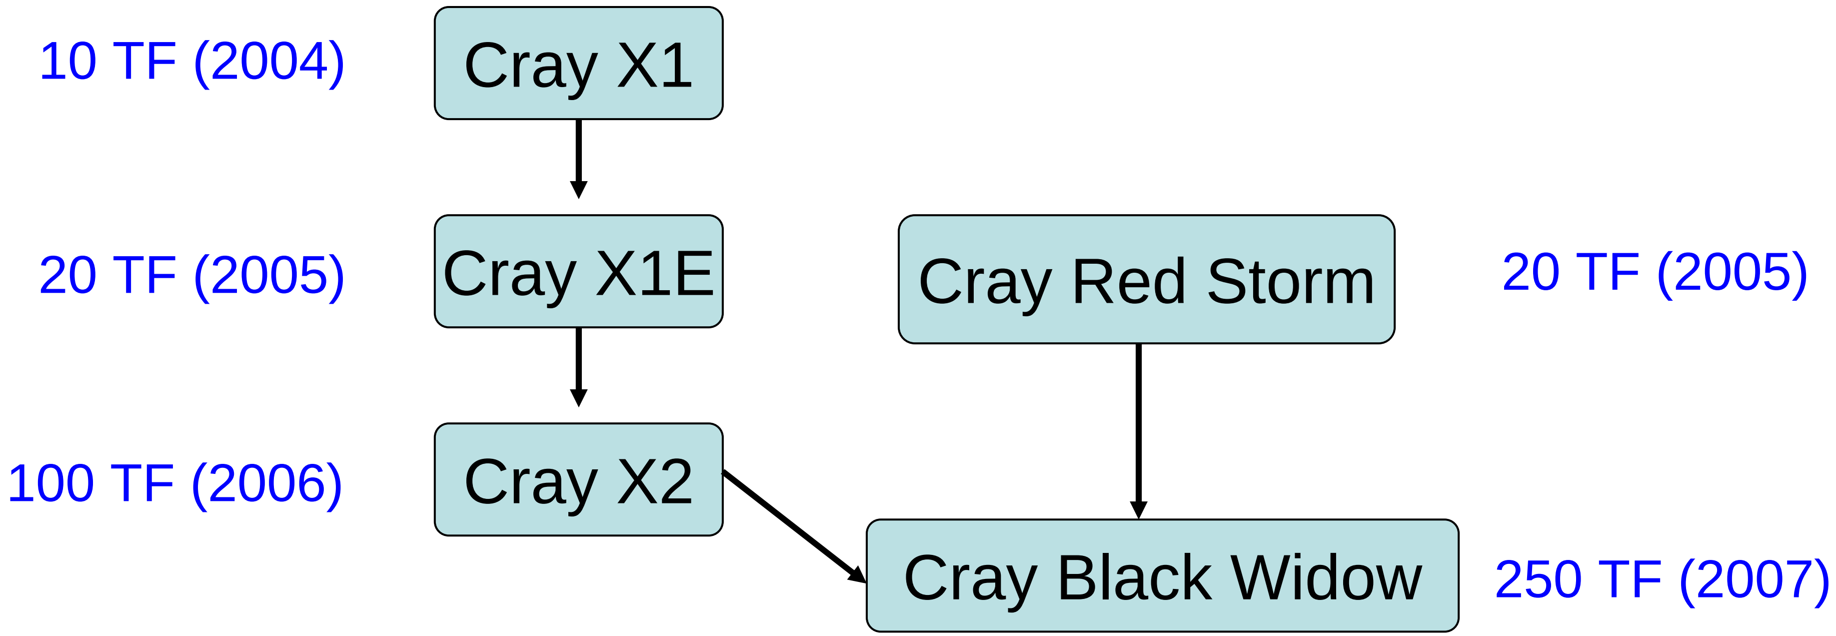

Our goals/motivations

- Adapt/optimize M3D for Cray systems
  - Opportunity to be involved again in MHD physics
  - Support future stellarator experiment (QPS)
  - ITER physics
- Adapt/optimize stellarator Monte Carlo code (DELTA5D) to Cray systems
  - Self-consistent non-local transport studies
  - Incorporate viscosity-base methods
- Develop improved particle-based closure relations for study of NTM's
- Programmatic: prepare for future involvement in:
  - Fusion scientific end station
  - Fusion Simulation Project
  - Multiscale Mathematics Initiative

M3D optimization issue for vector machines: sparse matrix solver

- Development of efficient sparse matrix solvers important to Cray X1, Red Storm
  - Common issue for PDE's based on finite elements/differences
  - Addressed by specialized storage schemes
    - E.g., jagged diagonal scheme successfully used on the Earth Simulator
    - Methods we have tested (code from SPARSKIT2 (by Yosef Saad, Univ. Minnesota - performs matrix multiply 1000 times)
    - Compressed sparse matrix row storage
    - ELLPACK
    - Diagonal
    - Jagged diagonal

M3D sparse matrix structure

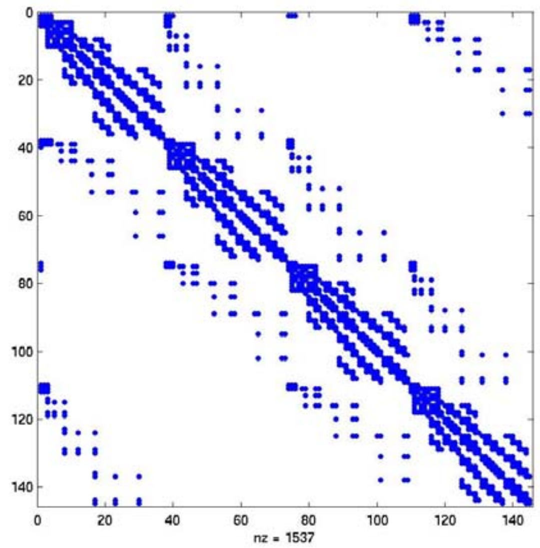

parallel scaling on NERSC Seaborg

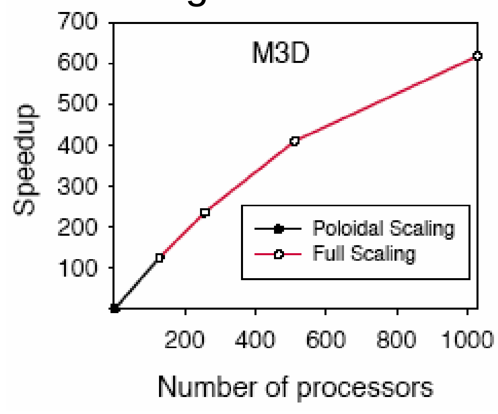

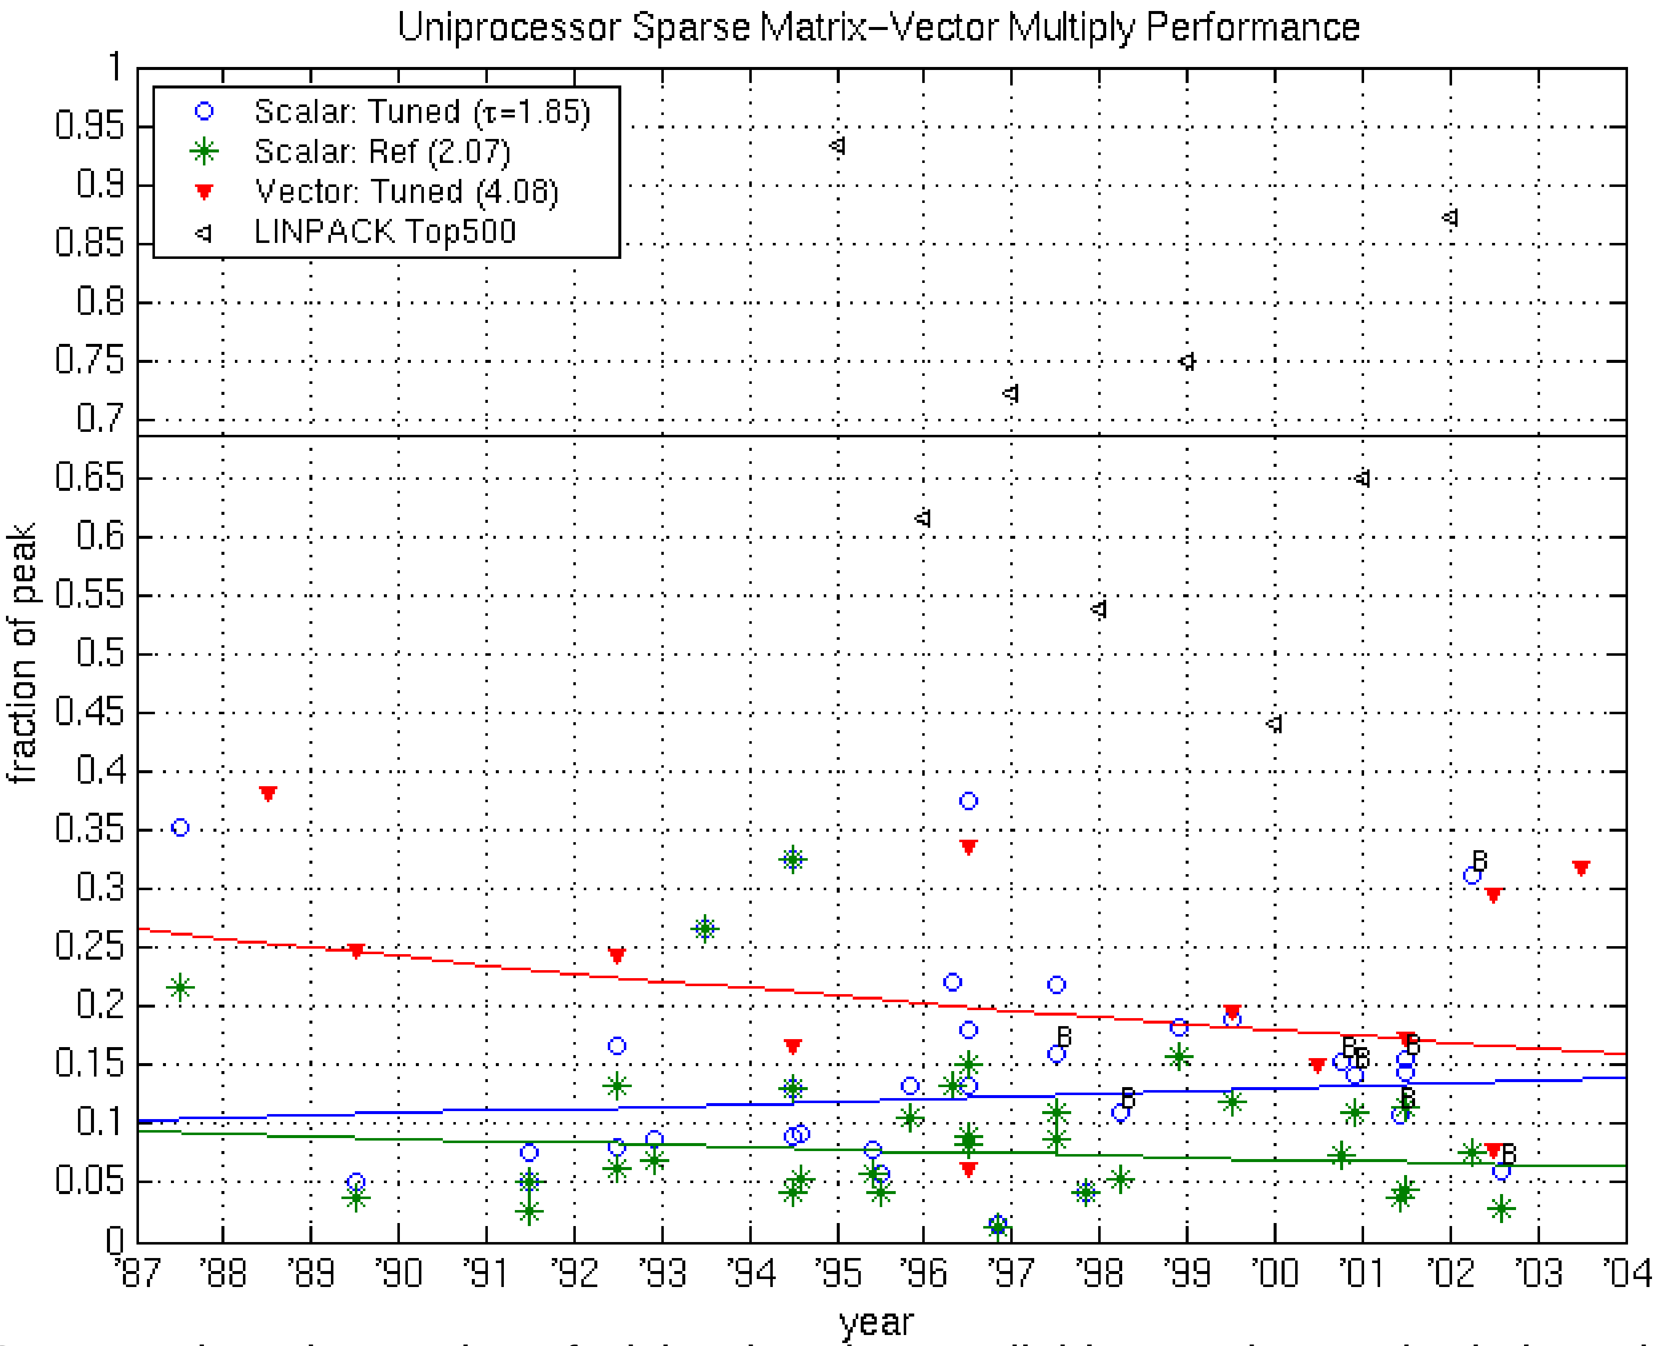

Source: PhD Dissertation of Richard Vuduc, available at Bebop.cs.berkeley.edu.

ELLPACK Format

- Assume each row has similar number of nonzeros
- Store and operate on extra zeros
- Do j=1,ncol
  - Do i=1,n
    - y(i) = y(i) + a(i,j) * x(ja(i,j))
- Regular computation is good for vectorization
- Bad if a few rows have many nonzeros

Jagged Diagonal Format

- Pre-sort rows by number of nonzeros (require permutation in result y(:) vector)
- Long vectors good for vectorization
- Do ii=1,jdiag
  - ! compute length and offset k1
  - Do j=1,len
    - y(j) = y(j) + a(k1+j) * x(ja(k1+j))
- Ellpack like computation for multiple diagonals with same length

7pt Stencil 50x50x50

- N=110592, nnz=760320
- Minnz=6, maxnz=7
- Ellpack 86% filled
- 7 jagged diagonals
- 7 diagonals

	IBM Power4	SGI Altix	Cray X1
CSR	4.08	16.74	21.90
JAG	5.04	13.06	1.60
Ellpack	5.58	13.59	1.67
Diag	4.51	8.69	0.97
Mflops	372	175	1569

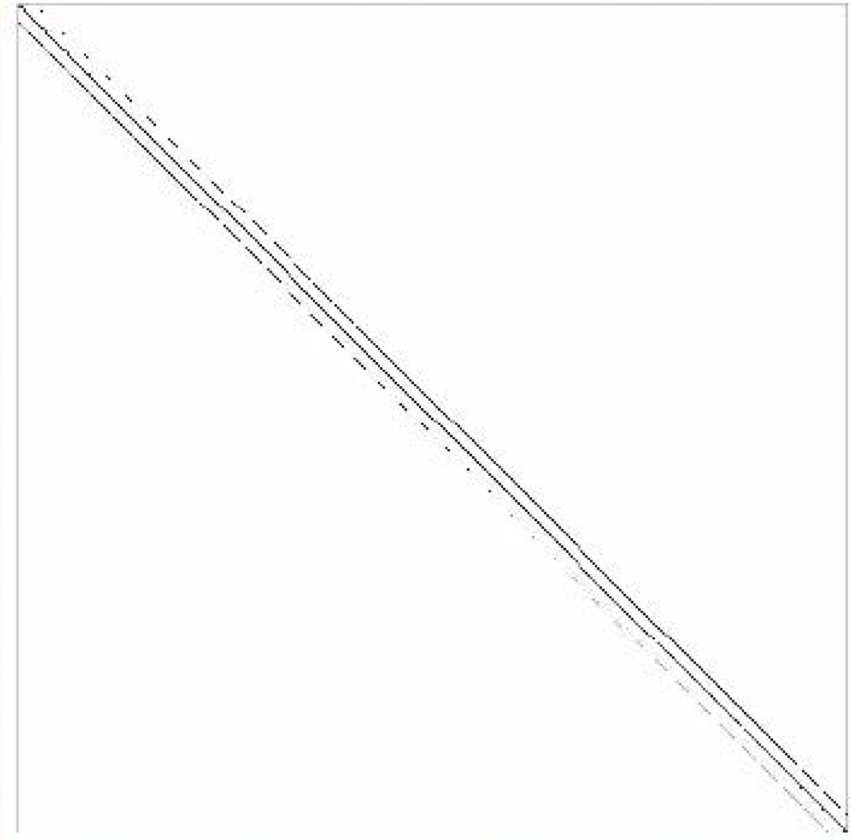

BCSSTK18 (Structural Engineering)

- **N=11948, nnz=80519**
- **Minnz=1, maxnz=31**
- **Ellpack 21% filled**
- **31 jagged diagonals**
- **1243 diagonals**

	IBM Power4	SGI Altix	Cray X1
CSR	0.37	0.55	1.75
JAG	0.41	0.28	0.13
Ellpack	1.54	2.58	0.63
Diag	57.39	211.58	20.42
Mflops	435	575	1239

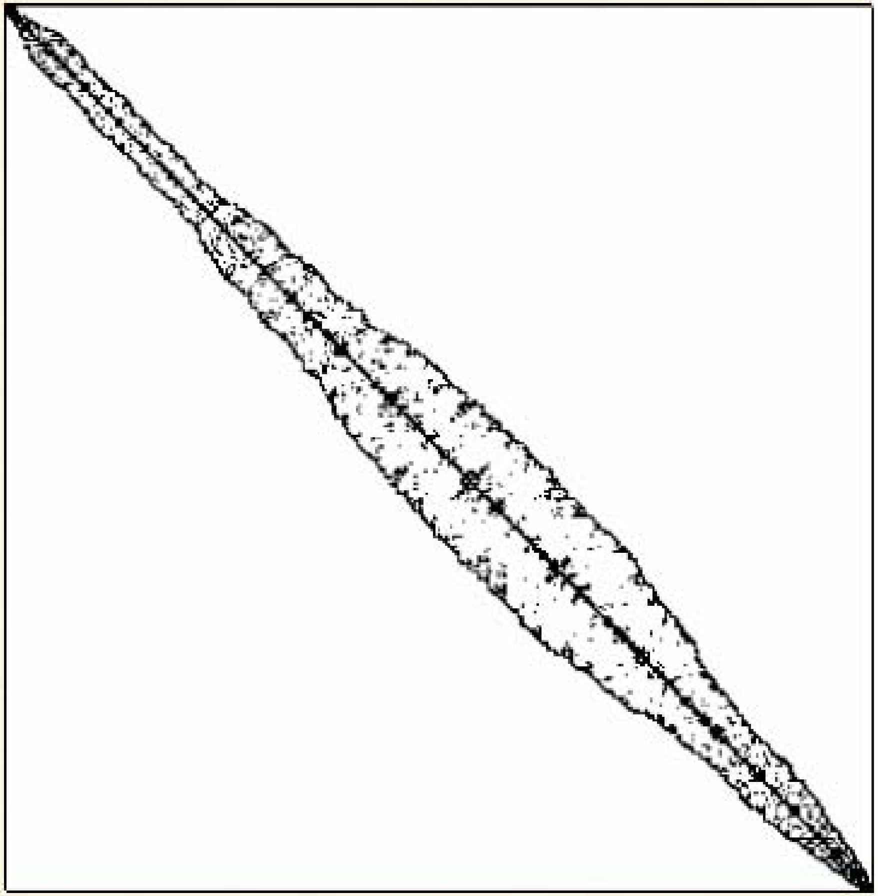

Particle-based closure relations

- We are interested both in utilizing the existing M3D hybrid method and using techniques from DELTA5D, as appropriate
- Computational issues
  - Global memory access (co-array Fortran or other methods)
  - Parallelization over particles or over particles in volumetric regions
  - Implicit stepping methods for particles - averaging over fast time scale bounce and transit motions
    - NTM: Ohm's law -> electrons
- Physics issues
  - Viscosity based calculation of bootstrap current
  - *l* = 2 Legendre harmonic moment
    - Not as sensitive to neglect of field-on-particle collisions, momentum restoring terms in collision operator, avoids large canceling terms of *l* = 1 moments
  - Such methods have recently been adapted/applied to 3D equilibria (H. Sugama, S. Nishimura, POP-2002, D. Spong Tues. morning invited talk)

Current Status

- DELTA5D ported to Cray X1 - factor of 10 faster than similar clock speed IBM-SP scalar systems (Seaborg)
  - Further optimization should be possible
- M3D partially ported
  - Runs under some domain decomposition options, but not others
- Potential performance increase ~100 in going from existing IBM-SP to Cray 100 TF system
  - TFlops: ~5 to 100
  - Efficiency: ~6% to 35%