

Schedule-Extended Synchronous Dataflow Graphs

Morteza Damavandpeyma, Sander Stuijk, Twan Basten, Marc Geilen and Henk Corporaal

This report is an extension of the following journal paper to be published. It adds the proofs omitted from the publication. If you want to cite this report, please refer to the paper instead.

M. Damavandpeyma, S. Stuijk, T. Basten, M.C.W. Geilen and H. Corporaal, “Schedule-Extended Synchronous Dataflow Graphs”. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 2013.


ES Reports

ISSN 1574-9517

ESR-2013-01

17 May 2013

Eindhoven University of Technology
Department of Electrical Engineering
Electronic Systems



Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending an email to pubs-permissions@ieee.org.

<http://www.es.ele.tue.nl/esreports>
esreports@es.ele.tue.nl

Eindhoven University of Technology
Department of Electrical Engineering
Electronic Systems
PO Box 513
NL-5600 MB Eindhoven
The Netherlands

Schedule-Extended Synchronous Dataflow Graphs

Morteza Damavandpeyma, *Student Member, IEEE*, Sander Stuijk, Twan Basten, *Senior Member, IEEE*,
 Marc Geilen, *Member, IEEE*, and Henk Corporaal, *Member, IEEE*

Abstract—Synchronous dataflow graphs (SDFGs) are used extensively to model streaming applications. An SDFG can be extended with scheduling decisions, allowing SDFG analysis to obtain properties like throughput or buffer sizes for the scheduled graphs. Analysis times depend strongly on the size of the SDFG. SDFGs can be statically scheduled using static-order schedules. The only generally applicable technique to model a static-order schedule in an SDFG is to convert it to a homogeneous SDFG (HSDFG). This may lead to an exponential increase in the size of the graph and to sub-optimal analysis results (e.g., for buffer sizes in multi-processors). We present techniques to model two types of static-order schedules, i.e., periodic schedules and periodic single appearance schedules, directly in an SDFG. Experiments show that both techniques produce more compact graphs compared to the technique that relies on a conversion to an HSDFG. This results in reduced analysis times for performance properties and tighter resource requirements.

Index Terms—Synchronous dataflow graphs, periodic schedules, single appearance schedules, schedule modeling.

I. INTRODUCTION

Synchronous dataflow graphs (SDFGs) are widely used to model digital signal processing (DSP) and multimedia applications [1]–[4]. Model-based design-flows (e.g., [1], [5]–[8]) model binding and scheduling decisions into the SDFG. This enables the analysis of performance properties (e.g. throughput [9]) or resource requirements (e.g., buffer sizes [10]) under resource constraints. Figure 1 shows an example of an SDFG with four *actors* and three *channels*. An essential property of SDFGs is that every time an actor *fires* (executes) it consumes a fixed number of tokens from its input edges and produces a fixed number of tokens on its output edges. These numbers are called the *rates* (indicated next to the channel ends when the rates are larger than 1). The fixed port rates make it possible to statically schedule SDFGs.

Many SDFG analysis algorithms, e.g., throughput calculation or buffer sizing, are straightforward when a single processor platform is used. For instance, the buffer sizes can be determined by executing the SDFG according to a given schedule. However, in a multi-processor environment, SDFG analysis algorithms are not trivial because of the inter-processor communication, amongst other reasons. An SDFG can be bound to a multi-processor platform. Each processor in the platform executes a set of actors from the SDFG; the firings of actors bound to a processor are required to be sequentialized. For this purpose, a finite periodic schedule can




Fig. 1. An example SDFG.

be constructed. Such a schedule is called a *periodic static-order schedule (PSOS)*. PSOSs only specify the firing order of actors. This separates them from fully static schedules, which determine absolute start times of actors (e.g., schedules generated using the technique of [11]). Traditionally, for DSP software synthesis, a sub-set of all periodic static-order schedules is considered. This sub-set contains so-called *single appearance schedules (SAS)* [1]. In a SAS, the functional code of the actors is included in a nested loop structure such that each piece of code occurs only once. This minimizes the code size potentially at the cost of additional buffer memory needed to implement the channels. A model-based design-flow usually uses PSOSs (or a sub-set of PSOSs such as SASs) for an application modeled with an SDFG. In this way timing (throughput) and memory usage (buffers) can be analyzed.

There is only one technique [12] known to model PSOSs in an SDFG. This technique requires a conversion of an SDFG to a so-called *homogeneous SDFG (HSDFG)* in which all rates are one [2]. Figure 2 (without the colored edges) shows the equivalent HSDFG of the SDFG in Figure 1. The technique of [12] sequentializes the actor firings by inserting a channel between each pair of consecutive actors in a schedule. At the end of a schedule, it adds a channel with one initial token from the last to the first actor in the schedule. This ensures an indefinite execution of the graph according to the schedule. To model PSOSs $s_0 = \langle a_0(a_2)^2 \rangle^*$ and $s_1 = \langle (a_1)^5(a_3)^3 a_1(a_3)^3 \rangle^*$, the technique of [12] adds in total 15 channels to the HSDFG of the example graph (the green edges for s_0 and the blue edges for s_1 in Figure 2). For example, s_0 indicates an indefinite sequence of one firing of a_0 followed by two firings of a_2 . This order is enforced in the HSDFG of Figure 2 by the green edges between the actors a_{0_1} , a_{2_1} , and a_{2_2} .

The SDFG to HSDFG conversion can lead to an exponential increase in the size of the graph. For example, such a conversion for an H.263 decoder [10] (with QCIF resolution) increases the graph size from 4 actors to 1190 actors. Note that the number of actors in the resulting HSDFG highly depends on how the application is modeled in the original SDFG. The run-time of SDFG analysis algorithms depends amongst others on the size of the graph. As a result, the run-time of many SDFG analysis algorithms may increase drastically when modeling PSOSs in the graph using the technique from [12]. For example, the buffer sizing algorithm from [10] takes less

The authors are with the Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands (e-mail: {m.damavandpeyma, s.stuijk, a.a.basten, m.c.w.geilen, h.corporaal}@tue.nl).

T. Basten is also with TNO, Embedded Systems Innovation, Eindhoven, The Netherlands.




Fig. 2. PSOSs $s_0 = \langle a_0(a_2)^2 \rangle^*$ and $s_1 = \langle (a_1)^5(a_3)^3 a_1(a_3)^3 \rangle^*$ modeled in the SDFG of Figure 1 using the technique from [12]; each $a_{i,j}$ actor in the HSDFG is an instance of SDFG actor a_i .

than 1 ms on the SDFG of an H.263 decoder. Modeling a schedule into this SDFG using the technique from [12], the same analysis lasts 1330 ms. SDFG analysis algorithms are usually repeated more than once in an iterative design-flow. As an example, for the SDFG of an H.263 decoder, the design-flow from [6] performs 8 throughput calculations on the SDFG to obtain the desired binding. Hence, it is vital to maintain a compact *schedule-extended graph*, i.e., a graph in which schedules are modeled explicitly, to provide a fast and practical design flow. There is a second drawback to the technique from [12]. The original graph structure is lost due to the conversion to an HSDFG. A single channel in an SDFG corresponds to a set of channels in the HSDFG. In Figure 2, for example, the six edges between actor $a_{0,1}$ and the $a_{1,j}$ actors correspond to the single edge between a_0 and a_1 in the SDFG of Figure 1. As a result, common buffer sizing techniques cannot find the minimal buffer size for the original SDFG. The H.263 decoder buffer sizes are for example overestimated by 49% when applying the technique of [10] to the HSDFG.

A novel technique is needed to model PSOSs in an SDFG. This technique should limit the increase in the number of actors such that analysis times do not increase too much when analyzing the SDFG with its schedules. The technique should also preserve the original graph structure as this enables accurate analysis of graph properties such as buffer sizes. In [13], we presented a schedule modeling technique, called DSM, to model any PSOS directly in an SDFG. In this paper, DSM is discussed in more detail. In addition, a second schedule modeling technique, called SASM, is introduced that is limited to SASs, but that results in more compact models compared to the first technique when modeling SASs. Correctness of the two approaches is formalized. Extensive experiments are carried out for evaluation purposes.

DSM and SASM can be used in any model-based design-flow that models PSOSs into the SDFG (e.g., [1], [5]–[8]). Conversion to an HSDFG may be inevitable at some steps of a design trajectory. For example, multi-processor scheduling may require such conversions, although some techniques exist

that can find schedules for SDFGs without any conversion to HSDFGs. For example, the technique presented in [14] solves the buffer sizing and scheduling problems simultaneously at the SDFG level. It is not the conversion from SDFG to HSDFG itself that is problematic though. The problem is that analyses or optimizations on large HSDFGs may be time consuming (e.g. throughput analysis) or inaccurate (e.g. buffer sizing). With our techniques, obtained schedules can be annotated back to the original SDFG; hence, the later analysis and optimization can be performed on the schedule-extended SDFG. Besides the already mentioned analyses, also for example dynamic voltage scaling can be directly applied to a schedule-extended SDFG model of an application mapped to a multi-processor platform [15]. Note that code generation is another step which requires an SDFG to HSDFG conversion; this conversion can be delayed until all (or most of the) prior analyses are carried out on the SDFG. Ultimately, the proposed techniques may save significant amounts of analysis time in a multi-processor design flow and they may lead to more accurate results.

The remainder of the paper is structured as follows. The next section discusses related work. Section III introduces SDFGs. Section IV formalizes SDFG schedules. Sections V and VI present our techniques to model PSOSs and SASs in an SDFG. Section VII contains the theorems related to the correctness of the presented techniques. We evaluate our technique by applying it to several realistic applications in Section VIII. Section IX concludes.

II. RELATED WORK

The technique from [12] is the only available technique to model PSOSs in an SDFG, through a conversion to an HSDFG. As already explained, this technique may result in a long run-time for analysis algorithms and/or inaccurate results from these algorithms. Our techniques alleviate both shortcomings of the technique from [12].

The work in [16] models the effect of a budget scheduler or preemptive TDMA scheduler on the temporal behavior of the SDFG, either by computing an accurate worst-case response time or, more precisely, by introducing additional actors to model the timing impact as a latency-rate model. In contrast, for non-preemptive schedules, such as PSOSs, we focus on the ordering of actor firings; their execution time remains the same. We force an SDFG to follow the PSOSs selected for each processor. This allows SDFG analysis to obtain properties like throughput or buffer sizes for the scheduled SDFG. This is also true for the models of [16]. However, for non-preemptive schedules, our results are tighter and our techniques require less analysis time. The authors of [17] have shown that an SDFG can be used to consider an application with resource sharing possibilities; they perform buffer sizing under a throughput constraint considering a given schedule for the actors using a shared resource. For shared resource analysis, they use event-models [18] which is based on realtime-calculus [19]. Our approach differs from [17], since modeling schedules directly into SDFGs enables us to use dedicated analysis for dataflow graphs. Moreover, the

technique of [17] can only handle an SDFG with limited types of cycles, such as cycles formed by self-edges or the back edges modeling the buffer capacity between two actors. However, staying in dataflow domain, as is done in our technique, does not impose such a limitation on the graph structure.

Ref [20] uses some new (custom) components, e.g., *if – then – else*, to model schedules in an SDFG. These components are not supported by the common model-based design-flows using SDFGs (e.g., [1], [5]–[8]) and cannot be modeled in an SDFG using the basic elements of an SDFG (i.e., actors and channels). Our techniques eliminate the need for any new (custom) component. As a result, any analysis technique for SDFGs is directly applicable to the schedule-extended SDFG.

III. SYNCHRONOUS DATAFLOW GRAPHS

Let \mathbb{N} denote the positive natural numbers and $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. Consider *Ports* as a set that contains all ports; each port $p \in Ports$ has a finite rate $Rate(p) \in \mathbb{N}$. An actor a_i is a tuple (In, Out) consisting of a set $In \subseteq Ports$ of input ports and a set $Out \subseteq Ports$ of output ports with $In \cap Out = \emptyset$.

Definition 1. (SDFG) *An SDFG is a tuple (A, C) consisting of a finite set A of actors and a finite set $C \subseteq Ports^2$ of channels. The channel source is an output port of an actor and the channel destination is an input port of an actor. Each port of an actor is connected to only one channel and each channel end is connected to a single port. For every actor $a_i = (In, Out) \in A$, $InC(a_i)$ represents all channels connected to the ports in In and $OutC(a_i)$ represents all channels connected to the ports in Out .*

The SDFG of Figure 1 has four actors ($A = \{a_0, a_1, a_2, a_3\}$) and three channels ($C = \{c_0, c_1, c_2\}$). Actors communicate with *tokens* sent from one actor to another over the channels. Tokens are depicted with a solid dot (and an attached number in case of multiple tokens). An essential property of SDFGs is that every time an actor *fires* (executes) it consumes a fixed number of tokens from its input edges and produces a fixed number of tokens on its output edges. These numbers are called the *rates* (indicated next to the channel ends when the rates are larger than 1). The rates determine how often actors have to fire with respect to each other such that the distribution of tokens over all channels is not changed. This property is captured in the *repetition vector* [1] of an SDFG. The repetition vector determines the number of times each actor should be fired in order to bring the SDFG back to its initial token distribution. Notation $\gamma(a)$ refers to the repetition vector value of actor a . The repetition vector of the SDFG shown in Figure 1 is $\gamma = [1 \ 6 \ 2 \ 6]^T$. It corresponds to 1 firing of a_0 , 6 firings of a_1 , 2 firings of a_2 and 6 firings of a_3 . Channels can contain different numbers of tokens. A *state* of an SDFG (represented by ω) is described by the number of tokens in all channels of the SDFG. We assume that the initial state of an SDFG is given by an initial token distribution ω_0 . An actor $a_i \in A$ is *enabled* in an SDFG state ω_j iff $\omega_j(c) \geq Rate(q)$ for each channel $c = (p, q) \in InC(a_i)$. When an actor a_i starts its firing, it removes $Rate(q)$ tokens

from all $(p, q) \in InC(a_i)$ and when it ends, it produces $Rate(p)$ tokens on every $(p, q) \in OutC(a_i)$. Consistency (i.e., the existence of a repetition vector) and absence of deadlock are necessary in practice for SDFGs and can be verified efficiently [21], [22]. Any SDFG which is not consistent requires unbounded memory to execute or deadlocks. Therefore, we limit ourselves to consistent and deadlock-free SDFGs.

IV. SDFG STATIC-ORDER SCHEDULING

Assume that the initial state ω_0 for the SDFG of Figure 1 is equal to $(c_0, c_1, c_2) \rightarrow (0, 0, 0)$. Actor a_0 is enabled in ω_0 . Firing a_0 results in a transition from state ω_0 to a state $(6, 0, 0)$. We use this concept of states and transitions to formalize the execution of an SDFG.

Definition 2. (EXECUTION) *An execution σ of an SDFG is an infinite alternating sequence of states and transitions $\omega_0 \xrightarrow{a_i} \omega_1 \xrightarrow{a_j} \dots$ starting from a designated initial state ω_0 .*

In a multi-processor system, multiple actors may be bound to the same processor. These actors may be enabled at the same time. In such a situation, a schedule is needed to order the firings of the enabled actors on the processor. The fixed port rates make it possible to statically schedule SDFGs with a finite schedule per processor. Such a schedule orders the actor firings on the underlying processor. This type of schedules, which are called *periodic static-order schedules (PSOSs)*, can be repeated indefinitely. A separate PSOS should be constructed for each processor. Each PSOS only includes actors bound to this specific processor. The following definition is used to formally specify a PSOS.

Definition 3. (PERIODIC STATIC-ORDER SCHEDULE (PSOS)) *A PSOS is a finite ordered list of (a sub-set of) actors in an SDFG (A, C) . A PSOS is denoted by $s_i = \langle \alpha_1 \alpha_2 \dots \alpha_n \rangle^*$ where each $\alpha_j |_{1 \leq j \leq n}$ is a sub-schedule that represents an actor from A and $n \in \mathbb{N}$ is the length of the schedule s_i , represented by $n = |s_i|$. The set A_i contains all actors that appear at least once in s_i ($A_i \subseteq A$).*

A PSOS can be represented in a compact format, called a looped schedule (LS).

Definition 4. (LOOPED SCHEDULE (LS)) *A looped schedule, $s_i = \langle (\alpha_1)^{\beta_1} (\alpha_2)^{\beta_2} \dots (\alpha_m)^{\beta_m} \rangle^*$, is defined as a successive execution of α_1 repeated β_1 times followed by α_2 repeated β_2 times and so on, where each α_j is either an actor firing or a (nested) looped schedule and $\beta_j \in \mathbb{N}$.*

Definition 5. (SINGLE APPEARANCE SCHEDULE (SAS)) *A LS in which each actor appears only once is called a single appearance schedule (SAS).*

Assume that the SDFG of Figure 1 is mapped to a platform with two processors (P_0 and P_1). Actors a_0 and a_2 are mapped to P_0 with the PSOS $s_0 = \langle a_0(a_2)^2 \rangle^*$ and actors a_1 and a_3 are mapped to P_1 with the PSOS $s_1 = \langle (a_1)^5(a_3)^3 a_1(a_3)^3 \rangle^*$. PSOS s_0 is a SAS while PSOS s_1 is not. PSOS $s'_1 = \langle (a_1)^3(a_3)^3 \rangle^*$ can be used as a SAS for actors a_1 and a_3 .

Definition 6. (SDFG ITERATION) *Assume SDFG (A, C) has*

repetition vector γ . An SDFG iteration is a set of actor firings such that for each $a \in A$, the set contains $\gamma(a)$ firings of a .

Definition 7. (PSOS ITERATION) Let $s_i = \langle \alpha_1 \alpha_2 \dots \alpha_n \rangle^*$ be a PSOS that schedules actors in $A_i \subseteq A$. A PSOS iteration is a sequence of actor firings following the actor order specified in s_i starting from actor α_1 and ending with actor α_n with a length equal to $|s_i|$ and including only actors from A_i .

The actor firing order in an execution $\sigma = \omega_0 \xrightarrow{a_x} \omega_1 \xrightarrow{a_y} \dots$ can be captured using a list $\langle a_x, a_y, \dots \rangle$ where the j^{th} element in this list is the actor which is fired in the transition from ω_{j-1} to ω_j . The notation $orderList(\sigma, A_i)$ represents the mentioned list where actors which do not belong to A_i are omitted. For example, in the SDFG of Figure 1, consider an execution σ that results in list $\langle a_0, a_1, a_1, a_1, a_2, a_1, a_1, a_3, a_3, a_3, a_1, a_2, a_3, a_3, a_3 \rangle$ with $A_1 = \{a_1, a_3\}$; then $orderList(\sigma, A_1) = \langle a_1, a_1, a_1, a_1, a_1, a_3, a_3, a_3, a_1, a_3, a_3, a_3 \rangle$. We say that the corresponding execution of an SDFG satisfies a PSOS when the SDFG is executed according to the PSOS. We use the following to formalize this term.

Definition 8. (SATISFACTION) Let σ be an execution of an SDFG (A, C) and s_i a PSOS for actors $A_i \subseteq A$. If it exists, let σ' be the prefix of σ such that it contains exactly $\gamma(a_i)$ occurrences of actor $a_i \in A_i$; σ' covers σ precisely up to the point that one PSOS iteration is executed. Execution σ satisfies PSOS s_i iff σ' exists and the ordered list $orderList(\sigma', A_i)$ corresponds to the order specified in s_i .

When an execution of a consistent and deadlock-free SDFG satisfies the specified PSOSs, the channels of the SDFG need bounded memories (according to Theorem 1 from [23]). The number of actor appearances in the PSOS is a fraction or multiple of its repetition vector entry. Formally, each actor a_i in the PSOS should appear $r \cdot \gamma(a_i)$ times in the PSOS (with $r = \frac{u}{v}$ where $u, v \in \mathbb{N}$) and the value r is identical for all actors in the PSOS [9]. This follows from the SDFG property that firing each actor as often as indicated in the repetition vector results in a token distribution that is equal to the initial token distribution. In the paper, the term *normalized PSOS* is used to refer to a PSOS with r equal to 1.

Definition 9. (NORMALIZED PSOS) A PSOS s_i is called *normalized* iff each actor $a_j \in A_i$ appears $\gamma(a_j)$ times in one iteration of the PSOS s_i .

We limit ourselves in the remainder to PSOSs in which r is a *unit fraction* (i.e., $r = \frac{1}{v}$ with $u = 1$ and $v \in \mathbb{N}$), although our technique can also be directly applied to model other PSOSs (i.e., in which $u \in \mathbb{N}$).

V. MODELING PERIODIC STATIC-ORDER SCHEDULES

In this section, we introduce a technique to model PSOSs in an SDFG. We first illustrate all ingredients of our approach through a running example, and then discuss the algorithm and the main steps in the algorithm in more detail. Note that a schedule is correctly modeled if and only if any execution of the schedule-extended graph satisfies the schedule and if

any execution of the original graph that satisfies the schedule is still feasible in the schedule-extended graph.

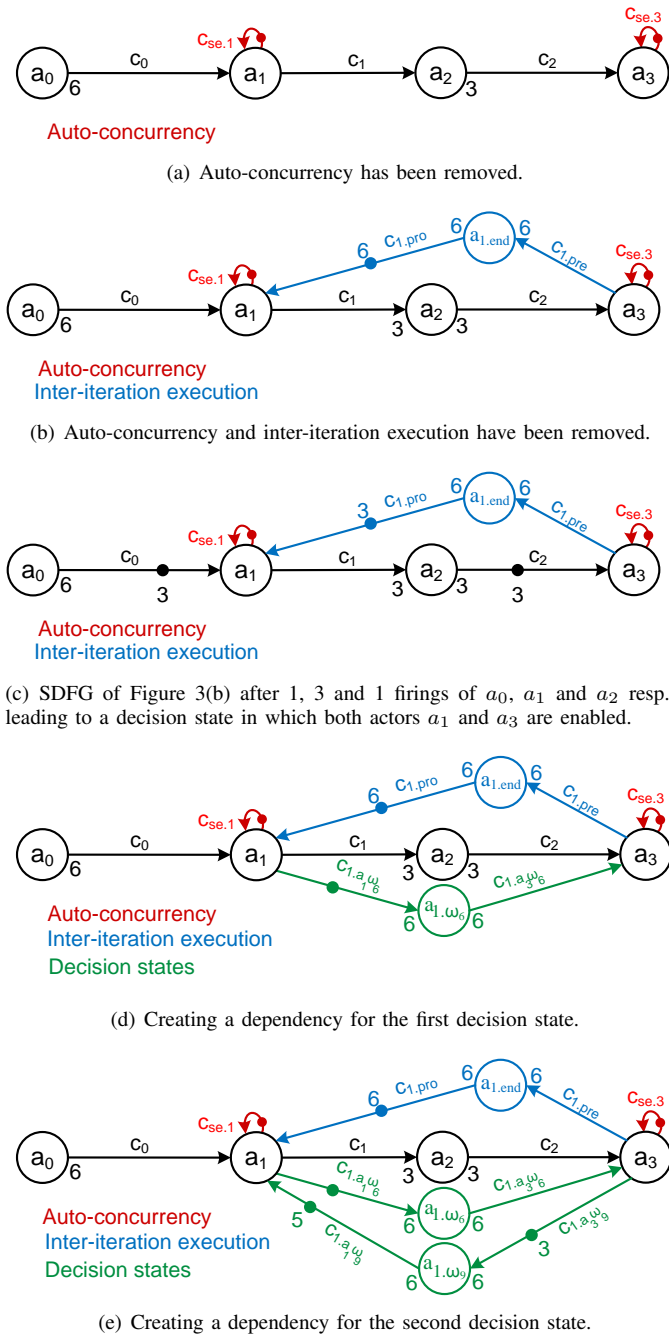
A. Running example

Here we briefly introduce our approach in modeling a PSOS in an SDFG using a running example. For this purpose consider the example SDFG shown in Figure 1 and the PSOS $s_1 = \langle (a_1)^5 (a_3)^3 a_1 (a_3)^3 \rangle^*$ which is a schedule for a_1 and a_3 . Our approach captures this schedule in the SDFG in three steps, that (i) remove auto-concurrency, (ii) avoid inter-iteration execution, and (iii) enforce the correct scheduling decisions.

In the example SDFG, a single firing of a_0 produces 6 tokens in channel c_0 ; then 6 firings of a_1 can be performed simultaneously. This simultaneous firing of an actor is called *auto-concurrency*; in practice, this corresponds to executing multiple instances of a function (task) in parallel. Auto-concurrency for an actor cannot be handled in a real hardware platform, unless more than one processor is allocated for that actor. In this work, we focus on the case that an actor is mapped to one processor. Hence, auto-concurrency must be removed from the model. To sequentialize firings of an actor, a self-edge with one initial token must be added to that actor; this way auto-concurrency can be removed for that actor. Figure 3(a) shows the example SDFG of Figure 1 in which any auto-concurrency related to a_1 and a_3 is removed by adding two self-edges (shown in red).

In one PSOS iteration, each actor must fire a specific number of times. Actors must not be able to get fired more than the number of times indicated by the PSOS. Consider the following situation in the example SDFG of Figure 1. Two firings of a_0 provide 12 tokens in channel c_0 ; this number of tokens is enough for 12 firings of the actor a_1 . The first 6 firings of a_1 belong to the first iteration and the second 6 firings belong to the second iteration. The second 6 firings of a_1 can occur before the completion of the first PSOS iteration of s_1 ; in this case, the resulting execution does not satisfy the given PSOS s_1 . This situation is called *inter-iteration execution*. To prevent inter-iteration execution related to s_1 , we create a dependency from the last actor appearing in s_1 to the first actor appearing in s_1 ; see the blue elements in Figure 3(b). This dependency limits the firing of the first actor to a number of firings required in one PSOS iteration.

In the SDFG of Figure 3(b), consider the case that actors a_0 , a_1 and a_2 have fired 1, 3 and 1 times, respectively. The initial token distribution is changed to the distribution shown in Figure 3(c). In this graph, both a_1 and a_3 from PSOS s_1 are enabled; but, only the firing of a_1 must be granted at this point to form an execution that satisfies the given PSOS s_1 . We call such a state in which several actors are enabled a *decision state*. Later on, a precise definition is given for this concept. According to schedule s_1 , actor a_3 must get enabled after 5 firings of a_1 . For this purpose, a dependency is created from a_1 to a_3 (shown with green actor and channels in Figure 3(d)); this new dependency prevents a_3 from getting enabled unless a_1 has completed 5 firings. Another decision state can be found after a_1 has completed 5 firings; once again, both a_1 and a_3


 Fig. 3. Step-by-step modeling of the PSOS s_1 in the SDFG of Figure 1.

from PSOS s_1 are enabled at this point; but, the firing of a_3 must take place. For this purpose, a dependency is created from a_3 to a_1 (see Figure 3(e)). This new dependency prevents a_1 from getting enabled from the current state onwards unless a_3 has completed 3 firings. The 5 initial tokens on the added input channel to a_1 ensure that the first 5 firings of a_1 can take place as planned. The SDFG of Figure 3(e) shows the final solution that models the PSOS s_1 in the SDFG of Figure 1.

B. The algorithm

Algorithm 1 captures our technique, called decision state modeling (DSM). DSM accepts an SDFG and one or several

PSOSs as its input. In the algorithm $n + 1$ ($n \in \mathbb{N}_0$) is the number of processors (or input PSOSs). DSM ensures that actor firings of each PSOS follow the specified order in that PSOS; the output of DSM is a new SDFG that models the provided PSOSs in the input SDFG. Figure 4 depicts the corresponding SDFG of Figure 1 which models the PSOSs s_0 and s_1 using DSM. The remainder of this section discusses the three main steps of the algorithm - removing auto-concurrency, avoiding inter-iteration execution, enforcing correct decisions in decision states - in detail.


The description of some basic functions used in Algorithm 1 is as follows. The function $AA(G, a_{new})$ is responsible to include the actor a_{new} in the SDFG G . The function $AC(G, c_{new}, a_{src}, a_{dst}, srcRate, dstRate, initTok)$ adds the channel c_{new} from source actor a_{src} to destination actor a_{dst} ; the production rate of a_{src} on this channel is equal to $srcRate$ and the consumption rate of a_{dst} on this channel is equal to $dstRate$; this channel is initialized with $initTok$ initial tokens. The function $CNT(a_j, s_i)$ returns the number of times that the actor a_j is fired in one iteration of PSOS s_i . The function $BEF(a_k, j, s_i)$ returns the number of times that a_k appears from the first position in the PSOS s_i to and including the j^{th} position in the PSOS s_i ; the function $AFT(a_k, j, s_i)$ returns the number of times that a_k appears from the $(j + 1)^{th}$ position in the PSOS s_i to the last position in the PSOS s_i .

C. Auto-concurrency

An actor $a_i \in A$ can be enabled multiple times simultaneously in an SDFG state ω_j if $\omega_j(c) \geq k \cdot Rate(q)$ for each channel $c = (p, q) \in InC(a_i)$ where $k \in \mathbb{N}, k \geq 2$. This is called auto-concurrency. The firings of actor a_i should occur sequentially according to the PSOS to which a_i belongs. This sequential execution can be enforced by adding a self-edge with one initial token to actor a_i (Line 1 in Algorithm 1). In Figure 4, channels $c_{se.0} - c_{se.3}$ (shown in red) are used to prevent any auto-concurrency in the SDFG of Figure 1.

D. Inter-iteration execution

Consider actor a_0 in the SDFG of Figure 1; the 1^{st} firing of a_0 belongs to the 1^{st} PSOS iteration of s_0 and the 2^{nd} firing of a_0 belongs to the 2^{nd} PSOS iteration of s_0 and so on. Since a_0 has no input channel, it can always be fired regardless of other actors in the graph. This behavior, which is called inter-iteration execution, can prevent an SDFG execution from satisfying the given PSOSs. Inter-iteration execution happens when one PSOS iteration has not been completed and an actor


 Fig. 4. PSOSs s_0 and s_1 modeled in the SDFG of Figure 1 using DSM.

Algorithm 1: Decision State Modeling (DSM)

```

input : SDFG  $G(A, C)$ , PSOSs  $\{s_0, \dots, s_n\}$ 
output:  $G$  extended with schedules  $\{s_0, \dots, s_n\}$ 

1 add a self edge with 1 initial token for each  $a \in A$ 
2  $\{s'_0, \mu_0, \dots, s'_n, \mu_n\} \leftarrow \text{normalize}(G, \{s_0, \dots, s_n\})$ 
3 for  $i \leftarrow 0$  to  $n$  do
  /* To control inter-iteration execution */
4   $a_L := \text{last actor in } s_i$ 
5   $a_F := \text{first actor in } s_i$ 
6   $\mathbf{AA}(G, a_{i.end})$ 
7   $\mathbf{AC}(G, c_{i.pre}, a_L, a_{i.end}, 1, \mathbf{CNT}(a_L, s_i), 0)$ 
8   $\mathbf{AC}(G, c_{i.pro}, a_{i.end}, a_F, \mathbf{CNT}(a_F, s_i), 1, \mathbf{CNT}(a_F, s_i))$ 
  /* To control decision states */
9   $\Omega, pos \leftarrow \text{getDecisionStates}(G, s'_i, \{s'_0, \dots, s'_n\} \setminus s'_i)$ 
10  $\Omega \leftarrow \text{reduceDecisionStates}(\Omega)$ 
11  $\Omega \leftarrow \text{foldDecisionStates}(\Omega, \mu_i)$ 
12 foreach  $\omega_j \in \Omega$  do
13    $\mathbf{AA}(G, a_{i.\omega_j})$ 
14   foreach  $a_k \in \Delta_j$  do
15     if  $a_k$  is the actor of choice then
16        $\mathbf{AC}(G, c_{i.a_k.\omega_j}, a_k, a_{i.\omega_j}, 1, \mathbf{CNT}(a_k, s_i),$ 
17          $\mathbf{AFT}(a_k, pos[\omega_j], s_i))$ 
18     else
19        $\mathbf{AC}(G, c_{i.a_k.\omega_j}, a_{i.\omega_j}, a_k, \mathbf{CNT}(a_k, s_i), 1,$ 
20          $\mathbf{BEF}(a_k, pos[\omega_j], s_i))$ 

```





Fig. 5. An example SDFG.

from that PSOS can proceed its firings beyond the current PSOS iteration. Lines 4-8 in Algorithm 1 are used to control this undesirable actor enabling. This part of the algorithm adds (per PSOS) one actor and two channels to create a dependency between the last and first actor appearing in the PSOS. The added components limit, within one PSOS iteration, the firing of the first actor in the PSOS (i.e., a_F) to the count of actor a_F (i.e., $\mathbf{CNT}(a_F, s_i)$) in one iteration of the PSOS s_i . The next iteration of the PSOS s_i can only commence if the last actor in PSOS s_i (i.e., a_L) fires $\mathbf{CNT}(a_L, s_i)$ times in one iteration of the PSOS s_i . In other words, the next iteration of a PSOS can only commence after the completion of the current iteration of this PSOS. In Figure 4, actor $a_{0.end}$ and channels $c_{0.pre}$ and $c_{0.pro}$ are added to prevent any inter-iteration execution in PSOS s_0 . Actor $a_{1.end}$ and channels $c_{1.pre}$ and $c_{1.pro}$ are added to prevent any inter-iteration execution in schedule s_1 . These elements are shown in our example in blue in Figure 4.

E. Decision states

This sub-section presents the third step of DSM. It first defines the concept of a decision state and then proceeds with the algorithm for identifying decision states; after explaining two optimization steps, it ends with the technique to enforce the appropriate schedule decisions.

1) *Concept*: Multiple different actors mapped to a single processor may be enabled in a specific state. Here, we describe


 Fig. 6. The state space of the SDFG of Figure 5 when PSOSs $s'_0 = \langle a_0 \rangle^*$ and $s'_1 = \langle (a_1)^2 a_2 a_1 a_2 \rangle^*$ are used.

such situations in an SDFG execution. Consider the SDFG in Figure 5. Assume that a_0 is mapped to processor P_0 with PSOS $s'_0 = \langle a_0 \rangle^*$ and a_1 and a_2 are mapped to processor P_1 with PSOS $s'_1 = \langle (a_1)^2 a_2 a_1 a_2 \rangle^*$. For brevity, we only discuss the actors mapped to processor P_1 . The corresponding state space - for one SDFG iteration - when executing our example SDFG using the PSOSs s'_0 and s'_1 is visualized in Figure 6. In this figure, the actors mapped to processor P_0 (P_1) are surrounded by a square (circle). Auto-concurrency and inter-iteration execution are excluded using the constructs introduced in Section V-C and Section V-D resp. The periodic behavior of the PSOSs is obvious from the state space where one SDFG iteration moves the graph to its initial state, i.e., ω_0 (see Figure 6). There are some states in Figure 6 in which more than one actor is enabled (e.g., $\omega_1 - \omega_5$) on processor P_1 . In such a situation, the execution related to those actors can deviate from the specified PSOS. We use the following definition to formalize such a situation.

Definition 10. (DECISION STATE) Consider the PSOS s_i as a schedule for actors $A_i \subseteq A$ and an execution σ of an SDFG (A, C) which satisfies PSOS s_i . A state $\omega_j \in \sigma$ is a decision state iff multiple different actors from A_i are enabled in ω_j .

We use Ω to refer to the finite set containing all decision states occurring in the execution of an SDFG up-to one iteration of a PSOS. The following terminology describes the enabled actors in a decision state.

Definition 11. (OPPONENT ACTOR SET) Let $\omega_j \in \Omega$ be a decision state for PSOS s_i . The opponent actor set Δ_j of the decision state ω_j is a finite set which contains all actors that are enabled in decision state ω_j and that belong to A_i .

The finite set Δ_j represents the opponent actors in the decision state $\omega_j \in \Omega$. One of the enabled actors in a decision state ω_j , in line with the given PSOS s_i , should be selected to fire. The following is used to describe such an actor.

Definition 12. (ACTOR OF CHOICE) Consider the PSOS s_i which schedules actors $A_i \subseteq A$ and the opponent actor set Δ_j of the decision state ω_j in an execution σ of the SDFG $G(A, C)$ which satisfies s_i . An actor $a_c \in \Delta_j$ is called the actor of choice of the decision state ω_j iff the firing of actor a_c in state ω_j is a necessity for the execution σ in order to satisfy the PSOS s_i . Since a PSOS specifies a fixed firing order, there can only be a single actor of choice in any decision state.

Lines 9-18 in DSM show how we deal with non-deterministic execution due to decision states. DSM models the given PSOSs one-by-one iteratively. The ordering of

Algorithm 2: Get Decision States

input : SDFG G , PSOS s_c , PSOSs $\{s_{o1}, \dots, s_{on}\}$
output: Decision state set Ω
output: relative positions pos

```

1  $\omega_j \leftarrow$  the initial state of  $G$ 
2 for  $i \leftarrow 1$  to  $|s_c|$  do
3    $\omega_j \leftarrow \text{maxExec}(G, \omega_j, \{s_{o1}, \dots, s_{on}\})$ 
4   if  $\text{sizeof}(\text{enabledActors}(G, \omega_j, s_c)) > 1$  then
5      $pos[\omega_j] \leftarrow i$ 
6      $\Omega \leftarrow \Omega \cup \{\omega_j\}$ 
7      $\Delta_j \leftarrow \text{enabledActors}(G, \omega_j, s_c)$ 
8    $\omega_j \leftarrow \text{fire}(G, \omega_j, s_c[i])$ 
    
```




Fig. 7. The state space of the SDFG of Figure 1 when the PSOS $s_1 = \langle (a_1)^5 (a_3)^3 a_1 (a_3)^3 \rangle^*$ is the schedule of interest (i.e., s_c) in Algorithm 2.

PSOSs in DSM does not have any impact on the final behavior. In each iteration of the for-loop in line 3, we enforce the execution of the actors in the current schedule of interest (i.e., schedule s_i) to follow schedule s_i . The next sub-section explains how decision states of the schedule of interest are extracted. At the same time, a value is preserved for each decision state that captures the relative position of that decision state with respect to the beginning of the schedule of interest; the notation $pos[\omega_j]$ refers to this position for ω_j . For example, in the SDFG of Figure 1, $pos[\omega_6] = 5$ since the relative position of ω_6 with respect to the beginning of the schedule of interest (i.e., s_1) is 5 (in Figure 7 consider 4 firings related to s_1 have been occurred before ω_6 and the 5th actor firing related to s_1 is going to happen in ω_6). For each $\omega_j \in \Omega$ extracted for s_i , DSM adds an actor (a_{i,ω_j} in line 13) and one channel between the new actor a_{i,ω_j} and each opponent actor in the set Δ_j (lines 14-18 in Algorithm 1). The elements added in each decision state (e.g., ω_j) postpone the execution of the actors in $\Delta_j \setminus \{a_c\}$ to the state after decision state ω_j . Hence, a_c (i.e., the actor of choice) is the only actor which can be fired in the state ω_j .

2) *Decision state identification*: Algorithm 2 shows our technique to detect all decision states within PSOS s_c . Assume s_c is a PSOS for the actors mapped to processor P_c . Schedules $s_{o1} \dots s_{on}$ are PSOSs for the other actors of the SDFG mapped to the other processors (denoted by $P_{o1} \dots P_{on}$). The output of Algorithm 2 is a set that contains all decision states for PSOS s_c . This algorithm also returns the relative positions of decision states with respect to the beginning of s_c . In Algorithm 2, the input schedules are normalized PSOSs. The function *normalize* (in line 2 of Algorithm 1) normalizes the input PSOSs. The function returns the normalized PSOSs along with their normalization factors. The normalized PSOS s'_x can be achieved by repeating μ_x times the input PSOS s_x (i.e., $s'_x = (s_x)^{\mu_x}$). μ_x is the normalization factor of s_x and can be calculated by dividing the repetition vector entry of an

arbitrary actor in s_x by the count of that actor in the PSOS s_x (in our running example, μ_0 and μ_1 are 1).

After normalizing the input schedules, all situations that can lead to multiple actors (mapped to the same processor) being ready to fire must be discovered. An actor in the schedule of interest s_c could be affected by the execution of an actor in the other schedules as well as by another actor in s_c . Processors can run at different clock rates; these differences and inter-processor dependencies cause variation in the number of tokens on the *inter-processor channels* originating from the actors mapped to the other processors to the actors mapped to the processor of interest (i.e., P_c). The number of tokens on the input channels of an actor determines whether an actor is enabled or not. To determine any possible actor enabling within s_c , a *necessary and sufficient number of tokens* on all inter-processor channels entering to the actors mapped to processor P_c must be considered. We will now explain what necessary and sufficient number of tokens means in our algorithm. Each iteration of the schedule of interest s_c requires that the actors mapped to the other processors are fired up to at most their repetition vector entry values. Hence, only executing one iteration of the other schedules $s_{o1} \dots s_{on}$ is enough to provide a *necessary number of tokens* on inter-processor channels entering to the actors mapped to processor P_c . More than one iteration for the other schedules $s_{o1} \dots s_{on}$ may be feasible; this may cause an actor in s_c to be enabled more than its count in one iteration of s_c . The inter-iteration prevention constructs introduced in Section V-D control this undesired actor enabling. So, we only extract decision states within one iteration of the normalized schedule to provide a *sufficient number of tokens*.

Also, DSM does not impose any limitation between PSOSs since no dependency is created between actors mapped to different processors. PSOSs can independently be iterated if the dependencies in the SDFG allow that. We allow the actors on the other processors to be executed (according to their schedules) as much as they can; the corresponding execution is named *maximal execution*. The maximal execution will stop at one point either due to a dependency on the actors on the processor P_c or because one PSOS iteration is completed. The SDFG state (denoted by ω_j in Algorithm 2) should be kept during the operation of the algorithm. The maximal execution is represented by the function *maxExec* in Algorithm 2. After one maximal execution, the number of tokens on the inter-processor channels entering into the actors on the processor P_c determines any possible enabled actor. The preserved state (i.e., ω_j) will be added to the decision state set (Ω) if more than one actor on the processor P_c is enabled at this state (line 6 in Algorithm 2). The current position (i.e., i) in the schedule of interest s_c is also preserved for the discovered decision state (see line 5 in Algorithm 2). All enabled actors will be recorded as opponent actors of the state ω_j (line 7 in Algorithm 2). The execution of the actors on the processor P_c is continued by executing the enabled actor in line with s_c to determine all possible decision states (line 8 in Algorithm 2). The function *fire*($G, \omega_j, s_c[i]$) fires the actor at the i th position in the PSOS s_c . The process is repeated until a full iteration of the PSOS has been examined. In the end, the set




Fig. 8. A third example SDFG.

Ω contains all possible decision states when executing s_c .

Figure 7 depicts the resulting state space from applying Algorithm 2 on the SDFG of Figure 1 when s_1 is the schedule of interest (i.e., s_c). In the SDFG of Figure 1, five consecutive decision states ($\Omega = \{\omega_5 \dots \omega_9\}$) have been found for s_1 and no decision state has been found for s_0 (see Figure 7).

3) *Redundant decision states*: Here, we explain an optimization proposed in DSM to remove unnecessary decision states. DSM adds some components (per decision state) to create a dependency from the actor of choice of a decision state to the other opponent actors of that decision state. Such a dependency delays the firing of those opponent actors to the state after the decision state. The added components are explained in detail in Section V-E5.

It is possible to have several consecutive decision states which are delaying the firing of an actor to several states later. For example, three consecutive decision states ($\omega_7 - \omega_9$) exist in Figure 7 that all delay the firing of a_1 ; the added components in ω_7 delay the sixth firing of a_1 to ω_8 ; the added components in ω_8 delay the sixth firing of a_1 to ω_9 ; and so on. The latest decision state in the sequence of decision states $\omega_7 - \omega_9$ is enough to delay the firing of a_1 to ω_{10} . Hence, the decision states $\omega_7 - \omega_8$ are redundant and can be removed from the decision state set Ω . The function *reduceDecisionStates* is responsible for removing redundant decision states. Note that it would be possible to perform this reduction during the decision state identification step. This optimization can significantly reduce the number of extra components in the final SDFG. Decision state ω_5 is also redundant according to our optimization. So, only two decision states ω_6 and ω_9 are necessary to model s_1 in the SDFG of Figure 1.

4) *Decision state folding*: In Algorithm 1, the input PSOSs are normalized to find all decision states. The normalization of PSOSs is required to explore sufficient states of an SDFG. Consider PSOSs $s_2 = \langle a_0 \rangle^*$ and $s_3 = \langle a_2 a_1 \rangle^*$ for our second example SDFG in Figure 8. To obtain normalized PSOSs, μ_2 and μ_3 must be equal to 3 and 4 respectively. This leads to the following normalized PSOSs: $s_2' = \langle (a_0)^3 \rangle^*$ and $s_3' = \langle (a_2 a_1)^4 \rangle^*$. Decision state identification for s_3' results in 5 decision states. $\begin{pmatrix} a_2 \\ - \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} \begin{pmatrix} a_2 \\ a_1 \end{pmatrix} \begin{pmatrix} a_2 \\ a_1 \end{pmatrix} \begin{pmatrix} a_2 \\ a_1 \end{pmatrix} \begin{pmatrix} a_2 \\ a_1 \end{pmatrix} \begin{pmatrix} a_2 \\ - \end{pmatrix} \begin{pmatrix} a_1 \\ - \end{pmatrix}$ shows

the corresponding execution of s_3' . In construct $\begin{pmatrix} a_x \\ a_y \end{pmatrix}$, a_x is the enabled actor in line with the PSOS and a_y is the other enabled actor if any at all. In this execution, the 1st, 3rd, 5th and 7th states are similar in behavior since the same actor (i.e., a_2) is expected to fire in all of those states.

Modeling a repetitive behavior for a PSOS s_i , also models this behavior for its normalized PSOS (i.e., $s_i' = (s_i)^{\mu_i}$). Using this property, we can merge decision states appearing in all μ_i repetitions of the PSOS s_i . We call this optimization *decision state folding* (line 11 in Algorithm 1). Folding groups the similar states. In our example, the 1st, 3rd, 5th and 7th states

are grouped and represented with one state. Similarly, the 2nd, 4th, 6th and 8th states are grouped. So, the above execution shrinks to $\begin{pmatrix} a_2 \\ a_1 \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}$. After grouping all similar states in the original execution into a representative state, it is considered a decision state if any of the group members is a decision state. In practice, a decision state in a state of the new folded execution will be considered as a decision state for each of the equivalent states in the original execution. This cannot violate the execution according to the input PSOS because DSM ensures that only the actor of choice executes in a decision state. This optimization could reduce the number of decision states up to μ_i times in a normalized PSOS s_i' . The firings related to those actors enabled in the last state except the actor of choice of that state are supposed to happen in subsequent PSOS iterations; this is already ensured by preventing inter-iteration execution (see Section V-D). Hence, after folding, the last state can be ignored as a decision state. In our third example, decision state folding reduces the number of decision states from 5 to 1 for the PSOS s_3 .

5) *Enforcing a schedule in decision states*: In our first example SDFG, only two actors are enabled in decision state ω_6 (i.e., $\Delta_6 = \{a_1, a_3\}$) (see Figure 7). Actor a_1 is the actor of choice in ω_6 and a_3 is the only opponent actor whose execution should be delayed to the state after ω_6 . To enforce firing of a_1 and to prevent firing of a_3 in ω_6 , DSM creates a dependency from a_1 to a_3 by adding actor $a_{1.\omega_6}$ and channels $c_{1.a_1\omega_6}$ and $c_{1.a_3\omega_6}$. The rates and initial tokens related to the new elements are set in such a way that the execution of the graph in other states are not affected. The actor $a_{1.\omega_6}$ is only responsible for decision state ω_6 . So, $a_{1.\omega_6}$ needs to only fire once in an iteration. For this purpose, the port rates of $a_{1.\omega_6}$ on its channels (i.e., $c_{1.a_1\omega_6}$ and $c_{1.a_3\omega_6}$) are set to 6. The added dependency channels from the newly added actor in decision state ω_j (e.g., $a_{1.\omega_6}$ in ω_6) to the opponent actors which are not the actor of choice (e.g., a_3 in ω_6) only provide enough tokens for their execution in states $\omega_0 - \omega_{j-1}$ (e.g., 0 tokens for a_3 in $\omega_0 - \omega_5$); these actors cannot be enabled due to the lack of initial tokens in the newly added channels in the corresponding decision state (e.g., there will be no token in $c_{1.a_3\omega_6}$ in ω_6). Hence, only the actor of choice amongst the opponent actors of a decision state will be enabled in that state (e.g., only a_1 can fire in ω_6). The delayed actors in a decision state (e.g., ω_j) will have sufficient tokens on the incoming channel from the newly added actor for that decision state (i.e., $a_{i.\omega_j}$) after firing of the actor of choice in ω_j . For example, there will be 6 tokens in channel $c_{1.a_1\omega_6}$ after the firing of a_1 (i.e., the actor of choice) in decision state ω_6 ; hence, the actor $a_{1.\omega_6}$ immediately fires and then provides sufficient tokens for later firings of a_3 . So, the delayed actor in decision state ω_6 (i.e., a_3) will no longer be blocked due to the absence of tokens in channel $c_{1.a_3\omega_6}$ after the decision state ω_6 . The firing of actor a_1 after decision state ω_6 produces 1 token in channel $c_{1.a_1\omega_6}$ and the firings of actor a_3 after decision state ω_6 consumes 6 tokens from channel $c_{1.a_3\omega_6}$; as a result, the number of tokens in the new channels are reset to the initial values at the end of one iteration of the schedule s_1 . Hence, the periodic behavior is also achievable for the added components.

DSM also adds actor $a_{1.\omega_9}$ and channels $c_{1.a_1\omega_9}$ and $c_{1.a_3\omega_9}$




Fig. 9. An example SDFG.

Algorithm 3: SAS Modeling (SASM)

```

input : SDFG  $G(A, C)$ , PSOS  $s_i = \{(\alpha_1)^{\beta_1}(\alpha_2)^{\beta_2} \dots (\alpha_n)^{\beta_n}\}$ 
output:  $G$  extended with schedule  $s_i$ 

1 add a self edge with 1 initial token for each  $a \in A_i$ 
2 for  $i \leftarrow 1$  to  $n$  do
3   if  $\alpha_i$  is not an actor then
4     SASM( $G, \alpha_i$ )
5   AA( $G, a_{cnt_i}$ )
6     /* adding a counter channel */
7     AC( $G, c_{cnt_i}, \text{rightMost}(\alpha_i), a_{cnt_i}, 1,$ 
8       RN(rightMost( $\alpha_i$ ),  $\alpha_i^{\beta_i}$ ), 0)
9     /* adding a limiter channel */
10    if  $i = n$  then
11      AC( $G, c_{lim_i}, a_{cnt_i}, \text{leftMost}(\alpha_1), \text{RN}(\text{leftMost}(\alpha_1),$ 
12         $\alpha_1^{\beta_1}), 1, \text{RN}(\text{leftMost}(\alpha_1), \alpha_1^{\beta_1}))$ 
13    else
14      AC( $G, c_{lim_i}, a_{cnt_i}, \text{leftMost}(\alpha_{n+1}),$ 
15        RN(leftMost( $\alpha_{n+1}$ ),  $\alpha_{n+1}^{\beta_{n+1}}), 1, 0)$ 

```

to the graph for the other decision state ω_9 . The components added in decision state ω_9 show similar behavior as the components added in ω_6 .

VI. MODELING SINGLE APPEARANCE SCHEDULES

A well-known class of scheduling techniques are single appearance schedules (SASs) in which each actor appears exactly once in the LS form. This aspect makes SASs suitable to minimize code memory size. $s_2 = \langle (a_0 a_1)^5 a_2 \rangle^*$ is a PSOS and SAS for part of the SDFG (i.e., actors a_0 - a_2) in Figure 9.

DSM is able to model any PSOS. However, more compact graphs are possible for SASs. Algorithm 3 presents our single appearance schedule modeling (SASM) technique. Similar to DSM, SASM adds some extra actors and channels to the original SDFG to model the given schedules. The original channels and actors in the schedule-extended SDFG are preserved and distinguishable from the newly added elements by any of our techniques. Hence, both our techniques preserve the original structure of an SDFG. This property can be beneficial when a resource allocation algorithm needs to be applied on the schedule-extended graph; a resource allocation algorithm can easily distinguish an original actor (or channel) from an actor (or channel) which is added to model the schedules.

We know that each actor appears only once in a SAS; this property can help us to model a SAS in an SDFG in a smarter way than DSM does. An actor (or a nested inner LS) should be executed a specific number of times before another actor (or another nested inner LS) starts executing. In $s_2 = \langle (a_0 a_1)^5 a_2 \rangle^*$, the nested inner LS $(a_0 a_1)$ must be executed 5 times before a_2 starts firing. This type of execution control can be handled using a counter construct. The key idea of SASM is to implement a counter concept in the graph.

Later, we explain how we implement such counters to model SASs in an SDFG. Similar to DSM, auto-concurrency can be eliminated by adding a self-edge with 1 initial token to each actor in the SDFG (see line 1 in Algorithm 3). The rest of Algorithm 3 deals with implementing the counter concept.

Figure 10 shows the graph of the SDFG in Figure 9 which models the schedule s_2 using SASM. Schedule s_2 has a nested $a_0 a_1$; this can be replaced with α_{01} to form a looped schedule representation (i.e. $s_2 = \langle (\alpha_{01})^5 a_2 \rangle^*$ where $\alpha_{01} = a_0 a_1$). A counter in SASM is implemented by one actor a_{cnt_i} (e.g., a_{cnt_3} in Figure 10), a *counter channel* c_{cnt_i} (e.g., c_{cnt_3}) and a *limiter channel* c_{lim_i} (e.g., c_{lim_3}). A counter in SASM has two properties: first, it counts the number of times that element α_i (e.g., α_{01} above) is being executed; second, it limits the element α_{i+1} (e.g., a_2) to be executed to β_{i+1} times (e.g., 1 as the number of repetitions of a_2 is one). The counter channel c_{cnt_i} (e.g. c_{cnt_3}) is placed from the rightmost actor in α_i (e.g., actor a_1 in α_{01}) to the actor a_{cnt_i} (e.g., a_{cnt_3}); the production rate of the rightmost actor in α_i on c_{cnt_i} is set to 1 and the consumption rate of a_{cnt_i} on c_{cnt_i} is set to the number of times that the rightmost actor in α_i can be executed in $\alpha_i^{\beta_i}$ (e.g., 5 as the number of times that a_1 can be executed in α_{01}^5 is five). In Algorithm 3, the function *rightMost*(α_i) (*leftMost*(α_i)) returns the rightmost (leftmost) actor in element α_i (e.g. *rightMost*($(a_0 a_1)^5$) returns actor a_1). The function *RN*($a, \alpha_i^{\beta_i}$) retrieves the count of a in element $\alpha_i^{\beta_i}$ (e.g. *RN*($a_0, (a_0 a_1)^5$) returns 5).

The limiter channel c_{lim_i} (e.g., c_{lim_3}) is placed from a_{cnt_i} (e.g., a_{cnt_3}) to the leftmost actor in the next element, i.e., α_{i+1} (e.g., a_2). SASM performs the placement of the counter constructs in a circular way. In other words, the next element of α_j is considered to be $\alpha_{(j+1) \bmod n}$ where $j \in \mathbb{N} \wedge j \leq n$ for a LS $\{(\alpha_1)^{\beta_1}(\alpha_2)^{\beta_2} \dots (\alpha_n)^{\beta_n}\}$. The production rate of a_{cnt_i} on c_{lim_i} is set to the number of times that the leftmost actor in α_{i+1} can be executed in $\alpha_{i+1}^{\beta_{i+1}}$ (e.g., 1 as the number of times that a_2 can be executed in element a_2) and the consumption rate of the leftmost actor in the next element, i.e., α_{i+1} , from c_{lim_i} is set to 1. So, element α_{i+1} depends on actor a_{cnt_i} (because of c_{lim_i}) and actor a_{cnt_i} depends on α_i (because of the c_{cnt_i}); the β_i executions of α_i produce enough tokens on the counter channel c_{cnt_i} and then actor a_{cnt_i} can be fired. The firing of actor a_{cnt_i} provides enough tokens on limiter channel c_{lim_i} to only allow β_{i+1} executions for the next element α_{i+1} . Hence, by adding these components we enforce that α_{i+1} can be executed β_{i+1} times after α_i is executed β_i times.

The limiter channel of the counter construct added for the last element (i.e., $(\alpha_n)^{\beta_n}$) in a LS $\{(\alpha_1)^{\beta_1}(\alpha_2)^{\beta_2} \dots (\alpha_n)^{\beta_n}\}$ is initialized with some initial tokens to prevent a deadlock in the graph. The number of initial tokens on the limiter channel is set to the count of the left most actor of α_1 in the first element of that LS (i.e., $(\alpha_1)^{\beta_1}$). Line 8 in SASM performs the token initialization. Inter-iteration execution cannot happen because SASM always creates a dependency from the last actor in the schedule to the first actor in the schedule.

In Figure 10, the actor a_{cnt_3} is added to count the number of times that the sequence $a_0 a_1$ is executed. The consumption rate of actor a_{cnt_3} on its input channel is 5; this means that




 Fig. 10. SDFG of Figure 9 extended with $s_2 = \langle (a_0 a_1)^5 a_2 \rangle^*$ using SASM.


Fig. 11. Optimization of SASM.

after 5 executions of sequence $a_0 a_1$ the next actor (i.e., a_2) can be enabled. Also, the actor a_{cnt_3} limits the number of times that actor a_2 should get enabled; this can be done by choosing the value 1 as production rate of actor a_{cnt_3} on its output channel. In other words, the actor a_2 can only fire once because of the limitation imposed by actor a_{cnt_3} . The actor a_{cnt_1} (a_{cnt_2}) is added to ensure the single execution of actors a_1 (a_0) after the single execution of actor a_0 (a_1). The actor a_{cnt_4} is added to ensure that the sequence $a_0 a_1$ can be executed 5 times after the actor a_2 is executed once.

SASM is applied recursively (line 4 in SASM) to model the nested LS α_i . For example, $SASM(G, a_0 a_1)$ will be called inside $SASM(G, (a_0 a_1)^5 a_2)$; the result of the recursive call is shown in Figure 10 with a rectangle marking the inner-loop.


Some of the elements added by SASM can be removed without affecting the outcome. Consider Figure 11(a) which contains a counter actor and two channels that can be discarded in the following cases:

- The counter actor a_{cnt_i} can be removed if rate p is equal to 1. The counter actor and two channels in the original form are replaced with channel c_{xy} (see Figure 11(b)).
- The counter actor a_{cnt_i} can be removed if rate q is equal to 1. The counter actor and two channels in the original form are replaced with channel c_{xy} (see Figure 11(c)).

The newly replaced channel c_{xy} is only necessary if there is no other equivalent channel in the original SDFG. The channels (p, q) and (p', q') which have the same source actor $a_x \in A$ and sink actor $a_y \in A$ are equivalent if the equation $\frac{Rate(p)}{Rate(p')} = \frac{Rate(q)}{Rate(q')} = \frac{\omega_0(c)}{\omega_0(c')}$ is true. Applying these optimizations on Figure 10 replaces all components added by SASM by channel c_{01} (see Figure 12). The SDFG which models schedule s_2 in the SDFG of Figure 9 with DSM and the HSDFG-based techniques result in a graph with 10 (26) and 13 (21) actors (channels) resp. The SDFG which models the same schedule with SASM only has 5 (9) actors (channels).

VII. CORRECTNESS OF THE PROPOSED TECHNIQUES

This section presents the theorems related to the correctness of our proposed techniques. A schedule is encoded correctly


 Fig. 12. SDFG of Figure 9 extended with $s_2 = \langle (a_0 a_1)^5 a_2 \rangle^*$ using optimized SASM.

if any execution of a schedule-extended graph satisfies the encoded schedule and if any execution of the original graph that satisfies the given schedule is still possible in the schedule-extended graph. The proofs of the theorems are given in the appendices of this report.

Assume that the actors and channels added by DSM (or SASM) to model schedule s_i in SDFG $G(A, C)$ are denoted by A_{s_i} and C_{s_i} , resp. $G'(A', C')$ is the SDFG that models s_i in G using DSM (or SASM) where $A' = A \cup A_{s_i}$ and $C' = C \cup C_{s_i}$.

Theorem 1 and Theorem 2 state the correctness of DSM in modeling a single PSOS for a sub-set of the actors of the SDFG. Applying the theorems once for each schedule to be encoded shows the correctness of Algorithm 1.

Note that actors and channels added to model a given schedule do not affect actors that are not part of the schedule. Theorem 1 shows that any execution of the schedule-extended graph satisfies the modeled schedule. Firings of the added actors need to be ignored, which is achieved by the stated condition on the ordered lists resulting from the executions in the schedule-extended and original graphs.

Theorem 1. Consider PSOS s_i as a schedule for actors $A_i \subseteq A$ from SDFG $G(A, C)$. Assume $G'(A', C')$ is the SDFG that models s_i in G using DSM. For any execution σ' of $G'(A', C')$ it holds that σ satisfies s_i where it is assumed that σ is the execution of $G(A, C)$ with $orderList(\sigma, A) = orderList(\sigma', A)$.

Theorem 2 shows that no execution of the original graph is unnecessarily excluded in the schedule-extended graph. In other words, any execution of the original graph that satisfies a given schedule is still feasible in the schedule-extended graph.

Theorem 2. Consider PSOS s_i as a schedule for actors $A_i \subseteq A$ from SDFG $G(A, C)$. Assume $G'(A', C')$ is the SDFG that models s_i in G using DSM. For any execution σ of G that satisfies s_i it holds that there is exactly one σ' that is an execution of $G'(A', C')$ such that $orderList(\sigma, A) = orderList(\sigma', A)$.

The following two theorems state the correctness of SASM.

They are very similar to the two theorems for DSM.

Theorem 3. Consider SAS $s_i = \{(\alpha_1)^{\beta_1}(\alpha_2)^{\beta_2} \dots (\alpha_n)^{\beta_n}\}$ as a schedule for actors $A_i \subseteq A$ from SDFG $G(A, C)$. Assume $G'(A', C')$ is the SDFG that models s_i in G using SASM. For any execution σ' of $G'(A', C')$ it holds that σ satisfies s_i where it is assumed that σ is the execution of $G(A, C)$ with $\text{orderList}(\sigma, A) = \text{orderList}(\sigma', A)$.

Theorem 4. Consider PSOS s_i as a schedule for actors $A_i \subseteq A$ from SDFG $G(A, C)$. Assume $G'(A', C')$ is the SDFG that models s_i in G using DSM. For any execution σ of G that satisfies s_i it holds that there is exactly one σ' that is an execution of $G'(A', C')$ such that $\text{orderList}(\sigma, A) = \text{orderList}(\sigma', A)$.

VIII. EXPERIMENTAL RESULTS

In this section, we evaluate our techniques experimentally. We first explain the experimental setup. We then evaluate our techniques in terms of the sizes of the schedule-extended graphs, comparing our techniques to that of [12]. We further consider the throughput analysis time when analyzing the schedule-extended graphs obtained by different techniques. Note that the throughput that is achievable for a given schedule is independent of the way it is encoded. It is the analysis time itself that is of interest. Finally, we look at the accuracy of buffer sizing analysis. The accuracy of obtained buffer requirements does depend on the way schedules are encoded.

A. Experimental setup

The DSM and SASM techniques have been integrated in the SDF³ [28] dataflow tool set, available at <http://www.es.ele.tue.nl/sdf3>. We use a set of DSP and multimedia applications (see the first column of Table I) to assess our DSM and SASM techniques.

In our experiments, applications are bound to a multi-processor platform using the technique of [6]. A PSOS determines the actor firing order and as such it influences the enabled actors in a state; as a result, the number of decision states can be different for different PSOSs. The size of the schedule-extended graph using DSM depends on the number of decision states in the given schedules. We use a list scheduling approach from [29] to determine PSOSs for the applications. We use two different variations to verify DSM in different situations. The first list schedule uses forward priorities (Lfp) and the second one uses reverse priorities (Lrp). Actors closer to the inputs of the graph have higher priority in the Lfp schedules compared to actors closer to the outputs of the graph and vice-versa in Lrp schedules.

The scheduling technique presented in [30] is used to derive SASs for our benchmark applications. The technique in [30] also minimizes the required buffer sizes when determining a SAS. However, the technique in [30] cannot directly be used for multi-processors. We have utilized the technique of [30] to find SASs for a multi-processor platform. Initially the binding technique from [6] is used to bind the SDFG to a multi-processor platform. Then, the technique of [30] is applied to the SDFG to derive a SAS for all actors in

the SDFG. This SAS is decomposed into some smaller SASs using the binding information; each of the smaller SASs is a schedule for one processor in the platform. Consider an example SDFG with 5 actors denoted by $a_0 - a_4$. Assume a_0, a_1 and a_3 are bound to processor P_1 and a_2 and a_4 are bound to processor P_2 . Applying the technique of [30] to this imaginary SDFG gives the SAS $s_0 = \langle (a_0(a_1^2 a_3^4)^3)^2 a_4^5 \rangle^*$ for the whole SDFG. This SAS can be decomposed using the binding information to form a SAS for each of the processor in the platform. Only considering actors bound to P_1 in s_0 results in $s_{01} = \langle a_0(a_1^2 a_3^4)^3 \rangle^*$ which is a SAS for the actors bound to P_1 . Similarly, a SAS $s_{02} = \langle a_2^6 a_4^5 \rangle^*$ can be extracted from s_0 to order actors bound to P_2 . This way we utilize the technique of [30] for multi-processor platforms. The optimality of the generated schedules from the performance or buffer sizing perspective is debatable. However, we use this adapted SAS technique merely to provide some near-optimal inputs to evaluate our SASM schedule encoding technique versus the existing technique. Our techniques do not affect the quality of the scheduling result itself.

B. Comparison on graph sizes

Table I contains the size of the schedule-extended graphs using the HSDFG-based and DSM techniques to model Lfp and Lrp schedules for a single core platform (see the first two rows of each application line). Using schedules generated by Lfp, the number of decision states is less than when Lrp is used, except in the channel equalizer and MP3 playback applications. By using Lfp scheduling, actors closer to inputs have higher priority compared to actors closer to outputs. This leads to consecutive execution of an actor followed by consecutive execution of another actor with lower priority and so on. Thanks to our optimization in DSM, considering only one decision state before a context switch will be sufficient (e.g., decision state ω_9 in Figure 7) and the number of decision states can be reduced significantly. Usually SDFG actors closer to outputs are dependent on actors closer to inputs; this dependency can prevent an actor from being executed consecutively in a graph scheduled by Lrp. As a result, the number of context switches in a graph scheduled by Lrp will typically be larger compared to Lfp. Hence, the effectiveness of the decision state optimization in DSM decreases and extra elements are required to model the schedules in the graph. The exceptions in the channel equalizer and MP3 playback are due to the existence of a cycle in the SDFG; the cycle can increase the number of context switches in the schedule and as a result, Lfp could result in the same or a higher number of decision states in DSM compared to Lrp. However, DSM always outperforms the HSDFG-based technique regardless of the input schedule in our experiments. The number of actors (channels) using DSM is 66% (71%) lower compared to the HSDFG-based technique on average and 22% (17%) lower in the worst-case observed in our experiments. The average case refers to the mean value of the obtained results and the worst-case reports the smallest graph size reduction (i.e., reduction in numbers of actors and channels compared to the HSDFG-based technique).

TABLE I
 SIZE OF THE SCHEDULE-EXTENDED SDFGS

Benchmark	Orig. size		# actors (# channels)			Reduction compared to the HSDFG-based technique		
	# actors (# channels)	Schedule type	HSDFG-based	DSM	SASM	DSM	SASM	
H.263 dec. [10]	4 (6)	Lfp	1190 (2973)	6 (14)	NA	99% (99%)	NA	
		Lrp	1190 (2972)	598 (1198)	NA	49% (59%)	NA	
		SAS	1190 (2972)	598 (1198)	4 (12)	49% (59%)	99% (99%)	
H.263 enc. [24]	5 (7)	Lfp	201 (596)	7 (16)	NA	96% (97%)	NA	
		Lrp	201 (499)	105 (212)	NA	47% (57%)	NA	
		SAS	201 (499)	106 (214)	5(15)	47% (57%)	97% (97%)	
MP3 dec. [10]	14 (21)	Lfp	911 (2849)	27 (61)	NA	97% (97%)	NA	
		Lrp	911 (2327)	400 (807)	NA	56% (65%)	NA	
		SAS	911 (2327)	400 (807)	14 (44)	56% (65%)	98% (98%)	
modem [1]	16 (35)	Lfp	48 (115)	22 (63)	NA	54% (45%)	NA	
		Lrp	48 (128)	35 (89)	NA	27% (30%)	NA	
		SAS	48 (128)	35 (89)	16 (61)	27% (30%)	66% (52%)	
samplerate conv. [1]	6 (11)	Lfp	612 (1639)	12 (29)	NA	98% (98%)	NA	
		Lrp	612 (1784)	157 (319)	NA	74% (82%)	NA	
		SAS	612 (1865)	238 (481)	12 (31)	61% (74%)	98% (98%)	
satellite rec. [25]	22 (48)	Lfp	4515 (11638)	41 (108)	NA	99% (99%)	NA	
		Lrp	4515 (12820)	1223 (2472)	NA	72% (80%)	NA	
		SAS	4515 (15270)	3673 (7372)	24(91)	18% (51%)	99% (99%)	
MP3 playback [26]	4 (8)	Lfp	10601 (37531)	5298 (10600)	NA	50% (71%)	NA	
		Lrp	10601 (37529)	5296 (10596)	NA	50% (71%)	NA	
		SAS	10601 (37530)	5297 (10598)	6 (16)	50% (71%)	99% (99%)	
bipartite [25]	4 (8)	Lfp	73 (341)	8 (20)	NA	89% (94%)	NA	
		Lrp	73 (359)	26 (56)	NA	64% (84%)	NA	
		SAS	73 (350)	17(38)	9(22)	76% (89%)	87% (93%)	
channel eq. [27]	21 (40)	Lfp	41 (100)	32 (83)	NA	22% (17%)	NA	
		Lrp	41 (95)	27 (73)	NA	34% (23%)	NA	
		SAS	41 (93)	30 (79)	21(65)	27% (15%)	48% (30%)	

 TABLE II
 THE THROUGHPUT ANALYSIS TIME (IN MILLISECONDS)

Benchmark	# of processors	Type of the schedules / The used schedule modeling technique						
		SAS			Lfp		Lrp	
		HSDFG-based	DSM	SASM	HSDFG-based	DSM	HSDFG-based	DSM
H.263 dec.	2	36 540	< 1	< 1	118 720	< 1	120 710	< 1
H.263 enc.	2	380	< 1	< 1	690	< 1	400	< 1
MP3 dec.	3	13 900	320	10	17 660	< 1	13 980	1 380
modem	3	10	< 1	< 1	< 1	< 1	< 1	< 1
samplerate conv.	3	3 970	110	< 1	3 140	< 1	3 880	< 1
satellite rec.	2	not finished in 3 days	12 414 400	130	not finished in 3 days	280	not finished in 3 days	675 700
MP3 playback	2	not finished in 3 days	< 1	< 1	not finished in 3 days	10 780	not finished in 3 days	< 1
bipartite	2	30	< 1	< 1	40	< 1	50	< 1
channel eq.	3	< 1	< 1	< 1	< 1	< 1	< 1	< 1

SASs are a suitable class of schedules that minimize code memory size. DSM is able to model any arbitrary schedule in an SDFG. SAS can be modeled using DSM; however, it is possible to consider the intrinsic property of SASs when modeling a SAS in an SDFG. Our second technique, SASM, uses the fact that each actor appears only once in the looped schedule form. SASM models the counter concept in the graph in order to force actors to be executed a specific number of times. The third row of each application line in Table I contains the size of the schedule-extended graphs using the HSDFG-based, DSM and SASM techniques to model SASs, generated by the technique developed in [30].

SASM results in a schedule-extended SDFG with a limited number of extra actors and channels. For example, SASM only adds 2 (8) extra actors (channels) to the original graph of the MP3 playback application in order to model a SAS, while the HSDFG-based technique adds 1057 (37522) extra actors (channels) to model the same schedule. The graphs obtained by SASM have 88% (85%) and 48% (30%) less actors (channels) compared to the HSDFG-based technique on average and in the worst-case among the benchmark applications. Using the DSM technique to model the same SASs results in 46% (57%) and 27% (15%) less actors (channels) compared to the

HSDFG-based technique on average and in the worst-case. Our results confirm that the techniques proposed in this paper achieve a more compact schedule-extended graph compared to the available technique.

C. Comparison on analysis times

The time required to perform an analysis on an SDFG depends on the size of the graph and the number of cycles in the graph. As an example, the throughput analysis of [9] is performed on the schedule-extended graphs using our techniques and the HSDFG-based technique. Our experiments are performed to evaluate the impact of the graph size on the analysis time of a common analysis technique. Note that other techniques (e.g., YTO [31]) can be employed to calculate throughput of HSDFGs, but the size of the schedule-extended HSDFGs is such that our conclusions remain the same. The benchmark graphs are mapped onto multi-processor platforms with two or three processors. Table II contains the throughput analysis times when SASs, list forward priority (Lfp) schedules and list reverse priority (Lrp) schedules are used as input schedules. The results show the superiority of SASM over DSM and the HSDFG-based technique. Note that the SDFG




Fig. 13. SDFG of H.263 decoder application.

to HSDFG conversion is fast; the numbers reported for the HSDFG-based technique in Table II are related to the run-time of the throughput analysis from [9]. In our experiment, the run-time of a throughput calculation for HSDFGs is long independent of the analysis technique used (i.e., state-space [9], YTO [31], etc.).

D. Comparison on buffer sizes

To further analyze the effectiveness of our techniques, the buffer sizing algorithm from [10] is applied to the schedule-extended SDFGs of the H.263 decoder and MP3 decoder. Figure 13 depicts the SDFG of the H.263 decoder. Besides the compactness of the schedule-extended graph, our techniques preserve the original structure of an SDFG (when ignoring the added actors and channels), allowing accurate buffer sizing, which is not guaranteed for the state of the art technique. The H.263 decoder is mapped to a platform with two processors. The actors *vld* and *iq* are mapped to the first processor with a PSOS $\langle vld(iq)^{594} \rangle^*$ and the actor *idct* and *mc* are mapped to the second processor with a PSOS $\langle (idct)^{594}mc \rangle^*$. The analysis time for buffer sizing on the schedule-extended H.263 decoder is less than 1 ms when using DSM (or SASM) to model the schedules. The same analysis takes 1330 ms when using the technique from [12] to model the same schedules in the same graph. Figure 14(a) shows the Pareto space of throughput and buffer size when modeling the schedules with DSM (or SASM) and the HSDFG-based technique [12]. In this experiment, the schedules are first modeled in the graph; then, the buffer sizing technique of [10] is applied. A single channel in an SDFG corresponds to a set of channels in the equivalent HSDFG. As a result, the buffer sizing technique cannot find the minimal buffer size when applying it on the equivalent HSDFG. Our experiments show these inaccuracies. Applying buffer sizing on the graph which models the schedules using the technique from [12] results in 49% overestimation in required buffers compared to applying the same buffer sizing technique on the graph which models the schedules with one of our techniques. Note that the maximal achievable throughput is independent of the way schedules are encoded. The analysis results confirm this. Only the computed buffer sizes differ. For instance in both cases of Figure 14, the maximal throughput for the given schedules is always achievable, also by using the HSDFG-based schedule modeling technique; the latter suggests the need for larger buffers though. Figure 14(b) shows results for the MP3 decoder. We use the mapping and scheduling from [32] for a platform with 3 processors. The analysis time on the graph which models the schedule using one of our techniques is 594 ms while 141610 ms is required to perform the same analysis on the graph using the technique from [12]. Using the technique from [12] results




Fig. 14. Pareto space of schedule-extended graphs modeled by DSM and HSDFG-based techniques (the scales of the two graphs are different).

in 226% overestimation in buffer size compared to using our techniques.

Modeling a PSOS in an SDFG using DSM requires execution of one complete SDFG iteration. The number of states in one iteration could be exponential in the number of actors in the graph. However, for all real-world SDFGs used in our experiments, the execution time of DSM is below 1 ms. SASM also models SASs based on the structure of the schedules in their looped form; as each actor appears once in a SAS, the complexity of SASM depends on the number of actors in the graph. Similar to DSM, the execution time of SASM is always below 1 ms in our experiments. The complexity of our techniques relates to the length of the SDFG iteration and the number of processors in the platform (i.e., $|P|$). Hence, the complexity of our techniques is bounded to $O(|P| \cdot \sum_{a \in A} \gamma(a))$.

IX. CONCLUSION

We present two techniques, DSM and SASM, to model PSOSs and SASs directly in an SDFG. The resulting graphs are much smaller (often much less than half the size) than graphs resulting from the state of the art technique that first converts an SDFG to an HSDFG. This results in a speed-up of analysis techniques. Computing the trade-off between buffering and throughput for multi-processor platforms, for example, becomes several orders of magnitude faster. Moreover, properties like buffer sizes can be analyzed more accurately. The techniques have been integrated in the SDF³ tool set available at <http://www.es.ele.tue.nl/sdf3>. This allows easy integration of the techniques in multi-processor design flows.

ACKNOWLEDGMENT

We thank the reviewers for their constructive comments, which helped improving the presentation of the paper. This work was supported in part by the Dutch Technology Foundation STW, project NEST 10346.

REFERENCES

- [1] S. S. Bhattacharyya *et al.*, "Synthesis of embedded software from synchronous dataflow specifications," *Journal of VLSI Signal Processing*, vol. 21, pp. 151–166, 1999.
- [2] S. Sriram *et al.*, *Embedded Multiprocessors: Scheduling and Synchronization*, 2nd ed. CRC Press, 2009.
- [3] P. Poplavko *et al.*, "Task-level timing models for guaranteed performance in multiprocessor networks-on-chip," CASES. ACM, 2003, pp. 63–72.
- [4] M.-Y. Ko *et al.*, "Compact procedural implementation in DSP software synthesis through recursive graph decomposition," SCOPES. ACM, 2004, pp. 47–61.
- [5] A. Bonfietti *et al.*, "Throughput constraint for synchronous data flow graphs," CPAIOR. Springer-Verlag, 2009, pp. 26–40.

- [6] S. Stuijk *et al.*, "Multiprocessor resource allocation for throughput-constrained synchronous dataflow graphs," in *DAC*, 2007, pp. 777–782.
- [7] W. Liu *et al.*, "Efficient SAT-based mapping and scheduling of homogeneous synchronous dataflow graphs for throughput optimization," *RTSS*. IEEE, 2008, pp. 492–504.
- [8] Y. Yang *et al.*, "Automated bottleneck-driven design-space exploration of media processing systems," *DATE*. ACM, 2010, pp. 1041–1046.
- [9] A. Ghamarian *et al.*, "Throughput analysis of synchronous data flow graphs," *ACSD*. IEEE, 2006, pp. 25–36.
- [10] S. Stuijk *et al.*, "Throughput-buffering trade-off exploration for cyclo-static and synchronous dataflow graphs," *IEEE Transaction on Computers*, vol. 57, no. 10, pp. 1331–1345, 2008.
- [11] M. Benazouz *et al.*, "A new method for minimizing buffer sizes for cyclo-static dataflow graphs," in *ESTIMedia'10*. IEEE, pp. 11–20.
- [12] N. Bambha *et al.*, "Intermediate representations for design automation of multiprocessor DSP systems," *Design Automation for Embedded Systems*, vol. 7, no. 4, pp. 307–323, 2002.
- [13] M. Damavandpeyma *et al.*, "Modeling static-order schedules in synchronous dataflow graphs," in *DATE'12*. ACM, pp. 775–780.
- [14] A.-P. Wang *et al.*, "Buffer optimization and dispatching scheme for embedded systems with behavioral transparency," *ACM Transactions on Design Automation of Electronic Systems*, vol. 17, no. 4, pp. 41:1–41:26, Oct. 2012.
- [15] M. Damavandpeyma *et al.*, "Throughput-constrained DVFS for scenario-aware dataflow graphs," in *RTAS'13*. IEEE.
- [16] M. H. Wiggers *et al.*, "Monotonicity and run-time scheduling," *EMSOFT*. ACM, 2009, pp. 177–186.
- [17] D. Thiele *et al.*, "Optimizing performance analysis for synchronous dataflow graphs with shared resources," in *DATE'12*. ACM, pp. 635–640.
- [18] R. Henia *et al.*, "System level performance analysis - the SymTA/S approach," *IEE Proceedings Computers and Digital Techniques*, vol. 152, no. 2, pp. 148–166, mar 2005.
- [19] L. Thiele *et al.*, "Real-time calculus for scheduling hard real-time systems," in *ISCAS'00*, vol. 4. IEEE, 2000, pp. 101–104.
- [20] H. H. Wu *et al.*, "A model-based schedule representation for heterogeneous mapping of dataflow graphs," *HCW*. IEEE, 2011, pp. 66–77.
- [21] S. Bhattacharyya *et al.*, *Software Synthesis from Dataflow Graphs*. Kluwer Academic Publishers, 1996.
- [22] E. Lee *et al.*, "Synchronous data flow," *Proceeding of the IEEE*, vol. 75, no. 9, pp. 1235–1245, 1987.
- [23] M. Geilen *et al.*, "Minimising buffer requirements of synchronous dataflow graphs with model checking," *DAC '05*. ACM, 2005, pp. 819–824.
- [24] H. Oh *et al.*, "Fractional rate dataflow model for efficient code synthesis," *Journal of VLSI Signal Processing*, vol. 37, pp. 41–51, 2004.
- [25] S. Ritz *et al.*, "Scheduling for optimum data memory compaction in block diagram oriented software synthesis," *ICASSP*. IEEE, 1995.
- [26] M. H. Wiggers *et al.*, "Efficient computation of buffer capacities for cyclo-static dataflow graphs," *DAC*. ACM, 2007, pp. 658–663.
- [27] A. Moonen *et al.*, "Practical and accurate throughput analysis with the cyclo static dataflow model," *MASCOTS*. IEEE, 2007, pp. 238–245.
- [28] S. Stuijk *et al.*, "SDF³: SDF for free," *ACSD*. IEEE, 2006, pp. 276–278.
- [29] G. De Micheli, *Synthesis and Optimization of Digital Circuits*, 1st ed. McGraw-Hill, 1994.
- [30] P. K. Murthy *et al.*, "Joint minimization of code and data for synchronous dataflow programs," *Formal Methods in System Design*, vol. 11, no. 1, pp. 41–70, 1997.
- [31] N. E. Young *et al.*, "Faster parametric shortest path and minimum balance algorithms," *CoRR*, vol. cs.DS/0205041, 2002.
- [32] M. Geilen *et al.*, "Worst-case performance analysis of synchronous dataflow scenarios," *CODES+ISSS*. ACM, 2010, pp. 125–134.



Morteza Damavandpeyma received his B.Sc. and M.Sc. degrees in computer engineering from the University of Tehran, Iran in 2006 and 2008, respectively. He is currently a Ph.D. candidate in the Department of Electrical Engineering at the Eindhoven University of Technology (TU/e). His research interests include modeling, analysis and synthesis of embedded systems and multiprocessor System-on-Chips.



Sander Stuijk received his M.Sc. degree (with honors) in Electrical Engineering in 2002 and his Ph.D. degree in 2007 from the Eindhoven University of Technology. He is currently an assistant professor in the Department of Electrical Engineering at the Eindhoven University of Technology. His research interests include modeling methods and mapping techniques for the design, specification, analysis and synthesis of predictable hardware/software systems.



Twan Basten is a professor in the Electrical Engineering department at Eindhoven University of Technology (TU/e), the Netherlands, where he chairs the Electronic Systems group. He is also a senior research fellow of TNO Embedded Systems Innovation in the Netherlands. He holds an M.Sc. and a Ph.D. in Computing Science from TU/e. His research interests include embedded and cyber-physical systems, dependable computing and computational models.



Marc Geilen is an assistant professor in the Department of Electrical Engineering at Eindhoven University of Technology. He holds an M.Sc. in Information Technology and a Ph.D. from the Eindhoven University of Technology. His research interests include modeling, simulation and programming of multimedia systems, multiprocessor systems-on-chip, networked embedded systems and cyber-physical systems, and multi-objective optimization and trade-off analysis.



Henk Corporaal has gained an M.Sc. in Theoretical Physics from the University of Groningen, and a Ph.D. in Electrical Engineering, in the area of Computer Architecture, from Delft University of Technology. He is a professor in Embedded System Architectures at the Eindhoven University of Technology (TU/e) in The Netherlands. His current research projects are on low power single and multiprocessor architectures, their programmability, and the predictable design of soft- and hard real-time systems.

APPENDIX A
CORRECTNESS OF DSM

This section discusses the correctness of DSM in modeling a single PSOS for a sub-set of the actors of the SDFG. The extra actors and channels added by our techniques to model a PSOS (e.g., s_i) are only placed between actors of that schedule (i.e., A_i) and this is not imposing any restriction on the other actors of the SDFG (i.e., actors in $A \setminus A_i$). Hence, if we can model a single PSOS in the SDFG, then we can model multiple PSOSs by applying the algorithm multiple times.

DSM adds some actors, denoted by A_{s_i} , and channels, denoted by C_{s_i} to model the PSOS s_i in SDFG $G(A, C)$. $G'(A', C')$ is the SDFG that models s_i in G using DSM where $A' = A \cup A_{s_i}$ and $C' = C \cup C_{s_i}$.

The existence of a repetition vector for an SDFG $G(A, C)$ ensures a balance between production and consumption rates. Hence, a *balance equation* can be defined as follow for each channel $c = (p, q) \in C$ where p and q are ports of actors a_p and a_q respectively: $Rate(p) \cdot \gamma(a_p) = Rate(q) \cdot \gamma(a_q)$. The existence of a repetition vector γ for an SDFG ensures that each balance equation related to a channel in the SDFG holds; under this situation the SDFG is consistent. The following proposition shows the consistency of the schedule-extended SDFG G' .

Proposition 1. *The SDFG $G'(A', C')$ which models PSOS s_i in the consistent SDFG $G(A, C)$ is consistent.*

Proof. The SDFG G is consistent. Hence, a non-trivial repetition vector γ exists for G . We need to show that a non-trivial repetition vector γ' exists for G' . The balance equation related to each self-loop added by DSM (in line 1 of Algorithm 1) to remove auto-concurrency is always valid because the source and destination actor of the self-loop channel are identical with production and consumption rate equal to 1. The rates of the other channels added by DSM (for decision states or inter-iteration execution) share the following properties: (1) the newly added channel $c \in C_{s_i}$, which is added by DSM (in all lines 7, 8, 16 or 18 of Algorithm 1), is between an actor $a_p \in A (= A' \setminus A_{s_i})$ and an actor $a_q \in A_{s_i}$; (2) the rate of the new channel on the side of the actor a_q is equal to the count of the actor a_p in one iteration of s_i (i.e., $CNT(a_p, s_i)$); (3) the rate of the new channel on the side of the actor a_p is equal to 1. From these properties we conclude that $a_q \in A_{s_i}$, which is added by DSM (in both line 6 and 13 of Algorithm 1), fires only once in each iteration of s_i . Consider a_q as the producer (see Figure 15(a)) of the newly added channel (i.e., $(a_q, a_p) \in C_{s_i}$); then, the only firing of a_q in one iteration of s_i provides $CNT(a_p, s_i)$ tokens for all firings of a_p in one iteration of s_i . Now reconsider actor a_q as the consumer (see Figure 15(b)) of the newly added channel (i.e., $(a_p, a_q) \in C_{s_i}$); as a result, all firings of actor a_p in one iteration of s_i provide $CNT(a_p, s_i)$ tokens for only one firing of a_q .

Let set A_i be the set of the actors in the PSOS s_i . Consider that each actor $a_p \in A_i$ appears $r \cdot \gamma(a_p)$ times in the PSOS s_i ($r = \frac{u}{v}$ where $u, v \in \mathbb{N}$) where the value r is identical for all actors in s_i . The appearance count of the actor $a_p \in$

A_i in s_i is represented by $CNT(a_p, s_i)$ and it is assumed to be equal to $\frac{u}{v} \cdot \gamma(a_p)$. We can write the above statement as follow: $CNT(a_p, s_i) \cdot v = \gamma(a_p) \cdot u$. From this equation we can conclude that u iterations of the SDFG G cause v iterations of the PSOS s_i , leading to v firings of each of the actors added by DSM (i.e., actors from the set A_{s_i}). So, in u iterations of the SDFG G (or v iterations of the PSOS s_i) the following equation holds for each channel $(a_p, a_q) \in C_{s_i}$ or $(a_q, a_p) \in C_{s_i}$ (where $a_p \in A_i$ and $a_q \in A_{s_i}$):

$$\underbrace{CNT(a_p, s_i)}_{Rate(a_q)} \cdot \underbrace{v}_{\gamma'(a_q)} = \underbrace{1}_{Rate(a_p)} \cdot \underbrace{\gamma(a_p) \cdot u}_{\gamma'(a_p)} \quad (1)$$

The entries of the repetition vector of the schedule-extended graph can be obtained using Equation 1. So a non-trivial repetition vector γ' can be found for the schedule-extended SDFG G' . The relation between the repetition vector of the schedule-extended SDFG and the repetition vector of the original graph is as follows: $\gamma'(a_q)$ (where $a_q \in A_{s_i}$) is equal to v and $\gamma'(a_p)$ (where $a_p \in A$) is equal to $\gamma(a_p) \cdot u$. \square

The following proposition states that inter-iteration execution for actors of a PSOS s_i is impossible in the SDFG G' which extends the original SDFG G with PSOS s_i using DSM.

Proposition 2. *PSOS inter-iteration execution is impossible for any actor appearing in PSOS $s_i = \langle \alpha_1 \alpha_2 \dots \alpha_n \rangle^*$ in the SDFG $G'(A', C')$, which models PSOS s_i in the SDFG $G(A, C)$ using DSM.*

Proof. Let set A_i be the set of the actors in the PSOS s_i . In one iteration of a PSOS s_i , an actor $a_p \in A_i$ could be enabled more often than its designated number (i.e., $CNT(a_p, s_i)$ times). DSM prevents this by creating a dependency from the last actor appearing in s_i (i.e., actor $a_L = \alpha_n$) to the first actor appearing in s_i (i.e., actor $a_F = \alpha_1$) (see lines 4-8 in DSM). This dependency is created by inserting a new actor $a_{i.end}$ and two channels $c_{i.pre}$ and $c_{i.pro}$. The source (destination) actor of the channel $c_{i.pre}$ is a_L ($a_{i.end}$) with rate 1 ($CNT(a_L, s_i)$). So, one firing of $a_{i.end}$ needs $CNT(a_L, s_i)$ tokens available in $c_{i.pre}$; for this purpose actor a_L should fire $CNT(a_L, s_i)$ times. The destination actor of the channel $c_{i.pro}$ is a_F ; the consumption rate of a_F from $c_{i.pro}$ is 1. The source actor of the channel $c_{i.pro}$ is $a_{i.end}$; the production rate of $a_{i.end}$ to $c_{i.pro}$ is $CNT(a_F, s_i)$. So the firing of a_F related to one iteration of PSOS s_i needs $CNT(a_F, s_i)$ tokens available in $c_{i.pro}$; for this purpose actor $a_{i.end}$ should fire once. The channel $c_{i.pro}$ is initialized with $CNT(a_F, s_i)$ number of tokens; this number of tokens provides enough tokens for actor a_F to fire $CNT(a_F, s_i)$ times in one iteration of PSOS s_i . The subsequent firing of actor a_F depends on the firing of the actor $a_{i.end}$ and the firing of the actor $a_{i.end}$ demands $CNT(a_L, s_i)$ times a firing of actor a_L . Hence, the firing of a_F belonging




Fig. 15. Extra actor a_q added by DSM in different situations.

to the subsequent iteration of s_i can be performed only after a_L finishes all of its firings belonging to the current iteration of s_i (i.e., after completion of the current iteration of s_i). This holds because actor a_F is the first actor which should be fired in an iteration of s_i and other actors in s_i cannot get enabled before the first firing of a_F . This second fact is guaranteed by adding decision state constructs (i.e., lines 12-18 in DSM) for any possible decision state and Proposition 6 below. \square

Even after eliminating inter-iteration execution from the SDFG, multiple actors from a schedule may be enabled in an SDFG iteration. We call such a state a *decision state*. The following proposition shows that analyzing only one SDFG iteration is enough to identify all possible decision states.

Proposition 3. *Executing an SDFG $G(A, C)$ for one iteration is enough to find all decision states within a PSOS s_i .*

Proof. Let the set A_i be the set of the actors in s_i and $A_o = A \setminus A_i$ the remaining actors in A . Consider inter-processor channels, which are defined as channel originating from the actors in A_o to the actors in A_i denoted by $C_{ipc} = \{(a_p, a_q) \in C \mid a_p \in A_o \wedge a_q \in A_i\}$. The execution of an actor $a_p \in A_o$ where $(a_p, a_q) \in C_{ipc}$ up-to its entry in the repetition vector of the SDFG produces $\gamma(a_p) \cdot \text{Rate}(a_p)$ tokens in the corresponding channel $(a_p, a_q) \in C_{ipc}$. Actor $a_q \in A_i$ consumes those produced tokens within one iteration of the normalized PSOS s_i (because $\gamma(a_p) \cdot \text{Rate}(a_p) = \gamma(a_q) \cdot \text{Rate}(a_q)$). Hence, actors in A_i can receive the required tokens from all inter-processor channels C_{ipc} for one iteration of the normalized PSOS s_i . This means that all possible decision states related to PSOS s_i are detectable.

Actors in A_o could possibly fire more than the number mentioned above (i.e., corresponding value in vector γ) if the channel dependencies in the SDFG allow additional firings of these actors. This could cause more than enough tokens (for one iteration of the normalized PSOS s_i) in channels C_{ipc} . This could enable an actor in A_i more than its designated number in one iteration of the normalized PSOS s_i . To avoid this undesirable actor enabling, some constructs are added to the graph to prevent inter-iteration execution (see Proposition 2). As a result, extra tokens produced by further firing of the actors A_o cannot enable any actor in A_i more than its designated value in one iteration of the normalized PSOS s_i . So, executing actors in the SDFG up-to their repetition vector entry (i.e., one SDFG iteration) is enough to find all decision states within s_i . \square

The identified decision states may be redundant. Proposition 4 discusses the proposed decision state reduction in the DSM.

Proposition 4. *Let σ be an execution of an SDFG (A, C) and a PSOS s_i which schedules actors $A_i \subseteq A$. In the execution σ , consider y consecutive decision states $\omega_{x+1}, \omega_{x+2}, \dots, \omega_{x+y}$. Assume that $a_o \in A_i$ is an opponent actor in each of these decision states but not the actor of choice in any of them. It is sufficient to only consider the last decision state ω_{x+y} to postpone the firing of the opponent actor a_o in those consecutive decision states to the state $\omega_{x+y+1} \in \sigma$.*

Proof. The purpose of the components added by DSM (i.e., lines 13-18 in Algorithm 1) in a decision state $\omega_j \in \Omega$ is to prevent any opponent actor a_o which is not the actor of choice in decision state ω_j from getting enabled in that state and as such to postpone that firing to the state ω_{j+1} . It is assumed that the opponent actor a_o is enabled in the consecutive decision states $\omega_{x+1}, \omega_{x+2}, \dots, \omega_{x+y}$. Suppose that the opponent actor a_o was fired e times before the first decision state (i.e., ω_{x+1}) where $0 \leq e < \gamma(a_o)$. The components added by DSM in the last decision state ω_{x+y} prevent the opponent actor a_o from getting enabled for the $(e+1)^{th}$ time in decision state ω_{x+y} . As a result, the opponent actor a_o can also not be enabled in states $\omega_{x+1}, \omega_{x+2}, \dots, \omega_{x+y-1}$ after adding DSM components for decision state ω_{x+y} . So, the components added by DSM in the last decision state of consecutive decision states are enough to prevent the firing of an opponent actor which is not the actor of choice in those consecutive decision states. \square

Decision state folding overlaps the consecutive repetitions of the designated PSOS in an SDFG iteration to reduce the number of decision states. The following proposition states that decision state folding does not dismiss any decision state.

Proposition 5. *Consider PSOS $s_i = \langle \alpha_1 \alpha_2 \dots \alpha_n \rangle^*$ for the sub-set of actors A_i from SDFG (A, C) ; assume s_i is repeated μ_i times to form the corresponding normalized PSOS $s'_i = \langle (s_i)^{\mu_i} \rangle^*$ to identify decision states related to PSOS s_i . After decision state folding all decision states are preserved.*

Proof. Normalization can be done by repeating s_i μ_i times (μ_i is the normalization factor of s_i). Decision state identification is applied to the normalized PSOS $s'_i = \langle (s_i)^{\mu_i} \rangle^* = \underbrace{\langle \alpha_1 \alpha_2 \dots \alpha_n \rangle}_{1^{st}} \underbrace{\langle \alpha_1 \alpha_2 \dots \alpha_n \rangle}_{2^{nd}} \dots \underbrace{\langle \alpha_1 \alpha_2 \dots \alpha_n \rangle}_{\mu_i^{th}}$. Decision point constructs are added based on the given PSOS s_i . Folding groups the identified decision states of the normalized PSOS s'_i in the following manner: the state related to the firing of an actor α_j in s_i is considered as decision state if a decision state is found at least in one of the μ_i states that α_j fires in the execution related to s'_i .

The corresponding state related to firing of the actor α_j in s_i which is considered as a decision state imposes a decision state to all μ_i corresponding states of actor α_j in s'_i . So, no decision state will be lost after decision state folding and the only effect is introducing unnecessary decision state controlling. We need to show that this extra controlling does not affect the execution of the SDFG according to the schedule. The construct added in a decision state is used to guarantee execution of the actor of choice of that decision state. Even if a state is not a decision state, the added constructs for that unnecessary decision state only ensure that the only enabled actor in that state (i.e., the actor of choice) can be fired. This does not violate the actor firing order according to the PSOS in that state. \square

DSM adds some components per decision state to enforce the firing of the enabled actor in a decision state which is in line with the given PSOS (i.e., actor of choice). Proposition 6 explains how those components can guarantee the firing of the actor of choice in the decision state.

Proposition 6. *The PSOS s_i is a schedule for actors $A_i \subseteq A$ from SDFG (A, C) . Let $\omega_j \in \Omega$ be a decision state within PSOS s_i and $\Delta_j \subseteq A_i$ the set of the opponent actors in decision state ω_j . The actor of choice in decision state ω_j (denoted by actor a_c) is the only actor which can fire in decision state ω_j among all actors in Δ_j after applying DSM. Also, one iteration of the resulting schedule-extended SDFG resets the tokens to their initial positions (i.e., the periodic behavior of the original SDFG is also preserved).*

Proof. We need to show that the opponent actors $\Delta_j \setminus \{a_c\}$ can not be enabled in the decision state ω_j after applying DSM. Accordingly, the actor of choice a_c in the decision state ω_j is the only actor which can fire in ω_j among all actors in Δ_j .

In the DSM technique, actor $a_{i.\omega_j}$ is added for each decision state $\omega_j \in \Omega$ within the PSOS s_i . An opponent actor $a_k \in \Delta_j \setminus \{a_c\}$ in the decision state ω_j is dependent on the new actor $a_{i.\omega_j}$ because of the added channel $c_{i.a_k\omega_j}$; this channel is initialized with $BEF(a_k, pos[\omega_j], s_i)$ tokens. The rate of the channel $c_{i.a_k\omega_j}$ on the side of a_k ($a_{i.\omega_j}$) is equal to one ($CNT(a_k, s_i)$). The new actor $a_{i.\omega_j}$ is also dependent on the actor of choice a_c of the decision state ω_j because of the added channel $c_{i.a_c\omega_j}$; this channel is initialized with $AFT(a_c, pos[\omega_j], s_i)$ tokens. The rate of the channel $c_{i.a_c\omega_j}$ on the side of a_c ($a_{i.\omega_j}$) is equal to one ($CNT(a_c, s_i)$).

An opponent actor $a_k \in \Delta_j \setminus \{a_c\}$ fires $BEF(a_k, pos[\omega_j], s_i)$ times before decision state ω_j and every time the opponent actor a_k consumes one token from $c_{i.a_k\omega_j}$. So, the $BEF(a_k, pos[\omega_j], s_i)$ firings of a_k before ω_j consume all tokens which were available in channel $c_{i.a_k\omega_j}$. Hence, the opponent actor $a_k \in \Delta_j \setminus \{a_c\}$ cannot fire in ω_j . Firings of the opponent actor $a_k \in \Delta_j \setminus \{a_c\}$ from decision state ω_j onward will be dependent on the firing of $a_{i.\omega_j}$ to provide the required tokens in channel $c_{i.a_k\omega_j}$. As mentioned before, $a_{i.\omega_j}$ depends on the actor of choice a_c . So, the opponent actor $a_k \in \Delta_j \setminus \{a_c\}$ cannot fire from decision state ω_j onward, unless the actor of choice a_c of the decision state ω_j fires. Hence, the actor of choice a_c is the only actor among the other opponent actors in ω_j which can fire. The actor of choice a_c is fired $BEF(a_c, pos[\omega_j], s_i)$ times by state ω_j and this results in $BEF(a_c, pos[\omega_j], s_i)$ tokens being produced in channel $c_{i.a_c\omega_j}$; as channel $c_{i.a_c\omega_j}$ is initialized with $AFT(a_c, pos[\omega_j], s_i)$ tokens, the number of tokens in this channel is $BEF(a_c, pos[\omega_j], s_i) + AFT(a_c, pos[\omega_j], s_i) = CNT(a_c, s_i)$ after firing a_c in decision state ω_j . So, there will be sufficient tokens (for one firing of $a_{i.\omega_j}$) on the only channel leading to $a_{i.\omega_j}$ after firing of the actor of choice a_c in decision state ω_j . Then, firing of $a_{i.\omega_j}$ consumes all $CNT(a_c, s_i)$ tokens that are present in channel $c_{i.a_c\omega_j}$ and it produces $CNT(a_k, s_i)$ tokens in channel $c_{i.a_k\omega_j}$ ($a_k \in \Delta_j \setminus \{a_c\}$); therefore, the opponent actors $\Delta_j \setminus \{a_c\}$ are not any more dependent on $a_{i.\omega_j}$ in the remainder of the current iteration of s_i .

The firings of the opponent actor $a_k \in \Delta_j \setminus \{a_c\}$ after decision state ω_j consumes $AFT(a_k, pos[\omega_j], s_i)$ tokens from channel $c_{i.a_k\omega_j}$; as a result, at the end of the PSOS iteration, the number of tokens on this channel returns to its initial value which is

$BEF(a_k, pos[\omega_j], s_i)$ (because $BEF(a_k, pos[\omega_j], s_i) = CNT(a_k, s_i) - AFT(a_k, pos[\omega_j], s_i)$). The actor of choice a_c fires $AFT(a_c, pos[\omega_j], s_i)$ times after decision state ω_j and the number of tokens in $c_{i.a_c\omega_j}$ returns to $AFT(a_c, pos[\omega_j], s_i)$. These initial token resettings at the end of the PSOS iteration ensure the periodic behavior for the added components in each decision state. Thus, in a decision state only the actor of choice (which is in line with the given schedule) amongst all opponent actors of the decision state can fire and this eliminates any non-determinism because of the decision state. \square

The following theorems state the correctness of DSM in modeling a single PSOS for a sub-set of the actors of the SDFG.

Theorem 1. *Consider PSOS s_i as a schedule for actors $A_i \subseteq A$ from SDFG $G(A, C)$. Assume $G'(A', C')$ is the SDFG that models s_i in G using DSM. For any execution σ' of $G'(A', C')$ it holds that σ satisfies s_i where it is assumed that σ is the execution of $G(A, C)$ with $orderList(\sigma, A) = orderList(\sigma', A)$.*

Proof. Proposition 6 states that in a decision state of s_i , an enabled actor of the decision state which is in line with s_i is the only actor able to fire in that state among all enabled actors in A_i . So, the order of s_i is the only possible order of actor firing for those actors of G' in the set A_i . Proposition 2 implies that the next PSOS iteration cannot interfere. Hence, for any execution σ' of $G'(A', C')$, $orderList(\sigma', A_i)$ has the form of $(s_i)^\kappa$ where $\kappa \in \mathbb{N}$ (i.e., infinite repetition of s_i). It is assumed that $orderList(\sigma, A) = orderList(\sigma', A)$; as $A_i \subseteq A$, we can conclude that $orderList(\sigma, A_i) = orderList(\sigma', A_i)$. Hence, $orderList(\sigma, A_i)$ also has the form of $(s_i)^\kappa$ and this form satisfies s_i ; in other words, σ satisfies s_i . \square

Theorem 2. *Consider PSOS s_i as a schedule for actors $A_i \subseteq A$ from SDFG $G(A, C)$. Assume $G'(A', C')$ is the SDFG that models s_i in G using DSM. For any execution σ of G that satisfies s_i it holds that there is exactly one σ' that is an execution of $G'(A', C')$ such that $orderList(\sigma, A) = orderList(\sigma', A)$.*

Proof. It is assumed that the firing order of actors belonging to the set A in execution σ' has the same actor firing order as in execution σ and execution σ satisfies s_i . We should show that there is precisely one execution with the property of execution σ' and that is a valid execution for G' . In a precise way, the actor firing order related to the set A in execution σ is a possible actor firing order for the original actors (i.e., actors not added by DSM) of G' when σ' is an execution of G' . Actor a_x from $orderList(\sigma, A)$ belongs either to A_i or to $A_o = A \setminus A_i$ ($x \in \mathbb{N}$). The state transition $\omega_x \xrightarrow{a_x} \omega_{x+1}$ in execution σ is related to the firing of a_x . The state transition $\omega'_y \xrightarrow{a_y} \omega'_{y+1}$ in execution σ' is related to the firing of a_y . The difference between states from execution σ and σ' is only in the extra channels added by DSM (i.e., C_{s_i}). Any channel from C_{s_i} is connected to an actor from A_i ; in other words, it is not connected to any actor from A_o . Incoming channels of an actor

determine whether that actor can fire or not. Consider ω'_y from execution σ' has the same content of state ω_x from execution σ for all channels in C . Hence, each $a_x \in A_o$ that fires in ω_x of execution σ can also fire in ω'_y of execution σ' (i.e., a_y is a_x) because the content of the state related to channels C_{s_i} has no influence on actor enabling for any actor from A_o . But, when an actor a_x belongs to A_i , the content of the state related to channels C' could manipulate the actor firing order. It is assumed that firings of actor a_x in execution σ satisfy the PSOS s_i ; actor a_x can also fire in the corresponding state (i.e., ω'_y) of execution σ' (i.e., a_y is a_x) because components added by DSM force the firing of the actor which is in line with the given s_i (i.e., a_x) among all actors in A_i (see Proposition 6). So, the firing order of actors from $A' \setminus A_{s_i}$ in execution σ' follows the same firing order as it is indicated in execution σ . Each actor $a \in A_{s_i}$ also has a single possible firing order in each PSOS iteration; if a is added to control the actor firing in a decision state, it fires before the actor of choice in the decision state (see Proposition 6) and if it is added for the inter-iteration prevention, it fires at the end of the PSOS iteration (see Proposition 2). Hence, there exists only one possible firing order in execution σ' for each actor $a \in A_{s_i}$. So, all actors from A' have exactly one firing order in execution σ' where $orderList(\sigma, A) = orderList(\sigma', A)$. \square

The size of the schedule-extended graph (e.g., G') is dependent on the number of decision states found in the given schedule (e.g., s_i). In this section, decision state identification for a sub-set of actors of G (i.e., $A_i \subseteq A$) which belong to the schedule of interest (i.e., s_i) is explained regardless of existing other schedules for the rest of the actors in the SDFG (i.e., actors in $A_o = A \setminus A_i$). In our implementation, we consider other possible schedules designated for the rest of the actors (i.e., actors in A_o) to reduce the number of the decision states and as a result the size of the schedule-extended graph. As we explained in Section V-E2, actors which do not belong to s_i should fire according to their schedule (if there is any) to perform their maximal execution. Firing of any actor $a_o \in A_o$, which belongs to another PSOS s_j ($j \neq i$), in function $maxExec$ of the DSM algorithm should satisfy the schedule to which a_o belongs. In other words, a_o in function $maxExec$ of DSM fires if (1) a_o is enabled and (2) its firing satisfies s_j . Without the second condition, firing of a_o could enable an actor from s_i and lead to unnecessary decision states. So, considering other schedules in the function $maxExec$ of DSM algorithm removes such redundant decision states.

APPENDIX B CORRECTNESS OF SASM

This section discusses the correctness of SASM in modeling a periodic static-order SAS for a sub-set of the actors of the SDFG. We can model multiple SASs by applying the algorithm multiple times.

SASM adds some actors, denoted by A_{s_i} , and channels, denoted by C_{s_i} to model the SAS s_i in SDFG $G(A, C)$. $G'(A', C')$ is the SDFG that models s_i in G using SASM where $A' = A \cup A_{s_i}$ and $C' = C \cup C_{s_i}$.

Between any (nested) LSs α_u and α_v , where $v = (u \bmod n) + 1$, in a LS SAS $s_i = \{(\alpha_1)^{\beta_1}(\alpha_2)^{\beta_2} \dots (\alpha_n)^{\beta_n}\}$, SASM adds one counter actor a_{cnt_u} and two channels c_{cnt_u} and c_{lim_u} . The following propositions shows how these components control the actor firings in α_u and α_v .

Proposition 7. *For any (nested) LSs α_u and α_v , where $v = (u \bmod n) + 1$, in LS $s_i = \{(\alpha_1)^{\beta_1}(\alpha_2)^{\beta_2} \dots (\alpha_n)^{\beta_n}\}$, SASM ensures that α_v can only be executed for β_v times after α_u is executed for β_u times.*

Proof. SASM creates a dependency from the right-most actor (RMA) in α_u to the left-most actor (LMA) in α_v by adding one counter actor a_{cnt_u} and two channels c_{cnt_u} and c_{lim_u} . The channel c_{cnt_u} is placed from the RMA in α_u to a_{cnt_u} ; the production rate of the RMA in α_u on channel c_{cnt_u} is set to one and consumption rate of a_{cnt_u} on channel c_{cnt_u} is set to the number of times that the RMA in α_u is needed to fire in β_u executions of α_u . Hence, actor a_{cnt_u} can fire once after the looped schedule α_u is executed for β_u times (the 1st behavior). The channel c_{lim_u} is placed from a_{cnt_u} to the LMA in α_v ; the production rate of a_{cnt_u} on channel c_{lim_u} is set to the number of times that the LMA in α_v is needed to be fired in β_v executions of α_v and the consumption rate of the LMA in α_v on channel c_{lim_u} is set to one. Hence, the looped schedule α_v can be executed for β_v times after one firing of a_{cnt_u} (the 2nd behavior). It is assumed that SASM is recursively applied to each (nested) looped schedules in s_i . Considering both of the explained 1st and 2nd behaviors provides that α_v can only be executed for β_v times after α_u is executed for β_u times. \square

SASM only places a specific number of initial tokens (i.e., the number of times that the LMA of α_1 should fire in $\alpha_1^{\beta_1}$) on the last limiter channel (i.e., c_{lim_n}) to prevent deadlock in execution; the following proposition explains how this prevents inter-iteration execution in the schedule-extended graph.

Proposition 8. *Inter-iteration execution is impossible for any actor appearing in SAS s_i in the SDFG $G'(A', C')$.*

Proof. Considering the behavior described in Proposition 7 for all pairs of consecutive (nested) looped schedules in s_i ensures that the next schedule iteration cannot be started before the current schedule iteration ends. The initial token placed on the last limiter channel (i.e., c_{lim_n}) only provides tokens for β_1 executions of the looped schedule α_1 ; the executions of α_1 related to the current schedule iteration consumes all tokens on the last limiter channel and this prevents the next iteration from being executed until completion of the current iteration (i.e., α_n is executed for β_n times). \square

The following theorems state the correctness of SASM. These theorems are similar to the theorems given in Appendix VII.

Theorem 3. *Consider SAS $s_i = \{(\alpha_1)^{\beta_1}(\alpha_2)^{\beta_2} \dots (\alpha_n)^{\beta_n}\}$ as a schedule for actors $A_i \subseteq A$ from SDFG $G(A, C)$. Assume $G'(A', C')$ is the SDFG that models s_i in G using SASM. For any execution σ' of $G'(A', C')$ it holds that σ satisfies s_i where it is assumed that σ is the execution of $G(A, C)$ with $orderList(\sigma, A) = orderList(\sigma', A)$.*

Proof. Considering Proposition 7 for all pairs of consecutive (nested) LSs in s_i guarantees that the order specified in s_i is the only possible order of actor firing for those actors of the SDFG G' in the set A_i . Proposition 8 implies that subsequent SAS iterations cannot interfere. Hence, for any execution σ' of SDFG $G'(A', C')$, $orderList(\sigma', A_i)$ has the form of $(s_i)^\kappa$ where $\kappa \in \mathbb{N}$ (i.e., infinite repetition of s_i). It is assumed that $orderList(\sigma, A) = orderList(\sigma', A)$; as $A_i \subseteq A$, we can conclude that $orderList(\sigma, A_i) = orderList(\sigma', A_i)$. Hence, $orderList(\sigma, A_i)$ also has the form of $(s_i)^\kappa$ and this form satisfies s_i ; in other words, σ satisfies s_i . \square

Theorem 4. Consider SAS s_i as a schedule for actors $A_i \subseteq A$ from SDFG $G(A, C)$. Assume $G'(A', C')$ is the SDFG that models s_i in G using SASM. For any execution σ of $G(A, C)$ that satisfies s_i it holds that there is exactly one σ' that is an execution of $G'(A', C')$ such that $orderList(\sigma, A) = orderList(\sigma', A)$.

Proof. It is assumed that the firing order of actors belonging to the set A in execution σ' is the same as in execution σ and σ satisfies s_i . We should show that there is precisely one execution with the property of execution σ' and that is a valid execution for G' . The actor firing order related to the set A in execution σ (i.e., $orderList(\sigma, A)$) is a possible actor firing order for the original actors (i.e., actors not added by SASM) of G' when σ' is an execution of G' . Actor a_x from $orderList(\sigma, A)$ belongs either to A_i or to $A_o = A \setminus A_i$ ($x \in \mathbb{N}$). The state transition $\omega_x \xrightarrow{a_x} \omega_{x+1}$ in execution σ is related to the firing of a_x . The state transition $\omega'_y \xrightarrow{a_y} \omega'_{y+1}$ in execution σ' is related to the firing of a_y . The difference between states from execution σ and σ' is only in the channels added by SASM (i.e., C_{s_i}). Any channel from C_{s_i} is connected to an actor from A_i ; in other words, it is not connected to any actor from A_o . Incoming channels of an actor determine whether that actor can fire or not. Consider state ω'_y from execution σ' has the same content of state ω_x from execution σ for all channels in C . Hence, each $a_x \in A_o$ that fires in ω_x of execution σ can also fire in ω'_y of execution σ' (i.e., a_y is a_x) because the content of the state related to channels C_{s_i} has no influence on actor enabling for any actor from A_o . But, when an actor a_x belongs to A_i , the content of the state related to channels C' could manipulate the actor firing order. As we assume firings of a_x in execution σ satisfy s_i , a_x can fire in the corresponding state (i.e., ω'_y) of execution σ' (i.e., a_y is a_x) because components added by SASM force the firing of the actor which is in line with the given SAS s_i among all actors in A_i (considering Proposition 7 for all pairs of consecutive nested LSs) and it is assumed that a_x is in line with s_i . So, the firing order of actors from $A' \setminus A_{s_i}$ in execution σ' follows the same firing order as indicated in σ . Each counter actor $a_{cnt_x} \in A_{s_i}$ also has a single possible firing order in each schedule iteration; it fires once after the LS a_x executed for β_x times (see proof of Proposition 7). Hence, there exists only one possible firing order in σ' for each actor $a \in A_{s_i}$. So, all actors from A' have exactly one firing order in σ' where $orderList(\sigma, A) = orderList(\sigma', A)$. \square