



Controllers for the Verification of Communicating Multi-Pushdown Systems

Cyriac Aiswarya, Paul Gustin, K. Narayan Kumar

► To cite this version:

Cyriac Aiswarya, Paul Gustin, K. Narayan Kumar. Controllers for the Verification of Communicating Multi-Pushdown Systems. 25th International Conference on Concurrency Theory (CONCUR'14), 2014, Rome, Italy. pp.297-311. hal-01057525

HAL Id: hal-01057525

<https://hal.science/hal-01057525>

Submitted on 22 Aug 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Controllers for the Verification of Communicating Multi-Pushdown Systems^{*}

C. Aiswarya¹, Paul Gastin², and K. Narayan Kumar³

¹ Uppsala University, Sweden

`aiswarya.cyriac@it.uu.se`

² LSV, ENS Cachan, CNRS & INRIA, France

`gastin@lsv.ens-cachan.fr`

³ Chennai Mathematical Institute, India

`kumar@cmi.ac.in`

Abstract. Multi-pushdowns communicating via queues are formal models of multi-threaded programs communicating via channels. They are turing powerful and much of the work on their verification has focussed on under-approximation techniques. Any error detected in the under-approximation implies an error in the system. However the successful verification of the under-approximation is not as useful if the system exhibits unverified behaviours. Our aim is to design *controllers* that observe/restrict the system so that it stays within the verified under-approximation. We identify some important properties that a good controller should satisfy. We consider an extensive under-approximation class, construct a distributed controller with the desired properties and also establish the decidability of verification problems for this class.

1 Introduction

Most of the critical hardware and software consists of several parallel computing units/components. Each of these may execute recursive procedures and may also have several unbounded data-structures to enhance its computing power. Several of such components may be running on the same processor giving rise to a multi-threaded system with many unbounded data-structures. Furthermore, such complex infinite state systems may communicate over a network and be physically distributed. The high computational power in combination with unconstrained interactions make the analysis of these systems very hard.

The verification of such systems is undecidable in general. Even the basic problem of control state reachability (or emptiness checking) is undecidable as soon a program has two stacks or a self queue. However, these systems are so important, that several under-approximation techniques have been invented for their verification. If the under-approximation fails to satisfy a requirement, that immediately indicates an error in the system. However, if the system is verified correct under such restrictions, the correctness is compromised if the system

^{*} Supported by LIA InForMel, and DIGITEO LoCoReP.

eventually exhibits behaviours outside the class. *Controlling* the system to only exhibit behaviours that have been verified to be correct is therefore crucial to positively use these under-approximation techniques. Alternately, we may use these controllers to raise a signal whenever the system behaviour departs from the verified class. For example, in the cruise control system of a car (or auto-pilot systems in trains/aircrafts), it will be useful to signal such a departure and switch from automatic to manual mode.

Our contributions We aim at obtaining a uniform controller for a class, which when run in parallel with the system, controls it so as to exhibit only those behaviours permitted by the class. Such a controller should possess nice properties like determinism, non-blocking, system independence etc. In Section 3, we identify and analyse such desirable features of a controller.

Our next contribution is to propose a very generous under-approximation class and to construct a controller satisfying all the desired properties. Our class bounds the number of phases – in a phase only one data-structure can be read in an unrestricted way though writes to all data-structures are allowed. But our notion of phases extends sensibly contexts of [13] and phases of [12]. In particular it permits *autonomous computations* within a phase instead of the well-queueing assumption. The latter corresponds to permitting reads from queues in the *main* program but not from any of the functions it calls. We permit recursive calls to be at any depth of recursion when reading from a queue. After such a read, however, returning from the function causes a phase change.

A concurrent system may be controlled in a global manner or in a distributed manner. If the concurrent processes are at a single location and communicate via shared variables, e.g., multi-threaded programs, a global controller is reasonable. We describe this sequential controller in Section 4. However, when these multi-threaded processes are physically distributed it is natural to demand a distributed controller. In Section 5, we illustrate the design of a controllable under-approximation class by extending our idea of phases to the distributed setting and constructing a distributed controller with all the desired properties.

Finally, we can prove using the split-width technique [5–7] that our generous under-approximation class can be model-checked against a wide variety of logics.

Related Work: In the study of distributed automata a number of difficult synthesis theorems [8–10, 18] have been proved. These theorems in conjunction with constructions for intersections yield controllers for these classes. Of particular interest is the theory of finite state machines communicating via queues, called message-passing automata (MPA). These have been well studied using labeled partial-orders (or graphs) called MSCs (Message-sequence charts) to represent behaviours. These systems are turing powerful and techniques restricting channel usage have been studied to obtain decidability. The most general class of this kind, called existentially k -bounded MSCs, consists of all behaviours (MSCs) that have at least one linearization in which the queue lengths are bounded by k at every point. A deep result of [8] shows that for each k there is an MPA which accepts precisely the set of existentially k -bounded MSCs. Thus, if one

uses such behaviours as an under-approximation class then this result implies the existence of a distributed controller. However, it is known that this controller cannot be made deterministic.

The bounding technique for verification has been extensively studied in the case of multi-pushdown systems (MPDS). For the restrictions studied in literature, bounded-context [17], bounded-phase [12], bounded-scope [15] and ordered stacks [2, 3], it is quite easy to construct deterministic controllers, though this question has not been addressed before. The context bounding technique is extended to pushdown systems communicating via queues under the restriction that queues may be read only when the stacks are empty (well-queuing) in [13], and under a dual restriction (on writes instead of reads) in [11]. Controllability is however not studied there. The k -Phase restriction we consider here is a natural joint generalization of these contexts (as well as the bounded-phase restriction for MPDS). In fact, for every bound k , there exist behaviours which are not captured by [12] and [13], but which are captured by our class with a bound of 3. (See Figure 1 for an example.)

2 Systems with stacks and queues

We provide a formal description of systems with data-structures and their behaviours. We restrict ourselves to systems with global states providing an (interleaved) sequential view. In Section 5 we extend this to the distributed case where there are a number of components each with their own collection of transitions. We consider a finite set $\mathbf{DS} = \mathbf{Stacks} \uplus \mathbf{Queues}$ of data-structures which are either stacks or queues and a finite set Σ of actions. Our systems have a finite set of control locations and use these (unbounded) stacks and queues. We obtain an interesting class of infinite state systems, providing an (interleaved) sequential view of multi-threaded recursive programs communicating via FIFO channels.

A stack-queue system (SQS) over data-structures \mathbf{DS} and actions from Σ is a tuple $\mathcal{S} = (\text{Locs}, \text{Val}, \text{Trans}, \text{in}, \text{Fin})$ where Locs is a finite set of locations, Val is a finite set of values that can be stored in the data-structures, $\text{in} \in \text{Locs}$ is the initial location, $\text{Fin} \subseteq \text{Locs}$ is the set of final locations, and Trans is the set of transitions which may write a value to, or read a value from, or do not involve a data-structure. For $\ell, \ell' \in \text{Locs}$, $a \in \Sigma$, $d \in \mathbf{DS}$ and $v \in \text{Val}$, we have

- internal transitions of the form $\ell \xrightarrow{a} \ell'$,
- write transitions of the form $\ell \xrightarrow{a, d!v} \ell'$, and
- read transitions of the form $\ell \xrightarrow{a, d?v} \ell'$.

Intuitively, an SQS consists of a finite state system equipped with a collection of stacks and queues. In each step, it may use an internal transition to merely change its state, or use a write transition to append a value to the tail of a particular queue or stack or use a read transition to remove a value from the head (or tail) of a queue (of a stack respectively). The transition relation makes explicit the identity of the data-structure being accessed and the type of the operation. As observed in [1, 6, 12, 16] it is often convenient to describe the

runs of such systems as a state-labeling of words decorated with a matching relation per data-structure instead of the traditional operational semantics using configurations and moves. This will prove all the more useful when we move to the distributed setting where traditionally semantics has always been given as state-labelings of appropriate partial orders [8, 10, 18].

A stack-queue word (SQW) over \mathbf{DS} and Σ is a tuple $\mathcal{W} = (w, (\triangleright^d)_{d \in \mathbf{DS}})$ where $w = a_1 a_2 \cdots a_n \in \Sigma^+$ is the sequence of actions, and for each $d \in \mathbf{DS}$, the matching relation $\triangleright^d \subseteq \{1, \dots, n\}^2$ relates write events to data-structure d to their corresponding read events. The following conditions should be satisfied:

- write events should precede read events: $e \triangleright^d f$ implies $e < f$,
- data-structure accesses are disjoint: if $e_1 \triangleright^d e_2$ and $e_3 \triangleright^{d'} e_4$ are distinct edges ($d \neq d'$ or $(e_1, e_2) \neq (e_3, e_4)$) then they are disjoint ($|\{e_1, e_2, e_3, e_4\}| = 4$),
- $\forall d \in \mathbf{Stacks}$, \triangleright^d conforms to LIFO: if $e_1 \triangleright^d f_1$ and $e_2 \triangleright^d f_2$ are different edges then we do not have $e_1 < e_2 < f_1 < f_2$.
- $\forall d \in \mathbf{Queues}$, \triangleright^d conforms to FIFO: if $e_1 \triangleright^d f_1$ and $e_2 \triangleright^d f_2$ are different edges then we do not have $e_1 < e_2$ and $f_2 < f_1$.

We let $\triangleright = \bigcup_{d \in \mathbf{DS}} \triangleright^d$ be the set of all matching edges and $\mathcal{E} = \{1, \dots, n\}$ be the set of events of \mathcal{W} . The set of all stack-queue words is denoted by SQW.

We say that an event e is a *read event* (on data-structure d) if there is an f such that $f \triangleright^d e$. We define *write events* similarly and an event is *internal* if it is neither a read nor a write. To define the run of an SQS over a stack-queue word \mathcal{W} , we introduce two notations. For $e \in \mathcal{E}$, we denote by e^- the immediate predecessor of e if it exists, and we let $e^- = \perp \notin \mathcal{E}$ otherwise. We let $\max(\mathcal{W})$ be the maximal event of \mathcal{W} .

A run of an SQS \mathcal{S} on a stack-queue word \mathcal{W} is a mapping $\rho: \mathcal{E} \rightarrow \text{Locs}$ satisfying the following consistency conditions (with $\rho(\perp) = \text{in}$):

- if e is an internal event then $\rho(e^-) \xrightarrow{\lambda(e)} \rho(e) \in \text{Trans}$,
- if $e \triangleright^d f$ for some data-structure $d \in \mathbf{DS}$ then for some $v \in \text{Val}$ we have both $\rho(e^-) \xrightarrow{\lambda(e), d!v} \rho(e) \in \text{Trans}$ and $\rho(f^-) \xrightarrow{\lambda(f), d?v} \rho(f) \in \text{Trans}$.

The run is accepting if $\rho(\max(\mathcal{W})) \in \text{Fin}$. The *language* $\mathcal{L}(\mathcal{S})$ accepted by an SQS \mathcal{S} is the set of stack-queue words on which it has an accepting run.

Notice that SQSs are closed under intersection, by means of the cartesian product. Let $\mathcal{S}_i = (\text{Locs}_i, \text{Val}_i, \text{Trans}_i, \text{in}_i, \text{Fin}_i)$ for $i \in \{1, 2\}$ be two SQSs. The cartesian product is $\mathcal{S}_1 \times \mathcal{S}_2 = (\text{Locs}_1 \times \text{Locs}_2, \text{Val}_1 \times \text{Val}_2, \text{Trans}, (\text{in}_1, \text{in}_2), \text{Fin}_1 \times \text{Fin}_2)$ where the set of transitions is defined by

- $(\ell_1, \ell_2) \xrightarrow{a} (\ell'_1, \ell'_2) \in \text{Trans}$ if $\ell_i \xrightarrow{a} \ell'_i \in \text{Trans}_i$ for $i \in \{1, 2\}$,
- $(\ell_1, \ell_2) \xrightarrow{a, d!(v_1, v_2)} (\ell'_1, \ell'_2) \in \text{Trans}$ if $\ell_i \xrightarrow{a, d!v_i} \ell'_i \in \text{Trans}_i$ for $i \in \{1, 2\}$,
- $(\ell_1, \ell_2) \xrightarrow{a, d?(v_1, v_2)} (\ell'_1, \ell'_2) \in \text{Trans}$ if $\ell_i \xrightarrow{a, d?v_i} \ell'_i \in \text{Trans}_i$ for $i \in \{1, 2\}$.

In fact, $\mathcal{S}_1 \times \mathcal{S}_2$ has an (accepting) run on a stack-queue word \mathcal{W} iff both \mathcal{S}_1 and \mathcal{S}_2 have an (accepting) run on \mathcal{W} . Therefore, $\mathcal{L}(\mathcal{S}_1 \times \mathcal{S}_2) = \mathcal{L}(\mathcal{S}_1) \cap \mathcal{L}(\mathcal{S}_2)$.

3 Controllers and Controlled Systems

SQSs are turing powerful as soon as **DS** contains two stacks or a queue, and hence their verification is undecidable. However, since it is an important problem, various under-approximation techniques have been invented in the recent years [2, 3, 12, 13, 15, 17], starting with the *bounded-context restriction* [17] for systems with only stacks. Here, the number of times the system switches from using one stack to another is bounded by a fixed number k . Reachability and many other properties become decidable when restricted to such behaviours.

A typical under-approximation technique describes a whole family of classes \mathbb{C}_k parametrized by an integer k which is proportional to the coverage: the higher the parameter, the more behaviours are covered. For example, the bound on number of context switches k serves as this parameter for the context bounding technique. Ideally, the under-approximations defined by the classes $(\mathbb{C}_k)_k$ should be *universal*, i.e., should cover all behaviours: every stack-queue word \mathcal{W} should be in \mathbb{C}_k for some k . This is true for the context bounding technique.

Traditionally under-approximations yield decidability for verification problems such as reachability [17] and model checking against linear time properties expressed in various logics upto MSO [16]. For such properties, if the model-checking problem yields a negative answer then this immediately means that the full system fails the verification as well.

However, assume that a system \mathcal{S} has been verified against some linear-time or reachability property (or properties) wrt. some under-approximation class \mathbb{C} . This give us little information on whether the full system satisfies these properties. Hence we need a mechanism, which we call a *controller*, to restrict the system so that it does not exhibit behaviours outside \mathbb{C} . Observe that w.r.t. linear-time properties restricting the system to even a proper subset of \mathbb{C} would still be acceptable though not desirable. However, for reachability properties a proper restriction might lead to a system that no longer satisfies the property. Therefore, a controller should allow all and only the behaviours of \mathbb{C} .

We now describe formally our notion of a controller for a class and examine some key properties that make it interesting.

A controller for a class $\mathbb{C} \subseteq \text{SQW}$ is an SQS \mathcal{C} such that $\mathcal{L}(\mathcal{C}) = \mathbb{C}$. We say that a class \mathbb{C} is *controllable* if it admits a controller.

Suppose the restriction of the behaviours of a system \mathcal{S} to a class \mathbb{C} has been verified against some linear-time or reachability property φ . Further suppose that \mathbb{C} admits a controller \mathcal{C} . Then, the *controlled system* $\mathcal{S}' = \mathcal{S} \times \mathcal{C}$ is such that $\mathcal{L}(\mathcal{S}') = \mathcal{L}(\mathcal{S}) \cap \mathbb{C}$, and therefore satisfies φ . Thus, a controller for a class is independent of the system \mathcal{S} as well as the property. Once we identify a controllable class with decidable verification we may verify and control any system in a completely generic and transparent manner without any additional work. Notice that we could have introduced more general controllability. For instance, a class \mathbb{C} is non-uniformly controllable if for each system \mathcal{S} , there exists another system \mathcal{S}' such that $\mathcal{L}(\mathcal{S}') = \mathcal{L}(\mathcal{S}) \cap \mathbb{C}$. While this would allow more classes to

be controllable, it would not be very useful since it does not yield an automatic way to build \mathcal{S}' from \mathcal{S} .

Using the cartesian product makes the controller integrable into the system. The controller, by definition, does not have its own auxiliary data-structures, but only shares the data-structures of the system. Moreover, it does not access a data-structure out of sync with the system. We could also give more intrusive power to a controller by allowing its transitions to depend on the current state of the system and on the current value read/written by the system on data-structures. But again, such a system would not be generic, and also, by its strong observation power, would compromise the privacy of the system.

We now consider other properties that a good controller must satisfy and use that to arrive at a formal definition of such a controller.

The under-approximation classes are often defined based on the data-structure accesses, and do not depend on the action labels/internal actions. Hence an ideal controller should be definable independent of the action labels and must be oblivious to the internal moves. This can be done as follows.

We omit action labels from read/write transitions of \mathcal{C} : an abstract transition $\ell \xrightarrow{d!v} \ell'$ stands for transitions $\ell \xrightarrow{a,d!v} \ell'$ for all $a \in \Sigma$ and similarly for read transitions. Also, we do not describe internal transitions and assume instead that there are self-loops $\ell \xrightarrow{a} \ell$ for all locations and actions.

This (abstract) controller should be deterministic and non-blocking, so that instantiating it with any alphabet will still be deterministic and non-blocking. Thus, the controller should have a unique run on any \mathcal{W} and moreover this run does not depend on the internal events / action labels along the run, but depends only on the sequence of reads/writes on the different data-structures that appear along \mathcal{W} . The state of the controller at any point along this run unambiguously indicates whether the current prefix can be extended to a word that belongs to the class \mathbb{C} . With this we are ready to formalize our notion of a good controller.

A DS-controller is an SQS \mathcal{C} which is oblivious to internal events and to action labels and which is deterministic and non-blocking. Formally, its (abstract) transitions should satisfy:

- for every $\ell \in \text{Locs}$ and $d \in \mathbf{DS}$ there exists exactly one $\ell' \in \text{Locs}$ and $v \in \text{Val}$ such that $\ell \xrightarrow{d!v} \ell'$,
- for every $\ell \in \text{Locs}$, $d \in \mathbf{DS}$ and $v \in \text{Val}$ there exists exactly one $\ell' \in \text{Locs}$ such that $\ell \xrightarrow{d?v} \ell'$.

All that we said so far suffices for a global (or sequential) system. If the system to be verified and controlled is actually physically distributed, then a global sequential controller would not be integrable in the system. Instead we would need a distributed controller and this is much harder to achieve. We discuss this in Section 5.

Next we examine real examples of controllable under-approximations. While an under-approximation \mathbb{C}_k is *nicely controllable* if it admits a controller with the above features, the class itself should satisfy some other properties for it to be

useful. Firstly, \mathbb{C}_k should have a wide coverage over the set of possible behaviours. A useful feature is that all behaviours fall in the class for an appropriately chosen parameter. Second, the definition of the class should be easy to describe. Finally, the verification problem for the class should be decidable. For instance, considering the collection of behaviours with clique/split/tree-width bounded by k satisfies the first and third properties but does not satisfy the second property. But more importantly, it is not clear that they have nice controllers of the form described above. We propose a meaningful class which has more coverage than bounded phase of [12], and is nicely controllable. We show the decidability of this class by demonstrating a bound on split-width.


4 Class and Controller: Sequential case

We begin by identifying a class of behaviours, called k -Phase behaviours, which is verifiable and admits a **DS**-controller. Roughly speaking, a phase is a segment of the run where the reads are from a fixed data-structure. However, between successive reads, read-free recursive computations are permitted which may write to all data-structures, including their own call-stack. We formalize this below.

An autonomous computation involves a single recursive thread executing a recursive procedure without reading any other data structure. All read events are from a single stack while there is no restriction placed on the writes. We say that an edge $e \triangleright f$ is *autonomous* if $e \triangleright^s f$ for some $s \in \mathbf{Stacks}$ and all in-between read events are from the same stack s : if $e' \triangleright^d f'$ with $e \leq f' \leq f$ then $d = s$. We shall write \triangleright_a for the subset of \triangleright consisting of the autonomous edges and \triangleright_{na} for $\triangleright \setminus \triangleright_a$ and refer to them as the non-autonomous edges. If $e \triangleright_a f$ then e and f are called *autonomous* write and read events respectively.

A **d -phase** is a sequence of consecutive events in which all non-autonomous reads are from the data-structure $d \in \mathbf{DS}$. Writes to all data-structures are permitted. Moreover, a phase must not break an autonomous computation. Formally, a d -phase is identified by a pair of events $e \leq f$ (the first and the last events in the sequence) such that, if $e' \triangleright_{na} f'$ with $e \leq f' \leq f$ then $e' \triangleright^d f'$ and if $e' \triangleright_a f'$ with $e \leq f' \leq f$ or $e \leq e' \leq f$ then $e \leq e' \leq f' \leq f$.

Example 1. Suppose $\mathbf{DS} = \{q, s_1, s_2\}$. A q -phase is depicted on the right. Straight lines (resp. curved lines) represent \triangleright^d edges from queues (resp. stacks). Autonomous computations are highlighted in white.



Remark 2. Permitting autonomous (recursive) computations during a phase is a natural generalization of *well-queueing* assumption of [13] where reads from queues are permitted only when the stack associated with a process is empty. The latter corresponds to permitting reads from queues in the *main* program

but not from any of the functions it calls. We permit recursive calls to be at any depth of recursion when reading from the queue. After such a read, however, returning from the function causes a phase change.

Our aim is to obtain a decidable and controllable class by bounding the number of phases. In the presence of queues, reading and writing on a queue during a phase can be used to simulate a turing machine using just 1-phase computations. Allowing autonomous computations on one stack while reading and writing on another also results in the same effect. This motivates the following definition which rules out such *self-loops*.

A phase identified by a pair (e, f) has a *self-loop* if it contains a non-autonomous edge: $e \leq e' \triangleright_{na} f' \leq f$.

A phase decomposition is a partition of the set of events into phases with no self-loops. A k -phase decomposition is a phase decomposition with at most k phases. We denote by k -Phase the class of stack-queue words that admit a k -phase decomposition.

Remark 3. Observe that by freely allowing autonomous computations (as opposed to well-queuing), every stack-queue word is in k -Phase for some k .

Remark 4. When restricted to systems with only stacks, k -Phase subsumes the k bounded phase restriction for multi-pushdown systems [12]. It also subsumes the k bounded context restriction for systems with stacks and queues [13]. In fact, for every bound k , there exist stack-queue words which are not captured by [12] and [13], but which are in 3-Phase. (See Figure 1.)

A phase with no self-loops identified by (e, f) is *upper-maximal* if it cannot be extended upwards in a phase with no self-loops: if (e, g) is a phase with no self-loops then $g \leq f$. Given any k phase decomposition, we may extend the first phase to be upper-maximal and then extend the next (remaining) phase to be upper-maximal and so on till all the phases are upper maximal.

Remark 5. A phase (e, f) with no self-loops is upper-maximal iff the successor f' of f is a non-autonomous read on some data-structure d' ($e' \triangleright_{na}^{d'} f'$) with either $e \leq e'$ (self-loop) or d' is not the data-structure d associated with phase (e, f) , i.e., there is a non-autonomous read $e'' \triangleright_{na}^d f''$ with $e \leq f'' \leq f$.

Lemma 6. *Every stack-queue word in k -Phase admits a maximal k -phase decomposition in which all phases are upper-maximal.*

Proof. We start with an arbitrary k -phase decomposition. Let (e, f) be the first phase (if any) in this decomposition which is not upper-maximal. Let $g > f$ be maximal such that (e, g) is a phase with no self-loops. We replace in the decomposition phase (e, f) by (e, g) . Notice that (e, g) may completely subsume some of the phases following (e, f) which are therefore removed from the decomposition. It may also partially subsume one phase (e', f') if $e' \leq g < f'$. This phase is trimmed in order to start at the successor of g . We repeat this procedure until all phases are upper-maximal. Note that the number of phases do not increase so that we get a maximal k -phase decomposition. \square




Fig. 1: A stack-queue word over two stacks and its maximal phase decomposition.

Now we take up the task of constructing a **DS**-controller for the class k -Phase. A crucial step towards this end is to identify autonomous reads. We show below that this can be achieved with a multi-pushdown automaton \mathcal{B} observing the data-structure access. When the system \mathcal{S} writes/reads some value on a stack s the automaton \mathcal{B} will simultaneously write/read a bit on the same stack. \mathcal{B} is obtained as a cartesian product of automata \mathcal{B}_s ($s \in \mathbf{Stacks}$) identifying the autonomous reads on stack s (described in Figure 2).

Here, $s!b$ (resp. $s?b$) means that the system \mathcal{S} writes/reads on stack s and b is the tag bit that is simultaneously written/read by \mathcal{B}_s on stack s . The other events do not change stack s . Moreover, $\bar{s}?$ is the observation of a read event of \mathcal{S} which is *not* on stack s , and *else* means any event which is not explicitly specified.




Fig. 2: The automaton \mathcal{B}_s .

We say that e is a *possibly autonomous write* to stack s at event g if $e \triangleright^s f$ and $e \leq g < f$ and $e' \triangleright^d f'$ with $e \leq f' \leq g$ implies $d = s$. Intuitively \mathcal{B}_s will be in state 1 iff in the current prefix there is an unmatched write event e to stack s which is possibly autonomous. On a write to s the automaton moves from state 0 to 1 since this write is possibly autonomous, and pushes 1 on the stack to indicate that it is the first possibly autonomous write in the past. Then, as long as it does not read from a data-structure $d \neq s$, it stays in state 1, pushing 0 on the stack on a write to s and reading 0 from the stack on a read from s . If it reads 1 from the stack, then it has matched the first possibly autonomous write in the past, hence it goes back to state 0. On a read from $d \neq s$ it goes to state 0 since there cannot be any possibly autonomous write to s at this read event.

Lemma 7. *The automaton \mathcal{B}_s is deterministic and non-blocking. Moreover, in the unique run of \mathcal{B}_s on a word, the state b_s before a read from stack s determines whether this read is autonomous ($b_s = 1$) or not ($b_s = 0$).*

Proof. Automaton \mathcal{B}_s is non-blocking, thanks to the *else* loops. Clearly, it is also deterministic.

Let $w = a_1 a_2 \dots a_n$ be a visible word, i.e., for each letter a_i we know whether it is a write on stack s ($a_i \in s!$) or a read from stack s ($a_i \in s?$) or a read from a data-structure $d \neq s$ ($a_i \in \bar{s}?$), or another type of event ($a_i \in \text{else}$). Let $(q_0, \gamma_0), (q_1, \gamma_1), \dots, (q_n, \gamma_n)$ be the sequence of configurations along the unique run of \mathcal{B}_s on w .

We can easily prove by induction that for all $1 \leq k \leq n$ we have $q_k = 1$ iff at event k there is a possibly autonomous write event $j \leq k$. To carry out this induction, we simultaneously prove that if $q_k = 1$ then the stack contents γ_k is of the form $\{0,1\}^*10^\ell$ where $\ell + 1$ is the number of possibly autonomous write events at k .

We deduce that a read event a_{k+1} from stack s is autonomous iff $q_k = 1$. \square

We now construct the deterministic **DS**-controller \mathcal{C}^k for k -Phase. This controller computes the maximal phase decomposition of a behaviour and uses the automaton \mathcal{B} to identify autonomous reads. We denote by $\mathbf{b} = (b_s)_{s \in \mathbf{Stacks}}$ a state of \mathcal{B} . In addition, a state of \mathcal{C}^k holds two other values:

- a counter $n \in \{1, \dots, k, \infty\}$ which indicates the current phase number. The counter starts from value 1 and is non-decreasing along a run. The ∞ indicates that the number of phases has exceeded k . We follow the convention that $i + 1$ has the usual meaning if $i < k$, $k + 1 = \infty$ and $\infty + 1 = \infty$.
- a value $d \in \mathbf{DS} \cup \{?\}$ which indicates that the current phase has non-autonomous reads from $d \in \mathbf{DS}$ or that only autonomous reads have occurred so far ($d = ?$). Note that in the first phase all reads are autonomous (a non-autonomous read would create a self-loop). Hence, $d = ?$ iff $n = 1$.

The initial state of the controller is $(1, ?, \mathbf{0})$. On an internal event, the state remains unchanged. When the system writes to a data-structure the controller \mathcal{C}^k writes its current phase number in addition to the bits written by \mathcal{B} .

$$(n, d, \mathbf{b}) \xrightarrow{d'!n} (n, d, \mathbf{b}) \quad \text{if } d' \in \mathbf{Queues} \quad (1)$$

$$(n, d, \mathbf{b}) \xrightarrow{d'!(n,c)} (n, d, \mathbf{b}') \quad \text{if } d' \in \mathbf{Stacks} \wedge \mathbf{b} \xrightarrow{d'!c} \mathbf{b}' \text{ in } \mathcal{B} \quad (2)$$

Notice that in the first case, $\mathbf{b} \xrightarrow{d'!} \mathbf{b}$ is a transition in \mathcal{B} . A read event from a queue d' will stay in the same phase if d' is the current data-structure and the matching write comes from a previous phase (to avoid self-loops): if $d' \in \mathbf{Queues}$ then we have the following transitions in \mathcal{C}^k

$$(n, d, \mathbf{b}) \xrightarrow{d'?m} (n, d', \mathbf{0}) \quad \text{if } d' = d \wedge m < n \quad (3)$$

$$(n, d, \mathbf{b}) \xrightarrow{d'?m} (n + 1, d', \mathbf{0}) \quad \text{otherwise} \quad (4)$$

Notice that in these cases, $\mathbf{b} \xrightarrow{d'?} \mathbf{0}$ is a transition in \mathcal{B} since no stack can be on an autonomous computation at a read event from a queue. Further if $d = ?$, reading from a queue forces a phase change. This is needed, as otherwise there will be a self-loop on the first phase.

Finally, a read event from a stack s will stay in the same phase if it is an autonomous read ($b_s = 1$), or $s = d$ is the current data-structure and this read does not create a self-loop: if $s \in \mathbf{Stacks}$ then in \mathcal{C}^k we have the transitions

$$(n, d, \mathbf{b}) \xrightarrow{s?(m,c)} (n, d, \mathbf{b}') \quad \text{if } (b_s = 1 \vee (s = d \wedge m < n)) \wedge \mathbf{b} \xrightarrow{s?c} \mathbf{b}' \text{ in } \mathcal{B} \quad (5)$$

$$(n, d, \mathbf{b}) \xrightarrow{s?(m,c)} (n + 1, s, \mathbf{0}) \quad \text{otherwise} \quad (6)$$

Notice that in the last case, $\mathbf{b} \xrightarrow{s?c} \mathbf{0}$ is a transition in \mathcal{B} and thus in all moves the third component stays consistent with moves of \mathcal{B} .




Fig. 3: A run of the deterministic sequential controller \mathcal{C}^k .

By construction the controller is deterministic and non-blocking. If the unique run of the controller on a \mathcal{W} does not use a state of the form (∞, d, \mathbf{b}) then \mathcal{W} is in k -Phase. The set of positions labeled by states of the form (i, d, \mathbf{b}) identify the i th phase in a k phase decomposition. Conversely, let \mathcal{W} be in k -Phase. Let (\mathbf{b}_e) be the state labeling position e in \mathcal{W} in the unique run of \mathcal{B} on \mathcal{W} . Let $(X_i)_{(i \leq l)}$ be the phases in the maximal decomposition of \mathcal{W} . It is easy to verify that the first position of X_i , $i \geq 2$ is a non-autonomous read and let d_i be the data-structure associated with this read. Then the labeling assigning $(1, ?, \mathbf{b}_e)$ to any position $e \in X_1$ and (i, d_i, \mathbf{b}_e) to any event e in X_i , $2 \leq i \leq l$ is an accepting run of the controller on \mathcal{W} .

Theorem 8. *The SQS \mathcal{C}^k is a DS-controller for the class k -Phase with $(|\mathbf{DS}| \cdot (k+1) + 1)2^{|\mathbf{Stacks}|}$ states.*

Proof. Let ρ be a run of \mathcal{C}_k controller on a stack-queue word \mathcal{W} . Let X_i be the set of events labeled by a state of the form (i, d, \mathbf{b}) for some $d \in \mathbf{DS} \cup \{?\}$ and \mathbf{b} some state of \mathcal{B} . Let $l \leq k$ be the maximum i such that X_i is not empty. Then, by the definition of the transition relation, X_j is nonempty for each $i \leq l$. We claim that $(X_i)_{i \leq l}$ is a (maximal) k -phase decomposition of \mathcal{W} .

Notice that if (i, d, \mathbf{b}) and (i, d', \mathbf{b}') are states labeling two events in X_i , then from the definition of the transition relation, $d = d'$. We let $d_i \in \mathbf{DS} \cup \{?\}$ to be the unique d such that some event in X_i is labeled by a state of the form (i, d_i, \mathbf{b}) for some \mathbf{b} .

By Lemma 7 the state of the \mathcal{B} automaton labelling any read event e determines correctly whether it is a autonomous or non-autonomous read. We first observe that each X_i contains non-autonomous read events on at most one data structure. This can be seen as follows.

In any state labelling and event in X_1 , the only read transitions permitted are from 5 and have to be autonomous transitions. Next we consider X_i with $i > 1$. From the definition of transitions 3, 4, 5 and 6, any non-autonomous read event e on a data-structure d has to be labeled by a state of the form (j, d, \mathbf{b}) . But, from the above, if $e \in X_i$ then $j = i$ and $d = d_i$ and thus all non-autonomous reads are on the same data-structure.

The only transitions that increment the first component are read transitions from 4, 6 and these cannot be autonomous. Thus if (e, f) defines an autonomous read-write pair then by definition any read event between e and f must also be

autonomous and by the correctness of \mathcal{B} , the transitions 4 and 6 will never be applicable. Thus if e is in X_i then so is f . Thus, each X_i is a phase.

We also observe that if $e \triangleright_{na} f$ and $e \in X_i$ then $f \notin X_i$. This is because, by the correctness of \mathcal{B} , one of the rules 4 or 6 is applicable at f . Thus, each X_i is acyclic as well and thus the word \mathcal{W} is in k -Phase and this completes one direction of the proof.

Next, we show that the controller has an accepting run on any \mathcal{W} that belongs to k -Phase. Let X_1, X_2, \dots, X_l be the partitions of \mathcal{W} in the maximal k -phase decomposition of \mathcal{W} . Let e_i be the first event in X_i , $2 \leq i \leq l$. By the property described in Remark 5, e_i is a read event and let d_i be the data-structure associated with e_i . In the unique run of the automaton \mathcal{B} on \mathcal{W} let (b_e) be the state labeling event e . We claim that the labeling ρ that labels event e in X_1 by $(1, ?, b_e)$ and every event e in X_i , $2 \leq i \leq l$ by (i, d_i, b_e) is an accepting run of the controller on \mathcal{W} . We only need to verify the accuracy of the first two components of the labeling as the correctness of the third component follows by definition.

Internal events pose no problems since \mathcal{B} does not change states on internal events. Let e be any event other than e_i , $1 < i \leq l$. If e is a write event then the first two components do not change according to transitions 1 and 2 as required by the labeling. If e is a read event then there are 3 cases: if it is an autonomous read then by the correctness of \mathcal{B} transition rule 5 is applicable and it leaves the first two components unchanged as required by the labeling. If it reads from a queue then since e is not one of e_i s the data-structure involved is d_i and by the acyclicity of the decomposition the corresponding write event occurs in a strictly earlier phase. Thus rule 3 will be applicable leaving the first two components unchanged. Finally, if e is a read event on a stack and is not autonomous then once again the associated data-structure must be d_i and by acyclicity the corresponding write has to be from an earlier phase and thus 5 is applicable and leaves the first two components unchanged. Thus, the labeling agrees with the transition relation at all events other than the e_i s ($2 \leq i \leq l$).

Let $f_{i-1} = e_i - 1$. Then $f_{i-1} \in X_{i-1}$. Clearly e_i cannot be an autonomous read. By Remark 5 there are two possibilities for e_i . If the data-structure d_i is different from the data-structure d_{i-1} , then one of the rules 4 or 6 is applicable and consistent with the labeling. If the data-structure d_i is the same as d_{i-1} then, by Remark 5 it must be the case that the corresponding write event must be in phase $i - 1$. Again one of the rules 4 or 6 is applicable consistent with the labels. This completes the proof. \square

5 Class and Controller: Distributed case

In this section we describe a model intended to capture collections of SQS communicating via reliable FIFO channels (or queues). Such systems are called Stack-Queue Distributed System (SQDS). A behaviour of an SQDS is a tuple of stack-queue words with additional matching relations describing the inter-process communication via queues. Such behaviours extend Message Sequence

Charts (MSCs) with matching relations for the internal stacks and queues. We call them stack-queue MSCs (SQMSC).

We then extend the notion of k -Phase to this distributed setting. We show that k -Phase enjoys a deterministic distributed controller with local acceptance conditions.

An architecture \mathfrak{A} is a tuple $(\mathbf{Procs}, \mathbf{Stacks}, \mathbf{Queues}, \text{Writer}, \text{Reader})$ consisting of a set of processes \mathbf{Procs} , a set of stacks \mathbf{Stacks} , a set of queues \mathbf{Queues} and functions Writer and Reader which assign to each stack/queue the process that will write (push/send) into it and the process that will read (pop/receive) from it respectively. We write \mathbf{DS} for $\mathbf{Stacks} \uplus \mathbf{Queues}$.

A stack d must be local to its process, so $\text{Writer}(d) = \text{Reader}(d)$. On the other hand, a queue d may be local to a process p if $\text{Writer}(d) = p = \text{Reader}(d)$, otherwise it provides a FIFO channel from $\text{Writer}(d)$ to $\text{Reader}(d)$.

A Stack-Queue Distributed System (SQDS) over an architecture \mathfrak{A} and an alphabet Σ is a tuple $\mathcal{S} = (\text{Locs}, \text{Val}, (\text{Trans}_p)_{p \in \mathbf{Procs}}, \text{in}, \text{Fin})$ where each $\mathcal{S}_p = (\text{Locs}, \text{Val}, \text{Trans}_p, \text{in}, \emptyset)$ is an SQS over \mathbf{DS} and Σ in which the transitions are compatible with the architecture: Trans_p may have a write (resp. read) transitions on data-structure d only if $\text{Writer}(d) = p$ (resp. $\text{Reader}(p) = d$). Moreover, $\text{Fin} \subseteq \text{Locs}^{\mathbf{Procs}}$ is the global acceptance condition. We say that the acceptance condition is local if $\text{Fin} = \prod_{p \in \mathbf{Procs}} \text{Fin}_p$ where $\text{Fin}_p \subseteq \text{Locs}$ for all $p \in \mathbf{Procs}$.

A stack-queue MSC (SQMSC) over architecture \mathfrak{A} and alphabet Σ is a tuple $\mathcal{M} = ((w_p)_{p \in \mathbf{Procs}}, (\triangleright^d)_{d \in \mathbf{DS}})$ where $w_p \in \Sigma^*$ is the sequence of events on process p and \triangleright^d is the relation matching write events on data-structure d with their corresponding read events. We let $\mathcal{E}_p = \{(p, i) \mid 1 \leq i \leq |w_p|\}$ be the set of events on process $p \in \mathbf{Procs}$. For an event $e = (p, i) \in \mathcal{E}_p$, we set $\text{pid}(e) = p$ and $\lambda(e)$ be the i th letter of w_p . We write \rightarrow for the successor relation on processes: $(p, i) \rightarrow (p, i+1)$ if $1 \leq i < |w_p|$ and we let $\triangleright = \bigcup_{d \in \mathbf{DS}} \triangleright^d$ be the set of all matching edges. We require the relation $< = (\rightarrow \cup \triangleright)^+$ to be a strict partial order on the set of events. Finally, the matching relations should comply with the architecture: $\triangleright^d \subseteq \mathcal{E}_{\text{Writer}(d)} \times \mathcal{E}_{\text{Reader}(d)}$. Moreover, data-structure accesses should be disjoint, stacks should conform to LIFO and queues should conform to FIFO (the formal definitions are taken verbatim from Section 2). An SQMSC is depicted in Figure 4.

As before, to define the run of an SQDS over a stack-queue MSC \mathcal{M} , we introduce two notations. For $p \in \mathbf{Procs}$ and $e \in \mathcal{E}_p$, we denote by e^- the unique event such that $e^- \rightarrow e$ if it exists, and we let $e^- = \perp_p \notin \mathcal{E}$ otherwise. We let $\max_p(\mathcal{M})$ be the maximal event of \mathcal{E}_p if it exists and $\max_p(\mathcal{M}) = \perp_p$ otherwise.

A run of an SQDS \mathcal{S} over a stack-queue MSC \mathcal{M} is a mapping $\rho: \mathcal{E} \rightarrow \text{Locs}$ satisfying the following consistency conditions (with $\rho(\perp_p) = \text{in}$):

- if e is an internal event then $\rho(e^-) \xrightarrow{\lambda(e)} \rho(e) \in \text{Trans}_{\text{pid}(e)}$,
- if $e \triangleright^d f$ for some data-structure $d \in \mathbf{DS}$ then for some $v \in \text{Val}$ we have both $\rho(e^-) \xrightarrow{\lambda(e), d!v} \rho(e) \in \text{Trans}_{\text{pid}(e)}$ and $\rho(f^-) \xrightarrow{\lambda(f), d?v} \rho(f) \in \text{Trans}_{\text{pid}(f)}$.




Fig. 4: A stack-queue MSC and its maximal phase decomposition.


The run is accepting if $(\rho(\max_p(\mathcal{M})))_{p \in \text{Procs}} \in \text{Fin}$. The *language* $\mathcal{L}(\mathcal{S})$ accepted by an SQDS \mathcal{S} is the set of stack-queue MSCs on which it has an accepting run.

Notice that SQDSs are closed under intersection, by means of the cartesian product. The construction is similar to the one for SQSs in Section 2.

Bounded Acyclic Phase SQMSCs We generalize the under-approximation class k -Phase to the distributed setting. We allow at most k phases per process. As in the sequential case, autonomous computations are freely allowed. However, cycles on phases can be caused by the richer structure of the SQMSC than simple self loops.

In the distributed setting, the definitions of **autonomous computations** and of **d -phases** are identical to the sequential case, cf. Section 4. Again, we write \triangleright_a for autonomous edges and \triangleright_{na} for non-autonomous edges. A phase, which is a sequence of consecutive events executed by a single process, is identified by a pair of events (e, f) such that $e \rightarrow^* f$.

A phase (e, f) has a cycle if there is a non-autonomous edge $e' \triangleright_{na} f'$ with $e \leq e'$ and $f' \rightarrow^* f$. Notice that e' needs not be in the phase. So a cycle starts from the phase at e then follows the partial order to some non-autonomous write e' whose read f' is in the phase. A phase is *acyclic* if it has no cycles. Notice that a non-autonomous edge within a phase induces a cycle (self-loop) whereas autonomous edges are freely allowed within phases. As a matter of fact, when there is exactly one process, a phase has a cycle iff it has a self-loop.



A phase decomposition of an SQMSC is a partition of its set of events into phases. A phase decomposition is *acyclic* if all phases are acyclic. It is a k -phase decomposition if there are at most k phases per process. We denote by k -Phase the set of SQMSCs that admits an acyclic k -phase decomposition.

An acyclic phase (e, f) is *upper-maximal* if extending it upwards would result in a cycle, i.e., for every other acyclic phase (e, f') , we have $f' \leq f$. See Figure 4 for an example. Lemma 6 easily lifts up to the distributed case as well.

Lemma 9. *Every SQMSC in k -Phase admits a maximal acyclic k -phase decomposition in which all phases are upper-maximal.*

Deterministic Distributed Controller We extend the notion of nice controllers to the distributed setting. That means controllers should be *distributed* and have *local* acceptance conditions. A local controller for one process should be able to control the behaviour of that process regardless of the states of the other local controllers. The communication between the local controllers is also only by means of overloading the actual messages sent between the processes. The local controllers are not allowed to send messages out of sync, as it would create new behaviours in the controlled system. Thus a controlled system should be again obtained as a cartesian product of the system with a controller where both are SQDS, but in addition the controller has local acceptance conditions.

Theorem 10. *The class k -Phase admits a deterministic distributed controller \mathcal{C}^k with $(|\mathbf{DS}| \cdot (k + 2)^{|\mathbf{Procs}|} + 1)2^{|\mathbf{Stacks}|}$ states.*

The distributed controller is a generalisation of the sequential controller of Section 4. The main difference is that the local controller of process p remembers not only its current phase number, but a tuple $\mathbf{n} = (n_q)_{q \in \mathbf{Procs}}$ of phase numbers for each process. The intuition is that n_q is the largest phase of process q that is known to process p ($n_q = 0$ if no events of process q are in the past of the current event of process p).

For each stack s , we use the automaton \mathcal{B}_s defined in Section 4 that identifies autonomous reads. For each process $p \in \mathbf{Procs}$, we let \mathcal{B}_p be the product of the automata \mathcal{B}_s where s is a stack of process p (i.e., $s \in \mathbf{Stacks}$ and $\text{Writer}(s) = p$).

A state of the local controller \mathcal{C}_p^k for process p is a tuple $(\mathbf{n}, d, \mathbf{b}_p)$ where $\mathbf{n} = (n_q)_{q \in \mathbf{Procs}}$ is the phase vector with $n_q \in \{0, 1, \dots, k, \infty\}$, $d \in \mathbf{DS} \cup \{?\}$ with $\text{Reader}(d) = p$ if $d \neq ?$, and \mathbf{b}_p is a state of \mathcal{B}_p . The initial state of \mathcal{C}_p^k is $\text{in}_p = (\mathbf{n}, ?, \mathbf{0})$ with $n_p = 1$ and $n_q = 0$ for $q \neq p$. The local acceptance condition Fin_p is given by the set of states $(\mathbf{n}, d, \mathbf{b}_p)$ with $n_q \neq \infty$ for all $q \in \mathbf{Procs}$.

We describe now the local transitions of \mathcal{C}_p^k . They are similar to the transitions of the sequential controller given in Section 4. We start with write transitions, so let $d' \in \mathbf{DS}$ be such that $\text{Writer}(d') = p$. On write events, the current phase vector is written on to the data-structure (in addition to the autonomous bit where needed).

$$(\mathbf{n}, d, \mathbf{b}_p) \xrightarrow{d'!\mathbf{n}} (\mathbf{n}, d, \mathbf{b}_p) \quad \text{if } d' \in \mathbf{Queues} \quad (7)$$

$$(\mathbf{n}, d, \mathbf{b}_p) \xrightarrow{d'!(\mathbf{n}, c)} (\mathbf{n}, d, \mathbf{b}'_p) \quad \text{if } d' \in \mathbf{Stacks} \wedge \mathbf{b}_p \xrightarrow{d'!c} \mathbf{b}'_p \text{ in } \mathcal{B}_p \quad (8)$$

Let $d' \in \mathbf{Queues}$ be such that $\text{Reader}(d') = p$. The transitions of \mathcal{C}_p^k that read queue d' are given below. We should switch to the next phase 1) if $m_p = n_p$ since otherwise this non-autonomous read would close a cycle, 2) or if $d' \neq d \neq ?$ since in a phase all non-autonomous reads should be from the same data-structure.

$$(\mathbf{n}, d, \mathbf{b}_p) \xrightarrow{d'?\mathbf{m}} (\mathbf{n}', d', \mathbf{0}) \quad \text{if } m_p = n_p \vee (d' \neq d \neq ?) \quad (9)$$

$$\begin{aligned} & \text{with } n'_p = n_p + 1 \wedge n'_q = \max(n_q, m_q) \text{ for } q \neq p \\ (\mathbf{n}, d, \mathbf{b}_p) & \xrightarrow{d'?\mathbf{m}} (\mathbf{n}', d', \mathbf{0}) \quad \text{otherwise, with } \mathbf{n}' = \max(\mathbf{n}, \mathbf{m}) \end{aligned} \quad (10)$$

Similarly, we give below read transitions from $d' \in \mathbf{Stacks}$ with $\text{Reader}(d') = p$. Here a switch of phase is required under the same conditions but only when the read is not autonomous.

$$(\mathbf{n}, d, \mathbf{b}_p) \xrightarrow{d'?(m,c)} (\mathbf{n}', d', \mathbf{0}) \quad \text{if } b_{d'} = 0 \wedge (m_p = n_p \vee (d' \neq d \neq ?)) \quad (11)$$

with $n'_p = n_p + 1 \wedge n'_q = \max(n_q, m_q)$ for $q \neq p$

$$(\mathbf{n}, d, \mathbf{b}_p) \xrightarrow{d'?(m,c)} (\mathbf{n}', d, \mathbf{b}'_p) \quad \text{otherwise,} \quad (12)$$

with $\mathbf{n}' = \max(\mathbf{n}, \mathbf{m}) \wedge \mathbf{b}_p \xrightarrow{d'?c} \mathbf{b}'_p$ in \mathcal{B}_p

One of the differences of a local controller from a sequential controller is that the first phase may also perform non-autonomous reads. However, in such case, it must be from a queue.

On read transitions (10 and 12) which stay in the same phase, the phase vector is updated by taking the maximum between the current phase vector and the read-phase vector ($\mathbf{n}' = \max(\mathbf{n}, \mathbf{m})$). On a phase switch, a similar update is performed but the current phase number of process p is incremented.

Remark 11. The phase vectors of these transitions when projected to the p th component, gives a transition of the sequential controller. That is, for $d' \in \mathbf{Queues}$, $(n_p, d, \mathbf{b}_p) \xrightarrow{d'?m_p} (n'_p, d', \mathbf{b}'_p)$ in \mathcal{C}^k and for $d' \in \mathbf{Stacks}$, $(n_p, d, \mathbf{b}_p) \xrightarrow{d'?(m_p,c)} (n'_p, d', \mathbf{b}'_p)$ in \mathcal{C}^k .

Proof of Correctness of the Distributed Controller

Soundness ($\mathcal{L}(\mathcal{C}^k) \subseteq k\text{-Phase}$): An accepting run induces a phase decomposition. Let X_i^p be the set of events on process p which are labelled by states of the form $(\mathbf{n}, d, \mathbf{b}_p)$ with $n_p = i$. By definition of the controller, each X_i^p is a phase as it reads non-autonomously from only one data-structure. Thus, the phases are valid and since the run is accepting there are at most k of them on each process. The autonomous reads are computed and handled correctly, thanks to \mathcal{B}_p (see Lemma 7).

It remains to show that this induced phase decomposition is acyclic. Notice that the current phase vector is always propagated along a \triangleright edge and that the phase vector at the read event is updated by \max . We deduce that for every process $p \in \mathbf{Procs}$ and event g , the value n_p of the phase vector at g is the phase number of the maximal event $e \leq g$ which is on process p .

Now, if there is a cycle, say in the i th phase (e, f) of process p , then by definition there exists $e' \triangleright_{na} f'$ with $e \leq e'$ and $f' \rightarrow^* f$. By the above observation, since $e \leq e' \leq f$, the value of n_p at e' is i . Hence the transition taken at f' must be either (9) or (11) which gives the phase number $i + 1$ to f' . This is a contradiction since (e, f) has phase number i and $f' \leq f$. Hence there cannot be cycles, which completes the soundness proof.

Completeness ($k\text{-Phase} \subseteq \mathcal{L}(\mathcal{C}^k)$): Let \mathcal{M} be an SQMSC in $k\text{-Phase}$. Consider the maximal phase decomposition as per Lemma 9. We define the following for every event with respect to this maximal phase decomposition.

- $\text{phase-nbr}(e) = i \in \{1, \dots, k\}$ if e is on the i th phase on $\text{pid}(e)$. Moreover, we set $\mathbf{n}^e = (n_q^e)_{q \in \text{Procs}}$ where $n_q^e = \max\{\text{phase-nbr}(f) \mid f \leq e \wedge \text{pid}(f) = q\}$. As usual, we assume that $\max(\emptyset) = 0$.
- $\text{data-struct}(e) = d$ if there is an earlier event in the phase of e which is a non-autonomous read from d , i.e., there is some $e' \triangleright_{na}^d f'$ such that $f' \rightarrow^* e$ and $\text{phase-nbr}(f') = \text{phase-nbr}(e)$. We set $\text{data-struct}(e) = ?$ otherwise. Notice that this last case may occur only if $\text{phase-nbr}(e) = 1$.
- \mathbf{b}^e : the bit vector computed by the automaton $\mathcal{B}_{\text{pid}(e)}$ on the sequence of events up to e on its process. Notice that \mathcal{B}_p is deterministic and non-blocking (Lemma 7) and hence provides a unique value of \mathbf{b}^e for every e .

This allows us to extract a labelling $\rho : \mathcal{E} \rightarrow \text{Locs}$ of events by the states of the controller in the obvious manner: $\rho(e) = (\mathbf{n}^e, \text{data-struct}(e), \mathbf{b}^e)$. We prove below that ρ is the unique run of the controller \mathcal{C}^k on \mathcal{M} . Since none of the phase-numbers exceed k , the run is accepting, which proves that $k\text{-Phase} \subseteq \mathcal{L}(\mathcal{C}^k)$.

We prove that the labelling ρ on \mathcal{M} conforms to the transitions of \mathcal{C}^k by induction on the set \mathcal{E} of events of \mathcal{M} ordered by $<$.

Let $e \in \mathcal{E}_p$ be an event on process p and assume that for all events $f < e$ the labelling ρ conforms to the transitions of \mathcal{C}^k . Recall that we denote by e^- the immediate predecessor of e on process p if it exists and otherwise $e^- = \perp_p$. By convention, we let $\rho(\perp_p) = \text{in}_p$ be the initial state of \mathcal{C}_p^k . This means that $\text{data-struct}(\perp_p) = ?$, $\mathbf{b}^{\perp_p} = \mathbf{0}$ and $n_p^{\perp_p} = 1 = \text{phase-nbr}(\perp_p)$ and $n_q^{\perp_p} = 0$ for $q \neq p$.

Notice that \mathbf{b}^e is obtained from \mathbf{b}^{e^-} via the computation of \mathcal{B}_p in the definition of the labelling ρ and this conforms to the transitions of \mathcal{C}_p^k (7–12).

If e is an internal event, then we can easily check that $\rho(e) = \rho(e^-)$ (even if $e^- = \perp_p$), and hence conforms to the internal transitions which are self-loops.

If e is a write event to a queue, then again, from the definition of ρ , we can check that $\rho(e) = \rho(e^-)$, which is as dictated by transition (7).

If e is writing to a stack, then we see that $\rho(e)$ differs from $\rho(e^-)$ only for the bits \mathbf{b}^e due to the computations of \mathcal{B}_p . This conforms to transition (8).

For the remaining cases, e is a read event from some data-structure d' . Hence, let $f \triangleright^{d'} e$.

If e is an autonomous read then $\text{phase-nbr}(e) = \text{phase-nbr}(e^-)$, $\text{data-struct}(e) = \text{data-struct}(e^-)$ and $\mathbf{n}^e = \max(\mathbf{n}^{e^-}, \mathbf{n}^f)$. Hence, transition (12) validates $\rho(e)$.

Assume that e is a non-autonomous read (from a stack or a queue) and that $\text{phase-nbr}(e) = \text{phase-nbr}(e^-)$. Let e_0 be the minimal event in the phase of e with respect to the maximal phase decomposition. Then, $e_0 \not\leq f$ (otherwise there would be a cycle). Hence it follows that $n_p^f < n_p^e$. Moreover, $d' = \text{data-struct}(e) = \text{data-struct}(e^-)$ or $\text{data-struct}(e^-) = ?$ (otherwise it would have non-autonomous reads from different data-structures contradicting the definition of a phase). Moreover, $\mathbf{n}^e = \max(\mathbf{n}^{e^-}, \mathbf{n}^f)$ and we can check that transition (10) or (12) validates $\rho(e)$.

Finally, assume that $\text{phase-nbr}(e) = \text{phase-nbr}(e^-) + 1$. Then, e is not autonomous, and either $n_p^f = n_p^{e^-}$ (change of phase to avoid a cycle) or $d' =$

$\text{data-struct}(e) \neq \text{data-struct}(e^-) \neq ?$ (change of phase to avoid non-autonomous reads from different data-structures). We can check that transition (9) or (11) validates $\rho(e)$. This concludes the proof.

6 Decidability

In this section we explain why k -Phase is a verifiable under-approximation for SQDS. Consider the reachability problem which is equivalent to asking if given an SQDS \mathcal{S} and $k \in \mathbb{N}$ whether \mathcal{S} accepts at least one \mathcal{M} from k -Phase. A non-trivial extension of the technique of [13] allows to reduce the reachability problem of SQDS restricted to k -Phase to the reachability problem of multi-pushdown systems for bounded phase.

A more general question is to model-check properties expressed in linear time logics ranging from temporal logics to $\text{MSO}(\rightarrow, \triangleright^d)$. Given a formula φ we have to determine whether every $\mathcal{M} \in k$ -Phase that is accepted by \mathcal{S} satisfies φ . Observe that we may equivalently ask whether *every* behaviour of the controlled system \mathcal{S}' satisfies φ . Using a slightly different approach we can obtain decidability not only for reachability but also for the linear-time model-checking problems.

In this approach we show that every behaviour in k -Phase has *split-width* [5–7] or tree-width [16] or clique-width [4] (measures of the complexity of graphs that happen to be equivalent for our class of graphs) bounded by some function $f(k)$. Here, we show an exponential bound on the split-width. Then, results from [6, 7, 16] imply that MSO model-checking for \mathcal{S}' is decidable and results from [5, 7] imply that model-checking linear-time temporal logic formulas can be solved in double exponential time. This is optimal, since reachability of k -phase multi-pushdown systems is double exponential time hard [14].

The rest of this section is devoted to proving the bound on split-width. We first perform two refinements on the phase decomposition prior to bounding the split-width.

6.1 Acyclic to Strongly Acyclic Phase Decomposition

The *residual graph* of a phase decomposition is a graph whose vertices are the phases. There is an edge from phase u to phase v if there are events e in u and f in v with $e \triangleright_{na} f$. A phase decomposition is *strongly acyclic* if its residual graph is acyclic. We write k -SAPD to denote a strongly acyclic k phase decomposition.

Interestingly we show that any acyclic phase decomposition can be turned into a strongly acyclic phase decomposition with just a quadratic blow up in the number of phases. Let χ be the set of phases of a decomposition of an SQMSC \mathcal{M} . For an event $e \in \mathcal{E}$, let $\chi(e)$ denote the phase in χ containing e . For a process $p \in \mathbf{Procs}$, let χ_p denote the phases involving process p in χ . Thus, $\chi = \biguplus \chi_p$. The elements of χ_p are linearly ordered in the obvious manner ($X \leq Y$ iff $e \leq f$ for some $e \in X$ and $f \in Y$ or equivalently, if $e \leq f$ for all $e \in X$ and $f \in Y$).

We say that a phase X in χ is *part of a weak cycle* if the corresponding vertex is part of cycle in the residual graph. Our aim is to refine the phases in

such a way that none of the phases is part of a weak cycle. We fix an acyclic phase decomposition χ for the rest of this section. To begin with we divide χ into three parts \mathcal{G} , \mathcal{O} and \mathcal{R} (Green, Orange and Red). A phase $X \in \chi_p$ is in \mathcal{O} iff X is part of a weak cycle and there is no $Y < X$ in χ_p that is part of a cycle. Thus, \mathcal{O} consists of the minimal phase from each process if any, that is part of at least one weak cycle. A phase $X \in \chi_p$ is in \mathcal{R} iff $Y < X$ for some $Y \in \chi_p \cap \mathcal{O}$. All other phases are in \mathcal{G} . Observe that any phase in \mathcal{G} is not part of any weak cycle. We say that a phase (or an event) is green (respectively orange or red) if it (or its phase) belongs to \mathcal{G} (respectively \mathcal{O} or \mathcal{R}).

We write \mathcal{O}_p (resp. $\mathcal{G}_p, \mathcal{R}_p$) for the restriction of \mathcal{O} (resp. \mathcal{G}, \mathcal{R}) to the phases of process p . Clearly \mathcal{O}_p is either empty or contains a single phase which we denote as O_p . For any event e the *past of e* , written $\downarrow e$, is the set $\{f \mid f \leq e\}$. The past of a phase X is given by $\downarrow X = \{f \mid f \in \downarrow e \text{ for some } e \in X\}$. The following proposition identifies a special orange phase.

Proposition 12. *If \mathcal{O} is nonempty then there is a process p such that $\downarrow O_p$ contains no red events.*

Proof. Towards a contradiction, we assume that $\downarrow O_p$ contains a red event for all orange phases. Pick some p for which \mathcal{O}_p is nonempty and let e_0 be the maximal event in O_p . We construct a sequence $f_1, e_1, f_2, e_2, \dots$ as follows. f_{i+1} is some red event in $\downarrow e_i$ and e_{i+1} is the maximal orange event in the process of f_{i+1} . Clearly such an e_{i+1} must exist and $e_{i+1} < f_{i+1}$. Thus we have, $e_0 > f_1 > e_1 > f_2 > e_2 \dots$. Some e_i and e_j have to be on the same process with $i < j$. But then e_i and e_j are the same event, by definition, and this contradicts the fact that the sequence is strictly decreasing. \square

Using the phase O_p identified by Proposition 12, we refine χ by subdividing some of the orange phases so that the number of phases that are not green reduces by at least 1.

Let $X_0 = O_p$ and let e_0 be the maximal non-autonomous read event in X_0 (such an event must exist since O_p is part of a weak cycle). We construct a sequence of *orange* phases X_1, X_2, \dots and events $f_i, e_i \in X_i$ such that, $f_i \triangleright_{na} e_{i-1}$ and $e_i \rightarrow_\chi^* f_i$ is the maximal (if any) non-autonomous read below f_i in the phase X_i . The sequence terminates when there is no non-autonomous read $e_i \rightarrow_\chi^* f_i$ or if f_{i+1} is a green event (f_i and X_i cannot be red since $\downarrow X_0$ contains no red events). Notice that f_1 is orange since X_0 is part of a weak cycle and e_0 is the *maximal* non-autonomous read in X_0 .

Since, $e_0 \succ f_1 \succ e_1 \dots$ and each process has only one orange phase, if this sequence enters the same process for the second time that would yield a cycle, contradicting the assumption that χ is an acyclic phase decomposition. Thus the length of the sequence X_0, X_1, X_2, \dots is bounded by $|\mathbf{Procs}|$, the number of processes. Let the sequence of phases obtained as above be X_0, X_1, \dots, X_m where either e_m does not exist or f_{m+1} is a green event.

We now divide each of the phases X_i ($1 \leq i \leq m$) into (possibly) two parts, X_i^g and X_i' with $X_i^g \leq X_i'$ as follows. Set $X_i^g = X_i$ if there is no non-autonomous

read above f_i in X_i . If g_i is the first non-autonomous read above f_i in X_i then $X_i^g = \{e \in X_i \mid e < g_i\}$ and $X_i' = X_i \setminus X_i^g$. We also let $X_0^g = X_0$.

By construction, the parts constructed from X_i are all phases. Let χ' be this new phase decomposition. We shall show that, for $0 \leq i \leq m$, X_i^g is not part of any weak cycle in χ' .

Claim. If X_i^g is part of a weak cycle in the decomposition χ' then the previous phase on this cycle must be X_{i+1}^g .

This claim immediately implies that X_i^g is not part of a weak cycle in χ' since the processes of the phases X_0, X_1, \dots, X_m are pairwise distinct.

To prove the claim, let Y be the phase before X_i^g on a weak cycle in χ' . We find $u \triangleright_{na}^d v$ with $u \in Y$, $v \in X_i^g$ and d being the data-structure read non-autonomously by the phase X_i . Since X_i^g has a non-autonomous read, e_i is defined and is the maximal non-autonomous read in X_i^g . Hence, we have $u < v \rightarrow^* e_i \leq e_0$. Since there are no red events in the past of e_0 and Y is part of a weak cycle, Y and u must be orange. Now, d cannot be a stack or self-queue. Otherwise, Y and X_i^g would be on the same process and since both are orange we would get $Y = X_i^g$. But in this case $u \triangleright_{na} v$ induces a self-loop on phase X_i , a contradiction with acyclicity of χ . Therefore, d is a queue and since $f_{i+1} \triangleright_{na}^d e_i$ we deduce that $u \rightarrow^* f_{i+1}$. Since $u \in Y$ is orange, so is $f_{i+1} \in X_{i+1}^g$ and since they are both on the same process we deduce that $Y = X_{i+1}^g$, which concludes the proof of the claim.

The phases that were left untouched by the subdivision which further were not part of a weak cycle in χ continue to enjoy the same property in χ' . This is a simple graph theoretic property. In any graph, if we divide a vertex v as v_1 and v_2 , assign every edge incident on v to v_1 or v_2 or both, any cycle in the new graph is also a cycle in the old graph (simply replace v_i by v).

In summary, in the new acyclic decomposition χ' , phase $X_0 = X_0^g$ that was orange in χ and becomes green in χ' , green phases of χ are unchanged and continue to be green in χ' , some of the orange phases X_1, \dots, X_m are split into 2 phases, of which the first one is green in χ' and the others are orange or green in χ' . Thus, total number of non-green phases in χ' is strictly less than the number of non-green phases in χ . Notice that χ' has at most $|\mathbf{Procs}| - 1$ more phases than χ since $m < |\mathbf{Procs}| - 1$. Starting with a k -acyclic decomposition, with at most $k \cdot |\mathbf{Procs}|$ (non-green) phases, and iterating this procedure at most $k \cdot |\mathbf{Procs}|$ times, we end up with a strongly acyclic decomposition with at most $k \cdot |\mathbf{Procs}|^2$ phases. We have proved:

Theorem 13. *Let \mathcal{M} be a acyclic k -phase SQMSC. Then, \mathcal{M} admits a strongly acyclic phase decomposition of size at most $k \cdot |\mathbf{Procs}|^2$.*

6.2 Strongly Acyclic to Tree-like Phase Decompositions

A *tree-like* phase decomposition of an SQMSC \mathcal{M} is a strongly acyclic phase decomposition such that the in-degree of any vertex in the residual graph is at

most 1. Next, we refine any strongly acyclic phase decomposition into a tree-like phase decomposition. This transformation results in an exponential increase in the number of phases.

Lemma 14. *If \mathcal{M} admits a k -SAPD then it also admits a tree-like 2^{k-1} -APD.*

Proof. Let χ be a phase decomposition whose residual graph is strongly acyclic. Let X_1, \dots, X_k be a topological ordering of the phases.

We will split each phase X_i into $f(i)$ many new components (some of them may be empty), call them $X_{i,1}, \dots, X_{i,f(i)}$, in such a way that all non-autonomous reads of the new phase $X_{i,j}$ have matching writes from a single (new) phase. Here, $f(1) = 1$ and $f(i) = 2^{i-2}$ for $i > 1$. Notice that the number of new phases preceding X_i in this procedure is $\sum_{j=1}^{i-1} f(j) = f(i)$.

This splitting is possible because all the non-autonomous reads are from the same data-structure, which is either a stack or a queue.

Suppose phase X_i has non-autonomous reads from a queue. Then, we split X_i into components $X_{i,1}X_{i,2} \dots X_{i,f(i)}$ in such a way that $X_{i,j}$ contains all non-autonomous reads from the j -th *new* component (which is X_1 if $j = 1$ and $X_{k,\ell}$ with $j = 2^{k-2} + \ell$ if $j > 1$). We may also assume that autonomous phases of X_i are not split. This is achieved for instance if we inductively define $X_{i,j}$ as the smallest factor following $X_{i,j-1}$ containing all non-autonomous reads from the j -th new component. Notice that each new component is indeed a phase.

Suppose phase X_i has non-autonomous reads from a stack. Due to the LIFO policy on stacks, we split X_i into components $X_{i,f(i)} \dots X_{i,2}X_{i,1}$ in such a way that $X_{i,j}$ contains all non-autonomous reads from the j -th *new* component. Again, each new component is a phase.

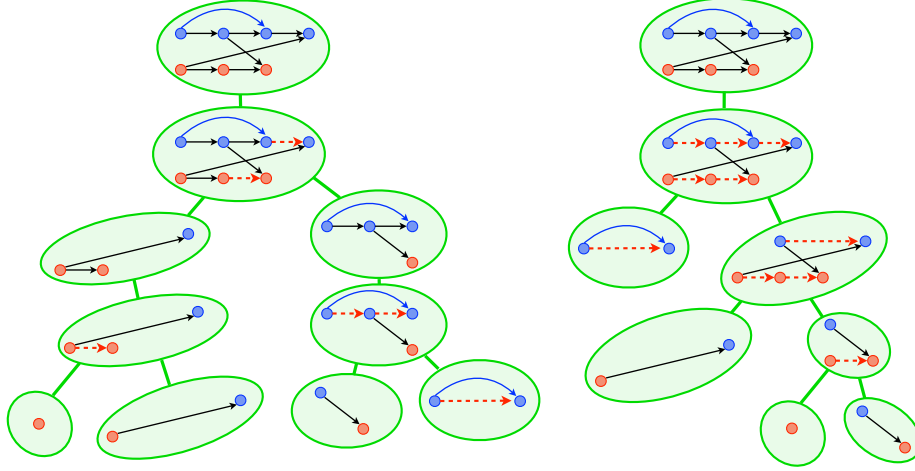
We can easily check that the associated residual graph is acyclic, has in-degree one, and has at most $f(k+1) = 2^{k-1}$ vertices. \square

6.3 Split-width and Decidability

We briefly recall the notion of split-width here. The idea of split-width was first described in [6] for multiply nested words and subsequently extended to include behaviours of systems with stacks and queues in [5, 7].

The idea is to decompose each SQMSC (in [5, 7] the name MSCNs is used instead) into *atomic pieces*. We begin by removing some of the \rightarrow edges to create holes which we call *elastic edges*. This operation is called *split*. We call an SQMSC with elastic edges as a *split-SQMSC*. The number of elastic edges in split-SQMSC its *elasticity*.

After removing some \rightarrow edges it is possible that the entire split-SQMSC consists of two disjoint parts with only elastic edges connecting them. At this point we may break-up this split-SQMSC into these two parts and then continue decomposing them separately. This operation is called *divide*. Our aim is to use split and divide repeatedly until we are left with the atomic parts, which are either singleton internal events (with elasticity 0) or an edge of the form $e \triangleright_d f$ (with elasticity 1 if d is a stack and 0 if d is a queue). Figure 5a describes this



(a) – of width 2.

(b) – of width 5.

Fig. 5: Two decompositions of an SQMSC into atomic splits

process on a sample SQMSC where the elastic edges are dashed (and colored red).

For any such complete break up of an SQMSC, its width is the maximum elasticity of all the split-SQMSCs produced in this procedure. The break-up described in Figure 5a has width 2. There may be several ways of starting with an SQMSC and breaking it down into its atomic components. A different and somewhat more trivial decomposition with width 5 is described in Figure 5b.

The *split-width* of an SQMSC is the minimum width among all possible ways of breaking it up into its atomic parts using split and divide. It turns out that in [5–7] the formal definitions are actually in dual. Instead of decomposing we compose split SQMSCs. We start with the atomic split SQMSCs and then put them together using *shuffle* (written \sqcup) which is the dual of divide and *merge* (written \mathbb{M}) which is the dual of *split* to generate the SQMSC we would like to. This makes it possible to define all this as an algebra. In particular, we write $\overline{\mathcal{M}} = \mathbb{M}\overline{\mathcal{M}}'$ to mean that $\overline{\mathcal{M}}'$ is obtained by dropping some \rightarrow edges in $\overline{\mathcal{M}}$ and $\overline{\mathcal{M}} = \overline{\mathcal{M}}_1 \sqcup \overline{\mathcal{M}}_2$ to mean that we can divide $\overline{\mathcal{M}}$ into $\overline{\mathcal{M}}_1$ and $\overline{\mathcal{M}}_2$.

A class C of SQMSCs have split-width k if each of its members has split-width at most k . Any class with bounded-split width that is MSO definable (or equivalently recognized by an SQDS) can be model-checked against MSO and other weaker logics. The precise complexity depends on the alternation depth of the MSO formula defining the class or the size of the SQDS (if any) recognizing it. In our case, the class k -Phase is recognizable by the controller of exponential size. Combining this with exponential bound on the split-width that we establish in the next section, we obtain double exponential complexity for the decidability for different temporal logics and (non-elementary) decidability for MSO. In particular, the reachability problem is decidable in time double exponential in

the parameter k , $|\mathbf{DS}|$ (number of the data-structures) and $|\mathbf{Procs}|$ (number of processes), but only polynomial in the size of the system. The precise complexity upperbounds of a variety of verification problems follow Table 4.5, and their optimality is argued in Section 4.4.1 and Section 11.4 of [5].

6.4 Split-width of SQMSCs in k -Phase

We complete our proof of the model-checking problems for k -Phase by proving the following theorem.

Theorem 15. *Any SQMSC \mathcal{M} in k -Phase has split-width at most $2^k - 1$.*

Given a k -phase decomposition of an SQMSC we may obtain a k -split SQMSC by dropping the \rightarrow edges between events in two adjacent phases. In the rest of this section, every component of the split-SQMSCs we construct will be phases and the associated residual graph is tree-like. Thus, we say $\overline{\mathcal{M}}$ is a tree-like m -APD to mean that it has at most m components (phases) per process. With this notation, from Theorem 13 and Lemma 14, it follows proving the following Lemma establishes Theorem 15.

Lemma 16. *Let $\overline{\mathcal{M}}$ be a tree-like m -APD, then the split-width of $\mathcal{M} = \mathbb{M}\overline{\mathcal{M}}$ is at most $2m - 1$.*

Proof. The proof is by induction on the size of \mathcal{M} . The base case is when all components of $\overline{\mathcal{M}}$ are singletons. In this case, $\overline{\mathcal{M}}$ is obtained by shuffling several basic split-SQMSCs. Doing so, the number of components never exceeds k . Hence the split-width in this case is at most $k - 1$.

For the inductive case, suppose $\overline{\mathcal{M}}$ has at least one non-trivial component. We will identify two nonempty split-SQMSCs $\overline{\mathcal{M}}_1$ and $\overline{\mathcal{M}}_2$ which are both tree-like m -APD and such that $\overline{\mathcal{M}} = \mathbb{M}(\overline{\mathcal{M}}_1 \sqcup \overline{\mathcal{M}}_2)$. The total number of components of the shuffle is at most $2m$, hence its elasticity is at most $2m - 1$. We can then easily conclude by induction.

Now we define the two split-SQMSCs $\overline{\mathcal{M}}_1$ and $\overline{\mathcal{M}}_2$. Let x_1, \dots, x_k with $k \leq m$ be a topological sorting of the components of $\overline{\mathcal{M}}$. Let x_i be the first non-trivial component. We have two cases to consider.

First, if there are two events e_1 and e_2 such that $e_1 \triangleright e_2$, e_1 is the first event of x_i and e_2 is the last event of x_i , then we let $\overline{\mathcal{M}}_1$ be $e_1 \triangleright e_2$ and $\overline{\mathcal{M}}_2$ be $\overline{\mathcal{M}}$ without $e_1 \triangleright e_2$. Clearly, $\overline{\mathcal{M}} = \mathbb{M}(\overline{\mathcal{M}}_1 \sqcup \overline{\mathcal{M}}_2)$ and both $\overline{\mathcal{M}}_1$ and $\overline{\mathcal{M}}_2$ define tree-like m -APD.

Otherwise, we show that x_i can be split into two nonempty components $x_i = x_i^1 \cdot x_i^2$ with no \triangleright -edges from x_i^1 to x_i^2 . Consider the first event e_1 of x_i . If e_1 is not a write event or if it is a write event with matching read not in x_i then we let $x_i^1 = e_1$ and x_i^2 be the corresponding suffix.

Otherwise, e_1 is a write event with matching read e_2 which is in x_i . Since we are not in the first case above, e_2 is not the last event of x_i . We let x_i^1 be the prefix of x_i up to and including e_2 , and x_i^2 be the corresponding suffix. Towards a contradiction, assume that there is \triangleright -edge from x_i^1 to x_i^2 . Let $f_1 \triangleright f_2$ with f_1 in

x_i^1 and $f_2 \in x_i^2$. Since the quotient graph of $\overline{\mathcal{M}}$ is acyclic, there are no self-loop on the component x_i in the quotient graph. Since x_i is a phase, the factor of x_i between f_1 and f_2 must be an autonomous computation. This is not possible since it contains the read event e_2 which is from a different stack.

We now propagate the splitting $x_i = x_i^1 x_i^2$ to the other components. For $j < i$, component x_j is a singleton, so it will not be split. If x_j is a write event with matching read in x_i^2 , we let $x_j^2 = x_j$ to put this component in $\overline{\mathcal{M}}_2$. Otherwise, we set $x_j^1 = x_j$ and it will be part of $\overline{\mathcal{M}}_1$.

We now explain inductively how to handle component x_j for $j > i$. By induction, we assume that for each $\ell < j$, either x_ℓ was not split in which case $x_\ell = x_\ell^b$ for some $b \in \{1, 2\}$, or x_ℓ was split in $x_\ell^b \cdot x_\ell^c$ with $\{b, c\} = \{1, 2\}$ and there are no \triangleright -edges from x_ℓ^b to x_ℓ^c .

If x_j has no incoming \triangleright -edges from a previous component, then we let $x_j^1 = x_j$ to put it in $\overline{\mathcal{M}}_1$. Otherwise, let x_ℓ with $\ell < j$ be the unique previous component having \triangleright -edges to x_j (ℓ is unique since we started with a *tree-like* APD). If x_ℓ was not split, i.e., $x_\ell = x_\ell^b$ with $b \in \{1, 2\}$, then we put x_j in $\overline{\mathcal{M}}_b$ by setting $x_j^b = x_j$. If x_ℓ was split but x_j has incoming edges only from one of the two new components, say x_ℓ^b , we proceed as above setting $x_j = x_j^b$.

The difficult case is when x_ℓ was split in $x_\ell^b \cdot x_\ell^c$ with $\{b, c\} = \{1, 2\}$ and x_j has incoming \triangleright -edges from both x_ℓ^b and x_ℓ^c . There are two cases to consider, depending on whether the data structure d of the \triangleright -edges from x_ℓ to x_j is a queue or a stack.

If $d \in \mathbf{Queues}$ then we split x_j in $x_j^b \cdot x_j^c$ with x_j^b being the shortest prefix of x_j containing all read events on queue d coming from x_ℓ^b . Note that there are no \triangleright -edges from x_j^b to x_j^c otherwise such a \triangleright -edge should be in an autonomous subphase, which is not possible since it contains the last event of x_j^b which is a read from d .

If $d \in \mathbf{Stacks}$ then we split x_j in $x_j^c \cdot x_j^b$ with x_j^c being the shortest prefix of x_j containing all read events on stack d coming from x_ℓ^c . As above, we can show that there are no \triangleright -edges from x_j^c to x_j^b .

Finally, for $b \in \{1, 2\}$ we let $\overline{\mathcal{M}}_b$ be the split-SQMSC with x_j^b as its components. Observe that there are no \triangleright -edges between components of $\overline{\mathcal{M}}_b$ and those of $\overline{\mathcal{M}}_{3-b}$ and that $\overline{\mathcal{M}} = \mathbb{M}(\overline{\mathcal{M}}_1 \sqcup \overline{\mathcal{M}}_2)$. Also, $\overline{\mathcal{M}}_1$ and $\overline{\mathcal{M}}_2$ define tree-like APDs with at most k components each. This concludes the proof. \square

References

1. R. Alur and P. Madhusudan. Adding nesting structure to words. *Journal of the ACM*, 56:16:1–16:43, 2009.
2. M. F. Atig, B. Bollig, and P. Habermehl. Emptiness of multi-pushdown automata is 2ETIME-Complete. In *DLT*, LNCS 5257, pages 121–133. Springer, 2008.
3. L. Breveglieri, A. Cherubini, C. Citrini, and S. Crespi-Reghizzi. Multi-pushdown languages and grammars. *Int. J. Found. Comput. Sci.*, 7(3):253–292, 1996.
4. B. Courcelle. The expression of graph properties and graph transformations in monadic second-order logic. In G. Rozenberg, editor, *Handbook of Graph Grammars*, pages 313–400. World Scientific, 1997.

5. A. Cyriac. *Verification of Communicating Recursive Programs via Split-width*. PhD thesis, ENS Cachan, 2014. http://www.lsv.ens-cachan.fr/~cyriac/download/Thesis_Aiswarya_Cyriac.pdf.
6. A. Cyriac, P. Gustin, and K. Narayan Kumar. MSO decidability of multi-pushdown systems via split-width. In *CONCUR*, volume 7454 of *LNCS*, pages 547–561. Springer, 2012.
7. A. Cyriac, P. Gustin, and K. Narayan Kumar. Verifying Communicating Multi-pushdown Systems. Technical report, Jan. 2014. <http://hal.archives-ouvertes.fr/hal-00943690>.
8. B. Genest, D. Kuske, and A. Muscholl. A Kleene theorem and model checking algorithms for existentially bounded communicating automata. *Information and Computation*, 204(6):920–956, 2006.
9. B. Genest, A. Muscholl, H. Seidl, and M. Zeitoun. Infinite-state high-level MSCs: Model-checking and realizability. *Journal of Computer and System Sciences*, 72(4):617–647, 2006.
10. J. G. Henriksen, M. Mukund, K. Narayan Kumar, M. A. Sohoni, and P. S. Thiagarajan. A theory of regular MSC languages. *Inf. Comput.*, 202(1):1–38, 2005.
11. A. Heußner, J. Leroux, A. Muscholl, and G. Sutre. Reachability analysis of communicating pushdown systems. In *FOSSACS*, volume 6014 of *LNCS*, pages 267–281. Springer, 2010.
12. S. La Torre, P. Madhusudan, and G. Parlato. A robust class of context-sensitive languages. In *LICS*, pages 161–170. IEEE Computer Society, 2007.
13. S. La Torre, P. Madhusudan, and G. Parlato. Context-bounded analysis of concurrent queue systems. In *TACAS*, volume 4963 of *LNCS*, pages 299–314. Springer, 2008.
14. S. La Torre, P. Madhusudan, and G. Parlato. An infinite automaton characterization of double exponential time. In *CSL*, volume 5213 of *LNCS*, pages 33–48. Springer, 2008.
15. S. La Torre and M. Napoli. Reachability of multistack pushdown systems with scope-bounded matching relations. In *CONCUR*, volume 6901 of *LNCS*, pages 203–218. Springer, 2011.
16. P. Madhusudan and G. Parlato. The tree width of auxiliary storage. In T. Ball and M. Sagiv, editors, *POPL*, pages 283–294. ACM, 2011.
17. S. Qadeer and J. Rehof. Context-bounded model checking of concurrent software. In *TACAS*, volume 3440 of *LNCS*, pages 93–107. Springer, 2005.
18. W. Zielonka. Notes on finite asynchronous automata. *R.A.I.R.O. — Informatique Théorique et Applications*, 21:99–135, 1987.