

On the Periodicity of Random Walks in Dynamic Networks

Bernard Chazelle 

Abstract—We investigate random walks in graphs whose edges change over time as a function of the current probability distribution of the walk. We show that such systems can be chaotic and can exhibit “hyper-torpid” mixing. Our main result is that, if each graph is strongly connected, then the dynamics is asymptotically periodic almost surely.

Index Terms—Random walks, temporal networks, Markov influence systems, hyper-torpid mixing, chaos.

I. INTRODUCTION

THERE is a growing body of literature on dynamic networks [1], [3], [7], [10], [17]–[19], [22], [24], [25], [29], [30] and random walks in such structures (also called time-varying or dynamic graphs/networks) [2], [12]–[15], [19], [20], [24], [26]–[28], [31], [33]. Starting from a random node in a graph g_1 , the walk moves to a random neighbor in g_1 , then a random neighbor in some graph g_2 , and iterates in this fashion through g_3, g_4 , etc. All the graphs are directed and share the same set of vertices. The walk is called *temporal* or *time-respecting* because it must traverse an edge of graph g_t at time t . In this work, the sequence of graphs is not given in advance but, rather, is specified endogenously by a dynamical system: thus, g_t is chosen from a finite set of graphs as a function of the current probability distribution on the vertices at time $t - 1$. For applications, one can think of a transportation network where roads are opened or closed depending on the current traffic at the intersections.

We formalize the model within the framework of *Markov influence systems* [11]. Under mild conditions, classical random walks always mix to their stationary distribution in time at most exponential in the size of the network. Markov influence systems display a far richer range of dynamics: they can be periodic, chaotic, or exhibit hyper-torpid mixing. Our main result is that, if every graph g_t is strongly connected, then a lazy random walk is almost surely (asymptotically) periodic. Informally, this means that, under a small random perturbation, with

probability one, there is a finite number of stable periodic orbits to which all orbits are attracted.

Markov influence systems are piecewise-linear dynamical systems with nonpositive Lyapunov exponents. Such systems are notoriously tricky to analyze and sometimes violate basic intuitions [4]–[6], [8], [9]: for example, piecewise isometries have zero topological entropy [6] whereas, remarkably, piecewise contractions can be chaotic [23]. Although deterministic, Markov influence systems are akin to certain spin systems from condensed matter physics in that their dynamics is driven by a tension between *order* (caused by the paracontractivity of the maps) and *disorder* (caused by their discontinuities). This work shows that, at least in strongly connected case, order almost always prevails; specifically, the disordered regime can be “covered” by a Cantor set of nonfull Hausdorff dimension in perturbation space.

II. MARKOV INFLUENCE SYSTEMS

Let \mathbb{S}^{n-1} (or \mathbb{S} when the dimension is understood) be the standard simplex $\{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{x} \geq \mathbf{0}, \|\mathbf{x}\|_1 = 1\}$ and let \mathcal{S} denote set of all n -by- n rational stochastic matrices with positive diagonals. A *Markov influence system* (*MIS*) is a discrete-time dynamical system with phase space \mathbb{S} , which is defined by the map $f: \mathbf{x} \mapsto f(\mathbf{x}) := \mathbf{x}^\top S(\mathbf{x})$, where $\mathbf{x} \in \mathbb{S}$ and S is a function $\mathbb{S} \rightarrow \mathcal{S}$ that is constant over the pieces of a finite polyhedral partition $\mathcal{P} = \{P_k\}$ of \mathbb{S} (fig. 1); we define f as the identity over the discontinuities of the partition. Algebraic discontinuities can be allowed as well; in fact, the map S can be defined by any sentence in the first-order theory of the reals [11]. For simplicity, we restrict our discussion to the case of linear discontinuities.

For fixed \mathbf{x} , the matrix $S(\mathbf{x})$ defines a lazy random walk. Thus, an *MIS* can be interpreted as a lazy temporal random walk with transition probabilities defined endogenously. The *MIS* is called *irreducible* if the graph defined by each $S(\mathbf{x})$ is strongly connected; this means that $S(\mathbf{x})$ is the stochastic matrix of an aperiodic, irreducible (hence ergodic) random walk. The *orbit* of $\mathbf{x} \in \mathbb{S}$ is the infinite sequence $(f^t(\mathbf{x}))_{t \geq 0}$ and its *itinerary* is the corresponding sequence of P_k 's visited in the process. The orbit is *periodic* if $f^t(\mathbf{x}) = f^s(\mathbf{x})$ for any $s = t$ modulo a fixed integer. It is asymptotically periodic if it gets arbitrarily close to a periodic orbit over time. The discontinuities in \mathcal{P} are formed by hyperplanes in \mathbb{R}^n of the form $\mathbf{a}_i^\top \mathbf{x} = 1 + \delta$, where $\mathbf{a}_i \in \mathbb{Q}^n$ and $\delta \in \Omega := [-\omega, \omega]$. Assuming general position, we can pick a small positive $\omega < 1/2$ so that \mathcal{P} remains (topologically)

Manuscript received March 28, 2019; revised May 20, 2019; accepted June 18, 2019. Date of publication June 27, 2019; date of current version September 2, 2020. This work was supported in part by the Army Research Office, in part by the Defense Advanced Research Projects Agency, and in part by Grant W911NF-17-1-0078. Recommended for acceptance by X. Cao.

The author is with the Department of Computer Science, Princeton University, Princeton, NJ 08540, USA (e-mail: chazelle@cs.princeton.edu).

Digital Object Identifier 10.1109/TNSE.2019.2924921

2327-4697 © 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See <https://www.ieee.org/publications/rights/index.html> for more information.




Fig. 1. A Markov influence system: Each polytope P_k is associated with a stochastic matrix. The first five steps of an orbit are shown to visit P_1, P_2, P_4, P_5, P_3 in this order.

invariant over Ω ; in this way, the MIS remains well-defined for all $\delta \in \Omega$.¹ We are now in a position to state our main new result:

Theorem 2.1. Typically, every orbit of an irreducible Markov influence system is asymptotically periodic.

Put more formally, we prove the existence of a set Λ of one-dimensional Lebesgue measure zero such that, for any $\delta \in \Omega \setminus \Lambda$, there is a finite set of stable periodic orbits such that every orbit is asymptotically attracted to one of them. The set Λ can be covered by a Cantor set of Hausdorff dimension less than one.

This article also describes Markov influence systems that are chaotic and others that robustly mix in time equal to a tower-of-twos in the dimension n . One lesson to draw from this work is that Markov influence systems can behave very differently from standard random walks. For example, here is an instance of two random walks that, individually, mix very quickly, but, once assembled as an irreducible MIS, exhibit an arbitrarily long mixing time:

$$S(\mathbf{x}) = \frac{1}{4} \begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix} \text{ if } 2x > 1 + a \text{ and } S(\mathbf{x}) = \frac{1}{3} \begin{pmatrix} 1 & 2 \\ 1 & 2 \end{pmatrix} \text{ else.}$$

Starting from $\mathbf{x} = (1, 0)$, ie, $x = 1$, the temporal random walk uses the first matrix until $x(t) \leq (1 + a)/2$, at which point it switches to the other matrix, which reaches its stationary distribution instantly. Prior to that point, we have $x(t) = 1/2 + 2^{-t-1}$ so the mixing time is about $\log(1/a)$, which can be arbitrarily large.

A. Proving Theorem 2.1

We build a more general framework from which Theorem 2.1 emerges as a special case. Recall that the *coefficient of ergodicity* $\tau(M)$ of a matrix M is defined as half the maximum ℓ_1 -distance between any two of its rows or, equivalently, as $1 - \min_{i,j} \sum_k \min\{M_{ik}, M_{jk}\}$ [32]. It is submultiplicative for stochastic matrices, which is a direct consequence of the identity $\tau(M) = \max\{\|\mathbf{x}^\top M\|_1 : \mathbf{x}^\top \mathbf{1} = 0 \text{ and } \|\mathbf{x}\|_1 = 1\}$. Given

¹Let H_δ be the set consisting of the hyperplane $\sum_i x_i = 1$ together with those used to define \mathcal{P} . We assume that H_0 is in general position; hence so is H_δ for any $\delta \in \Omega$, where $0 < \omega < 1/2$ is smaller than a value that depends only on the set of vectors $\{\mathbf{a}_i\}$. Note that this problem arises only because of the constraint $\sum_i x_i = 1$, since otherwise the hyperplane arrangement is central around $(0, \dots, 0, -1)$.

$\Delta \subseteq \Omega$, let L_Δ^t denote the set of t -long prefixes of any itinerary for any starting position $\mathbf{x} \in \mathbb{S}$ and any $\delta \in \Delta$. We define the *ergodic scale* η as the smallest integer such that, for any $t \geq \eta$ and any matrix sequence S_1, \dots, S_t associated with an element of L_Δ^t , the product $S_1 \cdots S_t$ is primitive (ie, some high enough power is a positive matrix) and its coefficient of ergodicity is less than $1/2$.

Let D be the union of the hyperplanes from \mathcal{P} in \mathbb{R}^n (where δ is understood). We define $Z_t = \bigcup_{0 \leq k \leq t} f^{-k}(D)$ and $Z = \bigcup_{t \geq 0} Z_t$. Since the domain of f is \mathbb{S} , we have $Z \subseteq \mathbb{S}$ by compactness. Remarkably, for almost all $\delta \in \Omega$, Z_t becomes strictly equal to Z in a finite number of steps. (Note that both Z and Z_t depend on $\delta \in \Omega$.)

Lemma 2.2. Assume that the ergodic scale η is finite. There is a constant $c > 0$ such that, for any $\varepsilon > 0$, there exist an integer $v \leq 2^{\eta c} \log(1/\varepsilon)$ and a finite union K of intervals of total length less than ε such that $Z_v = Z_{v-1}$, for any $\delta \in \Omega \setminus K$.

Note that the lemma does not assume irreducibility but only the finiteness of η . All the constants used in this work may depend on the system's parameters such as n, \mathcal{P} (but not on δ). Note that the lemma does not assume irreducibility but only the finiteness of η . Dependency on other parameters is indicated by a subscript. Note that $Z_v = Z_{v-1}$ implies that $Z = Z_v$. Indeed, suppose that $Z_{t+1} \supset Z_t$ for $t \geq v$; then, $f^{t+1}(\mathbf{y}) \in D$ but $f^t(\mathbf{y}) \notin D$ for some $\mathbf{y} \in \mathbb{S}$; in other words, $f^v(\mathbf{x}) \in D$ but $f^{v-1}(\mathbf{x}) \notin D$ for $\mathbf{x} = f^{t-v+1}(\mathbf{y})$, which contradicts the equality $Z_v = Z_{v-1}$.

Corollary 2.3. For δ almost everywhere in Ω , every orbit is asymptotically periodic.

Proof. The polytopes, called *cells*, defined by the connected components of the complement of $Z = Z_v$ form the continuity pieces of f^{v+1} : by continuity, each one of them maps, via f , not simply to within a single cell of D but actually to within a single cell of Z .² This in turn implies the eventual periodicity of the symbolic dynamics. The period cannot exceed the number of cells. Once an itinerary becomes periodic at time t_0 with period σ , the map f^t can be expressed locally by matrix powers. Indeed, divide $t - t_0$ by σ and let q be the quotient and r the remainder; then, locally, $f^t = g^q \circ f^{t_0+r}$, where g is specified by a stochastic matrix with a positive diagonal, which implies convergence to a periodic point at an exponential rate. In the case of an irreducible MIS, the matrix specifying g corresponds to a random walk that mixes to a unique stationary distribution: it follows that the attracting periodic orbits are stable and there are only a finite number of them.

Finally, apply Lemma 2.2 repeatedly, with $\varepsilon = 2^{-l}$ for $l = 1, 2, \dots$ and denote by K_l be the corresponding union of “forbidden” intervals. Define $K^l = \bigcup_{j \geq l} K_j$ and $K^\infty = \bigcap_{l > 0} K^l$; then $\text{Leb}(K^l) \leq 2^{1-l}$ and hence $\text{Leb}(K^\infty) = 0$. The corollary follows from the fact that any $\delta \in \Omega$ outside of K^∞ lies outside of K^l for some $l > 0$. ■

²Indeed, suppose that is not the case; then some \mathbf{x} in a cell of Z , thus outside of Z , would be such that $f(\mathbf{x}) \in Z = Z_{v-1}$. It would follow that $f^k(\mathbf{x}) \in D$, for $k \leq v$; hence $\mathbf{x} \in Z_v = Z$, which a contradiction.

We restate the corollary and add a few facts. For any δ in Ω outside a *critical set* of one-dimensional Lebesgue measure zero, every orbit is attracted to one from a finite set of stable periodic orbits. Within each cell defined by Z , any point produces an orbit that is attracted to the same periodic orbit. We prove that the critical set can be covered by a Cantor set of Hausdorff dimension strictly less than 1.

Corollary 2.3 implies Theorem 2.1. To see why, we observe that the product of any set of n stochastic matrices for lazy irreducible random walks is a positive matrix. Since the set of matrices $S(\mathbf{x})$ is finite, each entry of the product is at least c^n , for some $c > 0$, so the coefficient of ergodicity is at most $1 - nc^n$. By submultiplicativity, $\eta = 2^{O(n)}$, which proves Theorem 2.1.

B. Proving Lemma 2.2

We prove the lemma in three stages. First, we investigate the prefix products of a sequence of well-behaved stochastic matrices and we define a notion of “general position” for them: roughly, our aim is to show that not too many prefix products can map a given point to the same (typical) hyperplane. Second, we establish that, typically, the nested sequence $Z_1 \subseteq Z_2 \subseteq \dots$ cannot grow too rapidly. This, in turn, allows us to bound the topological entropy of the symbolic dynamics of a Markov influence system. Finally, we prove that, for large t , the iterated function f^t maps any of its continuity pieces to a cell so small that it is unlikely to fall in the neighborhood of a discontinuity. We can then use a union bound to ensure that the nesting of the Z_t 's is no longer strict.

1) *General Position for Matrix Products:* Let M_1, \dots, M_T be n -by- n matrices from a finite set \mathcal{M} of primitive stochastic rational matrices with positive diagonals, and assume that their coefficients of ergodicity satisfy $\tau(M) < 1/2$ for any $M \in \mathcal{M}$; hence $\tau(M_1 \cdots M_k) < 2^{-k}$. Because each product $M_1 \cdots M_k$ is a primitive matrix, it can be expressed as $\mathbf{1} \pi_k^\top + Q_k$ (by Perron-Frobenius), where π_k is its (unique) stationary distribution; note that positive diagonals matter here because primitiveness is not closed under multiplication.³ If π is a stationary distribution for a stochastic matrix S , then its j -th row s_j satisfies $s_j - \pi^\top = s_j - \pi^\top S = \sum_i \pi_i (s_j - s_i)$; hence, by the triangular inequality, $\|s_j - \pi^\top\|_1 \leq \sum_i \pi_i \|s_j - s_i\|_1 \leq 2\tau(S)$. This implies that

$$\begin{cases} M_1 \cdots M_k = \mathbf{1} \pi_k^\top + Q_k \\ \|Q_k\|_\infty \leq 2\tau(M_1 \cdots M_k) < 2^{1-k}. \end{cases} \quad (1)$$

We mention a useful geometric interpretation of the previous inequality:

$$\text{diam}_{\ell_\infty}(\mathbb{S}M_1 \cdots M_k) \leq 2\|Q_k\|_\infty. \quad (2)$$

This follows from (1) and, using $Q_k^{(j)}$ to denote the j -th column vector of Q_k ,

³For example, $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ and $\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$ are both primitive but their product is not.

$$\begin{aligned} \text{diam}_{\ell_\infty}(\mathbb{S}M_1 \cdots M_k) &= \text{diam}_{\ell_\infty}(\mathbb{S}Q_k) \\ &= \max_j \left\{ \max_{\mathbf{x} \in \mathbb{S}} \mathbf{x}^\top Q_k^{(j)} - \min_{\mathbf{y} \in \mathbb{S}} \mathbf{y}^\top Q_k^{(j)} \right\} \\ &= \max_j \left\{ \max_i (Q_k)_{ij} - \min_i (Q_k)_{ij} \right\}. \end{aligned}$$

Definition 2.4. Fix a vector $\mathbf{a} \in \mathbb{Q}^n$, and denote by $M^{(\theta)}$ the n -by- m matrix with the m column vectors $M_1 \cdots M_{k_i} \mathbf{a}$, where $\theta = (k_1, \dots, k_m)$ is an increasing sequence of integers in $[T]$. Property **U** is said to hold if there exists a rational vector $\mathbf{u} = \mathbf{u}(\theta)$ such that $\mathbf{1}^\top \mathbf{u} = 1$ and $\mathbf{x}^\top M^{(\theta)} \mathbf{u}$ does not depend on the variable $\mathbf{x} \in \mathbb{S}$.

Observe that, because \mathbf{x} is a probability distribution, property **U** does *not* imply that $M^{(\theta)} \mathbf{u} = \mathbf{0}$; for example, we have $\mathbf{x}^\top (\mathbf{1} \mathbf{1}^\top) \mathbf{u} = 1$ for $\mathbf{u} = \frac{1}{n} \mathbf{1}$. Property **U** is a quantifier elimination device useful for expressing a notion of “general position” for an *MIS*. To see why, consider a simple statement such as “the three points (x, x^2) , $(x+1, (x+1)^2)$, and $(x+2, (x+2)^2)$ cannot be collinear for any value of x .” This can be expressed by saying that a certain determinant polynomial in x is constant. Likewise, the vector \mathbf{u} manufactures a quantity, $\mathbf{x}^\top M^{(\theta)} \mathbf{u}$, that “eliminates” the variable \mathbf{x} . Some condition on \mathbf{u} is needed since otherwise we could pick $\mathbf{u} = \mathbf{0}$. Note that property **U** would be obvious if all the matrices Q_k in (1) were null: indeed, we would have $\mathbf{x}^\top M^{(\theta)} = \mathbf{x}^\top \mathbf{1}(b_1, \dots, b_m) = (b_1, \dots, b_m)$, where $b_i = \pi_{k_i}^\top \mathbf{a}$. This suggests that property **U** rests on the decaying properties of Q_k .

General position played an important role in proving the periodicity of planar piecewise contractions [5] and the same is true here. Whereas in [5], the claim of general position follows from a simple dimensionality argument, it is here the heart of the proof and requires ideas from linear algebra and Ramsey-type extremal set theory. To see the relevance of general position to the dynamics of an *MIS*, consider the iterates of a small ball through the map f . To avoid chaos, it is intuitively obvious that these iterated images should not fall across discontinuities too often. Fix such a discontinuity: if we think of the ball as being so small it looks like a point, then the case we are trying to avoid consists of many points (the ball's iterates) lying on (or near) a given hyperplane. This is similar to the definition of general position, which requires that a large enough set of points should not lie on the same hyperplane.

Lemma 2.5. There exists a constant $b > 0$ (linear in n) such that, given any integer $T > 0$ and any increasing sequence θ in $[T]$ of length at least $T^{1-\alpha}/\alpha$, property **U** holds, where $\alpha := \mu^{-b}$ and μ is the maximum number of bits needed to encode any entry of M_k for any $k \in [T]$.

Proof. By choosing b large enough, we can ensure that T is as big as we want. The proof is a mixture of algebraic and combinatorial arguments. ■

Fact 2.6. There is a constant $d > 0$ such that, if the sequence θ contains j_0, \dots, j_n with $j_i \geq d\mu j_{i-1}$ for each $i \in [n]$, then property **U** holds.

Proof. By (1), $\|Q_k \mathbf{a}\|_\infty < c_0 2^{-k}$ for constant $c_0 > 0$. Note that Q_k has rational entries over $O(\mu k)$ bits: the bound follows from the fact that the stationary distribution π_k has rational

coordinates over $O(\mu k)$ bits; as noted earlier, the constant factors may depend on n . We write $M^{(\theta)} = \mathbf{1}\mathbf{a}^\top \Pi^{(\theta)} + Q^{(\theta)}$, where $\Pi^{(\theta)}$ and $Q^{(\theta)}$ are the n -by- m matrices formed by the m column vectors π_{k_i} and $Q_{k_i}\mathbf{a}$, respectively, for $i \in [m]$; recall that $\theta = (k_1, \dots, k_m)$. The key fact is that the dependency on $\mathbf{x} \in \mathbb{S}$ is confined to the term $Q^{(\theta)}$: indeed,

$$\mathbf{x}^\top M^{(\theta)}\mathbf{u} = \mathbf{a}^\top \Pi^{(\theta)}\mathbf{u} + \mathbf{x}^\top Q^{(\theta)}\mathbf{u}. \quad (3)$$

This shows that, in order to satisfy property **U**, it is enough to ensure that $Q^{(\theta)}\mathbf{u} = 0$ has a solution such that $\mathbf{1}^\top \mathbf{u} = 1$. Let $\sigma = (j_0, \dots, j_{n-1})$. If $Q^{(\sigma)}$ is nonsingular then, because each one of its entries is a rational over $O(\mu j_{n-1})$ bits, we have $|\det Q^{(\sigma)}| \geq c_1^{\mu j_{n-1}}$, for constant $c_1 > 0$. Let R be the $(n+1)$ -by- $(n+1)$ matrix derived from $Q^{(\sigma)}$ by adding the column $Q_{j_n}\mathbf{a}$ to its right and then adding a row of ones at the bottom. If R is nonsingular, then $R\mathbf{u} = (0, \dots, 0, 1)^\top$ has a (unique) solution in \mathbf{u} and property **U** holds (after padding \mathbf{u} with zeroes). Otherwise, we expand the determinant of R along the last column. Suppose that $\det Q^{(\sigma)} \neq 0$. By Hadamard's inequality, all the cofactors are at most a constant $c_2 > 0$ in absolute value; hence, for d large enough,

$$0 = |\det R| \geq |\det Q^{(\sigma)}| - nc_2 \|Q_{j_n}\mathbf{a}\|_\infty \geq c_1^{\mu j_{n-1}} - nc_2 c_0 2^{-j_n} > 0.$$

This contradiction implies that $Q^{(\sigma)}$ is singular, so (at least) one of its rows can be expressed as a linear combination of the others. We form the n -by- n matrix R' by removing that row from R , together with the last column, and setting $u_{j_n} = 0$ to rewrite $Q^{(\theta)}\mathbf{u} = 0$ as $R'\mathbf{u}' = (0, \dots, 0, 1)^\top$, where \mathbf{u}' is the restriction of \mathbf{u}' to the columns indexed by R' . Having reduced the dimension of the system by one variable, we can proceed inductively in the same way; either we terminate with the discovery of a solution or the induction runs its course until $n = 1$ and the corresponding 1-by-1 matrix is null, so that the solution 1 works. Note that \mathbf{u} has rational coordinates over $O(\mu T)$ bits. ■

We are now in a position to prove Lemma 2.5. Let $N(T)$ be the largest sequence θ in $[T]$ such that property **U** does not hold. Divide $[T]$ into bins $[(d\mu)^k, (d\mu)^{k+1} - 1]$ for $k \geq 0$. By Fact 2.6, the sequence θ can intersect at most $2n$ of them; thus, if $T > t_0$, for some large enough $t_0 = (d\mu)^{O(n)}$, there is at least one empty interval in T of length $T/(d\mu)^{2n+3}$. This gives us the recurrence $N(T) \leq T$ for $T \leq t_0$ and $N(T) \leq N(T_1) + N(T_2)$, where $T_1 + T_2 \leq \beta T$, for a positive constant $\beta = 1 - (d\mu)^{-2n-3}$. The recursion to the right of the empty interval, say, $N(T_2)$, warrants a brief discussion. The issue is that the proof of Fact 2.6 relies crucially on the property that Q_k has rational entries over $O(\mu k)$ bits—this is needed to lower-bound $|\det Q^{(\sigma)}|$ when it is not 0. But this is not true any more, because, after the recursion, the columns of the matrix $M^{(\theta)}$ are of the form $M_1 \cdots M_k \mathbf{a}$, for $T_1 + L < k \leq T$, where L is the length of the empty interval and $T = T_1 + L + T_2$. Left as such, the matrices use too many bits for the recursion to go through. To overcome this obstacle, we observe that the recursively transformed $M^{(\theta)}$ can be factored as AB , where $A = M_1 \cdots M_{T_1+L}$ and B consists of the column vectors $M_{T_1+L+1} \cdots M_k \mathbf{a}$. The key observation now is that, if $\mathbf{x}^\top B\mathbf{u}$

does not depend on \mathbf{x} , then neither does $\mathbf{x}^\top M^{(\theta)}\mathbf{u}$, since it can be written as $\mathbf{y}^\top B\mathbf{u}$ where $\mathbf{y} = A^\top \mathbf{x} \in \mathbb{S}$. In this way, we can enforce property **U** while having restored the proper encoding length for the entries of $M^{(\theta)}$.

Plugging in the ansatz $N(T) \leq t_0 T^\gamma$, for some unknown positive $\gamma < 1$, we find by Jensen's inequality that, for all $T > 0$, $N(T) \leq t_0(T_1^\gamma + T_2^\gamma) \leq t_0 2^{1-\gamma} \beta^\gamma T^\gamma$. For the ansatz to hold true, we need to ensure that $2^{1-\gamma} \beta^\gamma \leq 1$. Setting $\gamma = 1/(1 - \log \beta) < 1$ completes the proof of Lemma 2.5.

2) *Symbolic Dynamics*: The growth exponent of a language is defined as $\lim_{n \rightarrow \infty} \frac{1}{n} \max_{k \leq n} \log N(k)$, where $N(k)$ is the number of words of length k ; for example, the growth exponent of $\{0, 1\}^*$ is 1 (all logarithms taken to the base 2). The language consisting of all the itineraries of a Markov influence system forms a shift space and its growth exponent is the topological entropy of its symbolic dynamics [21], [34]. We show that, typically, it is zero.

Define $\phi^k(\mathbf{x}) = \mathbf{x}^\top M_1 \cdots M_k$ for $\mathbf{x} \in \mathbb{R}^n$ and $k \leq T$; and let $h_\delta: \mathbf{a}^\top \mathbf{x} = 1 + \delta$ be some hyperplane in \mathbb{R}^n . We consider a set of canonical intervals of length ρ (or less): $\mathcal{D}_\rho = \{[k\rho, (k+1)\rho] \cap \Omega \mid k \in \mathbb{Z}\}$, where $\rho > 0$ (specified below), $\Omega = \omega\mathbb{I}$, $\mathbb{I} := [-1, 1]$, and $0 < \omega < 1/2$. Roughly, the ‘‘general position’’ lemma below says that, for most δ , the ϕ^k -images of any ρ -wide cube centered in the simplex \mathbb{S} cannot come very near the hyperplane h_δ for most values of $k \leq T$.

This may be counterintuitive. After all, if the stochastic matrices M_i are the identity, the images stay put, so if the initial cube collides then all of the images will! The point is that M_i is primitive so it cannot be the identity. The low coefficients of ergodicity will also play a key role. The crux of the lemma is that the exclusion set U does not depend on the choice of $\mathbf{x} \in \mathbb{S}$.

Lemma 2.7. For any real $\rho > 0$ and any integer $T > 0$, there exists $U \subseteq \mathcal{D}_\rho$ of size $c_T = 2^{O(\mu T)}$, where c_T is independent of ρ , such that, for any $\Delta \in \mathcal{D}_\rho \setminus U$ and $\mathbf{x} \in \mathbb{S}$, there are at most $T^{1-\alpha}/\alpha$ integers $k \leq T$ such that $\phi^k(X) \cap h_\Delta \neq \emptyset$, where $X = \mathbf{x} + \rho\mathbb{I}^n$ and $h_\Delta := \bigcup_{\delta \in \Delta} h_\delta$.

Proof. A point of notation: α refers to its use in Lemma 2.5; also, b_0, b_1, \dots refer to suitably large positive constants (which, we shall recall, may depend on n, \mathbf{a} , etc). We assume the existence of more than $T^{1-\alpha}/\alpha$ integers $k \leq T$ such that $\phi^k(X) \cap h_\Delta \neq \emptyset$ for some $\Delta \in \mathcal{D}_\rho$ and draw the consequences: in particular, we infer certain linear constraints on δ ; by negating them, we define the forbidden set U and ensure the conclusion of the lemma. Let $k_1 < \dots < k_m$ be the integers in question, where $m > T^{1-\alpha}/\alpha$. For each $i \in [m]$, there exists $\mathbf{x}(i) \in X$ and $\delta_i \in \Delta$ such that $|\mathbf{x}(i)^\top M_1 \cdots M_{k_i} \mathbf{a} - 1 - \delta_i| \leq \rho$. Note that $|\delta_i - \delta| \leq \rho$ for some $\delta \in \Delta$ common to all $i \in [m]$. By the stochasticity of the matrices, it follows that $|\mathbf{x}(i) - \mathbf{x}^\top M_1 \cdots M_{k_i} \mathbf{a}| \leq b_0 \rho$; hence $|\mathbf{x}^\top M_1 \cdots M_{k_i} \mathbf{a} - 1 - \delta| \leq (b_0 + 2)\rho$. By Lemma 2.5, there is a rational vector \mathbf{u} such that $\mathbf{1}^\top \mathbf{u} = 1$ and $\mathbf{x}^\top M^{(\theta)}\mathbf{u} = \psi(M^{(\theta)}, \mathbf{a})$ does not depend on the variable $\mathbf{x} \in \mathbb{S}$; on the other hand, $|\mathbf{x}^\top M^{(\theta)}\mathbf{u} - (1 + \delta)| \leq b_1 \rho$. Two remarks: (i) the term $1 + \delta$ is derived from $(1 + \delta)\mathbf{1}^\top \mathbf{u} = 1 + \delta$; (ii) $b_1 \leq (b_0 + 2)\|\mathbf{u}\|_1$, where \mathbf{u} is a rational over $O(\mu T)$ bits. We invalidate the condition on k_1, \dots, k_m by keeping δ

outside the interval $\psi(M^{(\theta)}, \mathbf{a}) - 1 + b_1\rho\mathbb{I}$, which rules out at most $2(b_1 + 1) = 2^{O(\mu T)}$ intervals from \mathcal{D}_ρ . Repeating this for all sequences (k_1, \dots, k_m) raises the number of forbidden intervals by a factor of at most 2^T . ■

We identify the family \mathcal{M} with the set of all matrices of the form $S_1 \cdots S_k$, for $\eta \leq k \leq 3\eta$, where the matrix sequence S_1, \dots, S_k matches some element of L_Ω^k . By definition of the ergodic scale, any $M \in \mathcal{M}$ is primitive and $\tau(M) < 1/2$; furthermore, both μ and $\log|\mathcal{M}|$ are in $O(\eta)$. Our next result implies that the topological entropy of the shift space of itineraries vanishes.

Lemma 2.8. For any real $\rho > 0$ and any integer $T > 0$, there exist $t_\rho = O(\eta|\log\rho|)$ and an exclusion set $V \subseteq \mathcal{D}_\rho$ of size $d_T = 2^{O(T)}$ such that, for any $\Delta \in \mathcal{D}_\rho \setminus V$, any integer $t \geq t_\rho$, and any $\sigma \in L_\Delta^t$, $\log|\{\sigma' \mid \sigma \cdot \sigma' \in L_\Delta^{t+T}\}| \leq \eta^b T^{1-\eta^{-b}}$, for constant $b > 0$.

In the lemma, t_ρ (resp. d_T) is independent of T (resp. ρ). The main point is that the exponent of T is strictly bounded above by 1.

Proof. We define V as the union of the sets U formed by applying Lemma 2.7 to each one of the hyperplanes h_δ involved in \mathcal{P} and every possible sequence of T matrices in \mathcal{M} . This increases c_T to $2^{O(\eta T)}$. Fix $\Delta \in \mathcal{D}_\rho \setminus V$ and consider the (lifted) phase space $\mathbb{S} \times \Delta$ for the dynamical system induced by the map $f_\uparrow: (\mathbf{x}^\top, \delta) \mapsto (\mathbf{x}^\top S(\mathbf{x}), \delta)$. The system is piecewise-linear with respect to the polyhedral partition \mathcal{P}_\uparrow of \mathbb{R}^{n+1} formed by treating δ as a variable in h_δ . Let Υ_t be a continuity piece for f_\uparrow^t , ie, a maximal polyhedron within $\mathbb{S} \times \Delta$ over which the t -th iterate of f_\uparrow is linear. Reprising the argument leading to (1), any matrix sequence S_1, \dots, S_k matching an element of L_Δ^k is such that $S_1 \cdots S_k = \mathbf{1}\boldsymbol{\pi}^\top + Q$, where

$$\|Q\|_\infty < 2^{2-t/\eta}. \quad (4)$$

Thus, by (2), there exists $t_\rho = O(\eta|\log\rho|)$ such that, for any $t \geq t_\rho$, $f_\uparrow^t(\Upsilon_t) \subseteq (\mathbf{x} + \rho\mathbb{I}^n) \times \Delta$, for some $\mathbf{x} = \mathbf{x}(t, \Upsilon_t) \in \mathbb{S}$. Consider a nested sequence $\Upsilon_1 \supseteq \Upsilon_2 \supseteq \dots$. Note that Υ_1 is a cell of \mathcal{P}_\uparrow , $f_\uparrow^k(\Upsilon_{k+1}) \subseteq f_\uparrow^k(\Upsilon_k)$, and S_k is the stochastic matrix used to map $f_\uparrow^{k-1}(\Upsilon_k)$ to $f_\uparrow^k(\Upsilon_k)$ (ignoring the δ -axis). We say there is a *split* at k if $\Upsilon_{k+1} \subset \Upsilon_k$, and we show that, given any $t \geq t_\rho$, there are only $O(\eta T^{1-\alpha}/\alpha)$ splits between t and $t + \eta T$, where $\alpha = \eta^{-b}$, for constant b .⁴ We may confine our attention to splits caused by the same hyperplane h_δ since \mathcal{P} features only a constant number of them. Arguing by contradiction, we assume the presence of at least $6\eta T^{1-\alpha}/\alpha$ splits, which implies that at least $N := 2T^{1-\alpha}/\alpha$ of those splits occur for values of k at least 2η apart. This is best seen by binning $[t + 1, t + \eta T]$ into T intervals of length η and observing that at least $3N$ intervals must feature splits. In fact, this proves the existence of N splits at positions separated by a least two consecutive bins. Next, we use the same binning to produce the matrices M_1, \dots, M_T , where $M_j = S_{t+1+(j-1)\eta} \cdots S_{t+j\eta}$.

Suppose that all of the N splits occur for values k of the form $t + j\eta$. In this case, a straightforward application of

⁴We may have to scale b up by a constant factor since $\mu = O(\eta)$ and, by Lemma 2.5, $\alpha = \mu^{-b}$.

Lemma 2.7 is possible: we set $X \times \Delta = f_\uparrow^t(\Upsilon_t)$ and note that the functions ϕ^k are all products of matrices from the family \mathcal{M} , which happen to be η -long products. The number of splits, $2T^{1-\alpha}/\alpha$, exceeds the number allowed by the lemma and we have a contradiction. If the splits do not fall neatly at endpoints of the bins, we use the fact that \mathcal{M} includes matrix products of any length between η and 3η . This allows us to reconfigure the bins so as to form a sequence M_1, \dots, M_T with the splits occurring at the endpoints: for each split, merge its own bin with the one to its left and the one to its right (neither of which contains a split) and use the split's position to subdivide the resulting interval into two new bins; we leave all the other bins alone.⁵ This leads to the same contradiction, which implies the existence of fewer than $O(\eta T^{1-\alpha}/\alpha)$ splits at $k \in [t, t + \eta T]$; hence the same bound on the number of strict inclusions in the nested sequence $\Upsilon_t \supseteq \dots \supseteq \Upsilon_{t+\eta T}$. The set of all such sequences forms a tree of depth ηT , where each node has at most a constant number of children and any path from the root has $O(\eta T^{1-\alpha}/\alpha)$ nodes with more than one child. Rescaling T to ηT and raising b completes the proof. ■

3) *Putting Everything Together:* We show that the excluded intervals in Ω can be covered by a Cantor set of Hausdorff dimension less than one. All the parameters below refer to Lemma 2.8. We fix $\varepsilon > 0$ and $\Delta \in \mathcal{D}_\rho \setminus V$ and assume that $\delta \in \Delta$. Set $T = 2^{\eta^{2b}}$, $\rho = \varepsilon/(2d_T)$, and $v = t_\rho + kT$, where $k = c\eta\log(1/\rho)$ for a large enough constant $c > 0$. Since $t_\rho = O(\eta|\log\rho|)$ and $d_T = 2^{O(T)}$, we have

$$v = 2^{\eta^{O(1)}} \log(1/\varepsilon). \quad (5)$$

Let M be the matrix $S_1 \cdots S_v$, where S_1, \dots, S_v the matrix sequence matching an element of L_Δ^v . By (2, 4), $\text{diam}_{\ell_\infty}(\mathbb{S}M) \leq 2^{3-v/\eta}$. There exists a point \mathbf{x}_M such that, given any point $\mathbf{y} \in \mathbb{S}$ whose v -th iterate $f^v(\mathbf{y}) = \mathbf{z}^T$ is specified by the matrix M , that is, $\mathbf{z}^T = \mathbf{y}^T M$, we have $\|\mathbf{x}_M - \mathbf{z}\|_\infty \leq 2^{3-v/\eta}$. Consider a discontinuity $h_\delta: \mathbf{a}_\delta^\top \mathbf{x} = 1 + \delta$ of the system. Testing which side of it the point \mathbf{z} lies is equivalent to checking the point \mathbf{x}_M instead with respect to $h_{\delta'}$ for some δ' that differs from δ by $O(2^{-v/\eta})$. It follows that adding an interval of length $O(2^{-v/\eta})$ to the exclusion set V ensures that all the v -th iterates $f^v(\mathbf{y})$ (specified by M) lie strictly on the same side of h_δ for all $\delta \in \Delta$. Repeating this for every string L_Δ^v and every $\Delta \in \mathcal{D}_\rho \setminus V$ increases the length covered by V from its original $d_T\rho = \varepsilon/2$ to at most $d_T\rho + O(|L_\Delta^v|2^{-v/\eta}/\rho) < \varepsilon$. This last bound follows from the consequence of Lemma 2.8 that $\log|L_\Delta^v| \leq k\eta^b T^{1-\eta^{-b}} + O(t_\rho)$. Thus, for any $\delta \in \Omega$ outside a set of intervals covering a length less than ε , no $f^v(\mathbf{x})$ lies on a discontinuity. It follows that, for any such δ , we have $Z_v = Z_{v-1}$ and, by (5), the proof of Lemma 2.2 is complete.

⁵We note the possibility of an inconsequential decrease in T caused by the merges. Also, we can now see clearly why Lemma 2.7 is stated in terms of the slab h_Δ and not the hyperplane h_δ . This allows us to express splitting caused by the hyperplane $\mathbf{a}^\top \mathbf{x} = 1 + \delta$ in lifted space \mathbb{R}^{n+1} .

III. HIGHER TIMESCALES AND CHAOS

Among the *MIS* that converge to a single stationary distribution, some of them feature super-exponential mixing time. Very slow clocks can be designed in the same manner: the *MIS* is periodic with a period of length equal to a tower-of-twos of height linear in the dimension. The creation of new timescales is what most distinguishes *MIS* from standard Markov chains. As we mentioned earlier, the systems can be chaotic as well.

A. Hyper-Torpid Mixing

How can reaching a fixed point distribution take so long? Before we answer this question formally, we provide a bit of intuition. Imagine having three unit-volume water reservoirs A, B, C alongside a clock that rings at 12 every day (that is, twice a day, at noon and at midnight). Initially, the clock is at 1 and A is full while B and C are empty. Reservoir A transfers half of its content to B and repeats this each hour until the clock rings 12. At this point, reservoir A empties into C the little water that it has left and B empties its content into A . At 1, we repeat the previous action: A transfers half of its water content to B , etc. This goes on until some day, the clock rings and the reservoir C finds its more than half full: this is sure to happen since the water level of C rises by about 10^{-3} at each cycle. At this point, both B and C transfer all their water back to A , so that when the clock is at 1, we are back to square one. The original 12-step clock has been extended into a new clock of period roughly 1,000. The proof below shows how to simulate this iterative process with an *MIS*.

Theorem 3.1. *There exist Markov influence systems that mix to a stationary distribution in time equal to a tower-of-twos of height linear in the number of states.*

Proof. We construct an *MIS* with a periodic orbit of length equal to a power-of-twos of height proportional to n ; this is the function $f(n) = 2^{f(n-1)}$, with $f(1) = 1$. It is easy to turn it into one with an orbit that is attracted to a stationary distribution (a fixed point) with an equivalent mixing rate, and we omit this part of the discussion. Assume, by induction, that we have a Markov influence system M cycling through states $1, \dots, p$, for $p \geq 4$. We build another one with period at least 2^p by adding a “gadget” to it consisting of a graph over the vertices $1, 2, 3$ with probability distribution $(x, y, z) \in \mathbb{S}$. We initialize the system by placing M in state 1 (ie, 1pm in our clock example) and setting $x = 1$. The dynamic graph is specified by these rules:

- 1) Suppose that M is in state $1, \dots, p - 1$. The graph has the edge $(1, 2)$, which is assigned probability $1/2$, as is the self-loop at 1. There are self-loops at 2 and 3.
- 2) Suppose that M is in state p . If $z \leq 1/2$, then the graph has the edge $(2, 1)$ and $(1, 3)$, both of them assigned probability 1, with one self-loop at 3. If $z > 1/2$, then the graph has the edge $(2, 1)$ and $(3, 1)$, both of them assigned probability 1, with one self-loop at 3.

Suppose that M is in state 1 and that $y = 0$ and $z \leq 1/2$. When M reaches state $p - 1$, then $x = (1 - z)2^{1-p}$ and $y = (1 - z)(1 - 2^{1-p})$. Since $z \leq 1/2$, the system cycles back

to state 1 with the updates $z \leftarrow x + z$ and $x \leftarrow y$. Note that z increases by a number between 2^{-p} and 2^{1-p} . Since z begins at 0, such increases will occur consecutively at least $p2^p/4 \geq 2^p$ times, before x is reset to 1. The construction on top of M adds three new vertices so we can push this recursion roughly $n/3$ times to produce a Markov influence system that is periodic with a period of length equal to a tower-of-twos of height roughly $n/3$.

We need to tie up a few loose ends. The construction needs to recognize state $p - 1$ by a polyhedral cell; in fact, any state will do. The easiest choice is state 1, which corresponds to $x \geq 1$ (to express it as an inequality). The base case of our inductive construction consists of a two-vertex system of period $p = 4$ with initial distribution $(1, 0)$. If $x > 2^{1-p}$, the graph has an edge from 1 to 2 and a self-loop at 1, both of them assigned probability $1/2$; else an edge from 2 to 1 given probability 1 to reset the system. Finally, the construction assumes probabilities summing up to 1 within each of the $\lfloor (n - 2)/3 \rfloor + 1$ gadgets, which is clearly wrong: we fix this by dividing the probability weights equally among the gadget and adjusting the linear discontinuities appropriately. ■

B. A Chaotic Markov Influence System

We give a simple 5-state construction with chaotic symbolic dynamics. The idea is to build an *MIS* to simulate the classic baker’s map. Given $\mathbf{x} \in \mathbb{S}^4$, $S(\mathbf{x}) = A$ if $x_1 + x_2 > x_4$, and $S(\mathbf{x}) = B$ otherwise, with

$$A = \frac{1}{3} \begin{pmatrix} 2 & 1 & 0 & 0 & 0 \\ 0 & 1 & 2 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 \\ 0 & 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 & 3 \end{pmatrix}$$

and

$$B = \frac{1}{3} \begin{pmatrix} 1 & 0 & 2 & 0 & 0 \\ 1 & 2 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 \\ 0 & 0 & 0 & 2 & 1 \\ 0 & 0 & 0 & 0 & 3 \end{pmatrix}.$$

We focus our attention on $\Sigma = \{(x_1, x_2, x_4) \mid 0 < x_1 \leq x_4/2 \leq x_2 < x_4\}$, and easily check that it is an invariant manifold. At time 0, we fix $x_4 = 1/4$ and $x_5 = 0$; at all times, of course, $x_3 = 1 - x_1 - x_2 - x_4 - x_5$. The variable $y := (2x_2 - x_4)/(2x_1 - x_4)$ is always nonpositive over Σ . It evolves as follows:

$$y \leftarrow \begin{cases} \frac{1}{2}(y + 1) & \text{if } y < -1 \\ \frac{2y}{y+1} & \text{if } -1 \leq y \leq 0. \end{cases}$$

Writing $z = (y + 1)/(y - 1)$, we note that $-1 \leq z < 1$ and it evolves according to $z \mapsto 2z + 1$ if $z \leq 0$, and $z \mapsto 2z - 1$ otherwise, a map that conjugates with the baker’s map and is known to be chaotic [16].

IV. CONCLUDING REMARKS


We have established the typical asymptotic periodicity of irreducible Markov influence systems. Informally, this means that a random walk over a graph changing endogenously as a function of the current probability distribution will almost surely converge to a periodic orbit. Our proof assumes that all the graphs are strongly connected. We conjecture that this assumption can be relaxed. To do so, however, would seem to require an understanding of graph renormalization [11] that is beyond our reach at the moment. We leave this as an exciting open problem.

ACKNOWLEDGEMENT

The author would like to thank the referees for their helpful comments and suggestions. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office, the Defense Advanced Research Projects Agency, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

REFERENCES

- [1] A. Anagnostopoulos, R. Kumar, M. Mahdian, E. Upfal, and F. Vandin, "Algorithms on evolving graphs," in *Proc. 3rd Innov. Theor. Comput. Sci.*, 2012, pp. 149–160.
- [2] C. Avin, M. Koucký, and Z. Lotker, "How to explore a fast-changing world (cover time of a simple random walk on evolving graphs)," in *Proc. 35th Int. Colloq. Automata Lang. Program.*, 2008, pp. 121–132.
- [3] A. Barrat, M. Barthélemy, and A. Vespignani, *Dynamical Processes on Complex Networks*. Cambridge, U.K.: Cambridge Univ. Press, 2008.
- [4] J. Brémont, "Dynamics of injective quasi-contractions," *Ergodic Theory Dyn. Syst.*, vol. 26, pp. 19–44, 2006.
- [5] H. Bruin and J. H. B. Deane, "Piecewise contractions are asymptotically periodic," *Proc. Amer. Math. Soc.*, vol. 137, no. 4, pp. 1389–1395, 2009.
- [6] J. Buzzi, "Piecewise isometries have zero topological entropy," *Ergodic Theory Dyn. Syst.*, vol. 21, pp. 1371–1377, 2001.
- [7] C. Castellano, S. Fortunato, and V. Loreto, "Statistical physics of social dynamics," *Rev. Mod. Phys.*, vol. 81, pp. 591–646, 2009.
- [8] E. Catsigeras and R. Budelli, "Topological dynamics of generic piecewise continuous contractive maps in n dimensions," *Int. J. Pure Appl. Math.*, vol. 68, pp. 61–83, 2011.
- [9] E. Catsigeras, P. Guiraud, A. Meyroneinc, and E. Ugalde, "On the asymptotic properties of piecewise contracting maps," *Dyn. Sys.*, vol. 31, no. 2, pp. 107–135, 2016.
- [10] B. Chazelle, "Diffusive influence systems," *SIAM J. Comput.*, vol. 44, pp. 1403–1442, 2015.
- [11] B. Chazelle, "Toward a theory of Markov influence systems and their renormalization," in *Proc. 9th Innov. Theor. Comput. Sci.*, 2018, pp. 58:1–58:18.
- [12] A. Condon and D. Hernek, "Random walks on colored graphs," *Random Struct. Algorithms*, vol. 5, pp. 285–303, 1994.
- [13] A. Condon and R. J. Lipton, "On the complexity of space bounded interactive proofs," in *Proc. 30th IEEE Symp. Found. Comput. Sci.*, 1989, pp. 462–267.
- [14] O. Denysyuk and L. Rodrigues, "Random walks on directed dynamic graphs," in *Proc. 2nd Int. Workshop Dyn. Netw., Algorithms Secur. (DYNASIO)*, Bordeaux, France, Jul. 2010.
- [15] O. Denysyuk and L. Rodrigues, "Random walks on evolving graphs with recurring topologies," in *Proc. 28th Int. Symp. Distrib. Comput. (DISC)*, Austin, Texas, USA, Oct. 2014, pp. 333–345.
- [16] R. L. Devaney, *An Introduction to Chaotic Dynamical Systems*, 2nd ed. Boulder, CO, USA: Westview, 2003.
- [17] F. Fagnani and P. Frasca, *Introduction to Averaging Dynamics Over Networks* (Lecture Notes in Control and Information Sciences), vol. 472. New York, NY, USA: Springer, 2018.
- [18] P. Holme, "Modern temporal networks theory: A colloquium," *Eur. Physical J. B*, vol. 88, no. 9, p. 1, 2015.
- [19] P. Holme and J. Saramäki, "Temporal networks," *Phys. Rep.*, vol. 519, pp. 97–125, 2012.
- [20] G. Iacobelli and D. R. Figueiredo, "Edge-attractor random walks on dynamic networks," *J. Complex Netw.*, vol. 5, pp. 84–110, 2017.
- [21] A. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*. Cambridge, U.K.: Cambridge Univ. Press, Dec. 28, 1996.
- [22] D. Kempe, J. Kleinberg, and A. Kumar, "Connectivity and inference problems for temporal networks," *J. Comput. Syst. Sci.*, vol. 64, pp. 820–842, 2002.
- [23] B. Kruglikov and M. Rypdal, "A piecewise affine contracting map with positive entropy," *Discrete Continuous Dyn. Syst.*, vol. 16, no. 2, pp. 393–394, 2006.
- [24] O. Michail, "An introduction to temporal graphs: An algorithmic perspective," *Internet Math.*, vol. 12, pp. 239–280, 2016.
- [25] O. Michail and P. G. Spirakis, "Elements of the theory of dynamic networks," *Commun. ACM*, vol. 61, no. 2, pp. 72–81, Feb. 2018.
- [26] G. H. Nguyen, J. B. Lee, R. A. Rossi, N. Ahmed, E. Koh, and S. Kim, "Dynamic network embeddings: From random walks to temporal random walks," in *Proc. IEEE Int. Conf. Big Data*, 2018, pp. 1085–1092.
- [27] N. Perra, A. Baronchelli, D. Mocanu, B. Gonçalves, R. Pastor-Satorras, and A. Vespignani, "Random walks and search in time-varying networks," *Phys. Rev. Lett.*, vol. 109, 2012, Art. no. 238701.
- [28] J. Petit, M. Gueuning, T. Carletti, B. Lauwens, and R. Lambiotte, "Random walk on temporal networks with lasting edges," *Phys. Rev. E*, vol. 98, 2018, Art. no. 052307.
- [29] A. V. Proskurnikov and R. Tempo, "A tutorial on modeling and analysis of dynamic social networks. Part I," *Annu. Rev. Control*, vol. 43, pp. 65–79, 2017.
- [30] A. V. Proskurnikov and R. Tempo, "A tutorial on modeling and analysis of dynamic social networks. Part II," *Annu. Rev. Control*, vol. 45, pp. 166–190, 2018.
- [31] V. Ramiro, E. Lochin, P. Sénac, and T. Rakotoarivelo, "Temporal random walk as a lightweight communication infrastructure for opportunistic networks," in *Proc. IEEE Int. Symp. (WoWMoM)*, 2014, pp. 1–6.
- [32] E. Seneta, *Non-Negative Matrices and Markov Chains*, 2nd ed. New York, NY, USA: Springer, 2006.
- [33] M. Starnini, A. Baronchelli, A. Barrat, and R. Pastor-Satorras, "Random walks on temporal networks," *Phys. Rev. E*, vol. 85, no. 5, 2012, Art. no. 056115.
- [34] S. Sternberg, *Dynamical Systems*. New York, NY, USA: Dover, 2010.



Bernard Chazelle received the Ph.D. degree in computer science from Yale University, New Haven, CT, USA, in 1980. He is currently an Eugene Higgins Professor in computer science with Princeton University, Princeton, NJ, USA, where he has been in the faculty since 1986. He was a Professor with the Collège de France, Paris, France, in recent years as well as a member of the Institute for Advanced Study, Princeton, NJ, USA. He has authored several books, including *The Discrepancy Method* (Cambridge University Press, 2001). His current research focuses on natural algorithms and computation. He is a fellow of the American Academy of Arts and Sciences and a member of the European Academy of Sciences, as well as an ACM fellow and a former Guggenheim fellow.