

The stability of saturated linear dynamical systems is undecidable

Vincent D. Blondel¹, Olivier Bournez², Pascal Koiran³, and John N. Tsitsiklis⁴

¹ Department of Mathematical Engineering, Université catholique de Louvain,
Avenue Georges Lemaitre 4, B-1348 Louvain-la-Neuve, Belgium,
blondel@inma.ucl.ac.be, <http://www.inma.ucl.ac.be/~blondel/>

² LORIA, Campus Scientifique BP 239, 54506 Vandoeuvre-les-Nancy, France,
Olivier.Bournez@loria.fr

³ LIP, ENS Lyon, 46 allée d'Italie, F-69364 Lyon Cedex 07, France,
Pascal.Koiran@ens-lyon.fr

⁴ LIDS, MIT, Cambridge, MA 02139, USA, jnt@mit.edu

Abstract. We prove that several global properties (global convergence, global asymptotic stability, mortality, and nilpotence) of particular classes of discrete time dynamical systems are undecidable. Such results had been known only for point-to-point properties. We prove these properties undecidable for saturated linear dynamical systems, and for continuous piecewise affine dynamical systems in dimension three. We also describe some consequences of our results on the possible dynamics of such systems.

1 Introduction

This paper studies problems such as the following: given a discrete time dynamical system of the form $x_{t+1} = f(x_t)$, where $f : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is a saturated linear function or, more generally, a continuous piecewise affine function, decide whether all trajectories converge to the origin.

We show in our main theorem that this global convergence problem is undecidable. The same is true for three related problems: Stability (is the dynamical system globally asymptotically stable?), Mortality (do all trajectories go through the origin?), and Nilpotence (does there exist an iterate f^k of f such that $f^k \equiv 0$?).

It is well-known that various types of dynamical systems, such as hybrid systems, piecewise affine systems, or saturated linear systems, can simulate Turing machines, see e.g., [15, 12, 16, 18]. In these simulations, a machine configuration is encoded by a point in the state space of the dynamical system. It then follows that *point-to-point* properties of such dynamical systems are undecidable. For example, given a point in the state space, one cannot decide whether the trajectory starting from this point eventually reaches the origin. The results described in this contribution are of a different nature since they deal with *global* properties of dynamical systems.

Related undecidability results for such global properties have been obtained in our earlier work [3], but for the case of *discontinuous* piecewise affine systems. The additional requirement of continuity imposed in this paper is a severe restriction, and makes undecidability much harder to establish. Surveys of decidability and complexity results for dynamical systems are given in [1], [12] and [8].

Our main result (Theorem 1) is a proof of Sontag’s conjecture [7, 19] that global asymptotic stability of saturated linear systems is not decidable. Saturated linear systems are systems of the form $x_{t+1} = \sigma(Ax_t)$ where x_t evolves in the state space \mathbf{R}^n , A is a square matrix, and σ denotes componentwise application of the saturated linear function $\sigma : \mathbf{R} \rightarrow [-1, 1]$ defined as follows: $\sigma(x) = x$ for $|x| \leq 1$, $\sigma(x) = 1$ for $x \geq 1$, $\sigma(x) = -1$ for $x \leq -1$. These dynamical systems occur naturally as models of neural networks [17, 18] or as models of simple hybrid systems [20, 5, 2].

Theorem 1 is proved in three main steps. First, in Section 4, we prove that any Turing machine can be simulated by a saturated linear dynamical system with a strong notion of simulation. (Turing machines are defined in Section 3.) Then, in Section 5, using a result of Hooper, we prove that there is no algorithm that can decide whether a given continuous piecewise affine system has a trajectory contained in a given hyperplane. Finally, we prove Theorem 1 in Section 6.

In light of our undecidability result, any decision algorithm for the stability of saturated linear systems will be able to handle only special classes of systems. In the full version of this paper [4] we consider two such classes: systems of the form $x_{t+1} = \sigma(Ax_t)$ where A is a nilpotent matrix, or a symmetric matrix. We show that stability remains undecidable for the first class, but is decidable for the second.

Saturated linear systems fall within the class of continuous piecewise affine systems and so our undecidability results equally apply to the latter class of systems. More precise statements for continuous piecewise affine systems are given in Section 7. Finally, some suggestions for further work are made in Section 8.

For some of our results we give complete proofs. For others we provide only a sketch, or we refer to the full version of the paper [4].

2 Dynamical systems

In the sequel, X denotes a metric space and 0 some arbitrary point of X , to be referred to as the *origin*. When $X \subseteq \mathbf{R}^n$, we assume that 0 is the usual origin of \mathbf{R}^n . A *neighborhood* of 0 is an open set that contains 0 . Let $f : X \rightarrow X$ be a function such that $f(0) = 0$. We say that f is:

- (a) *globally convergent* if for every initial point $x_0 \in X$, the trajectory $x_{t+1} = f(x_t)$ converges to 0 .
- (b) *locally asymptotically stable* if for any neighborhood U of 0 , there is another neighborhood V of 0 such that for every initial point $x_0 \in V$, the trajectory $x_{t+1} = f(x_t)$ converges to 0 without leaving U (i.e., $x(t) \in U$ for all $t \geq 0$ and $\lim_{t \rightarrow \infty} x_t = 0$).

- (c) *globally asymptotically stable* if f is globally convergent and locally asymptotically stable.
- (d) *mortal* if for every initial point $x_0 \in X$, there exists $t \geq 0$ with $x_t = 0$. The function f is called *immortal* if it is not mortal.
- (e) *nilpotent* if there exists $k \geq 1$ such that the k -th iterate of f is identically equal to 0 (i.e., $f^k(x) = 0$ for all $x \in X$).

Nilpotence obviously implies mortality, which implies global convergence; and global asymptotic stability also implies global convergence. In general, this is all that can be said of the relations between these properties. Note, however, the following simple lemma, which will be used repeatedly.

Lemma 1. *Let X be a metric space with origin 0, and let $f : X \rightarrow X$ be a continuous function such that $f(0) = 0$. If f is nilpotent, then it is globally asymptotically stable. Moreover, if X is compact and if there exists a neighbourhood O of 0 and an integer $j \geq 1$ such that $f^j(O) = \{0\}$, the four properties of nilpotence, mortality, global asymptotic stability, and global convergence are equivalent.*

Proof. Assume that f is nilpotent and let k be such that $f^k \equiv 0$. Let U and V be two neighborhoods of 0. A trajectory starting in V never leaves $\bigcup_{i=0}^{k-1} f^i(V)$. By continuity, for any U one can choose V so that $f^i(V) \subseteq U$ for all $i = 0, \dots, k-1$. A trajectory originating in such a V never leaves U . This shows that f is globally asymptotically stable.

Next, assume that X is compact and that $f^j(O) = \{0\}$ for some neighborhood O of 0 and some integer $j \geq 1$. It suffices to show that if f is globally convergent, then it is nilpotent. If f is globally convergent, then $X = \bigcup_{i \geq 0} f^{-i}(O)$. By compactness, there exists $p \geq 0$ such that $X = \bigcup_{i=0}^p f^{-i}(O)$. We conclude that $f^{p+j}(X) = \{0\}$. \square

A function $f : \mathbf{R}^n \rightarrow \mathbf{R}^{n'}$ is *piecewise affine* if \mathbf{R}^n can be represented as the union of a finite number of subsets X_i where each set X_i is defined by the intersection of finitely many open or closed halfspaces of \mathbf{R}^n , and the restriction of f to each X_i is affine. Let $\sigma : \mathbf{R} \rightarrow \mathbf{R}$ be the continuous piecewise affine function defined by: $\sigma(x) = x$ for $|x| \leq 1$, $\sigma(x) = 1$ for $x \geq 1$, $\sigma(x) = -1$ for $x \leq -1$. Extend σ to a function $\sigma : \mathbf{R}^n \rightarrow \mathbf{R}^n$, by letting $\sigma(x_1, \dots, x_n) = (\sigma(x_1), \dots, \sigma(x_n))$. A *saturated affine function* (σ -function for short) $f : \mathbf{R}^n \rightarrow \mathbf{R}^{n'}$ is a function of the form $f(x) = \sigma(Ax + b)$ for some matrix $A \in \mathbf{Q}^{n' \times n}$ and vector $b \in \mathbf{Q}^{n'}$. Note that we are restricting the entries of A and b to be rational numbers so that we can work within the Turing model of digital computation. A *saturated linear function* (σ_0 -function for short) is defined similarly except that $b = 0$. Note that the function $\sigma : \mathbf{R} \rightarrow \mathbf{R}$ is piecewise affine, with the polyhedra X_i corresponding to the different faces of the unit cube $[-1, 1]^n$, and so is the linear function $f(x) = Ax$. It is easily seen that the composition of piecewise affine functions is also piecewise affine and therefore σ -functions are piecewise affine.

Our main result is the following theorem.

Theorem 1. *The problems of determining whether a given saturated linear function is (i) globally convergent, (ii) globally asymptotically stable, (iii) mortal, or (iv) nilpotent, are all undecidable.*

Notice that deciding the global asymptotic stability of a saturated linear system is *a priori* no harder than deciding its global convergence, because the local asymptotic stability of saturated linear systems is decidable. (Indeed, a system $x_{t+1} = \sigma(Ax_t)$ is locally asymptotically stable if and only if the system $x_{t+1} = Ax_t$ is, since these systems are identical in a neighborhood of the origin. Furthermore, a linear system is locally asymptotically stable if and only if all of its eigenvalues have magnitude less than one [21].) In fact, we conjecture that for saturated linear systems, global convergence is equivalent to global asymptotic stability. This equivalence is proved for symmetric matrices in the full version of the paper. If this conjecture is true, it is not hard to see that the equivalence of mortality and nilpotence also holds.

Theorem 1 has some “purely mathematical” consequences. For instance:

Corollary 1. *For infinitely many integers n , there exists a nilpotent saturated linear function $f : \mathbf{R}^n \rightarrow \mathbf{R}^n$ such that $f^{2^n} \neq 0$.*

Of course, in this corollary, 2^n can be replaced by any recursive function of n . In contrast, if $f : \mathbf{R}^n \rightarrow \mathbf{R}^n$ is a nilpotent linear function, then $f^n \equiv 0$. As a side remark, we note that it can be shown that this is not only true for linear functions, but also for polynomials and even more generally for real analytic functions.

We conclude this section with two positive results: globally asymptotically stable saturated linear systems are recursively enumerable and so are saturated linear systems that have a nonzero periodic trajectory. The first observation is due to Eduardo Sontag, the second is due to Alexander Megretski.

Theorem 2. *The set of saturated linear systems that are globally asymptotically stable is recursively enumerable.*

Theorem 3. *The set of saturated linear systems that have a nonzero periodic trajectory is recursively enumerable.*

The proofs of these results are based on elementary arguments, they can be found in the full version of the paper. Combining these two observations with Theorem 1, we deduce that there exist saturated linear systems that are not globally asymptotically stable and have no nonzero periodic trajectories.

Corollary 2. *There exist saturated linear systems that are not globally asymptotically stable and have no nonzero periodic trajectory.*

3 Turing machines

A *Turing machine* M [14, 13] is an abstract deterministic computer with a finite set Q of internal states. It operates on a doubly-infinite tape over some finite

alphabet Σ . The tape consists of squares indexed by an integer i , $-\infty < i < \infty$. At any time, the Turing machine scans the square indexed by 0. Depending upon its internal state and the scanned symbol, it can perform one or more of the following operations: replace the scanned symbol with a new symbol, focus attention on an adjacent square (by shifting the tape by one unit), and transfer to a new state.

The instructions for the Turing machine are quintuples of the form

$$[q_i, s_j, s_k, D, q_l]$$

where q_i and s_j represent the present state and scanned symbol, respectively, s_k is the symbol to be printed in place of s_j , D is the direction of motion (left-shift, right-shift, or no-shift of the tape), and q_l is the new internal state. For consistency, no two quintuples can have the same first two entries. If the Turing machine enters a state-symbol pair for which there is no corresponding quintuple, it is said to *halt*.

Without loss of generality, we can and will assume that $\Sigma = \{0, 1, \dots, n-1\}$, $Q = \{0, 1, \dots, m-1\}$, $n, m \in \mathbf{N}$, and that the Turing machine halts if and only if the internal state q is equal to zero. We refer to $q = 0$ as the *accepting* state.

The tape contents can be described by two infinite words $w_1, w_2 \in \Sigma^\omega$, where Σ^ω stands for the set of infinite words over the alphabet Σ : w_1 consists of the scanned symbol and the symbols to its right; w_2 consists of the symbols to the left of the scanned symbol, excluding the latter. The tape contents (w_1, w_2) , together with an internal state $q \in Q$, constitute a *configuration* of the Turing machine. If a quintuple applies to a configuration (that is, if $q \neq 0$), the result is another configuration, a successor of the original. Otherwise, if no quintuple applies (that is, if $q = 0$), we have a *terminal* configuration. We thus obtain a successor function $\vdash: C \rightarrow C$, where $C = \Sigma^\omega \times \Sigma^\omega \times Q$ is the set of all configurations (the configuration space). Note that \vdash is a partial function, as it is undefined when $q = 0$. A configuration is said to be *mortal* if repeated application of the function \vdash eventually leads to a terminal configuration. Otherwise, the configuration is called *immortal*. We shall say that a Turing machine M is *mortal* if all configurations are mortal, and that it is *nilpotent* if there exists an integer k such that M halts in at most k steps starting from any configuration.

Theorem 4. *A Turing machine is mortal if and only if it is nilpotent.*

Proof. A nilpotent Turing machine is mortal, by definition. The converse will follow from Lemma 1. In order to apply that lemma, we endow the configuration space of a Turing machine with a topology which makes its successor function \vdash continuous, and its configuration space (X, d) compact. This is a fairly standard construction and we refer the reader to the full version of the paper for a complete description. The constructed function \vdash is identically equal to 0 in a neighborhood of 0. We therefore conclude from Lemma 1 that if M is mortal, then it must be nilpotent. \square

The next result is due to Hooper and will play a central role in the sequel.

Theorem 5 ([13]). *The problem of determining whether a given Turing machine is mortal is undecidable.*

In other words, one cannot decide whether a given Turing machine halts for every initial configuration. Equivalently, one cannot decide whether there exists an immortal configuration.

4 Turing machine simulation

A σ^* -function is a function obtained by composing finitely many σ -functions. It is well known that Turing machines can be simulated by piecewise affine dynamical systems [15, 16, 18]. Moreover, this simulation can be performed with a σ^* -function (see the full version of the paper for the details of the construction of this function).

Lemma 2 ([15, 16, 18]). *Let M be a Turing machine and let $C = \Sigma^\omega \times \Sigma^\omega \times Q$ be its configuration space. There exists a σ^* -function $g_M : \mathbf{R}^2 \rightarrow \mathbf{R}^2$ and an encoding function $\nu : C \rightarrow [0, 1]^2$ such that the following diagram commutes:*

$$\begin{array}{ccc} C & \xrightarrow{\vdash} & C \\ \nu \downarrow & & \downarrow \nu \\ \mathbf{R}^2 & \xrightarrow{g_M} & \mathbf{R}^2 \end{array}$$

(i.e. $g_M(\nu(c)) = \nu(c')$ for all configurations $c, c' \in C$ with $c \vdash c'$).

We extend this results by proving that any Turing machine can be simulated by a dynamical system in a stronger sense.

Lemma 3. *Let M be a Turing machine and let $C = \Sigma^\omega \times \Sigma^\omega \times Q$ be its configuration space. Then, there exists a σ^* -function $g_M : \mathbf{R}^2 \rightarrow \mathbf{R}^2$, a decoding function $\nu' : [0, 1]^2 \rightarrow C$, and some subsets $\mathcal{N}^\infty \subset \mathcal{N}^1 \subset [0, 1]^2$, $\mathcal{N}_{-acc}^1 \subset \mathcal{N}^1$ such that the following conditions hold:*

1. $g_M(\mathcal{N}^\infty) \subseteq \mathcal{N}^\infty$ and $\nu'(\mathcal{N}^\infty) = C$.
2. \mathcal{N}_{-acc}^1 (respectively \mathcal{N}^1) is the Cartesian product of two finite unions of closed intervals in \mathbf{R} . \mathcal{N}_{-acc}^1 is at a positive distance from the origin $(0, 0)$ of \mathbf{R}^2 .
3. For $x \in \mathcal{N}^1$, the configuration $\nu'(x)$ is nonterminal if and only if $x \in \mathcal{N}_{-acc}^1$.
4. The following diagram commutes:

$$\begin{array}{ccc} C & \xrightarrow{\vdash} & C \\ \nu' \uparrow & & \uparrow \nu' \\ \mathcal{N}_{-acc}^1 & \xrightarrow{g_M} & [0, 1]^2 \end{array}$$

(i.e. $\nu'(x) \vdash \nu'(g_M(x))$ for all $x \in \mathcal{N}_{-acc}^1$).

Intuitively, ν' is an inverse of the encoding function ν of Lemma 2, in the sense that $\nu'(\nu(c)) = c$ holds for all configurations c . The set \mathcal{N}^∞ is the image of the function ν , consisting of those points $x \in [0, 1]^2$ that are unambiguously associated with valid configurations of the Turing machine. The set \mathcal{N}^1 consists of those points that lie in some set $B_{\alpha, \beta, q}$ and therefore encode an internal state q , a scanned symbol α , and a symbol β to the left of the scanned one. (However, not all points in \mathcal{N}^1 are images of valid configurations. Once it encounters a “decoding failure” our decoding function ν' sets the corresponding tape square, and all subsequent ones to the zero symbol.) Finally, \mathcal{N}_{-acc}^1 is the subset of \mathcal{N}^1 associated with the nonterminal internal states $q \neq 0$. See the full paper for complete details.

Using Lemma 3 and Theorem 5, we can now prove:

Theorem 6. *The problems of determining whether a given (possibly discontinuous) piecewise affine function in dimension 2 is (i) globally convergent, (ii) globally asymptotically stable, (iii) mortal, or (iv) nilpotent, are all undecidable.*

The undecidability of the first three properties was first established in [3]. That proof was based on an undecidability result for the mortality of counter machines, instead of Turing machines.

Proof. We use a reduction from the problem of Theorem 5. Suppose that a Turing machine M is given. Denote by g'_M the discontinuous function which is equal to the function g_M of Lemma 3 on \mathcal{N}_{-acc}^1 , and which is equal to 0 outside of \mathcal{N}_{-acc}^1 .

Since 0 is at a positive distance from \mathcal{N}_{-acc}^1 , we have a neighborhood O of 0 such that $g'_M(O) = \{0\}$. By Lemma 1, all four properties in the statement of the theorem are equivalent.

Assume first that M is mortal. By Theorem 4, there exists k such that M halts on any configuration in at most k steps. We claim that $g'_M{}^{k+1}([0, 1]^2) = \{0\}$. Indeed, assume, in order to derive a contradiction, that there exists a trajectory $x_{t+1} = g'_M(x_t)$ with $x_{k+1} \neq 0$. Since g'_M is zero outside \mathcal{N}_{-acc}^1 , we have $x_t \in \mathcal{N}_{-acc}^1$ for $t = 0, \dots, k$. By the commutative diagram of Lemma 3, the sequence $c_t = \nu'(x_t)$ ($t = 0, \dots, k+1$) is a sequence of successive configurations of M . This contradicts the hypothesis that M reaches a terminal configuration after at most k steps. It follows that g'_M satisfies properties (i) through (iv).

Conversely, suppose that M has an immortal configuration: there exists an infinite sequence c_t of non-terminal configurations with $c_t \vdash c_{t+1}$ for all $t \in \mathbf{N}$. By condition 1 of Lemma 3, there exists $x_0 \in \mathcal{N}^\infty$ with $\nu'(x_0) = c_0$. We claim that the trajectory $x_{t+1} = g'_M(x_t)$ is immortal: using condition 2 of Lemma 3, it suffices to prove that $x_t \in \mathcal{N}_{-acc}^1$ for all t . Indeed, we prove by induction on t that $x_t \in \mathcal{N}_{-acc}^1 \cap \mathcal{N}^\infty$ and $\nu'(x_t) = c_t$ for all t . Using condition 3 of Lemma 3, the induction hypothesis is true for $t = 0$. Assuming the induction hypothesis for t , condition 1 of Lemma 3 shows that $x_{t+1} \in \mathcal{N}^\infty$. Now, the commutative diagram of Lemma 3 shows that $\nu'(x_{t+1}) = c_{t+1}$, and condition 3 of Lemma 3 shows that $x_{t+1} \in \mathcal{N}_{-acc}^1$. This completes the induction. Hence, g'_M is not mortal, and therefore does not satisfy any of the properties (i) through (iv).

5 The hyperplane problem

We now reach the second step of our proof. Using the undecidability result of Hooper for the mortality of Turing machines, we prove that it cannot be decided whether a given piecewise affine system has a trajectory that stays forever in a given hyperplane.

Theorem 7. *The problem of determining if a given σ^* -function $f : \mathbf{R}^3 \rightarrow \mathbf{R}^3$ has a trajectory $x_{t+1} = f(x_t)$ that belongs to $\{0\} \times \mathbf{R}^2$ for all t is undecidable*

Proof. We reduce the problem of Theorem 5 to this problem.

Suppose that a Turing Machine M is given. Consider the σ^* -function $f : \mathbf{R}^3 \rightarrow \mathbf{R}^3$ defined by

$$f(x_1, x_2, x_3) = \begin{pmatrix} \sigma(\sigma(Z_{\mathcal{N}_{-acc}^1}(x_2, x_3))) \\ g_M(x_2, x_3) \end{pmatrix}$$

where g_M is the function constructed in Lemma 3 and $Z_{\mathcal{N}_{-acc}^1}$ is a σ^* -function that is equal to zero for $x \in \mathcal{N}_{-acc}^1$ and is otherwise positive (an explicit construction of this function is provided in the full version of the paper). Note that in the definition of the function f we use a nested application of the function σ . This is to ensure that the definition of f involves an equal number of applications of the σ function on all its components.

Write (x^1, \dots, x^d) for the components of a point x of \mathbf{R}^d .

We prove that f has a trajectory $x_{t+1} = f(x_t)$ with $x_t^1 = 0$ for all t , if and only if Turing machine M has an immortal configuration.

Suppose that f has such a trajectory. Since $Z_{\mathcal{N}_{-acc}^1}$, and hence $\sigma(\sigma(Z_{\mathcal{N}_{-acc}^1}))$, is strictly positive outside of \mathcal{N}_{-acc}^1 , we must have $(x_t^2, x_t^3) \in \mathcal{N}_{-acc}^1$ for all $t \geq 0$. By the commutative diagram of Lemma 3, the sequence $\nu'(x_t^2, x_t^3)$, $t \in \mathbf{N}$, is a sequence of successive configurations of M . By condition 3 of Lemma 3, none of these configurations is terminal, i.e. $c_0 = \nu'(x_0^2, x_0^3)$ is an immortal configuration of M .

Conversely, assume that M has an immortal configuration, that is, there exists an infinite sequence of nonterminal configurations with $c_t \vdash c_{t+1}$. The argument here is the same as in the proof of Theorem 6. By condition 1 of Lemma 3, there exists a point $(x_0^2, x_0^3) \in \mathcal{N}^\infty$ with $\nu'(x_0^2, x_0^3) = c_0$. Consider the sequence defined by $(x_{t+1}^2, x_{t+1}^3) = g_M(x_t^2, x_t^3)$ for all t . Since $g_M(\mathcal{N}^\infty) \subseteq \mathcal{N}^\infty$, we have $(x_t^2, x_t^3) \in \mathcal{N}^\infty$ for all $t \geq 0$. Using the assumption that configuration c_t is nonterminal and condition 3 of Lemma 3, we deduce that $(x_t^2, x_t^3) \in \mathcal{N}_{-acc}^1$ for all $t \geq 0$, which means precisely that the sequence $x_t = (0, x_t^2, x_t^3)$, $t \in \mathbf{N}$, is a trajectory of f . \square

6 Proof of the main theorem

We now reach the last step in the proof, which consists of reducing the problem of Theorem 7 to the problems of Theorem 1.

Recall that a σ -function is a function of the form $f(x) = \sigma(Ax + b)$ and a σ_0 -function is a function of the form $f(x) = \sigma(Ax)$. A composition of finitely many σ_0 -functions is called a σ_0^* -function.

Lemma 4. *The problems of determining whether a given σ_0^* -function $\mathbf{R}^4 \rightarrow \mathbf{R}^4$ is (i) globally convergent, (ii) globally asymptotically stable, (iii) mortal, or (iv) nilpotent, are all undecidable.*

Proof. The problem of Theorem 7 can be reduced to the mortality problem for σ_0^* -functions. The construction is such that the σ_0^* -function is equal to zero in an neighborhood of the origin (see the full paper for the construction of the function). It therefore follows from Lemma 1 that for this function, the properties (i)-(iv) are equivalent. These four properties are therefore undecidable. \square

We can now prove Theorem 1.

Proof. (of Theorem 1) We reduce the problems in Lemma 4 to the problems in Theorem 1.

Let $f : \mathbf{R}^4 \rightarrow \mathbf{R}^4$ be a σ_0^* -function of the form $f = f_k \circ f_{k-1} \circ \dots \circ f_1$ for some σ_0 -functions $f_j(x) = \sigma(A_j x)$, where $f_j : \mathbf{R}^{d_{j-1}} \rightarrow \mathbf{R}^{d_j}$ with $d_0, d_1, \dots, d_k \in \mathbf{N}$, and $d_0 = d_k = 4$.

Let $d = d_0 + d_1 + \dots + d_k$, and consider the saturated linear function $f' : \mathbf{R}^d \rightarrow \mathbf{R}^d$ defined by $f'(x) = \sigma(Ax)$ where

$$A = \begin{pmatrix} 0 & 0 & \dots & 0 & A_k \\ A_1 & 0 & \dots & 0 & 0 \\ 0 & A_2 & \dots & 0 & 0 \\ \vdots & \vdots & & 0 & 0 \\ 0 & 0 & \dots & A_{k-1} & 0 \end{pmatrix}$$

Clearly, the iterates of this function simulate the iterates of the function f .

Suppose that f' is mortal (respectively nilpotent, globally convergent, globally asymptotically stable). Then, the same is true for f : indeed, when $x_{t+1} = f(x_t)$ is a trajectory of f , the sequence $(x_t, f_1(x_t), \dots, f_{k-1} \circ \dots \circ f_1(x_t))$ is a subsequence of a trajectory of f' .

Conversely, let $x'_{t+1} = f'(x'_t)$ be a trajectory of f' . Write $x'_t = (y_t^1, \dots, y_t^k)$ with each of the y^j in $\mathbf{R}^{d_{j-1}}$. For every $t_0 \in \{0, \dots, k-1\}$ and $j \in \{1, \dots, k\}$, the sequence $t \mapsto y'_{t_0+kt}$ is a trajectory of f . This implies that the sequence y'_t , $t \in \mathbf{N}$ is eventually null (respectively, converges to 0) if f is mortal (respectively, globally convergent). For the same reason, the global asymptotic stability of f implies that of f' ; and if $f^m \equiv 0$ for some integer m , we have $(f')^{km} \equiv 0$. \square

7 Continuous piecewise affine systems

We proved in Theorem 6 that it cannot be decided whether a given *discontinuous* piecewise affine system of dimension 2 is globally convergent, globally asymptotically stable, mortal, or nilpotent. We do not know whether these problems

remain undecidable when the systems are of dimension 1.

For continuous systems, we can prove the following.

Theorem 8. *For continuous piecewise affine systems in dimension 3, the four properties of global convergence, global asymptotic stability, mortality, and nilpotence are undecidable.*

Proof. The system built in the proof of Lemma 4 is of dimension 4. The construction can be adapted to a system of dimension 3. See the full paper. \square

The following proposition is proved in [3].

Theorem 9. *For continuous piecewise affine systems in dimension 1, the properties of global convergence, global asymptotic stability, and mortality are decidable.*

One can also show that nilpotence is decidable for this class of systems. Thus, all properties are decidable for continuous piecewise affine systems in dimension 1, and are undecidable in dimension 3. The situation in dimension 2 has not been settled.

Global properties of $f : \mathbf{R}^n \rightarrow \mathbf{R}^n$	$n = 1$	$n = 2$	$n = 3$
Piecewise affine	?	Undecidable	Undecidable
Continuous piecewise affine	Decidable	?	Undecidable

8 Final remarks

In addition to the two question marks in the table of the previous section, several questions which have arisen in the course of this work still await an answer:

1. Does there exist some fixed dimension n such that nilpotence (or mortality, global asymptotic stability and global convergence) of saturated linear systems of dimension n is undecidable? A negative answer would be somewhat surprising since there would be in that case a decision algorithm for each n , but no single decision algorithm working for all n .
2. It would be interesting to study the decidability of these four properties for other special classes of saturated linear systems, as we have already done for nilpotent and symmetric matrices. For instance, is global convergence or global asymptotic stability decidable for systems with invertible matrices? (Note that such a system cannot be nilpotent or mortal.) Are some of the global properties decidable for matrices with entries in $\{-1, 0, 1\}$?
3. For saturated linear systems, is mortality equivalent to nilpotence? Is global convergence equivalent to global asymptotic stability? (This last equivalence is conjectured in Section 2.) We show in the full version of the paper that these equivalences hold for systems with symmetric matrices.

4. For a polynomial map $f : \mathbf{R}^n \rightarrow \mathbf{R}^n$ mortality is equivalent to nilpotence; these properties are equivalent to the condition $f^n \equiv 0$, and hence decidable. It is however not clear whether the properties of global asymptotic stability and global convergence are equivalent, or decidable.
5. Does there exist a dimension n such that for any integer k there exists a nilpotent saturated linear system $f : \mathbf{R}^n \rightarrow \mathbf{R}^n$ such that $f^k \neq 0$? Note that this question (and some of the other questions) still makes sense if we allow matrices with arbitrary real (instead of rational) entries.

Acknowledgments

Thanks are due to Eduardo Sontag for pointing out Theorem 2 and formulating his undecidability conjecture. P. K. also acknowledges a useful discussion with Eduardo Sontag on the nilpotence of analytic functions; J. T. and V. B. would like to thank Alexander Megretski for helpful comments.

This work was supported in part by the NATO under grant CRG-961115, by the National Science Foundation under grant ECS-9873451 and by the European Commission under the TMR Alapedes network and the Esprit working group NeuroCOLT2.

References

1. R. Alur, C. Courcoubetis, N. Halbwachs, T. A. Henzinger, P. H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138(1):3–34, 6 February 1995.
2. E. Asarin, O. Maler, and A. Pnueli. Reachability analysis of dynamical systems having piecewise-constant derivatives. *Theoretical Computer Science*, 138(1):35–65, February 1995.
3. V. D. Blondel, O. Bournez, P. Koiran, C. Papadimitriou, and J. N. Tsitsiklis. Deciding stability and mortality of piecewise affine dynamical systems. Submitted. Available from <http://www.inma.ucl.ac.be/~blondel/publications>.
4. V. D. Blondel, O. Bournez, P. Koiran, and J. N. Tsitsiklis. The stability of saturated linear dynamical systems is undecidable. Submitted. Available from <http://www.inma.ucl.ac.be/~blondel/publications>.
5. V. D. Blondel and J. N. Tsitsiklis. Complexity of stability and controllability of elementary hybrid systems. *Automatica*, 35(3), 479–489, 1999.
6. V. D. Blondel and J. N. Tsitsiklis. Overview of complexity and decidability results for three classes of elementary nonlinear systems. In Y. Yamamoto and S. Hara, editors, *Learning, Control and Hybrid Systems*, pages 46–58. Springer-Verlag, London, 1999.
7. V. D. Blondel and J. N. Tsitsiklis. Three problems on the decidability and complexity of stability. In V. D. Blondel, E. D. Sontag, M. Vidyasagar, and J. C. Willems, editors, *Open Problems in Mathematical Systems and Control Theory*, pages 45–52. Springer-Verlag, London, 1999.
8. V. D. Blondel and J. N. Tsitsiklis. A survey of computational complexity results in systems and control. To appear in *Automatica*, 1999.

9. O. Bournez. Complexité algorithmique des systèmes dynamiques continus et hybrides. PhD Thesis, Ecole Normale Supérieure de Lyon, 1999.
10. L. Blum, F. Cucker, M. Shub, and S. Smale. *Complexity and Real Computation*. Springer-Verlag, 1998.
11. L. Blum, M. Shub, and S. Smale. On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bulletin of the American Mathematical Society*, 21(1):1–46, July 1989.
12. T. A. Henzinger, P. W. Kopke, A. Puri, and P. Varaiya. What’s decidable about hybrid automata? *Journal of Computer and System Sciences*, 57(1):94–124, August 1998.
13. P. K. Hooper. The undecidability of the Turing machine immortality problem. *The Journal of Symbolic Logic*, 31(2): 219–234, June 1966.
14. J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory Languages and Computation*. Addison-Wesley, October 1979.
15. P. Koiran, M. Cosnard, and M. Garzon. Computability with low-dimensional dynamical systems. *Theoretical Computer Science*, 132(1-2):113–128, September 1994.
16. C. Moore. Generalized shifts: unpredictability and undecidability in dynamical systems. *Nonlinearity*, 4:199–230, 1991.
17. H. T. Siegelmann and E. D. Sontag. Analog computation via neural networks. *Theoretical Computer Science*, 131(2):331–360, September 1994.
18. H. T. Siegelmann and E. D. Sontag. On the computational power of neural nets. *Journal of Computer and System Sciences*, 50(1):132–150, February 1995.
19. E. D. Sontag. From linear to nonlinear: Some complexity comparisons. In *IEEE Conference on Decision and Control*, pages 2916–2920, New Orleans, December 1995.
20. E. D. Sontag. Interconnected automata and linear systems: A theoretical framework in discrete time. In R. Alur, T. A. Henzinger, and E. D. Sontag, editors, *Hybrid Systems III*, volume 1066 of *Lecture Notes in Computer Science*. Springer-Verlag, 1996.
21. E. D. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Springer-Verlag, New York, second edition, August 1998.