HQS@HPC: Investigation of Stripe Formation in Hubbard Ladders Using Parallel DMRG

G. Hager	HPC Services, Computing Center Erlangen, Germany
E. Jeckelmann	Theoretical Physics, Univ. Mainz, Germany
H. Fehske	Theoretical Physics, Univ. Greifswald, Germany
G. Wellein	HPC Services, Computing Center Erlangen, Germany

CX HPC

Motivation – Microscopic Models

- **Microscopic Hamiltonians in second quantization**
  e.g. Hubbard model

$$H = -t \sum_{\langle ij \rangle, \sigma} \left[c^\dagger_{i\sigma} c_{j\sigma} + \mathrm{H.c.} \right] + U \sum_i n_{i\uparrow} n_{i\downarrow}$$

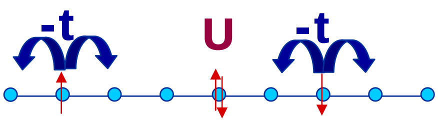

e.g. Holstein-Hubbard model (HHM)

$$H = -t \sum_{\langle ij \rangle, \sigma} \left[c^\dagger_{i\sigma} c_{j\sigma} + \mathrm{H.c.} \right] + U \sum_i n_{i\uparrow} n_{i\downarrow} + g\omega_0 \sum_{i,\sigma} (b_i^\dagger + b_i) n_{i\sigma} + \omega_0 \sum_i b_i^\dagger b_i$$

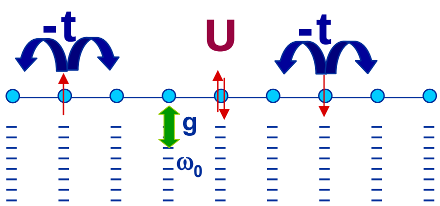

Hilbert space / #quantum states growth exponentially

HHM using an N-site lattice: 4^N * $(M+1)^N$ (N~10-100; M~10)

Electrons   Phonons: Max. M per Site

Motivation – Numerical Approaches to Ground-State Properties

Traditional Approaches

- **Quantum Monte Carlo (QMC)**
- **Exact Diagonalizaton (ED):** Massively Parallel Codes on Supercomputers

New Approach

- **Density Matrix Renormalization Group (DMRG) Method**
  - **Originally introduced by White in 1992**
  - **Large sequential C++ package is in wide use (quantum physics and quantum chemistry)**
  - **Elapsed Times: hours to weeks with desktop CPUs**
  - **No parallel implementation available to date**

DMRG Algorithm

- **Basic Idea: Find an appropriate (reduced) basis set describing the ground-state of *H* with high accuracy**
- **Basic Quantities:**
  - **Superblock = system & environment**

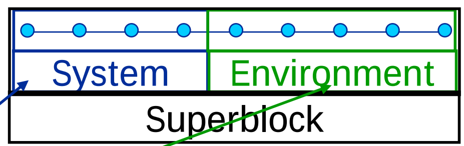

  - **Superblock state (product of system & environment states)**

$$|\psi\rangle = \sum_{ij} \psi_{ij} \, |i\rangle \, |j\rangle$$

  - **Reduced density matrix (DM): summation over environment states**

$$\rho_{ii'} = \sum_{j} \psi^*_{ij} \psi_{i'j}$$

- **Eigenstates of DM with largest eigenvalues have most impact on observables!**

DMRG Algorithm

DMRG algorithm (finite size; left to right sweep)

1. **Diagonalize the reduced DM for a system block of size *l* and extract the *m* eigenvectors with largest eigenvalue**
2. **Construct all relevant operators (system block & environment,…) for a system block of size *l+1* in the reduced density matrix eigenbasis**
3. **Form a superblock Hamiltonian from system & environment Hamiltonians plus two single sites**
4. **Diagonalize the new superblock Hamiltonian**

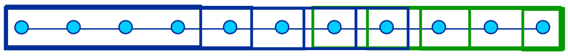

Accuracy depends mainly on *m* (*m*~100 – 10000)

DMRG Algorithm

Implementation

- **Start-Up with infinite-size algorithm**
- **DM diagonalization: LAPACK (dsyev) costs about 5 %**
- **Superblock diagonalization costs about 90 % (Davidson algorithm)**
- **Most time-consuming step: Sparse matrix-vector multiply (MVM) in Davidson (costs about 85 %)**
- **Sparse matrix *H* is constructed by the transformations of each operator in *H*:**

$$H_{ij;i'j'} = \sum_{\alpha} A^{\alpha}_{ii'} B^{\alpha}_{jj'}$$

Contribution from system block and from environment

DMRG Algorithm: Parallelization

Implementation of sparse MVM

- **Sparse MVM: Sum over dense matrix-matrix multiplies!**

$$\sum_{i'j'} H_{ij;i'j'}\psi_{i'j'} = \sum_{\alpha}\sum_{i'} A^{\alpha}_{ii'} \sum_{j'} B^{\alpha}_{jj'}\psi_{i'j'}$$

- **However *A* and *B* may contain only a few nonzero elements, e.g. if conservation laws (quantum numbers) have to be obeyed**
- **To minimize overhead an additional loop (running over nonzero blocks only) is introduced**

$$H\psi = \sum_{\alpha}\sum_{k} (H\psi)^{\alpha}_{L(k)} = \sum_{\alpha}\sum_{k} A^{\alpha}_{k}\psi_{R(k)}\left[B^{\mathrm{T}}\right]^{\alpha}_{k}$$

DMRG Algorithm: Parallelization

Implementation of sparse MVM - pseudocode

$$H\psi = \sum_{\alpha}\sum_{k} A^{\alpha}_{k}\psi_{R(k)}\left[B^{\mathrm{T}}\right]^{\alpha}_{k}$$

```
// W: wavevector ; R: result
for (α=0; α < number_of_hamiltonian_terms; α++) {

       term = hamiltonian_terms[α];

       for (k=0 ; k < term.number_of_blocks; k++) {

              li = term[k].left_index;
              ri = term[k].right_index;

              temp_matrix = term[k].B.transpose() * W[ri];

              R[li] += term[k].A * temp_matrix;

}}
```

Parallel loop !?

Parallel loop !?

Matrix-Matrix-Multiply (Parallel DGEMM ?!)

Data dependency !

DMRG: OpenMP Parallelization

Implementation of parallel sparse MVM – pseudocode (main loop)

```
// W: wavevector ; R: result
#pragma omp parallel private(mymat, li, ri, myid, ics)
{
       myid = omp_get_thread_num();
       mytmat = mm[myid]; // temp thread local matrix

#pragma omp for
       for (ics=0; ics< icsmax; ics++) {

       li = block_array[ics]->left_index;
       ri = block_array[ics]->right_index;

       mytmat = block_array[ics]->B.transpose() * W[ri];

       omp_set_lock(locks[li]);
       R[li] += block_array[ics]->A * mytmat;
       omp_unset_lock(locks[li]);
       }
}
```

Fused (α,k) loop

Protect each block of result vector **R** with locks

Parallel DMRG: References

- G. Hager, E. Jeckelmann, H. Fehske, and G. Wellein: *Parallelization Strategies for Density Matrix Renormalization Group Algorithms on Shared-Memory Systems.* J. Comp. Phys. 194, 795 (2004)

- G. Hager, E. Jeckelmann, H. Fehske, and G. Wellein: *Exact Numerical Treatment of Finite Quantum Systems using Leading-Edge Supercomputers.* To be published in: Proceedings of the International Conference on High Performance Scientific Computing, March 10-14 2003, Hanoi, Vietnam (Springer)

- **Goals**
  - **Show capabilities and limits of the parallel DMRG approach**
    - **Ground-state properties of 4x4 and 6x6 Hubbard systems at half-filling**
  - **Shed some light on issue of ground-state stripe formation in Hubbard ladders with open BCs in long (leg) direction and closed BCs in short (rung) direction**
    - **Stripes in ground-state hole density of 7x6, 11x6, 14x6 OxP with 4, 6 and 8 holes, respectively**

      Ref.: S. R. White and D. J. Scalapino: *Stripes on a 6-leg Hubbard Ladder.* Phys. Rev. Lett. 91, 136403 (2003)

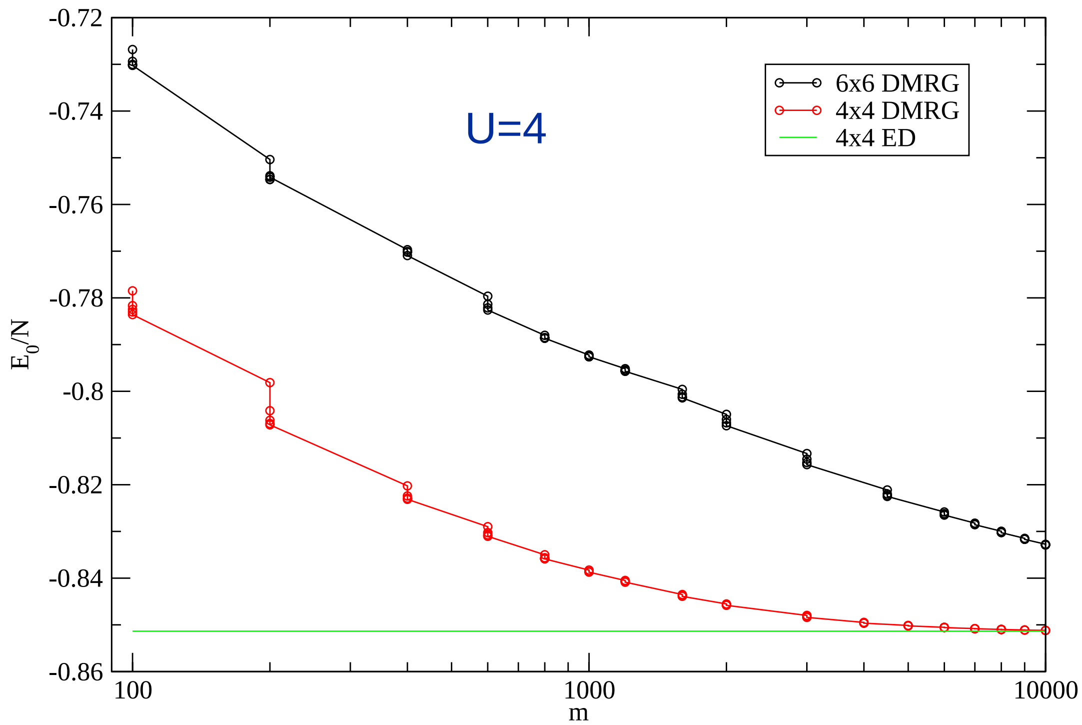

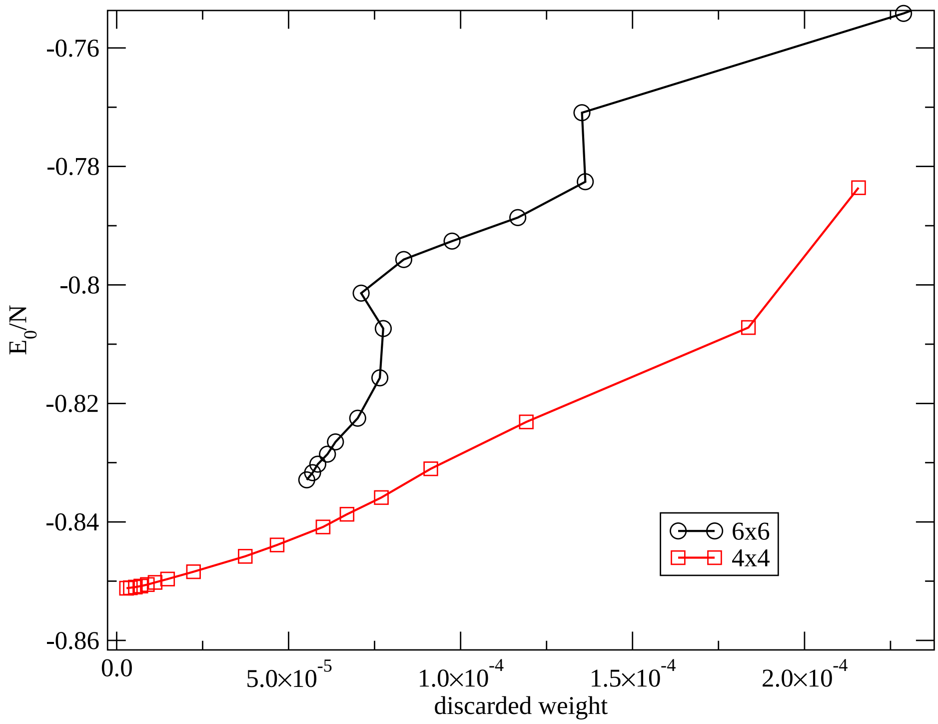
-0.76
-0.78
-0.8
-0.82
-0.84
-0.86
E_0/N
0.0
5.0×10^{-5}
1.0×10^{-4}
1.5×10^{-4}
2.0×10^{-4}
discarded weight
6x6
4x4

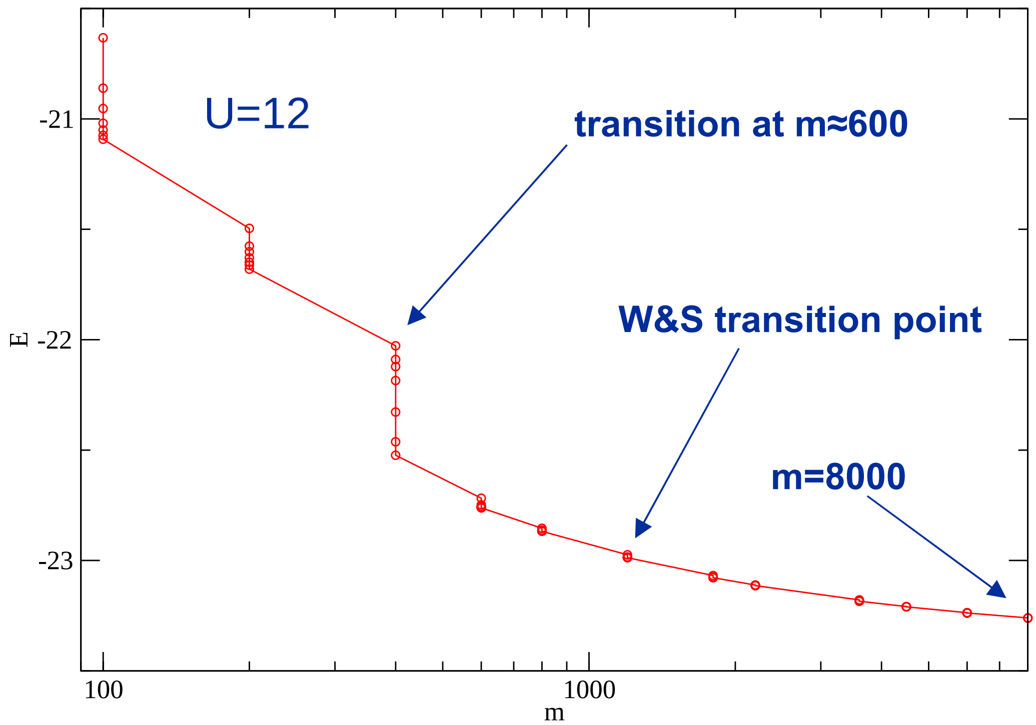
U=12
transition at m≈600
W&S transition point
m=8000
-21
-22
-23
E
100
1000
m

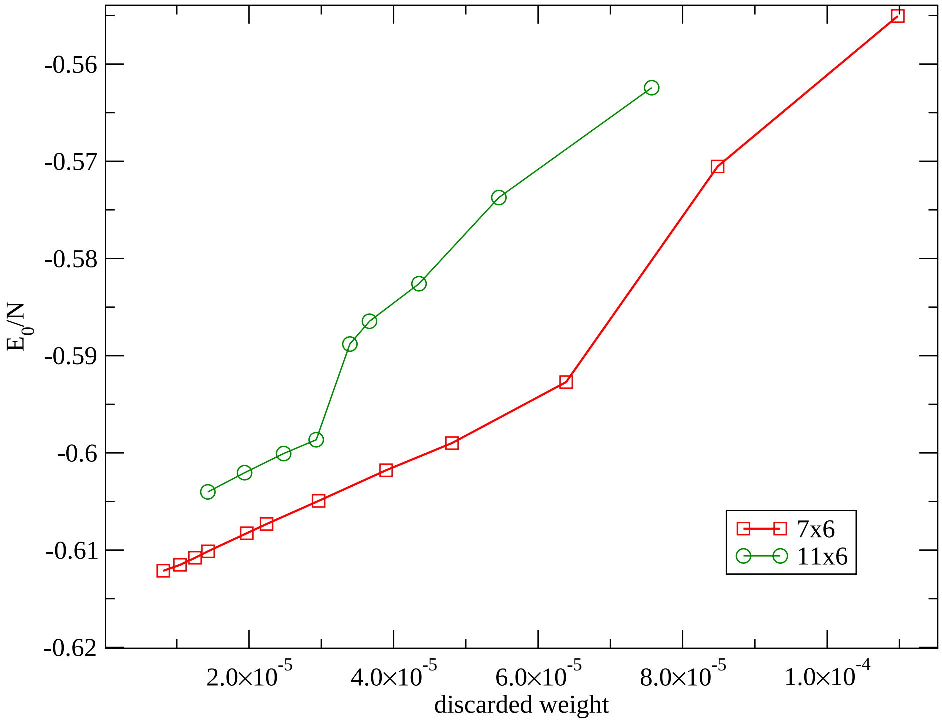
E0/N
-0.56
-0.57
-0.58
-0.59
-0.6
-0.61
-0.62
2.0×10-5
4.0×10-5
6.0×10-5
8.0×10-5
1.0×10-4
discarded weight
7x6
11x6

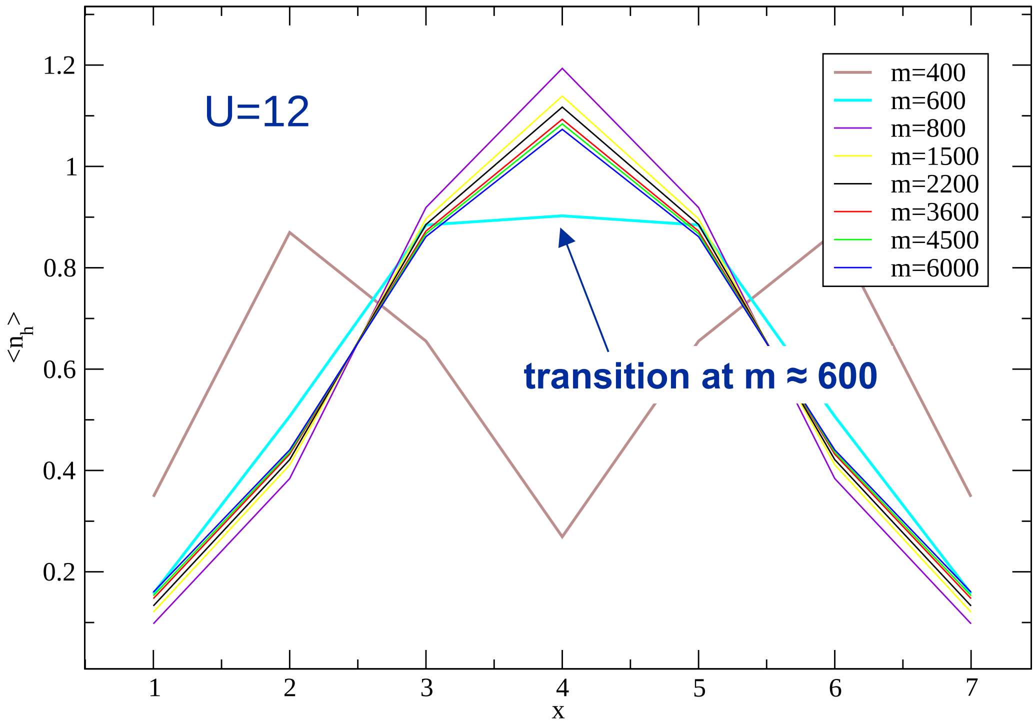
U=12
<nh>
1.2
1
0.8
0.6
0.4
0.2
1
2
3
4
5
6
7
x
m=400
m=600
m=800
m=1500
m=2200
m=3600
m=4500
m=6000
transition at m ≈ 600

7x6 OxP Hubbard ladder

- **Reproduced W&S results for U ∈ [3,12]**
  - **using reflection symmetry**
  - **transition to striped state already at m ≈ 600**
  - **about 6 CPU weeks (Power4 1.3 GHz) for getting to m = 8000**

- **What's next?**
  - **Is stripe formation an artifact from the boundaries?**
  - **Is stripe formation strongly dependent on the length of the ladder?**
  - **Can we answer those questions by going to larger systems (in x direction) in order to shift the boundaries away from interesting structures?**
  - **How far can we get with the current DMRG code?**

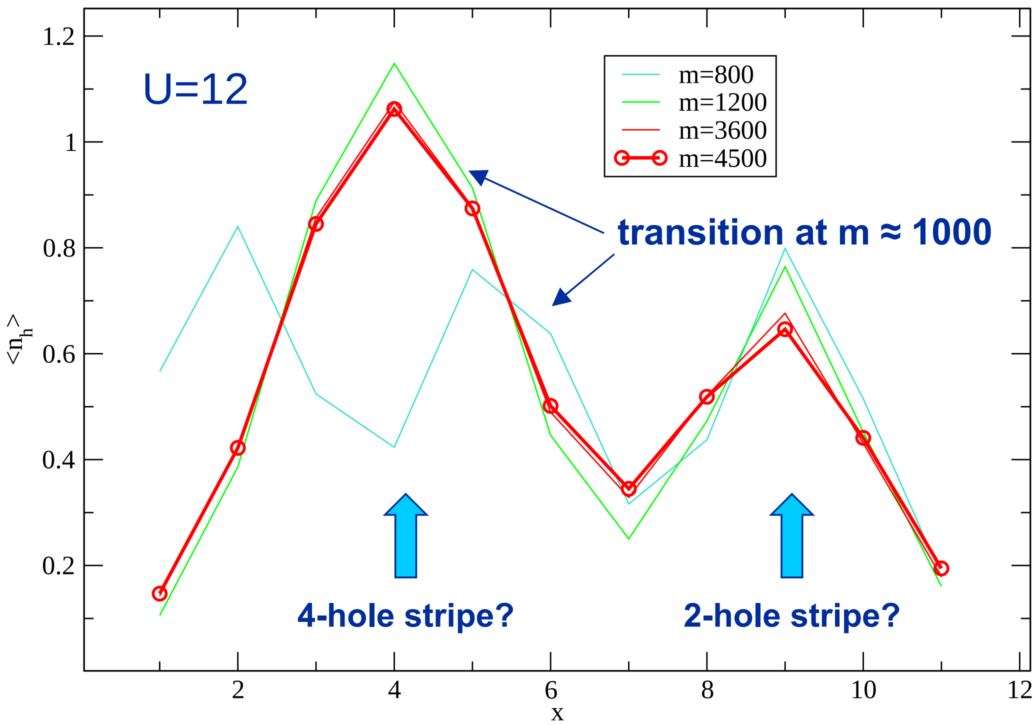

11x6 OxP Hubbard ladder with 6 holes: 2D hole and spin densities

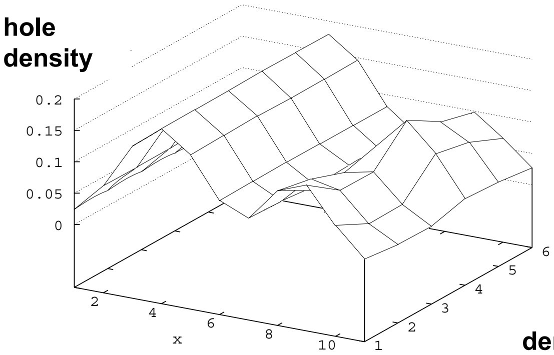

- Hole densities alone are not a reliable signature for stripe formation
- π phase shifts of the magnetization density across the stripe provide additional evidence

Stripe signatures: Zero crossings of magnetization density (antiferromagnetic!)

$$S_z(x,y)(-1)^{x+y}$$

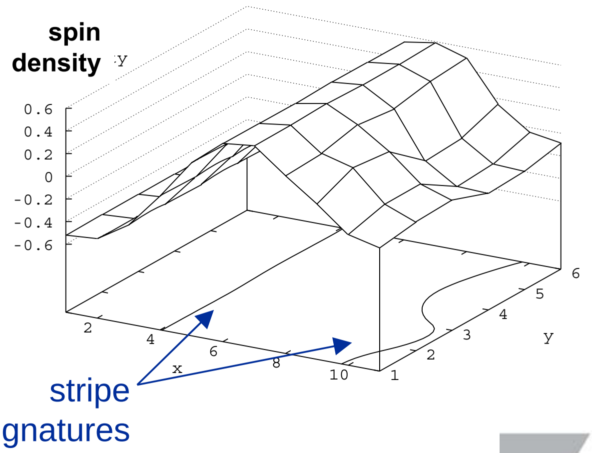

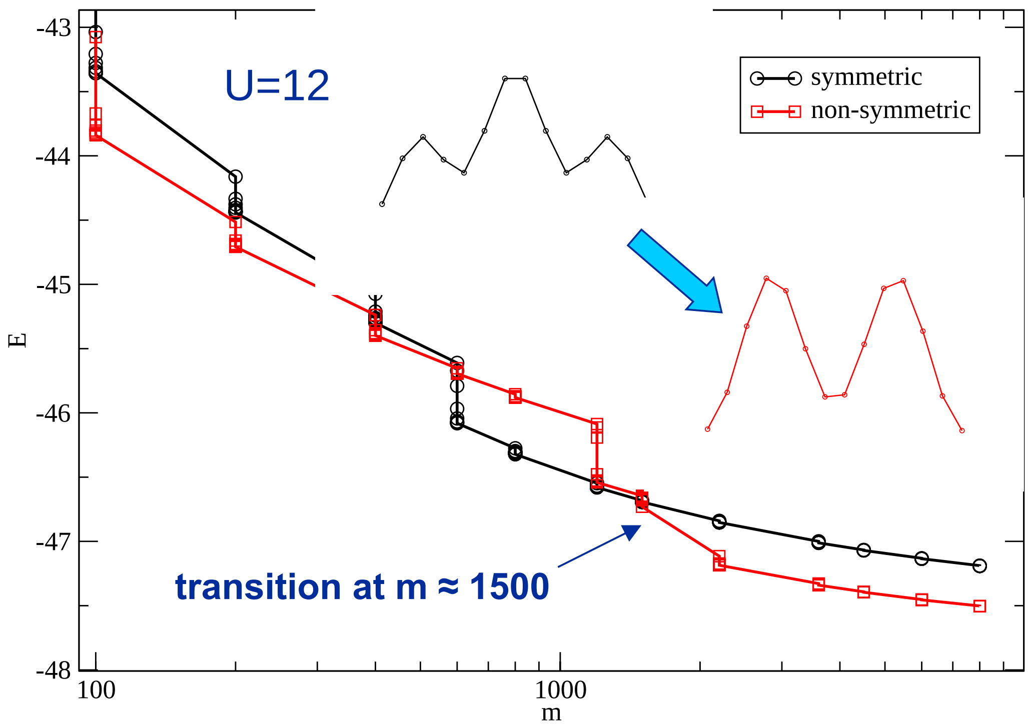

14x6 OxP Hubbard ladder with 8 holes: 2D hole and spin densities

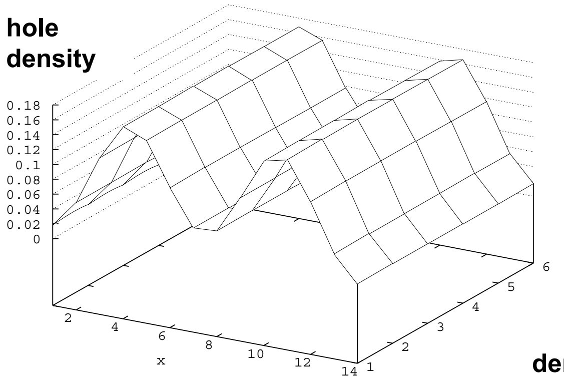

Solid stripe formation!

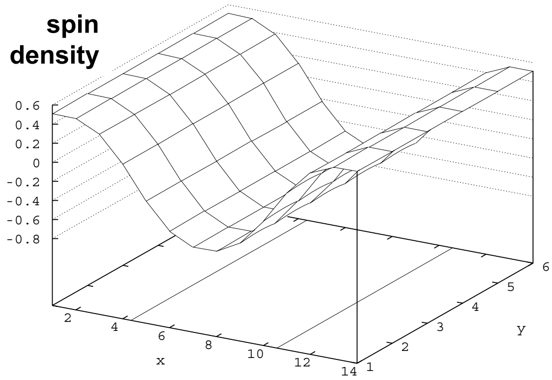

Summary (1)

- **Porting & parallelization of existing DMRG code from quantum physics/chemistry:**
  - **Kernel: sparse Matrix-Vector-Multiply (MVM)**
  - **Fusing inner & outer loop allows a scalable OpenMP implementation for MVM routine with a parallel efficiency of 98% for MVM**
  - **Good fraction of peak performance for whole application on modern SMP nodes**
- **Limits even of parallel DMRG show up clearly with periodic 6x6 Hubbard system at half filling**
  - **Convergence expected at m ≈ 10^5**

Summary (2)

- Reproduced 7x6 results by White & Scalapino
- Manageable system sizes for 6-leg OxP Hubbard ladders pushed as far as 21x6
  - Significantly earlier transition to striped state than W&S
  - Strange behaviour at 11x6, clear stripes at 14x6
  - 21x6 still inconclusive (lack of consistent theory)
- Using reflection symmetry is not always feasible
- Transition point to striped state is located at successively larger m with increasing system size
- Solid signatures for convergence are still difficult to establish in DMRG

Acknowledgement