

Tensor methods for the computation of MTTA in large systems of loosely interconnected components

Giulio Masetti^a, Leonardo Robol^{b,a,1,*}

^a*Institute of Information Science and Technologies “A. Faedo”, ISTI-CNR, Pisa, Italy.*

^b*Dipartimento di Matematica, Università di Pisa, Italy.*

Abstract

We are concerned with the computation of the mean-time-to-absorption (MTTA) for a large system of loosely interconnected components, modeled as continuous time Markov chains. In particular, we show that splitting the local and synchronization transitions of the smaller subsystems allows to formulate an algorithm for the computation of the MTTA which is proven to be linearly convergent. Then, we show how to modify the method to make it quadratically convergent, thus overcoming the difficulties for problems with convergent rate close to 1.

In addition, it is shown that this decoupling of local and synchronization transitions allows to easily represent all the matrices and vectors involved in the method in the tensor-train (TT) format — and we provide numerical evidence showing that this allows to treat large problems with up to billions of states — which would otherwise be unfeasible.

Keywords: Tensor trains, Kronecker structure, Reliability, Mean-time-to-absorption, Mean-time-to-failure.

2010 MSC: 15A60, 15A69, 60J22, 65F10, 65F60

1. Introduction

Model-based analysis of large and complex systems is considered of fundamental importance to tackle the increasing complexity of our society; at the same time, it is challenging, because of the great number of technical issues that need to be overcome in order to accomplish the task.

On the most prolific areas in stochastic modeling is represented by Markov chains because of the trade-off between representativeness and availability of

*Corresponding author

Email addresses: giulio.masetti@isti.cnr.it (Giulio Masetti),
leonardo.robol@unipi.it (Leonardo Robol)

¹The research of the author was partially supported by the INdAM/GNCS project “Metodi di proiezione per equazioni di matrici e sistemi lineari con operatori definiti tramite somme di prodotti di Kronecker, e soluzioni con struttura di rango”. The author is a member of the INdAM Research group GNCS.

solution techniques that offers to the modeling community. Nevertheless, when the number of interacting components in a system increases, even solution techniques that are known to scale well can suffer for the largeness problem. Of particular interest for this paper is the case when the modeler is asked to deal with reward structures on Continuous Time Markov Chains (CTMCs) [31], where the model is represented by an infinitesimal generator matrix Q , an initial probability vector π_0 and a reward vector r . Standard approaches to tackle model largeness, e.g., lumping and symbolic representation [13] of Q , might not be sufficient to address extremely large system models. The symbolic representation of Q by itself can in fact reduce the storage of information related to the chain, but cannot reduce the memory footprint of the vectors involved in the computations.

An example can be obtained by combining smaller subsystems into a larger one; if we suppose to consider k components with only two possible states (for instance, working and failed), which are combined in a single Markov chain, the state space could be as large as 2^k , including all the possible combination of states in the components. For instance, if $k = 50$, we might have up to 2^{50} chain states; storing every single vector, whose entry are assumed to occupy 8 bytes, would require more than 8 Petabyte, making unfeasible to store it in RAM.

To overcome this problem, one may consider a clever use of a combination of RAM and disk storage or considering approximations of the chain [5]; nevertheless, the gain is still not completely satisfactory, and indeed this is considered one of the main obstacle to the scalability of analytical methods for the evaluation of CTMC properties.

Thus, a symbolic representation also of the vectors involved in the computations can be considered a break-through for the modeling community to be able to analyze chains with huge state space. Two recent papers have pursued this direction for *irreducible chains*, where the steady-state probability vector π is computed. In [21], this is achieved exploiting tensor trains [28], whereas in [4] the authors use Hierarchical Tucker Decomposition [14, 22].

The focus of this paper, instead, is on performance, dependability and performability properties of CTMC with *absorbing states*. A new symbolic representation of both matrices and vectors is proposed to enable the assessment of huge models. In particular, a first step forward with respect to the available numerical techniques will be detailed for the case of the Mean Time To Absorption (MTTA) evaluation in the context of reliability modeling [31], where the CTMC has an unique absorbing state and we are interested in computing the MTTA as a cumulative reward measure. This measure sheds light on limit behaviors of the chain² but does not describe a steady-state property of the chain.

The new method is presented formally and applied to a simple but representative case study, where the technique is proved to be feasible for CTMC with up to $3^{32} \approx 10^{15}$ (potential) states, and is supposed to scale even further.

²the MTTA is the limit of the integral of a function of the probability vector $\pi(t)$, for $t \rightarrow \infty$, as detailed in Equation (5).

Storing a vector of double floating point numbers of this length would require more than 13 PetaByte of memory. A recent application of the method we are going to present can be found in [25]. The paper is structured as follows: In Section 2 notation for the Stochastic Automata Network (SAN) formalism will be recalled. Section 3 introduces the new low-rank representation for both the matrices and vectors. Section 4 specializes the new general method to the evaluation of the MTTA, providing all the details of the mathematical steps. In Section 5 very important computational remarks are discussed: the feasibility of the method strongly relies on the application of few key steps, detailed in this section. Section 6 presents the case study and in Section 7 numerical results are shown to demonstrate the feasibility of the method. Finally, in Section 8 conclusions are drawn.

2. Stochastic Automata Networks

As described in [5], it is possible to address the study of large CTMC defining symbolically the infinitesimal generator matrix Q , thus avoiding a complete state-space exploration. In particular, the Stochastic Automata Network (SAN) formalism [30] allows to represent the CTMC as a set of stochastic automata M_1, \dots, M_k , each having a (small) reachable set of states \mathcal{RS}_i , where transitions are of two kinds: *local* and of *synchronization*. Transitions t that are local to M_i , written $t \in \mathcal{LT}_i$, have impact only on \mathcal{RS}_i and indicate the switch from a state $s \in \mathcal{RS}_i$ to a state $s' \in \mathcal{RS}_i$, in the following written $s \xrightarrow{t} s'$. Synchronization transitions $t \in \mathcal{ST}$, instead, can appear in more than one automaton. In particular, if $t \in M_{i_1}$, $t \in M_{i_2}$, \dots , and $t \in M_{i_h}$ then $s_{i_j} \xrightarrow{t} s'_{i_j}$ can fire only if the automaton M_{i_1} is in state s_{i_1} , and the automaton M_{i_2} is in state s_{i_2} , \dots , and the automaton M_{i_h} is in state s_{i_h} , at the same time. The overall CTMC is then the orchestration of local stochastic automata where the director is represented by \mathcal{ST} . The infinitesimal generator matrix Q is not assembled explicitly, and its compressed representation is called *descriptor matrix* and is formally defined by

$$Q = R + W + \Delta, \quad (1)$$

i.e., the sum of local contributions, called R , and synchronization contributions, called W , where

$$R = \bigoplus_{i=1}^k R^{(i)}, \quad W = \sum_{t_j \in \mathcal{ST}} \bigotimes_{i=1}^k W^{(t_j, i)}, \quad (2)$$

$R^{(i)}$ and $W^{(t_j, i)}$ are $|\mathcal{RS}^{(i)}| \times |\mathcal{RS}^{(i)}|$ matrices, and the diagonal matrix Δ is defined as $\Delta = -\text{diag}((R + W)e)$ where e is the vector with all the entries equal to 1. The operator \oplus is the *Kronecker sum*, as formally described in Section 3.1. The matrices $R^{(i)}$ and $W^{(t_j, i)}$ are assembled exploring $\mathcal{RS}^{(i)}$ and can be specified through an high level formalism such as GSPN [13, 6] or PEPA [18]. In

particular, $W^{(t_j, i)} = \lambda_{t_j} \tilde{W}^{(t_j, i)}$ where $\tilde{W}^{(t_j, i)}$ is a $\{0, 1\}$ -matrix defined as follows:

$$\tilde{W}_{s_i, s'_i}^{(t_j, i)} = \begin{cases} 1 & \text{if } t_j \text{ is enabled in } s_i \text{ inside } M_i \text{ and } s_i \xrightarrow{t_j} s'_i \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where λ_{t_j} is the constant rate associated with t_j , equal in every M_i . In particular, if the transition t_j has no effect on the component M_i , we have $\tilde{W}^{(t_j, i)} = I$. In the following we will call

$$\mathcal{PS} = \mathcal{RS}^{(1)} \times \dots \times \mathcal{RS}^{(k)}$$

the *potential state space* and the $|\mathcal{PS}| \times |\mathcal{PS}|$ descriptor matrix Q will be treated implicitly.

Remark 2.1. In general, the set of reachable states given an initial probability distribution π_0 , might be a strict subset of \mathcal{PS} . Some techniques exploit the fact that the reachable state is smaller to achieve a higher efficiency. However, this hides the tensorized structure of Q and of the probability vectors to be computed; therefore, in this work the focus will be on \mathcal{PS} .

3. Low-rank tensors

To overcome the exponential explosion of memory requirements, sometimes called *curse-of-dimensionality* [28], there has been a recent trend in exploiting the structure of Q , which can be recognized from Equation (1), in the setting where the model describes the interaction of loosely interconnected component; in fact, in this case Q can be efficiently stored by only memorizing factors of Kronecker products [5]. This enables a reduction in storage and an acceleration of the matrix-vector operator required in the development of most of the algorithms for computing steady-state probabilities and performance and reliability measures.

However, with the exponential growth of the state space which can happen when combining several systems, even storing vectors with as many components as the cardinality of the (potential) state space can quickly become unfeasible. For this reason, there has been recently a shift in developing “symbolic” representations for the vectors under consideration as well. This turns out to require considerably more effort. Recent promising developments leverage the use of low-rank tensor formats, namely hierarchical Tucker decompositions [4], and Tensor Trains [21] (in this work considered for the steady-state analysis).

For the sake of self-completeness, we briefly review the theory of low-rank tensor operators and Tensor Trains, that will be the building block for the compressed representation proposed in this work. We refer the interested reader to [28] for further details.

3.1. Kronecker sums

As we discussed in Section 4, there are some kind of structures that appear in the definition of the infinitesimal generator Q . In particular, we may define

the *Kronecker sum* $\mathcal{A} := A_1 \oplus \dots \oplus A_k$ as follows:

$$\mathcal{A} = A_1 \otimes I \otimes \dots \otimes I + I \otimes A_2 \otimes I \otimes \dots \otimes I + \dots + I \otimes \dots \otimes I \otimes A_k.$$

Similarly structured matrix arise in other contexts as well, such as the discretization of high-dimensional PDEs, and solution of matrix and tensor equations.

Equation 1 shows that this is the form of the R matrix in the definition of the infinitesimal generator Q of the Markov chains under consideration. The matrix Q is then obtained adding W , which models the weak interaction between the components, and has a similar structure. Even if Q is not exactly in the form of a Kronecker sum, but it has a low-rank tensorial structure.

The term “tensor rank” does not have a single universally accepted meaning. In fact, unlike in the matrix case (which is obtained by setting $k = 2$), several different ranks can be defined — and they have different computational properties. The most classical definition is the CP rank, linked to the Canonical Polyadic Decomposition (also sometimes called PARAFAC — see [20] and the references therein for more details); this is linked to the definition of rank as sum of rank 1 terms, which are in turn defined as outer product $v_1 \otimes \dots \otimes v_k$. However, using this low-rank format is inherently difficult and unstable. For instance, the set of rank R tensors is not closed if $k > 2$, and this makes the low-rank approximation problem ill-posed [9]. Moreover, the computation of the best rank r approximation of a tensor is a difficult (indeed, NP-hard, [17]), and the solution can only be approximated by carefully adapted optimization algorithms, see [21] for a review.

For this reason, there has been interest in finding alternative low-rank representation of tensors. A very robust possibility that is well-understood is the Tucker decomposition, linked to the Higher Order SVD (HOSVD) [8]. However, this approach requires to store a k -dimensional tensor (even though of smaller sizes), and so is only suited for small values of k .

When one is faced with the task of working with high values of k (say, $k > 5$), and a small number of entries for each mode – a natural choice are instead tensor trains [28] or the hierarchical Tucker decomposition [22].

We shall concentrate on the former choice, and in the next section we briefly recall the main tools that we use in the rest of the paper.

3.2. Tensor-train format

Tensor trains are a technology aimed at treating high-dimensional problems: they have already been successfully applied to Markov chains (see [2, 21]) and to the numerical solution of high-dimensional PDEs ([11, 19] and [23]).

Let us consider a large system composed by k smaller components, each with n_i states, $i = 1, \dots, k$. The potential state space \mathcal{PS} can then be written as

$$\mathcal{PS} := \{1, \dots, n_1\} \times \dots \times \{1, \dots, n_k\}.$$

At each time t , the probability vector $\pi(t)^T = \pi_0^T e^{tQ}$ can be expressed in tensor form, as an array with k indices $\pi(t) = \pi(i_1, \dots, i_k)$. A tensor train representation of a tensor v is a collection of order 3 tensors M_i of size $r_i \times n_i \times r_{i+1}$ such

that³ $r_1 = r_k = 1$, and

$$v_{i_1, \dots, i_k} = \sum_{t_1, \dots, t_{k-1}} M_1(1, i_1, t_1) M_2(t_1, i_2, t_2) \dots M_k(t_{k-1}, i_k, 1) \quad (4)$$

The tensors M_i are called *carriages*, hence the name *tensor train* [28]. The tuple (r_2, \dots, r_{k-1}) is called the TT-rank of the tensor v . Similarly, matrices $m \times n$ can be represented as tensors by subdividing the row and column indices. More precisely, a matrix $A \in \mathbb{C}^{n_1 \dots n_k \times n_1 \dots n_k}$ can be viewed (up to permuting the indices) as a larger vector of the vector space $\mathbb{C}^{m_1 n_1 \times \dots \times m_k n_k}$. This vector can be stored in the TT-format as described in (4). Such arrangement makes computing matrix-vector product and matrix-matrix product relatively simple to implement; we refer the reader to [28] for further details.

On the software side, a well-established framework [27] is available for Python and MATLAB [26] We rely on the latter for our numerical experiments.

3.3. Exponential sums

Given a Kronecker sum $\mathcal{A} := A_1 \oplus \dots \oplus A_k$, in a few cases of interest one can devise an efficient strategy for evaluating its inverse (or the action of the inverse on some vector). Given the relevance of this problem in high-dimensional PDEs and various other settings of applied mathematics, several approaches have been devised over the years. In this section we briefly recall the one known after the name of *exponential sums* [3].

For the sake of self-completeness, we briefly recall the main important facts related to this topic. The idea behind exponential sums is to rephrase the inverse as a combination of matrix exponentials. The latter are much easier to compute for a Kronecker sum, as stated by the next Lemma.

Lemma 3.1. *Let $\mathcal{A} = A_1 \oplus \dots \oplus A_k$ be a Kronecker sum. Then, the matrix exponential $e^{\mathcal{A}}$ is given by*

$$e^{\mathcal{A}} = e^{A_1} \otimes e^{A_2} \otimes \dots \otimes e^{A_k}.$$

Proof. It suffices to recall that $e^{A+B} = e^A e^B$ whenever A and B commute; clearly, all the addends in the sum defining \mathcal{A} commute, and the result follows by $e^{A \otimes B} = e^A \otimes e^B$. \square

In view of the previous result, assume we know an expansion of $\frac{1}{x}$ of the following form:

$$\frac{1}{x} = \sum_{j=1}^{\infty} \alpha_j e^{-\beta_j x}, \quad \forall x \in \Lambda(\mathcal{A}),$$

³In particular, M_1 and M_k are matrices, instead of order 3, because they have one dimension with only 1 index.

where $\Lambda(\cdot)$ denote the spectrum of the operator. Then,

$$\mathcal{A}^{-1} = \sum_{j=1}^{\infty} \alpha_j e^{-\beta_j \mathcal{A}}.$$

Truncating the above series yields an approximation of the inverse, and the matrix exponentials are very cheap to compute if one knows the factors A_i , since these are relatively small matrices that can be handled with dense linear algebra techniques. It remains to construct a method to efficiently compute α_j and β_j of such an expansion. We say that an exponential sum has accuracy $\epsilon > 0$ on the interval $[a, b]$ if, for any $x \in [a, b]$, we have $\left| \frac{1}{x} - \sum_{j=1}^k \alpha_j e^{-\beta_j x} \right| \leq \epsilon$.

Lemma 3.2 ([16]). *Let A be a diagonalizable matrix, $f(z)$ a function, Then,*

$$\|f(A)\| \leq C \max_{z \in \Lambda(A)} |f(z)|,$$

where $\|\cdot\|$ is any induced norm.

Lemma 3.3. *Let α_j, β_j be the coefficients of an exponential sum with accuracy ϵ over $[1, R]$. Then, if A is a diagonalizable matrix with spectrum contained in $[1, R]$, we have*

$$\left\| A^{-1} - \sum_{j=1}^k \alpha_j e^{-\beta_j A} \right\| \leq C \epsilon$$

, where $C = \|V\| \|V^{-1}\|$, where V is the matrix of eigenvectors of A , and $\|\cdot\|$ is any induced norm. In particular, if $\|\cdot\|$ is the euclidean norm and A is normal, then $C = 1$.

Proof. The result is a straightforward application of Lemma 3.2 to the function $f(z) = \frac{1}{z} - \sum_{j=1}^k \alpha_j e^{-\beta_j z}$ over the domain $[1, R]$. \square

A similar result for the non-diagonalizable case can be obtained relying on the field of values, for which we refer to [7].

Lemma 3.3 implies that, given a (normal) matrix in Kronecker sum form $\mathcal{A} = A_1 \oplus \dots \oplus A_k$, to achieve an accuracy ϵ in the computation of \mathcal{A}^{-1} or, equivalently, in the solution of the linear system $\mathcal{A}x = b$, we shall obtain an exponential sum with coefficients α_j, β_j achieving that accuracy ϵ over the eigenvalues of \mathcal{A} . More precisely, we can state the following result that related the norm a matrix function to the spectral properties of the matrices A_i .

Lemma 3.4. *Let $\mathcal{A} = A_1 \oplus \dots \oplus A_k$, where A_i are diagonalizable for $i = 1, \dots, k$. Then, given a function $f(z)$ defined on $\Lambda(A_1) + \dots + \Lambda(A_k)$, we have*

$$\|f(\mathcal{A})\|_2 \leq \left(\prod_{1 \leq i \leq k} \|V_i\|_2 \|V_i^{-1}\|_2 \right) \cdot \max \{ |f(\lambda_{i_1} + \dots + \lambda_{i_k})|, \lambda_{i_j} \in \Lambda(A_j) \},$$

where V_i is the matrix of the eigenvectors of A_i .

Proof. We note that \mathcal{A} is diagonalizable if and only if A_i are, and in this case the matrix of eigenvectors is given by $V_1 \otimes \dots \otimes V_k$. Applying Lemma 3.3 yields the bound

$$\|f(\mathcal{A})\| \leq \|V\| \cdot \|V^{-1}\|_2 \cdot \max\{|f(\lambda_{i_1} + \dots + \lambda_{i_k})|, \lambda_{i_j} \in \Lambda(A_j)\}.$$

The conclusion follows noting that $\|V\|_2 = \prod_{1 \leq i \leq k} \|V_i\|_2$, and the similar result for its inverse. \square

The construction of the coefficients α_j, β_j is beyond the scope of this paper; our approach relies on [3], to which we refer for further details. In particular, we rely on the construction of the exponential sums using sinc quadrature; since we know the exact condition number, we may use the precomputed tables available at [15], which cannot be easily computed on the fly, since they require a specially adapted Newton method with extended precision to reach high accuracy. This would potentially add another speed up to the code, since they have a faster convergence rate.

4. Computing the MTTA

In [24] it has been shown that the computation of several performability measures can be recast as the evaluation of a matrix function. In most cases, one has to compute $w^T f(Q)v$ for appropriate vectors v, w , and a certain function $f(z)$. In this work, we focus on the computation of the *mean-time-to-absorption*.

We assume that the states of the Markov chain are labeled with the integers from 1 to $N = |\mathcal{PS}|$, and that there is a single absorbing state, which can be assumed to have index N . Then, we compute the quantity

$$\text{MTTA} = \int_0^\infty \mathbb{P}\{X(\tau) < N\} d\tau = \mathbb{E} \left[\int_0^\infty \mathbb{1}_{\{1, \dots, N-1\}}(X(\tau)) d\tau \right]. \quad (5)$$

Often, we are interested in the case where the absorbing state corresponds to the failure state of the system. In this case, we use the name *mean-time-to-failure* and the notation MTTF. We assume that the matrix Q is partitioned as follows:

$$Q = \left[\begin{array}{c|c} \hat{Q} & \begin{matrix} v_1 \\ \vdots \\ v_{N-1} \end{matrix} \\ \hline 0 \quad \dots \quad 0 & 0 \end{array} \right],$$

where the last row is forced to be zero because the state N is absorbing. Following [32], we know that

$$\text{MTTA} = -\hat{\pi}_0^T \hat{Q}^{-1} e = \pi_0^T f(Q)(e - e_N e_N^T), \quad f(z) = \begin{cases} -\frac{1}{z} & z \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where e is the vector of all ones, and $\hat{\pi}_0$ contain the first $N - 1$ entries of π_0 . The aim is to consider the case where Q can be efficiently represented in the TT format.

Remark 4.1. The fact that Q has a low-rank tensorial structure (in the TT sense) does not imply any particular structure for \hat{Q} . Indeed, the tensor structure requires making use of the isomorphism $\mathbb{C}^N = \mathbb{C}^{n_1} \times \dots \times \mathbb{C}^{n_k}$, which in turn is related to factorizing $N = n_1 \cdots n_k$. Knowing the factors n_i does not give any information on a similar factorization for $N - 1$, nor on the low-rank carriers that might be used to represent \hat{Q} .

In particular, a well-known method for transforming the MTTA problem into the computation of a steady state vector of a irreducible Markov chain, is to add a transition from the absorbing state back to the starting one [31]. However, this transforms the problem into an eigenvector computation of $Q + \delta Q$, where δQ has rank 1. This problem might be addressed directly, exploiting the techniques presented in [21] for computing steady state probabilities in the TT-format; recasting it into solving an augmented linear system (adding a row of ones to ensure that the computed eigenvector is a probability) is instead undesirable, because it would lead to a loss of tensor structure as reported in Remark 4.1.

Instead, here an alternative strategy is presented: we introduce an auxiliary matrix S which allows to rephrase the measure using the inverse of a low-rank perturbation of Q .

Lemma 4.2. *Let Q the infinitesimal generator of a continuous time Markov chain with exponential rates with N states; assume that the state N is the only failure state, and let S be the rank 1 matrix defined as*

$$S = (Qe_N)e_N^T - e_Ne_N^T.$$

Then, if π_0 has the N -th component equal to 0, we have $\text{MTTA} = -\pi_0^T(Q - S)^{-1}e$, where e is the vector of all ones.

Proof. By construction, we have that $Q - S$ is block diagonal and therefore

$$(Q - S)^{-1} = \begin{bmatrix} \hat{Q} & \\ & 1 \end{bmatrix}^{-1} = \begin{bmatrix} \hat{Q}^{-1} & \\ & 1 \end{bmatrix},$$

and $-\pi_0^T(Q - S)^{-1}e = -\hat{\pi}_0^T\hat{Q}^{-1}e - [\pi_0]_N = \text{MTTA} - [\pi_0]_N$. We conclude noting that $[\pi_0]_N = 0$. \square

Remark 4.3. Note that the TT-rank of S is $(1, \dots, 1)$, since it is a matrix of rank 1 [28], and if Q has a low TT-rank the same holds for $Q - S$.

The important consequence of Lemma 4.2 is that, even though we cannot extract a submatrix from Q to compute the MTTA, we can make a rank 1 (both in CP and in the TT sense) perturbation S to Q , such that $Q - S$ is invertible, and allows to easily obtain the same result.

Lemma 4.4. *Let $v = e + \gamma e_N$, for any $\gamma \in \mathbb{R}$. Then, with the notation of Lemma 4.2, we have*

$$\text{MTTA} = -\pi_0^T(Q - S)^{-1}e = -\pi_0^T(Q - S)^{-1}v.$$

Proof. Note that, $[\pi_0]_N = 0$, and therefore

$$-\pi_0^T(Q - S)^{-1}v = - \begin{bmatrix} \hat{\pi}_0^T \hat{Q}^{-1} & 0 \end{bmatrix} \begin{bmatrix} 1 \\ \vdots \\ 1 \\ 1 + \gamma \end{bmatrix} = -\hat{\pi}_0^T \hat{Q}^{-1}e = \text{MTTA}.$$

□

We have recast the problem to solving a linear system $(Q - S)x = e$, where the matrix $Q - S$ is expressed in TT format. Unfortunately, as we will see later on, a few problem of interest for us do not play very well together with the more widespread tensor train system solvers (such as AMEN [12] or DMRG [29]). For this reason, we propose a different solution scheme based on the *Neumann expansion*. In particular, let M be any matrix with spectral radius strictly smaller than 1. Then,

$$(I - M)^{-1} = I + M + M^2 + M^3 + \dots = \sum_{j=0}^{\infty} M^j \quad (7)$$

If we partition Q as $Q = Q_1 + Q_2$, we can write

$$(Q - S)^{-1} = (Q_1 + Q_2 - S)^{-1} = (I + Q_1^{-1}(Q_2 - S))^{-1}Q_1^{-1}.$$

Setting $M = -Q_1^{-1}(Q_2 - S)$, assuming that $\rho(M) < 1$ and using the Neumann expansion (7) we obtain

$$(Q - S)^{-1} = \sum_{j=0}^{\infty} (-1)^j (Q_1^{-1}(Q_2 - S))^j Q_1^{-1}. \quad (8)$$

The above formula can be used to approximate $x = (Q - S)^{-1}e$ as needed for (6) by truncating the infinite sum to ℓ terms:

$$x_\ell = \sum_{j=0}^{\ell} (-1)^j (Q_1^{-1}(Q_2 - S))^j Q_1^{-1}e, \quad \|x - x_\ell\|_\infty \leq C\rho(M)^{\ell+1},$$

for an appropriate constant C . We give an explicit method to construct the additive splitting $Q = Q_1 + Q_2$ so that M is guaranteed to have spectral radius less than 1. This will be achieved in Theorem 4.8. The pseudocode describing the resulting method is presented in Algorithm 1.

This method has a linear convergence rate [10], which is given by $\rho(M)$. However, it can be accelerated to obtain a quadratically convergent method by a simple modification. Note that we can refactor (7) as follows:

$$(I - M)^{-1} = (I + M)(I + M^2)(I + M^4) \dots (I + M^{2^j}) \dots \quad (9)$$

Algorithm 1 Neumann series (7) to approximate $x = (Q - S)^{-1}e$

```

1: procedure NEUMANSERIES( $Q_1, Q_2, \ell$ )
2:    $y \leftarrow Q_1^{-1}e$ 
3:    $x \leftarrow y$ 
4:   for  $j = 1, \dots, \ell$  do
5:      $y \leftarrow -Q_1^{-1}(Q_2 - S)y$ 
6:      $x \leftarrow x + y$ 
7:   end for
8:   return  $x$ 
9: end procedure

```

Truncating the above equation and permuting the factors $(I + M^{2^j})$ yields another method to approximate $x = (Q - S)^{-1}e$, which has a much faster convergence, and is described by the equation:

$$(I - M)^{-1}Q_1^{-1}e \approx x_{2^{\ell+1}-1} = (I + M^{2^\ell})(I + M^{2^{\ell-1}}) \cdots (I + M^2)(I + M)Q_1^{-1}e$$

In particular, the result of ℓ steps of this method gives the same result of $2^{\ell+1} - 1$ of Algorithm 1. The pseudocode for this approach is given in Algorithm 2. As visible on line 5, this method required to store the repeated squares of the matrix M .

Algorithm 2 Neumann series (9) to approximate $x = (Q - S)^{-1}e$

```

1: procedure NEUMANSERIES( $Q_1, Q_2, \ell$ )
2:    $M \leftarrow -Q_1^{-1}(Q_2 - S)$ 
3:    $x \leftarrow Q_1^{-1}e + MQ_1^{-1}e$ 
4:   for  $j = 2, \dots, \ell$  do
5:      $M \leftarrow M^2$ 
6:      $x \leftarrow x + Mx$ 
7:   end for
8:   return  $x$ 
9: end procedure

```

Remark 4.5. A favorable property of both approaches is that the convergence of x_ℓ to x is monotonically decreasing and non-positive. That is, for each $\ell' \leq \ell$ we have $x_{\ell'} \geq x_\ell$. Since the MTTA is equal to $-\pi_0^T x$, the estimates $-\pi_0^T x_\ell$ of the MTTA obtained in the intermediate steps are guaranteed lower bounds.

We note that Algorithm 1 can be slightly modified to compute $\pi_0^T(Q - S)^{-1}$ instead of $(Q - S)^{-1}e$. Both vectors can then be used to compute the MTTA through a dot product. However, the former choice has the advantage that $\pi_0(s) = 0$ implies that $x(s) = 0$ throughout the iterations for every $s \in \mathcal{PS} \setminus \mathcal{RS}$. In particular, the non-reachable part of the chain has no effect on the computed tensor, and this helps to keep the TT-rank low during the iterations. The

Algorithm 3 Neumann series (7) to approximate $\pi^T = \pi_0^T(Q - S)^{-1}$

```

1: procedure NEUMANSERIES( $Q_1, Q_2, \ell$ )
2:    $y \leftarrow \pi_0^T$ 
3:    $x \leftarrow y$ 
4:   for  $j = 1, \dots, \ell$  do
5:      $y \leftarrow -yQ_1^{-1}(Q_2 - S)$ 
6:      $x \leftarrow x + y$ 
7:   end for
8:    $x \leftarrow xQ_1^{-1}$ 
9:   return  $x$ 
10: end procedure

```

modified algorithm is reported for completeness in Algorithm 3, where now x, y are row vectors.

Lemma 4.6. *Let $A \geq 0$ be a non-negative $N \times N$ matrix, and e the vector of all ones. Then,*

$$\rho(A) \leq \|A\|_\infty = \max_{1 \leq i \leq N} (Ae)_i$$

Moreover, if there is at least one component of Ae strictly smaller than $\|A\|_\infty$, then $\rho(A) < \|A\|_\infty$.

Proof. The first statement is the definition of infinity norm, whereas the second follows directly by the first Gerschgorin theorem. \square

The next result provides a technique for splitting an infinitesimal general Q (i.e., up to a change of sign, any M -matrix with zero row sums) in a way that allow to perform the Neumann expansion. Let us first recall a few properties of diagonally dominant matrices.

This observation implies that, even though in view of Remark 4.1 we have not restricted the problem to the set of reachable states, when running Algorithm 3 the iteration is implicitly restricted to this set, as x_ℓ are vectors with positive components only for indices in \mathcal{RS} .

The next results are aimed at constructing the additive splitting for Q that satisfies the hypotheses for the Neumann expansion.

Lemma 4.7. *Let $A = D - B$, with $D < 0$ and diagonal, $B \geq 0$, and $Be < -De$, where e is the vector with all components equal to 1. Then, A is invertible and $A^{-1} \leq 0$.*

Proof. This fact can be easily prove using the tools from theory of nonnegative matrices, since A is an M -matrix and $D - B$ is a regular splitting (see for instance [1]). We provide a simple proof for the sake of completeness. Since $D < 0$, the condition $A^{-1} \leq 0$ is equivalent to $(D^{-1}A)^{-1} \geq 0$. Moreover, the strict row diagonal dominance implies that $\|D^{-1}A\|_\infty < 1$, so we have

$$(D^{-1}A)^{-1} = (I - (-D^{-1}A))^{-1} = \sum_{j \geq 0} (-D^{-1}A)^j \geq 0,$$

where we have used that $(-D^{-1}A) \geq 0$. \square

Theorem 4.8. *Let $A = D + A_1 + A_2$ any $N \times N$ matrix such as D is diagonal and non-positive, A_1, A_2 are non negative, $e_N^T(D + A_1 + A_2) = e_N^T A_1 = 0$, and $(D + A_1 + A_2)e = 0$. Then, if we define $S = (A_1 + A_2)e_N e_N^T$, $(D + A_1)$ is invertible and $\min_{i=1, \dots, N-1} A_{iN} > 0$, we have that $\|(D + A_1)^{-1}(A_2 - S)\|_\infty < 1$.*

Proof. Let us denote with $M := (D + A_1)^{-1}(A_2 - S)$. All the columns of this matrix are non-positive (see Lemma 4.7), with the only exception of the last one, which is non-negative. This is a consequence of the fact that $(D + A_1)^{-1}$ is non-positive, and $(A_2 - S)$ has the first $N - 1$ columns with positive entries, and the last one with negative ones.

Therefore, it is clear that we have $\|M\|_\infty = \|M(e - 2e_N)\|_\infty$, and the vector $M(e - 2e_N)$ is element-wise non-positive by construction. We have

$$M(e - 2e_N) = (D + A_1)^{-1}(A_2 - S)(e - 2e_N).$$

Using the relations $(D + A_1)^{-1}A_2e = (D + A_1)^{-1}(D + A_1 + A_2)e - e = -e$ and $Se = Se_N = (A_1 + A_2)e_N$ we get

$$\begin{aligned} M(e - 2e_N) &= -e - 2(D + A_1)^{-1}A_2e_N + (D + A_1)^{-1}(A_1 + A_2)e_N \\ &= -e + (D + A_1)^{-1}(A_1 - A_2)e_N \\ &= -e + e_N - (D + A_1)^{-1}(D + A_2)e_N. \end{aligned}$$

By construction, we know that the last row of $D + A_2$ is equal to $-e_N^T A_1$ and is therefore zero. The first $N - 1$ entries in $(D + A_2)e_N$ are taken from A_2 and therefore they are (strictly) positive. Since $(D + A_1)^{-1}$ is entry-wise strictly negative, we have that $v = -(D + A_1)^{-1}(D + A_2)e_N$ has the first $N - 1$ components strictly positive. Therefore, we have that $M(e - 2e_N) > -1$ element-wise, and on the other hand we knew that $M(e - 2e_N)$ is non-positive. This implies that $\|M\|_\infty < 1$, as claimed. \square

Remark 4.9. The previous results are closely related with the theory of M -matrices. Indeed, the decomposition $A = M - N$ with $M = (D + A_1)$ and $N = -A_2 + S$ is almost a regular splitting, because M^{-1} is negative, and N is positive, with the only exception of the last column. If it were a regular splitting, then this would automatically imply that $\rho(M^{-1}N) < 1$ - in view of the theory of nonsingular M -matrices [1].

Lemma 4.10. *Let y be any vector. Then, using the notation of Theorem 4.8, $(D + A_1)^{-1}(A_2 - S)y$ has the last component equal to zero. Moreover, let y_0 be any vector, and define*

$$y_{\ell+1} = -(D + A_1)^{-1}(A_2 - S)y_\ell, \quad \ell \geq 1.$$

Then, if y_ℓ for $\ell > 0$ is non-negative we have $y_{\ell'} \geq 0$ for any $\ell' \geq \ell$.

Proof. We start showing that $e_N^T(D + A_1)^{-1}(A_2 - S) = 0$, which proves the first claim. We have

$$\begin{aligned} e_N^T(D + A_1)^{-1}(A_2 - S) &= D_{NN}^{-1}e_N^T(A_2 - S) = e_N^T A_2 - e_N^T A_1 e_N e_N^T - e_N^T A_2 e_N e_N^T \\ &= e_N^T A_2 - e_N^T A_2 e_N e_N^T = e_N^T A_2 (I - e_N e_N^T) \\ &= -e_N^T D (I - e_N e_N^T) = 0, \end{aligned}$$

where we have used the properties $e_N^T(D + A_1 + A_2) = e_N^T + A_1 = 0$, and the definition of $S = (A_1 + A_2)e_N e_N^T$.

Assume now that $e_N^T y = 0$. Then, $z := (A_2 - S)y = A_2 y \geq 0$, since $Sy = 0$. Moreover, $(D + A_1)^{-1}$ is non-positive in view of Lemma 4.7, and therefore $-(D + A_1)^{-1}z \geq 0$, concluding the proof. \square

Remark 4.11. Note that choosing $\gamma = -1$ in the notation of Lemma 4.4 provides a starting vector for the Neumann iteration Equation (7) that satisfies the hypotheses of Lemma 4.10. Therefore, in this case the iteration to approximate the MTTA is monotonically increasing, and at the step ℓ gives a lower bound for the final value of the MTTA.

Theorem 4.12. *Let $Q = \Delta + R + W$ be an infinitesimal generator of a Markov chain as described in (1), with*

$$R = R^{(1)} \oplus \dots \oplus R^{(k)}, \quad \Delta = -\text{diag}((W + R)e),$$

and $\gamma \geq \|\Delta\|_\infty$. Then, if we define $D := -\gamma I$, $A_1 = R$, and $A_2 = W + (\Delta - \gamma I)$, these matrices satisfy the hypotheses of Theorem 4.8 and there exists α_j, β_j such that

$$D + A_1 = \left(R^{(1)} - \frac{\gamma}{k}I\right) \oplus \left(R^{(2)} - \frac{\gamma}{k}I\right) \oplus \dots \oplus \left(R^{(k)} - \frac{\gamma}{k}I\right),$$

and therefore

$$(D + A_1)^{-1}(A_2 - S) \approx \sum_{j=1}^{\ell} \alpha_j \left(e^{\beta_j(R_1 - \frac{\gamma}{k}I)} \otimes \dots \otimes e^{\beta_j(R_k - \frac{\gamma}{k}I)} \right) (A_2 - S).$$

Remark 4.13. We note that the choice of γ allows to control the condition number of the matrix $D + A_1$; our experience shows that larger values for γ (which give lower condition numbers), provide slower convergence with ρ approaching 1, but also lower TT-ranks during the Neumann iteration. This choice is discussed in further detail in Section 5.2.

5. Computational remarks

In this section we report a few computational remarks concerning our implementation. The variants of the Neumann expansion described in Algorithm 2 and 3 have been implemented in the toolbox `kaes`, which is freely available⁴.

⁴<https://github.com/numpi/kaes/>.

The toolbox is implemented in MATLAB, and given cell-arrays \mathbf{R} , \mathbf{W} containing the factors defining Q , one may compute the value of the MTTA by calling $\mathbf{m} = \text{eval_measure('inv', pi0, r, R, W)}$ where $\mathbf{pi0}$ and \mathbf{r} contain the initial probability distribution and the reward vector. The function has some optional parameters, that allows to tune the required tolerance and the value of γ .

A few considerations can be helpful in trying to obtain maximum performances from the implementation.

5.1. Ordering of the subsystems

Since the TT representation represents the interaction between between the subsystem i and $i + 1$ in each carriage, we have found that it is beneficial to reorder the topology so that few nodes are linked to far ones.

In particular, given the adjacency matrix \mathcal{T} that represents the connection graph (i.e., $\mathcal{T}_{ij} = 1$ if and only if there is an edge in graph from the node i to the node j), it can be helpful to reorder the subsystems to make this matrix as banded as possible. To this end, we have employed the reverse Cuthill-McKee ordering implemented in MATLAB in the function `symrcm`.

5.2. The choice of γ

The choice of the parameter γ in Theorem 4.12 can have important effects on the performance of the algorithm.

We have verified that choosing γ relatively large, for instance $\gamma \gg \|\Delta\|_\infty$, can be helpful. This reduces the conditioning of the matrix to invert to a small number, and thus very few exponential sums are needed to achieve a very high accuracy. More importantly, this helps to keep the TT-ranks low during the iteration, especially when applying (9), which in turn suffers very mildly from having the spectral radius close to 1 (thanks to the quadratic convergence rate). On the other hand, when applying the linearly convergence iteration (7), the minimal choice $\gamma = \|\Delta\|_\infty$ is often advisable. Indeed, for this iteration a ρ close to 1 is much more harmful, and in general it can be quite memory efficient (it only work with compressed vectors).

We do not have a “universal recipe” for these choices, so it might be helpful to do some preliminary parameter tuning on small problems of a given class before tackling the large scale cases. We plan to further investigate this matter in the future.

6. Case study

Consider a cyber-physical system comprising k components, consisting each of a mechanical object and a Monitoring and Control Unit (MCU). The mechanical objects are independent one from the other whereas the working status of the MCU software on component j depends on data produced by local sensors and can depend also on data coming from the MCU of component i . Thus, it is possible to define a topology of interactions among component MCUs: define

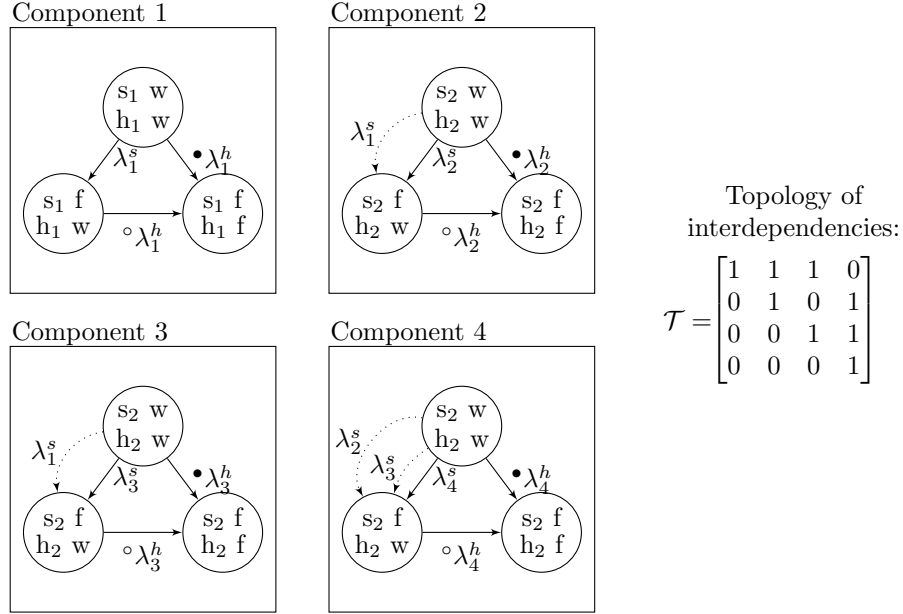


Figure 1: Example SAN for the case study where there are 4 components. Dotted arrows are synchronization transitions.

\mathcal{T} the $k \times k$ matrix as $\mathcal{T}(i, j) = 1$ if $i = j$ or the j -th MCU consumes data produced by the i -th MCU. At every time instant, the MCU and the mechanical object on each component can be *working* or *failed*. If the mechanical object on component i fails then instantaneously also the MCU on component i fails. If $\mathcal{T}(i, j) = 1$ then the failure of the i -th MCU implies an instantaneous failure of the j -th MCU. The failure time of the software running on the i -th MCU is assumed to be exponentially distributed with rate λ_i^s .

The MCU on component i can modify the behaviour of the mechanical object on component i , so the failure time of the mechanical object on component i is exponentially distributed with rate $\bullet\lambda_i^h$ if the MCU on component i is working, and $\circ\lambda_i^h$ if the MCU is already failed. No repair is considered.

We are interested in evaluating the Mean Time to System Failure, where the system is considered failed when all the mechanical objects are failed. Figure 1 depicts the SAN model for a simple case where there are 4 components. In particular, the state of component i , represented with a circle, is defined by software status (s_i is **w** if working, **f** if failed) and the hardware status (h_i is **w** if working, **f** if failed). Transitions can be local or synchronized, represented as arrows and dotted arrows, respectively, and labelled by their rate. Each component model has 3 states, so that $|\mathcal{PS}| = 3^k$, and the cardinality of \mathcal{RS} depends on \mathcal{T} . Notice that, if $\mathcal{T} = I$ then $\mathcal{RS} = \mathcal{PS}$, whereas the size of \mathcal{RS} decreases as the number of interactions increases. The local and synchronization

contribution matrices are then obtained as in Equation (10).

$$R^{(i)} = \begin{bmatrix} 0 & 0 & \bullet\lambda_i^h \\ 0 & 0 & \circ\lambda_i^h \\ 0 & 0 & 0 \end{bmatrix}, \quad W^{(t_j, i)} = \begin{cases} \begin{bmatrix} 0 & \lambda_j^s & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} & \text{if } i = j \\ \begin{bmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} & \text{if } i \neq j \text{ and } \mathcal{T}(j, i) = 1 \\ \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} & \text{otherwise} \end{cases} \quad (10)$$

7. Experimental results

The case study model presented in Section 6 has been implemented in MATLAB [26] and studied applying the method discussed so far. In particular, we consider the following set of parameters:

$$\bullet\lambda_i^h = \frac{i}{10}, \quad \circ\lambda_i^h = i, \quad \lambda_i^s = i,$$

and the topology \mathcal{T} has been chosen at random with the following constraints:

- each component has impact on itself, i.e., $\mathcal{T}(i, i) = 1$,
- For each $i \neq j$, the entry $\mathcal{T}(i, j)$ is set to 1 with probability $\frac{1}{2k}$.

More precisely, the sparse matrix \mathcal{T} has been generated using the MATLAB command `T = (speye(k)+sprand(k,k,.5/k)) > 0`.

We have tested the values of $k = \{10, 12, \dots, 32\}$. For each value of k , we have run 100 tests for Algorithm 3 and for Algorithm 2, generating random topologies \mathcal{T} . The tests have been performed on a node of a cluster with two Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz processors each. The processes have been limited to 20 GB of RAM and 4 threads each, with a time limit of 600 hours.

The results for what concern memory usage are reported in Table 1, and for runtime in Table 2. We note that, despite Algorithm 2 and 3 being equivalent (in the sense given in Section 4), the quadratic convergence of Algorithm 2 makes it the best choice on all the tests.

We note that Algorithm 3 has more predicable runtimes. In Figure 2, it is visible that they appear to have a cubic dependency on k for the case study under consideration. Algorithm 2, on the other hand, has timings with a weaker correlation with k ; from our observations, they appear to be related to the topology, which influences the growth of the TT-ranks during the iterations.

k	Algorithm 2		Algorithm 3		AMEn			DMRG		
	Avg	Max	Avg	Max	Avg	Max	OOM	Avg	Max	OOM
2	0.53	0.54	0.53	0.54	0.53	0.54	0%	0.53	0.54	0%
4	0.56	0.6	0.53	0.55	0.57	0.6	0%	0.54	0.57	0%
6	0.57	0.6	0.55	0.59	0.58	0.61	0%	0.57	0.6	0%
8	0.58	0.6	0.57	0.6	0.59	0.63	0%	0.61	0.7	0%
10	0.58	0.64	0.57	0.59	0.61	0.66	0%	0.65	0.72	0%
12	0.63	0.85	0.61	0.69	0.65	0.76	10%	0.68	0.73	0%
14	0.66	1.05	0.67	0.83	0.71	0.78	63%	0.71	0.76	0%
16	0.72	1.12	0.74	0.86	0.74	0.76	98%	0.73	0.77	0%
18	0.82	2.51	0.84	1.49			100%	0.75	0.79	10%
20	0.87	2.51	0.93	1.48			100%	0.76	0.79	66%
22	0.94	1.7	1.04	2.01			100%			100%
24	1.02	7.25	1.19	5.78			100%			100%
26	0.99	3.34	1.33	5.87			100%			100%
28	1.06	8.68	1.59	4.22			100%			100%
30	0.97	1.95	1.74	5.08			100%			100%
32	1.14	4.06	2.1	5.94			100%			100%

Table 1: Average and maximum memory usage (in GB) required by Algorithm 3 and 2, AMEn, and DMRG for computing the MTTA of the case study in Section 6. The percentage of tests that ran out of memory is reported in the column OOM. The results are obtained by running 100 random tests for each k , and taking average and maximum of the tests which did not encounter an OOM condition. This situation was never encountered for Algorithm 2 and 3.

This can be advantage, in the sense that even large scale cases might be treatable, or a disadvantage, because it makes very hard to predict how long the algorithm will need to give an answer. In particular, we have observed that for these problems the choice of a good parameter γ (as described in Section 5.2) is more important. In the case study under consideration, the dependency of Algorithm 3 on k appears to have an asymptotic behavior close to $\mathcal{O}(k^{3.5})$, as reported in Figure 2.

To show a typical behavior of the TT-ranks during the iteration of Algorithm 2 we have reported two runs for $k = 24$, whose evolution of the maximum TT-rank of M is reported in Figure 3. The two examples have been chosen one below and the other above the average runtime for this value of k . It is visible how the slowest of the two runs reaches a higher rank (40) than the other (which only gets up to 18).

We compared the results with the AMEn solver [11] and the DMRG algorithm [29], both available in the TT-Toolbox [27]. The AMEn solver has been proved to be quite effective for the computation of the steady-state vector of irreducible Markov chains in [21]. The solver can be used to compute $(Q - S)^{-1}v$, or to solve the normal equations $(Q - S)^T(Q - S)x = (Q - S)^T v$. The former problem is better conditioned, but the latter is symmetric positive definite, which guarantees convergence for the AMEn iteration. We have compared both choices, and we found that for this case study the second performs slightly better. However, the method stagnates on an increasing number of cases

k	Algorithm 2	Algorithm 3	AMEn		DMRG	
	Avg	Avg	Avg	OOM	Avg	OOM
2	0.4	0.55	0.38	0%	0.13	0%
4	0.57	2.03	0.73	0%	0.23	0%
6	0.74	5.71	2.86	0%	0.31	0%
8	0.99	12.9	26.16	0%	0.57	0%
10	1.34	25.53	217.83	0%	1	0%
12	2.26	45.99	1,943.46	10%	1.51	0%
14	2.87	75.92	4,865.38	63%	2.17	0%
16	3.86	116.69	7,811.5	98%	4.07	0%
18	6.21	179.51		100%	36.61	10%
20	7.89	259.49		100%	151.04	66%
22	9.16	361.71		100%		100%
24	13.88	516.21		100%		100%
26	11.69	667.32		100%		100%
28	16.97	935.96		100%		100%
30	10.72	1,135.37		100%		100%
32	18.57	1,584.34		100%		100%

Table 2: Average wall-clock time (in seconds) required to compute the MTTA for the case study with k components. OOM denotes the percentage of tests which were stopped because they ran out of memory. Each run was allocated 20GB of RAM and 4 logical cores. The average are computes on the tests which did not run out of memory. For Algorithm 2 and 3, this situation was never encountered.




Figure 2: Average timings as a function of k for Algorithm 3. For this case study, the timings appear to depend on k polynomially, with exponent close to 3.5.




Figure 3: Maximum TT-Rank of M during the iteration of Algorithm 2 for two runs with $k = 24$. The runs with smaller ranks took 4.55s, while the other needed 72.30s.

when $k > 10$, so we could only make a direct comparison in Table 1 and 2 for small values of k . DMRG, on the other hand, performed more favorably on the case study, and we have been able to solve problems (quite) reliably for k up to 18.

8. Conclusions

We have shown that tensor trains are a powerful tool for the analysis of performance and reliability measures (in this case, the mean time to failure) of large systems, when the interconnection between the smaller subsystems that compose them is sufficiently weak.

We have presented a theoretical analysis of an iteration that is easily applicable in the tensorized format, and with guaranteed convergence. A quadratically convergent variation has been shown as well, and the performances have been tested on a representative set of examples.

Several lines of research remain open: the connection between the weak connections and the TT-rank in the iteration needs to be studied further, in order to understand the relation more in depth. Moreover, several more measures are of interest in this context, and the application of tensor techniques for this task could lead to faster and reliable methods for their computation.

We have introduced some techniques for accelerating the operations in tensor formats for the iterations that arise from the Markovian context, and that will be subject of future study.

References

- [1] Berman, A., Plemmons, R. J., 1994. Nonnegative matrices in the mathematical sciences. Vol. 9 of Classics in Applied Mathematics. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, revised reprint of the 1979 original.
URL <https://doi.org/10.1137/1.9781611971262>
- [2] Bolten, M., Kahl, K., Kressner, D., Macedo, F., Sokolović, S., 2016. Multi-grid methods combined with low-rank approximation for tensor structured markov chains. arXiv preprint arXiv:1605.06246.
- [3] Braess, D., Hackbusch, W., 2005. Approximation of $1/x$ by exponential sums in $[1, \infty)$. IMA journal of numerical analysis 25 (4), 685–697.
- [4] Buchholz, P., Dayar, T., Kriege, J., Orhan, M. C., 2017. On compact solution vectors in Kronecker-based Markovian analysis. Performance Evaluation 115, 132–149.
- [5] Buchholz, P., Kemper, P., 2004. Kronecker based matrix representations for large Markov models. Springer, pp. 256–295.
- [6] Ciardo, G., Miner, A. S., 1999. A data structure for the efficient Kronecker solution of GSPNs. In: Proceedings 8th International Workshop on Petri Nets and Performance Models (Cat. No.PR00331). pp. 22–31.
- [7] Crouzeix, M., Palencia, C., 2017. The numerical range as a spectral set. arXiv preprint arXiv:1702.00668.
- [8] De Lathauwer, L., De Moor, B., Vandewalle, J., 2000. A multilinear singular value decomposition. SIAM Journal of Matrix Analysis and Applications 21 (4), 1253–1278.
URL <https://doi.org/10.1137/S0895479896305696>
- [9] de Silva, V., Lim, L.-H., 2008. Tensor rank and the ill-posedness of the best low-rank approximation problem. SIAM Journal of Matrix Analysis and Applications 30 (3), 1084–1127.
URL <https://doi.org/10.1137/06066518X>
- [10] Demmel, J. W., 1997. Applied numerical linear algebra. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA.
URL <https://doi.org/10.1137/1.9781611971446>
- [11] Dolgov, S. V., Khoromskij, B. N., Oseledets, I. V., 2012. Fast solution of parabolic problems in the tensor train/quantized tensor train format with initial application to the Fokker-Planck equation. SIAM Journal of Scientific Computing 34 (6), A3016–A3038.
URL <https://doi.org/10.1137/120864210>

- [12] Dolgov, S. V., Savostyanov, D. V., 2014. Alternating minimal energy methods for linear systems in higher dimensions. *SIAM Journal on Scientific Computing* 36 (5), A2248–A2271.
- [13] Donatelli, S., 1993. Superposed stochastic automata: A class of stochastic petri nets with parallel solution and distributed state space. *Performance Evaluation* 18 (1), 21–36.
- [14] Grasedyck, L., 2010. Hierarchical singular value decomposition of tensors. *SIAM Journal on Matrix Analysis and Applications* 31 (4), 2029–2054.
- [15] Hackbusch, W., 2019. Computation of best l^∞ exponential sums for $1/x$ by remez algorithm. *Computing and Visualization in Science* 20 (1-2), 1–11.
- [16] Higham, N. J., 2008. *Functions of matrices. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, theory and computation.* URL <https://doi.org/10.1137/1.9780898717778>
- [17] Hillar, C. J., Lim, L.-H., 2013. Most tensor problems are np-hard. *Journal of the ACM (JACM)* 60 (6), 45.
- [18] Hillston, J., 1996. *A Compositional Approach to Performance Modelling.* Cambridge University Press, New York, NY, USA.
- [19] Kazeev, V. A., Khoromskij, B. N., 2012. Low-rank explicit QTT representation of the Laplace operator and its inverse. *SIAM Journal of Matrix Analysis and Applications* 33 (3), 742–758. URL <https://doi.org/10.1137/100820479>
- [20] Kolda, T. G., Bader, B. W., 2009. Tensor decompositions and applications. *SIAM Rev.* 51 (3), 455–500. URL <https://doi.org/10.1137/07070111X>
- [21] Kressner, D., Macedo, F., 2014. Low-rank tensor methods for communicating markov processes. In: *International Conference on Quantitative Evaluation of Systems.* Springer, pp. 25–40.
- [22] Kressner, D., Tobler, C., 2014. Algorithm 941: `htucker`—a Matlab toolbox for tensors in hierarchical Tucker format. *ACM Trans. Math. Software* 40 (3), Art. 22, 22. URL <https://doi.org/10.1145/2538688>
- [23] Lubich, C., Oseledets, I. V., Vandereycken, B., 2015. Time integration of tensor trains. *SIAM Journal of Numerical Analysis* 53 (2), 917–941. URL <https://doi.org/10.1137/140976546>
- [24] Masetti, G., Robol, L., 2018. Computing performability measures in markov chains by means of matrix functions. arXiv preprint arXiv:1803.06322.

- [25] Masetti, G., Robol, L., Chiaradonna, S., Di Giandomenico, F., 2019. Stochastic evaluation of large interdependent composed models through kronecker algebra and exponential sums. In: Application and Theory of Petri Nets and Concurrency. pp. 47–66.
- [26] MathWorks, 2018. MATLAB R2018a. The Mathworks, Inc., Natick, Massachusetts.
- [27] Oseledets, I., Dolgov, S., Kazeev, V., Lebedeva, O., Mach, T., 2019. MATLAB TT-Toolbox. <https://github.com/oseledets/TT-Toolbox>.
- [28] Oseledets, I. V., 2011. Tensor-train decomposition. SIAM Journal of Scientific Computing 33 (5), 2295–2317.
URL <https://doi.org/10.1137/090752286>
- [29] Oseledets, I. V., Dolgov, S. V., 2012. Solution of linear systems and matrix inversion in the tt-format. SIAM Journal on Scientific Computing 34 (5), A2718–A2739.
- [30] Plateau, B., Stewart, W. J., 2000. Stochastic Automata Networks. Springer US, Boston, MA, pp. 113–151.
- [31] Trivedi, K. S., Bobbio, A., 2017. Reliability and Availability Engineering: Modeling, Analysis, and Applications. Cambridge University Press.
- [32] Trivedi, K. S., Bobbio, A., 2017. Reliability and Availability Engineering: Modeling, Analysis, and Applications. Cambridge University Press.