ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

CSCS
Swiss National Supercomputing Centre

The DCA++ Story

How new algorithms, new computers, and innovative software design allow us to solve real simulation problems of high temperature superconductivity

Thomas C. Schulthess
schulthess@cscs.ch

Cray User Group Meeting, Atlanta, May 4-7 2009

Oak Ridge
National Laboratory

Superconductivity: a state of matter with zero electrical resistivity

Discovery 1911

Heike Kamerlingh Onnes (1853-1926)

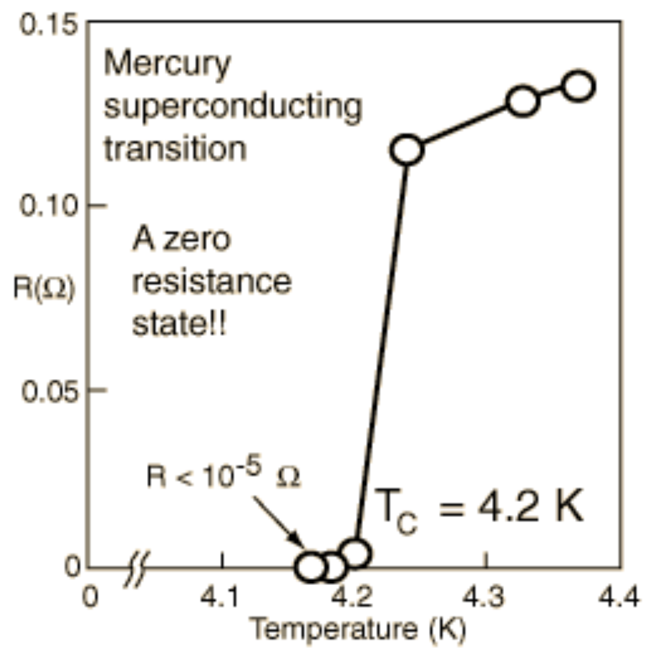

Superconductor repels magnetic field
Meissner and Ochsenfeld, Berlin 1933

Microscopic Theory for Superconductivity 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 electrons 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 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.

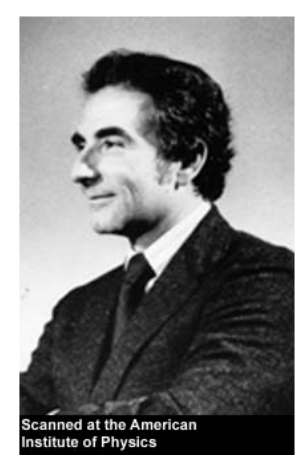

BCS Theory generally accepted in the early 1970s

Fermions, Bosons, and Cooper Pairs

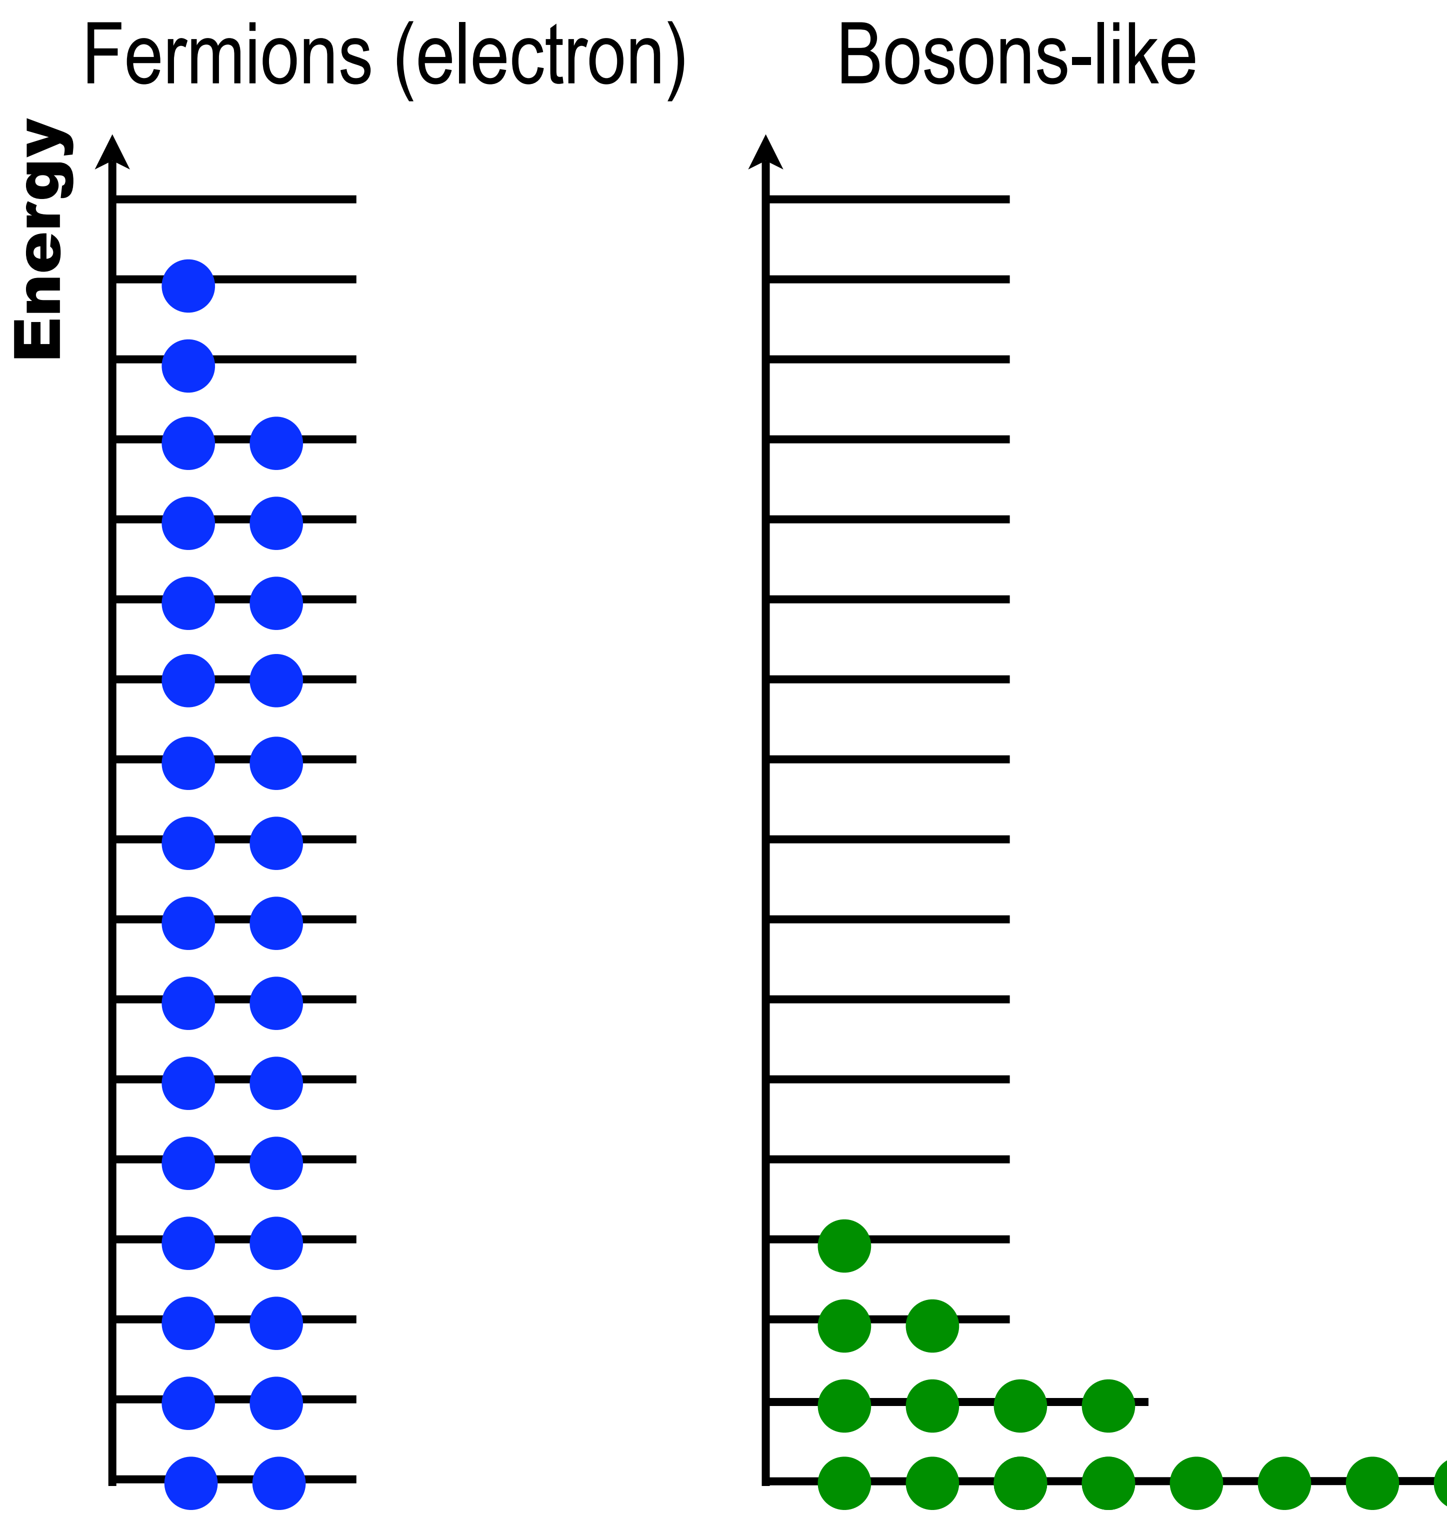

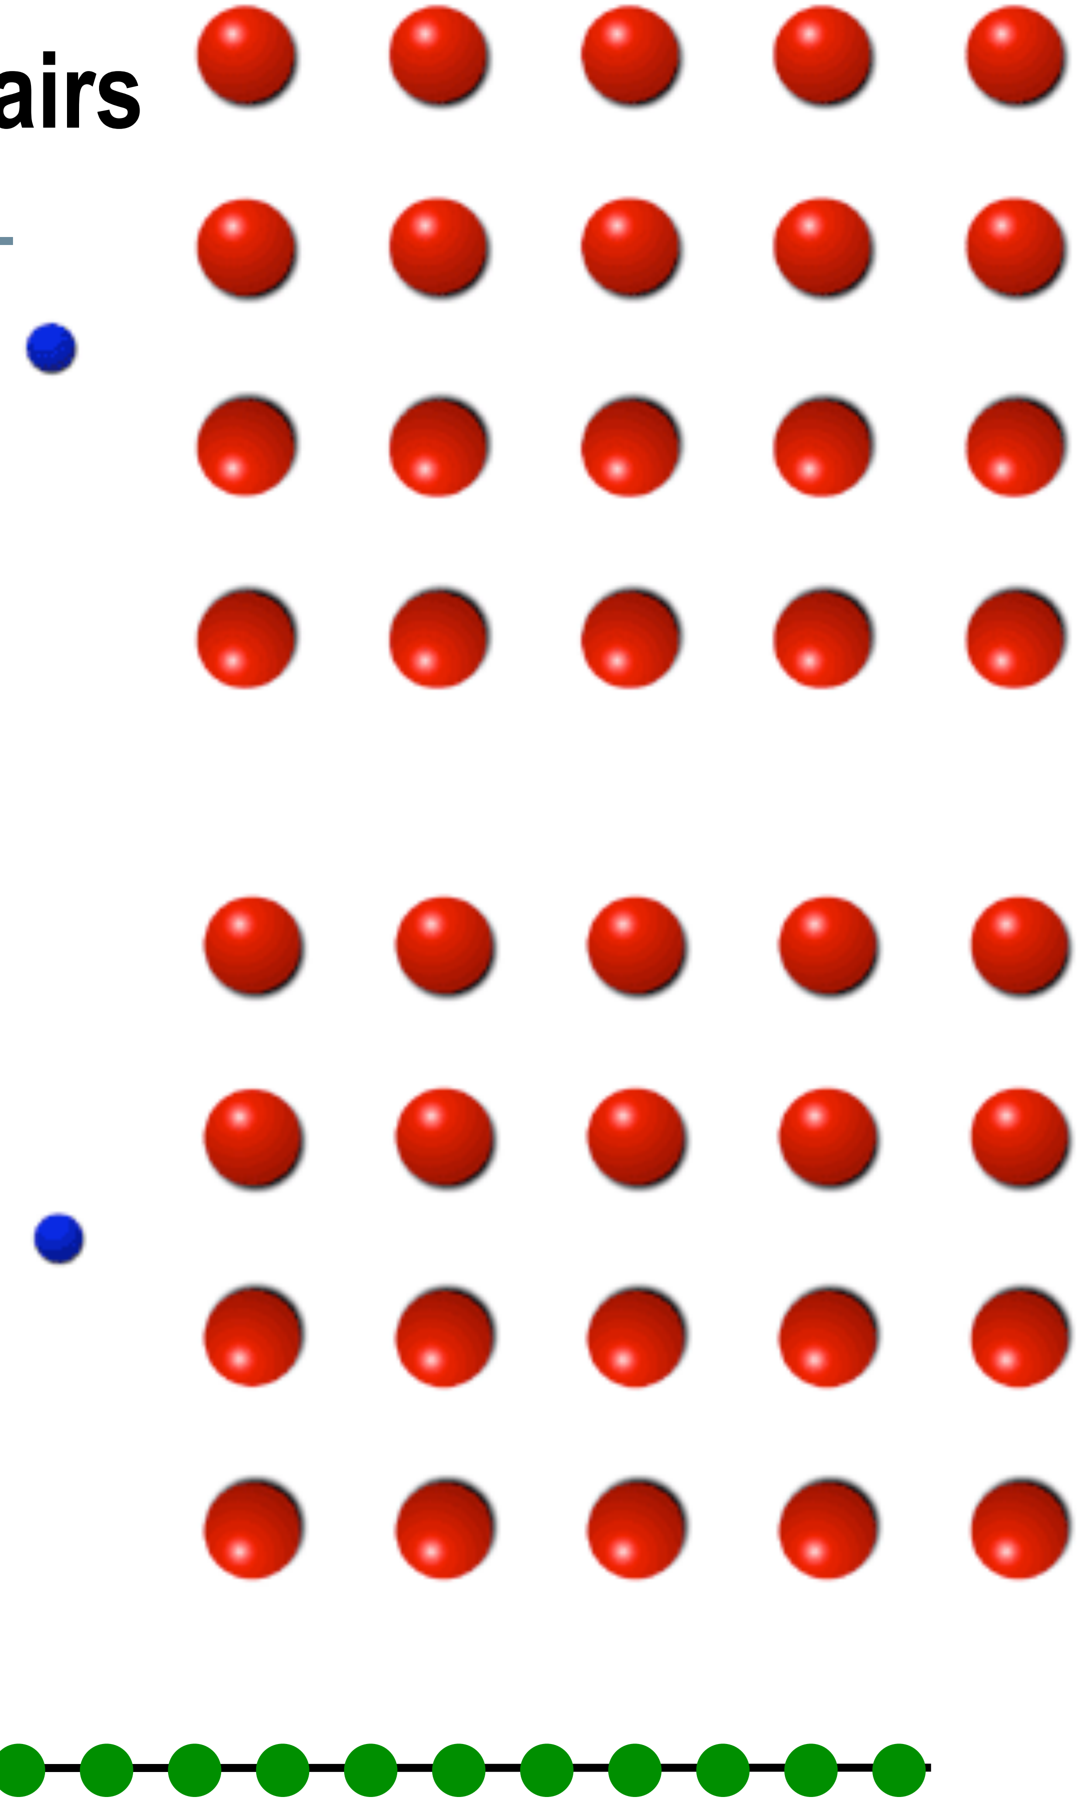

Superconductivity in the cuprates

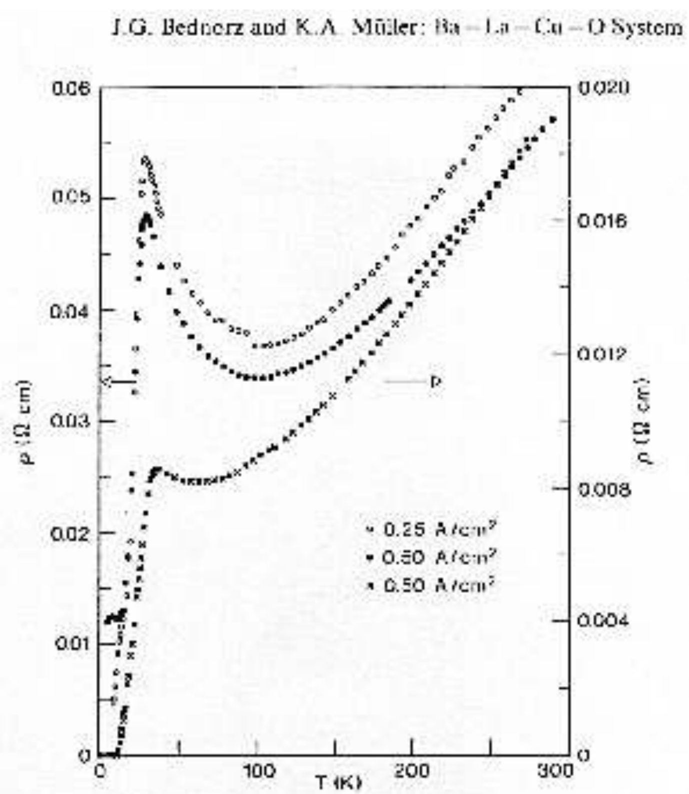

Discovery 1986

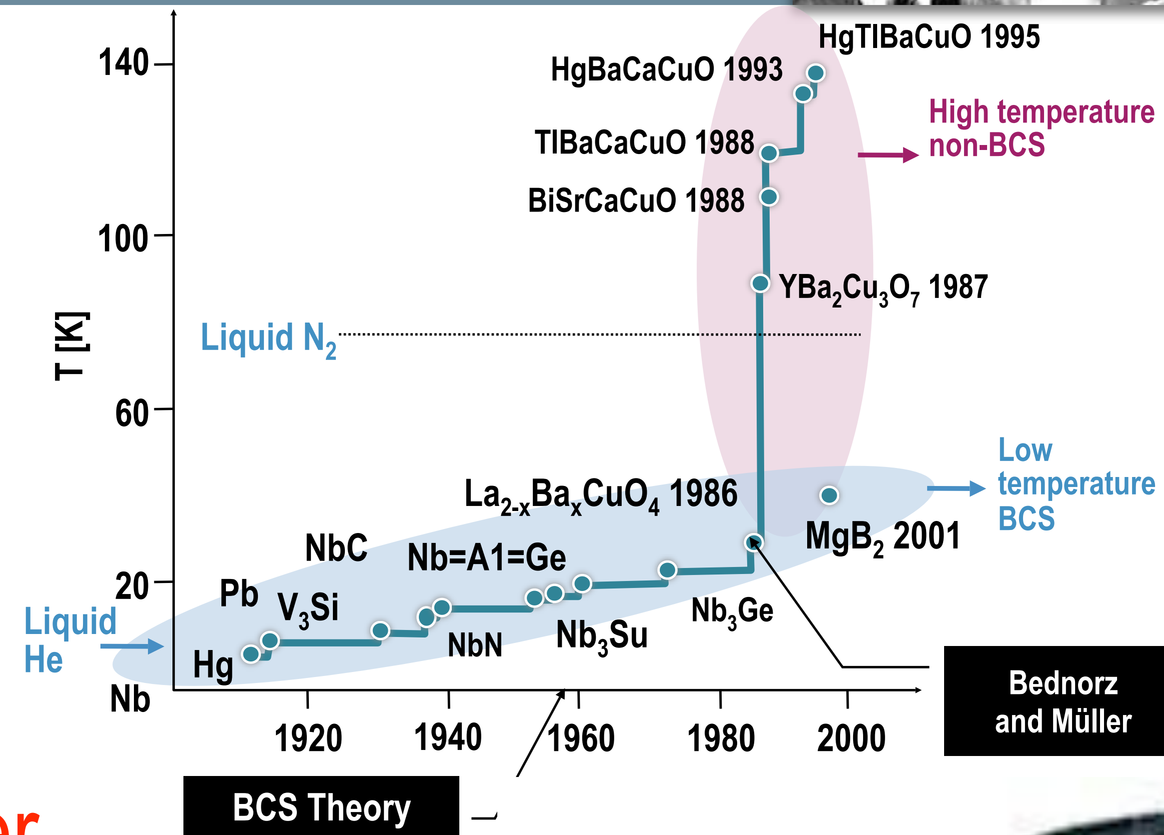

Two decades later

- Progress has been made in numerous areas relevant to applications
- Highest transition temperature (T_c) observed in a superconduc
- No predictive power for T_c in known materials
- No predictive power for design of new SC materials
- No explanation for other unusual properties of cuprates (pseudogap, transport, ...)
- Only partial consensus on which materials aspects are essential for high-T_c superconductivity
- No controlled solution for proposed models

The role of inhomogeneities

Stripes in neutron scattering:
Tranquada *et al.* ’95,
Mook *et al.*, ’00, ...

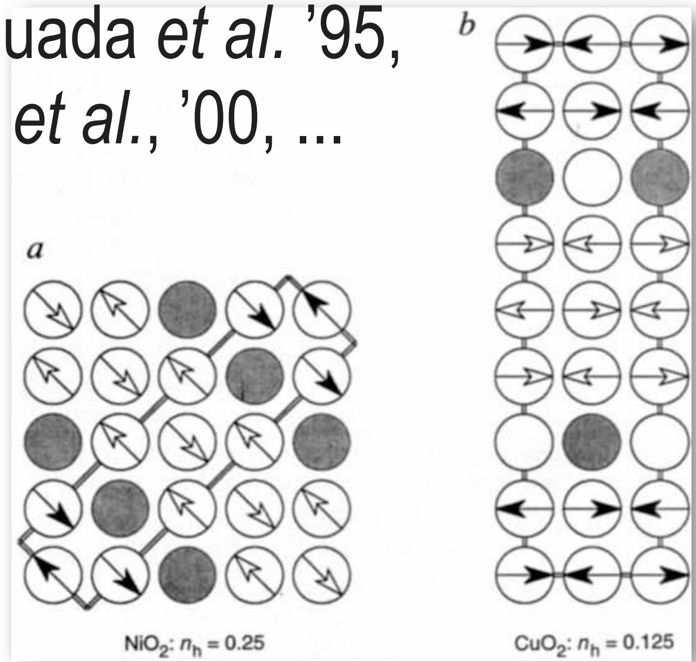

Random SC gap modulations in STM (BSCCO):
Lang *et al.* ‘02

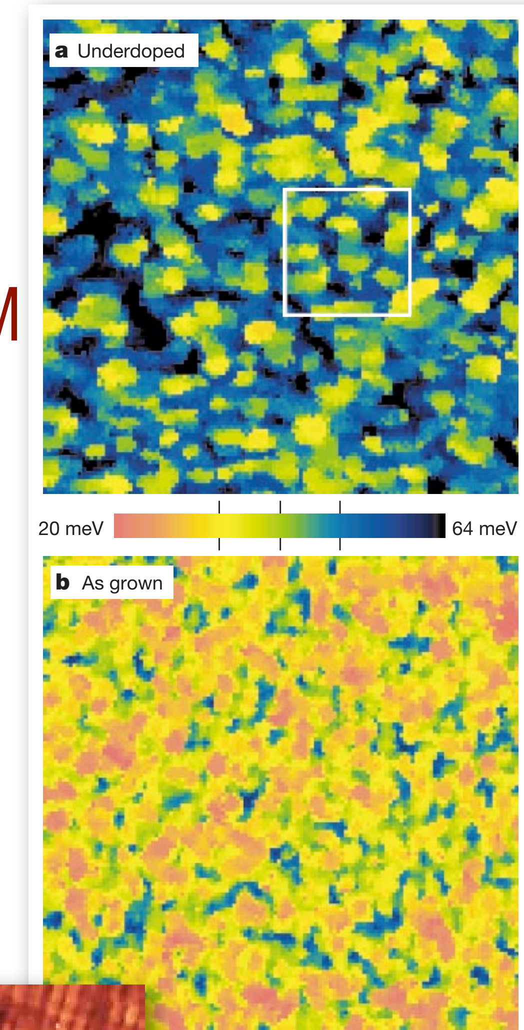

Charge ordered “checkerboard” state (Na doped cuprates):
Hanaguri *et al.* ‘04

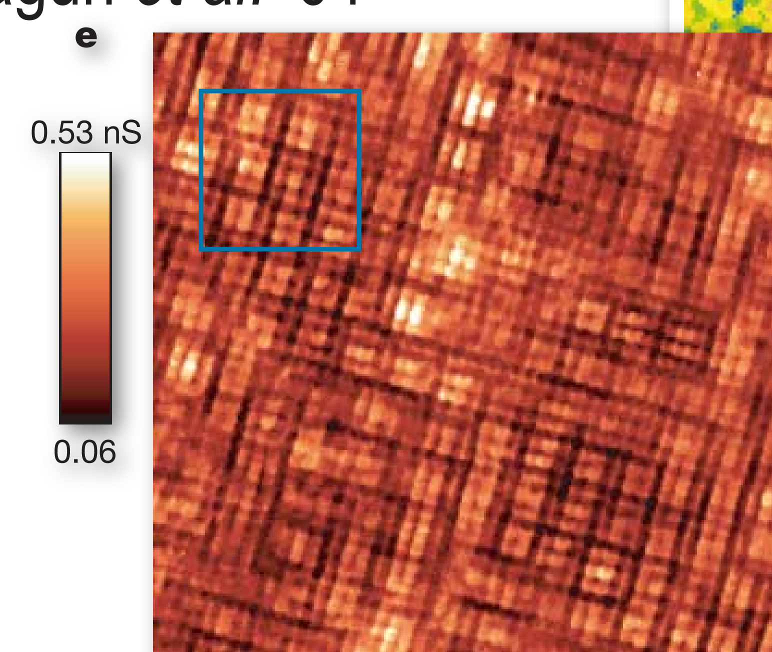

Random gap modulations above T_c (BSCCO):
Gomes *et al.* ‘07

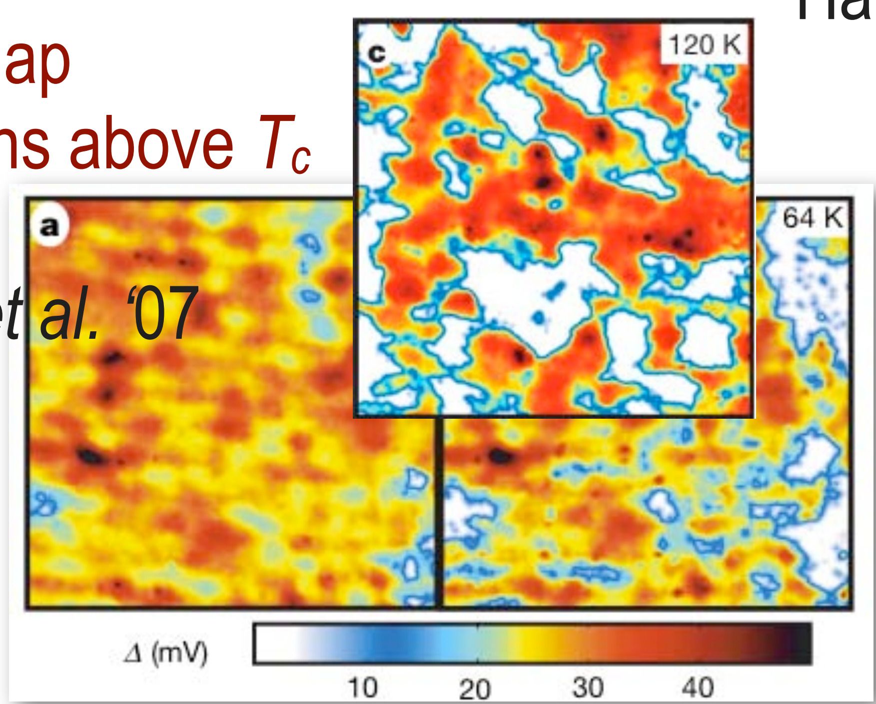

From cuprate materials to the Hubbard model

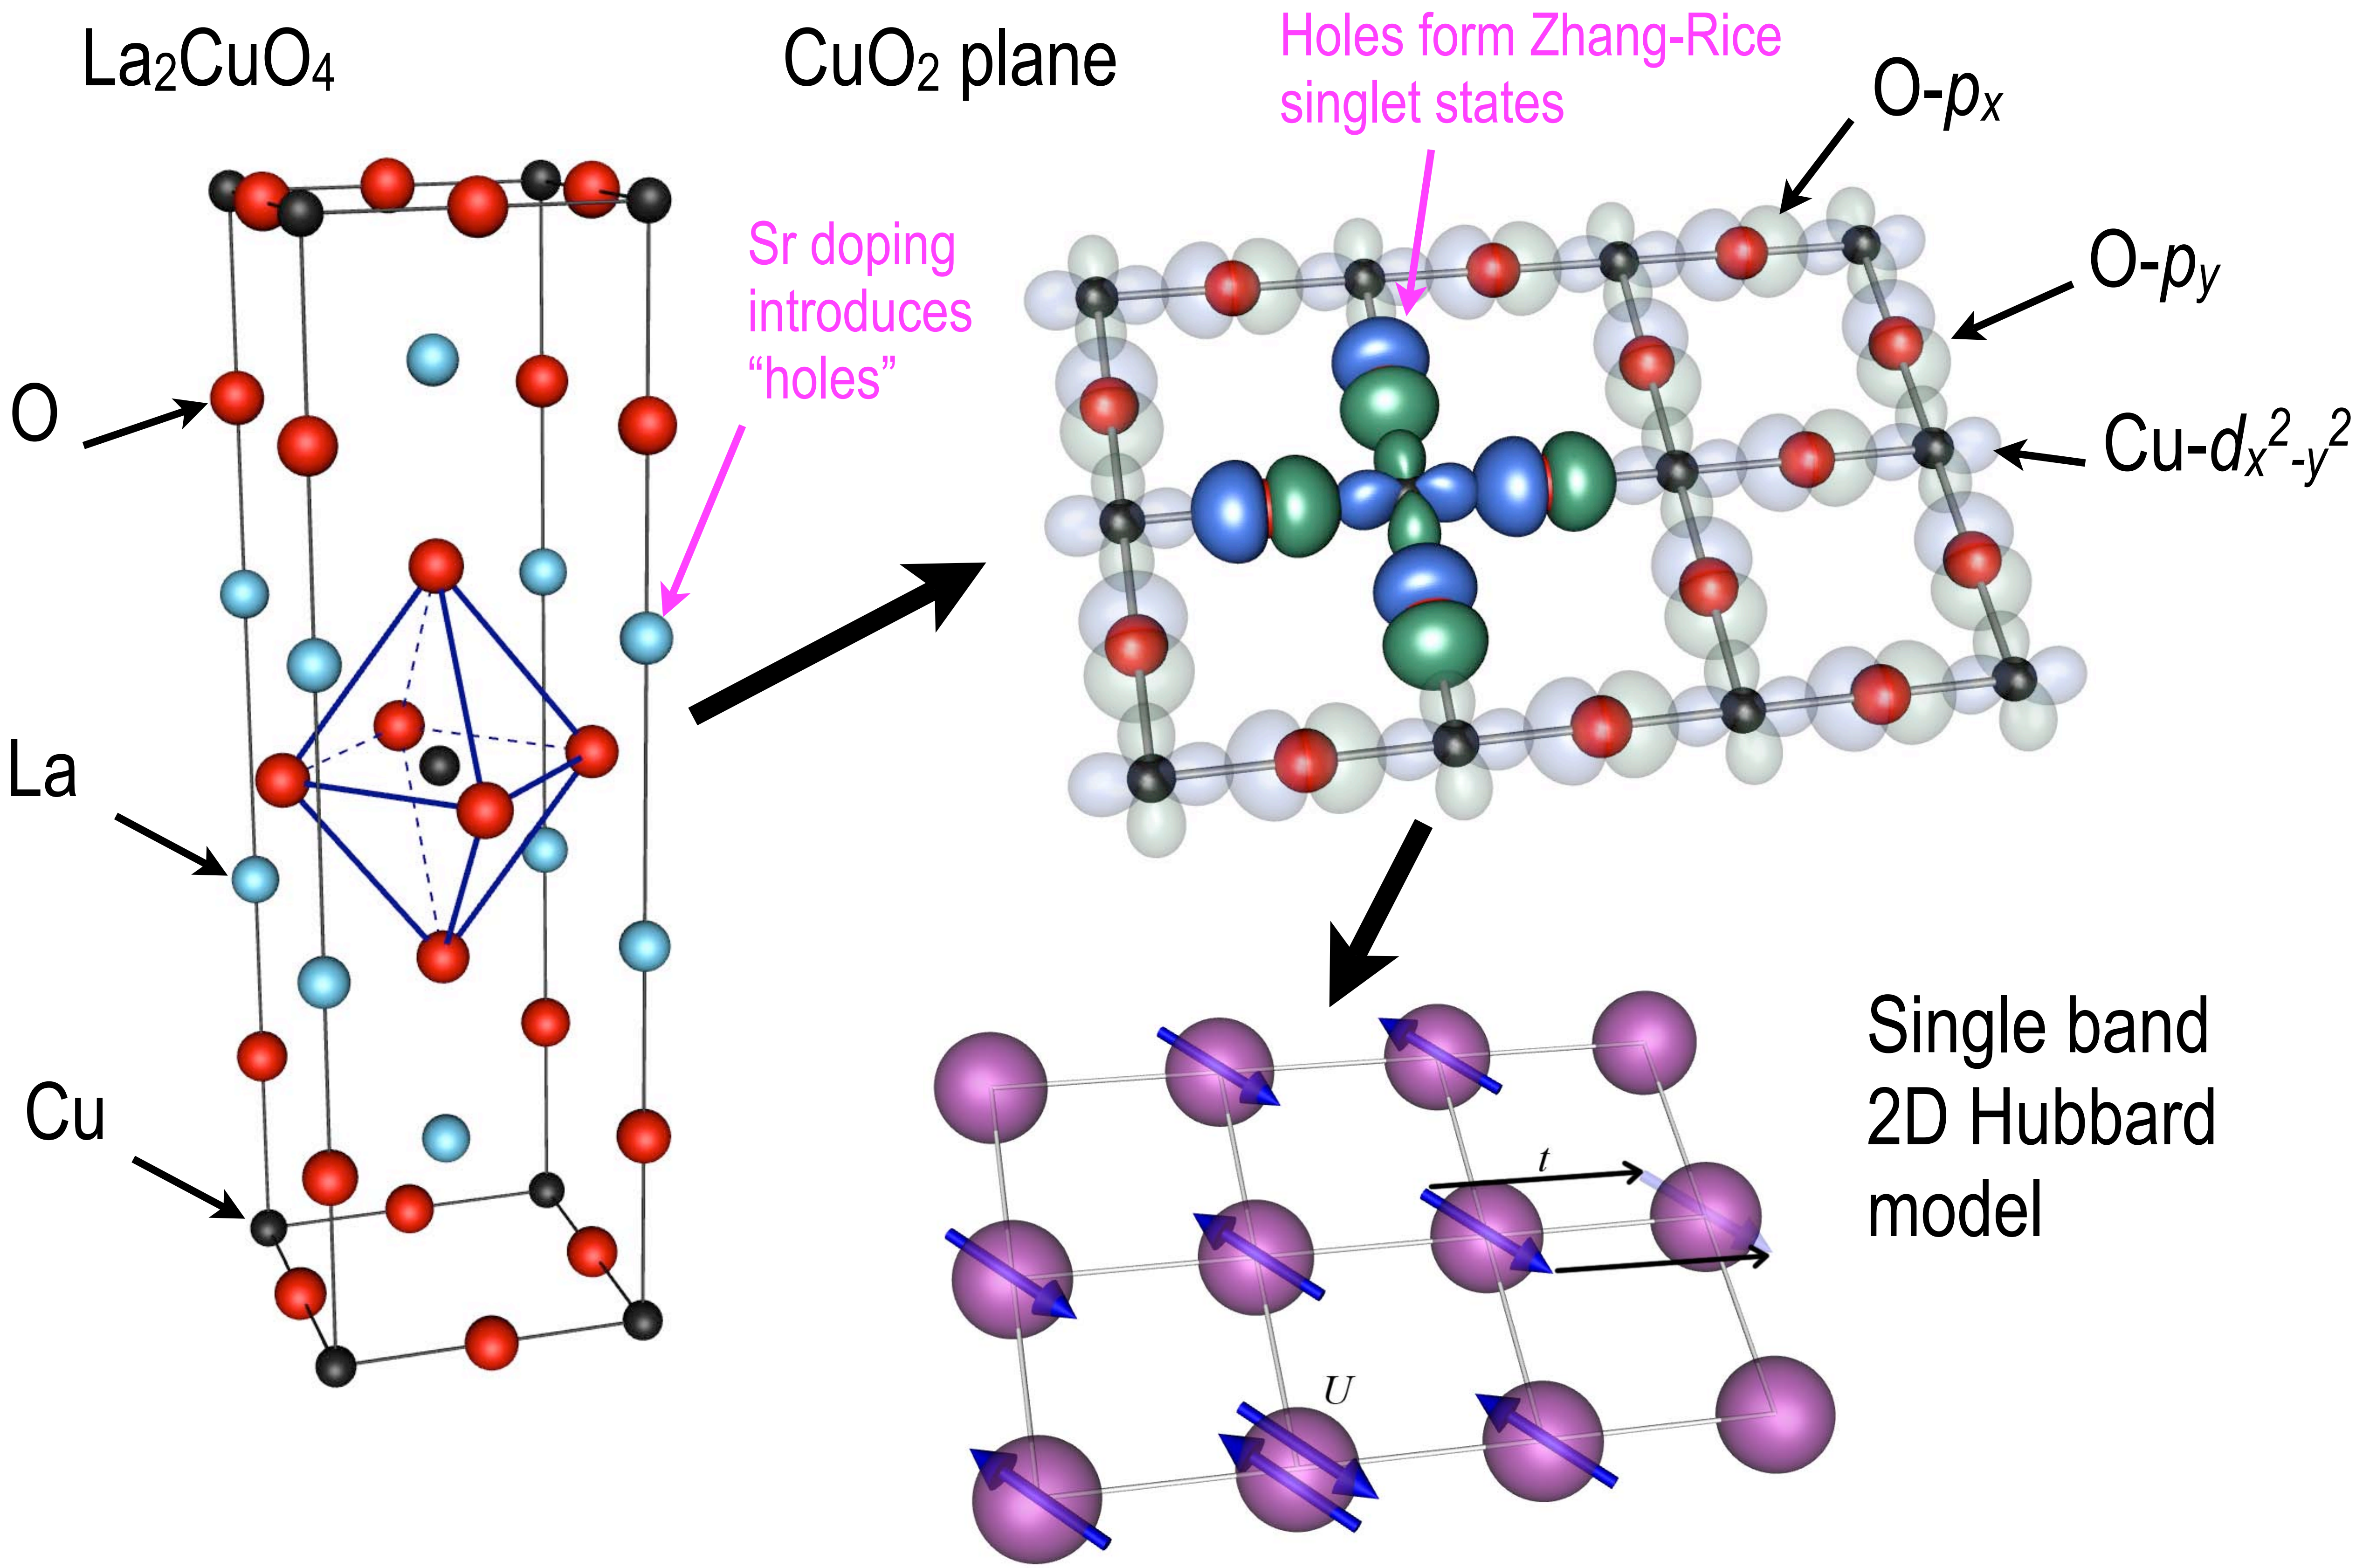

2D Hubbard model and its physics

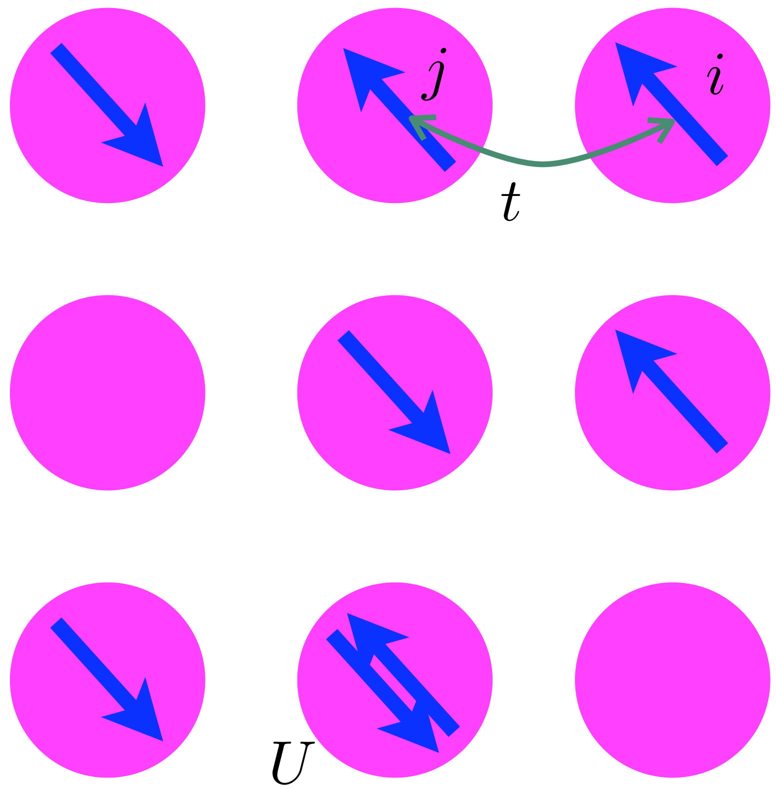

Half filling: number of carriers = number of sites

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

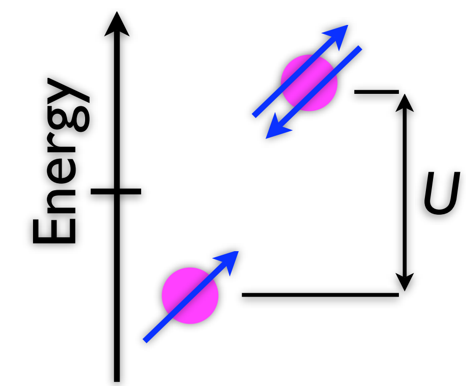

Antiferromagnetic alignment of neighboring moments

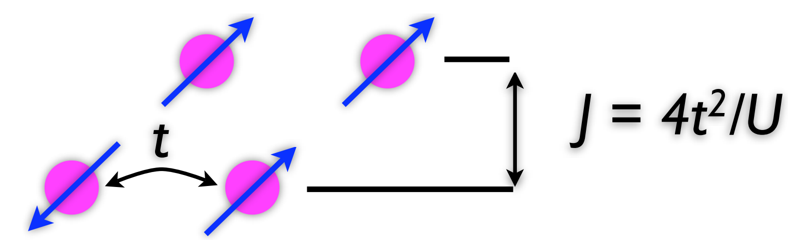

1. When $t >> U$:

Model describes a metal with band width W=8t

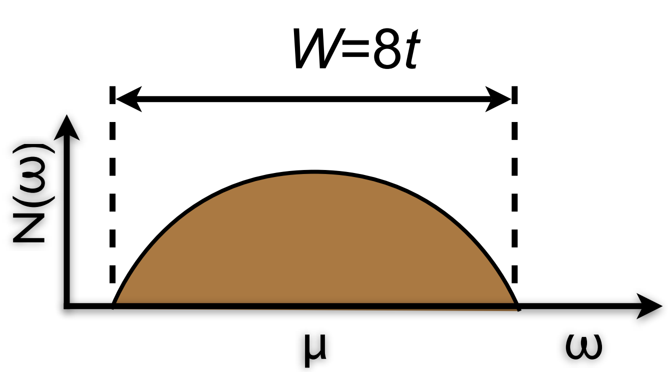

2. When $U >> 8t$ at half filling (not doped)

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

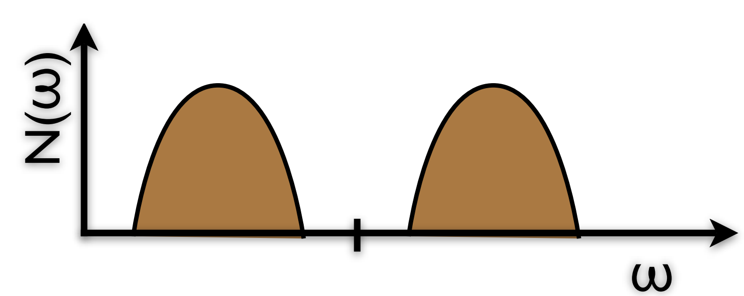

Hubbard model for the cuprates

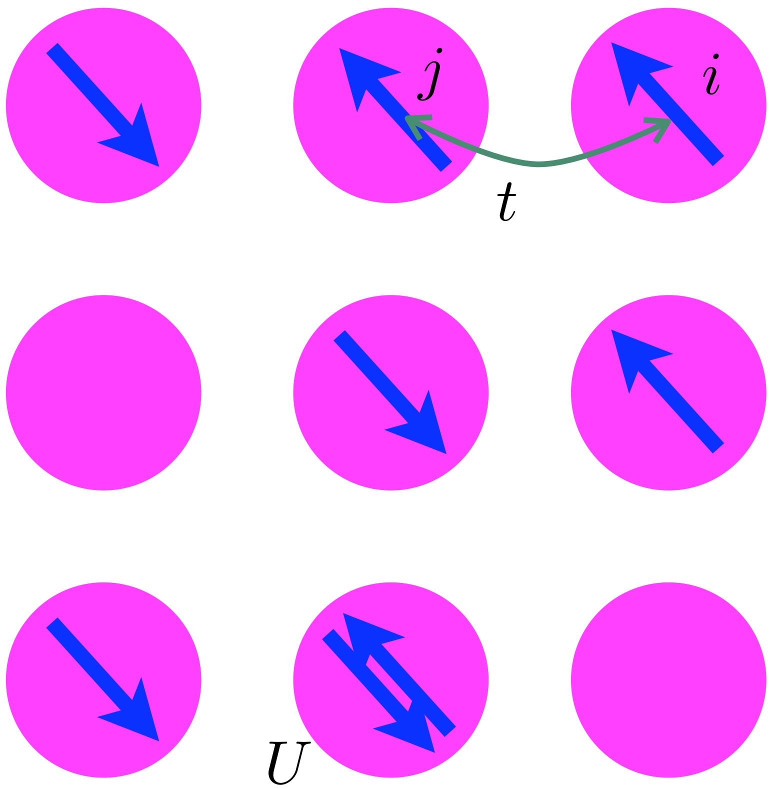

Half filling: number of carriers = number of sites

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

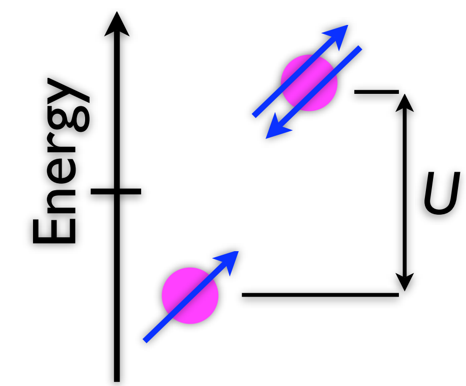

Antiferromagnetic alignment of neighboring moments

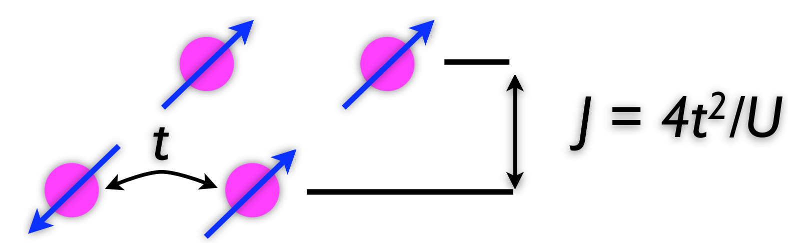

3. Parameter range relevant for superconducting cuprates

$U \approx 8t$

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

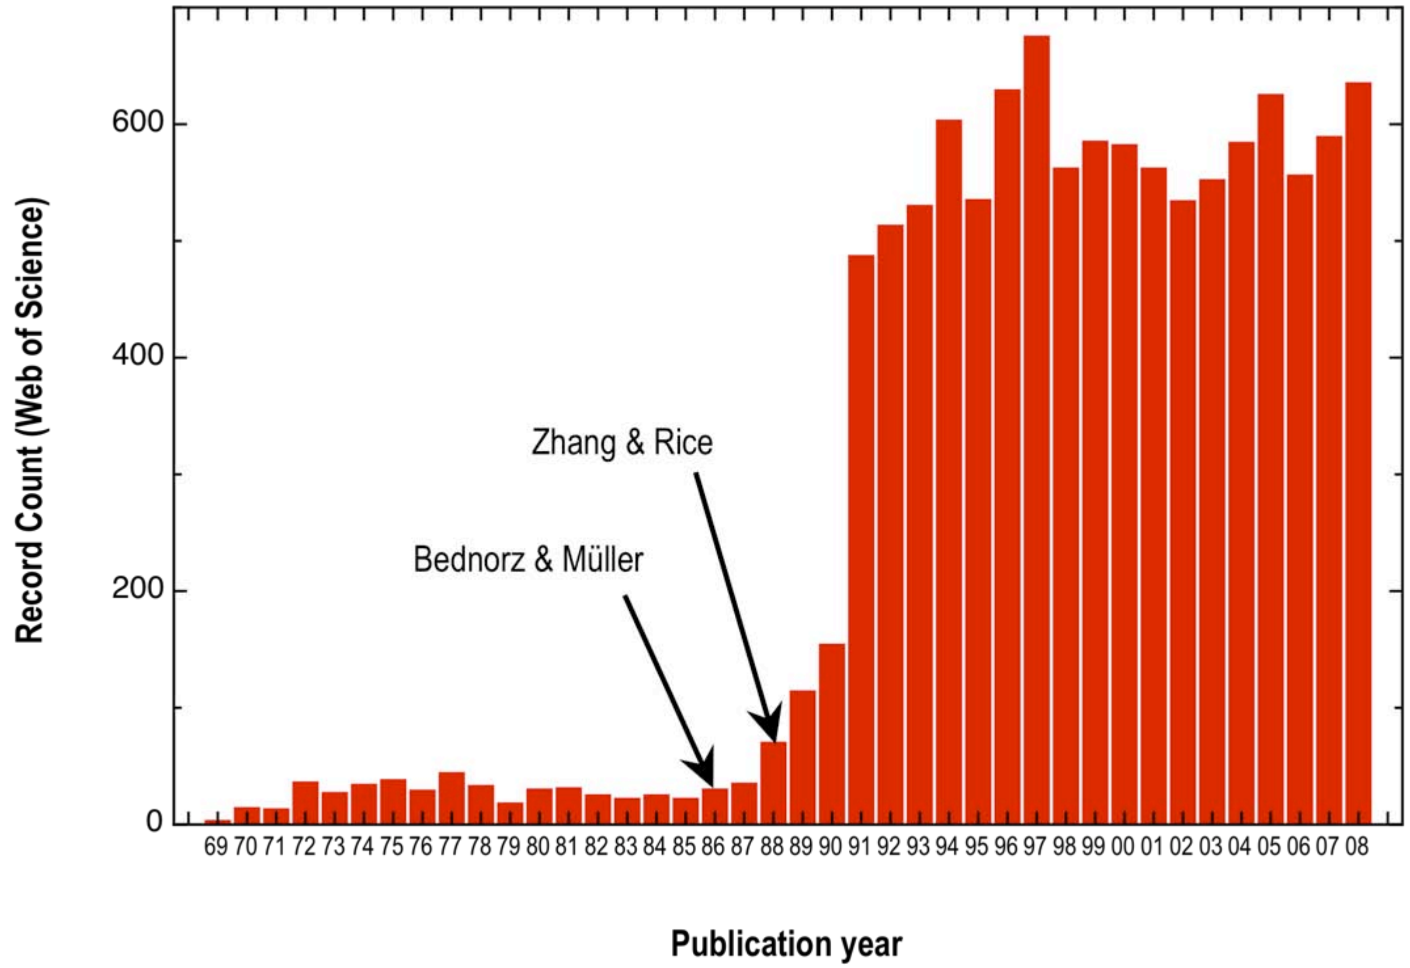

3. Parameter range relevant for superconducting cuprates

$U \approx 8t$

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

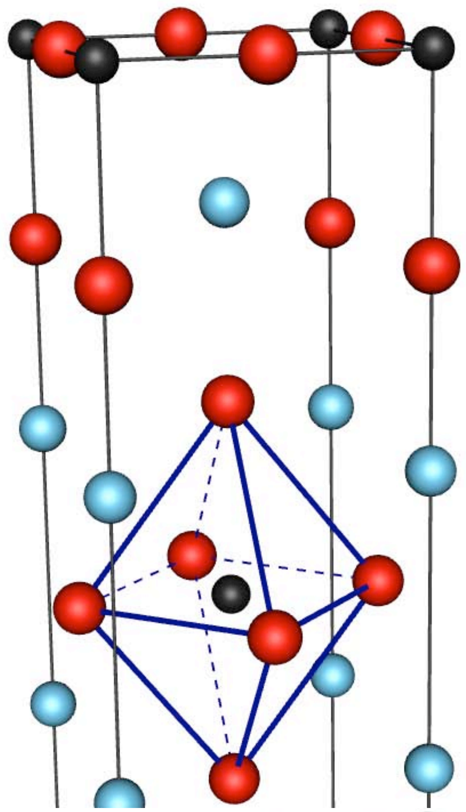

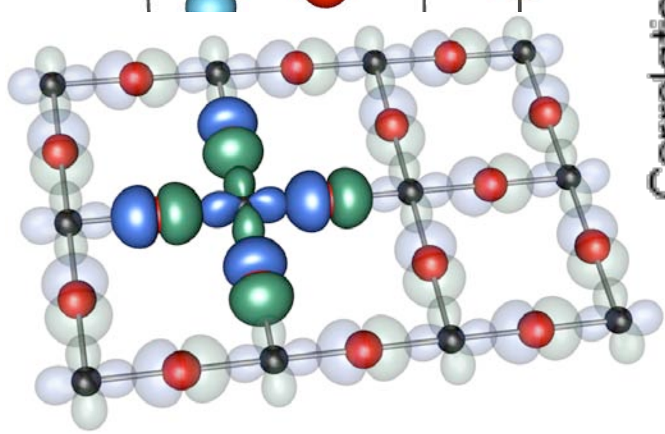

On-site Coulomb repulsion (~A)

Antiferromagnetic correlations / nano-scale gap fluctuations

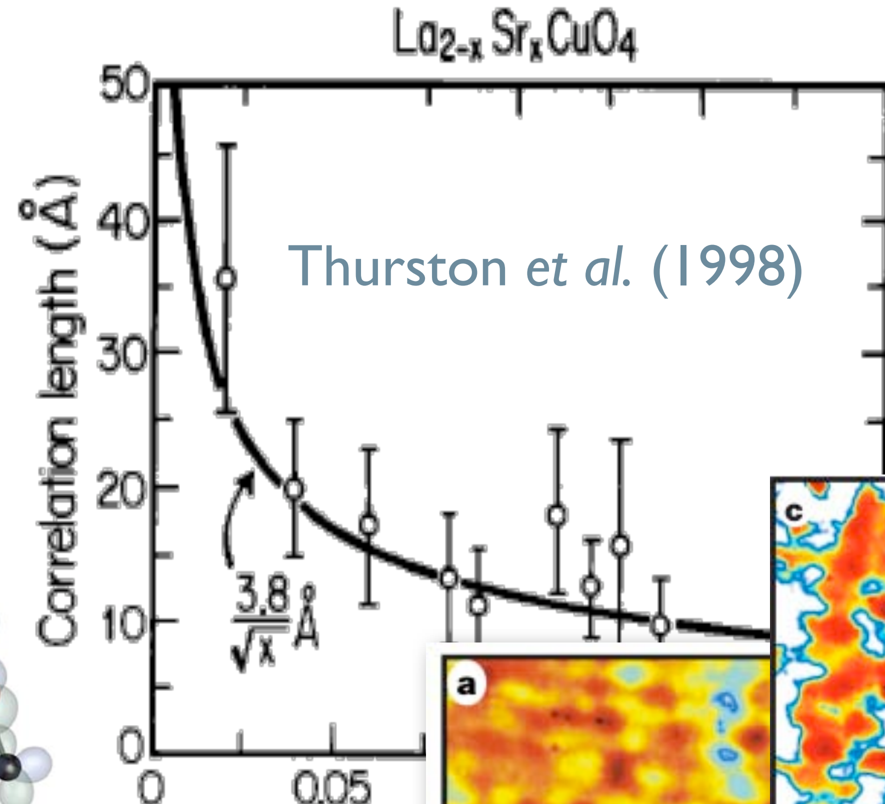

Thurston *et al.* (1998)

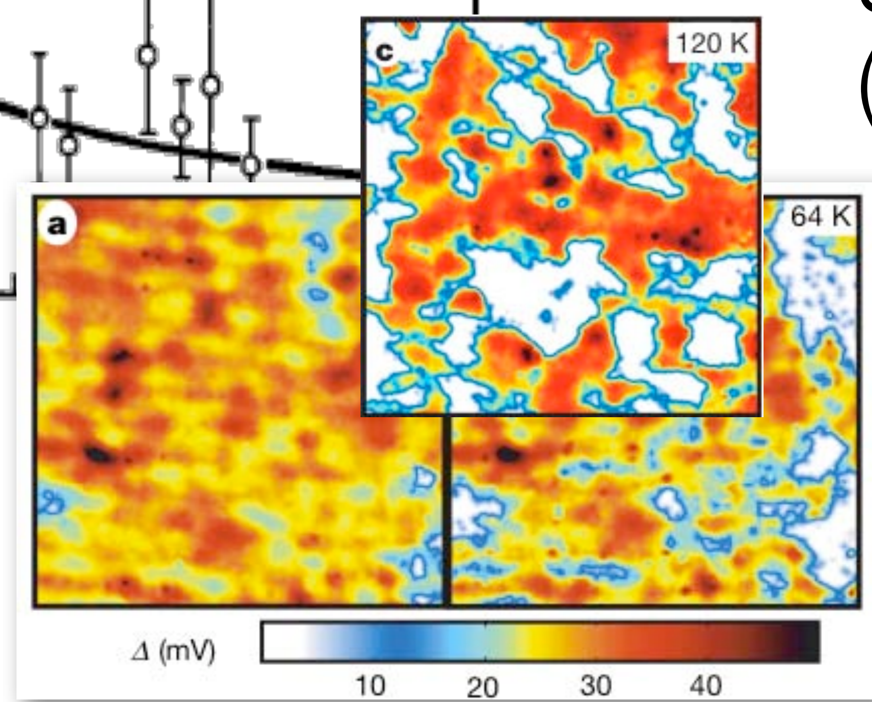

Gomes et al. (2007)

Superconductivity (macroscopic)

$N \sim 10^{23}$

complexity $\sim 4^N$

Quantum cluster theories

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

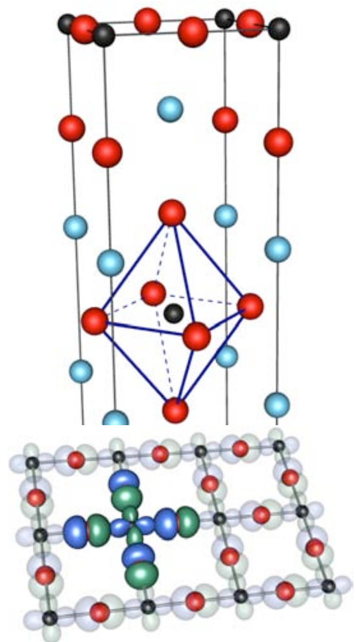

On-site Coulomb repulsion (~A)

Antiferromagnetic correlations / nano-scale gap fluctuations

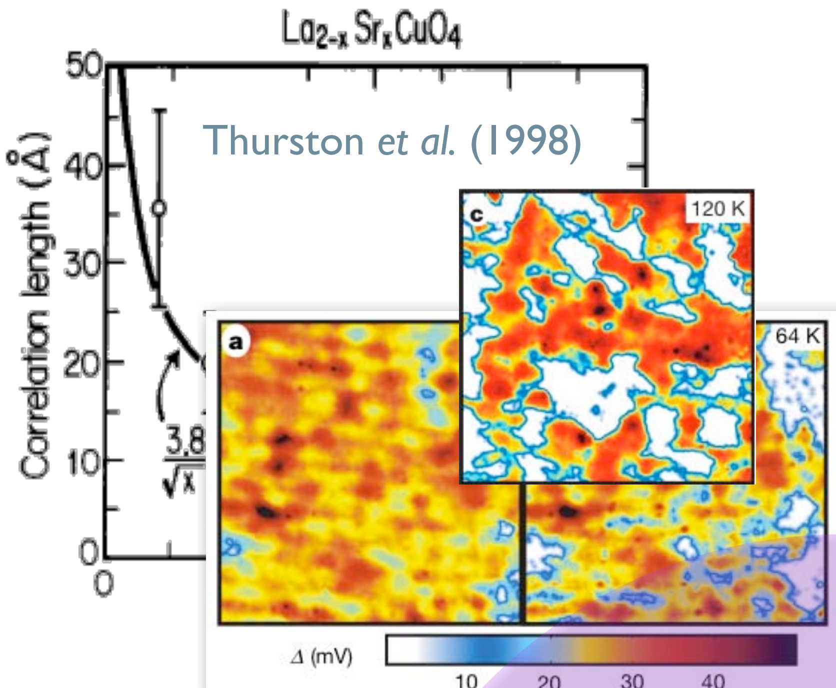

Gomes et al. (2007)

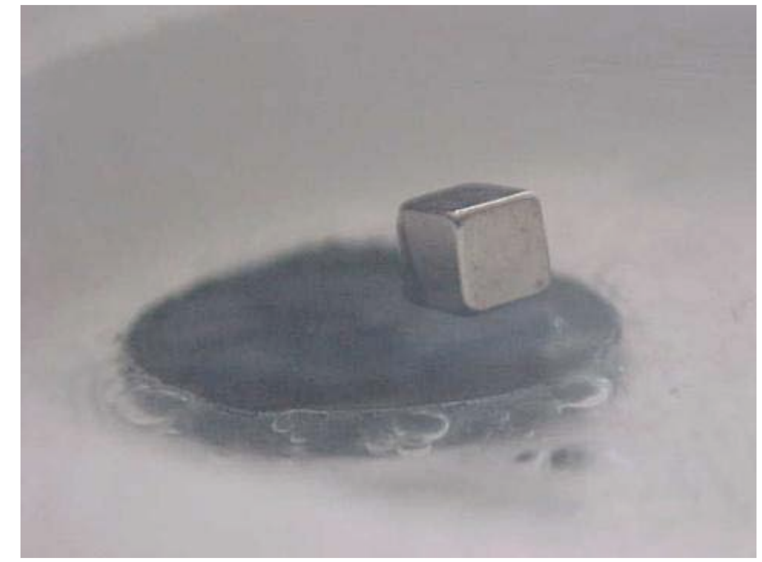

Superconductivity (macroscopic)

Explicitly treat correlations within a localized cluster

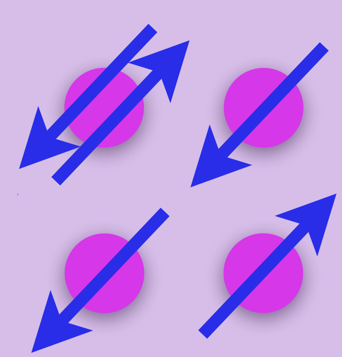

Treat macroscopic scales within mean-field

Coherently embed cluster into effective medium

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

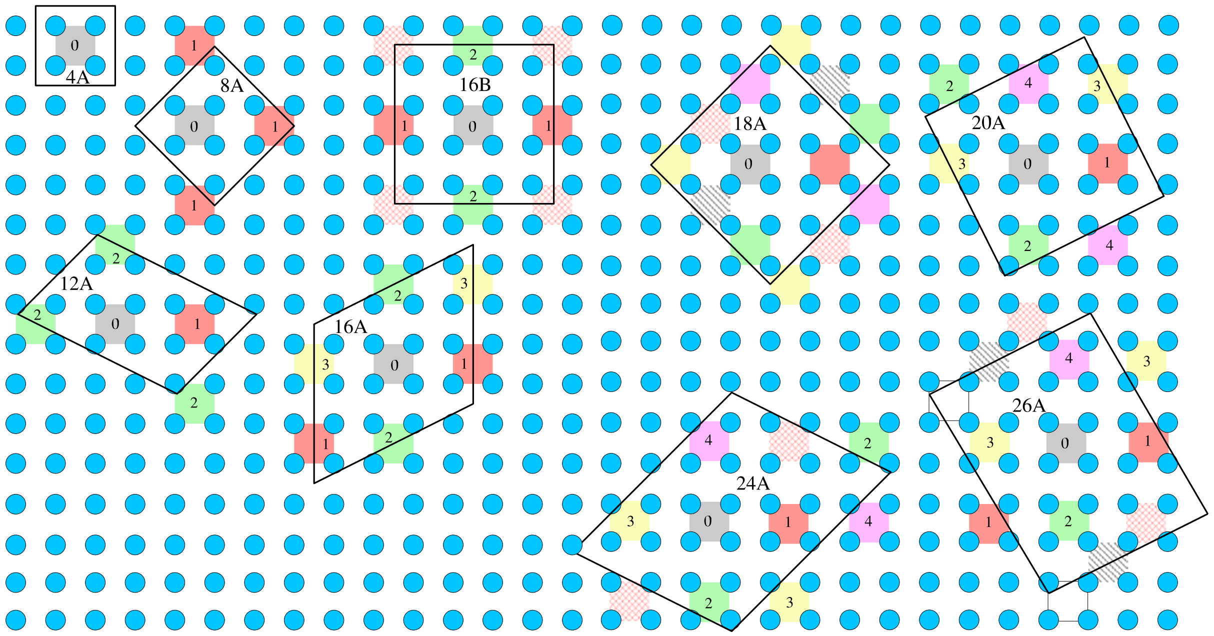

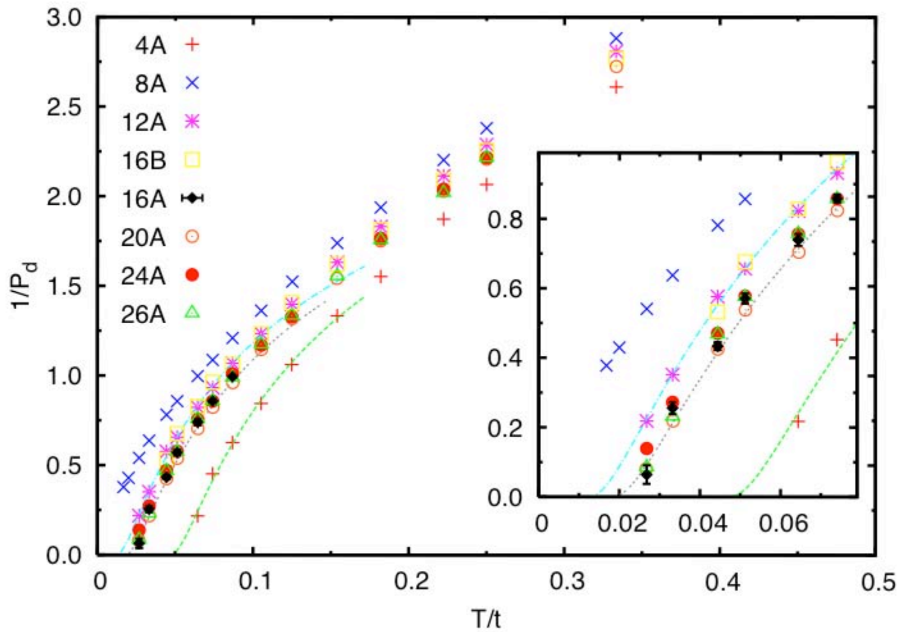

- Study the mechanism responsible for pairing in the model
  - Analyze the particle-particle vertex
  - Pairing is mediated by spin fluctuations
    Maier, et al., Phys. Rev. Lett. **96** 47005 (2006)

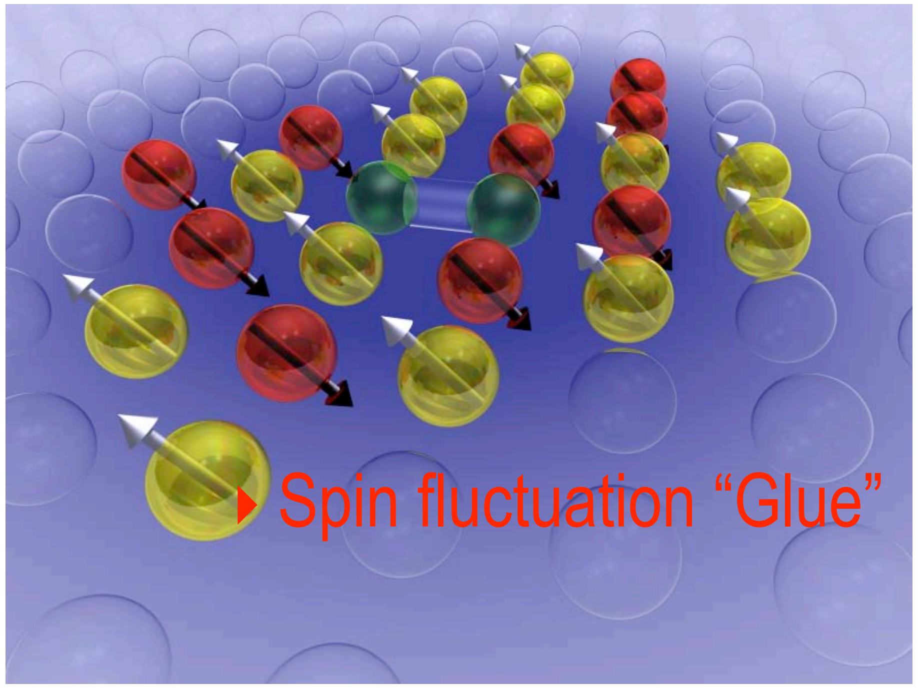

Moving toward a resolution of the debate over the pairing mechanism in the 2D Hubbard model

- "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 mechanisms
  - Maier et al., Phys. Rev. Lett. **100** 237001 (2008)

Fraction of superconducting gap arising from frequencies ≤ Ω

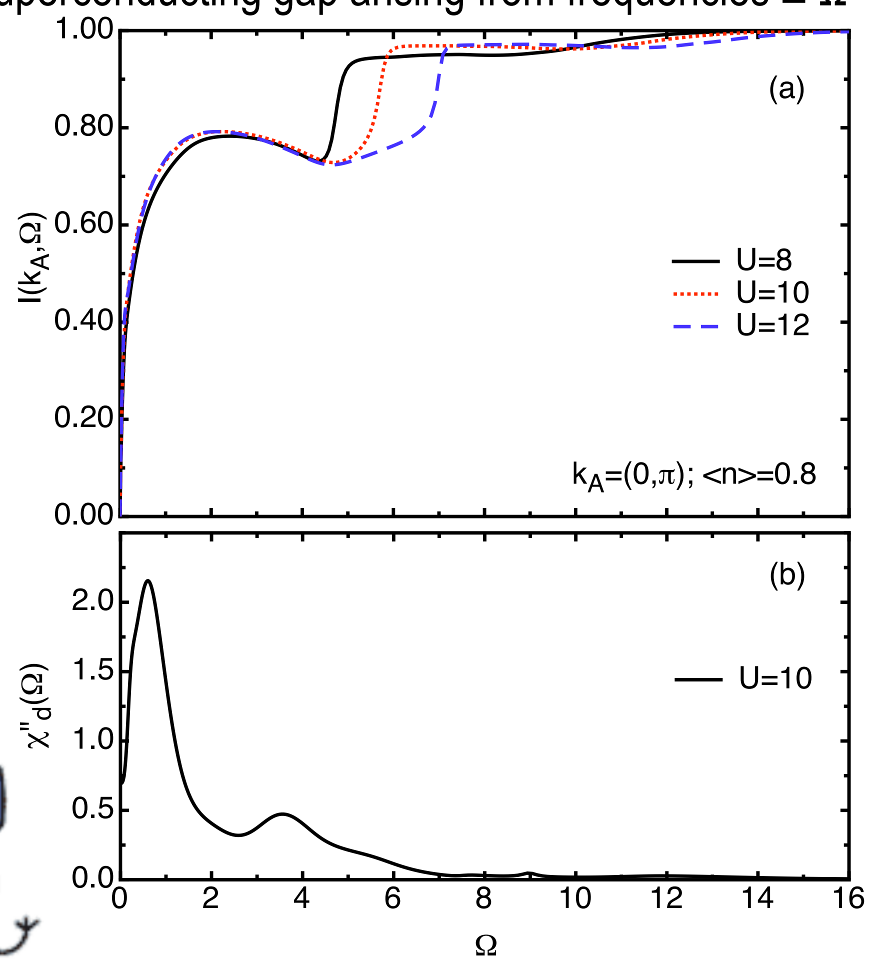

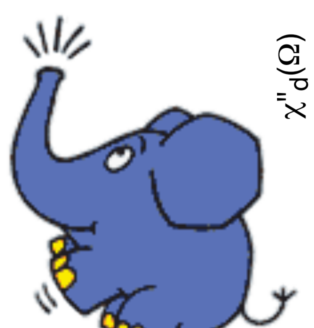

Both retarded spin-fluctuations and non-retarded exchange interaction J contribute to the pairing interaction

Dominant contribution comes from spin-fluctuations!

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

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

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

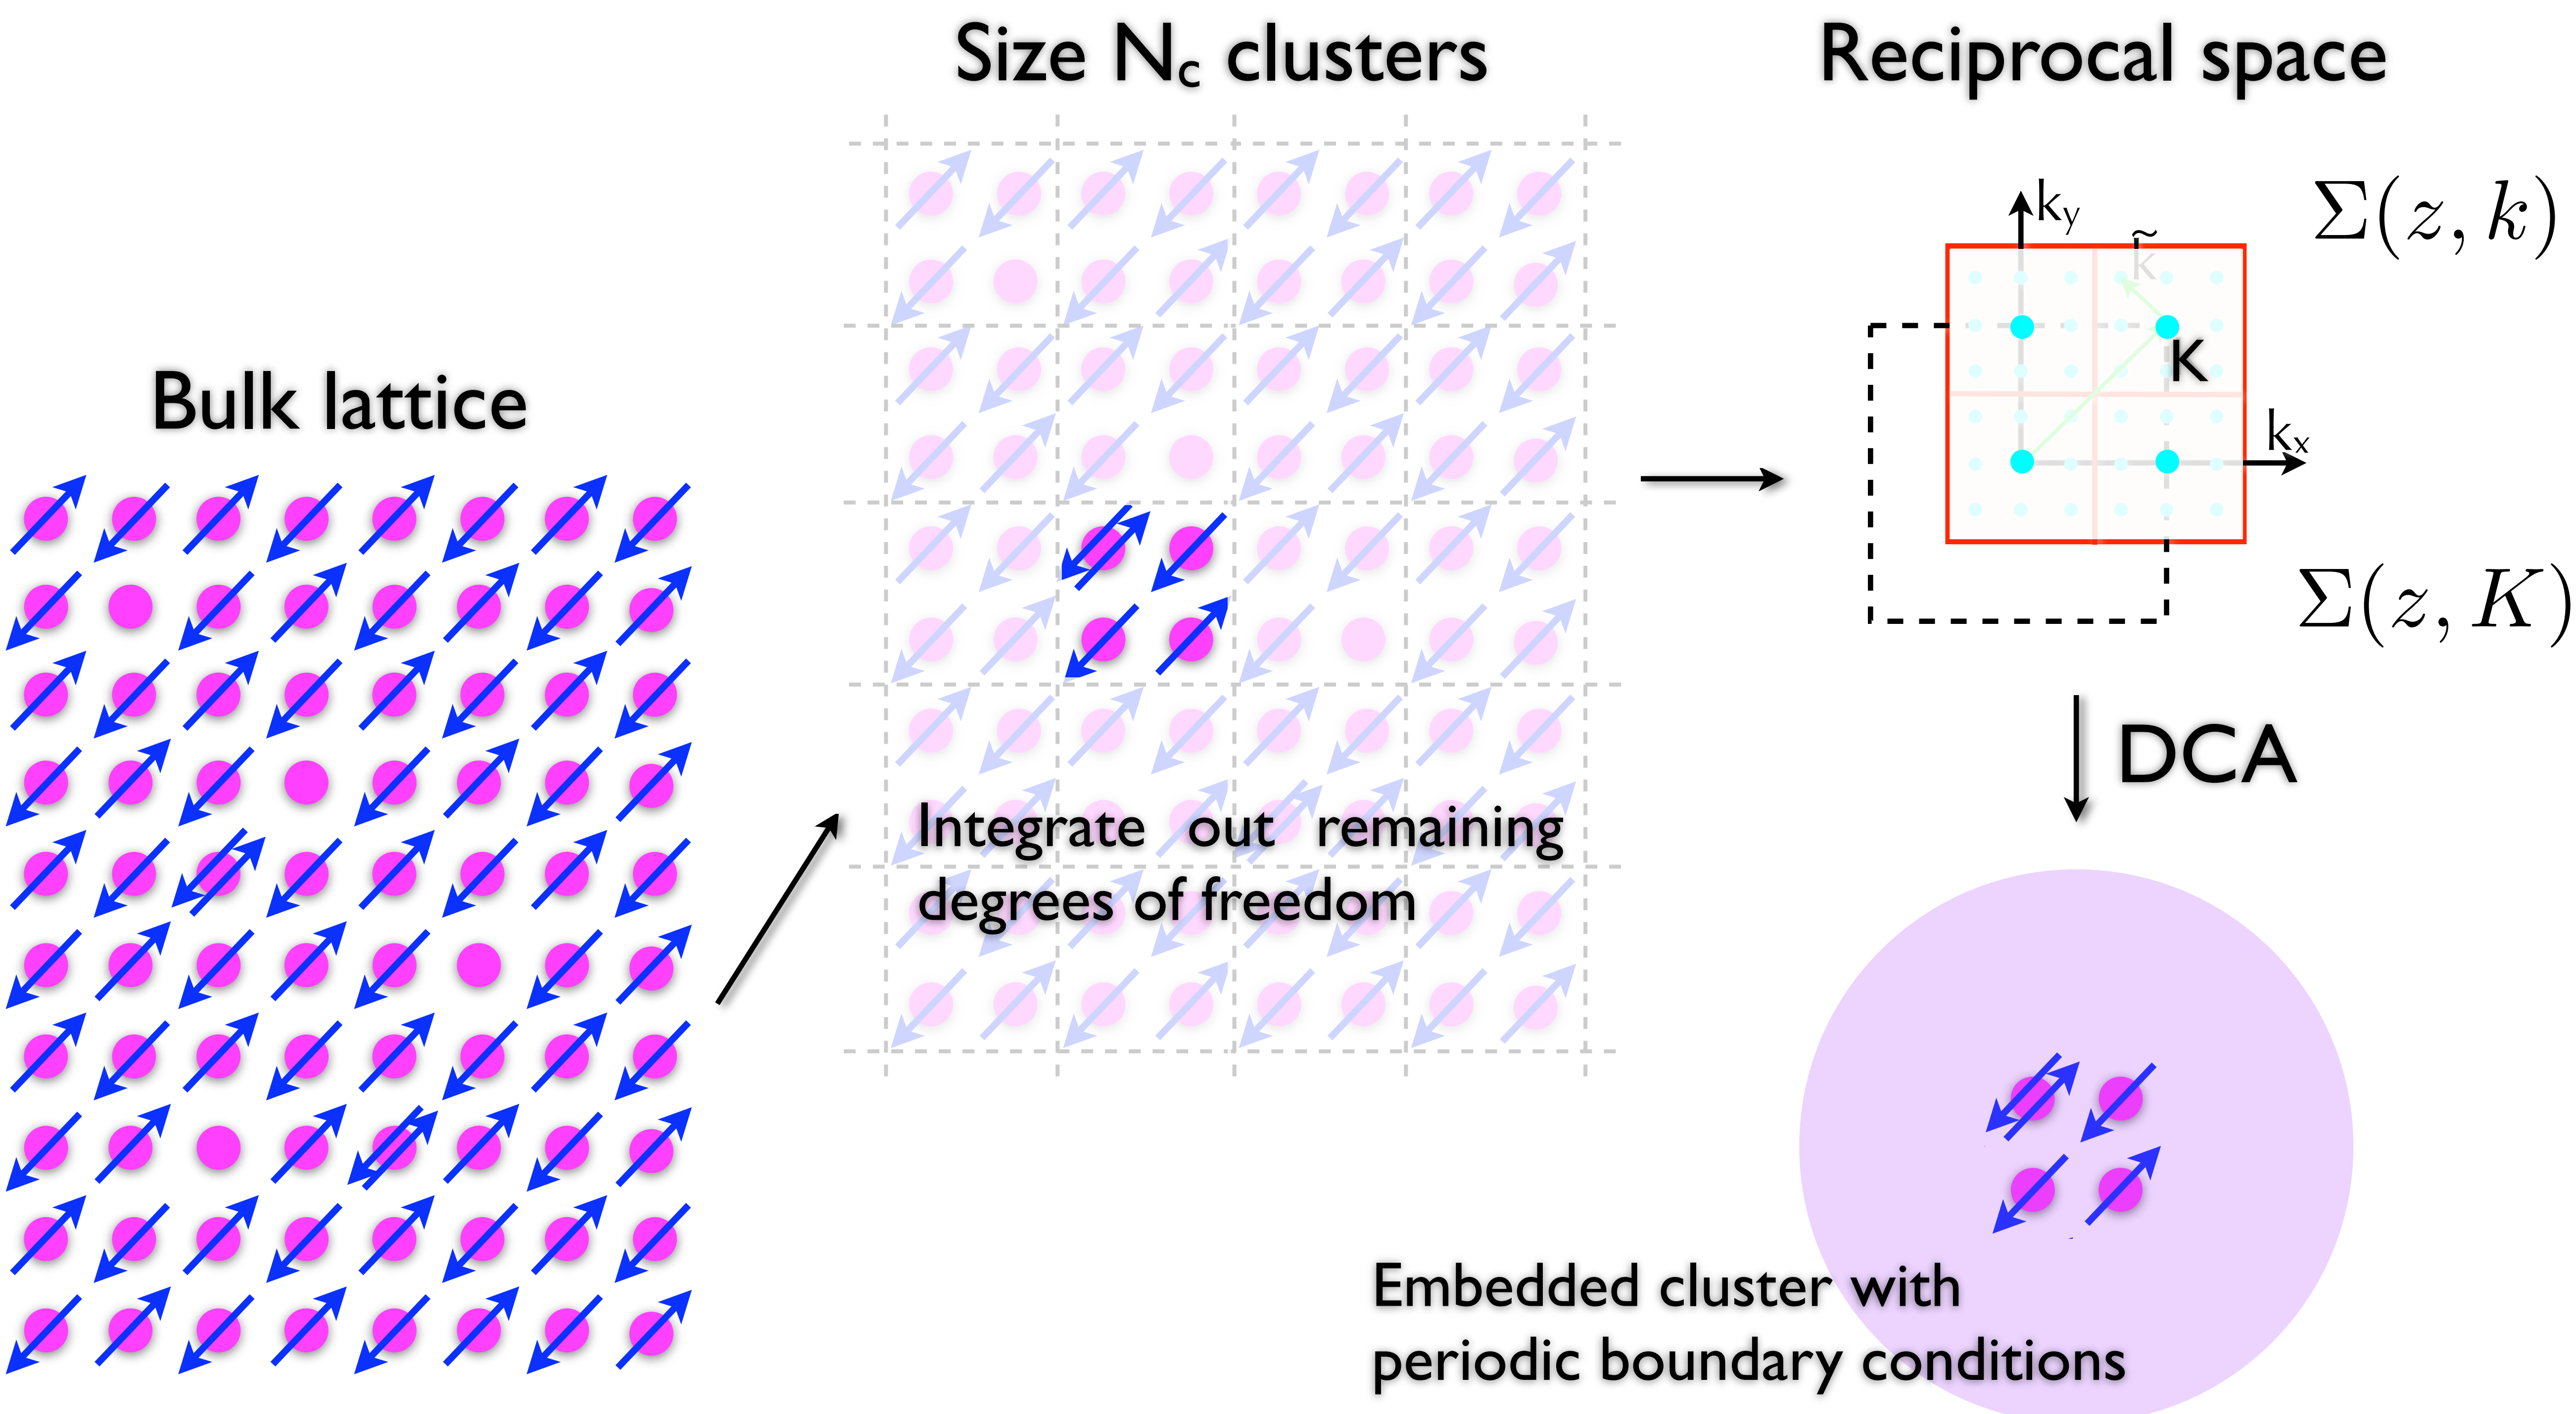

Solve many-body problem with quantum Monte Carole on cluster

➢Essential assumption: Correlations are short ranged

DCA method: self-consistently determine the “effective” medium

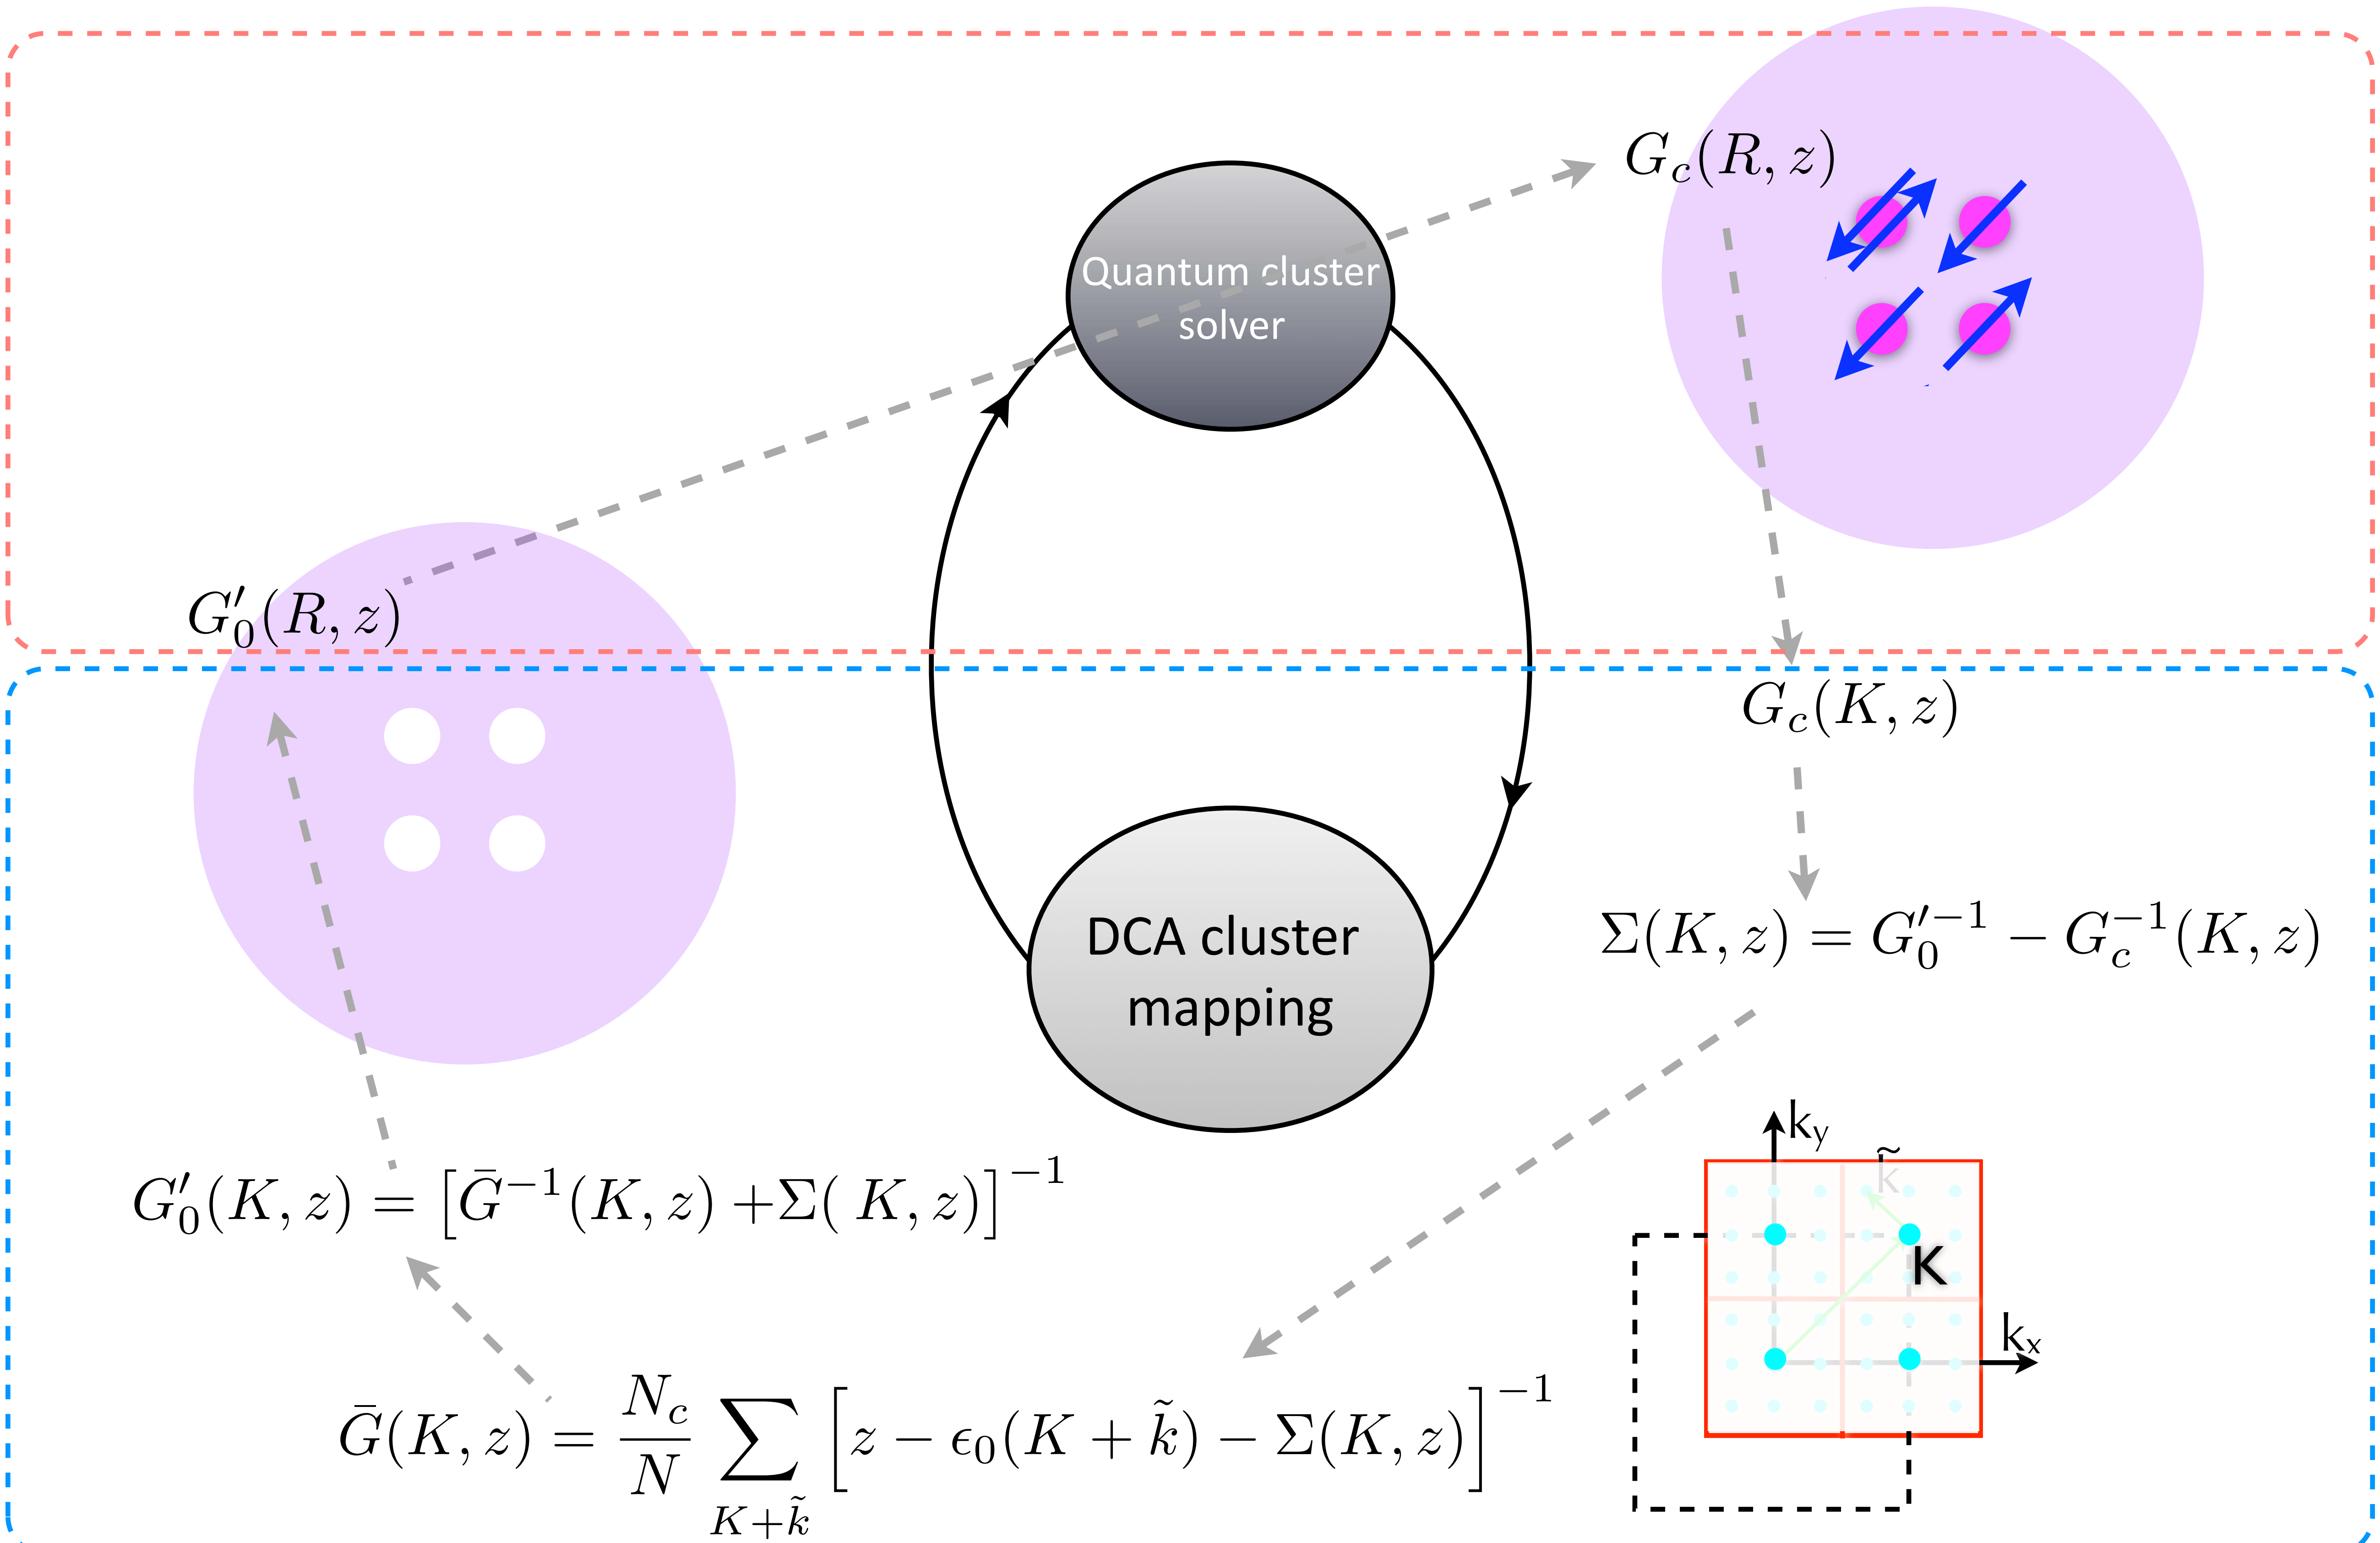

Structure of DCA++ code: generic programming

DCA++

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

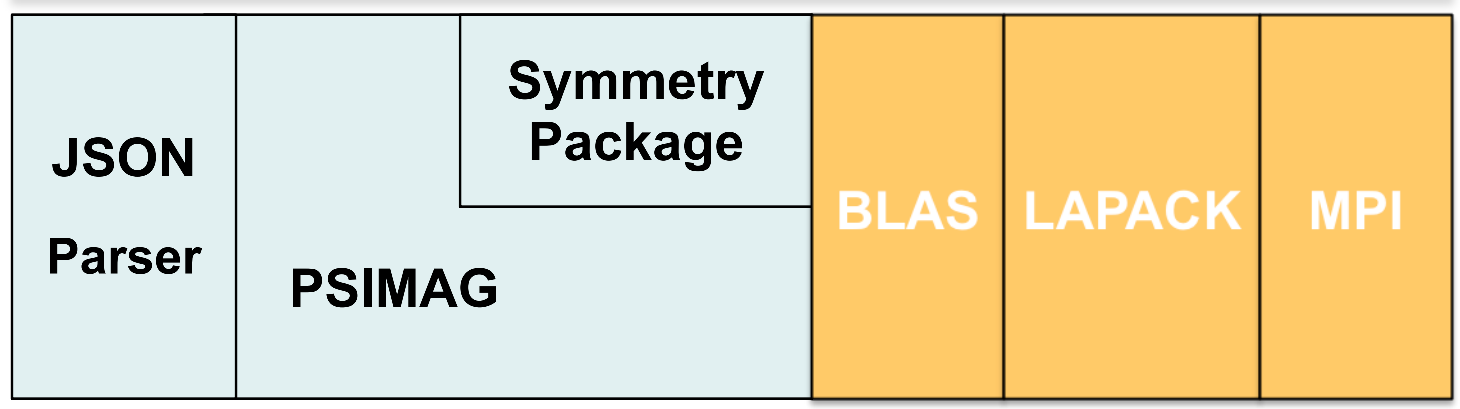

<u>PsiMag Implementation philosophy:</u>
Consider PsiMag as a systematic extension to the C++ Standard Template Library (STL) using as much as possible the **generic programming paradigm**

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 $N_l \approx 10^2$

matrix of dimensions $N_t \times N_t$

$N_t = N_c \times N_l \approx 2000$

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

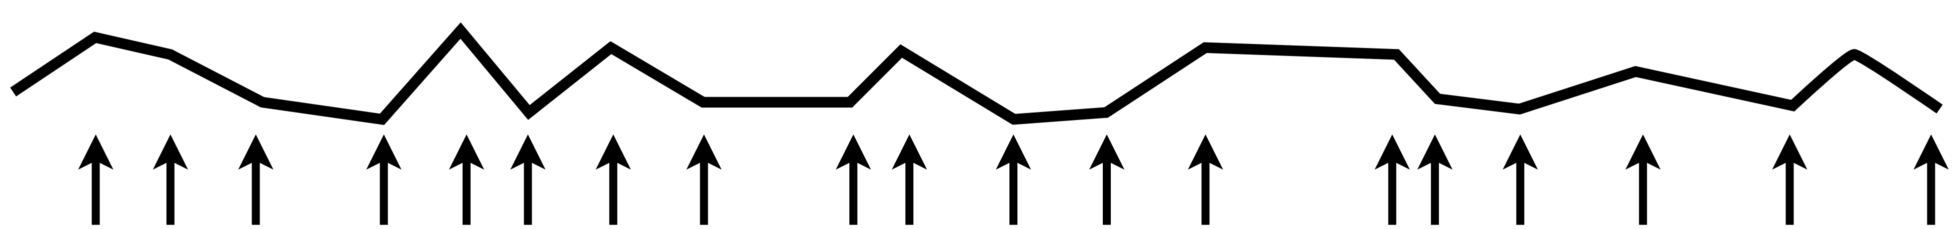

Update of accepted Green's function:

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

HF-QMC with Delayed updates (or Ed updates)

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

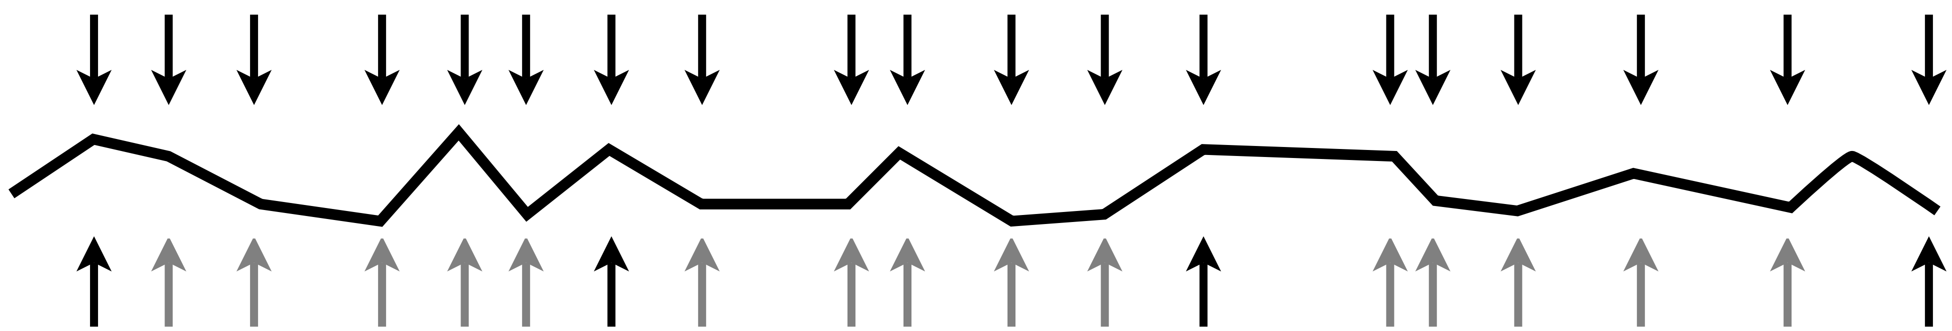

$$\mathbf{G}_c(\{s_i,l\}_{k+1}) = \mathbf{G}_c(\{s_i,l\}_0) + [\mathbf{a}_0|\mathbf{a}_1|...|\mathbf{a}_k] \times [\mathbf{b}_0|\mathbf{b}_1|...|\mathbf{b}_k]^t$$

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

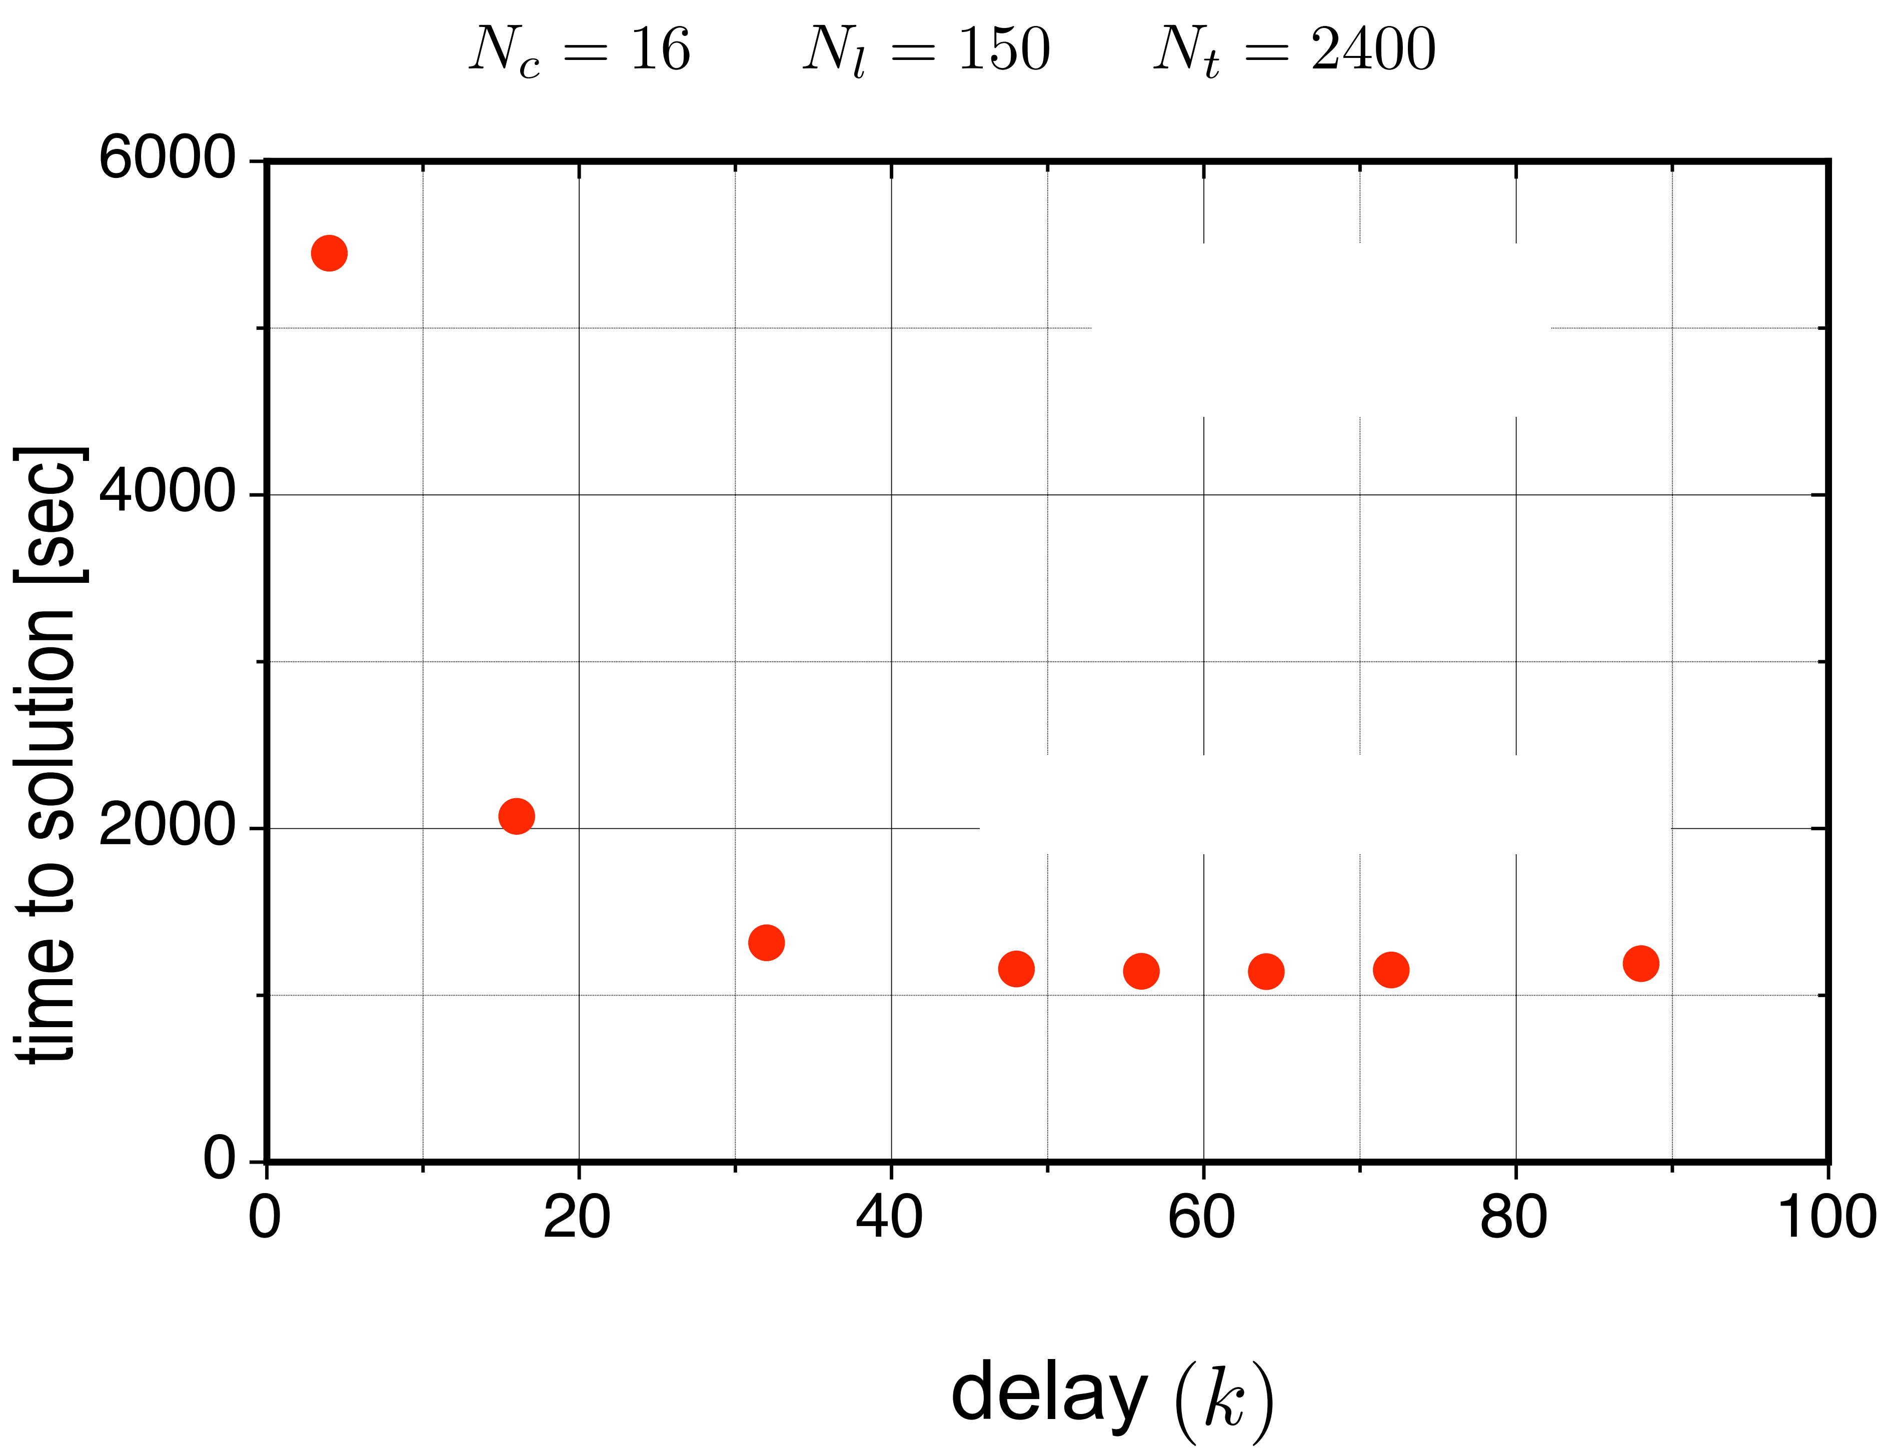

MultiCore/GPU/Cell: threaded programming

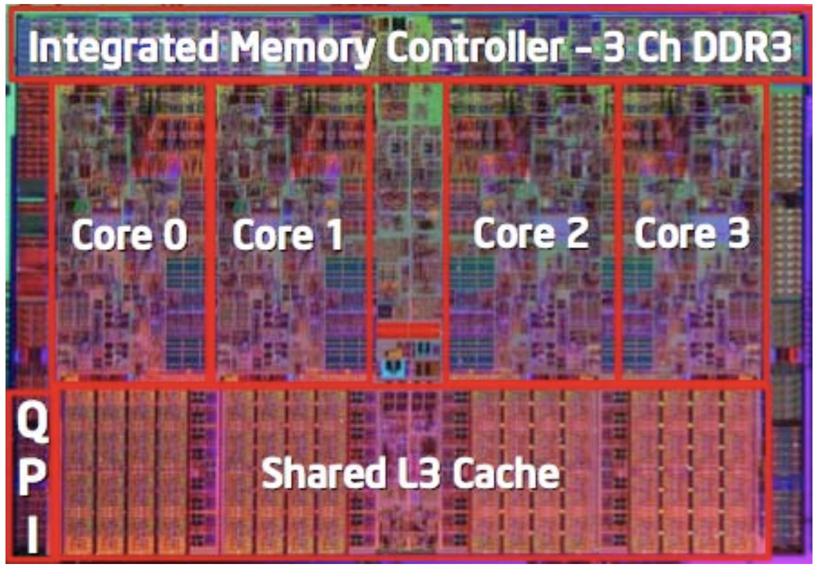

Multi-core processors: OpenMP (or just MPI)

NVIDIA G80 GPU: CUDA, cuBLAS

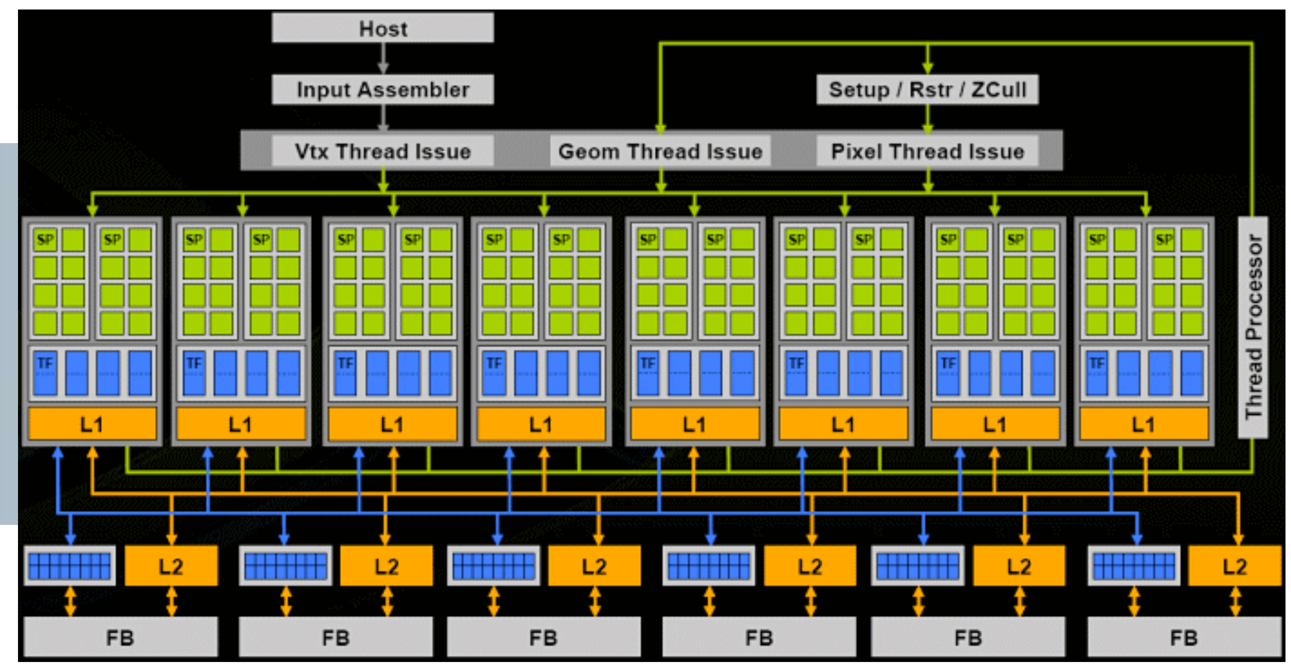

128 streaming processors
350 usable GFlop/s at 575 MHz
100 GB/s internal memory bandwidth
CUDA runtime API
cuBLAS (single precision)

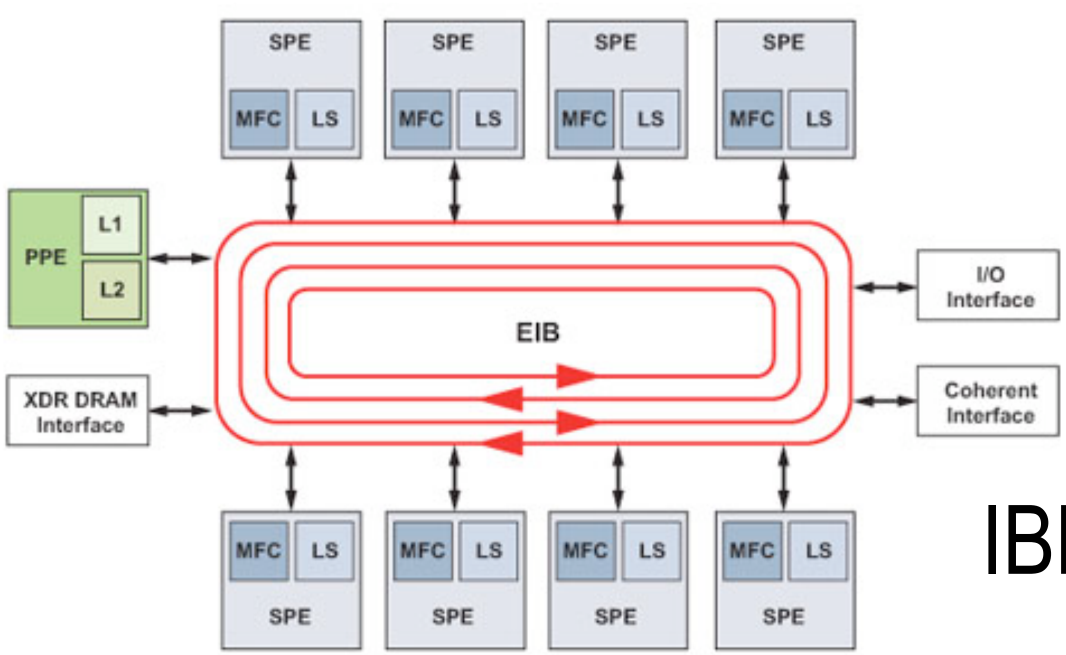

IBM Cell BE: SIMD, threaded prog.

GPU Programming Concepts

- "Streaming"
  - input and output arrays differ
- Data Parallel (SIMD)
  - same code, many times
- Threads to Hide Latency
  - $\sim 10^5$ threads in flight at once
- Gather Semantics
  - Required for good performance

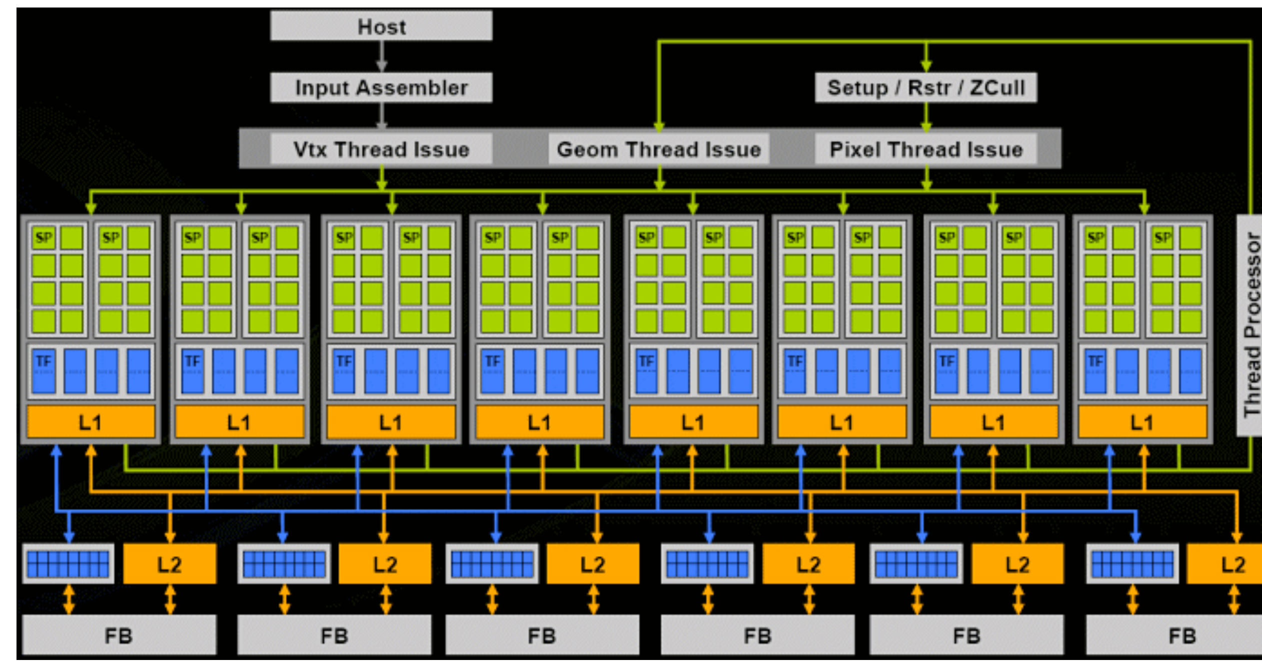

System layout for GPU

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

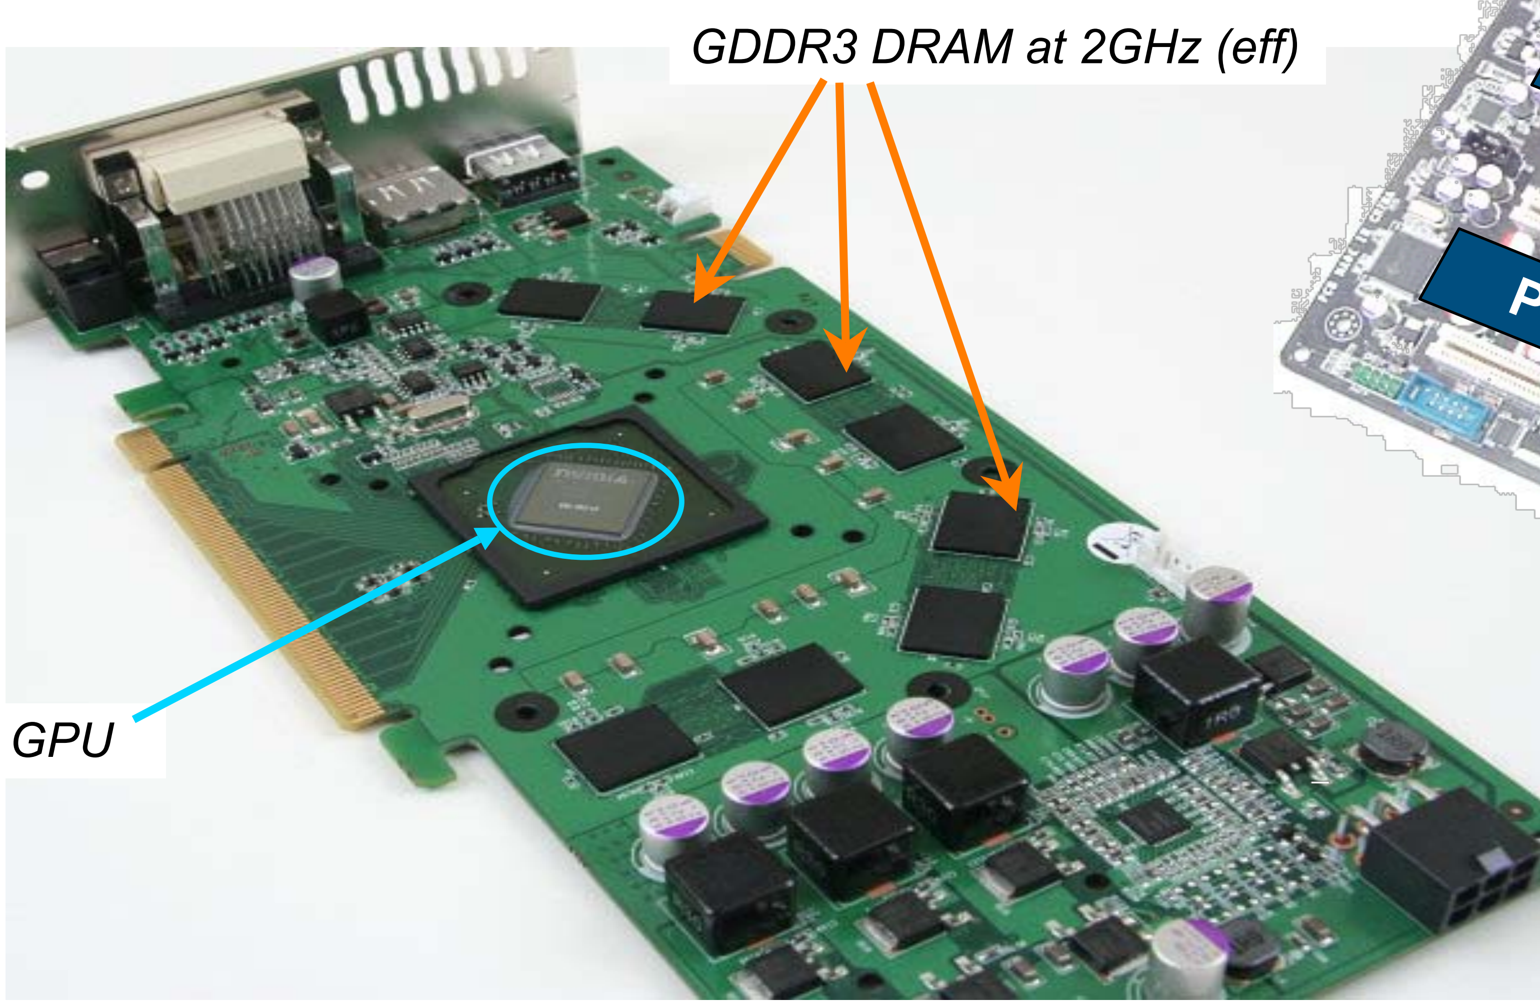

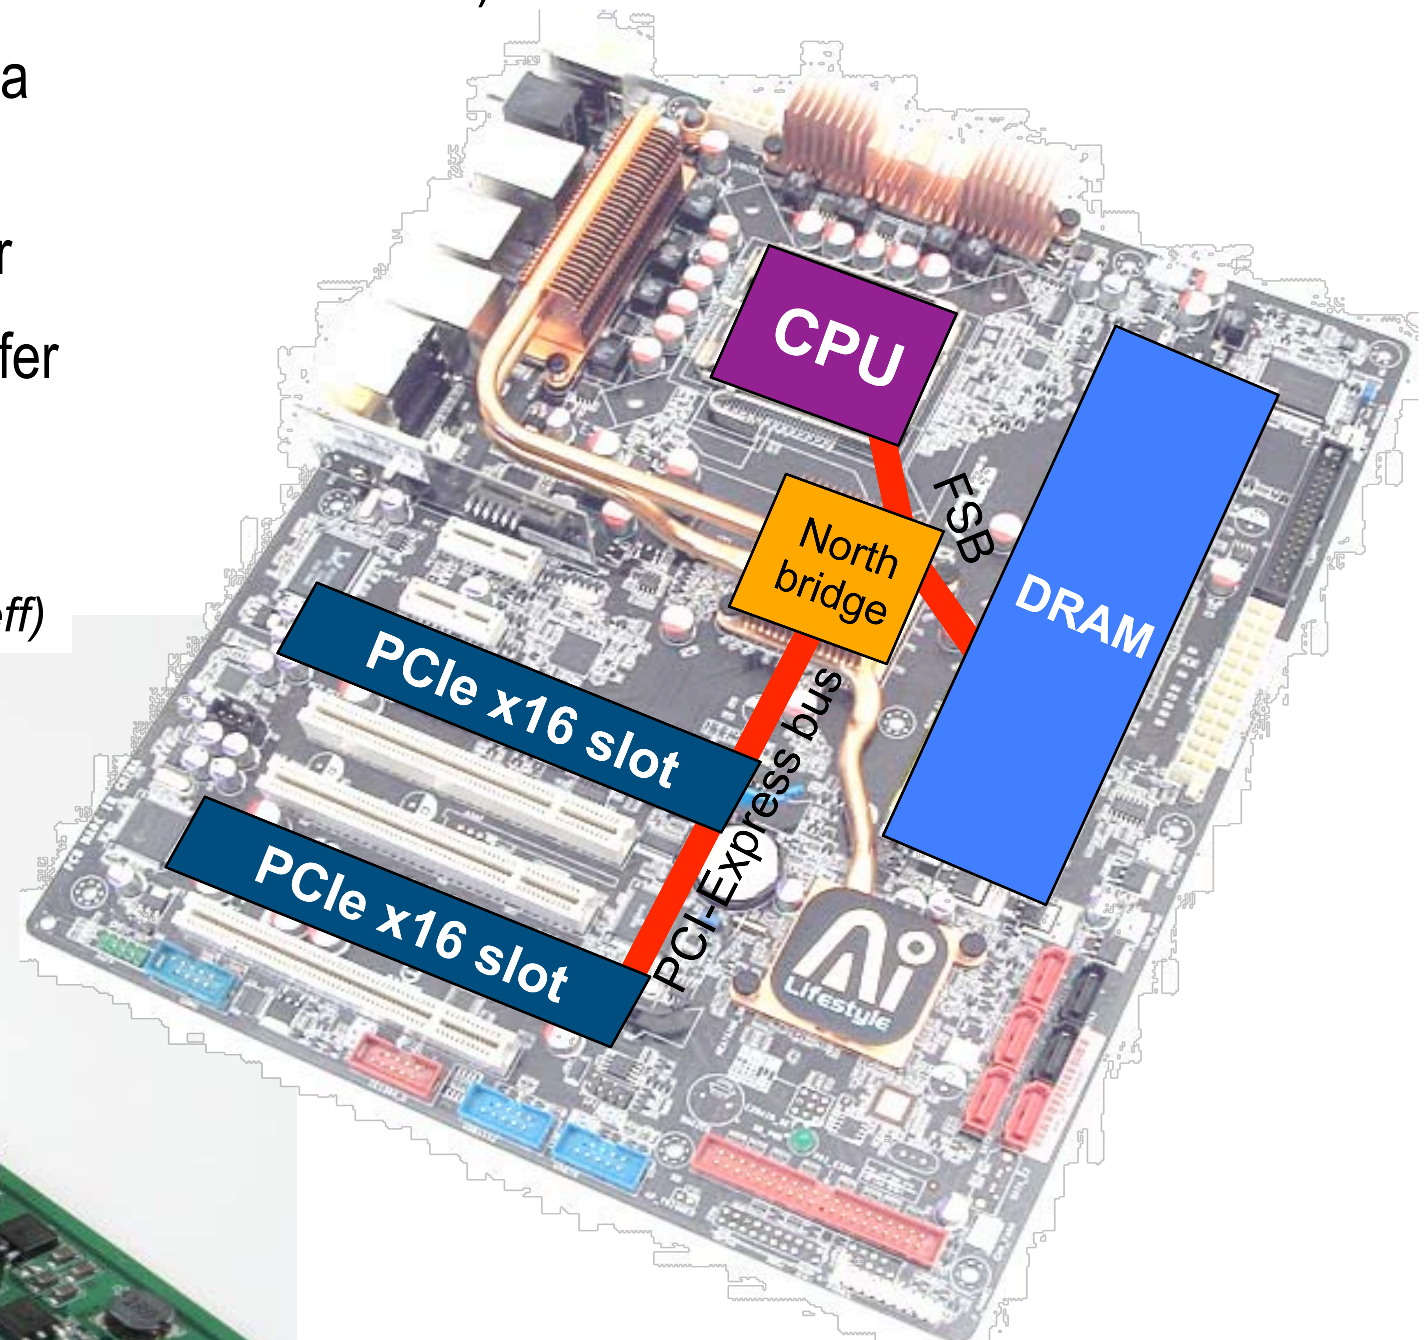

DCA++ with mixed precision

Run HF-QMC in single precision

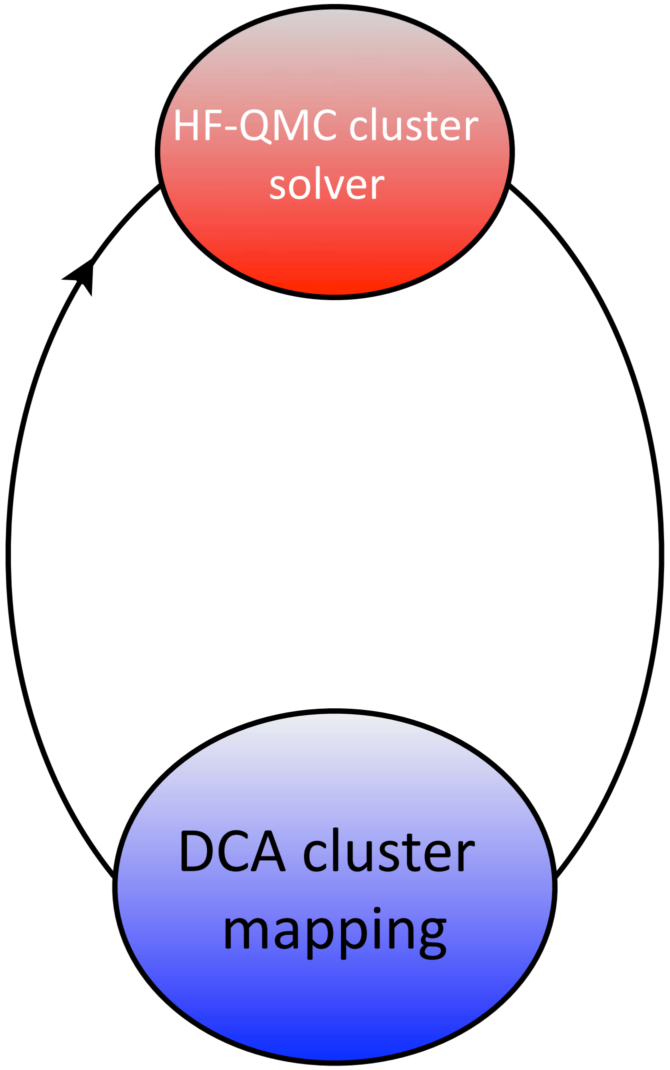

Keep the rest of the code, in particular cluster mapping in double precision

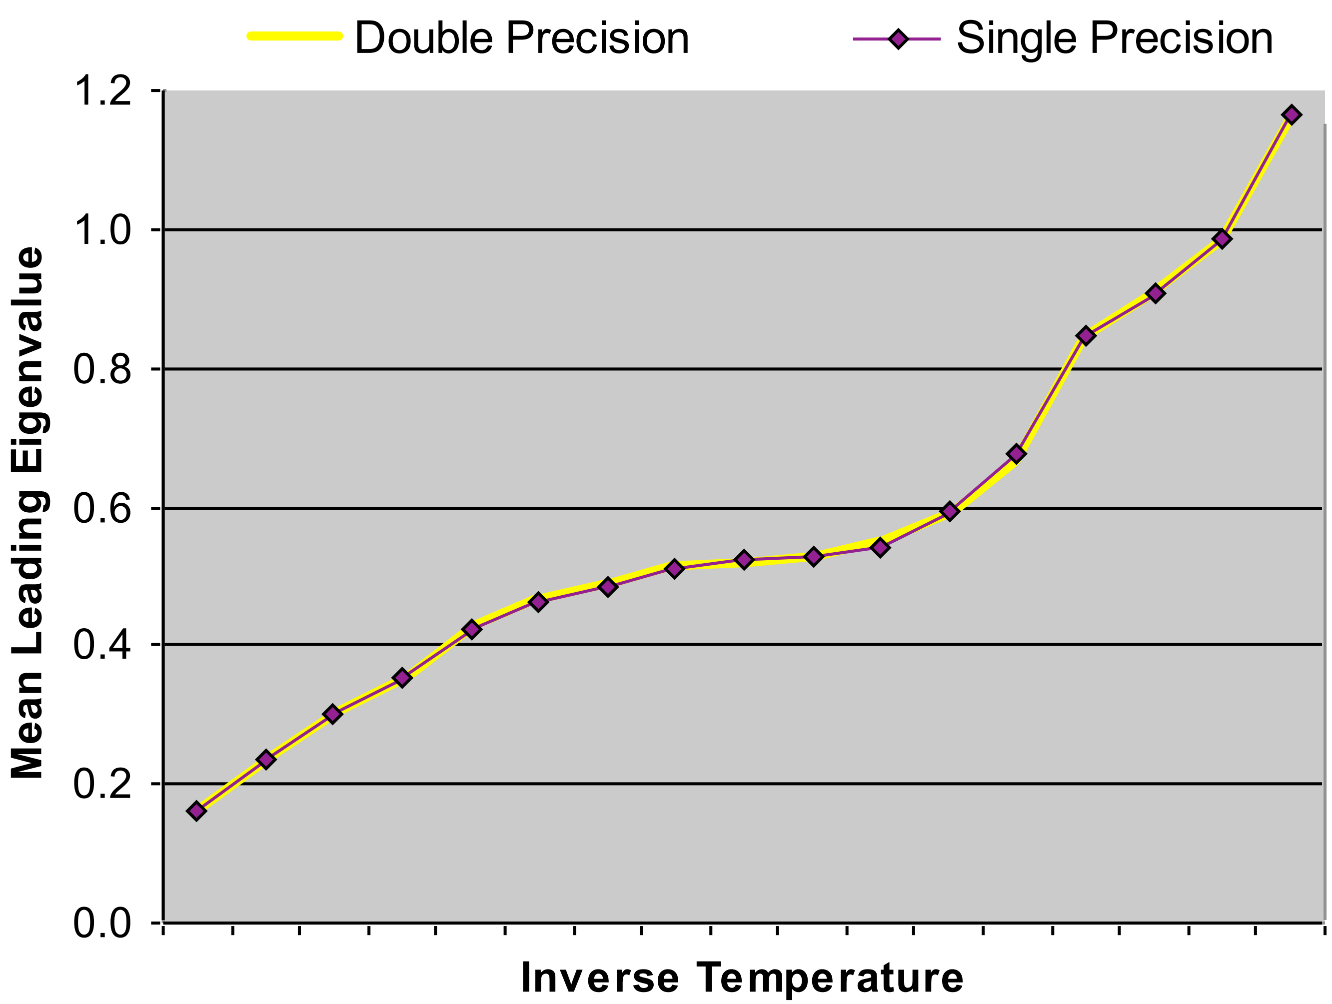

DCA++ with mixed precision

Run HF-QMC in single precision

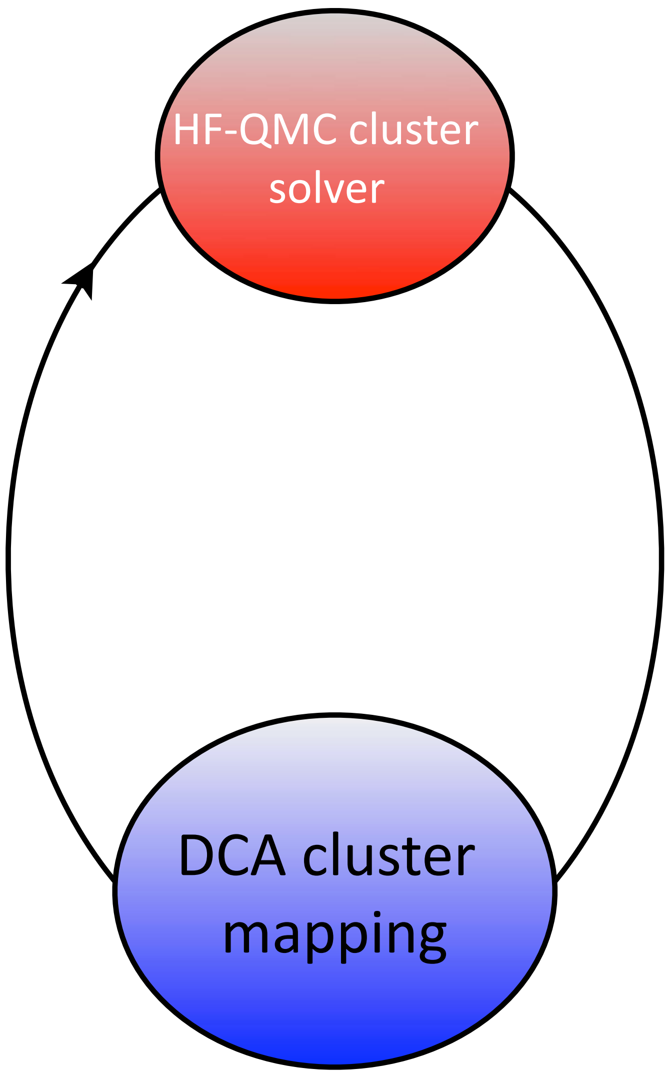

Keep the rest of the code, in particular cluster mapping in double precision

Multiple runs to compute T_c:

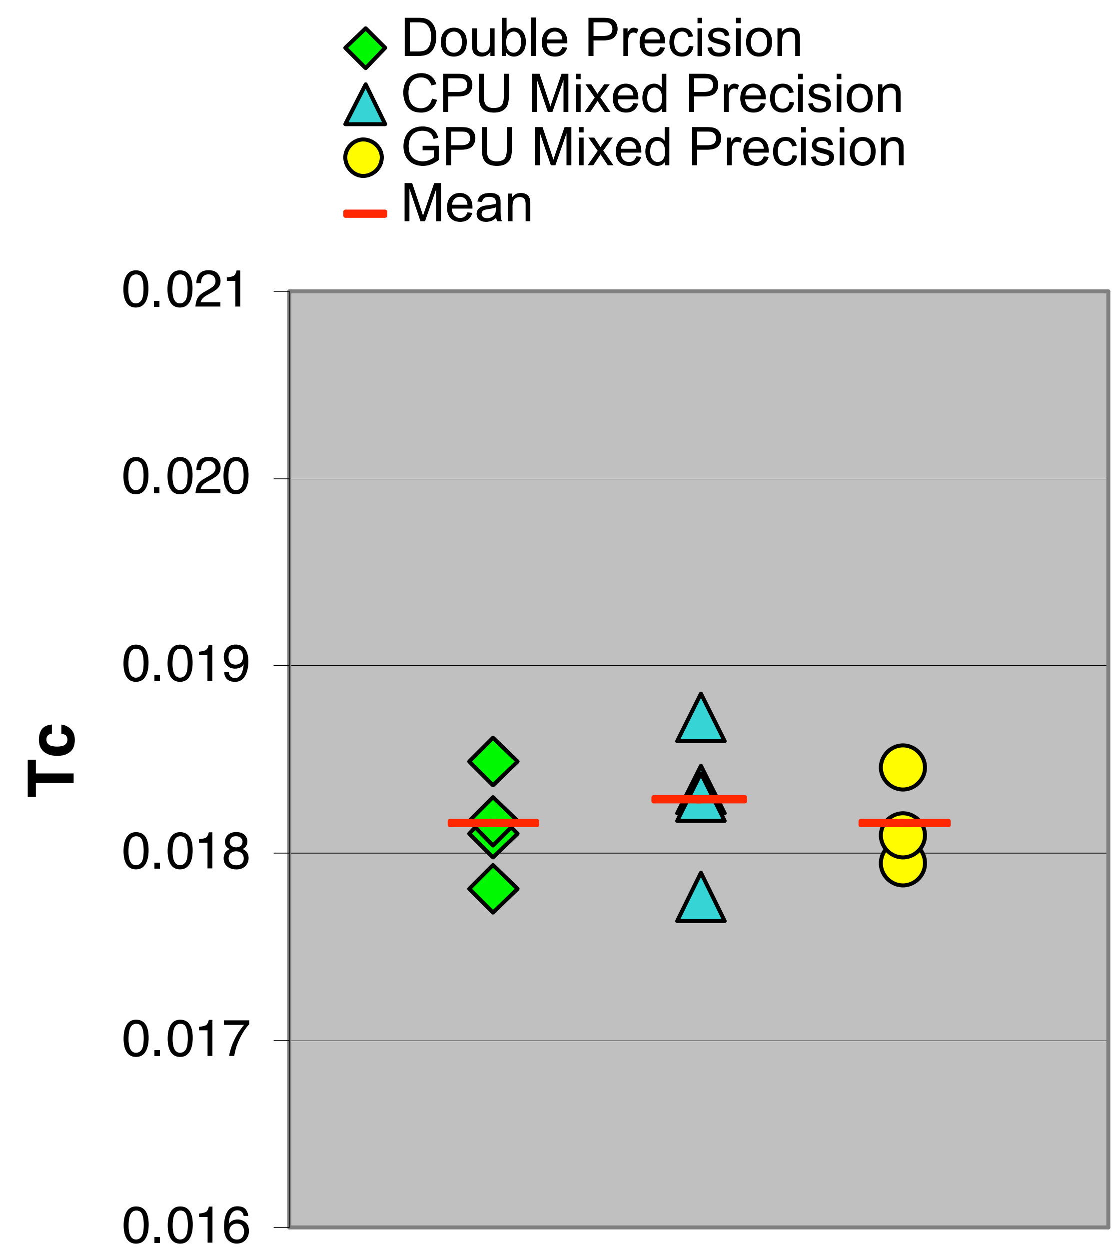

Performance improvement with delayed and mixed precision updates

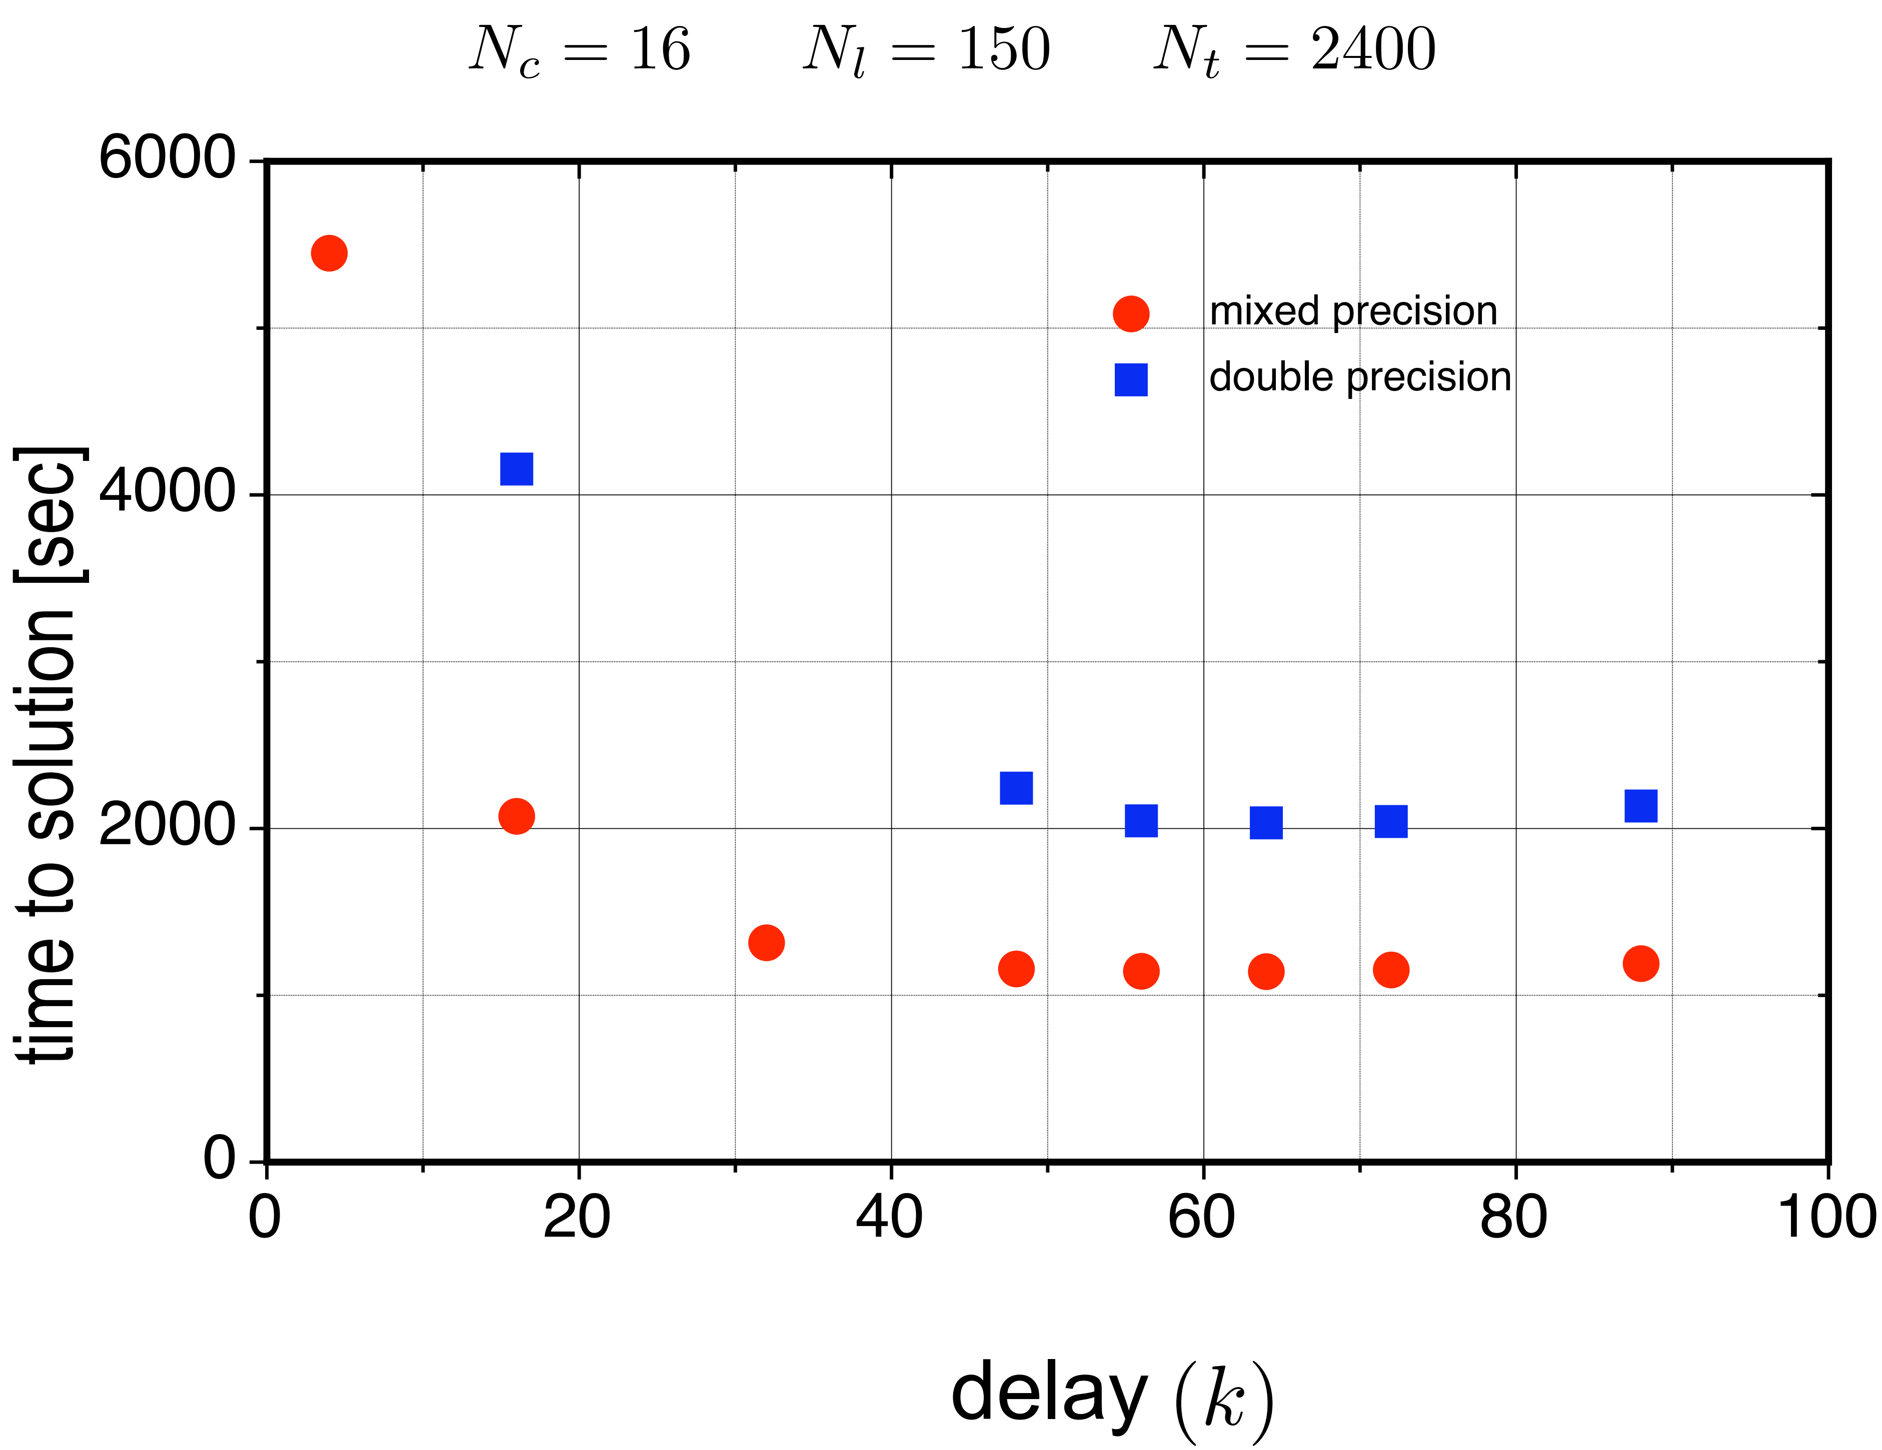

Disorder and inhomogeneities

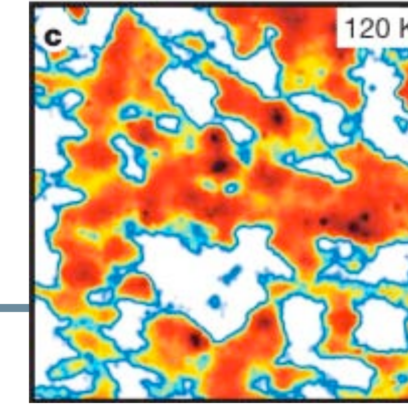

Hubbard Model with random disorder (eg. in U)

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

$$U_i^{(\nu)} \in \{U, 0\}; N_c = 16 \rightarrow N_d = 2^{16}$$

... need to disorder-average cluster Green function

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

Algorithm 1 DCA/QMC Algorithm with QMC cluster solver (lines 5-10), disorder averaging (lines 4, 11-12), and DCA cluster mapping (line 3, 13)

1: Set initial self-energy
2: **repeat**
3: Compute the coarse-grained Green Function
4: **for** Every disorder configuration (in parallel) **do**
5: Perform warm-up steps
6: **for** Every Markov chain (in parallel) **do**
7: Update auxiliary fields
8: Measure Green Function and observables
9: **end for**
10: Accumulate measurements over Markov chains
11: **end for**
12: Accumulate measurements over disorder configurations.
13: Re-compute the self-energy
14: **until** self consistency is reached

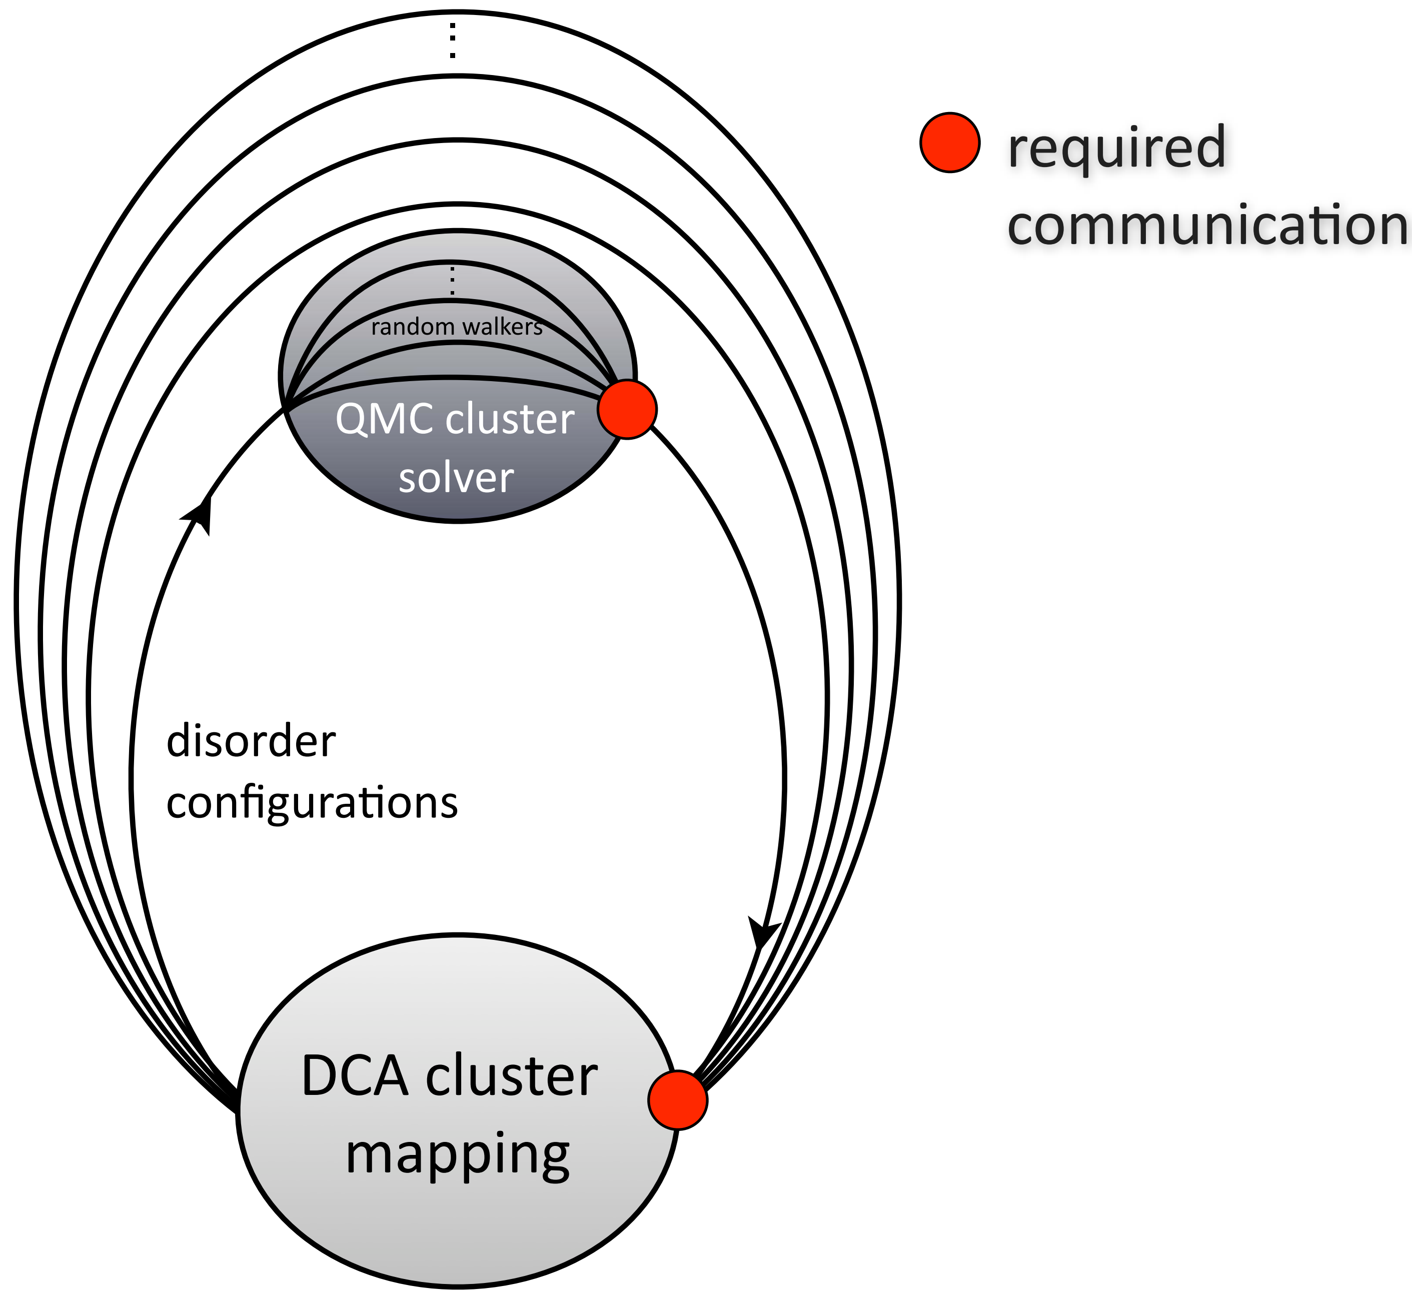

DCA++ code from a concurrency point of view

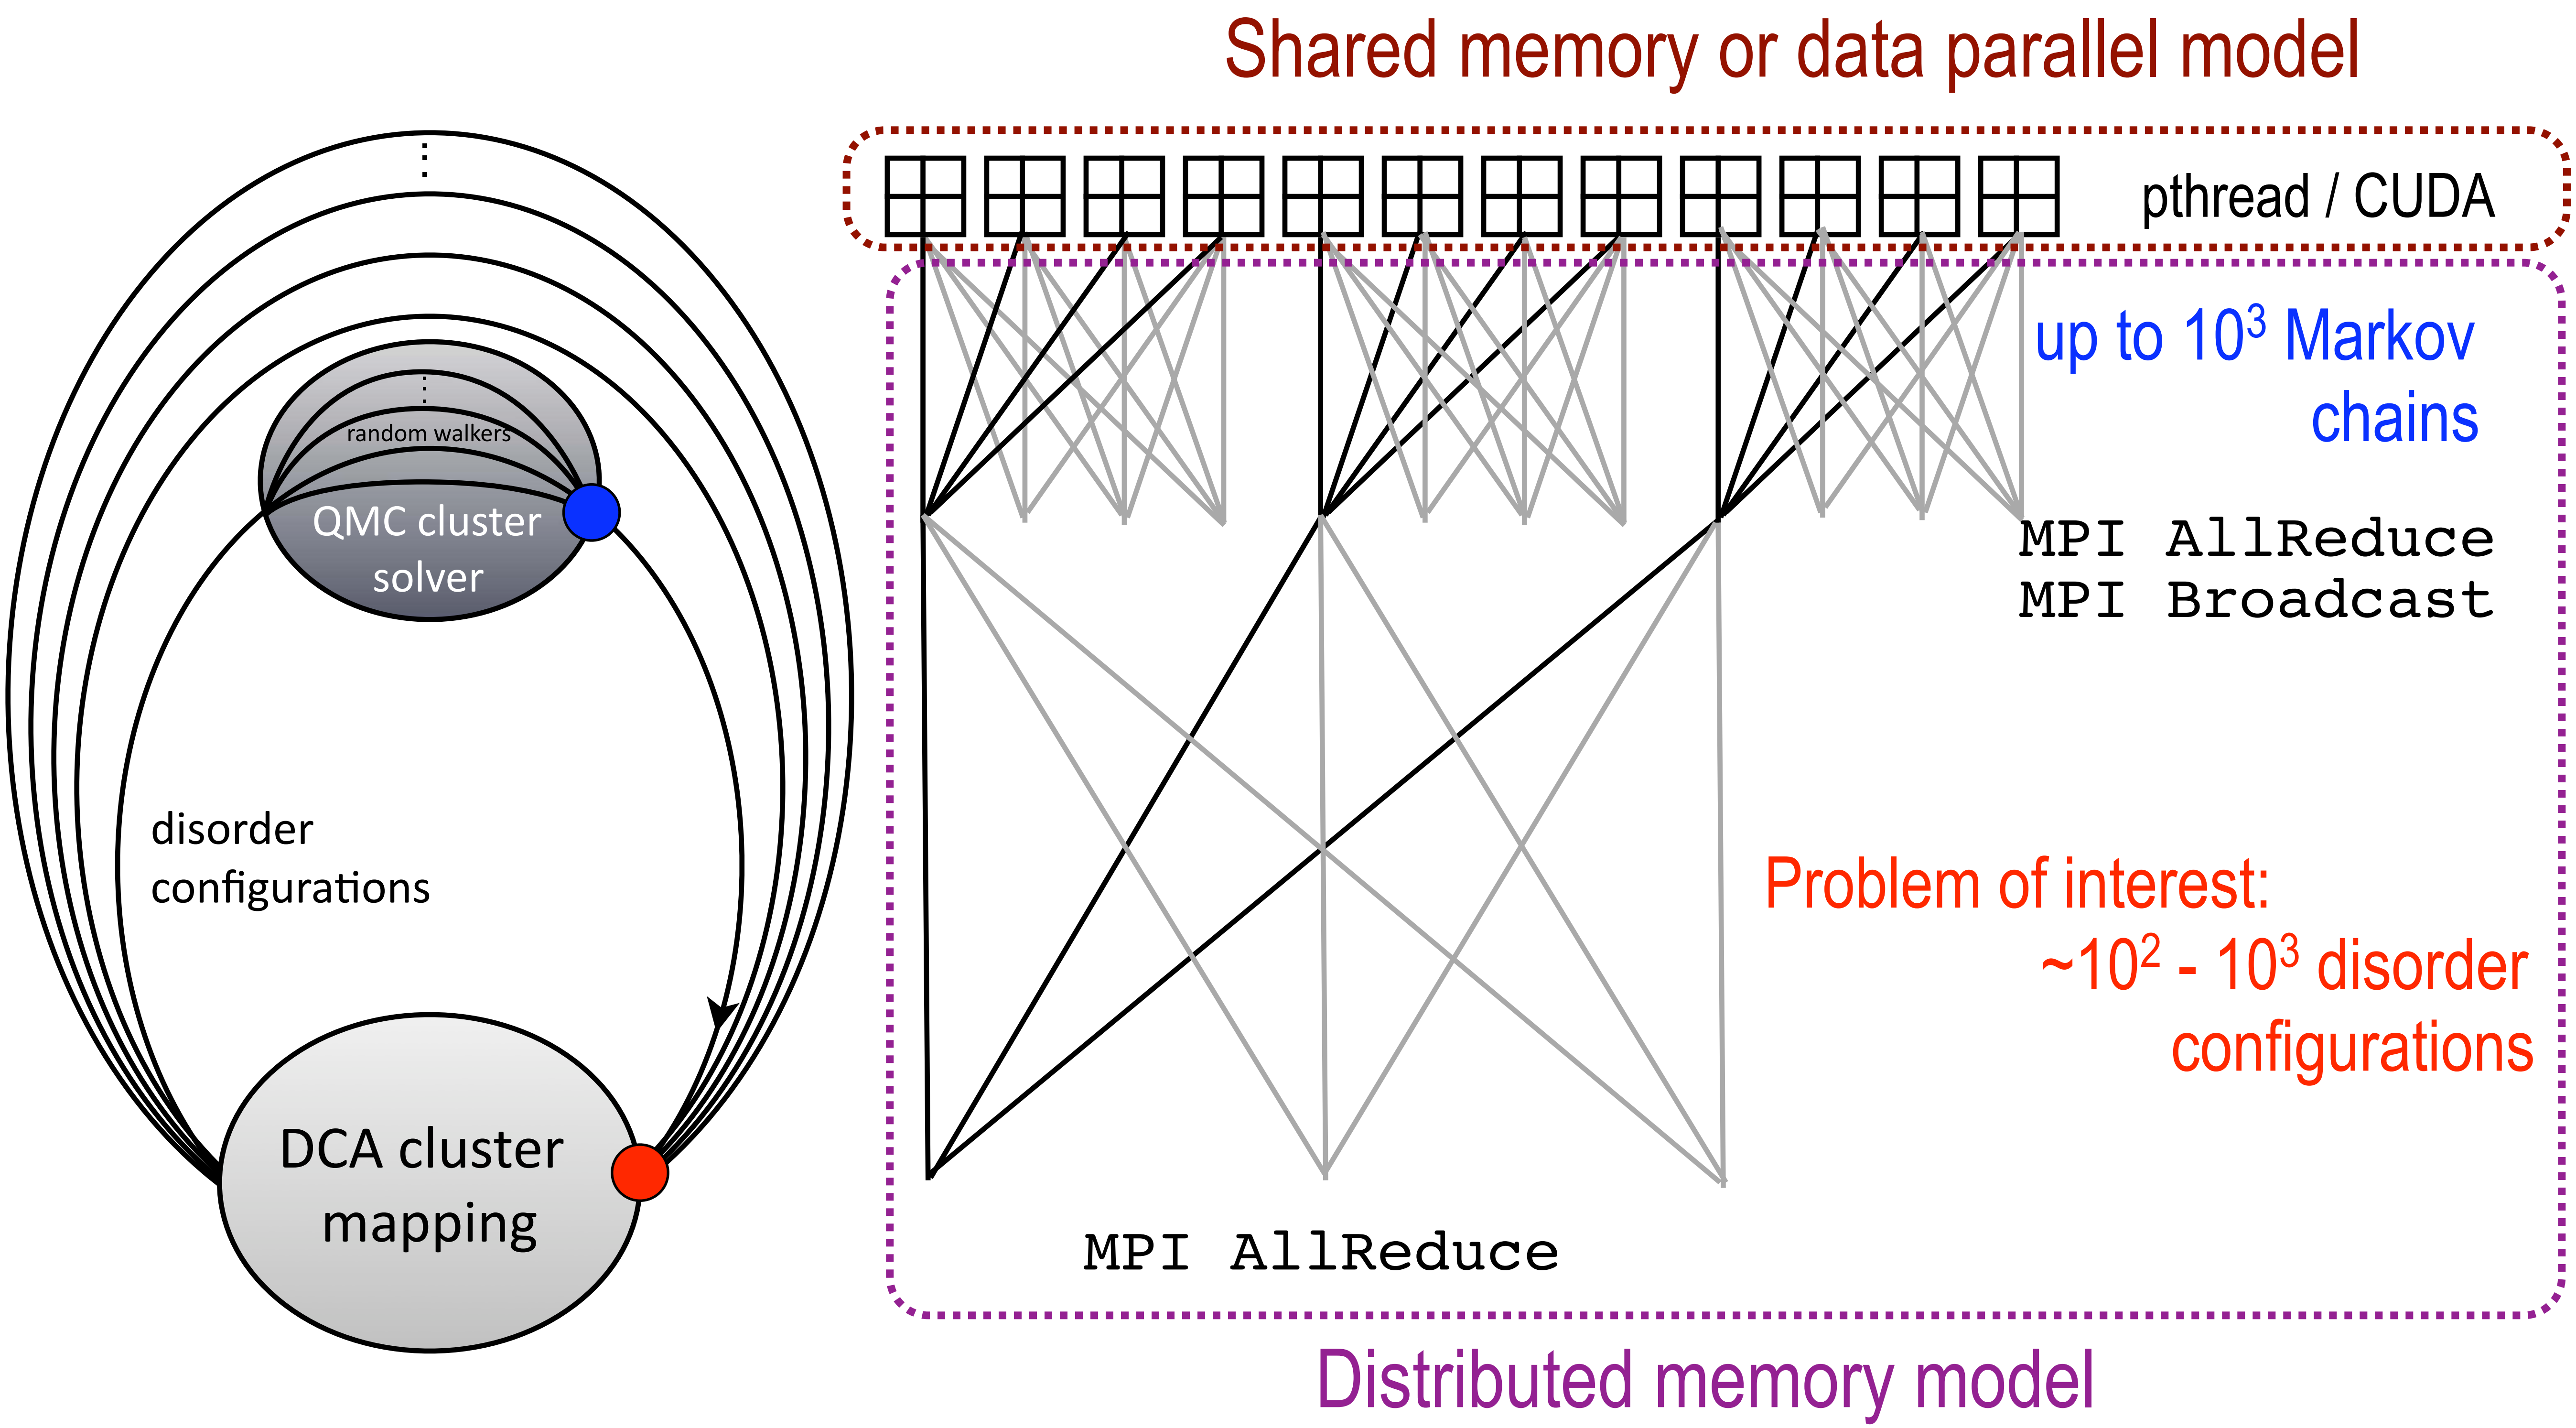

DCA++: strong scaling on HF-QMC

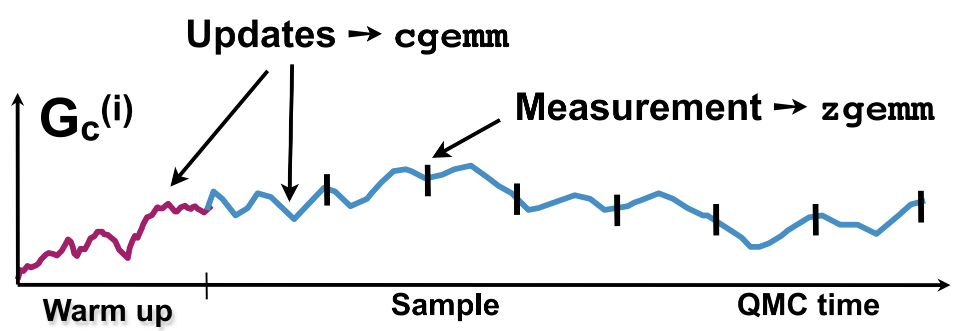

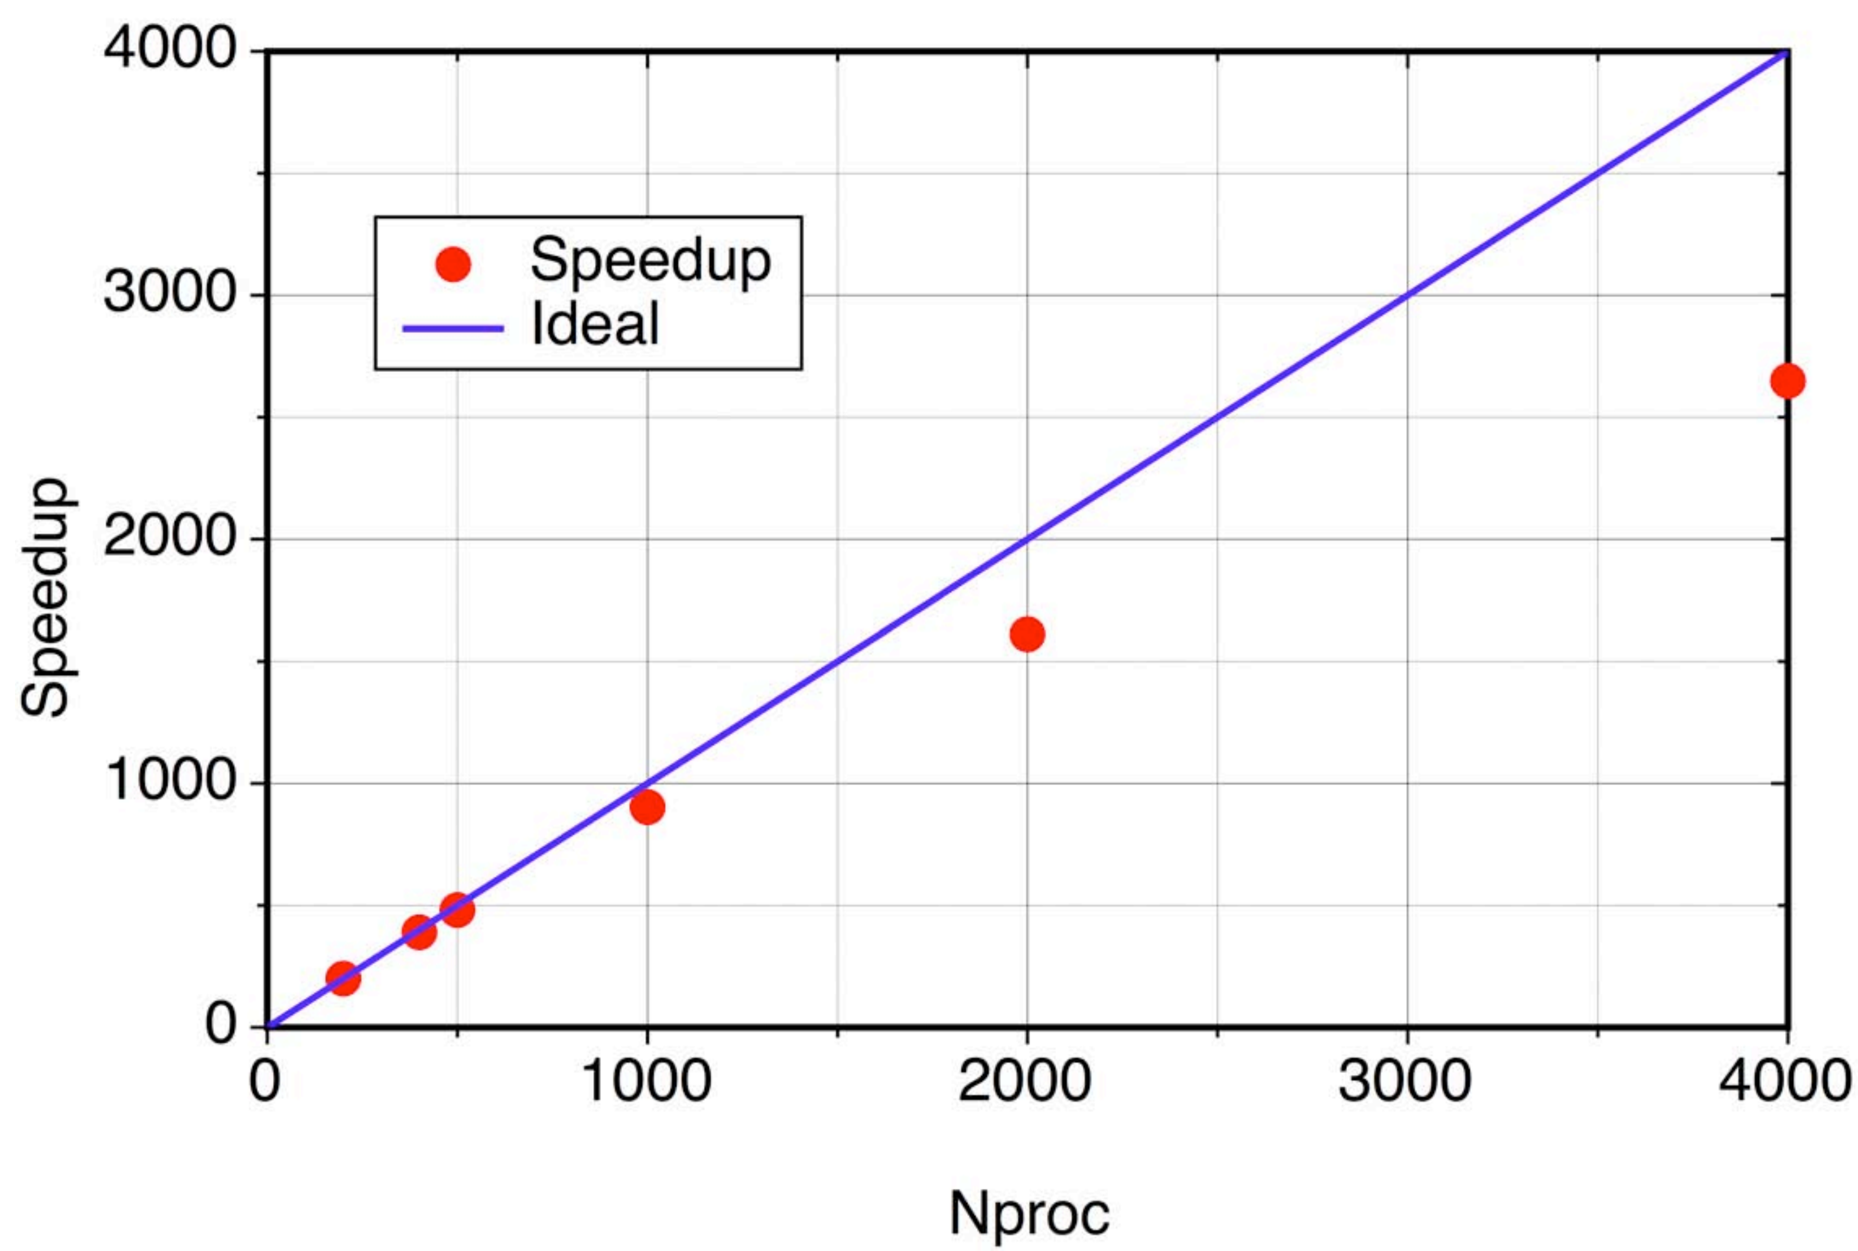

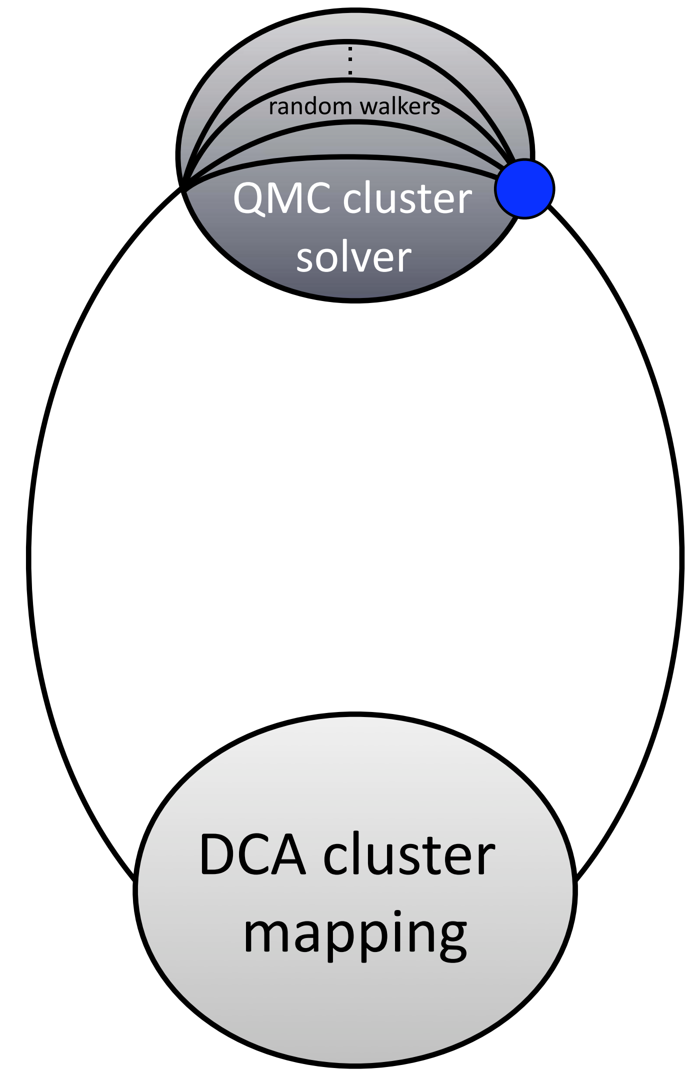

Weak scaling on Cray XT4

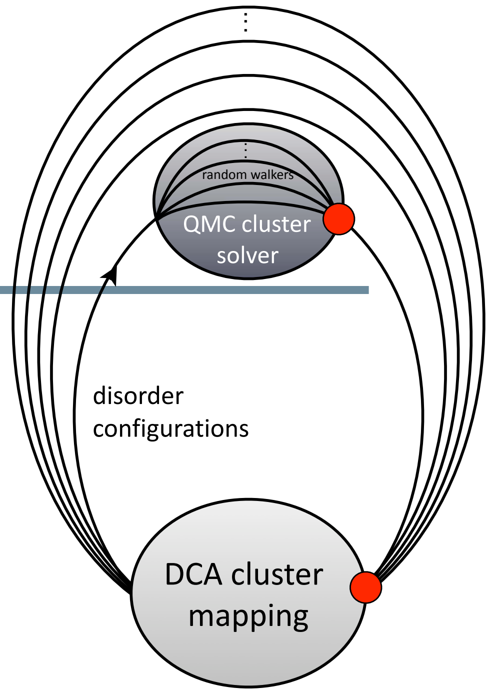

- HF-QMC: 122 Markov chains on 122 cores
- Weak scaling over disorder configurations

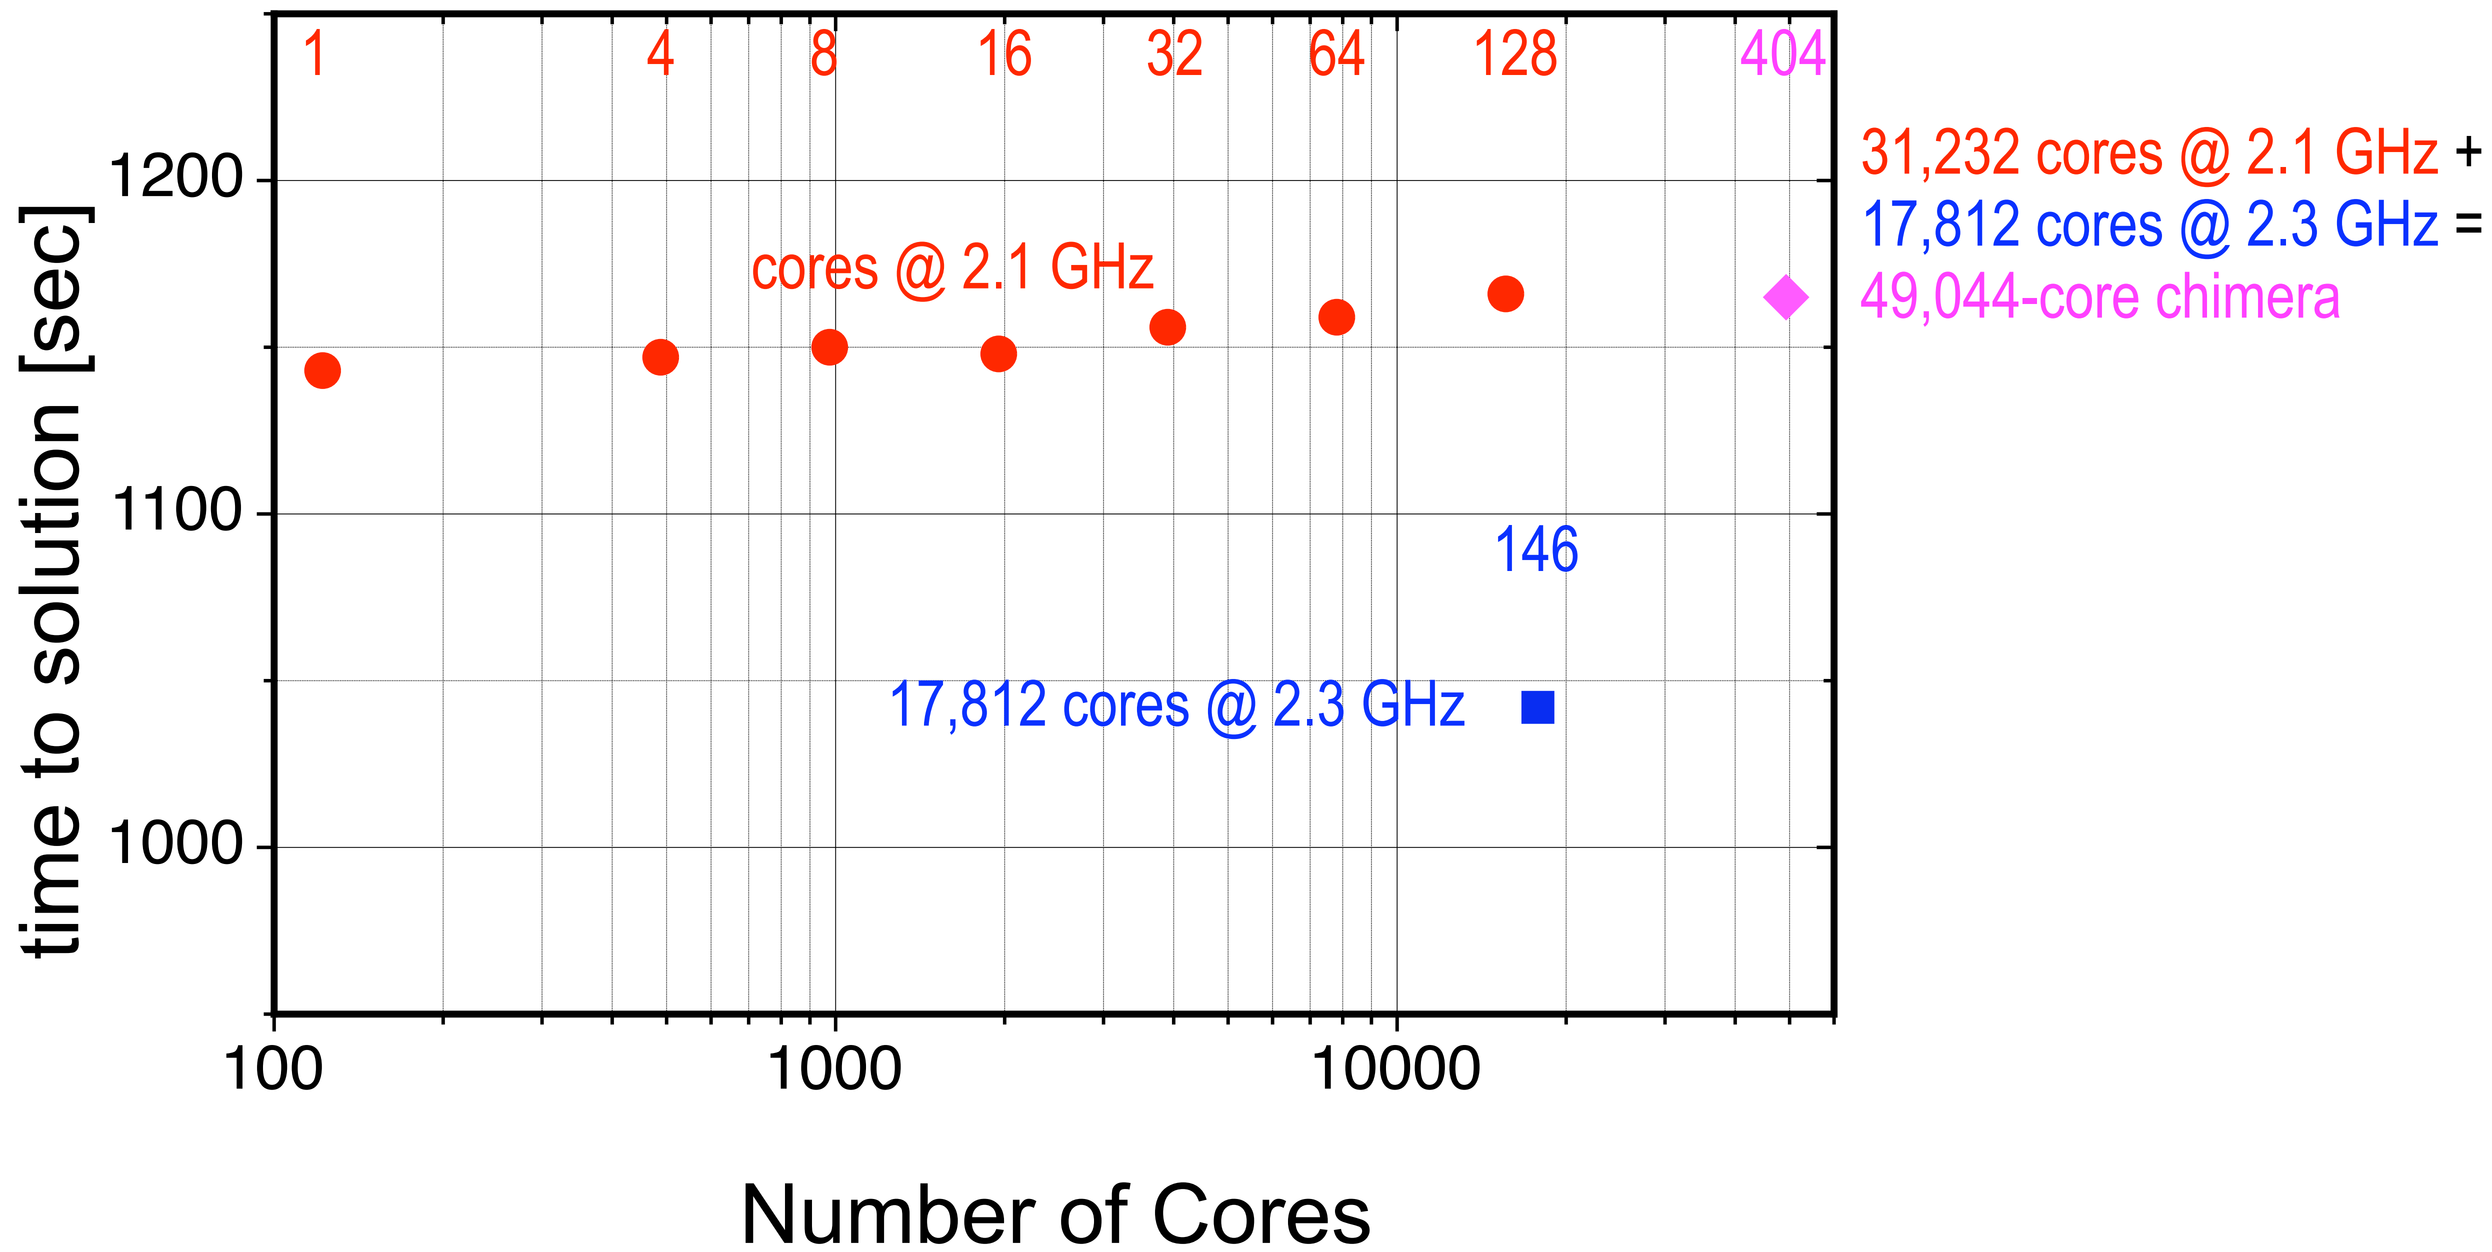

Cray XT5 portion of Jaguar @ NCCS

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

Weak scaling on Cray XT4/XT5 (with buggy use of MPI AllReduce

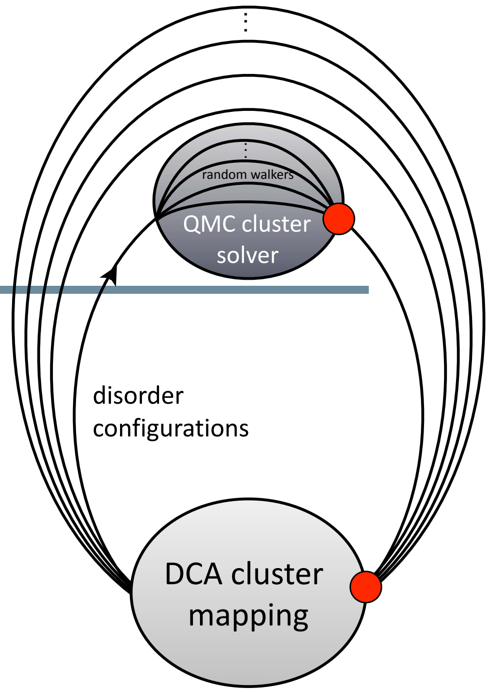

- HF-QMC: 122 Markov chains on 122 cores
- Weak scaling over disorder configurations

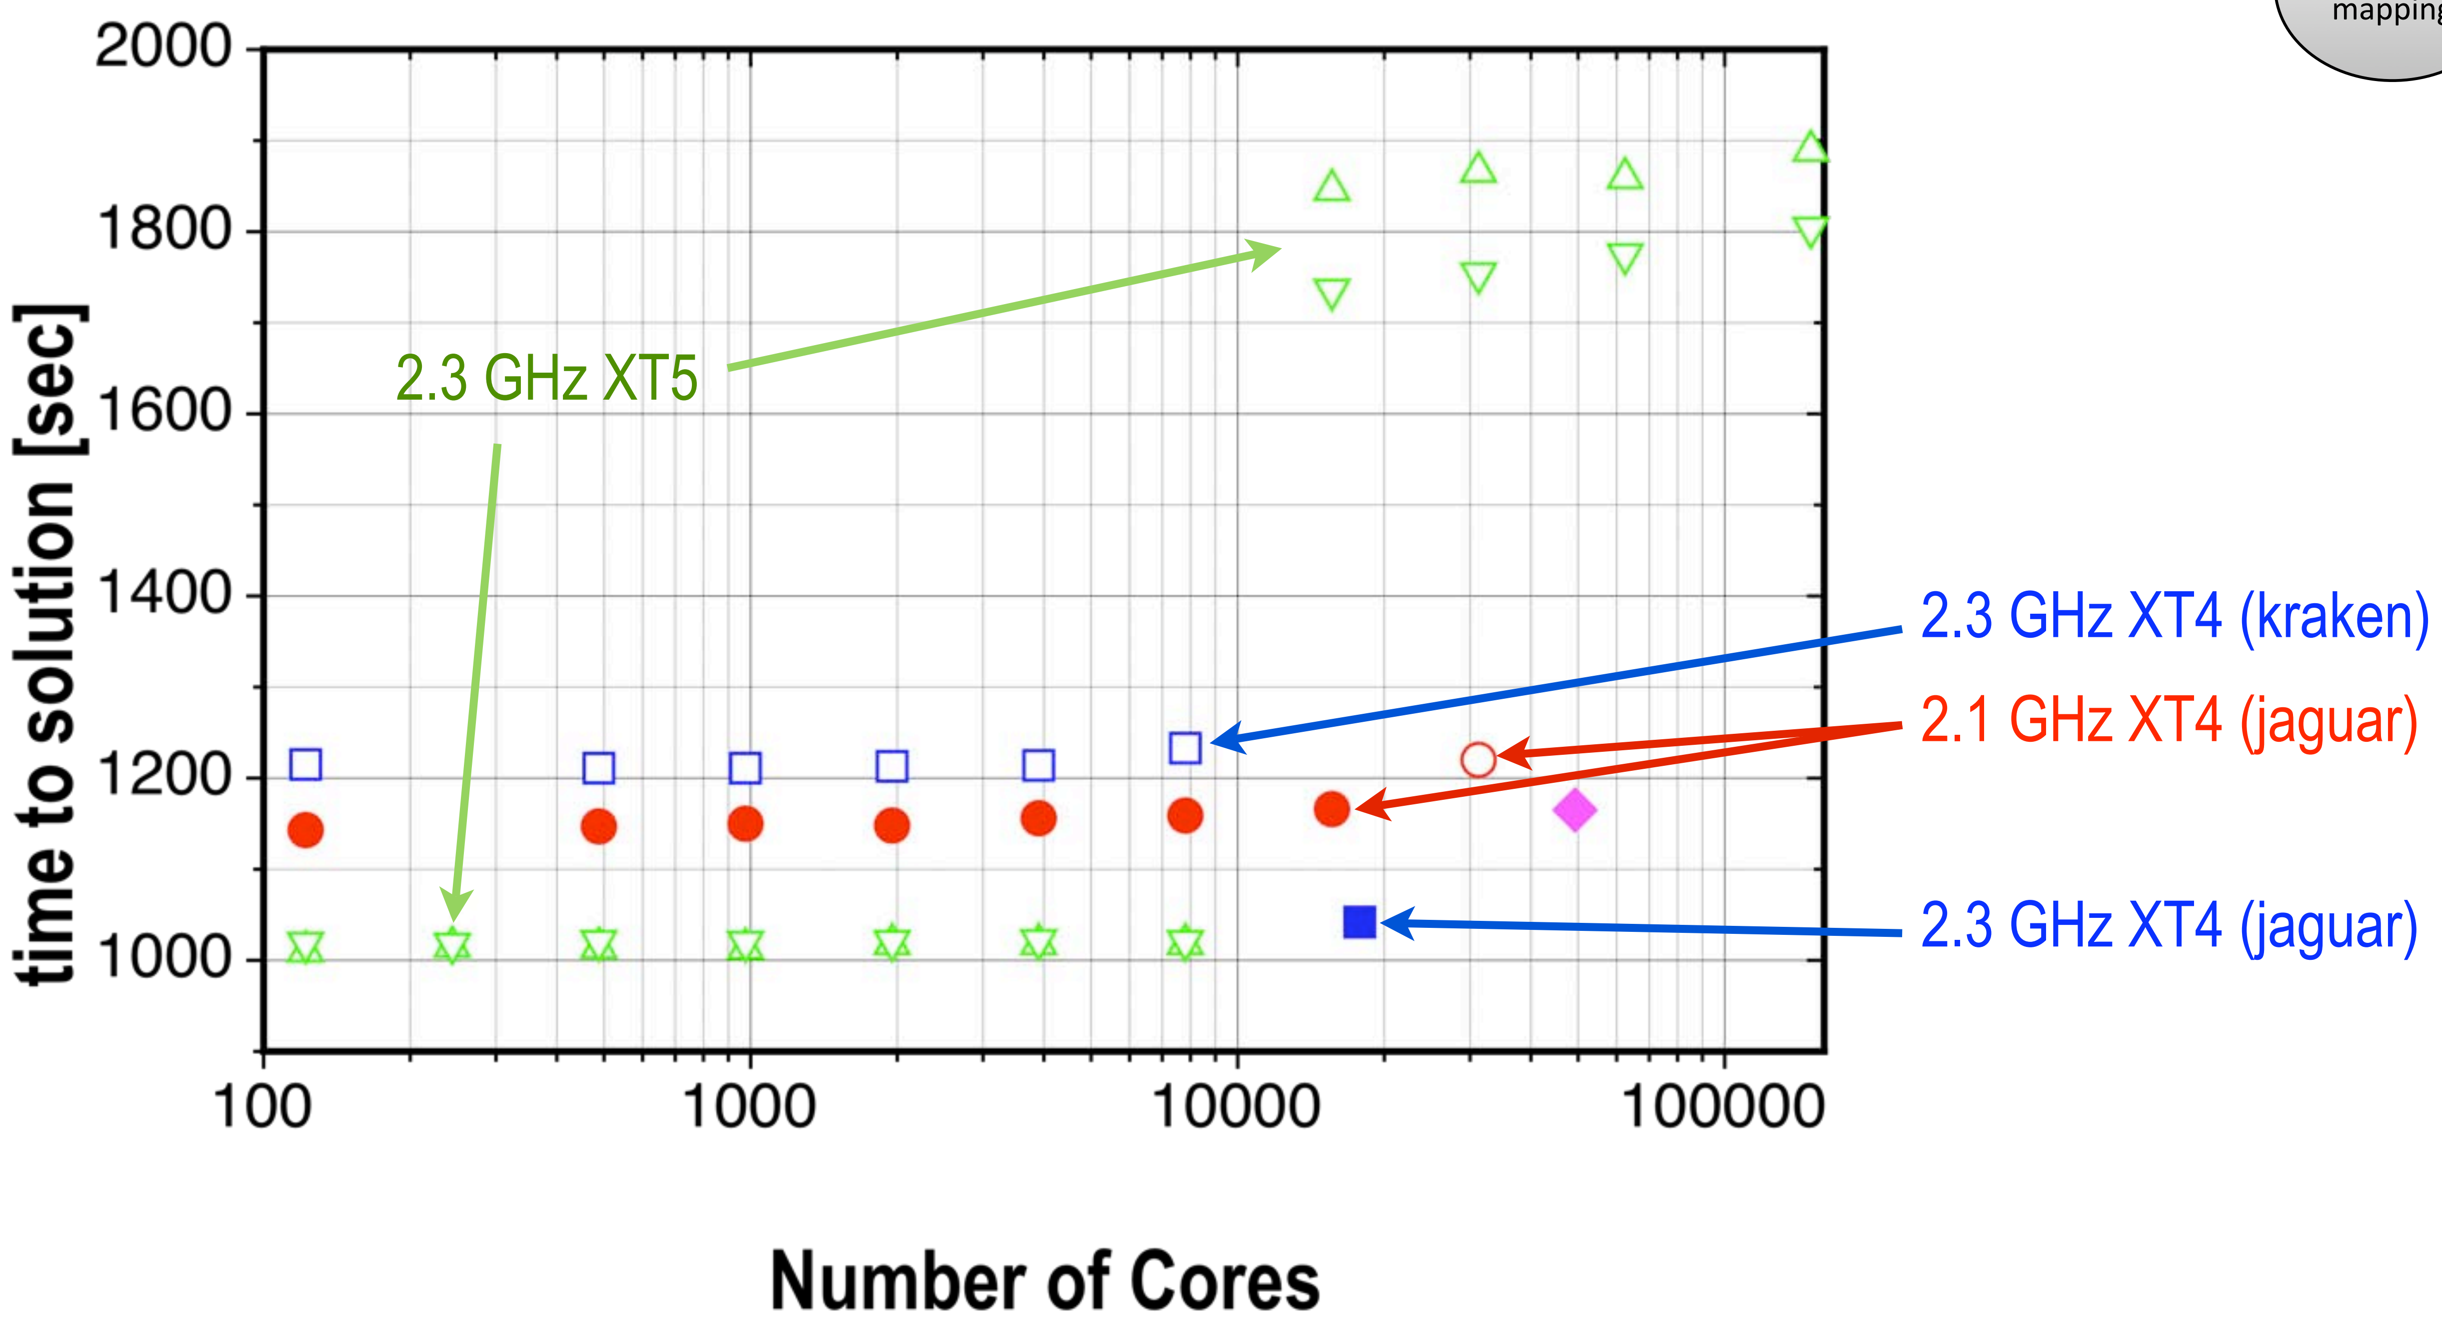

DCA++ code from a concurrency point of view

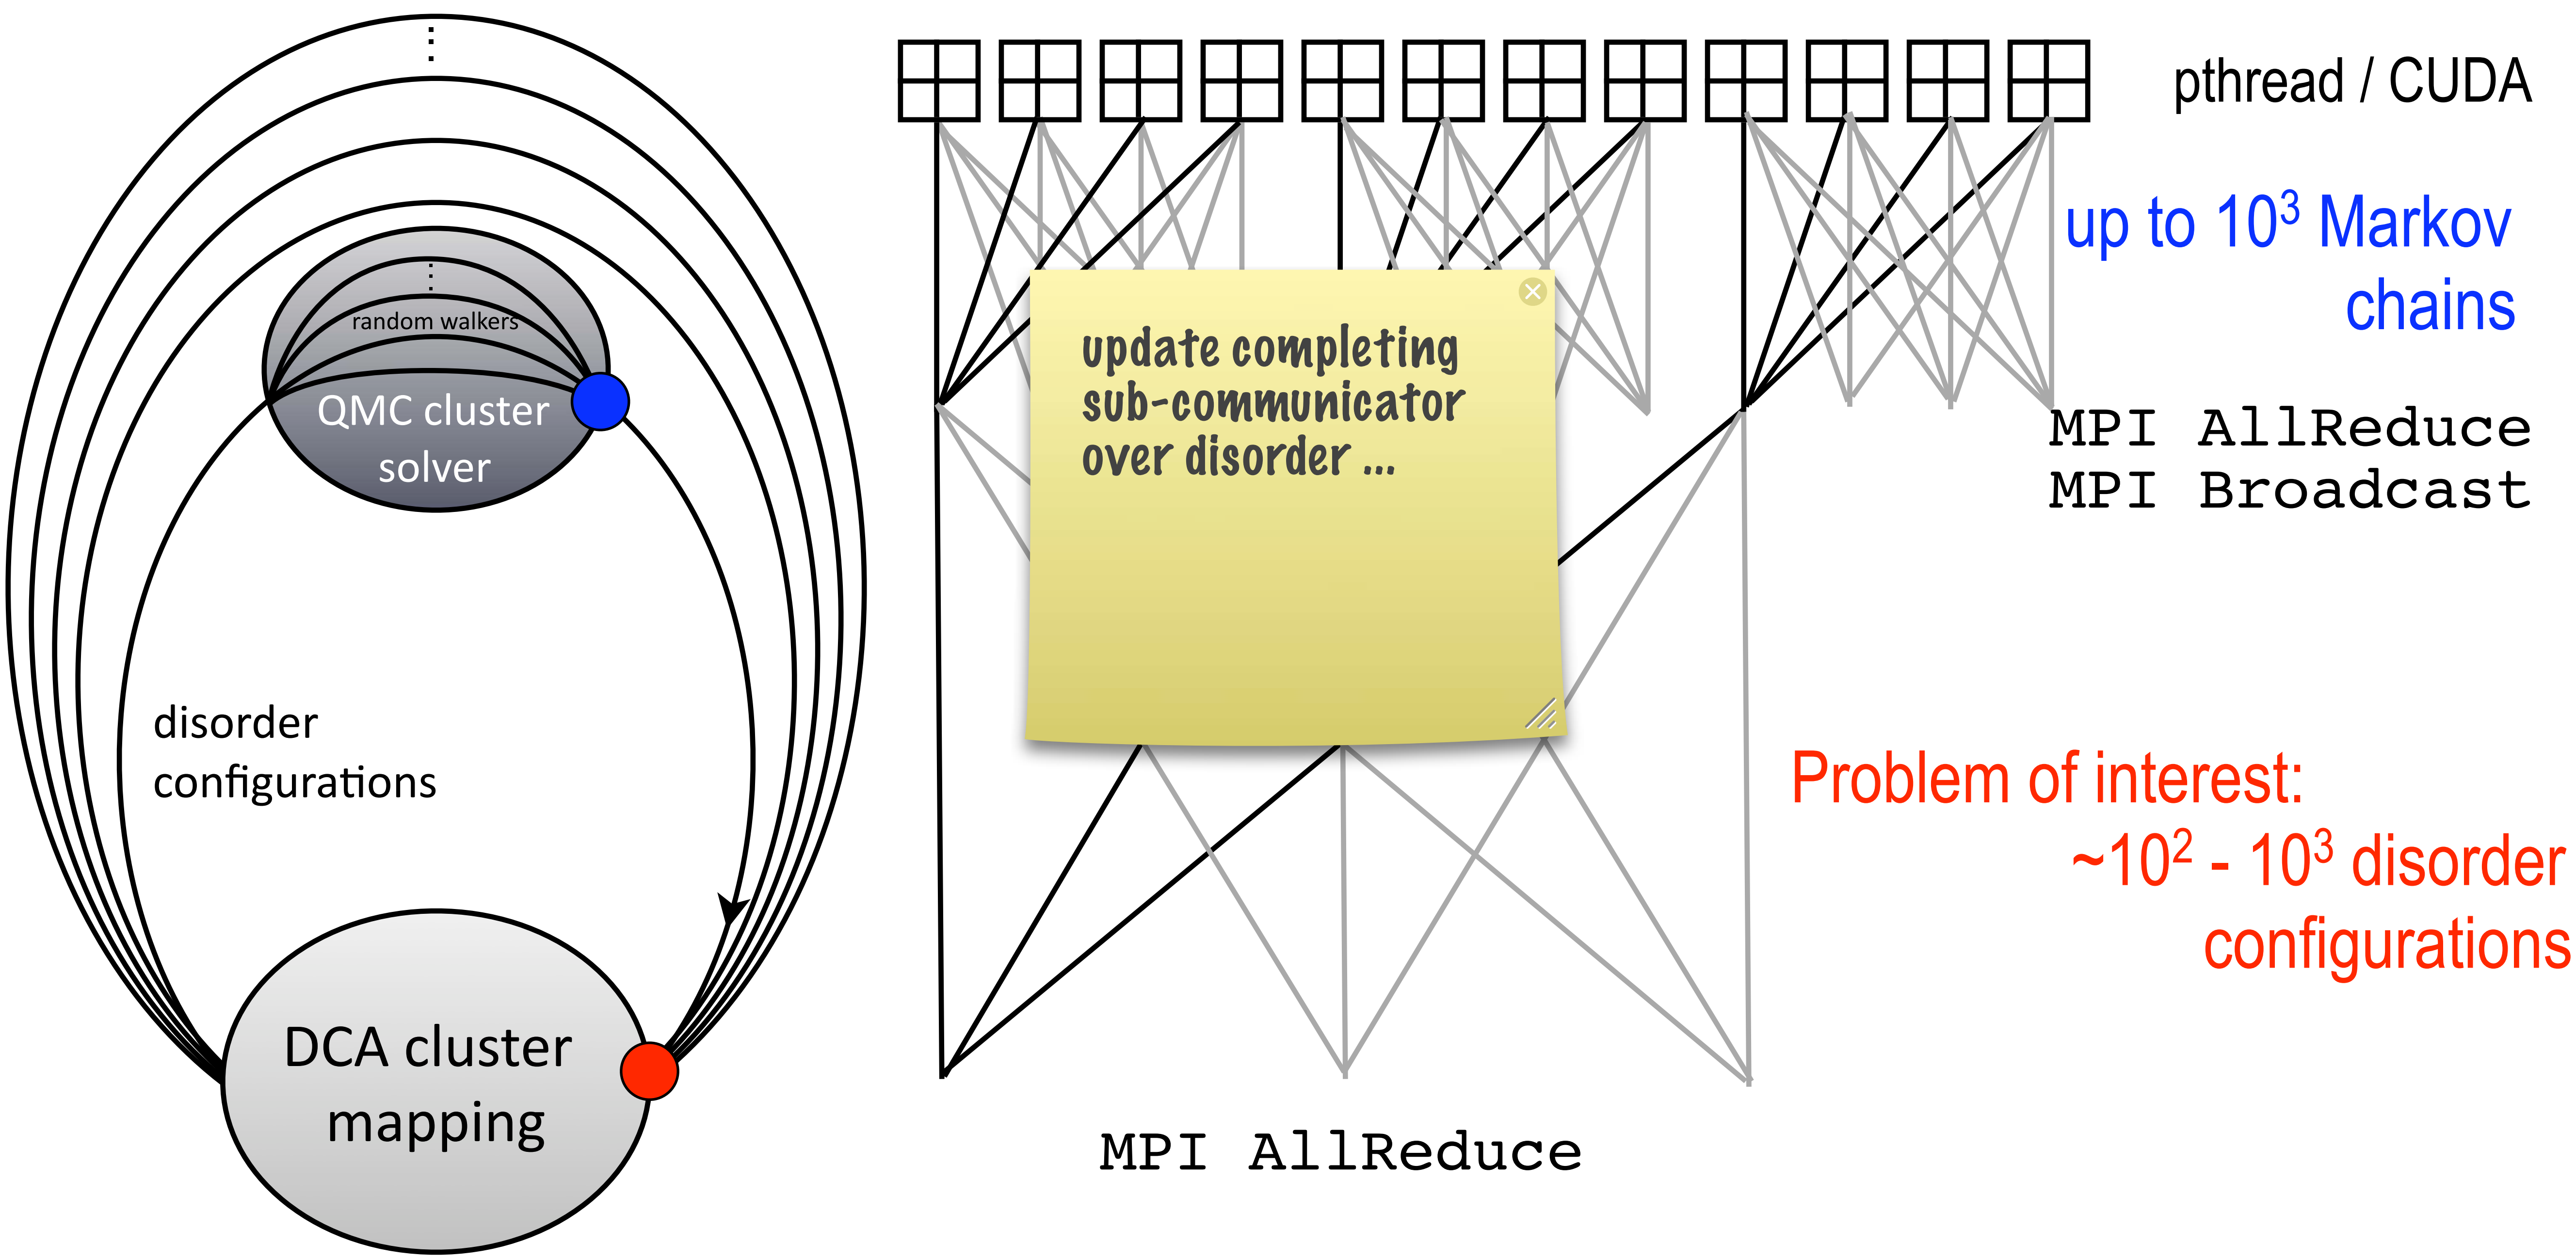

Sustained performance of DCA++ on Cray XT5

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

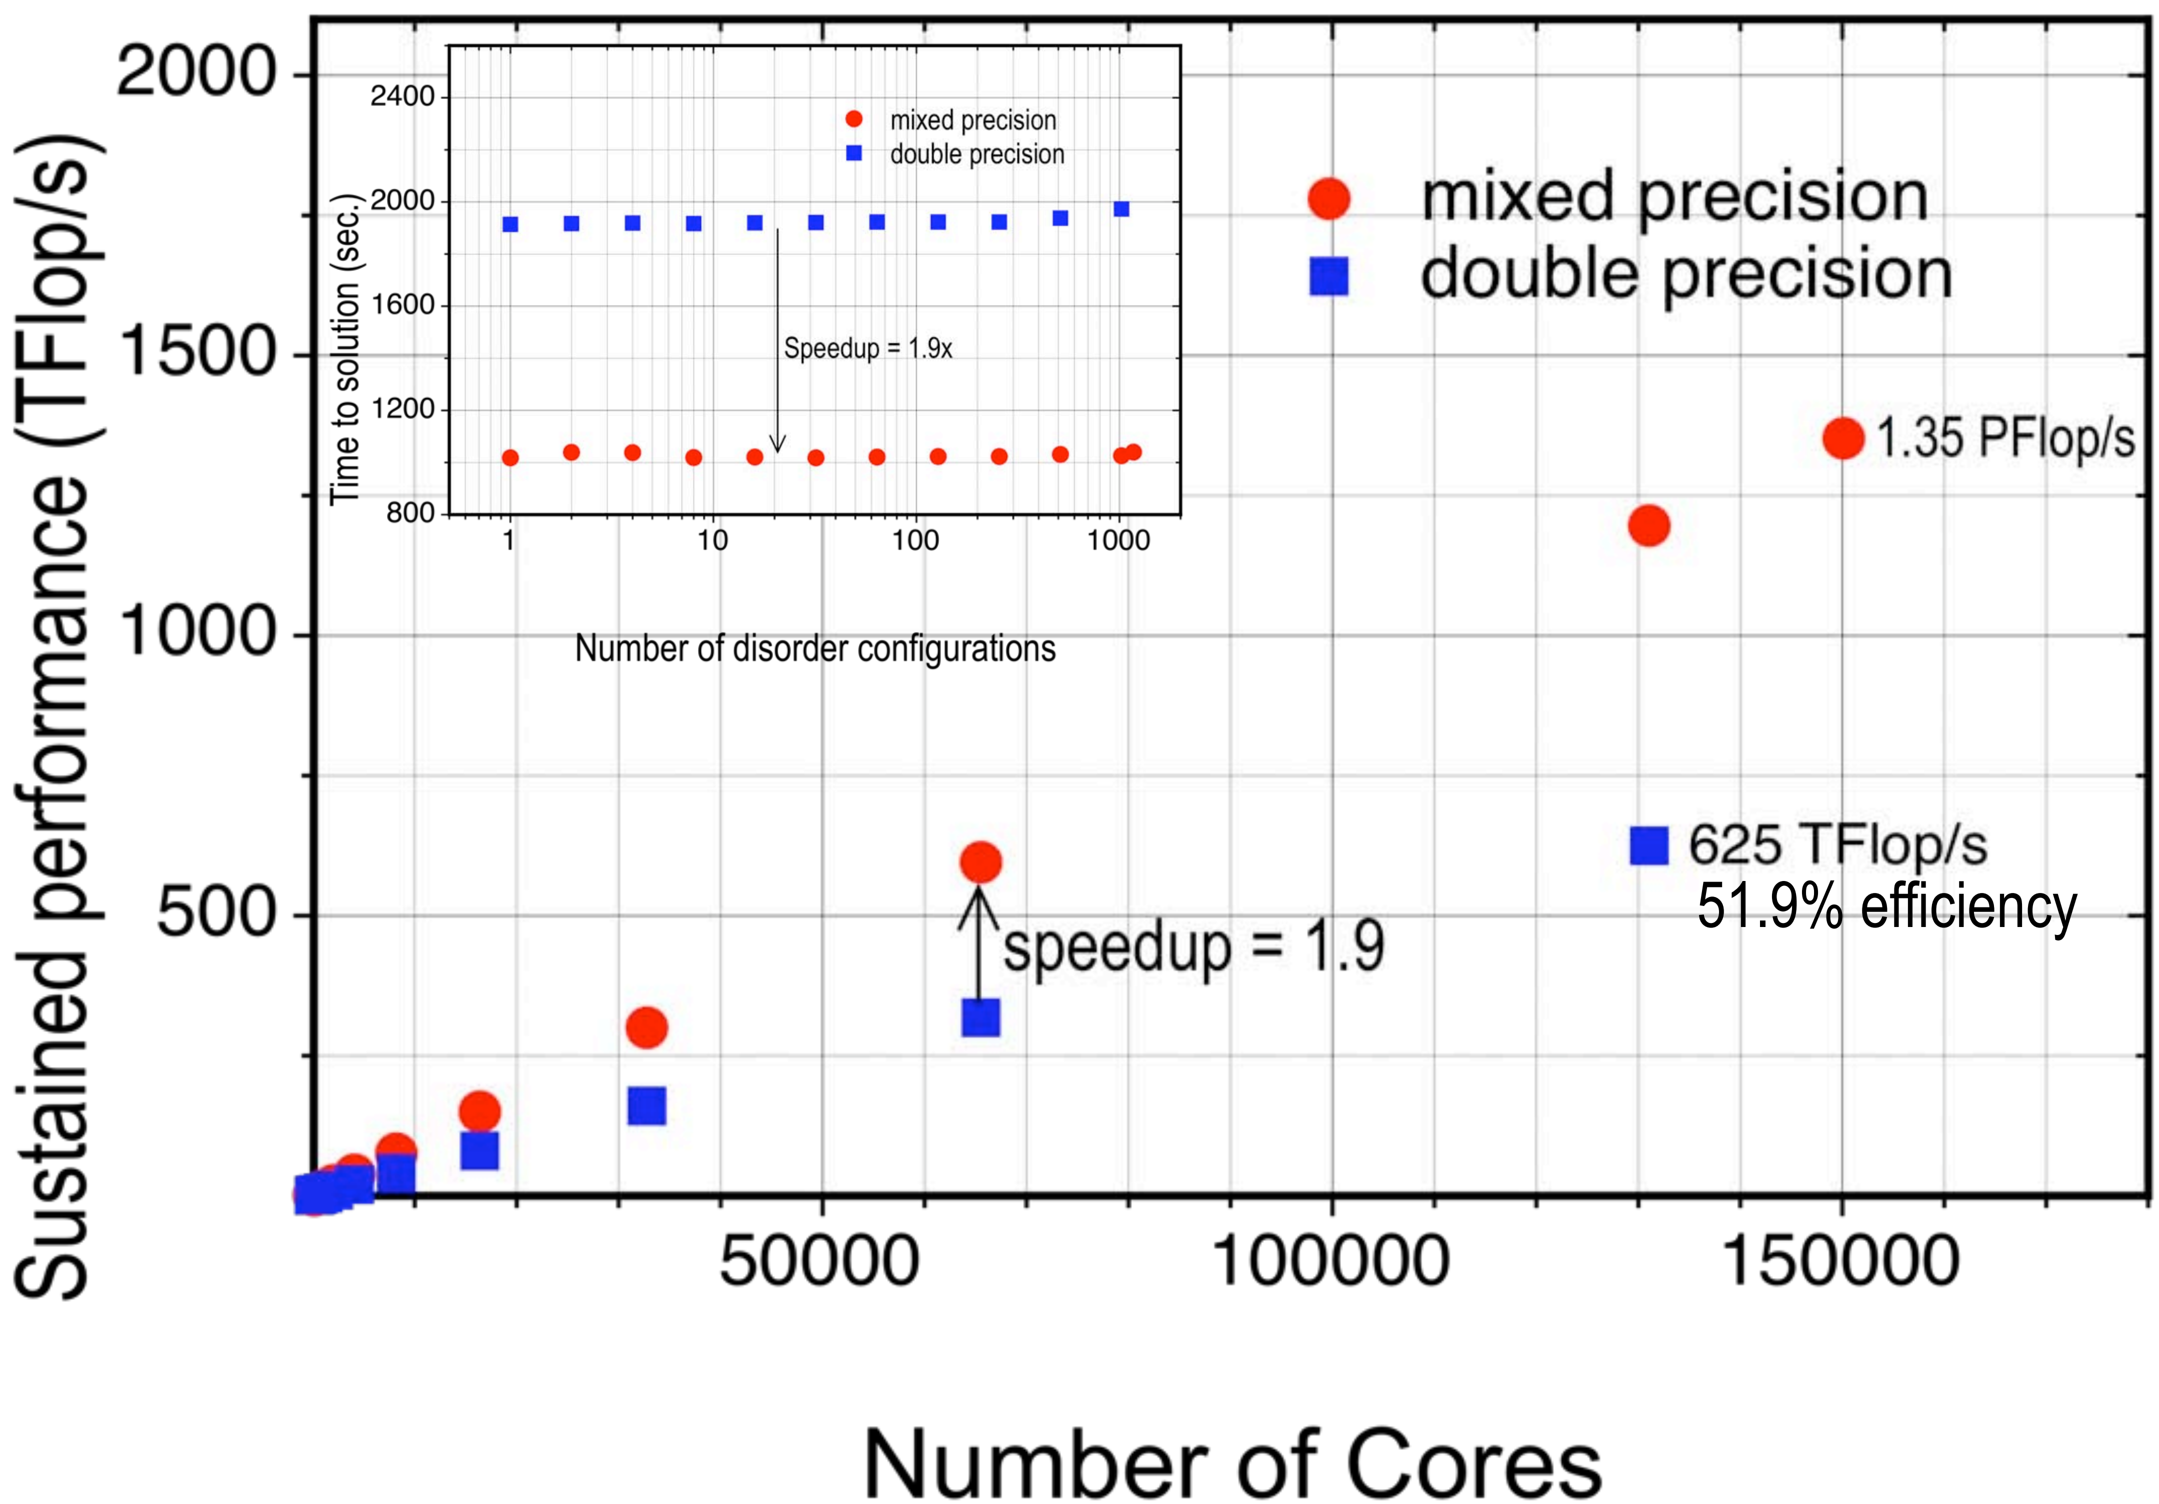

Enhancement of simulation capability since 2003

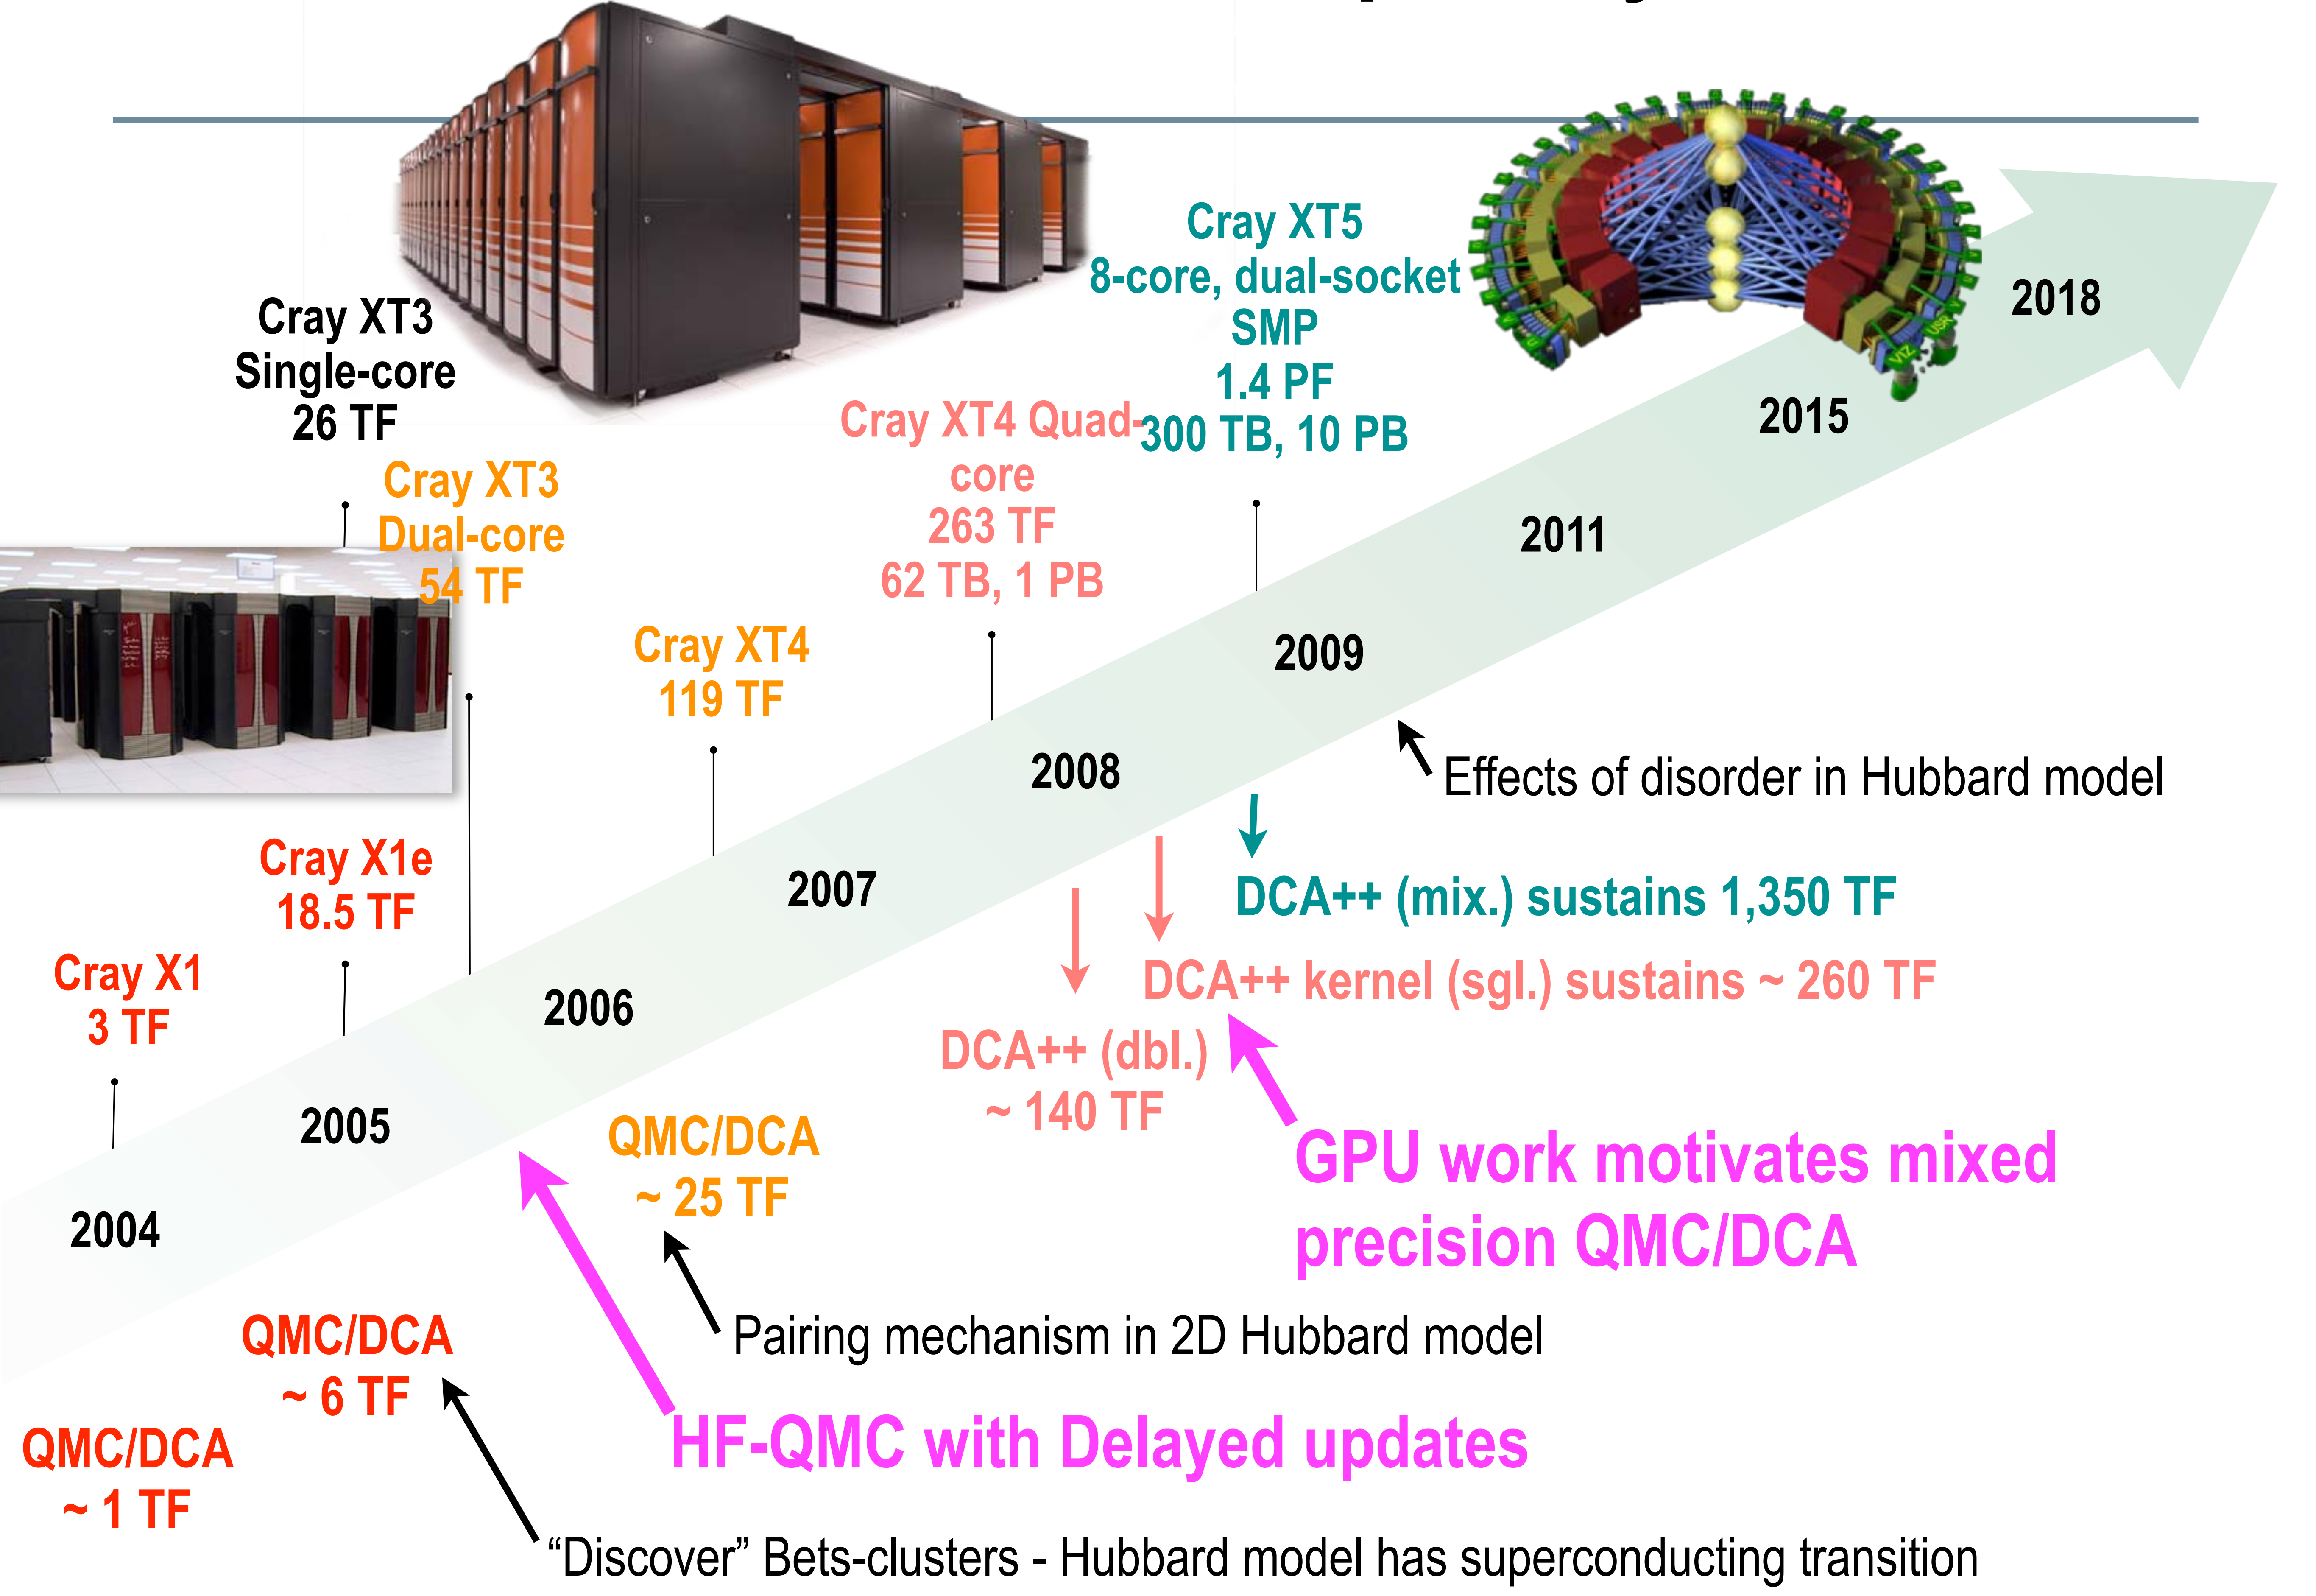

From sustained gigaflop/s to teraflop/s to petaflop/s and beyond

Evolution of the fastest sustained performance in real simulations

~1 Exaflop/s
~10^7 processing units

?

1.3 Gigaflop/s
Cray YMP
8 processors

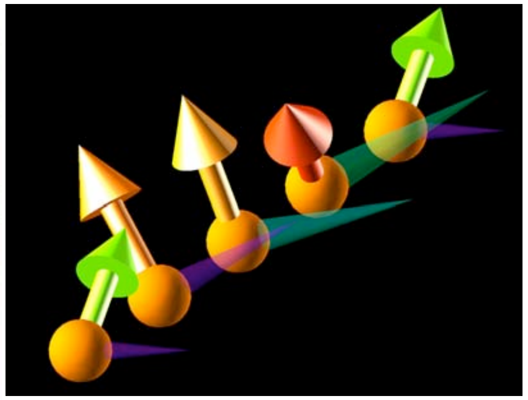

1.02 Teraflop/s
Cray T_{3E}
1.5 10^3 processors

1.35 Petaflop/s
Cray XT5
1.5 10^5 processor cores

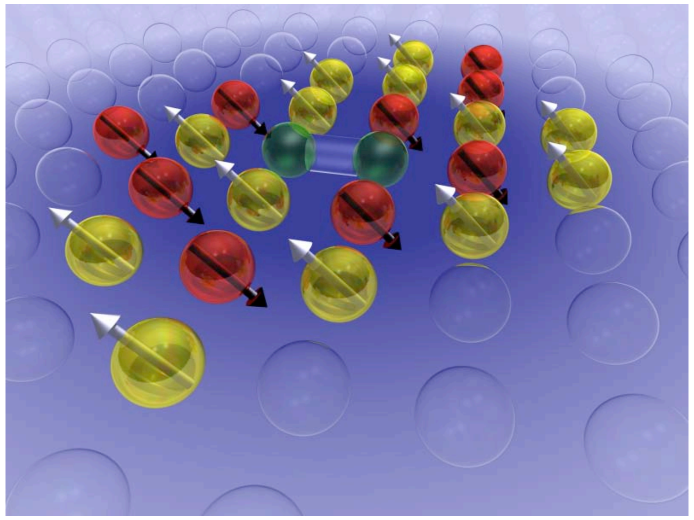

1989 | 1998 | 2008 | 2018

One of seven Gigaflop Award winners in 1989

First sustained TFlop/s Gordon Bell Prize 1998

First sustained PFlop/s Gordon Bell Prize 2008

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

CSCS
Swiss National Supercomputing Centre

DCA++ Story: team*, collaborators, resources, and funding

Thomas Maier
Paul Kent
T Schulthess
Gonzalo Alvarez
Mike Summers
Ed D'Azevedo
Jeremy Meredith
Markus Eisenbach
Don Maxwell
Jeff Larkin
John Levesque

Physics
Application software
Comp. mathematics
Computer Science
Computer Center
Hardware vendor

D. Scalapino
M. Jarrell
J. Vetter
Trey White
staff at NCCS & Cray
& many others

Computing resources:
NCCS @ ORNL

Funding:
ORNL-LDRD,
DOE-ASCR,
DOE-BES

*names order according to background

Oak Ridge National Laboratory

Questions / Comments?