Tight bounds of performance metrics computed on Markovian queuing models

Gerardo Rubino

Tight bounds of performance metrics computed on Markovian queuing models

G. Rubino

Denver, JMM'20

Outline

1. Introduction
2. Initial technique
3. Improvements

Outline

Context

- performance (or dependability, or performability) analysis
- systems in equilibrium
- continuous time Markov rewards models
- the underlying Markov chain is homogeneous and irreducible

Context

- performance (or dependability, or performability) analysis
- systems in equilibrium
- continuous time Markov rewards models
- the underlying Markov chain is homogeneous and irreducible

- the state space is very large, possibly infinite
- there is no closed-form expression of π

Context

- performance (or dependability, or performability) analysis
- systems in equilibrium
- continuous time Markov rewards models
- the underlying Markov chain is homogeneous and irreducible

- the state space is very large, possibly infinite
- there is no closed-form expression of π

- **Goal:** accurate numerical analysis in spite of previous points

Main assumption

Transitions are either *slow* or *fast*. Two examples:

- models of highly dependable systems:
  - failures $\longrightarrow$ slow transitions,
  - repairs $\longrightarrow$ fast transitions;
- queueing models in light traffic environments:
  - arrivals $\longrightarrow$ slow transitions,
  - departures $\longrightarrow$ fast transitions.

Main assumption

Transitions are either *slow* or *fast*. Two examples:

- models of highly dependable systems:
  - failures $\longrightarrow$ slow transitions,
  - repairs $\longrightarrow$ fast transitions;
- queueing models in light traffic environments:
  - arrivals $\longrightarrow$ slow transitions,
  - departures $\longrightarrow$ fast transitions.
- We want to evaluate

$$R = \mathbb{E}(r_{X_\infty}) = \sum_x r_x \pi_x,$$

where X_∞ is a stationary version of the ergodic stochastic process X,
$r_x \geq 0$ is the reward associated with state x,
π is the distribution of X_∞.

Intuition

- Because of previous setting, X spends "most of its life" in a "small" part of the state space, "close to 0", "around the left side of the space".
- Idea: try to keep this "left side" of the state space and replace the rest by "a few states", obtaining some new model such that evaluating the target on it we will get a good approximation of R.
- Actually, **this leads to several such auxiliary small models, that allow to build some lower and upper bounds of** R.
- To denote lower bounds of some real w we will use the notation $[w]_{lb}$, and $[w]_{ub}$ for an upper bound.

On the target

- Examples: if the system is a queue in equilibrium, or some subset of a queueing network, R can be the mean number of customers in it, or the probability that it has more than N customers; in a dependability model, R can be the asymptotic availability, or some mean cost, etc.
- The bounding approach described here can handle many large (sometimes infinite) Markov models with very general structures.
- It can solve for models where matrix geometric techniques cannot be applied.
- We assume that we cannot compute $\sum_x r_x \pi_x$ (too costly, for instance, because computing π_x for all x is too costly itself) but that upper and lower bounds of the rewards are available.

Example 1

- Consider first the $E_2/M/1/H$ queue, with Erlang-2 inter-arrival distribution with parameter η and service rate μ.
- This is easy to solve numerically, of course. The phased representation of this model is a Markov chain with $2(H+1)$ states,

Example 1

- Consider first the $E_2/M/1/H$ queue, with Erlang-2 inter-arrival distribution with parameter η and service rate μ.
- This is easy to solve numerically, of course. The phased representation of this model is a Markov chain with $2(H+1)$ states,
- Just for checking purposes, if I take a light traffic example using $\eta = 9.0$, $\mu = 8.0$ and $H = 100$, and if the target is $M =$ mean number of customers, in equilibrium, we obtain with our method, for instance, $[M]_{lb} = 1.0271$ and $[M]_{ub} = 1.3253$ solving models having approximately 25% of states than X (exact value: $M \approx 1.0272$).
- With models having around $1/3$ of X's states, we get $[M]_{lb} = [M]_{ub}$, that is, $= 1.0272$.

Example 2

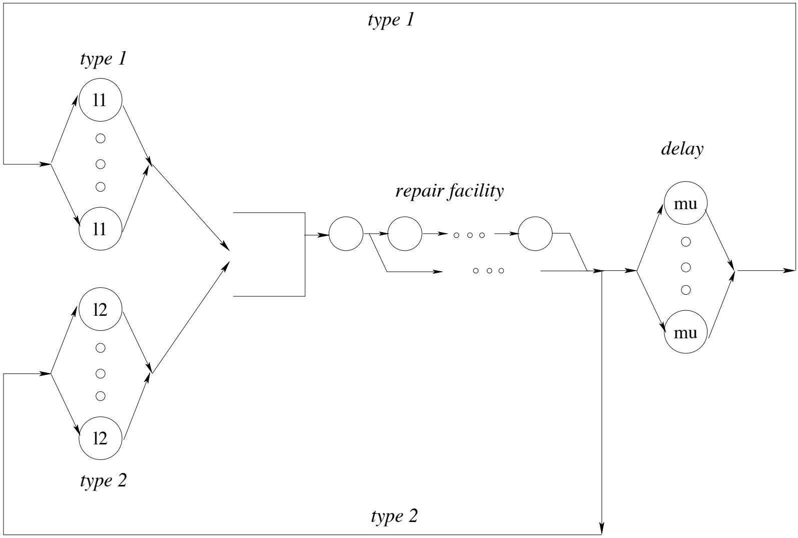
type 1
type 1
l1
l1
l2
l2
type 2
repair facility
delay
mu
mu
type 2

- Previous picture shows a queueing network seen as a generalization of the Machine Repair Model. It has two classes of customers, and a Coxian server, making that, in the general case, the network is not a product-form one
- Putting, for instance, 80 type 1 units and 120 type 2, a 6-phases Cox distribution for the repair time of type-1 machines, and a 5-phases one for type 2 ones, we obtain a model with 4,344,921 states.
- Assume that the system works iff at least 79 type 1 units and 115 type 2 ones are alive.

- Previous picture shows a queueing network seen as a generalization of the Machine Repair Model. It has two classes of customers, and a Coxian server, making that, in the general case, the network is not a product-form one
- Putting, for instance, 80 type 1 units and 120 type 2, a 6-phases Cox distribution for the repair time of type-1 machines, and a 5-phases one for type 2 ones, we obtain a model with 4,344,921 states.
- Assume that the system works iff at least 79 type 1 units and 115 type 2 ones are alive.
- If the target is the asymptotic availability, look at the following results:

# of aux. states	lower bound	upper bound
226	0.9997597**121**	0.9997597**466**
1826	0.9997597349	0.9997597349

Example 3

- This $M/E_2/1 -> ./E_2/1$ tandem has an infinite state space and no analytic solution.
- For instance, if we set $\lambda = 0.18$ and $\nu = 1$ we obtain a mean number of customers $= 1$ with error less than 0.01, solving a linear system having around 200 elements.

Example 4

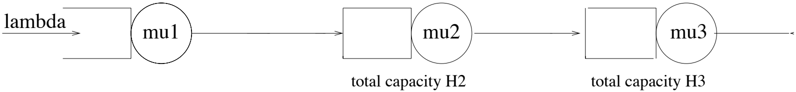

- Here, servers 1 and 2 block after service if next node is full, and when a position is freed they instantaneously send the served (and blocked) customer.
- Again, we have an infinite state space and no analytic solution.
- Using $\lambda = 0.2$, $\mu_1 = 0.7$, $\mu_2 = 1.5$, $\mu_3 = 0.2$, $H_1 = 18$, $H_2 = 10$, we obtain the metric $\mathbb{E}(N_\infty)$ with error $< 10^{-10}$ solving a linear system with size 3839.

Outline

Transformation

Transformation
$X \to X^{(j)}$

G, j: arbitrary

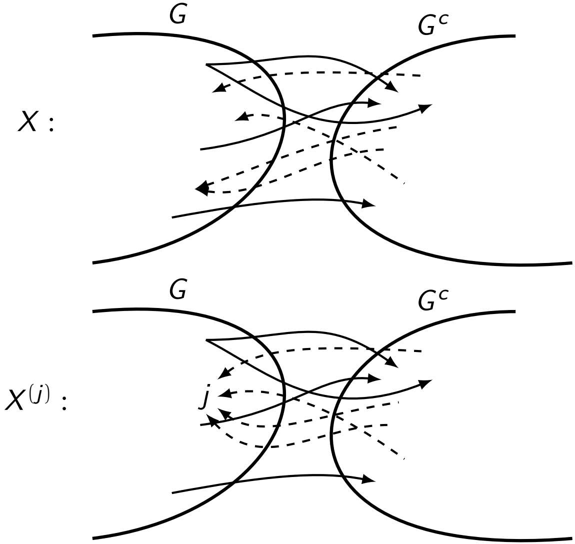

Key result

- That is, the state space of X is partitioned into (G, G^c).
- In $X^{(j)}$ we force all transitions from G^c to G to enter G by the same state j. The stationary distribution of $X^{(j)}$ ix $\pi^{(j)}$.
- $G, G^c, j \in G$ are arbitrary.

Key result

- That is, the state space of X is partitioned into (G, G^c).
- In $X^{(j)}$ we force all transitions from G^c to G to enter G by the same state j. The stationary distribution of $X^{(j)}$ ix $\pi^{(j)}$.
- $G, G^c, j \in G$ are arbitrary.

Starting result:

- There exists $\{\beta_j\}_{j\in G}$, $\beta_j \geq 0$, $\sum_{j\in G} \beta_j = 1$, s.t.

$$\pi = \sum_{j\in G} \beta_j \pi^{(j)}.$$

- Proved by Courtois & Semal. We improved the result slightly in previous joint work with S. Mahévas, then again recently when preparing this paper.
- In this presentation, we will give a more explicit statement.

Consequences

- Let $R^{(j)} = \mathbb{E}(r_{X_\infty^{(j)}}) = \sum_x r_x \pi_x^{(j)}$. Then,

$$R = \sum_{j \in G} \beta_j R^{(j)}.$$

- $\Longrightarrow \min_{j \in G} R^{(j)} \le R \le \max_{j \in G} R^{(j)}$
- Bounding technique (proposed by Muntz, De Souza e Silva and Goyal, see next slide): use *some computable* lower and upper bounds of previous bounds:

$$\left[\min_{j \in G} R^{(j)}\right]_{lb} \le R \le \left[\max_{j \in G} R^{(j)}\right]_{ub}.$$

Sources for the initial model

- Muntz, De Souza e Silva and Goyal, *Bounding Availability of Repairable Computer Systems*, IEEE Trans. Comp. 38 (12), 1989

  Nice use of the results of Courtois and Semal to obtain an operational procedure to compute tight bounds. Some assumptions are strong.

- Mahévas and Rubino, *Bound computation of dependability and performability measures*, IEEE Trans. Comp. 50 (5), 2001

  Some improvements to previous ideas, with much weaker applicability conditions, plus integrating the case of infinite state spaces, and opening the path to the new recents advances.

Pseudo-aggregation

- For instance, consider $G^c = C_1 + C_2$ (partition of G^c into two classes).
- The *pseudo-aggregation* of X (called "exact aggregation" by Courtois), denoted X' here, obtained by collapsing C_1 (resp. C_2) into state c_1 (resp. c_2), is a new Markov chain that
  - has state space $G + \{c_1, c_2\}$ (that is, $|G| + 2$ states),
  - with, for instance,

$$\text{rate}(c_1, c_2) = \frac{\sum_{x \in C_1} \pi_x \sum_{y \in C_2} Q_{xy}}{\sum_{x \in C_1} \pi_x},$$

  etc.
- Observe that X' is Markov *by construction*.
  Also observe that to build X' we need π.

Main property

- Let π' be the stationary distribution of X'.
- We have:

$$\pi'_{c_1} = \sum_{x \in C_1} \pi_x, \qquad \pi'_{c_2} = \sum_{x \in C_2} \pi_x,$$

and for all $x \in G$, $\quad \pi'_x = \pi_x$.

Who is G?

- Consider, for instance, that X **models a queueing network**.
- Let $\nu(x) = \#$ of customers in some subnetwork of the model when the state, in equilibrium, is x. We want to know, for instance, the mean number of customers in that subnetwork, in equilibrium, that is, $\mathbb{E}(\nu(X_\infty))$.

Who is G?

- Consider, for instance, that X **models a queueing network**.
- Let $\nu(x) = \#$ of customers in some subnetwork of the model when the state, in equilibrium, is x. We want to know, for instance, the mean number of customers in that subnetwork, in equilibrium, that is, $\mathbb{E}(\nu(X_\infty))$.
- Define $C_L = \{x : \nu(x) = L\}$.
- We set

$$G = \sum_{L<K} C_L, \qquad G^c = \sum_{L \geq K} C_L.$$

- Assume (only to simplify the presentation) that jumps are only possible inside the classes, or from C_L to C_{L-1} and to C_{L+1}.

- In a **dependability context, X can represent the configuration of some complex system having many components belonging to M different types, subject to failures and repairs**.
- A state could be a vector $x = \big(x_1, \ldots, x_M\big)$ where x_m is the number of failed components of type m.
- Transitions of X correspond to failures of some components, possibly repairs, to some configuration changes in the system, etc.
- We have a function Φ defined on the state space of X, where $\Phi(x) = 1(\text{ sys. is up when in state } x \text{ })$. The goal could be, for instance, to bound the system's asymptotic availability:

$$\mathbb{P}(\text{ sys. up in equilibrium }) = \mathbb{P}(\Phi(X) = 1) = \mathbb{E}\big(\Phi(X)\big).$$

- If $\nu(x)$ is the number of failed components when the state is x, a class of states could be $C_L = \{x \colon \nu(x) = L\}$.

Next step

- Once K chosen, we set $G = C_0 + C_1 + \cdots + C_{K-1}$ (we denote by '+' the union of disjoint sets).
- Let us denote by X' the pseudo-aggregation of X with respect to the partition $G + \{c_K, c_{K+1}, \ldots\}$, where class C_L is collapsed into a single state c_L.
- If we consider the same transformation as before, forcing the transitions in X' that go from c_K into G to enter by an arbitrary chosen state $j \in G$ (necessarily $j \in C_{K-1}$), it is easy to see that

$$X'^{(j)} \equiv X^{(j)\,\prime}.$$

Pseudo-aggregations

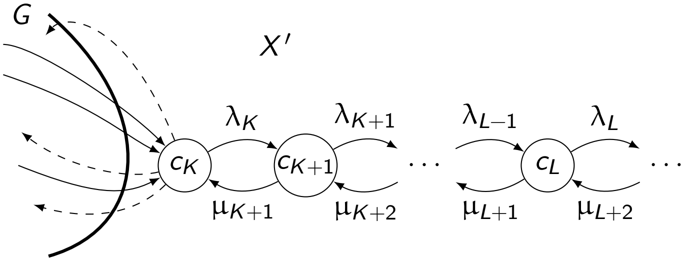

Pseudo-aggregations

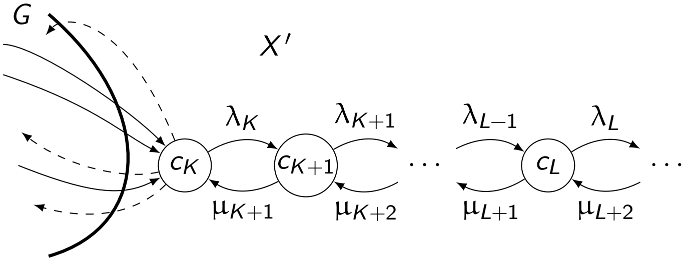

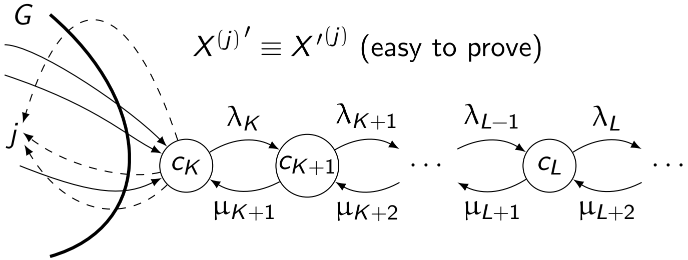

The last picture can be simplified as follows:

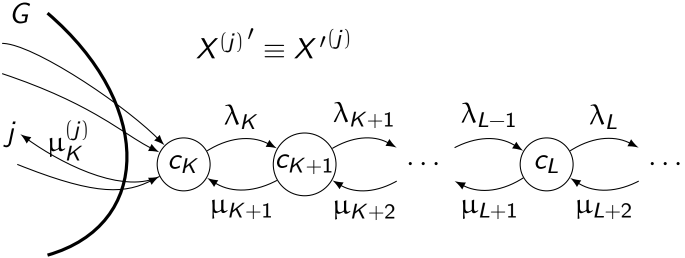

We have

$$\mu_K^{(j)} = \sum_{x \in G} \mathrm{rate}(c_K, x),$$

where "rate()" refers to the transitions in X'.

The last picture can be simplified as follows:

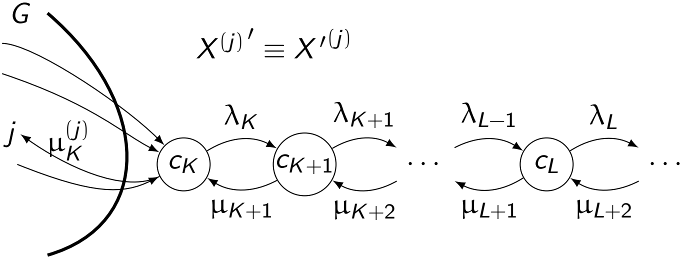

We have

$$\mu_K^{(j)} = \sum_{x \in G} \text{rate}(c_K, x),$$

where "rate()" refers to the transitions in X'.

Observe that the rates $\lambda_K, \lambda_{K+1}, \ldots, \mu_{K+1}, \ldots$ don't change if $j \in G$ changes.

Second main idea

- Assume first that two reals ρ_1 and ρ_2 are available, such that for all state x we have
$$0 \le \rho_1 \le r_x \le \rho_2 < \infty.$$
- Assume also that we can build upper and lower bounds of the exact (and unknown) rates λ_L and μ_L, $L \ge K$.
- This allows us to compute
$$\left[\min_{j\in G} R^{(j)}\right]_{lb} \quad \text{and} \quad \left[\max_{j\in G} R^{(j)}\right]_{ub}$$
in the following way:

Key results

- For each $j \in G$ define the Markov chain $Y^{(j)}$ by

$$Y^{(j)} = X'^{(j)}\Big|_{\lambda_L=[\lambda_L]_{ub}, \quad \mu_L=[\mu_L]_{lb}}$$

  with stationary distribution $y^{(j)}$.
- Let us denote $\pi^{(j)} = \left(\pi_G^{(j)} \;\; \pi_{G^c}^{(j)} \right)$ and $y^{(j)} = \left(y_G^{(j)} \;\; y_{G^c}^{(j)} \right)$.
- Now, on G,
  - $\tilde{\pi}_G^{(j)} = \dfrac{1}{\mathbb{P}(X_\infty^{(j)} \in G)} \pi_G^{(j)}$,
  - $\tilde{y}_G^{(j)} = \dfrac{1}{\mathbb{P}(Y_\infty^{(j)} \in G)} y_G^{(j)}$.
- Then, we have the two key following results:
  - $\tilde{y}_G^{(j)} = \tilde{\pi}_G^{(j)}$;
  - $\mathbb{P}(Y_\infty^{(j)} \in G) \le \mathbb{P}(X_\infty^{(j)} \in G)$.

First case

First case: finite models, or infinite models with bounded rewards.

- Using previous results and some algebra, we obtain

$$\rho_1 + \min_{j \in G} \sum_{x \in G} (r_x - \rho_1) y_x^{(j)} \leq R \leq \rho_2 - \max_{j \in G} \sum_{x \in G} (\rho_2 - r_x) y_x^{(j)}.$$

Second case

Second case: infinite models with unbounded rewards (from Mahévas and Rubino, *Bound computation of dependability and performability measures*, IEEE Trans. Comp. 50 (5), 2001).

- $[R]_{lb}$: basically as before
- for $[R]_{ub}$, let $X''^{(j)}$ denote the pseudo-aggregation of $X^{(j)}$ when G^c is entirely collapsed into a single state c.
- Let $Z^{(j)}$ be the pseudo-aggregation of $Y^{(j)}$ when the set of states $\{c_K, c_{K+1}, \ldots\}$ is collapsed into a single state c, with stationary distribution $z^{(j)}$.
- Let $\Theta_K = 1$ and for $I > K$,

$$\Theta_I = \prod_{L=K+1}^{I} \frac{[\lambda_{L-1}]_{ub}}{[\mu_L]_{lb}}.$$

- The only parameter of $Z^{(j)}$ that must be computed is $\nu_j = \mathrm{rate}(c, j)$; it is given by

$$\nu_j = \frac{[\mu_K]_{lb}}{\sum_{I \geq K} \Theta_I}.$$

- We need to use some $\overline{r}_L \geq r_i$, for all $i \in C_L$.
- Define

$$r_c = \max \left\{ \max_{i \in G} r_i \, , \, \frac{\sum_{L \geq K} \Theta_L \overline{r}_L}{\sum_{L \geq K} \Theta_L} \right\}$$

- Then, we have

$$R \leq \max_{j \in G} \left(\sum_{x \in G} r_x z_x^{(j)} + r_c z_c^{(j)} \right).$$

Outline

New results

- A critical performance issue of this approach is the number of $Y^{(j)}$ (or $Z^{(j)}$) chains to be solved.
- Let $M = -A_G^{-1}$ where A_G is the restriction of the infinitesimal generator of any of the considered chains to the states in G.
- Let s denote the mean sojourn time of X in G^c.
- Denote by $t_j = \mathbb{E}_j(\text{sojourn time of } X \text{ in } G)$.
  It is computable: $t_j = 1_j M 1^\mathsf{T}$, where 1_j is the jth vector of the canonical base (that is, $1_j M$ is the jth row of M).

Basically only one linear system

- The following result then holds:

$$y^{(j)} = \frac{1}{s+t_j}\left(1_j M, s, s\frac{\Theta_{K+1}}{D}, s\frac{\Theta_{K+2}}{D}, \cdots\right)$$

  where $D = \sum_{I \geq K} \Theta_I$.
- In the same way,

$$\pi''^{(j)} = \frac{1}{s+t_j}\left(1_j M, s\right).$$

(cont.)

To see how to apply this, consider, for instance, the bounds in the first case.

- To obtain $[R]_{lb}$ and $[R]_{ub}$, we must compute something of the form
$$\min_{j \in G} \sum_{i \in G} \gamma_j y_i^{(j)} \quad \text{and} \quad \max_{j \in G} \sum_{i \in G} \eta_j y_i^{(j)}.$$
- This basically reduces to compute $1_j M \gamma^{\mathsf{T}}$ and $1_j M 1^{\mathsf{T}}$.
  This means a fixed number of linear systems to solve, all involving the same matrix A_G.
  The exact scheme depends on the solving technique we use.

Second improvement

We can go a step further and prove that

$$\beta_j = \alpha_j \frac{s + t_j}{s + t}$$

where:

- $\alpha_j = \dfrac{\mu_{K,j}}{\sum_{h \in G} \mu_{K,h}}$,
- $\mu_{K,h} = \text{rate}(c_K, h)$ in X'
- and $t = \sum_{j \in G} \alpha_j t_j$.

$\rightsquigarrow$ new approach

- The starting point was $\min_{j\in G} R^{(j)} \le R \le \max_{j\in G} R^{(j)}$
- and then, to look for some $\left[\min_{j\in G} R^{(j)}\right]_{lb}$ and $\left[\max_{j\in G} R^{(j)}\right]_{ub}$.
- Now, we can use the obtained explicit expressions to be more precise.
- Let us consider, for instance, the case of $[R]_{ub}$ when $\rho_2 < \infty$.
- Using an upper bound ρ of the rewards on the set G^c and working on the state space $G + c$,

$$R^{(j)} \le \frac{1_j M r_G^{\mathrm{T}} + s\rho}{s + t_j}.$$

where r_G is the (row) vector of rewards inside G

(cont.)

- $\rightsquigarrow R \leq \sum_{j\in G} \frac{\alpha_j (1_j M r_G^{\mathrm{T}})}{s+t} + \frac{s}{s+t}\rho.$
- Looking at the sign of the derivatives of this upper bound w.r.t. the different parameters, we obtain that replacing the exact and unknown rates by upper of lower bounds, we obtain a new and computable bound.
- For instance,
$$\frac{\partial}{\partial s}\left(\frac{s}{s+t}\right) = \frac{t}{(s+t)^2} > 0,$$
so, using the previous auxiliary chains that "push the mass to the right" we obtain $s^* \geq s$ and this implies
$$\frac{s}{s+t} \leq \frac{s^*}{s^*+t}.$$

(cont.)

- We also can write

$$\frac{s^*}{s^* + t} \le \frac{s^*}{s^* + \min_{j \in G} t_j}$$

- A more complex expression can be derived in the same way for the second term of the bound.

Conclusions

- This method improves the efficiency (precision$\times$cost) of the previously published techniques.
- The key result is to make explicit some key quantities in the analysis, leading to a more precise bounding technique.
- Weakest point: the same as in our previous work, the computation of lower bounds of transition rates between lumped states (in a pseudo-aggregation).
- Future work: explore models where the mass is not necessarily "in the left of the space".