

GPU-accelerated algorithms for many-particle continuous-time quantum walks[☆]



Enrico Piccinini^{a,c,*}, Claudia Benedetti^b, Iliaria Siloi^c, Matteo G.A. Paris^{b,d},
Paolo Bordone^{c,d}

^a Dipartimento di Ingegneria dell'Energia Elettrica e dell'Informazione "Guglielmo Marconi" - DEI, Università di Bologna, I-40136 - Bologna, Italy

^b Quantum Technology Lab, Dipartimento di Fisica, Università degli Studi di Milano, I-20133 - Milano, Italy

^c Dipartimento di Scienze Fisiche, Informatiche e Matematiche – FIM, Università di Modena e Reggio Emilia, I-41125 - Modena, Italy

^d Centro S3, CNR-Istituto di Nanoscienze, I-41125 - Modena, Italy

ARTICLE INFO

Article history:

Received 2 December 2016

Received in revised form 31 January 2017

Accepted 10 February 2017

Available online 24 February 2017

Keywords:

GPU

CUDA

Continuous-time quantum walks

ABSTRACT

Many-particle continuous-time quantum walks (CTQWs) represent a resource for several tasks in quantum technology, including quantum search algorithms and universal quantum computation. In order to design and implement CTQWs in a realistic scenario, one needs effective simulation tools for Hamiltonians that take into account static noise and fluctuations in the lattice, i.e. Hamiltonians containing stochastic terms. To this aim, we suggest a parallel algorithm based on the Taylor series expansion of the evolution operator, and compare its performances with those of algorithms based on the exact diagonalization of the Hamiltonian or a 4th order Runge–Kutta integration. We prove that both Taylor-series expansion and Runge–Kutta algorithms are reliable and have a low computational cost, the Taylor-series expansion showing the additional advantage of a memory allocation not depending on the precision of calculation. Both algorithms are also highly parallelizable within the SIMT paradigm, and are thus suitable for GPGPU computing. In turn, we have benchmarked 4 NVIDIA GPUs and 3 quad-core Intel CPUs for a 2-particle system over lattices of increasing dimension, showing that the speedup provided by GPU computing, with respect to the OPENMP parallelization, lies in the range between 8x and (more than) 20x, depending on the frequency of post-processing. GPU-accelerated codes thus allow one to overcome concerns about the execution time, and make it possible simulations with many interacting particles on large lattices, with the only limit of the memory available on the device.

Program summary

Program Title: cuQuWa

Licensing provisions: GNU General Public License, version 3

Program Files doi: <http://dx.doi.org/10.17632/vjpnjgycdj.1>

Programming language: CUDA C

Nature of problem: Evolution of many-particle continuous-time quantum-walks on a multidimensional grid in a noisy environment. The submitted code is specialized for the simulation of 2-particle quantum-walks with periodic boundary conditions.

Solution method: Taylor-series expansion of the evolution operator. The density-matrix is calculated by averaging multiple independent realizations of the system.

External routines: cuBLAS, cuRAND

Unusual features: Simulations are run exclusively on the graphic processing unit within the CUDA environment. An undocumented misbehavior in the random-number generation routine (cuRAND package) can corrupt the simulation of large systems, though no problems are reported for small and medium-size systems. Compiling the code with the `-arch=sm_30` flag for compute capability 3.5 and above fixes this issue.

© 2017 Elsevier B.V. All rights reserved.

[☆] This paper and its associated computer program are available via the Computer Physics Communication homepage on ScienceDirect (<http://www.sciencedirect.com/science/journal/00104655>).

* Corresponding author at: Dipartimento di Ingegneria dell'Energia Elettrica e dell'Informazione "Guglielmo Marconi" – DEI, Università di Bologna, I-40136 - Bologna, Italy

E-mail address: enrico.piccinini@unibo.it (E. Piccinini).

1. Introduction

Quantum walks (QWs) are a generalization of classical random walks to the quantum regime. They were first introduced in the discrete-time version [1] and later as continuous-time quantum walks (CTQWs) in the context of quantum computation and decision trees [2]. In this framework, it has been shown that single-particle quantum walk-based algorithms may outperform the classical counterpart in terms of traveling time through a graph. Since then, QWs, both in the continuous- and discrete-time versions, have been the subject of extensive studies. Besides, QWs have been generalized to many-particle quantum walks, where the time evolution of the walkers depends upon their statistics, indistinguishability and kind of interaction [3–7]. CTQWs on more complex structures, e.g., complex graphs, have been also the focus of more recent analysis [8–11]. Overall, CTQWs have been proved a useful tool in a variety of contexts, ranging from transport through a graph [12], to quantum search algorithms [13,14], graph isomorphism testing [15–17] and universal quantum computation [18,19].

In realistic experimental scenarios, imperfections in the fabrication of the lattice may induce Anderson localization of the walkers [20–23], whereas stochastic fluctuations of the environment may come into play destroying the *quantumness* of the system, i.e., the superposition of states and the phase coherence, and, in turn, its peculiar propagation features characterized, for instance, by the single-particle position variance [24–27]. A more realistic description for noisy quantum walks should therefore take into account the noise that may affect the evolution of the walkers. A convenient way to describe noise is to introduce suitable stochastic terms in the Hamiltonian, in order to model static or dynamical fluctuations that may affect both the on-site energies or the tunneling amplitudes of the walkers [25,26,28]. The dynamical evolution of the QW is then obtained as the ensemble average over all possible realizations of the stochastic processes mimicking the noise. In practice, the ensemble average is computed numerically as an average over a finite number of realizations: the larger the number of the realizations, the more accurate the simulation of the CTQW.

Evaluating the dynamics of a many-particle state over a noisy lattice requires the numerical solution of a set of differential equations that include stochastic terms [29]. The total number of equations to solve grows rapidly as long as the numbers of nodes, particles, and realizations increase, thus making the problem more and more computationally demanding with longer execution times. In fact, codes for simulating many-particle CTQWs have been developed for high-performance clusters with distributed memory [30]. On the other hand, the evolution of computer architectures towards multicore processors even in stand-alone workstations enabled important cuts of the execution time by introducing the possibility of running multiple threads in parallel and spreading the workload among cores. This possibility was boosted up by the general purpose parallel computing architectures of modern graphic cards (GPGPUs). In the latter, hundreds or thousands of computational cores in the same single chip are able to process simultaneously a very large number of data. It should also be noted that an impressive computational power is present not only in dedicated GPUs for high-performance computing, but also in commodity graphic cards, which make modern workstations suitable for numerical analyses. In order to exploit such a huge computational power, algorithms must be first redesigned and adapted to the SIMT (Single Instruction Multiple Thread) and SIMD (Single Instruction Multiple Data) paradigms and translated then into programming languages with hardware-specific subsets of instructions. Among them, one of the most diffuse is CUDA-C, a C extension for the Compute Unified Device Architecture (CUDA)

that represents the core component of NVIDIA GPUs. As a matter of fact, the use of GPUs for scientific analysis, which dates back to mid and late 2000s [31–35], dramatically boosted with a two-digit yearly increasing rate since 2010. Just looking at the computational physics realm, several GPU-specific algorithms have been proposed in the last three years, e.g., for stochastic differential equations [36], molecular dynamics simulations [37,38], fluid dynamics [39,40], Metropolis Monte Carlo [41] simulations, quantum Monte Carlo simulations [42], and free-energy calculations [43].

The evolution of many-particle QWs in a noisy environment can be classified as an *embarrassing parallel* problem, since there is little to none communication among realizations. Problems of this kind take great advantage of GPGPU computing, since the solving algorithms can be designed to run directly on the GPU in such a way that communications are implemented via shared memory on the device (graphic card) and data transfer between the host (CPU) and the device and v.v. is limited to unavoidable input/output operations.

In this paper, we have compared parallel algorithms for CTQWs evolution in a noisy environment based on the exact diagonalization of the Hamiltonian, the 4th order Runge–Kutta integration method and the Taylor-series expansion of the evolution operator. Solutions that avoid the diagonalization of the Hamiltonian (besides those implemented in this work, see also, e.g., Refs. [44–47]) result in a lower computational cost and pave the way to highly parallelizable algorithms within the SIMT paradigm, thus leading to a straightforward implementation directly on the GPU. We have then benchmarked 4 NVIDIA GPUs and 3 quad-core Intel CPUs for a 2-particle system over a lattice of increasing dimensions and have shown that the GPU speedup with respect to the OPENMP parallelization fluctuates from 8x to more than 20x, depending on the frequency of post-processing. Thus, GPU-accelerated codes allow the design of simulations involving many particles or large lattices, with the only limit of the memory available on the device.

The paper is organized as follows: In Section 2 we discuss and derive efficient algorithms for the dynamics of CTQWs; in Section 3 we provide the main details on their implementation and in Section 4 we compare the performances of the algorithms on different CPUs and GPUs. Section 5 closes the paper with some concluding remarks.

2. Algorithms for quantum walks in a noisy environment

Let us consider a q -dimensional regular lattice hosting m quantum particles, and let N_i and $2k_i$ be the numbers of mesh elements (nodes) and of neighbors to be considered along the i th direction. The system in hand is described by an $N^m \times N^m$ matrix, storing the elements of the Hamiltonian \mathcal{H} and by a N^m vector for the wavefunction Ψ , where $N = \prod_i N_i$ is the total number of mesh nodes. When $k = \sum_i 2k_i \ll N$, the Hamiltonian \mathcal{H} is largely sparse with a maximum filling factor $(mk + 1)/N^m$.

Since we are interested in quantum walks in a noisy environment, transitions from node α to node β are ruled by a deterministic ($c_{\alpha\beta}$) and a stochastic ($\xi_{\alpha\beta}$) parameter, independently of which of the m particle jumps between them. Let $\vec{x} = \{x_1, x_2, \dots, x_m\}$ and $\vec{y} = \{y_1, y_2, \dots, y_m\}$, where $(x_s, y_s) \in [1, N]$, be the initial and final sets of nodes of a transition. The deterministic coefficients c read: $c_{x_s y_s} = c_{\alpha, \beta}$, for any $s \in [1, m]$ such that $x_s = \alpha$ and $y_s = \beta$. The same applies for the stochastic coefficients ξ . These terms switch between multiple values at random times during the simulation (switching times) in order to describe (generally time-dependent) fluctuations induced by lattice imperfections and/or external sources of noise. Let

$$\hat{x} = \hat{x}(x_1, x_2, \dots, x_m) = 1 + \sum_{s=1}^m N^{s-1}(x_s - 1),$$

$$\hat{y} = \hat{y}(y_1, y_2, \dots, y_m) = 1 + \sum_{s=1}^m N^{s-1}(y_s - 1);$$

in the position basis, the elements of the Hamiltonian $\mathcal{H}_{\hat{x}\hat{y}}$ read

$$\mathcal{H}_{\hat{x}\hat{y}} = \begin{cases} c_{\hat{x}\hat{x}} + \xi_{\hat{x}\hat{x}}, & \text{if } \hat{y} = \hat{x} \\ c_{x_s y_s} + \xi_{x_s y_s} & \text{if } \exists(x_s, y_s) \text{ such that } x_s \\ & \text{is connected to } y_s \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

The terms $\mathcal{H}_{\hat{x}\hat{x}}$ quantify both the on-site energies of the walkers and the interaction energy among particles, whereas the terms $\mathcal{H}_{\hat{x}\hat{y}}$ with $x \neq y$ describe the tunneling amplitudes between neighboring sites.

The interaction energy among particles is included in the diagonal deterministic terms $c_{\hat{x}\hat{x}}$. A significant example concerning the case of nearest neighbor interaction is reported in Ref. [26]. The repulsive or attractive nature of the interaction does not play a role on the dynamics of the particles here, as also noted in [48]. The applicability and the performances of the algorithm are also independent of the specific kind of interaction, since dealing with a different process only changes the numerical entries on the diagonal elements of the Hamiltonian and does not affect the features of our code. Moreover, our code is independent on the statistical nature of the particles, since Hamiltonian (1) conserves the statistics. This implies that we are able to describe both bosons and fermions dynamics by a suitable input parameter setting either symmetrized or anti-symmetrized initial quantum state. Then, the time evolution preserves the initial symmetry.

Hamiltonian (1) allows us to describe a large variety of physical systems, ranging from many-particle QWs [25,26], to Bose-Hubbard and Fermi-Hubbard models [49–51], where the dynamical noise has been a very challenging issue to address so far, due to the computational complexity of the problem. The ability to explore in a systematic and detailed way the effect of the stochastic noise on the dynamics of many-particle systems is a relevant tool to better describe realistic systems and to open up the road to the full understanding of the mechanisms of transport and diffusion over complex networks.

In particular, the decoherent dynamics of two interacting particles hopping on a one-dimensional noisy lattice has been recently addressed [26], showing that accelerated codes allow one to explore very different dynamical regimes upon tuning the ratio between the time scale of the noise and the interaction among the walkers. Under appropriate initial conditions, the dynamics in the presence of fast noise (e.g., fast decaying autocorrelation function) leads to a faster propagation with respect to the unitary evolution, as detected by the variance of the single particle probability distribution $\sigma^2(\tau)$; on the other hand, in the slow noise regime (e.g., slow decaying autocorrelation function) the system displays an Anderson-like localization, and propagation is suppressed. Results are reported in Fig. 1. As it is apparent from the plots, the behavior of the two fermions with next-neighbor interaction (and starting from next-neighbor sites) is very close, at least qualitatively, to that of two bosons with on-site interaction (and starting from the same site).

The time evolution of the system is provided by the Schrödinger equation

$$i\hbar \frac{d|\Psi\rangle}{dt} = \hat{\mathcal{H}}|\Psi\rangle, \quad (2)$$

where \hbar is the reduced Planck constant; the knowledge of $|\Psi(t)\rangle$ at each time step yields the $N^m \times N^m$ density matrix $\rho(t) = |\Psi(t)\rangle\langle\Psi(t)|$, which is used to evaluate the average over realizations $\langle\rho(t)\rangle$ and eventually further post-processed to calculate any desired observable quantity.

Consequently to the introduction of random terms, in order to avoid overweighting of outliers and produce a reliable ensemble




Fig. 1. Single particle variance σ^2 as a function of the dimensionless time τ for two interacting fermions (left) and bosons (right) starting from next-neighbors sites and same site, respectively. In the present case fermions interact when occupying neighboring sites, while bosons the same site. Each panel considers a different interaction strength V , whose value is normalized to the next-neighbors hopping probability. The noiseless evolution (solid black line) is here compared with the one obtained in fast noise regime (dashed red line), where the stochastic terms $\xi_{\alpha\beta}$ ($\alpha \neq \beta$) switch between values ± 0.9 with a switching rate (normalized to next-neighbors hopping probability) equal to 10.0.

average it is required to run a sufficiently large number R of simulations (a.k.a. realizations, usually $R \geq 1000$), and then averaging the density matrix. In order to speed-up the calculation and significantly cut the execution time, realizations can be run in parallel, as they are independent from each other. However, in the parallel execution, memory usage rapidly increases because at least a Hamiltonian matrix \mathcal{H}_i and a wave-function Ψ_i must be stored for each realization i . As a matter of fact, memory occupancy may become quickly an issue when the grid size or the number of particles increases.

2.1. Diagonalization of the Hamiltonian

If we suppose that the Hamiltonians \mathcal{H}_i do not change significantly within the time-step δt , Eq. (2) can be solved in the quasi-static approximation. The exact time evolution of a QW is provided by the well-known eigenproblem

$$(\hat{\mathcal{H}}_i - \varepsilon_i)|\Psi_i\rangle = 0, \quad (3)$$

that yields the eigenvalues ε_{ij} and the eigenvectors \mathbf{w}_{ij} of the i th Hamiltonian. The evolution of the wave function is then given by

$$|\Psi_i(t + \delta t)\rangle = \sum_j \exp\left(-\frac{i}{\hbar}\varepsilon_{ij}\delta t\right) |\mathbf{w}_{ij}\rangle \langle \mathbf{w}_{ij} | \Psi_i(t)\rangle. \quad (4)$$

The pseudocode for the parallel implementation is given in Algorithm 1.

Algorithm 1 Pseudocode for solving the CTQW dynamics via diagonalization of the Hamiltonian matrix

```

1: Initialize Hamiltonians  $\mathcal{H}_i$ 
2: Initialize switching times
3: while time  $t < t_{\max}$  do
4:   for all realizations do ▷ Begin Parallel Section
5:     Diagonalize  $\mathcal{H}_i \rightarrow \{\varepsilon_{ij}, \mathbf{w}_{ij}\}$ 
6:      $|\Psi_i(t + \delta t)\rangle \leftarrow \sum_j e^{-\frac{i}{\hbar}\varepsilon_{ij}\delta t} |\mathbf{w}_{ij}\rangle \langle \mathbf{w}_{ij} | \Psi_i(t)\rangle$ 
7:     Update switching times
8:      $\mathcal{H}_i \leftarrow \mathcal{H}_i(t + \delta t)$ 
9:   end for ▷ End Parallel Section,  $\sim O(N^{3m})$ 
10:   $t \leftarrow t + \delta t$ 
11:  if postprocessing then
12:     $\langle \rho(t) \rangle \leftarrow \frac{1}{R} \sum_i |\Psi_i(t)\rangle \langle \Psi_i(t)|$  ▷  $\sim O(N^{2m})$ 
13:    Post-process  $\langle \rho(t) \rangle$ 
14:  end if
15: end while

```

It is worth noticing that (a) this algorithm requires a large number of computationally intensive events of the order $\sim O(N^{3m})$ and (b) it is necessary to store N^m eigenvectors of N^m components per realization, which is exactly the same memory space that the dense Hamiltonian matrix would occupy. As a matter of fact, this issue may jeopardize the efficiency of the code, even in the case of a parallel implementation.

2.2. Integration of ordinary differential equations

Going back to the general solution of Eq. (2), we may directly tackle the time-dependent Schrödinger equation as a set of ordinary differential equations for the vector $|\Psi_i\rangle$ and solve it by means of standard integration techniques that dispose of the calculation of the eigenstates. A widely-used integration scheme is represented by the 4th-order Runge–Kutta method.

In this case, there is no need of allocating a memory space as large as a dense Hamiltonian would require. The Hamiltonian topology, i.e., how nodes are connected to each other, is known *a-priori* from the definition of the mesh, and holds true for all of the realizations. In principle, up to $mk + 1$ non-null elements are present in each row of the Hamiltonian. As a consequence, each of the $N^m \times N^m$ Hamiltonians \mathcal{H}_i can be stored as a $N^m \times (mk + 1)$ reduced matrix $\tilde{\mathcal{H}}_i$. A common $N^m \times (mk + 1)$ topology matrix holding the indexes of non-null elements also adds. Since transitions from node α to node β and v.v. share the same rate, the symmetry of $\tilde{\mathcal{H}}_i$ allows for further memory savings down to $N^m \times (mk/2 + 1)$ elements. These relationships hold true for a regular lattice; in the case of a general graph, where each site is connected to a variable number of other nodes, the approach is still applicable

with the only difference that the number k of non-null elements per row in the topology matrix is replaced by the number of connections.

The 4th-order Runge–Kutta procedure lets the wave-functions $|\Psi_i\rangle$ evolve by means of the linear combination of 4 intermediate states $|K_i^{(j)}\rangle$, $j = 1 \dots 4$. The evaluation of any component belonging to the j th intermediate state requires only the knowledge of the wave-function at the current time step, the reduced Hamiltonian and the $(j - 1)$ -th state at the indexes stored in the corresponding row of the topology matrix. Since nodes are topologically equivalent to each other, SIMD and SIMT paradigms apply, allowing for a second degree of parallelization over nodes. The parallelizations over realizations and over nodes can be collapsed into a larger loop (RN^m steps), which may better balance the computational burden assigned to each computing unit. The pseudocode for the implementation of the 4th-order Runge–Kutta method is reported in Algorithm 2.

Algorithm 2 Pseudocode for solving the CTQW dynamics via integration of the Schrödinger equation using the 4th-order Runge–Kutta method

```

1: Define Hamiltonian topology
2: Initialize reduced Hamiltonians  $\tilde{\mathcal{H}}_i$ 
3: Initialize switching times
4: while time  $t < t_{\max}$  do
5:   for all realizations do ▷ Begin Parallel/SIMT Section
6:     for  $j = 1 \rightarrow 4$  do
7:        $(|\Psi_i\rangle, |K_i^{(j-1)}\rangle, \tilde{\mathcal{H}}_i) \rightarrow |K_i^{(j)}\rangle$ 
8:     end for
9:      $|\Psi_i(t + dt)\rangle \leftarrow \sum_{j=1}^4 \mu_j |K_i^{(j)}\rangle$ 
10:    Check norm of  $|\Psi_i(t + dt)\rangle$ 
11:    Update switching times
12:     $\tilde{\mathcal{H}}_i \leftarrow \tilde{\mathcal{H}}_i(t + \delta t)$ 
13:   end for ▷ End Parallel/SIMT Section,  $\sim O(RN^m)$ 
14:    $t \leftarrow t + \delta t$ 
15:   if postprocessing then
16:      $\langle \rho(t) \rangle \leftarrow \frac{1}{R} \sum_i |\Psi_i(t)\rangle \langle \Psi_i(t)|$  ▷  $\sim O(RN^{2m})$ 
17:     Post-process  $\langle \rho(t) \rangle$ 
18:   end if
19: end while

```

The scheme in Algorithm 2 requires a single loop of sums and products; the algorithmic complexity of time evolution is thus reduced to the order $\sim O(RN^m)$, with a large speedup compared to the case discussed in Section 2.1. The most computationally intensive routine is now represented by the calculation of the average density-matrix, order $\sim O(RN^{2m})$, whose number of calls may vary depending on the desired precision of the output.

The 4th order Runge–Kutta method does not conserve the norm, and intermediate checks and corrective actions are required to avoid unphysical outcomes. It may also happen that the norm of $|\Psi_i(t)\rangle$ strongly deviates from its theoretical value within a single time step. In order to fix this issue two strategies may be devised. In the first one, higher-order Runge–Kutta methods can similarly be implemented to reach a better accuracy within the same time step, but memory allocation would grow since a larger number of intermediate states are required. On the other hand, one could reduce the time step in such a way that the cumulative error does not drive the simulation far away from its correct path. The immediate shortcoming is the increase of the running time inversely proportional to the step reduction; nonetheless, this solution becomes mandatory if it is not possible nor convenient to increase the memory allocation.

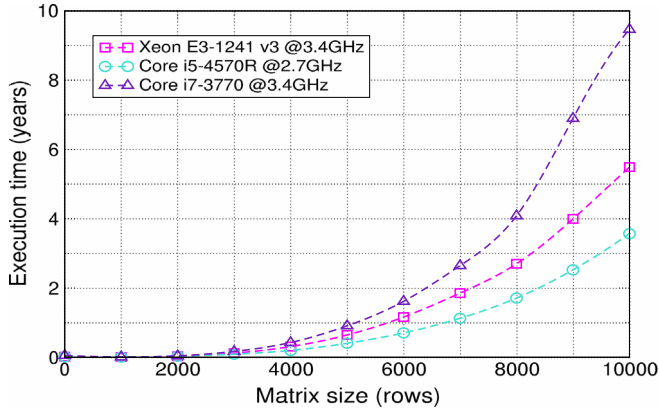


Fig. 2. Projected execution time required for diagonalizing 10^6 times a symmetric matrix in single precision with the MKL 11.2 library. The tests have been performed on 4 CPU cores, as this configuration preliminarily proved to maximize the overall performance.

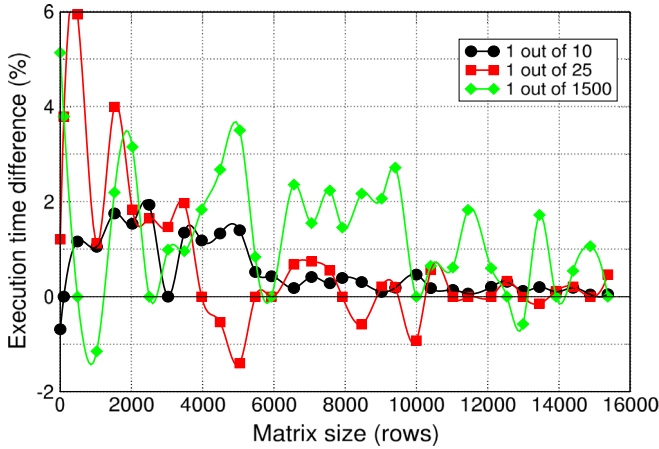


Fig. 3. Runtime difference between Algorithms 2 and 3 on a K40 board as a function of the post-processing rate. The 4th order Runge–Kutta method is on average 2 s faster, which becomes a negligible time as long as the size of the problem increases.

2.3. Series expansion of the evolution operator

Algorithm 2 may be modified in order to make the memory allocation independent of the required precision and slightly reduced with respect to the Runge–Kutta integration method. Upon introducing the evolution operator $\hat{U}_i(\delta t)$, such that

$$|\Psi_i(t + \delta t)\rangle = \hat{U}_i(\delta t)|\Psi_i(t)\rangle,$$

we may rewrite Eq. (2) in terms of $\hat{U}_i(\delta t)$ instead of $|\Psi_i(t)\rangle$, i.e.,

$$i\hbar \frac{d\hat{U}_i(\delta t)}{dt} = \hat{\mathcal{H}}_i \hat{U}_i(\delta t). \quad (5)$$

The formal solution is given by

$$\hat{U}_i(\delta t) = \exp\left(-\frac{i}{\hbar}\delta t \hat{\mathcal{H}}_i\right).$$

Upon expanding $\hat{U}_i(\delta t)$ in Taylor series we have

$$\hat{U}_i(\delta t) = \mathbb{1} + \left(-\frac{i}{\hbar}\delta t \hat{\mathcal{H}}_i\right) + \frac{1}{2}\left(-\frac{i}{\hbar}\delta t \hat{\mathcal{H}}_i\right)^2 + \dots + \frac{1}{n!}\left(-\frac{i}{\hbar}\delta t \hat{\mathcal{H}}_i\right)^n + o\left(\left(-\frac{i}{\hbar}\delta t \hat{\mathcal{H}}_i\right)^n\right); \quad (6)$$

the wave-function can be recast as

$$|\Psi_i(t + \delta t)\rangle = \sum_{j=0}^n |\Phi_i^{(j)}(t)\rangle + o\left(\left(-\frac{i}{\hbar}\delta t \hat{\mathcal{H}}_i\right)^n |\Psi_i(t)\rangle\right), \quad (7)$$

where

$$|\Phi_i^{(0)}(t)\rangle = |\Psi_i(t)\rangle,$$

and

$$|\Phi_i^{(j)}(t)\rangle = -\frac{1}{j} \frac{i}{\hbar} \delta t \hat{\mathcal{H}}_i |\Phi_i^{(j-1)}(t)\rangle.$$

The pseudocode for the evolution of the wave-functions by means of the expansion of the evolution operator in Taylor series is shown in Algorithm 3. In order to understand the similarities and the differences between the two methods, let us remind that the coefficients μ_j in the Runge–Kutta expansion are, in general, determined by an educated fitting of a formal Taylor series expansion of the unknown functions in such a way that the truncation error is the same. Due to the exponential form of the evolution operator, there is also a perfect coincidence between the n th order Taylor series expansion and the n th order Runge–Kutta method [52]. The advantage of the Taylor expansion is represented by the progressive updating of $|\Psi_i(t)\rangle$ with the help of the auxiliary vector $|\Phi_i(t)\rangle$, which is overwritten at each step of the expansion loop. Thus, the memory allocation of auxiliary variables does not depend any more on the precision of the calculation, without increasing the algorithmic complexity. Notice that at the same time, all the arguments discussed in Section 2.2 about the heaviest routines (and the influence of the time step on the results) still hold true.

Algorithm 3 Pseudocode for solving the CTQW dynamics via Taylor series expansion of the evolution operator

```

1: Define Hamiltonian topology
2: Initialize reduced Hamiltonians  $\tilde{\mathcal{H}}_i$ 
3: Initialize switching times
4: while time  $t < t_{\max}$  do
5:   for all realizations do ▷ Begin Parallel/SIMT Section
6:      $|\Phi_i^{(0)}(t)\rangle \leftarrow |\Psi_i(t)\rangle$ 
7:      $|\Psi_i(t + \delta t)\rangle \leftarrow |\Psi_i(t)\rangle$ 
8:     for  $j = 1 \rightarrow n$  do
9:        $|\Phi_i^{(j)}(t)\rangle \leftarrow -\frac{1}{j} \frac{i}{\hbar} \delta t \hat{\mathcal{H}}_i |\Phi_i^{(j-1)}(t)\rangle$ 
10:       $|\Psi_i(t + \delta t)\rangle \leftarrow |\Psi_i(t + \delta t)\rangle + |\Phi_i^{(j)}(t)\rangle$ 
11:    end for
12:    Check norm of  $|\Psi_i(t + \delta t)\rangle$ 
13:    Update switching times
14:     $\tilde{\mathcal{H}}_i \leftarrow \tilde{\mathcal{H}}_i(t + \delta t)$ 
15:  end for ▷ End Parallel/SIMT Section,  $\sim O(RN^m)$ 
16:   $t \leftarrow t + \delta t$ 
17:  if postprocessing then
18:     $\langle \rho(t) \rangle \leftarrow \frac{1}{R} \sum_i |\Psi_i(t)\rangle \langle \Psi_i(t)|$  ▷  $\sim O(RN^{2m})$ 
19:    Post-process  $\langle \rho(t) \rangle$ 
20:  end if
21: end while

```

3. Implementation

Algorithms 1–3 have been implemented to run on multicore shared-memory workstations and graphic accelerators, making use for linear algebra of the BLAS and LAPACK or the cuBLAS and CULA [53] libraries on the host system and on the device, respectively. We have not tackled any advanced memory optimization: as it will be discussed in Section 4, benefits brought in by a highly-optimized code are not expected to further increase the performance gain significantly.

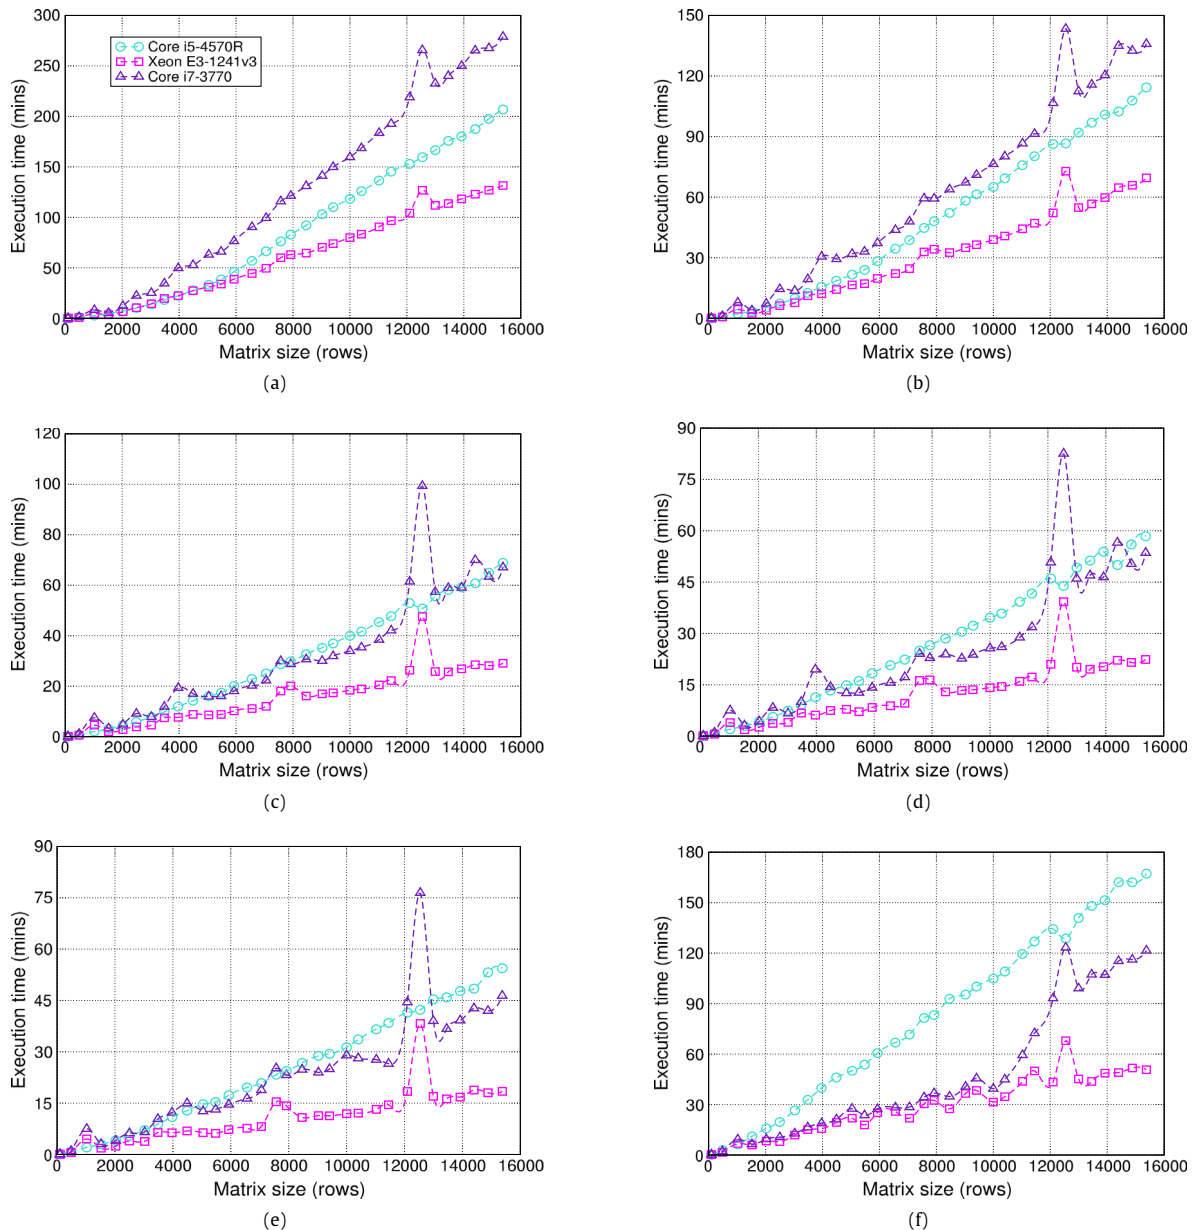


Fig. 4. Execution time of the OPENMP code for different rates of post-processing. 4(a): 1/10 time steps; 4(b): 1/25 time steps; 4(c): 1/100 time steps; 4(d): 1/250 time steps; 4(e) and 4(f): 1/1500 time steps. Panel 4(f) refers to the execution time of the single-core case.

As far as Algorithm 1 is concerned, we envisage two workflows for parallel execution. On the one hand, would memory not be an issue, one can split realizations among non-communicating cores in such a way that, even though the single realization is serialized, a number of realizations are handled at the same time. On the other hand, it may be convenient to serialize realizations and decrease the single-realization running time by spreading the matrix diagonalizations and the matrix-matrix products on multiple cooperating cores. In principle, the latter solution can be pushed farther if a large number of computing cores are available to the programmer, as it is the case of GPUs.

The execution times required by the diagonalization of symmetric matrices with single precision data (`ssyev` function of the Intel MKL 11.2 library) have preliminarily been measured for 3 Intel processors, then the outcomes have been projected over 10^6 calls, which is the typical number of diagonalizations required for the problem in hand. As shown in Fig. 2, a simulation may last for years, which is a virtually infinite time for a computational physics

problem. According to CULA white-papers [53], the corresponding routine ported to GPUs may achieve a speedup ranging from 3x to 10x, a condition that still prevents any investigation from completing within an affordable time.

Algorithms 2 and 3 have been implemented by means of 15 kernels directly on the GPU, then the corresponding OPENMP versions have been derived by replacing kernel invocations with loops. This approach allows for a direct execution time comparison since the number of floating-point operations is basically the same between host and device execution.

The two algorithms share the same 4-stage workflow (1. *initialization*; 2. *wave-function evolution*; 3. *Hamiltonian update*; 4. *density-matrix calculation and post-processing*) and approximately 90% of the code. Contrarily to Algorithm 1, where the limiting factor is primarily represented by time, the limiting factor of Algorithms 2 and 3 is given by the memory required to store the (symmetric, complex) density-matrix $\langle \rho(t) \rangle$ and the wave-functions $|\Psi_i(t)\rangle$. Top level, high-performance solutions for GPGPU

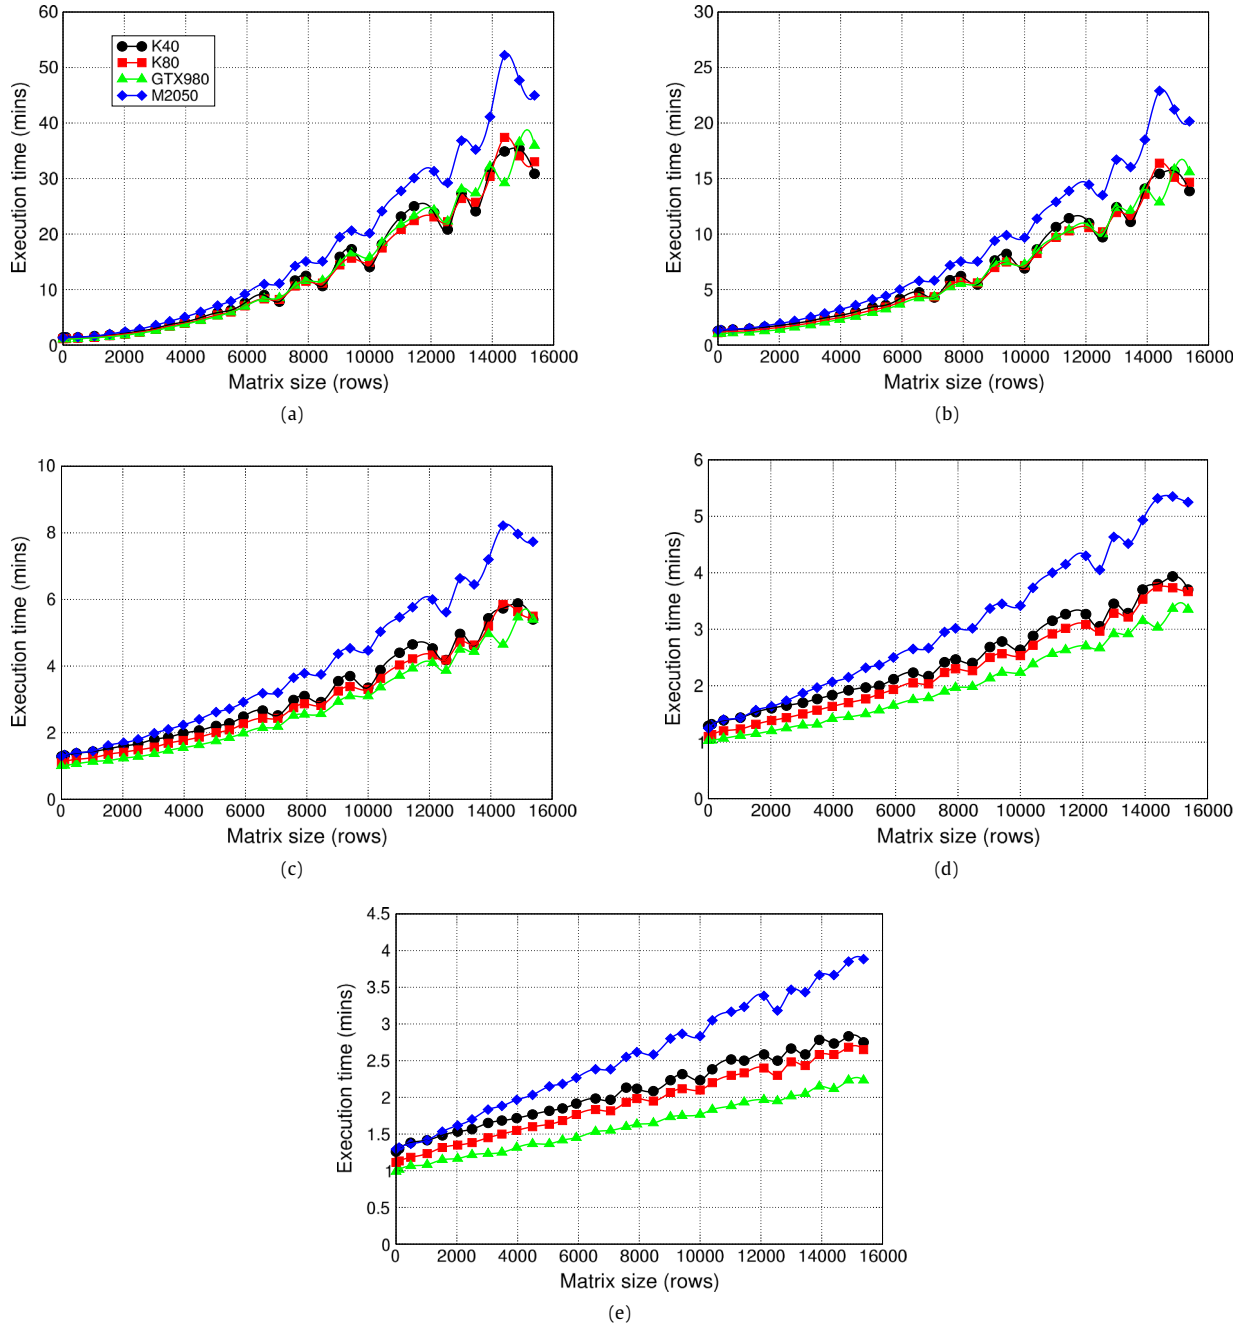


Fig. 5. Execution time of the GPU code for different rates of post-processing. **5(a):** 1/10 time steps; **5(b):** 1/25 time steps; **5(c):** 1/100 time steps; **5(d):** 1/250 time steps; **5(e):** 1/1500 time steps.

computing like NVIDIA Tesla K80 offer up to 24 GB of GPU-RAM, which can cap the maximum size around 51 000 rows (e.g., $q = 2$, $m = 2$, $N_1 \cdot N_2 = 225$).

4. Performance evaluation

In order to evaluate the performance of Algorithms 2 and 3 we tested the case of unidimensional, 2-particle, nearest-neighbor CTQWs with periodic boundary conditions (i.e., $q = 1$, $m = 2$, $k = 2$) and random noise on the tunneling energies. Simulations of 1500 time steps for $R = 1000$ realizations, with different rates for post-processing (from 1 out of 1500 to 1 out of 10 time steps) have been run on the following hardware:

- **Intel CPU:** Core i5-4570R @ 2.7 GHz and 8 GB RAM (4 cores), OS X 10.10.5
- **Intel CPU:** Core i7-3770 @ 3.4 GHz and 4 GB RAM (4 cores), 64-bit Linux OS
- **Intel CPU:** Xeon E3-1241 v3 @ 3.5 GHz and 16 GB RAM (4 cores), 64-bit Linux OS
- **NVIDIA GPU:** Tesla M2050 with 3 GB VRAM, ECC enabled, Compute capability 2.0, CUDA Toolkit 5.0
- **NVIDIA GPU:** Tesla K40 with 12 GB VRAM, ECC enabled, Compute capability 3.5, CUDA Toolkit 7.5
- **NVIDIA GPU:** Tesla K80 with 24 GB VRAM, ECC enabled, Compute capability 3.7, CUDA Toolkit 7.5
- **NVIDIA GPU:** GeForce GTX980 with 4 GB VRAM, no ECC, Compute capability 5.2, CUDA Toolkit 6.5.

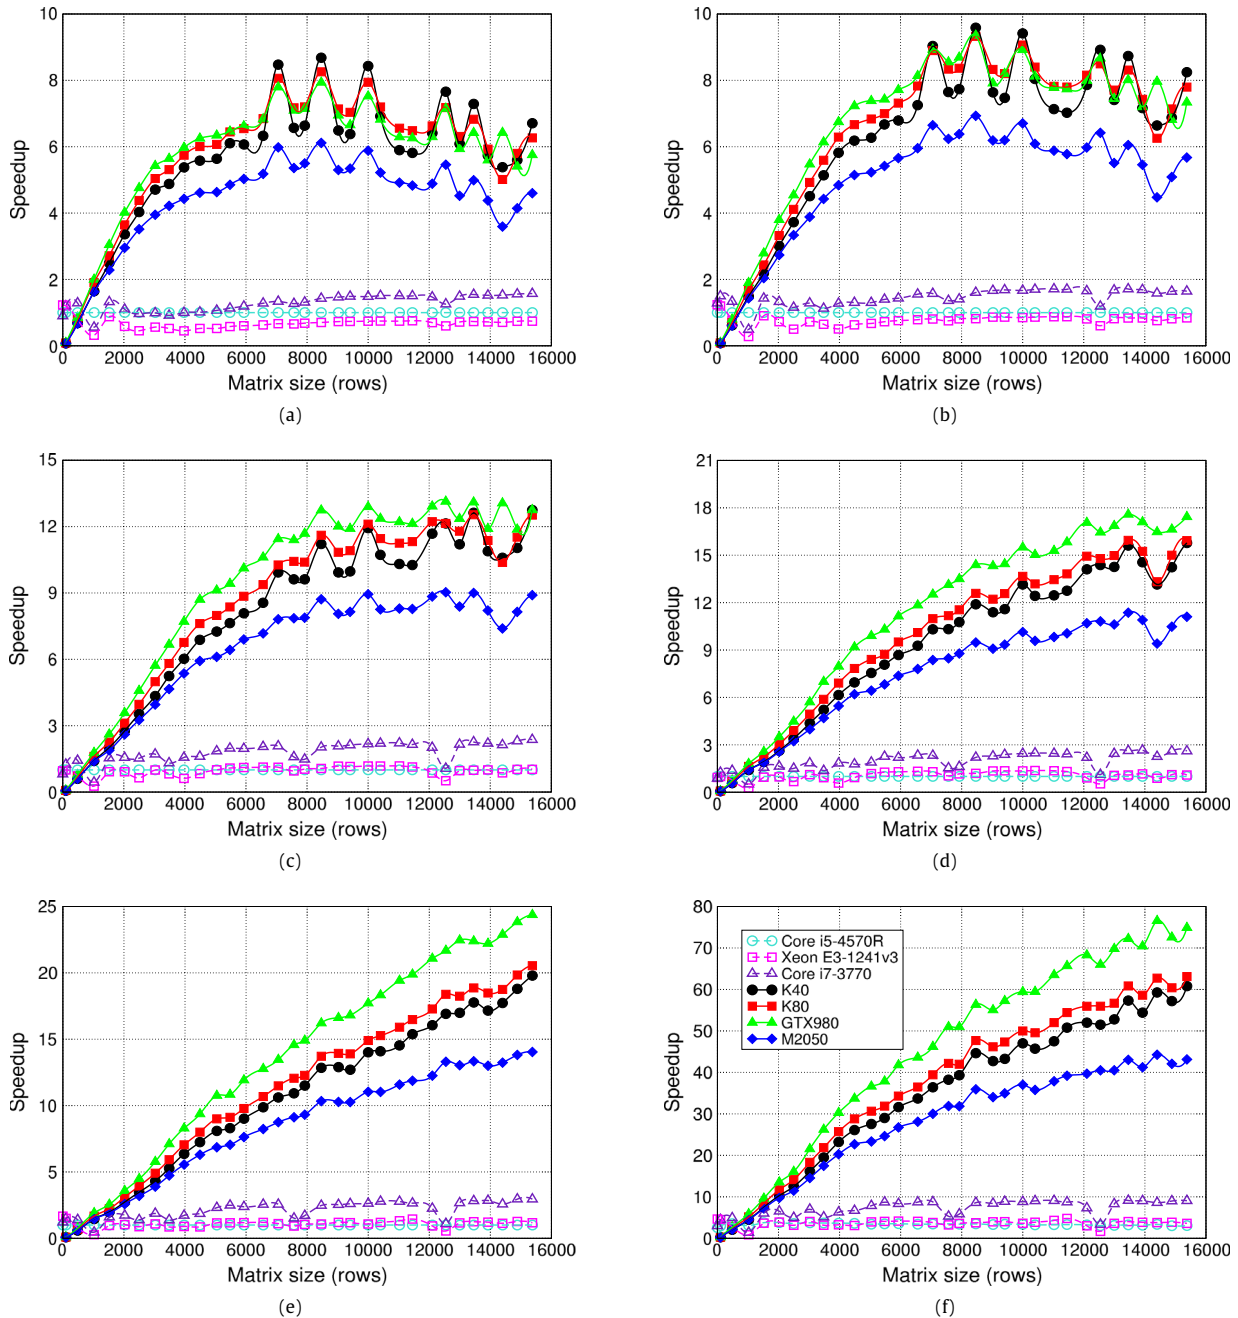


Fig. 6. Performance comparison for different rates of post-processing. 6(a): 1/10 time steps; 6(b): 1/25 time steps; 6(c): 1/100 time steps; 6(d): 1/250 time steps; 6(e) and 6(f): 1/1500 time steps. In panel 6(f) the comparison at post-processing rate 1/5000 time steps refers to the single-core execution.

The OPENMP source code has been compiled with the Intel C++ Compiler (ICC) version 15.0.3 for Linux and version 15.0.7 for OS X; the CUDA source code with the NVIDIA CUDA Compiler (NVCC), with no further optimizations other than those provided by default. Preliminary runs on GPUs proved that 256 threads per block maximize the efficiency.

The execution times of the 4th order Runge–Kutta and of the series expansion methods are basically the same. Depending on the hardware, very few seconds in favor of one algorithm or the other are reported; differences become negligible as long as the size of the mesh increases (see Fig. 3 for tests performed on the Tesla K40). Therefore, we proceed in the analysis only with Algorithm 3 and assume that the same conclusions hold true also for Algorithm 2.

Fig. 4 illustrates the execution time of the Series expansion algorithms as a function of the problem size for the three CPUs

under test. The running times for Core-i7 and Xeon E3 processors exhibit a similar qualitative behavior: clear performance losses around 7500 and 12500 rows, and, in general, a very similar shape of the execution time vs. matrix size curve. These evidences are lacking in the Core-i5 case, where the execution time steadily increases as a function of the matrix size without any particular gain or loss. Since the main difference between the Core-i5 and the Core-i7/Xeon E3 cases regards the compiler and the associated mathematical library (ICC 15.0.3 for Linux instead of ICC 15.0.7 for OS X), we attribute the underperformance to a failing code optimization or bad memory handling specific of the compiler-mathematical library bundle.

The analogous execution time comparison for the codes running on the four GPUs is shown in Fig. 5. All tests are completed in less than 1 h, with running times very similar to each other.

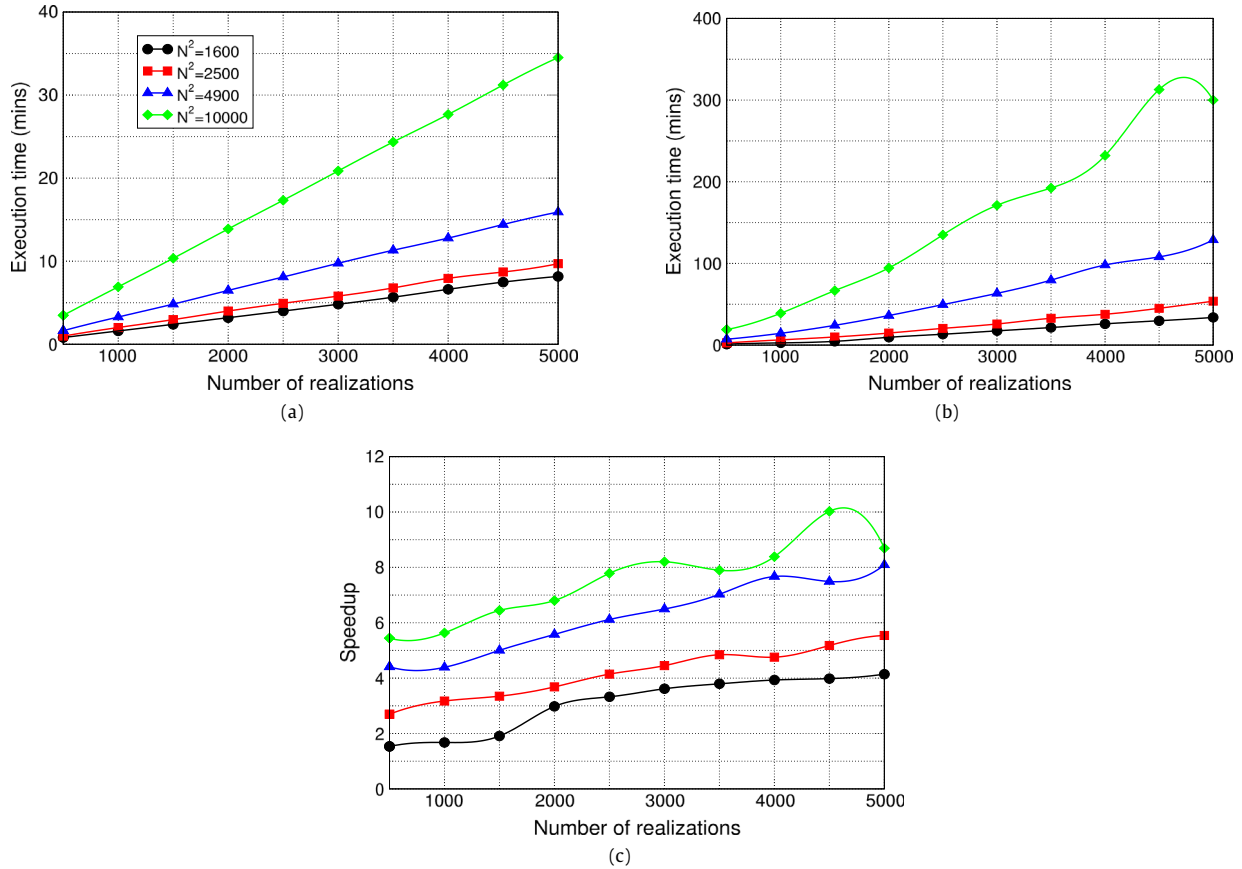


Fig. 7. Scaling behavior on Tesla K40 GPU 7(a) and Xeon E3 CPU 7(b) for different sizes of the mesh (from 1600 to 10000 rows), as a function of the number of realizations. The corresponding speedup is illustrated in panel 7(c).

The only exception is represented by the Tesla M2050 board that underperforms its competitors, though retaining a substantial gain over any OPENMP execution. In general, the fastest runs are achieved with the GeForce GTX980 board thanks to a superior clock rate. Notice that GeForce boards are not certified for GPGPU computing due to lack of ECC memory, and uncontrolled bit-flips or erratic bits in the memory locations devoted to the storage of $|\Psi_i(t)\rangle$ can jeopardize the reliability of the outcomes of the simulations. Though uncommon, this aspect deserves care and double checks are mandatory in presence of odd results.

In order to provide an overall comparative review of the performances, we chose the Core-i5 as the reference processor and we calculate the simulation speedup as

$$\text{Speedup} = \frac{\text{CPU or GPU-under-test execution time}}{\text{Core i5-4570R 4-thread execution time}}. \quad (8)$$

Data are shown in Fig. 6. GPU implementation becomes efficient, i.e., with a speedup greater than 1, roughly about 1000 rows, when the workload starts to fill completely the computational power of the GPUs, and the higher clock rate of the CPU does not compensate any longer for the reduced number of computing units. Since a matrix size greater than 1000 rows is a very common case for many-particle CTQWs (i.e., a lattice as small as $N \sim 32$ for the two-particle case), GPU computing sounds a viable and efficient option to pursue in order to reduce the execution time down to the minute-to-few-hour range. It is important to stress that the simulation speedup strongly depends on the post-processing rate. For an output generation as frequent as 1 out of 10 time steps (panel a) a gain about 5x-7x is obtained; the gain rises up to 8x-9x for an average post-processing rate of 1 out of 25 time steps (panel b) and up to a 10x-13x for a moderate output generation around

1 out of 100 time steps (panel c). Panels d and e refer instead to cases where the calculation of the density matrix is progressively reduced down to a single time per simulation. In other words, this is the speedup achievable for the pure evolution of the wave-functions, which settles in the 20x range and more. By comparing data reported in panels 4(e) and 4(f) of Fig. 4, the OPENMP parallelization introduces a further 2.5-3x gain with respect to the single-core execution, this boosting up the speedup at a minimum gain around 60x for the pure evolution of the wave functions, as shown in Fig. 6(f), and around 15x when a high post-processing frequency is required.

The speedup depends also on the number of realizations considered during parallel execution. About a +3x gain is observed when the number of realizations increases from 500 to 5000 (Fig. 7), irrespectively of the size of the mesh. While the GPU codes scale with the number of realizations (as should be according to the discussion of Sections 2.2 and 2.3), a performance loss is found for the OPENMP implementation. For the sake of truth, we recall that the OPENMP code was derived from the CUDA code with the strict constraint of adhering as much as possible to it and allowing a fair direct comparison of the computational burden, without introducing any further memory or algorithmic optimizations. Since the performance loss does not significantly depend on the size of the problem in hand, but only to the number of realizations, this poor behavior can primarily be ascribed to the larger number of calls to the BLAS functions.

The influence of the post-processing rate in GPU execution is even more evident from the shape of the curves of Fig. 5 that changes from parabolic to linear. Though not immediate at first sight, the same also applies for the curves of Fig. 4 and is validated by numerical regression. Further information stem from

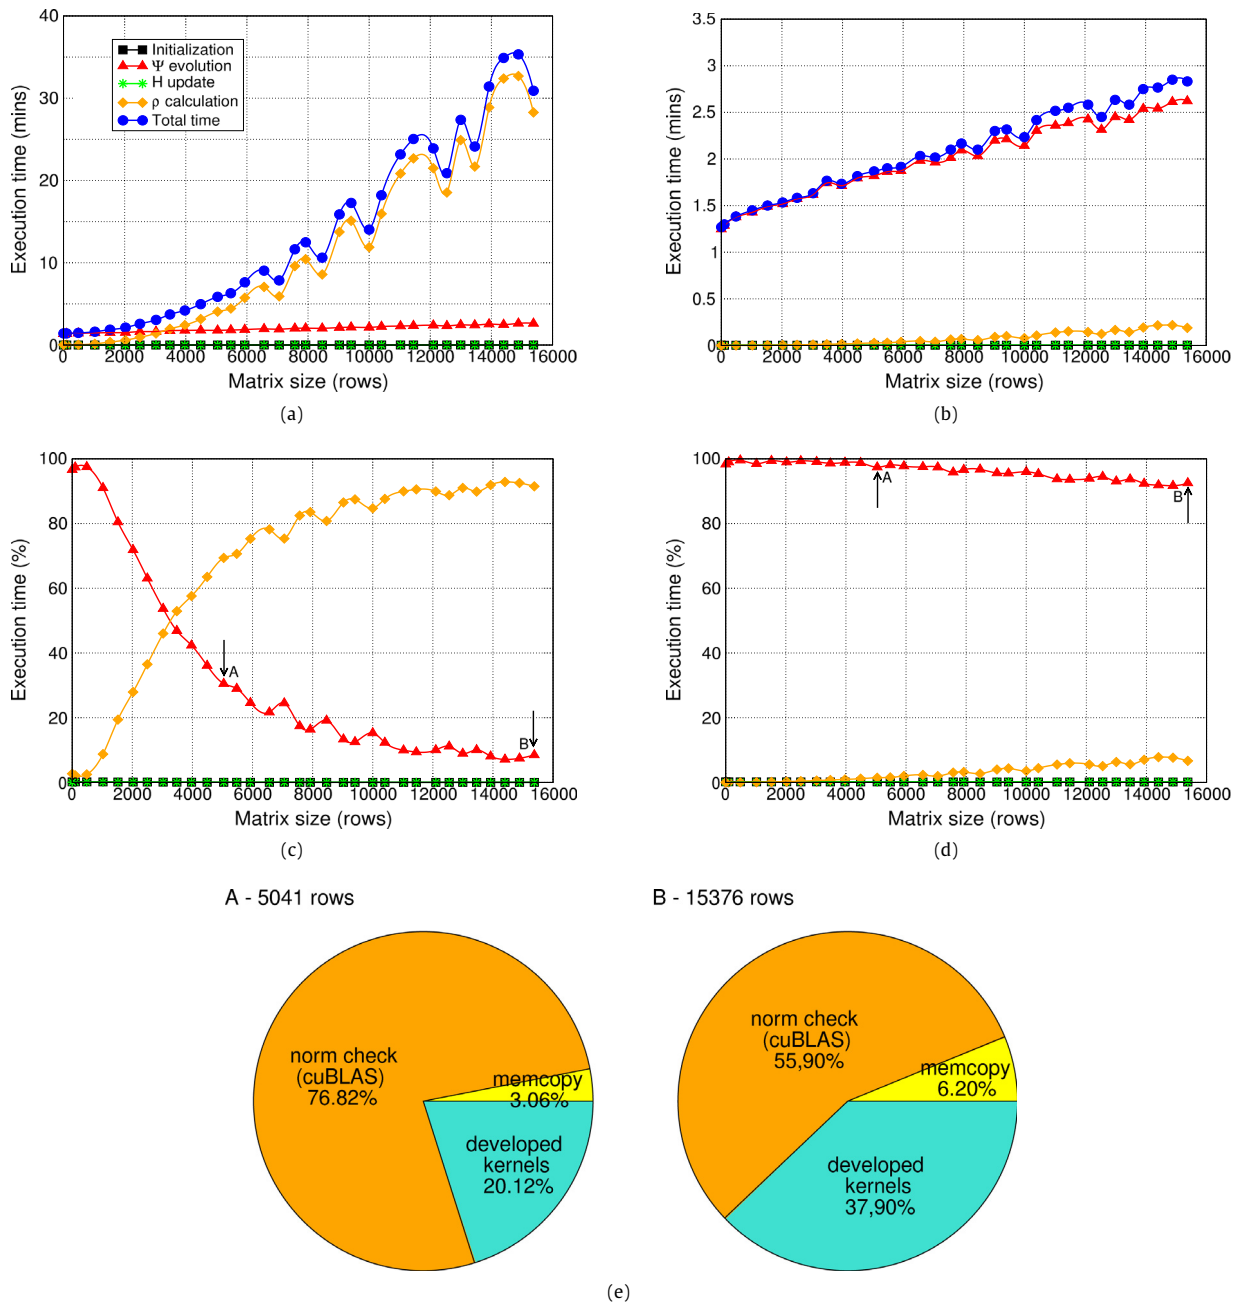


Fig. 8. Code profiling and relative weight of the four execution stages for a frequent (8(a) and 8(c)) and for the tiniest (8(b) and 8(d)) output generation rate. The pie-charts 8(e) show the time required by the sub-components of the wave-function evolution stage for matrix sizes identified by letters A and B. No substantial difference is found between the two cases.

code profiling. We have tracked the execution time of the four stages composing the software for the two opposite cases of very frequent and tiny output generation on the K40 board (Fig. 8): the *initialization* and the *Hamiltonian update* stages contribute with a negligible running time (less than 0.3% in total), while the *wave-function evolution* and the *density-matrix calculation and post-processing* stages largely prevail.

The running time in case of a very limited output generation is substantially dictated by the *wave-function evolution* stage, which grows linearly with the size of the problem as discussed in Section 2. On the contrary, in case of a frequent output generation, the heaviest stage is represented by *density-matrix calculation and post-processing*, whose influence quickly grows up and saturates about 90% of the total execution time. Going into details, more than the 99.3% of the time spent for post-processing is required

by the *cuBLASCher* library function that builds up the average density-matrix ($\rho(t)$). As a consequence, the peaks in panels a, b and c of Fig. 6 are due to outperforming conditions of the *cuBLAS* library. Also in the *wave-function evolution* stage most of the time is spent in calls to system or library functions (see Fig. 8(e)). As a matter of fact, even for large matrices (i.e., $N^m > 10\,000$), only up to approximately one third of the time is dedicated to the series expansion, whereas the remainder is due to device-to-device memcopy and norm evaluation (*cuBLASscnrm2* and *cuBLAScscal*). As we already pointed out, memory optimization for speed using, e.g., shared memory on the device, was not a goal of the present work. From the time profiling above, we do not believe it worth the effort: highly-optimized solutions able to cut the execution time of the computational kernels by a factor of 2 or 3 would only bring a very modest benefit around 1 min or less.

To obtain a further significant speedup it is instead mandatory to implement new kernels for linear algebra, other than those provided by the cuBLAS library.

5. Conclusions

The availability of a simulation tool for evolving many-particle CTQWs in a noisy environment represents a crucial prerequisite for the investigation of quantum many-body systems and for the implementation of effective quantum algorithms in realistic situations. In essence, the dynamics of a many-particle state over a noisy lattice can be associated with the solution of a set of stochastic differential equations. However, the need to post-process a large number of data in order to achieve information for any measurable quantity makes the problem much more resource-demanding. In fact, as long as the number of particles and/or the dimensionality of the domain increase, limiting factors such as the memory occupancy and the time required to run the simulations quickly become very challenging issues and determine whether a simulation scheme can or cannot provide results within the available computational power.

Though numerically accurate, the standard Hamiltonian diagonalization method is not feasible even for small systems and alternative numerical solutions must be sought. Among them, we have shown that the 4th order Runge–Kutta integration method and the Taylor-series expansion of the evolution operator have a low computational cost and provide reliable data. Moreover, they are highly parallelizable within the SIMT paradigm, and this allows the straightforward, direct implementation on GPUs.

After developing the codes, we have benchmarked four NVIDIA GPUs and three quad-core Intel CPUs for a 2-particle system over a lattice of increasing dimensions. GPU execution enables significant cuts of the running time of batches of thousands of simulations down to the minute-to-few-hour range. The speedup with respect to OPENMP parallelization stays in the range from 8x to more than 20x, depending on the frequency of post-processing. Our results show that GPU-accelerated codes allow one to overcome concerns about the execution time and make it possible to design simulations involving many particles or large lattices, whose only limit is dictated by the memory available on the device.

Acknowledgment

This work has been supported by EU through the Collaborative Project QuProCS (Grant Agreement 641277), by UniMI through the H2020 Transition Grant 15-6-3008000-625, and by UniMoRe through FAR2014.

References

- [1] Y. Aharonov, L. Davidovich, N. Zagury, *Phys. Rev. A* 48 (1993) 1687–1690.
- [2] E. Farhi, S. Gutmann, *Phys. Rev. A* 58 (1998) 915–928.
- [3] Y. Omar, N. Paunković, L. Sheridan, S. Bose, *Phys. Rev. A* 74 (2006) 042304.
- [4] Y. Lahini, M. Verbin, S.D. Huber, Y. Bromberg, R. Pugatch, Y. Silberberg, *Phys. Rev. A* 86 (2012) 011603.
- [5] C. Benedetti, F. Buscemi, P. Bordone, *Phys. Rev. A* 85 (2012) 042314.
- [6] L. Wang, L. Wang, Y. Zhang, *Phys. Rev. A* 90 (2014) 063618.
- [7] X. Qin, Y. Ke, X. Guan, Z. Li, N. Andrei, C. Lee, *Phys. Rev. A* 90 (2014) 062301.
- [8] M. Faccin, T. Johnson, J. Biamonte, S. Kais, P. Migdał, *Phys. Rev. X* 3 (2013) 041007.
- [9] F. Caruso, A. Crespi, A.G. Ciriolo, F. Sciarrino, R. Osellame, *Nat. Comm.* 7 (2016) 11682.
- [10] A. Schreiber, A. Gábris, P.P. Rohde, K. Laiho, M. Štefaňák, V. Potoček, C. Hamilton, I. Jex, C. Silberhorn, *Science* 336 (2012) 55–58.
- [11] A. Makmal, M. Tiersch, C. Ganahl, H.J. Briegel, *Phys. Rev. A* 93 (2016) 022322.
- [12] O. Mülken, A. Blumen, *Phys. Rep.* 502 (2011) 37–87.
- [13] A.M. Childs, J. Goldstone, *Phys. Rev. A* 70 (2004) 022314.
- [14] S. Chakraborty, L. Novo, A. Ambainis, Y. Omar, *Phys. Rev. Lett.* 116 (2016) 100501.
- [15] B.L. Douglas, J.B. Wang, *J. Phys. A* 41 (2008) 075303.
- [16] J.K. Gamble, M. Friesen, D. Zhou, R. Joynt, S.N. Coppersmith, *Phys. Rev. A* 81 (2010) 052313.
- [17] S.D. Berry, J.B. Wang, *Phys. Rev. A* 83 (2011) 042317.
- [18] A.M. Childs, *Phys. Rev. Lett.* 102 (2009) 180501.
- [19] A.M. Childs, D. Gosset, Z. Webb, *Science* 339 (2013) 791–794.
- [20] Y. Lahini, Y. Bromberg, D.N. Christodoulides, Y. Silberberg, *Phys. Rev. Lett.* 105 (2010) 163905.
- [21] A. Schreiber, K.N. Cassemiro, V. Potoček, A. Gábris, I. Jex, Ch. Silberhorn, *Phys. Rev. Lett.* 106 (2011) 180403.
- [22] A. Crespi, R. Osellame, R. Ramponi, V. Giovannetti, R. Fazio, L. Sansoni, F. De Nicola, F. Sciarrino, P. Mataloni, *Nat. Phot.* 7 (2013) 322–328.
- [23] J. Ghosh, *Phys. Rev. A* 89 (2014) 022309.
- [24] F. De Nicola, L. Sansoni, A. Crespi, R. Ramponi, R. Osellame, V. Giovannetti, R. Fazio, P. Mataloni, F. Sciarrino, *Phys. Rev. A* 89 (2014) 032322.
- [25] C. Benedetti, F. Buscemi, P. Bordone, M.G.A. Paris, *Phys. Rev. A* 93 (2016) 042313.
- [26] I. Siloi, C. Benedetti, E. Piccinini, J. Piilo, S. Maniscalco, M.G.A. Paris, P. Bordone, *Phys. Rev. A* 95 (2017) 022106.
- [27] A. Beggi, F. Buscemi, P. Bordone, *Quantum Inf. Process.* 15 (2016) 3711–3743.
- [28] C. Lee, A. Rai, C. Noh, D.G. Angelakis, *Phys. Rev. A* 89 (2014) 023823.
- [29] P. Hänggi, H. Thomas, *Phys. Rep.* 88 (1982) 207–319.
- [30] J.A. Izaac, J.B. Wang, *Comput. Phys. Comm.* 186 (2015) 81–92.
- [31] A.G. Anderson, W.A. Goddard III, P. Schröder, *Comput. Phys. Comm.* 177 (2007) 298–306.
- [32] J.A. Anderson, C.D. Lorenz, A. Travesset, *J. Comput. Phys.* 227 (2008) 5342–5359.
- [33] J. Tölke, M. Krafczyk, *Int. J. Comput. Fluid Dyn.* 22 (2008) 443–456.
- [34] T. Preis, P. Vírna, W. Paul, J.J. Schneider, *J. Comput. Phys.* 228 (2009) 4468–4477.
- [35] M. Januszewski, M. Kostur, *Comput. Phys. Comm.* 181 (2010) 183–188.
- [36] J. Spiechowicz, M. Kostur, L. Machura, *Comput. Phys. Comm.* 191 (2015) 140–149.
- [37] D. Tamascelli, F.S. Dambrosio, R. Conte, M. Ceotto, *J. Chem. Phys.* 140 (2014) 174109.
- [38] J. Glaser, T.D. Nguyen, J.A. Anderson, P. Lui, F. Spiga, J.A. Millan, D.C. Morse, S.C. Glotzer, *Comput. Phys. Comm.* 192 (2015) 97–107.
- [39] L.S. Smith, Q. Liang, *Comput. & Fluids* 88 (2013) 334–343.
- [40] M. Januszewski, M. Kostur, *Comput. Phys. Comm.* 185 (2014) 2350–2368.
- [41] J.A. Anderson, M.E. Irrgang, S.C. Glotzer, *Comput. Phys. Comm.* 204 (2016) 21–30.
- [42] Y. Lutsyshyn, *Comput. Phys. Comm.* 187 (2015) 162–174.
- [43] M. Januszewski, A. Ptok, D. Crivelli, B. Gardas, *Comput. Phys. Comm.* 192 (2015) 220–227.
- [44] M.D. Feit, J.A. Fleck Jr., A. Steiger, *J. Comput. Phys.* 47 (1982) 412–433.
- [45] H. De Raedt, *Comput. Phys. Rep.* 7 (1987) 1–72.
- [46] J.L. Richardson, *Comput. Phys. Comm.* 63 (1991) 84–94.
- [47] C. Leforestier, R.H. Bisseling, C. Cerjan, M.D. Feit, R. Friesner, A. Gulberg, A. Hammerich, G. Jolicard, W. Karrlein, H.-D. Meyer, N. Lipkin, O. Roncero, R. Kosloff, *J. Comput. Phys.* 94 (1991) 59–80.
- [48] K. Winkler, G. Thalhammer, F. Lang, R. Grimm, J.H. Denschlag, A. Daley, A. Kantian, H. Buchler, P. Zoller, *Nature* 441 (2006) 853.
- [49] M. Valiente, D. Petrosyan, *J. Phys. B* 41 (2008) 161002.
- [50] F. Hofmann, M. Potthoff, *Phys. Rev. B* 85 (2012) 205127.
- [51] D. Jaksch, P. Zoller, *Ann. Phys.* 315 (2005) 52–79.
- [52] J. Butcher, *Numerical Methods for Ordinary Differential Equations*, second ed., Wiley, Chichester, 2008.
- [53] J.R. Humphrey, D.K. Price, K.E. Spagnoli, A.L. Paolini, E.J. Kelmelis, *Proc. SPIE* 7705 (2010) 770502.