Decentralized Formation of Random Regular Graphs for Robust Multi-Agent Networks

A. Yasin Yazıcıoğlu, Magnus Egerstedt, and Jeff S. Shamma
Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, GA 30332, USA
yasin@gatech.edu, magnus@gatech.edu, shamma@gatech.edu

***Abstract*— Multi-agent networks are often modeled via interaction graphs, where the nodes represent the agents and the edges denote direct interactions between the corresponding agents. Interaction graphs have significant impact on the robustness of networked systems. One family of robust graphs is the random regular graphs. In this paper, we present a locally applicable reconfiguration scheme to build random regular graphs through self-organization. For any connected initial graph, the proposed scheme maintains connectivity and the average degree while minimizing the degree differences and randomizing the links. As such, if the average degree of the initial graph is an integer, then connected regular graphs are realized uniformly at random as time goes to infinity.**

I. Introduction

Multi-agent networks are often represented via their interaction graphs, where the nodes correspond to the agents and the edges exist between the agents having some direct interaction. Interaction graphs play a significant role in the overall behavior and performance of multi-agent networks. System properties such as *robustness* (how severely the system is influenced by perturbations of process/local components) and *mixing time* (how rapidly information spreads throughout the network) are often analyzed through the topology of the interaction graph (e.g., [1], [2], [3]). Robust interaction graphs with fast mixing times are desirable in numerous multi-agent applications including, but not limited to, flocking and swarming (e.g., [4]), sensor coverage (e.g., [5]), distributed estimation (e.g., [6]), and distributed control of robotic networks (e.g., [7], [8]).

In many applications, multi-agent networks face various perturbations such as component failures, noise, or malicious attacks. The impact of such perturbations on the overall system significantly depends on the robustness of the interaction graph. Node (or edge) connectivity is one of the fundamental robustness measures in graph theory. A graph is said to be *k-node (-edge) connected* if at least k nodes (or edges) must be removed to disconnect the graph. In general, graphs with higher connectivity have higher robustness to random failure of its components (e.g., [1]). An arguably richer measure of robustness is the *expansion ratio*. Expansion ratio is quantified in terms of node expansion or edge expansion (*isoperimetric number*), which are refined notions of connectivity. The isoperimetric number is also closely related to the *algebraic connectivity*, which is given by the second-smallest eigenvalue of the graph Laplacian, i.e. each of these measures is upper and lower bounded through the other (e.g., [9]). Graphs with high expansion ratios (*expanders*) are robust structures with fast mixing times. A detailed overview of expanders along with their numerous applications are presented in [10].

One class of expanders is Ramanujan graphs [11], which are contained within the family of regular graphs. A graph is called a m-regular graph if each node has m edges incident to itself. A m-regular graph is Ramanujan if the second largest (in absolute value) eigenvalue of its adjacency matrix is at most $2\sqrt{m-1}$. As such, the algebraic connectivity of a Ramunjan graph is at least $m-2\sqrt{m-1}$. For $m \geq 3$, almost every m-regular graph is Ramanujan [12], [13]. Hence, for $m \geq 3$, a random m-regular graph with n nodes is a Ramanujan graph with a probability approaching 1 as n goes to infinity.

In this paper, we present a locally applicable scheme to build random regular interaction graphs through self-organization. The proposed method extends the decentralized degree regularization scheme in [14] by incorporating a locally applicable neighborhood randomization rule introduced in [15]. The resulting dynamics minimize the degree differences and randomize the local neighborhoods simultaneously while maintaining the graph connectivity and the total number of edges. As such, it transforms any connected interaction graph with an integer average degree into a connected random regular graph. We also present a distributed implementation, which leads to a uniform limiting distribution over all the connected regular graphs.

The organization of this paper is as follows: Section II presents some preliminaries. Section III presents the proposed method for decentralized formation of random regular graphs. Section IV provides some simulation results. Finally, Section V concludes the paper.

II. Preliminaries

An undirected graph, $\mathcal{G} = (V, E)$, consists of a set of nodes, V, and a set of edges, E, given by unordered pairs of nodes. A graph is connected if there exists a path between any pair of nodes. A *path* is a sequence of nodes such that an edge exists between any two consecutive nodes in the sequence. Any two nodes are said to be adjacent if an edge exists between them. We refer to the set of nodes adjacent

to any node, $i \in V$, as its *neighborhood*, $\mathcal{N}_i$, defined as

$$\mathcal{N}_i = \{j \mid (i,j) \in E\}. \tag{1}$$

For any node i, the number of nodes in its neighborhood is called its *degree*, d_i, i.e.,

$$d_i = |\mathcal{N}_i|, \tag{2}$$

where $|\mathcal{N}_i|$ denotes the cardinality of $\mathcal{N}_i$. For any graph $\mathcal{G}$, we use $\delta(\mathcal{G})$, $\Delta(\mathcal{G})$ and $\bar{d}(\mathcal{G})$ to denote the minimum, the maximum, and the average degrees, respectively. A graph is said to be m-*regular*, if all the entries of its degree vector are equal to m. We refer to the difference of the maximum and the minimum node degrees in a graph as the *degree range* of the graph, $f(\mathcal{G})$, i.e.

$$f(\mathcal{G}) = \Delta(\mathcal{G}) - \delta(\mathcal{G}). \tag{3}$$

For any undirected graph, $\mathcal{G} = (V, E)$, the graph Laplacian, L, is a symmetric matrix whose entries are given as

$$[L]_{ij} = \begin{cases} d_i & \text{if } i = j, \\ -1 & \text{if } (i,j) \in E, \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

The second-smallest eigenvalue of the graph Laplacian is known as the algebraic connectivity of the graph, $\alpha(G)$.

One systematic method of representing locally applicable graph transformations is to use graph grammars (e.g., [16]). A graph grammar, $\Phi = \{r_1, r_2, \ldots\}$, is a set of rules. Each rule is represented as an ordered pair of labeled graphs, $r = (g_l, g_r)$, where the labels represent the node states. As such, a rule defines a change in the edge set that transforms graphs isomorphic to g_l to graphs isomorphic to g_r. A pair of graphs, g and g_l, are said to be *isomorphic* ($g \simeq g_l$) if there exists an edge-preserving bijection between their node sets. A rule is said to be applicable to a graph $\mathcal{G}$, if $\mathcal{G}$ has a subgraph isomorphic to g_l. An initial graph $\mathcal{G}(0)$ along with a grammar Φ defines a system represented as $(\mathcal{G}(0), \Phi)$.

III. Decentralized Formation of Random Regular Graphs

In multi-agent networks, connectivity is essential for various applications since information and interactions cannot spread throughout the network without connectivity. One way to improve the connectivity of a network is to introduce more edges to the interaction graph. However, each edge typically implies some power consumption, communications, sensor measurements, or a physical link. As such, sparsity (having a small number of edges) is also an important feature of multi-agent networks. One family of sparse yet well-connected graphs is the random regular graphs. In this section, we present a decentralized scheme for transforming interaction graphs into random regular graphs while ensuring that the graph remains connected and the average degree is maintained.

The proposed scheme can be represented as a graph grammar, Φ^*. In this setting, each node is labeled with its degree, and $\Phi^* = \{r_1, r_2\}$ is defined as

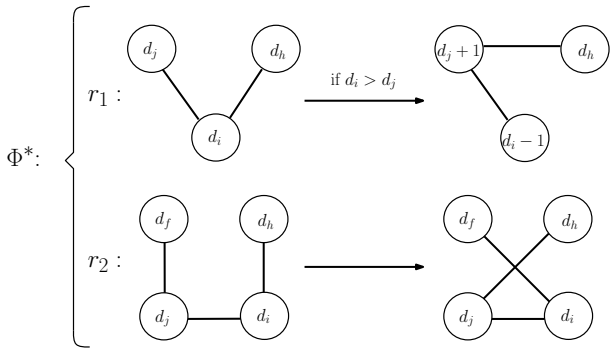

In accordance with Φ^*, if a node has more links than one of its neighbors, then it rewires one of its other neighbors to the less-connected node (r_1). Also, adjacent nodes exchange their exclusive neighbors (r_2). As such, while r_1 aims to balance the degree distribution, r_2 aims to randomize the links between the nodes.

In [14], $\{r_1\}$ was proposed as a grammar to transform any connected graph with an integer average degree into a connected regular graph with the same number of edges. Note that any m-regular graph is stationary under the dynamics induced by $\{r_1\}$ since all of its nodes have the same degree. As such, although almost every m-regular graph is Ramanujan for $m \geq 3$, $\{r_1\}$ can still result in a configuration with an arbitrarily small expansion rate, with a probability depending on the initial graph. For instance, Fig. 1 illustrates a pair of 3-regular graphs on 30 nodes, which are both stationary under $\{r_1\}$ whereas the graph in Fig. 1a is a poorly-connected configuration that can have half of the network disconnected from the rest due the removal of a single edge.

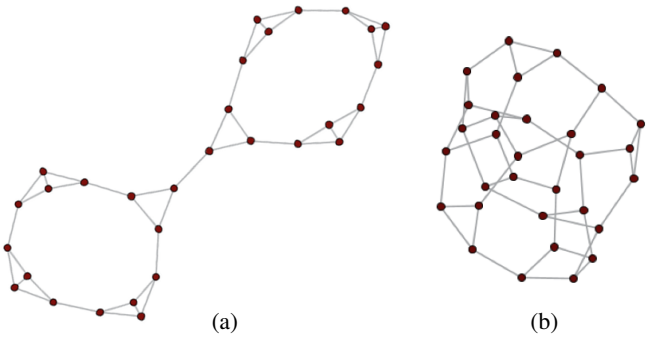

Fig. 1. A poorly-connected 3-regular graph (a) and a robust 3-regular graph (b).

The proposed grammar, $\Phi^* = \{r_1, r_2\}$, extends [14] by incorporating a locally applicable randomization rule introduced in [15], i.e. r_2. Under the resulting dynamics, a connected graph is never stationary (unless it is a complete graph). Furthermore, any connected graph is transformed into a connected random regular graph as time goes to infinity, if its average degree is an integer. In the remainder of this paper, we analyze the dynamics induced by Φ^*.

Lemma 3.1 *Graph connectivity and the average degree are maintained under* Φ^*.

Proof: Both r_1 and r_2 preserve the number edges, so the number of edges and the average degree are invariant to the applications of Φ^*. Furthermore, both rules preserve the connectivity of the local structures. Hence, the global connectivity is also maintained under any concurrent application on disjoint subgraphs. ■

Lemma 3.2 *The degree range,* $f(\mathcal{G})$*, monotonically decreases under* Φ^*.

Proof: Node degrees are invariant to the applications of r_2, and the degree range can only change due to the applications of r_1. In the applications of r_1, the degree of a node, j, can increase only if it has a neighbor, i, with a higher degree ($d_i > d_j$ required). Similarly, its degree can decrease only if it has a neighbor with fewer connection. So neither the minimum degree in the system can decrease, nor the maximum degree in the graph can increase. Hence, the degree range is monotonically decreasing under Φ^*. ■

In general, there may be many feasible applications of Φ^* on an interaction graph. In such cases, the agents need to randomly execute Φ^* on disjoint subgraphs without any global coordination. We propose Algorithm I as a distributed implementation of Φ^* such that any feasible transformation occurs with a non-zero probability.

In accordance with Algorithm I, the nodes behave as follows: At each iteration, each node is inactive with a small probability ϵ. The inactivation probability, ϵ, ensures that any feasible application of Φ^* can be realized through Algorithm I, as it will be shown in Lemma 3.3. Inactive nodes do not participate in any rule execution in that time step. Each active agent, i, picks one of its neighbors, $j \in \mathcal{N}_i$, uniformly at random, and it communicates its degree to that neighbor. Based on these communications, each active agent, i, checks the list of neighbors that picked itself, $R_i \subseteq \mathcal{N}_i$, to see if $j \in R_i$. If that is not the case, then i is a follower in that time step, i.e. it will not initiate a rule execution but it will participate if j wants to rewire i to some other node. If $j \in R_i$, then i and j are matched. Each matched pair randomly picks a candidate rule, $r \in \Phi^*$, that they will potentially execute. In Algorithm I, the candidate rule is picked by the agent with the larger node ID, i.e. $\max\{i, j\}$. If r_1 is picked, and one of the nodes, say i, has higher degree than the other, and i is picked by at least one other neighbor ($|R_i| \geq 2$), then i chooses a neighbor, $h \neq j \in R_i$, uniformly at random. If $h \notin \mathcal{N}_j$, then r_1 is executed by rewiring h to j. If r_2 is to be executed and $|R_i|, |R_j| \geq 2$, both i and j choose one neighbor, $h \neq j \in R_i$ and $f \neq i \in R_j$, uniformly at random. If neither h nor f is linked to both i and j, then r_2 is executed by rewiring h to j and f to i. A feasible iteration of the algorithm on a small network is illustrated in Fig. 2.

Algorithm I: Distributed Implementation of Φ^*

initialize: $\mathcal{G} = (V, E)$ connected, $\epsilon \in (0, 1)$ small
repeat
  Each agent, i, is active with probability $1 - \epsilon$.
  Each active i picks a random $j \in \mathcal{N}_i$.
  For each i, $R_i = \{i' \in \mathcal{N}_i \mid i' \text{ picked } i\}$.
  for (each (i, j) s.t. $i \in R_j, j \in R_i$, $d_i \geq d_j$)
   $\max\{i, j\}$ picks a random $r \in \Phi^*$.
    if ($r = r_1$, $d_i > d_j$, $|R_i| \geq 2$)
      i picks a random $h \in R_i \setminus \{j\}$.
      if ($(j, h) \notin E$)
        $E = (E \setminus \{(i, h)\}) \cup \{(j, h)\}$.
      end if
    else if ($r = r_2$, $|R_i| \geq 2$, $|R_j| \geq 2$)
      i picks a random $h \in R_i \setminus \{j\}$.
      j picks a random $f \in R_j \setminus \{i\}$.
      if ($(i, f) \notin E$, $(j, h) \notin E$)
        $E = (E \setminus \{(i, h), (j, f)\}) \cup \{(i, f), (j, h)\}$.
      end if
    end if
  end for
end repeat

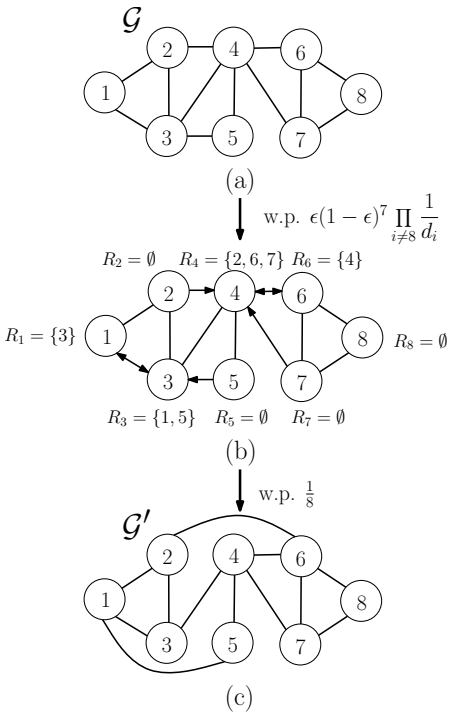

Fig. 2. A feasible iteration of Algorithm I on $\mathcal{G}$ in (a) resulting in $\mathcal{G}'$ in (c) along with the probabilities of the corresponding random events. In this example, each node other than 8 is active and picks a neighbor as illustrated in (b), where each arrow is pointed from a node to its chosen neighbor. Accordingly, (1,3) and (4,6) are the matched pairs satisfying $d_3 > d_1$ and $d_4 > d_6$. With probability 0.25, nodes 3 and 6 both pick r_1 as the candidate rule for their respective matchings. Furthermore, since $R_3 \setminus \{1\} = \{5\}$ and $R_4 \setminus \{6\} = \{2, 7\}$, node 3 picks node 5 to rewire with probability 1, and node 4 picks node 2 to rewire with probability 0.5. Hence, given the configuration in (b), $\mathcal{G}'$ can emerge with a probability of 0.125.

Algorithm I is memoryless since each iteration only depends on the current graph, and the probability of any feasible transition is independent of the past events. Furthermore, the average degree and the connectivity are maintained due to Lemma 3.1. As such, if all agents follow Algorithm I, then a Markov chain is induced over the state space of simple connected graphs having the same number of nodes and the same average degree, i.e.

$$\mathbb{G}_{n,m} = \{\mathcal{G} = (V, E) \mid |V| = n, \bar{d}(\mathcal{G}) = m\}. \quad (5)$$

Let $\mu(t)$ denote the probability distribution over $\mathbb{G}_{n,m}$ at time t. Then, $\mu(t)$ satisfies

$$\mu^T(t+1) = \mu^T(t) P_{\Phi^*}, \quad (6)$$

where P_{Φ^*} is the corresponding probability transition matrix. Accordingly, the probability of transition from any $\mathcal{G}$ to any $\mathcal{G}'$ is denoted by $P_{\Phi^*}(\mathcal{G}, \mathcal{G}')$.

Lemma 3.3 *Let $\mathcal{G}, \mathcal{G}' \in \mathbb{G}_{n,m}$ be any pair of graphs. $P_{\Phi^*}(\mathcal{G}, \mathcal{G}') > 0$ if and only if $\mathcal{G}'$ can be reached from $\mathcal{G}$ in one step via Φ^*.*

Proof: $\Rightarrow$: Since any possible transformation in Algorithm I (lines 11 and 17) satisfies a rule in Φ^*, $P_{\Phi^*}(\mathcal{G}, \mathcal{G}') > 0$ implies that $\mathcal{G}'$ can be reached from $\mathcal{G}$ in one step via Φ^*.

$\Leftarrow$: Let $\mathcal{G}'$ be reachable from $\mathcal{G}$ in one step via Φ^*, and let $G = \{g_1, g_2, \ldots\}$ denote the corresponding set of disjoint subgraphs of $\mathcal{G}$ to be transformed to reach $\mathcal{G}'$. Note that for each $g \in G$ there is a $r = (g_l, g_r) \in \Phi^*$ satisfying $g \simeq g_l$. Here, we present a feasible flow of Algorithm I that transforms $\mathcal{G}$ by applying the corresponding r to each $g \in G$. For each $g \in G$, let the nodes in g be active at that time step, and pick a neighbor as illustrated in Fig. 3. Furthermore, let any node that is not included in any $g \in G$ be inactive, which ensures that only the subgraphs in G will be transformed. Finally, let each g pick the corresponding r as the candidate rule to execute. In that case, the agents are guaranteed to only apply the corresponding $r \in \Phi^*$ to each $g \in G$. Hence, the corresponding transformation has a non-zero probability in P_{Φ^*}, i.e. $P_{\Phi^*}(\mathcal{G}, \mathcal{G}') > 0$.

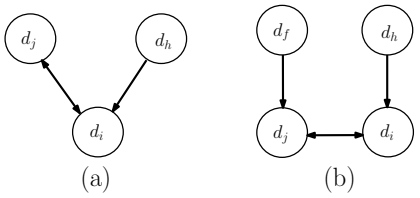

Fig. 3. An arrow is pointed from each agent to the neighbor it picked. For each $g \in G$, the nodes in g have non-zero probability to pick their neighbors as shown in (a) if $r = r_1$, and as shown in (b) if $r = r_2$.

■

In general, $\mathbb{G}_{n,m}$ can be represented as the union of two disjoint sets, $\mathbb{G}^0_{n,m}$ (regular graphs) and $\mathbb{G}^+_{n,m}$ (non-regular graphs), defined as

$$\mathbb{G}^0_{n,m} = \{\mathcal{G} \in \mathbb{G}_{n,m} \mid f(\mathcal{G}) = 0\}, \quad (7)$$
$$\mathbb{G}^+_{n,m} = \mathbb{G}_{n,m} \setminus \mathbb{G}^0_{n,m}. \quad (8)$$

Note that if $m \in \mathbb{N}$, then $\mathbb{G}^0_{n,m} \neq \emptyset$ and it is possible to build a regular graph with the available number of edges. We will show that if $\mathbb{G}^0_{n,m} \neq \emptyset$, then P_{Φ^*} has a unique limiting distribution that is uniform over $\mathbb{G}^0_{n,m}$. First, it is shown that all the graphs in $\mathbb{G}^+_{n,m}$ are transient states of P_{Φ^*} if $m \in \mathbb{N}$, i.e. whenever the current graph is non-regular, there is a non-zero probability that the system will leave that configuration and never return.

Lemma 3.4 *[14] Let $\mathcal{G}$ be a connected graph such that $\bar{d}(\mathcal{G}) = m \in \mathbb{N}$. Then $(\mathcal{G}, \{r_1\})$ almost surely converges to an m-regular graph.*

Proof: This is proved in [14] by showing that the degree range is monotonically decreasing under $\{r_1\}$, and for any graph with an integer average degree, there is always a finite sequence of r_1 applications leading to a regular graph. ■

Lemma 3.5 *Let $\mathcal{G} \in \mathbb{G}_{n,m}$, and let $m \in \mathbb{N}$. If $\mathcal{G} \in \mathbb{G}^+_{n,m}$, then $\mathcal{G}$ is a transient state of P_{Φ^*} .*

Proof: Let $\mathcal{G} \in \mathbb{G}^+_{n,m}$, and let $m \in \mathbb{N}$. In light of Lemma 3.4, there exists a finite sequence of r_1 applications transforming $\mathcal{G}$ to some $\mathcal{G}' \in \mathbb{G}^0_{n,m}$. Due to Lemma 3.3, the corresponding trajectory is also feasible under P^*_Φ. Furthermore, due to Lemma 3.2, it is not possible to return to $\mathcal{G}$ once a regular graph $\mathcal{G}'$ is reached. Hence, $\mathcal{G}$ is transient. ■

Next, we show that if $\mathbb{G}^0_{n,m} \neq \emptyset$, then $\mathbb{G}^0_{n,m}$ is a closed communicating class of P_{Φ^*}. To this end, we first provide a result from [15], where the authors proposed an operation called Random 1-Flipper, which is a randomized execution of r_2.

Random 1-Flipper [15]:

1 : Choose a random edge $(i, j) \in E$.
2 : Choose a random node $h \in \mathcal{N}_i \setminus \{j\}$.
3 : Choose a random node $f \in \mathcal{N}_j \setminus \{i\}$.
4 : **if** $((i, f) \notin E, (j, h) \notin E)$
5 : $\quad E = (E \setminus \{(i, h), (j, f)\}) \cup \{(i, f), (j, h)\}$.
6 : **end if**

Lemma 3.6 *[15] Let $\mathcal{G}$ be a connected m-regular graph with n nodes and $m > 2$. Then in the limit the Random 1-Flipper operation constructs all connected m-regular labeled graphs with the same probability.*

Proof: This is proved in [15] by showing that a repetitive application of Random 1-Flipper operation induces an aperiodic irreducible Markov chain with a doubly stochastic probability transition matrix over the set of all connected m-regular labeled graphs on n nodes. ■

Lemma 3.7 *If $\mathbb{G}^0_{n,m} \neq \emptyset$, then $\mathbb{G}^0_{n,m}$ is a closed communicating class of P_{Φ^*}.*

Proof: In light of Lemma 3.6, any $\mathcal{G} \in \mathbb{G}^0_{n,m}$ can be reached from any other $\mathcal{G}' \in \mathbb{G}^0_{n,m}$ through a sequence of r_2 applications. Due to Lemma 3.3, the corresponding sequence of r_2 applications is also feasible under P_{Φ^*}. Hence, $\mathbb{G}^0_{n,m}$ is a communicating class. Furthermore, $\mathbb{G}^0_{n,m}$ is closed due to Lemma 3.2. ∎

Lemma 3.8 *For any pair of regular graphs, $\mathcal{G}, \mathcal{G}' \in \mathbb{G}^0_{n,m}$,*

$$P_{\Phi^*}(\mathcal{G}, \mathcal{G}') = P_{\Phi^*}(\mathcal{G}', \mathcal{G}). \tag{9}$$

Proof: For any regular graph, $\mathcal{G} \in \mathbb{G}^0_{n,m}$, r_1 is not applicable on $\mathcal{G}$ since all the nodes have equal degrees. Hence, any transition from $\mathcal{G}$ to another $\mathcal{G}' \in \mathbb{G}^0_{n,m}$ is only via r_2. Note that r_2 is a reversible rule, i.e. if two nodes, i and j, exchange their neighbors, $h \in \mathcal{N}_i$ and $f \in \mathcal{N}_j$, in accordance with r_2, then swapping those neighbors back is also a feasible application of r_2.

Let us consider an arbitrary execution of Algorithm I, where $\mathcal{G} = (V, E) \in \mathbb{G}^0_{n,m}$ is transformed into $\mathcal{G}' = (V, E') \in \mathbb{G}^0_{n,m}$. Let u be the corresponding vector of randomly picked neighbors in line 4 of Algorithm I (let $u_i = null$ if i was inactive). For each node, i, let $R_i(u)$ be the set of nodes that picked i, and let $M(u) = \{(i,j) \mid i \in R_j(u), j \in R_i(u)\}$ be the set of matched pairs. For $\mathcal{G}'$, consider the vector, u', whose entries are

$$u'_i = \begin{cases} u_i & \text{if } u_i = null \text{ or } (i, u_i) \in E', \\ u_{u_i} & \text{otherwise.} \end{cases} \tag{10}$$

Note that $Pr[u] = Pr[u']$ since the inactive nodes will remain inactive with the same probability, and each active node picks a neighbor uniformly at random. Furthermore, $M(u) = M(u')$. Also, for every $(i,j) \in M(u)$, we have $|R_i(u)| = |R_i(u')|$ and $|R_j(u)| = |R_j(u')|$. Hence, each $(i,j) \in M(u)$ will reverse the neighbor-swapping in the transition from $\mathcal{G}$ to $\mathcal{G}'$ with the same probability (lines 13-19 in Algorithm I), so $Pr[\mathcal{G} \to \mathcal{G}'; u] = Pr[\mathcal{G}' \to \mathcal{G}; u']$. Consequently, $P_{\Phi^*}(\mathcal{G}, \mathcal{G}') = P_{\Phi^*}(\mathcal{G}', \mathcal{G})$. ∎

Theorem 3.9 *Let $\mathbb{G}_{n,m}$ satisfy $m \in \mathbb{N}$. Then, P_{Φ^*} has a unique limiting distribution, μ^*, satisfying*

$$\mu^*(\mathcal{G}) = \begin{cases} 1/|\mathbb{G}^0_{n,m}| & \text{if } \mathcal{G} \in \mathbb{G}^0_{n,m}, \\ 0 & \text{otherwise.} \end{cases} \tag{11}$$

Proof: For $m \in \mathbb{N}$, $\mathbb{G}^0_{n,m} \neq \emptyset$ and $\mathbb{G}^0_{n,m}$ is a closed communicating class due to Lemma 3.7. Furthermore it is the only closed communicating class since all the other states (graphs in $\mathbb{G}^+_{n,m}$) are transient due to Lemma 3.5. As such, P_{Φ^*} has a unique stationary distribution, μ^*, whose support is $\mathbb{G}^0_{n,m}$. In the remainder of the proof, the behavior of the Markov chain within $\mathbb{G}^0_{n,m}$ is inspected to show that μ^* is a limiting distribution and it is uniform over $\mathbb{G}^0_{n,m}$.

Let $P^0_{\Phi^*}$ be the $|\mathbb{G}^0_{n,m}|$ by $|\mathbb{G}^0_{n,m}|$ probability transition matrix that only represents the transitions within $\mathbb{G}^0_{n,m}$. Due to Lemma 3.7, $P^0_{\Phi^*}$ is irreducible. Also $P^0_{\Phi^*}$ is aperiodic since $P^0_{\Phi^*}(\mathcal{G}, \mathcal{G}) > 0$ for every $\mathcal{G}$ (for instance, there is a non-zero probability of all the nodes being inactive). As such, $P^0_{\Phi^*}$ has a unique limiting distribution, μ^{*0}. Furthermore, due to Lemma 3.8, $P^0_{\Phi^*}$ is symmetric, and it is consequently doubly stochastic. As a result, μ^{*0} is uniform over $\mathbb{G}^0_{n,m}$. ∎

IV. Simulation Results

In this section, we present some simulation results for the proposed scheme. In the first simulation, we pick an arbitrary $\mathcal{G}(0) \in \mathbb{G}^+_{8,3}$, which is illustrated in Fig. 2a, and the interaction graph is evolved using Algorithm I. In light of Theorem 3.9, the system is expected to approach a uniform limiting distribution over $\mathbb{G}^0_{8,3}$. Note that $\mathbb{G}^0_{8,3}$ is a very large set, hence it is not feasible to track the limiting probability corresponding to each individual graph in $\mathbb{G}^0_{8,3}$. However, there are only 5 non-isomorphic structures in $\mathbb{G}^0_{8,3}$ as depicted in Fig. 4, and we will present the proportions of time spent in each isomorphism class instead.

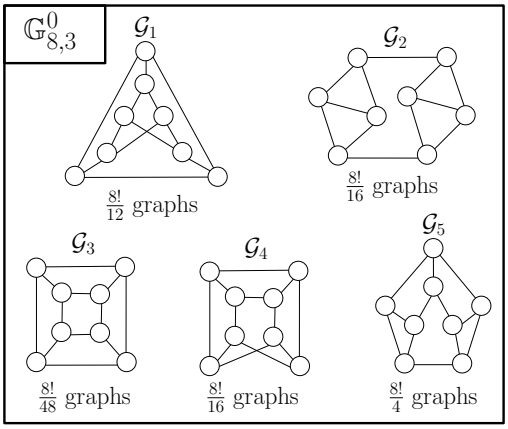

Fig. 4. Non-isomorphic graph structures, $\mathcal{G}_1, \ldots, \mathcal{G}_5$, in $\mathbb{G}^0_{8,3}$, and the number of labeled graphs isomorphic to each structure.

Note that the cardinality of the isomporphism classes in $\mathbb{G}^0_{8,3}$ are not equal since each structure has different symmetry properties. In the simulation, proportions of the time spent at each of the 5 isomorphism classes are recorded in a vector $v(t)$ defined as

$$v_i(t) = \frac{|\{0 \leq \tau \leq t \mid \mathcal{G}(\tau) \simeq \mathcal{G}_i\}|}{t+1}, \forall i = 1, \ldots, 5. \tag{12}$$

As the Markov chain approaches the limiting distribution in (11), $v(t)$ is expected to be in alignment with the sizes of isomorphism classes in $\mathbb{G}^0_{8,3}$, i.e. it should approach

$$v^* = [\frac{4}{23}, \frac{3}{23}, \frac{1}{23}, \frac{3}{23}, \frac{12}{23}]. \tag{13}$$

Fig. 5 depicts $||v(t) - v^*||_2$ throughout the simulation. Since $\mathcal{G}(0) \in \mathbb{G}^+_{8,3}$, $||v(t) - v^*||_2$ is stationary at $||v^*||$ for a short period ($t \leq 74$) until the degree range drops to 0, i.e. the system enters $\mathbb{G}^0_{8,3}$. After that, $||v(t) - v^*||_2$ rapidly decreases and approaches 0, as expected from a uniform limiting distribution over $\mathbb{G}^0_{8,3}$.

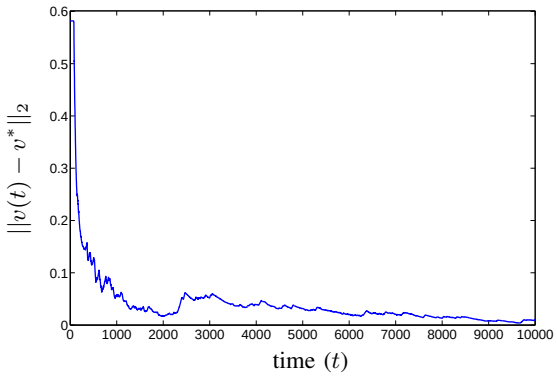

Fig. 5. $||v(t) - v^*||_2$ as a function of time. $v(t)$ approaches v^* in accordance with the uniform limiting distribution over $\mathbb{G}^0_{8,3}$.

In the second simulation, we consider a larger network, $\mathcal{G}(0) \in \mathbb{G}^+_{100,3}$, and illustrate how the robustness of the interaction graph changes under Algorithm I during a period of 10000 time seps. Particularly, we inspect the change in the algebraic connectivity. Fig. 6 illustrates $\mathcal{G}(0)$ and $\mathcal{G}(10000)$, which have algebraic connectivities 0.032 and 0.195 respectively. On the corresponding trajectory, the degree range drops from 5 to 0 within 2942 steps. After that, the system keeps randomizing within $\mathbb{G}^0_{100,3}$. The evolution of the algebraic connectivity throughout this simulation is shown in Fig. 7. Note that almost every 3-regular graph is Ramanujan and has an algebraic connectivity at least $3 - 2\sqrt{2}$. Accordingly, the algebraic connectivity of the simulated network is at least $3 - 2\sqrt{2}$ with a very high probability after a sufficient amount of time.

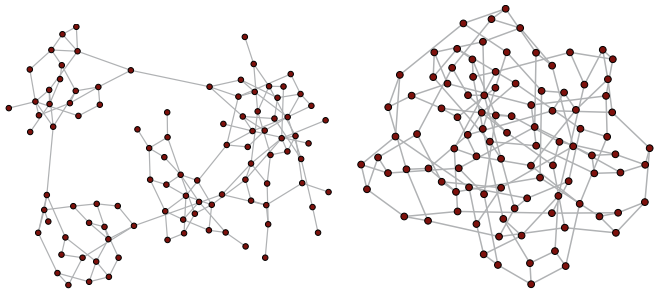

Fig. 6. Agents follow Algorithm I so that the initial graph in (a) is transformed into a robust interaction structure such as the one in (b).

V. Conclusion

In this paper, we presented a decentralized scheme to build robust interaction graphs. In particular, random regular graphs are obtained through self-organization. For $m \geq 3$, random m-regular graphs are almost surely Ramanujan, and as such they are sparse yet well-connected structures.

The proposed scheme is represented as a graph grammar, $\Phi^* = \{r_1, r_2\}$. The first rule minimizes the degree differences in the system, whereas the second rule randomizes the links. The average degree and the graph connectivity

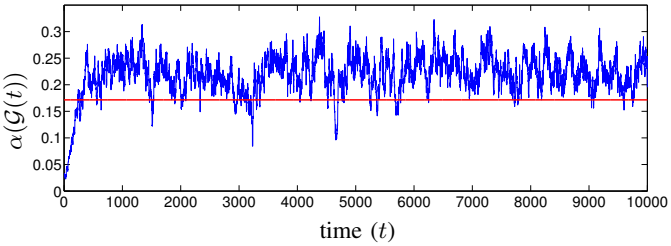

Fig. 7. The algebraic connectivity, $\alpha(\mathcal{G}(t))$, as the initial graph in Fig. 6a evolves via Algorithm I. After sufficiently large time, $\alpha(\mathcal{G}(t))$ rarely drops below $3 - 2\sqrt{2}$ (marked with a solid line), since the corresponding 3-regular graphs are Ramanujan with a very high probability.

are maintained under the resulting dynamics. Furthermore, a distributed implementation of Φ^* was presented. It was shown that, if the average degree of the initial graph is an integer, then the proposed algorithm leads to a limiting distribution that is uniform over all the connected regular graphs. Some simulations were also presented to complement the theoretical results.

References

[1] A. H. Dekker and B. D. Colbert, "Network robustness and graph topology," in *Australasian Conference on Computer Science*, pp. 359–368, 2004.

[2] A. Jamakovic and S. Uhlig, "On the relationship between the algebraic connectivity and graph's robustness to node and link failures," in *EuroNGI Conference on Next Generation Internet Networks*, pp. 96–102, 2007.

[3] R. Olfati-Saber, "Ultrafast consensus in small-world networks," in *American Control Conference*, pp. 2371–2378, 2005.

[4] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Flocking in fixed and switching networks," *IEEE Transactions on Automatic Control*, vol. 52, no. 5, pp. 863–868, 2007.

[5] J. Cortés, S. Martínez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.

[6] R. Olfati-Saber and J. S. Shamma, "Consensus filters for sensor networks and distributed sensor fusion," in *IEEE Conference on Decision and Control*, pp. 6698–6703, 2005.

[7] F. Bullo, J. Cortés, and S. Martinez, *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*. Princeton University Press, 2009.

[8] M. Mesbahi and M. Egerstedt, *Graph theoretic methods in multiagent networks*. Princeton University Press, 2010.

[9] B. Mohar, "Isoperimetric numbers of graphs," *Journal of Combinatorial Theory, Series B*, vol. 47, no. 3, pp. 274–291, 1989.

[10] S. Hoory, N. Linial, and A. Wigderson, "Expander graphs and their applications," *Bulletin of the American Mathematical Society*, vol. 43, no. 4, pp. 439–561, 2006.

[11] M. R. Murty, "Ramanujan graphs," *Journal of the Ramanujan Mathematical Society*, vol. 18, no. 1, pp. 33–52, 2003.

[12] N. Alon, "Eigenvalues and expanders," *Combinatorica*, vol. 6, no. 2, pp. 83–96, 1986.

[13] J. Friedman, "A proof of alon's second eigenvalue conjecture," in *ACM Symposium on Theory of Computing*, pp. 720–724, 2003.

[14] A. Y. Yazıcıoğlu, M. Egerstedt, and J. S. Shamma, "Decentralized degree regularization for multi-agent networks," in *IEEE Conference on Decision and Control*, pp. 7498–7503, 2013.

[15] P. Mahlmann and C. Schindelhauer, "Peer-to-peer networks based on random transformations of connected regular undirected graphs," in *ACM Symposium on Parallelism in Algorithms and Architectures*, pp. 155–164, 2005.

[16] E. Klavins, R. Ghrist, and D. Lipsky, "Graph grammars for self assembling robotic systems," in *IEEE International Conference on Robotics and Automation*, pp. 5293–5300, 2004.