GRZEGORZ MAZUR[*], MARCIN MAKOWSKI[**], DAWID KUNA[***]

HYBRID-PARALLEL FORMULATION OF FUNDAMENTAL QUANTUM-CHEMICAL ALGORITHMS

Hybrid-parallel variants of Hartree-Fock, Kohn-Sham and Møller-Plesset second-level perturbation theory are described. Their efficiency with respect to the serial and MPI-based parallel implementations are measured and briefly analyzed. It is shown that while hybrid parallelization provide increased efficiency in all cases, the magnitude of the effect strongly depends on the features of the particular algorithm.

Keywords: *computational chemistry, parallelization*

HYBRYDOWE ZRÓWNOLEGLENIE PODSTAWOWYCH ALGORYTMÓW KWANTOWO-CHEMICZNYCH

Przedstawiono hybrydowo zrównoleglone warianty metod Hartreego-Focka, Kohna-Shama i rachunku zaburzeń Møllera-Plesseta drugiego rzędu. Porównano ich wydajność względem implementacji szeregowej i implementacji zrównoleglonej za pomocą mechanizmu przekazywania komunikatów (MPI). Pokazano, że hybrydowe zrównoleglenie zapewnia zwiększoną wydajność we wszystkich analizowanych przypadkach, przy czym wielkość uzyskanego przyspieszenia silnie zależy od cech danego algorytmu.

Słowa kluczowe: *chemia obliczeniowa, zrównoleglenie*

1. Introduction

Significantly growing demand for high quality computational chemistry results for large molecular systems cannot go unnoticed. As a result, the chemistry community needs for computational power are constantly growing. This trend is accompanied by quick development of high performance hardware solutions. The most abundant computing systems used for the purposes of computational chemistry are clusters,

[*] Department of Computational Methods in Chemistry, Jagiellonian University, ul. Ingardena 3, 30-060 Krakow, Poland, `mazur@chemia.uj.edu.pl`

[**] Department of Theoretical Chemistry, Jagiellonian University, ul. Ingardena 3, 30-060 Krakow, Poland, `makowskm@chemia.uj.edu.pl`

[***] AGH University of Science and Technology, ACC CYFRONET AGH, ul. Nawojki 11, 30-950 Krakow, Poland, `dawid.kuna@cyfronet.krakow.pl`

which, owing to the specificity of quantum-chemical calculations, provide the best efficiency to price ratio.

The multicore architecture is now well established as the most typical hardware solution for cluster nodes. While such systems can be very well served by the message passing approach usually found in parallel formulation of the computational chemistry algorithms, the shared memory and fast synchronization mechanism available within a node can be exploited to improve the efficiency. This gives rise to the so-called hybrid parallelism, where two or more of MPI[2], OpenMP[3], POSIX Threads (PT)[1] and/or General-Purpose computing on Graphics Processing Units (GPGPU) techniques are combined to utilize the computational resources more efficiently and effectively than any one of them on their own[7].

We designed such a hybrid approach for Hartree-Fock (HF)[6], Configuration Interaction Singles (CIS)[6], Kohn-Sham (KS)[8], Time-Dependent Density Functional Theory (TD-DFT)[5], and Møller-Plesset second order perturbation theory (MP2)[6] calculations. The algorithms are implemented within the niedoida[10] computational chemistry package. Building upon optimized message-passing based parallel implementation of these quantum-chemical methods[9], we introduced the next level of hierarchy employing PT-based parallelism within a batch of tasks scheduled for given node by the MPI dispatcher.

The paper is structured as follows. The core computational quantum chemistry algorithms are introduced in Section 2. In the next section our implementation of the hybrid parallel algorithms is briefly presented. Benchmark results are shown and analyzed in Section 4. Conclusions are drawn in Section 5.

2. Core Algorithms of Computational Quantum Chemistry

Large part of the computational quantum chemistry can be considered to be just a glorified name for simple tensor algebra. Specifically, these are various transformations of the two-electron integrals (electron repulsion integrals, ERI) tensor which consitute the rate-determining step of the calculations. The ERI are defined as

$$(\mu\nu|\kappa\lambda) = \iint d\boldsymbol{r}_1 d\boldsymbol{r}_2 \chi_\mu(\boldsymbol{r}_1)\chi_\nu(\boldsymbol{r}_1)\frac{1}{|\boldsymbol{r}_1 - \boldsymbol{r}_2|}\chi_\kappa(\boldsymbol{r}_2)\chi_\lambda(\boldsymbol{r}_2) \tag{1}$$

where χ denote the atomic orbitals. It has to be noted that the sheer size of the ERI tensor and its sparsity pattern prevent standard algebraic treatment.

While the actual resource requirements depend strongly on the type of tensor contraction performed by a specific algorithm, we can distinguish two main types of them, the two-index contraction

$$G_{\mu\nu} = \sum_{\kappa\lambda} P_{\kappa\lambda}\left[2(\mu\nu|\kappa\lambda) - (\mu\lambda|\kappa\nu)\right] \tag{2}$$

and the four-index transformation

$$(ia|jb) = \sum_{\mu\nu\lambda\sigma} C_{\mu i}C_{\nu a}C_{\kappa j}C_{\lambda b}(\mu\nu|\kappa\lambda). \tag{3}$$

The former consitutes the core of the Hartree-Fock[11, 6] and hybrid Kohn-Sham[4] methods. The latter is crucial in the post-Hartree-Fock methods for which we used the Moller-Plesset second order perturbation theory (MP2)[6] as a representative example.

Apart from the complexity of the ERI tensor transformations, the cost of calculating the integrals is not negligible. Hence, usually efforts are made to cache at least part of the tensor and reuse them in subsequent computations.

The Kohn-Sham method additionally requires calculating of the exchange-correlation potential matrix. Typically this is done by numerical integration, which introduces additional cost. On the other hand, pure (non-hybrid) DFT does not actually require full-fledged ERI tensor, because the Coulombic contribution to the total potential energy can be obtained by density-fitting procedure[12]. Such approach is usually denoted by DF-DFT. The analysis does not hold for hybrid potentials, which include part of the Hartree-Fock potential contribution, and therefore involve the ERI tensor.

The specifics of the methods may have strong impact on the efficiency of parallel implementation. Hence, we include in the analysis representative algorithms from all of the above methodologies. To keep the paper concise, we decided to focus on the HF, DFT, DF-DFT and MP2 methods. The structure of CIS and TDDFT implementation is analogous to HF and DFT, respectively. Hence, they do not require specific analysis.

3. Implementation details

The parallel variants of the core quantum-chemical algorithms are formulated by splitting the calculations into tasks which are executed concurrently. To fully exploit the typical cluster architecture specifics, the concurrency is introduced at two different levels.

At the cluster level, the parallelism is implemented using MPI. The computational problem is divided into tasks of different sizes. The tasks are stored in a task queue. A node gets the next task from the queue as soon as it completes the previous one. The splitting of the problem into tasks is organized as follows. A fraction $1/f$ of the original problem is divided into n tasks, where f is the splitting factor and n stands for the number of nodes. Then the procedure is repeated recursively for the remaining part of the problem. The recurrence is stopped when the size of the remaining part is smaller then the threshold t. The algorithm is parametrized by f and t. For the details of the procedure see Ref. [9].

At the node level, POSIX threads are employed to exploit the multicore architecture. The task assigned to the node by the MPI dispatcher is further divided into subtasks, which are put in a subtask queue. Then group of threads is woken up. Each thread gets a subtask from the queue and executes it. The procedure is repeated until the subtask queue is empty. Then the next task is requested from the MPI dispatcher. The operation is repeated until the whole computational problem is completed.

4. Results

To assess the usefulness of the proposed approach, a series of calculations were run in parallel mode and their relative performance with respect to single-core execution times were measured.

All calculations were performed for a molecule of 1,3-diphenylisobenzofuran (see Fig. 1), as a representative case for medium-sized molecular systems. At the HF and MP2 levels of theory the cc-pVTZ atomic basis was used. Given that the quality of the DFT results depends to much smaller extent on the size of the applied basis set, cc-pVDZ basis was used for DFT methods.
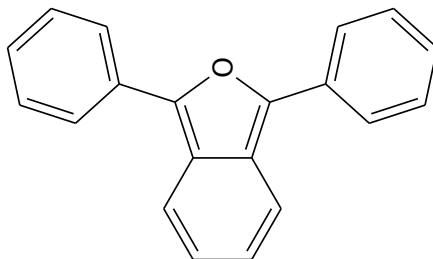

Fig. 1. Structural formula of 1,3-diphenylisobenzofuran

In order to allow for performance analysis, we run all the calculations on two different computer systems. The first was a cluster of 4-core nodes connected *via* Gigabit Ethernet. Each of the nodes was equipped with a Xeon X3450 CPU running at 2.67 GHz and 8 GB of memory. The other system was a 12-core fat node comprising two 6-core Xeon X5670 CPUs running at 2.93 GHz and 70 GB of memory. To avoid any interference with other processes runnning on the computer, we elected to use only 10 of the cores for the benchmark jobs.

The multicore fat node results (see Tab. 1) show that the difference between the MPI and thread-based parallelism is not very significant for CPU-bound methods. This is not the case for MP2, which seems to profit greatly from the shared memory present in the multithreaded version.

Speedup data obtained for a cluster (see Tab. 2) to some extent follow the results obtained for the single fat node. However, the quantitative analysis differ slightly. First of all, the advantage of the hybrid approach is more pronounced for the CPU-bound methods than in the fat node case. This is attributed mainly to more flexible dispatching of the subtasks, smaller communication overhead and better integral cache utilization with respect to purely message passing-based parallelism. On the other hand, the MP2 speedup, while significant, is not as large as in the case of the single-node multicore system. This can be easily explained by the fact that the hybrid approach results in the shared memory size being in-between the multithreaded and MPI cases.

Table 1

Relative speedup of selected parallel calculations on 12-core fat node (only 10 cores were used). See text for details

Method	Speedup	
	MPI	PT
HF	9.2	9.2
DFT	9.4	9.9
DF-DFT	9.0	9.9
MP2	8.5	40.2

Table 2

Relative speedup of selected parallel calculations on cluster. See text for details

Method	Number of cores	Speedup	
		MPI	Hybrid (MPI/PT)
HF	64	55.0	56.9
DFT	64	35.7	42.9
DF-DFT	64	22.8	35.1
MP2	24	7.6	21.7

5. Conclusions

The hierarchical hybrid implementation proved to provide increased efficiency in all cases we investigated. However, the magnitude of the effect strongly depends on the features of the particular algorithm. Among the algorithms under study, the most pronounced effect is observed in the case of MP2. In general, the algorithms which can effectively exploit shared memory are those profiting the most.

Acknowledgements

References

[1] IEEE Std 1003.1.

[2] MPI-2: Extensions to the message-passing interface.
`http://www-unix.mcs.anl.gov/mpi/mpi-standard/mpi-report-2.0/`
`mpi2-report.htm`, ?

[3] Openmp application programming interface. `http://www.openmp.org/drupal/`
`mp-documents/spec25.pdf`.

[4] Becke A.D.: *Density-functional exchange-energy approximation with correct asymptotic behaviour.* Phys. Rev. A, 38(6), 1988, pp. 3098–3100.

[5] Casida M.E.: *Time-dependent density functional response theory of molecular systems: Theory, computational methods, and functionals.* [in:] J.M. Seminario, editor, Recent developments and applications in density functional theory. Elsevier, Amsterdam, 1996.

[6] Helgaker T., Jorgensen P., Olsen J.: *Molecular Electronic Structure Theory.* John Wiley and Sons, ltd., 2000.

[7] Ishimura K., Kuramoto K., Ikuta Y., Hyodo S.: *Mpi/openmp hybrid parallel algorithm for hartree-fock calculations.* J. Chem. Th. Comp., 6, 2010, p. 1075.

[8] Kohn W., Sham L.J.: *Self-consistent equations including exchange and correlation effects.* Phys. Rev. A, 140(4), 1965, pp. 1133–1138.

[9] Mazur G., Makowski M.: *Effective resource allocation in parallel quantum-chemical calculations.* Computing and Informatics, in press.

[10] Mazur G., Makowski M.: *Development and optimization of computational chemistry algorithms.* Computing and Informatics, 28, 2009, pp. 115–125.

[11] Roothaan C.C.J.: *New developments in molecular orbital theory.* Rev. Mod. Phys., 23, 1951, pp. 69–89.

[12] Skylaris C.-K., Gagliardi L., Handy N.C., Ioannou A.G., Spencer S., Willetts A.: *On the resolution of identity coulomb energy approximation in density functional theory.* J. Mol. Struc. (Theochem), 501–502, 2000, pp. 229–239.