

# Scalable Supercomputers Solving Superconductivity

James B. White III (Trey)  
whitejbiii@ornl.gov

# Acknowledgement

Research sponsored by the  
Mathematical, Information, and  
Computational Sciences Division, Office  
of Advanced Scientific Computing  
Research, U.S. Department of Energy,  
under Contract No. DE-AC05-  
00OR22725 with UT-Battelle, LLC.

# The Science

- Simulating high-temperature “cuprate” superconductors
- New algorithm by Mark Jarrell *et al.*
  - Quantum Monte Carlo
  - Dynamic Cluster Approximation
- Now trying to reproduce experiments
- Goal of predicting new materials

# The Application

- Collaboration of U.Cincinnati and ORNL
  - Mark Jarrell and Jim Hague of UC
  - Thomas Maier and Thomas Schulthess of ORNL
- Previous HPC systems provided capability for only small 2D clusters (2-4 atoms)
  - Did not reproduce proper physics
- Need to look at larger 2D clusters (32-64 atoms)
- Then multiple 2D clusters in 3D layers
- Then *many* runs to solve inverse problem of predicting new materials

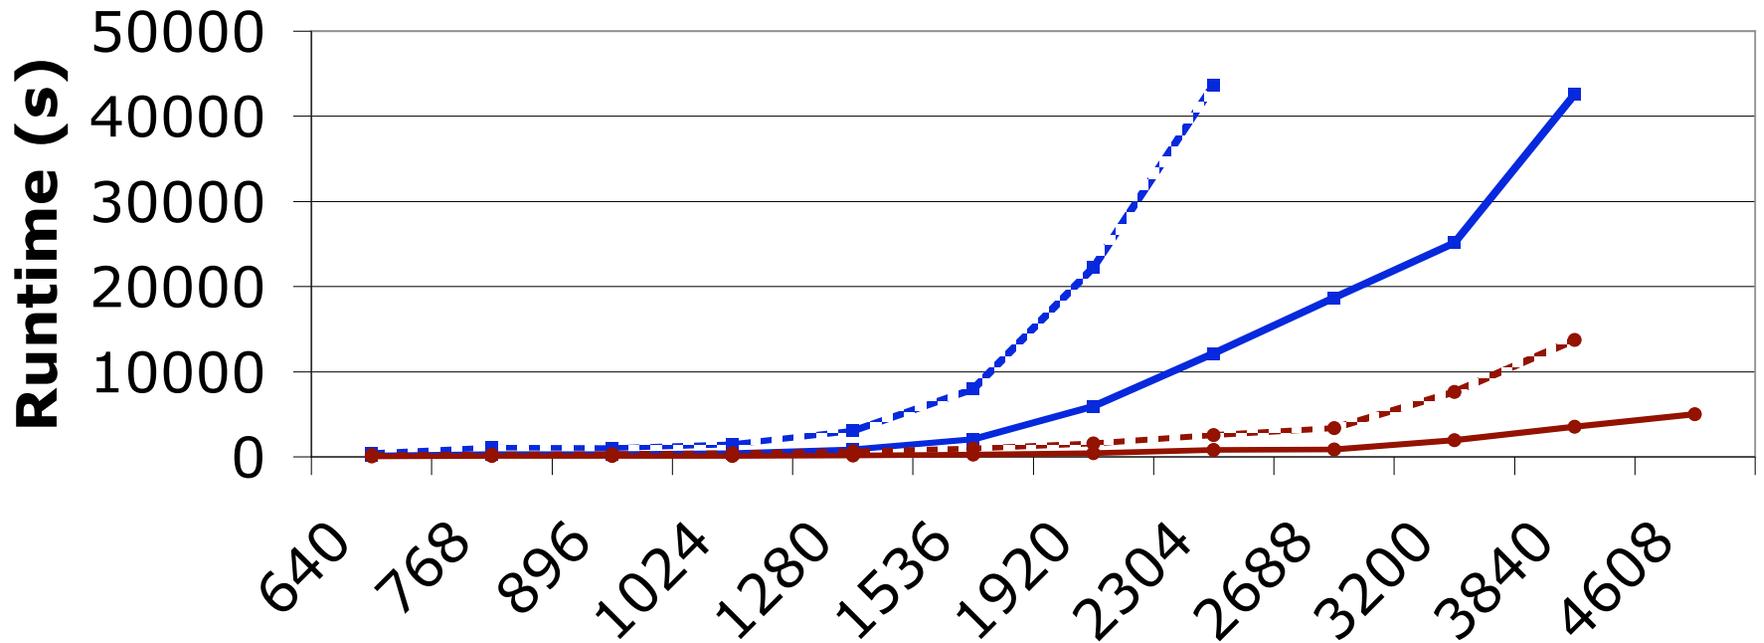
# Scaling

- Quantum Monte Carlo
  - Highly parallel, easy to scale, right?
  - No, each process has a significant fixed startup
  - Favors fewer, faster processors
- Dominated by  $N^3$  operations
  - CGEMM - level 3 BLAS, easy on memory
  - DGER - level 2 BLAS, needs memory bandwidth
  - $O(1)$  CGEMMs and  $O(N)$  DGERs per step
  - DGERs dominate

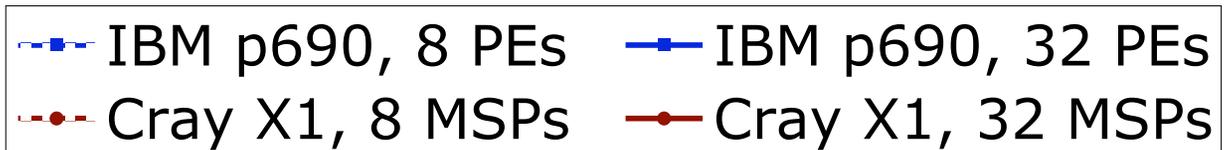
# Performance Experiment

- Production simulation
- $N = \text{cluster size} * \text{time slices}$
- Cluster size of 64
- Series of runs
  - Increasing numbers of time slices (10-70)
  - Decreasing numbers of Monte-Carlo samples
  - Runs use different parameters but generally get more expensive
  - Lines connecting points are for clarity

# DCA-QMC Runtime



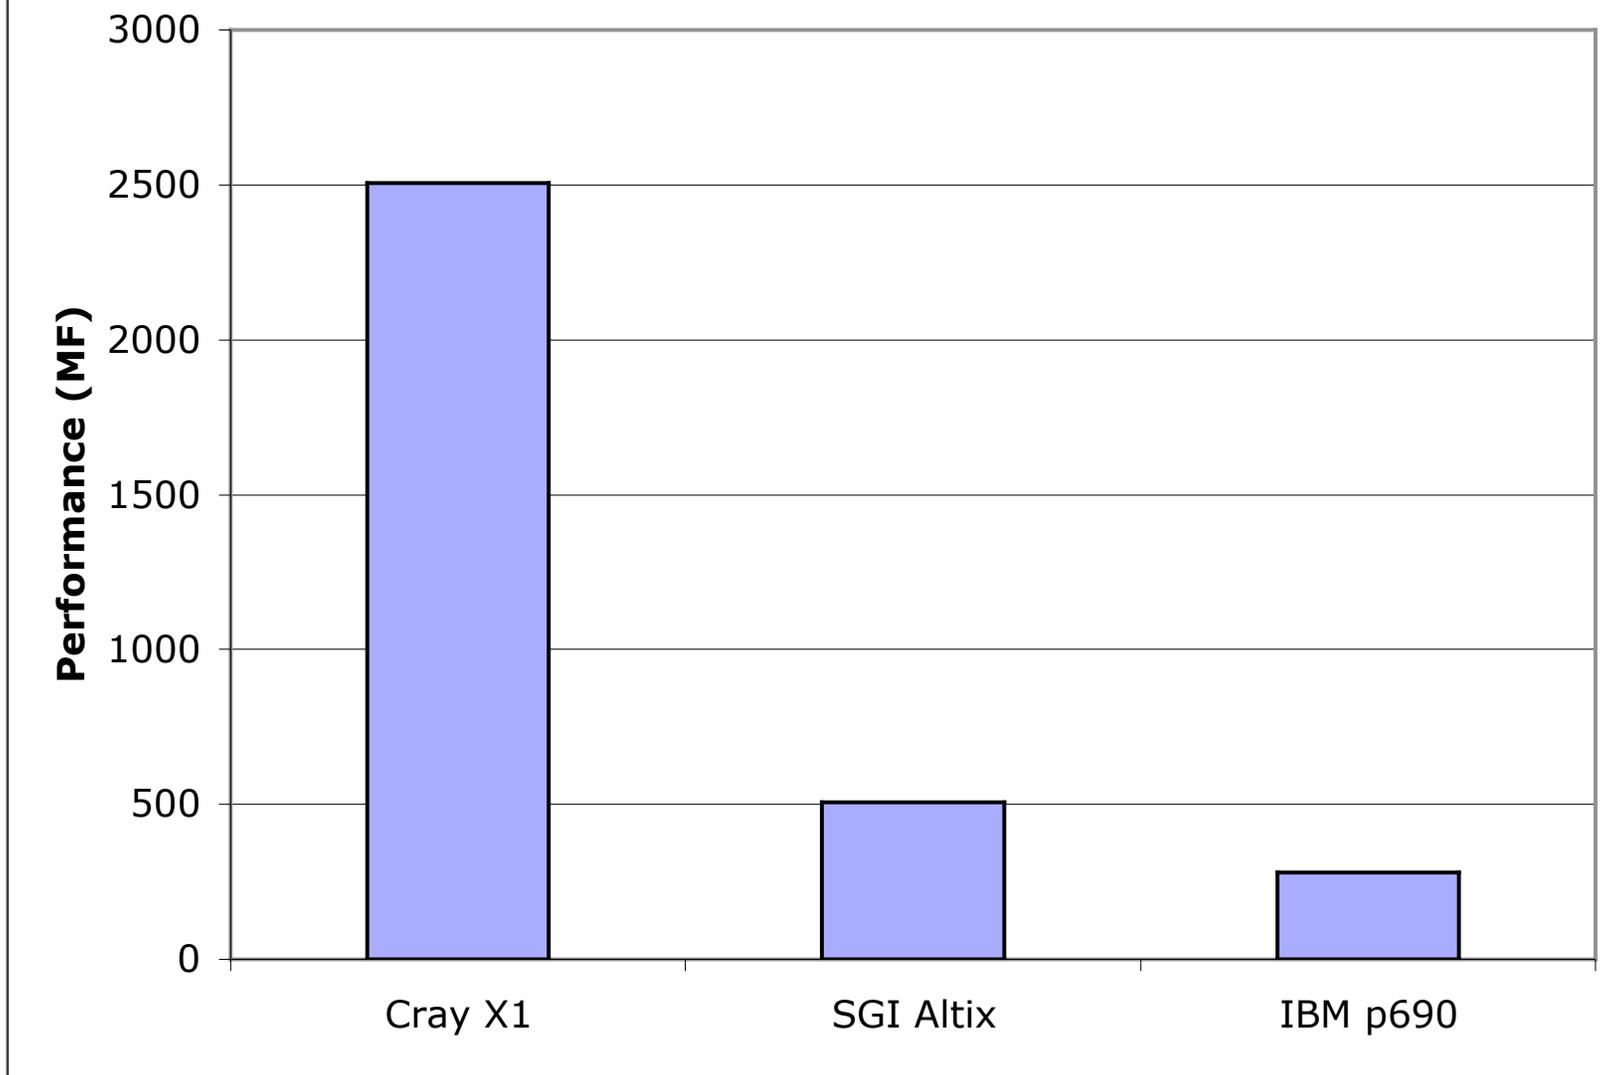
## Problem Size



# DGER Experiment

- Rank-1 matrix update
  - Few floating-point ops per memory op
  - Needs memory bandwidth
- $N = 64 * 70 = 4480$ 
  - Real application performs many moderate DGERs, not a few large ones
  - Cluster size of 64 with 70 time steps
  - Representative of current X1 runs

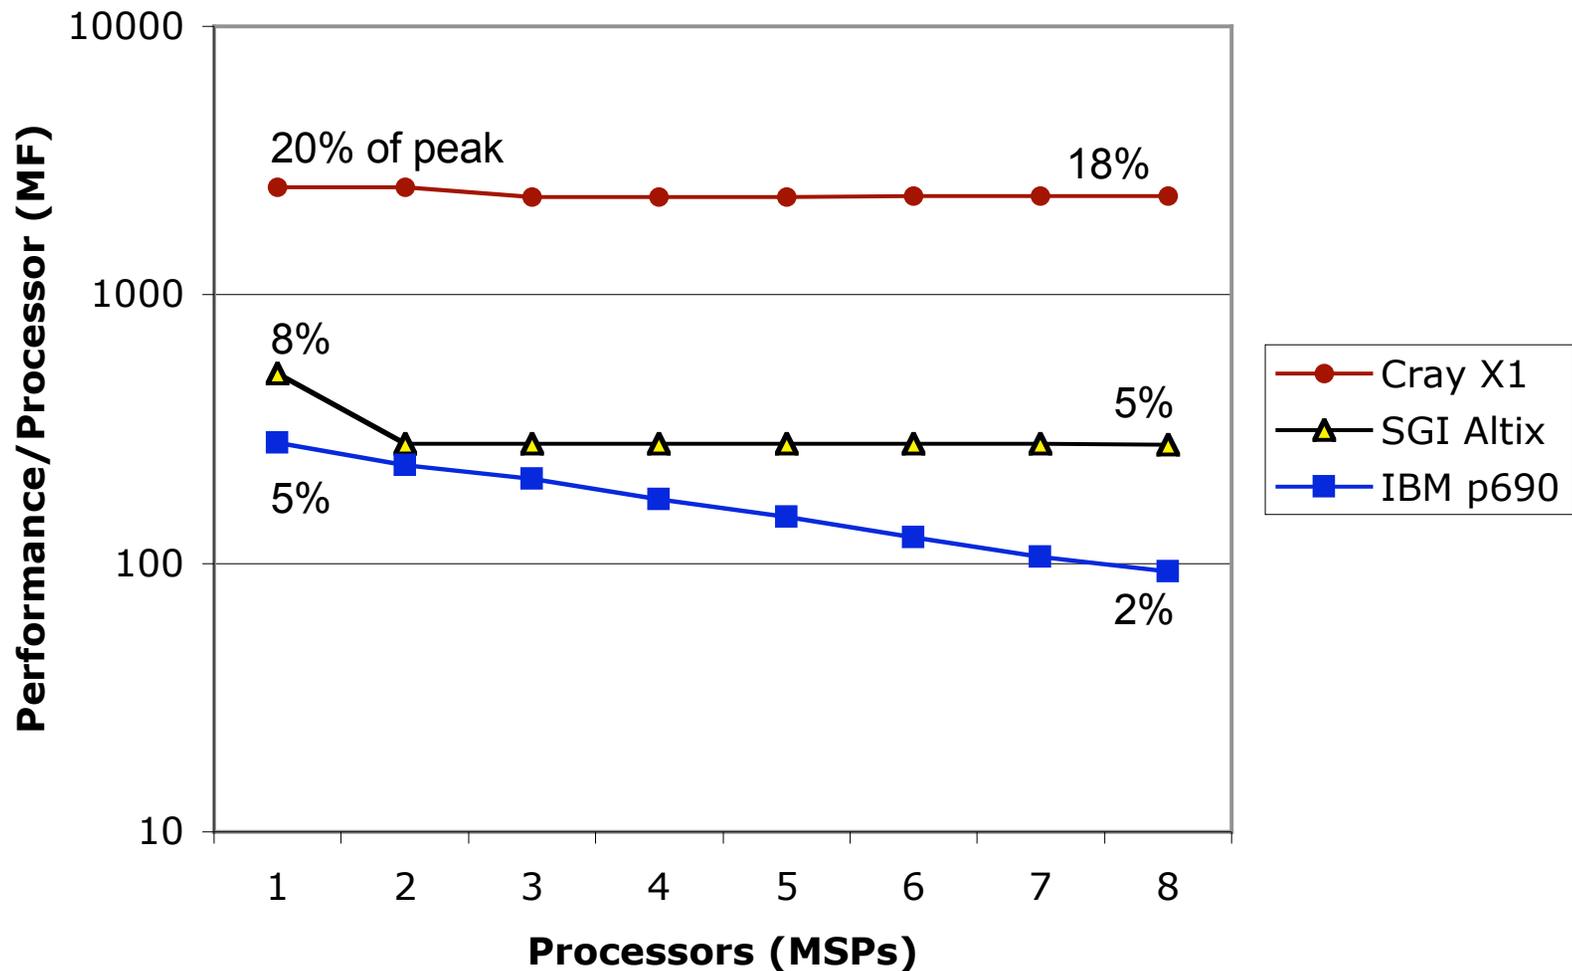
## DGER Performance (N=4480)



# Concurrent-DGERS Experiment

- Real application loads processors with DGERS
- Perform concurrent DGERS, one per processor
- Does memory bandwidth scale?
- Does performance scale?

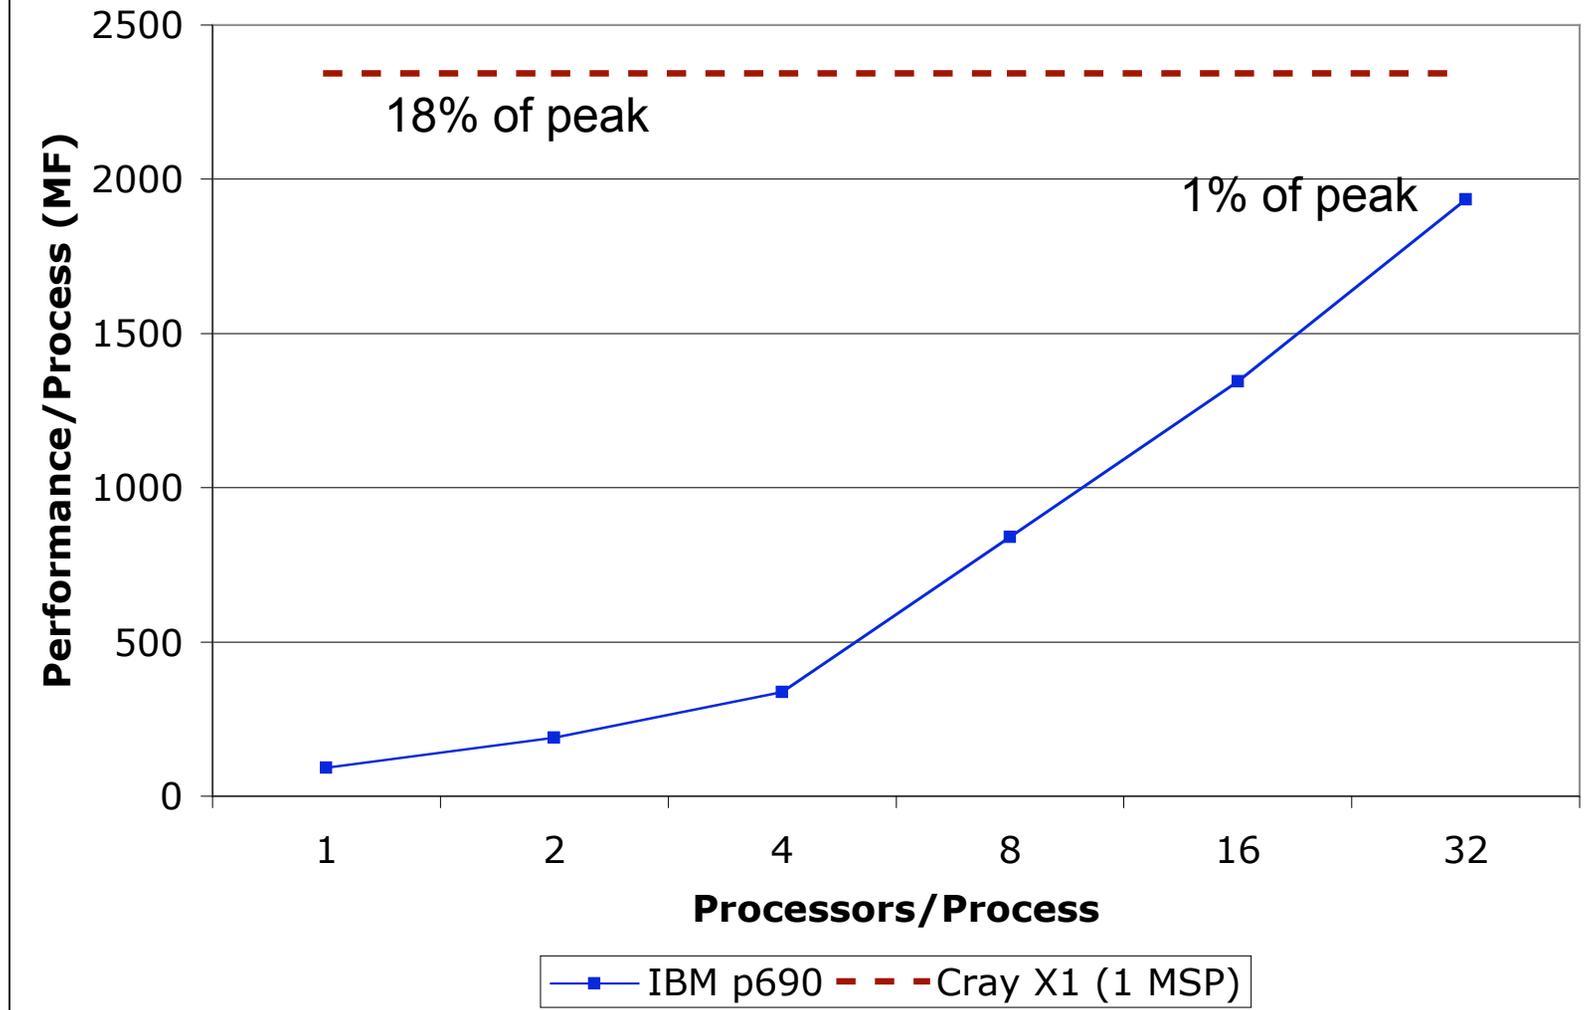
## Performance of Concurrent DGERs (N=4480)



# Threaded-DGER Experiment

- Developers could parallelize each Monte-Carlo process with OpenMP
  - Use SMPs to improve scaling
- Test possibilities using threaded DGER
  - Provided with IBM p690
  - Not yet available for SGI Altix, Cray X1
- Load 32-processor node with DGERs
  - Try different mixes of processes and threads
  - How fast can  $N=4480$  DGER be?

## Performance of Concurrent Threaded DGERS (N=4480)



# Scalability = Capability

- Previous runs used small clusters (2-4)
  - Failed to produce predicted physics
- Cray X1 enables larger clusters (32-64)
  - Powerful processors, scalable memory bandwidth
  - Predicted physics has been restored
- Plans for full 3D structures
  - Should validate or refute full theory
- Larger Crays may enable prediction of new materials

# Early Evaluation of the Cray X1

Pat Worley *et al.*

**Today at 2 PM Room 38-39**