

# A Rigorous, Compositional, and Extensible Framework for Dynamic Fault Tree Analysis

Hichem Boudali, *Member, IEEE*, Pepijn Couzen, and Mariëlle Stoelinga

**Abstract**—Fault trees (FT) are among the most prominent formalisms for reliability analysis of technical systems. Dynamic FTs extend FTs with support for expressing dynamic dependencies among components. The standard analysis vehicle for DFTs is state-based, and treats the model as a CTMC, a continuous-time Markov chain. This is not always possible, as we will explain, since some DFTs allow multiple interpretations.

This paper introduces a rigorous semantic interpretation of DFTs. The semantics is defined in such a way that the semantics of a composite DFT arises in a transparent manner from the semantics of its components. This not only eases the understanding of how the FT building blocks interact. It also is a key to alleviate the state explosion problem. By lifting a classical aggregation strategy to our setting, we can exploit the DFT structure to build the smallest possible Markov chain representation of the system. The semantics - as well as the aggregation and analysis engine is implemented in a tool, called CORAL. We show by a number of realistic and complex systems that this methodology achieves drastic reductions in the state space.

**Index Terms**—Fault trees, Reliability, Compositionality, Formal models, Framework.

## I. INTRODUCTION

RELIABILITY ENGINEERING is an important activity in the design of today's computer and communication systems. For safety critical systems, such as airplanes and nuclear power plants, failures can be life threatening; for other applications, such as online ticket vending systems, failures often incur a high cost.

One of the most popular formalisms to model and analyze systems' reliability is the fault tree (FT) formalism [27]. Dynamic fault trees (DFT) [13], [8], [26] extend standard (or static) FTs by defining additional gates called dynamic gates. These gates allow the modeling of complex system components' behaviors and interactions which greatly increases the modeling capabilities of standard FTs. Like standard FTs, dynamic fault trees are a high-level formalism for computing reliability measures of computer-based systems. For over a decade now, DFTs have been experiencing a growing success among reliability engineers.

DFTs, like FTs, describe the system failure in terms of the failure of its components. A DFT is a tree (or rather a directed acyclic graph (DAG), since subtrees can be shared) in which the leaves are *basic events (BEs)* and the other elements are *gates*. A BE typically models the failure of a physical component and is governed by a probability distribution. In this paper, we consider exponential distributions and phase-type distributions, the latter allowing to approximate other probability distributions with arbitrary precision. Gates express how component failures induce system failures and are either static (AND, OR, and the

K/M voting gate) or dynamic (Priority AND, SPARE, and the Functional Dependency gate). DFTs are typically used to compute *system unreliability*, that is, the probability that the system fails during a specified period of time (usually called mission time) and under given conditions. Other measures, such as the average time until a failure occurs can be computed as well.

Despite their success, current DFT analysis methods have several (mutually related) drawbacks.

- 1) Existing analysis methods (most notably, the DIFTree method [21] implemented in analysis tools like Galileo [25] and Relex [23]) typically convert a DFT into a continuous-time Markov chain (CTMC) whose states are vectors of modes (up, failed, active, inactive) for each BE. Hence, the size of the state space is exponential in the number of basic events.
- 2) These methods impose rather severe syntactic restrictions on DFTs, greatly diminishing the modeling flexibility and power of DFTs. Most notably, DFT spare components must be BEs, whereas spare components in practice are often entire subsystems.
- 3) The DFT semantics is rather imprecise and the lack of formality has, in some cases, led to undefined behavior and misinterpretation of the DFT model.
- 4) DFTs lack comprehensive modular analysis. DIFTree uses a limited form of compositional analysis: it solves in a separate way all stochastically independent subtrees of a static gate, provided none of its ancestors in the tree is a dynamic gate. Then it combines, using Binary Decision Diagrams, their analysis results to obtain the result for the entire DFT. However, this method is not applicable to dynamic gates. In particular, those DFTs whose top-node is a dynamic gate can not be analyzed compositionally.
- 5) The current methods are difficult to extend or to modify.

In this paper, we present a framework for DFT analysis based on I/O-IMCs that greatly alleviates these drawbacks. I/O-IMCs are a powerful and versatile formalism to model and analyze stochastic system behavior, and have been used in an number of applications, ranging from telecommunication systems [18] to railway networks [3] and multiprocessor arrays [11]. I/O-IMCs extend CTMCs with input, output and internal actions, used for communication between several I/O-IMCs. They are equipped with a parallel composition operator, allowing one to build larger I/O-IMCs from smaller ones, and with powerful minimization (a.k.a. lumping) techniques to reduce the state space of an I/O-IMC.

The core of our methodology is a compositional semantics of DFTs in terms of I/O-IMCs. That is, we translate each DFT element (i.e. gate or BE) into one or more I/O-IMCs — obtaining these semantics turned out to be non-trivial, and required a careful re-examination and generalization of the concept of spare

This research has been partially funded by the Netherlands Organisation for Scientific Research (NWO) under FOCUS/BRICKS grant number 642.000.505 (MOQS); the EU under grant IST-004527 (ARTIST2) and FP7-ICT-2007-1 grant 214755 (QUASIMODO); and by the DFG/NWO bilateral cooperation programme under project number DN 62-600 (VOSS2)

activation. Then, the semantics of an entire DFT is obtained as the parallel composition of all DFT element I/O-IMCs. Since these I/O-IMCs semantics pin down the meaning of a DFT in a mathematically precise way, we lift drawback 3 mentioned above. Relatedly [10] presented a formal semantics in Z. The main difference between the formal specification in [10] and the formal specification used in this paper is that in our framework we use a process algebra-like formalism (i.e. I/O-IMC) which comes with two very powerful concepts, namely parallel composition and aggregation/minimization.

Since composing all element I/O-IMCs at once would give the same blow up as the DIFTree method, we use the compositional aggregation method to reduce the size of the models. That is, we compose two I/O-IMCs, hide actions that are no longer needed for communication with other components, and minimize them. We repeat this process until all element I/O-IMCs have been composed. While this method is still exponential in worst case, our experiments show that serious reductions (one or two orders of magnitude) are realized in practice. Thus, our methodology relieves drawback 1 mentioned above. Since this analysis method is fully compositional, that is, the analysis of a DFT (i.e. its underlying I/O-IMC) is obtained from analysis results of its components (i.e. lumped I/O-IMCs of sub-models), we also lift drawback 4.

We note that the order in which we compose these I/O-IMCs matters for the size of the intermediate I/O-IMC models, but not for the final result. We employ heuristics, based on the way individual I/O-IMC models communicate with each other, to obtain smart composition orders. These models have specific properties (acyclicity) that we exploit by tailored algorithms. Also, it turns out that our techniques require much lighter syntactic restrictions than existing DFT methodologies. Any subsystem can now be used as a dependent event and any activation-independent (see Section IV) can be used as a spare component. Hence, drawback 2 is alleviated. Finally, we show how the current DFT semantics can readily be extended or modified; we present extensions with inhibition, mutual exclusion and repair, thus addressing drawback 5.

We have implemented our DFT framework in a tool called *CORAL*. The tool derives all I/O-IMC models and composes them using the CADP tool-set [15], which is also used to compute the system reliability. We used *CORAL* to analyze nine case studies, including ones showing systems with spare and dependent event subsystems which are currently not supported by any other DFT tool and a system showing the need for non-determinism. We have compared our tool with Galileo and our experiments show that, in almost all cases, our tool is much faster and generates significantly smaller models (where we consider the largest model encountered during analysis).

This paper combines and extends the work previously carried out in [7] and [5]. In particular, our contributions over these papers are

- 1) a complete semantics of all DFT elements: Whereas [7] and [5] describe the semantics of DFT gates for a specific number of inputs, we cover here the general case, employing the IOIML notation. Moreover, we allow phase-type distributions as failure distributions of basic events;
- 2) a complete proof for the congruence theorem;
- 3) a more extensive set of case studies including systems with spare and dependent event subsystems which are not

currently supported by any other DFT tool and examples showing the need of non-determinism;

- 4) *CORAL*, our prototype tool for analyzing DFT using the I/O-IMC semantics. It employs the specialized I/O-IMC minimization algorithm from [12], yielding much faster computation times than [7] and [5].

a) *Organization of the paper:* The remainder of the paper is organized as follows: In Section II we introduce DFTs and in Section III we discuss I/O-IMCs. In Sections IV and V we present the formal DFT syntax and semantics respectively. In Section VI, we show, through three examples, how one can readily extend the existing DFT formalism. Section VII presents the compositional aggregation technique and Section VIII describes the *CORAL* tool. Finally, in Section IX, we present a number of case studies and Section X concludes the paper.

## II. DYNAMIC FAULT TREES

As described in Section I, DFTs and FTs are directed acyclic graphs describing the system failure in terms of the failure of its components. Their leaves are labeled with *basic events*, and non-leaves with *gates*.

1) *BE*: A BE, graphically depicted by a circle (see Figure 1(g)), typically represents the failure of a basic system component; its failure behavior is governed by a probability distribution. In order to describe these distributions using I/O-IMCs, this paper considers exponential and (acyclic) phase-type distributions, the latter allowing to approximate other probability distributions with arbitrary precision.

An exponential distribution has a parameter  $\lambda$  that represents the component's failure rate (i.e. number of failures per time unit). A BE has three modes of operation: *dormant*, *active*, and *failed*. In dormant (or standby) mode, the BE failure rate  $\lambda$  is reduced by a factor  $\alpha \in [0, 1]$  called *dormancy factor*. Thus, the BE failure rate in standby mode is  $\mu = \alpha\lambda$ . In active mode, the failure rate is unchanged and equals to  $\lambda$ . The dormancy is relevant when the BE is used as a spare (more details on spare BEs is provided below). In failed mode, the BE, as the name suggests, has failed and remains in that state (i.e. we do not consider repairable systems at this point).

Phase-type basic events (PHBE) are basic events that fail after a delay governed by a phase-type (PH) distribution [22] with a finite number of phases. The passive behavior of a PHBE is also described by a PH distribution with a finite number of phases. Activation of a PHBE is described by a function which links passive phases to active phases. When a PHBE is activated it moves from its current passive phase to the associated active phase. Note that a BE (with exponential distribution) is a special case of a PHBE where both active and passive distributions have only one phase.

2) *Gates*: Non-leaf elements are called *gates* and express how component failures induce system failures. Their graphical representation is given in Figures 1(a)-(f). Each gate has one or more inputs, corresponding to outputs of other elements, and exactly one output. It often represents or maps to a subsystem contained in the whole system, the top element representing the system failure. When the failure event of a BE or a gate occurs, we use the terms *failing*, *occurring*, or *firing* interchangeably.

Gates can either be static (AND, OR gate, and VOTING (also called K/M gate) or dynamic. Static gates (which are the only gates in static fault trees) are combinatorial: they are only sensitive

to the combinations of failures of their inputs and not to their order.

Dynamic gates allow the modeling of sequence dependencies (via the priority AND (PAND) gate), functional dependencies (functional dependency (FDEP) gate) and spare management and allocation (via the SPARE gate)<sup>1</sup>. Thus, DFTs enrich the FT formalism with powerful and yet easy-to-use modeling capabilities.

Below, we describe all DFT gates.

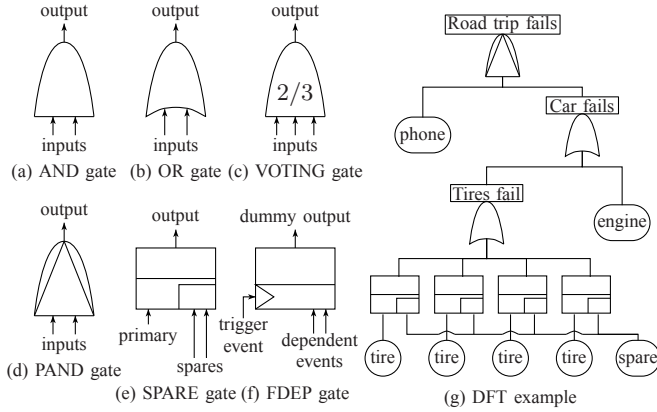


Fig. 1. DFT gates and example.

- 3) *AND gate*: The AND gate fails when all of its inputs fail.
- 4) *OR gate*: The OR gate fails when at least one of its inputs fails.
- 5) *VOTING gate*: A  $K/M$  VOTING gate fails when at least  $K$  (called the threshold) out of its  $M$  inputs fail.
- 6) *PAND gate*: The PAND gate fails when all its inputs fail and fail from left to right (as depicted in the figure) order.
- 7) *FDEP gate*: The functional dependency gate consists of a trigger event (i.e. a failure event) and a set of dependent events. When the trigger event occurs, it causes all the dependent components to become inaccessible or unusable. Essentially, once a dependent component is triggered, it is assumed to have failed. Dependent events, as originally defined in [13], need to be BEs. This restriction will be later lifted in our framework. All dependent events and the trigger event are considered to be inputs to the FDEP gate. The FDEP gate's output is a 'dummy' output (i.e. it is not taken into account during the calculation of the system failure probability).

8) *SPARE gate*: The SPARE gate has one primary input and zero (which is a degenerated case) or more alternate inputs called *spares*. The primary input of a SPARE gate is initially powered on (i.e. in active mode) and the alternate inputs are in standby mode. When the primary fails, it is replaced by the first available alternate input (which then switches from standby mode to active mode). This operation is called *spare activation* and causes the spare to switch from dormant to active mode. In turn, when this alternate input fails, it is replaced by the next available alternate input, and so on and so forth. Note that multiple spare gates can *share* a pool of spares. When the primary unit of any of the spare gates fails, it is replaced by the first available (i.e. not failed or not already taken by another spare gate) spare unit; which becomes, in turn, the active unit for that spare gate. The SPARE gate fails when the primary fails and all its spares are failed or unavailable.

If all SPARE inputs are BEs, two special cases arise depending on the spare's dormancy factor  $\alpha$ . If  $\alpha = 0$ , the spare is called

a *cold spare* and can not by definition fail before the primary. When  $\alpha = 1$ , the spare is called a *hot spare* and its failure rate is the same whether in standby or in active mode. If  $0 < \alpha < 1$ , the spare is called a *warm spare*.

**Example 1** Figure 1(g) shows a DFT modeling a road trip. Looking at the top PAND gate, we see that the road trip fails (i.e. we are stuck on the road) if the car fails after the mobile phone has failed; if the car fails first, then we can call the road services to tow the car and continue our journey. The car subsystem fails if either the engine fails or the tires subsystem fails. The car is equipped with a spare tire, which can be used to replace any of the primary tires; when a second tire fails, the tires subsystem fails, causing in turn a car failure. Thus, we model the tires subsystem by four spare gates, each having a primary tire and all sharing a spare tire. The spare tire is a cold spare, i.e. its failure rate is zero in standby mode.

#### A. Simultaneity and non-determinism

In earlier development of the DFT modeling formalism, the semantics (i.e. the model interpretation) of some DFT configurations, where FDEP gates are used, remained unclear. For instance, in Figure 2, the FDEP gate triggers (in both configurations) the failures of two basic events. Does this mean that the dependent events fail simultaneously and if so what is the state of the PAND gate in the left configuration and which spare gate gets the shared spare  $S$  in the right configuration? These examples were also discussed in [10], and we believe that this is an inherent non-determinism in these models. Whereas in [10], these special cases are dealt with by systematically removing the non-determinism by transforming it into a probabilistic (or deterministic) choice. In our framework, we allow non-determinism should this be intentional or unintentional. If the non-determinism was not intended, then its presence (which is easily detected) indicates that an error occurred during the model specification. Non-determinism could also be an inherent characteristic of the system being analyzed, and should therefore be explicitly modeled.




Fig. 2. The occurrence of non-determinism.

In the I/O-IMC formalism, the DFT configurations depicted in Figure 2 will be interpreted as follows: Whenever the dependent events failure has been triggered, then the trigger event (the cause) happened first and was then immediately (with no time elapsing) followed by the failure of the dependent events (the effect). This adheres to the classical *notion of causality*. Moreover, the dependent events fail in a non-deterministic order (i.e. essentially considering all combinations of ordering). In this case, the final I/O-IMC model is not a continuous-time Markov chain but rather a continuous-time Markov decision process (CTMDP), which can be analyzed by computing bounds of the performance measure of interest [2]. As an example, we have modeled and analyzed a simple non-deterministic case study (see Subsection IX) using the MRMC model-checker [19]. However, the conversion of I/O-IMCs to CTMDPs, which closely follows [17], has not yet been automated in the tool-chain.

<sup>1</sup>A fourth gate called 'Sequence Enforcing' gate introduced in [13] can be emulated using a cold spare gate.



### B. Lifting DFT restrictions

Previously, DFT required all inputs to a SPARE gate and all dependent events of an FDEP gate to be basic events. This restriction greatly diminishes the modeling power of DFTs, since it is very natural to have spare components that are comprised of multiple components or subsystems. To lift this restriction, we need to carefully reexamine the notion of spare activation.

For primaries and spares that are complex systems, we say that a BE  $b$  is a primary-BE (or just primary) of a SPARE gate  $G$  if  $b$  is contained in the subtree that constitutes the primary of  $G$ . This is the case if there exists a path from  $b$  to  $G$  whose last edge ends in the first input of  $G$ . Spare-BEs are defined analogously.

The basic idea behind spare activation is that all BEs that are primary-BEs of some SPARE gate are activated from the beginning. A BE that is a spare-BE of some SPARE gate gets activated as soon as one of its SPARE parents is activated. Since spares can be shared, a BE can have multiple SPARE parents.

When a BE is both a primary-BE and a spare-BE, activation is unclear: is this BE activated from the beginning or through the SPARE gate? To rule out such situations we require all primaries and all spares to be activation-independent subtrees. This means that primaries and spares are disjoint subtrees and that spares can only be shared via their top-node.




Fig. 3. Complex spare modules.

To illustrate the activation, consider Figure 3 (left). Here, the activation of module ‘spare’ simply means the activation of the BEs  $C$  and  $D$ . The AND gate has the same behavior whether ‘spare’ is active or not. In fact, whenever the SPARE gate (i.e. ‘system’) is activated, it activates BEs  $C$  and  $D$ .

The behavior of all the non-SPARE gates is unchanged whether they are used as spares or not; the SPARE gate does behave differently when used as a spare. To illustrate this, consider in Figure 3 (right). When ‘spare’ is not activated (i.e. ‘primary’ has not failed), BEs  $C$  and  $D$  are dormant; and even if  $C$  (being a warm spare) fails,  $D$  remains dormant. This is the same behavior as for the ‘spare’ AND gate in Figure 3 (left). If now ‘spare’ is activated, the activation signal is only used to activate the primary  $C$  and  $D$  remains dormant (this is clearly different from the AND gate ‘spare’ where both BEs are activated). Should  $C$  fail and ‘spare’ being in its active mode, then  $D$  is activated. Thus, ‘system’ activates ‘spare’, while ‘spare’ activates  $D$ .

## III. INPUT/OUTPUT INTERACTIVE MARKOV CHAINS

### A. The I/O-IMC model

Input/output interactive Markov chains (I/O-IMCs) are a combination of input/output automata (I/O-automata) [20] and interactive Markov chains (IMCs) [16].

I/O-IMCs distinguish two types of transitions: (1) *Interactive transitions* labeled with actions; (2) *Markovian transitions* labeled

with rates  $\lambda$ , indicating that the transition can only be taken after a delay that is governed by an exponential distribution with parameter  $\lambda$ . Inspired by I/O-automata, actions can be further partitioned into:

- 1) *Input actions* (denoted  $a?$ ) are controlled by the environment. They can be *delayed*, meaning that a transition labeled with  $a?$  can only be taken if another I/O-IMC performs an output action  $a!$ . A feature of I/O-IMCs is that they are *input-enabled*, i.e., in each state they are ready to respond to any of their inputs  $a?$ . Hence, each state has an outgoing transition labeled with  $a?$ .
- 2) *Output actions* (denoted  $a!$ ) are controlled by the I/O-IMC itself. In contrast to input actions, output actions cannot be delayed, i.e., transitions labeled with output actions must be taken immediately. An *observable* action is either an input or an output action.
- 3) *Internal actions* (denoted  $a;$ ) are not visible to the environment. Like output actions, internal actions cannot be delayed.

States are depicted by circles, initial states have an incoming arrow without origin, Markovian transitions are denoted by dotted lines, and interactive transitions by solid lines. Figure 4 shows an I/O-IMC  $B$  with two Markovian transitions: one from state 1 to state 2 and one from 3 to 4, both transitions with rate  $\lambda$ . The I/O-IMC has one input action  $a?$ . To ensure input-enabling, we specify  $a?$ -self-loops in states 3, 4, and  $5^2$ . Note that state 1 exhibits a race between the input and the Markovian transition: in 1, the I/O-IMC delays for a time that is governed by an exponential distribution with parameter  $\lambda$ , and moves to state 2. If however, before that delay ends, an input  $a?$  arrives, then the I/O-IMC transitions to 3. The only output action  $b!$  leads from 4 to 5.

Formally, an I/O-IMC is defined as follows.

**Definition 1 (I/O-IMC)** An input/output interactive Markov chain  $\mathcal{P}$  is a tuple  $\langle S, s^0, A, \rightarrow, \rightarrow^M \rangle$ , where:

- $S$  is a set of states,
- $s^0 \in S$  is the initial state,
- $A$  is a set of discrete actions, where  $A = (A^I, A^O, A^{int})$  is partitioned into a set of input actions  $A^I$ , output actions  $A^O$  and internal actions  $A^{int}$ . This partition is called the action signature of  $\mathcal{P}$ . We write  $A^V = A^I \cup A^O$  for the set of visible actions of  $\mathcal{P}$ .
- $\rightarrow \subseteq S \times A \times S$  is a set of interactive transitions. We write  $s \xrightarrow{a} s'$  for  $(s, a, s') \in \rightarrow$ . We require that I/O-IMCs are input-enabled:  $\forall s \in S, a? \in A^I \cdot (\exists s' \in S \cdot s \xrightarrow{a?} s')$ .
- $\rightarrow^M \subseteq S \times \mathbb{R}_{>0} \times S$  is a set of Markovian transitions. We write  $s \xrightarrow{\lambda} s'$  for  $(s, \lambda, s') \in \rightarrow^M$ .

We denote the components of  $\mathcal{P}$  by  $S_{\mathcal{P}}, s_{\mathcal{P}}^0, A_{\mathcal{P}}, \rightarrow_{\mathcal{P}}, \rightarrow_{\mathcal{P}}^M$  and omit the subscript  $\mathcal{P}$  whenever clear from the context.

### B. Parallel composition and hiding

The parallel composition operator allows one to build larger I/O-IMCs out of smaller ones. We say that two I/O-IMCs *synchronize* if either (1) they are both ready to accept the same input action or (2) one is ready to output an action  $a!$  and the other is ready to receive that same action (i.e., has input action  $a?$ ).

<sup>2</sup>In the sequel we often omit these self-loops for the sake of clarity and simplicity of the I/O-IMC representation.




Fig. 4. Two examples of I/O-IMCs.




Fig. 5. I/O-IMCs  $A||B$ , and  $(\text{hide } a \text{ in } A||B)$ .

I/O-IMCs are also equipped with a parallel composition operator “ $||$ ”, to build larger I/O-IMCs out of smaller ones. The behavior of  $\mathcal{P} = \mathcal{Q}||\mathcal{R}$ , i.e., the parallel composition of I/O-IMCs  $\mathcal{Q}$  and  $\mathcal{R}$ , is the joint behavior of its constituent I/O-IMCs and can be described as follows:

- 1) If an action does not require synchronization (i.e. it belongs to only one of the I/O-IMCs) then  $\mathcal{Q}$  and  $\mathcal{R}$  can evolve independently, i.e., if  $\mathcal{Q}$  (resp.  $\mathcal{R}$ ) can make any transition (interactive or Markovian) and behaves afterwards as  $\mathcal{Q}'$  (resp.  $\mathcal{R}'$ ), the same behavior is possible in the parallel context, i.e.,  $\mathcal{Q}||\mathcal{R}$  can evolve to  $\mathcal{Q}'||\mathcal{R}$  (resp.  $\mathcal{Q}||\mathcal{R}'$ ).
- 2) If an action of an interactive transition requires synchronization, then both I/O-IMCs  $\mathcal{Q}$  and  $\mathcal{R}$  must be able to perform that action at the same time, i.e.,  $\mathcal{Q}||\mathcal{R}$  evolves simultaneously into  $\mathcal{Q}'||\mathcal{R}'$ . Note that when an output and an input action synchronize the result is an output action.

Figure 5 illustrates the parallel composition of I/O-IMCs  $A$  and  $B$  where synchronization is on the shared action  $a$ . Formally, we have the following.

**Definition 2 (parallel composition)** Let  $\mathcal{P}$  and  $\mathcal{Q}$  be two I/O-IMCs.

- 1)  $\mathcal{P}$  and  $\mathcal{Q}$  are composable if  $A_{\mathcal{P}}^O \cap A_{\mathcal{Q}}^O = A_{\mathcal{P}}^{\text{int}} \cap A_{\mathcal{Q}} = A_{\mathcal{P}} \cap A_{\mathcal{Q}}^{\text{int}} = \emptyset$ .
- 2) If  $\mathcal{P}$  and  $\mathcal{Q}$  are composable, their composition  $\mathcal{P}||\mathcal{Q}$  is the I/O-IMC  $(S_{\mathcal{P}} \times S_{\mathcal{Q}}, (s_{\mathcal{P}}^0, s_{\mathcal{Q}}^0), ((A_{\mathcal{P}}^I \cup A_{\mathcal{Q}}^I) \setminus (A_{\mathcal{P}}^O \cup A_{\mathcal{Q}}^O), (A_{\mathcal{P}}^O \cup A_{\mathcal{Q}}^O), (A_{\mathcal{P}}^{\text{int}} \cup A_{\mathcal{Q}}^{\text{int}})), \rightarrow_{\mathcal{P}||\mathcal{Q}}, \rightarrow_{\mathcal{P}||\mathcal{Q}}^M)$ , where:

$$\begin{aligned} \rightarrow_{\mathcal{P}||\mathcal{Q}} &= \{(s, t) \xrightarrow{a} \mathcal{P}||\mathcal{Q}(s', t) \mid s \xrightarrow{a} \mathcal{P}s' \wedge a \in A_{\mathcal{P}} \setminus A_{\mathcal{Q}}\} \\ &\quad \cup \{(s, t) \xrightarrow{a} \mathcal{P}||\mathcal{Q}(s, t') \mid t \xrightarrow{a} \mathcal{Q}t' \wedge a \in A_{\mathcal{Q}} \setminus A_{\mathcal{P}}\} \\ &\quad \cup \{(s, t) \xrightarrow{a} \mathcal{P}||\mathcal{Q}(s', t') \mid s \xrightarrow{a} \mathcal{P}s' \wedge t \xrightarrow{a} \mathcal{Q}t' \wedge \\ &\quad \quad \quad a \in A_{\mathcal{P}} \cap A_{\mathcal{Q}}\} \\ \rightarrow_{\mathcal{P}||\mathcal{Q}}^M &= \{(s, t) \xrightarrow{\lambda} \mathcal{P}||\mathcal{Q}(s', t) \mid s \xrightarrow{\lambda} \mathcal{P}s'\} \\ &\quad \cup \{(s, t) \xrightarrow{\lambda} \mathcal{P}||\mathcal{Q}(s, t') \mid t \xrightarrow{\lambda} \mathcal{Q}t'\} \end{aligned}$$

Like in process algebras, the hiding operator  $\text{hide } B$  in  $\mathcal{P}$  makes internal all actions in a set  $B$  of output actions, such that no further synchronization is possible over actions in  $B$  (e.g. in Figure 5, we hide action  $a$ ).




Fig. 6. Aggregation of  $\text{hide } a$  in  $A||B$ .

**Definition 3 (Hiding)** Let  $B \subseteq A_{\mathcal{P}}^O$  be a set of output actions. We define  $\text{hide } B$  in  $\mathcal{P}$  as the I/O-IMC given by  $(S_{\mathcal{P}}, s_{\mathcal{P}}^0, (A_{\mathcal{P}}^I, A_{\mathcal{P}}^O \setminus B, A_{\mathcal{P}}^{\text{int}} \cup B), \rightarrow_{\mathcal{P}}, \rightarrow_{\mathcal{P}}^M)$ .

### C. Weak bisimilarity

State equivalences, such as bisimulation relations, are crucial in reducing the size of the model to be analyzed. By grouping together equivalent states, one obtains a model that is equivalent but smaller. This operation is called *aggregation*, *lumping* or *minimization*. For two states  $s, t$  to be bisimilar, one requires that all  $a$ -transitions in state  $s$  can be mimicked in state  $t$ . Weak bisimulations abstract from internal computation, thus the matching transition in  $t$  may be a weak transition, consisting of some internal steps, an  $a$  step (omitted if  $a$  is internal), and some more internal steps. For Markovian transitions, we compare the accumulated rates in  $s$  and  $t$ .

In this way, bisimilar states have the same observable behavior, and in particular, bisimilar states exhibit the same performance properties.

Our notion of weak bisimilarity for I/O-IMCs generalizes the one for IMCs [16]. Apart from the distinction between input and output transitions, an important difference between our approach and [16] is that we ignore Markovian self-loops (as in [9]), which drastically reduces the sizes of the I/O-IMC models.

Let  $s$  be a state and  $C \subseteq S$  be a subset of states in an I/O-IMC  $\mathcal{P}$ . We use the following notation.

- The accumulated rate from  $s$  into the set of states  $C$  is denoted  $\gamma_{\mathcal{P}}(s, C) = \sum\{\lambda \mid s \xrightarrow{\lambda} \mathcal{P}s' \wedge s' \in C\}$ , where  $\{\dots\}$  denotes a multiset of transition rates.
- State  $s$  *stable* is if it has no outgoing internal or output transitions.
- $\xrightarrow{\text{int}}$  is the *internal transition relation*, i.e. we have  $s \xrightarrow{\text{int}} t$  if  $s \xrightarrow{a} t$  for some  $a \in A^{\text{int}}$ . The *weak transition relation*  $\Longrightarrow$  arises from  $\rightarrow$  by abstracting from external steps. Thus, we have  $s \Longrightarrow t$  if there is a sequence  $s \xrightarrow{\text{int}} \dots \xrightarrow{\text{int}} t$ . We have  $s \xRightarrow{a} s'$  if there exist  $t, t'$  such that (1)  $s \Longrightarrow t$ ,  $t \xrightarrow{a} t'$  and  $t' \Longrightarrow s'$  or (2)  $a \in A^{\text{int}} \wedge s \xRightarrow{a} s'$ .
- The set  $C^{\text{int}} = \{s' \mid \exists s \in C \cdot s \xRightarrow{a} s'\}$  contains all states with a weak step into set  $C$ .

**Definition 4 (Weak bisimulation)** Let  $\mathcal{P} = \langle S, s^0, A, \rightarrow, \rightarrow^M \rangle$  be an I/O-IMC. Let  $R$  be an equivalence relation on  $S$ . Then  $R$  is a weak bisimulation iff for all  $(s, t) \in R$ ,  $a \in A$

- 1)  $s \xRightarrow{a} s'$  implies that there is a weak transition  $t \xRightarrow{a} t'$  with  $(s', t') \in R$ .
- 2)  $s \xRightarrow{a} s'$  and  $s'$  stable imply that there is a  $t'$  such that  $t \xRightarrow{a} t'$  and  $t'$  stable and  $\gamma_{\mathcal{P}}(s', C^{\text{int}}) = \gamma_{\mathcal{P}}(t', C^{\text{int}})$ , for all equivalence classes  $C \in (S/R) \setminus \{[s']_R\}$

The states  $s$  and  $t$  in  $\mathcal{P}$  are weakly bisimilar, notation  $s \approx_{\mathcal{P}} t$ , if and only if there exists a weak bisimulation  $R$  with  $(s, t) \in R$ . Weak bisimilarity for an I/O-IMC  $\mathcal{P}$  is defined as the union of all weak bisimulations on  $\mathcal{P}$ :  $\approx_{\mathcal{P}} = \bigcup\{R \mid R \text{ is a weak bisimulation on } \mathcal{P}\}$ . We often omit the name of the I/O-IMC if it is clear from context.

The following theorem states that our notion of weak bisimilarity enjoys the expected properties:  $\approx_{\mathcal{P}}$  is largest weak bisimulation relation on  $\mathcal{P}$  and weak bisimilarity is a congruence with respect to parallel composition and hiding. Its proof can be found in Appendix I.

**Theorem 1** *Let  $\mathcal{P}$  and  $\mathcal{Q}$  be two I/O-IMCs with identical action signatures, let  $\mathcal{R}$  be an I/O-IMC composable with  $\mathcal{P}$  and  $\mathcal{Q}$  and let  $B \subseteq A_{\mathcal{P}}^O$ :*

- 1)  $\approx_{\mathcal{P}}$  is a weak bisimulation on  $\mathcal{P}$  and it is the largest weak bisimulation on  $\mathcal{P}$ ,
- 2)  $\mathcal{P} \approx \mathcal{Q}$  implies  $\mathcal{P} \parallel \mathcal{R} \approx \mathcal{Q} \parallel \mathcal{R}$ ,
- 3)  $\mathcal{P} \approx \mathcal{Q}$  implies  $\mathcal{R} \parallel \mathcal{P} \approx \mathcal{R} \parallel \mathcal{Q}$ ,
- 4)  $\mathcal{P} \approx \mathcal{Q}$  implies  $\text{hide } B \text{ in } \mathcal{P} \approx \text{hide } B \text{ in } \mathcal{Q}$ .

Figure 6 shows the result after applying weak bisimulation on the I/O-IMC resulting from the composition of  $A$  and  $B$  and the hiding (i.e. made internal) of action  $a$ .

#### D. IOIML

IML[16] (IMC modeling language) provides a process algebra based syntax for specifying IMCs in an easy and concise way. We extend IML to IOIML (I/O-IMC modeling language), that provides a similar syntax for specifying I/O-IMCs. We use IML to describe the semantics of DFT elements in a parametric way.

We assume there is a countable set of *process variables*  $V$  and a countable action signature  $A = (A^I, A^O, A^V)$ .

**Definition 5 (IOIML)** *Let  $\lambda \in \mathbb{R}_{>0}$ ,  $a \in A$  and  $X \in V$ . We define the language IOIML as the set of expressions given by the following grammar:*

$$E ::= \mathbf{0} \mid a.E \mid (\lambda).E \mid E + E \mid X \mid \underline{x:=E} \mid \perp$$

The intuitive meaning of the language constructs is described below:

- The terminal symbol  $\mathbf{0}$  describes a *terminated* behaviour i.e. the process  $\mathbf{0}$  cannot perform any output or internal actions and absorbs all inputs of the I/O-IMC.
- The expression  $a.E$  may interact on action  $a$  and afterwards behave as expression  $E$ . We say that  $E$  is *action prefixed* by  $a$ . As before, we post-fix actions with "??", "!" or ":", according to their role as inputs, outputs or internal actions.

The remaining constructs are identical to their IML counterparts.

- The expression  $(\lambda).E$ , a *delay prefix* expression, describes a behaviour that will behave as expression  $E$  after a delay that is governed by an exponential distribution with a mean duration of  $1/\lambda$  time units.
- The expression  $E + F$  describes two alternatives. It may either exhibit the behaviour of expression  $E$  or the behaviour of expression  $F$ .
- The expression  $\underline{x:=E}$  describes a recursively defined behaviour. Assuming that the variable  $X$  appears somewhere inside expression  $E$ , the meaning is as follows. Whenever the variable  $X$  is encountered during the evolution of the expression, the expression will reinitialise its behaviour to  $\underline{x:=E}$ .
- The symbol  $\perp$  is intended to represent an *ill-defined* behaviour. We will not use this symbol, but it is included for completeness.

The formal semantics of an IOIML expression, i.e. its underlying I/O-IMC, can be obtained in a way similar to the semantics for IML (see [16]). Since the I/O-IMC  $\mathcal{P}$  obtained in this way need not to be input-enabled, we complete the expression by adding self-loops  $s \xrightarrow{a?} s$  whenever  $a?$  is not enabled from state  $s$ . An IOIML expression therefore must be accompanied by the action signature of the I/O-IMC it describes to be meaningful.

The IOIML description of the I/O-IMC  $B$  in Figure 4 is

$$\begin{aligned} P1 &= (\lambda).P2 + a?.P3 & P3 &= (\lambda).P4 \\ P2 &= a?.P4 & P4 &= b!.0 \end{aligned}$$

#### E. IMCs vs I/O-IMCs.

IMCs only distinguish between observable and internal actions. All observable actions are delayable and communication is a hand-shake, i.e. synchronization on action  $a$  only occurs when both IMCs involved are ready to perform the  $a$  action. While IMCs could in principle be used to model DFTs, we obtain more natural and more concise models by introducing an I/O distinction: it is always the failing DFT element that takes the initiative to notify its failure to its parents in the DFT.

## IV. DFT SYNTAX

To formalize the syntax of a DFT, we first define the set  $\mathcal{E}$ , characterizing each DFT element by its type, number of inputs and possibly some other parameters. We use the following notation. Given a set  $X$ , we denote by  $\mathcal{P}(X)$  the power set over  $X$  and by  $X^*$  the set of all sequences over  $X$ . For a sequence  $x \in X^*$ , we denote by  $|x|$  the length of the sequence (also called list), and by  $(x)_i$  the  $i^{\text{th}}$  element in  $x$ .

**Definition 6** *The set  $\mathcal{E}$  of DFT elements consists of the following tuples. Here,  $k, n \in \mathbb{N}$  are natural numbers with  $1 \leq k \leq n$  and  $\lambda, \mu \in \mathbb{R}_{>0}$  are rates.*

- $(OR, n)$ ,  $(AND, n)$ ,  $(PAND, n)$  represent respectively *OR*, *AND* and *PAND* gates with  $n$  inputs.
- $(VOT, n, k)$  represent a *voting gate* with  $n$  inputs and threshold  $k$ .
- $(SPARE, n)$  represent a *SPARE* gate with one primary and  $n - 1$  spares. By convention, the first non-dummy input to the SPARE gate is the primary component.
- $(FDEP, n)$  represents an *FDEP* gate with 1 trigger input event and  $n - 1$  dependent input events. By convention, the first non-dummy input to the FDEP gate is the trigger event.
- $(BE, 0, \lambda, \mu)$ , represents a *BE*, which has no inputs (i.e.  $n = 0$ ), an active failure rate  $\lambda$  and a dormant failure rate  $\mu$ .
- $(PHBE, 0, \phi_A, Q_A, \phi_P, Q_P, \psi)$ , represents a *phase-type BE*, which has no inputs (i.e.  $n = 0$ ), an active failure distribution with  $\phi_A \in \mathbb{N}$  phases and generator matrix  $Q_A \in \mathbb{R}^{\phi_A \cdot \phi_A}$  and a dormant failure distribution with  $\phi_P \in \mathbb{N}$  phases and generator matrix  $Q_P \in \mathbb{R}^{\phi_P \cdot \phi_P}$ . The activation of the PHBE is described by the function  $\psi : [1, \dots, \phi_P] \rightarrow [1, \dots, \phi_A]$ .

Given a tuple  $e \in \mathcal{E}$ , we write  $\text{type}(e)$  for the first item in  $e$ , and  $\text{arity}(e)$  for the second.

We introduce several notions for graphs (potentially with cycles) whose nodes are labeled with DFT elements. An edge in such graphs from  $v$  to  $w$  means that the output of the DFT element associated with  $v$  is an input to the DFT element of  $w$ .



Since the order of inputs to a gate matters (e.g. for a PAND gate), the inputs to  $v$  are given as a list  $in(v)$ , rather than as a set.

**Definition 7** An element-labeled graph is a triple  $\mathcal{D} = (V, in, l)$ , where

- $V$  is a set of vertices,
- $in : V \rightarrow V^*$  is an input function that assigns to each vertex a list of inputs.
- $l : V \rightarrow \mathcal{E}$  is a labeling function, that assigns to each vertex a DFT element.

We write  $type(v)$  for  $type(l(v))$  and  $arity(v)$  for  $arity(l(v))$ .

Given  $in$ , we define the set of edges  $E_{in}$  by  $\{(v, w) \in V^2 \mid \exists i . v = (in(w))_i\}$ . Thus,  $E_{in}$  contains all pairs  $(v, w)$  such that  $v$  appears as an input of  $w$ . We also define the *pruned* input function  $in'$  that contains only the non-dummy connections between vertices (recall that the outputs of FDEP gates are dummy outputs): Thus,  $in'(v) : V \rightarrow V^*$  is the function  $in'(v)$  that arises from  $in(v) = v_1 v_2 \dots v_n$  by removing all elements  $v_i$  s.t.  $type(v_i) = FDEP$ . Consequently, the set of edges  $E_{in'}$  is the set  $\{(v, w) \in V^2 \mid type(v) \neq FDEP \wedge \exists i . v = (in(w))_i\}$  containing all pairs  $(v, w)$  such that  $v$  appears as a non-dummy input of  $w$ . Finally, for the sake of spare activation, we define another pruned input function  $in''$  that ignores all inputs, except the trigger input, of all  $FDEP$  gates: Thus,  $in''(v) : V \rightarrow V^*$  is the function  $in''(v)$  that arises from  $in(v) = v_1 v_2 \dots v_n$  by keeping only the first element (i.e. the trigger)  $v_1$  for all  $v \in V$  s.t.  $type(v) = FDEP$ . Consequently, the set of edges  $E_{in''}$  is the set  $\{(v, w) \in V^2 \mid \exists i . (type(w) \neq FDEP \wedge v = (in(w))_i) \vee (type(w) = FDEP \wedge v = (in(w))_1)\}$ . We write  $E$  for  $E_{in}$ ,  $E'$  for  $E_{in'}$ , and  $E''$  for  $E_{in''}$  if  $in$  is clear from the context.

Given a DFT node  $v$ , the subtree below  $v$  consists of all vertices with a path in  $E''$  leading to  $v$ . Node  $v$  is *activation-independent* if there are no edges leading from  $stb(v)$  to a node outside  $stb(v)$ , except for the outgoing edges of  $v$ .

**Definition 8** Let  $\mathcal{D}$  be a DFT and  $v \in V$  a node in  $\mathcal{D}$ .

- Then subtree below  $v$ , denoted  $stb(v)$ , is the set  $\{w \mid \exists v_0, v_1 \dots v_n \in V, n \geq 0 . v_0 = w, v_n = v \wedge \forall 0 \leq i < n . (v_i, v_{i+1}) \in E''\}$ .
- Vertex  $v$  is activation-independent if  $\forall w \in stb(v), w' \in V \setminus stb(v) . (w, w') \in E'' \implies w = v$ .

Note that in the definition of an activation-independent vertex, we ignore the inputs, except the trigger input, to  $FDEP$  gates as by convention activation signals do not propagate through these edges. In the sequel, we will generally refer to an activation-independent vertex as simply an independent vertex.

Finally, we define a DFT as an element-labeled graph  $\mathcal{D}$  with several restrictions. These restrictions, which are checked syntactically by our tool, ensure that the DFT contains no anomalies and that it has a well-defined semantics.

**Definition 9** A DFT is an element-labeled graph  $\mathcal{D}$  with the following restrictions.

- $(V, E')$  forms a directed acyclic graph,
- All inputs to a DFT element must be connected to some node in  $\mathcal{D}$ , i.e. for all  $v \in V$ , we have  $arity(v) = |in(v)|$ ,
- All DFT gates must have at least one non-dummy input<sup>3</sup>:

<sup>3</sup>Otherwise, the gate should be removed.

for all  $v \in V$  with  $type(l(v)) \neq BE$  and  $type(l(v)) \neq PHBE$  we have  $|in'(v)| \geq 1$ ,

- There is a unique top element in  $\mathcal{D}$ , i.e. a non-FDEP element whose output is not connected. That is, there exists a unique  $v \in V$ ,  $type(v) \neq FDEP$  such that there is no  $w \in V$  with  $(v, w) \in E$ . This unique  $v$  is denoted by  $T_{\mathcal{D}}$ ; or by  $T$  if  $\mathcal{D}$  is clear from the context.
- The first non-dummy input of a SPARE gate (i.e. its primary) cannot be an input to another SPARE gate, i.e. primary components cannot be shared: If  $v = (in'(w))_1 = (in'(w'))_1$  and  $type(w) = type(w') = SPARE$ , then  $w = w'$ .
- Non-dummy inputs (primary and spare components) to a SPARE gate must be outputs coming from activation-independent vertices (see Section V for details): for all  $(v, w) \in E'$  with  $type(w) = SPARE$  we have that  $v$  is activation-independent,
- An output cannot be twice or more the input of the same gate: For all  $w \in V$  and  $1 \leq i, j \leq |in(w)|$  with  $(in(w))_i = (in(w))_j$ , we have  $i = j$ .

## V. DFT SEMANTICS

In this section, we first define the semantics of the DFT elements by giving the I/O-IMC for each of the tuples in  $\mathcal{E}$ . We also need two auxiliary I/O-IMCs: the activation auxiliary, that activates BEs and SPARE gates when they change from dormant to active mode and the firing auxiliary that handles the dependencies between events as modeled by the FDEP gate. Then, we obtain the semantics of the whole DFT from the parallel composition of the semantics of its elements and the auxiliaries.

The semantics of each non-FDEP element in  $\mathcal{E}$  (denoted  $[[\dots]]_{ELT}$ ) is a function which takes as input a number of actions and returns an I/O-IMC. The FDEP gate is handled through the use of firing auxiliaries. We present the graphical descriptions for BEs and gates with two or three inputs and we use the language IOIML to specify the semantics for the general case.

**Basic event I/O-IMC model.** As pointed out in Section II, a BE




Fig. 7. The I/O-IMCs  $[[BE, 0, \lambda, 0]]_{ELT}(a, f)$ ,  $[[BE, 0, \lambda, \mu]]_{ELT}(a, f)$ , and  $[[BE, 0, \lambda, \lambda]]_{ELT}(a, f)$ , modeling the semantics of a cold, warm and hot BE.

has a different failing behavior depending on its dormancy factor. Figure 7 shows the (parametrized) I/O-IMCs associated to a cold, warm, and hot BE<sup>4</sup>, i.e. it shows the functions  $[[BE, 0, \lambda, \mu]]_{ELT} : A^2 \rightarrow IOIMC$  taking as arguments an activation signal  $a?$  and a firing signal  $f!$ .

In IOIML the I/O-IMC  $[[BE, 0, \lambda, \mu]]_{ELT}(a, f)$  has action signature  $(\{a\}, \{f\}, \emptyset)$  and is described by the following expression  $E_0$ :

$$E_0 = \begin{cases} a?.E_1 + (\mu).E_2, & \text{if } \mu > 0 \\ a?.E_1, & \text{otherwise} \end{cases}$$

$$E_1 = (\lambda).E_2$$

$$E_2 = f!.0$$

<sup>4</sup>The hot BE I/O-IMC can be reduced to:  $\odot \xrightarrow{\lambda} \odot \xrightarrow{f!} \odot$

**Phase-type basic event I/O-IMC model.** A phase-type basic event (PHBE) does not fail after an exponential delay, but rather after a delay governed by a phase-type (PH) distribution [22]. Here, the phase-type distributions for failure in either passive or active mode are described by absorbing CTMCs. A PHBE is described by the tuple  $(PHBE, 0, \phi_A, Q_A, \phi_P, Q_P, \psi)$  where  $\phi_A, \phi_P \in \mathbb{N}$  denote the number of phases of the active and passive PH distribution, respectively. Matrices  $Q_A : [1, \phi_A] \times [1, \phi_A] \rightarrow \mathbb{R}$  and  $Q_P : [1, \phi_P] \times [1, \phi_P] \rightarrow \mathbb{R}$  are the generator matrices of the PH distributions<sup>5</sup>. Finally the function  $\psi : [1, \phi_P] \rightarrow [1, \phi_A]$  matches passives phases to active phases. If the basic event is activated while its passive failure distribution is in phase  $i$  then the I/O-IMC will move to phase  $\psi(i)$  of the active failure distribution. In the case of a cold spare the number of passive phases is set to one with the only entry in  $Q_P$  being 0. This is interpreted as being the PH representation with a single state that cannot reach the absorbing state. This representation in fact does not represent a true PH distribution, but the semantics is clear: the spare can never fail when it is in passive mode. For other PHBEs the generator matrices must have strictly negative numbers on the diagonal and positive numbers elsewhere. Furthermore the sum of each row must be negative.

In IOIML we find action signature  $(\{a\}, \{f\}, \emptyset)$  for the I/O-IMC  $\llbracket (PHBE, 0, \phi_A, Q_A, \phi_P, Q_P, \psi) \rrbracket_{\text{ELT}}(a, f)$ . The I/O-IMC is described by the expression  $E_{P,i}$ , below  $i \in [1, \phi_P]$  and  $k \in [1, \phi_A]$ :

$$E_{P,i} = \begin{cases} a?.E_{A,\psi(i)}, & \text{if } Q_P(i, i) = 0 \\ a?.E_{A,\psi(i)} + \\ \sum_{1 \leq j \leq \phi_P \wedge i \neq j} (Q_P(i, j)).E_{P,j} + \\ (-\sum_{1 \leq j \leq \phi_P} Q_P(i, j)).E_F, & \text{otherwise} \end{cases}$$

$$E_{A,k} = \sum_{1 \leq j \leq \phi_A \wedge k \neq j} (Q_A(k, j)).E_{A,j} + (-\sum_{1 \leq j \leq \phi_A} Q_A(k, j)).E_F$$

$$E_F = f!.0$$

**VOTING gate I/O-IMC model.** Figure 8 shows the semantics of the voting gate  $(VOT, 3, 2)$  element, i.e. the function  $\llbracket (VOT, 3, 2) \rrbracket_{\text{ELT}} : A^4 \rightarrow \text{IOIMC}$ , taking as arguments the output and three input signals of the VOTING gate. The voting gate fires (action  $f_1$ ) when at least 2 of its inputs fire (actions  $f_2, f_3$ , and  $f_4$ ).




Fig. 8. The I/O-IMC  $\llbracket (VOT, 3, 2) \rrbracket_{\text{ELT}}(f_1, f_2, f_3, f_4)$ .

To define the semantics of a  $(VOT, n, k)$  gate with  $n$  inputs and threshold  $k$ , we use the process variables  $P_V(I, U, f, k)$ , that depend on three parameters; a set  $I$  containing the firing signals of all inputs to the  $VOT$  gate, a set  $U$  containing the firing signals

<sup>5</sup>As the initial distribution of our phase-type representation we always use the vector  $[1, 0, \dots, 0]$ , i.e. a single starting state. This is not a problem since any PH representation can easily be transformed to a PH representation with the same number of phases and a single starting state.

of inputs that are still operational, and an action  $f, f \notin I \cup U$ , being the  $VOT$  gate's own output firing signal. We set

$$P_V(I, U, f, k) = f!.0 \quad \text{if } |I \setminus U| \geq k$$

$$P_V(I, U, f, k) = \sum_{a \in I} a?.P_V(I, U \setminus \{a\}, f, k) \quad \text{if } |I \setminus U| < k.$$

Thus,  $P_V(I, U, f, k)$  emits the failure signal  $f!$  after having received  $k$  failure signals. The I/O-IMC of an  $n$ -input voting gate is:  $\llbracket (VOT, n, k) \rrbracket_{\text{ELT}}(f_o, f_1, \dots, f_n) = P_V(\{f_1, \dots, f_n\}, \{f_1, \dots, f_n\}, f_o, k)$  with action signature  $(\{f_1, \dots, f_n\}, \{f_o\}, \emptyset)$ . Note that the  $VOT$  gate<sup>6</sup> does not have an activation signal as this element does not exhibit a dormant or active behavior as such.

**AND gate I/O-IMC model.** Figure 9(a) shows the semantics of the  $(AND, 2)$  gate, i.e. the function  $\llbracket (AND, 2) \rrbracket_{\text{ELT}} : A^3 \rightarrow \text{IOIMC}$ , taking as arguments the output and two input signals of the AND gate. This I/O-IMC models the fact that the AND gate fires (action  $f_1$ ) after it receives firing signals from both its inputs (actions  $f_2$  and  $f_3$ ).




Fig. 9. (a)  $\llbracket (AND, 2) \rrbracket_{\text{ELT}}(f_1, f_2, f_3)$ , (b)  $\llbracket (OR, 2) \rrbracket_{\text{ELT}}(f_1, f_2, f_3)$ , (c)  $\llbracket (PAND, 2) \rrbracket_{\text{ELT}}(f_1, f_2, f_3)$ .

The semantics of an  $(AND, n)$  gate with  $n$  inputs is defined as a special case of the  $VOT$  gate, where the threshold is equal to the number of inputs. The I/O-IMC associated to an  $n$ -ary AND gate is then given by:  $\llbracket (AND, n) \rrbracket_{\text{ELT}}(f_o, f_1, \dots, f_n) = P_V(\{f_1, \dots, f_n\}, \{f_1, \dots, f_n\}, f_o, |\{f_1, \dots, f_n\}|)$  with action signature  $(\{f_1, \dots, f_n\}, \{f_o\}, \emptyset)$ .

**OR gate I/O-IMC model.** Figure 9(b) shows the semantics of the OR gate  $(OR, 2)$  element, i.e. the function  $\llbracket (OR, 2) \rrbracket_{\text{ELT}} : A^3 \rightarrow \text{IOIMC}$ , taking as arguments the output and two input signals of the OR gate. The OR gate fires (action  $f_1$ ) after it receives one of its input firing signals (actions  $f_2$  or  $f_3$ ).

The semantics of an  $(OR, n)$  gate with  $n$  inputs is defined as a special case of the  $VOT$  gate with threshold equal to one. The I/O-IMC associated to an  $n$ -ary OR gate is then given by:  $\llbracket (OR, n) \rrbracket_{\text{ELT}}(f_o, f_1, \dots, f_n) = P_V(\{f_1, \dots, f_n\}, \{f_1, \dots, f_n\}, f_o, 1)$  with action signature  $(\{f_1, \dots, f_n\}, \{f_o\}, \emptyset)$ .

**PAND gate I/O-IMC model.** Figure 9(c) shows the semantics of the PAND gate  $(PAND, 2)$  element, i.e. the function  $\llbracket (PAND, 2) \rrbracket_{\text{ELT}} : A^3 \rightarrow \text{IOIMC}$ , taking as arguments the output and two input signals of the PAND gate. The PAND gate fires (action  $f_1$ ) after all its inputs (actions  $f_2$  or  $f_3$ ) fire from left to right order. If the inputs fire in the wrong order, the PAND gate moves to an operational absorbing state (denoted with an **X**). The semantics of a  $(PAND, n)$  gate with  $n$  inputs is defined by means of the process variables  $P_P(U, f)$ . However now,  $U$  is given as

<sup>6</sup>This is true for all gates except the SPARE gate.



a sequence of firing signals of operational inputs, rather than a set and  $f$  is the *PAND* gate's own firing signal. The actions in  $U$  must occur in the correct order for the *PAND* gate to fail. We write  $U = a_1 a_2 \dots a_n$ . We set

$$P_P(U, f) = f! \cdot \mathbf{0} \quad \text{if } U = \epsilon$$

$$P_P(U, f) = a_k? \cdot P_P(U \setminus \{a_k\}, f) + \sum_{a \in U \setminus \{a_k\}} a? \cdot \mathbf{0} \quad \text{if } U = a_k a_{k+1} \dots a_n$$

Now  $P_P(U, f)$  emits the failure signal  $f!$  after having received failure signals from all its inputs, which can only happen if they occurred in the specified order, since deviations from this order end in  $\mathbf{0}$ . We set  $\llbracket (PAND, n) \rrbracket_{\text{ELT}}(f_o, f_1, \dots, f_n) = P_P(f_1 \dots f_n, f_o)$  with action signature  $(\{f_1, \dots, f_n\}, \{f_o\}, \emptyset)$ .

**FDEP gate I/O-IMC model.** An FDEP gate does not have a semantics in itself, but instead it is used in combination with the semantics of its dependent events. To model a functional dependency, we define the *firing auxiliary* function  $FA : A^2 \times \mathcal{P}(A) \rightarrow \text{IOIMC}$ . This (parametric) I/O-IMC ensures that a dependent event fires either when the event fails by itself, or when its failure is triggered by the FDEP gate trigger: Figure 10 (left) shows the  $FA$  to be applied in combination with an event that is functionally dependent on  $n$  triggers. Signal  $f_2$  corresponds to the failure of the dependent event by itself; signals  $f_3, f_4, \dots, f_{n+2}$  correspond to the failures of any of the triggers; and  $f_1$  corresponds to the failure of the dependent event when also considering its functional dependency upon the triggers. Hence,  $f_1$  is emitted as soon as any signal from  $\{f_2, f_3, \dots, f_{n+2}\}$  occurs. Thus, the  $FA$  takes as arguments two firing signals and a set of firing signals (corresponding to all triggers of the dependent event).




Fig. 10.  $FA(f_1, f_2, \{f_3, f_4, \dots, f_{n+2}\})$  (left) and an example of the FDEP gate extension (right).

The I/O-IMC  $FA(f_1, f_2, T)$  can in fact be interpreted as an *OR* gate:  $FA(f_1, f_2, T) = \llbracket (OR, |T| + 1) \rrbracket_{\text{ELT}}(f_1, f_2, t_1, \dots, t_n)$ , where  $T = \{t_1, \dots, t_n\}$ . The I/O-IMC  $FA(f_1, f_2, T)$  has the following action signature:  $(\{f_2, t_1, \dots, t_n\}, \{f_1\}, \emptyset)$ .

Note that the FDEP gate can trigger the failure of any gate (representing a subsystem) and not only BE as originally defined in Galileo [25]. Indeed, this extension comes at no extra cost, and the I/O-IMC used in this case is still the same as the one shown in Figure 10 (left). Figure 10 (right) shows such a configuration where  $T$  triggers the failure of the subsystem  $A$ . Note that subsystem  $A$  does not need to be an independent module. Note also that the trigger  $T$  only affects the failure of the gate  $A$  and none of its elements below it such as the basic event  $C$ .

When  $l(v)^7$ , an element of the DFT, is triggered by multiple FDEP gates, then we define  $T_v = \{f_t \mid \exists w \in V. (v, w) \in E' \wedge \text{type}(w) = \text{FDEP} \wedge t = (\text{in}'(w))_1\}$  as the set of trigger signals of FDEP gates on which  $l(v)$  is dependent.

**SPARE gate I/O-IMC model.** Given the discussion in Section II-B, Figure 11 shows the I/O-IMC of a SPARE gate (the spare gate

on the left side) sharing a spare with another SPARE gate. When the SPARE gate is active, the state reached after the primary fails is of particular interest. In this state, a non-deterministic situation arises where the spare can be activated by either of the SPARE gates (signals  $a_{S,A}!$  and  $a_{S,B}?$ ). This matches exactly the non-deterministic choice described in Subsection II-A.

The semantics of a SPARE gate having  $n-1$  spares is a function  $A^3 \times (A^2 \times \mathcal{P}(A))^{n-1} \rightarrow \text{IOIMC}$  that takes as inputs the firing signal and the activation signal of the SPARE gate, the firing signal of its primary and a sequence of spare-tuples containing, for each spare, its firing signal, its activation signal (output by the SPARE gate in question) and a list of spare activation signals of the other SPARE gates sharing that spare.




Fig. 11. The semantics  $\llbracket (SPARE, 2) \rrbracket_{\text{ELT}}(f_A, a_A, f_P, (f_S, a_{S,A}, \{a_{S,B}\}))$  of (left) SPARE gate.

We now look at the IOIML definition of a spare gate with  $n-1$  (possibly shared) spares  $\llbracket (SPARE, n) \rrbracket_{\text{ELT}}(f_1, a_1, f_2, S)$ , where  $f_1$  is the failure signal of the spare gate,  $a_1$  is the activation signal of the spare gate,  $f_2$  is the failure signal of the primary and  $S = (f_3, a_{3,1}, P_3), \dots, (f_{n+1}, a_{n+1,1}, P_{n+1})$  and  $P_i = \{a_{i,2}, \dots, a_{i,m}\}$ . The set  $P_i$  is in fact the set of all activation signals of the  $i$ -th spare by other spare gates. The signals  $a_{x,y}$  then correspond to the activation of spare  $x$  by spare gate  $y$ . We separate the state space of the spare gate into four distinct sets:

- DO The spare gate is dormant and its primary is operational.
- DN The spare gate is dormant and its primary is not operational.
- AO The spare gate is active and its primary is operational.
- AN The spare gate is active and its primary is not operational.

We now define the functions  $DO, DN, AO, AN$ :

$$DO(f, a, fp, S) = a? \cdot AO(f, a, fp, S) + fp? \cdot DN(f, a, fp, S) + \sum_{(k,l,M) \in S} \sum_{x \in \{k\} \cup M} x? \cdot DO(f, a, fp, S - (k, l, M))$$

$$DN(f, a, fp, S) = f! \cdot \mathbf{0} \quad , \text{ if } S = \emptyset$$

$$DN(f, a, fp, S) = a? \cdot AN(f, a, fp, S) + \sum_{(k,l,M) \in S} \sum_{x \in \{k\} \cup M} x? \cdot DN(f, a, fp, S - (k, l, M)) \quad , \text{ if } S \neq \emptyset$$

<sup>7</sup>Does not apply to FDEP gates.

$$\begin{aligned}
 AO(f, a, fp, S) &= fp?.AN(f, a, fp, S) + \\
 &\sum_{(k,l,M) \in S} \sum_{x \in \{k\} \cup M} x?.AO(f, a, fp, S - (k, l, M)) \\
 AN(f, a, fp, S) &= f!.0 \quad , \text{ if } S = \emptyset \\
 AN(f, a, fp, S) &= q!.AO(f, a, p, S - (p, q, R)) + \\
 &\sum_{(k,l,M) \in S} \sum_{x \in \{k\} \cup M} x?.AN(f, a, fp, S - (k, l, M)) \\
 &\quad , \text{ if } S = (p, q, R), \dots
 \end{aligned}$$

Now the IOIML definition of  $\llbracket (SPARE, n) \rrbracket_{\text{ELT}}(f_1, a_1, f_2, S) = DO(f_1, a_1, f_2, S)$ . The action signature of  $\llbracket (SPARE, n) \rrbracket_{\text{ELT}}(f_1, a_1, f_2, S)$  is:  $(\{a_1, f_2\} \cup \{k\} \cup M \mid (k, l, M) \in S, \{f_1\} \cup \{l \mid (k, l, M) \in S\}, \emptyset)$ .

**The activation auxiliary.** BEs and SPARE gates have a distinct input activation signal. When more than one SPARE gate can activate any of these two elements it becomes convenient to carry out this activation through an intermediate I/O-IMC model called *activation auxiliary*.

Activating a BE (or a SPARE gate)  $l(v)$  is done by composing  $\llbracket l(v) \rrbracket_{\text{ELT}}$  in parallel with an activation auxiliary I/O-IMC model, where the latter outputs the activation signal  $a_v$  of  $l(v)$ . The activation auxiliary I/O-IMC model is obtained through a function  $AA : A \times \mathcal{P}(A) \rightarrow \text{IOIMC}$  that takes as arguments an output activation signal and a set of input activation signals (emitted by some SPARE gates). The activation auxiliary behaves similarly to an OR gate: It outputs the activation signal as soon as it receives an activation signal emitted by one of the SPARE gates.

The general form of  $v$ 's activation auxiliary function  $AA$  is  $AA(a_v, Atv_v)$ , where  $a_v$  is  $v$ 's activation signal and  $Atv_v = \{a_{w,sp} \mid \text{type}(sp) = SPARE \wedge w \in (\text{in}'(sp) \setminus (\text{in}'(sp))_1) \wedge v \in \text{stb}(w) \wedge (\#v_0, v_1 \dots v_n \in V, n \geq 0 . v_0 = v, v_n = w \wedge \forall 0 \leq i < n . (v_i, v_{i+1}) \in E'' \wedge \text{type}(v_{i+1}) = SPARE \wedge v_i \in (\text{in}'(v_{i+1}) \setminus (\text{in}'(v_{i+1}))_1))\}$  is the set of activation signals emitted by all SPARE gates sharing  $v$ . The last clause simply ensures that there is no directed path from  $v$  to  $w$  containing an edge that is a spare input (i.e. non-primary) to a SPARE gate. It is important to note that activation does not propagate through an FDEP dependent event input.

Thus we can write  $AA(a_v, Atv_v) = \llbracket (OR, n) \rrbracket_{\text{ELT}}(a_v, (Atv_v)_1, \dots, (Atv_v)_n)$ , given that  $|Atv_v| = n$  ( $n > 0$ ). The action signature of  $AA(a_v, Atv_v)$  is  $(Atv_v, a_v, \emptyset)$ . If  $n = 0$  (i.e. no explicit activation by a SPARE gate and therefore activated when system starts at time  $t = 0$ ), then  $AA(a_v, \emptyset) = a_v!.0$ .

**Complete semantics of a DFT.** To obtain the semantics of a DFT from the semantics of its elements, we need to appropriately instantiate the parameters of  $\llbracket l(v) \rrbracket_{\text{ELT}}$  (we sometimes use  $\llbracket v \rrbracket_{\text{ELT}}$  for short) of each node  $v$ . We use the following notation: (1) The firing signal  $f_v$  of element  $l(v) \in \mathcal{E}$  denotes the failure of  $v$ , (2) the activation signal  $a_v$  denotes its<sup>8</sup> activation, and (3)  $a_{v,u}$  denotes the activation signal output by a SPARE gate  $u$  to activate spare  $v$ . We also introduce the following notation:  $V_{BE}$  is the set of all nodes  $v \in V$  s.t.  $\text{type}(v) = BE$  or  $\text{type}(v) = PHBE$ ,  $V_{AOVP}$  is the set of all nodes  $v \in V$  s.t.  $\text{type}(v) = AND \vee OR \vee VOT \vee PAND$ , and  $V_{SPARE}$  is the set of all nodes  $v \in V$  s.t.

<sup>8</sup>Only for BEs and SPARE gates.

$\text{type}(v) = SPARE$ . Now, the semantics of a DFT is obtained by parallel composing the semantics of all (non-FDEP) nodes.

**Definition 10** *The semantics of a DFT  $\mathcal{D} = (V, in, l)$  is the I/O-IMC*

$$\begin{aligned}
 \llbracket \mathcal{D} \rrbracket &= \parallel_{v \in V_{BE}} \llbracket v \rrbracket_{\text{ELT}}(a_v, f_v^*) \parallel FA(f_v, f_v^*, T_v) \parallel AA(a_v, Atv_v) \\
 &\parallel_{v \in V_{AOVP}} \llbracket v \rrbracket_{\text{ELT}}(f_v^*, f_{w_1}, f_{w_2}, \dots, f_{w_n}) \parallel FA(f_v, f_v^*, T_v) \\
 &\parallel_{v \in V_{SPARE}} \llbracket v \rrbracket_{\text{ELT}}(f_v^*, a_v, f_{w_1}, S_2, \dots, S_n) \parallel FA(f_v, f_v^*, T_v) \parallel \\
 &\quad AA(a_v, Atv_v)
 \end{aligned}$$

where,  $\text{in}'(v) = w_1 w_2 \dots w_n$  and  $S_i = (f_{w_i}, a_{w_i, v}, P_{w_i})$  with  $i > 1$ , is a tuple which gives, for spare  $l(w_i)$  the failure signal ( $f_{w_i}$ ), the activation signal by SPARE gate  $l(v)$  ( $a_{w_i, v}$ ) and the set of activation signals emitted by all SPARE gates (except  $v$ ) sharing spare  $l(w_i)$ :  $P_{w_i} = \{a_{w_i, g} \mid (w_i, g) \in E' \wedge g \neq v \wedge \text{type}(g) = SPARE\}$ .

To compute the reliability of  $\mathcal{D}$ , we are only interested in the failure of the top node  $T$ . Hence, we hide all signals except  $f_T$ , i.e. we compute  $M_{\mathcal{D}} = \text{hide } A_{\mathcal{D}} \setminus f_T \text{ in } \llbracket \mathcal{D} \rrbracket$ ; recall that  $A_{\mathcal{D}}$  denotes the set of all actions in  $\mathcal{D}$ . The compositional aggregation technique described in Section VII is an efficient way to derive  $M_{\mathcal{D}}$ .

**Example 2** Figure 12 shows the I/O-IMC semantics of a DFT consisting of a SPARE gate  $A$  having a primary  $B$  and a spare  $C$ . Since the DFT contains no FDEP gates, we ignore all firing auxiliaries. The I/O-IMC of the DFT is obtained by parallel composing  $\llbracket A \rrbracket$ ,  $\llbracket B \rrbracket$ , and  $\llbracket C \rrbracket$ .

$$\begin{aligned}
 \llbracket A \rrbracket &= \llbracket (SPARE, 2) \rrbracket_{\text{ELT}}(f_A, a_A, f_B, (f_C, a_{C,A}, \emptyset)) \parallel AA(a_A, \emptyset) \\
 \llbracket B \rrbracket &= \llbracket (BE, 0, \lambda, 0) \rrbracket_{\text{ELT}}(a_B, f_B) \parallel AA(a_B, \emptyset) \\
 \llbracket C \rrbracket &= \llbracket (BE, 0, \lambda, \mu) \rrbracket_{\text{ELT}}(a_C, f_C) \parallel AA(a_C, \{a_{C,A}\})
 \end{aligned}$$

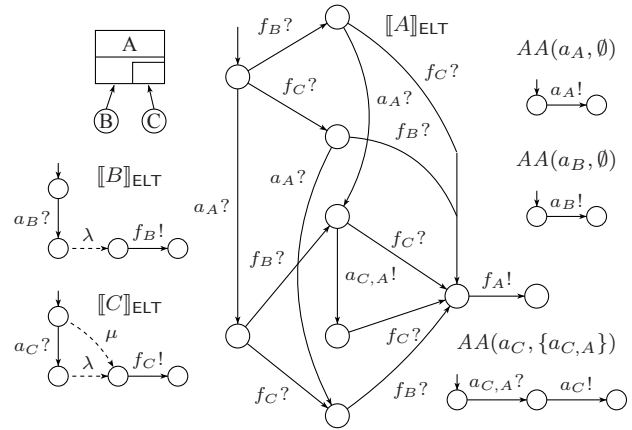


Fig. 12. A DFT example and the six I/O-IMCs that model its behavior.

## VI. DFT ELEMENTS EXTENSION

In this section, we show, through three examples, how one can readily extend the DFT elements within the I/O-IMC framework. These extensions concern the modeling of *inhibition*, *mutually exclusive events*, and *repair*.

Adding/modifying elements is done at the level of the elementary I/O-IMC models. Moreover, adding/modifying one element

does not affect the remainder of the elements (i.e. their corresponding I/O-IMC models). This is indeed a desirable property of the I/O-IMC framework, where the behavioral details and interactions of any element is kept as local as possible. These extensions only affect Step 1 of the DFT conversion/analysis algorithm laid out in Section VII, leaving the other five steps unchanged.

**Inhibition and mutual exclusion.** We say that event  $A$  *inhibits* the failure of  $B$  if the failure of  $B$  is prevented when  $A$  fails before  $B$ . Following the idea of the firing auxiliary (cf. Section V), this could be modeled by simply adding an *inhibition auxiliary* (IA). Figure 13 shows the configuration of such inhibition and the corresponding I/O-IMC model of the IA of  $B$ .  $f_B^*$  corresponds to the failure signal of  $B$  taken in isolation, i.e. without  $A$ 's inhibition. Note that, as with the FA, any element which has  $B$  as input has to now interface with  $B$ 's IA rather than directly with  $B$ .

If  $A$  inhibits the failure of  $B$  and  $B$  also inhibits the failure of  $A$ , the failure of  $A$  and the failure of  $B$  become two mutually exclusive events. Clearly, this can be modeled in our framework by adding IAs for both  $A$  and  $B$ . Mutual exclusion is very useful for modeling different failures of a single component with different effects (e.g. a valve being stuck open or closed).




Fig. 13. The I/O-IMC model of the IA.

**Repair.** Adding a notion of repair is somewhat more complicated as every DFT element can now fail or be repaired. Thus, no longer only a ‘failed event’ should be signaled but also a ‘repaired event’. However, as mentioned above, we only need to modify ‘locally’ the elementary I/O-IMC corresponding to each DFT element behavior. Here we will only discuss the new I/O-IMC for the BE and the AND gate (other elements are treated in the same fashion). The repairable cold BE’s I/O-IMC is shown in Figure 14. Here,  $\mu$  denotes the BE repair rate and  $r!$  is a signal output by the BE notifying the rest of the elements that it has been repaired. The repairable AND gate I/O-IMC model is shown




Fig. 14. The repairable BE I/O-IMC model.

in Figure 15. The AND gate has its own repair output signal (i.e.  $r!$ ) and needs to consider both failure ( $f_A?$  and  $f_B?$ ) and repair ( $r_A?$  and  $r_B?$ ) signals coming from its inputs  $A$  and  $B$ . Compared to the unreparable AND gate (Figure 9), Figure 15 has 3 extra states. If we consider a very simple repairable system composed of an AND gate with two BE  $A$  and  $B$  (Figure 16 (left)), then the resulting I/O-IMC after automatic composition, hiding of all signals and aggregation is, as expected, the CTMC shown in Figure 16 (right). At this point, one can perform analysis on the CTMC such as computing the system unavailability.




Fig. 15. The repairable AND gate I/O-IMC model.




Fig. 16. A simple repairable system. The gray state denotes the state in which the DFT has failed.

## VII. COMPOSITIONAL AGGREGATION APPROACH

The technique of *compositional aggregation* consists of composing a large model out of smaller ones and aggregating sub-models after each compositional step. This approach is to be contrasted with a more classical approach of model generation, such as the one used by Galileo DFTree [21], where the model of a system is generated at once and as a whole and then eventually aggregated at the end. Compositional aggregation is very effective in combating the state-space explosion problem and has been already successfully used on a number of case studies, most notably in [18].

Once the DFT elements have been converted into a set of I/O-IMCs, the compositional aggregation methodology can be applied to combine the set into a single I/O-IMC. The final I/O-IMC reduces in many cases to a CTMC. This CTMC can then be solved using standard methods [24] to compute performance measures such as system unreliability. The conversion/analysis algorithm<sup>9</sup> is as follows:

- 1) Map each DFT element to its corresponding (aggregated) I/O-IMC and match all inputs and outputs. The result of this step is a set of I/O-IMCs.
- 2) Pick two I/O-IMC and parallel compose them.
- 3) Hide output signals that won't be subsequently used (i.e. synchronized on).
- 4) Aggregate (using weak bisimulation) the I/O-IMC obtained from the composition of the two I/O-IMC picked in Step 2 and the hiding of the output signals in Step 3.
- 5) Go to Step 2 if more than one I/O-IMC is left, otherwise go to Step 6.
- 6) Analyse the aggregated CTMC.

The choice of I/O-IMCs made in step 2 is important as it influences the size of the generated state-space during the intermediate steps. If no non-determinism is present in the DFT model, then the algorithm yields a CTMC. However, in some cases where non-determinism arises, the result is a continuous-time Markov decision process, which can be analyzed by computing bounds on the performance measure of interest [2].

<sup>9</sup>Note that this algorithm is amenable to parallelization.



### VIII. TOOL SUPPORT

In this section we describe the CORAL [6] tool chain which supports our DFT analysis methodology. The tools presented in this section use the CADP tool-set [15] for many operations on I/O-IMCs, such as composition, aggregation and CTMC analysis. Before discussing the various tools in detail we will first give an overview of the tool chain.

**Tool chain overview.** Figure 17 shows an overview of the tool chain for our DFT analysis methodology. The user must supply

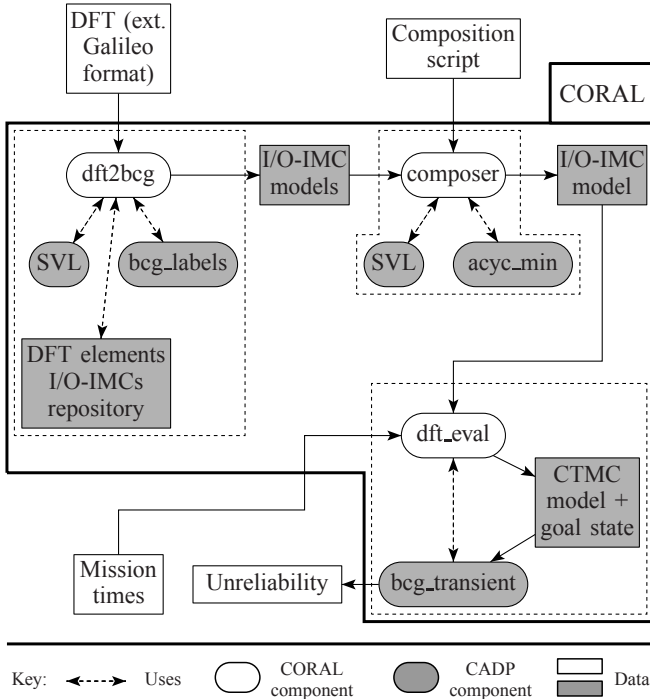


Fig. 17. An overview of the CORAL tool chain.

the following input: the DFT in a file using the Galileo format, a composition-script denoting the order of composition used in the compositional aggregation phase and the mission times for the unreliability analysis.

A DFT can be analyzed by performing the following three steps, which are elaborated below.

- 1) Call the *dft2bcg* tool with as input the DFT file in extended Galileo format<sup>10</sup>,
- 2) Call the *composer* tool with as input a composition script, and
- 3) Call the *dft\_eval* tool with as input a number of mission-times. The *dft\_eval* tool then calculates the unreliability of the system modeled by the original DFT for the given mission-times. It generates a CTMC model describing the exact failure distribution of the system if there is no non-determinism.

**Generating I/O-IMC models: dft2bcg.** The *dft2bcg* tool generates a number of I/O-IMC models that describe the behaviour of a given DFT. To be more exact, each of the I/O-IMCs describes one element in the DFT (see Section V). These I/O-IMC models are

<sup>10</sup>The standard Galileo format is extended to allow complex spares and dependent events

stored in binary coded graph (BCG) format supported by CADP. The *dft2bcg* tool uses a number of script verification language (SVL) scripts to generate BCG files. These scripts are interpreted by the SVL tool, which is part of the CADP tool set, to perform generation, parallel composition, hiding and minimization of BCG files. *dft2bcg* also uses the *bcg\_labels* tool of the CADP tool-set which allows the renaming of the actions of an I/O-IMC.

The *dft2bcg* tool performs the following steps to generate the I/O-IMC models:

- 1) Parses the DFT input file,
- 2) Checks the validity of the DFT (i.e. syntactic check),
- 3) Calls the SVL tool to generate I/O-IMC models with generic action names ( $f_1, f_2, \dots$ ) using the DFT SVL scripts, and
- 4) Calls the *bcg\_labels* tool to create the I/O-IMC models by renaming the generic actions to the specific actions derived from the names of the DFT elements.

All generic models generated in step 3 are stored in a DFT repository for reuse in later calls of the *dft2bcg* tool. For instance if we need three different 3-input AND-gates we can simply use the same generic 3-input AND gate, renaming it differently in step 4 of the *dft2bcg* tool. The repository also holds a number of basic I/O-IMC models which are used to generate the models of all DFT elements.

**Compositional aggregation: composer.** We have seen above that the *dft2bcg* tool generates a number of I/O-IMC models. The *composer* tool uses as input these I/O-IMC models and a composition script supplied by the user. The composition script describes the order in which the I/O-IMC models should be composed. The *composer* tool executes the commands in the script, composing the I/O-IMC models into a single I/O-IMC model which represents the stochastic behaviour of the entire DFT. Our choice of composition script is based on heuristics, such as maximizing the number of transitions that will be hidden and minimizing the number of actions that are not synchronized. After each composition the resulting I/O-IMC model is minimized using the *acyc\_min* tool [12]. The *acyc\_min* tool minimizes acyclic (except for self-loops with input-actions) I/O-IMCs with respect to weak bisimulation for I/O-IMCs (see Section III).

**Calculating measures: dft\_eval.** In many cases the stochastic behaviour of the system described by a DFT can be modelled as a CTMC. To be more specific: if there is no non-determinism present in the DFT model of the system the I/O-IMC generated in our approach reduces to a CTMC. See Section II-A for a detailed discussion on the occurrence of non-determinism in DFT analysis. The *dft\_eval* tool first reduces the I/O-IMC representation of the DFT into a CTMC and then invokes the CADP tool *bcg\_transient* to find the unreliability of the DFT for a set of mission-times supplied by the user.

### IX. CASE STUDIES

We have assessed the efficiency of our compositional aggregation approach by performing nine case studies from different application areas. We analyzed a cascaded PAND system (CPS), two versions of a cardiac assist system (CAS), five versions of a fault tolerant parallel processors (FTPP), and finally a pump system with inherent non-determinism. We systematically compare our results (using the *CORAL* tool) to the Galileo DIFTree tool [14] results, see Table I for an overview. Here, the number

of states/transitions corresponds to the largest I/O-IMC or CTMC encountered during analysis. All experiments were run on an AMD Athlon XP 2600+ running at 1.9 GHz with 1 GB memory.

**The cascaded PAND system.** This system, taken from [7] and shown on Figure 18 (left), illustrates the enhanced modularity of our methodology compared to Galileo DIFTree. In fact, given that the top node of the tree is a PAND dynamic gate, Galileo DIFTree can only consider the tree as a whole when generating/solving its corresponding CTMC. In our compositional aggregation approach, we realize that there are independent modules, in particular A, C, and D are all identical (all BEs have a failure rate equals to 1) and independent modules. In fact it suffices to generate and aggregate the I/O-IMC of one of these three modules and reuse the result, given some renaming of signals, for the remaining two modules. In this way, A, C, and D have each an aggregated I/O-IMC of 7 states only. The CTMC generated by Galileo has 4113 states. This result is to be compared with the largest I/O-IMC of 113 states obtained during the compositional aggregation.




Fig. 18. DFTs for CPS (left) and CAS (right) case studies

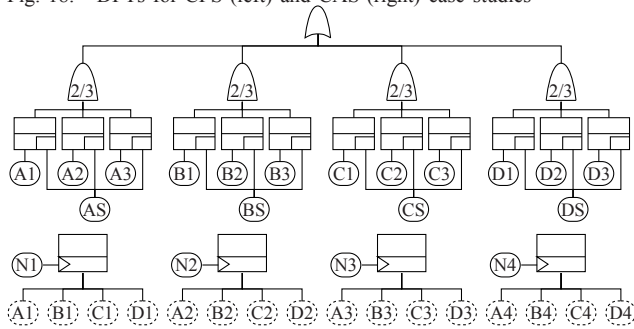


Fig. 19. DFT for the FTTP-4 case study

Case study	Approach	Max # of states	Max # of transitions	Unreliability (time = 1)	Time (sec)
CPS	Galileo	4113	24608	0.00135	490
	CORAL	133	465	0.00135	67
CAS	Galileo	8	10	0.65790	1
	CORAL	36	119	0.65790	94
CAS-PH	CORAL	40052	265442	0.112826	231
FTTP-4	Galileo	32757	426826	0.01922	13111
	CORAL	1325	13642	0.01922	65
FTTP-5	CORAL	43105	643339	0.00306	309
FTTP-6	CORAL	1180565	22147378	0.000453	1989
FTTP-C	CORAL	653303	12220653	0.02136	1806
FTTP-A	Galileo	32757	426826	0.0167	13111
	CORAL	19367	154566	0.0167	240
NDPS	CORAL	61	169	[0.00586, 0.00598]	266

TABLE I  
RESULTS OF THE CASE STUDIES.

**The cardiac assist system.** This system, taken from [5] and shown in Figure 18 (right), consists of three separate modules (i.e. CPU, motors, and pump units). Table II shows the failure rates of the various components. In addition, B is a warm spare

with a dormancy factor  $\alpha = 0.5$ , and MB and PS are cold spares (i.e.  $\alpha = 0$ ). During analysis, Galileo DIFTree modularizes the DFT into three independent modules (namely CPU, motors, and pump units) and generates a separate CTMC for each one of them. The biggest CTMC has 8 states. The CTMCs' results are then combined (through the top OR gate) using BDDs. Using the compositional-aggregation approach, and without modularization, the biggest I/O-IMC encountered has 36 state. The results of CAS are summarized in Table I. Here, clearly Galileo outperforms CORAL because it uses modularization which has not been implemented yet in CORAL. If we switch off modularization in Galileo (i.e. generate a single CTMC for the whole system) then it produces a CTMC with 85 states.

To illustrate the possibility of using phase-type distributions we have modified the CAS case study by replacing BEs with PHBEs (case CAS-PH). In this case all basic events occur after a delay governed by an Erlang distribution with four phases and the same expectation as the exponential delay from the CAS case study (for instance, PHBE CS is governed by a 4-phase Erlang with rate parameter 0.8 instead of an exponential distribution with rate 0.2). For the passive delay of warm spare B we also use an Erlang distribution with four phases and the same expectation as the passive exponential delay.

Component	CS	SS	P	B	MA	MB	MS	PA	PB	PS
Rate	0.2	0.2	0.5	0.5	1	1	0.01	1	1	1

TABLE II  
FAILURE RATES FOR CAS.

**The fault tolerant parallel processors.** This system, taken from [5], consists of 16 processors divided into 4 logical groups. In each group, a processor is used as a shared cold spare. A network element (NE) physically connects 1 processor in each group (thus there are 4 NEs) to the rest of the system. The failure of an NE makes the 4 processors connected to it unavailable (i.e. essentially failed). The requirement is to have at least two processors operational in each group. The DFT is shown in Figure 19, where the processors are denoted with  $A_i, B_i, C_i, D_i$  and the network elements with  $N$ . All network elements have a failure rate equals to 0.017, and all processors have a failure rate equals 0.11. The four spare processors are cold spares (i.e.  $\alpha = 0$ ).

To illustrate even further the state-space explosion problem, we took the FTTP system and made it larger by considering 5 (FTTP-5) respectively 6 (FTTP-6) processors in each group. For these case studies the Galileo tool runs out of memory. The (FTTP-C) case study is an extension to the FTTP system where each of the processors (e.g.  $A_1$ ) is replaced by a complex element which consists of an OR-gate with two inputs: a BE 'CPU' and a spare gate 'memory' with primary BE 'M1' and cold spare BE 'M2'. All these BEs have failure rate 0.11. No results are available from Galileo DIFTree as this kind of extended DFTs can not be handled by Galileo. It is interesting to investigate how dependent the success of the compositional aggregation technique is on the symmetries in the DFT models. In the case study FTTP-A we have therefore considered a variant of the FTTP case study where almost all basic events have different rates. The rates used are given in Table III. We see that indeed the models encountered are larger by a factor of 15, but still the computational time required is manageable at under 5 minutes.

**Non-deterministic pump system (NDPS).** This case study il-

BE	Rate	BE	Rate	BE	Rate	BE	Rate	BE	Rate
$N_1$	3.4	$A_1$	5.5	$B_1$	11	$C_1$	6.6	$D_1$	7.7
$N_2$	2.9	$A_2$	7.7	$B_2$	15.4	$C_2$	99	$D_2$	11
$N_3$	2.4	$A_3$	14	$B_3$	28	$C_3$	16	$D_3$	22
$N_4$	1.7	$A_4$	18	$B_4$	36	$C_4$	22	$D_4$	30

TABLE III

FAILURE RATES FOR FTTP-A. RATES ARE SCALED UP BY A FACTOR  $10^5$ .

illustrates how we deal with non-determinism in DFT models. The I/O-IMC model of the DFT is generated as usual and then converted into a CTMDP following [17]. This CTMDP is then analyzed using the MRMC [19] tool to obtain maximum and minimum unreliability. The DFT is a simple pumping system with five pumps: main pumps  $A$  and  $B$ , auxiliary pumps  $C$ ,  $D$  and spare auxiliary pump  $E$  (see Figure 20). BE  $A$  has rate 0.05 and BE  $B$  has rate 0.1; BEs  $C$  and  $D$  both have rate 0.2 and BE  $E$  has rate 0.3 with dormancy  $\alpha = 0$ . Finally BE  $X$  has rate 0.002. This DFT is non-deterministic since  $C$  and  $D$  may fail at the same time due to their dependence on BE  $X$ . It is then undetermined whether spare  $E$  replaces  $C$  or  $D$ . We can see that the maximum and minimum unreliabilities are not far apart, this is caused by the low probability of the ‘problematic’ BE  $X$ .




Fig. 20. DFT for NDPS case study

#### X. CONCLUSION AND FUTURE WORK

In this paper we have formalized the syntax and semantics of DFTs and introduced a DFT analysis framework based on I/O-IMCs, increasing the DFT modularity both at the analysis level and the model-building level. We have also demonstrated the ease with which one can define new DFT elements and provided examples of such extensions. Finally, we have built a prototype tool for analysing DFTs named CORAL and run some experiments to compare our methodology and results to the Galileo DFT tool.

Areas of future research include: (1) From a process algebra point of view, we would like to achieve even more drastic state-space reduction using more suitable aggregation techniques. (2) Adapt the I/O-IMC approach to other formalisms, such as the Architecture Analysis and Design Language [1]. A first step towards this goal has been made in [4].

**Acknowledgements:** We thank Holger Hermanns and the anonymous reviewers of this paper for their helpful comments.

#### REFERENCES

[1] As-2 Embedded Computing Systems Committee. Architecture analysis & design language (AADL), January 2009. Number: AS5506, Revision: A.  
 [2] C. Baier, H. Hermanns, J. P. Katoen, and B. R. Haverkort. Efficient computation of time-bounded reachability probabilities in uniform continuous-time markov decision processes. *Theor. Comput. Sci.*, 345(1):2–26, 2005.

[3] E. Böde, M. Herbstritt, H. Hermanns, S. Johr, T. Peikenkamp, R. Pulungan, J. Rakow, R. Wimmer, and B. Becker. Compositional dependability evaluation for statemate. *IEEE Trans. Software Eng.*, 35(2):274–292, 2009.  
 [4] H. Boudali, P. Crouzen, B. R. Haverkort, M. Kuntz, and M. Stoelinga. Architectural dependability evaluation with arcade. In *Proceedings of the 38th IEEE/IFIP International Conference on Dependable Systems and Networks*, pages 512–521, 2008.  
 [5] H. Boudali, P. Crouzen, and M. Stoelinga. A compositional semantics for Dynamic Fault Trees in terms of Interactive Markov Chains. In *Proceedings of the 5th International Symposium on Automated Technology for Verification and Analysis*, pages 441–456. LNCS, 2007.  
 [6] 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.  
 [7] H. Boudali, P. Crouzen, and M. Stoelinga. Dynamic fault tree analysis using input/output interactive markov chains. In *Proceedings of the 37th IEEE/IFIP International Conference on Dependable Systems and Networks*, pages 708–717, June 2007.  
 [8] M. A. Boyd. *Dynamic fault tree models: techniques for analyses of advanced fault tolerant computer systems*. PhD dissertation, Dept. of Computer Science, Duke University, 1991.  
 [9] M. Bravetti and R. Gorrieri. The theory of interactive generalized semi-markov processes. *Theoretical Computer Science*, 282(1):5–32, 2002.  
 [10] D. Coppit, K. J. Sullivan, and J. B. Dugan. Formal semantics of models for computational engineering: A case study on dynamic fault trees. In *Proceedings of the International Symposium on Software Reliability Engineering*, pages 270–282. IEEE, Oct 2000.  
 [11] N. Coste, H. Garavel, H. Hermanns, R. Hersemeule, Y. Thonnart, and M. Zidouni. Quantitative evaluation in embedded system design: Validation of multiprocessor multithreaded architectures. In *Design, Automation and Test in Europe*, pages 88–89, 2008.  
 [12] P. Crouzen, H. Hermanns, and L. Zhang. On the minimisation of acyclic models. In *Proceedings of the 19th International Conference on Concurrency Theory*, pages 295–309, 2008.  
 [13] J. B. Dugan, S. J. Bavuso, and M. A. Boyd. Dynamic fault-tree models for fault-tolerant computer systems. *IEEE Transactions on Reliability*, 41(3):363–377, September 1992.  
 [14] J. B. Dugan, B. Venkataraman, and R. Gulati. DIFTree: a software package for the analysis of dynamic fault tree models. In *Reliability and Maintainability Symposium*, pages 64–70, Jan 1997.  
 [15] H. Garavel, F. Lang, R. Mateescu, and W. Serwe. Cadp 2006: A toolbox for the construction and analysis of distributed processes. In *Proceedings of the 19th International Conference on Computer Aided Verification (CAV)*, 2006.  
 [16] H. Hermanns. *Interactive Markov Chains*, volume 2428/2002 of *Lecture Notes in Computer Science*. Heidelberg: Springer Berlin, 2002.  
 [17] H. Hermanns and S. Johr. Uniformity by construction in the analysis of nondeterministic stochastic systems. In *Proceedings of the 37th IEEE/IFIP International Conference on Dependable Systems and Networks*, pages 718–728, 2007.  
 [18] H. Hermanns and J. P. Katoen. Automated compositional Markov chain generation for a plain-old telephone system. *Science of Computer Programming*, 36(1):97–127, 2000.  
 [19] J.-P. Katoen, M. Khattri, and I. S. Zapreev. A markov reward model checker. In *Proceedings of the 2nd Second International Conference on the Quantitative Evaluation of Systems*, pages 243–244, 2005.  
 [20] N. A. Lynch and M. R. Tuttle. An introduction to input/output automata. *CWI Quarterly*, 2(3):219–246, 1988.  
 [21] R. Manian, J. B. Dugan, D. Coppit, and K. J. Sullivan. Combining various solution techniques for dynamic fault tree analysis of computer systems. *IEEE International High-Assurance Systems Engineering Symposium*, 3:21–28, 1998.  
 [22] M. F. Neuts. *Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach*. Dover, 1981.  
 [23] Relex, Fault Tree and Event Tree software. <http://www.relex.com/products/ftaeta.asp>. Relex Website.  
 [24] W. J. Stewart. *Introduction to the Numerical Solution of Markov Chains*. Princeton University Press, 1994.  
 [25] K. J. Sullivan, J. B. Dugan, and D. Coppit. The Galileo fault tree analysis tool. In *Proceedings of the 29th Annual International Symposium on Fault-Tolerant Computing*, pages 232–235. IEEE Computer Society, June 1999.  
 [26] K. K. Vemuri, J. B. Dugan, and K. J. Sullivan. Automatic synthesis of fault trees for computer-based systems. *IEEE Transactions on Reliability*, 48(4):394–402, December 1999.



- [27] W. E. Vesely, F. F. Goldberg, N. H. Roberts, and D. F. Haasl. Fault tree handbook, NUREG-0492. Technical report, NASA, 1981.

APPENDIX I  
 PROOFS

A. Theorem 1.1

For weak bisimilarity on an I/O-IMC  $\mathcal{P}$ ,  $\approx_{\mathcal{P}}$ , we find:

- 1)  $\approx_{\mathcal{P}}$  is a weak bisimulation on  $\mathcal{P}$  and it is the largest weak bisimulation on  $\mathcal{P}$ .

*Proof:* We now prove that for any I/O-IMC  $\mathcal{P} \approx_{\mathcal{P}}$  is the largest weak bisimulation on  $\mathcal{P}$ . In the following we refer to  $\approx_{\mathcal{P}}$  simply as  $\approx$ . To do this we must prove that  $\approx$  is an equivalence relation by proving that it is reflexive, symmetric and transitive. The proofs for reflexivity and symmetry are quite simple, but the proof for transitivity is somewhat more complex.

For any I/O-IMC  $\mathcal{P}$  it is trivial that the identity relation  $\text{Id}_{\mathcal{P}}$  is a weak bisimulation. For any state  $s$  in  $S_{\mathcal{P}}$  we find that  $(s, s) \in \text{Id}_{\mathcal{P}}$  and thus  $s \approx s$ . In other words  $\approx$  is reflexive. If  $s \approx t$  then there is a weak bisimulation  $R$  with  $(s, t) \in R$ . Because  $R$  is an equivalence relation we also find that  $(t, s) \in R$  and so  $t \approx s$ . We now find that  $s \approx t$  implies  $t \approx s$ . This means that  $\approx$  is symmetric.

Let  $R_1$  and  $R_2$  be two weak bisimulations on  $\mathcal{P}$  and let  $R_1 R_2^*$  be defined as their *recursive composition*:

$$R_1 R_2^* = \{(s, t) \mid \exists n \in \mathbb{N}, x_1, \dots, x_n \in S \cdot (s, x_1) \in R_1 \cup R_2 \wedge \dots \wedge (x_n, t) \in R_1 \cup R_2\}$$

Because  $R_1$  and  $R_2$  are equivalence relations it is trivial that  $R_1 R_2^*$  is reflexive and symmetric. From the definition of  $R_1 R_2^*$  we can also immediately deduce that it is a transitive relation.

We will now prove that  $R_1 R_2^*$  is a weak bisimulation relation. Recall that if  $(s, t)$  is in  $R_1 R_2^*$  then there exist a series of states  $x_1, \dots, x_n$  such that they are stepwise related through either  $R_1$  or  $R_2$ :  $s R_i x_1 R_i x_2 \dots x_n R_i t$  where  $R_i$  is  $R_1$  or  $R_2$ . We find that for any action  $a \in A_{\mathcal{P}}^{\text{all}}$  with  $s \xrightarrow{a} s'$  there is an  $x'_1$  with  $x_1 \xrightarrow{a} x'_1$  and  $(s', x'_1) \in R_1 \cup R_2$ . This in turn means that there is an  $x'_2$  with  $x_2 \xrightarrow{a} x'_2$  and  $(s', x'_2) \in R_1 \cup R_2$ . So we find a series of states  $x'_1, \dots, x'_n, t'$  such that  $(s', x'_1) \in R_1 \cup R_2 \wedge \dots \wedge (x'_n, t') \in R_1 \cup R_2$ . Thus  $t \xrightarrow{a} t'$  with  $(s', t') \in R_1 R_2^*$ , which means that the first clause of Definition 4 holds. To prove that the second clause holds we must first consider the equivalence classes of  $R_1 R_2^*$ .

From the definition we know that  $R_1 \subseteq R_1 R_2^*$  and  $R_2 \subseteq R_1 R_2^*$ . This means that any equivalence class of  $R_1 R_2^*$  is exactly the union of one or more equivalence classes of  $R_1$  and it is also exactly the union of one or more equivalence classes of  $R_2$ . So for any equivalence class  $C$  of  $R_1 R_2^*$  we find that there is a set  $C'_1, C'_2, \dots, C'_n$  of equivalence classes of  $R_1$  and a set  $C''_1, C''_2, \dots, C''_m$  of equivalence classes of  $R_2$  with  $n, m > 0$  such that:

$$C = C'_1 \cup C'_2 \cup \dots \cup C'_n = C''_1 \cup C''_2 \cup \dots \cup C''_m$$

It is trivial that this also holds for the internal backwards closure of an equivalence class of  $R_1 R_2^*$ :

$$C^{\text{int}} = C_1^{\text{int}} \cup C_2^{\text{int}} \cup \dots \cup C_n^{\text{int}} = C_1^{\text{int}} \cup C_2^{\text{int}} \cup \dots \cup C_m^{\text{int}}$$

For a state  $s$  we find that it's equivalence class for  $R_1 R_2^*$ :  $[s]_{R_1 R_2^*}$  contains all states  $t$  such that  $(s, t) \in R_1 R_2^*$ . Because  $(s, t) \in R_1$  implies  $(s, t) \in R_1 R_2^*$  we find that  $[s]_{R_1} \subseteq [s]_{R_1 R_2^*}$  and similarly  $[s]_{R_2} \subseteq [s]_{R_1 R_2^*}$ . Because equivalence classes are disjoint we now find that if an equivalence class of  $R_1 R_2^*$ :  $C \neq [s]_{R_1 R_2^*}$  then  $C \cap [s]_{R_1} = \emptyset$  and also  $C \cap [s]_{R_2} = \emptyset$ .



**Hichem Boudali** received a Ph.D. degree in Computer Engineering from the School of Engineering and Applied Science at the University of Virginia (Charlottesville, VA, USA), in 2005. He received a Computer Science Engineering degree (M.Sc. equivalent) from the Polytechnic School and Applied Sciences at the Universit Libre de Bruxelles (Brussels, Belgium), in 2002. The same year, he obtained a M.Sc. degree in Electrical Engineering from the University of Virginia. He is also a member of IEEE.



**Pepijn Couzen** received a Computer Science Engineering degree (M.Sc. equivalent) from the University of Twente (Enschede, The Netherlands) in 2006. He is now a PhD student in the Dependable Systems and Software group of the Computer Science faculty of Saarland University (Saarbrücken, Germany).



**Mariëlle Stoelinga** received her Ph.D. in Computer Science from the University of Nijmegen, the Netherlands, in 2002. She is currently an assistant professor at the Formal Methods & Tools group at the University of Twente, the Netherlands.

Let  $(s, t) \in R_1 R_2^*$  and let there be a  $s'$  such that  $s \Longrightarrow s'$  and  $s'$  is stable. If  $(s, t) \in R_1 \cup R_2$  then the second clause of Definition 4 holds. Otherwise there exists an  $x_1$  such that  $(s, x_1) \in R_1 \cup R_2$ , so either  $(s, x_1) \in R_1$  or  $(s, x_1) \in R_2$ . For the first case we find that there is an  $x'_1$  such that  $x_1 \Longrightarrow x'_1$  and  $x'_1$  stable and  $\gamma_M(s', C'^{int}) = \gamma_M(x'_1, C'^{int})$ , for all equivalence classes  $C'$  of  $R_1$  except  $[s']_{R_1}$ . For the second case we find that there is an  $x'_1$  such that:  $x_1 \Longrightarrow x'_1$  and  $x'_1$  stable and  $\gamma_M(s', C'^{int}) = \gamma_M(x'_1, C'^{int})$ , for all equivalence classes  $C''$  of  $R_2$  except  $[s']_{R_2}$ . So we find that there is always an  $x'_1$  such that  $x_1 \Longrightarrow x'_1$  and  $x'_1$  stable and  $\gamma_M(s', C'^{int}) = \gamma_M(x'_1, C'^{int})$ , for all equivalence classes  $C$  of  $R_1 R_2^*$  except  $[s']_{R_1 R_2^*}$ . Note that the equivalence classes  $[s']_{R_1}$  and  $[s']_{R_2}$  must be contained in  $[s']_{R_1 R_2^*}$  and not in any other equivalence class of  $R_1 R_2^*$ . The same holds for the pairs  $(x_1, x_2), (x_2, x_3), \dots, (x_n, t)$ , so in the end we find that the second clause of Definition 4 holds for  $R_1 R_2^*$  so  $R_1 R_2^*$  is a weak bisimulation.

If we now have that  $s \approx_{\mathcal{P}} x$  and  $x \approx_{\mathcal{P}} t$  then there exist two weak bisimulations  $R_1$  and  $R_2$  such that  $(s, x) \in R_1$  and  $(x, t) \in R_2$ . Then  $(s, t) \in R_1 R_2^*$  and because  $R_1 R_2^*$  is a weak bisimulation on  $P$  we find that  $s \approx_{\mathcal{P}} t$  which means that  $\approx_{\mathcal{P}}$  is an equivalence relation.

Because  $\approx_{\mathcal{P}}$  is the union of all weak bisimulations on  $\mathcal{P}$  it is trivial that  $\approx_{\mathcal{P}}$  itself is a weak bisimulation now that we have proved that it is an equivalence relation. That  $\approx$  is the largest weak bisimulation on  $\mathcal{P}$  follows immediately from its definition. ■

### B. Theorem 1.2 and 1.3

Weak bisimilarity is substitutive with parallel composition.

$$\begin{aligned} P_1 \approx P_2 & \text{ implies } P_1 || P_3 \approx P_2 || P_3 \\ P_1 \approx P_2 & \text{ implies } P_3 || P_1 \approx P_3 || P_2 \end{aligned}$$

*Proof:* We now prove that weak bisimilarity is substitutive with parallel composition. Let  $P_1$  and  $P_2$  be I/O-IMCs with identical action signatures, let  $P$  be their union and let  $x \approx_P y$ . This means that there is a weak bisimulation  $R$  on  $P$  such that  $(x, y) \in R$ . We now define the relation  $R'$  as follows:

$$R' = \{(x || z, y || z) \mid (x, y) \in R \wedge z \in S_{P_3}\}$$

Because  $R$  is an equivalence relation on  $P$  it is trivial that  $R'$  is also an equivalence relation on  $P || P_3$ .

Let  $(s || u, t || u) \in R'$ , then  $(s, t)$  is in  $R$ . Now let  $s || u \xrightarrow{a} s' || u'$  with  $a \in A^{all}$ . This means that there exist  $s'' || u''$  and  $s''' || u'''$  such that  $s \Longrightarrow s''$ ,  $u \Longrightarrow u''$ ,  $s'' || u'' \xrightarrow{a} s''' || u'''$ ,  $s''' \Longrightarrow s'$  and  $u''' \Longrightarrow u'$ . From the definition of parallel composition we can deduce that  $s'' || u'' \xrightarrow{a} s''' || u'''$  implies that one of the following holds:  $s'' \xrightarrow{a} s''' \wedge u'' = u'''$  or  $s'' \xrightarrow{a} s''' \wedge u'' \xrightarrow{a} u'''$  or  $s'' \xrightarrow{a} s''' \wedge u'' \xrightarrow{a} u''' \wedge a \in A(P) \wedge a \in A(P_3)$ . In other words either  $s \xrightarrow{a} s' \wedge u \Longrightarrow u' \wedge a \in A(P) \wedge a \notin A(P_3)$  or  $u \xrightarrow{a} u' \wedge s \Longrightarrow s' \wedge a \in A(P_3) \wedge a \notin A(P)$  or  $s \xrightarrow{a} s' \wedge u \Longrightarrow u' \wedge a \in A(P) \wedge a \in A(P_3)$ . Because  $(s, t) \in R$  we find that there is a  $t'$  with  $(s', t') \in R$  such that:  $t \xrightarrow{a} t' \wedge u \Longrightarrow u' \wedge a \in A(P) \wedge a \notin A(P_3)$  or  $u \xrightarrow{a} u' \wedge t \xrightarrow{a} t' \wedge a \in A(P_3) \wedge a \notin A(P)$  or  $t \xrightarrow{a} t' \wedge u \Longrightarrow u' \wedge a \in A(P) \wedge a \in A(P_3)$ . From the definition of parallel composition we now know that:  $t || u \xrightarrow{a} t' || u'$  and from the definition of  $R'$  we know that  $(s' || u', t' || u')$  is in  $R'$ . This means that the first clause of Definition 4 holds for  $R'$ .

From the definition of  $R'$  we can deduce that it has the following equivalence classes:

$$S/R' = \{\{x || y \mid x \in C\} \mid C \in S_P/R, y \in S_{P_3}\}$$

Because  $(s, s') \in R$  if and only if  $\forall u \in S_{P_3} \cdot (s || u, s' || u) \in R'$  we find that for an equivalence class  $[s || u]_{R'}: [s || u]_{R'} = \{s' || u \mid s' \in [s]_R\}$ .

Now let  $(s || u, t || u) \in R'$  with  $s || u \Longrightarrow s' || u'$  and  $s' || u'$  stable and let  $C$  be an equivalence class of  $R'$  with  $C \neq [s' || u']_{R'}$ . From the definition of parallel composition we know that this means that  $s \Longrightarrow s'$  and  $u \Longrightarrow u'$  and both  $s'$  and  $u'$  stable. We know that each equivalence class of  $R'$  is derived from an equivalence class of  $R$  and a state in  $S_{P_3}$ . Let  $D$  be that equivalence class of  $R$  and let  $y$  be that state. So  $C = \{s || y \mid s \in D\}$  and  $u' = y$  implies  $D \neq [s']_R$ . Note that if  $u' \neq y$ ,  $s' || u'$  cannot be in  $C$ . Now we find the following for the internal backward closure of  $C$ :

$$\begin{aligned} C^{int} &= \{x' || y' \mid \exists x || y \in C \cdot x' || y' \Longrightarrow x || y\} \\ &= \{x' || y' \mid \exists x \in D \cdot x' \Longrightarrow x \wedge y' \Longrightarrow y\} \end{aligned}$$

If we now look at the cumulative rate of  $s' || u'$  to  $C^{int}$  we find:

$$\begin{aligned} \gamma_M(s' || u', C^{int}) &= \sum \{|\lambda \mid \exists x \in D \cdot s' \xrightarrow{\lambda} x' \wedge x' \Longrightarrow x \wedge u' \Longrightarrow y\}| + \\ & \quad \sum \{|\lambda \mid \exists x \in D \cdot u' \xrightarrow{\lambda} y' \wedge s' \Longrightarrow x \wedge y' \Longrightarrow y\}| \\ &= \sum \{|\lambda \mid \exists x \in D \cdot s' \xrightarrow{\lambda} x' \wedge x' \Longrightarrow x \wedge u' = y\}| + \\ & \quad \sum \{|\lambda \mid \exists x \in D \cdot u' \xrightarrow{\lambda} y' \wedge s' = x \wedge y' \Longrightarrow y\}| \\ & \quad \text{(Because } s' \text{ and } u' \text{ are stable)} \\ &= \begin{cases} \gamma_M(s', D^{int}) + \gamma_M(u', \{y\}^{int}) & , \text{ if } u' = y \wedge s' \in D \\ \gamma_M(s', D^{int}) & , \text{ if } u' = y \wedge s' \notin D \\ \gamma_M(u', \{y\}^{int}) & , \text{ if } u' \neq y \wedge s' \in D \\ 0 & , \text{ if } u' \neq y \wedge s' \notin D \end{cases} \\ &= \begin{cases} \gamma_M(s', D^{int}) & , \text{ if } u' = y \wedge s' \notin D \\ \gamma_M(u', \{y\}^{int}) & , \text{ if } u' \neq y \wedge s' \in D \\ 0 & , \text{ if } u' \neq y \wedge s' \notin D \end{cases} \\ & \quad \text{(Because } u' = y \text{ implies } D \neq [s']_R.) \end{aligned}$$

Because  $(s, t) \in R$  we know that  $s \Longrightarrow s'$  and  $s'$  stable imply that there is a  $t'$  such that  $t \Longrightarrow t'$  and  $t'$  stable with  $(s', t') \in R^{11}$  and  $\gamma_M(s', D^{int}) = \gamma_M(t', D^{int})$  for every equivalence class  $D$  of  $R$  except  $[s']_R$ . It is furthermore trivial that  $s'$  and  $t'$  are in the same equivalence class of  $R$ . We now conclude that:

$$\begin{aligned} \gamma_M(s' || u', C^{int}) &= \begin{cases} \gamma_M(s', D^{int}) & , \text{ if } u' = y \wedge s' \notin D \\ \gamma_M(u', \{y\}^{int}) & , \text{ if } u' \neq y \wedge s' \in D \\ 0 & , \text{ if } u' \neq y \wedge s' \notin D \end{cases} \\ &= \begin{cases} \gamma_M(t', D^{int}) & , \text{ if } u' = y \wedge t' \notin D \\ \gamma_M(u', \{y\}^{int}) & , \text{ if } u' \neq y \wedge t' \in D \\ 0 & , \text{ if } u' \neq y \wedge t' \notin D \end{cases} \\ & \quad \text{(Because } s' \notin D \text{ implies that } D \neq [s']_R.) \\ &= \gamma_M(t' || u', C^{int}) \end{aligned}$$

<sup>11</sup>Because  $(s, t) \in R$  and  $s \Longrightarrow s'$  there is a  $y$  such that  $t \Longrightarrow y$  and  $(s', y) \in R$ . For this pair clause 1 and 2 of weak bisimulation must hold once more. Eventually we will find  $t'$  this way.

So clause two of Definition 4 holds for  $R'$ . For any state  $z \in S_{P_3}$  we find that if  $(x, y) \in R$  then  $(x||z, y||z) \in R'$  and since  $R'$  is a weak bisimulation  $x||z \approx y||z$  thus proving Theorem 1.2. We find a similar proof for Theorem 1.3 because parallel composition is completely symmetric. This then proves that weak bisimilarity is substitutive with parallel composition. ■

### C. Theorem 1.4

*Weak bisimilarity is substitutive with hiding.*

$P_1 \approx P_2$  implies  $\text{hide } a_1, \dots, a_n \text{ in } P_1 \approx \text{hide } a_1, \dots, a_n \text{ in } P_2$

*Proof:* Let  $P_1$  and  $P_2$  be I/O-IMCs with identical action signatures, let  $P$  be their union. Then if  $P_1 \approx P_2$  there is a weak bisimulation  $R$  on  $P$  such that  $\forall s \in S_{P_1} \cdot \exists t \in S_{P_2} \cdot (s, t) \in R$ . Let  $B \subseteq A^V(P) = \{a_1, \dots, a_n\}$  be a subset of the visual actions of  $P$  and let  $R'$  be a binary relation defined as follows:

$$R' = \{(\text{hide } B \text{ in } s, \text{hide } B \text{ in } t) \mid (s, t) \in R\}$$

It is trivial that  $R'$  is an equivalence relation on  $\text{hide } B \text{ in } P$ . Let  $\text{hide } B \text{ in } s \xrightarrow{a} \text{hide } B \text{ in } s'$ . From the definition of hiding we can immediately deduce that there exists an  $n \in \mathbb{N}$  with  $n \geq 0$ , a series of states  $x_1, \dots, x_n$  and a series of actions  $b_1, \dots, b_n \in B$  such that:

$$s \xrightarrow{b_1} x_1 \xrightarrow{b_2} \dots x_k \xrightarrow{a} x_{k+1} \dots \xrightarrow{b_{n-1}} x_n \xrightarrow{b_n} s'$$

Because  $(s, t) \in R$  we find a series of states  $x'_1, \dots, x'_{n-1}$  such that:

$$t \xrightarrow{b_1} x'_1 \xrightarrow{b_2} \dots x'_k \xrightarrow{a} x'_{k+1} \dots \xrightarrow{b_{n-1}} x'_n \xrightarrow{b_n} t'$$

It follows that  $\text{hide } B \text{ in } t \xrightarrow{a} \text{hide } B \text{ in } t'$  so the first clause of Definition 4 holds for  $R'$ .

Let  $\text{hide } B \text{ in } s \xrightarrow{} \text{hide } B \text{ in } s'$  and  $\text{hide } B \text{ in } s'$  stable. This means that there is an  $n \in \mathbb{N}$ , a series of actions  $b_1, \dots, b_n$  and a series of states  $x_1, \dots, x_n$  such that  $s \xrightarrow{b_1} x_1 \dots x_n \xrightarrow{} s'$  with  $b_1, \dots, b_n \in B$ . For  $t$  we then find the same series of actions and a weakly bisimilar set of states. We particularly find an  $x'_n$  such that  $(x_n, x'_n) \in R$ . Because  $x_n \xrightarrow{} s'$  and  $s'$  stable we find a  $t'$  such that  $x'_n \xrightarrow{} t'$  and  $t'$  stable with  $\gamma_M(s', D) = \gamma_M(t', D)$  for all equivalence classes  $D$  of  $R$  except for  $[s']_R$  and  $(s', t') \in R$ . Because  $\text{hide } B \text{ in } s'$  is stable  $s'$  and  $t'$  have no outgoing transitions labelled with an action in  $B$ .

From the definition of  $R'$  we see that its equivalence classes can be defined as follows:

$$\{\{\text{hide } B \text{ in } x \mid x \in D\} \mid D \in S/R\}$$

Note that for any state  $x$  in  $S$  we find that  $[\text{hide } B \text{ in } x]_{R'} = \{\text{hide } B \text{ in } y \mid y \in [x]_R\}$ .

Let  $C$  be an equivalence class of  $R'$  with  $C \neq [\text{hide } B \text{ in } s']_{R'}$  and let  $D$  be the corresponding equivalence class of  $R$ , thus  $D \neq [s']_R$ . We now find the following for the cumulative rate

of  $\text{hide } B \text{ in } s'$  to  $C^{int}$ :

$$\begin{aligned} & \gamma_M(\text{hide } B \text{ in } s', C^{int}) \\ &= \{|\lambda \mid \exists(\text{hide } B \text{ in } x') \in S, (\text{hide } B \text{ in } x) \in C \cdot \\ & \quad \text{hide } B \text{ in } s' \xrightarrow{\lambda}^M \text{hide } B \text{ in } x' \wedge \\ & \quad \text{hide } B \text{ in } x' \xrightarrow{} \text{hide } B \text{ in } x\}\} \\ &= \{|\lambda \mid \exists(\text{hide } B \text{ in } x') \in S, x \in D \cdot \\ & \quad \text{hide } B \text{ in } s' \xrightarrow{\lambda}^M \text{hide } B \text{ in } x' \wedge \\ & \quad \text{hide } B \text{ in } x' \xrightarrow{} \text{hide } B \text{ in } x\}\} \\ &= \{|\lambda \mid \exists x' \in S, x \in D \cdot s' \xrightarrow{\lambda}^M x' \wedge \\ & \quad \text{hide } B \text{ in } x' \xrightarrow{} \text{hide } B \text{ in } x\}\} \\ &= \gamma_M(s', D^{int+}) \end{aligned}$$

Here  $D^{int+} = \{x' \mid \exists x \in D \cdot \text{hide } B \text{ in } x' \xrightarrow{} \text{hide } B \text{ in } x\}$ . Note that  $D^{int+} \supseteq D^{int}$ . For some  $n \in \mathbb{N}$  let  $E_1, \dots, E_n$  be the equivalence classes of  $R$  that can reach  $D$  with a number of weak transitions labelled with actions in  $B$  or with a weak move (note that this means that  $D$  itself is always in  $E_1, \dots, E_n$ ). In other words we find that:

$$\forall E_x \cdot (\forall e \in E_x \cdot \exists m \in \mathbb{N}, f_1, \dots, f_m \in S, b_1, \dots, b_m \in B \cdot m \geq 0 \wedge e \xrightarrow{b_1} f_1 \dots f_{m-1} \xrightarrow{b_m} f_m \wedge f_m \in D)$$

When we hide the set of actions  $B$  in the states of the equivalence classes  $E_x$  we find that they move internally to states in equivalence class  $C$ . Note that because  $\text{hide } B \text{ in } s'$  is stable and  $C \neq [\text{hide } B \text{ in } s']_{R'}$ ,  $[s']_R$  can not be in  $E$  (that would either mean that  $\text{hide } B \text{ in } s'$  moves internally to a state in  $C$  making it unstable or that  $\text{hide } B \text{ in } s'$  is in  $C$  but this is not possible because  $C \neq [\text{hide } B \text{ in } s']_{R'}$ ). So  $D^{int+} = \bigcup_{E_i} E_i^{int}$ . Now we find that  $\gamma_M(s', D^{int+}) = \gamma_M(t', D^{int+})$  (Recall that we found earlier that  $\gamma_M(s', D) = \gamma_M(t', D)$  for all equivalence classes  $D$  of  $R$  except  $[s']_R$  which is not in  $E$ ). In the same way we showed that  $\gamma_M(\text{hide } B \text{ in } s', C)$  equals  $\gamma_M(s', D^{int+})$ , we can now show that  $\gamma_M(t', D^{int+})$  equals  $\gamma_M(\text{hide } B \text{ in } t', C^{int})$ . This proves that the second clause of Definition 4 holds for  $R'$ . This means that  $\text{hide } B \text{ in } s \approx \text{hide } B \text{ in } t$  when  $s \approx t$  and thus weak bisimilarity is substitutive with hiding. ■