

# Evaluating Network Rigidity in Realistic Systems: Decentralization, Asynchronicity, and Parallelization

Ryan K. Williams, *Student Member, IEEE*, Andrea Gasparri, *Member, IEEE*, Attilio Priolo, and Gaurav S. Sukhatme, *Fellow, IEEE*

**Abstract**—In this paper, we consider the problem of evaluating the rigidity of a planar network, while satisfying common objectives of real-world systems: decentralization, asynchronicity, and parallelization. Motivating our investigation are the implications that rigidity has in fundamental multi-robot problems, e.g. guaranteed formation stability and relative localizability. We propose the decentralization of the pebble game algorithm of Jacobs et. al., an  $O(n^2)$  method that determines the generic rigidity of a planar network. Our decentralization is based on asynchronous inter-agent messaging and a distributed memory architecture, coupled with distributed auctions for electing network leaders to arbitrate rigidity evaluation. Further, we provide a parallelization of our methods that takes inspiration from gossip algorithms to yield significantly reduced execution time and messaging burden. A thorough analysis of the correctness, finite termination, and complexity of our propositions is given, along with a simulated application in decentralized rigidity control. Finally, we provide Monte Carlo analysis of our algorithms in a *Contiki* networking environment, illustrating the real-world applicability of our methods, and yielding a bridge between rigidity theory and realistic interacting systems.

**Index Terms**—Networked Robots; Distributed Robot Systems; Asynchronous and Parallel Communication; Graph Rigidity.

## I. INTRODUCTION

MULTI-ROBOT networks remain among the areas of interest at the forefront of robotics research, particularly given the steady advancement of wireless communication, embedded computation, and hardware platforms (e.g. quadcopters). Intuitively, networks of intelligently interacting robots provide significant advantages over the single-agent alternative; for example scalability, failure robustness, spatiotemporal efficiency, heterogeneity, etc. As recent work has demonstrated, multi-robot investigations are far-reaching across various disciplines, ranging from sampling, tracking, and coverage [1]–[3], mobility and topology control [4]–[6], to general agent agreement problems [7]–[9].

In modeling and analyzing multi-robot networks, research balances between accuracy in approximating realistic systems, and ease of technical analysis, most typically in understanding mobility, communication, and sensing. Here we take the former approach, considering a problem that underlies *fundamental* objectives in multi-robot research, while operating under the commonly desired parameters of real-world systems: *decentralized* implementation where information is exchanged only in local neighborhoods, *asynchronicity* in interaction (communication),

and *parallelization* of agent actions to maximize efficiency. Our problem of interest in this work is the evaluation of the *rigidity* property of an interconnected system of intelligent agents (e.g. robots, sensors, etc.). A relatively under-explored topic in the area of multi-agent systems, rigidity has important implications particularly for mission objectives requiring collaboration. For example, rigidity is vital for guaranteeing stability in controlling formations of mobile vehicles, when only relative inter-agent information is available [10]–[15]. Further, when a global frame of reference is inaccessible, rigidity becomes a necessary (and under certain conditions sufficient) condition for localization tasks with distance or bearing-only measurements [16]–[19]. Rigidity is also a necessary component of *global rigidity* [20]–[22], which can further strengthen the guarantees of formation stability and localizability. We point out that it is typical in the literature to *assume* rigidity properties of the network in order to achieve multi-agent behaviors, however few works provide means of evaluating or achieving network rigidity in a dynamic manner, or under the network conditions considered here.

The general study of rigidity has a rich history in various contexts of science, mathematics, and engineering [21]–[28]. In [27] combinatorial operations are defined which preserve rigidity, with works such as [10], [12] extending the ideas to multi-robot formations. In [29] an algorithm is proposed for generating rigid graphs in the plane based on the Henneberg construction [27], however from a centralized (or all-to-all) perspective. Similarly, [30] defines decentralized rigid constructions that are edge length optimal, however provide no means of determining an unknown graph’s rigidity properties. The work [31] defines a rigidity eigenvalue for infinitesimal rigidity evaluation and control, however such efforts remain centralized and require continuous communication and computational resources.

As opposed to previous work, we propose a *decentralized* method of evaluating *generic* graph rigidity in the plane, *without* a priori topological information, to our knowledge the first such effort, particularly in a multi-agent context. Specifically, we decentralize in an *asynchronous* manner the *pebble game* proposed by Jacobs and Hendrickson in [28], an algorithm that determines in  $O(n^2)$  time the combinatorial rigidity of a network, and a spanning edge set defining the minimally rigid subcomponent of the graph<sup>1</sup>. Specifically, we propose a leader election procedure based on distributed auctions that manages the sequential nature of the pebble game in a decentralized setting, together with a distributed memory architecture. Further, an asynchronous messaging scheme preserves local-only agent interaction, as well as robustness to delays, failures, etc. Towards network effi-

R. K. Williams and G. S. Sukhatme are with the Departments of Electrical Engineering and Computer Science at the University of Southern California, Los Angeles, CA 90089 USA (rkwillia@usc.edu; gaurav@usc.edu).

A. Gasparri and A. Priolo are with the Department of Engineering, University of Roma Tre, Via della Vasca Navale, 79. Roma, 00146, Italy (gasparri@dia.uniroma3.it; priolo@dia.uniroma3.it).

<sup>1</sup>Such an edge set is not obtained by infinitesimal approaches, e.g. [31]

ciency, we extend our decentralization by parallelizing a portion of the rigidity evaluation, taking inspiration from local *gossip* messaging, yielding significant improvements in execution time and communication. To illustrate our contributions, we provide a thorough analysis of the correctness, finite termination, and complexity of our propositions, along with an illustration of decentralized rigidity control. Finally, we provide Monte Carlo analysis of our algorithms in a *Contiki* networking environment, illustrating the real-world applicability of our methods. Although a few recent works have begun to investigate rigidity evaluation or control (e.g. [29]–[32]), they provide graph constructions or centralized control relying on expensive estimation techniques. Our contributions therefore aim to bridge the gap between fundamentally important multi-agent behaviors and realistic rigidity evaluation in networked systems, ultimately moving towards robotic/sensor systems with achievable rigidity-based behaviors.

In summary, the major contributions of this paper are as follows:

- A leader-based, asynchronous decentralization of the classically centralized and serial pebble game algorithm, yielding a decentralized method for planar rigidity evaluation.
- A case study of exploiting problem structure (i.e. independence preserving operations) for parallelization of our decentralized algorithm.
- A characterization of our algorithms in the real-world *Contiki* networking environment, with a full codebase release for use in the robotics community.

A preliminary portion of this paper will appear in [33], compared to which we provide expanded algorithm pseudocode, analysis, and correctness proofs, a complete treatment of parallelization, expanded rigidity control simulations, and a Monte Carlo analysis in the *Contiki* environment demonstrating applicability under realistic conditions.

The outline of the paper is as follows. In Section II, we provide preliminary materials including agent and network models, a model of algorithm execution, and primers on rigidity theory and the pebble game. A decentralization of the pebble game is presented in Section III, with a parallelization given in Section IV. Simulation results are provided in Section V, with concluding remarks as well as directions for future work are stated in Section VI.

## II. PRELIMINARIES AND FORMULATION

### A. Agent, Network, and Execution Models

Consider a system composed of  $n$  agents indexed by  $\mathcal{I} = \{1, \dots, n\}$  operating in  $\mathbb{R}^2$ , each possessing computation and communication capabilities, denoting by  $(i, j)$  a bi-directional communication link between agents  $i$  and  $j$ . The agent may also be mobile, and thus we assume basic continuous dynamics

$$\dot{x}_i = f(\mathbf{x}) \quad (1)$$

where  $x_i, f(x_i) \in \mathbb{R}^2$  are the position and the velocity control input for an agent  $i \in \mathcal{I}$ , respectively, and  $\mathbf{x} \in \mathbb{R}^{2n}$  is the vector of stacked agent positions. For the purposes of integration into a motion control architecture (Section V-A) it is further assumed that each robot can sense other nearby robots and obstacles, yielding the displacement  $d_{ij} \in \mathbb{R} \triangleq \|x_{ij}\| \triangleq \|x_i - x_j\|$ .

To describe the interconnected system formally, we define *undirected graph*  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , having vertices (nodes)  $\mathcal{V} = \{v_1, \dots, v_n\}$  associated with each agent  $i \in \mathcal{I}$ , and edge set  $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$  with members  $(i, j)$ , where by definition  $(i, j) \in \mathcal{E} \Leftrightarrow (j, i) \in \mathcal{E}, \forall i \neq j \in \mathcal{I}$ , excluding the possibility for self loops,  $(i, i) \notin \mathcal{E}, \forall i \in \mathcal{I}$ . Agents  $i$  and  $j$  with an edge  $(i, j) \in \mathcal{E}$  are referred to as *neighbors*, where the set of neighbors for the  $i$ th agent is given by  $\mathcal{N}_i = \{v_j \in \mathcal{V} \mid (i, j) \in \mathcal{E}\}$ .

*Assumption 1 (Connectedness)*: We assume the network topology  $\mathcal{G}$  is *connected*<sup>2</sup> for all time to guarantee all agents can participate in rigidity evaluation (Sections III and IV), that is for every pair of nodes  $i, j$  there exists a sequence of nodes in  $\mathcal{G}$  that are adjacent (a *path*) and connect  $i, j$ .

As our concern in this work is to operate under the typical parameters of realistic interacting systems, we assume an *asynchronous* model of time, where we assume each agent  $i \in \mathcal{I}$  has a clock which ticks at a Poisson rate of 1, independently of the clocks of the other agents [34] (allowing also for the possibility of delayed communication over links  $(i, j) \in \mathcal{E}$ ). Equivalently, this corresponds to a global clock which ticks according to a rate  $n$  Poisson process at times  $t_k, k \geq 1$ , having time-slots  $[t_k, t_{k+1})$  which discretizes system time according to clock ticks (for convenience we will use simply  $t$ ) [35]. Such assumptions induce asynchronicity in both the execution of local agent computations (i.e. algorithms) and the broadcast and reception of inter-agent messages. First, we make the following assumptions concerning agent execution:

*Assumption 2 (Agent execution)*: Each agent  $i \in \mathcal{I}$  executes according to an algorithm (which may be state dependent and change over time) on ticks of their local clock, handling messages from  $j \in \mathcal{N}_i$  and sending messages if dictated by the execution. *Local execution* is assumed to consist of atomic logic (i.e. the algorithm) and message handling, that is all local algorithmic state, denoted  $\mathbf{X}_i$ , is assumed to be devoid of race conditions due to asynchronicity.

A coordinated algorithm execution with associated stopping condition can then be defined as:

*Definition 2.1 (Coordinated execution)*: A coordinated *algorithm execution* is given as a sequence of ticks  $t_k$  and therefore local execution and asynchronous messaging, yielding a *terminal state* upon some discrete network stopping condition

$$f_{stop} \in \{0, 1\} \triangleq f(\{\mathbf{X}_1, \dots, \mathbf{X}_n\}) \quad (2)$$

dependent on the execution states of the network agents. It is assumed that (2) can be computed using distributed techniques (such as agreement [36]), as will be demonstrated in our proposed algorithms. After the stopping condition is observed the agents enter into an *idle* state where no execution occurs.

Finally, to guarantee soundness with respect to network communication and asynchronicity, we make the following assumptions:

*Assumption 3 (Asynchronous messaging)*: We assume each agent  $i \in \mathcal{I}$  treats messages received from neighbors  $j \in \mathcal{N}_i$  in a *first-in-first-out* (FIFO) manner, guaranteeing soundness with respect to our proposed algorithm executions. Further, the possibility of communication failure is handled with *best-effort*

<sup>2</sup>Note that this assumption is trivially satisfied in rigid networks.

messaging, i.e. there exists an underlying communication control layer where a best effort is made to deliver packets in the network. Specifically, we assume that the best effort guarantees message reception in finite time, or equivalently a message failure can be handled appropriately with respect to the algorithms that will be discussed.

### B. Rigidity Theory

The primary concern of this work is the *rigidity* property of the underlying graph  $\mathcal{G}$  describing the network topology<sup>3</sup>, specifically as rigid graphs imply guarantees for example in both localizability and formation stability of multi-robot systems [12]. To begin we require the notion of a graph embedding in the plane, captured by the *framework*  $\mathbb{F}_p \triangleq (\mathcal{G}, p)$  comprising graph  $\mathcal{G}$  together with a mapping  $p: \mathcal{V} \rightarrow \mathbb{R}^2$ , assigning to each node in  $\mathcal{G}$ , a location in  $\mathbb{R}^2$ . The natural embedding for us is to assign each node the position  $x_i$  associated with each agent, defined by the mapping  $p(i) = x_i$ , otherwise known as a *realization* of  $\mathcal{G}$  in  $\mathbb{R}^m$ . The *infinitesimal motion* of  $\mathbb{F}_p$  can be described by assigning to the vertices of  $\mathcal{G}$ , a velocity  $\dot{p}_i \triangleq \dot{x}_i \in \mathbb{R}^2$  such that

$$(\dot{x}_i - \dot{x}_j) \cdot (x_i - x_j) = 0, \quad \forall (i, j) \in \mathcal{E} \quad (3)$$

where  $\cdot$  is the standard dot product over  $\mathbb{R}^m$ . That is, the edge lengths (inter-agent distance) over  $\mathcal{G}$  are preserved in time, in other words no edge is compressed or stretched over time. The framework is said to undergo a *finite flexing* if  $p_i$  is differentiable and edge lengths are preserved, with *trivial flexings* defined as translations and rotations of  $\mathbb{R}^2$  itself. If for  $\mathbb{F}_p$  all infinitesimal motions are trivial flexings, then  $\mathbb{F}_p$  is said to be *infinitesimally rigid*. Otherwise, the framework is called *infinitesimally flexible* (Fig. 1a,  $v_1$  and  $v_3$  can move inward while  $v_2$  and  $v_4$  move outward, while preserving edge lengths) [27]. Intuitively, the concept of rigidity can be thought of in a physical way, that is if the graph were a bar and joint framework, it would be mechanically rigid against external and internal forces.

Clearly the infinitesimal rigidity of  $\mathbb{F}_p$  is intimately tied to the specific embedding of  $\mathcal{G}$  in  $\mathbb{R}^2$ , however it has been shown that the notion of rigidity is a *generic* property of  $\mathcal{G}$ , specifically as *almost all* realizations of a graph are either infinitesimally rigid or flexible (i.e. they form a dense open set in  $\mathbb{R}^2$ ) [37]. Thus, we can treat rigidity from the perspective of  $\mathcal{G}$ , abstracting away the necessity to check every possible realization, and motivating directly our contributions in this work (i.e. it can be shown that (3) constitutes a matrix rank condition, a difficult and costly decentralization). The first such *combinatorial* characterization of graph rigidity was described by Laman in [23], and is summarized as follows (also called *generic rigidity*)<sup>4</sup>:

**Theorem 2.1 (Graph rigidity, [23]):** A graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with realizations in  $\mathbb{R}^2$  having  $n \geq 2$  nodes is rigid if and only if there exists a subset  $\hat{\mathcal{E}} \subseteq \mathcal{E}$  consisting of  $|\hat{\mathcal{E}}| = 2n - 3$  edges satisfying the property that for any non-empty subset  $\hat{\mathcal{E}}' \subseteq \hat{\mathcal{E}}$ , we have  $|\hat{\mathcal{E}}'| \leq 2k - 3$ , where  $k$  is the number of nodes in  $\mathcal{V}$  that are endpoints of  $(i, j) \in \hat{\mathcal{E}}'$ .

<sup>3</sup>We provide a brief overview of rigidity theory here. We direct the reader to [23]–[27] for an in depth review of the subject.

<sup>4</sup>The extension of Laman’s conditions to higher dimensions is at present an unresolved problem in rigidity theory.




Fig. 1. Example graphs demonstrating several embodiments of rigidity, where dashed links indicate edges that have been added to form new networks. Notice that all links in graphs (a), (b), and (d) are independent, while edge  $(v_1, v_3)$  in (c) is redundant.

We refer to the above as the *Laman conditions*, and denote by  $\mathcal{G}_{\mathbb{R}}$  the set of all rigid graphs in  $\mathbb{R}^2$ . We call the graph  $\mathcal{S} = (\mathcal{V}, \hat{\mathcal{E}})$  satisfying Theorem 2.1 a *Laman subgraph* of  $\mathcal{G}$ , where it follows that any rigid graph in the plane must then have  $|\mathcal{E}| \geq 2n - 3$  edges, with equality for *minimally rigid* graphs. The impact of each edge on the rigidity of  $\mathcal{G}$  is captured in the notion of *edge independence*, a direct consequence of Theorem 2.1:

**Definition 2.2 (Edge independence, [28]):** Edges  $(i, j) \in \mathcal{E}$  of a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  are *independent* in  $\mathbb{R}^2$  if and only if no subgraph  $\hat{\mathcal{G}} = (\hat{\mathcal{V}}, \hat{\mathcal{E}})$  has  $|\hat{\mathcal{E}}| > 2|\hat{\mathcal{V}}| - 3$ . A set of independent edges will be denoted by  $\mathcal{E}^*$ , while the graph over  $\mathcal{E}^*$  is denoted by  $\mathcal{G}^*$ .

The above conditions imply that a graph is rigid in  $\mathbb{R}^2$  if and only if it possesses  $|\mathcal{E}^*| = 2n - 3$  independent edges, where edges that do not meet the conditions of Definition 2.2 are called *redundant* (see Fig. 1 for a depiction of graph rigidity). Thus, in determining the rigidity of  $\mathcal{G}$ , we must verify the Laman conditions to discover a suitable set of independent edges  $\mathcal{E}^*$ . In a static sense, by examining  $\mathcal{G}$  as opposed to infinitesimal methods, we gain advantages in complexity ( $O(n^2)$  vs.  $O(n^3)$ ), and amenability to decentralization. Further, when considering dynamic networks with changing topology, checking and controlling the rigidity of  $\mathcal{G}$  must only be performed at times  $t$  where a transition in  $\mathcal{G}(t)$  occurs, as opposed to a continuous check and control required for infinitesimal rigidity (such as in [31]), providing a clear advantage both in complexity and robustness to network failure.

### C. A Pebble Game for Evaluating Generic Rigidity

To lessen the exponential complexity of the Laman conditions we consider the *pebble game* proposed by Jacobs and Hendrickson in [28], which admits a natural form in the context of interconnected systems, with amenability to decentralization.

A brief overview of the centralized pebble game will be given here, beginning with a restatement of the Laman conditions:

*Theorem 2.2 (Laman restated, [28]):* For graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , the following statements are equivalent:

- All  $(i, j) \in \mathcal{E}$  are independent in  $\mathbb{R}^2$ .
- For each  $(i, j) \in \mathcal{E}$ , the graph formed by *quadrupling*  $(i, j)$ , i.e. adding 4 copies of  $(i, j)$  to  $\mathcal{E}$ , has no subgraph  $\bar{\mathcal{G}} = (\bar{\mathcal{V}}, \bar{\mathcal{E}})$  in which  $|\bar{\mathcal{E}}| > 2|\bar{\mathcal{V}}|$ .

Now, applying Theorem 2.2 we have a simple graph operation useful in determining rigidity:

*Lemma 2.1 (Edge quadrupling, [28]):* Given an independent edge set  $\mathcal{E}^*$ , an edge  $(i, j) \notin \mathcal{E}^*$  is independent of  $\mathcal{E}^*$  if and only if the graph formed by quadrupling  $(i, j)$  has no subgraph  $\bar{\mathcal{G}} = (\bar{\mathcal{V}}, \bar{\mathcal{E}})$  in which  $|\bar{\mathcal{E}}| > 2|\bar{\mathcal{V}}|$ .

The generic rigidity test of [28] operates by applying Lemma 2.1 to grow members of  $\mathcal{E}^*$  *incrementally*, terminating when  $|\mathcal{E}^*| = 2n - 3$ . In growing  $\mathcal{E}^*$  it can be shown that every independence check of  $(i, j) \notin \mathcal{E}^*$  requires only  $O(n)$  operations (as there are at most  $2n - 3$  members of  $\mathcal{E}^*$ ) [28], and thus the exponential testing of subgraphs has been reduced to simply quadrupling every new edge and checking that the induced subgraphs have a cardinality respecting Lemma 2.1. A natural simplification to this process is found in the following *pebble game*:

*Definition 2.3 (The pebble game, [28]):* Considering a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , give to each agent associated with  $v_i \in \mathcal{V}$  two pebbles, each of which can be assigned to an edge in  $\mathcal{G}$ . Our goal in the pebble game is to assign the pebbles in  $\mathcal{G}$  such that all edges are covered, i.e. a *pebble covering*. In finding a pebble covering, we allow the assignment of pebbles by agent  $i$  *only* to edges incident to  $v_i$  in  $\mathcal{G}$ . Further we allow pebbles to be rearranged only by removing pebbles from edges which have an adjacent vertex with a free pebble, such that the free pebble is shifted to the assigned pebble, freeing the assigned pebble for assignment elsewhere. Thus, if we consider pebble assignments as directed edges exiting from an assigning agent  $i$ , when a pebble is needed in the network to cover an edge  $(i, j)$ , a *pebble search* over a directed network occurs. If a free (unassigned) pebble is found, the local assignment and rearrangement rules then dictate the pebble's return and assignment to  $(i, j)$ .

*Lemma 2.2 (Pebble covering, [28]):* In the context of the pebble game of Definition 2.3, if there exists a pebble covering for an independent edge set  $\mathcal{E}^*$  with a quadrupled edge  $(i, j) \notin \mathcal{E}^*$ , there is no subgraph violating the conditions of Lemma 2.1, and the set  $\mathcal{E}^* \cup (i, j)$  is independent.

Intuitively, an agent's pebbles represent its possible commitments to the network's subgraphs, while maintaining the conditions of Lemma 2.1, or in a physical way the degrees of freedom of motion in  $\mathbb{R}^2$ . Thus, the rigidity evaluation pebble game of [28] operates as follows: every edge  $e \in \mathcal{E}$  is quadrupled, and an attempt to expand the current pebble covering for  $\mathcal{E}^*$  to each copy of  $e$  is made, with success resulting in  $\mathcal{E}^* \leftarrow \mathcal{E}^* \cup e$  and termination coming when  $|\mathcal{E}^*| = 2n - 3$ . The *centralized* pebble game of Jacobs is depicted in Algorithm 1, with an illustration of the quadrupling and pebble search procedure depicted in Fig. 2.




Fig. 2. An example of the pebble game for a rigid graph with  $n = 3$  with progression from left to right. Pebbles are given by black dots, quadrupled edges by thick links (blue), pebble assignments by dashed arrows, and the assignment of pebbles by black arrows. We have here  $|\mathcal{E}^*| = 3$ .

### III. AN ASYNCHRONOUS DECENTRALIZED PEBBLE GAME

The primary considerations in decentralizing the pebble game of [28] lie in the sequential nature of edge quadrupling, the storage of  $\mathcal{E}^*$  and associated pebble assignments over a distributed network, and the search and rearrangement of pebbles under the asynchronicity of realistic interacting systems. Our proposition in this work is to handle such issues through the election of *leaders* in the network (Section III-A), each expanding  $\mathcal{E}^*$  by examining their incident edges, querying the network for free pebbles, and then exchanging leadership when their local neighborhood has been evaluated. The independent edges and pebble assignments are thus localized to each agent, distributing network storage and relying on inter-agent messaging to support pebble searching under asynchronicity (Section III-C). In this way, our contributions in decentralization will be not only the leader election procedure, but the logic by which each leader must operate to accommodate asynchronicity, and the structure of the pebble game over distributed interacting agents. Intuitively, our proposed leader-based decentralization is effectively an incremental rooting of pebble searches at appropriately elected network leaders. For convenience we will denote by  $\mathbb{S}$  our decentralization of the serial pebble game of Algorithm 1.

To begin, we associate with each agent  $i \in \mathcal{I}$  the following variables (with initialization indicated by  $\leftarrow$ ):

- $\mathcal{P}_i \leftarrow \emptyset$ : Pebble assignment set containing at most two edges  $\{(i, j) \in \mathcal{E} \mid j \in \mathcal{N}_i\}$ , that is incident edges  $(i, j)$  to which a pebble is associated.
- $\mathcal{E}_i^* \leftarrow \emptyset$ : Local independent edge storage, containing edges  $\{(i, j) \in \mathcal{E} \mid j \in \mathcal{N}_i\}$  for which quadrupling and pebble covering succeeds. By construction  $\mathcal{E}^* = \bigcup_i \mathcal{E}_i^*$ .
- $isLeader(i) \in \{0, 1\} \leftarrow 0$ : Current leadership status.
- $beenLeader(i) \in \{0, 1\} \leftarrow 0$ : Prior leadership status.
- $requester(i) \in \mathcal{I} \cup \{0\} \leftarrow 0$ : Stores the agent from which a pebble request has been received.
- $pathsSearched(i) \in \{0, 1, 2\} \leftarrow 0$ : Enumerates the local paths searched along incident edges  $(i, j) \in \mathcal{P}_i$  to which a pebble is assigned.

For convenience, we let  $p_i = 2 - |\mathcal{P}_i| \in \{0, 1, 2\}$  denote agent  $i$ 's free pebble count, and use shorthand notation  $\mathcal{P}_i(k, l)$

**Algorithm 1** The centralized pebble game [28].

---

```

1: procedure PEBBLEGAME( $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ )
2:   Assign each  $v_i$  two pebbles,  $\forall i \in \mathcal{I}$ 
3:    $\mathcal{E}^* \leftarrow \emptyset$ 
4:   for all  $(i, j) \in \mathcal{E}$  do
5:     Quadruple  $(i, j)$  over  $\mathcal{G}$ 
6:     Search for free pebbles, originating from  $v_i$  and  $v_j$ 
7:     if found then
8:       Rearrange pebbles to cover quadrupled  $(i, j)$ 
9:        $\triangleright$  Expand independent set, check rigidity:
10:       $\mathcal{E}^* \leftarrow \mathcal{E}^* \cup (i, j)$ 
11:      if  $|\mathcal{E}^*| = 2|\mathcal{V}| - 3$  then
12:        return  $\mathcal{E}^*$ 
13:      end if
14:    end if
15:  end for
16: end procedure

```

---

for the  $l$ th element of the  $k$ th edge in  $\mathcal{P}_i$ , with  $k, l \in \{1, 2\}$ . The given initialization is executed at the beginning of any evaluation of the  $\mathbb{S}$  algorithm.

#### A. Leader Election

An execution of the  $\mathbb{S}$  algorithm begins with an agent that initiates execution in response to network conditions or mission objectives (e.g. verifying link deletion in  $\mathcal{G}$  for rigidity preservation). The initiating agent begins by triggering an *auction* for electing an agent in the network to become the *lead agent* (or *leader*); an auction also runs after each leader considers all members of  $\mathcal{E}_i$  for quadrupling (detailed in Section III-B). Specifically, to each agent  $i \in \mathcal{I}$  we associate a *bid* for leadership  $r_i = [i, b_i]$  with  $b_i \in \mathbb{R}_{\geq 0}$  indicating the agent's *fitness* in becoming the new leader, with  $b_i = 0$  if  $\text{beenLeader}(i) = 1$  and  $b_i \in \mathbb{R}_+$  otherwise. Denoting the local bid set by  $\mathcal{R}_i = \{r_j \mid j \in \mathcal{N}_i \cup \{i\}\}$ , the leadership auction then operates according to the following agreement process:

$$r_i(t^+) = \underset{r_j \in \mathcal{R}_i}{\operatorname{argmax}}(b_j) \quad (4)$$

where the notation  $t^+$  indicates a transition in  $r_i$  after all neighboring bids have been collected through asynchronous messaging. As  $\mathcal{G}$  is assumed connected for all time, the auction (4) converges *uniformly* to the largest leadership bid

$$r_i = \underset{r_j(0)}{\operatorname{argmax}}(b_j(0)), \quad \forall i, j \in \mathcal{I} \quad (5)$$

after some finite time [38], [39]. After convergence of (4) the winning agent then takes on the leadership role ( $\text{isLeader} \leftarrow 1$ ), with the previous leader relinquishing its status ( $\text{isLeader} \leftarrow 0$ ). The proposed auction mechanism allows us to decentralize the sequential nature of the pebble game by assigning to each leader the responsibility of locally expanding  $\mathcal{E}_i^*$ , noting that previous leaders are never reelected in a given execution due to  $b_i = 0$  for such agents<sup>5</sup>.

<sup>5</sup>Further, the condition  $b_i = 0, \forall i \in \mathcal{I}$  allows termination of the algorithm.

It is important to note that the order of leader election (and thus bid selection) is immediately relevant:

*Proposition 3.1 (Initial leader edges):* All incident edges  $\{(i, j) \in \mathcal{E} \mid j \in \mathcal{N}_i\}$  belonging to an initial leader  $i$  are members of the independent set  $(i, j) \in \mathcal{E}^*$ .

*Proof:* For each edge, a new node  $j \neq i$  must be considered as no two edges of  $i$  can have the same endpoint and  $\mathcal{E}^*$  is empty due to  $i$  being the initial leader. Therefore, every subgraph containing the edges and nodes incident to  $i$  must have  $|\mathcal{E}_s| \leq 2|\mathcal{V}_s| - 3$  edges, where  $\mathcal{V}_s$  are the nodes of the considered subgraph and  $|\mathcal{E}_s| = |\mathcal{V}_s| - 1$  due to the subgraph's implicit tree structure. Thus, as there exists no subgraph violating Definition 2.2, the result follows. ■

*Remark 3.1:* The above reasoning can be extended to subsequent leaders as well, specifically when incident nodes are not endpoints of  $(i, j) \in \mathcal{E}^*$ , independence is clearly implied. It then follows that leader ordering is intimately tied to the composition of  $\mathcal{E}^*$  and thus appropriate bids  $b_i$  can regulate directly the resultant Laman subgraph. Bids can be chosen for each agent to reflect availability, processing capability, or the cardinality of incident edges, or they can be leveraged in terms of metrics related to mission objectives. For example in a sensing regime, bids could reflect the aggregate information value of an agent's incident edges (e.g. [40]), thereby selecting a rigid spanning subgraph with maximal value. The proposed auction technique therefore affords us control over  $\mathcal{E}^*$  that goes beyond simply discovering the network's rigidity property<sup>6</sup>.

The elected leader,  $i \in \mathcal{I}$ , maintains the following variables for managing the local expansion of  $\mathcal{E}_i^*$ :

- $\text{requestWait}(i) \in \{0, 1\} \leftarrow 0$ : Waiting condition when a pebble has been requested.
- $\mathcal{E}_i \leftarrow \{(i, j) \in \mathcal{N}_i \mid \neg \text{beenLeader}(j)\}$ : Contains the incident edges for quadrupling consideration.
- $e_i \in \mathcal{E} \leftarrow (i, j) \in \mathcal{E}_i$ : The current edge being quadrupled.
- $\text{quad}(i) \in \{1, \dots, 4\} \leftarrow 1$ : Edge quadrupling progress for  $e_i$ .

The leader variables are initialized as above, together with  $\text{isLeader}(i) \leftarrow 1$  and  $\text{beenLeader}(i) \leftarrow 1$ , after winning auction (4). In initializing  $\mathcal{E}_i$ , incident edges  $(i, j)$  are considered only when the neighbor  $j \in \mathcal{N}_i$  has not yet been a leader, as edges incident to a previous leader  $j$  have already been checked. Such an initialization guarantees that network edges are considered only *once* for quadrupling and pebble covering. Also, note that each leader receives the current size of the independent edge set  $|\mathcal{E}^*(t)|$  in initialization<sup>7</sup>, allowing termination when  $2n - 3$  independent edges are identified.

#### B. Leader Execution Logic

After election and initialization, the task of the leader  $i$  is to continue the expansion of  $\mathcal{E}^*$  by evaluating the independence of each  $(i, j) \in \mathcal{E}_i$ . The leader executes the procedure LEADERRUN depicted in Algorithm 2 to accomplish this task,

<sup>6</sup>In this paper, we concentrate on the construction of a decentralized pebble game. Our future work focuses on bid selection and related optimality properties.

<sup>7</sup>We can simply embed  $|\mathcal{E}^*(t)|$  in the leadership auction to facilitate this transfer.



where we assume such execution occurs relative to the leader's local clock. First, recall that in checking independence a pebble covering for each quadrupled edge  $e_i \in \mathcal{E}_i$  must be determined. As the pebble information is distributed across the network, the lead agent must therefore *request* pebbles through messaging in an attempt to assign pebbles to  $e_i$  (messaging is treated in Section III-C). In making such requests, the lead agent sets  $requestWait(i) \leftarrow \text{Yes}$ , allowing the agent to wait for pebble responses before continuing; often referred to as *blocking*, lines 2-4.

When there exists no unfulfilled pebble request,  $requestWait(i) = \text{No}$ , the lead agent resumes the quadrupling procedure on the current incident edge  $e_i \in \mathcal{E}_i$ , with progress indicated by  $quad(i)$ . For each step of the quadrupling (lines 7-16), the leader searches for a pebble to cover  $e_i$ , first by looking locally for free pebbles, assigning  $e_i$  to  $\mathcal{P}_i$  if found. If no local pebbles are available,  $p_i = 0, |\mathcal{P}_i| = 2$ , the agent then sends a PEBBLEREQUESTMSG to agent  $\mathcal{P}_i(1, 2)$  requesting a free pebble. Note that in sending requests only along  $(i, j) \in \mathcal{P}_i$ , we properly evaluate independence with respect to  $\mathcal{E}^*$ , as each  $(i, j) \in \mathcal{E}^*$  must have an assigned pebble.

For any edge  $e_i$  with a pebble covering, obtained through a combination of local assignment and pebble responses (more on this later), the following actions are taken (lines 18-32). First, we return 3 pebbles to the endpoints of  $e_i$  leaving a single pebble on  $e_i$  to establish independence, and then add  $e_i$  to  $\mathcal{E}_i^*$ . If in adding  $e_i$ ,  $2n - 3$  independent edges have been identified, the leader sends a simple message to the network indicating that the graph is rigid, and the algorithm terminates. Otherwise, the leader moves to the next member of  $\mathcal{E}_i$ , where we use the notation  $e_i = \mathbf{0}$  to indicate  $|\mathcal{E}_i| = 0$ . When all members of  $\mathcal{E}_i$  have been evaluated, the leader initiates the auction (4), embedding  $|\mathcal{E}^*|$  to the next leader. The process of edge quadrupling, pebble requests and covering, and independence expansion then continues from leader to leader until either the network is found to be rigid, or  $beenLeader(i) = 1, \forall i \in \mathcal{I}$ , indicating non-rigidity.

### C. Inter-Agent Messaging

As each leader attempts to expand  $\mathcal{E}_i^*$  through quadrupling, free pebbles are needed to establish a pebble covering. We accommodate the pebble search by defining asynchronous message PEBBLEREQUESTMSG, accompanied by response messages PEBBLEFOUNDMSG and PEBBLENOTFOUNDMSG, indicating the existence of free pebbles. Note that no knowledge of the global topology is held by any single agent, only local neighbors, with no explicit routing. The arrival of these messages is asynchronous at the recipient and triggers message handlers that form the foundation of the pebble search mechanism, as described in the following.

The reception of a PEBBLEREQUESTMSG initiates the handler HANDLEPEBBLEREQUEST depicted in Algorithm 3. Each pebble request is marked with a unique identifier, e.g.  $uID \in \mathbb{R}$ , originating from the lead agent, defining the pebble search to which the request is a member (lines 2-6). This allows for any given agent to participate simultaneously in multiple searches<sup>8</sup>,

<sup>8</sup>In this work, we do not explicitly investigate such a possibility. This is the focus of future work.

### Algorithm 2 Leader execution logic.

---

```

1: procedure LEADERRUN( $i$ )
2:   if  $requestWait(i)$  then           ▷ Wait for pebble response
3:     return
4:   end if
5:   while  $e_i \triangleq (i, j) \neq \mathbf{0}$  do           ▷ Continue covering
6:     while  $quad(i) \leq 4$  do
7:       if  $p_i > 0$  then           ▷ Assign local pebble
8:          $\mathcal{P}_i \leftarrow \mathcal{P}_i \cup e_i$ 
9:          $p_i \leftarrow p_i - 1$ 
10:         $quad(i) \leftarrow quad(i) + 1$ 
11:       else           ▷ Request pebble along first edge
12:         PEBBLEREQUESTMSG( $i, \mathcal{P}_i(1, 2)$ )
13:          $requestWait(i) \leftarrow \text{Yes}$ 
14:          $pathsSearched(i) \leftarrow 1$ 
15:       return
16:     end if
17:   end while
18:   ▷ Quadrupling success, return 3 pebbles:
19:    $\mathcal{P}_i \leftarrow \emptyset$ 
20:    $p_i \leftarrow 2$ 
21:   Return 1 pebble to  $e_i(2)$ 
22:   ▷ Add independent edge and check rigidity:
23:    $\mathcal{E}_i^* \leftarrow \mathcal{E}_i^* \cup e_i$ 
24:   if  $|\mathcal{E}^*| = 2n - 3$  then
25:     Send network rigidity notification
26:   return
27:   end if
28:   ▷ Go to next incident edge:
29:    $\mathcal{E}_i \leftarrow \mathcal{E}_i - e_i$ 
30:    $e_i \leftarrow (j, k) \in \mathcal{E}_i$ 
31:    $quad(i) \leftarrow 1$ 
32: end while
33:   ▷ All local edges checked:
34:   Initiate leadership auction
35: end procedure

```

---

and guarantees that each search does not revisit nodes, guaranteeing termination (i.e. the leader *will* receive a response). For unique requests, the agent first attempts to assign local pebbles to the edge connecting the pebble requester, sending a PEBBLEFOUNDMSG in response, allowing the requester to free an assigned pebble for either local assignment (a leader) or to itself respond to a pebble request (non-leader). If  $p_i = 0$  (lines 11-14) the agent forwards the request to agent  $\mathcal{P}_i(1, 2)$ , recording the requester such that pebble responses can be properly returned<sup>9</sup>.

When the response PEBBLEFOUNDMSG to a pebble request is received it triggers the handler HANDLEPEBBLEREQUEST depicted in Algorithm 4. Congruent to the shifting action of HANDLEPEBBLEREQUEST, the agent first frees the local pebble assigned to the edge connecting the responder (line 2), and applies the pebble relative to leader status. If  $isLeader(i) = 1$  (lines 3-6), the freed pebble is assigned locally to  $e_i$ , continuing

<sup>9</sup>Notice that as opposed to explicit message routing, the local nature of agent interaction is preserved by tracking requesters in this manner.

**Algorithm 3** Pebble request handler for agent  $i$ .

---

```

1: procedure HANDLEPEBBLEREQUEST( $from, i$ )
2:   if Already Requested then
3:     PEBBLENOTFOUNDMSG( $i, from$ )
4:     return
5:   end if
6:   if  $p_i > 0$  then ▷ Local pebble available
7:      $\mathcal{P}_i \leftarrow \mathcal{P}_i \cup (i, from)$  ▷ Shift free pebble
8:      $p_i \leftarrow p_i - 1$ 
9:     PEBBLEFOUNDMSG( $i, from$ )
10:  else ▷ Request along first assigned edge
11:    PEBBLEREQUESTMSG( $i, \mathcal{P}_i(1, 2)$ )
12:     $pathsSearched(i) \leftarrow 1$ 
13:     $requester(i) \leftarrow from$ 
14:  end if
15: end procedure

```

---

**Algorithm 4** Pebble found handler for agent  $i$ .

---

```

1: procedure HANDLEPEBBLEFOUND( $from, i$ )
2:    $\mathcal{P}_i \leftarrow \mathcal{P}_i - (i, from)$  ▷ Free local pebble
3:   if isLeader( $i$ ) then
4:      $\mathcal{P}_i \leftarrow \mathcal{P}_i \cup e_i$  ▷ Expand covering
5:      $quad(i) \leftarrow quad(i) + 1$  ▷ Go to next edge
6:      $requestWait(i) \leftarrow \text{No}$ 
7:   else ▷ Give free pebble to requester
8:      $\mathcal{P}_i \leftarrow \mathcal{P}_i \cup (i, requester(i))$ 
9:     PEBBLEFOUNDMSG( $i, requester(i)$ )
10:  end if
11: end procedure

```

---

the edge quadrupling process and relieving the request blocking condition. Otherwise, the non-lead agent performs a pebble shift to again free a pebble for the requesting agent, indicating the shift by returning a PEBBLEFOUNDMSG to the requester (lines 7-9).

Finally, the response PEBBLENOTFOUNDMSG to a pebble request initiates the handler HANDLEPEBBLENOTFOUND depicted in Algorithm 5. For both leaders and non-leaders the non-existence of a free pebble initiates a further search in the network, along  $\mathcal{P}_i(2, 2)$  (lines 2-4). However, if both available search paths are exhausted (lines 6-13), the leadership status of the receiver dictates the action taken. In the simple case of a non-leader (line 13) the response is simply returned to the original requester in order to initiate further search. For the leader, the lack of free pebbles in the network indicates precisely that the conditions of Lemma 2.1 do not hold, implying the currently considered edge  $e_i$  is *redundant*. The edge  $e_i$  is removed from consideration by returning all pebbles assigned during the covering attempt to the endpoints of  $e_i$ , and the process is moved to the next incident edge (lines 7-11). A basic illustration of the  $\mathbb{S}$  algorithm is given in Fig. 3.

**D.  $\mathbb{S}$  Correctness, Termination, and Cost Analysis**

To conclude our discussion of the  $\mathbb{S}$  algorithm, we analyze the correctness, finite termination, and cost properties of our

**Algorithm 5** Pebble not found handler for agent  $i$ .

---

```

1: procedure HANDLEPEBBLENOTFOUND( $from, i$ )
2:   if  $pathsSearched(i) < 2$  then ▷ Search other path
3:     PEBBLEREQUESTMSG( $i, \mathcal{P}_i(2, 2), uID$ )
4:      $pathsSearched(i) \leftarrow 2$ 
5:   else ▷ Search failed, no free pebbles
6:     if isLeader( $i$ ) then ▷  $e_i$  is redundant
7:       Return pebbles assigned to  $e_i$ 
8:       ▷ Go to next incident edge:
9:        $\mathcal{E}_i \leftarrow \mathcal{E}_i - e_i$ 
10:       $e_i \leftarrow (i, j) \in \mathcal{E}_i$ 
11:       $quad(i) \leftarrow 1$ 
12:       $requestWait(i) \leftarrow \text{No}$ 
13:    else
14:      PEBBLENOTFOUNDMSG( $i, requester(i)$ )
15:    end if
16:  end if
17: end procedure

```

---

propositions. First, we formally establish the stopping condition for the  $\mathbb{S}$  algorithm:

*Definition 3.1 ( $\mathbb{S}$  stopping condition):* As previously discussed, the  $\mathbb{S}$  algorithm terminates upon satisfaction of the following condition:

$$f_{stop}^{\mathbb{S}} \triangleq \left\{ \left( \sum_{i=1}^n b_i = 0 \right) \vee \left( \sum_{i=1}^n |\mathcal{E}_i^*| = 2n - 3 \right) \right\} \quad (6)$$

where the  $\sum_i b_i = 0$  indicates that all agents have been a leader (easily detected during auction (4), and  $\sum_i |\mathcal{E}_i^*| = 2n - 3$  is detected by the lead agent on line 24 of Algorithm 2.

Next, we verify that our formulation guarantees the *entire* network is evaluated for rigidity, with *no* edge reconsideration (i.e. no edge is checked by multiple agents), and further that mutual exclusion of local independent sets holds:

*Proposition 3.2 (Edge consideration, mutual exclusion):*

Disregarding algorithm termination when  $|\mathcal{E}^*| = 2n - 3$ , every  $(i, j) \in \mathcal{E}$  is eligible to be considered for independence. Further,  $\mathcal{E}_i^* \cap \mathcal{E}_j^* = \emptyset$  holds for all  $i \neq j \in \mathcal{I}$ .

*Proof:* These results are a simple consequence of the guaranteed convergence of auction (4),  $b_i = 0$  for all *beenLeader*( $i$ ) = 1 (guaranteeing no reelection), and the initialization of  $\mathcal{E}_i$  with edges *not* shared with previous leaders. ■

To ensure timely results, we must also have finite termination of  $\mathbb{S}$ :

*Proposition 3.3 ( $\mathbb{S}$  termination):* Consider the execution of the  $\mathbb{S}$  algorithm as described in Sections III-A, III-B, and III-C. By construction, it follows that the stopping condition  $f_{stop}^{\mathbb{S}}$  of (6) is satisfied in *finite* time.

*Proof:* We can guarantee no message-induced race conditions by Assumptions 2 and 3, and that there exist no algorithmic race conditions due to the internal blocking on line 2 of Algorithm 2. From the request checking mechanism (line 2 Algorithm 3) and the guaranteed delivery of inter-agent messages by the best-effort Assumption 3, we have that all pebble request messaging rooted at agent  $i$  with *isLeader*( $i$ ) = 1, is finite, i.e. every




Fig. 3. Illustration of the first two pebble covering attempts for a quadrupling on edge (3, 4). Agents  $v_1$  and  $v_2$  have previously been a leader, while agent  $v_3$  (blue) is the current leader. Pebbles are depicted by solid black dots, requests by inter-agent arrows, and responses and shifts by dashed arrows.

pebble request receives a response. Now, the finiteness of execution is a direct consequence of the finiteness of each  $\mathcal{E}_i$ ,  $\forall i \in \mathcal{I}$  and the finite convergence of auction (4) [39], as there exists no leader reelection by construction. ■

Now we come to our primary result concerning the correctness of the  $\mathbb{S}$  algorithm:

**Proposition 3.4 ( $\mathbb{S}$  correctness):** Consider an execution of  $\mathbb{S}$ , applied to a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ . It follows that by construction we are guaranteed to identify  $|\mathcal{E}^*| = 2n - 3$  independent edges (c.f. Definition 2.2) when  $\mathcal{G} \in \mathcal{G}_{\mathbb{R}}$ , and  $|\mathcal{E}^*| < 2n - 3$  otherwise, i.e.  $\mathbb{S}$  properly identifies (without false positives) the generic rigidity of  $\mathcal{G}$ .

*Proof:* First, notice that by Proposition 3.2, we can ensure that every  $(i, j) \in \mathcal{E}$  is eligible for quadrupling and pebble covering as dictated by the original pebble game [28], and further that  $\mathcal{E}_i^* \cap \mathcal{E}_j^* = \emptyset$  holds for all  $i \neq j \in \mathcal{I}$  ensures that  $|\mathcal{E}^*|$  is properly tracked by our distributed storage (such that the rigidity condition is correctly identified). Thus, correctness is shown by arguing that our leadership and messaging formulation is faithful to the rules of the pebble game (c.f. Definition 2.3). This result follows by observing that pebble assignments and shift operations are only made locally (line 8 of Algorithm 2, line 7 of Algorithm 3, and lines 2, 4, and 8 of Algorithm 4), and that the pebble search mechanism respects the network’s distributed pebble assignments  $\mathcal{P}_i \forall i \in \mathcal{I}$  (line 12 of Algorithm 2, line 11 of Algorithm 3, and line 3 of Algorithm 5). Finally, as there is only one leader active at any time, each quadrupling operation (lines 6-17, Algorithm 2) is sound with respect to the current set  $\mathcal{E}^*$ , and the result follows. ■

The above result demonstrates that  $\mathbb{S}$  is sound in terms of planar rigidity evaluation. The complexity of  $\mathbb{S}$  is promising for realistic decentralized operation:

**Proposition 3.5 ( $\mathbb{S}$  complexity):** By construction, executions of the  $\mathbb{S}$  algorithm have *worst-case*  $O(n^2)$  messaging complexity (and thus execution complexity), and  $O(n)$  storage scaling.

*Proof:* As the pebble game exhibits  $O(n^2)$  complexity [28], our pebble messaging scales like  $O(n^2)$ . In applying leader auction (4) we incur  $O(n^2)$  as we expend  $O(n)$  auction messaging for  $O(n)$  leaders<sup>10</sup>. The per-agent storage complexity scales like

<sup>10</sup>Of course, complexity can be mitigated by choosing an  $O(1)$  leader election, such as ordering by label  $i \in \mathcal{I}$ , sacrificing control over  $\mathcal{E}^*$ .

$O(n)$ , the maximal cardinality of  $\mathcal{N}_i$ , as assignments to  $\mathcal{E}_i$  occur over only edges incident to  $i$ , line 23 Algorithm 2. ■

The analysis of this section demonstrates that our proposed  $\mathbb{S}$  algorithm represents a fully decentralized and efficient solution to the planar generic rigidity evaluation problem, providing the opportunity to exploit the vast advantages of network rigidity in realistic robotic networks.

#### IV. EXPLOITING STRUCTURE TOWARDS PARALLELIZATION

Although the proposed decentralized pebble game provides a unique method of evaluating the topological rigidity of a distributed network, the fundamental execution remains serial in nature, i.e. ultimately each  $(i, j) \in \mathcal{E}$  is considered individually and the pebble search messaging involves at any time step communication between only a single pair of neighbors in  $\mathcal{G}$ . Thus, to fully exploit a distributed multi-agent system, we seek a parallelization of the algorithm proposed in Section III, such that edge consideration occurs in simultaneity, with the goal of reducing the overall execution time of the rigidity evaluation, and rendering real-world application feasible. As is required in parallelization schemes, our approach in this work focuses on identifying elements of the rigidity problem that are approachable *without* globally inspecting  $\mathcal{G}$ , i.e. we identify local, independence preserving operations.

##### A. Independence Preserving Operations

Let us begin by formally defining addition and subtraction operations for graph edges as follows.

**Definition 4.1 (Edge addition/subtraction [10]):** Consider a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  and let the graph augmented with an edge  $e$  be denoted  $\mathcal{G}^+ = (\mathcal{V}, \mathcal{E} \cup \{e\})$ . Similarly, the graph  $\mathcal{G}$  with  $e$  removed is denoted by  $\mathcal{G}^- = (\mathcal{V}, \mathcal{E} \setminus \{e\})$ . We refer to the operation  $[\cdot]_e^+$  such that  $\mathcal{G}^+ = [\mathcal{G}]_e^+$  as *edge addition*. Likewise, the operation  $[\cdot]_e^-$  such that  $\mathcal{G}^- = [\mathcal{G}]_e^-$  as *edge subtraction* (or deletion).

Now we are prepared to consider *independence preserving* graph operations. Specifically:

**Definition 4.2 (Independence preserving operations):** We call the edge operations  $[\cdot]_e^+$  and  $[\cdot]_e^-$  over graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  having independent edges  $\mathcal{E}$  satisfying Definition 2.2 *independence preserving* (or an IP operation) if  $\mathcal{E} \cup \{e\}$  and  $\mathcal{E} \setminus \{e\}$  are themselves independent, respectively. Clearly, all operations  $[\cdot]_e^-$  over independent edges  $\mathcal{E}$  are independence preserving. Also, we have that addition operations  $[\cdot]_e^+$  that are *not* independence preserving imply  $e$  is redundant with respect to  $\mathcal{E}$ .

Much in the vein of the original pebble game we thus seek IP edge addition operations that are readily defined by local-only agent interaction, enabling the construction of  $\mathcal{E}^*$  in a parallel fashion. Then, given an initial independent set  $\mathcal{E}^*(0) = \emptyset$  with associated graph  $\mathcal{G}^* = (\mathcal{V}, \mathcal{E}^*)$ , we can generate the sequence

$$\mathcal{G}^*(0) = (\mathcal{V}, \emptyset) \quad \mathcal{G}^*(k) = [\mathcal{G}^*(k-1)]_e^+, \quad k = 1, \dots, m \quad (7)$$

where if each  $[\cdot]_e^+$  is independence preserving, the resulting graph  $\mathcal{G}^*(m)$  possesses independent edges, and if  $[\cdot]_e^+$  is local to the endpoints of edge  $e$ , such a sequence can be achieved by parallel application of each operation.






Fig. 4. Illustration of a sequence of the endpoint expansion rule (a) and the two incident edge rule (b). Edge addition operations to  $\mathcal{E}^*(k)$  are given by red links, endpoint membership in  $\mathcal{G}^*(k+1)$  is depicted by blue nodes, and  $\mathcal{E}^*(k)$  is shown by black links. Notice in the illustrated minimally rigid graph, a combination of EER and TIER operations identifies fully the independent edge set  $\mathcal{E}^*$ .

Taking inspiration from the minimal rigidity construction of the *Henneberg operations*<sup>11</sup> [27], edge addition rules will be defined that require for a considered edge  $(i, j) \in \mathcal{E}$  and  $(i, j) \notin \mathcal{E}^*$ , the inspection only of graph properties local to  $v_i$  and  $v_j$ . We begin with a simple rule based on the membership of  $v_i$  or  $v_j$  as endpoints in  $\mathcal{G}^*(k)$ :

**Definition 4.3 (Endpoint expansion rule):** Consider the graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  and the associated node set  $\hat{\mathcal{V}} = \{v_i \in \mathcal{V} \mid \exists j \in \mathcal{I}, (i, j) \in \mathcal{E}\}$ , containing the nodes in  $\mathcal{V}$  which are endpoints of edges in  $\mathcal{E}$ . Define the *endpoint expansion rule* (EER) as the edge addition operation  $[\mathcal{G}]_e^+$  possessing the property that  $|\hat{\mathcal{V}}_e^+| > |\hat{\mathcal{V}}|$ , where  $[\hat{\mathcal{V}}]_e^+$  are the endpoints of  $\mathcal{G}^+$ . Notice that for an EER operation it trivially holds that  $2 \geq |[\hat{\mathcal{V}}]_e^+| - |\hat{\mathcal{V}}| \geq 1$ .

Clearly the endpoint expansion rule 4.3 is limiting in terms of the identified set  $\mathcal{E}^*$ , specifically as the identified set can be described as the union of local spanning trees over  $\mathcal{G}$  (a direct consequence of expanding endpoints in a graph, e.g. Fig. 4a). Thus, we can further consider a two edge rule that is also independence preserving:

**Definition 4.4 (Two incident edge rule):** Consider the graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  and the augmented graph  $\mathcal{G}^+$  through addition of edge  $e \triangleq (i, j)$ . The *two incident edge rule* (TIER) is an edge addition operation  $[\mathcal{G}]_e^+$  where it holds that  $(\mathcal{N}_i^+ \leq 2) \vee (\mathcal{N}_j^+ \leq 2)$  over  $\mathcal{G}^+$ , with  $(\mathcal{N}_i \geq 1) \wedge (\mathcal{N}_j \geq 1)$  over  $\mathcal{G}$  (otherwise the edge operation would constitute an endpoint expansion).

We illustrate the rules of Definitions 4.3 and 4.4 in Figs. 4a and 4b, respectively. Notice that as expected each edge addition operation requires only local information to implement. The independence preservation of the proposed rules is shown below.

<sup>11</sup>A sequence of Henneberg operations generates a minimally rigid graph as the operations preserve minimal rigidity. Each operation itself is a node/edge addition to the previous graph, yielding an iterative construction.

**Proposition 4.1 (Independence preservation):** Consider the graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  having edges  $\mathcal{E}$  forming an independent set according to Definition 2.2. The edge addition operations  $[\cdot]_e^+$  over  $\mathcal{G}$  abiding by the EER and TIER requirements of Definitions 4.3 and 4.4 are independence preserving in the sense of Definition 4.2, respectively. Further, considering a sequence of graphs  $\{\mathcal{G}(0), \dots, \mathcal{G}(m)\}$  generated by

$$\mathcal{G}(0) = \mathcal{G} \quad \mathcal{G}(k) = [\mathcal{G}(k-1)]_e^+, \quad k = 1, \dots, m \quad (8)$$

over EER and TIER operations yields graph  $\mathcal{G}^m$  having independent edges  $\mathcal{E}^m$ .

*Proof:* First, consider the case of an EER operation over  $\mathcal{G}$ . By the independence of  $\mathcal{E}$ , we have by Definition 2.2 that for every subgraph  $\bar{\mathcal{G}} = (\bar{\mathcal{V}}, \bar{\mathcal{E}})$ ,  $|\bar{\mathcal{E}}| \leq 2|\bar{\mathcal{V}}| - 3$ . In the augmented graph edge  $e$  introduces expanded subgraphs containing  $e$  all having the property

$$|\bar{\mathcal{E}}| + 1 \leq 2|\bar{\mathcal{V}}| - 3 + 1 \leq 2(|\bar{\mathcal{V}}| + 1) - 3 \quad (9)$$

due to the node expansion property of the EER, all of which therefore abide by the independence subgraph property. The remaining subgraphs of  $\mathcal{G}^+$  are independent by assumption. Thus, we conclude that the EER operation according to Definition 4.3 is independence preserving.

Now, consider the application of the TIER operation over  $\mathcal{G}$ . By Definition 4.4 there must exist an endpoint of  $e \triangleq (i, j)$ , indexed by  $i \in \mathcal{I}$ , with exactly  $\mathcal{N}_i = 1$  over  $\mathcal{G}$ . Thus, we can view the edges  $(i, j)$  and  $(i, k)$  with  $k \in \mathcal{N}_i$ , and the node  $i$  as members of a two edge Henneberg operation, as in Section 3 of [27] (e.g. adding vertex  $v_4$  with edges (4,2) and (4,1) in Fig. 4b). As the edge subtraction operation  $[\cdot]_e^-$  is independence preserving, the graph  $\mathcal{G}^+$  described by applying the previous two edge Henneberg operation to  $[\mathcal{G}]_{(i,k)}^-$  has independent edges by Proposition 3.1 of [27]<sup>12</sup>. Thus, the TIER operation is also independence preserving.

Finally, the independence preservation of a sequence of EER and TIER operations is a trivial consequence of the initial independence of  $\mathcal{E}$  and the IP properties of each edge augmentation. ■

The above result implies that we can apply EER and TIER operations to build an independent set incrementally as in (7), much like the original pebble game. However, notice that the operations take on a distinctly local nature, making them amenable to distributed and parallel implementation.

### B. Gossip-like Messaging for Parallelization

As the proposed independence preserving operations require only information local to neighbors  $i$  and  $j$ , and as our network operates under the assumption of asynchronicity, we take inspiration from the *communication* scheme typical of *gossip* algorithms (see [41], [42]) to define the execution and messaging structure for partial parallelization of rigidity evaluation. As opposed to the serialized execution of the algorithm presented in Section III (i.e. there is only a single leader from which messages originate), here we assume all agents execute simultaneously according to ticks of their own local clocks. Each agent  $i$  exchanges

<sup>12</sup>Briefly, this result follows from the independence of  $\mathcal{E}$  and the relationships,  $|\mathcal{E}| = |\mathcal{E}^+| - 2$  and  $|\mathcal{V}| = |\mathcal{V}^+| - 1$ .

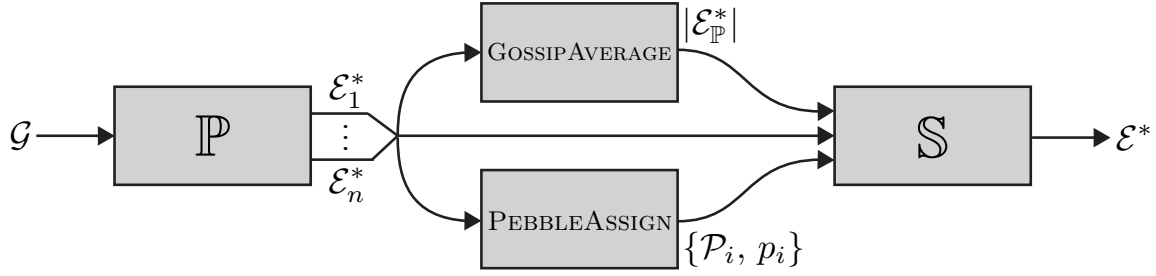


Fig. 5. Block diagram of the parallelized pebble game algorithm for decentralized planar rigidity evaluation. The parallel phase algorithm  $\mathbb{P}$  acts as input for the serial phase  $\mathbb{S}$ , through localized pebble assignment and a gossip-based averaging.

inter-neighbor messages in an attempt to assign incident edges  $(i, j)$ ,  $\forall j \in \mathcal{N}_i$  to  $\mathcal{E}_i^*$  according to the EER and TIER rules. As the messaging is purely local (i.e. no pebble searches) and concurrent, such a construction, denoted as algorithm  $\mathbb{P}$ , allows the network to determine a subset of independent edges  $\mathcal{E}_{\mathbb{P}}^* \subseteq \mathcal{E}^*$  with significantly reduced execution time and messaging, as will be verified in Section V-B.

To begin, the parallel phase algorithm  $\mathbb{P}$  requires three additional variables stored locally by each agent  $i$ , described and initialized ( $\leftarrow$ ) as follows:

- $committed(i) \in \mathbb{Z}_{\geq 0} \leftarrow 0$ : The cardinality of  $i$  as an endpoint of edges in the distributed independent edge set, or more formally, the node degree of  $v_i$  in the graph  $\mathcal{G}^* = (\mathcal{V}, \cup_i \mathcal{E}_i^*)$ . Given the asynchronicity of agent execution and messaging, this local variable operates similar to a traditional *semaphore* in computer science, guaranteeing each agent's commitment to the independent set follows the EER and TIER operation requirements.
- $requestedFrom(i) \in \mathcal{I} \cup \{0\} \leftarrow 0$ : The index  $j \in \mathcal{I}$  of the latest agent to which a `EDGEREQUESTMSG`( $i, j$ ) was sent, where 0 represents no messages having been sent. Such bookkeeping allows conflict resolution when neighbors compete over a shared edge.
- $stop(i) \in \{0, 1\} \leftarrow 0$ : A stopping condition indicating all  $(i, j) \in \mathcal{E}_i$  have been exhausted, that is the local independent edge search has completed for agent  $i$ .

Again, in opposition to the leader-based execution of algorithm  $\mathbb{S}$ , here every agent executes concurrently according to Algorithm 6, with initialization

$$\mathcal{E}_i(0) \leftarrow \mathcal{N}_i, \quad e_i \leftarrow (i, j) \in \mathcal{E}_i(0) \quad \forall i \in \mathcal{I} \quad (10)$$

where we use notation  $\leftarrow$  to represent a random assignment of a link  $(i, j) \in \mathcal{E}_i$ . Each agent attempts to expand its local independence set  $\mathcal{E}_i^*$  by exploiting messaging to determine a neighbor  $j$ 's feasibility as an endpoint in  $\mathcal{G}^*(k)$ , ensuring that the EER and TIER conditions are fulfilled. Thus, for each edge  $e_i \triangleq (i, j)$  chosen randomly from  $\mathcal{E}_i$ , agent  $i$  requests neighbor  $j$ 's commitment to  $\mathcal{G}^*$  through message `EDGEREQUESTMSG`( $i, j$ ) (Algorithm 6, lines 5-10) and blocks further edge consideration (Algorithm 6, lines 2-4), similar to the logic described in Section III-B. When all edges in  $\mathcal{E}_i$  have been considered (Algorithm 6, line 12), the stopping condition  $stop(i)$  is set and the agent enters into the idle state (continuing to handle messages) until

the network stopping condition

$$f_{stop}^{\mathbb{P}} \in \{0, 1\} \triangleq \left( \bigwedge_i stop(i) = 1 \right) \quad (11)$$

is satisfied. The remaining logic for independent edge selection and rule checking resides in the message handlers related to `EDGEREQUESTMSG`( $i, j$ ), as will be discussed in the following.

The reception of an edge request message by agent  $i$  from neighbor  $j \in \mathcal{N}_i$  triggers the message handling logic `HANDLEEDGEREQUEST`( $i, j$ ) depicted in Algorithm 7. Recall that when an agent  $i$  receives such a request from  $j$  it indicates an attempt by  $j$  to make the assignment  $\mathcal{E}_j(t+1) = \mathcal{E}_j(t) \cup (j, i)$ . Requests for an edge assignment must thus be tested for fulfillment of the EER and TIER rules. First, receiving agent  $i$  determines whether the request represents a conflict over edge  $(i, j)$ , due to  $i$ 's current interest in the assignment  $\mathcal{E}_i(t+1) = \mathcal{E}_i(t) \cup (i, j)$ , i.e.  $requestedFrom(i) = j$ , line 3 of Algorithm 7. If there is a conflict over  $(i, j)$ , agents  $i$  and  $j$  resolve the conflict as discussed in the following:

*Remark 4.1 (Conflict resolution)*: We define an *edge conflict* in the execution of algorithm  $\mathbb{P}$  as the state in which an agent  $i \in \mathcal{I}$  has received an `EDGEREQUESTMSG`( $i, j$ ) from  $j \in \mathcal{N}_i$ , where for  $i$  it holds that  $requestedFrom(i) = j$ . Intuitively, this state indicates that both  $i$  and  $j$  are attempting to add edge  $(i, j)$  to  $\mathcal{E}^*$ , a condition that would introduce inconsistencies in  $|\mathcal{E}^*|$ . Thus, in such scenarios we assume that  $i$  and  $j$  exploit local communication to resolve conflicts such that for the function `RESOLVECONFLICT`( $i, j$ ) :  $\mathcal{E} \rightarrow \mathcal{I}$  indicating the conflict *winner*, it holds that

$$\text{RESOLVECONFLICT}(i, j) = \text{RESOLVECONFLICT}(j, i) \quad (12)$$

for all  $i \neq j \in \mathcal{I}$ . In achieving condition (12), we could consider various options, e.g. balancing  $|\mathcal{E}_i^*|$  and  $|\mathcal{E}_j^*|$  in choosing the winner, applying a simple predetermined condition such as agent label, or perhaps a more complex method such as considering optimal assignments based on some cost or utility function over  $(i, j)$ .

Following Remark 4.1, in lines 4-6 of Algorithm 7, conflicts are resolved simply by artificially saturating the commitments of the winning agent, such that the edge request cannot be fulfilled due to violation of the EER and TIER rules. After properly handling potential conflicts over  $(i, j)$ , receiver  $i$  then simply responds to  $j$  via `EDGERESPONSEMSG`( $i, j, response$ ) indicating its current commitment to  $\mathcal{E}^*$ , increases its own commitment

**Algorithm 6** Parallel execution logic for agent  $i$ .

---

```

1: procedure PARALLELRUN( $i$ )
2:   if  $requestWait(i)$  then      ▷ Wait for request response
3:     return
4:   end if
5:   if  $e_i \triangleq (i, j) \neq \mathbf{0}$  then      ▷ Next incident edge
6:     EDGEREQUESTMSG( $i, j$ )
7:      $requestWait(i) \leftarrow \text{Yes}$ 
8:      $requestedFrom(i) \leftarrow j$ 
9:     return
10:  end if
11:  ▷ All local edges checked:
12:   $stop(i) \leftarrow \text{Yes}$ 
13: end procedure

```

---

**Algorithm 7** Parallel edge request handler for agent  $i$ .

---

```

1: procedure HANDLEEDGEREQUEST( $i, j$ )
2:    $response \leftarrow committed(i)$ 
3:   if  $requestedFrom(i) = j$  then      ▷ Edge contention
4:     if  $RESOLVECONFLICT(i, j) = i$  then
5:        $response \leftarrow 2$       ▷ Ensure  $i$  wins edge
6:     end if
7:   end if
8:   EDGERESPONSEMSG( $i, j, response$ )
9:   if  $response < 2$  then      ▷ Max of 2 incident edges
10:     $committed(i) \leftarrow committed(i) + 1$ 
11:  end if
12:  ▷ Do not double check  $(i, j)$ :
13:   $\mathcal{E}_i \leftarrow \mathcal{E}_i - (i, j)$ 
14: end procedure

```

---

count to guarantee independence preservation, and removes  $(i, j)$  from  $\mathcal{E}_i$  to avoid double checking.

In coherence with the edge request logic, in receiving an EDGERESPONSEMSG( $i, j, response$ ), an agent  $i$  acts according to Algorithm 8. Again, as  $response$  indicates agent  $j$ 's commitment to  $\mathcal{E}^*$ , it is validated against the EER and TIER rules (line 2, Algorithm 8), with success resulting in the assignment of  $(i, j)$  to  $\mathcal{E}_i$  and an incrementing of  $committed(i)$  as  $i$  is now an endpoint of  $\mathcal{G}^*$ . After independence preserving assignment, agent  $i$  simply moves to consider its next incident edge by choosing a random  $(i, j) \in \mathcal{E}_i$ , lines 7-9 of Algorithm 8.

Finally, at the conclusion of algorithm  $\mathbb{P}$  consisting of local execution Algorithm 6 and messaging Algorithms 7 and 8, a distributed independent edge set  $\mathcal{E}_{\mathbb{P}}^*$  has been computed across each  $\mathcal{E}_i^*$ . However, by the EER and TIER definitions, it must hold generally that  $\mathcal{E}_{\mathbb{P}}^* \subseteq \mathcal{E}^*$ . Thus, to fully determine  $\mathcal{E}^*$  we generate a composite algorithm by passing the terminal state of  $\mathbb{P}$  to the serialized algorithm  $\mathbb{S}$  in order to apply global network information in completing  $\mathcal{E}^*$ . To derive a valid initial condition for  $\mathbb{S}$ , we first apply a network averaging algorithm, such as gossip averaging [41], to yield  $|\mathcal{E}_{\mathbb{P}}^*|$  (i.e. multiplying by  $n$  the network average of  $|\mathcal{E}_i^*| \forall i \in \mathcal{I}$ ). Then, local pebble assignments are determined by applying Algorithm 9 to yield  $\mathcal{P}_i, p_i \forall i \in \mathcal{I}$ . This composite construction, which we will denote  $\mathbb{P} + \mathbb{S}$ , is illustrated in Fig. 5.

**Algorithm 8** Parallel edge response handler for agent  $i$ .

---

```

1: procedure HANDLEEDGERESPONSE( $i, j, response$ )
2:   if  $response < 2$  then      ▷ Independence guaranteed
3:      $\mathcal{E}_i^* \leftarrow \mathcal{E}_i^* \cup (i, j)$ 
4:      $committed(i) \leftarrow committed(i) + 1$ 
5:   end if
6:   ▷ Go to next incident edge:
7:    $\mathcal{E}_i \leftarrow \mathcal{E}_i - (i, j)$ 
8:    $e_i \leftarrow (i, j) \in \mathcal{E}_i$ 
9:    $requestWait(i) \leftarrow \text{No}$ 
10: end procedure

```

---

**Algorithm 9** Local pebble assignment for agent  $i$ .

---

```

1: procedure PEBBLEASSIGN( $i$ )
2:   for  $(i, j) \in \mathcal{E}_i^*$  do
3:     if  $p_i > 0$  then      ▷ Free local pebble
4:        $\mathcal{P}_i \leftarrow \mathcal{P}_i \cup (i, j)$ 
5:        $p_i \leftarrow p_i - 1$ 
6:     else      ▷ Use neighbor pebble
7:        $\mathcal{P}_j \leftarrow \mathcal{P}_j \cup (i, j)$ 
8:        $p_j \leftarrow p_j - 1$ 
9:     end if
10:  end for
11: end procedure

```

---

*C. Parallel Correctness, Termination, and Cost Analysis*

To conclude our discussion of the  $\mathbb{P} + \mathbb{S}$  algorithm, we analyze the correctness, finite termination, and cost properties of our propositions. First, we verify that  $\mathcal{E}_{\mathbb{P}}^* \triangleq \bigcup_i \mathcal{E}_{\mathbb{P}, i}^*$  has a valid distributed construction:

*Proposition 4.2 (Parallel mutual exclusion):* Consider the application of the parallel  $\mathbb{P}$  algorithm to a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ . It follows that upon termination we have mutual exclusion  $\mathcal{E}_i^* \cap \mathcal{E}_j^* = \emptyset, \forall i \neq j \in \mathcal{I}$ .

*Proof:* To show that mutual exclusion holds we consider the varying scenarios of message reception and local edge consideration for a pair of agents  $i, j \in \mathcal{I}$  sharing an edge  $(i, j) \in \mathcal{E}$ . In each scenario we will then reason that mutual exclusion is satisfied. Given the assumptions of asynchronicity in messaging and the FIFO queuing of received messages (Assumption 3) and execution devoid of race conditions (Assumption 2), the following scenarios must be considered, viewed from the instant when agent  $i$  handles an EDGERESPONSEMSG from  $j$ , implying that  $e_i = (i, j)$  and  $(i, j) \notin \mathcal{E}_i^*$  (line 2, Algorithm 8):

- $(i, j) \notin \mathcal{E}_j^*$ : We must consider two cases here, either  $e_j = (i, j)$  or  $e_j \neq (i, j)$ . First, in the trivial case of  $e_j \neq (i, j)$ , it follows from reception of an EDGERESPONSEMSG from  $j$ , the atomic nature of execution, and line 13 of Algorithm 7, that regardless of assignment to  $\mathcal{E}_i^*$ ,  $(i, j) \notin \mathcal{E}_j^*$  for all execution  $t > 0$ . When  $e_j = (i, j)$ , the conflict resolution of line 3-7 in Algorithm 7, together with line 2 of Algorithm 8 ensures that simultaneous requests made over  $(i, j)$  agree on assignment, specifically as by assumption  $RESOLVECONFLICT(i, j) = RESOLVECONFLICT(j, i)$ .
- $(i, j) \in \mathcal{E}_j^*$ : Here it is implied that at some previous time

agent  $i$  received and responded to a EDGEREQUESTMSG from  $j$ . As agent  $i$  being in a state of response reception over  $(i, j)$  is contradictory given line 13 of Algorithm 7, it must be the case that requests over  $(i, j)$  have been made in concert. However, as previously stated, the conflict resolution ensures  $\mathcal{E}_i^* \cap \mathcal{E}_j^* = \emptyset$  in such scenarios.

Notice that due to the uniformity of execution and messaging logic across  $i \in \mathcal{I}$ , the previous scenarios hold equivalently from the perspective of agent  $j$ , and thus for all pairs  $\{i \neq j \mid (i, j) \in \mathcal{E}\}$ , and the result follows. ■

Of course,  $\mathcal{E}_{\mathbb{P}}^*$  must also fulfill the independence requirements of Definition 2.2, as is shown below.

*Proposition 4.3 (Parallel Correctness):* Consider the algorithm  $\mathbb{P}$  applied to a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ . For all execution  $t > 0$  it follows that edge addition operations,  $\mathcal{E}_i^*(t^+) = \mathcal{E}_i^*(t) \cup (i, j)$  (line 3 Algorithm 8), are independence preserving and  $\bigcup_i \mathcal{E}_i^*$  is independent with  $|\bigcup_i \mathcal{E}_i^*| \leq 2n - 3$ .

*Proof:* Resting again on Assumptions 2 and 3, and the uniform conflict resolution of  $\text{RESOLVECONFLICT}(i, j)$ , this result is a consequence of the commitment counting rules (lines 10 and 5 of Algorithms 7 and 8), and the condition on line 2 of Algorithm 8 that enforces the cardinality of endpoint  $j$  in  $\bigcup_i \mathcal{E}_i^*(t^+)$ . In particular, when  $j \notin \bigcup_i \mathcal{E}_i^*(t)$  ( $\text{committed}(j) = 0$ ),  $\mathcal{E}_i^*(t^+) = \mathcal{E}_i^*(t) \cup (i, j)$  constitutes an EER operation (c.f. Definition 4.3), otherwise when  $j \in \bigcup_i \mathcal{E}_i^*(t)$  ( $\text{committed}(j) = 1$ ),  $\mathcal{E}_i^*(t^+) = \mathcal{E}_i^*(t) \cup (i, j)$  constitutes a TIER operation (c.f. Definition 4.4). As sequences of EER and TIER operations preserve independence by Proposition 4.1,  $\bigcup_i \mathcal{E}_i^*$  is independent, with  $|\bigcup_i \mathcal{E}_i^*| \leq 2n - 3$  following directly from the Laman conditions Theorem 2.1. ■

To ensure timely results, we must also have finite termination of  $\mathbb{P} + \mathbb{S}$ :

*Proposition 4.4 ( $\mathbb{P}$  termination):* Consider the execution of the  $\mathbb{P}$  algorithm as described in Section IV-B. By construction, it follows that the stopping condition  $f_{stop}^{\mathbb{P}}$  of (11) is satisfied in finite time.

*Proof:* We can again guarantee no message-induced race conditions by Assumptions 2 and 3. Thus, the finiteness of execution is a direct consequence of the finiteness of each  $\mathcal{E}_i$ ,  $\forall i \in \mathcal{I}$ , the internal blocking on line 2 of Algorithm 6, the guaranteed delivery of inter-agent messages by the best-effort Assumption 3, and finally the symmetric conflict resolution of Remark 4.1, disallowing conflict based race conditions. ■

To constitute a valid initial condition for  $\mathbb{S}$ , the pebble assignments applied to the terminal state  $\mathcal{E}_{\mathbb{P}}^*$  of  $\mathbb{P}$  must be sound:

*Proposition 4.5 (Pebble Assignments):* Consider an execution of  $\mathbb{P}$ , applied to a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ . There must exist a local pebble covering for every  $(i, j) \in \mathcal{E}_i^* \forall i \in \mathcal{I}$ , that is an assignment of a pebble by either  $i$  or  $j$  to  $(i, j)$ . In particular, this assignment is achieved as depicted in Algorithm 9.

*Proof:* First notice that Algorithm 9 simply dictates that each agent  $i$  negotiates the assignment of a free pebble to  $(i, j) \in \mathcal{E}_i^*$  using local communication. In guaranteeing that such an assignment exists, we rely on the properties of the EER and TIER operations. As each operation  $\mathcal{E}_i^*(t^+) = \mathcal{E}_i^*(t) \cup (i, j)$  respects Proposition 4.1 by Proposition 4.3, we have that for any  $(i, j) \in \bigcup_i \mathcal{E}_i^*$  there must exist an endpoint  $i$  or  $j$  with at most two incident edges. From this endpoint we can thus always




Fig. 6. Example parallel messaging for a minimally rigid graph with  $n = 4$ . Inter-agent requests are denoted by solid arrows, responses by dashed arrows, and pebble assignments by solid dots. Notice that conflicts occur in (a) over  $(v_1, v_2)$  and in (c) over  $(v_2, v_4)$ , with resolution dictated by agent label  $i < j$  for  $j \in \mathcal{N}_i$ .

select a pebble to cover  $(i, j)$  as each agent  $i \in \mathcal{I}$  is initially assigned two pebbles. ■

Now we come to our primary result concerning the correctness and finite termination of the  $\mathbb{P} + \mathbb{S}$  algorithm:

*Proposition 4.6 ( $\mathbb{P} + \mathbb{S}$  correctness and termination):* Consider an execution of  $\mathbb{P} + \mathbb{S}$ , applied to a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ . It follows that the terminal system state:

$$\mathbb{P}(\bar{t}_{\mathbb{P}}) \triangleq \left\{ \bigcup_i \mathcal{E}_i^*(\bar{t}_{\mathbb{P}}), \mathcal{P}_i, p_i \right\} \quad \forall i \in \mathcal{I} \quad (13)$$

is a valid initial condition for the  $\mathbb{S}$  algorithm, the execution of  $\mathbb{P} + \mathbb{S}$  terminates in finite time, and properly identifies the generic topology rigidity of  $\mathcal{G}$ .

*Proof:* First, we have directly from Propositions 3.3 and 4.4, the known finite convergence of gossip averaging [41], and the trivial finiteness of Algorithm 9, that the composite execution of  $\mathbb{P} + \mathbb{S}$  terminates in finite time. Now, from Propositions 4.2 and 4.3 it follows that the state  $\mathbb{P}(\bar{t}_{\mathbb{P}})$  represents a properly distributed and independent edge set, and from Proposition 4.5 that there must exist pebble assignments  $\mathcal{P}_i$  that are local shift operations relative to Definition 2.3, and thus constitute a valid pebble covering. Thus the application of  $\mathbb{S}$  with input  $\mathbb{P}(\bar{t}_{\mathbb{P}})$  is correct by Proposition 3.4, and the result follows. ■

The above result demonstrates that  $\mathbb{P} + \mathbb{S}$  is indeed a sound rigidity evaluation technique for planar networks. Further, towards the real-world applicability of our propositions, we have complexity that scales well with network size:

*Proposition 4.7 ( $\mathbb{P} + \mathbb{S}$  complexity):* By construction, executions of the  $\mathbb{P} + \mathbb{S}$  algorithm have worst-case  $O(n^2)$  messaging



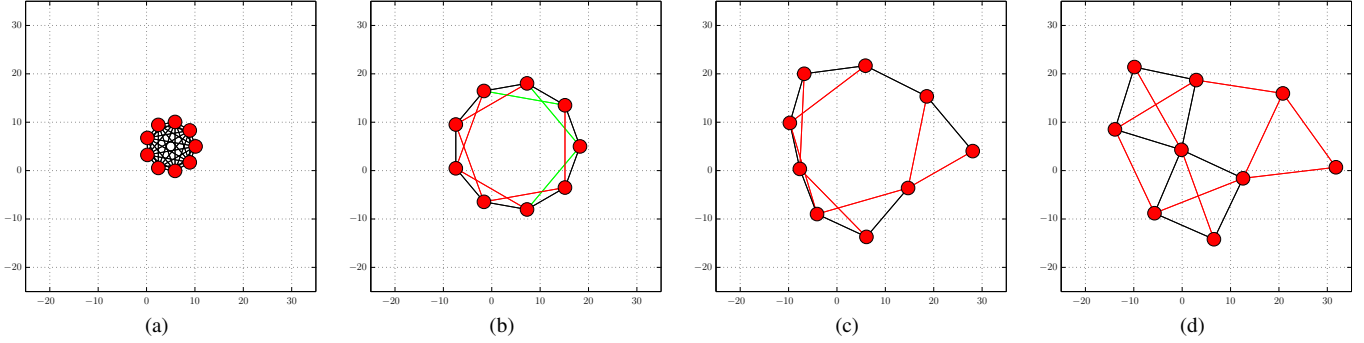


Fig. 7. Rigidity control simulation for  $n = 9$  mobile agents applying a dispersive objective with a rigidity maintenance constraint: (a) initial rigid configuration; (b) and (c) intermediate configurations; (d) final converged and rigid configuration. Green edges are not independent and thus can be lost, while red edges are actively retained due to their independence.

complexity (and thus execution complexity), and  $O(n)$  storage scaling.

*Proof:* For the  $\mathbb{P}$  portion of the  $\mathbb{P} + \mathbb{S}$  algorithm, we have that each of  $n$  agents communicates with at most  $n - 1$  neighbors, yielding  $O(n^2)$  messaging. The gossip averaging for determining  $|\mathcal{E}_{\mathbb{P}}^*|$  exhibits  $O(n \log n)$  messaging [41], while the pebble assignment Algorithm 9 requires  $O(n^2)$  by inspection. Thus we have an overall worst-case message complexity of  $O(n^2)$  for a  $\mathbb{P} + \mathbb{S}$  as  $\mathbb{S}$  also has  $O(n^2)$  messaging by Proposition 3.5. The overall storage complexity follows directly from the fact that assignments to  $\mathcal{E}^*$  can be made only locally by each agent  $i \in \mathcal{I}$ , and thus scales like  $O(n)$ . ■

Finally, to close we can also roughly evaluate the expected improvement provided by the partial parallelization due to  $\mathbb{P}$ :

*Proposition 4.8 (Parallel identification):* An execution of  $\mathbb{P}$ , applied to a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  (where we assume every  $v_i \in \mathcal{V}$  is an endpoint in  $\mathcal{E}$ ), with  $n \geq 3$  must result in the terminal condition

$$|\mathcal{E}_{\mathbb{P}}^*(\bar{t}_{\mathbb{P}})| \triangleq \left| \bigcup_i \mathcal{E}_{\mathbb{P},i}^*(\bar{t}_{\mathbb{P}}) \right| \geq \left\lceil \frac{n-1}{2} \right\rceil + 1 \quad (14)$$

where  $\lceil \cdot \rceil$  is the standard *ceiling operator*, yielding a lower bound on the independent edges identified by the  $\mathbb{P}$  algorithm.

*Proof:* First, notice that we disregard the case of  $n = 1$  (as  $|\mathcal{E}| = 0$ ), and for  $n = 2$  we can always identify the single member of  $\mathcal{E}^*$  due to symmetric conflict resolution. Now, observing that the single  $n = 2$  graph is a worst-case in terms of detectable independent edges, with  $|\mathcal{E}^*| = \lceil n/2 \rceil$ , for  $n \geq 3$  we can construct a similar worst-case graph by appending a single node and edge to the  $n - 1$  worst-case, as the appended edge will always be detectable by an EER operation. As we add a single detectable edge to the previous worst-case, a simple inductive argument yields the result. ■

The above result states directly that we can always detect a spanning tree over  $\mathcal{G}$  using our proposed parallel and distributed interactions. Extrapolations of Proposition 4.8 to the complexity of  $\mathcal{P} + \mathcal{S}$  are not trivial given the nature of the pebble search mechanism of  $\mathcal{S}$  and are beyond the scope of this work. However, as will be demonstrated in Section V-B, the proposed parallelization does indeed admit significant improvements in messaging and execution under realistic conditions.




Fig. 8. Swarm size and edge set cardinality for the rigidity control simulation.

## V. SIMULATION RESULTS

### A. A Rigidity Control Scenario

Although identifying rigid network elements has strong implications in static networks (e.g. self-discovery of minimally localizable sensor network topologies), we wish to also demonstrate a scenario where network rigidity can be controlled in a dynamic multi-robot system. To begin, consider that in controlling generic rigidity we need only to disallow the loss of independent edges  $(i, j) \in \mathcal{E}^*$ . For this purpose, we can employ the *constrained interaction framework* proposed by Williams and Sukhatme in [4], a very brief overview of which will be given here<sup>13</sup>. Assuming proximity-limited sensing and communication (e.g. disk-based sensing and fading wireless channels), together with agent dynamics  $\dot{x}_i = u_i$ , the constrained interaction framework regulates link addition and deletion *spatially*, through a switching combination of hysteresis, attraction, and repulsion, in order to satisfy a desired set of constraints. In particular, each agent is assigned *predicates*  $P_{ij}^a, P_{ij}^d : \mathcal{V} \times \mathcal{V} \leftrightarrow \{0, 1\}$  that indicate constraint violations if link  $(i, j)$  were gained or lost. Thus, in applying our propose decentralized pebble game to identify local sets  $\mathcal{E}_i^* \forall i \in \mathcal{I}$ , we arrive directly at rigidity preserving predicates. That is,

$$P_{ij}^a \triangleq 0, \quad P_{ij}^d \triangleq (i, j) \in (\mathcal{E}_i^* \cup \mathcal{E}_j^*) \quad (15)$$

<sup>13</sup>For brevity, we omit certain details related to framework assumptions, constructions, and stability. We urge the reader to consult [4] for analysis.



where we point out that link addition has no impact on network rigidity<sup>14</sup>, and that  $P_{ij}^d$  can easily be determined through local communication between agents  $i$  and  $j$ . Notice that the proposed algorithms of Sections III and IV, in conjunction with the predicates (15) of the constrained interaction framework, results in a decentralized and efficient solution to dynamic network rigidity control. In comparison to the state of the art, e.g. [31], which provides to date only centralized treatments, relying on constant communication and spectral analysis techniques that bear significantly higher cost.

Now, we are prepared to present our rigidity control simulation results. We assume a system of  $n = 9$  mobile agents, each with proximity-limited communication and sensing, applying for the sake of link deletion, a dispersive objective controller, yielding agent controllers with generic form:

$$u_i = u_{CI} - \nabla_{x_i} \left( \sum_{i \in \mathcal{N}_i} \frac{1}{\|x_i - x_j\|^2} \right) \quad \forall i \in \mathcal{I} \quad (16)$$

where  $u_{CI}$  is the control contribution due to the constrained interaction framework and predicates (15). The agents begin in the fully connected initial configuration given by the *ring network* depicted in Fig. 7a, satisfying the initial condition  $\mathcal{G}(0) \in \mathcal{G}_{\mathbb{R}}$ . Through controllers (16) the agents reach intermediate configurations given by Figs. 7b and 7c, ultimately terminating in the final configuration in Fig. 7d. For  $n = 9$ , a rigid network configuration has  $|\mathcal{E}(t)| \geq |\mathcal{E}^*(t)| \triangleq 2n - 3 = 15, \forall t > 0$ , and thus we can conclude from the swarm size and edge counts depicted in Fig. 8 that network rigidity control under the desired dispersion objective is in fact achieved throughout the simulation.

*Remark 5.1 (Simulation videos):* We direct the reader to **Video 1** and **Video 2** for extended simulation videos illustrating rigidity control and a rigidity recovery scenario, each exploiting the techniques discussed in this work. Explanations can be found **Here**.

### B. Contiki Implementation: Real-World Feasibility Results

To determine the performance of our algorithms under realistic networking conditions, we consider the *Contiki* operating system, together with the *Cooja* network simulator<sup>15</sup>. Contiki and Cooja provide a vast selection of hardware compatibility, and an array of other networking and power management features that make it an ideal choice for evaluating a networked algorithm under realistic conditions. We implemented both the serial and parallel decentralized pebble game for Contiki and tested our codebase against a range of emulated hardware platforms and communication stacks for correctness. In particular, the primary settings of our simulations are as follows:

- *MAC driver:* Carrier Sense Multiple Access (CSMA)
- *RDC driver:* NullRDC
- *Networking protocol:* RIME with best effort delivery, [43]
- *Emulated hardware:* Tmote Sky platform

We simulated in a Monte Carlo fashion by generating rigid and non-rigid networks for  $n \in \{5, 29\}$ , generating the results

<sup>14</sup>Note however that as link addition does impact redundancy, the constrained interaction framework could be exploited to provide redundancy control as well as rigidity guarantees.

<sup>15</sup>See <http://www.contiki-os.org>.




Fig. 9. Monte Carlo results from a Contiki networking environment demonstrating execution time and message complexity for the  $\mathbb{S}$  and  $\mathbb{P} + \mathbb{S}$  algorithms. Solid lines are average values, while dashed lines are maximum and minimums.

depicted in Fig. 9. Specifically, Fig. 9 (top) compares the execution time (seconds) for the serial and parallel algorithms, while Fig. 9 (bottom) shows the per-agent messaging burden. It is clear that both of our algorithms exhibit feasible and efficient performance, with actual scaling that is approximately  $O(n)$  in both execution and messaging. The parallel version however represents our goal of real-world capability (outperforming the serial version by a factor 2), as even in reasonably sized networks, execution times for evaluation are under 1 second. When coupled with the necessity to check network rigidity only on transitions in  $\mathcal{G}(t)$ , we suggest that the parallel pebble game algorithm provides a fully decentralized, time and communication efficient bridge between rigidity theory and practicing roboticists.

Finally, to commit to providing a realistic rigidity evaluation method, we have released an initial version of the base code for our proposed algorithm for the community to investigate, apply, and extend in the ongoing effort to move multi-robot systems out of the laboratory and into our world<sup>16</sup>.

## VI. CONCLUSIONS AND FUTURE WORK

In this paper, we considered the problem of evaluating the rigidity of a planar network under the requirements of decentralization, asynchronicity, and parallelization. We proposed the decentralization of the pebble game algorithm of Jacobs et. al., based on asynchronous inter-agent message-passing and distributed auctions for electing network leaders. Further, we provided a parallelization of our methods that yielded significantly reduced execution time and messaging burden. Finally, we provided a simulated application in decentralized rigidity control, and Monte Carlo analysis of our algorithms in a *Contiki*

<sup>16</sup>See [http://github.com/Attilio-Priolo/Rigidity\\_Check\\_Contiki](http://github.com/Attilio-Priolo/Rigidity_Check_Contiki)

networking environment, illustrating the real-world applicability of our methods.

Directions for future work include demonstrating our methods on robotic hardware in field experiments (our current focus), extensions to global rigidity for even stronger guarantees in multi-robot applications, and the inclusion of network utility metrics to yield decentralized rigidity evaluation and control with provable optimality conditions.

## REFERENCES

- [1] N. E. Leonard, D. A. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, and R. E. Davis, "Collective Motion, Sensor Networks, and Ocean Sampling," *Proceedings of the IEEE*, 2007.
- [2] M. A. Hsieh, E. Forgoston, T. W. Mather, and I. B. Schwartz, "Robotic manifold tracking of coherent structures in flows," in *IEEE International Conference on Robotics and Automation*, 2012.
- [3] J. Cortes, S. Martínez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, 2004.
- [4] R. K. Williams and G. S. Sukhatme, "Constrained Interaction and Coordination in Proximity-Limited Multi-Agent Systems," *IEEE Transactions on Robotics*, 2013.
- [5] —, "Topology-Constrained Flocking in Locally Interacting Mobile Networks," in *IEEE International Conference on Robotics and Automation*, 2013.
- [6] S. Martínez, J. Cortes, and F. Bullo, "Motion Coordination with Distributed Information," *Control Systems Magazine, IEEE*, 2007.
- [7] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, 2004.
- [8] S. Kar and J. Moura, "Distributed Consensus Algorithms in Sensor Networks With Imperfect Communication: Link Failures and Channel Noise," *IEEE Transactions on Signal Processing*, 2009.
- [9] W. Ren and R. W. Beard, "Consensus seeking in multiagent systems under dynamically changing interaction topologies," *IEEE Transactions on Automatic Control*, 2005.
- [10] R. Olfati-Saber and R. M. Murray, "Graph rigidity and distributed formation stabilization of multi-vehicle systems," in *IEEE Conference on Decision and Control*, 2002.
- [11] T. Eren, P. N. Belhumeur, and A. Morse, "Closing ranks in vehicle formations based on rigidity," in *IEEE Conference on Decision and Control*, 2002.
- [12] B. Anderson, C. Yu, B. Fidan, and J. Hendrickx, "Rigid graph control architectures for autonomous formations," *Control Systems, IEEE*, 2008.
- [13] A. N. Bishop, I. Shames, and B. D. Anderson, "Stabilization of rigid formations with direction-only constraints," in *IEEE Conference on Decision and Control and European Control Conference*, 2011.
- [14] S. Motevallian, C. Yu, and B. D. O. Anderson, "Multi-agent rigid formations: A study of robustness to the loss of multiple agents," in *IEEE Conference on Decision and Control and European Control Conference*, 2011.
- [15] T. Eren, "Formation shape control based on bearing rigidity," *International Journal of Control*, 2012.
- [16] T. Eren, W. Whiteley, A. Morse, P. N. Belhumeur, and B. D. O. Anderson, "Sensor and network topologies of formations with direction, bearing, and angle information between agents," in *IEEE Conference on Decision and Control*, 2003.
- [17] T. Eren, W. Whiteley, and P. N. Belhumeur, "Using Angle of Arrival (Bearing) Information in Network Localization," in *Decision and Control, 2006 45th IEEE Conference on*, 2006.
- [18] J. Aspnes, T. Eren, D. K. Goldenberg, A. Morse, W. Whiteley, Y. R. Yang, B. D. O. Anderson, and P. N. Belhumeur, "A Theory of Network Localization," *IEEE Transactions on Mobile Computing*, 2006.
- [19] I. Shames, A. N. Bishop, and B. D. O. Anderson, "Analysis of Noisy Bearing-Only Network Localization," *IEEE Transactions on Automatic Control*, 2013.
- [20] A. R. Berg and T. Jordan, "A proof of Connelly's conjecture on 3-connected circuits of the rigidity matroid," *J. Comb. Theory Ser. B*, 2003.
- [21] B. Jackson and T. Jordan, "Connected rigidity matroids and unique realizations of graphs," *J. Comb. Theory Ser. B*, 2005.
- [22] B. Hendrickson, "Conditions for unique graph realizations," *SIAM J. Comput.*, 1992.
- [23] G. Laman, "On graphs and rigidity of plane skeletal structures," *Journal of Engineering Mathematics*, 1970.
- [24] L. Asimow and B. Roth, "The rigidity of graphs, II," *Journal of Mathematical Analysis and Applications*, 1979.
- [25] B. Roth, "Rigid and Flexible Frameworks," *The American Mathematical Monthly*, 1981.
- [26] L. Lovasz and Y. Yemini, "On Generic Rigidity in the Plane," *Siam Journal on Algebraic and Discrete Methods*, 1982.
- [27] T.-S. Tay and W. Whiteley, "Generating Isostatic Frameworks," *Structural Topology*, 1985.
- [28] D. J. Jacobs and B. Hendrickson, "An algorithm for two-dimensional rigidity percolation: the pebble game," *J. Comput. Phys.*, 1997.
- [29] D. Zelazo and F. Allgower, "Growing optimally rigid formations," in *American Control Conference*, 2012.
- [30] R. Ren, Y.-Y. Zhang, X.-Y. Luo, and S.-B. Li, "Automatic generation of optimally rigid formations using decentralized methods," *Int. J. Autom. Comput.*, 2010.
- [31] D. Zelazo, A. Franchi, F. Allgower, H. H. Bulthoff, and P. R. Giordano, "Rigidity Maintenance Control for Multi-Robot Systems," in *Robotics: Science and Systems*, 2012.
- [32] B. Fidan, J. M. Hendricks, and B. D. O. Anderson, "Edge contraction based maintenance of rigidity in multi-agent formations during agent loss," in *Control and Automation, 2009. MED '09. 17th Mediterranean Conference on*, 2009.
- [33] R. K. Williams, A. Gasparri, A. Priolo, and G. S. Sukhatme, "Decentralized Generic Rigidity Evaluation in Interconnected Systems," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (to appear)*, 2013.
- [34] A. Nedic, "Asynchronous Broadcast-Based Convex Optimization over a Network," *IEEE Transactions on Automatic Control*, 2010.
- [35] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Gossip algorithms: design, analysis and applications," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, 2005.
- [36] R. Olfati-Saber, J. Fax, and R. M. Murray, "Consensus and Cooperation in Networked Multi-Agent Systems," *Proceedings of the IEEE*, 2007.
- [37] H. Gluck, "Almost all simply connected closed surfaces are rigid," in *Geometric Topology*, 1975.
- [38] D. P. Bertsekas, "The auction algorithm: A distributed relaxation method for the assignment problem," *Annals of Operations Research*, 1988.
- [39] M. M. Zavlanos, L. Spesivtsev, and G. J. Pappas, "A distributed auction algorithm for the assignment problem," in *IEEE Conference on Decision and Control*, 2008.
- [40] D. Zelazo and M. Mesbahi, "Graph-Theoretic Analysis and Synthesis of Relative Sensing Networks," *IEEE Transactions on Automatic Control*, 2011.
- [41] A. G. Dimakis, S. Kar, J. M. F. Moura, M. G. Rabbat, and A. Scaglione, "Gossip Algorithms for Distributed Signal Processing," *Proceedings of the IEEE*, 2010.
- [42] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *Information Theory, IEEE Transactions on*, 2006.
- [43] A. Dunkels, F. Osterlind, and Z. He, "An adaptive communication architecture for wireless sensor networks," in *Proceedings of the 5th international conference on Embedded networked sensor systems*, 2007.