

A Monte Carlo method for computing the action of a matrix exponential on a vector

Juan A. Acebrón^{a,b}

^a*Dept. Information Science and Technology, ISCTE-University Institute of Lisbon, Portugal*

^b*INESC-ID, Instituto Superior Técnico, Universidade de Lisboa, Portugal*

Abstract

A Monte Carlo method for computing the action of a matrix exponential for a certain class of matrices on a vector is proposed. The method is based on generating random paths, which evolve through the indices of the matrix, governed by a given continuous-time Markov chain. The vector solution is computed probabilistically by averaging over a suitable multiplicative functional. This representation extends the existing linear algebra Monte Carlo-based methods, and was used in practice to develop an efficient algorithm capable of computing both, a single entry or the full vector solution. Finally, several relevant benchmarks were executed to assess the performance of the algorithm. A comparison with the results obtained with a Krylov-based method shows the remarkable performance of the algorithm for solving large-scale problems.

Key words: Monte Carlo methods, matrix functions, network analysis, communicability

PACS: 65C05, 65C20, 65N55, 65M75, 65Y20

1 Introduction

Computing the action of a matrix function on a vector is experiencing these days a reborn interest. This is not because of the absence of relevant applications in science and engineering in the past, but rather because the improvement in the numerical methods, along with the advent of highly massive parallel computers, now allows one to be able to attack more realistic problems on a large scale, far beyond the merely academic problems capable of

Email address: juan.acebron@iscte-iul.pt (Juan A. Acebrón).

being simulated in the past. These applications include circuit simulations [36]; power grid simulation [37,33]; nuclear reaction simulations [31]; analysis of transient solutions in Markov chains [34]; simulations of quantum systems [32]; numerical solution of partial differential equations (PDEs) [24,27]; and analysis of complex networks [5].

Circuit and power grid simulation play an important role during the design of integrated circuits, being in general a heavy computationally task of the whole design process. In the analysis of a reactor fuel, a computational heavy task of the analysis consists in solving the burnup equations describing the rates of the concentration of the different nuclides of the nuclear fuel. Computing the action of a matrix exponential over the initial state appears as an important numerical alternative among the different available techniques, such as integrating the Chapman-Kolmogorov system of differential equations for obtaining the transient solution of homogeneous irreducible Markov chains. In the field of partial differential equations, numerically solving a boundary-value PDE problem by means of the method of lines requires in practice to compute the action of a matrix exponential over the initial condition. Finally, in network analysis, determining some important metrics of the network, such as for instance the total communicability which characterizes the importance of the nodes inside the network, entails computing the exponential of the adjacency matrix of the network.

Over the last few years many numerical methods have been proposed for computing the action of a matrix function over a vector. There are already excellent reviews in the literature describing the different numerical methods proposed so far (see [10,20,22,23], e.g.); therefore it is not intended to go into any details here. Instead we will briefly describe some of them for those readers not familiar with the topic. Essentially we can classify these methods as follows: Krylov-based subspace methods, contour integration methods, ordinary differential equation methods, and polynomial or rational methods. One of the most studied methods in theory, and used in practice, are those based on the Krylov subspace method. The idea behind the method consists in projecting the given (typically large) matrix onto a Krylov subspace. For the specific case of an exponential function, through a basis of the subspace constructed using the Arnoldi process, the exponential of the projected matrix is computed by using a standard technique based typically on the squaring and scaling method [22].

The idea of using probabilistic methods based on Monte Carlo simulations for computing functions of matrices goes back to the pioneering work of von Neumann and Ulam during the 1940's [18]. Although initially the method was proposed merely for computing the inverse of a matrix, it was later generalized for solving linear algebra problems in a series of seminal papers, see [13] e.g., and [12] for further references. Briefly the underlying idea consists in

generating a discrete Markov chain which evolves by random paths through the different indices of the matrix. Mathematically, the method can be seen in a way as a Monte Carlo sampling of the Neumann series of the matrix. The convergence of the method was rigorously established in [26], and improved further more recently (see for instance [15], and [7] just to cite a few references). More recently, and for the specific case of computing the action of a Hermitian matrix exponential over a vector, which is of interest in Quantum Physics, it has been proposed in [32] an efficient algorithm based on a novel randomized linear algebra technique known in the literature as the Nyström method.

These probabilistic methods offer important computational advantages. Furthermore, the algorithms are much simpler to code than their deterministic counterparts, which impact positively in promoting an easy further optimization of the codes; it turns out that they are specially suited for parallel computing. This is because the solution is often computed through an expectation value of a given finite sample, the simulations are independent from each other. This is of paramount importance because it allows for the development of parallel codes with extremely low communication overhead among processors, and has a positive effect on properties such as scalability and fault-tolerance. For the parallel implementation of the Monte Carlo method for solving linear algebra problems see [14] e.g.

Another important advantage of the probabilistic methods consists in the feasibility of computing the solution of the problem at specific chosen points, without the need for solving globally the entire problem. This remarkable feature offers important advantages in dealing with some specific applications found in science and engineering, and it has been explored for efficiently solving continuous problems such as boundary-value problems for PDEs in [1–3], and references therein. However, this is not an exclusive advantage of the probabilistic methods. In fact, for the specific problem of matrix functions, it has been proposed in the literature several methods [19] capable also of estimating individual entries of the matrix function. The main idea consists in applying several quadrature rules along with a single iteration step of the Lanczos algorithm to obtain *a priori* lower and upper bounds. Moreover the bounds can be further improved *a posteriori* using several iterations of the Lanczos algorithm in the quadrature formula. This idea has been applied successfully to the problem of estimating different metrics of complex networks in [8]. The computational cost has been estimated to grow linearly with the matrix size in the best case, since for each iteration of the Lanczos algorithm is required to compute a matrix-vector multiplication.

The purpose of this paper is to extend the existing aforementioned Monte Carlo methods for dealing with other functions of matrices, such as the matrix exponential, and more specifically for the problem of computing the action of

a matrix exponential on a vector for a certain class of matrices. This is done by resorting to a probabilistic representation of the vector solution based on generating random paths corresponding to samples of a suitable continuous-time Markov chain. The convergence of the method was conveniently analyzed, as well as the computational cost estimated. In addition, several relevant numerical examples, extracted from network analysis are given, focusing on both, the accuracy and performance of the method.

The paper is organized as follows. The probabilistic representation of the vector solution is presented in Sec. 2. In Sec. 3, it is explained how the probabilistic representation can be implemented in practice. Secs. 4 and 5 are devoted to the analysis of the computational cost of the algorithm, and the associated numerical errors of the method, respectively. Finally in Sec. 6 several benchmarks are executed to assess the performance of the method in comparison with the performance obtained by the classical Krylov-based method. To conclude we summarize the main results and discuss potential directions for future research.

2 The numerical method

Let $A = \{a_{ij}\}_{i,j=1}^n$ a given sparse n -by- n symmetric matrix, v an n -dimensional vector, and x an n -dimensional vector solution of evaluating $e^{\beta A} v$. We assume that A can be decomposed as $A = D - L$, where L is the Laplacian matrix symmetric and irreducible [28], and D a diagonal matrix with entries $d_i, i = 1, \dots, n$. Since in general both matrices do not commute, it does not hold that $e^{\beta A} v = e^{-\beta L} e^{\beta D} v$. However, an approximation of the matrix exponential can be easily obtained by resorting to the exponential Lie splitting, and yields

$$\bar{x}^L = (e^{-\Delta t L} e^{\Delta t D})^N v \approx e^{\beta A} v. \quad (1)$$

Here $\Delta t = \beta/N$, and in the following for convenience it will be termed as the time step. It is known that the local error $\varepsilon_L = x - \bar{x}^L$ of the Lie splitting after one time step is given by

$$\varepsilon_L = \frac{\Delta t^2}{2} [D, L] v + O(\Delta t^3), \quad (2)$$

being in general of order of $O(\Delta t)$ for the global error. A higher order approximation does exist, and in view of being the matrix D diagonal, it can be computed without any additional computational cost. In fact, the well known Strang splitting yields,

$$\bar{x}^S = \left(e^{\Delta t D/2} e^{-\Delta t L} e^{\Delta t D/2} \right)^N v. \quad (3)$$

The local error after one time step $\varepsilon_S = x - \bar{x}^S$ of this approximation is known [25] to be

$$\varepsilon_S = \Delta t^3 \left(\frac{1}{12} [D, [D, L]] - \frac{1}{24} [L, [L, D]] \right) v + O(\Delta t^4), \quad (4)$$

and globally of order of $O(\Delta t^2)$. The next lemma will be used to derive a probabilistic representation for the vector solution \bar{x}^S , in the following denoted as \bar{x} for simplicity. To this purpose we need first to prove the following useful fact for the partial solution $e^{\Delta t D/2} e^{-\Delta t L} e^{\Delta t D/2} v$.

Lemma 1 *Assume j is a discrete random variable that takes values on $S = \{1, 2, \dots, n\}$ with probability $p_{ij}(t)$ given by the transition probabilities of a continuous-time Markov chain generated by the infinitesimal generator $Q = -(L)_{ij}$ and evaluated at time Δt . Then, any entry i of the vector*

$$y = e^{\Delta t D/2} e^{-\Delta t L} e^{\Delta t D/2} v, \quad (5)$$

can be represented as $y_i = e^{\Delta t d_i/2} \mathbf{E}[\eta]$, with $\eta = e^{\Delta t d_j/2} v_j$, and $\mathbf{E}[\eta]$ its expected value.

Proof. Since D is a diagonal matrix, y_i can be computed as follows

$$y_i = \sum_{j=1}^n e^{\Delta t d_i/2} (e^{-\Delta t L})_{ij} e^{\Delta t d_j/2} v_j. \quad (6)$$

By the definition of the unnormalized Laplacian matrix of a graph G , $L(G)$ is a matrix with diagonal elements L_{ii} equal to the degree of each vertex d_i , and the off-diagonal L_{ij} , -1 if (i, j) is an edge, or 0 otherwise. Therefore, it follows that $\sum_j L_{ij} = 0$, and $L_{ii} > 0$, and hence the matrix $Q = -L$ can be assumed to be a generator of a suitable continuous-time Markov chain on the state space $S = \{1, 2, \dots, n\}$. Then,

$$y_i = e^{\Delta t d_i/2} \sum_{j=1}^n p_{ij}(\Delta t) e^{\Delta t d_j/2} v_j, \quad (7)$$

where $p_{ij}(t)$ are the corresponding transition probabilities of the Markov chain evaluated at time Δt , solution of the Kolmogorov's backward equations,

$$P'(t) = Q P(t), \quad P(0) = \mathbb{1} \quad (t \geq 0), \quad (8)$$

for the matrix transition probability $P = (p_{ij})$. \square

Note that such a probabilistic representation allows in practice to compute a single entry i of the vector solution. This is done by generating suitable random paths, corresponding to a continuous-time Markov chain, which evolve backward in time from the state i at $t = \Delta t$ to a final state on S for $t = 0$.

Finally, the chosen entry is computed by averaging the functional η over the sample. Such a functional depends on the initial vector v and the diagonal matrix D . Moreover, applying this Lemma to Eq. (3) allow us to derive the following general theorem.

Theorem 2 *Let $i_k, k = 1, \dots, N$, a sequence of N discrete random variables with outcomes on $S = \{1, 2, \dots, n\}$. The probabilities $p_{i_{k-1}i_k}(t), k = 2, \dots, N$, and $p_{i_{i_1}}(t)$ for $k = 1$, correspond to the transition probabilities of a continuous-time Markov chain generated by the same infinitesimal generator $Q = -L$ and evaluated at time Δt for each k . Then, we have that any entry i of the vector solution \bar{x} in Eq. (3) can be represented probabilistically as*

$$\bar{x}_i = e^{\Delta t d_i/2} \mathbf{E} \left[\prod_{k=1}^N \eta_k \right], \quad (9)$$

where $\eta_k = e^{\Delta t d_{i_k}}, k = 1, \dots, N - 1$, and $\eta_N = e^{\Delta t d_{i_N}/2} v_{i_N}$.

Proof. In view of D being a diagonal matrix, from Eq. (3) the entry i of the vector \bar{x} is given by

$$\begin{aligned} \bar{x}_i = e^{\Delta t d_i/2} \sum_{i_1=1}^n \sum_{i_2=1}^n \dots \sum_{i_N=1}^n (e^{-\Delta t L})_{ii_1} e^{\Delta t d_{i_1}/2} e^{\Delta t d_{i_1}/2} (e^{-\Delta t L})_{i_1 i_2} e^{\Delta t d_{i_2}/2} \dots \\ \times e^{\Delta t d_{i_{N-1}}/2} (e^{-\Delta t L})_{i_{N-1} i_N} e^{\Delta t d_{i_N}/2} v_{i_N}. \end{aligned} \quad (10)$$

As proved in the Lemma above, the matrix $-L$ can be assumed to be the generator of a continuous-time Markov chain with transition probability matrix $P(t)$. Therefore, the equation above can be rewritten as

$$\begin{aligned} \bar{x}_i = e^{\Delta t d_i/2} \sum_{i_1=1}^n \sum_{i_2=1}^n \dots \sum_{i_N=1}^n p_{i i_1} e^{\Delta t d_{i_1}} p_{i_1 i_2} e^{\Delta t d_{i_2}} \dots \\ \times p_{i_{N-1} i_N}(\Delta t) e^{\Delta t d_{i_N}/2} v_{i_N}. \end{aligned} \quad (11)$$

□

Similarly to Lemma 1, a neat picture of this general probabilistic representation can be described as follows: A random path starting at the chosen entry i is generated according to the continuous-time Markov chain governed by the generator Q , and evolves in time by jumping randomly from i to any state on S . Along this process, N given functions η_k are evaluated, and the solution is obtained through the expected value of a suitable multiplicative functional.

The Lemma 1 and Theorem 2 can be conveniently modified to represent probabilistically the complete vector solution \bar{x} . In fact, it is worth observing that the transpose of the generator Q, Q^\top , corresponds to the generator of the

continuous-time Markov chain generated rather forward in time, being in this case the matrix transition probability P solution of Kolmogorov's forward equations,

$$P'(t) = Q^\top P(t), \quad P(0) = \mathbb{1} \quad (t \geq 0). \quad (12)$$

This can be used to generate instead a random path that starts at a given state according to a specific initial distribution, and evolves forward in time governed by a continuous-time Markov chain generated by a suitable generator, which in practice corresponds to the transpose of the generator of the backward equation. However, note that $Q = Q^\top$, since the matrix A is symmetric, and therefore it holds that the generator of the continuous Markov chain forward in time coincides with the generator backward in time. This is mathematically formalized through the following Lemma:

Lemma 3 *Let i and j be discrete random variables on the state space $S = \{1, 2, \dots, n\}$. The random variable j is governed by the probability function $p_j = v_j / \sum_{l=1}^n v_l$ provided $v_j \geq 0$, while the random variable i by the probability function $p_{ij}(t)$ being the transition probabilities of a continuous-time Markov chain generated by the infinitesimal generator $Q = -L$ and evaluated at time Δt . Then, the vector $y = e^{\Delta t D/2} e^{-\Delta t L} e^{\Delta t D/2} v$, can be represented probabilistically as*

$$y_i = e^{\Delta t d_i/2} V \mathbf{E}[\eta], \quad (13)$$

where $V = \sum_{l=1}^n v_l$, $\eta = e^{\Delta t d_j/2}$, and $\mathbf{E}[\eta]$ its expected value.

Proof. Since the proof is similar of the previous Lemma's proof, for completeness we sketch here only the main differences. From Eq. (6) in Lemma 1, y_i can be rewritten as follows

$$y_i = \left(\sum_{l=1}^n v_l \right) \sum_{j=1}^n e^{\Delta t d_i/2} (e^{-\Delta t L})_{ij} e^{\Delta t d_j/2} \frac{v_j}{\sum_{l=1}^n v_l}. \quad (14)$$

Whenever $v_j \geq 0$, $v_j / \sum_{l=1}^n v_l$ can be defined as a suitable probability function p_j for the discrete random variable j on S . Similarly than in the previous Lemma, P is the transition probability matrix for the continuous-time Markov chain generated by $-L$, therefore it holds

$$y_i = \left(\sum_{l=1}^n v_l \right) \sum_{j=1}^n e^{\Delta t d_i/2} p_{ji} e^{\Delta t d_j/2} p_j, \quad (15)$$

and hence y_i can be represented as $e^{\Delta t d_i/2} (\sum_{l=1}^n v_l) \mathbf{E}[e^{\Delta t d_j/2}]$. \square

As before, we used the Lemma 1 to formulate a general theorem, which allows in practice to represent probabilistically the vector solution \bar{x} .

Theorem 4 *Let i_k , $k = 1, \dots, N$, and j , $N + 1$ be discrete random variables with outcomes on*

$S = \{1, 2, \dots, n\}$. The probabilities $p_{i_{k-1}i_k}(t)$, $k = 2, \dots, N$, and $p_{i_1}(t)$ for $k = 1$ correspond to the transition probabilities of a continuous-time Markov chain generated by the same infinitesimal generator $-L$ for each Δt and k , while for the random variable j the probability p_j is given by $p_j = v_j / \sum_{l=1}^n v_l$ provided $v_j \geq 0$. Then, we have that the vector solution \bar{x} in Eq. (3) can be represented as

$$\bar{x}_i = e^{\Delta t d_i/2} V \mathbf{E} \left[\prod_{k=1}^N \eta_k \right], \quad (16)$$

where $\eta_k = e^{\Delta t d_{i_k}}$, $k = 1, \dots, N-1$, and $\eta_N = e^{\Delta t d_{i_N}/2}$, and $V = \sum_{l=1}^n v_l$.

3 The algorithm

To implement the theorems above in order to obtain numerically a single entry or rather the full vector solution, we need to choose a finite sample of given size M , replacing in such a way the expected value by the arithmetic mean. Accordingly this entails a statistical error that it will be discussed in the next section. Here we present the algorithm implemented so far to compute the solution, but before doing that we describe the numerical method used to generate in practice the continuous-time Markov chain. Let $p_{ij}(t)$ be the transition probability matrix, then the Kolmogorov backward equation in Eq. (8) can be equivalently represented as the following system of integral equations

$$p_{ij}(t) = \delta_{ij} e^{-d_i t} + \sum_{j \neq i} \int_0^t ds d_i e^{-d_i s} k_{ij} p_{ij}(t-s), \quad (17)$$

where $k_{ij} = L_{ij}/d_i$. Let S_0, S_1, \dots be a sequence of independent exponential random times picked up from the exponential probability density $p(S_i) = d_i e^{-d_i S_i}$. The integral equations above, along with the sequences of random times, can be used to simulate a path according to the following recursive algorithm: Generate a first random time S_0 that obeys the exponential density function. Then, depending on whether $S_0 < t$ or not, two different routes are taken. If $S_0 > t$, the algorithm is stopped, and no jump from the state i to a different state is taken. If, on the contrary, $S_0 < t$, then the state i jumps to a different state j according to the probability function k_{ij} , and a new second random number exponentially distributed S_1 is generated. If $S_1 < (t - S_0)$ the algorithm proceeds repeating the same elementary rules, otherwise it is stopped. To illustrate the procedure above, in Fig. 1 we show two random paths of a continuous-time Markov chain corresponding to an adjacency matrix of a small world network of size 100.

Algorithm 1 describes a pseudocode corresponding to the implementation of Theorem 2, which allows one to compute single entries of the vector solution,




Fig. 1. Two different sampled paths obtained for computing the first entry of the vector $e^{4A}\mathbf{1}$. The paths correspond to random jumps to different states of a small-world network. The size of the network is 100, the initial state is $i = 1$, and the discretization parameter $\Delta t = 0.25$.

Algorithm 1. Algorithm to compute a single entry i of the vector solution x

Require: $i, \Delta t, N, M$

$x_i = 0$

for $l = 1, M$ **do**

$\eta = 1, j = i$

for $n = 1, \dots, N$ **do**

$\eta = \eta e^{d_j \Delta t / 2}$

generate(τ)

while $\tau < \Delta t$ **do**

generate(S), generate(j)

$\tau = \tau + S$

end while

$\eta = \eta e^{d_j \Delta t / 2}$

end for

$x_i = x_i + (v_j \eta) / N$

end for

while Algorithm 2 describes the pseudocode for computing probabilistically the complete vector solution \bar{x} , mathematically formalized in Theorem 4.

Algorithm 2. Algorithm to compute the complete vector solution x

Require: $\Delta t, N, M, n$

$x_i = 0, V = \sum_{j=1}^n v_l$

for $l = 1, M$ **do**

 generate(i)

$\eta = V, j = i$

for $n = 1, \dots, N$ **do**

$\eta = \eta e^{d_j \Delta t / 2}$

 generate(τ)

while $\tau < \Delta t$ **do**

 generate(S), generate(j)

$\tau = \tau + S$

end while

$\eta = \eta e^{d_j \Delta t / 2}$

end for

$x_i = x_i + \eta / N$

end for

4 Computational complexity of the Monte Carlo algorithm

To estimate properly the computational cost of the algorithms above, the tasks inside the two nested loops were analyzed separately from those that were composed inside the *do-while* loop, and those outside it. Hence,

$$T_{CPU} = T_{in} + T_{out} = \sum_{l=1}^M \sum_{n=1}^N (t_{ln}^{in} + t_{ln}^{out}). \quad (18)$$

Note that the computational cost T_{in} of the inside task, which accounts for the time spent in generating the sequences of random times for the evolving paths, depends on the given matrix, while the cost of the outside tasks T_{out} are totally independent. In fact, the computational time requires to generate a random paths is random, and in practice depends on the specific entries of the row of the matrix j randomly visited by the paths, including the connectivity and degree d_j of the adjacency matrix of the associated graph. More specifically, the degree affects directly the time spent by the algorithm inside the *do-while* loop, since the mean time $\langle S_j \rangle$ of the exponential probability density governing the random time S_j is given by $1/d_j$. In view of the random jumping through the rows of the matrix, in the following let us assume that the computational time spent in total for all random paths can be reasonably approximated as

$$T_{in} = \alpha_{in} \frac{\Delta t}{\bar{\tau}} N M. \quad (19)$$

Here $\bar{\tau} = 1/\bar{d}$ is the corresponding mean time value obtained for a suitable matrix, and its associated adjacency matrix, with an average degree \bar{d} given

by

$$\bar{d} = \frac{1}{n} \sum_{i=1}^n d_i, \quad (20)$$

and α_{in} the corresponding proportionality constant. Such a constant takes into account the computational cost for evaluating the two functions ($generate(S)$, $generate(j)$) in Algorithms 1 and 2. These functions are responsible for generating both, a random time for the evolving paths ($generate(S)$), and a random number j governed by the probability function k_{ij} defined in Eq. (17), which determines the row of the matrix where to jump ($generate(j)$). Note that the cost for generating the exponential random time using the function $generate(S)$ is fully independent of the matrix size, which is not the case when generating the random jumping using the function $generate(j)$. In fact, for a matrix with arbitrary different matrix coefficients the probability function described by k_{ij} is in general nonuniform, and therefore the cost for generating a random j increases at most linearly with the matrix size. However, to improve the performance of this function more efficient algorithms can be implemented, such as the row-searching method based in practice in a binary search tree method as proposed in [32]. Nevertheless, interestingly, for the specific case of the adjacency matrix of undirected networks it turns out that such a probability function becomes uniform, since all nodes are equally probable to jump, and therefore the computational cost for generating the random j becomes fully independent of the matrix size. Indeed such a random j can be trivially generated by simply multiplying a random number uniformly distributed between 0 and 1 by the degree of the i node, and finally rounding to the nearest integer.

Concerning the time spent by the remaining outside tasks, it can be readily estimated as

$$T_{out} = \alpha_{out} N M, \quad (21)$$

where α_{out} is the corresponding proportionality constant. Recall that $\Delta t = \beta/N$, therefore it holds that

$$T_{CPU} = \alpha_{in} \beta \bar{d} M + \alpha_{out} \frac{\beta}{\Delta t} M. \quad (22)$$

Note that for a sufficiently small Δt , the predominant term is the second one, which appears to be almost independent of the matrix, while for a large Δt the opposite behavior occurs. This is in good agreement with the results shown in Table 1, corresponding to the CPU time spent by the Monte Carlo algorithm when computing the total communicability of two different networks characterized by different average degrees.

Δt	CPU Time SM (s)	CPU Time SC (s)
0.5	2.08	2.91
0.25	2.88	3.85
0.125	4.11	4.96
0.0625	6.33	6.86
0.03125	10.71	11.05
0.0156	19.51	19.59

Table 1

CPU time spent by the Monte Carlo algorithm for computing the total communicability of a small-world network (SM), and a scale-free network (SC) of size 10^6 nodes. The average degree of the small-world network is $\bar{d} = 2.4$, while for the scale-free network is $\bar{d} = 4$. The sample size is $M = 10^6$

5 Numerical errors

In computing a single entry of the vector solution or the complete vector, we should consider, in practice, two sources of numerical error. In fact, we have to face the error due to the splitting in Eq. (2) and (4), and the error coming from necessarily replacing the expected value in (16) by a finite sum over a given finite sample of size M . In the following we focus exclusively on the numerical scheme for computing the full vector solution, since the analysis of the error for the companion method for computing a single entry turns out to be identical. To be more precise, the global error made in computing probabilistically the vector solution can be evaluated as

$$\varepsilon = x - e^{\Delta t d_i/2} V \frac{1}{M} \sum_{l=1}^M [\prod_{k=1}^N \eta_k^l] = \varepsilon_1 + \varepsilon_2, \quad (23)$$

where η_k^l corresponds to the l realization of the η_k random variable defined in Theorem 4, and

$$\varepsilon_1 = x - \left(e^{\Delta t D/2} e^{-\Delta t L} e^{\Delta t D/2} \right)^N v \quad (24)$$

$$\varepsilon_2 = \left(e^{\Delta t D/2} e^{-\Delta t L} e^{\Delta t D/2} \right)^N v - e^{\Delta t d_i/2} V \frac{1}{M} \sum_{l=1}^M [\prod_{k=1}^N \eta_k^l] \quad (25)$$

As mentioned already in Sec. 2, the first error ε_1 is due merely to the splitting procedure, and the error is of the order of $O(\Delta t)$ or $O(\Delta t^2)$ depending on whether the Lie or the Strang splitting is used.

The second error, ε_2 , is the pure Monte Carlo statistical error, and of order

of $O(M^{-1/2})$. In fact, it is well known that the arithmetic mean appearing in (23) provides the best unbiased estimator for the expected value in (16). In practice, one should simulate on the computer the random variables, based on generating random numbers. By doing so, the error made in replacing the expected value with the mean over a finite size sample is statistical in nature. More precisely, ε_2 turns out to be, for a large M value, approximately a random Gaussian variable with standard deviation proportional to $M^{-1/2}$, i.e.,

$$\varepsilon_2 \approx \frac{\sigma \nu}{M^{1/2}}, \quad (26)$$

where σ denotes the square root of the variance, and ν is a standard normal (i.e., $N(0, 1)$) random variable. All this clearly shows that the proposed Monte Carlo method could in principle have a poor numerical performance, and also that the error is merely statistical, so it can only be bounded by some quantity with a certain degree of confidence. However, there already exist many available statistical techniques, such as variance reduction, multilevel Monte Carlo, and quasi-random numbers, that can be used, in practice, to improve greatly the order of the global error, and consequently the overall performance of the algorithm.

To illustrate the global error of the numerical method, and its convergence, several examples were run to examine the specific problem of computing the total communicability of a complex network (see [6] e.g.) for different network sizes. The error was computed assuming the solution obtained using the built-in function *expm* of Matlab as if it were the theoretical solution. The underlying algorithm consists essentially of a rational approximation by means of the Padé approximation of an underlying series expansion of the matrix, along with a direct method for computing the inverse of a suitable linear algebra problem, which in practice entails an LU decomposition. Therefore, since it is based on more highly accurate methods, we can assume it to be a highly accurate approximation of the generally unavailable theoretical solution. In Fig. 2 the absolute error is plotted as a function of the time step Δt for a network of 100 nodes in (a), and 1000 in (b), and the two different splitting methods. The solid line corresponds in practice to the purely splitting error ε_1 , while the dashed line corresponds to the error of the complete Monte Carlo method ε . Surprisingly, both the Lie and Strang splitting seem to show the same order of the error (which is 2 as can be seen from the slope of the ancillary function plotted in the figure to help the reader). This can be readily explained from Eq. (2) and the definition of total communicability. Indeed, by definition of the total communicability of a network [6]

$$TC = (\mathbf{1}, e^A \mathbf{1}), \quad (27)$$

where $\mathbf{1}$ is a vector of ones, (\cdot, \cdot) the scalar product, and since $L\mathbf{1} = \mathbf{0}$, it holds that

$$[(\mathbf{1}, D L \mathbf{1}) - (\mathbf{1}, L D \mathbf{1})] = 0. \quad (28)$$

Therefore, from (2) the order of the local error for the Lie splitting turns out to be of order $O(\Delta t^3)$, and in particular of order $O(\Delta t^2)$ as the global error, as it happens for the Strang splitting. However, quantitatively the global error for the Strang splitting appears to be smaller than the error obtained with the Lie splitting. The reason can be found in that the proportionality constant multiplying Δt^2 for the Strang splitting is smaller. Rather, this does not occur when computing the communicability of a single node, which in practice entails computing a single entry of the vector solution. In fact, Fig. 3 shows that the error of the Strang splitting is one order of magnitude larger than for the Lie splitting, being in both cases the order of convergence as was theoretically expected.

Moreover, it is worth observing that the absolute error tends to a constant value for sufficiently small values of Δt , being such a critical value larger when a smaller sample size is used. This is because for this range of values of Δt the global error comes predominantly from the statistical error (which is independent of Δt), since the splitting error is already much smaller than the statistical error. In practice this means that from a given value of Δt it becomes useless to reduce further the time step Δt , and consequently increase the computational cost. From that point the error becomes mostly statistical, being therefore required rather to increase the sample size M in order to continue reducing the global error. In fact, in Fig. 4 the time step Δt was chosen to be sufficiently small, 10^{-3} . This makes the splitting error negligible compared with the statistical error, and therefore the absolute error shown is mostly statistical, decreasing as $M^{-1/2}$ as expected by theory.

For the specific problem of computing the total communicability of a network in Eq. (27), it can be analytically estimated how the numerical error depends on the network size. Concerning the error due to exponential splitting, from Eq. (4) and since $L\mathbf{1} = \mathbf{0}$, we obtain

$$\epsilon_{TC} \leq |(\mathbf{1}, D L D, \mathbf{1})| \Delta t^2. \quad (29)$$

By using the Cauchy-Schwarz inequality, the error can be bounded as follows

$$\epsilon_{TC} \leq \|D\|_\infty^2 \|L\|_\infty \Delta t^2 = 2d_{max}^3 \Delta t^2, \quad (30)$$

where d_{max} corresponds to the maximum degree of the network. Therefore, for those networks characterized by having a maximum degree almost independent of the network size, such as the small-world networks, the error is almost independent of the size. In fact, this is in agreement with the results plotted in Fig. 2, where the total communicability has been computed for two different network sizes. Rather for those other networks where the maximum degree increases with the network size, the error increases with the network size requiring therefore to reduce Δt accordingly to keep constant the error. Concerning the other source of errors, as it was mentioned above this concerns

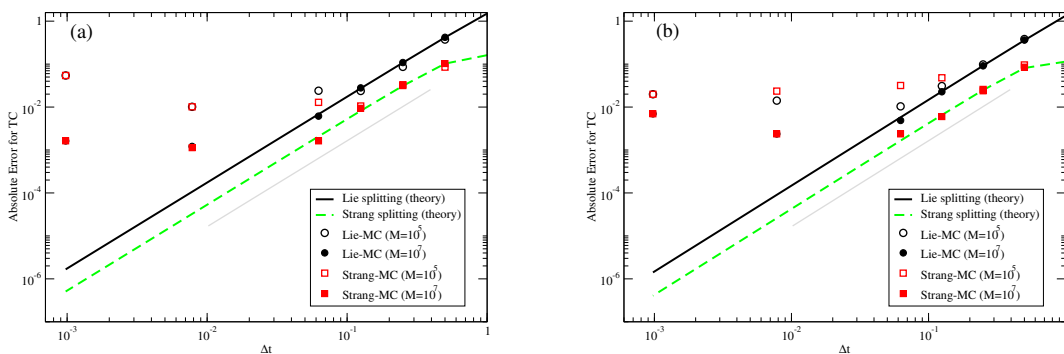


Fig. 2. Absolute numerical error obtained when computing the total communicability of a small-world network for different values of the time step Δt . The networks are composed of (a) 100 nodes and (b) 1000 nodes, respectively. The solid and dashed lines correspond to the theoretical solution obtained with the Lie splitting, and the Strang splitting, respectively, while the points denote the errors obtained when simulating using Monte Carlo. The gray line corresponds to an ancillary function of slope 2.

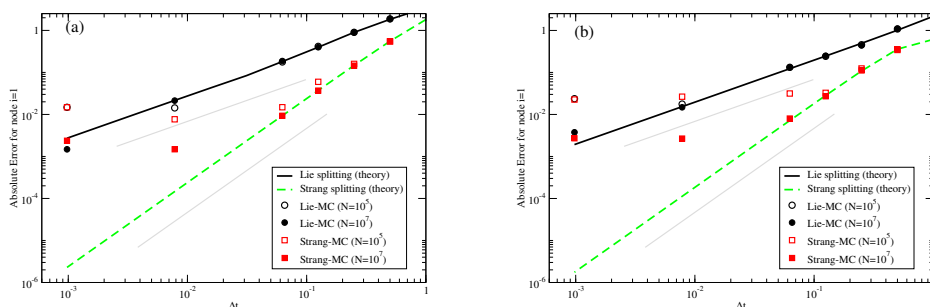


Fig. 3. Absolute numerical error obtained when computing the communicability of a single node ($i = 1$) of a small-world network for different values of the discretization step Δt . The networks are composed of (a) 100 nodes and (b) 1000 nodes, respectively. The solid and dashed lines correspond to the theoretical solution obtained with Lie splitting, and Strang splitting, respectively, while the points denote the errors obtained when simulating using Monte Carlo. The gray lines correspond to ancillary functions of slope 1 and 2.

the statistical error due to the finite sample of the Monte Carlo simulations. However, this error turns out to be independent of the network size as it is shown in Fig. 4. This should not be so surprising since as it happens for Monte Carlo integration of definitive integrals, the underlying error when computing numerically the expected value of the function does not depend on the number of dimensions of the integral. In fact, this is the main advantage of Monte Carlo integration against most deterministic methods, which it is known to grow exponentially with the dimensions.




Fig. 4. Statistical error obtained when computing the total communicability of a small-world network for different values of the sample size M . The networks are composed of 100 nodes and 1000 nodes, respectively. The time step is kept fixed to $\Delta t = 10^{-3}$. The gray line corresponds to an ancillary function of slope $-1/2$.

6 Some results and benchmarks

To illustrate the numerical method and performance of the underlying algorithm, in the following we show the results corresponding to several benchmarks run so far. They concern the numerical computation of the metric communicability in both synthetic and real complex networks. For comparison with other methods, and to estimate the numerical errors, the Matlab toolbox *funm kryl* developed in [4], and freely available in [21], has been used. Such a code implements a Krylov subspace method with deflated restarting for matrix functions. Concerning the Monte Carlo algorithm, it was implemented in Fortran 90, and for a fair comparison with the performance obtained with the Matlab code, no further optimization of the code or the Fortran compiler was performed. Moreover, Fortran is purely sequential, while Matlab by default employs multi-threading architecture for running simulations. Therefore, to make a fair comparison between the systems, Matlab's multi-threading feature was completely disabled. The simulations were run on a computer equipped with an Intel Xeon CPU E5620 at 2.40 GHz and 96 GB of RAM.

Small-world networks. These networks were generated in Matlab using the function *smallw*, freely available through the toolbox CONTEST [11]. In Table 2 the computational time to compute the total communicability of the network for about the same error is shown for different network sizes. Up to a network size of 10,000 nodes, the error was estimated using Matlab’s built-in function *expm* as if it were the theoretical solution. For increasingly larger networks however, the high computational cost of this function makes its computation a formidable task, making it necessary therefore to resort to other methods to estimate the numerical error. For such a purpose, the aforementioned Krylov-based method was used by setting a very small value of the stopping-accuracy parameter, 10^{-16} , as well as the restart parameter to 40. These are also the parameters that have been modified in order to obtain a similar error with the Monte Carlo simulations.

Table 2

CPU time spent for computing the total communicability of a small-world network. For the Monte Carlo method the parameters are $M = 10^6$ and $\Delta t = 0.03125$. The error was kept fixed to 10^{-3} .

Network size	CPU Time MC (s)	CPU Time Matlab (s)
10^3	0.52	0.009
10^4	0.83	0.011
10^5	0.85	0.050
10^6	0.95	0.491
10^7	0.98	6.036
10^8	1.07	73.22

It is remarkable to note that the computational cost of the Monte Carlo method appears to be almost independent of the size of the network, while increasing almost linearly for the Krylov-based method. For the Monte Carlo method this can be explained by the fact that the maximum degree of the network is almost independent of the network size, and consequently as explained in Sec. 5, the numerical error. Therefore, to compute the solution within a given prescribed accuracy it is not required to modify the values of the sample size M , and the time step Δt for increasingly larger sizes, thereby ensuring the same computational cost of the algorithm for any network size. Such a feature allows the Monte Carlo method to achieve a computational performance which is notably higher than the classical counterpart, based on the Krylov-based method, for large scale problems. In fact, it has been pointed out in the literature [8,32] and more specifically in [30] through a suitable Theorem, that there exists an algorithm based on the Lanczos method capable of computing the vector solution of the action of a matrix exponential over a given vector in a time that grows linearly with the matrix size. This is mostly due to the spar-

sity of the adjacency matrix of the network, which simplifies considerably the cost of the matrix-vector products associated to each Lanczos iteration. These theoretical findings are therefore numerically confirmed in Table 2, where it can be seen indeed that the CPU time for the Krylov-based method increases almost linearly with the network size.

As discussed in Sec. 3, the Monte Carlo method can be used as well to compute the full vector solution $e^{\beta A} v$. In the particular case of complex networks, and when v is the vector with all entries equal to 1, the vector solution represents the total subgraph communicability of each node of the network [6]. More important than the quantitative values of the entries of the vector, it is the insight obtained through the ranking of the network organized by the importance of its nodes in terms of being more or less communicable inside the network, which could be of primary importance in the field of complex networks. For this purpose, and to evaluate the similarities between the rankings obtained with the Monte Carlo method and the Krylov-based method, we use the intersection distance method [9] on both the full set of nodes of the network, and on 10% of them. The intersection similarity distance for the top K nodes of two vectors x and y is defined as

$$isim_K(x, y) := \frac{1}{K} \sum_{i=1}^K \frac{|x_i \Delta y_i|}{2i}. \quad (31)$$

Here Δ is the symmetric difference operator between the two vectors. In practice, small values of the intersection distance denote large similarities between the vectors, while the limiting value of 1 suggests vectors that are totally disjointed. Since computing the intersection distance could be computationally costly for increasingly large networks, only relatively small sizes of the network were analyzed so far. In Table 3 the results corresponding to networks composed of 1000 and 10,000 nodes are shown. Note that the two ranked vectors show strong similarities, being even stronger for larger network sizes.

Table 3

Similarity results of the two computed ranked communicability vectors obtained with the Krylov-based method, and the Monte Carlo method for different sample sizes. The network is a small-world network of size a) 1000 nodes, and b) 10,000 nodes. For the Monte Carlo method $\Delta t = 0.03125$.

	(a)		(b)	
Sample size M	isim (100%)	isim (10%)	isim (100%)	isim (10%)
10^5	0.0037	0.0279	4.89×10^{-4}	0.00379
10^6	0.0022	0.020	3.39×10^{-4}	0.0021
10^7	0.0014	0.011	2.24×10^{-4}	0.0014

Scale-free networks. Such networks have been generated using the function *pref* belonging to the aforementioned toolbox CONTEST. In contrast to the

small-world network, these networks are characterized by the presence of hubs, which in practice entail a much larger maximum degree, and correspondingly larger maximum eigenvalue than for the small-world networks. For this reason, and to avoid dealing with very large values when computing the total network communicability for large networks, it is more convenient instead to analyze the so-called normalized total communicability [6], which corresponds in practice to the average total communicability of the network per node. This metric can be readily obtained dividing the total network communicability by the network size, that is $TC_n = TC/n$. Since the value of the maximum degree increases with the network size, then in order to keep constant the numerical error it may be necessary for the Monte Carlo method to reduce the time step Δt (or equivalently increasing the parameter N) accordingly. From Eq. (30), and assuming $d_{max} \approx n$ as an upper-bound approximation, it holds that the time step Δt should reduce, at most, as n^{-1} (or equivalently the parameter N increase linearly with n), to ensure a constant numerical error when computing the normalized total network communicability for arbitrary large scale-free networks. As a result, the computational cost of the algorithm, estimated in Sec. 4 as being $T_{CPU} \approx T_{out}$ for sufficiently small Δt , increases therefore linearly with the network size n for these type of networks.

To avoid such a computationally costly procedure, a reasonable alternative relies on computing a generalization of the communicability, that is $e^{\beta A}$, where β is typically interpreted as an effective "temperature" of the network (see [17], e.g.). Essentially the idea is to use the inverse of the maximum eigenvalue as the value of the parameter β , which in practice will control the rapid growth of the norm of the matrix A with the size of the network. To ensure fast convergence of the Monte Carlo solution, using $\beta = 1/\lambda_{max}$, where λ_{max} is the maximum eigenvalue of A , should suffice. However, finding the maximum eigenvalue for large networks is computationally costly, and in the following a faster alternative, based on computing the maximum degree of the network, d_{max} , was used instead as an upper bound value. Note that in doing that the numerical error obtained when computing the normalized total network communicability becomes independent of the network size. This can be proved readily as follows: From Eq. (30), the error to compute the normalized total network communicability is given by,

$$\epsilon_{TC_n} \leq 2d_{max}^3 \Delta t^2 / n = \frac{2d_{max}}{N^2 n}. \quad (32)$$

where the time-step Δt defined in Eq. (1) was replaced by $1/(d_{max} N)$. Now using $d_{max} \approx n$ as an upper-bound approximation, then it follows that the error becomes indeed totally independent of the network size, and hence the computational cost of the algorithm.

However, different values of β could have a direct impact not only on the entries of the communicability vector, but also on the ranking of the nodes

according to their communicability values, and therefore it becomes essential to analyze at least qualitatively such an issue. In Table 4 the similarity results of two ranked communicability vectors are shown for two different network sizes. All the vector solutions are computed this time using the function *expm* of Matlab to minimize the error, and the comparison is done by choosing as the reference vector the ranked communicability vector with $\beta = 1$.

Table 4

Similarity results of two computed ranked communicability vectors obtained for different values of β . The reference vector used for comparison corresponds to $\beta = 1$. The network is a scale-free network of size a) 1000 nodes, and b) 10,000 nodes.

(a) $\lambda_{max} = 10.22, d_{max} = 69$			(b) $\lambda_{max} = 19.52, d_{max} = 357$	
β	isim (100%)	isim (10%)	isim (100%)	isim (10%)
1	0	0	0	0
0.5	0.0025	0.0137	3.74×10^{-4}	0.0026
0.125	0.0026	0.0145	4.01×10^{-4}	0.0029
$1/\lambda_{max}$	0.0026	0.0145	4.01×10^{-4}	0.0029
$1/d_{max}$	0.0026	0.0145	4.27×10^{-4}	0.00315

From Table 4 it is worthwhile to observe the close similarity of the ranked vectors for different values of β , being even closer for larger values of the network size. Recall that for the typical accuracy asked for Monte Carlo simulations, the error is already higher (10^{-3} in the previous examples) than the values obtained for the intersection similarities. This fact can be exploited in practice to choose values of β smaller, and consequently Δt larger, and still be able to characterize properly the communicability of the network, being indeed indistinguishable within the prescribed accuracy for the Monte Carlo simulations.

In Table 5 the times spent to compute the normalized total communicability of scale-free networks of different sizes are shown. As for small-world networks, in all Monte Carlo simulations the error has been kept fixed to 10^{-3} . Similar to the results obtained for the small-world network, the Monte Carlo method outperforms the Krylov-based method for large size networks, having also a computational time independent of the network size, in agreement with the theoretical considerations discussed above.

Real networks. In Table 6 the results corresponding to a few real networks of arbitrary large size are shown. These networks were downloaded from the freely available sparse matrix repository *SuiteSparse Matrix Collection* [35], and correspond to undirected graphs describing the largest strongly connected components of the corresponding Open Street Map road networks in Europe (Europe OSM), the USA roads (USA roads), and finally a directed graph

Table 5

CPU time spent for computing the normalized total communicability of a scale-free network. For the Monte Carlo method the parameters are $M = 10^6$ and $\Delta t = 0.03125$. The error was kept fixed to 10^{-3} .

Network size	CPU Time MC (s)	CPU Time Matlab (s)
10^3	0.62	0.0068
10^4	0.71	0.016
10^5	0.68	0.038
10^6	0.71	0.469
10^7	0.75	4.08
10^8	0.76	52.59

corresponding to Wikipedia. The latter network was conveniently symmetrized following the procedure described in [5]. As in the previous examples, the performance of the Monte Carlo method is notably superior to the Krylov-based method, being that the differences are even more pronounced for large network sizes.

Table 6

CPU time spent for computing the normalized total communicability of real complex networks. For the Monte Carlo method the parameters are $M = 10^6$ and $\Delta t = 0.03125$. The error in the norm L^∞ was kept fixed to 10^{-3} .

Network type	Size	CPU Time MC (s)	CPU Time Matlab (s)
Wikipedia	7,133,814	0.75	5.104
USA roads	23,947,347	0.77	9.925
Europe OSM	50,912,018	0.79	11.79

7 Conclusion

A new Monte Carlo method for computing the action of an exponential matrix on a vector has been proposed. The method is based on generating suitable random paths corresponding to a continuous-time Markov chain governed by the associated Laplacian matrix. It extends the existing Monte Carlo methods discussed so far in the literature for solving linear algebra problems, for dealing now with more involved functions of matrices such as the matrix exponential. An important advantage of the Monte Carlo method is that the probabilistic representation of the solution allows for efficiently computing single entries of the vector solution, along with global metrics involving the full matrix, such as the total communicability in the field of complex networks. Moreover, since the solution is obtained through averaging independent calculations it is espe-

cially well suited for parallel computation. In fact, it is known that the Monte Carlo algorithm turns out to be fully scalable and naturally fault tolerant. To test the performance of the algorithm, several benchmarks have been used, consisting of a variety of complex networks (real and synthetic) for computing the communicability of the network. The numerical errors of the method have been analyzed through the paper. The results have been compared with a classical Krylov-based method, showing a notably superior performance of the algorithm for large-scale matrices, both in terms of computational time and memory requirements.

Acknowledgments

This work was supported by Fundação para a Ciência e a Tecnologia under Grant No. UID/CEC/50021/2019.

References

- [1] J.A. Acebrón, M.P. Busico, P. Lanucara, and R. Spigler, *Domain decomposition solution of elliptic boundary-value problems*, SIAM J. Sci. Comput., 27 (2005) pp. 440-457.
- [2] J.A. Acebrón, and A. Rodríguez-Rozas, *Highly efficient numerical algorithm based on random trees for accelerating parallel Vlasov-Poisson simulations*, J. Comput. Phys., 250 (2013) pp. 224-245.
- [3] S. Mancini, F. Bernal, and J.A. Acebrón, *An Efficient Algorithm for Accelerating Monte Carlo Approximations of the Solution to Boundary Value Problems*, J. Sci. Comput., 66 (2016) pp. 577-597.
- [4] M. Afanasjew, M. Eiermann, O.G. Ernst, and S. Güttel, *Implementation of a restarted Krylov subspace method for the evaluation of matrix functions*, Linear Algebra Appl., 429 (2008) pp. 2293-2314
- [5] M. Benzi, E. Estrada, and C. Klymko, *Ranking hubs and authorities using matrix functions*, Linear Algebra Appl., 438 (2013) pp. 2447-2474.
- [6] M. Benzi, and C. Klymko, *Total communicability as a centrality measure*, J. Complex Networks, 1 (2013) pp. 124149.
- [7] M. Benzi, T.M. Evans, S.P. Hamilton, M.L. Pasini, and S.R. Slattery, *Analysis of Monte Carlo accelerated iterative methods for sparse linear systems*, Numerical Linear Algebra with Appl., 24 (2017).
- [8] M. Benzi, and P. Boito, *Quadrature rule-based bounds for functions of adjacency matrices*, Linear Algebra Appl., 433 (2010) pp. 637-652.

- [9] P. Boldi, *TotalRank: Ranking without Damping*, Tracks and Posters of the 14th International Conference on World Wide Web (2005), New York: ACM Press, pp. 898899.
- [10] M. Caliari, P. Kandolf, A. Ostermann, and S. Rainer, *The Leja Method Revisited: Backward Error Analysis for the Matrix Exponential*, SIAM J. Sci. Comput., 38 (2016), A1639A1661.
- [11] http://www.maths.strath.ac.uk/research/groups/numerical_analysis/contest
- [12] I.T. Dimov, *Monte Carlo Methods for Applied Scientists*, World Scientific, (2008).
- [13] I. T. Dimov, and V.N. Alexandrov, *A new highly convergent Monte Carlo method for matrix computations*, Math. Comput. Simulation, 47 (1998) pp. 165-181.
- [14] I. T. Dimov, V.N. Alexandrov, and A. Karaivanova, *Parallel resolvent Monte Carlo algorithms for linear algebra problems*, Math. Comput. Simulation, 55 (2001) pp. 25-35.
- [15] I. Dimov, S. Maire, and J.M. Sellier, *A new Walk on Equations Monte Carlo method for solving systems of linear algebraic equations*, Appl. Math. Model., 39 (2015) pp. 4494-4510.
- [16] M. Evans, and T. Swartz, *Approximating Integrals Via Monte Carlo and Deterministic Methods*, Oxford University Press, (2000).
- [17] E. Estrada, N. Hatano, and M. Benzi, *The physics of communicability in complex networks*, Phys. Rep., 514 (2012) pp. 89-119.
- [18] G. Forsythe, and R. Leibler, *Matrix inversion by a Monte Carlo method*, Math. Tables Other Aids Comput., 4 (1950) pp. 127129.
- [19] G.H. Golub, and G. Meurant, *Matrices, Moments and Quadrature with Applications*, Princeton University Press, (2010).
- [20] N.J. Higham, and A. H. Al-Mohy, *Computing matrix functions*, Acta Numer., 19 (2010) pp. 159208.
- [21] http://guettel.com/funm_kryl/
- [22] N.J. Higham, and A. H. Al-Mohy, *Functions of matrices: Theory and Computation*, SIAM , (2008)
- [23] A. H. Al-Mohy, and N.J. Higham, *Computing the Action of the Matrix Exponential with an Application to Exponential Integrators*, SIAM J. Sci. Comput., 3 (2011) pp. 488511.
- [24] M. Hochbruck, and A. Ostermann, *Exponential integrators*, Acta Numer., 19 (2010) pp. 209286.

- [25] T. Jahnke, and C. Lubich, *Error bounds for exponential operator splittings*, BIT,40 (2000) pp. 735744.
- [26] H. Ji, M. Mascagni, and Y. Li, *Convergence Analysis of Markov Chain Monte Carlo Linear Solvers Using Ulam–von Neumann Algorithm*, SIAM J. Numer. Anal., 51 (2013) pp. 21072122.
- [27] R.M.M. Mattheij, S.W. Rienstra, and J.H.M. ten Thije Boonkamp, *Partial Differential Equations: Modeling, Analysis, Computation*, SIAM monographs, (2005).
- [28] R. Merris, *Laplacian matrices of graphs: A survey*, Linear Algebra Appl., 197 (1994) pp. 143-176.
- [29] G. Ökten, *Solving linear equations by Monte Carlo simulation*, SIAM J. Sci. Comput., 27 (2005) pp. 511531.
- [30] L. Orecchia, S. Sachdeva, and N. K. Vishnoi, *Approximating the exponential, the lanczos method and an $O(m)$ -time spectral algorithm for balanced separator*, Proceedings of the forty-fourth annual ACM symposium on Theory of computing. ACM. (2012), pp. 11411160.
- [31] M. Pusa, and J. Leppänen, *Computing the Matrix Exponential in Burnup Calculations*, Nuclear Sci. and Eng.,164 (2010) pp. 140150.
- [32] A. Rudi, L. Wossnig, C. Ciliberto, A. Rocchetto, M.Pontil, and Simone Severini, *Approximating Hamiltonian dynamics with the Nyström method*, arXiv:1804.02484 [quant-ph].
- [33] M.N.O. Sadiku, *Monte Carlo methods for electromagnetics*, CRC press, (2009).
- [34] R.B. Sidjea, and W.J. Stewart, *A numerical study of large sparse matrix exponentials arising in Markov chains*, Comput. Stat. Data Anal., 29 (1999) pp. 345-368.
- [35] <https://sparse.tamu.edu/>
- [36] S.H. Weng, Q. Chen, and C.K. Cheng, *Circuit Simulation by Matrix Exponential Method*, IEEE ASIC Conference, (2011) pp. 369-372.
- [37] H. Zhuang, S.H. Weng, and C.K Cheng, *Power Grid Simulation using Matrix Exponential Method with Rational Krylov Subspaces*, IEEE ASIC Conference, (2013).