Algorithms for simulating quantum mechanics

Bryan Clark

CSE Seminar

Why care about algorithms for quantum mechanics?

Accurate and efficient quantum mechanical simulations would have revolutionary changes to industrial applications, what we know about our universe and the way physics is done.

We know the rules.

Just need to figure out how to solve them.

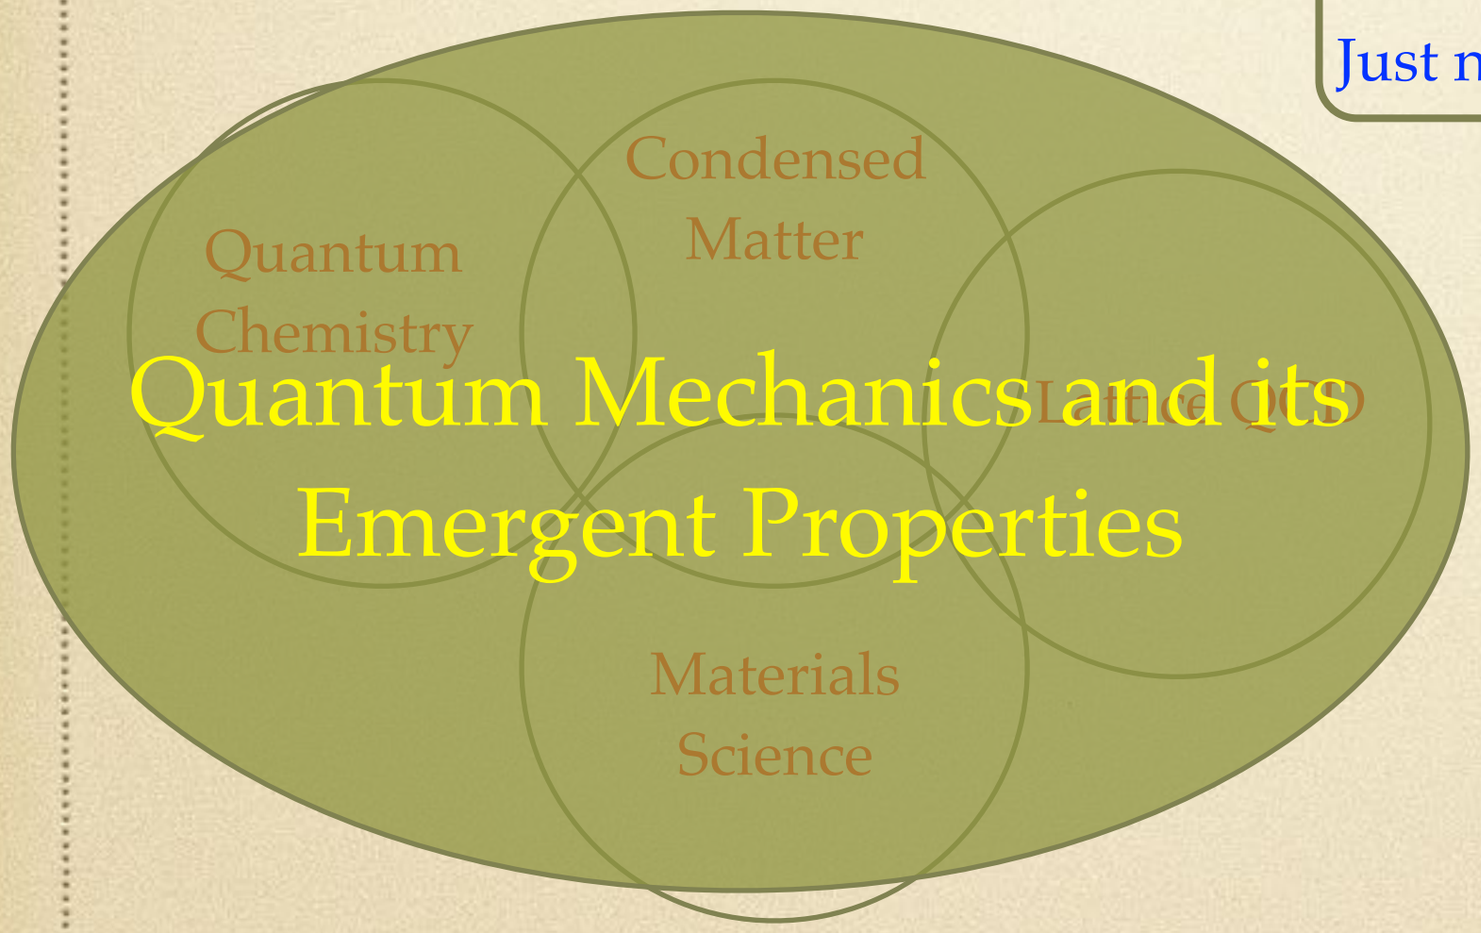

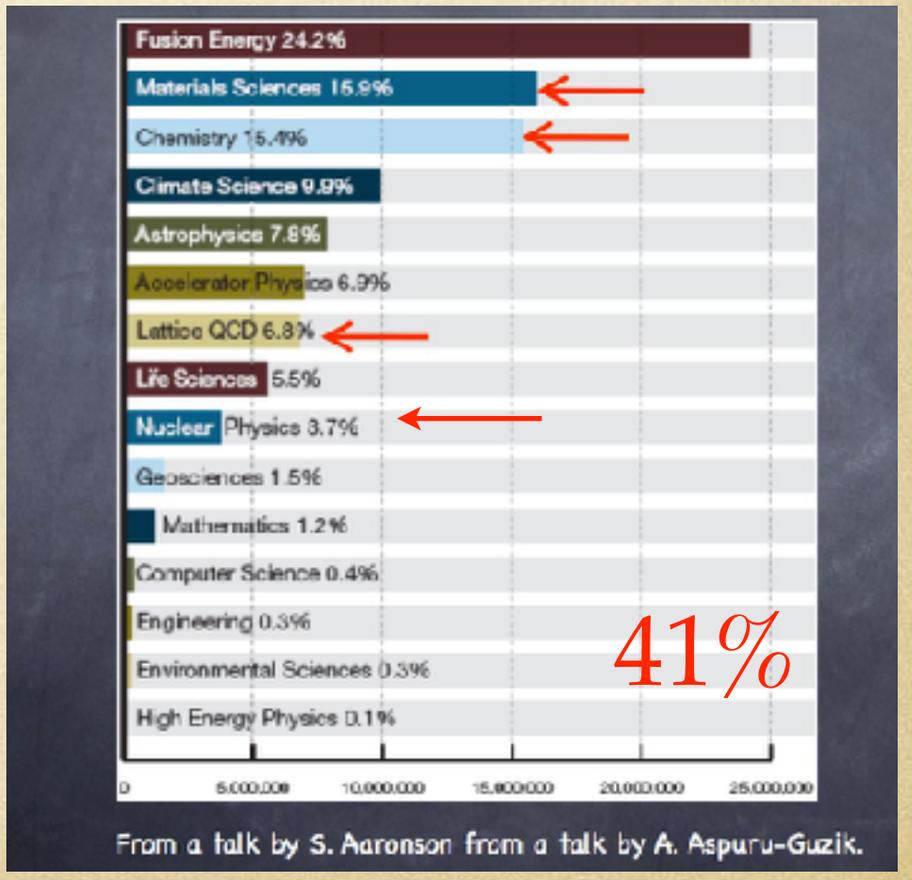

The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble. - Paul Dirac (1929)

It's lucky for us that they solved this problem a long time ago...

Equation of State Calculations by Fast Computing Machines

THE purpose of this paper is to describe a general method, suitable for fast electronic computing machines, of calculating the properties of any substance which may be considered as composed of interacting individual molecules. Classical statistics is assumed,

(Received March 6, 1953)

NICHOLAS METROPOLIS, ARIANNA W. ROSENBLUTH, MARSHALL N. ROSENBLUTH, AND AUGUSTA H. TELLER,
Los Alamos Scientific Laboratory, Los Alamos, New Mexico

AND

EDWARD TELLER,* *Department of Physics, University of Chicago, Chicago, Illinois*

Unfortunately almost everything is a fermion.

Markov Chain Monte Carlo techniques still amongst the best and most used techniques.

Even water is hard…

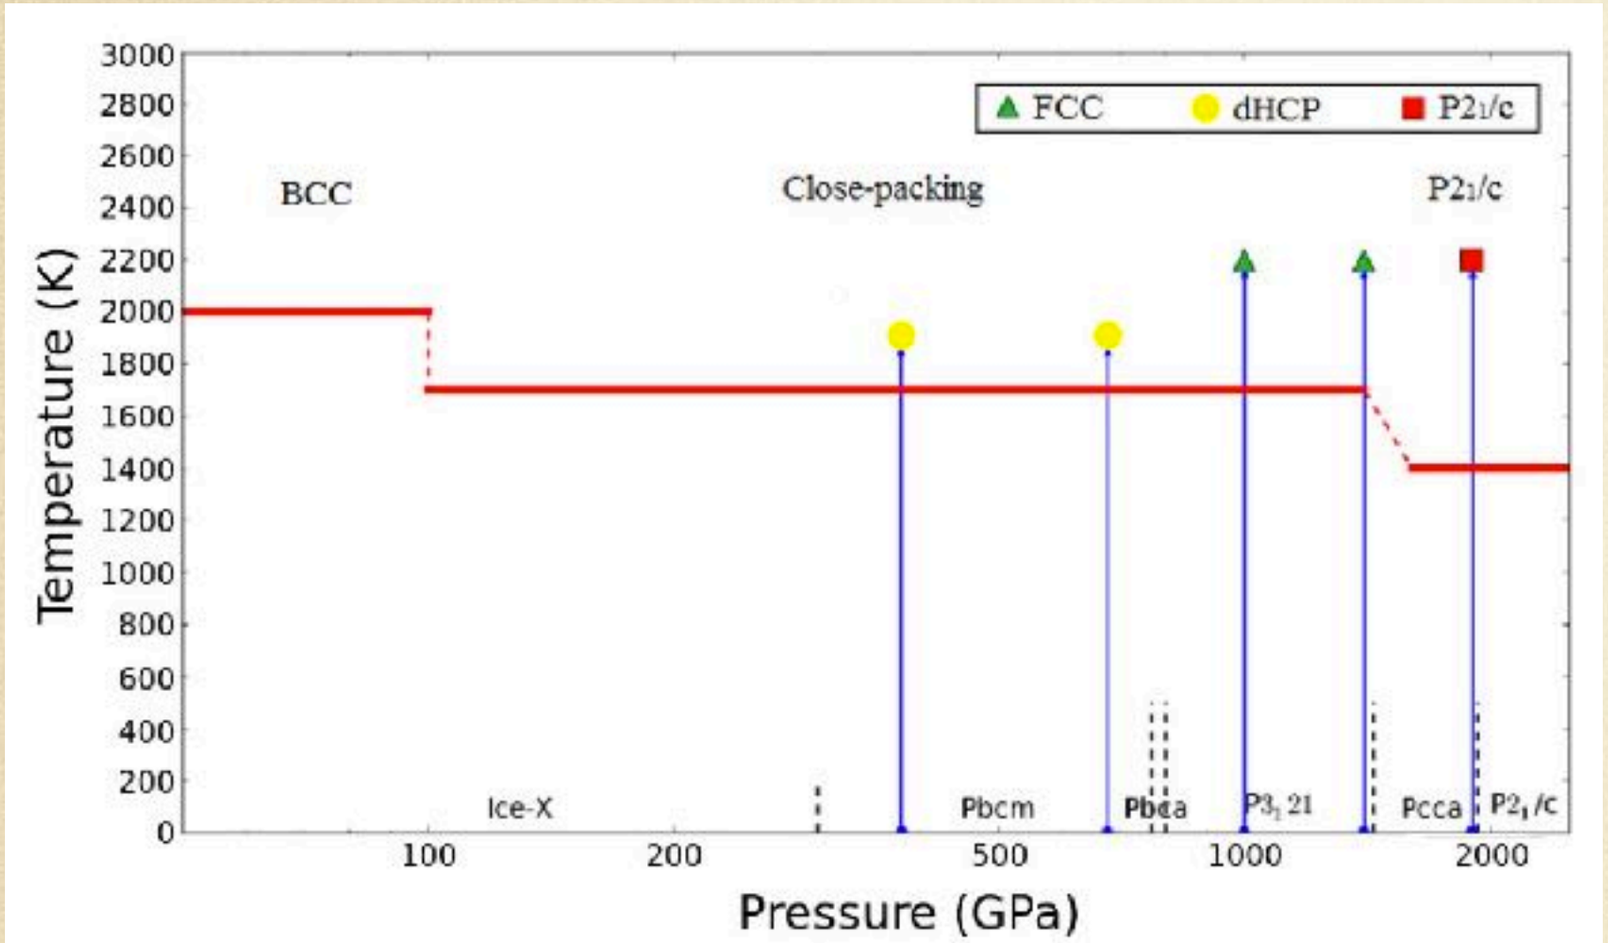

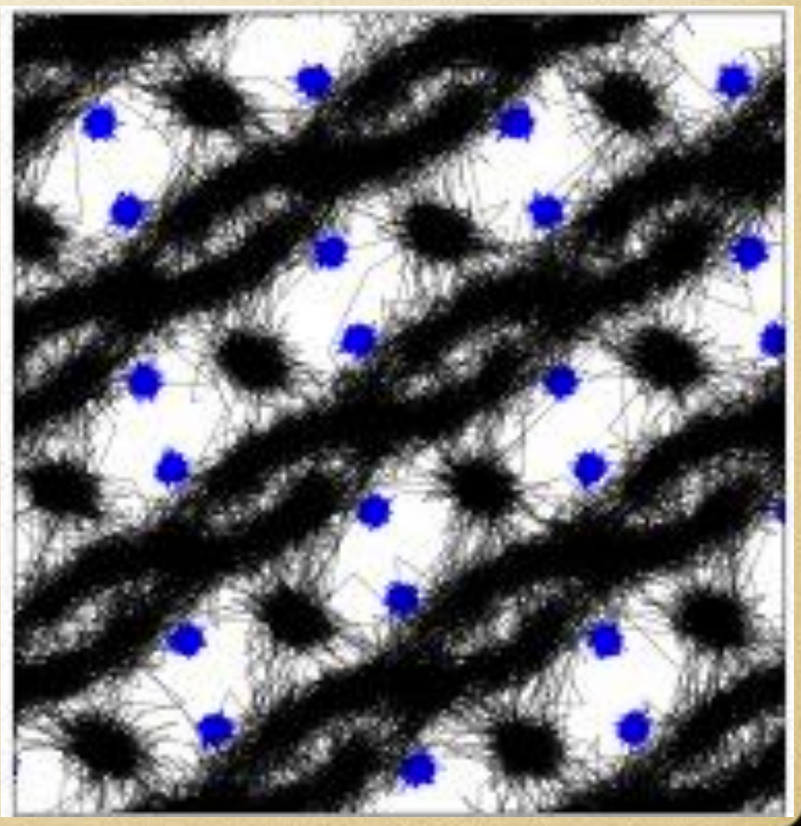

Since that time, we've seen that there appears to be an exponential wall to simulating quantum systems

Why should that be?

2^n possible electron configurations

Quantum mechanics tells us we are simultaneously in a superposition of electronic configurations.

$$\frac{1}{\sqrt{2}}$$ 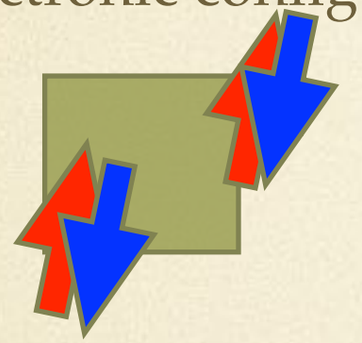 $$+ \quad \frac{1}{\sqrt{2}}$$ 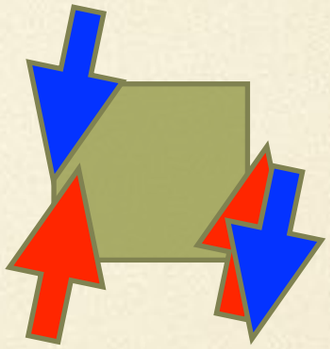

Superconductor: electrons like to be on top of each other .

Our goal: Find out where the electrons want to be.

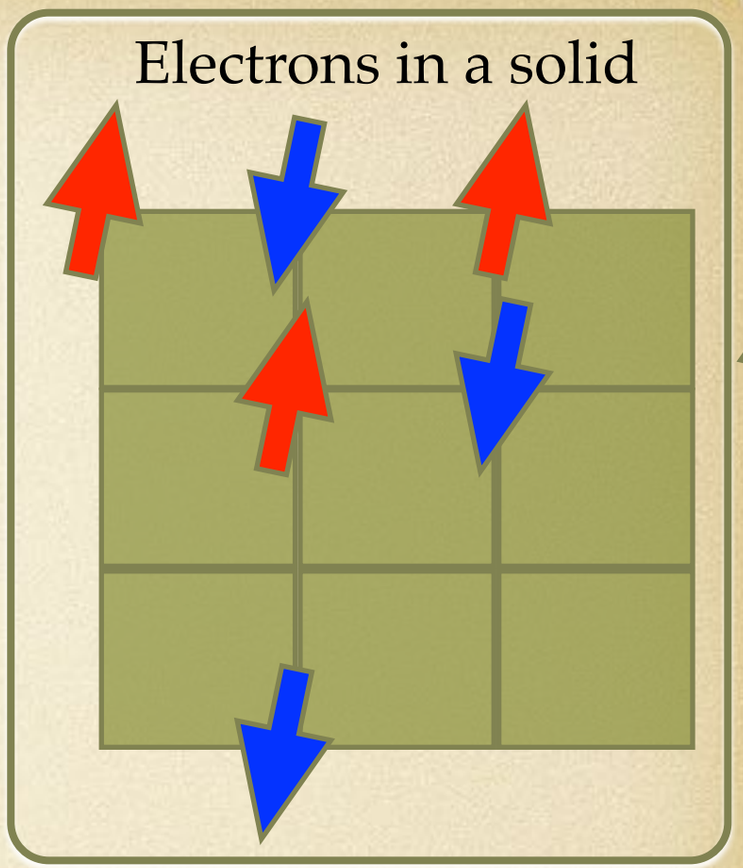

Minimal eigenstate of $H = -t \sum_{ij} c_i^\dagger c_j + U \sum_i c_{i\uparrow}^\dagger c_{i\uparrow} c_{i\downarrow}^\dagger c_{i\downarrow}$

2^n x 2^n matrix

$$\begin{pmatrix} \cdot & \cdot & \cdot \end{pmatrix}$$

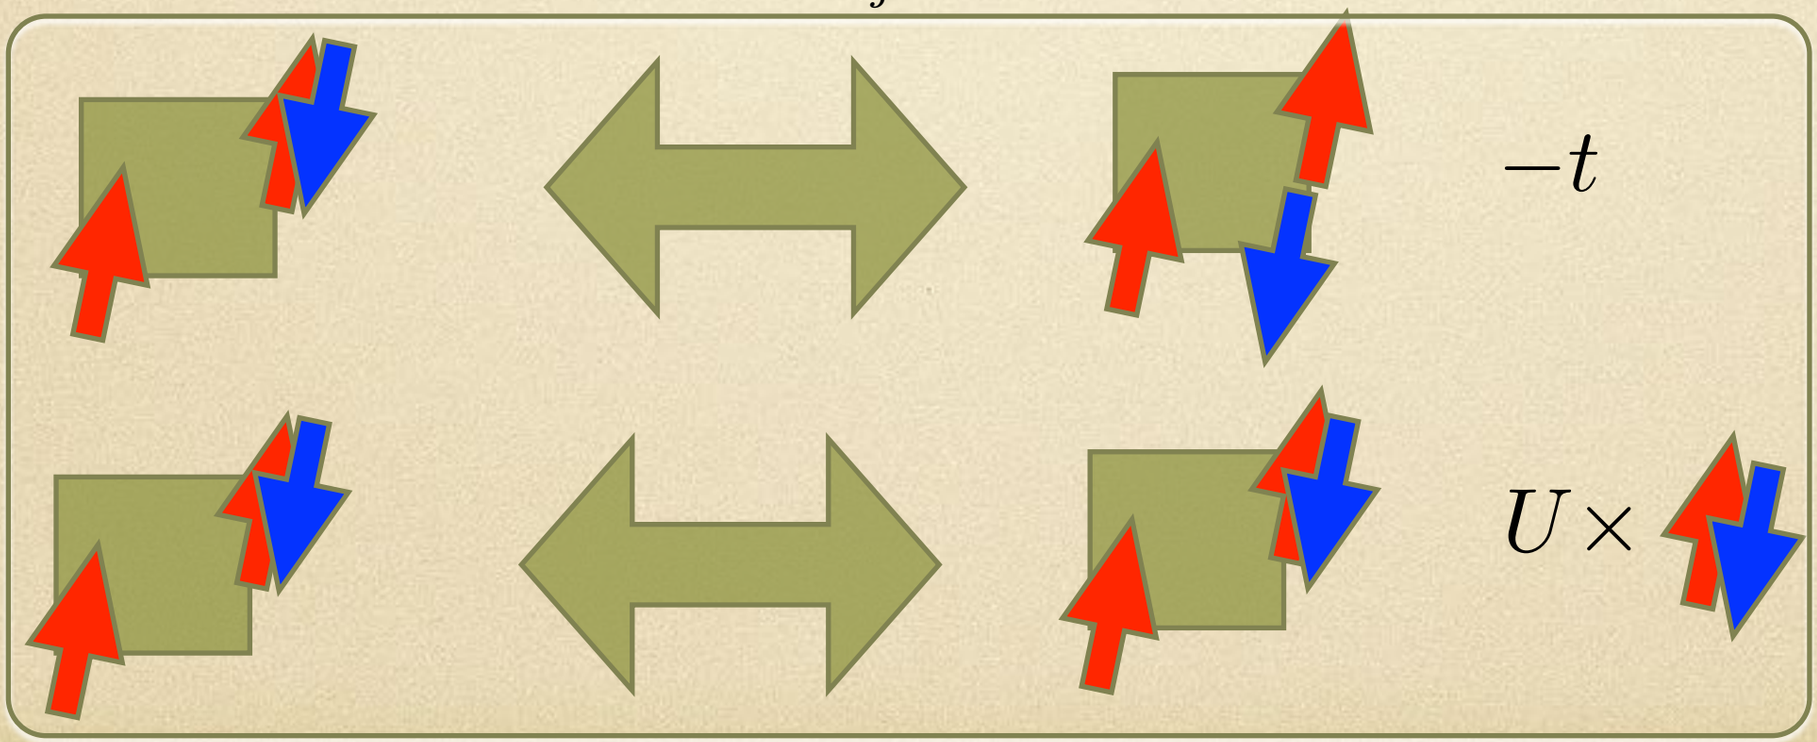

Linear Algebra

$\underline{H}\Psi = E\Psi$ via Lanczos

2^n x 2^n matrix

State of the art: 48 spins

24 electrons

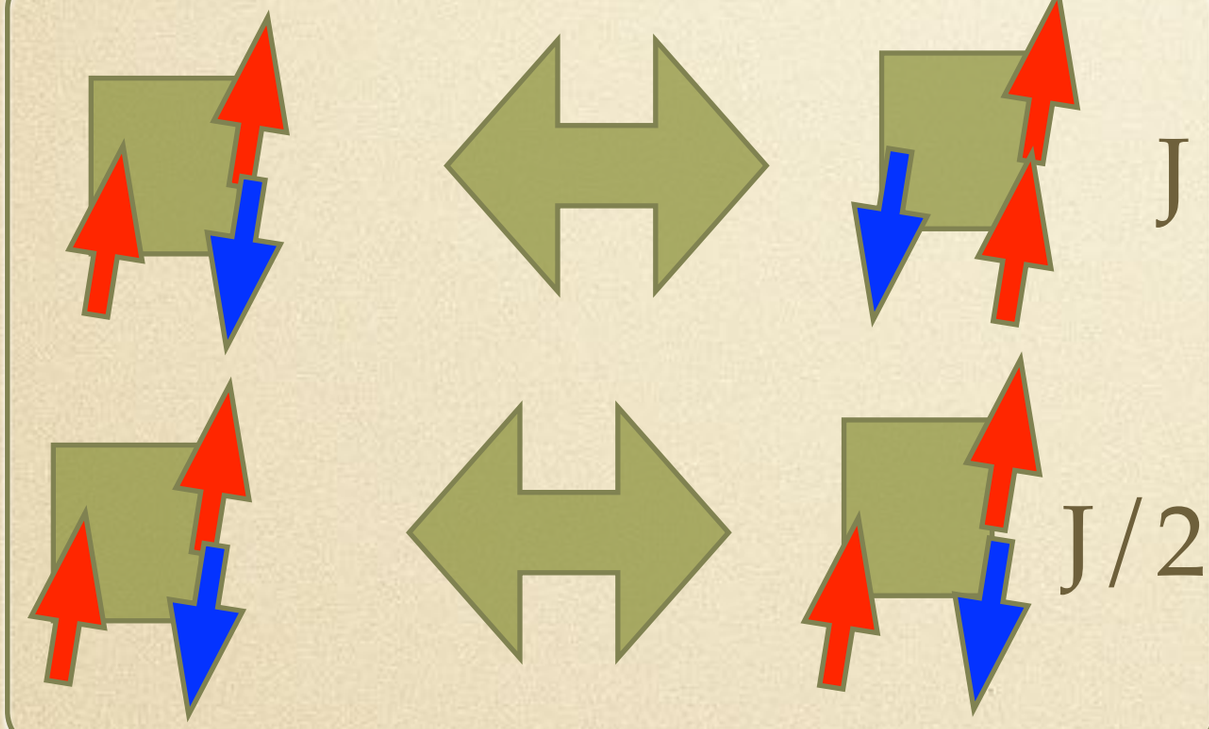

Each step is exponential

Spin liquid?

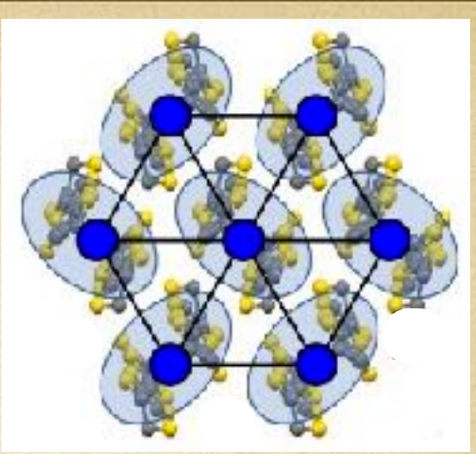

$LiZn_2Mo_3O_8$

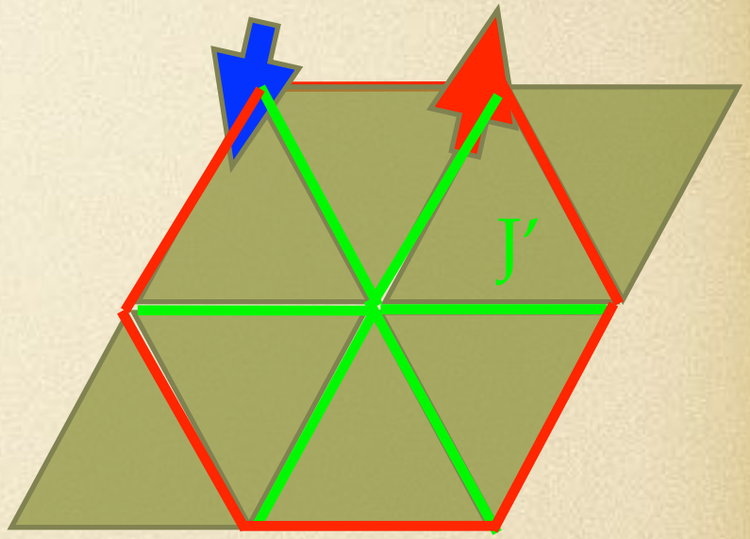

$\frac{\partial^2 E}{\partial^2 J'}$

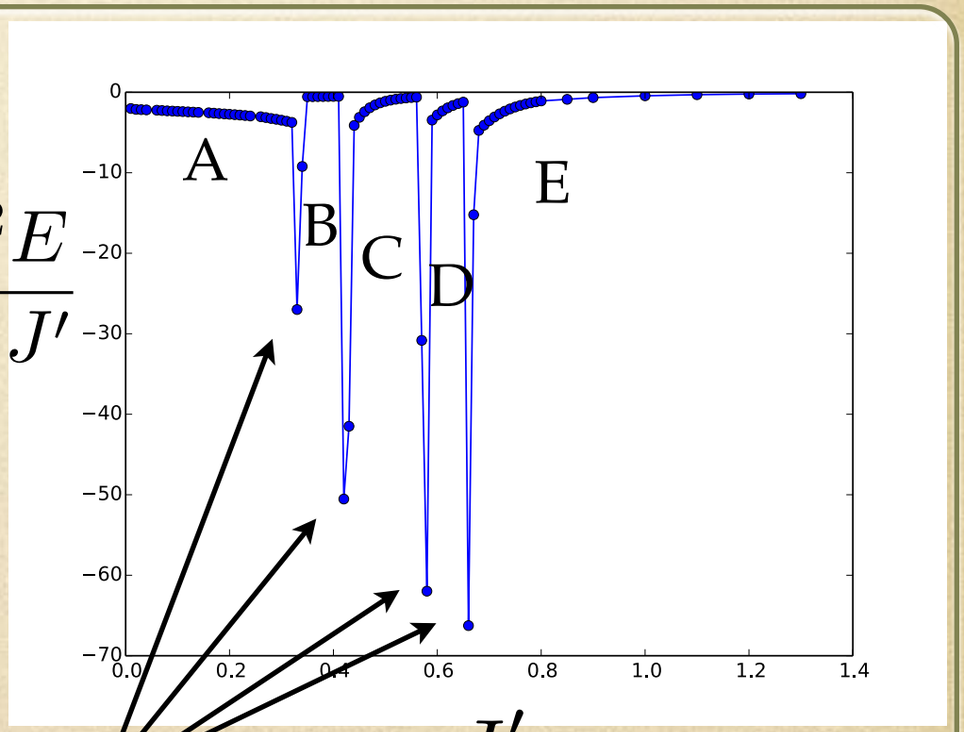

Phase transitions

Aside: What's a Spin Liquid

- Insulator
- Topological: degeneracy that depends on manifold
- Anyonic Excitations

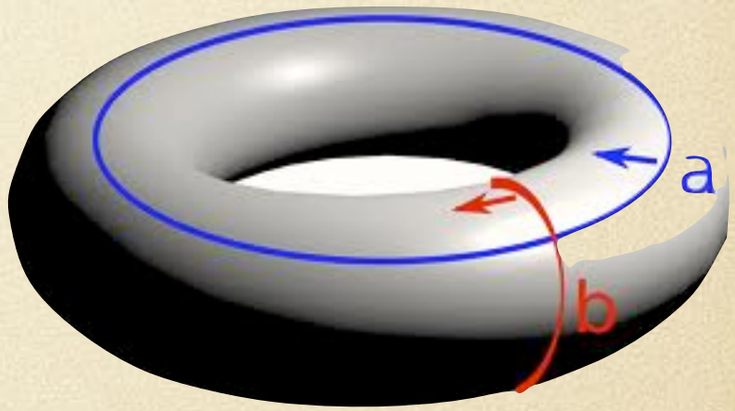

Guess a Wave-function

> The true Ψ has the property that $E = \Psi^T H \Psi$ is minimal over all Ψ.
>
> We need a compact representation of a 2^n state vector. (Important open question)
>
> In some cases, this needs to be an antisymmetric function on the electron positions

Venerable history: BCS Superconductivity

Quantum Hall Effect

Model Wave-functions

People such as Bardeen and Laughlin guessed these wave-functions without a computer by sheer genius (and were rewarded nobel prizes for it). We want to replace nobel prize winner with computers

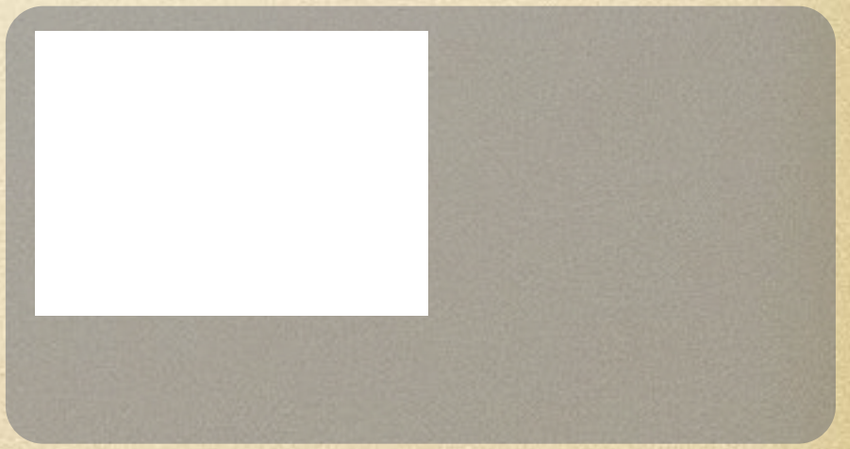

HILBERT SPACE IS A BIG PLACE

Protoypical Wave-functions:

$$\Psi(\uparrow,\downarrow,\uparrow,\uparrow) = 0.3 \quad V_{\uparrow,\downarrow,\uparrow,\uparrow} \qquad 2^n$$

Huse-Elser/CPS/EPS

$$v^a_{12} v^b_{13} v^c_{23} v^d_{14} v^e_{24} \qquad 4n^2$$

$$v^a_{\uparrow\downarrow} v^b_{\uparrow\uparrow} v^c_{\downarrow\uparrow} v^d_{\uparrow\uparrow} v^e_{\downarrow\uparrow}$$

MPS

$$M^a_1 M^b_2 M^c_3 M^d_4 \qquad nD^2$$

$$M^a_\uparrow M^b_\downarrow M^c_\uparrow M^d_\uparrow$$

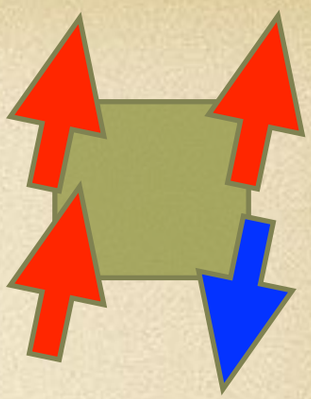

v : scalar

M : matrix

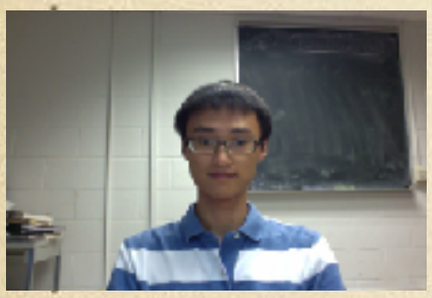

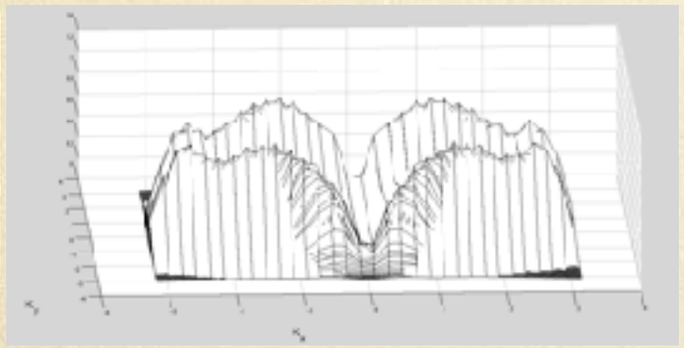

Unfortunately, these wave-functions are not compact enough. For the most interesting physical systems you need an exponential number of parameters (turn up the matrix size, the number of determinants, etc.) to ge converged energies.

$$|\Psi\rangle = \exp[-J(R)] \sum_k \alpha_k \det M_{\uparrow,k} \det M_{\downarrow,k}$$

$$J(R) = \sum_{ij} j(r_i - r_j)$$

$$M_{ij} = \phi_i(r_j)$$

3d function $\quad R \equiv \{r_1, r_2, \ldots, r_n\}$

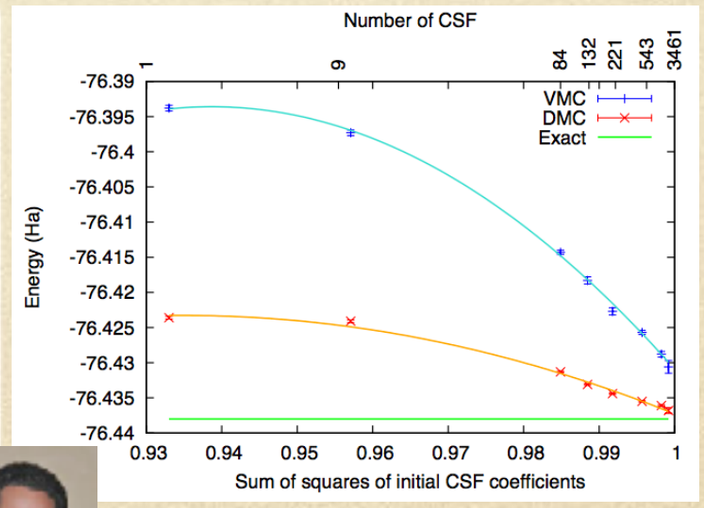

Beyond compactness, even for a small number of parameters (1000) we have a hard optimization problem:

> Given a parameterization for $\phi_i(r_j)$
> or v_{ij}^A how do you find (even local) minima.
>
> You only have *stochastic* access to the energy
> and each energy evaluation is slow!
>
> Essentially an online learning problem.

(Important open question)

Currently: Stochastic Gradient Descent $\vec{\alpha} \to \vec{\alpha} + \delta \nabla_{\vec{\alpha}} E$

Stochastic Reconfiguration $\alpha_i \to \alpha_i + \delta \mathrm{sign}(\partial E / \partial \alpha)$

'Time Evolution' $\Psi = P(1 - \tau H)\Psi$

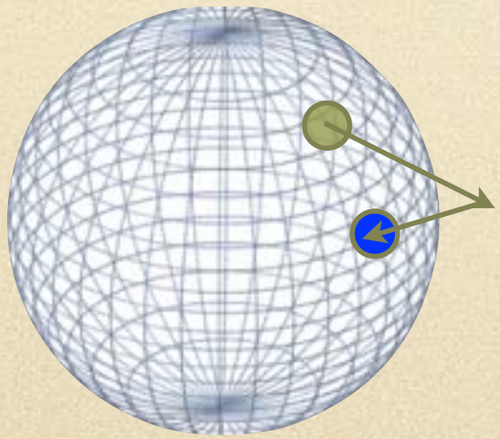

Striped Spin Liquid Crystal

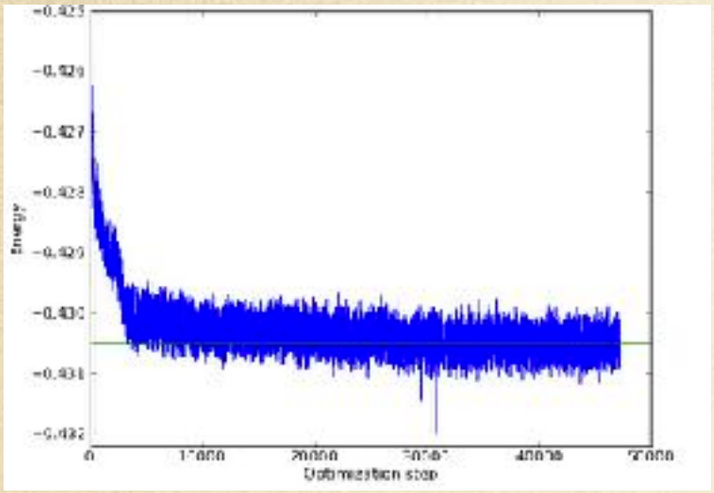

Topological + Nematic!

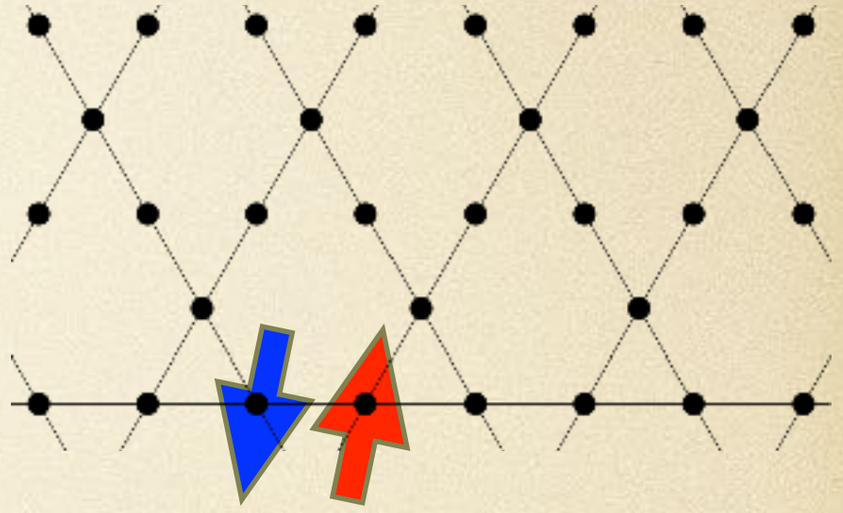

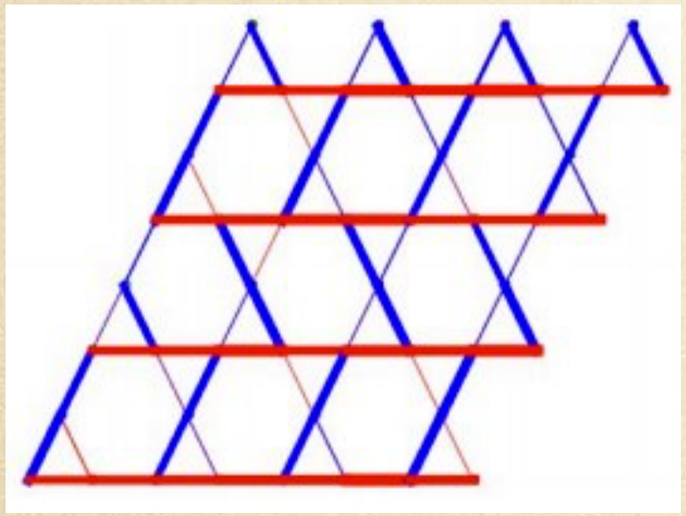

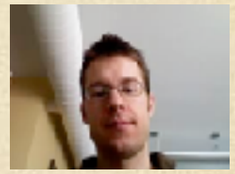

- Herbertsmithite

- Volborthite

- Zn-Paratacamite

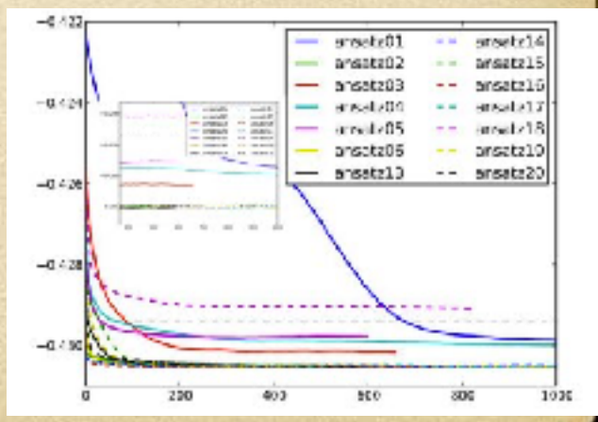

An aside on DMRG

The most important (physics) algorithms still running on one node.

Every other wavefunction gets optimized by stochastic variational means.

DMRG is optimized by alternating least squares.

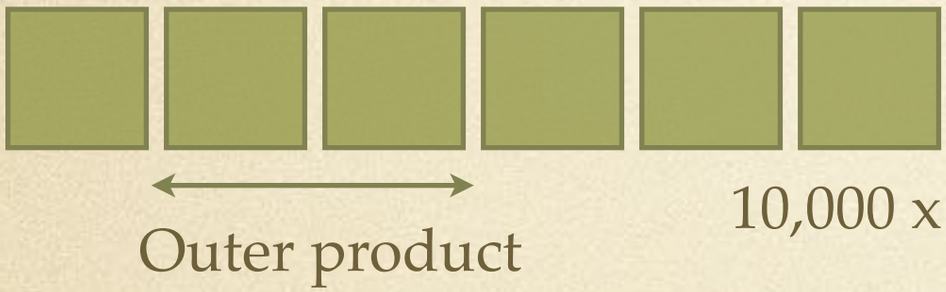

Diffusion Monte Carlo

(Stochastic Power Series)

$$\lim_{M\to\infty}(1-\tau H)^M\Psi = \Psi_0$$

$$\lim_{M\to\infty} G^M\Psi = \Psi_0$$

$$\sum_{ijklmn} G_{ij}G_{jk}G_{kl}\ldots.G_{mn}\Psi_n = (\Psi_0)_i$$

$$P(i\to j) = \frac{|G_{ij}|}{\sum_j |G_{ij}|} \qquad w(i\to j) = \sum_j |G_{ij}|\text{sign}(G_{ij})$$

Optimization: Walk w/ $(1-\tau H)$ on manifold of parameterized states.

DMC: Walk w/ $(1-\tau H)$

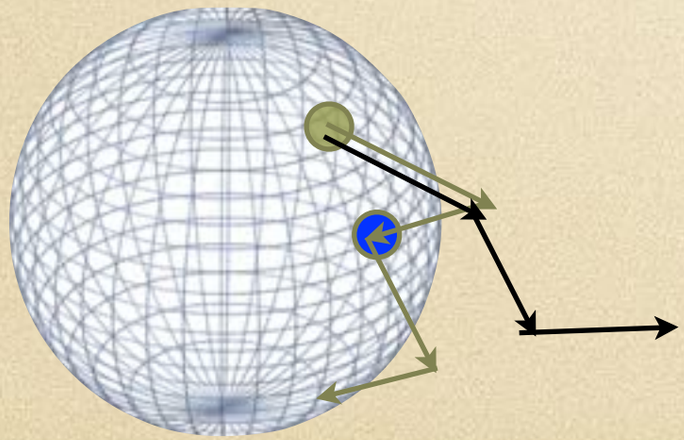

Quantum

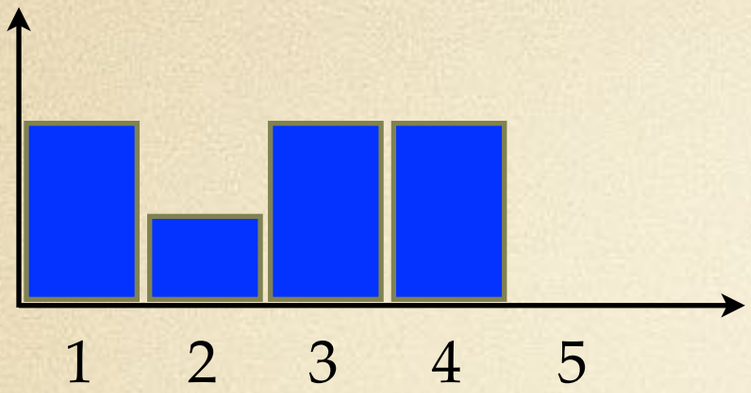

The Exciton Bose Liquid

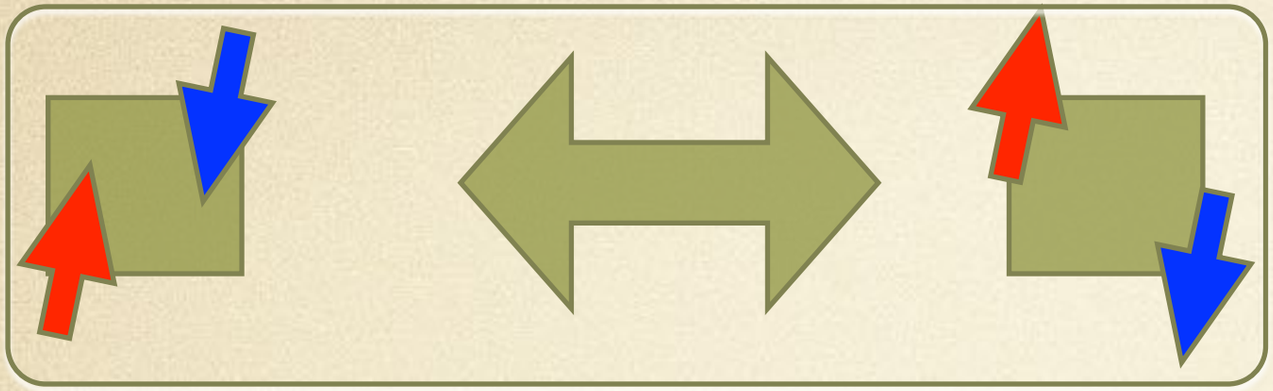

High Tc superconductors have a bad metal.

Can we find a bad metal in a simple Hamiltonian.

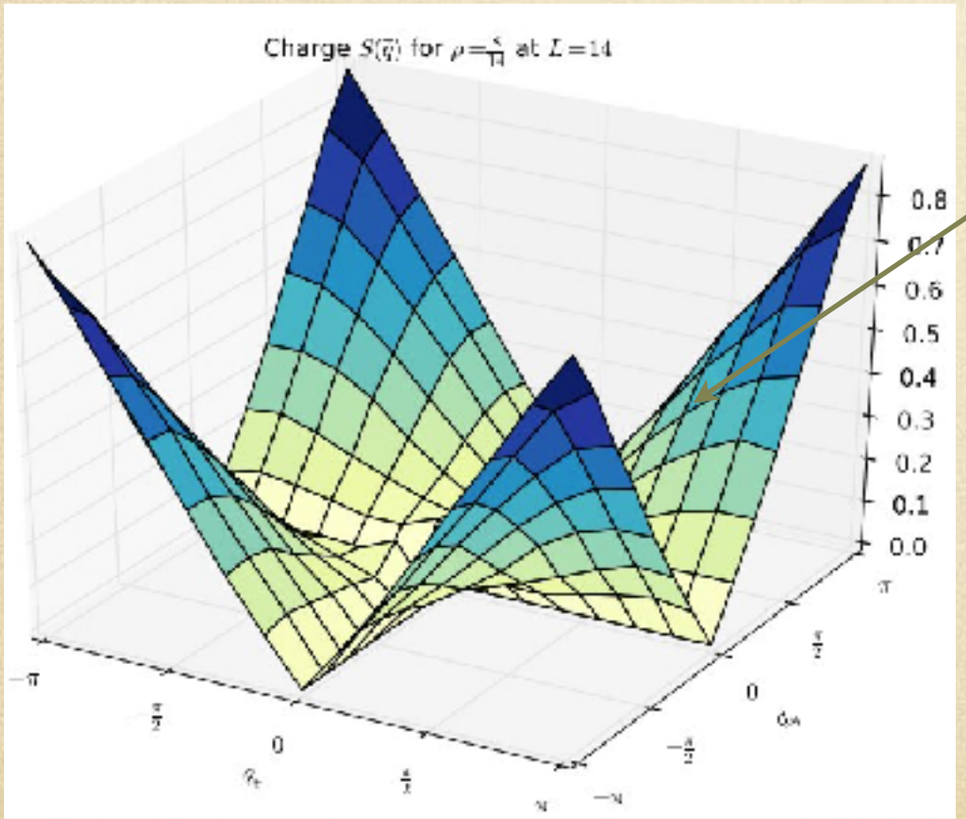

Signs of bad metal - X in the structure factor.

Parallelization

Unlike DMRG, diffusion Monte Carlo parallelizes extremely well.

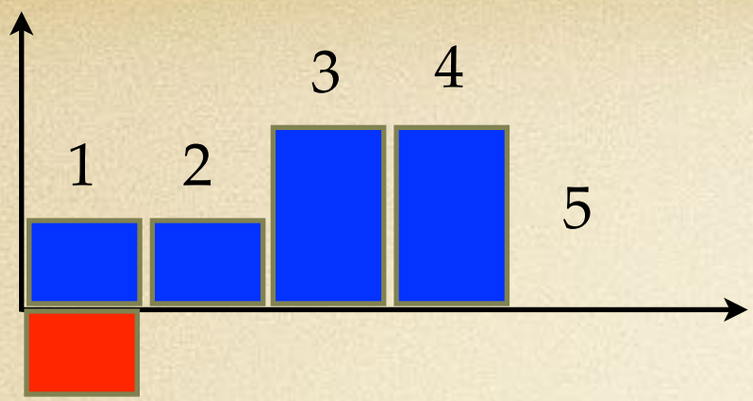
1
2
3
4
5

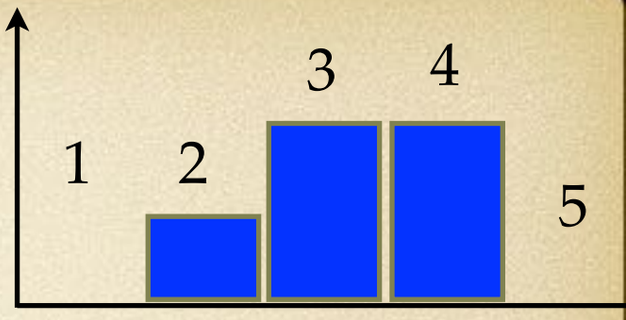
1
2
3
4
5

A Sign Problem

The "only" problem in physics

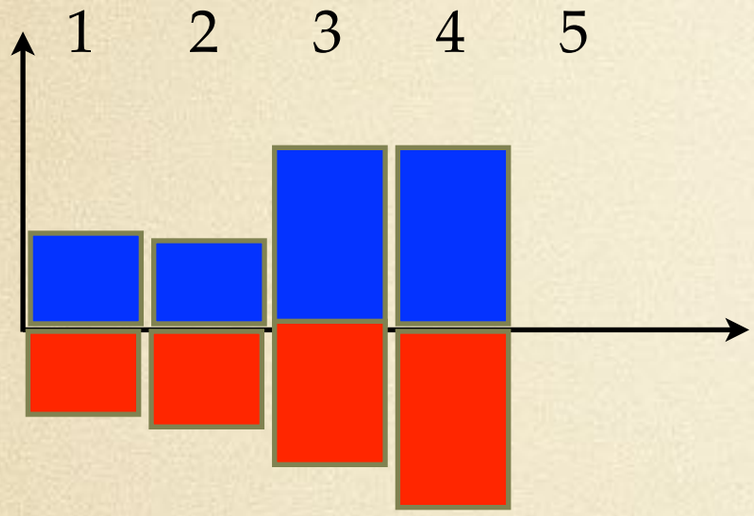

Your errors on this are good

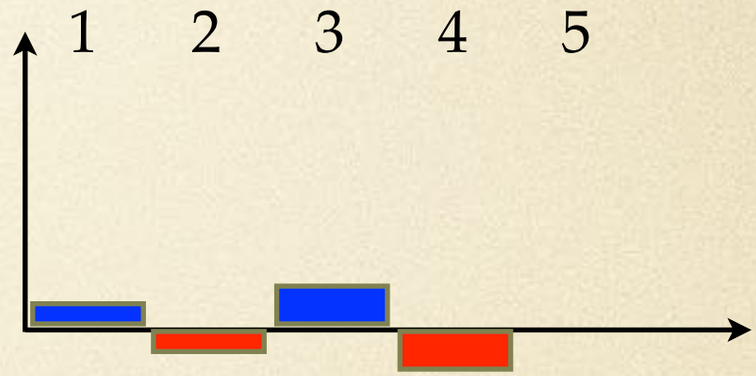

but your errors on this are horrible.

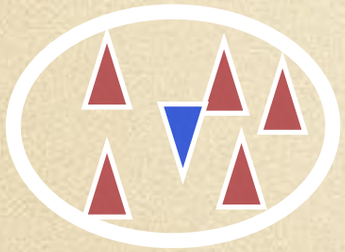

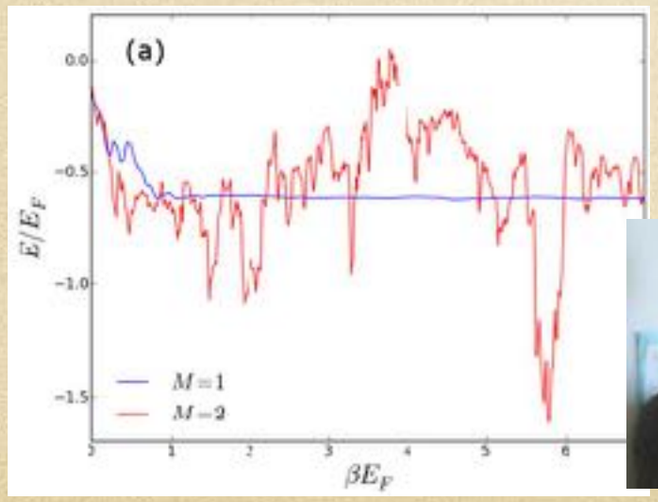

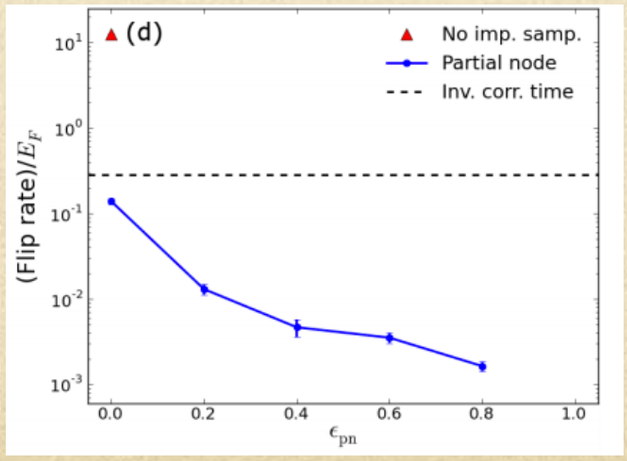

Exponential Relative Variance

Annihilation

Without

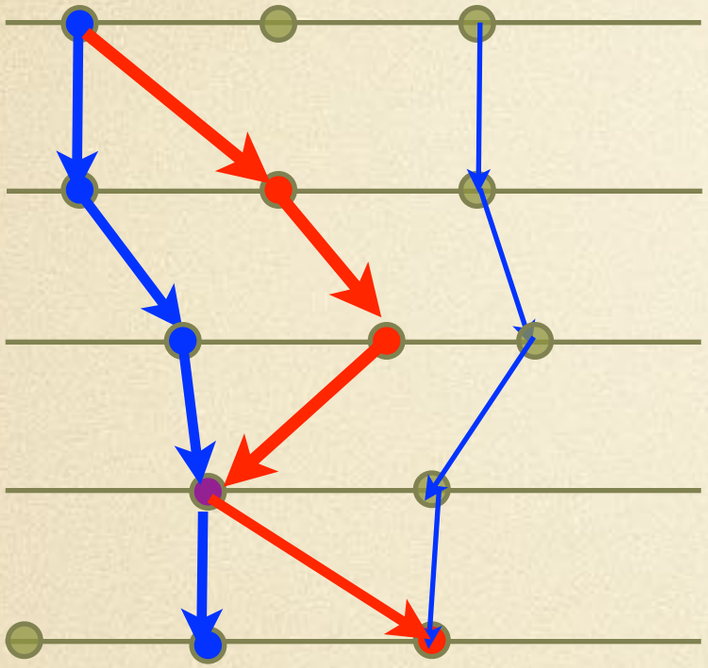

- Annihilation helps because paths of different signs cancel.
- Annihilation fails because you can't keep enough walkers to get cancellation of all paths.

With

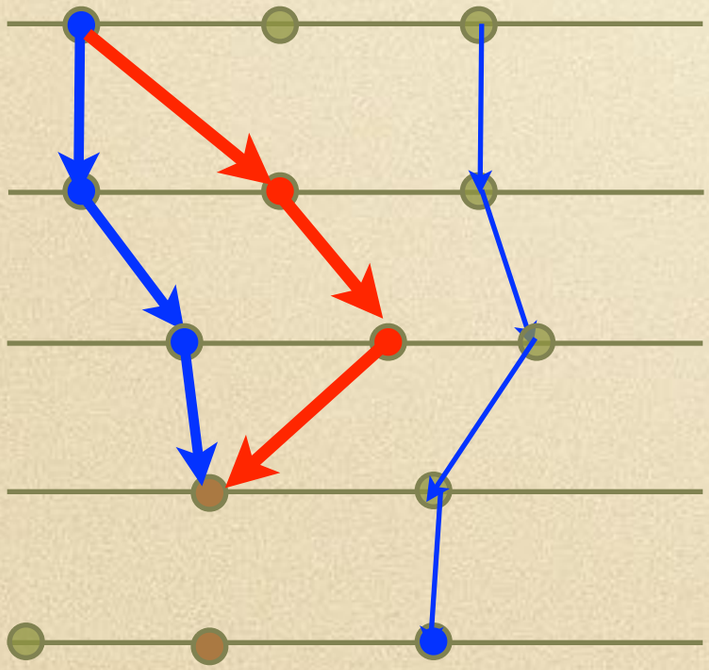

With but too few walkers

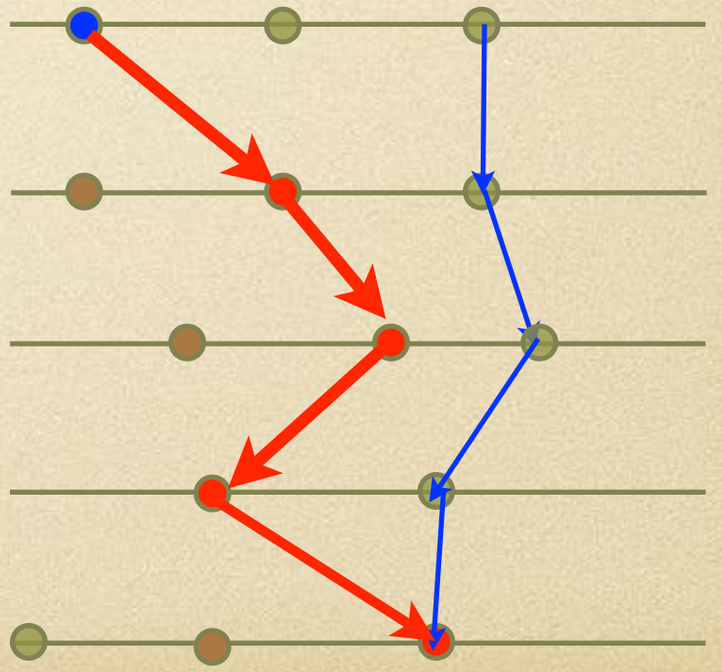

Variational Wave Function

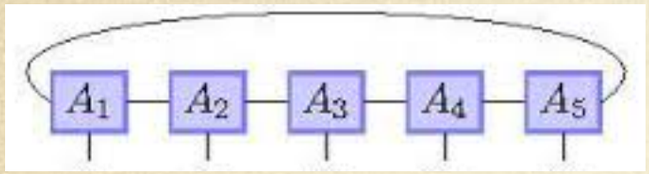

Product states

$= 0.01$ 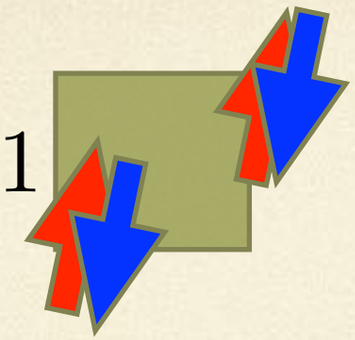 $+ 0.03$ 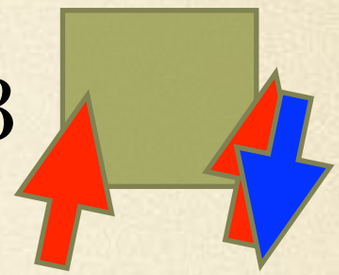 $+ \ldots$

Sample

$|\langle MPS_1|C\rangle|^2$

weight

$$\frac{1}{\langle MPS_1|C\rangle}$$

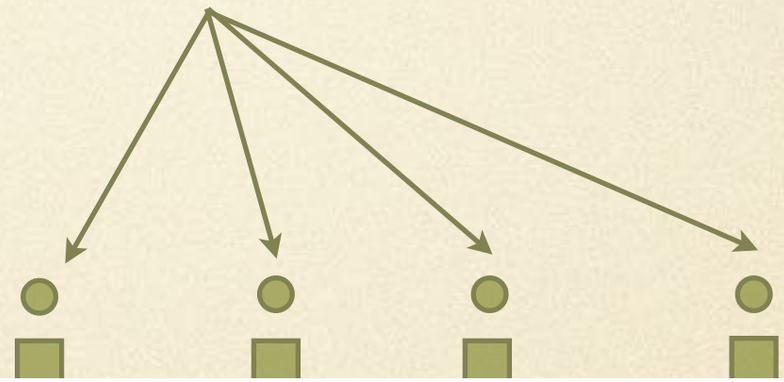

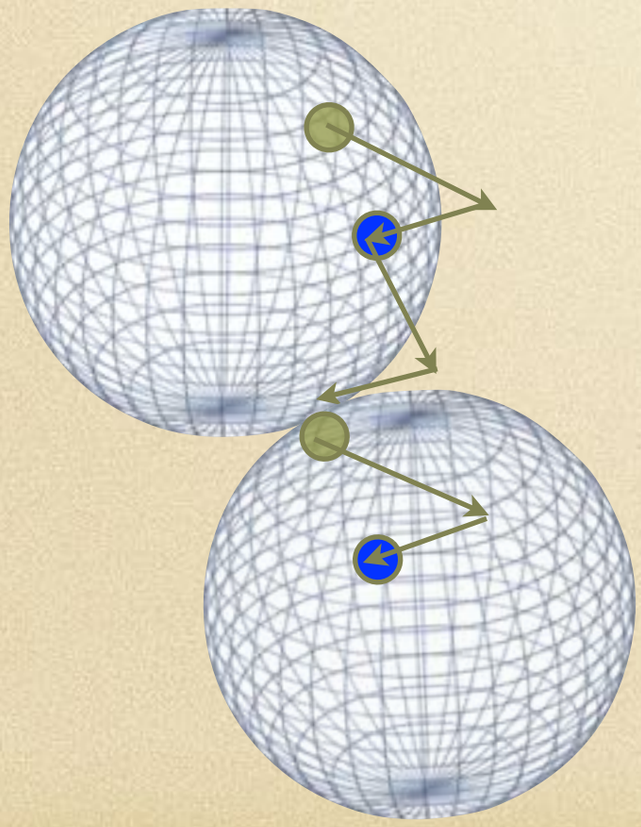

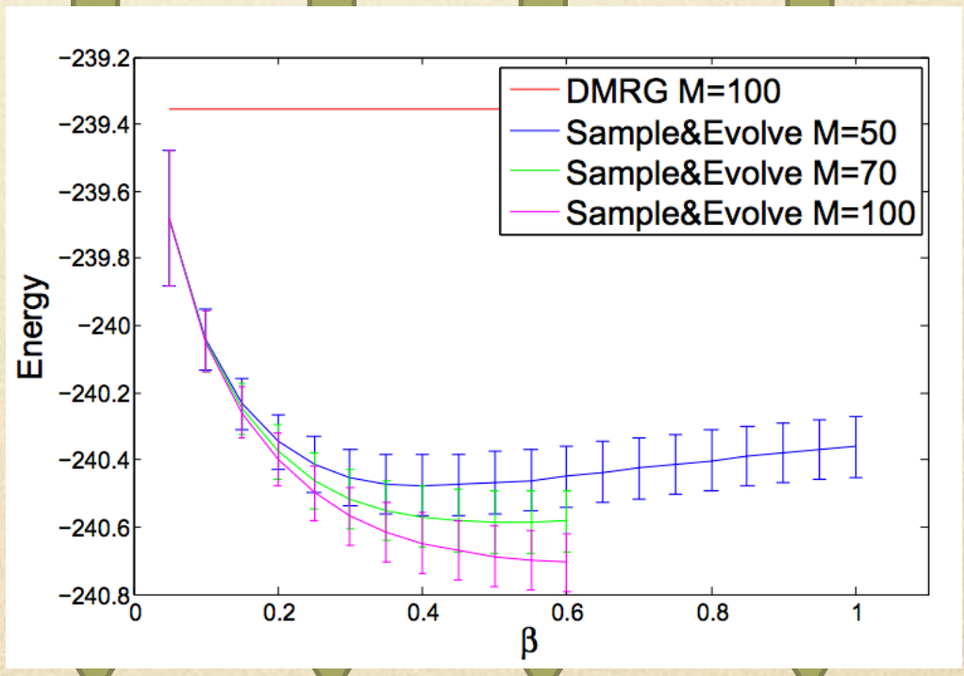

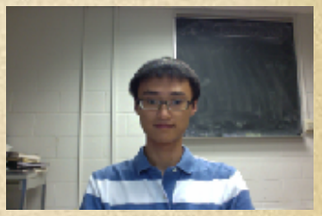

Quantum Computing

Will quantum computers solve our problems?

Maybe...

Modified Church-Turing Thesis

All computers (physical systems) are essentially equivalent to your laptop.

Quantum Mechanics broke this

Quantum simulations are quantum computers kill application, not factoring!

Quantum computers can (probably) compute g.s. in poly-time.

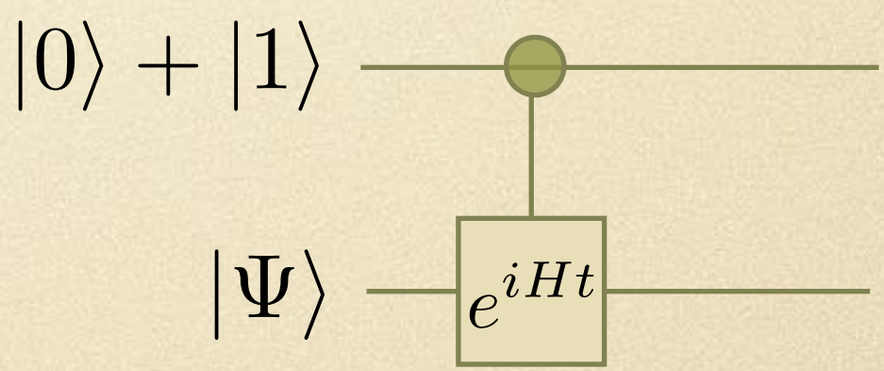

Two (minor) problems:

- No quantum computers (16 qubits)
- 100 qubit simulation -> 10^{16} gates

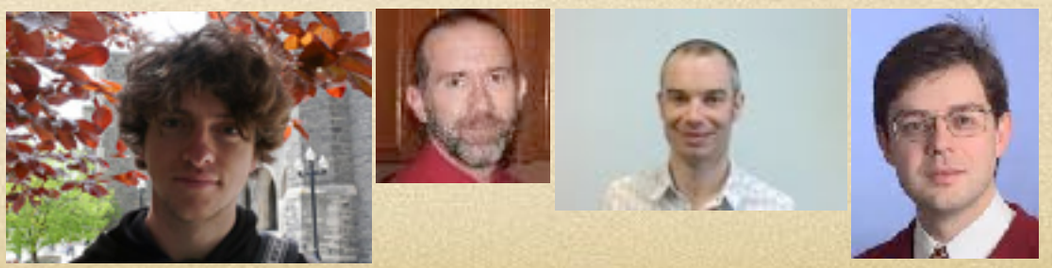

"A method is more important than a discovery, since the right method will lead to new and even more important discoveries."

-Lev Landau

Conclusions

- Quantum simulations are important:
  - and you've seen some exciting physics we've already discovered
- There are many beautiful algorithms
  - some of which we've developed.
- But we have a long way to go.
- Deep and interesting questions...
  - algorithms: FOCS/STOCS/ICML style-questions
  - More traditional CSE focus on parallelization
- (Also interested in quantum computing!)
- Interested in collaboration; come find me!

- Striped Spin Liquid

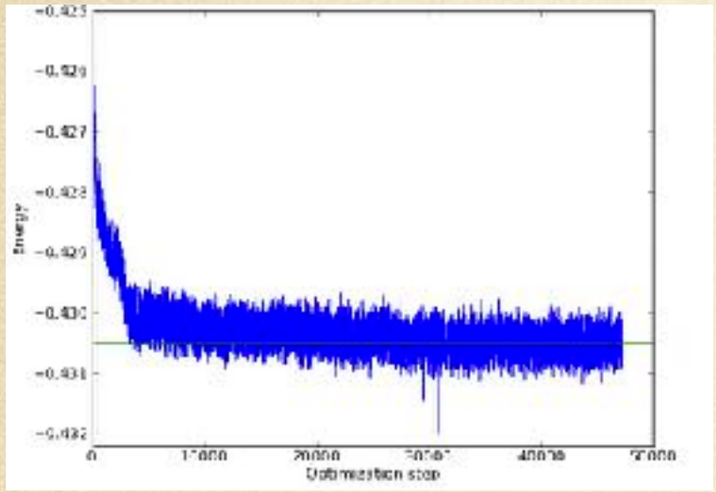

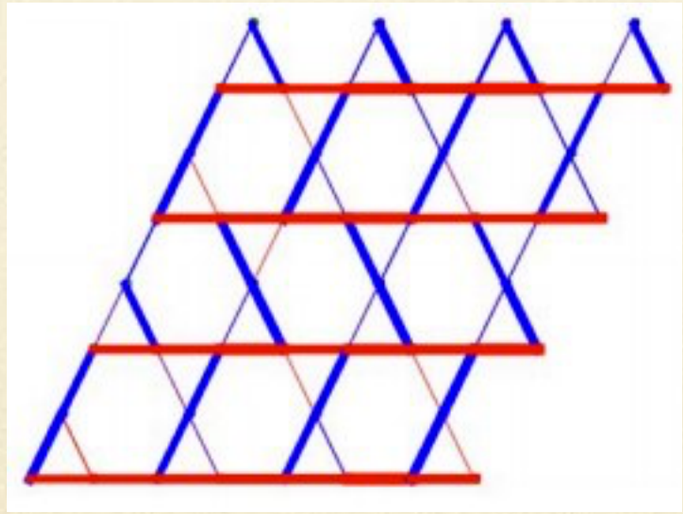

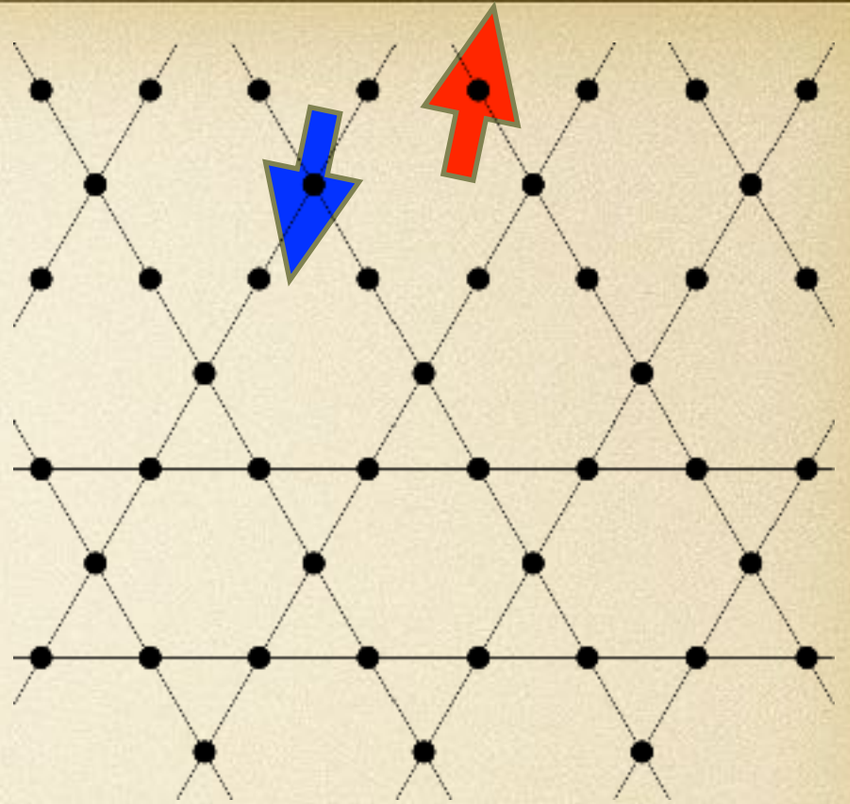

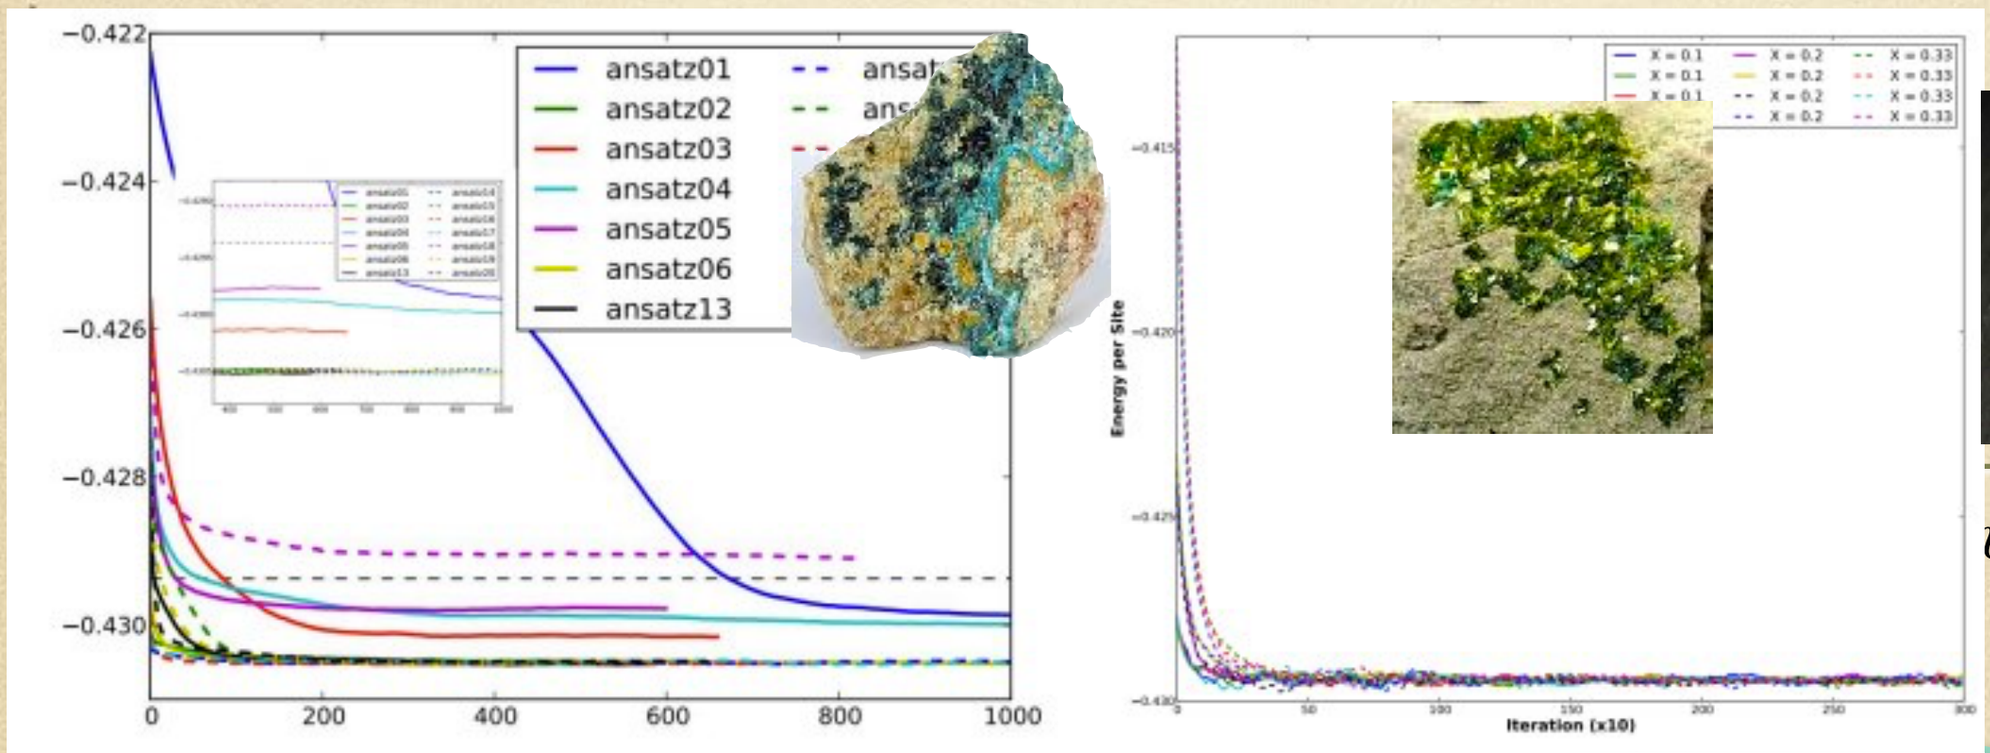

$$\upsilon_k c^\dagger_{k,\uparrow} c^\dagger_{-k,\downarrow})|0\rangle$$

$$M_{ij} = \phi(\vec{r}_{\uparrow,i} - \vec{r}_{\downarrow,j}) \equiv \phi(\vec{r}_{ij})$$

(Modified) Church-Turing Thesis

Quantum computers can compute g.s. in polytime

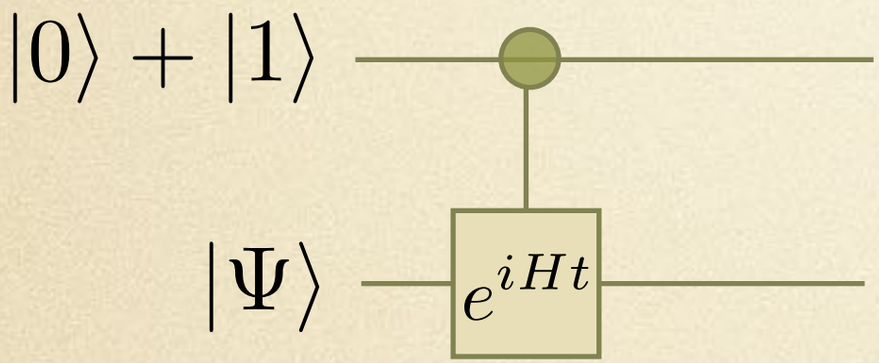

$$|0\rangle|\Psi\rangle + |1\rangle(e^{iE_0t}|\Psi_0\rangle + e^{iE_1t}|\Psi_1\rangle + e^{iE_2t}|\Psi_2\rangle)$$

- Minor Problems:
  - No quantum computers (16 qubits)
  - 100 qubit simulation -> 1016 gates

Everything is quantum mechanics

41%

- ‘Condensed Matter Physics’
- ‘Chemistry’
- ‘Materials’
- ‘Nuclear Physics’
- ‘Water’
- ‘Biology’

Accurate and efficient quantum mechanical simulations would have revolutionary changes to industrial applications, what we know about our universe and the way physics is done.

We know the rules.

Just need to figure out how to solve them.

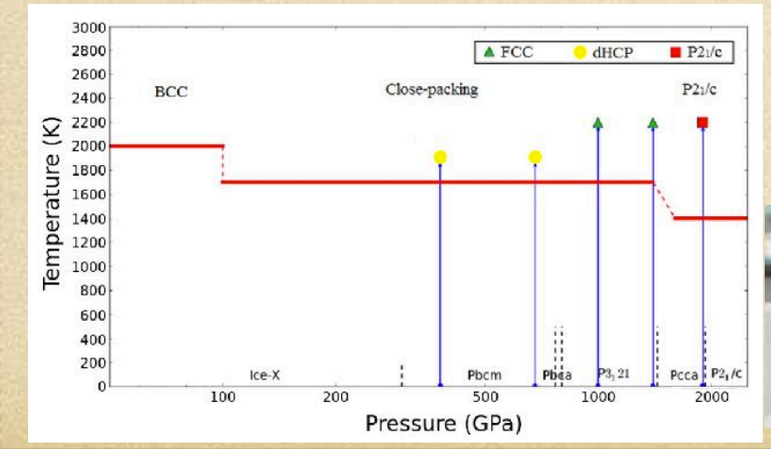

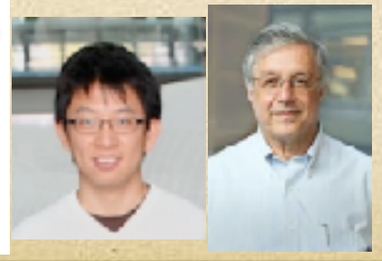

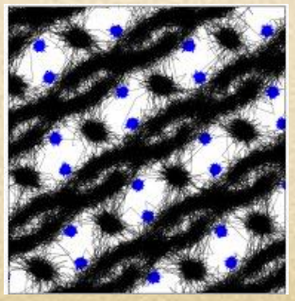

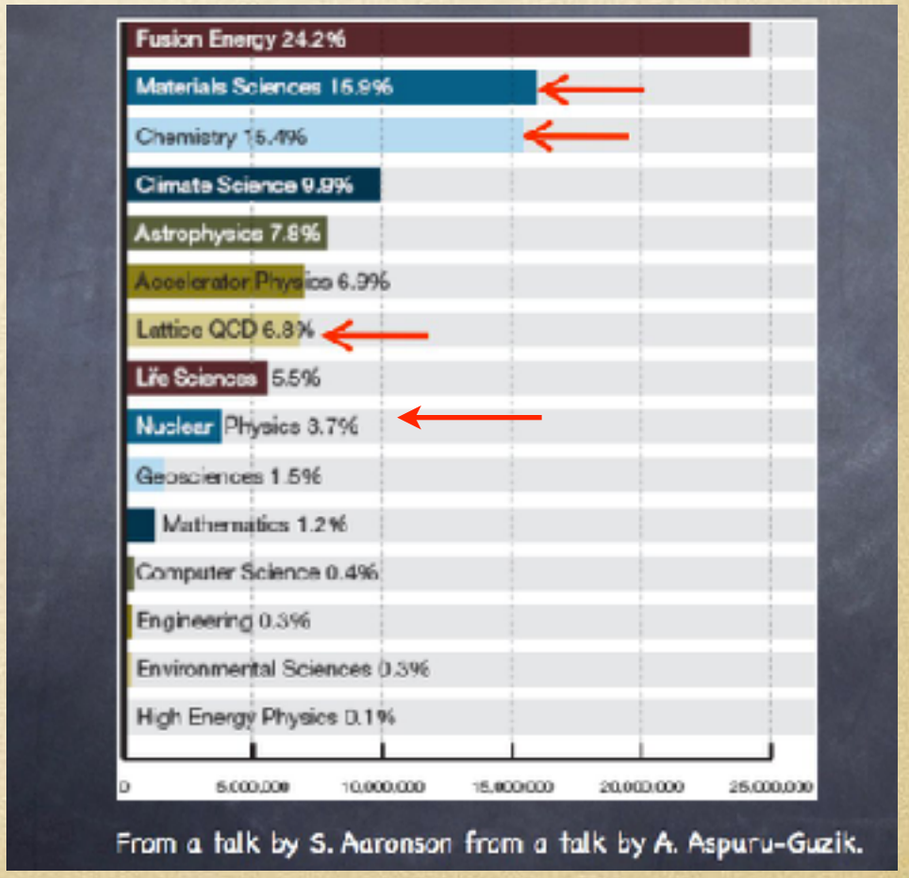

From a talk by S. Aaronson from a talk by A. Aspuru-Guzik.

Variational Monte Carlo

Variational Principle: $E_0 = \langle \Psi_0 | H | \Psi_0 \rangle \leq \langle \Psi_T | H | \Psi_T \rangle$

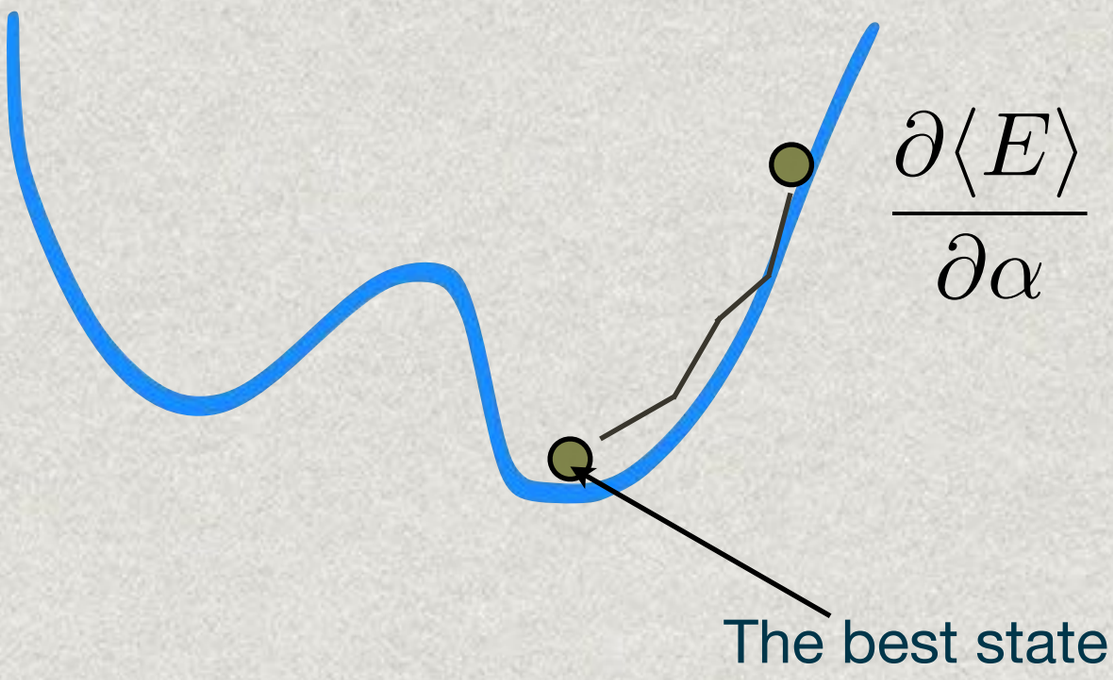

- Choose set of $\Psi[\vec{\alpha}]$
- Find the best one in set

Highly *nonlinear* optimization with an objective function $\langle E[\Psi[\vec{\alpha}]] \rangle$ and derivatives $\partial \langle E \rangle / \partial \alpha_i$ which can only be evaluated *noisily* and *slowly*.

Aside: Analogous to online learning.

Church-Turing Thesis

All computers (physical systems) are essentially equivalent to your laptop.

Quantum Mechanics broke this

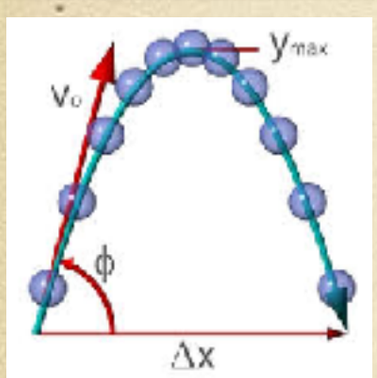

Simulating quantum mechanics is hard

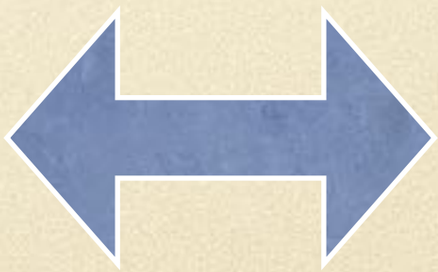

Quantum Computing

Quantum information has started to play an important role in condensed matter.

Guessing a WaveFunction

- Striped Spin Liquid

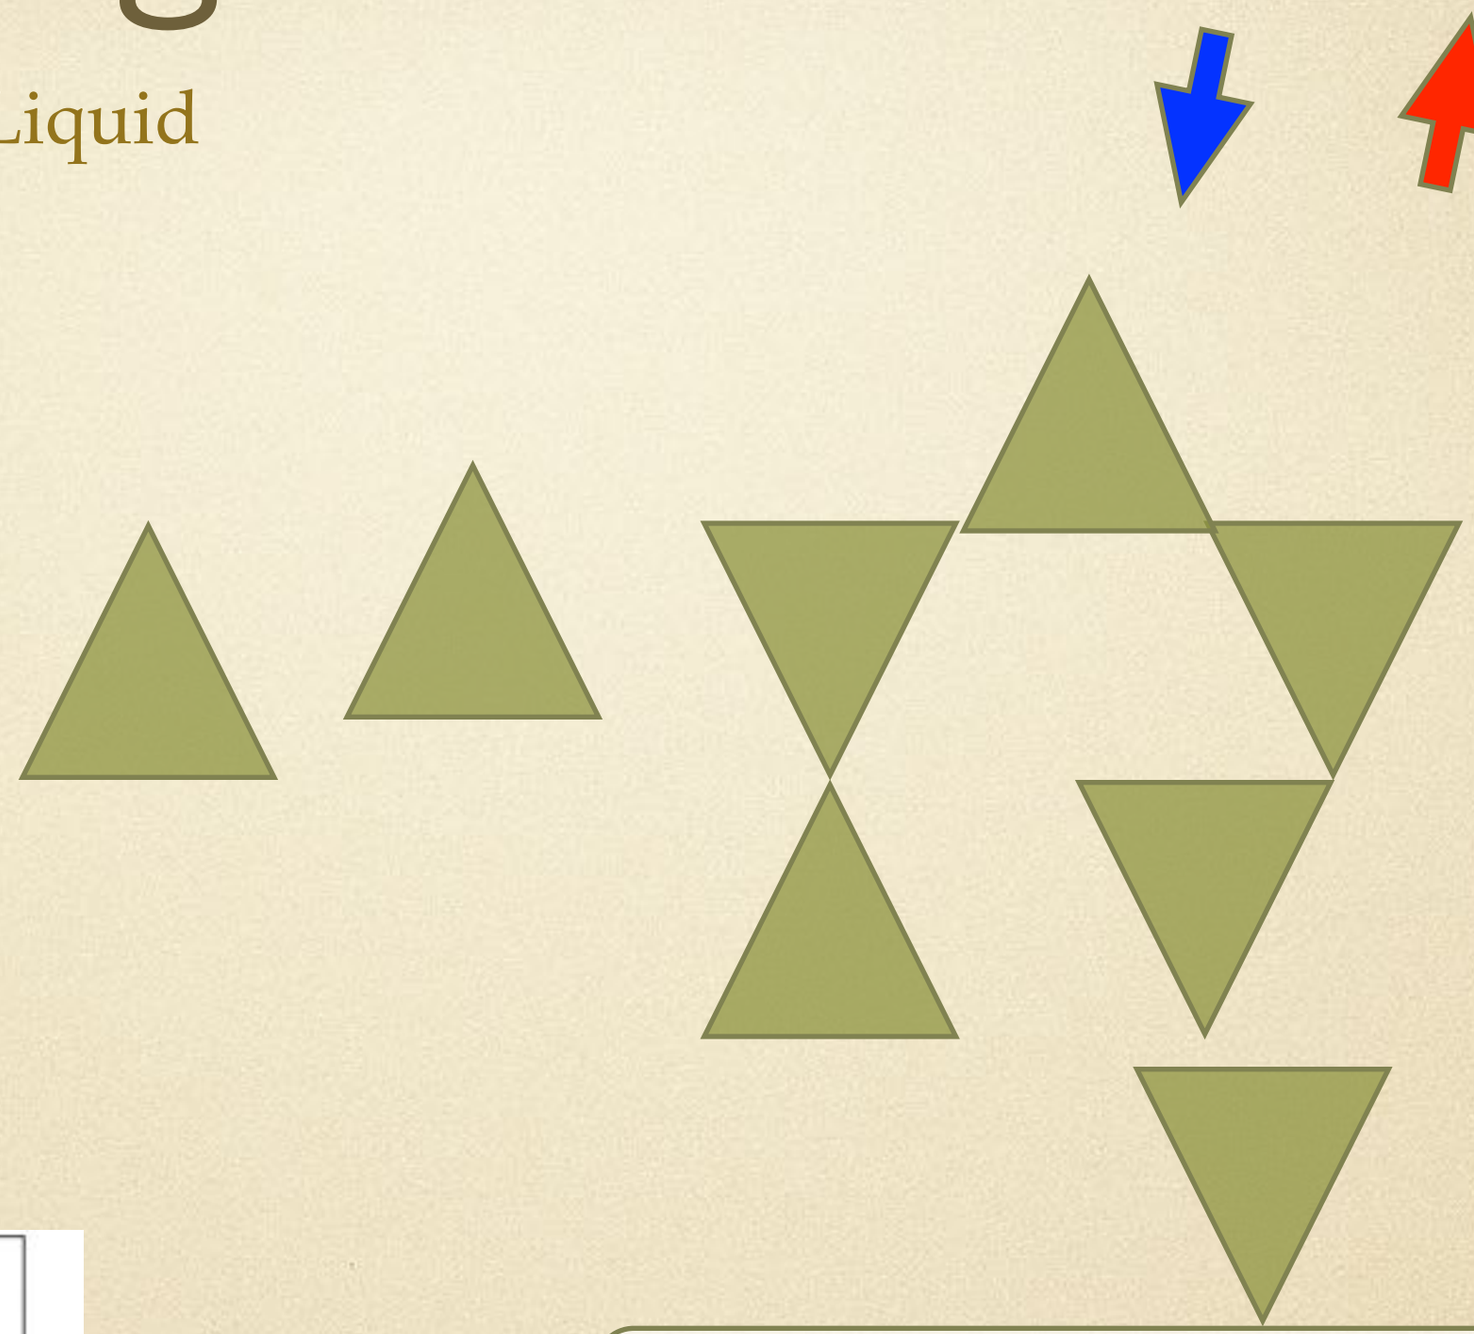

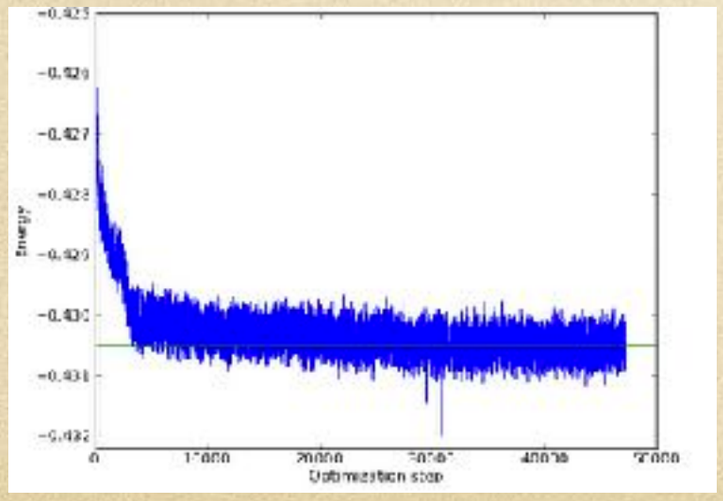

$$\Psi_{PBCS} = \underline{P} \prod_{k} (u_k + v_k c^{\dagger}_{k,\uparrow} c^{\dagger}_{-k,\downarrow})|0\rangle$$

$$\langle R|\Psi_{PBCS}\rangle = \det M$$

$$M_{ij} = \phi(\vec{r}_{\uparrow,i} - \vec{r}_{\downarrow,j}) \equiv \phi(\vec{r}_{ij})$$

Some wave-functions

$$M(i_1, i_2, .., i_n)$$

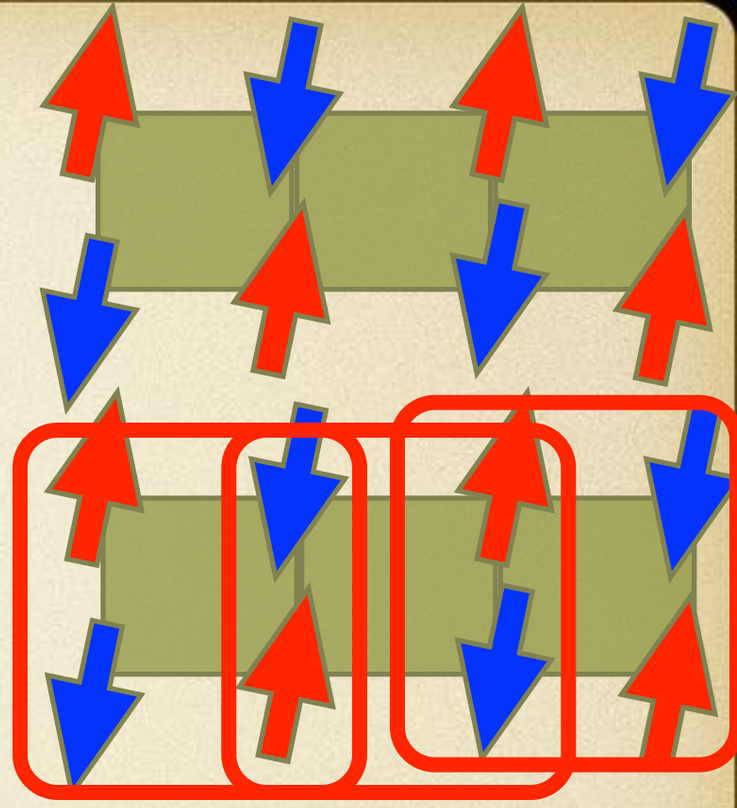

$$A(i_1, i_2, i_3, i_4)B(i_3, i_4, i_5, i_6)C(i_5, i_6, \dots)$$

$$M(i_1)M(i_2)M(i_3)M(i_4)$$

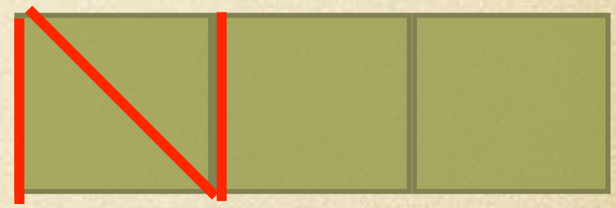

Optimization by alternating least squares

Very hard to parallelize.

Thinking about stochastic SVD

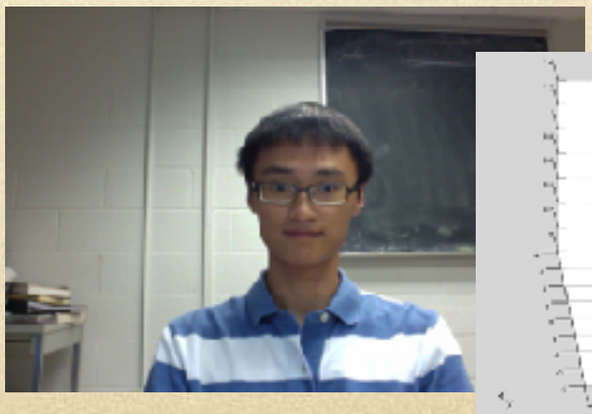

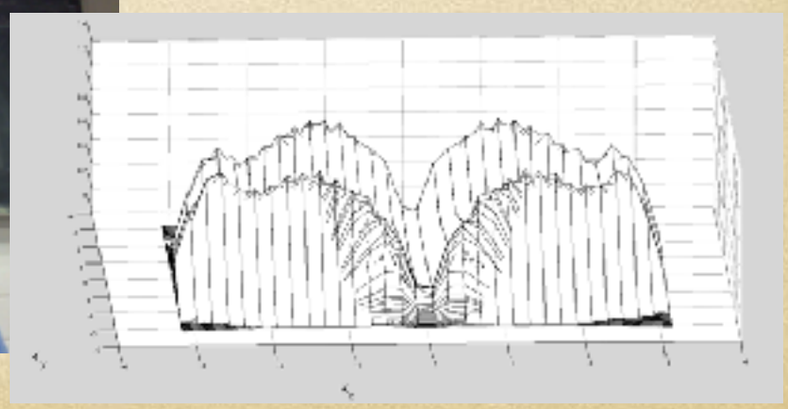

Our group works on simulating strongly correlated systems. We are currently attacking the Hubbard model as a stepping stone toward more sophisticated models.

Finite T

DQMC

METTS

PIMC

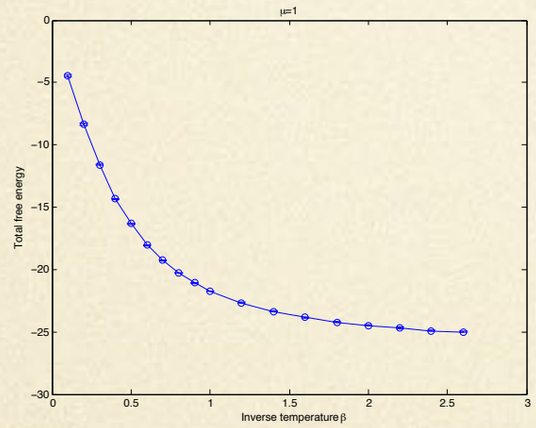

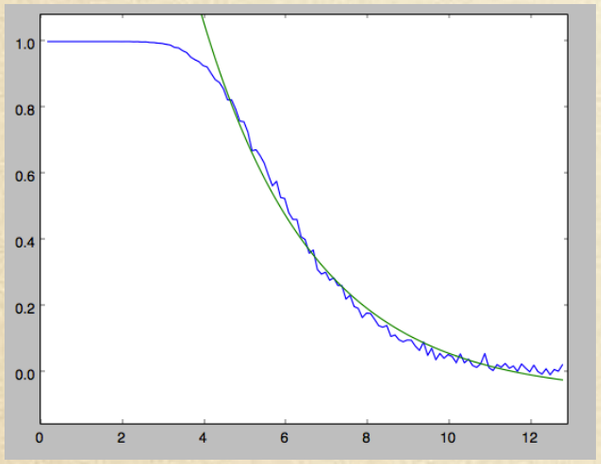

FT PMC

FT VMC

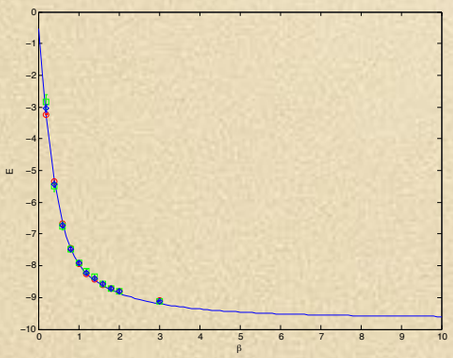

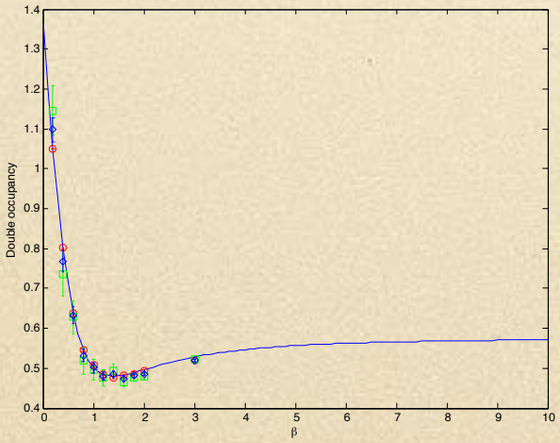

Ground State

VMC

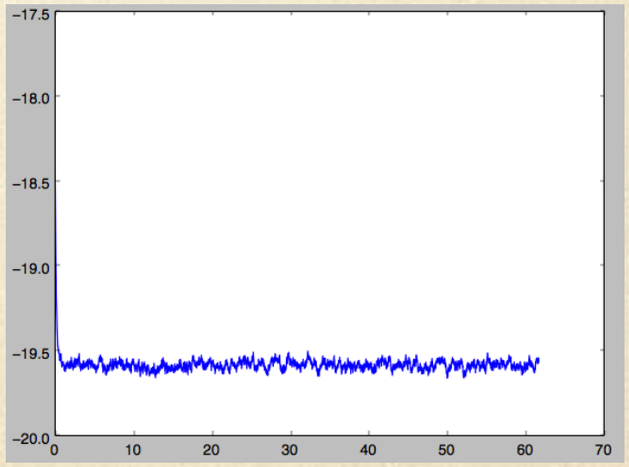

DMRG
PEPS
TN
AFQMC
DMC
FCIQMC

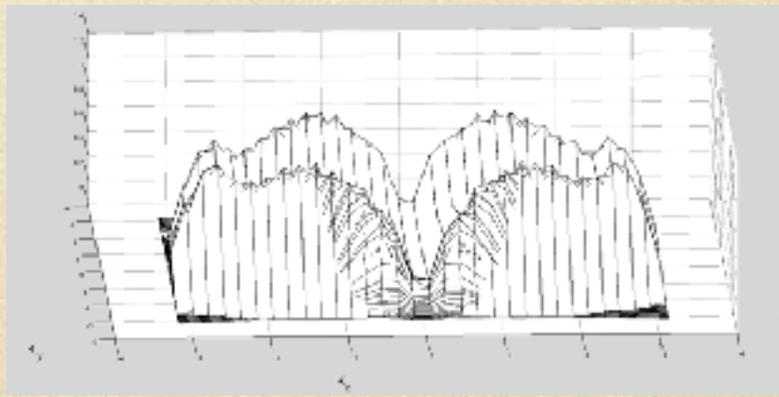

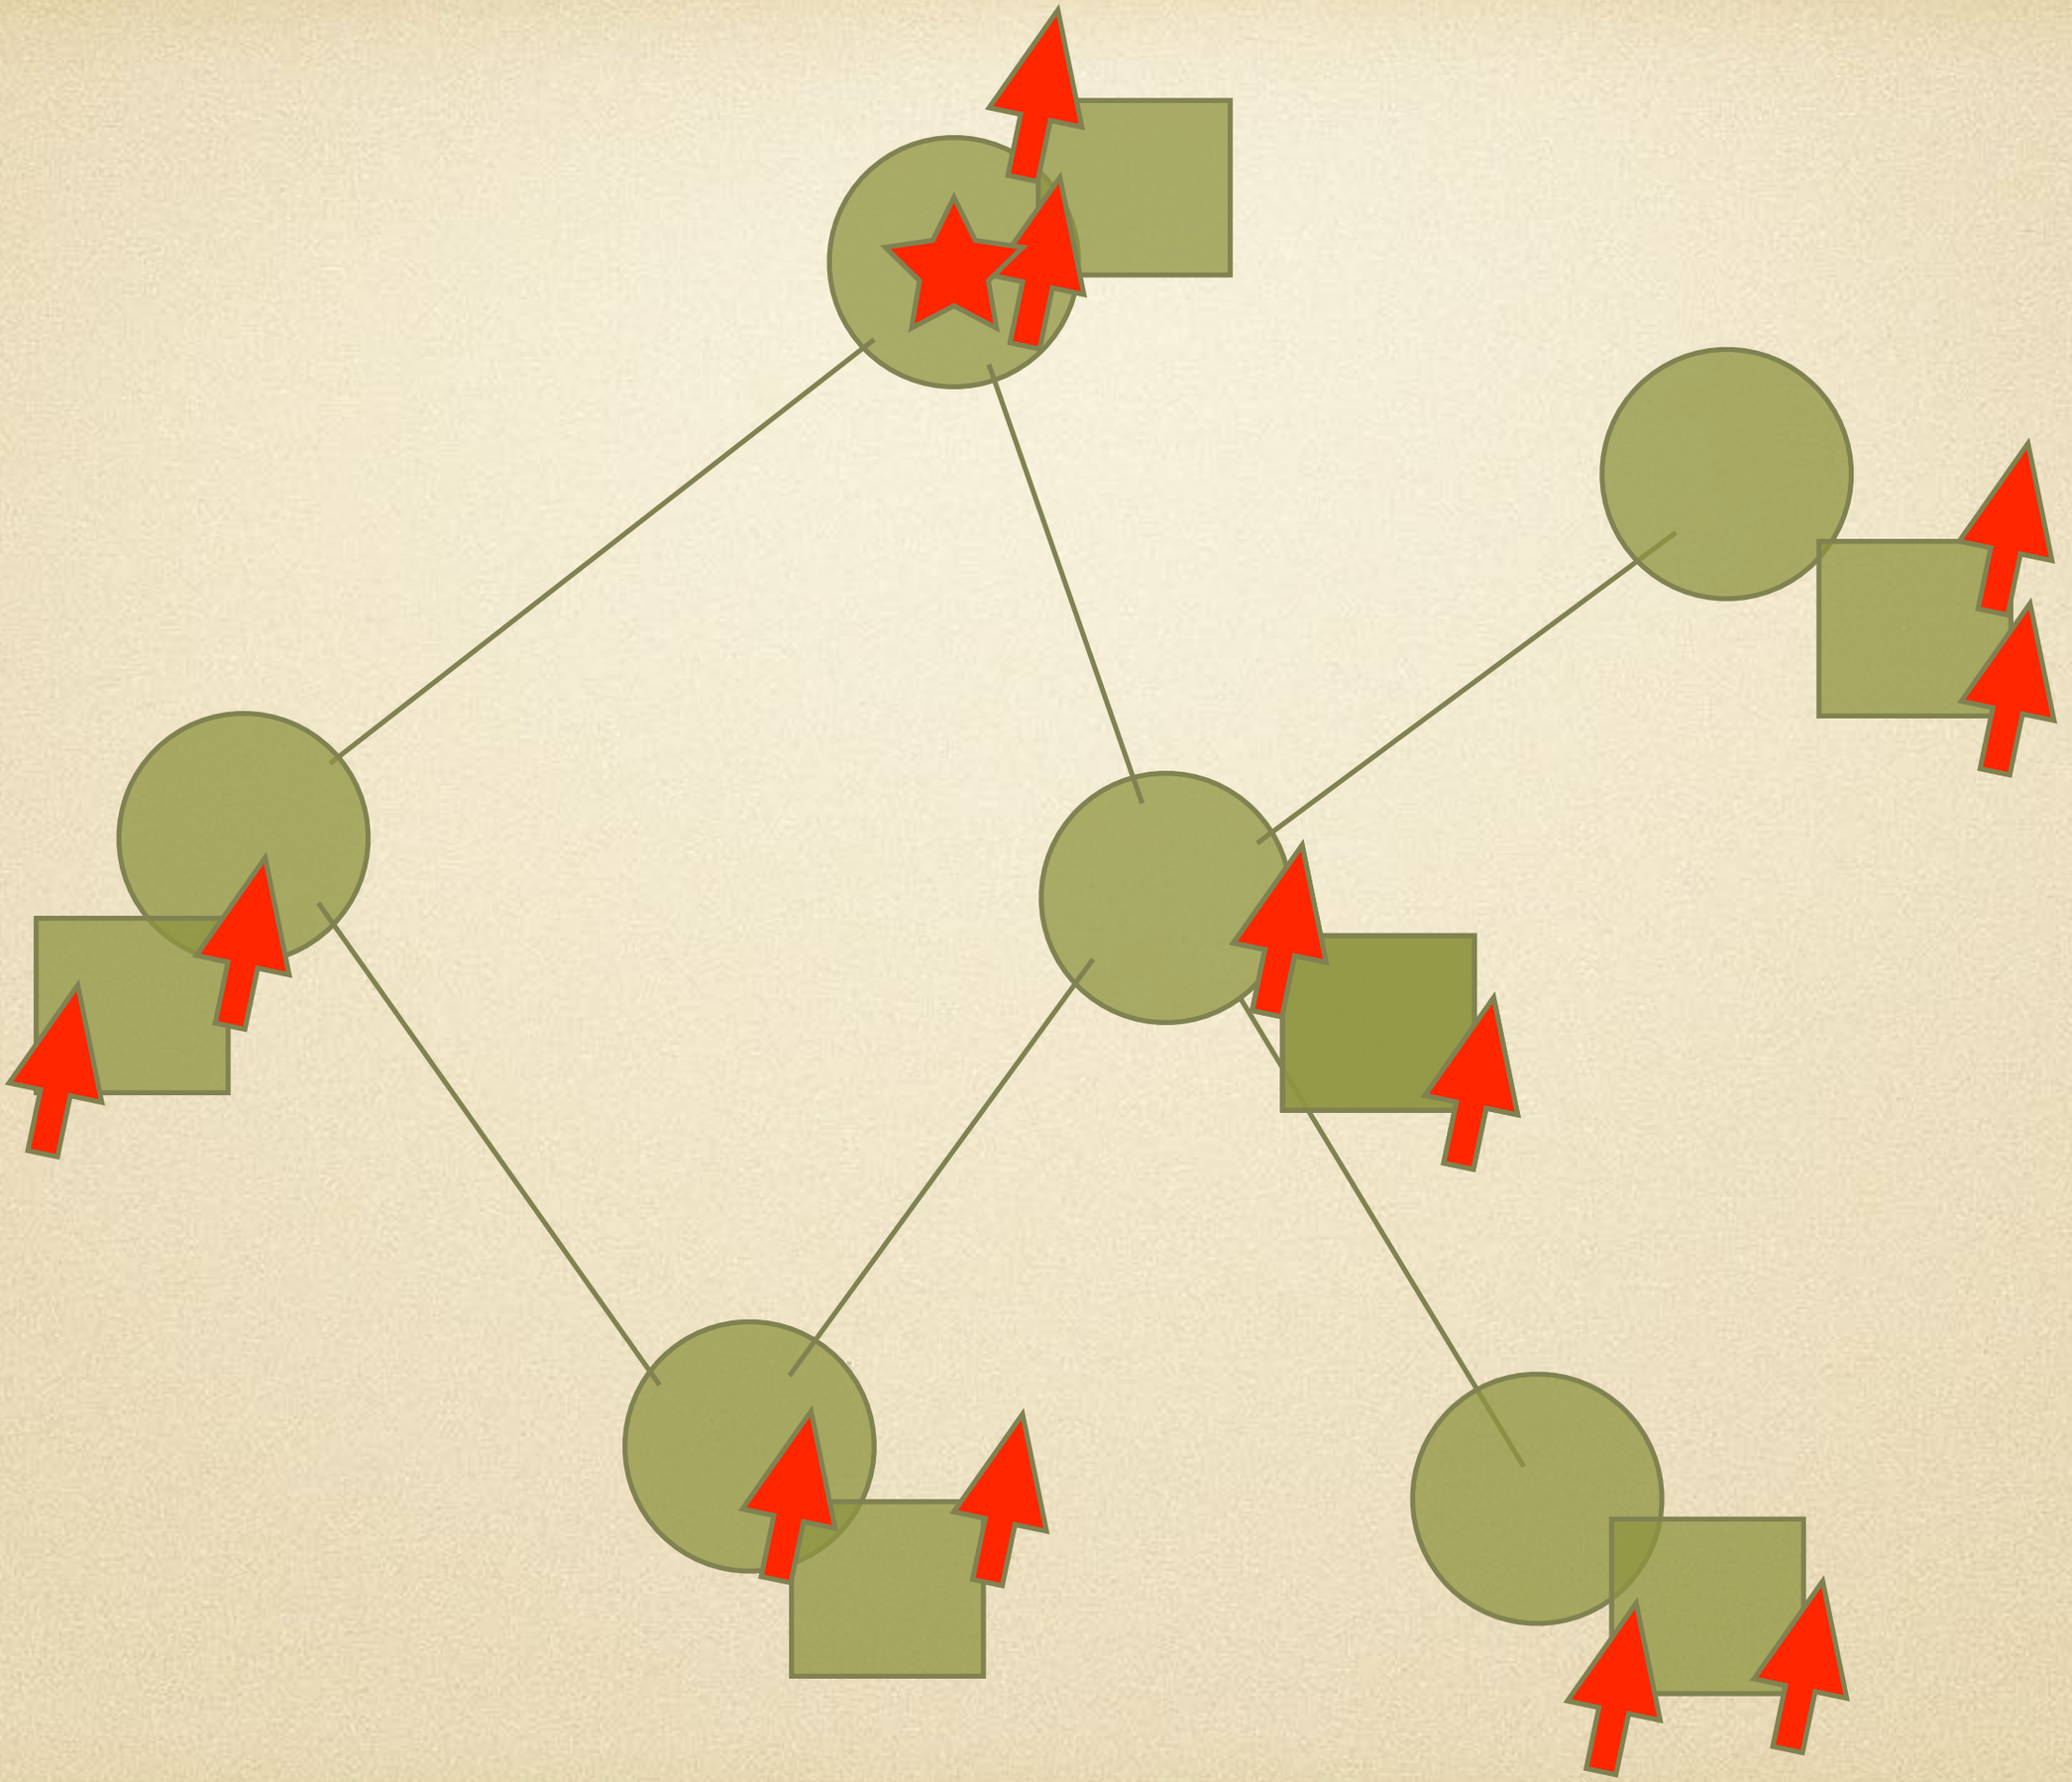

An Exponential Problem

There's an exponential wall to simulating quantum systems.

A (only slightly) biased view on the state of the art to attenuate it.

With a few new algorithms

- Partial Node FCIQMC
- Release FCIQMC
- Release + FN MPS
- Efficient Multi-MPS
- SEMPS

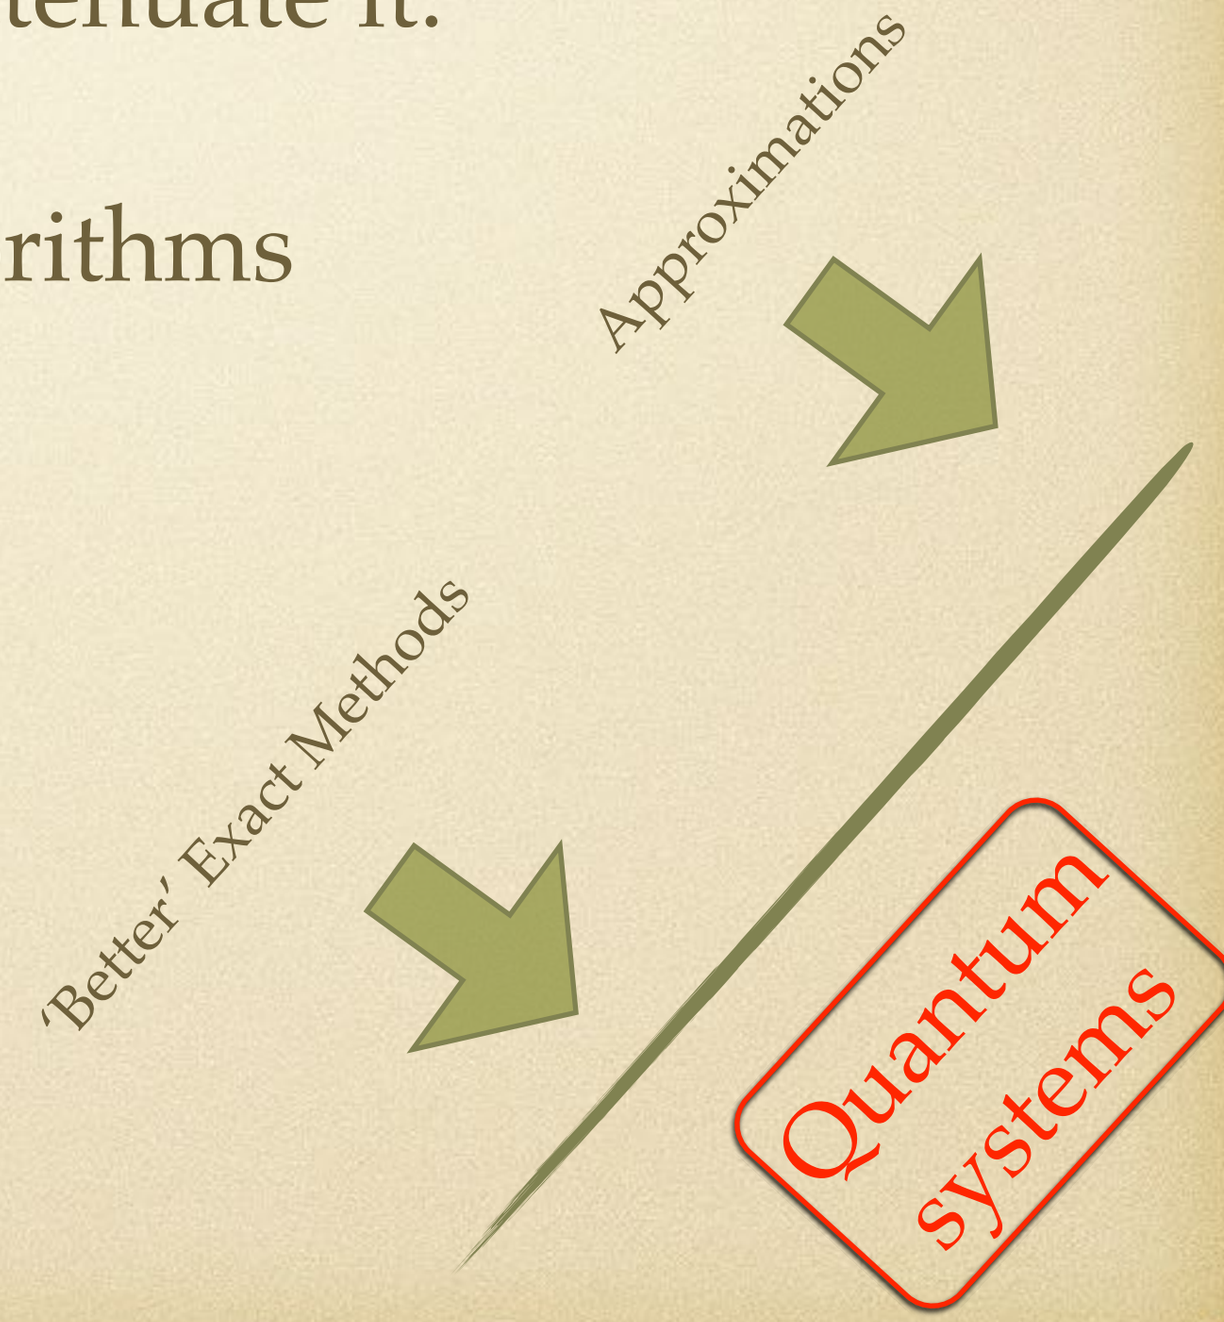

Variational

Approach I: Just write down the wave-function

Exponential number of terms

- Multislater -Jastrow++

$$|\Psi\rangle = \exp[-J(R)] \sum_k \alpha_k \det M_{\uparrow,k} \det M_{\downarrow,k}$$

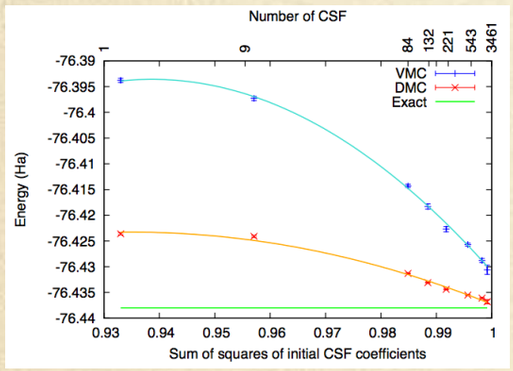

exponential number of determinants

No sign problem but "bond-dimension" problem.

- PEPS or Huse-Elser or MERA
- Multi-non-orthogonal SD + symmetry projection

- MPS

Optimize without quantum numbers and project afterwards gains non-trivial energy. On triangular lattice, ~10%

- **Multi-MPS**

exponential in width

Multi-MPS

$$\alpha|\Psi_{MPS1}\rangle + \beta|\Psi_{MPS2}\rangle + \gamma|\Psi_{MPS3}\rangle$$

How do we choose the MPS

Optimize?

Faster approach to get reasonable states...

Exact: $\{|\Psi_{MPS}\rangle, H|\Psi_{MPS}\rangle, H^2|\Psi_{MPS}\rangle, \dots\}$

Approx: $\{|\Psi_{MPS}\rangle, PH|\Psi_{MPS}\rangle, PHPH|\Psi_{MPS}\rangle, \dots\}$

Better: Let $H=h_1+h_2+h_3+h_4+h_5$

$$\{|\Psi_{MPS}\rangle, h_i|\Psi_{MPS}\rangle, h_i h_j|\Psi_{MPS}\rangle, \dots\}$$

4x8 Hubbard Model:

5 MPO's of size 6

1 MPO of size 18

For n=3, factor of 2000x faster!

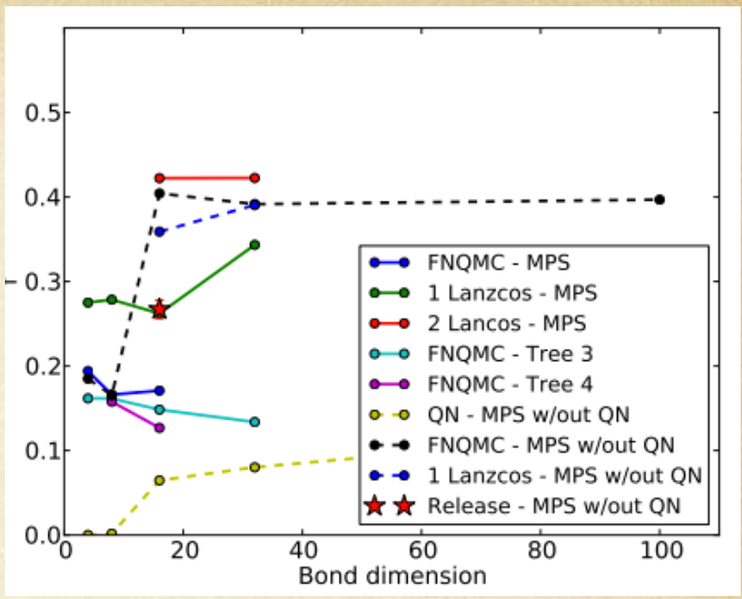

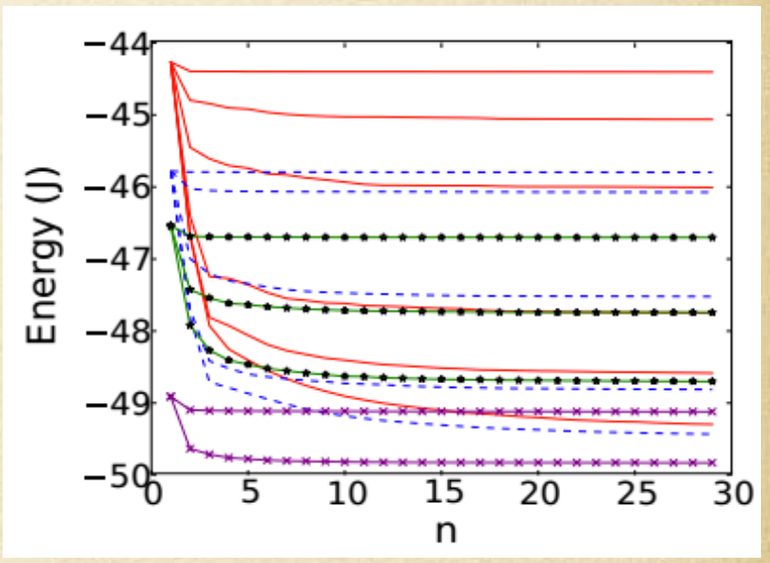

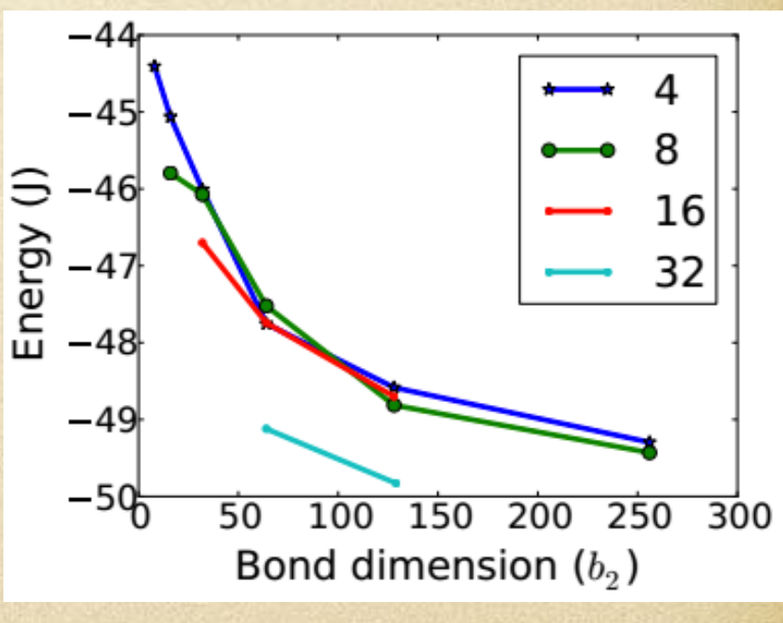

Variational ++

- Fixed Node: A (stochastic) sample of the w.f.

Two recent `improvements':

Fixed node for less-local Hamiltonians

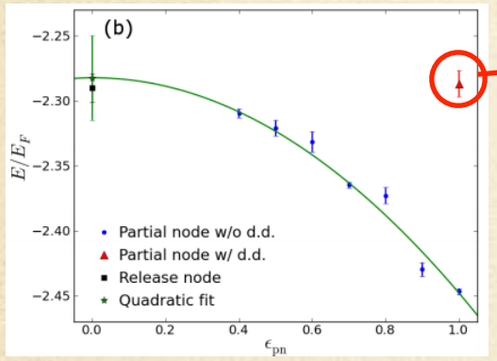

10^6 connections

Fixed node on tensor networks

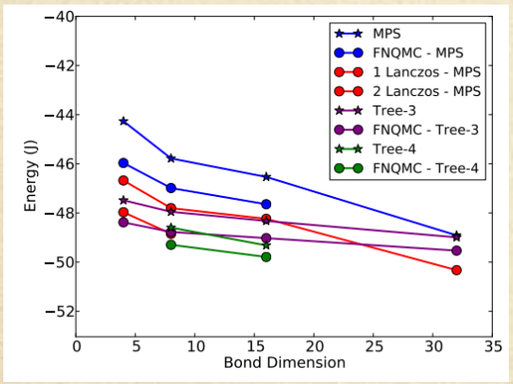

- Constrained Path:

Shiwei: Determinants

Garnet: MPS

Approach II: Sample

Sign Problem - Efficiency as $\exp[-\beta \Delta E]$

Sample

- PQMC + Annihilation

  Brings up Delta E

  - Kalos
  - + initiator: Ali Alavi

- AFQMC Free projection

Approach II: Sample

Sign Problem - Efficiency as $\exp[-\beta \Delta E]$

Sample ++

- PQMC + Annihilation
  - Brings up Delta E
    - Kalos
    - + initiator: Ali Alavi
- AFQMC Free projection
- SEMPS
- AFQMC release
- Partial Node FCIQMC

  Importance Sample + Partial Fixed-Node + Annihilation

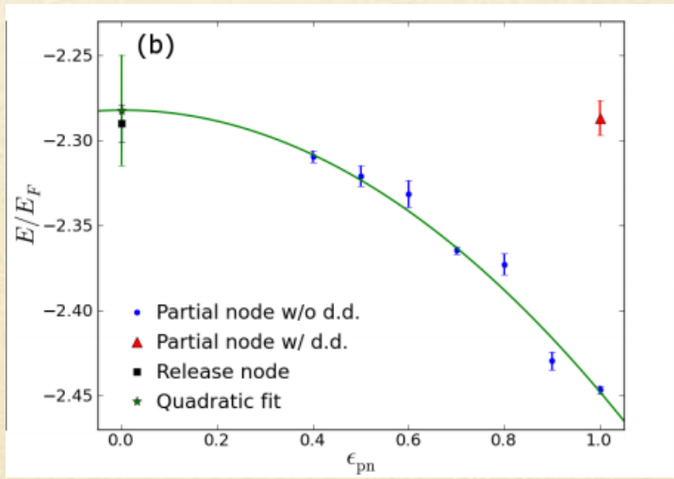

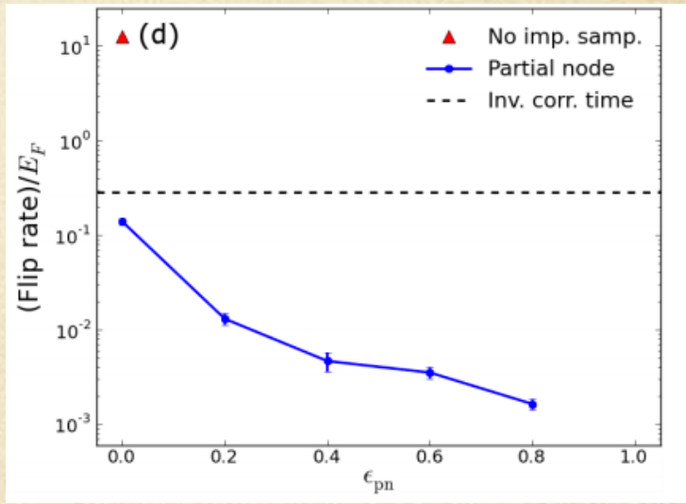

- Sample from Tensor Networks + Annihilation

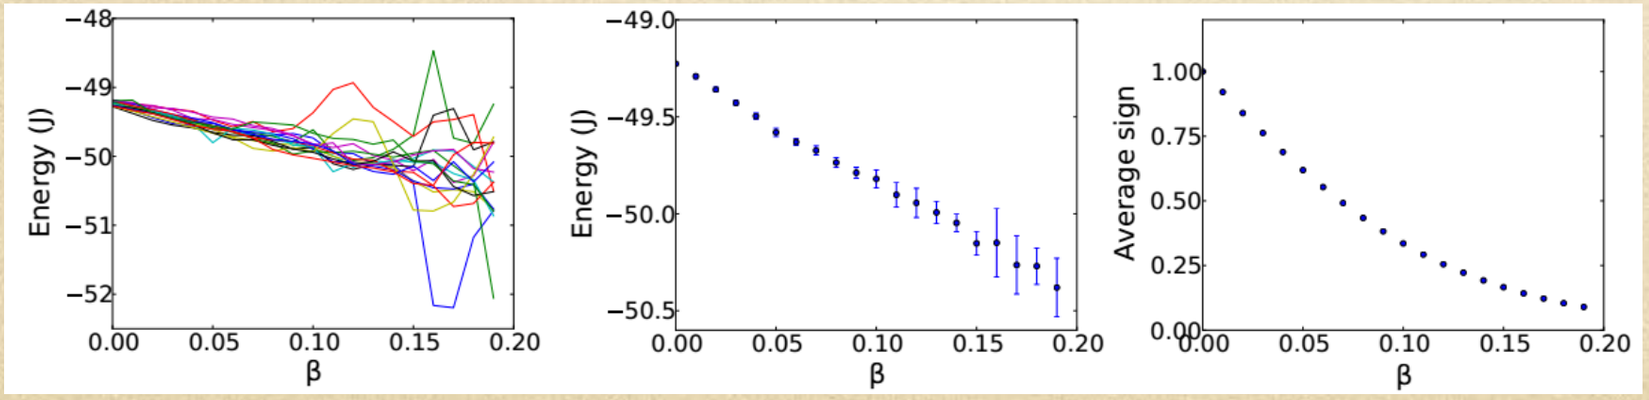

- RFCIQMC

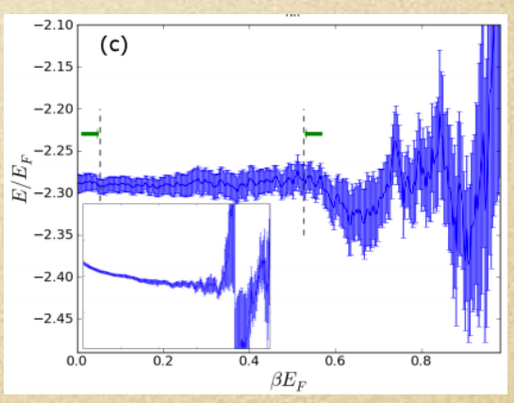

Approach II: Sample

Sign Problem - Efficiency as $\exp[-\beta\Delta E]$

Sample ++

- Annihilation + QMC
  - Brings up Delta E
    - Kalos
    - + initiator: Ali Alavi
- AFQMC Free projection
- SEMPS
- AFQMC release
- Partial Node FCIQMC
  - Importance Sample + Partial Fixed-Node + Annihilation

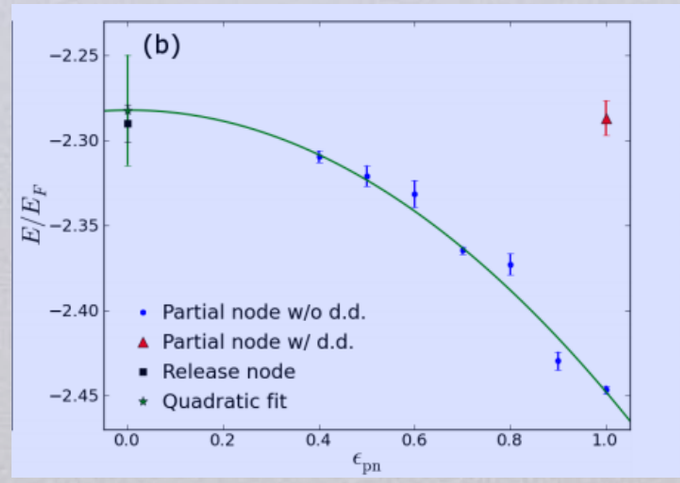

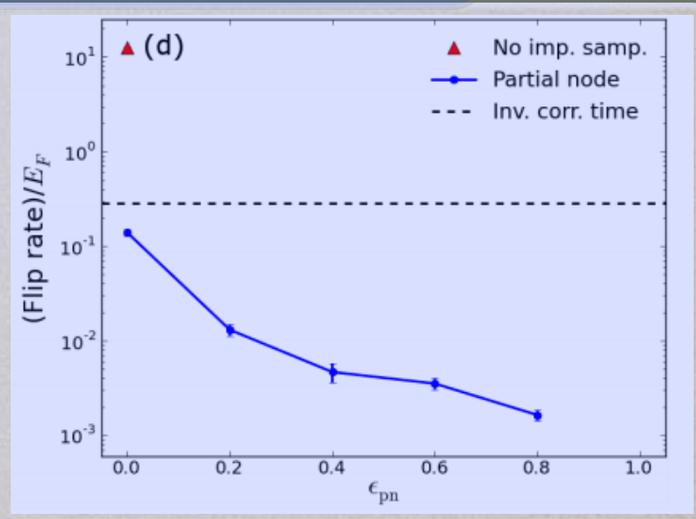

- Sample from Tensor Networks + Annihilation

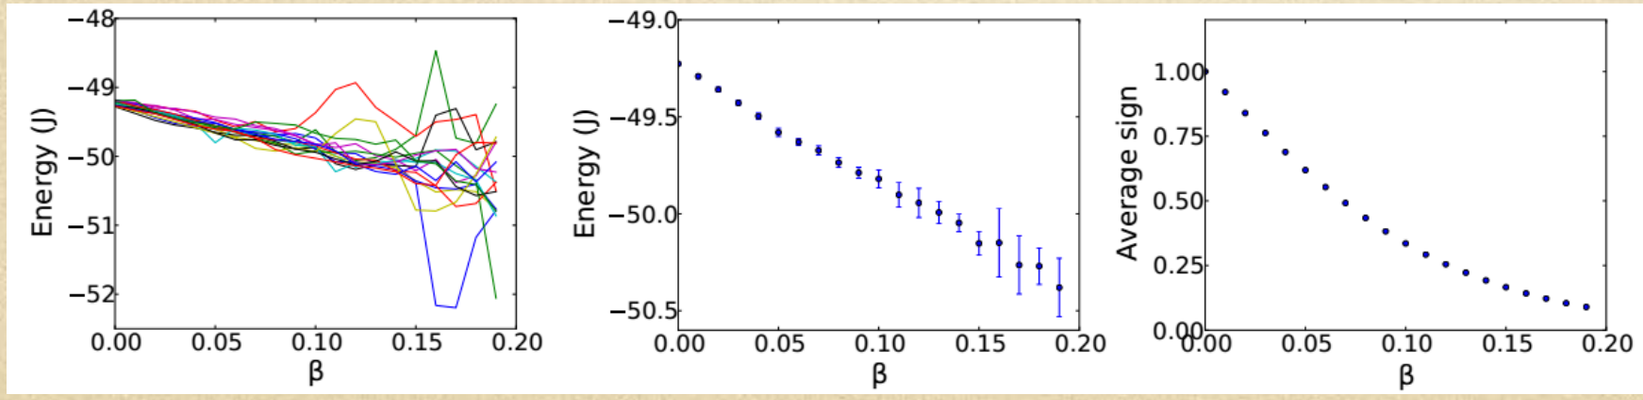

- RFCIQMC

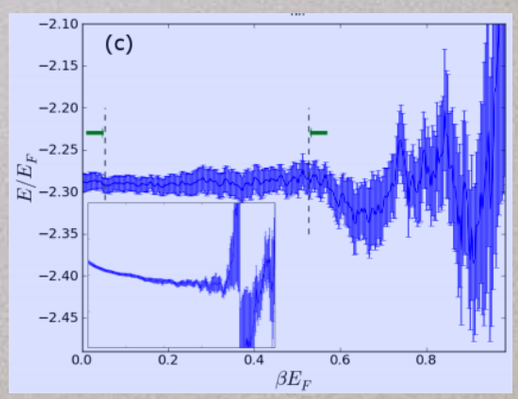

QMC: A Sign Problem

DMRG: A bond dimension problem.

Our goal is to write down an algorithm that has both a sign problem and a bond dimension problem.

he worst of bot s!

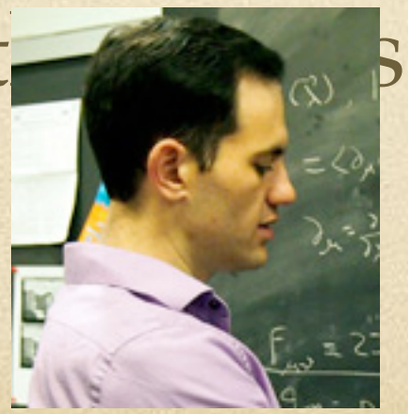

MPS

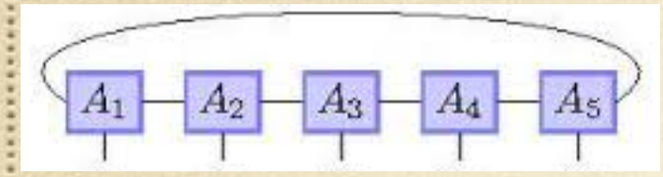

Product states

$=$ $\sum$

Sample

$|\langle MPS_1|C\rangle|^2$

weight

$\frac{1}{\langle MPS_1|C\rangle}$

Bond dimension 1

Imaginary Time Evolve

Bond dimension small

Imaginary Time Evolve

Bond dimension bigger

Imaginary Time Evolve

Energy: $\frac{\sum_i \langle \Psi_i | H | \Psi_T \rangle w_i}{\sum_i w_i}$

We'd really like perfect annihilation through all these paths.

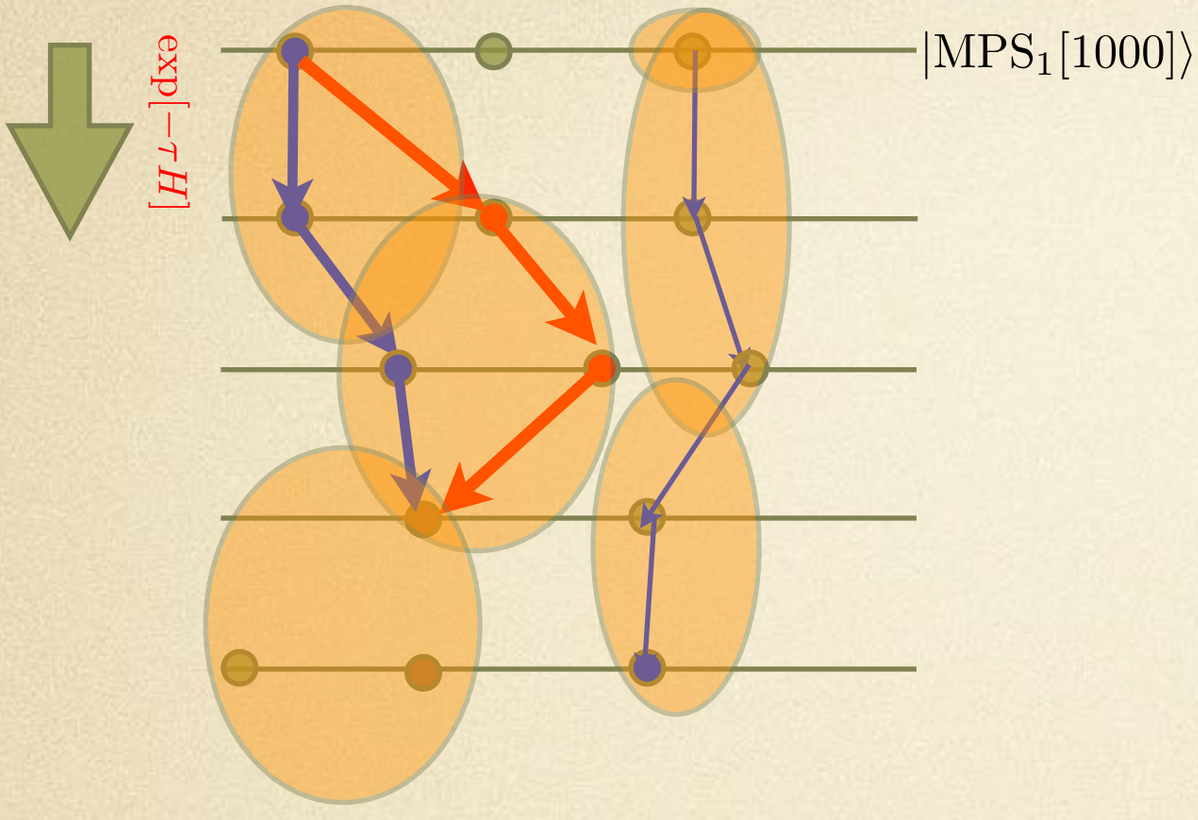

How can we do this?

We'd really like effectively higher bond dimension.

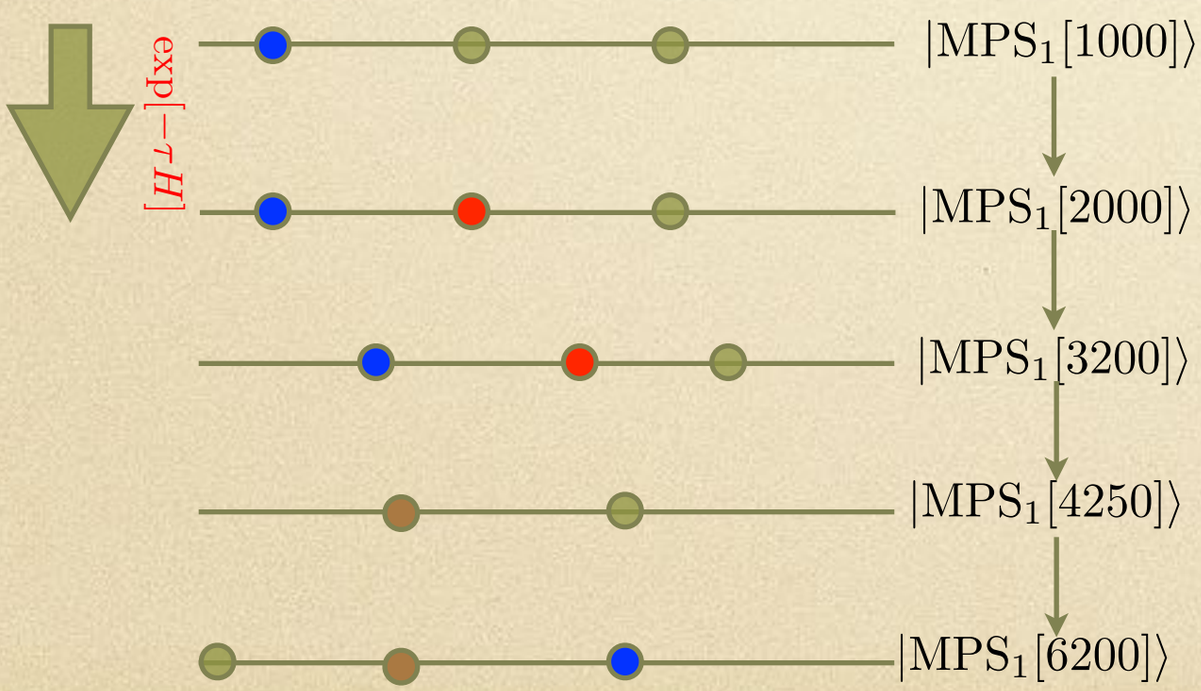

How can we do this?

Venerable history: BCS Superconductivity

Quantum Hall Effect

Model Wave-functions

Particularly valuable if the wave-function is conceptually simple and connects to analytical theory

Replace nobel prize winner with computers

HILBERT SPACE IS A BIG PLACE

$$|\Psi\rangle = \exp[-J(R)] \sum_k \alpha_k \det M_{\uparrow,k} \det M_{\downarrow,k}$$

$$J(R) = \sum_{ij} j(r_i - r_j)$$

$$M_{ij} = \phi_i(r_j)$$

3d function

$$R \equiv \{r_1, r_2, \ldots, r_n\}$$

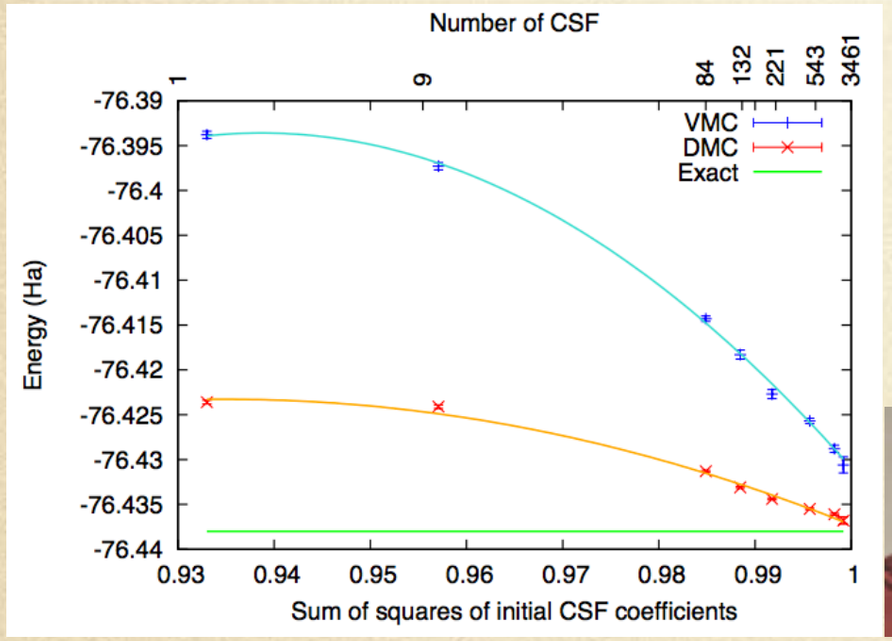

Unfortunately, these wave-functions are not compact enough. For the most interesting physical systems you need an exponential number of parameters (turn up the matrix size, the number of determinants, etc.) to ge converged energies.

We'd really like perfect annihilation through all these paths.

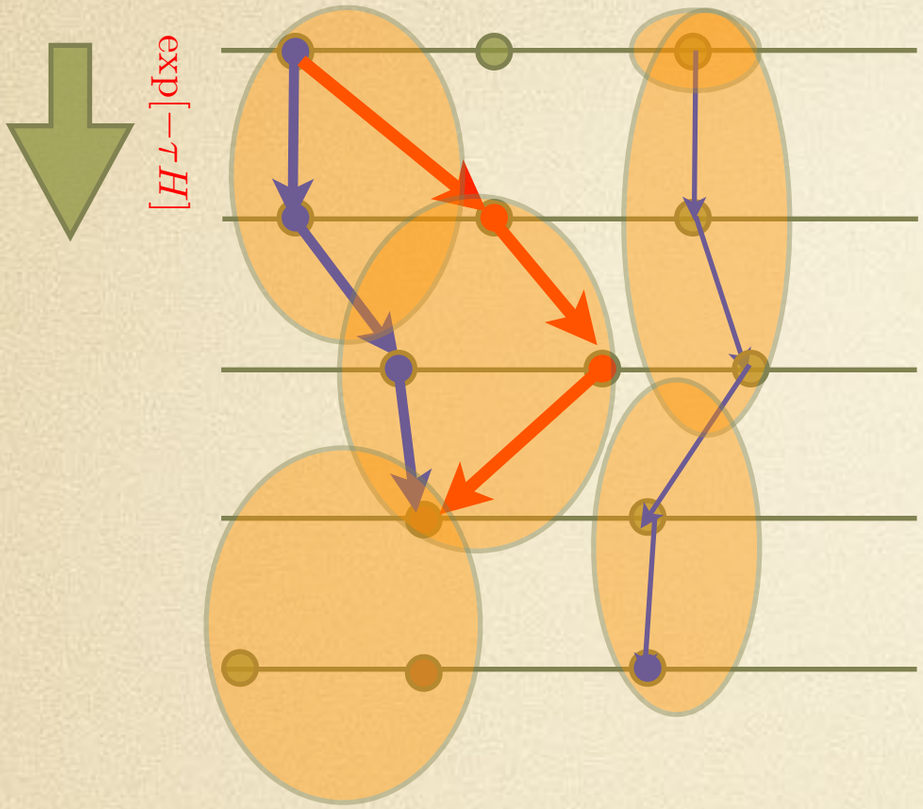

$$|MPS_1\rangle \approx |D_1\rangle + |D_3\rangle + |D_{20}\rangle + \ldots$$

$$\exp[-\tau H]|D_1\rangle + \exp[-\tau H]|D_3\rangle + \exp[-\tau H]|D_{20}\rangle + \ldots$$

Represented 'exactly' by MPS of small bond-dimension.

Sample $|\langle MPS_1|C\rangle|^2$

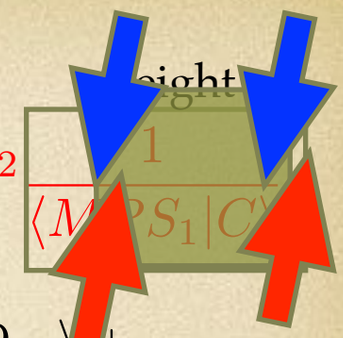

You run out of bond-dimension much slower.

You're already starting at the best MPS you can get for your bond dimension. You're guaranteed to be better.

Massively Parallel

You do have a bond-dimension problem.

If $\text{Sign}(\langle MPS_1|C\rangle) \neq \text{Sign}(\langle \Psi_0|C\rangle)$, you have a weak sign problem.

Proof of Principle

$\beta = 0.08$

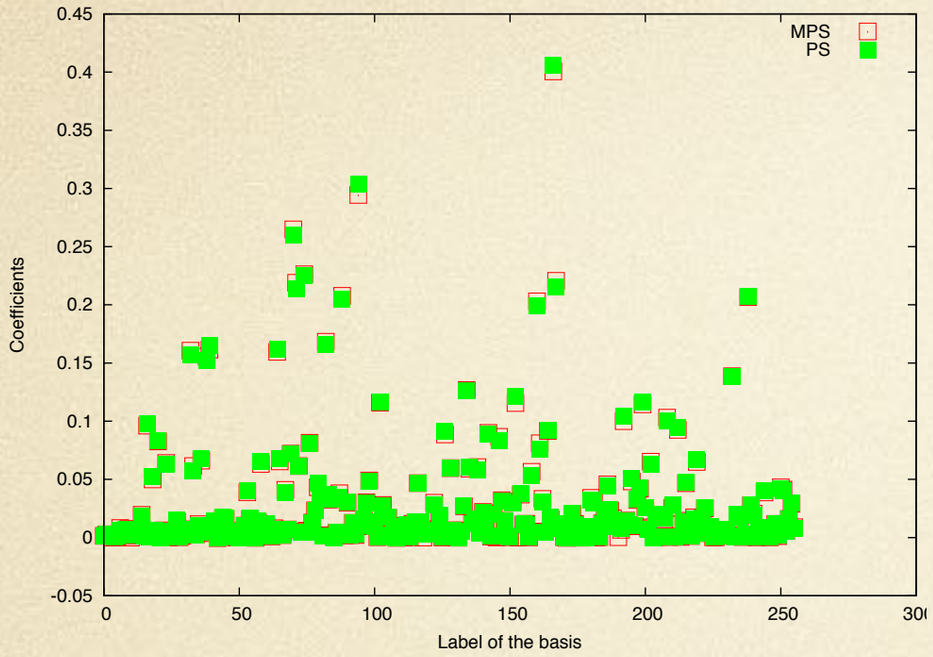

$\beta = 4.0$

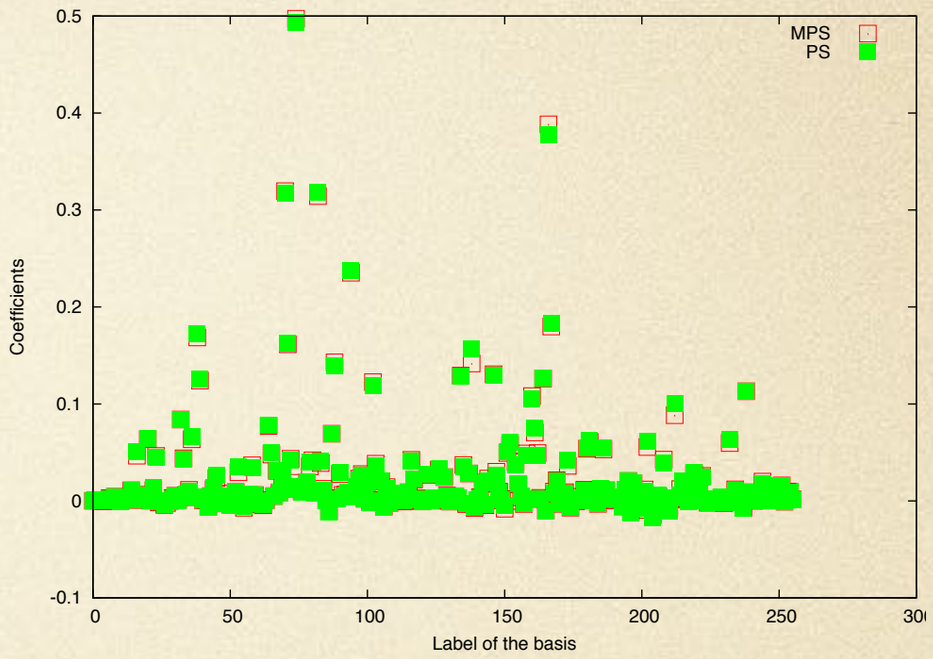

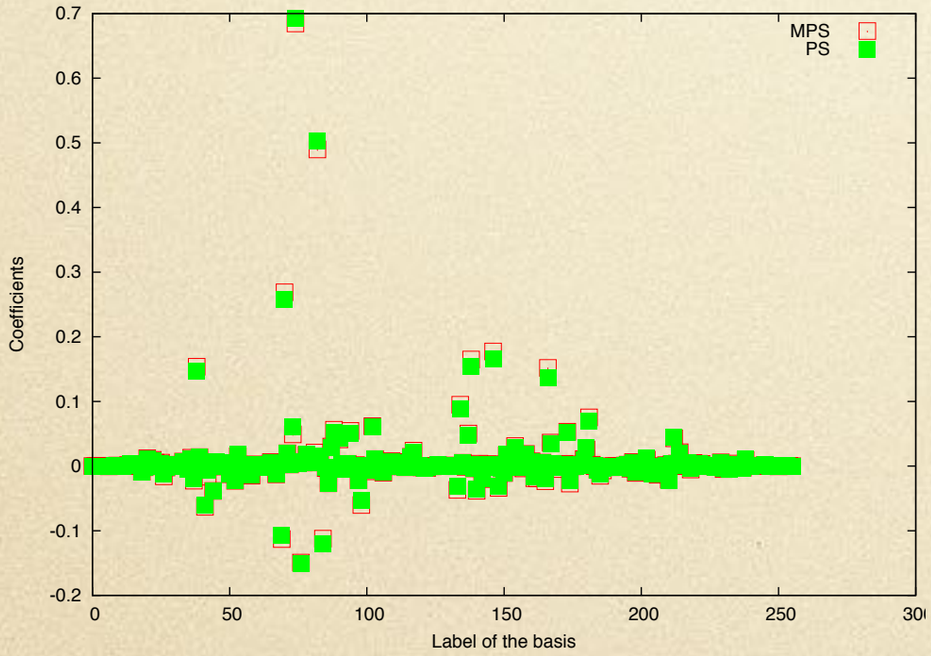

Random MPS (M=8),

Proof of Principle

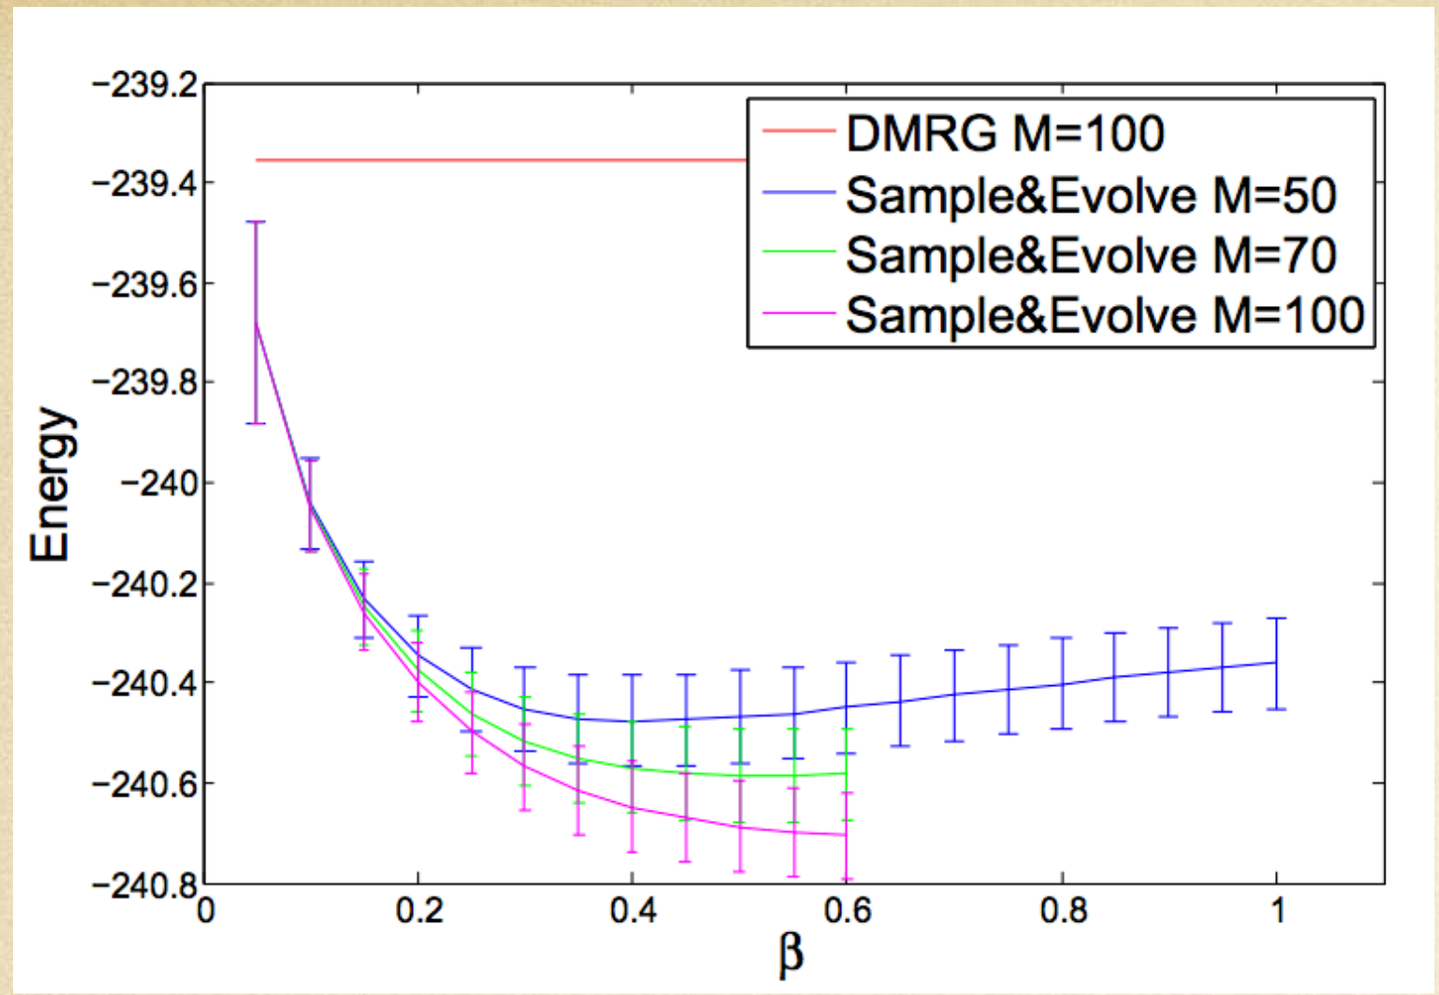

4 x 32 hubbard model

What to do when you run out of bond dimension?

SEMPS MC

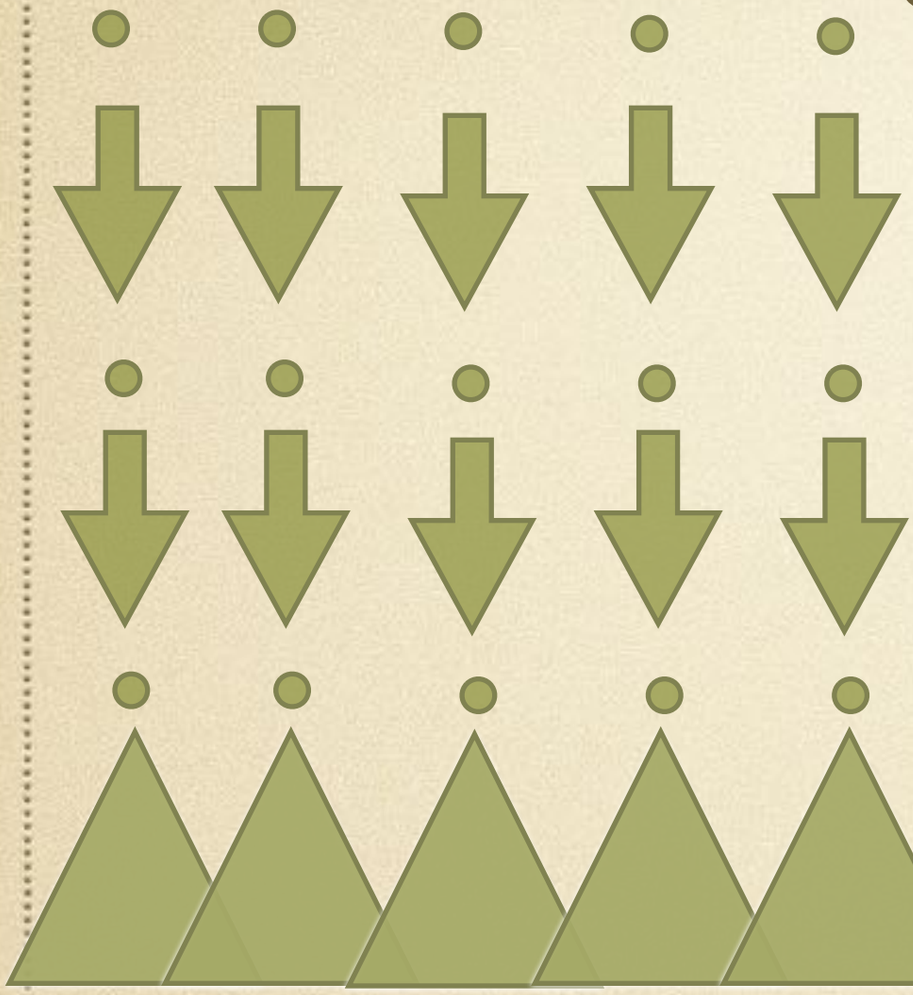

Resample here

Cone annihilation

Exact annihilation

A much smaller sign problem.

The best (or worst) of both worlds: SEMPS

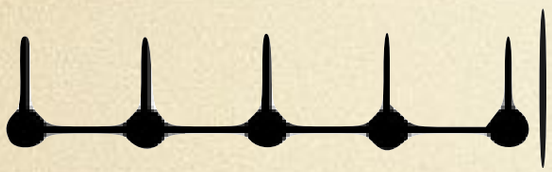

Sample the state

Release on MPS

Stochastically evolve

iTEBD

imaginary time evolve

SEMPS

imaginary time evolve

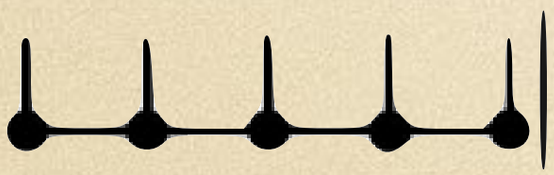

QMC to evaluate DMRG

Sample the state

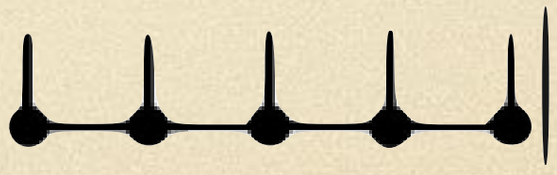

SEMPS MC

Stochastically evolve

Conclusions

DRMG+QMC gives us powerful new algorithms including

- Multi-MPS
- SEMPS
- Fixed-Node w/ MPS
- Release w/ MPS

<u>Pareto-Optimal:</u>

Multi-MPS SEMPS

Partial Node FCIQMC on Multi-MPS or Multi-Slater Jastrow

Release of CP AFQMC + SEMPS