Exploring complex quantum systems with a hybrid CPU-GPU computing platform

A. Yu. Polyakov[1,2], S. Denisov[2], and P. Hänggi[2]

[1]Sumy State University, Sumy, Ukraine
[2]Institute of Physics, University of Augsburg, Universitätstr. 1, D-86159 Augsburg, Germany

alexander.p1987@gmail.com

Abstract. One of the most striking features of quantum mechanics is the exponential growth of resources, required to find the states of a composite system, with the size of the system. This also is the origin of the two main bottlenecks in numerical studies of complex quantum systems, that are (i) diagonalizations of big matrices and (ii) propagations of large systems of linear differential equations with global symplectic structure. Operations of the first type are purely scalable, while most of the propagation algorithms allow for the high degree of parallelism. Here we show how the workload of finding Floquet eigenstates of an ac-driven non-integrable quantum system can be shared between a general-purpose central processing unit (CPU) and a graphic processing unit (GPU), when both are working within one computing platform. Namely, diagonalization steps are delegated to the CPU, while the time propagation is performed on the GPU. This strategy led to a computational time speed-up of several order of magnitude as compared to the performance of the CPU alone.

Keywords

GPU computing, CPU/GPU computing platforms, exact diagonalization, Floquet theory, computational quantum dynamics

1 Introduction

Fast progress in manipulations with cold and ultra-cold atoms, quantum optics and nanoscale fabrication techniques has brought quantum physics in touch with technology [1, 2]. No surprise that computational quantum physics plays an ever increasing role in suggesting, explaining and guiding new experiments. However, the main curse of quantum physics (some might like to call it 'peculiarity'), the exponential growth of the size of system's Hilbert space upon the deviation from the exactly solvable limit and increase of the number of the system components, imposes strong restrictions on the complexity of models. For example, there is no way to find the ground-state of an arbitrary quantum network consisting of 50 interacting spins on a classical computer; that would mean a (partial) diagonalization of a matrix $2^{50} \times 2^{50}$, which task is beyond the limit chip-based computers could grasp even in far future [3]. On the general-purpose CPU clusters, installed in most of the German universities nowadays, the full diagonalization of an arbitrary $N \times N$ matrix is limited by $N \approx 150,000$ [4]. Aside of the memory problems caused by the size of the matrices, diagonalization algorithms are poorly scalable and do not allow for the advanced parallelization. Recent developments, such as MAGMA [5], help to speed-up the diagonalization process through GPU-acceleration of some involved routines. Therefore, even during the diagonalization stage, one can tangible benefit from the use of hybrid CPU-GPU computing platforms.

When a complex quantum system is periodically modulated in time, the search for the system eigenstates, called *Floquet states* [6], is further complicated by the need to propagate a set of basis states to obtain the propagator matrix, whose eigenvectors are the Floquet states of the system. Practically, that means numerical integration of a large set of coupled differential equations for N complex variables. Taking into account that the chosen propagating scheme should be accurate enough to preserve the norm with a reasonable accuracy, which, in turn, means that the time step should be small and the complexity of every integration step is high, it is evident that the realization of this routine on a CPU will take a while. In fact, the time needed for the consecutive propagation of N vectors over one period of the driving, T, overweights the time needed for the

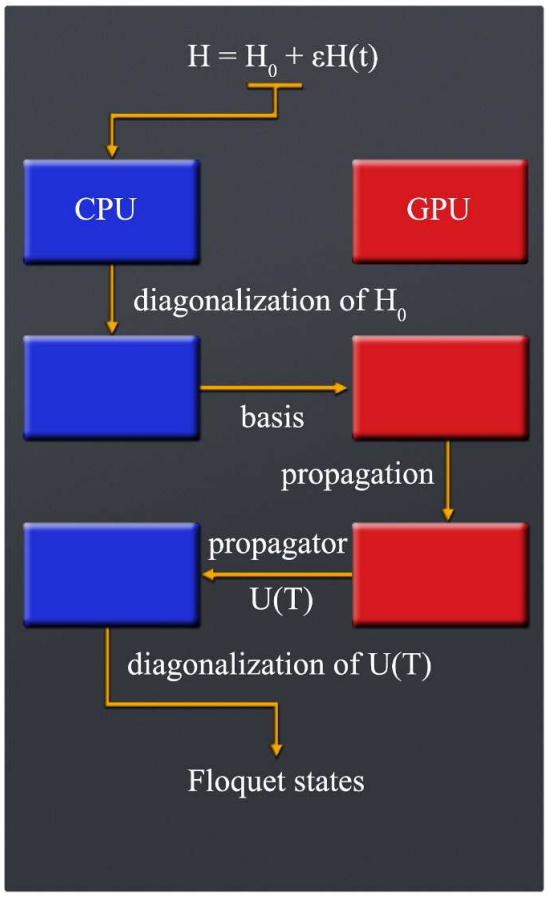

Figure 1. Workflow of the standard Floquet diagonalization routine [6] realized on a CPU-GPU computing platform. The first step, finding of the eigenbasis of a stationary Hamiltonian, H_0, is performed on a CPU. The obtained eigenbasis is then used to calculate the coupling matrix for the set of N linear differential equations for complex variables $b_i(t)$, $i = 1, ..., N$. This set is propagated over the time span $[0, T]$ N times, for the initial conditions $b_i(0) = \delta_{ki}$, $k = 1, .., N$. Once filled up, the propagator matrix $U(T)$ is sent back to the CPU where it is finally diagonalized. This step yields the full set of Floquet eigenstates of a Hamiltonian H.

diagonalization of a complex $N \times N$ matrix severely. However, (i) the propagation of the basis vectors can be performed in parallel and (ii) the components of a propagated vector can be updated simultaneously. The propagation routine is a highly scalable process, which is a clear call for the use of GPUs.

In this paper, we demonstrate that calculations of the Floquet eigenvectors of an ac-driven non-integrable quantum system of the form $\tilde{H}(t) = H_0 + \varepsilon H(t)$, where H_0 is a stationary opeartor and $H(t)$ is time-periodic operator, $H(t+T) = H(t)$, can be significantly accelerated by implementing a GPU as a coprocessor of a hybrid CPU-GPU computing platform. We outline the corresponding strategy and show, by using a driven quantum Duffing oscillator [2, 7] as a benchmark, that the overall computational time for a system with $N = 524$ states can be reduced by more than two orders of magnitude as compared to runs on a single CPU core.

2 Workflow and algorithms

The key idea is to perform diagonalization steps on a CPU while a GPU should be used for the time propagation of the coefficients $b_i(t)$. The corresponding workflow is sketched on Fig. 1. To integrate the system of linear time-dependent differential equations we use the symplectic integrators (SIs) proposed in Ref. [8]. The corresponding

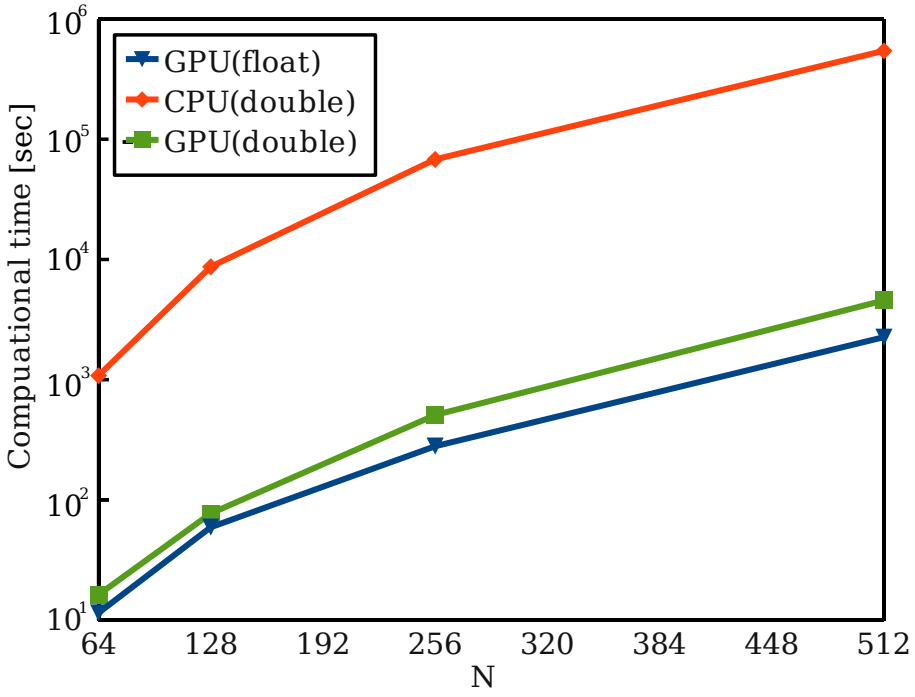

Figure 2. Finding Floquet states of the ac-driven quantum Duffing oscillator with N basis states. Computational times of a single CPU and a CPU-GPU platform as functions of the matrix size N.

schemes are straightforward in time and good in preservation of the symplectic structure of the quantum evolution. Although being over-performed by more sophisticated schemes, the SI algorithms allow for the high parallelism since the coefficients $b_i(t)$ can be updated in parallel.

3 Results

The GPU-located part of the workload can be performed with a high device occupancy, due to the small number of conditional statements that leads to extremely small thread divergence and rather good memory access pattern. We used Intel Xeon X5670 @ 2.93GHz (only one core was involved) and Tesla M2050 GPU (448 CUDA cores, 3 Gb of global memory) to compare the performance of a hybrid GPU-CPU platform and of a single CPU, as functions of the matrix size. A standard PCIe x16 Gen2e was used as a connector between the CPU and GPU processors.

To achieve higher accuracy, sometimes it is reasonable to use `double` type for floating-point numbers. However, this will take twice more memory, which might be undesirable in the context of rather limited capabilities of CUDA fast memory types. In addition, the use of the `float` type variables allows to gain high speed-up on non-specialized GPUs. The obtained results of the performance of the CPU-GPU platform and the single CPU are presented with Fig. 3.

4 Conclusion

In this paper we reported the implementation of a GPU-accelerated propagation routine for calculations of Floquet states of non-integrable ac-driven quantum systems. The performance gains increase with the size of the system, thus providing with the possibility to get a glance into Floquet spectra of complex quantum systems. The only restriction on the system size is imposed by the direct diagonalization step. A further seed-up can be readily obtained by using a multi-GPU coprocessor.

The used CPU-GPU computing platform is also suitable to out-shape quantum attractors of open ac-driven systems. In the framework of Floquet-Markov approach [6], this would mean an additional round of propagation, this time of the elements of the corresponding density matrix, ϱ_{ij}, until the latter relax to their stationary values.

The strategy sketched on Fig. 1 naturally allows for this extension: the set of Floquet states should be used to construct a system of equations for the coefficients ϱ_{ij}, which is then propagated on the GPU co-processor.

5 Acknowledgment

S.D. and P.H. acknowledge the support by the cluster of excellence Nanosystems Initiative Munich (NIM) (a seed-funding project "High-performance computational many-body quantum physics using Graphic Processing Units").

References

[1] I. Bloch, J. Dalibard, and W. Zwerger: Many-Body Physics with Ultracold Gases, *Rev. Mod. Phys.* 80, 885-964, 2008.

[2] M. Poot and H. S. J. van der Zant: Mechanical Systems in the Quantum Regime, *Phys. Rep.* 511, 273–335, 2012.

[3] When the topology of the interaction is simple and represent, for example, a regular one-dimensional chain, the groundstate of the system can be found by using the density-matrix renormalization group methods, see U. Schollwoeck: The Density-matrix Renormalization Group, *Rev. Mod. Phys.* 77: 259-315, 2005. This is indeed one of the greatest breakthrough of modern computational quantum physics.

[4] This is the biggest matrix we were able to diagonalize on the computational cluster of the Max-Planck Institute for the Physics of Complex System (Dresden, Germany). This job fully occupied a multi-core CPU with $1,5$ Tb of RAM for 44 days.

[5] F. Song, S. Tomov, J. Dongarra, "Enabling and Scaling Matrix Computations on Heterogeneous Multi-Core and Multi-GPU Systems," *26th ACM International Conference on Supercomputing (ICS 2012)*,San Servolo Island, Venice, Italy, June, 2012.

[6] M. Grifoni and P. Hänggi: Driven Quantum Tunneling, *Phys. Rep.* 304, 229-358, 1998.

[7] F. Grossmann, T. Ditrich, P. Jung, and P. Hänggi: Coherent Destruction of Tunneling, *Phys. Rev. Lett.* 67, 516-519, 1991; V. Peano and M. Thorwart: Macroscopic Quantum Effects in a Strongly Driven Nanomechanical Resonator, *Phys. Rev. B* 70, 235401, 2004.

[8] S. K. Gray and D. Manolopoulos: Symplectic Integrators Tailored to the Time-Dependent Schrödinger Equation, *J. Chem. Phys.* 104, 7099-7112, 1996.