


# Task-Based Parallel Computation of the Density Matrix in Quantum-Based Molecular Dynamics using Graph-Partitioning

## Motivation

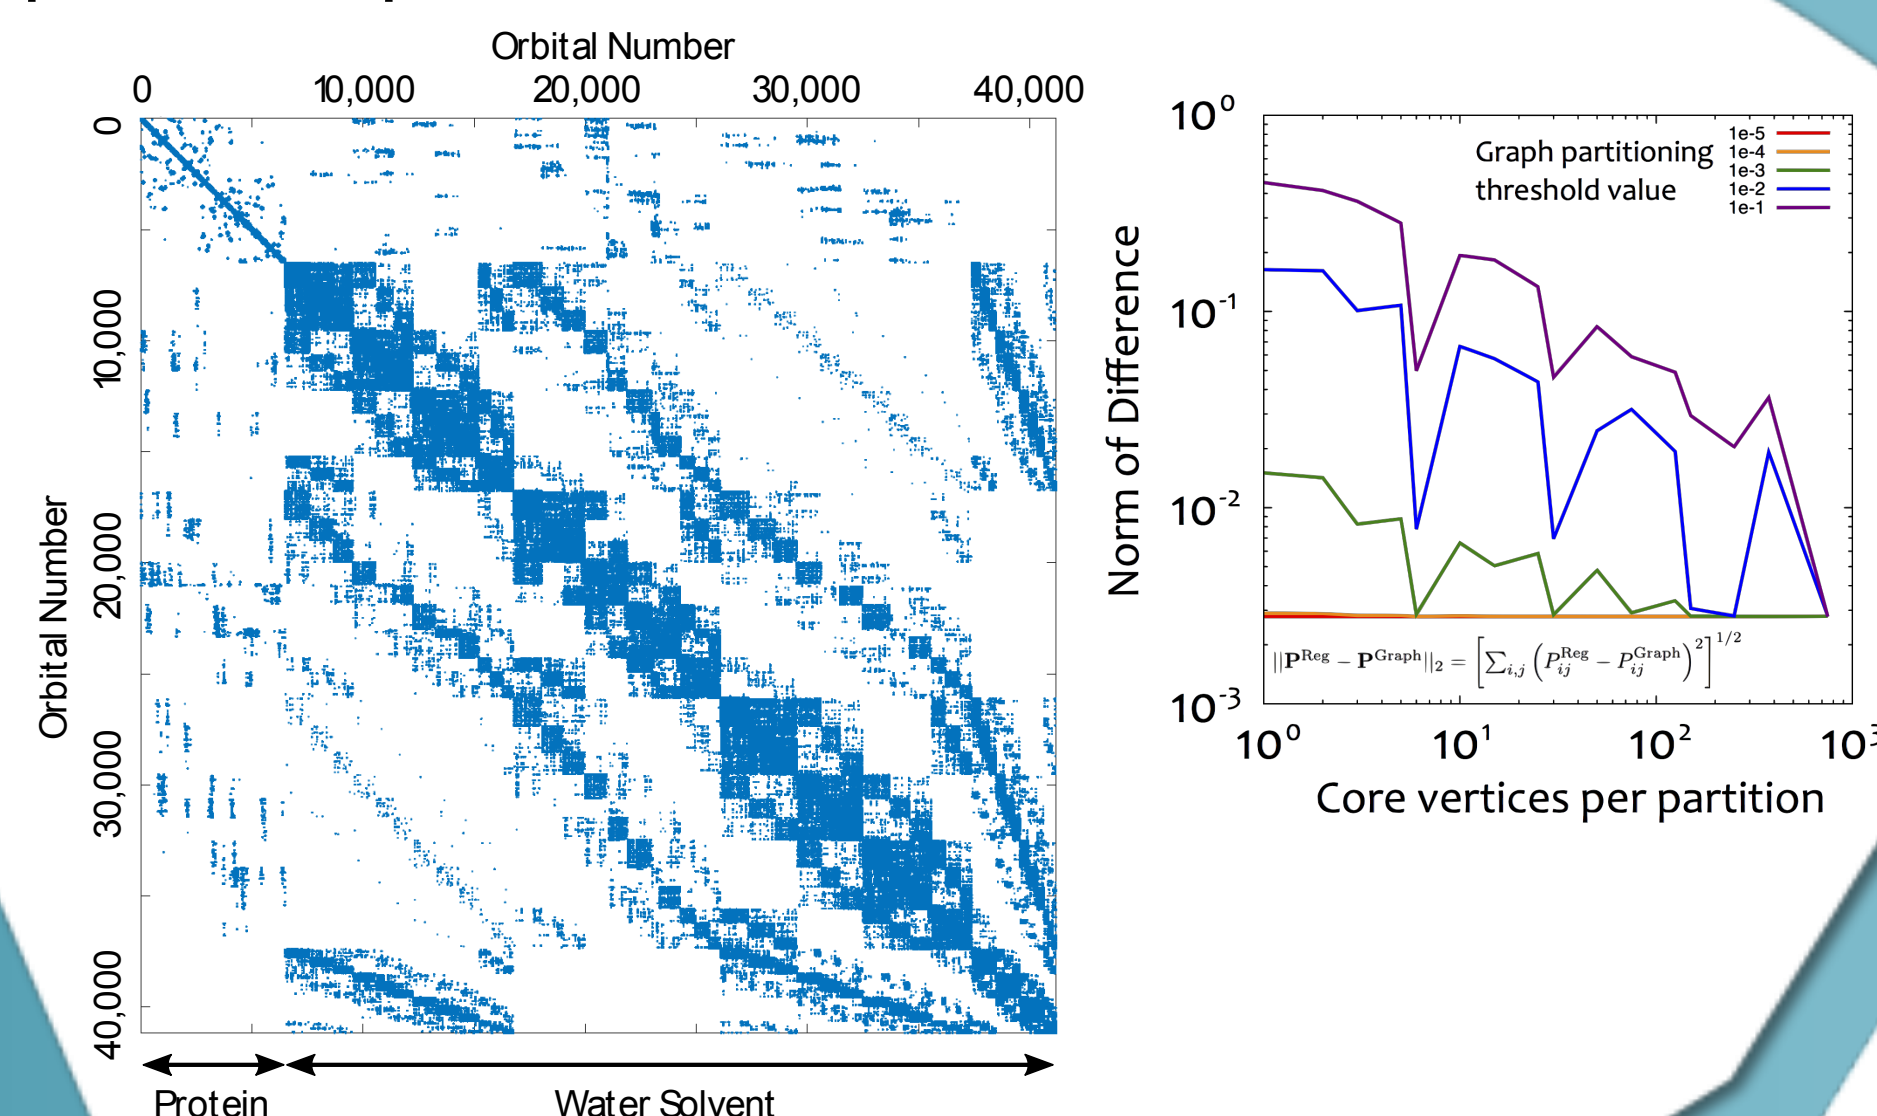
Quantum-based molecular dynamics (QMD) simulations are substantially more accurate than those using classical models. QMD trajectories evolve on a potential energy surface  $U$

$$U = 2\text{Tr}[(\mathbf{P} - \mathbf{P}^0)\mathbf{H}] - \frac{1}{2} \sum_{i,j=1}^N \gamma_{ij} q_i q_j + E_{\text{pair}}$$

that depends on the Hamiltonian matrix  $\mathbf{H}$ , which represents chemical bonding, and the corresponding electron density matrix  $\mathbf{P}$ .<sup>[1]</sup> Obtaining  $\mathbf{P}$  from  $\mathbf{H}$  is the most expensive computational step. Traditional  $O(N^3)$  matrix diagonalization limits the number of simulated atoms  $N$  to  $\sim 1000$ . Our goal is to use QMD for large systems ( $N > 10,000$ ), such as solvated proteins.



QMD simulations depend on quickly computing  $\mathbf{P}$  from  $\mathbf{H}$ . We investigate graph partitioning to obtain  $\mathbf{P}$  in a data parallel implementation of SP2.



## Second-Order Spectral Projection (SP2)

SP2 is an efficient,  $O(N)$  method to obtain  $\mathbf{P}$  for non-metallic systems, replacing  $O(N^3)$  diagonalization with a polynomial expansion of  $\mathbf{H}$ .<sup>[2]</sup> SP2 scales best when  $\mathbf{H}$  and  $\mathbf{P}$  are sparse.

$$\mathbf{P} = \theta [\mu \mathbf{I} - \mathbf{H}] = \lim_{i \rightarrow \infty} f_i [f_{i-1} [\dots f_0 [\mathbf{X}_0] \dots]]$$

$$f_i [\mathbf{X}_i] = \begin{cases} \mathbf{X}_i^2 & \text{if } 2\text{Tr}[\mathbf{X}_i] \geq N_e \\ 2\mathbf{X}_i - \mathbf{X}_i^2 & \text{if } 2\text{Tr}[\mathbf{X}_i] < N_e \end{cases}$$

$$\mathbf{X}_0 = \frac{\epsilon_{\text{max}} \mathbf{I} - \mathbf{H}}{\epsilon_{\text{max}} - \epsilon_{\text{min}}}$$

## Graph Partitioning

Traditional graph partitioning approaches minimize edge-cut between partitions.

To account for inter-partition interactions, we extend existing approaches to reduce the size of partitions with respect to both the number of vertices per partition as well as the number of neighbor vertices in adjacent partitions.

Denoting the number of nodes per partition as  $c$  and the number of neighbor vertices as  $h$ , we reduce the matrix operation costs for the SP2 method summed over all the partitions  $P$ :

$$\sum_{i \in P} (c_i + h_i)^3$$

## Acknowledgments



We wish to thank our mentors: Ben Bergen, Nick Bock, Marc Cawkwell, Hristo Djidjev, Christoph Junghans, Sue Mniszewski, Christian Negre, Anders Niklasson, Robert Pavel, and Ping Yang

[1] M. J. Cawkwell and A. M. N. Niklasson, J. Chem. Phys. **137**, 134105 (2012)  
 [2] A. M. N. Niklasson, Phys. Rev. B **66**, 155115 (2002)  
 [3] G. Karypis et al., Proc. 34th Design Automat. Conf. (1997)  
 [4] Z. Budimic et al., Sci. Program. **18**, 3 (2010)  
 [5] L. V. Kale and S. Krishnan, Charm++ (MIT Press, 1996)  
 The poster design was adapted from the designs of Felix Breuer.  
 This work has been approved for unlimited release under LA-UR-15-25700

## ISTI/ASC Co-Design Summer School


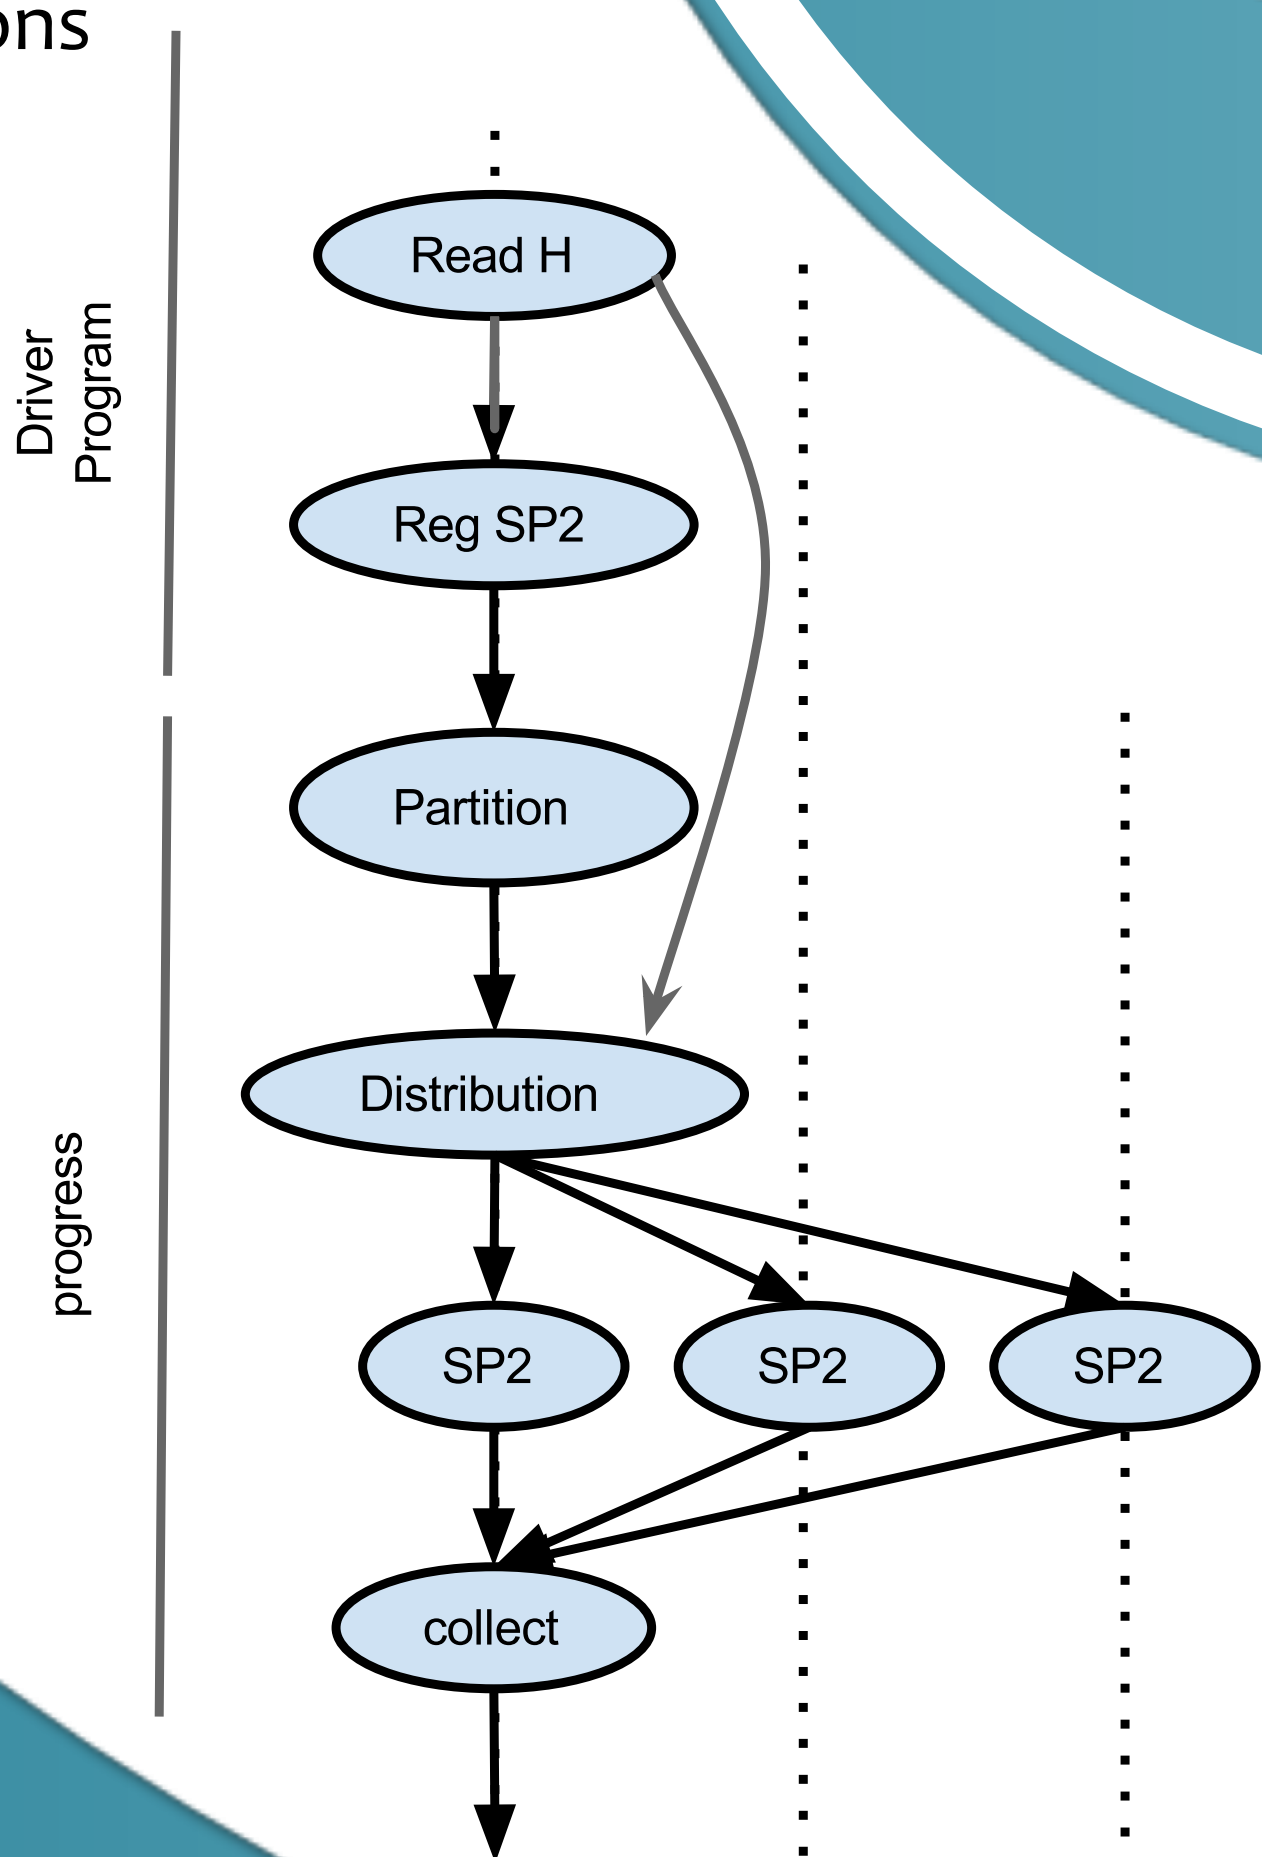
Purnima Ghale, Georg Hahn, Matthew Kroonblawd, Sergio Pino, Vivek Sardeshmukh, and Jerry Shi  
 Computer, Computational, and Statistical Sciences Division  
 Los Alamos National Laboratory

## Task-Based Parallel Implementation



Undesirable load-imbalances can arise as the partitions are not of equal size. Asynchronous task-based programming models, such as CnC and Charm++, mitigate load-balancing problems typical of MPI/OpenMP implementations by efficiently scheduling computations at runtime.

Intel Concurrent Collections<sup>[4]</sup> (CnC) and Charm++<sup>[5]</sup> implementations exploit the independent sub-problems resulting from our graph partitioned approach to SP2 and are integrated into an existing QMD code.



A naive approach to divide  $\mathbf{H}$  into sub-problems (blocking) is not efficient. Applying classical graph partitioning tools ( $hMETIS$ )<sup>[3]</sup> to  $\mathbf{P}$  performs better, but is not optimal for our purpose. Post-processing with simulated annealing (SA) can refine this result.

Protein: 31941 orbitals, 64 partitions.

method	$\sum_{i \in P} (c_i + h_i)^3$
blocking	1791576117859
$hMETIS$	346007029558
$hMETIS+SA$	311427370057