

Two Methods for Computing Bounds for the Distribution of Cumulative Reward for Large Markov Models

Juan A. Carrasco

Departament d'Enginyeria Electrònica

Universitat Politècnica de Catalunya

Diagonal 647, plta. 9

08028 Barcelona, Spain

Except for formatting details, this version matches exactly the version published with the same title and authors in *Performance Evaluation*, vol. 63, no. 12, 2006, pp. 1165–1195

Abstract

Degradable fault-tolerant systems can be evaluated using rewarded continuous-time Markov chain (CTMC) models. In that context, a useful measure to consider is the distribution of the cumulative reward over a time interval $[0, t]$. All currently available numerical methods for computing that measure tend to be very expensive when the product of the maximum output rate of the CTMC model and t is large and, in that case, their application is limited to CTMC models of moderate size. In this paper, we develop two methods to compute bounds for the cumulative reward distribution of CTMC models with reward rates associated with states: BT/RT (Bounding Transformation/Regenerative Transformation) and BT/BRT (Bounding Transformation/Bounding Regenerative Transformation). The methods require the selection of a regenerative state, are numerically stable and compute the bounds with well-controlled error. For a class of rewarded CTMC models, class C_1''' , and a particular, natural selection for the regenerative state the BT/BRT method allows to trade off bounds tightness with computational cost and will provide bounds at a moderate computational cost in many cases of interest. For a class of models, class C_1'' , slightly wider than class C_1''' , and a particular, natural selection for the regenerative state, the BT/RT method will yield tighter bounds at a higher computational cost. Under additional conditions, the bounds obtained by the less expensive version of BT/BRT and BT/RT seem to be tight for any value of t or not small values of t , depending on the initial probability distribution of the model. Class C_1'' and class C_1''' models with those additional conditions include both exact and bounding typical failure/repair performability models of fault-tolerant systems with exponential failure and repair time distributions and repair in every state with failed components and a reward rate structure which is a non-increasing function of the collection of failed components. We illustrate both the applicability and the performance of the methods using a large CTMC performability example of a fault-tolerant multiprocessor system.

Keywords: Fault-tolerant computer systems; Degradable systems; Continuous-time Markov chains; Distribution of cumulative reward; Bounds; Model transformation

1 Introduction

Fault-tolerant parallel and distributed systems typically exhibit a degradable performance in the presence of faults. Examples include multiprocessors, grid cluster computing systems and distributed systems built over local or wide area networks. A combined evaluation of the dependability and the performance of such systems can be performed by associating reward rates with the states of a “structure” continuous-time Markov chain (CTMC) model, capturing the failure/repair behavior of the fault-tolerant system, where the reward rate associated with a state is a measure of the performance rate of the system in that state. Those models have been called performability models because they capture the general performability concept introduced by Meyer [13, 14]. Rewarded CTMC models have become very popular and several well-known tools supporting their specification and solution have been developed [1, 4, 6, 7, 8, 27]. A recent, sophisticated performability modeling case-study of a fault-tolerant distributed system using the UltraSAN tool can be found in [15].

In the context of performability models, a useful measure to consider is the distribution of the cumulative reward over a time interval. With reward rates being performance rates, that measure is the distribution of the performance accumulated over a time interval. Formally, let $X = \{X(t); t \geq 0\}$ be a CTMC with state space Ω and let $r_i, i \in \Omega$ be the reward rate structure imposed over X . The quantity r_i has the meaning of rate at which reward is earned while X is in state i . The cumulative reward complementary distribution, $\text{CRCD}(t, s)$, is defined as the complementary probability distribution function of the reward earned in the time interval $[0, t]$, i.e.

$$\text{CRCD}(t, s) = P \left[\int_0^t r_{X(\tau)} d\tau > s \right].$$

The $\text{CRCD}(t, s)$ measure has interesting applications. Thus, X could model the failure/repair behavior of a multiprocessor and $r_i, i \in \Omega$, could be the speedup of the multiprocessor in state i . In that context, $\int_0^t r_{X(\tau)} d\tau$ would have the meaning of accumulated performance of the multiprocessor in the time interval $[0, t]$ (normalized with respect to the performance rate of a single processor) and $\text{CRCD}(t, s)$ would be the probability that that accumulated performance is above a given threshold s . Several numerical methods have been developed to compute the $\text{CRCD}(t, s)$ measure. Most of them assume Ω finite. Some of the methods allow impulse rewards associated with transitions which are earned every time a transition is followed. In our review, we will restrict our attention to general-purpose methods which, besides (possibly) Ω being finite, do not impose any restrictions on X . Smith et al. [28] developed a method with time complexity $O(M^3)$, where $M = |\Omega|$, which is based on the inversion of a double Laplace transform. Another method with time complexity $O(M^3)$ using Laguerre functions was developed by Islam and Ammar [11]. De Souza e Silva and Gail [29] developed a randomization-based method with time complexity exponential on the number of different reward rates. An extension of that method to cover models with both reward rates and impulse rewards using a path pruning technique was performed later by Qureshi and Sanders [21]. A method based on both randomization and Laplace transform techniques was developed by Donatiello and Grassi [9]. That method has time complexity $O(TmN^2)$, where T is the number of transitions of the randomized discrete time Markov chain (DTMC) of X with randomization

rate equal to Λ , the maximum output rate of X , $m + 1$ is the number of different reward rates of the model, and N is the truncation parameter of a Poisson process with rate Λ , which for large Λt has value approximately equal to Λt . Pattipati et al. [20] have formalized the computation of $\text{CRCD}(t, s)$ in terms of the solution of a linear hyperbolic partial differential equation which, by discretization, can be transformed into an ordinary differential equation (ODE), and have proposed to solve it using a standard stiff ODE solver. However, for large Λt , the number of discretization steps required to have good accuracy will be large and the method will be expensive. Nabli and Sericola [16, 17] developed another randomization-based method which for $s \geq r_{\text{fmax}} t$, r_{fmax} being the second largest reward rate of the model, has time complexity $O(T[C(N - C) + mC^2/2])$, where C is a truncation parameter which is small when s is close to $r_{\text{max}} t$, r_{max} being the largest reward rate of the model. The method is completely general and covers also the case $s < r_{\text{fmax}} t$, but with a time complexity $O(TmN^2)$. In addition, the method is guaranteed to be numerically stable. Queshi and Sanders [22] developed a method which is based on the enumeration of sample paths and is also numerically stable. Later, de Souza e Silva and Gail [30] developed a randomization-based method with time complexity $O(T\theta N^2)$ for the case in which only reward rates are present, where θ is a positive integer no greater than m . Finally, Rácz et al. [23] have developed a method to compute bounds for the $\text{CRCD}(t, s)$ measure based on the computation, using randomization [31], of the first moments of the cumulative reward. The computational cost of the method is basically the cost of computing the moments, which is $O(TqN')$, where q is the number of computed moments and, for large Λt , N' has value approximately equal to Λt . The bounds are very coarse for s around the expected cumulative reward but get tighter as s gets apart from that value. All currently available general-purpose methods have high computational cost when Λt is large and, in that case, their applicability is limited to models of moderate size.

In this paper, we develop two methods for computing bounds for the $\text{CRCD}(t, s)$ measure. The main advantage of the methods is that, for certain classes of rewarded CTMC models of interest, they can be relatively inexpensive and yet provide quite tight bounds, allowing a numerical analysis with error bounds in reasonable CPU times of very large models which were previously out of the scope of numerical techniques. The methods are numerically stable and compute the bounds with well-controlled error. The methods build upon recently developed methods for computing both the interval availability distribution [2] and bounds for that measure [3]. The rest of the paper is organized as follows. Section 2 describes the models which can be handled by the methods and defines the model classes C_1'' and C_1''' to which the methods are primarily intended to be applied. The methods are developed in Section 3. That section also discusses the computational costs of the methods for models in classes C_1'' and C_1''' . Section 4 illustrates the performance of the methods using a large performability example of a fault-tolerant multiprocessor for which bounding models belong to classes C_1'' and C_1''' and identifies under which additional conditions the bounds obtained by the methods seem to be tight. Model classes C_1'' and C_1''' with those additional conditions include both exact and bounding typical failure/repair performability models of fault-tolerant systems with exponential failure and repair time distributions and repair in every state with failed components and a reward rate structure which is a non-increasing function of the collection of failed components. Finally, Section 5 presents the conclusions. The Appendix includes a technical lemma which can be

used to circumvent some of the conditions imposed by the methods to the rewarded CTMC model and the (long) proof of the theoretical result on which the methods are based.

2 Classes of Models

Let $X = \{X(t); t \geq 0\}$ be a CTMC with state space Ω and reward rate structure $r_i, i \in \Omega$. Assuming those quantities defined, let $r_{\max} = \max_{i \in \Omega} r_i$, $r_{\text{fmax}} = \max_{i \in \Omega: r_i < r_{\max}} r_i$ and $r_{\min} = \min_{i \in \Omega} r_i$. In this paper we target the computation of bounds for the cumulative reward complementary distribution measure

$$\text{CRCDD}(t, s) = P \left[\int_0^t r_{X(\tau)} d\tau > s \right],$$


where $t > 0$ and $r_{\min} t < s < r_{\max} t$. In addition, we will assume $s > r_{\text{fmax}} t$ whenever an upper bound for $\text{CRCDD}(t, s)$ has to be computed.

Let $\alpha_i = P[X(0) = i], i \in \Omega$, let $\lambda_{i,j}, i, j \in \Omega, j \neq i$ denote the transition rates of X and let $\lambda_i = \sum_{j \in \Omega - \{i\}} \lambda_{i,j}, i \in \Omega$ denote the rates of output of X . Also, given $B \subset \Omega$, let $\alpha_B = \sum_{i \in B} \alpha_i$ and, given $i \in \Omega$ and $B \subset \Omega - \{i\}$, let $\lambda_{i,B} = \sum_{j \in B} \lambda_{i,j}$. The BT/RT (*Bounding Transformation/Regenerative Transformation*) method will require the selection of a regenerative state r . Letting $\Omega_{\max} = \{i \in \Omega : r_i = r_{\max}\}$, $\Omega_{\text{fmax}} = \{i \in \Omega : r_i = r_{\text{fmax}}\}$, $\Omega_{\min} = \{i \in \Omega : r_i = r_{\min}\}$, $\bar{\Omega} = \Omega - \Omega_{\max} - \Omega_{\text{fmax}} - \Omega_{\min}$, $S_{\max} = S \cap \Omega_{\max}$, $S_{\text{fmax}} = S \cap \Omega_{\text{fmax}}$, $S_{\min} = S \cap \Omega_{\min}$, $\bar{S} = S \cap \bar{\Omega}$, $S'_{\max} = S_{\max} - \{r\}$, $S'_{\text{fmax}} = S_{\text{fmax}} - \{r\}$, $S'_{\min} = S_{\min} - \{r\}$, and $\bar{S}' = \bar{S} - \{r\}$, the method will cover rewarded CTMCs X and selections for r satisfying the following conditions:

- C1. Ω is finite.
- C2. The reward rates $r_i, i \in \Omega$ are all finite and take at least three different values.
- C3. $\Omega = S$ or $\Omega = S \cup \{f\}$, where f is an absorbing state.
- C4. $|S| \geq 2$.
- C5. Either all states in S are transient or X has a single recurrent class of states $C \subset S$.
- C6. All states are reachable (from some state with nonnull initial probability).
- C7. $\max_{i \in \Omega_{\max}} \lambda_i > 0$ and $\max_{i \in \Omega_{\text{fmax}} \cup \bar{\Omega} \cup \Omega_{\min}} \lambda_i > 0$.
- C8. $r \in S$ and, if X has a single recurrent class of states $C \subset S, r \in C$.
- C9. If $S'_{\max} \neq \emptyset, \lambda_{r, S'_{\max}} > 0$.
- C10. If $S'_{\max} \neq \emptyset, \alpha_{S'_{\text{fmax}} \cup \bar{S}' \cup S'_{\min}} > 0$ and $\alpha_{S'_{\max}} = 0, \lambda_{i, S'_{\max}} > 0$ for some $i \in S'_{\text{fmax}} \cup \bar{S}' \cup S'_{\min}$ with $\alpha_i > 0$.

Given a selection for the regenerative state r , there exists a unique subset of states S for which conditions C3, C5 and C8 can all be satisfied: S must be Ω if X has no absorbing state or X has a single absorbing state a and $r = a$; S must be $\Omega - \{a\}$ if X has a single absorbing state a and $r \neq a$ or X has two absorbing states a, b and $b = r$; and, in any other case, no S exists for which conditions C3, C5 and C8 can all be satisfied. This makes it easy to check whether a given rewarded CTMC X with a given selection for the regenerative state r is covered by BT/RT. Conditions C2 and C7, except for the fact that reward rates be finite, are mild, in the sense that, when these conditions are not satisfied, computation of $\text{CRCD}(t, s)$ or of bounds for $\text{CRCD}(t, s)$ can be reduced to simpler problems. Thus, when the reward rates of X are finite but take only two different values, r_{\max} and r_{\min} , $\text{CRCD}(t, s)$ can be formulated in terms of the simpler interval availability complementary distribution measure, $\text{IAVCD}(t, p) = P[(1/t) \int_0^t I_{X(\tau) \in U} d\tau > p]$ (I_c denotes the indicator function returning value 1 if condition c is satisfied and value 0 otherwise), with subset of up states $U = \Omega_{\max}$ using $\text{CRCD}(t, s) = \text{IAVCD}(t, (s/t - r_{\min})/(r_{\max} - r_{\min}))$ (see Proof of Theorem 2). When condition C2 is satisfied but $\max_{i \in \Omega_{\max}} \lambda_i = 0$ the rewarded CTMCs X^{lb} and X^{ub} to be defined in Section 3 will be such that all their states with reward rate equal to r_{\max} will be absorbing and the remaining states will have reward rate equal to, respectively, r_{\min} and r_{fmax} , and, according to the previous discussion and the discussion performed in [2] concerning the $\text{IAVCD}(t, p)$ measure when $\max_{i \in U} \lambda_i = 0$, lower and upper bounds for $\text{CRCD}(t, s)$ can be computed as $P[X^{\text{lb}}((1 - (s/t - r_{\min})/(r_{\max} - r_{\min}))t) \in \Omega_{\max}]$ and $P[X^{\text{ub}}((1 - (s/t - r_{\text{fmax}})/(r_{\max} - r_{\text{fmax}}))t) \in \Omega_{\max}]$. Similarly, assuming C2 satisfied but $\max_{i \in \Omega_{\text{fmax}} \cup \bar{\Omega} \cup \Omega_{\min}} \lambda_i = 0$, the rewarded CTMC X^{lb} and X^{ub} to be defined in Section 3 will be such that all states with reward rate equal to, respectively, r_{\min} and r_{fmax} will be absorbing and the remaining states will have reward rate equal to r_{\max} , and, according to the previous discussion and the discussion performed in [2] concerning the $\text{IAVCD}(t, p)$ measure when $\max_{i \in D} \lambda_i = 0$, lower and upper bounds for $\text{CRCD}(t, s)$ can be computed as $P[X^{\text{lb}}(((s/t - r_{\min})/(r_{\max} - r_{\min}))t) \in \Omega_{\max}]$ and $P[X^{\text{ub}}(((s/t - r_{\text{fmax}})/(r_{\max} - r_{\text{fmax}}))t) \in \Omega_{\max}]$. Condition C6 can be trivialized by deleting non-reachable states. Finally, conditions C9 and C10 can be circumvented by adding to X a tiny transition rate $\lambda \leq 10^{-10} \varepsilon / (2t_{\max})$, where ε is the allowed error and t_{\max} is the largest time t at which bounds for $\text{CRCD}(t, s)$ have to be computed, with, according to Lemma 1 in the Appendix, a negligible impact on $\text{CRCD}(t, s)$, $0 < t \leq t_{\max}$ no greater than $10^{-10} \varepsilon$. The possibility $\Omega = S \cup \{f\}$, where f is an absorbing state, allows us to cover bounding models, which are useful for systems for which an exact model would have a state space of unmanageable size. In a bounding model, S would include a strict subset of the state space of the exact model and the bounding model would enter state f when the exact model would exit S . Assigning to the states in S the same reward rates as in the exact model and to state f a lower bound for all reward rates of the exact model, the $\text{CRCD}(t, s)$ measure of the bounding model would lower bound the $\text{CRCD}(t, s)$ measure of the exact model; if a reward rate upper bounding all reward rates of the exact model is assigned instead to state f , the $\text{CRCD}(t, s)$ measure of the bounding model would upper bound the $\text{CRCD}(t, s)$ measure of the exact model.

The BT/BRT (*Bounding Transformation/Bounding Regenerative Transformation*) method also requires the selection of a regenerative state r and covers a subset of the rewarded CTMC mod-



$$r_1 = r_2 = 4 \text{ h}^{-1}, r_3 = 3 \text{ h}^{-1}, r_4 = 2 \text{ h}^{-1}, r_5 = 1 \text{ h}^{-1},$$

$$r_6 = r_{1'} = r_{2'} = r_{3'} = r_{4'} = 0$$

Figure 1: Small rewarded CTMC X which with the selection $r = 1$ is covered by BT/RT and BT/BRT.

els covered by BT/RT, which can be described as the set including all rewarded CTMCs X and selections for r satisfying conditions C1–C10 and

$$\text{C11. } S'_{\max} \neq \emptyset.$$

Note that conditions C9 and C11 imply that r cannot be absorbing and, then, according to the discussion regarding the possibilities for S in BT/RT, in BT/BRT S must include precisely the non-absorbing states.

To illustrate the classes of models covered by the methods, Fig. 1 shows a small rewarded CTMC model of a repairable multiprocessor system with five processors, one of which is spare, in which active processors fail with rate $\lambda = 10^{-4} \text{ h}^{-1}$ and failed processors are repaired by a single repairer at rate $\mu = 0.1 \text{ h}^{-1}$. A processor fault is covered with probability $C = 0.98$. The initial state is the state 1 in which no processor is failed. Since the model has no absorbing state, $S = \Omega$. With the selection $r = 1$, that model is covered by both BT/RT and BT/BRT.

The BT/RT method is primarily intended to be used for a class of rewarded CTMC models, class C''_1 , with a given “natural” selection for the regenerative state. Model class C''_1 includes all rewarded CTMCs X satisfying conditions C1–C7 and the condition:

$$\text{C12. A partition } S_0 \cup S_1 \cup \dots \cup S_{N_C} \text{ for } S_{\max} \text{ exists satisfying the properties:}$$

- P1. $S_0 = \{o\}$ (i.e. $|S_0| = 1$).
- P2. If X has a single recurrent class of states $C \subset S$, $o \in C$.
- P3. If $|S_{\max}| \geq 2$, $\lambda_{o, S_1 \cup \dots \cup S_{N_C}} > 0$.
- P4. If $|S_{\max}| \geq 2$, $\alpha_{S_{\max} \cup \bar{S} \cup S_{\min}} > 0$ and $\alpha_{S_1 \cup \dots \cup S_{N_C}} = 0$, $\lambda_{i, S_1 \cup \dots \cup S_{N_C}} > 0$ for some $i \in S_{\max} \cup \bar{S} \cup S_{\min}$ with $\alpha_i > 0$.
- P5. If $|S_{\max}| \geq 2$, $\max_{0 \leq k \leq N_C} \max_{i \in S_k} \lambda_{i, S_k - \{i\} \cup S_{k+1} \cup \dots \cup S_{N_C} \cup S_{\max} \cup \bar{S} \cup S_{\min}}$ is significantly smaller than $\min_{0 < k \leq N_C} \min_{i \in S_k} \lambda_{i, S_0 \cup \dots \cup S_{k-1}} > 0$ if $\Omega = S$ or $\min_{0 < k \leq N_C} \min_{i \in S_k} \lambda_{i, S_0 \cup \dots \cup S_{k-1} \cup \{f\}} > 0$ if $\Omega = S \cup \{f\}$.

The natural selection for the regenerative state for class C_1'' models is $r = o$. With that natural selection, properties P2, P3 and P4 imply the fulfillment of, respectively, conditions C8, C9 and C10 and, then, the models will be covered by BT/RT.

The BT/BRT method is primarily intended to be used for a class of rewarded CTMC models, class C_1''' , with a given “natural” selection for the regenerative state. Model class C_1''' includes all rewarded CTMCs X satisfying conditions C1–C7 and the conditions:

- C13. $|S_{\max}| \geq 2$.
- C14. A partition $S_0 \cup S_1 \cup \dots \cup S_{N_C}$ for S_{\max} exists satisfying the properties:
 - P1. $S_0 = \{o\}$ (i.e. $|S_0| = 1$).
 - P2. If X has a single recurrent class of states $C \subset S$, $o \in C$.
 - P6. $\lambda_{o, S_1 \cup \dots \cup S_{N_C}} > 0$.
 - P7. If $\alpha_{S_{\max} \cup \bar{S} \cup S_{\min}} > 0$ and $\alpha_{S_1 \cup \dots \cup S_{N_C}} = 0$, $\lambda_{i, S_1 \cup \dots \cup S_{N_C}} > 0$ for some $i \in S_{\max} \cup \bar{S} \cup S_{\min}$ with $\alpha_i > 0$.
 - P8. $\max_{0 \leq k \leq N_C} \max_{i \in S_k} \lambda_{i, S_k - \{i\} \cup S_{k+1} \cup \dots \cup S_{N_C} \cup S_{\max} \cup \bar{S} \cup S_{\min}}$ is significantly smaller than $\min_{0 < k \leq N_C} \min_{i \in S_k} \lambda_{i, S_0 \cup \dots \cup S_{k-1}} > 0$ if $\Omega = S$ or $\min_{0 < k \leq N_C} \min_{i \in S_k} \lambda_{i, S_0 \cup \dots \cup S_{k-1} \cup \{f\}} > 0$ if $\Omega = S \cup \{f\}$.
 - P9. $\lambda_o \leq \min_{i \in S_1 \cup \dots \cup S_{N_C}} \lambda_i$.

The natural selection for the regenerative state for class C_1''' models is $r = o$. With that natural selection, condition C13 and properties P2, P6 and P7 imply the fulfillment of, respectively, conditions C11, C8, C9, and C10 and, then, the models will be covered by BT/BRT. Class C_1''' is a subset of class C_1'' .

The rewarded CTMC model described in Fig. 1 belongs to model classes C_1'' and C_1''' . A partition for $S_{\max} = \{1, 2\}$ showing that is $S_{\max} = S_0 \cup S_1$ with $S_0 = \{1\}$ and $S_1 = \{2\}$.

3 The Methods

The BT/RT and BT/BRT methods are based on the following result.

Theorem 1. *Let $W = \{W(t); t \geq 0\}$ be a conservative, uniformizable CTMC with denumerable state space Ω ,¹ uniformly upper bounded reward rate structure r_i , $i \in \Omega$ with $|r_i| < \infty$, $i \in \Omega$ and transition rates $\lambda_{i,j}$, $i, j \in \Omega$, $j \neq i$. Let $r_{\text{sup}} = \sup_{i \in \Omega} r_i$, let $x \in \Omega$ with $r_x < r_{\text{sup}}$ and let r'_x such that $r_x \leq r'_x < r_{\text{sup}}$. Let $W' = \{W'(t); t \geq 0\}$ be another rewarded conservative, uniformizable CTMC identical to W except that state x has reward rate r'_x and the rates of transition from state x have values $\lambda'_{x,j} = \beta \lambda_{x,j}$, $j \in \Omega$, $j \neq x$, $\beta = (r_{\text{sup}} - r'_x)/(r_{\text{sup}} - r_x)$. Let $\text{CRCD}(t, s) = P[\int_0^t r_{W(\tau)} d\tau > s]$ be the cumulative reward complementary distribution of X in the time interval $[0, t]$, $t > 0$ and let $\text{CRCD}'(t, s) = P[\int_0^t r_{W'(\tau)} d\tau > s]$ be the cumulative reward complementary distribution of W' in the time interval $[0, t]$, $t > 0$. Then $\text{CRCD}'(t, s) \geq \text{CRCD}(t, s)$.*

Proof. See the Appendix. □

Essentially, the reasons why the result holds are that: (1) the scaling of rates of transition from state x keeping their relative values will not modify the embedded discrete time Markov chain DTMC Π of W , (2) the scaling factor β is chosen in such a way that the reward lost by W' in each visit to state x with respect to the reward that would be earned had state x reward rate r_{sup} has an exponential distribution with parameter $\beta \lambda_x / (r_{\text{sup}} - r'_x)$, identical to the parameter, $\lambda_x / (r_{\text{sup}} - r_x)$, of the exponential distribution of the reward lost by W in each visit to state x . Then, since (see, for instance, [12]) both W and W' can be interpreted in terms of Π by associating with the states visited by Π independent exponential holding times with parameter equal to the rate of output from the visited state, the rates of output from and the rates of reward of the states different from x are equal in W and W' , and, with $0 < \beta \leq 1$, the rate of output from state x in W' is non-greater than the rate of output from state x in W , each realization of W will have a corresponding realization of W' with the same sequence of visited states, same holding times in states different from x , greater holding times in x , and same reward lost in each state visit with respect to the reward that would be earned had the state a reward rate r_{sup} . As Fig. 2 illustrates, this will make the “loss” of reward accumulated by W' in the time interval $[0, t]$ not greater than the “loss” of reward accumulated by W in the same time interval and, therefore, the reward accumulated by W' in the time interval $[0, t]$ will be not smaller than the reward accumulated by W in the same time interval. As the reward accumulated in the time interval $[0, t]$ by W' for a realization of W' is not smaller than the reward accumulated in the time interval $[0, t]$ by W for the corresponding realization of W , the probability that W' will accumulate a reward in the time interval $[0, t]$ greater than s will be not smaller than the probability that W will accumulate a reward in the time interval $[0, t]$ greater than s . Note that Theorem 1 also holds if x is an absorbing state. In that case, both W and W' will remain indefinitely in x once they enter that state and the result simply holds because $r'_x \geq r_x$.

¹See, for instance, [12] for the definitions of conservative, uniformizable CTMCs with denumerable state space. Basically, they are CTMCs with denumerable state space in which the output rate from any state i is equal to the sum of the transition rates from i and in which the output rates are uniformly upper bounded.




Figure 2: Comparison of corresponding realizations of W and W' satisfying the conditions of Theorem 1 regarding the loss of accumulated reward in the time interval $[0, t]$ with respect to the reward that W and W' would earn had all the states reward rate r_{sup} .

The BT/RT method allows us to compute a lower bound for $\text{CRCD}(t, s)$, an upper bound for $\text{CRCD}(t, s)$, or both. The bounds are computed with error upper bounded by ε , where ε is a user-provided error control parameter. The method combines a model transformation phase (*Bounding Transformation*) with the RT (*Regenerative Transformation*) method described in [2]. When a lower bound for $\text{CRCD}(t, s)$ has to be computed, the model transformation is performed so that the resulting rewarded CTMC, X^{lb} , has only two different reward rates and, according to Theorem 1, has a cumulative reward complementary distribution measure, $\text{CRCD}^{\text{lb}}(t, s)$, which lower bounds $\text{CRCD}(t, s)$. When an upper bound for $\text{CRCD}(t, s)$ has to be computed, the model transformation is performed so that the resulting rewarded CTMC, X^{ub} , has only two different reward rates and, according to Theorem 1, has a cumulative reward complementary distribution measure, $\text{CRCD}^{\text{ub}}(t, s)$, which upper bounds $\text{CRCD}(t, s)$. With X^{lb} and X^{ub} having two different reward rates, the computation of $\text{CRCD}^{\text{lb}}(t, s)$ and $\text{CRCD}^{\text{ub}}(t, s)$ can be mapped into the computation of the simpler interval availability complementary distribution measures of, respectively, X^{lb} and X^{ub} , with subset of up states Ω_{max} , and those computations are performed using RT with regenerative state r and allowed error ε .

The lower bounding rewarded CTMC model, X^{lb} , has the same state space Ω as X , same initial probability distribution as X , reward rates r_{max} associated with the states in Ω_{max} , reward rates r_{min} associated with the states in $\Omega_{\text{fmax}} \cup \bar{\Omega} \cup \Omega_{\text{min}}$, and transition rates:

$$\lambda_{i,j}^{\text{lb}} = \frac{r_{\text{max}} - r_{\text{min}}}{r_{\text{max}} - r_i} \lambda_{i,j}, \quad i \in \Omega_{\text{fmax}} \cup \bar{\Omega},$$

$$\lambda_{i,j}^{\text{lb}} = \lambda_{i,j}, \quad i \in \Omega_{\text{max}} \cup \Omega_{\text{min}}.$$

The upper bounding rewarded CTMC model, X^{ub} , has same state space Ω as X , reward rates r_{max} associated with the states in Ω_{max} , reward rates r_{fmax} associated with the states in $\Omega_{\text{fmax}} \cup \bar{\Omega} \cup \Omega_{\text{min}}$, and transition rates:

$$\lambda_{i,j}^{\text{ub}} = \frac{r_{\text{max}} - r_{\text{fmax}}}{r_{\text{max}} - r_i} \lambda_{i,j}, \quad i \in \bar{\Omega} \cup \Omega_{\text{min}}.$$

$$\lambda_{i,j}^{\text{ub}} = \lambda_{i,j}, \quad i \in \Omega_{\text{max}} \cup \Omega_{\text{fmax}}.$$

That $\text{CRCD}^{\text{ub}}(t, s)$ upper bounds $\text{CRCD}(t, s)$ follows by considering, for each $i \in \bar{\Omega} \cup \Omega_{\text{min}}$ in turn, the scaling by the factor $(r_{\text{max}} - r_{\text{fmax}})/(r_{\text{max}} - r_i)$ yielding the transition rates from state i in X^{ub} from the transition rates from state i in X and applying Theorem 1 to each scaling step, noting that $r_i < r_{\text{fmax}} < r_{\text{max}}$, $i \in \bar{\Omega} \cup \Omega_{\text{min}}$. That $\text{CRCD}(t, s)$ upper bounds $\text{CRCD}^{\text{lb}}(t, s)$ and, therefore, $\text{CRCD}^{\text{lb}}(t, s)$ lower bounds $\text{CRCD}(t, s)$ follows by considering, for each $i \in \Omega_{\text{fmax}} \cup \bar{\Omega}$ in turn, the scaling by the factor $(r_{\text{max}} - r_i)/(r_{\text{max}} - r_{\text{min}})$ yielding the transition rates from state i in X from the transition rates from state i in X^{lb} and applying Theorem 1 to each scaling step, noting that $r_{\text{min}} < r_i < r_{\text{max}}$, $i \in \Omega_{\text{fmax}} \cup \bar{\Omega}$.

The mappings of $\text{CRCD}^{\text{lb}}(t, s)$ and $\text{CRCD}^{\text{ub}}(t, s)$ into the interval availability complementary distributions of, respectively, X^{lb} and X^{ub} with subset of up states Ω_{max} are given by the following theorem.

Theorem 2. Let $\text{IAVCD}^{\text{lb}}(t, p) = P[(1/t) \int_0^t I_{X^{\text{lb}}(\tau) \in \Omega_{\text{max}}} d\tau > p]$ and $\text{IAVCD}^{\text{ub}}(t, p) = P[(1/t) \int_0^t I_{X^{\text{ub}}(\tau) \in \Omega_{\text{max}}} d\tau > p]$. Then, for $t > 0$,

$$\text{CRCD}^{\text{lb}}(t, s) = \text{IAVCD}^{\text{lb}}\left(t, \frac{s/t - r_{\text{min}}}{r_{\text{max}} - r_{\text{min}}}\right)$$

and

$$\text{CRCD}^{\text{ub}}(t, s) = \text{IAVCD}^{\text{ub}}\left(t, \frac{s/t - r_{\text{fmax}}}{r_{\text{max}} - r_{\text{fmax}}}\right).$$

Proof. Denoting by $r_i^{\text{lb}}, i \in \Omega$ the reward rate structure of X^{lb} , since $r_i^{\text{lb}} = r_{\text{max}}$ for $i \in \Omega_{\text{max}}$ and $r_i^{\text{lb}} = r_{\text{min}}$ for $i \in \Omega_{\text{fmax}} \cup \bar{\Omega} \cup \Omega_{\text{min}}$,

$$\begin{aligned} \text{CRCD}^{\text{lb}}(t, s) &= P\left[\int_0^t r_{X^{\text{lb}}(\tau)}^{\text{lb}} d\tau > s\right] \\ &= P\left[r_{\text{max}} \int_0^t I_{X^{\text{lb}}(\tau) \in \Omega_{\text{max}}} d\tau + r_{\text{min}} \int_0^t I_{X^{\text{lb}}(\tau) \in \Omega_{\text{fmax}} \cup \bar{\Omega} \cup \Omega_{\text{min}}} d\tau > s\right] \\ &= P\left[r_{\text{min}} t + (r_{\text{max}} - r_{\text{min}}) \int_0^t I_{X^{\text{lb}}(\tau) \in \Omega_{\text{max}}} d\tau > s\right] \\ &= P\left[\int_0^t I_{X^{\text{lb}}(\tau) \in \Omega_{\text{max}}} d\tau > \frac{s - r_{\text{min}} t}{r_{\text{max}} - r_{\text{min}}}\right] = P\left[\frac{1}{t} \int_0^t I_{X^{\text{lb}}(\tau) \in \Omega_{\text{max}}} d\tau > \frac{s/t - r_{\text{min}}}{r_{\text{max}} - r_{\text{min}}}\right] \\ &= \text{IAVCD}^{\text{lb}}\left(t, \frac{s/t - r_{\text{min}}}{r_{\text{max}} - r_{\text{min}}}\right). \end{aligned}$$

The result for $\text{CRCD}^{\text{ub}}(t, s)$ can be obtained similarly, the only difference being that X^{lb} has to be replaced by $X^{\text{ub}}, r_i^{\text{lb}}$ by $r_i^{\text{ub}}, r_i^{\text{ub}}$ being the reward rate of X^{ub} in state i , and r_{min} by r_{fmax} , because the reward rate of X^{ub} in the states $i \in \Omega_{\text{fmax}} \cup \bar{\Omega} \cup \Omega_{\text{min}}$ is $r_i^{\text{ub}} = r_{\text{fmax}}$. \square

The method BT/BRT differs from BT/RT in that, instead of computing $\text{CRCD}^{\text{lb}}(t, s) = \text{IAVCD}^{\text{lb}}(t, (s/t - r_{\text{min}})/(r_{\text{max}} - r_{\text{min}}))$ with error $\leq \varepsilon$ using the RT method with regenerative state r , a lower bound for it, $\text{CRCD}^{\text{lb,lb}}(t, s)$, is obtained by computing a lower bound for $\text{IAVCD}^{\text{lb}}(t, (s/t - r_{\text{min}})/(r_{\text{max}} - r_{\text{min}}))$ using the BRT (*Bounding Regenerative Transformation*) method [3] with regenerative state r and allowed error ε . Similarly, instead of computing $\text{CRCD}^{\text{ub}}(t, s) = \text{IAVCD}^{\text{ub}}(t, (s/t - r_{\text{fmax}})/(r_{\text{max}} - r_{\text{fmax}}))$ with error $\leq \varepsilon$ using RT with regenerative state r , an upper bound for it, $\text{CRCD}^{\text{ub,ub}}(t, s)$, is obtained by computing an upper bound for $\text{IAVCD}^{\text{ub}}(t, (s/t - r_{\text{fmax}})/(r_{\text{max}} - r_{\text{fmax}}))$ using BRT with regenerative state r and allowed error ε . Thus, BT/BRT will yield less tighter bounds than BT/RT. The BRT method [3] has a control parameter D_C which allows to trade-off bounds tightness with computational cost. The BT/BRT method has also a control parameter D_C and the BRT method is invoked with its D_C control parameter equal to the D_C control parameter of BT/BRT. The D_C control parameter has to be selected [3] so that it satisfies $1 \leq D_C < \lambda_{\text{max}}/\lambda_{\text{min}}$, $\lambda_{\text{max}} = \max_{i \in S'_{\text{max}}} \lambda_i$, $\lambda_{\text{min}} = \min_{i \in S'_{\text{max}}} \lambda_i$. Note that, as we have discussed, conditions C9 and C11 imply that S must include precisely the non-absorbing states and, therefore, $\lambda_{\text{max}} \geq \lambda_{\text{min}} > 0$. In the possible case $\lambda_{\text{max}} = \lambda_{\text{min}}$ no selection for D_C would be possible. In that case, BT/RT should be used. Since [3], as D_C increases, the bounds for the interval availability complementary distribution obtained by BRT get arbitrarily tighter; as

D_C increases, the bounds obtained by BT/BRT approach those obtained by BT/RT. Furthermore, as we shall discuss, for class C_1''' models with the natural selection $r = o$, the computational cost of BT/BRT should increase with D_C and, therefore, the D_C parameter allows to trade-off bounds tightness with computational cost.

The BRT method [3] can be described as the succession of a model transformation phase, in which transition rates from “up” states different from the regenerative state and, if existent, the absorbing state, are scaled, and the RT method with the same regenerative state as BRT, which is used to solve the transformed CTMC model. When, in BT/BRT, the lower bound for CRCD(t, s) has to be computed, the CTMC model, $X^{\text{lb,lb}}$, which has to be solved by the RT method has [3] same state space Ω as X^{lb} (and X), same initial probability distribution as X^{lb} (and X), subset of up states Ω_{max} and, calling $\lambda_i^{\text{lb}}, i \in \Omega$, the output rates of X^{lb} and noting that $\lambda_i^{\text{lb}} = \lambda_i, i \in \Omega_{\text{max}}$ and, therefore, $\max_{i \in S'_{\text{max}}} \lambda_i^{\text{lb}} = \lambda_{\text{max}}$, transition rates related to the transition rates of X^{lb} as

$$\lambda_{i,j}^{\text{lb,lb}} = \frac{\max\{\lambda_i^{\text{lb}}, \lambda_{\text{max}}/D_C\}}{\lambda_i^{\text{lb}}} \lambda_{i,j}^{\text{lb}}, \quad i \in S'_{\text{max}},$$

$$\lambda_{i,j}^{\text{lb,lb}} = \lambda_{i,j}^{\text{lb}}, \quad i \notin S'_{\text{max}}.$$

When, in BT/BRT, the upper bound for CRCD(t, s) has to be computed, the CTMC model, $X^{\text{ub,ub}}$, which has to be solved by the RT method has [3] same state space Ω as X^{ub} (and X), same initial probability distribution as X^{ub} (and X), subset of up states Ω_{max} and, calling $\lambda_i^{\text{ub}}, i \in \Omega$, the output rates of X^{ub} and noting that $\lambda_i^{\text{ub}} = \lambda_i, i \in \Omega_{\text{max}}$ and, therefore, $\min_{i \in S'_{\text{max}}} \lambda_i^{\text{ub}} = \lambda_{\text{min}}$, transition rates related to the transition rates of X^{ub} as

$$\lambda_{i,j}^{\text{ub,ub}} = \frac{\min\{\lambda_i^{\text{ub}}, D_C \lambda_{\text{min}}\}}{\lambda_i^{\text{ub}}} \lambda_{i,j}^{\text{ub}}, \quad i \in S'_{\text{max}},$$

$$\lambda_{i,j}^{\text{ub,ub}} = \lambda_{i,j}^{\text{ub}}, \quad i \notin S'_{\text{max}}.$$

We also point out that RT uses Algorithm A of [25] as a back-end. That back-end is preceded by a transformation phase in which the behavior of the CTMC model (X^{lb} or X^{ub} in BT/RT and $X^{\text{lb,lb}}$ or $X^{\text{ub,ub}}$ in BT/BRT) from $S' = S - \{r\}$ until either a hit of state r or, if existing, hit of the absorbing state f , and from r until either next hit of state r or, if existing, hit of the absorbing state is characterized while keeping track of the amount of time spent in up states in S (states in S_{max}). That characterization is done through a truncated transformed CTMC model having, with some small error, the same interval availability distribution as the original CTMC model (X^{lb} or X^{ub} in BT/RT and $X^{\text{lb,lb}}$ or $X^{\text{ub,ub}}$ in BT/BRT). Fig. 3 clarifies the decomposition of the BT/RT method in terms of the RT method and Algorithm A of [25] and the decomposition of the BT/BRT method in terms of the BRT method, the RT method, and Algorithm A of [25]. In the figure, the CTMC model which is solved by Algorithm A of [25] when in BT/RT the lower bound for CRCD(t, s) has to be computed is called V_T^{lb} , the CTMC model which is solved by Algorithm A of [25] when in BT/BRT the lower bound for CRCD(t, s) has to be computed is called $V_T^{\text{lb,lb}}$, the CTMC model which is solved by Algorithm A of [25] when in BT/RT the upper bound for CRCD(t, s) has to be computed is called V_T^{ub} , and the CTMC model which is solved by Algorithm A of [25] when in BT/BRT the upper bound for CRCD(t, s) has to be computed is called $V_T^{\text{ub,ub}}$.

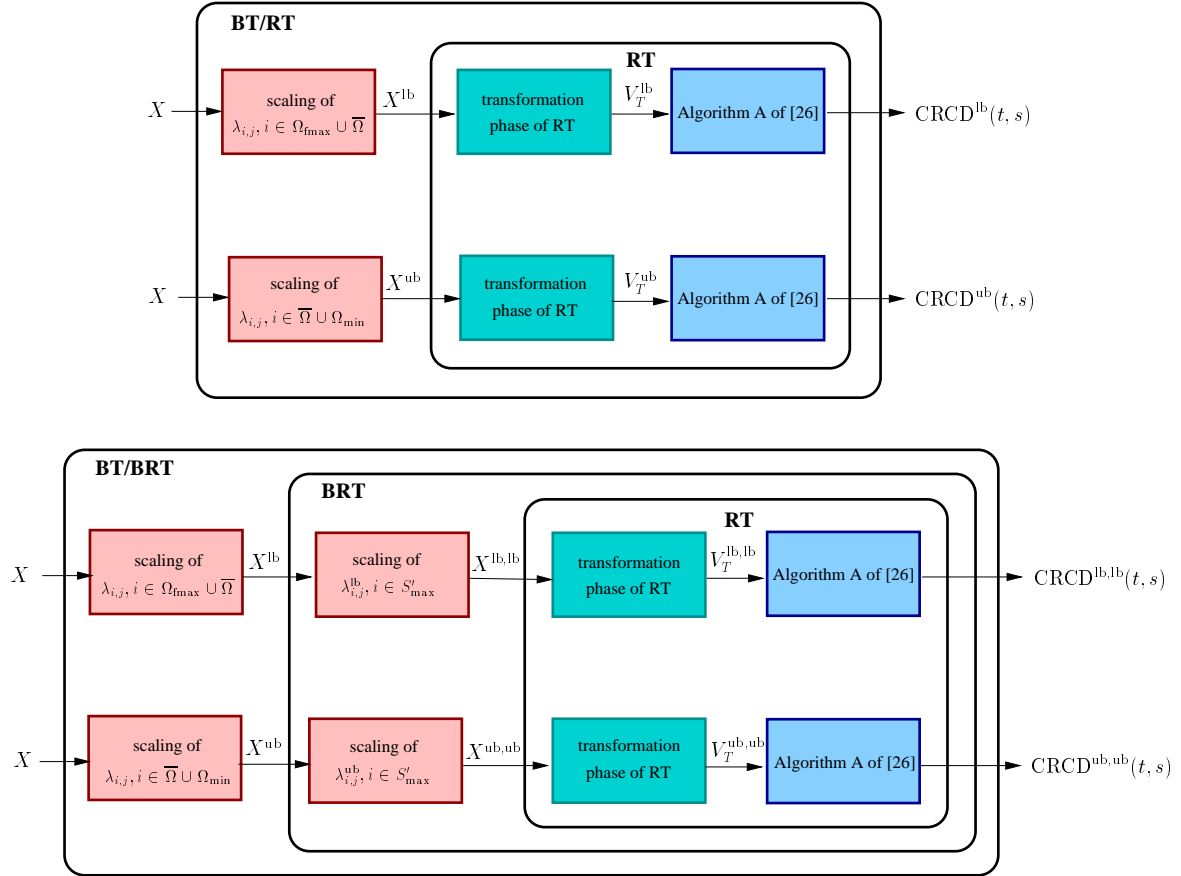


Figure 3: Schematic description of the BT/RT method in terms of the RT method and Algorithm A of [25] and schematic description of the BT/BRT method in terms of the BRT method, the RT method, and Algorithm A of [25].

Since the RT method is numerically stable, the BT/RT and BT/BRT methods will be numerically stable. In addition, since the computation error in RT is well-controlled, both BT/RT and BT/BRT will compute the bounds with well-controlled error. The assumed conditions in the BT/RT method for the rewarded CTMC X and the selection for the regenerative state r (conditions C1–C10) guarantee that the CTMCs X^{lb} and X^{ub} with subset of up states Ω_{max} and regenerative state r will be covered by the RT method (see [2]). Similarly, the assumed conditions in the BT/BRT method for the rewarded CTMC X and the regenerative state r (conditions C1–C11) guarantee that the CTMCs X^{lb} and X^{ub} with subset of up states Ω_{max} and regenerative state r will be covered by the BRT method (see [3]).

A possible alternative would be to use Algorithm A of [25] to solve the CTMC models X^{lb} and X^{ub} or $X^{\text{lb,lb}}$ and $X^{\text{ub,ub}}$. That alternative would be, however, expensive when Λt is large and the original rewarded CTMC model X is large.

The transformation phases involving the scaling of transition rates yielding X^{lb} or X^{ub} in BT/RT and $X^{\text{lb,lb}}$ or $X^{\text{ub,ub}}$ in BT/BRT have, in practice, negligible computational cost compared with the computational costs of the transformation phase of the RT method and the application of Algorithm A of [25] to the solution of the resulting transformed model. The flop count of the transformation phase of the RT method is, for large CTMC models X and $S'_{\text{max}} \neq \emptyset$, $CK(2T' + M|\Omega|) + I_{\alpha_{S'} > 0}CL(2T' + M|\Omega|)$, where T' is the number of transitions of X , $M = 11$ if $\Omega = S \cup \{f\}$ and $M = 9$ if $\Omega = S$, and C , K and L are truncation parameters defining the size of the transformed CTMC model. The flop count of the application of Algorithm A of [25] to the solution of the transformed model can be estimated as $2NC'T''$, where T'' is the number of transitions of the randomized DTMC of the transformed model with randomization rate equal to the maximum output rate of that model and N and C' are truncation parameters. An approximate estimate for T'' for the case $S'_{\text{max}} \neq \emptyset$ is $M'CK + I_{\alpha_{S'} > 0}M'CL$, where $M' = 9$ if $\Omega = S \cup \{f\}$ and $M' = 7$ if $\Omega = S$. Let ε be the absolute error with which the bounds have to be computed. The value of the truncation parameters C , N and C' can be easily estimated. For the BT/RT method, C , N and C' have values²

$$\begin{aligned} C^{\text{lb}} &= \min \left\{ c \geq 1 : \sum_{k=c+1}^{\infty} e^{-\Lambda^{\text{lb}} x_{\text{max}}^{\text{lb}}} \frac{(\Lambda^{\text{lb}} x_{\text{max}}^{\text{lb}})^k}{k!} \leq \frac{\varepsilon}{4} \right\}, \\ N^{\text{lb}} &= \min \left\{ n \geq 0 : \sum_{k=n+1}^{\infty} e^{-\Lambda^{\text{lb}} t_{\text{max}}} \frac{(\Lambda^{\text{lb}} t_{\text{max}})^k}{k!} \leq \frac{\varepsilon}{4} \right\}, \\ C'^{\text{lb}} &= \min \left\{ c \geq 0 : \sum_{k=c+1}^{\infty} e^{-\Lambda^{\text{lb}} x_{\text{max}}^{\text{lb}}} \frac{(\Lambda^{\text{lb}} x_{\text{max}}^{\text{lb}})^k}{k!} \leq \frac{\varepsilon}{4} \right\} \end{aligned}$$

when the lower bound is computed and values

$$\begin{aligned} C^{\text{ub}} &= \min \left\{ c \geq 1 : \sum_{k=c+1}^{\infty} e^{-\Lambda^{\text{ub}} x_{\text{max}}^{\text{ub}}} \frac{(\Lambda^{\text{ub}} x_{\text{max}}^{\text{ub}})^k}{k!} \leq \frac{\varepsilon}{4} \right\}, \\ N^{\text{ub}} &= \min \left\{ n \geq 0 : \sum_{k=n+1}^{\infty} e^{-\Lambda^{\text{ub}} t_{\text{max}}} \frac{(\Lambda^{\text{ub}} t_{\text{max}})^k}{k!} \leq \frac{\varepsilon}{4} \right\}, \end{aligned}$$

²Strictly speaking C' could be slightly larger in some cases.

$$C^{\text{ub}} = \min \left\{ c \geq 0 : \sum_{k=c+1}^{\infty} e^{-\Lambda^{\text{ub}} x_{\max}^{\text{ub}}} \frac{(\Lambda^{\text{ub}} x_{\max}^{\text{ub}})^k}{k!} \leq \frac{\varepsilon}{4} \right\}$$

when the upper bound is computed, where, with (t_i, s_i) , $1 \leq i \leq n$ being the (t, s) pairs at which the bounds for $\text{CRCDC}(t, s)$ have to be computed and with θ being a positive quantity $\ll 1$ (in our implementations, 10^{-4}), Λ^{lb} and Λ^{ub} , which are the maximum output rates of, respectively, X^{lb} and X^{ub} affected by the factor $(1 + \theta)$, have values

$$\Lambda^{\text{lb}} = (1 + \theta) \max \left\{ \max_{i \in \Omega_{\max} \cup \Omega_{\min}} \lambda_i, \max_{i \in \Omega_{\text{fmax}} \cup \bar{\Omega}} \frac{r_{\max} - r_{\min}}{r_{\max} - r_i} \lambda_i \right\}$$

and

$$\Lambda^{\text{ub}} = (1 + \theta) \max \left\{ \max_{i \in \Omega_{\max} \cup \Omega_{\text{fmax}}} \lambda_i, \max_{i \in \bar{\Omega} \cup \Omega_{\min}} \frac{r_{\max} - r_{\text{fmax}}}{r_{\max} - r_i} \lambda_i \right\},$$

$$x_{\max}^{\text{lb}} = \max_{1 \leq i \leq n} \frac{r_{\max} t_i - s_i}{r_{\max} - r_{\min}},$$

$$x_{\max}^{\text{ub}} = \max_{1 \leq i \leq n} \frac{r_{\max} t_i - s_i}{r_{\max} - r_{\text{fmax}}},$$

and

$$t_{\max} = \max_{1 \leq i \leq n} t_i.$$

For the BT/BRT method, C , N and C' have the values given by the previous expressions with (we remember, $\lambda_{\max} = \max_{i \in S'_{\max}} \lambda_i$ and $\lambda_{\min} = \min_{i \in S'_{\max}} \lambda_i$)

$$\Lambda^{\text{lb}} = (1 + \theta) \max \left\{ \lambda_{\max}, \max_{i \in S'_{\min}} \lambda_i, \max_{i \in S'_{\text{fmax}} \cup \bar{S}'} \frac{r_{\max} - r_{\min}}{r_{\max} - r_i} \lambda_i, \right. \\ \left. I_{r \in S_{\max} \cup S_{\min}} \lambda_r + I_{r \in S_{\text{fmax}} \cup \bar{S}} \frac{r_{\max} - r_{\min}}{r_{\max} - r_r} \lambda_r \right\},$$

$$\Lambda^{\text{ub}} = (1 + \theta) \max \left\{ D_C \lambda_{\min}, \max_{i \in S'_{\text{fmax}}} \lambda_i, \max_{i \in \bar{S}' \cup S'_{\min}} \frac{r_{\max} - r_{\text{fmax}}}{r_{\max} - r_i} \lambda_i, \right. \\ \left. I_{r \in S_{\max} \cup S_{\text{fmax}}} \lambda_r + I_{r \in \bar{S} \cup S_{\min}} \frac{r_{\max} - r_{\text{fmax}}}{r_{\max} - r_r} \lambda_r \right\},$$

and same values for x_{\max}^{lb} , x_{\max}^{ub} and t_{\max} . The values of Λ^{lb} and Λ^{ub} for the BT/BRT method are the maximum output rates of, respectively, $X^{\text{lb,lb}}$ and $X^{\text{ub,ub}}$, affected by the factor $(1 + \theta)$.

The truncation parameters K and L are known to be smooth functions of t (they are $O(\log(t/\varepsilon))$), but their actual values can be large, depending on the characteristics of the original rewarded CTMC model X and the selection for the regenerative state. Stronger results are available regarding the values of K and L for class C''_1 and class C'''_1 models with the natural selection $r = o$ for the regenerative state. Assume that X is a class C''_1 model with³ $|S_{\max}| \geq 2$ and let $R_{\max} =$

³For $|S_{\max}| = 1$, with the selection $r = o$ $S'_{\max} = \emptyset$ and [2] the RT method for computing the interval availability complementary distribution $\text{IAVCD}(t, p)$ will be highly efficient for, in X^{lb} , $p = (s/t - r_{\min})/(r_{\max} - r_{\min})$ will be close to 1 and, for X^{ub} , $p = (s/t - r_{\text{fmax}})/(r_{\max} - r_{\text{fmax}})$ will be close to 1, implying that BT/RT will be highly efficient for s close to $r_{\max} t$.

$\max_{i \in S_{\max}} \lambda_i / \min_{i \in S_{\max} - \{o\}} \lambda_i$. Then, since for class C_1'' models X with $|S_{\max}| \geq 2$, both X^{lb} and X^{ub} belong to the model class C_1 defined in [2] with the same o state, for class C_1'' models with $|S_{\max}| \geq 2$ and the natural selection $r = o$, the truncation parameters K and L when BT/RT is used to compute both the lower and the upper bound should increase with R_{\max} and, for $R_{\max} \gg 1$, can be roughly upper bounded by $30R_{\max}$. Assume X is a class C_1''' model. Since for class C_1''' models X , both X^{lb} and X^{ub} belong to the model class C_1' defined in [3] with the same o state, for class C_1''' models with the natural selection $r = o$, the truncation parameters K and L when BT/BRT is used to compute both the lower and the upper bound should be very small for $D_C = 1$, should increase with D_C and, for $D_C \gg 1$, can be roughly upper bounded by $30D_C$. Thus, for class C_1''' models with the natural selection $r = o$, in BT/BRT, the D_C control parameter allows to trade-off bounds tightness with computational cost. Since the truncation parameters C^{lb} and C^{ub} will have moderate values when s is close to $r_{\max} t$ for every (t, s) pair of interest, for large C_1''' models, we should expect the BT/BRT method with $r = o$ and $D_C = 1$ to be relatively inexpensive when s is close to $r_{\max} t$ for every (t, s) pair of interest. Whether the BT/RT with $r = o$ is also relatively inexpensive in that case depends on how large R_{\max} is.

4 Analysis

In this section we analyze the performances of the BT/RT and BT/BRT methods using a large performability example. The example is a model of a fault-tolerant multiprocessor including 16 processors interconnected by a 8-node hypercube, as shown in Fig. 4. Processors fail with rate λ_P ; nodes of the hypercube fail with rate λ_N ; links of the hypercube fail with rate λ_L . A fault of a processor is covered with probability C_P ; a fault of a node of the hypercube is covered with probability C_N . Coverage to link faults is assumed perfect. There is an unlimited number of repairmen to repair components in covered failure. The repair rate is μ_P for processors, μ_N for nodes, and μ_L for links. A completely down system because there was an uncovered fault is brought to a fully operational state without failed components at rate μ_G . Components do not fail when the system is completely down. Unless otherwise stated, we will use the set of model parameters $\lambda_P = 2 \times 10^{-5} \text{h}^{-1}$, $\lambda_N = 10^{-5} \text{h}^{-1}$, $\lambda_L = 5 \times 10^{-6} \text{h}^{-1}$, $C_P = 0.99$, $C_N = 0.995$, $\mu_P = 0.1 \text{h}^{-1}$, $\mu_N = 0.05 \text{h}^{-1}$, $\mu_L = 0.05 \text{h}^{-1}$, and $\mu_G = 0.2 \text{h}^{-1}$. We assume the availability of diagnosis and reconfiguration procedures to determine a subset of interconnected unfailed processors of maximal size and to reconfigure the multiprocessor so that it works using such a maximal subset. As reward rates, we use the speedup function of the number of connected processors in the healthy subset in which the system is configured described by Table 1. Thus, r_{\max} will be equal to 12h^{-1} . The cumulative reward will be then the accumulated performance (normalized with respect to the performance rate of a single processor) and we should expect it to be close to $12t$ with high probability.

An exact rewarded CTMC model of the multiprocessor system has an unmanageable state space. Instead, we will use bounding models with state space $S \cup \{f\}$, where f is an absorbing state in which the bounding models enter when the exact model would exit subset S and S includes the states with up to four covered faults and the state in which the system is down due to an uncovered




Figure 4: Architecture of the fault-tolerant multiprocessor system.

fault. By assigning to the absorbing state a reward rate $r_f = 0$ we obtain a lower bounding model; by assigning to the absorbing state a reward rate $r_f = 12 \text{ h}^{-1}$ we obtain an upper bounding model. The bounding models have 213,055 states and 2,072,658 transitions. The bounds given by those bounding models are very tight. This can be justified by the fact that the probability that the bounding models have entered state f by time $t = 20,000 \text{ h}$ (the largest time we will consider) is very small, e.g., assuming that the initial state is the state without failed components, that probability is only 8.764×10^{-10} , implying that the CRCD(t, s) measures given by the bounding models would differ in less than 8.764×10^{-10} , for $t \leq 20,000 \text{ h}$. We will use the methods to compute a lower bound for the lower bounding model and an upper bound for the upper bounding model. Since the solutions of the bounding models are extremely tight, virtually all the difference between the lower and upper bounds for the CRCD(t, s) measure of the exact model thus obtained are attributable to the limited tightness of the bounds given by the BT/RT and BT/BRT methods. Both bounding models belong to model classes C_1'' and C_1''' , being state o the state without failed components, which will be taken as regenerative state in both BT/RT and BT/BRT. All methods are run with a single (t, s) target and an error requirement $\varepsilon = 10^{-10}$. CPU times are measured/estimated in a workstation with a Sun-Blade 1000 processor and 4 GB of memory (significantly larger than the memory consumption for all methods). We will start assuming that initially the multiprocessor is in the state without failed components.

Tables 2 and 3 summarize the performances of, respectively, BT/BRT with $D_C = 1$ and BT/RT in terms of bounds tightness and values of the truncation parameters C and K of the model transformation step implicit to the RT method (the truncation controlled by the parameter L is not performed because the initial probability distribution of the models in S is concentrated in the regenerative state) when both bounds are computed for increasing values of t and $s/t = 11.99, 11.999$, both of which are reasonable choices, since the steady-state reward rate of the exact model can be estimated to be equal to 11.9966. That estimation was obtained by computing the expected transient reward rates at time t of the bounding models, which bound the expected transient reward rate at time t of the exact model, for increasing values of t and taking the

Table 1: Speedups of the multiprocessor as a function of the number of connected operational processors.

processors	speedup
1	1
2	1.96667
3	2.9
4	3.8
5	4.66667
6	5.5
7	6.3
8	7.06667
9	7.8
10	8.5
11	9.16667
12	9.8
13	10.4
14	10.96667
15	11.5
16	12

value at which they “stabilize” (before the probability of the absorbing state is significant). We call C^{lb} and K^{lb} the values of the truncation parameters C and K corresponding to the application of the RT method to X^{lb} in BT/RT and $X^{\text{lb,lb}}$ in BT/BRT and C^{ub} and K^{ub} the values of the truncation parameters C and K corresponding to the application of the RT method to X^{ub} in BT/BRT and to $X^{\text{ub,ub}}$ in BT/BRT. We can first note that the bounds obtained by the methods are quite tight for all values of t and s considered. For BT/RT we define the relative error as $\max\{(\text{CRCD}^{\text{ub}}(t, s) - \text{CRCD}^{\text{lb}}(t, s))/2 / ((\text{CRCD}^{\text{lb}}(t, s) + \text{CRCD}^{\text{ub}}(t, s))/2), (\text{CRCD}^{\text{ub}}(t, s) - \text{CRCD}^{\text{lb}}(t, s))/2 / (1 - (\text{CRCD}^{\text{lb}}(t, s) + \text{CRCD}^{\text{ub}}(t, s))/2)\}$, i.e. as the maximum of the relative errors over $\text{CRCD}(t, s)$ and $1 - \text{CRCD}(t, s)$ when $(\text{CRCD}^{\text{lb}}(t, s) + \text{CRCD}^{\text{ub}}(t, s))/2$ is taken as estimate for $\text{CRCD}(t, s)$. The relative error for BT/BRT is defined similarly. The relative error varies from 0.105 % to 2.97 % for the BT/BRT method with $D_C = 1$ and from 0.0808 % to 2.72 % for the BT/RT method. The relative error is significantly smaller for s closer to $r_{\max} t = 12t$. Regarding the truncation parameters C and K , we can note that, as predicted theoretically, K is always very small in the BT/BRT method with $D_C = 1$. The value of the truncation parameter K is significantly larger for the BT/RT method and is upper bounded by $30R_{\max} \approx 120$. In both cases, the truncation parameter K increases logarithmically with t and is independent of s . The truncation parameter C increases with t and as s gets apart from $r_{\max} t = 12t$.

Fig. 5 plots the CPU times consumed by the methods for the computation of the lower bound for $\text{CRCD}(t, s)$ (lb) and the upper bound for $\text{CRCD}(t, s)$ (ub) as a function of t for $s = 11.99t$ and $s = 11.999t$. We can note that the CPU times increase as s gets apart from $r_{\max} t = 12t$ and increase fast with t . The latter is due to two reasons. The first one is that the computational cost

Table 2: Results for BT/BRT with $D_C = 1$ when the initial state is the state without failed components.

t (h)	s/t	$\text{CRCD}^{\text{lb,lb}}(t, s)$	$\text{CRCD}^{\text{ub,ub}}(t, s)$	error (%)	C^{lb}	K^{lb}	C^{ub}	K^{ub}
1	11.99	0.99960102	0.99960874	0.977	4	5	4	4
10	11.99	0.99607366	0.99615272	1.02	5	7	5	6
100	11.99	0.96665320	0.96733703	1.04	10	9	10	9
1,000	11.99	0.90000197	0.90216084	1.09	26	10	26	10
10,000	11.99	0.95260672	0.95454693	2.09	103	11	103	11
20,000	11.99	0.98111980	0.98220912	2.97	173	11	173	11
1	11.999	0.99960011	0.99960095	0.105	2	5	2	4
10	11.999	0.99601099	0.99602264	0.146	4	7	4	6
100	11.999	0.96136877	0.96150542	0.177	5	9	5	9
1,000	11.999	0.71162283	0.71251906	0.156	10	10	10	10
10,000	11.999	0.19927493	0.20078590	0.378	26	11	26	11
20,000	11.999	0.08224490	0.08336233	0.675	37	11	37	11

Table 3: Results for BT/RT when the initial state is the state without failed components.

t (h)	s/t	$\text{CRCD}^{\text{lb}}(t, s)$	$\text{CRCD}^{\text{ub}}(t, s)$	error (%)	C^{lb}	K^{lb}	C^{ub}	K^{ub}
1	11.99	0.99960108	0.99960874	0.969	4	7	4	7
10	11.99	0.99607762	0.99615272	0.967	5	13	5	13
100	11.99	0.96670875	0.96733698	0.953	10	38	10	38
1,000	11.99	0.90018604	0.90216058	0.999	26	79	26	79
10,000	11.99	0.95277489	0.95454669	1.91	103	88	103	88
20,000	11.99	0.98121513	0.98220898	2.72	173	91	173	91
1	11.999	0.99960018	0.99960095	0.0964	2	7	2	7
10	11.999	0.99601501	0.99602264	0.0958	4	13	4	13
100	11.999	0.96143288	0.96150536	0.0941	5	38	5	38
1,000	11.999	0.71205362	0.71251843	0.0808	10	79	10	79
10,000	11.999	0.20000281	0.20078479	0.195	26	88	26	88
20,000	11.999	0.08278256	0.08336150	0.348	37	91	37	91




Figure 5: CPU times in seconds consumed by the BT/BRT method with $D_C = 1$ and the BT/RT method for $s = 11.99 t$ (left) and $s = 11.999 t$ (right) as a function of t .

of the transformation step implicit to the RT method is approximately proportional to the truncation parameters C and K , and C increases considerably with t . The second one is that the computational cost of Algorithm A of [25] is approximately proportional to C and K (since the size of truncated transformed model which is solved by the algorithm is proportional to C and K) and, for large t , increases, approximately quadratically with t , since, for large t , the truncation parameter N increases approximately linearly with t and, for large t (typically much larger), the truncation parameter C' also increases approximately linearly with t . The higher computational cost of the lower bound is due to the fact that the maximum output rate, Λ^{lb} , of the model (X^{lb} in BT/RT and $X^{\text{lb,lb}}$ in BT/BRT) built during the computation of the lower bound is significantly larger than the maximum output rate, Λ^{ub} , of the model (X^{ub} in BT/RT and $X^{\text{ub,ub}}$ in BT/BRT) built during the computation of the upper bound and this makes the computational cost of the application of Algorithm A of [25] substantially higher. For the largest t considered ($t = 20,000$ h) the resulting CPU times are considerable, specially when the lower bound is computed and $s = 11.99 t$. Thus, for those t and s , computation of the lower bound consumed 231,000 s (about 64 hours) under the BT/RT method and 14,500 s (about 4.0 hours) under the BT/BRT method with $D_C = 1$. Since the bounds achieved by BT/BRT with $D_C = 1$ are only slightly less tight than the bounds achieved by BT/RT, the significantly smaller computational cost of the BT/BRT method with $D_C = 1$ makes that method more attractive than the BT/RT method, when both are applicable. For class C_1'' models with $S_{\text{max}} = \{o\}$ only the BT/RT method is applicable and, since for those models BT/RT with $r = o$ will be relatively inexpensive for s close to $r_{\text{max}} t$, that method should be used.

To illustrate the relative importance of the two components of the computational cost of the methods: model transformation (up to V_T^{lb} and V_T^{ub} in BT/RT and up to $V_T^{\text{lb,lb}}$ and $V_T^{\text{ub,ub}}$ in BT/BRT) and application of Algorithm A of [25], Fig. 6 displays the breakdown into these two components of the CPU times consumed by the methods when the lower bound is computed for $s = 11.99 t$ and $s = 11.999 t$, as a function of t . We label the CPU times consumed in the model transformation phase by “trans”, the CPU times consumed in the application of Algorithm A of [25] to the transformed model by “sol”, and the total CPU times by “tot”. We remember that the model




Figure 6: Breakdown of the CPU times consumed by the BT/BRT method with $D_C = 1$ and the BT/RT method when the lower bound is computed for $s = 11.99 t$ (left) and $s = 11.999 t$ (right) as a function of t .

transformation phases involving scaling of transition rates (up to X^{lb} and X^{ub} in BT/BRT and up to $X^{\text{lb,lb}}$ and $X^{\text{ub,ub}}$ in BT/BRT) have negligible relative computational cost and, then, virtually all CPU times labeled by “trans” are due to the transformation phase implicit to the RT method. We can note that for large t and $s = 11.99 t$, almost all the computational cost of the methods is due to the application of Algorithm A. For $s = 11.999 t$, the computational cost due to the application of Algorithm A is relatively smaller but also dominant for large t . The computational cost of the application of Algorithm A is approximately proportional to the size of the transformed model handled by the method (V_T^{lb} in BT/RT and $V_T^{\text{lb,lb}}$ in BT/BRT) and the truncation parameters N and C' of the method. Table 4 gives the size of the truncated transformed model and the truncation parameters N and C' corresponding to the application of the methods for computing the lower bound, for increasing t and the two values considered for s/t . The approximately linear dependence of N on t and the increase of the size of the transformed model and C' with t explain the fast growth of the computational cost of the application of Algorithm A with t and the fast growth for large t of the computational cost of the BT/RT and BT/BRT methods. The relative importance of that second component of the computational cost (application of Algorithm A) decreases as the size of the original model increases and, for large enough X , the second component should be negligible. The experiments seem to indicate, however, that X has to be extremely large for that to happen.

We compare next the computational costs of the methods developed in this paper with that of the method described in [16, 17]. The considered values for s satisfy $s \geq r_{\text{max}} t$ and, then, the method described in [16, 17] works in its most favorable case. The CPU times consumed by the method described in [16, 17] for $t = 20,000$ h are very large and we estimated them based on measured CPU times for $t = 1,000$ h and the rule, that we found accurate, that, for a given, large model and $s \geq r_{\text{max}} t$, the CPU times are proportional to $C(N - C) + mC^2/2$, where $m + 1$ is the number of different reward rates and N and C are the truncation parameters of the method. Table 5 gives actual speedups and potential speedups (neglecting for BT/RT and BT/BRT the computational cost of the application of Algorithm A of [25]). Potential speedups reflect the speedups which would be

Table 4: Sizes of the transformed models in terms of numbers of states and numbers of transitions and truncation parameters of their solution by Algorithm A of [25] when computing the lower bound using the BT/BRT method with $D_C = 1$ and the BT/RT method.

t (h)	s/t	BT/BRT				BT/RT			
		states	trans.	N^{lb}	C^{lb}	states	trans.	N^{lb}	C^{lb}
1	11.99	39	109	28	4	57	180	28	4
10	11.99	69	224	118	5	135	488	118	5
100	11.99	171	612	769	10	780	3,048	769	10
1,000	11.99	480	1,784	6,529	26	4,137	16,412	6,529	26
10,000	11.99	2,073	7,848	61,741	104	18,012	71,604	61,741	104
20,000	11.99	3,473	13,168	122,530	174	31,233	124,208	122,530	174
1	11.999	23	55	28	2	33	94	28	2
10	11.999	57	180	118	4	111	396	118	4
100	11.999	91	312	769	5	410	1,588	769	5
1,000	11.999	192	696	6,529	10	1,641	6,492	6,529	10
10,000	11.999	533	1,996	61,741	26	4,614	18,320	61,741	26
20,000	11.999	753	2,832	122,530	37	6,753	26,832	122,530	37

obtained for large enough X . They also reflect the speedups which would be achieved if Algorithm A of [25] were replaced by a substantially more efficient back-end, a direction we are pursuing. We can note that, for large t , the actual speedups are important for the BT/BRT method with $D_C = 1$ and significant for the BT/RT method. In practical terms, the speedups achieved by the BT/BRT method with $D_C = 1$ make affordable the computation of bounds for very large class C_1''' models out of hand for “exact” available methods (the method described in [16, 17] can be considered the state-of-the-art exact general-purpose method for rewarded CTMC models X for values of t and s for which $\max_{i \in \Omega} \lambda_i t$ is large and s is $\geq r_{\text{fmax}} t$). To illustrate the point, for $t = 20,000$ h and $s = 11.99 t$, the CPU time of the BT/BRT method with $D_C = 1$ is 14,500 s (about 4.0 hours) for the lower bound and 1440 s (about 24 minutes) for the upper bound, implying that both bounds are obtained in 15,940 s (about 4.4 hours), while our estimate for the CPU time for the method described in [16, 17] is 1,083,000 s (about 13 days). If Algorithm A of [25] could be replaced by a substantially more efficient back-end, the CPU times of the BT/BRT method with $D_C = 1$ would be reduced up to 858 s (about 14 minutes) for the lower bound and up to 841 s (about 14 minutes) for the upper bound, implying that both bounds would be obtained in 1699 s (about 28 minutes).

For large t , the tightness of the bounds obtained by BT/RT and BT/BRT can be, intuitively, explained by the fact that most of the time spent by all X , X^{lb} , X^{ub} , $X^{\text{lb,lb}}$ and $X^{\text{ub,ub}}$ in S is spent in state o and the rewarded CTMC models only differ in that the holding times in states in $S - \{o\}$ are different (the reward lost with respect to the reward that would have been earned had the state a reward rate r_{max} in each visit to states $i \in S_{\text{max}} - \{o\} \cup S_{\text{fmax}} \cup \overline{S} \cup S_{\text{min}}$ is, by construction,

Table 5: Actual/potential speedups of the BT/BRT method with $D_C = 1$ and the BT/RT method over the method described in [16, 17].

method	t (h)	s/t	lower bound	upper bound
BT/BRT	1,000	11.99	36.6/38.1	39.3/39.5
BT/BRT	20,000	11.99	74.7/1,261	753/1,287
BT/BRT	1,000	11.999	17.0/17.2	17.5/17.5
BT/BRT	20,000	11.999	242/636	619/666
BT/RT	1,000	11.99	9.07/10.4	10.5/10.6
BT/RT	20,000	11.99	4.69/187	69.6/190
BT/RT	1,000	11.999	7.15/7.50	7.49/7.51
BT/RT	20,000	11.999	26.9/145	121/145

identical in all models and the reward rate associated with the absorbing state f in the bounding models is identical to the reward rate associated with state f in X). This will be the case for any class C'_1 or any class C'''_1 model for which the following additional conditions are satisfied:

C15. The partition $S_0 \cup S_1 \cup \dots \cup S_{N_C}$ for S_{\max} can be extended to a partition $S'_0 \cup S'_1 \cup \dots \cup S'_{N'_C}$ for S satisfying the properties:

P10. $S'_0 = S_0 = \{o\}$.

P11. For each $i \in S'_k$, $0 < k \leq N'_C$, $\lambda_{i, S'_k - \{i\} \cup S'_{k+1} \cup \dots \cup S'_{N'_C}}$ if $\Omega = S$ or

$\lambda_{i, S'_k - \{i\} \cup S'_{k+1} \cup \dots \cup S'_{N'_C} \cup \{f\}}$ if $\Omega = S \cup \{f\}$ is significantly smaller than $\lambda_{i, S'_0 \cup \dots \cup S'_{k-1}}$.

P12. $\lambda_o \ll \min\{\min_{i \in S_{\max} - \{o\} \cup S_{f_{\max}}} \lambda_i, \min_{i \in \bar{S} \cup S_{\min}} ((r_{\max} - r_{f_{\max}})/(r_{\max} - r_i)) \lambda_i\}$.

C16. If $\Omega = S \cup \{f\}$, $f \in \Omega_{\max} \cup \Omega_{\min}$ if the lower bound has to be computed and $f \in \Omega_{\max} \cup \Omega_{f_{\max}}$ if the upper bound has to be computed.

The reasons are that P10 and P11 imply that from any state $i \in S - \{o\}$ the embedded DTMC of all X , X^{lb} , X^{ub} , $X^{\text{lb,lb}}$, and $X^{\text{ub,ub}}$ will go towards state o with almost 1 probability, P12 implies that the holding time in each visit to a state $i \in S - \{o\}$ will be much smaller than the holding time in each visit to state o in all X , X^{lb} , X^{ub} , $X^{\text{lb,lb}}$, and $X^{\text{ub,ub}}$, and C16 implies that, in the case $\Omega = S \cup \{f\}$, the reward rate of state f will not be modified by the model transformation. Class C'_1 and C'''_1 models with the additional conditions include typical failure/repair performability models of fault-tolerant systems with exponential failure and repair time distributions and repair in every state with failed components and a reward rate structure which is a non-increasing function of the collection of failed components. Partitions for, respectively, S_{\max} and S showing that would be $S_k = \{\text{states with maximum reward rate and } k \text{ failed components}\}$ and $S'_k = \{\text{states with } k \text{ failed components}\}$. Then, properties P5 and P8 of the partition for S_{\max}

would follow from the fact that $\lambda_{i, S_k - \{i\} \cup S_{k+1} \cup \dots \cup S_{N_C} \cup S_{\text{fmax}} \cup \bar{S} \cup S_{\text{min}}}$, $i \in S_k$, $0 \leq k \leq N_C$, only collects failure transitions while $\lambda_{i, S_0 \cup \dots \cup S_{k-1}}$, $i \in S_k$, $0 < k \leq N_C$, only collects repair transitions; similarly, property P11 of the partition for S would follow from the fact that $\lambda_{i, S'_k - \{i\} \cup S'_{k+1} \cup \dots \cup S'_{N'_C}}$, $0 < k \leq N'_C$, only collects failure transitions while $\lambda_{i, S'_0 \cup \dots \cup S'_{k-1}}$, $0 < k \leq N'_C$, only collects repair transitions; finally, property P12 of the partition for S would follow from the fact that λ_o only collects failure transitions while every λ_i , $i \in S_{\text{max}} - \{o\} \cup S_{\text{fmax}} \cup \bar{S} \cup S_{\text{min}}$, collects repair transitions. Bounding models of that type would also be covered, since in those models the absorbing state would be entered following failure transitions, implying that property P11 of the partition for S would continue to follow, a lower bounding model of that type would naturally assign to the absorbing state capturing the pruned behavior the minimum reward rate of the exact model, and an upper bounding model of that type would naturally assign to that state the maximum reward rate of the exact model. Then, since when a lower bounding model is used only the lower bound would be typically computed, condition C16 would be satisfied; similarly, since when an upper bounding model is used only the upper bound would be typically computed, condition C16 would also be satisfied. As will be illustrated next using the example, uncovered faults taking the system to down states with null reward rate which are recovered into up states by fast recovery activities are not a problem: it suffices to consider a partition for S in which every down state is after every state from which the down state can be reached and after the up state to which the system is recovered from that down state.

The bounding rewarded CTMC models under consideration satisfy conditions C15 and C16. A partition for S showing that is $S'_0 \cup S'_1 \cup S'_2 \cup S'_3 \cup S'_4 \cup S'_5$, where S'_0 includes the single state o , S'_k , $1 \leq k \leq 4$ includes the states with k covered faults and S'_5 includes the state in which the multiprocessor is completely down due to an uncovered fault. The induced partition on S_{max} , $S_0 \cup S_1 \cup S_2 \cup S_3 \cup S_4$, $S_0 = S'_0$, $S_1 = S'_1 \cap S_{\text{max}} \neq \emptyset$, $S_2 = S'_2 \cap S_{\text{max}} \neq \emptyset$, $S_3 = S'_3 \cap S_{\text{max}} \neq \emptyset$, $S_4 = S'_4 \cap S_{\text{max}} \neq \emptyset$ shows that the bounding rewarded CTMC models belong to model classes C''_1 and C'''_1 . Properties P11 and P12 of the partition are satisfied moderately by the bounding rewarded CTMC models since $\max_{0 < k \leq 5} \max_{i \in S'_k} \lambda_{i, S'_k - \{i\} \cup S'_{k+1} \cup \dots \cup S'_5 \cup \{f\}} = 5.15 \times 10^{-4} \text{ h}^{-1}$, $\min_{0 < k \leq 5} \min_{i \in S'_k} \lambda_{i, S'_0 \cup \dots \cup S'_{k-1}} = 0.05 \text{ h}^{-1}$, $\lambda_o = 5.2 \times 10^{-4} \text{ h}^{-1}$, and $\min\{\min_{i \in S_{\text{max}} - \{o\} \cup S_{\text{fmax}}} \lambda_i, \min_{i \in \bar{S} \cup S_{\text{min}}} ((r_{\text{max}} - r_{\text{fmax}})/(r_{\text{max}} - r_i)) \lambda_i\} \approx 0.0227 \text{ h}^{-1}$. We should expect the bounds to be tighter were these properties satisfied more strongly. The fact that the bounds in BT/BRT are also tight for small t has to do with the fact that all the initial probability distribution of X in S_{max} is concentrated in state o . Table 6 gives the bounds obtained by BT/BRT with $D_C = 1$ when the initial state is the state in which the link from node N0 to node N1 is in covered fault (which belongs to $S_{\text{max}} - \{o\}$). We can note that, in that case, the bounds are not tight for small values of t . Having some initial probability distribution in $S_{\text{max}} - \{o\}$ does not seem to degrade the quality of the bounds obtained by BT/RT for small values of t as Table 7 illustrates. We should, however, expect a degradation of the quality of the bounds obtained by BT/RT for small values of t when the model has some initial probability distribution in $S_{\text{fmax}} \cup \bar{S} \cup S_{\text{min}}$. This is because transition rates from $S_{\text{fmax}} \cup \bar{S}$ are scaled when constructing X^{lb} and transition rates from $\bar{S} \cup S_{\text{min}}$ are scaled when constructing X^{ub} .

Table 6: Results for BT/BRT with $D_C = 1$ when the initial state is the state in which the link from node N0 to node N1 is in covered fault.

t (h)	s/t	$\text{CRCD}^{\text{lb,lb}}(t, s)$	$\text{CRCD}^{\text{ub,ub}}(t, s)$	error (%)
1	11.99	0.99852766	0.99960875	58.0
10	11.99	0.99104228	0.99615350	39.9
100	11.99	0.96177512	0.96734418	7.86
1,000	11.99	0.89846302	0.90216318	1.86
10,000	11.99	0.95246687	0.95454714	2.24
20,000	11.99	0.98108013	0.98220917	3.08
1	11.999	0.99852441	0.99960096	57.4
10	11.999	0.99090241	0.99602346	39.2
100	11.999	0.95574350	0.96151370	6.97
1,000	11.999	0.70801574	0.71252460	0.778
10,000	11.999	0.19866954	0.20078683	0.530
20,000	11.999	0.08202169	0.08336267	0.811

Table 7: Results for BT/RT when the initial state is the state in which the link from node N0 to node N1 is in covered fault.

t (h)	s/t	$\text{CRCD}^{\text{lb}}(t, s)$	$\text{CRCD}^{\text{ub}}(t, s)$	error (%)
1	11.99	0.99960108	0.99960874	0.969
10	11.99	0.99607762	0.99615272	0.967
100	11.99	0.96670875	0.96733698	0.952
1,000	11.99	0.90018604	0.90216058	0.999
10,000	11.99	0.95277489	0.95454669	1.91
20,000	11.99	0.98121513	0.98220898	2.72
1	11.999	0.99960018	0.99960095	0.0964
10	11.999	0.99601501	0.99602264	0.0958
100	11.999	0.96143288	0.96150536	0.0941
1,000	11.999	0.71205362	0.71251843	0.0808
10,000	11.999	0.20000281	0.20078479	0.195
20,000	11.999	0.08278256	0.08336150	0.348




Figure 7: $\text{CRCD}(t, (12 - p)t)$ as a function of p for several values of t (h).

We will end this section using the BT/BRT method with $D_C = 1$ to analyze the performability of the fault-tolerant multiprocessor. We plot in all figures $(\text{CRCD}^{\text{lb,lb}}(t, s) + \text{CRCD}^{\text{ub,ub}}(t, s))/2$ and do not plot explicitly the bounds, since they are close enough to consider the estimate $(\text{CRCD}^{\text{lb,lb}}(t, s) + \text{CRCD}^{\text{ub,ub}}(t, s))/2$ almost exact at the plot resolution. Fig. 7 analyzes how the shape of $\text{CRCD}(t, s)$, as a function of s , depends on t . Let $\text{SSRR} = \lim_{t \rightarrow \infty} E[r_{X_{\text{exact}}}(t)]$, where X_{exact} is the rewarded CTMC modeling exactly the multiprocessor system (with state space of unmanageable size). SSRR has a value approximately equal to 11.9966. Using renewal reward process and regenerative process theories (see, for instance, [24]), $\text{CRCD}(t, s)$ for the exact rewarded CTMC model X_{exact} has an asymptotic shape with $\text{CRCD}(t, s) = 1$ for $s/t < \text{SSRR}$ and $\text{CRCD}(t, s) = 0$ for $s/t > \text{SSRR}$. However, that asymptotic shape is reached very slowly, implying that the $\text{CRCD}(t, s)$ measure could be of interest for very large values of t and stressing the need for efficient methods to compute the measure for large t . In Fig. 8 we perform a sensitivity analysis. More specifically, we analyze how $\text{CRCD}(t, s)$ for $t = 2$ years is improved when we improve the repair actions in three different ways: (1) faster repair of processors ($\mu_P = 0.2 \text{ h}^{-1}$), (2) faster repair of components of the hypercube ($\mu_N = \mu_L = 0.1 \text{ h}^{-1}$), and (3) faster repair of down systems due to an uncovered fault ($\mu_G = 0.4 \text{ h}^{-1}$). In all cases, we double the corresponding repair rates with respect to the baseline values. We can note that the most efficient way of improving the performability depends on the required probability level. When the accumulated performance has to be guaranteed with very high probability, improving the repair of down systems due to an uncovered fault is the most efficient alternative; when the accumulated performance has to be guaranteed with moderate probability, the most efficient alternative is to improve the repair of processors; for intermediate values of the probability with which the accumulated performance has to be guaranteed, the most efficient alternative is to improve the repair of components of the hypercube. Using the simpler $\text{EARR}(t) = E[(1/t) \int_0^t r_{X(\tau)} d\tau]$ measure would have led to the conclusion that the most efficient alternative is to improve the repair of components of the hypercube, as Table 8 illustrates. Thus, use of the more detailed $\text{CRCD}(t, s)$ measure provides interesting information to guide the maintenance of the fault-tolerant multiprocessor system.




Figure 8: $\text{CRCD}(t, (12 - p)t)$ as a function of p for $t = 2$ years and the baseline repair rates, a set of repair rates with $\mu_P = 0.2 \text{ h}^{-1}$, a set of repair rates with $\mu_N = \mu_L = 0.1 \text{ h}^{-1}$, and a set of repair rates with $\mu_G = 0.4 \text{ h}^{-1}$.

Table 8: $\text{EARR}(t)$ measure for $t = 2$ years and the baseline repair rates, a set of repair rates with $\mu_P = 0.2 \text{ h}^{-1}$, a set of repair rates with $\mu_N = \mu_L = 0.1 \text{ h}^{-1}$, and a set of repair rates with $\mu_G = 0.4 \text{ h}^{-1}$.

case	$\text{EARR}(t)$
baseline	11.996558
$\mu_P = 0.2 \text{ h}^{-1}$	11.997350
$\mu_N = \mu_L = 0.1 \text{ h}^{-1}$	11.997379
$\mu_G = 0.4 \text{ h}^{-1}$	11.996666

5 Conclusions

We have developed two methods, BT/RT and BT/BRT, for the computation of bounds for the cumulative reward complementary distribution measure $\text{CRCD}(t, s)$ for CTMC models with reward rates associated with states. Both methods require the selection of a regenerative state, are numerically stable and compute the bounds with well-controlled error. For a class of rewarded CTMC models, class C_1''' , and a particular, natural selection for the regenerative state the BT/BRT method allows to trade off bounds tightness with computational cost and will provide bounds at moderate computational costs for s close to $r_{\max} t$, where r_{\max} is the largest reward rate of the model, allowing in such a case a rigorous analysis, with error bounds, of some large models which were beyond the scope of previously proposed “exact” methods. For a class of rewarded CTMC models, class C_1'' , slightly wider than class C_1''' , and a particular, natural selection for the regenerative state, the BT/RT method will yield tighter bounds at a higher computational cost. When the rewarded CTMC model satisfies additional conditions, the bounds obtained by the less expensive version of the BT/BRT method seem to be tight for any value of t or not small values of t , depending on whether the initial probability distribution of the CTMC model in the subset of states with maximum reward rate different, if existing, from the absorbing state is concentrated in the natural selection for the regenerative state or not, and the bounds obtained by the BT/RT method seem to be tight for any value of t or not small values of t , depending on whether the initial probability distribution of the CTMC model in the subset of states different, if existing, from the absorbing state is concentrated in the states with maximum reward rate or not. Class C_1''' and class C_1'' models with those additional conditions include both exact and bounding typical failure/repair performability models of fault-tolerant systems with exponential failure and repair time distributions and repair in every state with failed components and a reward rate structure which is a non-increasing function of the collection of failed components. Combined with bounding techniques, the methods developed in the paper should allow the analysis of performability models of quite complex degradable and repairable fault-tolerant parallel and distributed systems in affordable CPU times, with the obvious implications. In the future, we are planning to develop more efficient methods than Algorithm A of [25] to act as a back-ends of BT/RT and BT/BRT. For large t , this would yield significant reductions in the CPU times of those methods.

In Section 1, we concentrated our review of numerical methods for computing the $\text{CRCD}(t, s)$ measure to general purpose methods which, besides (possibly) Ω being finite, do not impose any restrictions on X . None of those methods is able to handle in reasonable CPU times when Λt is large large rewarded CTMC models of the type for which BT/RT and BT/BRT seem to give tight bounds at a relatively moderate computational cost when s is close to $r_{\max} t$. Numerical methods exist which can be (are) more efficient than the method described in [16, 17] for rewarded acyclic CTMC models [18] (rewarded block acyclic CTMC models [19]). Rewarded acyclic and rewarded block acyclic CTMC models naturally arise when modeling, respectively, non-repairable fault-tolerant systems and fault-tolerant systems in which some components are repairable and some are not.

Acknowledgments

The author wants to acknowledge Prof. Víctor Suñé for his help in the implementation of the methods. This research work has been supported in part by the Comisión Interministerial de Ciencia y Tecnología (CICYT) of the Ministry of Science and Technology of Spain under research grant TAP99-0443-C05-05.

Appendix

Lemma 1. *Let $W = \{W(t); t \geq 0\}$ be a CTMC with finite state space Ω and uniformly bounded reward rate structure r_i , $i \in \Omega$ with $|r_i| < \infty$, $i \in \Omega$ and let $W' = \{W'(t); t \geq 0\}$ be a rewarded CTMC differing from W only in that the transition rate of W from some state i to some state j has been increased by $\lambda > 0$. Let $\text{CRCD}(t, s) = P[\int_0^t r_{W(\tau)} d\tau > s]$ be the cumulative reward complementary distribution of W in the time interval $[0, t]$, $t > 0$ and let $\text{CRCD}'(t, s) = P[\int_0^t r_{W'(\tau)} d\tau > s]$ be the cumulative reward complementary distribution of W' in the time interval $[0, t]$, $t > 0$. Then $|\text{CRCD}'(t, s) - \text{CRCD}(t, s)| \leq 2\lambda t$.*

Proof. The proof is based on the formulation for $\text{CRCD}(t, s)$ used in the algorithm described in [16, 17]. Let α denote the initial probability distribution row vector of W , let $r^m > r^{m-1} > \dots > r^0$ be the $m + 1$ different reward rates of W , and assume⁴ $m \geq 1$. Let Ω be partitioned as $B_0 \cup B_1 \cup \dots \cup B_m$, where $B_j = \{i \in \Omega : r_i = r^j\}$. Let $\Lambda \geq \max_{i \in \Omega} \lambda_i + \lambda$, where λ_i , $i \in \Omega$ are the output rates of W , and let \mathbf{P} be the transition probability matrix of the randomized DTMC of W with randomization rate Λ : $\mathbf{P} = \mathbf{I} + \mathbf{A}/\Lambda$, where \mathbf{A} is the transition rate matrix (infinitesimal generator) of W and \mathbf{I} is an identity matrix. Let \mathbf{P}_{B_i, B_j} denote the subblock of \mathbf{P} including the elements with subindices $(l, m) \in B_i \times B_j$ (transition probabilities from states in B_i to states in B_j), given a (row or column) vector \mathbf{x} with subindices in Ω , let \mathbf{x}_{B_i} denote the restriction of \mathbf{x} to the subindices in B_i , and let $\mathbf{0}_{B_i}$ and $\mathbf{1}_{B_i}$ denote a vector of size $|B_i|$ with all its elements equal to, respectively, 0 and 1. The following formulation for $\text{CRCD}(t, s)$ has been obtained in [16, 17]:

$$\text{CRCD}(t, s) = \sum_{n=0}^{\infty} e^{-\Lambda t} \frac{(\Lambda t)^n}{n!} \sum_{k=0}^n \sum_{j=1}^m I_{r^{j-1}t \leq s < r^j t} \binom{n}{k} s_j^k (1 - s_j)^{n-k} b^{(j)}(n, k), \quad (1)$$

where I_c denotes the indicator function with value 1 when condition c is satisfied and value 0 otherwise,

$$s_j = \frac{s - r^{j-1}t}{(r^j - r^{j-1})t},$$

the coefficients $b^{(j)}(n, k)$ are given by

$$b^{(j)}(n, k) = \alpha \mathbf{b}^{(j)}(n, k),$$

⁴In the case $m = 0$, the cumulative reward is equal to $r^0 t$ with probability 1 and the result is trivial.

and the column vectors $\mathbf{b}_{B_l}^{(j)}(n, k)$ satisfy the following set of recursive expressions:

$$\mathbf{b}_{B_l}^{(1)}(n, 0) = \mathbf{1}_{B_l}, \quad 1 \leq l \leq m, \quad (2)$$

$$\mathbf{b}_{B_l}^{(m)}(n, n) = \mathbf{0}_{B_l}, \quad 0 \leq l < m, \quad (3)$$

$$\mathbf{b}_{B_l}^{(j)}(n, 0) = \mathbf{b}_{B_l}^{(j-1)}(n, n), \quad 1 < j \leq l \leq m, \quad (4)$$

$$\mathbf{b}_{B_l}^{(j)}(n, n) = \mathbf{b}_{B_l}^{(j+1)}(n, 0), \quad 0 \leq l < j < m, \quad (5)$$

$$\begin{aligned} \mathbf{b}_{B_l}^{(j)}(n, k) &= \frac{r^l - r^j}{r^l - r^{j-1}} \mathbf{b}_{B_l}^{(j)}(n, k-1) + \frac{r^j - r^{j-1}}{r^l - r^{j-1}} \sum_{i=0}^m \mathbf{P}_{B_l, B_i} \mathbf{b}_{B_i}^{(j)}(n-1, k-1), \\ & \quad j \leq l \leq m, \quad 1 \leq k \leq n, \end{aligned} \quad (6)$$

$$\begin{aligned} \mathbf{b}_{B_l}^{(j)}(n, k) &= \frac{r^{j-1} - r^l}{r^j - r^l} \mathbf{b}_{B_l}^{(j)}(n, k+1) + \frac{r^j - r^{j-1}}{r^j - r^l} \sum_{i=0}^m \mathbf{P}_{B_l, B_i} \mathbf{b}_{B_i}^{(j)}(n-1, k), \\ & \quad 0 \leq l < j, \quad 0 \leq k < n. \end{aligned} \quad (7)$$

It is also proved in [16, 17] that $\mathbf{0}_{B_l} \leq \mathbf{b}_{B_l}^{(j)}(n, k) \leq \mathbf{1}_{B_l}$, where the inequality among the vectors means that the inequality stands element by element.

In the following we will denote using a prime the quantities referred to the CTMC W' , i.e. \mathbf{P}' will denote the transition probability matrix of the randomized DTMC of W' with randomization rate Λ .

We will start by showing $\|\mathbf{b}'^{(j)}(n, k) - \mathbf{b}^{(j)}(n, k)\|_\infty \leq n(2\lambda)/\Lambda$, $1 \leq j \leq m$, $n \geq 0$, $0 \leq k \leq n$. The proof is by induction on n .

For $n = 0$, (2)–(5) imply $\mathbf{b}_{B_l}^{(j)}(0, 0) = \mathbf{1}_{B_l}$ for $1 \leq j \leq m$, $j \leq l \leq m$ and $\mathbf{b}_{B_l}^{(j)}(0, 0) = \mathbf{0}_{B_l}$ for $1 \leq j \leq m$, $0 \leq l < j$; similarly, we have $\mathbf{b}'_{B_l}{}^{(j)}(0, 0) = \mathbf{1}_{B_l}$ for $1 \leq j \leq m$, $j \leq l \leq m$ and $\mathbf{b}'_{B_l}{}^{(j)}(0, 0) = \mathbf{0}_{B_l}$ for $1 \leq j \leq m$, $0 \leq l < j$, implying $\|\mathbf{b}'^{(j)}(0, 0) - \mathbf{b}^{(j)}(0, 0)\|_\infty = 0$, $1 \leq j \leq m$. This shows the base case $n = 0$.

Consider now that the result holds for $n = \nu - 1$ and let us show that the result holds for $n = \nu$, $\nu \geq 1$. Since $\|\mathbf{b}'^{(j)}(\nu, k) - \mathbf{b}^{(j)}(\nu, k)\|_\infty = \max_{0 \leq l \leq m} \|\mathbf{b}'_{B_l}{}^{(j)}(\nu, k) - \mathbf{b}_{B_l}^{(j)}(\nu, k)\|_\infty$, it is enough to show: (1) $\|\mathbf{b}'_{B_l}{}^{(j)}(\nu, k) - \mathbf{b}_{B_l}^{(j)}(\nu, k)\|_\infty \leq \nu(2\lambda)/\Lambda$, $1 \leq j \leq m$, $0 \leq k \leq \nu$, $j \leq l \leq m$, and (2) $\|\mathbf{b}'_{B_l}{}^{(j)}(\nu, k) - \mathbf{b}_{B_l}^{(j)}(\nu, k)\|_\infty \leq \nu(2\lambda)/\Lambda$, $1 \leq j \leq m$, $0 \leq k \leq \nu$, $0 \leq l < j$. To that end we can use $\|\mathbf{b}'_{B_l}{}^{(j)}(\nu - 1, k) - \mathbf{b}_{B_l}^{(j)}(\nu - 1, k)\|_\infty \leq \|\mathbf{b}'^{(j)}(\nu - 1, k) - \mathbf{b}^{(j)}(\nu - 1, k)\|_\infty \leq (\nu - 1)(2\lambda)/\Lambda$, $1 \leq j \leq m$, $0 \leq k \leq \nu - 1$, $0 \leq l \leq m$. The proof of both (1) and (2) will be done by complete induction. For (1), the base case will be $j = 1$, $k = 0$ and the induction step will prove that the result for $1 \leq j \leq \iota - 1$, $0 \leq k \leq \nu$ and $j = \iota$, $0 \leq k \leq \kappa - 1$ implies the result for $j = \iota$ and $k = \kappa$, with $\iota = 1$, $0 < \kappa \leq \nu$ or $1 < \iota \leq m$, $0 \leq \kappa \leq \nu$. For (2), the base case will be $j = m$, $k = \nu$ and the induction step will prove that the result for $\iota + 1 \leq j \leq m$, $0 \leq k \leq \nu$ and $j = \iota$, $\kappa + 1 \leq k \leq \nu$ implies the result for $j = \iota$ and $k = \kappa$, with $\iota = m$, $0 \leq \kappa < \nu$ or $0 \leq \iota < m$,

$0 \leq \kappa \leq \nu$. Let s_f be the state “from” and let s_t be the state “to” of the transition whose rate has been increased by λ in W' . Let n_f and n_t such that $s_f \in B_{n_f}$ and $s_t \in B_{n_t}$.

Proof of (1): The base case ($j = 1, k = 0$) follows immediately from (2). For the induction step, we will consider two cases: (a) $1 < \iota \leq m, \kappa = 0$, (b) $1 \leq \iota \leq m, 0 < \kappa \leq \nu$. In case (a), from (4):

$$\mathbf{b}_{B_i}^{(\iota)}(\nu, 0) - \mathbf{b}_{B_i}^{(\iota)}(\nu, 0) = \mathbf{b}_{B_i}^{(\iota-1)}(\nu, \nu) - \mathbf{b}_{B_i}^{(\iota-1)}(\nu, \nu),$$

and using the (inner) induction hypothesis:

$$\|\mathbf{b}_{B_i}^{(\iota)}(\nu, 0) - \mathbf{b}_{B_i}^{(\iota)}(\nu, 0)\|_\infty = \|\mathbf{b}_{B_i}^{(\iota-1)}(\nu, \nu) - \mathbf{b}_{B_i}^{(\iota-1)}(\nu, \nu)\|_\infty \leq \nu(2\lambda)/\Lambda.$$

In case (b), we will consider three subcases: (b1) $l \neq n_f$, (b2) $l = n_f \neq n_t$, and (b3) $l = n_f = n_t$.

In case (b1), from (6), noting that $\mathbf{P}'_{B_l, B_i} = \mathbf{P}_{B_l, B_i}$ for all i :

$$\begin{aligned} \mathbf{b}_{B_i}^{(\iota)}(\nu, \kappa) - \mathbf{b}_{B_i}^{(\iota)}(\nu, \kappa) &= \frac{\gamma^\iota - \gamma^\kappa}{\gamma^\iota - \gamma^{\iota-1}} \left(\mathbf{b}_{B_i}^{(\iota)}(\nu, \kappa - 1) - \mathbf{b}_{B_i}^{(\iota)}(\nu, \kappa - 1) \right) \\ &\quad + \frac{\gamma^\iota - \gamma^{\iota-1}}{\gamma^\iota - \gamma^{\iota-1}} \sum_{i=0}^m \mathbf{P}_{B_i, B_i} \left(\mathbf{b}_{B_i}^{(\iota)}(\nu - 1, \kappa - 1) - \mathbf{b}_{B_i}^{(\iota)}(\nu - 1, \kappa - 1) \right), \end{aligned}$$

and using the induction hypotheses and the fact that the rows of the matrix \mathbf{P} add up 1:

$$\begin{aligned} \|\mathbf{b}_{B_i}^{(\iota)}(\nu, \kappa) - \mathbf{b}_{B_i}^{(\iota)}(\nu, \kappa)\|_\infty &\leq \frac{\gamma^\iota - \gamma^\kappa}{\gamma^\iota - \gamma^{\iota-1}} \|\mathbf{b}_{B_i}^{(\iota)}(\nu, \kappa - 1) - \mathbf{b}_{B_i}^{(\iota)}(\nu, \kappa - 1)\|_\infty \\ &\quad + \frac{\gamma^\iota - \gamma^{\iota-1}}{\gamma^\iota - \gamma^{\iota-1}} \|\mathbf{b}^{(\iota)}(\nu - 1, \kappa - 1) - \mathbf{b}^{(\iota)}(\nu - 1, \kappa - 1)\|_\infty \\ &\leq \frac{\gamma^\iota - \gamma^\kappa}{\gamma^\iota - \gamma^{\iota-1}} \nu \frac{2\lambda}{\Lambda} + \frac{\gamma^\iota - \gamma^{\iota-1}}{\gamma^\iota - \gamma^{\iota-1}} (\nu - 1) \frac{2\lambda}{\Lambda} < \nu \frac{2\lambda}{\Lambda}. \end{aligned}$$

In case (b2), using (6), noting that in that case the blocks \mathbf{P}'_{B_l, B_i} which are different from \mathbf{P}_{B_l, B_i} are the blocks $\mathbf{P}'_{B_{n_f}, B_{n_f}}$ and $\mathbf{P}'_{B_{n_f}, B_{n_t}}$:

$$\begin{aligned} \mathbf{b}_{B_{n_f}}^{(\iota)}(\nu, \kappa) - \mathbf{b}_{B_{n_f}}^{(\iota)}(\nu, \kappa) &= \frac{\gamma^{n_f} - \gamma^\kappa}{\gamma^{n_f} - \gamma^{\iota-1}} \left(\mathbf{b}_{B_{n_f}}^{(\iota)}(\nu, \kappa - 1) - \mathbf{b}_{B_{n_f}}^{(\iota)}(\nu, \kappa - 1) \right) \\ &\quad + \frac{\gamma^\iota - \gamma^{\iota-1}}{\gamma^{n_f} - \gamma^{\iota-1}} \sum_{i=0}^m \mathbf{P}'_{B_{n_f}, B_i} \left(\mathbf{b}_{B_i}^{(\iota)}(\nu - 1, \kappa - 1) - \mathbf{b}_{B_i}^{(\iota)}(\nu - 1, \kappa - 1) \right) \\ &\quad + \frac{\gamma^\iota - \gamma^{\iota-1}}{\gamma^{n_f} - \gamma^{\iota-1}} \left(\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}} \right) \mathbf{b}_{B_{n_f}}^{(\iota)}(\nu - 1, \kappa - 1) \\ &\quad + \frac{\gamma^\iota - \gamma^{\iota-1}}{\gamma^{n_f} - \gamma^{\iota-1}} \left(\mathbf{P}'_{B_{n_f}, B_{n_t}} - \mathbf{P}_{B_{n_f}, B_{n_t}} \right) \mathbf{b}_{B_{n_t}}^{(\iota)}(\nu - 1, \kappa - 1), \end{aligned}$$

and using the induction hypotheses, the fact that the rows of \mathbf{P}' add up 1, $\|\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}}\|_\infty = \lambda/\Lambda$, $\|\mathbf{P}'_{B_{n_f}, B_{n_t}} - \mathbf{P}_{B_{n_f}, B_{n_t}}\|_\infty = \lambda/\Lambda$, and $0 \leq \|\mathbf{b}_{B_i}^{(\iota)}(\nu - 1, \kappa - 1)\|_\infty \leq 1$:

$$\begin{aligned} \|\mathbf{b}_{B_{n_f}}^{(\iota)}(\nu, \kappa) - \mathbf{b}_{B_{n_f}}^{(\iota)}(\nu, \kappa)\|_\infty &\leq \frac{\gamma^{n_f} - \gamma^\kappa}{\gamma^{n_f} - \gamma^{\iota-1}} \|\mathbf{b}_{B_{n_f}}^{(\iota)}(\nu, \kappa - 1) - \mathbf{b}_{B_{n_f}}^{(\iota)}(\nu, \kappa - 1)\|_\infty \\ &\quad + \frac{\gamma^\iota - \gamma^{\iota-1}}{\gamma^{n_f} - \gamma^{\iota-1}} \|\mathbf{b}^{(\iota)}(\nu - 1, \kappa - 1) - \mathbf{b}^{(\iota)}(\nu - 1, \kappa - 1)\|_\infty \\ &\quad + \frac{\gamma^\iota - \gamma^{\iota-1}}{\gamma^{n_f} - \gamma^{\iota-1}} \|\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}}\|_\infty \|\mathbf{b}_{B_{n_f}}^{(\iota)}(\nu - 1, \kappa - 1)\|_\infty \end{aligned}$$

$$\begin{aligned}
& + \frac{r^\ell - r^{\ell-1}}{r^{n_f} - r^{\ell-1}} \|\mathbf{P}'_{B_{n_f}, B_{n_t}} - \mathbf{P}_{B_{n_f}, B_{n_t}}\|_\infty \|\mathbf{b}_{B_{n_t}}^{(\ell)}(\nu-1, \kappa-1)\|_\infty \\
\leq & \frac{r^{n_f} - r^\ell}{r^{n_f} - r^{\ell-1}} \nu \frac{2\lambda}{\Lambda} + \frac{r^\ell - r^{\ell-1}}{r^{n_f} - r^{\ell-1}} (\nu-1) \frac{2\lambda}{\Lambda} + \frac{r^\ell - r^{\ell-1}}{r^{n_f} - r^{\ell-1}} \frac{\lambda}{\Lambda} + \frac{r^\ell - r^{\ell-1}}{r^{n_f} - r^{\ell-1}} \frac{\lambda}{\Lambda} \\
= & \nu \frac{2\lambda}{\Lambda}.
\end{aligned}$$

In case (b3), using (6), noting that in that case the only block \mathbf{P}'_{B_l, B_i} which is different from \mathbf{P}_{B_l, B_i} is the block $\mathbf{P}'_{B_{n_f}, B_{n_f}}$:

$$\begin{aligned}
\mathbf{b}_{B_{n_f}}'^{(\ell)}(\nu, \kappa) - \mathbf{b}_{B_{n_f}}^{(\ell)}(\nu, \kappa) &= \frac{r^{n_f} - r^\ell}{r^{n_f} - r^{\ell-1}} \left(\mathbf{b}_{B_{n_f}}'^{(\ell)}(\nu, \kappa-1) - \mathbf{b}_{B_{n_f}}^{(\ell)}(\nu, \kappa-1) \right) \\
& + \frac{r^\ell - r^{\ell-1}}{r^{n_f} - r^{\ell-1}} \sum_{i=0}^m \mathbf{P}'_{B_{n_f}, B_i} \left(\mathbf{b}_{B_i}^{(\ell)}(\nu-1, \kappa-1) - \mathbf{b}_{B_i}^{(\ell)}(\nu-1, \kappa-1) \right) \\
& + \frac{r^\ell - r^{\ell-1}}{r^{n_f} - r^{\ell-1}} \left(\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}} \right) \mathbf{b}_{B_{n_f}}^{(\ell)}(\nu-1, \kappa-1),
\end{aligned}$$

and using the induction hypotheses, the fact that the rows of \mathbf{P}' add up 1, $\|\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}}\|_\infty = (2\lambda)/\Lambda$, and $0 \leq \|\mathbf{b}_{B_i}^{(\ell)}(\nu-1, \kappa-1)\|_\infty \leq 1$:

$$\begin{aligned}
\|\mathbf{b}_{B_{n_f}}'^{(\ell)}(\nu, \kappa) - \mathbf{b}_{B_{n_f}}^{(\ell)}(\nu, \kappa)\|_\infty &\leq \frac{r^{n_f} - r^\ell}{r^{n_f} - r^{\ell-1}} \|\mathbf{b}_{B_{n_f}}'^{(\ell)}(\nu, \kappa-1) - \mathbf{b}_{B_{n_f}}^{(\ell)}(\nu, \kappa-1)\|_\infty \\
& + \frac{r^\ell - r^{\ell-1}}{r^{n_f} - r^{\ell-1}} \|\mathbf{b}'^{(\ell)}(\nu-1, \kappa-1) - \mathbf{b}^{(\ell)}(\nu-1, \kappa-1)\|_\infty \\
& + \frac{r^\ell - r^{\ell-1}}{r^{n_f} - r^{\ell-1}} \|\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}}\|_\infty \|\mathbf{b}_{B_{n_f}}^{(\ell)}(\nu-1, \kappa-1)\|_\infty \\
\leq & \frac{r^{n_f} - r^\ell}{r^{n_f} - r^{\ell-1}} \nu \frac{2\lambda}{\Lambda} + \frac{r^\ell - r^{\ell-1}}{r^{n_f} - r^{\ell-1}} (\nu-1) \frac{2\lambda}{\Lambda} + \frac{r^\ell - r^{\ell-1}}{r^{n_f} - r^{\ell-1}} \frac{2\lambda}{\Lambda} \\
= & \nu \frac{2\lambda}{\Lambda}.
\end{aligned}$$

This completes the proof of (1).

Proof of (2): The base case ($j = m, k = \nu$) follows immediately from (3). For the induction step, we will consider two cases: (a) $1 \leq \ell < m, \kappa = \nu$, (b) $1 \leq \ell \leq m, 0 \leq \kappa < \nu$. In case (a), from (5):

$$\mathbf{b}_{B_l}^{(\ell)}(\nu, \nu) - \mathbf{b}_{B_l}^{(\ell)}(\nu, \nu) = \mathbf{b}_{B_l}'^{(\ell+1)}(\nu, 0) - \mathbf{b}_{B_l}^{(\ell+1)}(\nu, 0),$$

and using the (inner) induction hypothesis:

$$\|\mathbf{b}_{B_l}^{(\ell)}(\nu, \nu) - \mathbf{b}_{B_l}^{(\ell)}(\nu, \nu)\|_\infty = \|\mathbf{b}_{B_l}'^{(\ell+1)}(\nu, 0) - \mathbf{b}_{B_l}^{(\ell+1)}(\nu, 0)\|_\infty \leq \nu(2\lambda)/\Lambda.$$

In case (b), we will consider three subcases: (b1) $l \neq n_f$, (b2) $l = n_f \neq n_t$, and (b3) $l = n_f = n_t$.

In case (b1), from (7), noting that $\mathbf{P}'_{B_l, B_i} = \mathbf{P}_{B_l, B_i}$ for all i :

$$\begin{aligned}
\mathbf{b}_{B_l}'^{(\ell)}(\nu, \kappa) - \mathbf{b}_{B_l}^{(\ell)}(\nu, \kappa) &= \frac{r^{\ell-1} - r^\ell}{r^\ell - r^l} \left(\mathbf{b}_{B_l}'^{(\ell)}(\nu, \kappa+1) - \mathbf{b}_{B_l}^{(\ell)}(\nu, \kappa+1) \right) \\
& + \frac{r^\ell - r^{\ell-1}}{r^\ell - r^l} \sum_{i=0}^m \mathbf{P}_{B_l, B_i} \left(\mathbf{b}_{B_i}'^{(\ell)}(\nu-1, \kappa) - \mathbf{b}_{B_i}^{(\ell)}(\nu-1, \kappa) \right),
\end{aligned}$$

and using the induction hypotheses and the fact that the rows of the matrix \mathbf{P} add up 1:

$$\begin{aligned} \|\mathbf{b}'_{B_i}(\nu, \kappa) - \mathbf{b}^{(\iota)}_{B_i}(\nu, \kappa)\|_\infty &\leq \frac{r^{\iota-1} - r^\iota}{r^\iota - r^\iota} \|\mathbf{b}'_{B_i}(\nu, \kappa + 1) - \mathbf{b}^{(\iota)}_{B_i}(\nu, \kappa + 1)\|_\infty \\ &\quad + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^\iota} \|\mathbf{b}'_{B_i}(\nu - 1, \kappa) - \mathbf{b}^{(\iota)}_{B_i}(\nu - 1, \kappa)\|_\infty \\ &\leq \frac{r^{\iota-1} - r^\iota}{r^\iota - r^\iota} \nu \frac{2\lambda}{\Lambda} + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^\iota} (\nu - 1) \frac{2\lambda}{\Lambda} < \nu \frac{2\lambda}{\Lambda}. \end{aligned}$$

In case (b2), using (7), noting that in that case the blocks \mathbf{P}'_{B_l, B_i} which are different from \mathbf{P}_{B_l, B_i} are the blocks $\mathbf{P}'_{B_{n_f}, B_{n_f}}$ and $\mathbf{P}'_{B_{n_f}, B_{n_t}}$:

$$\begin{aligned} \mathbf{b}'_{B_{n_f}}(\nu, \kappa) - \mathbf{b}^{(\iota)}_{B_{n_f}}(\nu, \kappa) &= \frac{r^{\iota-1} - r^{n_f}}{r^\iota - r^{n_f}} \left(\mathbf{b}'_{B_{n_f}}(\nu, \kappa + 1) - \mathbf{b}^{(\iota)}_{B_{n_f}}(\nu, \kappa + 1) \right) \\ &\quad + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^{n_f}} \sum_{i=0}^m \mathbf{P}'_{B_{n_f}, B_i} \left(\mathbf{b}'_{B_i}(\nu - 1, \kappa) - \mathbf{b}^{(\iota)}_{B_i}(\nu - 1, \kappa) \right) \\ &\quad + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^{n_f}} \left(\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}} \right) \mathbf{b}^{(\iota)}_{B_{n_f}}(\nu - 1, \kappa) \\ &\quad + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^{n_f}} \left(\mathbf{P}'_{B_{n_f}, B_{n_t}} - \mathbf{P}_{B_{n_f}, B_{n_t}} \right) \mathbf{b}^{(\iota)}_{B_{n_t}}(\nu - 1, \kappa), \end{aligned}$$

and using the induction hypotheses, the fact that the rows of \mathbf{P}' add up 1, $\|\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}}\|_\infty = \lambda/\Lambda$, $\|\mathbf{P}'_{B_{n_f}, B_{n_t}} - \mathbf{P}_{B_{n_f}, B_{n_t}}\|_\infty = \lambda/\Lambda$, and $0 \leq \|\mathbf{b}^{(\iota)}_{B_i}(\nu - 1, \kappa)\|_\infty \leq 1$:

$$\begin{aligned} \|\mathbf{b}'_{B_{n_f}}(\nu, \kappa) - \mathbf{b}^{(\iota)}_{B_{n_f}}(\nu, \kappa)\|_\infty &\leq \frac{r^{\iota-1} - r^{n_f}}{r^\iota - r^{n_f}} \|\mathbf{b}'_{B_{n_f}}(\nu, \kappa + 1) - \mathbf{b}^{(\iota)}_{B_{n_f}}(\nu, \kappa + 1)\|_\infty \\ &\quad + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^{n_f}} \|\mathbf{b}'_{B_i}(\nu - 1, \kappa) - \mathbf{b}^{(\iota)}_{B_i}(\nu - 1, \kappa)\|_\infty \\ &\quad + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^{n_f}} \|\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}}\|_\infty \|\mathbf{b}^{(\iota)}_{B_{n_f}}(\nu - 1, \kappa)\|_\infty \\ &\quad + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^{n_f}} \|\mathbf{P}'_{B_{n_f}, B_{n_t}} - \mathbf{P}_{B_{n_f}, B_{n_t}}\|_\infty \|\mathbf{b}^{(\iota)}_{B_{n_t}}(\nu - 1, \kappa)\|_\infty \\ &\leq \frac{r^{\iota-1} - r^{n_f}}{r^\iota - r^{n_f}} \nu \frac{2\lambda}{\Lambda} + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^{n_f}} (\nu - 1) \frac{2\lambda}{\Lambda} + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^{n_f}} \frac{\lambda}{\Lambda} + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^{n_f}} \frac{\lambda}{\Lambda} \\ &= \nu \frac{2\lambda}{\Lambda}. \end{aligned}$$

In case (b3), using (7), noting that in that case the only block \mathbf{P}'_{B_l, B_i} which is different from \mathbf{P}_{B_l, B_i} is the block $\mathbf{P}'_{B_{n_f}, B_{n_f}}$:

$$\begin{aligned} \mathbf{b}'_{B_{n_f}}(\nu, \kappa) - \mathbf{b}^{(\iota)}_{B_{n_f}}(\nu, \kappa) &= \frac{r^{\iota-1} - r^{n_f}}{r^\iota - r^{n_f}} \left(\mathbf{b}'_{B_{n_f}}(\nu, \kappa + 1) - \mathbf{b}^{(\iota)}_{B_{n_f}}(\nu, \kappa + 1) \right) \\ &\quad + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^{n_f}} \sum_{i=0}^m \mathbf{P}'_{B_{n_f}, B_i} \left(\mathbf{b}'_{B_i}(\nu - 1, \kappa) - \mathbf{b}^{(\iota)}_{B_i}(\nu - 1, \kappa) \right) \\ &\quad + \frac{r^\iota - r^{\iota-1}}{r^\iota - r^{n_f}} \left(\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}} \right) \mathbf{b}^{(\iota)}_{B_{n_f}}(\nu - 1, \kappa), \end{aligned}$$

and using the induction hypotheses, the fact that the rows of \mathbf{P}' add up 1, $\|\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}}\|_\infty = (2\lambda)/\Lambda$, and $0 \leq \|\mathbf{b}^{(\iota)}_{B_i}(\nu - 1, \kappa)\|_\infty \leq 1$:

$$\|\mathbf{b}'_{B_{n_f}}(\nu, \kappa) - \mathbf{b}^{(\iota)}_{B_{n_f}}(\nu, \kappa)\|_\infty \leq \frac{r^{\iota-1} - r^{n_f}}{r^\iota - r^{n_f}} \|\mathbf{b}'_{B_{n_f}}(\nu, \kappa + 1) - \mathbf{b}^{(\iota)}_{B_{n_f}}(\nu, \kappa + 1)\|_\infty$$

$$\begin{aligned}
& + \frac{r^\ell - r^{\ell-1}}{r^\ell - r^{n_f}} \|\mathbf{b}'^{(\ell)}(\nu - 1, \kappa) - \mathbf{b}^{(\ell)}(\nu - 1, \kappa)\|_\infty \\
& + \frac{r^\ell - r^{\ell-1}}{r^\ell - r^{n_f}} \|\mathbf{P}'_{B_{n_f}, B_{n_f}} - \mathbf{P}_{B_{n_f}, B_{n_f}}\|_\infty \|\mathbf{b}_{B_{n_f}}^{(\ell)}(\nu - 1, \kappa)\|_\infty \\
\leq & \frac{r^{\ell-1} - r^{n_f}}{r^\ell - r^{n_f}} \nu \frac{2\lambda}{\Lambda} + \frac{r^\ell - r^{\ell-1}}{r^\ell - r^{n_f}} (\nu - 1) \frac{2\lambda}{\Lambda} + \frac{r^\ell - r^{\ell-1}}{r^\ell - r^{n_f}} \frac{2\lambda}{\Lambda} \\
= & \nu \frac{2\lambda}{\Lambda}.
\end{aligned}$$

This completes the proof of (2) and, therefore, the proof of $\|\mathbf{b}'^{(j)}(n, k) - \mathbf{b}^{(j)}(n, k)\|_\infty \leq n(2\lambda)/\Lambda$, $1 \leq j \leq m$, $n \geq 0$, $0 \leq k \leq n$.

But

$$b^{(j)}(n, k) - b^{(j)}(n, k) = \boldsymbol{\alpha} \mathbf{b}'^{(j)}(n, k) - \boldsymbol{\alpha} \mathbf{b}^{(j)}(n, k) = \boldsymbol{\alpha} \left(\mathbf{b}'^{(j)}(n, k) - \mathbf{b}^{(j)}(n, k) \right).$$

Then, using $\|\boldsymbol{\alpha}\|_\infty = 1$, we have

$$\begin{aligned}
|b^{(j)}(n, k) - b^{(j)}(n, k)| & \leq \|\boldsymbol{\alpha}\|_\infty \|\mathbf{b}'^{(j)}(n, k) - \mathbf{b}^{(j)}(n, k)\|_\infty \\
& = \|\mathbf{b}'^{(j)}(n, k) - \mathbf{b}^{(j)}(n, k)\|_\infty \leq n \frac{2\lambda}{\Lambda}, \quad 1 \leq j \leq m, n \geq 0, 0 \leq k \leq n.
\end{aligned}$$

Let l be the j such that $r^{j-1}t \leq s < r^j t$, we can write (1) as

$$\text{CRCD}(t, s) = \sum_{n=0}^{\infty} e^{-\Lambda t} \frac{(\Lambda t)^n}{n!} \sum_{k=0}^n \binom{n}{k} s_l^k (1 - s_l)^{n-k} b^{(l)}(n, k).$$

Then, noting that $0 \leq s_l < 1$:

$$\begin{aligned}
|\text{CRCD}'(t, s) - \text{CRCD}(t, s)| & \leq \sum_{n=0}^{\infty} e^{-\Lambda t} \frac{(\Lambda t)^n}{n!} \sum_{k=0}^n \binom{n}{k} s_l^k (1 - s_l)^{n-k} |b^{(l)}(n, k) - b^{(l)}(n, k)| \\
& \leq \sum_{n=0}^{\infty} e^{-\Lambda t} \frac{(\Lambda t)^n}{n!} \sum_{k=0}^n \binom{n}{k} s_l^k (1 - s_l)^{n-k} n \frac{2\lambda}{\Lambda} = \frac{2\lambda}{\Lambda} \sum_{n=0}^{\infty} n e^{-\Lambda t} \frac{(\Lambda t)^n}{n!} \\
& = \frac{2\lambda}{\Lambda} \sum_{n=1}^{\infty} e^{-\Lambda t} \frac{(\Lambda t)^n}{(n-1)!} = \frac{2\lambda}{\Lambda} \Lambda t \sum_{n=0}^{\infty} e^{-\Lambda t} \frac{(\Lambda t)^n}{n!} = 2\lambda t. \quad \square
\end{aligned}$$

Proof of Theorem 1. As theoretical background for measure theory and Lebesgue integration we use [10]. The characterization of the probability space underlying a discrete time Markov chain with denumerable state space is discussed in [5]. A recent short proof of the existence of arbitrary product probability measures can be found in [26]. Let $\lambda_i = \sum_{j \in \Omega - \{i\}} \lambda_{i,j}$ denote the output rate of W from state i . Let $\Pi = \{\Pi_n; n = 0, 1, 2, \dots\}$ be the embedded discrete-time Markov chain of W (see, for instance, [12]). Π has the same state space and initial probability distribution as W and transition probabilities $\psi_{i,j} = P[\Pi_{n+1} = j | \Pi_n = i] = \lambda_{i,j}/\lambda_i$, $j \neq i$, $\psi_{i,i} = P[\Pi_{n+1} = i | \Pi_n = i] = 0$ for the states i with $\lambda_i = \sum_{j \in \Omega - \{i\}} \lambda_{i,j} > 0$ and $\psi_{i,j} = P[\Pi_{n+1} = j | \Pi_n = i] = 0$, $j \neq i$, $\psi_{i,i} = P[\Pi_{n+1} = i | \Pi_n = i] = 1$ for the states i with $\lambda_i = 0$. The embedded DTMC of W' has same state space, initial probability distribution and transition probabilities as Π and, therefore, is probabilistically identical to Π . Both W and W' can be interpreted in terms of Π : Π gives the

sequence of states visited by W (W') and each state visit has a duration given by an independent holding time variable with exponential distribution with parameter equal to the output rate from the visited state.

We start by constructing a common probability space $(\mathcal{E}, \mathcal{A}, Q)$ in terms of which both W and W' can be defined. This is done by combining the probability space underlying Π with the probability space underlying a set of exponentially distributed independent random variables which will account (with scaling in the case of W') for the holding times. To simplify the proof, we will associate with absorbing states exponentially distributed holding times with finite parameter. Let $(\mathcal{E}_\Pi, \mathcal{A}_\Pi, Q_\Pi)$ be the probability space underlying Π : \mathcal{E}_Π is the set of infinite sequences $\pi = (s_0^\pi, s_1^\pi, \dots)$, $s_i^\pi \in \Omega$, \mathcal{A}_Π is the σ -algebra generated by the collection of subsets $\mathcal{E}_\Pi^{s_0, \dots, s_n} = \{\pi = (s_0^\pi, s_1^\pi, \dots) : s_0^\pi = s_0 \wedge s_1^\pi = s_1 \wedge \dots \wedge s_n^\pi = s_n\}$, $(s_0, \dots, s_n) \in \Omega^{n+1}$, $n = 0, 1, \dots$, and $Q_\Pi[\mathcal{E}_\Pi^{s_0, \dots, s_n}] = P[\Pi_0 = s_0] \psi_{s_0, s_1} \cdots \psi_{s_{n-1}, s_n}$. Let $H_{n,s}$, $n = 0, 1, \dots$, $s \in \Omega$ be independent exponential random variables with parameter Λ_s , where $\Lambda_s = \lambda_s$ if $\lambda_s > 0$ and $\Lambda_s = \Lambda^* > 0$ if $\lambda_s = 0$. For each random variable $H_{n,s}$, $n = 0, 1, \dots$, $s \in \Omega$ let $([0, \infty), \mathcal{B}_{[0, \infty)}, \mu_{n,s})$ be the underlying probability space: $\mathcal{B}_{[0, \infty)}$ is the Borel σ -algebra on $[0, \infty)$ and $\mu_{n,s}$ is the Borel probability measure defined by the distribution function of the random variable $H_{n,s}$. Let $(\mathcal{E}_H, \mathcal{A}_H, \mu)$ be the product of the probability spaces $([0, \infty), \mathcal{B}_{[0, \infty)}, \mu_{n,s})$, $n = 0, 1, \dots$, $s \in \Omega$. The probability space $(\mathcal{E}, \mathcal{A}, Q)$ is the product of the probability spaces $(\mathcal{E}_\Pi, \mathcal{A}_\Pi, Q_\Pi)$ and $(\mathcal{E}_H, \mathcal{A}_H, \mu)$. With respect to \mathcal{E}_H , given a $\omega \in \mathcal{E}_H$, $h_{n,s}(\omega)$ will denote the coordinate of ω equal to the realization of the random variable $H_{n,s}$. With respect to \mathcal{E} , given a $\omega \in \mathcal{E}$, $\pi(\omega) = (s_0^{\pi(\omega)}, s_1^{\pi(\omega)}, \dots)$ will denote the \mathcal{E}_Π coordinate of ω and $h_{n,s}(\omega)$ will denote the coordinate of ω equal to the realization of the random variable $H_{n,s}$, $n = 0, 1, \dots$, $s \in \Omega$.

The CTMC W can be defined in terms of $(\mathcal{E}, \mathcal{A}, Q)$ as follows. Each $\omega \in \mathcal{E}$ gives a realization, $W(\omega, t)$, of W :

$$\begin{aligned} W(\omega, t) &= s_0^{\pi(\omega)}, & 0 \leq t < h_{0, s_0^{\pi(\omega)}}(\omega), \\ W(\omega, t) &= s_1^{\pi(\omega)}, & h_{0, s_0^{\pi(\omega)}}(\omega) \leq t < h_{0, s_0^{\pi(\omega)}}(\omega) + h_{1, s_1^{\pi(\omega)}}(\omega), \\ &\vdots \\ W(\omega, t) &= s_m^{\pi(\omega)}, & \sum_{n=0}^{m-1} h_{n, s_n^{\pi(\omega)}}(\omega) \leq t < \sum_{n=0}^m h_{n, s_n^{\pi(\omega)}}(\omega), \\ &\vdots \end{aligned}$$

Let $L(\omega)$, $\omega \in \mathcal{E}$ be the random variable defined as $L(\omega) = \min\{l \geq 0 : \sum_{n=0}^l h_{n, s_n^{\pi(\omega)}}(\omega) > t\}$. It is well known (see, for instance, [12]) that, H_0, H_1, \dots being independent exponential random variables with parameters $\lambda_0, \lambda_1, \dots$ such that $\sup_{i \geq 0} \lambda_i < \infty$, $\lim_{n \rightarrow \infty} H_0 + H_1 + \dots + H_n = \infty$ with probability 1, implying that $\min\{n \geq 0 : H_0 + H_1 + \dots + H_n > t\}$ is defined with probability 1. Then:

$$Q[L \text{ is defined}] = Q \left[\left\{ \omega \in \mathcal{E} : \min \left\{ l \geq 0 : \sum_{n=0}^l h_{n, s_n^{\pi(\omega)}}(\omega) > t \right\} \text{ is defined} \right\} \right]$$

$$\begin{aligned}
&= \int_{\mathcal{E}_\Pi} \mu \left[\left\{ \omega' \in \mathcal{E}_H : \min \left\{ l \geq 0 : \sum_{n=0}^l h_{n, s_n^\pi}(\omega') > t \right\} \text{ is defined} \right\} \right] dQ_\Pi(\pi) \\
&= \int_{\mathcal{E}_\Pi} dQ_\Pi(\pi) = 1.
\end{aligned}$$

Let A be the subset of \mathcal{E} :

$$\begin{aligned}
A = \left\{ \omega \in \mathcal{E} : L(\omega) \text{ is defined} \wedge \sum_{n=0}^{L(\omega)-1} r_{s_n^{\pi(\omega)}} h_{n, s_n^{\pi(\omega)}}(\omega) \right. \\
\left. + r_{s_{L(\omega)}^{\pi(\omega)}} \left(t - \sum_{n=0}^{L(\omega)-1} h_{n, s_n^{\pi(\omega)}}(\omega) \right) > s \right\}.
\end{aligned}$$

Since A collects, except for a subset with probability 0, all realizations of W for which the cumulative reward in the time interval $[0, t]$ is $> s$,

$$\text{CRCD}(t, s) = Q[A].$$

Since, given $\beta > 0$ and being H an exponential random variable with parameter $\lambda > 0$, H/β is an exponential random variable with parameter $\beta\lambda$, the CTMC W' can be defined in terms of $(\mathcal{E}, \mathcal{A}, Q)$ as follows. Each $\omega \in \mathcal{E}$ gives a realization, $W'(\omega, t)$, of W' :

$$\begin{aligned}
W'(\omega, t) &= s_0^{\pi(\omega)}, \quad 0 \leq t < I_{s_0^{\pi(\omega)}=x} \frac{h_{0, s_0^{\pi(\omega)}}(\omega)}{\beta} + I_{s_0^{\pi(\omega)} \neq x} h_{0, s_0^{\pi(\omega)}}(\omega), \\
W'(\omega, t) &= s_1^{\pi(\omega)}, \quad I_{s_0^{\pi(\omega)}=x} \frac{h_{0, s_0^{\pi(\omega)}}(\omega)}{\beta} + I_{s_0^{\pi(\omega)} \neq x} h_{0, s_0^{\pi(\omega)}}(\omega) \\
&\leq t < I_{s_0^{\pi(\omega)}=x} \frac{h_{0, s_0^{\pi(\omega)}}(\omega)}{\beta} + I_{s_0^{\pi(\omega)} \neq x} h_{0, s_0^{\pi(\omega)}}(\omega) \\
&\quad + I_{s_1^{\pi(\omega)}=x} \frac{h_{1, s_1^{\pi(\omega)}}(\omega)}{\beta} + I_{s_1^{\pi(\omega)} \neq x} h_{1, s_1^{\pi(\omega)}}(\omega) \\
&\vdots \\
W'(\omega, t) &= s_m^{\pi(\omega)}, \quad \sum_{n=0}^{m-1} \left(I_{s_n^{\pi(\omega)}=x} \frac{h_{n, s_n^{\pi(\omega)}}(\omega)}{\beta} + I_{s_n^{\pi(\omega)} \neq x} h_{n, s_n^{\pi(\omega)}}(\omega) \right) \\
&\leq t < \sum_{n=0}^m \left(I_{s_n^{\pi(\omega)}=x} \frac{h_{n, s_n^{\pi(\omega)}}(\omega)}{\beta} + I_{s_n^{\pi(\omega)} \neq x} h_{n, s_n^{\pi(\omega)}}(\omega) \right). \\
&\vdots
\end{aligned}$$

Let $L'(\omega)$, $\omega \in \mathcal{E}$ be the random variable defined as $L'(\omega) = \min\{l \geq 0 : \sum_{n=0}^l \left(I_{s_n^{\pi(\omega)}=x} h_{n, s_n^{\pi(\omega)}}(\omega)/\beta + I_{s_n^{\pi(\omega)} \neq x} h_{n, s_n^{\pi(\omega)}}(\omega) \right) > t\}$. It can be proved that L' is defined with probability 1 as it was proved that L was defined with 1. Let A' be the subset of \mathcal{E} :

$$A' = \left\{ \omega \in \mathcal{E} : L'(\omega) \text{ is defined} \wedge \right.$$

$$\begin{aligned}
& \sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} r'_x \frac{h_{n,s_n^{\pi(\omega)}}(\omega)}{\beta} + I_{s_n^{\pi(\omega)} \neq x} r_{s_n^{\pi(\omega)}} h_{n,s_n^{\pi(\omega)}}(\omega) \right) \\
& + \left(I_{s_{L'(\omega)}^{\pi(\omega)}=x} r'_x + I_{s_{L'(\omega)}^{\pi(\omega)} \neq x} r_{s_{L'(\omega)}^{\pi(\omega)}} \right) \\
& \left(t - \sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} \frac{h_{n,s_n^{\pi(\omega)}}(\omega)}{\beta} + I_{s_n^{\pi(\omega)} \neq x} h_{n,s_n^{\pi(\omega)}}(\omega) \right) \right) > s \}.
\end{aligned}$$

Since A' collects, except for a subset with probability 0, all realizations of W' for which the cumulative reward in the time interval $[0, t]$ is $> s$,

$$\text{CRCD}'(t, s) = Q[A'].$$

To prove the theorem it suffices to show that $A \subset A'$. In that proof, we will use the shorthand $h_n^{\pi(\omega)}$ for $h_{n,s_n^{\pi(\omega)}}(\omega)$.

First note that, as $0 < \beta \leq 1$, $\sum_{n=0}^l \left(I_{s_n^{\pi(\omega)}=x} h_n^{\pi(\omega)} / \beta + I_{s_n^{\pi(\omega)} \neq x} h_n^{\pi(\omega)} \right) \geq \sum_{n=0}^l h_n^{\pi(\omega)}$, implying that $L'(\omega)$ is defined when $L(\omega)$ is and that, in that case, $L'(\omega) \leq L(\omega)$. Assuming $L(\omega)$ and $L'(\omega)$ defined, let

$$B(\omega) = \sum_{n=0}^{L(\omega)-1} r_{s_n^{\pi(\omega)}} h_n^{\pi(\omega)} + r_{s_{L(\omega)}^{\pi(\omega)}} \left(t - \sum_{n=0}^{L(\omega)-1} h_n^{\pi(\omega)} \right)$$

and let

$$\begin{aligned}
B'(\omega) &= \sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} r'_x \frac{h_n^{\pi(\omega)}}{\beta} + I_{s_n^{\pi(\omega)} \neq x} r_{s_n^{\pi(\omega)}} h_n^{\pi(\omega)} \right) \\
&+ \left(I_{s_{L'(\omega)}^{\pi(\omega)}=x} r'_x + I_{s_{L'(\omega)}^{\pi(\omega)} \neq x} r_{s_{L'(\omega)}^{\pi(\omega)}} \right) \left(t - \sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} \frac{h_n^{\pi(\omega)}}{\beta} + I_{s_n^{\pi(\omega)} \neq x} h_n^{\pi(\omega)} \right) \right).
\end{aligned}$$

It suffices to show $B'(\omega) \geq B(\omega)$. Since

$$B(\omega) = r_{\text{sup}} t - C(\omega)$$

with

$$C(\omega) = \sum_{n=0}^{L(\omega)-1} \left(r_{\text{sup}} - r_{s_n^{\pi(\omega)}} \right) h_n^{\pi(\omega)} + \left(r_{\text{sup}} - r_{s_{L(\omega)}^{\pi(\omega)}} \right) \left(t - \sum_{n=0}^{L(\omega)-1} h_n^{\pi(\omega)} \right) \quad (8)$$

and, it can be checked using $\beta = (r_{\text{sup}} - r'_x) / (r_{\text{sup}} - r_x)$ that

$$B'(\omega) = r_{\text{sup}} t - C'(\omega)$$

with

$$\begin{aligned}
C'(\omega) &= \sum_{n=0}^{L'(\omega)-1} \left(r_{\text{sup}} - r_{s_n^{\pi(\omega)}} \right) h_n^{\pi(\omega)} \\
&+ \left(I_{s_{L'(\omega)}^{\pi(\omega)}=x} (r_{\text{sup}} - r'_x) + I_{s_{L'(\omega)}^{\pi(\omega)} \neq x} (r_{\text{sup}} - r_{s_{L'(\omega)}^{\pi(\omega)}}) \right) \\
&\left(t - \sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} \frac{h_n^{\pi(\omega)}}{\beta} + I_{s_n^{\pi(\omega)} \neq x} h_n^{\pi(\omega)} \right) \right), \quad (9)
\end{aligned}$$

it suffices to show that, assuming $L(\omega)$ and $L'(\omega)$ defined and $L'(\omega) \leq L(\omega)$, $C'(\omega) \leq C(\omega)$. Two cases will be considered: (a) $L'(\omega) = L(\omega)$, and (b) $L'(\omega) < L(\omega)$.

In case (a), using (8), (9), $r_{\text{sup}} - r'_x \leq r_{\text{sup}} - r_x$, and $h_n^{\pi(\omega)}/\beta \geq h_n^{\pi(\omega)}$:

$$\begin{aligned}
C'(\omega) &= \sum_{n=0}^{L(\omega)-1} \left(r_{\text{sup}} - r_{s_n^{\pi(\omega)}} \right) h_n^{\pi(\omega)} \\
&\quad + \left(I_{s_{L(\omega)}^{\pi(\omega)}=x} (r_{\text{sup}} - r'_x) + I_{s_{L(\omega)}^{\pi(\omega)} \neq x} \left(r_{\text{sup}} - r_{s_{L(\omega)}^{\pi(\omega)}} \right) \right) \\
&\quad \left(t - \sum_{n=0}^{L(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} \frac{h_n^{\pi(\omega)}}{\beta} + I_{s_n^{\pi(\omega)} \neq x} h_n^{\pi(\omega)} \right) \right) \\
&\leq \sum_{n=0}^{L(\omega)-1} \left(r_{\text{sup}} - r_{s_n^{\pi(\omega)}} \right) h_n^{\pi(\omega)} + \left(r_{\text{sup}} - r_{s_{L(\omega)}^{\pi(\omega)}} \right) \left(t - \sum_{n=0}^{L(\omega)-1} h_n^{\pi(\omega)} \right) \\
&= C(\omega).
\end{aligned}$$

In case (b), assuming $s_{L'(\omega)}^{\pi(\omega)} \neq x$, $\sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} h_n^{\pi(\omega)}/\beta + I_{s_n^{\pi(\omega)} \neq x} h_n^{\pi(\omega)} \right) + h_{L'(\omega)}^{\pi(\omega)} > t$ implies $\left(t - \sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} h_n^{\pi(\omega)}/\beta + I_{s_n^{\pi(\omega)} \neq x} h_n^{\pi(\omega)} \right) \right) < h_{L'(\omega)}^{\pi(\omega)}$ and $\left(r_{\text{sup}} - r_{s_{L'(\omega)}^{\pi(\omega)}} \right) \left(t - \sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} h_n^{\pi(\omega)}/\beta + I_{s_n^{\pi(\omega)} \neq x} h_n^{\pi(\omega)} \right) \right) \leq \left(r_{\text{sup}} - r_{s_{L'(\omega)}^{\pi(\omega)}} \right) h_{L'(\omega)}^{\pi(\omega)}$. Assuming $s_{L'(\omega)}^{\pi(\omega)} = x$, $\sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} h_n^{\pi(\omega)}/\beta + I_{s_n^{\pi(\omega)} \neq x} h_n^{\pi(\omega)} \right) + h_{L'(\omega)}^{\pi(\omega)}/\beta > t$ implies $\left(t - \sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} h_n^{\pi(\omega)}/\beta + I_{s_n^{\pi(\omega)} \neq x} h_n^{\pi(\omega)} \right) \right) < h_{L'(\omega)}^{\pi(\omega)}/\beta$ and $(r_{\text{sup}} - r'_x) \left(t - \sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} h_n^{\pi(\omega)}/\beta + I_{s_n^{\pi(\omega)} \neq x} h_n^{\pi(\omega)} \right) \right) \leq (r_{\text{sup}} - r'_x) h_{L'(\omega)}^{\pi(\omega)}/\beta = (r_{\text{sup}} - r_x) h_{L'(\omega)}^{\pi(\omega)}$. Thus, we always have

$$\begin{aligned}
&\left(I_{s_{L'(\omega)}^{\pi(\omega)}=x} (r_{\text{sup}} - r'_x) + I_{s_{L'(\omega)}^{\pi(\omega)} \neq x} \left(r_{\text{sup}} - r_{s_{L'(\omega)}^{\pi(\omega)}} \right) \right) \\
&\left(t - \sum_{n=0}^{L'(\omega)-1} \left(I_{s_n^{\pi(\omega)}=x} \frac{h_n^{\pi(\omega)}}{\beta} + I_{s_n^{\pi(\omega)} \neq x} h_n^{\pi(\omega)} \right) \right) \\
&\leq \left(r_{\text{sup}} - r_{s_{L'(\omega)}^{\pi(\omega)}} \right) h_{L'(\omega)}^{\pi(\omega)}.
\end{aligned}$$

Using that result, (8) and (9), for the case (b),

$$\begin{aligned}
C'(\omega) &\leq \sum_{n=0}^{L'(\omega)-1} \left(r_{\text{sup}} - r_{s_n^{\pi(\omega)}} \right) h_n^{\pi(\omega)} + \left(r_{\text{sup}} - r_{s_{L'(\omega)}^{\pi(\omega)}} \right) h_{L'(\omega)}^{\pi(\omega)} \\
&\leq \sum_{n=0}^{L(\omega)-1} \left(r_{\text{sup}} - r_{s_n^{\pi(\omega)}} \right) h_n^{\pi(\omega)} \\
&\leq \sum_{n=0}^{L(\omega)-1} \left(r_{\text{sup}} - r_{s_n^{\pi(\omega)}} \right) h_n^{\pi(\omega)} + \left(r_{\text{sup}} - r_{s_{L(\omega)}^{\pi(\omega)}} \right) \left(t - \sum_{n=0}^{L(\omega)-1} h_n^{\pi(\omega)} \right) \\
&= C(\omega).
\end{aligned}$$

It remains to check that $E_n, E'_n \in \mathcal{A}$, $n = 0, 1, \dots$, where $E_n = \{\omega \in \mathcal{E} : L(\omega) = n\}$ and $E'_n = \{\omega \in \mathcal{E} : L'(\omega) = n\}$, and that $A, A' \in \mathcal{A}$.

We start by checking that $E_n \in \mathcal{A}$, $n = 0, 1, \dots$. Let $F_n = \{\omega \in \mathcal{E} : \sum_{m=0}^n h_{m, s_m^{\pi(\omega)}}(\omega) > t\}$, $n = 0, 1, \dots$. Since $E_0 = F_0$ and, for $n \geq 1$, $E_n = F_n \cap F_{n-1}^c$, it suffices to check that $F_n \in \mathcal{A}$, $n = 0, 1, \dots$. Let $F^{s_0, \dots, s_n} = \{\omega \in \mathcal{E}_H : \sum_{m=0}^n h_{m, s_m}(\omega) > t\}$. Since $F_n = \cup_{(s_0, \dots, s_n) \in \Omega^{n+1}} \mathcal{E}_{\Pi}^{s_0, \dots, s_n} \times F^{s_0, \dots, s_n}$, Ω^{n+1} is denumerable and $\mathcal{E}_{\Pi}^{s_0, \dots, s_n} \in \mathcal{A}_{\Pi}$, it suffices to check that $F^{s_0, \dots, s_n} \in \mathcal{A}_H$, $(s_0, \dots, s_n) \in \Omega^{n+1}$, $n = 0, 1, \dots$. This follows if $H_n = \{(h_0, \dots, h_n) \in [0, \infty)^{n+1} : \sum_{m=0}^n h_m > t\} \in \otimes_{m=0}^n \mathcal{B}_{[0, \infty)} = \mathcal{B}_{[0, \infty)^{n+1}}$, $n = 0, 1, \dots$, which can be proved by induction on n as follows. The case $n = 0$ is trivial since $H_0 = (t, \infty) \in \mathcal{B}_{[0, \infty)}$. Assume the result holds for $n = i \geq 0$. We have $H_{i+1} = G_{i+1} \cup \cup_{j=0}^{i+1} J_j$, where $G_{i+1} = \{(h_0, \dots, h_{i+1}) \in [0, \infty)^{i+2} : h_0 > 0 \wedge \dots \wedge h_{i+1} > 0 \wedge \sum_{m=0}^{i+1} h_m > t\}$ and $J_j = \{(h_0, \dots, h_{i+1}) \in [0, \infty)^{i+2} : h_j = 0 \wedge \sum_{\substack{m=0 \\ m \neq j}}^{i+1} h_m > t\}$. But $G_{i+1} \in \mathcal{B}_{[0, \infty)^{i+2}}$, since G_{i+1} is an open subset of $[0, \infty)^{i+2}$ and $J_j \in \mathcal{B}_{[0, \infty)^{i+2}} = \mathcal{B}_{[0, \infty)} \otimes \mathcal{B}_{[0, \infty)^{i+1}}$, since $\{0\} \in \mathcal{B}_{[0, \infty)}$ and, by the induction hypothesis, $H_i \in \mathcal{B}_{[0, \infty)^{i+1}}$. That $E'_n \in \mathcal{A}$, $n = 0, 1, \dots$ can be checked similarly, the only difference being that H_n has to be replaced by $H'_n = \{(h_0, \dots, h_n) \in [0, \infty)^{n+1} : \sum_{m=0}^n h_m / \alpha_m > t\}$, $0 < \alpha_m \leq 1$, which can be easily shown to belong to $\mathcal{B}_{[0, \infty)^{n+1}}$.

Let us check now that $A \in \mathcal{A}$. Let

$$A_n = \left\{ \omega \in \mathcal{E} : \sum_{m=0}^{n-1} r_{s_m^{\pi(\omega)}} h_{m, s_m^{\pi(\omega)}}(\omega) + r_{s_n^{\pi(\omega)}} \left(t - \sum_{m=0}^{n-1} h_{m, s_m^{\pi(\omega)}}(\omega) \right) > s \right\}.$$

Since $A = \cup_{n=0}^{\infty} (E_n \cap A_n)$, it suffices to check that $A_n \in \mathcal{A}$, $n = 0, 1, \dots$. Let

$$A^{s_0, \dots, s_n} = \left\{ \omega \in \mathcal{E}_H : \sum_{m=0}^{n-1} r_{s_m} h_{m, s_m}(\omega) + r_{s_n} \left(t - \sum_{m=0}^{n-1} h_{m, s_m}(\omega) \right) > s \right\}.$$

Since $A_n = \cup_{(s_0, \dots, s_n) \in \Omega^{n+1}} \mathcal{E}_{\Pi}^{s_0, \dots, s_n} \times A^{s_0, \dots, s_n}$, Ω^{n+1} is denumerable and $\mathcal{E}_{\Pi}^{s_0, \dots, s_n} \in \mathcal{A}_{\Pi}$, it suffices to check that $A^{s_0, \dots, s_n} \in \mathcal{A}_H$, $(s_0, \dots, s_n) \in \Omega^{n+1}$, $n = 0, 1, \dots$. Since A^{s_0, \dots, s_n} can be expressed as

$$A^{s_0, \dots, s_n} = \left\{ \omega \in \mathcal{E}_H : \sum_{m=0}^{n-1} (r_{s_m} - r_{s_n}) h_{m, s_m}(\omega) > s - r_{s_n} t \right\},$$

it suffices to check that $K_n(\gamma_0, \dots, \gamma_n, \delta) = \{(h_0, \dots, h_n) \in [0, \infty)^{n+1} : \sum_{m=0}^n \gamma_m h_m > \delta\} \in \otimes_{m=0}^n \mathcal{B}_{[0, \infty)} = \mathcal{B}_{[0, \infty)^{n+1}}$, $-\infty < \gamma_m < \infty$, $\gamma_m \neq 0$, $0 \leq m \leq n$, $-\infty < \delta < \infty$, $n = 0, 1, \dots$, which can be proved by induction on n as follows. The base case $n = 0$ is trivial, since $K_0(\gamma_0, \delta) = (\delta/\gamma_0, \infty)$ if $\gamma_0 > 0$ and $\delta \geq 0$, $K_0(\gamma_0, \delta) = [0, \infty)$ if $\gamma_0 > 0$ and $\delta < 0$, $K_0(\gamma_0, \delta) = \emptyset$ if $\gamma_0 < 0$ and $\delta \geq 0$, and $K_0(\gamma_0, \delta) = [0, \delta/\gamma_0)$ if $\gamma_0 < 0$ and $\delta < 0$ and $\emptyset, [0, \infty) \in \mathcal{B}_{[0, \infty)}$ and, for $a \geq 0$, $[0, a), (a, \infty) \in \mathcal{B}_{[0, \infty)}$. Assume the result holds for $n = i \geq 0$ and let us prove that the result holds for $n = i + 1$. We have $K_{i+1}(\gamma_0, \dots, \gamma_{i+1}, \delta) = M_{i+1}(\gamma_0, \dots, \gamma_{i+1}, \delta) \cup \cup_{j=0}^{i+1} N_j(\gamma_0, \dots, \gamma_{i+1}, \delta)$, where $M_{i+1}(\gamma_0, \dots, \gamma_{i+1}, \delta) = \{(h_0, \dots, h_{i+1}) \in [0, \infty)^{i+2} : h_0 > 0 \wedge \dots \wedge h_{i+1} > 0 \wedge \sum_{m=0}^{i+1} \gamma_m h_m > \delta\}$ and $N_j(\gamma_0, \dots, \gamma_{i+1}, \delta) = \{(h_0, \dots, h_{i+1}) \in [0, \infty)^{i+2} : h_j = 0 \wedge \sum_{\substack{m=0 \\ m \neq j}}^{i+1} \gamma_m h_m > \delta\}$. But $M_{i+1}(\gamma_0, \dots, \gamma_{i+1}, \delta) \in \mathcal{B}_{[0, \infty)^{i+2}}$, since $M_{i+1}(\gamma_0, \dots, \gamma_{i+1}, \delta)$

is an open subset of $[0, \infty)^{i+2}$ and $N_j(\gamma_0, \dots, \gamma_{i+1}, \delta) \in \mathcal{B}_{[0, \infty)^{i+2}} = \mathcal{B}_{[0, \infty)} \otimes \mathcal{B}_{[0, \infty)^{i+1}}$, since $\{0\} \in \mathcal{B}_{[0, \infty)}$ and, by the induction hypothesis, $K_i(\gamma_0, \dots, \gamma_{j-1}, \gamma_{j+1}, \dots, \gamma_{i+1}, \delta) \in \mathcal{B}_{[0, \infty)^{i+1}}$.

To prove that $A' \in \mathcal{A}$, let

$$A'_n = \left\{ \omega \in \mathcal{E} : \sum_{m=0}^{n-1} \left(I_{s_m(\omega)=x} r'_x \frac{h_{m, s_m(\omega)}(\omega)}{\beta} + I_{s_m(\omega) \neq x} r_{s_m(\omega)} h_{m, s_m(\omega)}(\omega) \right) + \left(I_{s_n(\omega)=x} r'_x + I_{s_n(\omega) \neq x} r_{s_n(\omega)} \right) \left(t - \sum_{m=0}^{n-1} \left(I_{s_m(\omega)=x} \frac{h_{m, s_m(\omega)}(\omega)}{\beta} + I_{s_m(\omega) \neq x} h_{m, s_m(\omega)}(\omega) \right) \right) > s \right\}.$$

Since $A' = \cup_{n=0}^{\infty} (E'_n \cap A'_n)$, it suffices to check that $A'_n \in \mathcal{A}$, $n = 0, 1, \dots$. Let

$$A'^{s_0, \dots, s_n} = \left\{ \omega \in \mathcal{E}_H : \sum_{m=0}^{n-1} \left(I_{s_m=x} r'_x \frac{h_{m, s_m}(\omega)}{\beta} + I_{s_m \neq x} r_{s_m} h_{m, s_m}(\omega) \right) + \left(I_{s_n=x} r'_x + I_{s_n \neq x} r_{s_n} \right) \left(t - \sum_{m=0}^{n-1} \left(I_{s_m=x} \frac{h_{m, s_m}(\omega)}{\beta} + I_{s_m \neq x} h_{m, s_m}(\omega) \right) \right) > s \right\}.$$

Since $A'_n = \cup_{(s_0, \dots, s_n) \in \Omega^{n+1}} \mathcal{E}_{\Pi}^{s_0, \dots, s_n} \times A'^{s_0, \dots, s_n}$, Ω^{n+1} is denumerable and $\mathcal{E}_{\Pi}^{s_0, \dots, s_n} \in \mathcal{A}_{\Pi}$, it suffices to check that $A'^{s_0, \dots, s_n} \in \mathcal{A}_H$, $(s_0, \dots, s_n) \in \Omega^{n+1}$, $n = 0, 1, \dots$. Since A'^{s_0, \dots, s_n} can be expressed as

$$A'^{s_0, \dots, s_n} = \left\{ \omega \in \mathcal{E}_H : \sum_{m=0}^{n-1} \left(I_{s_m=x \wedge s_n \neq x} (r'_x - r_{s_n}) \frac{h_{m, s_m}(\omega)}{\beta} + I_{s_m \neq x \wedge s_n=x} (r_{s_m} - r'_x) h_{m, s_m}(\omega) + I_{s_m \neq x \wedge s_n \neq x} (r_{s_m} - r_{s_n}) h_{m, s_m}(\omega) \right) > s - (I_{s_n=x} r'_x + I_{s_n \neq x} r_{s_n}) t \right\},$$

the result follows from $K_n(\gamma_0, \dots, \gamma_n, \delta) = \{(h_0, \dots, h_n) \in [0, \infty)^{n+1} : \sum_{m=0}^n \gamma_m h_m > \delta\} \in \otimes_{m=0}^n \mathcal{B}_{[0, \infty)} = \mathcal{B}_{[0, \infty)^{n+1}}$, $-\infty < \gamma_m < \infty$, $\gamma_m \neq 0$, $0 \leq m \leq n$, $-\infty < \delta < \infty$, $n = 0, 1, \dots$, which was proved previously. \square

References

- [1] C. Béounes, M. Aguéra, J. Arlat, S. Bachman, C. Bourdeu, J. E. Doucet, K. Kanoun, J. C. Laprie, S. Metge, J. Moreira de Souza, D. Powell and P. Spiesser, “SURF-2: A Program for Dependability Evaluation of Complex Hardware and Software Systems,” in *Proc. 23rd IEEE Int. Symp. on Fault-Tolerant Computing (FTCS-23)*, Toulouse, June 1993, pp. 142–150.
- [2] J. A. Carrasco, “Solving Large Interval Availability Models using a Model Transformation Approach,” *Computers and Operations Research*, vol. 31, no. 6, May 2004, pp. 1226–1235.

- [3] J. A. Carrasco, "An Efficient and Numerically Stable Method for Computing Bounds for the Interval Availability Distribution," Technical Report DMSD_2005_1, Departament d'Enginyeria Electrònica, Universitat Politècnica de Catalunya, July 2005, available at <ftp://ftp-ee1.upc.es/techreports>.
- [4] G. Chiola, G. Franceschinis, R. Gaeta, and M. Ribaud, "GreatSPN 1.7: Graphical Editor and Analyzer for Timed and Stochastic Petri Nets," *Performance Evaluation*, vol. 24, no. 1–2, November 1995, pp. 47–68.
- [5] K. L. Chung, *Markov Chains with Stationary Transition Probabilities*, Springer-Verlag, 1960.
- [6] G. Ciardo and K. Trivedi, "SPNP: The Stochastic Petri Net Package (version 3.1)," in *Proc. 1st IEEE Int. Workshop on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS'93)*, January 1993, pp. 390–391.
- [7] G. Ciardo and A. S. Miner, "SMART: Simulation and Markovian Analyzer for Reliability and Timing," in *Proc. 7th IEEE Int. Workshop on Petri Nets and Performance Models (PNPM'97)*, June 1997, pp. 41–43.
- [8] D. D. Deavours, G. Clark, T. Courtney, D. Daly, S. Derisavi, J. M. Doyle, W. H. Sanders, and P. G. Webster, "The Möbius Framework and Its Implementation," *IEEE Trans. on Software Engineering*, vol. 28, no. 10, October 2002, pp. 956–969.
- [9] L. Donatiello and V. Grassi, "On evaluating the cumulative performance distribution of fault-tolerant computer systems," *IEEE Trans. on Computers*, vol. 40, no. 11, November 1991, pp. 1301–1307.
- [10] G. B. Folland, *Real Analysis. Modern Techniques and Their Applications*, John Wiley & Sons, 2nd edition, 1999.
- [11] S. M. R. Islam and H. H. Ammar, "Performability of the Hypercube," *IEEE Trans. on Reliability*, vol. 38, no. 5, December 1989, pp. 518–526.
- [12] M. Kijima, *Markov Processes for Stochastic Modeling*, Chapman & Hall, 1997.
- [13] J. F. Meyer, D. G. Furchtgott, and L. T. Wu, "Performability evaluation of the SIFT computer," *IEEE Trans. on Computers*, vol. C-29, June 1980, pp. 501–509.
- [14] J. F. Meyer, "On Evaluating the Performability of Degradable Computing Systems," *IEEE Trans. on Computers*, vol. C-29, August 1980, pp. 720–731.
- [15] J. F. Meyer, "Performability of an Algorithm for Connection Admission Control," *IEEE Trans. on Computers*, vol. 50, no. 7, July 2001, pp. 724–733.
- [16] H. Nabli and B. Sericola, "Performability Analysis: A New Algorithm," *IEEE Trans. on Computers*, vol. 46, no. 4, April 1996, pp. 491–494.
- [17] H. Nabli and B. Sericola, "Performability Analysis of Fault-Tolerant Computer Systems," Technical Report 2254, 1994, INRIA, Rennes, France, available at <http://www.inria.fr/rrrt/index.en.html>.
- [18] H. Nabli, "Performability Measure for Acyclic Markovian Models," *Computers and Mathematics with Applications*, vol. 35, no. 8, April 1998, pp. 41–51.
- [19] H. Nabli and B. Sericola, "Performability Analysis for Degradable Computer Systems," *Computers and Mathematics with Applications*, vol. 39, no. 3–4, February 2000, pp. 217–234.

- [20] K. R. Pattipati, Y. Li and H. A. P. Blom, "A Unified Framework for the Performability Evaluation of Fault-Tolerant Computer Systems," *IEEE Trans. on Computers*, vol. 42, no. 3, March 1993, pp. 312–326.
- [21] M. A. Qureshi and W. H. Sanders, "Reward model solution methods with impulse and rate rewards: an algorithm and numerical results," *Performance Evaluation*, vol. 20, 1994, pp. 413–436.
- [22] M. A. Qureshi and W. H. Sanders, "A New Methodology for Calculating Distributions of Reward Accumulated during a Finite Interval," in *Proc. 26th IEEE Int. Symp. on Fault-Tolerant Computing*, 1996, pp. 116–125.
- [23] S. Rácz, A. Tari, and M. Telek, "MRMSolve: Distribution Estimation of Large Markov Reward Models," in *Proc. Tools Conf.*, 2002, pp. 71–81.
- [24] S. Ross, *Stochastic Processes*, John Wiley & Sons, New York, 1983.
- [25] G. Rubino and B. Sericola, "Interval Availability Analysis Using Denumerable Markov Processes: Application to Multiprocessor Subject to Breakdowns and Repair," *IEEE Trans. on Computers*, vol. 44, no. 2, February 1995, pp. 286–291.
- [26] S. Saeki, "A Proof of the Existence of Infinite Product Probability Measures," *American Mathematical Monthly*, vol. 103, no. 8, October 1996, pp. 682–683.
- [27] W. H. Sanders, W. D. Obal II, M. A. Qureshi, and F. K. Widjanarko, "The UltraSAN modeling environment," *Performance Evaluation*, vol. 24, 1995, pp. 89–115.
- [28] R. Smith, K. S. Trivedi, and A. V. Ramesh, "Performability analysis: Measures, an algorithm, and a case study," *IEEE Trans. on Computers*, vol. 37, no. 4, April 1988, pp. 406–417.
- [29] E. de Souza e Silva and H. R. Gail, "Calculating Availability and Performability Measures of Repairable Computer Systems using Randomization," *Journal of the ACM*, vol. 36, no. 4, January 1989, pp. 171–193.
- [30] E. de Souza e Silva and H. R. Gail, "An Algorithm to Calculate Transient Distributions of Cumulative Rate and Impulse based Reward," *Communications in Statistics—Stochastic Models*, vol. 14, no. 3, 1998, pp. 509–536.
- [31] M. Telek and S. Rácz, "Numerical Analysis of Large Markov Reward Models," *Performance Evaluation*, no. 36–37, 1999, pp. 95–114.