The Convergence of Bird Flocking

BERNARD CHAZELLE, Princeton University

We bound the time it takes for a group of birds to stabilize in a standard flocking model. Each bird averages its velocity with its neighbors lying within a fixed radius. We resolve the worst-case complexity of this natural algorithm by providing asymptotically tight bounds on the time to equilibrium. We reduce the problem to two distinct questions in computational geometry and circuit complexity.

1. INTRODUCTION

Bird flocking has received considerable attention in the scientific and engineering literature, beginning with the classic *Boids* framework of Reynolds [Camazine et al. 2003; Reynolds 1987; Tanner et al. 2003a, 2003b, 2007]. All models seem to follow a subset of these three rules: the birds try constantly to (i) align their headings, (ii) stay grouped together, and (iii) avoid collision. For reasons that are not yet completely understood, this model produces spectacular visuals; as a result, most bird flocking animations in Hollywood are based on it. Of all three rules, the first one drives the dynamics, so it is customary for theoretical investigations to ignore the other two. Indeed, for simplicity, the bulk of the mathematical work on bird flocking has focused on variants of rule (i) [Chazelle 2009; Cucker and Smale 2007; Hendrickx and Blondel 2006; Jadbabaie et al. 2003; Ji and Egerstedt 2007; Li and Wang 2004; Moreau 2005; Moshtagh et al. 2005; Olfati-Saber 2006; Shi et al. 2005; Tahbaz-Salehi and Jadbabaie 2007; Tang and Guo 2007; Tanner et al. 2003a, 2003b, 2007; Toner and Tu 1995; Vicsek et al. 1995]. Continuing in that vein, we prove convergence and bound the time to equilibrium. The distinctive feature of our result is that it requires no assumption about the connectivity of the flocking network.

Our model is a variant of the one proposed by Cucker and Smale [2007], which is itself a holonomic variant of the classical Vicsek model [Vicsek et al. 1995]. Given n birds in $\mathbb{R}^3$ represented by their positions and velocities, that is, by a total of $6n$ numbers, the *flocking network* has a node for each bird and an edge between any pair
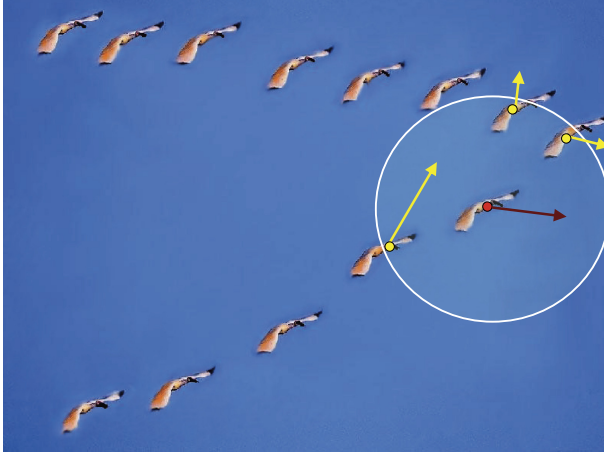
Fig. 1. Each bird updates its velocity by averaging it with those of its neighbors within a unit-radius circle.

at distance one or less (Figure 1). The connected components of the graph (undirected, with self-loops) are called the *flocks* of the system. The time t is discrete, so the position of a bird i at time t is given by its position x_i at time $t-1$ shifted by its current velocity v_i; that is, x_i becomes $x_i + v_i$. To update v_i, the new velocity of bird i is formed as the average of those of its neighbors in the flocking network. The averaging can be weighted, with the weights changing over time if so desired; even a moderate amount of decaying noise can be tolerated as well without affecting our results.

The model is a particular case of a class of natural algorithms known as nondiffusive influence systems [Chazelle 2012a, 2012b]. We show that it always converges and we bound the maximum time it takes to reach equilibrium. We do this in two steps. First (the difficult part), we prove that the flocking network converges to a fixed graph in finite time; second (the easy part), we show why this implies equilibrium in the physical sense: each flock flying at constant speed in a fixed direction and forming a damped coupled oscillator decaying exponentially fast. To prove that the flocking network stabilizes is the heart of this article. We also do it in two steps. We show that the network must at some stage cease to fragment. From that point on, the flocks can no longer split but they can coalesce (obviously at most $n-1$ times). The last such coalescing event, also the last nonlinearity in the system, occurs within a number of steps equal to a tower-of-twos of height logarithmic in n. We first bound the height by $O(n)$ using tools from computational geometry. Then we use circuit complexity arguments to reduce the bound to $O(\log n)$. We have shown that a tower-of-twos of logarithmic height is essentially optimal [Chazelle 2009]. While the flocking network may take very long to stabilize, we show that the birds themselves reach their final flying direction within a number of steps that is a single exponential in their number.

The model requires infinite precision. This may be seen as unrealistic but the astronomical time scale of equilibrium reflects a real phenomenon, called a *spectral shift*. As a result of an external event (here a near-collision), the spectrum of a flock finds itself shifting sideways and losing the energy of its lowest mode. Traces of this phenomenon have been reported in condensed matter physics, for example, in the transport of impurities in an ultra-cold Fermi sea [Kim and Huse 2012]. The infinite precision assumption is an idealization which does not necessarily invalidate the results in practice. A good analogy would be chaos, which mathematically requires infinite precision, yet models physical phenomena of high relevance even with finite precision.

Previous Work. The choice of distance-based criteria was validated recently by the empirical findings of the STARFLAG project, the most comprehensive experimental bird flocking investigation to date [Ballerini et al. 2008; Cavagna et al. 2008a, 2008b, 2010]. Computer simulations also support the intuitive belief that, by repeated averaging, each bird should eventually converge to a fixed speed and heading. To prove this, until now, had required various assumptions: recurrent connectivity, for example, which stipulates that, over any time interval of a fixed length, every pair of birds should be able to communicate with each other, directly or indirectly via other birds. Jadbabaie et al. [2003] proved the first of several convergence results under that assumption (or related ones [Moshtagh et al. 2005; Olfati-Saber 2006; Shi et al. 2005; Tanner et al. 2003b]). Several authors have extended these results to variable-length intervals [Hendrickx and Blondel 2006; Li and Wang 2004; Moreau 2005]. Others have shown how to do away with the recurrent connectivity assumption by changing the model suitably. Tahbaz-Salehi and Jadbabaie [2007], for example, assume that the birds fly on the surface of a torus to achieve spatial ergodicity. Cucker and Smale [2007] use a broadcast model that extends a bird's influence to the entire group while scaling it down as a function of distance. In a similar vein, Ji and Egerstedt [2007] introduce a hysteresis rule to ensure that connectivity increases over time. Tang and Guo [2007] prove convergence in a high-density probabilistic model.

The Model and The Results. The n birds $B_1, \ldots, B_n$ are represented at time t by their position $x(t) = (x_1(t), \ldots, x_n(t))$ and velocity $v(t) = (v_1(t), \ldots, v_n(t))$. We view each $x_i(t)$ as a three-dimensional row vector so that $x(t)$ is an n-by-3 real-valued matrix and the same is true of $v(t)$. The input consists of the initial position $x(0)$ and velocity $v(1)$, all of whose entries are rationals over $O(\log n)$ bits. The *flocking network* G_t has a vertex for each bird and an edge between any two of them within unit distance of each other (Figure 1).[1] The velocity matrices obey the dynamics

$$v_i(t+1) - v_i(t) = c_i(t) \sum_{(i,j) \in G_t} (v_j(t) - v_i(t)),$$

with $c_i(t) d_i(t) < 1$ for $d_i(t) = |\{j \neq i | (i,j) \in G_t\}|$. (This equation can be derived from Langevin dynamics by defining a suitable Hamiltonian [Bialek et al. 2012].) We assume that $c_i(t)$ may vary only when G_t does. The standard setting is $c_i(t) = 1/(d_i(t)+1)$ but we prefer to keep this choice open. In matrix form, for any $t \geq 1$,

$$\begin{cases} x(t) = x(t-1) + v(t); \\ v(t+1) = P_t v(t), \end{cases} \tag{1}$$

where $P_t = \mathbf{I}_n - C_t L_t$, $C_t = \text{diag}(c_i(t))$, and L_t denotes the graph Laplacian.[2] We define the tower-of-twos function as follows: $2 \uparrow\uparrow n = 2^{2 \uparrow\uparrow (n-1)}$ for $n > 1$ and $2 \uparrow\uparrow 1 = 2$.

THEOREM 1.1. *A group of n birds reaches equilibrium in fewer than $2 \uparrow\uparrow (4 \log n)$ steps. The velocity of each bird converges to a stationary vector with rational coordinates. After $n^{O(n^2)}$ steps, network edges can no longer disappear.*

For simplicity, the stated height on the tower-of-twos assumes that n exceeds a suitable constant. Equilibrium means that the flocking network has become static and

---

[1]We should be writing G_t as $G(x(t))$ but fixing the initial conditions once and for all allows us to simplify the notation.

[2]The Laplacian of an n-node undirected graph G is the n-by-n matrix L, where L_{ii} is the number of neighbors of node i distinct from itself and, for $i \neq j$, L_{ij} is -1 if $(i,j) \in L$ and 0, otherwise. We assume that all entries are rationals encoded over $O(\log n)$ bits.

that the birds' deviation from a constant-speed, straightline trajectory decays exponentially fast.[3] We note that convergence *requires* a perturbation rule to prevent a flock from breaking apart because of microscopic motion. Formally, an edge (i,j) of G_t remains in G_{t+1} if the distance between B_i and B_j changes by less than n^{-bn^2} at step t, for some large $b > 0$. (Any small enough threshold works.) The rule ensures *completeness*: any two birds within distance 1 are always linked in G_t. As we show here, it also guarantees *soundness*: no birds are ever joined by an edge if their distance exceeds $1 + o(1)$. We prove that the perturbation rule is necessary for convergence.

This article is organized into three parts. We investigate the case of a single time-invariant flock in Section 2; we bound the number of network switches in Section 3 and prove both the soundness and necessity of the perturbation rule; we investigate the geometry of bird flocking in Section 4; we bound the flocking time by a tower-of-twos of linear height in Section 5; finally, we prove in Section 6 that the height is actually logarithmic.

2. STATIONARY FLOCKS

Birds spend most of their time in time-invariant flocks. These are damped coupled oscillators which converge exponentially fast: the principal mode sets the flock's limiting velocity while the subdominant eigenvalues quickly squelch the oscillations. We indicate time invariance by writing $G_t = G$ and $P = P_t = P$. Without loss of generality, we can assume that G has a single connected component. The transition matrix P has the simple dominant eigenvalue 1 with right and left eigenvectors $\mathbf{1}$ and $C^{-1}\mathbf{1}$, respectively. We can normalize the latter to form the probability distribution

$$\pi = \frac{1}{\operatorname{tr} C^{-1}} C^{-1} \mathbf{1}.$$

We define

$$M = C^{-1/2}PC^{1/2} = C^{-1/2}(\mathbf{I}_n - CL)C^{1/2} = \mathbf{I}_n - C^{1/2}LC^{1/2}. \tag{2}$$

Being symmetric, M can be diagonalized as $\sum_{k=1}^n \lambda_k u_k u_k^T$, where the u_k's are orthonormal eigenvectors and the eigenvalues are real. It follows that P can be diagonalized as well, with the same eigenvalues. By Perron-Frobenius theory and standard properties of random walks [Chung 1997; Seneta 2006], $1 = \lambda_1 > \lambda_2 \geq \cdots \geq \lambda_n > -1$ and $u_1 = (\sqrt{\pi_1}, \ldots, \sqrt{\pi_n})^T$. Since $\sum_k u_k u_k^T = \mathbf{I}_n$, the following identity holds for all nonnegative t:

$$P^t = C^{1/2}M^tC^{-1/2} = \mathbf{1}\pi^T + \sum_{k=2}^n \lambda_k^t C^{1/2}u_k u_k^T C^{-1/2}. \tag{3}$$

Let $\mu = \max_{k>1} |\lambda_k|$ denote the magnitude of the second largest eigenvalue. Using classical spectral gaps from algebraic graph theory [Landau and Odlyzko 1981], we find that

$$\mu \leq 1 - n^{-O(1)}. \tag{4}$$

We give a self-contained proof in Chazelle [2009]. By (3), for all $i, j, t > 0$,

$$(P^t)_{ij} \geq \pi_j - \sum_{k>1} |\lambda_k|^t \sqrt{c_i/c_j}\, |(u_k)_i (u_k)_j| \geq \pi_j - n^{O(1)}\mu^t.$$

---

[3]The same results apply in a noisy model whose details are given in full in our original manuscript [Chazelle 2009]. The matching lower bound on the convergence time is also given there.

It follows from (4) that[4]

$$\|P^t - \mathbf{1}\pi^T\|_F \le e^{-tn^{-O(1)}+O(\log n)}.$$ (5)

The locations of the birds remain rational at all times. We must show that this remains true in the limit. For $t > 0$, we define

$$\Gamma_t = -\mathbf{1}\pi^T t + \sum_{s=0}^{t-1} P^s.$$ (6)

It is immediate that Γ_t converges to some matrix Γ, as t goes to infinity. Indeed, by (3),

$$\Gamma = \sum_{s \ge 0}(P^s - \mathbf{1}\pi^T) = \sum_{k>1} \tfrac{1}{1-\lambda_k} C^{1/2} u_k u_k^T C^{-1/2}.$$

What is perhaps less obvious is why the limit is rational (although the eigenvalues need not be). We begin with a simple characterization of Γ, which we derive from standard arguments about the fundamental matrix of a Markov chain [Kemeny and Snell 1983]. We use the notation $(Y \,|\, y)$ to refer to the n-by-n matrix derived from Y by replacing its last column with the vector y.

LEMMA 2.1. $\Gamma = -\mathbf{1}\pi^T + (\mathbf{I}_n - P + \mathbf{1}\pi^T)^{-1} = (\mathbf{I}_n - \mathbf{1}\pi^T \,|\, \mathbf{0})\,(\mathbf{I}_n - P \,|\, \mathbf{1})^{-1}$.

PROOF. Because $\mathbf{1}$ and π are, respectively, right and left eigenvectors of P for the eigenvalue 1, for any integer $s > 0$,

$$(P - \mathbf{1}\pi^T)^s = P^s - \mathbf{1}\pi^T,$$ (7)

which follows from the commutativity of the product $P \times \mathbf{1}\pi^T$ and the identity

$$(P - \mathbf{1}\pi^T)^s = P^s + \sum_{k=0}^{s-1}(-1)^{s-k}\binom{s}{k}P^k(\mathbf{1}\pi^T)^{s-k}$$

$$= P^s + (\mathbf{1}\pi^T)\sum_{k=0}^{s-1}(-1)^{s-k}\binom{s}{k} = P^s - \mathbf{1}\pi^T.$$

And so, for $t > 1$,

$$\Gamma_t + \mathbf{1}\pi^T = \mathbf{I}_n + \sum_{s=1}^{t-1}(P^s - \mathbf{1}\pi^T) = \sum_{s=0}^{t-1}(P - \mathbf{1}\pi^T)^s.$$

Premultiplying this identity by the "denominator" that we expect from the geometric sum, ie, $\mathbf{I}_n - P + \mathbf{1}\pi^T$, we simplify the telescoping sum, using (7) again,

$$(\mathbf{I}_n - P + \mathbf{1}\pi^T)(\Gamma_t + \mathbf{1}\pi^T) = (\mathbf{I}_n - P + \mathbf{1}\pi^T)\sum_{s=0}^{t-1}(P - \mathbf{1}\pi^T)^s$$

$$= \mathbf{I}_n - (P - \mathbf{1}\pi^T)^t = \mathbf{I}_n - (P^t - \mathbf{1}\pi^T).$$

---

[4]The Frobenius norm $\|M\|_F$ of a matrix is the Euclidean norm of the vector formed by its elements. It is submultiplicative and $\|Mu\|_2 \le \|M\|_F \|u\|_2$.

By (3), P^t converges to $\mathbf{1}\pi^T$ as t goes to infinity, so $(\mathbf{I}_n - P + \mathbf{1}\pi^T)(\Gamma_t + \mathbf{1}\pi^T)$ converges to the identity. This implies that, for t large enough, the matrix cannot be singular and, hence, neither can $\mathbf{I}_n - P + \mathbf{1}\pi^T$. This allows us to write:

$$\Gamma + \mathbf{1}\pi^T = (\mathbf{I}_n - P + \mathbf{1}\pi^T)^{-1},$$

which proves the lemma's first identity. For the second one, we argue that, since π is a left eigenvector of P for 1, $\mathbf{1}\pi^T(\mathbf{I}_n - P) = \mathbf{0}$; hence, for $t > 0$,

$$\mathbf{I}_n - P^t = (\mathbf{I}_n + P + \cdots + P^{t-1})(\mathbf{I}_n - P) = (\Gamma_t + \mathbf{1}\pi^T t)(\mathbf{I}_n - P) = \Gamma_t (\mathbf{I}_n - P).$$

As $t \to \infty$, $P^t \to \mathbf{1}\pi^T$; therefore, $\Gamma (\mathbf{I}_n - P) = \mathbf{I}_n - \mathbf{1}\pi^T$. Since $\mathbf{1}$ lies in the kernel of Γ_t, and hence of Γ, the latter matrix satisfies the relation

$$\Gamma (\mathbf{I}_n - P \,|\, \mathbf{1}) = (\mathbf{I}_n - \mathbf{1}\pi^T \,|\, \mathbf{0}). \tag{8}$$

The simplicity of P's dominant eigenvalue 1 implies that $\mathbf{I}_n - P$ is of rank $n - 1$. Since $\mathbf{1} \in \ker (\mathbf{I}_n - P)$, the last column of $\mathbf{I}_n - P$ is the negative sum of the others; so to get the correct rank the first $n - 1$ columns of $\mathbf{I}_n - P$ must be independent. Note that the vector $\mathbf{1}$ is not in the space they span: if, indeed, it were, we would have $\mathbf{1} = (\mathbf{I}_n - P)y$, for some $y \in \mathbb{R}^n$. Since $\pi^T(\mathbf{I}_n - P) = \mathbf{0}$, this would imply that $1 = \pi^T\mathbf{1} = \pi^T(\mathbf{I}_n - P)y = 0$, a contradiction. This is evidence that $(\mathbf{I}_n - P \,|\, \mathbf{1})$ is of full rank, which, by (8), completes the proof. □

The equation of motion (1) becomes, for $t \geq 1$,

$$x(t) = x(0) + \sum_{s=0}^{t-1} P^s v(1) \tag{9}$$

or, equivalently, by (6),

$$x(t) = x(0) + t\,\mathbf{1}\pi^T v(1) + \Gamma_t v(1). \tag{10}$$

We call $\mathbf{m}_\pi[x(t)] \overset{\text{def}}{=} \pi^T x(t)$ the *mass center* of the flock and the vector $\mathbf{m}_\pi[v(1)]$ its *stationary velocity*; both are three-dimensional row vectors. The mass center drifts in space at constant speed along a fixed line: indeed, $\pi^T \Gamma_t = \mathbf{0}$; so by (10),

$$\mathbf{m}_\pi[x(t)] = \mathbf{m}_\pi[x(0)] + t\,\mathbf{m}_\pi[v(1)]$$

and

$$x(t) = \underbrace{x(0)}_{start} + \underbrace{t\,\mathbf{1}\mathbf{m}_\pi[v(1)]}_{linear\ drift} + \underbrace{\Gamma_t v(1)}_{damped\ oscillator}. \tag{11}$$

The oscillations are damped at a rate of $e^{-tn^{-O(1)}}$. Moving the origin to the mass center of the birds, we express $x(t)$, relative to this moving frame, as

$$x^r(t) \overset{\text{def}}{=} x(t) - \mathbf{1}\,\mathbf{m}_\pi[x(t)];$$

therefore,

$$x(t) = x^r(t) + \mathbf{1}\pi^T x(0) + t\,\mathbf{1}\pi^T v(1); \tag{12}$$

and, by (10),

$$x^r(t) = (\mathbf{I}_n - \mathbf{1}\pi^T)x(0) + \Gamma_t v(1)$$

and, by Lemma 2.1,

LEMMA 2.2. *If G is connected, as t tends to infinity,*

$$x^r(t) \to x^r \overset{\text{def}}{=} (\mathbf{I}_n - \mathbf{1}\pi^T)x(0) + \Gamma\, v(1).$$

The mass center of the flock moves in $\mathbb{R}^3$ at constant speed in a fixed direction.

LEMMA 2.3. *The elements of Γ and the coordinates of the limit configuration x^r can be expressed as $O(n \log n)$-bit rationals with a common denominator.*

PROOF. Let C_b denote the $O(n \log n)$-bit long product of all the denominators in the diagonal matrix C. The determinant of $(CL \,|\, \mathbf{1})$ can be expressed as C_b^{-1} times the determinant N of an n-by-n matrix with $O(\log n)$-bit integer elements. By the Hadamard bound [Yap 2000], N is an $O(n \log n)$-bit integer. For the same reason, each element of $\text{adj}\,(CL \,|\, \mathbf{1})$ is also the product of C_b^{-1} with an $O(n \log n)$-bit integer; therefore,

$$(\mathbf{I}_n - P \,|\, \mathbf{1})^{-1} = (CL \,|\, \mathbf{1})^{-1} = \frac{\text{adj}\,(CL \,|\, \mathbf{1})}{\det\,(CL \,|\, \mathbf{1})}$$

is of the form N^{-1} times an $O(n \log n)$-bit integer matrix (since the two appearances of C_b^{-1} cancel out). The same is true of $(\mathbf{I}_n - \mathbf{1}\pi^T \,|\, \mathbf{0})$, which follows from looking at the product $\pi^T(\mathbf{I}_n - P \,|\, \mathbf{1})$ and observing that $\pi^T = (0, \ldots, 0, 1)(\mathbf{I}_n - P \,|\, \mathbf{1})^{-1}$. Therefore, both $(\mathbf{I}_n - \mathbf{1}\pi^T \,|\, \mathbf{0})$ and $(\mathbf{I}_n - P \,|\, \mathbf{1})^{-1}$ can be expressed with entries that are $O(n \log n)$-bit rationals sharing a common denominator. The proof follows from Lemmas 2.1 and 2.2. □

3. SWITCHES AND PERTURBATIONS

We define the *relative displacement* $\Delta_{ij}(t) = |\text{DIST}_t(B_i, B_j) - \text{DIST}_{t-1}(B_i, B_j)|$ between the two birds B_i and B_j, where $\text{DIST}_t(B_i, B_j) = \|x_i(t) - x_j(t)\|_2$ denotes the distance between them. By the triangle inequality,

$$\|x_i(t) - x_j(t)\|_2 \leq \|x_i(t-1) - x_j(t-1)\|_2 + \|x_i(t) - x_i(t-1) - (x_j(t) - x_j(t-1))\|_2 .$$

Reversing the roles of t and $t-1$ gives us a similar inequality, from which we find that

$$|\text{DIST}_t(B_i, B_j) - \text{DIST}_{t-1}(B_i, B_j)| \leq \|x_i(t) - x_i(t-1) - (x_j(t) - x_j(t-1))\|_2 ;$$

hence, for $t \geq 1$,

$$\Delta_{ij}(t) \leq \|v_i(t) - v_j(t)\|_2. \tag{13}$$

Consider the line embedding of the graph sequence $(G_t)_{|t \geq 1}$ obtained by placing node i at time t at position $(v_i(t))_1$. The number of times the embedded graph G_t has at least one edge of length greater than λ is called the *communication count* C_λ: by appealing to the *total s-energy*,[5] we have shown [Chazelle 2011] that $C_\lambda \leq (\log(M/\lambda))^{n-1}\rho^{-n^2 - O(1)}$, where ρ is a lower bound on the nonzero entries of P_t and M is the initial diameter of the embedded nodes; here, $\rho = n^{-O(1)}$ and $M = n^{O(1)}$. By embedding the graph sequence along all three coordinate axes, it follows that the number of times two birds are joined in G_t with a velocity difference of ℓ_∞-norm in excess of $\frac{1}{2}n^{-bn^2}$ is $(bn)^{O(n^2)}$. At all other times, by (13), any two birds joined in G_t are subject to a relative displacement at most n^{-bn^2}; hence, by the perturbation rule, remain linked at the next step. Since the number of edge losses cannot differ from the number of edge gains by more than $\binom{n}{2}$, we conclude:

---

[5]The total s-energy is the sum of the sth powers of the edge lengths over all G_t ($t > 0$). It is defined for any real $s > 0$.

LEMMA 3.1. *The number of network switches, that is, times t at which the flocking network G_t changes, is bounded by $(bn)^{O(n^2)}$.*

Why Perturbations Are Necessary. The perturbation rule can err only in one direction: by keeping birds linked when they should not. We show this is not an issue: edges can link only birds that are apart by at most $1 + o(1)$. But first we prove that the perturbation is indeed required for convergence. Our example consists of four birds flying in a plane, which maintain constant unit velocity in the Y-direction. It thus suffices to specify the dynamics along the X-axis.

$$\begin{cases} x(0) = \frac{1}{16}(0, 8, 21, 29); \\ v(1) = \frac{1}{8}(1, -1, 1, -1). \end{cases}$$

The flocking network alternates between a pair of 2-bird edges and a single 4-bird path, whose respective transition matrices are given by setting the diagonal of C_t to $\frac{2}{3}(1, 1, 1, 1)$ and $\frac{1}{3}(2, 1, 1, 2)$, which gives us

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

The initial X-velocity $v(1)$ is a right eigenvector (for $-1/3$) for both flocking networks; therefore, for $t > 0$, $v(t) = (-3)^{1-t}v(1)$ and, by (1),

$$x(t) = x(0) + \sum_{s=1}^{t} v(s) = x(0) + \frac{3}{4}\left(1 - (-\frac{1}{3})^t\right)v(1);$$

hence,

$$x_{i+1}(t) - x_i(t) = \begin{cases} \frac{1}{16}(5 - (-\frac{1}{3})^{t-1}) & \text{if } i = 1, 3; \\ 1 + \frac{1}{16}(-\frac{1}{3})^{t-1} & \text{if } i = 2. \end{cases}$$

The distance between the middle birds B_2 and B_3 oscillates around 1, so the network forever alternates between one and two connected components. A perturbation rule is therefore needed for the flocking network to stabilize.

Why the Perturbation Rule Is Sound. We now show that only birds within distance $1 + o(1)$ of each other can be linked to each other in the flocking network.

LEMMA 3.2. *Any two birds adjacent in G_t are a distance $1 + n^{-bn^2/2}$.*

PROOF. Write $\varepsilon = n^{-bn^2}$ and $\delta = \sqrt{\varepsilon}$. Assume by contradiction that two birds B_i and B_j are at distance greater than $1 + \delta$ at time $t_1 > t_0$ and that, at any $t_0 \leq t \leq t_1$, the distance remains in the interval $(1, 1 + \delta]$ while the two birds are joined in G_t (Figure 2). The perturbation rule implies that $\Delta_{ij}(t) < \varepsilon$, for $t_0 < t \leq t_1$. Consider the $t_1 - t_0$ relative displacements in the time interval $[t_0, t_1]$. Together, they create a displacement in excess of δ. Mark the unit-time intervals within $[t_0, t_1]$ that are
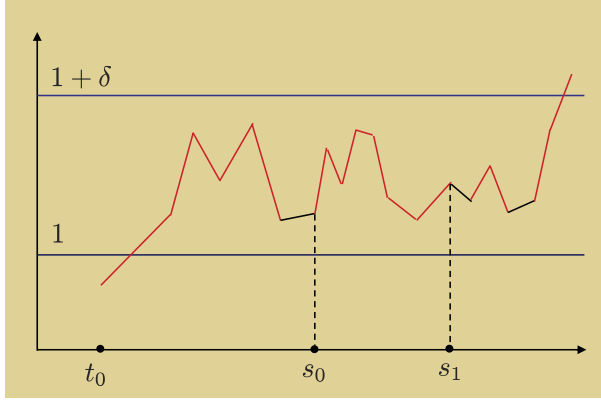
Fig. 2. The soundness of the perturbation rule.

associated with network switches and let N be their number: N runs of consecutive unmarked unit-time intervals together contribute a relative displacement in excess of $\delta - \varepsilon N$, so one of them must contribute at least $(\delta - \varepsilon N)/(N + 1)$. If $[s_0, s_1]$ denotes the corresponding time interval ($t_0 \leq s_0 \leq s_1 \leq t_1$), then G_t remains invariant for all $s_0 \leq t \leq s_1$ and, by Lemma 3.1,

$$\sum_{t=s_0+1}^{s_1} \Delta_{ij}(t) \geq \frac{\delta - \varepsilon N}{N+1} \geq \delta (bn)^{-O(n^2)}. \tag{14}$$

We now show that this displacement is too large for two birds trapped in the same time-invariant flock. The edge (i,j) is in the network G_t for all $t \in [s_0, s_1]$, so the two birds B_i and B_j are in the same flock during that time period. By (5, 13), for $s_0 < t \leq s_1$,

$$\Delta_{ij}(t) \leq \|v_i(t) - v_j(t)\|_2 \leq 2^{-(t-s_0)n^{-O(1)}+O(\log n)}.$$

By (14) and the perturbation rule,

$$\delta (bn)^{-O(n^2)} \leq \sum_{t=s_0+1}^{s_1} \Delta_{ij}(t) \leq \sum_{t=s_0+1}^{s_1} \min \left\{ \varepsilon, 2^{-(t-s_0)n^{-O(1)}+O(\log n)} \right\}$$

$$\leq \min_{T>0} \{ T\varepsilon + 2^{-Tn^{-O(1)}+O(\log n)} \}.$$

Setting $T = n^c \log \frac{1}{\varepsilon}$ for a large enough constant c leads to a contradiction:

$$\delta \leq (bn)^{O(n^2)} \varepsilon \log \frac{1}{\varepsilon} + (bn)^{O(n^2)} \varepsilon < \sqrt{\varepsilon}. \qquad \square$$

4. THE GEOMETRY OF FLOCKING

We have bounded the number of network switches, proved the necessity and soundness of the perturbation rule, and resolved the case of a single flock. What is left to do? The remaining issue is to show that distinct flocks cannot take too long to collide if they must. To begin with, we show that, after enough time has elapsed, two birds can be newly joined only if they fly almost parallel to each other. We also prove that they cannot stray too far from each other if they want to get together again in the future.

Fig. 3. The flight net is formed by joining together the convex polytopes associated with the birds' updated velocities.

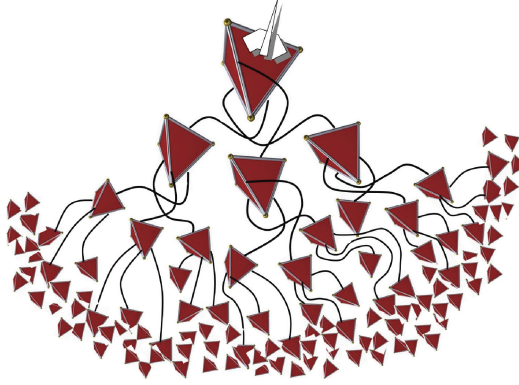We investigate the geometric structure of flocking and, to help us do so, we introduce a useful device, the *flight net* (Figure 3). It is convenient to lift the birds to $\mathbb{R}^4$ by adding time as an extra dimension,

$$x_i(t) = (x_{i,1}(t), x_{i,2}(t), x_{i,3}(t), t),$$

and doing the same for the velocities. The *angular offset* $\angle(x_i(t), v_i(t))$, denoted by $\omega_i(t)$, plays an important role in the analysis: we will show that it decreases roughly as $(\log t)/t$.

Instead of following a given bird over time and investigating its trajectory locally, we track an imaginary bird that has the ability to switch identities with its neighbors: this *virtual bird* could be B_i for a while and then decide, at any time, to become any B_j adjacent to it in the flock. Think of a bird passing a *baton* to any of its neighbors: whoever holds the baton is the virtual bird. This gives us a language in which we can define a baton-passing protocol of our choice. Our goal will be to design a protocol that keeps the baton along a trajectory that is nearly straight. A key idea is to trace the flight path of virtual birds *backwards* in time. This is how we are able to translate stochasticity into convexity and express the problem in geometric terms. This is made possible by the velocity dynamics, $v(t) = P(t-1)v(t-1)$, which implies that

$$v_i(t) \in \text{Conv}\,\{\, v_j(t-1) \mid (i,j) \in G_{t-1} \,\}.$$

By iterating in this fashion, we create the *flight net* $N_i(t)$ of bird B_i at time $t > 0$. It is a connected collection of line segments (i.e., a 1-skeleton): $N_i(t) = N_i(t, K_t)$, where K_t is a large integer parameter. Specifically, we set

$$K_t = \lceil n^{b_0} \log t \rceil \tag{15}$$

for a big enough constant b_0. The power of the flight net comes from its ability to deliver both kinetic and positional information about the "genealogy" of a bird's current state. It is a geometric analog of the "path integrals" of quantum mechanics, which allows us to make statements about an exponential number of virtual paths. Let K be an arbitrary positive integer; we define $N_i(t, K)$ inductively as follows. The case $t = 1$ is straightforward: $N_i(t, K)$ consists of the single line segment $x_i(0)x_i(1)$ in $\mathbb{R}^4$. Suppose
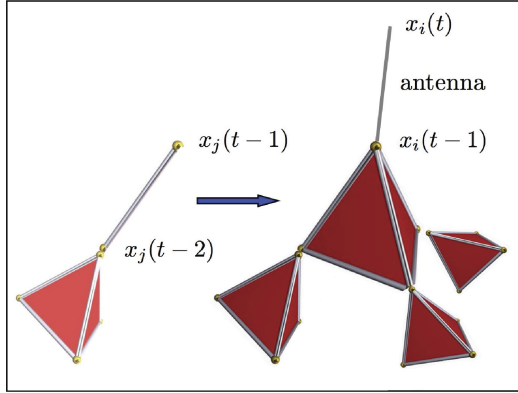
Fig. 4. In the critical case, the flight net for $t-1$ is shifted from bird B_j to bird B_i by the baton-passing drift.

that $t > 1$. We say that time s is *critical* if $s \leq K$ or if, during the time interval $[s-K, s]$, there is a network switch, i.e., $G_u \neq G_{u+1}$ for some u $(s - K \leq u \leq s)$.

— If t is critical, then $N_i(t, K)$ consists of the segment $x_i(t-1)x_i(t)$, together with the translates $N_j(t-1, K) + x_i(t-1) - x_j(t-1)$, for all $(i, j) \in G_{t-1}$ and $j = i$.
— If t is noncritical, then $N_i(t, K)$ consists of the segment $x_i(t-1)x_i(t)$, together with $N_i(t-1, K)$.

Every flight net has an *antenna* sitting on top, which is a line segment extending from $X_4 = t-1$ to $X_4 = t$ in the case of $N_i(t, K)$ (Figure 4). In the noncritical case, the antenna is connected on top of the previous one, ie, the one for $N_i(t-1, K)$. Otherwise, we slide the time-$(t-1)$ flight nets of the adjacent birds so that their antennas join with the bottom vertex of the new antenna: this shift is called the *baton-passing drift*.

Here is the intuition. Flying down the top antenna of the net, the virtual bird hits upon another antenna: either there is only one to choose from, in which case it is almost collinear (because of noncriticality, the corresponding random walk is thoroughly mixed) or else the virtual bird discovers a whole bouquet of antennas and picks one of them. Because the old antenna (marching backwards in time) is a convex combination of the new ones, the virtual bird can continue its backward flight by choosing from a convex cone of directions. This is when the baton is passed: the virtual bird changes its correspondence with an actual bird as it chooses one of these directions. Because of the translation by $x_i(t-1) - x_j(t-1)$, this change of correspondence is accompanied by a shift of length at most one, what we dub the baton-passing drift.

Viewed from a suitable perspective, the flight net provides a quasi-convex structure from which all sorts of metric information can be inferred. Most important, it yields the crucial *Escape Lemma*, which implies that, as time goes by, it becomes increasingly easy to predict the velocity of a bird from its location, and vice-versa. We begin with a simple observation. For any time $t > 0$, the 4-dimensional vector

$$w_i(t) \overset{\text{def}}{=} \tfrac{1}{t} x_i(t) \tag{16}$$

represents the constant velocity that bird B_i would need to have if it were to leave the origin at time 0 and be at position $x_i(t)$ at time t while flying in a fixed direction. We show that the angular offset $\omega_i(t) = \angle(x_i(t), v_i(t))$ cannot deviate too much from the *velocity offset* $\|v_i(t) - w_i(t)\|_2$.

LEMMA 4.1. *For any $t > 0$, $n^{-O(1)} \|v_i(t) - w_i(t)\|_2 \leq \omega_i(t) \leq O(\|v_i(t) - w_i(t)\|_2)$.*
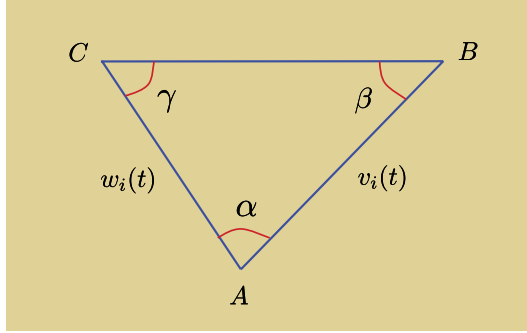
Fig. 5. Angular and velocity offsets are closely aligned.

PROOF. Consider the triangle ABC formed by identifying $\overrightarrow{AB}$ with $v_i(t)$ and $\overrightarrow{AC}$ with $w_i(t)$, and let α, β, γ be the angles opposite BC, CA, AB, respectively (Figure 5). Note that $\alpha = \omega_i(t)$ and $\|v_i(t) - w_i(t)\|_2 = |BC|$. Assume that $\beta \leq \gamma$; we omit the other case, which is virtually identical. All velocities maintain polynomially-bounded magnitudes throughout the flight,[6] so AB and AC have length between 1 and $n^{O(1)}$; therefore, if $\alpha \neq 0$ then $n^{-O(1)} \leq \beta < \pi/2$. The proof follows from the law of sines, $|BC|^{-1} \sin \alpha = |AC|^{-1} \sin \beta$.   □

LEMMA 4.2 (ESCAPE LEMMA). *For any bird B_i, at any time $t > 0$,[7]*

$$\omega_i(t) \leq \frac{\log t}{t} n^{O(n^2)}.$$

PROOF. From the initial conditions, we derive a trivial upper bound of $n^{O(1)}$ for constant t, so we may assume that t is large enough and $\omega_i(t) > 0$. The line passing through $x_i(t)$ in the direction of $v_i(t)$ intersects the hyperplane $X_4 = 0$ in a point p at distance from the origin, $\|p\|_2 = \Omega(t\omega_i(t))$. Recall that the bird B_i started its journey at distance $n^{O(1)}$ from the origin. If it had flown in a straight line, then we would have $p = x_i(0)$, hence $\omega_i(t) = \frac{1}{t} n^{O(1)}$, and we would be done. Chances are the bird did not fly straight, however, so to rescue the argument we exhibit a virtual bird that (almost) does. The idea is to use the flight net to follow the trajectory of a virtual bird that closely mimics a straight flight from p to $x_i(t)$.

Some words of intuition. If all times were critical, then it would be easy to prove by backward induction that, for all $0 \leq s < t$, the segment $px_i(t)$ intersects each hyperplane $X_4 = s$ in a point that lies within the convex hull of $N_i(t) \cap \{X_4 = s\}$. This would imply that p lies in the convex hull of the birds at time 0, which again would give us the same lower bound on $\omega_i(t)$ as previously mentioned (modulo the baton-passing drift). In fact, it would be possible to trace a *shadow path* from $x_i(t)$ down the flight net that leads to a virtual bird at time 0 that is even further away from the origin than p (Figure 6); we use the fact that no point can be further to a point in a convex polytope than to all of its vertices. Unfortunately, this convexity argument breaks down because of the net's jagged paths over noncritical time periods. The jaggedness is so small, however, that it provides us enough "quasi-convexity" to salvage the argument.

––––––––––
[6]Because they start out with that property and are modified only by taking averages among themselves.
[7]The exponent $O(n^2)$ hides a factor of $\log b$.
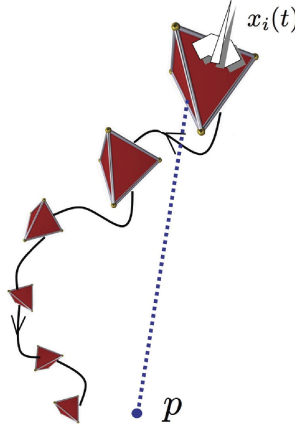
Fig. 6. The shadow path attempts to follow the segment $px_i(t)$ closely. Polytopes have been inverted to highlight the velocity formation.

First, we describe the shadow path; then we show why it works. Instead of handling convexity in $\mathbb{R}^4$, we will find it easier to do this in projection. By Lemma 4.1, there exists a coordinate axis, say X_1, such that

$$0 < \omega_i(t) = O(v_i(t)_1 - w_i(t)_1). \tag{17}$$

Note that we may have to reverse the sign of $v_i(t)_1 - w_i(t)_1$ of it is negative, but this is immaterial. The shadow path $x_i^{\mathrm{v}}(t), x_i^{\mathrm{v}}(t-1), \ldots, x_i^{\mathrm{v}}(0)$ describes the flight of the virtual bird B_i^{v} backwards in time. The first two vertices are $x_i^{\mathrm{v}}(t) = x_i(t)$ and $x_i^{\mathrm{v}}(t-1) = x_i(t-1)$. This means the virtual bird flies down the topmost edge of $N_i(t)$, ie, in the negative X_4 direction. Next, the following rule applies for $s = t, t-1, \ldots, 2$.

— If s is noncritical, $N_i(t)$ has a single edge $y_{s-2}y_{s-1}$, with $(y_{s-2})_4 = s-2$. The virtual bird flies down $y_{s-2}y_{s-1}$ and we set $x_i^{\mathrm{v}}(s-2) = y_{s-2}$ accordingly.
— If s is critical, $N_i(t)$ has one or several edges $y_{s-2}^k y_{s-1}$ ($k = 1, 2, \ldots$), with $(y_{s-2}^k)_4 = s-2$. The virtual bird follows the edge with maximum X_1-extant, that is, the one that maximizes $(y_{s-1})_1 - (y_{s-2}^k)_1$. (Recall that, although neither y_{s-1} nor y_{s-2}^k might be the position of any actual bird, their difference $y_{s-1} - y_{s-2}^k$ is the velocity vector $v_j(s-1)$ of some B_j.) We set $x_i^{\mathrm{v}}(s-2) = y_{s-2}^k$.

The virtual bird thus moves down the flight net back in time until it lands at $X_4 = 0$. We will show that it must then be close to both p and the origin, which will give us an upper bound on the length of p, hence on $\omega_i(t)$. The resulting collection of $t+1$ vertices forms the shadow path of the virtual bird B_i^{v} at time t. Naturally, we define the velocity of B_i^{v} at time $s > 0$ as $v_i^{\mathrm{v}}(s) = x_i^{\mathrm{v}}(s) - x_i^{\mathrm{v}}(s-1)$. Note that $v_i^{\mathrm{v}}(t) = v_i(t)$. To prove that the shadow path does not stray far from the straightline flight from $x_i(t)$ to p, we focus on the difference

$$V_s = v_i^{\mathrm{v}}(s)_1 - w_i(t)_1, \tag{18}$$

for $s \geq 1$ (Figure 7). Consider an interval $[r, s]$ consisting entirely of noncritical times (hence, $r > K_t$). The flock that contains the virtual bird B_i^{v} is invariant between times $r - K_t$ and s; furthermore, B_i^{v} has the same incarnation as some fixed B_j during the time

Fig. 7.   Following the shadow path.

period $[r-1, s]$. If $\chi(j)$ denotes the n-dimensional vector with all coordinates equal to 0, except for $\chi(j)_j = 1$, then, for $r-1 \leq u \leq s$,

$$v_j(u) = \chi(j)^T P^{u-r+K_t} v(r - K_t).$$

We abuse notation and restrict P and $v(r - K_t)$ to the flock of B_j and not to all of $G_{r-K_t} = \cdots = G_s$. By (5), we find that

$$
\begin{aligned}
|v_j(u)_1 - (\mathbf{m}_\pi[v(r-K_t)])_1| &\leq \|(\chi(j)^T P^{u-r+K_t} - \pi^T)v(r-K_t)\|_F \\
&\leq \|P^{u-r+K_t} - \mathbf{1}\pi^T\|_F \|v(r-K_t)\|_F \\
&\leq e^{-(u-r+K_t)n^{-O(1)} + O(\log n)}.
\end{aligned}
$$

We conclude that

$$
\begin{aligned}
|V_{r-1} - V_s| = |v_i^{\mathrm{v}}(r-1)_1 - v_i^{\mathrm{v}}(s)_1| = |v_j(r-1)_1 - v_j(s)_1| \\
\leq |v_j(r-1)_1 - (\mathbf{m}_\pi[v(r-K_t)])_1| + |v_j(s)_1 - (\mathbf{m}_\pi[v(r-K_t)])_1| \qquad (19) \\
\leq e^{-K_t n^{-O(1)} + O(\log n)}.
\end{aligned}
$$

Suppose now that $s > 1$ is critical. Then, $v_i^{\mathrm{v}}(s)$ is a convex combination of the vectors of the flight net joining $X_4 = s - 2$ to $X_4 = s - 1$. It follows that $v_i^{\mathrm{v}}(s-1)_1 \geq v_i^{\mathrm{v}}(s)_1$; hence $V_{s-1} \geq V_s$. Let C be the number of critical times. For all $1 \leq s \leq t$,

$$V_s \geq V_t - \mathrm{C}\, e^{-K_t n^{-O(1)} + O(\log n)}.$$

Summing over all s,

$$\sum_{s=1}^{t} V_s \geq t V_t - (t-1)\mathrm{C}\, e^{-K_t n^{-O(1)} + O(\log n)}.$$

By (17), $V_t = v_i^{\mathrm{v}}(t)_1 - w_i(t)_1 = \Omega(\omega_i(t))$; therefore,

$$\omega_i(t) = O(V_t) = \frac{O(1)}{t}\left|\sum_{s=1}^{t} V_s\right| + \mathrm{C}\, e^{-K_t n^{-O(1)} + O(\log n)}. \qquad (20)$$

By (16) and (18),

$$\sum_{s=1}^{t} V_s = \sum_{s=1}^{t} \left\{ x_i^{\mathrm{v}}(s)_1 - x_i^{\mathrm{v}}(s-1)_1 - w_i(t)_1 \right\} = x_i^{\mathrm{v}}(t)_1 - x_i^{\mathrm{v}}(0)_1 - tw_i(t)_1$$
$$= x_i^{\mathrm{v}}(t)_1 - x_i(t)_1 - x_i^{\mathrm{v}}(0)_1 = -x_i^{\mathrm{v}}(0)_1 .$$

Since $x_i^{\mathrm{v}}(0)_1$ is the position of a virtual bird at time 0, it is tempting to infer that it is also the position of some actual bird at that time; hence, $|x_i^{\mathrm{v}}(0)_1| = n^{O(1)}$. This is not quite true because adding together the velocity vectors ignores the baton-passing drift, ie, the displacements caused by switching birds. At critical times, the virtual bird gets assigned a new physical bird that is adjacent to its currently assigned feathered creature. Since $(i,j) \in G_{t-1}$, this causes a displacement of at most 1 and highlights the role of the flight net as both a kinetic and a positional object. Summing up, we find that $|x_i^{\mathrm{v}}(0)_1| \le C + n^{O(1)}$; hence

$$\left| \sum_{s=1}^{t} V_s \right| \le C + n^{O(1)}. \tag{21}$$

A time is critical if a network switch takes place within the past K_t steps; so, by (15) and the bound on the number N of switches (Lemma 3.1),

$$C \le K_t N = n^{O(n^2)} \log t$$

and the lemma follows from (20) and (21). □

COROLLARY 4.3. *At any time $t > 1$, a bird turns by an angle* $(v_i(t), v_i(t+1))$ *at most* $\frac{\log t}{t} n^{O(n^2)}$.

PROOF. No bird can take a step bigger than $n^{O(1)}$, so the angle between the vectors $x_i(t)$ and $x_i(t+1)$ is at most $\frac{1}{t} n^{O(1)}$. As a result,

$$(v_i(t), v_i(t+1)) \le (v_i(t), x_i(t)) + (x_i(t), x_i(t+1)) + (x_i(t+1), v_i(t+1))$$
$$\le \omega_i(t) + (x_i(t), x_i(t+1)) + \omega_i(t+1),$$

and the proof follows from the Escape Lemma. The property we are using here is the triangle inequality for angles: equivalently, the fact that, among the three angles around a vertex of a tetrahedron in $\mathbb{R}^3$, none can exceed the sum of the others. Even though the birds were lifted to four dimensions, the argument involves only three points at a time and so can be carried out in $\mathbb{R}^3$. □

We can show that two adjacent birds can never have strayed too far from each other in the past.

COROLLARY 4.4. *If two birds are adjacent in the flocking network at time $t > 1$, their distance prior to t always remains within $n^{O(n^2)} \log t$.*

PROOF. We may assume that t is large enough. Consider the time s that maximizes the distance R_s, for all $s \in [0, t-1]$, between the points $x_i(s)$ and $p = (s/t)x_i(t)$ in the

Fig. 8. Two birds cannot stray too far from each other if they are ever to meet again.

hyperplane $X_4 = s$. For the same reason, we may assume that $s > 1$. By the Escape Lemma and Corollary 4.3,

$$(x_i(s), v_i(s+1)) \leq \omega_i(s) + \quad (v_i(s), v_i(s+1)) \leq \frac{\log s}{s} n^{O(n^2)}. \tag{22}$$

Set up an orthogonal coordinate system in the plane spanned by $O, p, x_i(s)$ (Figure 8): O is the origin; the X-axis lies in the hyperplane $X_4 = 0$ and runs in the direction from p to $x_i(s)$; the Y-axis is normal to OX in the $O, p, x_i(s)$ plane. The Y-coordinate p_Y of p satisfies $s \leq p_Y \leq sn^{O(1)}$. Let $Y = X \tan \alpha$ and $Y = X \tan \beta$ be the two lines through the origin passing through $x_i(t)$ and $x_i(s)$, respectively. Setting $Y = p_Y$, we find that $p_X = p_Y / \tan \alpha$ and $x_i(s)_X = p_X / \tan \beta$; therefore

$$R_s \leq \left| \frac{1}{\tan \beta} - \frac{1}{\tan \alpha} \right| sn^{O(1)} \leq \frac{\sin(\alpha - \beta)}{(\sin \alpha)(\sin \beta)} sn^{O(1)}.$$

Be definition of s, the velocity $v_i(s+1)$ cannot take the bird B_i outside the elliptical cylinder that is centered at the line $(O, x_i(t))$ with the point $x_i(s)$ on its boundary and that intersects $X_4 = 0$ in a disk of radius $R_s = |px_i(s)|$. It follows that the normal projection w of $v_i(s+1)$ on the (X, Y)-plane forms an angle γ with $x_i(s)$ at least equal to the angle between the two lines $Y = X \tan \alpha$ and $Y = X \tan \beta$, which is $\alpha - \beta$. By (22), therefore,

$$\alpha - \beta \leq \gamma \leq \quad (x_i(s), v_i(s+1)) \leq \frac{\log s}{s} n^{O(n^2)}.$$

Birds are at most $n^{O(1)}$ away from the origin at time 0 and take no step larger than that bound. It follows that both α and β are at least $n^{-O(1)}$; hence

$$R_s \leq n^{O(n^2)} \log t.$$

If two birds B_i and B_j share an edge in a flock at time t, then $\|x_i(t) - x_j(t)\|_2 \leq 1$; so, by the triangle inequality, at any time $1 < t' \leq t$,

$$\|x_i(t') - x_j(t')\|_2 \leq 2n^{O(n^2)} \log t + \frac{t'}{t} \|x_i(t) - x_j(t)\|_2. \qquad \square$$

Suppose that birds B_i and B_j are distance at most D at time $t > 0$. (No assumption is made whether they belong to the same flock.) By (13) and (16) and Lemmas 4.1 and 4.2,

$$\begin{aligned}
\Delta_{ij}(t) &\le \|v_i(t) - v_j(t)\|_2 \\
&\le \|v_i(t) - \tfrac{1}{t}x_i(t)\|_2 + \|v_j(t) - \tfrac{1}{t}x_j(t)\|_2 + \tfrac{1}{t}\|x_i(t) - x_j(t)\|_2 \quad (23)\\
&\le (\omega_i(t) + \omega_j(t))n^{O(1)} + \tfrac{D}{t} \, .
\end{aligned}$$

COROLLARY 4.5. *At any time $t > 1$, the difference in stationary velocities between two distinct flocks joining into a common one at time $t + 1$ has Euclidean norm at most $\frac{\log t}{t}\, n^{O(n^2)}$.*

PROOF. The stationary velocity of a flock is a convex combination of its constituents' individual velocities, so the difference in stationary velocities cannot exceed, lengthwise, the maximum difference between individual ones. By the connectivity of flocks, the distance at time t between any two birds in the common flock at time $t + 1$ cannot exceed $D = n^{O(1)}$. The lemma follows from (23) and Lemma 4.2. □

We define the *fragmentation breakpoint t_f* as

$$t_f = n^{(bn)^2}, \quad (24)$$

where b is the constant used in the perturbation rule. Setting $D = 1$ in (23), we find that, by the perturbation rule and the Escape Lemma (including its footnote), the edges of G_t can break only if $t < t_f$. Past the fragmentation breakpoint, flocks can never fragment.

LEMMA 4.6. *At any time $t \ge t_f$, the flocking network G_t may gain new edges but never lose any.*

The Escape Lemma tells us that, after the fragmentation breakpoint, birds fly almost in a straight line and both their positions and velocities can be predicted with low relative error. From a physical standpoint, they have already converged. The flocking network may still change, however. It may keep doing so even after an astronomical amount of time. This is what we show in the next two sections.

5. FLOCKING TIME: COARSE ANALYSIS
We show that the flocking network becomes static after a number of steps equal to a tower-of-twos of linear height. This allows us to present some of the main ideas and prepare the grounds for the more difficult proof of the logarithmic height in Section 6. The main tools we use in this section are the rationality of limit configurations and root separation bounds from elimination theory. Our investigation focuses on the post-fragmentation phase, ie, with $t > t_f$.

LEMMA 5.1. *Consider two birds adjacent at time t but not $t - 1$. Assume that the flocks that contain them remain invariant during the period $[t_1, t - 1]$, where $t_f < t_1 < t - 1$. If, at time $t - 1$, the birds are in different flocks with distinct stationary velocities, then $t \le n^{O(t_1 n)}$; otherwise, $t \le t_1 2^{n^{O(1)}}$.*

PROOF. Assume that the flocking network G_t stays invariant during the period $[t_1, t - 1]$. Consider two birds B_i and B_j that are adjacent in G_t but not during $[t_1, t - 1]$. The two birds may or may not be in the same flock at time $t - 1$. Let the flock for B_i (respectively, B_j) consist of m (respectively, m') birds: $m = m'$ if the birds are in the same flock, else $m + m' \le n$. By abuse of notation, we use the terminology of (3), that is, $P, \pi, C, u_k, \lambda_k$, as well as $v(t)$, to refer to the flock of m birds, and we add primes to

distinguish it from the flock of B_j. We wish to place an upper bound on $t - t_1$. Let $\chi(i)$ denote the m-dimensional vector with all coordinates equal to 0, except for $\chi(i)_i = 1$. By (3) and (9), for $t > t_1$,

$$x_i(t) = x_i(t_1) + \sum_{s=0}^{t-t_1-1} \chi(i)^T P^s v(t_1 + 1)$$

$$= x_i(t_1) + (t - t_1)y + \sum_{k=2}^{m} \frac{1 - \lambda_k^{t-t_1}}{1 - \lambda_k} \Phi_k,$$

where the three-dimensional row vectors y and Φ_k are defined by

$$\begin{cases} y = \pi^T v(t_1 + 1) = \mathbf{m}_\pi [v(t_1 + 1)] ; \\ \Phi_k = \chi(i)^T C^{1/2} u_k u_k^T C^{-1/2} v(t_1 + 1). \end{cases} \tag{25}$$

Note that, by (3) and (6),

$$\sum_{k=2}^{m} \frac{1}{1 - \lambda_k} \Phi_k = \lim_{t \to \infty} \sum_{s=0}^{t-1} \sum_{k=2}^{m} \lambda_k^s \chi(i)^T C^{1/2} u_k u_k^T C^{-1/2} v(t_1 + 1)$$

$$= \sum_{s=0}^{\infty} \chi(i)^T (P^s - \mathbf{1}\pi^T) v(t_1 + 1) = \chi(i)^T \Gamma v(t_1 + 1) ;$$

therefore,

$$x_i(t) = x_i(t_1) + \chi(i)^T \Gamma v(t_1 + 1) + (t - t_1)y - \sum_{k=2}^{m} \lambda_k^{t-t_1} \frac{\Phi_k}{1 - \lambda_k} .$$

Adding primes to distinguish between the flocks of B_i and B_j (if need be), we find that

$$x_i(t) - x_j(t) = A + (t - t_1)B - \sum_{k=1}^{m_0} \Psi_k \mu_k^{t-t_1}, \tag{26}$$

where the following hold.

(i) $A = x_i(t_1) - x_j(t_1) + \chi(i)^T \Gamma v(t_1 + 1) - \chi'(j)^T \Gamma' v'(t_1 + 1)$. The vectors $v(t_1 + 1)$, $v'(t_1 + 1)$, $x_i(t_1)$, and $x_j(t_1)$ have rational coordinates over $O(t_1 n \log n)$ bits that can be assumed to share a common denominator. In view of Lemma 2.3, this implies that the same is true of the three-dimensional vector A.

(ii) $B = y - y'$. The stationary distribution $\pi = (\operatorname{tr} C^{-1})^{-1} C^{-1} \mathbf{1}$ is a rational vector over $O(n \log n)$ bits, which implies that B has rational coordinates over $O(t_1 n \log n)$ bits; hence either $B = (0, 0, 0)$ or $\|B\|_2 \geq n^{-O(t_1 n)}$.

(iii) $\mu_1 \geq \cdots \geq \mu_{m_0}$. Each μ_k is an eigenvalue λ_l or λ_l' ($l, l' > 1$) and $|\mu_k| < 1$. Their number m_0 is either $m - 1$ (if the two birds B_i and B_j belong to the same flock) or $m + m' - 2$, otherwise.

(iv) Each Ψ_k is a three-dimensional row vector of the form $\Phi_l/(1 - \lambda_l)$ or $-\Phi_l'/(1 - \lambda_l')$. By (4), the eigenvalues are bounded away from 1 by $n^{-O(1)}$, so it follows from (25) that $\|\Psi_k\|_2 = n^{O(1)}$. Likewise, we note for future reference that

$$\| \sum_{k=1}^{m_0} \Psi_k \mu_k^{t-t_1} \|_2 \leq e^{-(t-t_1)n^{-O(1)} + O(\log n)} = n^{O(1)}. \tag{27}$$

We distinguish among three cases:

Case I. $B \neq \mathbf{0}$. The two stationary velocities are distinct; hence, so are the corresponding flocks. By (i) and (ii), $\|A\|_2 \leq n^{O(t_1 n)}$ and $\|B\|_2 \geq n^{-O(t_1 n)}$. If the two birds are to be joined in G_t, then $\text{DIST}_t(\mathcal{B}_i, \mathcal{B}_j) = \|x_i(t) - x_j(t)\|_2 \leq 1$. It follows from (26) that $t - t_1 \leq n^{O(t_1 n)}$. Note that, for the lower bound of $n^{-O(t_1 n)}$ on $\|B\|_2$ to be tight, the flock would have to be able to generate numbers almost as small as allowed by the bit-length growth. For this to happen, energy must shift toward the dominant eigenvalue. This spectral shift occurs only in a specific context, which we examine in detail in the next section.

Case II. $B = \mathbf{0}$ and $\|A\|_2 \neq 1$. The two stationary velocities are identical but the distance between the two birds does not tend to 1. By (i), $\|A\|_2$ is bounded away from 1 by $n^{-O(t_1 n)}$. It follows from (26) and (27) and the triangle inequality that

$$
\begin{aligned}
| \, \|x_i(t) - x_j(t)\|_2 - 1 \, | &\geq | \, \|A\|_2 - 1 \, | - | \, \|x_i(t) - x_j(t)\|_2 - \|A\|_2 \, | \\
&\geq n^{-O(t_1 n)} - \| \textstyle\sum_k \Psi_k \, \mu_k^{t - t_1} \|_2 \\
&\geq n^{-O(t_1 n)} - e^{-(t - t_1) n^{-O(1)} + O(\log n)}.
\end{aligned}
$$

For a large enough constant b_0, the distance between the two birds remains bounded away from 1 by $n^{-O(t_1 n)}$ at any time $s \geq t_1 n^{b_0}$. Not only that, but the sign of $\text{DIST}_s(\mathcal{B}_i, \mathcal{B}_j) - 1$ can no longer change after time $t_1 n^{b_0}$. Indeed, for any $s \geq t_1 n^{b_0}$, the distance between times $s - 1$ and s varies by an increment of $\Delta_{ij}(s)$, where, by (27),

$$
\begin{aligned}
\Delta_{ij}(s) &= | \, \|x_i(s) - x_j(s)\|_2 - \|x_i(s-1) - x_j(s-1)\|_2 | \\
&\leq \| \textstyle\sum_k \Psi_k \, \mu_k^{s-1-t_1} \|_2 + \| \textstyle\sum_k \Psi_k \, \mu_k^{s-t_1} \|_2 \\
&\leq e^{-(s - t_1) n^{-O(1)} + O(\log n)} \leq e^{-t_1 n^{b_0/2}}.
\end{aligned}
$$

With b_0 assumed large enough, this ensures that, past time $t_1 n^{b_0}$, the distance can never cross the value 1. Thus, if the two birds have not gotten within distance 1 of each other by time $t_1 n^{b_0}$, they never will—at least while their respective flocks remain invariant. We conclude that $t \leq t_1 n^{O(1)}$.

Case III. $B = \mathbf{0}$ and $\|A\|_2 = 1$. The distance between the two birds tends to 1. The concern is that they might stay safely away from each other for a long period of time and then suddenly get close enough to share an edge. The rationality of the limit configuration alone is insufficient to prevent this. Only a local analysis of the convergence can show that a long-delayed pairing is impossible. We wish to prove that, if $\text{DIST}_s(\mathcal{B}_i, \mathcal{B}_j)$ is to fall below 1 for $s > t_1$, this must happen relatively soon. Recall that, by (26),

$$
x_i(s) - x_j(s) = A - \sum_{k=1}^{m_0} \Psi_k \, \mu_k^{s - t_1},
$$

where A is a unit vector. We investigate the behavior of the birds' distance locally around 1. Keeping in mind that we are working with row vectors,

$$
\|x_i(s) - x_j(s)\|_2^2 = 1 - 2 \sum_k \Psi_k A^T \mu_k^{s - t_1} + \sum_{k,k'} \Psi_k \Psi_{k'}^T (\mu_k \mu_{k'})^{s - t_1}.
$$

Let $1 > \rho_1 > \cdots > \rho_N > 0$ be the distinct nonzero values among $\{|\mu_k|, |\mu_k \mu_{k'}|\}$ ($N < n^2$). These absolute values may appear with a plus or minus sign in the previous expression, so we rewrite it as

$$\|x_i(s) - x_j(s)\|_2^2 - 1 = \sum_{k=1}^{N} \Upsilon_k \, \rho_k^{s-t_1}, \tag{28}$$

where each

$$\Upsilon_k = \Upsilon_k^+ + (-1)^s \, \Upsilon_k^-$$

corresponds to a distinct ρ_k. We distinguish between odd and even values of s so as to keep each Υ_k time-invariant. We assume that s is even and skip the odd case because it is similar. Of course, we may also assume that each $\Upsilon_k = \Upsilon_k^+ + \Upsilon_k^-$ is nonzero. We know that $\sum_k \Upsilon_k \rho_k^{s-t_1}$ tends to 0 as s goes to infinity, but the issue is how so. To answer this question, we need bounds on eigenvalue gaps and on $|\Upsilon_k|$. Tighter results could be obtained from current spectral technology, but they would not make any difference for our purposes, so we settle for simple, conservative estimates. $\square$

LEMMA 5.2. *For all $k > 1$ and $k \geq 1$, respectively,*

$$\rho_k \leq (1 - 2^{-n^{O(1)}})\rho_1 \qquad and \qquad 2^{-t_1 2^{n^{O(1)}}} \leq |\Upsilon_k| = n^{O(1)}.$$

PROOF. We begin with the eigenvalue gap.[8] For this, we use a conservative version of Canny's root separation bound [Canny 1988; Yap 2000]: Given a system of m integer-coefficient polynomials in m variables with a finite set of complex solution points, any nonzero coordinate has modulus at least

$$2^{-\ell D^{O(m)}}, \tag{29}$$

where $D - 1$ is the maximum degree of any polynomial and ℓ is the number of bits needed to represent any coefficient. Any difference $\rho_k - \rho_l$ can be expressed by a quadratic polynomial, $z = z_1 z_2 - z_3 z_4$, where each z_i is either 1 or the root of the characteristic polynomial $\det(P - \lambda \mathbf{I}_n)$. The elements of P are rationals over $O(n \log n)$ bits, so by the Hadamard bound [Yap 2000] the roots of $\det(P - \lambda \mathbf{I}_n)$ are also those of a polynomial of degree n with integer coefficients over $O(n^2 \log n)$ bits; therefore, $m \leq 5$; $D = n + 1$; and $\ell = n^{O(1)}$. This proves that the minimum gap between two ρ_k's is $2^{-n^{O(1)}}$. Since $\rho_1 < 1$, we find that, for $k > 1$,

$$\rho_k \leq (1 - 2^{-n^{O(1)}})\rho_1,$$

which proves the first part of the lemma.

By (iv), $\|\Psi_l\|_2 = n^{O(1)}$; therefore, by Cauchy-Schwarz and $\rho_k < 1$, $|\Upsilon_k| = n^{O(1)}$, which proves the second upper bound of the lemma. We now prove that $|\Upsilon_k|$ cannot be too small. Recall that it is the sum/difference of inner products between vectors in $\{A, \Psi_h\}$. We know from (iv) that Ψ_h is of the form $\Phi_l/(1 - \lambda_l)$ or $-\Phi'_l/(1 - \lambda'_l)$. We assume the former without loss of generality. By (3) and (6),

$$\Gamma = \sum_{r=2}^{n} \sum_{s \geq 0} \lambda_r^s C^{1/2} u_r u_r^T C^{-1/2}.$$

---

[8]For the purpose of this lemma, we again abuse notation by letting P and n pertain to the flock of either one of the two birds.

In view of (iv) and (25), it then follows that

$$
\begin{aligned}
\Psi_h &= \frac{\Phi_l}{1-\lambda_l} = \frac{1}{1-\lambda_l}\chi(i)^T C^{1/2} u_l u_l^T C^{-1/2} v(t_1+1) \\
&= \frac{1}{1-\lambda_l}\chi(i)^T C^{1/2} u_l u_l^T C^{-1/2} C^{1/2} u_l u_l^T C^{-1/2} v(t_1+1) \\
&= \sum_{r=2}^n \frac{1}{1-\lambda_r}\chi(i)^T C^{1/2} u_r u_r^T C^{-1/2} C^{1/2} u_l u_l^T C^{-1/2} v(t_1+1) \\
&= \sum_{r=2}^n \sum_{s\geq 0} \chi(i)^T \lambda_r^s C^{1/2} u_r u_r^T C^{-1/2} C^{1/2} u_l u_l^T C^{-1/2} v(t_1+1) \\
&= \chi(i)^T \Gamma C^{1/2} u_l u_l^T C^{-1/2} v(t_1+1) = \chi(i)^T \Gamma W,
\end{aligned}
$$

where $W = C^{1/2} u_l u_l^T C^{-1/2} v(t_1+1)$ and $v(t_1+1)$ is a vector with (same-denominator) rational coordinates over $O(t_1 n \log n)$ bits. By Lemma 2.3, the elements of Γ are rationals encoded over $O(n \log n)$ bits. Any coordinate of Ψ_h can thus be written as a sum $\sum_i$ of at most $n^{O(1)}$ terms of the form $R_i \alpha_i y_i z_i$, where:

— all the R_i's are products of the form $\Gamma_{\star\star} v_\star(t_1+1)$, hence rationals over $O(t_1 n \log n)$ bits;
— α_i is the square root of a rational $c_\star/c_\star$ over $O(\log n)$ bits;
— y_i, z_i are two coordinates of u_l. Recall that, by (2), u_l is a unit eigenvector of $C^{-1/2} P C^{1/2}$.

By (i), A is a vector with rational coordinates over $O(t_1 n \log n)$ bits. It follows that Υ_k is a sum $\sum_i$ of $n^{O(1)}$ terms of the form $S_i \gamma_i y_i z_i y_i' z_i'$:

— all the S_i's are rationals over $O(t_1 n \log n)$ bits;
— γ_i is the square root of an $O(\log n)$-bit rational, that is, of the form $\sqrt{(c_\star/c_\star)(c_\star/c_\star)}$;
— y_i, z_i, y_i', z_i' are coordinates of the eigenvectors of $C^{-1/2} P C^{1/2}$ (or 1, to account for $A^T \Psi_h$).

It is straightforward (but tedious) to set up an integer-coefficient algebraic system over $m = n^{O(1)}$ variables that includes Υ_k as one of the variables. The number of equations is also m and the maximum degree is n. All the coefficients are integers over $O(t_1 n \log n)$ bits. Rather than setting up the system in full, let us briefly review what it needs to contain.

(1) Υ_k is a sum of $n^{O(1)}$ quintic monomials $S_i \gamma_i y_i z_i y_i' z_i'$; where the S_i's are rationals over $O(t_1 n \log n)$ bits.
(2) Each γ_i is of the form $\sqrt{a/b}$, where a, b are $O(\log n)$-bit integers. We express it by the equation $b\gamma_i^2 = a$. (This yields two roots, but any solution set is fine as long as it is finite and contains those we want.)
(3) The y_i, z_i, y_i', z_i' are coordinates of the eigenvectors u_l of $C^{-1/2} P C^{1/2}$. We specify them by first defining the eigenvalues $\lambda_1, \ldots, \lambda_n$ and

$$
\begin{cases}
\det(P - \lambda_i \mathbf{I}_n) = 0; \\
C^{-1/2} P C^{1/2} u_i = \lambda_i u_i; \qquad (1 \leq i < j \leq n) \\
\|u_i\|_2^2 = 1, \text{ and } u_i^T u_j = 0.
\end{cases}
$$

The issue of eigenvalue multiplicity arises. If the kernels of the various $P - \lambda_i \mathbf{I}_n$ are not of dimension 1, we throw in cutting planes to bring it down (details omitted). We rewrite each eigensystem as $P v_i = \lambda_i v_i$, where $v_i = C^{1/2} u_i$, and again we square the latter set of equations to bring them in polynomial form.

Once we reduce all the rational coefficients to integers, we can use the separation bound (29), for $m = n^{O(1)}$, $D = n + 1$, and $\ell = O(t_1 n \log n)$. This gives us a bound on the modulus of any nonzero coordinate of the solution set; hence on $|\Upsilon_k|$.

By (28), it follows from the lemma that $\|x_i(s) - x_j(s)\|_2^2 - 1 = \Upsilon_1 \rho_1^{s-t_1}(1 + \zeta)$, where

$$|\zeta| \leq e^{-(s-t_1)2^{-n^{O(1)}} + t_1 2^{n^{O(1)}}} = o(1),$$

for $s \geq t_1 2^{n^{b_1}}$, with b_1 being a large enough constant. The same argument for odd values of s shows that, after $t_1 2^{n^{b_1}}$, either $\|x_i(s) - x_j(s)\|_2^2$ stays on one side of 1 forever or it constantly alternates (at odd and even times). Since the birds are joined in G_t but not in G_s ($t_1 \leq s < t$), it must be the case that $t \leq t_1 2^{n^{O(1)}}$. This concludes Case III.

Putting all three results together, we find that the bound from Case I is the most severe, $t \leq n^{O(t_1 n)}$: the collision is driven only by the lowest mode and not by the dissipative part of the spectrum (the damped oscillator). □

Lemmas 4.6 and 5.1 show that all network switches take place within the first $t_\infty = 2 \uparrow\uparrow O(n)$ steps. After time t_∞, the flocking network remains invariant. By virtue of (12), the limit trajectory of the birds within a given flock is expressed as

$$x(t) = x^r + \mathbf{1}\pi^T x(t_\infty) + (t - t_\infty)\mathbf{1}\pi^T v(t_\infty + 1),$$

where the stationary distribution π refers to the bird's flock.

6. FLOCKING TIME: FINE ANALYSIS

We prove that the tower-of-twos has height not just $O(n)$ but $O(\log n)$, which is the optimal bound. The improvement rests on a more careful analysis of the flock merges and edge additions subsequent to the fragmentation breakpoint t_f. Note that in the proof of Lemma 5.1 the bottleneck lies in Case I, specifically in the bounds on $\|A\|_2$ and $\|B\|_2$. The former can be improved easily by invoking the Escape Lemma. To get around $\|B\|_2$ requires more work. Recall from (11) that the position vector of one flock is given by

$$x(t) = a + bt + \Gamma_t v,$$

where the matrix Γ_t specifies a damped oscillator. The stationary velocity b is formed by the first spectral coordinates, one for each dimension, associated with the eigenvalue 1.

The Combinatorics of the Spectral Shift. The reason flocks take longer to merge into larger flocks is that they fly in formations increasingly parallel to one another. The term bt grows linearly in t, so an iterated exponential growth can only come from the oscillator. The angle between the flight directions of two flocks is given by the stationary velocities; therefore, for the angles to inherit an exponentially decaying growth, it is necessary to *transfer* the exponentially decaying energy of the oscillators to the stationary velocities themselves. In other words, the collision between two flocks must witness a *spectral shift* from the "subdominant" eigenspace to the stationary velocities. Small angles are achieved by getting two stationary velocities to be very close to each other. The spectral shift does not cause a decay of the velocities themselves but of pairwise differences. Let b be the stationary velocity of the new flock formed by two flocks
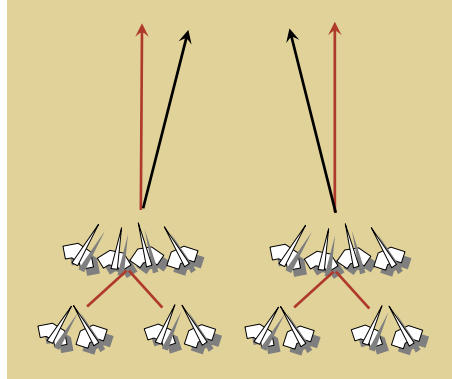
Fig. 9. Without spectral shift, the difference between stationary velocities becomes null and the two flocks never meet. The spectral shift resupplies the stationary velocities with the fast-decaying energy located in the subdominant part of the spectrum. This causes a slight inflection of the trajectory (black lines).

joining together after flying on their own during t steps. Let b' be the stationary velocity resulting from two other flocks flying in similar conditions. The spectral shift will ensure that the difference $b - b'$ has Euclidean norm $e^{-tn^{-O(1)}}$, that is, exponentially small in the flight time (Figure 9).

We model the sequence of post-fragmentation breakpoint edge inserts/deletes by a forest F: each internal node a corresponds to a flock F_a of n_a birds formed at time $t_a > t_f$. If a is a leaf of F, then by definition its formation time t_a is at most t_f. A node with at least two children is called *branching*. A nonbranching node represents a network switch formed by the addition of edges within the same flock. Our analysis will focus on branching nodes with no more than two children. In general, of course, this number can be arbitrarily high, as several flocks may come together to merge simultaneously. We will see later how to break down multiple aggregation of this form into pairwise merges. Late merging requires many birds, so let $L(t)$ denote the minimum value of n_a, the number of birds in F_a, over all branching nodes a such that $t_a \geq t$. Our previous upper bound shows that $L(t) = \Omega(\log^* t)$. We now strengthen this.

LEMMA 6.1. $L(t) \geq x_0^{\log^* t - O(\log \log n)}$, *where* $x_0 \approx 1.1938$ *is the unique real root of* $x^5 - x^2 - 1$.

This implies that the last merge must take place before time t such that $L(t) \leq n$; hence $t \leq 2 \uparrow\uparrow (3.912 \log n)$. By Lemma 5.1, multiplying this quantity by $2^{n^{O(1)}}$ suffices to account for the network switches following the last merge. This proves Theorem 1.1.

The merges occur in tree-like fashion. We show that the worst case is given by a balanced tree. But if the merging time grows exponentially at each step up the tree, wouldn't a fishbone-like tree be preferable by providing a long path? The answer is no: it will barely get us past exponential (depending on how we define the fishbone). Only two flocks with roughly the same characteristics can produce, by merging together, the sort of spectral shift causing an exponential growth in the next merge. Informally, the reason is that the spectral shift requires "clean" collisions if we want to iterate it. Indeed, poorly matched merges might momentarily cause an exponential burst, but they will accumulate *residues* that must be eliminated before the spectral shift can be iterated. Think of it as a sort of *lactic acid* for flocks. What is a residue? While a formal definition must await the introduction of the proper notation, it suffices to say that it

is the presence of undesired lower-order terms. Here is an example. Subtracting $A = 0.10110100000$ from $B = 0.10110100001$ yields the very small number 0.00000000001. Now turn the third bit of B to 1. The difference becomes 0.01000000001, which is not so small. This new third bit of B is a residue. If we want to produce small numbers again by subtraction, we will need somehow to turn that bit back to 0: this is what "clearing residues" is about.

Clearing Residues. Recall that $t_a > t_f$ is the time at which the flock F_a is formed at node a of F after the fragmentation breakpoint t_f given in (24). With the usual notational convention, it follows from (1) and (3) that, for $t \geq t_a$,

$$v_a(t) = P^{t-t_a} v_a(t_a) = \mathbf{1}_{n_a} \mathbf{m}_a + \sum_{k>1} \lambda_k^{t-t_a} C^{1/2} u_k u_k^T C^{-1/2} v_a(t_a),$$

where $\mathbf{m}_a = \pi_a^T v_a(t_a)$ is the stationary velocity of the flock F_a. As usual, it is understood that P, C, λ_k, u_k, etc., are all defined with respect to the specific flock F_a and not the whole group of n birds. We subscript $\mathbf{1}$ with the flock size for clarity. By (5),

$$\|v_a(t) - \mathbf{1}_{n_a} \mathbf{m}_a\|_2 \leq e^{-(t-t_a)n^{-O(1)} + O(\log n)}. \tag{30}$$

By the general form of the stationary distribution π_a as $(\operatorname{tr} C^{-1})^{-1} C^{-1} \mathbf{1}_{n_a}$, its coordinates are rationals over $O(n \log n)$ bits. Each coordinate of $\mathbf{m}_a$ is an irreducible rational p_a/q_a, where the number of bits needed for p_a and q_a is $O(t_a n \log n)$. We denote the maximum bit-length over all three coordinates by $\ell(\mathbf{m}_a)$:

$$\ell(\mathbf{m}_a) = O(t_a n \log n). \tag{31}$$

Consider a flock F_c associated with a branching node c of F: let a and b be the two children of c in F (hence $n_c = n_a + n_b$) and assume that $t_a \geq t_b$ and that no node of the forest F has more than two children, that is, flocks merge only two at a time.[9] By Corollary 4.5, the difference in stationary velocities between F_a and F_b satisfies

$$\|\mathbf{m}_a - \mathbf{m}_b\|_2 \leq \frac{\log t_c}{t_c} n^{O(n^2)}. \tag{32}$$

If the difference is null, then by Cases II, III of the previous analysis ($B = 0$), $t_c = t_a 2^{n^{O(1)}}$. Otherwise, by (31), and the equivalent bound for $\|\mathbf{m}_b\|_2$,

$$\|\mathbf{m}_a - \mathbf{m}_b\|_2 \geq n^{-O(t_a n)}. \tag{33}$$

The two inequalities (32) and (33) yield an upper bound on t_c. By our treatment of Cases II and III in the proof of Lemma 5.1, we conclude that, whether $\mathbf{m}_a = \mathbf{m}_b$ or not,

$$t_c \leq n^{O(t_a n)}. \tag{34}$$

This leads again to our earlier $\Omega(\log^* t)$ bound on $L(t)$. We show how to improve it. Let F_o be the forest derived from F by removing all nonbranching internal nodes and merging the adjacent edges in the obvious way. Our earlier assumption implies that each internal node of F_o has exactly two children. Let $a_0, \ldots, a_k$ ($k > 1$) be an ascending path in F_o and let b_i denote the unique sibling of a_i. We postpone the proof of the next result.

---

[9]The simultaneous merging of more than two flocks can be dealt with by breaking ties arbitrarily. Since there are fewer than n merges, this means that in our calculations time might be off by at most an additive term less than n. One can verify that this discrepancy has no real effect on any of the derivations and conclusions presented here.

LEMMA 6.2. *Assume that $2^{2^{t_f}} < \log\log\log t_{a_k} < t_{a_0}^4 < t_{a_1} < \log t_{a_k}$. Then, $t_{b_{i_0}} \geq \sqrt{\log\log t_{a_0}}$, for some $0 \leq i_0 < k$.*

The Recurrence. We set up a recurrence relation on $L(t)$ to prove the lower bound of Lemma 6.1, that is, $L(t) \geq (1.1938)^{\log^* t - O(\log\log n)}$. Write $t_0 = 2 \uparrow\uparrow \lfloor \log\log n \rfloor$. We may assume that n is large enough. For $t \leq t_0$, we have the trivial lower bound $L(t) \geq 1$ (choose the constant in the big-oh to be larger than 1), so we may assume that $t > t_0$. The child b of a node c (both defined with respect to $\mathcal{F}_o$) is called *near* if $t_b > (\log t_c)^{2/3}$.

LEMMA 6.3. *Any internal node c of $\mathcal{F}_o$ such that $t_c \geq 2^{2^{t_f}}$ has at least one near child.*

PROOF. By (34), we know that c has a child b_0 in the original forest $\mathcal{F}$ such that $t_c = n^{O(t_{b_0} n)}$. We exhibit a near child b for c. If b_0 is branching, set $b = b_0$; otherwise, set b to the nearest branching descendant of b_0. By Lemma 5.1, the formation times of any node in $\mathcal{F}$ and its nonbranching parent differ by at most a factor of $2^{n^{O(1)}}$. Since $\mathcal{F}$ has fewer than n^2 nodes and $t_c \geq 2^{2^{t_f}}$, with $t_f = n^{\Theta(n^2)}$,

$$t_b \geq 2^{-n^{O(1)}} t_{b_0} \geq 2^{-n^{O(1)}} \log t_c > (\log t_c)^{2/3}. \qquad \square$$

Let c_0 be an arbitrary node of $\mathcal{F}_o$ such that

$$t_{c_0} \geq t > t_0 = 2 \uparrow\uparrow \lfloor \log\log n \rfloor. \tag{35}$$

By the previous lemma, we can follow a descending path in $\mathcal{F}_o$ of near children $c_0, c_1, \ldots, c_l$, where $t_{c_l} < 2^{2^{t_f}} \leq t_{c_{l-1}}$. Because t_0 is so much greater than t_{c_l}, the path has more than a constant number of nodes—in fact, at least on the order of $\log\log n$. For future use, we note that

$$2^{2^{t_f}} < \log\log\log t_{c_0}. \tag{36}$$

LEMMA 6.4. *There exists $k > 1$ such that*

$$\log\log\log t_{c_0} < t_{c_k}^4 < t_{c_{k-1}} < \log t_{c_0}.$$

PROOF. By (36) and Lemma 6.3, there exists some c_j in $\mathcal{F}_o$ such that

$$(\log\log t_{c_0})^{2/3} < t_{c_j} < \log t_{c_0}.$$

Suppose now that all the nodes c_i, for $i = j+1, j+2, \ldots, l$, satisfy $t_{c_i}^4 \geq t_{c_{i-1}}$. Since there are most n nodes along the path from c_0 to c_l in $\mathcal{F}_o$, then, by (36) again,

$$2^{2^{t_f}} > t_{c_l} \geq t_{c_j}^{4^{-n}} > (\log\log t_{c_0})^{4^{-n-1}} > 2^{2^{2^{t_f/2}}}. \tag{37}$$

This contradiction proves the existence of some node c_k $(j < k \leq l)$ such that

$$t_{c_k}^4 < t_{c_{k-1}} < \log t_{c_0}.$$

The argument used in (37) shows that the smallest such k satisfies, via (36),

$$t_{c_{k-1}} \geq t_{c_j}^{4^{-n}} > (\log\log t_{c_0})^{4^{-n-1}} > 2^{2^{t_f}}.$$

Another application of the inequality above, $t_{c_{k-1}} > 2^{2^{t_f}}$, allows us to invoke
Lemma 6.3. By virtue of t_{c_0} being so big (36) and c_k being a near child of c_{k-1} (by
construction),

$$t_{c_k}^4 > (\log t_{c_{k-1}})^{8/3} > 4^{-8n}(\log\log\log t_{c_0})^{8/3} > \log\log\log t_{c_0}. \qquad \square$$

We now prove Lemma 6.1. Setting $a_i = c_{k-i}$ for $i = 0,\ldots,k$, together with (36),
the previous lemma sets the conditions of Lemma 6.2. This shows that $t_{a_0} >
(\log\log\log t_{a_k})^{1/4}$ and, conservatively,

$$t_{b_{i_0}} > (\log\log\log\log\log t_{a_k})^{1/3}.$$

Nodes a_0 and b_{i_0} are roots of disjoint subtrees, so the number of leaves below a_k is at
least the number of them below a_0 added to those below b_{i_0}. Since L is a monotone
function and, by (35), a_k is an arbitrary node such that $t_{a_k} \geq t$,

$$L(t) \geq L((\log\log\log t)^{1/4}) + L((\log\log\log\log t)^{1/3}), \qquad (38)$$

for $t > t_0 = 2 \uparrow\uparrow \lfloor \log\log n \rfloor$, and $L(t) \geq 1$ for $t \leq t_0$. We solve the recurrence without the
exponents, and then show that ignoring them makes no asymptotic difference. Define
$L^*(t) = 1$ for $t \leq t_0$ and, for any $t > t_0$,

$$L^*(t) = L^*(\log\log\log t) + L^*(\log\log\log\log\log t).$$

Given the bound we are aiming for, we can round off t down to the next tower-of-twos.
If $L^*(t) = M(\sigma)$, where $\sigma = \log^* t$, we can rewrite the recurrence relation as

$$M(\sigma) = M(\sigma - 3) + M(\sigma - 5),$$

where $M(\sigma) = 1$ for $\sigma \leq \log^* t_0$. Quite clearly, $M(\sigma)$ bounds the maximum number n_s
of leaves in a binary tree T^* where: (i) each left edge is labeled 3 and each right edge
5; and (ii) the sum of the labels along any path is at most $s = \log^* t - \log^* t_0$. We seek
a lower bound of the form cx^s. This means that $x^s \geq x^{s-3} + x^{s-5}$, for $s \geq 5$ and $cx^s \leq 1$
else. The characteristic equation is

$$x^5 - x^2 - 1 = 0.$$

We choose the unique real root $x_0 \approx 1.1938$; this leads to $c = x_0^{-5}$. This shows that
$n_s \geq x_0^{s-5}$; hence,

$$L^*(t) \geq x_0^{\log^* t - \log\log n - 5}.$$

It is obvious that the binary tree T associated with the recurrence for $L(t)$ embeds in
T^* with the same root. We show here that ignoring the exponents in (38) is harmless;
specifically, we prove that no leaf in T has more than a constant number of descendants
in T^*. This implies immediately that

$$L(t) \geq x_0^{\log^* t - O(\log\log n)},$$

which proves Lemma 6.1.

To prove our claim, we show that no path in T^* extends past its counterpart in T by
more than a constant number of nodes. We model simultaneous, parallel walks down
the trees as a collaborative game between two players, Bob and Alice, who take turns.
Initially, both of them share the same value

$$t_A = t_B = t > t_0.$$

In one round, Bob modifies his current value by taking iterated logs. He is entitled to up to five logarithm iterations; in other words, he can set

$$t_B \leftarrow \log t_B \quad \text{or} \quad t_B \leftarrow \log \log \log \log \log t_B,$$

or anything in-between. Alice mimics Bob's move but then completes it by taking a fractional power; for example, if Bob opts for, say, $\log \log t_B$, then Alice resets her value to $(\log \log t_A)^\alpha$, where α is a number between $\frac{1}{4}$ and 1. To summarize, Bob chooses the number of log iterations and Alice chooses α: they can change these parameters at each round. A player's *score* is the number of rounds before his or her value falls below (or at) t_0. Alice's score cannot be higher than Bob's, so the latter is expected to play the last rounds on his own. We show that his score never exceeds Alice's by more than a constant. This follows directly from the next two lemmas, whose proofs we postpone.

LEMMA 6.5. *The score differential is maximized when Bob always selects the single-iterated log rule and Alice follows suit with $\alpha = \frac{1}{4}$; in other words, $t_B \leftarrow \log t_B$ and $t_A \leftarrow (\log t_A)^{1/4}$.*

With the strategy of the lemma, Bob's score is $\log^* t - \log^* t_0$. Within an additive constant, Alice's score is at least the minimum h such that $c_h \geq t$, where c_i is defined by $c_0 = t_0^4$ and, for $i > 0$, $c_i = 2^{4c_{i-1}}$. To see why, note that the inverse of the function $z \mapsto (\log z)^{1/4}$ is $z \mapsto 2^{z^4}$; taking logarithms on both sides gives the recurrence on c_i.

LEMMA 6.6. *For $t > t_0$, $\min\{h \,|\, c_h \geq t\} \geq \log^* t - \log^* t_0 - O(1)$.*

This validates our claim that no path in T^* extends past its counterpart in T by more than a constant number of nodes. This fills in the missing part in the proof of Lemma 6.1.

PROOF OF LEMMA 6.2. We begin with a few technical facts. Recall from the "Clearing Residues" section that the flock F_c is associated with a branching node c of F and that a and b are its two children in F; furthermore, $t_a \geq t_b$ and $t_a > t_f$, where $t_f = n^{\Theta(n^2)}$. Assume that the velocity vector of F_a at time t_a can be written as

$$v_a(t_a) = \mathbf{1}_{n_a} \widetilde{\mathbf{m}}_a + u_a \mu_a + \zeta_a, \tag{39}$$

where $u_a \in \mathbb{R}^{n_a}$, $\widetilde{\mathbf{m}}_a$ and μ_a are three-dimensional row vectors, ζ_a is an n_a-by-3 matrix, and, for some real τ,

$$\begin{cases} 2^{t_f} \leq \tau \leq t_a^{1/3}; \\ \ell(\widetilde{\mathbf{m}}_a) = O(\log \log \tau); \\ \|u_a\|_\infty = 1 \ \& \ u_a \geq \mathbf{0}; \\ e^{-\tau n^{O(1)}} \leq \|\mu_a\|_2 \leq \frac{1}{\tau}; \\ \|\zeta_a\|_F \leq e^{-\tau^2 n^{-O(1)} + n^{O(1)}}. \end{cases} \tag{40}$$

Before justifying this assumption, we explain why it is useful. Note that the 3-dimensional rational vector $\widetilde{\mathbf{m}}_a$ is not defined as the stationary velocity $\mathbf{m}_a$ of F_a, though it plays essentially the same role. The term $u_a \mu_a$ creates the *residue* $\|\mu_a\|_2$ of F_a. Unless F_b can "destroy" this residue when it joins with F_a, one should not expect the flock formation time to grow exponentially. The crux is then to show that only a flock F_b with many birds can perform such a task. The following result says that, if the flock F_b settles too early, its effect on the residue of F_a is negligible. The conditions on F_c stated here differ slightly from those for F_a to make them closed under composition. This lemma below also covers the case $n_b = 0$, when the transition from F_a to F_c is involves the addition of an edge within the same flock. (Here, too, we assume, without

loss of generality, that these additions occur only one at a time within the same flock.) We postpone the proof of this result.   □

LEMMA 6.7.   *If node b is well defined, then assume that $t_b < \log\log\tau$. Whether node b exists or not, given* (40),

$$v_c(t_c) = \mathbf{1}_{n_c}\widetilde{\mathbf{m}}_a + u_c\mu_c + \zeta_c\,,$$

where

$$\begin{cases} \|u_c\|_\infty = 1 \ \& \ u_c \geq \mathbf{0}\,; \\ \|\mu_a\|_2\,n^{-O(1)} \leq \|\mu_c\|_2 \leq \|\mu_a\|_2\,; \\ \|\zeta_c\|_F \leq n\|\zeta_a\|_F + e^{-\tau^2}\,. \end{cases}$$

Furthermore, if node b is well defined, then $\mathbf{m}_b = \widetilde{\mathbf{m}}_a \neq \mathbf{m}_a$.

Remark. It might be helpful to explain, at an intuitive level, the meaning of the three terms in the expression for $v_a(t_a)$, or equivalently $v_c(t_c)$: $\widetilde{\mathbf{m}}_a$ is a low-precision approximation of the stationary velocity $\mathbf{m}_a$; the vector $u_a\,\mu_a$ creates the residue; the remainder ζ_a is an error term. The term $\widetilde{\mathbf{m}}_a$ is a low-resolution component of the velocity that any other flock F_b has to share if it is to create small angles with F_a (the key to high flock formation times). Think of it as a shared velocity caused by, say, wind affecting all flocks in the same way. This component must be factored out from the analysis since it cannot play any role in engineering small angles. It reflects the "relativity principle" that only velocity *differences* matter. To create small angles with F_a, incoming flocks F_b must "attack" the residue vector $u_a\,\mu_a$. Of course, they could potentially take turns doing so; therefore, one should read the inequalities of the lemma as a repeat of (40). The lemma states a closure property: unless F_b brings many bits to the table (via a formation time at least $\log\log\tau$), conditions (40) will still hold. These conditions prevent the creation of small angles between flocks, and hence of huge formation times. In other words, flocks that settle too early cannot hope to dislodge the residue $\|\mu_a\|_2$. The reason is that this residue is shielded in three ways.

— It is too big for the error term ζ_a to interfere with it: compare $e^{-\tau n^{O(1)}}$ with $e^{-\tau^2 n^{-O(1)} + n^{O(1)}}$.
— It is too small to be affected by $\widetilde{\mathbf{m}}_a$: compare $\frac{1}{\tau}$ with a rational over $O(\log\log\tau)$ bits.
— All of its coordinates have the same sign ($u_a \geq \mathbf{0}$), so taking averages among them cannot cause any cancellations. This form of "enduring" positivity is the most remarkable aspect of residues.

By (40), the lemma's bounds imply that

$$e^{-\tau n^{O(1)}} \leq \|\mu_c\|_2 \leq \frac{1}{\tau} \quad \& \quad \|\zeta_c\|_F \leq e^{-\tau^2 n^{-O(1)} + n^{O(1)}}\,,$$

which brings us back to (40). If c has a parent c' and sibling b', then we can apply the lemma again. Note that composition will always be applied for the same value of τ, that is, one is that is not updated at each iteration. In other words, the first three lines of (40) are closed under composition. This closure property does not hold for the last two lines, however: we lose a polynomial factor at each iteration, which is conveniently hidden in the big-oh notation. So we may compose the lemma only $n^{O(1)}$ times if we are to avoid any visible loss in the bounds of (40). Since the forest has fewer than n^2 nodes, this means that, as long as its conditions are met, we can compose the lemma with ancestors of c to our heart's content and still get the full benefits of (40).

The provision that b might not be well defined allows us to handle nonbranching switch nodes with equal ease. A related observation is that nowhere do we use the

fact that t_a is the actual formation time of F_a. It could be replaced in (39) by any t'_a strictly between t_a and t_c. We thus trivially derive a "delayed" version of Lemma 6.7. We summarize its two features: (i) Lemma 6.7 can be composed iteratively as often as we need to; (ii) node a need not be an actual node of F but one introduced artificially along an edge of F.

LEMMA 6.8. *Let* $c_0, \ldots, c_l$ *be an ascending path in* F *and let* d_i *be the sibling, if any, of* c_i. *Assume that* c_0, *possibly an artificial node, satisfies the conditions of node* a *in* (40) *and that* $t_{d_i} < \log \log \tau$ *for all* d_i. *Then,*

$$v_{c_i}(t_{c_i}) = \mathbf{1}_{n_{c_i}} \widetilde{\mathbf{m}}_{c_i} + u_{c_i} \mu_{c_i} + \zeta_{c_i},$$

where

$$\begin{cases} \ell(\widetilde{\mathbf{m}}_{c_i}) = O(\log \log \tau); \\ \|u_{c_i}\|_\infty = 1 \ \& \ u_{c_i} \geq \mathbf{0}; \\ e^{-\tau n^{O(1)}} \leq \|\mu_{c_i}\|_2 \leq \frac{1}{\tau}; \\ \|\zeta_{c_i}\|_F \leq e^{-\tau^2 n^{-O(1)} + n^{O(1)}}. \end{cases}$$

For all d_i, $\mathbf{m}_{d_i} = \widetilde{\mathbf{m}}_{c_i} \neq \mathbf{m}_{c_i}$.

We are now equipped with the tools we need to prove Lemma 6.2. Recall that $a_0, \ldots, a_k$ ($k > 1$) is an ascending path in F_o and b_i denotes the unique sibling of a_i. (Note that $a_0 \cdots a_k$ is a path in F_o whereas, in Lemma 6.8, $c_0 \cdots c_l$ is a path in F.) Also,

$$2^{2^{t_f}} < \log \log \log t_{a_k} < t_{a_0}^4 < t_{a_1} < \log t_{a_k}.$$

Assume, by contradiction, that $t_{b_i} < \sqrt{\log \log t_{a_0}}$ for $i = 0, \ldots, k-1$. As we observed earlier, Lemma 5.1 ensures that the ratio between the formation times of any node in F and that of its parent, if nonbranching, is at least $2^{-n^{O(1)}}$. Since there are fewer than n^2 switches, this implies that F_{a_0} can undergo switches only between t_{a_0} and $t_{a_0} 2^{n^{O(1)}}$. Because $t_{a_1} > t_{a_0}^4 > 2^{2^{t_f}}$, with $t_f = n^{\Theta(n^2)}$ for a large enough constant in the exponent, this shows that the entire time interval $[\frac{1}{2} t_{a_1}, t_{a_1})$ is free of switches. Let a be the last node in F from a_0 to a_1 and let c_0 be the artificial parent of a corresponding to the flock F_a at time $t_{a_1} - 1$: we set $n_{c_0} = n_{a_0}$ and $t_{c_0} = t_{a_1} - 1$. The bound in (30) ensures that the oscillations in the flock F_{c_0} are heavily damped. Indeed,

$$v_{c_0}(t_{c_0}) = \mathbf{1}_{n_{c_0}} \mathbf{m}_{c_0} + \zeta_{c_0}, \tag{41}$$

where, because of the magnitude of t_{a_1},

$$\|\zeta_{c_0}\|_F \leq e^{-(t_{a_1}/2 - 1)n^{-O(1)} + O(\log n)} \leq e^{-t_{a_1} n^{-O(1)}} \leq e^{-\tau^2}, \tag{42}$$

where $\tau \overset{\text{def}}{=} \frac{1}{2} t_{a_1}^{1/3}$. The rest of the sequence $\{c_i\}$ is now entirely specified. In particular, $c_1 = a_1$, $d_0 = b_0$, and $c_l = a_k$ ($l \geq k$). By extension, $\mathbf{m}_{c_0} = \mathbf{m}_a$; so, by (32),

$$\|\mathbf{m}_{c_0} - \mathbf{m}_{b_0}\|_2 \leq \frac{\log t_{a_1}}{t_{a_1}} n^{O(n^2)} < \frac{1}{\tau}.$$

Therefore, $\mathbf{m}_{c_0} = \mathbf{m}_{b_0} + \mu_{c_0}$, where

$$\|\mu_{c_0}\|_2 < \frac{1}{\tau}. \tag{43}$$

As we shall see, the presence of the square τ^2 in the exponent of (42) ensures that the oscillations of F_{c_0} are too small to interfere with the residue $\|\mu_{c_0}\|_2$. Writing $\widetilde{\mathbf{m}}_{c_0} = \mathbf{m}_{b_0}$, it follows from (41) that

$$v_{c_0}(t_{c_0}) = \mathbf{1}_{n_{c_0}} \widetilde{\mathbf{m}}_{c_0} + \mathbf{1}_{n_{c_0}} \mu_{c_0} + \zeta_{c_0},$$

which matches (39), with $u_{c_0} = \mathbf{1}_{n_{c_0}}$. Since all the nodes d_i are of the form b_{j_i},

$$t_{d_i} < \sqrt{\log\log t_{a_0}} < \log\log\tau.$$

Thus, we will be able to apply Lemma 6.8 once we verify that all conditions in (40) are met.

—[$2^{t_f} \le \tau \le t_{c_0}^{1/3}$]. This follows from our setting $\tau = \frac{1}{2}(t_{c_0}+1)^{1/3}$ and our assumption that $t_{a_1} > 2^{2^{t_f}}$.

—[$\ell(\widetilde{\mathbf{m}}_{c_0}) = O(\log\log\tau)$]. Because $\tau > 2^{2^{t_f-2}}$,

$$\sqrt{\log\log t_{a_0}}\, n \log n < (\log\log t_{a_0})^{2/3} = o(\log\log\tau).$$

The desired bound follows from (31):

$$\ell(\widetilde{\mathbf{m}}_{c_0}) = \ell(\mathbf{m}_{b_0}) = O(t_{b_0} n \log n) = O(\sqrt{\log\log t_{a_0}}\, n \log n) < \log\log\tau.$$

—[$e^{-\tau n^{O(1)}} \le \|\mu_{c_0}\|_2 \le \frac{1}{\tau}$]. The upper bound comes from (43). For the lower bound, note that $\mathbf{m}_{c_0} = \mathbf{m}_a$, with $t_a \le t_{a_0} 2^{n^{O(1)}}$. Another application of (31) shows that

$$\ell(\mathbf{m}_{c_0}) = O(t_a n \log n) < t_{a_0}^{7/6} < \tau.$$

We just saw that $\ell(\mathbf{m}_{b_0}) < \log\log\tau$, so $\mu_{c_0} = \mathbf{m}_{c_0} - \mathbf{m}_{b_0}$ is a three-dimensional row vector with rational coordinates over fewer than 2τ bits. The lower bound follows from the fact that $\mu_{c_0} \ne 0$. Indeed, the stationary velocities $\mathbf{m}_{c_0}$ and $\mathbf{m}_{b_0}$ cannot be equal as we now show. By $t_{b_0} < \sqrt{\log\log t_{a_0}} < t_a$ and the fact that, as observed earlier, F_{a_0} can undergo switches only between t_{a_0} and $t_{a_0} 2^{n^{O(1)}}$, we know that $\max\{t_{b_0}, t_a\} \le t_{a_0} 2^{n^{O(1)}}$. But, by Lemma 5.1, the switch at t_{a_1} must take place before time

$$\max\{t_{b_0}, t_a\} 2^{n^{O(1)}} = t_{a_0} 2^{n^{O(1)}} < t_{a_1},$$

which is a contradiction.

—[$\|u_{c_0}\|_\infty = 1$ & $u_{c_0} \ge \mathbf{0}$ & $\|\zeta_{c_0}\|_F \le e^{-\tau^2 n^{-O(1)} + n^{O(1)}}$]: The bounds follow from (42) and $u_{c_0} = \mathbf{1}_{n_{a_0}}$.

Recall that $c_l = a_k$. By applying Lemma 6.8 at c_l, we find that $\mathbf{m}_{b_{k-1}} = \widetilde{\mathbf{m}}_{c_{l-1}}$. Applying the same lemma now at node c_{l-1} shows that

$$v_{c_{l-1}}(t_{c_{l-1}}) = \mathbf{1}_{n_{c_{l-1}}} \widetilde{\mathbf{m}}_{c_{l-1}} + u_{c_{l-1}} \mu_{c_{l-1}} + \zeta_{c_{l-1}},$$

where

$$\begin{cases} \|\mu_{c_{l-1}}\|_2 \ge e^{-\tau n^{O(1)}} \\ \|u_{c_{l-1}}\|_\infty = 1 \ \& \ u_{c_{l-1}} \ge \mathbf{0}; \\ \|\zeta_{c_{l-1}}\|_F \le e^{-\tau^2 n^{-O(1)} + n^{O(1)}}. \end{cases}$$

The lemma also allows us to express the stationary velocity at c_{l-1}:

$$
\begin{aligned}
\mathbf{m}_{c_{l-1}} &= \pi_{c_{l-1}}^T v_{c_{l-1}}(t_{c_{l-1}}) \\
&= \pi_{c_{l-1}}^T (\mathbf{1}_{n_{c_{l-1}}} \widetilde{\mathbf{m}}_{c_{l-1}} + u_{c_{l-1}} \mu_{c_{l-1}} + \zeta_{c_{l-1}}) \\
&= \mathbf{m}_{b_{k-1}} + \pi_{c_{l-1}}^T u_{c_{l-1}} \mu_{c_{l-1}} + \pi_{c_{l-1}}^T \zeta_{c_{l-1}}.
\end{aligned}
$$

By the triangle inequality, it follows that

$$
\begin{aligned}
\|\mathbf{m}_{c_{l-1}} - \mathbf{m}_{b_{k-1}}\|_2 &\geq \|\pi_{c_{l-1}}^T u_{c_{l-1}} \mu_{c_{l-1}}\|_2 - \|\pi_{c_{l-1}}^T \zeta_{c_{l-1}}\|_2 \\
&\geq \pi_{c_{l-1}}^T u_{c_{l-1}} \|\mu_{c_{l-1}}\|_2 - \|\pi_{c_{l-1}}\|_2 \|\zeta_{c_{l-1}}\|_F \\
&\geq \min_i \{(\pi_{c_{l-1}})_i\} e^{-\tau n^{O(1)}} - e^{-\tau^2 n^{-O(1)} + n^{O(1)}} \geq e^{-\tau n^{O(1)}}.
\end{aligned}
$$

By (32),

$$
\|\mathbf{m}_{c_{l-1}} - \mathbf{m}_{b_{k-1}}\|_2 \leq \frac{\log t_{a_k}}{t_{a_k}} n^{O(n^2)};
$$

therefore, since $t_{a_k} > 2^{t_f}$,

$$
t_{a_k} \leq \|\mathbf{m}_{c_{l-1}} - \mathbf{m}_{b_{k-1}}\|_2^{-2} \leq e^{\tau n^{O(1)}} \leq e^{\tau^{1.5}},
$$

which contradicts our assumption that $\tau = \frac{1}{2} t_{a_1}^{1/3} < (\log t_{a_k})^{1/3}$.

PROOF OF LEMMA 6.7. Using the shorthand $u^a = P_a^{t_c - t_a} u_a$ and $\zeta^a = P_a^{t_c - t_a} \zeta_a$, we express the velocity of the flock F_a at time t_c. From

$$
v_a(t_c) = P_a^{t_c - t_a} v_a(t_a),
$$

we find that, by (39),

$$
\begin{aligned}
v_a(t_c) &= P_a^{t_c - t_a} \mathbf{1}_{n_a} \widetilde{\mathbf{m}}_a + P_a^{t_c - t_a} u_a \mu_a + \zeta^a \\
&= \mathbf{1}_{n_a} \widetilde{\mathbf{m}}_a + u^a \mu_a + \zeta^a.
\end{aligned}
\tag{44}
$$

Because P_a is an averaging operator, $\|P_a^{t_c - t_a} u_a\|_\infty \leq \|u_a\|_\infty = 1$. The vector u_a is non-negative, so

$$
\begin{aligned}
\|P_a^{t_c - t_a} u_a\|_\infty &\geq \frac{1}{n_a} \|P_a^{t_c - t_a} u_a\|_1 = \frac{1}{n_a} \mathbf{1}_{n_a}^T P_a^{t_c - t_a} u_a \geq \frac{1}{n_a} \pi_a^T P_a^{t_c - t_a} u_a \\
&\geq \frac{1}{n_a} \pi_a^T u_a \geq \frac{1}{n_a} \min_i \{(\pi_a)_i\} \|u_a\|_\infty \geq n^{-O(1)}.
\end{aligned}
$$

Similarly, by stochasticity and submultiplicativity,

$$
\|P_a^{t_c - t_a} \zeta_a\|_F \leq \sqrt{n_a}\, \|\zeta_a\|_F ;
$$

therefore,

$$
\begin{cases}
n^{-O(1)} \leq \|u^a\|_\infty \leq 1 \ \& \ u^a \geq \mathbf{0}; \\
\|\zeta^a\|_F \leq n \|\zeta_a\|_F .
\end{cases}
\tag{45}
$$

$\square$

Case I. Node b is well defined and $t_b < \log\log \tau$: Since, by (40), $t_c > t_a \geq \tau^3 \geq 8^{t_f}$, with $t_f = n^{(bn)^2}$,

$$
-(t_c - t_b) n^{-O(1)} + \Theta(n^3) \leq -\tau^2;
$$

so, by applying (30) to the flock F_b, we find that

$$\|v_b(t_c) - \mathbf{1}_{n_b}\mathbf{m}_b\|_F \leq e^{-(t_c-t_b)n^{-O(1)}+O(\log n)};$$

hence

$$v_b(t_c) = \mathbf{1}_{n_b}\mathbf{m}_b + e^{-\tau^2}z_c,$$

where $\|z_c\|_F \leq 1$. It follows from (44) that

$$v_c(t_c) = \begin{pmatrix} v_a(t_c) \\ v_b(t_c) \end{pmatrix} = \begin{pmatrix} \mathbf{1}_{n_a}\widetilde{\mathbf{m}}_a \\ \mathbf{1}_{n_b}\mathbf{m}_b \end{pmatrix} + \begin{pmatrix} u^a \\ \mathbf{0} \end{pmatrix}\mu_a + \begin{pmatrix} \zeta^a \\ e^{-\tau^2}z_c \end{pmatrix}. \qquad (46)$$

By (39), the stationary velocity of F_a is equal to

$$\begin{aligned} \mathbf{m}_a &= \pi_a^T v_a(t_a) = \pi_a^T(\mathbf{1}_{n_a}\widetilde{\mathbf{m}}_a + u_a\mu_a + \zeta_a) \\ &= \widetilde{\mathbf{m}}_a + (\pi_a^T u_a)\mu_a + \pi_a^T\zeta_a. \end{aligned} \qquad (47)$$

By the triangle inequality, it follows that

$$\begin{aligned} \|\mathbf{m}_a - \widetilde{\mathbf{m}}_a\|_2 &\geq \pi_a^T u_a\|\mu_a\|_2 - \|\pi_a\|_2\|\zeta_a\|_F \\ &\geq \min_i\{(\pi_a)_i\}e^{-\tau n^{O(1)}} - e^{-\tau^2 n^{-O(1)}+n^{O(1)}} \geq e^{-\tau n^{O(1)}}; \end{aligned}$$

which shows that

$$\mathbf{m}_a \neq \widetilde{\mathbf{m}}_a. \qquad (48)$$

Note also that, by (47),

$$\begin{aligned} \|\widetilde{\mathbf{m}}_a - \mathbf{m}_b\|_2 &\leq \|\widetilde{\mathbf{m}}_a - \mathbf{m}_a\|_2 + \|\mathbf{m}_a - \mathbf{m}_b\|_2 \\ &\leq \pi_a^T u_a\|\mu_a\|_2 + \|\pi_a\|_2\|\zeta_a\|_F + \|\mathbf{m}_a - \mathbf{m}_b\|_2. \end{aligned}$$

We bound each term on the right-hand side: by (40) and Cauchy-Schwarz,

$$\pi_a^T u_a\|\mu_a\|_2 \leq \tfrac{1}{\tau}\|\pi_a\|_2\|u_a\|_2 \leq \tfrac{1}{\tau}\sqrt{n_a}\,\|u_a\|_\infty \leq \tfrac{1}{\tau}\sqrt{n}.$$

By (32) and $t_c > \tau^3 \geq 8^{t_f}$,

$$\|\mathbf{m}_a - \mathbf{m}_b\|_2 \leq \tfrac{1}{\tau}.$$

Also, $\|\pi_a\|_2 = O(1)$ and, by (40), $\|\zeta_a\|_F \leq e^{-\tau^2 n^{-O(1)}+n^{O(1)}}$; therefore

$$\|\widetilde{\mathbf{m}}_a - \mathbf{m}_b\|_2 < \sqrt{\tfrac{1}{\tau}}.$$

By (31), our assumption that $t_b < \log\log\tau$ implies that

$$\ell(\mathbf{m}_b) = O(n(\log n)\log\log\tau) < (\log\log\tau)^2.$$

Since, by (40), $\ell(\widetilde{\mathbf{m}}_a) = O(\log\log\tau)$, the squared distance $\|\widetilde{\mathbf{m}}_a - \mathbf{m}_b\|_2^2$ is a rational over $O(\log\log\tau)^2$ bits: being less than $1/\tau$ implies that it is actually zero; hence $\widetilde{\mathbf{m}}_a = \mathbf{m}_b$, as claimed in the lemma. We verify from (45) that

$$\mu_c \stackrel{\text{def}}{=} \mu_a\|u^a\|_\infty \qquad \text{and} \qquad u_c \stackrel{\text{def}}{=} \begin{pmatrix} u^a \\ \mathbf{0} \end{pmatrix}\|u^a\|_\infty^{-1}$$

satisfy the conditions of the lemma. By (46),

$$v_c(t_c) = \mathbf{1}_{n_c}\widetilde{\mathbf{m}}_a + u_c\mu_c + \zeta_c,$$

where

$$\zeta_c = \begin{pmatrix} \zeta^a \\ e^{-\tau^2} z_c \end{pmatrix}.$$

By (45) and $\|z_c\|_F \le 1$, the lemma's condition on ζ_c is trivially satisfied.

Case II. Node b is not defined: We set $\zeta_c = \zeta^a$; $\mu_c = \mu_a \|u^a\|_\infty$; and $u_c = u^a \|u^a\|_\infty^{-1}$. This matches the identity (44) with the one claimed in the lemma.

PROOF OF LEMMA 6.5. Suppose that Bob does not always follow the single-iterated log rule. We show how to force him to do so without decreasing the score differential. If Bob uses the rule $t_B \leftarrow \log \log t_B$, then Alice follows suit with $t_A \leftarrow (\log \log t_A)^\alpha$. We break this round into two parts:

(1) $t_B \leftarrow \log t_B$ and $t_A \leftarrow \log t_A$;
(2) $t_B \leftarrow \log t_B$ and $t_A \leftarrow (\log t_A)^\alpha$.

We proceed similarly for higher log-iterations and apply the modification systematically. This transformation increases the scores of the players but it does not change their difference. Finally, we apply one last transformation to the new game, which is to convert all of Alice's moves into $t_A \leftarrow (\log t_A)^{1/4}$. This brings Bob's and Alice's moves in conformity with Lemma 6.5 without decreasing the score differential. □

PROOF OF LEMMA 6.6. Consider the two recurrence relations:

$$a_0(x) = b_0(x) = x,$$

and, for $h > 0$,

$$\begin{cases} a_h(x) = 2^{a_{h-1}(x)} \\ b_h(x) = 2^{b_{h-1}(x)} + 2. \end{cases}$$

Recall that c_h is defined by $c_0 = t_0^4$ and, for $h > 0$, $c_h = 2^{4c_{h-1}}$. We verify by induction that, for any $h > 0$,

$$c_h = 2^{2^{b_{h-1}(4\log t_0 + 2)}}.$$

To prove the inequality, we seek,

$$\min\{ h \mid c_h \ge t \} \ge \log^* t - \log^* t_0 - O(1),$$

where $t > t_0$, we may assume that $t > 2^{t_0}$, otherwise the result is trivial. The assumption implies that the minimum h is positive; therefore it suffices to prove that, for all $h \ge 0$,

$$b_h(4\log t_0 + 2) \le a_h(4\log t_0 + 4). \tag{49}$$

We prove by induction that, for all $h \ge 0$, $x \ge 2$, and $\varepsilon > 0$,

$$a_h(x) + \varepsilon \le a_h(x + \varepsilon 2^{-h}). \tag{50}$$

The case $h = 0$ is obvious, so consider $h > 0$. Note that, for any $y \ge 2$,

$$2^y + \varepsilon \le 2^{y + \varepsilon/2},$$

which follows from

$$\ln(1 + \varepsilon 2^{-y}) \le \varepsilon 2^{-y} \le \tfrac{\ln 2}{2}\, \varepsilon.$$

Since $a_{h-1}(x) \ge 2$, this shows that

$$a_h(x) + \varepsilon = 2^{a_{h-1}(x)} + \varepsilon \le 2^{a_{h-1}(x) + \varepsilon/2} \le 2^{a_{h-1}(x + \varepsilon 2^{-h})} = a_h(x + \varepsilon 2^{-h}),$$

which proves (50). Next, we show by induction that, for all $h \geq 0$ and $x \geq 2$,

$$b_h(x) \leq a_h(x + 2 - 2^{1-h}). \tag{51}$$

The case $h = 0$ again being obvious, assume that $h > 0$. By (50),

$$b_h(x) = 2^{b_{h-1}(x)} + 2 \leq 2^{a_{h-1}(x+2-2^{2-h})} + 2$$
$$\leq a_h(x + 2 - 2^{2-h}) + 2 \leq a_h(x + 2 - 2^{1-h}),$$

which establishes (51); and hence (49). □

ACKNOWLEDGMENTS

REFERENCES

Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., Lecomte, V., Orlandi, A., Parisi, G., Procaccini, A., Viale, M., and Zdravkovic, V. 2008. Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proc. Nat. Acad. Sci. 105*, 1232–1237.

Bialek, W., Cavagna, A., Giardina, I., Mora, T., Silvestri, E., Viale, M., and Walczak, A. 2012. Statistical mechanics for natural flocks of birds. *Proc. Nat. Acad. Sci. 109*, 4786–4791.

Camazine, S., Deneubourg, J.-L., Franks, N., Sneyd, J., Theraulaz, G., and Bonabeau, E. 2003. *Self-Organization in Biological Systems*. Princeton University Press, Princeton, N.J.

Canny, J. 1988. *Complexity of Robot Motion Planning*. The MIT Press.

Cavagna, A., Cimarelli, A., Giardina, I., Orlandi, A., Parisi, G., Santagati, R., Stefanini, F., and Viale, M. 2010. Scale-free correlations in starling flocks. *Proc. Nat. Acad. Sci. 107*, 11865–11870.

Cavagna, A., Giardina, I., Orlandi, A., Parisi, G., and Procaccini, A. 2008a. The starflag handbook on collective animal behaviour: 2. Three-dimensional analysis. *Animal Behav. 76*, 237–248.

Cavagna, A., Giardina, I., Orlandi, A., Parisi, G., Procaccini, A., Viale, M., and Zdravkovic, V. 2008b. The starflag handbook on collective animal behaviour: 1. empirical methods. *Animal Behav. 76*, 217–236.

Chazelle, B. 2009. The convergence of bird flocking. arXiv:0905.4241v1, May 2009. Preliminary version in *Proceedings SODA'09*, 422–431.

Chazelle, B. 2011. The total s-energy of a multiagent system. *SIAM J. Control Optim. 49*, 1680–1706.

Chazelle, B. 2012a. The dynamics of influence systems. In *Proceedings of the 53rd IEEE FOCS*, 311–320.

Chazelle, B. 2012b. Natural algorithms and influence systems. *Commun. ACM 55*, 101–110.

Chung, F. 1997. *Spectral Graph Theory*. CBMS Regional Conference Series in Mathematics. American Mathematical Society, Providence, R.I.

Cucker, F. and Smale, S. 2007. Emergent behavior in flocks. *IEEE Trans. Automatic Control 52*, 852–862.

Hendrickx, J. and Blondel, V. 2006. Convergence of different linear and non-linear vicsek models. In *Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems (MTNS'06)*. 1229–1240.

Jadbabaie, A., Lin, J., and Morse, A. 2003. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Trans. Automatic Control 48*, 988–1001.

Ji, M. and Egerstedt, M. 2007. Distributed coordination control of multi-agent systems while preserving connectedness. *IEEE Trans. Robot. 23*, 693–703.

Kemeny, J. and Snell, J. 1983. *Finite Markov Chains*. Springer.

Kim, H. and Huse, D. 2012. Superdiffusive nonequilibrium motion of an impurity in a fermi sea. *Phys. Rev. A 85*.

Landau, H. and Odlyzko, A. 1981. Bounds for eigenvalues of certain stochastic matrices. *Lin. Algebra Its Appl. 38*, 5–15.

Li, S. and Wang, H. 2004. Multi-agent coordination using nearest neighbor rules: Revisiting the Vicsek model. arXiv:cs/0407021v2.

Moreau, L. 2005. Stability of multiagent systems with time-dependent communication links. *IEEE Trans. Automatic Control 50*, 169–182.

Moshtagh, N., Jadbabaie, A., and Daniilidis, K. 2005. Distributed geodesic control laws for flocking of non-holonomic agents. In *Proceedings of the 44th IEEE Conference on Decision and Control and European Control Conference (ECC-CDC'05)*.

Olfati-Saber, R. 2006. Flocking for multi-agent dynamic systems: Algorithms and theory. *IEEE Trans. Automatic Control 51*, 401–420.

Reynolds, C. 1987. Flocks, herds, and schools: A distributed behavioral model. *Comput. Graph. 21*, 25–34.

Seneta, E. 2006. *Non-Negative Matrices and Markov Chains*. 2nd Ed., Springer.

Shi, H., Wang, L., and Chu, T. 2005. Coordination of multiple dynamic agents with asymmetric interactions. In *Proceedings of the IEEE International Symposium on Intelligent Control*. 1423–1428.

Tahbaz-Salehi, A. and Jadbabaie, A. 2007. Bounds for eigenvalues of certain stochastic matrices. On recurrence of graph connectivity in Vicsek's model of motion coordination for mobile autonomous agents. In *Proceedings of the American Control Conference (ACC'07)*. 699–704.

Tang, G. and Guo, L. 2007. Convergence of a class of multi-agent systems in probabilistic framework. *J. Syst. Sci. Complex. 20*, 173–197.

Tanner, H., Jadbabaie, A., and Pappas, G. 2003a. Stable flocking of mobile agents, Part I: Fixed topology. In *Proceedings of the IEEE Conference on Decision and Control*. 2010–2015.

Tanner, H., Jadbabaie, A., and Pappas, G. 2003b. Stable flocking of mobile agents, Part II: Dynamic topology. In *Proceedings of the IEEE Conference on Decision and Control*, Maui, Hawaii. 2016–2021.

Tanner, H., Jadbabaie, A., and Pappas, G. 2007. Flocking in fixed and switching networks. *IEEE Trans. Automatic Control 52*, 863–868.

Toner, J. and Tu, Y. 1995. Long range order in a twodimensional dynamical xy model: How birds fly together. *Phys. Rev. Lett. 75*, 4326–4329.

Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., and Shochet, O. 1995. Novel type of phase transition in a system of self-driven particles. *Phys. Rev. Lett. 75*, 1226–1229.

Yap, C. 2000. *Fundamental Problems of Algorithmic Algebra*. Oxford University Press, Inc., New York.