Linear Algebra and its Applications 619 (2021) 176–209

Contents lists available at ScienceDirect

Linear Algebra and its Applications

www.elsevier.com/locate/laa

Triangulated Laman graphs, local stochastic matrices, and limits of their products

Mohamed-Ali Belabbas [a], Xudong Chen [b,*]

[a] Coordinated Science Laboratory, University of Illinois, Urbana-Champaign, United States of America

[b] Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, United States of America

ARTICLE INFO

Article history:
Received 1 November 2020
Accepted 21 February 2021
Available online 27 February 2021
Submitted by R. Brualdi

MSC:
15B51
40A20
52C25
60J05

Keywords:
Stochastic matrices
Infinite products
Rigid graphs
Consensus

ABSTRACT

We derive conditions on the products of stochastic matrices guaranteeing the existence of a unique limit invariant distribution. Belying our approach is the hereby defined notion of restricted triangulated Laman graphs. The main idea is the following: to each triangle in the graph, we assign a stochastic matrix. Two matrices can be adjacent in a product only if their corresponding triangles share an edge in the graph. We provide an explicit formula for the limit invariant distribution of the product in terms of the individual stochastic matrices.
© 2021 Elsevier Inc. All rights reserved.

1. Introduction

The issue of convergence of infinite product of (row) stochastic matrices arises naturally in the study of finite-state Markov chains and in the design of consensus algorithms.

* Corresponding author.
E-mail addresses: belabbas@illinois.edu (M.-A. Belabbas), xudong.chen@colorado.edu (X. Chen).

https://doi.org/10.1016/j.laa.2021.02.016
0024-3795/© 2021 Elsevier Inc. All rights reserved.

As a result, it has been widely investigated in the past decades [1–7] from a variety of perspectives. The main problem investigated in the above works is whether the limit of a left product $\lim_{k\to\infty} A_k \cdots A_2 A_1$ converges to a rank one matrix $\mathbf{1}w^\top$, where $\mathbf{1}$ is a vector of all ones and w is a probability vector, i.e., entries of w are nonnegative and sum to 1.

A less studied, yet critical, problem is to characterize the limit beyond the fact that it is rank one. This amounts to the characterization of the probability vector w. In the context of Markov chain, w is the limiting distribution while in the context of (weighted) consensus, entries of w are the averaging weights in the convex combination. The problem is hard to tackle. Indeed, barring simple cases such as using only commuting matrices, the limit depends on the order in which the stochastic matrices appear in the infinite products. This is true even if the matrices appearing in the product are chosen from a finite set. See [3] for some illustrations of the above mentioned dependence. Thus, without knowing the entire sequence a priori, it is in general infeasible to characterize the limit (provided that it exists). In fact, even if one knows the order of the entire sequence, the analysis for obtaining an explicit formula of the limit is often intractable.

In this paper, we address this latter problem, i.e., we characterize limits of certain products of stochastic matrices. We elaborate below on the type of products considered in the paper. As is usually done, we use a graph $G = (V, E)$ to represent the states of Markov chain and the allowable transitions between these states. In the context of consensus, the graph represents the information-flow topology between different agents. We introduce a class of graphs, termed triangulated Laman graphs (TLGs), and use their structure to define sets of stochastic matrices and the orders in which we can take their products. Specifically, given any TLG, we assign a stochastic matrix to each triangle in the graph. The matrix can be obtained by starting with the identity matrix and, then, replacing the principal submatrix corresponding to the nodes in the triangle with an arbitrary 3×3 rank-one stochastic matrix. We call these matrices "local stochastic matrices" as the transition probabilities (or the communications in the context of consensus) involve only the nodes in that triangle. To describe the allowable products of those local stochastic matrices, we introduce the notion of derived graph associated with a TLG. It is a graph whose nodes are the triangles of a TLG, and whose edges capture a notion of adjacency between these triangles — two triangles are adjacent if they share a common edge. The allowable products are then the ones for which adjacent matrices correspond to adjacent nodes in the derived graph.

A major contribution of the paper is to show that if a walk in the derived graph visits every node infinitely often, then the limit of the associated product is a rank-one matrix. Moreover, the limit depends only on the *first node* of the walk. Because the derived graph is finite, there can only be finitely many different limits. The result is formulated in Theorem 3.1, and a complete characterization of these limits is provided in Sec. 3.2.

There are several implications of the above result. For example, any simple random walk on these derived graphs yields a convergent product of local stochastic matrices with probability one and the limits are independent of the sample paths but for their starting nodes (Corollary 3.2). Another consequence of the result concerns absolute probability vectors (APVs), which were introduced in [8] to study the convergence of products of stochastic matrices (We recall its definition in Definition 3.2). Generically, the sequence of APVs depends on a particular convergent product, and, moreover, takes infinitely many different values, even when only a finite number of distinct matrices appear in the product. In contrast, we show in Corollary 3.4 that one can assign to a TLG, together with a set of local stochastic matrices, a *finite* set of vectors such that the sequence of APVs attached to *any* allowable, convergent product of these local stochastic matrices takes values only from that finite set.

A large part of the novelty of this work lies in the introduction of TLGs and their derived graphs, as one may observe from the above description. To characterize their properties, we will obtain a recursive construction for them. This construction is akin to the celebrated Henneberg sequence that appears in rigidity theory [9], and we thus call it Restricted Henneberg Construction (RHC). We prove that any TLG can be obtained by an RHC and, reciprocally, any RHC yields a TLG. The proof may be of independent interest—indeed, TLGs have also appeared in earlier work on formation control [10]—but because it uses a set of ideas distinct from the ones used in the main part of the paper, we relegate it to the Appendix.

What is perhaps the closest line of work, in spirit, to the present is the work on gossiping [11–13]. A gossip can be described, in terms of message passing, as an operation in which two agents communicate their values to each other and take the average. When described in terms of stochastic matrices, this yields a matrix which is the identity save for a 2-by-2 principal submatrix whose entries are 1/2. It is shown that the left-product of such stochastic matrices converges, under some conditions, to the matrix with all entries $\frac{1}{n}$. More recently, it has been extended to clique gossiping [14], where k agents in a clique perform an averaging operation. In these works, the convergence to the averaging matrix is a by-product of the fact that the matrices involved are in fact *doubly-stochastic*, i.e., all the row sums and column sums of the matrix are one. For an application of some of the ideas developed in this work to gossip processes, see [15].

In terms of applications to consensus, besides the fact that our work allows for a control of the limiting probability vector while requiring minimal information about the allowable sequence (namely, only its starting node), it also enables the implementation of simple *secure-by-design* consensus algorithms. Indeed, small networks are by nature more secure than larger networks, since by definition they contain fewer possible points of failure or attack. The smallest meaningful network in our case is the triangle. The local stochastic matrices are so that after each iteration, each node of the triangle has to agree on the same value. Furthermore, the adjacency rule is so that the next triangle to update has *two* nodes in common with the previous triangle. Hence the third node in the triangle can verify that it receives the same value from the other two nodes. This built-in redundancy adds an obvious layer of security to the updates and complements some existing secure consensus algorithm, e.g. [16], but of course does not make them entirely impervious to tampering.

The remainder of the paper is organized as follows: we end this section by introducing key notations and terminologies used throughout the paper. In Sec. 2, we introduce the basic objects used in the paper: namely, triangulated Laman graphs, their derived graphs, and local stochastic matrices. Several key properties will be established in the section as well. Next, in Sec. 3, we state the main results of the paper, including an explicit formula for the limits of allowable convergent products. In Sec. 4, we prove the main results, save for Theorem 2.1 concerning the construction of TLGs, which we relegate to the appendix. Numerical studies are provided in Sec. 5, validating the main results and showing that they do not hold if some of the assumptions are broken. The paper ends with conclusions.

Notations and conventions. We denote by $G = (V, E)$ be a graph, with node set V and edge set E. All graphs considered in the paper are simple, i.e., there have no self-arc. We use v_i to denote a node of G. If G is *undirected*, we denote an edge by (v_i, v_j), and if G is *directed*, we denote an edge *from v_i to v_j* by v_iv_j. We refer to $|V|$ as the size of G. Given a subset of nodes $V' \subseteq V$, the subgraph of G induced by V' is defined as $G' = (V', E')$ where $E' = \{(v_i, v_j) \mid v_i, v_j \in V' \text{ and } (v_i, v_j) \in E\}$ (resp. $E' = \{v_iv_j \mid v_i, v_j \in V' \text{ and } v_iv_j \in E\}$).

We call a sequence of nodes $\gamma = v_1 \cdots v_k$ a *walk* in G if (v_i, v_{i+1}) (resp. v_iv_{i+1}) is an edge of G, for all $1 \leq i \leq k-1$. We denote by $\gamma \vee v_*$ the walk $v_1 \cdots v_kv_*$ where (v_k, v_*) (resp. v_kv_*) needs to be an edge in G for the operation to be well-defined. We denote by γ^{-1} the reverse walk $v_k \cdots v_1$.

We say that γ is a *closed walk* if γ is a walk with $\gamma_1 = \gamma_k$. We emphasize that for our purpose, a closed walk has a well-defined starting node. A path is a walk without repetition of nodes. A cycle is a closed path, i.e., only the starting node and ending node are repeated. The *length* of γ is the number of edges traversed by γ, counted with multiplicity. The *cardinality* of γ, denoted by $|\gamma|$, is the number of nodes in γ, counted with multiplicity as well.

A triangle in a graph is a cycle of length 3. We denote triangles using the letter Δ, and describe them as the sets of their constituent nodes, e.g., $\Delta = \{v_i, v_j, v_k\}$ and we can write $v_i \in \Delta$.

We denote by $\{e_1, \ldots, e_n\}$ the standard basis in $\mathbb{R}^n$. Denote by $\mathbf{1}_n$ the vector of all 1's of dimension n. We omit the index when the dimension is clear from the context. For any vector $w = (w_1, \ldots, w_n)$, we use shorthand notation $\min w := \min_{1 \leq i \leq n} w_i$. We call w a positive (resp. nonnegative) vector if each entry w_i is positive (resp. nonnegative), and w a *probability vector* if it is nonnegative and its entries sum to 1. We denote by $\operatorname{splx}(n-1)$ the standard simplex in $\mathbb{R}^n$, which is comprised of all probability vectors.

2. Triangulated Laman graphs

We present in this section a class of graphs, termed *Triangulated Laman Graphs* (TLGs), as well as a simple iterative algorithm to construct them, termed *restricted Henneberg construction* (RHC).

In order to introduce the TLGs, we first recall that a graph $G = (V, E)$ is said to be *triangulated* if for every cycle of length strictly greater than 3, there is an edge joining two nonconsecutive vertices of the cycle. We call any cycle of length 3 a *triangle*. Any edge e that belongs to only *one* triangle is called *simple*. TL graphs are also minimally rigid, see the Appendix or [9] for a formal definition. We thus include "Laman" explicitly in the definition:

Definition 2.1 *(Triangulated Laman Graphs (TLGs)).* A graph G is a **triangulated Laman Graph** ***(TLG)*** if it is both triangulated and minimally rigid.

Because all minimally rigid graphs can be obtained by a so-called Henneberg construction [9], so can be all TLGs. However, not every Henneberg construction gives rise to a TLG. We now introduce below restricted Henneberg constructions that produce TLGs and have the property that *all* TLGs can be obtained by such construction:

Initialization: Start from a graph $G_3 = (V_3, E_3)$ with $V_3 = \{v_1, v_2, v_3\}$ and $E_3 = \{(v_1, v_2), (v_2, v_3), (v_1, v_3)\}$. It consists of one triangle.

Inductive step: Suppose that a subgraph G_k of k nodes $v_1, \ldots, v_k$ has been constructed. Pick an edge $e = (v_l, v_m)$ in G_k. Add a node v_{k+1} and two edges $(v_l, v_{k+1}), (v_m, v_{k+1})$ to G_k to obtain G_{k+1}

Definition 2.2. We refer to the above construction as a **restricted Henneberg construction** ***(RHC)***.

We have the following result:

Theorem 2.1. *A graph is a TLG if and only if it can be constructed by a restricted Henneberg construction.*

A proof of the theorem is provided in the Appendix.

Derived graphs and their properties. We now introduce the notion of derived graph associated with a triangulated graph G. Roughly speaking, the derived graph is used to reflect the adjacency of triangles in G, see Fig. 1 for an illustration.

Definition 2.3 *(Derived graph).* Let G be a triangulated graph. The derived graph D_G of G is an undirected graph defined as follows: Each node Δ_i of D_G corresponds to a triangle of G. If two distinct triangles corresponding to Δ_i and Δ_j share a common edge in G, then an edge (Δ_i, Δ_j) is in D_G.

Throughout the paper, we will view Δ_i both as a node of D_G, and as a subgraph of G—more precisely, a subgraph induced by three adjacent nodes. We will write $v \in \Delta_i$ (resp. $e \in \Delta_i$) to denote that the vertex v (resp. edge e) is in the subgraph Δ_i.

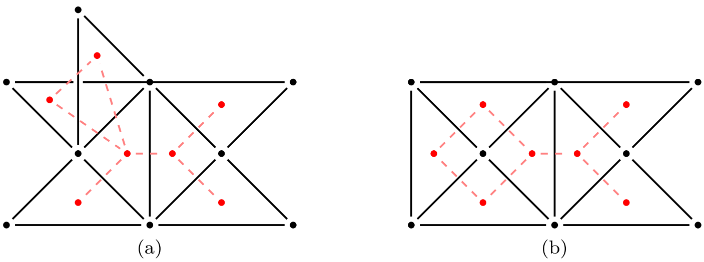

Fig. 1. Two triangulated graphs and their derived graphs (dashed). The graph on the left is not a TLG, as it contains 8 nodes and 14 edges. The graph on the right is a TLG.

We next establish a few relevant properties for the derived graphs of TLGs. We start with the following fact:

Proposition 2.2. *Let G be a TLG on n nodes. Then, there are $(n-2)$ triangles in G and the derived graph D_G is a connected, triangulated graph on $(n-2)$ nodes.*

Proof. It should be clear from the RHC that G has $(n-2)$ triangles and that D_G is connected. We show below that D_G is triangulated. The proof will be carried out by induction on the number of nodes in G. For the base case $n = 3$, G contains one triangle and D_G is comprised of a single node. This proves the base case.

For the inductive step, we assume that the statement holds for $(n-1)$ and prove it for n. Let G be a TLG on n nodes. By Theorem 2.1, it admits an RHC. Following this RHC up to step $(n-3)$ yields a TLG G' on $(n-1)$ nodes, which is a *subgraph* of G. By the induction hypothesis, the derived graph $D_{G'}$ is triangulated. We now focus on the last step of the RHC, yielding G from G'. Denote by $e = (v_l, v_m)$ the edge in G' selected, and by v_n the newly added node. Denote by $\Delta_{i_1}, \ldots, \Delta_{i_p}$ the triangles in G' that contain the edge e. Then, the subgraph of $D_{G'}$ induced by these nodes is the complete graph K_p.

The newly added triangle $\Delta_{n-2} = \{v_l, v_m, v_n\}$ is a node in D_G. It is connected in D_G to *all* the nodes $\Delta_{i_1}, \ldots, \Delta_{i_p}$. We thus conclude that the subgraph of D_G induced by $\Delta_{i_1}, \ldots, \Delta_{i_p}\Delta_n$ is a *complete graph on* $(p+1)$ *nodes*. We denote by K_{p+1} the clique. We now show that G is triangulated. By the induction hypothesis, it suffices to show that cycles of length greater than 3 containing Δ_{n-2} have a chord. To this end, observe that if Δ_{n-2} is in a cycle of length greater than 3, then Δ_{n-2} has 2 distinct neighbors in the cycle. Denote them by Δ_{i_j} and Δ_{i_k}. Then, necessarily, both of them belong to K_{p+1}. Hence, the edge $(\Delta_{i_j}, \Delta_{i_k})$ is a chord of the cycle. This completes the proof. □

In the sequel, we will require an RHC that yields a given TLG G with a particular initialization. We thus show the following:

Proposition 2.3. *Let G be an TLG on n nodes with triangles $\Delta_i, 1 \leq i \leq n-2$. Then, for any Δ_i, there exists an RHC starting with Δ_i that yields G.*

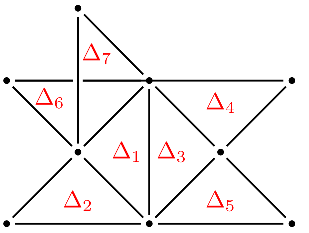

Fig. 2. The TLG from Fig. 1b is reproduced and the are triangles labeled in the order of appearance with respect to a certain RHC.

Proof. The proof will be carried out by induction on the number of nodes in G. The base case of $n = 3$ is trivially true. We thus assume that the result holds for any TLG on $(n - 1)$ nodes and prove that it holds for TLGs on n nodes.

Let G be an TLG on n nodes with $(n - 2)$ triangles. Then, there is an RHC that builds G by Theorem 2.1. Without loss of generality, we let $\Delta_{n-2} = \{v_1, v_2, v_n\}$ (resp. v_n) be the last triangle (resp. node) appearing in the RHC. Then, the degree of v_n is 2 and (v_1, v_2) is a common edge shared by Δ_{n-2} with at least one another triangle, say $\Delta_j = \{v_1, v_2, v_k\}$ for some $k \leq n - 1$.

Now, let G' be the subgraph of G induced by the nodes $v_1, \ldots, v_{n-1}$. Then, G' is constructed by stopping an RHC construction after $(n - 3)$ steps, and is thus a TLG graph on $(n - 1)$ nodes. By the induction hypothesis, for each triangle $\Delta_i \subset G'$, there exists an RHC starting with Δ_i that produces G'. Note that Δ_i is a also a triangle of G. Continuing the above RHC by one step joining node v_n to nodes v_1 and v_2 yields an RHC that builds G.

It remains to show that there is an RHC that produces G starting with triangle Δ_{n-2}. This a two-step construction: First, starting from Δ_{n-2}, we add node v_k and connect it to nodes v_1 and v_2, thus obtaining a graph with 4 nodes and 2 triangles (Δ_{n-2} and Δ_j). This graph is clearly a TLG. For the second step, we appeal again to the induction hypothesis, to obtain an RHC that builds G' starting with Δ_j. Since the concatenation of two RHCs is an RHC, using the two steps above, we have obtained an RHC that produces G from Δ_{n-2}. $\square$

Example 2.1. We illustrate here Proposition 2.3. Consider the TLG of Fig. 2. From Theorem 2.1, we know that there exists an RHC producing it. We label the triangles in the order of appearance with respect to the RHC. The proposition says that one can find an RHC yielding the same G starting from *any* Δ_i. Starting from Δ_4, a valid RHC is, e.g., $\Delta_4\Delta_3\Delta_5\Delta_1\Delta_2\Delta_6\Delta_7$.

The next few propositions shed more light on the structure of the derived graph D_G. In particular, both the triangulated and Laman character of G will come into play to show the existence of so-called bottleneck nodes in D_G (see Definition 2.4 below). These bottleneck nodes will in turn be essential ingredients in obtaining the limits of the products of local stochastic matrices.

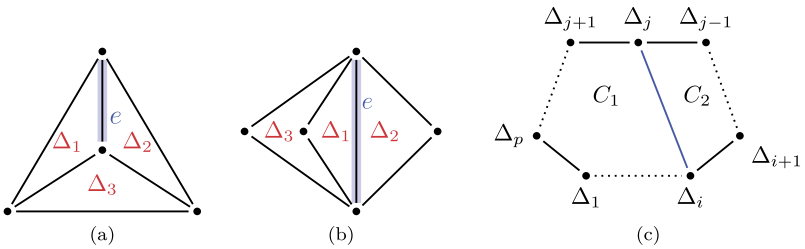

Fig. 3. The three triangles Δ_1, Δ_2, and Δ_3 in 3a and 3b are pairwise adjacent, so the corresponding derived graphs are triangles. However, the graph in 3a cannot be a subgraph of a TLG because it violates the Laman condition for minimal rigidity. Fig. 3c illustrates the two cycles C_1 and C_2 introduced in (1).

Proposition 2.4. *Let $\Delta_1\Delta_2\cdots\Delta_p\Delta_1$ be a cycle in D_G of length greater than 3. Then, all of these triangles in G share a common edge. In particular, the subgraph of D_G induced by nodes $\Delta_1,\ldots,\Delta_p$ is a complete graph.*

Proof. The proof is carried out by induction on the length p of the cycle.

For the base case $p = 3$, we first note that two distinct triangles can share at most one edge. Assume, without loss of generality, that $\Delta_i = \{v_1, v_2, v_3\}$ and $\Delta_j = \{v_1, v_2, v_4\}$, i.e., (v_1, v_2) is the edge shared by Δ_i and Δ_j. If the same edge (v_1, v_2) is also shared by Δ_k, then we are done. Suppose not, say Δ_i and Δ_k share edge (v_1, v_3); then Δ_j and Δ_k must share edge (v_1, v_4) (it cannot be (v_2, v_4) because otherwise, Δ_k has four distinct nodes $v_1, \ldots, v_4$.) But, then, the subgraph G' of G induced by $v_1, \ldots, v_4$ is K_4. The total number of edges in K_4 is 6, which violates the Laman condition, which states that the number of edges of any induced subgraph on k nodes does not exceed $(2k - 3)$. This proves the base case. See Fig. 3a and Fig. 3b for illustration.

For the inductive step we assume that the statement holds for any $p' \le p-1$ and prove for p. Since $p \ge 4$, by Proposition 2.2, there is a chord (Δ_i, Δ_j), with $1 \le i < j \le p$, in the cycle. Using this chord, we obtain the following two cycles:

$$\begin{aligned} C_1 &:= \Delta_1 \cdots \Delta_i \Delta_j \Delta_{j+1} \cdots \Delta_p \Delta_1 \\ C_2 &:= \Delta_i \Delta_{i+1} \cdots \Delta_{j-1} \Delta_j \Delta_i \end{aligned} \tag{1}$$

of lengths strictly less than p. See Fig. 3c for illustration. By the induction hypothesis, the triangles in each cycle C_k, for $k = 1, 2$, share a common edge e_k. Furthermore, note that nodes Δ_i and Δ_j appear in both C_1 and C_2 and, hence, e_1 and e_2 are shared by both Δ_i and Δ_j. If e_1 and e_2 are distinct, then $\Delta_i = \Delta_j$, which is a contradiction. We thus conclude that $e_1 = e_2 =: e$, i.e., the common edge e is shared by all of the triangles in the original cycle. □

Let G be a TLG with derived graph D_G. For a node v in G, we denote by $D_G(v)$ the subgraph of D_G induced by the triangles that contain v. See Fig. 4 for illustration.

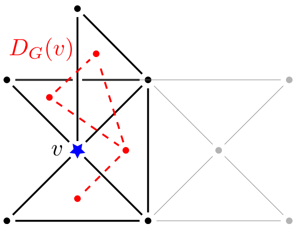

Fig. 4. We depict the triangles that contain the starred node v and the corresponding subgraph they induce in D_G, denoted by $D_G(v)$ (dashed).

Proposition 2.5. *Let G be an arbitrary TLG and v be a node of G. Then, $D_G(v)$ is a connected subgraph of D_G.*

Proof. We proceed by induction on the number of triangles in G that contain v. The base case is such that v belongs to exactly one triangle, say Δ_1, in G. The subgraph of D_G induced by Δ_1 is a single node and thus connected. This proves the base case.

For the inductive step, we assume that the statement holds for $(k-1)$ and prove it for k. Choose an RHC that builds G. Let ℓ (resp. m) be the step such that the RHC stopped right after step ℓ (resp. step m) yields a subgraph $G_\ell \subset G$ with exactly $(k-1)$ triangles containing v (resp. a subgraph $G_m \subset G$ with k triangles containing v). Since G_ℓ is a TLG, by the induction hypothesis, $D_{G_\ell}(v)$ is a connected graph on $(k-1)$ nodes. Label these nodes as $\Delta_{j_1}, \ldots, \Delta_{j_{k-1}}$. At step m, the RHC chooses an existing edge $(v, v') \in G_{m-1}$ and adds a node to form a new triangle Δ_{j_k} that contains v. Without loss of generality, we assume that Δ_{j_1} is another triangle that contains the edge (v, v'). As a consequence, $(\Delta_{j_1}, \Delta_{j_k})$ is an edge in D_{G_m} that connects Δ_{j_k} with $D_{G_\ell}(v)$. In other words, the subgraph $D_{G_m}(v)$ is connected. Finally, observe that $D_{G_m}(v)$ and $D_G(v)$ have the same node set by assumption. Since the RHC does not remove existing nodes or edges out of G (and, hence, D_G as well) along the construction process, $D_G(v)$ is connected as well. $\square$

We now introduce the notion of bottleneck nodes, see Fig. 5 for an illustration.

Definition 2.4. Let D be an undirected graph, α be a node of D, and D' be a subgraph of D. A node $\alpha^* \in D'$ is a *bottleneck in D' for α* if every walk from any node in D' to α contains α^*.

If $\alpha \in D'$, then clearly α is its own bottleneck in D', i.e., $\alpha^* = \alpha$. In most of the time, we are interested in the case where $\alpha \notin D'$. We establish below some relevant properties for bottlenecks. We start with the following one:

Lemma 2.1. *If a bottleneck $\alpha^* \in D'$ for α exists, then it is unique.*

Proof. The proof is carried out by contradiction. Suppose that there exist two distinct bottlenecks α_1^* and α_2^* in D' for α. Let $\gamma = \gamma_1 \ldots \gamma_p$ be an arbitrary finite walk with

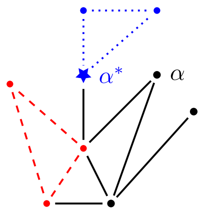

Fig. 5. Let α be a node in the graph depicted above. The subgraph depicted in blue (dotted) has a bottleneck for α, denoted by α^*. The subgraph depicted in red (dashed) does not have a bottleneck node for α.

starting node $\gamma_1 = \alpha_1^*$ and $\gamma_p = \alpha$. Since α_2^* is a bottleneck distinct from α_1^*, there exists $k_1 > 1$ such that $\gamma_{k_1} = \alpha_2^*$. But, then, $\gamma' := \gamma_{k_1} \cdots \gamma_p$ is a walk from α_2^* to α. Note that $|\gamma'| < p$. Similarly, since α_1^* is a bottleneck, there exists another integer k_2, with $k_2 > k_1$, such that $\gamma_{k_2} = \alpha_1^*$. Define $\gamma'' := \gamma_{k_2} \cdots \gamma_p$, which is a walk from α_1^* to α. By repeatedly applying the above arguments, we obtain an infinite integer sequence $k_1 < k_2 < k_3 < \cdots$, such that $\gamma_{k_{2i+1}} = \alpha_2^*$ and $\gamma_{k_{2i}} = \alpha_1^*$. However, the original walk γ has finite length, which is a contradiction. We thus have to conclude that $\alpha_1^* = \alpha_2^*$. $\square$

It should be clear that given the subgraph D' and the node α, if a bottleneck exists, then it is unique. The next proposition shows the existence of bottleneck nodes for *any* subgraph $D_G(v)$ of D_G and any node outside $D_G(v)$:

Proposition 2.6. *Let Δ_0 be a triangle and v be a node in G. Let $D_G(v)$ be the subgraph of D_G induced by triangles that contain v in G. Then, there exists a bottleneck $\Delta^* \in D_G(v)$ for Δ_0.*

Proof. If $D_G(v)$ contains one node Δ_0 or if $v \in \Delta_0$, then clearly Δ_0 is the bottleneck. Hence we assume it contains at least two nodes and, moreover, $v \notin \Delta_0$, so $\Delta_0 \notin D_G(v)$.

The remainder of the proof is carried out by contradiction. Suppose that there is no bottleneck. By Proposition 2.2, one can find two paths, γ and γ', that start with nodes in $D_G(v)$ and end at Δ_0. Moreover, by our assumption, γ and γ' can be chosen with the property that they exit the subgraph $D_G(v)$ through two distinct nodes.

More precisely, let γ_p (resp. γ'_p) the pth node in γ (resp. γ'). A node γ_p is called the exiting node of γ if $\gamma_p \in D_G(v)$ and $\gamma_q \notin D_G(v)$ for any $q > p$. Similarly, we let $\gamma'_{p'}$ be the exiting node of γ'. Then, by the hypothesis, we can find γ and γ' such that the two exiting nodes γ_p and $\gamma'_{p'}$ are distinct. For convenience, we assume, by truncating the two paths (if necessary), that the first nodes γ_1 and γ'_1 of the two paths are the exiting nodes.

We will now construct a cycle in D_G that contains nodes γ_1, γ'_1, and at least one node not in $D_G(v)$. To this end, since $D_G(v)$ is connected by Proposition 2.5, there exists a path ω in $D_G(v)$ from γ_1 to γ'_1.

Next, we let Δ be the first node that belongs to both γ and γ'. Since γ and γ' have the same ending node Δ_0, the node Δ always exists (and it could be Δ_0). By concatenating

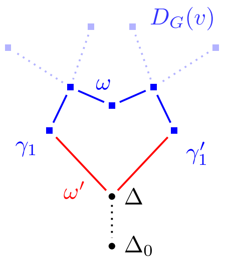

Fig. 6. This figure shows that if the bottleneck does not exist, then there would be a cycle C formed by the conjunction of two paths ω and ${\omega'}^{-1}$ as shown in the proof. The nodes in blue and square are nodes of $D_G(v)$. The node Δ is outside $D_G(v)$ but belongs to C, which leads to a contradiction.

the subpath $\gamma_1 \cdots \Delta$ of γ with the subpath $\Delta \cdots \gamma_1'$ of $(\gamma')^{-1}$, we obtain a new path ω' joining γ_1 to γ_1'.

Note, in particular, that only the starting and the ending nodes of ω' belong to $D_G(v)$. By concatenating ω with ${\omega'}^{-1}$, we obtain the desired cycle. See Fig. 6 for illustration.

Denote the cycle by C. By Proposition 2.4, the triangles of G that correspond to the nodes of C share a common edge, which we denote by e. Because both γ_1 and γ_1' belong to $D_G(v)$ and because $\gamma_1 \neq \gamma_1'$, the edge e must contain the node v. Since the triangle Δ is also a node of C, it contains the edge e and, hence, the node v. On the other hand, Δ does not belong to $D_G(v)$, which is a contradiction. $\square$

3. Main results

In the section, we state three main results concerning products of local stochastic matrices (which will be introduced below), namely Theorems 3.1, 3.6, and 3.7. Relying on properties of TLGs and their derived graphs established in the previous section, we (1) characterize the limits of these products, (2) make connections between these limits and the so-called absolute probability vectors [8], and (3) show that for any given target limit, one can find a set of local stochastic matrices so that their infinite products converge to the target one.

In the sequel, we will view D_G as a *directed graph* by replacing an undirected edge (v_i, v_j) with two directed ones, namely v_iv_j and v_jv_i. The purpose of doing so is to emphasize the direction in which an edge of D_G is traveled.

3.1. Local stochastic matrices and their infinite products

Local stochastic matrices. We now show how to attach a set of stochastic matrices A_i to a given TLG, and how the same graph can be used to generate an infinite family of products of A_i with a known limit distribution. Throughout this section, G is a TLG on n nodes.

There are $(n-2)$ triangles in G by Proposition 2.2 and we denote them by $\Delta_1, \ldots, \Delta_{n-2}$. For each $\Delta_i = \{v_j, v_k, v_l\}$, with $j < k < l$, we assign weights to its three nodes. We denote these weights by $a_{i,j}$, $a_{i,k}$, and $a_{i,l}$ respectively. We emphasize that a given node does not necessarily have a unique weight assigned, but has one weight assigned per triangle to which it belongs. On occasion, we will use a_{Δ_i,v_j}, instead of $a_{i,j}$, to denote the weight assigned to node v_j in triangle Δ_i.

We call $a_i = (a_{i,j}, a_{i,k}, a_{i,l})$ the **local weight vector** associated with Δ_i. We next define for each Δ_i a stochastic matrix as follows:

$$A_i := \sum_{v_j, v_k \in \Delta_i} a_{i,k} e_j e_k^\top + \sum_{v_j \notin \Delta_i} e_j e_j^\top . \tag{2}$$

The structure of the A_i is easy to state in words: The principal submatrix of A_i corresponding to Δ_i is a 3-by-3 rank-one stochastic matrix while the remainder of the matrix is simply the identity matrix. We illustrate this below a simple example:

Example 3.1 *(Local stochastic matrices with $n = 5$).* Consider a graph on $n = 5$ nodes consisting of the three triangles $\Delta_1 = \{v_1, v_2, v_3\}$, $\Delta_2 = \{v_1, v_2, v_4\}$, and $\Delta_3 = \{v_2, v_3, v_5\}$. In this case, we have the following local stochastic matrices:

$$A_1 = \begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,3} & 0 & 0 \\ a_{1,1} & a_{1,2} & a_{1,3} & 0 & 0 \\ a_{1,1} & a_{1,2} & a_{1,3} & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad A_2 = \begin{bmatrix} a_{2,1} & a_{2,2} & 0 & a_{2,4} & 0 \\ a_{2,1} & a_{2,2} & 0 & a_{2,4} & 0 \\ 0 & 0 & 1 & 0 & 0 \\ a_{2,1} & a_{2,2} & 0 & a_{2,4} & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix},$$

and

$$A_3 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & a_{3,2} & a_{3,3} & 0 & a_{3,5} \\ 0 & a_{3,2} & a_{3,3} & 0 & a_{3,5} \\ 0 & 0 & 0 & 1 & 0 \\ 0 & a_{3,2} & a_{3,3} & 0 & a_{3,5} \end{bmatrix}.$$

For a later purpose, we need a mild assumption on the local stochastic matrices:

Assumption 3.1. *For each triangle Δ_i and each nonsimple edge (v_j, v_k) in Δ_i, $a_{i,j} + a_{i,k} > 0$.*

Products of local stochastic matrices. We now describe allowable products of local stochastic matrices. Let γ be a walk in D_G. We say that the walk γ is *infinite* if $|\gamma| = \infty$. To any walk $\gamma := \Delta_{i_1} \Delta_{i_2} \cdots \Delta_{i_k}$ in the derived graph D_G, we associate the product of stochastic matrices

$$P_\gamma := A_{i_k} \cdots A_{i_2} A_{i_1} . \tag{3}$$

We will mostly be interested in the case of infinite walks and, in particular, determining the corresponding product P_γ.

The problem, which is twofold in nature, is well-known to be difficult. First, one has to guarantee that the infinite products exists (i.e., in the limit $|\gamma| \to \infty$). Second, provided that the limit exists, it usually depends on the complete sequence γ, making its characterization generically intractable. While the first problem has been the subject of many investigations, as mentioned in the Introduction, the second problem has received much less attention so far.

Surprisingly, under certain mild assumptions on the infinite walks (which we introduce in Definition 3.1), a complete characterization of P_γ can be obtained. We state the results below. To proceed, we first introduce the following definition:

Definition 3.1 *(Exhaustive walk).* A finite walk γ in D_G is **exhaustive** if it visits every node of D_G at least once. An infinite walk γ in D_G is **exhaustive** if it visit every node of D_G infinitely often.

With the above definition, we now state the first main result, which says that P_γ exists if the infinite walk γ is exhaustive and, moreover, there exists a *finite* set of rank one matrices to which the limit P_γ can belong.

Theorem 3.1. *Let G be a TLG on n nodes, with triangles $\Delta_1, \dots, \Delta_{n-2}$. Let $\{A_1, \dots, A_{n-2}\}$ be an arbitrary set of local stochastic matrices that satisfy Assumption 3.1. Then, there exist $(n-2)$ probability vectors $\overline{w}_1, \dots, \overline{w}_{n-2}$ such that for every infinite exhaustive walk γ with starting node Δ_i, $1 \le i \le n-2$,*

$$P_\gamma = \mathbf{1}\overline{w}_i^\top.$$

We introduce below a few corollaries of Theorem 3.1.

Randomized scheduling. An infinite exhaustive walk γ in D_G can be obtained easily by periodic extension of a finite exhaustive walk whose starting and ending nodes are adjacent. It can also be obtained via random walks as we describe below. Given a node $\Delta_i \in D_G$, we denote by $N(\Delta_i)$ the set of neighbors of Δ_i (the in-neighbors and the out-neighbors of Δ_i are the same). We call γ a *simple random walk* in D_G, if γ is an infinite walk and the transition probability $\mathbb{P}(\gamma_{t+1} = \Delta_j \mid \gamma_t = \Delta_i)$ is given by

$$\mathbb{P}(\gamma_{t+1} = \Delta_j \mid \gamma_t = \Delta_i) = \begin{cases} \frac{1}{|N(\Delta_i)|} & \text{if } \Delta_j \in N(\Delta_i), \\ 0 & \text{otherwise.} \end{cases}$$

Because D_G is connected, it is well known that a simple random walk visits every node of D_G infinitely often (and, hence, it is infinite exhaustive) with probability 1. The following fact is then an immediate consequence of Theorem 3.1:

Corollary 3.2. *Let γ be a simple random walk with starting node Δ_i. Then, $P_\gamma = \mathbf{1}\overline{w}_i^\top$ with probability one.*

Connection to absolute probability vectors. Theorem 3.1 has a few notable consequences. Let γ and P_γ be as in the theorem's statement. We use the common notation that for a pair $0 \leq s < t$ of positive integers, the partial product corresponding to the indices is

$$P_\gamma(t:s) = A_{\gamma_t} A_{\gamma_{t-2}} \cdots A_{\gamma_{s+1}}.$$

With the above notation, we can write, e.g., $P_\gamma(t:s)P_\gamma(s:r) = P_\gamma(t:r)$.

Kolmogorov introduced in [8] the **absolute probability vectors** associated with a product P_γ (see also [17,18]):

Definition 3.2 *(Absolute probability vectors).* A sequence of vectors $\{x_s\}_{s=0}^{\infty}$ are **absolute probability vectors** ***(APVs)*** for P_γ if every x_s is a probability vector and if for every pair (s,t) of integers, with $0 \leq s < t$, $x_t^\top P_\gamma(t:s) = x_s^\top$.

APVs are tightly related to the existence of the limit $\lim_{t\to\infty} P_\gamma(t:0)$. For example, it is known that the limit exists if and only if there is a *unique* set of APVs for P_γ and, moreover,

$$\lim_{t\to\infty} P_\gamma(t:s) = \mathbf{1}x_s^\top, \tag{4}$$

for any given $s \geq 0$. We refer the reader to the recent work [6] for more details on the use of the APVs (note that the author uses "absolute probability sequence" instead of APVs). As an immediate consequence of Theorem 3.1, we have the following corollary:

Corollary 3.3. *If γ is an infinite exhaustive walk, then there is a unique sequence of APVs for P_γ.*

Furthermore, a complete characterization of the values of the APVs can be obtained using Theorem 3.1:

Corollary 3.4. *Let γ be an arbitrary infinite exhaustive walk and $\{x_s\}_{s=0}^{\infty}$ be the unique sequence of APVs for P_γ. Let $\overline{w}_1, \ldots, \overline{w}_{n-2}$ be as in Theorem 3.1. Then, the image of the map $s \mapsto x_s$ is $\{\overline{w}_1, \ldots, \overline{w}_{n-2}\}$.*

Proof. For any $s \geq 0$, we consider the sequence $\gamma' := \gamma_{s+1}\gamma_{s+2}\cdots$, i.e., γ' is obtained from γ by omitting its first s nodes. If γ is exhaustive, then so is γ'. Then, by Theorem 3.1, $P_{\gamma'} = \mathbf{1}\overline{w}_{\gamma_{s+1}}$. On the other hand, by (4), we have that $P_{\gamma'} = \mathbf{1}x_s^\top$. It follows that $x_s = \overline{w}_{\gamma_{s+1}}$. This shows that the image of $s \mapsto x_s$ has finite cardinality. Finally, because γ is exhaustive, for every Δ_i, there exists an s such that $\gamma_s = \Delta_i$. □

3.2. Characterization of the product limits

Theorem 3.1 states that $\lim_{t\to\infty} P_\gamma(t:0)$ exists and can only take value in a finite set. We describe this set below.

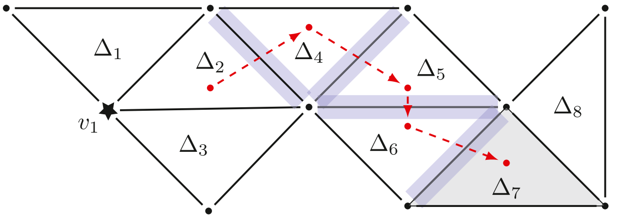

Fig. 7. We illustrate the construction of the entry $w_{i,j}$ for the case where $v_j \notin \Delta_i$. The graph G in the figure is a TLG, and the triangles are labeled in order of appearance in a valid RHC starting at Δ_1. We choose $\Delta_i = \Delta_7$ (the shaded triangle) and $v_j = v_1$ (the star node). There are three triangles Δ_1, Δ_2, and Δ_3 that contain v_1, and Δ_2 is the bottleneck for Δ_7. The directed path $\gamma = \Delta_2\Delta_4\Delta_5\Delta_6\Delta_7$ is shown in red (dashed). The path crosses four edges in G, which are highlighted in blue. Each of these is shared by two adjacent triangles along the path γ. To these edges correspond the four ratios $r_{\gamma_p,\gamma_{p+1}}$ defined in (5). The entry $w_{i,j}$ is then the product of these four ratios and the local weight a_{Δ_2,v_1} from Δ_2.

Unnormalized APVs.

For each node Δ_i in D_G, we define a positive vector $w_i \in \mathbb{R}^n$ according to the following construction. Let $w_{i,j}$ be the jth entry of w_i and recall that $a_i \in \mathbb{R}^3$ is the local weight vector of triangle Δ_i, introduced at the beginning of Section 3.1.

Let $\Delta_{j_1}, \dots, \Delta_{j_m}$ be the triangles that contain v_j. By Proposition 2.6, there is a unique Δ_{j_k}, for some $k = 1, \dots, m$, such that Δ_{j_k} is the bottleneck in $D_G(v_j)$ for Δ_i. We have two cases depending on whether $v_j \in \Delta_i$ or $v_j \notin \Delta_i$:

Case 1: $v_j \in \Delta_i$**:** In this case, $\Delta_{j_k} = \Delta_i$. We then set $w_{i,j} := a_{i,j}$.

Case 2: $v_j \notin \Delta_i$**:** In this case, we let γ be a finite walk in D_G from Δ_{j_k} to Δ_i. Let γ_p be the pth node in the walk, with $\gamma_1 = \Delta_{j_k}$. For each $p = 1, \dots, |\gamma| - 1$, we let $e_p = (v_{p_1}, v_{p_2})$ be the unique edge in G that is shared by triangles γ_p and γ_{p+1}. Define the ratio

$$r_{\gamma_p,\gamma_{p+1}} := \frac{a_{\gamma_{p+1},v_{p_1}} + a_{\gamma_{p+1},v_{p_2}}}{a_{\gamma_p,v_{p_1}} + a_{\gamma_p,v_{p_2}}}. \tag{5}$$

Now, define a product of the above ratios along the path γ as follows:

$$R_\gamma := \prod_{p=1}^{|\gamma|-1} r_{\gamma_p,\gamma_{p+1}}.$$

Apparently, R_γ depends on γ. However, we will show in Proposition 3.5 that the choice of the walk does not affect R_γ, as long as the starting and ending nodes are fixed.

With R_γ defined above, we set

$$w_{i,j} := a_{\gamma_1,j} R_\gamma, \tag{6}$$

where $a_{\gamma_1,j}$ is the weight of node v_j for $\gamma_1 = \Delta_{j_k}$. See Fig. 7 for illustration.

In words, $r_{\gamma_p,\gamma_{p+1}}$ is a ratio of the sums of the local weights assigned to the nodes incident to e_p: in the numerator, we take the weights from γ_p, and in the denominator, the weights from the next triangle in γ, namely, γ_{p+1}. We emphasize that the *order* of the two subindices of r matters and it reflects the orientation of the edge $\gamma_p\gamma_{p+1}$ in γ. By Assumption 3.1, the denominator in the ratio is strictly positive.

Note that the above two cases can actually be unified if one extends the definition of R_γ to allow for γ a path of cardinality 1 (i.e., a path comprised of a single node). Specifically, we set $R_\gamma = 1$ for any such path. Then, we can express $w_{i,j}$, for $v_j \in \Delta_i$, as $w_{i,j} = a_{i,j} R_{\Delta_i} = a_{i,j}$.

We now prove the above claim about the *independence* of R_γ on the particular walk chosen. It is a consequence of the following fact:

Proposition 3.5. *Let ω be a closed walk in D_G. Then, $R_\omega = 1$.*

Proof. We first assume that ω is a cycle and write $\omega = \omega_1 \cdots \omega_k \omega_1$. By Proposition 2.4, all the triangles $\omega_1, \ldots, \omega_k$ share a common edge in G, which we denote by (v_1, v_2). Clearly, (v_1, v_2) is also the only edge shared by two distinct ω_i and ω_j. It then follows that

$$\begin{aligned} R_\omega &= r_{\omega_1,\omega_2} r_{\omega_2,\omega_3} \cdots r_{\omega_k,\omega_1} \\ &= \frac{a_{\omega_2,v_1} + a_{\omega_2,v_2}}{a_{\omega_1,v_1} + a_{\omega_1,v_2}} \frac{a_{\omega_3,v_1} + a_{\omega_3,v_2}}{a_{\omega_2,v_1} + a_{\omega_2,v_2}} \cdots \frac{a_{\omega_1,v_1} + a_{\omega_1,v_2}}{a_{\omega_k,v_1} + a_{\omega_k,v_2}} \\ &= 1. \end{aligned}$$

We now assume that ω is a closed walk. One can always decompose ω edge-wise and use the directed edges in ω to form multiple cycles. Label these cycles as $C_1, \ldots, C_m$. Then, it should be clear that $R_\omega = R_{C_1} \cdots R_{C_m}$. Because $R_{C_i} = 1$ for each $i = 1, \ldots, m$, we have that $R_\omega = 1$. □

Returning to the independence of R_γ on a particular path, assume that γ and γ' are two walks in D_G with the same starting and ending nodes. Then, by concatenating γ with γ'^{-1}, we obtain a closed walk ω. By Proposition 3.5, $R_\omega = R_\gamma R_{\gamma'^{-1}} = 1$. On the other hand, $R_{\gamma'^{-1}} = 1/R_{\gamma'}$. It then follows that $R_\gamma = R_{\gamma'}$, as is claimed above.

Example 3.2 *(Unnormalized APVs in case $n = 5$).* We return to the case $n = 5$ illustrated in Example 3.1, with three triangles Δ_1, Δ_2, and Δ_3. We denote the associated local weight vectors a_1, a_2, and a_3. Then, the vectors w_1, w_2, and w_3 obtained according to the above construction are

$$w_1 = \begin{pmatrix} a_{1,1} & a_{1,2} & a_{1,3} & a_{2,4}\frac{a_{1,1}+a_{1,2}}{a_{2,1}+a_{2,2}} & a_{3,5}\frac{a_{1,2}+a_{1,3}}{a_{3,2}+a_{3,3}} \end{pmatrix}^\top$$

$$w_2 = \begin{pmatrix} a_{2,1} & a_{2,2} & a_{1,3}\frac{a_{2,1}+a_{2,2}}{a_{1,1}+a_{1,2}} & a_{2,4} & a_{3,5}\frac{a_{1,2}+a_{1,3}}{a_{3,2}+a_{3,3}}\frac{a_{2,1}+a_{2,2}}{a_{1,1}+a_{1,2}} \end{pmatrix}^\top$$

$$w_3 = \begin{pmatrix} a_{1,1}\frac{a_{3,2}+a_{3,3}}{a_{1,2}+a_{1,3}} & a_{3,2} & a_{3,3} & a_{2,4}\frac{a_{1,1}+a_{1,2}}{a_{2,1}+a_{2,2}}\frac{a_{3,2}+a_{3,3}}{a_{1,2}+a_{1,3}} & a_{3,5} \end{pmatrix}^\top \quad \square$$

It should be clear that each vector w_i defined above is nonnegative and nonzero (the entries $w_{i,j}$, for $v_j \in \Delta_i$, defined in case 1 cannot all be zero because they sum to one). However, each w_i is not a probability vector because the sum of its entries is in general greater than one. The following theorem explains why these vectors are called unnormalized APVs:

Theorem 3.6. *Let $w_1, \dots, w_{n-2}$ be the positive vectors defined above. Then, the vectors $\overline{w}_1, \dots, \overline{w}_{n-2}$ in Theorem 3.1 are given by normalization:*

$$\overline{w}_i := \frac{w_i}{\sum_{j=1}^n w_{i,j}}.$$

3.3. Surjectivity of left-eigenvector map

The remainder of the section concerns the design of stochastic matrices that yield a desired rank one limit for the product P_γ. We have the following theorem:

Theorem 3.7. *Let $\overline{w}^*$ be a positive probability vector. Then, for any given Δ_i, there exist local weight vectors $a_j \in \mathbb{R}^3$, for $\Delta_j \in D_G$, such that $\overline{w}_i = \overline{w}^*$.*

The statement is not surprising when one compares the dimensions of the local weight vectors (totaling $2(n-2)$) with the dimension of a probability vector (which is $(n-1)$). Nevertheless, the proof we provide is constructive. Since the proof relies on a different set of arguments from the ones needed for the previous theorems and since it is relatively simpler, we provide it here.

Proof. We proceed by induction on the number of nodes in the graph G. For the base case $n = 3$, we follow Theorem 3.6 and set $a = \overline{w}^*$.

Now, let us assume that the statement holds for all TLGs G' on n nodes. Given a TLG G on $(n+1)$ nodes it can be obtained from a TLG graph G' on n nodes by performing one step of RHC. By Proposition 2.3, we can assume that the subgraph G' contains Δ_i (specifically, by Proposition 2.3, we can choose an RHC that starts with Δ_i). We denote by v_{n+1} the newly added node going from G' to G.

Let $\overline{w}^* = (\overline{w}_1^*, \dots, \overline{w}_{n+1}^*) \in \mathbb{R}^{n+1}$ be an arbitrary positive probability vector. By the induction hypothesis, we can choose local weights a_j, $1 \le j \le n-2$, so that the unnormalized vector $(w_{i,1}, \dots, w_{i,n})$ satisfies

$$(w_{i,1}, \dots, w_{i,n}) \propto \frac{(\overline{w}_1^*, \dots, \overline{w}_n^*)}{\sum_{i=1}^n \overline{w}_i^*},$$

i.e., the right hand side is realized as the probability vector for G'.

Let γ be a path from Δ_{n+1} to Δ_i. Without loss of generality, we can assume that Δ_n is the second node in the path, so Δ_n and Δ_{n+1} are adjacent, and we can label the nodes of G so that (v_1, v_2) is the common edge shared by Δ_n and Δ_{n+1}. Let γ' be a subpath of γ that starts from Δ_n and ends at Δ_i. Then, from the construction of w_i in Eq. (6), we have that

$$\begin{aligned} w_{i,n+1} &= a_{n+1,n+1} R_\gamma \\ &= a_{n+1,n+1} \frac{a_{n,1} + a_{n,2}}{a_{n+1,1} + a_{n+1,2}} R_{\gamma'} \\ &= \frac{a_{n+1,n+1}}{a_{n+1,1} + a_{n+1,2}} (a_{n,1} + a_{n,2}) R_{\gamma'}, \end{aligned} \tag{7}$$

where the second equality follows by unwrapping the definition of r_{γ_1,γ_2} and the third equality is just a rearrangement. Now choose the three entries $a_{n+1,1}$, $a_{n+1,2}$, and $a_{n+1,n+1}$ of the local weight vector a_{n+1} so that the last entry $w_{i,n+1}$ of w_i satisfies

$$w_{i,n+1} = \frac{w_{i,n}}{\overline{w}^*_n} \overline{w}^*_{n+1}.$$

This can done since $(a_{n,1} + a_{n,2}) R_{\gamma'}$ is independent of the local weight vector a_{n+1}. By normalizing w_i, we obtain $\overline{w}_i = \overline{w}^*$. □

4. Proof of main results

This section is devoted to the proofs of the main theorems stated in Sec. 3.

4.1. Properties of unnormalized APVs

We start by deriving some key properties of the unnormalized APVs w_i of the products P_γ introduced above.

Proposition 4.1. *Let w_i be defined in Sec. 3.2 and A_i be the corresponding local stochastic matrix. Then, $w_i^\top A_i = w_i^\top$ for any $i = 1, \ldots, n-2$.*

Proof. The result is a direct computation. Assume without loss of generality that Δ_i is comprised of the vertices v_1, v_2, and v_3. Then, the matrix A_i takes the form $A_i = \operatorname{diag}(\mathbf{1}_3 a_i^\top, I_{n-3})$ where $a_i := (a_{i,1}, a_{i,2}, a_{i,3})$. It thus suffices to show that

$$(w_{i,1}, w_{i,2}, w_{i,3})^\top \mathbf{1}_3 a_i^\top = (w_{i,1}, w_{i,2}, w_{i,3})^\top,$$

which follows from the fact that $(w_{i,1}, w_{i,2}, w_{i,3}) = a_i$ by construction of the vector w_i. □

The following Proposition is a major building block in the proofs of the main theorems.

Proposition 4.2. *Let Δ_i and Δ_j be any two adjacent triangles in G. Then,*

$$w_j^\top A_i = r_{i,j} w_i^\top,$$

where $r_{i,j}$ is defined in Eq. (5).

Proof. We prove the proposition by showing that $w_{j,k} = r_{i,j} w_{i,k}$ for all $k = 1, \ldots, n$. To do so, we first relabel the vertices (if necessary) so that Δ_i (resp. Δ_j) is comprised of vertices v_1, v_2, and v_3 (resp. v_1, v_2, and v_4). Then, we have that $A_i = \operatorname{diag}(\mathbf{1}_3 a_i^\top, I_{n-3})$ with $a_i = (a_{i,1}, a_{i,2}, a_{i,3})$. We consider below two cases for the subindex k:

Case 1: $k = 1, 2, 3$. In this case, it suffices to show that

$$(w_{j,1}, w_{j,2}, w_{j,3})^\top \mathbf{1}_3 a_i^\top = r_{i,j}(w_{i,1}, w_{i,2}, w_{i,3})^\top.$$

First, it should be clear that

$$(w_{i,1}, w_{i,2}, w_{i,3}) = (a_{i,1}, a_{i,2}, a_{i,3}).$$

Next, note that Δ_i is the bottleneck of $D_G(v_3)$ for Δ_j because Δ_i and Δ_j are adjacent. The construction of w_j described in Sec. 3 yields that

$$(w_{j,1}, w_{j,2}, w_{j,3}) = (a_{j,1}, a_{j,2}, a_{i,3} r_{i,j}).$$

It then follows that

$$(w_{j,1}, w_{j,2}, w_{j,3})^\top \mathbf{1}_3 = a_{j,1} + a_{j,2} + a_{i,3} \frac{a_{j,1} + a_{j,2}}{a_{i,1} + a_{i,2}} = \frac{a_{j,1} + a_{j,2}}{a_{i,1} + a_{i,2}} = r_{i,j},$$

where we have used the fact that $a_{i,1} + a_{i,2} + a_{i,3} = 1$.

Case 2: $k = 4, \ldots, n$. It suffices to show that $w_{j,4} = r_{i,j} w_{i,4}$ (because the principal submatrix of A formed by the last $(n-3)$ rows/columns is the identity matrix). From Proposition 2.6, there is a unique bottleneck Δ^* in $D_G(v_k)$ for Δ_i. Because Δ_i and Δ_j are adjacent, the same node Δ^* is also the bottleneck in $D_G(v_k)$ for Δ_j. To see this, let γ be an arbitrary walk from a node in $D_G(v_k)$ to Δ_j. Then, $\gamma' := \gamma \vee \Delta_i$ is a walk from the same node in $D_G(v_k)$ to Δ_i. Since Δ^* is the bottleneck for Δ_i, it is necessarily contained in γ' and, hence, γ. Thus, all walks from $D_G(v_k)$ to Δ_j contain Δ^*. Since by Lemma 2.1 the bottleneck is unique, it has be to Δ^*.

Now, let γ be a walk from the bottleneck Δ^* to Δ_j and $\gamma' := \gamma \vee \Delta_i$ as above. Then, $R_{\gamma'} = r_{j,i} R_\gamma$. On the other hand, we have $w_{i,k} = a_{\Delta^*, v_k} R_{\gamma'}$ and $w_{j,k} = a_{\Delta^*, v_k} R_\gamma$. It follows that $w_{i,k} = r_{j,i} w_{j,k}$. Because $r_{i,j} r_{j,i} = 1$, we have $w_{j,k} = r_{i,j} w_{i,k}$.

The proof is now complete. □

The above proposition has important implications as we state below:

Corollary 4.3. *Let γ be a finite walk in D_G that starts at Δ_i and ends at either Δ_i or at a node adjacent to Δ_i. Then, $w_i^\top P_\gamma = w_i^\top$.*

Proof. Let Δ_j be an arbitrary node *adjacent* to Δ_i, and $\gamma = \gamma_1 \cdots \gamma_k$ be a finite walk in D_G, with $\gamma_1 = \Delta_i$ and $\gamma_k = \Delta_j$. By adding a node Δ_i to the end of γ, one obtains the closed walk $\gamma' := \gamma \vee \Delta_i$. We show that the statement holds for both γ and γ'.

For γ, we repeatedly apply Proposition 4.2 to obtain that

$$\begin{aligned} w_i^\top P_\gamma &= w_i^\top A_{\gamma_k} \cdots A_{\gamma_1} \\ &= r_{\gamma_k,\gamma_1} w_{\gamma_k} A_{\gamma_{k-1}} \cdots A_{\gamma_1} \\ &\vdots \\ &= r_{\gamma_k,\gamma_1} r_{\gamma_{k-1},\gamma_k} \cdots r_{\gamma_1,\gamma_2} w_i^\top \\ &= R_{\gamma'} w_i^\top. \end{aligned} \tag{8}$$

Because γ' is a closed walk, $R_{\gamma'} = 1$ by Proposition 3.5.

Next, for γ', we note that by Proposition 4.1, $w_i^\top A_{\gamma_1} = w_i^\top A_i = w_i^\top$, so

$$w_i^\top P_{\gamma'} = w_i^\top A_{\gamma_1} A_{\gamma_k} \cdots A_{\gamma_1} = w_i^\top A_{\gamma_k} \cdots A_{\gamma_1} = w_i^\top P_\gamma.$$

It is shown above that $w_i^\top P_\gamma = w_i^\top$. This completes the proof. □

4.2. Exhaustive walks and contraction property

We develop in this section the necessary tools to show that the limit P_γ exists when γ is an infinite exhaustive walk in D_G.

Contraction property. We start by introducing the following semi-norm [1]: Let $A = [a_{ij}]$ be an arbitrary $n \times m$ matrix. We set

$$\|A\|_S := \max_{1 \le j \le m} \max_{1 \le i_1, i_2 \le n} |a_{i_1 j} - a_{i_2 j}| \tag{9}$$

It is known [2, Theorem 1] that if the matrices in the product P_γ are stochastic matrices and if $\lim_{t\to\infty} \|P_\gamma(t:0)\|_S = 0$, then there exists a probability vector w so that $\lim_{t\to\infty} P_\gamma(t:0) = \mathbf{1}_n w^\top$.

We introduce below another known fact:

Lemma 4.1. *Let $A \in \mathbb{R}^{n\times n}$ be a stochastic matrix. Suppose that A has a positive column z with $\min z \ge \epsilon > 0$; then, for any matrix $B \in \mathbb{R}^{n\times m}$,*

$$\|AB\|_S < (1-\epsilon)\|B\|_S.$$

We do not provide a proof here, but refer to the proof of a similar statement in [19, Lemma 3]. There, the author also assumes that B is a stochastic matrix. However, the proof provided does not rely on this assumption.

Exhaustive walks and products with positive columns. We show that if G is a TLG with derived graph D_G and if γ is a finite exhaustive walk in D_G, then P_γ has a positive column. For that, we first have the following fact:

Lemma 4.2. *Let A be a stochastic matrix and z be a nonnegative vector. Then,* $\min(Az) \geq \min z$.

Proof. Since A is row stochastic, every entry of the vector Az is a convex combination of the entries of z, and thus larger than $\min z$. □

We now establish the following result, as announced above:

Proposition 4.4. *There is an $\epsilon \in (0,1)$ such that for any finite, exhaustive walk γ in D_G, the stochastic matrix P_γ has a positive column z with* $\min z \geq \epsilon$.

Proof. Let γ be an arbitrary finite, exhaustive walk. Let μ_p be the number of distinct nodes in $\gamma_1 \cdots \gamma_p$. Then

$$1 = \mu_1 \leq \mu_2 \leq \cdots \leq \mu_{|\gamma|} = n-2.$$

Because μ_p is an integer-valued increasing sequence and because $\mu_{p+1} - \mu_p \leq 1$ by construction, there exist $(n-2)$ time steps $t_1 < t_2 < \cdots < t_{n-2} := |\gamma|$ with the property that $\mu_{t_k+1} - \mu_{t_k} = 1$, for $1 \leq k \leq n-3$. Note that t_1 has to be 1 since D_G is simple and γ is a walk in D_G.

For ease of analysis, but without loss of generality, we label the nodes of G such that the triangles $\gamma_1, \ldots, \gamma_{t_k}$, for all $k = 1, \ldots, (n-2)$, cover nodes $v_1, \ldots, v_{k+2}$. Then, by the definition (2) of A_i, the matrix

$$P_\gamma(t_k : 0) = A_{\gamma_{t_k}} \cdots A_{\gamma_1} =: \operatorname{diag}[Q_k, I], \tag{10}$$

for $k = 1, \ldots, n-2$, is block diagonal with Q_k being an $(k+2) \times (k+2)$ stochastic matrix and I being the identity matrix of dimension $(n-k-2) \times (n-k-2)$.

We show below that for every $k = 1, \ldots, n-2$, there exists an $\epsilon_k > 0$, independent of γ, such that Q_k has a column z_k with $\min z_k \geq \epsilon_k$. The proof is carried out by induction on k. To proceed, we first define the minimum *non-zero* entry over all local weight vectors of triangles:

$$\underline{a} := \min \left\{ a_{\Delta, v_i} \mid a_{\Delta, v_i} \neq 0, v_i \in \Delta, \Delta \in D_G \right\}. \tag{11}$$

It should be clear that $\underline{a} \in (0, 1]$.

For the base case $k = 1$, we have that $t_1 = 1$ and $P_\gamma(t_1 : 0) = A_{\gamma_1} = \text{diag}[Q_1, I]$. It follows that $Q_1 = (\mathbf{1}_3 a_1^\top) = \mathbf{1}_3 a_1^\top$, where $a_1 \in \mathbb{R}^3$ is the local weight vector of triangle Δ_1. It should be clear that Q_1 has a positive column, which we denote by z_1. Setting $\epsilon_1 := \underline{a}$, we obtain that $\min z_1 \geq \epsilon_1$.

For the inductive step, we assume that ϵ_{k-1} exists and prove the existence of ϵ_k. By (10), we have that

$$P_\gamma(t_{k-1} : 0) = \text{diag}[Q_{k-1}, I],$$

where $\dim Q_{k-1}$ is $(k+1) \times (k+1)$. Now, consider the sub-walk $\gamma_1 \cdots \gamma_{t_k}$ and the corresponding product $P_\gamma(t_k : 0)$. We can express $P_\gamma(t_{k-1} + 1 : 0) = \text{diag}[Q'_k, I]$, where Q'_k is of dimension $(k+2) \times (k+2)$.

By an earlier assumption, triangle $\gamma_{t_{k-1}+1}$ contains the node v_{k+2}. Moreover, by relabeling nodes $v_1, \ldots, v_{k+1}$, we can arrange matters so that $\gamma_{t_{k-1}+1}$ has nodes $\{v_k, v_{k+1}, v_{k+2}\}$. Denote by $\tilde{Q}_{k-1}$ the matrix Q_{k-1} after the relabeling, then $\tilde{Q}_{k-1}$ is obtained by a permutation of rows/columns of Q_{k-1}. It should be clear that if Q_{k-1} has a column z_{k-1} such that $\min z_{k-1} \geq \epsilon$ for some $\epsilon > 0$, then so does $\tilde{Q}_{k-1}$. For ease of notation, we will still write Q_{k-1} instead of $\tilde{Q}_{k-1}$. Using this labeling, we can compute Q'_k from Q_{k-1} via the following expression:

$$Q'_k = \text{diag}[I, \mathbf{1}_3 a^\top] \, \text{diag}[Q_{k-1}, 1],$$

where $a \in \mathbb{R}^3$ is the local weight vector corresponding to the triangle $\gamma_{t_{k-1}+1}$ (the sub-index of a has been omitted for simplicity).

By the induction hypothesis, there is a positive column z_{k-1} of Q_{k-1} such that $\min z_{k-1} \geq \epsilon_{k-1}$. Then, $[z_{k-1}; 0]$ is a column of $\text{diag}[Q_{k-1}, 1]$. Let $z_{k-1,j}$ (resp. a_j) be the jth entry of z_{k-1} (resp. a). We compute below the column vector $z'_k := \text{diag}[I, \mathbf{1}_3 a^\top][z_{k-1}; 0]$:

$$z'_{k,j} = \begin{cases} z_{k-1,j} & \text{for } 1 \leq j \leq k-1, \\ a_1 z_{k-1,k} + a_2 z_{k-1,k+1} & \text{for } k \leq j \leq k+2. \end{cases} \tag{12}$$

From Assumption 3.1, a_1 and a_2 cannot both be zero, so Q'_k has a positive column z'_k. Moreover, because $\min z_{k-1} \geq \epsilon_{k-1}$ and $a_1, a_2 \geq \underline{a}$ (note that $\underline{a} < 1$),

$$\min z'_k \geq \underline{a}\epsilon_{k-1} =: \epsilon'_k \in (0, 1).$$

Note that ϵ'_k is independent of γ and, hence, a uniform lower bound.

By the definition of t_k, the sub-walk $\gamma_1 \cdots \gamma_{t_k}$ covers the same set of nodes as the sub-walk $\gamma_1 \cdots \gamma_{t_{k-1}+1}$ does. In particular, using Eq. (10), the matrix $P_\gamma(t_k : 0)$ takes the form $\text{diag}[Q_k, I]$ with Q_k and Q'_k of the same dimension. Moreover, Q_k can be expressed as $Q_k = S_k Q'_k$ where S_k is the $(k+2) \times (k+2)$ leading principal submatrix of

$P_\gamma(t_k : t_{k-1}+1)$. Note that S_k is a stochastic matrix. By Lemma 4.2, $\min(S_k z'_k) \geq \min z'_k$, which implies that Q_k has a positive column z_k and, moreover, $\min z_k \geq \epsilon_k := \epsilon'_k$. □

Remark 4.1. A close inspection of the above arguments yields that ϵ can be chosen to be $\epsilon = \underline{a}^{n-2}$.

4.3. Proofs of main theorems

We now put the above results together to prove Theorems 3.1 and 3.6.

Let γ be an infinite, exhaustive walk starting at node Δ_i. Let t_k, for $k \geq 0$, be a monotonically increasing sequence of time steps at which the walk first revisits Δ_i after having visited all other nodes since t_{k-1}, i.e., $\gamma_{t_k} = \Delta_i$ and $\gamma_{t_k} \cdots \gamma_{t_{k+1}-1}$ is a finite, exhaustive walk. We set $t_0 := 1$. By assumption, $(n-2) \leq t_{k+1} - t_k < \infty$. For convenience, we introduce $s_k := t_k - 1$, so $s_0 = 0$.

From Lemma 4.1, the semi-norm $\|P_\gamma(s : 0)\|_S$ is *non-increasing* in s and since it is obviously lower-bounded, it converges to a limit. We now show that this limit is 0. To this end, we use Proposition 4.4 to obtain an $\epsilon \in (0, 1)$ such that for each $k \geq 0$, $P_\gamma(s_{k+1} : s_k)$ has a positive column z_k with $\min z_k \geq \epsilon$. Then, by repeatedly applying Lemma 4.1, we have that

$$\begin{aligned}\|P_\gamma(s_k : s_0)\|_S &= \|P_\gamma(s_k : s_{k-1}) P_\gamma(s_{k-1} : s_0)\|_S \\ &\leq (1-\epsilon)\|P_\gamma(s_{k-1} : s_0)\|_S \leq (1-\epsilon)^k. \quad (13)\end{aligned}$$

It then follows that

$$\|P_\gamma\|_S = \lim_{k\to\infty} \|P_\gamma(s_k : s_0)\| = \lim_{k\to\infty} (1-\epsilon)^k = 0,$$

as claimed. Using [2, Theorem 1], we conclude that P_γ converges to a rank-one stochastic matrix.

Let $\hat{w}$ be such that $P_\gamma = \mathbf{1}\hat{w}^\top$. We next show that $\hat{w} = \overline{w}_i$ as is in the statement of Theorem 3.6. On the one hand, the convergence of $P_\gamma(s_k : s_0)$ to P_γ implies the convergence of the row vectors $w_i^\top P_\gamma(s_k : s_0)$ to the row vector $w_i^\top P_\gamma = w_i^\top \mathbf{1}\hat{w}^\top = \hat{w}^\top$. On the other hand, each finite walk $\gamma_{t_0} \cdots \gamma_{t_k - 1}$ starts with γ_i and ends with a node *adjacent* to γ_i. Thus, by Corollary 4.3, $w_i^\top P_\gamma(s_k : s_0) = w_i^\top$. Combining the above arguments, we have that

$$w_i^\top = \lim_{k\to\infty} w_i^\top P_\gamma(s_k : s_0) = w_i^\top P_\gamma = \hat{w}^\top.$$

This completes the proofs of both Theorems 3.1 and 3.6. □

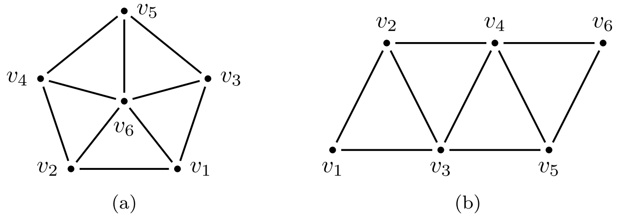

Fig. 8. The graph in Fig. 8a is a triangulated rigid graph, but is not Laman. The graph in Fig. 8b is a TLG.

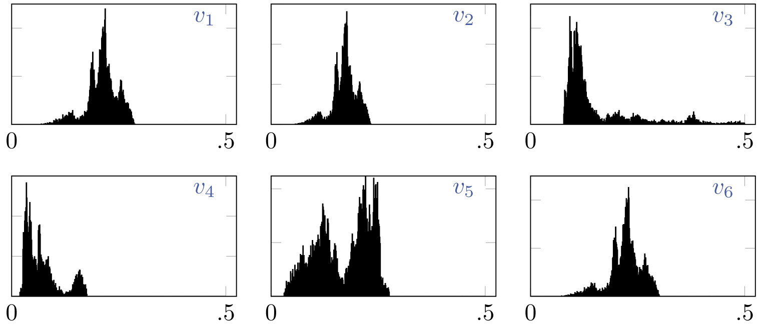

Fig. 9. For random walks γ, with $|\gamma| = 10^4$, in the derived graph of Fig. 8a, we evaluate the corresponding left eigenvector w of P_γ. We plot the empirical distribution for each entry of w.

5. Numerical studies

We present here simulation results showing the validity of the main theorems and the importance of the adjacency rule and the condition that the underlying graph is a TLG. Precisely, we present three sets of experiments. In the first one, we show that if G is not a TLG, but simply a triangulated graph, then the conclusions do not hold. In the second set, we explore the importance of the structure of the local stochastic matrices. In the last set, we show that the adjacency rule for the product, i.e., that γ is a walk in D_G, is critical as well.

Experiment 1: On triangulated Laman graphs. We consider in Fig. 8a a triangulated graph with 5 triangles. The derived graph is a cycle of length 5. The graph is not Laman because it violates the Laman condition, but it is rigid. To each triangle, we assign a local weight vector, which yields the associated local stochastic matrix. These local weight vectors are *i.i.d* random variables uniformly drawn from splx(2).

We then sample $N_r = 50 \cdot 10^3$ random walks γ in D_G. The cardinality of each walk is $N = 10^4$. This length was sufficient to guarantee converge of the product P_γ to a rank-one matrix, as was observed in the simulation (we took the absolute values of the eigenvalues of the products and verified that there was only one nonzero value, namely value one).

We denote by w the left eigenvector of P_γ corresponding to eigenvalue 1, i.e., $P_\gamma = \mathbf{1}(w)^\top$. Then, w is a random variable taking value in splx(5). We plot in Fig. 9 the

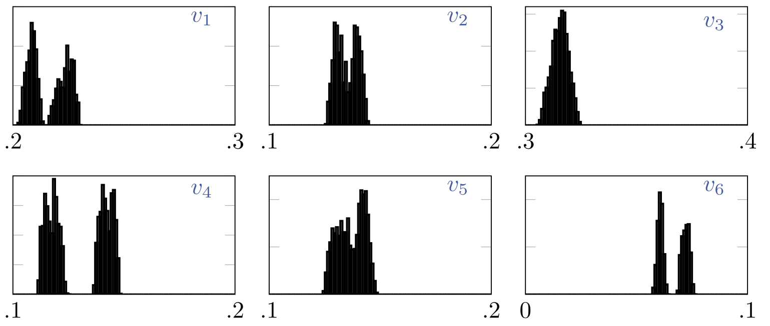

Fig. 10. For random walks γ, with $|\gamma| = 10^4$, in the derived graph of Fig. 8b, we evaluate the corresponding left eigenvector w of P_γ. We plot the empirical distribution for each entry of w.

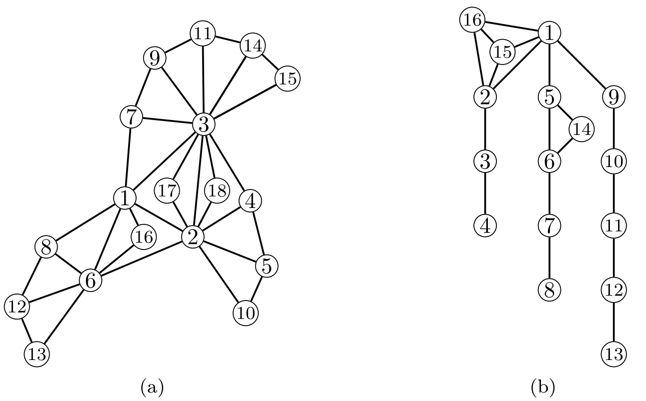

Fig. 11. The "horse" graph G in Fig. 11a is a TLG on 18 nodes. We plot its derived graph D_G in Fig. 11b. We provide here a few (but not all) correspondences between triangles in G and nodes in D_G: $\Delta_1 = \{v_1, v_2, v_3\}$, $\Delta_2 = \{v_2, v_3, v_4\}$, $\Delta_5 = \{v_1, v_2, v_6\}$, and $\Delta_9 = \{v_1, v_3, v_7\}$.

histograms for the 6 entries of w. We observe that the support of these empirical distributions has non-zero measure, indicating that there is a continuum of limits, associated to the chosen set of local weight vectors.

Experiment 2: On local stochastic matrices. We consider in Fig. 8b a TLG with 4 triangles whose derived graph is line graph. To each triangle Δ_i, we assign a random 3×3 stochastic matrix, realized as the principal submatrix of A_i corresponding to the nodes of Δ_i. All of the row vectors of these 3×3 matrices are *i.i.d* random variables uniformly drawn from splx(2). This construction of local weight matrices violates our assumption on the A_i's, which requires these principal submatrices to have identical rows.

Similarly, we sample $N_r = 50 \cdot 10^3$ random walks γ of cardinality $N = 10^4$ in D_G (for which we observe convergence of every P_γ) and plot the empirical distribution of the entries of the limiting left-eigenvector of P_γ (Fig. 10). We again observe that the support of these empirical distributions have non-zero measures, indicating that there is a continuum of limits.

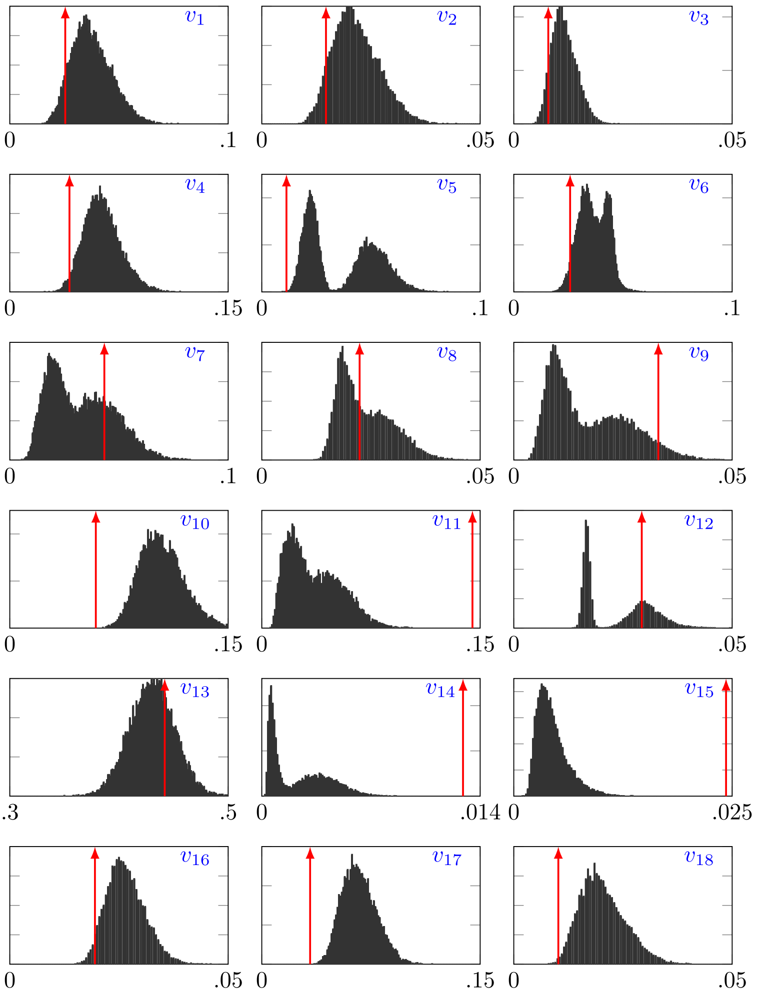

Fig. 12. For random sequences (not necessarily walks) γ, with $|\gamma| = 30 \cdot 10^3$, of nodes in the derived graph shown in Fig. 11b, we evaluate the corresponding left eigenvector w of P_γ. We plot the empirical distribution for each entry of w. The delta functions in red are distributions that correspond to random walks γ in D_G. The distributions in black correspond to sequences γ of nodes in D_G chosen uniformly at random.

Experiment 3: On adjacency rules. In this last set of experiments, we verify that if the assumptions are met, P_γ converges to a rank-one matrix whose value does not depend on the walk γ in D_G. We also verify that if all assumptions are met but γ is a random sequence of triangles, and thus does not respect the adjacency rules afforded by D_G, the conclusions do not hold. The simulations shown in Fig. 12 are based on a larger TLG G with 18 nodes, depicted in Fig. 11a, with derived graph D_G on 16 nodes shown in Fig. 11b. To each triangle, we assign a local weight vector, which yields the associated local stochastic matrix. These local weight vectors are *i.i.d* random variables uniformly drawn from splx(2).

We first sampled $N_r = 50 \cdot 10^3$ random walks γ of cardinality $N = 30 \cdot 10^3$ in D_G. Every walk starts at node $\Delta_1 = \{v_1, v_2, v_3\}$. We observe convergence of every P_γ to a common rank-one matrix. This provides numerical support of Theorem 3.1. Denote by w the common left-eigenvector of P_γ corresponding to eigenvalue 1, and w_i its entry. We plot the delta functions in Fig. 12 in red at these w_i.

We then sampled $N_r = 50 \cdot 10^3$ random sequences γ of cardinality $N = 30 \cdot 10^3$ in D_G (for which we observe again convergence of every P_γ to a rank-one matrix). In this case, each element of the sequence is a randomly chosen node in D_G. We plot in Fig. 12 in black the empirical distributions for the entries of the left-eigenvector of P_γ corresponding to the eigenvalue 1. Again, observe that the support of these empirical distributions have non-zero measures, indicating that there is a continuum of limits.

6. Summary and outlook

We have shown how to construct sets of stochastic matrices (called local stochastic matrices) and adjacency rules for taking product of these matrices that guarantee, under some mild assumptions, convergence of the product to one out of a finite number of possible limits. These limits are all rank-one matrices and, knowing the first matrix in the product is enough to determine which limit the product will converge to.

Underlying our work is the notion of triangulated Laman graph (TLG), where we recall that a Laman graph is a minimally rigid graph. The local stochastic matrices are in one-to-one correspondence with the triangles of this graph, and the adjacency rules for their product are encoded in the hereby defined derived graph of a TLG.

The connections between the minimal rigidity of the underlying graph and the convergence of the product appear at several point in the proof: first, and foremost, in the construction of the unnormalized APVs w_i, where properties of the derived graph of a TLG are integral to the argument; second, in the proof of convergence itself. These connections are either direct, using characterizations of Laman graphs, or rely on the existence of a Restricted Henneberg Construction for the graphs, a fact we proved in the appendix and that requires the graph to be minimally rigid.

We have provided simulations showing that departing from our assumptions—e.g., using a rigid, but non-minimally rigid, graph, or changing the structure of the local stochastic matrices, or disrespecting the adjacency rules in the products— would generically break the conclusions of the Theorems, in particular the conclusion about the number of possible limits of the product.

Beyond their application in the present paper, we believe that some of the novel ideas introduced can form the basis of a broader set of results. In particular, one of the key facts for proving the finite cardinality of the set of possible limits is Proposition 4.2. There, we have established the relation $w_j^\top A_i = r_{i,j} w_i^\top$, where i and j correspond to *adjacent nodes* in the derived graph; indeed, the fact that the limits of the convergent products are exactly $\mathbf{1}\overline{w}_i^\top$ follow as a consequence of the proposition as shown in Sec. 4.3.

The above relation motivates us to consider the following problem: Suppose that one is given a finite set of stochastic matrices $\{A_1, \ldots, A_k\}$ and a set of nonnegative vectors $\{w_1, \ldots, w_k\}$. Define a directed graph D on k nodes as follows: there is an edge from node i to node j if $w_j^\top A_i \propto w_i$. Then, under what circumstances is the graph D strongly connected? If it is, when is every P_γ a rank-one matrix for infinite exhaustive walk γ in D? If we meet conditions so that the answer to the above questions are positive, then P_γ is a rank-one matrix and if the product starts with matrix A_i, then the rank-one matrix has to be $\mathbf{1}\overline{w}_i^\top$, where $\overline{w}_i$ is again the normalized version of w_i. The associated sequence of absolute probability vectors take values $\overline{w}_1, \ldots, \overline{w}_k$.

A few simple examples fitting the above framework are the case of all matrices A_i being the same, or all matrices A_i commuting with each other. In both cases, it is easy to see that D is the complete graph and a positive answer to the above questions is obtained whenever the A_i's are irreducible. A more involved example is the one of gossiping: in this case, the matrices A_i are doubly stochastic matrices, one-to-one correspondent to the edges of a connected undirected graph. All the vectors w_i are chosen to be $\mathbf{1}$. The directed graph D is again the complete graph. The uniqueness of the limit for any infinitely exhaustive walk in D is a consequence of the doubly-stochastic nature of the A_i's. Finally, in the case of the present paper, the A_i's are local stochastic matrices in one-to-one correspondence with triangles of a TLG, and the directed graph D is the directed version of the derived graph D_G. Not much is known beyond these cases. By proposing the framework, and the attendant questions raised above, we look for solutions that generalize and unify the existing results and the results established in the present paper.

Declaration of competing interest

There is no competing interest.

Acknowledgement

This work was supported by the National Science Foundation [ECCS-1809076, ECCS-1809315] and the Air Force Office of Scientific Research [FA9550-20-1-0076, FA9550-20-1-0333].

Appendix A. Proof of Theorem 2.1

On rigidity theory. A graph is rigid if, upon embedding the graph in a Euclidean space $\mathbb{R}^k$, fixing all edge lengths precludes motions of the vertices, save for translations and rotations of the embedded graph. A graph is *minimally rigid* if it is rigid, and no edge can be removed without losing that property. Rigidity in dimension two (i.e., $k = 2$) is relatively well-understood, but many basic questions remain open in dimensions three and above. Minimally rigid graphs in dimension two are called *Laman graphs*. We refer the reader to [9] for formal definitions.

A major result in rigidity theory is the so-called Laman condition, which completely characterizes minimally rigid graphs in dimension two.

Lemma A.1 *(Laman's condition). An undirected graph $G = (V, E)$ on n nodes is minimally rigid if and only if:*

1. *There are $(2n - 3)$ edges in G;*
2. *Every induced subgraph of G on k nodes, for $2 \leq k \leq n - 1$, has at most $(2k - 3)$ edges.*

Henneberg construction. A basic tool in rigidity theory is the so-called Henneberg construction. It is known that every minimally rigid graph admits a *Henneberg construction* and, reciprocally, every Henneberg construction yields a minimally rigid graph. We describe the Henneberg construction below: Starting with an edge, the Henneberg construction iteratively adds a node by applying one of the following two operations at each stage:

1. Node-add: Select two nodes v_i, v_j in G_{n-1}, add a node v_n and the edges (v_i, v_n) and (v_j, v_n) to obtain G_n.
2. Edge-split: Select an edge (v_i, v_j) and a node v_k in G_{n-1}, add a node v_n and edges $(v_i, v_n), (v_j, v_n)$ and (v_k, v_n) remove edge (v_i, v_j).

The sequence of graphs obtained following the construction is called a *Henneberg sequence*. Each graph in a Henneberg sequence is minimally rigid. It should be clear that the RHC introduced in Sec. 2 is a type of Henneberg construction that starts with a triangle.

A Henneberg construction for TLGs. We introduce below a few preliminary results that are needed for the proof of Theorem 2.1.

Lemma A.2. *Let G be a TLG and $G_3, \cdots, G_n$ be a sequence of graphs obtained by following the steps of an RHC, with G_3 a triangle. Then, every graph in the sequence is a TLG.*

Proof. Since the RHC is a type of Henneberg construction, each G_i in the sequence is Laman. We show that it is also triangulated. We proceed by induction of the number of nodes n in the graph. For $n = 3$, G_3 is a triangle and the statement holds. Now assume that G_{n-1} is a TLG. Let v_n, (v_n, v_i), (v_n, v_j) be the node and edges newly added to G_{n-1} to obtain G_n. Consider any cycle γ of length greater than 3 in G_n. Either γ does not contain v_n: it is then included in G_{n-1} and since G_{n-1} is triangulated, it contains a chord. Otherwise, γ contains v_n: then it necessarily contains v_i and v_j as they are the

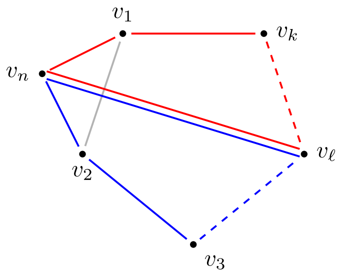

Fig. 13. In the figure, $v_1 v_2 \cdots v_k$ is the chord-free cycle γ inside G_{n-1}. We perform the edge-split operation with (v_1, v_2) the selected edge and v_ℓ the selected node that belongs to the cycle. Then, there are two chord-free cycles after the operation, namely $v_1 v_n v_\ell v_{\ell-1} \cdots v_1$ in red and $v_2 v_3 \cdots v_\ell v_n v_2$ in blue.

only two nodes adjacent to v_n. Since (v_i, v_j) is an edge in G_{n-1} and thus in G_n, the cycle γ has a chord, namely (v_i, v_j). □

In the following Lemma, we show that any *Henneberg* construction for a TLG yields a sequence in which each graph is also a TLG.

Lemma A.3. *Let G be a TLG. Fix a Henneberg construction for G and denote by $G_3, \cdots, G_n = G$ the Laman graphs obtained in that construction, with G_3 a triangle. Then, each G_i is triangulated.*

Proof. We first show that G_{n-1} is triangulated. Assume, by contradiction, that $\gamma = v_1 \cdots v_k v_1$, $k \geq 4$, is a chord-free cycle in G_{n-1} of length greater than 3. On the one hand, if G_n is obtained from G_{n-1} with a node-add operation, then clearly γ is also a chord-free cycle in G_n. On the other hand, assume that G_n is obtained using an edge-split operation. If the selected edge of the edge-split operation is not part of γ, then γ is a chord-free cycle in G_n. We thus assume that the selected edge is part of γ, say (v_1, v_2), and denote by v_ℓ the selected node of G_{n-1}. Note that $v_\ell \neq v_1, v_2$. If $v_\ell \notin \gamma$, then $v_1 v_n v_2 v_3 \cdots v_k v_1$ is a chord-free cycle of length $(k+1)$ in G_n. If $v_\ell \in \gamma$, then $v_1 v_n v_\ell v_{\ell-1} \cdots v_1$ and $v_2 v_3 \cdots v_\ell v_n v_2$ are two distinct chord-free cycles in G_n (see Fig. 13 for illustration). The sum of the lengths of the two cycles is $(k+3) \geq 7$, so at least one of them is of length greater than 3. We have thus shown that if G_{n-1} has a chord-free cycle of length greater than 3, then so does G_n. It thus contradicts the assumption that G_n is triangulated.

We have just shown that if G_n is triangulated, then so is G_{n-1}. Applying the above arguments iteratively, we obtain that $G_{n-2}, \ldots, G_3$ (in a reversed order) are all TLGs as announced. □

The next result indicates which operations of a Henneberg construction create chord-free cycles of lengths greater than 3. Owing to the previous Lemma, these operations cannot be used to construct a TLG.

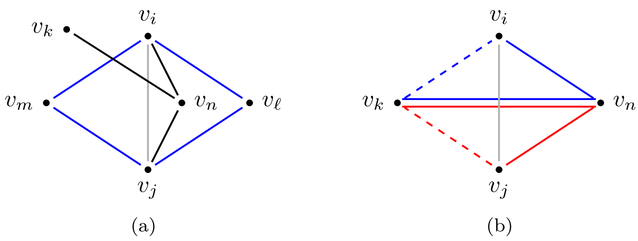

Fig. 14. In Fig. 14a, (v_i, v_j) is a non-simple edge shared by two triangles in G'. Performing edge-split operation on (v_i, v_j) with any node v_k in G' (v_k can be v_m or v_l) results in a chord-free cycle $v_i v_m v_j v_\ell$ of length 4. In Fig. 14b, the dashed lines indicate the shortest path γ from v_i to v_j in $G'' = G' - (v_i, v_j)$. Performing edge-split operation on the edge (v_i, v_j) and a node v_k that belong to γ yields two chord-free cycles depicted in blue $(v_i v_n v_k \cdots v_i)$ and red $(v_j v_n v_k \cdots v_j)$, respectively. The total lengths of the two cycles is the length of γ plus 4.

Lemma A.4. *Let G' be a Laman graph. Let G be the graph obtained by performing on G'* one *Henneberg step taken from the following options:*

1. *Node-add operation connecting a new node to two non-adjacent nodes;*
2. *Edge-split operation splitting a non-simple edge;*
3. *Edge-split operation performed on an edge $e = (v_i, v_j)$ and a node v_k such that the three nodes $\{v_i, v_j, v_k\}$ do not form a triangle in G'.*

Then, the resulting graph G is not triangulated.

Proof. We deal with the three options individually.

Option 1: Denote by v_n the new node and v_i, v_j the existing nodes in G' that are connected to v_n via the one-step Henneberg construction. By assumption, v_i and v_j are not adjacent. Let $\gamma = v_i \cdots v_j$ be the shortest path in G' joining these two nodes. Then, $\omega := \gamma \vee v_n v_i$ is a cycle in G. This cycle is chord-free because γ is a shortest path. Furthermore, since v_i and v_j are not adjacent, the length of ω is greater than 3.

For the remaining two options, we let (v_i, v_j) and v_k be the selected edge and node of G', respectively, for the edge-split operation.

Option 2: Since (v_i, v_j) is not simple, there exist two distinct triangles $\Delta = \{v_i, v_j, v_\ell\}$ and $\Delta' = \{v_i, v_j, v_m\}$ in G' that share the edge. Because G' is a Laman graph, (v_ℓ, v_m) cannot be an edge of G'. To see this, note that if (v_ℓ, v_m) is an edge, then there will be 6 edges in the subgraph of G' induced by the four nodes v_i, v_j, v_ℓ, v_m, which violates Laman's condition (Lemma A.1). But, then, after the edge-split operation, the edge (v_i, v_j) is removed and, hence, $v_i v_m v_j v_\ell v_i$ is a chord-free cycle in G of length 4; see Fig. 14a.

Option 3: Since G' is Laman, it is two-edge-connected. Let G'' be obtained from G' by removing the edge (v_i, v_j). Then, G'' is connected. Denote by $\gamma = v_i \cdots v_j$ the shortest path in G'' joining v_i to v_j. The length of γ is at least 2. If v_k does not belong to γ, then $\gamma \vee v_n v_i$ is chord-free (because γ is a shortest path) and its length is at least 4. We now assume that v_k belongs to γ. Note that $\gamma \neq v_i v_k v_j$ because otherwise these three nodes formed a triangle in G', contradicting our assumption. Hence, the length of γ is at least 3. In this case, the cycle $\omega := \gamma \vee v_n v_i$ has a single chord, namely (v_k, v_n). Indeed, on the one hand, v_n is only connected to v_k by construction, so no other chord is incident to v_n; on the other hand, the fact that γ is a shortest path precludes the existence of a chord between any two nodes in the path. This shows that (v_k, v_n) is the only chord in ω, which can thus be split into two cycles of smaller lengths given by $v_i \cdots v_k v_n v_i$ and $v_k \cdots v_j v_n v_k$. Moreover, the two cycles are chord-free. The sum of the lengths of these two cycles is the length of γ plus 4, which is at least 7. Thus, at least one of the two cycles has its length greater than 3. See Fig. 14b for illustration. □

With the above preliminaries, we now prove Theorem 2.1:

Proof of Theorem 2.1. From Lemma A.2, every graph obtained by an RHC is a TLG. We now show that the converse is also true, i.e., every TLG G can be obtained by an RHC. Let H be a Henneberg construction for G. Note that H can be described by either a sequence of graphs $G_3, \ldots, G_n = G$ along the construction or by a sequence of operations $H_3, \ldots, H_{q-1}$ applied to these graphs, i.e., operation H_p is applied to G_p to obtain G_{p+1}. In the sequel, we will use both descriptions. To keep the notation simple, we do not make explicit the argument of the operations H_i. The arguments are an edge for a node-add operation, and an edge and a node for an edge-split operation. Note that an operation H_p can be applied to any G_q as long as G_q contains the selected edge (and node if H_p is an edge-split).

By Lemma A.3, the Henneberg construction H can only contain operations of the following two types: (1) node-add operation as described in the RHC, or (2) edge-split operation performed on a simple edge (v_i, v_j) and a node v_k so that $\{v_i, v_j, v_k\}$ is a triangle. We now prove that the operations of type (2) can be translated into operations of type (1), thus showing that any Henneberg construction yielding a TLG can be replaced by an RHC.

Let $q \geq 4$ be the smallest integer such that an edge-split operation H_{q-1} as in (2) above is used on G_{q-1} to obtain G_q. Then, the triangle $\{v_i, v_j, v_k\}$ belongs to G_{q-1} and, furthermore, G_{q-1} is obtained by using only the node-add operation. Hence, each H_p, for $p = 3, \ldots, q-2$, is necessarily of type (1) and the truncated sequence $H_3 \cdots H_{q-2}$ is in fact an RHC. The starting triangle of the RHC might not be $\{v_i, v_j, v_k\}$. However, by Proposition 2.3, we can always find another RHC that starts with $\{v_i, v_j, v_k\}$ and yields G_{q-1}. We can thus assume, without loss of generality, that $G_3 = \{v_i, v_j, v_k\}$ is the starting triangle. It is important to note that because (v_i, v_j) is *simple* in G_{q-1}, no operation H_p, for $3 \leq p \leq q-2$, selects the edge (v_i, v_j), since it already belongs to the triangle $G_3 = \{v_i, v_j, v_k\}$.

Next, we will exhibit an RHC $H' = H'_3 \cdots H'_{q-1}$ that yields G_q. Starting from $G'_3 = \{v_i, v_k, v_q\}$, the operation H'_3 simply adds the node v_j and edges (v_j, v_k) and (v_j, v_q), so G'_4 is comprised of two triangles, namely $\{v_i, v_k, v_q\}$ and $\{v_j, v_k, v_q\}$. Note that G'_4 can be also obtained by applying the sequence $H_3 H_{q-1}$ to the triangle $\{v_i, v_j, v_k\}$. Now,

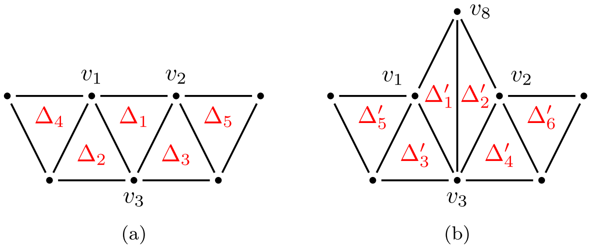

Fig. 15. The TLG depicted in Fig. 15a corresponds to G_{q-1} (here, $q = 8$) in the proof and is obtained via an RHC starting with Δ_1. The indices of the triangles reflect the order of which they have been added into the graph. The operation H_4 selects (v_1, v_3), and H_5 selects (v_2, v_3), etc. The operation H_7 applied to G_7 is an edge-split operation on the edge (v_1, v_2) and the node v_3. It yields the TLG G_8 depicted in Fig. 15b. Starting with $G'_3 = \{v_1, v_3, v_8\}$, we apply the node-add operation H'_3 by adding node v_2 and edges (v_2, v_8) and (v_3, v_8) to obtain G'_4. Then, applying the sequence of operations $H_4 \cdots H_7$, we add triangles $\Delta'_3, \cdots, \Delta'_6$ into the graph and obtain an RHC for G_8.

since H_3 was applied to the triangle $\{v_i, v_j, v_k\}$, but did not select edge (v_i, v_j), we can apply the same operation to G'_4 to obtain G'_5, i.e., we let $H'_4 := H_3$. Next, observe that the edge selected by H_4 belongs to G_4 and it is not (v_i, v_j). Because G_4 (resp. G'_5) is obtained from G_3 (resp. G'_4) using the same operation H_3, the edge selected by H_4 belongs to G'_5. We can thus set $H'_5 := H_4$. Furthermore, note that G'_5 can be obtained by applying the sequence $H_3 H_4 H_{q-1}$ to the triangle $\{v_i, v_j, v_k\}$. Applying the above arguments iteratively, we conclude that for any $p = 4, \ldots, q-1$, the graph G'_p obtained by applying $H'_3 H'_4 \cdots H'_p = H'_3 H_3 \cdots H_{p-1}$ to the triangle $G'_3 = \{v_i, v_k, v_q\}$ is the same as the graph obtained by applying $H_3 \cdots H_{p-1} H_{q-1}$ to the triangle $G_3 = \{v_i, v_j, v_k\}$. In particular, for $p = q-1$, we obtain that $G'_q = G_q$.

Now, replace the original Henneberg construction $H_3 \cdots H_{n-1}$ with the one $H'_3 \cdots H'_{q-1} H_q \cdots H_{n-1}$. By doing so, we reduce by one the number of edge-split operations. One can repeatedly apply the above arguments until all the edge-split operations in the original Henneberg construction are removed. This process ends with an RHC that yields the graph G. □

References

[1] J. Wolfowitz, Products of indecomposable, aperiodic, stochastic matrices, Proc. Am. Math. Soc. 14 (5) (1963) 733–737.

[2] S. Chatterjee, E. Seneta, Towards consensus: some convergence theorems on repeated averaging, J. Appl. Probab. (1977) 89–97.

[3] I. Daubechies, J.C. Lagarias, Sets of matrices all infinite products of which converge, Linear Algebra Appl. 161 (1992) 227–263.

[4] J. Lorenz, A stabilization theorem for dynamics of continuous opinions, Phys. A, Stat. Mech. Appl. 355 (1) (2005) 217–223.

[5] S. Bolouki, R.P. Malhamé, On consensus with a general discrete time convex combination based algorithm for multi-agent systems, in: 2011 19th Mediterranean Conference on Control & Automation (MED), IEEE, 2011, pp. 668–673.

[6] B. Touri, Product of Random Stochastic Matrices and Distributed Averaging, Springer Science & Business Media, 2012.

[7] P.-Y. Chevalier, Convergent products of stochastic matrices: Complexity and algorithms, Ph.D. thesis, Université catholique de Louvain, 2018.

[8] A. Kolmogoroff, Zur theorie der markoffschen ketten, Math. Ann. 112 (1) (1936) 155–160.
[9] J.E. Graver, B. Servatius, H. Servatius, Combinatorial Rigidity, Vol. 2, American Mathematical Soc., 1993.
[10] X. Chen, M.-A. Belabbas, T. Başar, Global stabilization of triangulated formations, SIAM J. Control Optim. 55 (1) (2017) 172–199.
[11] S. Boyd, A. Ghosh, B. Prabhakar, D. Shah, Randomized gossip algorithms, IEEE Trans. Inf. Theory 52 (6) (2006) 2508–2530.
[12] F. He, A.S. Morse, J. Liu, S. Mou, Periodic gossiping, IFAC Proc. Vol. 44 (1) (2011) 8718–8723.
[13] J. Liu, S. Mou, A.S. Morse, B.D. Anderson, C. Yu, Deterministic gossiping, Proc. IEEE 99 (9) (2011) 1505–1524.
[14] Y. Liu, B. Li, B.D. Anderson, G. Shi, Clique gossiping, IEEE/ACM Trans. Netw. 27 (6) (2019) 2418–2431.
[15] X. Chen, M.-A. Belabbas, J. Liu, Gossip over holonomic graphs, available on arXiv:2102.08853, submitted for publication.
[16] X. Wang, S. Mou, S. Sundaram, A resilient convex combination for consensus-based distributed algorithms, Numer. Algebra Control Optim. 9 (3) (2019) 269.
[17] D. Blackwell, Finite non-homogeneous chains, Ann. Math. (1945) 594–599.
[18] J.L. Doob, Kolmogorov's early work on convergence theory and foundation, Ann. Probab. 17 (3) (1989) 815–821.
[19] J. Hajnal, M.S. Bartlett, Weak Ergodicity in Non-homogeneous Markov Chains, Mathematical Proceedings of the Cambridge Philosophical Society, vol. 54, Cambridge University Press, 1958, pp. 233–246.