New approaches to strongly interacting Fermi gases

Joaquín E. Drut

The Ohio State University

INT Program
Simulations and Symmetries
Seattle, March 2010

In collaboration with
Timo A. Lähde
Aalto University, Finland

Overview

- Strongly interacting Fermi gases
  - The unitary Fermi gas

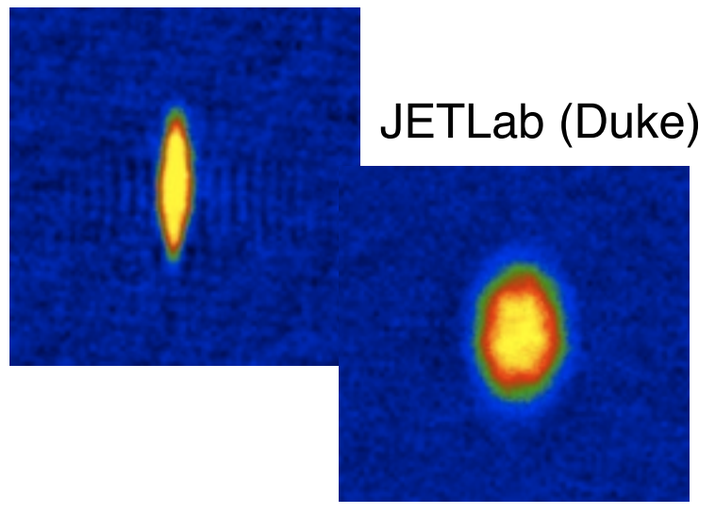

JETLab (Duke)

- New approaches
  - Determinantal MC, Hybrid MC & beyond
  - Beyond conventional hardware: GPUs
- Where do we go from here?

nVidia

The unitary Fermi gas

Spin 1/2 fermions, attractive interaction

$$r_0 \to 0 \ll n^{-1/3} \ll |a| \to \infty$$

Range of the interaction

Inter-particle distance

S-wave scattering length

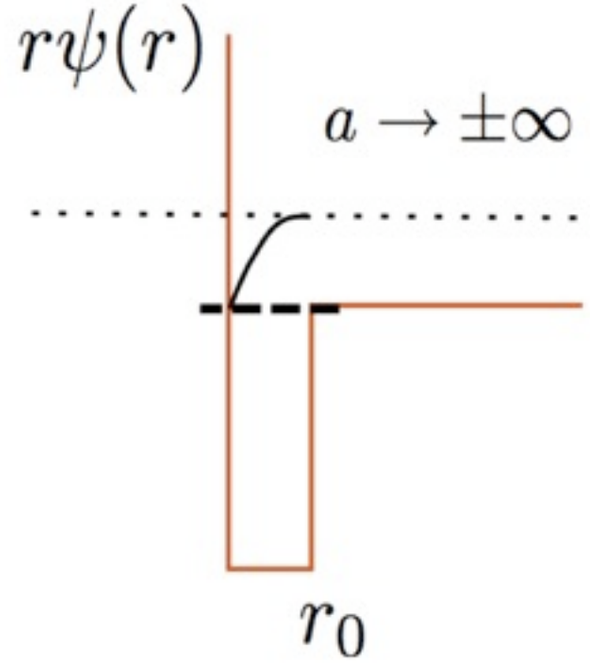

- As many scales as a free gas!

$$k_F = \hbar(3\pi^2 n)^{1/3} \qquad \varepsilon_F = \frac{\hbar^2}{2m}(3\pi^2 n)^{2/3}$$

- Qualitatively

Every dimensionful quantity should come as a power of ε_F times a **universal** constant/function.

- Quantitatively **?**

The BCS-BEC Crossover

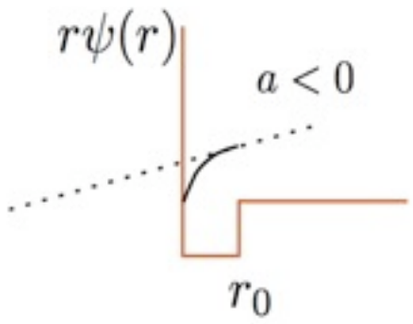

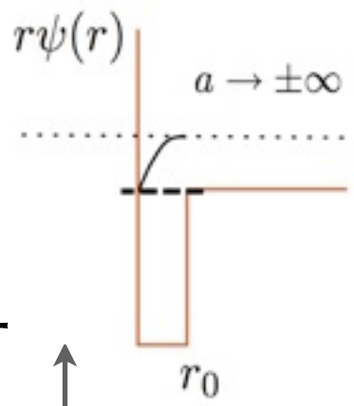

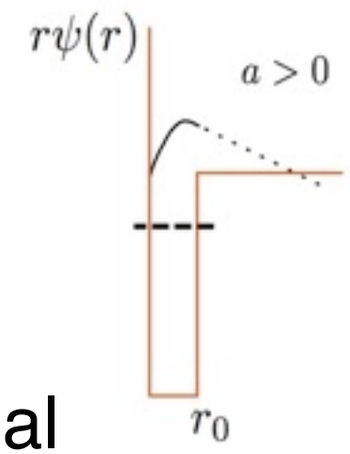

T

Normal

T_c

?

?

Normal

T_c

BCS
Superfluid

BEC
Superfluid

$1/k_F a$

$1 \ll k_F|a|$

Unitary regime

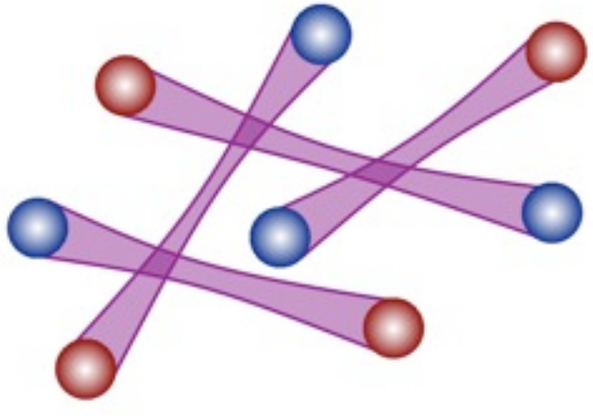

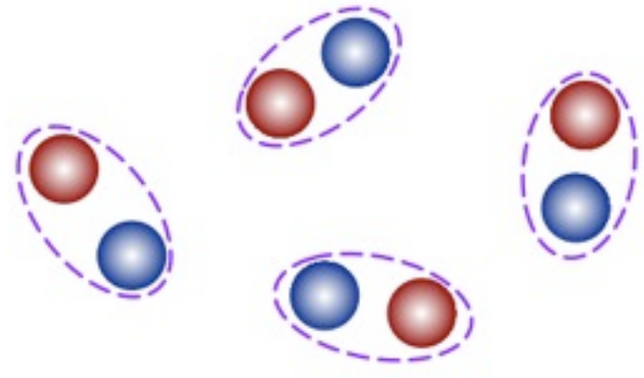

What do we know so far?

- Ground state energy per particle

J. Carlson *et al.*, Phys. Rev. Lett. **91**, 050401 (2003). (GFMC)

G. Astrakharchik *et al.*, Phys. Rev. Lett. **93**, 200404 (2004). (DMC)

"Bertsch parameter"

$$\xi \equiv \frac{E}{\frac{3}{5}\varepsilon_F N} \simeq 0.4$$

Also computed in various expansions:
1/d, 4-d, d-2, 1/N

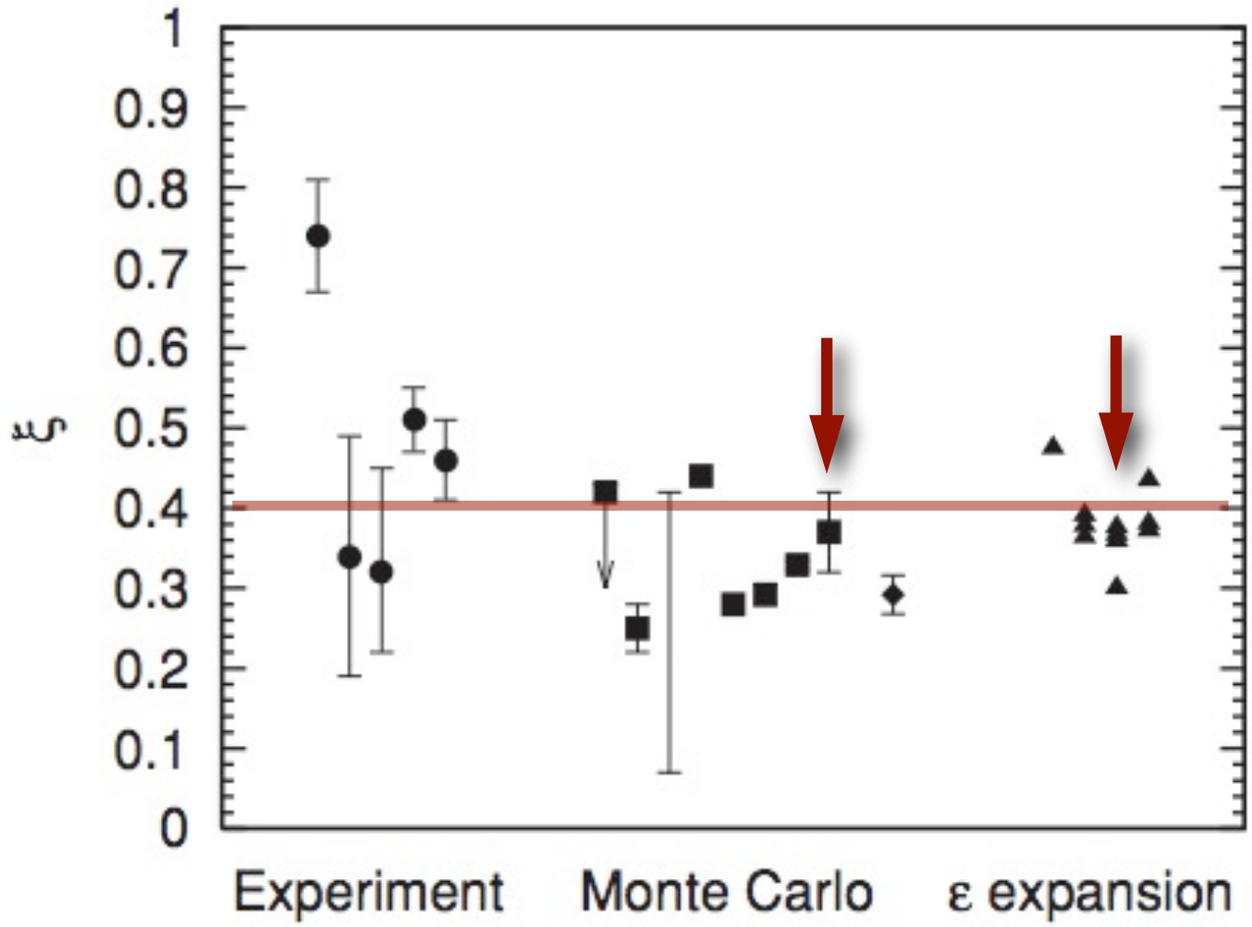

Abe & Seki,
Phys. Rev. C **79**, 054003 (2009).

What do we know so far?

- Equation of state

A. Bulgac, J. E. Drut, and P. Magierski,
Phys. Rev. Lett. **96**, 090404 (2006).

$$\xi(T/\varepsilon_F) \qquad \mu(T/\varepsilon_F)$$

Auxiliary Field
Determinantal Monte Carlo

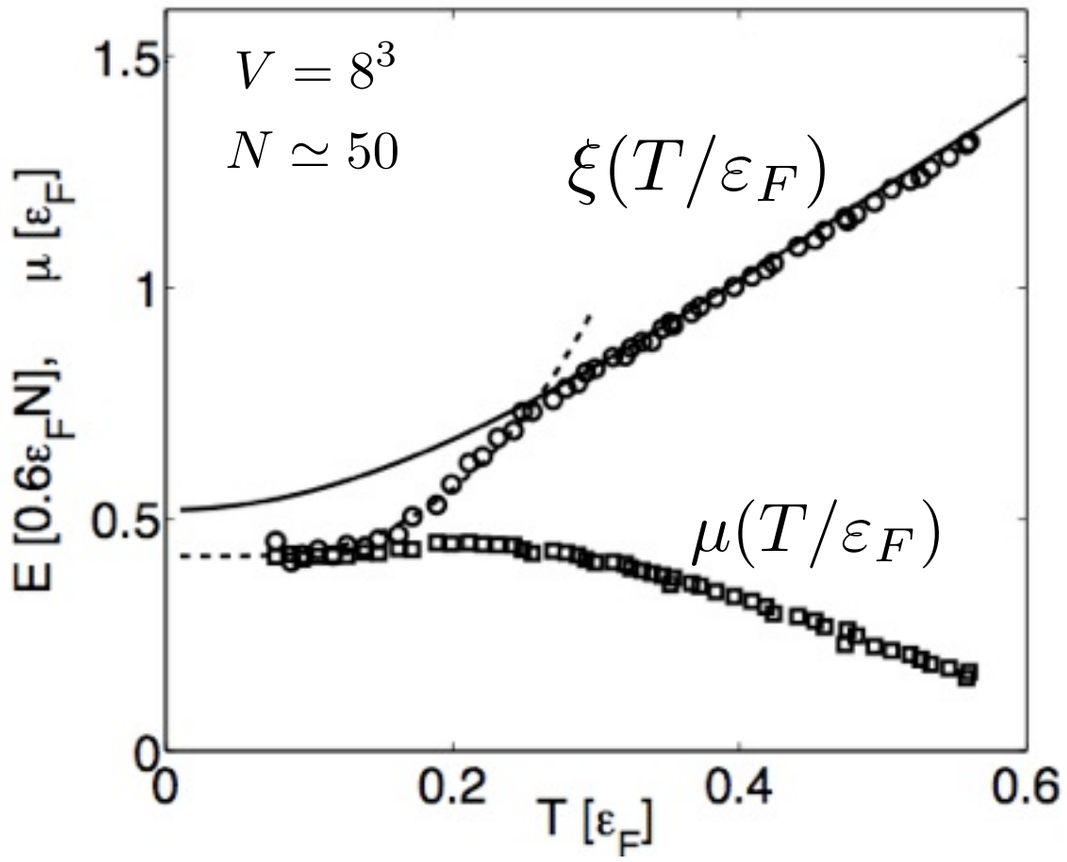

- Critical temperature

E. Burovski *et al.*, Phys. Rev. Lett. **96**, 160402 (2006).

$$T_c/\varepsilon_F \simeq 0.15$$

Diagrammatic Monte Carlo

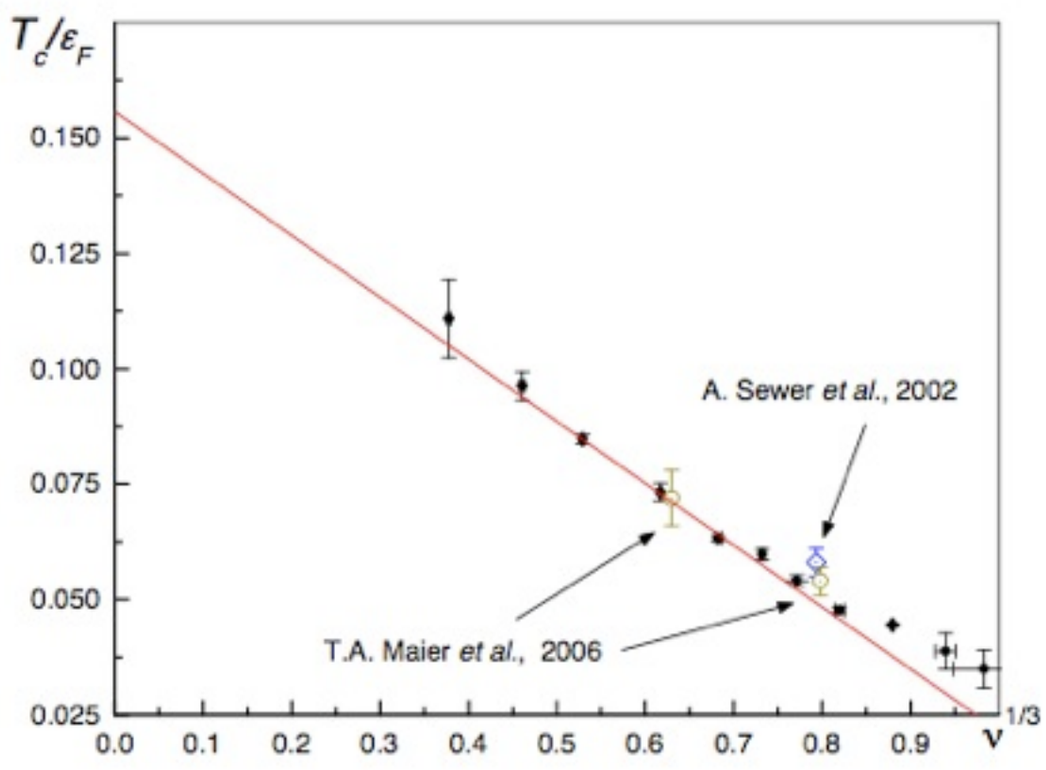

What do we know so far?

- Caloric curve in a trap (via Local Density Approximation)

A. Bulgac, J. E. Drut, and P. Magierski,
Phys. Rev. Lett. **99**, 120401 (2007).

L. Luo et al.
Phys. Rev. Lett. **98**, 080402 (2007).

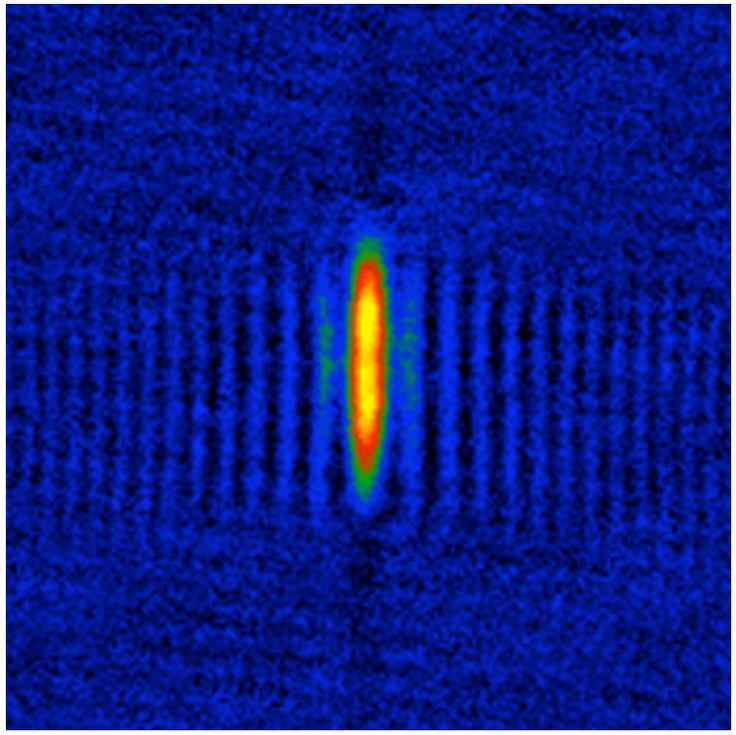

JETLab (Duke)

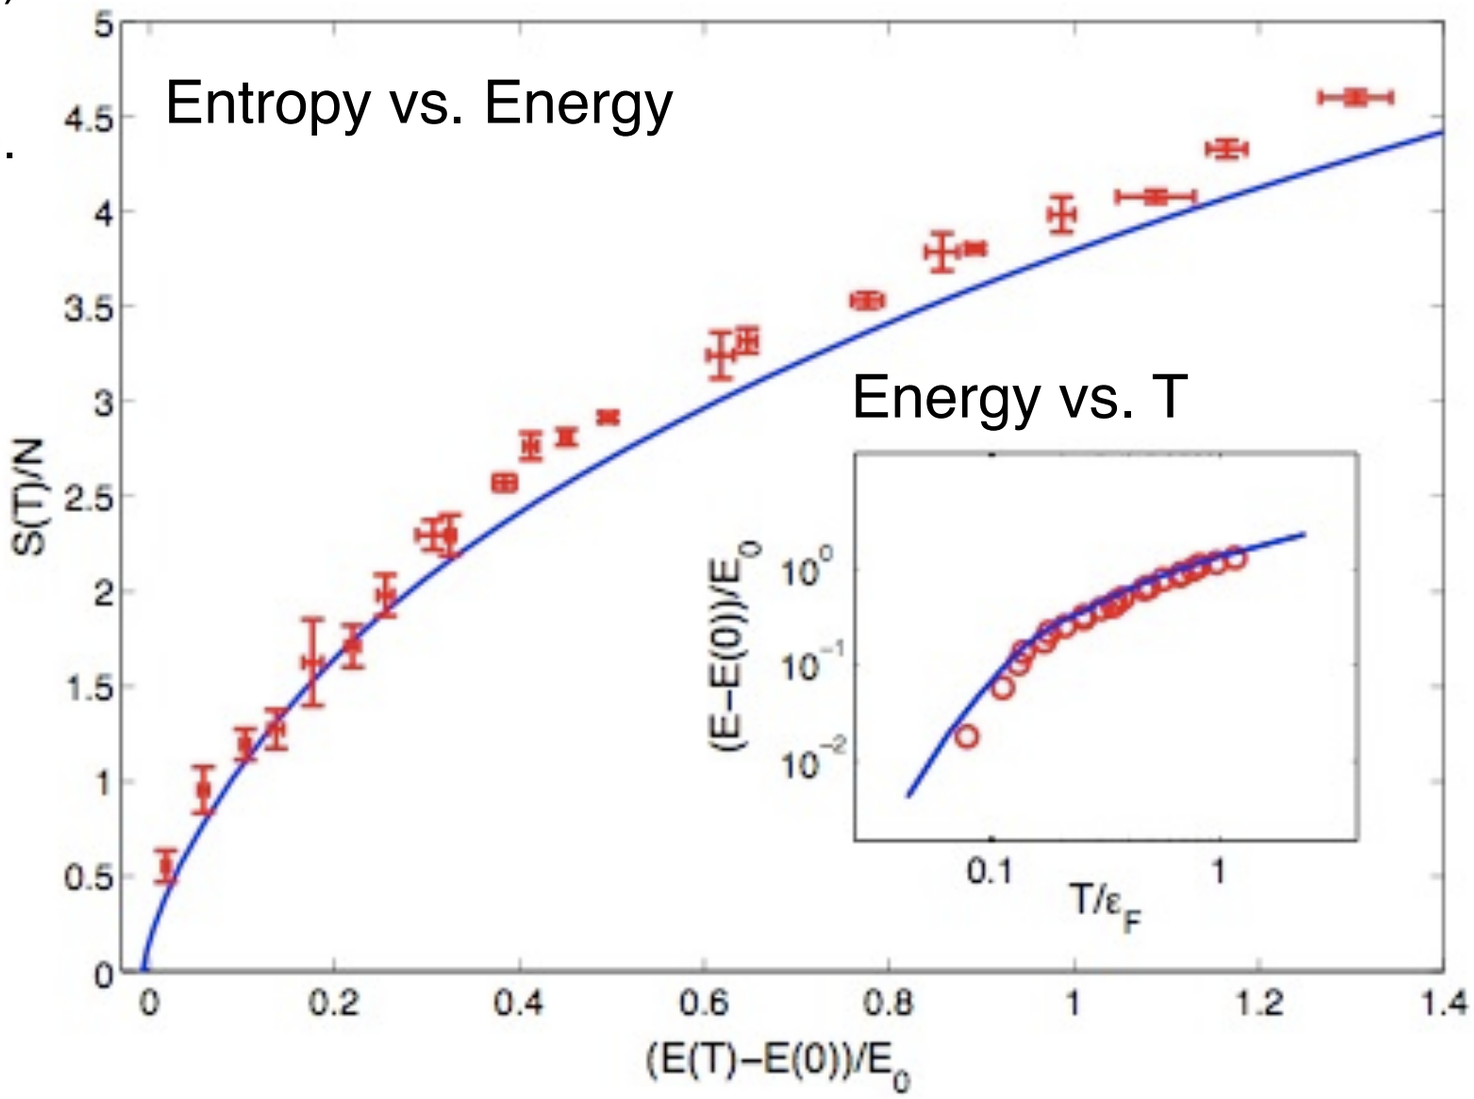

What do we know so far?

- Momentum distribution

G. E. Astrakharchik *et al.*, Phys. Rev. Lett. **95**, 230405 (2005).

S. Tan, Annals of Physics **323**, 2952 (2008).

E. Braaten and L. Platter, Phys. Rev. Lett. **100**, 205301 (2008).

$$n_k \underset{k \to \infty}{\longrightarrow} C/k^4$$

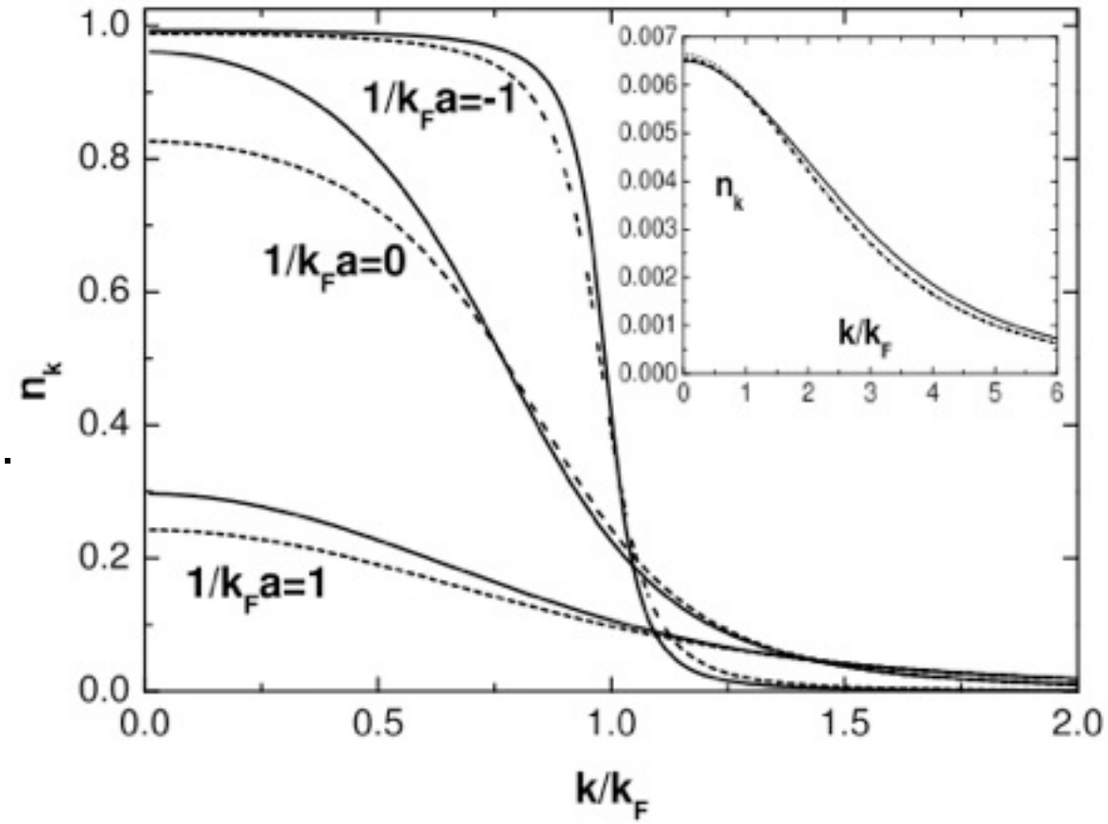

- Quasiparticle spectrum $E(p)$

J. Carlson and S. Reddy, Phys. Rev. Lett. **95**, 060401 (2005).

$$\Delta/\varepsilon_F \simeq 0.5$$ (GFMC)

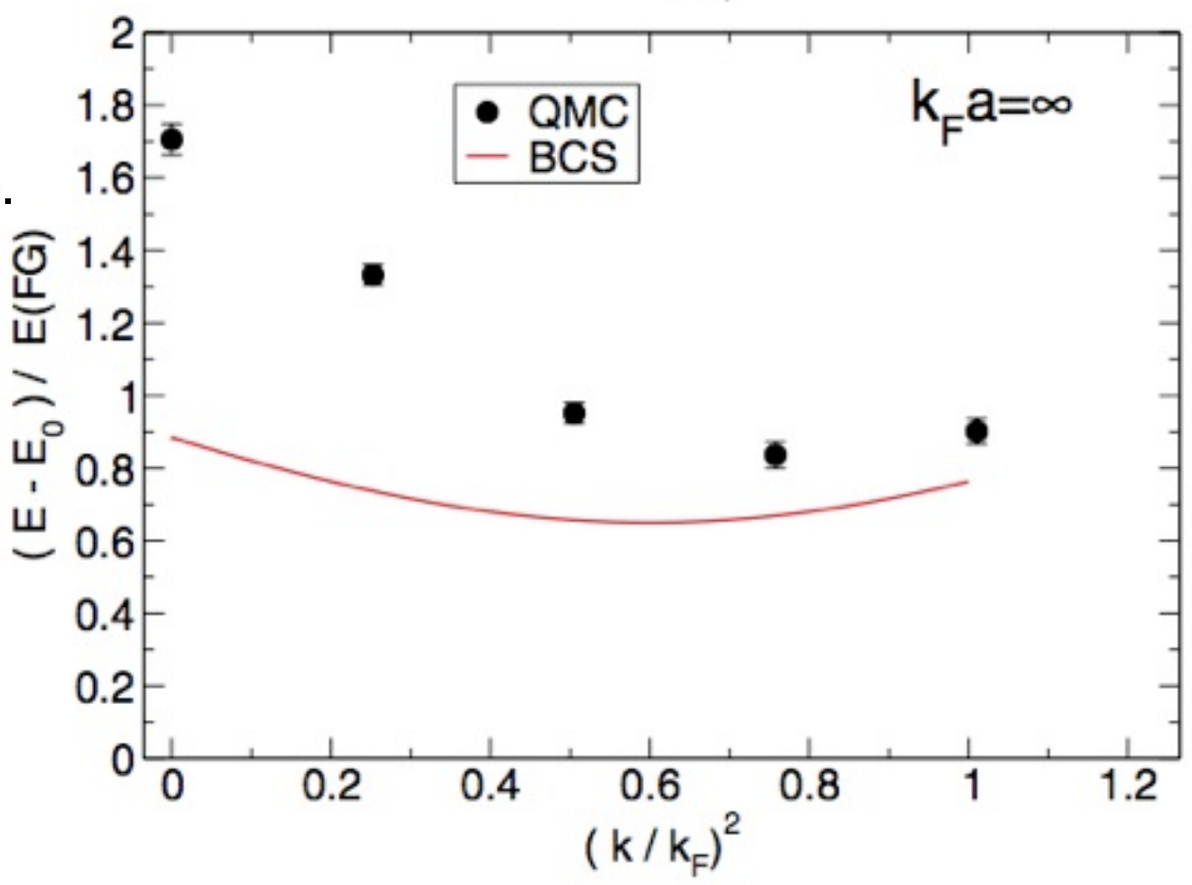

What do we know so far?

- Finite T quasiparticle spectrum

Magierski et al.
Phys. Rev. Lett. **103**, 210403 (2009).

- Compute finite T single-particle propagator using QMC

- Extract spectral density →

- Improvement on the way

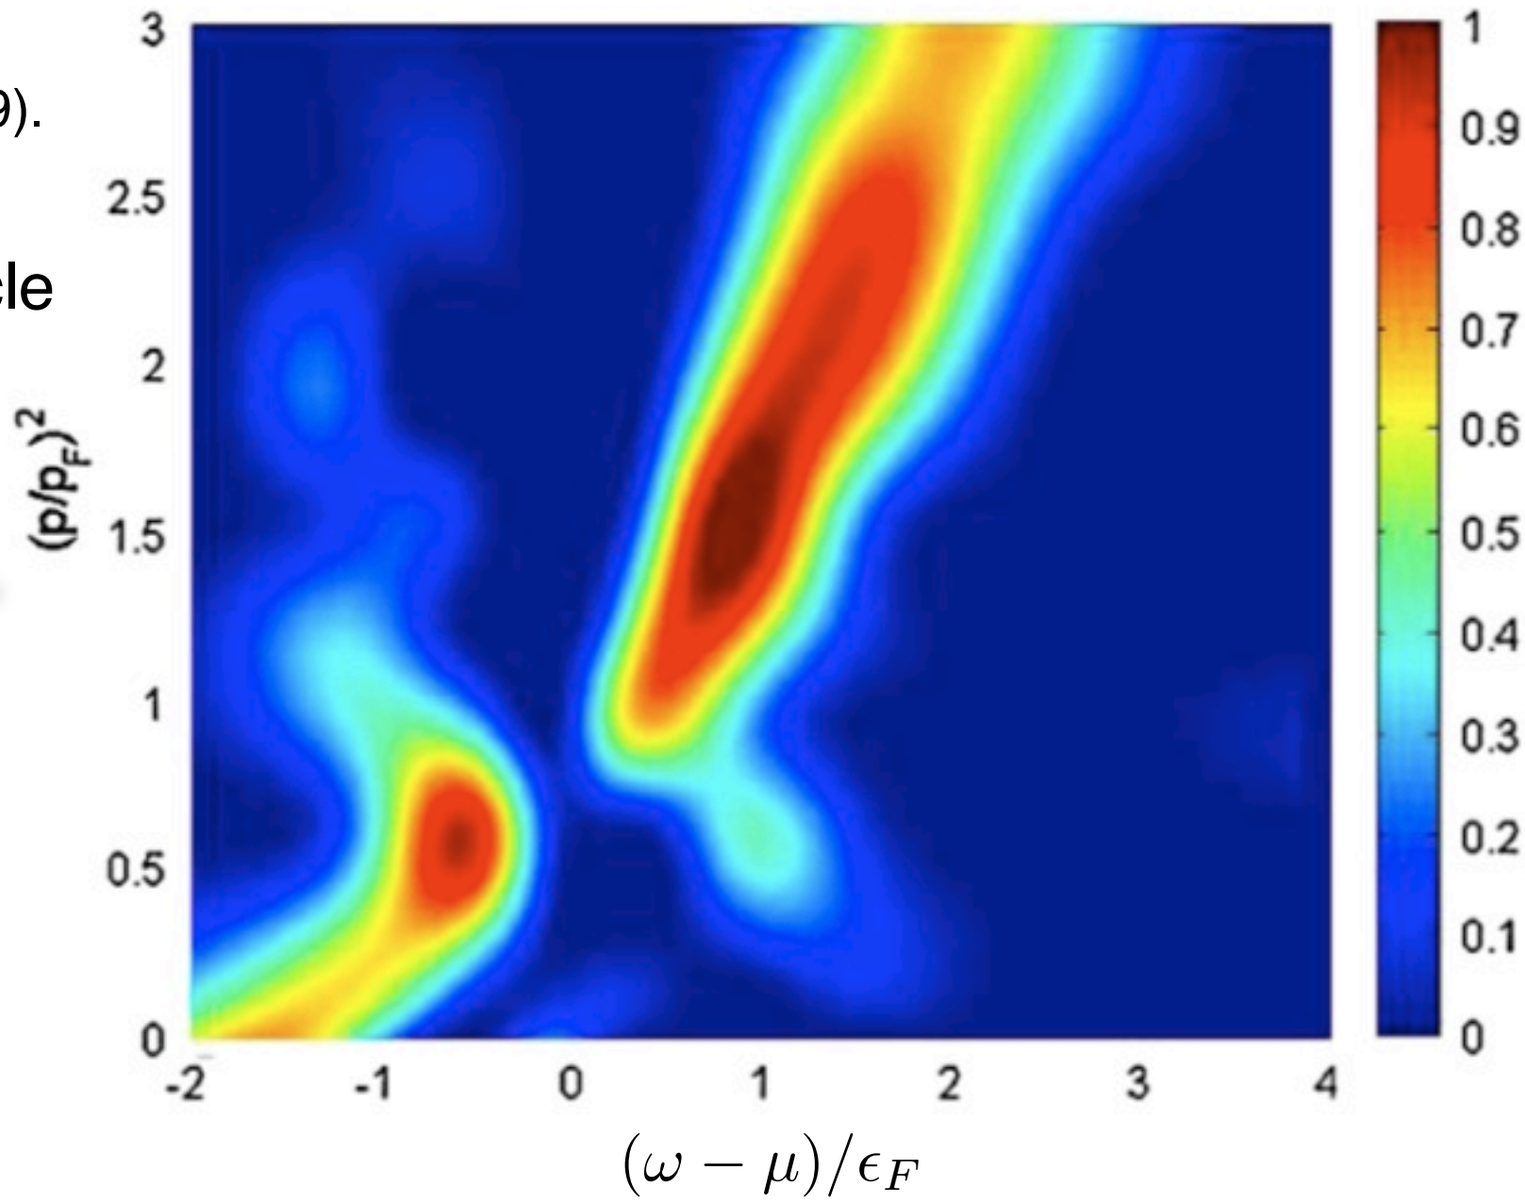

What about away from equilibrium?

- A hydrodynamic description (beyond the ideal case) requires knowledge of:
  - Equation of state: $P(T,n) = \frac{2}{3}\epsilon(T/\epsilon_F)$ ✔
  - Transport coefficients ✘
    - Bulk viscosity ζ
    - Shear viscosity η

    } → Momentum flow
    - Thermal conductivity κ → Heat flow

$$\frac{\eta}{s} \overset{?}{\geq} \frac{\hbar}{4\pi k_B}$$

Does the unitary gas saturate the KSS bound?

Shear & bulk viscosities

- Superfluid phase

  2 bulk viscosities vanish

  Son, PRL 020604 (2007)

  Phonon shear viscosity ◆

  Schäfer, PRA 063618 (2007)

  Rupak & Schäfer, PRA 053607 (2007)

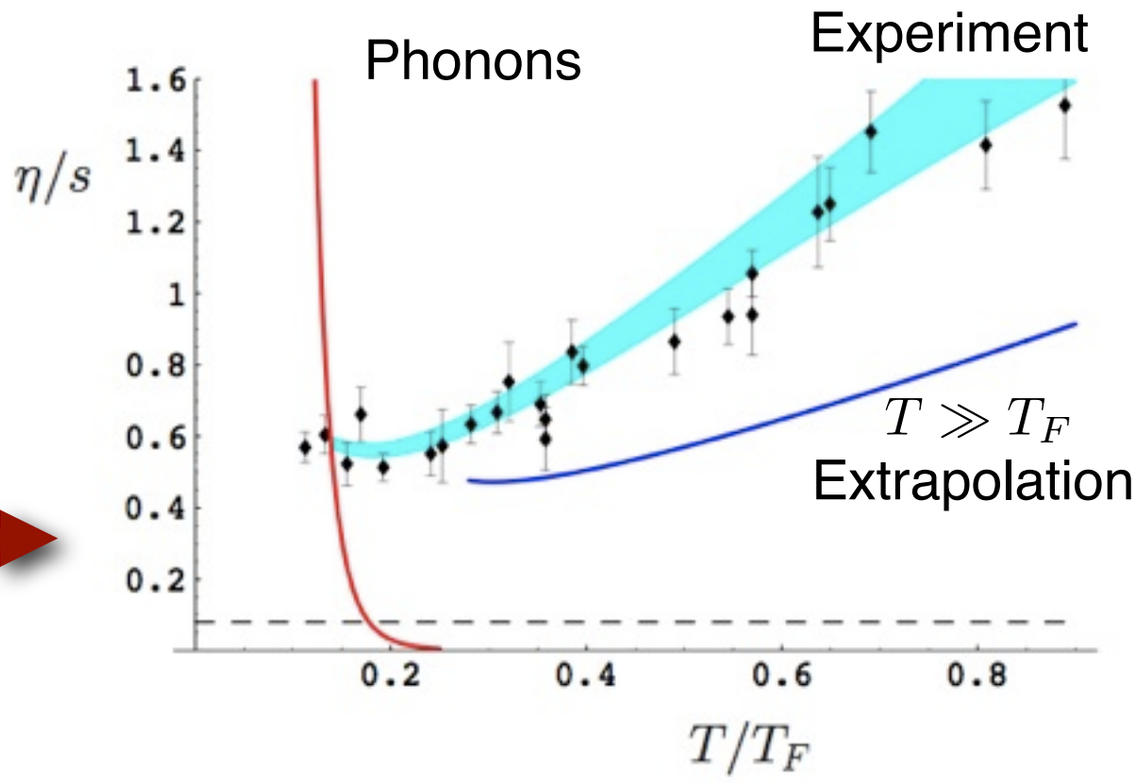

- Normal phase

  Bulk viscosity vanishes due to conformal invariance

  Son, PRL 020604 (2007)

  Shear viscosity $T \gg T_F$ ◆

  Rupak & Schäfer, PRA 053607 (2007)

  Massignan, Bruun & Smith, PRA 033607 (2005)

◆ Using kinetic theory

Thermal conductivity

Matt Braby, Jingyi Chao and Thomas Schäfer

arXiv:1003.2601v1

- Is this picture qualitatively correct around T_c ?

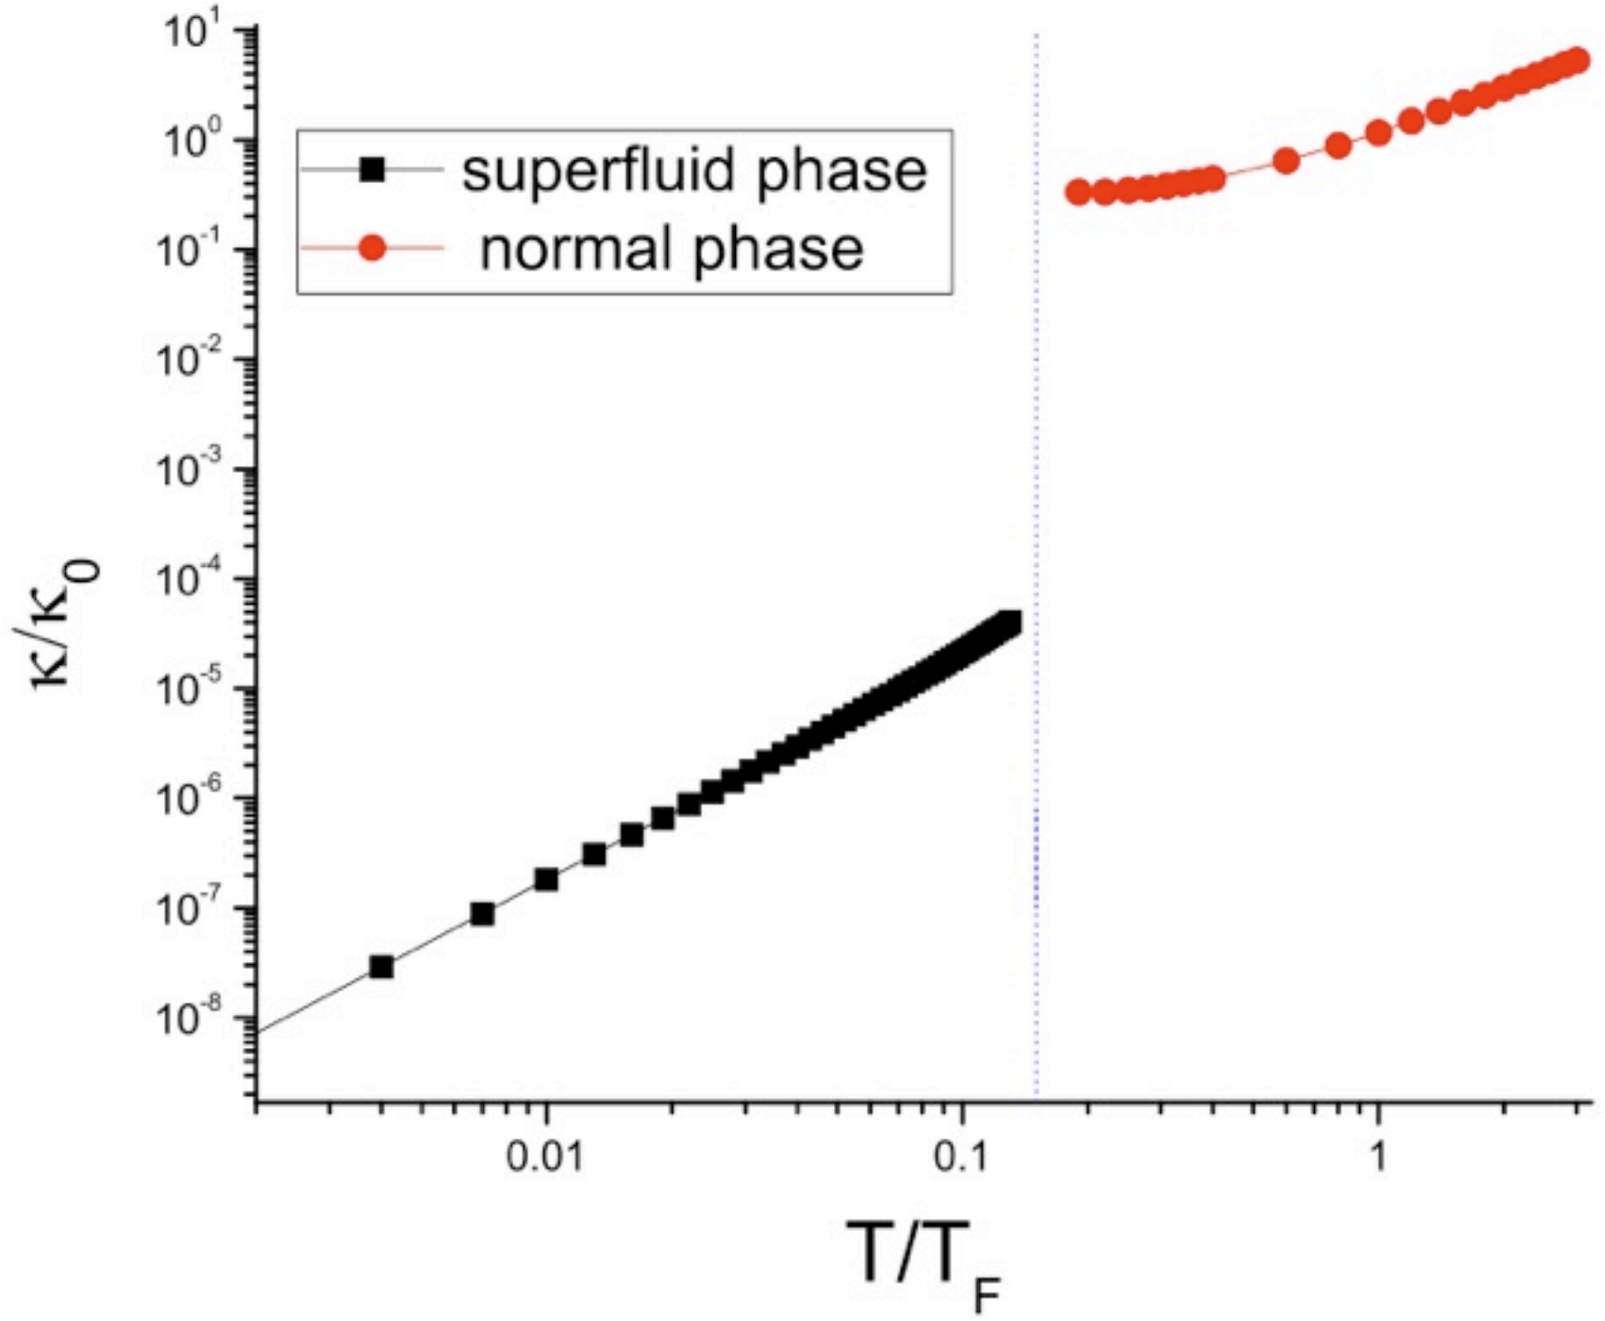

Current challenge: Universal transport coefficients

- How to proceed? We need
  - A code that can calculate the unitary gas at finite temperature.

    Challenge: larger system sizes required!
  - Ability to compute the stress-energy tensor Π_{ij}
  - Ability to compute correlations of Π_{ij}
  - Ability to determine spectral density from discrete data.

    Challenge: ill-defined inversion problem!

Our original approach: Determinantal Monte Carlo

- Discretize imaginary time → $e^{-\beta(H-\mu N)} = e^{-\tau(H-\mu N)} e^{-\tau(H-\mu N)} \dots e^{-\tau(H-\mu N)}$

$$\beta = \tau N_\tau$$

$$e^{-\tau/2(T-\mu N)} e^{-\tau V} e^{-\tau/2(T-\mu N)} + O(\tau^3)$$

Our original approach: Determinantal Monte Carlo

- Discretize imaginary time → $e^{-\beta(H-\mu N)} = e^{-\tau(H-\mu N)} e^{-\tau(H-\mu N)} \dots e^{-\tau(H-\mu N)}$

$$\beta = \tau N_\tau$$

$$e^{-\tau/2(T-\mu N)} e^{-\tau V} e^{-\tau/2(T-\mu N)} + O(\tau^3)$$

- Hubbard-Stratonovich transformation

$$e^{-\tau V(\mathbf{r})} = \int d\sigma \; e^{-v([\sigma],\mathbf{r},\tau) n_\uparrow(\mathbf{r})} e^{-v([\sigma],\mathbf{r},\tau) n_\downarrow(\mathbf{r})}$$

For each point in space-time!

Two-body interaction → Sum over all possible configurations of the "auxiliary field" $\sigma(\mathbf{r}, \tau)$

Determinantal Monte Carlo

- Partition function (collecting all the factors from the previous slide)

$$\mathcal{Z} = \mathrm{Tr}\ e^{-\beta(H-\mu N)} = \int \mathcal{D}\sigma \det[1 + U_\uparrow[\sigma]] \det[1 + U_\downarrow[\sigma]]$$

Fermions have been “integrated out”!

$$U_s[\sigma] = e^{-\tau/2(T_s - \mu N_s)} \times e^{-v([\sigma],\mathbf{r},\tau) n_s(\mathbf{r})} \times e^{-\tau/2(T_s - \mu N_s)} \times \ldots$$

Determinantal Monte Carlo

- Partition function (collecting all the factors from the previous slide)

$$\mathcal{Z} = \mathrm{Tr}\ e^{-\beta(H-\mu N)} = \int \mathcal{D}\sigma \det[1+U_\uparrow[\sigma]] \det[1+U_\downarrow[\sigma]]$$

Fermions have been "integrated out"!

$$U_s[\sigma] = e^{-\tau/2(T_s-\mu N_s)} \times e^{-v([\sigma],\mathbf{r},\tau)n_s(\mathbf{r})} \times e^{-\tau/2(T_s-\mu N_s)} \times \ldots$$

FFT FFT FFT

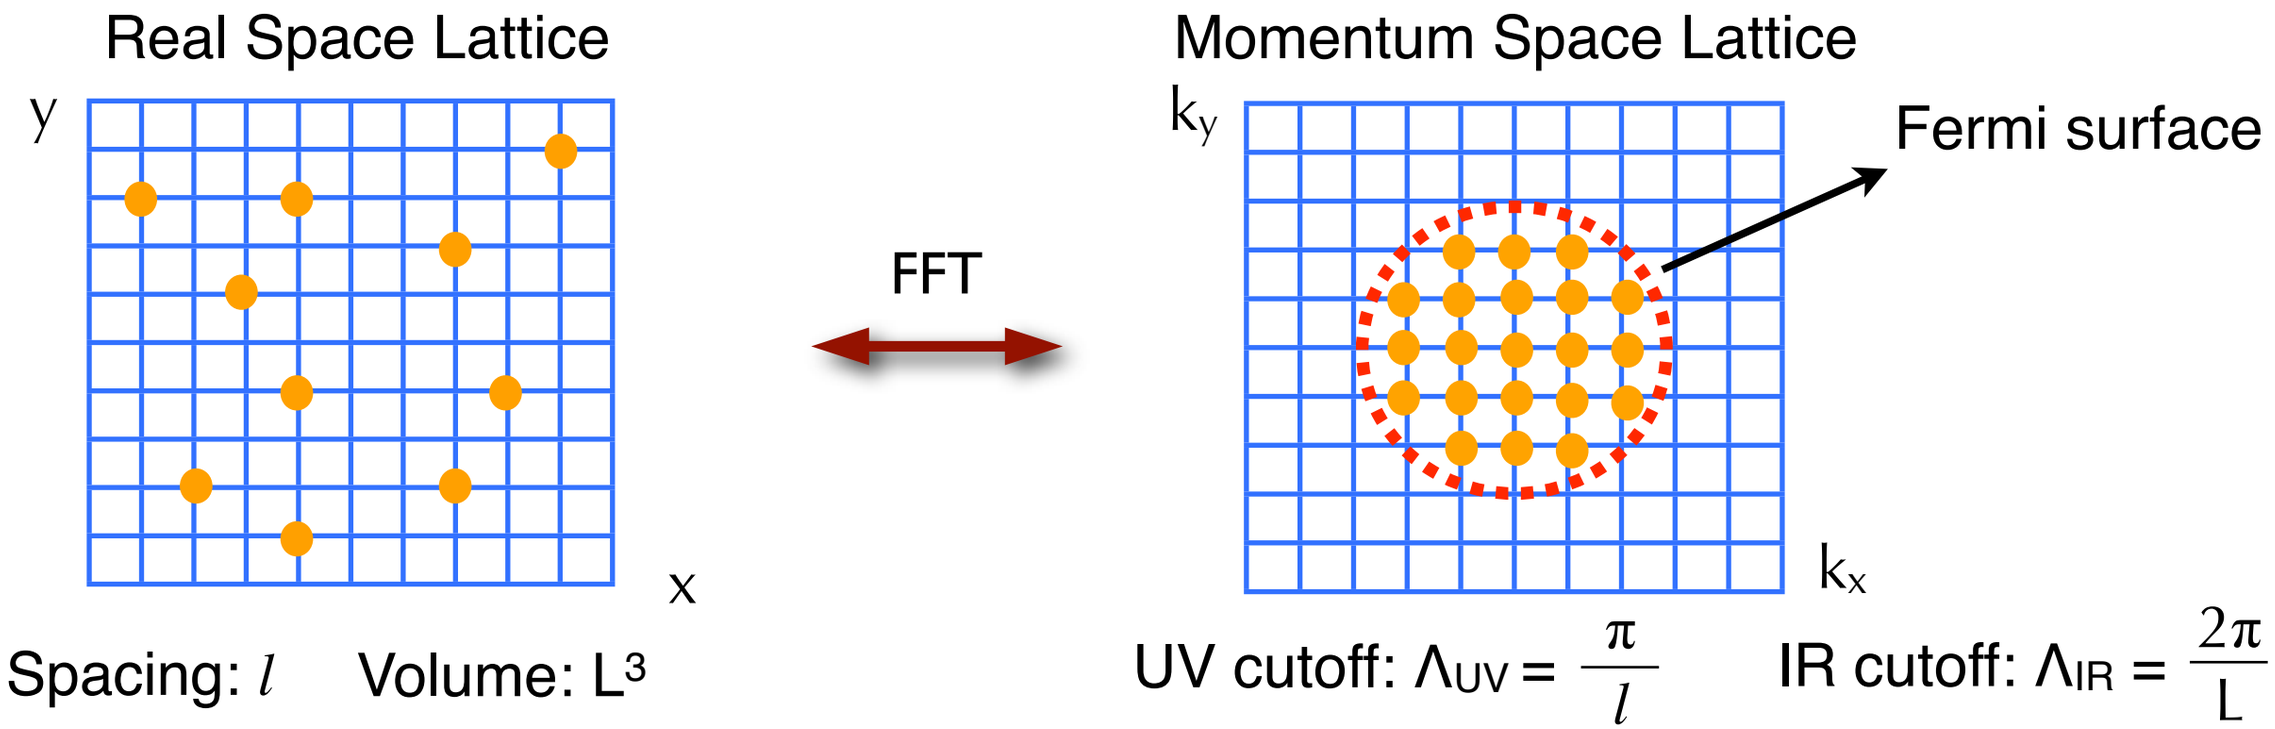

Spacing: l Volume: L^3

UV cutoff: $\Lambda_{UV} = \frac{\pi}{l}$ IR cutoff: $\Lambda_{IR} = \frac{2\pi}{L}$

How efficient is this algorithm?

Determinantal Monte Carlo

- Scaling of computational cost

$$L^3 \times N_\tau \times \underline{L^3 \times \log L^3} \times \underline{L^3 \times N_\tau}$$

Full basis

Time evolution

FFT

Full update (sweep) via local updates

Determinantal Monte Carlo

- Scaling of computational cost

$$L^3 \times N_\tau \times \underline{L^3 \times \log L^3} \times \underline{L^3 \times N_\tau}$$

Full basis

Time evolution

FFT

Full update (sweep) via local updates

It doesn't matter how big your computer is...

A more efficient way

- Scaling of computational cost

$$\cancel{L^3} \times N_\tau \times \underbrace{L^3 \times \log L^3}_{\text{FFT}} \times \underbrace{\cancel{L^3} \times N_\tau}_{\text{Full update (sweep) via local updates}}$$

Full basis

Time evolution

Use **pseudofermions**!

Stochastic calculation of the determinant

Use **Hybrid Monte Carlo**!

Global updates of the auxiliary field via molecular dynamics

A more efficient way

- Scaling of computational cost

$$\cancel{L^3} \times N_\tau \times \underbrace{L^3 \times \log L^3}_{\text{FFT}} \times \cancel{\underbrace{L^3 \times N_\tau}_{\text{Full update (sweep) via local updates}}}$$

Full basis

Time evolution

Use **pseudofermions**!

Use **Hybrid Monte Carlo**!

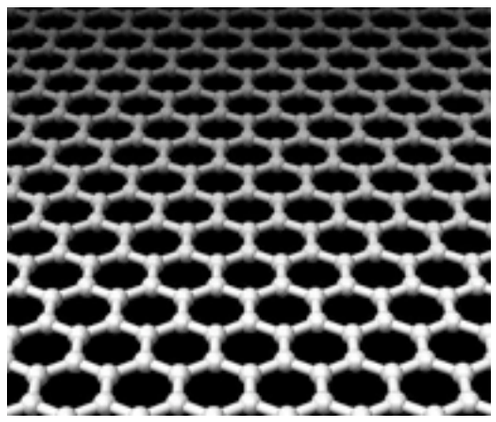

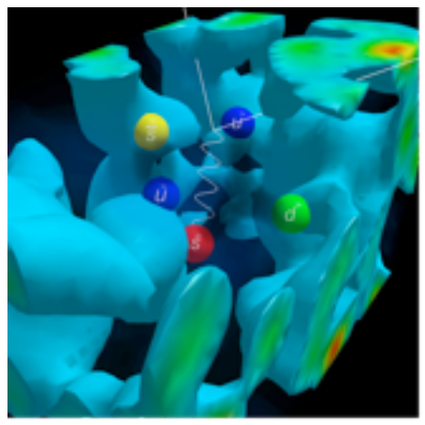

This is how state-of-the-art lattice QCD is done! (and graphene!)

Drut & Lähde, Phys. Rev. Lett. **102**, 026802 (2009)

...and now you are in better shape to use big computers!

A few words about HMC...

- Pseudofermions

$$\mathcal{Z} = \int \mathcal{D}\sigma \ \mathrm{Det}\left[M^T M(\sigma)\right] = \int \mathcal{D}\sigma \mathcal{D}\varphi^\dagger \mathcal{D}\varphi \ \exp\left(-\varphi^\dagger \left[M^T M(\sigma)\right]^{-1} \varphi\right)$$

  - φ is the **pseudofermion** field
  - M is sparse, of size $(VN_\tau) \times (VN_\tau)$

$$M \equiv \begin{pmatrix} 1 & 0 & 0 & 0 & \dots & U_{N_\tau} \\ -U_1 & 1 & 0 & 0 & \dots & 0 \\ 0 & -U_2 & 1 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & -U_{N_\tau-2} & 1 & 0 \\ 0 & 0 & \dots & 0 & -U_{N_\tau-1} & 1 \end{pmatrix}$$

$$\mathrm{Det}[M^T M] = \det(1+U)^2$$

Space-time formulation **Spatial formulation**

A few words about HMC...

- Pseudofermions

$$\mathcal{Z} = \int \mathcal{D}\sigma \ \mathrm{Det}\left[M^T M(\sigma)\right] = \int \mathcal{D}\sigma \mathcal{D}\varphi^\dagger \mathcal{D}\varphi \ \exp\left(-\varphi^\dagger \left[M^T M(\sigma)\right]^{-1} \varphi\right)$$

  - φ is the **pseudofermion** field
  - M is sparse, of size $(VN_\tau) \times (VN_\tau)$

Problem: How do we change σ_i as much as possible without obtaining an extremely improbable configuration?

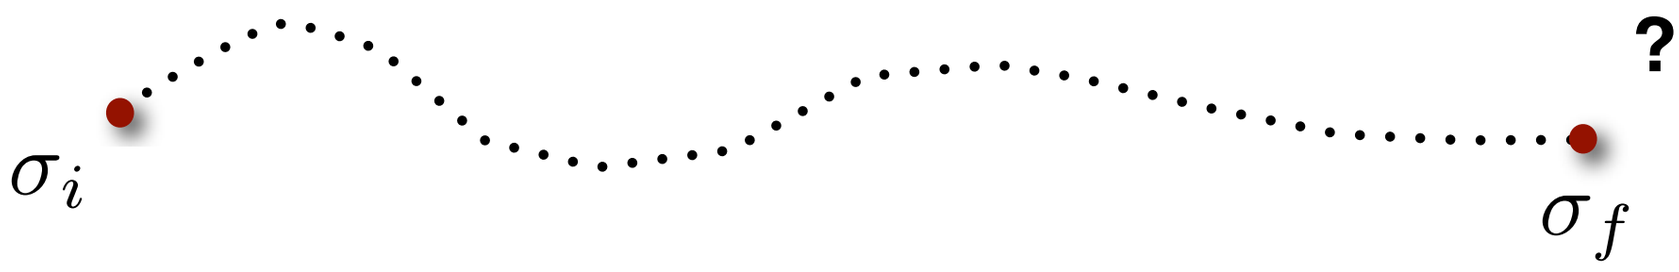

Determinantal MC: The determinant is a very non-local and non-linear object; we can only perform local changes, or else we obtain a very improbable configuration

A few words about HMC...

- Pseudofermions

$$\mathcal{Z} = \int \mathcal{D}\sigma \ \mathrm{Det}\left[M^T M(\sigma)\right] = \int \mathcal{D}\sigma \mathcal{D}\varphi^\dagger \mathcal{D}\varphi \ \exp\left(-\varphi^\dagger \left[M^T M(\sigma)\right]^{-1} \varphi\right)$$

  - φ is the **pseudofermion** field
  - M is sparse, of size $(VN_\tau) \times (VN_\tau)$

- Molecular dynamics

$$\mathcal{Z} = \int \mathcal{D}\pi \mathcal{D}\sigma \mathcal{D}\varphi^\dagger \mathcal{D}\varphi \ \exp\left(\underbrace{-\pi^2/2 - \varphi^\dagger \left[M^T M(\sigma)\right]^{-1} \varphi}_{-\mathcal{H}}\right)$$

  - Enables **global updates!**

$$\dot{\sigma} = \frac{\delta \mathcal{H}}{\delta \pi} = \pi$$

$$\dot{\pi} = -\frac{\delta \mathcal{H}}{\delta \sigma} = F[\sigma, \varphi]$$

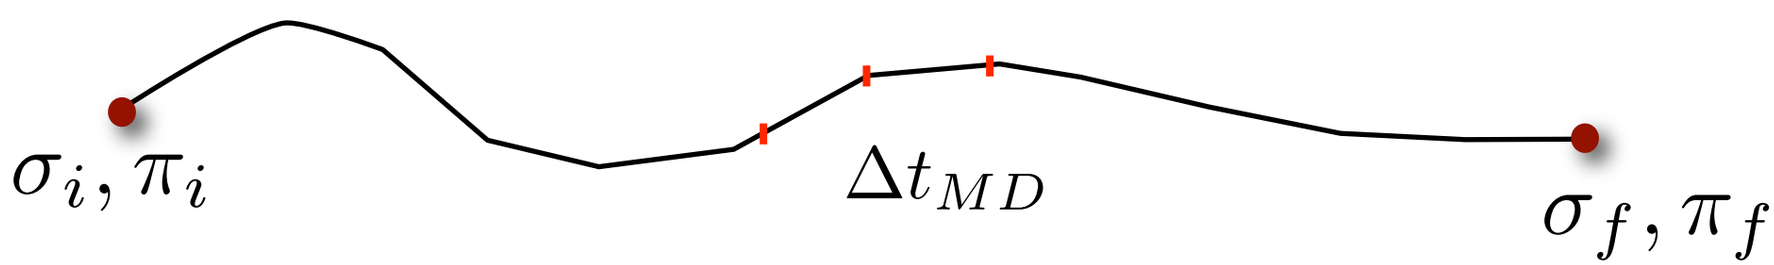

A few words about HMC...

- Pseudofermions

$$\mathcal{Z} = \int \mathcal{D}\sigma \ \mathrm{Det}\left[M^T M(\sigma)\right] = \int \mathcal{D}\sigma \mathcal{D}\varphi^\dagger \mathcal{D}\varphi \ \exp\left(-\varphi^\dagger \left[M^T M(\sigma)\right]^{-1} \varphi\right)$$

  - φ is the **pseudofermion** field
  - M is sparse, of size $(VN_\tau) \times (VN_\tau)$

- Molecular dynamics

$$\mathcal{Z} = \int \mathcal{D}\pi \mathcal{D}\sigma \mathcal{D}\varphi^\dagger \mathcal{D}\varphi \ \exp\left(\underbrace{-\pi^2/2 - \varphi^\dagger \left[M^T M(\sigma)\right]^{-1} \varphi}_{-\mathcal{H}}\right)$$

  - Enables **global updates!**
  - But requires **frequent linear solves!**

$$\dot{\sigma} = \frac{\delta \mathcal{H}}{\delta \pi} = \pi$$

$$\dot{\pi} = -\frac{\delta \mathcal{H}}{\delta \sigma} = F[\sigma, \varphi]$$

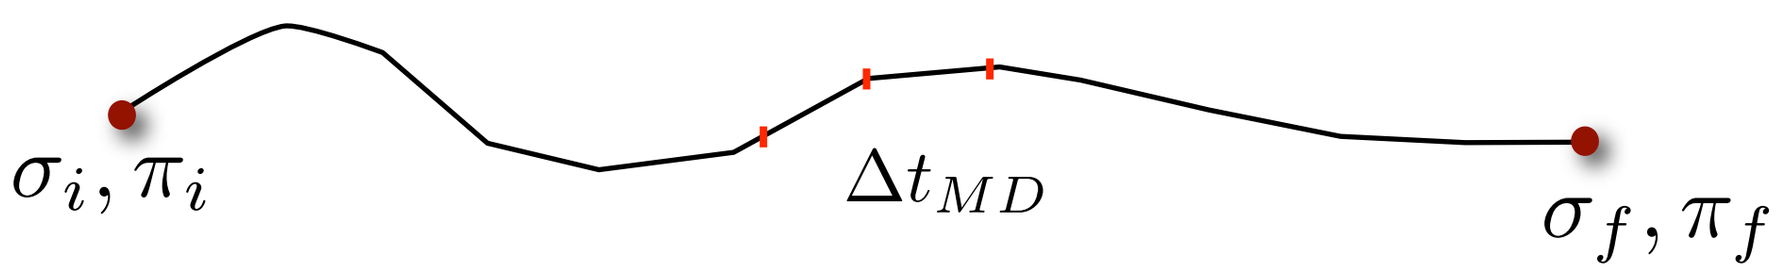

First scaling tests

- Determinantal MC ~ V^3

  $(V = N_x^3)$

- Determinantal MC with "worm" updates ~ V^2

- Same picture for different N_τ but different constant in front.

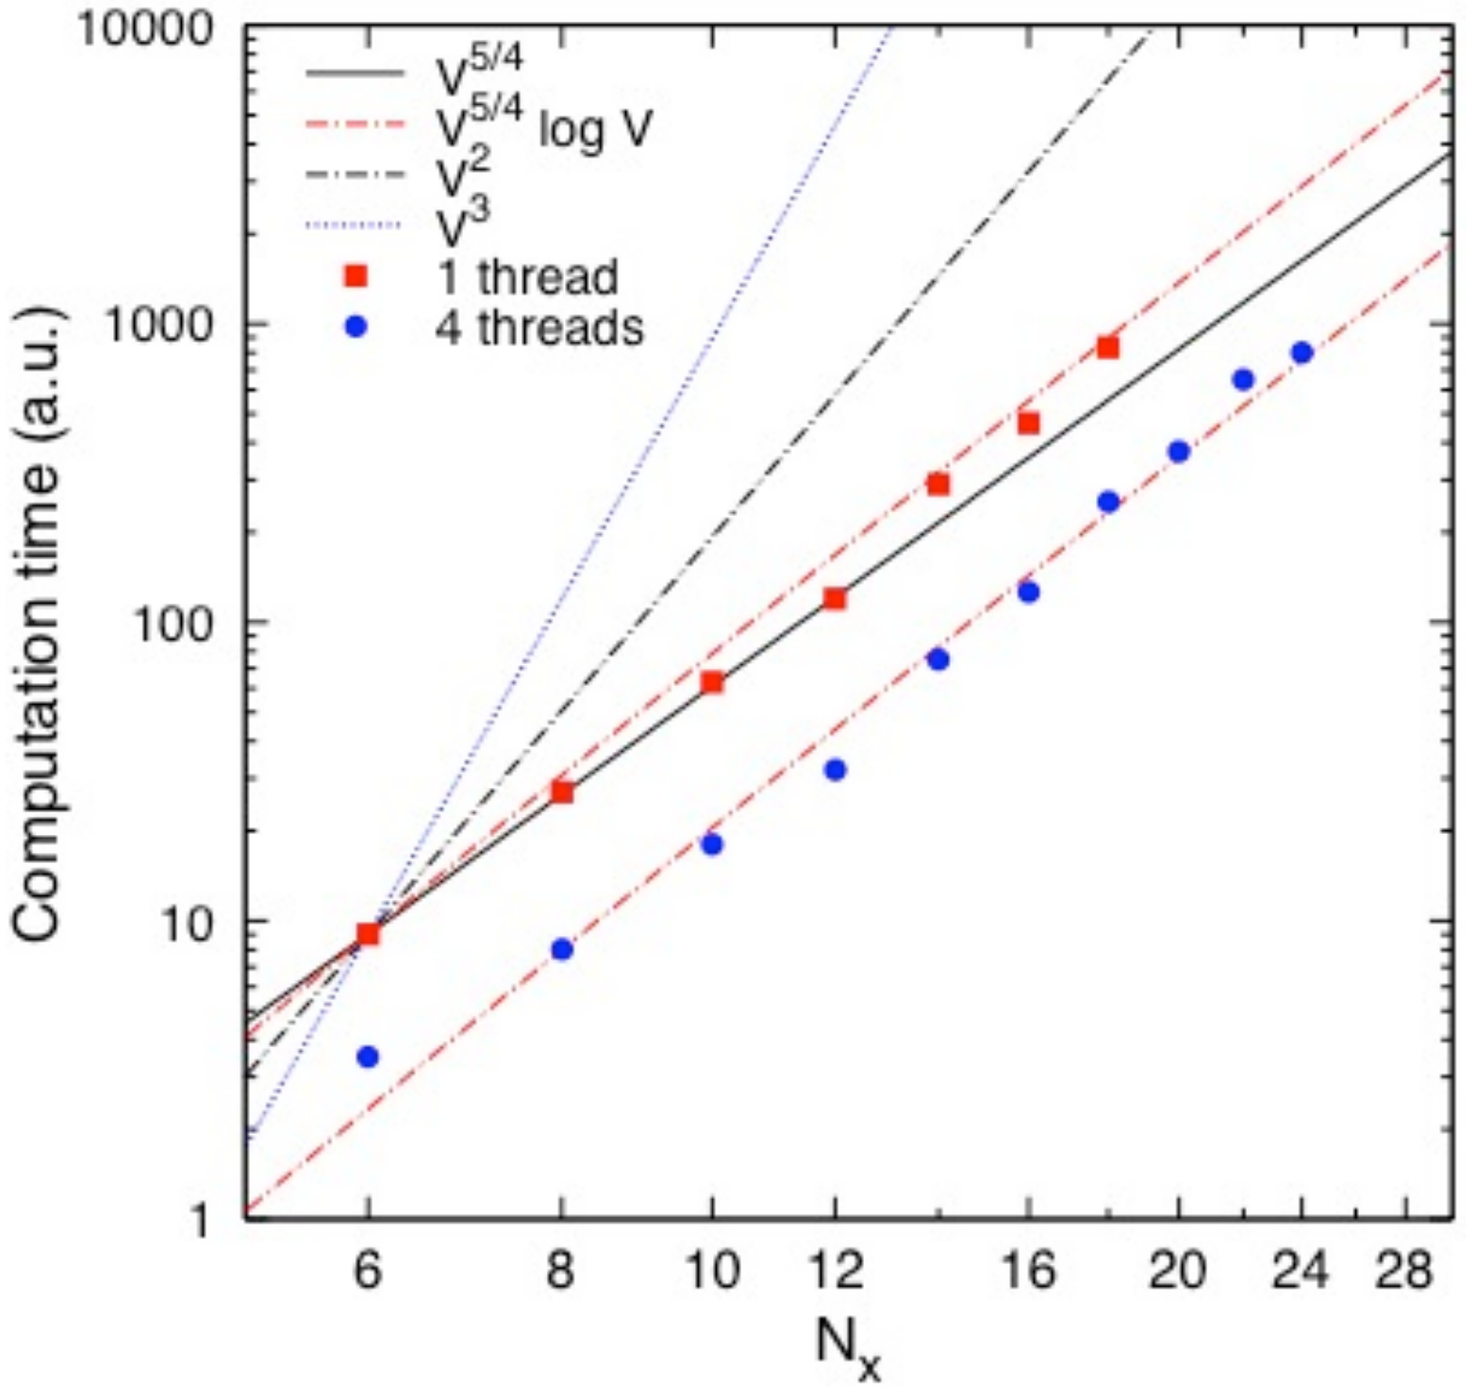

First correctness tests

- Small differences remain, due to <u>finite time-step effects</u>

  Continuous time and HMC ?

- High T is under control

- Low T is more <u>challenging</u> but <u>doable</u>

  “Berlin wall” problem

  You can still run the calculation but it slows down dramatically as T is lowered due to **ill-conditioned** matrix inversion

  Polynomial HMC

  Use polynomials to filter the problematic modes and treat those differently...

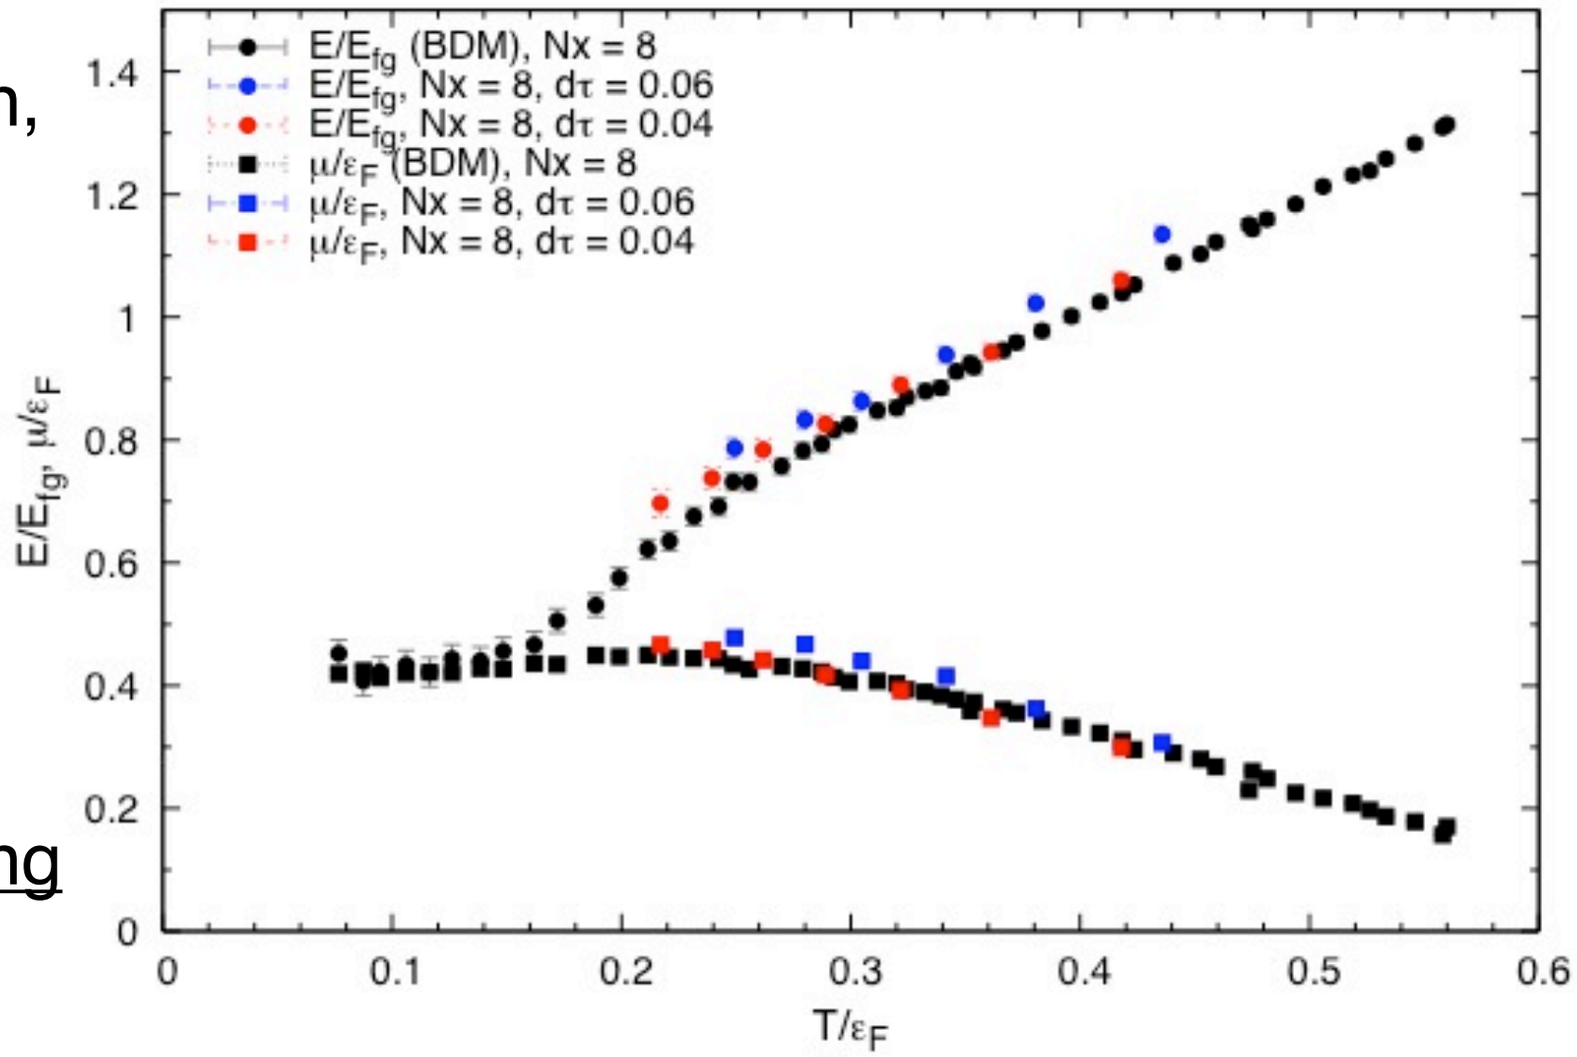

Beyond conventional HMC

- HMC requires the inversion of an **ill-conditioned** matrix...

  ➡ **Preconditioning is essential!**

Beyond conventional HMC

- HMC requires the inversion of an **ill-conditioned** matrix...

→ **Preconditioning is essential!**

If you have an **approximation to the inverse**, you can help the calculation converge faster (or converge at all!)

The **range of the eigenvalues** can be larger than the machine's precision

Beyond conventional HMC

- HMC requires the inversion of an **ill-conditioned** matrix...

➡ **Preconditioning is essential!**

- Strong coupling expansion
- Chebyshev polynomials
- Domain decomposition
- ...

Spectrum of M^TM and P M^TM

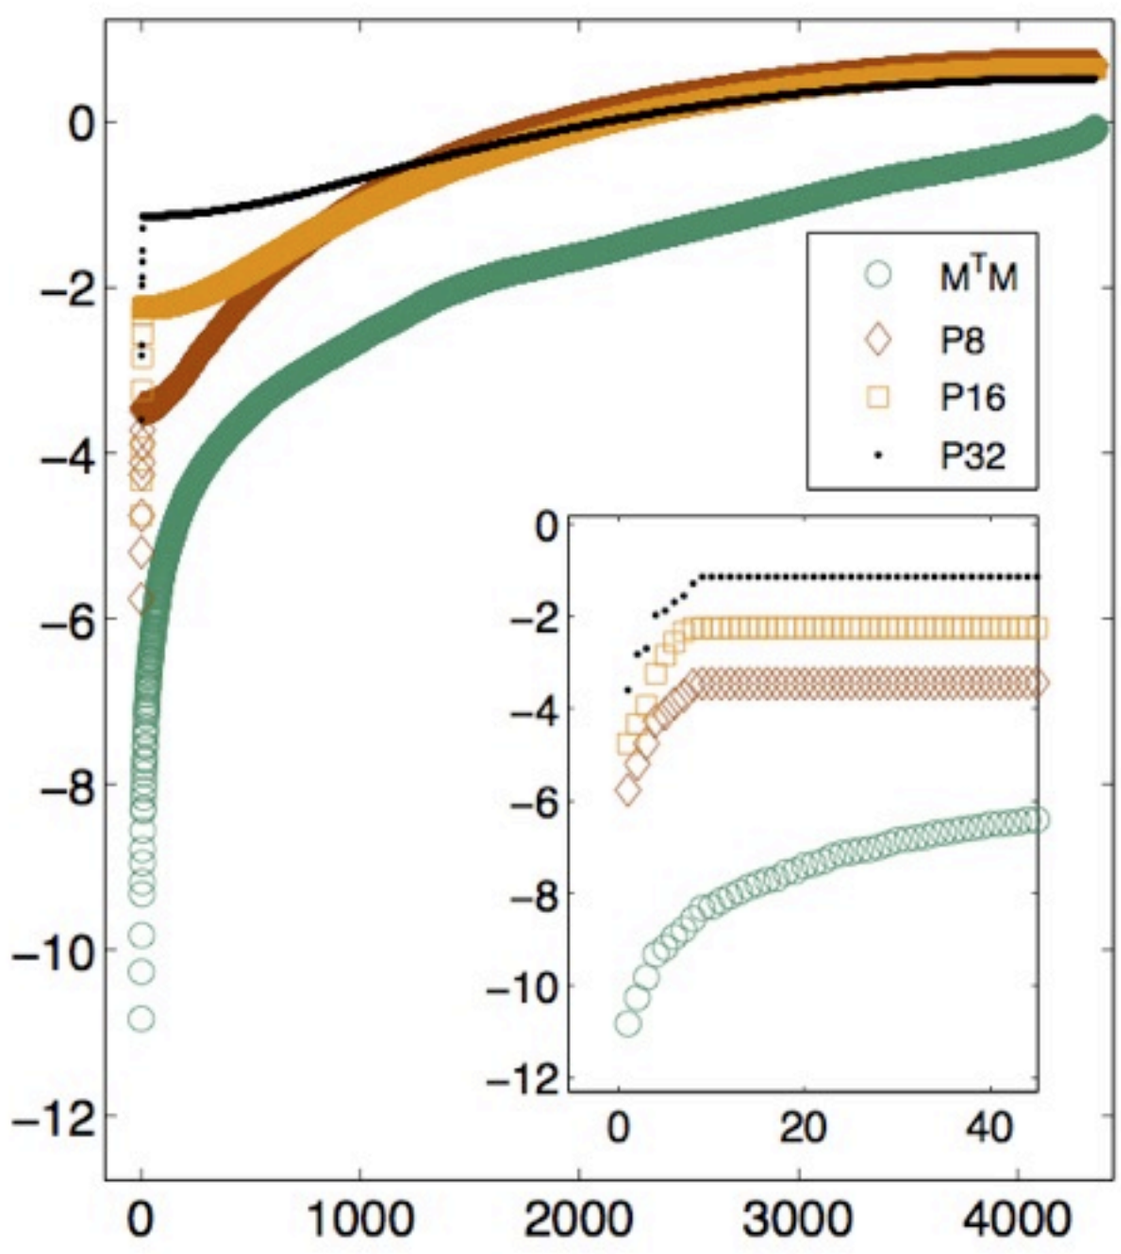

Beyond conventional HMC

- HMC requires the inversion of an **ill-conditioned** matrix...

➡ **Preconditioning is essential!**

- Strong coupling expansion
- Chebyshev polynomials
- Domain decomposition
- ...

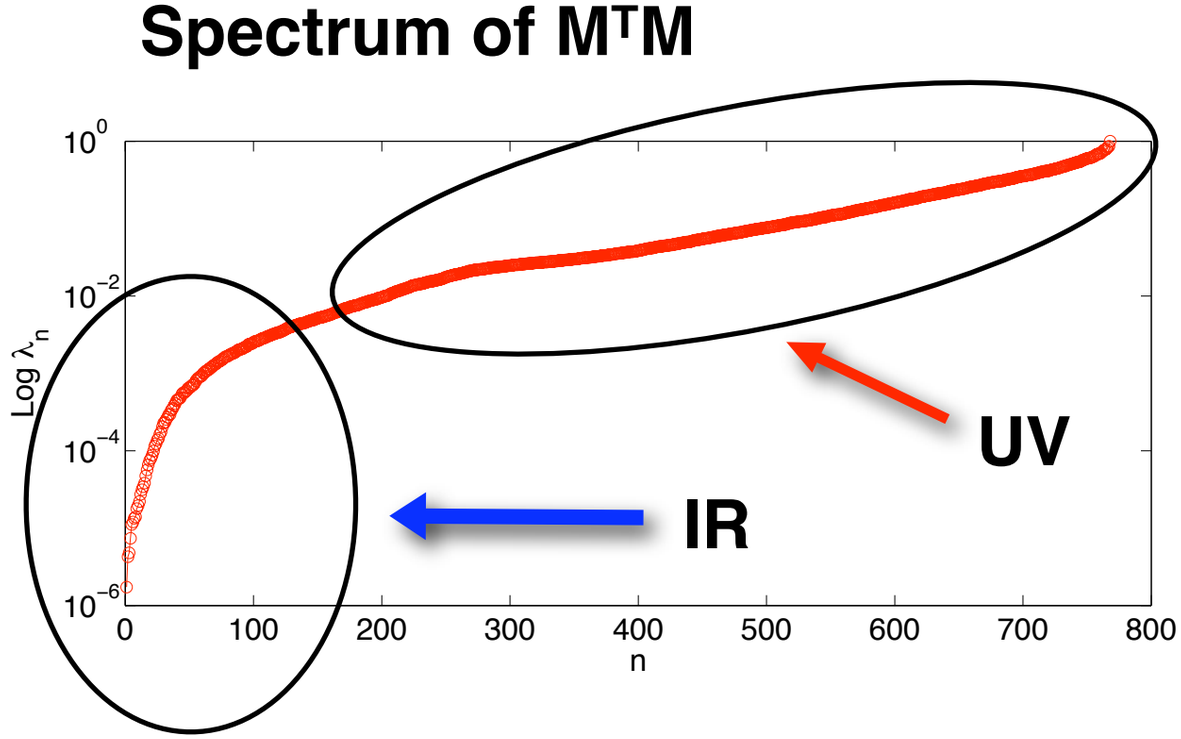

Beyond conventional HMC

- HMC requires the inversion of an **ill-conditioned** matrix...

➡ **Preconditioning is essential!**

- Strong coupling expansion
- Chebyshev polynomials
- Domain decomposition
- ...

Spectrum of M^TM

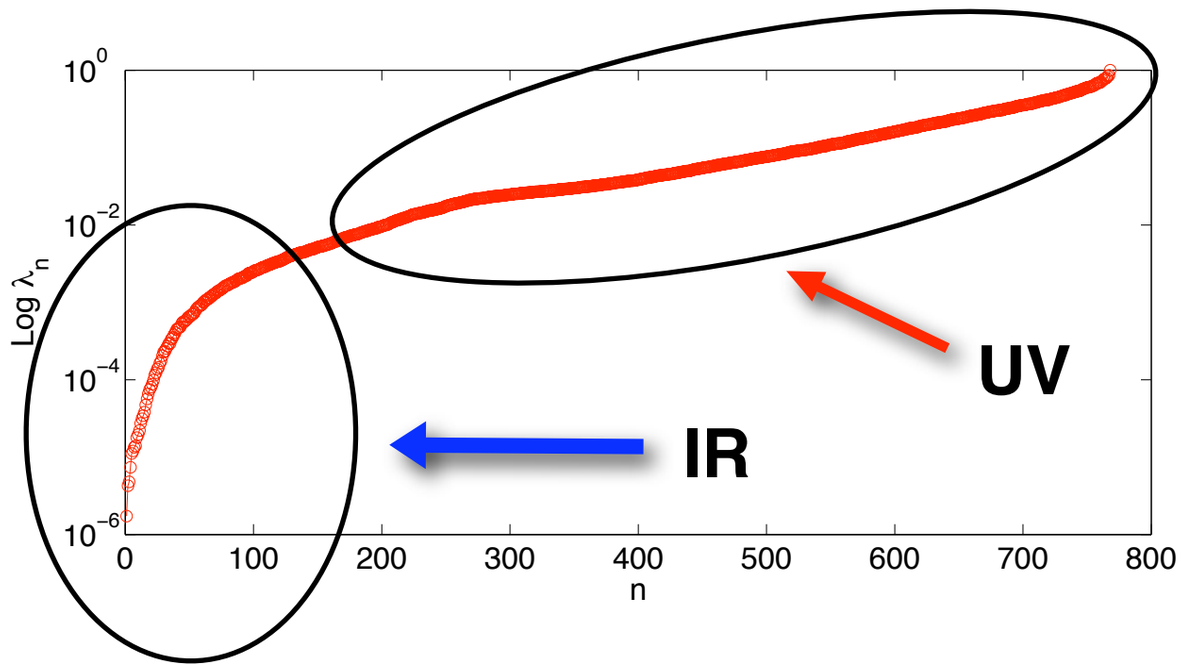

➡ **Multiscale MD integration!**

$F_{\mathrm{IR}} \times \Delta t_{MD}$

Small and expensive
Seldom computed!

$F_{\mathrm{UV}} \times {}^{\Delta t_{MD}}$

Large and cheap
Often computed

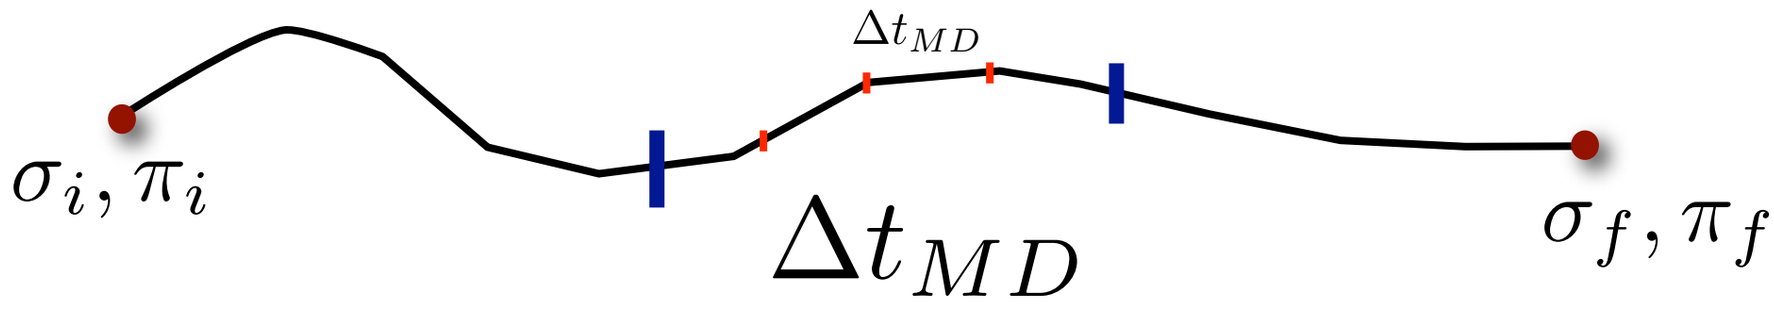

A few words about big computers...

IBM Blue Gene/L

IBM Roadrunner at LANL

- Performance: 5 - 1000 TFlops
- Power consumption: ~ 10^6 Watt
- Cost: as much as 100 M$

Cray XT5 Jaguar at ORNL

Graphics Processing Units (GPUs)

- GPUs are massively multithreaded many-core chips
  - nVidia Tesla c1060 has 240 scalar processors!
  - Can sustain over 12,000 threads concurrently
  - Over 900 GFlops (SP) of processing performance!
  - Cost: $500 ！ (nVidia’s special offer)

nVidia Tesla c1060

CPUs vs. GPUs

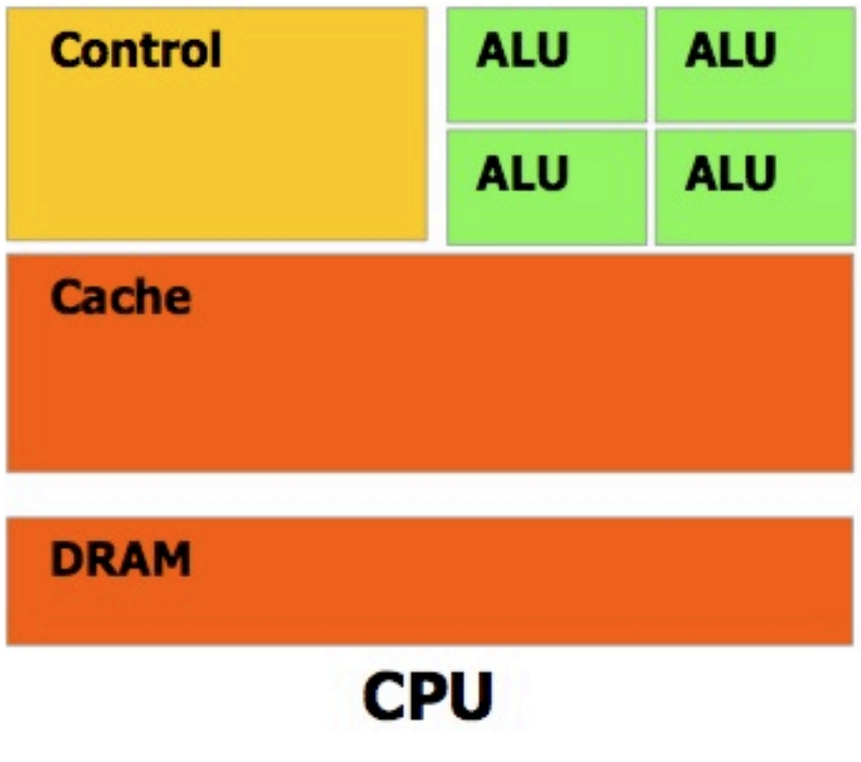

CPU

- Lots of memory per core
- Can handle some limited threading (both heavy and light)
- Easy to program

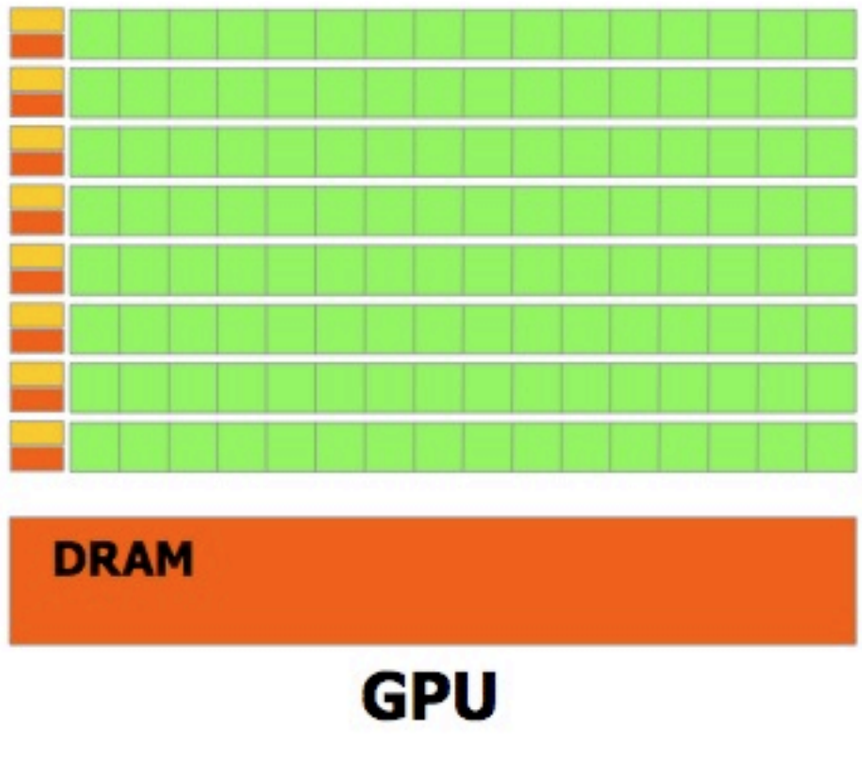

GPU

- Little memory per core
- Best if used with thousands of light threads
- Not-so-easy to program

  but it's getting easier!
  See e.g. http://www.nvidia.com/cuda

CPUs vs. GPUs

- First results with nikola

nikola

1+1 dimensions

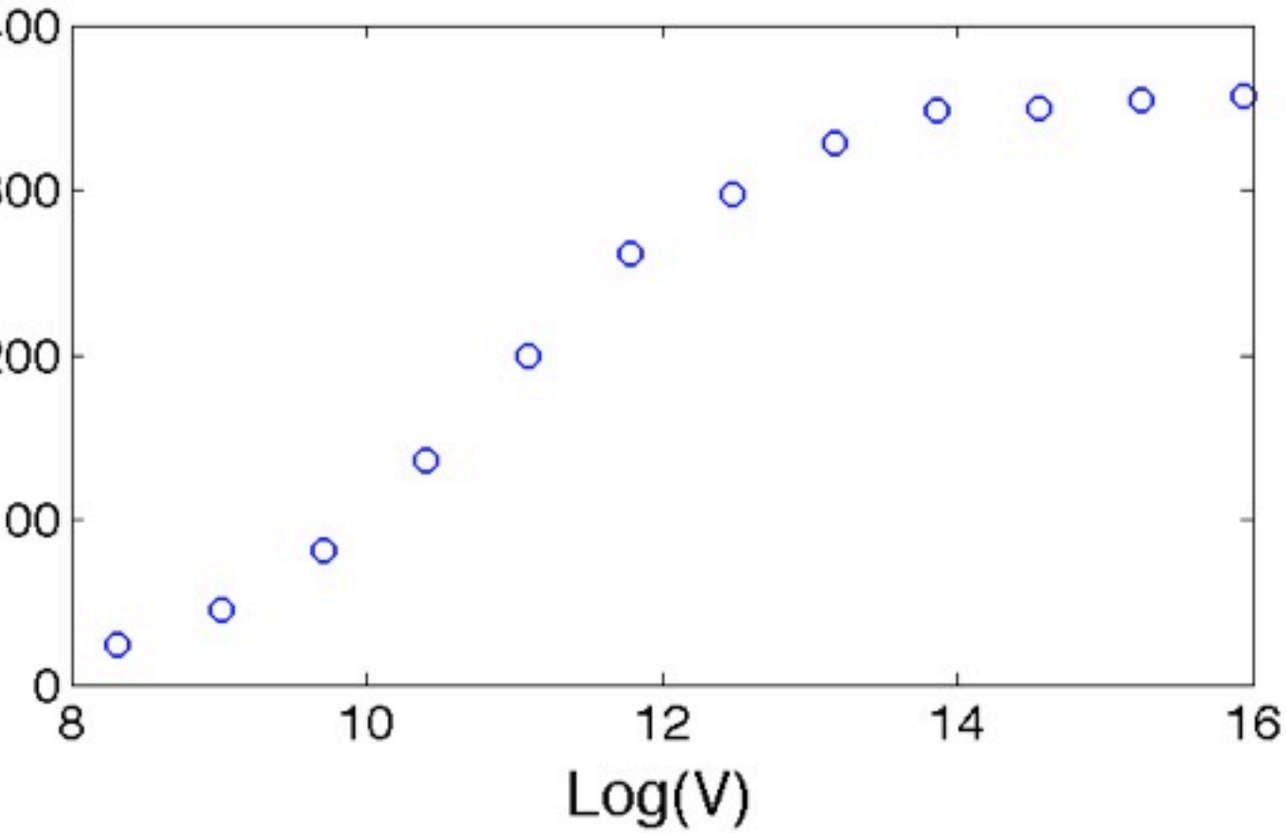

K. A. Wendt, J. E. Drut, T. A. Lähde,
in preparation.

→ ~ 30-350x speedup

CPUs vs. GPUs

- First results with nikola

nikola

3+1 dimensions

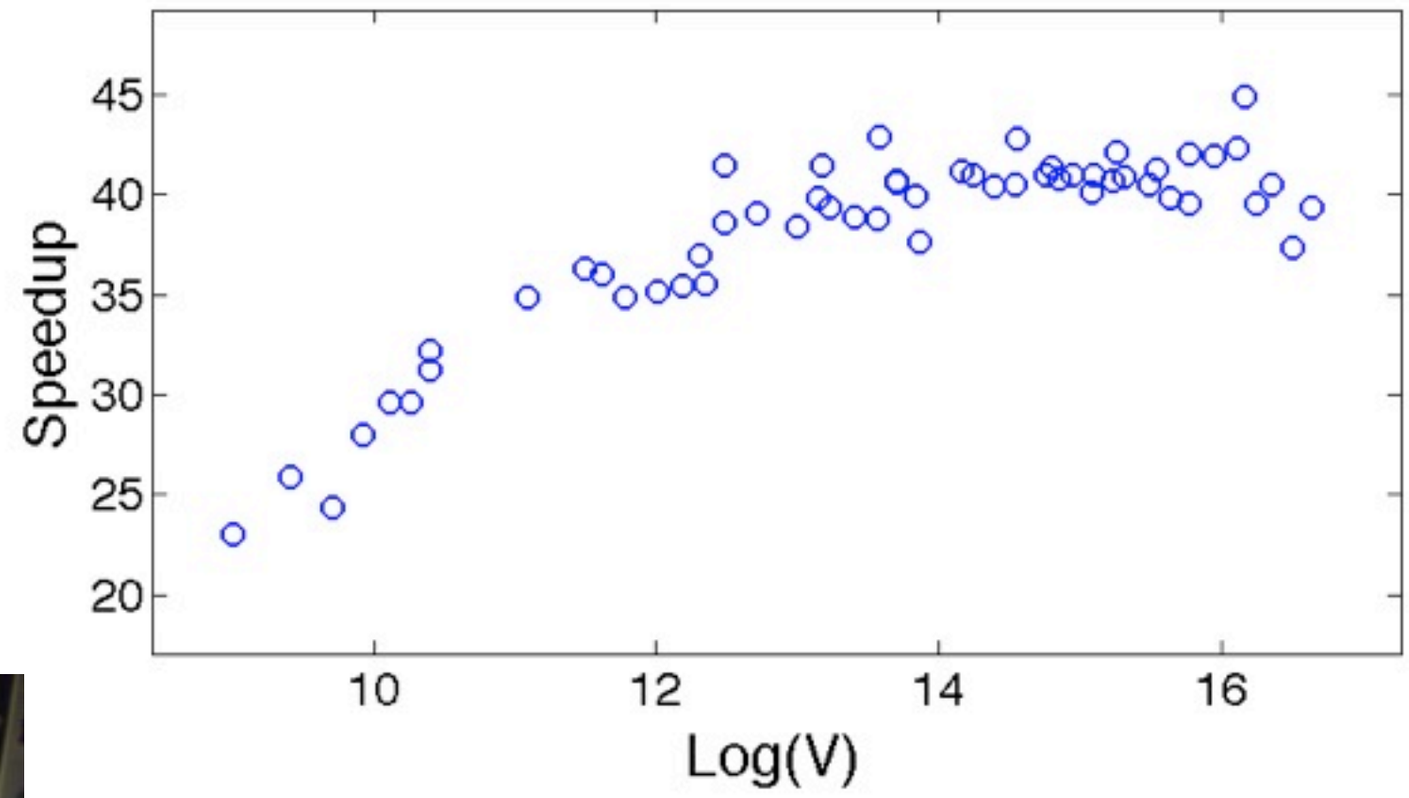

K. A. Wendt, J. E. Drut, T. A. Lähde, in preparation.

→ ~ 20x-40x speedup

enrico

- 2 Tesla c1060 (240 processors each)

  2 Tesla c2050 (upcoming, 512 processors each)

  code-named “Fermi” 1040 GFlops (SP)
  520 GFlops (DP)

Supported in part (mostly) by the Waldemar von Frenckell Foundation (Finland)

Summary

- We know quite a bit about the **equilibrium properties** of Fermi gases at and around unitarity.
  - Conventional algorithms are **simple**, but **scale badly** with system size
  - **Larger system sizes** → **New algorithms!**
  - **GPUs** provide an **efficient** and **inexpensive** way to perform these calculations.
- Our main goal: **universal transport coefficients**

The road is long...

- We know quite a bit about the **equilibrium properties** of Fermi gases at and around unitarity.

- Our main goal: **universal transport coefficients**

So... what else will you do with all this?

- Superfluid density → needed as input for Gross-Pitaevskii calculations
- Static thermal response → compressibility, specific heat, etc.
- Dynamic thermal response → e.g. structure factor

- Neutron matter → Finite range needed

↓

Continuous space HMC?

- The formalism is much more user-friendly in the space-time formulation (used in HMC), in that it yields simpler expressions for the observables than the purely spatial formulation (used in DMC).

To be continued...

Early timing tests

- Determinantal MC vs. Hybrid MC on a single processor

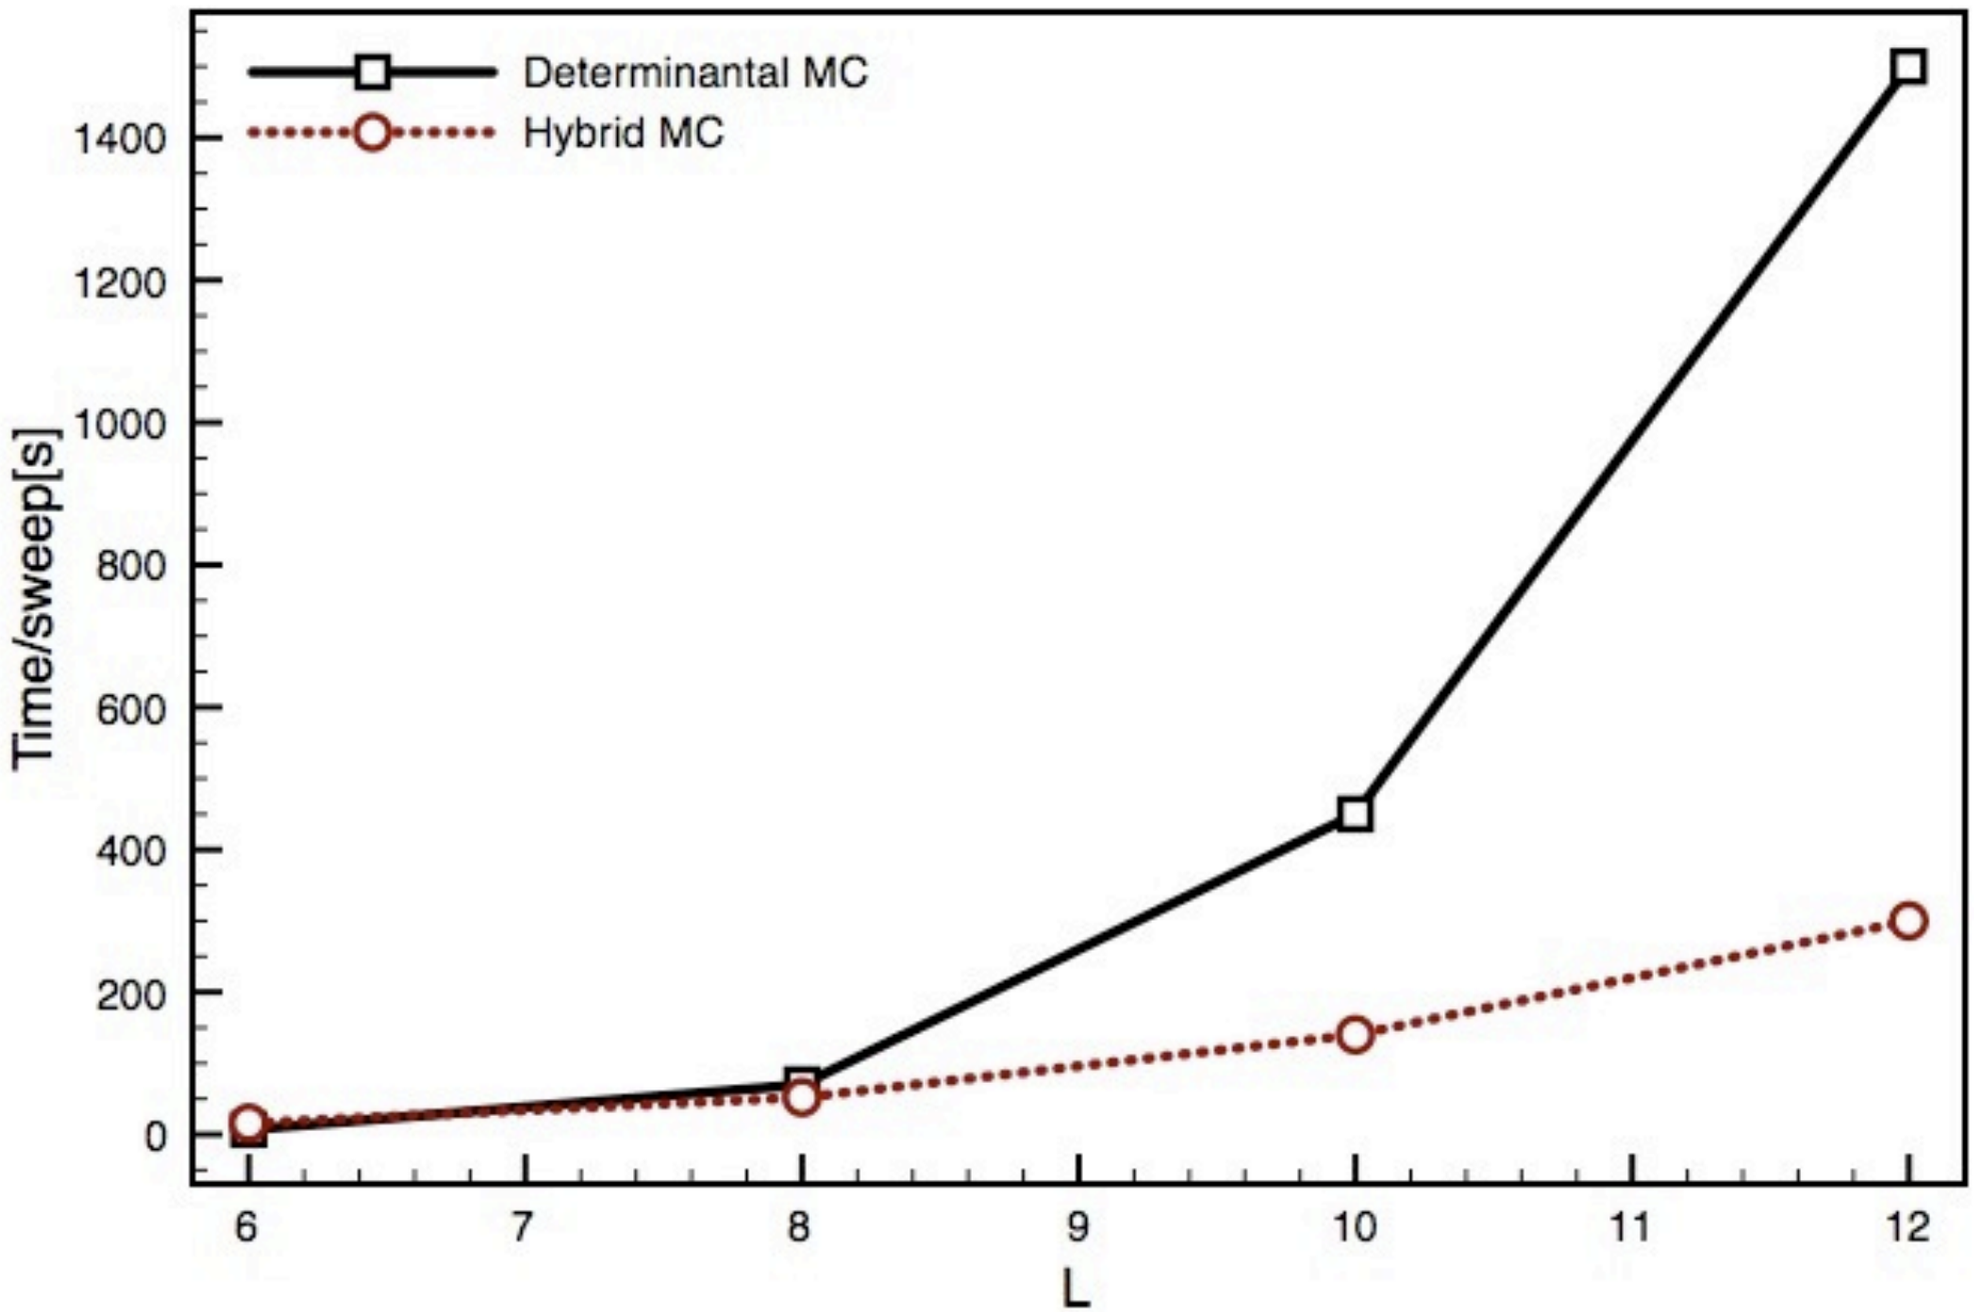

Lattice size