Stochastic Evaluation of Large Interdependent Composed Models Through Kronecker Algebra and Exponential Sums

G. Masetti, L. Robol, S. Chiaradonna and F. Di Giandomenico

ISTI-CNR Pisa, University of Pisa

Presentation outline

- Context (state of the art)
- Focus of the analysis and KAES methodology
- Case study and numerical evaluation
- Conclusions and future work

Focus of the analysis

- Stochastic Petri Nets context
- Irreducible CTMCs
- CTMCs with absorbing states

Focus of the analysis

- Stochastic Petri Nets context
- Irreducible CTMCs
- CTMCs with absorbing states
    - Definition of reward structure is crucial
    - System comprising a large number of weakly interdependent components
- Goal 1: extend the class of systems for which Mean Time To Absorption can be evaluated
- Goal 2: safe evaluation of MTTA (lower bound)
- Goal 3: foresee a road map to evaluate also transient measures

Structure of this presentation

- List of problems we encountered when adapting techniques for the evaluation of
  - steady-state probability vector (irreducible CTMCs)
  - transient measures (relatively small CTMCs)

  to the solution of CTMCs with absorbing states
- Often solving a given problem produces a new set of problems
- For each problem I will discuss the proposed solution, observing that a complete comparison among available solutions is out of the paper scope

Context

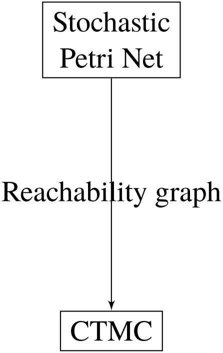
Stochastic
Petri Net
Reachability graph
CTMC

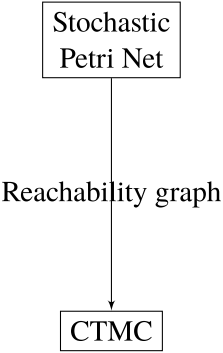

finite state space $\mathcal{S}$

infinitesimal generator $Q \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$

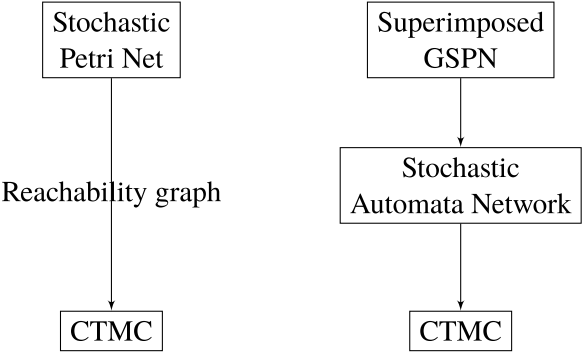

finite state space $\mathcal{S}$

infinitesimal generator $Q \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$

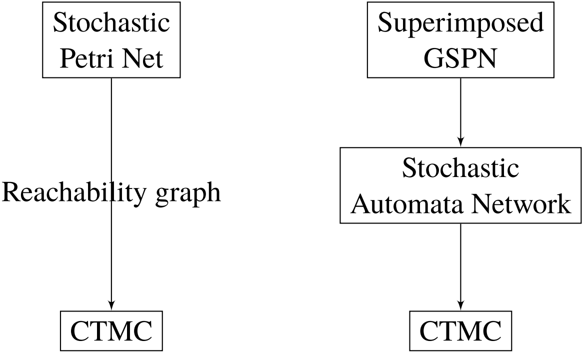

finite state space $\mathcal{S}$

small state spaces $\mathcal{S}^{(1)}, \ldots, \mathcal{S}^{(n)}$

potential states $\mathcal{PS} = \mathcal{S}^{(1)} \times \cdots \times \mathcal{S}^{(n)}$

infinitesimal generator $Q \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$

infinitesimal generator $Q \in \mathbb{R}^{|\mathcal{PS}| \times |\mathcal{PS}|}$

Lumping, MDD, Kronecker, Reachable

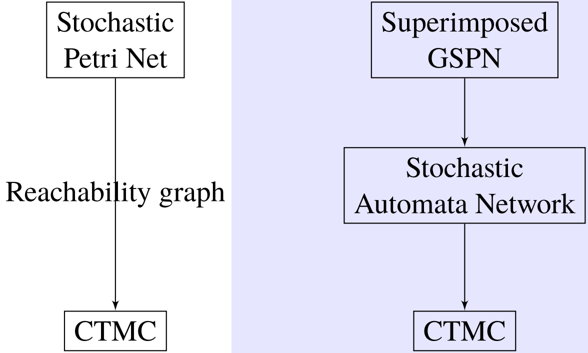

finite state space $\mathcal{S}$

infinitesimal generator $Q \in \mathbb{R}^{|\mathcal{S}|\times|\mathcal{S}|}$

small state spaces $\mathcal{S}^{(1)}, \ldots, \mathcal{S}^{(n)}$

potential states $\mathcal{PS} = \mathcal{S}^{(1)} \times \cdots \times \mathcal{S}^{(n)}$

infinitesimal generator $Q \in \mathbb{R}^{|\mathcal{PS}|\times|\mathcal{PS}|}$

Descriptor matrix

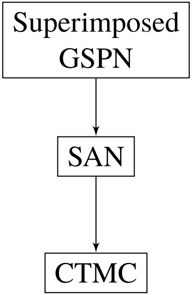

Descriptor matrix

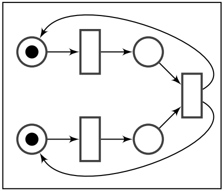

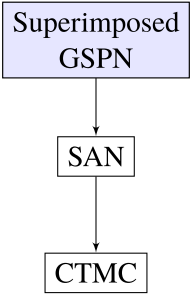

Descriptor matrix

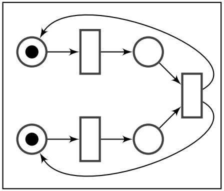

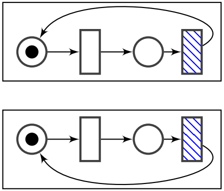

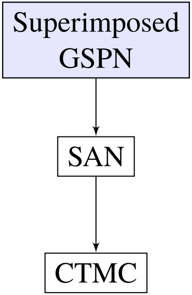

Descriptor matrix

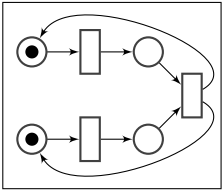

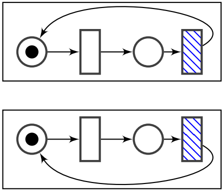

$$s_1^1 = (1,0)$$
$$s_2^1 = (0,1)$$

$$s_1^2 = (1,0)$$
$$s_2^2 = (0,1)$$

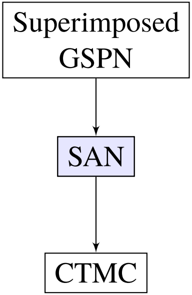

Descriptor matrix

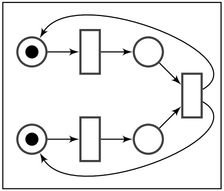

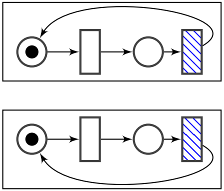

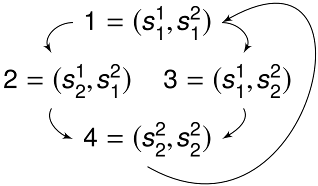

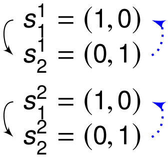

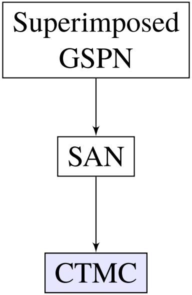

Descriptor matrix

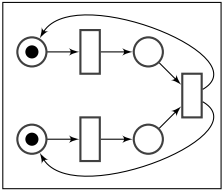

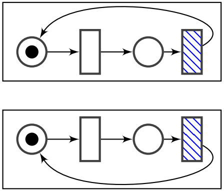

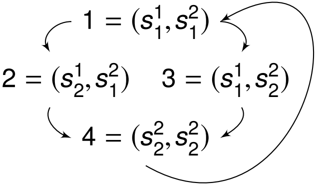

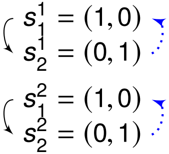

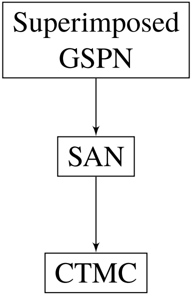

Descriptor matrix

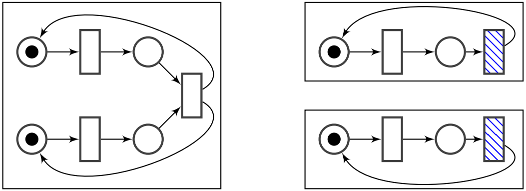

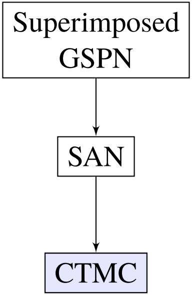

$$\begin{bmatrix} & \times & \times & \\ & & & \times \\ & & & \times \\ \times & & & \end{bmatrix} = \begin{bmatrix} & \times \\ & \end{bmatrix} \oplus \begin{bmatrix} & \times \\ & \end{bmatrix} + \begin{bmatrix} & \\ \times & \end{bmatrix} \otimes \begin{bmatrix} 1 & \\ & 1 \end{bmatrix} + \begin{bmatrix} 1 & \\ & 1 \end{bmatrix} \otimes \begin{bmatrix} & \\ \times & \end{bmatrix}$$

Descriptor matrix — R local — W synchronization

Descriptor matrix

steady-state π s.t. $\begin{cases} \pi Q & = 0 \\ \pi \mathbb{1}^T & = 1 \end{cases}$

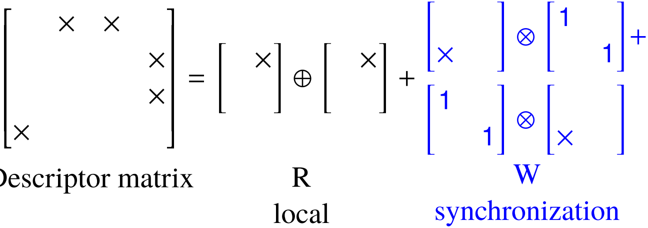

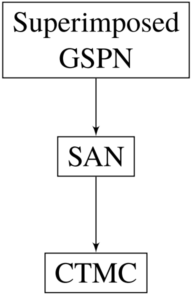

Descriptor vectors

- Recent development: compress both Q and π exploiting the Kronecker structure
  - Kressner, 2014: Tensor-Trains
  - Buchholz, 2017: Hierarchical Tucker Decomposition
- We followed Kressner

$$\mathcal{A}_{i_1 \dots i_d} = \underbrace{G_1[i_1]}_{1\times r} \underbrace{G_2[i_2]}_{r\times r} \dots \underbrace{G_d[i_d]}_{r\times 1}$$

An example of computing one element of 4-dimensional tensor:

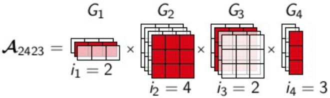

- Standard iterative methods to evaluate π fail because in $\pi^{k+1} = \pi^k + \delta\pi^k$ the TT-ranks can grow too quickly
- Thus, ad hoc solution methods have been designed

CTMCs with an absorbing state

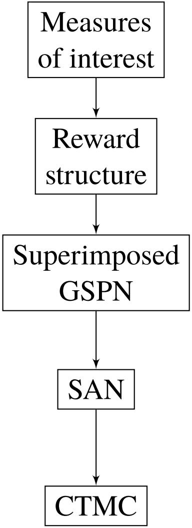

CTMCs with an absorbing state

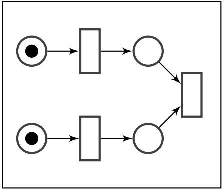

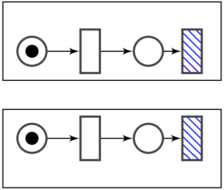

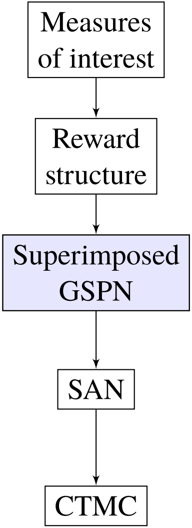

CTMCs with an absorbing state

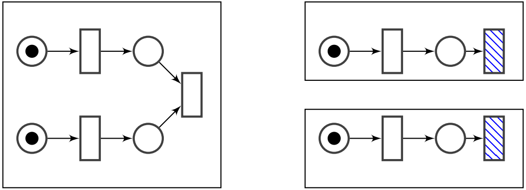

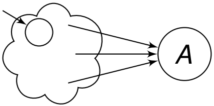

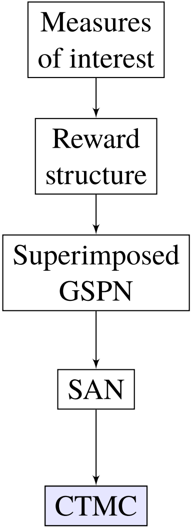

CTMCs with an absorbing state

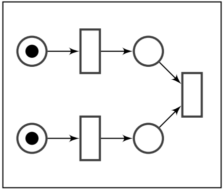

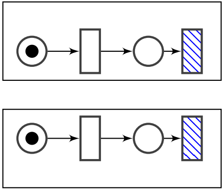

$$Q = \left[\begin{array}{ccc|c} & & & v_1 \\ & \hat{Q} & & \vdots \\ & & & v_{N-1} \\ \hline 0 & \ldots & 0 & 0 \end{array}\right]$$

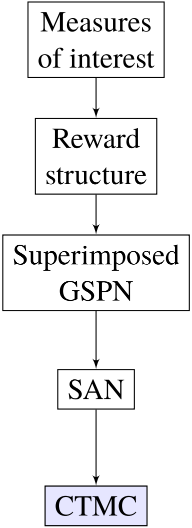

CTMCs with an absorbing state

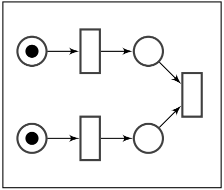

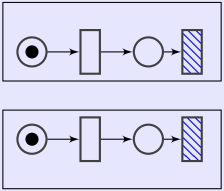

$$Q = \left[\begin{array}{ccc|c} & & & v_1 \\ & \hat{Q} & & \vdots \\ & & & v_{N-1} \\ \hline 0 & \dots & 0 & 0 \end{array}\right] \qquad r = \begin{bmatrix}1\\1\\1\end{bmatrix} \otimes \begin{bmatrix}1\\1\\1\end{bmatrix} - \begin{bmatrix}0\\0\\1\end{bmatrix} \otimes \begin{bmatrix}0\\0\\1\end{bmatrix}$$

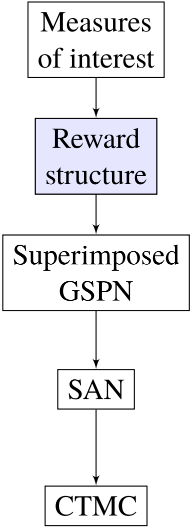

CTMCs with an absorbing state

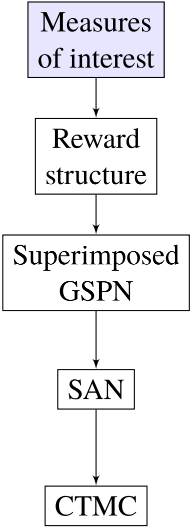

$$\text{MTTA} = -\hat{\pi}_0 \, \hat{Q}^{-1} \, \mathbb{1}^T$$

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

Focus of the analysis

Problems identification

- Shifting Q to obtain an irreducible CTMC as

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

  assuming $(\pi_0)_i \neq 0$ only for $i = 1$, introduces many difficulties
- If Q has a nice TT decomposition it is not guaranteed that $\hat{Q}$ has a nice TT decomposition too
- Even a good compression (small TT-ranks) for $\hat{Q}$ does not guarantee a good compression for the vectors involved in iterative solution methods

Proposed solutions and new problems (I)

Proposed solution:

- Define the new shift $Q - S$ as

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

- $Q - S$ and Q have the same Kronecker structure and the TT-ranks of S are $(1, 1, \dots, 1)$
- We can solve the system with $Q - S$ because

$$MTTA = -\hat{\pi}_0 \hat{Q}^{-1} \mathbb{1}^T = -\pi_0 (Q - S)^{-1} r^T,$$

provided that $(\pi_0)_N = 0$

Problem:

- The methods studied by Kressner and Buchholz do not work for $Q - S$ because the TT-ranks grow too much anyway

Proposed solutions and new problems (II)

Proposed solution:

- Define a new split $Q = Q_1 + Q_2$ such that

$$Q - S = Q_1 + Q_2 - S = \underbrace{R + \Delta'}_{Q_1} + \underbrace{W + \Delta - \Delta'}_{Q_2} - S$$

$$= Q_1\Big(I + Q_1^{-1}(Q_2 - S)\Big) = Q_1\Big(I - M\Big)$$

where $M = -Q_1^{-1}(Q_2 - S)$ and Δ' is a suitable Kronecker sum of diagonal matrices (details in the paper)

- Exploit the Neumann series

$$(Q-S)^{-1} = (I-M)^{-1}Q_1^{-1} = Q_1^{-1} + MQ_1^{-1} + M^2Q_1^{-1} + M^3Q_1^{-1} + \ldots = \sum_{j=0}^{\infty} M^j Q_1^{-1}$$

Problem:

- Is this a convergent series?

Proposed solutions and new problems (III)

Proposed solution:

- Of course it is a convergent series (Theorem in the paper) because the spectral radious of M, called $\rho(M)$, is strictly less than 1

Problem:

- Is the inversion of Q_1 computationally cheaper than inverting $Q - S$?

Proposed solution:

- Consider the exponential sums approximation

$$\frac{1}{x} \approx \sum_{j=1}^{\ell} \alpha_j e^{-\beta_j x},$$

  which has a controlled error bound on $[1, +\infty)$
- Exploit the matrix exponential $e^{Q_1} = I + \frac{1}{2}Q_1^2 + \frac{1}{6}Q_1^3 + \ldots$
- Exploit the property $e^{Q_1^1 \oplus \cdots \oplus Q_1^n} = e^{Q_1^1} \otimes \ldots \otimes e^{Q_1^n}$

Putting all together

- The core of the computation is

$$MTTA = -\pi_0(Q-S)^{-1}r^T \approx -\pi_0 x^k, \text{ where } \begin{cases} x^0 & = Q_1^{-1}r^T \\ x^{k+1} & = x^k + Mx^k \end{cases}$$

- The matrix-vector product is evaluated through

$$Mx = -Q_1^{-1}(Q_2 - S)x = -(Q_1^1 \oplus \cdots \oplus Q_1^n)^{-1}(\Delta - \Delta' + W - S)x \approx$$
$$\approx \sum_{j=1}^{\ell} \alpha_j \left(e^{\beta_j(R^1+\Delta_1')} \otimes \ldots \otimes e^{\beta_j(R^n+\Delta_n')} \right) (\Delta - \Delta' + W - S)x,$$

where all the matrices and x are in TT format.

Safe approximation

Remark

Notice that, defining $z^{k+1} = Q_1^{-1}(Q_2 - S)z^k$ *and* $z^0 = Q_1^{-1}(\mathbb{1} - e_N^T Q_1^{-1} \mathbb{1} \cdot e_N)$*, we obtain*

$$MTTA = -\pi_0^T \cdot z^k + \mathcal{O}(\rho(M)^{k+1}),$$

where $z^{k+1} \geq z^k$ *for all* $k = 0, 1, \ldots$ *because* $e_N^T z^0 = 0$ *and both* Q_1^{-1} *and* Q_2 *are non-negative matrices.*

This means that the MTTA can be computed in a safe way, being the approximation $-\pi_0^T \cdot z^k$ a lower bound

computations saving and acceleration techniques

- Instead of implementing Mx we can work with M^T exploiting the fact that
  - from a state in $\mathcal{S}$ we cannot reach a state in $\mathcal{PS} \setminus \mathcal{S}$
  - from a state in $\mathcal{PS} \setminus \mathcal{S}$ we can reach a state in $\mathcal{S}$
- Consider the following series

$$(I - M)^{-1} = (I + M)(I + M^2) \cdots (I + M^{2^k}) \cdots$$

  it is equivalent to the Neumann series, but converges quadratically
- Here we have to implement the matrix-matrix product instead of the matrix-vector product, and the TT-ranks can grow quickly

Case study

GSPN for C_i

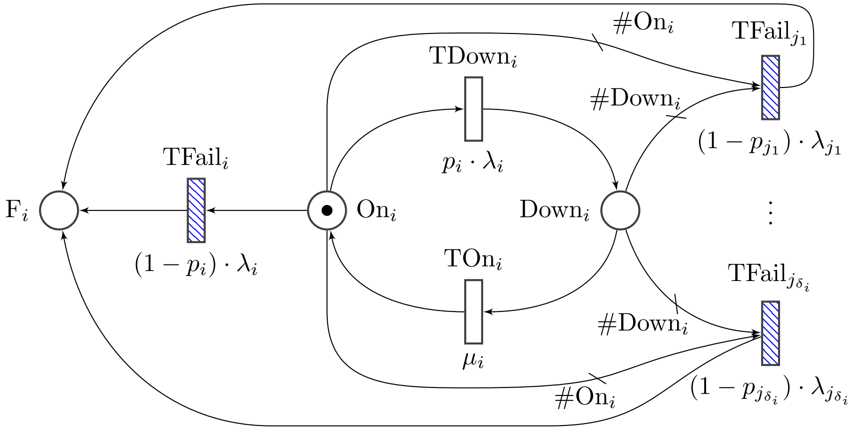

The failure of C_i can impact on $C_{j_1},\ldots,C_{j_{\delta_i}}$

Evaluation results

- Consider $n = 10, 20, \ldots, 50$
- Number of potential states: $|\mathcal{PS}| = 3^n$
- Topology of interactions is obtained as follows:
  - a star topology is constructed, where, labeling the nodes from 1 to n, there exist $n - 1$ edges connecting 1 to j, for $j = 2, \ldots, n$
  - for each node with index greater than 1, another edge connecting it to a random node is added with probability 0.2

  Such topologies are good representatives of topologies addressed by KAES: large number of components, loosely interconnected
- The remaining parameters are chosen at random within the following intervals

$$\lambda_i \in [0.5, 1.5], \quad \mu_i \in [2000, 3000], \quad p \in [0.95, 1],$$

so that there are 4 orders of magnitude among the parameters. The tests have been repeated 100 times for each value of n, using the randomized topology described above

Table: Potential spaces dimensions, memory consumption, time and number of cases where the KAES approach was successful, where μ reports the average over the 100 runs and σ is the standard deviation.

n	$\|\mathcal{PS}\|$	memory (Gb)		time (s)		% solved cases
		μ	σ	μ	σ	
10	59049	0.90	0.08	1.17	0.81	100%
20	$3.49 \cdot 10^{9}$	3.07	9.68	65.83	346.24	100%
30	$2.06 \cdot 10^{14}$	8.31	19.40	193.29	619.63	91%
40	$1.22 \cdot 10^{19}$	4.42	9.97	140.89	477.67	91%
50	$7.18 \cdot 10^{23}$	7.79	17.27	299.44	840.78	84%

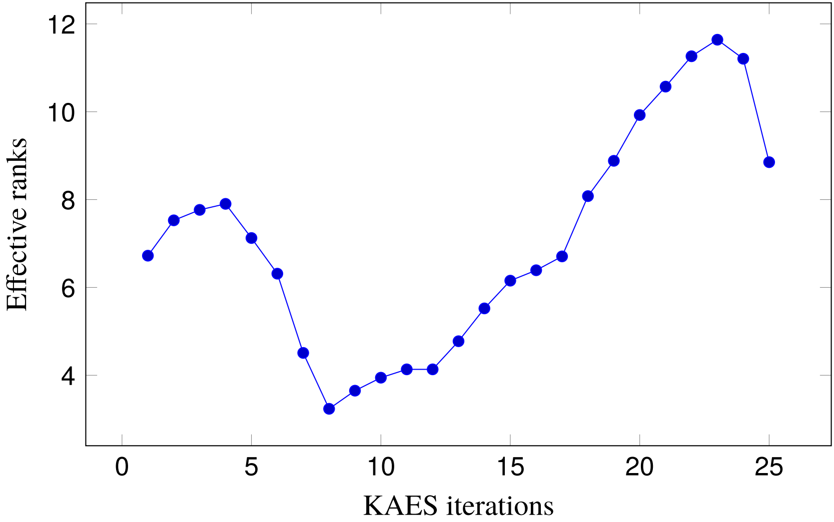

Figure: Evolution of the effective ranks, representing an average of the TT-ranks of the carriages, for each iteration of KAES, with $n = 20$ and a specific topology.

Conclusions and future work

Conclusions

- Analytical modeling of large, interconnected systems by developing a new numerical evaluation approach called KAES
- Focus on Mean Time To Absorption
- Symbolic representation of both the descriptor matrix and the descriptor vector to mitigate the state space explosion
- Although symbolic representation has been already applied in existing studies, such previous works focus on steady-state analysis
- KAES targets limiting analysis in presence of absorbing states
- The way MTTA is computed guarantees a safe assessment, which is relevant when dealing with dependability critical applications
- We started a numerical evaluation campaign, where the presented case study is the first step

Future work

- More experiments are needed to better understand strengths and limitations of this new technique
- A deeper understanding of the link between TT-ranks and the topology of interactions among system components would be desirable
- The powerfulness of the adopted techniques make this method not restricted to the evaluation of MTTA only, but adaptable to evaluate general performability related indicators
- We are working on a general treatment of reward vectors to allow the modeler to define them at SGSPN level maintaining both the Kronecker structure and a good compression

Future work

- We have recently published a new way of interpreting performability measures in terms of matrix functions
- Being able to evaluate efficiently $f(\tilde{Q} - S)$ where

$$f(z) = \begin{cases} \frac{1}{z} & \text{if } z \neq 0, \\ 0 & \text{otherwise,} \end{cases}$$

is the first step to evaluate transient measures that can be expressed through

$$\varphi_j(z) = \begin{cases} \frac{\varphi_{j-1}(z)-1}{z} & \text{if } j > 1, \\ \frac{e^z-1}{z} & \text{if } j = 1. \end{cases}$$

Thank you

Questions?