

Towards Full Simulations of High-Temperature Superconductors

T.A. Maier, J.B. White, T.C. Schulthess

Oak Ridge National Laboratory

M. Jarrell

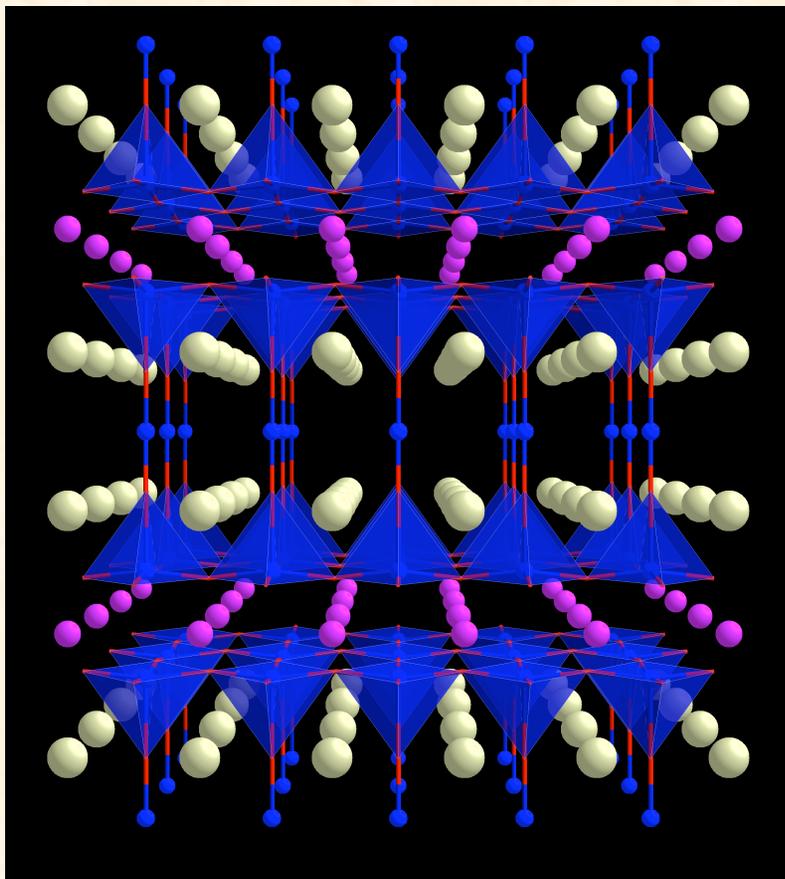
University of Cincinnati

Acknowledgements

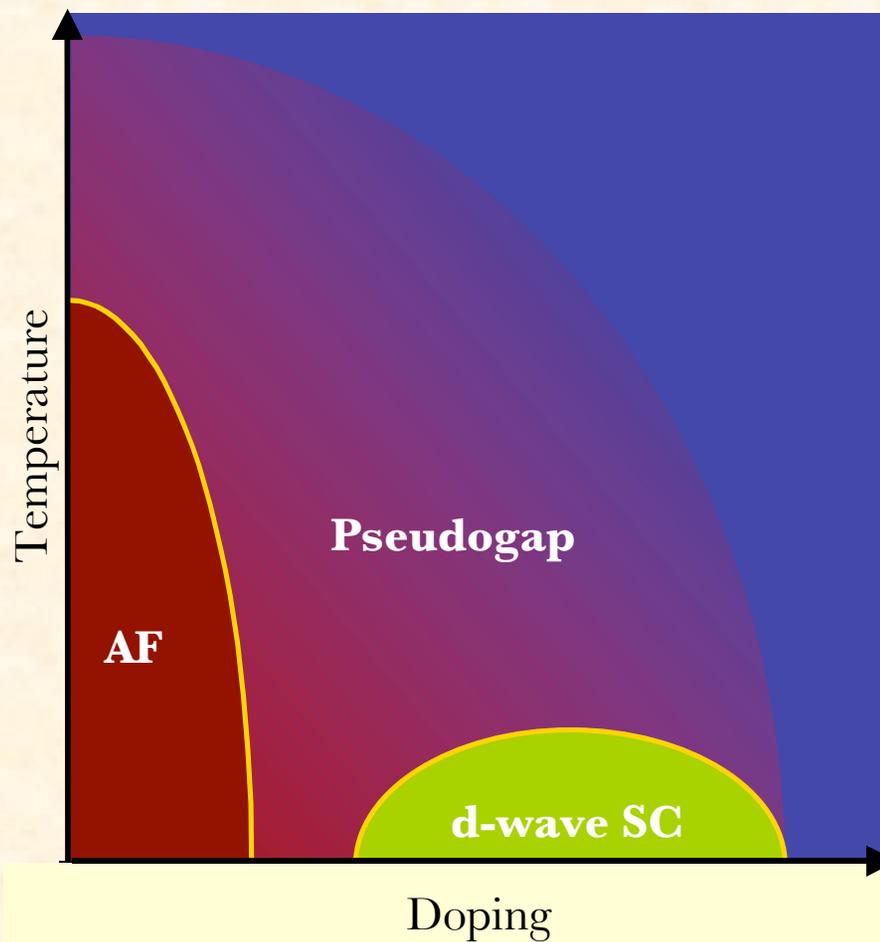
This research used resources of the Center for Computational Sciences and was sponsored in part by the offices of Advance Scientific Computing Research and Basic Energy Sciences, U.S. Department of Energy. The work was performed at Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725, and where TM is a Eugene P. Wigner Fellow. Work at Cincinnati was supported by the NSF Grant No. DMR-0113574.

High- T_c cuprate superconductors

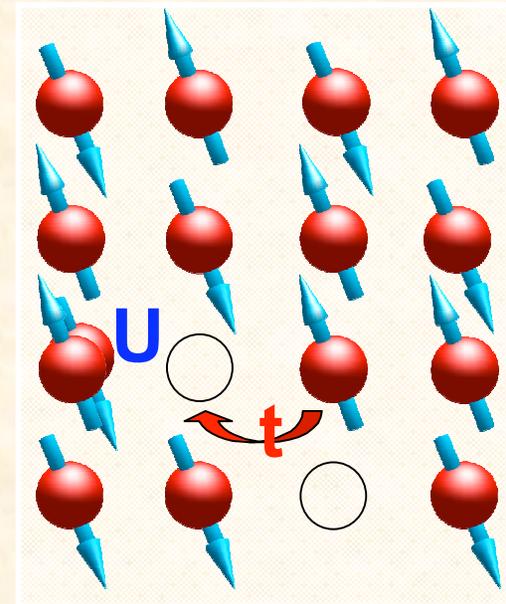
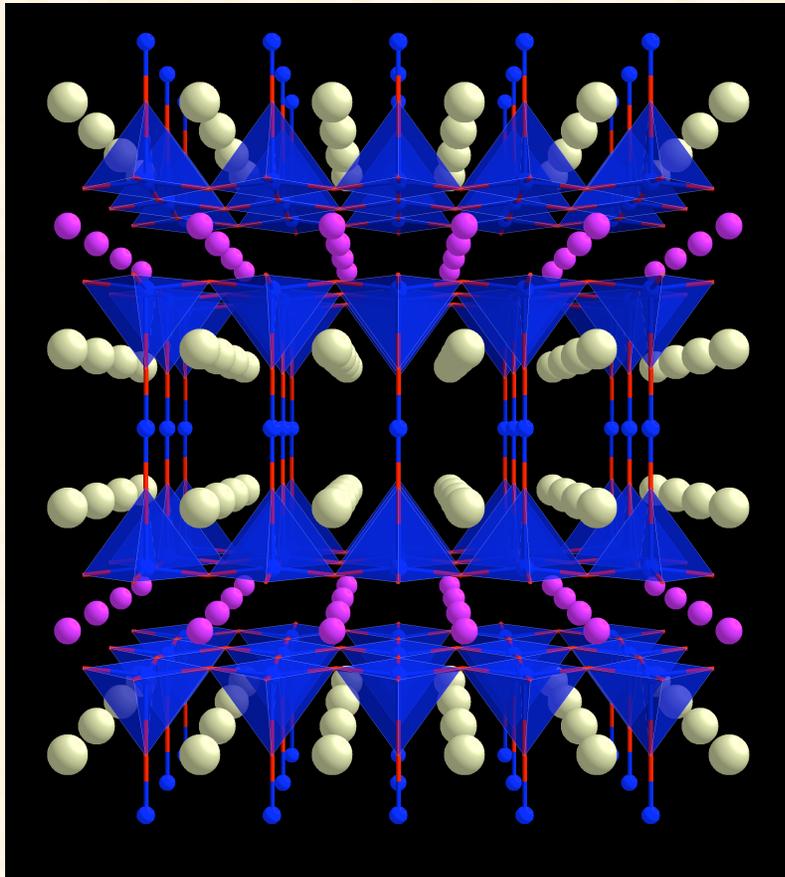
Layered structure



Generic HTSC Phase diagram



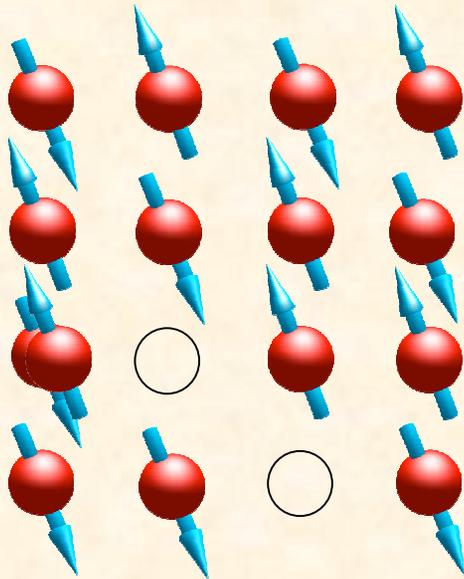
Hubbard model of HTSC



- N interacting electrons on lattice
- Problem: $N \approx 10^{23}$
- Solvable only in 1D, not 2D

Approximative Methods

- **Finite size simulations**



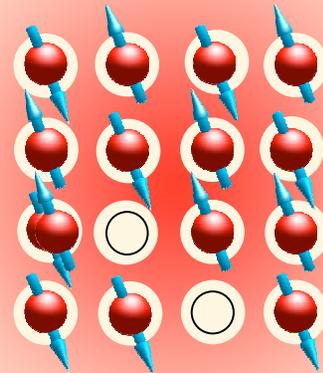
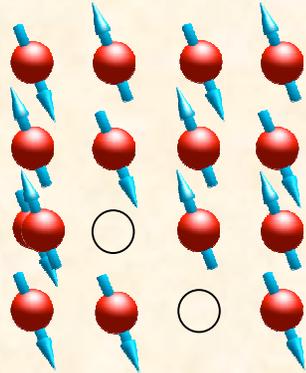
- ✓ **Exact solution for**
 $N \approx 50$

- **Mean-field approach**



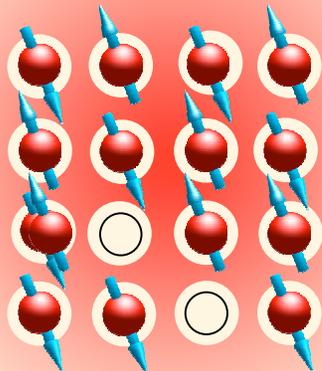
- ✓ **Thermodynamic limit**
 $(N = \infty)$

Dynamical Cluster Approximation



Dynamical Cluster Approximation

$$G_0(K, \omega) = \left[\left(\frac{N_c}{N} \sum_{\vec{k}} \frac{1}{\omega + \mu - \epsilon_{K+\vec{k}} - \Sigma_c(K, \omega)} \right)^{-1} + \Sigma_c(K, \omega) \right]^{-1}$$



QMC $\rightarrow \Sigma_c[G_0(K, \omega), U]$

- ✓ Thermodynamic limit
- ✓ Non-local correlations

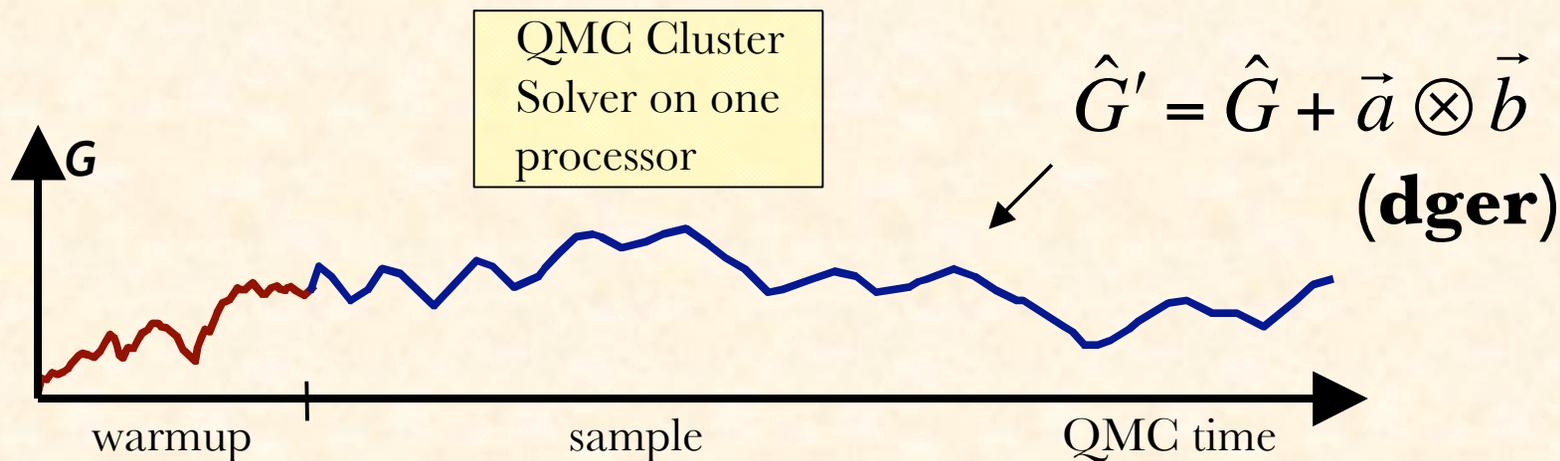
Cluster in reciprocal space



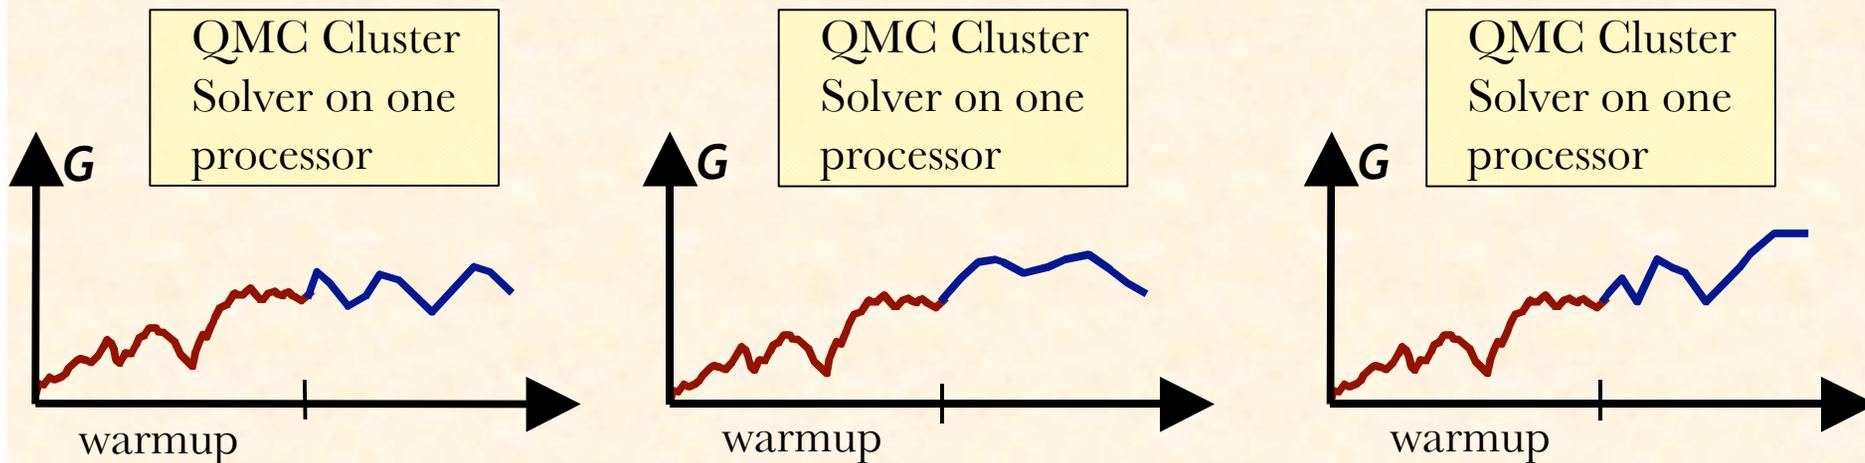
- ✓ Translational symmetry

DCA/QMC Cluster solver

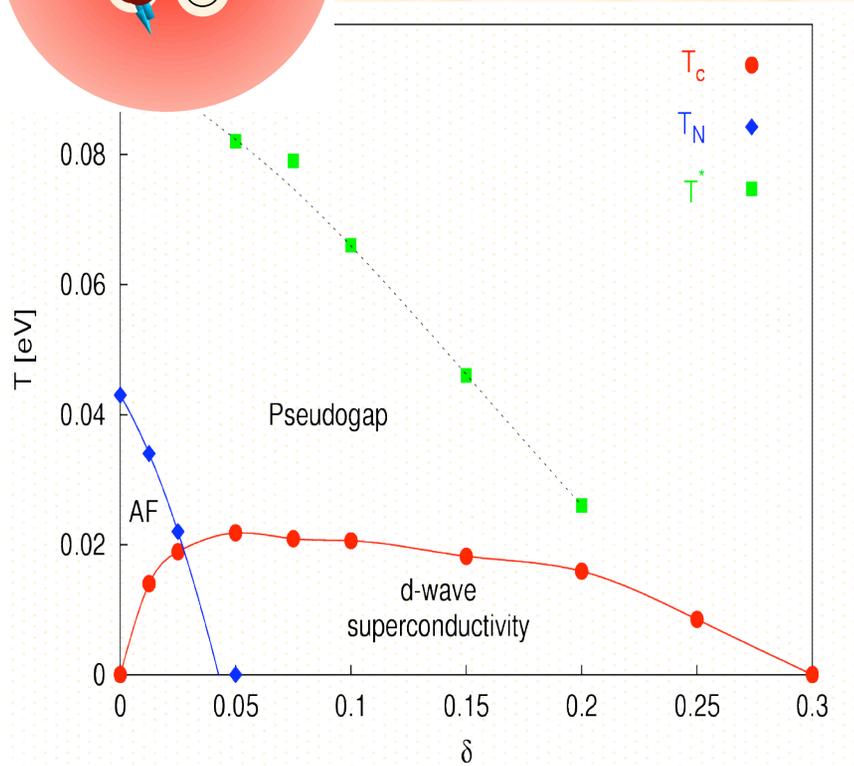
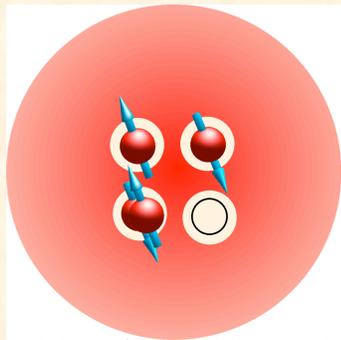
Serial:



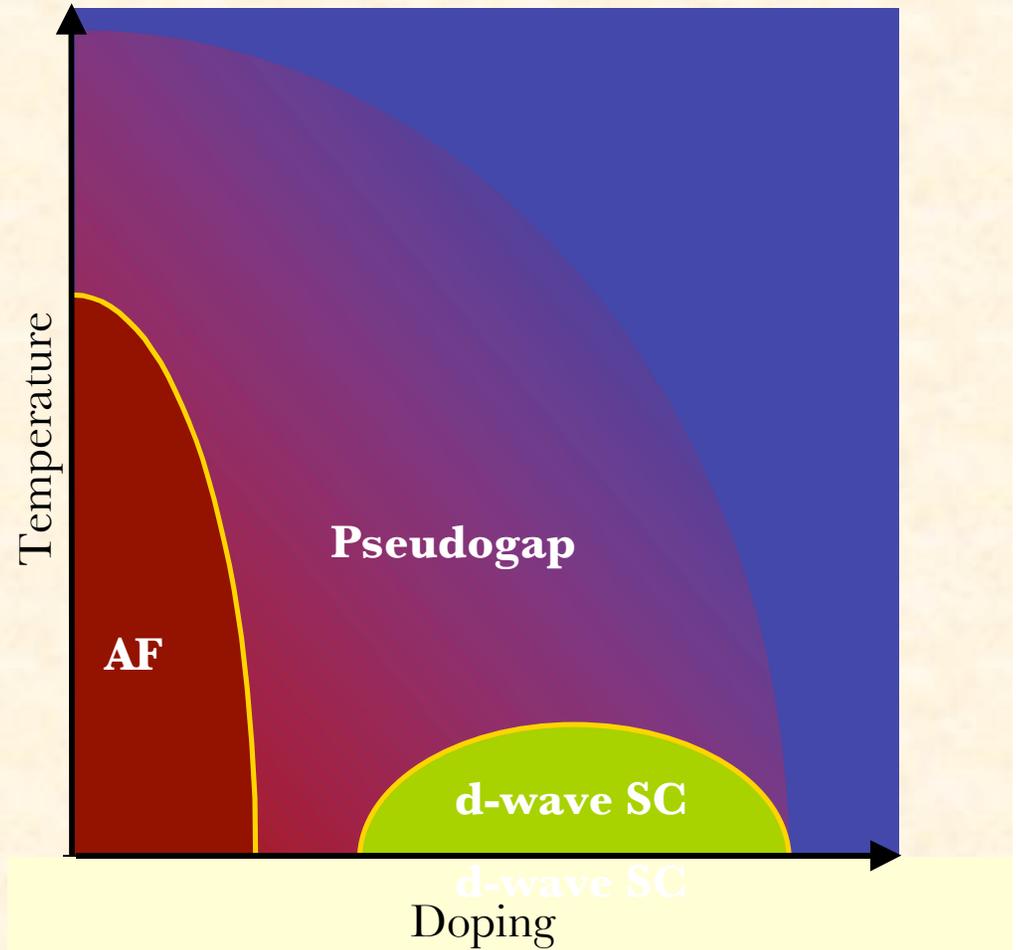
Perfectly parallel:



4-site cluster



Generic HTSC Phase diagram



Problem ...

- **Results in contradiction with Mermin-Wagner theorem:**
 - *No long-range order at finite temperatures in 2D systems if broken symmetry is continuous.*
- **Consequences:**
 - $T_N = 0$
 - Superconductivity only possible as Kosterlitz-Thouless topological order
- **Violation caused by small cluster**
- **Cure: Simulate larger clusters**
 - Computational cost grows like $\sim N_c^3$

Porting and tuning on Cray X1

- **Easy port**
 - **Modifications to Makefile**
- **Tuning**
 - **Performance profile**
 - **Loopmarks**
 - **Unvectorized nested loops with indirect addressing (all in one file)**
 - **Figure out which loops are independent**
 - **!dir\$ concurrent**

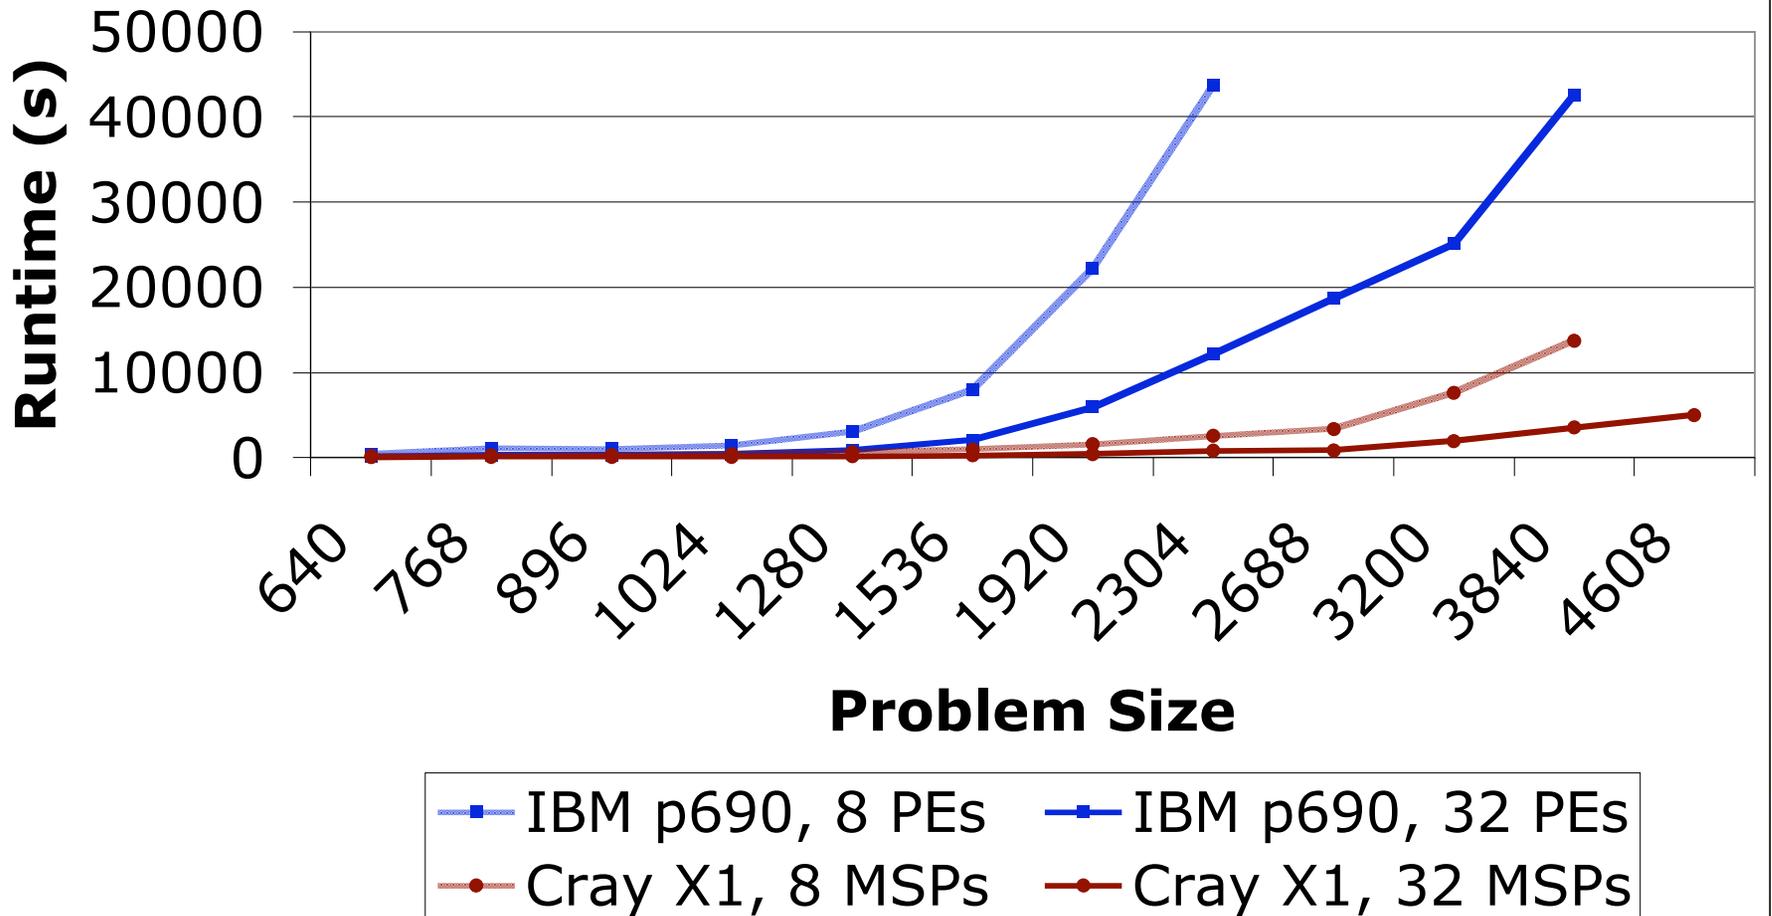
Performance

- **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**

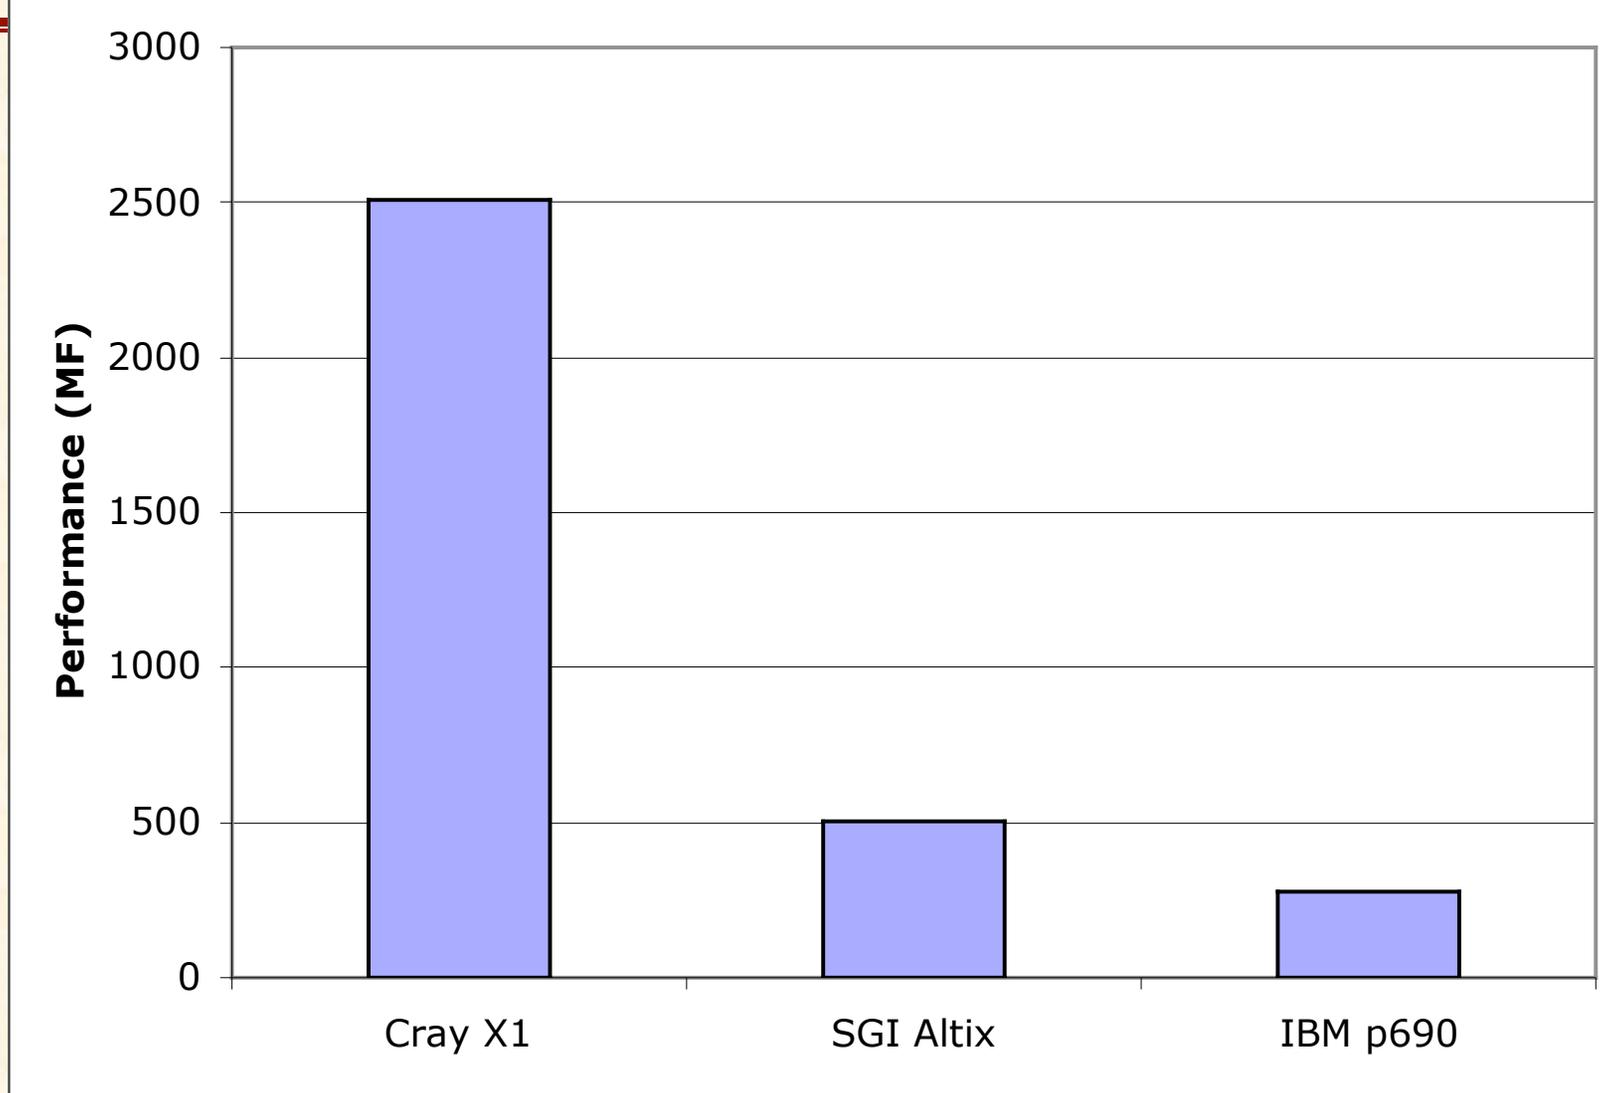
DCA-QMC Runtime



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**

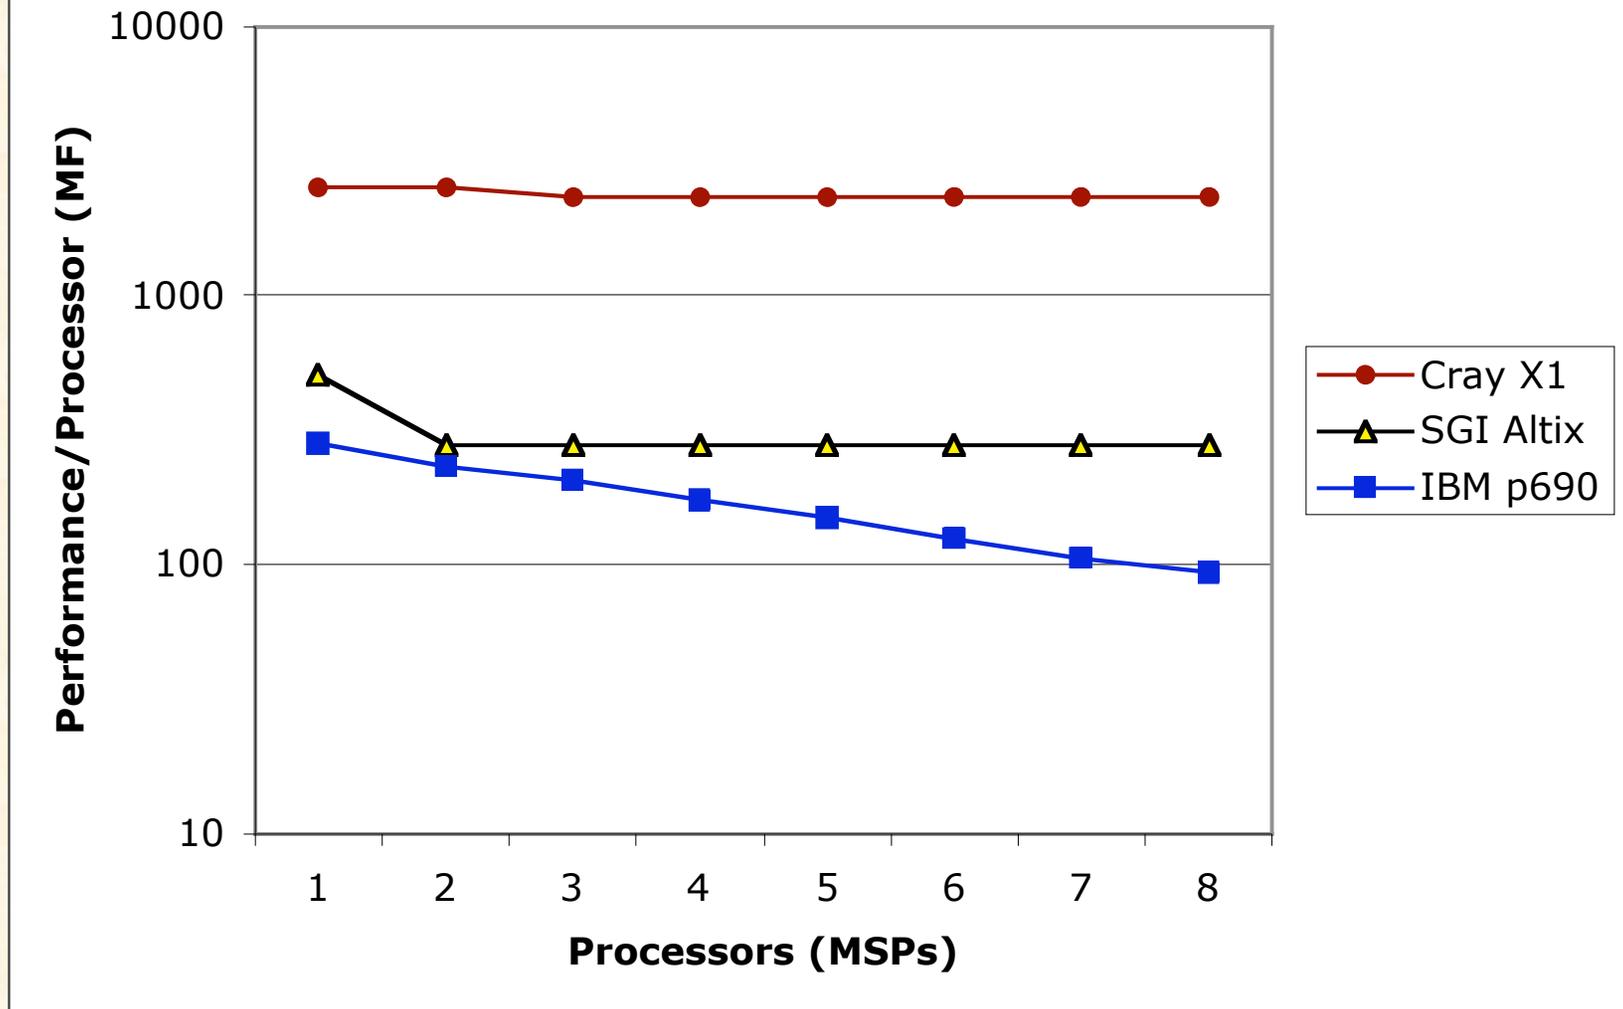
DGER Performance (N=4480)



Concurrent-DGERs Experiment

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

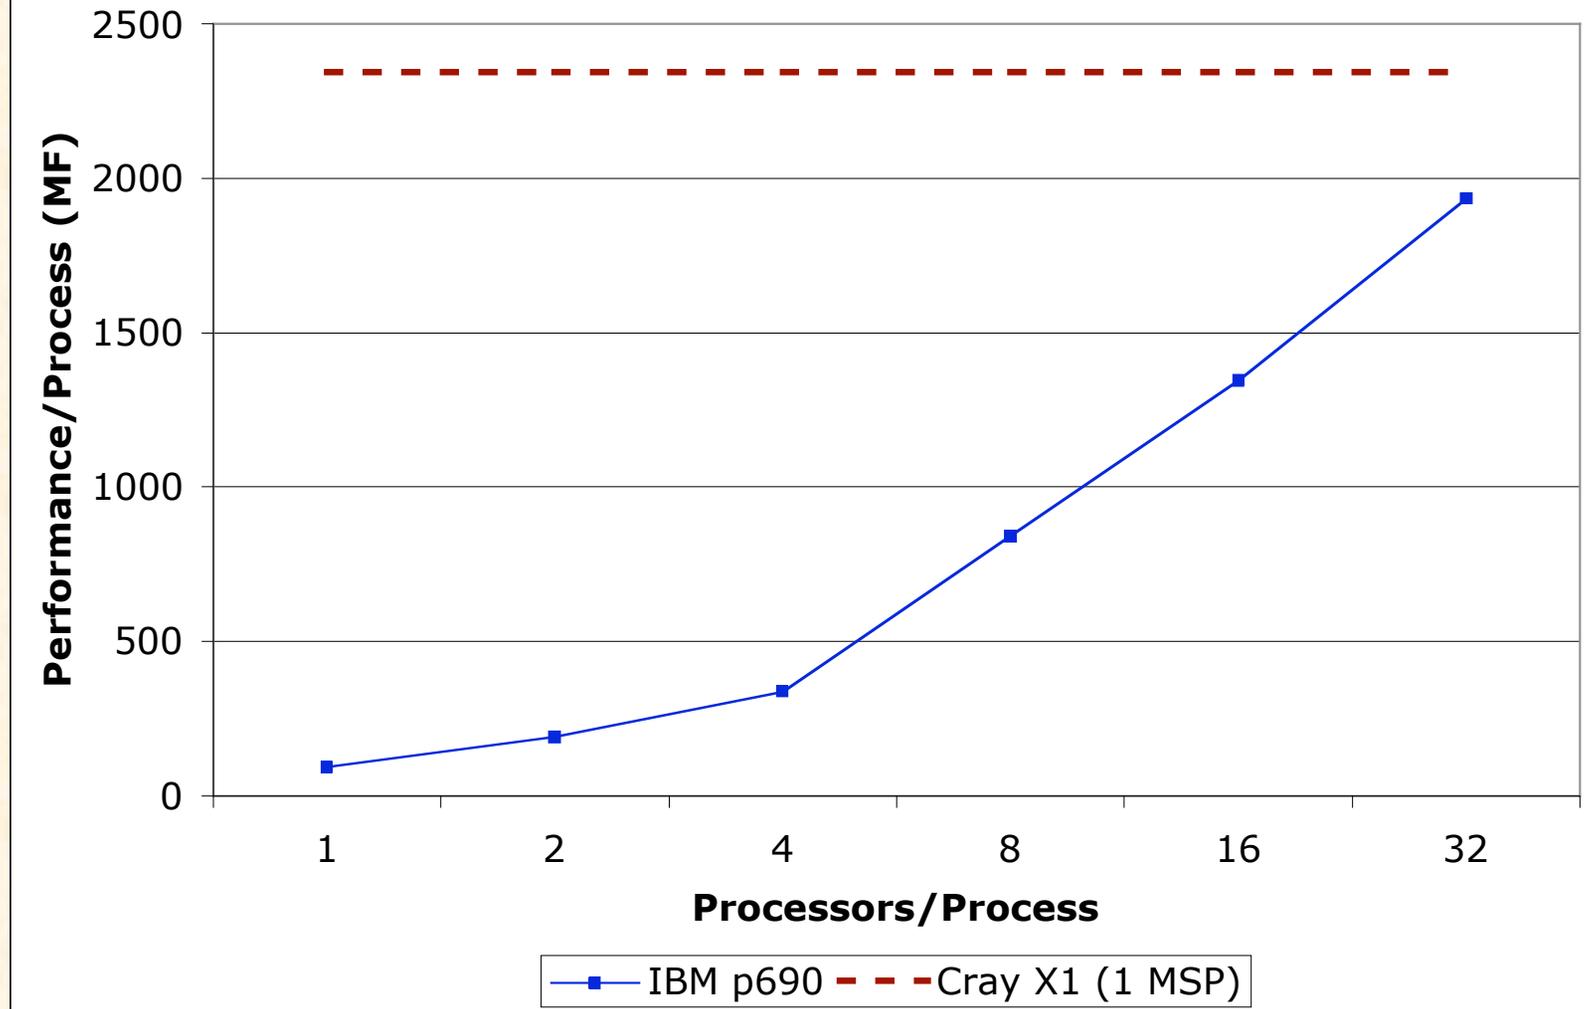
Performance of Concurrent DGERs (N=4480)



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?

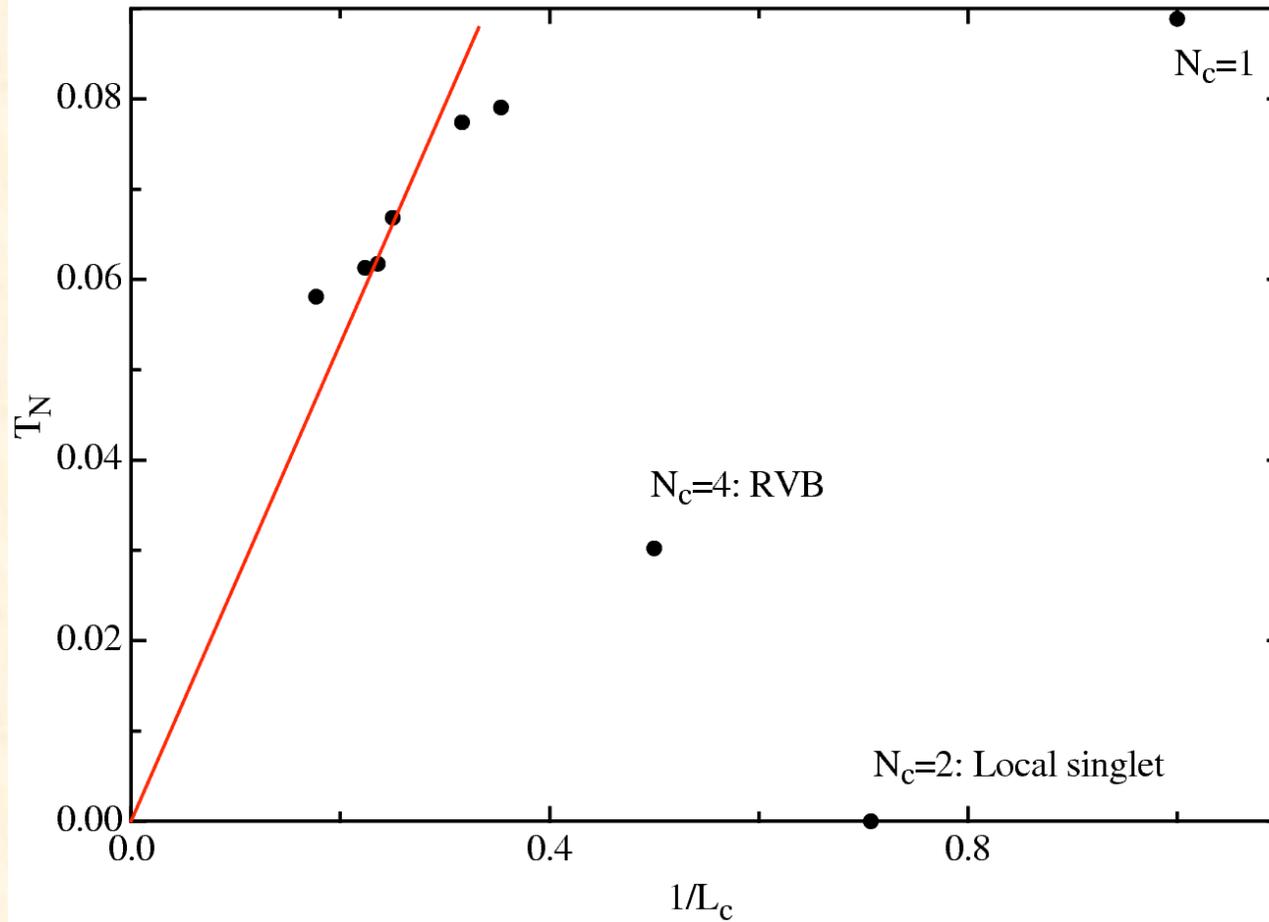
Performance of Concurrent Threaded DGERs (N=4480)



Scalability = Capability

- **Cray X1 enables larger clusters (32-64)**
 - **Powerful processors**
 - **scalable memory bandwidth**
- **Provides capability**
 - **To check compliance with Mermin-Wagner theorem**
 - **To study possible Kosterlitz-Thouless transition to phase with superconducting topological order**

Larger clusters - Antiferromagnetism

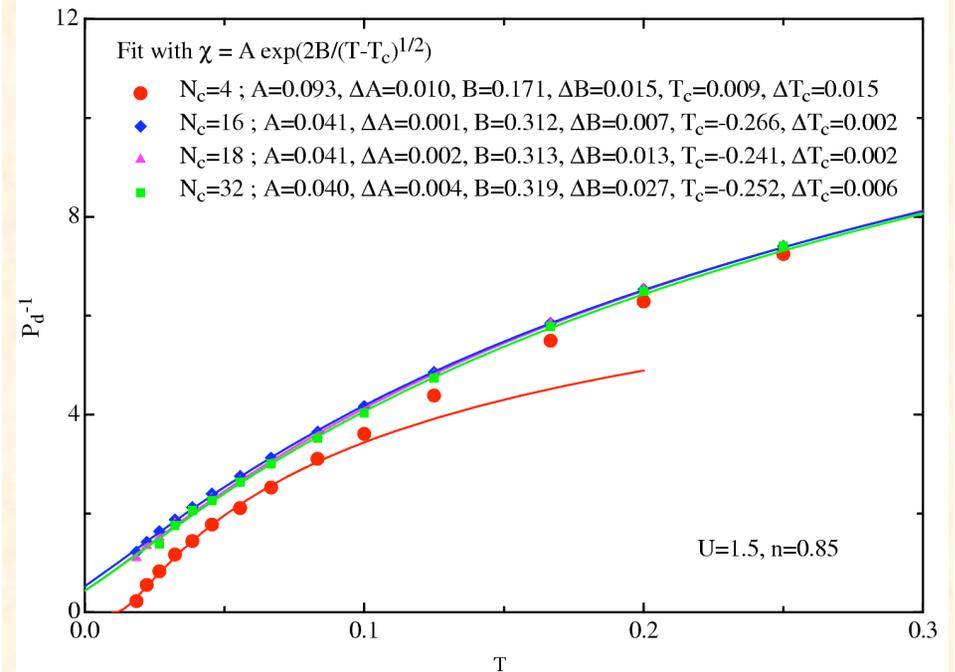
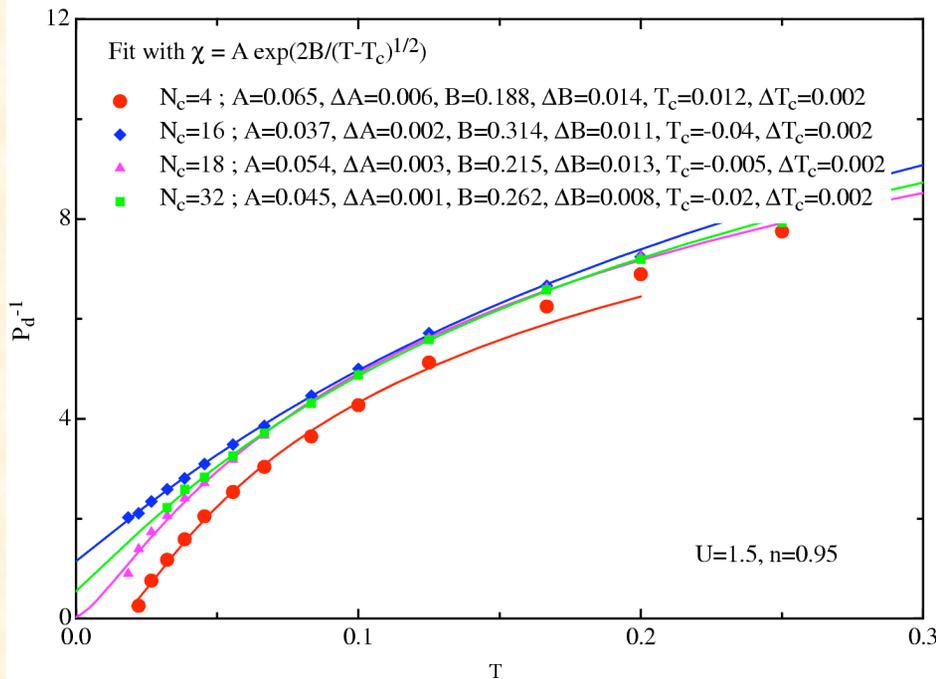


✓ **Mermin-Wagner theorem recovered**

OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY

UT-BATTELLE

Larger clusters - Superconductivity



- **Superconductivity suppressed at larger clusters**
- **No Kosterlitz-Thouless transition**

Conclusions: Scalability = Capability

- **Previous runs used small clusters (2-4)**
 - Violate fundamental theorem
- **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**