VASS reachability in three steps

S. Lasota

November 28, 2016

This note presents a digested version of the decidability proof by Kosaraju [1]. Very roughly, the overall idea is to provide a decidable condition Θ on a VASS such that Θ implies reachability and $\neg\Theta$ implies that the size of VASS can be reduced. With these two properties, the size of input can be incrementally reduced until the problem becomes trivial. We first formulate the condition Θ for plain VASSes, then incrementally generalize it to VASSes with unconstrained coordinates, and finally to generalized VASSes of [1].

1 The reachability problem

A *vector addition system with states* (VASS) consists of a finite set of control states Q and a finite set $E \subseteq Q \times \mathbb{Z}^d \times Q$ of arcs. The number $d \geqslant 1$ is the *dimension* of a VASS. A pseudo-configuration is a pair $(q, v) \in Q \times \mathbb{Z}^d$; it is a *configuration* if $v \in \mathbb{N}^d$. An arc $e(q, z, q')$ induces a step

$$(q, v) \overset{e}{\rightsquigarrow} (q', v + z)$$

between pseudo-configurations. We write $q, v \dashrightarrow q', v'$ if there is a sequence of steps from (q, v) to (q', v'); every such sequence we call *pseudo-run*. If all vectors appearing in a pseudo-run belong to $\mathbb{N}^d$ we call it *run*, and write $q, v \longrightarrow q', v'$; this implies in particular that (q, v) and (q', v') themselves are configurations. The aim of this note is to describe an algorithm for

VASS REACHABILITY PROBLEM:

Input: a VASS (d, Q, E) and two configurations $(q, v), (q', v')$.

Question: does $q, v \longrightarrow q', v'$ hold?

Sufficient condition

As a warm-up, we prove a sufficient condition for reachability. For a VASS and two configurations (q, v), (q', v'), define the following two conditions:

Θ_1: For every $m \geqslant 1$, $q, v \dashrightarrow q', v'$ by a pseudo-run that uses every arc at least m times.

Θ_2: There are vectors $\Delta, \Delta' \geqslant \vec{1}$ such that

$$q, v \longrightarrow q, v + \Delta$$
$$q', v' \longleftarrow q', v' + \Delta'$$

Proposition 1. $\Theta_1 \wedge \Theta_2$ *implies* $q, v \longrightarrow q', v'$.

Proof. We will use the following claim, to be proved later:

Claim 1. $q', \Delta \dashrightarrow q', \Delta'$.

Here is a shape of a required run from (q, v) to (q', v'):

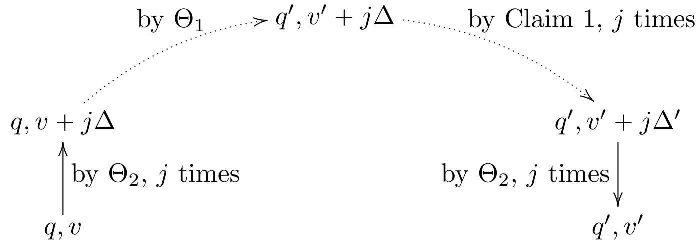

Observe that when j increases, the three intermediate points also increase on all coordinates. Therefore, for a sufficiently large j, the two pseudo-runs become runs. ☐

Proof of Claim 1

Consider the underlying graph of the VASS, whose vertices are control states and edges are arcs (note that there may be parallel edges). Every pseudo-run induces a path in the graph. For a pseudo-run from (p, w) to (p', w'), we shortly speak of a pseudo-run from p to p' when vectors w, w' are irrelevant. Let E denote the set of arcs. By the *folding* of a pseudorun π we mean the vector $\text{fold}(\pi) \in \mathbb{N}^E$ that says how many times every arc is used by π. The following lemma, roughly speaking, allows us to substract one pseudo-run from another:

Lemma 1. *Let τ, ρ be two pseudo-runs from p to p' such that*[1]

$$\textit{fold}(\tau) - \textit{fold}(\rho) \geqslant \vec{1}_E.$$

For every non-isolated control state p'' there is a pseudo-run σ from p'' to p'' with $\textit{fold}(\sigma) = \textit{fold}(\tau) - \textit{fold}(\rho)$.

By $\text{shift}(\pi) \in \mathbb{Z}^d$ we mean the effect of a pseudo-run π, namely the difference between its final vector and its initial one. Note that the shift of a pseudo-run is completely determined by its folding. To prove Claim 1, we need to show that there is a pseudo-run from q' to q' with shift $\Delta' - \Delta$.

Basing on condition Θ_1, we know that we can pick two pseudo-runs τ, ρ from (q, v) to (q', v') with arbitrarily large difference $\text{fold}(\tau) - \text{fold}(\rho)$. Fix a run π from (q, v) to $(q, v + \Delta)$, and a run π' from $(q', v' + \Delta')$ to (q', v'). Then fix two pseudoruns τ, ρ from (q, v) to (q', v') such that

$$\text{fold}(\tau) - \text{fold}(\rho) - \text{fold}(\pi) - \text{fold}(\pi') \geqslant \vec{1}.$$

Finally, apply Lemma 1 three times in a sequence, to deduce that there is a pseudo-run ν from q' to q' satisfying

$$\text{fold}(\nu) = \text{fold}(\tau) - \text{fold}(\rho) - \text{fold}(\pi) - \text{fold}(\pi').$$

Indeed, $\text{shift}(\nu) = \text{shift}(\tau) - \text{shift}(\rho) - \text{shift}(\pi) - \text{shift}(\pi') = \Delta' - \Delta$ as required.

[1] We write $\vec{m}_C$ for a constant vector in $\mathbb{Z}^C$ having m on all coordinates. We prefer to omit the subscript C and write simply $\vec{m}$ whenever this does not lead to confusion.

2 Partially unconstrained reachability problem

We now slightly generalize the reachability problem, and a sufficient condition. In the next section we will provide a yet further generalization that will be finally suitable for designing a decision procedure for reachability.

We will need a bit of concise notation. From now on we identify $\mathbb{Z}^d$ and $\mathbb{Z}^{\{1 \dots d\}}$; for instance, the set of configurations is $Q \times \mathbb{N}^{\{1 \dots d\}}$. For two disjoint subsets $C, B \in \{1 \dots d\}$ and two vectors $v \in \mathbb{Z}^C$ and $w \in \mathbb{Z}^B$, we write $v \oplus w$ for the unique vector in $\mathbb{Z}^{C \cup B}$ obtained by glueing together v and w. Formally:

$$(v \oplus w)(i) = \begin{cases} v(i) & \text{if } i \in C \\ w(i) & \text{if } i \in B. \end{cases}$$

From now on, by convention $\bar{C}$ will always denote the complement $\{1 \dots d\} - C$.

The generalization of the reachability problem amounts to considering only some subset $C \subseteq D$ of coordinates as *constrained*, while the remaining coordinates in $\bar{C}$ are considered *unconstrained*. The input and output configuration is specified only on constrained coordinates, and left unspecified on the remaining ones. Nevertheless, a run we ask for should remain nonnegative on all coordinates. Here is a precise formulation:

Partially unconstrained VASS reachability problem:

Input: a VASS (d, Q, E), two subsets $C, C' \subseteq \{1 \dots d\}$, $(q, v) \in Q \times \mathbb{N}^C$ and $(q', v') \in Q \times \mathbb{N}^{C'}$.

Question: does $q, v \oplus \bar{v} \longrightarrow q', v' \oplus \bar{v}'$ hold for some vectors $\bar{v} \in \mathbb{N}^{\bar{C}}$, $\bar{v}' \in \mathbb{N}^{\bar{C}'}$?

We do not assume $C = C'$. The setting of the previous section is the special case $C = C' = \{1 \dots d\}$.

Sufficient condition

Here is a generalization of Θ_1 and Θ_2 to the more general setting. We write $q, v \overset{C}{\dashrightarrow} q', v'$, for $C \subseteq \{1 \dots d\}$, to say that there is a pseudo-run from (q, v) to (q', v') whose all vectors are non-negative on coordinates from C. We write shortly $\mathbb{N}_{\geqslant m}$ for $\mathbb{N} - \{0 \dots m-1\}$.

Θ_1: For every $m \geqslant 1$, there are some vectors $\bar{v} \in (\mathbb{N}_{\geqslant m})^{\bar{C}}$, $\bar{v}' \in (\mathbb{N}_{\geqslant m})^{\bar{C}'}$ such that $q, v \oplus \bar{v} \dashrightarrow q', v' \oplus \bar{v}'$ by a pseudo-run that uses every arc at least m times.

Θ_2: There are vectors $\Delta \in (\mathbb{N}_{\geqslant 1})^C$, $\Delta' \in (\mathbb{N}_{\geqslant 1})^{C'}$, $\bar{\Delta} \in \mathbb{Z}^{\bar{C}}$ and $\bar{\Delta}' \in \mathbb{Z}^{\bar{C}'}$ such that

$$q, v \oplus \vec{0} \overset{C}{\dashrightarrow} q, (v + \Delta) \oplus \bar{\Delta}$$

$$q', v' \oplus \vec{0} \overset{C'}{\dashleftarrow} q', (v' + \Delta') \oplus \bar{\Delta}'$$

Proposition 2. $\Theta_1 \wedge \Theta_2$ *implies* $q, v \oplus \bar{v} \longrightarrow q', v' \oplus \bar{v}'$ *for some vectors* $\bar{v} \in \mathbb{N}^{\bar{C}}$, $\bar{v}' \in \mathbb{N}^{\bar{C}'}$.

Proof. The general idea of the proof is similar to the previous section, namely pumping up by a multiplicity of Δ (and de-pumping down by the same multiplicity of Δ') in order to make some pseudorun $q, v \oplus \bar{v} \dashrightarrow q', v' \oplus \bar{v}'$ a run. The new difficulty is that pumping involves $\Delta \oplus \bar{\Delta}$, with $\bar{\Delta}$ possibly negative on some coordinates (and likewise for de-pumping). This issue is solved by starting from $v \oplus \bar{v}$, for a sufficiently large $\bar{v} \geqslant \vec{m}$.

We will need a couple of facts. The first one easily follows from Θ_1:

Claim 2. *For some vectors* $\bar{v} \in \mathbb{N}^{\bar{C}}$, $\bar{\delta} \in (\mathbb{N}_{\geqslant 1})^{\bar{C}}$, $\bar{v}' \in \mathbb{N}^{\bar{C}'}$ *and* $\bar{\delta}' \in (\mathbb{N}_{\geqslant 1})^{\bar{C}'}$,

$$q, v \oplus \bar{v} \dashrightarrow q', v' \oplus \bar{v}'$$
$$q, v \oplus (\bar{v} + \bar{\delta}) \dashrightarrow q', v' \oplus (\bar{v}' + \bar{\delta}')$$

by two pseudo-runs π_0 *and* π_1 *with* $\mathit{fold}(\pi_1) - \mathit{fold}(\pi_0) \geqslant \vec{1}$.

Applying Lemma 1 to the two pseudo-runs of Claim 2 yields a pseudo-run τ from q' to q' satisfying $\mathrm{fold}(\tau) \geqslant \vec{1}$ and $\mathrm{shift}(\tau) = \vec{0} \oplus (\bar{\delta}' - \bar{\delta})$. Thus we can state:

Claim 3. $q', \vec{0} \oplus \bar{\delta} \dashrightarrow q', \vec{0} \oplus \bar{\delta}'$.

The next claim generalizes Claim 1 in the previous section; we use the fact that Θ_1 allows us to choose $\bar{\delta}$ and $\bar{\delta}'$ arbitrarily large.

Claim 4. *The vectors* $\bar{\delta}$, $\bar{\delta}'$ *can be chosen large enough so that additionally:*

(a) $\bar{\delta} + \bar{\Delta} \in (\mathbb{N}_{\geqslant 1})^{\bar{C}}$

(b) $\bar{\delta}' + \bar{\Delta}' \in (\mathbb{N}_{\geqslant 1})^{\bar{C}'}$

(c) $q', \Delta \oplus (\delta + \bar{\Delta}) \dashrightarrow q', \Delta' \oplus (\delta' + \bar{\Delta}')$

Fix two pseudo-runs π, π' as guaranteed by Θ_2. We have $\mathrm{shift}(\pi) = \Delta \oplus \bar{\Delta}$ and $\mathrm{shift}(\pi') = -\Delta' \oplus -\bar{\Delta}'$. By Θ_1 we can choose $\bar{\delta}$ and $\bar{\delta}'$ large enough so that the first two points hold, and moreover

$$\mathrm{fold}(\tau) - \mathrm{fold}(\pi) - \mathrm{fold}(\pi') \geqslant \vec{1}.$$

Then apply Lemma 1 to the pseudo-run τ repeated a times, two times in a sequence, to deduce that there is a pseudo-run ν from q' to q' satisfying

$$\mathrm{fold}(\nu) = \mathrm{fold}(\tau) - \mathrm{fold}(\pi) - \mathrm{fold}(\pi').$$

We check: $\mathrm{shift}(\nu) = \mathrm{shift}(\tau) - \mathrm{shift}(\pi) - \mathrm{shift}(\pi') = (\Delta' - \Delta) \oplus (\bar{\delta}' - \bar{\delta} + \bar{\Delta}' - \bar{\Delta})$, as required.

The last claim says that the pseudo-runs π and π', when lifted by sufficiently large multiplicities of $\vec{0} \oplus \bar{\delta}$ and $\vec{0} \oplus \bar{\delta}'$, respectively, become runs.

Claim 5. *For sufficiently large* $m \geqslant 0$ *it holds*

$$q, v \oplus m\bar{\delta} \longrightarrow q, (v + \Delta) \oplus (m\bar{\delta} + \bar{\Delta})$$
$$q', v' \oplus m\bar{\delta}' \longleftarrow q', (v' + \Delta') \oplus (m\bar{\delta}' + \bar{\Delta}')$$

We are now prepared to draw a shape of a required run (for readability, the primed items are depicted in blue):

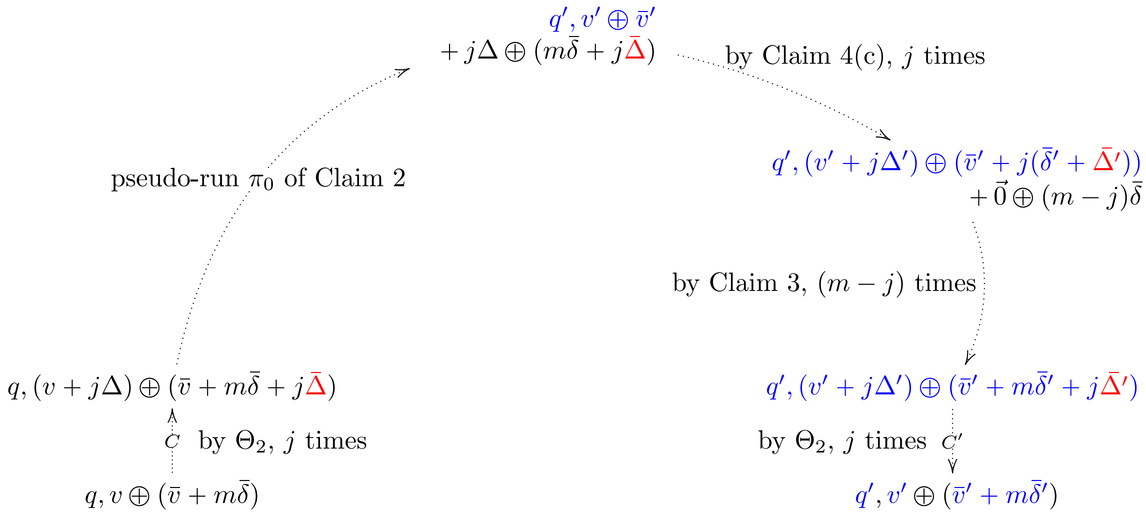

There are five consecutive pseudo-runs in the picture. When m increases, each of the six points increases on coordinates in $\bar{C}$, or on coordinates in $\bar{C}'$, and does not change on coordinates in C or C', respectively. Thus the first and the last pseudo-run become runs for a sufficiently large m, by Claim 5. Moreover, j and m increase simultaneously, each of the four intermediate points increases on all coordinates. As a conclusion, for sufficiently large j and m, all the five pseudo-runs become runs. □

Remark. For the purpose of the next section it is important to note that we have actually shown $q, v \oplus m\bar{\delta} \longrightarrow q', v' \oplus m\bar{\delta}'$ for every sufficiently large m.

3 Generalized reachability problem

In this section we do the last generalization in order to complete the decidability proof. By a *component* we mean a VASS (d, Q, E) together with the following data:

- initial and final state $q, q' \in Q$;
- subset of rigid coordinates $R \subseteq \{1 \dots d\}$; we assume that all arcs in E have 0 on all coordinates in R and hence, intuitively speaking, $d - |R|$ may be considered as the actual dimension of $\mathcal{V}$;
- rigid vector $r \in \mathbb{N}^R$;
- two partitions $\{1 \dots d\} - R = C \cup U = C' \cup U'$ of non-rigid coordinates into initial constrained coordinates C and initial unconstrained coordinates U, and into final constrained coordinates C' and final unconstrained coordinates U';
- initial and final vector $v \in \mathbb{N}^C$, $v' \in \mathbb{N}^{C'}$.

Note that component does not essentially differ from the input of the partially unconstrained reachability problem from the previous section. The *generalized VASS* (GVASS) $\mathcal{G}$ consists of $l \geqslant 1$ components

$$\mathcal{V}_i = (d, E_i, Q_i, q_i, q'_i, R_i, r_i, C_i, U_i, C'_i, U'_i, v_i, v'_i)$$

of the same dimension d, with pairwise disjoint state sets Q_i, and $l-1$ arcs of the form $e_i = (q'_i, z_i, q_{i+1})$, where $z_i \in \mathbb{Z}^d$, for $i \in \{1 \dots i-1\}$. We will be interested in pseudo-runs from q_1 to q'_l of the following form:

$$\begin{aligned} q_1, r_1 \oplus v_1 \oplus u_1 &\dashrightarrow q'_1, r_1 \oplus v'_1 \oplus u'_1 \overset{e_1}{\rightsquigarrow} \\ q_2, r_2 \oplus v_2 \oplus u_2 &\dashrightarrow q'_2, r_2 \oplus v'_2 \oplus u'_2 \overset{e_2}{\rightsquigarrow} \quad \dots \\ \dots \overset{e_{l-1}}{\rightsquigarrow} \quad q_l, r_l \oplus v_l \oplus u_l &\dashrightarrow q'_l, r_l \oplus v'_l \oplus u'_l \end{aligned} \tag{1}$$

for $u_1 \in \mathbb{N}^{U_1}, u'_1 \in \mathbb{N}^{U'_1}, \dots, u_l \in \mathbb{N}^{U_l}, u'_l \in \mathbb{N}^{U'_l}$. Each such pseudo-run π passes through every arc e_i exactly once, and thus splits into l pseudo-runs $\pi = \pi_1 e_1 \pi_2 e_2 \dots e_{l-1} \pi_l$, each π_i being a pseudo-run in $\mathcal{V}_i$. When each of π_i is a run, we say that $\mathcal{G}$ *admits reachability.*

> GENERALIZED VASS REACHABILITY PROBLEM:
>
> *Input:* a GVASS $\mathcal{G}$.
> *Question:* does $\mathcal{G}$ admit reachability?

The setting of the previous section is the special case of one component without rigid coordinates: $l = 1$, $R_1 = \varnothing$.

Sufficient condition

The condition Θ_2 below is essentially the conjunction of conditions Θ_2 of the previous section for each of the VASSes $\mathcal{V}_i$ separately; the only difference is taking rigid coordinates into account. On the other hand, the condition Θ_1 below speaks jointly about all the VASSes $\mathcal{V}_i$.

- Θ_1: For every $m \geqslant 1$, there is a pseudo-run from q_1 to q'_l of the form (1) that visits every arc in every E_i at least m times, for some $u_1 \in (\mathbb{N}_{\geqslant m})^{U_1}, u'_1 \in (\mathbb{N}_{\geqslant m})^{U'_1}, \dots, u_l \in (\mathbb{N}_{\geqslant m})^{U_l}, u'_l \in (\mathbb{N}_{\geqslant m})^{U'_l}$.
- Θ_2: For every $i \in \{1 \dots l\}$ there are vectors $\Delta \in (\mathbb{N}_{\geqslant 1})^{C_i}$, $\Delta' \in (\mathbb{N}_{\geqslant 1})^{C'_i}$, $\bar{\Delta} \in \mathbb{Z}^{U_i}$ and $\bar{\Delta}' \in \mathbb{Z}^{U'_i}$ such that

$$q_i, r_i \oplus v_i \oplus \vec{0} \overset{C_i}{\dashrightarrow} q_i, r_i \oplus (v_i + \Delta) \oplus \bar{\Delta} \tag{2}$$

$$q'_i, r_i \oplus v'_i \oplus \vec{0} \overset{C'_i}{\dashleftarrow} q'_i, r_i \oplus (v'_i + \Delta') \oplus \bar{\Delta}' \tag{3}$$

Observe that Θ_1 implies $C'_i = C_{i+1}$ for $i \in \{1 \dots l-1\}$. The sufficient condition for reachability is proved similarly as in the previous section:

Proposition 3. *If $\mathcal{G}$ satisfies $\Theta_1 \wedge \Theta_2$ then $\mathcal{G}$ admits reachability.*

In Claim 2 one should consider all components simultaneously; for the other claims and the run one can consider the components separately. Furthermore, the sufficient condition can be effectively tested.

Proposition 4. *Both* Θ_1 *and* Θ_2 *are decidable.*

Proof. Pseudo-runs (1) can be encoded as the set of nonegative solutions of a system of linear diophantine equations. Then condition Θ_1 can be decided by inspecting the set of solutions (cf. Claim 7 below). (Alternatively, Presburger arithmetic can be used.) Condition Θ_2 can be decided using coverability trees. □

Refinement

Let $\text{size}(\mathcal{V}_i) = (d - |R_i|, |E_i|, |U_i| + |U'_i|) \in \mathbb{N}^3$. Thus the size of $\mathcal{V}_i$ is a triple consisting of: the number of non-rigid coordinates, the number of arcs, the number of unconstrained coordinates. For a GVASS $\mathcal{G}$, we define $\text{size}(\mathcal{G})$ as the multiset of sizes of all components $\mathcal{V}_i$.

Order triples in $\mathbb{N}^3$ lexicographically. For two finite multisets of triples m and m', we say that m' *refines* m if m' is obtained by removing one triple from m, and replacing it by a finite number of lexicographically strictly smaller triples.

Claim 6. *The refinement relation is well-founded.*

We shortly say that $\mathcal{G}'$ refines $\mathcal{G}$ when $\text{size}(\mathcal{G}')$ refines $\text{size}(\mathcal{G})$. For *trivial* $\mathcal{G}$, whose size contains only zero triples $(0, 0, 0)$, the reachability problems trivializes. Otherwise, either $\mathcal{G}$ satisfies $\Theta_1 \wedge \Theta_2$ and thus admits reachability, or $\mathcal{G}$ can be refined:

Proposition 5. *If a non-trivial* $\mathcal{G}$ *violates* Θ_1 *then one can compute* $\mathcal{G}_1 \dots \mathcal{G}_n$ *refining* $\mathcal{G}$ *such that* $\mathcal{G}$ *admits reachability if, and only if some of* $\mathcal{G}_1 \dots \mathcal{G}_n$ *does.*

Proof. Wlog. assume that the underlying graphs of all VASSes $\mathcal{V}_i$ are strongly connected (otherwise $\mathcal{G}$ can be easily refined along the lines described below). Let $k = \sum_{i=1}^{l} |U_i| + E_i + |U'_i|$. Consider the set $L \subseteq \mathbb{N}^k$ of all vectors

$$(u_1, f_1, u'_1, \dots, u_l, f_l, u'_l) \in \mathbb{N}^k$$

such that there is a pseudo-run $\pi = \pi_1 e_1 \pi_2 e_2 \dots e_{l-1} \pi_l$ of the form (1) with $\text{fold}(\pi_1) = f_1 \geqslant \vec{1}, \dots, \text{fold}(\pi_l) = f_l \geqslant \vec{1}$. The set L is the set of nonnegative solutions of a system of linear diophantine equations, and thus we have:

Claim 7. *One can compute finite sets* $B, P \subseteq \mathbb{N}^k$ *such that* $L = B + P^*$.

Suppose $\mathcal{G}$ does not satisfy Θ_1. Hence all vectors in P have zero on some coordinate in $\{1 \dots k\}$. This *zero coordinate* corresponds either to some arc, or to some unconstrained (input or output) coordinate.

Suppose the first case holds, and let $e \in E_i$ be the arc corresponding to the zero coordinate. By Claim 7 one can compute a number c such that every pseudo-run (1) passes through e at most c times. We refine $\mathcal{G}$ by $c+1$ GVASSes $\mathcal{G}_0 \dots \mathcal{G}_c$, each $\mathcal{G}_m$ obtained by replacing $\mathcal{V}_i$ by a sequence of $m+1$ copies of $\mathcal{V}_i - \{e\}$, i.e. of $\mathcal{V}_i$ without the arc e. The rigid coordinates and rigid vector of all copies are as in $\mathcal{V}_i$. The initial constrained coordinates of the first copy are

C_i, the final constrained coordinates of the last copy are C'_i, and the remaining initial or final constrained coordinates of all copies are empty sets. The initial vector of the first copy is v_1 and the final vector of the last copy is v'_l; all other initial and final vectors are empty ones.

Now suppose the second case holds, i.e. the zero coordinate corresponds to some, say, initial unconstrained coordinate $j \in U_i$ (final unconstrained coordinate is treated symmetrically). By Claim 7 one can compute a number c such that the value on coordinate j in $\mathcal{V}_i$ is bounded by c in every pseudo-run (1). We refine $\mathcal{G}$ by constraining the coordinate j to some value in $\{0 \ldots c\}$. We define $c + 1$ refining GVASSes $\mathcal{G}_0 \ldots \mathcal{G}_c$, where $\mathcal{G}_m$ differs from $\mathcal{G}$ only by making the coordinate j in $\mathcal{V}_i$ an initial constrained coordinate, with value m. □

Proposition 6. *If a non-trivial $\mathcal{G}$ satisfies Θ_1 but violates Θ_2 then one can compute $\mathcal{G}_1 \ldots \mathcal{G}_n$ refining $\mathcal{G}$ such that $\mathcal{G}$ admits reachability if, and only if some of $\mathcal{G}_1 \ldots \mathcal{G}_n$ does.*

Proof. Suppose that $\mathcal{G}$ does satisfy Θ_1 (in particular, its underlying graph is strongly connected) but not satisfy Θ_2, i.e. condition (2) fails for some i (condition (3) is treated symmetrically). From the coverability tree for $\mathcal{V}_i$ one can extract a number c such that in every pseudo-configuration reachable in $\mathcal{V}_i$ from $v_i \oplus \vec{0}$ via the relation $\overset{C_1}{\dashrightarrow}$, some of initial constrained coordinates $j \in C_i$ is bounded by c. (Otherwise, all input constrained coordinates could be simultaneously increased arbitrarily; this, together with strong connectedness of the underlying graph, would imply (2).) We refine $\mathcal{G}$ by a finite family of GVASSes. For every $j \in C_i \cap C'_i$ the family contains one GVASS $\mathcal{G}_j$, and for every $j \in C_i \cap U'_i$ the family contains $c + 1$ GVASSes $\mathcal{G}_{j,0} \ldots \mathcal{G}_{j,c}$, as outlined below:

$j \in C_i \cap U'_i$: Thus j is a final unconstrained coordinate. We define GVASSes $\mathcal{G}_{j,0} \ldots \mathcal{G}_{j,c}$, where $\mathcal{G}_{j,m}$ differes from $\mathcal{G}$ only by making the coordinate j a final constrained coordinate in $\mathcal{V}_i$, and fixing its value to m.

$j \in C_i \cap C'_i$: Thus j is a final constrained coordinate. Let a and a' be the values of initial and final vectors v_i, v'_i on coordinate j. We define $\mathcal{G}_j$ by replacing $\mathcal{V}_i$ with two VASSes $\mathcal{V}'$ and $\mathcal{V}''$. $\mathcal{V}'$ behaves exactly as $\mathcal{V}_i$ with the only exception that the value of the jth coordinate is kept between 0 and c. This can be achieved using a cross-product of $\mathcal{V}_i$ with a finite state automaton, with states $\{0, \ldots, c\}$, the initial state a, the final state a', and transitions induced by the jth coordinate of arcs in E_i. This allows to set the jth coordinate of all arcs in $\mathcal{V}'$ to 0; in consequence, the coordinate j can be moved to rigid coordinates of $\mathcal{V}'$. Thus $\mathcal{V}'$ has $(c + 1)$ times more states and arcs than $\mathcal{V}_i$ but one less non-rigid coordinate. The rigid vector of $\mathcal{V}'$ is set to a on coordinate j. The difference $a' - a$ is easily compensated by adding one arc-less VASS $\mathcal{V}''$ to $\mathcal{G}_j$, connected to $\mathcal{V}'$ by an arc that adds $a' - a$ on coordinate j and preserves all other coordinates. □

References

[1] S. Rao Kosaraju. Decidability of reachability in vector addition systems (preliminary version). In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing*, pages 267–281, 1982.