Acyclic Phase-Type Distributions in Fault Trees

Pepijn Crouzen
Department of Computer Science
Saarland University
Saarbrücken, Germany
Email: crouzen@cs.uni-sb.de

Reza Pulungan
Jurusan Ilmu Komputer
Universitas Gadjah Mada
Yogyakarta, Indonesia
Email: pulungan@ugm.ac.id

***Abstract*—Acyclic phase-type distributions can be used to describe the time until a basic event in a fault tree occurs. We show in this paper that the top event of a fault tree built from such basic events is also acyclic phase-type distributed. We then apply a recently developed acyclic phase-type minimization algorithm to effectively combat the state-space explosion problem in a dynamic fault tree analysis.**

I. Introduction

Fault trees (FT) [1] are a popular and well-established formalism used for dependability modeling. A FT describes how different combinations of component failures lead to a system failure. A FT is a directed acyclic graph in which the sources are *basic events*, which generally denote component failures, and the other vertices are *gates*, which connect the basic events and generally describe the failures of subsystems. A FT has one sink, the *top event*, which describes the failure of the entire system. The most important measure of interest for a FT is the *unreliability*: the probability that the top event occurs given specified probabilities that the basic events occur.

A FT with m basic events encodes a boolean function $f : \{0,1\}^m \to \{0,1\}$, where boolean vector $\vec{b} = (b_1, \cdots, b_m)$ describes the operational status of the basic events, and $f(\vec{b})$ describes the operational status of the top event. When b_i is equal to 0, the i-th basic event has not yet occurred; when it is equal to 1, it has occurred. Similarly, $f(\vec{b}) = 0$ denotes that the system is operational, whereas $f(\vec{b}) = 1$ means that the top event has happened, i.e. the system has failed. In this paper, we consider only *coherent* FTs, namely FTs where system failure is irreversible: the failure of a basic event cannot cause the entire system to go from failed to operational. In terms of boolean functions, this means that the function f is monotonically increasing with respect to the partial order $(\{0,1\}^m, \succeq)$ with $\vec{a} \succeq \vec{b} \Leftrightarrow \forall i \in [1,m] \cdot b_i = a_i \vee b_i$. The most effective analysis technique for FTs is based on an encoding of FTs in binary decision diagrams (BDD) [2].

Instead of considering FTs where each basic event $B_i, i \in [1,m]$ has a fixed failure probability $P(B_i = 1)$, we can also consider FTs where basic events are random processes $(B_i^{(t)})_{t \in \mathbb{R}_{\geq 0}}$, which fail after a randomly distributed delay. The top event of the FT is then also described by a random process $(T^{(t)})_{t \in \mathbb{R}_{\geq 0}}$. The analysis strategy remains the same: for a particular time point t, we calculate $P(B_i^{(t)} = 1)$ for each basic event and use these probabilities in the standard way to find the probability $P(T^{(t)} = 1)$ of the system having failed at time point t. However, we may now also consider systems where the order in which basic events occur matters for the system reliability. To allow such modelling, [3] introduced *dynamic fault trees* (DFTs), which extend FTs with several *dynamic gates*, which take into account the order in which events happen. DFT analysis is more difficult than FT analysis because we must consider the random process $(T^{(t)})_{t \in \mathbb{R}_{\geq 0}}$, which described the failure of the top event.

In [3], a continuous-time Markov chain (CTMC) representation $X^{(t)}$ of the random process $T^{(t)}$ for a DFT is created monolithically from the DFT description, assuming that the basic events fail after exponentially distributed delays. The states of the CTMC describe the operational status of the basic events and, if necessary, the order in which they occurred. Using the DFT description, we can find the set of states $\mathcal{S}_Y$ of X in which the top event has occurred. We then find the reliability of the system by numerically solving [4] the CTMC to find the probability $P(X^{(t)} \in \mathcal{S}_Y) = P(T^{(t)} = 1)$ for some time-point t.

Because we must record the order in which basic events occur, the state space of the CTMC grows exponentially in the number of basic events. This problem is also known as the *state-space explosion*. One way of mitigating the state-space explosion is by using *compositional aggregation* [5]. Here the CTMC is constructed by composing interactive models, which represent the syntactic elements of the DFT. After each compositional step, the resulting model is minimized, i.e. a smaller model that is still equivalent to the original model is identified. The equivalence used is *weak bisimulation*, which considers only the structure of the models. Although compositional aggregation mitigates the problem of state-space explosion, it does not avoid it.

II. Contribution

Instead of assuming that basic events are exponentially distributed, we may also use acyclic phase-type (APH) distributions [6] as suggested in [5]. APH distributions have three useful properties: (1) they can be represented by acyclic CTMCs, (2) they are topologically dense and can be used to

This work is supported by the German Research Council (DFG) as part of the Transregional Collaborative Research Center "Automatic Verification and Analysis of Complex Systems" SFB/TR 14 AVACS). See www.avacs.org for more information. The research leading to these results has received funding from the EuropeanCommunity's Seventh Framework Programme under grant agreement n° 214755.

The work was done while Reza Pulungan was with Saarland University.

approximate any probability distribution, and (3) we can generally find the minimal representation of an APH distribution [7]. Now we are equipped to model any possible basic event. However, as the size of the CTMC representation of APH distributions may be large, the state-space explosion problem is worsened by using APH distributions. We give a short introduction into APH distributions and their representations in Section III.

We will show that we can combine compositional DFT analysis with APH minimization to fight back against the state-space explosion. In DFTs, we can often find independent subtrees (or modules) [8] that consist only of non-dynamic gates. We show in Section IV that the top event of such a subtree, where the basic events occur after APH-distributed delays, occurs also after an APH-distributed delay. We can then replace the subtree by a basic event and, more importantly, minimize the associated APH representation. A short overview of the APH minimization algorithm is given in Section V.

Finally we show in Section VI, that applying APH minimization instead of weak-bisimulation minimization in the analysis of DFTs can have a huge impact on memory requirements and computation times. In essence, we are improving the compositional aggregation technique of [5] by using a better minimization strategy whenever possible. In this case, APH minimization is better than weak-bisimulation minimization as APH minimization computes the smallest representation regardless of structural properties, while weak-bisimulation minimization only computes the smallest representation that is structurally equivalent to the original.

III. Acyclic Phase-Type Distributions

Consider a continuous-time Markov chain $(X^{(t)})_{t\in\mathbb{R}_{\geq 0}}$ with finite state space $\mathcal{S} = \{1, \ldots, n, n+1\}$ for some $n \in \mathbb{N}$, initial distribution $\vec{\pi}$, and infinitesimal generator matrix $\mathbf{Q}$. The non-diagonal entries of $\mathbf{Q}$ are such that, for an infinitesimal time interval Δ, for all $t \in \mathbb{R}_{\geq 0}$, and for all $i, j \in \mathcal{S}$ with $i \neq j$, we have $P(X^{(t+\Delta)} = j \mid X^{(t)} = i) = \Delta \cdot \mathbf{Q}(i,j)$. For diagonal entries of $\mathbf{Q}$ we have $\mathbf{Q}(i,i) = -\sum_{j\in\mathcal{S}, i\neq j} \mathbf{Q}(i,j)$. The derivatives of the state probabilities are given by the Chapman-Kolmogorov equation:

$$\frac{d}{dt}P(X^{(t)} = j) = \sum_{i\in\mathcal{S}} P(X^{(t)} = i)\mathbf{Q}(i,j). \tag{1}$$

Given the initial distribution, the solution of the system of differential equations (1) is, for $t \in \mathbb{R}_{\geq 0}$:

$$\Big(P(X^{(t)} = 1), \ldots, P(X^{(t)} = n+1)\Big) = \vec{\pi}e^{\mathbf{Q}t}.$$

We can interpret a CTMC as an edge-labelled graph with vertices $\mathcal{S}$ and an edge from state $i \in \mathcal{S}$ to $j \in \mathcal{S}$ labelled with $\mathbf{Q}(i,j)$, whenever $\mathbf{Q}(i,j) > 0$.

Assume now that state $n+1$ is absorbing, i.e. for all $t, t' \in \mathbb{R}_{\geq 0}$, $t < t'$: $P(X^{(t')} = n+1 \mid X^{(t)} = n+1) = 1$, and all other states are transient, i.e. there is a path from each transient state to the absorbing state $n+1$. Let $\mathbf{A}$ be the $n \times n$ submatrix of $\mathbf{Q}$ that corresponds to the transient states, and $\vec{\alpha}$ be the row vector that corresponds to the initial probabilities of the transient states. Let $\vec{e}$ be a column vector of appropriate dimension whose entries are all equal to 1.

Definition 1: A random variable Z is distributed according to a *phase-type (PH) distribution* with *representation* $(\vec{\alpha}, \mathbf{A})$ if its distribution function is given by

$$F(t) = \Pr(Z \leq t) = \begin{cases} 1 - \vec{\alpha}e^{\mathbf{A}t}\vec{e}, & t \in \mathbb{R}_{\geq 0}, \\ 0, & \text{otherwise}, \end{cases} \tag{2}$$

A matrix of the form of $\mathbf{A}$ is called a *PH-generator*, and $\mathrm{PH}(\vec{\alpha}, \mathbf{A})$ denotes the PH distribution of $(\vec{\alpha}, \mathbf{A})$. The *size of the representation* $(\vec{\alpha}, \mathbf{A})$ is the dimension of $\mathbf{A}$. From now on, we only deal with irreducible representations, i.e. those that do not contain superfluous states, which means that the vector $\vec{\alpha}e^{\mathbf{A}t}$ is strictly positive for all $t > 0$. A PH distribution has uncountably many different representations [6]. We associate a CTMC X with each PH representation $(\vec{\alpha}, \mathbf{A})$ by adding the absorbing state $n+1$ to $\vec{\alpha}$ and $\mathbf{A}$ to obtain $\vec{\pi}$ and $\mathbf{Q}$ respectively. We then have from (2), for $t \in \mathbb{R}_{\geq 0}$:

$$F(t) = 1 - \vec{\alpha}e^{\mathbf{A}t}\vec{e} = \left[\vec{\pi}e^{\mathbf{Q}t}\right](n+1) = P(X^{(t)} = n+1).$$

Here $\left[\vec{\pi}e^{\mathbf{Q}t}\right](n+1)$ denotes the $(n+1)$-th entry of vector $\vec{\pi}e^{\mathbf{Q}t}$.

Beside its distribution function, a PH distribution can be characterized by its Laplace-Stieltjes transform (LST):

$$\tilde{f}(s) = -\vec{\alpha}(s\mathbf{I} - \mathbf{A})^{-1}\mathbf{A}\vec{e} + 1 - \vec{\alpha}\vec{e} = \frac{p(s)}{q(s)}, s \in \mathbb{R}_{\geq 0}, \tag{3}$$

where $\mathbf{I}$ is the n-dimensional identity matrix. The transform is a rational function for some polynomials $p(s)$ and $q(s) \neq 0$. When the LST is expressed in irreducible ratio, the degree of the numerator $p(s)$ is no more than the degree of the denominator $q(s)$. The degrees of the two polynomials are equal only when $1 - \vec{\alpha}\vec{e} > 0$ [9].

An *acyclic phase-type (APH) distribution* has at least one representation that, under some permutation of its state space, has a triangular representation matrix. Such representations are called APH representations. A *triangular minimal representation* of an APH distribution is an APH representation with the least number of states. The *triangular order* of an APH distribution [10] is the size of its triangular minimal representation. The degree of the denominator polynomial of the LST of a PH distribution, expressed in irreducible ratio, is called the *algebraic degree of the distribution*. The zeros of the denominator polynomial are called the *poles* of the LST. APH distributions are exactly those distributions whose LST is rational and has only real poles [11].

For $\lambda_n \geq \lambda_{n-1} \geq \cdots \geq \lambda_1 > 0$, let PH-generator

$$\begin{bmatrix} -\lambda_1 & \lambda_1 & 0 & \cdots & 0 \\ 0 & -\lambda_2 & \lambda_2 & \cdots & 0 \\ 0 & 0 & -\lambda_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & -\lambda_n \end{bmatrix}$$

be denoted by $\mathbf{Bi}(\lambda_1, \lambda_2, \cdots, \lambda_n)$. A representation with PH-generator of this form is called an *ordered bidiagonal representation* (OBR).

Theorem 2 ([12]): Let $(\vec{\alpha}, \mathbf{A})$ be an acyclic phase-type representation and let $-\lambda_1, -\lambda_2, \cdots, -\lambda_n$ be the eigenvalues of $\mathbf{A}$. W.l.o.g., assume $\lambda_n \geq \lambda_{n-1} \geq \cdots \geq \lambda_1 > 0$. Then there exists a unique representation $(\vec{\beta}, \mathbf{Bi}(\lambda_1, \lambda_2, \cdots, \lambda_n)$ such that both represent the same phase-type distribution.

The *spectral polynomial algorithm* (SPA) [13] is an efficient algorithm to transform a given APH representation to an OBR that has at most the size of the original APH representation. The SPA has complexity $\mathcal{O}(n^3)$, where n is the size of the given APH representation.

IV. Probability Distribution of a Fault Tree

We now consider m independent PH distributions running in parallel with CTMC representations $X_1, \ldots, X_m$, which in turn have state spaces $\mathcal{S}_1 = \{1, \ldots, n_1 + 1\}, \ldots, \mathcal{S}_m = \{1, \ldots, n_m + 1\}$, initial distributions $\vec{\pi}_1, \ldots, \vec{\pi}_m$, and infinitesimal generator matrices $\mathbf{Q}_1, \ldots, \mathbf{Q}_m$, respectively. The probability that each of the CTMCs is in a particular state is described by a CTMC with state space $\mathcal{S} = \prod_{i=1}^m \mathcal{S}_i$, initial distribution $\vec{\pi} = \bigotimes_{i=1}^m \vec{\pi}_i$, and infinitesimal generator matrix $\mathbf{Q} = \bigoplus_{i=1}^m \mathbf{Q}_i$ [14][1]. We then have, for $t \geq 0$ and for all $1 \leq i \leq m$, $x_i \in \mathcal{S}_i$:

$$P(X_1^{(t)} = x_1, \ldots, X_m^{(t)} = x_m) = P(X^{(t)} = (x_1, \ldots, x_m)).$$

Each state of the combined CTMC X is then a vector of states of the CTMCs of the composite PH representations. Let $\vec{b} : \mathcal{S} \to \{0,1\}^m$ be a function that, for each state $s = (x_1, \ldots, x_m)$, returns a boolean vector that specifies which component PH distributions are in the absorbing state and which are in a transient state. If the i-th component distribution has absorbed, i.e. if $x_i = n_i + 1$, then $[\vec{b}(s)](i) = 1$ otherwise $[\vec{b}(s)](i) = 0$.

We have seen that a fault tree with m basic events encodes a boolean function $f : \{0,1\}^m \to \{0,1\}$. Assume that these m basic events occur after delays that are distributed according to the m PH distributions discussed above. For a state s of the combined CTMC X the vector $\vec{b}(s)$ then encodes which basic events have occurred and which haven't. We then find that the top event of the fault tree has occurred in a state s if and only if $f(\vec{b}(s)) = 1$. We define the function $g : \mathcal{S} \to \{0,1\}$, which describes for each state of X whether the top event has failed or not, as $g(s) = f(\vec{b}(s))$.

Let $\mathcal{S}_N$ be the set of states where the top event has not yet occurred, i.e. $\mathcal{S}_N = \{s \in \mathcal{S} \mid g(s) = 0\}$. Naturally, $\mathcal{S}_Y = \mathcal{S} \setminus \mathcal{S}_N$ is the set of all states where the top event has occurred. Because we only consider coherent fault trees, the probability to move from a state where the top event has occurred to a state where it has not occurred is always zero. System failure is irreversible. Thus,

$$\forall s \in \mathcal{S}_Y, s' \in \mathcal{S}_N \cdot \mathbf{Q}(s, s') = 0, \tag{4}$$

$$\forall s \in \mathcal{S}_Y \cdot \sum_{s' \in \mathcal{S}_Y} \mathbf{Q}(s, s') = \sum_{s' \in \mathcal{S}} \mathbf{Q}(s, s') = 0. \tag{5}$$

Here (5) is derived from the fact that the rows of $\mathbf{Q}$ add up to zero.

We are interested in the probability that the top event occurs within time $t \geq 0$: $P(g(X^{(t)}) = 1) = P(X^{(t)} \in \mathcal{S}_Y)$. From (1), (4) and (5) we can derive the following system of ordinary differential equations for the top event probability and state probabilities for states $s' \in \mathcal{S}_N$:

$$\begin{aligned}
\frac{d}{dt} P(g(X^{(t)}) = 1) &= \sum_{s \in \mathcal{S}, s' \in \mathcal{S}_Y} P(X^{(t)} = s) \mathbf{Q}(s, s'), \\
&= \sum_{s \in \mathcal{S}} P(X^{(t)} = s) \sum_{s' \in \mathcal{S}_Y} \mathbf{Q}(s, s'), \\
&= \sum_{s \in \mathcal{S}_N} P(X^{(t)} = s) \sum_{s' \in \mathcal{S}_Y} \mathbf{Q}(s, s'). \qquad (6) \\
\frac{d}{dt} P(X^{(t)} = s') &= \sum_{s \in \mathcal{S}} P(X^{(t)} = s) \mathbf{Q}(s, s'), \\
&= \sum_{s \in \mathcal{S}_N} P(X^{(t)} = s) \mathbf{Q}(s, s'). \qquad (7)
\end{aligned}$$

The initial condition is derived from the initial probability distribution of X:

$$P(g(X^{(0)}) = 1) = \sum_{s \in \mathcal{S}_Y} \vec{\pi}(s) = 1 - \sum_{s \in \mathcal{S}_N} \vec{\pi}(s), \tag{8}$$

$$P(X^{(0)} = s') = \vec{\pi}(s'). \tag{9}$$

From (6), (7), (8), and (9) we see that it is not necessary to consider individual states within set $\mathcal{S}_Y$. In fact we can construct a smaller CTMC $Y^{(t)}$ from $X^{(t)}$ by collapsing all states in $\mathcal{S}_Y$ to one state s_y, while still preserving the top event probability. The state space of Y is then $\mathcal{S}_N \cup \{s_y\}$. The initial probability distribution $\vec{\pi}_Y$ is given by (8) and (9), i.e. for $s \in \mathcal{S}_N$, $\vec{\pi}_Y(s) = \vec{\pi}(s)$, and for s_y we have $\vec{\pi}_Y(s_y) = \sum_{s' \in \mathcal{S}_Y} \vec{\pi}(s')$. The infinitesimal generator matrix $\mathbf{Q}_Y$ is derived from (6) and (7):

$$\mathbf{Q}_Y(s, s') = \begin{cases} \mathbf{Q}(s, s'), & \text{if } s, s' \in \mathcal{S}_N, \\ \sum_{s'' \in \mathcal{S}_Y} \mathbf{Q}(s, s''), & \text{if } s \in \mathcal{S}_N, s' = s_y, \\ 0, & \text{otherwise.} \end{cases}$$

Applying (1) to CTMC Y we find, for all states $s' \in \mathcal{S}_N$:

$$\begin{aligned}
\frac{d}{dt} P(Y^{(t)} = s_y) &= \sum_{s \in \mathcal{S}_N \cup \{s_y\}} P(Y^{(t)} = s) \mathbf{Q}(s, s_y), \\
&= \sum_{s \in \mathcal{S}_N} P(Y^{(t)} = s) \mathbf{Q}(s, s_y), \\
&= \sum_{s \in \mathcal{S}_N} P(X^{(t)} = s) \sum_{s' \in \mathcal{S}_Y} \mathbf{Q}(s, s'), \\
&= \frac{d}{dt} P(g(X^{(t)}) = 1),
\end{aligned}$$

[1] $\bigotimes$ and $\bigoplus$ are the Kronecker product and the Kronecker sum operators, respectively.

$$\begin{aligned}\frac{d}{dt}P(Y^{(t)}=s') &= \sum_{s\in\mathcal{S}_N\cup\{s_y\}} P(Y^{(t)}=s)\mathbf{Q}(s,s'),\\ &= \sum_{s\in\mathcal{S}_N} P(Y^{(t)}=s)\mathbf{Q}(s,s')\\ &+\ P(Y^{(t)}=s_y)\sum_{s\in\mathcal{S}_Y}\mathbf{Q}(s,s'),\\ &= \sum_{s\in\mathcal{S}_N} P(Y^{(t)}=s)\mathbf{Q}(s,s'),\\ &= \frac{d}{dt}P(X^{(t)}=s').\end{aligned}$$

The initial distributions of X and Y trivially match as well.

Theorem 3: The top event of a coherent fault tree whose basic events occur after delays governed by independent PH distributions itself occurs after a delay governed by a PH distribution.

We have shown how to construct the CTMC Y above. It is straightforward that state s_y is absorbing and that the states in $\mathcal{S}_N$ are transient. Hence, Y is a CTMC representation of a PH distribution and the probability to be in state s_y at time t is the same as the probability that the top event of the fault tree happened within time t.

Corollary 4: The top event of a coherent fault tree whose basic events occur after APH distributed delays itself occurs after an APH distributed delay.

Corollary 4 follows directly from the fact that the Kronecker sum of several acyclic CTMCs is itself acyclic. Hence the CTMC X is acyclic, and then so is CTMC Y.

V. Minimization

In this section, we discuss an algorithm for reducing the size of APH representations [7]. The reduction procedure is roughly as follows: (1) Given an APH distribution with representation $(\vec{\alpha},\mathbf{A})$, it is transformed into an OBR $(\vec{\beta},\mathbf{Bi}(\lambda_1,\cdots,\lambda_n))$ by using SPA, without increasing its size. (2) A smaller representation is obtained by eliminating unnecessary states from the OBR. The unnecessary states can be decided as explained below.

The Laplace-Stieltjes transform (LST) of an exponential distribution with rate λ is $\tilde{f}(s)=\frac{\lambda}{s+\lambda}$. Let $L(\lambda)=\frac{s+\lambda}{\lambda}$, then the LST of OBR $(\vec{\beta},\mathbf{Bi}(\lambda_1,\cdots,\lambda_n))$ can be written as

$$\tilde{f}(s)=\frac{\vec{\beta}(1)+\vec{\beta}(2)L(\lambda_1)+\cdots+\vec{\beta}(n)L(\lambda_1)\cdots L(\lambda_{n-1})}{L(\lambda_1)L(\lambda_2)\cdots L(\lambda_n)}.$$

Theorem 5 ([7]): If for some $1\le i\le n$, polynomial $\vec{\beta}(1)+\vec{\beta}(2)L(\lambda_1)+\cdots+\vec{\beta}(i)L(\lambda_1)\cdots L(\lambda_{i-1})$ is divisible by $L(\lambda_i)$, then there exists a unique vector $\vec{\delta}$ such that

$$\begin{aligned}&\mathrm{PH}(\vec{\beta},\mathbf{Bi}(\lambda_1,\cdots,\lambda_n))\\ &\quad=\mathrm{ME}(\vec{\delta},\mathbf{Bi}(\lambda_1,\cdots,\lambda_{i-1},\lambda_{i+1},\cdots,\lambda_n),\vec{e}).\end{aligned}$$

If vector $\vec{\delta}$ is substochastic (i.e. $\vec{\delta}\vec{e}\le 1$), then

$$\begin{aligned}&\mathrm{PH}(\vec{\beta},\mathbf{Bi}(\lambda_1,\cdots,\lambda_n))\\ &\quad=\mathrm{PH}(\vec{\delta},\mathbf{Bi}(\lambda_1,\cdots,\lambda_{i-1},\lambda_{i+1},\cdots,\lambda_n)).\end{aligned}$$

$\mathrm{ME}(\cdot,\cdot,\cdot)$ in the theorem is a matrix-exponential distribution.

Based on Theorem 5, an algorithm to reduce the size of a given APH representation can be constructed. For a more exhaustive discussion of the algorithm, we refer to [7].

Two relevant properties of the algorithm are the minimality on triangular-ideal APH distributions and almost minimality on minimum and maximum operations.

An APH distribution is triangular ideal if and only if its triangular order is equal to its algebraic degree. Given an APH representation whose APH distribution is triangular ideal—no matter how large the size of the representation is—the reduction algorithm is certain to produce a minimal representation [7].

Minimum and maximum operations are two of the most common operations on probability distributions. The set of PH distributions is closed under each of these operations. Neuts [6] describes the CTMC representations of the PH distributions produced by the operations. In the context of this paper, the important fact is that these operations correspond precisely to OR and AND gates, respectively, in FTs.

Now, applied on the results of minimum and maximum operations on triangular-ideal APH distributions, the reduction algorithm almost always produces minimal representations [7], in the sense that the measure of APH distributions for which it doesn't produce minimal representations is zero. Hence, given two arbitrary triangular-ideal APH distributions, regardless of the size of their representations, we are almost certain that the APH distribution of their minimum or maximum is also triangular ideal, and hence its APH representation is reducible to minimal size by applying the reduction algorithm. We conjecture that almost minimality applies to all top events of FTs whose basic events are APH distributed. The implementation of the algorithm is tentatively called APHMIN.

VI. Case Study

As a case study, we consider a fault tree with three basic events connected to a 2-out-of-3 voting gate, which is also the top event (see Fig. 1 (left)). The three basic events occur after delays governed by Erlang distributions with shape 2 and rates 1, 2, and 3, respectively. Erlang distributions are a subset of APH distributions (see Fig. 1 (right)).

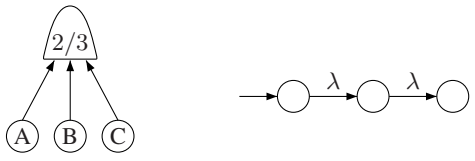

Fig. 1. A fault tree (left) and a CTMC representation of an Erlang distribution with shape 2 and rate λ (right).

The PH distribution of the top event of the fault tree as well as its CTMC representation can be constructed as described in Section IV. The constructed CTMC representation has 21 states and 45 transitions. This CTMC is minimal with respect to weak bisimulation, however by using APH minimization

we can obtain a minimal representation, which consists of 14 states and 22 transitions. In either case, generating the CTMC representation takes less than one second.

Although the difference in the size of the state spaces seems small, we will see that even such a modest difference can have a huge impact on memory and time requirements in a compositional setting. To illustrate this we consider the fault-tolerant parallel processor (FTPP) case study from [5]. We use the CORAL tool [15] to analyze the DFT models of the FTPP.

We investigate what happens if some of the exponentially distributed basic events in the FTPP case study are replaced by APH-distributed basic events that correspond to the fault tree in Fig. 1 (left). We replace one, two, or three of the *network element* basic events by the voting gate in the FTPP DFT. In the first experiment, we run CORAL without modification. In the second experiment, we replace the voting gate with a basic event that occurs after the same APH distribution. However, this time the associated APH representation is first minimized by using APHMIN and then the CORAL tool is run on the resulting DFT. The results are shown in Table I, which shows the state space of the largest model encountered during compositional aggregation and the total analysis time. The times reported here are faster than the ones in [15] since the CORAL tool has since been improved.

TABLE I
RESULT OF THE CASE STUDY

#	Tool	States	Transitions	Time (s)	Unreliability
1	CORAL	1,672	12,303	10.37	$1.13 \cdot 10^{-7}$
	APHMIN	1,119	7,410	10.42	$1.13 \cdot 10^{-7}$
2	CORAL	59,739	598,524	24.52	$3.21 \cdot 10^{-4}$
	APHMIN	26,006	219,310	14.14	$3.21 \cdot 10^{-4}$
3	CORAL	1,777,955	21,895,068	14,047.99	0.209
	APHMIN	507,067	5,010,000	367.71	0.209

From the table, we can observe that using the larger PH representation greatly deteriorates the performance of the CORAL tool. When the minimized PH representation is used, we see that the 7-state difference in the basic event model can lead to a difference of 1.2 million states in the last DFT model and a reduction of computation time by a factor greater than 38, from almost four hours to just over six minutes.

VII. FUTURE WORK

As future work, we would like to fully integrate the APHMIN tool and the CORAL tool, identify dynamic fault trees that are APH distributed, and prove the conjecture stated in Section V.

REFERENCES

[1] W. E. Veseley, F. F. Goldberg, N. H. Roberts, and D. F. Haasl, *Fault Tree Handbook*. United States Nuclear Regulatory Commision, 1981, vol. (NUREG-0492).

[2] O. Coudert and J. Madre, "Fault tree analysis: 10^{20} prime implicants and beyond," in *Annual Proceedings on Reliability and Maintainability Symposium, 1993*. IEEE Computer Society, 1993, pp. 240–245.

[3] J. B. Dugan, S. J. Bavuso, and M. A. Boyd, "Dynamic fault-tree models for fault-tolerant computer systems," *IEEE Transactions on Reliability*, vol. 41, no. 3, pp. 363–377, 1992.

[4] W. J. Stewart, *Introduction to the Numerical Solution of Markov Chains*. Princeton University Press, 1994.

[5] H. Boudali, P. Crouzen, and M. Stoelinga, "A compositional semantics for Dynamic Fault Trees in terms of Interactive Markov Chains," in *Proc. of the 5th International Symposium on Automated Technology for Verification and Analysis*. LNCS, 2007, pp. 441–456.

[6] M. F. Neuts, *Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach*. Dover, 1981.

[7] R. Pulungan and H. Hermanns, "Acyclic minimality by construction—almost," in *Fifth International Conference on the Quantitative Evaluaiton of Systems (QEST 2009), 13-16 September, 2009, Budapest, Hungary*. IEEE Computer Society, 2009.

[8] Y. Dutuit and A. Rauzy, "A linear-time algorithm to find modules of fault trees," *IEEE Transactions on Reliability*, vol. 45, no. 3, pp. 422–425, 1996.

[9] C. A. O'Cinneide, "Characterization of phase-type distributions," *Communications in Statistics: Stochastic Models*, vol. 6, no. 1, pp. 1–57, 1990.

[10] ——, "Triangular order of triangular phase-type distributions," *Communications in Statistics: Stochastic Models*, vol. 9, no. 4, pp. 507–529, 1993.

[11] ——, "Phase-type distributions and invariant polytopes," *Advances in Applied Probability*, vol. 23, no. 43, pp. 515–535, 1991.

[12] A. Cumani, "Canonical representation of homogeneous Markov processes modelling failure time distributions," *Microelectronics and Reliability*, vol. 2, no. 3, pp. 583–602, 1982.

[13] Q.-M. He and H. Zhang, "Spectral polynomial algorithms for computing bi-diagonal representations for phase type distributions and matrix-exponential distributions," *Stochastic Models*, vol. 2, no. 2, pp. 289–317, 2006.

[14] W. J. Stewart, K. Atif, and B. Plateau, "The numerical solution of stochastic automata networks," *European Journal of Operational Research*, vol. 86, no. 3, pp. 503 – 525, 1995.

[15] H. Boudali, P. Crouzen, and M. Stoelinga, "Coral – a tool for compositional reliability and availability analysis," in *ARTIST WS: Tool Platforms for ES Modelling, Analysis and Validation*, 2007.