

# An Algorithmic Approach to Collective Behavior

Bernard Chazelle

Received: 24 February 2014 / Accepted: 11 October 2014 / Published online: 21 October 2014  
© Springer Science+Business Media New York 2014

**Abstract** The emergence of collective structure from the decentralized interaction of autonomous agents remains, with notable exceptions, a mystery. While powerful tools from dynamics and statistical mechanics have been brought to bear, sometimes with great success, an algorithmic perspective has been lacking. Viewing collective behavior through the lens of *natural algorithms* offers potential benefits. This article examines the merits and challenges of an algorithmic approach to the emergence of collective order.

**Keywords** Collective behavior · Natural algorithms · Influence systems · s-Energy · Dynamic networks · Renormalization

## 1 Introduction

How much, if any, does the study of collective behavior stand to gain from a focus on its algorithmic nature? Leaving the precise meaning of “collective behavior” aside for now, what we have in mind is a group of agents communicating with one another and taking action on the basis of what they “hear.” These could be computers spreading a virus, neurons responding to stimuli, birds flocking along a migration route, fish swarming into a baitball, bacteria engaged in quorum sensing, or fireflies synchronizing their flashes. In all cases, the agents are equipped with their own (possibly distinct) procedures to help them decide what to do under what conditions and, crucially, whom to “listen” to. Out of a plethora of local interactions, striking patterns will emerge: birds will form V shapes; fireflies will flash in synchrony; ants will find shortest paths; and brains will become minds. How do we study emergent self-organization of this sort? Is each case its own story or can universal principles arise from the bewildering diversity on offer?

Given a complex multiagent system, the first challenge is to find an effective *natural algorithm* for it, which is to say, a model rich enough to track reality yet amenable to analysis—

---

B. Chazelle (✉)  
Department of Computer Science, Princeton University, Princeton, NJ, USA  
e-mail: chazelle@cs.princeton.edu

**Fig. 1** Once we know the birds' natural algorithm, then what?




the second one is to figure out what to do with it (Fig. 1). The two challenges are related, yet distinct. We focus on the latter. Our approach is premised on the hypothesis that natural algorithms, rather than differential equations, form the native idiom of biological systems. Why the skepticism about PDEs, the workhorses of modern physics? If their “unreasonable effectiveness” in physics is any indication, shouldn’t we expect them to work their magic in the life sciences as well? Probably not. Planetary motion has its equations but cancer most likely never will. Blame it on the historicity of biology and its unending trail of broken symmetries, if you will. Whatever the reason might be, calculus is too rigid and symmetry-hungry to be the language of life. At its core, the issue is, indeed, linguistic. Suppose that, instead of following the same laws of physics, each particle in a gas were free to obey its own distinct laws. Would physics still have general principles? While not nearly as dire, the reality of living organisms is one of tremendous diversity—not just of states but of rules. Endowed with their own “algorithms,” *particles* become *agents* and the need for a new syntax arises.

The study of complex systems stands on a conceptual tripod. One leg is planted in the field of dynamics, which itself traces its roots to celestial mechanics; another one draws its methodology from statistical physics; the third leg is the computational paradigm pioneered by Turing and von Neumann. Beware of the computing hardware metaphor, however: machines speak in the tongue of Boolean processes but living systems do not. Deprived of the language of algorithms, indeed, computational modeling risks dissolving into mere simulation and data analysis. To avoid the reductionist trap of equating computing with circuitry, we need a formalism that appeals to the classic tropes of algorithmic theory: typing, abstraction, persistence, composition, amortization, fault tolerance, encapsulation, etc. The hope is that modeling living systems as *natural algorithms* will allow us to supplement (not supplant) the toolkit of dynamics and statistical physics with a new set of analytical instruments [14, 17]. Algorithms share with equations a certain closure property: just as we study equations by introducing more equations, we turn to algorithms to analyze natural algorithms. The idea is not new. Perhaps the best known example of an algorithmic argument in science is the optimality proof for the Carnot cycle.

One final point to address before we get to the heart of the matter. *Big Data* is frequently hailed as the future of science, perhaps even the “end of theory” [2]. A common line of reasoning is that the search for causation is passé and that raw correlation, freed from the pesky demands of theory, will soon drive scientific progress. To appreciate the “wisdom” of that insight, consider the RSA chip in Fig. 2, an electronic circuit for data encryption. Could the proverbial Martian ever guess what it does by probing it with electrodes and running

**Fig. 2** Would a Martian ever guess that this chip computes RSA?



massive statistics on the measurements? Not a chance. It might be countered that nature does not design number-theoretic circuits. True, but evolution does engineer its own “logic,” much of it of extreme sophistication. From single organisms to populations, ingenious error-correction, checkpoint, counting, and repair mechanisms can be found at all scales. Life does not come with an owner’s manual, however, and before we can decrypt its algorithmic narrative, we need to acquire the ability to read it. Most likely, reading and deciphering must go hand in hand. This is a key point. In English, syntax and semantics are decoupled, so one can take dictation without understanding a single word. In biology, the notions of genotype and phenotype are so intertwined as to blur the contrast. Building on Turing’s profound insight, namely the duality between program and data, natural algorithms elide this distinction altogether. They do not segregate recipes from ingredients but rather view collective behavior as the interplay among memories with different timescales.

We recently introduced a general model for multiagent dynamics, called *influence systems* [18]. We did it for two reasons: one was to study the emergence of collective behavior in specific systems (eg, consensus, swarming, synchronization, opinion dynamics); the other was to provide a minimalist theoretical platform within which one could test the utility of the natural algorithms paradigm. Influence systems raise questions that the standard tools of dynamics seem ill-equipped to tackle. After describing in broad strokes what this is about, we turn our attention in Sect. 2 to influence systems of the *diffusive* type, which generalize the notion of diffusion to heterogeneous systems [34]. We first examine the restriction to bidirectional communication in Sect. 3 and then take on the general case in Sect. 4, showing in Sect. 5 how it can be resolved via renormalization in dynamic networks. We conclude with a brief discussion of bird flocking in Sect. 6.

To clear up any possible confusion, we add that our aims run opposite to those of *bio-inspired computing*. We do not look to nature for tips on how to compute better. Rather, we seek algorithmic ideas to understand living systems. Our approach is unabashedly theoretical and our work makes no claim of practical relevance. Its ambition is more modest: it lies in the hope that an algorithmic approach to collective behavior can enrich the current analytical toolkit. Recent examples of this approach from computer science include memory bounds for ant navigation [22], shortest-path computation for slime molds [9], and maximal independent set methods in use during fly brain development [1]. All of these cases highlight algorithmic features of specific living organisms [11]. Rather than hopping from one model to the next with a different story each time, we focus our discussion on a single abstract framework in the pursuit of general principles. In the process we draw various connections to statistical

mechanics and highlight the hidden role of pseudorandomness in deterministic processes. While abstract, combinatorial models can be surprisingly accurate in statistical physics, the same cannot be said of biology. Our working models for bird flocking, for example, may have little to do with the behavior of actual birds. Why bother then? The answer lies in the dearth of analytical techniques for studying nonequilibrium systems in biology. Our objective is not so much to understand, say, how birds flock, something currently out of reach, but to discover new conceptual tools that one day will allow us to do so. Different levels of reality call for their own modes of abstraction. More, indeed, is different [3]. Social dynamics, in particular, needs to find its own mathematical language. This is the theme of this work.

## 2 Influence Systems

An *influence system* consists of  $n$  agents, labeled  $1, \dots, n$  [18]. We fix the dimension  $d$  of the ambient space and denote by  $x_i$  the position of agent  $i$  in  $\mathbb{R}^d$ . The  $d$  coordinates of  $x_i$  encode the state of agent  $i$ : these could include its physical location, momentum, temperature, or any pertinent variable associated with the agent. We represent the overall state of the system by an  $n$ -by- $d$  matrix  $\mathbf{x} = (x_1, \dots, x_n)$ , where the entry  $\mathbf{x}_{ij}$  indicates the  $j$ -th coordinate of  $x_i$ . The dynamics is specified by a function  $f$  from  $\mathbb{R}^{n \times d}$  to itself. If  $\mathbf{x}$  is the system's state at (discrete) time  $t$ , then  $f(\mathbf{x})$  indicates its state at time  $t + 1$ . To understand an influence system is to understand the geometry of the orbit  $\mathbf{x}, f(\mathbf{x}), f^2(\mathbf{x}), f^3(\mathbf{x}), \dots$ , for *any* initial state  $\mathbf{x}$  in  $\mathbb{R}^{n \times d}$ . The salient feature of an influence system is the reliance of the map  $f$  on an intermediate function  $G$  to specify the communication channels among agents. Roughly speaking, each agent is endowed with its own, distinct procedure to decide, at any step, which other agents get to be its neighbors: these are the communication rules. The agent also has access to another procedure (action rules) which tells it where to move next on the basis of the information it has just gathered from its neighbors.

### 2.1 The Model


We begin with the *communication rules*. The function  $G$  maps any state  $\mathbf{x}$  to a directed graph whose nodes are in bijection with the agents. The existence of an edge  $(i, j)$  in  $G(\mathbf{x})$  is determined by the truth condition of a first-order sentence over the reals:

$$\Phi_{i,j}(\mathbf{x}) \triangleq \square y_1 \cdots \square y_k \mathcal{L}_{i,j}(\mathbf{x}, y_1, \dots, y_k), \tag{1}$$

where  $\square \in \{\exists, \forall\}$  and  $\mathcal{L}_{i,j}$  is a quantifier-free logical proposition whose clauses are (strict and nonstrict) algebraic inequalities with rational coefficients. The pair  $(i, j)$  is an edge of  $G(\mathbf{x})$  if and only if  $\Phi_{i,j}(\mathbf{x})$  is true. Here are two examples:

1. Suppose that  $d = 3$  so each agent  $i$  is positioned at  $x_i \in \mathbb{R}^3$ . The graph  $G(\mathbf{x})$  links any agent  $i$  to its 7 nearest neighbors. This rule, taken from a recently proposed bird flocking model [4], can be expressed by a quantifier-free formula  $\Phi_{i,j}(\mathbf{x})$  consisting of the disjunction, over all subsets  $S$  of  $\{1, \dots, n\}$  of size  $n - 7$ , of the conjunction over all  $k \in S$  of the proposition "agent  $i$  is closer to  $j$  than to  $k$ ." Informally, bird  $j$  is among the 7 nearest neighbors of  $i$  if and only if there are  $n - 7$  birds further away than  $j$  from  $i$ . Note the paradoxical situation that this rule, though meant to express the locality of the birds' sensory inputs, requires global knowledge of all the birds' positions. Physics sometimes allows global information to be inferred locally. Being a nearest neighbor, for example, entails a logical formula involving all the agents. By contrast, the rule linking any bird

**Fig. 3** The communication graph is formed by joining each agent to its three nearest neighbors including itself. It is not bidirectional. The action rule of this diffusive system consists of moving every agent to the mass center of its neighbors simultaneously. Only the motion of the leftmost agent is shown in the picture



to any other within a fixed distance is local both physically and mathematically—the logical formula involves only the two agents in question.

2. To dispel the suspicion that the full theory of the reals might be overkill, we note that full quantification is there for a reason. In a popular model of sensor networks [8], robots can only communicate with their neighbors if there is no obstacle between them. This is expressed by allowing only communication channels along the so-called constrained Delaunay triangulation. Agent  $i$  links to  $j$  if there exists a sphere passing through both of them that is free of any other agent as well as any obstacles (e.g., trees, cars, buildings). We leave it as an exercise to check that this property requires two alternating runs of quantifiers:

$$\Phi_{i,j}(\mathbf{x}) = \exists y_1 \cdots \exists y_l \forall y_{l+1} \cdots \forall y_k \mathcal{L}_{i,j}(\mathbf{x}, y_1, \dots, y_k).$$

Moving on to the *action rules*, the next position of an agent is uniquely specified by the labels and the positions of its neighbors. It is required that, for any  $i$ , the  $i$ -th row of the  $n$ -by- $d$  matrix  $f(\mathbf{x})$  can be evaluated uniquely on the sole basis of the information contained in  $\{(j, x_j) \mid (i, j) \in G(\mathbf{x})\}$ . The influence system is called *diffusive* if the next position of any agent is a convex combination of its neighbors' positions. In other words,  $f(\mathbf{x}) = P(\mathbf{x})\mathbf{x}$  for any state  $\mathbf{x} \in \mathbb{R}^{n \times d}$ , where  $P(\mathbf{x})$  is a stochastic matrix such that  $P(\mathbf{x})_{ij} > 0$  if and only if  $(i, j)$  is an edge of  $G(\mathbf{x})$ . The system is called *bidirectional* if the graph  $G(\mathbf{x})$  is undirected for all  $\mathbf{x}$ .

We illustrate these definitions with the help of a few examples:

1. *Generalized HK systems* Figure 3 illustrates a nonbidirectional diffusive system inspired by the Hegselmann-Krause model (more on this below). In this example,  $\mathbf{x}$  is an  $n$ -by- $d$  matrix, where  $n = 9, d = 2$ , and each row indicates the placement of an agent. The communication rule stipulates that each agent should be joined to its two nearest neighbors plus itself (each node has outdegree 3, including a self-loop). The action is determined by simultaneously moving each of the agents to the centroid of the triangle formed by itself and its neighbors. The matrix  $P(\mathbf{x})$  coincides with the adjacency matrix of the graph  $G(\mathbf{x})$ , with its entries divided by 3 to make it stochastic.
2. *Ising model* Influence systems can be viewed as a grand generalization of spin systems. In the case of the Ising model, the communication graph  $G(\mathbf{x})$  is a fixed grid. The action is implied by the Metropolis–Hastings algorithm: each agent computes the effect of flipping its spin on the energy and biases a coin accordingly; it flips its spin if the move is energetically beneficial or if the coin toss comes up heads. To resolve synchronization issues, one can schedule the action by using the coloring of a checkerboard to ensure

that each agent acts only at every other step; other parallel updates are also possible. One dimension is needed for the spin and another one to keep track of time parity; hence  $d = 2$ . (There is no need for positional variables since the particles are fixed.)

3. *Flocking* For a simple example of a nondiffusive influence system, we can turn to a classical model for bird flocking. The bird labeled  $i$  has position  $p_i = (x_i, y_i, z_i)$  and velocity  $v_i = (\bar{x}_i, \bar{y}_i, \bar{z}_i)$ . At the next step,  $p_i$  is updated to  $p_i + v_i$  and  $v_i$  becomes  $\frac{1}{k} \sum_j v_j$ , where the sum extends over the  $k$  nearest birds  $j$  (for some small  $k$ ). The state  $\mathbf{x}$  is the  $n$ -by-6 matrix  $(x_1, \dots, x_n)$ , where  $x_i$  is the row vector  $(x_i, y_i, z_i, \bar{x}_i, \bar{y}_i, \bar{z}_i)$ . The map of the system cannot be expressed by a stochastic matrix so the system is not diffusive. Suppose we rearrange  $\mathbf{x}$  by stacking the leftmost half on top of the rightmost half, so that  $\mathbf{x}$  becomes the  $2n$ -by-3 matrix whose first and last groups of  $n$  rows become  $(x_i, y_i, z_i)$  and  $(\bar{x}_i, \bar{y}_i, \bar{z}_i)$ , respectively, for  $i = 1, \dots, n$ . It follows that

$$P(\mathbf{x}) = \begin{pmatrix} \mathbf{I}_n & \mathbf{I}_n \\ \mathbf{0} & Q_{\mathbf{x}} \end{pmatrix}, \tag{2}$$

where  $Q_{\mathbf{x}}$  is a stochastic matrix with  $k$  positive entries per row and  $\mathbf{I}_n$  is the  $n$ -by- $n$  identity matrix. The velocity part by itself forms a diffusive system. Note that the communication graph, being based on distances, cannot be inferred from the velocity in Markovian fashion (velocities must be summed over all previous steps). It is a general rule that no diffusive influence system can escape to infinity. Diffusivity, indeed, perpetually confines the mass center of the agents to the box enclosing their initial configuration.

In biology, dynamical systems are always open and reactive [23,33]: by exploiting free-energy gradients supplied by the input of energy and nutrients, they are able to “build” order while dissipating heat. Influence systems are subject to endogenous dynamics but are, in fact, sufficiently expressive to model environmental inputs (via the introduction of dedicated agents). Likewise, one could extend the model to open systems and allow for self-replication.

## 2.2 What is the Model Trying to Accomplish?

Given their generality, influence systems would appear ideally suited to model all sorts of collective behaviors. Though an important part of the plot, this is not the main story line. The principal motivation is to provide a formal context to the following line of inquiry: How would a system of  $n$  particles behave if each particle were allowed to “choose” its own laws of physics? Short of settling the question, the study of influence systems can, one hopes, uncover hidden principles. At the very least, it can test the hypothesis that insufficient attention has been given to heterogeneity in dynamics.

Physics is ruled by symmetry and to focus on heterogeneous agency is not without risk. When it is up to the modeler, and not nature, to break symmetries, it is tempting to embed into the model the answers we want. It becomes easy to “discover” emergent functions that were in fact injected as hidden assumptions. To put it plainly, there is no such thing as a minimalistic high-level abstraction. Once we decide to model the heart as a pump, whether right or wrong, we are stuck with plumbing. Ideally, influence systems should not be hand-coded but produced algorithmically. This is a delicate issue and one this discussion will tiptoe around.

By distinguishing communication from action, i.e.,  $G$  from  $f$ , the model separates the syntactic (who is talking to whom?) from the semantic (who is doing what?). If the communication function  $G$  were to ensure that every pair of agents keeps a communication channel open at all times, the influence system would become indistinguishable from a general discrete-time dynamical system and would lose its distinctive flavor. Intuitively,  $G$  is

there to enforce locality of communication. How the agents are left free to choose how to do so is part of the model's appeal. It is also what makes it hard to tame. With information traveling across each edge in a single step, regardless of the distance traveled, the metrics induced by the ambient space and the communication graphs might be unrelated and such basic notions as correlation function and characteristic length risk losing their meaning.

The philosophy behind influence systems is to keep the action rule simple so that emergent phenomena arise not so much from the sheer power of individual agents, but from the information flowing across the communication network. Our sights are on the emergence of features beyond the reach of a single agent. In biology, at least, the motivation is obvious. Why would a bird bother to flock with its friends if it could get by alone? Why would an agent seek computational help from a distributed system if it had the power of a Turing machine? Collective action might speed things up but computational efficiency is not the sort of emergence we're after. Because we look for behaviors that only collective action can engineer, it is sensible to limit the power of any given agent. In the examples we discuss, the actions of the agents are limited to integration and averaging. Needless to say, most applications in biology will require more sophisticated action rules.

Besides one-shot random perturbations, we also confine the action to deterministic rules. On the face of it, this is a severe blow: why deprive our systems of randomness when it is so obviously essential to physics and biology? The reason is to disentangle the two sources of unpredictability commonly encountered in stochastic systems. One of them results from the noise itself, with thermal energy acting as an engine of uncertainty. The other one reflects the nonlinearities of the dynamics: think of deterministic chaos. In our case, quite remarkably, this source is strong enough to produce not only chaos, but Turing completeness. Though deterministic, an influence system can behave very much like a contraction lifted from the statistical mechanics toolbox. The agents' freedom to decide deterministically their communication channels as they please confer upon the system all the trappings of stochasticity. From this emerges a notion of entropy central to the analysis. In the diffusive case, we identify the averaging performed by the agents as dissipation.<sup>1</sup> Much as in statistical mechanics, we can then interpret the dynamics near criticality as a tug-of-war between entropy and energy.

### 3 Bidirectional Communication

In anticipation of our treatment of general influence systems, we begin with the simpler case of bidirectional communication. In this restricted model, information between two agents always flows both ways, so that no agent can learn anything about another one without revealing to it something about itself. Two-way communication suffices to supply all orbits with fixed-point attractors. Such systems do not admit of quadratic Lyapunov functions [35], however, and any attempt to show "progress" by local means runs into all sorts of difficulties. We bypass this minefield altogether by introducing a *global* measure of convergence, the *s-energy*, whose interpretation as a partition function can be harnessed to yield convergence results via large-deviation arguments.

During the evolution of an influence system, at each step every agent passes on information about itself to all the agents that link to it. In the diffusive case, the information is averaged out and then passed on to other agents. Note that information diffuses outward in the *reverse* direction of the edges: if  $i$  links to  $j$ , information about  $j$  is passed on to  $i$ .

<sup>1</sup> None of the Lyapunov exponents at the regular points are positive.

If the communication graph is fixed, the transformation is equivalent (technically, dual) to a random walk through the graph, a process by now well understood.<sup>2</sup> Including an agent as a self-neighbor prevents (trivial) periodicity and ensures that the system always reaches equilibrium at a rate given by the spectral gap. If, in addition, detailed balance holds, the relaxation time (or, dually, the mixing time) is typically polynomial in the number of agents; otherwise, it can be exponential, a sign of physical irreversibility.

Granting the agents freedom to update their communication channels dynamically changes everything. The analogy with random walks is a scenario where edges are repeatedly activated and disabled as a function of the current probability distribution. Imagine a diffusion process in a medium whose topology changes in real time with the concentration. The technical challenge we face is that, with rare exceptions, spectral techniques no longer work, as eigenmodes become incoherent. The failure is structural. To illustrate the centrality of this issue, we examine a “model organism” for opinion dynamics: *HK systems*.

### 3.1 Opinion Dynamics

Also known as the *bounded-confidence* model, the *HK system* was introduced by Hegselmann and Krause [24] in an attempt to understand the emergence of polarization in opinion dynamics [12]. The communication graph links every pair of agents within unit distance of each other. The action rule moves each agent to the mass center of its neighbors. In other words, the following assignment is performed by all the agents simultaneously and it is repeated forever (Fig. 4):<sup>3</sup>

$$x_k \leftarrow \frac{1}{|\mathcal{N}_k|} \sum_{i \in \mathcal{N}_k} x_i, \quad \text{where } \mathcal{N}_k = \{1 \leq i \leq n : \|x_k - x_i\|_2 \leq 1\}. \tag{3}$$


It has been known for some time that *HK systems* always converge to a fixed configuration [7,27,29]. In fact, it was shown recently that the convergence time is polynomial in both  $n$  (the number of agents) and  $d$  (the dimension) [5]. Numerical simulations in low dimensions suggest that, if the initial position of the agents is given by a Poisson process of constant rate, the number of clusters is independent of the number of agents. This is still a conjecture awaiting a proof.

The dynamics of *HK systems* is driven by products of stochastic matrices. Indeed, we can rewrite (3) in matrix form as  $\mathbf{x} \leftarrow P(\mathbf{x})\mathbf{x}$ , where  $\mathbf{x}$  is the  $n$ -by- $d$  matrix of agent positions and  $P(\mathbf{x})$  is the stochastic matrix derived from the adjacency matrix of the communication graph by rescaling each row to sum up to 1. This constitutes a bidirectional diffusive influence system, arguably the simplest nontrivial example of its kind. The stochastic matrices induced by the system do not commute so, absent a common eigenbasis, spectral considerations fall short. A whole menagerie of *coefficients of ergodicity* [37] has been introduced to remedy this loss in the hope that, if eigenvalues won’t work, perhaps some other parameters can be found with which to build a Lyapunov function. Since *HK systems* always have a fixed-point attractor, it is indeed natural to look for something like a free energy to help us measure “progress” at each step. This works well for *HK systems* but often fails when we extend the model. It is important to understand why before turning to the general case.

<sup>2</sup> This is the difference—which we brush over—between studying  $P^t x$  (primal) and  $y^T P^t$  (dual) for  $t = 0, 1, \dots$

<sup>3</sup> We use the computer programming notation  $y \leftarrow x$  to indicate that the variable  $y$  takes on the value  $x$  at the current step.





**Fig. 4** Clockwise, 20,000 randomly placed agents move under *HK* dynamics to form shapes of decreasing “dimension”

### 3.2 The *s*-Energy

Since the dynamics evolves as  $Q_t \mathbf{x}$ , where  $Q_t$  is a product of stochastic matrices, it is sensible to step back from *HK* systems and begin with a more general line of inquiry. Let  $(P_t)_{t \geq 0}$  be an infinite sequence of  $n$ -by- $n$  stochastic matrices. Given  $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$ , we consider the orbit defined by  $\mathbf{x}(0) = \mathbf{x}$  and, for  $t > 0$ ,


$$\mathbf{x}(t) \triangleq P_{t-1} \cdots P_0 \mathbf{x}. \tag{4}$$

Under what conditions is the sequence  $\mathbf{x}(0), \mathbf{x}(1), \mathbf{x}(2), \dots$  attracted to a fixed point? If  $n = 2$  and  $P_t = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$  for all  $t$ , then  $\mathbf{x}(t)$  is stuck in a 2-cycle if  $x_1 \neq x_2$ , so we need to enforce positivity along the diagonal entries of each  $P_t$  to avoid (trivial) periodicity. This is not sufficient, however, since composing the two matrices

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 2/3 & 1/3 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 2/3 & 0 & 1/3 \end{pmatrix}$$

in alternation exchanges the vectors  $(0, 1, 1/4)$  and  $(0, 1, 3/4)$  endlessly. This time, the periodicity is caused by the lack of bidirectionality: indeed, the  $n$ -node graph  $G_t$  whose edges correspond to the positive entries in  $P_t$  is *directed*. To keep with the spirit of *HK* systems, we enforce bidirectionality by requiring that  $P_t$  be *type-symmetric*, meaning that any two diagonally opposite entries of any  $P_t$  must be either both zero or both positive.

**Fig. 5** The deflating matrix polytope



With these conditions in place, quite remarkably, every orbit is attracted asymptotically to a fixed point [25,29,32].<sup>4</sup> This can be proven with standard techniques from dynamics. Bounding the convergence time requires a new idea, however. Here is why. If  $\text{conv } P$  denotes the convex hull of the points formed by the rows of the matrix  $P$ , we have the nesting structure (Fig. 5):

$$\text{conv } P_t \cdots P_0 \subseteq \cdots \subseteq \text{conv } P_1 P_0 \subseteq \text{conv } P_0.$$


Whether it is done with eigenvalues or, should that fail, any of the other coefficients of ergodicity defined in the literature, convergence is tied to the shrinking of the nested polytopes. Specifically, one picks a measure such as width, diameter, or volume and then tries to show that it decreases by a certain amount at each step: a lower bound on the rate of decay yields a commensurate bound on the attraction rate. The problem is that, by seeking progress *at each step*, this approach cannot cope with “stalling”, which refers to long periods of little or no shrinking. Since we make no assumption on the sequence  $(P_t)_{t \geq 0}$ , this can certainly happen: imagine long runs of the identity matrix thrown into the mix. Rather than aiming for stepwise progress toward a fixed-point attractor, therefore, we change our perspective and look at entire orbits at once: we then define a parameterized series for each one. The approach is akin to taking the Fourier series of a periodic signal: the goal is to switch the meaning of local and global. We define the *s-energy* [16] for any real  $s \geq 0$  as

$$E(s) = \sum_{t \geq 0} \sum_{(i,j) \in G_t} |x_i(t) - x_j(t)|^s. \tag{5}$$

Technically, this defines the total *s-energy*. A close cousin, the *kinetic s-energy*, measures the displacement of the agents [16]. The (total) 0-energy counts the number of edges in  $G_t$  for all  $t \geq 0$  and is usually infinite. The 1-energy indicates how much ink is needed to draw all the graphs. As defined, it is not even clear that the *s-energy* converges even for *any* value of  $s$ . If it does for  $s_0 > 0$ , however, it also does over the semi-infinite interval  $[s_0, \infty)$ . The *s-energy* can be specialized to rederive classic mixing times for Markov chains; its true power, however, lies in its ability to handle time-varying chains. Mathematically, the function is a generalized Dirichlet series: it is invertible and forms a lossless encoding of the edge lengths. We have shown that the series is uniformly convergent over any finite region  $\mathcal{D}$  of the

<sup>4</sup> In the special case where  $P_t$  is constant, we recover a basic fact about aperiodic Markov chains. Keeping the diagonal positive to avoid periodicity defines what are known as *lazy random walks*.

**Fig. 6** The analytic continuation of  $|E_2(s)|$



complex plane within  $\Re(s) \geq r$ , for any  $r > 0$ , and it defines an analytic function over  $\mathcal{D}$ .<sup>5</sup> Like the Riemann zeta function  $\sum n^{-s}$ , the system’s underlying structure is multiplicative: indeed, just as  $n$  is a product of primes,  $x_i(t) - x_j(t)$  is a product of the form  $v^T P_{t-1} \cdots P_0 \mathbf{x}$ , for some suitable vector  $v$ .

Let  $E_n(s)$  denote the maximum value of  $E(s)$  over all  $\mathbf{x} \in [0, 1]^n$ . (Obvious scaling laws allow us to confine the phase space to the unit cube.) One should expect the function  $E_n(s)$  to encode all sorts of symmetries. This can be shown for  $n = 2$  by proving that it can be continued into a meromorphic function over the whole complex plane, with all the poles along the imaginary axis (Fig. 6). By analogy with the Riemann zeta function, generalizing this result to all values of  $n$  would likely uncover hidden symmetries in the dynamics generated by products of stochastic matrices.

We have shown [16] that, if  $\rho$  is the smallest nonzero entry among the matrices  $P_t$ , then  $E_n(1) = \Omega(1/\rho)^{\lfloor n/2 \rfloor}$  and  $E_n(s) = s^{1-n} (1/\rho)^{\Omega(n)}$ , for  $n$  large enough, any  $s \leq s_0$ , and any fixed positive  $s_0 < 1$ . These lower bounds are tight except for  $s < 1$ , where the best known exponent is quadratic in  $n$ . Regarding upper bounds,

$$E_n(s) \leq \begin{cases} \left(\frac{1}{\rho}\right)^{O(n)} & \text{for } s = 1; \\ s^{1-n} \left(\frac{1}{\rho}\right)^{n^2 + O(1)} & \text{for } 0 < s < 1. \end{cases} \tag{6}$$

Recall that the sequence of type-symmetric matrices  $P_t$ , hence of graphs  $G_t$ , is left completely arbitrary, except for the requirement that the diagonals be positive and the nonzero entries exceed a fixed positive value. Remarkably, the vector  $\mathbf{x}(t)$  always converges to a fixed point. The attraction is asymptotic, so to measure it requires picking some small  $\varepsilon > 0$  and asking how many steps of length at least  $\varepsilon$  are needed to reach equilibrium. (It is easy to see why asking for the relaxation time is the wrong question.) If the *communication count*  $C_\varepsilon$  denotes the maximum number of times  $t$  at which  $G_t$  has at least one edge of length  $\varepsilon$  or higher, then the bounds on the  $s$ -energy given in (6) imply:

$$C_\varepsilon(n) \leq \min \left\{ \frac{1}{\varepsilon} \left(\frac{1}{\rho}\right)^{O(n)}, \left(\log \frac{1}{\varepsilon}\right)^{n-1} \left(\frac{1}{\rho}\right)^{n^2 + O(1)} \right\}. \tag{7}$$

We assume that the initial diameter of the system is 1, which we can always do in view of the linear scaling laws of the dynamics. These bounds are optimal for  $\varepsilon > \rho^{O(n)}$  where they become  $C_\varepsilon(n) \leq \rho^{-O(n)}$ . For reversible systems (a notion defined below), the bounds become polynomial:  $C_\varepsilon \leq O\left(\frac{1}{\rho} n^2 \log \frac{1}{\varepsilon}\right)$  for  $\varepsilon < \rho/n$ . As one might expect, reversibility plays a crucial role in the communication count.

<sup>5</sup> Trivially, the continuation cannot always extend over the whole complex plane.

### 3.3 An Algorithmic Proof

We prove the bound on the 1-energy given in (6) by algorithmic means. What does that mean? The details are fairly technical, but here is the gist of the argument. (We omit the case  $s < 1$ , which is handled by setting up a recurrence relation—see [16] for all the proofs.) To begin with, one can show that it suffices to focus on the kinetic 1-energy. Think of the  $n$  agents as cars initially positioned at  $x_1, \dots, x_n$ . Suppose that, for every car, we knew an upper bound on how many miles it would drive. We could then fill each tank with enough gas so the car could complete its journey. At the end, we could see how much gas was left in each car and from that information infer the total mileage of all the cars, ie, the kinetic 1-energy. The reason this setup is interesting is that we can use it even when we do not know ahead of time how much gas each car will need.

Suppose that we only knew an upper bound on how much gas the *total* fleet will use. We could then divide that amount evenly and fill each tank with the same number of gallons. When a car completes its ride, it deposits its remaining gas into a tanker. It can also do so when it estimates that its current supply exceeds its needs. When a car runs out of gas, on the other hand, it should be able to refill by using the tanker. Our task is now to design a trading procedure that instructs under which circumstances drivers should take or deposit gas and how much. The existence of such a protocol, together with the guarantee that no car will ever run out of gas with an empty tanker, gives us a handle on the kinetic 1-energy.

We thus reduce the entire analysis to the design of a gas-trading algorithm. The details do not matter here: we only wish to make a point of methodology. The proof goes beyond the design of a Lyapunov function (for example, the amount of unused gas): it requires an actual trade protocol. The power of the framework is to allow credit (i. e., borrowing from the future). We will encounter more examples of algorithmic proofs below. The need for them reflects the loss of symmetry incurred by bringing heterogeneity into the dynamics.

### 3.4 The View from Statistical Mechanics

With one foot in physics and the other one in probability theory, statistical mechanics can provide a conceptual framework for problems with no obvious physical reality. A good example is the recent flurry of work on constraint satisfaction, error correction, belief propagation, and social dynamics [12,31]. In fact, many statistical mechanics techniques do not even require randomness per se: they can be viewed as powerful counting tools. In particular, they can yield tail estimates for positive convergent series. In that spirit, we show how a standard large-deviation argument from statistical mechanics can be used to bound the communication count  $C_\varepsilon$ . (In the language of probabilities, the same argument can be interpreted as a Markov bound applied to the moment generating function.) We interpret the  $s$ -energy as the partition function<sup>6</sup>

$$E(s) = \sum_{k \geq 0} e^{-s \mathcal{E}_k},$$

where  $s = 1/T$  is the (normalized) inverse temperature and  $\mathcal{E}_k = \ln(1/d_k)$  is the energy (Hamiltonian) of the edge  $(i, j)$  of  $G_t$ , with  $d_k = |x_i(t) - x_j(t)|$ . Perhaps a more physical interpretation would redefine  $\mathcal{E}_k$  in terms of the Dirichlet sum

<sup>6</sup> Statistical physicists might cringe at the terminology, but the  $s$ -energy derives its name from discrepancy theory.

$$-\ln \left\{ \sum_{(i,j) \in G_t} |x_i(t) - x_j(t)|^2 \right\},$$

but this makes little difference for our purposes, so we stick to our original definition. Assuming that  $Z \triangleq E(s)$  is bounded, we define the Boltzmann distribution  $p_k(s) = e^{-s \mathcal{E}_k} / Z$ . As usual, we can express basic thermodynamic quantities in term of the free energy  $F = -T \ln Z$ . For example, the mean energy  $U = \langle \mathcal{E}_k \rangle$  and the entropy  $S = -\langle \ln p_k \rangle$  become

$$U = -T^2 \frac{\partial(F/T)}{\partial T} \quad \text{and} \quad S = -\frac{\partial F}{\partial T} \quad \text{and} \quad F = U - TS.$$

It is also useful to express the mean energy directly as a function of the inverse temperature:

$$U = -\frac{\partial \ln E(s)}{\partial s}.$$

Truncating the infinite series formed by the  $s$ -energy gives a uniform upper bound on the communication cost:

$$C_\varepsilon \leq E(s)e^{-s \ln \varepsilon}.$$

Setting the inverse temperature  $s$  to minimize the right-hand side above leads to the Legendre transform and the rate function

$$\inf_{s>0} \left\{ \ln E(s) - s \ln \varepsilon \right\}.$$

Taking derivatives yields  $\partial \ln E(s) / \partial s = \ln \varepsilon$ . It follows that we must set the equilibrium inverse temperature  $s_0$  so that the mean energy  $U$  is equal to  $\ln \frac{1}{\varepsilon}$ . By (6), we easily check that  $s_0 \approx n / \ln \frac{1}{\varepsilon}$ , which proves the right-hand side of (7): this states that, for small enough  $\varepsilon > 0$  (to ensure  $s_0 < 1$ ),

$$C_\varepsilon \leq \left( \log \frac{1}{\varepsilon} \right)^{n-1} \left( \frac{1}{\rho} \right)^{n^2 + O(1)}.$$

### 3.5 Reversibility and the Attraction Rate

Reversible Markov chains are commonly used to model thermodynamic systems in equilibrium. In the ergodic case (ie, irreducible and aperiodic), the stationary distribution is unique and particles hopping along the chain from an initial state sampled from it will exhibit no “arrow of time.” This means that videos of the trajectory played forward and backward will be statistically indistinguishable. In the ergodic case, the dominant eigenvalue has multiplicity one and the spectral gap is large, which in turn implies rapid mixing. In time-varying Markov chains, by contrast, transitions lose their spectral coherence. As a result, it is possible to compose reversible chains to produce exponentially long mixing times, something impossible to do with a single chain. This raises the question: Can one still define a notion of reversibility for the dynamics defined in (4)? The answer is yes. At each step, every agent moves toward the mass center of its neighbors so as to leave the overall center of gravity of the whole system unchanged. The latter is computed by weighting each agent with its stationary probability. Algebraically, this assumes that all matrices  $P_t$  share a common dominant left eigenvector. Geometrically, this means that the agents are chasing a fixed target and not a point that keeps wiggling around. Assuming  $G_t$  remains connected, the agents reach consensus, ie, fall within an interval of length  $\varepsilon < \rho/n$  in time  $O(n^2 \rho^{-1} \log \frac{1}{\varepsilon})$ , where  $\rho$  is the smallest nonzero entry

among the matrices  $P_i$ . This result, which is optimal, follows from the fact [16] that, for all  $0 < s \leq 1$ , the  $s$ -energy is of the form (for constant  $c > 0$ ):

$$E_n(s) \leq \frac{n^c}{s} \left(\frac{1}{\rho}\right)^{s/2+1}.$$

#### 4 Diffusive Influence Systems

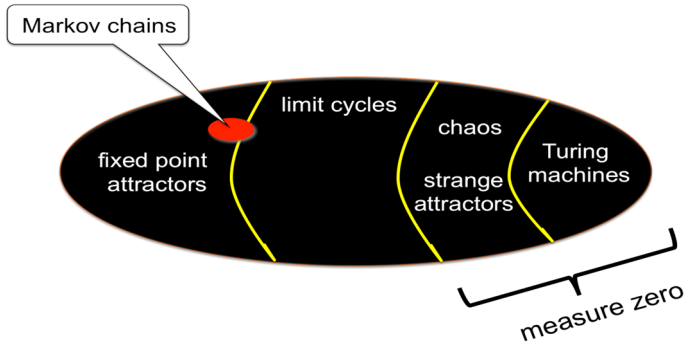
The previous section confined its attention to bidirectional systems: the flow of information between two agents always went both ways (eg, cell phones allowed but no radios). It is this feature that ensured convergence. Relaxing the bidirectionality constraint (eg, allowing a bird to see another one without being seen by it) adds considerable expressivity to the dynamics. Recall that a diffusive influence system is specified by the map,  $\mathbf{x} \mapsto P(\mathbf{x})\mathbf{x}$ , where  $\mathbf{x}$  is an  $n$ -by- $d$  matrix whose  $i$ -th row indicates the position of agent  $i$  in  $\mathbb{R}^d$  and  $P(\mathbf{x})$  is a stochastic matrix specified by a first-order sentence parameterized by  $\mathbf{x}$ . We have proven the surprising fact that the dynamics, though almost always asymptotically periodic, has a critical regime of arbitrarily high complexity [18]. Among many uses, diffusive influence systems can be enlisted to model how people change opinions over time as a result of human interaction and knowledge acquisition. Unless people keep varying the modalities of their interactions, as mediated by such things as trust levels, self-confidence, etc, after a learning phase, they will be caught forever recycling the same opinions in the same order. If the interactions are finetuned to fall within the critical regime, however, opinions can grow to solve any problem.<sup>7</sup>

The result, which we state formally below, reveals the existence of a phase transition between periodic and chaotic dynamics. The number of critical “temperatures” is infinite and the critical region, indeed, is likely to be fractal (somewhat reminiscent of the logistic map). As physical objects, influence systems operate out of equilibrium: they entail the constant injection of energy and dissipation of heat. The real point of departure from statistical mechanics, however, is the ever-changing topology of interactions, which blurs even such a basic concept as the correlation length. Perhaps the most surprising fact about these systems is that anything nontrivial can be proven at all. We state the main result, explain what it means, and sketch the arguments behind the proof. We omit the precise perturbation assumptions behind this result and refer the reader to [18] for all the technical details. We also skip the explicit constructions of chaotic and Turing-universal diffusive influence systems, which on their own shed little light on the issue of criticality. Our goal is to provide an intuitive feel for the hurdles to clear in order to prove the claim below. We show how a form of algorithmic renormalization can help.

**Claim** *Given any initial state, the orbit of a perturbed diffusive influence system is asymptotically periodic almost surely. At the critical values of the perturbation parameter, the system can be chaotic or Turing-complete.*

The phase space is broken up into basins of attraction leading to limit cycles of varying sizes. The number of periodic orbits is finite up to foliation. As the system approaches criticality, the periods increase and the effects of tiny fluctuations become more dramatic. Eventually the behavior can become indistinguishable from random. It is always the case that at least two agents have a fixed-point attractor. More generally, the number of stabilizing agents always exceeds the dimension of the attracting manifold (if any).

<sup>7</sup> Needless to say, this must be taken metaphorically and not as a sociological claim.



**Fig. 7** Diffusive influence systems range from Markov chains to Turing machines. The critical region of measure zero is (probably) fractal-like and spread all across the space of systems. (The figure is just a Venn diagram with no relation to the shapes and dimensions of the various regions.)

#### 4.1 Perturbation and Criticality

We do not perturb the initial state of the system but, rather, the communication channels. Specifically, we add an arbitrarily small random number  $\delta$  to the algebraic constraints used to define the communication rules: all of them appear in the formula  $\mathcal{L}_{i,j}$  in (1). For example, assuming that the agents all lie on the real line (i.e.,  $d = 1$ ), a communication rule stipulating that agent  $i$  should listen to  $j$  whenever  $|x_i - x_j| \leq 1$  is perturbed by substituting the condition  $|x_i - x_j| \leq 1 + \delta$ , for a random  $\delta$  picked in  $[0, \delta_0]$  and a fixed  $\delta_0 > 0$  chosen as small as we want. This is the sole source of true randomness for the system. The nonlinearities provide the pseudorandomness. To avoid scaling issues, we can always assume that each of the agents lies in  $[0, 1]^d$ . The parameter  $\delta$  plays the role of the “temperature.” As  $\delta$  approaches the critical region, limit cycles get longer and longer, until they become infinite and periodicity is lost.


At criticality, chaos appears and even Turing universality, meaning that all dynamic regimes of complex systems can be achieved. This requires finetuning because the critical region is of measure zero. In fact, it lies within a Cantor set of Hausdorff dimension less than 1 (Fig. 7). Very little is known about the geometry of the critical region or for that matter the growth of the periods near criticality. Whether the systems fall within universality classes is an open question.<sup>8</sup> The claim above requires another sort of perturbation rule regarding the status of infinitesimally close agents and vanishing edges. Pinpointing the minimal set of perturbation constraints is still open. It is of little concern to us here, as it hinges on rather technical issues about the ranks of certain matrices that are unrelated to the main proof technique.

#### 4.2 Doing Away with Logic

The first-order theory of the reals comes with a syntax rich with quantifiers, connectives, and algebraic inequalities—hardly the language of dynamics. It has been known since Tarski, however, that the quantifiers can be eliminated: this is how the decidability of Euclidean geometry was established. Collins provided an effective procedure for doing this via his celebrated “cylindrical algebraic decomposition” [19]. In simple terms, all this means is the

<sup>8</sup> As is, for that matter, defining precisely what we mean by that.

**Fig. 8** The margin  $\mathcal{M}_\delta$  subdivides the phase space  $[0, 1]^n$  into atoms. The restriction of  $f$  to each atom is specified by a stochastic matrix



existence of a set of polynomials whose variables are the coordinates of the agents' positions (all  $dn$  of them) and whose signs specify the graph  $G(\mathbf{x})$ . Picture a set of algebraic surfaces cutting up the phase space  $[0, 1]^{dn}$  into *cells*. Each such cell  $c$  is associated with its own graph  $G(\mathbf{x})$ , which is the same for all  $\mathbf{x} \in c$ . Navigating across algebraic varieties may seem inauspicious but diffusivity allows us to linearize all the constraints, thus turning surfaces into hyperplanes. The idea behind the transformation is simple: replace, say,  $x^2 + 3xy$  by  $z_1 + 3z_2$ , where  $z_1 = x^2$  and  $z_2 = xy$ , and observe that if  $(x, y, \dots)$  is multiplied by a stochastic matrix  $P$ , then  $(z_1, z_2, \dots)$  is multiplied by the Kronecker product  $P \otimes P$ , which is also a stochastic matrix. This form of tensor lift establishes a conjugacy with a diffusive influence system in a higher-dimensional space whose dimension is a function of  $d, n$ , and the degree of the polynomials.

The next step is to create a “clone” of every agent for each one of its coordinates and lower the dimension of the ambient space to  $d = 1$ . Being peripheral to the main argument, the details can be skipped. To summarize, the system is piecewise-linear in the sense of [38]. The dynamics is entirely specified by assigning a stochastic matrix to each cell of an arrangement<sup>9</sup> of hyperplanes in dimension  $n$ . These cells are the basic building blocks of the system and are called the *atoms*. Each hyperplane is of the form  $\mathbf{a}^T \mathbf{x} = 1 + \delta$ , where the “temperature”  $\delta$  is picked randomly in  $[0, \delta_0]$ , for a small fixed  $\delta_0 > 0$ . The union of all the hyperplanes is called the *margin* and is denoted by  $\mathcal{M}_\delta$ . The atoms are the connected components of the complement of the margin within the phase space  $[0, 1]^n$ . Given a starting point  $\mathbf{x}$ , the orbit is formed by mapping the current point via the linear map associated with the atom to which it belongs. In other words,  $f(\mathbf{x}) = P(\mathbf{x})\mathbf{x}$ , where  $P(\mathbf{x}) = P_c$  is a stochastic matrix invariant over each atom  $c$  (Fig. 8); for technical reasons, we also assume a positive diagonal.

### 4.3 Energy Versus Entropy

Although neither concept arises in obvious fashion from the model, the tension between energy and entropy is what drives the dynamics and it is crucial to understand their interplay. To do this, we turn again to Fig. 8. The atom  $c$  is mapped to a polytope  $f(c)$  (a triangle in the picture) that is cut through by some hyperplane of the margin. Each slice, in turn, is mapped by a distinct stochastic matrix, which disconnects the image. Indeed, in the picture,  $f^2(c) \triangleq f(f(c))$  consists of two disjoint polytopes. Because of the curse of dimensionality,

<sup>9</sup> An arrangement of hyperplanes divides  $n$ -space into cells of all dimensions between 0 and  $n$ . We need only focus on the full-dimensional cells for all our purposes.



the number of slices could be as high as exponential in  $n$ . Smashing bodies into little pieces does not augur well for the rest of the analysis since, like fireworks, this could happen again and again.

To ease our way into this complication, let us begin with a make-believe scenario where the iterated images of  $c$  under  $f$  never intersect the margin and the map  $f$  itself is piecewise-isometric. If so, not only does  $f^t(c)$  remain connected for all  $t$  but, inevitably,  $f^{t_1}(c)$  must intersect  $f^{t_2}(c)$  for some  $t_1, t_2$  ( $t_1 < t_2$ ), meaning that  $b = f^{t_1}(c) \cup f^{t_2}(c)$  is connected. Let  $c_1, \dots, c_{t_2-t_1}, c_1$  be the atoms hit by the iterated images of  $c$  between times  $t_1$  and  $t_2$ . We now follow the orbit of  $b$ . Having assumed away any contact with the margin,  $f^k(b)$  is connected and bounces from atom to atom as  $k$  grows from 1 to  $t_2 - t_1$ . Which atoms? Obviously,  $c_1, \dots, c_{t_2-t_1}, c_1$ . Furthermore,  $f^{t_2-t_1}(b)$  intersects  $b$ , so we can repeat the same argument and, by induction, infer that the iterated images of  $b$  traverse atoms cyclically. The symbolic dynamics of the system is eventually periodic. Of course, this world is fiction. Some iterated image of  $c$  is likely to hit the margin at some point. Indeed, the margin and the stochastic matrices associated with the atoms are chosen independently so, from the perspective of the latter, the piecewise-linear map  $f$  might as well act quasi-randomly. In Fig. 8, imagine picking up the triangle  $c$  and forming its image under  $f$  by throwing it at random, then taking the image and again throwing it at random. Sooner or later, an iterated image is bound to hit the margin and cause a split of the sort shown in the figure. Intuitively, such a system is likely to be mixing, as it keeps breaking up the phase space into little pieces and throwing them at random. This causes the phase space to be endlessly fragmented and shuffled around, a sure sign of chaos.<sup>10</sup>


What could possibly save the system from this fate? Perhaps the fact that the shuffling draws from a finite number of linear maps, hence can't be all that random. This life preserver won't float. If there is one thing theoretical computer science has taught us, it's how easy it is to produce fake randomness that is undistinguishable from the real thing. Betting on the nonrandomness of the map sequences is a fool's errand. In fact, the pseudorandomness is so effective we should conceptually think of the orbits as random hops from atom to atom. As time elapses, the iterated images of the phase space might get increasingly splintered, fragmented, and mixed, which causes the entropy to grow. What might keep chaos at bay is that the map  $f$  is not volume-preserving. We explore this thread in a few elementary steps.

Imagine that the linear maps associated with the atoms were strictly contractive, meaning that they bring any two points closer together (something patently false<sup>11</sup>), then perhaps the iterated images of  $c$  would shrink fast enough to avoid crashing into the (perturbed) margin at a sufficiently high rate. Indeed, the probability that throwing  $c$  at random hits the margin decreases if we scale  $c$  down; in fact, we can drive the hitting probability to zero by squeezing  $c$  down to a single point. Contractivity corresponds to dissipation, so we may interpret the shrinking as a loss of energy. The tug-of-war between entropy and energy is now clear: the increase of entropy pushes the system toward chaos (or even computational universality) while an energy that decays faster than the entropy grows pushes the system toward an ordered state. Before we address the serious issue that the system is actually far from contractive, let's keep pretending that it is and see how energy dissipation might create limit cycles.

<sup>10</sup> Somewhat paradoxically, piecewise isometries cannot be chaotic [10] but piecewise contractions can [28]. This hinges on technical facts about multiplicities in polyhedral cell complexes and is not germane to the basic intuition we're trying to build.

<sup>11</sup> Stochastic matrices can actually increase distances by a factor as large as  $\sqrt{n}$ .

**Fig. 9** The coding tree  $\mathcal{T}$



### 4.4 In Search of Limit Cycles


We appealed to the intuition that, if the system were sufficiently piecewise-contractive (which, of course, it is not), perturbation would make it asymptotically periodic. We flesh out this idea. Contractivity implies that the restriction of  $f$  to any atom maps any segment to a shorter one. It will then be the case that, at some time  $t$  large enough, the set of all atoms will be refined as a disjoint collection of cells whose  $t$ -th iterated images are very small. The notion of refinement is crucial so we formalize it. For any  $t$ , we refine the set of atoms into a collection  $\mathcal{S}_t$  of cells whose  $k$ -th iterated images stay clear of the margin for  $k = 0, \dots, t - 1$ : formally,  $\mathcal{S}_1$  is the union of all atoms and, for  $t > 1$ ,

$$\mathcal{S}_t = [0, 1]^n \setminus \bigcup_{k=0}^{t-1} f^{-k}(\mathcal{M}_\delta).$$

This decomposition induces an arborecence in the obvious way. Figure 9 illustrates (a fragment of) the first three levels of the coding tree  $\mathcal{T}$ : the root is associated with the phase space  $[0, 1]^n$ ; any child of the root denotes the image of an atom under  $f$ . In the figure, the (rectangular) atom  $c$  is mapped to the rectangle appearing next to the label  $v$ . This rectangle is cut into two parts by the margin, which in turn induces a splitting of  $c$  by the preimage of the line cutting through  $f(c)$  at node  $v$ . The figure depicts only one child of the root but there is one for each atom, bringing the total to  $\#\mathcal{S}_1 = m^{O(n)}$ , where  $m$  is the number of hyperplanes in the margin. At any level  $t$ , the nodes represent the  $t$ -th iterated images of the connected components of  $\mathcal{S}_t$ . The coding tree encodes the symbolic dynamics of the system, as all valid atom trajectories can be read off the tree's paths. It is not only a combinatorial object, however. The branches of the tree have a thickness which relates to the energy while the branching rate determines the entropic growth.


By (our false assumption of) contractivity, branches get thinner and thinner as we go down the tree. Equivalently, if we wait long enough, the  $t$ -th iterated image of any cell  $c$  of  $\mathcal{S}_t$  becomes as small as we want. Indeed, for any  $\varepsilon > 0$ , we can find some  $t$  large enough so that  $f^t(c)$  fits within a ball of radius  $\varepsilon$ . At this point, we have another way to avoid collision with the margin besides refinement: we can perturb the margin. As Fig. 10 indicates, choosing a random  $\delta$  in  $[0, \delta_0]$  ensures that the margin  $\mathcal{M}_\delta$  avoids all the balls. This will happen with high probability as long as  $\delta_0$  sufficiently larger than  $\varepsilon$ . (This argument is not quite right but we return to clean it up at the end of this section.)

Of course, this hinges on the hope that the contraction rate outpaces the growth of  $\#\mathcal{S}_t$ . It is not enough that the balls in Fig. 10 should be small enough (via high energy dissipation): they should also not be too numerous, a consequence of low entropic growth. If the energy loss “beats” the entropy gain, a union bound will work and the  $t$ -iterated images of the cells



**Fig. 10** A random shift of the margin prevents cutting through the  $\varepsilon$ -balls

**Fig. 11** The exclusion zone prevents collisions and is enclosed within a fractal



of  $S_t$  will avoid the margin  $\mathcal{M}_\delta$  for a random choice of  $\delta$  in the small interval  $[0, \delta_0]$ . Let  $t_\nu$  be the earliest so-called *nesting time* at which this happens (assuming that it does):

$$f^{t_\nu}(c) \cap \mathcal{M}_\delta = \emptyset,$$

for all  $c \in S_{t_\nu}$ . If the nesting time is finite, then something remarkable happens. The image of any cell of  $S_{t_\nu}$  under  $f$  lies entirely within a single cell of  $S_{t_\nu}$ . An immediate consequence is that the symbolic dynamics of the system is eventually periodic: if we annotate an orbit with the labels of the cells of  $S_{t_\nu}$  it visits along the way, the sequence will inevitably repeat one label twice and from then on lock into a cycle. How we go from periodic symbolic dynamics to asymptotic periodicity is a question we can postpone. First, we need to address our bogus contractivity assumption.

But even before we get to this, let us return to our earlier comment that the perturbation argument was “not quite right.” Recall that  $S_t$  refines the set of atoms into a collection of maximal cells  $c$  such that  $f^k(c) \cap \mathcal{M}_\delta = \emptyset$  for  $k = 0, \dots, t - 1$ . We pick  $t \geq t_\nu$  large enough so that  $f^t(c)$  is so small that the tiniest perturbation of the margin will make it collision-free. This argument would work if the perturbation affected the margin but *not*  $f^t(c)$ . But since the cell  $c$  is defined by preimages of the margin’s hyperplanes, shifting a hyperplane away from  $f^t(c)$  might potentially move  $f^t(c)$  along with it and keep collisions unresolved. The solution is to include  $\delta$  in the phase space itself, replacing  $[0, 1]^n$  by  $[0, 1]^n \times [0, \delta_0]$  and the map  $f : \mathbf{x} \mapsto f(\mathbf{x})$  by  $g : (\mathbf{x}, \delta) \mapsto (f(\mathbf{x}), \delta)$ . Note that the lift keeps the system diffusive. Figure 11 illustrates the case where the phase space is one-dimensional. The boxes enclose the  $t$ -th iterated images of the cells of  $S_t$ . Contractivity ensures their narrowness. There is

no contraction along the  $\delta$  axis so their height can be as high as  $\delta_0$ . For fixed  $\delta$ , the margin  $\mathcal{M}_\delta$  is a finite set of points on the line  $Y = \delta$  but it becomes a collection of lines when we let  $\delta$  vary. Projecting onto the  $\delta$  axis the portions of these lines that lie within the boxes gives us an *exclusion zone*: a collection of intervals that must be removed from  $[0, \delta_0]$  in order to prevent any collision between  $f^t(c)$  and the margin ( $c \in S_t$ ). Once all the exclusion zones are assembled together, we end up with a mysterious geometric object: to bound its size we enclose it in a fractal (a Cantor set) of Hausdorff dimension less than 1.

### 4.5 Entropic Growth

We claimed earlier that a tiny perturbation will bring about periodicity as long as the energy dissipation outpaces the entropic growth. The idea was that, in Fig. 10, the number of balls should be small enough that a random shift will likely remove all the collisions with the margin. In the worst case, every node in the coding tree has up to  $\# S_1 = m^{O(n)}$  children. This is the well-known curse of dimensionality: a set of hyperplanes cutting up space into a number of cells exponential in the dimension. If  $m^{O(n)}$  were indeed the degree of every node of the coding tree, then the cross-section of the branches (of the tree) would have to shrink by at least a factor of roughly  $1/m$  in all directions. This is clearly impossible: not only does a stochastic matrix cause no shrinking whatsoever along the principal eigenvector but its average eigenvalue can easily be of magnitude bounded away from 0 by a constant anyway. To assume the worst-case entropic growth and still count on the energy dissipation to do the job, therefore, is doomed. One must argue that the entropy itself does not grow too fast. (Of course, we need to show that the topological entropy is zero but our broader definition of entropy will imply that as well.)

Geometrically, the task is to show that, given any ball  $b$  in  $[0, 1]^n$ , its iterated images  $f(b), f^2(b), \dots$  cannot hit the margin too frequently—or, if they do, then perturbing the margin takes care of it. This may sound like a straightforward “general position” feature but it is actually highly nontrivial. For one thing it must use special properties of  $f$ . If, indeed,  $f$  were a translation by  $\vec{v}$ , for example, perturbation would do no good. To see why, think of a margin in two dimensions consisting of parallel lines  $\ell, \ell + \vec{v}, \ell + 2\vec{v}$ , etc. Placing the ball  $b$  outside the margin, say on  $\ell - \vec{v}$ , will always cause the same sequence of collisions regardless of perturbation. Since perturbation involves translations, one must use the fact that stochastic matrices rescale objects in a manner that does *not* respect translations.

The issues at stake are technical but it is still possible to give a rough sketch of what needs to be done. The remainder of this section is mostly of technical interest and can be skipped by the reader more interested in the big picture. Fix an integer  $D$  large enough; for  $k = 0, \dots, D$ , let  $a_k$  be a row vector in  $\mathbb{R}^n$  and  $A_k$  be an  $n$ -by- $n$  nonnegative matrix. Write  $v_k = a_k A_k \cdots A_1$  and assume that the maximum row-sum is less than 1:

$$\alpha \triangleq \max_{k>0} \|A_k \mathbf{1}_n\|_\infty < 1. \tag{8}$$

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. For reasons we sketch below, we need to show that, for some  $D_0$  sufficiently smaller than  $D$  and any  $I \subseteq \{0, \dots, D\}$  of size at least  $D_0$ , the vector  $\mathbf{1}_{|I|}$  is not only outside the column space of  $V_I$  but “far enough” from it: technically, we must show the existence of a unit vector  $u \in \mathbb{R}^{|I|}$  such that

$$u^T V_I = \mathbf{0} \quad \text{and} \quad u^T \mathbf{1}_{|I|} \geq n^{-O(n^3 D)}.$$

We have established the upper bound:

$$D_0 \leq D^{1-\beta}, \quad (9)$$

for some small  $\beta = \beta(n, \alpha) > 0$ . The proof being only of technical interest, we omit it, but we say a few words about the motivation. Branching in the tree reflects the collision of orbits with the margin. As we explain in the next section, we may assume that some coordinates are held fixed ( $y_i$ ) so that any such collision can be expressed via its preimage by a linear equation of the form:  $\sum_i c_i x_i + \sum_j c'_j y_j = 1 + \delta$ . If we pick out nodes of a path of the coding tree where branching occurs, we can then regroup these equations into a linear system:

$$M\mathbf{x} + M'\mathbf{y} = (1 + \delta)\mathbf{1}. \quad (10)$$

To eliminate  $\mathbf{x}$  from the system above, we need to find a vector  $u$  such that  $u^T M = \mathbf{0}$ , so that we are left with conditions on the perturbation alone:  $u^T M'\mathbf{y} = (1 + \delta)u^T \mathbf{1}$ . Because  $M$  is chosen to be much higher than it is wide, such a vector must exist. The problem is that it might be orthogonal to  $\mathbf{1}$ , in which case premultiplying each side of (10) by  $u$  would give us  $u^T M'\mathbf{y} = 0$ , which would make the perturbation useless. The challenge is how to eliminate  $\mathbf{x}$  without getting rid of  $\delta$  in the process. This is where (8) comes into play, as we exploit the energy dissipation in the system, which translates into exponential decay for the rows of  $M$ . Of course, collisions may involve different values of  $\mathbf{x}$  so we need to add slack variables to (10). This allows us to define exclusion zones for  $\delta$  outside of which the coding tree does not branch too often (most nodes having a single child). To make this work, we need to break up the coding tree into three parts, but these are merely technical issues we can skip here. It is time to switch topics and get to the heart of the analysis: algorithmic renormalization.

## 5 A New Kind of Renormalization

Our line of argumentation so far has relied heavily on a strong contractivity assumption which happens to be invalid. After being constantly reminded of the issue, the reader must surely be wondering what we're going to do about it. On the face of it, it seems a minor technical issue: the Lyapunov exponents at the regular points are all nonpositive, so factoring out the noncontractive modes and applying the previous line of reasoning ought to do the work. Not quite. In fact, this is the biggest thorn in the analysis of diffusive influence systems. We can appreciate why this is no mere technicality from several angles. The most intuitive take on the problem is based on communication. When the graph  $G_t = G(f^t(\mathbf{x}))$  is fixed, linear algebra can model precisely how and when information propagates across the network. When the graph is not fixed but bidirectional (as in Sect. 3), the timing cannot be predicted but progress toward convergence can be measured every time information propagates. The absence of bidirectionality changes everything. We discuss why and what to do about it. We begin with the limitations of linear algebra when the system's eigenmodes are determined by ever-changing subsets of agents. We then highlight the knotty issue of mixed timescales, a phenomenon pervading the dynamics of influence systems. This points the way toward renormalization techniques capable of breaking a system into a (dynamic) hierarchy of subsystems. Operationally, the renormalization is carried out over an infinite sequence of changing networks.

## 5.1 Hopping Modes

It is useful to picture the system as a message-passing mechanism aimed at dissipating energy (an algorithmic variant of the Dirichlet principle). When agent  $i$  averages its position with its neighbors, it receives information about each one. Recall that the information flows *against* the direction of the edges in the communication graph, ie, from  $j$  to  $i$  along the edge  $(i, j)$  in  $G_t$ . Over time, information diffuses among agents across the communication channels. It's a tricky kind of diffusion, however: (i) the domain is a network of arbitrary topology; (ii) it varies with time; (iii) it does so endogenously (via a feedback loop).

Even two agents that never come in contact can communicate by routing the information through intermediate nodes: the communication can be one-way or two-way. Note that even if the graph is not bidirectional, it is still possible for all pairs of agents to be in perpetual two-way communication: take a fixed directed cycle  $G_t = G$ , for example. The main message from Sect. 3 was that bidirectionality enforced fixed-point attraction. Could the condition be relaxed by requiring only perpetual pairwise two-way communication? The answer is no. It is possible for a nonbidirectional system to be trapped in a long periodic orbit although all pairs of agents stay in perpetual two-way communication. Doing away with bidirectionality indeed comes at a price.

To build some intuition, we begin with the case of a fixed graph  $G_t = G$ . Looking at the corresponding stochastic matrix  $P$  as modeling a Markov chain, we partition the agents into essential communicating classes and inessential ones. The first kind form aperiodic, irreducible, hence ergodic, subchains. Their number coincides with the dimension of the principal eigenspace of  $P$  (for the eigenvalue 1), which is also the rank of  $P^\infty$ . We call these agents *dominant*. Each class of dominant agents collapses to single point. Thinking of them as pins attached to the wall, the other agents (the *subdominant* ones) wind up resting in the convex hull of the pins as though attached to them (and possibly one another) by “one-way” springs (ie, pulling only on one end). When the communication graph is fixed, a diffusive influence system behaves like a system of one-way springs. The subdominant agents become “amnesic,” meaning that their final position contains no information about their initial state: indeed, the final barycentric coordinates of these agents can be computed without knowing their initial position. The loss of information coincides with the dissipation of energy. We illustrate this with a concrete example.

*Example 1* We embed the agents in two dimensions ( $d = 2$ ) for visual clarity but we reason with respect to a single set of coordinates so, in effect,  $d = 1$ . Figure 12 shows the initial state  $\mathbf{x}$  of the system (a) and its final configuration (b). There are three essential classes of dominant agents: (1, 2, 4), (5, 7, 9) and the static agent 6, all of which are attracted to the vertices of a triangle (b). When the system reaches equilibrium, the two agents 3, 8 become stable convex combinations of the triangle's vertices. The system is of rank three in the limit: the principal eigenspace is three-dimensional. Agents 3 and 8 are *subdominant*; the seven other agents are *dominant*. In the case of a fixed  $G_t$ , the coding tree consists of a trunk with a nine-dimensional cross-section and no branches. As we go down the tree, the trunk gets squeezed along six dimensions (and not nine) to become three-dimensional: the system is not globally contractive.

If the noncontracting directions were fixed once and for all, we could try to factor them out from the dynamics. Unfortunately, as  $G_t$  changes, so do the names and number of the dominant agents. The principal eigenmodes can hop from agents to agents over time. Our goal is to show that, in the end, the hopping almost surely comes to a halt. The dominant agents end up pinned to the ground asymptotically while, floating within their convex hull, the




Fig. 12 A fixed-graph influence system in its initial and final states

subdominant agents are attracted to limit cycles. Think of a tennis racket being perpetually rewired, where each crossing between two strings forms an agent. This pertains to the system in steady state. Before it gets there, the frame of the racket, under the control of the dominant agents, will undergo all sorts of transformations. These are driven endogenously by changes in status (between dominant and subdominant) among the agents. It is this kind of social mobility that the renormalization algorithm seeks to monitor.

### 5.2 Mixing Timescales


Renormalization consists in decomposing the communication graph dynamically and treating each part as an “independent” dynamical system to be treated recursively. The goal is to break the original influence system into a hierarchy of smaller ones that we can analyze separately. The challenge is how to do this when the hierarchy itself requires perpetual updating. Nothing crystallizes the difficulty better than the thorny issue of mixed timescales. Two decoupled subsystems can operate on widely different relaxation scales so that one is still humming along semi-randomly long after the other has reached equilibrium. This would not be an issue if not for the possibility that the first one should “reawaken” the second one just as it is about to reach equilibrium. The problem is that some agents can serve as long-term memory devices for other ones. The ability of agents to archive past information about other agents and then let them recover it can seriously delay the onset of asymptotic periodicity. We give a simple illustration of this phenomenon.

*Example 2* The picture in the upper-left corner of Fig. 13 shows the communication graph and the agents’ positions at time  $t_a = 0$ . Agent 10 moves according to the rule,

$$x_{10} \leftarrow \frac{1}{2}(x_{10} + x_8), \tag{11}$$

which means that it acquires half of the memory contents of agent 8 in its initial state.<sup>12</sup> From that point on, agent 10 does not communicate with anyone until the remainder of the system has come to rest at time  $t_c$ . Let  $t_b = t_a + 1$  and assume that  $t_c - t_b$  is large enough that agents 1 through 9 have reached equilibrium at time  $t_c$ , ie, their fluctuations are by then minuscule;

<sup>12</sup> For the sake of this simplified explanation, think of granting the same, arbitrarily large number of bits to each agent. Since agent 10 is made to store as much information about agent 8 as about itself, it must dedicate half its information storage to each.



**Fig. 13** By recovering ten percent of its lost memory, agent 8 delays the onset of equilibrium. (The figure does not show the motion of all the agents.)

agent 10 is static. Because agent 8 is subdominant (not being in the support of the principal eigenmode), the information contained in its initial position has evaporated. In other words, the agent is amnesic at time  $t_c$ . Imagine now that agent 8, though barely moving, still manages to cross the margin, which causes it to connect to agent 10 and update itself as

$$x_8 \leftarrow \frac{1}{5}(x_8 + x_3 + x_6 + x_9 + x_{10}).$$

By acquiring one-fifth of the information held by agent 10, agent 8 thus recovers one-tenth of its own lost memory, thereby kicking the system out of equilibrium. Amnesia is the price subdominant agents must pay for falling into an attracting orbit, so recovering memory ruins it all. We might think we are back to square one but that is not quite true: first, agent 8 didn't recover all of its memory but only a fraction of it; second, by acting as an archiving device for agent 8, agent 10 had to move toward the 9-agent subsystem ever so slightly in step  $t_a$ : see (11). If this process were to happen too often, the agent would end getting absorbed into the other agents' orbits. The difficulty is that archiving information for future use in this way can be nested recursively. To deal with this kind of timescale mixing in a principled manner is the motivation behind renormalization.

The key event in Fig. 13 is at time  $t_c$  when agent 8 crosses the margin to change the communication graph. Between  $t_b$  and  $t_c$ , agent 10 is decoupled from the rest and the dynamics of the subsystem 1, . . . , 9 can be resolved separately (and, in fact, recursively, as renormalization dictates). Doing so should yield an upper bound on  $t_c - t_b$  or indicate that it is infinite. We call  $t_c$  a *coupling time*: it indicates a change in the principal modes, with agent 10 being added to the subsystem and promoted to dominant status. This suggests a way to proceed: resolve



subsystem 1, . . . , 9 recursively and establish an upper bound on how long it might take it to communicate again with agent 10 if ever. Note that the two subsystems are only decoupled at the level of the action rule. The communication rules among the first 9 agents between  $t_b$  and  $t_c$  might require knowledge of the position of agent 10: indeed, recall that no locality should be imposed on the communication rules. But at least no state information can be exchanged between the two subsystems. Of course, the graph might remain weakly connected at all times,<sup>13</sup> pointing to the need to broaden our scope further. This is our next topic.

### 5.3 Algorithmic Renormalization

The renormalization is driven by a simple “flooding” procedure: we pick an arbitrary agent and call it *wet*, all the others being *dry*. At the next step, any dry agent that links to it is marked wet—water flows against the direction of the edges. In subsequent steps, any dry agent pointing to a wet one is made wet as well. As soon as all the agents become wet (if ever), they all become dry again and we iterate on the same process. Recall that the graph may change at each step, so the order in which agents become wet may be highly unpredictable. Suppose that, at time  $t_0$ , the set  $A$  of wet agents has no outside agent pointing to them until time  $t_1 > t_0$ . The complement subset  $B$  acquires no information from  $A$  during that period (and so stays dry); therefore, its dynamics can be resolved in isolation from the rest. We extract  $B$  and treat it recursively. This means we apply the flooding algorithm to it in parallel. The case of  $A$  is more complicated because its agents might be acquiring information from  $B$  between  $t_0$  and  $t_1$ , but we also handle it recursively, as shown below. The flooding never ends, so the number of recursive calls is infinite.

To formalize this process, we classify the graphs by their “block-directionality”:  $G(\mathbf{x})$  is of type  $m \rightarrow n - m$  if the agents can be partitioned into  $A, B$  ( $|A| = m$ ) so that no  $B$ -agent ever links to an  $A$ -agent (at least during the time interval of interest); if, in addition, no  $A$ -agent links to any  $B$ -agent,  $G(\mathbf{x})$  is of type  $m \parallel n - m$ . Of course, types change over time. We start out with the  $A$ -agents dry and the  $B$ -agents wet (the latter for notational convenience). Whenever an edge of the (time-varying) communication graph links a dry agent to a wet one, the dry one becomes wet and stays so until explicitly dried. As soon as all the agents become wet (if ever), the  $A$ -agents are dried and the process repeats itself; the  $B$ -agents always remain wet. The interesting events are the long periods in which wetness does not propagate: the wet agents of  $A$  then play the role of subdominant agents being pulled in all directions by both  $B$  and the dry agents of  $A$ .

**Flow tracker**

[1]  $t_0 \leftarrow 0$ .

[2] Repeat:

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

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

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

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

<sup>13</sup> A graph is strongly (weakly) connected if a directed (undirected) path can take us from any node to any other one.

The initial state  $\mathbf{x}$  is assumed fixed for the purpose of this discussion, so we use the shorthand  $G_t = G(f^t(\mathbf{x}))$  to designate the communication graph at time  $t$ . We denote by  $W_t$  the set of wet agents at that time. The flow tracker monitors communication among the agents. 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. Each epoch is itself divided into subintervals by the *coupling times*  $t_1 < \dots < t_\ell$ , such that  $W_{t_k} \subset W_{t_{k+1}}$ . Here is an illustration of this process:

*Example 3* The four-agent system is of type  $3 \rightarrow 1$ , with  $d$  the sole  $B$ -agent and the only one initially wet. We write  $abc \rightarrow d$  instead of  $3 \rightarrow 1$  to make the transitions more transparent.

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

The last column indicates the recursive calls. In row 3 ( $t = 4, 5$ ) we wrote the type  $a \rightarrow bcd$  and not  $ad \rightarrow bc$ . Likewise, in row 5 ( $t = 7, 8, 9$ ), we chose  $ab \rightarrow cd$  over  $abd \rightarrow c$ . The alternative forms are technically correct so what was behind our choice? Inductive soundness is the answer. In statistical mechanics, induction is straightforward. It is based on scales and these point only in one direction. Not so with algorithmic renormalization. For example, it would be erroneous to write  $abd \rightarrow c$  instead of  $ab \rightarrow cd$  because it would renormalize the original type  $3 \rightarrow 1$  into another subsystem *of the same type*, a clear violation of inductive soundness. Whenever wetness propagates (as indicated by the superscript  $w$ ), no decoupling can be assumed and we treat the system as a whole: this lasts only a single step, so these are no more than punctuation marks in the renormalization narrative. The system  $abc \rightarrow d$  is thus renormalized as an infinite sequence of subsystems, each one unfolding between two coupling times. Highlighting the first wave of wetness, we rewrite the system  $abc \rightarrow d$  as

$$(d \parallel abc)_{|3} \circ abcd \circ (a \rightarrow bcd)_{|2} \circ abcd \circ (ab \rightarrow cd)_{|3} \circ abcd \circ \dots \quad (12)$$


The time lengths are successively  $t_1 = 3, t_2 - t_1 - 1 = 2$ , and  $t_3 - t_2 - 1 = 3$ .

If we define the renormalization scale  $w_k = |W_{t_{k+1}}| - n + m$  for  $k = 1, \dots, \ell - 1$ , by analogy with (12), we can express a system of type  $m \rightarrow n - m$  as

$$(m \parallel n - m)_{|t_1} \circ n_{|1} \circ \left\{ \bigcirc_{k=1}^{\ell-1} \left( (w_k \rightarrow n - w_k)_{|t_{k+1}-t_k-1} \circ n_{|1} \right) \right\} \circ \dots, \quad (13)$$

where subscripts indicate the time length of each recursive call and each occurrence of  $n_{|1}$  denotes a case of single-step wetness propagation. Expression (13) is an infinite chronological sequence of subsystems defined recursively. Its parse tree is not of the sort commonly

**Fig. 14** The direct sum/product



encountered in computer science: each node has an infinite number children and every path is infinitely long (there are no leaves). Our main claim implies that, when the system is perturbed, then almost surely the parse tree is not only effectively finite (ie, after pruning it of its trivial parts) but eventually periodic along each (infinite) path. The critical region of “bad” perturbations is likely to be extremely complex, however, and requires a global perspective, one that fixing the initial state  $x$  does not allow. In particular, renormalizing a given orbit as we did in (13) is futile because, as we saw earlier, periodicity is inferred as a global property of the coding tree: fixing our gaze on any one orbit leads nowhere. We must renormalize the *entire* coding tree all at once. In other words, we must look at the geometry of the space of parse trees over the entire phase space and all perturbations  $\delta \in [0, \delta_0]$ . To do this, we need a new language.

A useful analogy might be the set of all sentences (finite and infinite) in the English language. Any given sentence can be parsed recursively so as to yield its own parse tree. On the other hand, the set of all sentences can be represented as a tree (with paths from the root corresponding to common prefixes): that would be the coding tree. Can one use the coding tree to structure the infinite set of parse trees. Note that we are facing two completely different kinds of recursion: the challenge is to unite them seamlessly into one renormalization framework.

Expression (13) tells us how to renormalize a single path of the coding tree. To extend this over the whole coding tree, we need a language that allows us to cobble together the coding tree of a system by assembling the coding trees of its subsystems. Along a given path, the only operation we needed was the concatenation operation  $\circ$ . We now need tensor operations over trees (Fig. 14). The *direct sum* (which we won’t use here) models two decoupled systems operating in parallel in total independence: we take the two coding trees and form a new tree by making new paths out of every pair of paths in the respective trees. The *direct product* models two systems that come in contact and whose coding trees can be glued together with proper pruning. We omit the formal definitions but note that the direct sum is commutative while the direct product is not. The purpose of these building devices is to give us a handle on the entropic growth in the tree.

With each tensor operation comes a formula giving a bound on its entropic effect. In this way, renormalization gives us a way to assemble the coding tree as we might assemble lego blocks while, at each step, monitoring the growth in entropy caused by the corresponding operation. Since this has any meaning only in relation to the energy dissipation of the system, we must also keep track of the coupling times, which tell us how much the subdominant agents

lose memory in the process. We have referred to entropy several times without defining it: the *word-entropy*  $h(\mathcal{T})$  is the growth rate of the language induced by the labels along the paths of the tree (the symbolic dynamics of the system). We can now translate expression (13) in the language of tensors. In the formula below, we must keep in mind that the coupling times are variables that depend on the path followed down the tree. In other words, the direct product  $\mathcal{T} \otimes \mathcal{T}'$  must be interpreted as a set  $\mathcal{T} \otimes \{\mathcal{T}'\}$ , since different coding trees  $\mathcal{T}'$  get attached  $\mathcal{T}$ . By analogy with (13), we can express  $\mathcal{T}_{m \rightarrow n-m}$  as

$$\underline{\mathcal{T}_{m \parallel n-m}}_{|t_1} \otimes \underline{\mathcal{T}_n}_{|1} \otimes \left\{ \bigotimes_{k=1}^{\ell-1} \left( \underline{\mathcal{T}_{w_k \rightarrow n-w_k}}_{|t_{k+1}-t_{k-1}} \otimes \underline{\mathcal{T}_n}_{|1} \right) \right\} \otimes \mathcal{T}_{m \rightarrow n-m}. \tag{14}$$

Note that the factors are optional: for example,  $\mathcal{T}_{m \parallel n-m}$  is empty if the first step shows an edge linking an *A*-agent to a *B*-agent. The case  $m = n$  (no *B*-agent) gets started by designating as wet any one of the *A*-agents. We cannot write these factors as direct sums because the communication rules for the various subgraphs are not decoupled—only the action rules are. The rewriting rule for  $\mathcal{T}_{m \parallel n-m}$  is similar to (14), with the factors  $\mathcal{T}_{w_k \rightarrow n-w_k}$  becoming

$$\mathcal{T}_{(w_k \rightarrow m-w_k) \parallel (w'_k \rightarrow n-m-w'_k)}.$$

Proceeding in this vein, the general factor takes on the more complicated form

$$\mathcal{T}_{(w_{k,1} \rightarrow z_{k,1}) \parallel (w_{k,2} \rightarrow z_{k,2}) \parallel \dots \parallel (w_{k,l} \rightarrow z_{k,l})},$$

where  $\sum_{i=1}^l (w_{k,i} + z_{k,i}) = n$ . Expression (14) describes a single path of the coding tree. The coupling times  $t_k$  and renormalization scales  $w_k$  must be understood as variables since their values depend on the paths with which they are associated. Now the hard work can begin. By combining this recursive expression with the subadditivity of the word-entropy and (9) we can set recurrence relations that allow us to bound the overall entropic growth. It is then just a matter of comparing this growth with the energy loss. We have skipped over the issue of bounding the dissipation rate [18]. This formalism allows us to establish conditions for the asymptotic periodicity of diffusive influence systems under certain perturbative conditions. Algorithmic renormalization plays a key role in the proof and, in light of its relevance to dynamic networks in general, warrants further investigation.

### 6 Bird Flocking

Many models have been proposed to study the flocking of birds in mid-air, most of which fall in the category of nondiffusive influence systems. Typically, the velocity component of the system is diffusive but the positional information, necessary for specifying the communication graph, is not. Although relatively simple flocking models are strikingly realistic to the human eye, visualization hardly counts as serious validation—perhaps we could force birds to watch videos of our simulations and ask them to rate them for realism... Approaching the problem from a theoretical perspective, our goal is not so much to figure out how birds really flock but to discover new analytical methods for “natural algorithms.” Gathering tools is our only goal at this juncture. To do that, we settle on a highly simplified model and focus on an aspect of the dynamics rarely addressed in the flocking literature: geometry.

Theoretical work on flocking draws heavily from statistical physics and agent-based modeling [4, 6, 15, 20, 21, 26, 36, 39, 40]. It is natural to think that size might help. The now-extinct

passenger pigeon was known to form giant flocks of more than two billion birds.<sup>14</sup> In such cases, flocks can show fluid-like behavior and it is natural to turn to hydrodynamics for modeling purposes. At the mesoscopic scale, when the bird group is merely large, techniques from statistical mechanics can be brought to bear to resolve the short timescale behavior. Moving to an agent-based (Lagrangian) model, the idea is to look at birds as moving magnets trying to align their headings under conditions of environmental noise. Statistical mechanics helps us capture an intriguing feature of flocking: how locally-induced communication topology mediates the tension between a bird's tendency to behave like its neighbors and the contrarian effect of its own entropic forces (noise, cognitive errors, random environmental cues, etc). The framework requires the usual assumptions: (i) the topology must be fixed (to have a Hamiltonian); (ii) the system must be at or near equilibrium (to have the Boltzmann distribution); and (iii) it must have an induced metric (to have a correlation function).

Over a short time span, a flock of starlings can appear to meet these criteria. After factoring out the broken symmetry of their common direction of flight, one is left with angular fluctuations reminiscent of a Heisenberg spin system. Indeed, an impressive set of experiments even suggests the appearance of criticality [6, 13, 30]: the idea is that perhaps the presence of long-range scale-free correlations allows for rapid coordination across the entire flock. Unlike magnets, however, the state of a bird requires the integration of all its past velocities. As we raise the "temperature," the system will begin to diffuse apart, rearrange its topology, and break out of equilibrium. Somehow, one must come to grip with the constraints of a dynamic communication network, which is, of course, the province of influence systems.

Our approach to flocking is to focus on this particular aspect in the hope of uncovering some general principles. In that sense, we use flocking as an "excuse" rather than a subject of scientific inquiry. The risk is to end up with little more than a theoretical exercise. While this risk is not negligible, it is tempered by the realization that the field of collective behavior is bereft of native tools: so far most of the techniques used have been borrowed from somewhere else; so perhaps, in this case, the path to discovery begins with theoretical exercises. Our objective is to explore the feedback loop:

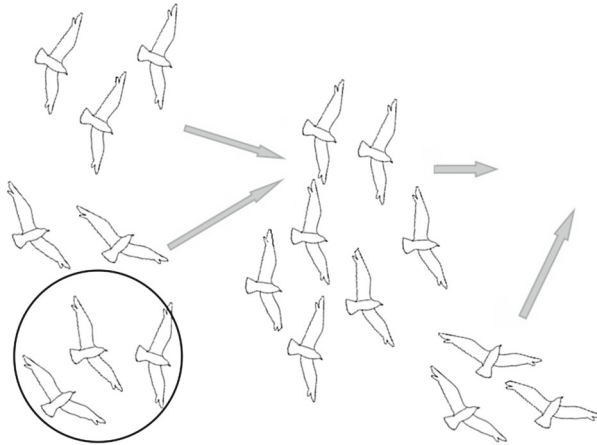
$$\text{velocity} \rightarrow \text{position} \rightarrow \text{topology} \rightarrow \text{communication} \rightarrow \text{velocity}. \quad (15)$$

We have been able to show that the birds always reach equilibrium and to pinpoint the exact relaxation time [15]. The proof is long and we will not attempt to summarize it here. Rather, we wish to focus on two specific aspects related to the geometry of the dynamics. We use the model described in (2), a variant of Vicsek et al. [40] and Cucker-Smale [21]. Birds are joined if their distance is below a fixed threshold. Given  $n$  birds in  $\mathbb{R}^3$ , if  $\mathcal{N}_k$  denotes the set of birds to which bird  $k$  currently links, the dynamics is given by this rule: for  $k = 1, \dots, n$  simultaneously,

$$\begin{cases} v_k \leftarrow \sum_{i \in \mathcal{N}_k} v_i / |\mathcal{N}_k| \\ x_k \leftarrow x_k + v_k. \end{cases}$$

The birds in a connected component form a *flock*. After an initial, usually short period of time during which flocks successively form and split in unpredictable fashion, the group settles on a common velocity and fragmentation stops. From that point on, flocks can only coalesce and eventually that, too, comes to a halt: equilibrium has been reached and the communication graph remains forever fixed. The dynamics of each flock is then that of a damped coupled oscillator with fixed coupling topology. The birds' velocities fluctuate (in both speed and heading) around a mean value associated with their flock. Whereas the fragmentation time

<sup>14</sup> Not a typo.



**Fig. 15** The bird at the center of the circle is influenced by its two neighbors in it. Flock fragmentation stops fairly quickly but flocks can keep coalescing for a long time


is reasonably short, the relaxation time at which all coalescence stops can be very large (yet always finite). It is derived by viewing the group of birds as a distributed computing device and bounding its *busy-beaver* function (Fig. 15).<sup>15</sup> Numerical simulations suggest that convergence is often quite fast. The geometry of bird flocking is fascinating and we briefly discuss two aspects of it.

### 6.1 The Flight Net

Velocities alone mean nothing. Indeed, that they might converge asymptotically toward a fixed vector might not prevent the flocks from getting constantly reconfigured. Discounting the forces pulling nearby birds together, any two straightline trajectories will always diverge unless they are perfectly parallel. Microscopic differences in velocity can thus produce macroscopic positional effects. A key lemma for bounding the relaxation time is a geometric fact: roughly, it asserts that, if two birds drift apart by a distance exceeding by a certain threshold factor the initial diameter of the bird group, they will never rejoin into the same flock in the future. Intuitively, any proof of relaxation must include or imply such a statement. Picture two flocks of birds flying away from each other, with each bird's velocity fluctuating mildly from the mean of its own flock. If the two convex hulls of the sets of velocities (viewed as vectors sharing a common origin) do not intersect then clearly the two flocks can never meet again. The converse is not necessarily true, however: the hulls might overlap at all times while the distance between the two flocks increases to infinity. To see why this cannot be the case is not obvious.

The idea is to argue that the *flight net* always contains a piecewise-linear path that is almost straight. What is the flight net? Add time as an extra dimension so birds now fly in four-dimensional spacetime. Every time a bird averages its velocity with its neighbors, its new velocity lies within the convex hull of these neighbors' velocities. The flight net is a polyhedral complex formed by the entire history of these convex hulls. By going back in time, the flight net grows exponentially: suppose that at time 2, bird 1 had 5 neighbors (including

<sup>15</sup> The busy-beaver function indicates the maximum number of steps a computer can take for a given input and memory size before it halts, assuming that it does.



**Fig. 16** The flight net in spacetime and the combinatorial explosion of its path integrals


itself) and each one had 7 neighbors at time 1; then the flight net of bird 1 at time 2 would consist of  $1 + 5 + 5 \times 7 = 41$  polytopes attached at their vertices by shifting them suitably. Note that a distinct flight net is defined for each pair (bird, time). At time  $t$ , a flight net consists of at most  $(n^{t+1} - 1)/(n - 1)$  polytopes glued together at their vertices (Fig. 16).

Because of the shifts, all flight nets might be wildly different: they correspond to path integrals tracing all “influence” sequences. There is a twist: once a bird is sufficiently settled within its flock, meaning past the mixing time of its corresponding Markov chain, we skip the convex hull construction and encode the trajectory of the bird in the flight net directly. This is a key part of the construction and reflects the fact that birds flying almost in a straight line should stop contributing meaningfully to the path integrals. Figure 16 depicts the mixture of polytopes and polygonal lines constituting a flight net. Not only is it an exponential-size repository of a bird entire velocity history, but it provides a geometric encoding of the locally convex structure of the dynamics. For example, the approach to equilibrium can be read from the curvature of certain paths along the flight net. This plays an important role in bounding the relaxation time [15].

## 6.2 The Spectral Shift

The busy-beaver function of a group of  $n$  birds is a tower-of-twos of height logarithmic in  $n$ : this is also the relaxation time of the system.<sup>16</sup> Astonishingly, this bound is optimal. Of course, no flock will ever behave anywhere near that limit. Nevertheless, this astronomical growth is indicative of a general phenomenon of *spectral shift* which can happen at all time scales. The situation can be compared to the role of chaos in dynamics or the thermodynamic limit in statistical mechanics, where a limit case must be assumed in order to work out the mathematics since, technically, there is no chaos over bounded length scales and no criticality in finite spin systems. These limit assumptions open a window into complex phenomena otherwise beyond reach. Likewise, the spectral shift illustrates how the interaction between position and velocity can cause sudden shifts in scaling.

<sup>16</sup> The tower-of-twos function is defined recursively by:  $f(0) = 1$  and, for  $n > 0$ ,  $f(n) = 2^{f(n-1)}$ .



**Fig. 17** The spectral shift in action with its Goldilocks mixture of symmetries

Here is the basic idea. We say that a flock is *well-mixed* if its communication graph no longer changes and enough time has elapsed for its spectrum to be mostly concentrated around its principal eigenmode (the mean velocity). A flock is a connected component of a graph and the Markov chain it induces has self-loops (since a bird is its own nearest neighbor), so it is ergodic. By Perron-Frobenius, there is a single eigenvalue 1 and all the others are of magnitude less than 1 (the chain is reversible and the spectrum has a full set of real eigenvalues). This unique principal mode corresponds to the mean velocity of the flock. The other eigenmodes represent the damped part of the oscillator and decay at least as fast as  $\lambda^t$ , for  $\lambda < 1 - 1/\text{poly}(n)$ .

Let's set our goal as getting a flock to align itself with the  $Y$  axis as much as possible yet not fly *strictly* parallel to it: for the purpose of this scenario, the birds will fly in a plane  $(X, Y)$ . We start with a well-mixed  $k$ -bird flock having a mean velocity forming an angle  $\theta$  with the  $Y$  axis ( $k = 4$ ). To proceed, we might consider forming the mirror-image of that flock and having the two of them run into each as shown in Fig. 17a. Unfortunately, the symmetry will cause the  $X$  components of the velocities to cancel so that the mean velocity of the new  $2k$ -bird flock will be strictly parallel to the  $Y$  axis: precisely what we tried to avoid! Assuming that the  $Y$  component of each bird's velocity is 1, it remains so at all times so we need only focus on its  $X$  component. For  $i = 1, \dots, k$ , the  $i$ -th bird from the left flies with  $X$ -velocity

$$v_i(t) = c_{i,1} + \sum_{j=2}^k c_{i,j} \lambda_j^t,$$

where each  $\lambda_j$  is real and the spectral gap  $1 - |\lambda_j|$  is at least inverse polynomial in  $k$ . The birds share the same principal eigenmode,  $c_{1,1} = \dots = c_{k,1} = \tan \theta$ , while their mirror-image counterparts on the right have their sign reversed:  $c_{k+1,1} = \dots = c_{2k,1} = -\tan \theta$ . The mirror-image symmetry extends to the other modes:  $c_{i,j} + c_{2k+1-i,j} = 0$ . If the flocks start out separated by a suitable constant, they will coalesce at time roughly  $t_1 = 1/\tan \theta \approx 1/\theta$  for  $\theta > 0$  small enough. The principal left eigenvector  $(u_1, \dots, u_{2k})$  for the  $2k$ -bird flock inherits the mirror-image symmetry, so  $u_i = u_{2k+1-i}$ . This implies that, as shown in Fig. 17a, the lateral ( $X$  component) of the mean velocity of the new flock vanishes:



$$c' = \sum_{i=1}^{2k} u_i v_i(t_1) = 0.$$

To ensure that the new flock flies with a mean velocity nearly—but not perfectly—aligned to the  $Y$  axis, we must ensure that the mean velocities of the two  $k$ -bird flocks cancel but that the other modes do not:

$$c_{1,1} + c_{2k,1} = 0 \quad \text{and} \quad c' = \sum_{i=1}^{2k} u_i v_i(t_1) = \sum_{i=1}^k u_i \sum_{j=2}^k (c_{i,j} + c_{2k+1-i}) \lambda_j^{t_1} > 0. \quad (16)$$

If these conditions holds, as illustrated in Fig. 17b, the angle of the new mean velocity will form a positive angle with the  $Y$  axis at most  $\lambda^{1/\theta}$ , for some  $\lambda < 1 - k^{-b}$  and constant  $b > 0$ . This is an exponential decay. If two such  $k$ -bird flocks are placed at distance 1 of each other, roughly in mirror image position, they will coalesce after a number of steps

$$t_2 \approx 1/\tan \lambda^{1/\theta} \approx e^{t_1/k^b}. \quad (17)$$

Iterating in this process to form flocks of size 4, 8, 16, . . . produce delays growing as a tower of exponentials. We can repeat this at most  $\log n$  times. We have shown that a logarithmic height is the maximum that can be achieved. Coalescing flocks of different sizes cannot increase the delay further because it then becomes impossible to cancel the principal eigenmodes. To prove this entails a highly technical circuit complexity argument [15].

Of course, the entire construction is predicated on the ability to achieve condition (16) again and again as we grow the flocks. Why this is called a *spectral shift* should be clear by now. The coalescence of the two  $k$ -bird flocks wipes out the principal eigenmodes and shift the remaining part of the spectrum to form the modes of the  $2k$ -bird flock. The drastic effect of this maneuver is that by the time ( $t_1$ ) this happens these lower eigenmodes will have lost most of their energy: what is left is exponentially small in  $t_1$ , thus delaying future collisions exponentially (17). To achieve (16) is not obvious, however. Randomization cannot be part of the answer: indeed, this would keep  $c' > 0$  but destroy the cancellation of the principal modes:  $c_{1,1} + c_{2k,1} = 0$ . Instead, we need to inject some symmetry yet not too much. It's a fine energy balance which we achieve with an *origami-like* construction [15].

We note that even a single instance of spectral shift would have dramatic effects on the relaxation time. While being slowed down by a whole tower of exponentials is highly unrealistic (if anything, noise will take care of it), it seems plausible that fragments of the spectral shift occasionally appear in real-life flocking. After all, the phenomenon merely reflects the extreme sensitivity of position to velocity fluctuations. How this mix of geometric and algorithmic ideas extends to more expressive flocking models is an open question.

**Acknowledgments** I wish to thank the anonymous referees for many useful comments and suggestions. This work was supported in part by NSF Grants CCF-0832797, CCF-0963825, and CCF-1016250.

**References**

1. Afek, Y., Alon, N., Barad, O., Hornstein, E., Barkai, N., Bar-Joseph, Z.: A biological solution to a fundamental distributed computing problem. *Science* **331**, 183–185 (2011)
2. Anderson, C.: The end of theory: the data deluge makes the scientific method obsolete. *Wired Magazine*, 23 June 2008
3. Anderson, P.W.: More is different. *Science* **177**, 4047 (1972)
4. 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. Natl. Acad. Sci. USA* **105**, 1232–1237 (2008)
5. Bhattacharyya, A., Braverman, M., Chazelle, B., Nguyen, H.L.: On the convergence of the Hegselmann-Krause system. In: *Proceedings of 4th ITCS*, 2013
  6. Bialek, W., Cavagna, A., Giardina, I., Mora, T., Silvestri, E., Viale, M., Walczak, A.: Statistical mechanics for natural flocks of birds. *Proc Natl. Acad. Sci. USA* **109**, 4786–4791 (2012)
  7. Blondel, V.D., Hendrickx, J.M., Tsitsiklis, J.N.: On Krause's multi-agent consensus model with state-dependent connectivity. *IEEE Trans. Autom. Control* **54**(11), 2586–2597 (2009)
  8. Bullo, F., Cortés, J., Martinez, S.: *Distributed Control of Robotic Networks*. Applied Mathematics Series. Princeton University Press, Princeton (2009)
  9. Bonifaci, V., Mehlhorn, K., Varma, G.: Physarum can compute shortest paths. In: *Proceedings of 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 233–240 (2012)
  10. Buzzi, J.: Piecewise isometries have zero topological entropy. *Ergod. Theory Dyn. Syst.* **21**, 1371–1377 (2001)
  11. Camazine, S., Deneubourg, J.-L., Franks, N.R., Sneyd, J., Theraulaz, G., Bonabeau, E.: *Self-Organization in Biological Systems*. Princeton University Press, Princeton (2003)
  12. Castellano, C., Fortunato, S., Loreto, V.: Statistical physics of social dynamics. *Rev. Mod. Phys.* **81**, 591–646 (2009)
  13. Cavagna, A., Cimarelli, A., Giardina, I., Parisi, G., Santagati, R., Stefanini, F., Viale, M.: Scalefree correlations in starling flocks. *Proc. Natl. Acad. Sci. USA* **107**, 11865–11870 (2010)
  14. Chazelle, B.: Natural algorithms. In: *Proceedings of 20th SODA*, pp. 422–431 (2009)
  15. Chazelle, B.: The convergence of bird flocking. *J. ACM* **61**, 21 (2014)
  16. Chazelle, B.: The total  $s$ -energy of a multiagent system. *SIAM J. Control Optim.* **49**, 1680–1706 (2011)
  17. Chazelle, B.: Natural algorithms and influence systems. *CACM* **55**, 101–110 (2012)
  18. Chazelle, B.: The dynamics of influence systems. *SIAM J. Comput.* (2014). [arXiv:1204.3946v2](https://arxiv.org/abs/1204.3946v2). Preliminary version in *Proceedings of 53rd IEEE FOCS*, 2012
  19. Collins, G.E.: Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In: *Proceedings of 2nd GI Conference in Automata Theory and Formal Languages*, pp. 134–183. Springer-Verlag, New York (1975)
  20. Couzin, I.D., Krause, J.: Self-organization and collective behavior in vertebrates. *Adv. Study Behav.* **32**, 1–75 (2003)
  21. Cucker, F., Smale, S.: Emergent behavior in flocks. *IEEE Trans. Autom. Control* **52**, 852–862 (2007)
  22. Feinerman, O., Korman, A.: Memory lower bounds for randomized collaborative search and implications to biology. In: *DISC*. Springer, Berlin (2012)
  23. Fisher, J., Harel, D., Henzinger, T.A.: Biology as reactivity. *Commun. ACM* **54**, 72–82 (2011)
  24. Hegselmann, R., Krause, U.: Opinion dynamics and bounded confidence models, analysis, and simulation. *J. Artif. Soc. Soc. Simul.* **5**, 3 (2002)
  25. Hendrickx, J.M., Blondel, V.D.: Convergence of different linear and non-linear Vicsek models. In: *Proceedings of 17th International Symposium on Mathematical Theory of Networks and Systems (MTNS2006)*, Kyoto, Japan, July 2006, pp. 1229–1240.
  26. Jadbabaie, A., Lin, J., Morse, A.S.: Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Trans. Autom. Control* **48**, 988–1001 (2003)
  27. Krause, U.: A discrete nonlinear and non-autonomous model of consensus formation. In: *Proceedings of Communications in Difference Equations*, pp. 227–236 (2000)
  28. Kruglikov, B., Rypdal, M.: A piecewise affine contracting map with positive entropy. *Discrete Continuous Dyn. Syst.* **16**(2), 393–394 (2006)
  29. Lorenz, J.: A stabilization theorem for dynamics of continuous opinions. *Phys. A* **355**, 217–223 (2005)
  30. Mora, T., Bialek, W.: Are biological systems poised at criticality? *J. Stat. Phys.* **144**, 268–302 (2011)
  31. Mézard, M., Montanari, A.: *Information, Physics, and Computation*. Oxford University Press, New York (2009)
  32. Moreau, L.: Stability of multiagent systems with time-dependent communication links. *IEEE Trans. Autom. Control* **50**, 169–182 (2005)
  33. Navlakha, S., Bar-Joseph, Z.: Algorithms in nature: the convergence of systems biology and computational thinking. *Mol. Syst. Biol.* **7**, 546 (2011)
  34. Okubo, A., Levin, S.A.: *Diffusion and Ecological Problems*. Modern Perspectives, 2nd edn. Springer, Berlin (2001)
  35. Olshevsky, A., Tsitsiklis, J.N.: On the nonexistence of quadratic Lyapunov functions for consensus algorithms. *IEEE Trans. Autom. Control* **53**, 2642–2645 (2008)
  36. Reynolds, C.W.: Flocks, herds, and schools: a distributed behavioral model. *Comput. Graph.* **21**, 25–34 (1987)

37. Seneta, E.: *Non-negative Matrices and Markov Chains*, 2nd edn. Springer, New York (2006)
38. Sontag, E.D.: Nonlinear regulation: the piecewise linear approach. *IEEE Trans. Automat. Control* **26**, 346–358 (1981)
39. Toner, J., Tu, Y.: Flocks, herds, and schools: a quantitative theory of flocking. *Phys. Rev. E* **58**, 4828–4858 (1998)
40. Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., Shochet, O.: Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett.* **75**, 1226–1229 (1995)