

A Grammatical Approach to Self-Organizing Robotic Systems

Eric Klavins* Robert Ghrist† David Lipsky†

*Electrical Engineering
University of Washington
Seattle, WA 98195

†Mathematics
University of Illinois
Urbana, IL 61801

Abstract—In this paper we define a class of graph grammars that can be used to model and direct concurrent robotic self-assembly and similar self-organizing processes. We give several detailed examples of the formalism and then focus on the problem of synthesizing a grammar so that it generates a given, prespecified assembly. In particular, to generate an acyclic graph we synthesize a binary grammar (rules involve at most two parts), and for a general graph we synthesize a ternary grammar (rules involve at most three parts). In both cases, we characterize the number of concurrent steps required to achieve the assembly. We also show a general result that implies that no binary grammar can generate a unique stable assembly. We conclude the paper with a discussion of how graph grammars can be used to direct the self-assembly of robotic parts.

I. INTRODUCTION

Engineering processes in the realm of the very small presents us with the daunting problem of manipulating and coordinating vast numbers of objects so that they perform some global task. Because of the potentially enormous quantities of objects involved, uniquely addressing and manipulating each one is impossible. There are of course examples of sophisticated machines, such as the ribosome or the mechanical motor in the bacterial flagellum, that seem to be built *in bulk* spontaneously out of large numbers of simple interacting components. One hypothesis for how this occurs is that the simple small components *self-assemble* into more complex aggregates which, in turn, self assemble into larger aggregates or function as simple machines. The goal of the present work is to understand how the interactions between the components of a system can give rise to coordinated behavior in a self-organizing (i.e. decentralized and unsupervised) way. In particular, we suppose that the components themselves are (very) simple robots, whose capabilities we will specify in the body of this work.

Our starting point is the idea of *conformational switching* [1]: Each component (robot, molecule, particle, etc.) exists in one of several *conformations* or shapes. When two components come into close proximity, they attach or not based on whether their conformations are complimentary. If they do attach, their conformations change (mechanically for example), thereby determining in what future interactions the components may partake.

As in other work [2], [3], we consider the conformation of a part or *robot* as corresponding to a discrete symbol, and we

model an assembly as a simple graph labeled by such symbols. Vertices in these graphs represent robots, and the presence of an edge between two robots represents the fact they are attached. A *rule* is a pair of labeled graphs and a grammar is a set of such rules interpreted as follows. If a subset of robots together with their labels and edges matches the first part some rule, then the subset may be replaced by the second part of the rule to achieve a new state of the system. That is, we use the rule to *rewrite* a part of the current graph. Continuing to apply rules one may produce a class of trajectories whose properties one might want to guarantee.

We are mainly interested in the situation where the parts decide in a distributed fashion whether and how to execute an assembly rule. In this sense, a graph grammar defines a *distributed algorithm*, but not one that works on a fixed-topology network. Instead, the fact that the processors are located on robotic parts means that the network topology changes as the robots move through their environment.

One motivation for the theory presented in this paper is a robotic testbed under construction at the University of Washington wherein “programmable parts” execute local rules to form global assemblies [4], [5]. Figure 1 shows a photograph of several programmable parts. On each edge of a part is a magnetic latching mechanism and an infrared (IR) transceiver. Each part is controlled by an on-board micro-controller programmed according to rules of the sort presented in this paper. Many such parts are floated on an air table as in Figure 2(a) and experience essentially random motions. When two parts collide, their default behavior is to attach to each other. Only then may they communicate over their IR link. By sharing their internal states, the parts may decide whether to remain attached or to detach. Using the theory presented in this paper, this basic interaction protocol can be used to form assemblies of parts, such as those shown in Figure 2(b) (see also Example IV-C). The point of the testbed is to demonstrate the feasibility of the approach described in this paper. With relatively simple devices and a formal algorithmic and synthesis approach, we are able to control this system to self-assemble into a variety of forms.

The main problem we consider in this paper is the *synthesis* problem, defined roughly as follows. Given a specification S (i.e., a logical statement about trajectories) and an initial configuration of robots G_0 , define a grammar Φ so that all trajectories arising from G_0 via the application of rules in Φ



Fig. 1. Robotic “programmable parts” [4]. Each edge of a part may attach via a magnetic latch to another part. Once attached, the parts may communicate and decide to remain attached or may dis-attach by rotating a permanent magnet, which disengages the latch.




Fig. 2. (a) An initial configuration of programmable parts on an air table. The parts are stirred by air jets (not shown) so as to produce collisions. (b) An assembly of programmable parts after the assembly process has completed. A complete final assembly and several partial assemblies are shown.

meet the specification S . The specification may require that some cyclic process occurs among the robots (as for example in the locomotion of *Metamorphic Robots* [6], [7], [8]) or that some undesirable set of configurations be avoided. In particular, we address the *Self-Assembly Synthesis Problem*: Given an arbitrary graph H define Φ so that only copies of H emerge as stable components of the system. For example, one might desire that the robots all form rings of 10 robots, or arrange themselves into a long line.

Specifically, the contributions of this paper are as follows. In Section III we introduce a class of graph grammars suitable for describing distributed assembly. We give a set of detailed examples that elucidate the definitions and demonstrate the capabilities of the formalism in Section IV. We next examine some basic properties and limitations of graph grammars. In particular, in Section V we prove a theorem about the impossibility of assembling a unique stable graph using rules consisting of acyclic graphs (to be defined). We also explain how to view graph grammars as describing a concurrent process in Section VI and associate to each trajectory of a system a canonical trajectory that captures the number of *concurrent* steps required to assemble a given component. This number is used to evaluate the rule sets described later in the paper. Then in Section VII we describe two algorithms that take a graph G and produce a grammar Φ such that the only *stable* assemblies of Φ are isomorphic to G . The first algorithm requires that G is acyclic, and the second algorithm (which uses the first) works for any given graph. We show the

algorithms are correct and also explain how quickly the desired graphs can be assembled. In Section VIII we discuss how assembly rules can be implemented in a distributed fashion by simple communication protocols and finally, in Section IX, we describe physical models that can be used to implement a graph grammar.

II. PREVIOUS AND RELATED WORK

Conformational switching was described as a symbolic process for self-assembly by Saitou [2], who considered the assembly of strings in one dimension. Berger et al. showed that a conformational switching model could mimic the assembly of virus capsids in simulation [9]. The addition of simple processing to each part, similar in capability to that assumed in the present paper, is considered in models of the assembly of the T4 bacteriophage [10]. Self assembly as a graph process has been described by Klavins [3], although the graph grammar formalism is new to this paper, and a rule synthesis procedure for trees was given that is somewhat more complex than the one described in Section VII. A method for using potential fields and a certain deadlock avoidance scheme to implement local rules with a group of mobile robots was also described [3]. The proof of Theorem 5.1 is topological, utilizing tools from covering space theory (see e.g., [11]). The notion of concurrency we use in this paper is similar to that of processes as described by Reisig [12] and our notion of concurrent equivalence is directly related to Abrams’ and Ghrist’s work on State Complexes [13].

Graph grammars were introduced [14], [15] more than two decades ago and have been used to describe a broad array of systems, from data structure maintenance to mechanical system synthesis. Graph grammars come in many flavors. The variety we use in this paper is called *Graph Rewriting Systems on Induced Subgraphs* (IGRS), investigated by Litovsky et al. [16] (The requirement that rules are applied on induced subgraphs is equivalent to our requirement in Definition 3.2 that witnesses be monomorphisms). Our results about the reachable and stable sets being closed under covers are related in spirit to the results of Courcelle that the classes of graphs that can be *recognized* by local computations (i.e. by an appropriate graph grammar) must be closed under covers [17].

The study of graph grammars usually takes place in a much different context than in the present paper, as can be seen by examining the somewhat different names usually given to the objects we define in this paper: The initial graph is often called the *axiom*, trajectories are called *derivations*, rewrite rules are called *productions* and so on. The terminology in the present work is meant to correspond more closely to the dynamical systems literature. Graph grammars are a generalization of the standard “linear” grammars used in automata theory and linguistics and (incidentally) can perform arbitrary computation. The use of graph grammars to model distributed assembly, to the best of our knowledge, is new.

There are many other models of self assembly besides graph grammars. For example, several groups [18], [19] have explored self assembly using passive *tiles* floating in liquid. The tiles attach along complimentary edges (due, for example, to capillary forces or the hybridization of complimentary strands of DNA) upon random collisions. Somewhat similar to

the *stable set* in this paper (Definition 3.8), the identification of “unique” assemblies has been explored [20]. There is also other work on supplying assembling robots with state information [21], [22], [23], although not in a graph-theoretic context, and also initial work by other groups [24], [25] on building a self-assembling robot similar to that described in the introduction to this paper.

III. DEFINITIONS

A. Rules and Systems

In this section we provide the basic definitions related to our notion of a graph grammar. We save examples for Section IV. The reader may wish to read the examples in Section IV-A concurrently with this section to help ground the definitions.

A *simple labeled graph* over an alphabet Σ is a triple $G = (V, E, l)$ where V is a set of *vertices*, E is a set of pairs or *edges* from V , and $l : V \rightarrow \Sigma$ is a labeling function. We restrict our discussion to simple labeled graphs and thus simply use the term *graph*. We denote an edge $\{x, y\} \in E$ by xy . We denote by V_G , E_G and l_G the vertex set, edge set and labeling function of the graph G by V , E and l when there is no danger of confusion. We usually use the alphabet $\Sigma = \{a, b, c, \dots\}$.

In this paper, a graph is a model of the *network topology* of an interconnected collection of robots, vehicles or particles. In most examples, a vertex x corresponds to a robot, or, more exactly, the *index* of a robot in an indexed set. The presence of an edge xy corresponds to an attachment (via a physical and/or communication link) between robots x and y . The label $l(x)$ of robot x corresponds to the *state* of the robot and will be used to keep track of local information.

Definition 3.1: A rule is a pair of graphs $r = (L, R)$ where $V_L = V_R$. The graphs L and R are called the *left hand side* and *right hand side* of r respectively. The *size* of r is $|V_L| = |V_R|$. Rules whose vertex sets have one, two and three vertices are called *unary*, *binary* and *ternary*, respectively.

If (L, R) is a rule, we sometimes denote it by $L \rightarrow R$ to emphasize its use as a *rewrite rule* or *grammatical production*. We also represent rules graphically as in

$$\begin{array}{c} b \\ / \quad \backslash \\ a \quad c \end{array} \rightarrow \begin{array}{c} e \\ / \quad \backslash \\ d \text{---} f \end{array},$$

where the relative locations of the vertices represent their identities. In the above rule, for example, $V = \{1, 2, 3\}$ and vertex 1 is labeled by $l_L(1) = a$ in the left hand side and by $l_R(1) = d$ in the right hand side.

We may refer to rules as being *constructive* ($E_L \subset E_R$), *destructive* ($E_L \supset E_R$) or *mixed* (neither constructive or destructive). A rule is called *acyclic* if its *right hand side* contain no cycles (n.b., the left hand side of an acyclic rule may contain cycles).

A rule describes how a (small) sub-collection of robots can change its local network topology and states. In particular, if an induced subgraph of a graph G matches the left hand side L of a rule (L, R) , that subgraph may be replaced by the right hand side R (for simplicity, we suppose this transition occurs *instantaneously* in this paper). The size (number of vertices) of a rule is a measure of “how local” it is: A unary

rule can be executed by a robot without any communication. A binary rule requires the cooperation of two robots via one communication event. Larger rules require even more communication, since a larger group of robots must somehow realize that they collectively match the left hand side of a rule. Typically, one prefers small rules. Also notice that the rules depend *only* on the labels (or states) of the robots and not on their underlying indices in the graph. That is, the indices we use for underlying vertex set V do not give particular robots any special properties. The following definitions make formal the notion of rule application.

A *homomorphism* between graphs G_1 and G_2 is a function $h : V_{G_1} \rightarrow V_{G_2}$ which preserves edges: $xy \in E_{G_1} \Leftrightarrow h(x)h(y) \in E_{G_2}$. A *monomorphism* is an injective homomorphism. An *isomorphism* is a surjective monomorphism: in this case, G_1 and G_2 are said to be *isomorphic*. A homomorphism h is said to be *label preserving* if $l_{G_1} = l_{G_2} \circ h$.

Definition 3.2: A rule $r = (L, R)$ is *applicable* to a graph G if there exists a label-preserving monomorphism $h : V_L \rightarrow V_G$. In this case, the function h is called a *witness*. An *action* on a graph G is a pair (r, h) such that r is applicable to G with witness h .

Definition 3.3: Consider a graph $G = (V, E, l)$ and an action (r, h) on G with $r = (L, R)$. The *application* of (r, h) to G yields a new graph $G' = (V, E', l')$ defined by

$$\begin{aligned} E' &= (E - \{h(x)h(y) \mid xy \in E_L\}) \\ &\quad \cup \{h(x)h(y) \mid xy \in E_R\} \\ l'(x) &= \begin{cases} l(x) & \text{if } x \notin h(V_L) \\ l_R \circ h^{-1}(x) & \text{otherwise.} \end{cases} \end{aligned}$$

We write $G \xrightarrow{r, h} G'$ to denote that G' was obtained from G by the application of (r, h) .

Remark: The vertex set V of the the graph G in the above definition remains the *same* upon the application of a rule. Only the presence or absence of edges and the labels on the vertices change.

A set of rules (i.e. a *grammar*) defines a concurrent (see Section VI) algorithm for a group of robots to follow. We suppose that each robot has a copy of the rule set initially and that by communicating with other nearby robots they can decide in a distributed fashion if there is some applicable rule and, if so, apply it. Continuing this process changes the states of the robots and their connections to other robots. The result is a system with a well defined set of behaviors, or trajectories. These are the objects we examine in this paper.

Definition 3.4: A *system* is a pair (G_0, Φ) where G_0 is the *initial graph* of the system and Φ is a set of rules (called the *rule set* or *grammar*).

We sometimes refer to a system simply by its rule set Φ and frequently suppose that the initial graph is the infinite graph defined by

$$G_0 \triangleq (\mathbb{N}, \emptyset, \lambda x.a) \quad (1)$$

where $a \in \Sigma$ is the *initial symbol* (here $\lambda x.a$ is the *lambda calculus* notation for the function assigning the label a to all vertices). The idea is to model systems with vast numbers of robots or parts all of which have the same initial internal state.

Definition 3.5: A *trajectory* of a system (G_0, Φ) is a (finite or infinite) sequence

$$G_0 \xrightarrow{(r_1, h_1)} G_1 \xrightarrow{(r_2, h_2)} G_2 \xrightarrow{(r_3, h_3)} \dots$$

If the sequence is finite, then we require that there is no rule in Φ applicable to the terminal graph. We denote the set of trajectories of a system by $\mathcal{T}(G_0, \Phi)$.

We denote a trajectory by, for example, $\sigma \in \mathcal{T}(G_0, \Phi)$ and use the notation σ_j to mean the j^{th} graph in the trajectory.

Remark: Many authors call G_0 the *axiom* and $\sigma \in \mathcal{T}(G_0, \Phi)$ a *derivation* when the goal is to study the language generated by the grammar. By our choice of terminology in the present paper we intend to emphasize the dynamical properties of grammars. For example, the grammars we define can exhibit non-trivial limit-cycles.

Definition 3.5 defines what many authors call the *interleaving semantics* of a concurrent system [26]. In this model, actions cannot occur simultaneously. However, a set of successive actions that operate on disjoint parts of a graph may be interleaved in any order and still produce essentially the same behavior. In Section VI we provide an alternate characterization of trajectories that better accounts for concurrency. In either case, we see that the pair (G_0, Φ) defines a non-deterministic transition system whose states are the labeled graphs over V_{G_0} . Non-determinism arises due to the fact that, at any given step, several rules in Φ may be simultaneously applicable, each possibly via several different witnesses. Our goal is to reason about all such behaviors of a given system.

B. Reachable Graphs

Given a system (G_0, Φ) of robots and rules, one naturally wishes to know (1) what types of aggregates can be built and (2) which of these are “finished products” with regards to the rule set.

Definition 3.6: A graph G is *reachable* by the system (G_0, Φ) if there exists a trajectory $\sigma \in \mathcal{T}(G_0, \Phi)$ with $G = \sigma_k$ for some k . The set of all such reachable graphs is denoted $\mathcal{R}(G_0, \Phi)$.

For assembly problems, we are particularly interested in the connected components of reachable graphs, as these correspond to aggregates of robots that are connected by physical or communication links. Recall that a graph G is *connected* if any pair of vertices can be connected by a sequence of edges in G . The connectivity relation partitions any graph into connected *components*.

Definition 3.7: A connected graph H is a *reachable component* of a system (G_0, Φ) if there exists a graph $G \in \mathcal{R}(G_0, \Phi)$ such that H is a component of G . The set of all such reachable components is denoted $\mathcal{C}(G_0, \Phi)$.

A reachable component may be temporary. That is, there may be some rule in Φ that operates on part of it. Other components are permanent: once they show up in a trajectory, they remain forever.

Definition 3.8: A component $H \in \mathcal{C}(G_0, \Phi)$ is *stable* if, whenever H is a component of $G_k \in \mathcal{R}(G_0, \Phi)$ via a monomorphism f , then H is also a component via f of every graph in $\mathcal{R}(G_k, \Phi)$. The stable components are denoted

$\mathcal{S}(G_0, \Phi) \subseteq \mathcal{C}(G_0, \Phi)$. Reachable components that are not stable are called *transient*.

In other words, the stable components are those that no applicable rule can change. Note, however, this does not mean that a stable component may not take part in an action, merely that it is left unchanged by it.

We illustrate these definitions with several examples in Section IV. In Section VII we will describe rule-synthesis algorithms that solve the problem of how to assemble a unique stable component:

Problem 3.1: (Self-Assembly Synthesis) Given any graph H and an initial graph G_0 , find a set of (preferably small) rules Φ such that $\mathcal{S}(G_0, \Phi) = \{H\}$.

IV. EXAMPLES

A. Paths and Cycles

We illustrate the definitions in Section III by examining a simple system that assembles paths and cycles from individual parts. Define a constructive rule set by

$$\Phi_1 = \begin{cases} a \ a \ \rightarrow \ b - b, & (r_1) \\ a \ b \ \rightarrow \ b - c, & (r_2) \\ b \ b \ \rightarrow \ c - c. & (r_3) \end{cases}$$

We have named the rules r_1 , r_2 and r_3 . Recall that we use the position of the vertices in the presentation of the rules to denote the re-labeling. For example, the first rule in Φ corresponds to

$$\begin{aligned} L &= (\{1, 2\}, \emptyset, \lambda x.a) \text{ and} \\ R &= (\{1, 2\}, \{12\}, \lambda x.b). \end{aligned}$$

We suppose that the rule set is used by a very large set of robots all initially labeled by the symbol a . Thus, the initial graph is defined by Equation (1). At first, the only applicable rule is r_1 since initially no vertices are labeled by either b or c . After r_1 is applied, both r_1 and r_2 become applicable. The rule r_3 eventually becomes applicable as soon as two unconnected vertices labeled b arise.

An example trajectory of (G_0, Φ_1) is shown in Figure 3. In the figure, the relative positions of the nodes denote the identities of the nodes and do not change after each rule application. The names of the rules used in each action are shown above the arrows. Witnesses are not shown because they can be inferred from the diagram.

It is apparent that the grammar produces paths P_k ($k \geq 1$) starting and ending with vertices labeled b and with internal vertices labeled c , except for the length-one path, which is labeled a . The grammar also produces cycles C_k ($k \geq 3$) with all vertices labeled by c . This can be proved by explicitly constructing a sequence of actions that realizes any given component. No other components are reachable, which can be shown by induction: The initial graph contains only P_1 . If at any point in a trajectory only paths and cycles have been produced then the application of any rule in Φ can only: make P_2 from two copies of P_1 (using r_1); make a path of length P_{j+1} from copies of P_j and P_1 (using r_2); make a path of length P_{j+k} from copies of P_j and P_k (using r_3); make a cycle C_k from a copy of P_k (using r_3). Thus we have:

Proposition 4.1: $\mathcal{C}(G_0, \Phi_1) = \{P_1, P_2, P_3, \dots\} \cup \{C_3, C_4, C_5, \dots\}$.




Fig. 3. A trajectory of the system defined in Example IV-A. The relative positions of the nodes denote the identities of the nodes and do not change after each rule application. The rules used are shown above the arrows. The monomorphisms witnessing the actions should be clear from the figure. For convenience (to the discussion in Section VI-C) the actions are numbered sequentially by the integers 1, ..., 8 appearing below the arrows.




Fig. 4. An example trajectory of the system defined in Section IV-B.

The stable components of (G_0, Φ) are the cycles, since no rule in Φ has a left hand side with a vertex labeled c . The paths are all transient.

Proposition 4.2: $\mathcal{S}(G_0, \Phi_1) = \{C_3, C_4, C_5, \dots\}$.

B. A Uniquely Stable Cycle

In the previous example, all cycles were stable components. There is, in fact, a fundamental limitation, as Theorem 5.1 implies (Section V): a rule set containing only binary rules cannot have a stable set consisting of exactly one cycle. However, using larger rules we can “stabilize” a cycle of a particular size. For example, consider the following rule set.

$$\Phi_2 = \begin{cases} a \ a \ \rightarrow \ b - c, & (r_1) \\ a \ c \ \rightarrow \ e - d, & (r_2) \\ a \ e \ \rightarrow \ g - f, & (r_3) \\ \begin{array}{c} d \\ \diagdown \quad \diagup \\ b \quad f \end{array} \ \rightarrow \ \begin{array}{c} d_1 \\ \diagdown \quad \diagup \\ b_1 \text{---} f_1 \end{array}, & (r_4) \\ \begin{array}{c} f_1 \\ \diagdown \quad \diagup \\ b_1 \quad g \end{array} \ \rightarrow \ \begin{array}{c} f_2 \\ \diagdown \quad \diagup \\ b_2 \text{---} g_1 \end{array}, & (r_5) \\ b_2 - f_2 \ \rightarrow \ b_3 \ f_3. & (r_6) \end{cases}$$

Once again, we take the initial graph to be G_0 , defined in Equation (1).

An example trajectory for Φ_2 is shown in Figure 4. The three constructive binary rules yield chains of length 4. The two ternary rules “triangulate” the cycle. The last rule removes the first triangulating edge to yield a copy of the cycle C_4 , which is the unique stable graph of the system.

Proposition 4.3: $\mathcal{S}(G_0, \Phi_2) = \{C_4\}$.

If instead of the “scaffolding” rules above, we simply used the rule

$$b \ g \ \rightarrow \ b_1 - g_1,$$

to close P_4 , then C_4 would indeed be stable, but so would C_8 , C_{12} and so on. This is because, for example, two copies of P_4 could combine to form C_8 via two applications of the above rule. In Section VII, we describe a general procedure for building cyclic graphs via triangulation.

C. Rules for Programmable Parts

The triangular programmable parts shown in Figures 1 and 2 each have three latches with which they can connect to other parts. We associate a label with each of these latches. The *state* stored by the micro-controller is then a triple of labels, and communication between parts must compare a pair of triples against a rule set. As an example, define a simple rule set with the following rules:

$$\begin{array}{ccc} \begin{array}{c} a \\ \diagdown \quad \diagup \\ a \end{array} \ \begin{array}{c} a \\ \diagdown \quad \diagup \\ a \end{array} & \rightarrow & \begin{array}{c} x \\ \diagdown \quad \diagup \\ b \end{array} \ \begin{array}{c} c \\ \diagdown \quad \diagup \\ c \end{array} \ \begin{array}{c} b \\ \diagdown \quad \diagup \\ x \end{array} \\ \begin{array}{c} a \\ \diagdown \quad \diagup \\ a \end{array} \ \begin{array}{c} b \\ \diagdown \quad \diagup \\ x \end{array} & \rightarrow & \begin{array}{c} x \\ \diagdown \quad \diagup \\ b \end{array} \ \begin{array}{c} c \\ \diagdown \quad \diagup \\ c \end{array} \ \begin{array}{c} * \\ \diagdown \quad \diagup \\ x \end{array} \\ \begin{array}{c} x \\ \diagdown \quad \diagup \\ * \end{array} \ \begin{array}{c} b \\ \diagdown \quad \diagup \\ x \end{array} & \rightarrow & \begin{array}{c} x \\ \diagdown \quad \diagup \\ * \end{array} \ \begin{array}{c} c \\ \diagdown \quad \diagup \\ c \end{array} \ \begin{array}{c} \# \\ \diagdown \quad \diagup \\ x \end{array} \end{array}$$

This presentation uses the “wildcard” symbols $*$ and $\#$ to stand for either b or c . Thus, the second line above is a schema representing two rules and the third line is a schema representing four rules. The symbols a , b and c are used in a way similar to that in Example IV-A. The symbol x is used to effectively “turn off” an edge, in that this symbol does not appear in the left hand side of any rule.

In Figure 5 we shown two reachable components of this grammar, with the actual geometry of the parts shown as well. Note that the orientation of the graphs in the above rules is not specified (i.e. the x can be either immediately clockwise of the binding edge or counter-clockwise (part




Fig. 5. Two reachable components of the grammar described in Section IV-C.




Fig. 6. An example trajectory of the ratchet defined in Section IV-D.

(a) of the figure) because graphs define only the connection topology and not the geometry of the assemblies. In an actual implementation, however, one can easily guarantee that the x is always clockwise of the binding edge (part (b) of the figure). In this paper we do not address this problem or the problem of geometry in general.

D. A Ratchet

A rule set need not define an assembly process. For example, consider the following rule set.

$$\Phi_3 = \begin{cases} a-c \rightarrow d e, & (r_1) \\ e \begin{array}{c} \nearrow b \\ \searrow g \end{array} \rightarrow c \begin{array}{c} \nearrow b \\ \searrow h \end{array}, & (r_2) \\ b-h \rightarrow f-b, & (r_3) \\ d-f \rightarrow g-a. & (r_4) \end{cases}$$

A trajectory of Φ_3 is shown in Figure 4. The sequence starts with a cycle of robots labeled by a , b and c attached to substrate of robots labeled g (i.e. this structure defines G_0 for this example). As illustrated in the figure, the rule set “ratchets” the cycle along the substrate. The ternary rule is used to prevent the “loose” stage of the ratchet (in the second graph in the trajectory) from attaching to the wrong substrate vertex (i.e. the rule forces the vertex labeled e to attach to the next g in the sequence).

If the substrate of vertices labeled g is infinite (or circular), then $S(G_0, \Phi) = \emptyset$. The reachable components are all isomorphic to those shown in Figure 6.

E. Other Examples Not Covered in This Paper

Other examples to which the grammatical formulation is applicable include: Graph recognizers [16]; Distributed algorithms, such as consensus and leader election; Grammars that




Fig. 7. Two examples (above) of covers of a graph (below). The cover on the left is a 3-fold cover; that on the right is 2-fold.

describe geometry, such as tilings of the plane [27]; Distributed self-replication [28]; and Grammars for directing distributed tasks in cooperative control of vehicle swarms [29].

V. TOPOLOGICAL PROPERTIES OF GRAPH GRAMMARS

It is intuitively clear that a grammar containing only binary rules cannot distinguish between a cycle of length $2N$, and two identical cycles of length N (cf. [16]). This has implications to the problem of constructing uniquely stable components. We extend this intuition in to a much larger class of rule sets. Given a system (G_0, Φ) , we bound the size of the reachable and stable sets using a basic topological tool: covering space theory.¹ A clear and comprehensive introduction to these classical techniques can be found in the text by Hatcher [11, pp. 60-68].

Definition 5.1: Given a graph G , an n -fold cover of G is a graph \tilde{G} such that, equivalently:

- 1) There exists a label-preserving n -to-1 homomorphism $p : V(\tilde{G}) \rightarrow V(G)$ that preserves degree (i.e., the image of an degree k vertex is a degree k vertex).
- 2) There exists a label-preserving n -to-1 continuous map $p : \tilde{G} \rightarrow G$ which is a local homeomorphism.

The latter definition is the standard one in covering space theory; the former is particular to graphs and easier to verify. It is straightforward to demonstrate that these two definitions are equivalent. Examples appear in Figure 7. Note that in the case of a *trivial* 1-fold cover, p is an isomorphism.

Theorem 5.1: For (G_0, Φ) an acyclic rule set, $\mathcal{C}(G_0, \Phi)$ is closed under covers. In particular, $\mathcal{C}(G_0, \Phi)$ contains infinitely many isomorphism types of graphs if it contains any graph with a cycle.

Proof: Since we use tools from covering space theory in this proof, we will consider graphs to be topological objects: 1-d cell complexes in particular. All maps between graphs will be continuous maps, and covering projections will be local homeomorphisms as in the second half of Definition 5.1.

Assume that $H \in \mathcal{C}(G_0, \Phi)$ is a component of σ_k for some trajectory $\sigma = (\sigma_i, (r_i, h_i))$. Consider any n -fold covering projection $p : \tilde{H} \rightarrow H$ of this component. We will reverse the trajectory σ and lift this disassembly procedure to build a trajectory $\tilde{\sigma}$ with \tilde{H} a component of $\tilde{\sigma}_{nk}$.

Denote by $\tilde{\sigma}_{nk}$ the graph consisting of the disjoint union of \tilde{H} with n disjoint copies of the complement $\sigma_k - H$ (see

¹We have simplified the definitions and suppressed most of the explicit topological terminology for the sake of clarity. However, for the tools to apply, we need to view graphs as one-dimensional *cell complexes* or *simplicial complexes*. In the 1-D case, a cell complex is a graph where each edge is viewed explicitly as a copy of the unit interval $[0, 1] \subset \mathbb{R}$.




Fig. 8. Part of a trajectory and its lifting via a 2-fold covering for the system defined in Example IV-A. In this figure, the symbol the * after the yields relation denotes two applications of the rule r_i .

$$b \text{ --- } e \text{ --- } f \text{ --- } b \text{ --- } e \text{ --- } f$$

Fig. 9. A rule with this left hand side can de-stabilize the covers in Figure 7 even though it does not apply to the (stable) graph being covered.

Figure 8). Define the projection $p_k : \tilde{\sigma}_{nk} \rightarrow \sigma_k$ via p on \tilde{H} and via the obvious projection of the complementary copies.

The image of the right hand side of the k^{th} rule $h_k(R_k)$ is a subtree of σ_k . Since the rule set is acyclic, the *lifting criterion* [11, pp. 61-62] is automatically satisfied, and it follows that the inverse image $\tilde{R}_k := p_k^{-1}(R_k)$ is a disjoint union of isomorphic copies of R_k .

Replace each copy of R_k in \tilde{R}_k with its left hand side L_k , reversing the assembly step. There are n such replacements to be performed; any order is acceptable.² Denote by $\tilde{\sigma}_{(k-1)n}$ the graph obtained after all n replacements, and denote by \tilde{L}_k the disjoint union of n copies of L_k within $\tilde{\sigma}_{(k-1)n}$. Define a new projection map $p_{k-1} : \tilde{\sigma}_{(k-1)n} \rightarrow \sigma_{k-1}$ to be (1) p_k on the complement of \tilde{L}_k ; and (2) the natural projection $\tilde{L}_k \rightarrow L_k$ identifying the disjoint copies. This graph $\tilde{\sigma}_{(k-1)n}$ is clearly an n -fold cover of σ_{k-1} via p_{k-1} since labels and indices are preserved.

Continue this procedure inductively, lifting R_i to \tilde{R}_i , replacing it in parallel with copies of L_i , and then defining the projection map $p_{i-1} : \tilde{\sigma}_{(i-1)n} \rightarrow \sigma_{i-1}$. This terminates in a covering projection $p_0 : \tilde{\sigma}_0 \rightarrow \sigma_0$. Since $\sigma_0 = G_0$ and the lift of any discrete set is again a discrete set, we have that $\tilde{\sigma}_0$ is isomorphic to G_0 . Thus, $\tilde{\sigma} \in \mathcal{T}(G_0, \Phi)$ and $H \in \mathcal{C}(G_0, \Phi)$.

In the case where H possesses a cycle, one has that there are infinitely many non-isomorphic covers (corresponding to subgroups of the fundamental group of H : see [11, Thm. 1.38]). \square

The stable set, $\mathcal{S}(G_0, \Phi) \subset \mathcal{C}(G_0, \Phi)$, is, however, not necessarily closed under n -fold covers. Indeed, there may be additional rules in Φ which have “large” connected regions of \tilde{H} in their left hand sides. If the left hand sides are sufficiently large, then these rules can apply to covers even though H itself is inert. See Figure 9.

The following is a sample of the type of result that can be obtained using covering space theory.

Theorem 5.2: Assume that Φ is an acyclic rule set, and that the stable set contains a component $H \in \mathcal{S}(G_0, \Phi)$ but not any

of its covers. Then for each edge $e \in E_H$, there exists a rule in Φ whose left hand side contains a copy of every edge of some cycle in H passing through e .

Proof: Consider the 2-fold cover of H lifted along the edge e ; that is, take two copies of H , snip each copy of e , and chain them end-to-end to form the connected graph \tilde{H} , as in Figure 7[right] using the subgraph $e - f$. One checks that $p : \tilde{H} \rightarrow H$ is indeed a covering projection.

Theorem 5.1 implies that \tilde{H} is a component of some $G_k \in \mathcal{R}(G_0, \Phi)$. By hypothesis, this component is not stable; hence there is some rule $(r, h) \in \Phi$ applicable to \tilde{H} . Consider the image $p(h(L))$ of L in H . If this image were isomorphic to $h(L)$, then (r, h) would be applicable to H , contradicting the fact that H is stable. But it follows from the lifting criterion that the only subgraphs of H that do not lift to isomorphic copies in \tilde{H} are those having a loop in H passing through the edge e . \square

Corollary 5.1: If Φ is an acyclic rule set all of whose left hand sides have no edges, then $\mathcal{S}(G_0, \Phi)$ is closed under covers.

Proof: Since the left hand sides have no edges, there can be no long chains to wrap around loops in a cover. \square

These results are best interpreted as topological bounds on the maximal amount of communication required to build a unique stable graph. Additional results and extensions to non-acyclic rules or to initial graphs different than G_0 are possible if one carefully tracks cycles.

VI. CONCURRENCY

A. Interleaved Versus Concurrent Trajectories

A graph grammar essentially describes a (non-deterministic) *parallel* dynamical system. Two actions can be executed in parallel if they operate on disjoint parts of whatever is the current graph: that is, if they are physically independent of one another. In Definition 3.5, however, we give what is traditionally called an *interleaved semantics* for trajectories [26] where we suppose nothing truly parallel ever happens, but rather that any two actions that could occur³ in parallel are ordered in time somehow to create two different trajectories. From the point of view of concurrency, then, two trajectories as we have defined them could really be representatives of the *same* behavior.

In this section, we relate the interleaved semantics with a *concurrent* or *partial order* semantics [30]. Concurrency allows us to reason more naturally about the number of steps required to assemble a structure with a given graph grammar. The definition we give is similar to that found in the Petri Net literature [12], Higher Dimensional Automata [31] and State Complexes [13].

Definition 6.1: Let $A = \{a_1, \dots, a_k\}$ with $a_i = ((L_i, R_i), h_i)$ be a set of actions each applicable to a graph G . The set A is called *commutative* if for all i and j

$$h_i(L_i) \cap h_j(L_j) = \emptyset.$$

If A is commutative with respect to G_0 and $G_0 \xrightarrow{r_1, h_1} \dots \xrightarrow{r_k, h_k} G_k$, then we write $G_0 \xrightarrow{A} G_k$.

³We assume, in this paper, that actions occur essentially instantaneously or each take unit time to complete.

²These n rules are commutative, see Section VI.

When A is commutative, we are justified in writing $G \xrightarrow{A} G'$ since the graph G' is independent of the order in which the actions in A are applied. Using this notion, we form the set of *concurrent trajectories* of a system.

Definition 6.2: A *concurrent trajectory* of a system (G_0, Φ) is a (finite or infinite) sequence

$$G_0 \xrightarrow{A_1} G_1 \xrightarrow{A_2} G_2 \xrightarrow{A_3} \dots$$

where A_{i+1} is any (not necessarily maximal) commutative set of actions applicable to G_i for each i . If the sequence is finite, then we require that there is no rule in Φ applicable to the terminal graph.

One can compress an interleaved trajectory by applying several consecutive but physically independent actions simultaneously (under the working assumption that all rules execute in unit time). Conversely, any concurrent trajectory can be transformed into an interleaved trajectory by simply choosing some ordering of the elements within each commutative step A_i . We call two concurrent trajectories *similar* if they can be transformed into the same interleaved trajectory, and hence model essentially the same behavior. Clearly, similarity is an equivalence relation. We argue that within each equivalence class, there is a canonical representative which is optimal with respect to parallelizability.

B. Left-Greedy Concurrent Trajectories

Recall that a *partial order* (P, \preceq) is a set P with an order relation \preceq that is reflexive, antisymmetric and transitive [32, ch. 1]. An element $x \in P$ is *minimal* if there does not exist an element $y \in P$ with $y \preceq x$ and $y \neq x$. The set of all minimal elements of (P, \preceq) is denoted $\min(P, \preceq)$.

There is a natural partial order associated with any trajectory $\sigma \in \mathcal{T}(G_0, \Phi)$. Suppose that

$$\sigma_0 \xrightarrow{a_1} \sigma_1 \xrightarrow{a_2} \dots \xrightarrow{a_k} \sigma_k$$

is an interleaved trajectory where $a_i = ((L_i, R_i), h_i)$. Let $P = \{1, \dots, k\}$ and define a relation \sqsubset by

$$i \sqsubset j \Leftrightarrow i \leq j \text{ and } h_i(L_i) \cap h_j(L_j) \neq \emptyset.$$

The partial order relation \preceq is then given by the reflexive and transitive closure of \sqsubset . This partial order gives rise to a canonical concurrent trajectory as follows.

Definition 6.3: Let $\sigma \in \mathcal{T}(G_0, \Phi)$ denote a trajectory with partial order (P, \preceq) . Define

$$\begin{aligned} A_1 &= \{a_i \mid i \in \min(P, \preceq)\} \\ A_2 &= \{a_i \mid i \in \min(P - A_1, \preceq)\} \\ &\vdots \\ A_{j+1} &= \{a_i \mid i \in \min(P - \bigcup_{i=1}^j A_i, \preceq)\} \\ &\vdots \end{aligned}$$

Then the *left-greedy concurrent trajectory* arising from σ is defined to be $\bar{\sigma}$ where $\bar{\sigma}_0 = \sigma_0$ and

$$\bar{\sigma}_0 \xrightarrow{A_1} \bar{\sigma}_1 \xrightarrow{A_2} \dots \xrightarrow{A_j} \bar{\sigma}_j.$$

We denote the set of all left-greedy trajectories of a system by $\bar{\mathcal{T}}(G_0, \Phi)$.

Note that the left-greedy concurrent trajectory obtained from a given interleaved trajectory is certainly not the *only* concurrent trajectory that can be formed. However, these trajectories perform as many commutative steps as possible as early as possible, hence the name. They are both ‘canonical’ and ‘optimal’ as we demonstrate below.

Lemma 6.1: An arbitrary concurrent trajectory

$$\bar{\sigma}_0 \xrightarrow{A_1} \bar{\sigma}_1 \xrightarrow{A_2} \dots$$

is left-greedy if and only if for every i and every $a \in A_{i+1}$, the collection $A_i \cup \{a\}$ is not commutative.

Proof: This follows directly from Definition 6.3. \square

Theorem 6.1: Any finite concurrent trajectory (or prefix of an infinite trajectory) is similar to a unique left-greedy concurrent trajectory.

Proof: We begin with existence. Let

$$\bar{\sigma}_0 \xrightarrow{A_1} \bar{\sigma}_1 \xrightarrow{A_2} \dots \xrightarrow{A_n} \bar{\sigma}_n$$

be a concurrent trajectory. If $\bar{\sigma}$ is not in left-greedy form, then there is some step i and some action $a \in A_{i+1}$ for which $A_i \cup \{a\}$ is commutative, by Lemma 6.1. Form a trajectory $\bar{\sigma}'$ by shifting the action a from A_{i+1} to A_i . (If A_{i+1} is now empty, then it is deleted, and the sequence is re-indexed.) Shifting an action clearly does not change the similarity class of the trajectory. This shifting strictly decreases the positive quantity $\sum_i i \cdot |A_i|$; thus, repeating the process terminates in a trajectory which, by Lemma 6.1 is left-greedy.

To prove uniqueness, suppose $\bar{\sigma}$ has two different left-greedy forms

$$\bar{\beta} = \bar{\beta}_0 \xrightarrow{A_1} \bar{\beta}_1 \xrightarrow{A_2} \dots \quad \text{and} \quad \bar{\rho} = \bar{\rho}_0 \xrightarrow{B_1} \bar{\rho}_1 \xrightarrow{B_2} \dots$$

and suppose that $A_i = B_i$ for $i = 1, \dots, k$ but that $A_k \neq B_k$. Then, without loss of generality, there is some action $a \in A_k - B_k$. Because the two trajectories are similar and because $A_{k-1} = B_{k-1}$, it must be the case that $a \in B_{k+1}$. But then $B_k \cup \{a\}$ is commutative (since it is a subset of A_k), which is a contradiction. \square

Corollary 6.1: The left-greedy form of a trajectory has the minimal number of steps (commutative sets of actions) among all trajectories in its similarity class.

Proof: The left-greedy path minimizes the quantity $\sum_i i \cdot |A_i|$ within its similarity class. \square

Remark: These two results together are a restatement, in the language of partial orders, of geometric results about paths in cubical complexes [13]. The notion of similarity corresponds to homotopy of paths in a certain cubical complex, and the left-greedy trajectories correspond normal to forms for geodesics on these spaces. Furthermore, the reduction in the number of states seen when transforming from interleaved to concurrent trajectories is essentially the process of partial order reduction [33, Ch. 10], which may render graph grammars amenable to model checking.

This motivates the following:

Definition 6.4: For $H \in \mathcal{C}(G_0, \Phi)$ a reachable component, the *assembly time* $\tau_H(\bar{\sigma})$ of H in the trajectory $\bar{\sigma} \in \bar{\mathcal{T}}(G_0, \Phi)$ is the smallest integer i such that H is a component of $\bar{\sigma}_i$ (or




Fig. 10. (a) The partial order associated with the trajectory in Figure 3. The sets of commutative actions A_1 through A_3 as in Definition 6.2. (b) The left-greedy concurrent trajectory associated with the trajectory in Figure 3.

∞ if there is no such integer). The *best case assembly time* of H is the minimum of the set

$$\{\tau_H(\bar{\sigma}) \mid \bar{\sigma} \in \bar{T}(G_0, \Phi)\}.$$

The *worst case assembly time* of H is the maximum of the set

$$\{\tau_H(\bar{\sigma}) \mid \bar{\sigma} \in \bar{T}(G_0, \Phi) \text{ and } H \text{ occurs in } \bar{\sigma}\}.$$

Note that a given reachable component H may not occur in every trajectory. If this is the case, we could have declared its worst case assembly time to be ∞ . However, in the above definition, we only consider the worst case among those trajectories that assemble H , and so perhaps we should use the term “worst case realizable assembly time.” We believe this is a more useful and informative number for characterizing the assembly time of a component with respect to a given grammar. Of course, an even more useful characterization is the expected assembly time as a function of some stochastic interpretation of rule application. We leave this issue to a future paper.

C. Chains and Cycles Revisited

To illustrate the notions of concurrency and assembly time, we revisit the example system that assembles paths and cycles, described in Section IV-A. In Figure 3, the actions in a trajectory of this system are numbered 1, 2, ..., 8.

The partial order associated with the trajectory in Figure 3 is shown (as a *Hasse Diagram*) in Figure 10(a). The commutative sets A_1 , A_2 and A_3 from Definition 6.2 are also shown. The associated concurrent trajectory is shown in Figure 10(b). If this trajectory is called $\bar{\sigma}$, then from the figure it can be seen that, for example, the assembly time $\tau_{P_3}(\bar{\sigma})$ of P_3 is 2. Also, the assembly time $\tau_{C_6}(\bar{\sigma})$ of C_6 is 3.

In general, the best case assembly time for P_k (with $k > 2$) is 2 steps since we could use rule r_1 in parallel to make $k/2$ copies of P_2 and then use rule r_3 in parallel to join these copies of P_2 into a copy of P_k . The best case assembly time for C_k is similarly 2 steps or 3 steps depending on whether k is even or odd, respectively.

If we consider only those trajectories in which P_k occurs, then we see that the worst case assembly time is k : Rule r_1 is used to make P_2 and then r_2 is used repeatedly. Similarly, the worst case assembly time for C_k , in those trajectories in which it occurs, is also k steps.

Proposition 6.1: In the system (G_0, Φ_1) , the best case assembly time for P_k is 2, the best case assembly time for C_k

Algorithm 1 *MakeTree*(V, E)

Require: $T = (V, E)$ is an unlabeled tree

- 1: **if** $V = \{x\}$ **then**
 - 2: **return** $(\emptyset, \{(x, a)\})$
 - 3: **else**
 - 4: **choose** $xy \in E$
 - 5: **let** (V_1, E_1) be the component of $(V, E - xy)$ containing x
 - 6: **let** (V_2, E_2) be the component of $(V, E - xy)$ containing y
 - 7: **let** $(\Phi_i, l_i) = \text{MakeTree}(V_i, E_i)$ for $i = 1, 2$
 - 8: **let** u, v be new labels
 - 9: **let** $\Phi = \Phi_1 \cup \Phi_2 \cup \{l_1(x) \ l_2(y) \rightarrow u - v\}$
 - 10: **let** $l = (l_1 - \{(x, l_1(x))\}) \cup (l_2 - \{(y, l_2(y))\}) \cup \{(x, u), (y, v)\}$
 - 11: **return** (Φ, l)
 - 12: **end if**
-

is 2 if k is even and 3 if k is odd. The worst case assembly time for both P_k and C_k is k .

VII. SYNTHESIS ALGORITHMS

In this section we consider the problem of *Self-Assembly Synthesis* (Problem 3.1). We first describe an algorithm that constructs a rule set to assemble a given tree (i.e. an acyclic graph) and then use it to build an algorithm that constructs a rule set to assemble an arbitrary graph.

A. Trees

We define in Algorithm 1 a recursive procedure *MakeTree* that, given any tree T , produces a set of binary rules Φ_T so that $\mathcal{S}(G_0, \Phi_T) = \{T\}$ (up to isomorphism and not including labels). The procedure takes as an argument an unlabeled tree (V, E) and returns a pair (Φ, l) where Φ is a rule set and l is a labeling function on V . For convenience, the function l is here denoted by subset of $V \times \Sigma$.

The base case of the procedure (lines 1-2) labels the singleton graph with the label a . For the recursive step, an edge (x, y) is chosen and *MakeTree* is called recursively on the two components that result from removing (x, y) from E . The recursive calls return rule sets Φ_1 and Φ_2 and labeling functions l_1 and l_2 . The rule set Φ for (V, E) is constructed from these in line 9 and the labeling function l is constructed in line 10. The construction uses two new labels u and v that we suppose have not been used before by any recursive call to *MakeTree*.

Theorem 7.1: Let $(\Phi, l) = \text{MakeTree}(V, E)$ where (V, E) is an unlabeled tree. Then $\mathcal{S}(G_0, \Phi) = \{T\}$.

Proof: By general induction on the size of T . \square

Notice that the *MakeTree* procedure defines an acyclic assembly graph as shown, for example, in Figure 11. The structure of this graph essentially defines the partial order associated with any trajectory of the resulting system (G_0, Φ) . In some cases, the assembly graph is balanced and the concurrent assembly time of (V, E) , using Φ , is $O(\log |V|)$. However, in other cases the assembly graph cannot be balanced (as with, for example, a *star graph* consisting of a root vertex connected




Fig. 11. An example assembly graph produced by the *MakeTree* procedure.

directly to all other vertices). In this case, the assembly time of (V, E) is $O(n)$ since none of the rules produced by *MakeTree* commute. In either case, we may state:

Theorem 7.2: Let $(\Phi, l) = \text{MakeTree}(V, E)$ where (V, E) is an unlabeled tree. The best and worst case assembly times of T in the system $\mathcal{S}(G_0, \Phi)$ are both $O(|V|)$.

Of course, the best and worst case assembly times may be considerably better, depending on the structure of the input graph. For example, *MakeTree* will produce a rule set that assembles the chain C_n in $O(\log n)$ steps as long as the edge chosen in step 4 partitions the tree into approximately equal parts.

Note that the procedure $\text{MakeTree}(V, E)$ produces exactly $|V| - 1$ rules and uses $2|V| - 2$ labels.

B. Arbitrary Graphs

Given a graph G , Algorithm 2 defines a function MakeGraph that produces a rule set Φ_G such that $\mathcal{S}(\Phi_G) = \{G\}$ using rules that are *at most* ternary in size. A priori, this would appear to be difficult, given the constraints of Theorem 5.2. The ingredient that makes this low-communication synthesis possible involves building a temporary “scaffold” to close each cycle.

The procedure begins in lines 1-2 by finding a maximal spanning tree $T = (V, E_T)$ for (V, E) and constructing a rule set Φ for this tree using *MakeTree*. We choose r (for root) to be a vertex on the first edge chosen in the call to *MakeTree*. In a trajectory, when r is eventually labeled by $l(r)$, we can be certain that the tree is assembled and that we may begin to close the cycles.

The loop in lines 6-13 adds rules to Φ that, in effect, triangulate the spanning tree until the resulting graph contains a subgraph isomorphic to G . At each step, E_S denotes the edges that have been used in this triangulation.

The triangulation procedure, defined by *Triangulate* and called in lines 7-8 of *MakeGraph*, is as follows. For each chord $uv \in E - E_T$ we find a shortest path in (V, E_S) from r to u , and form ternary rules that add edges from r to each successive vertex along this path. We then update E_S to include these edges, and repeat the procedure for a minimal path from r to v . Then in lines 9-10 of *MakeGraph* we add an edge from u to v , which can be done with a ternary rule since r will be connected via the scaffolding to both u and v . Note that if the distance from r to v in *Triangulate* less

Algorithm 2 $\text{MakeGraph}(V, E)$

Require: $G = (V, E)$ is an unlabeled connected graph

- 1: **let** $T = (V, E_T)$ be a maximal spanning tree of (V, E)
- 2: **let** $(\Phi, l) = \text{MakeTree}(V, E_T)$
- 3: **let** r be a vertex on the first edge chose by *MakeTree*
- 4: **let** $x = l(r)$
- 5: **let** $E_S = E_T$
- 6: **for all** $uv \in E - E_T$ **do**
- 7: **let** $(\Psi_1, E_S, x) = \text{Triangulate}((V, E_S, l), r, u, x)$
- 8: **let** $(\Psi_2, E_S, x) = \text{Triangulate}((V, E_S, l), r, v, x)$
- 9: **let** y be a new label

- 10: **let** $\Phi = \Phi \cup \Psi_1 \cup \Psi_2 \cup \left\{ \begin{array}{c} a \qquad b \\ \swarrow \quad \searrow \quad \rightarrow \quad \swarrow \quad \searrow \\ l(u) \quad l(v) \quad l(u) \quad l(v) \end{array} \right\}$

- 11: **let** $E_S = E_S \cup \{uv\}$
 - 12: **let** $x = y$
 - 13: **end for**
 - 14: **for all** $v \in V$ **do**
 - 15: **if** $rv \in E_S - E$ **then**
 - 16: **let** $\Phi = \Phi \cup \{x - l(v) \rightarrow x \quad l(v)\}$
 - 17: **end if**
 - 18: **end for**
 - 19: **return** Φ
-

than 3, the loop on line 5 will not execute so that no rules are added.

After each new rule is added, we change the label on the root node, ensuring that the rules can only be applied in the order given. The triangulation for one chord cannot begin until the triangulation for the previous one is complete (i.e. none of the triangulation rules commute). After the loop on lines 6-14 has finished, the resulting rule set will have a single graph in its stable set, and this graph will contain a subgraph isomorphic G . The rules added by the loop on lines 14-18 serve to remove the excess edges between the root node r and the other vertices. The left side of each of these rules contains a label that only appears on the root node after the entire triangulation process is complete. In effect, the label on the root node tracks the progress toward completion of the triangulation. After all the rules given in lines 14-18 have been applied, the resulting graph will be isomorphic to G . As shown above, this graph will be the only element of the stable set for Φ . An example trajectory illustrating this procedure is shown in Figure 12. This discussion constitutes the following:

Theorem 7.3: Let $(\Phi, l) = \text{MakeGraph}(V, E)$ where $G = (V, E)$ is an unlabeled graph. Then $\mathcal{S}(G_0, \Phi) = \{G\}$.

The *MakeGraph* rules produce an essentially a serial process, once the spanning tree has been assembled. If there are $|E|$ edges in the input graph, then rules for adding $c \triangleq |E| - n + 1$ of them will be created by *MakeGraph* (the rest are created by the *MakeTree* rules). The end-vertices of each chord are, in the worst case, a distance $O(|V|)$ apart in the graph (V, E_S) in lines 7-8 of *MakeGraph*. Thus, $O(c|V|)$ rules will be added for each chord. Since these rules do not commute, the assembly times of the rules Φ produced by *MakeGraph* are easily characterized:

Theorem 7.4: Let $\Phi = \text{MakeGraph}(V, E)$ where (V, E) is




Fig. 13. The partial orders associated with the two possible trajectories from Φ_r .

The set S in the new labels denotes the state of a protocol for executing r . For example, the symbol $b:k:\emptyset$ is the state of a robot k labeled by b that has not started to execute r . The symbol $b:k:j$ is the state of a robot k that has initiated execution of r by communicating with a robot j labeled by $a:***$. The symbol $b:k:ij$ denotes a robot that has proceeded through the execution of r by communicating with robots i and j labeled by $a:***$ and $c:***$ respectively. Notice that in the rule (S5), where the edge between $c:***$ and $b:***$ is actually added, we require that the robots labeled $c:***$ and $b:***$ are both connected to some other robot j labeled with $a:***$. This restriction is possible since the parts have unique identifiers. It follows that for any G

$$\mathcal{R}(G, \Phi_r) = \mathcal{R}(G, \{r\}) \quad \text{and} \quad \mathcal{S}(G, \Phi_r) = \mathcal{S}(G, \{r\}).$$

This does not contradict Theorem 5.1 since, by assigning unique identifiers to robots, we violate the hypotheses of having an infinite supply of identically labeled vertices as the start graph G_0 . In fact, in the initial graph for the new system, we would suppose that each robot i is initially labeled by $a:i:\emptyset$.

There are two possible concurrent trajectories that can arise from the rules in Φ applied to the appropriate initial graph (one matching the left hand side of the original ternary rule). Either rule (S1) can be applied first or rule (S2) can be applied first. The resulting concurrent trajectories, shown in Figure 13 both have length four. Thus, we can conclude that *it is possible to implement a ternary rule in four concurrent binary steps*. This (the number four) is in a sense the *cost* of a ternary rule in terms of the basic currency of binary rules.

IX. PHYSICAL MODELS OF ASSEMBLY

We describe briefly several physical robotic systems, besides the programmable parts discussed in the introduction, that can implement graph grammars to self-organize. Each of the following sub-sections is only an overview meant to convince the reader that graph grammars can be easily employed in a variety of robotic systems.

A. Self-Motive Robots

In previous work [3] we showed how to use graph grammars as a basis for multi-robot formation forming. To each robot i we associate two sets: The set \mathcal{A}_i of robots j such that either i is “attached” to j or the graph $l(i)$ $l(j)$ matches a rule in

Φ , and the set \mathcal{R}_i of the robots where no rule matches. We then define an *artificial potential function* U_i of the form

$$U_i = \sum_{j \in \mathcal{A}_i} U_{\text{attract}}(x_i, x_j) + \sum_{j \in \mathcal{R}_i} U_{\text{repel}}(x_i, x_j)$$

where x_k is the position of robot k . The function U_{attract} is defined so that $-\nabla_{x_i} U(x_i, x_j)$ has the set $\|x_i - x_j\| = R$ as an attractor. The function U_{repel} is defined so that $-\nabla_{x_i} U(x_i, x_j)$ has the set $\|x_i - x_j\| = 0$ as a repeller. Here R is the desired distance between “attached” robots. Each robot i simply follows the negative gradient of U_i to move toward robots in \mathcal{A}_i and away from robots in \mathcal{R}_i . When two attracting robots come in close proximity, they execute the appropriate rule and change their states. This has the effect of changing \mathcal{A}_k and \mathcal{R}_k , and therefore U_k , for each k .

If the number of robots is finite, then deadlock can occur in the above system due to groups forming incompatible sub-assemblies. We thus add the rule

$$(V, E, l) \Rightarrow (V, \emptyset, \lambda x.a)$$

for each $(V, E, l) \in \mathcal{C}(G_0, \Phi) - \mathcal{S}(G_0, \Phi)$. With low probability, robots randomly choose to disassociate using this rule. Note that this rule may be large. A simple consensus algorithm can be implemented by the robots in each component to decide whether to execute the rule.

B. Stirred Robotic Parts

Now suppose that a large number of robotic parts *float* in a stirred fluid, instead of moving themselves. Upon colliding (by chance), two parts will latch onto each other or not based on whether their current labels match the left hand side of a rule in Φ (See Figure ??). If they do latch together, then they change their labels according to the rule. A similar scheme works for mixed or destructive rules (see Section VIII).

To model this system, we suppose that each part i with position x_i has a latch variable $L_{i,j} \in \{0, 1\}$ associated with every other robot j . When the parts come in close proximity, they communicate their current labels to each other in an attempt to find an applicable rule. If one is found, they both set their latch variable for the other to 1 and change their labels. Otherwise the latch variable is set to 0 and the parts bounce off each other. The dynamics of robot i are

$$m\ddot{x} = F_1(x_i, t) - c\dot{x}_i + \sum_{j \neq i} L_{i,j} L_{j,i} F_2(x_i, x_j) \quad (3)$$

where F_1 is a time varying force field that models the effect of the fluid on the part and

$$F_2(x_i, x_j) \triangleq -\nabla_{x_i} U(x_i, x_j) - b \frac{\dot{x}_i(x_j - x_i)}{\|x_j - x_i\|} (x_j - x_i)$$

is a damped nonlinear spring, with spring potential U , modeling the latching mechanism. We suppose that U has a minimum at $\|x_i - x_j\| = R$ as before. In other work [34] we describe how such a mechanism would work using capillary forces or with robots [4], [5].

⁴We use the notation ijk (or jki , etc) to represent the set $\{i, j, k\}$.

X. DISCUSSION

We have defined a class of graph grammars that describe self-organizing robotic systems. We focused on the properties of the reachable and stable graphs that rule sets produce. We presented two algorithms to synthesize rule sets for arbitrary trees and graphs as unique stable outputs. Topological results inform these algorithms: cyclic rules are used for the general case. Using the notions of concurrency described in Section VI, we determined the number of concurrent time-steps required to assemble the desired graph using these synthesized rule sets. In Section VIII we introduced infinite rule sets that can, with binary rules, implement ternary cyclic rules. To work around Corollary 5.1, we added a unique identifier to each node, in a similar fashion to symmetry-breaking methods found in distributed algorithms [35]. Finally, we discussed how graph grammar rules can be used as a basis for robot assembly by describing two physical models appropriate to the task.

The methods we propose show that algorithms for collective tasks, here distributed self-assembly of a prespecified structure, can be *engineered*. The model is one of controlling the local interactions of a system so that global properties result. We believe that similar methods will be applicable to other systems and may potentially provide a predictive model for distributed tasks observed in nature. There are many natural problems in our model to be explored more completely in the future such as: the complexity of assemblies in terms of the minimum number of concurrent steps required to build them; optimality of rules sets in terms of number of rules, number of labels, etc.; the synthesis of behaviors other than assembly; the inclusion of a stochastic model of rule application (this is introduced in [5]); the relationship between rule sets and the possible geometries they describe; and the exploration of physical instantiations of the idea, such as the programmable part [4].

Acknowledgments

Klavins is supported in part by NSF CAREER grant number #0347955. Ghrist is supported in part by NSF PECASE grant number DMS-0337713. Lipsky is supported in part by NSF VIGRE grant number DMS-9983160.

REFERENCES

- [1] K. Saitou and M. Jakiela, "Automated optimal design of mechanical conformational switches," *Artificial Life*, vol. 2, no. 2, pp. 129–156, 1995.
- [2] K. Saitou, "Conformational switching in self-assembling mechanical systems," *IEEE Transactions on Robotics and Automation*, vol. 15, no. 3, pp. 510–520, 1999.
- [3] E. Klavins, "Automatic synthesis of controllers for distributed assembly and formation forming," in *Proceedings of the IEEE Conference on Robotics and Automation*, Washington DC, May 2002.
- [4] J. Bishop, S. Burden, E. Klavins, R. Kreisberg, W. Malone, N. Napp, and T. Nguyen, "Self-organizing programmable parts," in *International Conference on Intelligent Robots and Systems*. IEEE/RSJ Robotics and Automation Society, 2005.
- [5] N. Napp, S. Burden, and E. Klavins, "The statistical dynamics of programmed robotic self-assembly," in *International Conference on Robotics and Automation*, 2006, submitted.
- [6] G. Chirikjian, "Kinematics of a metamorphic robotic system," in *International Conference on Robotics and Automation*, 1994.
- [7] K. Kotay, D. Rus, M. Vona, and C. McGray, "The self-reconfiguring robotic molecule: Design and control algorithms," in *Algorithmic Foundations of Robotics*, P. Agrawal, L. Kavraki, and M. Mason, Eds. A. K. Peters, 1998.
- [8] A. Nguyen, L. Guibas, and M. Yim, "Controlled module density helps reconfiguration planning," in *Workshop on the Algorithmic Foundations of Robotics*, Dartmouth, NH, March 2000.
- [9] B. Berger, P. Shor, L. Tucker-Kellogg, and J. King, "Local rule-based theory of virus shell assembly," *Proceedings of the National Academy of Science, USA*, vol. 91, no. 6, pp. 7732–7736, August 1994.
- [10] R. L. Thompson and N. S. Goel, "Movable finite automata (MFA) models for biological systems I: Bacteriophage assembly and operation," *Journal of Theoretical Biology*, vol. 131, pp. 152–385, 1988.
- [11] A. Hatcher, *Algebraic Topology*. Cambridge University Press, 2001.
- [12] W. Reisig, *Petri Nets: An Introduction*. Springer-Verlag, 1985.
- [13] A. Abrams and R. Ghrist, "State complexes for metamorphic robot systems," *International Journal of Robotics Research*, vol. 23, no. 7–8, pp. 809–824, 2004.
- [14] H. Ehrig, "Introduction to the algebraic theory of graph grammars," in *Graph-Grammars and Their Application to Computer Science and Biology*, ser. Lecture Notes in Computer Science, V. Claus, H. Ehrig, and G. Rozenberg, Eds., vol. 73, 1979, pp. 1–69.
- [15] B. Courcelle, *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*. MIT Press, 1990, ch. On Graph Rewriting: An Algebraic and Logic Approach, pp. 193–242.
- [16] I. Litovsky, Y. Métevier, and W. Zielonka, "The power and limitations of local computations on graphs and networks," in *Graph-theoretic Concepts in Computer Science*, ser. Lecture Notes in Computer Science, vol. 657. Springer-Verlag, 1992, pp. 333–345.
- [17] B. Courcelle and Y. Métevier, "Coverings and minors: Application to local computations in graphs," *European Journal of Combinatorics*, vol. 15, pp. 127–138, 1994.
- [18] E. Winfree, "Algorithmic self-assembly of DNA: Theoretical motivations and 2D assembly experiments," *Journal of Biomolecular Structure and Dynamics*, vol. 11, no. 2, pp. 263–270, May 2000.
- [19] N. Bowden, A. Terfort, J. Carbeck, and G. M. Whitesides, "Self-assembly of mesoscale objects into ordered two-dimensional arrays," *Science*, vol. 276, no. 11, pp. 233–235, April 1997.
- [20] Y. S. Smentanich, Y. B. Kazanovich, and V. V. Kornilov, "A combinatorial approach to the problem of self assembly," *Discrete Applied Mathematics*, vol. 57, pp. 45–65, 1995.
- [21] I. Chen and J. W. Burdick, "Determining task optimal modular robot assembly configurations," in *International Conference on Robotics and Automation*, 1995.
- [22] C. Jones and M. J. Matarić, "From local to global behavior in intelligent self-assembly," in *International Conference on Robotics and Automation*, Taipei, Taiwan, 2003.
- [23] W.-M. Shen, P. Will, and B. Khoshnevis, "Self-assembly in space via self-reconfigurable robots," in *International Conference on Robotics and Automation*, Taiwan, 2003.
- [24] P. J. White, K. Kopanski, and H. Lipson, "Stochastic self-reconfigurable cellular robotics," in *Proceedings of the International Conference on Robotics and Automation*, New Orleans, LA, 2004.
- [25] P. White, V. Zykov, J. Bongard, and H. Lipson, "Three dimensional stochastic reconfiguration of modular robots," in *Proceedings of Robotics Science and Systems*, Boston, MA, June 2005.
- [26] K. M. Chanday and J. Misra, *Parallel Program Design: A Foundation*. Addison-Wesley, 1988.
- [27] R. Ghrist and D. Lipsky, "Grammatical self assembly for planar tiles," in *The 2004 International Conference on MEMs, NANO and Smart Systems*, Banff, Canada, 2004.
- [28] E. Klavins, "Universal self-replication using graph grammars," in *The 2004 International Conference on MEMs, NANO and Smart Systems*, Banff, Canada, 2004.
- [29] J.-M. McNew and E. Klavins, "A grammatical approach to cooperative control: The wanderers and scouts example," in *Cooperative Control*, S. Butenko, R. Murphey, and P. Pardalos, Eds. Kluwer, 2005, in Press.
- [30] W. Reisig, "Partial order semantics versus interleaving semantics for cps-like languages and its impact on fairness," in *Automata, Languages and Programming*, ser. Lecture Notes in Computer Science, J. Paradaens, Ed. Springer, 1984, vol. 172, pp. 403–413.
- [31] V. Pratt, "Modeling concurrency with geometry," in *Proc. 18th ACM Symposium on Principles of Programming Languages*, 1991.
- [32] B. A. Davey and H. A. Priestley, *Introduction to Lattices and Order*. Cambridge University Press, 1990.
- [33] E. M. Clarke, O. Grumberg, and D. A. Peled, *Model Checking*. MIT Press, 2000.
- [34] E. Klavins, "Toward the control of self-assembling systems," in *Control Problems in Robotics*, A. Bicchi, H. Christensen, and D. Prattichizzo, Eds. Springer Verlag, 2002, pp. 153–168.
- [35] N. Lynch, *Distributed Algorithms*. Morgan Kaufmann, 1996.