Approximation of First Passage Time Distributions of Compositions of Independent Markov Chains

András Horváth[1], Marco Paolieri[2], and Enrico Vicario[3]

[1] Department of Computer Science, University of Turin, Italy
horvath@di.unito.it
[2] Department of Computer Science, University of Southern California, USA
paolieri@usc.edu
[3] Department of Information Engineering, University of Florence
enrico.vicario@unifi.it

Abstract. To improve performance or reliability, systems frequently include multiple components that operate in parallel or with limited interaction, e.g., replicated components for triple modular redundancy. We consider components modeled by independent and possibly different continuous-time Markov chains and propose an approach to estimate the distribution of first passage times for a combination of component states (e.g., a system state where all components have failed) without generating the joint state space of the underlying Markov chain nor evaluating probabilities for each of its states. Our results highlight that the approach leads to accurate approximations with significant reductions of computational complexity.

Keywords: First Passage · Bounded Reachability · CTMC · Markov Chain · Replicated · Modular Redundancy

1 Introduction

Continuous-time Markov chains (CTMCs) are a class of stochastic processes that has found broad application in models of system performance and reliability. Many high-level modeling formalisms such as stochastic Petri nets [17], stochastic process algebras [9], and queuing networks [15] define CTMC processes that can then be analyzed using dedicated tools [1,12,16] to evaluate transient or steady-state metrics.

To improve performance or reliability, systems frequently include replicated components operating in parallel or with limited interaction, e.g., components replicated for triple modular redundancy (TMR). When the system includes many replicated components, each with a state evolving over time, the large number of states of the resulting CTMC process presents major challenges due to memory and computation requirements. These issues are exacerbated by the necessity of using phase-type (PH) distributions [18] to model activities with non-exponential durations. PH distributions can be introduced in the model as a sequence of

intermediate states (phases) for a non-exponential activity, where rates between states are selected to fit the original distribution (by matching moments [4], tail behavior [11], or both [13], or by maximizing likelihood [3]). While the use of additional intermediate states allows more accurate approximations, it also increases the number of system states, especially when many such activities can execute in parallel.

Several approaches were proposed to analyze large CTMCs with multiple components. These include *binary decision diagrams* (BDDs) [16], which represent the CTMC transition matrix as a directed graph where paths select rate values, and *structured analysis* [6], which represents the CTMC transition matrix as a sum of Kronecker products of the small transition matrices of individual components. Either approach can be used to evaluate transient or steady-state probabilities of each system state through iterative methods.

In this work, we propose an *approximate solution method* to evaluate the cumulative distribution function (cdf) of the first passage time of a combination of component states (e.g., a system state where all components have failed) *without enumerating system states* (in contrast with alternative approaches computing probabilities for each system state). The approximation will be developed by time discretization assuming that the system can be in the concerned combination of component states at the end of a small time interval only if all but one of the components are already there at the beginning of the time interval. The state distribution of the components at the beginning of each interval is assumed to be independent, irrespective of first passage events in previous time intervals. Time discretization will be removed then leading to a differential equation whose numerical integration provides the approximation.

In contrast with classical work on replicated components [2,5,19], we consider possibly different CTMC components with an arbitrary number of states (instead of one up and one down state) and evaluate the cdf (instead of the moments) of the first passage time.

The paper is organized as follows. We recall background and define our problem in Section 2, and we present our approximation method in Section 3. In Section 4, we provide numerical results and discuss some implementation issues, drawing our conclusions in Section 5.

2 Background and Problem Definition

We consider a system composed of n independent components modeled as CTMCs $\{Y_k(t)\}_{1\leq k\leq n}$ with finite state spaces S_k and infinitesimal generators $Q_k = (q_{kij})_{i,j\in S_k}$ with $k = 1, ..., n$. We denote by $F_k \subseteq S_k$ the set of failure states of component k. In general, the components can be repaired (i.e., the states in F_k are not absorbing). The transient probabilities of component k are denoted by

$$p_{kij}(t) = P(Y_k(t) = j \mid Y_k(0) = i) \quad \text{for} \quad i, j \in S_k, k = 1, ..., n,$$

which can be calculated, for example, by uniformization.

The system can be modeled as a CTMC $\{X(t)\}$ with state space denoted by $S = S_1 \times S_2 \times ... \times S_n$, where $X(t) = (Y_1(t), \ldots, Y_n(t))$. The set of states in which all components are in a failure state is denoted by $F = F_1 \times F_2 \times ... \times F_n \subseteq S$.

Let T be the first time when all the components are in a failure state, that is,

$$T = \min \{t \geq 0 \mid X(t) \in F\} \ ,$$

and let $F_T(t)$ indicate its cdf, that is, $F_T(t) = P(T \leq t)$. The cdf $F_T(t)$ is known as *first passage time distribution* to reach a state in F. It is also referred to as *time-bounded reachability* in probabilistic model checking.

Since the components are independent, the transient probabilities of the composed system can easily be computed based on the transient probabilities of the components in product form as

$$P(X(t) = (y_1, ..., y_n) \mid X(0) = (s_1, ..., s_n)) = \prod_{k=1}^{n} p_{k s_k y_k}(t) \, . \tag{1}$$

where $(s_1, ..., s_n)$ is the initial state. Accordingly, also $P(X(t) \in F)$, that is, the probability that all components are in a failure state at a given time t (which is not equal to $F_T(t)$), can be computed in product form considering the components in isolation.

On the contrary, $F_T(t)$, i.e., the first passage time cdf to reach a state in F, cannot be obtained in product form multiplying first passage time cdfs of the components. Indeed, the product of first passage time cdfs of the components yields the probability that all components have been in a failure state before time t at least once (and not the probability that all components have been in a failure state at the same time).

In order to determine $F_T(t)$ exactly, one has to consider a variant of $\{X(t)\}$, that we denote by $\{X'(t)\}$, in which states in F are made absorbing. In this modified CTMC we have

$$F_T(t) = P(X'(t) \in F \mid X'(0) = (s_1, ..., s_n)) \, .$$

Making states in F absorbing couples the behavior of the components and they are not independent anymore in a probabilistic sense. Consequently, the transient probabilities of $\{X'(t)\}$ are not in product form and we need to consider a CTMC with $|S| - |F|$ states (it is not necessary to represent states in F during the calculations), with exponential growth of the state space with respect to the number of components, quickly making the analysis unfeasible.

The aim of this paper is to propose an approximation of $F_T(t)$, denoted by $\hat{F}_T(t)$, that is based on the individual behavior of the components and hence does not require analyzing $\{X'(t)\}$. By doing so, the computational complexity is kept linear with respect to the number of components.

3 Approximation Method

In order to provide the stochastic interpretation of the proposed approximation method, we give first a description in which time is discretized. This discretized

version proceeds in time by taking steps of length δ and calculating $\hat{F}_T(i\delta), i = 0, 1, 2, \ldots$, the approximation of $F_T(i\delta)$ for $i = 0, 1, 2, \ldots$. We assume that δ is such that there is negligible probability that more than one component makes a transition in an interval of length δ.

The assumption on δ implies that the system enters a state in F in $(t, t+\delta]$ only if at time t the number of failed components is $n-1$. For this reason, at each step we calculate the probability that in $\{X(t)\}$ at time t all components other than component k are failed by

$$U_k(t) = \prod_{1 \le i \le n, i \neq k} P(Y_i(t) \in F_i)\,. \tag{2}$$

The intensity with which component k moves from up states to down states at time t can be calculated as

$$D_k(t) = \sum_{i \notin F_k} \left(P(Y_k(t) = i) \sum_{j \in F_k} q_{kij} \right). \tag{3}$$

The probability itself that component k is up at time t and makes a transition from an up state to a down state in $(t, t+\delta]$ can be approximated by $D_k(t)\delta$.

In order to easily consider all components together, we introduce also

$$G(t) = \sum_{i=1}^{n} U_i(t) D_i(t)\,. \tag{4}$$

The approximation starts with $\hat{F}_T(0) = F_T(0) = P(X(0) \in F)$, which can be easily calculated given the initial distribution of the components, and proceeds according to

$$\hat{F}_T((i+1)\delta) = \hat{F}_T(i\delta) + (1 - \hat{F}_T(i\delta)) G(i\delta)\delta\,, \tag{5}$$

where we multiply by $(1 - \hat{F}_T(i\delta))$ to consider the probability that a state in F was already reached.

The approximation in Eq. (5) relies on the fact that in $\{X'(t)\}$ the components evolve independently up to the moment in which the system reaches a state in F. This suggests that the product-form probabilities in Eq. (1) multiplied by the probability that a state in F has not been reached give a good approximation of the probabilities of the states in $S \setminus F$, that is,

$$P(X'(t) = (y_1, ..., y_n)) \approx (1 - F_T(t)) \prod_{i=1}^{n} P(Y_i(t) = y_i) \quad \text{for} \quad (y_1, ..., y_n) \notin F\,,$$

where in our numerical scheme $F_T(t)$ is substituted by $\hat{F}_T(t)$. For the states in F we have $P(X'(t) \in F) = F_T(t) \approx \hat{F}_T(t)$.

To remove time discretization, note that Eq. (5) can be reorganized as

$$\frac{\hat{F}_T((i+1)\delta) - \hat{F}_T(i\delta)}{\delta} = (1 - \hat{F}_T(i\delta)) G(i\delta)\,,$$

and by taking the limit $\delta \to 0$, we obtain the differential equation

$$\hat{F}'_T(t) = (1 - \hat{F}_T(t))G(t) \tag{6}$$

with initial condition $\hat{F}_T(0) = F_T(0)$. The approximation $\hat{F}(t)$ can then be calculated by numerical integration of Eq. (6).

Note that if the components are identical and share also the initial distribution, then

$$U(t) := U_1(t) = ... = U_n(t) = P(Y_1(t) \in F_1)^{n-1},$$

$$D(t) := D_1(t) = ... = D_n(t) = \sum_{i \notin F_1} \left(P(Y_1(t) = i) \sum_{j \in F_1} q_{1ij} \right),$$

and Eq. (4) simplifies to $G(t) = nU(t)D(t)$.

4 Numerical Experiments and Implementation Issues

We present two sets of experiments. First, we use a model composed of identical components described by a CTMC with a small state space. This case allows us to compare the results obtained by the proposed approximation method against exact results. Next, a model composed of different CTMCs with large state spaces is analyzed. In this case we compare the approximation against simulation.

After the experiments, we briefly discuss implementation issues and provide execution times.

4.1 Identical Components with Small State Space

The infinitesimal generator of the components is

$$Q = \begin{pmatrix} -\alpha & \alpha & 0 & 0 \\ \alpha & -2\alpha - \beta & \alpha & \beta \\ 0 & \alpha & -\alpha - \beta & \beta \\ \gamma & 0 & 0 & -\gamma \end{pmatrix},$$

which describes a four state system in which (i) the first three states correspond to normal operation states among which there are transitions with intensity α, (ii) from the second and third up states, the down state can be reached by a transition with intensity β, and (iii) repair takes to the first state with intensity γ.

The parameters (α, β, γ) allow us to calibrate the steady-state probability of being in the failure state in a component considered in isolation which, as we will see, has an impact on the accuracy of the approximations. The following sets of parameters (α, β, γ) will be used: $(1, 1, 12)$ (Case 1), $(2, 2, 3)$ (Case 2), $(4, 4, 0.75)$ (Case 3). Steady-state probabilities of the failure state are $\frac{1}{33}$, $\frac{1}{5}$, $\frac{2}{3}$, respectively.

We calculated exact and approximate first passage distributions, i.e., $F_T(t)$ and $\hat{F}_T(t)$, for several values of n in the above cases. The initial state s is the state in which all components are in the first state.

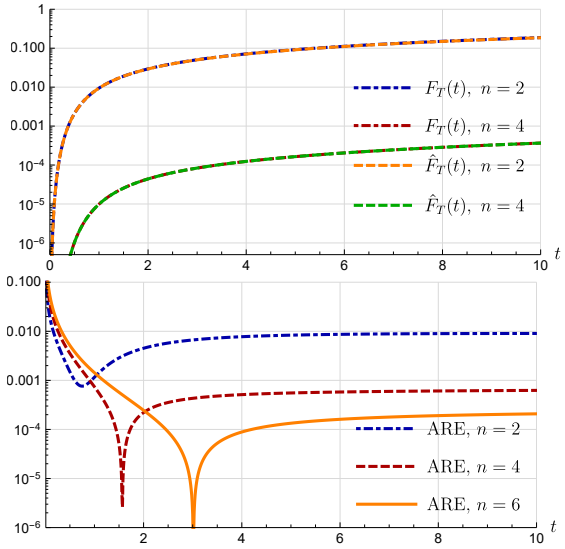

Fig. 1. Case 1: Exact $F_T(t)$ and approximate $\hat{F}_T(t)$ (top); approximation ARE (bottom).

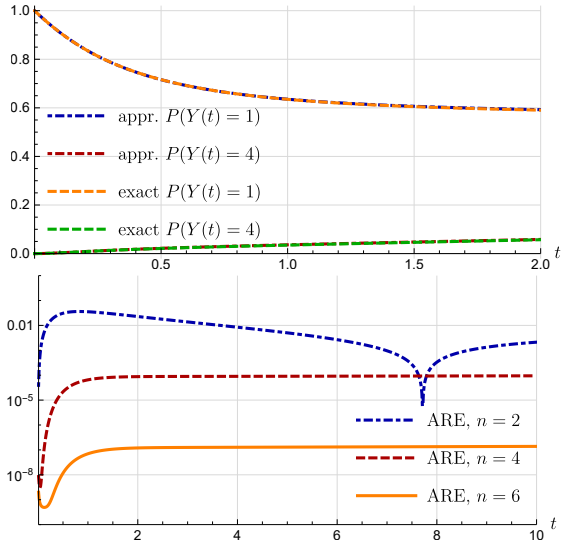

Fig. 2. Case 1: Approximate and exact state probabilities with $n = 2$ (top); ARE of the probability of the fourth state (i.e., the failure state) for different values of n (bottom).

Fig. 1 shows $F_T(t)$ and $\hat{F}_T(t)$ and the absolute relative error (ARE) of $\hat{F}_T(t)$. Visually, the approximation cannot be distinguished from the exact values. The ARE shows that the approximation error is low and decreases as we increase the number of components. The downward spikes in the ARE are due to the points where exact and approximate curves cross each other leading to a point

where the ARE is zero. In order to investigate the source of the approximation error of $\hat{F}_T(t)$, we calculated state probabilities of the components based on the approximation and exactly. In Fig. 2 there is no visible error in the state probabilities with $n = 2$ and the ARE of the probability of the fourth state 4 decreases as n is increased. (Note that since the initial state is the same for all components, they have the same transient probabilities; hence we use $Y(t)$ without specifying the component.)

In Fig. 3, we observe that for Case 3, where the failure state is reached with higher probability, we obtain greater absolute error and ARE of the approximation. Similarly to Fig. 1, the error decreases as we increase the number of components. In Fig. 4, we can see that (as expected) the errors in the state probabilities and ARE of the failure state are also higher.

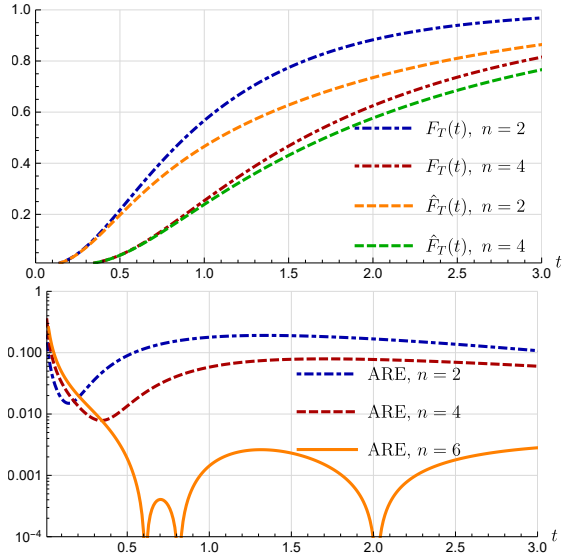

Fig. 3. Case 3: Exact $F_T(t)$ and approximate $\hat{F}_T(t)$ (top); approximation ARE (bottom).

As a comparison of the three cases, in Fig. 5, we depicted for $n = 4$ the ARE of $\hat{F}_T(t)$ and the ARE of the probability of the fourth state. In the range $t \in [0, 2]$, as expected, the higher the probability of the failure state, the worse approximation we obtain. For the third considered case, after $t = 2$ the approximation gets better. This is due to the fact that $F_T(t)$ is getting closer and closer to one making it easier to obtain good ARE values. In the other two cases, we have $F_T(10) < 0.2$ and the approximation error remains stable.

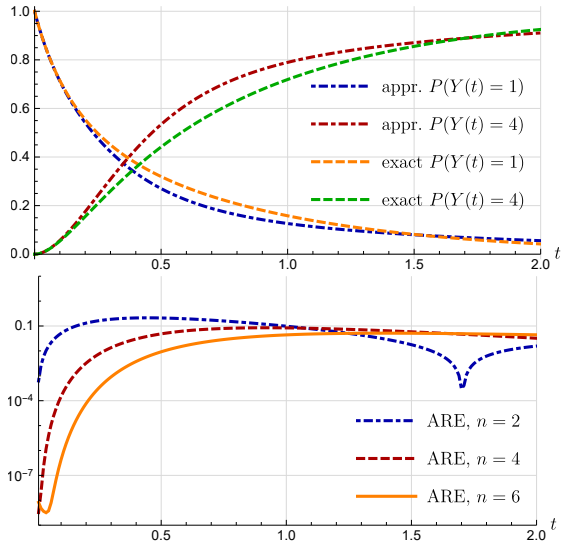

Fig. 4. Case 3: Approximate and exact state probabilities with $n = 2$ (top); ARE of the probability of the fourth state (i.e., the failure state) for different values of n (bottom).

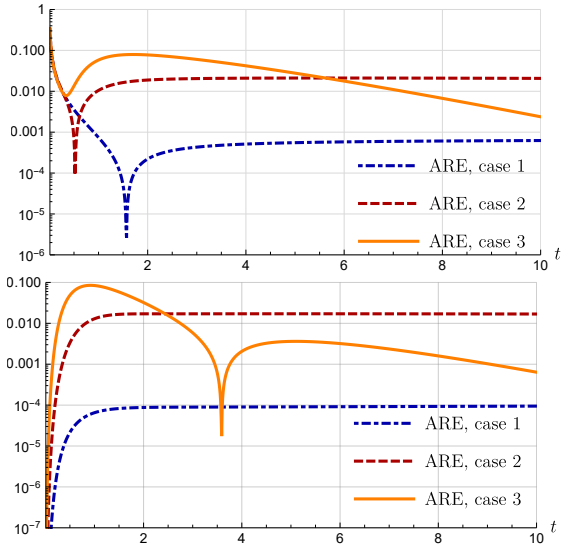

Fig. 5. ARE of $\hat{F}_T(t)$ (top) and ARE of the probability of the fourth state (bottom) with $n = 4$ for all three considered cases.

4.2 Different Components with Large State Spaces

In order to describe the components, we use the Petri net (PN) depicted in Fig. 6 which models a rejuvenation mechanism [8]. The system is in one of four states:

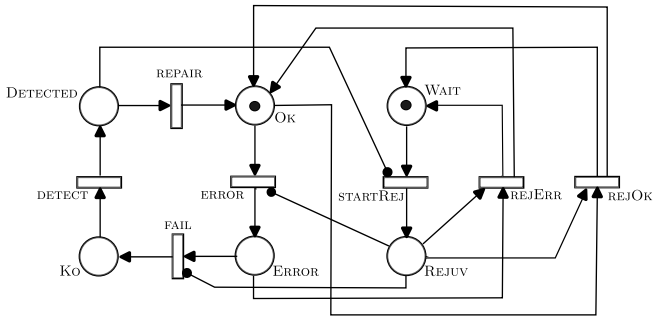

Fig. 6. PN model of a rejuvenation mechanism.

OK is a safe operational state; ERROR corresponds to an aged operational state that may lead to a failure; in state KO the system is down due to a failure; the fact that a failure has been detected is modeled by state DETECTED. Transitions among these four states are modeled by PN transitions ERROR, FAIL, DETECT, and REPAIR. The system is complemented by a rejuvenation mechanism with a timer that controls when the next rejuvenation takes place. The rejuvenation mechanism is either in state WAIT, where it waits for the timer to run out, or in state REJUV, where rejuvenation is carried out. The transition from WAIT to REJUV is called STARTREJ. Rejuvenation can be completed while the system is in state OK or ERROR and it takes the system back to its initial state through either transition REJOK or REJERR, respectively. In state KO rejuvenation is not possible but the timer is stopped only when a failure is detected. When rejuvenation is in progress, the system cannot degrade, that is, transitions ERROR and FAIL are inhibited when place REJUV has a token.

The time to fire distribution of the transitions will be defined through PH distributions [18]. An order r PH distribution is given through the time to absorption in a CTMC with r transient states, called phases, and one absorbing state. Accordingly, it is determined by the initial probability vector, denoted by a, and the infinitesimal generator of its CTMC, denoted by A. We will use two subclasses of PH distributions. The first one is the family of Erlang distributions. In terms of PH distributions, an Erlang distribution with shape parameter[4] r and mean equal to m is obtained by

$$a = (1\ 0\ ...\ 0),\ A = \begin{pmatrix} -\frac{r}{m} & \frac{r}{m} & & & \\ & -\frac{r}{m} & \frac{r}{m} & & \\ & & \ddots & \ddots & \\ & & & -\frac{r}{m} & \frac{r}{m} \\ & & & & -\frac{r}{m} \end{pmatrix},$$

[4] The shape parameter is equal to the number of phases, that is, the number of transient states.

where only the parameters of the transient states are given in a and A (the others can be deduced from these). An Erlang distribution with r phases and mean equal to m will be denoted by $\mathrm{Erl}(r, m)$. The second subclass is a mixture of Erlang distributions with common intensity and uniform mixing probability. In terms of PH distributions, this subclass has

$$a = \left(\underbrace{\frac{1}{k} \; \frac{1}{k} \; \ldots \; \frac{1}{k}}_{k} \; \underbrace{0 \; \ldots \; 0}_{r-k} \right), \; A = \begin{pmatrix} -\lambda & \lambda & & & \\ & -\lambda & \lambda & & \\ & & \ddots & \ddots & \\ & & & -\lambda & \lambda \\ & & & & -\lambda \end{pmatrix},$$

which are determined by three parameters: the total number of phases, r, the number of phases with non-zero initial probability, k, and the intensities in A denoted by λ. We will refer to this family of distributions as Erlang mixture (EM); an EM distribution will be denoted by $\mathrm{EM}(r, k, \lambda)$. In Fig. 7 we show a few examples of Erlang and EM distributions through their probability density functions (pdf).

When time to fire distributions in a PN are PH distributions, the underlying stochastic process is a CTMC whose infinitesimal generator can be built by Kronecker operations (see, e.g., [7]). The CTMC is subject to the so-called state space explosion problem. This is due to the fact that, given a marking, the number of states corresponding to the marking in the underlying CTMC is equal to the product of the number of phases of the PH distributions of the enabled transitions in the marking.

As Case 1, we consider a system composed of three components. Their PH distributions together with the number of states of the resulting CTMC are reported in Table 1. Note that the components are not identical. Every state of the CTMC corresponding to a marking in which there is a token in place KO or DETECTED is considered as a failure state. These markings are (KO,WAIT), (KO,REJUV), (DETECTED,WAIT), and (DETECTED,REJUV). For example, for what concerns the second component, the above four markings correspond to 600,

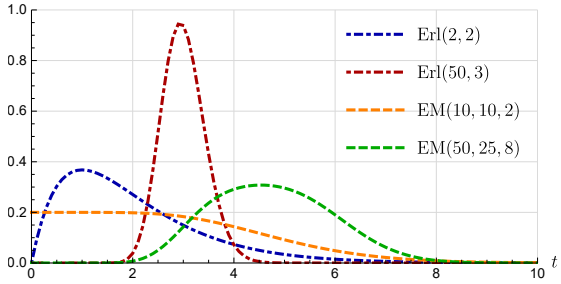

Fig. 7. Pdf of Erlang and EM distributions with various parameters.

	ERROR	FAIL	DETECT	REPAIR	STARTREJ	REJERR	REJOK	# st.	# up st.
comp. 1	Erl(2,40)	Erl(2,50)	EM(4,2,1)	EM(40,20,1)	Erl(100,60)	EM(8,6,1)	EM(6,3,1)	898	414
comp. 2	Erl(2,40)	Erl(2,25)	EM(4,2,1)	EM(40,20,1)	Erl(150,50)	EM(8,6,2)	EM(6,3,1)	1298	614
comp. 3	Erl(2,40)	Erl(2,50)	EM(4,2,2)	EM(40,20,2)	Erl(100,50)	EM(8,6,2)	EM(6,3,1)	898	414

Table 1. PH distributions of the components of Case 1, the number of states of the resulting CTMC and the number of up states in the CTMC.

4, 40 and 40 states, respectively, for a total of 684 failure states. In Table 1 we reported also the number of up states for each component. The product of these numbers provides the number of states of the CTMC that we should analyze in order to calculate the first passage time distribution exactly. For Case 1, it is 105,237,144 (i.e., approximately 10^8).

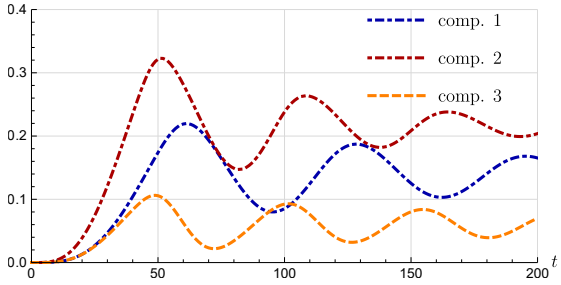

Fig. 8. Total probability of failure states as function of time for the three components of Case 1 in isolation.

To give an idea of the behavior of the components in isolation, in Fig. 8 we show the probability of being in a failure state as function of time for the three components.

The approximate first passage time distribution calculated by the proposed method is depicted in Fig. 9 together with an empirical cdf obtained by simulation. The number of simulation runs was set to $5 \cdot 10^5$. The figure also shows the 95% confidence band constructed around the empirical, simulation-based cdf using the Dvoretzky-Kiefer-Wolfowitz inequality [10,21]. The approximate first passage time distribution cannot be distinguished from the one obtained by simulation. The width of the confidence band is 0.00384. Fig. 10 shows the function $G(t)$ defined in Eq. (4) which is used to obtain the approximation $\hat{F}_T(t)$ through numerical integration of Eq. (6).

As Case 2, we consider a system obtained by small modifications of the system of Case 1 in such a way that failure probabilities become smaller. Specifically, the time to fire distribution of transition REPAIR of components 1 and 2 of Case 2 is EM(40,20,5) (instead of EM(40,20,1)) which means that these two components get back to state OK from state DETECTED five times faster. This change does not modify the state space, that is, the number of states and the number of

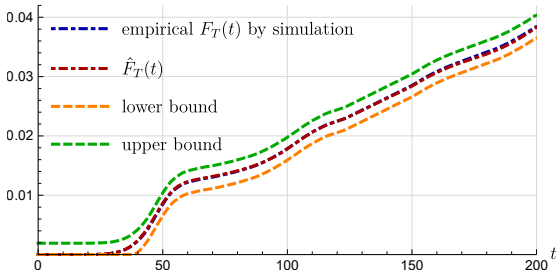

Fig. 9. Case 1: First passage time distribution by approximation and by simulation with 95% confidence bounds based on simulation.

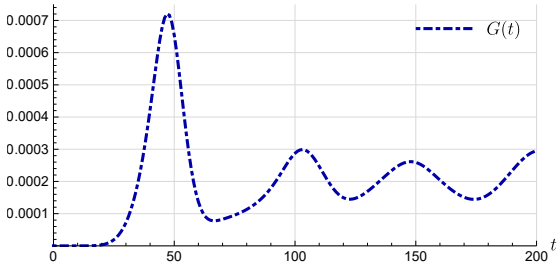

Fig. 10. Case 1: The function $G(t)$ defined in Eq. (4) which is the base of the approximation procedure.

up states is the same as in Case 1, reported in Table 1. The approximate first passage time distribution together with results obtained by 10^6 simulation runs are provided in Fig. 11. The approximate cdf deviates very slightly from the simulation-based cdf from about $t = 150$. Note however that, as indicated by the confidence band whose width is 0.00272, a much larger number of simulation runs would be necessary to estimate such small probabilities with high confidence. That is, the visible but very small difference can be due to chance.

Case 3 is obtained from Case 2 by increasing the number of phases of the applied Erlang distributions. Specifically, we double the number of phases of transition WAIT and change the number of phases of transitions ERROR and FAIL to 5 (in the two cases before it was 2). These modifications have a twofold impact. First, failure probabilities become even smaller because the Erlang distribution with shape parameter equal to 5 has a smaller variability and it becomes less likely that transitions ERROR and FAIL fire before a rejuvenation. Second, the state space becomes larger. The number of states of the three components is 2,898, 4,298 and 2,898, respectively. The number of up states in the components is 2,014, 3,014 and 2,014, respectively, meaning that exact analysis would require dealing with a CTMC with $12,225,374,744 \approx 1.2 \cdot 10^{10}$ states. Results are shown in Fig. 12 with $2 \cdot 10^6$ simulations runs. Also in this case there is a small but visible difference between the approximation and the simulation-based empirical

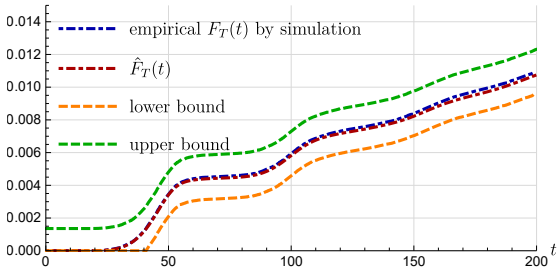

Fig. 11. Case 2: First passage time distribution by approximation and by simulation with 95% confidence bounds based on simulation.

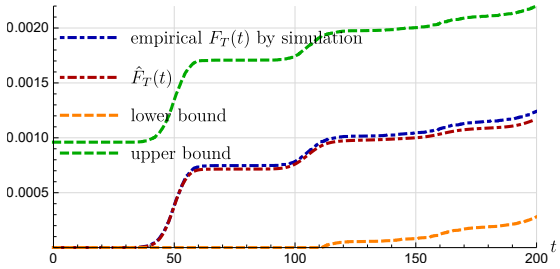

Fig. 12. Case 3: First passage time distribution by approximation and by simulation with 95% confidence bounds based on simulation.

cdf but, as indicated by the confidence band with width equal to 0.00192, it can very well be due to chance and only an extremely large number of simulation traces could verify the precision of the approximation.

The intended use of the proposed approximation is the analysis of reliability of systems composed of independent components. Consequently, we are interested in approximating relatively small probabilities. In Fig. 13 we show that the approximation can result in good precision even in case of much larger probabilities by evaluating Case 1 up to $t = 10000$.

4.3 Implementation Issues

The presented numerical results were calculated by a prototype implementation of the method using Wolfram Mathematica [14]. The transient probabilities of the components, which are necessary to compute $G(t)$ defined in Eq. (4), were determined by uniformization (see, e.g., [20]) with precision 10^{-8}, representing the infinitesimal generators of the components by sparse matrices. The approximate first passage time distribution was calculated by numerical integration of the differential equation Eq. (6) applying the NDSolve function of Mathematica. The sought relative precision was set to 10^{-8}. NDSolve evaluates $G(t)$ at several time points in order to compute $\hat{F}_T(t)$. Calculation of the transient probabilities of the components by uniformization is efficient if these time point are in increasing

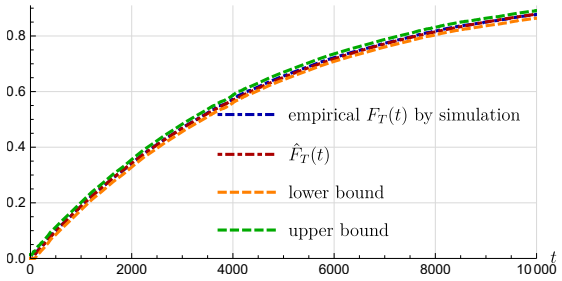

Fig. 13. Case 1: First passage time distribution by approximation and by simulation with 95% confidence bounds based on simulation up to $t = 10000$.

order. This is however not always the case since NDSolve, in order to guarantee precision, takes steps also backward in time. For this reason, for a time-efficient computation it is convenient to store the transient probabilities at a few recently used time points, allowing for not starting uniformization from $t = 0$.

The execution time of calculating the approximation for the complex system referred to as Case 1 in Section 4.2 (described in Table 1) up to $t = 200$ took 0.7 seconds. Case 2 required the same amount of time. Case 3 required instead about 4 seconds. This is because in this system both the state spaces and the intensities in the infinitesimal generators are larger (meaning that more steps are required for the uniformization to guarantee the same precision).

Simulation was also carried out in Mathematica based on the Petri net (that is, not the underlying CTMC) generating firing times according to the PH distributions of the transitions. A standard laptop was used parallelizing the generation of the simulation traces among 14 processor cores. Generating the $5 \cdot 10^5$ traces for Case 1 required around 30 minutes of computation. Case 2 required about an hour because generating a trace requires roughly the same amount of time but we generated twice as many traces. Simulation of Case 3 took about 5 hours because a single trace requires more time and we also generated more traces.

5 Conclusions

We presented an approximate solution to compute the cdf of the first passage time of a combination of component states without enumerating system states. The approximation achieves high accuracy when failures are rare or when the system includes many components. In future work, we plan to extend the approach to analyze systems with m-out-of-n failure conditions. We also plan to investigate how far $\{X'(t)\}$ is from a product form.

References

1. Amparore, E.G.: Stochastic modelling and evaluation using GreatSPN. SIGMETRICS Perform. Evaluation Rev. **49**(4), 87–91 (2022)
2. Barlow, R.E., Proschan, F.: Theory of maintained systems: Distribution of time to first system failure. Math. Oper. Res. **1**(1), 32–42 (1976)
3. Bobbio, A., Cumani, A.: ML estimation of the parameters of a PH distribution in triangular canonical form. Comp. Perf. Eval. (22), 33–46 (1992)
4. Bobbio, A., Horváth, A., Telek, M.: Matching three moments with minimal acyclic phase type distributions. Stoch. Models **21**, 303–326 (2005)
5. Brown, M.: The first passage time distribution for a parallel exponential system with repair. In: Proceedings of the Conference on Reliability and Fault Tree Analysis, pp. 365–396. Society for Industrial and Applied Mathematics (1975)
6. Buchholz, P., Katoen, J., Kemper, P., Tepper, C.: Model-checking large structured Markov chains. J. Log. Algebraic Methods Program. **56**(1-2), 69–97 (2003)
7. Buchholz, P., Kemper, P.: Kronecker based matrix representations for large Markov models. In: Validation of Stochastic Systems - A Guide to Current Research. LNCS, vol. 2925, pp. 256–295. Springer (2004)
8. Carnevali, L., Paolieri, M., Reali, R., Scommegna, L., Vicario, E.: A Markov regenerative model of software rejuvenation beyond the enabling restriction. In: Proceedings of IEEE ISSRE Workshops (WOSAR). pp. 138–145. IEEE (2022)
9. Clark, A., Gilmore, S., Hillston, J., Tribastone, M.: Stochastic process algebras. In: Int. School on Formal Methods. LNCS, vol. 4486, pp. 132–179. Springer (2007)
10. Dvoretzky, A., Kiefer, J., Wolfowitz, J.: Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. Ann. Math. Stat. **27**(3), 642 – 669 (1956)
11. Feldman, A., Whitt, W.: Fitting mixtures of exponentials to long-tail distributions to analyze network performance models. Perf. Eval. **31**, 245–279 (1998)
12. Hensel, C., Junges, S., Katoen, J., Quatmann, T., Volk, M.: The probabilistic model checker Storm. Int. J. Softw. Tools Technol. Transf. **24**(4), 589–610 (2022)
13. Horváth, A., Telek, M.: Approximating heavy tailed behavior with phase-type distributions. In: Proc. of 3rd International Conference on Matrix-Analytic Methods in Stochastic Models. Leuven, Belgium (June 2000)
14. Inc., W.R.: Mathematica, Version 14.0, `https://www.wolfram.com/mathematica`, Champaign, IL, 2024
15. Kleinrock, L.: Queueing Systems, Vol. 1: Theory. Wiley and Sons (1975)
16. Kwiatkowska, M.Z., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: Proceedings of CAV 2011. LNCS, vol. 6806, pp. 585–591. Springer (2011)
17. Marsan, M.A., Conte, G., Balbo, G.: A class of generalized stochastic Petri nets for the performance evaluation of multiprocessor systems. ACM Trans. Comput. Syst. **2**(2), 93–122 (1984)
18. Neuts, M.: Probability distributions of phase type. In: Liber Amicorum Prof. Emeritus H. Florin. pp. 173–206. University of Louvain (1975)
19. Ross, S.M.: On the time to first failure in multicomponent exponential reliability systems. Stochastic Process. Appl. **4**(2), 167–173 (1976)
20. Stewart, W.J.: Probability, Markov Chains, Queues, and Simulation: The Mathematical Basis of Performance Modeling. Princeton University Press (2009)
21. Wasserman, L.: All of Statistics. New York, Springer (2004)