Making good use of HPC in Materials Science

Thomas C. Schulthess

Computational Materials Science

Computer Science and Mathematics

schulthesstc@ornl.gov

HPC Centers: the reality today

User Community

mostly individual researchers, small groups

(exceptions: climate)

- **Commodity hardware**
- **Users develop/port/run codes with minimal support**
- **Efficiency often very low**
  - **Codes run below 5% of peak**
  - **Researchers spend most of their time developing / porting / optimizing codes**

HPC in the coming years

~10^3 vector proc.

~10^5 super scalar proc.

- **Deliver 100-1000 times more performance with specialized hardware**
  - **not commonly available**
- **Capability Computing (few use entire resource)**
  - **Applications *must* scale and run with high efficiency**

Role of HPC from apps point of view

- **Misreading the reality:**
  - **Machines built to solve LINPACK**
  - **Processor count going from 10, to 100, to 1000, to 10,000, to 100,000, to ???**
  - **Programming model never changed (goes back to the 60s)**
- **What I try to show here:**
  - **HPC is badly needed to solve problems in computational science**
  - **There should not be a barrier between low end and high end computing**
  - **What an adequate programming model may look like**

Outline

- **Introduction**
- **Application example 1: A prototype grand challenge team (High-T_c superconductivity)**
- **Application example 2: Simulations and experiments in nanoscience**
- **Programming model for materials, nano-, and condensed-matter science**

Two Dimensional Hubbard model for High-T_c superconductivity

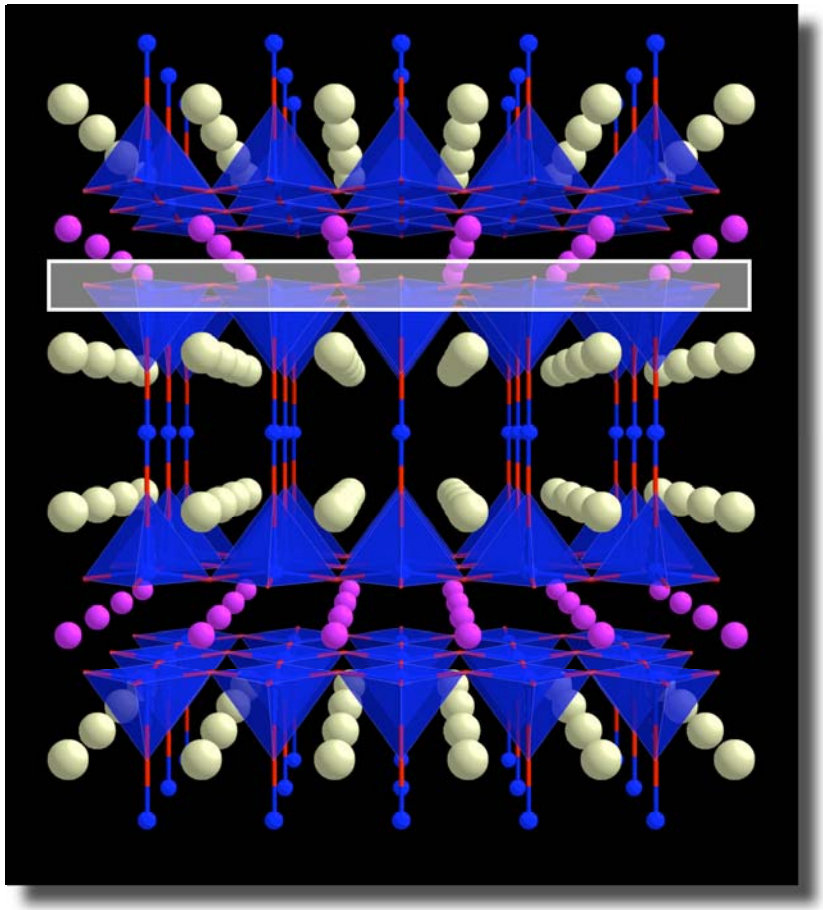

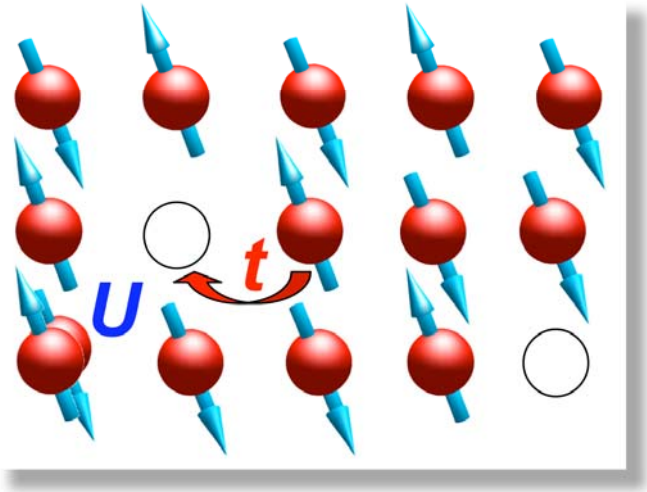

(Zhang and Rice, PRB 1988, P.W. Anderson)

$$\mathcal{H} = -t \sum_{\langle ij \rangle, \sigma} c^{\dagger}_{i\sigma} c_{j\sigma} + U \sum_{i} n_{i\uparrow} n_{i\downarrow}$$

- **Simples model**
- **Not solvable**
- **~1000 papers/year**

4-site cluster QMC/DCA code (IBM Power 4)

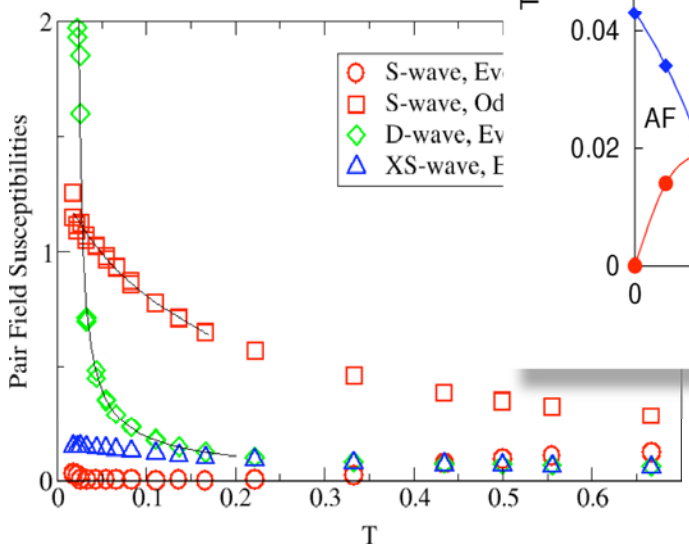

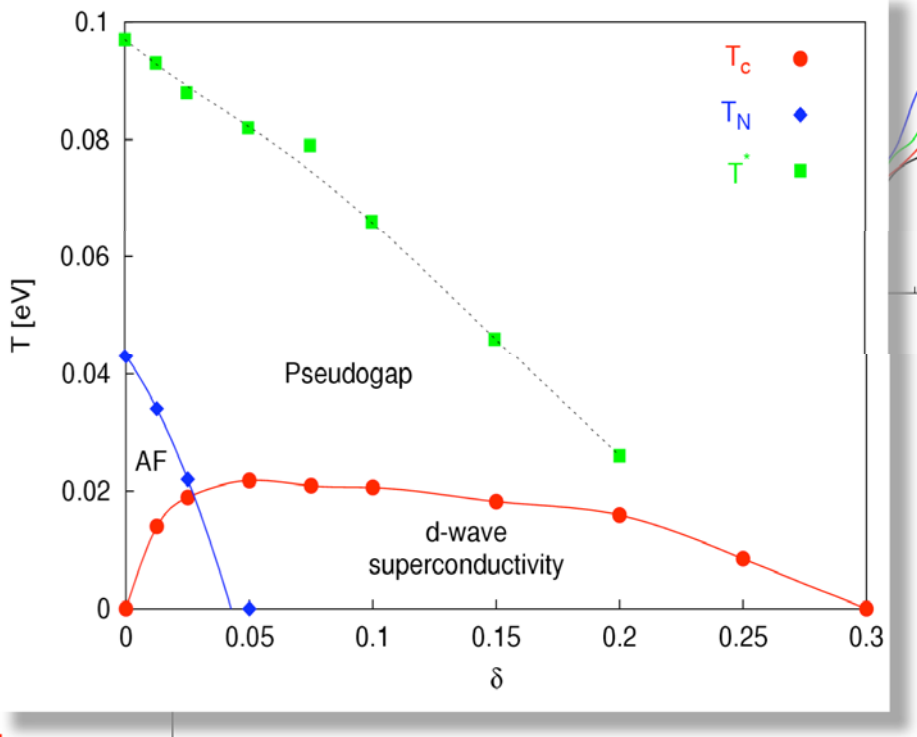

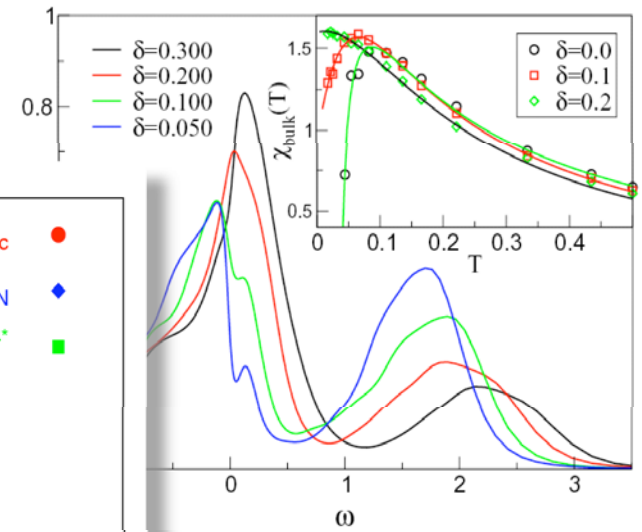

Phase diagram in remarkable agreement with experiment!

But ...

- **Antiferromagnetism**
  - **Finite *T* order in contradiction to Mermin-Wagner theorem**

- **Superconductivity**
  - **N_c=4 results represent mean-field result for d-wave order**
  - **No fluctuations of d-wave order parameter included**

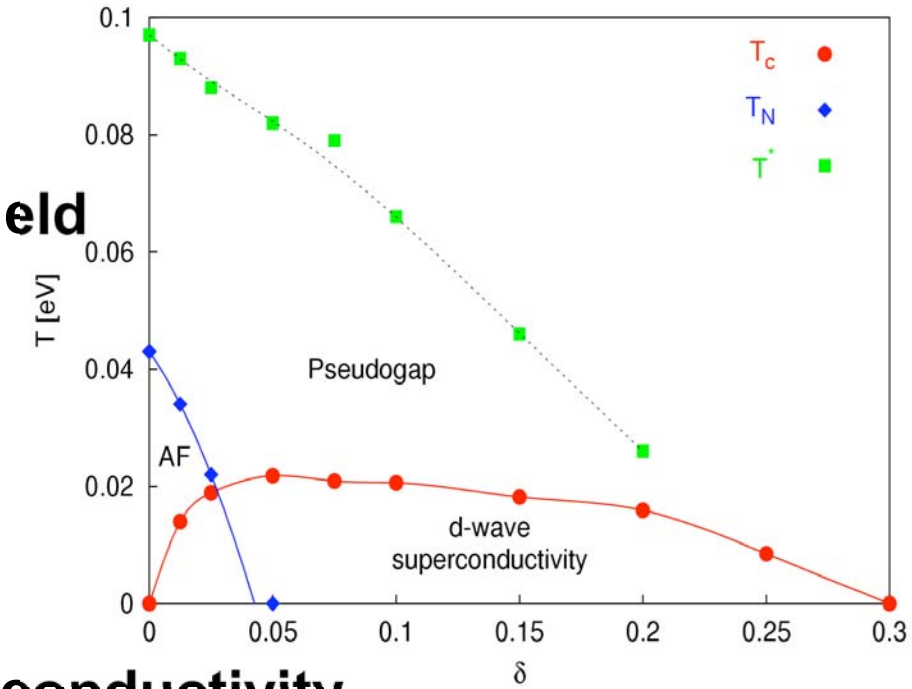

- **Questions:**
  - **KT transition to d-wave super-conductivity in larger clusters ?**
  - **Exact limit $N_c \rightarrow \infty$?**

Antiferromagnetism: Cluster size dependence

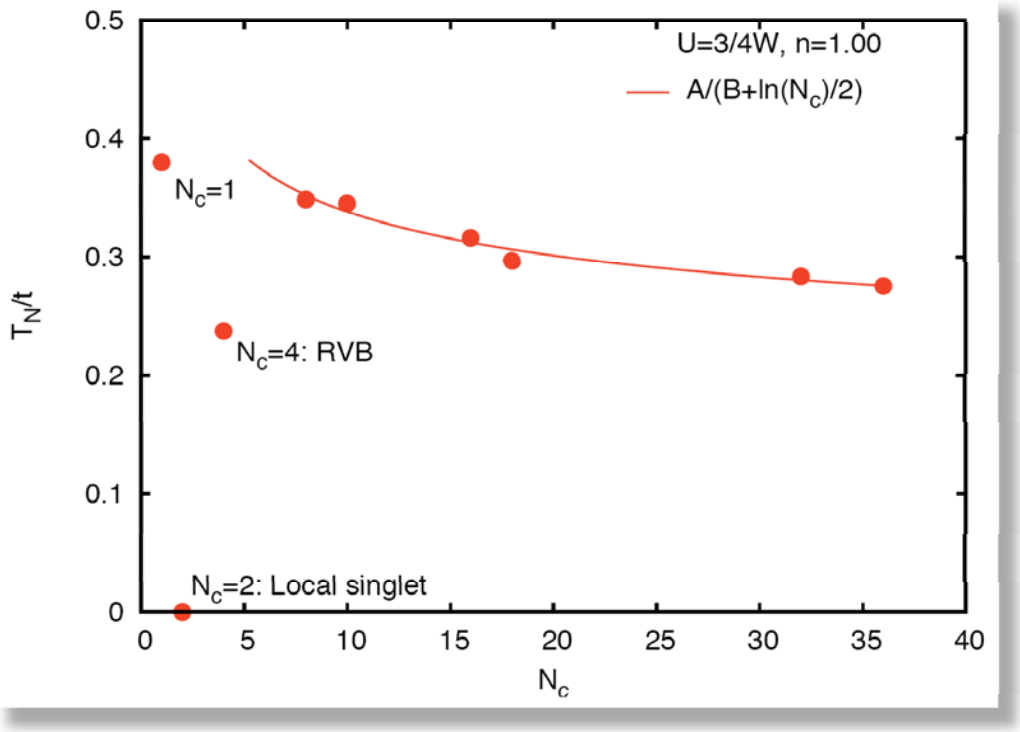

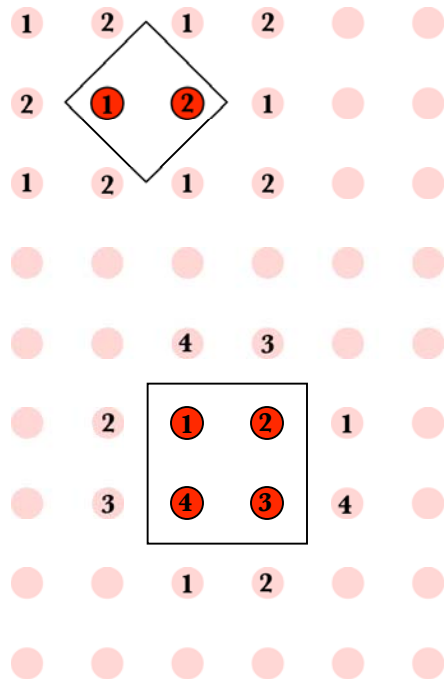

Nc=2:
1 neighbor

Nc=4:
2 neighbors

- ***$T_N \rightarrow 0$* logarithmically with *$N_c \rightarrow \infty$***
  (since AF correlations build exponentially with decreasing *T*)

d-wave pairing susceptibility (U=4t; n=0.90)

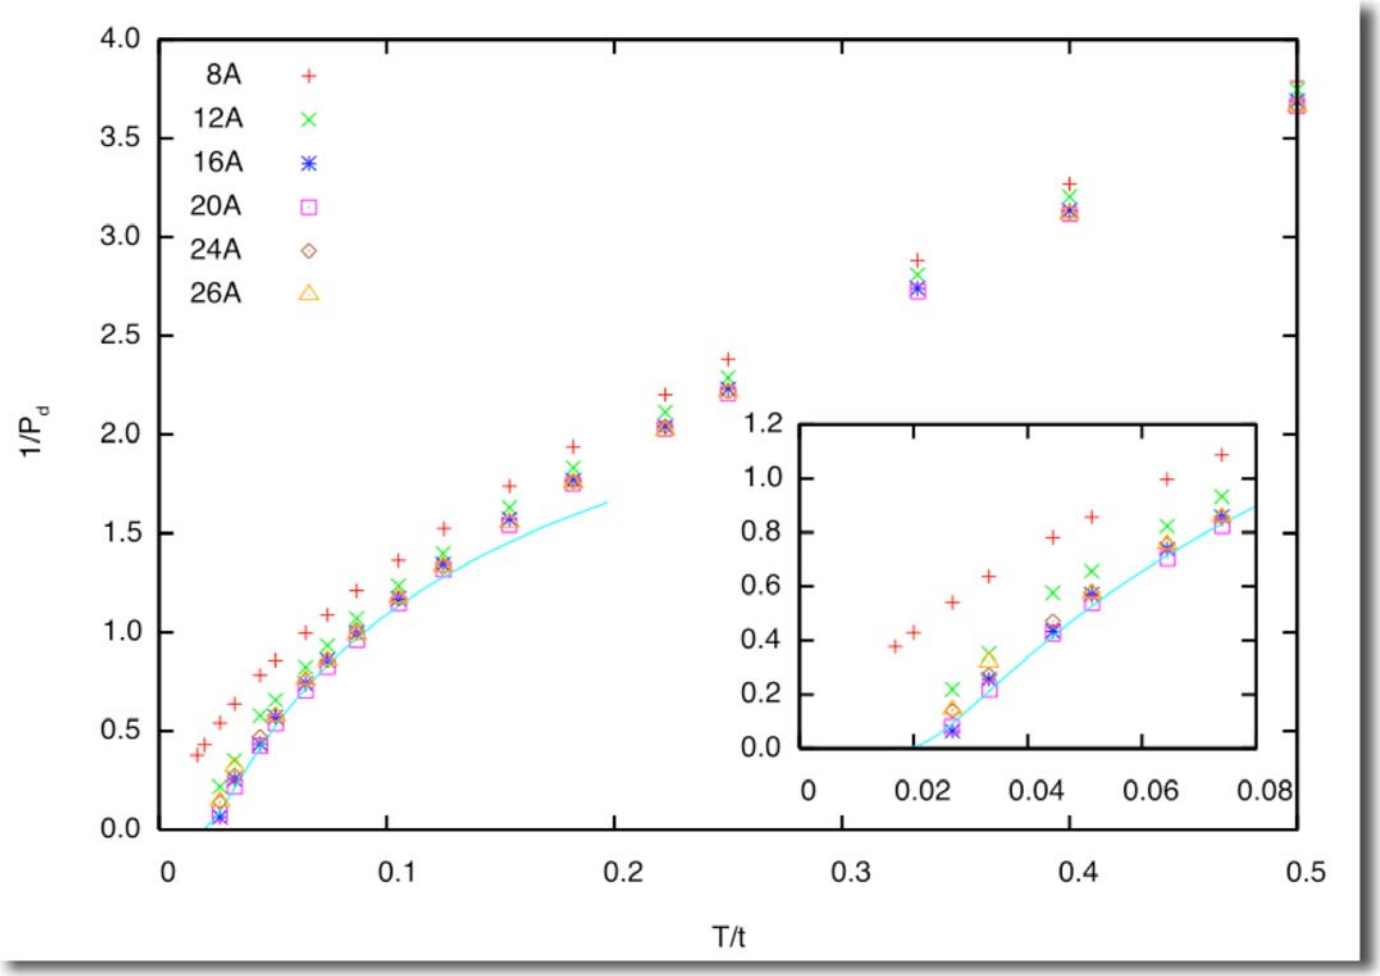

Cluster	Z_d
8A	1
12A	2
16B	2
16A	3
20A	4
24A	4
26A	4

$T_c \approx 0.025t$

High-T_c Work: Summary

- **What we accomplished:**
  - **Big step towards a solution of the most studied model in condensed matter science**
  - **Will guide the future work on theory of high temperature superconductivity**
- **Where it was done:**
  - **Small calculations (IBM Power 4, SGI Altix)**
  - **Breakthrough calculations: Cray X1 (production runs of 12-24h at ~50% of peak on up to 500 nodes)**
- **Who did the work:**
  - **Thomas Maier (ORNL Wigner Fellow)**
  - **Trey White (ORNL / CCS)**
  - **Mark Jarrell (Univ. of Cincinnati) & TCS (ORNL)**
- **Invited talk at upcoming APS March Meeting, manuscript in final stage of preparation**

Fe nanoparticle in Y-stabilized Zirconia

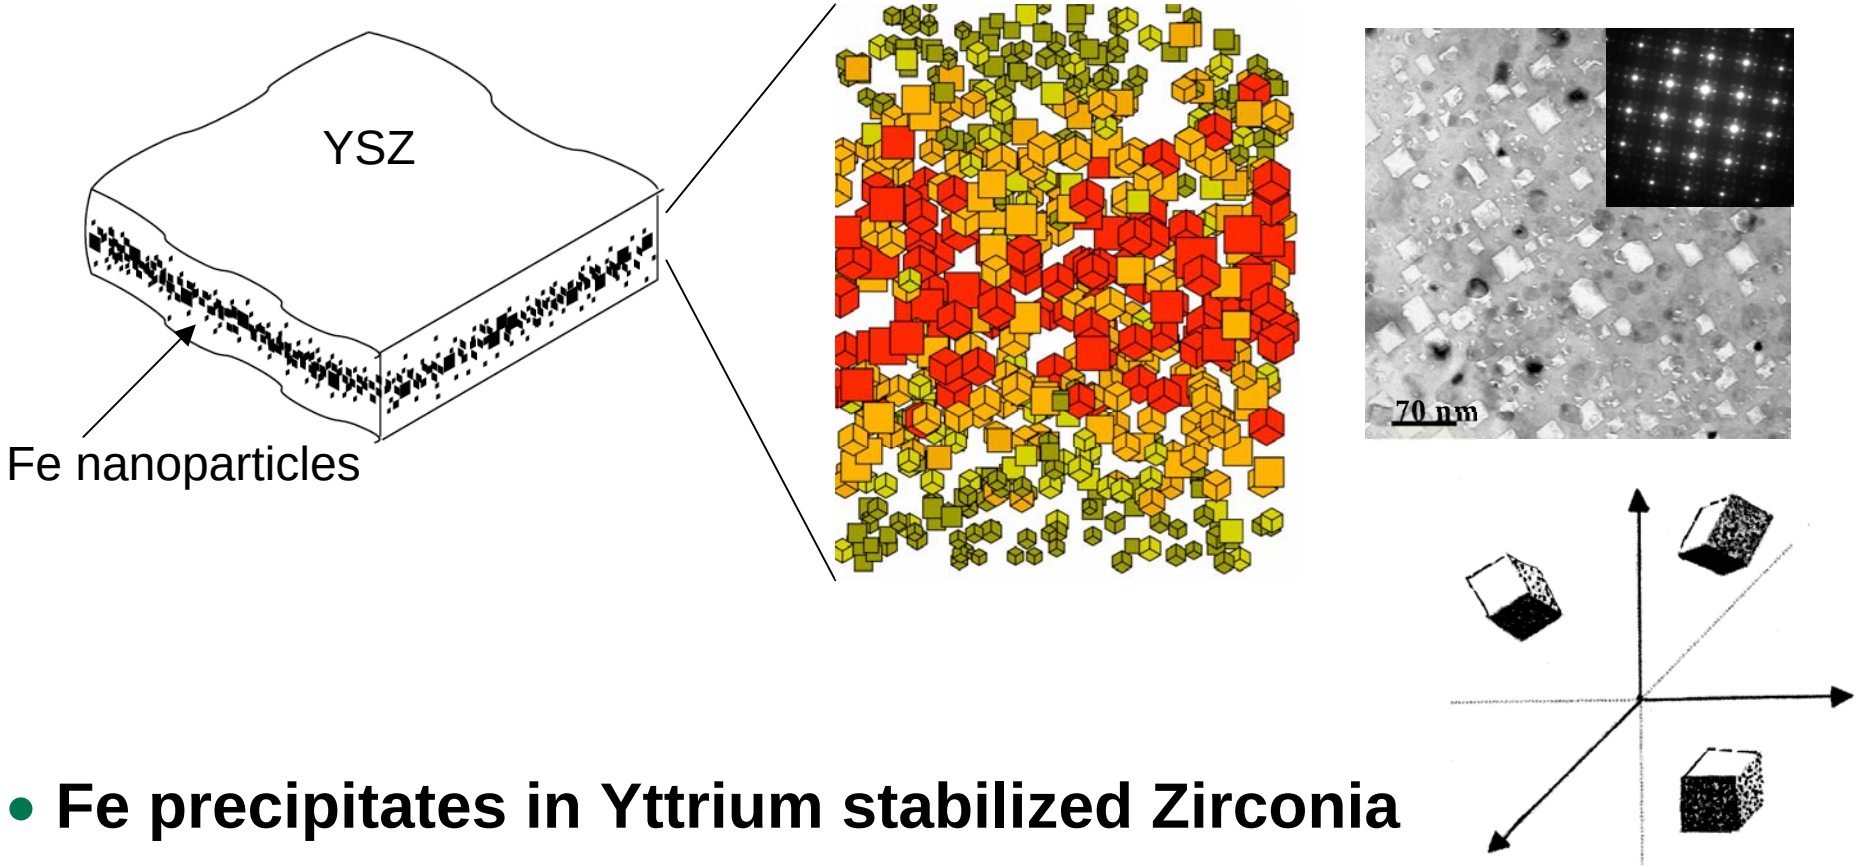

- **Fe precipitates in Yttrium stabilized Zirconia**
- **Particles are coherent with matrix**
- **Packing fraction, orientation, and size distribution known**

Constructing the Model

Insert particles uniformly in to a slab oriented along three directions

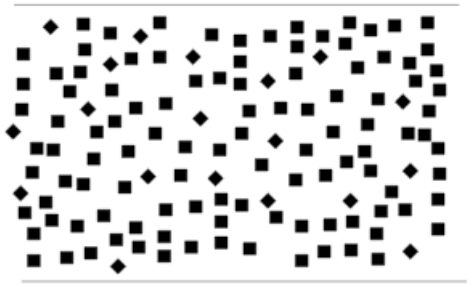

Voronoi construction to determine relative volume

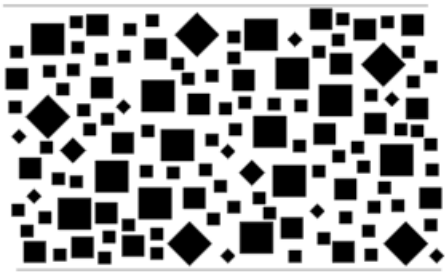

Scale with
Gaussian profile
in perp. dirrection

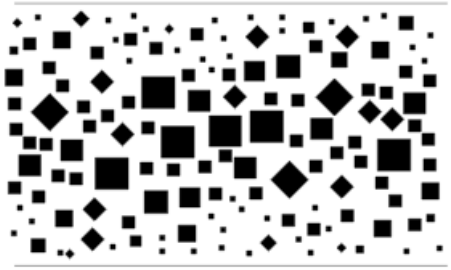

Hamiltonian (no fitting parameters)

$$H = -\frac{1}{2}\sum_{i \neq j} \vec{m}_i \left(\frac{3\hat{r}_{ij}\hat{r}_{ij} - \hat{\hat{1}}}{\left|\vec{r}_{ij}\right|^3} \right) \vec{m}_j +$$

$$K_0 \sum_i V_i \left[(\hat{K}_1^i.\hat{m})^2 (\hat{K}_2^i.\hat{m})^2 + (\hat{K}_2^i.\hat{m})^2 (\hat{K}_3^i.\hat{m})^2 + (\hat{K}_3^i.\hat{m})^2 (\hat{K}_1^i.\hat{m})^2 \right]$$

Bulk BCC Fe

Temperature dependence of remnant magnetization

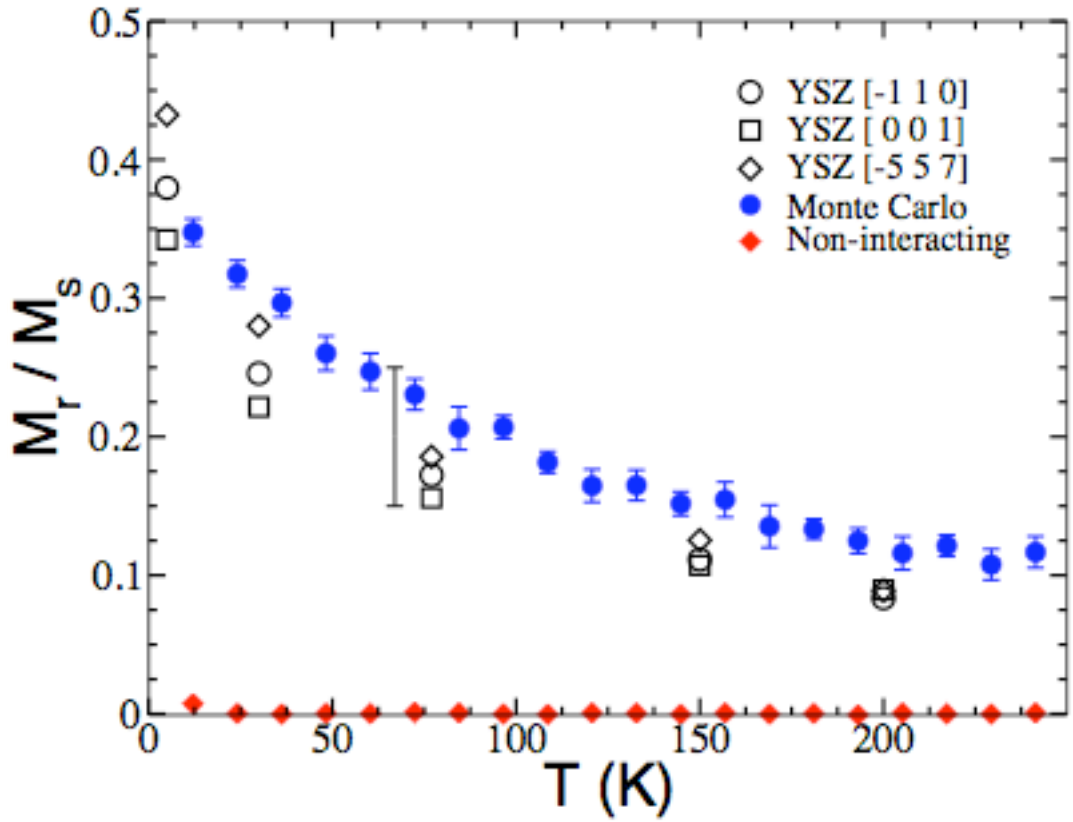

Fe nano-particles on Cu (111)

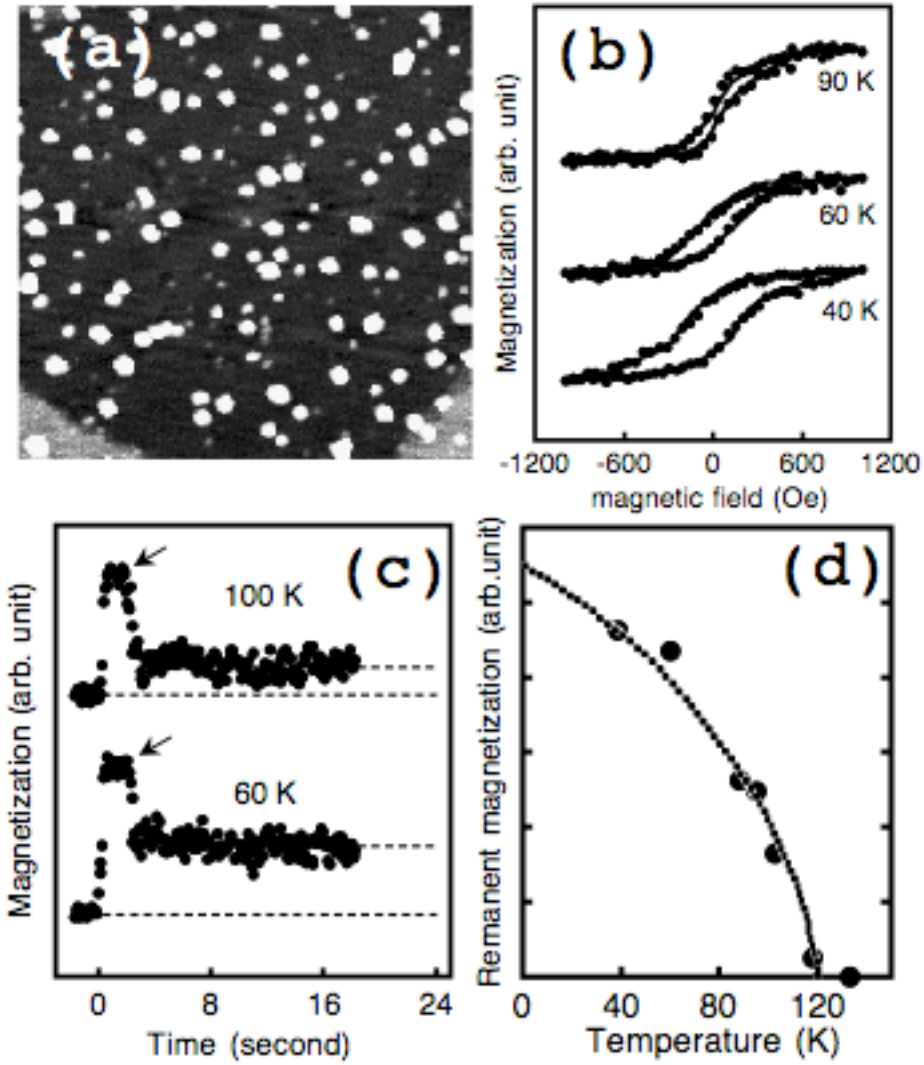

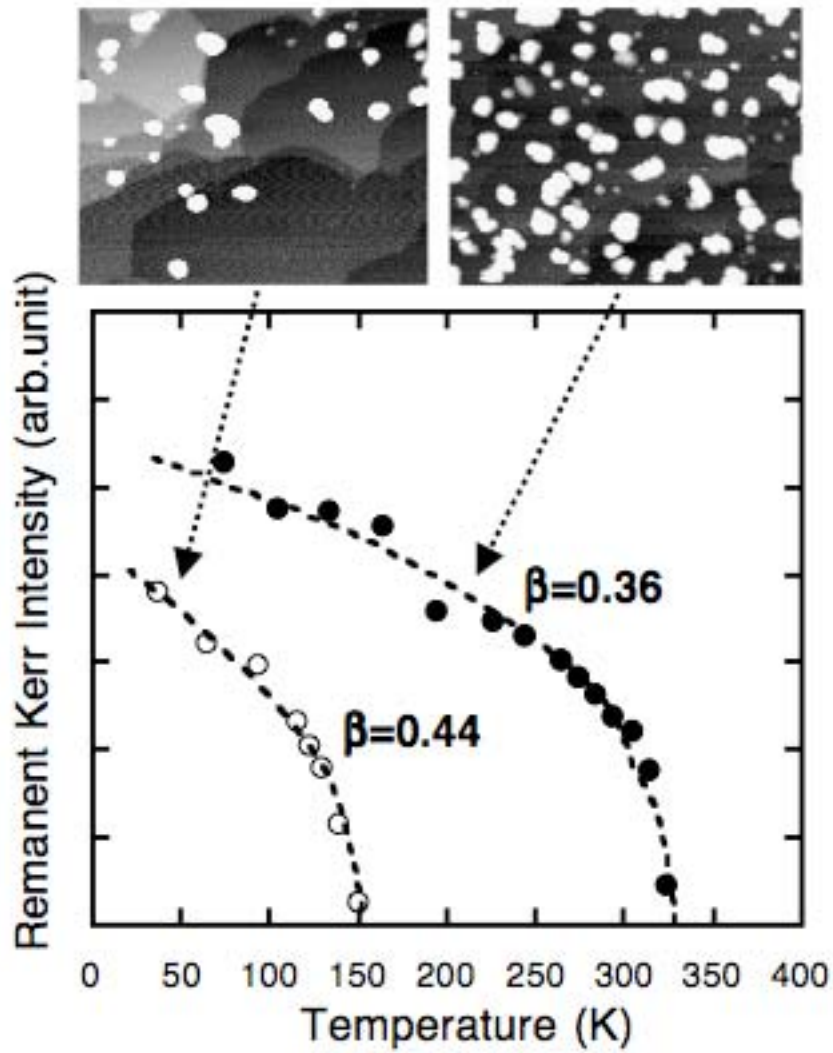

Constructing the model in this case

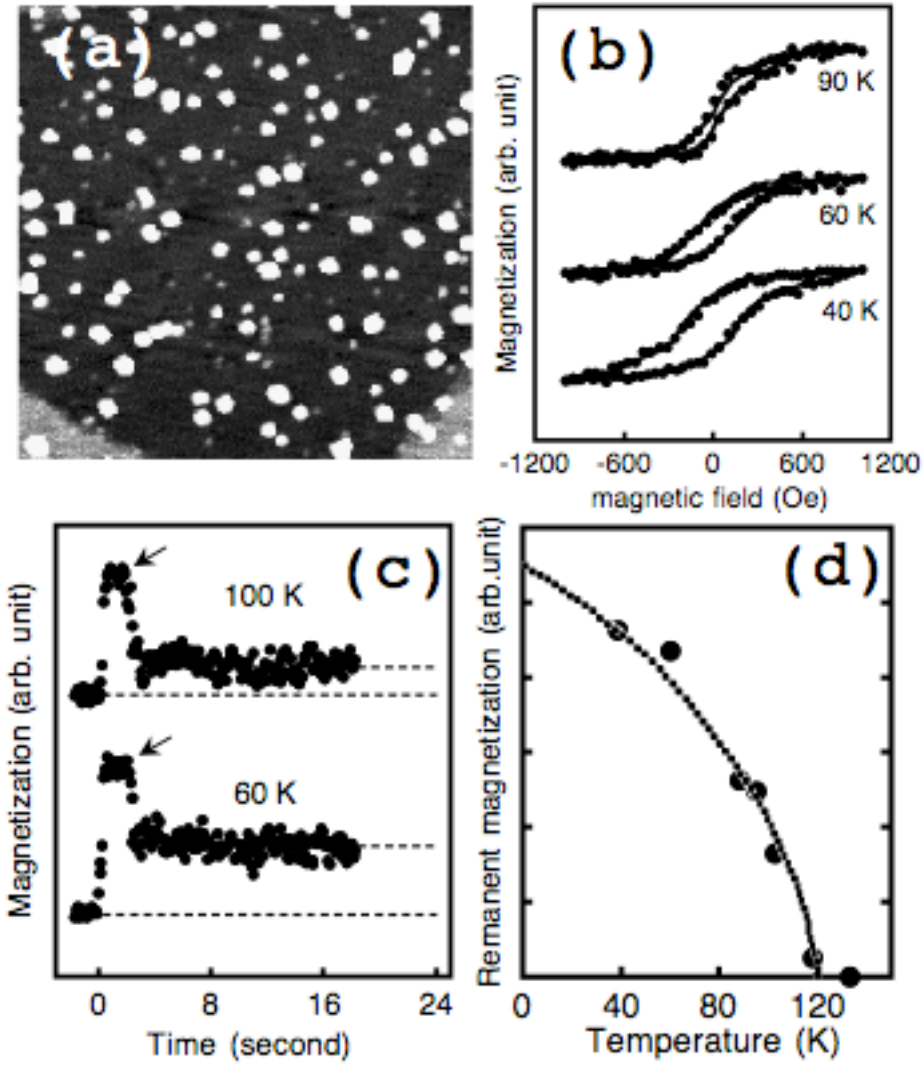

- **Positions and sizes known from experiment**
- **Hamiltonian**

$$H = -\frac{1}{2}\sum_{i \neq j} \vec{m}_i \left(\frac{3\hat{r}_{ij}\hat{r}_{ij} - \hat{\hat{1}}}{\left|\vec{r}_{ij}\right|^3} \right) \vec{m}_j + K \sum_i V_i \cos^2 \vartheta_i$$

- **Uniaxial anisotropy with only free parameter *K***
- **Monte Carlo simulations starting from ordered state**

Monte Carlo Simulations

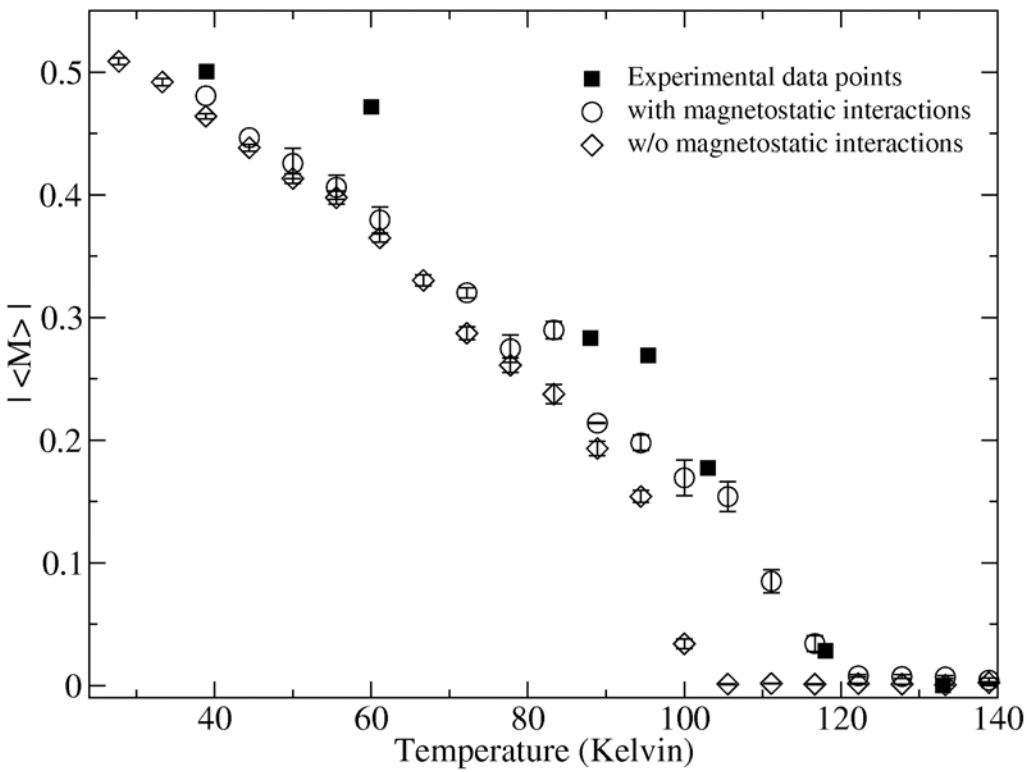

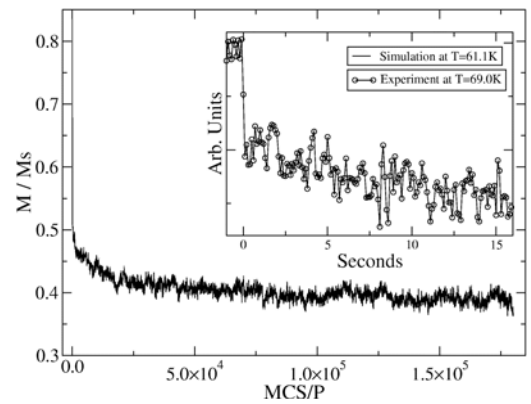

$$H = -\frac{1}{2}\sum_{i \neq j} \vec{m}_i \left(\frac{3\hat{r}_{ij}\hat{r}_{ij} - \hat{\hat{1}}}{\left|\vec{r}_{ij}\right|^3} \right) \vec{m}_j + K \sum_i V_i \cos^2 \vartheta_i$$

- **Magnetostatic interaction play minor role**
- **Need to increase anisotropy by order of magnitude compare to bulk bcc Fe (?)**
- **But particles are hemispherical**
  - **High anisotropy is hard to justify**

Are we missing something in the model?

Discovering a new mechanism

- **Cu (111) has a well known surface state**
- **RKKY like interaction between Fe nanoparticles**
  - **2D electron gas, $1/r^2$ decay for atoms**
- **Check hypothesis experimentally by changing substrates**

Pierce et. al.,
Phys. Rev. Lett., (2004)

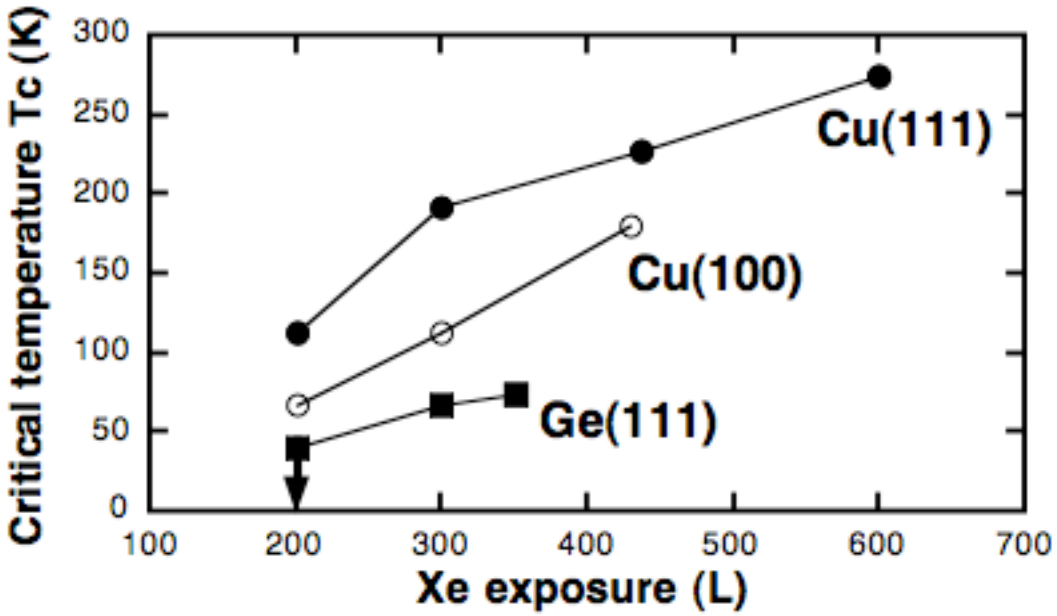

- **Check Hypothesis in computational model**
  - **Will require more computing power (…)**

Extensible "Heisenberg" Model

$$E\left[\left\{\vec{m}_i\right\}\right] = -\sum_{i \neq j} J_{ij} \vec{s}_i \vec{s}_j - \frac{1}{2} \sum_{i \neq j} \vec{m}_i \frac{3\hat{r}_{ij}\hat{r}_{ij} - \hat{\hat{1}}}{\left|\vec{r}_{ij}\right|^3} \vec{m}_j + \sum_i K_i \sin\theta_i + \cdots$$

- **Flexibly add / remove energy terms**
- **Efficient algorithms to evaluate energy terms**
  - **Long range interactions (FFT, FMM, Ewald)**
- **Various types of Monte Carlo sampling schemes**
- **Spin dynamics simulations**
  - **Regular spin dynamics; Trotter dynamics; Stochastic dynamics**

Needed Extensions

$$E\left[\{\vec{m}_i\}\right] = -\sum_{i\neq j} J_{ij}\vec{s}_i\vec{s}_j - \frac{1}{2}\sum_{i\neq j}\vec{m}_i \frac{3\hat{r}_{ij}\hat{r}_{ij} - \hat{\hat{1}}}{\left|\vec{r}_{ij}\right|^3}\vec{m}_j + \sum_i K_i \sin\theta_i + \cdots$$

- **Spin Fermion Models (quantum)**

$$H = -t \sum_{<ij>\sigma} c^{\dagger}_{i\sigma} c_{j\sigma} - J_H \sum_{i\alpha\beta} c^{\dagger}_{i\alpha} c_{i\beta} \vec{S}_i \cdot \vec{\tau}_{\alpha,\beta}$$

- **First principles exchange functional**

$$E\left[\{\vec{m}_i\}\right] \rightarrow E\left[n(\vec{r}), \vec{m}(\vec{r})\right]$$

- **Will need to scale to HPC systems to run!**
- **But do we have to rewrite a new code?**

Implementation: The *Ψ-Mag* Toolkit

- **C++ library for computational magnetism and serves as a prototype for a more general library for computational materials science.**
  - **New effort: expand *Ψ-Mag* for molecular simulations**

- **Design is modeled and inspired by the generic programming techniques of the C++ standard template library (STL).**

Language: C++ and the STL

- **C++ programming language:**
  - **Object oriented "extension" of C**
    - →*Can be as efficient at C*
  - **Support of *templates* (replacement of pre-processor)**
    - →*Templates allow for compile time polymorphism*
- **Standard Template Library (STL)**
  - **Generic algorithms and data structures (lists, sort, …)**
  - **Designed for extensibility: "to use STL is to extend it"**
  - **Algorithms are decoupled from containers**
  - **Extensible and customizable without inheritance**
  - **Abstractions based on compile time polymorphism**
- **C++ (and STL) exist on HPC just like C & Fortran**

Ψ-Mag Toolkit: an extension of STL

- **No claim for completeness:**
  - **Implement carefully what is needed**
  - **Extensible design makes adding tools rather simple**
- **Layered software structure with minimal interdependence**
  - **no horizontal dependence**
  - **Individual classes are simple**

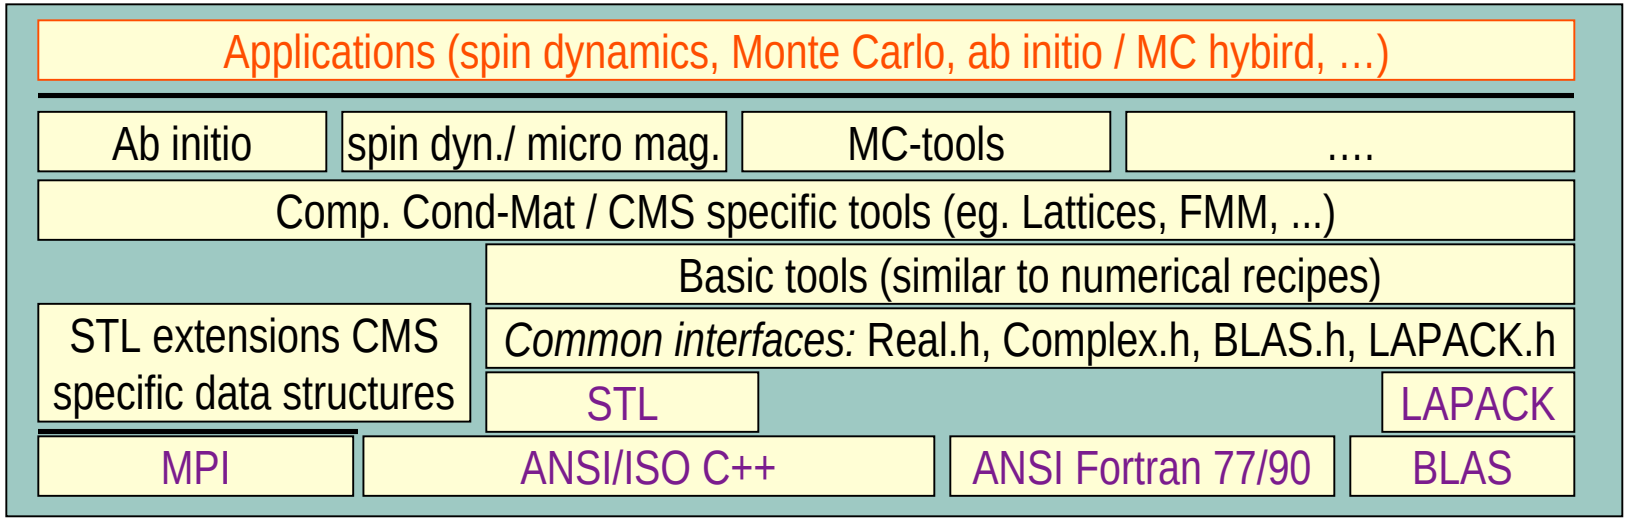

- **Applications using *Ψ-Mag*:**
  - **Can access all layers directly**
  - **Small and therefore flexible / maintainable**

Basic Libraries
(required, not part of Ψ-*Mag*)

- **C++ Standard Template Library (STL)**
  - **Always required**
- **Basic Linear Algebra Subroutines (BLAS)**
  - **Look for optimized implementation**
- **Linear Algebra Package (LAPACK)**
- **Fast Fourier Transforms (FFT):**
  - **FFTW (default), or try native implementation (interface in toolkit still a research issue)**
- **Message Passing Interface (MPI)**

A Meaningful Evolution of Software?

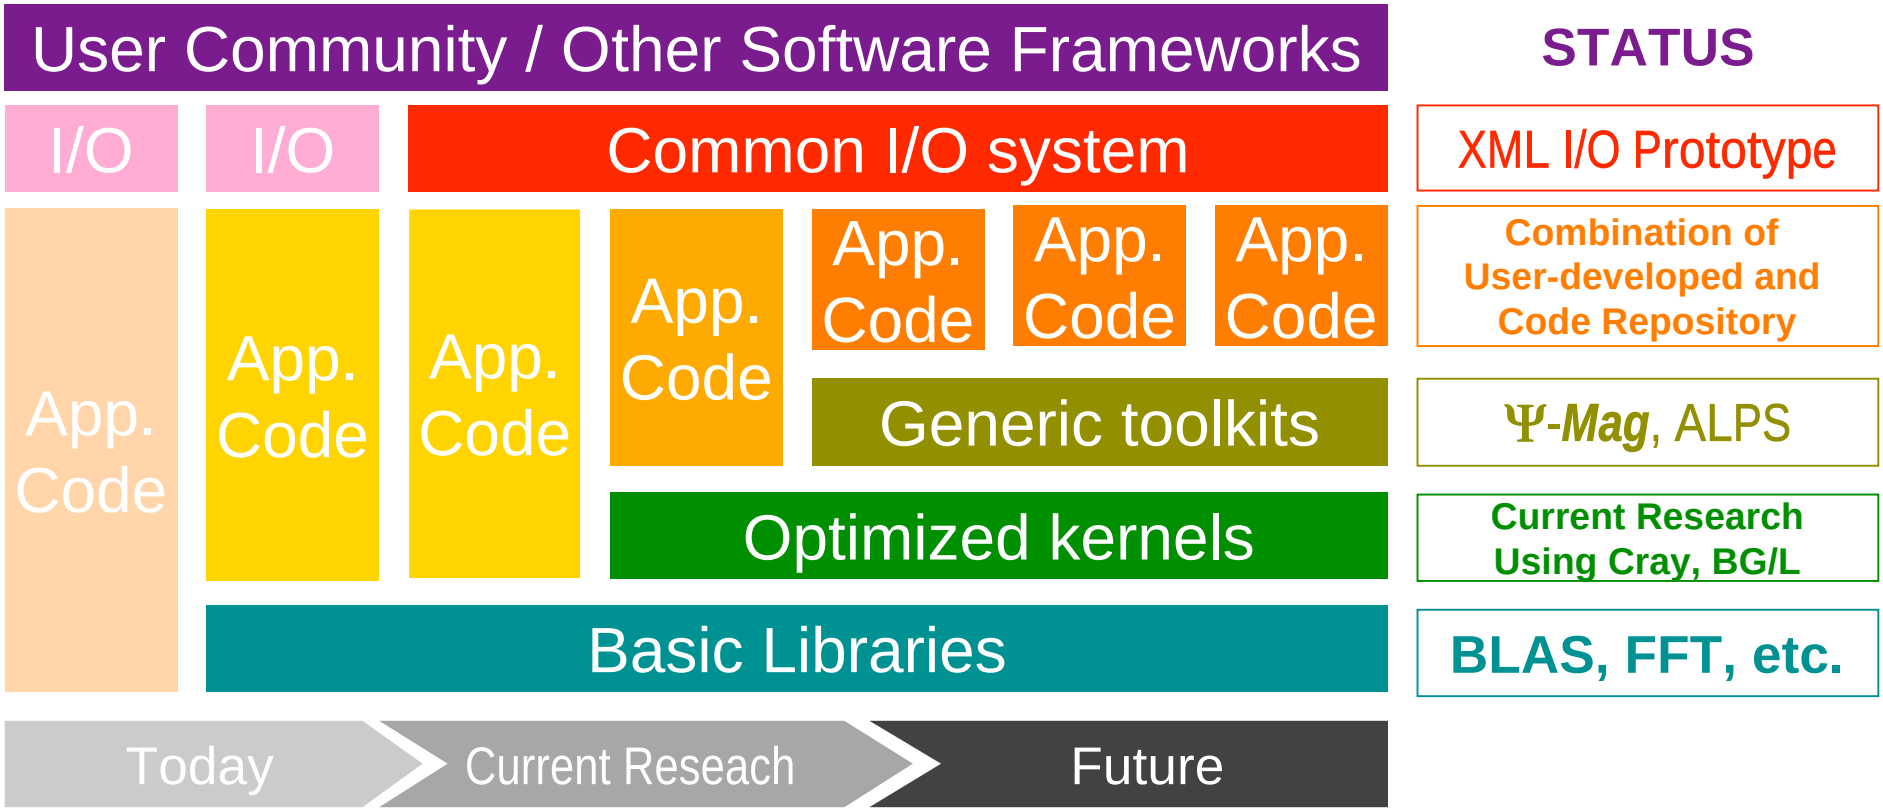

- **Reduce size of application code**
- **Improve usability of codes**
- **Enhance productivity / creativity**

- **Optimize performance / efficiency**
- **Simplify user access to machines**
- **Prioritization of resources**

New Approach: facility analogy

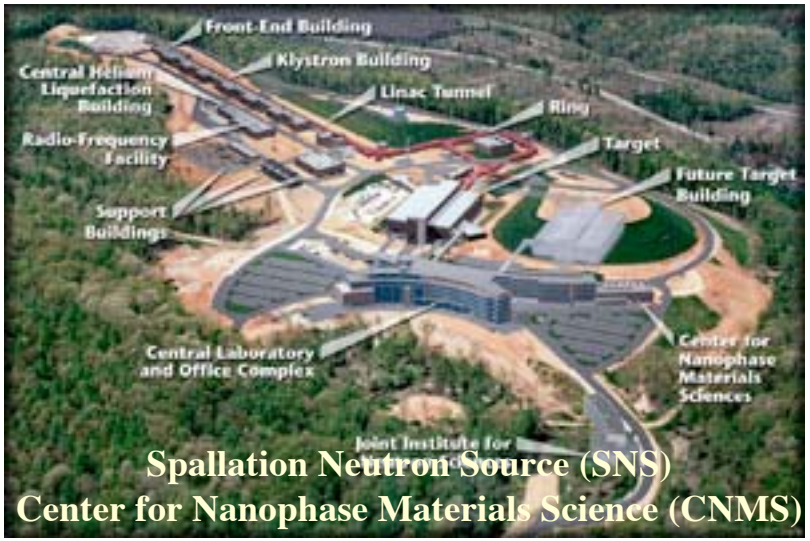

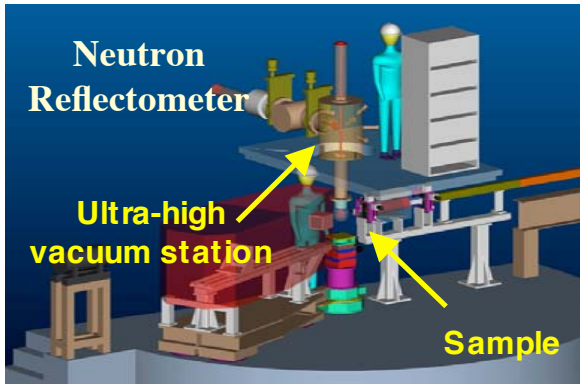

User Community

Facility

Instrumentation

Materials Science Virtual User Center

Materials : Math : Computer Scientists

- *Open Source Repository*
- *Object Oriented Tool Kit*
- *Focused User Laboratories*
- *Education*

User Community

Conclusions

- **Leadership computing (or HPC) is a great opportunity for materials, nano-, and condensed-matter science**
  - **HPC has moved well beyond proof of principle!**
- **Leadership computing projects have to be integrated with science projects**
  - **Integrated teams consisting of Scientists (computational and non-computational) and HPC support (computer science, numerical analysts, etc.)**
- **Software development should not be a byproduct!**
  - **Properly designed software enables science**

Collaborators

- **Application example 1: A prototype grand challenge team (High-T_c superconductivity)**
  Thomas Maier and Trey White (ORNL); Mark Jarrell (U. of Cincinnati)
  - **Using QMC-DCA to solve the problem**
  - **QMC-DCA on scalar and vector machines**
  - **Solving the 2D Hubbard model with CCS' leadership computing capabilities**
- **Magnetic nano-particles: Discovering physics and solving the software crisis**
  Greg Brown, Hwee-Kuan Lee, Markus Esibenbach, Chengang Zhou (CMSN postdoc)
  Korye Sorge and Jim Thompson (UT/ORNL) et al.; D. Pierce and Jian Shen (ORNL) et al.;
  Malcolm Stocks (ORNL), David Landau, and Xiuping Tao (UGA)
  Michael Summers (ORNL), Tom Swain (UT), and Gonzalo Alvarez (ORNL)
  - **Experiment and simulations of magnetic properties of nanoparticle assemblies**
  - **Extensible Heisenberg Model and the *Ψ-Mag* toolkit**
  - **A meaningful evolution of software for materials science**
- **Summary and Conclusions**

Acknowledgments

This research used resources of the Center for Computational Sciences and was sponsored in part by the offices of Basic Energy Sciences and Advance Scientific Computing Research, 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.