

A Connectivity-Sensitive Approach to Consensus Dynamics

Bernard Chazelle 

Department of Computer Science, Princeton University, NJ, USA

Kritikorn Karntikoon 

Department of Computer Science, Princeton University, NJ, USA

Abstract

The paper resolves a long-standing open question in network dynamics. Averaging-based consensus has long been known to exhibit an exponential gap in relaxation time between the connected and disconnected cases, but a satisfactory explanation has remained elusive. We provide one by deriving nearly tight bounds on the *s-energy* of disconnected systems. This in turn allows us to relate the convergence rate of consensus dynamics to the number of connected components. We apply our results to opinion formation in social networks and provide a theoretical validation of the concept of an *Overton window* as an attracting manifold of “viable” opinions.

2012 ACM Subject Classification Mathematics of computing → Combinatoric problems; Theory of computation → Dynamic graph algorithms

Keywords and phrases *s-energy*, dynamic networks, relaxation time, multiagent systems

Digital Object Identifier 10.4230/LIPIcs.SAND.2023.10

Funding This work was supported in part by NSF grant CCF-2006125.

1 Introduction

Consensus dynamics based on local averaging has been the object of considerable attention [5, 8, 9, 10, 11, 21, 26, 31, 32, 33, 34, 36, 37, 39, 40, 41, 42, 45, 48, 47, 50, 54]. This owes largely to the ubiquity of these systems, from flocking and swarming to synchronization, social epistemology, and opinion dynamics [7, 12, 18, 21, 22, 26, 32, 34, 37, 55]. Agents interact across a network by averaging their state variables with those of their neighbors. Under mild conditions, such systems are known to converge to a fixed-point attractor.

When the network is fixed, the dynamics is dual to a Markov chain, which puts a wealth of analytical tools at our disposal. It is well known that convergence within ε is reached in time $C \log(1/\varepsilon)$, for some parameter C depending only on the graph’s size and topology. This bound still holds for time-varying graphs as long as they remain connected at all times. When connectivity is not guaranteed, however, the convergence time shoots up to $C \log(1/\varepsilon)^{n-1}$. This exponential jump has been a puzzling mystery in the field of time-varying network dynamics [13, 26, 34, 50, 51]. Recent works on oblivious message adversaries also exhibit exponential gaps in the time complexity of certain broadcast and consensus problems [17, 23, 24, 25, 57]. The gap is also behind the emergence of hyper-torpid mixing in *Markov influence systems* and the *slow-clock* phenomenon [14].

This paper explains the exponential jump in consensus dynamics by relating it to the number of connected components in the system. The convergence rate is shown to be of the form $C(\log 1/\varepsilon)^m$, where $m < n$ is the maximum number of connected components at any time. We derive quasi-optimal bounds on the parameter C . In addition, we look at three important special cases – *reversible*, *expanding*, *random* – and we discuss applications to opinion formation in social networks. The results in this work rely on new *s-energy* bounds of independent interest. The *s-energy* is a generating function designed specifically



© Bernard Chazelle and Kritikorn Karntikoon;

licensed under Creative Commons License CC-BY 4.0

2nd Symposium on Algorithmic Foundations of Dynamic Networks (SAND 2023).

Editors: David Doty and Paul Spirakis; Article No. 10; pp. 10:1–10:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

for the analysis of networked averaging systems [12, 13]. Its main purpose is to overcome the technical difficulties one encounters when systems become disconnected and one does not have apriori bounds on how long they might stay so. The difficulty is fundamental: With changing topologies, networked systems cease to have coherent eigenmodes and spectral techniques break down. In other words, when linear algebra fails, the s -energy gives us a way out.

Averaging dynamics

Let $(G_t)_{t=1}^\infty$ be an infinite sequence of graphs over the vertex set $[n]$. Each graph has a self-loop. Let P_t be the stochastic matrix of a random walk over G_t . By construction, a matrix entry is positive if and only if it corresponds to an edge of G_t . Each row sums up to 1 and the diagonal is positive. We assume that the nonzero entries in P_t are at least some fixed $\rho \in (0, 1/2]$.¹ Let $P_{\leq t}$ denote the product $P_t \cdots P_1$. The set of orbits $(P_{\leq t}x)_{t \geq 0}$, over all $x \in \mathbb{R}^n$, forms an *averaging system*, often called *consensus dynamics* in the literature [26]. When all the matrices $P_t = P$ are identical, the map $x \mapsto Px$ is the dual map of the Markov chain $(y \mapsto yP)$ and its convergence time is the chain's *mixing time*. The case of a fixed matrix has been studied exhaustively, so the novelty of the paper comes from the dynamic nature of the networks.

- A *general* averaging system (*genS*) assumes only that the graphs G_t are undirected.
- A *reversible* averaging system (*revS*) is a *genS* whose individual Markov chains P_t are reversible and share the same stationary distribution. This means that $P_t = \text{diag}(q)^{-1} M_t$, where M_t is symmetric with nonzero entries at least 1 and $q = M_t \mathbf{1} \preceq \mathbf{1}/\rho$.
- An *expanding* averaging system (*expS*) is a *revS* where $q = d\mathbf{1}$ and the connected components of each G_t are d -regular *expanders*. Recall that a d -regular expander is a graph of degree d such that, for any set X of at most half of the vertices, we have $|\partial X| \geq h|X|$, where ∂X is the set of edges with exactly one vertex in X ; the factor h is called the Cheeger constant.
- A *random* averaging system (*ranS*) assumes that the graphs G_t are d -regular and random.

Our results

Given $x \in [0, 1]^n$, each point of the orbit $(P_{\leq t}x)_{t \geq 0}$ corresponds to an embedding of the graph G_t over the reals. Let T_ε be the number of timesteps t at which G_t has an embedded edge of length at least $\varepsilon > 0$. We denote by $T_{m,\varepsilon}$ the maximum value of T_ε over all graph sequences $(G_t)_{t \geq 0}$ such that no graph has more than m connected components.² We use superscripts to distinguish among the general, reversible, expanding, and random cases: $T^{\text{GEN}}, T^{\text{REV}}, T^{\text{EXP}}$ and T^{RAN} respectively.

For simplicity, $T^{\text{REV}}, T^{\text{EXP}}$ and T^{RAN} do not assume that the initial diameter is bounded by 1 but, rather, that the initial variance is. Here we define the (scaled) variance as $\|x - \hat{x}\|_q^2$, where x is shorthand for $x(1)$ and (i) $\hat{x}$ is the mean initial position $\|q\|_1^{-1} \langle x, \mathbf{1} \rangle_q \mathbf{1}$; (ii) $\|x\|_q^2 := \langle x, x \rangle_q$; and (iii) $\langle x, y \rangle_q := \sum_i q_i x_i y_i$.

¹ If $\rho > 1/2$, the only edges of G_t have to be self-loops, which is of no interest.

² Since the systems always converge to a fixed-point attractor, the reader might wonder why we do not define T_ε as the time past which no edge length exceeds ε . This would not work because an adversary could always insert the identity matrix repeatedly to delay convergence at will and push T_ε to infinity.

Note that unit diameter implies a variance of at most n/ρ and, conversely, unit variance implies a diameter bounded by 2. Thus, the following bounds can be easily scaled to accommodate either assumption about initial conditions.

► **Theorem 1.** *For some constant $c > 0$ and any positive ε small enough,*

$$\left\{ \begin{array}{l} T_{m,\varepsilon}^{\text{GEN}} \leq c(1/\rho)^{n-1} \left(mn \log \frac{1}{\varepsilon} \right)^m \\ T_{m,\varepsilon}^{\text{REV}} \leq c \left(\frac{n^2}{\rho} \log \frac{1}{\varepsilon} \right)^m \quad [15] \\ T_{m,\varepsilon}^{\text{EXP}} \leq c \left(\frac{d^3 mn}{h^2} \log \frac{1}{\varepsilon} \right)^m \quad \& \quad T_{1,\varepsilon}^{\text{EXP}} \leq \frac{cd^3}{h^2} \log \frac{1}{\varepsilon} \\ \mathbb{E} T_{1,\varepsilon}^{\text{RAN}} \leq c \log \frac{1}{\varepsilon}. \end{array} \right.$$

Proof. The upper bound on $T_{m,\varepsilon}^{\text{REV}}$ was proven in [15] and is mentioned here for completeness. Note that the case $m = 1$ of *genS* and *revS* recover the classic mixing times for Markov chains ($P_t = P$), in particular the polynomial vs. exponential gap between general and reversible chains. We also rediscover the logarithmic bound for expanders ($m = 1$). Previous work addressed only the cases $m = 1$ or $m = n - 1$. What made any connectivity-sensitive extension challenging is that the proof techniques for the s -energy do not seem to generalize. It is often possible to set up recurrence relations but these are too coarse to deliver good upper bounds. Intricate multiscale amortization arguments were used to overcome these limitations. In a surprising turn, we show how to rescue the divide-and-conquer approach via a new linearization technique. Before we turn to the s -energy $\mathcal{E}_{m,s}$, we need to mention its relevance: Theorem 1, indeed, follows from combining the corresponding s -energy bounds with the inequality

$$T_{m,\varepsilon} \leq \inf_{0 < s \leq 1} \varepsilon^{-s} \mathcal{E}_{m,s}. \quad (1)$$

The idea is to provide upper bounds on $\mathcal{E}_{m,s}$ for each of the four cases: general, reversible, expanding, and random. Proving Theorem 1 is then a matter of choosing s to minimize the right-hand side in (1). This step is straightforward calculus and, hence, omitted. ◀

The bounds in Theorem 1 are very general and can be applied to countless instances of real-world dynamics (swarming, flocking, polarization, power grid sync, firefly flashing, etc. [12]). We conclude this work in Section 3 with an application of our ideas to opinion formation in social networks. We extend the model to include directed edges so as to capture both evolving and fixed sources of information. We show that, while all opinions might keep changing forever, they will inevitably land in the convex hull of the fixed sources. Furthermore, we bound the time at which this must happen. Our result is a quantitative validation of the *Overton window* as an attracting manifold of “viable” opinions [4, 6, 19, 27, 46].

2 New Bounds on the s -Energy

Let $(G_t)_{t=1}^\infty$ be an infinite sequence of graphs over the vertex set $[n]$. A vertex is also called an *agent*. Each graph is embedded in $\mathbb{R}$ and we denote by $x_i(t)$ the position of agent i . The union of the embedded edges of G_t forms disjoint intervals, called *blocks*. Let $l_1, \dots, l_k$ be the lengths of these blocks and put $E_{s,t} = \sum_{i=1}^k l_i^s$, with $s \in (0, 1]$.³ We define the s -energy $E_s = \sum_{t \geq 1} E_{s,t}$ and we denote by $\mathcal{E}_{m,s}$ the supremum of E_s , over all initial agent positions $x \in [0, 1]^n$, under the constraint that G_t should have at most m connected components.

³ For example, if G_t consists of three edges embedded as $[0.1, 0.3]$, $[0.2, 0.4]$ and $[0.7, 0.8]$, then there are two blocks $[0.1, 0.4]$, $[0.7, 0.8]$ and $E_{s,t} = (0.3)^s + (0.1)^s$.

2.1 General averaging systems

We begin by stating our bounds on the s -energy of any $genS$ with at most m connected components. As stated above, we assume that the initial diameter Δ of the vertex positions is 1. If it is not, it suffices to multiply the bounds by a factor of Δ^s .

► **Theorem 2.** $\mathcal{E}_{m,s}^{\text{GEN}} \leq (c/s)^m (1/\rho)^{n-1}$, for any $s \in (0, 1]$, where $c = O(mn)$. For $s = 1$, the bound can be improved to $\mathcal{E}_{m,1}^{\text{GEN}} \leq 3en(1/\rho)^{\lfloor n/2 \rfloor}$.

Twist systems

A $genS$ is a special case of a *twist system* [13]. The latter is easier to analyze so we turn our attention to it. Relabel the agents so their positions $x_1 \leq \dots \leq x_n$ appear in sorted order at time t . A twist system moves them to positions $y_1 \leq \dots \leq y_n$ at time $t + 1$ in such a way that

$$(1 - \rho)x_u + \rho x_{\min\{i+1, v\}} \leq y_i \leq \rho x_{\max\{i-1, u\}} + (1 - \rho)x_v, \quad (2)$$

for any i in $[u, v]$ and $y_i = x_i$ otherwise. We repeat this step indefinitely. Twist systems are highly nondeterministic. At each step, a new interval $[u, v] \subseteq [n]$, called a *block*, is picked and the agents' motion is only constrained by (2) and the need to maintain their ranks (ie, agents never cross).

For the purposes of this work, we extend the concept to *m-twist systems* by stipulating, at each time t , a partition of $[n]$ into up to m_t blocks $[u_{t,l}, v_{t,l}]$ ($1 \leq l \leq m_t \leq m$). Each agent is now subject to (2) within its own enclosing block. We define the s -energy E_s^{TW} of a twist system as we did with a $genS$ by adding together the s -th powers of all the block lengths. We use the same notation with the addition of the superscript TW.

► **Lemma 3.** A $genS$ with at most m connected components at any time can be interpreted as an *m-twist system* with the same s -energy.

Proof. Fix a $genS$ and let $(x_i)_{i=1}^n$ and $(y_i)_{i=1}^n$ be the positions of the agents at times t and $t + 1$, given in nondecreasing order. We denote by x'_i the position of agent i at time $t + 1$. Let $[x_u, x_v]$ be a block of the $genS$ at time t . Pick $k < v$ and write $z = \rho x_k + (1 - \rho)x_v$. All the diagonal elements of P_t are at least ρ ; hence $x'_i \leq z$, for all $i \leq k$, and $y_k \leq z$. In fact, the inequality even holds for $i = k + 1$: Indeed, the embedded edges of G_t cover all of $[x_u, x_v]$, so at least one of them, call it (l, r) , must join $[u, k]$ to $[k + 1, v]$; hence $x'_r \leq z$. Our claim follows. This proves that, for all $i \in (u, v]$, $y_i \leq \rho x_{\max\{i-1, u\}} + (1 - \rho)x_v$. We omit the case $i = u$ and the mirror-image inequality, which repeat the same argument. Summing up all the powers $(x_v - x_u)^s$ shows the equivalence between the two s -energies. ◀

Proof of Theorem 2. We may assume that the agents stay within $[0, 1]$. We begin with showing the bound $\mathcal{E}_{m,1}^{\text{GEN}} \leq \mathcal{E}_{m,1}^{\text{TW}} \leq 3en(1/\rho)^{\lfloor n/2 \rfloor}$.

$$\boxed{\text{Case } s = 1}$$

We prove a stronger result by bounding $K_t(z) := \sum_{k=1}^n (x_{v(k)} - x_k) z^k$, where $v(k) = v_{t,l}$ for l such that $k \in [u_{t,l}, v_{t,l}]$. As usual, $0 \leq x_1 \leq \dots \leq x_n \leq 1$ denotes the sorted positions of the agents at time t ; we omit t for convenience but it is understood throughout. We define the *weighted 1-energy* $K(z) = \sum_{t \geq 0} K_t(z)$ and, finally, $\mathcal{K}(z) = \sup K(z)$. As long as $z \geq 1$, the 1-energy is obviously dominated by its weighted version. We improve this crude bound via a symmetry argument:

► **Lemma 4.** For any $z \geq 1$, $\mathcal{E}_{m,1}^{\text{TW}} \leq 2z^{-\nu} \mathcal{K}(z)$, where $\nu = \lceil n/2 \rceil$.

Proof. We define the mirror image of K_t as $\bar{K}_t(z) = \sum_{k=1}^n (x_k - x_{u(k)}) z^{n-k+1}$, where $u(k)$ is the left counterpart of $v(k)$. We have

$$\begin{aligned} E_{1,t}^{\text{TW}} &\leq \sum_{k \leq \nu} (x_k - x_{u(k)}) + \sum_{k \geq \nu} (x_{v(k)} - x_k) \\ &\leq z^{-\nu} \sum_{k \leq \nu} (x_k - x_{u(k)}) z^{n-k+1} + z^{-\nu} \sum_{k \geq \nu} (x_{v(k)} - x_k) z^k \\ &\leq z^{-\nu} (\bar{K}_t(z) + K_t(z)). \end{aligned}$$

Because $\mathcal{K}(z) = \sup K(z)$, the lemma then follows by summing up all $t > 0$. ◀

We define the polynomial $P_t(z) = \sum_{k=1}^n x_k z^k$ for $z > 1/\rho$ and exploit two simple but surprising facts: $P_t(z)$ cannot increase over time;⁴ and, at each step, the drop from $P_t(z)$ to $P_{t+1}(z)$ is at least proportional to $K_t(z)$. Thus, we develop a discrete version of the inference: $dP_t/dt \leq -cK_t$ implies

$$\int_{t \geq 1} cK_t \leq - \int_{t \geq 1} \frac{dP_t}{dt} \leq P_1.$$

► **Lemma 5.** For any $z > 1/\rho$, $P_t(z) - P_{t+1}(z) \geq (\rho z - 1)K_t(z)$.

Proof. The inequality is additive in the number of blocks so we can assume there is a single one $[u, v]$ at time t . Using the notation of (2), we have $y_k \leq \rho x_{\max\{k-1, u\}} + (1 - \rho)x_v$; hence

$$\begin{aligned} P_t(z) - P_{t+1}(z) &= \sum_{k=u}^v (x_k - y_k) z^k \geq \sum_{k=u}^v (x_k - \rho x_{\max\{k-1, u\}} - (1 - \rho)x_v) z^k \\ &\geq (\rho - 1)(x_v - x_u) z^u + \sum_{k=u+1}^v \rho(x_v - x_{k-1}) z^k - \sum_{k=u+1}^v (x_v - x_k) z^k \\ &\geq (\rho - 1)(x_v - x_u) z^u + \sum_{k=u}^{v-1} \rho z(x_v - x_k) z^k - \sum_{k=u+1}^v (x_v - x_k) z^k \\ &\geq (\rho - 1)(x_v - x_u) z^u + \sum_{k=u}^v (\rho z - 1)(x_v - x_k) z^k + (x_v - x_u) z^u \\ &\geq (\rho z - 1)K_t(z) + \rho(x_v - x_u) z^u. \end{aligned}$$

The lemma implies that

$$(\rho z - 1)K(z) = \sum_{t \geq 0} (\rho z - 1)K_t(z) \leq P_1(z) \leq \sum_{k=1}^n z^k = \frac{z^{n+1} - z}{z - 1}. \quad (3)$$

⁴ Recall that x_k depends on t . Note also that, among the n agents, rightward motion within $[0, 1]$ might greatly outweigh the leftward kind. Thus, if most of the x_i 's keep growing, how can $P_t(z)$ not follow suit? The point is that $P_t(z)$ puts weights exponentially growing on the right, so their leftward motion, outweighed as it might be, will always dominate with respect to $P_t(z)$. This balancing act between left and right motion is the core principle of twist systems.

10:6 A Connectivity-Sensitive Approach to Consensus Dynamics

With $z = (1 + \varepsilon)/\rho$, $\varepsilon = 1/(n - \nu + 1)$, and $n > 2$, we find that

$$z^{-\nu}K(z) \leq \frac{z^n - 1}{(z - 1)(\rho z - 1)z^{\nu-1}} \leq 2\rho^{\nu-n}e^{(n-\nu+1)\varepsilon}/\varepsilon \leq 2e(n - \nu + 1)\rho^{\nu-n} \leq \frac{3en}{2}\rho^{\nu-n}.$$

The case $s = 1$ of Theorem 2 follows immediately from Lemma 4. Finally, for $n = 2$, we verify that $\mathcal{E}_{m,1}^{\text{TW}} = \sum_{k \geq 0} (1 - 2\rho)^k = 1/2\rho$.

Case $s < 1$

The previous argument relied crucially on the linearity of the 1-energy. If $s < 1$, the s -energy gives more relative weight to small lengths, so we need a different strategy to keep the scales separated. We omit the superscript TW below but it is understood. We use a threshold δ which, though set to $1/3$, is best kept as δ in the notation.

A recurrence relation. Let T_δ be the number of steps at which the diameter remains above $1 - \delta$; note that these steps are consecutive and T_δ might be infinite. By scaling, we find that $E_s \leq F_s + (1 - \delta)^s E_s$, where $F_s = \sum_{t \leq T_\delta} E_{s,t}$. Since $(1 - \delta)^s \leq 1 - \delta s$, for $\delta, s \in [0, 1]$, we have

$$E_s \leq (\delta s)^{-1} F_s. \quad (4)$$

If $m = 1$, then $(1 - \delta)F_s \leq (1 - \delta)T_\delta \leq F_1 \leq \mathcal{E}_{1,1}$. Thus, by (4) and the previous section,

$$\mathcal{E}_{1,s} \leq \frac{3en}{\delta(1 - \delta)} (1/s)(1/\rho)^{\lfloor n/2 \rfloor}. \quad (5)$$

The case $m > 1$. Fix $t \leq T_\delta$; if $m_t > 1$, let j maximize $x_{j+1}(t) - x_j(t)$ over all $j = v_i$ and $i < m_t$ (break ties by taking the smallest j). This corresponds to the maximum distance between consecutive blocks. For this reason, we call $(j, j + 1)$ the *max-gap* at time t . We say that t is *ungapped* if $m_t = 1$ or $x_{j+1}(t) - x_j(t) \leq \delta/m$; it is *gapped* otherwise. Assuming that t is gapped, let $(j, j + 1)$ be its max-gap and write $\zeta_t = \min_k \{ t < k \leq T_\delta \mid \exists l : u_{k,l} \leq j < v_{k,l} \}$: If the set is empty, we set $\zeta_t = T_\delta$; else l is unique and we denote it by l_t . We call the interval $[t, \zeta_t]$ a *span* and the block l_t , if it exists, its *cap*. We note that $x_{j+1}(k) - x_j(k)$ cannot decrease during the times $k = t, \dots, \zeta_t$. This shows that a cap covers a length greater than δ/m .

We begin with a few words of intuition. The energetic contribution of an ungapped time t is easy to account for: It is at most m . On the other hand, the 1-energy is at least the diameter minus the added length of the gaps between blocks, which amounts to at least $1 - 2\delta \geq 1/3$; in other words, $E_{s,t} \leq 3mE_{1,t}$. Summing up over all ungapped times and plugging in our bound for $s = 1$ gives us the desired result. Accounting for gapped times is more difficult, as it requires dealing with small scales. If we had only one span, we could simply split the system into two decoupled subsystems and set up a recurrence relation. The problem is that the presence of k capped spans would force us to repeat the recursion $k - 1$ times. With no apriori bound on k , this approach is not too promising. Instead, we make a bold move: We argue that, because a cap is longer than δ/m , its own 1-energy contribution (ie, its length) is large enough to “pay” for the s -energy of its entire span. This is not quite right, of course, but one can fix the argument by using the weighted 1-energy of the cap and upscaling it suitably. Once again, this reduces the problem to the case $s = 1$, so our method is, in effect, a linearization. Here is the proof.

Proof. We partition the times between 1 and T_δ into two subsets G and $U := [1, T_\delta] \setminus G$, each one supplied with its own energetic accounting scheme. We form G by greedily extracting a maximal set of nonoverlapping spans and taking their union.

1. $G \leftarrow \emptyset$ and $t' \leftarrow 1$;
2. **if** $t \leftarrow \min \{ \text{gapped } i \mid t' \leq i \leq T_\delta \}$ exists
3. **then** $G \leftarrow G \cup \{ i \mid t \leq i \leq \zeta_t \}$;
4. **if** $\zeta_t < T_\delta$ **then** $t' = \zeta_t + 1$; go to 2;

We postulate that, for any $0 < s \leq 1$, and any number of agents $j \leq n$,

$$\mathcal{E}_{m,s} \leq c_m (1/s)^m (1/\rho)^{j-1}, \quad (6)$$

and we derive a recurrence relation for c_m (for given n).

- *Accounting for G :* In line 3, let $k = \zeta_t$ and $(j, j+1)$ be the corresponding max-gap. Suppose that the span $[t, k]$ is capped. The absence of an interval including j and $j+1$ during $[t, k-1]$ implies that $\sum_{t \leq l \leq k} E_{s,l} \leq L + R + m$, where L and R denote the s -energy of systems with at most $m-1$ connected components. For reasons we address below, we may assume that L is dominant; hence $R \leq L \leq \mathcal{E}_{m-1,s}$.⁵ It follows that

$$\sum_{l=t}^k E_{s,l} \leq 2\mathcal{E}_{m-1,s} + m \leq 3c_{m-1} (1/s)^{m-1} (1/\rho)^{j-1}. \quad (7)$$

Note that we (safely) assume $c_{m-1} \geq m$. Using the shorthand v for $v_{k,l}$, we have $x_v(k) - x_j(k) \geq x_{j+1}(k) - x_j(k) \geq x_{j+1}(t) - x_j(t) > \delta/m$. We add the artificial multiplier $x_v(k) - x_j(k)$ to (7) to make the right-hand side resemble $K(z)$. Recall that $\delta = 1/3$; assuming that $z > 1/\rho$ from now on, we have

$$\sum_{l=t}^k E_{s,l} \leq B(x_v(k) - x_j(k))z^j, \quad \text{with } B = 9c_{m-1}m\rho(1/s)^{m-1} \quad (8)$$

The set G is a union of spans. If $\zeta_t = T_\delta$, the last span might not be capped. If so, remove it from G and call the resulting set G' . Summing up, we find that $\sum_{t \in G'} E_{s,t} \leq B \sum_{t \in G'} K_t(z)$. If the last span is uncapped then $\zeta_t = T_\delta$ and no block contains both j and $j+1$ in the span $[t, T_\delta]$. The s -energy expended in that span is thus of the form $L + R \leq 2\mathcal{E}_{m-1,s}$.

- *Accounting for U :* Only ungapped times belong to U , so the 1-energy at time $t \in U$ is at least $1 - \delta - (m_t - 1)\delta/m \geq 1/3$. On the other hand, $E_{s,t} \leq m \leq 3mE_{1,t} \leq 3m\rho K_t(z) \leq BK_t(z)$.

Set $z = (1 + \varepsilon)/\rho$, for $\varepsilon > 0$. Putting all of our bounds together, we have

$$\begin{aligned} F_s &= \sum_{t \leq T_\delta} E_{s,t} \leq B \sum_{t \in G'} K_t(z) + 2\mathcal{E}_{m-1,s} + B \sum_{t \in U} K_t(z) \\ &\leq 2c_{m-1} (1/s)^{m-1} (1/\rho)^{n-1} + BK(z). \end{aligned}$$

⁵ The inequality relies on the (easy) fact that the maximum s -energy grows monotonically with the number of agents. This is not even needed, however, if we redefine $\mathcal{E}_{m,s}$ as the maximum s -energy over all systems with at most n agents and then reason with the value $n' \leq n$ that achieves the maximum.

Actually, the exponent to $1/\rho$ can be reduced to $n - 2$, but this is immaterial. By (3),

$$BK(z) \leq 18c_{m-1}(1/s)^{m-1}(1/\rho)^{n-1}e^{(n+1)\varepsilon}(m/\varepsilon).$$

Setting $\varepsilon = 1/(n + 1)$ gives us, for some constant $d > 0$,

$$F_s \leq c_{m-1}(dmn)(1/s)^{m-1}(1/\rho)^{n-1}.$$

We tie up the loose ends by arguing that it was legitimate to assume that $L \geq R$. The point is that individual values of L and R do not matter: only their sums do. Thus, if the R s outweigh the L s, we restore the dominance of the L s by flipping the system around. Finally, by (4), $E_s \leq (3/s)F_s$; and so, by (5), Theorem 2 follows from the recurrence: $c_1 = O(n)$ and $c_m \leq 3dmnc_{m-1}$, for $m > 1$. ◀

Lower bounds for twist systems

We begin with the case $m = 1$. Assume that $n = 2k + 1$. At time $t = 1$, we have $x_k = -x_{-k} = 1/2$ and $x_i = -x_{-i} = \frac{1}{2}(1 - \rho^i)$, for $0 \leq i < k$ and ρ small enough. For $t > 1$, we set $x_i(t) = (1 - \rho^k)x_i(t - 1)$. The agents are labeled $-k, \dots, k$ from left to right. It is easily verified that this constitutes a twist system for the block $[-k, k]$ with initial unit diameter. The s -energy E is $(1 - \rho^k)^s E + 1$ so, for constant $c > 0$,

$$E \geq (c/s)(1/\rho)^{\lfloor n/2 \rfloor}. \quad (9)$$

If $n = 2k$, we set $x_k = -x_{-k} = 1/2$ and $x_i = -x_{-i} = \frac{1}{2}(1 - 2\rho^i)$, for $1 \leq i < k$. For $t > 1$, we set $x_i(t) = (1 - 2\rho^k)x_i(t - 1)$ and rederive (9).

For the general case, we describe the evolution of an m -block twist system with n agents, and denote its s -energy by $F(n, m)$: It is assumed that $n - 1$ agents are positioned at 0 at time 1 and the last one is at position 1. If $m = 1$, we apply the previous construction after shifting the initial interval from $[-0.5, 0.5]$ to $[0, 1]$. The initial positions still do not match, but we note that, in a single step, we can move the agents anywhere we want in the interval $[\rho, 1 - \rho]$ while respecting the constraints of a twist system. This gives us $F(n, 1) = 1 + (1 - 2\rho)^s E$. By adjusting the constant c in (9), the same lower bound still holds.

For $m > 1$, at time 1, we move the agents $n - 1$ and n to positions ρ and $1 - \rho$, respectively, and we leave the others (if any) at position 0. We then use an $(m - 1)$ -block twist system recursively for the agents $1, \dots, n - 1$. This brings these agents to a common position⁶ in $[0, \rho]$. This gives us the recurrence relation: $F(n, m) \geq 1 + \rho^s F(n - 1, m - 1) + (1 - 2\rho)^s F(n, m)$; hence, by induction, for constant $c > 0$,

$$F(n, m) \geq (c/s)^m \rho^{(m-1)s} (1/\rho)^{\lfloor (n+m-1)/2 \rfloor}. \quad (10)$$

The s -energy is often used for small s , so we state the case of $s = O(1/m \log \frac{1}{\rho})$, which matches the bound of Theorem 2 for $m = n - 1$.

► **Theorem 6.** $\mathcal{E}_{m,s}^{\text{TW}} \geq (c/s)^m (1/\rho)^{\lfloor (n+m-1)/2 \rfloor}$, for constant $c > 0$, small enough ρ and $s = O(1/m \log \frac{1}{\rho})$.

⁶ To keep the time finite, we can always force completion in a single step once the agents are sufficiently close to each other.

2.2 Expanding averaging systems

In a *revS*, the stochastic matrices P_t are of the form $P_t = \text{diag}(q)^{-1}M_t$, where M_t is symmetric with nonzero entries at least 1 and $q = M_t \mathbf{1} \preceq \mathbf{1}/\rho$. We verify that q is the common (dominant) left-eigenvector. We revisit the definition of $\mathcal{E}_{m,s}$ to include only the reversible averaging systems of unit variance $\|x - \hat{x}\|_q^2 = 1$, where $\hat{x} = q\langle x, \mathbf{1} \rangle$. We denote the s -energy by $\mathcal{E}_{m,s}^{\text{REV}}$. The following result is already known. (Note that, if the variance is not one, it suffices to multiply the upper bound by $\|x - \hat{x}\|_q^s$.)

► **Theorem 7** ([15]). $\mathcal{E}_{m,s}^{\text{REV}} \leq (cn^2/\rho s)^m$, for any $s \in (0, 1]$ and constant $c > 0$.

A d -regular expander with Cheeger constant h is a graph of degree d such that, for any set X of at most half the vertices, we have $|\partial X| \geq h|X|$, where ∂X is the set of edges with exactly one vertex in X . We say that $G = (V, E)$ is a d -regular m -expander if it has at most m connected components. Recall that an expanding averaging system (*expS*) is a *revS* consisting of d -regular m -expanders. Each nonzero entry in M_t is equal to 1 and $q = d\mathbf{1}$. We redefine the s -energy to include only *expS* of unit variance with at most m connected components and denote it by $\mathcal{E}_{m,s}^{\text{EXP}}$. Adding the expanding assumptions cancels the dependency on n in the case $m = 1$. More generally, we prove the following:

► **Theorem 8.** $\mathcal{E}_{m,s}^{\text{EXP}} \leq (c/s)^m$, for any $s \in (0, 1]$, where $c = O(d^3/h^2)$ for $m = 1$ and $c = O(d^3mn/h^2)$ for $m > 1$.

We begin the proof with a lower bound on the Dirichlet form that exploits the expansion of a d -regular expander with Cheeger constant h . This is known as *Cheeger's inequality*. We include the proof below for completeness.

► **Lemma 9.** If $G = (V, E)$ is connected, then $\sum_{(i,j) \in E} (x_i - x_j)^2 \geq b(h/d)^2 \|x - \hat{x}\|_q^2 \geq b(h\Delta)^2/2d$, for constant $b > 0$, where Δ is the diameter of the agent positions $x_1, \dots, x_n$.

Proof. All of the ideas in this proof come from [3, 52]. The inequality is invariant under shifting and scaling, so we may assume that $\hat{x} = \mathbf{0}$ and $\|x\|_2 = 1$. Relabel the coordinates of x so they appear in nonincreasing order, and define $y \in \mathbb{R}^n$ such that $y_i = \max\{x_i, 0\}$. Let $\alpha = \arg\max_k (y_k > 0)$ and $\beta = \min\{\alpha, \lfloor n/2 \rfloor\}$. By switching x into $-x$ if necessary,⁷ we can always assume that $\|y\|_2^2 > c := 1/6$ if $\alpha = \beta$, and $\|y\|_2^2 \geq 1 - c$ if $\alpha > \beta$. By Cauchy-Schwarz, $(y_i + y_j)^2 \leq 2(y_i^2 + y_j^2)$; hence,

$$\begin{aligned} \sum_{(i,j) \in E} |y_i^2 - y_j^2| &= \sum_{(i,j) \in E} (y_i + y_j)|y_i - y_j| \leq \sqrt{\sum_{(i,j) \in E} (y_i + y_j)^2} \sqrt{\sum_{(i,j) \in E} (y_i - y_j)^2} \\ &\leq \sqrt{\sum_i 2dy_i^2 \sum_{(i,j) \in E} (y_i - y_j)^2} \leq \sqrt{\sum_{(i,j) \in E} 2d(y_i - y_j)^2} \leq \sqrt{\sum_{(i,j) \in E} 2d(x_i - x_j)^2}. \end{aligned} \quad (11)$$

By the expansion property of G , summation by parts yields

$$\sum_{(i,j) \in E} |y_i^2 - y_j^2| \geq \sum_{k=1}^{\lfloor n/2 \rfloor} hk(y_k^2 - y_{k+1}^2) + \sum_{k=\lfloor n/2 \rfloor + 1}^{n-1} h(n-k)(y_{k+1}^2 - y_k^2) = h(\|y\|_2^2 - ny_{\lfloor n/2 \rfloor + 1}^2). \quad (12)$$

⁷ Intuitively, by changing all signs if necessary, we force the minority sign among the coordinates of x to be positive unless their contribution to the norm of x is too small.

Suppose that $\alpha > \beta = \lfloor n/2 \rfloor$. It follows from $\sum_{i=1}^n x_i = 0$ that $\sum_{i=\alpha+1}^n |x_i| = \|y\|_1 \geq (\beta+1)y_{\beta+1}$. By Cauchy-Schwarz, this yields $ny_{\beta+1}^2/4 \leq \sum_{i=\alpha+1}^n x_i^2 = 1 - \|y\|_2^2$, and by (12), $\sum_{(i,j) \in E} |y_i^2 - y_j^2| \geq (1-5c)h = ch$. If $\alpha = \beta$, then $y_{\lfloor n/2 \rfloor + 1} = 0$ and $\sum_{(i,j) \in E} |y_i^2 - y_j^2| \geq h\|y\|_2^2 > ch$. Applying (11) shows that $\sum_{(i,j) \in E} (x_i - x_j)^2 \geq (ch)^2/2d$, which gives us the first inequality of the lemma. The second one follows from the fact that the interval $[a, b]$ enclosing the vertex positions contains 0. By Cauchy-Schwarz, $1 = \|x\|_2^2 \geq a^2 + b^2 \geq \frac{1}{2}(b+|a|)^2 = \Delta^2/2$, and the proof is complete. $\blacktriangleleft$

Let $G_{\leq t}$ be the graph obtained by adding all the edges from $G_1, \dots, G_t$. Let $m_t \leq m$ be the number of connected components in G_t , and let $\Delta_{t,i}$ denote the diameter of the i -th component of G_t (labeled in any order). Let $t_1, \dots, t_c$ be the times $t > 1$ at which the addition of G_t reduces the number of components in $G_{\leq t-1}$. If no such times exist, set $c = t_c = 1$.

► **Lemma 10.** *If $G_{\leq t_c}$ is connected, then $\sum_{t \leq t_c} \sum_{i=1}^{m_t} \Delta_{t,i}^2 \geq \frac{1}{2dmn} \|x - \hat{x}\|_q^2$.*

Proof. At any time $t_k > 1$, the drop d_k in the number of components can be achieved by d_k (or fewer) components in G_{t_k} . We collect the intervals spanned by these components into a set F , to which we add the intervals for the components of G_1 ; thus $|F| < 2m$. A simple convexity argument (omitted) shows that the union of the intervals in F coincides with the interval $[a, b]$ enclosing the n vertices at time 1; so the lengths $l_1, \dots, l_{|F|}$ of the intervals in F sum up to at least $b-a$. By Cauchy-Schwarz, $\sum_{t \leq t_c} \sum_{i=1}^{m_t} \Delta_{t,i}^2 \geq \sum_{i=1}^{|F|} l_i^2 \geq (b-a)^2/(2m-1)$. The lemma follows from $\|x - \hat{x}\|_q^2 \leq \|q\|_1(b-a)^2 \leq dn(b-a)^2$.

If we define the variant of the Dirichlet form, $D_t = \sum_i \max_{j: (i,j) \in E_t} (x_i(t) - x_j(t))^2$, we know from [15] that, for any $x = x(1) \in \mathbb{R}^n$,

$$\|P_t x\|_q^2 \leq \|x\|_q^2 - \frac{D_t}{2}.$$

It follows that $\|x\|_q^2 - \|x(t_c + 1)\|_q^2 \geq \frac{1}{2} \sum_{t \leq t_c} D_t \geq \frac{1}{d} \sum_{t \leq t_c} \sum_{(i,j) \in E_t} (x_i(t) - x_j(t))^2$. Assuming that $G_{\leq t_c}$ is connected, Lemmas 9 and 10 imply that

$$\|x\|_q^2 - \|x(t_c + 1)\|_q^2 \geq \frac{bh^2}{2d^2} \sum_{t \leq t_c} \sum_{i=1}^{m_t} \Delta_{t,i}^2 \geq \frac{bh^2}{4d^3 mn} \|x - \hat{x}\|_q^2. \quad (13)$$

Let $A(n, m)$ be the maximum s -energy of an $\exp S$ with at most n vertices and m connected components at any time, subject to the initial condition $\|x - \hat{x}\|_q^2 \leq 1$ and, without loss of generality, $\hat{x} = \mathbf{0}$. By (13), $\|x(t)\|_q^2$ shrinks by at least a factor of $\alpha := 1 - bh^2/(4d^3 mn)$ by time $t_c + 1$. By scaling, we see that the s -energy expanded after t_c is at most $\alpha^{s/2} A(n, m)$. While $t < t_c$ (or if $G_{\leq t_c}$ is not connected), the system can be decoupled into two $\exp S$ with fewer than m components. Since $\|x\|_q = 1$, the diameter of the system is at most $2 \max_i |x_i| \leq 2/\sqrt{d}$; therefore $A(n, m) \leq \alpha^{s/2} A(n, m) + 2A(n, m-1) + m(2/\sqrt{d})^s$. It follows that

$$A(n, m) \leq \frac{2}{1 - \alpha^{s/2}} (A(n, m-1) + m). \quad (14)$$

If $m = 1$ then $t_c = 1$, so we can bypass Lemma 10 and its reliance on the diameter. Instead, we use the connectedness of the graphs to derive from Lemma 9:

$$\|x\|_q^2 - \|x(2)\|_q^2 \geq \frac{1}{d} \sum_{(i,j) \in E_1} (x_i - x_j)^2 \geq \frac{bh^2}{d^3} \|x - \hat{x}\|_q^2.$$

Setting $\alpha = 1 - bh^2/d^3$ proves the first part of Theorem 8. The second part follows from (14) and the boundary case $m = 1$ we just derived. $\blacktriangleleft$

2.3 Random averaging systems

It is assumed here that each graph G_t is picked independently, uniformly from the set of simple $(d-1)$ -regular graphs with n vertices, with $d > 3$ [58]. (We use $d-1$ because d must account for the self-loops.) The system is a special case of a *revS*, so we use the same notation. The stochastic matrix P_t for G_t is $\frac{1}{d}M_t$, where M_t is a random symmetric 0/1 matrix with a positive diagonal and all row sums equal to d ; we have $q = d\mathbf{1}$. We define the s -energy $\mathcal{E}_s^{\text{RAN}}$ for systems with unit variance $\|x - \hat{x}\|_q^2 = 1$, where $\hat{x} = \frac{1}{d}\langle x, \mathbf{1} \rangle$.

► **Theorem 11.** $\mathbb{E} \mathcal{E}_s^{\text{RAN}} \leq c/s$, for any $s \in (0, 1]$, where c is an absolute constant.

The term absolute refers to the fact that c is independent of the problem's size and parameters; we assume that d is fixed. Let x_i, y_i be the positions of agent i at step t and $t+1$ respectively. As usual, we may place the center of gravity $\hat{x}$ at the origin at time 1, where it will remain forever; that is, $\sum_{i=1}^n x_i = 0$.

► **Lemma 12.** $\mathbb{E} \sum_{i=1}^n y_i^2 = (1-b) \sum_{i=1}^n x_i^2$, where $b = \frac{(d-1)n}{d(n-1)}$.

Proof. Write $\delta_{ij} = x_i - x_j$ and $M_t = (m_{ij})$. With all sums extending from 1 to n , we have

$$\begin{aligned} \sum_i x_i^2 - \sum_i y_i^2 &= \sum_i x_i^2 - \sum_i \left(x_i - \frac{1}{d} \sum_j m_{ij} \delta_{ij} \right)^2 = \frac{2}{d} \sum_{i,j} m_{ij} x_i \delta_{ij} - \frac{1}{d^2} \sum_{i,j,k} m_{ij} m_{ik} \delta_{ij} \delta_{ik} \\ &= \frac{1}{d} \sum_{i,j} m_{ij} \delta_{ij}^2 - \frac{1}{2d^2} \sum_{i,j,k} m_{ij} m_{ik} (\delta_{ij}^2 + \delta_{ik}^2 - \delta_{jk}^2) = \frac{1}{2d^2} \sum_{i,j,k:i \neq j} m_{ik} m_{jk} \delta_{ij}^2, \end{aligned} \quad (15)$$

with the last equality following from $\sum_{k=1}^n m_{ik} = d$ and $\delta_{ii} = 0$. By symmetry, $\Pr[m_{ij} = 1] = (d-1)/(n-1)$ and $\Pr[m_{ij} m_{ik} = 1] = \binom{d-1}{2} / \binom{n-1}{2}$, for any pairwise distinct i, j, k . For any $i \neq j$, we have $\sum_k m_{ik} m_{jk} = 2m_{ij} + \sum_{k:k \neq i, k \neq j} m_{ik} m_{jk}$; hence

$$\sum_k \mathbb{E}[m_{ik} m_{jk}] = 2\mathbb{E}[m_{ij}] + \sum_{k:k \neq i, k \neq j} \mathbb{E}[m_{ik} m_{jk}] = \frac{d(d-1)}{n-1}.$$

Since $\sum_i x_i = 0$, we have $\sum_{i,j:i \neq j} \delta_{ij}^2 = 2n \sum_i x_i^2$. By (15), it follows that

$$\mathbb{E} \sum_{i=1}^n y_i^2 = \sum_i x_i^2 - \frac{1}{2d^2} \sum_{i,j,k:i \neq j} \delta_{ij}^2 \mathbb{E}[m_{ik} m_{jk}] = \sum_i x_i^2 - \frac{d-1}{2d(n-1)} \sum_{i,j:i \neq j} \delta_{ij}^2.$$

Markov's inequality tells us that $\sum_i y_i^2 \geq (1-b/3) \sum_i x_i^2$ holds with probability at most $\mathbb{E}[\sum_i y_i^2] / [(1-b/3) \sum_i x_i^2] \leq 1-b/2$. Since $\|x\|_q = 1$, the diameter of the system is at most $2/\sqrt{d}$; by the usual scaling law, it follows that

$$\mathbb{E} \mathcal{E}_s^{\text{RAN}} \leq 2^s \mathbb{E}K + \frac{b}{2} (1-b/3)^{s/2} \mathbb{E} \mathcal{E}_s^{\text{RAN}} + (1-b/2) \mathbb{E} \mathcal{E}_s^{\text{RAN}},$$

where K is the number of connected components in G_1 . It is known [58] that, for $d > 3$, the probability that the graph is not connected is $O(n^{3-d})$; hence $\mathbb{E}K = O(1)$. Since $b \geq 1/2$, we conclude that $\mathbb{E} \mathcal{E}_s^{\text{RAN}} = O(1/(1-(5/6)^{s/2})) = O(1/s)$; hence Theorem 11. ◀

3 The Overton Window Attractor

Following in a long line of opinion dynamics models [9, 20, 21, 28, 31], we consider a collection of n agents, each one holding an opinion vector $x_i(t) \in [0, 1]^d$ at time t ; we denote by $x(t)$ the n -by- d matrix whose i -th row corresponds to $x_i(t)$. Given a stochastic matrix P_t , the agents update their opinion vectors at time $t > 0$ according to the evolution equation $x(t+1) = P_t x(t)$. We assume that the last k agents $n-k+1, \dots, n$ are *fixed* in the sense that $x_i(t)$ remains constant at all times $t > 0$. Algebraically, the square block of P_t corresponding to the k fixed agents is set to the identity matrix $\mathbb{I}_k$. The fixed agents can influence the mobile ones, but not the other way around. The presence of fixed agents (also called “stubborn,” “forceful” or “zealots” in the literature) has been extensively studied [2, 1, 30, 43, 44, 53, 60, 59].

In the context of social networks, the fixed sources may consist of venues with low user influence, such as news outlets, wiki pages, influencers, TV channels, political campaign sites, etc. [16, 29, 35, 38, 49, 56, 61]. We show how the mobile agents migrate to the convex hull of the fixed agents; crucially, we bound the rate of attraction. This provides both a quantitative illustration of the famous *Overton window* phenomenon as well as a theoretical explanation for why the window acts as an attracting manifold [4, 6, 19, 27, 46]. Interestingly, the emergence of a global attractor does not imply convergence (ie, fixed-point attraction). The mobile agents might still fluctuate widely in perpetuity. The point is that they will always do so within the confines of the global attractor.

To reflect the stochasticity inherent in the choice of sources visited by a user on a given day, we adopt a classic “planted” model: Fix a connected n -vertex graph G and two parameters $p \in (0, 1]$ and $\rho \in (0, 1/2]$. At each time $t > 0$, G_t is defined by picking every edge of G with probability at least p . (No independence is required and n self-loops are included.) We define an n -by- n stochastic matrix P_t by setting every entry to 0 and updating it as follows:

1. For $i > n - k$, $(P_t)_{ii} = 1$.
2. For $i \leq n - k$, set $(P_t)_{ij} \geq \rho$ for any j such that (i, j) is an edge of G_t .

Note that the update is highly nondeterministic. The only two conditions required are that (i) nonzero entries be at least ρ and (ii) each row sum up to 1.

► **Theorem 13.** *For any $\delta, \varepsilon > 0$, with probability at least $1 - \delta$, all of the agents fall within distance ε of the convex hull of the fixed agents after a number of steps at most*

$$\frac{1}{p\delta} \left(\frac{c}{\rho} \log \frac{dn}{\varepsilon} \right)^{2(n-1)}$$

for constant $c > 0$.

Proof. Let Q_t be the h -by- h upper-left submatrix of P_t , where $h = n - k$. Note that $Q_{\leq t} := Q_t \cdots Q_1$ coincides with the h -by- h upper-left submatrix $P_{\leq t}$. Thus, to show that the mobile agents are attracted to the convex hull of the fixed ones, it suffices to prove that $Q_{\leq t}$ tends to $\mathbf{0}_{h \times h}$. To do that, we create an agreement system consisting of $h + 1$ agents embedded in $[0, 1]$ and evolving as $y(t+1) = A_t y(t)$, where: $y(t) \in \mathbb{R}^{h+1}$; $y_{h+1}(1) = 0$; $v = (\mathbb{I}_h - Q_t)\mathbf{1}_h$; and

$$A_t = \begin{pmatrix} Q_t & v \\ \mathbf{0}_h^T & 1 \end{pmatrix}.$$

The system lacks the requisite zero-symmetry to qualify as a *genS*, so we use symmetrization [12] by duplicating the h mobile agents and initializing the embedding of the two copies as mirror-image reflections about the origin. The new evolution matrix is now ν -by- ν , where $\nu = 2h + 1$:

$$B_t = \begin{pmatrix} Q_t & v & \mathbf{0}_{h \times h} \\ u & 1 - 2\|u\|_1 & u \\ \mathbf{0}_{h \times h} & v & Q_t \end{pmatrix}.$$

We define the row vector $u \in \mathbb{R}^h$ by setting its i -th coordinate to ρ if $v_i > 0$ and 0 otherwise. We require that $1 - 2\|u\|_1 \geq \rho$; hence $\rho \leq 1/(2d_t + 1)$, where d_t is the number of mobile agents (among the h of them) adjacent in G_t to at least one fixed agent. This condition is easily satisfied by setting $\rho \leq 1/2n$. The evolution follows the update: $z(t+1) = B_t z(t)$, where $z(t) \in [-1, 1]^\nu$ and $z_{h+1}(1) = 0$.

Let G^* be the augmented ν -vertex graph formed from G and let G_t^* be its subgraph selected at time t . Note that, via $z(t)$, these graphs are embedded in $[-1, 1]^\nu$. If Δ_t denotes the length of the longest edge of G^* at time t and T_α is the last time at which the diameter of the system is at least α , then $\Delta_t \geq \alpha/\nu$ for all $t \leq T_\alpha$ because G^* is connected. The longest edge in G^* (with ties broken alphabetically) appears in G_t^* with probability at least p . Fix $s \in (0, 1]$ and define the random variable χ_t to be Δ_t^s if the longest edge of G at time t is in G_t and 0 otherwise. By [13], the maximum s -energy satisfies $\mathcal{E}_s \leq 2^s(3/\rho s)^{\nu-1}$; hence

$$\mathbb{E} T_\alpha \leq \left(\frac{\nu}{\alpha}\right)^s \mathbb{E} \sum_{t \geq 0} \Delta_t^s \leq \frac{1}{p} \left(\frac{\nu}{\alpha}\right)^s \mathbb{E} \sum_{t \geq 0} \chi_t \leq \frac{1}{p} \left(\frac{\nu}{\alpha}\right)^s \mathcal{E}_s \leq \frac{2}{p} \left(\frac{\nu}{\alpha}\right)^s \left(\frac{3}{\rho s}\right)^{\nu-1}.$$

Minimizing the right-hand side over all $s \in (0, 1]$ yields

$$\mathbb{E} T_\alpha \leq \frac{4}{p} \left(\frac{3}{\rho} \log \frac{2n-1}{\alpha}\right)^{2(n-1)}.$$

By Markov's inequality, $\Pr[T_\alpha \geq t_\delta] \leq \delta$, where

$$t_\delta := \frac{4}{p\delta} \left(\frac{3}{\rho} \log \frac{2n-1}{\alpha}\right)^{2(n-1)}. \quad (16)$$

This implies that $\|Q_{\leq t} \mathbf{1}_h\|_\infty \leq \alpha$, for all $t > t_\delta$, with probability at least $1 - \delta$. In other words, for any such t , it holds that, for $i \leq h$,

$$q := \sum_{j=1}^h (P_{\leq t})_{ij} = \sum_{j=1}^h (Q_{\leq t})_{ij} \leq \alpha.$$

Trivially, $x_i(t+1) = qu + (1-q)v$, where

$$u = \frac{1}{q} \sum_{j=1}^h (P_{\leq t})_{ij} x_j(1) \quad \text{and} \quad v = \frac{1}{1-q} \sum_{j=h+1}^n (P_{\leq t})_{ij} x_j(1).$$

Observing that v lies in the convex hull of the fixed agents, we form the difference $x_i(t+1) - v = q(u - v)$ and note that the distance from $x_i(t+1)$ to the hull is bounded by $q\|u - v\|_2 \leq \alpha\sqrt{d}$. Setting $\alpha = \varepsilon/\sqrt{d}$ completes the proof. $\blacktriangleleft$

We can extend this result so as to relate convergence to connectivity. We now produce the random graph G_t from fixed connected G as we did above, but if this results in a graph with more than m connected components, we add random edges picked uniformly from G until the number of components drops to m . Using Theorem 2 in the proof above leads to a more refined bound:

► **Theorem 14.** For any $\delta, \varepsilon > 0$, with probability at least $1 - \delta$, all of the agents fall within distance ε of the convex hull of the fixed agents in time bounded by

$$\frac{1}{p\delta} \left(\frac{1}{\rho} \right)^{2(n-1)} \left(cmn \log \frac{dn}{\varepsilon} \right)^{2m-1},$$

for constant $c > 0$. This assumes that no graph used in the process has more than m connected components.

Proof. By Theorem 2, we know that $\mathcal{E}_{m,s}^{\text{GEN}} \leq (b/s)^m (1/\rho)^{n-1}$, for any $s \in (0, 1]$, where $b = O(mn)$. The previous proof leads us to update (16) into:

$$t_\delta := \frac{1}{p\delta} \left(\frac{1}{\rho} \right)^{2(n-1)} \left(cmn \log \frac{n}{\alpha} \right)^{2m-1},$$

for constant $c > 0$, from which the theorem follows. ◀

References

- 1 Olle Abrahamsson, Danyo Danev, and Erik G. Larsson. Opinion dynamics with random actions and a stubborn agent. In Michael B. Matthews, editor, *53rd Asilomar Conference on Signals, Systems, and Computers, ACSCC 2019, Pacific Grove, CA, USA, November 3-6, 2019*, pages 1486–1490. IEEE, 2019. doi:10.1109/IEEECONF44664.2019.9048901.
- 2 Daron Acemoglu, Asuman E. Ozdaglar, and Ali ParandehGheibi. Spread of (mis)information in social networks. *Games Econ. Behav.*, 70(2):194–227, 2010. doi:10.1016/j.geb.2010.01.005.
- 3 Noga Alon. Eigenvalues and expanders. *Comb.*, 6(2):83–96, 1986. doi:10.1007/BF02579166.
- 4 Emmi Bevensee and Alexander Reid Ross. The alt-right and global information warfare. In Naoki Abe, Huan Liu, Calton Pu, Xiaohua Hu, Nesreen K. Ahmed, Mu Qiao, Yang Song, Donald Kossmann, Bing Liu, Kisung Lee, Jiliang Tang, Jingrui He, and Jeffrey S. Saltz, editors, *IEEE International Conference on Big Data (IEEE BigData 2018), Seattle, WA, USA, December 10-13, 2018*, pages 4393–4402. IEEE, 2018. doi:10.1109/BigData.2018.8622270.
- 5 Vincent D. Blondel, Julien M. Hendrickx, and John N. Tsitsiklis. On krause’s multi-agent consensus model with state-dependent connectivity. *IEEE Trans. Autom. Control.*, 54(11):2586–2597, 2009. doi:10.1109/TAC.2009.2031211.
- 6 George-Daniel Bobric. The overton window: A tool for information warfare. In *ICCWS 2021 16th International Conference on Cyber Warfare and Security*, pages 20–27. Academic Conferences Limited, 2021.
- 7 Francesco Bullo, Jorge Cortés, and Sonia Martinez. *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*, volume 27. Princeton University Press, 2009.
- 8 Ming Cao, A. Stephen Morse, and Brian D. O. Anderson. Reaching a consensus in a dynamically changing environment: Convergence rates, measurement delays, and asynchronous events. *SIAM J. Control. Optim.*, 47(2):601–623, 2008. doi:10.1137/060657029.
- 9 Claudio Castellano, Santo Fortunato, and Vittorio Loreto. Statistical physics of social dynamics. *Reviews of modern physics*, 81(2):591–646, 2009.
- 10 Bernadette Charron-Bost and Patrick Lambein-Monette. Computing outside the box: Average consensus over dynamic networks. In James Aspnes and Othon Michail, editors, *1st Symposium on Algorithmic Foundations of Dynamic Networks, SAND 2022, March 28-30, 2022, Virtual Conference*, volume 221 of *LIPIcs*, pages 10:1–10:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.SAND.2022.10.
- 11 Samprit Chatterjee and Eugene Seneta. Towards consensus: Some convergence theorems on repeated averaging. *Journal of Applied Probability*, 14(1):89–97, 1977.
- 12 Bernard Chazelle. The total s-energy of a multiagent system. *SIAM J. Control. Optim.*, 49(4):1680–1706, 2011. doi:10.1137/100791671.

- 13 Bernard Chazelle. A sharp bound on the s-energy and its applications to averaging systems. *IEEE Trans. Autom. Control.*, 64(10):4385–4390, 2019. doi:10.1109/TAC.2019.2899509.
- 14 Bernard Chazelle. On the periodicity of random walks in dynamic networks. *IEEE Trans. Netw. Sci. Eng.*, 7(3):1337–1343, 2020. doi:10.1109/TNSE.2019.2924921.
- 15 Bernard Chazelle and Kritkorn Karntikoon. Quick relaxation in collective motion. In *61st IEEE Conference on Decision and Control, CDC 2022, Cancun, Mexico, December 6-9, 2022*, pages 6472–6477. IEEE, 2022. doi:10.1109/CDC51059.2022.9992475.
- 16 Uthsav Chitra and Christopher Musco. Analyzing the impact of filter bubbles on social network polarization. In James Caverlee, Xia (Ben) Hu, Mounia Lalmas, and Wei Wang, editors, *WSDM '20: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, February 3-7, 2020*, pages 115–123. ACM, 2020. doi:10.1145/3336191.3371825.
- 17 Étienne Coulouma, Emmanuel Godard, and Joseph G. Peters. A characterization of oblivious message adversaries for which consensus is solvable. *Theor. Comput. Sci.*, 584:80–90, 2015. doi:10.1016/j.tcs.2015.01.024.
- 18 Felipe Cucker and Steve Smale. Emergent behavior in flocks. *IEEE Trans. Autom. Control.*, 52(5):852–862, 2007. doi:10.1109/TAC.2007.895842.
- 19 CIH David Dyjack DrPH. The overton window. *Journal of Environmental Health*, 82(7):54–53, 2020.
- 20 Guillaume Deffuant, David Neau, Frédéric Amblard, and Gérard Weisbuch. Mixing beliefs among interacting agents. *Adv. Complex Syst.*, 3(1-4):87–98, 2000. doi:10.1142/S0219525900000078.
- 21 Morris H DeGroot. Reaching a consensus. *Journal of the American Statistical association*, 69(345):118–121, 1974.
- 22 Matthew G Earl and Steven H Strogatz. Synchronization in oscillator networks with delayed coupling: A stability criterion. *Physical Review E*, 67(3):036204, 2003.
- 23 Antoine El-Hayek, Monika Henzinger, and Stefan Schmid. Brief announcement: Broadcasting time in dynamic rooted trees is linear. In Alessia Milani and Philipp Woelfel, editors, *PODC '22: ACM Symposium on Principles of Distributed Computing, Salerno, Italy, July 25–29, 2022*, pages 54–56. ACM, 2022. doi:10.1145/3519270.3538460.
- 24 Antoine El-Hayek, Monika Henzinger, and Stefan Schmid. Asymptotically tight bounds on the time complexity of broadcast and its variants in dynamic networks. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*, volume 251 of *LIPIcs*, pages 47:1–47:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.ITCS.2023.47.
- 25 Antoine El-Hayek, Monika Henzinger, and Stefan Schmid. Time complexity of broadcast and consensus for randomized oblivious message adversaries. *CoRR*, abs/2302.11988, 2023. doi:10.48550/arXiv.2302.11988.
- 26 Fabio Fagnani and Paolo Frasca. *Introduction to averaging dynamics over networks*. Springer, 2018.
- 27 Denis N Fedyanin. On control for reaching a consensus in multiagent systems with bounded opinion updates. In *2021 14th International Conference Management of large-scale system development (MLSD)*, pages 1–5. IEEE, 2021.
- 28 Noah E Friedkin and Eugene C Johnsen. Social influence and opinions. *Journal of Mathematical Sociology*, 15(3-4):193–206, 1990.
- 29 Jason Gaitonde, Jon M. Kleinberg, and Éva Tardos. Polarization in geometric opinion dynamics. In Péter Biró, Shuchi Chawla, and Federico Echenique, editors, *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, pages 499–519. ACM, 2021. doi:10.1145/3465456.3467633.
- 30 Javad Ghaderi and R. Srikant. Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate. *Autom.*, 50(12):3209–3215, 2014. doi:10.1016/j.automatica.2014.10.034.

- 31 Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence: models, analysis and simulation. *J. Artif. Soc. Soc. Simul.*, 5(3), 2002. URL: <http://jasss.soc.surrey.ac.uk/5/3/2.html>.
- 32 Julien Hendrickx and Vincent Blondel. Convergence of different linear and non-linear vicsek models. In *MTNS 2006*, pages 1229–1240, 2006.
- 33 Julien M. Hendrickx and John N. Tsitsiklis. Convergence of type-symmetric and cut-balanced consensus seeking systems. *IEEE Trans. Autom. Control.*, 58(1):214–218, 2013. doi:10.1109/TAC.2012.2203214.
- 34 Ali Jadbabaie, Jie Lin, and A. Stephen Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Trans. Autom. Control.*, 48(6):988–1001, 2003. doi:10.1109/TAC.2003.812781.
- 35 Jing Jiang, Christo Wilson, Xiao Wang, Wenpeng Sha, Peng Huang, Yafei Dai, and Ben Y. Zhao. Understanding latent interactions in online social networks. *ACM Trans. Web*, 7(4):18:1–18:39, 2013. doi:10.1145/2517040.
- 36 Fabian Kuhn, Nancy A. Lynch, and Rotem Oshman. Distributed computation in dynamic networks. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 513–522. ACM, 2010. doi:10.1145/1806689.1806760.
- 37 Fadel Lashhab. *Building Efficiency Control; Dynamic Consensus Networks*. LAP Lambert Academic Publishing, 2018.
- 38 Kevin Lewis, Marco Gonzalez, and Jason Kaufman. Social selection and peer influence in an online social network. *Proceedings of the National Academy of Sciences*, 109(1):68–72, 2012.
- 39 Jan Lorenz. A stabilization theorem for dynamics of continuous opinions. *Physica A: Statistical Mechanics and its Applications*, 355(1):217–223, 2005.
- 40 Jan Lorenz. Heterogeneous bounds of confidence: Meet, discuss and find consensus! *Complex.*, 15(4):43–52, 2010. doi:10.1002/cplx.20295.
- 41 Anahita Mirtabatabaei and Francesco Bullo. Opinion dynamics in heterogeneous networks: Convergence conjectures and theorems. *SIAM J. Control. Optim.*, 50(5):2763–2785, 2012. doi:10.1137/11082751X.
- 42 Ciamac Cyrus Moallemi and Benjamin Van Roy. Consensus propagation. *IEEE Transactions on Information Theory*, 52(11):4753–4766, 2006.
- 43 Mauro Mobilia. Does a single zealot affect an infinite group of voters? *Physical review letters*, 91(2):028701, 2003.
- 44 Mauro Mobilia, Anna Petersen, and Sidney Redner. On the role of zealotry in the voter model. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(08):P08029, 2007.
- 45 Luc Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Trans. Autom. Control.*, 50(2):169–182, 2005. doi:10.1109/TAC.2004.841888.
- 46 Daniel J Morgan. The overton window and a less dogmatic approach to antibiotics. *Clinical Infectious Diseases*, 70(11):2439–2441, 2020.
- 47 Angelia Nedic, Alex Olshevsky, Asuman E. Ozdaglar, and John N. Tsitsiklis. On distributed averaging algorithms and quantization effects. *IEEE Trans. Autom. Control.*, 54(11):2506–2517, 2009. doi:10.1109/TAC.2009.2031203.
- 48 Angelia Nedic. Convergence rate of distributed averaging dynamics and optimization in networks. *Found. Trends Syst. Control.*, 2(1):1–100, 2015. doi:10.1561/26000000004.
- 49 Tien T. Nguyen, Pik-Mai Hui, F. Maxwell Harper, Loren G. Terveen, and Joseph A. Konstan. Exploring the filter bubble: the effect of using recommender systems on content diversity. In Chin-Wan Chung, Andrei Z. Broder, Kyuseok Shim, and Torsten Suel, editors, *23rd International World Wide Web Conference, WWW '14, Seoul, Republic of Korea, April 7-11, 2014*, pages 677–686. ACM, 2014. doi:10.1145/2566486.2568012.
- 50 Alex Olshevsky and John N. Tsitsiklis. Convergence speed in distributed consensus and averaging. *SIAM J. Control. Optim.*, 48(1):33–55, 2009. doi:10.1137/060678324.

- 51 Alex Olshevsky and John N. Tsitsiklis. Degree fluctuations and the convergence time of consensus algorithms. *IEEE Trans. Autom. Control.*, 58(10):2626–2631, 2013. doi:10.1109/TAC.2013.2257969.
- 52 Alistair Sinclair. *Algorithms for random generation and counting – A Markov chain approach*. Progress in theoretical computer science. Birkhäuser, 1993.
- 53 Ye Tian and Long Wang. Opinion dynamics in social networks with stubborn agents: An issue-based perspective. *Autom.*, 96:213–223, 2018. doi:10.1016/j.automatica.2018.06.041.
- 54 John Tsitsiklis, Dimitri Bertsekas, and Michael Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE transactions on automatic control*, 31(9):803–812, 1986.
- 55 Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. Novel type of phase transition in a system of self-driven particles. *Physical review letters*, 75(6):1226–1229, 1995.
- 56 Bimal Viswanath, Alan Mislove, Meeyoung Cha, and P. Krishna Gummadi. On the evolution of user interaction in facebook. In Jon Crowcroft and Balachander Krishnamurthy, editors, *Proceedings of the 2nd ACM Workshop on Online Social Networks, WOSN 2009, Barcelona, Spain, August 17, 2009*, pages 37–42. ACM, 2009. doi:10.1145/1592665.1592675.
- 57 Kyrill Winkler, Ami Paz, Hugo Rincon Galeana, Stefan Schmid, and Ulrich Schmid. The time complexity of consensus under oblivious message adversaries. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*, volume 251 of *LIPIcs*, pages 100:1–100:28. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.ITCS.2023.100.
- 58 Nicholas C Wormald et al. Models of random regular graphs. *London Mathematical Society Lecture Note Series*, pages 239–298, 1999.
- 59 Ercan Yildiz, Daron Acemoglu, Asuman E Ozdaglar, Amin Saberi, and Anna Scaglione. Discrete opinion dynamics with stubborn agents. *Available at SSRN 1744113*, 2011.
- 60 Mehmet Ercan Yildiz, Asuman E. Ozdaglar, Daron Acemoglu, Amin Saberi, and Anna Scaglione. Binary opinion dynamics with stubborn agents. *ACM Trans. Economics and Comput.*, 1(4):19:1–19:30, 2013. doi:10.1145/2538508.
- 61 Reza Zafarani, Mohammad Ali Abbasi, and Huan Liu. *Social Media Mining: An Introduction*. Cambridge University Press, 2014. doi:10.1017/CB09781139088510.