The stability of saturated linear dynamical systems is undecidable

Vincent Blondel[1], Olivier Bournez[2], Pascal Koiran[3] and John N. Tsitsiklis[4]

[1] *CESAME, Université catholique de Louvain, 4 avenue Georges Lemaitre, B-1348 Louvain-la-Neuve, Belgium; Email: blondel@inma.ucl.ac.be.*

[2] *LORIA and INRIA-Lorraine, Campus Scientifique, 615 rue du Jardin Botanique BP-101, F-54602 Villers-lès-Nancy, France; Email: Olivier.Bournez@loria.fr*

[3] *LIP, ENS Lyon, 46 allée d'Italie, F-69364 Lyon Cedex 07, France; Email: Pascal.Koiran@ens-lyon.fr*

[4] *LIDS, MIT, Cambridge, MA 02139, USA; Email: jnt@mit.edu*

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.

Keywords: Dynamical systems, saturated linear systems, piecewise affine systems, hybrid systems, mortality, stability, decidability.

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 $f : \mathbf{R}^n \to \mathbf{R}^n$ 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., [18, 15, 19, 21]. 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 [5], 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], [15] and [9].

Our main result (Theorem 2.1) is a proof of Sontag's conjecture [8, 22] 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 $\sigma : \mathbf{R} \to [-1, 1]$ 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 are often used as artificial neural network models [20, 21] or as models of simple hybrid systems [23, 6, 2].

Theorem 2.1 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 2.1 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 Section 6 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.

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 $f : X \to X$ be a function such that $f(0) = 0$. We say that f is:

- *globally convergent* if for every initial point $x_0 \in X$, the trajectory $x_{t+1} = f(x_t)$ converges to 0.

- *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 $\lim_{t\to\infty} x_t = 0$).

- *globally asymptotically stable* if f is globally convergent and locally asymptotically stable.

- *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.

- *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 2.1. *Let X be a metric space with origin 0, and let $f : X \to X$ 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 $i = 0, \ldots, k-1$. 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 $f^{p+j}(X) = \{0\}$. ■

Figure 1: Graph of the function σ.

A function $f : \mathbf{R}^n \to \mathbf{R}^{n'}$ 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 $\sigma : \mathbf{R} \to \mathbf{R}$ 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 $\sigma : \mathbf{R}^n \to \mathbf{R}^n$, by letting $\sigma(x_1, \ldots, x_n) = (\sigma(x_1), \ldots, \sigma(x_n))$. A *saturated affine function* (*σ-function* for short) $f : \mathbf{R}^n \to \mathbf{R}^{n'}$ 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. Using arbitrary real entries would give rise to systems whose computational power is related to non-uniform complexity classes: see [13].

A *saturated linear function* (*σ_0-function* for short) is defined similarly except that $b = 0$. Note that the function $\sigma : \mathbf{R}^n \to \mathbf{R}^n$ is piecewise affine 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 2.1. *The problems of determining whether a given saturated linear function is*

(i)globally convergent,
(ii)globally asymptotically stable,
(iii)mortal,
(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, the system $x_{t+1} = Ax_t$ is locally asymptotically stable if and only if the matrix A is stable, i.e., all its eigenvalues have magnitude less than one. Matrix stability can be decided by solving Lyapunov equations and is therefore decidable. For a stability checking algorithm see, e.g., [24].) 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 Theorem 6.2. If this conjecture is true, it is not hard to see that the equivalence of mortality and nilpotence also holds.

Theorem 2.1 has some "purely mathematical" consequences. For instance:

COROLLARY 2.1. *For infinitely many integers n, there exists a nilpotent saturated linear function $f : \mathbf{R}^n \to \mathbf{R}^n$ such that $f^{2^n} \not\equiv 0$.*

Proof. Assume that there exists an integer k such that for all $n \geq k$, a saturated linear function $f : \mathbf{R}^n \to \mathbf{R}^n$ is nilpotent if and only if $f^{2^n} \equiv 0$. The following algorithm solves the nilpotence problem for saturated linear functions, which is in contradiction with Theorem 2.1.

Let $f : \mathbf{R}^n \to \mathbf{R}^n$ be a saturated linear function. If $n \geq k$, declare f nilpotent if and only if $f^{2^n} \equiv 0$. If $n < k$, let $g : \mathbf{R}^k \to \mathbf{R}^k$ be the saturated linear function such that $g_i(x_1, \ldots, x_k) = f_i(x_1, \ldots, x_n)$ for $i \leq n$, and $g_i(x_1, \ldots, x_k) = 0$ for $n + 1 \leq i \leq k$. This transformation brings us back to the preceding case since the nilpotence of f is equivalent to the nilpotence of g. ■

Of course, in this corollary, 2^n can be replaced by any recursive function of n. In contrast, if $f : \mathbf{R}^n \to \mathbf{R}^n$ is a nilpotent linear function, then $f^n \equiv 0$. As shown in Theorem 2.2, this is not only a property of linear maps, but also of polynomial maps. For the proof of this theorem we need some basic notions from semi-algebraic geometry [3, 4]. In particular, we will use the fact that there is a well-defined notion of dimension for semi-algebraic sets. Those are the subsets of $\mathbf{R}^n$ defined by boolean combinations of polynomial inequalities.

LEMMA 2.2. *Let $f : \mathbf{R}^n \to \mathbf{R}^m$ be a polynomial map and $X = f(\mathbf{R}^n)$. For any polynomial map $g : \mathbf{R}^m \to \mathbf{R}$, if $\dim X = \dim X \cap \{g = 0\}$ then $g = 0$ on X.*

Proof. Let $Y = f^{-1}(X \cap \{g = 0\})$. Assume for a moment that $\dim Y < n$. Let $Z = \{x \in \mathbf{R}^n;\ g \circ f(x) \neq 0\}$ be the complement of Y. This set must be dense in $\mathbf{R}^n$ by the assumption $\dim Y < n$; $f(Z) = X \cap \{g \neq 0\}$ is therefore dense in X. For any nonempty semi-algebraic set S, the closure $\overline{S}$ of S satisfies $\dim \overline{S} \setminus S < \dim S$ ([4], Proposition 2.8.13). Here we use the convention $\dim \emptyset = -\infty$. Applying this observation to $S = X \cap \{g \neq 0\}$, we obtain the contradiction that

$$\dim X \cap \{g = 0\} < \dim X.$$

We conclude that in fact $\dim Y = n$, i.e., Y has nonempty interior. The polynomial function $g \circ f$ is null on an open set, and is therefore null on $\mathbf{R}^n$. ■

THEOREM 2.2. *Let $f : \mathbf{R}^n \to \mathbf{R}^n$ be a nilpotent polynomial map. For any $j \geq 0$, if $f^j \neq 0$ then $\dim f^{j+1}(\mathbf{R}^n) < \dim f^j(\mathbf{R}^n)$. As a consequence, $f^n = 0$.*

Proof. Let k be the smallest integer such that $f^k = 0$. The fact that $k \leq n$ follows immediately from the first part of the theorem. Let us therefore fix an integer $j < k$, and assume by contradiction that $\dim f^{j+1}(\mathbf{R}^n) = \dim f^j(\mathbf{R}^n)$. Since f^{k-j-1} is null on $f^{j+1}(\mathbf{R}^n)$, by Lemma 2.2 f^{k-j-1} is also null on $f^j(\mathbf{R}^n)$, i.e., $f^{k-1} = 0$. This is in contradiction with the minimality of k. ■

The statement of this theorem remains correct if we only assume that f is mortal. Indeed, for polynomial maps mortality is equivalent to nilpotence by, e.g., the Baire category theorem.

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. Combining these two observations with Theorem 2.1, we deduce that there exist saturated linear systems that are not globally asymptotically stable and have no nonzero periodic trajectories. We start with a lemma.

LEMMA 2.3. *Let X be a compact metric space with origin 0, and let $f : X \to X$ be a continuous function such that $f(0) = 0$. Then, the following two properties are equivalent:*

(i)f is globally asymptotically stable.

(ii)For every neighborhood U of 0, there exists an integer $k \geq 1$ such that $f^k(X) \subseteq U$.

Proof. If (ii) holds, it is clear that f is globally convergent. In order to show that f is also locally asymptotically stable, take any neighborhood U of 0 and let k be such that $f^k(X) \subseteq U$. By continuity, there exists another neighborhood V of 0 such that $\bigcup_{j=0}^{k-1} f^j(V) \subseteq U$. A trajectory of f originating in V never leaves U.

Assume now that f is globally asymptotically stable, and let U be a neighborhood of 0. By the definition of local asymptotic stability, there exists a neighborhood V of 0 such that a trajectory of f originating in V never leaves U. By global convergence, $X = \bigcup_{i \geq 0} f^{-i}(V)$. By compactness, this implies the existence of an integer $k \geq 1$ such that $X = \bigcup_{i=0}^{k} f^{-i}(V)$. This integer satisfies $f^k(X) \subseteq U$. ■

Our recursive enumerability result relies on our definition of saturated linear systems in terms of rational matrices A, which allows us to work within the Turing model of computation. The same argument applies to matrices with real entries, if we work in the real number model of computation [11, 10], and establishes that the set of globally asymptotically stable saturated linear systems is a countable union of semi-algebraic sets.

THEOREM 2.3. *The set of saturated linear systems that are globally asymptotically stable is recursively enumerable.*

Proof. Let $f(x) = \sigma(Ax)$. Consider the following algorithm:

1. Decide whether A is a stable matrix. If not, enter an infinite loop. Otherwise, go to Step 2.
2. Compute the sets $f^k([-1,1]^n)$ for $k = 1, 2, 3, \ldots$. Halt if an integer k such that $f^k([-1,1]^n) \subseteq]-1,1[^n$ is found.

We claim that this algorithm halts if and only if f is globally asymptotically stable.

Suppose that f is globally asymptotically stable. As pointed out earlier, A is a stable matrix. Consequently, the algorithm does not enter the infinite loop of Step 1. The algorithm must then halt at Step 2, according to Lemma 2.3.

Assume now that the algorithm halts. Since A must be a stable matrix, f is locally asymptotically stable. It remains to show that f is globally convergent. For any starting point $x_0 \in [-1,1]^n$, we have $f^j(x_0) \in]-1,1[^n$ for all $j \geq k$, i.e., the system never saturates after k steps, and $f^j(x_0) = A^{j-k}(f^k(x_0))$. Since A is stable, we conclude that $f^j(x_0) \to 0$ as $j \to \infty$. ■

THEOREM 2.4. *The set of saturated linear systems that have a nonzero periodic trajectory is recursively enumerable.*

Proof. Let $f(x) = \sigma(Ax)$. For any given positive integer k, it is straightforward to check whether there exists some nonzero x_0 such that $f^k(x_0) = x_0$, by solving a number (exponential in k) of linear systems of equations. Thus, the set of saturated linear systems that have a nonzero trajectory with period k is recursive. The set of saturated linear systems that have a nonzero periodic trajectory is the countable union of these recursive sets, hence the set is recursively enumerable. ■

COROLLARY 2.2. *There exist saturated linear systems that are not globally asymptotically stable and have no nonzero periodic trajectory.*

Proof. Assume by contradiction that the saturated linear systems that are not globally asymptotically stable always have a nonzero periodic trajectory. Then, by Theorem 2.4, these systems are recursively enumerable. But, by Theorem 2.3, the complement of this set is also recursively enumerable and so this would make global asymptotic stability a decidable property for saturated linear systems; a contradiction to Theorem 2.1. ■

3. TURING MACHINES

A *Turing machine* [17] is a deterministic model of computation. A given Turing machine M has a finite set Q of internal states and 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 $\Sigma = \{0, 1, \ldots, n-1\}$, $Q = \{0, 1, \ldots, m-1\}$, $n, m \in \mathbf{N}$, 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 *halting* 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 $\vdash: C \to C$, 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 3.1. *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 2.1. In order to apply that lemma, we endow the configuration space of a Turing machine with a topology which makes its successor function continuous, and its configuration space compact.

This is a fairly standard construction: let M be a Turing machine, C its configuration space, and $\vdash$ its successor function. Since $\vdash$ is not defined everywhere on C, we shall work on the space $X = C \cup \{0\}$, where 0 denotes a new, "final", configuration. We extend $\vdash$ to all of X by setting $c \vdash 0$ for every terminal configuration in C, and $0 \vdash 0$. Let d be a metric on X, defined by the following conditions:

(a) $d(0, c) = 1$ for every $c \in C$.

(b) for any two distinct configurations $c = (u, v, q)$ and $c' = (u', v', q')$, we have $d(c, c') = 1$ if $q \neq q'$; otherwise, $d(c, c') = 1/2^k$ where k is the largest integer such that u coincides with u' on the first k letters, and v coincides with v' on the first k letters.

It is clear that $\vdash$ is continuous with respect to the topology induced by d. One can also check that (X, d) is compact (for instance, one can use König's lemma on infinite trees to show that a convergent subsequence can be extracted from any sequence of points of X). Moreover, $\vdash$ is identically 0 in a neighborhood of 0 since this point is isolated in X. We therefore conclude from Lemma 2.1 that if M is mortal, then it must be nilpotent. ■

This theorem states that for mortal Turing machines, there is a uniform upper bound on the halting time of configurations. It follows from the next result that this upper bound is not computable. This result is due to Hooper and will play a central role in the sequel.

THEOREM 3.2 ([16]). *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

It is well known that Turing machines can be simulated by piecewise affine dynamical systems [18, 19, 21].

LEMMA 4.1. *Let M be a Turing machine and let $C = \Sigma^\omega \times \Sigma^\omega \times Q$ be its configuration space. There exists a piecewise affine function $g_M : \mathbf{R}^2 \to \mathbf{R}^2$ and an encoding function $\nu : C \to [0,1]^2$ such that the following diagram commutes:*

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

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

Proof. We define $\nu : C \to [0,1]^2$ as follows. Consider a configuration (p_1, p_2, q) of M, where $p_i = a_i^0 a_i^1 a_i^2 \ldots$, and each a_i^j is an element of Σ. We encode p_i in a real number x_i given by

$$x_i = \sum_{j=0}^{\infty} \frac{2a_i^j}{(2n)^{j+1}},$$

and we finally let

$$\nu(p_1, p_2, q) = \left(\frac{q}{m} + \frac{x_1}{m}, x_2\right).$$

For any $\alpha, \beta \in \Sigma$, and $q \in Q$, define the disjoint subsets $B_{\alpha,\beta,q}$ of $\mathbf{R}^2$ by

$$B_{\alpha,\beta,q} = \left[\frac{q}{m} + \frac{2\alpha}{2mn}, \frac{q}{m} + \frac{2\alpha+1}{2mn}\right] \times \left[\frac{2\beta}{2n}, \frac{2\beta+1}{2n}\right].$$

By the definition of ν, a configuration of the form $(\alpha p'_1, \beta p'_2, q)$, with $p'_1, p'_2 \in \Sigma^\omega$, $q \in Q$, has an image under ν which is a point in $B_{\alpha,\beta,q}$. Therefore, the same quintuple of the Turing machine M applies to all configurations that are mapped by ν to the same subset $B_{\alpha,\beta,q}$ (assuming $q \neq 0$; otherwise, no quintuple applies).

Such a quintuple has the effect of replacing the currently scanned symbol α by a new symbol α', of moving (or not) the tape to the right or to the left, and of changing the internal state q into a new internal state q'. Accordingly, we define the function g_M on the subset $B_{\alpha,\beta,q}$, $q \neq 0$, by $g_M(q/m+x_1/m, x_2) = (q'/m+x'_1/m, x'_2)$, where $x'_1 = ax_1 + b$, $x'_2 = cx_2 + d$, with:

- $a = 2n, b = -2\alpha, c = 1/(2n), d = (2\alpha')/(2n)$, if the tape is moved to the left.
- $a = 1/(2n), b = (2\beta)/(2n) + 2(\alpha' - \alpha)/(2n)^2, c = 2n, d = -2\beta$, if the tape is moved to the right.
- $a = 1, b = 2(\alpha' - \alpha)/(2n), c = 1, d = 0$, if the tape is not moved.

We then have $g_M(\nu(c)) = \nu(c')$ for all configurations $c, c' \in C$ with $c \vdash c'$. ■

A *closed box* is a Cartesian product of closed intervals in $\mathbf{R}$. A *σ^*-function* is a function obtained by composing finitely many σ-functions. For instance,

$$x \mapsto \sigma(\sigma(x) + \sigma(2\sigma(x+1)) \tag{1}$$

is a σ^*-function. In order to emphasize the structure of this function as a composition of three σ-functions (from $\mathbf{R}$ to $\mathbf{R}^2$, from $\mathbf{R}^2$ to $\mathbf{R}^2$ and from $\mathbf{R}^2$ to $\mathbf{R}$) we prefer to write

$$x \mapsto \sigma(\sigma(\sigma(x)) + \sigma(2\sigma(x+1))$$

instead of (1).

LEMMA 4.2. *Let P be a finite union of disjoint closed boxes of $\mathbf{R}^2$. Let $f : P \to [-1,1]^2$ be a function that is affine on each of the boxes in P. Then, there exists a σ^*-function $g : \mathbf{R}^2 \to \mathbf{R}^2$ that is equal to f on P.*

Proof. When ϵ is a positive real number and a is a real number, observe that the function $h^+_\epsilon(x,a) = \sigma(1 + 2(x-a)/\epsilon))$ has value 1 for $x \geq a$, and value -1 for $x \leq a - \epsilon$. Also, the function $h^-_\epsilon(x,a) = h^+_\epsilon(-x,-a)$ has value 1 for $x \leq a$, and value -1 for $x \geq a + \epsilon$. Write $P = \cup_{i=1}^n B_i$ with $B_i = [a_i^1, b_i^1] \times [a_i^2, b_i^2]$. Let d be the Euclidean distance between the closest two boxes. Consider the σ^*-function $\psi : \mathbf{R}^2 \to \mathbf{R}$ given by

$$\psi(x_1,x_2) = h^+_1(h^+_{d/2}(x_1,a_i^1) + h^-_{d/2}(x_1,b_i^1) + h^+_{d/2}(x_2,a_i^2) + h^-_{d/2}(x_2,b_i^2), 4)$$

and note that it takes the value 1 on B_i, and the value -1 on the $B_j, j \neq i$. Consider also the function $\chi_i : \mathbf{R}^2 \to \mathbf{R}^2$ defined by $\chi_i(x_1,x_2) = (\psi_i(x_1,x_2), \psi_i(x_1,x_2))$, and note that it is a σ^*-function. Define now the function g by

$$g(x_1,x_2) = \sigma\left(n - 1 + \sum_{i=1}^{n} \sigma\Big(\sigma(\sigma(f_i(x_1,x_2))) + \chi_i(x_1,x_2) - 1\Big)\right)$$

where f_i denotes the affine function that coincides with f on B_i. (In the above formula, the terms "1" and "n" stand for the vectors $(1,1)$ and (n,n) in $\mathbf{R}^2$.) ■

COROLLARY 4.1. *We can assume that the function g_M of Lemma 4.1 is a σ^*-function.*

Proof. The piecewise affine function g_M built in the proof of Lemma 4.1 is affine on a finite number of disjoint closed boxes in $\mathbf{R}^2$, namely, the sets $B_{\alpha,\beta,q}$ for $q \neq 0$. Furthermore, it can be checked that the image of each set $B_{\alpha,\beta,q}$ is contained in $[-1,1]^2$. Therefore, by Lemma 4.2, it can be extended to a σ^*-function defined on all of $\mathbf{R}^2$. ■

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

LEMMA 4.3. *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 $g_M : \mathbf{R}^2 \to \mathbf{R}^2$, a decoding function $\nu' : [0,1]^2 \to C$, and some subsets $\mathcal{N}^\infty \subset \mathcal{N}^1 \subset [0,1]^2$, $\mathcal{N}^1_{\neg term} \subset \mathcal{N}^1$ such that the following conditions hold:*

1.*$g_M(\mathcal{N}^\infty) \subseteq \mathcal{N}^\infty$ and $\nu'(\mathcal{N}^\infty) = C$.*

2.*$\mathcal{N}^1_{\neg term}$ (respectively $\mathcal{N}^1$) is the Cartesian product of two finite unions of closed intervals in $\mathbf{R}$. $\mathcal{N}^1_{\neg term}$ 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 $x \in \mathcal{N}^1_{\neg term}$.*

4.*The following diagram commutes:*

$$\begin{tikzcd} C \arrow[r, "\vdash"] & C \\ \mathcal{N}^1_{\neg term} \arrow[u, "\nu'"] \arrow[r, "g_M"] & {[0,1]^2} \arrow[u, "\nu'"'] \end{tikzcd}$$

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

Intuitively, ν' is an inverse of the encoding function ν of Lemma 4.1, in the sense that $\nu'(\nu(c)) = c$ holds for all configurations c. The set $\mathcal{N}^\infty$ 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, $\mathcal{N}^1_{\neg term}$ is the subset of $\mathcal{N}^1$ associated with the nonterminal internal states $q \neq 0$.

Proof. We use the notation and the functions ν and g_M introduced in the proof of Lemma 4.1. Using Corollary 4.1, we can assume that g_M is a σ^*-function.

We wish to define the function $\nu' : [0,1]^2 \to C$ in such a way that $\nu'(\nu(c)) = c$ holds for all $c \in C$. Toward this goal, we define $pop : [0,1] \times \mathbf{N} \to \Sigma$ by

$$pop(x,j) = \begin{cases} k & \text{if there exist } l \in \mathbf{N} \text{ and } k \in \Sigma \\ & \text{with } x - \frac{l}{(2n)^j} \in [\frac{(2k)}{(2n)^{j+1}}, \frac{(2k+1)}{(2n)^{j+1}}] \\ 0 & \text{otherwise} \end{cases}$$

Observe that if $x_i = \sum_{j=0}^{\infty} (2a_i^j)/(2n)^{j+1}$, then $pop(x_i, j) = a_i^j$, for all $j \in \mathbf{N}$. We then define $\nu' : [0,1]^2 \to C$ by $\nu'(y_1/m, y_2) = (p_1, p_2, int(y_1))$, where $p_i = a_i^0 a_i^1 a_i^2 \dots$ and the a_i^j are defined by $a_i^j = pop(fract(y_i), j)$. Here, int and $fract$ denote the integer part and fractional part, respectively. We then have $\nu'(\nu(c)) = c$ for all $c \in C$.

Define $\mathcal{N}^\infty$ as the union of the boxes $B_{\alpha,\beta,q}$, for $\alpha, \beta \in \Sigma, q \in Q$. Define $\mathcal{N}^1_{\neg term}$ as the union of the boxes $B_{\alpha,\beta,q}$ for $\alpha, \beta \in \Sigma$ and for $q \in Q$ not equal to the halting state 0 of M.

It can be verified that $\nu'(x) \vdash \nu'(g_M(x))$ for all $x \in \mathcal{N}^1_{\neg term}$.

Now, set $\mathcal{N}^\infty = \nu(C)$. Since $\nu'(\nu(c)) = c$, it follows that $\nu'(\mathcal{N}^\infty) = C$ holds. Furthermore, we have $g_M(\mathcal{N}^\infty) \subseteq \mathcal{N}^\infty$. Finally, the origin $(0,0)$ does not belong to $\mathcal{N}^1_{\neg term}$, and hence, is at a strictly positive distance from this set. ■

Using Lemma 4.3 and Theorem 3.2, we can now prove:

THEOREM 4.1. *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,

(iv)nilpotent,

are all undecidable.

The undecidability of the first three properties was first established in [5]. 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 Turing machine immortality problem (Theorem 3.2). 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 4.3 on $\mathcal{N}^1_{\neg term}$, and which is equal to 0 outside of $\mathcal{N}^1_{\neg term}$.

Since 0 is at a positive distance from $\mathcal{N}^1_{\neg term}$, 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 3.1, 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 $\mathcal{N}^1_{\neg term}$, we have $x_t \in \mathcal{N}^1_{\neg term}$ for $t = 0, \dots, k$. By the commutative diagram of Lemma 4.3, 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 4.3, 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 4.3, it suffices to prove that $x_t \in \mathcal{N}^1_{\neg term}$ for all t. Indeed, we prove by induction on t that $x_t \in \mathcal{N}^1_{\neg term} \cap \mathcal{N}^\infty$ and $\nu'(x_t) = c_t$ for all t. Using condition 3 of Lemma 4.3, the induction hypothesis is true for $t = 0$. Assuming the induction hypothesis for t, condition 1 of Lemma 4.3 shows that $x_{t+1} \in \mathcal{N}^\infty$. Now, the commutative diagram of Lemma 4.3 shows that $\nu'(x_{t+1}) = c_{t+1}$, and condition 3 of Lemma 4.3 shows that $x_{t+1} \in \mathcal{N}^1_{\neg term}$. 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. We start with a lemma.

Lemma 5.1. *Let P be a subset of $\mathbf{R}^2$ equal to the Cartesian product of two finite unions of closed intervals of $\mathbf{R}$. Then there exists a σ^*-function $Z_P : \mathbf{R}^2 \to \mathbf{R}$ that satisfies*

(i)$Z_P(x) = 0$ for all $x \in P$,
(ii)$Z_P(x) > 0$ for all $x \notin P$.

Proof. As in Lemma 4.2, when ϵ is a positive real number and a is a real number, denote by $h^+_\epsilon(x,a)$ the function defined by $h^+_\epsilon(x,a) = \sigma(1 + 2(x-a)/\epsilon))$, and by $h^-_\epsilon(x,a)$ the function defined by $h^-_\epsilon(x,a) = h^+_\epsilon(-x,-a)$.

Let I be an open interval of the form $I =]a,b[$. The function $\chi_I(x) = -h^+_{(b-a)/2}(x,b) - h^-_{(b-a)/2}(x,a)$ is zero for $x \notin I$, and strictly positive for $x \in I$. Let I be an open interval of the form $I =]a,\infty[$. The function $\chi_I(x) = 1 + h^+_1(x, a+1)$ is zero for $x \notin I$, and strictly positive for $x \in I$. Let I be an open interval of type $I =]-\infty, a[$. The function $\chi_I(x) = 1 + h^-_1(x, a-1)$ is zero for $x \notin I$, and strictly positive for $x \in I$.

When $J = \cup_{i=1}^n I_i$ is a finite union of closed intervals of $\mathbf{R}$, the complement J^c of J in $\mathbf{R}$ can be written as a finite union of open intervals: say $J^c = \cup_{i=1}^n I_i$. Define the function Z_J by $Z_J(x) = \sum_{i=1}^n \chi_{I_i}(x)$. This function is zero for $x \in J$, and is strictly positive for $x \notin J$. Finally, if $P = J_1 \times J_2$, let $Z_P(x_1,x_2) = \sigma(Z_{J_1}(x_1) + Z_{J_2}(x_2))$. ■

Theorem 5.1. *The following decision problem is undecidable:*

•*Instance: a σ^*-function $f : \mathbf{R}^3 \to \mathbf{R}^3$.*

•*Question: does there exist a trajectory $x_{t+1} = f(x_t)$ that belongs to $\{0\} \times \mathbf{R}^2$ for all t?*

Proof. We reduce the Turing machine immortality problem (Theorem 3.2) to this problem.

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

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

where the functions g_M and $Z_{\mathcal{N}^1_{\neg term}}$ are defined in Lemma 4.3 and Lemma 5.1, with $P = \mathcal{N}^1_{\neg term}$. 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 $x^1_t = 0$ for all t, if and only if Turing machine M has an immortal configuration.

Suppose that f has such a trajectory. Since $Z_{\mathcal{N}^1_{\neg term}}$, and hence $\sigma(\sigma(Z_{\mathcal{N}^1_{\neg term}}))$, is strictly positive outside of $\mathcal{N}^1_{\neg term}$, we must have $(x^2_t, x^3_t) \in \mathcal{N}^1_{\neg term}$ for all $t \geq 0$. By the commutative diagram of Lemma 4.3, the sequence $\nu'(x^2_t, x^3_t)$, $t \in \mathbf{N}$, is a sequence of successive configurations of M. By condition 3 of Lemma 4.3, none of these configurations is terminal, i.e. $c_0 = \nu'(x^2_0, x^3_0)$ 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 4.3, there exists a point $(x^2_0, x^3_0) \in \mathcal{N}^\infty$ with $\nu'(x^2_0, x^3_0) = c_0$. Consider the sequence defined by $(x^2_{t+1}, x^3_{t+1}) = g_M(x^2_t, x^3_t)$ for all t. Since $g_M(\mathcal{N}^\infty) \subseteq \mathcal{N}^\infty$, we have $(x^2_t, x^3_t) \in \mathcal{N}^\infty$ for all $t \geq 0$. Using the assumption that configuration c_t is nonterminal and condition 3 of Lemma 4.3, we deduce that $(x^2_t, x^3_t) \in \mathcal{N}^1_{\neg term}$ for all $t \geq 0$, which means precisely that the sequence $x_t = (0, x^2_t, x^3_t)$, $t \in \mathbf{N}$, is a trajectory of f. ∎

6. PROOF OF THE MAIN THEOREM

We now reach the last step in the proof, which consists of reducing the problem of Theorem 5.1 to the problems of Theorem 2.1.

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*.

We start with some preliminary technical results.

LEMMA 6.1 (Function *Abs*). *There exists a σ^*_0-function $Abs : \mathbf{R}^2 \to \mathbf{R}$, which is zero in some neighborhood of* 0, *and satisfies*

1. $Abs(1, u) \geq 0$ *for all* $u \in \mathbf{R}$*;*
2. $Abs(1, u) = 0$ *if and only if* $u = 0$*;*
3. $Abs(z, u) \geq 0$ *for all* $z \in [0, 1], u \in \mathbf{R}$.

Proof. Define $Abs(z, u) = \sigma(\sigma(u - z) - \sigma(u + z) + 2\sigma(z))$. ∎

LEMMA 6.2 (Function *Sel*). *There exists a σ^*_0-function $Sel : \mathbf{R}^2 \to \mathbf{R}$, which is zero in some neighborhood of* 0, *and satisfies*

1.$Sel(1, e) = e$
2.$Sel(0, e) = 0$

for all $e \in [-1, 1]$.

Proof. Define $h(x) = \sigma(2\sigma(x) - \sigma(2x))$ (see Figure 2) and $Sel(z, u) = \sigma(2h(3z/4 + u/4) - h(z))$. ■

Figure 2: Graph of the function $h(x) = \sigma(2\sigma(x) - \sigma(2x))$.

The construction that follows is the key to reducing the undecidable problem of Theorem 5.1 to the problems of Theorem 2.1.

LEMMA 6.3. *There exists a* σ_0^**-function* $Stab : \mathbf{R}^2 \to \mathbf{R}$*, null on some neighborhood of* 0*, with the following property. For all* $z_0 \in \mathbf{R}$ *and for all functions* $e : \mathbf{N} \to \mathbf{R}$*, the sequence* $z_{t+1} = Stab(z_t, e_t)$ *falls into one of the following three mutually exclusive cases:*

1.*The sequence* z_t*,* $t \geq 1$ *is constant, always equal to* 1*. This case happens only when* $\sigma(z_0) = 1$ *and when* $e_t = 0$ *for all* t*.*
2.*The sequence* z_t*,* $t \geq 1$ *is constant, always equal to* -1*. This case happens only when* $\sigma(z_0) = -1$ *and when* $e_t = 0$ *for all* t*.*
3.*The sequence* z_t *is eventually null: there exists* t_0 *with* $z_t = 0$ *for all* $t \geq t_0$*.*

Proof. Define $Stab$ for all z, e by $Stab(z, e) = h(\sigma(\sigma(z)) - Abs(z, e)/2)$, where h denotes the function $h(x) = \sigma(2\sigma(x) - \sigma(2x))$ shown in Figure 2.

Since $Abs(z, e)$ always belongs to the interval $[-1, 1]$, we have $\sigma(\sigma(z)) - Abs(z, e)/2 \geq -1/2$ for $z \geq 0$, and $\sigma(\sigma(z)) - Abs(z, e)/2 \leq 1/2$ for $z \leq 0$. We deduce that $0 \leq Stab(z, e)$ for $0 \leq z$, and $Stab(z, e) \leq 0$ for $z \leq 0$; hence the sign of z_t is constant. Assume without loss of generality that $0 \leq z_t$ for all t: if not, observing that $Stab(-z, -e) = -Stab(z, e)$, consider the sequences $-z_t$ and $-e_t$.

We now observe that $0 \leq Stab(z,e) \leq h(z)$ for all $e \in \mathbf{R}, z \in [0,1]$: indeed h is a nondecreasing function and we have $Abs(z,e) \geq 0$ for all $z \in [0,1], e \in \mathbf{R}$, from Lemma 6.1.

The function h has -1 and 1 as unstable fixed points, and 0 as a stable fixed point. Moreover, every sequence of the form $x_{t+1} = h(x_t)$ with $\sigma(x_0) \notin \{-1,1\}$ eventually reaches the stable fixed point 0. It follows that if there exists some t with $0 \leq z_t < 1$, then the sequence z_t eventually becomes zero. Now, if $z_t = 1$ for all t, and since the function h has value 1 only for $z \geq 1$, we must have $\sigma(\sigma(z_t)) - Abs(z_t, e_t)/2 \geq 1$ for all t, from which we deduce that $Abs(1, e_t) = 0$, and $e_t = 0$ for all t. ■

Lemma 6.4. *The problems of determining whether a given σ_0^*-function $\mathbf{R}^4 \to \mathbf{R}^4$ is*

(i)globally convergent,
(ii)globally asymptotically stable,
(iii)mortal,
(iv)nilpotent,

are all undecidable.

Proof. We first reduce the problem of Theorem 5.1 to the mortality problem for σ_0^*-functions.

Suppose that a σ^*-function $f : \mathbf{R}^3 \to \mathbf{R}^3$ is given. Thus, f is of the form $f = f_k \circ f_{k-1} \circ \cdots \circ f_1$ for some σ-functions $f_j = \sigma(A_j x + b_j)$. Define $f' : \mathbf{R}^4 \to \mathbf{R}^3$ by $f' = f'_k \circ f'_{k-1} \circ \ldots \circ f'_1$ where $f'_j(x,z) = \sigma(A_j x + b_j z)$ (x is a vector in $\mathbf{R}^3$ and z is scalar) so that $f(x) = f'(x,1)$ holds for all x.

Consider the σ_0^*-function $f'' : \mathbf{R}^4 \to \mathbf{R}^4$ defined for all $x^1, x^2, x^3, z \in \mathbf{R}$, by

$$f''(x^1,x^2,x^3,z) = \begin{pmatrix} Sel(\sigma^{(k)}(z), f'(x^1,x^2,x^3,z)) \\ Stab(\sigma^{(k-1)}(z), \sigma^{(k-1)}(x^1)) \end{pmatrix} \tag{2}$$

Here, the function Sel is applied componentwise, that is, $Sel(a, e_1, \ldots, e_4) = (Sel(a,e_1), \ldots, Sel(a,e_4))$.

We claim that f'' has an immortal trajectory $x''_{t+1} = f''(x''_t)$ (i.e. with $x''_t \neq 0$ for all t) if and only if f has a trajectory $x_{t+1} = f(x_t)$ with $x^1_t = 0$ for all t. Indeed, we argue as follows.

Suppose that f has a trajectory $x_{t+1} = f(x_t)$ with $x^1_t = 0$ for all t. Then $(x^1_t, x^2_t, x^3_t, 1)$ is a trajectory of f'', since

$$\begin{aligned} f''(x^1_t,x^2_t,x^3_t,1) &= (Sel(\sigma^{(k)}(1), f'(x^1_t,x^2_t,x^3_t,1)), Stab(\sigma^{(k-1)}(1), \sigma^{(k-1)}(x^1_t))) \\ &= (Sel(1, f'(x^1_t,x^2_t,x^3_t,1)), Stab(1, x^1_t)) \\ &= (f(x^1_t,x^2_t,x^3_t), 1) \\ &= (x^1_{t+1}, x^2_{t+1}, x^3_{t+1}, 1) \end{aligned}$$

This trajectory is immortal because its last component is constant and equal to 1.

Conversely, suppose that f'' has an immortal trajectory $x''_{t+1} = f''(x''_t)$. Denote $x''_t = (x''^1_t, x''^2_t, x''^3_t, x''^4_t)$. By Lemma 6.3, the sequence x''^4_t is either constant with value 1, or constant with value -1, or eventually null. The last case cannot happen because if there exists a t with $x_t^{''4} = 0$, then

$$\begin{aligned} x''_{t+1} &= (Sel(\sigma^{(k)}(0), f'(x_t^{''1}, x_t^{''2}, x_t^{''3}, 0)), Stab(\sigma^{(k-1)}(0), \sigma^{(k-1)}(x_t^{''1}))) \\ &= (Sel(0, f'(x_t^{''1}, x_t^{''2}, x_t^{''3}, 0)), Stab(0, \sigma^{(k-1)}(x_t^{''1}))) \\ &= (0, 0) \end{aligned}$$

Therefore, the sequence $x_t^{''4}$ is constant with value 1 or -1 and, by Lemma 6.3, we must have $x_t^{''1} = 0$ for all t. We can assume without loss of generality that $x_t^{''4} = 1$ for all t (otherwise, consider the sequence $-x''_t$ instead of x''_t which is also a trajectory of f'' since every σ_0^*-function, and hence f'', is odd). The sequence $x_t = (x^{''1}, x^{''2}, x^{''3})$ is a trajectory of f with $x_t^1 = 0$ for all t: indeed, $x_{t+1} = Sel(\sigma^{(k)}(x^{''4}), f'(x_t^1, x_t^2, x_t^3, x^{''4})) = Sel(1, f(x_t^1, x_t^2, x_t^3)) = f(x_t)$ and $x_t^1 = x^{''1}$ is zero for all $t \geq 0$.

We have just shown that the mortality problem for σ_0^*-functions is undecidable. Since Sel and $Stab$ are zero in a neighborhood of 0, the same is true of f''. It therefore follows from Lemma 2.1 that for f'', properties (i)-(iv) are equivalent. These four properties are therefore undecidable. ∎

We can now prove Theorem 2.1.

Proof. (of Theorem 2.1) We reduce the problems in Lemma 6.4 to the problems in Theorem 2.1.

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

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

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

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 $(x_t, f_1(x_t), \ldots, f_{k-1} \circ \ldots \circ f_1(x_t))$ is a subsequence of a trajectory of f'.

Conversely, let $x'_{t+1} = f'(x'_t)$ be a trajectory of f'. Write $x'_t = (y_t^1, \ldots, y_t^k)$ with each of the y^j in $\mathbf{R}^{d_{j-1}}$. For every $t_0 \in \{0, \ldots, k-1\}$ and $j \in \{1, \ldots, k\}$, the sequence $t \mapsto y_{t_0+kt}^j$ is a trajectory of f. This implies that the sequence y_t^j, $t \in \mathbf{N}$ 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 $(f')^{km} \equiv 0$. ∎

We now prove that all four properties remain undecidable for saturated linear systems of the form $x_{t+1} = \sigma(Ax_t)$ when A is a nilpotent matrix.

LEMMA 6.5. *A saturated linear system $x_{t+1} = \sigma(Ax_t)$ is nilpotent if and only if it is globally convergent and A is nilpotent.*

Proof. Let $f(x) = \sigma(Ax)$, and assume that $f^m \equiv 0$ for some $m \geq 1$. For all x in a suitably small neighborhood of 0, we have $f^m(x) = A^m x$. By linearity, this implies that $A^m x = 0$ for all $x \in \mathbf{R}^n$.

Conversely, assume that f is globally convergent and that A is nilpotent. By Lemma 2.1, f must be nilpotent. ■

THEOREM 6.1. *For a saturated linear system $x_{t+1} = \sigma(Ax_t)$ with A nilpotent, the properties of global convergence, global asymptotic stability, mortality, and nilpotence are all undecidable.*

Proof. When A is nilpotent, Lemma 2.1 shows that these properties are in fact equivalent. It is therefore sufficient to show that nilpotence is undecidable. Assume, to derive a contradiction, that we have a decision algorithm $\mathcal{A}$ for this problem. By the preceding lemma, one could then decide whether an arbitrary saturated linear system $x_{t+1} = \sigma(Ax_t)$ is nilpotent: if A is not nilpotent, output "system not nilpotent", otherwise call $\mathcal{A}$. This contradicts Theorem 2.1. (For a more direct proof, one can also check directly that the matrices constructed in the proof of that theorem are nilpotent). ■

THEOREM 6.2. *For a saturated linear system $x_{t+1} = \sigma(Ax_t)$ with A symmetric, mortality and nilpotence are both equivalent to the condition $A = 0$. Moreover, the properties of global convergence and global asymptotic stability are equivalent and decidable.*

Proof. For a nilpotent system, 0 is the only possible eigenvalue of A. If A is symmetric, this is equivalent to $A = 0$.

The decision algorithm for global asymptotic stability works as follows. We first decide whether A is a stable matrix. If it isn't, then the system cannot be globally asymptotically stable. If it is, then $\|Ax\| \leq \lambda \|x\|$, where $\lambda < 1$ is the spectral radius of A and $\|\cdot\|$ stands for the Euclidean norm. It follows that $\|\sigma(Ax)\| \leq \|Ax\| \leq \lambda \|x\|$, which implies that the system is globally asymptotically stable.

Next we show that global convergence implies global asymptotic stability. We shall use the existence of an "energy function" $E : [-1,1]^n \to \mathbf{R}$ satisfying the following property [14]: For any trajectory of the system, we have $E(x_{t+1}) < E(x_t)$ except if $x_t = x_{t+2}$, in which case $E(x_{t+1}) = E(x_t)$.

By compactness of $[-1,1]^n$, E achieves its minimum at some point a. This implies by the above property that a is a periodic point. For the system to be globally convergent we must therefore have $a = 0$ (and it is the only point where E achieves its minimum). To complete the proof we need to use the specific form of E:

$$E(x_t) = -x_t^T A x_{t+1} + \sum_{i=1}^{n} [(x_t)_i^2 + (x_{t+1})_i^2]/2.$$

Let λ be any eigenvalue of A and x_0 an eigenvector associated to λ. If x_0 is of sufficiently small norm we have $x_1 = \lambda x_0$ so that $E(x_0) = \frac{1-\lambda^2}{2}|x_0|^2$. Since E achieves its minimum only in 0 and $E(0) = 0$ we conclude that $|\lambda| < 1$. As we have seen previously, this implies that the system is globally asymptotically stable.

The proof that mortality implies $A = 0$ is now easy. As we have just shown, for a mortal system any eigenvalue λ of A must satisfy $|\lambda| < 1$. If $\lambda \neq 0$, a trajectory starting at an eigenvector $x_0 \neq 0$ is therefore not mortal. We conclude that 0 is the only eigenvalue of A, whence $A = 0$. ■

7. CONTINUOUS PIECEWISE AFFINE SYSTEMS

We proved in Theorem 4.1 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 7.1. *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 6.4 is of dimension 4. However, if in the right hand side of equation (2), we replace x^1 by $\sigma(\sigma(Z_{\mathcal{N}^1_{\neg term}}(x_2, x_3)))$ and x^2, x^3 by $g_M(x^2, x^3)$, then, from Theorem 5.1 and Lemma 6.4, we obtain a function in dimension 3, which gives a direct reduction from the problem of Theorem 3.2 to the problems of Lemma 6.4. ■

The following proposition is proved in [5].

THEOREM 7.2. *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 $f : \mathbf{R}^n \to \mathbf{R}^n$	$n = 1$	$n = 2$	$n = 3$
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 have seen in Theorem 6.2 that these equivalences hold for systems with symmetric matrices.

4. For a polynomial map $f : \mathbf{R}^n \to \mathbf{R}^n$ mortality is equivalent to nilpotence; these properties are equivalent to the condition $f^n \equiv 0$ and hence decidable (here f^n denotes the n-th iterate of f, as in the rest of the paper). 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 $f : \mathbf{R}^n \to \mathbf{R}^n$ such that $f^k \not\equiv 0$? 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 2.3 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, by the European Commission under the TMR Alapedes network, and by 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. R. Benedetti and J.-J. Risler. *Real algebraic and semi-algebraic sets.* Hermann, Paris, 1990.
4. J. Bochnak, M. Coste and M.-F. Roy. *Real algebraic geometry.* Springer-Verlag, 1998.

5. V. D. Blondel, O. Bournez, P. Koiran, C. Papadimitriou, and J. N. Tsitsiklis. Deciding stability and mortality of piecewise affine dynamical systems. Technical Report 1999-05, Laboratoire de L'Informatique du Parallélisme, Ecole Normale Supérieure de Lyon, 1999. To appear in *Theoretical Computer Science*.

6. V. D. Blondel and J. N. Tsitsiklis. Complexity of stability and controllability of elementary hybrid systems. *Automatica*, 35(3), 479–489, 1999.

7. 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*, 46–58. Springer-Verlag, 1999.

8. 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*, 45–52. Springer-Verlag, London, 1999.

9. V. D. Blondel and J. N. Tsitsiklis. A survey of computational complexity results in systems and control. *Automatica*, 36(9), 1249–1274, 2000.

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. O. Bournez. Complexité algorithmique des systèmes dynamiques continus et hybrides. PhD Thesis, Ecole Normale Supérieure de Lyon, 1999.

13. O. Bournez, and M. Cosnard. On the computational power of dynamical systems and hybrid systems. *Theoretical Computer Science*, 168(2), 417–459, 1996.

14. F. Fogelman Soulié, C. Mejia, Eric Goles, and S. Martinez. Energy functions in neural networks with continuous local functions. *Complex Systems*, 3, 269–293, 1989.

15. 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.

16. P. K. Hooper. The undecidability of the Turing machine immortality problem. *The Journal of Symbolic Logic*, 31(2), 219–234, June 1966.

17. J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory Languages and Computation*. Addison-Wesley, October 1979.

18. P. Koiran, M. Cosnard, and M. Garzon. Computability with low-dimensional dynamical systems. *Theoretical Computer Science*, 132(1-2), 113–128, September 1994.

19. C. Moore. Generalized shifts: unpredictability and undecidability in dynamical systems. *Nonlinearity*, 4, 199–230, 1991.

20. H. T. Siegelmann and E. D. Sontag. Analog computation via neural networks. *Theoretical Computer Science*, 131(2), 331–360, September 1994.

21. 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.

22. E. D. Sontag. From linear to nonlinear: Some complexity comparisons. In *IEEE Conference on Decision and Control*, 2916–2920, New Orleans, December 1995.

23. 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.

24. E. D. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Springer-Verlag, New York, second edition, August 1998.