

Chapter 11

Parallel Quantum Chemistry at the Crossroads

Hubertus J. J. van Dam

Pacific Northwest National Laboratory, USA

ABSTRACT

Quantum chemistry was a compute intensive field from the beginning. It was also an early adopter of parallel computing, and hence, has more than twenty years of experience with parallelism. However, recently parallel computing has seen dramatic changes, such as the rise of multi-core architectures, hybrid computing, and the prospect of exa-scale machines requiring 1 billion concurrent threads. It is doubtful that current approaches can address the challenges ahead. As a result, the field finds itself at a crossroads, facing the challenge to successfully identify the way forward. This chapter tells a story in two parts. First, the achievements to date are considered, offering insights learned so far. Second, we look at paradigms based on directed acyclic graphs (DAG). The computer science community is strongly advocating this paradigm, but the quantum chemistry community has no experience with this approach. Therefore recent developments in that area will be discussed and their suitability for future parallel quantum chemistry computing demands considered.

INTRODUCTION

Quantum chemistry was formulated with the Schrödinger equation in 1926 as the science that seeks to explain chemistry, chemical processes and properties of molecules from solutions of this

equation. As a result this field of science was a compute intensive enterprise from the start. Since the beginning, when the word “computer” referred to a job description rather than a machine (Kopplin, 2002), there have been dramatic developments both on the theory side and as well as the practical side. On the theory side the formulation of the problem in a basis set, the formulation of

DOI: 10.4018/978-1-61350-116-0.ch011

theories to deal with the complex electron–electron interaction, such as Density Functional Theory (DFT) (Hohenberg & Kohn, 1964), Many Body Perturbation Theory (MBPT) (Møller & Plesset, 1934), and Coupled Cluster (CC) (Čížek, 1966), provided important handles on the problem. On the practical side the rapid development of computer technology and the design of efficient algorithms delivered access to the amount of computation needed to deliver useful answers.

In particular the development of computers followed initially a straightforward path of ever increasing compute power from single processing units. Most recently this trend has been broken as practical limits were hit. As it became impossible to simply crank up the clock speed even higher processor engineers opted instead for building chips with multiple processing units, the so called multi-core processors. This is a significant change. Although parallelism, as we currently know it, was introduced in quantum chemistry as early as 1988 (Clementi et al., 1988) and continuous and significant investment in parallel codes was made there was always the excuse that most machines were actually single processor machines to avoid becoming embroiled in code parallelization. However, now that every new machine is effectively a parallel machine it is clear that only parallel applications can utilize the available hardware effectively. But the changes go far beyond this point. Alongside the development of multi-core processors the development of Graphics Processing Units (GPU) towards General Purpose computing (GPGPU) created a new level of parallelism (Ufimtsev et al., 2008) (Vogt et al., 2008) (Genovese et al., 2009). The GPGPU's of today offer the potential of using 240 cores concurrently within a single physical machine. This is paired with the development of off-the-shelf solutions to connect multiple machines to build compute clusters. As a result relatively powerful machines can be built for low costs at present.

Nevertheless, the development towards exa-scale computers is likely going to be disruptive.

The reason for this lies in the energy requirements for such a machine. Having that the performance is specified and thereby the CPU power consumption to a large extent fixed it will be essential to save on memory and network power consumption to deliver these machines. This, however, raises the bar for writing scalable applications as there will be less space to keep data but it will also become more expensive to move data.

Against this back-drop of computer developments let us consider how close we have come to delivering a dream that was once captured in the title of a Europort2 project “Interactive Molecular Modeling through Parallelism” (Colbrook et al., 1995). What has been learned about parallelism in quantum chemistry and do we currently have a way to tackle the next level of parallelism?

With respect to answering these questions we find the field currently at a crossroads. Over the next decade machines will be developed that are a thousand times more powerful than the top end machines of today. These machines can be developed only by increasing the concurrency. However, many believe that the software development paradigms adopted so far cannot be stretched to such high levels of concurrency. Instead, computer scientists propose to write codes using directed acyclic graphs (DAG) that make the data dependencies explicit. This, they claim, allows parallelism to be exploited much more flexibly. In quantum chemistry there currently is no experience with paradigms of this kind. Hence, the field is faced with the question in which direction to proceed. In this chapter the way parallelism has been approached to date will be considered first. This provides insight into what has been achieved, how parallelism has been addressed, and what has been learned along the way. Next, looking forward, DAG based approaches are considered. In particular a number of existing frameworks and programming models are discussed to provide a flavor of the direction in which parallel programming may develop in the years to come.

26 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:
www.igi-global.com/chapter/parallel-quantum-chemistry-crossroads/60363

Related Content

Kansei Database and AR*-Tree for Speeding up the Retrieval

Yaokai Feng (2011). *Kansei Engineering and Soft Computing: Theory and Practice* (pp. 111-125).
www.irma-international.org/chapter/kansei-database-tree-speeding-retrieval/46394

AI-Based DBMS Controlled Speech Recognition Model for Some Common Computing Commands

Mrinmoy Sen, Sunanda Jana, Swarnajit Bhattacharya and Gitika Maity (2023). *Novel Research and Development Approaches in Heterogeneous Systems and Algorithms* (pp. 143-154).
www.irma-international.org/chapter/ai-based-dbms-controlled-speech-recognition-model-for-some-common-computing-commands/320128

Online Testing of Nondeterministic Systems with the Reactive Planning Tester

Jüri Vain, Marko Kääramees and Maili Markvardt (2012). *Dependability and Computer Engineering: Concepts for Software-Intensive Systems* (pp. 113-150).
www.irma-international.org/chapter/online-testing-nondeterministic-systems-reactive/55327

Formal Stepwise Development of Scalable and Reliable Multiagent Systems

Denis Grotsev, Alexei Iliasov and Alexander Romanovsky (2012). *Dependability and Computer Engineering: Concepts for Software-Intensive Systems* (pp. 58-74).
www.irma-international.org/chapter/formal-stepwise-development-scalable-reliable/55324

Virtualizing Network Functions in Software-Defined Networks

Vishal Kaushik, Ajay Sharma and Ravi Tomar (2018). *Innovations in Software-Defined Networking and Network Functions Virtualization* (pp. 26-51).
www.irma-international.org/chapter/virtualizing-network-functions-in-software-defined-networks/198192