

Probabilistic Model Checking for Uncertain Scenario-Aware Data Flow

JOOST-PIETER KATOEN and HAO WU, Software Modelling and Verification Group,
RWTH Aachen University

The Scenario-Aware Dataflow (SADF) model is based on concurrent actors that interact via channels. It combines streaming data and control to capture scenarios while incorporating hard and soft real-time aspects. To model data-flow computations that are subject to uncertainty, SADF models are equipped with random primitives. We propose to use probabilistic model checking to analyze uncertain SADF models. We show how measures such as expected time, long-run objectives like throughput, as well as timed reachability—can a given system configuration be reached within a deadline with high probability?—can be automatically determined. The crux of our method is a *compositional* semantics of SADF with exponential agent execution times combined with *automated abstraction* techniques akin to partial-order reduction. We present the semantics in detail and show how it accommodates the incorporation of execution platforms, enabling the analysis of energy consumption. The feasibility of our approach is illustrated by analyzing several quantitative measures of an MPEG-4 decoder and an industrial face recognition application.

CCS Concepts: • **Mathematics of computing** → **Markov processes**; • **Software and its engineering** → **Data flow architectures**; **Model checking**; • **General and reference** → *Performance*; • **Hardware** → *Power estimation and optimization*

Additional Key Words and Phrases: Data-flow computation, digital signal processing, embedded systems, energy consumption, Markov (reward) automata, model checking

ACM Reference Format:

Joost-Pieter Katoen and Hao Wu. 2016. Probabilistic model checking for uncertain scenario-aware data flow. ACM Trans. Des. Autom. Electron. Syst. 22, 1, Article 15 (September 2016), 27 pages.
DOI: <http://dx.doi.org/10.1145/2914788>

1. INTRODUCTION

Synchronous Data Flow. Embedded systems such as smart phones, TV sets, and modern printing devices typically involve intensive multimedia processing. Current applications require massive data signal processing facilities like photo editing, face recognition, audio and video streaming, and need to adhere to demanding quality of service (QoS) requirements. Such data processing facilities are adequately described in data-flow languages such as Synchronous Dataflow (SDF) [Lee and Messerschmitt 1987], a language which is equally expressive as weighted marked graphs. SDF has been used extensively [Bhattacharyya et al. 2013; Eker et al. 2003; Sriram and

This work is supported by the European Union's Seventh Framework Programme for Research (FP7), under grant no. 318490 (FP7-ICT-2011-8, the SENSATION project). This work is partially based on the authors' previous work, "Exponentially timed SADF: Compositional semantics, reductions, and analysis," published at EMSOFT 2014 in New Delhi, India.

Authors' addresses: J.-P. Katoen and H. Wu, Software Modelling and Verification Group, Computer Science Department, RWTH Aachen University, D-52066, Aachen, Germany; emails: {katoen, hao.wu}@cs.rwth-aachen.de.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2016 ACM 1084-4309/2016/09-ART15 \$15.00

DOI: <http://dx.doi.org/10.1145/2914788>

Bhattacharyya 2009] and originates from the field of digital signal processing, where a coherent set of interacting tasks with different execution frequencies are to be performed in a distributed and pipelined fashion by a number of parallel computing resources as provided, e.g., by Multi-Processor Systems-on-Chips (MPSoC). Modern embedded applications are very dynamic in the sense that their execution costs, such as memory usage and energy consumption, vary substantially depending on their context (like input data and quality level).

Uncertainty in Scenario-Aware Data Flow. Data-flow languages [Buck and Lee 1993; Geilen and Basten 2004; Kahn 1974] are intended as an implementation framework for such dynamic applications. However, these languages lack facilities for predicting performance during embedded system design—how to deal with uncertain fluctuations in data availability, data quality, and so on—and have rather limited means to describe dynamic dataflow behavior. The Scenario-Aware Dataflow (SADF) language [Theelen et al. 2006, 2008; Bhattacharyya et al. 2013] extends SDF so as to develop adequate scenario-aware performance models of specifications expressed in other data-flow formalisms. Like most data-flow languages, SADF is based on (asynchronously concurrent) actors that interact via (unbounded) First-in, First-out (FIFO) channels. The novelty of SADF is its combination of streaming data and control to capture scenarios, as well as combining both hard and soft real-time aspects. A recent survey of data-flow formalisms, such as SDF and SADF, as well as their relationship, is given in Bhattacharyya et al. [2013]. This article focuses, in particular, on the possibility in SADF to model *uncertainty*, both in the execution times of processes (called actors), as well as in scenario generation.

Exponentially Timed SADF. This article considers *Exponentially Timed SADF* (called eSADF), a version of SADF in which the duration of all firings of actors are governed by negative exponential probability distributions. eSADF can be considered as an extension of exponentially timed SDF, as originally proposed in Sriram and Bhattacharyya [2009]. An example of an eSADF which models the MPEG-4 decoder is shown in Figure 1. Although the restriction to exponential distributions seems a severe restriction at first sight, there are (at least) three good reasons to consider it. First, this assumption is a rather adequate abstraction when considering that actor firings are typically subject to random fluctuations (e.g., in hardware) and only mean durations are known. Technically speaking, the exponential distribution maximizes the entropy under these assumptions. That is to say, if only mean values are known, the statistically most neutral assumption is to have these phenomena exponentially distributed. Secondly, series-parallel combinations of exponential distributions (so-called phase-type distributions) can approximate any arbitrary continuous probability distribution with arbitrary precision. Our semantic model and analysis algorithms support the analysis of these phase-type distributions. Finally, the use of exponential distributions enables the usage of modern probabilistic model-checking tools for the quantitative analysis of SADF models.

Compositional Semantics. As eSADF is based on asynchronously communicating actors, firings have exponential durations, and sub-scenario selection is based on discrete-time Markov chains, *Markov Automata* (MA) [Deng and Hennessy 2013; Eisentraut et al. 2010] are a natural choice for capturing the operational semantics of eSADF. These automata are transition systems in which the target of an interactive transition is a distribution over states (rather than a single state), and that incorporates random delay transitions to model firing durations. Non-determinism occurs if several interactive transitions emanate from a given state. This article provides *compositional semantics* of eSADF using MA. The compositional aspect naturally reflects the logical structure of the eSADF graph, is easily amenable to single actor replacements—as

just the semantics of that actor is to be adapted while the other automata remain the same—and finally, enables component-wise reduction. The latter technique is important in case the state space of the eSADF graph is too large to be handled. Our compositional semantics allows for replacing the automaton for a few actors by an equivalent, but much smaller, automaton. This technique has, e.g., been exploited in Theelen et al. [2012]. Compositionality in SDF has recently also been exploited for modular code generation [Tripakis et al. 2013]. Our semantics is defined using a succinct process algebra for describing MA [Timmer et al. 2012] in a textual way. As a result, the semantics is relatively easy to comprehend and modular.

Automated Abstraction. MA of realistic, industrially sized eSADF graphs can be huge and too large to be handled. One of the main causes for this state-space explosion is the highly concurrent character of typical data-flow computations in which many agents run in parallel. To diminish this effect on the state-space growth, we use *confluence reduction* [Timmer et al. 2013]—a technique akin to partial-order reduction—that allows for an on-the-fly reduction of the state space. The key of confluence reduction is to detect independent concurrent transitions such that for the analysis only one ordering of concurrent transitions needs to be considered (instead of all possible orderings). We show that all non-determinism in the MA-semantics of eSADF arises from the independent execution of actors, and can (in theory) be eliminated using confluence reduction. As confluence reduction is performed at the language level (i.e., the process algebra) using conservative heuristics, non-determinism may persist after reduction; but if so, it does not influence quantitative measures.

Quantitative Analysis of eSADF Graphs. To analyze eSADF graphs, we exploit recently developed algorithms and software tool-support for verifying MA. The main challenge in MA analysis is the intricate interplay between exponential durations and non-determinism. The latter arises naturally by the concurrent execution of the different actors. It was recently shown that several measures-of-interest, such as expected time and long-run average objectives, can be reduced to efficient analysis techniques for Markov decision processes [Guck et al. 2014a]. In addition, timed reachability objectives—can a certain system configuration be reached within a deadline with a high probability?—were shown to be computable by appropriate discretization techniques [Guck et al. 2014a], and extensions towards the treatment of costs (modeling energy consumption) in MA have become available [Guck et al. 2014b]. Our MA semantics facilitates the quantitative analysis of eSADF graphs using these novel and efficient analysis techniques. To sum up, our semantics is conceptually simple, compositional, and yields a rigorous framework for quantitative analysis of eSADF.

Case Studies. We have developed a prototypical implementation of our approach that maps eSADF graphs (expressed in the XML-format as supported by the SDF³ tool¹) onto MA. These MA can then be analyzed by the analysis tool presented in Guck et al. [2014a]. An extension with confluence reduction enables the minimization of MA. This article shows the feasibility of our approach by presenting two case studies. The *MPEG-4 decoder* is a benchmark SADF case from the literature [Theelen et al. 2006]. We show the effect of confluence reduction and analyze several measures of interest of the MPEG-4 decoder, such as buffer occupancy, throughput, and probability to reach certain buffer occupancies within a given deadline. We compare the results to the analysis results obtained using the SDF³ tool for SADF and to the work by Theelen et al. [2012]. As a second case study, we present the analysis of an industrial *face recognition* application. This is a simplified anonymized case study from the company Recore Systems. This model is substantially larger than the MPEG-4 decoder as it

¹See <http://www.es.ele.tue.nl/sadf/xml.php>.

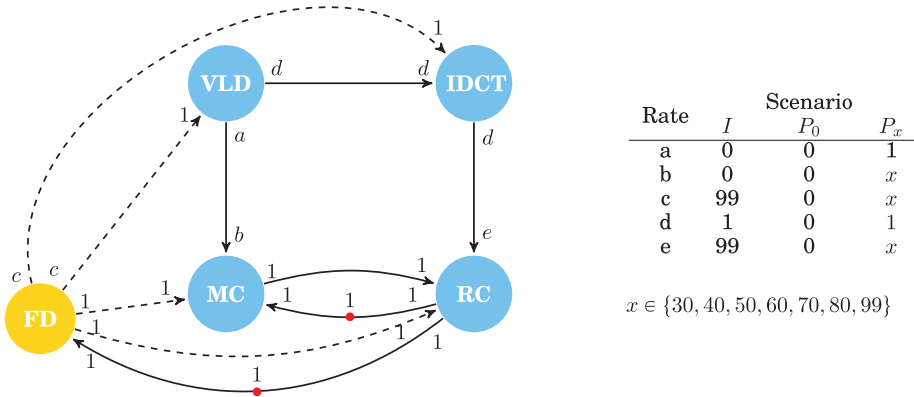


Fig. 1. An SADF model for an MPEG-4 decoder taken from Theelen et al. [2006] where control channels are indicated by dashed arrows, data channels by solid arrows, blue circles represent kernels, and yellow circles represent detectors.

exhibits a high degree of parallelism. We study the quantitative effect of including auto-concurrency, and analyze several metrics that are similar to those for the MPEG-4 decoder. We conclude by presenting an extension of our framework by incorporating execution platforms on which the eSADF agents are supposed to be executed. This allows for studying the quantitative effect of exploiting multi-core platforms, as well as the quantitative impact of Dynamic Voltage and Frequency Scaling (DVFS) and Dynamic Power Management (DPM). We present an *energy analysis* of the MPEG-4 decoder for a simple execution platform.

Organization of the Article. Sections 2 and 3 introduce eSADF and MA, the operational model used for our semantics, respectively. Section 4 presents the compositional eSADF semantics in detail. Section 5 discusses non-determinism and the quantitative evaluation of eSADF using MA. Section 6 presents the MPEG-4 decoder and the face recognition case studies. Section 7 presents our approach to incorporate execution platforms. Section 8 discusses related work and Section 9 concludes. This article is based on the conference paper [Katoen and Wu 2014] and extends it with the industrial face recognition case study, the incorporation of the executing platform, and an energy analysis of the MPEG-4 decoder benchmark.

2. SCENARIO-AWARE DATAFLOW

In this section, we will give a formal definition of an exponentially-timed SADF (eSADF) graph which is based on Theelen [2007]. Figure 1 illustrates an SADF graph of an MPEG-4 decoder. We will first define the channels (depicted as directed arrows) in an eSADF graph, which can transfer information between the agents (called processes and represented as circles). Based on the types of the elements that the channels can transfer, the scenarios are defined. Afterwards, we introduce two distinct kinds of processes, i.e., kernels and detectors. In order to build an eSADF graph correctly, the processes are connected by channels which are assumed to be “type-consistent” with the processes’ ports. In the end, we give the formal definition of an eSADF graph. Let \mathcal{C} denote the set of channels.

Definition 1 (Channel). A channel $c \in \mathcal{C}$ is a (possibly unbounded) FIFO queue to carry information modeled by tokens of a certain type (e.g., none, integers, Boolean, symbols, etc.). Based on these types, we distinguish channels by:

- data* channels whose type is none, i.e., its tokens are placeholders and not valued, and
- control* channels whose type is not none.

The justification of having two types of channels is that data channels are like the channels in traditional SDF [Lee and Messerschmitt 1987], whereas control channels are key features in (e)SADF to carry the scenario (control) information between processes. We, therefore, have $\mathcal{C} = \mathcal{DC} \cup \mathcal{CC}$ with \mathcal{DC} the set of data channels, \mathcal{CC} the set of control channels, and $\mathcal{DC} \cap \mathcal{CC} = \emptyset$. For a control channel $cc \in \mathcal{CC}$, we denote by Σ_{cc} the type of cc , which is the set of elements (values of tokens) that can be carried in cc , e.g., $\Sigma_{cc} = \{a, \dots, z\}$ or $\Sigma_{cc} = \mathbb{B} = \{0, 1\}$. Based on these types of channels, we can now define the scenario for an ordered set C of control channels as follows.

Definition 2 (Scenario). For an ordered set $C = (cc_1, \dots, cc_k)$ of control channels, a *scenario* is an ordered k -tuple $(\sigma_1, \dots, \sigma_k)$, where σ_i is a value of type Σ_{cc_i} for channel $cc_i \in C$. The set of possible scenarios for C is denoted by $\Sigma_C = \prod_{cc \in C} \Sigma_{cc}$.

If we consider $C = (cc_1, cc_2)$ with $k = 2$, $\Sigma_{cc_1} = \{s, t\}$ and $\Sigma_{cc_2} = \{1, 2\}$, for example, the possible scenario set $\Sigma_C = \{(s, 1), (s, 2), (t, 1), (t, 2)\}$. The set of scenarios can be used to transfer the control information.

Now we define kernels and detectors, which can be treated as the “executable” processes in eSADF. Processes are connected with each other by channels through their ports. We denote the set of kernels as \mathcal{K} and the set of detectors as \mathcal{D} , with $\mathcal{K} \cap \mathcal{D} = \emptyset$. The set of processes $\mathcal{P} = \mathcal{K} \cup \mathcal{D}$.

Definition 3 (Kernel). A *kernel* K is a pair (P, S) with

- P is a *set of ports* partitioned into the set P_I, P_O , and P_C of input, output, and control ports, respectively;
- S is the *scenario setting* (Σ, R, E) (a triple) with:
 - ▷ $\Sigma = \prod_{p_c \in P_C} \Sigma_{p_c}$ is the *set of scenarios* in which kernel K can operate, where Σ_{p_c} is the type (set of values) that is allowed to pass through port p_c ,
 - ▷ $R: \Sigma \times P \rightarrow \mathbb{N}$, the *consumption/production rate function*, such that $R(\sigma, p_c) = 1$ for any control port p_c and scenario σ ,
 - ▷ $E: \Sigma \rightarrow \mathbb{R}_{>0}$, the *execution rate function*, i.e., $E(\sigma) = r$ means that the execution time of scenario σ by kernel K is exponentially distributed with mean $1/r$.

Intuitively, for a kernel K , the possible scenarios in which it can operate are given by the values of ordered tuples from its control ports. The consumption/production rate function gives the number of tokens to be consumed or produced after K 's execution in such scenarios into K 's ports, accordingly. Since at each time we only need one token from each control port to select the scenario, $R(\sigma, p_c)$ equals 1 for every control channel. For data channels, the rate function may be arbitrarily large (but finite).

Detectors are not only an “executable” component similar to a kernel, but are equipped with more features, such as capturing the (sub-)scenario occurrence and sending control information by the generation of “scenario” tokens. Detectors do not operate in scenarios but in sub-scenarios which are determined by scenarios received via the control ports.

Definition 4 (Detector). A *detector* D is pair (P, S) with

- P is a *set of ports* partitioned into P_I, P_O , and P_C (as for kernels). P_O is partitioned into P_{O_d} and P_{O_c} , the data and control output ports;
- S is the *sub-scenario setting* $(\Sigma, \Omega, F, R, E, t)$ with
 - ▷ $\Sigma = \prod_{p_c \in P_C} \Sigma_{p_c}$ is the set of scenarios of D ,

- ▷ Ω , a non-empty finite set of sub-scenarios values,
- ▷ $F : \Sigma \rightarrow M$, the random decision function. M is the set of DTMCs $(\mathbb{S}, \iota, \mathbb{P}, \Phi)$ ². Function F associates a DTMC to each scenario $\sigma \in \Sigma$. Here, \mathbb{S} , ι , and \mathbb{P} are the finite state space, initial state, and one-step transition probability function, respectively, and $\Phi : \mathbb{S} \rightarrow \Omega$, the sub-scenario decision function,
- ▷ $R : \Omega \times P \rightarrow \mathbb{N}$, the consumption/production rate function, such that $R(\omega, p_c) = 1$ for any sub-scenario ω and control port p_c ,
- ▷ $E : \Omega \rightarrow \mathbb{R}_{>0}$, the execution rate function as defined for kernels,
- ▷ $t : \Omega \times P_{O_c} \rightarrow \Sigma_{p_{O_c}}$, the sub-scenario token production function. $t(\omega, p_{O_c})$ determines the value of a token to be produced into an output port $p_{O_c} \in P_{O_c}$ after D 's execution in sub-scenario ω . Note that $R(\omega, p_{O_c})$ is the number of valued tokens to be produced.

Intuitively, a detector works in two phases. In the first phase, scenarios of the detector are determined in the same way as for kernels. After the scenario σ is determined, i.e., there is at least one scenario token in each control channel, the DTMC of this scenario is entered, which is defined by $F(\sigma) \in M$. The entry point of such DTMC is its last occupied state. After moving from this state to one of its successor states, the current sub-scenarios and probabilities of new sub-scenarios can be computed by the functions Φ and \mathbb{P} . This is an important difference with SDF: the next sub-scenario in SADF is determined probabilistically, so as to model uncertainty. (This should not be confused with the execution times of actors being randomly distributed.) In the second phase, D will execute in these sub-scenarios. The rate function R and expected execution time function E are then defined for such execution. After the execution of the sub-scenario, the detector D consumes/produces the tokens of its ports. In order to send control information, the valued tokens defined by function t are produced into control output ports and the number of such tokens to be produced is defined by R .

In order to define the eSADF graph, we need a consistent way to connect the processes by using channels. We define the consistency between the channels and a process's ports, which determines whether a channel and two ports are compatible.

Definition 5 (Type Consistency). The channel $c \in \mathcal{C}$ is “type consistent” with two ports in a (two) process(es), if

- c is a data channel connecting a (data) output port of a kernel (detector) with an input port of the same/another process;
- c is a control channel connecting a control output port of a detector with a control port of the same/another process, where the types of the channel and both ports coincide.

An eSADF graph is *type consistent* if all its channels are type consistent.

Definition 6 (eSADF Graph). An eSADF graph $\mathcal{G} = (\mathcal{P}, \mathcal{C}^*)$, where \mathcal{P} is a finite set of processes (vertices), \mathcal{C}^* is the set of “type consistent” channels (edges) of the form $(src(c), tgt(c))$, where $src(c), tgt(c)$ are the source and target ports of channel $c \in \mathcal{C}$, with proper initialization resulting in a (rate) consistent eSADF.³

Example 2.1. Figure 1 illustrates an eSADF graph of an MPEG-4 decoder (for the simple profile [Theelen et al. 2006]). It consists of the kernels VLD, IDCT, MC, and RC, and the detector FD. The decoder can process I and P frames in video streams, which consist of different numbers of macro blocks. This involves operations like *Variable Length Decoding* (VLD), *Inverse Discrete Cosine Transformation* (IDCT), *Motion*

²DTMC = Discrete-Time Markov Chain. This is a finite-state automaton in which all transitions are equipped with a discrete probability.

³Note that “rate” consistent eSADF can be defined as in Theelen et al. [2006]. The initialization of the channels corresponds to the initial number of tokens carried by a channel is declared.

Compensation (MC), and *Reconstruction* (RC). The kernels VLD and IDCT fire once per macro block in a decoded frame, while MC and RC fire once per frame. The detector FD detects the frame-type occurrences in video streams. If it detects an *I* frame, all macro blocks are decoded by VLD and IDCT, and RC will reconstruct the image (without MC). We assume an image size of 176×144 pixels, i.e., an *I* frame consists of 99 macro blocks (cf. the right table in Figure 1). If FD detects a *P* frame, then MC computes the correct position of macro blocks from the previous frame out of motion vectors. We assume that the number of processing motion vectors equals the number of decoding macro blocks and is 0 or $x \in \{30, 40, 50, 60, 70, 80, 99\}$ representing conservative approximations for a range of actual number of motion vectors [Theelen et al. 2006]. This is captured by the scenarios P_0 and P_x , as indicated in Figure 1 (right). All parameters are taken from Theelen et al. [2006] and were obtained by a specific profiling tool.

3. MARKOV AUTOMATA

This section introduces MA, the operational model for eSADF. Stated briefly, an MA is a state-transition system equipped with both stochastic timed, as well as (untimed) non-deterministic transitions. MA are sufficiently expressive to provide a semantics of modeling languages such as dynamic fault trees [Boudali et al. 2009], the domain-specific language Architecture Analysis & Design Language (AADL) [Bozzano et al. 2011], and Generalized Stochastic Petri Nets (GSPNs) [Marsan et al. 1984; Eisentraut et al. 2013]. The treatment of MA is kept brief here; for a full description, we refer to Eisentraut et al. [2010] and Deng and Hennessy [2013]. The crucial observation is that MA are a very natural semantic model of eSADF: concurrency can be modeled by non-determinism (interleaving), exponential delays by Markovian transitions, and the probabilistic selection of sub-scenarios by discrete probabilistic branching. In addition, MA are compositional. To the best of our knowledge, there is no other model that has all these characteristics.

The Syntax of MA. A probability distribution μ over a countable set S is a function $\mu : S \rightarrow [0, 1]$ such that $\sum_{s \in S} \mu(s) = 1$; $\mu(s)$ is the probability assigned by μ to s . Let $\text{Distr}(S)$ be the set of all distributions over set S .

Definition 7 (MA). An MA is a tuple $\mathcal{M} = (S, s_0, \text{Act}, \xrightarrow{\cdot}, \Rightarrow)$, where

- S is a countable set of *states* with *initial state* $s_0 \in S$,
- Act is a countable set of *actions*,
- $\xrightarrow{\cdot} \subseteq S \times \text{Act} \times \text{Distr}(S)$ is the *interactive probabilistic transition relation*,
- $\Rightarrow \subseteq S \times \mathbb{R}_{>0} \times S$ is the *Markovian transition relation*.

The Semantics of MA. We let $\tau \in \text{Act}$ denote the invisible internal action and abbreviate $(s, a, \mu) \in \xrightarrow{\cdot}$ as $s \xrightarrow{a} \mu$, which means if we are at the state s , the probability of reaching state s' by taking action a is $\mu(s')$. The Markovian transition $(s, \lambda, s') \in \Rightarrow$ means that the probability of moving from state s to s' within time t is exponentially distributed and equals $1 - e^{-\lambda t}$. We call λ the *rate* of transition $s \xrightarrow{\lambda} s'$.

Furthermore, we say a state s is *Markovian* iff s has only outgoing Markovian transitions; it is *interactive* iff it has only outgoing interactive probabilistic transitions; it is a *hybrid* state, otherwise. For states s, s' , we let $\mathbf{R}(s, s') = \sum\{\lambda \mid (s, \lambda, s') \in \Rightarrow\}$ be the *rate* between states s and s' , and let $\mathbf{E}(s) = \sum_{s' \in S} \mathbf{R}(s, s')$ be the *exit rate* of s . The probability of leaving the state s within time t is given by $1 - e^{-\mathbf{E}(s)t}$. If $\mathbf{R}(s, s') > 0$ for more than one state s' , a *race* exists between such states after leaving s . For a particular state s' , the probability of winning the race is given by the *branching probability distribution* $\mathbf{P}(s, s') = \frac{\mathbf{R}(s, s')}{\mathbf{E}(s)}$.

Maximal Progress. If in a state both Markovian transitions and internal transitions τ are enabled, the maximal progress assumption asserts that internal transitions τ take precedence over Markovian transitions. This is justified by the fact that the probability of taking a Markovian transition immediately is zero (as $Pr(\text{delay} \leq 0) = 1 - e^{-\lambda \cdot 0} = 0$), whereas the internal transition τ happens immediately (since they cannot be delayed).

Closed MA. A *closed MA* is an MA which is self-contained and has no further synchronization with other components. For a closed MA, since all interactive probabilistic transitions are not subject to any other transitions for synchronization purpose, they cannot be delayed. Therefore, we can safely hide them and turn them into internal transitions τ . A closed MA has no *hybrid* state due to maximal progress. All outgoing transitions of a state in a closed MA are either interactive probabilistic transitions labeled by τ or Markovian transitions (cf. Figure 4(c)). Note that non-determinism exists when there are multiple internal probabilistic transitions emanating from one state.


The sample MA depicted on the right above has seven states (indicated as dots), two Markovian transitions (indicated as double arrows, with rates λ and ν , respectively), and seven interactive transitions. All except one interactive transition yield a successor state with probability one.

MA Process Algebra (MAPA). To generate an MA, a language based on μCRL [Groote and Ponse 1995] called MA Process Algebra (MAPA) was introduced in Timmer et al. [2012]. We use MAPA to define our MA semantics for eSADF.

Definition 8 (Process Terms). A process term in MAPA adheres to the syntax:

$$p ::= Y(\mathbf{t}) \mid c \Rightarrow p \mid p + p \mid \sum_{\mathbf{x}:\mathbf{D}} p \mid a(\mathbf{t}) \sum_{\mathbf{x}:\mathbf{D}} f : p \mid (\lambda) \cdot p$$

Let Prc denote the set of process names, Act denote a countable universe of actions, and \mathbf{x} denote a vector of variables ranging over a (possibly infinite) type domain \mathbf{D} . If the cardinality of \mathbf{x} exceeds one, \mathbf{D} is a Cartesian product. Observe that this matches the scenario definition based on types in eSADF graphs and simplifies our MA definition for eSADF graphs. In the process term $Y(\mathbf{t})$, $Y \in Prc$ is a process name, \mathbf{t} is a vector of data expressions, and $Y(\mathbf{t})$ expresses a process Y initialized by setting its variables to \mathbf{t} . $c \Rightarrow p$ is a guarded expression, which asserts if the condition c (a boolean expression) is satisfied, then the process will behave like the process p , otherwise, it will do nothing. The operator $p_1 + p_2$ expresses a non-deterministic choice between the left process p_1 and the right process p_2 . Further, if there is a (possibly infinite) non-deterministic choice over a data type \mathbf{D} , this is denoted by the term $\sum_{\mathbf{x}:\mathbf{D}} p$. The term $a(\mathbf{t}) \sum_{\mathbf{x}:\mathbf{D}} f : p$ states that the process can perform an $a(\mathbf{t})$ action (an action based on the vector \mathbf{t}) and then resolves a probabilistic choice over \mathbf{D} determined by a predefined function f . The function application $f[\mathbf{x} := \mathbf{d}]$ returns the probability when the variables are evaluated as $\mathbf{d} \in \mathbf{D}$. The term $(\lambda) \cdot p$ expresses that after an exponentially distributed delay with rate $\lambda \in \mathbb{R}_{\geq 0}$, the process will behave like p . We will see later that the MAPA language is concise and expressive enough to handle our eSADF semantics, since it allows processes with different data types and is equipped with both interactive probabilistic transitions and Markovian transitions. MA can now be obtained by the modular construction using MAPA processes through parallel composition, encapsulation, hiding, and renaming operators [Timmer et al. 2012].






Fig. 2. A simplified setting for detectors.

Definition 9 (Parallel Process Terms). An initial process in parallel MAPA is a term obeying the syntax:

$$q ::= Y(\mathbf{t}) \mid q \parallel_I q \mid \partial_E(q) \mid \tau_H(q) \mid \rho_R(q)$$

Here, $Y \in \text{Prc}$ is a process name, \mathbf{t} is a vector of data expressions, $I, E, H \subseteq \text{Act}$ are sets of actions, and the function $R : \text{Act} \setminus \{\tau\} \rightarrow \text{Act} \setminus \{\tau\}$. A *parallel MAPA specification* is a set of MAPA process equations with an initial process q defined by the above rules.

In an initial MAPA process generated by the rules above, $q_1 \parallel_I q_2$ is *parallel composition* w.r.t. the action set I (defined in Timmer et al. [2012] with $\gamma(a, a) = a, a \in I$) and $\tau_H(q)$ *hides* the actions in H , i.e., turns all the actions in H into τ and removes their parameters. Since we only use these two operators later in our definition, we refer for further information about MA and the MAPA language to Timmer et al. [2012].

4. COMPOSITIONAL SEMANTICS OF eSADF

This section presents an MA semantics for eSADF. The first consideration is to model the channels and the processes in eSADF separately. This approach is simple and the MAs for processes are finite, whereas the MAs for channels are countably infinite as channels are unbounded. However, the drawback is the intermediate state space caused by the communication between the processes and the channels. For example, a data channel must either know the current operating scenario of the process which connects with it through its output port and then notify the process whether the tokens are available in such operating scenario or constantly send its channel status to the process. Hence, we model the control channels (as FIFO buffers) and the data channel (as natural numbers) as part of the process's definition. As the process's behavior merely depends on the token availabilities and/or the contents of these channels, no further status synchronization between other components is needed. After all channels are integrated into the corresponding processes, the MAs for processes take care of the channel's status update. This is done by using the action synchronization between processes.

From now on, we will only consider the MA semantics of a detector, since the MA semantics of a kernel can be easily derived from the detector's semantics:

Remark 1. A kernel is a special kind of detector, in which

- (1) no output channel is a control channel,
- (2) its sub-scenario set is identical with its scenario set, and
- (3) for each scenario σ , $F(\sigma)$ is a DTMC with one state $s \in \mathbb{S}$, $\mathbb{P}(s, s) = 1$ and $\Phi(s) = \sigma$.

Recall that a detector can have more than one (data/control) channel connected through its input ports (input ports and/or control ports) of a single kernel (detector). Moreover, since we integrate the input channels as variables into their process's definition, we only consider these channels and the corresponding ports. For simplicity's sake we assume that the detector, say D , has only one such data channel and control channel from one detector, say D' (see Figure 2). This is easily generalized to several channels.

Detector D has an input data channel $DC_{(D',D)}$ from D' which is connected through D' 's output port $oi_{D',D}$ and D 's input port $i_{D',D}$, and a control channel $CC_{(D',D)}$ from D' which is connected through D' 's output port $oc_{D',D}$ and D 's control port $c_{D',D}$. As mentioned earlier, the channels in eSADF are modeled as variables. Let the variable $dc_{(D',D)} \in \mathbb{N}$ represent the current number of tokens in data channel $DC_{(D',D)}$. For control channels, we let the variable $cc_{(D',D)}$ represent the current status of control channel $CC_{(D',D)}$, where $\Sigma_{(D',D)}$ is the set of values of the scenario token which can be stored in $CC_{(D',D)}$. Hence, we have $cc_{(D',D)} \in \Sigma_{cc_{(D',D)}}^*$. If clear from the context, we omit the subscript (D',D) . Various operators on these variables are defined, for instance, $|cc|$ returns the length of the sequence cc ; $\text{head}(cc)$ returns the first element (head) of the sequence; $\text{tail}(cc)$ is only defined for non-empty channels and returns the remaining string by removing $\text{head}(cc)$, and $cc \ ++ \ \omega$ for $\omega \in \Sigma^*$ denotes concatenation. Note that the general definition for detectors with several data and control channels can be easily derived from this simplified definition, since we can just replace the single variable by a vector of variables.

The MA semantics of a detector D consists of two modules: the (*sub*)*Scenario Module* (SM) and the *Function Module* (FM). First, we give the MA definition of scenario module SM . Since SM only communicates (synchronizes) with FM via requesting and waiting for sub-scenario decisions, no variable is used in SM . Recall that for each scenario $\sigma \in \Sigma$, $(\mathbb{S}, \iota, \mathbb{P}, \Phi)$ is a non-empty finite Markov chain $(\mathbb{S}, \iota, \mathbb{P})$ associated with a function $\Phi : \mathbb{S} \rightarrow \Omega$. Now we define an MA for $(\mathbb{S}, \iota, \mathbb{P}, \Phi)$ for each scenario $\sigma \in \Sigma$, and then compose them in parallel. For scenario σ , we assume that $\mathbb{S} = \{S_0, \dots, S_n\}$, $n \in \mathbb{N}$, and $\iota = S_0$. We also let the set $\text{Post}(S) = \{T \mid \mathbb{P}(S, T) > 0\}$ for $S \in \mathbb{S}$.

Semantics of Detectors. For a better understanding, we distinguish input and output actions: if the synchronization actions are initiated by an MA, these actions are overlined by “—”, and the synchronization actions waiting for the synchronization are denoted as usual. Note that this notation will not affect the original MA semantics.


Definition 10 (Semantics of a Scenario). The semantics of the *scenario module* SM^σ of detector D in scenario σ is defined by:

$$SM(S : \mathbb{S}) := \text{req}(\sigma) \ \sum_{T: \text{Post}(S)} \mathbb{P}(S, T) : SM'(T)$$

$$SM'(S : \mathbb{S}) := \overline{\text{subsc}(\omega)}.SM(S) \ \text{where } \omega = \Phi(S) \in \Omega$$

The task of SM is to simulate the DTMC embedded in scenario σ to return the sub-scenario decision which the detector D is going to execute. Since in the original SADF semantics the sub-scenario [Theelen 2007] is selected by both making a one-step transition in the DTMC and checking the function Φ , we use an intermediate state for every original state in the DTMC. To this end, we let $SM(S)$ represent the behavior of the original state and $SM'(S)$ represent the behavior of the intermediate state of S . First, if SM receives the sub-scenario decision request from function module FM in σ (i.e., action $\text{req}(\sigma)$), and the last left state in the DTMC for σ is S , we make a one-step to the intermediate state of S 's successor states (i.e., the intermediate states of $\text{Post}(S)$) with probabilities determined by \mathbb{P} . The behavior of the intermediate state just returns the sub-scenario decision using $\Phi(S)$, and then behaves like state $SM(S)$.

Example 1. A simple example shows how to translate the DTMC and sub-scenario decision function (left figure below) in eSADF to an MA (right).



We now initialize the scenario module of D in scenario σ by setting $SM^\sigma = SM(S := S_0)$, where $S_0 = \iota^\sigma$ is the initial state of the DMTC of σ .

Definition 11 (Scenario Module). The scenario module of detector D is:

$$SM_D = SM^{\sigma_1} ||| SM^{\sigma_2} ||| \dots ||| SM^{\sigma_n},$$

where $|||$ equals $\|\emptyset$ (as there is no synchronization between the MAs for the scenarios).

Definition 12 (Functional Module). The functional module FM of a detector D is defined by:

$FM(dc : \mathbb{N}, cc : \Sigma^*, subsc : \Omega_D)$

$$:= (|cc| \geq 1 \wedge \text{head}(cc) = \sigma \wedge subsc = \perp) \Rightarrow \overline{\text{req}(\sigma)}. \sum_{\omega \in \Omega} subsc(\omega). FM(dc, cc, subsc := \omega)$$

$$+ \sum_{\omega \in \Omega_D} (subsc = \omega \wedge dc \geq R_D(\omega, i))$$

$$\Rightarrow (\lambda_D^\omega). \overline{\text{exe}_D(\omega)}. FM(dc - R_D(\omega, i), \text{tail}(cc), subsc := \perp)$$

$$+ \sum_{\omega' \in \Omega_{D'}} \text{exe}_{D'}(\omega'). FM(dc + R_{D'}(\omega', oi), cc ++ t_{D'}(\omega', oc), subsc)$$

The tasks of FM are manifold. FM has three parameters: the number of tokens in data channel $dc_{(D', D)}$, the content of control channel $cc_{(D', D)}$, and the current operating sub-scenario $subsc$. One task of FM is to infer both the current scenario of D from the content of the first scenario token in each control channel and whether the sub-scenario is already available (i.e., equal to ω) or undefined (\perp). If the sub-scenario is undefined and the current scenario can be determined ($\text{head}(cc) = \sigma$), FM will synchronize with SM to determine the operating sub-scenario in σ (by action $\overline{\text{req}(\sigma)}$). After SM returns the sub-scenario, say ω , FM writes ω into variable $subsc$. After the sub-scenario is available (i.e., $subsc \neq \perp$), FM can execute as soon as there are enough data tokens in dc , which is checked by inspecting the rate function $R_D(\omega, i)$ for port i . If there are enough tokens, FM can execute, and the execution time is exponentially distributed with a mean duration of $1/\lambda^\omega$. After the execution, FM will synchronize with another process to update the corresponding channel status (i.e., process the tokens to its output channels by action $\overline{\text{exe}_D(\omega)}$ action) and consumes the tokens from the input channels (i.e., updates its own variables' values). The last task of FM is to let other processes update their input channel status (i.e., to produce tokens onto the channels which are the input channels of D) after their executions.

We are now in a position to define the MA semantics of a detector D as the composition of its functional and scenario module.

Definition 13 (Detector Semantics). The semantics of detector D is given by:

$$\mathcal{M}_D := \tau_{I_D}(FM_D ||_{I_D} SM_D),$$




Fig. 3. A sample eSADF graph and its (somewhat simplified) MA.

where $I_D = \{\text{req}(\sigma) \mid \sigma \in \Sigma\} \cup \Omega$ and

$$FM_D = FM(dc := \phi^*(DC), cc := \psi^*(CC), \text{subsc} := \perp),$$

where ϕ^* , ψ^* are the initial content of data channels and control channels in C^* , respectively.

The action set I_D includes the sub-scenario request action req with all scenario values and the sub-scenario return action with all sub-scenario values. As the actions in I_D are only used for synchronizing FM_D and SM_D , all these actions are made unobservable by applying $\tau_{I_D}(\cdot)$.

Semantics of Kernels. The semantics of kernel K is obtained by simplifying the semantics of a detector D_K . First, the set of scenarios Σ_{D_K} and the set of sub-scenarios Ω_{D_K} of D_K are identical to K 's set of scenarios Σ_K (i.e., $\Sigma_{D_K} = \Omega_{D_K} = \Sigma_K$). Similarly, the consumption/production rate function R of D_K equals \bar{R} of K . As K has no control outputs (so does D_K), the function t and rate function R are not defined in D_K . The scenario module SM of D_K in scenario ω is defined as described in Remark 1 by:

$$\begin{aligned} SM(S_0) &:= \text{req}(\omega).SM'(S_0) \\ SM'(S_0) &:= \overline{\text{subsc}(\omega)}.SM(S_0) \end{aligned}$$

Semantics of eSADF Graphs. Finally, we define the MA semantics for an eSADF graph. We assume that an eSADF graph consists of a set of detectors $\{D_1, \dots, D_n\}$, $n \in \mathbb{N}$. For each detector D_i ($1 \leq i \leq n$) let MA \mathcal{M}_{D_i} be its semantics and Act_{D_i} the set of interactive actions in D_i .

Definition 14 (eSADF Semantics). The MA for eSADF graph $\mathcal{G} = (\mathcal{P}, C^*)$ is:

$$\mathcal{M}_{\mathcal{G}} := \tau_H(\mathcal{M}_{D_1} \parallel_{I_1} \dots \parallel_{I_{n-1}} \mathcal{M}_{D_n}),$$

where \parallel is left-associative, $I_i = Act_{D_{i+1}} \cap (Act_{D_1} \cup \dots \cup Act_{D_i})$ for $1 \leq i \leq n$, and $H = Act_{D_1} \cup \dots \cup Act_{D_n}$.

Example 2. The eSADF graph in Figure 3 (left) consists of detector A and kernel B , control channel $CC_{(A,B)}$, and data channels $DC_{(B,B)}$ and $DC_{(B,A)}$. Production and consumption rates equal to 1 are omitted, and the red numbered points indicate the number of initial tokens in these channels (control channels are initially empty). Kernel B can execute in scenarios I and P . The execution time of I is exponentially distributed with mean duration one; P has mean duration $\frac{1}{2}$. The scenario occurrence is decided by A based on the embedded DTMC (cf. Example 1) and sent to B via the scenario tokens

valued with I and P through channel $CC_{(A,B)}$. Since there is no input control channel for A , A always executes in a default scenario ϵ . Here, we assume the sub-scenario decision procedure in A will be done immediately.

5. QUANTITATIVE ANALYSIS

Now that we have seen that the behaviors of an eSADF graph and its actors can be adequately described by means of a closed MA, we address how quantitative measures on eSADF graphs, e.g., expected time and long-run objectives, as well as probabilities of certain events happening unto a certain deadline, can be obtained. A detailed treatment of the algorithms is outside the scope of this article; a full explanation can be found in Guck et al. [2014a].

5.1. Analyzing Markov Automata

Given a closed finite-state MA, we consider three quantitative objectives: expected time, long-run average, and timed (interval) reachability. *Expected time* objectives focus on determining the minimal (or maximal) expected time to reach a given set of states. *Long-run* objectives determine the fraction of time to be in a set of states when considering an infinite time horizon. *Timed reachability* objectives are focused on the probability to reach a set of states within a given time interval. As MA exhibit non-determinism, we focus on maximal and minimal values for all three objectives. These correspond to the best and worst possible resolution of the non-determinism present in the MA. Expected-time and long-run average objectives can be efficiently reduced to well-known problems on Markov decision processes such as stochastic shortest path, maximal end-component decomposition, and long-run ratio objectives. As shown in Guck et al. [2014a], the reduction to these well-investigated problems enables efficient analysis techniques such as linear programming, value iteration, and maximal end-component decomposition, which all have a polynomial time-complexity in the size of the MA. This all relies on the fact that for optimal expected time or long-run objectives, the resolution of the non-determinism by extremely simple (so-called memoryless) policies suffices. Timed reachability objectives, however, are harder to obtain. The main technical complication here is that the optimal way of resolving the non-determinism requires policies with infinite—even uncountably large—memory. Intuitively, this can be seen as follows. Assume that in an MA there is a non-deterministic choice between two options: either reaching a target state slowly, but almost surely (i.e., with probability one), or reaching a target soon, but with the risk that the target is not reached at all (so the target is reached with a probability strictly smaller than one). Then it makes a difference whether ample time remains to reach the target, or whether there is almost no time left. In the latter case, the fast but unsafe option is optimal, whereas in the first case, the slow and safe option is optimal. The way around this is to resort to *discretization*. Here, the time interval is split into equally-sized discretization steps, each of length δ . The discretization step is assumed to be small enough such that with high probability it carries at most one Markovian transition. This yields a discretized MA (dMA), a variant of a semi-MDP, obtained by summarizing the behavior of the MA at equidistant time points. The analysis of the dMA yields an approximation of the true timed reachability probability in the MA where the error is bounded (depending on the discretization step δ and the largest rate occurring in the MA). As we will see in the case studies (Section 6), the analysis of timed reachability objectives, therefore, is time-consuming.

5.2. Confluence Reduction and Non-Determinism

Prior to analysing an MA, we employ a state-space reduction. This is an on-the-fly technique, meaning that it can be used while generating the MA from a given eSADF




Fig. 4. (a) a sample GSPN, (b) its MA semantics, and (c) its reduced MA.

graph. In addition, it is a symbolic technique; that is, it is directly applicable on the MAPA terms that describe the behavior of an eSADF graph. Besides being an effective and efficient state-space reduction technique, confluence reduction—in theory—can reduce *all* non-determinism in the MA semantics of an eSADF graph.

Confluence Reduction. Confluence reduction [Timmer et al. 2013] reduces the state space based on commutativity of transitions, removing non-deterministic transitions caused by parallel composition of independent components. It is similar in spirit as partial-order reduction. The reduction preserves the three quantitative metrics of interest described above. Based on heuristics to detect confluence in MAPA terms, the state space is reduced in an *on-the-fly* manner. Its effect is illustrated in Figure 4, which gives the MA semantics (b) for a simple stochastic Petri net with immediate transitions (solid bars) and timed transitions (open bars) (a) and afterwards reduces the state space (c) by applying confluence reduction. The key observation is that the commutativity of the immediate transitions t_1 and t_4 , and t_2 and t_4 is exploited in the reduction. Confluence reduction yields a reduction of seven to four states.

Confluence Reduction in a Nutshell. The basic idea of confluence reduction is to determine the confluent sets of transitions. To obtain these, groups consisting of only confluent interactive probabilistic transitions should satisfy the following conditions: (1) all transitions are τ -transitions with Dirac distribution, (2) all transitions enabled prior to a transition in this group are still enabled after taking such transition. The diagram right above illustrates the latter constraint. If transition $s \xrightarrow{\tau} t$ is in a group, say T , and if

$$\begin{array}{ccc}
 s & \xrightarrow{\tau} & t \\
 a \downarrow & & \downarrow a \\
 \mu & \equiv_R & \nu
 \end{array}$$

$s \xrightarrow{a} \mu$, then $t \xrightarrow{a} \nu$ must exist such that μ and ν are related, i.e., all states in the support of μ and ν are connected by transitions from T . Timmer et al. proved that the transitions satisfying the conditions above connect *divergence-sensitive branching bisimilar* states [Timmer et al. 2013]. Hence, it is safe to prioritize confluent transitions. As the intermediate states on a confluent path are bisimilar, they can be aggregated. Confluence reduction is applied on syntactic MAPA terms in an *on-the-fly* manner, thus, avoiding a full state-space generation prior to the reduction (as opposed to bisimulation reduction [Theelen et al. 2012], which requires the construction of a full state space prior to reduction). Case studies show a state-space reduction from 26% to 83% [Timmer et al. 2013]; for the MPEG-4 decoder and the face recognition case study, this is about 66% (cf. Table I).

Non-Determinism in eSADF. Non-determinism in our MA semantics only arises from the execution of independent concurrent actors in eSADF graphs:

Table I. MA Size for Sample eSADF Graphs

		before red.	with conf. red
Example 2	#states	19	7
	#transitions	19	7
MPEG-4	#states	61918	20992
	#transitions	81847	40910
	#non-det. state	4	1
Face recognition	#states	106784	29440
	#transitions	154688	77344

THEOREM 1. *Non-determinism only occurs between sub-scenario decision actions from different, independent processes. All these transitions are confluent, i.e., they yield the same (Markovian) state. The probability distribution to reach such states is independent from the resolution of the non-determinism.*

The proof sketch of this result is provided in Katoen and Wu [2014] and confirms a similar result in Theelen et al. [2006] for an alternative SADF semantics. The key point is that, since the enabled probabilistic choice among sub-scenarios is instantaneous, the transitions representing the time progress can not be enabled before all such enabled probabilistic transitions⁴. Since these enabled probabilistic transitions are independent, they are confluent to the state where the Markovian transitions are enabled. Whereas in SADF [Theelen et al. 2006], timed transitions are deterministic (since it always takes the earliest finished execution time of a process), in eSADF they are probabilistic and resolved by the race condition. Thanks to the above result, confluence reduction can potentially reduce all non-determinism from the MA semantics of an eSADF graph. As heuristics are used, this is not always established in practice, but then the above result guarantees that worst and best case quantities coincide.

6. CASE STUDIES

In this section, we provide two case studies: an eSADF of the MPEG-4 decoder benchmark and an eSDF⁵ of a real-life face recognition application (that we obtained from an industrial partner). Different system metrics, such as throughput of different kernels (actors), channel buffer occupancies, and response delay are obtained from the quantitative analysis of the resulting MA.

6.1. The MPEG-4 Decoder Benchmark

We consider the eSADF graph of an MPEG-4 decoder, as given in Figure 1. Applying confluence reduction to the MA of the MPEG-4 decoder reduces the state space by about a factor of 3:

Note that non-determinism for the MPEG-4 decoder is not completely eliminated by confluence reduction, as heuristics are used to detect confluent transitions.

Our MA semantics allows for determining several performance metrics for the MPEG-4 decoder in a fully automated manner. We illustrate this for various quantitative measures of the MPEG-4 decoder example. We will also provide the verification time for all the properties in Table IV. For the sake of convenience, the properties are labeled (numbered) p1 through p11, as indicated at the top of the corresponding Tables II and III or throughout the text.

⁴In SADF, this is ensured by action-urgency; here, by maximal progress.

⁵Since SDF is a sub-class of SADF, our eSADF semantics is readily applicable to eSDF.

Table II. Average Buffer Occupancy in Each MPEG-4 Channel

	p1	p2	p3	p4	p5
av. #tokens	idct-rc	vld-mc	vld-idct	mc-rc	rc-mc
Our approach	38.2584	27.7168	1.3118	0.2071	0.7929
SDF ³	42.2215	24.2215	1.18198	0.4412	0.8574

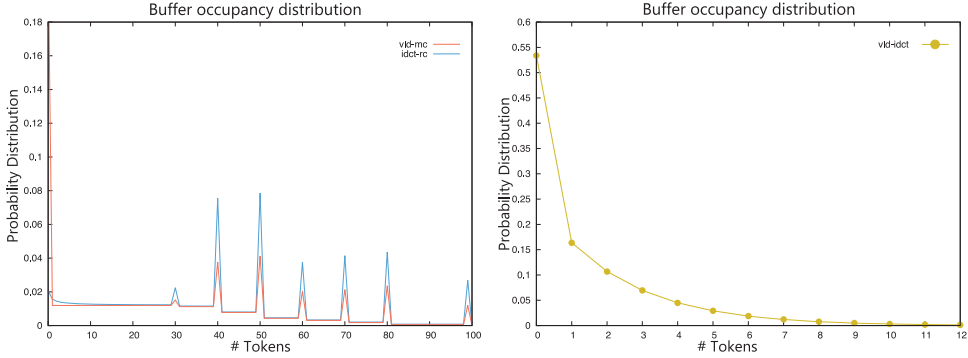


Fig. 5. Token distribution in three MPEG-4 channels.

Buffer Occupancy. We consider two long-run properties: *the average number of tokens and the probability distribution of tokens in a channel*. The average buffer occupancy of each channel of the MPEG-4 decoder is shown in Table II. This includes (second row) a comparison with the results obtained for the MPEG-4 decoder with the SDF³ tool, using deterministic execution times. The first row indicates the buffer occupancies obtained with the approach in this article. From the results, we observe that the channels of IDCT-RC and VLD-MC have a much higher average occupancy than the other three. This is due to the fact that, in case of the I frame and P_x frames, IDCT and VLD need to execute only one time, whereas RC and MC need to compute 99 and x times, respectively. Hence, the tokens can accumulate in both channels waiting for processing. The average number of tokens in channels MC-RC and RC-MC together is one. This is due to the cyclic channels between MC and RC, which guarantees that MC and RC execute at the same rate.

The token distributions of three data channels are shown in Figure 5. The peak values for channels VLD-MC and IDCT-RC are caused by conservatively approximating the motion vectors (cf. Section 6.1) by fixed numbers (i.e., $x \in \{30, 40, 50, 60, 70, 80, 99\}$). Further, the average number of tokens in channel VLD-IDCT is 0.57, which is possibly due to that VLD will produce at most one token to the channel VLD-IDCT at each time. Thus, the probability of having more than three tokens in channel VLD-IDCT is very low; whereas, the probability of VLD-IDCT being empty is quite high (≥ 0.65).

Expected Time. The second property considered is *the expected time and timed reachability probability for a channel to reach its 50% ($p6$) and 90% capacity*, respectively. In our evaluation, we let the maximal number of tokens in the channel to be the capacity of that channel. Afterwards, we mark such states as target states, where the current number of tokens in the channel is more than 50% and 90% of its capacity, respectively. Here, we only show the result in 50% case of channel VLD-MC in Figure 6.

Response Delay. We exploit timed reachability objectives for *response time estimation*. This allows us to answer questions such as “what is the expected time until a process finishes its first execution?” or “what is the probability of a process responses for the

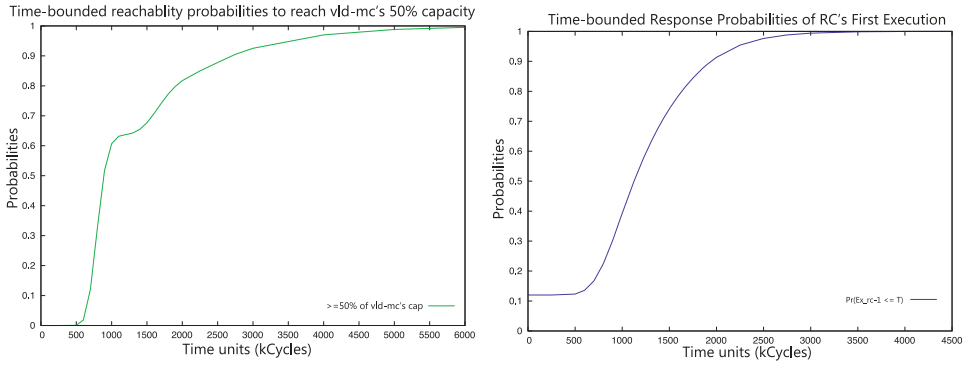


Fig. 6. Timed reachability probabilities to reach VLD-MC's 50% capacity (left)/first response probabilities of RC within time t (right).

Table III. Throughput of Each Kernel in MPEG-4 Decoder

	p8	p9	p10	p11
Throughput	IDCT	VLD	MC	RC
Our approach	0.0423732	0.0423732	0.000746128	0.000746128
SDF ³	0.0437919	0.0437919	0.000745268	0.000745268

first time within time t ?”. Computing response time estimates boils down to computing the expected time or timed reachability probabilities from the initial state to those states where the process has just finished its first execution. Taking RC for example, we get 1,152.617 (kCycle) as the answer to the first question and Figure 6 to the second (p7). As a P_0 frame (which means a still video frame) occurs with probability 0.12 and RC just copies it from MC, the probability of RC finishing its first execution within time 0 is 0.12.

Throughput and Inter-Firing Latency. We compute the *throughput* of a kernel by the following approach. First, we compute the long-run average probability (P_σ) of a kernel executing in scenario $\sigma \in \Sigma$. This can be done by adding a Boolean variable to the kernel's MAPA definition. It is set to true when the execution condition is satisfied and set to false when the execution finishes. Since the expected execution time (E_σ) of a kernel in scenario σ is known, the throughput of this kernel is computed as the sum of the long-run average probability P_σ divided by the expected time E_σ for each scenario σ :

$$Tr = \sum_{\sigma \in \Sigma} \left(\lim_{t \rightarrow \infty} \frac{P_\sigma \cdot t}{E_\sigma} \cdot \frac{1}{t} \right) = \sum_{\sigma \in \Sigma} \frac{P_\sigma}{E_\sigma} = \sum_{\sigma \in \Sigma} \lambda_\sigma \cdot P_\sigma.$$

The results for the MPEG-4 decoder are shown in Table III. The throughput of IDCT and VLD, as well as of MC and RC coincide.

Analogously, the average inter-firing delay (In) can be computed as:

$$In = \sum_{\sigma \in \Sigma} \left(\frac{(1 - P_\sigma)}{\sum_{\sigma \in \Sigma} P_\sigma \cdot \lambda_\sigma} \cdot \frac{P_\sigma}{\sum_{\sigma \in \Sigma} P_\sigma} \right).$$

We take MC as an example and compute the In of MC as 1,460.5 (1,341.8 in SDF³) kCycles.

Table IV. Verification Time of Properties (M = Minutes, S = Seconds)⁷

p1	p2	p3	p4	p5	p6
489m (6s)	592m (3.4s)	141m (11.4s)	11m (2.7s)	10.4m (12.7s)	14.7h (n.a)
p7	p8	p9	p10	p11	
456m (n.a)	24.87m (4.76s)	18.48m (6.56s)	4.45m (13.02s)	5.1m (3.2s)	




Fig. 7. The SDF model for face recognition application.

The verification time to check the various properties is indicated in Table IV. The numbers in brackets are the run times of SDF³ (where n.a. stands for not applicable). All experiments have been conducted on a machine with a 48-core CPU at 2.1GHz and 192GB memory.

6.2. Face Recognition

The Face Recognition Example. In this case study, we apply our approach to an SDF model (see Figure 7) of a real-life application of face recognition. This case study is a simplified anonymized case study from the company Recore Systems. The SDF models the processing of a single frame in a two-dimensional picture. The SDF consists of 25 kernels. For the sake of readability, production/consumption rates equal to 1 are omitted; furthermore, the initial number of tokens in a channel is marked by a black dot attached to a number. Note that an SDF graph can be considered as an SADF graph without detector and every kernel executes in a default scenario (no sub-scenarios and no probabilistic selection of such sub-scenarios). We can, thus, apply our approach without any problems to SDF graphs.

There is, however, one difference: SADF does not allow auto-concurrency, whereas SDF does. In Figure 7, the actors in which auto-concurrency is disallowed (hence, we omitted the self-loop with rate one in such actors for simplicity) are marked as blue, while the actors (i.e., actor D, H, and X) in which auto-concurrency is allowed are marked red. Auto-concurrency amounts to the simultaneous firing of an actor (provided the firing condition is multiply enabled). Auto-concurrency can be expressed in our MA semantics as the parallel composition of multiple *enabled* copies of the actor process, due to the fact that a process with multiple (say n) Markovian delays λ (representing

⁷These verification times seem prohibitive, but exploit algorithms that allow for analyzing MA with non-determinism.




Fig. 8. Treating auto-concurrency in the MA semantics of SDF.

Table V. Throughput of Each Actor in the Face Recognition Application

Throughput	A-C/F/G/I-W/Y	D & E	H	X
Auto-Concurrency	0.1226	18.390	7.8469	3.187
No-Auto-Concurrency	0.08805	13.20900	5.6361	2.289

the firings) enabled in parallel is equivalent—in the sense of branching bisimulation—to a process which evolves with rate $n \cdot \lambda$ to its successor, and then evolves subsequently with rate $(n-1) \cdot \lambda$ and so forth (see Figure 8). Thus, the execution rate in the original MA definition of such actor with auto-concurrency becomes multiple times (= number of enabled copies of the actor) the original rate. This allows us to quantitatively assess the impact of auto-concurrency on the face recognition example in Figure 7. In our experiments, we evaluate the SDF without auto-concurrency for actors D, H, and X, and provide a quantitative comparison between the two versions. The state space of the MA is about 1,000,000 states and can be reduced by confluence reduction by a factor of 66% (see Table I). The computation time of the throughput of each actor varies from 3 to 10 minutes, and the computation of the token distribution for each case takes about 1 to 2 minutes. As indicated before, the most time-consuming property to compute is the probability of reaching a (critical) situation within a given time-bound, which varies from minutes to a few hours, depending on the given time-bounds.

Throughput. We first compute the throughput of each actor in the SDF graph. We observe that all the actors, except the “auto-concurrent” actors D, H, X and E, have the same throughput. This throughput is the base throughput of the SDF graph. It follows that the throughput of D, E is 150 times, H 64 times, and X 26 times the base throughput, respectively. On the other hand, the SDF’s throughput is increased by a factor of 1.5 when auto-concurrency is exploited. Thus, auto-concurrency increases the throughput by about 50%. The results are listed in Table V. Note that since in SDF the actors have only one default scenario, the inter-firing latency of an actor is simply the reciprocal of its throughput.

Buffer Occupancy. As for the MPEG-4 decoder, we determine the average number of tokens and the distribution of tokens in different representative channels. The outcomes are shown in Table VI and Figure 9.

Observe that channel c-d and d-f have almost the same average number of tokens, and the probability of having one token in channel d-e is much higher than for e-d (one token is either in e-d or d-e). This is due to the heavy workload at actor D. Moreover, we observe that without auto-concurrency, more tokens will accumulate in some channels in average, but this does not always occur. For some channels, such as

Table VI. Average Tokens and Token Distribution in Different Channels

Average Buffer Occupancy	c-d	d-f	d-e	g-h	h-i	v-x	u-x
Auto-Concurrency	82.16925	82.05888	0.09196	62.35874	63.57629	0.06694	25.57622
No-Auto-Concurrency	101.28171	101.20245	0.06605	54.37433	54.10945	0.64911	25.23185

c-d/d-f

Token distribution	0	1-149	150
Auto-Concurrency	0.00466	0.00601	0.10025
No-Auto-Concurrency	0.00335	0.00431	0.35377

d-e

Token distribution	0	1
Auto-Concurrency	0.908041	0.09196
No-Auto-Concurrency	0.93395	0.066048




Fig. 9. Token distribution in channels v-x & h-i when using auto-concurrency.

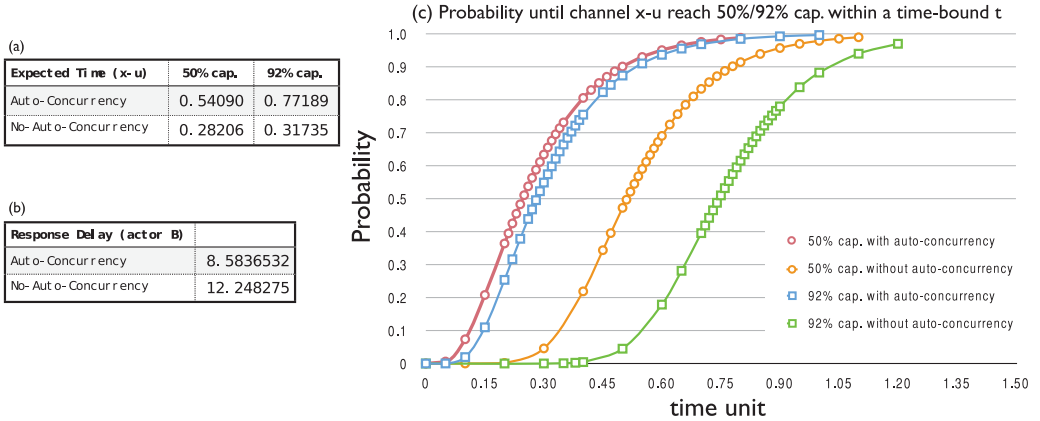


Fig. 10. The expected time (a), timed reachability to critical situation (b), and response delay of actor B (c).

u-x, there is almost no impact on the average number of tokens. The distribution of different number of tokens in channel c-d, d-f, d-e, v-x, and d-e is also given. We see that the probability distribution is monotonically decreasing or increasing, and the large probability (such as 0.99 in v-x and 0.969 in h-i) is at marginal values (0 or 64).

Expected Time. We now consider the expected time and time-bounded reachability until certain critical situations occur, such as 50% or 90% channel's capacity. Figure 10(c) indicates that, in absence of auto-concurrency, there is a large difference between the probability of reaching 50% and 90% within a given time-bound. If auto-concurrency is used, the curves of reaching 50% and 92% channel capacity are close, which means that if a channel reaches 50% capacity, it will reach 92% of its capacity

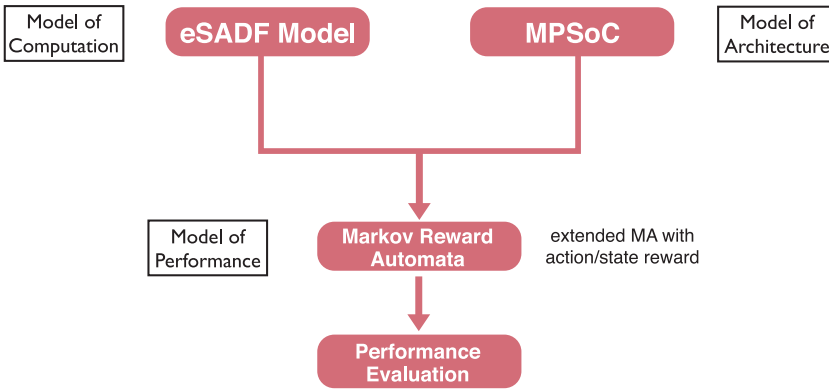


Fig. 11. An overview of our approach extended to hardware platform.

soon thereafter. If no auto-concurrency is used, the period between these two situations becomes significantly larger.

Response Delay. In the property of response delay, we consider the first response until the actor B, which is the most significant actor in the SDF, since it is the “last” actor in the SDF to finish one cycle of all the tasks of actors. From Figure 10(b), we can conclude that auto-concurrency improves the response delay by about 30%.


7. DEPLOYING eSADF GRAPHS ON HARDWARE PLATFORMS

In this section, we consider deploying eSADF graphs on hardware platforms. Adopting eSADF as a Model of Computation (MoC) as before, we extend our framework to two additional models: the Model of Architecture (MoA) and the Model of Performance (MoP). Our approach is illustrated in Figure 11 and details are provided below.

7.1. The Hardware Platform (MoA)

Since eSADF abstracts the software applications as MoC, the MoA is used to model the hardware platform. There are various approaches [Castrillon and Leupers 2013] that can be used. For instance, if the communication overhead in the system needs to be taken into account, we can adopt the approach of CA-MPSoC, introduced in Shabbir et al. [2010] as the MoA. Since we do not consider the communication overhead in our model, the MoA can be seen as an MPSoC with individual cores. To investigate the energy consumption, we also need to define an energy model in the MoA.

The Energy Model of MoA. In our model, we consider modern processing elements (PEs) equipped with DVFS and/or DPM to reduce both dynamic and static power consumption. DVFS reduces the energy consumption of a PE by dynamically scaling its voltage and its frequency. We consider in our model *inter-task DVFS*, which allows the frequency to be changed only after a task is finished on this PE. We do not take the intra-task DVFS into account to keep things simple. DPM reduces the energy consumption by allowing the PE to get into sleep mode after a certain idle period. The sleep mode has a much lower static power consumption than the idle mode. Note that the *break-even* time is the key factor in DPM, since only the idle period is longer than such break-even time, the system will get energy benefit. An example of our energy model with only DVFS based on the Samsung Exynos 4210 processor can be found in Figure 13, which is used in our case study later.

Fig. 12. The γ function in MRA.

7.2. Markov Reward Automata (MoP)

In order to model energy consumption in a quantitative manner, we consider an extension of MA with *rewards*, known as Markov Reward Automata (MRA) [Guck et al. 2014b]. MRA equip states and transitions with rewards. Transition rewards are instantaneous and incurred on taking the transition. The reward earned in state s when residing t time units in that state equals the state reward of s multiplied by t . Briefly speaking, an MRA is an extended labeled transition system:

- equipped with both continuous time stochastic and non-deterministic transitions,
- for each state, a cost/gain reward is assigned, which indicates the cost/gain per time unit, if the system stays in that state,
- for each transition, a transition reward is assigned in order to indicate the cost/gain if this transition is taken.

As we only consider energy consumption (and not energy recovery), rewards are non-negative.

Definition 15 (MRA). An MRA is a tuple $(S, s_0, Act, \longrightarrow, \Rightarrow, \rho, \eta)$, where $(S, s_0, Act, \longrightarrow, \Rightarrow)$ is an MA, and

- $\rho : S \rightarrow \mathbb{R}_{\geq 0}$ is the *state-reward function*,
- $\eta : S \times Act \cup \{\chi\} \times Distr(S) \rightarrow \mathbb{R}_{\geq 0}$ is the *transition-reward function*.

Note that, since we also need to define the transition reward on outgoing Markovian transitions, a “equivalent” rewriting function γ is defined in order to translate the outgoing Markovian rates of a state to be an action followed by a distribution (similar to interactive probabilistic transition). The γ function takes a Markovian state and attaches it with an action $\chi(E(s))$ together with the exit rate $E(s)$ in it and then followed by a distribution which is identical to \mathbb{P}_s . An example of rewriting function γ over the outgoing Markovian transitions of a state is shown in Figure 12.

7.3. Deploying the MPEG-4 Decoder

To exhibit our extended approach, we conduct a case study on the MPEG-4 decoder introduced earlier together with a hardware model of MPSoC, based on the Samsung Exynos 4210 processor. Since we are mostly interested in energy consumption, we make some simplifications in our model: (1) when modeling DVFS, we assume two voltage-frequency levels, i.e., the lowest voltage of 1V with 1032.7MHz and the highest voltage of 1.2V with 1400MHz, and (2) the DPM mechanism is omitted, since Samsung Exynos 4210 does not support this feature. In our platform, we assume there are only two processors, and each processor can run with a different frequency independently. The energy model is shown in Figure 13, and the energy consumption of the hardware platform of Samsung Exynos 4210 is computed by utilizing the following formula given in Park et al. [2013] as:

$$P_{\text{cpu}} = 0.446 v_{\text{proc}} \cdot V_{\text{cpu}}^2 \cdot f_{\text{cpu}} + 0.1793 V_{\text{cpu}} - 0.1527,$$

where v_{proc} is the total utilization with $v_{\text{proc}} = 0.2$ in the idle state, and $v_{\text{proc}} = 1$ in the running state.

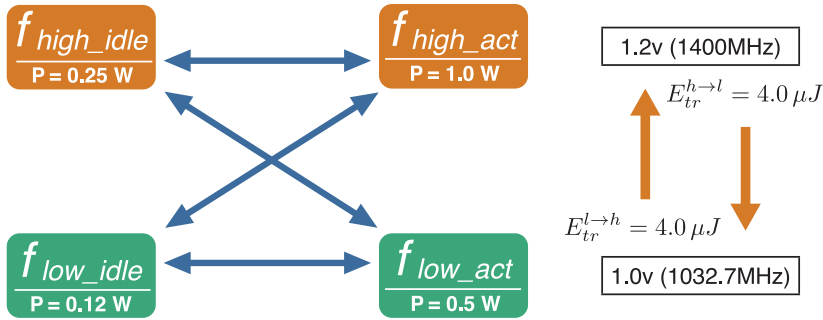


Fig. 13. A simplified energy model of a PE based on the Samsung Exynos 4210 processor.

Table VII. Energy Consumption of Deploying the MPEG-4 Decoder on the Samsung Exynos 4210 Processor

#states	#transitions	Power cons.	Thr. IDCT	Thr. RC	Exp. energy 1 iter.
215851	314609	Min 0.746337 (W)	21.975093	0.415335	1.581342 (mJ)
		Max 1.528917 (W)	30.279699	0.563065	2.423743 (mJ)

Each processor's behavior is modeled by an MRA, in which the energy consumption is modeled by a state reward, and the energy consumed by downscaling/upscaling in DVFS is modeled by the transition reward. The deployment of an eSADF process of the MPEG-4 decoder to the processors is modeled by using synchronization (i.e., parallel composition). Note that we assume a fully dynamic scheduling rather than a static one. So as to obtain manageable state spaces, we restrict the scenarios in the MPEG-4 decoder to I , P_0 , and P_{50} -frames.

The experimental results are summarized in Table VII. As before, we apply confluence reduction. First, we evaluate the maximal (and minimal) power consumption by computing the maximal (and minimal) long-run reward in the resulting MA. Then, we compute the maximal and minimal throughput of actors IDCT and RC (note that VLD has the same throughput as IDCT, and MC has the same throughput as RC). The throughput here is the number of firing times in one ms of such actor. From the result that the difference between the minimal and maximal throughputs of RC is much smaller than the difference between the maximal and minimal energy consumption, there should exist a balanced scheduler which can keep the energy consumption reasonably low, while the throughput is still acceptable (this is due to the fact that doubling the power consumption does not cause a doubling of the throughput, therefore relaxing the requirement on the throughput may cause a more significant reduction of power consumption, w.r.t. the throughput, and archive more energy efficiency). The last property we compute is the energy consumed when the MPEG-4 decoder finishes its first iteration. This is computed by the expected cumulative reward from the initial state to the goal states. We observe that substantial energy savings (up to 50%) can be achieved by a dedicated deployment and DVFS.

8. RELATED WORK

SADF Semantics. Whereas we consider exponentially timed SADF—akin to exponentially timed SDF [Sriram and Bhattacharyya 2009]—and use MA as semantic models, the original works on SADF focus on a *real-time* semantics using so-called Timed Probabilistic Systems (TPS) [Theelen et al. 2008; Theelen 2007; Theelen et al. 2006]. TPS have deterministic delays and discrete probabilistic branching. A direct comparison of

analysis results is, thus, not possible. The compositional nature of our semantics, together with the memoryless property of exponential distributions, yields a simple and lean semantics. In contrast, the TPS semantics has to account for actors that are enabled at the same time; this occurs in our framework with probability zero. In addition, the probability measure of Zeno paths for MA obtained from eSADF graphs is zero. This avoids a special (and typically intricate) treatment, so as to exclude Zeno paths from the analysis. The simplicity of our semantics allows for considering kernels as simplified detectors, and providing a relatively straightforward formal proof (sketch) of the absence of non-determinism—confirming the result in Theelen et al. [2006] for their TPS semantics. Finally, confluence reduction allows for an on-the-fly state-space reduction, which (to the best of our knowledge) does not exist for the TPS semantics. As time in deterministically timed systems has a global synchronizing character, efforts to apply partial-order reduction to timed systems have not been very successful so far.

Model Checking SADF. Earlier work [Theelen et al. 2012] exploited the Construction and Analysis of Distributed Processes (CADP) tool-set [Garavel et al. 2013] for model checking eSADF. There are various benefits and differences with the approach in this article. First, we provide a *full formal* definition of the eSADF semantics. Secondly, the operational model in Theelen et al. [2012] is better suited for SDF than for SADF. In particular, it does not natively support probabilistic choices as needed for random sub-scenario selection in SADF. Using MA, there is no need for awkward—and incomplete—transformations [Rettelbach 1995] to delete probabilistic branching, as applied in Theelen et al. [2012]. This results in smaller models. In addition, using MA, a much richer palette of quantitative measures, can be supported, whereas CADP only supports transient and steady-state measures. In fact, the absence of non-determinism allows for a full-fledged model checking of stochastic versions of computation tree logic (CTL). Finally, confluence reduction is an *on-the-fly* technique, whereas bisimulation reduction (as applied in Theelen et al. [2012]) is not. As shown in the following table

	no red.	with red.	red. factor
[Theelen et al. 2012]	121,430	20,664	5.88
Our work	47,266	16,042	2.95

the use of MA yields smaller models (without reduction), whereas confluence reduction outperforms branching bisimulation used in Theelen et al. [2012] while preserving the quantitative measures addressed in this article.

Energy Analysis by Model Checking. Several models and model-checking approaches have been extended towards the treatment of costs, or dually: rewards. Prominent examples are priced timed automata and Markov reward chains. Like in MA, states are equipped with a reward that grows linearly depending on the state residence time. Markov reward chains are extensions of DTMCs in which a reward is earned on visiting a state (no dependency on state residence times). Norman et al. [2005] applied this model class so as to quantify the impact of various DPM schemes. This has been done using the PRISM model checker. This article applies a similar analysis on *continuous-time* probabilistic models that *include non-determinism*. Our semantics and analysis algorithms allow similar analyzes for all eSADF models. Recently, Ahmad et al. [2015] provided a mapping from SADF (together with an execution platform) onto timed automata, and using the timed automata model checker UPPAAL, the authors showed how to determine the schedulability of SADF graphs. An extension of this approach with prices enables the verification of timed energy usage. Schedulability analysis is complementary to the techniques of this article; our approach does not

support schedulability analysis for hard, real-time deadlines. Instead, the techniques presented in this article can be used to obtain deployments of agents onto execution platforms that optimize *soft* deadlines. This includes, e.g., obtaining the employment such that the maximum expected time to finish all activities is minimized.

Beyond Exponential Distributions. This article considers SADF graphs in which all actor execution times are exponentially distributed. Evidently, this is an abstraction from reality. As argued in the introduction, exponential distributions are convenient if only the mean of an uncertain event is known. Stated more precisely, exponential distributions are the best possible statistical abstraction if the only information about the actor execution times is their expected value. (In case more information is known, such as upper and lower bounds, then uniform distributions are statistically more neutral.) Exponential distributions are mathematically elegant and enable the usage of numerical analysis techniques, as exploited in this article. More general distributions (such as Weibull, gamma, normal distributions, etc.) can be treated by resorting to *simulative* techniques, such as Monte Carlo simulation. Simulation has the advantage of treating such distributions without further ado, and by an on-the-fly analysis, i.e., without the need to generate the entire state space prior to the analysis. Simulation provides statistical estimates of the measure-of-interest, such as buffer occupancy or expected completion time of all SADF actors. These outcomes typically come with confidence intervals for a given allowed inaccuracy. In contrast, the numerical analysis techniques yield exact results (up to the allowed inaccuracy). Younes et al. [2006] provide an extensive comparison between simulation and probabilistic model checking techniques, both from a conceptual and an empirical perspective. An alternative to simulation is to *approximate* the general distributions by (e.g., acyclic) Continuous-time Markov Chains (CTMCs). There are automated algorithms that, given a distribution, yield a minimal representation in terms of a CTMC [Pulungan and Hermanns 2015]. The resulting MA is amenable to the numerical analysis algorithms as advocated in this article, but at the expense of a state-space increase. The growth of the state space strongly depends on the distribution at hand and the required precision; e.g., smooth distributions typically yield reasonably small CTMCs; deterministic distributions (as used to model, e.g., hard deadlines and time-outs) can be exponential in size for small precisions.

9. CONCLUSION AND FUTURE WORK

We presented a compositional semantics of eSADF, SADF in which all executions take exponential time. The semantics is provided in terms of MA, a formal model that naturally fits all the ingredients of eSADF. Two case studies have been provided that illustrate achievable state-space reductions using confluence reduction and obtaining quantitative assertions about eSADF graphs in a fully automated manner. The incorporation of the execution platform into our framework is shown to enable energy analysis.

Future work consists of considering more realistic execution platforms, the comparison of different deployment strategies of SADF actors, and the use of parametric verification to synthesize maximal (or minimal) execution times from high-level specifications.

ACKNOWLEDGMENTS

We thank Recore Systems for providing us with the face recognition case study and, in particular, Kim Sunesen and Timon ter Braak for fruitful discussions about the case study and its analysis.

REFERENCES

- Waheed Ahmad, Philip K. F. Hölzenspies, Mariëlle Stoelinga, and Jaco van de Pol. 2015. Green computing: Power optimisation of VFI-based real-time multiprocessor dataflow applications. In *Euromicro Conf. on Digital System Design (DSD)*. IEEE Computer Society, 271–275.
- Shuvra S. Bhattacharyya, Ed F. Deprettere, and Bart D. Theelen. 2013. Dynamic dataflow graphs. In *Handbook of Signal Processing Systems*. Springer, 905–944.
- Hichem Boudali, A. P. Nijmeijer, and Mariëlle Stoelinga. 2009. DFTSim: A simulation tool for extended dynamic fault trees. In *SpringSim*. SCS/ACM.
- Marco Bozzano, Alessandro Cimatti, Joost-Pieter Katoen, Viet Yen Nguyen, Thomas Noll, and Marco Roveri. 2011. Safety, dependability and performance analysis of extended AADL models. *Computer Journal* 54, 5 (2011), 754–775.
- J. T. Buck and Edward A. Lee. 1993. Scheduling dynamic dataflow graphs with bounded memory using the token flow model. In *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, Vol. 1. 429–432.
- Jeronimo Castrillon and Rainer Leupers. 2013. *Programming Heterogeneous MPSoCs: Tool Flows to Close the Software Productivity*. Springer.
- Yuxin Deng and Matthew Hennessy. 2013. On the semantics of Markov automata. *Information and Computation* 222 (2013), 139–168.
- Christian Eisentraut, Holger Hermanns, Joost-Pieter Katoen, and Lijun Zhang. 2013. A semantics for every GSPN. In *Petri Nets (LNCS)*, Vol. 7927. Springer, 90–109.
- Christian Eisentraut, Holger Hermanns, and Lijun Zhang. 2010. On probabilistic automata in continuous time. In *IEEE Symp. on Logic in Computer Science (LICS)*. IEEE, 342–351.
- Johan Eker, Jörn Janneck, Edward Lee, Jie Liu, Xiaojun Liu, Jozsef Ludvig, Stephen Neuendorffer, Sonia Sachs, and Yuhong Xiong. 2003. Taming heterogeneity - The ptolemy approach. *Proceedings of the IEEE* 91, 1 (2003), 127–144.
- Hubert Garavel, Frédéric Lang, Radu Mateescu, and Wendelin Serwe. 2013. CADP 2011: A toolbox for the construction and analysis of distributed processes. *STTT* 15, 2 (2013), 89–107.
- Marc Geilen and Twan Basten. 2004. Reactive process networks. In *EMSOFT*. ACM, 137–146.
- Jan Friso Groote and Alban Ponse. 1995. The syntax and semantics of μ CRL. In *Algebra of Communicating Processes*. Springer, 26–62.
- Dennis Guck, Hassan Hatefi, Holger Hermanns, Joost-Pieter Katoen, and Mark Timmer. 2014a. Analysis of timed and long-run objectives for Markov automata. *Logical Methods in Computer Science* 10 (2014), 1–29.
- Dennis Guck, Mark Timmer, Hassan Hatefi, Enno Ruijters, and Mariëlle Stoelinga. 2014b. Modelling and analysis of Markov reward automata. In *Int. Symp. on Automated Technology for Verification and Analysis (ATVA) (LNCS)*, Vol. 8837. Springer, 168–184.
- Gilles Kahn. 1974. The semantics of simple language for parallel programming. In *IFIP Congress*. 471–475.
- Joost-Pieter Katoen and Hao Wu. 2014. Exponentially timed SADF: Compositional semantics, reductions, and analysis. In *Int. Conf. on Embedded Software (EMSOFT)*. ACM.
- Edward A. Lee and David G. Messerschmitt. 1987. Synchronous data flow: Describing signal processing algorithm for parallel computation. In *COMPCON*. IEEE, 310–315.
- Marco Ajmone Marsan, Gianni Conte, and Gianfranco Balbo. 1984. A class of generalized stochastic petri nets for the performance evaluation of multiprocessor systems. *ACM TOCS* 2, 2 (1984), 93–122.
- Gethin Norman, David Parker, Marta Z. Kwiatkowska, Sandeep K. Shukla, and Rajesh Gupta. 2005. Using probabilistic model checking for dynamic power management. *Formal Aspects of Computing* 17, 2 (2005), 160–176.
- Sangyoung Park, Jaehyun Park, Donghwa Shin, Yanzhi Wang, Qing Xie, Massoud Pedram, and Naehyuck Chang. 2013. Accurate modeling of the delay and energy overhead of dynamic voltage and frequency scaling in modern microprocessors. *IEEE Transactions on CAD of Integrated Circuits and Systems* 32, 5 (2013), 695–708.
- Reza Pulungan and Holger Hermanns. 2015. A construction and minimization service for continuous probability distributions. *STTT* 17, 1 (2015), 77–90.
- Michael Rettelbach. 1995. Probabilistic branching in Markovian process algebras. *Computer Journal* 38, 7 (1995), 590–599.
- Ahsan Shabbir, Akash Kumar, Sander Stuijk, Bart Mesman, and Henk Corporaal. 2010. CA-MPSoC: An automated design flow for predictable multi-processor architectures for multiple applications. *Journal of Systems Architecture - Embedded Systems Design* 56, 7 (2010), 265–277.

- Sundararajan Sriram and Shuvra S. Bhattacharyya. 2009. *Embedded Multiprocessors: Scheduling and Synchronization*. CRC Press.
- Bart D. Theelen. 2007. A performance analysis tool for scenario-aware streaming applications. In *QEST*. 269–270.
- Bart D. Theelen, Marc Geilen, Twan Basten, Jeroen Voeten, Stefan Valentin Gheorghita, and Sander Stuijk. 2006. A scenario-aware data flow model for combined long-run average and worst-case performance analysis. In *MEMOCODE*. IEEE, 185–194.
- Bart D. Theelen, Marc Geilen, Sander Stuijk, Stefan Valentin Gheorghita, Twan Basten, Jeroen Voeten, and A. Ghamarian. 2008. *Scenario-Aware Dataflow*. Technical Report ESR-2008-08. TU Eindhoven.
- Bart D. Theelen, Joost-Pieter Katoen, and Hao Wu. 2012. Model checking of scenario-aware dataflow with CADP. In *DATE*. IEEE, 653–658.
- Mark Timmer, Joost-Pieter Katoen, Jaco van de Pol, and Mariëlle Stoelinga. 2012. Efficient modelling and generation of Markov automata. In *Int. Conf. on Concurrency Theory (CONCUR) (LNCS)*, Vol. 7454. Springer, 364–379.
- Mark Timmer, Jaco van de Pol, and Mariëlle Stoelinga. 2013. Confluence reduction for Markov automata. In *FORMATS (LNCS)*, Vol. 8053. Springer, 243–257.
- Stavros Tripakis, Dai N. Bui, Marc Geilen, Bert Rodiers, and Edward A. Lee. 2013. Compositionality in synchronous data flow: Modular code generation from hierarchical SDF graphs. *ACM Transactions on Embedded Computing Systems* 12, 3 (2013), 83:1–83:26.
- Håkan L. S. Younes, Marta Z. Kwiatkowska, Gethin Norman, and David Parker. 2006. Numerical vs. statistical probabilistic model checking. *STTT* 8, 3 (2006), 216–228.

Received January 2016; revised March 2016; accepted April 2016