The Invariance Problem for Matrix Semigroups

Klaus Dräger(✉)

EECS, Queen Mary University of London, London, UK
klaus.draeger@gmail.com

Abstract. The question of whether a given subspace of $\mathbb{Q}^d$ can be reached from a starting vector using linear transformations from a given finite set is well known to be undecidable in dimension 3 and above. We show that, in contrast, the *invariance* problem, i.e. the question of whether it is possible to remain inside a given subspace indefinitely using linear transformations from a given finite set, is decidable.

1 Introduction

The classic subspace reachability problem for matrix semigroups is the following.

- Given: A finite set $A \subseteq M_d(\mathbb{Q})$ of matrices, a linear subspace $U \subseteq \mathbb{Q}^d$, and a vector $x_0 \in \mathbb{Q}^d$.
- Question: Starting from x_0, can we reach U by applying a finite sequence of matrices from A? That is, do there exist $M_1, \ldots, M_k \in A$ such that $M_k \cdots M_1 x_0 \in U$?

This problem is known to be undecidable for $d \geq 3$ [11]. In this paper, we consider a temporal dual, the *invariance* problem:

- Given: A finite set $A \subseteq M_d(\mathbb{Q})$ of matrices, a linear subspace $U \subseteq \mathbb{Q}^d$, and a vector $x_0 \in \mathbb{Q}^d$.
- Question: Starting from x_0, can we remain in U indefinitely, using matrices from A? That is, does there exist a sequence $M_1, M_2, \ldots$ such that, for all $k \in \mathbb{N}$, $M_k \in A$ and $M_k \cdots M_1 x_0 \in U$?

The main result of this paper is that, unlike the reachability problem, the invariance problem is decidable.

The temporal duality becomes clearer when thinking about the negation of the properties: Consider the infinite-state transition system $S = (\mathbb{Q}^d, x_0, A)$ with transitions $\mathbb{Q}^d \to \mathbb{Q}^d$ defined by the given set A of matrices, and the formula $\varphi_U \equiv r_1 \cdot x = \ldots = r_k \cdot x = 0$ defining the subspace U in terms of a basis $(r_1, \ldots, r_k)$ of its orthogonal complement, then a solution to the reachability problem is a counterexample to the LTL assertion $S \vDash \Box\neg\varphi_U$ (stating that S will *always* be outside U), while a solution to the invariance problem is a counterexample to the LTL assertion $S \vDash \Diamond\neg\varphi_U$ (stating that S will *eventually* be outside U). Note the existential character of the problem: we are asking about *satisfiability* of the invariant. The universal version (the question of *validity*) can easily be shown to be also decidable; see Sect. 4.3.

B. Jacobs and C. Löding (Eds.): FOSSACS 2016, LNCS 9634, pp. 479–492, 2016.
DOI: 10.1007/978-3-662-49630-5_28

As a technical aside, when we talk about a matrix semigroup in this paper, we mean a sub-semigroup G of $M_n(\mathbb{Q})$ *equipped with a particular generating set* A. For the reachability problem, this detail is irrelevant (all that is required is the existence of some $M \in G$ with $Mx_0 \in U$), but it matters for the invariance problem. Consider the case $x_0 = (1,1,0)^T$, $U = \{(x,y,0)^T \mid x,y \in \mathbb{Q}\}$, and $A_1 = \{M_1, M_3\}, A_2 = \{M_2, M_3\}$, where

$$M_1 = \begin{pmatrix} 0&1&0 \\ 1&0&0 \\ 0&0&1 \end{pmatrix}, M_2 = \begin{pmatrix} 0&0&1 \\ 0&1&0 \\ 1&0&0 \end{pmatrix}, M_3 = \begin{pmatrix} 1&0&0 \\ 0&0&1 \\ 0&1&0 \end{pmatrix}.$$

Both A_1 and A_2 generate the natural permutation representation of the symmetric group S_3, but the invariance problem for x_0, U, A_1 has a solution $M_1^\omega \in A_1^\omega$, while there is no solution for x_0, U, A_2.

1.1 Attacking the Invariance Problem

There is an intuitively obvious approach to the invariance problem, which proceeds as follows:

- Expand the tree of words $w \in A^*$.
- Abort any branch $M_1 \dots M_k$ with $M_k \cdots M_1 x_0 \notin U$; if there are no more branches left, then there is no solution.
- If for some branch $M_1 \dots M_k$ we have $M_k \cdots M_1 x_0 = 0$, we have a solution, since $0 \in U$ is fixed by any $M \in A$ (i.e. $M_1 \dots M_k w$ will do for any $w \in A^\omega$).

However, this method is not complete. The result which allows us to fix it is a pumping lemma for solution prefixes: there is a bound N depending only on $|A|$ and $\dim U$ such that any word $w \in A^\omega$ has a prefix $uv \trianglelefteq w$ with

- $|uv| \leq N$, and
- (u,v) is *dominating* in a sense defined below, which in particular implies that v is nonempty and if w is a solution to the invariance problem, then so is uv^ω.

Once we have this, the decision problem reduces to checking finitely many pairs (u,v), as in the algorithm in Fig. 1.

Example 1. Let $x_0 = (1,0,0,0)^T \in \mathbb{Q}^4$, $U = \{(x,y,z,0)^T \mid x,y,z \in \mathbb{Q}\}$, and $A = \{P,Q,R\}$ with

$$P = \begin{pmatrix} 1&2&0&0 \\ 0&0&1&2 \\ 2&1&0&0 \\ 0&0&2&1 \end{pmatrix}, Q = \begin{pmatrix} 0&-1&0&2 \\ -1&0&2&0 \\ 2&0&-1&0 \\ 0&2&0&-1 \end{pmatrix}, R = \begin{pmatrix} 0&0&-1&2 \\ -1&0&2&0 \\ 0&1&0&0 \\ 2&0&0&-1 \end{pmatrix}.$$

Figure 2 shows the beginning of a breadth-first exploration for this example; one solution to the invariance problem is the periodic sequence $PQR(RPPQPQ)^\omega$, corresponding to the dominating prefix $(PQR, RPPQPQ)$.

Data: finite set $A \subseteq M_d(\mathbb{Q})$ of matrices,
initial vector $x_0 \in \mathbb{Q}^d$,
invariant subspace $U \subseteq \mathbb{Q}^d$
Result: A pair $(u, v) \in A^* \times A^*$ such that uv^ω solves the invariance problem for A, x_0, U, or FAIL if no such pair exists

```
if A = ∅ or x_0 ∉ U then return FAIL;
branches := {ϵ};
while branches not empty do
    pop first w from branches;
    if w(x_0) = 0 then
        pick M ∈ A;
        return (w, M) ;                    /* wM^ω is solution for any M */
    if w(x_0) ∈ U then
        foreach split w = uv with v ≠ ϵ do
            if v is pumpable then
                return (u, v) ;  /* w is dominating prefix for all wM */
        foreach M ∈ A do
            append wM to branches;
return FAIL;
```

Fig. 1. Finding a trace remaining inside U, starting from x_0. Exploration follows a breadth-first strategy in order to find a minimal-length solution. The details of the notions of *pumpability* and *dominating prefix*, and proof that such a prefix exists for any sufficiently long branch, are given in Sects. 2 and 3.

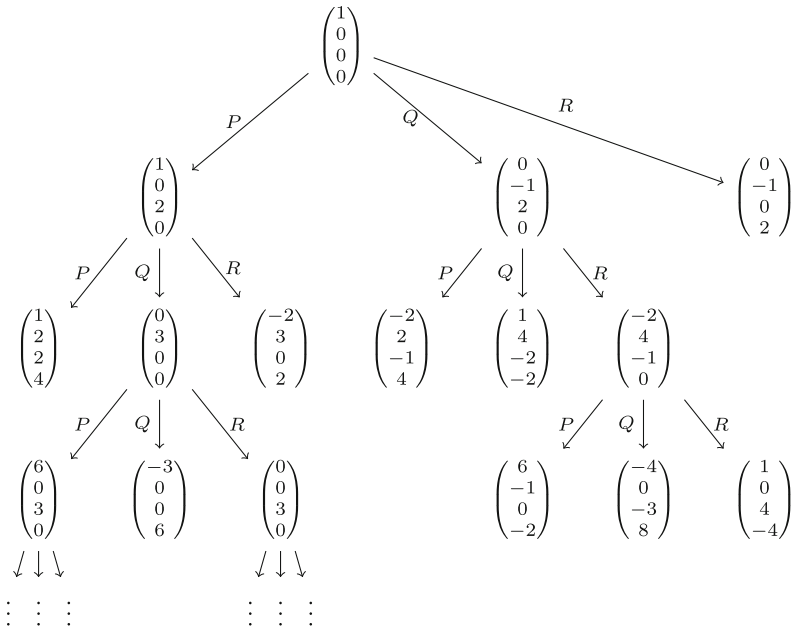

Fig. 2. The first few layers of the running example

2 Preliminaries

2.1 Setting

Throughout this paper, we work within $\mathbb{Q}^d$ for some $d \in \mathbb{N}$. We assume the linear subspace $U \subseteq \mathbb{Q}^d$, starting vector $x_0 \in U$, and finite non-empty set of matrices $A \subseteq M_d(\mathbb{Q})$ to be arbitrary but fixed, subject to the non-triviality condition

$$MU \nsubseteq U \text{ for all } M \in A.$$

If there is some M violating this condition, then M^ω is a trivial solution to the invariance problem.

We use A^* and A^ω for the sets of finite and infinite words over A, respectively, and ϵ for the empty word. For $x \in \mathbb{Q}^n, S \subseteq \mathbb{Q}^n$, and $w = M_1 \dots M_k \in A^*$, $w(x) := M_k \cdots M_1 x$ and $w(S) := \{w(x) \mid x \in S\}$. A word $p \in A^*$ is a prefix of $w = M_1 \dots M_k \in A^*$ (resp. $w = M_1 M_2 \dots \in A^\omega$), denoted by $p \trianglelefteq w$, iff $p = M_1 \dots M_j$ for some $j \leq k$ (resp. some $j \in \mathbb{N}$).

Definition 1. *Let $w = M_1 \dots M_k \in A^*$. Associated with w are the following subspaces of $\mathbb{Q}^n$:*

- *the* source space $S(w) = \{x \in \mathbb{Q}^n \mid p(x) \in U \text{ for each prefix } p \trianglelefteq w\}$, *and*
- *the* target space $T(w) = w(S(w))$.

The following elementary properties of $S(\_), T(\_)$ are used throughout the proof of the main result:

Lemma 1. *For all $u, v, w \in A^*$,*

(i) $S(uv) = \{x \in S(u) \mid u(x) \in S(v)\} \subseteq S(u)$, *with equality iff* $T(u) \subseteq S(v)$,
(ii) $T(uv) = v(T(u) \cap S(v)) \subseteq T(v)$, *with equality if* $T(u) \supseteq S(v)$,
(iii) $\dim T(uvw) \leq \dim T(v)$.

Proof

(i) Since any prefix of uv is either a prefix of u or a word of the form up, where p is a prefix of v, we have

$$\begin{aligned} S(uv) &= \{x \in \mathbb{Q}^n \mid p(x) \in U \text{ for all } p \trianglelefteq uv\} \\ &= \{x \in \mathbb{Q}^n \mid p(x) \in U \text{ for all } p \trianglelefteq u \text{ and } p(u(x)) \in U \text{ for all } p \trianglelefteq v\} \\ &= \{x \in \mathbb{Q}^n \mid x \in S(u) \text{ and } u(x) \in S(v)\}, \end{aligned}$$

which is obviously a subspace of $S(u)$. Equality holds iff for all $x \in S(u)$, $u(x)$ is in $S(v)$, iff $T(u) \subseteq S(v)$.

(ii) Using $uv(x) = v(u(x))$, we get

$$\begin{aligned} T(uv) &= v(u(S(uv))) \\ &= v(u(\{x \in S(u) \mid u(x) \in S(v)\})) \\ &= v(\{u(x) \mid x \in S(u) \text{ and } u(x) \in S(v)\}) \\ &= v(T(u) \cap S(v)) \\ &\subseteq v(S(v)) = T(v). \end{aligned}$$

If $T(u) \supseteq S(v)$, the inclusion in the last line holds with equality.

(iii) By (ii), $T(uvw) \subseteq T(vw)$, and we have

$$\begin{aligned}\dim T(uvw) &\leq \dim T(vw)\\ &= \dim w(T(v) \cap S(w))\\ &\leq \dim w(T(v))\\ &\leq \dim T(v),\end{aligned}$$

where in the last step we use that $\dim f(V) \leq \dim V$ for any linear map f and vector space V. □

If $w = M_1 M_2 \ldots \in A^\omega$ is an infinite word, then due to Lemma 1(i), the source spaces of its prefixes form a descending chain $S(\epsilon) \supseteq S(M_1) \supseteq \ldots$; since the lattice of subspaces satisfies the descending chain condition, this chain has a limit, $S(w)$, the space of all vectors whose w-orbit remains in U. The invariance problem obviously amounts to checking the existence of an infinite word w with $x_0 \in S(w)$.

Definition 2. *A word $q \in A^*$ is* pumpable *if $q \neq \epsilon$ and $T(q) \subseteq S(q)$.*

A dominating prefix *of a (finite or infinite) word w is a pair (p, q) such that pq is a proper prefix of w and q is pumpable.*

Any word which has a dominating prefix can then be disregarded due to the following lemma.

Lemma 2. *Let $p, q \in A^*$. If q is pumpable, then*

(i) $S(q^k) = S(q)$ for all $k \geq 1$,
(ii) $S(pq^k) = S(pq)$ for all $k \geq 1$, and
(iii) if pq is a prefix of a word $w \in A^\omega$, and w is a solution to the invariance problem for x_0, U, A, then so is pq^ω.

Proof

(i) We use induction on k. The assertions holds trivially for $k = 1$; for the step $k \to k+1$, from $T(q) \subseteq S(q) = S(q^k)$ we get, using Lemma 1(i), that $S(q) = S(qq^k) = S(q^{k+1})$.
(ii) From (i) and Lemma 1(i) we have
$S(pq) = \{x \in S(p) \mid p(x) \in S(q)\} = \{x \in S(p) \mid p(x) \in S(q^k)\} = S(pq^k)$.
(iii) $S(pq^\omega)$ is the limit of $(S(pq^k))_{k\geq 1}$, which by (ii) equals the constant sequence $(S(pq))_{k\geq 1}$, and therefore $S(pq^\omega) = S(pq)$. Since pq is a prefix of w, and $x_0 \in S(w)$, we get
$x_0 \in S(w) \subseteq S(pq) = S(pq^\omega)$. □

Our goal now is to show that there is a bound N such that any word w with $|w| > N$ has a dominating prefix (p, q) with $|pq| \leq N$.

Definition 3. *A word $w \in A^*$ is*

– essential *if it has no dominating prefix,*
– *S-minimal of dimension k if $S(w) \subsetneq S(p)$ for each proper prefix $p \triangleleft w$ and* $\dim S(w) = k$,

- *T-minimal of dimension k if $k = \dim T(w) < \dim T(p)$ for each proper prefix $p \lhd w$,*
- *S-essential (resp. T-essential) of dimension k if it is both essential and S-minimal (resp. T-minimal) of dimension k.*

We will omit the dimension for the last three properties when it is irrelevant.

Example. Continuing from the partial exploration shown in Fig. 2, Table 1 contains the source and target spaces for some selected words occurring in the tree, and whether or not they are S-minimal. All the words in the figure are essential.

Table 1. Some source and target spaces occurring in the running example

w	$S(w)$	$T(w)$	S-minimal
ϵ	$\{(x,y,z,0)^T\}$	$\{(x,y,z,0)^T\}$	yes, dimension 3
P	$\{(x,y,0,0)^T\}$	$\{(x+2y,0,2x+y,0)^T\}$	yes, dimension 2
Q	$\{(x,0,z,0)^T\}$	$\{(0,2z-x,2x-z,0)^T\}$	yes, dimension 2
R	$\{(0,y,z,0)^T\}$	$\{(-z,2z,y,0)^T\}$	yes, dimension 2
PQ	$\{(x,y,0,0)^T\}$	$\{(0,3x,3y,0)^T\}$	no
QR	$\{(x,0,z,0)^T\}$	$\{(z-2x,4x-2z,2z-x,0)^T\}$	no
PQP	$\{(x,0,0,0)^T\}$	$\{(6x,0,3x,0)^T\}$	yes, dimension 1
PQR	$\{(x,y,0,0)^T\}$	$\{(-3y,6y,3x,0)^T\}$	no
$PQRQ$	$\{(x,0,0,0)^T\}$	$\{(0,6y,-3x,0)^T\}$	yes, dimension 1

The plan now is to prove that there are only finitely many essential words, using the following *factorization*.

Definition 4. *The* S-factorization $F(w)$ *of a word* $w \in A^*$ *is a finite sequence of nonempty finite words defined as follows. If* $w = \epsilon$, *then* $F(w) = ()$. *For* $w \neq \epsilon$, *let* $p \trianglelefteq w$ *be the shortest prefix of* w *for which* $S(p) = S(w)$, *and* $(q_1, \ldots, q_k) = F(q)$, *where* q *is the suffix with* $w = pq$; *then* $F(w) = (p, q_1, \ldots, q_k)$. *Note that since we require* $MU \not\subseteq U$ *for all* $M \in A$, *we have* $U = S(\epsilon) \supsetneq S(w)$ *for all* $w \neq \epsilon$, *which ensures that* $p \neq \epsilon$.

For infinite $w \in A^\omega$, *we get an infinite version of this factorization by corecursively defining* $F(w) = (p_1, p_2, \ldots)$, *where* p_1 *is the shortest prefix of* w *with* $S(p_1) = S(w)$, *and* $(p_2, \ldots) = F(q)$ *for the infinite suffix* q *with* $w = pq$.

Note that actually computing the infinite factorization (even incrementally) would generally not be feasible (since a decrease in dimension $S(uM) \subsetneq S(u)$ could occur after an arbitrarily long prefix u, we could not determine p_1 based on any finite prefix u, unless $S(u)$ happens to be $\{0\}$), but we will just need its existence and some of its properties. However, in case w is periodic, we can compute $F(w)$:

Example 2. Consider the word $w = uv^\omega$ with $u = PQR, v = RPPQPQ$ from Example 1. A quick calculation gives that $S(v) = T(v) = S(v^\omega) = \{(0, 0, z, 0)^T \mid z \in \mathbb{Q}\}$. Since $u(x_0) \in S(v)$, we get that $x_0 \in S(p)$, and therefore $\dim S(p) \geq 1$, for any prefix $p \trianglelefteq w$.

On the other hand, since P, Q, R are invertible, we have that $\dim S(pv) = \dim T(pv) \leq \dim T(v) = 1$ for any word p, so that the dimension of each factor in $F(w)$ must be 1. Since w is periodic, the same will be true for the factorization. In fact, we get

$$F(PQR(RPPQPQ)^\omega) = (PQR, RP, PQ, PQR, PPQ, PQR, \ldots)$$

3 Deciding Satisfiability

We now show that there are only finitely many essential words, and derive upper bounds on their number and length, from which the main result follows.

3.1 Finiteness of the Set of Essential Words

Lemma 3. *If $w \in A^*$ is S-minimal of dimension k, then it is also T-minimal of some dimension $j \leq k$. As a direct corollary, if w is S-essential of dimension k, then it is also T-essential of some dimension $j \leq k$.*

Proof. That $\dim T(w) \leq \dim S(w)$ follows directly from $T(w) = w(S(w))$. If $w = \epsilon$, then it is trivially T-minimal.

Otherwise, $w = vM$ for some $v \in A^*$ and $M \in A$. By Lemma 1(iii), it suffices to show $\dim T(w) < \dim T(v)$. Since w is S-minimal, $S(w) = S(vM) \subsetneq S(v)$, and thus $T(v) \nsubseteq S(M)$ by Lemma 1(i). Therefore

$$\begin{aligned} \dim T(vM) &= \dim M(T(v) \cap S(M)) \qquad \text{by Lemma 1(ii)} \\ &\leq \dim\, (T(v) \cap S(M)) \\ &< \dim T(v), \end{aligned}$$

i.e. w is T-minimal.

The corollary follows by just adding non-existence of dominating prefixes of w to both the assumption and conclusion. □

Theorem 1. *Let $w = vM$ be a non-empty essential word, where $v \in A^*$ and $M \in A$. Let the S-factorization of v be $F(v) = (p_1, \ldots, p_m)$. Then*

(i) each p_i is T-essential,
(ii) if w is T-minimal of dimension k, then each p_i is T-minimal of some dimension $d_i > k$, and
(iii) $p_i \neq p_j$ for $i \neq j$.

Proof

(i) By construction of the S-factorization, each p_i is S-minimal. By Lemma 3, it is also T-minimal.
Suppose p_i is not essential, then it has a dominating prefix (r, s). This implies that $(p_1 \dots p_{i-1} r, s)$ is a dominating prefix for the essential word w, contradiction.

(ii) We already have that each p_i is S-minimal and thus T-minimal by Lemma 3; it remains to show $\dim T(p_i) > \dim T(w)$. This follows because p_i is a factor of v, so $\dim T(p_i) \geq \dim T(v)$ by Lemma 1(iii), and $\dim T(v) > \dim T(w)$ due to T-minimality of w.

(iii) Suppose $p_i = p_j$ for some $i < j$. By the definition of $F(v)$ we then have $S(p_i \dots p_m) = S(p_i) = S(p_j) = S(p_j \dots p_m)$; furthermore, since the spaces $(S(p_i \dots p_k))$ form a descending chain, we have in fact $S(p_i) = S(p_i \dots p_k)$ for all $i \leq k \leq m$, and in particular $S(p_i) = S(p_i \dots p_{j-1}) = S(p_i \dots p_m)$. This implies

$$\begin{aligned} T(p_i \dots p_{j-1}) &\subseteq S(p_j \dots p_m) && \text{by Lemma 1(i)} \\ &= S(p_j) \\ &= S(p_i) \\ &= S(p_i \dots p_{j-1}). \end{aligned}$$

Therefore $(p_1 \dots p_{i-1}, p_i \dots p_{j-1})$ is a dominating prefix for w, contradiction. □

This theorem allows us to prove finiteness of the set of essential words by induction on the *codimension* $c(w) := \dim U - \dim T(w)$. It also enables us to derive upper bounds on their number and length, for which we need the following definition.

Definition 5

– *The* arrangement function $a : \mathbb{N} \to \mathbb{N}$ *is given by*

$$a(n) = \sum_{i=0}^{n} \frac{n!}{i!};$$

it is the number of sequences from a set of n elements with no repeated element. Note that $a(n)/n!$ converges to Euler's number e from below, in particular $a(n) \leq 3n!$ for all n.

– *The numbers N_i for $i \in \mathbb{N}$ are defined by* $N_i = \begin{cases} 1 & \text{for } i = 0 \\ |A| \cdot a(N_{i-1}) & \text{otherwise.} \end{cases}$

– *The numbers L_i for $i \in \mathbb{N}$ are defined by* $L_i = \begin{cases} 0 & \text{for } i = 0 \\ N_{i-1} L_{i-1} + 1 & \text{otherwise.} \end{cases}$

Theorem 2. *We have the following bounds on the numbers and lengths of essential words w, based on their codimension $c(w) = \dim U - \dim T(w)$.*

(i) There are at most N_i T-essential words w of codimension $c(w) \leq i$.
(ii) A T-essential word of codimension i has length $\leq L_i$.
(iii) There at most $N_{1+\dim U}$ essential words, and none of them is longer than $L_{1+\dim U}$.

Proof

(i) We proceed by induction on i. For $i = 0$, the empty word ϵ is trivially T-essential. Since $T(\epsilon) = U$ and $\epsilon \trianglelefteq w$ for all $w \in A^*$, it is the only T-essential word of dimension $\dim U$, i.e. codimension 0.
For $i \to i+1$, by the induction hypothesis there are at most N_i T-essential words of codimension $c(w) \leq i$. By Theorem 1, any T-essential word w of codimension $i+1$ has a factorization $w = p_1 \dots p_n M$ in which the p_j are pairwise distinct and T-essential of codimension $\leq i$, and $M \in A$. The number of such factorizations is at most $|A| \cdot a(N_i) = N_{i+1}$, and they include the ones for words of codimension $\leq i$, giving the upper bound N_{i+1} for the number of T-essential words w of codimension $c(w) \leq i+1$.
(ii) As in (i), we argue inductively using Theorem 1. For $i = 0$, the only T-essential word ϵ of codimension 0 has length $0 = L_0$.
For $i \to i+1$, the decomposition in Theorem 1 consists of at most N_i words, each of length at most L_i, plus the final letter, giving a total length of at most $N_i L_i + 1 = L_{i+1}$.
(iii) As in (i) and (ii), we get from Theorem 1 that any non-empty essential word has a factorization into (at most $N_{\dim U}$) pairwise T-essential words of length at most $L_{\dim U}$ and a single trailing $M \in A$. The number of such factorizations is at most $N_{1+\dim U}$, and their length is bounded by $L_{1+\dim U}$, by the same argument as before. □

3.2 Decidability of the Invariance Problem

The main result now follows immediately from the bounds established in the previous section.

Theorem 3. *Algorithm 1 terminates. It returns FAIL if and only if there is no solution to the invariance problem.*

Proof. Let $w = M_1 M_2 \dots \in A^\omega$ be any infinite word. By Theorem 2, w has an essential prefix $m(w) = M_1 \dots M_k$ of maximal length $|m(w)| \leq L_{1+\dim U}$ (note that the empty word is always essential, so $m(w)$ exists). Then $M_1 \dots M_{k+1}$ has a dominating prefix (p, q). We must have that $pq = m(w)$ and there is no shorter dominating prefix, since otherwise $m(w)$ would also fail to be essential. Due to the properties of dominating prefixes, $x_0 \in S(m(w))$ if and only if pq^ω is a solution to the invariance problem. Therefore we have two cases.

If there is no solution to the invariance problem, then for any branch w, x_0 cannot be in $S(m(w))$, i.e. the branch will be discarded before reaching depth $|m(w)| \leq L_{1+\dim U}$. By König's lemma, only finitely many words are explored, and the algorithm returns FAIL.

If there is a solution w, then $x_0 \in S(p)$ for every prefix $p \trianglelefteq w$, and the algorithm keep exploring until it reaches depth $|m(w)|$ for one such w, at which point the dominating prefix is discovered and returned. □

4 Further Remarks and Variations

4.1 Computing All Possible Initial Vectors

The algorithm we gave can be easily adapted to solve the following, more general problem:

- Given: A finite set $A \subseteq M_d(\mathbb{Q})$ of matrices, and a subspace $U \subseteq \mathbb{Q}^d$.
- Question: From which starting vectors x_0 can we remain in U indefinitely, using matrices from A? That is, for which $x_0 \in \mathbb{Q}^d$ does there exist a sequence $M_1, M_2, \ldots$ such that, for all $k \in \mathbb{N}$, $M_k \in A$ and $M_k \cdot \cdots \cdot M_1 x_0 \in U$?

Essentially, all that has to be changed is to remove the special treatment of x_0 and collect all the pairs which would have been returned into a set P, as in Fig. 3.

Data: finite set $A \subseteq M_d(\mathbb{Q})$ of matrices, invariant subspace $U \subseteq \mathbb{Q}^d$
Result: A set P of pairs $(u, v) \in A^* \times A^*$ such that $\bigcup_{(u,v)\in P} S(uv^\omega)$ contains exactly the initial vectors x_0 from which the invariance problem for A, x_0, U has a solution
if $A = \emptyset$ **then return** $\emptyset$;
branches := $\{\epsilon\}$;
while *branches not empty* **do**
  pop first w from branches;
  foreach *split* $w = uv$ *with* $v \neq \epsilon$ **do**
    if v *is pumpable* **then**
      Add (u, v) to P;
  if *no split for* w *was added* **then**
    foreach $M \in A$ **do**
      append wM to branches;
return P;

Fig. 3. Finding all starting vectors from which the invariance problem has a solution.

4.2 Locations

The transition system $(\mathbb{Q}^d, x_0, A)$ can be extended using a finite set L of control locations, giving $(L \times \mathbb{Q}^d, (l_0, x_0), T)$ for a finite set $T \subseteq L \times M_n(\mathbb{Q}) \times L$. The main properties of words $w \in A^*$ can still be used as before, the main difference being that a pumpable word q is now additionally required to label a cycle in the location graph.

The algorithm then proceeds as before, except that nodes are labeled with states $(l, x) \in L \times \mathbb{Q}^d$, and only successors using matrices which are available in l are considered.

4.3 The Universal Version

The question whether the invariant given by U is *valid*, i.e. whether $w(x_0) \in U$ for *all* $w \in A^*$, is also decidable, and is in fact easier than the problem we have dealt with in the previous sections. Again, the basic idea is to expand the tree of words $w \in A^*$ in a breadth-first order. While doing so, we can

(i) as soon as we encounter a counterexample $w(x_0) \notin U$, return w;
(ii) cut any branch v if $v(x_0)$ is a linear combination of vectors we have previously explored.

The reason for (ii) is that, if $y := v(x_0) = \lambda_1 x_1 + \cdots + \lambda_n x_n$, then due to linearity for any $w \in A^*$, $w(y) = \lambda_1 w(x_1) + \cdots + \lambda_n w(x_n)$; in particular, if $w(y) \notin U$ for some $w \in A^*$, then there is some i for which $w(x_i)$ is also not in U. Since we use breadth-first search, $x_i = u_i(x_0)$ for some word u_i which is smaller than v in the length-lexicographic order. Thus for any counterexample which we lose by discarding y, there is a length-lexicographically shorter one. In particular, if there are any counterexamples, then the length-lexicographically minimal one among them cannot be lost and will be found by the search.

Data: finite set $A \subseteq M_d(\mathbb{Q})$ of matrices,
initial vector $x_0 \in \mathbb{Q}^d$,
invariant subspace $U \subseteq \mathbb{Q}^d$
Result: A $w \in A^*$ such that $w(x_0) \notin U$, or SUCCESS if no such pair exists
if $A = \emptyset$ **then return** SUCCESS;
branches := $\{\epsilon\}$;
basis := $\{x_0\}$;
while *branches not empty* **do**
  pick and remove w from branches;
  if $w(x_0) \notin U$ **then return** w;
  foreach $M \in A$ **do**
    $y := wM(x)$;
    if y *linearly independent of basis* **then**
      add wM to branches;
      add y to basis
return SUCCESS;

Fig. 4. Checking validity of the invariant U, starting from x_0.

5 Summary and Future Work

5.1 Summary

We presented a solution to the invariance problem for matrix semigroups, i.e. the question of an infinite sequence of matrices satisfying a given linear invariant.

We gave an algorithm to find a solution if one exists, and proved its termination. The latter relied on the analysis of various properties of words in matrix semigroups and their interaction. In particular, we showed that for any finite set A of matrices and subspace U, there is a bound N such that any word of length $> N$ has a *dominating prefix*, reducing the problem to a finite one (Fig. 4).

5.2 Related Work

Matrix semigroups are a rich source of problems [3,11], including many surprisingly complicated decision problems. Among these are the *scalar reachability problem* (undecidable in dimension 3 and above [3]), which is the problem of reaching a subspace of codimension 1 from a starting vector using matrices from a given semigroup, and the *vector reachability problem*, in which the target is a single vector. A special case of the latter, for a single generating matrix, is the *orbit problem* which was shown to be decidable by Kannan and Lipton [5]. A closely related problem to the universal version of the invariance problem (Sect. 4.3) is the question of *boundedness*, which is undecidable [1].

One source of the complexity of such problems is that, in contrast to similar models like *vector addition systems* [4,6,9], the behaviour of matrix semigroups has inherently nonlinear aspects; for example, the simple 3-dimensional matrix $\begin{pmatrix} 1 & 2 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}$ has the orbit $\{(n^2, n, 1)^T \mid n \in \mathbb{N}\}$. This ability to reflect polynomial relationships connects problems like scalar reachability to known undecidable ones, like solvability of diophantine equations [10].

This complexity extends to various related models. Polynomial Register Machines [2] generalize vector addition systems with polynomial update functions (with integer coefficients); while this leads to undecidability in most cases, in dimension 1 reachability turns out to be decidable (in fact PSPACE-complete). Iteration of piecewise affine maps [7,8], similarly to matrix semigroups, involves a choice between affine-linear transformations, but this choice is deterministic based on the current variable values; the (vector) reachability problem is undecidable in general, but decidability is still open in dimension 1.

5.3 Future Work

There are a number of interesting ways to extend these results. Among them are:

Complexity Bounds. From the proof of decidability we get the upper bound $L_\infty \leq 1 + N_{\dim U}(1 + N_{\dim U - 1}(\ldots(1 + N_0)\ldots)$ for the exploration depth, where $N_0 = 1$ and $N_{i+1} \leq 3 \cdot |A| \cdot N_i!$, for a total of $\dim U$ nested factorials. This implies a complexity upper bound of $O(2^{L_\infty})$ for the algorithm. It would be interesting to see how much this can be improved, and what lower bounds can be found.

Polyhedral Invariants. The invariant in this paper was given as a linear subspace U, corresponding to a conjunction of linear equations. Generalizing this to linear inequalities leads to the question of whether we can find a sequence of matrices from A which allows us to remain inside a given polyhedron. This adds significant complications; in particular, it is not the case that a descending sequence of polyhedra $P_0 \supseteq P_1 \supseteq \ldots$ stabilizes after finitely many steps, so notions like the S-factorization cannot simply be translated to this setting. It is not clear at all whether this more general question is decidable.

Guards. The control structure of the transition system $(\mathbb{Q}^d, x_0, A)$ is relatively simple in that any transition is enabled at any time. Adding control location already changes this, but the values of the vector x in a state (l, x) still have no influence on the control flow. This changes with the introduction of application conditions or *guards*: linear equations or inequalities which have to be satisfied before the associated transition can be taken. Note that equation guards can be translated into extra dimensions in such a way that guard violations translate into (extended) invariance violations, so that they don't increase expressivity; inequalities on the other hand make for an interesting addition.

Games. Since both the existential and (as seen in Sect. 4.3) the universal version of the invariance problem are decidable, it is natural to ask what can be done about an alternating version. This would amount to considering games in which two players $\Box, \Diamond$ take turns applying matrices from given sets $A_\Box, A_\Diamond$ to the current vector x, with the goal of preserving and violating the invariant U, respectively.

Acknowledgements. The author is part supported by EPSRC project EP/K032011/1.

References

1. Blondel, V.D., Tsitsiklis, J.N.: The boundedness of all products of a pair of matrices is undecidable. Syst. Control Lett. **41**(2), 135–140 (2000)
2. Finkel, A., Göller, S., Haase, C.: Reachability in register machines with polynomial updates. In: Chatterjee, K., Sgall, J. (eds.) MFCS 2013. LNCS, vol. 8087, pp. 409–420. Springer, Heidelberg (2013)
3. Halava, V., Harju, T., Hirvensalo, M.: Undecidability bounds for integer matrices using claus instances. Technical report 766 (2006)
4. Hopcroft, J., Pansiot, J.-J.: On the reachability problem for 5-dimensional vector addition systems. Theor. Comput. Sci. **8**(2), 135–159 (1979)
5. Kannan, R., Lipton, R.J.: Polynomial-time algorithm for the orbit problem. J. ACM **33**(4), 808–821 (1986)
6. Karp, R.M., Miller, R.E.: Parallel program schemata. J. Comput. Syst. Sci. **3**(2), 147–195 (1969)
7. Kurganskyy, O., Potapov, I.: Reachability problems for pams. CoRR, abs/1510.04121 (2015)

8. Kurganskyy, O., Potapov, I., Sancho-Caparrini, F.: Reachability problems in low-dimensional iterative maps. Int. J. Found. Comput. Sci. **19**(04), 935–951 (2008)
9. Leroux, J.: Vector addition systems reachability problem (a simpler solution). In: Voronkov, A. (ed.)Turing-100. The Alan Turing Centenary. EasyChair Proceedings in Computing, vol. 10, pp. 214–228. EasyChair (2012)
10. Matiyasevich, Y.: Hilbert's Tenth Problem. With a foreword by Martin Davis. MIT Press, Cambridge (1993)
11. Potapov, I.: From post systems to the reachability problems for matrix semigroups and multicounter automata. In: Calude, C.S., Calude, E., Dinneen, M.J. (eds.) DLT 2004. LNCS, vol. 3340, pp. 345–356. Springer, Heidelberg (2004)