Data Multi-Pushdown Automata

Parosh Aziz Abdulla
Uppsala University

C. Aiswarya
Chennai Mathematical Institute

Mohamed Faouzi Atig
Uppsala University

***Abstract*—We extend the classical model of multi-pushdown systems by considering systems that operate on a finite set of variables ranging over the natural numbers. The conditions on variables are defined via gap-order constraints that allow to compare variables for equality, or to check that the gap between the values of two variables exceeds a given natural number. Furthermore, each message inside a stack is equipped with a data item representing its value. When a message is pushed to the stack, its value may be defined by a variable. When a message is popped, its value may be copied to a variable. Thus, we obtain a system that is infinite in multiple dimensions, namely we have a number of stacks such that each stack may contain an unbounded number of messages each of which is equipped with a natural number. It is well-known that the verification of any non-trivial property of multi-pushdown systems is undecidable, even for two stacks and for a finite data-domain. In this paper, we show the decidability of the reachability problem for the classes of data multi-pushdown system that admit a bounded split-width (or equivalently a bounded tree-width). As an immediate consequence, we obtain decidability for several subclasses of data multi-pushdown systems. These include systems with single stacks, restricted ordering policies on stack operations, bounded scope, bounded phase, and bounded context switches.**

I. Introduction

In the last few years, a lot of efforts have been devoted to the verification of *discrete* program models that have *infinite* state spaces such as Petri nets, lossy channel machines and multi-pushdown systems. In particular, multi-pushdown systems have been extensively studied as a natural model for concurrent Boolean recursive programs. Unfortunately, multi-pushdown systems are in general Turing powerful, and hence all basic decision problems are undecidable for them [32]. To overcome the undecidability barrier, several subclasses of multi-pushdown systems have been proposed (e.g., [2], [5], [6], [10]–[13], [16], [25]–[27], [29]–[31], [34], [35]).

Bounded context-switch has been proposed in [31] as an adequate criterion for the verification of multi-pushdown systems. The idea is to restrict the analysis to executions that can be split into a given number of contexts where, in each context, pop and push operations are exclusive to one stack. This provides a good trade-off between the verification coverage and the computational complexity. One the one hand, the state space that can be explored is still unbounded. On the other hand, the context-bounded reachability problem is NP-complete.

In [26], La Torre et al. generalize the notion of context into phase. A phase is a sequence of operations in which at most one stack can be popped, though there is no restriction on the push operations. The bounded-phase restriction considers only executions of the system that can be split into a given number of phases. In this case, the phase-bounded reachability problem is decidable in double exponential time.

Another generalization of bounded context-switch is bounded scope [27] which restricts the analysis of the multi-pushdown system to those executions in which the number of context-switches between any push operation and its corresponding pop operation is bounded by a given number. This definition extends the notion of contexts in term of coverage while being orthogonal to the notion of phase. In [27], the scope-bounded reachability problem is shown to be PSPACE-complete.

Another way to obtain decidability is to impose a linear order on stacks [16]. Stack operations are constrained in such a way that any pop operation is only allowed on the first non-empty stack. In [10], the reachability problem is shown to be 2ETIME-complete when assuming this ordering policy on stack operations. Furthermore, imposing such a restriction strictly extends the notion of phases while being orthogonal to the notion of scope-boundedness.

In [22], [28], [30], a unified technique to reason about multi-pushdown systems under such restrictions is presented. The idea is to see an execution as a graph with extra edges relating push operations and their corresponding pop operations. Then, the authors prove that the graphs generated under these restrictions have bounded split-width (or equivalently bounded tree-width). As an immediate consequence of Courcelle's theorem [20], we get the decidability of the model-checking problem for multi-pushdown systems under these restrictions against monadic second order logic.

However, all these models assume a finite-state control, which means that the variables of the modelled programs are assumed to range over finite domains. Several extensions of (multi-)pushdown systems with data have been studied in e.g., [1], [8], [14], [17], [23]. Most of these extensions concern the case of multi-pushdown systems with one stack except the work presented in [14] where an extension of multi-pushdown systems with data has been proposed. In order to obtain decidability of the reachability problem, the model requires the strong assumption of *data freshness*, and the restriction of the stack accesses to the bounded phase policy. Furthermore, the variable operations are restricted to only checking equality or disequality.

In this paper, we consider an extension of multi-pushdown systems, which we call *Data Multi-Push-Down Automata* (DMPDA), that strengthens the classical model in two ways. First in addition to stacks, a DMPDA uses a finite set of variables ranging over the natural numbers. Moreover, each

message inside the stack is equipped with a natural number which represents its value. Thus, we obtain a model that is possibly unbounded in multiple dimensions, namely we have a number of stacks such that each stack may contain an unbounded number of messages each of which is equipped with a natural number. The operations allowed on variables are defined by the *gap-order* constraint system [18], [33]. More precisely, DMPDA allow to compare the values of variables for equality, or to check that the gap between the values of two variables exceeds a given natural number. Also, a variable may be assigned a new arbitrary value, the value of another variable, or a value that is larger than at least a given natural number than the current value of another variable. Furthermore, a push operation may copy the value of a variable to the pushed message, and a pop operation may copy the value attached to the popped message to a variable. In this manner, the model of DMPDA subsumes two basic models, namely multi-pushdown systems (that we get by removing the variables and neglecting the values associated to the pushed messages) and the model of *integral relational automata* [18] (that we get by removing all the stacks).

Our main result is the decidability of the reachability problem for the classes of DMPDA that admit a bounded split-width. To that aim, we solve a more general problem, namely we characterize the *reachability relation* on variables between each pair of control states. More precisely, we present an algorithm for computing a finite set of gap-order formulas whose denotations describe values of variables that allow to reach one state from another with empty stacks. The main ingredient of the algorithm is a symbolic representation, called *traces*, that encode certain transition sequences in the automaton. A trace represents a set of *partial runs*. A partial run does not record the contents of the stacks, but marks positions inside the run that correspond to matching push/pop operations. Furthermore, a partial run is not contiguous in the sense that it may contain a number of "holes". Our algorithm will characterize the relation between the variables at the points where the holes occur. In particular, a partial run with no holes corresponds to a concrete run that starts and ends with empty stacks. The definition of partial runs allows to extend naturally the notion of *split width* [22] that has been considered for the analysis of multi-pushdown systems (without data). Intuitively, a run has a bounded split width if it can be built from atomic runs by using a shuffle and a contraction operator without producing any intermediate runs with more holes than the given bound. An atomic run is one that consists either of a single transition, or a pair of matching push/pop transitions. We show that our algorithm is guaranteed to terminate for all classes of systems that generate runs with a bounded split width. As an immediate consequence, we obtain the decidability for several subclasses of multi-pushdown systems with data including the ones that restrict the ordering policy on stack operations, or bound the scope, the number of phases, or the number of context switches.

Related work. Several subclasses of multi-pushdown systems have been proposed in the literature including bounded-context [31], bounded-phase [26], bounded scope [27] and ordered multi-pushdown systems [16]. The reachability problem for these classes has been shown to be decidable under the assumption of finiteness of the set of control states. These classes are subclasses of our model DMPDA and our decidability result subsumes the decidability of the reachability problem for these models. In contrast, we do not provide any complexity results in this paper.

Split-width and tree-width[1] have been used for showing, in a unified way, the MSO decidability of several classes of multi-pushdown systems [22], [28], [30]. The method has been extended for message passing systems [6] and parameterized message passing systems [24]. However the considered models are restricted to the manipulation of variables over finite data domains while in our model, variables range over natural numbers. In fact the results presented in [22], [28], [30] are orthogonal to our result since we do not consider the model-checking problem against monadic second order logic.

Decidability of the reachability problem for pushdown systems (i.e., multi-pushdown systems with one stack) with data has been extensively studied in the literature (see e.g., [1], [3], [15], [17], [19], [23]). The closest work is pushdown systems with gap-order constraints [1], which is subsumed by our model. Furthermore, the used techniques to show the decidability of the reachability problem for pushdown systems with gap-order constrains are different from the ones used in this paper.

Extensions of multi-pushdown systems with data have been studied in [14]. These require the strong assumption of freshness of data, and bounded phase restriction on stack accesses for decidability.

In [8] the split-width technique is lifted to analyze timed multi-pushdown system. Timed systems give rise to an infinite data domain. However, reachability in this case can be reduced to MSO model checking of untimed systems with finite propositional labelling indicating timing constraints, since realizability of a word with timing constraints can be expressed in MSO [7]. The crux of the decidability proof in all these cases is the use of tree-automata. In contrast, the reachability problem of DMPDA under bounded split-width does not reduce to the Boolean case. Furthermore, our algorithm uses a fix-point computation which terminates, thanks to well quasi-ordering of gap-order formulas.

II. PRELIMINARIES

Let $\mathbb{N}$ denote the set of natural numbers. For sets A and B, we use $f : A \to B$ to denote that f is a function from A to B. We use $f[a \leftarrow a']$ to denote the function f' such that $f'(a) = a'$, and $f'(x) = f(x)$ if $x \neq a$. For $A' \subseteq A$, we use $f \odot A'$ to denote the restriction of f to A'. For sets A_1 and A_2 with $A_1 \cap A_1 = \varnothing$, and functions $f_1 : A_1 \to B$ and $f_2 : A_2 \to B$, we use $f_1 \cup f_2 : A_1 \cup A_2 \to B$ to denote the

[1]Split-width and tree-width are not identical, but one is bounded if and only if the other is. Further the bounds are related linearly [22].

function g such that $g(a) = f_1(a)$ if $a \in A_1$ and $g(a) = f_2(a)$ if $a \in A_2$. For a finite set A, we use $|A|$ to denote the size of A.

For a set A, we use A^* to denote the set of finite words over A. We use ϵ to denote the empty word. For words $w_1, w_2 \in A^*$, we use $w_1 \cdot w_2$ to denote the concatenation of w_1 and w_2.

Consider a set A and a total ordering $\preceq$ on A. We use $a_1 \prec a_1$ to denote that $a_1 \preceq a_2$ and $a_1 \neq a_2$. We use $\prec$ to denote the induced immediate successor relation, i.e., $a_1 \prec a_2$ iff $a_1 \preceq a_2$ and there is no a_3 such that $a_1 \prec a_3 \prec a_2$. We define the function $\texttt{rank}_{\preceq} : A \mapsto \mathbb{N}$ such that $\texttt{rank}_{\preceq}(a) := |\{b | b \prec a\}|$, i.e., it gives the position of a in the total ordering.

III. Model

In this section, we introduce *Data Multi-Pushdown Automata* (DMPDA). A DMPDA operates on *multiple* unbounded stacks each of which allows pushing (appending) and popping (removing) messages in a last-in-first-out manner. In addition to the stacks, a DMPDA uses a finite set of variables ranging over the natural numbers.

The allowed operations on variables are defined by the *gap-order* constraint system [18], [33]. More precisely, the model allows non-deterministic value assignment, copying the value of one variable to another, and assignment of a value v to some variable such that v is larger than at least a given natural number than the current value of another variable. The transitions may be conditioned by tests that compare the values of two variables for equality, or that give the smallest allowed gap between two variables. In addition to carrying a name (taken from a finite alphabet), each message inside a stack is equipped by a natural number that represents its "value". A *push* operation may copy the value of variable to the pushed message, and a *pop* operation may copy the value of the popped message to a variable. Notice that DMPDA extend the classical model of Push-Down Automata in three ways, namely they allow (i) multiple stacks, (ii) numerical variables, and (iii) an infinite (numerical) stack alphabet.

Syntax: In the rest of the paper, we assume a finite set of variables $\mathbb{X}$, a finite set of stacks Σ, and a finite stack alphabet Γ. A DMPDA $\mathcal{A}$ is a tuple $\langle Q, \Delta \rangle$ where Q is a finite set of states, and Δ is a finite set of transitions. A transition is a triple $\langle q_1, op, q_2 \rangle$ where $q_1, q_2 \in Q$ are states, and op is an operation of one of the following forms: (i) nop is the *empty* operation that does not change the values of the variables or the contents of the stacks. (ii) $x \leftarrow *$ assigns non-deterministically an arbitrary value in $\mathbb{N}$ to the variable x. (iii) $x \leftarrow y$ copies the value of variable y to x. (iv) $x \leftarrow (>_c y)$ assigns non-deterministically to x a value that exceeds the current value of y by c (so the new value of x is $> y + c$). (v) $x = y$ checks whether the value of x is equal to the value of y. (vi) $x <_c y$ checks whether the gap between the values of y and x is larger than c. (vii) $push(\sigma)(a)(x)$ pushes the symbol $a \in \Gamma$ to the stack $\sigma \in \Sigma$ and assigns to it the value of the variable x. (viii) $pop(x)(\sigma)(a)$ pops the symbol $a \in \Gamma$ (if a is the top-most symbol at the stack $\sigma \in \Sigma$) and assigns its value to the variable x. We define the *source* $\texttt{src}(t) := q_1$ and the *target* $\texttt{tgt}(t) := q_2$.

We define $\Delta^{\texttt{intern}}$ to be the set of *internal* transitions, i.e., those that do not perform push or pop operations. We define $\Delta^{\texttt{push}}_{\sigma,a}$ to be the set of transitions whose operations are of the form $push(\sigma)(a)(x)$ for some $x \in \mathbb{X}$. We define $\Delta^{\texttt{push}}_{\sigma} := \cup_{a \in \Gamma} \Delta^{\texttt{push}}_{\sigma,a}$. We define $\Delta^{\texttt{pop}}_{\sigma,a}$ and $\Delta^{\texttt{pop}}_{\sigma}$ analogously.

Semantics: A DMPDA induces a transition system as follows. A *configuration* c is a triple $\langle q, \alpha, \beta \rangle$ where $q \in Q$ is a state, $\alpha : \mathbb{X} \mapsto \mathbb{N}$ defines the values of the variables, and $\beta : \Sigma \mapsto (\Gamma \times \mathbb{N})^*$ defines, for each stack $\sigma \in \Sigma$, its content $\beta(\sigma)$. The content of a stack is a word whose elements are of the form $\langle a, c \rangle$ where a is a symbol and c is its value. In particular, we define β_ϵ such that $\beta_\epsilon(\sigma) = \epsilon$ for all $\sigma \in \Sigma$. We say that c is *plain* if $\beta = \beta_\epsilon$.

We define the transition relation $\longrightarrow := \cup_{t \in \Delta} \xrightarrow{t}$, where $\xrightarrow{t}$ describes the effect of the transition t. The semantics of the transition relation is presented through the inference rules of Fig. 1, explained below one by one.

- nop. The values of the variables and the stack contents are not changed.
- $x \leftarrow *$. The value of the variable x is changed non-deterministically to some natural number. The values of the other variables and the stack contents are not changed.
- $x \leftarrow y$. The value of the variable y is copied to the variable x. The values of the other variables and the stack contents are not changed.
- $x \leftarrow (>_c y)$. The variable x is assigned non-deterministically a value that exceeds the value of y by c. The values of the other variables and the stack contents are not changed.
- $x = y$. The transition is only enabled if the value of x is equal to the value of y. The values of the variables and the stack contents are not changed.
- $x <_c y$. The transition is only enabled if the value of y is larger than the value of x by more than c. The values of the variables and the stack contents are not changed.
- $push(\sigma)(a)(x)$. The symbol a is pushed onto the stack σ with a value equal to that of x.
- $pop(x)(\sigma)(a)$. The symbol a is popped from the stack σ (if it is the top-most symbol of σ), and its value is copied to the variable x.

We use $\xrightarrow{*}$ to denote the reflexive transitive closure of $\longrightarrow$. A *run* π is an alternating sequence $c_0 t_1 c_1 \cdots c_{n-1} t_n c_n$ of configurations and transitions such that $c_{i-1} \xrightarrow{t_i} c_i$ for all $i : 1 \leq i \leq n$. We say that π is *plain* if c_0 and c_n are plain. For configurations c and c', we write $c \xrightarrow{\pi} c'$ to denote that there is a run π of the above form such that $c_0 = c$ and $c_n = c'$. Notice that $c \xrightarrow{*} c'$ iff $c \xrightarrow{\pi} c'$ for some run π.

Reachability: In the *Reachability Problem*, we are given two plain configurations c_1 and c_2, and are asked whether $c_1 \xrightarrow{*} c_2$. Observe that requiring that the two configurations of the reachability problem are plain is not a restriction since we can easily reduce the reachability problem for arbitrary configurations to the current definition of the problem. In

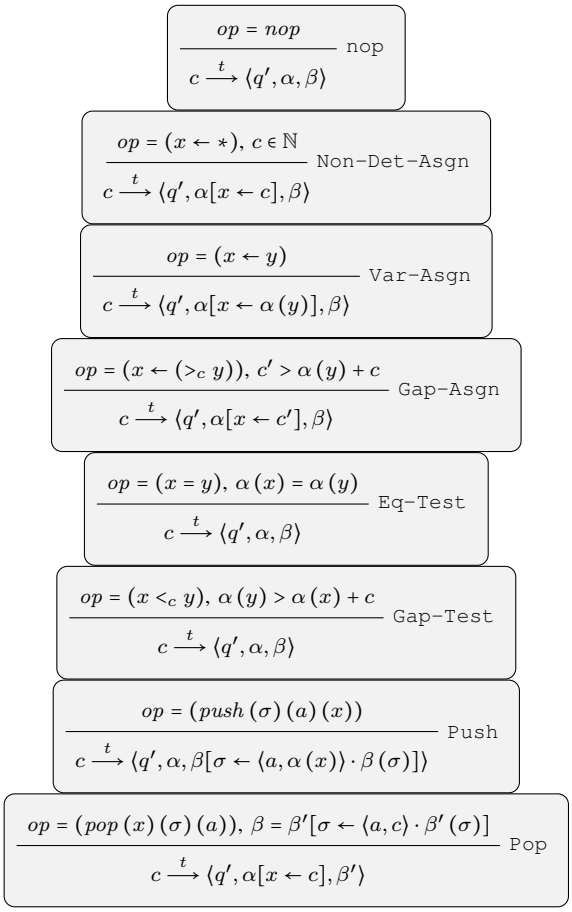

Fig. 1. Inference rules defining the relation $\xrightarrow{t}$, where $t = \langle q, op, q'\rangle$, starting from a configuration $c = \langle q, \alpha, \beta\rangle$.

order to solve the reachability problem we will study *reachability relations*. For states $q, q' \in Q$, we define $\mathbb{R}(q, q') := \left\{\langle \alpha, \alpha'\rangle \mid \langle q, \alpha, \beta_\epsilon\rangle \xrightarrow{*} \langle q', \alpha', \beta_\epsilon\rangle\right\}$. Intuitively, we summarize the values of the variables that allow us to move from q to q', starting and ending with empty stacks. More precisely, $\mathbb{R}(q, q')$ contains all pairs $\langle \alpha, \alpha'\rangle$ such that we can start from a configuration where the state is q, the values of the variables are given by α, and the stacks are empty to another configurations where the state is q', the values of the variables are given by α', and the stacks are empty again.

IV. Gap-Order Formulas

Fix a set $\mathbb{X}$ of variables. An *atomic gap-order formula* over $\mathbb{X}$ is either of the form $x = y$ or of the form $x <_c y$ where $x, y \in \mathbb{X}$ and $c \in \mathbb{N}$. A *gap-order formula* ϕ over $\mathbb{X}$ is a conjunction of atomic constraints over $\mathbb{X}$. Sometimes, we represent ϕ as a set (containing all its conjuncts). For a function $\mathtt{Val} : \mathbb{X} \mapsto \mathbb{N}$, we write $\mathtt{Val} \vDash \phi$ to denote that $\mathtt{Val}$ satisfies ϕ. We will also consider existentially quantified formulas of the form $\exists \mathbb{Y}.\ \phi$ where ϕ is a gap-order formula over $\mathbb{X}$, and $\mathbb{Y} \subseteq \mathbb{X}$. For $\mathtt{Val} : \mathbb{X} - \mathbb{Y} \mapsto \mathbb{N}$, we write $\mathtt{Val} \vDash \exists \mathbb{Y}.\ \phi$ to denote that there is a mapping $\mathtt{Val}' : \mathbb{Y} \mapsto \mathbb{N}$ such that $\mathtt{Val} \cup \mathtt{Val}' \vDash \phi$. For a (quantified) gap-order formula ϕ, we define its denotation $[\![\phi]\!] := \{\mathtt{Val} \mid \mathtt{Val} \vDash \phi\}$.

A gap-order formula ϕ over $\mathbb{X}$ is said to be in *normal form* if it satisfies the following conditions:

1) If $(x <_{c_1} y) \in \phi$ and $(y <_{c_2} z) \in \phi$ then $(x <_{c_3} z) \in \phi$ for some c_3 with $c_1 + c_2 < c_3$.
2) If $(x <_c y) \in \phi$ and $(y = z) \in \phi$ then $(x <_c z) \in \phi$.
3) If $(x <_c y) \in \phi$ and $(x = z) \in \phi$ then $(z <_c y) \in \phi$.
4) If $(x = y) \in \phi$ and $(y = z) \in \phi$ then $(x = z) \in \phi$.
5) For each $x, y \in \mathbb{X}$, there is at most one conjunct in ϕ containing both x and y.

Lemma 1 ([1], [4]): For each gap-order formula ϕ, we can effectively compute a gap-order formula ϕ' such that ϕ' is in normal form and $[\![\phi']\!] = [\![\phi]\!]$.

We obtain ϕ' from ϕ by repeatedly adding conjuncts to maintain properties 1-4 and removing conjuncts which violate property 5 (for instance, if we have both $(x <_{c_1} y) \in \phi$ and $(x <_{c_2} y) \in \phi$, with $c_1 \leq c_2$, then we can remove the former conjunct.) Normalization can be used to check consistency: the formula is consistent if and only if no inequalities of the form $x <_c x$ are generated.

Furthermore, we can use normalization to perform quantifier elimination as follows. For sets of variables $\mathbb{Y} \subseteq \mathbb{X}$ and a gap-order formula ϕ, let $\phi \ominus \mathbb{Y}$ to be the gap-order formula we get from ϕ by eliminating all conjuncts in which a variable $x \in \mathbb{Y}$ occurs.

Lemma 2 ([1], [4]): Suppose that ϕ is consistent and in normal form. Assume that $\mathtt{Val} \vDash \phi \ominus \mathbb{Y}$. There there is a $\mathtt{Val}' : \mathbb{Y} \to \mathbb{N}$ such that $\mathtt{Val} \cup \mathtt{Val}' \vDash \phi$

From Lemma 2 we get the following corollary.

Corollary 1: Suppose that ϕ is consistent and in normal form. Then $[\![\exists \mathbb{Y}.\ \phi]\!] = [\![\phi \ominus \mathbb{Y}]\!]$.

We write $\phi_1 \sqsubseteq \phi_2$ to denote that $[\![\phi_2]\!] \subseteq [\![\phi_1]\!]$. We can check $\phi_1 \sqsubseteq \phi_2$ as follows. By Lemma 1 we can assume that ϕ_1 and ϕ_2 are in normal form. Then the following conditions should be satisfied:

- If $(x <_{c_1} y) \in \phi_1$ then $(x <_{c_2} y) \in \phi_2$ for some $c_2 \geq c_1$.
- If $(x = y) \in \phi_1$ then $(x = y) \in \phi_2$.

V. Traces

We will introduce a data structure, called *traces*, that encode certain transition sequences, called *segments*. Roughly speaking, a segment is a sequence of consecutive transitions, i.e., the source of one transition in the sequence is identical to the target of the preceding transition in the segment. First, we give the definition of traces. Then, we define the set of atomic traces, and describe two operations that allow to build new traces from existing ones. Finally, we define an entailment relation on traces. In the rest of the section, fix a DMPDA $\mathcal{A} = \langle Q, \Delta\rangle$.

A. Definition

A trace τ is a tuple $\langle \mathbb{I}, \preceq, \mathtt{src}, \mathtt{tgt}, E, \phi\rangle$ defined as follows.

- $\mathbb{I}$ is a finite (index) set. Each index will be used to represent the summary of a segment. The summary is

given by the starting and the end states of each segment, i.e., the source of the first transition and the target of the last transition in the segment, and by a relation on the values of the variables before and after performing the different segments. For each variable $x \in \mathbb{X}$ and index $\mathtt{i} \in \mathbb{I}$, we will introduce two new variables $x^{\mathtt{i}}_{\mathtt{s}}$ and $x^{\mathtt{i}}_{\mathtt{t}}$ representing the *source* and *target* values of x, i.e., the value of x at the start and at the end of the segment represented by the index $\mathtt{i}$. We define the set $\mathbb{X}^{\mathtt{i}} := \{x^{\mathtt{i}}_{\mathtt{s}} | (x \in \mathbb{X}) \wedge (\mathtt{i} \in \mathbb{I})\} \cup \{x^{\mathtt{i}}_{\mathtt{t}} | (x \in \mathbb{X}) \wedge (\mathtt{i} \in \mathbb{I})\}$, and define $\mathbb{X}^{\mathbb{I}} := \cup_{\mathtt{i} \in \mathbb{I}} \mathbb{X}^{\mathtt{i}}$.
- $\preceq$ is a total ordering on $\mathbb{I}$ that gives the order in which the segments represented by the indices are performed. We let $\prec$ be the induced immediate successor relation (cf. Section II).
- $\mathtt{src} : \mathbb{I} \mapsto Q$ maps each index to a state representing the source of the corresponding segment, i.e., the state from which the segment starts (the source of the first transition in the segment.) Analogously, $\mathtt{tgt} : \mathbb{I} \mapsto Q$ defines the target of the segment, i.e., the state at the end of the segment (the target of the last transition in the segment.)
- $E : \mathbb{I} \times \mathbb{I} \mapsto 2^{\Sigma}$ is a function representing an "edge relation" between the indices. For indices $\mathtt{i}_1, \mathtt{i}_2$ the value of $E(\mathtt{i}_1, \mathtt{i}_2)$ gives the set of stacks such that there is a push operation in the segment represented by $\mathtt{i}_1$ whose corresponding pop operation is performed in the segment represented by $\mathtt{i}_2$. We impose two conditions on E. First, we require that $E(\mathtt{i}_1, \mathtt{i}_2) \neq \varnothing$ only if $\mathtt{i}_1 \prec \mathtt{i}_2$ since a pop operation can only occur after the corresponding push operation (and furthermore, we do not record push operations whose pop operations lie in the same segment). Second, we require that, for all stacks $\sigma \in \Sigma$, there are no indices $\mathtt{i}_1 \prec \mathtt{i}_2 \prec \mathtt{i}_3 \prec \mathtt{i}_4$ such that $\sigma \in E(\mathtt{i}_1, \mathtt{i}_3) \cap E(\mathtt{i}_2, \mathtt{i}_4)$. This condition ensures that we are consistent with the stack semantics since there is no overlap between two pairs of push/pop operations on the same stack.
- ϕ is gap-order formula over the set $\mathbb{X}^{\mathbb{I}}$, that defines the relation on values of the variables at the start and the end of the different segments.

We will equate traces that are equivalent modulo the renaming of the indices.

For a trace $\tau = \langle \mathbb{I}, \preceq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$, we define its *degree* $\#\tau = |\mathbb{I}|$, i.e., it is the size of the index set.

B. Atomic Traces

Atomic traces are built using the set of transitions. We will define two types of atomic traces, namely those induced by single internal transitions, and those that are induced by pairs of matching push/pop transitions.

Internal Transitions: Let $t = \langle q_1, op, q_2 \rangle \in \Delta^{\mathtt{intern}}$ be an internal transition. We will build a trace with a single segment which contains only one transition, namely t. Formally, we define $\mathtt{MkTrace}(t) := \langle \mathbb{I}, \preceq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$ where

- $\mathbb{I} = \{\mathtt{i}\}$, i.e., the set of indices is a singleton.
- $\preceq$ is trivial since the index set contains only one node.
- $\mathtt{src}(\mathtt{i}) = q_1$ and $\mathtt{tgt}(\mathtt{i}) = q_2$, i.e., we label the single index with the source and target states of t.
- $E = \varnothing$ reflecting the fact that the operation performed by t does not affect the stack.
- ϕ consists of all conjuncts of the following forms:
  - if $op = nop$ or $op = (x = y)$ or $op = (x <_c y)$ then $x^{\mathtt{i}}_{\mathtt{s}} = x^{\mathtt{i}}_{\mathtt{t}}$ for all $x \in \mathbb{X}$, i.e., the values of the variables are not changed during t.
  - if $op = (x = y)$ then $x^{\mathtt{i}}_{\mathtt{s}} = y^{\mathtt{i}}_{\mathtt{s}}$, and if $op = (x <_c y)$ then $y^{\mathtt{i}}_{\mathtt{s}} > x^{\mathtt{i}}_{\mathtt{s}} + c$. The values of the variables should satisfy the condition of the transition.
  - If $op = (x \leftarrow *)$ or $op = (x \leftarrow y)$ or $op = (x \leftarrow (>_c y))$ then $z^{\mathtt{i}}_{\mathtt{s}} = z^{\mathtt{i}}_{\mathtt{t}}$ for all $z \in \mathbb{X} - \{x\}$, i.e., the values of the variables different from x are not changed during t.
  - If $op = (x \leftarrow y)$ then $x^{\mathtt{i}}_{\mathtt{t}} = y^{\mathtt{i}}_{\mathtt{s}}$.
  - If $op = (x \leftarrow (>_c y))$ then $x^{\mathtt{i}}_{\mathtt{t}} > y^{\mathtt{i}}_{\mathtt{s}} + c$.

Stack Transitions: Consider transitions $t_1 = \langle q_1, op_1, q_2 \rangle \in \Delta^{\mathtt{push}}_{\sigma,a}$, $t_2 = \langle q_3, op_2, q_4 \rangle \in \Delta^{\mathtt{pop}}_{\sigma,a}$ where $op_1 = push(\sigma)(a)(x)$ and $op_2 = pop(y)(\sigma)(a)$. Notice that the two transitions push/pop the same symbol a to/from the same stack σ. We will build a trace with two segments containing t_1 resp. t_2. Formally, we define $\mathtt{MkTrace}(t_1, t_2) := \langle \mathbb{I}, \preceq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$, where:

- $\mathbb{I} = \{\mathtt{i}_1, \mathtt{i}_2\}$ i.e., the index set contains two elements. We use the indices $\mathtt{i}_1$ and $\mathtt{i}_2$ to represent two segments each containing a single transition, namely t_1 and t_2 respectively.
- $\mathtt{i}_1 \preceq \mathtt{i}_2$. We require that $\mathtt{i}_1$ is ordered before $\mathtt{i}_2$. This reflects the fact that a pop transition occurs after the matching push transition.
- $E(\mathtt{i}_1, \mathtt{i}_2) = \{\sigma\}$, i.e., we add an edge between $\mathtt{i}_1$ to $\mathtt{i}_2$ labeled with σ corresponding to the matching push/pop operations on σ performed by t_1 resp. t_2.
- $\mathtt{src}(\mathtt{i}_1) = q_1$, $\mathtt{tgt}(\mathtt{i}_1) = q_2$, $\mathtt{src}(\mathtt{i}_2) = q_3$, and $\mathtt{tgt}(\mathtt{i}_2) = q_4$. In other words, we label the new indices with the source and target states of t_1 resp. t_2.
- ϕ consists of all conjuncts of the following forms: (i) $z^{\mathtt{i}_1}_{\mathtt{s}} = z^{\mathtt{i}_1}_{\mathtt{t}}$ for all $z \in \mathbb{X}$, i.e., the values of the variables are not changed during t_1. (ii) $z^{\mathtt{i}_2}_{\mathtt{s}} = z^{\mathtt{i}_2}_{\mathtt{t}}$ for all $z \in \mathbb{X} - \{y\}$, i.e., the values of the variables, except y, are not changed during t_2. (iii) $y^{\mathtt{i}_2}_{\mathtt{t}} = x^{\mathtt{i}_1}_{\mathtt{s}}$. This condition corresponds to the fact the value of a when pushed to the stack during t_1 is equal to the value of variable x. This value is identical to the value stored in y after the pop operation of transition t_2.

C. Operations

We define two operations for building new traces.

Shuffling: Consider two traces $\tau_1 = \langle \mathbb{I}_1, \preceq_1, \mathtt{src}_1, \mathtt{tgt}_1, E_1, \phi_1 \rangle$, and $\tau_2 = \langle \mathbb{I}_2, \preceq_2, \mathtt{src}_2, \mathtt{tgt}_2, E_2, \phi_2 \rangle$, where $\mathbb{I}_1 \cap \mathbb{I}_2 = \varnothing$. We will build a new trace by shuffling the index sets of τ_1 and τ_2. We define $\tau_1 \otimes \tau_2$ to be the set of traces of the form $\langle \mathbb{I}, \preceq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$ satisfying the following conditions:

- $\mathbb{I} = \mathbb{I}_1 \cup \mathbb{I}_2$, i.e., the new trace contains exactly all the segments that are in τ_1 and τ_2.
- $\preceq=$ is a total ordering on $\mathbb{I}$ such that $\preceq_1 \subseteq \preceq$ and $\preceq_2 \subseteq \preceq$. We do not change the original orderings of the indices, but we do not constraint the places of the two sets of indices relative to each other.
- $\mathtt{src}(\mathtt{i}) = \mathtt{src}_1(\mathtt{i})$, and $\mathtt{tgt}(\mathtt{i}) = \mathtt{tgt}_1(\mathtt{i})$ for all $\mathtt{i} \in \mathbb{I}_1$. Furthermore, $\mathtt{src}(\mathtt{i}) = \mathtt{src}_2(\mathtt{i})$, and $\mathtt{tgt}(\mathtt{i}) = \mathtt{tgt}_2(\mathtt{i})$ for all $\mathtt{i} \in \mathbb{I}_2$. In other words, we keep the state labelings of the old indices.
- $E(\mathtt{i}_1, \mathtt{i}_2) = E_1(\mathtt{i}_1, \mathtt{i}_2)$ if $\mathtt{i}_1, \mathtt{i}_2 \in \mathbb{I}_1$, and $E(\mathtt{i}_1, \mathtt{i}_2) = E_2(\mathtt{i}_1, \mathtt{i}_2)$ if $\mathtt{i}_1, \mathtt{i}_2 \in \mathbb{I}_2$, i.e., all the edges in τ_1 and τ_2 are maintained in τ. Also, $E(\mathtt{i}_1, \mathtt{i}_2) = \varnothing$, if $\mathtt{i}_1 \in \mathbb{I}_1$ and $\mathtt{i}_2 \in \mathbb{I}_2$ or if $\mathtt{i}_1 \in \mathbb{I}_2$ and $\mathtt{i}_2 \in \mathbb{I}_1$, i.e., we do not add any edges between the two sets of indices. Furthermore, we require that there are no $\mathtt{i}_1, \mathtt{i}_2 \in \mathbb{I}_1$ and $\mathtt{i}_3, \mathtt{i}_4 \in \mathbb{I}_2$ such that (i) $\sigma \in E(\mathtt{i}_1, \mathtt{i}_2)$, (ii) $\sigma \in E(\mathtt{i}_3, \mathtt{i}_4)$, and (iii) either $\mathtt{i}_1 \prec \mathtt{i}_3 \prec \mathtt{i}_2 \prec \mathtt{i}_4$ or $\mathtt{i}_3 \prec \mathtt{i}_1 \prec \mathtt{i}_4 \prec \mathtt{i}_2$. This is to ensure that τ respects the stacks semantics.
- $\phi = \phi_1 \wedge \phi_2$. Notice that the values of the variables indexed by elements from $\mathbb{I}_1$ are not related to the values of the variables indexed by elements from $\mathbb{I}_2$.

Contraction: We define a *contraction operation* $\downarrow$ on indices that represents merging the corresponding segments. Consider $\mathtt{i}_1, \mathtt{i}_2 \in \mathbb{I}$ such that $\mathtt{i}_1 \prec \mathtt{i}_2$, i.e., $\mathtt{i}_2$ is the immediate successor of $\mathtt{i}_1$. Let $\mathtt{src}(\mathtt{i}_1) = q_1$, $\mathtt{tgt}(\mathtt{i}_1) = q_2$, $\mathtt{src}(\mathtt{i}_2) = q_2$, and $\mathtt{tgt}(\mathtt{i}_1) = q_3$, i.e., the target state of $\mathtt{i}_1$ is identical to the source state of $\mathtt{i}_2$. We will merge $\mathtt{i}_1$ and $\mathtt{i}_2$ to a new (single) index $\mathtt{j}$. We define $\tau \downarrow \langle \mathtt{i}_1, \mathtt{i}_2 \rangle := \langle \mathbb{I}', \preceq', \mathtt{src}', \mathtt{tgt}', E', \phi' \rangle$ defined as follows:

- $\mathbb{I}' = \mathbb{I} - \{\mathtt{i}_1, \mathtt{i}_2\} \cup \{\mathtt{j}\}$ where $\mathtt{j} \notin \mathbb{I}$, i.e., we replace the two merged indices by a new one.
- $\mathtt{k} \preceq' \mathtt{j}$ iff $\mathtt{k} \preceq \mathtt{i}_1$, and $\mathtt{j} \preceq' \mathtt{k}$ iff $\mathtt{i}_2 \preceq \mathtt{k}$ for all $\mathtt{k} \in \mathbb{I} - \{\mathtt{i}_1, \mathtt{i}_2\}$, i.e., in the new ordering, the new index $\mathtt{j}$ will take the places of the two (consecutive) indices $\mathtt{i}_1$ and $\mathtt{i}_2$. Furthermore, $\mathtt{k}_1 \preceq' \mathtt{k}_2$ iff $\mathtt{k}_1 \preceq \mathtt{k}_2$ for all $\mathtt{k}_1, \mathtt{k}_2 \in \mathbb{I} - \{\mathtt{i}_1, \mathtt{i}_2\}$, i.e., the relative orderings of the original indices are not changed.
- $\mathtt{src}'(\mathtt{j}) = q_1$ and $\mathtt{tgt}'(\mathtt{j}) = q_3$. Furthermore, $\mathtt{src}'(\mathtt{k}) = \mathtt{src}(\mathtt{k})$ and $\mathtt{tgt}'(\mathtt{k}) = \mathtt{tgt}(\mathtt{k})$ for all $\mathtt{k} \in \mathbb{I} - \{\mathtt{i}_1, \mathtt{i}_2\}$. In other words, we keep the state labelings of the old indices, while we take the source and target states of $\mathtt{j}$ to be the source state of $\mathtt{i}_1$ and the target state of $\mathtt{i}_2$ respectively.
- $E'(\mathtt{j}, \mathtt{k}) = E(\mathtt{i}_1, \mathtt{k}) \cup E(\mathtt{i}_2, \mathtt{k})$ and $E'(\mathtt{k}, \mathtt{j}) = E(\mathtt{k}, \mathtt{i}_1) \cup E(\mathtt{k}, \mathtt{i}_2)$ for all $\mathtt{k} \in \mathbb{I} - \{\mathtt{i}_1, \mathtt{i}_2\}$, i.e., we merge the edges originating from/to the two merged indices. Also, $E'(\mathtt{k}_1, \mathtt{k}_2) = E(\mathtt{k}_1, \mathtt{k}_2)$ for all $k_1, k_2 \in \mathbb{I} - \{\mathtt{i}_1, \mathtt{i}_2\}$, i.e., the edges to/from the other indices are maintained.
- $\phi' = \exists \left(\mathbb{X}^{\mathtt{i}_1} \cup \mathbb{X}^{\mathtt{i}_2}\right). \ \phi''$, where $\phi'' = \phi \wedge \left(\bigwedge_{x \in \mathbb{X}} \left(x_{\mathtt{t}}^{\mathtt{i}_1} = x_{\mathtt{s}}^{\mathtt{i}_2}\right)\right) \wedge \left(\bigwedge_{x \in \mathbb{X}} \left(x_{\mathtt{s}}^{\mathtt{i}_1} = x_{\mathtt{s}}^{\mathtt{j}}\right)\right) \wedge \left(\bigwedge_{x \in \mathbb{X}} \left(x_{\mathtt{t}}^{\mathtt{i}_2} = x_{\mathtt{t}}^{\mathtt{j}}\right)\right)$. We require that ϕ' is consistent. Since we merge the two segments corresponding to $\mathtt{i}_1$ and $\mathtt{i}_2$, we require the values of the variables at the end of the first segment to be consistent with the values of the variables at the start of the second segment. If these conditions are not satisfied, then the resulting formula ϕ' will not be consistent. Furthermore, the values of variables at the start of the new segment are defined to be the values of the variables at the start of the segment corresponding to $\mathtt{i}_1$. Analogously, the values of variables at the end of the new segment are defined to be the values of the variables at the end of the segment corresponding to $\mathtt{i}_2$. By Lemma 1, we know that ϕ'' can be transformed to an equivalent formula in normal form, and hence by Corollary 2, we can compute ϕ' as a gap-order formula.

We define $\tau \downarrow := \{\tau' \mid \exists \mathtt{i}_1, \mathtt{i}_2 \in \mathbb{I}.\ \mathtt{i}_1 \prec \mathtt{i}_2 \wedge \tau' \in \tau \downarrow \langle \mathtt{i}_1, \mathtt{i}_2 \rangle\}$.

D. Entailment

We define an entailment relation $\sqsubseteq$ on traces. Intuitively, a trace τ_1 is *weaker* than a trace τ_2, denoted $\tau_1 \sqsubseteq \tau_2$, if their graphs are isomorphic (equivalent up to the renaming of indices), but the gap-order formula of τ_1 is weaker than the one of τ_2. Later in the paper, when we compute reachability relations, we let τ_1 "subsume" τ_2 in the sense that if we encounter both τ_1 and τ_2 then we only include τ_1 (and discard τ_2), without suffering any loss of any precision in the analysis.

Formally, consider traces $\tau_1 = \langle \mathbb{I}_1, \preceq_1, \mathtt{src}_1, \mathtt{tgt}_1, E_1, \phi_1 \rangle$ and $\tau_2 = \langle \mathbb{I}_2, \preceq_2, \mathtt{src}_2, \mathtt{tgt}_2, E_2, \phi_2 \rangle$. Let $h : \mathbb{I}_1 \mapsto \mathbb{I}_2$ be a bijection. We write $\tau_1 \sqsubseteq_h \tau_2$ to denote that the following conditions are satisfied:

- $\mathtt{i}_1 \preceq_1 \mathtt{i}_2$ iff $h(\mathtt{i}_1) \preceq_2 h(\mathtt{i}_2)$.
- $\mathtt{src}_1(\mathtt{i}) = \mathtt{src}_2(h(\mathtt{i}))$ and $\mathtt{tgt}_1(\mathtt{i}) = \mathtt{tgt}_2(h(\mathtt{i}))$.
- $E_1(\mathtt{i}_1, \mathtt{i}_2) = E_2(h(\mathtt{i}_1), h(\mathtt{i}_2))$.
- $\phi_1 \sqsubseteq \phi_2^h$. We obtain ϕ_2^h from ϕ_2 through replacing each $x_{\mathtt{s}}^{\mathtt{i}}$ by $x_{\mathtt{s}}^{h(\mathtt{i})}$ and replacing each $x_{\mathtt{t}}^{\mathtt{i}}$ by $x_{\mathtt{t}}^{h(\mathtt{i})}$.

We use $\tau_1 \sqsubseteq \tau_2$ to denote that $\tau_1 \sqsubseteq_h \tau_2$ for some h.

VI. Algorithm

The algorithm inputs a DMPDA $\mathcal{A} = \langle Q, \Delta \rangle$ together with an upper limit θ on the degrees of the traces to be considered during the analysis. The algorithm maintains a set $\mathtt{W}$ of traces that have been detected but not analyzed, and a set $\mathtt{V}$ of traces that have been both detected and analyzed. Initially, the sets $\mathtt{W}$ and $\mathtt{V}$ are empty (Line 1–2). We add all the atomic traces induced by internal transitions (Line 3), and by matching push/pop transitions (Line 5) to the set $\mathtt{W}$. After the initialization phase, the algorithm performs a number of iterations using the repeat-loop of Line 8. In each iteration, we first select and remove an element τ from $\mathtt{W}$ (Lines 9–10). We check that τ satisfies two conditions (Line 11), namely: that (i) the degree of τ is within the allowed limit, and that (ii) τ is not subsumed by any trace already in the set $\mathtt{V}$. If the two conditions are satisfied, we use τ to generate new traces to analyze. These new traces are added to the set $\mathtt{W}$. First, we take the shuffle of τ with each member of the set $\mathtt{V}$ (Line 12). Then, we add all possible contractions of τ (Line 14). Finally, at Line 15, we add τ to $\mathtt{V}$, and at the same time remove all elements of $\mathtt{V}$ that are subsumed by τ. Notice that this means that all the traces in the set $\mathtt{V}$ will be pairwise incomparable

Algorithm 1: Computing the Reachability Relation.

Input: $\mathcal{A} = \langle Q, \Delta \rangle$: DMPDA,
θ: maximal index size
Output: characterization of the reachability relation

1 $\mathtt{V} \leftarrow \varnothing$;
2 $\mathtt{W} \leftarrow \varnothing$;
3 **for** *each* $t \in \Delta^{\mathtt{intern}}$ **do**
4 $\quad \mathtt{W} \leftarrow \mathtt{W} \cup \{\mathtt{MkTrace}(t)\}$
5 **for** *each* $t_1 \in \Delta^{\mathtt{push}}_{\sigma_1,a_1}$ *and* $t_2 \in \Delta^{\mathtt{pop}}_{\sigma_2,a_2}$ **do**
6 $\quad$ **if** $\sigma_1 = \sigma_2$ *and* $a_1 = a_2$ **then**
7 $\quad\quad \mathtt{W} \leftarrow \mathtt{W} \cup \{\mathtt{MkTrace}(t_1, t_2)\}$
8 **repeat**
9 $\quad$ select some $\tau \in \mathtt{W}$;
10 $\quad \mathtt{W} \leftarrow \mathtt{W} - \{\tau\}$;
11 $\quad$ **if** $(\#\tau \leq \theta) \wedge (\nexists\, \tau' \in \mathtt{V}.\ \tau' \sqsubseteq \tau)$ **then**
12 $\quad\quad$ **for** *each* $\tau' \in \mathtt{V}$ **do**
13 $\quad\quad\quad \mathtt{W} \leftarrow \mathtt{W} \cup (\tau \otimes \tau')$
14 $\quad\quad \mathtt{W} \leftarrow \mathtt{W} \cup \tau\downarrow$;
15 $\quad\quad \mathtt{V} \leftarrow \{\tau' \in \mathtt{V} \mid \tau \not\sqsubseteq \tau'\} \cup \{\tau\}$;
16 **until** $\mathtt{W} = \varnothing$;
17 **for** *each* $q, q' \in Q$ **do**
18 $\quad \mathcal{R}(q, q') \leftarrow \varnothing$
19 **for** *each* $\tau \in \mathtt{V}$ **do**
20 $\quad$ Let $\tau = \langle \mathbb{I}, \leq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$;
21 $\quad$ **if** $\#\tau = 1$ **then**
22 $\quad\quad$ let $\mathbb{I} = \{\mathtt{i}\}$, $\mathtt{src}(\mathtt{i}) = q$, $\mathtt{tgt}(\mathtt{i}) = q'$;
23 $\quad\quad \mathcal{R}(q, q') \leftarrow \mathcal{R}(q, q') \cup \{\phi\}$
24 **return** $\mathcal{R}$

wrt. $\sqsubseteq$. The iteration of the main loop is repeated until the set $\mathtt{W}$ becomes empty. After the termination of the loop, we build the reachability relations successively by going through all traces that have been added to $\mathtt{V}$ (Lines 17–23). For each pair of states q and q', we maintain a set $\mathcal{R}(q, q')$ of gap-order formulas. Each time a trace τ of degree one is encountered in $\mathtt{V}$, we add the gap-order formula of τ to the set $\mathcal{R}(q, q')$ corresponding the source state q and target state q' of the (only) index of τ. At the end of the algorithm, the reachability relation between any pair of states is characterized by the union of the denotations of all the gap-order formulas in the corresponding set.

VII. Partial Runs

In order to show correctness of our algorithm for computing reachability relations (cf. Section VI), we will introduce the notion of *partial runs*. Partial runs are different from (concrete) runs in two ways: (i) A configuration contains only the state and the values of the variables, but not the contents of the stacks. Instead, the partial run marks positions inside the run that correspond to matching push/pop transitions. (ii) A partial run is not contiguous in the sense that it consists of a sequence of *segments*. The target of the last transition in one segment need not be identical to the source of the first transition in the next segment. First, we give the formal definition of partial runs, and then we introduce a notion of *satisfiability* of a trace by a partial run, allowing us to view a trace as a symbolic encoding of sets of partial runs. In a similar manner to the case of traces (Section V), we define atomic partial runs, and define shuffling and contraction operations. Finally, we introduce a notion of *split width* for partial runs. In the rest of the section, fix a DMPDA $\mathcal{A} = \langle Q, \Delta \rangle$.

A. Definition

A *stackless configuration* b is a pair $\langle q, \alpha \rangle$ where $q \in Q$ is a state, and $\alpha : \mathbb{X} \mapsto \mathbb{N}$, i.e., it only defines the state and the values of the variables (without giving the stack contents). We define $\mathtt{state}(b) := q$. A partial run consists of a sequence of *segments*. A segment δ is a sequence $b_0 t_1 b_1 \cdots b_{m-1} t_m b_m$ of alternating elements, where b_i is a stackless configuration for all $i : 0 \leq i \leq m$, $t_i \in \Delta$ is a transition for all $i : 1 \leq i \leq m$, $\mathtt{state}(b_i) = \mathtt{src}(t_{i+1})$ for all $i : 0 \leq i < m$, and $\mathtt{state}(b_i) = \mathtt{tgt}(t_i)$ for all $i : 1 \leq i \leq m$. A *partial run* ρ is a sequence $[\delta_1][\delta_2]\cdots[\delta_n]$ where each δ_i is a segment. We require the segments to satisfy a number of conditions, as follows. Let δ_i be of the form $b_{i,0} t_{i,1} b_{i,1} \cdots b_{i,m_i-1} t_{i,m_i} b_{i,m_i}$, with $t_{i,j} = \langle q_{i,j-1}, op_{i,j}, q_{i,j} \rangle$ and $b_{i,j} = \langle q_{i,j}, \alpha_{i,j} \rangle$. We define $\mathtt{Dom}(\rho) := \{\langle i, j \rangle \mid (1 \leq i \leq n) \wedge (1 \leq j \leq m_i)\}$. Let $\leq_{\mathtt{lex}}$ be the lexicographic ordering on $\mathtt{Dom}(\rho)$, i.e., $\langle i_1, j_1 \rangle \leq_{\mathtt{lex}} \langle i_2, j_2 \rangle$ iff either $i_1 < i_2$ or $i_1 = i_2$ and $j_1 \leq j_2$. We require that the following two conditions satisfied:

- There exists a bijection $h : \mathtt{Dom}(\rho) \mapsto \mathtt{Dom}(\rho)$ such that:
  - If $t_{i_1,j_1} \in \Delta^{\mathtt{push}}_{\sigma,a}$, then $h(i_1, j_1) = \langle i_2, j_2 \rangle$ where $\langle i_1, j_1 \rangle <_{\mathtt{lex}} \langle i_2, j_2 \rangle$ and $t_{i_2,j_2} \in \Delta^{\mathtt{pop}}_{\sigma,a}$.
  - If $t_{i_2,j_2} \in \Delta^{\mathtt{pop}}_{\sigma,a}$, then $h(i_2, j_2) = \langle i_1, j_1 \rangle$ for some $\langle i_1, j_1 \rangle \in \Delta^{\mathtt{push}}_{\sigma,a}$ such $h(i_1, j_1) = \langle i_2, j_2 \rangle$.
  - If $t_{i,j} \in \Delta^{\mathtt{intern}}$, then $h(i, j) = \langle i, j \rangle$.
  - For each stack $\sigma \in \Sigma$, there are no $\langle i_1, j_1 \rangle <_{\mathtt{lex}} \langle i_2, j_2 \rangle <_{\mathtt{lex}} \langle i_3, j_3 \rangle <_{\mathtt{lex}} \langle i_4, j_4 \rangle$ such that $h(i_1, j_1) = \langle i_3, j_3 \rangle$ and $h(i_2, j_2) = \langle i_4, j_4 \rangle$.
- The stackless configurations satisfy the following conditions for each $i : 1 \leq i \leq n$ and $j : 1 \leq j \leq m_i$:
  - If $op_{i,j} = nop$ then $\alpha_{i,j} = \alpha_{i,j-1}$.
  - If $op_{i,j} = (x \leftarrow *)$ then $\alpha_{i,j}(y) = \alpha_{i,j-1}(y)$ for all $y \in \mathbb{X} - \{x\}$.
  - If $op_{i,j} = (x \leftarrow y)$ then $\alpha_{i,j} = \alpha_{i,j-1}[x \leftarrow \alpha_{i,j-1}(y)]$.
  - If $op_{i,j} = (x \leftarrow (>_c\ y))$ then $\alpha_{i,j} = \alpha_{i,j-1}[x \leftarrow \alpha_{i,j-1}(y) + c']$ for some $c' > c$.
  - If $op_{i,j} = (x = y)$ then $\alpha_{i,j-1}(x) = \alpha_{i,j-1}(y)$ and $\alpha_{i,j} = \alpha_{i,j-1}$.
  - If $op_{i,j} = (x <_c y)$ then $\alpha_{i,j-1}(y) > \alpha_{i,j-1}(x) + c$ and $\alpha_{i,j} = \alpha_{i,j-1}$.
  - If $op_{i,j} = (push(\sigma)(a)(x))$ then $\alpha_{i,j} = \alpha_{i,j-1}$.
  - If $op_{i,j} = (pop(y)(\sigma)(a))$ let $h(i, j) = \langle i', j' \rangle$. Let $op_{i',j'} = (push(\sigma)(a)(x))$. Then, $\alpha_{i,j} = \alpha_{i,j-1}[y \leftarrow \alpha_{i',j'}(x)]$.

Notice that the function h is unique. We define $\mathtt{Match}(\rho) := h$. We define *degree* of ρ by $\#\rho := n$, i.e., it is the number of segments in ρ.

B. Satisfiability

We relate traces and partial runs by defining a notion of *satisfiability*. Intuitively, a partial run ρ satisfies a trace τ if each segment of ρ can be "collapsed" to an index in τ. More precisely, each segment of ρ corresponds to an index in τ such that (i) the order of the indices in τ is identical to the order of the corresponding segments in ρ, (ii) the state labelings of the indices in τ are consistent with the states at the start and end of the corresponding segments in ρ, (iii) the placings of the matching push/pop pairs are consistent in τ and ρ, and (iv) the values of the variables at the end of the segments in ρ satisfy the gap-order formula in τ. Formally, consider a partial run ρ of the form of the previous subsection, and a trace $\tau = \langle \mathbb{I}, \leq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$ with $\#\tau = n$. Consider a bijection $h : \mathbb{I} \mapsto \{1, \ldots, n\}$ and a mapping $\mathtt{Val} : \mathbb{X}^{\mathbb{I}} \mapsto \mathbb{N}$ such that $\mathtt{Val} \vDash \phi$. We write $\rho \vDash_{h,\mathtt{Val}} \tau$ to denote that the following conditions are satisfied:

- $\mathtt{i} \lessdot \mathtt{j}$ iff $h(\mathtt{j}) = h(\mathtt{i}) + 1$.
- $\mathtt{src}(\mathtt{i}) = q_{h(\mathtt{i}),0}$.
- $\mathtt{tgt}(\mathtt{i}) = q_{h(\mathtt{i}),m_{h(\mathtt{i})}}$.
- $\sigma \in E(\mathtt{i},\mathtt{j})$ iff there are $k : 1 \leq k \leq m_{h(\mathtt{i})}$ and $\ell : 1 \leq \ell \leq m_{h(\mathtt{j})}$ such that $\mathtt{Match}(\rho)(h(\mathtt{i}), k) = \langle h(\mathtt{j}), \ell \rangle$ and $op_{h(\mathtt{i}),k} \in \Delta_\sigma^{\mathtt{push}}$ (equivalently $op_{h(\mathtt{j}),\ell} \in \Delta_\sigma^{\mathtt{pop}}$).
- $\mathtt{Val}\left(x_{\mathtt{s}}^{\mathtt{i}}\right) = \alpha_{h(\mathtt{i}),0}(x)$, and $\mathtt{Val}\left(x_{\mathtt{t}}^{\mathtt{i}}\right) = \alpha_{h(\mathtt{i}),m_{h(\mathtt{i})}}(x)$, for all $x \in \mathbb{X}$ and $\mathtt{i} \in \mathbb{I}$.

We write $\rho \vDash_h \tau$ to denote that $\rho \vDash_{h,\mathtt{Val}} \tau$ for some $\mathtt{Val}$, and write $\rho \vDash \tau$ to denote that $\rho \vDash_h \tau$ for some h. Furthermore, for $\mathtt{Val}$ with $\mathtt{Val} \vDash \phi$, we write $\rho \vDash_{\mathtt{Val}} \tau$ to denote that $\rho \vDash_{h,\mathtt{Val}} \tau$ for some h. Notice that if $\rho \vDash \tau$ then $\#\rho = \#\tau$.

C. Atomic Partial Runs

We build atomic partial runs either using single internal transitions or pairs of matching push/pop transitions.

Internal Transitions: Let $t = \langle q_1, op_1, q_2 \rangle \in \Delta^{\mathtt{intern}}$ be an internal transition. We define $\mathtt{MkPrun}(t)$ to be the set of partial runs of the form $[\langle q_1, \alpha_1 \rangle t \langle q_2, \alpha_2 \rangle]$ such that the following conditions are satisfied: (i) If $op = nop$ then $\alpha_2 = \alpha_1$. (ii) If $op = (x \leftarrow *)$ then $\alpha_2(y) = \alpha_1(y)$ for all $y \in \mathbb{X} - \{x\}$. (iii) If $op = (x \leftarrow y)$ then $\alpha_2 = \alpha_1[x \leftarrow \alpha_1(y)]$. (iv) If $op = (x \leftarrow (>_c y))$ then $\alpha_2 = \alpha_1[x \leftarrow c']$ for some $c' > \alpha_1(y) + c$. (v) If $op = (x = y)$ then $\alpha_1(x) = \alpha_1(y)$ and $\alpha_2 = \alpha_1$. (vi) If $op = (x <_c y)$ then $\alpha_1(y) > \alpha_1(x) + c$ and $\alpha_2 = \alpha_1$. In other words, the partial run consists of one segment. The segment contains two stackless configurations that precede and follow t, such that their states are consistent with the source and target states of t, and such that the values they assign to their variables are consistent with the operation of t.

Stack Transitions: Consider transitions $t_1 = \langle q_1, op_1, q_2 \rangle \in \Delta_{\sigma,a}^{\mathtt{push}}$, $t_2 = \langle q_3, op_2, q_4 \rangle \in \Delta_{\sigma,a}^{\mathtt{pop}}$ where $op_1 = push(\sigma)(a)(x)$ and $op_2 = pop(y)(\sigma)(a)$. We define $\mathtt{MkPrun}(t_1, t_2)$ to be the set of partial runs of the form $[\delta_1][\delta_2]$ such that the following conditions are satisfied: where:

- $\delta_1 = \langle q_1, \alpha_1 \rangle t_1 \langle q_2, \alpha_2 \rangle$.
- $\delta_2 = \langle q_3, \alpha_3 \rangle t_2 \langle q_4, \alpha_4 \rangle$.
- $\alpha_2 = \alpha_1$.
- $\alpha_4 = \alpha_3[y \leftarrow \alpha_2(x)]$.

In other words, the partial run consists of two segments. The first segment contains two stackless configurations that precede and follow t_1, such that their states are consistent with the source and target states of t_1. The second segment has a similar form with respect to transition t_2. Transition t_1 does not change the values of the variables. The difference between the values of the variables before/after t_2 is that variable y will get the same value as the value of variable x at the point when t_1 was executed, due to the matching push/pop operations.

D. Operations

We define two operations for building new partial runs.

Shuffling: Consider two partial runs $\rho_1 = [\delta_{1,1}] \cdots [\delta_{1,n_1}]$. and $\rho_2 = [\delta_{2,1}] \cdots [\delta_{2,n_2}]$. We define $\rho_1 \otimes \rho_2$ to be the set of partial runs of the form $[\delta_1] \cdots [\delta_{n_1+n_2}]$ such that there are functions $h_1 : \{\langle 1,1\rangle, \ldots, \langle 1,n_1\rangle\} \mapsto \{1, \ldots, n_1 + n_2\}$, and $h_2 : \{\langle 2,1\rangle, \ldots, \langle 2,n_2\rangle\} \mapsto \{1, \ldots, n_1 + n_2\}$ satisfying the following conditions:

- $h_1 \cup h_2$ is a bijection, i.e., each segment in ρ_1 or ρ_2 is represented by a unique segment in ρ.
- $h_k(i) \leq h_k(j)$ iff $i \leq j$ for $k = 1, 2$, i.e., we preserve the relative orderings of the segments in ρ_1 and ρ_2 but do not constraint the ordering between segments if they belong to different partial runs.

Notice that since we require that ρ is a partial run, the segments must be placed such that the stack semantics is preserved.

Contraction: We merge two consecutive segments in a partial run. Consider a partial run $\rho = [\delta_1] \cdots [\delta_n]$. For a partial run ρ', we write $\rho' \in \rho \downarrow$ to denote that ρ' is of the form $[\delta_1] \cdots [\delta_{k-1}][\delta'][\delta_{k+2}] \cdots [\delta_n]$ where (i) $\delta_k = \delta'_k \cdot b$, (ii) $\delta_{k+1} = b \cdot \delta'_{k+1}$, and (iii) $\delta' = \delta'_k \cdot b \cdot \delta'_{k+1}$. In other words, we require the consecutive segments k and $k+1$ to be consistent in the sense that the last stackless configuration in the former should be identical to the first stackless configuration in the latter. In such a case, we merge the two segments by merging these two configurations.

E. Split Width

We say that a partial run has *split width* θ if it can be derived by starting from the atomic partial runs and repeatedly applying the shuffling and contraction operations without letting any of the intermediately generated partial runs have a degree larger than θ.

A *proof tree* $\mathcal{T}$ of split width θ is a binary tree whose nodes are labeled with partial runs such that the following conditions are satisfied:

- The leafs are labeled with atomic partial runs.
- If an internal node, labeled with ρ, has two children then the children are labeled with partial runs ρ_1 and ρ_2 such that $\rho \in \rho_1 \otimes \rho_2$.
- If an internal node, labeled with ρ', has a single child then that child is labeled with a partial run ρ such that $\rho \in \rho' \downarrow$.
- For each label τ of a node in $\mathcal{T}$, we have $\#\tau \leq \theta$.

A partial run ρ is of *split width* θ if θ is the smallest number such that ρ is the root of a proof tree of split width θ. We use $\mathtt{SW}(\rho)$ to denote the split width of ρ.

We extend the notion of split width to (concrete) runs as follows. Consider a plain run $\pi = c_0 t_1 c_1 \cdots c_{n-1} t_n c_n$ where $c_i = \langle q_i, \alpha_i, \beta_i \rangle$. We define $\widehat{\pi}$ to be the partial run (of degree one) $[b_0 t_1 b_1 \cdots b_{n-1} t_n b_n]$ where $b_i = \langle q_i, \alpha_i \rangle$. We define $\mathtt{SW}(\pi) := \mathtt{SW}(\widehat{\pi})$. For states $q, q' \in Q$ and $\theta \in \mathbb{N}$, we define $\mathbb{R}^{\leq\theta}(q, q') := \left\{ \langle \alpha, \alpha' \rangle \mid \exists \pi. (\mathtt{SW}(\pi) \leq \theta) \wedge \left(\langle q, \alpha, \beta_\epsilon \rangle \xrightarrow{\pi} \langle q', \alpha', \beta_\epsilon \rangle \right) \right\}$.

VIII. Correctness

In this section, we show the correctness of the algorithm in Section VI (stated formally in Subsection VIII-E). We do that in several steps. We start by showing that the algorithm always terminates (Lemma 4). Then, we show soundness and completeness of the algorithm. More precisely, we will consider the set $\mathcal{R}(q, q')$ of gap-order formulas for each pair of states q and q', returned by the algorithm. We define a denotation function for these formulas, and relate them to the reachability relation $\mathbb{R}(q, q')$. We show that each member in the denotation corresponds to a concrete run (Lemma 8) implying the soundness of the algorithm. Conversely, we show that each run with split width θ belongs to the denotation (Lemma 19) implying the completeness of the algorithm.

In the rest of the section, we let $\mathtt{W}^i$ and $\mathtt{V}^i$ denote the values of the variables W and V at the start of the i^{th} iteration of the repeat loop of line 8 in Algorithm 1.

A. Termination

For a set A, a pre-order $\preceq$ on A is said to be a *Well Quasi-Ordering (WQO)* if, for each infinite sequence $a_0 a_1 a_2 \cdots$ of elements from A, there are $i < j$ such that $a_i \preceq a_j$. For a set of T of traces, we define its degree $\#T := \max_{\tau \in T}(\#\tau)$. Notice that $\#T$ need not exist in general.

Lemma 3: For any $k \in \mathbb{N}$ and any set of traces T with $\#T \leq k$, the entailment relation $\sqsubseteq$ is a WQO on T.

Proof: First, we recall that for any finite set of variables $\mathbb{X}$, the set of gap-order formulas over $\mathbb{X}$ is WQO under the $\sqsubseteq$ relation [1], [4]. Furthermore, we observe that, for traces τ_1 and τ_2, if $\tau_1 \sqsubseteq \tau_2$ then their graphs are equal up to renaming of the indices. Therefore, for any set T of traces such that $\#T \leq k$ for some k, the set of different trace graphs is finite. The result follows immediately. □

Lemma 4: The algorithm is guaranteed to terminate.

Proof: Assume that the algorithm does not terminate. Then, there is an infinite sequence of traces, with degrees bounded by θ, added to the set V. However, each time we add a new trace τ to V the test of Line 11 ensures that $\tau' \not\sqsubseteq \tau$ for every trace $\tau' \in \mathtt{V}$. Thus the sequence of elements added to V during the run of the algorithm will violate the WQO of sets of traces with bounded degree. □

B. Denotation

We provide a denotation for the sets in $\mathcal{R}$. A formula ϕ is said be *transitional* over $\mathbb{X}$ if ϕ is a gap-order formula over $\mathbb{X}^{\mathbb{I}}$ where $|\mathbb{I}| = 1$. Notice that, if $\mathbb{I} = \{\mathtt{i}\}$, then each variable in ϕ is either of the form $x^{\mathtt{i}}_{\mathtt{s}}$ or of the form $x^{\mathtt{i}}_{\mathtt{t}}$ where $x \in \mathbb{X}$. We define $\|\phi\|$ to be the set of pairs $\langle \alpha, \alpha' \rangle$ such that $\alpha : \mathbb{X} \mapsto \mathbb{N}$, $\alpha' : \mathbb{X} \mapsto \mathbb{N}$, and there is a $\mathtt{Val} : \mathbb{X}^{\mathbb{I}} \mapsto \mathbb{N}$ with $\mathtt{Val} \vDash \phi$, $\alpha = \mathtt{Val} \odot \left\{ x^{\mathtt{i}}_{\mathtt{s}} \mid x \in \mathbb{X} \right\}$ and $\alpha' = \mathtt{Val} \odot \left\{ x^{\mathtt{i}}_{\mathtt{t}} \mid x \in \mathbb{X} \right\}$. For a set of Φ of transitional gap-order formulas, we define $\|\Phi\| := \cup_{\phi \in \Phi} \|\phi\|$. Notice that all members of $\mathcal{R}(q, q')$ are transitional gap-order formulas.

C. Soundness

We show soundness of the algorithm by introducing a notion of *consistency* for traces. Consider a trace $\tau = \langle \mathbb{I}, \leq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$ and a function $\mathtt{Val} : \mathbb{X}^{\mathbb{I}}$ such that $\mathtt{Val} \vDash \phi$. We say that τ is *consistent* wrt. Val if there is a partial run ρ such that $\rho \vDash_{\mathtt{Val}} \tau$. We say that τ is *consistent* if τ is *consistent* wrt. all $\mathtt{Val} \in [\![\phi]\!]$. Notice that all atomic traces are consistent. The following two lemmas show that consistency is preserved when applying the shuffle and contraction operations.

Lemma 5: For traces τ_1, τ_2, and τ if τ_1 and τ_2 are consistent, and $\tau \in \tau_1 \otimes \tau_2$ then τ is consistent.

Lemma 6: For traces τ and τ', if τ is consistent and $\tau' \in \tau\downarrow$ then τ' is consistent.

We use Lemma 5 and Lemma 6 to show that all traces generated by the algorithm are consistent.

Lemma 7: For any trace τ and i, if $\tau \in \mathtt{W}^i \cup \mathtt{V}^i$, then τ is consistent.

Proof: By induction on i. If $\tau \in \mathtt{W}^0 \cup \mathtt{V}^0$ then, since $\mathtt{V}^0 = \varnothing$, we know that $\tau \in \mathtt{W}^0$. Therefore, by construction, τ is atomic, and hence τ is consistent. Assume that $\tau \in \left(\mathtt{W}^{i+1} \cup \mathtt{V}^{i+1}\right) - \left(\mathtt{W}^i \cup \mathtt{V}^i\right)$. This means that τ is added to W during the i^{th} iteration. We consider two cases. (i) If there are traces $\tau_1, \tau_2 \in \mathtt{W}^i \cup \mathtt{V}^i$ such that $\tau \in \tau_1 \otimes \tau_2$. By Lemma 5 it follows that τ is consistent. (ii) If there is a trace $\tau_1, \in \mathtt{W}^i \cup \mathtt{V}^i$ such that $\tau \in \tau_1\downarrow$. By Lemma 6 it follows that τ is consistent. □

We use Lemma 7 to show the soundness of our algorithm as follows.

Lemma 8: $\forall q, q' \in Q.\ \|\mathcal{R}(q, q')\| \subseteq \mathbb{R}(q, q')$

Proof: Suppose that $\langle \alpha, \alpha' \rangle \in \|\mathcal{R}(q, q')\|$, i.e., $\langle \alpha, \alpha' \rangle \in \|\phi\|$ for some $\phi \in \mathcal{R}(q, q')$. Since $\phi \in \mathcal{R}(q, q')$ we know that there is a trace $\tau = \langle \mathbb{I}, \leq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle \in \mathtt{V}$ at the termination point of the repeat loop of line 8, such that $\mathbb{I} = \{\mathtt{i}\}$ for some index $\mathtt{i}$. Define $\mathtt{Val} : \mathbb{X}^{\mathtt{i}} \mapsto \mathbb{N}$ such that $\mathtt{Val}(x^{\mathtt{i}}_{\mathtt{s}}) := \alpha(x)$ and $\mathtt{Val}(x^{\mathtt{i}}_{\mathtt{t}}) := \alpha'(x)$ for all $x \in \mathbb{X}$. It follows that $\mathtt{Val} \vDash \phi$. By Lemma 7, we know that τ is consistent. Therefore, there is a partial run ρ with $\rho \vDash_{\mathtt{Val}} \tau$. Notice that $\#\rho = \#\tau = 1$. This means that ρ is of the form $[b_0 t_1 b_1 \cdots b_{m-1} t_m b_m]$ where b_i is of the form $\langle q_i, \alpha_i \rangle$, $q_0 = q$, $\alpha_0 = \alpha$, $q_m = q'$, and $\alpha_m = \alpha'$. Define the run $\pi := c_0 t_1 c_1 \cdots c_{m-1} t_m c_m$ where $c_i := \langle q_i, \alpha_i, \beta_i \rangle$, and β_i is defined as follows. $\beta_0(\sigma) := \beta_\epsilon$. If $t_{i+1} \in \Delta^{\mathtt{intern}}$ then $\beta_{i+1} = \beta_i$. If $t_{i+1} = push(\sigma)(a)(x)$ then $\beta_{i+1} := \beta_i[\sigma \leftarrow \langle a, \alpha_i(x) \rangle \cdot \beta_i(\sigma)]$. If $t_{i+1} = pop(x)(\sigma)(a)$ then $\beta_{i+1} := \beta_i[\sigma \leftarrow w]$ where β_i is of the form $\langle a, c \rangle \cdot w$. This operation is well-defined by the definition of partial runs. Also by the definition of partial runs it follows that $\beta_m = \beta_\epsilon$. This implies that $\langle q, \alpha, \beta_\epsilon \rangle \xrightarrow{\pi} \langle q', \alpha', \beta_\epsilon \rangle$ and hence $\langle \alpha, \alpha' \rangle \in \mathbb{R}(q, q')$. □

D. Completeness

To prove completeness of our algorithm, we first show some properties of traces. The following lemma shows satisfiability of traces by partial runs is preserved by the shuffle operation.

Lemma 9: If $\rho_1 \vDash \tau_1$, $\rho_2 \vDash \tau_2$, and $\rho \in \rho_1 \otimes \rho_2$, then there is a trace $\tau \in \tau_1 \otimes \tau_2$ such that $\rho \vDash \tau$.

The following lemma shows a similar result for the contraction operation.

Lemma 10: If $\rho \vDash \tau$ and $\rho \in \rho' \downarrow$ then there is a trace $\tau' \in \tau \downarrow$ such that $\rho' \vDash \tau'$.

The following lemma shows a "reverse-monotonicity" property for traces with respect to the shuffle operation.

Lemma 11: For traces τ_1, τ_2, τ_3, τ_4, and τ_5 if $\tau_5 \in \tau_1 \otimes \tau_2$, $\tau_3 \sqsubseteq \tau_1$, and $\tau_4 \sqsubseteq \tau_2$, then there is a trace τ_6 such that $\tau_6 \in \tau_3 \otimes \tau_4$ and $\tau_6 \sqsubseteq \tau_5$.

The following lemma shows a similar property for the contraction operation.

Lemma 12: For traces τ_1, τ_2, and τ_3, if $\tau_3 \in \tau_1 \downarrow$ and $\tau_2 \sqsubseteq \tau_1$, then there is a trace τ_4 such that $\tau_4 \in \tau_2 \downarrow$ and $\tau_4 \sqsubseteq \tau_3$.

Next, we show some properties of the sets $\mathtt{W}$ and $\mathtt{V}$ during the run of the algorithm. Below, we show that once a trace is generated then a weaker trace will always remain in one of the sets $\mathtt{W}$ or $\mathtt{V}$.

Lemma 13: If $\#\tau \leq \theta$ and $\tau \in \mathtt{W}^i \cup \mathtt{V}^i$, then for all $j \geq i$ there is a trace τ' such that $\tau' \sqsubseteq \tau$ and $\tau' \in \mathtt{W}^j \cup \mathtt{V}^j$.

Proof: The lemma follows from the fact that once a trace τ is added to $\mathtt{W}$ or $\mathtt{V}$ then it can only be removed (i) at Line 11, but then, since $\#\tau \leq \theta$, the set $\mathtt{V}$ already contains some τ' with $\tau' \sqsubseteq \tau$, or (ii) at Line 15, but then the algorithm adds some τ' to $\mathtt{V}$ where $\tau' \sqsubseteq \tau$. □

Below, we show that once a trace is added to $\mathtt{V}$ then a weaker trace will always remain in $\mathtt{V}$.

Lemma 14: If $\tau \in \mathtt{V}^i$ then for all $j \geq i$ there is a trace τ' such that $\tau' \sqsubseteq \tau$ and $\tau' \in \mathtt{V}^j$.

Proof: The lemma follows from the fact that once a trace τ is added to $\mathtt{V}$ then it can only be removed at Line 15, but then the algorithm adds some τ' to $\mathtt{V}$ where $\tau' \sqsubseteq \tau$. □

Below, we show that if a trace is generated then a weaker trace (possibly the generated trace itself) will eventually be added to $\mathtt{V}$.

Lemma 15: If $\#\tau \leq \theta$ and $\tau \in \mathtt{W}^i \cup \mathtt{V}^i$, then there is a $j \geq i$ and a trace τ' such that $\tau' \sqsubseteq \tau$ and $\tau' \in \mathtt{V}^j$.

Proof: Suppose that $\tau \in \mathtt{W}^i \cup \mathtt{V}^i$. If $\tau \in \mathtt{V}^i$ then we are done. Otherwise, since the algorithm terminates (by Lemma 4) and since $\mathtt{W} = \varnothing$ when the algorithm terminates, we know that τ will be selected from $\mathtt{W}$ during some iteration $k \geq i$. Then, since $\#\tau \leq \theta$, either τ is added to $\mathtt{V}$ or $\mathtt{V}$ already contains some τ' with $\tau' \sqsubseteq \tau$. In both cases, the lemma follows. □

The following lemma shows that if two traces are generated by the algorithm, then their shuffle (or possibly some weaker trace) will also be generated.

Lemma 16: If $\#\tau_1 \leq \theta$, $\#\tau_2 \leq \theta$, $\tau_1 \in \mathtt{W}^{i_1} \cup \mathtt{V}^{i_1}$, $\tau_2 \in \mathtt{W}^{i_2} \cup \mathtt{V}^{i_2}$, and $\tau \in \tau_1 \otimes \tau_2$, then there is a j and a trace τ' such that $\tau' \sqsubseteq \tau$ and $\tau' \in \mathtt{W}^j \cup \mathtt{V}^j$.

Proof: Since $\#\tau_1 \leq \theta$, $\#\tau_2 \leq \theta$, $\tau_1 \in \mathtt{W}^{i_1} \cup \mathtt{V}^{i_1}$ and $\tau_2 \in \mathtt{W}^{i_2} \cup \mathtt{V}^{i_2}$, it follows by Lemma 15 that there are j_1 and j_2 and traces τ_1' and τ_2' such that $\tau_1' \sqsubseteq \tau_1$, $\tau_2' \sqsubseteq \tau_2$, $\tau_1' \in \mathtt{V}^{j_1}$, and $\tau_2' \in \mathtt{V}^{j_2}$. Without loss of generality, we can assume that (i) there are no $j_1' < j_1$ and τ_1' such that $\tau_1' \sqsubseteq \tau_1$, and $\tau_1' \in \mathtt{V}^{j_1'}$, (ii) there are no $j_2' < j_2$ and τ_2' such that $\tau_2' \sqsubseteq \tau_2$, and $\tau_2' \in \mathtt{V}^{j_2'}$, and (iii) $j_1 < j_2$. This means that during iteration $j_2 - 1$, we will select some τ_2' from $\mathtt{W}$ such that (i) $\tau_2' \sqsubseteq \tau_2$, (ii) by Lemma 14, there is a trace τ_1' in $\mathtt{V}$ with $\tau_1' \sqsubseteq \tau_1$, and (iii) there is no $\tau_2'' \in \mathtt{V}$ with $\tau_2'' \sqsubseteq \tau_2'$. Therefore, the algorithm will add all elements of $\tau_1' \otimes \tau_2'$ to $\mathtt{W}$. In particular by Lemma 11 the algorithm will add some τ' with $\tau' \sqsubseteq \tau$ to $\mathtt{W}$. □

The following lemma shows a similar result for the contraction operation.

Lemma 17: If $\#\tau_1 \leq \theta$, $\tau_1 \in \mathtt{W}^i \cup \mathtt{V}^i$, and $\tau \in \tau_1 \downarrow$, then there is a j and a trace τ' such that $\tau' \sqsubseteq \tau$ and $\tau' \in \mathtt{W}^j \cup \mathtt{V}^j$.

Proof: The proof is similar to Lemma 16. The difference is that we need to use Lemma 12 (instead of Lemma 11). □

The following lemma shows that, for each partial run with split width than θ, we will generate a trace that it satisfies and whose degree is smaller than θ.

Lemma 18: For each partial run ρ with $\mathtt{SW}(\rho) \leq \theta$ there is an i and a trace τ such that $\#\tau \leq \theta$, $\rho \vDash \tau$, and $\tau \in \mathtt{W}^i \cup \mathtt{V}^i$.

Proof: We know that there is a proof tree $\mathcal{T}$ of split width θ whose root is labeled with ρ. We use induction on the depth of $\mathcal{T}$. For the base case, we consider any atomic partial run ρ. Let τ be an atomic trace such that $\rho \vDash \tau$ (we know that at least one such a trace exists.) By construction, we have $\tau \in \mathtt{W}^0$.

For the induction, step we consider two cases. If the current node has two children labeled by partial runs ρ_1 and ρ_2 such that $\rho \in \rho_1 \otimes \rho_2$. By the induction hypothesis, we know that there are traces τ_1 and τ_2 with $\#\tau_1 \leq \theta$ and $\#\tau_2 \leq \theta$, together with i_1 and i_2 such that $\rho_1 \vDash \tau_1$, $\rho_2 \vDash \tau_2$, $\tau_1 \in \mathtt{W}^{i_1} \cup \mathtt{V}^{i_1}$, and $\tau_2 \in \mathtt{W}^{i_2} \cup \mathtt{V}^{i_2}$. By Lemma 9 it follows that there is a trace $\tau \in \tau_1 \otimes \tau_2$ such that $\rho \vDash \tau$. By Lemma 16 it follows that there is a j and a trace τ' such that $\tau' \sqsubseteq \tau$ and $\tau' \in \mathtt{W}^j \cup \mathtt{V}^j$. Since $\rho \vDash \tau$ it follows that $\#\tau = \#\rho$ and hence $\#\tau \leq \theta$.

If the current node has a single child labeled by ρ_1 where $\rho \in \rho_1 \downarrow$, then the proof follows in a similar fashion (now, using Lemma 17 instead of Lemma 16). □

Lemma 19: $\forall q, q' \in Q.\ \|\mathcal{R}(q, q')\| \supseteq \mathbb{R}^{\leq\theta}(q, q')$

Proof: Suppose that $\langle \alpha, \alpha' \rangle \in \mathbb{R}^{\leq\theta}(q, q')$. This means that there is a plain run $\pi = c_0 t_1 c_1 \cdots c_{n-1} t_n c_n$ where $c_i = \langle q_i, \alpha_i, \beta_i \rangle$, $q_0 = q$, $\alpha_0 = \alpha$, $q_n = q'$, and $\alpha_n = \alpha$. Consider the partial run $\rho := \widehat{\pi}$. By definition, we now that $\mathtt{SW}(\rho) \leq \theta$. By Lemma 18, there is and i and a trace τ such that $\#\tau \leq \theta$, $\rho \vDash \tau$, and $\tau \in \mathtt{W}^i \cup \mathtt{V}^i$. By Lemma 13, we know that for all $j \geq i$ there is a trace τ' such that $\tau' \sqsubseteq \tau$ and $\tau' \in \mathtt{W}^j \cup \mathtt{V}^j$. Since $\mathtt{W} = \varnothing$ when the algorithm terminates, we know, from Lemma 14 and Lemma 15, that there is a $\tau'' \in \mathtt{V}$ such that $\tau'' \sqsubseteq \tau'$ when the algorithm terminates. Since $\rho \vDash \tau$ and $\tau'' \sqsubseteq \tau' \sqsubseteq \tau$, it follows that $\rho \vDash \tau''$. Let τ'' be of the form $\langle \mathbb{I}, \leq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$ where $\mathbb{I} = \{\mathtt{i}\}$. Since $\#\tau'' = \#\tau' = \#\tau = \#\rho = 1$, we know that $\phi \in \mathcal{R}(q, q')$. Since $\rho \vDash \tau''$, we know that $\rho \vDash_{\mathtt{Val}} \tau''$ where $\mathtt{Val} : \mathbb{X}^{\mathbb{I}} \mapsto \mathbb{N}$ is defined by $\mathtt{Val}(x^{\mathtt{i}}_{\mathtt{s}}) = \alpha(x)$ and $\mathtt{Val}(x^{\mathtt{i}}_{\mathtt{t}}) = \alpha'(x)$ for all $x \in \mathbb{X}$. This implies that $\langle \alpha, \alpha' \rangle \in \|\phi\|$. □

E. Main Theorem

For a DMPDA $\mathcal{A}$, we define the split width $\mathtt{SW}(\mathcal{A})$ to be the largest k such that there is a plain run π in $\mathcal{A}$ with $\mathtt{SW}(\pi) = k$. For a class $\mathcal{C}$ of DMPDA, we define $\mathtt{SW}(\mathcal{C})$ to be the largest k such that there is an an $\mathcal{A} \in \mathcal{C}$ with $\mathtt{SW}(\pi) = k$. We say that $\mathcal{C}$ has *bounded split width* if $\mathtt{SW}(\pi) = k$ for some $k \in \mathbb{N}$. From Lemma 4 Lemma 8, and Lemma 19, we get the following theorem.

Theorem 1: The reachability problem is decidable for any class of DMPDA with bounded split width.

IX. Applications

A. Pushdown automata

A data push-down automata is a DMPDA with a single stack. All plain runs of a data-push-down automaton have split-width bounded by 3 [22]. Thus, it follows from Theorem 1 that

Corollary 2: The reachability problem is decidable for data push-down automata.

B. Multi-push-down systems

As already mentioned in the introduction, control state reachability is undecidable even in a finite data setting for multi-push-down systems. Several under-approximation classes (cf. Introduction) have been proposed in the literature for regaining decidability in the finite data case. We recall their definitions below.

- **Bounded context-switch** [31] A context is a sequence of operations in which at most one stack is touched. A run is k-context bounded if it is the concatenation of at most k contexts.
- **Bounded phase** [26] A phase is a sequence of operations in which at most one stack can be popped, though there are not restrictions on pushes. An run is k-phase bounded if it is the concatenation of at most k phases. This subsumes the k- context bounded runs.
- **Ordered stacks** [9], [10]. In an ordered multipushdown system, the stacks are ordered linearly by priority. Further a stack may be popped only if all the stacks with higher priority are empty. An ordered-stack run respects the priority ordering in its stack operations.
- **Bounded scope** [27] A run of a multipushdown system is k-scope bounded if the number of context switches between a push and the corresponding pop is always bounded by k. This definition is slightly general than the original round-based definition of [27]. This subsumes the k- context bounded runs,but is orthogonal to k-phase bounded runs.

Runs falling into any of the above class have bounded split-width.

Fact 1 ([22]):

- Split-width of k-context bounded runs is at most $k + 2$.
- Split-width of k-phase bounded runs is at most 2^k.
- Split-width of k-scope bounded runs is at most $k + 2$.
- Split-width of ordered-stack runs is at most $2^{|\Sigma|}$.

Let U be an under-approximation class above. A U-run of a DMPDA $\mathcal{A}$ is a run that is in U. For the under-approximation U, let $\mathcal{C}_U$ be the class of DMPDA such that all runs of any DMPDA $\mathcal{A} \in \mathcal{C}_U$ are U-runs. By Fact 1, $\mathcal{C}_U$ has bounded split-width. Hence by Theorem 1, we get

Corollary 3: Reachability problem is decidable for class $\mathcal{C}_U$.

Let U be an under-approximation class above. The U-reachability problem asks, given a DMPDA $\mathcal{A}$ and two configurations c_1 and c_2, whether it is possible to reach c_2 from c_1 by a U-run.

If the input DMPDA $\mathcal{A}$ belongs to $\mathcal{C}_U$, then U-reachability problem is same as reachability problem, which is decidable by Corollary 3. However, an arbitrary $\mathcal{A}$ may have runs outside of U in general, but still having bounded split-width. Thus running our algorithm naively on any input $\mathcal{A}$ could say "yes" to a pair of configurations c_1 and c_2 even when they are not U-reachable.

In order to decide the U-reachability problem, given the class U and DMPDA $\mathcal{A}$, we will construct a new DMPDA $\mathcal{A}' \in \mathcal{C}_U$. The DMPDA $\mathcal{A}'$ in effect will enforce the semantic restriction of U-runs syntactically into the automaton $\mathcal{A}$. Thus runs of $\mathcal{A}'$ are precisely U-runs of $\mathcal{A}$. We will thus reduce the U-reachability problem in $\mathcal{A}$ to the reachability problem in $\mathcal{A}' \in \mathcal{C}_U$ which is decidable by Corollary 3.

More formally, we reduce the U-reachability problem to reachability problem in $\mathcal{C}_U$. The reduction depends on the under-approximation U. On input DMPDA $\mathcal{A}$, plain configurations c_1 and c_2, we will construct a DMPDA $\mathcal{A}'$ belonging to $\mathcal{C}_U$. Then we will compute a finite set of pairs of plain configurations c_1' and c_2' from c_1 and c_2. We check whether c_2' is reachable from c_1' in $\mathcal{A}'$ for at least one pair of computed plain configurations c_1' and c_2'. If this is the case, we conclude that c_2 is U-reachable from c_1 in $\mathcal{A}$. Otherwise, c_2 is not U-reachable from c_1 in $\mathcal{A}$.

We will now describe the reduction in detail for each of the under-approximation class U above. Due to lack of space, we include only the case of bounded-phase in the main text. The other cases are given in the appendix.

1) Reduction for k-phase bounded.: Given a DMPDA $\mathcal{A} = \langle Q, \Delta \rangle$ and a bound k on the number phases, we construct a new one $\mathcal{A}' = \langle Q', \Delta' \rangle$ where $Q' = Q \times \{1 \dots k\} \times \Sigma \cup Q \times \{0\}$. The state remembers, in addition to the state of $\mathcal{A}$, how many phases have been used so far, and the stack that is being popped from in the current phase. Thus a state of the form (q, i, σ) means that currently $\mathcal{A}$ would have been in state q, and in the ith phase, which is allowed to pop only from stack σ. The state $(q, 0)$ is used to start off a computation, where no stack has been popped yet. The transitions Δ' lifts Δ to smoothly extend to Q' while maintaining the intended semantics. For instance, if $\langle q_1, op, q_2 \rangle \in \Delta$, then $\langle (q_1, i, \sigma), op, (q_2, i, \sigma) \rangle \in \Delta'$ and $\langle (q_1, 0), op, (q_2, 0) \rangle \in \Delta'$ if op is an operation of the forms (i) to (vii) (cf. Section III). If $\langle q_1, op, q_2 \rangle \in \Delta$ and if op is of the form (viii), i.e, $pop\,(x)\,(\sigma)\,(a)$, then we have i) $\langle (q_1, i, \sigma), op, (q_2, i, \sigma) \rangle \in \Delta'$, ii) $\langle (q_1, i, \sigma'), op, (q_2, i + 1, \sigma) \rangle \in \Delta'$ if $\sigma' \neq \sigma$ and $i < k$, and, iii) $\langle (q_1, 0), op, (q_2, 1, \sigma) \rangle \in \Delta'$.

The DMPDA $\mathcal{A}'$ exhibits all and only executions of $\mathcal{A}$ in which the number of phases is bounded by k. Given the pair of plain configurations $c_1 = \langle q_1, \alpha_1, \beta_\epsilon \rangle$ and $c_2 = \langle q_2, \alpha_2, \beta_\epsilon \rangle$ of $\mathcal{A}$, we obtain the set of pairs of the form $(\langle q_1', \alpha_1, \beta_\epsilon \rangle, \langle q_2', \alpha_2, \beta_\epsilon \rangle)$ where $q_1' = (q_1, 0)$ and q_2' is either $(q_2, 0)$ or of the form $q_2' = (q_2, i, \sigma)$ for some i and σ.

If $c_2' = \langle q_2', \alpha_2, \beta_\epsilon \rangle$ is reachable from $c_1' = \langle q_1', \alpha_1, \beta_\epsilon \rangle$ in $\mathcal{A}'$ for one such computed pair, then indeed, c_2 is k-phase reachable from c_1 in $\mathcal{A}$. Conversely, if c_2 is k-phase reachable from c_1 in $\mathcal{A}$, then there exists c_2' and c_1' of the form described above such that c_2' is reachable from c_1' in $\mathcal{A}'$.

This concludes our reduction.

Corollary 4: k-phase reachability problem is decidable for any DMPDA.

We have similar reductions for each of the under-approximations described above (cf. Appendix). We get

Corollary 5: U-reachability problem is decidable for any DMPDA, for the under-approximations U described above.

X. CONCLUSIONS

We have studied the reachability problem for multi-pushdown systems with gap-order constraints. We provide an algorithm for solving the reachability problem. The algorithm is sound and complete for the classes of automata that have a bounded split-width.

For future work, we plan to consider lifting the framework to a more general setting of auxiliary storages which include queues and multi-sets. Furthermore, it would be interesting to consider the case of distributed processes.

REFERENCES

[1] P. A. Abdulla, M. F. Atig, G. Delzanno, and A. Podelski. Push-down automata with gap-order constraints. In *FSEN*, volume 8161 of *LNCS*, pages 199–216. Springer, 2013.

[2] P. A. Abdulla, M. F. Atig, O. Rezine, and J.Stenman. Budget-bounded model-checking pushdown systems. *Formal Methods in System Design*, 45(2):273–301, 2014.

[3] P. A. Abdulla, M. F. Atig, and J. Stenman. The minimal cost reachability problem in priced timed pushdown systems. In *LATA*, volume 7183 of *LNCS*, 2012.

[4] P. A. Abdulla and G. Delzanno. On the coverability problem for constrained multiset rewriting. In *Proc. AVIS'06, The fifth Int. Workshop on on Automated Verification of Infinite-State Systems*, 2006.

[5] C. Aiswarya, P Gastin, and K. Narayan Kumar. Controllers for the verification of communicating multi-pushdown systems. In *CONCUR*, volume 8704 of *LNCS*, pages 297–311, 2014.

[6] C. Aiswarya, P. Gastin, and K. Narayan Kumar. Verifying communicating multi-pushdown systems via split-width. In *ATVA'14*, volume 8837 of *LNCS*. Springer, 2014. To appear.

[7] S. Akshay, P. Gastin, V. Juge, and S. N. Krishna. personal communication.

[8] S. Akshay, P. Gastin, and S. N. Krishna. Analyzing timed systems using tree automata. In *CONCUR'16*, volume 59 of *LIPICS*, pages 27:1–27:14. Leibniz-Zentrum für Informatik, 2016.

[9] M. F. Atig. Global Model Checking of Ordered Multi-Pushdown Systems. In *FSTTCS 2010*, volume 8, pages 216–227, 2010.

[10] M. F. Atig, B. Bollig, and P. Habermehl. Emptiness of multi-pushdown automata is 2ETIME-complete. In *Proceedings of DLT'08*, volume 5257 of *LNCS*, pages 121–133. Springer, 2008.

[11] M. F. Atig, K. Narayan Kumar, and P. Saivasan. Adjacent ordered multi-pushdown systems. *Int. J. Found. Comput. Sci.*, 25(8):1083–1096, 2014.

[12] M. F. Atig, K. Narayan Kumar, and P. Saivasan. Acceleration in multi-pushdown systems. In *TACAS 2016*, volume 9636 of *LNCS*, pages 698–714. Springer, 2016.

[13] Mohamed Faouzi Atig. Model-checking of ordered multi-pushdown automata. *Logical Methods in Computer Science*, 8(3), 2012.

[14] B. Bollig, A. Cyriac, P. Gastin, and K. Narayan Kumar. Model checking languages of data words. In *FoSSaCS'12*, volume 7213 of *LNCS*, pages 391–405. Springer, March 2012.

[15] A. Bouajjani, R. Echahed, and R. Robbana. On the automatic verification of systems with continuous variables and unbounded discrete data structures. In *Hybrid Systems II*, volume 999 of *LNCS*, pages 64–85. Springer, 1994.

[16] L. Breveglieri, A. Cherubini, C. Citrini, and S. Crespi Reghizzi. Multi-push-down languages and grammars. *International Journal of Foundations of Computer Science*, 7(3):253–292, 1996.

[17] X. Cai and M. Ogawa. Well-structured pushdown systems. In *CONCUR 2013*, volume 8052 of *LNCS*, pages 121–136. Springer, 2013.

[18] K. Čerāns. Deciding properties of integral relational automata. In Abiteboul and Shamir, editors, *ICALP 94*, volume 820 of *LNCS*, pages 35–46. Springer Verlag, 1994.

[19] Lorenzo Clemente and Slawomir Lasota. Reachability analysis of first-order definable pushdown systems. In *CSL 2015,*, volume 41 of *LIPIcs*, pages 244–259. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.

[20] B. Courcelle. The expression of graph properties and graph transformations in monadic second-order logic. In Grzegorz Rozenberg, editor, *Handbook of Graph Grammars and Computing by Graph Transformations, Volume 1: Foundations*, pages 313–400. World Scientific, 1997.

[21] 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.

[22] A. Cyriac, P. Gastin, and K. Narayan Kumar. MSO decidability of multi-pushdown systems via split-width. In *CONCUR'12*, volume 7454 of *LNCS*, pages 547–561. Springer, 2012.

[23] F. S. de Boer, M. M. Bonsangue, and J. Rot. It is pointless to point in bounded heaps. *Sci. Comput. Program.*, 112:102–118, 2015.

[24] M. Fortin and P. Gastin. Verification of parameterized communicating automata via split-width. In *FoSSaCS'16*, volume 9634 of *LNCS*, pages 197–213. Springer, 2016.

[25] A. Heußner, J. Leroux, A. Muscholl, and G. Sutre. Reachability analysis of communicating pushdown systems. *Logical Methods in Computer Science*, 8(3), 2012.

[26] S. La Torre, P. Madhusudan, and G. Parlato. A robust class of context-sensitive languages. In *LICS'07*, pages 161–170. IEEE Computer Society Press, 2007.

[27] S. La Torre and M. Napoli. Reachability of multistack pushdown systems with scope-bounded matching relations. In *CONCUR 2011*, volume 6901 of *LNCS*, pages 203–218. Springer, 2011.

[28] S. La Torre and G. Parlato. Scope-bounded multistack pushdown systems: Fixed-point, sequentialization, and tree-width. In *FSTTCS 2012*, volume 18 of *LIPIcs*, pages 173–184. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.

[29] A. Lal and T.W. Reps. Reducing concurrent analysis under a context bound to sequential analysis. In *CAV*, volume 5123 of *LNCS*, pages 37–51. Springer, 2008.

[30] P. Madhusudan and G. Parlato. The tree width of auxiliary storage. In Thomas Ball and Mooly Sagiv, editors, *POPL*, pages 283–294. ACM, 2011.

[31] S. Qadeer and J. Rehof. Context-bounded model checking of concurrent software. In N. Halbwachs and L.D. Zuck, editors, *TACAS 2005*, volume 3440 of *LNCS*, pages 93–107. Springer, 2005.

[32] G. Ramalingam. Context-sensitive synchronization-sensitive analysis is undecidable. *ACM Trans. Program. Lang. Syst.*, 22(2):416–430, 2000.

[33] P. Revesz. A closed form evaluation for datalog queries with integer (gap)-order constraints. *Theoretical Computer Science*, 116(1):117–149, 1993.

[34] A. Seth. Global reachability in bounded phase multi-stack pushdown systems. In *CAV'10*, LNCS, 2010.

[35] S. La Torre, P. Madhusudan, and G. Parlato. Reducing context-bounded concurrent reachability to sequential reachability. In *CAV*, volume 5643 of *LNCS*, pages 477–492. Springer, 2009.

APPENDIX

PROOFS OF LEMMAS

Lemma 5

Let $\tau_1 = \langle \mathbb{I}_1, \preceq_1, \mathtt{src}_1, \mathtt{tgt}_1, E_1, \phi_1 \rangle$, $\tau_2 = \langle \mathbb{I}_2, \preceq_2, \mathtt{src}_2, \mathtt{tgt}_2, E_2, \phi_2 \rangle$, and $\tau = \langle \mathbb{I}, \preceq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$. Assume that τ_1 and τ_2 are consistent. Consider $\mathtt{Val} : \mathbb{X}^{\mathbb{I}} \mapsto \mathbb{N}$ where $\mathtt{Val} \vDash \phi$. Define $\mathtt{Val}_1 := \mathtt{Val} \odot \mathbb{X}^{\mathbb{I}_1}$, and $\mathtt{Val}_2 := \mathtt{Val} \odot \mathbb{X}^{\mathbb{I}_2}$. We observe that $\mathtt{Val}_1 \vDash \phi_1$ and $\mathtt{Val}_2 \vDash \phi_2$. Since τ_1 and τ_2 are consistent, there are partial runs $\rho_1 = [\delta_{1,1}] \cdots [\delta_{1,n_1}]$ and $\rho_2 = [\delta_{2,1}] \cdots [\delta_{2,n_2}]$, and bijections $h_1 : \mathbb{I}_1 \mapsto \{1, \ldots, n_1\}$, and $h_2 : \mathbb{I}_1 \mapsto \{1, \ldots, n_2\}$ such that $\rho_1 \vDash_{h_1, \mathtt{Val}_1} \tau_1$ and $\rho_2 \vDash_{h_2, \mathtt{Val}_2} \tau_2$. Define the partial run $\rho := [\delta_1] \cdots [\delta_{n_1+n_2}]$ where $\delta_i = \delta_{1,j}$ if $\mathtt{rank}_{\preceq}\left(h_1^{-1}(i)\right) = j$, and $\delta_i = \delta_{2,j}$ if $\mathtt{rank}_{\preceq}\left(h_2^{-1}(i)\right) = j$. Define $h : \mathbb{I} \mapsto \{1, \ldots, n_1 + n_2\}$ such that $h'(\mathtt{i}) = \mathtt{rank}_{\preceq}(\mathtt{i})$. It follows that $\rho \vDash_{h, \mathtt{Val}} \tau$ and hence $\rho \vDash_{\mathtt{Val}} \tau$.

Lemma 6

Let $\tau = \langle \mathbb{I}, \preceq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$, and $\tau' = \langle \mathbb{I}', \preceq', \mathtt{src}', \mathtt{tgt}', E', \phi' \rangle$, where $\mathbb{I}' = \mathbb{I} - \{\mathtt{i}_1, \mathtt{i}_2\} \cup \{\mathtt{j}\}$. Assume that τ is consistent. Consider $\mathtt{Val}' : \mathbb{X}^{\mathbb{I}'} \mapsto \mathbb{N}$ where $\mathtt{Val}' \vDash \phi'$. By Lemma 1, Lemma 2, and Corollary 1, there exists a $\mathtt{Val}'' : \mathbb{X}^{\mathtt{i}_1} \cup \mathbb{X}^{\mathtt{i}_2} \mapsto \mathbb{N}$ such that $\mathtt{Val}' \cup \mathtt{Val}'' \vDash \phi$. Define $\mathtt{Val} := (\mathtt{Val}' \cup \mathtt{Val}'') \odot \mathbb{X}^{\mathbb{I}}$. Notice that $\mathtt{Val} \vDash \phi$. Since τ is consistent, there is a a partial run $\rho = [\delta_1] \cdots [\delta_n]$ and a bijection $h : \mathbb{I} \mapsto \{1, \ldots, n\}$ such that $\rho \vDash_{h, \mathtt{Val}} \tau'$. Let $k = \mathtt{rank}_{\preceq'}(\mathtt{j})$. Define the partial run $\rho' := [\delta_1] \cdots [\delta'] \cdots [\delta_n] \in \rho \downarrow$. Define $h' : \mathbb{I}' \mapsto \{1, \ldots, n+1\}$ such that $h'(\mathbb{k}) = h(\mathbb{k})$ if $\mathbb{k} \prec' \mathtt{j}$, $h'(\mathtt{j}) = k$, and $h'(\mathbb{k}) = h(\mathbb{k}) - 1$ if $\mathbb{k} \prec' \mathtt{j}$. It follows that $\rho' \vDash_{h', \mathtt{Val}'} \tau'$ and hence $\rho' \vDash_{\mathtt{Val}'} \tau'$.

Lemma 9

Let $\tau_1 = \langle \mathbb{I}_1, \preceq_1, \mathtt{src}_1, \mathtt{tgt}_1, E_1, \phi_1 \rangle$, $\tau_2 = \langle \mathbb{I}_2, \preceq_2, \mathtt{src}_2, \mathtt{tgt}_2, E_2, \phi_2 \rangle$, $\rho_1 = [\delta_{1,1}] \cdots [\delta_{1,n_1}]$, $\rho_2 = [\delta_{2,1}] \cdots [\delta_{2,n_2}]$, and $\rho_2 = [\delta_1] \cdots [\delta_{n_1+n_2}]$. Since $\rho \in \rho_1 \otimes \rho_2$, there are $h_1 : \{1, \ldots, n_1\} \mapsto \{1, \ldots, n_1 + n_2\}$ and $h_2 : \{1, \ldots, n_2\} \mapsto \{1, \ldots, n_1 + n_2\}$ where $h_1 \cup h_2$ is a bijection and $h_k(i) \leq h_k(j)$ iff $i \leq j$ for $k = 1, 2$. Since $\rho_1 \vDash \tau_1$, there is a bijection $h_3 : \mathbb{I}_1 \mapsto \{1, \ldots, n_1\}$ and $\mathtt{Val}_3 : \mathbb{X}^{\mathbb{I}_1} \mapsto \mathbb{N}$ such that $\rho_1 \vDash_{h_3, \mathtt{Val}_3} \tau_1$, and since $\rho_2 \vDash \tau_2$, there is a bijection $h_4 : \mathbb{I}_2 \mapsto \{1, \ldots, n_2\}$ and $\mathtt{Val}_4 : \mathbb{X}^{\mathbb{I}_2} \mapsto \mathbb{N}$ such that $\rho_2 \vDash_{h_4, \mathtt{Val}_4} \tau_1$. Define $\tau := \langle \mathbb{I}, \preceq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$ to be the unique trace such that $\tau \in \tau_1 \otimes \tau_2$ and $\mathtt{rank}_{\preceq}(\mathtt{i}) = h_1(h_3(\mathtt{i}))$ if $\mathtt{i} \in \mathbb{I}_1$ and $\mathtt{rank}_{\preceq}(\mathtt{i}) = h_2(h_4(\mathtt{i}))$ if $\mathtt{i} \in \mathbb{I}_2$. Define $h : \mathbb{I} \mapsto \{1, \ldots, n_1 + n_2\}$ such that $h(\mathtt{i}) := h_1(h_3(\mathtt{i}))$ if $\mathtt{i} \in \mathbb{I}_1$ and $h(\mathtt{i}) := h_2(h_4(\mathtt{i}))$ if $\mathtt{i} \in \mathbb{I}_2$. Define $\mathtt{Val} := \mathtt{Val}_3 \cup \mathtt{Val}_4$. It follows that $\rho \vDash_{h, \mathtt{Val}} \tau$ and hence $\rho \vDash \tau$.

Lemma 10

Let $\tau = \langle \mathbb{I}, \preceq, \mathtt{src}, \mathtt{tgt}, E, \phi \rangle$, $\rho = [\delta_1] \cdots [\delta_n]$, $\rho' = [\delta_1] \cdots [\delta_{k-1}] [\delta'] [\delta_{k+2}] \cdots [\delta_n]$. Since $\rho \vDash \tau$, there is a bijection $h : \mathbb{I} \mapsto \{1, \ldots, n\}$ and $\mathtt{Val} : \mathbb{X}^{\mathbb{I}} \mapsto \mathbb{N}$ such that $\rho \vDash_{h, \mathtt{Val}} \tau$. Define $\tau' = \langle \mathbb{I}', \preceq', \mathtt{src}', \mathtt{tgt}', E', \phi' \rangle$, to be the unique trace such that $\mathbb{I}' = \mathbb{I} - \{\mathtt{i}_1, \mathtt{i}_2\} \cup \{\mathtt{j}\}$, and $\mathtt{rank}_{\preceq'}(\mathtt{j}) = k$. Define $\mathtt{Val}' = (\mathtt{Val} \odot \{\mathtt{i}_1, \mathtt{i}_2\}) \cup \mathtt{Val}''$ where $\mathtt{Val}'' : \mathbb{X}^{\mathtt{j}} \mapsto \mathbb{N}$ is defined by $\mathtt{Val}''(x_{\mathtt{s}}^{\mathtt{j}}) = x_{\mathtt{s}}^{\mathtt{i}_1}$ and $\mathtt{Val}''(x_{\mathtt{t}}^{\mathtt{j}}) = x_{\mathtt{t}}^{\mathtt{i}_2}$ for all $x \in \mathbb{X}$. Define $h' : \mathbb{I}' \mapsto \{1, \ldots, n-1\}$ such that $h'(\mathbb{k}) = h(\mathbb{k})$ if $\mathbb{k} \prec \mathtt{i}_1$, $h'(\mathtt{j}) = k$, and $h'(\mathbb{k}) = h(\mathbb{k})$ if $\mathtt{i}_2 \prec \mathbb{k}$. It follows that $\rho' \vDash_{h', \mathtt{Val}'} \tau'$ and hence $\rho' \vDash \tau'$.

Lemma 11

Let $\tau_1 = \langle \mathbb{I}_1, \preceq_1, \mathtt{src}_1, \mathtt{tgt}_1, E_1, \phi_1 \rangle$, $\tau_2 = \langle \mathbb{I}_2, \preceq_2, \mathtt{src}_2, \mathtt{tgt}_2, E_2, \phi_2 \rangle$, $\tau_3 = \langle \mathbb{I}_3, \preceq_3, \mathtt{src}_3, \mathtt{tgt}_3, E_3, \phi_3 \rangle$, and $\tau_4 = \langle \mathbb{I}_4, \preceq_4, \mathtt{src}_4, \mathtt{tgt}_4, E_4, \phi_4 \rangle$, where $\phi_3 \sqsubseteq \phi_1$ and $\phi_4 \sqsubseteq \phi_2$. Since $\tau_3 \sqsubseteq \tau_1$ we know that $\tau_3 \sqsubseteq_{h_1} \tau_1$ for some $h_1 : \mathbb{I}_3 \mapsto \mathbb{I}_1$, and Since $\tau_4 \sqsubseteq \tau_4$ we know that $\tau_4 \sqsubseteq_{h_2} \tau_2$ for some $h_2 : \mathbb{I}_4 \mapsto \mathbb{I}_2$. Since $\tau_5 \in \tau_1 \otimes \tau_2$, τ_5 is of the form $\langle \mathbb{I}_5, \preceq_5, \mathtt{src}_5, \mathtt{tgt}_5, E_5, \phi_5 \rangle$, where $\mathbb{I}_5 = \mathbb{I}_1 \cup \mathbb{I}_2$, $\preceq_1 \subseteq \preceq_5$, $\preceq_2 \subseteq \preceq_5$, $\mathtt{src}_5 = \mathtt{src}_1 \cup \mathtt{src}_2$, $\mathtt{tgt}_5 = \mathtt{tgt}_1 \cup \mathtt{tgt}_2$, $E_5 = E_1 \cup E_2$, $\phi_5 = \phi_1 \wedge \phi_2$. Define $\tau_6 := \langle \mathbb{I}_6, \preceq_6, \mathtt{src}_6, \mathtt{tgt}_6, E_6, \phi_6 \rangle$, where $\mathbb{I}_6 = \mathbb{I}_3 \cup \mathbb{I}_4$, $\mathtt{src}_6 = \mathtt{src}_3 \cup \mathtt{src}_4$, $\mathtt{tgt}_6 = \mathtt{tgt}_3 \cup \mathtt{tgt}_4$, $E_6 = E_3 \cup E_4$, $\phi_6 = \phi_3 \wedge \phi_4$. Furthermore $\mathtt{i} \preceq_6 \mathtt{j}$ iff either $h_k(\mathtt{i}) \preceq_5 h_\ell(\mathtt{j})$, where $k = 1$ if $\mathtt{i} \in \mathbb{I}_3$ and $k = 2$ if $\mathtt{i} \in \mathbb{I}_4$, and where ℓ is defined in a similar manner. Define $h := h_1 \cup h_2$. It follows that $\tau_6 \in \tau_3 \otimes \tau_4$ and that $\tau_6 \sqsubseteq_h \tau_5$, i.e., $\tau_6 \sqsubseteq \tau_5$.

Lemma 12

Let $\tau_1 = \langle \mathbb{I}_1, \preceq_1, \mathtt{src}_1, \mathtt{tgt}_1, E_1, \phi_1 \rangle$, $\tau_2 = \langle \mathbb{I}_2, \preceq_2, \mathtt{src}_2, \mathtt{tgt}_2, E_2, \phi_2 \rangle$, and $\tau_3 = \langle \mathbb{I}_3, \preceq_3, \mathtt{src}_3, \mathtt{tgt}_3, E_3, \phi_3 \rangle$. Since $\tau_2 \sqsubseteq \tau_1$ we know that $\tau_2 \sqsubseteq_h \tau_1$ for some $h_1 : \mathbb{I}_2 \mapsto \mathbb{I}_1$. Since $\tau_3 \in \tau_1 \downarrow$, we know that τ_3 is of the form $\langle \mathbb{I}_3, \preceq_3, \mathtt{src}_3, \mathtt{tgt}_3, E_3, \phi_3 \rangle$, where $\mathbb{I}_3 = \mathbb{I}_1 - \{\mathtt{i}_1, \mathtt{i}_2\} \cup \{\mathtt{j}_1\}$ for some $\mathtt{i}_1, \mathtt{i}_2 \in \mathbb{I}_1$ and $\mathtt{j} \notin \mathbb{I}_1$. Define $\tau_4 := \langle \mathbb{I}_4, \preceq_4, \mathtt{src}_4, \mathtt{tgt}_4, E_4, \phi_4 \rangle$ to be the unique trace (up to renaming) such that $\tau_2 \in \tau_4 \downarrow$ and $\mathbb{I}_4 = \mathbb{I}_2 - \left\{h_1^{-1}(\mathtt{i}_1), h_1^{-1}(\mathtt{i}_2)\right\} \cup \{\mathtt{j}_2\}$ for some $\mathtt{j}_2 \notin \mathbb{I}_3$. Define $h_2 : \mathbb{I}_4 \mapsto \mathbb{I}_3$ such that $h_2(\mathtt{i}) = h_1(\mathtt{i})$ if $\mathtt{i} \in \mathbb{I}_2$ and $h_2(\mathtt{j}_2) = \mathtt{j}_1$. It follows that $\tau_4 \sqsubseteq_{h_2} \tau_3$ and hence $\tau_4 \sqsubseteq \tau_3$.

ENFORCING UNDER-APPROXIMATION ON THE AUTOMATON

Ordered stacks: Given a DMPDA $\mathcal{A} = \langle Q, \Delta \rangle$ and an ordering $\prec$ on the stacks Σ, we construct a new one $\mathcal{A}' = \langle Q', \Delta' \rangle$ where $Q' = Q \times 2^{\Sigma}$. The state remembers, in addition to the state of $\mathcal{A}$, the subset of stacks that are currently empty. The transitions Δ' lifts Δ to smoothly extend to Q' while maintaining the intended semantics. In particular a stack $\sigma \in \Sigma$ may be popped only if all the stacks σ' with $\sigma' \prec \sigma$ are empty. When pushing a symbl to a stack which is currently empty, the symbol that is pushed is tagged with flag. Further that stack is removed from the list of stacks that are currently empty. Eventually, when a stack is popped, if the symbol that is popped turns out be to be flagged, then that stack is added to the list of currently empty stacks. $\mathcal{A}'$ exhibits all and only executions of $\mathcal{A}$ which follows the ordering policy on stacks with respect to $\prec$. Thus the reachability relations wrt. $\prec$-ordered runs will exactly be the reachability relations of $\mathcal{A}'$. From the reachability relations of $\mathcal{A}'$ computed by the algorithm, we extract the reachability relations of $\mathcal{A}$ by projecting to Q.

Bounded scope: Given a DMPDA $\mathcal{A} = \langle Q, \Delta \rangle$ and an integer k which is the bound on scope, we construct a new one $\mathcal{A}' = \langle Q', \Delta' \rangle$ such that it exhibits all and only the k-scope

bounded runs. The construction of $\mathcal{A}'$ is more involved (cf. [21], section 12.3.3). We keep $Q' = Q \times \Sigma \times \{1 \dots k\}^{\Sigma}$. The state remembers, in addition to the state of $\mathcal{A}$, the stack that is being operated on in the current context, and for each stack a counter that tracks the number of context switches taken place after the bottom-most element of the stack has been inserted. The bottom-most element of a stack is tagged with a special symbol $\#$. If a $\#$ element is popped from the stack, then the corresponding counter is reset to 0. Further, each time a context switch is effectuated, all the counters are incremented. In this process, if any counter needs an increment beyond k, this run is aborted since such a transition is not defined. This case would arise only if there is some element in a stack such that after it has been pushed, there has been more than k context-switches. Thus $\mathcal{A}'$ exhibits all and only executions of $\mathcal{A}$ which are k-scope bounded. Thus the reachability relations wrt. k-scope bounded runs will exactly be the reachability relations of $\mathcal{A}'$, projecting the states to Q. From the reachability relations computed by the algorithm, we obtain the required reachability relations by projecting the states to Q.