Graph Partitioning Methods for Fast Parallel Quantum Molecular Dynamics

Hristo Djidjev, Georg Hahn, Sue Mniszewski

Christian Negre, Anders Niklasson, Vivek Sandeshmuk

Ocober 10, 2016

Talk outline

- Background and motivation of partitioning approach
  - Quantum MD background
  - Recursive polynomial expansion of Hamiltonian matrices
  - Partitioned evaluation of matrix polynomials
- Formulation of the GP problem and its application
  - CH-partitioning definition
  - Application to matrix polynomial evaluation
  - Correctness of approach
- Development of CH-partitioning algorithms
- Experimental analysis
- Conclusion

Quantum MD background

- Classical MD simulations
  - Atoms as bodies that move based on Newton's laws of motion
  - Forces between atoms calculated using interatomic potentials
  - Positions of atoms updated in small time steps
  - Interaction models use a priori knowledge of the system
  - Cannot explain events on atomic and subatomic level
- Quantum MD simulations
  - Based on laws of quantum mechanics
  - Density functional theory (DFT) most used model
  - Second-order spectral projection (SP2) approach
    - Density matrix as a function f of the Hamiltonian
    - Representing f as a recursive polynomial expansion

Recursive polynomial matrix expansion

- Given Hamiltonian *H*, compute density matrix *D*

$$D = \lim_{n\to\infty} f_n(f_{n-1}(\dots f_0(H)\dots))$$

$$f_0(X) = \alpha I - \beta X$$

$$f_i(X) = \begin{cases} X^2, & \text{if } Tr[X] > N_i \\ 2X - X^2, & \text{otherwise} \end{cases}$$

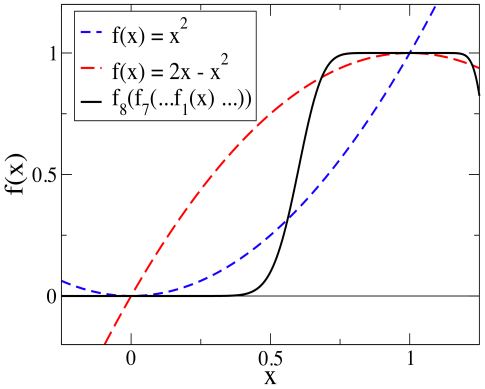

- The degree grows at an exponential rate, hence 20-30 iterations suffice
- Thresholding used to reduce MM complexity

$$D = \lim_{n\to\infty} f_n t_n(\dots f_0 t_0(H)\dots)$$

Parallel evaluation of matrix polynomial for *D*

- Large number of time steps (10^4-10^6) – need parallelism
- Bottleneck operation $Y = X^2$ for a sparse matrix X
- Sparse matrix algebra
  - Works well in sequential and shared-memory environment
  - Speedup of distributed implementation goes down with the # nodes due to communication overhead
- Partitioning based approach
  - Computational overhead (total number of operations higher)
  - Reduced communication overhead
  - Scalable parallelism

Partitioned evaluation

- Model the sparsity structure of H by a graph $G = G(H)$
- Partition G into (overlapping) graphs G_i
  - *core* vertices of $G_1, \ldots, G_p$ form a partition of $V(G)$
  - *halo* vertices are neighbors of core vertices & not in the core
  - CH-partitioning (core-halo)

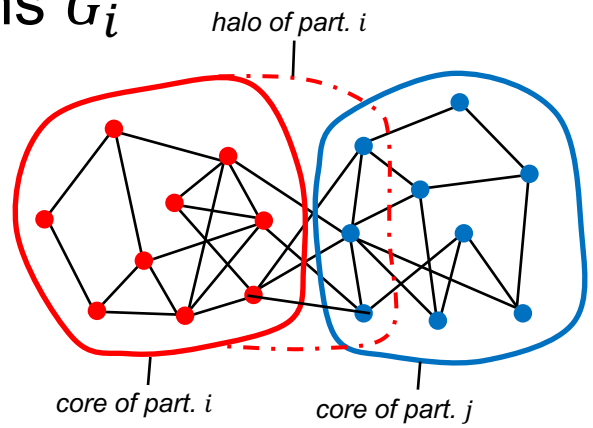

- Send submatrix H_i of H defined by G_i to node i
- Compute polynomial $P(H_i)$ by node i
- Copy core elements of $P(H_i)$ to $D := P(H)$

The CH-partitioning problem

- The partitioned algorithm correctly computes during the i-th iteration $D(H_i)$ assuming
  - Time step is small enough so that density matrix does not change a lot in one iteration
  - Graph used for partitioning is based on $(D_{i-1}+H_i)^2$
  - Thresholding is used after each matrix computation
- CH-partitioning problem formulation:

  Given an undirected graph G and $q \geq 2$, find a partition $C_1, \ldots, C_q$ of $V(G)$ with corr. halos $H_1, \ldots, H_q$ that minimizes

  $$\sum_i (|C_i| + |H_i|)^3 \quad (or, alternatively, \max_i \{|C_i| + |H_i|\}.$$

Partitioning algorithms

- Standard graph partitioning
  - Related, but different than CH-graph partitioning

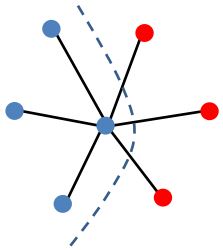

Standard graph partitioning

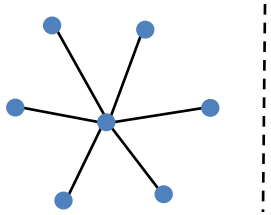

CH graph partitioning

  - Solvers Metis, hMetis, KaHIP
- New algorithms
  - Kernighan-Lin based
  - Simulated annealing
  - Metis+SA

Experimental setup

- Test cases motivated by physical systems

No.	Name	n	m	m/n	Description
1	polyethylene dense crystal	18432	4112189	223.1	crystal molecule in water low threshold
2	polyethylene sparse crystal	18432	812343	44.1	crystal molecule in water high threshold
3	phenyl dendrimer	730	31147	42.7	polyphenylene branched molecule
4	polyalanine 189	31941	1879751	58.9	poly-alanine protein solvated in water
5	peptide 1aft	385	1833	4.76	ribonucleoside-diphosphate reductase protein
6	polyethylene chain 1024	12288	290816	23.7	chain of polymer molecule, almost 1-d
7	polyalanine 289	41185	1827256	44.4	large protein in water solvent
8	peptide trp cage	16863	176300	10.5	small protein dissolved in H_2O molecules
9	urea crystal	3584	109067	30.4	organic compound

Test matrices

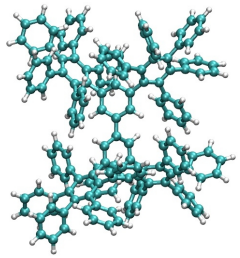

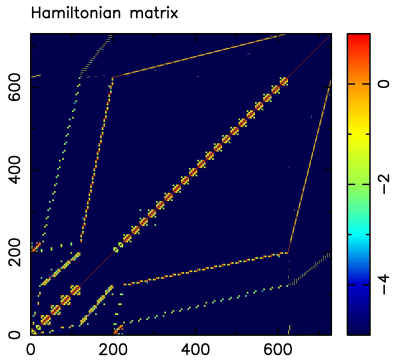

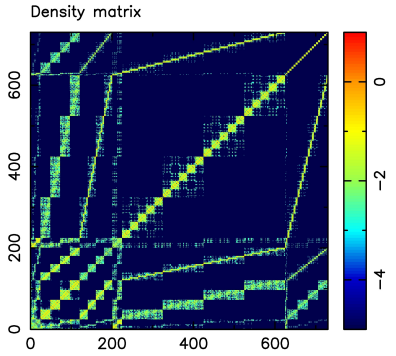

Phenyl dendrimer system with its molecular representation (left)

2D plot representation of the Hamiltonian (middle)

Thresholded density matrix (right)

Comparison of accuracies

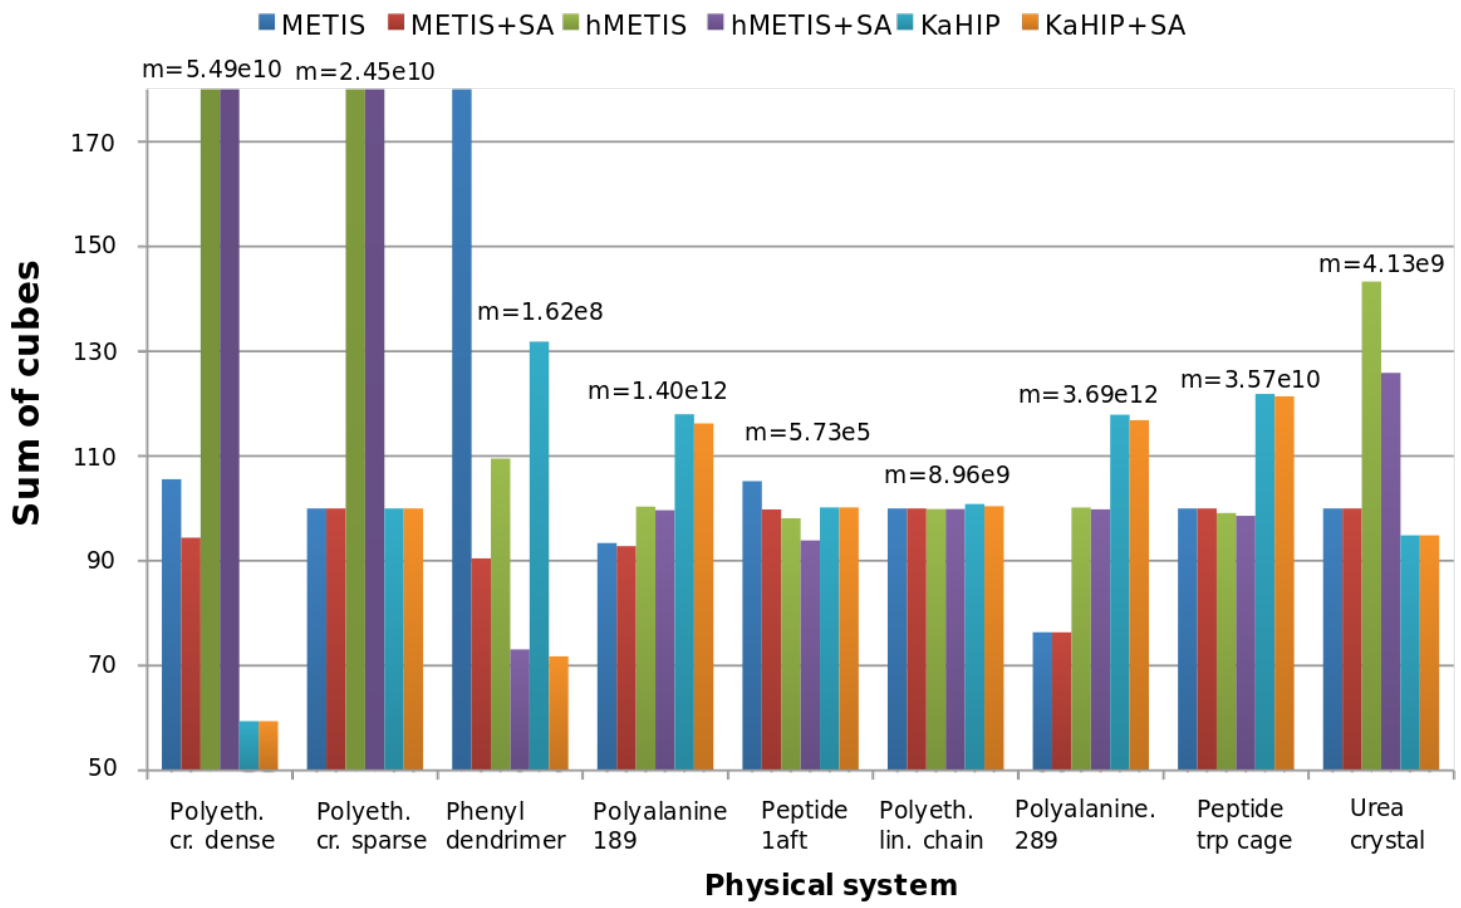

Comaprison of running times

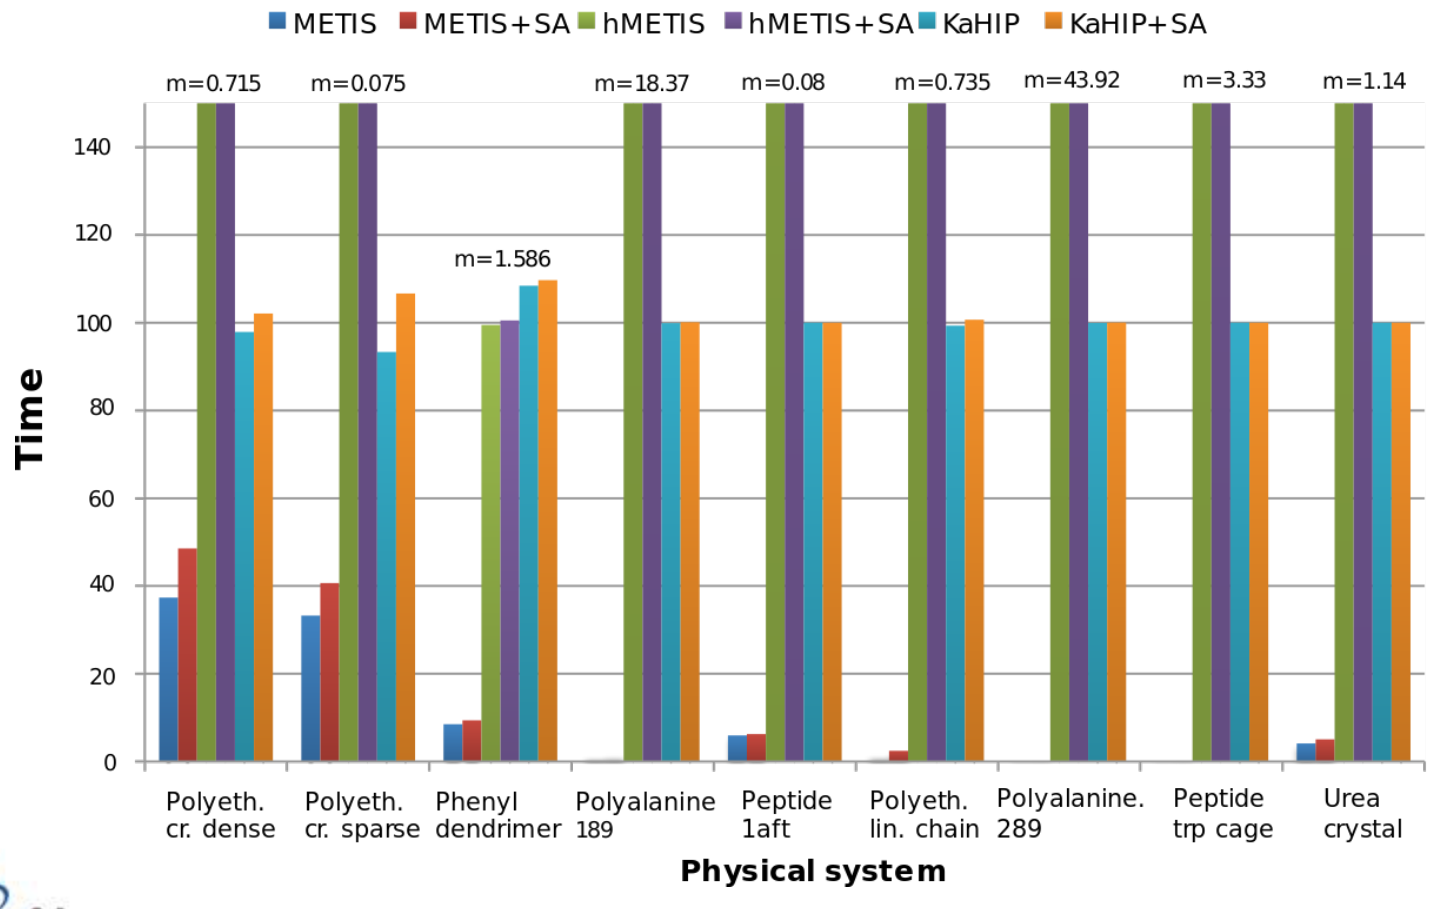

QMD running time comparison

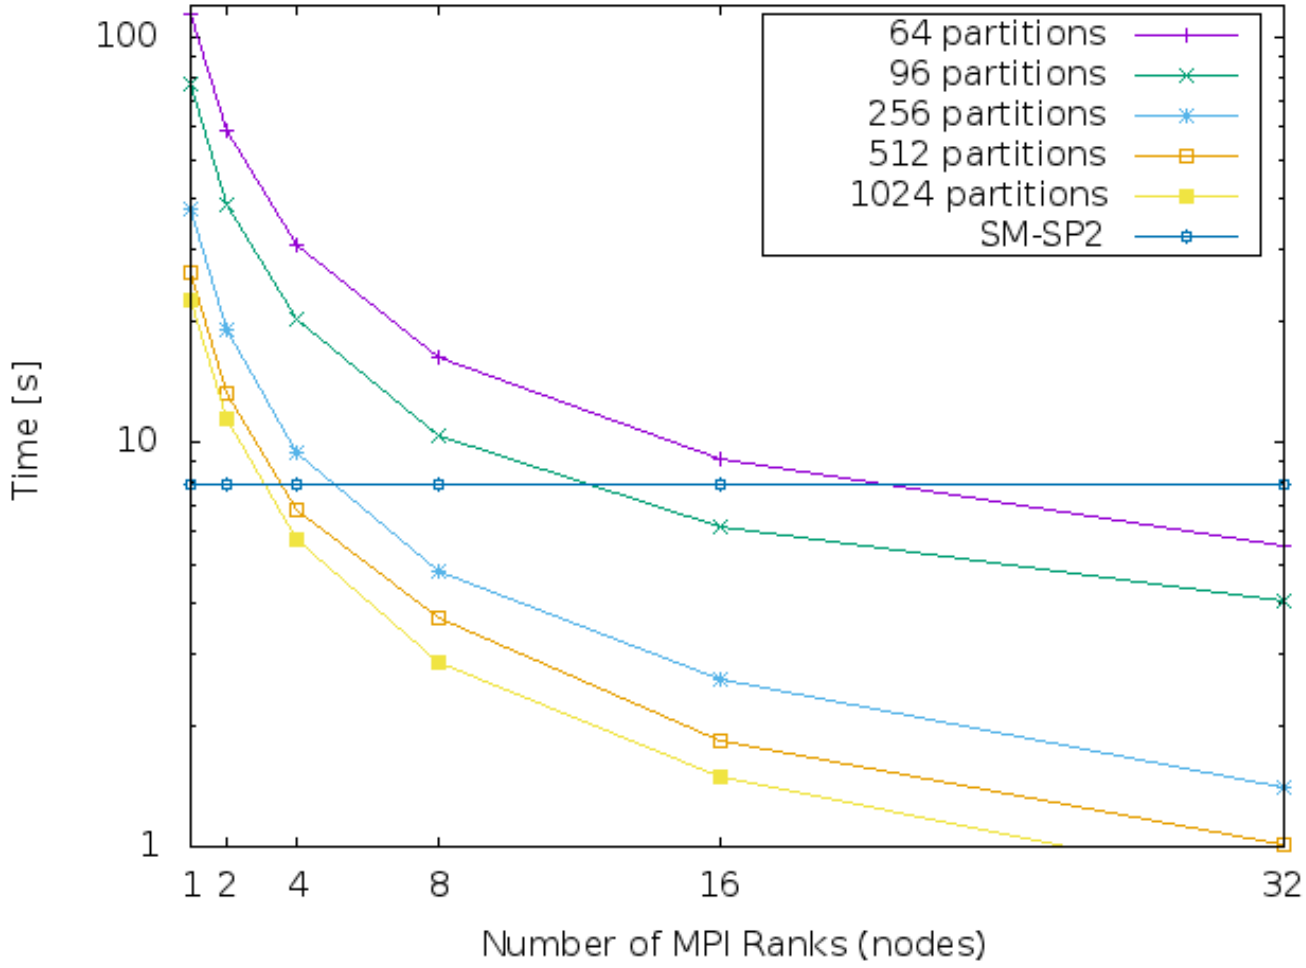

Conclusion

- New graph partitioning problem with applications in materials science and sparse matrix polynomials
  - Parts overlap
  - Objective function not directly related to edge cut
- Several implementations
  - Classical GP algorithms + SA postprocessing
  - KaHIP+SA gives best quality
  - Metis+SA best running time and best overall
- Parallel QMD implementation based on CHP runs about 10 times faster than SM based version