ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

CSCS
Swiss National Supercomputing Centre

High-Performance Computing & Simulations in Quantum Many-Body Systems – PART II

Thomas Schulthess

schulthess@phys.ethz.ch

Recommended reading (part I)

- PGAS languages
  - Tarek El-Ghazawi, Willam Carison, Thomas Sterling and Katherine Yelick, “UPC: Distributed Shared Memory Programming” – Wiley 2005

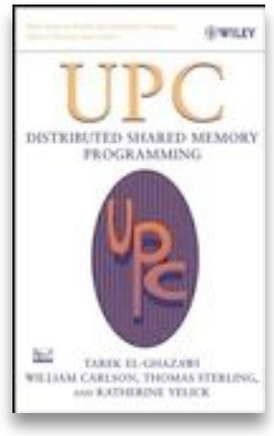

  - Co-Array Fortran (Fortran with co-arrays) – http://www.co-array.org/
- Distributed or shared memory model
  - POSIX Thread (pthread): https://computing.llnl.gov/tutorials/pthreads
  - OpenMP: https://computing.llnl.gov/tutorials/openMP/ & BOOST
  - MPI: https://computing.llnl.gov/tutorials/mpi/
- History of Computing
  - Paul E. Curezzi, “A History of Modern Computing” – MIT Press 2003

Relationship between simulations and supercomputer system (the old way of doing business)

Simulations

Science

Model & method of solution

?

Port codes developed on workstations
> vectorize codes
> parallelize codes
> petascaling and soon exascaling

Basic numerical libraries

Programming environment

Runtime system

Operating systems

Computer Hardware

Supercomputer

Applications running at scale on Jaguar @ ORNL

Fall 2009

Domain area	Code name	Institution	# of cores	Performance	Notes
Materials	DCA++	ORNL	213,120	**1.9 PF**	2008 Gordon Bell Prize Winner
Materials	WL-LSMS	ORNL/ETH	223,232	**1.8 PF**	2009 Gordon Bell Prize Winner
Chemistry	NWChem	PNNL/ORNL	224,196	1.4 PF	2008 Gordon Bell Prize Finalist
Materials	OMEN	Duke	222,720	860 TF	
Chemistry	MADNESS	UT/ORNL	140,000	550 TF	
Materials	LS3DF	LBL	147,456	442 TF	2008 Gordon Bell Prize Winner
Seismology	SPECFEM3D	USA (multiple)	149,784	165 TF	2008 Gordon Bell Prize Finalist
Weather	WRF	USA (multiple)	150,000	50 TF	
Combustion	S3D	SNL	147,456	83 TF	

Superconductivity: a state of matter with zero electrical resistivity

Heike Kamerlingh Onnes (1853-1926) Discovery 1911

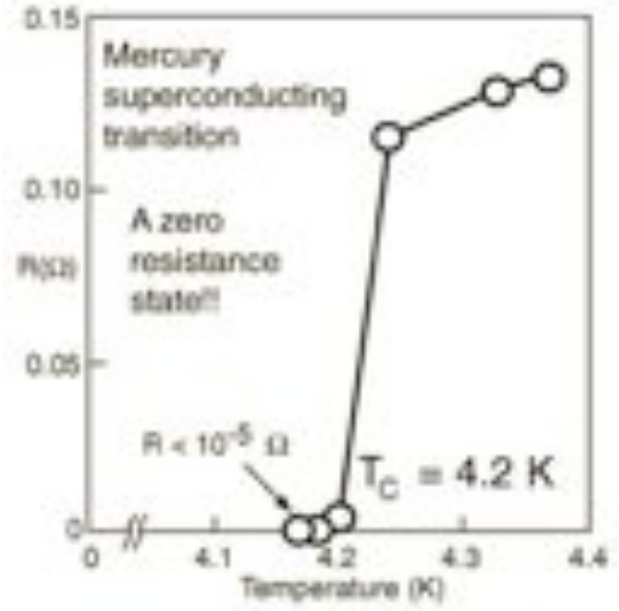

Microscopic Theory 1957


PHYSICAL REVIEW VOLUME 108, NUMBER 5 DECEMBER 1, 1957

Theory of Superconductivity*

J. BARDEEN, L. N. COOPER,† AND J. R. SCHRIEFFER‡
Department of Physics, University of Illinois, Urbana, Illinois
(Received July 8, 1957)

A theory of superconductivity is presented, based on the fact that the interaction between electrons resulting from virtual exchange of phonons is attractive when the energy difference between the electronic states involved is less than the phonon energy, $\hbar\omega$. It is favorable to form a superconducting phase when this attractive interaction dominates the repulsive screened Coulomb interaction. The normal phase is described by the Bloch individual-particle model. The ground state of a superconductor, formed from a linear combination of normal state configurations in which electrons are virtually excited in pairs of opposite spin and momentum, is lower in energy than the normal state by an amount proportional to an average $(\hbar\omega)^2$, consistent with the isotope effect. A mutually orthogonal set of excited states in one-to-one correspondence with those of the normal phase is obtained by specifying occupation of certain Bloch states and by using the rest to form a linear combination of virtual pair configurations. The theory yields a second-order phase transition and a Meissner effect in the form suggested by Pippard. Calculated values of specific heats and penetration depths and their temperature variation are in good agreement with experiment. There is an energy gap for individual-particle excitations which decreases from about $3.5kT_c$ at $T=0°K$ to zero at T_c. Tables of matrix elements of single-particle operators between the excited-state superconducting wave functions, useful for perturbation expansions and calculations of transition probabilities, are given.

Superconductivity in the cuprates 1986

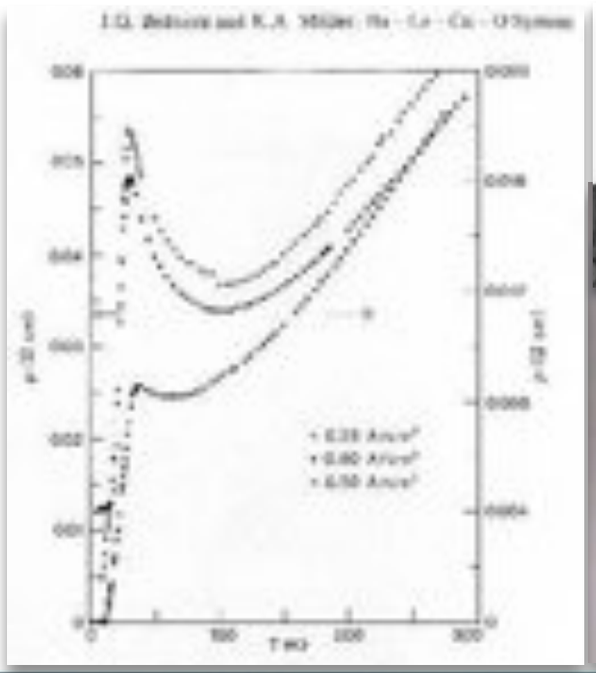

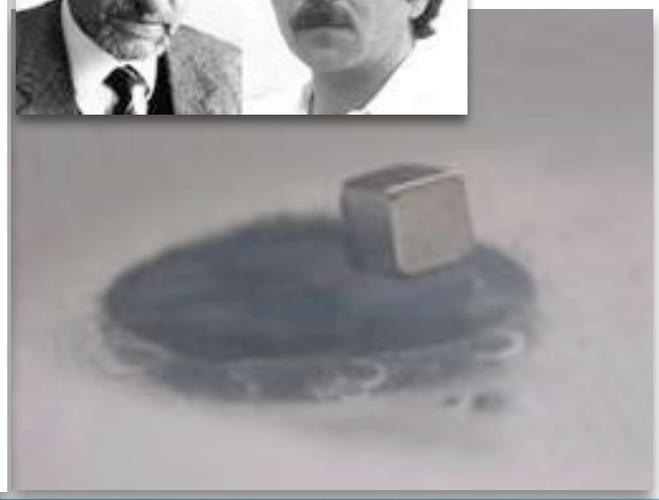

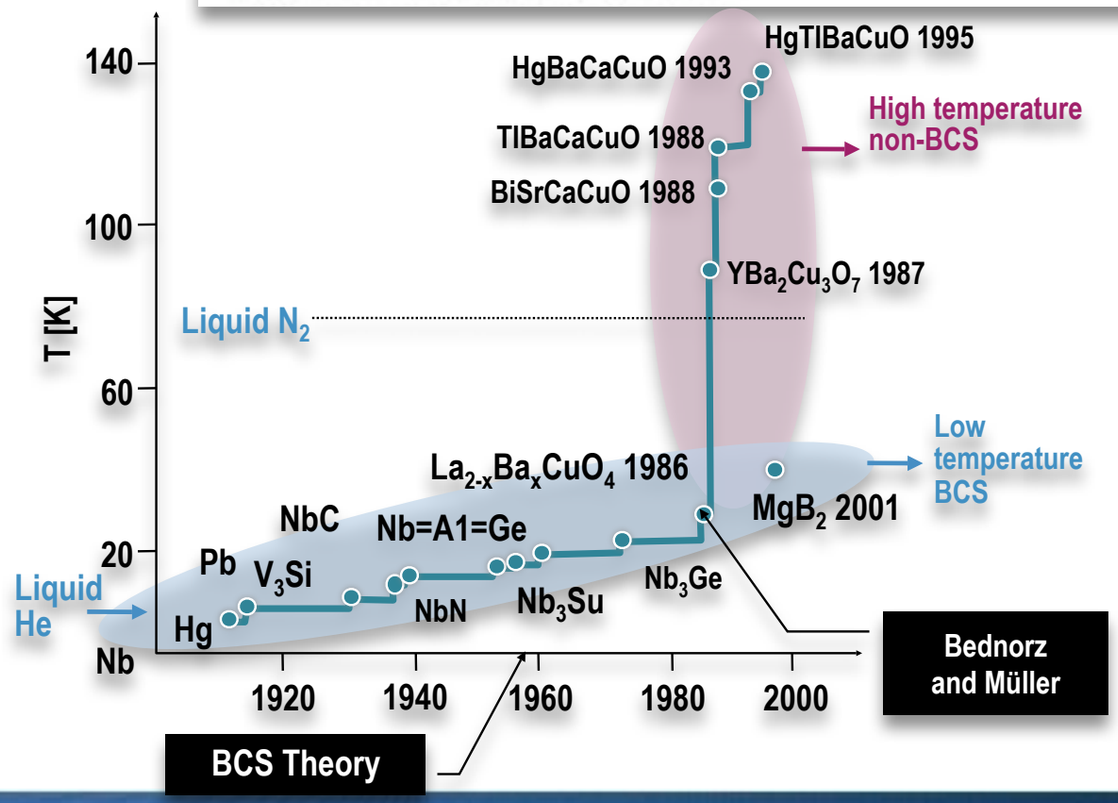

From cuprate materials to the Hubbard model

La_2CuO_4

CuO_2 plane

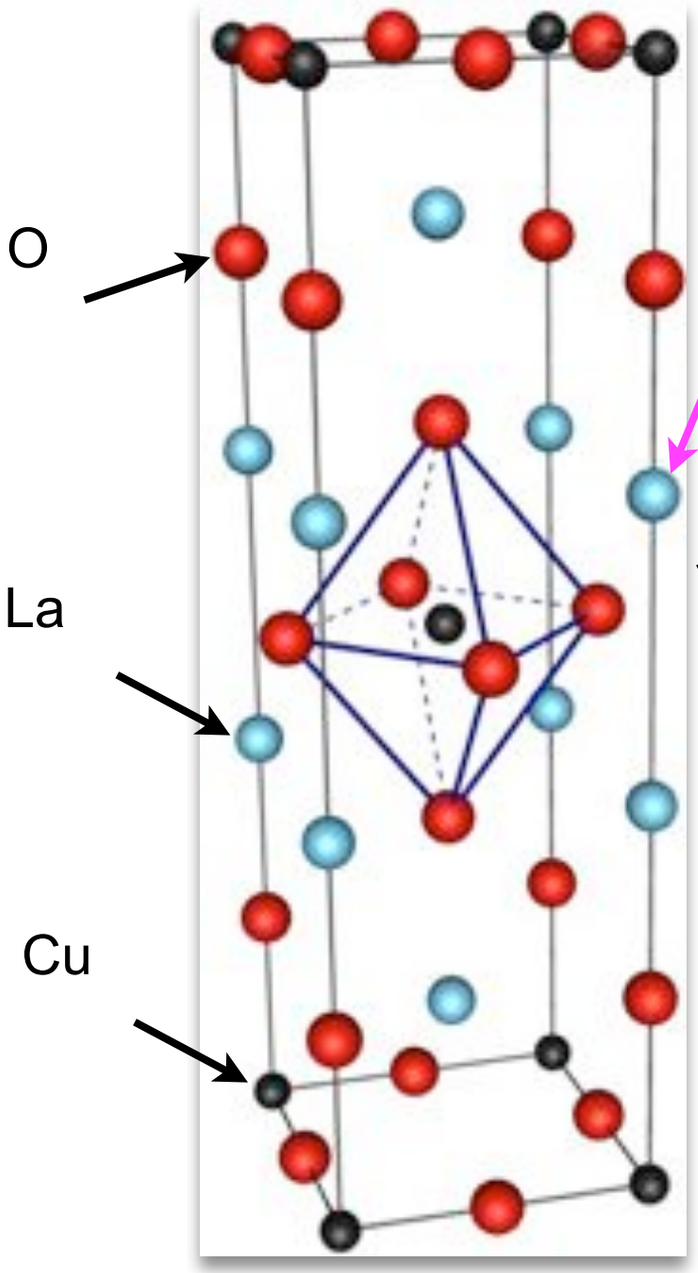

Sr doping introduces "holes"

Holes form Zhang-Rice singlet states

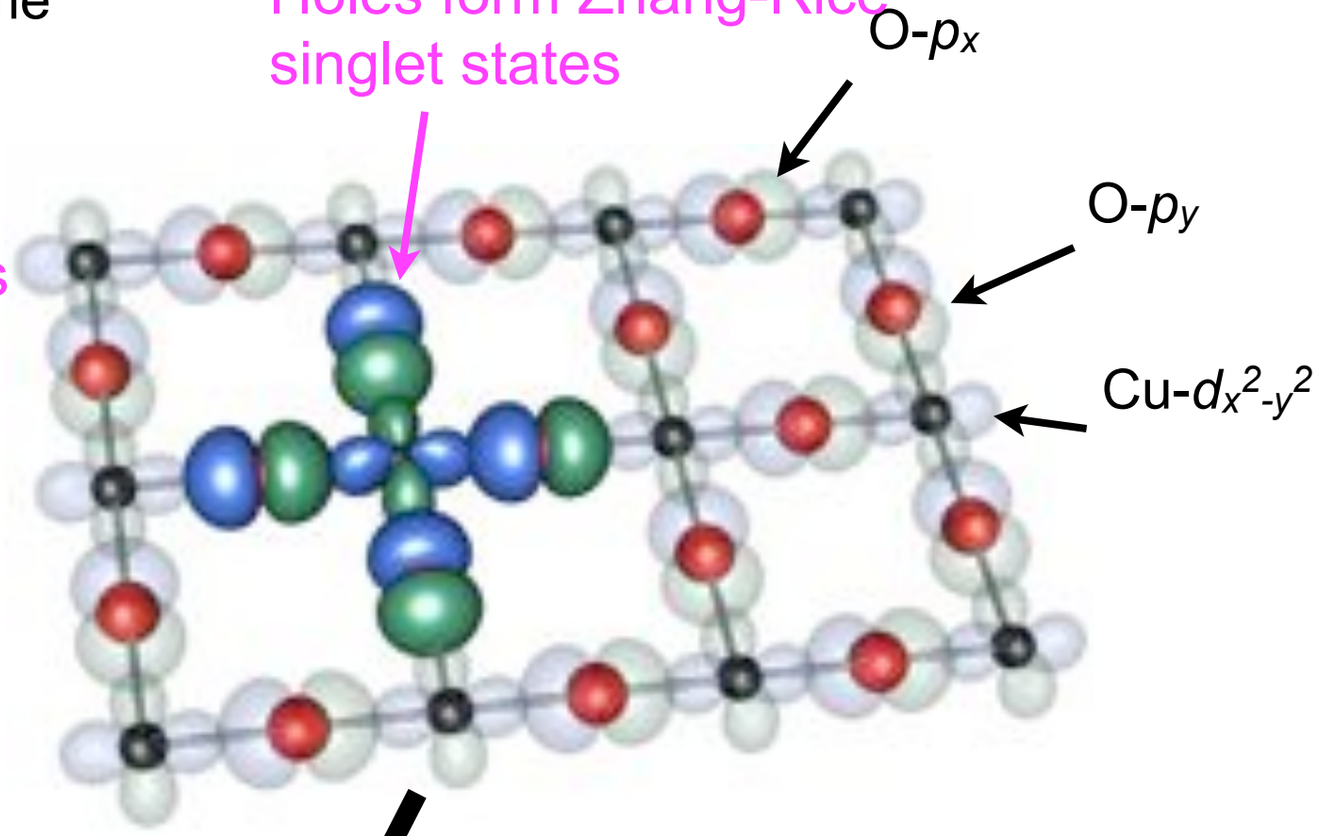

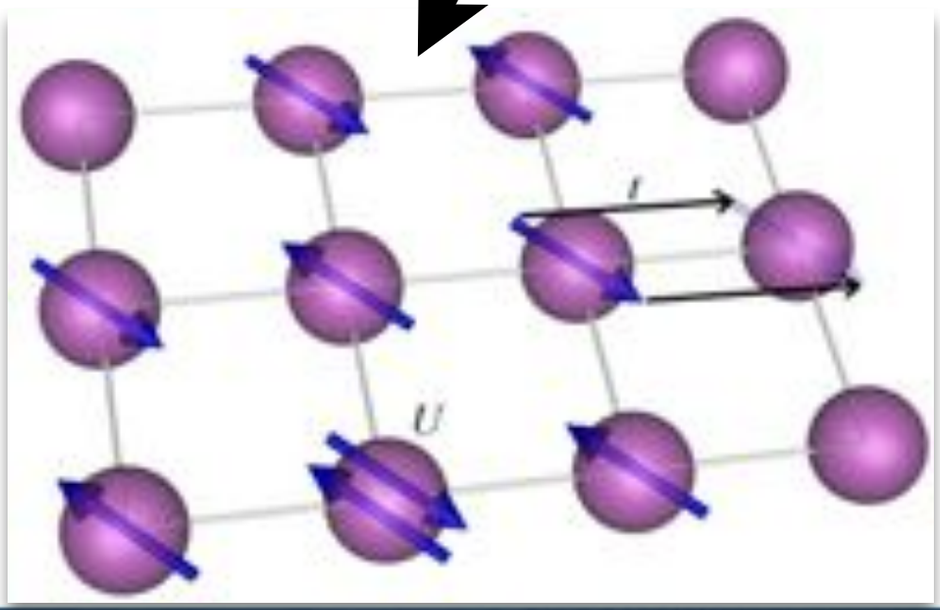

Single band 2D Hubbard model

2D Hubbard model and its physics

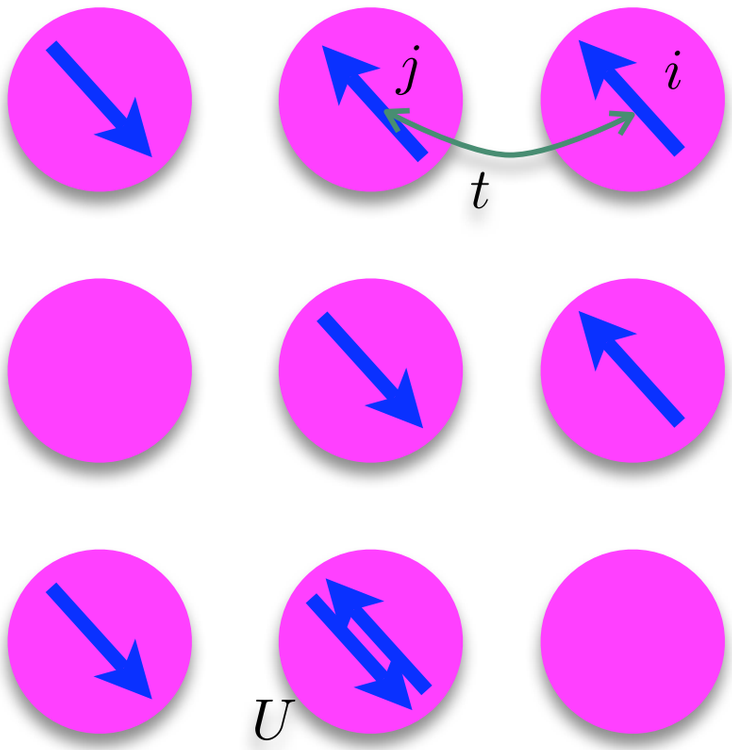

Half filling: number of carriers = number of sites

Formation of a **magnetic moment** when *U* is large enough

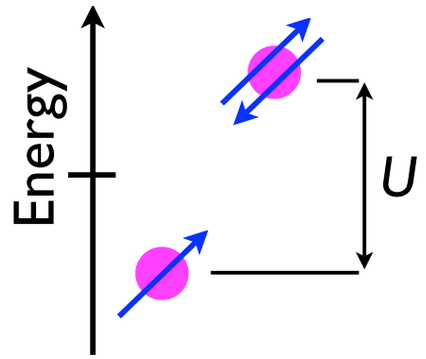

Antiferromagnetic alignment of neighboring moments

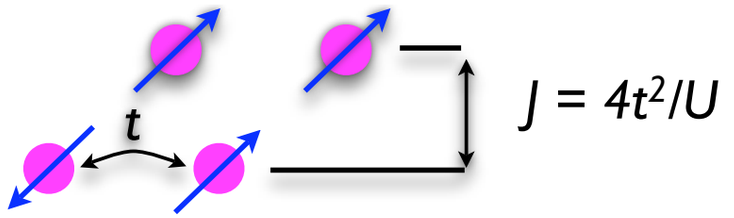

1. When *t* >> *U*:

Model describes a metal with band width *W*=8*t*

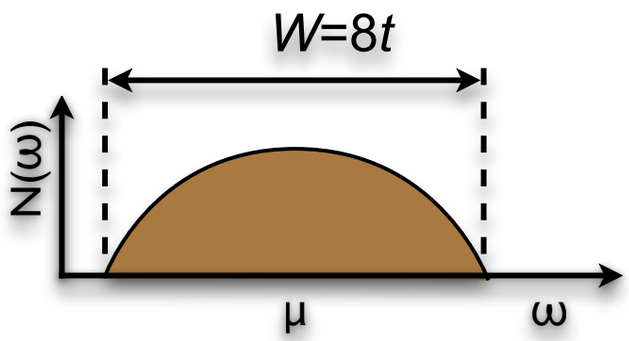

2. When *U* >> 8*t* at half filling (not doped)

Model describes a “Mott Insulator” with antiferromagnetic ground state (as seen experimentally seen in undoped cuprates)

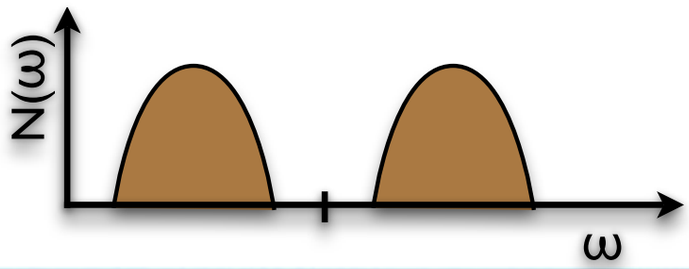

Hubbard model for the cuprates

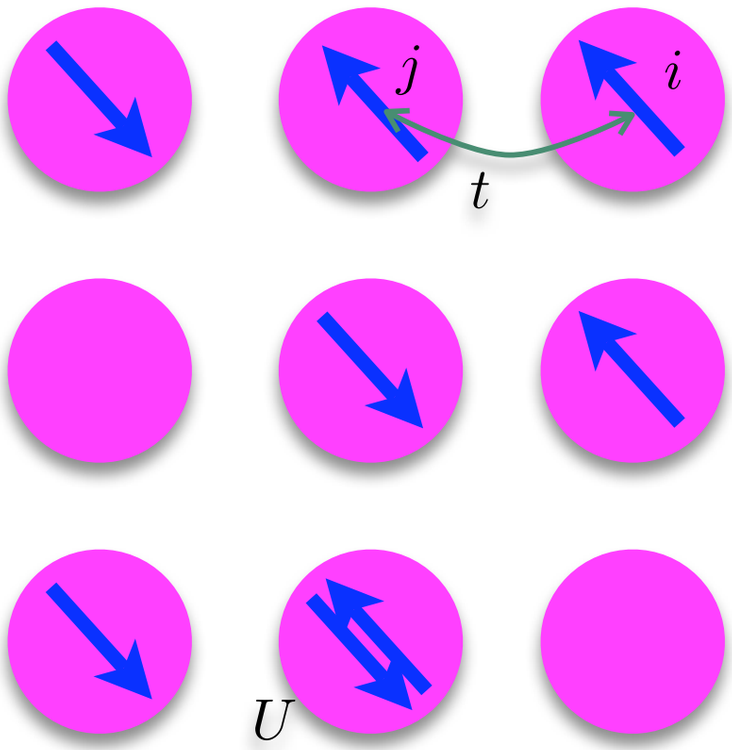

Half filling: number of carriers = number of sites

Formation of a **magnetic moment** when *U* is large enough

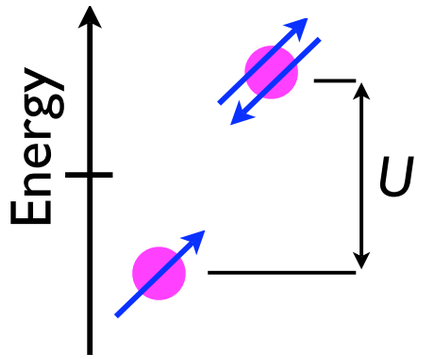

Antiferromagnetic alignment of neighboring moments

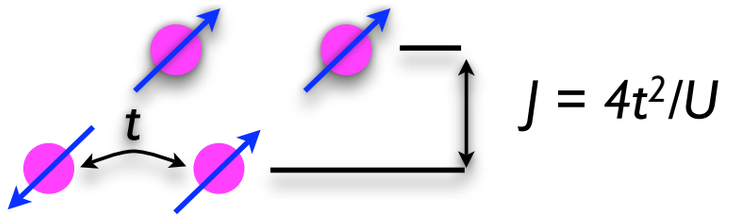

3. Parameter range relevant for superconducting cuprates

$U \approx 8t$

No simple solution!

Finite doping levels (0.05 – 0.25)

Typical values: U~10eV; t~0.9eV; J~0.2eV; (0.1eV ~ 10^3 Kelvin)

Hubbard model for the cuprates

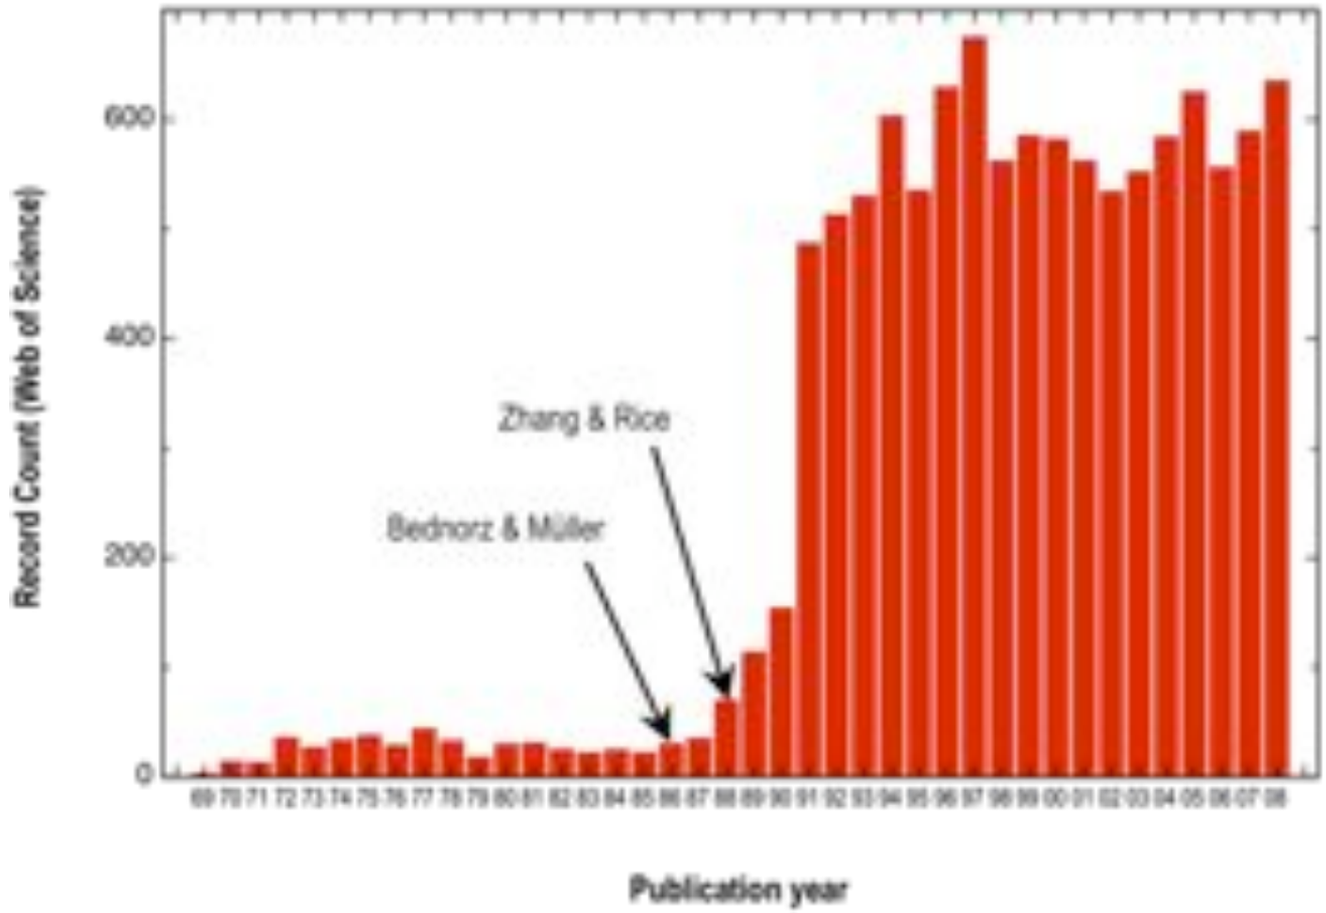

3. Parameter range relevant for superconducting cuprates

$U \approx 8t$

No simple solution!

Finite doping levels (0.05 – 0.25)

Typical values: U~10eV; t~0.9eV; J~0.2eV; (0.1eV ~ 10^3 Kelvin)

The challenge: a (quantum) multi-scale problem

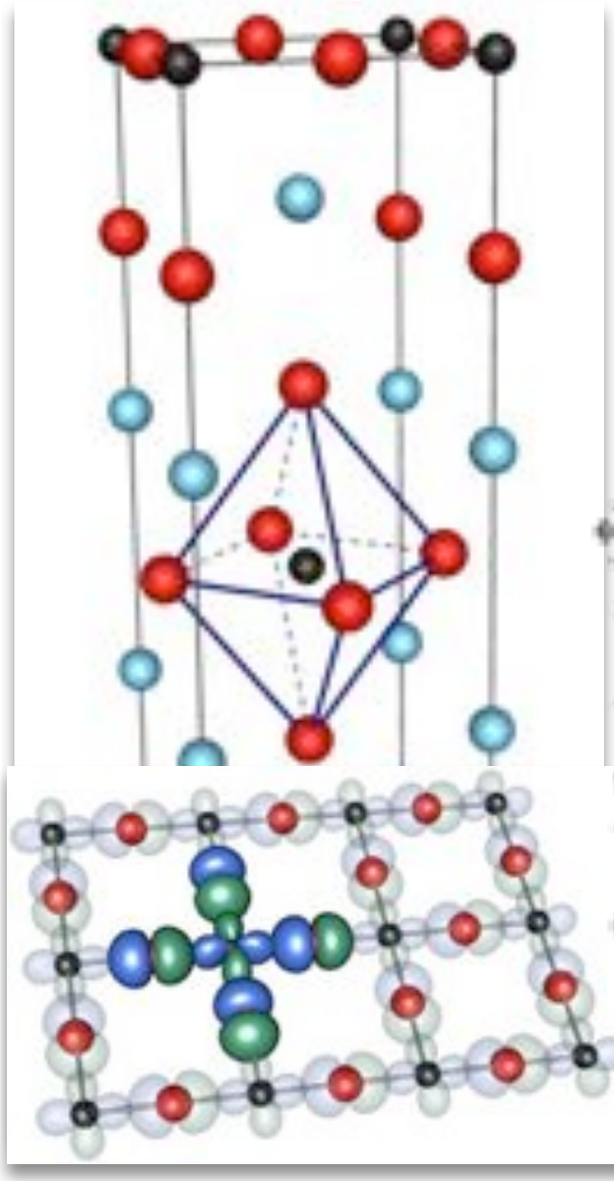

On-site Coulomb repulsion (~A)

Antiferromagnetic correlations / nano-scale gap fluctuations

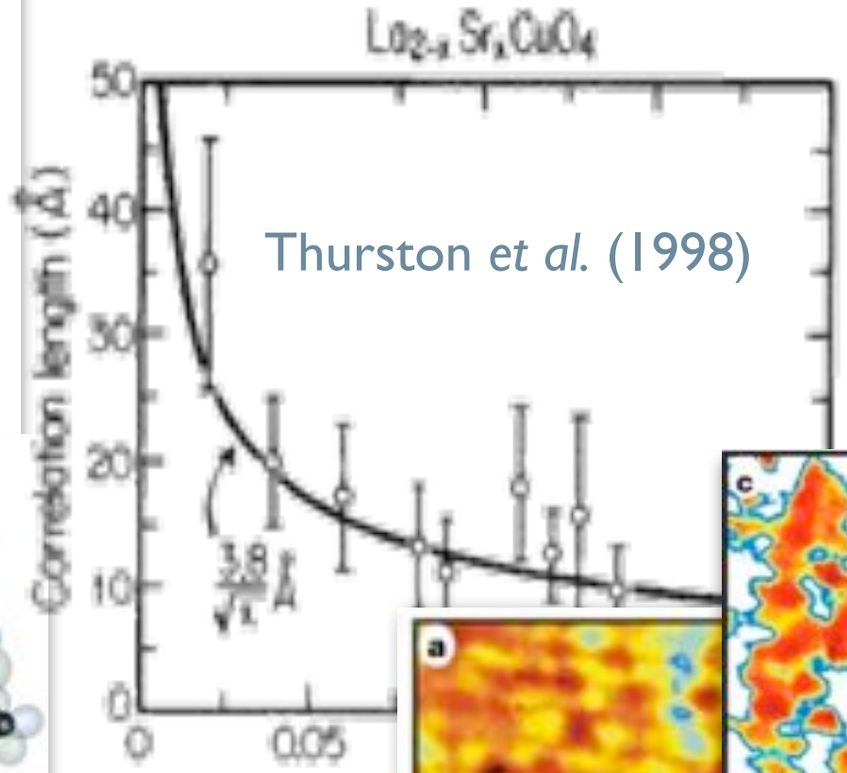

Thurston *et al.* (1998)

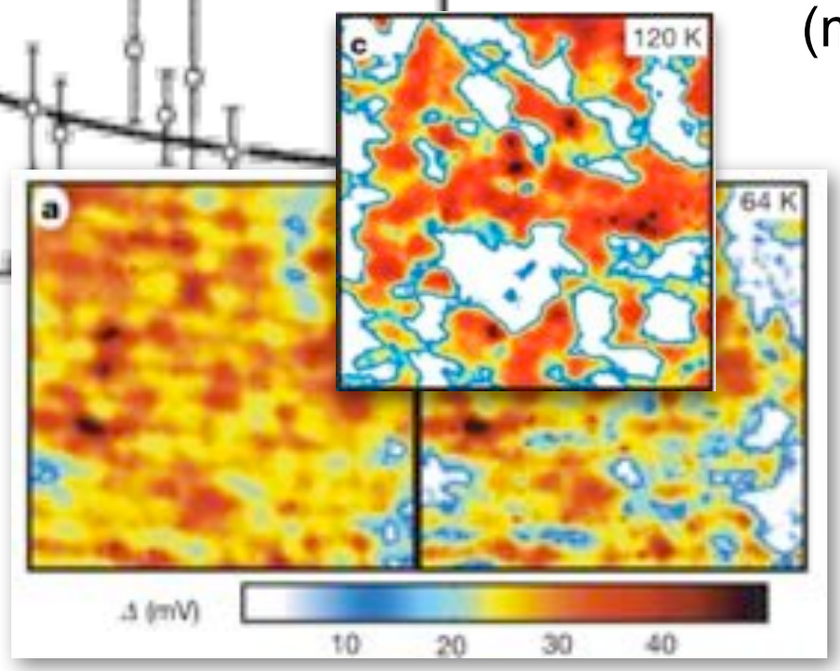

Gomes et al. (2007)

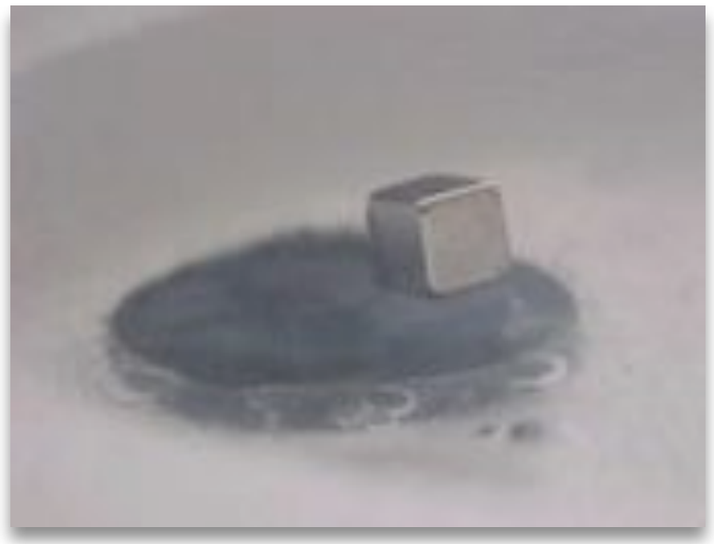

Superconductivity (macroscopic)

$N \sim 10^{23}$

complexity ~ 4^N

Quantum cluster theories

Maier *et al.*, Rev. Mod. Phys. '05

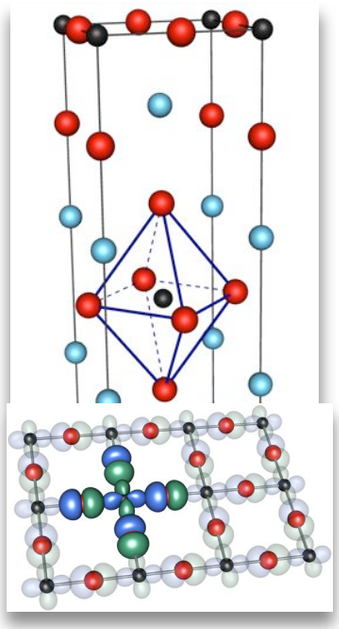

On-site Coulomb repulsion (~A)

Antiferromagnetic correlations / nano-scale gap fluctuations

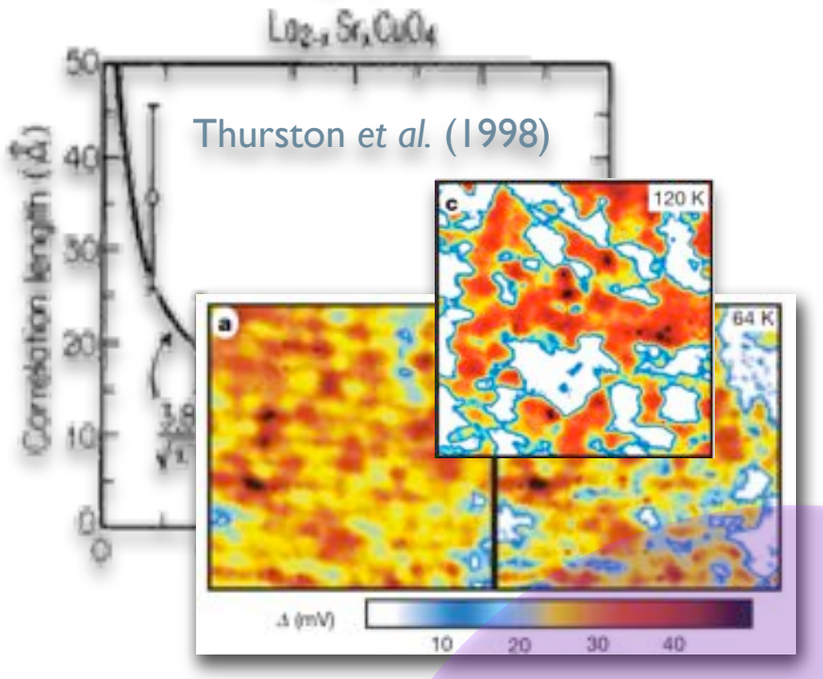

Thurston *et al.* (1998)

Gomes et al. (2007)

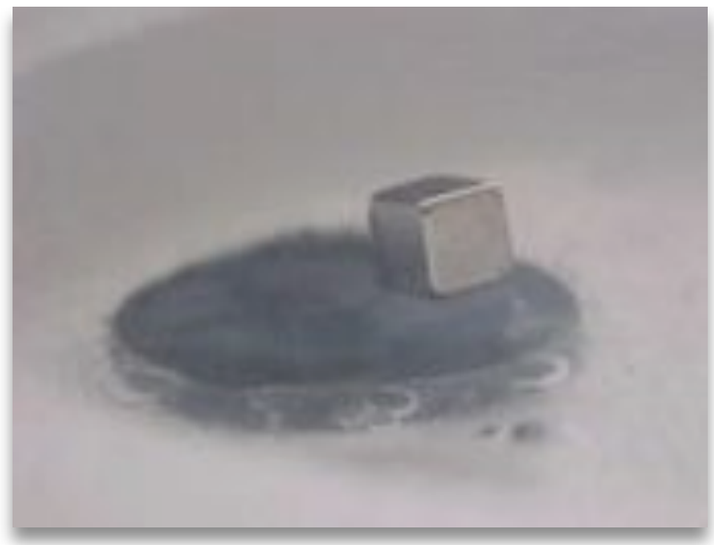

Superconductivity (macroscopic)

Explicitly treat correlations within a localized cluster

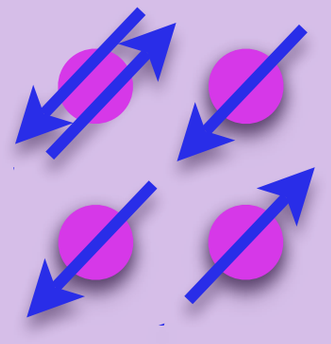

Treat macroscopic scales within mean-field

Coherently embed cluster into effective medium

QMC/DCA for a small cluster (4-site): phase diagram

M. Jarrell et al., EPL '01& var. other papers

4-site plaquette is the smallest cluster that describes *d*-wave pairing

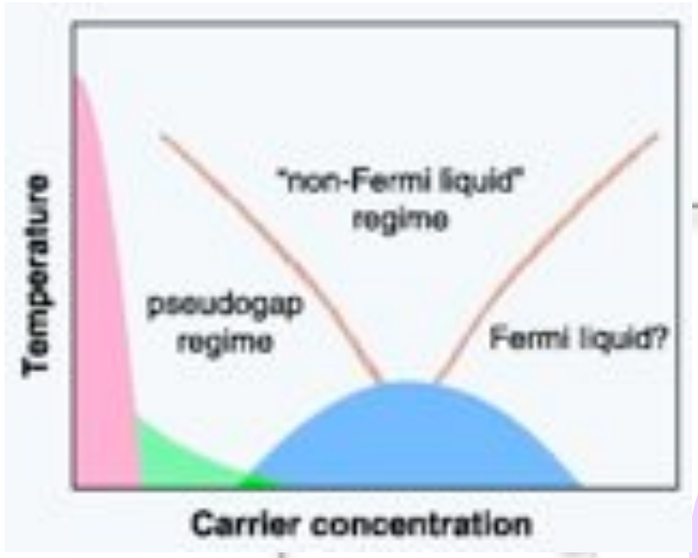

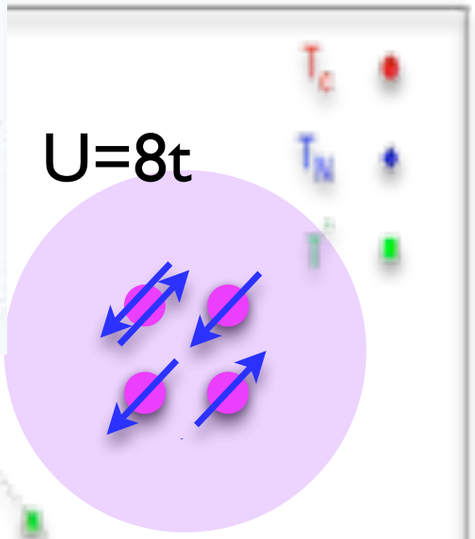

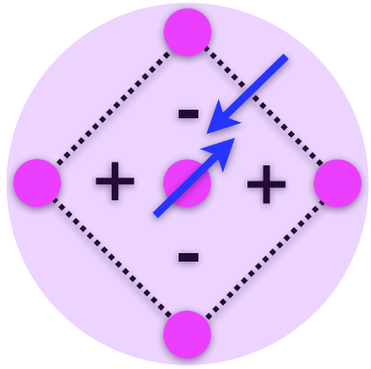

≡

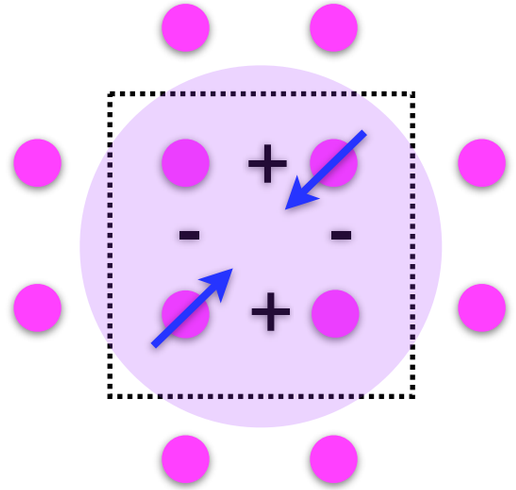

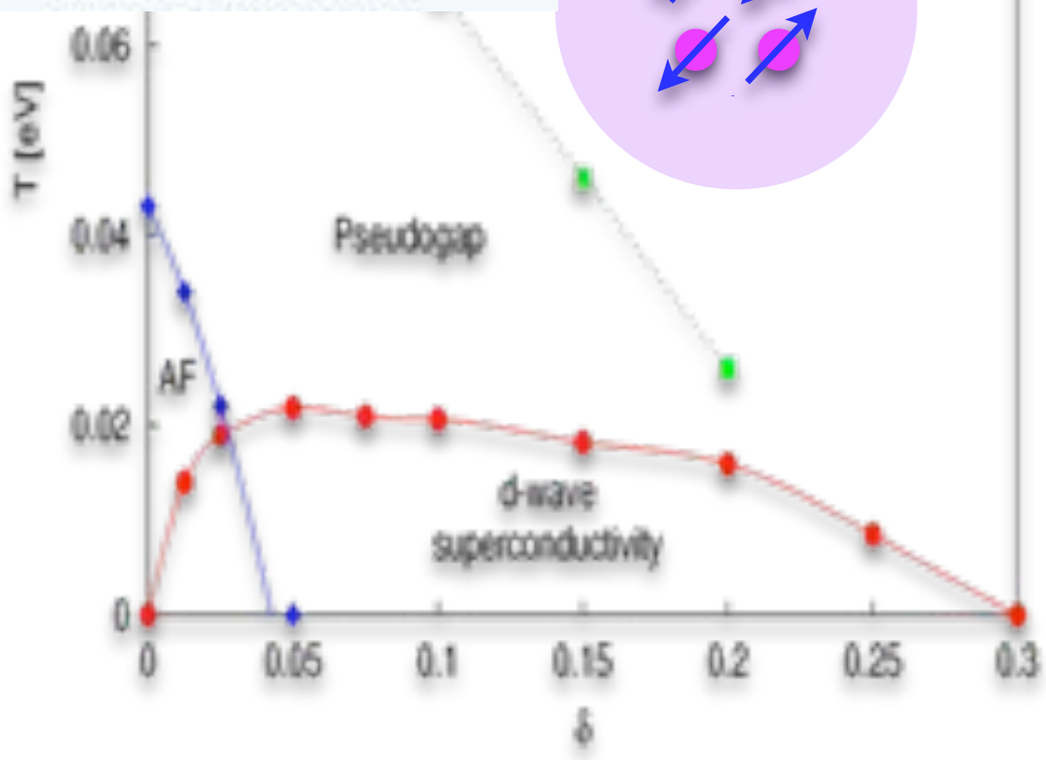

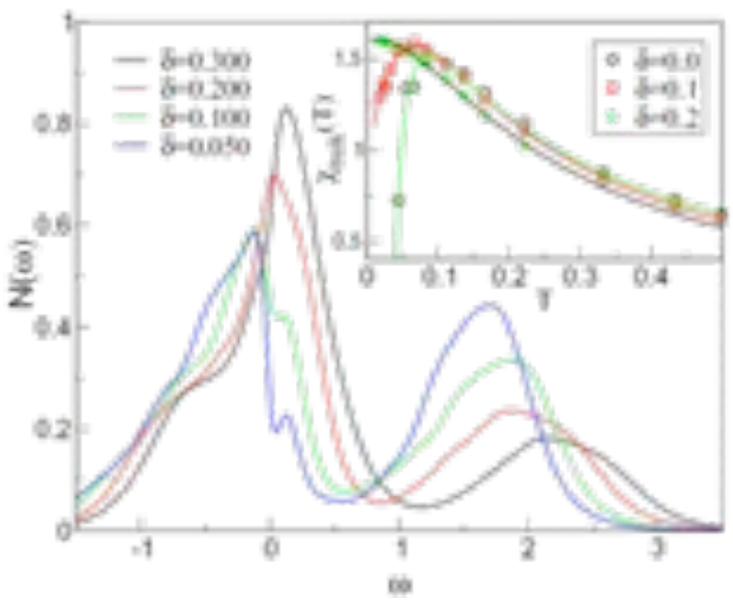

Nice, but ...

- Antiferromagnetism
  - Finite *T* order in 2D is in contradiction to Mermin-Wagner theorem
- Superconductivity
  - 4-site results represent mean-field results for *d*-wave order
  - No fluctuations on *d*-wave order parameter

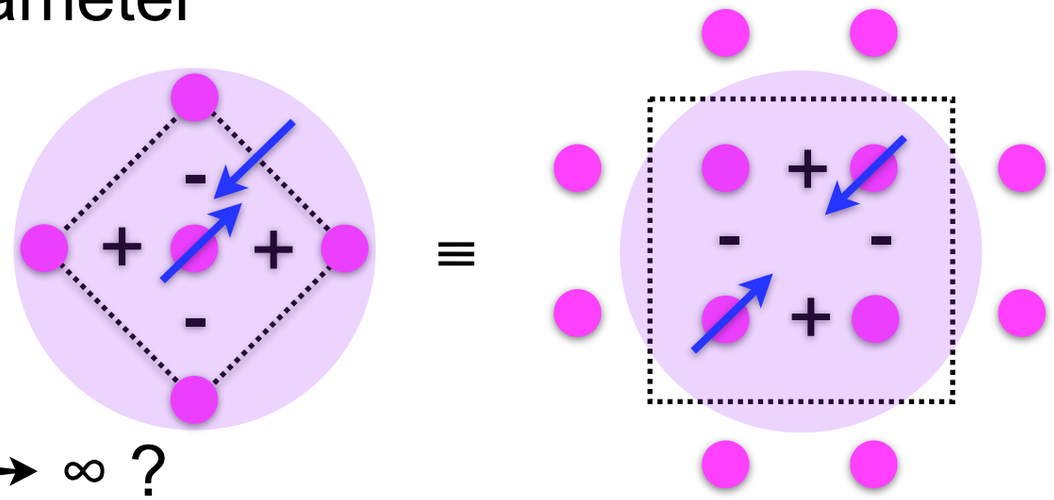

- Questions:
  - What happens in the exact limit, $N_c \rightarrow \infty$?
  - Kosterlitz-Thouless transition do *d*-wave superconductivity in large clusters?

Cluster size dependence of Néel temperature

No antiferromagnetic order in 2D
Néel temperature indeed vanishes logarithmically with cluster size
(Mermin Wagner Theorem satisfied)

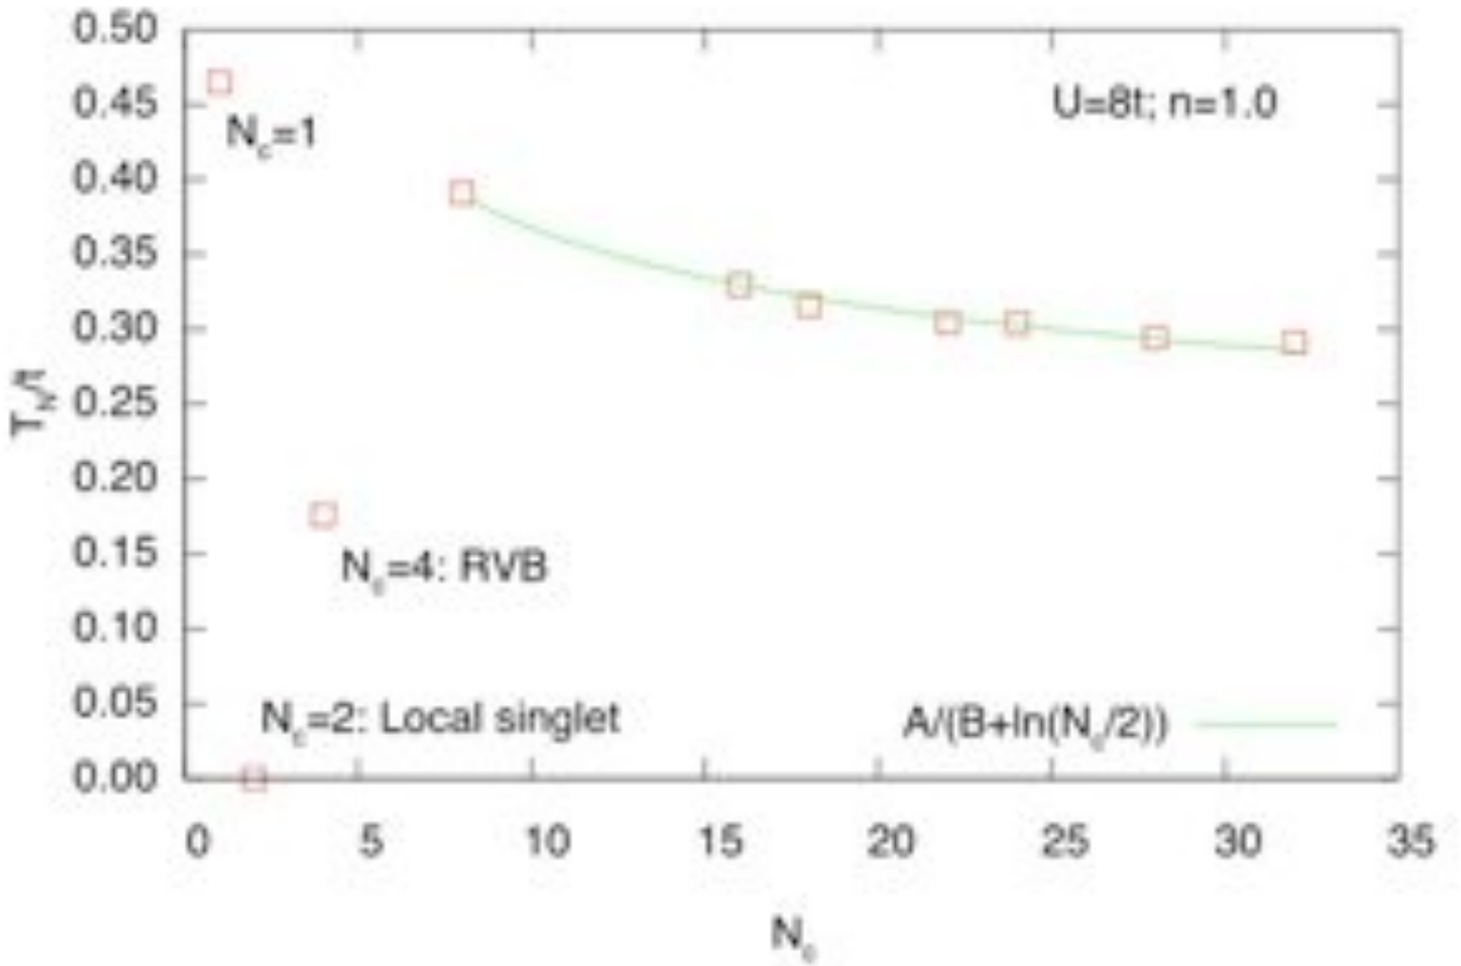

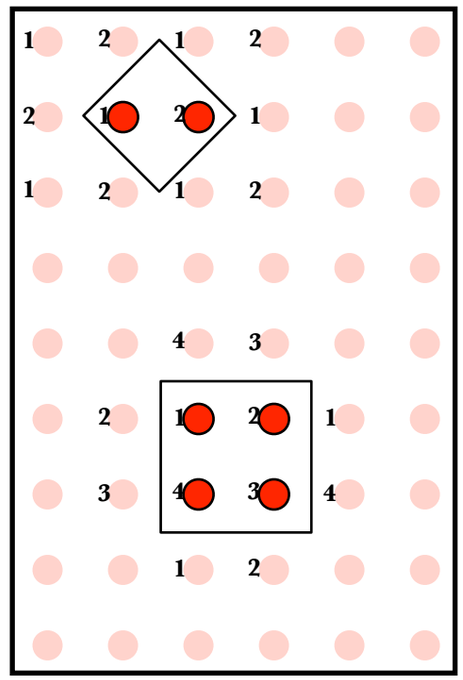

Nc=2:
1 neighbor

Nc=4:
2 neighbors

Maier et al., Phys. Rev. Lett. **95**, 237001 (2005)

Simulate larger clusters: Computational tour de force in 2004/2005 on Cray X1E @ NCCS

- Betts *et al.*, for 2D Heisenberg model: (Betts, Can. J. Phys. '99)
  - Selection criteria: symmetry, squareness, # of neighbors in a given shell
- Generalized for *d*-wave pairing in 2D Hubbard model:
  - Count number of neighboring independent 4-site *d*-wave plaquettes

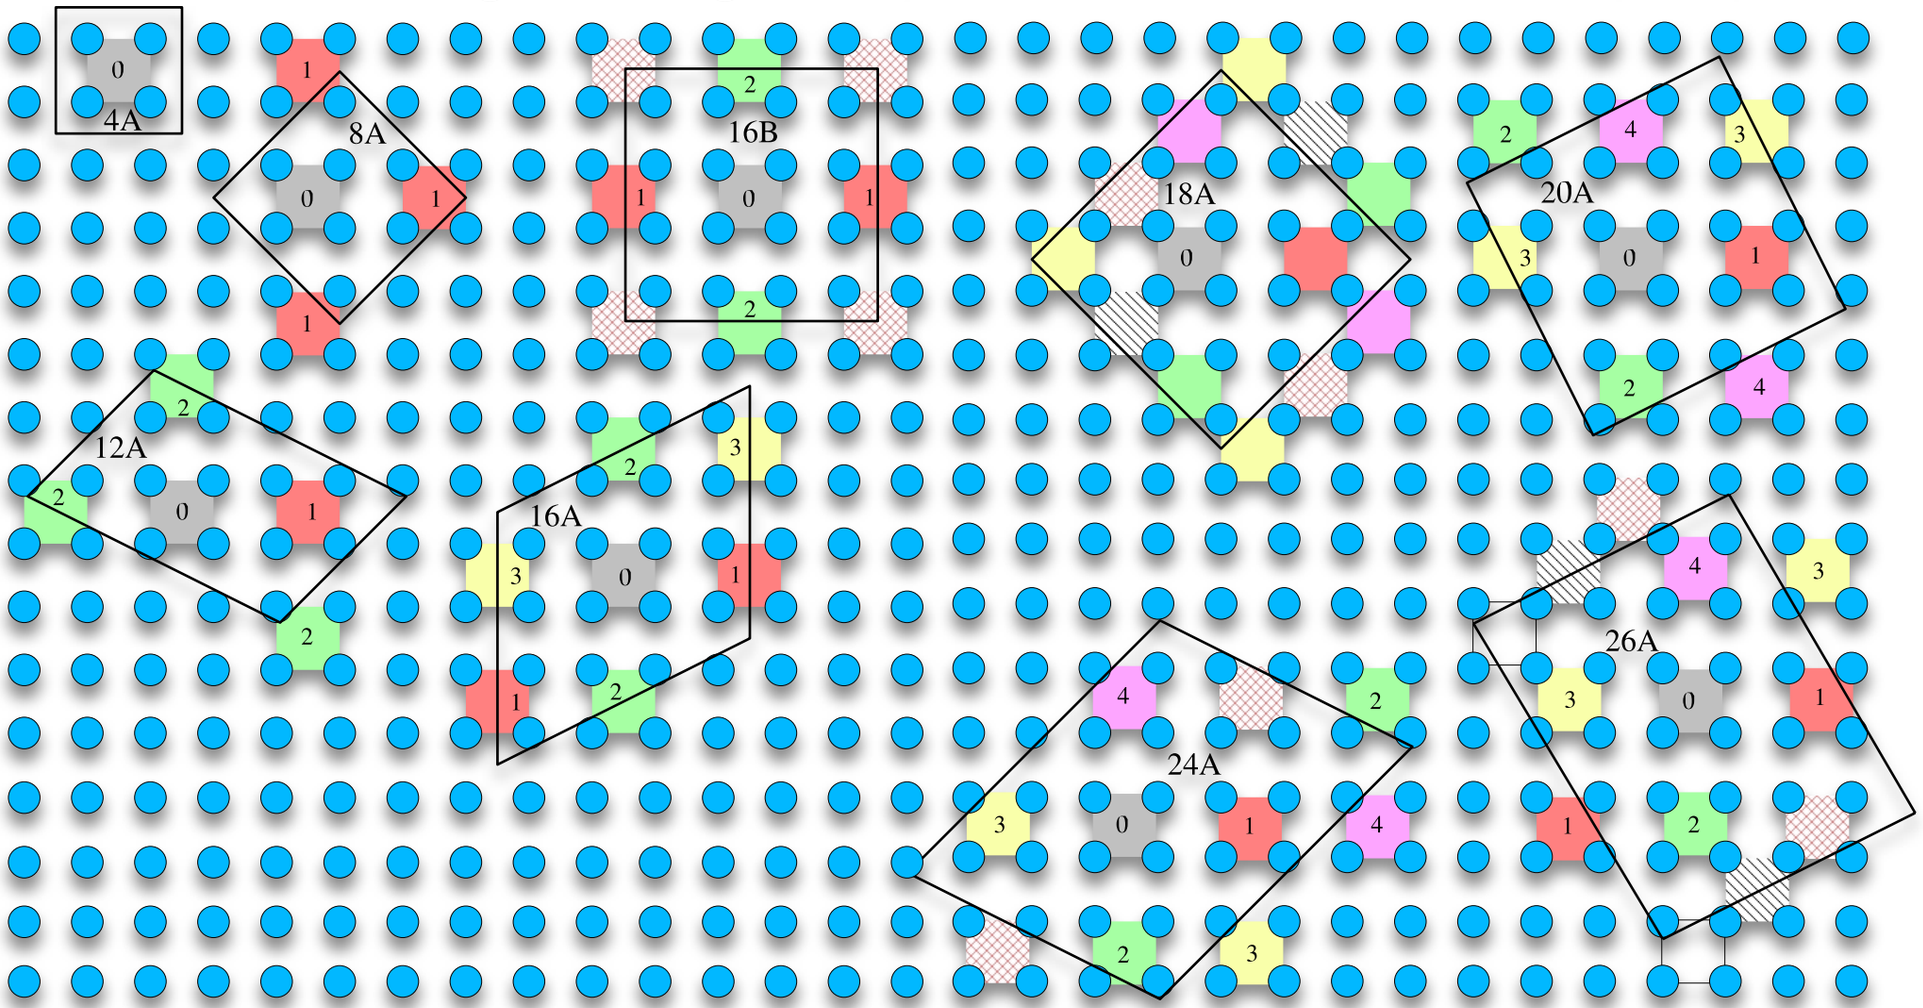

Superconducting transition as function of cluster size: Study divergence of pair-field susceptibility

Measure the pair-field susceptibility $P_d = \int_0^\beta d\tau \langle \Delta_d(\tau)\Delta_d^\dagger(0)\rangle$

$$\Delta_d^\dagger = \frac{1}{2\sqrt{N}} \sum_{l,\delta} (-1)^\delta c_{l\uparrow}^\dagger c_{l+\delta\downarrow}^\dagger$$

Inverse pair-field susceptiblity

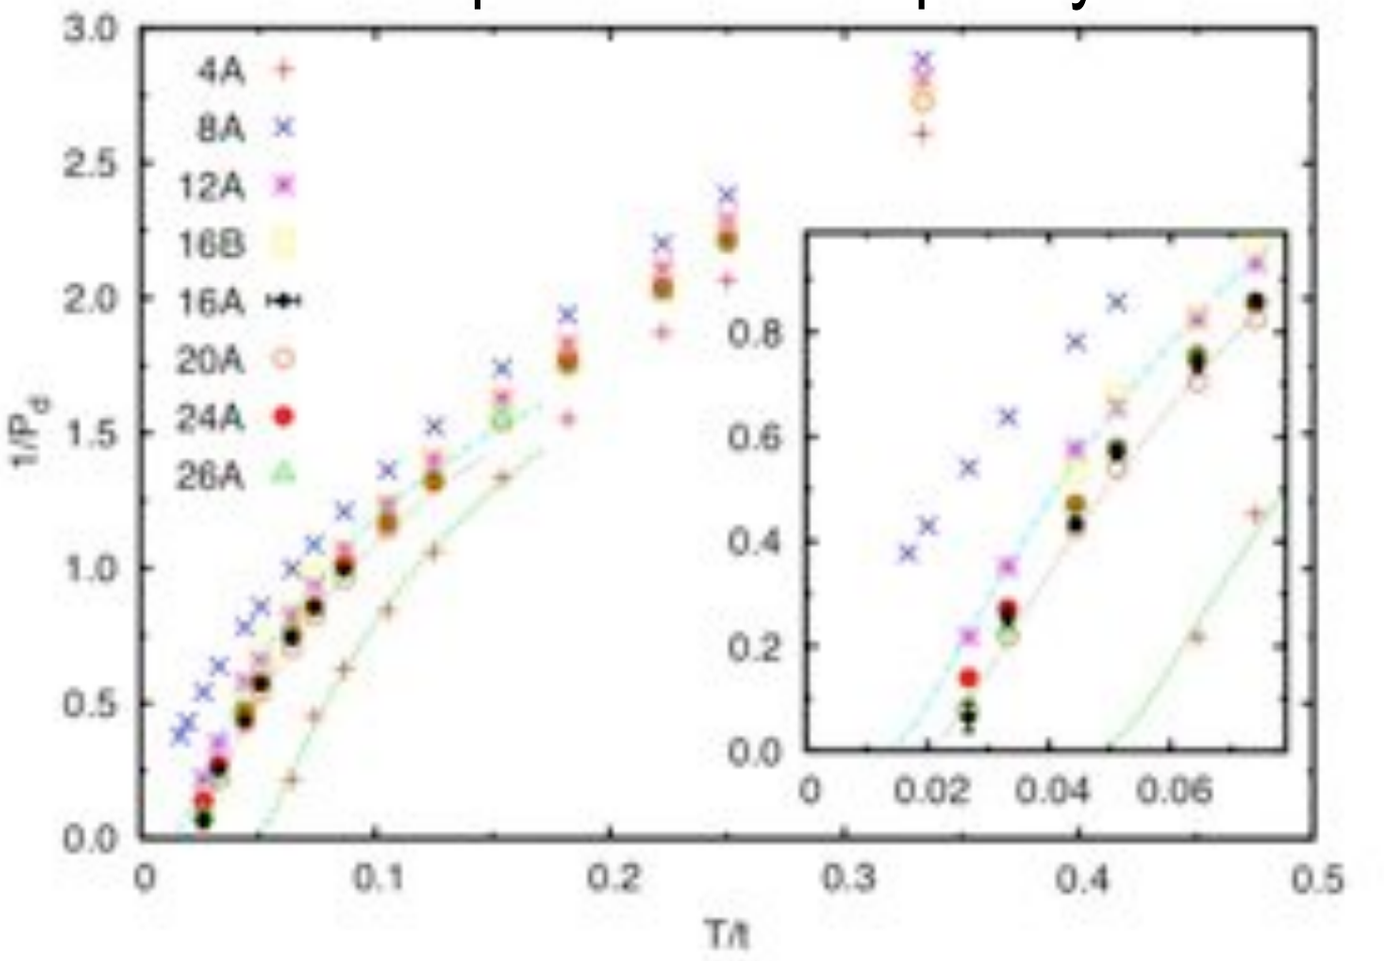

Maier, et al., Phys. Rev. Lett. **95**, 237001 (2005)

Cluster	Z_d
0(MF)	
8A	1
12A	2
16B	2
16A	3
20A	4
24A	4
26A	4

$T_c \approx 0.025t$

Second neighbor shell difficult due to QMC sign problem

Systematic solution and analysis of the pairing mechanism in the 2D Hubbard Model

- First systematic solution demonstrates existence of a superconducting transition in 2D Hubbard model

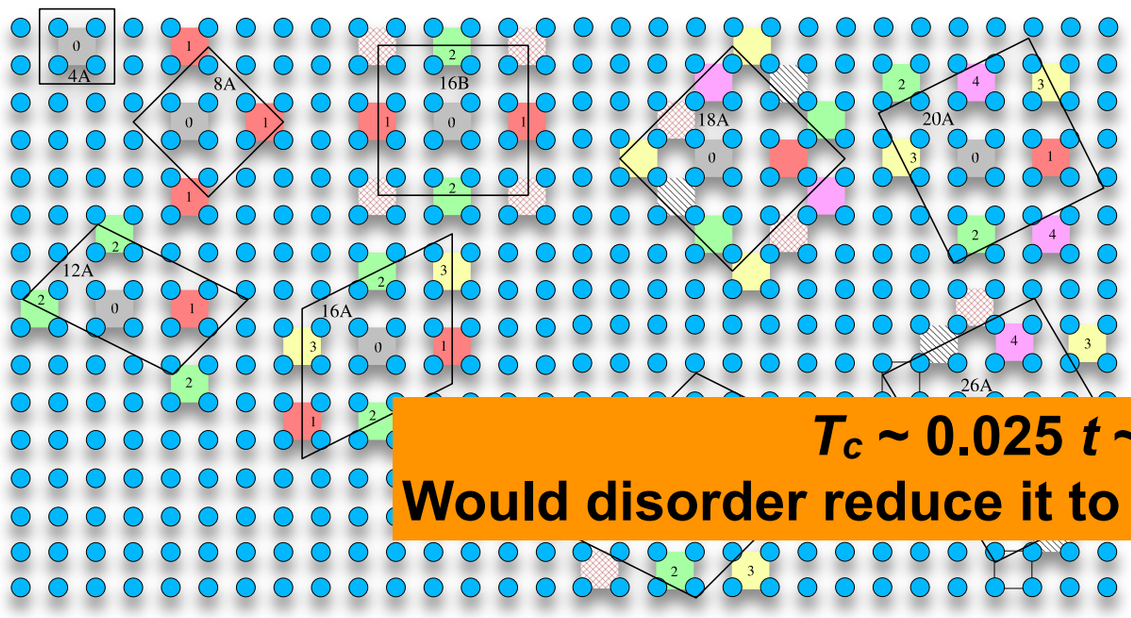

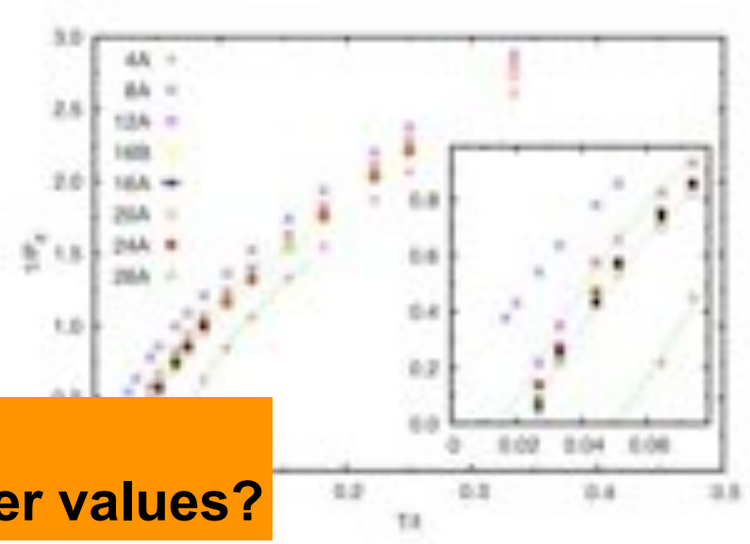

T_c ~ 0.025 t ~ 100 K
Would disorder reduce it to much smaller values?

- Study the mechanism responsible for pairing in the model
  - Analyze the particle-particle vertex
  - Pairing is mediated by spin fluctuations
    Maier, et al., Phys. Rev. Lett. **96** 47005 (2006)

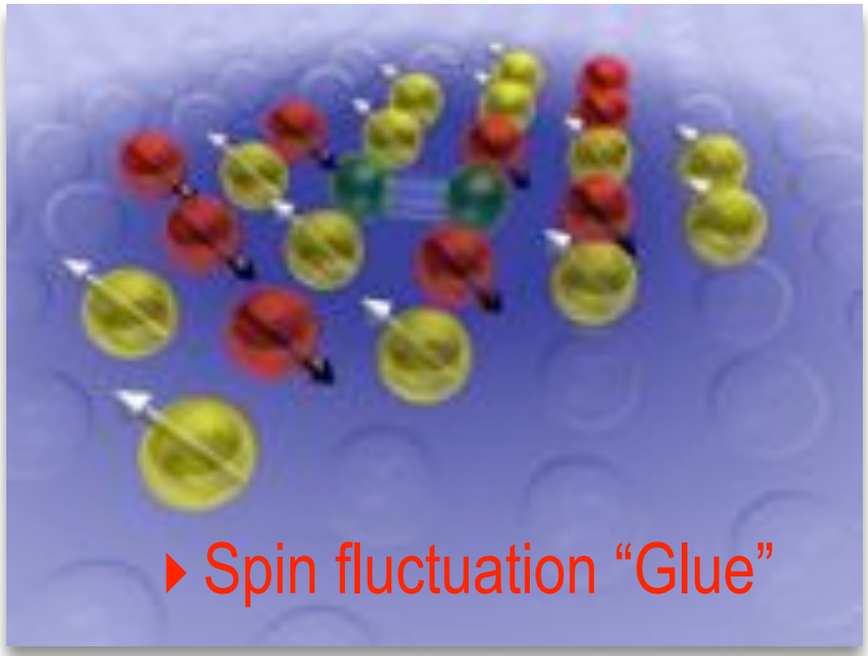

Moving toward a resolution of the debate over the pairing mechanism in the 2D Hubbard model

- "We have a mammoth (*U*) and an elephant (*J*) in our refrigerator - do we care much if there is also a mouse?"
  - P.W. Anderson, Science **316**, 1705 (2007)
  - see also www.sciencemag.org/cgi/eletters/316/5832/1705 "Scalapino is not a glue sniffer"
- Relative importance of resonant valence bond and spin-fluctuation mechanism
  - Maier et al., Phys. Rev. Lett. **100** 237001 (2008)

Fraction of superconducting gap arising from frequencies ≤ Ω

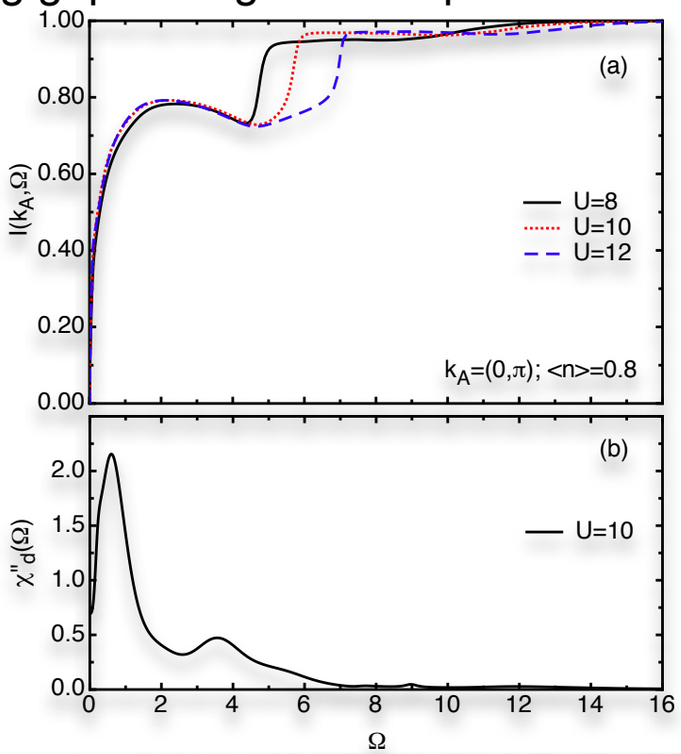

Both retarded spin-fluctuations and non-retarded exchange interaction J contribute to the pairing interaction

Dominant contribution comes from spin-fluctuations!

Nanoscale stripe modulations enhance superconducting transition temperature

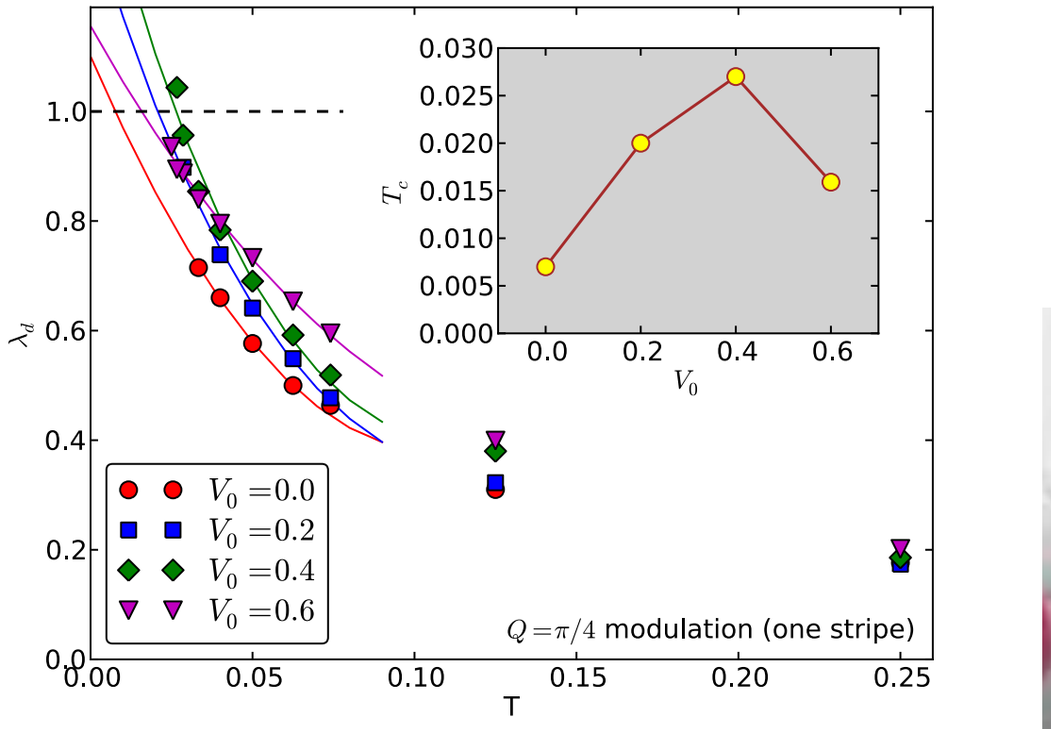

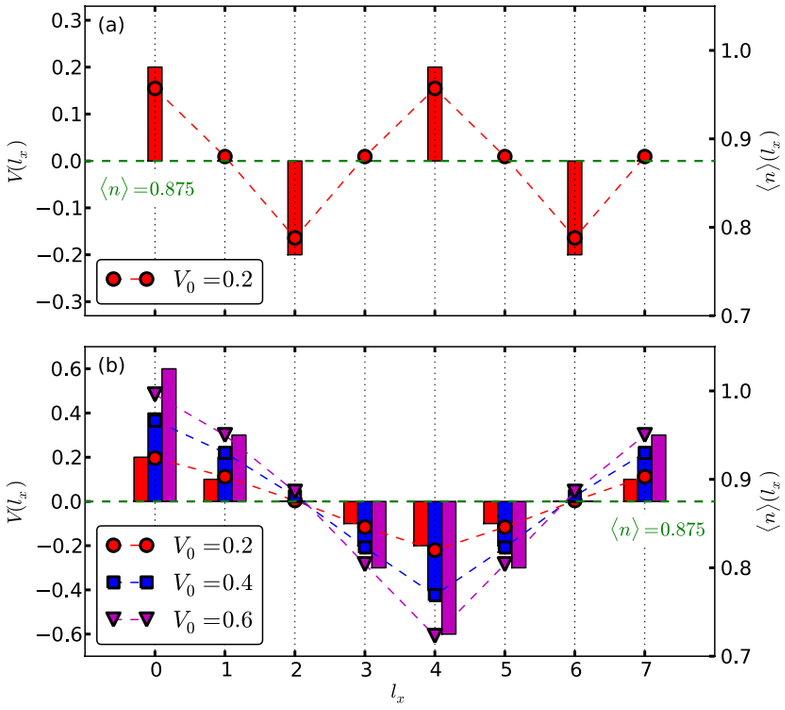

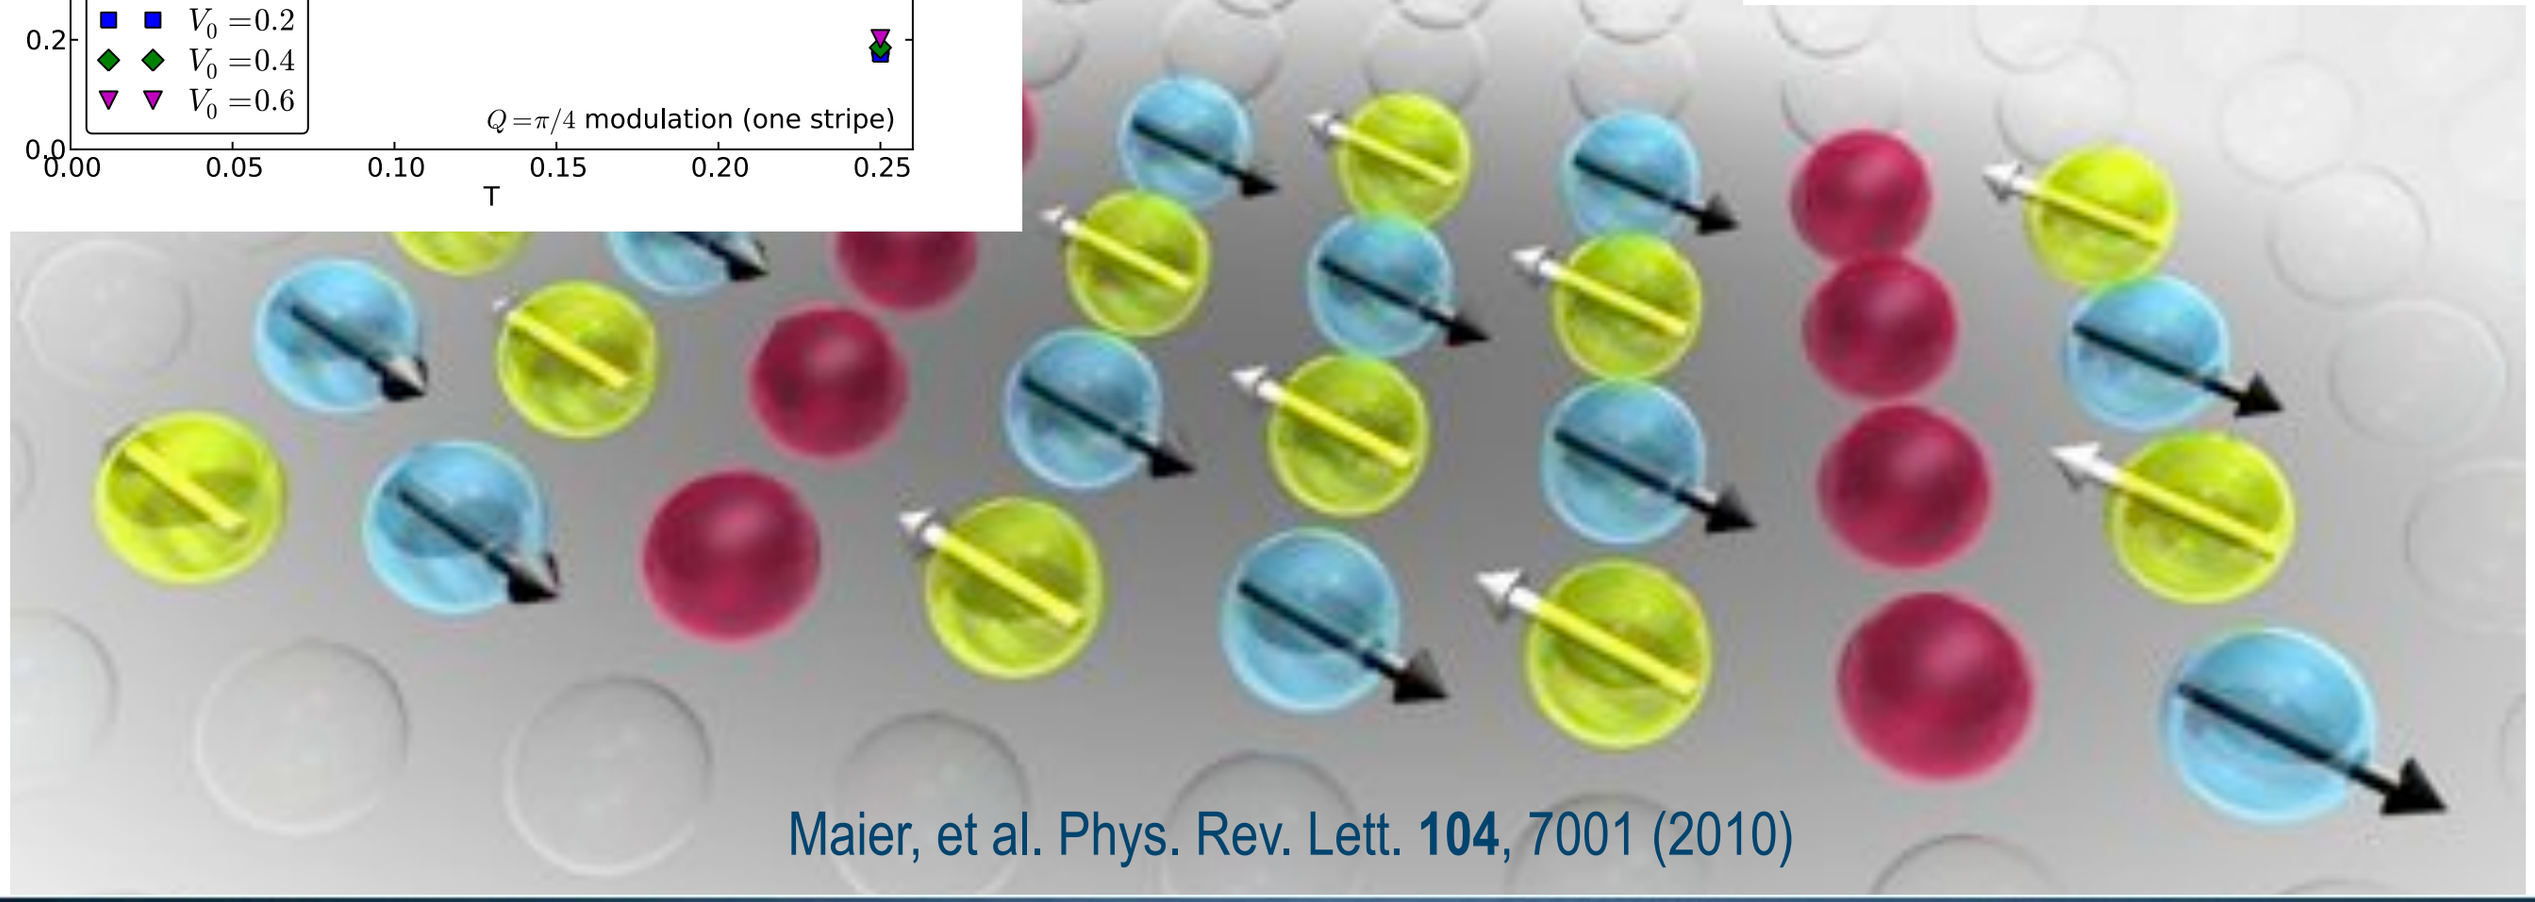

Maier, et al. Phys. Rev. Lett. **104**, 7001 (2010)

Green's functions in quantum many-body theory

Noninteracting Hamiltonian & $H_0 = \left[-\frac{1}{2}\nabla^2 + V(\vec{r})\right]$

Green's function $\left[i\frac{\partial}{\partial t} - H_0\right] G_0(\vec{r}, t, \vec{r}', t') = \delta(\vec{r} - \vec{r}')\delta(t - t)$

Fourier transform & analytic continuation: $z^{\pm} = \omega \pm i\epsilon$ $G_0^{\pm}(\vec{r}, z) = \left[z^{\pm} - H_0\right]^{-1}$

Hubbard Hamiltonian $H = -t \sum_{<ij>,\sigma} c_{i\sigma}^{\dagger} c_{j\sigma} + U \sum_i n_{i\uparrow} n_{i\downarrow}$ $n_{i\sigma} = c_{i\sigma}^{\dagger} c_{i\sigma}$

Hide symmetry in algebraic properties of field operators $c_{i\sigma} c_{j\sigma'} + c_{j\sigma'} c_{i\sigma} = 0$

$c_{i\sigma} c_{j\sigma'}^{\dagger} + c_{j\sigma'}^{\dagger} c_{i\sigma} = \delta_{ij}\delta_{\sigma\sigma'}$

Green's function $G_{\sigma}(r_i, \tau; r_j, \tau') = -\left\langle \mathcal{T} c_{i\sigma}(\tau) c_{j\sigma}^{\dagger}(\tau') \right\rangle$

Spectral representation $G_0(k, z) = [z - \epsilon_0(k)]^{-1}$

$$G(k, z) = [z - \epsilon_0(k) - \Sigma(k, z)]^{-1}$$

Sketch of the Dynamical Cluster Approximation

Bulk lattice

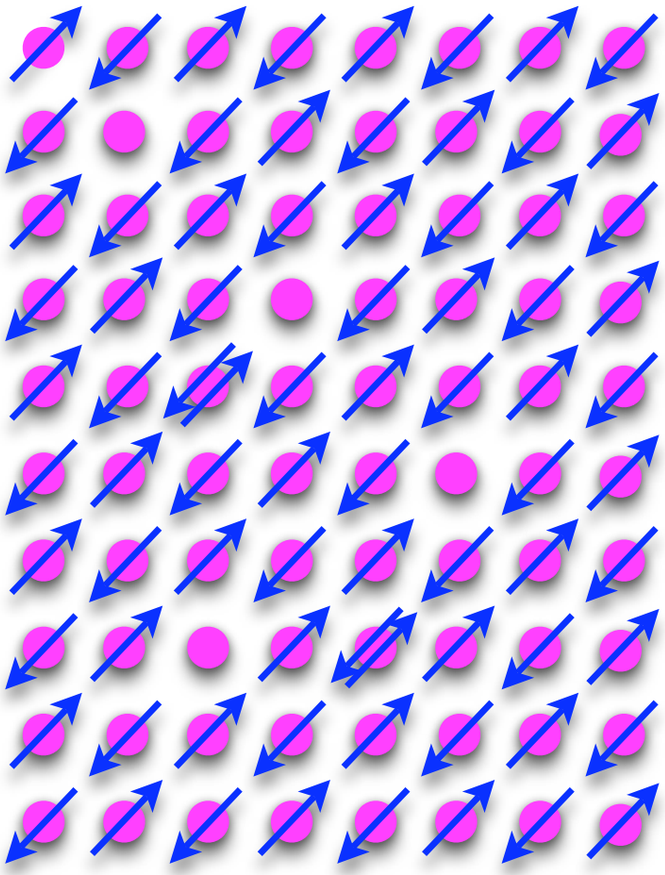

Size N_c clusters

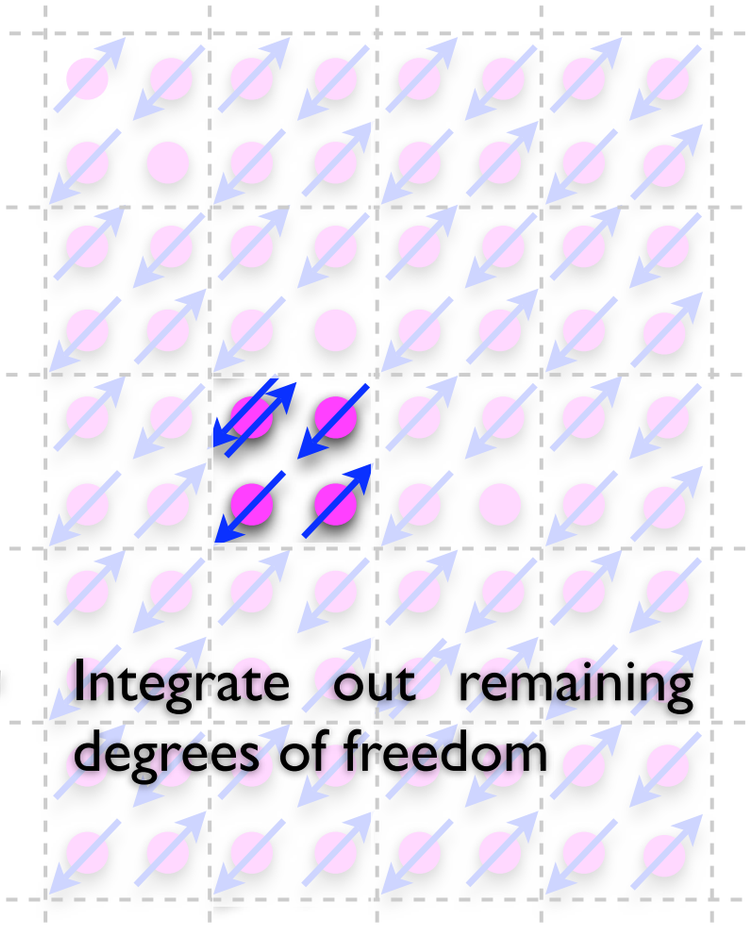

Integrate out remaining degrees of freedom

Reciprocal space

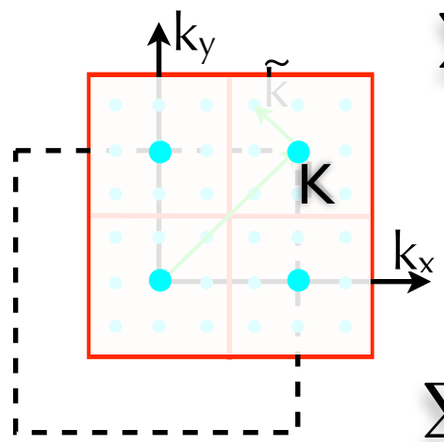

$\Sigma(z, k)$

$\Sigma(z, K)$

DCA

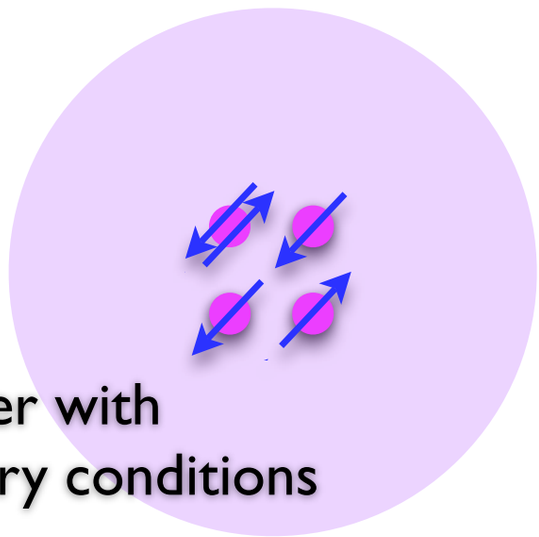

Embedded cluster with periodic boundary conditions

Solve many-body problem with quantum Monte Carole on cluster

➢Essential assumption: Correlations are short ranged

DCA method: self-consistently determine the "effective" medium

Maier *et al.*, Rev. Mod. Phys. '05

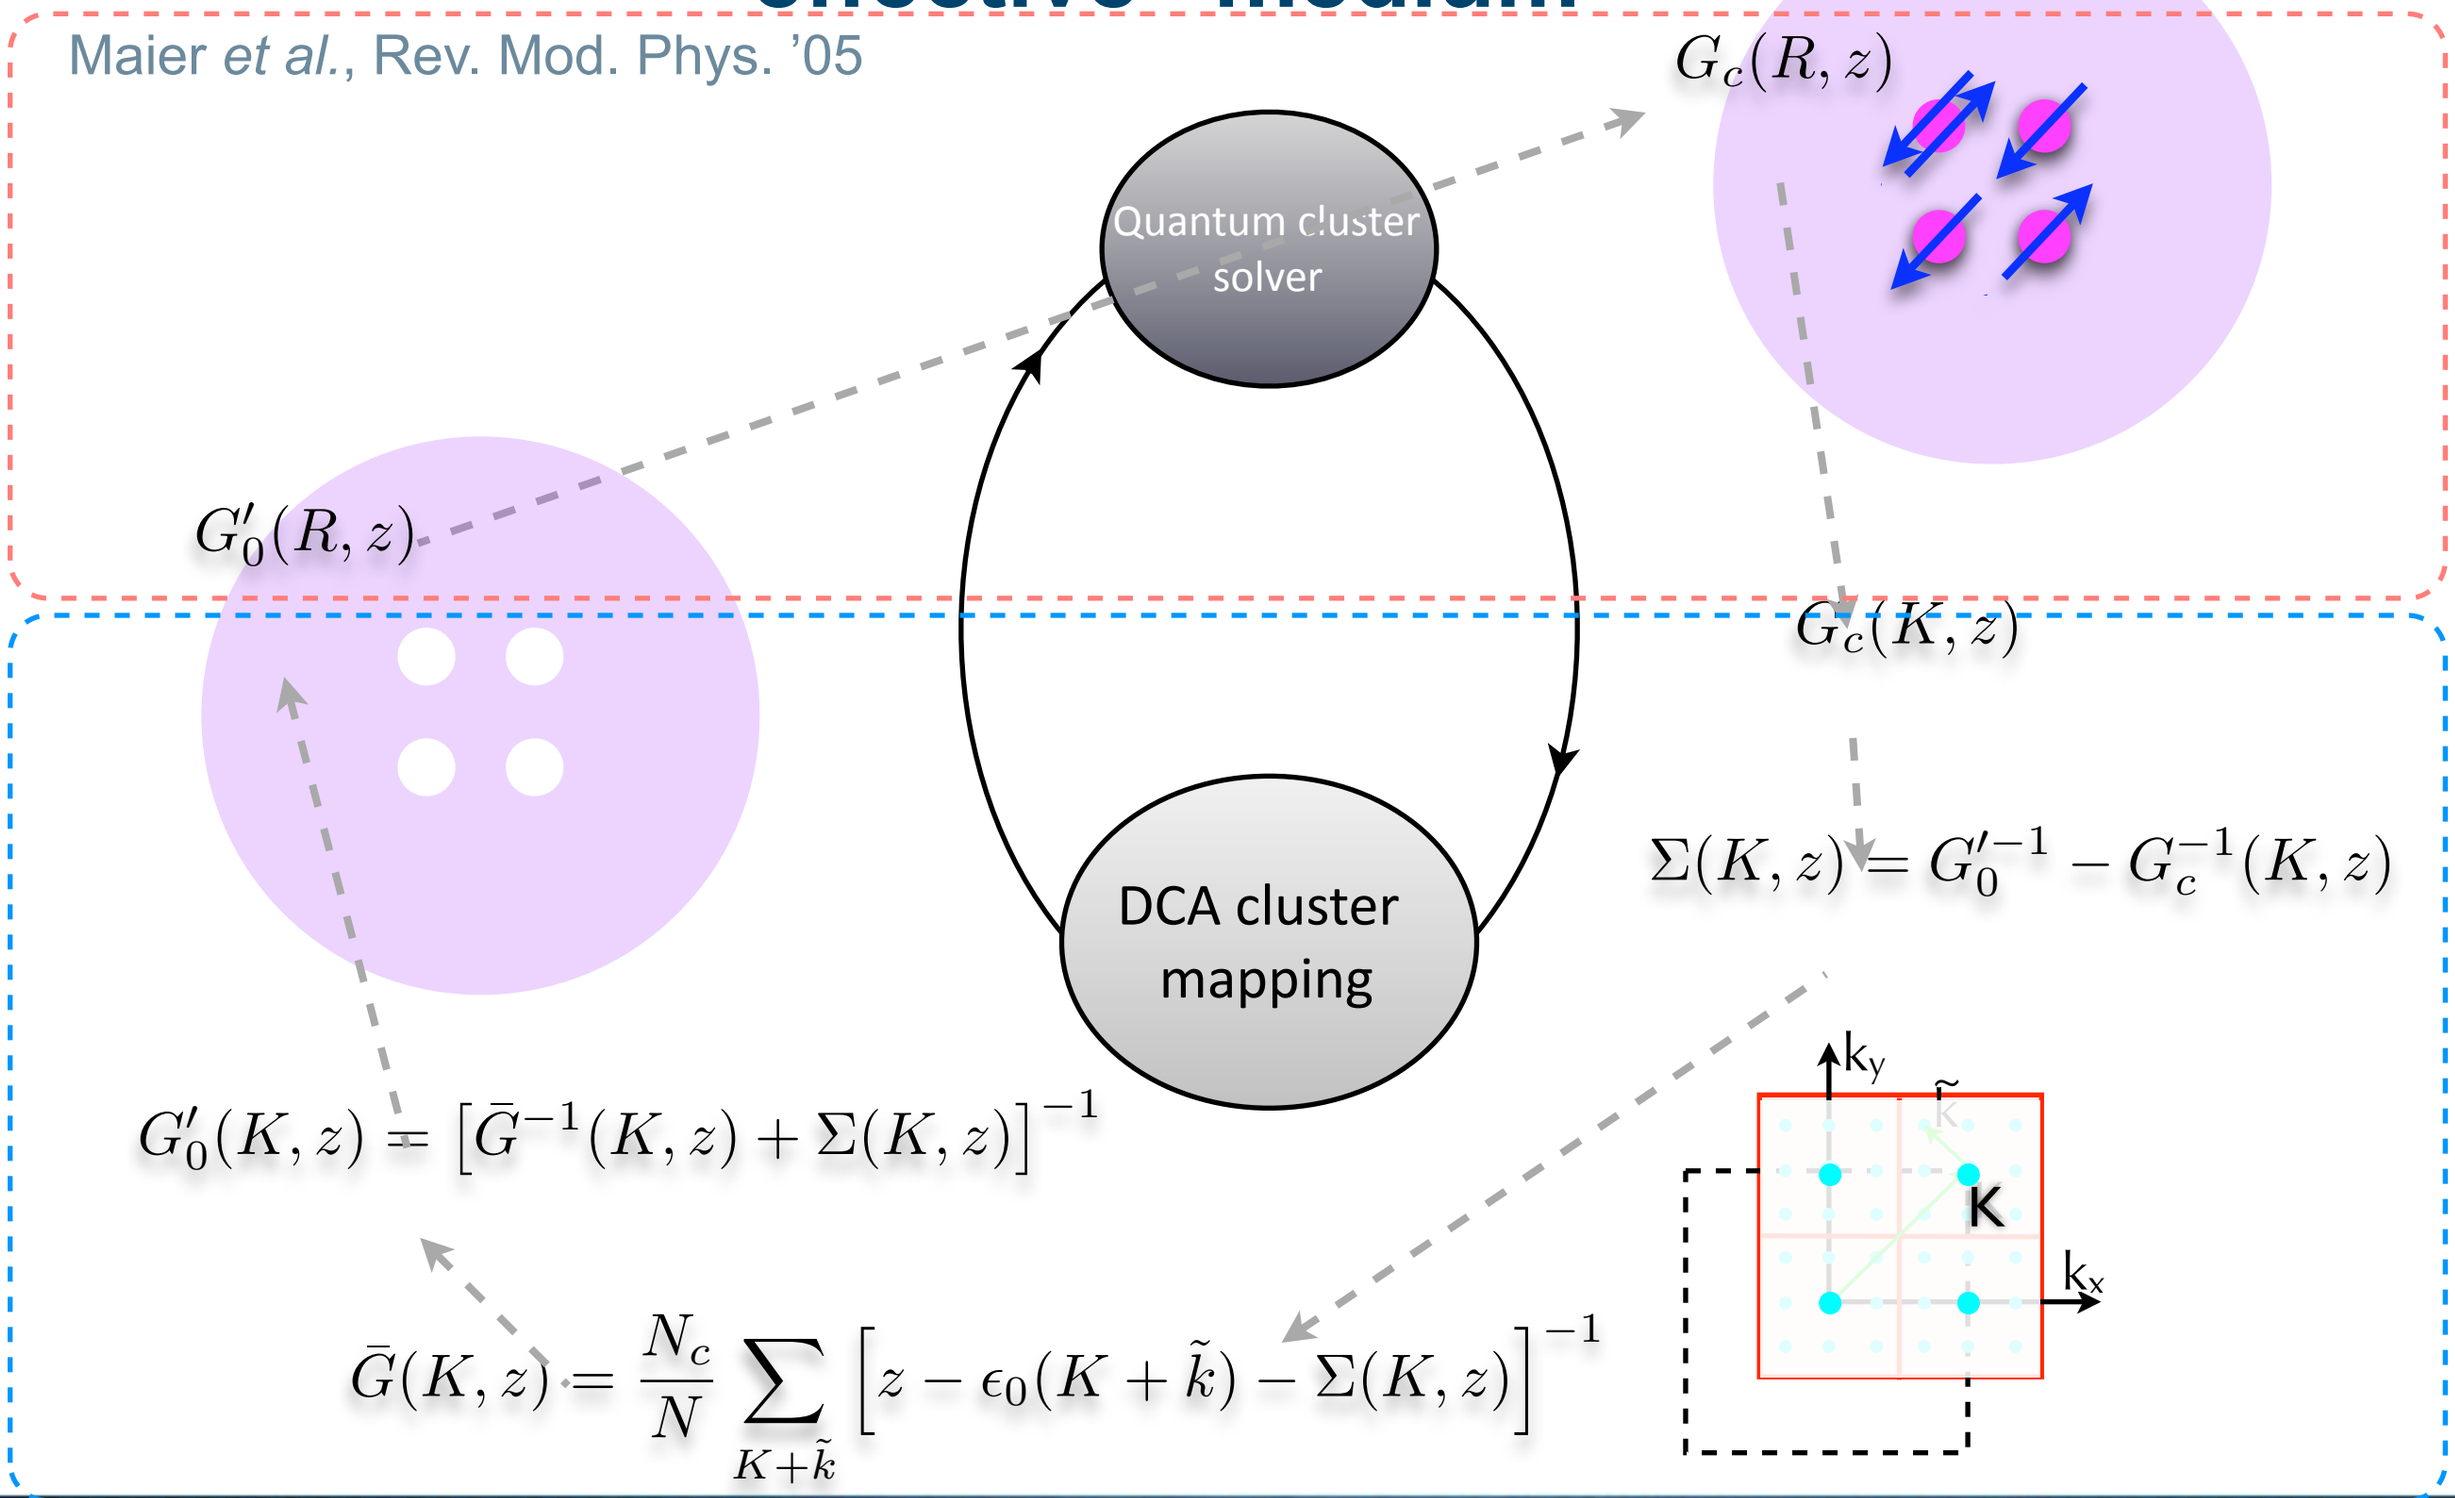

Introducing disorder into the calculations

$$H^{(\nu)} = -t \sum_{<ij>,\sigma} c^{\dagger}_{i\sigma} c_{j\sigma} + \sum_{i} U_i^{(\nu)} n_{i\uparrow} n_{i\downarrow} + \sum_{i,\sigma} V_i^{(\nu)} n_{i\sigma}$$

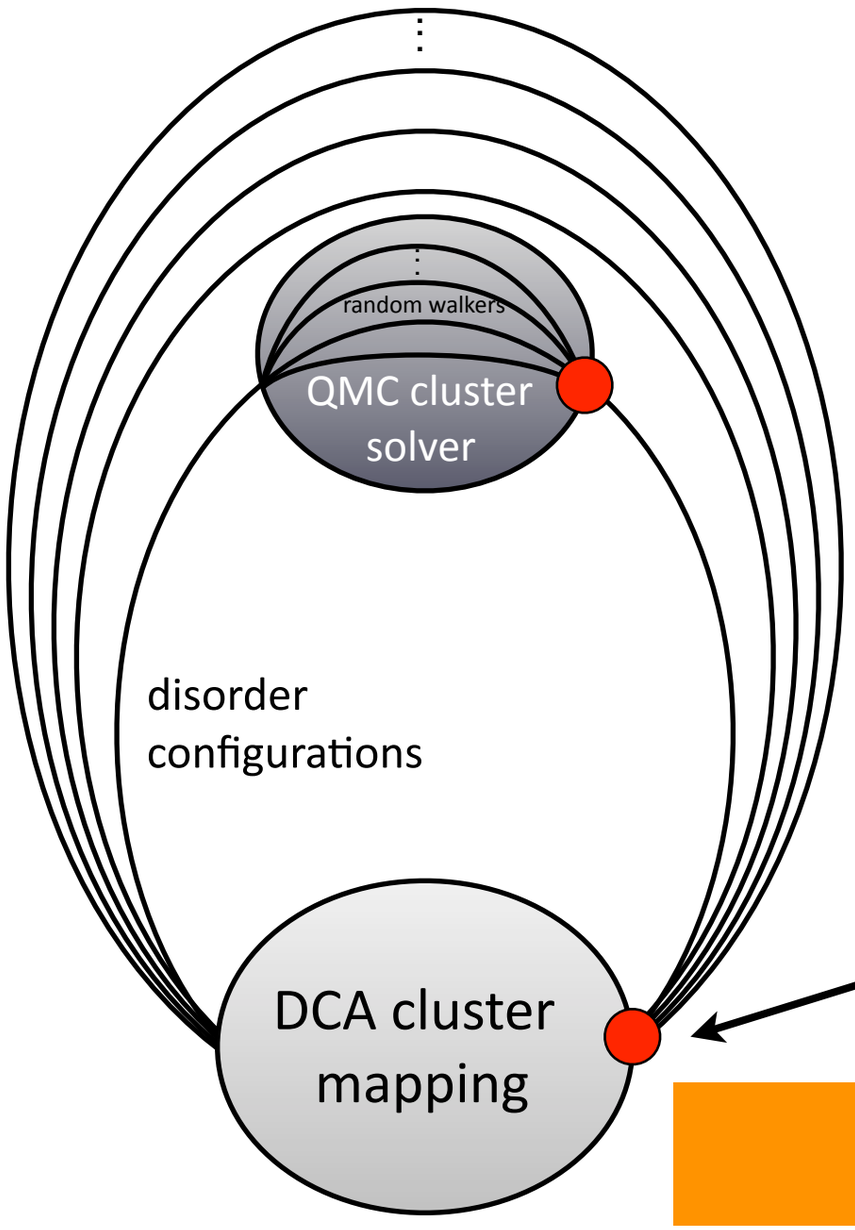

$$V_i \in \{V, 0\}$$
$$U_i \in \{U + \Delta U, U - \Delta U\}$$
$$N_c = 16 \rightarrow N_d = 2^{16}$$

required communication

$$G_c(X_i - X_j, z) = \frac{1}{N_c} \sum_{\nu=1}^{N_d} G_c^{\nu}(X_i, X_j, z)$$

Computations with disorder averaging requires factor 100-1000 more performance

Structure of DCA++ code: generic programming

Category	Number	Lines of Code
Functions	23	170
Operators	29	562
Generic Classes	171	23,185
Regular Classes	34	2,005
Total		25,922

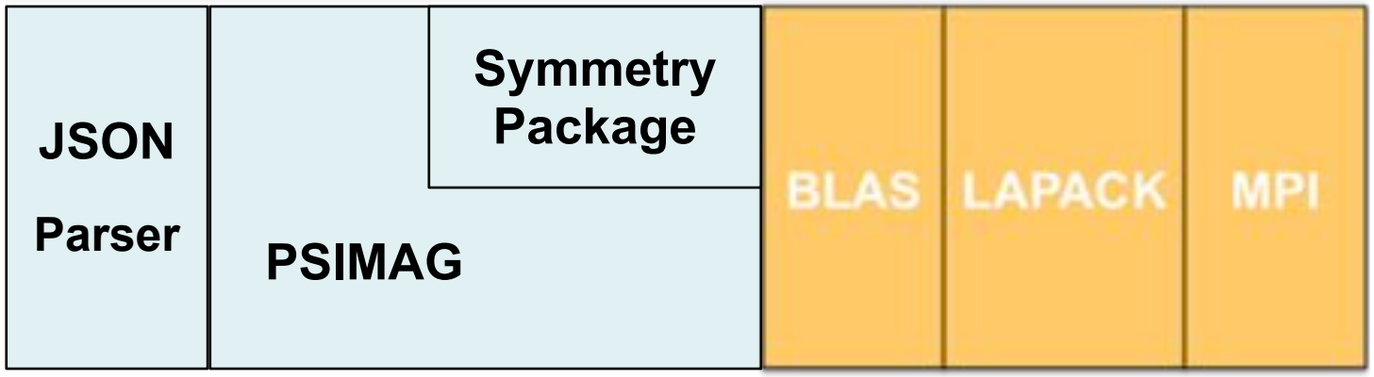

PsiMag Implementation philosophy:
Consider PsiMag as a systematic extension to the C++ Standard Template Library (STL) using as much as possible the **generic programming paradigm**

Hirsch-Fye Quantum Monte Carole (HF-QMC) for the quantum cluster solver

Hirsch & Fye, Phys. Rev. Lett. **56**, 2521 (1998)

Partition function & Metropolis Monte Carlo $Z = \int e^{-E[\mathbf{x}]/k_{\mathrm{B}}T} d\mathbf{x}$

Acceptance criterion for M-MC move: $\min\{1, e^{E[\mathbf{x}_k]-E[\mathbf{x}_{k+1}]}\}$

Partition function & HF-QMC: $Z \sim \sum_{s_i,l} \det[\mathbf{G}_c(s_i,l)^{-1}]$

N_c $N_l \approx 10^2$

Acceptance: $\min\{1, \det[\mathbf{G}_c(\{s_i,l\}_k)]/\det[\mathbf{G}_c(\{s_i,l\}_{k+1})]\}$

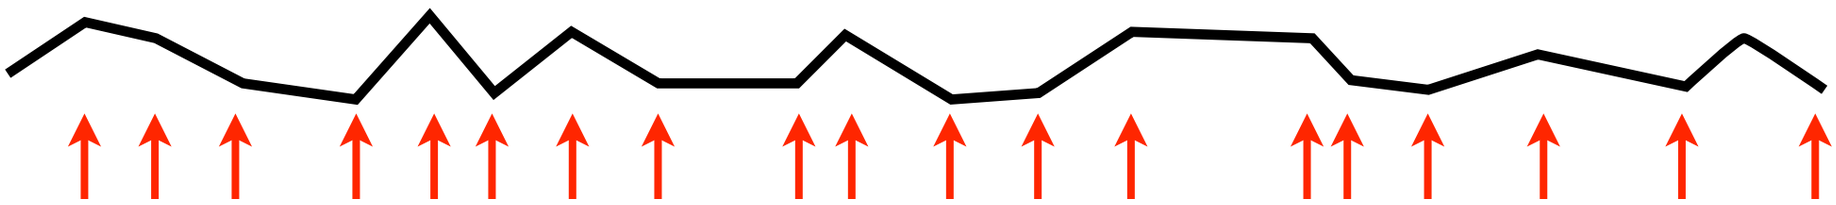

Update of accepted Green's function: matrix of dimensions $N_t \times N_t$

$$\mathbf{G}_c(\{s_i,l\}_{k+1}) = \mathbf{G}_c(\{s_i,l\}_k) + \mathbf{a}_k \times \mathbf{b}_k$$

$N_t = N_c \times N_l \approx 2000$

Take advantage of many-cores / shared L3 cash?

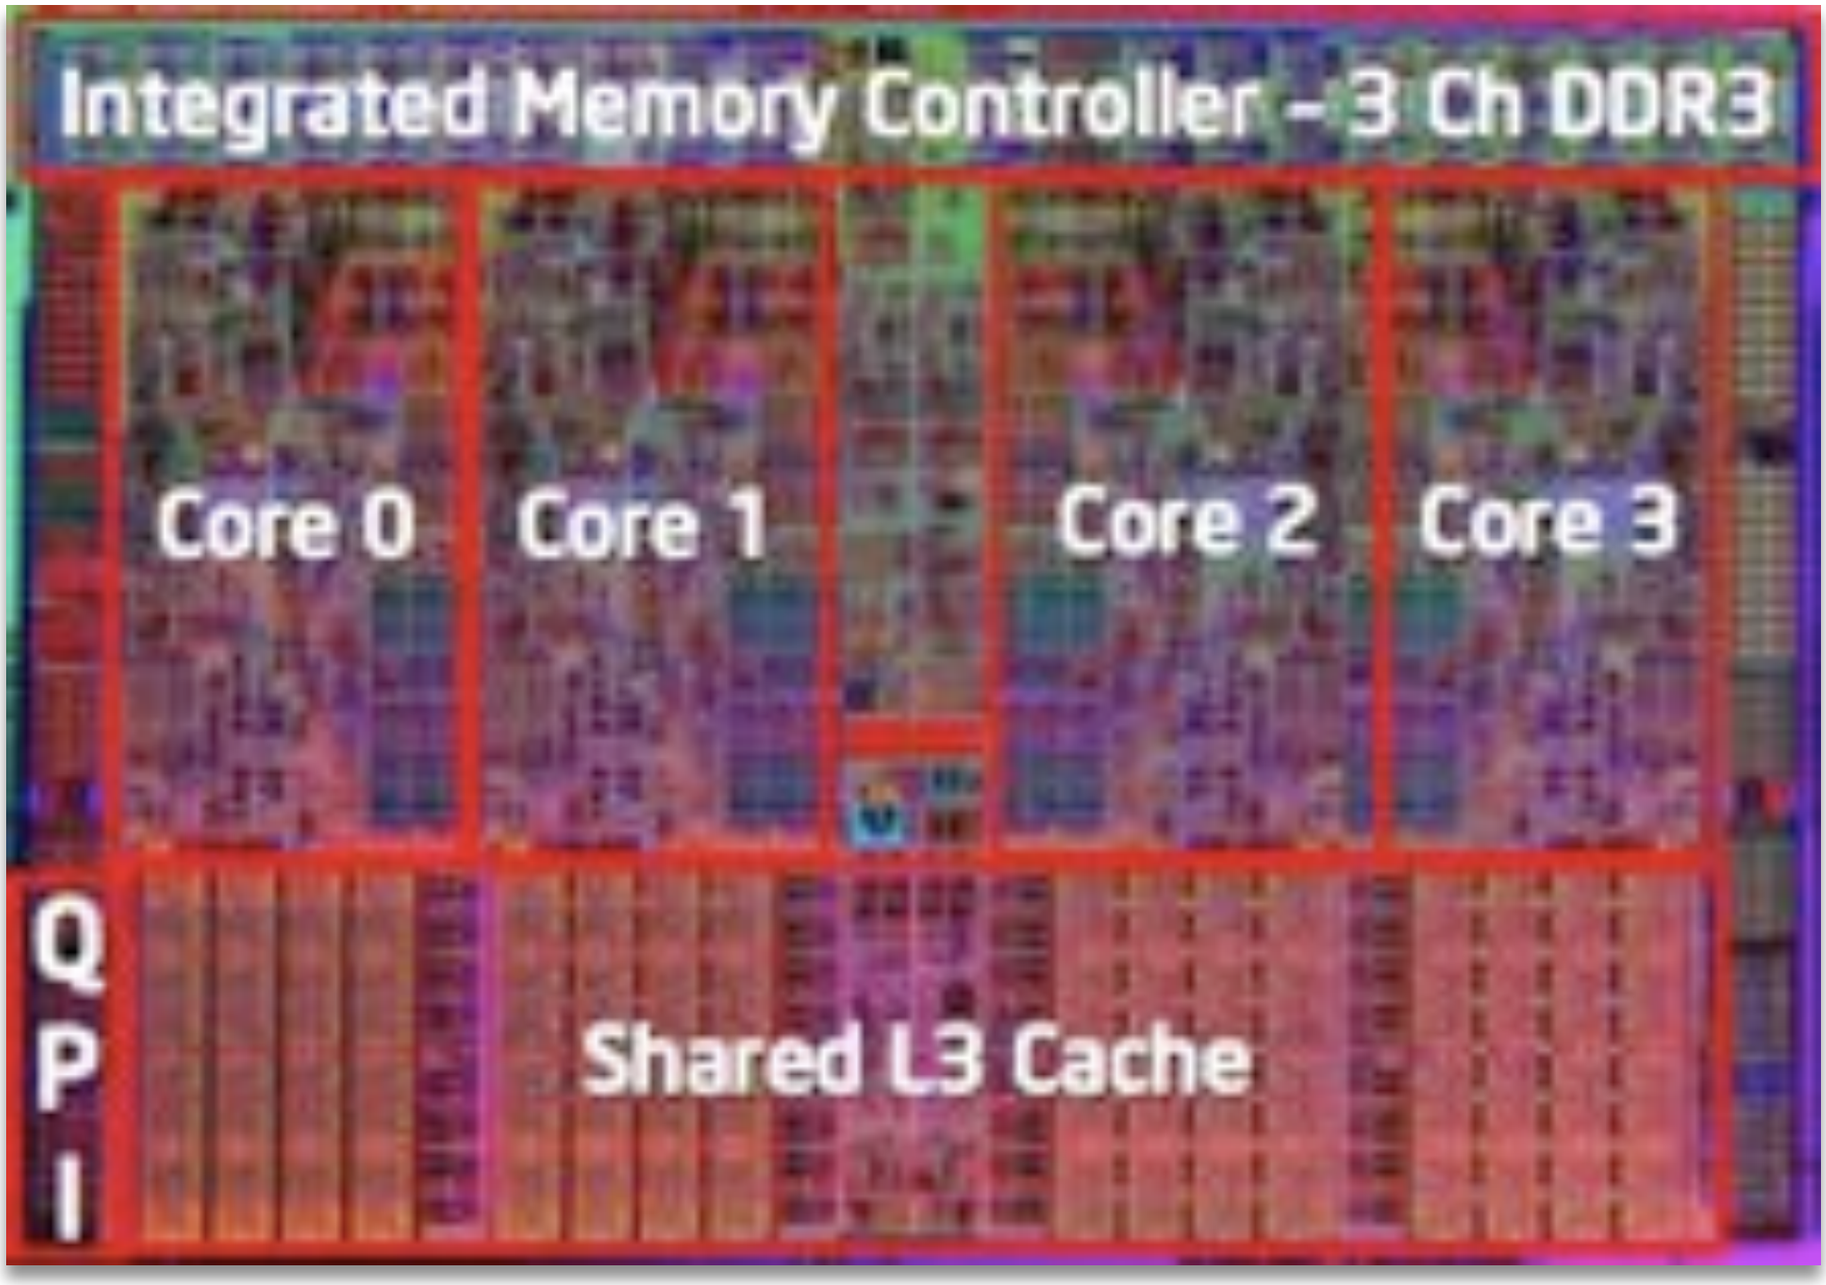

HF-QMC with Delayed updates (or Ed updates)

Ed D'Azevedo, ORNL

$$\mathbf{G}_c(\{s_i,l\}_{k+1}) = \mathbf{G}_c(\{s_i,l\}_k) + \mathbf{a}_k \times \mathbf{b}_k^t$$

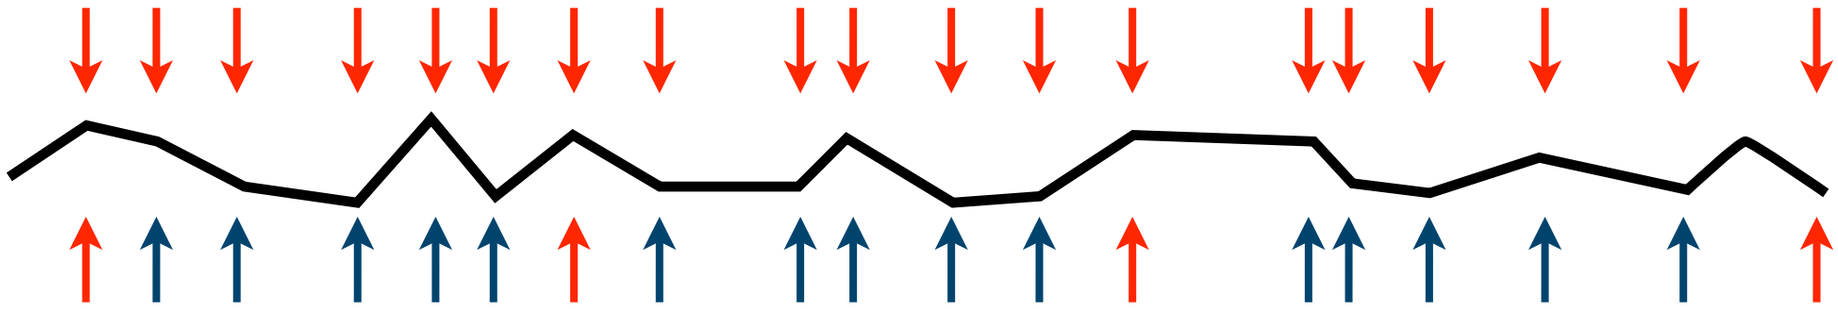

$$\mathbf{G}_c(\{s_i,l\}_{k+1}) = \mathbf{G}_c(\{s_i,l\}_0) + [\mathbf{a}_0|\mathbf{a}_1|...|\mathbf{a}_k] \times [\mathbf{b}_0|\mathbf{b}_1|...|\mathbf{b}_k]^t$$

Complexity for *k* updates remains $\mathcal{O}(kN_t^2)$

But we can replace *k* rank-1 updates with one matrix-matrix multiply plus some additional bookkeeping.

Performance improvement with delayed updates

$$N_c = 16 \qquad N_l = 150 \qquad N_t = 2400$$

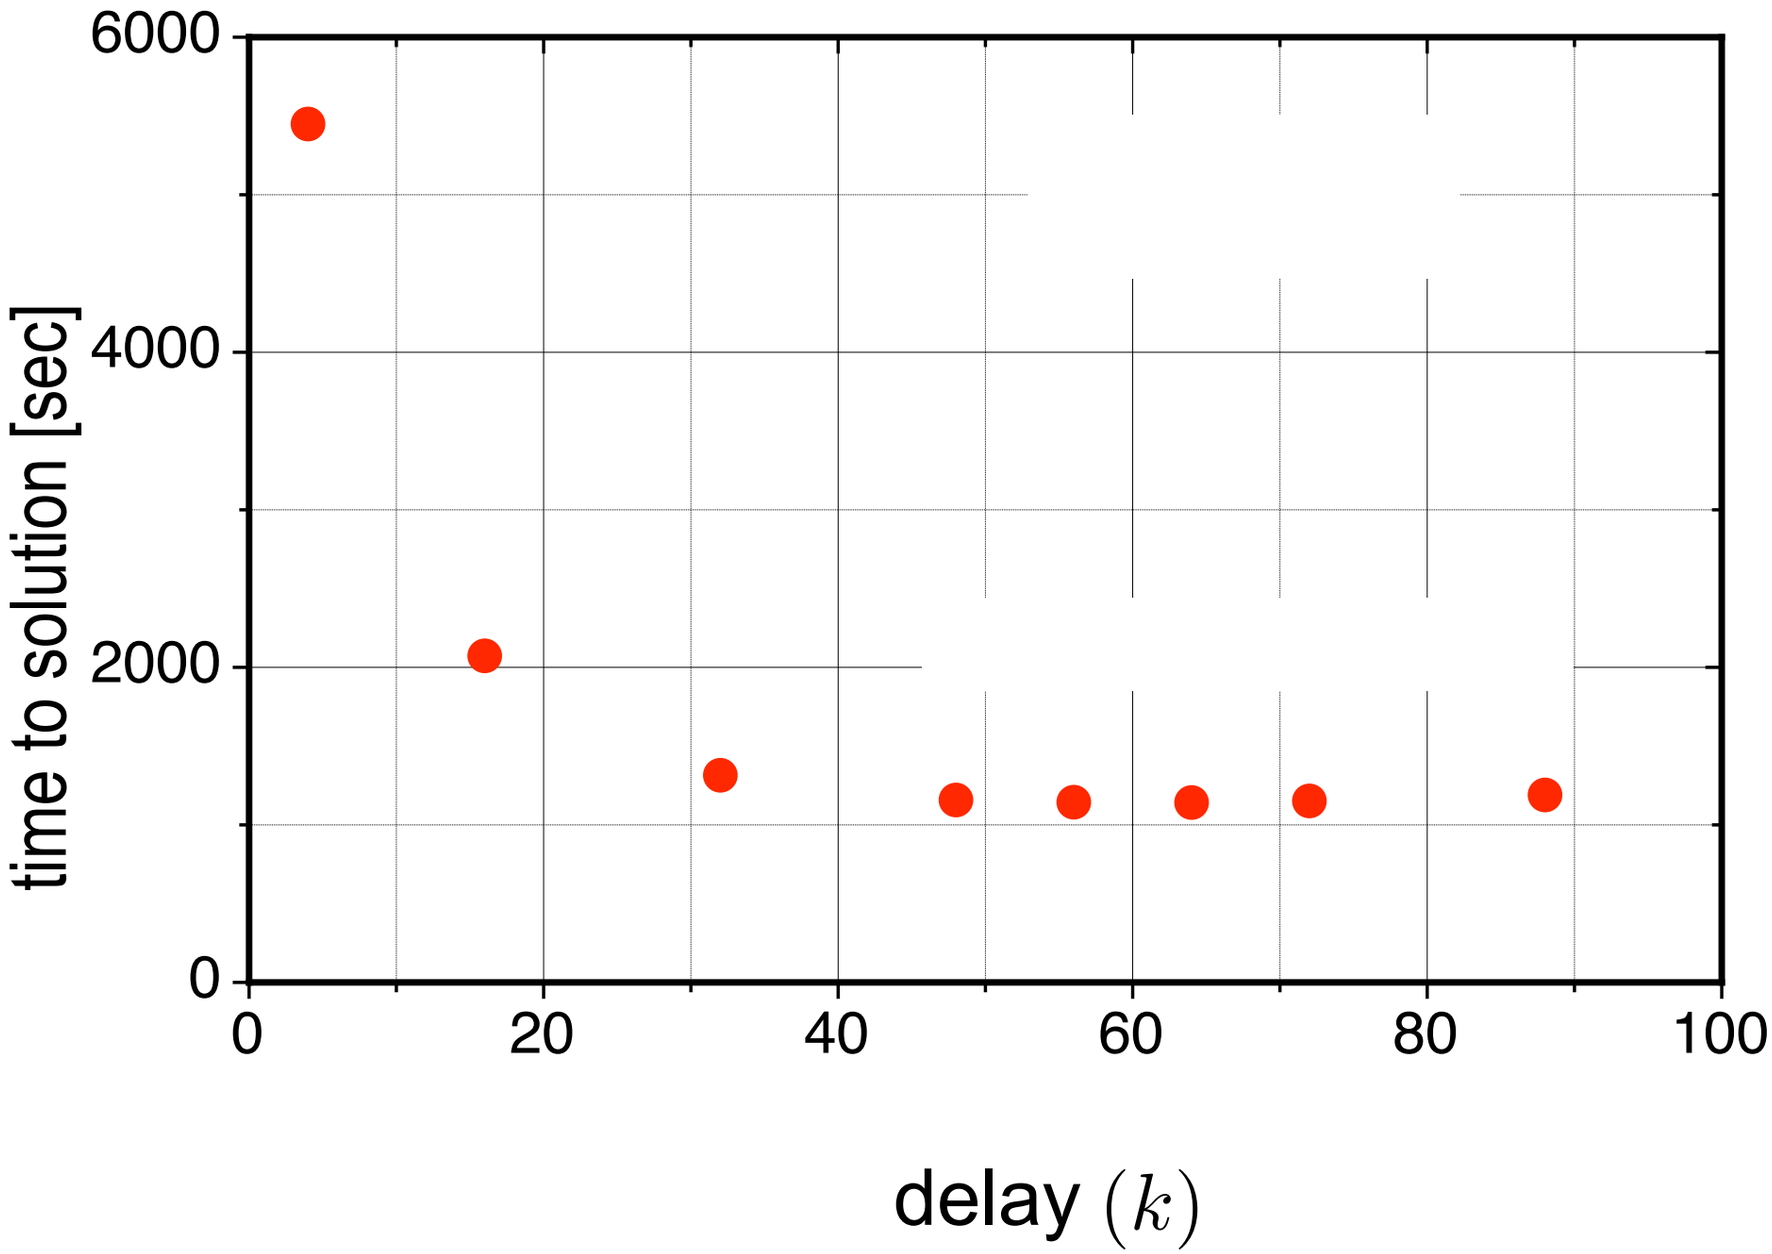

DCA++ speedup on GPU

Meredith et al., Par. Comp. **35**, 151 (2009)

Speedup of HF-QMC updates (2GHz Opteron vs. NVIDIA 8800GTS GPU):

- 9x for offloading BLAS to GPU & transferring all data (completely transparent to application code)
- 13x for offloading BLAS to GPU & lazy data transfer
- 19x for full offload HF-updates & full lazy data transfer

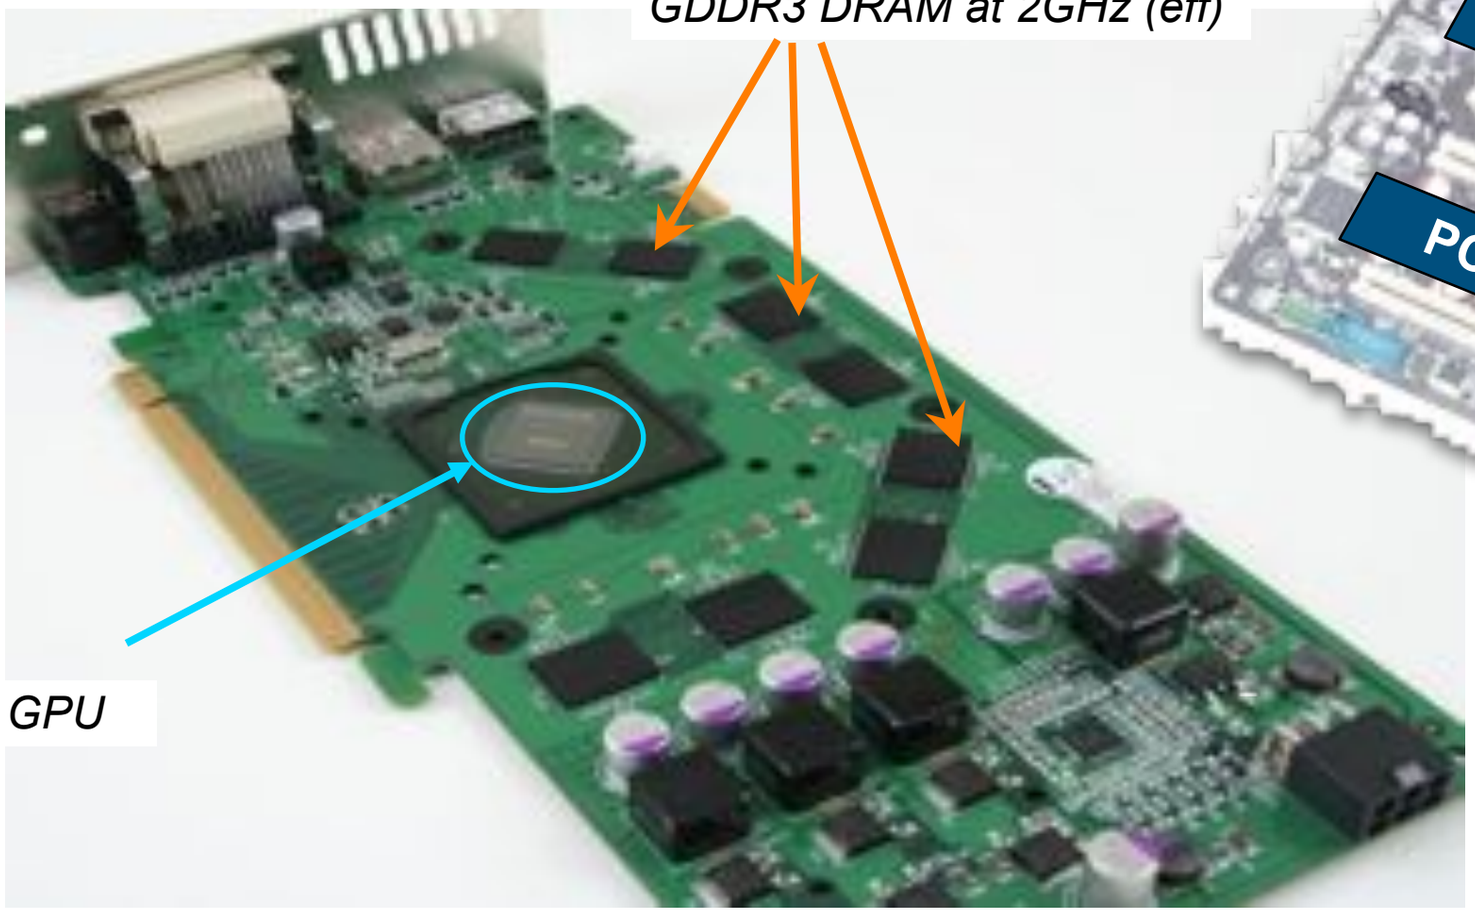

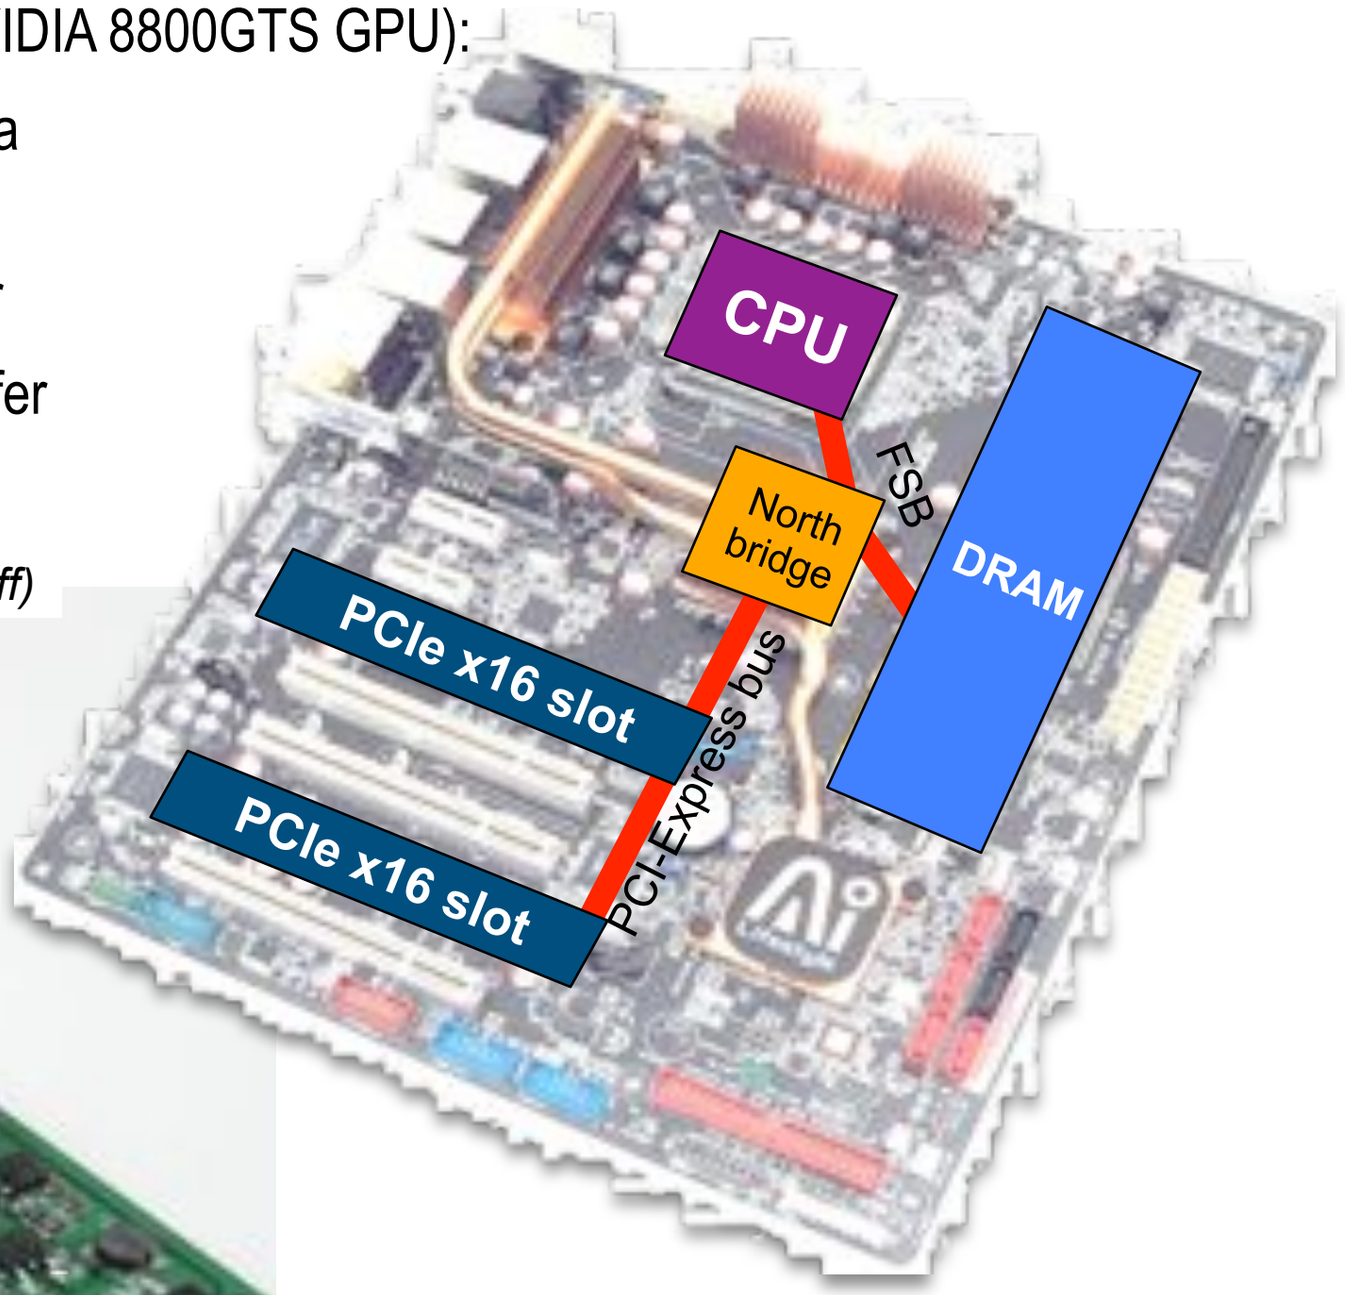

DCA++ with mixed precision

Run HF-QMC in single precision

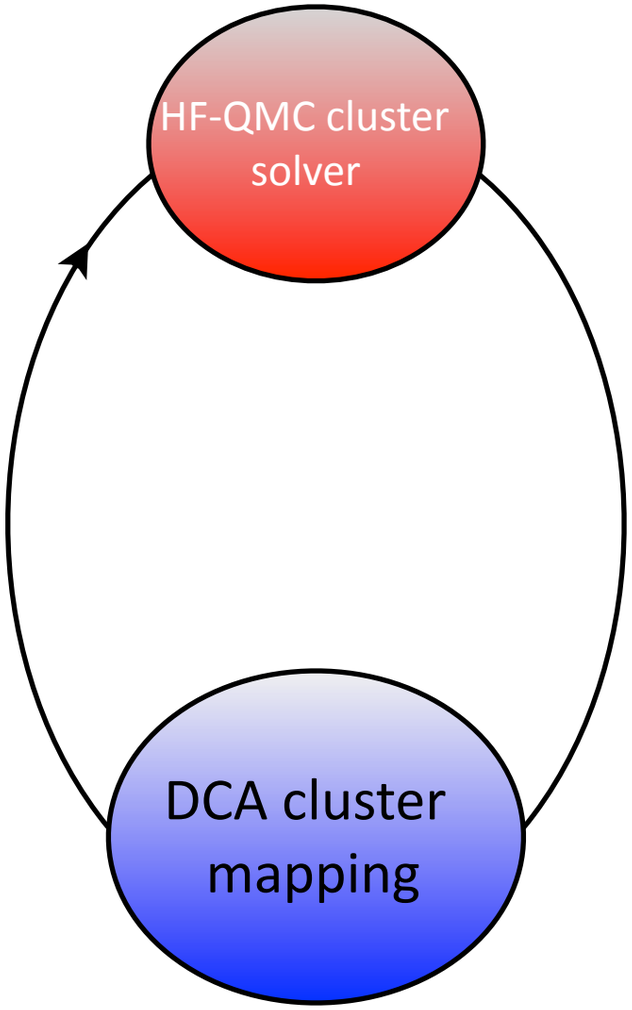

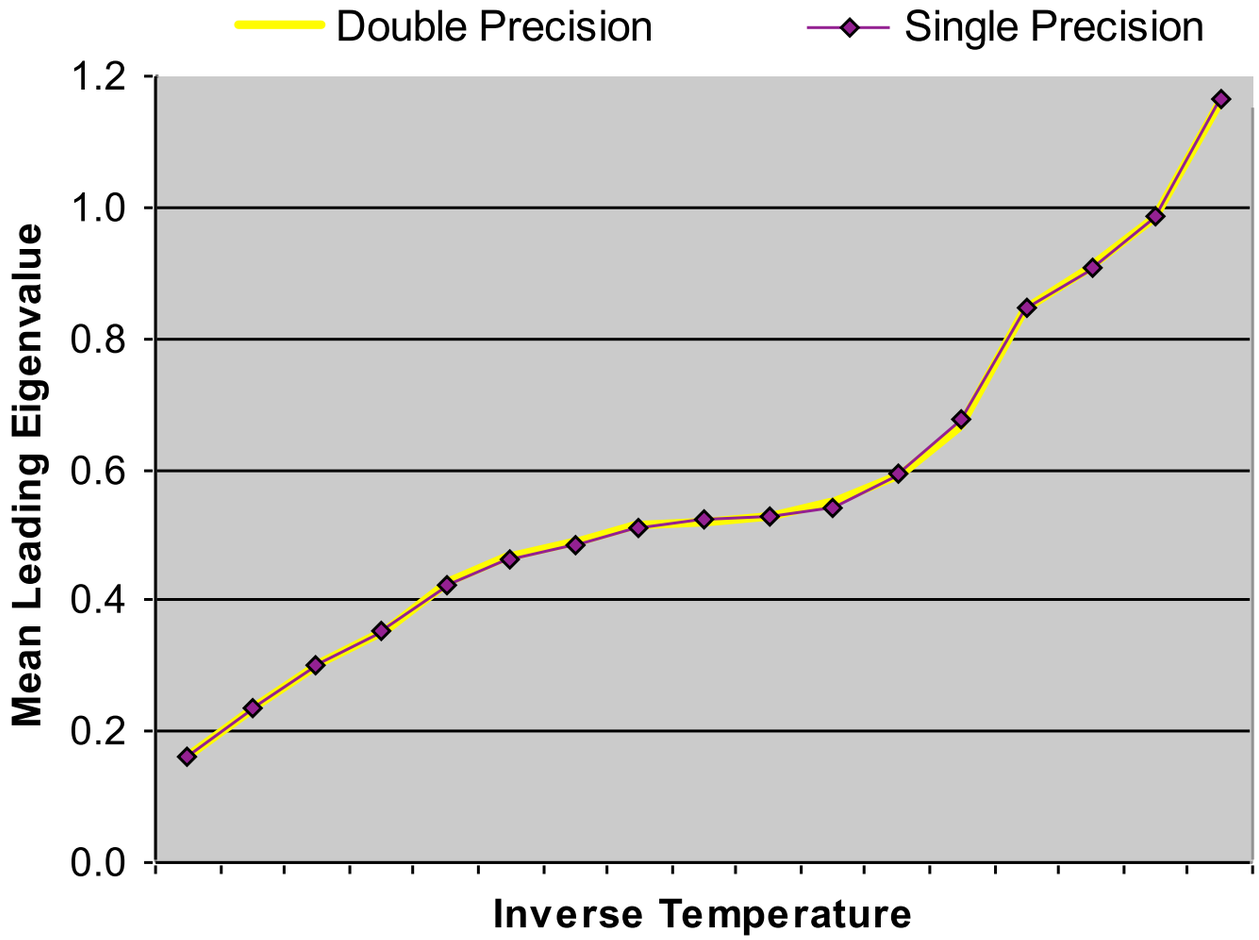

Keep the rest of the code, in particular cluster mapping in double precision

DCA++ with mixed precision

Run HF-QMC in single precision

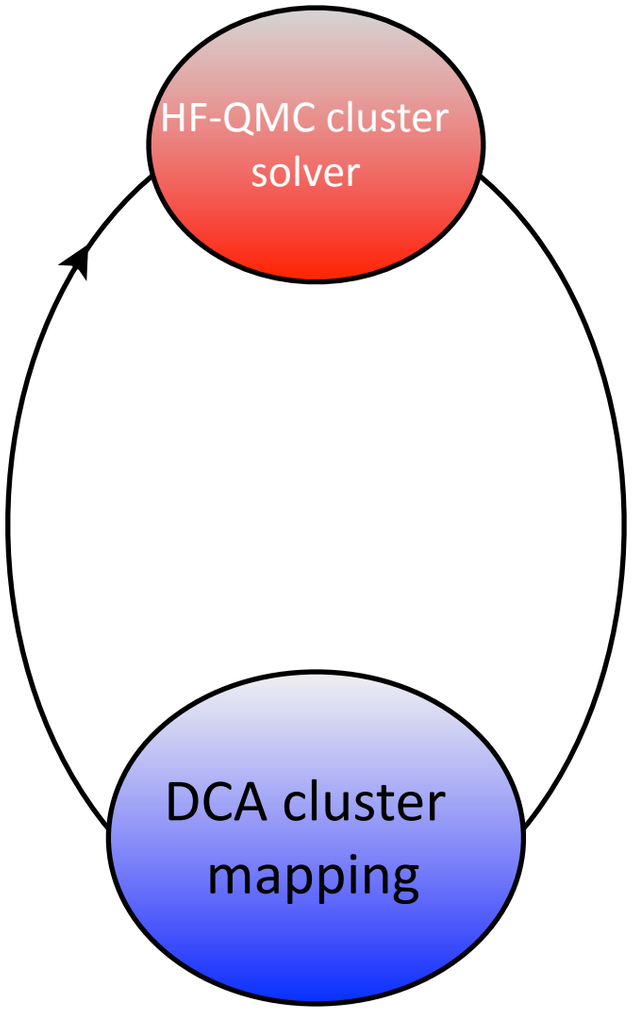

Keep the rest of the code, in particular cluster mapping in double precision

Multiple runs to compute T_c:

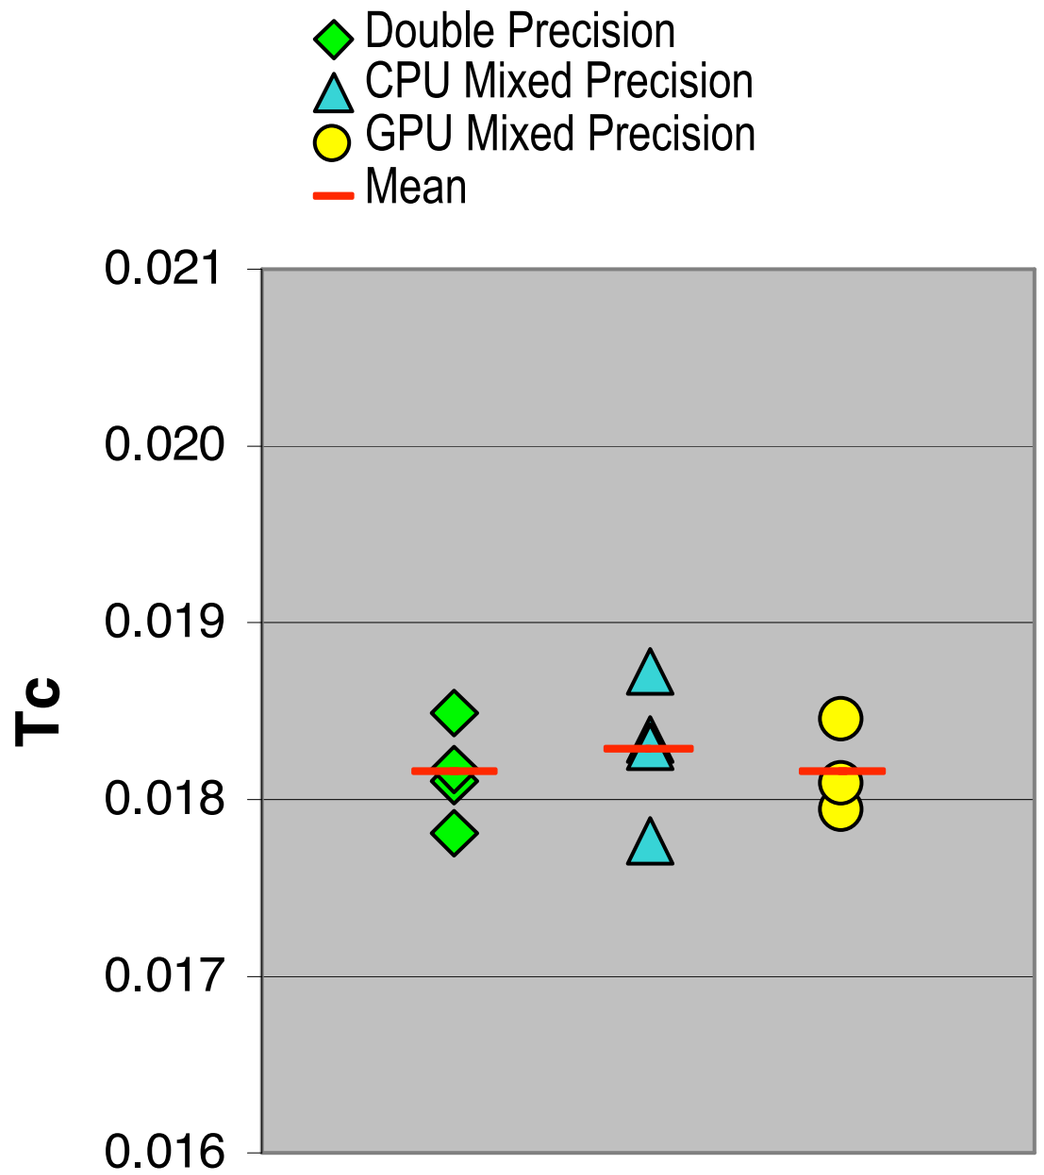

Performance improvement with delayed and mixed precision updates

$$N_c = 16 \qquad N_l = 150 \qquad N_t = 2400$$

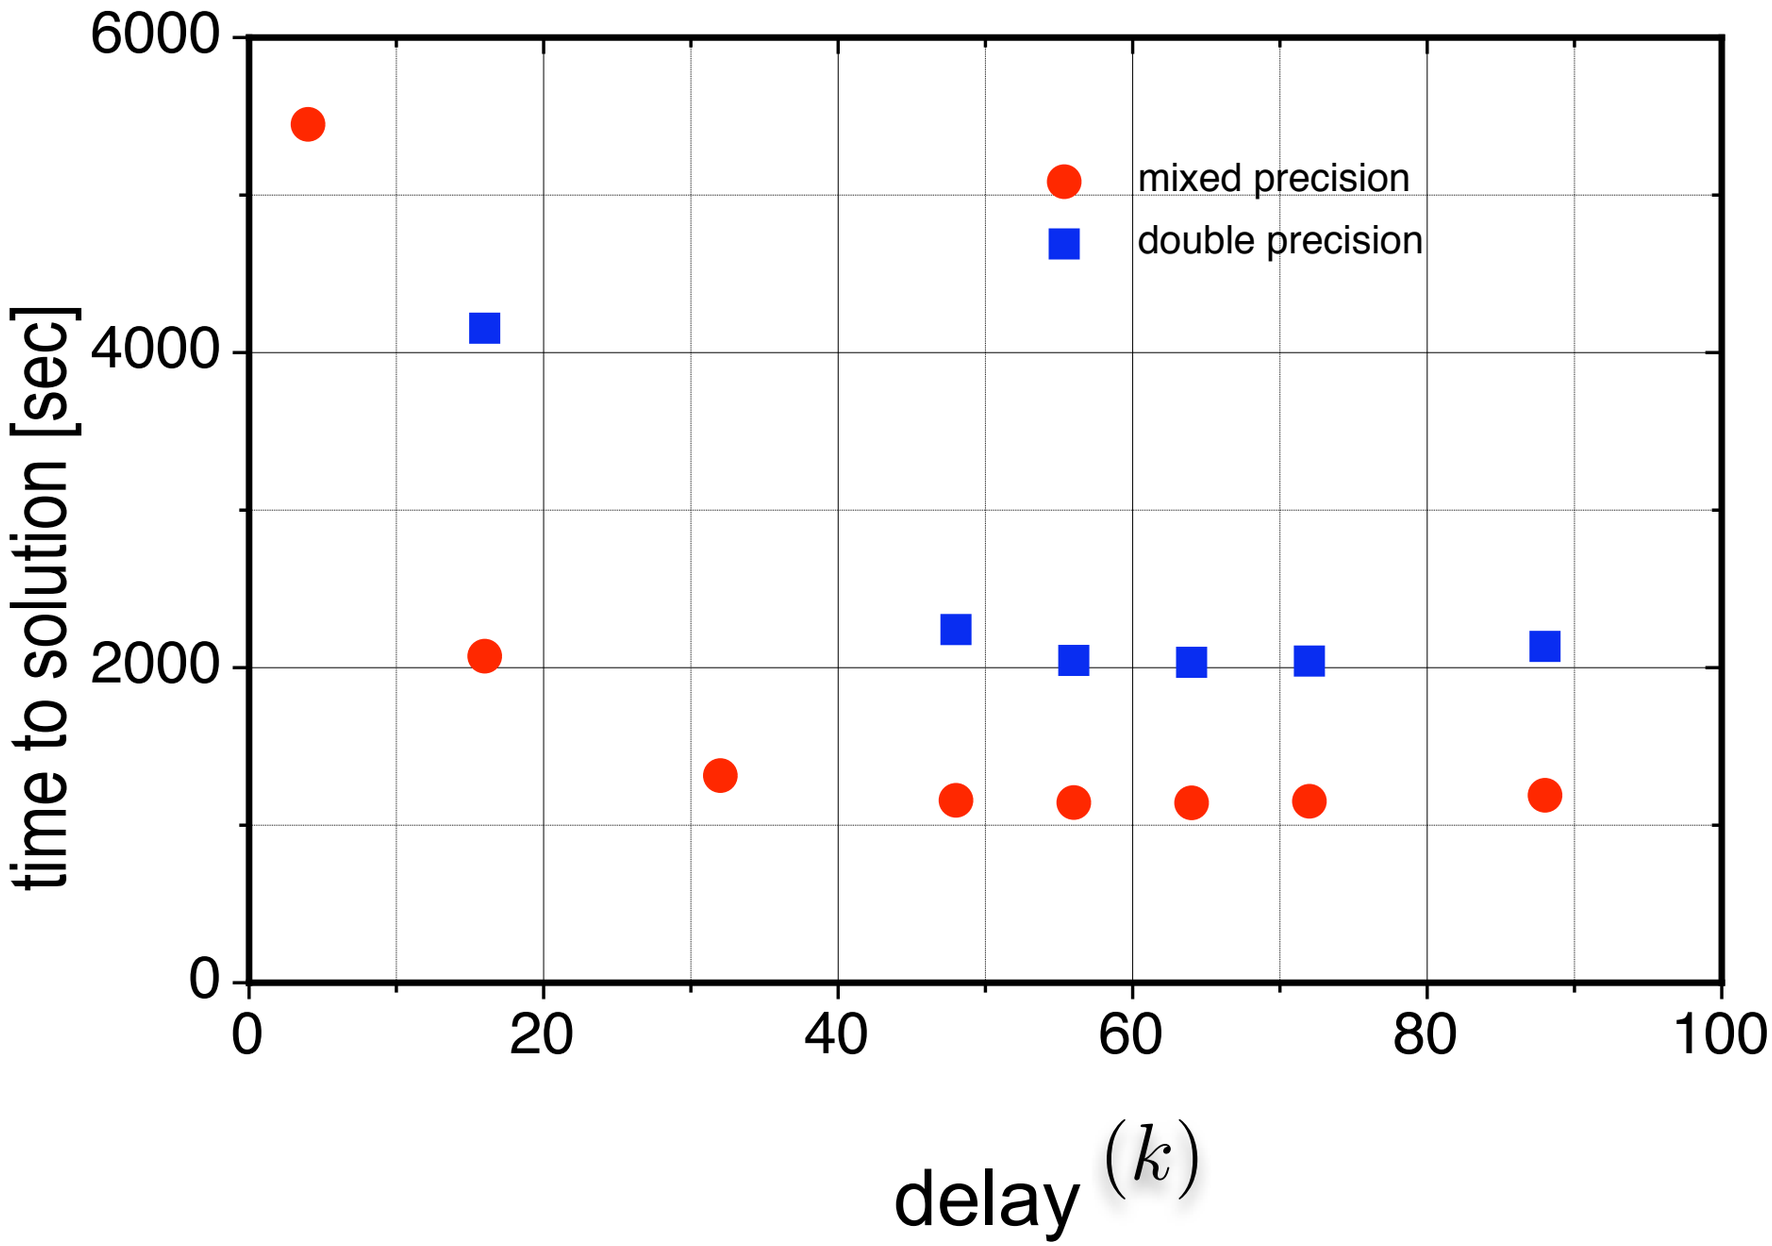

DCA++ code from a concurrency point of view

Shared memory or data parallel model

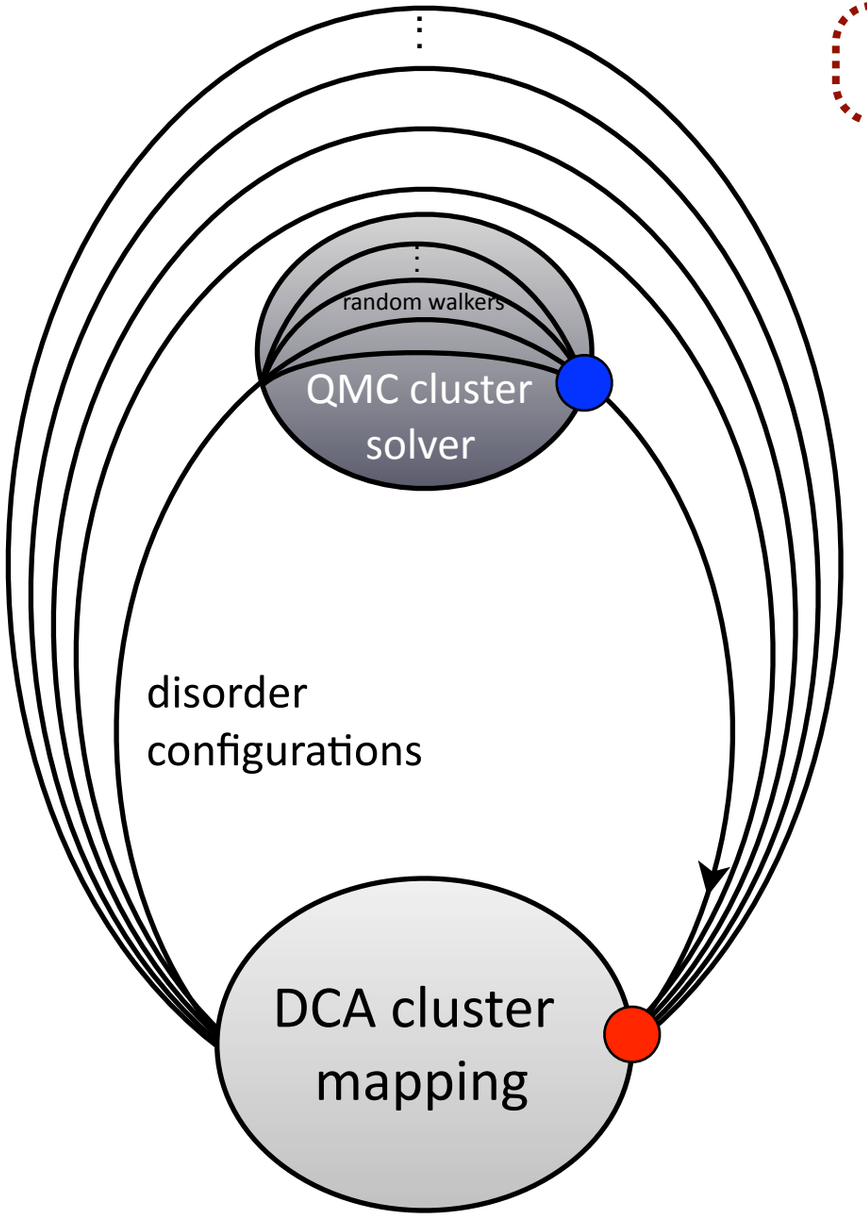

pthread / CUDA

up to 10^3 Markov chains

MPI AllReduce
MPI Broadcast

Problem of interest:
$\sim 10^2 - 10^3$ disorder configurations

MPI AllReduce

Distributed memory model

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

CSCS
Swiss National Supercomputing Centre

Cray XT5 Jaguar @ NCCS (in late 2008)

Peak: 1.382 TF/s
Quad-Core AMD
Freq.: 2.3 GHz
150,176 cores
Memory: 300 TB
For more details, see
www.nccs.gov

Sustained performance of DCA++ on Cray XT5

Weak scaling with number disorder configurations, each running on 128 Markov chains on 128 cores (16 nodes) - 16 site cluster and 150 time slices

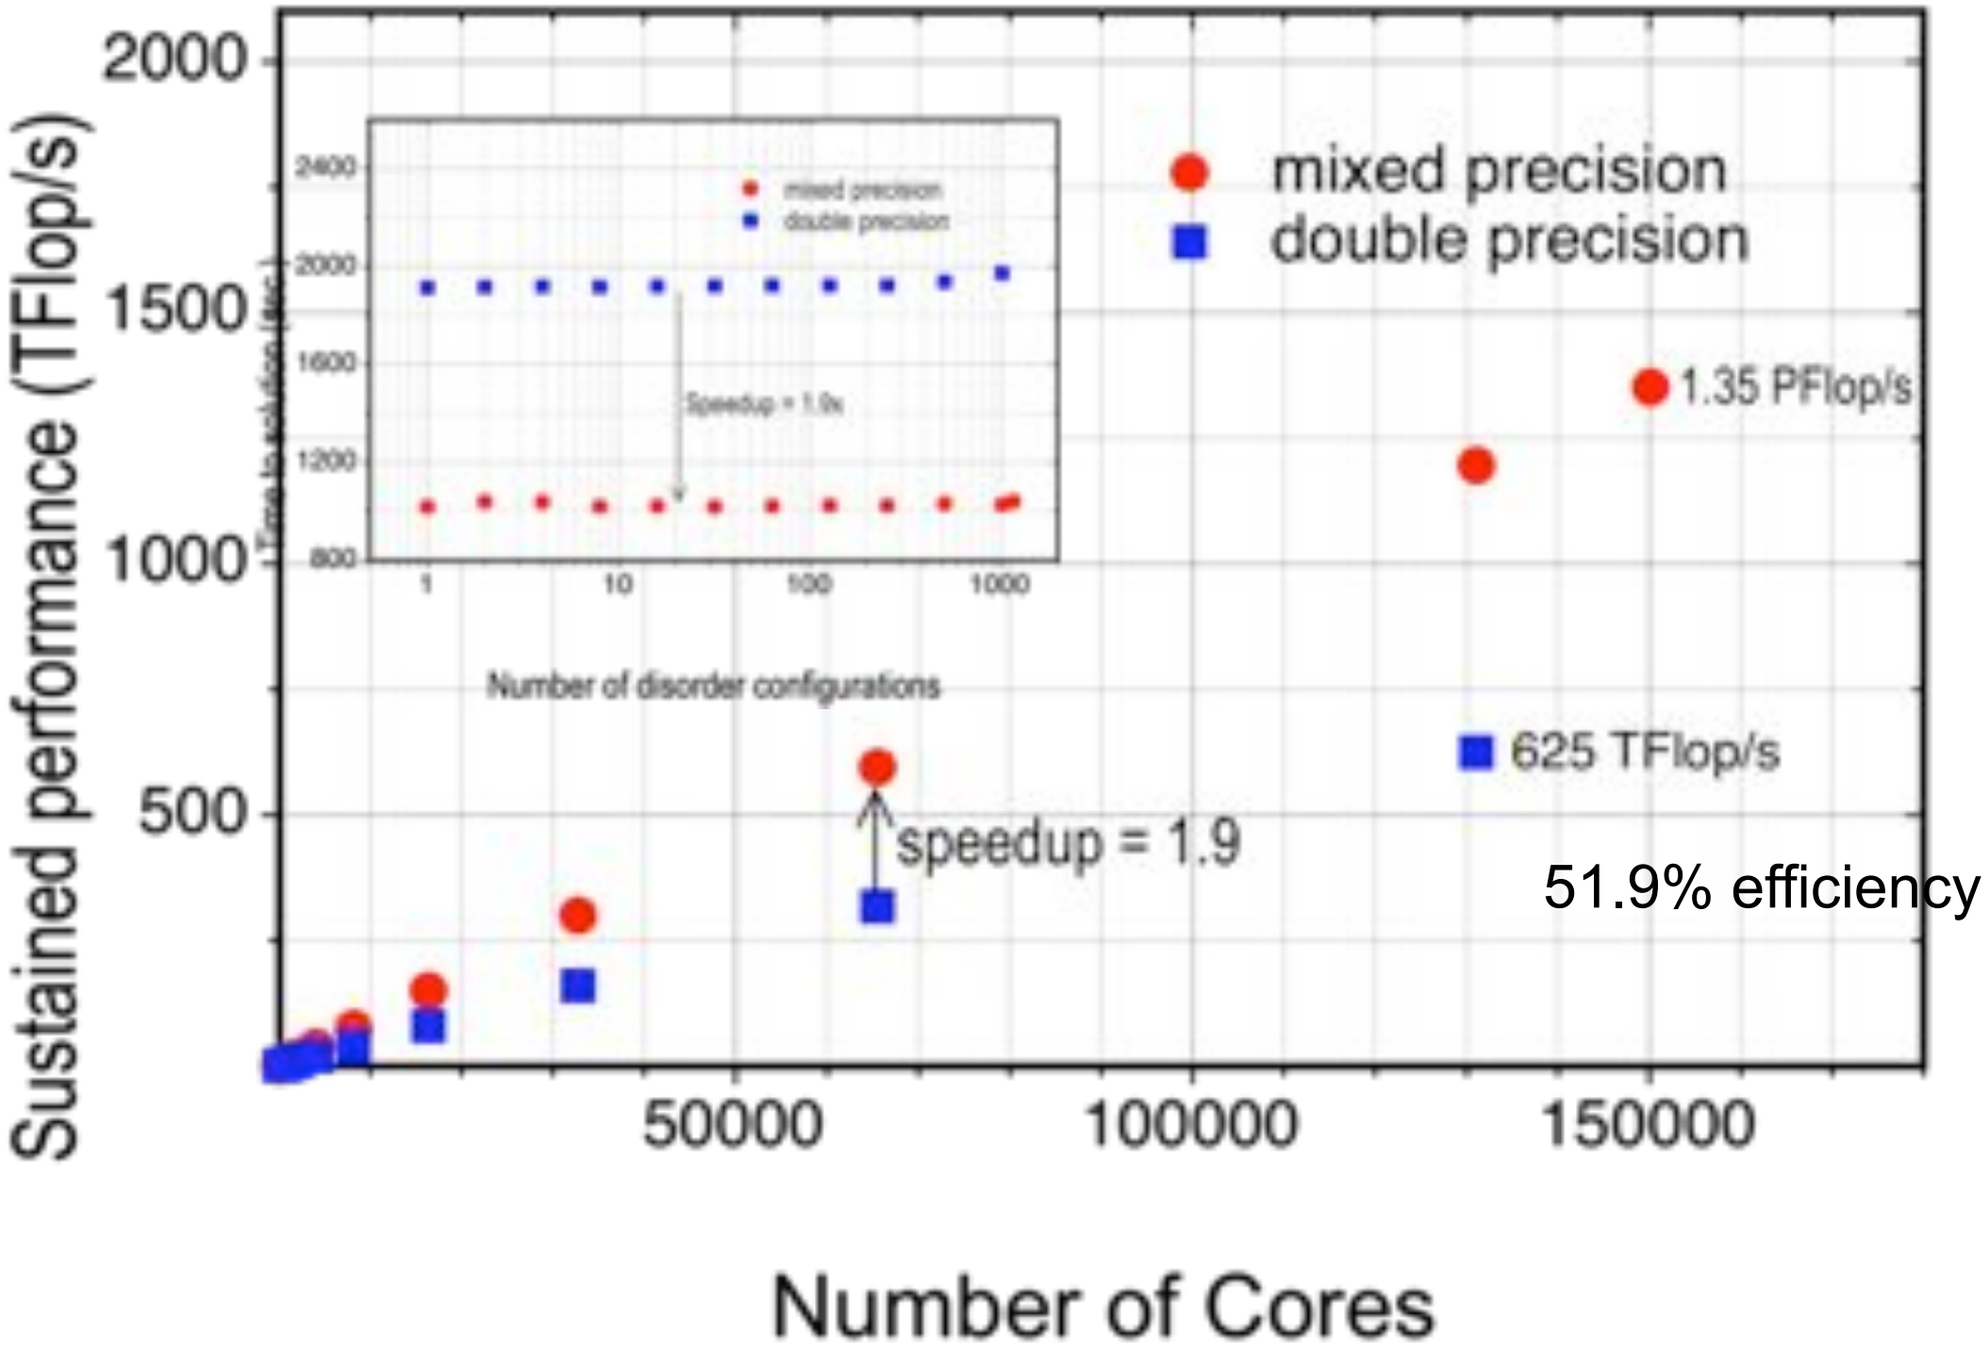

High-T_c superconductivity: DCA/QMC simulations of the 2D Hubbard model

First simulation proving model describes superconducting transition

High memory bandwidth helps performance of rank 1 update in DCA/QMC algortihm

Finally settled two-decade old debate over nature of pairing mechanism!

Realistic models of disorder and nanoscale inhomogeneities

In 5 years, factor
10^3 X compute
10^{-4} X time to solution

Sustained performance of DCA/QMC

2004 ~1 TF
2005 ~6 TF
2006 ~26 TF
2007 ~60 TF
2008 ~140 TF
2009 ~260 TF
1.34 PF
>2 PF

First sustained petaflop/s under production condition

Change DCA/QMC algorithm to mixed precision

Modify QMC algorithm: replaced (delayed) rank 1 updates by matrix multiply

2004 3 TF
2005 18.5 TF
2006 26 TF
2007 54 TF
2008 119 TF
2009 263 TF
1.4 PF
2.4 PF

Relationship between simulations and supercomputer system (how things should be)

Simulations + Theory + Experiment

Science

Model & method of solution

Mapping applications to computer system
> Algorithm re-engineering
> Software refactoring
> Domain specific libraries/languages, etc.

Basic numerical libraries
Programming environment
Runtime system
Operating systems
Computer Hardware

Co-Design

Supercomputer

Applications running at scale on Jaguar @ ORNL

Fall 2009

Domain area	Code name	Institution	# of cores	Performance	Notes
Materials	DCA++	ORNL	213,120	**1.9 PF**	2008 Gordon Bell Prize Winner
Materials	WL-LSMS	ORNL/ETH	223,232	**1.8 PF**	2009 Gordon Bell Prize Winner
Chemistry	NWChem	PNNL/ORNL	224,196	1.4 PF	2008 Gordon Bell Prize Finalist
Materials	OMEN	Duke	222,720	860 TF	
Chemistry	MADNESS	UT/ORNL	140,000	550 TF	
Materials	LS3DF	LBL	147,456	442 TF	2008 Gordon Bell Prize Winner
Seismology	SPECFEM3D	USA (multiple)	149,784	165 TF	2008 Gordon Bell Prize Finalist
Weather	WRF	USA (multiple)	150,000	50 TF	
Combustion	S3D	SNL	147,456	83 TF	

Algorithmic motifs and their arithmetic intensity

Arithmetic intensity: number of operations per word of memory transferred

Rank-1 update in HF-QMC → **Rank-N update in DCA++**

Finite difference / stencil → **QMR in WL-LSMS**

Linpack

Sparse linear algebra

BLAS1

BLAS2

FFT

Dense matrix-matrix

BLAS3

O(1) **O(log N)** **O(N)**

Metric for “super” in supercomputing?

- Is sustained floating point performance really the right metric?
  - Much better than peak performance but doesn’t do justice for large classes of applications
- There is probably no simple answer!
- Optimize simulation systems for things that matter for the scientists / engineers
- **Time to solution** (minimize of good enough)
- **Energy to solution** / total cost of ownership

Collaborator – DCA++ case study

- DCA++: Thomas Maier, Mike Summers, Gonzalo Alvarez, Paul Kent
- GPU: Jeremy Meredith and Jeff Vetter
- Applied math: Ed D'Azevedo
- Cray Inc.: Jeff Larkin and John Levesque
- NCCS: Markus Eisenbach, Don Maxwell (and many others)
- Doug Scalapino (UCSB) and Mark Jarrell (now at LSU)