Markov Modeling of Availability and Unavailability Data

Peter Buchholz, Jan Kriege
Informatik IV, TU Dortmund
D-44221 Dortmund, Germany
{peter.buchholz,jan.kriege}@cs.tu-dortmund.de

Author prepared version of a paper published in

Proc. of the Tenth European Dependable Computing Conference - EDCC 2014, IEEE, 2014

http://dx.doi.org/10.1109/EDCC.2014.22

Markov Modeling of Availability and Unavailability Data

Peter Buchholz and Jan Kriege
Informatik IV, TU Dortmund
D-44221 Dortmund, Germany
{peter.buchholz,jan.kriege}@cs.tu-dortmund.de

Abstract

Markov models are often used in performance and dependability analysis and allow a precise and numerically stable computation of many result measures including those which result from rare events. It is, however, known that simple exponential distributions, which are the base of Markov modeling, cannot adequately describe the duration of availability or unavailability intervals of components in a distributed system. Commonly used in modeling those durations are Weibull, log-normal or Pareto distributions that can also capture a possibly heavy tailed behavior but cannot be analyzed analytically or numerically. An alternative to applying the mentioned distributions in modeling availability or unavailability intervals are phase type distributions and Markovian arrival processes that still result in a Markov model. Based on experiments for a large number of publically available availability traces, we show that phase type distributions are a flexible alternative to other commonly known distributions and even more that Markov models can be easily extended to capture also correlation in the length of availability or unavailability intervals.

Keywords: Dependability Modeling, Phase Type Distributions, Availability Distributions, Parameter Fitting of Markov Models.

1 Introduction

With an increasing number of components in distributed systems, the probability of observing failures resulting in the unavailability of the whole system or parts of the system increases too. The consequences of unavailability of parts of a distributed system are often a degradation of system performance and sometimes even wrong results. This may then result in the loss of customers and profit. Several practical examples for the implications of failures in real systems can be found in [13]. The serious consequences of failures and the resulting unavailability of system components implies that these aspects have to be considered in system planing and control. Usually these tasks are performed model based such that adequate models which include detailed and realistic models of the duration of availability and unavailability intervals are a key step of model building.

Availability is analyzed using stochastic simulation models [9] or Markov models [27] which can be analyzed using simulation or numerical techniques. Simulation is very flexible and can be used independently of the model size and the used distributions. However, simulation is an approximation technique which only observes several sample paths such that confidence intervals have to be computed for the derived results. If models are large, result measures have a high variance or occur rarely, it is cumbersome to estimate or compute reliable and precise results from simulation. On today's computers very large Markov models can be analyzed numerically in a few seconds and the analysis is exact up to numerical errors such that steady state or interval availability of complex systems can be easily determined from Markov models [23].

Independently whether simulation or numerical techniques are applied, the central step of availability modeling is the modeling of the duration of availability and unavailability intervals. For simulation models *Weibull*, *Pareto*, *log-normal* or *Gamma* distributions are commonly used [16, 13]. For these distributions maximum likelihood estimators of the parameters are available [19] which allow an easy parameterization of the distributions from available measurements. Interestingly, only few studies exist which compare the modeling power of the different distributions on real data as already noticed in [16, 13]. If Markov models are built for availability modeling, then the exponential distribution is the base since holding times in states of Markov chains have to be memoryless [27]. However, it is also known that availability or unavailability interval lengths are rarely exponentially distributed such that the exponential distribution is only a rough approximation of the real behavior and modeling errors become large. It is possible to use Markov models in conjunction with more general distributions by using *phase type distributions* (PHDs) [21] which are a distribution type that can theoretically approximate any non-negative distribution arbitrarily close. Although PHDs have a few times been used to model availability or unavailability intervals [8, 22], we are not aware of any study that generates PHDs from measured availability data and compares the resulting PHDs systematically with other distributions. The major obstacle, that prohibited a wider use of PHDs in the past, lies in the complexity of finding adequate parameters such that the PHD matches the observed behavior. However, nowadays efficient algorithms are available to fit the parameters of a PHD according to measured data [25, 28] and it has been shown that the resulting PHDs are very accurate models for the data, often much better than standard distribution with a small number of parameters.

Another aspect which is rarely considered in modeling availability or unavailability is the autocorrelation in the process. Since failures in distributed systems have often similar reasons, it is very likely that the intervals are correlated and disregarding this correlation may result in an overestimation of system performance. Correlation can be considered according to different measures. It may occur in space since parts of a distributed system which are located nearby are more likely to fail simultaneously. It may also occur in time since a failure is usually more likely if a system is highly loaded and high loads are often observed during the day in the middle of the week, whereas the system is often lightly loaded over the weekend or the night. The aspect of time dependent failures has been analyzed in [29]. Finally, correlation of consecutive availability or unavailability intervals can be observed even if location and absolute time are neglected. We will consider this aspect in our models but the proposed approaches may be extended to take into account also locations or absolute time. In the Markovian setting, PHDs can be easily extended to *Markovian Arrival Processes* (MAPs) [20] which have been mainly applied to model correlated arrival streams in computer networks [4, 15] but similar process types can in principle also be used to model availability and unavailability intervals such that the resulting model is a Markov chain and describes correlated durations of availability and unavailability intervals.

In this paper we use the traces from the *failure trace archive* (FTA) [13, 16] and generate PHDs to model the length of availability and unavailability intervals. The resulting distributions are compared to standard distributions like Weibull or Gamma which have been used in the mentioned papers as a model for the data. Furthermore, we consider correlations in the length of unavailability and availability intervals and use Markov chains to model these correlations.

The paper is structured as follows. In the next section the basic definitions and notations are introduced. Section 3 describes the modeling of the unavailability and availability intervals for different traces from the FTA using PHDs. The quality of the resulting PHDs is compared with the quality of other distributions like which have shown in [16] to be adequate models for the data. In Section 4, the correlation in the length of availability and unavailability intervals is analyzed and modeled using a Markov model which can be interpreted as a MAP. The paper ends with the conclusions and an outlook of future work in the area.

2 Background and Basic Definitions

In the following paragraphs we introduce the basic definitions and notations that are needed for the experiments presented subsequently. Additionally, we give a brief overview of related approaches from the literature.

2.1 Distributions of Availability and Unavailability Intervals

We use the same basic concepts and terminology as in [13] which is commonly used for describing dependability of distributed systems. A *failure* is an event that implies that a component of a distributed system does not act as described in its specification. A time interval in which a component does not fulfill its specification is denoted as an unavailability interval, often this means that the component does not react at all, but it may also contain a behavior where the system is too slow, computes incorrect results or does not observe the required SLA. On the other hand, an availability interval is a period where the component behaves according to its specification. As indicated it is not always clear whether a component is available or not because it may only be too slow during some period of time and it is also not clear why the system is not available at some time. However, abstract models, as they are required for complex systems, only consider the length of availability and unavailability intervals of components to evaluate a system. For a realistic modeling, stochastic models have to be used that describe the observed behavior of real systems. Since different possibilities exist to model the length of availability and unavailability intervals, a fair comparison of the model types has to be based on real data and has to be done in a systematic way.

As already observed in [16, 13], most studies in the area use their own data which is only available in a specific format, if it is available at all. To provide a common base of availability data, the failure time archive (FTA) has been built [16, 13]. The FTA contains a large number of failure traces from different sources in a common format. Furthermore, it includes Matlab scripts to perform a basic analysis of the data. In the mentioned papers the length of the availability and unavailability intervals of the traces have been modeled using Weibull, Pareto, log-normal, Gamma and exponential distributions. It turns out that the Weibull, log-normal and Gamma distribution are often adequate which means that the Kolmogorov-Smirnov (KS) and Anderson-Darling (AD) tests give a high p-value (for details about the test see [16, 13] and the standard literature [19]).

The exponential distribution fails to model the data adequately. Even if the Gamma distribution can at least be approximated by a Markov model, the available results do not allow one to build a Markov model from the data such that simulation has to be used as analysis method which is done in [13].

To the best of our knowledge, no attempt has been made yet to model the data from the FTA using PHDs. This will be done in this paper and the results will compared to the above mentioned models from [16].

The data sets in the FTA contain for each component of a distributed system the absolute times when a session begins and ends which means that the component is available in this interval and unavailable in between the intervals. As described in [16], it is not always clearly defined when the system is available and when not. For the detailed traces an interpretation of availability and unavailability phases is not always obvious but has a significant impact on the observed behavior. However, we do not consider this interpretation in detail and use the data in its raw form, as also done in the basic modeling presented in [16].

If we consider one component k in a distributed system, then we obtain from a trace an ordered sequence $\mathcal{T}^{(k)} = (a_1^k, u_1^k, a_2^k, \ldots, u_{m^{(k)}-1}^k, a_{m^{(k)}}^k)$, where a_i^k is the length of the ith availability interval and u_i^k is the length of the ith unavailability interval. We define additionally two subtraces for component k, namely $\mathcal{A}^{(k)} = (a_1^k, \ldots, a_{m^{(k)}}^k)$ and $\mathcal{U}^{(k)}(u_1^k, \ldots, u_{m^{(k)}-1}^k)$. In some cases, traces may also start or end with an unavailability interval which implies that the following equations have to be slightly modified. If the distribution of the interval length is considered, it is assumed for abstract modeling that all components behave statistically identically such that $\mathcal{A} = (a_1^1, \ldots, a_{m^{(K)}}^K)$ and $\mathcal{U} = (u_1^1, \ldots, u_{m^{(K)}-1}^K)$ collect all durations of availability and unavailability intervals. Let $m_a = \sum_{k=1}^K m^{(k)}$ and $m_u = \sum_{k=1}^K (m^{(k)} - 1)$ be the number of elements in the traces $\mathcal{A}$ and $\mathcal{U}$ and denote the elements as a_i and u_i ($i = 1, \ldots, m_a$ or $i = 1, \ldots, m_u$), respectively. Then

$$\hat{\mu}_a^j = \frac{1}{m_a} \sum_{i=1}^{m_a} (a_i)^j \text{ and } \hat{\mu}_u^j = \frac{1}{m_u} \sum_{i=1}^{m_u} (u_i)^j \tag{1}$$

are estimates for the jth moments of the length of availability and unavailability intervals. For the first moments we write $\hat{\mu}_a$ and $\hat{\mu}_u$. Additionally,

$$\hat{\mu}_a^{(k)} = \frac{1}{m^{(k)}} \sum_{i=1}^{m^{(k)}} a_i^k \text{ and } \hat{\mu}_u^{(k)} = \frac{1}{m^{(k)} - 1} \sum_{i=1}^{m^{(k)}-1} u_i^k \tag{2}$$

are the first moments of the length of availability and unavailability intervals of component k.

2.2 Phase Type Distributions

Phase type distributions (PHDs) are the basic model to describe non-exponential behavior in a Markov setting. A PHD can be interpreted as an absorbing Markov chain with $n + 1$ states, generator matrix

$$\boldsymbol{Q} = \begin{pmatrix} \boldsymbol{D}_0 & \boldsymbol{d}_1 \\ \boldsymbol{0} & 0 \end{pmatrix} \tag{3}$$

and initial distribution $(\boldsymbol{\pi}, 0)$. $\boldsymbol{D}_0$ is a non-singular subgenerator (i.e., diagonal elements are negative, non-diagonal elements are non-negative and all row sums are smaller or equal 0), $\boldsymbol{d}_1 = \boldsymbol{D}_0 \mathbb{I}$ and $\boldsymbol{\pi}$ is a distribution vector. For details about PHDs we refer to the literature [8, 21, 25, 28].

The determination of the parameters of a PHD to match trace data is denoted as *parameter fitting*. There are two classes of approaches, namely moment and maximum likelihood based fitting. Moment based fitting tries to find a PHD such that lower order moments of the trace (see (1)) are exactly matched or approximated by a PHD. Corresponding approaches can be found in [3, 11, 14]. Most methods of this type are fairly efficient but allow one only to match or approximate a smaller number (at most 5 to 10) lower order moments. This is often not sufficient, in particular if the density function is multimodal.

The alternative to moment based parameter fitting are methods that maximize the likelihood which is defined for a trace $\mathcal{T}$ and a PHD $(\boldsymbol{\pi}, \boldsymbol{D}_0)$ as

$$\mathcal{L}_{(\boldsymbol{\pi},\boldsymbol{D}_0)}(\mathcal{T}) = \prod_{t\in\mathcal{T}} f_{(\boldsymbol{\pi},\boldsymbol{D}_0)}(t) = \prod_{t\in\mathcal{T}} \boldsymbol{\pi} e^{\boldsymbol{D}_0 t}\boldsymbol{d}_1. \tag{4}$$

Of course, the likelihood can be defined for any distribution with density $f(t)$. The maximization problem is then given by

$$\mathcal{L}^*(\mathcal{T}) = \max_{(\boldsymbol{\pi},\boldsymbol{D}_0)} \left(\mathcal{L}_{(\boldsymbol{\pi},\boldsymbol{D}_0)}(\mathcal{T})\right). \tag{5}$$

Maximum likelihood estimates are computed with *expectation maximization* (EM) algorithms [2]. Recent variants of these algorithms [25, 28] are very efficient and can be applied to large traces to generate PHDs with 10 through 20 states in at most a few minutes, which is usually sufficient for most data sets. Even if EM algorithms are only local optimization algorithms, the resulting PHDs are often a better approximation of a trace than those resulting from moment fitting.

To compare the quality of some distribution as a model of a trace, different methods exist. In [16, 13] the p values of the KS and AD test are used. Since the AD test is, as far as we know, not available for PHDs, we use only the p values of the KS tests in our comparisons in Section 3. The KS test can be used for PHDs since the distribution function is given by

$$F_{(\boldsymbol{\pi},\boldsymbol{D}_0)}(t) = 1 - \boldsymbol{\pi} e^{\boldsymbol{D}_0 t}\mathbb{1} \tag{6}$$

and can be easily computed. An alternative criterion is the comparison of the likelihood values of different distributions. It is obvious that a larger likelihood indicates a better fitting. In this way, distributions like Weibull or Pareto can be compared with PHDs.

The proposed fitting approaches are used to fit PHDs for the duration of availability and unavailability phases. Let $(\boldsymbol{\pi}^a, \boldsymbol{D}_0^a)$ and $(\boldsymbol{\pi}^u, \boldsymbol{D}_0^u)$ be the resulting PHDs. If we assume that a component is initially up, then the availability of the component can be described by a Markov chain with initial distribution $(\boldsymbol{\pi}^a, \boldsymbol{0})$ and generator matrix

$$\boldsymbol{Q} = \begin{pmatrix} \boldsymbol{D}_0^a & \boldsymbol{d}_1^a\boldsymbol{\pi}^u \\ \boldsymbol{d}_1^u\boldsymbol{\pi}^a & \boldsymbol{D}_0^u \end{pmatrix}. \tag{7}$$

From this model the transient and stationary availability can be easily analyzed. Let

$$\boldsymbol{p}_t = (\boldsymbol{p}_t^a, \boldsymbol{p}_t^u) = \boldsymbol{p}_0 e^{\boldsymbol{Q}t} \tag{8}$$

and $\boldsymbol{p} = (\boldsymbol{p}^a, \boldsymbol{p}^u) = \lim_{t\to\infty} \boldsymbol{p}_t$. Since $\boldsymbol{Q}$ is irreducible, $\boldsymbol{p}$ is the unique solution of $\boldsymbol{p}\boldsymbol{Q} = \boldsymbol{0}$ subject to $\boldsymbol{p}\mathbb{1} = 1$.

From the vectors certain availability measures can be derived. For example, the system is available at time t with probability $q_t^a = \boldsymbol{p}_t^a\mathbb{1}$ and the steady state availability equals $q^a = \boldsymbol{p}^a\mathbb{1}$. Since the use of distributions is based on the assumption of independent and identically distributed availability and unavailability intervals, the probability that L out of K ($L \leq K$) components are available equals

$$\binom{L}{K} (q_t^a)^L (1 - q_t^a)^{K-L} . \tag{9}$$

More complex measures can also be computed for the components. For example the mean time before the component becomes unavailable under the condition that is available at a random time point equals

$$\frac{\boldsymbol{p}^a}{\boldsymbol{p}^a \mathbb{1}} (-\boldsymbol{D}_0^a)^{-1} \mathbb{1}.$$

Similarly other quantities can be computed. We briefly describe below how the Markov model described in (7) can be used to compute the availability in a more general context.

2.3 Adding Correlations to PHDs

The assumption of independent and identically distributed durations of availability and unavailability intervals is usually only an approximation. In practice, failures resulting in unavailability of components have common reasons and are therefore often correlated. The traces from the FTA have been analyzed with respect to correlation in [29]. In this paper a time dependent correlation is analyzed and the behavior is modeled by a process with two phases, a phase with low and another phase with high failure rates. Here we use two different models for correlations. First, we consider the correlation of the durations of availability and unavailability intervals of single components. This viewpoint considers local sources for failures which may result in sequences of consecutive failures. Second, we consider the mean duration of availability intervals depending on the number of failed components in the system. This usually means, that the length of availability intervals is smaller, if less components are available such that the load of a component increases. Even if time dependent failures are not explicitly considered in this model, the behavior is similar to the behavior of the model described in [29] because intervals with many available components and low failure rates alternate with phase where many components are unavailable and failure rates are high. However, a detailed comparison of both modeling approaches is left for future research.

We begin with analysis of the dependence of component failures and define for a system of K components numbered 1 through K the following four quantities.

$$
\begin{aligned}
\hat{\rho}_a^{(k)} &= \frac{(m^{(k)}-1)\sum_{i=1}^{m^{(k)}-1}(a_i^k-\hat{\mu}_a^{(k)})(a_{i+1}^k-\hat{\mu}_a^{(k)})}{(m^{(k)}-2)\sum_{i=1}^{m^{(k)}}((a_i^k-\hat{\mu}_a^{(k)})^2} \\
\hat{\rho}_u^{(k)} &= \frac{(m^{(k)}-2)\sum_{i=1}^{m^{(k)}-2}(u_i^k-\hat{\mu}_u^{(k)})(u_{i+1}^k-\hat{\mu}_u^{(k)})}{(m^{(k)}-3)\sum_{i=1}^{m^{(k)}-1}((u_i^k-\hat{\mu}_u^{(k)})^2} \\
\hat{\rho}_{au}^{(k)} &= \frac{\sum_{i=1}^{m^{(k)}-1}(a_i^k-\hat{\mu}_a^{(k)})(u_i^k-\hat{\mu}_u^{(k)})}{\sqrt{\sum_{i=1}^{m^{(k)}-1}(a_i^k-\hat{\mu}_a^{(k)})^2\sum_{i=1}^{m^{(k)}-1}(u_i^k-\hat{\mu}_u^{(k)})^2}} \\
\hat{\rho}_{ua}^{(k)} &= \frac{\sum_{i=1}^{m^{(k)}-1}(u_i^k-\hat{\mu}_u^{(k)})(a_{i+1}^k-\hat{\mu}_a^{(k)})}{\sqrt{\sum_{i=1}^{m^{(k)}-1}(a_i^k-\hat{\mu}_a^{(k)})^2\sum_{i=1}^{m^{(k)}-1}(u_i^k-\hat{\mu}_u^{(k)})^2}}
\end{aligned}
\tag{10}
$$

which are the correlation coefficients of the length of availability, unavailability, availability-unavailability and unavailability-availability intervals for component k. Since the implicit assumption of the modeling is that all components are stochastically identical, the corresponding coefficients over all components can be computed from the following equations.

$$
\begin{aligned}
\hat{\rho}_a &= \sum_{k=1}^{K}\frac{m^{(k)}}{m_a}\rho_a^{(k)} & \hat{\rho}_u &= \sum_{k=1}^{K}\frac{m^{(k)}-1}{m_u}\rho_a^{(k)} \\
\hat{\rho}_{au} &= \sum_{k=1}^{K}\frac{m^{(k)}-1}{m_u}\rho_{au}^{(k)} & \hat{\rho}_{ua} &= \sum_{k=1}^{K}\frac{m^{(k)}-1}{m_u}\rho_{ua}^{(k)}
\end{aligned}
\tag{11}
$$

The quantities describe the correlation between availability and unavailability intervals of components. We now build models that consider this correlation and begin with $\hat{\rho}_{au}$ and $\hat{\rho}_{ua}$. The resulting Markov model is an extension of (7) with generator matrix

$$
\boldsymbol{Q} = \begin{pmatrix} \boldsymbol{D}_0^a & \boldsymbol{Q}_{au} \\ \boldsymbol{Q}_{ua} & \boldsymbol{D}_0^u \end{pmatrix}. \tag{12}
$$

Observe that (7) is just a special case of (12). Let $\boldsymbol{M}^a = (-\boldsymbol{D}_0^a)^{-1}$ and $\boldsymbol{M}^u = (-\boldsymbol{D}_0^u)^{-1}$. We assume that matrices $\boldsymbol{Q}_{au}$ and $\boldsymbol{Q}_{ua}$ are chosen such that $\boldsymbol{\pi}^a\boldsymbol{M}^a\boldsymbol{Q}_{au} = \boldsymbol{\pi}^u$ and $\boldsymbol{\pi}^u\boldsymbol{M}^u\boldsymbol{Q}_{ua} = \boldsymbol{\pi}^a$, then the sojourn time in the first subset of states (where the system is available) is distributed according to PHD $(\boldsymbol{\pi}^a, \boldsymbol{D}_0^a)$ and the sojourn time in the second subset of states (where the system is unavailable) is distributed according to PHD $(\boldsymbol{\pi}^u, \boldsymbol{D}_0^u)$.

The Markov chain in (12) can be interpreted as a Markovian Arrival Process (MAP). A MAP [20, 12, 17] can be interpreted as a irreducible Markov chain with generator matrix $\boldsymbol{Q} = \boldsymbol{D}_0 + \boldsymbol{D}_1$ such that $\boldsymbol{D}_0$ is a non-singular subgenerator and $\boldsymbol{D}_1$ is non-negative. The interpretation of the behavior of a MAP is as follows, transitions from $\boldsymbol{D}_0$ are silent whereas transitions in matrix $\boldsymbol{D}_1$ are related to events. The distribution of the time between two events is given by a PHD $(\boldsymbol{\pi}, \boldsymbol{D}_0)$ where $\boldsymbol{\pi}\,(-\boldsymbol{D}_0)^{-1}\,\boldsymbol{D}_1 = \boldsymbol{\pi}$ and $\boldsymbol{\pi}\mathbb{I} = 1$. The inter-event times of MAPs are usually correlated.

For the modeling of availability and unavailability interval length, we have two types of events, namely the end of an availability interval and the end of an unavailability interval. Since the duration of availability and unavailability intervals usually differs, it is not appropriate to model both events using a single process. Instead we use so called marked MAPs (MMAPs) [10] that describe the generation of K rather than only one event. In our case $K = 2$ is sufficient. An MMAP with two event types is described by three matrices $\boldsymbol{D}_0$, $\boldsymbol{D}_1$ and $\boldsymbol{D}_2$ such that $\boldsymbol{Q} = \boldsymbol{D}_0 + \boldsymbol{D}_1 + \boldsymbol{D}_2$ is an irreducible generator matrix, $\boldsymbol{D}_0$ is as above and $\boldsymbol{D}_1, \boldsymbol{D}_2$ are non-negative. Some results about results like conditional moments or conditional joint moments of MMAPs can be found in [6]. In our setting matrix $\boldsymbol{Q}$ from (12) is decomposed as follows.

$$\begin{gathered}
\boldsymbol{D}_0 = \begin{pmatrix} \boldsymbol{D}_0^a & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{D}_0^u \end{pmatrix}, \\
\boldsymbol{D}_1 = \begin{pmatrix} \boldsymbol{0} & \boldsymbol{Q}_{au} \\ \boldsymbol{0} & \boldsymbol{0} \end{pmatrix}, \ \boldsymbol{D}_2 = \begin{pmatrix} \boldsymbol{0} & \boldsymbol{0} \\ \boldsymbol{Q}_{ua} & \boldsymbol{0} \end{pmatrix}.
\end{gathered} \tag{13}$$

Events of type 1 indicate the end of an availability interval and events of type 2 the end of an unavailability interval. The parameter fitting of MMAPs can be done according to moments [6] or the likelihood value [15]. We consider a two phase approach based on the ideas proposed in [12] where first a PHD is fitted to match the distribution and afterwards the PHD is extended to an MMAP such that the distribution remains the same but correlation is added to match the coefficient of correlation. This approach can be extended to capture also other quantities describing the correlation as shown in [6]. However, we consider only the coefficient of correlation as estimated in (10).

For our MMAP we denote by ξ_{au} and ξ_{ua} the correlation coefficients between availability and unavailability interval length. The values are computed from the following equations.

$$\begin{aligned}
\xi_{au} &= \frac{\boldsymbol{\pi}^a(\boldsymbol{M}^a)^2\boldsymbol{Q}_{au}\boldsymbol{M}^u\,\mathbb{1} - (\boldsymbol{\pi}^a\boldsymbol{M}^a\,\mathbb{1})(\boldsymbol{\pi}^u\boldsymbol{M}^u\,\mathbb{1})}{\sqrt{\left(2\boldsymbol{\pi}^a(\boldsymbol{M}^a)^2\,\mathbb{1} - (\boldsymbol{\pi}^a\boldsymbol{M}^a\,\mathbb{1})^2\right)\left(2\boldsymbol{\pi}^u(\boldsymbol{M}^u)^2\,\mathbb{1} - (\boldsymbol{\pi}^u\boldsymbol{M}^u\,\mathbb{1})^2\right)}} \\
&= \frac{\boldsymbol{a}\boldsymbol{Q}_{au}\boldsymbol{b} - \psi}{\eta} \\
\xi_{ua} &= \frac{\boldsymbol{\pi}^u((\boldsymbol{M}^u)^2\boldsymbol{Q}_{ua}\boldsymbol{M}^a\,\mathbb{1} - (\boldsymbol{\pi}^a\boldsymbol{M}^a\,\mathbb{1})(\boldsymbol{\pi}^u\boldsymbol{M}^u\,\mathbb{1})}{\sqrt{\left(2\boldsymbol{\pi}^a(\boldsymbol{M}^a)^2\,\mathbb{1} - (\boldsymbol{\pi}^a\boldsymbol{M}^a\,\mathbb{1})^2\right)\left(2\boldsymbol{\pi}^u(\boldsymbol{M}^u)^2\,\mathbb{1} - (\boldsymbol{\pi}^u\boldsymbol{M}^u\,\mathbb{1})^2\right)}} \\
&= \frac{\boldsymbol{c}\boldsymbol{Q}_{au}\boldsymbol{d} - \phi}{\zeta}
\end{aligned} \tag{14}$$

Observe that the vectors $\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}, \boldsymbol{d}$ and the scalars ψ, ϕ, η, ζ result from the PHDs modeling the length of availability and unavailability intervals. Since the estimated autocorrelation coefficients should be matched or at least approximated by $\xi_{.}$, $(\xi_{au} - \hat{\rho}_{au})^2$ and $(\xi_{ua} - \hat{\rho}_{ua})^2$ have to be minimized under the constraints $\boldsymbol{\pi}^a\boldsymbol{M}^a\boldsymbol{Q}_{au} = \boldsymbol{\pi}^u$, $\boldsymbol{Q}_{au}\mathbb{1} = -\boldsymbol{D}_0^a\mathbb{1}$, $\boldsymbol{Q}_{au} \geq \boldsymbol{0}$ and $\boldsymbol{\pi}^u\boldsymbol{M}^u\boldsymbol{Q}_{ua} = \boldsymbol{\pi}^a$, $\boldsymbol{Q}_{ua}\mathbb{1} = -\boldsymbol{D}_0^u\mathbb{1}$, $\boldsymbol{Q}_{ua} \geq \boldsymbol{0}$. This is a non-negative least squares problem with linear constraints. If an exact solution exists, it can be computed from a set of linear equations of an order smaller or equal to $n_a \cdot n_u$ where n_a and n_u the number of phases of the PHDs modeling the availability and unavailability interval length, respectively. The solution is described in the context of parameter fitting of MAPs in [12] and for MMAPs in [6].

To match or approximate additionally $\hat{\rho}_a$ and $\hat{\rho}_u$, we show first the coefficients of correlation of the sojourn time in the availability and unavailability phases.

$$\begin{aligned}
\xi_a &= \frac{\boldsymbol{\pi}^{\boldsymbol{a}}(\boldsymbol{M}^a)^2\boldsymbol{Q}_{au}\boldsymbol{M}^u\boldsymbol{Q}_{ua}\boldsymbol{M}^a\,\mathbb{1} - (\boldsymbol{\pi}^a\boldsymbol{M}^a\,\mathbb{1})^2}{2\boldsymbol{\pi}^a(\boldsymbol{M}^a)^2\,\mathbb{1} - (\boldsymbol{\pi}^a\boldsymbol{M}^a\,\mathbb{1})^2} \\
&= \frac{\boldsymbol{a}\boldsymbol{Q}_{au}\boldsymbol{M}^u\boldsymbol{Q}_{ua}\boldsymbol{b} - \varphi}{\upsilon} \\
\xi_u &= \frac{\boldsymbol{\pi}^{\boldsymbol{u}}(\boldsymbol{M}^u)^2\boldsymbol{Q}_{ua}\boldsymbol{M}^a\boldsymbol{Q}_{au}\boldsymbol{M}^u\,\mathbb{1} - (\boldsymbol{\pi}^u\boldsymbol{M}^u\,\mathbb{1})^2}{2\boldsymbol{\pi}^u(\boldsymbol{M}^u)^2\,\mathbb{1} - (\boldsymbol{\pi}^u\boldsymbol{M}^u\,\mathbb{1})^2} \\
&= \frac{\boldsymbol{c}\boldsymbol{Q}_{ua}\boldsymbol{M}^a\boldsymbol{Q}_{au}\boldsymbol{d} - \chi}{\omega}
\end{aligned} \tag{15}$$

In the equations only $\boldsymbol{Q}_{au}$ and $\boldsymbol{Q}_{ua}$ are unknown. Unfortunately, the equations are no longer linear in $\boldsymbol{Q}_{au}$ and $\boldsymbol{Q}_{ua}$. However, if assume that $\boldsymbol{Q}_{au}$ is known, then the problem of computing $\boldsymbol{Q}_{ua}$ is again a non-negative least squares problem and if we assume that $\boldsymbol{Q}_{ua}$ is known, then computation of $\boldsymbol{Q}_{au}$ is a non-negative least squares problems. Thus, first matrices are computed according to (14) using a constrained non-negative least squares solution algorithm. Then the approach can be iterated until the correlation coefficients are matched or approximated adequately. The outlined approach is denoted as alternating least squares [7, 18]. If it is not possible to approximate the correlation coefficients close enough, one can try to perform equivalence transformations of the PHDs $(\boldsymbol{\pi}^a \boldsymbol{D}_0^a)$ and $(\boldsymbol{\pi}^u \boldsymbol{D}_0^u)$ to increase their flexibility as in [5].

After matrix $\boldsymbol{Q}$ is available, the component can be analyzed using (8). With the additional assumption that components are independent, the complete system can be analyzed using (9).

The dependencies considered in the previous equations are all based on local failure conditions in a component such that dependencies between successive availability and unavailability intervals are described using only informations about this component. Apart from these conditions there are also global conditions such that the availability and unavailability intervals of different components are correlated. This dependency has also been considered in [29] where two failure modes are introduced, a normal mode and a peak mode. In peak mode, failure rates are higher. The model considers the duration of peak and normal modes and the length of the availability and unavailability intervals in both modes. Here we use a different approach to model the same phenomena.

We denote by $\delta(a_k)$ for $a_k \in \mathcal{A}$ and $\delta(u_k)$ for $u_k \in \mathcal{U}$ the number of components that are available at the end of the corresponding interval excluding the component to which the interval belongs. Let $m_a(k)$ and $m_u(k)$ be the number of availability and unavailability intervals that end if k other components are available. Then the first moments conditioned on the number of available components at the end of the interval are given by

$$\begin{aligned} \hat{\mu}_a(k) &= \tfrac{1}{m_a(k)} \sum_{a_i \in \mathcal{A}, \delta(a_i)=k} a_i \quad \text{and} \\ \hat{\mu}_u(k) &= \tfrac{1}{m_u(k)} \sum_{u_i \in \mathcal{U}, \delta(u_i)=k} u_i . \end{aligned} \tag{16}$$

If some values $m_a(k)$ or $m_u(k)$ are 0, then the corresponding values are approximated by inter- or extrapolation. Then define

$$\alpha_k = \frac{\hat{\mu}_a}{\hat{\mu}_a(k)} \text{ and } \beta_k = \frac{\hat{\mu}_u}{\hat{\mu}_u(k)}. \tag{17}$$

Under the assumption that the length of intervals depends on the state of components but the distribution of the interval length remains similar, the following Markov chain describes the behavior of one component when k other components are available.

$$\boldsymbol{Q}[k] = \begin{pmatrix} \alpha_k \boldsymbol{D}_0^a & \alpha_k \boldsymbol{Q}_{au} \\ \beta_k \boldsymbol{Q}_{ua} & \beta_k \boldsymbol{D}_0^u \end{pmatrix} \tag{18}$$

The matrix describes the behavior of a single component which is now no longer independent such that for an analysis the state of other components has to be taken into account. However, this implies that for an exact analysis a Markov chain has to be built that considers the state of all components which results in the so called *state space explosion*, an in the number of components exponentially growing state space. The Markov chain usually cannot be analyzed numerically if K is large as in real systems. In this case simulation or approximation techniques have to be applied. Simulation of Markov models is straightforward, whereas approximation techniques have to be developed specifically for this type of system.

It is also possible to apply the presented Markov processes in more general models. In [13] the availability models are integrated in a grid simulator to analyze the effect of failures on grid workloads. It is, of course, also possible to include the Markov models in such an environment. In contrast to [13], where correlations are not considered, the Markov models include correlation. It would be interesting to analyze the effect of correlations in the simulation model.

3 Modeling Availability Distributions with PHDs

To analyze the modeling quality of our Markov models we build models for nine different traces from the FTA. In this section we consider only the computation of the distribution of the lengths of availability and unavailability intervals. In the following section, correlation is added. In [13, 16] the same traces have been used and exponential, Weibull, log-normal, Gamma and Pareto distributions have been fitted to the traces using maximum likelihood estimators for the parameters.

The set of traces contains data from 22 HPC systems at Los Alamos National Laboratory (*lanl05*), from a computational grid platform in France (*g5k06*), and from desktop PCs at Microsoft (*microsoft99*) and the University of Notre Dame (*nd07cpu*). Additional measurements cover HTTP file requests to different web servers (*websites02*), results from different DNS servers (*ldns04*), data measured between all pairs of PlanetLab nodes (*pl05*), data from the Overnet peer-to-peer file-sharing system (*overnet03*) and from the Skype superpeer network (*skype06*). For detailed information about the traces we refer to [13, 16] and the references given therein.

For PHD fitting we used the EM algorithm that is described in [28]. This algorithm fits the parameters of Hyper-Erlang distributions, a subclass of PHDs, and is very efficient such that for most traces PHDs are computed in a few seconds. We generated PHDs with 2, 3, 5, 8 and 10 phases to examine possible differences in the fitting quality when increasing the number of phases. The number of parameters of a PHD is too large to present the detailed distributions here. We will make them available on our web page [1]. The parameters of the PHDs have all been computed with the publically available software *gfit* which can also been downloaded form the mentioned web page. The time to compute the PHD parameters with *gfit* depends on the number of phases, the number of elements in the trace and the structure of the trace. If a trace contains too many elements, i.e. more than 10^5 entries, it is possible to perform trace aggregation and to consider only an aggregated trace with about 10^3 or 10^4 aggregated entries. The approach is described in detail in [26]. In this way the computation of the parameters for PHDs with about 10 phases needs less than a minute and sometimes only a few seconds. Therefore it is possible to compute PHDs with an increasing number of phases and stop the generation if the likelihood value, which is implicitly computed during parameter fitting, does not significantly grow after adding a new phase. This approach works as long as the number of phases becomes not too large and is also used in the following examples.

3.1 Comparison of Lower Order Moments

First we compare lower order moments of the traces and lower order moments of the fitted distributions. Observe that the parameters of the distributions are not explicitly fitted to match lower order moments. It is, for example, possible to match with PHDs lower order moments exactly but this does not mean that the PHD is a good model for the data. Usually maximum likelihood based as done in the fitting algorithms used here and also in the estimators for the distributions from [13, 16] are much better. This means that the comparison of moments of the trace and the fitted distributions is not necessarily a criterion for the quality of the fitted distribution it is only a first hint. We compare the fitted PHDs with exponential, Weibull, log-normal and Gamma distributions from [13, 16].

Table 1 contains for each trace and all distributions that have been generated for the trace, the first moment, the coefficient of variation and the skewness. In each column that is headed by the name of the trace, the values are given in the above order for the trace and nine distributions. The value that is reached by a distribution and is closest to the corresponding value of the trace is printed in bold. PHDs including the exponential distribution usually match the first moment exactly, even if they are fitted with an ME algorithm, as done here, that does not explicitly use the moments for fitting. For the trace *ldns04* this is not true for the exponential distribution which we took from [13, 16]. For the Weibull, log normal and Gamma distribution the first moment is only approximated. In some case, the difference between first moment of the trace and the distribution is fairly large. The coefficient of variation is not exactly matched by the PHDs but they provide a much better approximation of this quantity than the other distributions do. For the third measure, the skewness, in two traces, *websites02* and *nd07cpu*, the Weibull distribution gives the best approximation. For the seven remaining traces the PHDs show the best approximations. These examples show that the flexibility of PHDs, which is a consequence of the large number of parameters, results in a better approximation of moments even if maximum likelihood methods are applied.

3.2 Comparison using Statistical Tests and Likelihood Values

It is not easy to evaluate the quality of a distribution with respect to trace data. In [13, 16] the AD and KS test are used for this purpose. As already mentioned that AD test cannot be applied for PHDs because the corresponding test statistics is not available. The KS test can be used also for PHDs. However, since the traces are very long, an application of the test to the whole trace will usually result in a low p value. This is known to be a general problem of *goodness-of-fit tests*. Therefore, in [13, 16] the following approach is adopted from [24]. 30 samples from a trace are selected randomly and the corresponding p-value is computed. This is repeated 1000 times and the average p-value is used as p-value for the distribution. This step is by no means exact, i.e., we cannot prove that resulting p-value is really the probability of rejecting the hypothesis that the trace is drawn from the distribution. However, the value can be used as an indicator of the matching quality of the distribution according to the trace and it can therefore be used to compare different distributions. As a second measure we compare the log-likelihood values of the different distributions, as defined in (4). Again, a larger log-likelihood value indicates a better match of distribution and trace. Almost all of the traces contain intervals with length zero. To obtain feasible log-likelihood values we deleted those values for the likelihood computation, because most of the distributions have density functions that are only defined for $x > 0$.

Table 1: Mean, coefficient of variation and skewness for the availability traces and the fitted distributions

	lanl05			**g5k06**			**microsoft99**		
Trace	1779.99	1.95	3.09	32.41	2.91	15.06	67.01	2.07	3.40
Exponential	**1779.99**	1.00	2.00	**32.41**	1.00	2.00	**67.01**	1.00	2.00
Weibull	1766.1	2.37	7.24	31.08	2.37	7.24	60.10	1.97	5.43
Log-Normal	4519.0	17.36	5287.4	84.62	18.67	6560.80	74.65	5.34	168.46
Gamma	**1785.9**	1.69	3.38	32.08	1.72	3.43	66.50	1.56	3.12
PH(2)	**1780.0**	1.74	2.87	**32.41**	1.50	2.46	67.01	1.90	3.24
PH(3)	**1780.0**	1.90	**3.22**	**32.41**	2.28	4.78	**67.01**	2.12	4.08
PH(5)	**1780.0**	**1.99**	3.52	**32.41**	**2.71**	**7.86**	**67.01**	1.89	3.19
PH(8)	**1780.0**	**1.99**	3.51	**32.41**	2.69	7.62	**67.01**	**2.09**	3.88
PH(10)	**1780.0**	**1.99**	3.52	**32.41**	2.70	7.66	**67.01**	**2.14**	4.14

	websites02			**pl05**			**ldns04**		
Trace	11.85	3.38	9.02	159.48	2.91	4.91	140.93	1.37	1.24
Exponential	**11.85**	1.00	2.00	**159.49**	1.00	2.00	141.06	1.00	2.00
Weibull	8.69	2.52	**7.99**	120.62	4.45	20.25	153.01	2.18	6.34
Log-Normal	9.68	7.63	466.5	252.09	59.72	213160	389.34	15.06	3463.0
Gamma	11.99	1.80	3.59	157.61	2.24	4.47	141.35	1.60	3.20
PH(2)	**11.85**	2.18	3.49	**159.48**	2.19	3.36	**140.93**	1.55	2.54
PH(3)	**11.85**	2.97	5.66	**159.48**	2.66	4.14	**140.93**	1.60	2.65
PH(5)	**11.85**	2.41	3.96	**159.48**	2.97	**4.95**	**140.93**	1.47	1.90
PH(8)	**11.85**	**2.99**	5.72	**159.48**	2.63	3.84	**140.93**	1.44	1.71
PH(10)	**11.85**	2.92	5.12	**159.48**	**2.93**	4.61	**140.93**	**1.43**	**1.60**

	overnet03			**nd07cpu**			**skype06**		
Trace	2.29	2.02	8.03	13.73	4.37	25.49	16.27	2.12	4.81
Exponential	**2.29**	1.00	2.00	**13.73**	1.00	2.00	**16.27**	1.00	2.00
Weibull	2.22	1.18	2.56	10.31	2.61	**8.41**	15.10	1.62	4.08
Log-Normal	1.95	1.27	5.86	15.18	11.20	1439.0	16.93	3.28	45.15
Gamma	2.30	1.05	2.97	13.85	1.83	3.65	16.32	1.37	2.74
PH(2)	**2.29**	1.96	5.66	**13.73**	2.07	3.33	**16.27**	2.01	4.26
PH(3)	**2.29**	1.86	5.07	**13.73**	3.16	6.40	**16.27**	2.08	4.56
PH(5)	**2.29**	**2.03**	**8.24**	**13.73**	**3.31**	7.11	**16.27**	2.02	4.29
PH(8)	**2.29**	1.89	5.38	**13.73**	3.02	5.87	**16.27**	**2.13**	**4.93**
PH(10)	**2.29**	2.05	8.48	**13.73**	3.18	6.45	**16.27**	2.14	5.04

Table 2: p-Values and Log-likelihood for the availability traces

	lanl05		**g5k06**		**microsoft99**		**websites02**		**pl05**	
Exponential	0.007	−168618	0.015	−1318033	0.003	−2738184	0.000	−166122	0.000	−151356
Weibull	0.465	−155798	0.513	−1117466	0.318	−2493940	0.098	−119874	0.092	−101085
Log-Normal	0.498	−155941	0.421	−1123175	0.402	−2444803	0.211	−112550	0.190	−97496
Gamma	0.368	−156660	0.455	−1134256	0.212	−2534195	0.068	−127251	0.045	−105054
PH(2)	0.426	−156271	0.307	−1156746	0.484	−2433653	0.241	−113191	0.273	−96302
PH(3)	0.512	−155309	0.493	−1118104	0.487	−2428660	0.205	−109868	0.344	−93023
PH(5)	0.520	−155107	0.525	−1106988	0.485	−2416022	**0.253**	−106771	0.313	−92563
PH(8)	**0.521**	−155104	**0.530**	−1105644	**0.491**	−2408099	0.249	−105215	**0.379**	−91214
PH(10)	**0.521**	**−155103**	**0.530**	**−1105569**	0.490	**−2407495**	0.237	**−104139**	**0.379**	**−90672**

	ldns04		**overnet03**		**nd07cpu**		**skype06**			
Exponential	0.012	−1328996	0.053	−61094	0.000	−485623	0.044	−110711		
Weibull	0.360	−1227189	0.081	−59959	0.388	−349636	0.427	−103576		
Log-Normal	0.390	−1232037	0.168	−52889	0.456	−335876	0.516	−101241		
Gamma	0.329	−1227126	0.062	−60986	0.182	−369163	0.291	−105507		
PH(2)	0.449	−1206368	0.094	−55642	0.365	−349772	0.486	−101547		
PH(3)	0.487	−1202157	**0.204**	−51213	0.477	−335631	**0.530**	−101341		
PH(5)	0.487	−1197068	0.172	−50424	**0.516**	−333845	0.523	−100659		
PH(8)	0.497	−1192887	0.168	−50398	0.501	−332763	0.524	−100603		
PH(10)	**0.499**	**−1190530**	0.197	**−49886**	**0.516**	**−330375**	0.527	**−100576**		

Table 2 shows the results for the availability traces. For each distribution and trace the p-value (left entry) and the log-likelihood value (right entry) are listed. For both values a higher value indicates a better approximation. The largest values for each trace are printed in bold face. Note, that for exponential, Weibull, log-normal and Gamma the p-values slightly differ from the ones presented in [16] due to the random selection of the samples for the KS-test, although we used the same distributions. As already noticed in [13, 16] the exponential distribution fails to model the availability interval length. Even if the first moment is matched exactly, this is not sufficient for a good fitting. The Weibull, log-normal and Gamma distribution are much better than the exponential distribution and show a relatively high p-value in most cases. However, it can also be noticed that a PHD with only 2 phases is most of the time better than the other distributions in terms of the p-value and in terms of the log-likelihood value. Increasing the number of states of the PHD usually results in higher p-values and log-likelihood values. In all cases the highest log-likelihood value is achieved with 10 phases, the largest number of phases we consider in our study. A further increase of the number of phases would probably increase the likelihood-value but it can be seen that the difference between 8 and 10 phases is usually marginal such that a further increase in the number of states is probably not worth the additional effort for parameter fitting and later analysis.

We now perform the same steps for the length of unavailability intervals. Table 3 contains the mean, coefficient of variation and skewness of the traces and the corresponding distributions. From the moments of the traces it can be seen that some of the unavailability traces have a very irregular behavior, i.e., a high coefficient of variation and skewness. This indicate that the traces are harder to model which also comes out of our analysis and confirms the results from [13, 16]. As for the availability traces the PHDs behave better than the other distributions. The mean values are exactly matched and the coefficient of variation is also usually quite well approximated by PHDs. However, in some cases, a larger number of states are necessary. For trace like *lanl05*, *g5k06* or *websites02* PHDs with 2 or 3 states are not sufficient. For the other distributions even more problems become visible. Especially for the Weibull and log-normal distribution the maximum likelihood estimators sometimes result in a distribution with a mean that is far away from the mean of the trace which usually indicates a low fitting quality, even if the mean might not be the major measure for heavy tailed distributions. In general the comparison of the moments for the unavailability traces confirms our observations from the availability traces. PHDs are better than the other distributions.

Table 4 shows the p-values of the KS test and the log-likelihood values for the unavailability traces which also confirm the observations from the availability traces. Note, that some log-likelihood values are $-\infty$ because of some very large values in the trace for which the corresponding density of the distribution becomes zero during computation. As already indicated by the moments, the unavailability traces appear to be more difficult to model, e.g. for the traces *websites02* and *overnet03* none of the distributions reached a p-value above the significance level of 0.05. Of course, the same can be observed in [13, 16]. This does not necessarily mean that the modeling quality is bad, it can also imply that the test statistics is not appropriate. It should be noted that for the unavailability traces, the difference between the log-likelihood values of PHDs and the other distributions is sometimes very large, much larger than for the availability traces. In particular for the traces *g5k06* and *websites02* the PHDs achieve a much better log-likelihood value. Since we consider the logarithm of the likelihood value, the difference is really significant and should become visible if the distributions are applied as part of a larger model.

Table 3: Mean, coefficient of variation and skewness for the unavailability traces and the fitted distributions

	lanl05			**g5k06**			**microsoft99**		
Trace	5.88	13.32	43.96	7.41	8.13	26.26	16.49	2.82	8.52
Exponential	5.92	1.00	2.00	**7.41**	1.00	2.00	**16.49**	1.00	2.00
Weibull	3.43	1.83	4.90	2.36	3.97	**16.74**	14.05	1.76	4.59
Log-Normal	2.88	2.55	24.27	1.52	11.20	1439.0	13.54	3.12	39.63
Gamma	**5.87**	1.62	3.24	7.58	2.29	4.59	16.34	1.47	2.95
PH(2)	**5.88**	4.40	8.99	**7.41**	3.02	4.59	**16.49**	1.76	3.01
PH(3)	**5.88**	7.41	18.27	**7.41**	5.29	8.98	**16.49**	1.71	2.87
PH(5)	**5.88**	11.03	34.18	**7.41**	3.77	5.80	**16.49**	2.40	4.99
PH(8)	**5.88**	11.81	38.30	**7.41**	5.42	9.28	**16.49**	2.25	4.16
PH(10)	**5.88**	**11.92**	**39.02**	**7.41**	**5.69**	9.97	**16.49**	**2.91**	**9.20**

	websites02			**pl05**			**ldns04**		
Trace	1.18	19.46	111.03	49.61	5.44	15.10	8.61	2.40	8.62
Exponential	**1.18**	1.00	2.00	**49.61**	1.00	2.00	**8.61**	1.00	2.00
Weibull	0.83	1.59	3.97	25.55	3.77	**15.35**	7.96	1.66	4.20
Log-Normal	0.62	1.61	8.98	30.0	20.09	8163.6	9.53	3.70	61.97
Gamma	**1.19**	1.41	2.82	49.91	2.18	4.36	**8.60**	1.40	2.80
PH(2)	**1.18**	2.58	4.95	**49.61**	2.42	3.72	**8.61**	1.75	3.28
PH(3)	**1.18**	2.30	4.34	**49.61**	3.49	5.67	**8.61**	2.26	6.60
PH(5)	**1.18**	9.74	42.74	**49.61**	3.10	4.87	**8.61**	2.24	6.37
PH(8)	**1.18**	9.64	42.13	**49.61**	3.43	5.53	**8.61**	**2.42**	**9.29**
PH(10)	**1.18**	**15.77**	**83.72**	**49.61**	**4.53**	9.01	**8.61**	2.43	9.57

	overnet03			**nd07cpu**			**skype06**		
Trace	11.98	3.07	3.66	4.25	14.77	33.72	14.31	2.11	6.26
Exponential	12.00	1.00	2.00	**4.25**	1.00	2.00	**14.31**	1.00	2.00
Weibull	7.78	2.70	8.88	1.43	2.18	6.34	13.42	1.66	4.20
Log-Normal	5.47	4.95	136.4	0.81	2.00	14.08	18.11	4.35	95.51
Gamma	12.08	1.86	**3.71**	4.22	1.89	3.78	14.27	1.41	2.83
PH(2)	**11.98**	2.72	4.31	**4.25**	5.91	10.03	**14.31**	1.94	4.66
PH(3)	**11.98**	2.52	3.97	**4.25**	9.56	17.17	**14.31**	**2.07**	**5.32**
PH(5)	**11.98**	3.47	5.98	**4.25**	7.79	13.59	**14.31**	2.05	5.10
PH(8)	**11.98**	3.20	4.86	**4.25**	**10.56**	**19.35**	**14.31**	1.94	4.28
PH(10)	**11.98**	**3.18**	4.45	**4.25**	10.06	18.23	**14.31**	2.03	4.94

Table 4: p-Values and Log-likelihood for the unavailability traces

	lanl05		**g5k06**		**microsoft99**		**websites02**		**pl05**	
Exponential	0.000	$-\infty$	0.000	$-\infty$	0.005	-1877456	0.000	$-\infty$	0.000	-118859
Weibull	0.235	-47131	0.010	-141961	0.059	-1689962	0.001	-34361	0.043	-70312
Log-Normal	0.500	-42343	0.048	-59979	$\mathbf{0.088}$	-1613592	0.007	-20081	0.098	-65821
Gamma	0.054	-53716	0.005	-247270	0.063	-1741644	0.004	$-\infty$	0.023	-76303
PH(2)	0.439	-43297	0.097	-28955	0.035	-1642806	0.000	-25135	0.142	-65725
PH(3)	0.495	-42290	0.065	13855	0.059	-1581405	0.001	-15247	0.100	-61914
PH(5)	0.504	-42097	$\mathbf{0.214}$	49853	0.055	-1538586	0.001	-11705	$\mathbf{0.171}$	-60887
PH(8)	$\mathbf{0.520}$	-41647	0.208	65074	0.077	-1502368	0.002	-6372	0.151	-59677
PH(10)	$\mathbf{0.520}$	$\mathbf{-41635}$	0.192	$\mathbf{69969}$	0.073	$\mathbf{-1496229}$	0.003	$\mathbf{-3513}$	0.149	$\mathbf{-59086}$

	ldns04		**overnet03**		**nd07cpu**		**skype06**	
Exponential	0.047	-508926	0.000	-123349	0.000	$-\infty$	0.071	-99340
Weibull	0.434	-465792	0.004	-82798	0.033	-136837	0.307	-92225
Log-Normal	0.509	-455891	0.014	-73824	0.152	-85521	0.200	-91270
Gamma	0.325	-476444	0.006	-90446	0.004	-197933	0.302	-93535
PH(2)	0.519	-459526	0.002	-68353	0.135	-91851	0.119	-93205
PH(3)	$\mathbf{0.525}$	-455154	0.009	-63184	0.147	-81006	0.290	-89834
PH(5)	0.523	-454443	0.006	-60325	0.213	-66613	0.341	-87384
PH(8)	0.522	-454168	0.011	-58344	0.213	-64242	$\mathbf{0.354}$	-87009
PH(10)	0.523	$\mathbf{-453874}$	0.013	$\mathbf{-55797}$	$\mathbf{0.232}$	$\mathbf{-61733}$	0.351	$\mathbf{-86974}$

3.3 Analyzing Availability Models for Components

With the distributions simple availability models for components can be built as in Eq. (7). We use this approach to analyze components from the traces *g5k06* and *pl05*.

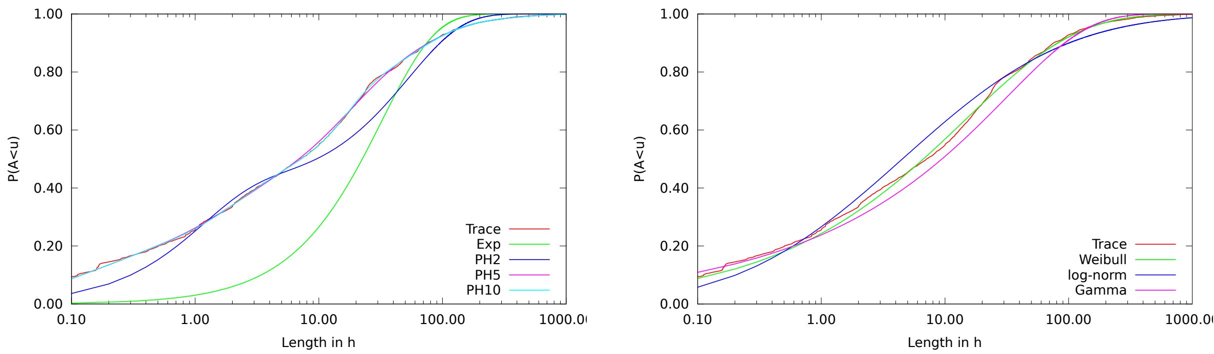

Figure 1: Cumulative distribution function of the availability trace *g5k06*, the fitted PHDs (left side) and the fitted remaining distributions (right side).

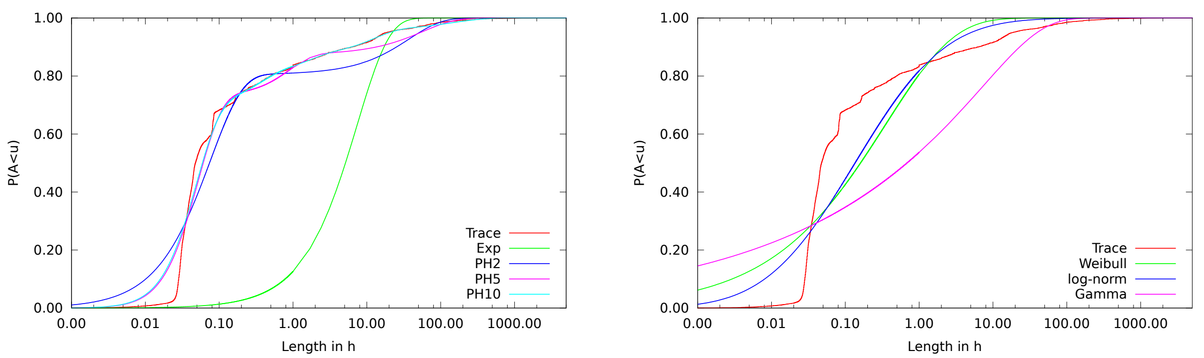

Figure 2: Cumulative distribution function of the unavailability trace *g5k06*, the fitted PHDs (left side) and the fitted remaining distributions (right side).

We begin with the trace *g5k06*. The Figures 1 and 2 show the distribution functions of the length of the availability phases for the trace and the different distributions. For the distribution of the length of availability phases, the exponential distribution is a bad approximation and the PHDs with 5 and 10 states and the Weibull distribution show an almost perfect matching with the trace. This, of course, corresponds to the values of the log-likelihood functions. For the unavailability intervals, the distribution is much harder to fit which can be seen by the varying log-likelihood values for the different distributions. The cumulative distribution functions in Figure 2 clearly show the problems. None of the distributions is able to match probability of very small durations of unavailability intervals. They are overestimated by all distributions except the exponential distribution which, however, differs completely from the distribution of the trace. After 0.1, the cumulative distribution functions of the trace and the PHDs with 5 and 10 states are very similar, but the difference for small values exists

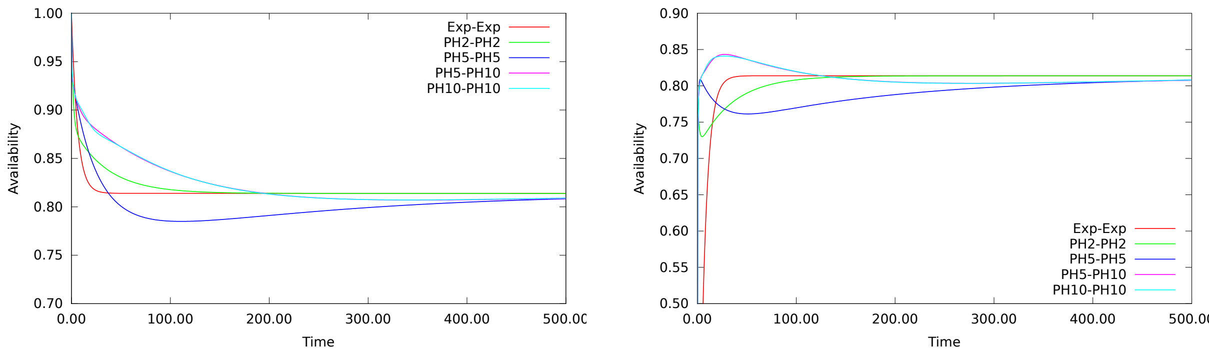

Figure 3: Availability of a component from the trace *g5k06* starting from at the beginning of an availability (left) or unavailability (right) phase.

also for these distributions.

The steady state availability of a component depends only on the mean duration of availability and unavailability phases. Since the PHDs, including the exponential distribution, match the mean of the trace exactly, they all result in a steady state availability of 0.814 for a component of *g5k06*. If the Weibull distribution is used to model the length of availability and unavailability intervals, the mean availability becomes 0.929, for the log-normal distribution it becomes 0.982, and for the Gamma distribution it equals 0.809. These results show significant differences in the computed steady state availability. The availability of a component during the measurement interval of the trace corresponds to the results for the PHDs which use the exact mean values from the trace. For transient analysis, the behavior of the PHDs, of course, differs, although they all reach eventually the same steady state value. In Figure 3 we show the availability of a component if the observation starts at the beginning of an availability or unavailability phase. *Exp-Exp* means that availability and unavailability phases are modeled by exponential distributions and *PHx-PHy* means that the availability is modeled by a PHD with x states and the unavailability by a PHD with y states. It can be seen that the behavior, in particular at the beginning of the interval, depends strongly on the distributions. This is not surprising since the log-likelihood values of the distributions differ significantly for the unavailability. The models *Exp-Exp* and *PH2-PH2* quickly reach the steady state availability, whereas in the other cases, the availability is below steady state availability after some time, even if started at the beginning of an availability interval. The influence of the model for the unavailability interval is larger than the influence of the model for the availability interval. The reason is that the log-likelihood value of the PHD with 10 phases is much larger than the log-likelihood value of all other distributions for modeling unavailability intervals. This can also bee seen by an inspection of the availability, the two cases *PH5-PH10* and *PH10-PH10* show almost the same behavior and differ from the other cases. Thus, for modeling the availability, a PHD with 5 phases is sufficient whereas the unavailability should be modeled with 10 or even more phases.

Results for the trace *pl05* are similar. In Figure 4 the cumulative distribution function of the availability interval length of the trace and the distributions is shown. Again it can be seen that most distributions have a larger probability of generating very short intervals. For larger times the PHDs with 5 and 10 states again show a very good matching, better than the remaining distributions. Exactly the same holds for the unavailability distribution which is shown in Figure 5.

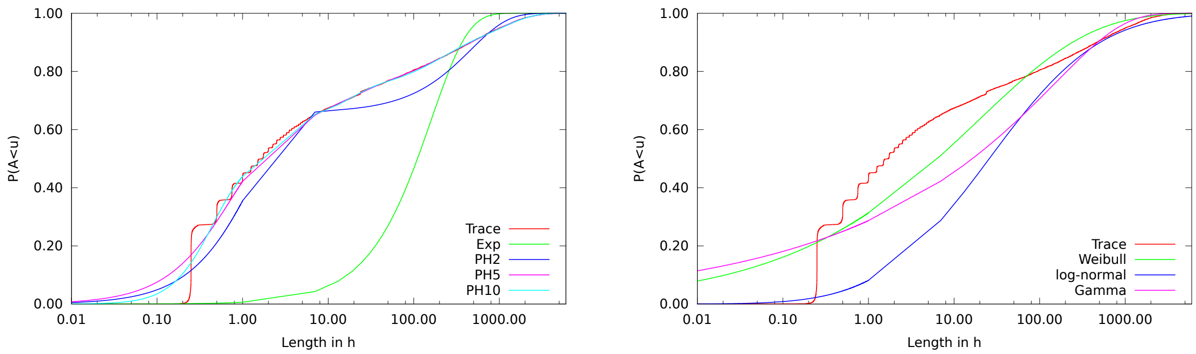

Figure 4: Cumulative distribution function of the availability trace *pl05*, the fitted PHDs (right side) and the fitted remaining distributions (left side).

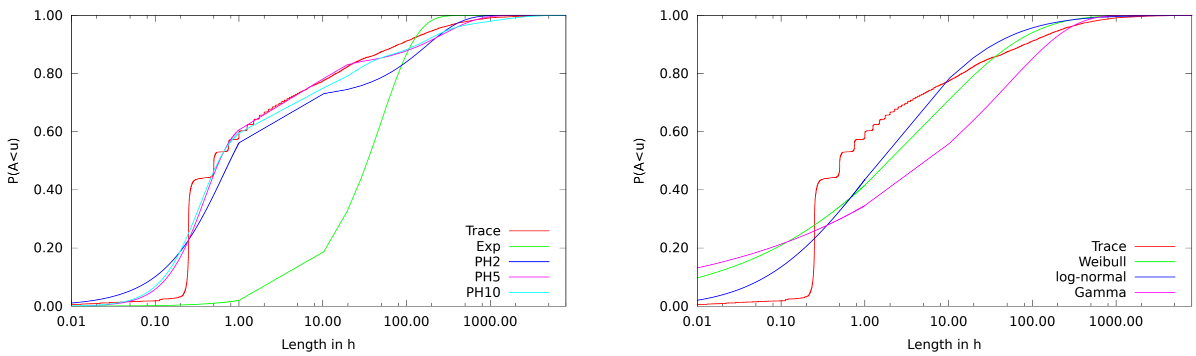

Figure 5: Cumulative distribution function of the unavailability trace *pl05*, the fitted PHDs (right side) and the fitted remaining distributions (left side).

4 Analyzing and Modeling Correlations

In Section 2 it is shown how four different correlation values from sequences of alternating availability and unavailability intervals can be computed. The correlation between succeeding availability times $\hat{\rho}_a$, the correlation between succeeding unavailability times $\hat{\rho}_u$ and correlation between intervals a_i, u_i ($\hat{\rho}_{au}$) and u_i, a_{i+1} ($\hat{\rho}_{ua}$) are considered.

Table 5 shows the different correlation values for the nine traces. Note, that the traces for some components start with an unavailability interval and not with an availability interval as defined in Sect. 2. In these cases we used pairs a_i, u_{i+1} for the computation of $\hat{\rho}_{au}$ and u_i, a_i for $\hat{\rho}_{ua}$. Furthermore, for some components only few, i.e. two or less, availability or unavailability intervals were available, which does not allow one to compute correlation values for those components. We ignored these components for the computation. From Table 5 it becomes visible that the correlation in the traces differs considerably which is not surprising because of the different systems the traces have been measured from. For the traces *g5k06* and *pl05* all four values are too large to be neglected in modeling. Consequently, we used the fitted PHDs from Sect. 3 and tried to incorporate the correlation as described in Sect. 2. In the first step matrix $\boldsymbol{Q}$ from (12) has to be constructed. The

Table 5: Correlation Coefficients according to Eq. (11)

Trace	$\hat{\rho}_{\mathbf{a}}$	$\hat{\rho}_{\mathbf{u}}$	$\hat{\rho}_{\mathbf{au}}$	$\hat{\rho}_{\mathbf{ua}}$
g5k06	0.211820	0.058601	0.139210	0.043502
lanl05	0.087681	-0.054249	-0.022377	-0.020657
ldns04	-0.063043	-0.136805	0.008968	0.008649
microsoft99	-0.001572	-0.066955	0.050020	0.023496
nd07cpu	0.155535	0.016765	0.042394	0.029314
overnet03	0.050644	-0.012634	0.024182	0.027758
pl05	0.166929	0.111189	0.055080	0.099938
skype06	0.006294	-0.067591	0.021445	0.033601
websites02	0.200295	0.038603	0.052859	0.023160

submatrices $\boldsymbol{D}_0^a$ and $\boldsymbol{D}_0^u$ are given by the two PHDs modeling the availability and unavailability times. $\boldsymbol{Q}_{au}$ and $\boldsymbol{Q}_{ua}$ have to be modeled such that the $\xi_.$ values from (14) (15) approximate the measured correlation values from Table 5. Additionally, $\boldsymbol{Q}_{au}$ and $\boldsymbol{Q}_{ua}$ have to observe the constraints on row sums and the initial vectors of the distributions mentioned in Section 2. For the construction of matrix $\boldsymbol{Q}$ we used PHDs of the same order, although other combination would be possible as well and obtain $\boldsymbol{Q}_{au}$ and $\boldsymbol{Q}_{ua}$ after solving an alternating least squares optimization problem as described in Section 2.

Table 6: Correlation Coefficients for fitted MMAPs

	$\hat{\rho}_{\mathbf{a}}$	$\hat{\rho}_{\mathbf{u}}$	$\hat{\rho}_{\mathbf{au}}$	$\hat{\rho}_{\mathbf{ua}}$
g5k06	0.211820	0.058601	0.139210	0.043502
PH(2)	0.045647	0.072663	0.142246	0.142199
PH(3)	0.211819	0.017055	0.108412	0.046229
PH(4)	0.211819	0.043256	0.139288	0.044761
PH(8)	0.181407	0.026113	0.107164	0.043751
pl05	0.166929	0.111189	0.055080	0.099938
PH(2)	0.155021	0.108354	0.245462	0.139143
PH(3)	0.299886	0.076783	0.106937	0.100046
PH(4)	0.104537	0.111739	0.054987	0.099937
PH(6)	0.135391	0.111083	0.055007	0.099937

As it turned out, the correlation values are difficult to approximate. Table 6 shows the values for the traces *g5k06* and *pl05* and MMAPs resulting from PHDs of different order. As one can see, it was not possible to obtain a good approximation for all correlation values in most cases. Two or three values could be approximated usually, but for the remaining values the approximation is only modest. The best results have been obtained for the trace *g5k06* using two PHDs of order 4. The resulting matrix is

$$
\begin{aligned}
\boldsymbol{Q} &= \begin{bmatrix} \boldsymbol{D}_0^a & \boldsymbol{Q}_{au} \\ \boldsymbol{Q}_{ua} & \boldsymbol{D}_0^u \end{bmatrix} \text{ where} \\
\boldsymbol{D}_0^a &= \begin{bmatrix} -0.006 & 0 & 0 & 0 \\ 0 & -0.044 & 0 & 0 \\ 0 & 0 & -0.645 & 0 \\ 0 & 0 & 0 & -7.913 \end{bmatrix} \\
\boldsymbol{Q}_{au} &= \begin{bmatrix} 0.001 & 0.005 & 0 & 0 \\ 0.001 & 0 & 0 & 0.043 \\ 0.043 & 0 & 0 & 0.602 \\ 0 & 0 & 7.461 & 0.452 \end{bmatrix} \\
\boldsymbol{Q}_{ua} &= \begin{bmatrix} 0.002 & 0 & 0.004 & 0 \\ 0.051 & 0 & 0.029 & 0 \\ 0 & 0 & 1.307 & 0 \\ 1.044 & 12.480 & 1.041 & 3.086 \end{bmatrix} \\
\boldsymbol{D}_0^u &= \begin{bmatrix} -0.006 & 0 & 0 & 0 \\ 0 & -0.080 & 0 & 0 \\ 0 & 0 & -1.307 & 0 \\ 0 & 0 & 0 & -17.651 \end{bmatrix}
\end{aligned}
\tag{19}
$$

that can be used for further analysis.

In a second step we further extend this matrix to describe the behavior of a component when k other components are available, i.e. we construct the matrices $\boldsymbol{Q}[k]$ from (18). The trace *g5k06* contains intervals for 1258 different components (recall, that we removed some components from the trace because only few intervals were available for those components), so we cannot present all matrices $\boldsymbol{Q}[k]$, but will only give an example. Let $k = 500$. Then, we obtain $m_a(k) = 690$ and $m_u(k) = 606$, i.e. 690 availability intervals and 606 unavailability intervals end when 500 components are available. For the conditional moments we get $\hat{\mu}_a(k) = 13.438$ and $\hat{\mu}_u(k) = 2.1933$, which results in factors $\alpha_k = 2.4115$ and $\beta_k = 3.2794$ which shows a significant reduction in the length of availability and unavailability intervals compared to the mean values. Combining these factors with the matrices from (19) we get

$$
\boldsymbol{Q}[k] = \begin{bmatrix} 2.4115 \cdot \boldsymbol{D}_0^a & 2.4115 \cdot \boldsymbol{Q}_{au} \\ 3.2794 \cdot \boldsymbol{Q}_{ua} & 3.2794 \cdot \boldsymbol{D}_0^u \end{bmatrix}.
$$

The computation of the matrices $\boldsymbol{Q}[k]$ is straightforward, however, for a real system like the one that generated the trace *g5k06* we obtain a huge number of matrices making a numerically analysis cumbersome.

5 Conclusion

In this paper we described the modeling of different availability traces by means of phase type distributions. One outcome of the experiments is that PHDs usually provide a better model for the data than other commonly used distributions like Weibull, log-normal or Gamma distributions do. Since nowadays efficient algorithm for the parameter fitting of phase type distributions are available, as long as the number of phases remains moderate, they are a good alternative. Another advantage of phase type distributions is the possibility to analyze the resulting model analytically or numerically which is not possible with most other distributions. This can be exploited to compute transient of stationary dependability measures. Furthermore, correlation can be integrated by extending phase type distributions to processes. Approaches how to fit the parameters to match the coefficient of correlation are presented in the paper. However, it also turns out that correlation in real traces is hard to analyze and match such that additional research is necessary to find reliable methods to match the measured correlations by a Markov model. The problems in correlation fitting mainly result from the solution of Eq. 15 in the paper which is non-linear in the elements of the matrices $\boldsymbol{Q}_{au}$ and $\boldsymbol{Q}_{ua}$. The alternating least squares solver we used for the computation of the matrices often seems to get stuck in some local optimum. However, there might be other optimization methods that are more reliable for this kind of problems.

The paper considers only the parameterization of phase type distributions and not the analysis of the resulting models. It would be interesting to compare the results of phase type distributions if they are used as parts of larger dependability models. Furthermore, the distributions can also be used in simple queueing models to analyze the performance under failure or the performability.

References

[1] PHDs for failure traces from the FTA. http://www-ls4.cs.tu-dortmund.de/failure, Feb. 2014.

[2] S. Asmussen, O. Nerman, and M. Olsson. Fitting phase-type distributions via the EM-algorithm. *Scand. J. Stat.*, 23(4):419–441, 1996.

[3] A. Bobbio, A. Horváth, and M. Telek. Matching three moments with minimal acyclic phase type distributions. *Stochastic Models*, 21(2-3):303–326, 2005.

[4] P. Buchholz. An EM-algorithm for MAP fitting from real traffic data. In P. Kemper and W. H. Sanders, editors, *Computer Performance Evaluation / TOOLS*, volume 2794 of *Lecture Notes in Computer Science*, pages 218–236. Springer, 2003.

[5] P. Buchholz, I. Felko, and J. Kriege. Transformation of acyclic phase type distributions for correlation fitting. In *ASMTA*, pages 96–111, 2013.

[6] P. Buchholz, P. Kemper, and J. Kriege. Multi-class Markovian arrival processes and their parameter fitting. *Perform. Eval.*, 67(11):1092–1106, 2010.

[7] P. Buchholz and J. Kriege. Aggregation of Markovian models - an alternating least squares approach-. In *QEST'12*, pages 43–52. IEEE Computer Society, 2012.

[8] M. J. Faddy. Phase-type distributions for failure times. *Mathematical and Computer Modelling*, 22(10):63–70, 1995.

[9] J. Faulin, A. A. Juan, S. S. M. Alsina, and J. E. Ramirez-Marquez, editors. *Simulation Methods for Reliability and Availability of Complex Systems*. Springer Series in Reliability Engineering. Springer, 2010.

[10] Q.-M. He and M. Neuts. Markov arrival processes with marked transitions. *Stochastic Processes and their Applications*, 74:37–52, 1998.

[11] A. Horváth and M. Telek. Matching more than three moments with acyclic phase type distributions. *Stochastic Models*, 23:167–194, 2007.

[12] G. Horváth, M. Telek, and P. Buchholz. A MAP fitting approach with independent approximation of the inter-arrival time distribution and the lag-correlation. In *QEST*, pages 124–133. IEEE CS Press, 2005.

[13] B. Javadi, D. Kondo, A. Iosup, and D. H. J. Epema. The failure trace archive: Enabling the comparison of failure measurements and models of distributed systems. *J. Parallel Distrib. Comput.*, 73(8):1208–1223, 2013.

[14] M. Johnson and M. Taaffe. Matching moments to phase distributions: Nonlinear programming approaches. *Stochastic Models*, 2(6):259–281, 1990.

[15] A. Klemm, C. Lindemann, and M. Lohmann. Modeling IP traffic using the batch Markovian arrival process. *Perform. Eval.*, 54(2):149–173, 2003.

[16] D. Kondo, B. Javadi, A. Iosup, and D. H. J. Epema. The failure trace archive: Enabling comparative analysis of failures in diverse distributed systems. In *CCGRID*, pages 398–407. IEEE, 2010.

[17] J. Kriege and P. Buchholz. An empirical comparison of MAP fitting algorithms. In *MMB/DFT*, volume 5987 of *Lecture Notes in Computer Science*, pages 259–273. Springer, 2010.

[18] W. P. Krijnen. Convergence of the sequence of parameters generated by alternating least squares algorithms. *Computational Statistics and Data Analysis*, 51:481–489, 2006.

[19] A. M. Law and W. D. Kelton. *Simulation modeling and analysis*. Wiley, 2000.

[20] M. F. Neuts. A versatile Markovian point process. *Journal of Applied Probability*, 16:764–779, 1979.

[21] M. F. Neuts. *Matrix-geometric solutions in stochastic models*. Johns Hopkins University Press, 1981.

[22] M. F. Neuts and K. S. Meier. On the use of phase type distributions in reliability modelling of systems with two components. *OR Spectrum*, 2(4):227–234, 1981.

[23] D. M. Nicol, W. H. Sanders, and K. S. Trivedi. Model-based evaluation: From dependability to security. *IEEE Trans. Dependable Sec. Comput.*, 1(1):48–65, 2004.

[24] D. Nurmi, J. Brevik, and R. Wolski. Modeling machine availability in enterprise and wide-area distributed computing environments. In *Euro-Par*, volume 3648 of *Lecture Notes in Computer Science*, pages 432–441. Springer, 2005.

[25] H. Okamura, T. Dohi, and K. S. Trivedi. A refined EM algorithm for PH distributions. *Perform. Eval.*, 68(10):938–954, 2011.

[26] A. Panchenko and A. Thümmler. Efficient phase-type fitting with aggregated traffic traces. *Perform. Eval.*, 64(7-8):629–645, 2007.

[27] R. A. Sahner, K. S. Trivedi, and A. Puliafito. *Performance and Reliability Analysis of Computer Systems*. Kluwer, 1996.

[28] A. Thümmler, P. Buchholz, and M. Telek. A novel approach for phase-type fitting with the EM algorithm. *IEEE Trans. Dep. Sec. Comput.*, 3(3):245–258, 2006.

[29] N. Yigitbasi, M. Gallet, D. Kondo, A. Iosup, and D. H. J. Epema. Analysis and modeling of time-correlated failures in large-scale distributed systems. In *11th IEEE/ACM International Conference on Grid Computing*, pages 65–72. IEEE, 2010.