

The Dynamics of Influence Systems *

BERNARD CHAZELLE †

Abstract

Influence systems form a large class of multiagent systems designed to model how influence, broadly defined, spreads across a dynamic network. We build a general analytical framework which we then use to prove that, while sometimes chaotic, influence dynamics of the diffusive kind is almost always asymptotically periodic. Besides resolving the dynamics of a popular family of multiagent systems, the other contribution of this work is to introduce a new type of renormalization-based bifurcation analysis for multiagent systems.

1 Introduction

The contribution of this paper is twofold: (i) to formulate an “algorithmic calculus” for continuous, discrete-time multiagent systems; and (ii) to resolve the behavior of a popular type of social dynamics that had long resisted analysis. In the process, we also introduce a new approach to time-varying Markov chains. *Diffusive influence systems* are piecewise-linear dynamical systems $\mathbf{x} \mapsto P(\mathbf{x})\mathbf{x}$, which are specified by a piecewise-constant function P mapping any $\mathbf{x} \in \mathbb{R}^n$ to an n -by- n stochastic matrix $P(\mathbf{x})$. We prove that, while sometimes chaotic, such systems are almost surely attracted to a fixed point or a limit cycle.

As in statistical mechanics, the difficulty of analyzing influence systems comes from the tension between two opposing forces: one, caused by the map’s discontinuities, is “entropic” and leads to chaos; the other one, related to the Lyapunov exponents, is “energetic” and pulls the system toward an attracting manifold within which the dynamics is periodic. The challenge is to show that, outside a vanishingly small critical region in parameter space, entropy always loses. Because the interaction topology changes all the time (endogenously), the proof relies heavily on an algorithmic framework to monitor the flow of information across the system. As a result, this work is, at its core, an algorithmic study in dynamic networks. Influence systems include finite Markov chains as a

*This work was supported in part by NSF grants CCF-0832797, CCF-0963825, and CCF-1016250. Prelim. version available in Proc. 53rd FOCS 2012.

†Department of Computer Science, Princeton University, chazelle@cs.princeton.edu

special case but the differences are deep and far-reaching: whereas Markov chains have predictable dynamics, influence systems can be chaotic even for small n ; whereas the convergence of a Markov chain can be checked in polynomial time, the convergence of an influence system is undecidable. Our main result is that this bewildering complexity is in fact confined to a vanishing region of parameter space. Typically, influence systems are asymptotically periodic.

Influence and social dynamics. There is a context to this work and this is where we begin. An overarching ambition of social dynamics is to understand and predict the collective behavior of agents influencing one another across an endogenously changing network [13]. *HK* systems have emerged in the last decade as a prototypical platform for such investigations [5, 6, 26, 27, 31, 33, 34, 36, 37, 39]. To unify its varied strands (eg, bounded-confidence, bounded-influence, truth-seeking, Friedkin-Johnsen type, deliberative exchange) into a single framework and supply closed-loop analogs to standard consensus models [4, 35, 40], we introduce *influence systems*. These are discrete-time dynamical systems $\mathbf{x} \mapsto f(\mathbf{x})$ in $(\mathbb{R}^d)^n$: each “coordinate” x_i of the state $\mathbf{x} = (x_1, \dots, x_n)$ is a d -tuple encoding the location of *agent* i as a point in \mathbb{R}^d ; with any state \mathbf{x} is associated a directed graph $\mathcal{G}(\mathbf{x})$ with the n agents as nodes. Each coordinate function f_i of the map $f = (f_1, \dots, f_n)$ takes as input the neighbors of agent i in $\mathcal{G}(\mathbf{x})$ and outputs the new location $f_i(\mathbf{x})$ of agent i in d -space. One should think of agent i as a “computer” and x_i as its “memory.” All influence systems in this work will be assumed to be *diffusive*, meaning that at each step an agent may move only within the convex hull of its neighbors.¹ Note that the system $\mathbf{x} \mapsto P(\mathbf{x})\mathbf{x}$ in the opening paragraph corresponds to the one-dimensional case.

Influence systems arise in processes as diverse as chemotaxis, synchronization, opinion dynamics, flocking, swarming, and rational social learning.² Typically, a natural algorithm directs n autonomous agents to obey two sets of rules: (i) one of them determines, on the basis of the system’s current state \mathbf{x} , which agent communicates with which one; (ii) the other one specifies how an agent updates its state by processing the information it receives from its neighbors. Diffusive influence systems are central to social dynamics insofar as they extend the fundamental concept of *diffusion* to autonomous agents operating within dynamic, heterogeneous environments.³ This stands in sharp contrast with the classic brand of diffusion found in physics and chemistry,

¹ This is a standard assumption meant to ensure that consensus is a fixed point.

² The states of an influence system can be: opinions [6, 23, 26, 31, 33, 35, 36], Bayesian beliefs [1], neuronal spiking sequences [15], animal herd locations [20], consensus values [12, 40], swarming trajectories [25], cell populations [47], schooling fish velocities [43, 45], sensor networks data [10], synchronization phases [24, 46, 48], heart pacemaker cell signals [52, 56], cricket chirpings [55], firefly flashings [38], yeast cell suspensions [47], microwave oscillator frequencies [52], or flocking headings [3, 16, 21, 28, 29, 54].

³ For a fanciful but illustrative example, imagine n insects on the ground ($d = 2$), each one moving toward the mass center of its neighbors. Each one gets to “choose” who is its neighbor: this cricket picks the five ants closest to it within its cone of vision; that spider goes for the ladybugs within two feet; these ants select the 10 furthest termites; etc. Once the insects have determined their neighbors, they move to their mass center (or a weighted version of it). This is repeated forever.

which assumes passive particles subject to identical laws. Naturally, influence systems are “downward-compatible” and can model standard (discrete) diffusion as well. They also allow exogeneities (eg, diffusion-reaction) via the addition of special-purpose agents. Autonomy and heterogeneity are the defining features of influence systems: they grant agents the freedom to have their own, distinct decision procedures to choose their neighbors as they please and act on the information collected from them. This explains their ubiquity among natural algorithms.

The model. In a diffusive influence system, $f(\mathbf{x}) = (P(\mathbf{x}) \otimes \text{Id}) \mathbf{x}$, where $P(\mathbf{x})$ is a stochastic matrix whose positive entries correspond to the edges of $\mathcal{G}(\mathbf{x})$ and are rationals larger than some arbitrarily small $\rho > 0$. We form the Kronecker product with the d -by- d identity matrix to perform the averaging along each coordinate axis. We grant the agents a measure of self-confidence by adding a self-loop to each node of $\mathcal{G}(\mathbf{x})$. Agent i computes the i -th row of $P(\mathbf{x})$ by means of its own algebraic decision tree; that is, on the basis of the signs of a finite number of dn -variate polynomials evaluated at the coordinates of \mathbf{x} . This high level of generality allows $\mathcal{G}(\mathbf{x})$ to be specified by any first-order sentence over the reals:⁴ In a recent bird flocking model [3], for instance, the communication graph joins every agent to its 7 nearest neighbors. We show below how to reduce the dimension to $d = 1$ and linearize the system so that $P(\mathbf{x}) = P_c$, for any $\mathbf{x} \in c$, where c is any *atom* (open n -cell) of an arrangement of hyperplanes in \mathbb{R}^n , called the *switching partition* (*SP*). An influence system is called *bidirectional* if $\mathcal{G}_{ij} \equiv \mathcal{G}_{ji}$ (with $\mathcal{G} = (\mathcal{G}_{ij})$), which implies that $\mathcal{G}(\mathbf{x})$ is undirected. Such a system is further called *metrical* if \mathcal{G}_{ij} is solely a function of $|x_i - x_j|$. Homogeneous *HK* systems [26, 27, 31] constitute the canonical example of a metrical system. We assume that all the relevant parameters (matrix entries, number and coefficients of hyperplanes, ρ , etc) can be encoded as rationals over $O(\log n)$ bits: this assumption can be freely relaxed—in fact, the bit lengths can be arbitrarily large as a function of n —and is only made to simplify the notation.

Past work and present contribution. Beginning with their introduction by Sontag [51], piecewise-linear systems have become the subject of an abundant literature, which we do not attempt to review here. Restricting ourselves to influence systems, we note that the bidirectional kind are known to be attracted to a fixed point while expending a *total s-energy* at most exponential in the number of agents and polynomial in the reversible case [17, 23, 28, 35, 40]. Convergence times are known only in the simplest cases [10, 17, 37]. In the nonbidirectional case, most convergence results are conditional [12, 14, 15, 29, 39–41, 44, 53].⁵ The standard assumption is that some form of joint connectivity property should hold in perpetuity. To check such a property is in general

⁴ This is the language of geometry and algebra with statements specified by any number of quantifiers and polynomial (in)equalities. It was shown to be decidable by Tarski and amenable to quantifier elimination and algebraic cell decomposition by Collins [18].

⁵ As they should be, since convergence is not assured. An exception is *truth-seeking HK systems*, which have been shown to converge unconditionally [17].

undecidable (see why below), so these convergence results are somewhat of a heuristic nature. A significant recent advance was Bruin and Deane’s unconditional resolution of planar piecewise contractions, which are special kinds of influence systems with a single mobile agent [9]. Our main result can be interpreted as a grand generalization of theirs.

THEOREM 1.1. *Given any initial state, the orbit of an influence system is attracted exponentially fast to a limit cycle almost surely under an arbitrarily small perturbation. The period and preperiod are bounded by a polynomial in the reciprocal of the failure probability. Without perturbation, the system can be Turing-complete. In the bidirectional case, the system is attracted to a fixed point in time $n^{O(n)} \log \frac{1}{\varepsilon}$ almost surely, where n is the number of agents and ε is the distance to the fixed point.*

The theorem bounds the convergence time of bidirectional systems by a single exponential and establishes the asymptotic periodicity of generic influence systems. These results are essentially optimal. We also estimate the attraction rate of general systems but the bounds we obtain are probably too conservative to be useful. Perturbing the system means replacing each hyperplane $\mathbf{a}^T \mathbf{x} = a_0$ of the *SP* by $\mathbf{a}^T \mathbf{x} = a_0 + \delta$, for some arbitrarily small random δ . Note that neither the initial state nor the transition matrices are perturbed.⁶ We enforce an *agreement rule*, which sets \mathcal{G}_{ij} to be constant over the microscopic slab $|x_i - x_j| \leq \varepsilon_0$, for an arbitrarily small $\varepsilon_0 > 0$.⁷ Intuitively, the agreement rule stipulates that minute fluctuations of opinion between two agents otherwise in full agreement should have no macroscopic effect on the system.⁸ We emphasize that *both* the perturbation and the agreement rule are necessary: without them, the attraction claims of Theorem 1.1 are provably false.⁹ We show that finely tuned influence systems are indeed Turing-complete.

Our work resolves the long-term behavior of a fundamental natural process which includes the extended family of *HK* systems as a special case. The high generality of our results precludes statements about particular restrictions which might be easier. A good candidate for further investigation is the heterogeneous *bounded-confidence* model, where each \mathcal{G}_{ij} is defined by a single interval, and which is conjectured to converge [39]. (We show below that this is false if the averaging is not perfectly uniform.) Such systems were not even known to be periodic, a feature that our result implies automatically. Generally, our work exposes a surprising gap in the expressivity of directed and undirected dynamic networks: while the latter always lead to stable agreement (of a consensual, polarized, or fragmented nature), directed graphs offer a much richer complexity landscape.

The second contribution of this work is the introduction of a new brand of bifurcation analysis based on algorithmic renormalization. In a nutshell, we use a graph algorithm

⁶ This is not a noise model [8]: the perturbation happens only once at the beginning.

⁷ Agent i is free to set the function \mathcal{G}_{ij} to either 0 or 1. For notational convenience, we set ε_0 to be $n^{-O(1)}$, but smaller values would work just the same.

⁸ Interestingly, this is precisely meant to prevent the “narcissism of small differences,” identified by Freud and others as a common source of social conflicts.

⁹ In the nonbidirectional case, agents are made to enforce a timeout mechanism to prevent edges from reappearing after an indefinite absence of unbounded length. While probably unnecessary, this minor technical feature seems to simplify the proof.

to decompose a dynamical system into a hierarchy of recursively defined subsystems. We then develop a tensor calculus to “compile” the graph algorithm into a bifurcation analysis. The tension between energy and entropy is then reduced to a question in matrix rigidity.

In the context of social dynamics, Theorem 1.1 might be somewhat disconcerting. Influence systems model how people change opinions over time as a result of human interaction and knowledge acquisition. Our results show that, unless people keep varying the modalities of their interactions, as mediated by trust levels, self-confidence, etc, *they will be caught forever recycling the same opinions in the same order*. The saving grace is that the period can be exponentially long, so the social agents might not even realize they have become little more than a clock...

2 The Complexity of Influence Systems

Piecewise-linear systems are known to be Turing-complete [2, 7, 30, 50]. A typical simulation relies on the existence of Lyapunov exponents of both signs, negative ones to move the head in one direction and positive ones to move it the other way. Influence systems have no positive exponents and yet are Turing-complete, as we show below. In dynamics, chaos is typically associated with positive topological entropy, which entails expansion, hence positive Lyapunov exponents. But piecewise linearity blurs this picture and produces surprises. For example, isometries (with only null Lyapunov exponents) are not chaotic [11] but, paradoxically, contractions (with only negative exponents) can be [32]. Influence systems, which, with only null and negative Lyapunov exponents, sit in the middle, can be chaotic. The spectral lens seems to break down completely in the face of piecewise linearity!

Exponential periods. It is an easy exercise to use higher bit lengths to increase the period of an oscillating influence system by any amount. More interesting is the observation that the period can be raised to exponential with only logarithmic bit length. We simulate a counter modulo 2 by building a system with $n = 3$: the first two agents are fixed at 0 and 3 while the third oscillates between positions 1 and 2; this is trivially achieved with a two-test linear decision tree. Add another mobile agent oscillating between 1 and 2 like the previous one, but which moves only when the first oscillating agent is at position 1. (Adding a single test makes this possible.) Iterating in this fashion produces an n -agent influence system with $O(n)$ tests whose period is exactly 2^{n-2} .

For a system where action and control are more closely mixed, consider implementing $\mathbb{Z}/2\mathbb{Z}$ as a 3-agent influence system by fixing the first two agents at positions 0 and 3, respectively, and then letting the third one oscillate between 1 and 2. By adding $O(q)$ discontinuities, we extend this scheme to keep an agent cycling through $1, \dots, q$, and then back to 1, thus implementing $\mathbb{Z}/q\mathbb{Z}$. Repeating this construction for the first k primes $p_1 < \dots < p_k$ allows us to implement the system based on the direct sum $\mathbb{Z}/p_1\mathbb{Z} \oplus \dots \oplus \mathbb{Z}/p_k\mathbb{Z}$, which has period of $\prod_{j \leq k} p_j$ for a total of $N = O(p_1 + \dots + p_k)$

agents and discontinuities. By the prime number theorem [42], this gives us a period of length $2^{\Omega(\sqrt{N \log N})}$. While the period of the system as a whole is huge, each agent cycles through a short periodic orbit: this is easily remedied by adding another agent attracted to the mass center of the k cycling agents. By the Chinese Remainder Theorem, that last agent has an exponential period and acts as a sluggish clock.

Chaos and perturbation. Perturbation is needed for several reasons, including uniform bounds on the time to stationarity. We focus here on the agreement rule and show why it is necessary by designing a chaotic system that is resistant to perturbation. We use a total of four agents. The first two agents stay on opposite sides of 0.5, with the one further from 0.5 moving toward it:

$$(x_1, x_2) \mapsto \frac{1}{2} \begin{cases} (2x_1, x_1 + x_2) & \text{if } x_1 + x_2 \geq 1 \\ (x_1 + x_2, 2x_2) & \text{else.} \end{cases}$$

The two agents converge toward 0.5 but the order in which they proceed (ie, their symbolic dynamics) is chaotic. To turn this into actual chaos, we introduce a third agent, which oscillates between a fourth agent fixed at $x_4 = 0$ and x_1 (which is roughly 0.5), depending on the order in which the first two agents move: $x_3 \mapsto \frac{1}{3}(x_3 + 2x_1)$ if $x_1 + x_2 \geq 1$ and $x_3 \mapsto \frac{1}{3}(x_3 + 2x_4)$ else. Assume that $x_1(0) < \frac{1}{2} < x_2(0)$ and consider the trajectory of a line $L: X_2 - \frac{1}{2} = u(X_1 - \frac{1}{2})$, for $u < 0$. If the point $(x_1(t), x_2(t))$ is on the line, then $x_1(t) + x_2(t) \geq 1$ implies that $u \leq -1$ and L is mapped to $X_2 - \frac{1}{2} = \frac{1}{2}(u + 1)(X_1 - \frac{1}{2})$; if $x_1(t) + x_2(t) < 1$, then $u > -1$ and L becomes

$$X_2 - \frac{1}{2} = \frac{2u}{u + 1} (X_1 - \frac{1}{2}).$$

The parameter u obeys the dynamics: $u \mapsto \frac{1}{2}(u + 1)$ if $u \leq -1$ and $u \mapsto 2u/(u + 1)$ if $-1 < u \leq 0$. Writing $u = (v + 1)/(v - 1)$ gives $v \mapsto 2v + 1$ if $v < 0$ and $v \mapsto 2v - 1$ else. (Geometrically, v is the tangent of the angle between L and the line $X + Y = 0$.) The system v escapes for $|v(0)| > 1$ and otherwise conjugates with the baker's map [22] via the variable change: $v = 2w - 1$. Agent 3 is either at most 1/6 or at least 1/3 depending on which of agent 1 or 2 moves. This implies that the system has positive topological entropy: to know where agent 3 is at time t requires on the order of t bits of accuracy in the initial state. The cause of chaos is the first two agents' convergence toward the SP discontinuity. It is immediate that no perturbation can prevent this, so the agreement rule is indeed needed.

Turing completeness. Absent perturbation and the agreement rule, an influence system can simulate a piecewise-linear system and hence a Turing machine. Here is how. Given a nonzero n -by- n real-valued matrix A , let A^+ (resp. A^-) be the matrix obtained by zeroing out the negative entries of A (resp. $-A$), so that $A = A^+ - A^-$. Define the matrices

$$B = \rho \begin{pmatrix} A^+ & A^- \\ A^- & A^+ \end{pmatrix} \quad \text{and} \quad C = \begin{pmatrix} B & (\text{Id} - B)\mathbf{1} & \mathbf{0} \\ \mathbf{0} & 1 & 0 \\ \mathbf{0} & 1 - \rho & \rho \end{pmatrix},$$

where ρ is the reciprocal of the maximum row-sum in the matrix derived from A by taking absolute values. It is immediate that C is stochastic and conjugates with the dynamics of A . Indeed, given $\mathbf{x} \in \mathbb{R}^n$, if $\bar{\mathbf{x}}$ denotes the $(2n + 2)$ -dimensional column vector $(\mathbf{x}, -\mathbf{x}, 0, 1)$, then $C\bar{\mathbf{x}} = \rho \overline{A\mathbf{x}}$; hence the commutative diagram:

$$\begin{array}{ccc} \mathbf{x} & \longrightarrow & A\mathbf{x} \\ \downarrow & & \downarrow \\ \bar{\mathbf{x}} & \longrightarrow & \rho^{-1}C\bar{\mathbf{x}}. \end{array}$$

Imagine now a piecewise-linear system consisting of a number of matrices $\{A\}$ and an *SP*. We add n negated clones to the existing set of n agents, plus a *stochasticity* agent permanently positioned at $x_{-1} = 0$ and a *projectivity* agent initially at x_0 . This allows us to form the vector $\bar{\mathbf{x}} = (\mathbf{x}, -\mathbf{x}, x_{-1}, x_0)$. The system scales down, so we must projectify the *SP* by rewriting with homogeneous coordinates any $\mathbf{a}^T \mathbf{x} = a_0$ as $\mathbf{a}^T \mathbf{x} = a_0 x_0$. We can use the same value of ρ throughout by picking the smallest one among all the matrices A used in the piecewise-linear system.

Koiran et al [30] have shown how to simulate a Turing machine with a 3-agent piecewise-linear system, so we set $n = 3$. We need an *output* agent to indicate whether the system is in an accepting state: this is done by pointing to one of two fixed agents. We can enlist one of the three original agents for that purpose, which brings the agent count up to $2n + 3 = 9$. Predicting basic state properties of an influence system is therefore undecidable. With a few more agents, we can easily encode as an undecidable question whether the communication graphs (or their union over bounded time windows) satisfy certain connectivity property infinitely often.

Linearization. We show how to linearize an influence system by tensor powering (for any d). Let \mathbf{d} be the maximum total degree of the polynomial tests used in the algebraic decision trees (recall that each agent comes equipped with its own). We can always assume the existence of an agent confined to position 1 with no in/out-link: we use it to homogenize the test polynomials, so that every monomial has degree exactly \mathbf{d} . Given $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$, we define the monomial $y_{k_1, \dots, k_d} = \prod_{i=1}^d x_{k_i}$ ($1 \leq k_1, \dots, k_d \leq n$) and, listing them in lexicographic order, form $\mathbf{y} = (y_{k_1, \dots, k_d}) \in \mathbb{R}^N$, where $N = n^d$; note that \mathbf{y} lies on a (real) algebraic variety \mathcal{V} smoothly parametrized injectively by \mathbf{x} . The map $\mathbf{x} \mapsto f(\mathbf{x})$ induces the lifted map $\mathbf{y} \mapsto g(\mathbf{y})$, where $g(\mathbf{y}) = P(\mathbf{x})^{\otimes d} \mathbf{y}$ and

$$P(\mathbf{x})^{\otimes d} = \overbrace{P(\mathbf{x}) \otimes \dots \otimes P(\mathbf{x})}^{\mathbf{d}}.$$

Being the Kronecker product of stochastic matrices, $P(\mathbf{x})^{\otimes d}$ is stochastic: its diagonal is positive and its nonzero entries all exceed ρ^d . Its associated graph, whose edges map out its nonzero entries, is the tensor graph product $\mathcal{G}(\mathbf{x})^{\otimes d}$. We use the term *ground agents* to refer to the n agents positioned at \mathbf{x} . Including all the test polynomials from all the ground agents' decision trees gives us as many hyperplanes in \mathbb{R}^N and the sign conditions of a cell c specify a unique stochastic matrix Q_c . This matrix is always a

tensor power $P^{\otimes d}$ but it is guaranteed to be of the form $P(\mathbf{x})^{\otimes d}$ only if c contains a point \mathbf{y} of V parametrized by \mathbf{x} .

Whereas a random shift produces affine forms $a_1 y_1 + \dots + a_N y_N + \delta$, the agreement rule acts in a more subtle way. While the whole point of the lifting is to forget about the variety \mathcal{V} , the tensor structure of the matrices Q_c brings benefits we can exploit. Given $K \subseteq \{1, \dots, n\}$, the *cluster* C_K refers to the subset of $|K|^d$ agents with labels in K^d . If all the agents of a cluster fit within a tiny interval then so do their ground agents; to see why, just expand $(x_i - x_j)^d$. By the agreement rule, therefore, the induced subgraph of the cluster cannot change until it is pulled apart by outside agents. Assume now that $d > 1$. We write

$$\mathbf{x} = (x_{1,1}, \dots, x_{1,d}, \dots, x_{n,1}, \dots, x_{n,d}),$$

with the homogeneizing agent 1 permanently positioned at $(x_{1,1}, \dots, x_{1,d}) = \mathbf{1}_d$. Next, we define $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_N)$, where $N = (dn)^d$ and $\mathbf{y}_l = \prod_{i=1}^d x_{k_i, j_i}$ with l denoting the lexicographic rank of the string $(k_1, j_1, \dots, k_d, j_d)$ for $k_i \in \{1, \dots, n\}$ and $j_i \in \{1, \dots, d\}$. The matrix Q_c associated with cell c is of the form $(P \otimes \mathbf{I}_d)^{\otimes d}$; furthermore, $P = P(\mathbf{x})$ whenever \mathbf{y} satisfies the N conditions $\mathbf{y}_l = \prod_{i=1}^d x_{k_i, j_i}$ for some $\mathbf{x} \in \mathbb{R}^{dn}$. The cluster C_K consists now of $(d|K|)^d$ agents.

Nonconvergent *HK* systems. Heterogeneous *HK* systems [26,27] are influence systems where each agent i is associated with a confidence value r_i and the communication graph links i to any agent j such that $|x_i - x_j| \leq r_i$. We design a periodic 5-agent system with period 2. We start with a 2-agent system with $r_1 = r_2 = 2$. Instead of uniform averaging, we decrease the self-confidence weight so that, if the two agents are linked, then $x_1 \mapsto \frac{1}{3}(x_1 + 2x_2)$ and $x_2 \mapsto \frac{1}{3}(2x_1 + x_2)$. Any self-confidence weight less than 0.5 would work, too; of course, a weight of zero makes the problem trivial and uninteresting. If the agents are initially positioned at -1 and 1 , they oscillate around 0 with $x_i = (-1)^{i+t} 3^{-t}$ at time t . Now place a copy of this system with its center at $X = 2$ and a mirror-image copy at $X = -2$; then place a fifth agent at 0 and link it to any agent at distance at most 2. As indicated in Fig. 1, even though the agents themselves converge, their communication graph does not.



Figure 1: The communication graph of the *HK* system alternates between these two configurations.

3 An Overview of the Proof

Given the challenge of presenting the multiple threads of the argument in digestible form, we begin with a bird’s eye view of the proof. The standard way to establish the convergence of an algorithm or a dynamical system is to focus on a single unknown input and track the rate at which the system expends a certain resource on its way toward equilibrium: a potential function in algorithms; a free energy in statistical physics; or a Lyapunov function in dynamics. This approach cannot work here. Instead, we need to study the system’s action on all inputs at once. This is probably the single most important feature distinguishing natural algorithms from the classical kind: because they run forever, qualitative statements about their behavior will sometimes require a global view of the algorithm’s actions with respect to *all* of its inputs. For this, we need a language that allows us to model the evolution of phase space as a single geometric object. This is our next topic. As explained earlier, we may assume that $d = 1$.

The coding tree. This infinite rooted tree encodes into one geometric object the set of all orbits and the full symbolic dynamics. It is the system’s “Rosetta stone,” from which everything of interest can be read off. The coding tree \mathcal{T} is embedded in $\Omega^n \times \mathbb{R}$, where $\Omega = (0, 1)$ and the last dimension represents time.¹⁰ Each child v of the root is associated with an atom U_v , while the root itself stands for the phase space Ω^n . The *phase tube* (U_v, V_v) of each child v is the “time cylinder” whose cross-sections at times 0 and 1 are U_v and $V_v = f(U_v)$, respectively. In general, a phase tube is a discontinuity-avoiding sequence of iterated images of a given cell in phase space. The tree is built recursively by subdividing V_v into the cells c formed by its intersection with the atoms, and attaching a new child w for each c : we set $V_w = f(c)$ and $U_w = U_v \cap f^{-t_v}(c)$, where t_v is the depth of v (Fig. 2). The phase tube (U_v, V_v) consists of all the cylinders whose cross-sections at $t = 0, \dots, t_v$ are, respectively, $U_v, f(U_v), \dots, f^{t_v}(U_v) = V_v$. Intuitively, \mathcal{T} divides up the phase space into maximal regions over which the iterated map is linear.

The coding tree has three structural parameters that require investigation. One of them is combinatorial. Label each node w of the tree by the unique atom that contains the cell c defined above. This allows us to interpret any path as a word of atom labels and define the language $L(\mathcal{T})$ of all such words: the word-length growth of $L(\mathcal{T})$ plays a central role, which we capture with the *word-entropy* (formal definitions below). The two other parameters are geometric: the *thinning rate* tells us how fast the tree’s branches thin out; the *attraction rate* tells us how close to “periodic” the branches become. Whereas the latter concerns the behavior of single orbits, the thinning rate indicates how quickly a ball in the space of orbits contracts with time, or equivalently how quickly the distribution of agent positions loses entropy.

How do we read periodicity off from the coding tree? Intuitively, one would expect that, at some time ν called the *nesting time*, for every v of depth $t_v = \nu$, there exists w at the same depth with $V_v \subseteq U_w$. In other words the bottom sections of the phase

¹⁰ By convexity, we can restrict the phase space to Ω^n .




Figure 2: Node w at depth $t_w = 2$ in the coding tree, with its phase tube (U_w, V_w) (curved for aesthetic reasons). The cell U_w lies within a single atom whereas V_w splits over two of them. Time points downwards. We have represented only two of the possibly many children of the root.

tubes will, suitably permuted, fit snugly within the top sections. This is not always true, however, and to find necessary conditions for it necessitates a delicate bifurcation analysis. Fig. 3 suggests a visual rule-of-thumb to guide our intuition in distinguishing between chaos and periodicity: the set \mathcal{R} consists of the points in phase space where the map f is not continuous.




Figure 3: Two scenarios for the iterated preimages of the SP discontinuities \mathcal{R} : the set $\mathcal{R} \cup f^{-1}(\mathcal{R}) \cup \dots \cup f^{-t}(\mathcal{R})$ depicted on the left seems to spread everywhere in phase space so as to cover all of it eventually, a symptom of chaos; the set on the right tends to fall into clusters or escape outside of Ω^n , a sign of periodicity.

The algorithmic pipeline. We assemble the coding tree by glueing together smaller coding trees defined recursively. We entrust this task to the *arborator*, an algorithm expressed in a language for “lego-like” assembly. The arborator needs two (infinite) sets of parameters to do its job, the *coupling times* and the *renormalization scales*. To produce these numbers, we use the *flow tracker*, an algorithm that, in the bidirectional

case, works roughly like this: (i) declare agent 1 *wet*; (ii) any dry agent becomes wet as soon as it links to a wet one; (iii) if all agents ever become wet, dry them all and go back to (i). The instants t_k at which wetness propagates constitute the coupling times; the renormalization scales are given by the number w_k of wet agents at time t_k . The key idea is that, between two coupling times t_k and t_{k+1} , the system breaks up into two subsystems with interaction between them going only in one direction: from wet to dry.¹¹ We denote by $\mathcal{A}(p \rightarrow q)$ an influence system that consists of two groups of size p and q , with none of the q agents ever linking up to any of the p agents. This allows us a recursive decomposition of the overall system:

$$\mathcal{A}(n \rightarrow 0)$$

For $k = 1, 2, \dots$

Run $\mathcal{A}(w_k \rightarrow n - w_k)$ and $\mathcal{A}(n - w_k \rightarrow 0)$ concurrently between times t_k and t_{k+1} .

This formulation is of interest only if we can bound $t_{k+1} - t_k$. This is done implicitly by recursively monitoring the long-term behavior of the two subsystems and inferring from it the possibility of further wetness propagation. The flow tracker is a *syntactical* device because it merely monitors the exchange of information among agents with no regard for what is done with it. By contrast, the arborator models the agents' *interpretation* of that information into a course of action. The arborator is assembled as a recursive arithmetic expression over four operations: \oplus , \otimes , **absorb**, and **renorm** (Fig. 4). It comes with a dictionary that spells out the effect of each term on the coding tree's structural parameters. Here is a quick overview:

- The direct sum \oplus models the parallel execution of two independent subsystems. Think of two agents, Bob and Alice, interacting with each other in one corner of the room while Carol and David are chatting on the other side. The coding tree of the whole is the (pathwise) Cartesian product of both two-agent coding trees.
- The direct product \otimes performs tree surgery. It calls upon another primitive, **absorb**, to prune the trees and prepare their phase tubes for "glueing." Imagine Alice suddenly turning to Carol and addressing her. The flow tracker records that the two groups, Bob-Alice and Carol-David, are no longer isolated. Since this

¹¹ This does not mean that the dynamics within the dry agents is not influenced by the wet ones: only that dry agents do not include wet ones in the averaging. The standout exception is the case of a metrical system, where the dry agents act entirely independently of the wet ones between t_k and t_{k+1} .




Figure 4: The two tensor products.

might not have happened had Alice been at a slightly different location, the phase tube leading to this event may well split into two parts: one bearing witness to the new interaction; and the other extending the direct sum unchanged. By analogy with the addition of an absorbing state to a Markov chain, the first operation is called **absorb**.¹²

- The primitive **renorm**, so named for its kinship with the renormalization group of statistical physics, uses the renormalization scales to compress subtrees into single nodes so as to produce (nonuniform) time rescaling.

Attraction and chaos. The occurrence of chaos is mediated by the tension between two forces: dissipation causes the phase tubes to become thinner, which favors periodicity; phase tube splitting produces a form of expansion conducive to chaos. Two arbitrarily close orbits can indeed diverge wildly once they fall on both sides of a discontinuity. The phase tubes snake around phase space while getting thinner at an exponential rate, so hitting *SP* discontinuities should become increasingly rare over time. The problem is that branching multiplies the chances of hitting discontinuities. For dissipation to overcome branching, the average node degree should be small. To show this is indeed the case requires a fairly technical rank argument about the linear constraints implied by the splitting of a phase tube.

The thinning rate is about contraction, not attraction. To see why, consider a trivial system with only self-loops: it is stuck at a fixed point, yet the agents' marginal distributions suffer no loss of entropy.¹³ The information-theoretic interpretation of thinning is illuminating. As agents are attracted to a limit cycle, they lose memory of

¹² The dynamics multiplies transition matrices to the left. Looking at it dually, the rightward products model a random walk over a time-varying graph. The operation **absorb** involves adding a new leaf, which is similar to adding an absorbing state; the direct product glues the root of another coding tree at that leaf.

¹³ This is not to be confused with the word-entropy or the topological entropy.




Figure 5: Given the specification of the natural algorithm, the *flow tracker* computes the coupling times and the renormalization scales, which are needed by the *arborator* to assemble the coding tree. A dictionary allows us to bound the coding tree’s structural parameters by examining the arborator one component at a time. Renormalization makes this a recursive process, hence the loop between the arborator and the parameters box.

where they came from, something that would not happen in a chaotic system. Paradoxically, interaction can then act as a memory recovery device and thus delay the onset of periodicity.

Say the group Alice-Bob-Carol is isolated from David, until the latter decides to interact with Alice, thus taking in a fixed fraction of her entropy. Fast-forward. Alice is now caught in a limit cycle with Bob and Carol, while David has yet to interact with anyone since his earlier contact with Alice. His isolation means that he has had no chance to shed any of Alice’s entropy. Although later caught in a periodic orbit, Alice might still be subject to tiny fluctuations, leading to a sudden interaction with David. When this happens, she will recapture part of the entropy she had lost: she will recover her memory! Happy as the news might be to her, this only delays the inevitable, which is being caught yet again in a limit cycle. Memory recovery cannot recur forever because David loses some of his own memory every time. In the end, because of dissipation, all the agents’ memory will be lost.

4 Algorithmic Dynamics

We flesh out the ideas above, beginning with a simple local characterization of periodicity. We then proceed to define the coding tree (§4.2), the arborator (§4.3), and the flow tracker (§4.4).

4.1 Conditions for asymptotic periodicity

It is convenient to thicken the discontinuities. This does not change the dynamics of the system and is used only as an analytical device. Fix a small parameter $\varepsilon > 0$ once and for all, and, for any $t \geq 0$, define the *margin* \mathcal{R}_ε , where

$$\mathcal{R}_\varepsilon = \bigcup_{SP} \left\{ \mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n : |a_0 + a_1 x_1 + \dots + a_n x_n + \delta| \leq \varepsilon \right\}, \quad (1)$$

where the union extends over all the *SP* discontinuities. The margin is made of $n^{O(1)}$ closed slabs of width at least $\varepsilon n^{-O(1)}$. It is useful to classify the initial states by how long it takes their orbits to hit \mathcal{R}_ε , if ever. With $f^0 = \text{Id}$ and $\min \emptyset = \infty$, we define the label of $\mathbf{x} \in \Omega^n$ as

$$\ell(\mathbf{x}) = \min \left\{ t \geq 0 \mid f^t(\mathbf{x}) \in \mathcal{R}_\varepsilon \right\}.$$

The point \mathbf{x} is said to *vanish* at time $\ell(\mathbf{x})$ if its label is finite. As we shall see, the analysis needs to focus only on the nonvanishing points. Write $\mathcal{S}_t = \{ \mathbf{x} \in \Omega^n = (0, 1)^n \mid \ell(\mathbf{x}) \geq t \}$ for the set of points that do not vanish before time t : \mathcal{S}_0 is Ω^n ; and, for $t > 0$,

$$\mathcal{S}_t = \Omega^n \setminus \bigcup_{k=0}^{t-1} f^{-k}(\mathcal{R}_\varepsilon).$$

Each of its connected components is specified by a set of strict linear inequalities in \mathbb{R}^n , so \mathcal{S}_t is a union of disjoint open n -cells, whose number we denote by $\#\mathcal{S}_t$. We redefine an atom to be a cell of \mathcal{S}_1 and restrict the domain of f to these new atoms. Each cell of \mathcal{S}_{t+1} lies within a cell of \mathcal{S}_t . The limit set $\mathcal{S}_\infty = \bigcap_{t \geq 0} \mathcal{S}_t$ collects the points that never vanish. Unlike those of \mathcal{S}_t , its cells may not be open or full-dimensional.

Periodic sofic shifts. Any cell c of $\mathcal{S}_\infty \subseteq \mathcal{S}_1$ lies within a single atom, so we can define $f|_c$ as the linear map corresponding to the transition matrix P_c . Since \mathcal{S}_∞ is an invariant set, the image $f(c)$ must, by continuity, lie entirely within a cell of \mathcal{S}_∞ . Suppose that $\#\mathcal{S}_\infty < \infty$, a fact we will prove shortly. We define a directed graph F , with each node labeled by a cell c of \mathcal{S}_∞ and with an edge (c, c') , labeled by $f|_c$, joining c to the unique cell c' of \mathcal{S}_∞ that contains $f(c)$. The system forms a *sofic shift* (ie, a regular language over the edge labels). Furthermore, F is *functional*, meaning that each node has exactly one outgoing edge (possibly a self-loop), so any infinite path ends up in a cycle. The *trajectory* of a point \mathbf{x} is the string $s(\mathbf{x}) = c_0 c_1 \dots$ of atoms that it visits: $f^t(\mathbf{x}) \in c_t$ for all $0 \leq t < \ell(\mathbf{x})$. It is infinite if and only if \mathbf{x} does not vanish, so all infinite trajectories are eventually periodic. The weakness of this result is that it might be a statement about the empty set. To strengthen it, we declare the system to be *nesting at t* if no cell c of \mathcal{S}_t contains more than one cell of \mathcal{S}_{t+1} . (This does not mean that $f(c)$ lies inside an atom.) The minimum value of t is called the *nesting time* ν of the system. Observe that $\#\mathcal{S}_\nu \geq \#\mathcal{S}_t$, for any $t \geq \nu$. We bound the nesting time and then proceed with an alternative characterization of nesting.

LEMMA 4.1. *Both the nesting time ν and the number of cells in \mathcal{S}_t are bounded by $(n/\varepsilon)^{O(n)}$, for $t = 0, 1, \dots, \infty$.*

Proof. We begin with the second claim. If, in (1), we replace \mathbf{x} by $P\mathbf{x}$, for a stochastic matrix P , the coefficients of the affine form remain polynomially bounded, so the cells of \mathcal{S}_t are separated from one another by slabs of thickness at least $\varepsilon n^{-O(1)}$. A simple volume argument implies an upper bound of $(n/\varepsilon)^{O(n)}$ on the number of such cells. To bound the nesting time, consider this procedure: suppose that we have placed a special point (called a *witness*) in each cell c of \mathcal{S}_t . If c contains only one cell of \mathcal{S}_{t+1} , we move its witness to that unique cell; if it contains more than one cell, then we move the witness to one of them and create new witnesses to supply the others; if c contains no cell of \mathcal{S}_{t+1} , we leave its witness in place. We carry out this process for $t = 0, 1, \dots$, beginning with a single witness in \mathcal{S}_0 . Witnesses may move around but never disappear; furthermore, by the previous argument, any two of them are separated by at least $\varepsilon n^{-O(1)}$, so their number is bounded by $(n/\varepsilon)^{O(n)}$. Any time t at which the system fails to be nesting sees the creation of at least one new witness, and the first claim follows. \square

LEMMA 4.2. *Given any cell c of \mathcal{S}_t and $k \leq t$, the function $f|_c^k$ is linear. Given any cell $b \subseteq \Omega^n$ and any linear function g , if $g(b) \setminus \mathcal{R}_\varepsilon$ is connected then so is $b \setminus g^{-1}(\mathcal{R}_\varepsilon)$.*

Proof. To call $f|_c^k$ linear is to say that f^k is described by a single stochastic matrix over all of c . We may assume that $t > 0$. Given a cell $c \subseteq \mathcal{S}_t$, none of the cells $c, f(c), \dots, f^{t-1}(c)$ intersect \mathcal{R}_ε , hence each one falls squarely within a single atom and $f|_c^k$ is linear for any $k \leq t$. For the second claim, note that, if the cell b intersects more than one connected component of $\Omega^n \setminus g^{-1}(\mathcal{R}_\varepsilon)$, then it contains a segment pq and a point $r \in pq$ such that g maps p and q outside of \mathcal{R}_ε and r inside of it. By linearity, $g(r)$ lies on the segment $g(p)g(q)$; therefore $g(b) \setminus \mathcal{R}_\varepsilon$ is nonconvex hence disconnected. \square

LEMMA 4.3. *The nesting time ν is the minimum t such that $f^t(c) \setminus \mathcal{R}_\varepsilon$ is connected for each cell c of \mathcal{S}_t ; as a corollary, if c is a cell of \mathcal{S}_ν , then $f(c)$ intersects at most one cell of \mathcal{S}_ν .*

Proof. The claims are trivial if $\nu = 0$, so assume that $\nu > 0$. For the first claim, it suffices to show that the system is nesting at time $t > 0$ if and only if $f^t(c) \setminus \mathcal{R}_\varepsilon$ is connected for each cell c of \mathcal{S}_t . For the “only” part, we show why $f^t(c) \setminus \mathcal{R}_\varepsilon$ must be connected. By Lemma 4.2, $f|_c^t$ is linear; therefore, since $c' = c \setminus f^{-t}(\mathcal{R}_\varepsilon)$ is connected so is $f^t(c') = f^t(c) \setminus \mathcal{R}_\varepsilon$. Conversely, assuming that each set $f^t(c) \setminus \mathcal{R}_\varepsilon$ is connected, then we identify the function g in Lemma 4.2 with $f|_c^t$ (in its linear extension) and conclude that $c \setminus g^{-1}(\mathcal{R}_\varepsilon) = c'$ is connected, hence constitutes the sole cell of \mathcal{S}_{t+1} lying within c . To prove the corollary, again we turn to Lemma 4.2 to observe that $f|_c, \dots, f|_c^\nu$ are all linear, hence so is $g = f|_b^{\nu-1}$, for $b = f(c)$. Our new characterization of nesting implies that $f^\nu(c) \setminus \mathcal{R}_\varepsilon = g(b) \setminus \mathcal{R}_\varepsilon$ is connected, hence so is $b \setminus g^{-1}(\mathcal{R}_\varepsilon) = f(c) \setminus f^{1-\nu}(\mathcal{R}_\varepsilon)$.

Since $f(c)$ lies entirely within a cell of $\mathcal{S}_{\nu-1}$, the labels of its points are all at least $\nu-1$. Removing from $f(c)$ the points of label $\nu-1$ leaves the connected set $f(c) \setminus f^{1-\nu}(\mathcal{R}_\varepsilon)$; therefore, $f(c)$ can intersect at most one cell of \mathcal{S}_ν . \square

We define the directed graph F with one node per cell c of \mathcal{S}_ν and an edge from c to c' , where c' is the unique cell of \mathcal{S}_ν , if it exists, that intersects $f(c)$. Every trajectory corresponds to a directed path in F . The main difference with the previous graph is that the converse is not true. Not only a node may lack an outgoing edge but, worse, nothing in this framework keeps an orbit from going around a cycle for a while only to vanish later. The previous lemma's failure to ensure that $f(c)$ lies strictly within another cell of \mathcal{S}_ν puts periodicity in jeopardy. Perturbation is meant to get around that difficulty. Periods and preperiods are defined with respect to the paths of F , not trajectories: since the correspondence from paths to trajectories is not injective, the latter may have shorter periods.

LEMMA 4.4. *The system is nesting at ν and any time thereafter. Any nonvanishing orbit is eventually periodic and the sum of its period and preperiod is bounded by $\#\mathcal{S}_\nu$.*

The attraction rate. Assume that $\nu > 0$ and let c be a cell of \mathcal{S}_ν . Identifying the nodes of F with their cells in \mathcal{S}_ν , we denote by $\sigma_0, \sigma_1, \dots$ the path from $c = \sigma_0$. Let j be the smallest index such that $\sigma_i = \sigma_j$ for some $i < j$. This defines the period $p = p(c) = j - i$ and the preperiod $q = q(c) = i$, with $p + q \leq \#\mathcal{S}_\nu$. Given any $\mathbf{x} \in c$, its trajectory $s(\mathbf{x}) = c_0 c_1 \cdots c_{\ell(\mathbf{x})-1}$ is such that c_k is the atom containing the cell σ_k . Furthermore, for any $q \leq t \leq \ell(\mathbf{x})$,

$$f^t(\mathbf{x}) = M_{t-q \pmod{p}} Q^{\lfloor (t-q)/p \rfloor} f^q(\mathbf{x}), \quad (2)$$

where $M_k = P_{c_{q+k-1}} \cdots P_{c_q}$, for $k = 0, \dots, p-1$, and $Q = M_p$, with M_0 the identity matrix.¹⁴ Because of the self-loops in the communication graphs, the powers of Q are known to converge to a matrix \tilde{Q} [49]. Given c and $t \geq 0$, we define

$$\Pi_t = M_{t-q \pmod{p}} \tilde{Q} P_{c_{q-1}} \cdots P_{c_0}.$$

The approximation Π_t is one of p matrices obtained by substituting \tilde{Q} for as many "chunks" $Q = M_p$ we can extract from the matrix product $P_{c_{t-1}} \cdots P_{c_0}$ that defines $f|_c$. Note that this includes the case $t < p + q$, where no such chunk is to be found. Given any real $\alpha > 0$, we define the *attraction rate* θ_α as the maximum value of $\theta_\alpha(c)$, over all cells c of \mathcal{S}_ν , where

$$\theta_\alpha(c) = \min \left\{ \theta \geq 0 : \|f^t(\mathbf{x}) - \Pi_{t \pmod{p}} \mathbf{x}\|_\infty \leq \alpha, \right. \\ \left. \text{for all } \mathbf{x} \in c \text{ and } \theta \leq t \leq \ell(\mathbf{x}) \right\}. \quad (3)$$

¹⁴ Note that $f^q(\mathbf{x}) = P_{c_{q-1}} \cdots P_{c_0} \mathbf{x}$, with the matrix denoting the identity if $q = 0$.

Suppose that Q can be written as

$$Q = \begin{pmatrix} A & C \\ \emptyset & B \end{pmatrix} \quad (4)$$

and assume the existence of a limit matrix \tilde{B} such that $\|B^t - \tilde{B}\|_{\max} \leq e^{-\gamma t}$, for some $\gamma > 0$. We tie the attraction rate to the maximum row-sum in A , which is itself related to the thinning rate (whose formal definition we postpone).

LEMMA 4.5. *Given Q as in (4) and an upper bound μ on $\|A\mathbf{1}\|_{\infty}$ such that $e^{-\gamma} \leq \mu < 1$, for any $0 < \alpha < 1 - \mu$,*

$$\theta_{\alpha} = O\left(\frac{\#\mathcal{S}_{\nu}}{1 - \mu}\right) \log \frac{n}{\alpha}.$$

Proof. For any $t > 0$,

$$Q^t = \begin{pmatrix} A^t & C_t \\ 0 & B^t \end{pmatrix}.$$

The matrix A is strictly substochastic ($\mu < 1$), so, by standard properties of a Markov chain's fundamental matrix, $\sum_{k \geq 0} A^k = (I - A)^{-1}$; therefore, for $t > 0$,

$$C_t - (I - A^t)(I - A)^{-1}C\tilde{B} = \sum_{k=0}^{t-1} A^{t-k-1}CB^k - \sum_{k=0}^{t-1} A^kC\tilde{B} = \sum_{k=0}^{t-1} A^{t-k-1}CD_k,$$

where $D_k = B^k - \tilde{B}$. Since C is substochastic, $\|CD_k\|_{\max} \leq e^{-\gamma k}$. From $\|A^k\mathbf{1}\|_{\infty} \leq \mu^k$, we derive

$$\|A^{t-k-1}CD_k\|_{\max} \leq \mu^{t-k-1}e^{-\gamma k}.$$

Since $\mu \geq e^{-\gamma}$, it follows that $\|C_t - (I - A^t)(I - A)^{-1}C\tilde{B}\|_{\max} \leq t\mu^{t-1}$; hence,

$$\tilde{Q} = \begin{pmatrix} 0 & (I - A)^{-1}C\tilde{B} \\ 0 & \tilde{B} \end{pmatrix},$$

where, by $\|A^t\|_{\max} \leq \mu^t$ and $\|(I - A)^{-1}\|_{\max} \leq 1/(1 - \mu)$, $\|Q^t - \tilde{Q}\|_{\max} = O(tn\mu^{t-1}/(1 - \mu))$. As a result, by Lemma 4.4, $\theta_{\alpha} \leq q + pt \leq (\#\mathcal{S}_{\nu})t$, if t satisfies $t\mu^{t-1} < \alpha(1 - \mu)n^{-b}$, for a large enough constant $b > 0$. \square

This next result argues that, although a vanishing point may take arbitrarily long to do so, it comes close to vanishing fairly early. This gives us a useful analytical device to avoid summing complicated series when estimating the probability that a point will eventually vanish under random margin perturbation.

LEMMA 4.6. *Given any finitely-labeled point \mathbf{x} in a cell c of \mathcal{S}_{ν} , there exists $t < \theta_{\alpha} + p(c) + q(c)$ such that $f^t(\mathbf{x}) \in \mathcal{R}_{2\varepsilon}$, for some $\alpha \geq \varepsilon n^{-O(1)}$.*

Proof. We can obviously assume that $\ell(\mathbf{x}) \geq \theta_\alpha + p(c) + q(c)$. For any t such that $\theta_\alpha \leq t \leq \ell(\mathbf{x})$, $f^t(\mathbf{x})$ lies in an ℓ_∞ ball of radius α centered at $\Pi_{t \bmod p(c)} \mathbf{x}$. This means that, between times $\theta_\alpha + q(c)$ and $\ell(\mathbf{x})$, the orbit of \mathbf{x} lies entirely in the union of $p(c)$ balls of radius α and, by periodicity, each ball is visited before time $\theta_\alpha + p(c) + q(c)$. Since \mathbf{x} vanishes at time $\ell(\mathbf{x})$, one of these $p(c)$ balls must intersect \mathcal{R}_ε . Thickening all the margin slabs by a width of $4\alpha\sqrt{n}$ is enough to cover that ball entirely. If $\alpha = \varepsilon n^{-b}$ for a large enough constant b , replacing ε by 2ε achieves the required thickening. \square

Although nesting occurs within finite time, the strict inclusion $\mathcal{S}_{k+1} \subset \mathcal{S}_k$ may occur infinitely often. We show why:

EXAMPLE 4.1: Vanishing can take arbitrarily long. Consider the two-agent influence system

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \xrightarrow{f} \frac{1}{3} \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix},$$

with the SP discontinuities formed by the single slab

$$\mathcal{R}_\varepsilon = \{ \mathbf{x} \in \mathbb{R}^3 : |x_1 - 1 + \delta| \leq \varepsilon \}.$$

For simplicity, assume the same linear map f in the two atoms. It follows that

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \xrightarrow{f^t} \frac{1}{2} \begin{pmatrix} 1 + 3^{-t} & 1 - 3^{-t} \\ 1 - 3^{-t} & 1 + 3^{-t} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}.$$

The set \mathcal{S}_∞ is the complement within Ω^2 (the effective phase space) of

$$\bigcup_{t=0}^{\infty} \left\{ |(1 + 3^{-t})X_1 + (1 - 3^{-t})X_2 - 2 + 2\delta| \leq 2\varepsilon \right\}.$$

Note that if $\varepsilon = 0$, the number of cells in \mathcal{S}_∞ is infinite: they are defined by an infinite number of lines passing through $(1 - \delta, 1 - \delta)$, with increasing slopes tending to -1 . As soon as we allow thickness $\varepsilon > 0$, however, the margin creates only $O(|\log \varepsilon|)$ cells. Not all of them are open. To see this, consider the point $2(1 - \delta + \varepsilon, 0)$. It never vanishes yet any neighborhood contains points that do. Some points take arbitrarily long to vanish. Thickening the SP discontinuities into slabs is a “finitizing” device meant to keep the number of ∞ -labeled cells bounded.

4.2 The coding tree

The richly decorated tree \mathcal{T} encodes the branching structure of the sets \mathcal{S}_k as a geometric object in “phase space \times time” $= \Omega^n \times [0, \infty)$. Recall that each atom c comes with its own transition matrix P_c . Unless specified otherwise, a fixed perturbation value of δ is assumed once and for all. Think of U_v and V_v as the end-sections at times 0 and t_v of a phase tube containing all the orbits originating from U_v . At time t_v , the SP discontinuities might split the tube. This happens only if V_v intersects the margin, which is the “branching condition” in the boxed algorithm. That intersection indicates the vanishing time of some points in U_v , so we place a leaf as an indicator, and call it a

vanishing node. Whereas U_v is an open n -cell, V_v can be a cell of any dimension; hence so can be the connected components of $V_v \setminus \mathcal{R}_\varepsilon$. For each one, c , we attach a new child w to v and denote by P_w the matrix of the map's restriction to c . The image of c at time t_v , ie, $P_w c$, forms the end-section V_w of a new phase tube from the root, whose starting section U_w is the portion of U_v mapping to c at time t_v (Fig. 2).¹⁵

Building \mathcal{T}

- [1] The root v has depth $t_v = 0$; set $U_v \leftarrow V_v \leftarrow \Omega^n$.
- [2] Repeat forever:
 - [2.1] For each newly created node v :
 - If $V_v \cap \mathcal{R}_\varepsilon \neq \emptyset$ [**branching condition**] then create a leaf and make it a child of v .
 - For each cell c of $V_v \setminus \mathcal{R}_\varepsilon$, create a child w of v and set $P_w \leftarrow P_c$; $V_w \leftarrow P_w c$; $U_w \leftarrow U_v \cap f^{-t_v}(c)$.

Let $ww'w'' \dots$ denote the upward, t_w -node path from w to the root (but excluding the root). Using the notation $P_{\leq w} = P_w P_{w'} P_{w''} \dots$, we have the identity $V_w = P_{\leq w} U_w$. No point in U_w vanishes before time t_w , and, in fact, $\mathcal{S}_k = \bigcup_w \{U_w \mid t_w = k\}$. The points of \mathcal{S}_∞ are precisely those whose orbits follow an infinite path $v_\infty = v_0, v_1, v_2, \dots$ down the coding tree. Each such path has its own limit cell $U_{v_\infty} = \bigcap_{t \geq 0} U_{v_t}$: collectively, these form the cells of \mathcal{S}_∞ . Example 4.1 features two infinite paths each of whose nodes has two children, one vanishing and one not.

- The *nesting time* $\nu = \nu(\mathcal{T})$ is the minimum depth at which any node has at most one nonvanishing child (Lemma 4.3); visually, below depth ν , the tree consists of single paths, some finite, others infinite, with vanishing leaves hanging off some of them. A node v is *deep* if $t_v > \nu$ and *shallow* otherwise.
- The *word-entropy* $h(\mathcal{T})$ is the logarithm of the number of shallow nodes.¹⁶ As we observed, $\mathcal{S}_\nu = \bigcup_v \{U_v \mid t_v = \nu\}$; therefore $\#\mathcal{S}_\nu \leq 2^{h(\mathcal{T})}$.
- The *period* $p(\mathcal{T})$ is the maximum value of $p(c)$ for all cells $c = U_v$, with $t_v = \nu$. The attraction rate $\theta_\alpha(\mathcal{T})$ is the maximum value of the attraction rate for any such c .

¹⁵ Note that U_w cannot be defined as the portion of U_v mapping to V_w at time t_w : the orbits must pass through c .

¹⁶ The trajectories form a language $L(\mathcal{T})$ over the alphabet of atom labels. Its growth rate plays a key role in the analysis and is bounded via the word-entropy.

The global coding tree. Let \mathbb{I} denote the interval $[-1, 1]$. Since not all perturbations δ are equally good, we must understand how the coding tree \mathcal{T} varies as a function of δ . To do that, a global approach is necessary: given $\Delta \subseteq \mathbb{I}$, we encode the coding trees for all $\delta \in \Delta$ into a single one, \mathcal{T}^Δ , which can be viewed as the standard coding tree for the augmented $(n + 1)$ -dimensional system $(\mathbf{x}, \delta) \mapsto (f(\mathbf{x}), \delta)$, with the phase space $\Omega^n \times \Delta$. The sets U_v and V_v are now cells in \mathbb{R}^{n+1} . In the branching condition, one should replace the margin \mathcal{R}_ε , as defined in (1), by the *global margin*:

$$\bigcup_{SP} \left\{ (\mathbf{x}, \delta) \in \mathbb{R}^{n+1} : |a_0 + a_1x_1 + \dots + a_nx_n + \delta| \leq \varepsilon \right\}. \quad (5)$$

The degree of any node is bounded by $n^{O(n)}$, which is the maximum number of cells in an arrangement of $n^{O(1)}$ hyperplanes in \mathbb{R}^{n+1} . The definition of nesting can be extended, unchanged, to this lifted system. Since a standard coding tree is just a “cross-section” of the global one, nesting in \mathcal{T} even for all δ does not imply nesting in \mathcal{T}^Δ .¹⁷ The global word-entropy $h(\mathcal{T}^\Delta)$ is defined in the obvious way.

4.3 The arborator

This algorithm assembles the coding tree by glueing smaller pieces together. It relies on a few primitives that we now describe. The direct sum and direct product are tensor-like operations used to attach coding trees together. The primitives **absorb** and **renorm** respectively prune and compress trees. We present these operations and assemble the dictionary that allows us to bound the coding tree’s parameters as we parse the arborator.

Direct sum. The coding tree $\mathcal{T} = \mathcal{T}_1 \oplus \mathcal{T}_2$ models two *independent* systems of size n_1 and n_2 . Independence means that the systems are *decoupled* (no edge joins agents from distinct groups) and *oblivious* (no *SP* discontinuity has nonzero coefficients from both groups): this implies that the two systems can be analyzed separately; decoupling alone is not sufficient. The phase space of the direct sum is of dimension $n = n_1 + n_2$. A path w_0, w_1, \dots of \mathcal{T} is a pairing of paths in the constituent trees: the node w_t is of the form (u_t, v_t) , where u_t (resp. v_t) is a node of \mathcal{T}_1 (resp. \mathcal{T}_2) at depth t ; it is a leaf if and only if u_t or v_t is one—the vanishing of one group implies the vanishing of the whole. If $w = (u, v)$ is not a leaf, then $U_w = U_u \times U_v$, and $V_w = V_u \times V_v$. The direct sum is commutative and associative. The name comes from the fact that P_w is the direct matrix sum of P_u and P_v :

$$P_w = P_u \oplus P_v = \begin{pmatrix} P_u & 0 \\ 0 & P_v \end{pmatrix}.$$

¹⁷ Just as a region in the (X, Y) -plane need not be connected simply because all of its horizontal cross-sections are.

- *Nesting time, period, and attraction rate.*

$$\nu(\mathcal{T}) \leq \max_{i=1,2} \nu(\mathcal{T}_i) \quad \text{and} \quad p(\mathcal{T}) \leq \prod_{i=1,2} p(\mathcal{T}_i) \quad \text{and} \quad \theta_\alpha(\mathcal{T}) \leq \max_{i=1,2} \theta_\alpha(\mathcal{T}_i). \quad (6)$$

The first two inequalities are obvious, so we focus on the last one. Consider a cell $c = c_1 \times c_2$ of \mathcal{S}_ν and follow the path of F emanating from it: this navigation corresponds to the parallel traversal of a path in F_i from c_i —we use the subscript $i = 1, 2$ to refer to either one of the subsystems. Assume without loss of generality that $q_1 \geq q_2$. By definition, to revisit an earlier node means doing likewise in each traversal; hence $q \geq q_1$. At time q_1 , however, both parallel traversals are already engaged in their own respective cycles, so the node pair at time q_1 will be revisited $\text{lcm}(p_1, p_2)$ steps later, the time span that constitutes the period p of the direct sum; it also follows that $q = q_1$. If $q_1 > q_2$, the traversals do not enter their cycles at the same time, so in general, referring to (2), the matrix Q is not the direct sum of Q_1^{p/p_1} and Q_2^{p/p_2} but, rather, of a shifted version $Q = Q_1^{p/p_1} \oplus (BA)^{p/p_2}$, where $Q_2 = AB$. We easily verify that

$$\tilde{Q} = \lim_{k \rightarrow \infty} Q^k = \tilde{Q}_1 \oplus (B\tilde{Q}_2A),$$

where, as before, $\tilde{Q}_i = \lim_{k \rightarrow \infty} Q_i^k$. The use of the ℓ_∞ norm allows us to verify the bound on the attraction rate of the direct sum by checking the accuracy of the approximation for each subsystem separately. It suffices to focus on the case of \mathcal{T}_2 , which presents the added difficulty that the approximation scheme delays the cycle entrance until q_1 . The other difference with the approximation scheme in the original system \mathcal{T}_2 is that, since the period can be much longer, so can the sequence $(M_2)_k$. In all cases, however, the approximation scheme in \mathcal{T} as it applies to \mathcal{T}_2 differs from the scheme in \mathcal{T}_2 in only one substantive way. Consider the language consisting of the words $(AB)^*$, $(AB)^*A$, $(BA)^*$, and $(BA)^*B$. One approximation scheme involves replacing any number of “ AB ”s by \tilde{Q}_2 , while the other scheme replaces any number of “ BA ”s by $B\tilde{Q}_2A$. Because $AB\tilde{Q}_2 = \tilde{Q}_2AB = \tilde{Q}_2$, any application of one scheme or the other produces the same matrix.

- *Word-entropy.* We prove (quasi) subadditivity. Assume without loss of generality that $\nu(\mathcal{T}_1) \geq \nu(\mathcal{T}_2)$. The word-entropy counts the number of shallow nodes $w = (u, v)$. This implies that $t_u \leq \nu(\mathcal{T}_1)$, which limits the number of such nodes u to $2^{h(\mathcal{T}_1)}$. If all the nodes v were shallow in \mathcal{T}_2 , the subadditivity of word-entropy would be immediate; but it need not be the case. If v is deep, let $s(v)$ be its deepest shallow ancestor. The function s may not be injective but it is at most two-to-one. Thus,

$$h(\mathcal{T}) \leq h(\mathcal{T}_1) + h(\mathcal{T}_2) + 1. \quad (7)$$

All of the relations in (6, 7) still hold when the superscript Δ is added to the coding trees. We discuss (7) to illustrate the underlying principle. First, we provide an independent

perturbation variable $\delta_i \in \Delta$ to \mathcal{T}_i ($i = 1, 2$) and add it as an extra coordinate to the state vector, thus lifting system i to dimension $n_i + 1$. By (7), the word-entropy of the joint system \mathcal{T}^* in dimension $n+2$ is at most $h(\mathcal{T}_1^\Delta) + h(\mathcal{T}_2^\Delta) + 1$. Second, we restrict the system \mathcal{T}^* to the invariant hyperplane $\delta_1 = \delta_2$, which cannot increase the word-entropy, hence $h(\mathcal{T}^\Delta) \leq h(\mathcal{T}^*)$, as claimed.

Absorption. The direct product, which we define below, requires an intermediate construction. The goal is to allow the selection of nodes for removal, with an eye toward replacing the subtrees they root by coding trees with different characteristics. The selection is carried out by an operation called `absorb`(\mathcal{T}), which replaces any deleted node by a leaf. For reasons that the *flow tracker* will soon clarify, such leaves are designated *wet*. An orbit that lands into one of these wet leaves is suddenly governed by a different dynamics, modeled by a different coding tree, so from the perspective of \mathcal{T} alone, wet leaves are where orbits come to a halt. While vanishing leaves signal the termination of an orbit (at least from the perspective of the analysis), the wet variety merely indicates a change of dynamics. Here is a simple illustration:

EXAMPLE 4.3: The system consists of two independent subsystems. Suppose we add a union of slabs, denoted by \mathcal{R}'_ε , to the original margin $\mathcal{R}_\varepsilon^o$, thus breaking the direct-sum nature of the coding tree. In Fig. 6, \mathcal{R}'_ε would consist of the two infinite strips bordering b . We keep the transition matrices unchanged everywhere except in cell b , which we call *wet*: all transition matrices are still direct (matrix) sums, with the possible exception of P_b . Suppose we had available the coding tree prior to the margin's augmentation. Let V_v denote the pentagon in the figure and w be the node associated with the trapezoid c that holds a, b, d . We need to replace w by three nodes: two of them for a, d and one, a *wet leaf*, for b . The transition matrices for a, d are both equal to the direct sum P_c , while P_b can be arbitrary. The idea is that b can then be made the region U_{root} of a new coding tree.

Minor technicality: usually, $U_{\text{root}} = \Omega^n$, so the coding tree must be *cropped* by substituting b for Ω^n ; note that b need not be an invariant set. Cropping might involve pruning the tree but it cannot increase any of the key parameters, such as the nesting time, the attraction rate, and the period. Absorption appeals to the fact that we can ignore b and its wet leaf until we have fully analyzed the direct sum. This separation is very useful, especially since absorption does not require a direct sum—we never used the fact that the old slabs were horizontal or vertical—and is therefore extremely general.

A crucial observation is that the nodes z created for the subcells c' of a given c (subcells a and d in Fig. 6) have the same matrix P_c . As a result of all the absorptions, the tube (U_v, V_v) is split up by up to t_v linearly transformed copies of the $n^{O(1)}$ margin slabs, hence into at most $t_v^n n^{O(n)}$ subcells. This compares favorably with the naive upper bound of $n^{O(nt_v)}$ based on the sole fact that absorption at each ancestor of v produces $n^{O(n)}$ children.




Figure 6: The original cell c of V_v (bottom-center trapezoid) splits up into the wet cell b and the dry cells a and d , both of which inherit the matrix P_c .

Absorption surgery

- [1] If v has no leaf, create a vanishing leaf and make it a child of v .
- [2] For each cell c of $V_v \setminus \mathcal{R}_\varepsilon^o$, let w be the child of v for c (ie, such that $f(c) = V_w$) and let T be the tree rooted at w . If $c \cap \mathcal{R}'_\varepsilon$, then remove T and, for each cell c' of $c \setminus \mathcal{R}'_\varepsilon$, create a node z and make it a child of v .
 - If c' is wet, make z a wet leaf.
 - If c' is dry, reattach to z a suitably cropped copy of T .
Set $P_z \leftarrow P_c$, $V_z \leftarrow P_z c$, and $U_z \leftarrow U_v \cap f^{-t_v}(c')$.

Direct product. The tree $\mathcal{T} = \mathcal{T}_1 \otimes \mathcal{T}_2$ models the concatenation of two systems. The direct product is associative but not commutative. It is always preceded by a round of absorptions at one or several nodes of \mathcal{T}_1 . We begin with a few words of intuition. Consider two systems S_1 and S_2 , governed by different dynamics yet evolving in the same phase space Ω^n . Given an arbitrary region $\Lambda \subset \Omega^n$, we define the hybrid system S with the dynamics of S_2 over Λ and S_1 elsewhere. Suppose we had complete knowledge of the coding tree \mathcal{T}_i for each S_i ($i = 1, 2$). Could we then combine them in some ways to assemble the coding tree \mathcal{T} of S ? To answer this question, we follow a three-step approach:

- (i) we absorb the tree \mathcal{T}_1 by creating wet leaves w for all the nodes v with $V_v \cap \Lambda \neq \emptyset$;
- (ii) we attach the roots of cropped copies of \mathcal{T}_2 at the wet leaves; and
- (iii) we iterate and glue \mathcal{T}_1 and \mathcal{T}_2 in alternation, as orbits move back and forth in and out of Λ .

Absorption, direct products, and the arborator address (i, ii, iii) in that order. The root of \mathcal{T}_2 is attached to w , but not until that tree itself has been properly cropped so

that $U_{\text{root}} = V_{\text{root}}(\mathcal{T}_2) = V_w(\mathcal{T}_1) = P_w c$, with P_w given by \mathcal{T}_2 and not \mathcal{T}_1 . To be fully rigorous, we should write a direct product as $\mathcal{T}_1 \otimes \{\mathcal{T}_2\}$ since the trees \mathcal{T}_2 we attach to the wet nodes might not all be the same: the cropping might vary, as might the wet regions.

- *Nesting time and attraction rate.* Bounding the nesting time of a direct product is not merely a combinatorial matter, as was the case for direct sums: the geometry of attraction plays a role. Even the case of $\text{absorb}(\mathcal{T}_1)$ demands some attention and this is where we begin. Adding \mathcal{R}'_ε to the margin cannot create arbitrarily deep wet nodes: specifically, no $v \in \text{absorb}(\mathcal{T}_1)$ of depth at least $\max\{\nu(\mathcal{T}_1), \theta_{\alpha_a}(\mathcal{T}_1) + p\}$ can have a wet child, where $p = p(U_v)$ and $\alpha_a = \varepsilon n^{-a}$ for a large enough constant a . Indeed, suppose there is such a node v . Pick $\mathbf{x} \in U_v$ such that $f^{t_v}(\mathbf{x})$ lies in a wet cell c within V_v and observe that

$$\|f^{t_v}(\mathbf{x}) - f^{t_v-p}(\mathbf{x})\|_\infty \leq \|f^{t_v}(\mathbf{x}) - \Pi_{t_v \pmod{p}} \mathbf{x}\|_\infty + \|f^{t_v-p}(\mathbf{x}) - \Pi_{t_v \pmod{p}} \mathbf{x}\|_\infty \leq 2\alpha_a.$$

By our choice of α_a , this implies that $f^{t_v}(\mathbf{x})$ and $f^{t_v-p}(\mathbf{x})$ are at a distance apart less than the width of the margin's slabs; therefore, $f^{t_v-p}(\mathbf{x})$ lies in a wet cell or in a slab. It follows that the orbit of \mathbf{x} either vanishes or comes to a (wet) halt at a time earlier than t_v , so $\mathbf{x} \notin U_v$ and we have a contradiction. It follows that all deep nodes of \mathcal{T}_1 deeper than $\theta_{\alpha_a}(\mathcal{T}_1) + p(\mathcal{T}_1)$ are also deep in $\text{absorb}(\mathcal{T}_1)$. With $\mathcal{T} = \mathcal{T}_1 \otimes \mathcal{T}_2$, therefore,

$$\begin{cases} \nu(\text{absorb}(\mathcal{T}_1)) \leq \max\{\nu(\mathcal{T}_1), \theta_{\alpha_a}(\mathcal{T}_1) + p(\mathcal{T}_1)\}, & \text{for some } \alpha_a \geq \varepsilon n^{-O(1)} \\ \nu(\mathcal{T}) \leq \nu(\text{absorb}(\mathcal{T}_1)) + \nu(\mathcal{T}_2) \\ \theta_\alpha(\mathcal{T}) \leq \max\{\theta_\alpha(\mathcal{T}_1), \nu(\text{absorb}(\mathcal{T}_1)) + \theta_\alpha(\mathcal{T}_2)\} \end{cases} \quad (8)$$

- *Word-entropy.* Absorption can occur only at nodes v of depth $t_v \leq \nu(\text{absorb}(\mathcal{T}_1))$. This means that the number of nodes where wet cells can emerge is at most $2^{h(\mathcal{T}_1)}(\theta_{\alpha_a}(\mathcal{T}_1) + p(\mathcal{T}_1))$. As we argued earlier, each such node v can give birth to at most $t_v^n n^{O(n)}$ new nodes, so the number of shallow nodes in \mathcal{T} is (conservatively) at most

$$\underbrace{2^{h(\mathcal{T}_1)}(\theta_{\alpha_a}(\mathcal{T}_1) + p(\mathcal{T}_1))}_{\#v \text{ with wet child}} \times \underbrace{(\nu(\mathcal{T}_1) + \theta_{\alpha_a}(\mathcal{T}_1) + p(\mathcal{T}_1))^n n^{O(n)}}_{\#\text{splits}/v} \times \underbrace{2^{h(\mathcal{T}_2)}}_{\#\mathcal{T}_2 \text{ nodes}}.$$

We use the fact that cropping cannot increase the word-entropy. Taking logarithms, we find that

$$h(\mathcal{T}) \leq h(\mathcal{T}_1) + h(\mathcal{T}_2) + (n+1) \log(\nu(\mathcal{T}_1) + \theta_{\alpha_a}(\mathcal{T}_1) + p(\mathcal{T}_1)) + O(n \log n). \quad (9)$$

Since both $\nu(\mathcal{T}_1)$ and $p(\mathcal{T}_1)$ are no greater than $2^{h(\mathcal{T}_1)}$, we can simplify the bound:

$$h(\mathcal{T}) \leq (n+2)h(\mathcal{T}_1) + h(\mathcal{T}_2) + (n+1) \log \theta_{\alpha_a}(\mathcal{T}_1) + O(n \log n). \quad (10)$$

We repeat our earlier observation that, by viewing the perturbation variable δ as an extra coordinate of the state vector, the relations above still hold for global coding trees with n incremented by one.

Renormalization. This operation is both the simplest and the most powerful in the arborator’s toolkit: the simplest because all it does is compress time by folding together consecutive levels of \mathcal{T} ; the most powerful because it reaches beyond lego-like assembly to bring in the full power of algorithmic recursion into the analysis. The primitive **renorm** takes disjoint subtrees of \mathcal{T} and regards them as nodes of the renormalized tree. This is done in the obvious way: if u is any node in \mathcal{T} with two children v_1, v_2 , each one with two children, v_{11}, v_{12} and v_{21}, v_{22} , then compressing the subtree u, v_1, v_2 means replacing it by a node z with the same parent as u ’s (if any) and the four children v_{ij} . We discuss this process in more detail below. Although inspired by the renormalization group of statistical physics, our approach is more general. For one thing, the compressed subtrees may differ in size, resulting in nonuniform rescaling across \mathcal{T} . This lack of uniformity rules out closed-form composition formulae for the nesting time, attraction rate, and word-entropy of renormalized coding trees, which must then be resolved algorithmically.

4.4 The flow tracker

We approach periodicity through the study of an important family, the *block-directional* influence systems, whose agents can be ordered so that

$$\mathcal{G} = \begin{pmatrix} \mathcal{G}_A & \mathcal{G}_C \\ 0 & \mathcal{G}_B \end{pmatrix}, \quad (11)$$

where 0 denotes the $(n-m)$ -by- m matrix whose entries are the constant function $\mathbf{x} \mapsto 0$; in other words, in a block-directional system, no B -agent ever links to an A -agent. Suppose that $m < n$. Wet the B -agents with water while keeping all the A -agents dry. Whenever an edge of the communication graph links a dry agent to a wet one, the former gets wet. Note how water flows in the *reverse* direction of the edges. As soon as all agents become wet (if ever), dry them but leave the B -agents wet, and repeat forever. The case $m = n$ is similar, with one agent designated wet once and for all. The sequence of times at which water spreads or drying occurs plays a key role in building the arborator.

Coupling times and renormalization scales. Let $\mathcal{T}_{m \rightarrow n-m}$ denote the coding tree of a block-directional system consisting of m (resp. $n-m$) A -agents (resp. B -agents). The arrow indicates that no B -agent can ever link to an A -agent: \mathcal{G}_{ij} is identically zero for any B -agent i and A -agent j . We use the notation $\mathcal{T}_{m \parallel n-m}$ for the decoupled case: no edge ever joins the two groups in either direction, but the discontinuities may still mix variables from both groups. Note that the metrical case implies full independence (§1), so that

$$\mathcal{T}_{m \parallel n-m} = \mathcal{T}_m \oplus \mathcal{T}_{n-m}.$$

Assume that $n > 1$ and $0 < m \leq n$. We write $\mathcal{T}_{m \rightarrow 0}$ as \mathcal{T}_m . Likewise, we can always express $\mathcal{T}_{m \rightarrow n-m}$ as \mathcal{T}_m , but doing so is less informative. When the initial state \mathbf{x} is undersood, we use the shorthand $G_t = \mathcal{G}(f^t(\mathbf{x}))$ to designate the communication graph at time t and we denote by W_t the set of wet agents at that time. The flow tracker is not

concerned with information exchanges among the B -agents: these are permanently wet and, should they not exist ($m = m$), agent 1 is kept wet at all times [2.1]. Thus the set W_t of wet agents is never empty. The assignments of t_0 in step [2.3] divide the timeline into *epochs*, time intervals during which either all agents become wet or, failing that, the flow tracker comes to a halt (breaking out of the repeat loop at “stop”). Each epoch is itself divided into subintervals by the *coupling times* $t_1 < \dots < t_\ell$, with $W_{t_k} \subset W_{t_{k+1}}$. The last coupling time t_ℓ marks either the end of the flow tracking (if not all A -agents become get) or one less than the next value of t_0 in the loop.

The notion of coupling is purely *syntactical*, being only a matter of information transfer. Our interest in it is *semantic*, however: as befits a dissipative system, a certain quantity, to which we shall soon return, can be bounded by a decreasing function of time. To get a handle on that quantity is the main purpose of the flow tracker.

Flow tracking in action. Suppose that, for a long period of time, the wet agents fail to interact with any dry one. The two groups can then be handled recursively. While this alone will not tell us whether dry-wet interaction is to occur ever again, it will reveal enough fine-grained information about the groups’ behavior to help us resolve that very question. Suppose that such interaction takes place, to be followed by another long period of interaction. Renormalization squeezes these “non-interactive” periods into single time units, thus providing virtual time scales over which information flows at a steady rate across the system. Thus, besides analyzing subsystems recursively, renormalization brings uniformity to the information transfer rate.

Flow tracker

[1] $t_0 \leftarrow 0$.

[2] Repeat forever:

[2.1] If $m < n$ then $W_{t_0} \leftarrow \{m + 1, \dots, n\}$ else $W_{t_0} \leftarrow \{1\}$.

[2.2] For $t = t_0, t_0 + 1, \dots, \infty$

$$W_{t+1} \leftarrow W_t \cup \{i \mid \exists (i, j) \in G_t \ \& \ j \in W_t\}.$$

[2.3] If $|W_\infty| = n$ then $t_0 \leftarrow \min\{t > t_0 : |W_t| = n\}$ else stop.

EXAMPLE 4.4: The third column below lists a graph sequence G_0, \dots, G_{11} in chronological order, with the superscript w indicating the edges through which water propagates to dry nodes. The system is block-directional with three A -agents labeled a, b, c and one B -agent labeled d . For clarity, we spell out the agents by writing the corresponding coding tree $\mathcal{T}_{3 \rightarrow 1}$ as $\mathcal{T}_{abc \rightarrow d}$, instead, thus indicating that no edge may link d to any of a, b, c .

| Flow tracking | | | |
|---------------|---------------------------|---|-----------------------------------|
| renorm | $W_0 = \{d\}$ | $d \quad a \rightarrow b \rightarrow c$ | $\mathcal{T}_{d \parallel abc}$ |
| | $W_1 = \{d\}$ | $d \quad a \leftarrow b \rightarrow c$ | |
| | $W_2 = \{d\}$ | $d \quad a \rightarrow b \leftarrow c$ | |
| $t_1 = 3$ | $W_3 = \{d\}$ | $d \xleftarrow{w} a \leftarrow b \leftarrow c$ | \mathcal{T}_{abcd} |
| renorm | $W_4 = \{a, d\}$ | $d \leftarrow a \rightarrow b \rightarrow c$ | $\mathcal{T}_{a \rightarrow bcd}$ |
| | $W_5 = \{a, d\}$ | $d \quad a \rightarrow b \rightarrow c$ | |
| $t_2 = 6$ | $W_6 = \{a, d\}$ | $d \leftarrow a \xleftarrow{w} b \leftarrow c$ | \mathcal{T}_{abcd} |
| renorm | $W_7 = \{a, b, d\}$ | $d \leftarrow a \rightarrow b \rightarrow c$ | $\mathcal{T}_{ab \rightarrow cd}$ |
| | $W_8 = \{a, b, d\}$ | $d \leftarrow a \leftarrow b \quad c$ | |
| | $W_9 = \{a, b, d\}$ | $d \leftarrow a \rightarrow b \rightarrow c$ | |
| $t_3 = 10$ | $W_{10} = \{a, b, d\}$ | $d \leftarrow a \rightarrow b \xleftarrow{w} c$ | \mathcal{T}_{abcd} |
| | $W_{11} = \{a, b, c, d\}$ | $d \quad a \leftarrow b \quad c$ | $\mathcal{T}_{d \parallel abc}$ |

In the first renormalized 3-step phase, the system “waits” for an edge from $\{a, b, c\}$ to d , and so can be modeled as $\mathcal{T}_{d \parallel abc}$. In the metrical case, this is further reducible to $\mathcal{T}_d \oplus \mathcal{T}_{abc}$. The times t_1, t_2, t_3 coincide with the growth of the wet set: these are one-step events, which are treated trivially as height-one absorbed trees. They entail no recursion, so inductive soundness is irrelevant and writing the uninformative \mathcal{T}_{abcd} is harmless. The other renormalized phases are counterintuitive and should be discussed. Take the last one: it might be tempting to renormalize it as $\mathcal{T}_{abd \rightarrow c}$ to indicate that the phase awaits the wetting of c (with a, b, d already wet). This strategy is inductively unsound, however, as it attempts to resolve a system $\mathcal{T}_{abc \rightarrow d}$ by means of another one, $\mathcal{T}_{abd \rightarrow c}$, of the same combinatorial type. Instead, we use the fact that not only no edge can link c to $\{a, b\}$ (by definition of the current phase) but no edge can link d to $\{a, b\}$ either (by block-directionality). This allows us to use $\mathcal{T}_{ab \rightarrow cd}$, instead, which is inductively sound.

Renormalization, which is denoted by underlining, compresses into single time units all the time intervals during which wetness does not spread to dry agents. With the subscripts (resp. superscript) indicating the time compression rates (resp. tree height), the 11-node path of $\mathcal{T}_{abc \rightarrow d}$ matching the graph sequence above can be expressed as

$$\underline{\mathcal{T}_{d \parallel abc}}_{|3} \otimes \mathcal{T}_{abcd}^1 \otimes \underline{\mathcal{T}_{a \rightarrow bcd}}_{|2} \otimes \mathcal{T}_{abcd}^1 \otimes \underline{\mathcal{T}_{ab \rightarrow cd}}_{|3} \otimes \mathcal{T}_{abcd}^1.$$

As the example above illustrates, the coupling time t_k is immediately followed by a renormalization phase of the form $\mathcal{T}_{w_k \rightarrow n-w_k}$, where $w_k = |W_{t_k+1}| - n + m$ is the renormalization scale ($k = 1, \dots, \ell - 1$). Thus, any path of the coding tree can be renormalized as

$$\mathcal{T}_{m \rightarrow n-m} \implies \underline{\mathcal{T}_{m \parallel n-m}}_{|t_1} \otimes \mathcal{T}_n^1 \otimes \left\{ \bigotimes_{k=1}^{\ell-1} \left(\underline{\mathcal{T}_{w_k \rightarrow n-w_k}}_{|t_{k+1}-t_k-1} \otimes \mathcal{T}_n^1 \right) \right\} \otimes \mathcal{T}_{m \rightarrow n-m}. \quad (12)$$

The recursion comes in two forms: as calls to inductively smaller subsystems $\mathcal{T}_{w_k \rightarrow n-w_k}$; and as a rewriting rule, $\mathcal{T}_{m \rightarrow n-m} \Rightarrow \dots \{ \} \otimes \mathcal{T}_{m \rightarrow n-m}$. It is the latter that makes the

arborator, if expanded in full, an infinitely long expression. We note that all these derivations easily extend to the global coding trees.

5 Bidirectional Systems

We begin our proof of the bidirectional case of Theorem 1.1 by establishing a weaker result for metrical systems: recall that these make the presence of an edge between two agents a sole function of their distance. The proof is almost automatic and a good illustration of the algorithmic machinery we have put in place. By appealing to known results on the total s -energy, we are able to improve the bounds and extend them to the nonmetrical case.

5.1 The metrical case

It is worth noting that, even for this special case, perturbations are required for any uniform convergence rate to hold.

EXAMPLE 5.1: Consider the 3-agent system:

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \mapsto \frac{1}{3} \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix},$$

with $x_3 \mapsto \frac{1}{2}(x_2 + x_3)$ if $x_3 - x_2 \geq 1$ and $x_3 \mapsto x_3$ else. Initialize the system with $x_2 = -x_1 = 1$ and x_3 slightly bigger than x_2 . The edge joining agents 2 and 3 will then appear only after on the order of $|\log(x_3 - x_2)|$ steps, which implies that the convergence time cannot be bounded uniformly without perturbation.

Fix δ in $\Delta = (n^{-b} \mathbb{I}) \setminus (n^{-2b} \mathbb{I})$, where $\mathbb{I} = [-1, 1]$ and b is a suitably large constant.¹⁸ The margin slabs of a metrical system are of the form $|a_0 + x_i - x_j + \delta| \leq \varepsilon$. Because a_0 is an $O(\log n)$ -bit rational, as long as $\varepsilon < n^{-3b}$, \mathbf{x} cannot lie in that slab if $|x_i - x_j| \leq n^{-3b}$. Let $\text{diam}(s)$ be the diameter of the system after the s -th epoch. From (14) in [17], we conclude that water propagation to all the agents entails the shrinking of the system's diameter by at least a factor of $1 - n^{-O(n)}$. Since an epoch witnesses the wetting of all the agents, repeated applications of this principle yields

$$\text{diam}(s) \leq e^{-sn^{-O(n)}}. \tag{13}$$

After n^{cn} epochs have elapsed (if ever), for a large enough constant c , the diameter of the system falls beneath n^{-3b} and, by convexity, never rises again. By our previous observation, the orbit can never hit a margin subsequently. The maximum time it takes for n^{cn} epochs to elapse, over all $\mathbf{x} \in \Omega^n$ and $\delta \in \Delta$, is an upper bound on the nesting time of the global coding tree. Furthermore, past that time, the communication graph is *frozen*, meaning that it can never change again.

¹⁸ Recall that ideally Δ should be $\{0\}$ so the more confined around 0 we can make it the better; thus a higher value of b is an asset, not a drawback.

LEMMA 5.1. *If P is the transition matrix associated with the undirected communication graph G , there is a matrix Π such that $\|P^k - \Pi\|_{\max} = e^{-kn^{-O(n)}}$, for any $k \geq 0$.*

Proof. By repeating the following argument for each connected component if needed, we can assume that G is connected. The positive diagonal ensures that P is primitive (being the stochastic matrix of an irreducible, aperiodic Markov chain), hence P^n , which we denote by M , is positive. Since each nonzero entry of P is at least $n^{-O(1)}$, the *coefficient of ergodicity* of M , defined as

$$\beta = \frac{1}{2} \max_{i,j} \sum_l |M_{il} - M_{jl}| = 1 - \min_{i,j} \sum_l \min\{M_{il}, M_{jl}\}$$

satisfies $\beta \leq 1 - n^{-O(n)}$. Two classic results from the theory of nonnegative matrices [49] hold that β is an upper bound on the second largest eigenvalue of M (in absolute value) and that β is submultiplicative.¹⁹ Given any probability distribution \mathbf{x} , if $\mathbf{y} = M^l \mathbf{x}$, then

$$\max_{i,j} |y_i - y_j| \leq \beta^l \max_{i,j} |x_i - x_j| \leq e^{-ln^{-O(n)}}. \quad (14)$$

By Perron-Frobenius and the ergodicity of P , its powers tend to the rank-one matrix $\mathbf{1}v^T$, where v is the dominant left-eigenvector of P with unit ℓ_1 -norm; furthermore,

$$\|P^k - v\mathbf{1}^T\|_{\max} = e^{-kn^{-O(n)}}.$$

Indeed, setting \mathbf{x} to the j -th basis vector $(0, \dots, 1, \dots, 0)^T$ in (14) shows that the j -th column of $M^l = P^{ln}$, for $l = \lfloor k/n \rfloor$, consists of identical entries plus or minus a term in $e^{-ln^{-O(n)}}$. By convexity, these near-identical entries cannot themselves oscillate as l grows. Indeed, besides (14), it is also true that $[\min y_i, \max y_j] \subseteq [\min x_i, \max x_j]$. \square

The next step in deriving the coding tree's parameters is to specialize the arborator's expression (12) to the metrical case. The outer product enumerates the first $n^{O(n)}$ epochs leading to the combinatorial (but not physical) "freezing" of the system. The coupling times and renormalization scales might vary from one epoch to the next; to satisfy the rewriting rule below, we set $w_0 = 1$ and $t_0 = -1$. The cropped coding tree \mathcal{T}_n^* models the post-freezing phase.

$$\mathcal{T}_n \implies \left\{ \bigotimes_{s=1}^{n^{O(n)}} \bigotimes_{k=0}^{\ell_s} \left(\mathcal{T}_{w_k} \oplus \mathcal{T}_{n-w_k} \Big|_{t_{k+1}-t_k-1} \otimes \mathcal{T}_n^{|1|} \right) \right\} \otimes \mathcal{T}_n^*. \quad (15)$$

The following derivations entail little more than looking up the dictionary compiled in §4.3.

¹⁹ The stochastic matrix P may not correspond to a reversible Markov chain and might not be diagonalizable. It is primitive, however; therefore, by Perron-Frobenius, it has unique left and right unit eigenvectors associated with the dominant eigenvalue 1.

- *Nesting time and attraction rate.* It is convenient to define

$$\mu_\alpha(\mathcal{T}) = \max\{\nu(\mathcal{T}), \theta_\alpha(\mathcal{T})\}.$$

If the coding trees T_1, \dots, T_k have period one then, by (8),

$$\mu_\alpha\left(\bigotimes_{i=1}^k T_i\right) \leq k + \sum_i \mu_{\alpha_a}(T_i) + \max_i \mu_\alpha(T_i). \quad (16)$$

The coding tree \mathcal{T}_n^* involves a single matrix whose powers converge to a fixed matrix Π and, by Lemma 5.1, $\mu_\alpha(\mathcal{T}_n^*) \leq n^{O(n)} \log \frac{1}{\alpha}$. The following bounds derive from monotonicity and successive applications of (6, 8). For some suitable $\alpha_a = \varepsilon n^{-O(1)}$ and any $\alpha \leq \alpha_a$,

$$\begin{aligned} \mu_\alpha(\mathcal{T}_n) &\leq \mu_{\alpha_a}(\mathcal{T}_n^*) + n^{O(n)} \sum_{k=1}^{n-1} \max\left\{\mu_{\alpha_a}(\mathcal{T}_k), \mu_{\alpha_a}(\mathcal{T}_{n-k})\right\} + \max_k \left\{\mu_\alpha(\mathcal{T}_k), \mu_\alpha(\mathcal{T}_n^*)\right\} \\ &\leq n^{O(n)} \mu_{\alpha_a}(\mathcal{T}_{n-1}) + \mu_\alpha(\mathcal{T}_{n-1}) + n^{O(n)} \log \frac{1}{\alpha} \\ &\leq n^{O(n^2)} \log \frac{1}{\varepsilon} + n^{O(n)} \log \frac{1}{\alpha}. \end{aligned} \quad (17)$$

In view of this last upper bound, the condition $\alpha \leq \alpha_a$ can be relaxed to $\alpha < 1$. Thus,

$$\nu(\mathcal{T}_n) \leq n^{O(n^2)} \log \frac{1}{\varepsilon} \quad \text{and} \quad \theta_\alpha(\mathcal{T}_n) \leq n^{O(n^2)} \log \frac{1}{\varepsilon} + n^{O(n)} \log \frac{1}{\alpha}. \quad (18)$$

- *Word-entropy.* By (9) and the attraction rate bound above, for $0 < \varepsilon < 1/2$,

$$\begin{aligned} h(T_1 \otimes T_2) &\leq h(T_1) + h(T_2) + (n+1) \log(2\mu_{\alpha_a}(T_1) + 1) + O(n \log n) \\ &\leq h(T_1) + h(T_2) + (n+1) \log \log \frac{1}{\varepsilon} + O(n^3 \log n). \end{aligned} \quad (19)$$

By (7) and $h(\mathcal{T}_n^*) = 0$, it follows that

$$\begin{aligned} h(\mathcal{T}_n) &\leq \sum_{s=1}^{n^{O(n)}} \sum_{k=1}^{n-1} \left\{ h((\mathcal{T}_k \oplus \mathcal{T}_{n-k}) \otimes \mathcal{T}_n^{\lfloor 1 \rfloor}) \right\} + h(\mathcal{T}_n^*) \\ &\quad + (n+1) \log(2\mu_{\alpha_a}(\mathcal{T}_n^*) + 1) + n^{O(n)} \log \log \frac{1}{\varepsilon} \\ &\leq n^{O(n)} h(\mathcal{T}_{n-1}) + n^{O(n)} \log \log \frac{1}{\varepsilon} \leq n^{O(n^2)} \log \log \frac{1}{\varepsilon}. \end{aligned}$$

Our earlier observation that such derivations apply to the global coding trees tells us that

$$h(\mathcal{T}_n^\Delta) \leq n^{O(n^2)} \log \log \frac{1}{\varepsilon}. \quad (20)$$

Note the crucial fact that, from the vantage point of (18, 20), the global word-entropy is lower than the nesting time, which shows that the coding tree's average node is less

than 2. By Lemmas 4.4 and 4.6, any vanishing point \mathbf{x} hits an enlarged margin fairly early: $f^{t_a}(\mathbf{x}) \in \mathcal{R}_{2\varepsilon}$ for $t_a \leq \theta_{\alpha_o} + \#\mathcal{S}_\nu$ and some $\alpha_o \geq \varepsilon n^{-O(1)}$; therefore,

$$t_a \leq \theta_{\alpha_o} + 2^{h(\mathcal{T}_n^\Delta)} \leq |\log \varepsilon|^{n^{O(n^2)}}. \quad (21)$$

For random $\delta \in \Delta$, a fixed point lies in a given slab of $\mathcal{R}_{2\varepsilon}$ with probability at most $4\varepsilon/(2n^{-b} - 2n^{-2b})$; by a union bound over the margin slabs, the probability of being in $\mathcal{R}_{2\varepsilon}$ does not exceed $\varepsilon n^{O(1)}$. Therefore, the probability that a fixed \mathbf{x} ever vanishes is at most $\varepsilon n^{O(1)}$ times the number of paths of depth at most t_a in the global coding tree \mathcal{T}_n^Δ , which, by (21), is

$$|\log \varepsilon|^{n^{O(n^2)}} 2^{h(\mathcal{T}_n^\Delta)}.$$

By (20), this puts the vanishing probability at

$$\varepsilon (\log \frac{1}{\varepsilon})^{n^{O(n^2)}} < \sqrt{\varepsilon},$$

for ε small enough, which means that it can be set arbitrarily low. Removing a small interval in the middle of $n^{-b}\mathbb{I}$ to form Δ was only useful for the analysis: in practice, we might as well pick the random perturbation uniformly in $n^{-b}\mathbb{I}$ since it would add only $2n^{-b}$ to the error probability. The merit of the proof is that it is a straightforward, automatic application of the arborator's dictionary. It illustrates the power of renormalization, which can be seen in the fact that no explicit bound on $t_{k+1} - t_k$ is ever needed. By appealing to known results about the *total s-energy* [17] we can both extend and improve the bound on the convergence rate.

5.2 The bidirectional case

To give up the metrical assumption means that the presence of an edge in the communication graph no longer depends on its two agents alone but possibly on all of them. In such a system, for example, two agents might be joined by an edge if and only if fewer than ten percent of them lie in between. We revisit the previous argument and show how to extend it to general bidirectional systems. We retain the ability of the communication graph to freeze when the agents' diameter becomes negligible by enforcing the *agreement rule*: \mathcal{G}_{ij} is constant over the slab $|x_i - x_j| \leq n^{-bn}$, for some suitably large constant b . The difficulty with nonmetrical dynamics is that, though decoupled, subsystems are no longer independent, so in (15) the direct sum $\mathcal{T}_{w_k} \oplus \mathcal{T}_{n-w_k}$ is no longer operative.

We set $\Delta = n^{-b}\mathbb{I}$ and fix $\mathbf{x} \in \Omega^n$ for the time being. This induces a length on each edge of any communication graph $\mathcal{G}(f^t(\mathbf{x}))$, so we can call a node v of \mathcal{T}_n^Δ *heavy* if its communication graph contains one or more edges of length at least n^{-2bn} . The number of times the communication graph has at least one edge of length λ or more is called the *communication count* C_λ : it has been shown, using the *total s-energy* [17], that $C_\lambda \leq \lambda^{-1} \rho^{-O(n)}$, where ρ is the smallest nonzero entry in the stochastic matrices; here $\rho \geq n^{-O(1)}$. It follows that, along any path of \mathcal{T}_n , the number of heavy nodes is $n^{O(n)}$. Let us follow one such path and let G^k denote the communication graph common to the subpath between the k -th and $(k+1)$ -st heavy nodes. To see why that graph is unique,

suppose two consecutive light-node graphs are different. Then some (i, j) is an edge of one but not the other. But, since the first graph only has edges of length less than n^{-2bn} , the locations of both i and j cannot vary by more than n^{-2bn} between the two graphs. It means that in both graphs their distance cannot exceed $3n^{-2bn} < n^{-bn}$; therefore, by the agreement rule, (i, j) is an edge in both graphs, which is a contradiction. We rewrite (12), for fixed \mathbf{x} , as

$$\mathcal{T}_n^\Delta \implies \left\{ \bigotimes_{k=1}^{n^{O(n)}} (\mathcal{T}_{|G^k}^\Delta \otimes \mathcal{T}_n^{\Delta|1}) \right\} \otimes \mathcal{T}_{|G^\infty}^\Delta,$$

where G^∞ is the final graph, which forms an infinite suffix of the graph sequence $\mathcal{G}(f^t(\mathbf{x}))_{t \geq 0}$. We reduce unnecessary branching as follows: whenever V_v (which, with \mathbf{x} fixed, is an interval along the δ -axis) is split into two or more cells by the switching partition, we give it two or more children (besides vanishing leaves) only if at least one of these cells corresponds to a heavy node. The reasoning is that, in the absence of heavy nodes, splitting U_v into subcells is pointless since the communication graphs of all the children are the same; so we might as well give v a single child and, if need be, a vanishing leaf. This ensures that the nesting time of $\mathcal{T}_{|G^k}^\Delta$ is 0. By Lemma 5.1, $\theta_\alpha(\mathcal{T}_{|G^k}^\Delta) \leq n^{O(n)} \log \frac{1}{\alpha}$, and, by (16),

$$\mu_\alpha(\mathcal{T}_n^\Delta) \leq n^{O(n)} + \mu_{\alpha_a}(\mathcal{T}_{|G^\infty}^\Delta) + \sum_{k=1}^{n^{O(n)}} \mu_{\alpha_a}(\mathcal{T}_{|G^k}^\Delta) + \max_k \{ \mu_\alpha(\mathcal{T}_{|G^k}^\Delta), \mu_\alpha(\mathcal{T}_{|G^\infty}^\Delta) \} \leq n^{O(n)} \log \frac{1}{\varepsilon \alpha}.$$

Since $\theta_\alpha(\mathcal{T}_n^\Delta) \leq n^{O(n)} \log \frac{1}{\varepsilon \alpha}$ and $h(\mathcal{T}_{|G^k}^\Delta) = 0$, by (9), inequality (19) becomes

$$h(T_1 \otimes T_2) \leq h(T_1) + h(T_2) + (n+1) \log \log \frac{1}{\varepsilon} + O(n^2 \log n);$$

therefore, $h(\mathcal{T}_n^\Delta) \leq n^{O(n)} \log \log \frac{1}{\varepsilon}$. Repeating the argument we used for the metrical case implies that the vanishing probability of \mathbf{x} is at most

$$\varepsilon (\log \frac{1}{\varepsilon})^{n^{O(n)}} < \sqrt{\varepsilon},$$

for $\varepsilon < 2^{-n^c}$ and constant c large enough. The attraction rate is at most $n^{O(n)} \log \frac{1}{\alpha}$, for any $\alpha < \varepsilon$, and the proof of the bidirectional case of Theorem 1.1 is complete. \square

6 General Influence Systems

We prove Theorem 1.1. The centerpiece of our proof is the bifurcation analysis of a certain non-Markovian extension of an influence system. We focus on that extension first and then show how it relates to the original system. We impose a timeout mechanism to prevent any edge from reappearing after an absence of t_o consecutive steps, for arbitrarily

large t_o . Fix a directed graph H with n nodes labeled 1 through n . Given any $\mathbf{x} \in \Omega^n$, as soon as either the communication graph $\mathcal{G}(f^t(\mathbf{x}))$ contains an edge not in H or some edge of H fails to appear in at least one of $\mathcal{G}(f^{t-t_o+1}(\mathbf{x})), \dots, \mathcal{G}(f^t(\mathbf{x}))$ for some $t \geq t_o$, set all future communication graphs to be the trivial graph consisting of n self-loops. This creates a new coding tree, still denoted \mathcal{T}_n for convenience, which has special *switching* leaves associated with the trivial communication graph. We show that, almost surely, the orbit of any point is attracted to a limit cycle or its path in the coding tree reaches a switching leaf.²⁰ As in the bidirectional case, we assume the agreement rule, which sets \mathcal{G}_{ij} to a constant function over the thin slab $|x_i - x_j| \leq n^{-bn}$.

What is H ? Any infinite graph sequence such as $\mathcal{G}(\mathbf{x}), \mathcal{G}(f(\mathbf{x})), \mathcal{G}(f^2(\mathbf{x}))$, etc, defines a unique *persistent graph*, which consists of all the edges that appear infinitely often. The timeout mechanism allows an equivalent characterization, which includes the edges appearing at least once every t_o steps. The persistent graph depends on the initial state and is unknown ahead of time, so our analysis must handle all possible such graphs. While it plays a key role in the analysis, it would be wrong to think of the persistent graph as determining the dynamics: influence systems can be chaotic and nontrivially periodic, two behaviors that can never be found in systems based on a single graph.

Consider the directed graph derived from H by identifying each strongly connected component with a single node. Let B_1, \dots, B_r be the components whose corresponding nodes are sinks and let n_i denote the number of agents in the group B_i ; write $n = m + n_1 + \dots + n_r$. (In Markov chain terminology, B_i is a closed communicating class.) The linear subspace spanned by the agents of each B_i is forward-invariant and, as we shall see, the phase space evolves toward a subspace of rank r . We reserve the indices $1, \dots, m$ to denote the agents outside of the B_i 's. Unless they hit a vanishing or switching node, the agents indexed $m + 1, \dots, n$ are expected to settle eventually, while the other agents orbit around them, being attracted to a limit cycle. We shall see that nontrivial periodicity is possible only if $r > 1$. We are left with a block-directional system with m (resp. $n - m$) A -agents (resp. B -agents), and the former exercising no influence on the latter (§4.4). It follows from (11) that, for each node v of the global coding tree $\mathcal{T}_{m \rightarrow n-m}$,

$$P_{\leq v} = \begin{pmatrix} A_{\leq v} & C_v \\ 0 & B_{\leq v} \end{pmatrix}. \quad (22)$$

To resolve the system requires a fairly subtle bifurcation analysis which, for convenience, we break down into four stages: in §6.1 we bound the thinning rate; in §6.2 we argue that, deep enough in the coding tree, perturbations keep the coding tree's expected (mean) degree below 1; in §6.3, we show how perturbed phase tubes avoid being split by SP discontinuities at large depths; finally, in §6.4, we show to remove the switching leaves and do away with the persistent graph assumption. We also explain why it is legitimate to ignore the non-Markovian aspect of the system in most of the discussion.

²⁰ Vanishing leaves and switching leaves are distinct: the former “cover” the chaotic regions of the system and are the places perturbations help us avoid; the switching leaves, on the other hand, represent a change in dynamics type and plug into the roots of other coding trees.

6.1 The thinning rate

We prove that, as the depth of a node v of the global coding tree grows, $A_{\leq v}$ and $B_{\leq v}$ tend to matrices of rank 0 and rank r , respectively, with the thinning rates γ and γ' telling us how quickly.

LEMMA 6.1. *Given a node v of $\mathcal{T}_{m \rightarrow n-m}$, there exist vectors $\mathbf{z}_i \in \mathbb{R}^{n_i}$ ($i = 1, \dots, r$), such that, for any $t_v \geq t_c = n^{c n t_o}$ and a large enough constant c ,*

$$(i) \|A_{\leq v} \mathbf{1}_m\|_\infty \leq e^{-\gamma t_v} \quad \text{and} \quad (ii) \left\| B_{\leq v} - \text{diag}(\mathbf{1}_{n_1} \mathbf{z}_1^T, \dots, \mathbf{1}_{n_r} \mathbf{z}_r^T) \right\|_{\max} \leq e^{-\gamma' t_v},$$

where $\gamma = 1/t_c$ and $\gamma' = n^{-cn}$.

Proof. We begin with (i). Consider the initial state $\mathbf{x} = (\mathbf{1}_m, \mathbf{0}_{n-m})$, with all the A -agents at 1 and the B -agents at 0, and let $\mathbf{y} = P_{\leq v} \mathbf{x}$; obviously, $\|A_{\leq v} \mathbf{1}_m\|_\infty = \|\mathbf{y}\|_\infty$. To bound the ℓ_∞ -norm of \mathbf{y} , we apply to \mathbf{x} the sequence of maps specified along the path of $\mathcal{T}_{m \rightarrow n-m}$ from the root to v .²¹ Referring to the arborator (12), let's analyze the factor

$$\underline{\mathcal{T}_{w_k \rightarrow n-w_k}}_{|t_{k+1}-t_k-1} \otimes \mathcal{T}_n^1.$$

The wait period $t_{k+1} - t_k$ before wetness propagates again at time t_{k+1} is at most t_o : indeed, by definition, any A -agent can reach some B -agent in H via a directed path, so all of them will eventually get wet. It follows that the set W_k cannot fail to grow in t_o steps unless it already contains all n nodes or the trajectory reaches a switching leaf. Assume that the agents of W_{t_k+1} , the wet agents at time $t_k + 1$ lie in $(0, 1 - \sigma]$. Because their distance to 1 can decrease by at most a polynomial factor at each step, they all lie in $(0, 1 - \sigma n^{-O(t_o)}]$ between times t_k and t_{k+1} . The agents newly wet at time $t_{k+1} + 1$, ie, those in $W_{t_{k+1}+1} \setminus W_{t_{k+1}}$, move to a weighted average of up to n numbers in $(0, 1)$, at least one of which is in $(0, 1 - \sigma n^{-O(t_o)}]$. This implies that the agents of $W_{t_{k+1}+1}$ lie in $(0, 1 - \sigma n^{-O(t_o)}]$. Since $\sigma \leq 1$, when all the A -agents are wet, which happens within nt_o steps, their positions are confined within $(0, 1 - n^{-O(nt_o)}]$. It follows that

$$\|\mathbf{y}\|_\infty \leq e^{-\lfloor t_v / (nt_o) \rfloor n^{-O(nt_o)}},$$

which proves (i). We establish (ii) along similar lines. Although B_i and B_j ($i \neq j$) are decoupled, they are not independent; so their joint coding tree cannot be expressed as a direct sum. The subgraph $H|_{B_i}$ of H induced by the agents of any given B_i is strongly connected, so viewed as a separate subsystem, the B -agents are newly wetted at least once every nt_o steps. By repeating the following argument for each B_i , we can assume, for the purposes of this proof, that $B = B_1$, $n_1 = n - m$ and $r = 1$.

Initially, place B -agent j at 1 and all the others at 0; then apply to it the sequence of maps leading to $B_{\leq v}$ (this may not be the actual trajectory of that initial state). The previous argument shows that the entries of the j -th column of $B_{\leq v}$, which denote the

²¹ The path need not track the orbit of \mathbf{x} .

locations of the agents at time t_v , are confined to an interval of length $e^{-\lfloor t_v/(nt_o) \rfloor n^{-O(nt_o)}}$. By the agreement rule, this implies that the communication subgraph among the B -agents must freeze at some time $t_c = n^{cnt_o}$ for a constant c large enough, hence become $H|_B$.²² Let $\{u_i\}$ be the $n^{O(nt_c)}$ nodes of the coding tree at depth t_c . Any deeper node v is such that $B_{\leq v} = Q^{t_v - t_{u_i}} B_{\leq u_i}$ for some i , where Q is the stochastic matrix associated with $H|_B$. Since that graph is strongly connected, the previous argument shows that the entries in column j of Q^k lies in an interval of length $e^{-kn^{-O(n)}}$. Since Q^{k+1} is derived from Q^k by taking convex combinations of the rows of Q^k , as k grows, these intervals are nested downwards and hence converge to a number z_j . It follows that Q^k tends to $\mathbf{1}_{n_1} \mathbf{z}^T$, with $\|Q^k - \mathbf{1}_{n_1} \mathbf{z}^T\|_{\max} \leq e^{-kn^{-O(n)}}$. Doubling the value of t_c yields part (ii) of the lemma. \square

The proof suggests that, for any node v deep enough in the coding tree, the matrix $A_{\leq v}$ becomes an error term while $B_{\leq v}$ tends to a matrix that depends only on the ancestor of v of depth t_c . The bifurcation analysis requires a deeper understanding of the error term and calls for more sophisticated arguments. We state the thinning bound in terms of the global coding tree for the perturbation interval $\mathbb{I} = [-1, 1]$.

LEMMA 6.2. *Any node v of $\mathcal{T}_{m \rightarrow n-m}^{\mathbb{I}}$ of depth $t_v \geq t_c$ has an ancestor u of depth t_c such that*

$$\left\| P_{\leq v} - \begin{pmatrix} 0 & C_v \\ 0 & D_u \end{pmatrix} \right\|_{\max} \leq e^{-\gamma t_v},$$

where D_u is a stochastic matrix of the form $D_u = \text{diag}(\mathbf{1}_{n_1} \mathbf{z}_1(u)^T, \dots, \mathbf{1}_{n_r} \mathbf{z}_r(u)^T)$.

6.2 Sparse branching

If we look deep enough in the coding tree for the thinning rate to “kick in,” we observe that, under random margin perturbation, the average branching factor is less than two. Bruin and Deane observed a similar phenomenon in single-agent contractive systems [9]. Their elegant dimensionality argument does not seem applicable in our case, so we follow a different approach, based on geometric considerations. We begin with some terminology: $\text{Lin}[x_1, \dots, x_n]$ refers to a real linear form over x_1, \dots, x_n , with $\text{Aff}[x_1, \dots, x_n]$ designating the affine version; in neither case may the coefficients depend on δ or on the agent positions.²³ With y_1, \dots, y_r understood, a *gap* of type ω denotes an interval of the form $a + \omega \mathbb{I}$, where $a = \text{Aff}[y_1, \dots, y_r]$. We define the set

$$\mathbb{C}[y_1, \dots, y_r] = \left\{ (\xi, \overbrace{y_1, \dots, y_1}^{n_1}, \dots, \overbrace{y_r, \dots, y_r}^{n_r}) \mid \xi \in \Omega^m \right\}.$$

²² We emphasize that we are making no heuristic *assumption* about the repeated occurrence of the edges of H : switching leaves are there precisely to allow violations of the rule.

²³ For example, we can express $y = \delta + x_1 - 2x_2$ as $y = \delta + \text{Lin}[x_1, x_2]$ and $y = \delta + x_1 - 2x_2 + 5$ as $y = \delta + \text{Aff}[x_1, x_2]$.

The variables y_1, \dots, y_r denote the limit positions of the B -agents: they are linear combinations of their initial positions x_{m+1}, \dots, x_n (but functions of the full initial state \mathbf{x}). Let v be a node of the global coding tree $\mathcal{T}_{m \rightarrow n-m}^{\mathbb{I}}$. The matrix $P_{\leq v}$ is a product of the form $P_{t_v} \cdots P_0$, with $P_0 = \text{Id}$ and P_0, \dots, P_{t_v} forming what we call a *valid matrix sequence*. Fix a parameter $\rho > 0$ and a point \mathbf{x} in \mathbb{R}^n . The phase tube formed by the cube $\mathbf{B} = \mathbf{x} + \rho \mathbb{I}^n$ and the matrix sequence P_0, \dots, P_{t_v} consists of the cells $P_0 \mathbf{B}, \dots, (P_{t_v} \cdots P_0) \mathbf{B}$. It might not track any orbit from \mathbf{B} and hence have little relation with a phase tube of the actual system. The phase tube *splits* at node v if the global margin \mathcal{R}_ε defined in (5) makes $(P_k \cdots P_0 \mathbf{B}) \setminus \mathcal{R}_\varepsilon$ disconnected. The following result is the key to sparse branching:

LEMMA 6.3. *Fix $\varepsilon, \rho > 0$, $D_0 \geq 2^{(1/\gamma)^{n+1}}$, and $(y_1, \dots, y_r) \in \mathbb{R}^r$, where $\gamma = n^{-cnt_0}$. There exists a union W of $n^{O(nD_0)}$ gaps of type $(\varepsilon + \rho)n^{O(n^5 D_0)}$ such that, for any interval $\Delta \subseteq \mathbb{I} \setminus W$ of length ρ and any $\mathbf{x} \in \mathbb{C}[y_1, \dots, y_r]$, the phase tube formed by the box $\mathbf{x} + \rho \mathbb{I}^n$ along any path of $\mathcal{T}_{m \rightarrow n-m}^\Delta$ of length at most D_0 cannot split at more than $D_0^{1-\gamma^{n+1}}$ nodes.²⁴*

Proof. We begin with a technical lemma which we prove later. For $k = 0, \dots, D$, let a_k be a row vector in \mathbb{R}^m with $O(\log n)$ -bit rational coordinates and A_k be an m -by- m nonnegative matrix whose entries are rationals over $O(\log N)$ bits, for $N > n$.²⁵ Write $v_k = a_k A_k \cdots A_0$, with $A_0 = \text{Id}$, and assume that the maximum row-sum $\alpha = \max_{k>0} \|A_k \mathbf{1}\|_\infty$ satisfies $0 < \alpha < 1$. Given $I \subseteq \{0, \dots, D\}$, denote by V_I the matrix whose rows are, from top to bottom, the row vectors v_k with the indices $k \in I$ sorted in increasing order. The following result is an elimination device meant to factor out the role of the A -agents. It is a type of matrix rigidity statement.

LEMMA 6.4. *Given any integer $D \geq 2^{(1/\beta)^{m+1}}$ and $I \subseteq \{0, \dots, D\}$ of size $|I| \geq D^{1-\beta^{m+1}}$, where $\beta = |\log \alpha| / (cm^3 \log N)$ for a constant c large enough, there exists a unit vector u such that*

$$u^T V_I = \mathbf{0} \quad \text{and} \quad u^T \mathbf{1} \geq N^{-cm^3 D}.$$

Although c is unrelated to its namesake in Lemma 6.1, we use the same constant by picking the larger of the two; in general, such constants are implied by the bit complexity of the transition matrices and the SP discontinuities. Note also that $\alpha \geq N^{-O(1)}$, so β can be assumed to be much less than 1. To prove Lemma 6.3, we first consider the case where the splitting nodes are well separated, which allows for Lemma 6.1 to be used; then we extend this result to all cases. Given a valid matrix sequence P_0, \dots, P_{D_0} , choose $D \geq 2^{(1/\beta)^{m+1}}$ and pick a sequence of $D + 1$ integers $0 = s_0 < \cdots < s_D \leq D_0$ such that

$$D \geq 2^{(1/\beta)^{m+1}} \quad \text{and} \quad 1/\gamma \leq s_k - s_{k-1} \leq 3/\gamma, \quad (23)$$

²⁴ The crux of the lemma is the *uniformity* over \mathbf{x} : only (y_1, \dots, y_r) needs to be fixed.

²⁵ The coefficients a_k express the discontinuities. Being extracted from the product of several transition matrices, A_k requires more bits.

for $k = 1, \dots, D$: we identify the matrix A_k of Lemma 6.4 with the m -by- m upper left principal submatrix of $P_{s_k} P_{s_k-1} \cdots P_{s_k-1+1}$; using the notation of (22), $A_k = A_{\leq w}$, for some node w (not necessarily an ancestor of v) of depth $t_w = s_k - s_{k-1} \geq 1/\gamma$. Thus, by Lemma 6.1, for $k > 0$, the maximum row-sum of any A_k satisfies $\alpha \leq 1/e$: each A_k is a submatrix of a product of at most $3/\gamma$ transition matrices, so each entry is an $O(\log N)$ -bit rational, with $N = n^{n^2/\gamma}$. What is the row vector a_k ? For $k = 0, \dots, D$, pick any one of the $n^{O(1)}$ margin slabs and denote by a_k the m -dimensional vector of $O(\log n)$ -bit rational coefficients indexed by the A -agents.²⁶ Fix $\delta \in \mathbb{I}$ and pick I in Lemma 6.4 to be of size $\lceil D^{1-\beta^{m+1}} \rceil$. Assume that, given $\mathbf{x} \in \mathbb{C}[y_1, \dots, y_r]$, the phase tube formed by the box $\mathbf{x} + \rho \mathbb{I}^n$ and P_0, \dots, P_{s_D} splits at every index of I along the chosen slabs.²⁷ In other words, for each $k \in I$, there exist a node z_k of depth $t_{z_k} = s_k$ and $\rho_i = \rho_i(k)$, for $i = 1, \dots, n$, such that $|\rho_i| \leq \rho$ and

$$\left| c_k + (a_k, b_k) \begin{pmatrix} A_{\leq z_k} & C_{z_k} \\ 0 & B_{\leq z_k} \end{pmatrix} (x_1 + \rho_1, \dots, x_n + \rho_n)^T + \delta \right| \leq \varepsilon,$$

where the chosen slab is of the form $|c_k + a_k(x_1, \dots, x_m)^T + b_k(x_{m+1}, \dots, x_n)^T + \delta| \leq \varepsilon$, with $b_k \in \mathbb{R}^{n-m}$ and $c_k \in \mathbb{R}$. Since $v_k = a_k A_{\leq z_k}$ and $\mathbf{x} \in \mathbb{C}[y_1, \dots, y_r]$, it follows that

$$\left| v_k(x_1 + \rho_1, \dots, x_m + \rho_m)^T + \text{Aff}[y_1, \dots, y_r, \rho_{m+1}, \dots, \rho_n] + \delta \right| \leq \varepsilon, \quad (24)$$

where the coefficients in the affine form are of magnitude $n^{O(1)}$.




Figure 7: The choice of slabs at the nodes causes the phase tube to split at the nodes indexed by $I = \{2, 4, 6\}$. The nodes of depth s_k for $k \notin I$ are represented as black dots: s_0, s_1, s_3, s_5, s_7 ($D = 7$). The other nodes in the paths are the white dots.

²⁶ With $m = 3$, $x_1 - x_3 + \delta = 0.2$ gives $a_k = (1, 0, -1)$ and $x_2 - x_4 + \delta = 0.7$ produces $a_k = (0, 1, 0)$.

²⁷ It is immaterial that $\mathbf{x} + \rho \mathbb{I}^n$ might slightly bulge out of the phase space Ω^n .

Lemma 6.4 allows us to eliminate the variables x_1, \dots, x_m : we premultiply $V|_I$ by the unit vector u to find that

$$|\text{Aff}[y_1, \dots, y_r] + \delta| \leq (\varepsilon + \rho)N^{O(cm^3D)}, \quad (25)$$

where the coefficients of the new affine form are bounded by $N^{O(cm^3D)}$. (We leave the constant c in the exponent to highlight its influence.) The remarkable fact is that the variable δ is assured not to vanish during the elimination. Thus, as long as δ remains outside a gap of type $(\varepsilon + \rho)N^{O(cm^3D)}$, the phase tube formed by $\mathbf{x} + \rho\mathbb{I}^n$ and P_0, \dots, P_D cannot split at every index of I . Counting the number of possible choices of slabs per node raises the number of gaps to $n^{O(|I|)}$. The argument assumes that δ has the same value in each of $|I|$ inequalities. It need not be so: each δ in (24) can be replaced by $\delta + \nu_k$ ($k \in I$), for $|\nu_k| \leq \rho$, and the new system of inequalities will still imply (25).²⁸ A combinatorial argument shows how adding more gaps to the “exclusion zone” keeps branching low. Before proceeding with that final part of the proof, we summarize our results, using the bound $|\log \alpha| \geq \log e > 1$.

LEMMA 6.5. *Let $N = n^{n^2/\gamma}$ and $\beta = 1/(cm^3 \log N)$, where c is the constant of Lemma 6.4. Fix a path in $\mathcal{T}_{m \rightarrow n-m}^{\mathbb{I}}$ from the root and pick $D + 1$ nodes on it of depth $0 = s_0 < \dots < s_D$ satisfying (23); out of these nodes, choose a subset I of size $\lceil D^{1-\beta^{m+1}} \rceil$. There exists an exclusion zone W consisting of the union of at most $n^{O(|I|)}$ gaps of type $(\varepsilon + \rho)N^{O(cm^3D)}$, such that, for any interval $\Delta \subseteq \mathbb{I} \setminus W$ of length ρ and any $\mathbf{x} \in \mathbb{C}[y_1, \dots, y_r]$, the phase tube formed by $\mathbf{x} + \rho\mathbb{I}^n$ cannot split at all the nodes of I in $\mathcal{T}_{m \rightarrow n-m}^{\Delta}$ (assuming they exist).*

The crux of the lemma is that it holds uniformly for all \mathbf{x} . To prove Lemma 6.3, we need to extend the previous lemma to all the paths of the coding tree of the prescribed length and remove from (23) the lower bound of $1/\gamma$ on the distance between consecutive splitting nodes. Fix $D_0 \geq 2^{(1/\gamma)^{n+1}}$, and let v be a node of $\mathcal{T}_{m \rightarrow n-m}^{\mathbb{I}}$ of depth $t_v = D_0$. Since the path is fixed, we can uniquely identify the node v and its ancestors by their depths and denote by P_t the transition matrix of the node at depth t . Define the node set $J = \{1/\gamma, 2/\gamma, \dots, D_0\}$, with $|J| = \lceil \gamma D_0 \rceil$; recall that $1/\gamma = t_c$ is an integer. Let K be the set of ancestors of v at which the phase tube formed by $\mathbf{x} + \rho\mathbb{I}^n$ and P_0, \dots, P_{D_0} splits (with respect to $\mathcal{T}_{m \rightarrow n-m}^{\mathbb{I}}$); assume that

$$|K| \geq D_0^{1-\gamma^{n+1}}. \quad (26)$$

We define I to be the largest subset of K with no two elements of $I \cup \{0\}$ at a distance less than $1/\gamma$; obviously, $|I| \geq \lfloor \gamma|K| \rfloor - 1$. To define s_1, \dots, s_D , we add all of J to I (to keep distances between consecutive nodes small enough) and then clean up the set to avoid distances lower than allowed: we define J' to be the smallest subset of J such that $L = I \cup (J \setminus J')$ contains no two elements at a distance less than $1/\gamma$. Each element of

²⁸ This observation is crucial for the degree structure analysis to come next and the need to randomize δ .

I can cause the disappearance of at most two elements in J for the addition of one into L , hence $|J|/2 \leq |L| \leq \gamma D_0 + 1$. By construction, consecutive elements of L are at most $3/\gamma$ away from each other, so we can identify L with the sequence $s_1 < \dots < s_D$. By $m < n$ and the specifications of γ in Lemma 6.1 and N, β in Lemma 6.5, we can verify that

$$(i) \ D_0 \geq 2^{(1/\gamma)^{n+1}} \geq \gamma^{-1} 2^{(1/\beta)^{m+1}+1} \quad \text{and} \quad (ii) \ D_0^{1-\gamma^{n+1}} \geq \frac{2}{\gamma} (\gamma D_0 + 1)^{1-\beta^{m+1}}. \quad (27)$$

Part (i) ensures (23). By Lemma 6.5, keeping δ outside the union W of at most $n^{O(|I|)}$ gaps of type $(\varepsilon + \rho)N^{O(m^3 D)}$ prevents I from witnessing a phase tube split at each of its nodes, and hence keeps $K \supseteq I$ from being, as claimed, made entirely of ‘‘splitting’’ nodes. For this, we need to ensure that $|I| \geq D^{1-\beta^{m+1}}$, which follows from: (26); $|I| \geq \lfloor \gamma |K| \rfloor - 1$; $D = |L| \leq \gamma D_0 + 1$; and part (ii) of (27).

We conclude that, as long as we choose an interval $\Delta \subseteq \mathbb{I} \setminus W$ of length ρ , the coding tree $\mathcal{T}_{m \rightarrow n-m}^\Delta$ cannot witness splits at all of the nodes of K (if they exist—their existence is ensured only in $\mathcal{T}_{m \rightarrow n-m}^\mathbb{I}$) for the phase tube formed by any box $\mathbf{x} + \rho \mathbb{I}^n$, where y_1, \dots, y_r are fixed and $\mathbf{x} \in \mathbb{C}[y_1, \dots, y_r]$. Note the order of the quantifiers: first, we fix the coordinates y_k and the target length D_0 , and we pick a large enough candidate splitting node set K in $\mathcal{T}_{m \rightarrow n-m}^\mathbb{I}$; these choices determine the exclusion zone W ; next, we pick a suitable Δ and then claim an impossibility result for *any* \mathbf{x} in $\mathbb{C}[y_1, \dots, y_r]$. To complete the proof of Lemma 6.3, we bound, by 2^{D_0} and $n^{O(nD_0)}$ respectively, the number of ways of choosing K (hence I, L) and the number of nodes v in $\mathcal{T}_{m \rightarrow n-m}^\mathbb{I}$ of depth $t_v = D_0$. \square

Proof of Lemma 6.4. We can make the assumption that I includes 0, since all cases easily reduce to it. Indeed, let l be the smallest index in I . If $l > 0$, subtract l from the indices of I to define $I' \supseteq \{0\}$. Form the matrix $V_{|I}'$ of vectors v'_k , where $v_{k+l} = v'_k A_l \dots A_0$. Rewriting $V_{|I}$ as $V_{|I}' A_l \dots A_0$ takes us to the desired case: we (cosmetically) duplicate the last matrix, P_D , l times to match the lemma’s assumptions and observe that, if $u^T V_{|I}' = \mathbf{0}$, then so does $u^T V_{|I}$. We may also assume that all v_k are nonzero since the lemma is trivial otherwise. All the coordinates of v_k can be expressed as $O(m^2(k+1) \log N)$ -bit rationals sharing a common denominator; therefore,

$$N^{-O((k+1)m^2)} \leq \|v_k\|_1 \leq 2^{-k|\log \alpha| + O(\log n)}. \quad (28)$$

The *affine hull* of $V_{|I}$ is the flat defined by $\{z^T V_{|I} : z^T \mathbf{1} = 1\}$: its dimension is called the *affine rank* of $V_{|I}$. Let $g(D, r)$ be the maximum value of $|I|$, for $\{0\} \subseteq I \subseteq \{0, \dots, D\}$, such that $V_{|I}$ has affine rank at most r and its affine hull does not contain the origin. Lemma 6.4 follows from this inequality, whose proof we postpone: for $r = 0, \dots, m-1$,

$$g(D, r) < D^{1-\beta^{m+1}}, \quad \text{for any } D \geq 2^{(1/\beta)^{m+1}}, \quad (29)$$

where $\beta = |\log \alpha| / (cm^3 \log N)$, for constant c large enough. Indeed, given any $\{0\} \subseteq I \subseteq \{0, \dots, D\}$ of size at least $D^{1-\beta^{m+1}}$, we have $|I| > g(D, m-1)$, so the affine hull of

$V_{|I}$ contains the origin. If r is its affine rank, then there exists $J \subseteq I$ of size $r + 1$ such that the affine rank of $V_{|J}$ is r and its affine hull contains the origin, hence coincides with the row space of $V_{|J}$,²⁹ which is therefore of dimension r . This implies the existence of r independent columns in $V_{|J}$ spanning its column space: add a column of $r + 1$ ones to the right of them to form the $(r + 1)$ -by- $(r + 1)$ matrix M . Since the affine hull of $V_{|J}$ contains the origin, there exists z such that $z^T V_{|J} = \mathbf{0}$ and $z^T \mathbf{1} = 1$, which in turn shows that $\mathbf{1}_{r+1}$ lies outside the column space of $V_{|J}$;³⁰ therefore M is nonsingular. Since each one of its rows consists of $O(m^2 D \log N)$ -bit rationals with a common denominator,

$$|\det M| \geq N^{-O(m^3 D)}. \quad (30)$$

Let ξ be the $(r + 1)$ -dimensional vector whose k -th coordinate is the cofactor of the k -th entry in the last column of ones in M . Simple determinant cofactor expansions show that

$$\xi^T M = (\overbrace{0, \dots, 0}^r, \det M).$$

Since the first r columns of M span the column space of $V_{|J}$, it follows that

$$\xi^T (V_{|J}, \mathbf{1}_{r+1}) = (\overbrace{0, \dots, 0}^m, \det M).$$

By Hadamard's inequality, each coordinate of ξ is at most $n^{O(m)}$ in absolute value, so straightforward rescaling and padding with zeroes turns ξ into a suitable vector u such that $u^T V_{|I} = \mathbf{0}$ and $u^T \mathbf{1} \geq N^{-c_1 m^3 D}$, for an absolute constant c_1 that does not depend on c . Replacing c by $\max\{c, c_1\}$ establishes Lemma 6.4.

It suffices now to prove (29), which we do by induction on r . If $V_{|I}$ has affine rank $r = 0$ and its affine hull does not contain the origin, then all the rows of $V_{|I}$ are equal and nonzero. Since $V_{|I}$ has the row v_0 , it follows from (28) that $|I| \leq 1 + \max\{k \in I\} = O(|\log \alpha|^{-1} m^2 \log N)$, hence

$$g(D, 0) \leq \beta^{-1}. \quad (31)$$

Assume now that $r > 0$ and that $V_{|I}$ has affine rank exactly r . Put $I = \{k_0, k_1, \dots, k_i\}$, with $k_0 = 0$, and consider the smallest j such that $V_{|J}$ has affine rank r , where $J = \{k_0, k_1, \dots, k_j\} \subseteq I$. Since the origin is not in the affine hull of $V_{|I}$ hence of $V_{|J}$, we can always pick a subset $K \subseteq J$ consisting of $r + 1$ independent rows: let $M = V_{|K \cup \{k_i\}}$ denote the $(r + 2)$ -by- m matrix formed by adding the row v_{k_i} at the bottom of $V_{|K}$.³¹ Since $V_{|I}$ has affine rank r , its rank is $r + 1$ (using once again the noninclusion of O in the affine hull of $V_{|I}$), hence so is the rank of M . We show that if k_i is large enough, the system below is feasible in $\xi \in \mathbb{R}^{r+2}$:

$$\xi^T M_+ = (\overbrace{0, \dots, 0}^m, 1), \quad (32)$$

²⁹ Because any $y^T V_{|J}$ can be written as $(y + (1 - y^T \mathbf{1})z)^T V_{|J}$, where $z^T V_{|J} = \mathbf{0}$ and $z^T \mathbf{1} = 1$.

³⁰ Otherwise, $\mathbf{1} = z^T \mathbf{1} = z^T V_{|J} y = 0$.

³¹ It may be the case that $i = j$ or $k_i \in K$. Since $r > 0$, we have $k_i \geq k_j \geq 1$ and $j > 0$.




Figure 8: Why a large value of k_i implies that the affine hull of $V_{|I}$, hence of M , contains the origin.

where M_+ is the $(r+2)$ -by- $(m+1)$ matrix $(M, \mathbf{1}_{r+2})$, which leads to a contradiction. This is the crux of the argument and makes essential use of the rapid decay of the vectors v_k . Assume that $k_i > ck_j |\log \alpha|^{-1} m^3 \log N$, for a large enough constant c . We first show that M_+ is of rank $r+2$. Pick $r+1$ independent columns of $V_{|K}$, which is possible since the latter has rank $r+1$, to form the full-ranked $(r+1)$ -by- $(r+1)$ matrix Q . Add a new row to it by fitting the relevant part of v_{k_i} (the last row of M) and call R the resulting $(r+2)$ -by- $(r+1)$ matrix (Fig. 8); consistent with our notation, R_+ will denote the matrix $(R, \mathbf{1})$. A cofactor expansion of the determinant of R_+ along the bottom row shows that

$$|\det R_+| \geq |\det Q| - \Delta \|v_{k_i}\|_1,$$

where Δ is an upper bound on the absolute values of the cofactors other than $\det Q$. In view of (28), the matrix entries involved in these cofactors are all in $n^{O(1)}$; by Hadamard's inequality, this shows that we can set $\Delta = n^{O(m)}$. Likewise, we find that

$$\|v_{k_i}\|_1 \leq 2^{-k_i} \log \alpha + O(\log n).$$

Since Q is nonsingular, we can adapt (30) to derive $|\det Q| \geq N^{-O(m^3 k_j)}$, hence $|\det R_+| > 0$. It follows that the linear system (32) is feasible if we replace M_+ by R_+ . As it happens, there is no need to do so since every column of M missing from R lies in the column space of the latter: thus the missing homogeneous equalities are automatically satisfied by the solution ξ . The feasibility of (32) contradicts our assumption that the origin is outside the affine hull of $V_{|I}$; therefore

$$k_j \geq \beta k_i > 0, \tag{33}$$

where $\beta = |\log \alpha| / (cm^3 \log N)$. The affine rank of $V_{\{k_0, \dots, k_{j-1}\}}$ is $r-1$ and its affine hull does not contain the origin, so $j \leq g(k_{j-1}, r-1)$, with $g(0, r-1) = 1$. Let $w_0 = a_{k_j}$ and, for $k > 0$, $w_k = a_{k_j+k} A_{k_j+k} \cdots A_{k_j+1}$, thus ensuring that $v_{k_j+k} = w_k A_{k_j} \cdots A_0$. Since the affine hull of $V_{|I}$ does not contain the origin, neither does that of the matrix W with rows $w_0, w_{k_{j+1}-k_j}, \dots, w_{k_i-k_j}$. It follows that the affine rank of W is less than m , so $i-j+1 \leq g(k_i-k_j, m-1)$, hence³² $i \leq g(k_{j-1}, r-1) + g(k_i-k_j, m-1) - 1$. By (33)

³² It would be nice to bound the affine rank as a function of r , but since we never perturb the transition matrices it is unclear how to do that.

and $i = |I| - 1$, we derive, by monotonicity,

$$|I| \leq g(k, r - 1) + g(D - k, m - 1),$$

where $\beta D \leq k \leq D$; hence, by (31), for $m > 0$ and $D \geq 0$:

$$g(D, r) \leq \begin{cases} 1 & \text{if } D = 0 \\ \beta^{-1} & \text{if } r = 0 \\ g(n_1, m - 1) + \cdots + g(n_r, m - 1) + \beta^{-1} & \text{if } 0 < r < m, \end{cases}$$

where $n_1 + \cdots + n_r \leq (1 - \beta^s)D$, with $s = |\{i \mid n_i > 0\}|$. Setting $\eta = \beta^m$, we check that, for all $D > 0$ and $m > 0$,

$$g(D, m - 1) \leq \beta^{-2}(2D^{1-\eta} - 1). \quad (34)$$

The case $m = 1$ follows from $g(D, 0) \leq \beta^{-1}$. For $m > 1$, we begin with the case $s = 0$, where

$$g(D, m - 1) \leq m - 1 + \beta^{-1} \leq \beta^{-2}(2D^{1-\eta} - 1).$$

This follows from $\alpha \geq N^{-O(1)}$, which implies that βm^3 can be made arbitrarily small by increasing c . For $s = 1$,

$$\begin{aligned} g(D, m - 1) &\leq \beta^{-2}(2(1 - \beta)^{1-\eta}D^{1-\eta} - 1) + m - 2 + \beta^{-1} \\ &\leq 2\beta^{-2}D^{1-\eta} - (2\beta^{-1}(1 - \eta) - O(1))D^{1-\eta} - \beta^{-2} + \beta^{-1} + m - 2 \\ &\leq \beta^{-2}(2D^{1-\eta} - 1). \end{aligned}$$

Assume that $s > 1$. Being concave, the function $x \mapsto x^{1-\eta}$ is subadditive for $x \geq 0$; therefore,

$$n_1^{1-\eta} + \cdots + n_r^{1-\eta} \leq (1 - \beta^s)^{1-\eta}D^{1-\eta}.$$

Setting $r = m - 1$, relation (34) follows from the inequality,

$$\begin{aligned} g(D, m - 1) &\leq \beta^{-2}(2(1 - \beta^s)^{1-\eta}D^{1-\eta} - s) + m - s - 1 + \beta^{-1} \\ &\leq 2\beta^{-2}(1 - \beta^{m-1})^{1-\eta}D^{1-\eta} - \frac{3}{2}\beta^{-2} \leq 2\beta^{-2}D^{1-\eta} - \beta^{-2}, \end{aligned}$$

which proves (34), hence (29) and Lemma 6.4. \square

6.3 The degree structure

We decompose the global coding tree into three layers: the top one has no degree constraints; the second has mean degree less than two; and the third has no branching. Consider a placement of the B -agents, such that the diameter of each B_i is less than n^{-bn} . By the agreement rule, the communication subgraph induced by the B -agents is frozen and its transition matrix Q is fixed and independent of the particular placement

of the B -agents.³³ By Perron-Frobenius, or simply by repeating the proof of Lemma 6.1, we derive the existence of a rank- r stochastic matrix

$$\tilde{Q} = \text{diag}(\mathbf{1}_{n_1}\chi_1^T, \dots, \mathbf{1}_{n_r}\chi_r^T)$$

such that $\chi_i \in \mathbf{R}^{n_i}$ and $\|Q^k - \tilde{Q}\|_{\max} \leq e^{-kn^{-O(n)}}$. The B -agents find themselves attracted to the fixed point $\mathbf{y} = \tilde{Q}\xi$, where $\xi \in \mathbb{R}^{n-m}$ is their initial state vector and

$$\mathbf{y} = (\overbrace{y_1, \dots, y_1}^{n_1}, \dots, \overbrace{y_r, \dots, y_r}^{n_r}).$$

Define $\Upsilon = \Omega^m \times (\Omega^{n-m} \cap \Upsilon_B)$, where

$$\Upsilon_B = \mathbf{y} + (n^{-2bn} \mathbb{I}^{n-m}) \cap \ker \tilde{Q}.$$

If $\mathbf{x} \in \Upsilon$, the diameter of any group B_i is at most $2n^{-2bn} < n^{-bn}$ so the communication graph induced by their agents is frozen and remains so. The B -agents are attracted to \mathbf{y} .³⁴ This follows easily, as does the next lemma, whose proof we omit, from the stochasticity of Q and the identities: $\tilde{Q}Q = Q\tilde{Q} = \tilde{Q}^2 = \tilde{Q}$.

LEMMA 6.6. *The set Υ is forward-invariant. Furthermore, any $\xi \in \mathbf{y} + n^{-2bn} \mathbb{I}^{n-m}$ belongs to Υ_B if and only if $\tilde{Q}\xi = \mathbf{y}$.*

We set ε, ρ, D_0 as in Lemma 6.3 and call an interval Δ *free* if it does not intersect the exclusion zone $W = W(\mathbf{y})$. As usual, we choose the perturbation sample space $n^{-b} \mathbb{I}$ to make perturbations inconsequential in practice. For counting purposes, it is convenient to partition $n^{-b} \mathbb{I}$ into *canonical* intervals of length ρ (with possibly a single smaller one). A gap of W can keep only $(1 + \varepsilon/\rho)n^{O(n^5 D_0)}$ canonical intervals from being free, so the Lebesgue measure of the free ones satisfies:

$$\text{Leb} \left\{ \bigcup \text{free canonical intervals} \right\} \geq 2n^{-b} - (\varepsilon + \rho)n^{O(n^5 D_0)}. \quad (35)$$

Fixing the B -agent attractor. With \mathbf{y} fixed, we pick a free canonical interval Δ and focus on the global coding tree $\mathcal{T}_{m \rightarrow n-m}^{\Delta|\Upsilon}$, with the superscripts indicating the perturbation and phase spaces, respectively.³⁵ For any node v of depth $t_v \geq t_c$, the limit matrix D_u in Lemma 6.2 is the same for all nodes u of depth t_c . Indeed,

$$\left\| P_{\leq v} - \begin{pmatrix} 0 & C_v \\ 0 & \tilde{Q} \end{pmatrix} \right\|_{\max} \leq e^{-\gamma t_v}.$$

³³ The system under consideration is the non-Markovian extension defined by the persistent graph H .

³⁴ Although the B -agents in Υ_B have been essentially immobilized around \mathbf{y} , they are not decoupled from the rest. Indeed, while the increasingly microscopic movement of the B -agents can no longer affect their own communication graph, it can still influence the communication among the A -agents: furthermore, this may still be true even if no edge is ever to join an A -agent to a B -agent.

³⁵ The reason we do not fix the perturbation δ is that it needs to be randomized and it is easier to avoid randomizing the coding tree itself.

Pick v of depth $t_v \geq 3t_c$ and let w be its ancestor at depth $t_w = \lfloor t_v/2 \rfloor$. Given $\mathbf{x} \in U_v \subseteq \Upsilon$,

$$\begin{aligned} \mathbf{x}' = f^{t_w}(\mathbf{x}) &= P_{\leq w} \mathbf{x} \in \begin{pmatrix} C_w \\ \tilde{Q} \end{pmatrix} (x_{m+1}, \dots, x_n)^T + ne^{-\gamma t_w} \mathbb{I}^n \\ &\in \begin{pmatrix} C_w(x_{m+1}, \dots, x_n)^T \\ \mathbf{y} \end{pmatrix} + ne^{-\gamma t_w} \mathbb{I}^n. \end{aligned}$$

By Lemma 6.6, $\mathbf{x}' \in \Upsilon$, so there exists a node v' of depth $t_{v'} = t_v - t_w \geq t_c$ such that,

$$f^{t_v}(\mathbf{x}) = f^{t_{v'}}(\mathbf{x}') = P_{\leq v'} \mathbf{x}' \in \begin{pmatrix} C_{v'} \\ \tilde{Q} \end{pmatrix} (\mathbf{y} + ne^{-\gamma t_w} \mathbb{I}^n) + ne^{-\gamma t_{v'}} \mathbb{I}^n \subseteq \begin{pmatrix} C_{v'} \mathbf{y} \\ \mathbf{y} \end{pmatrix} + 2ne^{-\gamma t_v/3} \mathbb{I}^n.$$

It is important to note that v' depends only on v and not on $\mathbf{x} \in U_v$: indeed, the phase tube from U_v between time t_w and t_v does not split; therefore $f^{t_w}(U_v) \subseteq U_{v'}$. It follows that, for $t_v \geq 3t_c$ and $v' = v'(v)$,

$$V_v \subseteq \begin{pmatrix} C_{v'} \mathbf{y} \\ \mathbf{y} \end{pmatrix} + 2ne^{-\gamma t_v/3} \mathbb{I}^n. \quad (36)$$

The A -agents evolve toward convex combinations of the B -agents, which themselves become static. The weights of these combinations (ie, the barycentric coordinates of the A -agents), however, might change at every node, so there is no assurance that the orbit is always attracted to a limit cycle. The layer decomposition of the coding tree, which we describe next, allows us to bound the nesting time while exhibiting weak yet sufficient conditions for periodicity.




Figure 9: The global coding tree is stratified into three layers, with decreasing branching rates.

To stratify the coding tree $\mathcal{T}_{m \rightarrow n-m}^{\Delta|\Upsilon}$ into layers, we set up three parameters D_0 , D_1 , and D_2 : the first targets the topological entropy; the second specifies the height of the first layer; the third indicates the nesting time. We examine each one in turn and indicate their purposes and requirements.

FIRST LAYER. By (36), the phase tubes get thinner over time at a rate of roughly $e^{-\gamma/3}$, while the tree is branching at a rate of $n^{O(n)}$. To ensure that the topological entropy is zero, the product of these two rates should be less than 1: with $\gamma < 1$, this is far from being the case, so we need a sparsification mechanism. This is where Lemma 6.3 comes in. Indeed, deep enough in $\mathcal{T}_{m \rightarrow n-m}^{\Delta|\Upsilon}$, the size of a subtree of height D_0 is at most³⁶

$$D_0(n^{O(n)})D_0^{1-\gamma^{n+1}},$$

while the tubes get thinner at a rate of $2ne^{-\gamma D_0/3}$ for every consecutive D_0 nodes: the choice of D_0 below ensures that the product is less than 1, as desired. We justify this choice formally below.

$$D_0 \geq 2^{(1/\gamma)^{n+2}} \quad [D_0 \text{ big enough for thinning to outpace branching}]. \quad (37)$$

SECOND LAYER. Technically, Lemma 6.3 addresses only the branching of the phase tube formed by a small box $\mathbf{x} + \rho \mathbb{I}^n$, for $\mathbf{x} \in \mathbb{C}[y_1, \dots, y_r]$, whereas we are concerned here with phase tubes originating at some cell V_v of $\mathcal{T}_{m \rightarrow n-m}^{\Delta|\Upsilon}$. To make V_v thin enough, we choose a node v deep in the tree.³⁷ By (36), $V_v \subseteq \mathbf{x} + \rho \mathbb{I}^n$, for $\mathbf{x} \in \mathbb{C}[y_1, \dots, y_r]$, provided that $t_v \geq D_1$ and

$$D_1 \geq \frac{3}{\gamma} \log \frac{2n}{\rho} \quad [D_1 \text{ big enough for tree branches to be thinner than } \rho]. \quad (38)$$

Note that the requirement in (36) that $t_v \geq 3t_c = 3/\gamma$ is implied by $t_v \geq D_1$. In view of Lemma 6.3, the number of nodes in $\mathcal{T}_{m \rightarrow n-m}^{\Delta|\Upsilon}$ of depth no greater than $t \geq D_1$, is bounded by

$$\underbrace{n^{O(nD_1)}}_{\text{depth } D_1} \times \underbrace{n^{O(nD_0^{1-\gamma^{n+1}} \lfloor (t-D_1)/D_0 \rfloor)}}_{\text{from } D_1 \text{ to } t \text{ in chunks of } D_0} \times \underbrace{n^{O(nD_0)}}_{\text{truncated chunk}} \times \underbrace{D_0}_{\text{single paths}};$$

hence, for any $t \geq D_1$,

$$\left| \{v \in \mathcal{T}_{m \rightarrow n-m}^{\Delta|\Upsilon} \mid t_v \leq t\} \right| \leq n^{O(nD_0 + nD_1 + ntD_0^{-\gamma^{n+1}})}. \quad (39)$$

³⁶ This assumes a thinness condition we discuss below. The factor D_0 comes from the nonbranching paths in the subtree spanned by the phase tubes from Υ .

³⁷ Factoring out the B -agents gives us the sort of fixed-point attraction that is required by Lemma 6.3: it is a dimension reduction device in attractor space.

THIRD LAYER. The bottom layer of the stratified global coding tree begins at a depth $D_2 \geq D_0 + D_1$. If the node v of depth $t_v \geq D_2$ has nontrivial branching,³⁸ then, by continuity, V_v contains a point right on the boundary of the global margin. By (36), this implies the existence of $\zeta \in \mathbb{R}^n$ such that $\|\zeta\|_\infty \leq 2ne^{-\gamma D_2/3}$ and $\text{Aff}[\mathbf{y} + \zeta] = \delta$, where the coefficients of the affine form are of magnitude $n^{O(1)}$ and depend only on the node v . It then follows from (39) that $\mathcal{T}_{m \rightarrow n-m}^{\Delta' \Upsilon}$ has no nontrivial branching at depth D_2 , provided that $\Delta' = \Delta \setminus W'$, where W' consists of gaps of type $n^{O(1)}e^{-\gamma D_2/3}$ numbering at most

$$\underbrace{n^{O(nD_0+nD_1+nD_2D_0^{-\gamma^{n+1}})}}_{\# \text{ nodes at depth } D_2} \times \underbrace{n^{O(1)}}_{\# \text{ margin slabs}}.$$

This calculation, in which ε played no role, puts a bound of D_2 on the nesting time. It follows that

$$\text{Leb}(W') \leq e^{-\gamma D_2/3} n^{O(nD_0+nD_1+nD_2D_0^{-\gamma^{n+1}})} \quad (40)$$

Pick an arbitrarily small $\varepsilon_o > 0$ and a large enough constant $d = d(b, c)$; recall that $\gamma = n^{-cnt_o}$. We set the parameters $\rho = \varepsilon_o^2 n^{-dn^5 D_0}$ and $\varepsilon \leq \min\{\rho, e^{-\gamma D_2}\}$, where, rounding up to the nearest integer,

$$\begin{cases} D_0 = 2^{d(1/\gamma)^{n+2}} \\ D_1 = \frac{d^2}{\gamma}(n^6 D_0 + |\log \varepsilon_o|) \\ D_2 = \frac{d}{\gamma}(n^2 D_1 + |\log \varepsilon_o|). \end{cases} \quad (41)$$

We verify that conditions (37, 38) are both satisfied and that

$$D_1 \geq D_2 D_0^{-\gamma^{n+1}}. \quad (42)$$

Thus the measure bound (40) implies that $\text{Leb}(W') \leq \rho 2^{-D_0}$. By making ε tend to 0, the point \mathbf{x} vanishes with arbitrarily small probability for random $\delta \in \Delta'$. By Lemma 4.4, this implies that, with probability at least $1 - 2^{-D_0}$, subjecting the system's margin to a perturbation δ chosen randomly in Δ makes the orbit of any $\mathbf{x} \in \Upsilon$ attracted to a limit cycle (or a switching leaf):³⁹ we call this *success*. The sum of the period and preperiod is bounded by the number of nodes of depth at most D_2 (the nesting time), which, by (39, 42), is no greater, conservatively, than

$$\bar{p} = n^{O(nD_1)} \leq (1/\varepsilon_o)^{O(\gamma^{-2})} 2^{D_0 \gamma^{-1} n^{O(1)}}. \quad (43)$$

We bound the attraction rate by appealing to Lemmas 4.5 and 6.2. Note that if p is the period then so is $p \lceil (\log 2n)/\gamma \rceil$. This choice of p still satisfies the upper bound (43) while ensuring that, at every period, the error bound in Lemma 6.2 is at most $\frac{1}{2n}$. The

³⁸ A node is branching nontrivially if it has at least two children neither of which is switching or vanishing.

³⁹ The regions W and W' , which make perturbation a requirement, depend only \mathbf{y} . But perturbation is also needed to avoid vanishing, which depends on the initial state \mathbf{x} .

row-sums in A in Lemma 4.5 are at most $1/2$, so we can set $\mu = e^{-\gamma}$. Since $\nu \leq D_2$ and $(\#\mathcal{S}_\nu)[(\log 2n)/\gamma] \leq \bar{p}$, it follows that

$$\theta_\alpha \leq D_0^{O(D_0)}(1/\varepsilon_o)^{O(\gamma^{-2})} \log \frac{1}{\alpha}, \quad (44)$$

for any $0 < \alpha < 1/2$. The perturbation space is not Δ but $n^{-b}\mathbb{I}$, so we apply the previous result to each free canonical interval and argue as follows. If Λ is the measure of the union of all the free canonical intervals, then the perturbations that do not guarantee success have measure at most $(2n^{-b} - \Lambda) + 2^{-D_0}\Lambda$. Dividing by $2n^{-b}$ and applying (35) shows that

$$\text{Prob}[\text{failure in } \mathcal{T}_{m \rightarrow n-m}^{n^{-b}\mathbb{I}|\Upsilon}] \leq 1 - (1 - 2^{-D_0})(1 - (\varepsilon + \rho)n^{O(n^5 D_0)}) \leq 2^{1-D_0}. \quad (45)$$

The nesting time is at most D_2 , which, by (39, 42), implies that

$$h(\mathcal{T}_{m \rightarrow n-m}^{n^{-b}\mathbb{I}|\Upsilon}) \leq O(D_1 n \log n) \leq \frac{1}{\gamma}(D_0 + |\log \varepsilon_o|)n^{O(1)}. \quad (46)$$

Freeing the B -agents. Set $D_3 = \lceil 3bt_c n \log n \rceil$ and fix \mathbf{x} in Ω^n . Let ξ denote the projection of $f^{D_3}(\mathbf{x})$ onto the last $n - m$ coordinate axes. By Lemma 6.2 and $t_c = 1/\gamma$ (Lemma 6.1), the coding tree $\mathcal{T}_{m \rightarrow n-m}^{n^{-b}\mathbb{I}}$ has $n^{O(nt_c)}$ nodes u such that $t_u = t_c$ and

$$\xi \in \mathbf{y} + ne^{-\gamma D_3} \mathbb{I}^{n-m} \subseteq \mathbf{y} + n^{-2bn} \mathbb{I}^{n-m},$$

where $\mathbf{y} = D_u(x_{m+1}, \dots, x_n)^T$. The state vector for the B -agents is ξ at time D_3 and $Q^{t-D_3}\xi$ at $t > D_3$, where Q is the transition matrix of the frozen communication subgraph joining the B -agents at time D_3 . By taking t to infinity, it follows that $\mathbf{y} = \tilde{Q}\xi$ and, by Lemma 6.6, $\xi \in \Upsilon_B$ hence $f^{D_3}(\mathbf{x}) \in \Upsilon$. We can then apply the previous result. Since \mathbf{x} is fixed, only the choice of random perturbation δ can change which path in $\mathcal{T}_{m \rightarrow n-m}^{n^{-b}\mathbb{I}}$ the orbit will follow. The failure probability of (45) needs to be multiplied by the number of nodes u , which yields an upper bound of $n^{O(nt_c)}2^{1-D_0}$; hence

$$\text{Prob}[\text{failure in } \mathcal{T}_{m \rightarrow n-m}^{n^{-b}\mathbb{I}}] \leq 2^{-D_0/2}. \quad (47)$$

If \mathcal{T}^* denotes the part of the global coding tree extending to depth D_3 , then

$$\mathcal{T}_{m \rightarrow n-m}^{n^{-b}\mathbb{I}} = \mathcal{T}^* \otimes \mathcal{T}_{m \rightarrow n-m}^{n^{-b}\mathbb{I}|\Upsilon}.$$

The tree \mathcal{T}^* has at most $n^{O(nt_c)}D_3$ nodes; therefore, by (46),

$$h(\mathcal{T}_{m \rightarrow n-m}^{n^{-b}\mathbb{I}}) \leq O(t_c n \log n + \log D_3) + h(\mathcal{T}_{m \rightarrow n-m}^{n^{-b}\mathbb{I}|\Upsilon}) \leq \frac{1}{\gamma}(D_0 + |\log \varepsilon_o|)n^{O(1)}. \quad (48)$$

6.4 Removing persistence

We use direct products to relax the condition that the permanent graph H be fixed once and for all. This touches on the non-Markovian nature of the system, a feature we chose to ignore in the previous section. We explain now why this was legitimate. Because the switching condition is about time differences and not absolute times, any subpath in the coding tree has an incarnation as a path from the root. Equivalently, any interval of a trajectory appears as the prefix of another trajectory. This property explains why, following (23), we could argue that A_k was of the form $A_{\leq w}$. Likewise, in the derivations leading to (36), we used the fact that $f^{t_v}(\mathbf{x}) = f^{t_v - t_w}(f^{t_w}(\mathbf{x}))$, an identity that might not always hold in a non-Markovian setting, but which, in this case, did. Finally, weren't we too quick to appeal to Lemma 4.4 for periodicity since its proof relied heavily on the Markov property? To see why the answer is no, observe that the argument did establish the periodicity of the “wrap-around” system derived from $\mathcal{T}_{m \rightarrow n-m}^{\Delta/\Upsilon}$ by redirecting any trajectory that reaches the nesting depth to the root. The only problem is that this system, being Markovian, is not the one modeled by the coding tree. Wrapping around resets the time to zero, which might cause switching conditions to be missed and trajectories to be continued when they should be stopped: none of this stops nonvanishing orbits from being periodic, however.

We now show how to relax the permanent graph assumption. The idea is to begin with H set as the complete graph and update it at each switching leaf by removing the edge(s) whose missing presence causes the node to be a switching leaf. We then append to each such leaf the coding tree, suitably cropped, defined with respect to the new value of H . We model this iteration by means of direct products, using m_k to denote the number of A -agents in the block-directional system used in the k -th product:

$$\mathcal{T}_n^{n-b \mathbb{I}} \implies \bigotimes_{k=1}^{k_0} \mathcal{T}_{m_k \rightarrow n-m_k}^{n-b \mathbb{I}}. \quad (49)$$

The upper limit k_0 is bounded by $n(n-1)$; note that each decision procedure \mathcal{G}_{ij} needs its own counter. To keep the failure probability small throughout the switching of dynamics, we need to update the value of D_0 in (41) at each iteration, so we define C_k as its suitable value for a persistent graph consisting of k (nonloop) directed edges and let ϕ_k denote the maximum failure probability for such a graph: $C_{n(n-1)} \geq D_0$ and $\phi_0 = 0$. The logarithm of the number of switching leaves is at most the word-entropy;⁴⁰ by (47, 48), for $k > 0$,

$$\phi_k \leq 2^{-C_k/2} + 2^{\gamma^{-1}(C_k + |\log \varepsilon_o|)n^a} \phi_{k-1},$$

for some constant $a > 0$. Setting

$$C_{n(n-1)-j} = \lceil \gamma^{-j} n^{2aj} (D_0 + 3|\log \varepsilon_o|) \rceil,$$

⁴⁰ No two switching leaves can have the same ancestor at a depth equal to the nesting time. Because we can bound the number of switching leaves, we may dispense with (10) altogether.

for $j = 0, \dots, n(n-1)$, we verify by induction that $\phi_k \leq 2^{1-C_k/2}$, for $k = 0, \dots, n(n-1)$; hence,

$$\text{Prob}[\text{failure in nonpersistent } \mathcal{T}_{m \rightarrow n-m}^{n-b\mathbb{I}}] \leq \phi_{n(n-1)} \leq 2^{1-\frac{1}{2}(D_0+3|\log \varepsilon_o|)} \leq \varepsilon_o.$$

The attraction rate is still exponential: using (44) yields a geometric series summing up to

$$\theta_\alpha \leq C_0^{O(C_0)} (1/\varepsilon_o)^{O(\gamma^{-2})} \log \frac{1}{\alpha} \leq O_{n,\varepsilon_o,t_o}(\log \frac{1}{\alpha}), \quad (50)$$

for any $0 < \alpha < 1/2$. By (43), the period and preperiod are bounded by

$$(1/\varepsilon_o)^{O(\gamma^{-2})} 2^{C_0\gamma^{-1}n^{O(1)}} \leq (1/\varepsilon_o)^{O_{n,\varepsilon_o,t_o}(1)},$$

which completes the proof of Theorem 1.1. \square

Acknowledgments

I wish to thank Pascal Koiran and John Tsitsiklis for helpful conversations.

References

- [1] Acemoglu, D., Dahleh, M.A., Lobel, I., Ozdaglar, A. *Bayesian learning in social networks*, LIDS report 2780, 2009.
- [2] Alur, R., Henzinger, T.A., Sontag, E.D. *Hybrid Systems III. Verification and Control*, Springer Verlag, Berlin, LNCS 1066, 1996.
- [3] Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., Lecomte, V., Orlandi, A., Parisi, G., Procaccini, A., Viale, M., Zdravkovic, V. *Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study*, Proc. National Academy of Sciences 105 (2008), 1232-1237.
- [4] Blondel, V.D., Hendrickx, J.M., Olshevsky, A., Tsitsiklis, J.N. *Convergence in multiagent coordination, consensus, and flocking*, Proc. 44th IEEE Conference on Decision and Control, Seville, Spain, 2005.
- [5] Blondel, V.D., Hendrickx, J.M., Tsitsiklis, J.N. *On the 2R conjecture for multi-agent systems*, Proc. European Control Conference 2007 (ECC 2007), July 2007, 874-881, Kos (Greece), 2007.
- [6] Blondel, V.D., Hendrickx, J.M., Tsitsiklis, J.N. *On Krause's multi-agent consensus model with state-dependent connectivity*, IEEE Transactions on Automatic Control 54, 11 (2009), 2586-2597.
- [7] Blondel, V.D., Tsitsiklis, J.N. *A survey of computational complexity results in systems and control*, Automatica 36 (2000), 1249-1274.
- [8] Braverman, M., Grigo, A., Rojas, C. *Noise vs computational intractability in dynamics*, Proc. 3rd ITCS (2012), 128-141.

- [9] Bruin, H., Deane, J.H.B. *Piecewise contractions are asymptotically periodic*, Proc. American Mathematical Society 137, 4 (2009), 1389–1395.
- [10] Bullo, F., Cortés, J., Martinez, S., *Distributed Control of Robotic Networks*, Applied Mathematics Series, Princeton University Press, 2009.
- [11] Buzzi, J. *Piecewise isometries have zero topological entropy*, Ergod. Theory Dyn. Syst. 21 (2001), 13711377. MR1855837 (2002f:37029)
- [12] Cao, M., Spielman, D.A. Morse, A.S. *A lower bound on convergence of a distributed network consensus algorithm*, 44th IEEE Conference on Decision and Control, and the European Control Conference 2005, Seville, Spain, 2005
- [13] Castellano, C., Fortunato, S., Loreto, V. *Statistical physics of social dynamics*, Rev. Mod. Phys. 81 (2009), 591–646.
- [14] Catsigeras, E., Budelli, R. *Topological dynamics of generic piecewise continuous contractive maps in n dimensions*, International Journal of Pure and Applied Mathematics, 68 (2011), 61-83.
- [15] Cessac, B. *A discrete time neural network model with spiking neurons: rigorous results on the spontaneous dynamics*, J. Math. Biol. 56 (2008), 311-345.
- [16] Chazelle, B. *The convergence of bird flocking*, arXiv:0905.4241v1, 2009. Prelim. version in ACM SoCG 2010.
- [17] Chazelle, B. *The total s -energy of a multiagent system*, SIAM J. Control Optim. 49 (2011), 1680–1706.
- [18] Collins, G. E. *Quantifier elimination for real closed fields by cylindrical algebraic decomposition*, Proc. 2nd GI Conf. Automata Theory and Formal Languages, New York, Springer-Verlag (1975), 134–183.
- [19] Condon, A., Lipton, R.J. *On the complexity of space bounded interactive proofs*, Proc. 30th IEEE Symp. on Foundations of Computer Science (1989), 462–267.
- [20] Conradt, L., Roper, T.J. *Group decision making in animals*, Nature 421 (2003), 155–158.
- [21] Cucker, F., Smale, S. *Emergent behavior in flocks*, IEEE Trans. Automatic Control 52 (2007), 852–862.
- [22] Devaney, R.L. *An Introduction to Chaotic Dynamical Systems*, 2nd edition, Westview Press, 2003.
- [23] Dittmer, J.C. *Consensus formation under bounded confidence*, Nonlinear Analysis, Theory, Methods & Applications, 47 (2001), 4615–4622.
- [24] Earl, M.G., Strogatz, S.H. *Synchronization in oscillator networks with delayed coupling: a stability criterion*, Phys. Rev. E, 67 (2003), 036204(1–4).
- [25] Gazi, V., Passino, K.M. *Stability analysis of swarms*, IEEE Transactions on Automatic Control 48 (2003), 692–697.
- [26] Hegselmann, R., Krause, U. *Opinion dynamics and bounded confidence models, analysis, and simulation*, J. Artificial Societies and Social Simulation 5, 3 (2002).

- [27] Hegselmann R, Krause U. *Truth and cognitive division of labor: first steps towards a computer aided social epistemology*, J. Artificial Societies and Social Simulation 9 (2006).
- [28] Hendrickx, J.M., Blondel, V.D. *Convergence of different linear and non-linear Vicsek models*, Proc. 17th International Symposium on Mathematical Theory of Networks and Systems (MTNS2006), Kyoto (Japan), July 2006, 1229–1240.
- [29] Jadbabaie, A., Lin, J., Morse, A.S. *Coordination of groups of mobile autonomous agents using nearest neighbor rules*, IEEE Trans. Automatic Control 48 (2003), 988–1001.
- [30] Koiran, P., Cosnard, M., Garzon, M. *Computability with low-dimensional dynamical systems*, Theoretical Computer Science A 132 (1994), 113–128.
- [31] Krause, U. *A discrete nonlinear and non-autonomous model of consensus formation*, Communications in Difference Equations (2000), 227236.
- [32] Kruglikov, B., Rypdal, M. *A piecewise affine contracting map with positive entropy*, Discrete and Continuous Dynamical Systems 16, 2 (2006), 393–394.
- [33] Kurz, S., Rambau, J. *On the Hegselmann-Krause conjecture in opinion dynamics*, Journal of Difference Equations and Applications (2009).
- [34] Lorenz, J. *Multidimensional opinion dynamics when confidence changes*, Economic Complexity, Aix-en-Provence, May 2003.
- [35] Lorenz, J. *A stabilization theorem for dynamics of continuous opinions*, Physica A: Statistical Mechanics and its Applications 355 (2005), 217–223.
- [36] Lorenz, J. *Bounds of confidence: Meet, discuss and find consensus!*, Complexity, 4 (2010), 43–52.
- [37] Martinez, S., Bullo, F., Cortés, J., Frazzoli, E. *On synchronous robotic networks – Part II: Time complexity of rendezvous and deployment algorithms*, IEEE Transactions on Automatic Control 52 (2007), 2214–2226.
- [38] Mirolo, R.E., Strogatz, S. *Synchronization of pulse-coupled biological oscillators*, SIAM J. Applied Math 50 (1990), 1645–1662.
- [39] Mirtabatabaei A., Bullo, F. *Opinion dynamics in heterogeneous networks: convergence conjectures and theorems*, SIAM Journal on Control and Optimization, March 2011.
- [40] Moreau, L. *Stability of multiagent systems with time-dependent communication links*, IEEE Transactions on Automatic Control 50 (2005), 169–182.
- [41] Nedic, A., Olshevsky, A., Ozdaglar, A., Tsitsiklis, J.N. *On distributed averaging algorithms and quantization effects*, IEEE Transactions on Automatic Control, 2009.
- [42] Niven, I., Zuckerman, H.S., Montgomery, H.L. *An Introduction to the Theory of Numbers*, Wiley, 5th edition, 1991.
- [43] Okubo, A., Levin, S.A. *Diffusion and Ecological Problems*, Springer, 2nd ed, 2002.
- [44] Olshevsky, A., Tsitsiklis, J.N. *Convergence speed in distributed consensus and averaging*, SIAM Journal on Control and Optimization 48 (2009), 33–55.

- [45] Parrish, J.K., Hamner, W.M. *Animal Groups in Three Dimensions*, Cambridge University Press, 1997.
- [46] Papachristodoulou, A., Jadbabaie, A. *Synchronization in oscillator networks: switching topologies and non-homogeneous delays*, Proc. IEEE CDC-ECC (2005), 5692–5697.
- [47] Pikovsky, A., Rosenblum, M., Kurths, J. *Synchronization: A Universal Concept in Nonlinear Sciences*, Cambridge University Press, 2001.
- [48] Scardovi, L., Sarlette, A., Sepulchre, R. *Synchronization and balancing on the N -torus*, Systems & Control Letters 56 (2007), 335–341.
- [49] Seneta, E. *Non-Negative Matrices and Markov Chains*, Springer, 2nd ed., 2006.
- [50] Siegelmann, H.T., Sontag, E.D. *On the computational power of neural nets*, J. Comput. System Sci., 50 (1995), 132–150.
- [51] Sontag, E.D. *Nonlinear regulation: the piecewise linear approach*, IEEE Trans. Automat. Control 26 (1981), 346–358.
- [52] Strogatz, S.H. *From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators*, Physica D 143 (2000), 1–20.
- [53] Tsitsiklis, J.N., Bertsekas, D.P., Athans, M. *Distributed asynchronous deterministic and stochastic gradient optimization algorithms*, IEEE Transactions on Automatic Control 31 (1986), 803–812.
- [54] Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., Shochet, O. *Novel type of phase transition in a system of self-driven particles*, Physical Review Letters 75 (1995), 1226–1229.
- [55] Walker, T.J. *Acoustic synchrony: two mechanisms in the snowy tree cricket*, Science 166 (1969), 891–894.
- [56] Winfree, A.T. *Biological rhythms and the behavior of populations of coupled oscillators*, J. Theoret. Bio. 16, 1 (1967), 15–42.