Automatic Configuration Recognition Methods in Modular Robots

Abstract— **Recognizing modular robot configurations composed of hundreds of homogeneous or heterogeneous modules is a significant challenge. Matching a new modular robot configuration to a library of known configurations helps in identifying and applying control schemes on the new configuration and is a step towards self-repair. We present three different algorithms to address the problem of (a) matching and (b) mapping a new robot configuration onto a known library of configurations. The first method solves the problem using graph isomorphisms and can identify configurations that share the same underlying graph structure, but have different port connections amongst the modules. The second technique utilizes graph spectral techniques and proves useful in finding approximate or near matches to configurations that may not be in the library. The third algorithm exploits the unique structure of the problem for the particular robots used in our research to achieve impressive gains in performance and speed over existing techniques, especially for larger configurations. Together, these three algorithms provide a complete solution to the problem of configuration matching and mapping. Results and examples are provided to compare the performance of the three algorithms and discuss their relative advantages.**

I. Introduction

A modular robotic system comprised of reconfigurable units has been an active area of research for the past two decades. Various groups have created an array of different models with diverse control schemes (master/slave, entirely local, genetic algorithms) [1], [2], [3], [4], [5], [6], [7]. These systems are made with a number of repeated units that can be rearranged to form different configurations, each of which can be used for different applications. For example, a snake-like configuration may be good at going through small holes, while a monkey-like configuration may be better at climbing. As the number of modules increases, the number of possible configurations rapidly increases.

Automatic configuration recognition is the process by which a modular system can determine its own configuration without having it explicitly programmed. This has a variety of uses including self-repair. Robustness becomes an increasing concern as the number of modules increases, since the average proportion of defective modules is constant or may increase with the addition or complex configurations. This is particularly problematic when the modules are early prototypes that are already not very reliable. At the same time, having more of the same unit becomes advantageous for failure situations. If one breaks, it is simply replaced with an identical unit. A few groups [8], [9] have demonstrated diagnosis and simple replacement of a broken module. However, self-repair has only been demonstrated in very simple, ad-hoc situations and often requires reprogramming of modules.

A process of self-repair typically includes several steps, the first of which is recognizing and diagnosing failures. The second step is determining a process for repair, and the final step is executing the computed process. Automatic configuration recognition is useful in the second step of the self-repair process. One method of self-repair is to find a new configuration that is functionally similar to the current one (so the new configuration can continue the task) that does not contain the failing modality.

Finding the set of all functionally similar configurations is difficult as an understanding and breakdown of the function of possible tasks would be required as well as reasoning about failure modes. It is easier (but not easy) to find the set of all kinematically similar, or isomorphic configurations which would be a step towards finding functionally similar configurations. This is the main focus of this paper.

In addition to self-repair where the configuration of working modules is dynamic, finding kinematically similar configurations is also useful in scenarios where modules are assembled by hand. Ideally, a quick manual repair would involve adding or replacing a module on the system without having to reprogram the new added module nor the other modules in the system.

Beyond self-repair and user interface, development architectures can be constructed around the idea that the morphology of the robot determines the software that runs on the robot - function follows form. For example, if the robot is constructed into the shape of a snake, it should behave like a snake. If it's in the form of a dog, it should behave like a dog. The basic underpinnings for an interface like this would require the ability to recognize configurations from a library of configurations.

This work presents a comparison of three methods of matching a new configuration to a set of known configurations and mapping their physical labeled modules to their logical position in a control scheme. The three methods are:

1) an automorphism grouping method using *nauty*,
2) a spectral decomposition approach,
3) a heuristic graph search called 3DLL.

An analysis of these approaches is given with insights into understanding the physical relationships between the scalability and execution times of the different approaches.

Chen and Burdick [10] published related work on enumeration of non-isomorphic configurations and identifying kinematically similar structures in this set. The work presented here is a departure from enumeration and instead focuses on algorithmic approaches to configuration recognition and mapping to configurations in a known database. A positive match and mapping provides the network structure needed for isomorphic control.

Castano and Will introduce the matching of a physical arrangement of modules to a known configuration as *configuration discovery* [11]. While they address hardware and software processes for building a representation of the configuration, they don't address the matching problem, instead referencing the graph automorphism and *nauty* as a possible approach. In contrast, we present complete implementation details for three different algorithms each of which has its

own relative advantages. This includes a new algorithm that exploits the special structure of our problem and is able to achieve impressive improvements in speed and performance for certain cases. The implementation of these algorithms and analysis of the results achieved with them provides *new* insights into the problem and the feasiblity of employing the different techniques we present here.

A. Automatic configuration detection

A configuration of a modular self-reconfigurable robot is defined as the arrangement (connectivity) of modules into a single connected component. There are many ways of representing this connectivity, but it is typically done with a graph where the nodes of the graph are the modules and the edges represent the connection between modules.

In a homogenous system, all modules are identical having a number of connection ports c with each pair of ports having w multiple ways (e.g., orientations) of connecting. The number of different ways that n modules can connect into one connected component is very large approaching $(cw)^n$. This can be seen since each module added to a configuration can be added to any of $c(n-1)$ connection ports in w different orientations. However, many of these configurations are morphologically identical (isomorphic). In addition, many modules have symmetries such that attaching a module results in functionally identical morphologies though the control may need to be modified (e.g., mirrored modules might need control to have an opposite sense).

Automatic configuration recognition has two functions, 1) identifying a configuration (for example checking to see if the configuration is isomorphic to one in a library of configurations) 2) mapping the labeled modules in the physical configuration with the logical arrangement of the known configurations to which it is isomorphic. We call the first part configuration *matching* and the second part configuration *mapping*.

B. Graph representation of modular robots

Graphs are concise representations for modular robots and readily allow application of techniques from graph theory. A few principal tools derived from ideas in graph theory are employed in our methods of configuration recognition and mapping. In this section, we introduce notation and definitions related to work in later sections.

A modular robot configuration is often represented as a graph $G = (V, E)$. The vertices ($V = (v_0, v_1, \ldots, v_n)$) of the graph represent the modules while the edges ($E = (e_1, e_2, \ldots, e_c)$) of the graph represent the connections between the modules. Here n is the number of modules in the robot while c is the number of connections between the modules. The labels for the vertices are the unique node IDs corresponding to each module. Further, a mapping $f_E : E \rightarrow E$ is used to assign a particular edge type to each edge. The edge type represents the type of connection between the two modules (which is based on the ports that are connected to each other). The graph representation can be converted to a matrix representation using a *adjacency matrix (M)*. The adjacency matrix for a graph with n vertices and c edges is a $n \times n$ matrix with $M_{ij} = 1$ implying that vertex v_i is adjacent to vertex v_j and 0 otherwise. A special version of the adjacency matrix called the *port adjacency matrix* is defined as a $n \times n$ matrix A where A_{ij} represents the port number on module i that module j connects to. Instead of 1 as in an adjacency matrix the element holds an integer from 0 to c.

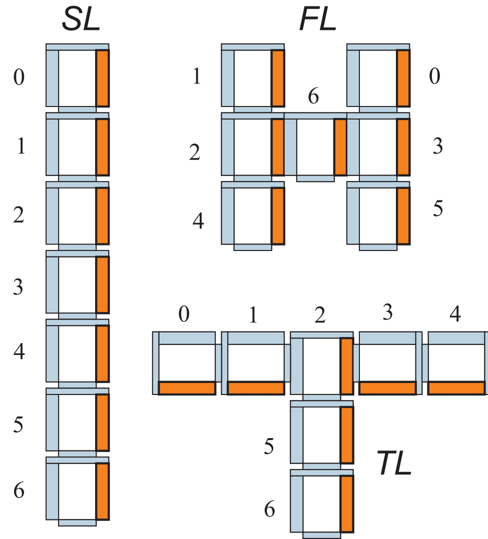

Fig. 1. A database of known configurations.

A graph G_1 is isomorphic to another graph G_2 if the adjacency matrices for the graphs are related by a set of row and column permutations. If γ_{12} is the isomorphism between the two graphs, then $M_{\gamma_{12}}(G_1) = G_2$. A graph may also be isomorphic to itself, i.e. there exists an isomorphism γ such that $M_\gamma(G_1) = M(G_1)$. The set of graph automorphisms is especially important to test against when trying to identify module configurations. The problem of modular robot configuration recognition and mapping can now be formally defined as follows,

Given a set of robot configurations $G = G_1, G_2, \ldots, G_n$ and a new modular robot configuration G_{new}, identify a robot configuration $G_i \in G$ that is isomorphic to G_{new}.

C. Organization

The paper is organized as follows. In Section II, we introduce a modular robotic system which serves as a good example system for connectivity analysis. Sections III, IV and V discuss the three implemented solutions with Section III serving to introduce the problem in more detail. These implementations were compared by testing on different sets of configurations including a set of randomly generated configurations. These tests and results are shown in each section and discussed in Section VI. Some experimental results of an embedded application are also shown in this section. Finally, Section VII talks about the implications of these results and describes future work.

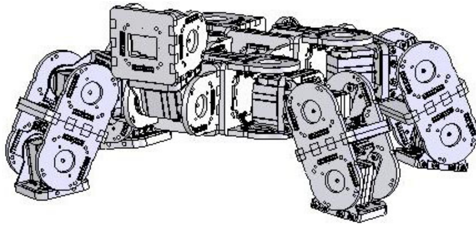

Fig. 2. *Surveyer Dog* 17 modules in a 4-legged configuration

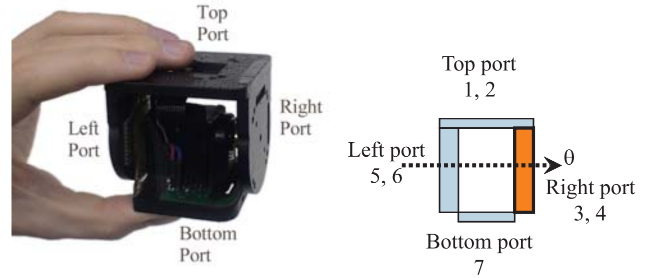

Fig. 3. One CKbot module with a schematic representation. The arrow indicates the rotational axis, the numbers are the port identification numbers assigned to each port

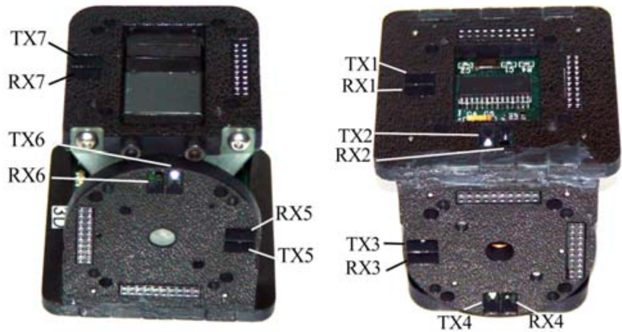

Fig. 4. Two IR transmitter and receiver pairs are on each side except the bottom port which has one pair.

II. CKBOT

While many of the current incarnations of modular robots are intended for self-reconfiguration, the work here applies to both self-reconfigured and human reconfigured modules. In both cases, the connection mechanisms, whether autonomous or manual, define the possible set of configurations.

A connector physically attaches two modules together and often connects power and/or communications as well. The connectors may be gendered (male and female) or they may be hermaphroditic (containing both male and female components). Genendered connectors connect only to the opposite gender whereas hermaphroditic connectors can be homogeneous, all connectors are the same and can connect to each other. For this paper we will consider only homogeneous hermaphroditic connectors.

A. CKbot Hardware

CKbot (Connector Kinetic roBot) is a modular reconfigurable robot that is used as the experimental platform in this work. The kinematics and connector strategy is typical for many chain style reconfigurable modular robots [2], [3], [5]. Figure 2 shows 17 modules in the configuration of a dog. Each module in the system consists of:

1) A laser cut plastic (ABS) body with a hobby servo actuator to control one rotational degree of freedom.
2) A controller (PIC18F2680) and associated hardware for implementing a Controller Area Network (CAN) communications protocol.
3) Four connector faces that pass the communications bus and power bus with an option of attaching at 90^o rotations. If power and communications are ignored, the faces have rotational symmetry order 4.

Two modules are attached together by using screws (4 pairs of holes are available, 4 threaded, 4 non-threaded). An electrical header is included in between the modules to facilitate the communications and electrical power bus. With this header the connection is homogeneous and hermaphroditic. Power can be supplied either from an external power supply or onboard Li-poly batteries that plug into the power/data ports on the module.

The module can be considered as a cube with connectors on top, bottom, left and right faces as in Figure 3. The top, left and right faces are rigidly mounted together, the bottom face is actuated to rotate up to 90^o to form the front or rear face of a perfect cube. Functionally the module has one symmetry where the module is rotated 180^o so left and right sides are swapped, though the actuator would need to be controlled in an opposite sense to be equivalent. The top and bottom connectors almost have the same symmetry, however the left and right faces are rigidly attached to the top, so a 180^o rotation swapping top and bottom results in a kinematic change (though not positional) in the left and right faces.

The system has a global communications bus allowing each module to talk to and discover which other modules are available. However, this bus does not indicate neighbor connectivity. Local communication allows each module to talk to its neighbors which inherently indicates the connectivity (presense or absence) of its neighbors. CKbot uses an infrared emitter/detector-based local communications mechanism.

Figure 3 shows a CKbot module with four connection ports and the 2D schematic representation of a module. Each port except the bottom port has two infra-red (IR) transmitter/receiver pairs. Figure 4 show the layout of the seven IR pairs. Note that when two faces are attached together, the transmitter LED (TX) faces directly on to the receiver photodiode (RX) on the opposing face and vice versa. Currently, 40 CKbot modules have been constructed and a variety of tasks have been demonstrated including moving like a snake, rolling like a tread, digging in sand and walking like a slinky toy. Section VI-A also describes a demo combining multiple gaits.

B. Low Level Software

The software used in CKbot builds upon the Robotics Bus [12]. The Robotics Bus was designed to be adaptable to a variety of applications based on the Controller Area Network (CAN) bus communication protocol. One feature of the Robotics Bus architecture is a broadcast *heartbeat* signal of each module about once every second. This signal contains a unique identifier allowing other modules on the bus to periodically identify the other modules present on the bus. Note that this necessitates the assignment of a unique **node ID** number to each module. The Robotics Bus also promotes *browseability*, *i.e.*, the ability of a central controller to query and identify the important attributes of a module. This includes the number and type of modules, important parameters like gain and scaling factors that can be set externally, the module's current state and sensor data associated with the module. For CKbot, the important information available on the bus includes the joint angle sensor for the module.

The unique node ID on CKbot means that the modules are labeled. Modular systems may either be labeled or unlabeled [3]. For many control schemes, module labels and configurations are coupled in a manner critical to the control. E.g. the methods rely on mapping of physical modules to logical configuration dependent locations. In one such method individual modules perform prescripted motions in a *gait control table* [13]. These control methods are applied only to specific configurations or configurations that can be extended in a regular pattern (e.g., making a snake longer, or increasing the number of legs in a centipede.) In most of these cases, the configuration of the robot and the position of each module in the robot is specified either manually, or determined using a set of rules for a small class of configurations.

Each CKbot module functions as an independent entity. Distributed control is possible with CKbot, though for this paper centralized control will be the focus as it is easier to implement and explain.

When power is applied to a CKbot system, the central controller logs *heartbeats* and additional information including positional feedback values and neighbor information of each module. The central controller can use this information to output a desired gait. Since the robots are sometimes hand assembled, the same configuration can be formed by putting together disparate sets of modules. The central controller must then recognize the manually assembled configuration and output the correct set of gait values to different individual modules. The *heartbeat* allows a central controller to recognize the addition of new modules or the removal (or failure) of modules in the robot while running.

C. Neighbor Discovery Subroutine

Neighbor discovery is the process where each module detects the presence or absence of neighboring modules and a representation of the connectivity of the whole system is generated. This was the central contribution in [11]. On CKbot, detection is done by the seven IR ports on each connection as shown in Figure 4.

There are 11 possible ways that the top, left, or right ports can be attached to another module and 7 possible ways for the bottom port. Each module has 7 IR pairs to determine each of these possibilities but one asynchronous serial communication device (on the PIC, a USART). The USART is multiplexed to each IR port.

The neighbor detection procedure can be distributed (as in [11]) or centralized. For CKbot a centralized controller is used. The controller designates each module (as discovered from heartbeats), one at a time as a transmitter and the others as receivers. The transmitter cycles through its 7 IR ports, waiting for a certain time on each port and transmitting its node ID on that port. The receivers cycle through checking their 7 ports at a rate 8 times faster than the transmitter. Any contact with the transmitter is logged in a neighbor table along with the node ID specified by the transmitter. Each module stores information about its own neighbors. A complete neighbor map can be constructed by the central controller by querying this information from the individual modules.

In this section we looked at some of the conventions and implementation details necessary to formulate the problem of configuration recognition and mapping. In the next three section, we will present the three algorithms used to solve this problem.

III. IDENTIFYING GRAPH ISOMORPHISMS

Traditional techniques from graph theory form the basis of the first approach to solving the configuration recognition problem. The process involves two steps, computing the graph isomorphism that relates the two configurations and then mapping the two configurations onto each other. The first problem is tackled by comparing a *canonical* representation for the underlying graph structure (the graph nodes without port information, i. e. the adjacency matrix, not the port adjaceny matrix) of the configurations. This pares the matching problem down to considering a fraction of configurations in the library. The second problem is tackled by comparing the complete automorphism group for the new configuration with the canonical representations in the library.

A. Example

Consider the set of known configurations in Figure 1. All the configurations have 7 modules and include a snake-like (*SL*), a four-limb (*FL*) and a three-limb configuration (*TL*). The goal is to match the new configuration *NL* in Figure 5 to one of the configurations in the database.

The *adjacency matrices* for the set of known configurations and the new configuration are generated by labeling the vertices in increasing order of node IDs. Vertex v_0 corresponds to the lowest node ID while vertex v_6 corresponds to the highest node ID. The adjacency matrices for the set of known

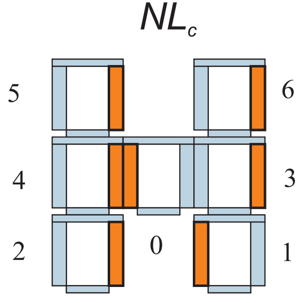

Fig. 5. A new configuration.

configurations in the database are given by:

$$M_{SL} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix},$$

$$M_{FL} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 \end{bmatrix}$$

$$M_{TL} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

The adjacency matrix for the new configuration is given by:

$$M_{NL} = \begin{bmatrix} 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix}$$

The graph isomorphism problem is approached using a software program called **nauty** [14]. **nauty** (**no aut**omorphisims, **y**es?), authored by Brendan McKay (and available electronically from `http://cs.anu.edu.au/~bdm/nauty/`), is a software program that determines a set of generators and size of the automorphism group of a graph. The input to **nauty** is the *adjacency matrix* corresponding to a particular robot configuration. **nauty** returns the automorphism group corresponding to the graph and a canonical representation for the input graph. Two graphs are isomorphic if they have the same canonical representation.

The canonical representations that are used here also represent only the connectivity of the underlying graphs corresponding to the different configuration, but do not contain any information about the port connections themselves. As we shall see later in this section, the port connection are essential in determining actual matches. However, the initial step of comparing canonical forms allows the search space for determining complete matches to be significantly narrowed down.

The first step, as mentioned earlier, is to find the canonical representations of known configurations. Let SL_c denote the canonical representation of a configuration *SL*.

The *canonical* representations for these configurations are computed using **nauty** and the mappings from the canonical representations to the configurations in the database are given below.

$$g_{SL} : \begin{pmatrix} 0 & 1 & 2 & 3 & 4 & 5 & 6 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 6 & 3 & 4 & 2 & 1 & 5 \end{pmatrix}$$

$$g_{FL} : \begin{pmatrix} 0 & 1 & 2 & 3 & 4 & 5 & 6 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 1 & 4 & 5 & 6 & 3 & 2 \end{pmatrix}$$

$$g_{TL} : \begin{pmatrix} 0 & 1 & 2 & 3 & 4 & 5 & 6 \end{pmatrix} \rightarrow \begin{pmatrix} 1 & 3 & 5 & 0 & 4 & 6 & 2 \end{pmatrix}$$

The gait for the reference configuration is also mapped to the gait for the canonical configuration. Given $(\phi_0(t), \ldots, \phi_i(t), \ldots, \phi_n(t))$ as the gait for the reference configuration *SL*, the gait for the canonical configuration of *SL* is given by $(\theta_0(t), \ldots, \theta_i(t), \ldots, \theta_n(t)) = (\phi_{g_{SL}(0)}, \ldots, \phi_{g_{SL}(i)}, \ldots, \phi_{g_{SL}(n)})$.

On examining the canonical configurations, *FL* and *NL* have the same canonical representation given by:

$$M_{FL_c} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 \end{bmatrix}$$

It is easy to see that the new configuration is a four-limb configuration. In addition to the canonical representation, a mapping g from the vertices of the canonical configuration FL_c to the vertices (and thus node IDs) of *NL* is also calculated by the algorithm. Thus, this part of the solution gives two results: (a) a match between *NL* and FL_c and (b) a mapping from the nodes of *NL* and FL_c. However, note from Figure 5 that the connections between nodes 0, 3 and 4 in configuration *NL* are different from the connections between nodes 4, 5 and 6 in configuration FL_c. This necessiates an additional test for complete matching where the symmetries of the individual modules are taken into account, using the *port adjacency matrix*.

B. Automorphisms and the port adjacency matrix

In the previous section, the isomorphism of the underlying graph structures was used to reduce the size of the search space to be considered for the matching problem. However, there are two additional issues in completely matching the new configuration to the reduced database of configurations. Consider, for

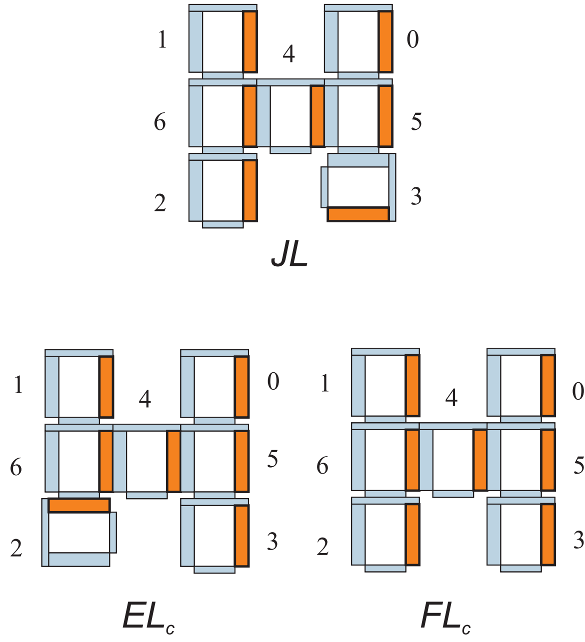

Fig. 6. Examining the automorphism group and the port adjacency matrix.

example, the set of configurations in Figure 6. Here, EL_c and FL_c are two canonical representations for configurations in the database while *JL* is a new configuration we are looking to identify. The graphs for all three configurations have the same canonical representation, yet only one pair match.

Looking at the figures it is clear that *JL* and FL_c are different configurations. The difference lies in the nature of the attachment between modules 3 and 5. Although this information is unavailable from just the adjacency matrices corresponding to these two configurations, it is available in the port adjacency matrices corresponding to these two configurations:

$$A_{FL_c} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 7 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 7 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 5 \\ 1 & 0 & 0 & 7 & 6 & 0 & 0 \\ 0 & 1 & 7 & 0 & 4 & 0 & 0 \end{bmatrix}$$

$$A_{JL} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 7 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 7 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 5 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 5 \\ 1 & 0 & 0 & 7 & 6 & 0 & 0 \\ 0 & 1 & 7 & 0 & 4 & 0 & 0 \end{bmatrix}$$

It is easy to see that the two configurations differ in row 4.

Now consider EL_c and *JL*. They have the same canonical representations, but different port adjacency matrices. However, with the left/right symmetries, they represent the same configuration. This can be seen on examining the *automorphism group* for *JL*. An automorphism denotes the isomorphism of the graph to itself and in this case EL_c and *JL* belong to the automorphism group for the underlying graph structure. Thus, during comparison, two configurations must be compared against each other and against the whole automorphism group for one of the configurations.

The last property we must consider in comparing the two configurations is the symmetry of the modules themselves. Again, this can be seen by comparing EL_c and *JL*. Modules 5 and 3 in *JL* map to modules 6 and 2 in EL_c. However, on matching the two configurations and aligning them with respect to each other, all the modules of *JL* are rotated 180^o such that the left and right ports are swapped. Thus, in addition to comparing the port adjacency matrices, we must also take into account symmetries in the module itself. This is easily done by listing all the symmetries explicitly and comparing against them when carrying out the port adjacency matrix comparison.

C. Final mapping

Using the techniques detailed in the previous sections, we can now find the mapping from configuration FL_c to configuration *NL* in our original example.

$$g : \begin{pmatrix} 0 & 1 & 2 & 3 & 4 & 5 & 6 \end{pmatrix} \rightarrow \begin{pmatrix} 6 & 5 & 2 & 1 & 0 & 3 & 4 \end{pmatrix}$$

Vertex v in the canonical configuration FL_c maps to $g(v)$ in the new configuration *NL*.

To find the correct mapping for the gait to the new configuration, align the two configurations using the graph isomorphism defined earlier and compute the direction of the joint axis of each module *with respect to* the joint axis of the module corresponding to vertex v_0 in the canonical configuration. Now, compare the direction of the joint axis for the vertex v_i (with respect to v_0) in the reference (database) configuration to the direction of the joint axis for the vertex $v_k = g(v_i)$ (with respect to $g(v_0)$) in the new configuration. The direction can either be parallel or anti-parallel. If the directions are parallel then θ_i in the gait maps to θ_k in the new configuration. However, if the directions are anti-parallel then θ_i in the gait maps to $-\theta_k$ in the new configuration. The desired goal of mapping the gait for the reference configuration $\begin{pmatrix} \theta_1(t) & \ldots & \theta_n(t) \end{pmatrix}$ onto the joint angles for the new configuration has been achieved.

Applying this procedure to *NL*, given the gait $(\theta_0(t), \ldots, \theta_i(t), \ldots, \theta_6(t))$ for the canonical configuration FL_c, the gait for *NL* is found to be by $(\phi_0(t), \ldots, \phi_6(t)) = (-\theta_4(t), -\theta_3(t), \theta_2(t), \theta_5(t), \theta_6(t), \theta_1(t), \theta_0(t))$.

D. General algorithm

The above procedure is presented as a general algorithm in Algorithm 1.

For our example problem from Figure 5, given a gait control table [13] for the canonical form of configuration *FL*,

$$\begin{bmatrix} -45 & -45 & 60 & 60 & 0 & 15 & 15 \\ -30 & -30 & 45 & 45 & 0 & 30 & 30 \\ -15 & -15 & 30 & 30 & 0 & 45 & 45 \\ 0 & 0 & 15 & 15 & 0 & 60 & 60 \\ 15 & 15 & 0 & 0 & 0 & 45 & 45 \end{bmatrix}$$

Algorithm 1 Nauty based robot configuration matching algorithm

Given a database of configurations $\mathbb{S}$, find the corresponding canonical representation of the database $\mathbb{S}_c$.
For new configuration E, compare the canonical representation of E (denoted by E_c) with all the elements of $\mathbb{S}_c$.
Store any matches in $\mathbb{M}_c$.
if $\mathbb{M}_c \neq \phi$ **then**
  Construct $\Phi_{\mathbb{M}_c}$ - the set of assembly port matrices for all elements of $\mathbb{M}_c$ under the action of the respective automorphism group.
  Compare Φ_{E_c} to $\Phi_{\mathbb{M}_c}$ to find a match, $F_c \in \mathbb{S}_c$, for E.
  if Match found **then**
    Generate the map for the robot from F to F_c.
    Align F_c and E and map orientation changes for F_c to E.

the mapping found earlier can be used to write a similar gait table for the new configuration *NL*:

$$\begin{bmatrix} 0 & -60 & 60 & 15 & 15 & -45 & -45 \\ 0 & -45 & 45 & 30 & 30 & -30 & -30 \\ 0 & -30 & 30 & 45 & 45 & -15 & -15 \\ 0 & -15 & 15 & 60 & 60 & 0 & 0 \\ 0 & 0 & 0 & 45 & 45 & 15 & 15 \end{bmatrix}$$

E. Implementation

The algorithm was implemented using MATLAB and a interface to *nauty*. The database of reference configurations was generated using a combination of known configurations and randomly generated configurations. Configurations with upto 200 modules were considered. The input to the algorithm is the port adjacency matrix for a new configuration and the output is the port adjacency matrix for a matching configuration from the database of known configurations. This is used to determine a mapping from the new configuration to the matching configuration in the database and thus the mapping for any gaits generated for that configuration.

In addition to matching random configurations, tests were also performed with specific configurations, specifically robot configurations that have been constructed and used in our research. Such robot configurations have been constructed with a maximum of 10 modules, however virtual configurations of the same form were created for testing for the higher number of modules. The configurations created included a snake-like serial line configuration, a centipede configuration, a loop configuration and a plane configuration.

F. Results

From the results in Figure 7 and Figure 8, it is clear that while the matching time increases with number of modules, there is considerable variation in matching times even for configurations with the same number of modules. The matching time is a function of the size of the automorphism group for the particular configuration to be matched. This can be clearly seen in Figure 10. Here, the matching times for configurations with 100 modules in each are plotted against the size of the

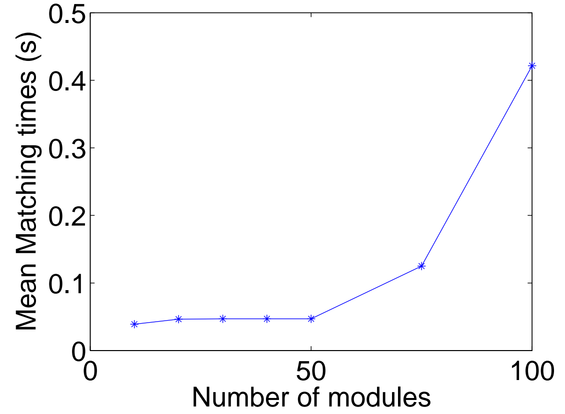

Fig. 7. Mean search times for nauty based algorithm vs. number of modules in robot configuration.

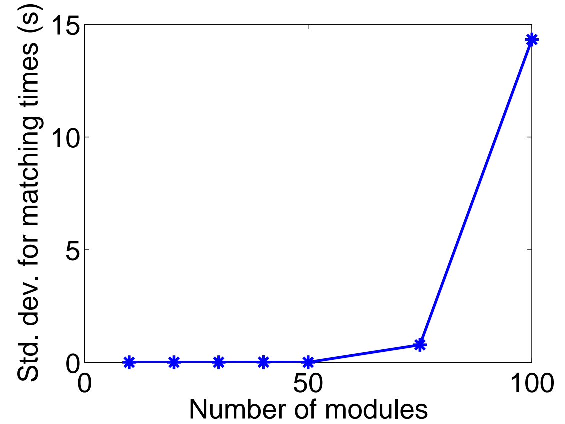

Fig. 8. Standard deviation of search times for nauty based algorithm vs. number of modules in robot configuration.

automorphism group for that particular configuration. It is obvious that the relationship between matching time and size of automorphism group is almost linear.

The size of the automorphism group is a reflection of the symmetries in the underlying graph structure of the robot. If the number of symmtries in the structure of the robot is high, the size of the automorphism group will be higher and the algorithm needs to check through more permutations of the graph to perform the matching. Thus, the algorithm will be slower for configurations that are very symmetric and fastest for asymmetric configurations where the only member of the automorphism group is the current configuration itself.

The results also bring out another advantage of using this particular method for lower number of modules (< 100). The canonical forms for the graphs are pre-computed and stored beforehand. This ability to preprocess the library greatly reduces the time required for matching since the canonical forms have to be computed only for the new configuration to be matched. The matching then only involves a matrix comparison followed by a permute and match operation for the isomorphic graphs. However, the method scales badly with increase in the number of modules in the configurations. In

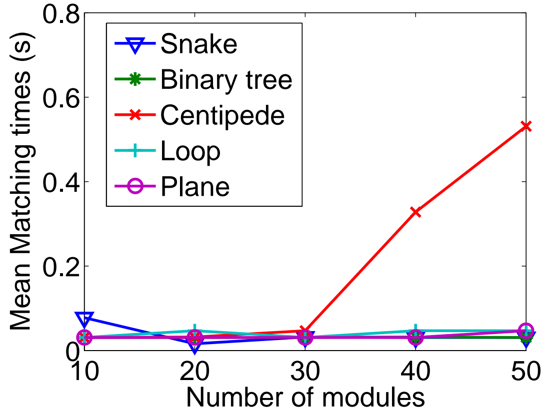

Fig. 9. Mapping times for specific configurations vs. number of modules in configuration.

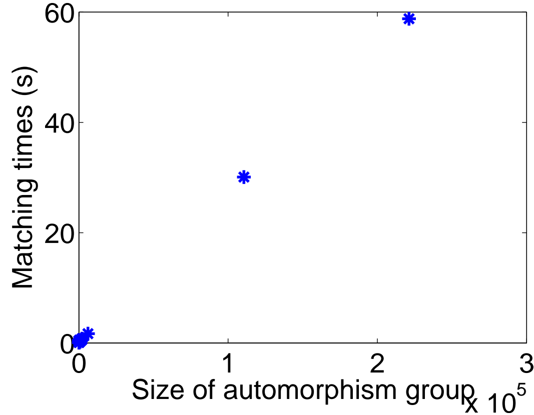

Fig. 10. Matching time vs. size of automorphism group.

fact, it was often difficult to obtain meaningful data with configurations of more than 200 modules. Since most hardware implementations so far have been less than 100 modules, this approach will still be useful for the near term.

The other advantage of this method is its ability to detect configurations whose port adjacency matrices may not match although they are isomorphic. Thus, it can detect kinematically equivalent configurations where the port numbers corresponding to connection between a pair of modules in the two configurations may not be the same. This is advantageous even in situations where the robots are put together simply manually since it spares the designer the tedious task of specifying the correct orientations for each module and modifying the gaits accordingly.

IV. Spectral Decomposition

The next approach we consider applies basic concepts from spectral graph theory as a means for determining configuration isomorphism and label mapping. A known method for checking for isomorphism between two graphs is through adjacency matrix spectral decomposition [15]. That is, if two graphs G_1 and G_2 are isomorphic, then the eigenvalues of the corresponding port adjacency matrices A_1 and A_2 are equal. The

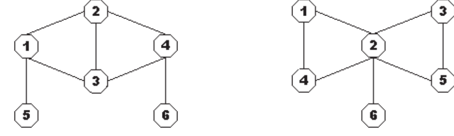

Fig. 11. Example of cospectral graphs with the same characteristic polynomial: $-1 + 4\lambda + 7\lambda^2 - 4\lambda^3 - 7\lambda^4 + 6\lambda^6$.

inverse, however, is not always true: adjacency matrices with identical arrays of eigenvalues are not necessarily isomorphic. To deal with this issue, permutations of eigenvector elements are employed as a confirmation of isomorphism and as a basis for finding the permutation mapping of module IDs.

The other methods considered in this paper use heuristics to search for matches between two graphs. This approach differs in its use of well-established ideas in spectral graph theory and its applicability to generate approximate methods [16] where the other techniques may fail.

Cospectral graphs such as the pair shown in Figure 11 [17] are rare and interesting cases that may arise in modular robotic configurations. The characteristic polynomial for these graphs are the same despite non-isomorphism. In these scenarios, although the eigenvalues are the same, the structures are not isomorphic since no relabeling of nodes maps one configuration to the other. A comparison of the eigenvectors for these graphs is required to find that permutation doesn't exist and confirm non-isomorphism.

Given two port adjacency matrices A_1 and A_2 with the same graph spectrum, we wish to find the *permutation matrix* P that reorders the rows and columns of A_1 so that they are identical to those of A_2:

$$A_2 = PA_1P^{-1}. \tag{1}$$

Note that P swaps the rows and P^{-1} swaps the columns of A_1 so that gaits for A_1 can be mapped onto corresponding gaits for A_2. The permutation matrix is composed of only one 1 across any row and column with the remaining entries as 0's. This gives the property that all permutation matrices are orthogonal, satisfying $P^T = P^{-1}$. The identity matrix is a permutation matrix that maps a configuration onto itself.

If A_1 and A_2 are decomposed into their Jordan canonical forms

$$A_1 = Q_1 \Lambda Q_1^{-1}$$
$$A_2 = Q_2 \Lambda Q_2^{-1}$$

where Λ is the diagonal eigenvalue matrix (note that they are same for both since A_1 and A_2 correspond to isomorphic graphs since the rows and columns are just interchanged) and Q_1 and Q_2 are the associated eigenvector matrices. This gives

$$\begin{aligned} Q_2 \Lambda Q_2^{-1} &= PQ_1 \Lambda Q_1^{-1} P^{-1} \\ &= (PQ_1)\Lambda(PQ_1)^{-1} \end{aligned}$$

which reduces to

$$Q_2 = PQ_1. \tag{2}$$

This shows that the permutation matrix also relates eigenvector elements of adjacency matrices of isomorphic configurations. Therefore, by mapping the columns of Q_2 to corresponding columns of Q_1, we can determine the permutation matrix that satisfies Equation 1.

To illustrate this method, consider configuration *FL* from Figure 1. Recall that the port adjacency matrix corresponding to this configuration is given by:

$$A_{FL_c} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 7 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 7 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 4 & 6 \\ 1 & 0 & 0 & 7 & 6 & 0 & 0 \\ 0 & 1 & 7 & 0 & 4 & 0 & 0 \end{bmatrix}$$

Now, consider another configuration with a port adjacency matrix:

$$A_2 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 7 & 0 \\ 0 & 0 & 0 & 1 & 7 & 0 & 4 \\ 0 & 0 & 7 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 7 & 1 & 0 & 0 & 0 & 0 & 6 \\ 0 & 0 & 6 & 0 & 0 & 4 & 0 \end{bmatrix}$$

. First, we note that A_{FL_c} and A_2 have the same characteristic polynomial:

$$\text{Det}(A_{FL_c} - \lambda I) = \text{Det}(A_2 - \lambda I) = \lambda^7 - 76\lambda^5 + 868\lambda^3.$$

This property suggests that these configurations are likely candidates for being isomorphic. To confirm this suggestion (and rule out that these structures are cospectral), we proceed further to find a permutation matrix that satisfies the property in Equation 2. We first compute the eigenvector matrices with columns ordered according to the roots of the characteristic polynomial (eigenvalue graph spectrum):

$$\lambda = \begin{bmatrix} 7.87 & 3.74 & -3.74 & -7.87 & 0 & 0 & 0 \end{bmatrix}$$

$$Q_1 = \begin{bmatrix} -0.47 & -0.72 & 0.72 & -0.47 & -0.97 & 0.32 & -0.58 \\ -0.31 & 0.48 & -0.48 & -0.31 & -0.18 & -0.93 & 0.12 \\ -0.04 & 0.06 & -0.69 & -0.04 & -0.02 & 0.11 & -0.35 \\ -0.06 & -0.10 & 0.10 & -0.06 & 0.06 & -0.07 & -0.42 \\ -0.53 & 0 & 0 & -0.53 & 0.09 & 0.02 & 0.58 \\ -0.52 & -0.38 & 0.38 & 0.52 & 0 & 0 & 0 \\ -0.34 & -0.25 & 0.25 & 0.34 & 0 & 0 & 0 \end{bmatrix}$$

$$Q_2 = \begin{bmatrix} -0.06 & -0.10 & 0.10 & -0.06 & 0.06 & -0.07 & -0.42 \\ -0.47 & 0.72 & -0.72 & 0.47 & -0.91 & -0.83 & 0.83 \\ -0.34 & -0.25 & -0.25 & -0.34 & 0 & 0 & 0 \\ -0.31 & -0.48 & 0.48 & 0.31 & 0.38 & 0.41 + 0.11i & 0.41 + 0.11i \\ -0.04 & -0.06 & 0.06 & 0.04 & -0.01 & -0.19 + 0.04i & -0.19 + 0.04i \\ -0.52 & 0.38 & 0.38 & -0.52 & 0 & 0 & 0 \\ -0.53 & 0 & 0 & 0.53 & -0.078 & 0.23 - 0.10i & 0.23 - 0.10i \end{bmatrix}$$

Note that the columns of Q_1 and Q_2 (the eigenvectors of A_{FL_c} and A_2) are both ordered so that they correspond to the same eigenvalue elements. The columns have been normalized so that the sum of the squares down any column equals one. Also, note that the absolute value of each eigenvalue element is of interest. To create the permutation matrix, note that Q_{2ij} (the element in the i^{th} row and the j^{th} column of Q_2) can be written as:

$$Q_{2ij} = P_{i1}Q_{11j} + P_{i2}Q_{12j} + \cdots + P_{i7}Q_{17j}.$$

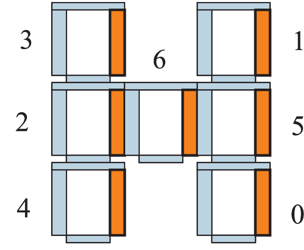

Fig. 12. Relabeling of A_{FL_c}.

The property that P has a single 1 across any column or row (with all other elements zero) allows us to build P simply by comparing the permutation of elements down corresponding columns in Q_1 and Q_2.

For instance, we see that $|Q_{111}| = |Q_{221}| = 0.470$. Consequently, $P_{21} = 1$ with all other elements in the rank and file of P_{21} equal to zero. Next, observe that $|Q_{121}| = |Q_{241}| = 0.31$. This gives us $P_{42} = 1$ with all other elements in the rank and file of P_{42} equal to 0. Similarly, $|Q_{131}| = |Q_{251}| = 0.04$ giving $P_{53} = 1$ with all other elements in the rank and file of P_{53} equal to 0. Continuing down the first columns of Q_1 and Q_2, we can construct the following P that confirms isomorphism and gives the desired module labels:

$$P = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}$$

The mapping is given by:

$$\pi_{1\to 2} = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 2 & 4 & 5 & 1 & 7 & 6 & 3 \end{pmatrix}$$

Generally, for a column k, when $|Q_{1jk}| = |Q_{2ik}|$, $P_{ij} = 1$ with all other elements across row i and down column j zero. The relabeled graph is shown in Figure 12.

It is evident that the choice of eigenvector for comparison is important. In the above example, if we had chosen any of the eigenvectors associated with the degenerate eigenvalue (zero), we would not have been able to build the permutation matrix. This redundancy in eigenvalues can be attributed to an algebraic regularity in the graph structure [18].

Structural symmetry creates interesting scenarios for this method to find the graph isomorphism mapping. Consider the following configurations in Figure 13.

The two rows of labels give the adjacency matrices

$$A_1 = \begin{bmatrix} 0 & 7 & 0 & 0 \\ 7 & 0 & 1 & 0 \\ 0 & 1 & 0 & 7 \\ 0 & 0 & 7 & 0 \end{bmatrix}$$

$$A_2 = \begin{bmatrix} 0 & 1 & 0 & 7 \\ 1 & 0 & 7 & 0 \\ 0 & 7 & 0 & 7 \\ 7 & 0 & 0 & 0 \end{bmatrix}$$

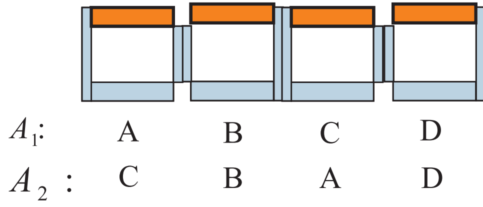

Fig. 13. Example of symmetric configuration with two permutation matrices that relabel modules in the same way.

with the following eigensystem:

$$\lambda = \begin{bmatrix} -7.52 & -6.52 & 6.52 & 7.52 \end{bmatrix}$$

$$Q_1 = \begin{bmatrix} -0.48 & 0.52 & -0.52 & 0.48 \\ 0.52 & -0.48 & -0.48 & 0.52 \\ -0.52 & -0.48 & 0.48 & 0.52 \\ 0.48 & 0.52 & 0.52 & 0.48 \end{bmatrix}$$

$$Q_2 = \begin{bmatrix} 0.52 & -0.48 & 0.48 & -0.52 \\ -0.52 & -0.48 & -0.48 & -0.52 \\ 0.48 & 0.52 & -0.52 & -0.48 \\ -0.48 & 0.52 & 0.52 & -0.48 \end{bmatrix}$$

Following the same approach as above, we end up with the following permutation matrix:

$$P^* = \begin{bmatrix} 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 \end{bmatrix}$$

The multiple 1's across the rows and columns occur because of the redundant elements in the eigenvectors. The reason this occurs is because two distinct permutation matrices both satisfy Equation 1, namely:

$$P^* = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & 7 & 0 & 0 \\ 7 & 0 & 1 & 0 \\ 0 & 1 & 0 & 7 \\ 0 & 0 & 7 & 0 \end{bmatrix} \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$
$$= \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 0 & 7 & 0 & 0 \\ 7 & 0 & 1 & 0 \\ 0 & 1 & 0 & 7 \\ 0 & 0 & 7 & 0 \end{bmatrix} \begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

where P^* is the union of P_1 and P_2 given by:

$$P_1 = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$P_2 = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix}$$

The reason two permutations occur for this configuration (and any isomorphic labeling of this graph), is because both matrices are in the graph's symmetry group. Also called automorphic group (as discussed earlier in Section III), we see that the symmetry is a 180^o rotation about the center of mass of the system. For the purpose of spectral decomposition approach, the algorithm tries valid combinations of P^* type unions (only one 1 in all rows and columns) until Equation 1 is satisfied. Therefore, this algorithm is slightly faster for asymmetric systems, as seen in Figure 16 which shows that the time to map random structures is, on average, less than the time of the more ordered structures (tree, snake, plane, centipede). The general algorithm is presented formally in Algorithm 2.

Algorithm 2 Configuration matching using spectral decomposition.

Compute the characteristic polynomial of the given structure.
if Characteristic polynomial matches any (pre-computed) characteristic polynomial in the library database **then**
  Decompose both adjacency matrices to give the eigenvector matrices.
  Compare the elements as described above to build the permutation matrix.
  (1) Do comparisons down all n eigenvector columns.
  (2) Record the column that produces a P with the minimum number of redundant ones.
  (3)
  if Any pair of eigenvector produces a permutation matrix that satisfies equation 1 **then**
    RETURN P.
  else
    Choose the optimal column (as determined in Step (2) above) and find a P^*.
    Choose the first of the $m!$ possibilities for each block (where m is the size of the block redundancy; the snake example above has two blocks of two).
    Compare the mapping using equation 1
    if mapping is valid **then**
      RETURN P.
    else
      Try the next of the $m!$ possibilities for each block.

Figure 14 shows the time to find a match using comparisons between graph spectra in a library of 200 random configurations (for each data point), for up to 1000 modules. The matlab function $eigs(A)$ was used to find the largest eigenvalues of the sparse matrices. In comparison with Figure 16, we see that the time to find the module mapping is more time-consuming. This figure compares the times to find the permutation matrix between two isomorphic configurations for up to 50 modules. Snake, centipede, plane, tree and random structures were tested. The random configuration times also include a negligible matching time to find the correct structure in a library of other random structures. Even so, the cumulative time to find the mapping is slightly less than the other structures. This can be attributed to the fact that there are less symmetries (on average) in random structures and this reduces the number redundant permutation elements that the algorithm must choose from in the method described above. For comparison, the brute-force $n!$ time is included for up to 12 modules. The data point itself is off the scale of the graph.

Some limiting considerations in this approach include nu-

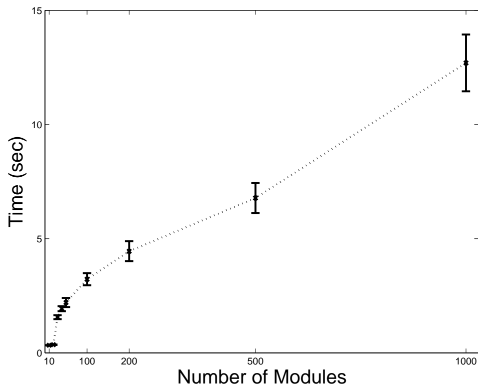

Fig. 14. Search times for the spectral based algorithm vs. number of modules in a random robot configuration.

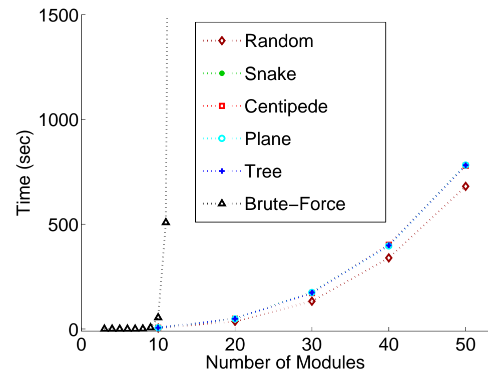

Fig. 16. Mapping times for specific configurations vs. number of modules in configuration.

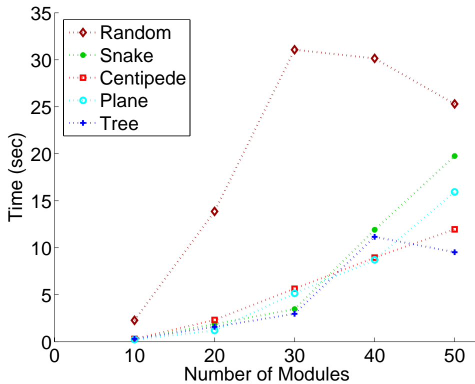

Fig. 15. Standard deviation of search times for spectral based algorithm vs. number of modules in robot configuration.

merical stability and structural symmetry of configurations. In particular, the number of elements of each eigenvector equals the number of modules. For very large matrices, these normalized eigenvectors are composed of small numbers that have accumulated rounding errors attributed to the LU-decomposition approach that the LAPACK matrix algebra package that Matlab employs. There are various methods that one can implement to deal with this issue (i.e., large normalizing factors, matrix balancing) but for numbers larger than 10^3, the problem becomes difficult to handle. Additionally, the time to compute eigenvectors becomes prohibitively large at that scale.

For highly symmetric structures (loops, planes, tori), the automorphic group is large and the algorithm has a large set of 1s to sort through to find an appropriate P that satisfies equation 1. Lastly, a small class of graphs called strongly regular graphs pose a problem in that they give many degenerate eigenvalues, each of which have eigenvectors that give no information as to the permutation of the structure [18].

V. 3DLL Approach

The approach presented in Section III attempts to use the underlying graph structure of a modular robot to match new configurations to the database. However, it is evident that there is more information in our problem that could be used to further reduce the search space. In particular, the port adjacency matrix stores additional information that can be easily incorporated into a heuristic approach.

Just as in Section III, where unique canonical forms of the automorphic group configurations in the library are precomputed, this method precomputes and stores a representation that exploits the physical nature of the modular robot configurations to generate nearly unique forms of the automorphic groups.

In this method, robots are represented by a 3-dimensional linked list of module objects (thus the name 3DLL), each of which contains a position and orientation in 3-space with respect to an *expansion origin* module. Each configuration can have multiple representations based on which expansion origin module is chosen. Expansion origins are distinguished from *comparison origins*, which are a subset of the potential expansion origins used as starting points when comparing configurations module-by-module. Two representations of the same configuration starting from two different expansion origins are related by a constant translation and rotation relating the two origins in a global reference frame.

To compare two configurations, several heuristic measures are evaluated first. If these fail to distinguish the two configurations, the algorithm proceeds by comparing the positions, orientations, and connection presences of each module in the configurations, starting from each comparison origin.

A. Library Creation

The first step in this approach is to create a representation of any new configuration. To do this, an expansion origin module must be selected. The selection process trims the number of possible expansion origin modules by allowing only modules with the degree that has the lowest multiplicity and then by giving priority to modules with the lowest combination of port values (i. e. an origin with connections on ports 1 and 3 takes priority over over an origin with connections on ports 2 and 4). If more than one expansion origin remains after this process, a module is randomly chosen from this set and designated as the expansion origin. The goal here is to restrict the number of possible origin modules as much as possible to limit the amount of time that must be spent comparing configurations.

The next step assigns the origin a default position and orientation and stores it in a 3-dimensional linked list. New modules are added in a depth-first fashion by choosing the unplaced module connected to the lowest port number of the most recently-placed module. At each step, the position and orientation of the next module is determined from the neighbor data, position, and orientation of the current module. A global list of modules that have been added is maintained to prevent loops from running indefinitely and to end the expansion when all the modules have been added. This expansion process is deterministic: it will always generate a representation which stores the modules in the same order when the choice of origin module is constant.

During the expansion process, several metrics are calculated for later use in comparing configurations. These include the center of mass of the system, degree multiplicity, and a port count (the total number of connections on each port). The final step determines the comparison origins of the configuration, which include only those expansion origin modules with the minimum distance to the center of mass. This minimum distance is stored for use as another heuristic check.

B. Matching and Mapping

Two identical configurations (i.e. configurations having the same geometrical shape, but not necessarily same module IDs) with the same choice of expansion origin module will have 3-dimensional linked lists of modules with the same positions, orientations, and connection presences. Therefore, comparing two representations is just a matter of stepping through the linked lists representing them and checking that these attributes match for each module in the list. The two representations are not the same if any two corresponding modules of the linked list do not match at any point. Note that this comparison has to be carried out in turn for each possible choice of origin module for one of the configurations until a match has been found or the set of origin modules is exhausted. Finding a match also finds the mapping between the two configurations since this is a simple assignment of node IDs stored in the linked lists for the two configurations.

To compare configurations, first several fast heuristics are evaluated in order of speed. These include graph invariants like the number of modules in a configuration and multiplicity of degrees of the configuration. Note that because of the special nature of our robots, the maximum degree is 4. If two configurations pass this first set of heuristic comparisons, then the linked list based representations of the two robots are compared starting from the comparison origin modules. The process of comparison of two configurations is formally presented in Algorithm 3.

C. Example

We illustrate this method further using the example considered earlier in Section III using configuration *FL* from Figure 1. Module 4 is the only module with degree 4 and is thus chosen as the expansion origin. The next module for addition to the linked list is 5, since it is attached to port 2 of the origin. Now, module 5 is expanded to obtain, in turn, modules 0 and 3. At each step, the relative position and orientation of each module with respect to the previous module is determined by the ports that attach the two modules together. The expansion process terminates when modules 6, 1 and 2 have been added, in turn, to the linked list representing the configuration. This procedure is shown in Figure 17.

Algorithm 3 3DLL configuration matching algorithm for a new configuration G_1 and a library representation G_2

```
Match = 0
if Number of modules in G_1 != Number of modules in G_2 then
    RETURN
Build the representation of G_1.
if Port counts of the two configurations are not equal then
    RETURN
if Degree multiplicities of the two configurations are not equal then
    RETURN
if Number of comparison origins n_co for the two configurations is not equal then
    RETURN
if Distance of the comparison origins to the center of mass for the two configurations is not close (within a threshold ϵ) then
    RETURN
for i = 1 to n_co do
    Set the current module in the new configuration to the first comparison origin and set the current module in the library configuration to be comparison origin i.
    Match = 1
    for j = 1 to n_G1 do
        if Current module positions, orientations, or connection presences do not match then
            Match = 0
            BREAK
        Update the module mapping with the two current modules.
        Set the current module for both configurations by choosing the unplaced module connected to the lowest port number of the most-recently-current module with ports that have not been used in this manner.
    if Match == 1 then
        BREAK
RETURN
```

D. Results

Tests were carried out for this approach using exactly the same procedure as the one used for the previous two approaches in Section III and Section IV. Results are presented here for both sets of tests, i.e. the test for configurations with different numbers of modules against a database of known configurations (Figure 18 and Figure 19) and the mapping test for specific robot configurations where they are mapped onto the same configuration with a different set of node IDs (Figure 20).

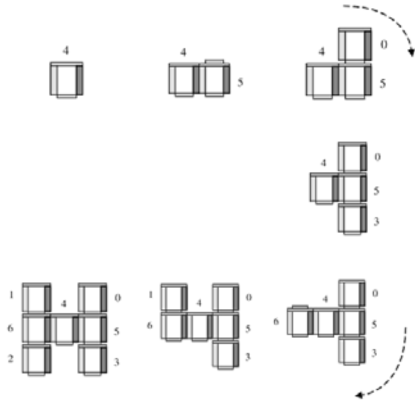

Fig. 17. Building a linked list based representation for a modular robot.

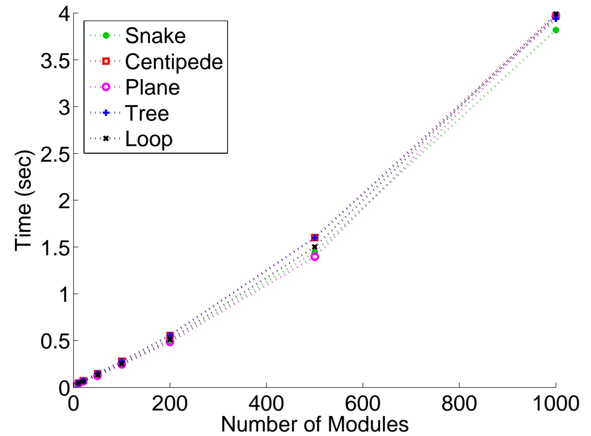

Fig. 20. Mapping times for specific configurations using 3DLL vs. number of modules in configuration.

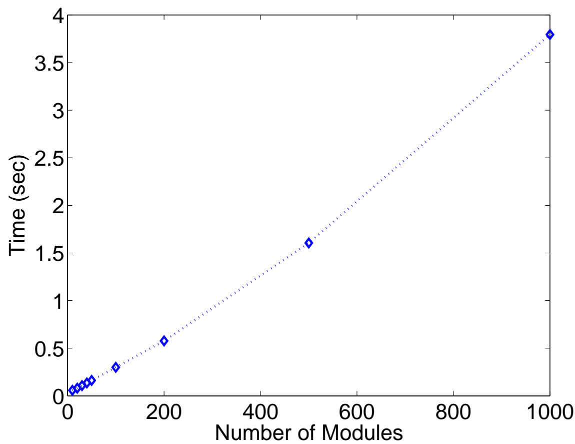

Fig. 18. Mean search times for 3DLL vs. number of modules in robot configuration.

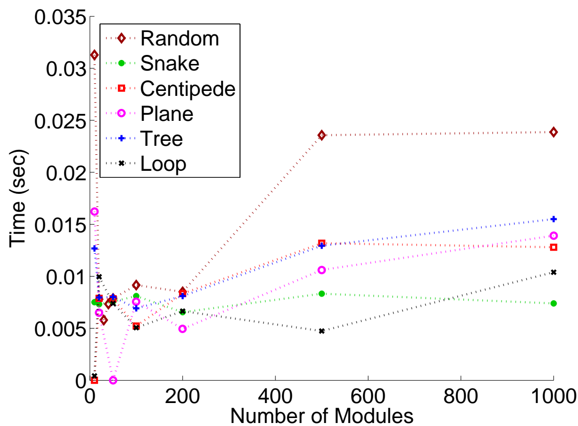

Fig. 19. Standard deviation of search times for 3DLL vs. number of modules in robot configuration.

Note that, for all five different configuration types, the scaling is nearly linear and the maximum amount of time is roughly four seconds. The time to search a library of 200 configurations and generate a mapping is nearly the same as the amount of time to just generate the mapping. While the standard deviation of the library test times is that of the times to find the 10th, 20th, ..., 200th configuration in the library, it is still similar to the mapping test standard deviations.

The mapping times for this algorithm are in general O($n \times n_{co}$) where n_{co} is the number of comparison origins. However, as can be seen from Figure 20, mapping times are in practice nearly O(n). Because of the constraint that origins must have the lowest multiplicity of degree, the number of origins is constant for the snakes, planes, and centipedes, and therefore the results scale linearly with the number of modules in the configuration. Similarly, because of the constraint that all comparison origins must have the minimum distance to the center of mass, the loops have only two possible comparison origins. The trees, while having anywhere from 1 to 63 modules with lowest multiplicity of degree, have only one possible comparison origin and therefore it too scales as O(n).

In practice, the heuristics filter out the majority of the configurations in the library, and step-by-step module comparison occurs primarily for identical configurations, i.e. the heuristics reduce the search space in the library to exactly one. Further, the number of potential comparison origins (those that must have the minimum distance to the center of mass) is nearly always one or two. These observations explain the relatively low standard deviation and the linearity of the library test results - the index of the matching configuration in the library is relatively unimportant because the majority of the time spent is building the representation of the new configuration and stepping through the modules of the new and matching configuration.

This method is able to recognize configurations in an amount of time reasonable even for configurations that are currently physically unrealizable because of the number of modules required. Further, it does so in a scalable way thanks to the extra information inherent in the port adjacency matrix.

Finally, because robot representations are stored in a form analogous to the actual robot, this method allows the inclusion of other robot features, such as different module types. They are simply added on as another check in the module-by-module comparisons in Algorithm 3.

VI. Discussion

The configuration matching and mapping problem can be presented as a task involving four steps:

1) Match the underlying graph, independent of ports.
2) Match the ports.
3) Choose an origin with which to start the mapping process.
4) Generate a module ID mapping between the matched configurations.

In the nauty-based method Step 3 is not needed as the canonical form includes the same origin for both the new and matched configuration. The first and last step are fast and linear in the number of modules. However Step 2 (matching ports) depends on the number of configurations in the automorphism group of the underlying structure which can sometimes be very large. As a library becomes larger, the size of the automorphism group can be exponential in the number of modules in a configuration worst case. As can be seen comparing Figure 7 with Figure 18 and Figure 14, the nauty-based method runs faster for lower numbers of modules ($<= 50$ modules). Some of this speed may be due to the implementation since the core nauty routines are compiled C rather than interpreted matlab code for the other two tests.

In the 3DLL method steps 1, 2 and 4 occur at the same time. Just as symmetries can cause the automorphic group to become large for the nauty-based methods, symmetries in configurations can cause ambiguities in finding an origin to start the matching process. However, in practice the center of mass heuristic works very well in reducing the candidate origins to one or two making this method very scalable. It should be noted that this method cannot recognize functionally identical configurations - that is, configurations with different module orientations but identical orientations of the axis of rotation for each module. This is because the port count heuristic is not invariant under functionally similar configurations. In the future this could be replaced with a different heuristic - *eg*, a count of the number of modules in each functional orientation rather than a count of the port connections.

While both the nauty-based method and the spectral decomposition method are very general and easily applied to any self-reconfiguring system, the 3DLL method is specific to CKBot and cube oriented modules.

In the spectral decomposition method steps 1 and 2 happen concurrently and quickly. For the 3rd and 4th steps symmetries in the configuration can lead to redundant eigenvector elements which require explicit disambiguation. This process can take a very long time.

Both the nauty-based and 3DLL approaches essentially precompute an approximate canonical form for the elements in the library. Nauty computes a canonical form with out the port information, while the 3DLL method generates a representation that is very likely to be unique exploiting the physical properties of configurations. This precomputation saves computation time with a nominal cost in memory, about half a megabyte for a 1000-module robot in the 3DLL representation and roughly the same for the nauty based method when stored as a sparse array. The spectral decomposition method does no precomputation, working directly with adjacency matrices. This allowed this method to be implemented on embedded controllers for a hardware demonstration.

A. Hardware demonstration

As a preliminary test, isomorphic gait control using graph spectra has been implemented on CKbot configurations containing up to 7 modules. A centralized controller containing 42 configurations with corresponding gaits ran a configuration detection scheme using communication architecture described in Section II.

The centralized controller compares the graph spectra of port adjacency matrices to see if a given structure is in the configuration library. If a match in the library was found, up to n! permutations of node labeling schemes are tried until the exact mapping between structures is found. Prohibitively slow speeds for structures containing 8 or more modules is a motivating factor for the work presented here. Configuration dependent gaits include those for snake, slinky, walking, rolling, lurching, and turning motions.

While the simplicity of graph spectra method allowed it to be implemented on the small embedded PIC controllers, converting the nauty and 3DLL based methods to run on small memory (e. g. 32K) embedded systems is future work.

VII. Conclusion

A comparison of the results for the three methods reveals the relative advantages and disadvantages of the three approaches. The spectral decomposition represents the most mathematically elegant of the three techniques but suffers from numerical issues. In addition, calculating the spectra, especially for larger number of modules is a computationally expensive process. However, this technique works very well in the recognition problem, i.e. it can be used to quickly identify the configuration in the database. It is however slower in the *mapping* problem.

The traditional graph isomorphism comparison method benefits from the ability to create a canonical configuration which pares the size of the search space for further matching. In addition, since it initially compares the underlying graph structure, it can match robot configurations whose port adjacency matrices may differ by symmetric rotations of the individual modules.

The 3DLL method using linked lists uses the extra information inherent in the port adjacency matrix but suffers from the need to run through every configuration in the library at runtime.

References

[1] G. S. Chirikjian, "Kinematics of a metamorphic robotic system," in *Proc. of the IEEE Intl. Conf. on Robotics and Automation*, 1996, pp. 1452–1457.

[2] M. Yim, D. Duff, and K. Roufas, "Polybot: a modular reconfigurable robot," in *Proc. of the IEEE Intl. Conf. on Robotics and Automation*, San Francisco, April 2000.

[3] A. Castano, W. M. Shen, and P. Will, "Conro: Towards deployable robots with inter-robots metamorphic capabilities," *Autonomous Robots*, vol. 8(3), pp. 309–324, 2000.

[4] T. Fukuda and Y. Kawauchi, "Cellular robotic system (cebot) as one of the realization ofself-organizing intelligent universal manipulator," in *Proc. IEEE Intl. Conf. on Robotics and Automation*, Cincinnati, OH, USA, May 1990, pp. 662–667.

[5] S. Murata, E. Yoshida, K. Tomita, H. Kurokawa, A. Kamimura, and S. Kokaji, "Hardware design of modular robotic system," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Takamatsu, Japan, October 2000.

[6] D. Rus and M. Vona, "Self-reconfiguration planning with compressible unit modules," in *Proc. of the IEEE Intl. Conf. on Robotics and Automation*, Detroit, 1999.

[7] C. Ünsal, H. Kýlýççöte, and P. K. Khosla, "I(ces)-cubes: A modular self-reconfigurable bipartite robotic system," in *Proc. of SPIE, Sensor Fusion and Decentralized Control in Robotic Systems II*, 1999, pp. 258–269.

[8] S. Murata, H. Kurokawa, and S. Kokaji, "Self-assembling machine," in *Proc. IEEE Int. Conf. on Robotics and Automation*, San Diego, California, May 1994.

[9] R. Fitch, D. Rus, and M. Vona, "Basis for self-repair robots using self-reconfiguring crystal modules," in *Proc. IEEE Int. Autonomous Systems 6*, 2000, pp. 903–910.

[10] I. M. Chen and J. Burdick, "Enumerating the non-isomorphic assembly configurations of a modular robotic system," *Intl. Journal of Robotics Research*, vol. 17(7), pp. 702–719, 1996.

[11] A. Castano and P. Will, "Representing and discovering the configuration of conro robots," in *Proc. of the IEEE Intl. Conf. on Robotics and Automation*, Seoul, Korea, May 2001, pp. 3503–09.

[12] D. Gomez-Ibanez, E. Stump, B. Grocholsky, V. Kumar, and C. J. Taylor, "The robotics bus: A local communications bus for robots," in *Proc. of the Society of Photo-Optical Instrumentation Engineers*, 2004.

[13] S. H. M. Yim and K. Roufas, "Climbing with snake-like robots," in *Proc. of the IFAC Workshopon Mobile Robot Technology*, Jejudo, Korea, May 2001.

[14] B. D. McKay, "Practical graph isomorphism," *Congressus Numerantium 30*, pp. 45–87, 1981, available at http://cs.anu.edu.au/~bdm/nauty/PGI/.

[15] F. R. K. Chung, *Spectral Graph Theory*. Providence: AMS, 1997.

[16] M. Zavlanos and G. Pappas, "A dynamical systems approach to weighted graph matching," in *The 45th IEEE Conference on Decision and Control*, San Diego, December 2006, (to appear).

[17] L. Hogben, "Spectral graph theory and the inverse eigenvalue problem of a graph," *International Linear Algebra Society*, vol. 14, pp. 12–31, 2005.

[18] D. A. Spielman, "Faster isomorphism testing of strongly regular graphs," *STOC 96: 28th Annual ACM Symposium on Theory of Computing*, pp. 576–584, 1996.