

Accomplishing High-Level Tasks with Modular Robots

Gangyuan Jing · Tarik Tosun · Mark Yim · Hadas Kress-Gazit

Received: date / Accepted: date

Abstract The advantage of modular self-reconfigurable robot systems is their flexibility, but this advantage can only be realized if appropriate configurations (shapes) and behaviors (controlling programs) can be selected for a given task. In this paper, we present an integrated system for addressing high-level tasks with modular robots, and demonstrate that it is capable of accomplishing challenging, multi-part tasks in hardware experiments. The system consists of four tightly integrated components: (1) A high-level mission planner, (2) A large design library spanning a wide set of functionality, (3) A design and simulation tool for populating the library with new configurations and behaviors, and (4) modular robot hardware. This paper build on earlier work by the authors [9], extending the original system to include environmentally adaptive parametric behaviors, which integrate motion planners and feedback controllers with the system.

1 Introduction

Modular self-reconfigurable robots (MSRR) are systems composed of repeated robot elements (called *modules*)

G. Jing
Cornell University
E-mail: gj56@cornell.edu

T. Tosun
University of Pennsylvania
E-mail: tarikt@seas.upenn.edu

M. Yim
University of Pennsylvania
E-mail: yim@seas.upenn.edu

H. Kress-Gazit
Cornell University
E-mail: hadaskg@cornell.edu



Fig. 1: Six configurations from the design library

that have the ability to connect together to form larger robotic structures. These systems distinguish themselves from traditional robot systems through their ability to *self-reconfigure*: changing the connective structure of the modules to assume different shapes that have different capabilities. Over the last three decades, many kinds of modular reconfigurable robots have been built [19], [13], [7],[14], and many different approaches have been introduced for controlling and programming them [21], [24], [35].

Robotics research is increasingly focused on deploying robots in real-world applications such as search and rescue. Operating in these scenarios entails handling an enormous amount of variability in task requirements and environment conditions. One approach to this problem is to build complex monolithic systems, such as large humanoids [12]. These systems can perform a wide range of actions, but are extremely complex. In a sense, their broad range of capability comes at the cost of having to solve each individual problem in a complicated way. For example, to pick up and move an object, a humanoid must balance on two legs while using a high degree of freedom (DOF) arm to manipulate the object.

In contrast, a robot that was purpose-built for that task could accomplish it with far fewer DOF, and would require less complicated control algorithms.

The strength of MSRR systems lies in their flexibility. In principal, self-reconfiguration will allow modular robots to transform into designs specifically tailored to the needs of each new task they encounter, allowing them to elegantly address a wide variety of tasks by re-configuring into a wide variety of simple solutions. However, this strategy poses an obvious challenge: given a task, is it possible to select an appropriate configuration (robot shape) and behavior (controlling program) to address it? This is an important unsolved problem in the field, and remains a significant barrier to the use of modular robots to solve real-world problems [33].

In this paper, we present a system capable of selecting appropriate modular robot configurations and behaviors to solve complex high-level tasks. Our system is library-driven: rather than attempting to generate new designs from scratch, users specify task requirements and a high-level controller retrieves designs satisfying the requirements from a library of existing designs. In addition to library management, the system integrates tools for low-level design creation, high-level mission planning, and physical modular robot hardware.

We leverage ideas from recent work on automatic controller synthesis with correctness guarantees from high-level task specification [2, 3, 10, 11, 20, 31]. These methods have proven effective for addressing high-level tasks with traditional robots, allowing users to specify task requirements at a high level using formal languages and then automatically synthesizing low-level robot controllers with performance guarantees. Applying these methods in the context of modular robotics introduces an additional layer of complexity due to the fact that the morphology of the robot is not fixed.

This paper builds upon our earlier work, presented in [9]. Specifically, we expand our system by introducing environmentally-adaptive parametric behaviors (EAP behaviors) that can leverage sophisticated motion planners and feedback controllers to continuously respond to environment conditions.

Through hardware experiments, we demonstrate that our system is capable of addressing challenging multi-part tasks. This paper presents the details of the system, discusses its strengths and weaknesses, and provides a roadmap forward to apply a similar system in a real-world setting. We hope others will be able to adopt our framework and utilize it to bring modular robots into real-world applications.

1.1 System Overview

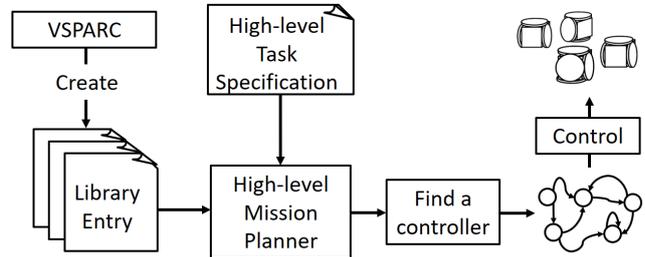


Fig. 2: System flowchart

Here, we provide a brief overview of the entire system. Figure 2 provides a visual companion to this section.

The system is built around a design library that spans a wide range of useful functionality. Library entries are configurations and behaviors for the SMORES-EP modular robot, which are designed in a physics-based simulator and design tool called VSPARC which we created for this purpose. Users build, program, and test modular robot designs through a graphical user interface, and can save their designs to a web server, allowing them to be shared with others. Any configuration or behavior created in the simulator can be directly ported to the hardware modular robot system, SMORES-EP.

Our system allows users to solve high-level tasks with modular robots. Tasks are specified in a mission planning tool using Structured English [6], a high-level language. Users do not specify which configurations and behaviors should be used to complete the task, but rather describe the required functionality. For example, the user might request that the robot perform a **drive** action in a **tunnel** environment labeled with the property `max_height = 3`.

To develop a solution to the task, the high-level mission planner fulfills each of the specified functionality by automatically selecting robot configurations and behaviors from the design library, generating a controller in the form of a finite-state automaton. In the above example, the system could select any configuration that is capable of executing a **drive** behavior while maintaining a maximum height of 3 modules or less. In a sense, the high-level planner treats the entire modular robot system as a single robot with a set of capabilities defined by the library. The mission planner can then execute the controller to complete the task, directly commanding hardware SMORES-EP robots based on environment information from sensors.

1.2 Contributions

This paper presents an integrated system capable of addressing high-level tasks with modular robots. The tasks it addresses are *reactive*: they require decision-making about what action to perform based on the sensed environment; *complex*: they include multiple sub-tasks with potentially very different requirements; and *high-level*: the task specification encodes the desired outcomes, and the system intelligently synthesizes a solution that results in those outcomes using the available configurations and behaviors from the library.

This system is one of the first to address these kinds of tasks with modular self-reconfigurable robots, which introduce an additional layer of complexity because they can assume many configurations. This represents a significant contribution to the field, because such systems will be necessary for modular robots to operate in realistic task scenarios. By providing this framework and demonstrating its success in the lab, we hope to lay the foundation for future modular robot systems to address tasks in the real world.

The system includes four tightly integrated components: (1) A high-level mission planner, (2) A large design library spanning a wide set of functionality, (3) A design and simulation tool for populating the library with new configurations and behaviors, and (4) modular robot hardware. Several of the subcomponents represent research contributions. Our novel design tool (VS-PARC) represents a novel contribution, as does our library of 52 configurations and 97 behaviors. We also introduce a minor theoretical contribution by checking the feasibility of robot behaviors prior to controller synthesis.

This paper builds on earlier work by the authors, presented in [9]. In Section 4.3.1, we present environmentally adaptive parametric behaviors, which are a major novel extension to the system presented in [9]. These behaviors allow sophisticated closed-loop behaviors to be developed, integrating motion planners and feedback control.

2 Related Work

Modular self-reconfigurable robots (MSRR) systems have the potential for great advantage over traditional robot systems in scenarios where flexibility is required: for example, search-and-rescue scenarios where the environment and task requirements may not be well-known *a priori*. Much of the existing research in MSRR systems has focused on establishing the the fundamental capabilities that differentiate these systems from traditional robots. Notably, MSRR systems have demon-

strated the ability to form a wide variety of physical morphologies capable of diverse modes of locomotion, suitable to a range of different terrains [32]. The ability to autonomously reconfigure has been demonstrated [34], and a number of reconfiguration planning algorithms have been developed [25].

Similarly, a great deal of work has been done to develop behaviors for MSRR. Many efforts focus on distributed control strategies, taking advantage of the distributed nature of MSRR hardware [30]. Distributed strategies include central pattern generators [23] and hormone-based control [21]. Genetic algorithms have been used to automatically generate both modular robot designs and behaviors [8].

It is clear that MSRR systems have demonstrated the ability to accomplish low-level tasks such as reconfiguration, locomotion, and manipulation. However, to truly live up to their promise of flexibility in real-world applications, systems must be developed that leverage these low-level capabilities to address complex, high-level, multi-part tasks.

While there is a robust body of research into addressing high-level tasks with traditional robots, little work has been done in this area with modular robots. High-level control of modular robots poses a unique challenge, because solving tasks involves selecting not only appropriate behaviors, but also appropriate configurations. This makes it all the more important to develop automated systems that can synthesize task-appropriate modular robot configurations and behaviors from high-level specifications.

In [4], Castro *et al.* introduce a high-level control framework for the CKBot modular robot. This framework lays the theoretical foundations for our high-level mission planner, one of the four major components of our system. We expand the framework into a larger system capable of addressing significantly more sophisticated tasks. In addition to the mission planner, we provide design and simulation tools for creating and testing modular robot configurations and behaviors, and a large library (52 designs, 97 behaviors, 19 properties) with designs capable of addressing a wide range of tasks. We expand the theoretical formalism introduced by Castro to include both behavior and environment properties, increasing the expressiveness of task specification, and introduce a performance improvement by grounding abstract action specifications in concrete configurations and behaviors prior to automata synthesis.

Tosun *et al.* [29] introduce a system that allows users to rapidly synthesize modular robot designs and behaviors by composition. The system includes a physics-based simulator and a hierarchically organized library of configurations and associated behaviors. The goal of

this work is to aid in the selection of modular robot configurations and behaviors appropriate to complex tasks, but it takes a very different approach than our automated system, instead providing tools for users to manually create new designs by combining library entries using series and parallel composition operations.

Outside the realm of modular robotics, systems have been developed that can synthesize rapidly manufacturable robot designs from high-level user specifications [17],[16], [22]. This work is similar to ours in the sense that high-level specifications from the user are interpreted to synthesize robot designs and behaviors from elements in a design library. The goal of these systems is to allow novice users to rapidly design and build functioning robots at low cost, using fabrication techniques such as 3D printing [17],[16] and origami folding [22]. Consequently, the scope of the tasks they address is very different from ours. In these systems, library entries are electromechanical components such as motors, motor drivers, and microcontrollers, and a high level task might be “Create a robot that can walk and turn.” In contrast, library entries in our system are whole robots with associated behaviors, and we address complex, multi-part tasks such as “Climb on top of the table, and move any debris you find into the trash bin.”

3 Background

In this section, we define modular robot systems and provide background on controller synthesis techniques.

3.1 Modular Robot Systems

Definition 1 (Module) A module is the fundamental unit of a modular robot system. Each module is a small robot that can receive and respond to commands, move, and connect to other modules. In this work, we consider only homogeneous modular robot systems, meaning that all modules in the system are identical.

We define a module as $m = (J, A)$. $J = \{J_1, \dots, J_d\}$ is the set of joints of the module with d degrees of freedom. $A = \{A_1, \dots, A_k\}$ is the set of *attachment points* where the module can connect to other similar modules. Each attachment point can only connect to one other module at a time. We denote the attachment point A_i of module m as $m.A_i$.

Definition 2 (Configuration) A configuration is a connected set of modules that acts together as a single robot. The smallest configuration is a single module. A configuration is denoted as $\mathcal{C} = (M, E)$, where $M = \{m_1, \dots, m_q\}$ is the set of connected modules that

form the configuration and E is the set of connections between modules, represented as pairs of attachment points. $(m_i.A_{a_1}, m_j.A_{a_2}) \in E$, where $m_i, m_j \in M$, and $m_i \neq m_j$. cannot have two disconnected sets of modules that move independently).

Definition 3 (Joint Command) Joint commands are used to control the joints of the modules. A command to a joint J_i is defined as $u_{J_i} = (\alpha, V, t)$, where $\alpha \in \{\text{Position, Velocity}\}$ is the type of command, $V \in \mathbb{R}$ is the value of the command, and $t \in \mathbb{R}$ is the time duration of the command. For example, $u_{J_i} = (\text{Position}, \frac{\pi}{2}, 2)$ commands joint J_i to hold the angle $\theta = \frac{\pi}{2}$ rad for 2 seconds. Similarly, $u_{J_i} = (\text{Velocity}, \pi, 3)$ will drive joint J_i with angular velocity of $\dot{\theta} = \pi \frac{\text{rad}}{\text{sec}}$ for 3 seconds. We assume there are low-level controllers (*e.g.* PID controllers) that can drive the corresponding joint to satisfy the command u_{J_i} .

Definition 4 (Behavior) For a configuration \mathcal{C} , we define a behavior $B_{\mathcal{C}} = \{b_1, \dots, b_n\}$ as a sequence of behavior states. Each behavior state is defined as $b_i = (U, T)$, where U is the set of joint commands for all joints of all modules in the configuration. The time duration T of each behavior state is equal to longest duration of its joint commands U , ensuring that behavior execution will move on to the next state only once all joint commands in the current behavior state have completed.

3.2 Controller Synthesis

In this work, we utilize existing work on controller synthesis [6, 11] to generate high-level controllers for modular robot systems. The process of controller synthesis consists of three main steps: (1) representing the robot and the environment using a discrete abstraction, (2) expressing desired robot tasks with a formal specification language, (3) searching for a control strategy that satisfies the given task specification, or determining that such a strategy does not exist.

Robot and Environment Abstraction: To represent the continuous environment state and robot actions as discrete models, we abstract the environment events and robot capabilities into sets of boolean variables. The value of each variable represents the sensed environment state or the current robot actions. For example, the environment variable **Cup** is **True** if and only if the robot is currently sensing a cup with its camera. Similarly, the robot variable **Push** is **True** if and only if the robot currently performing a pushing action.

Robot Task Specification: A wide range of robot tasks can be defined using a formal language called Linear Temporal Logic (LTL). In [6], authors introduce a tool called LTLMoP that allows users who are unfamiliar with LTL to specify robot tasks in a formal language called Structured English, which is closer to natural language. LTLMoP then automatically translates Structured English specifications into LTL formulas. The following is an example of a robot task specification written in the Structured English:

- visit **Classroom**
- if the robot senses **Student** then do **Greet**
- do **Pickup** if and only if the robot senses **Trash**

In these examples, **Classroom**, **Greet**, **Pickup** are robot action variables and **Student**, **Trash** are environment variables.

Controller Synthesis and Execution: Authors of [11] introduce a framework to automatically generate a high-level robot controller to satisfy a task specification, or decide such controller does not exist. The synthesized controller is a finite-state automaton, and specifies robot actions that satisfy the task. Each state in the controller is labeled with robot variables, and each transition is labeled with environment variables.

To accomplish a tasks, the synthesized controller is implemented continuously by mapping each robot variable to a low-level robot controller, and mapping each environment variable to a robot sensing function. With these mappings, the robot is able to detect the environment and perform desired actions to satisfy the task specification.

4 System

4.1 Modular Robot Hardware - SMORES-EP Robot

Our system is built around the SMORES-EP modular robot, but could be extended to other modular robot hardware systems. In this section we provide an overview of the capabilities of SMORES-EP.

Each SMORES-EP module (Figure 3) is the size of an 80mm cube. Four faces of the cube have magnetic connectors known as *EP-Faces* that allow them to connect to other modules. The EP-Face connector is an array of four electro-permanent magnets (EP magnets) embedded in a planar face, and provides fast, strong, energy-efficient connection between modules. Each EP magnet consist of an electromagnet coil wrapped around a core of two permanent magnet rods. Short pulses of current through the coil generate a magnetic field that re-polarizes one of the magnets in the core, allowing

the external force to be turned on (magnetically attractive) or off (neutral). Once polarized, the magnets will maintain either state indefinitely, so a pair of connected EP-Face sustains a connection strength of 88.4N without consuming any power. Each face requires 80ms and 2.5J of energy to switch states. The magnets are arrayed in a ring with south poles counterclockwise of north, making the connector hermaphroditic (any two faces can connect) and able to connect at 90° increments [27].

The modules are kinematically identical to their predecessor, the SMORES robot [5], and have four actuated joints. The *left* and *right* faces of the module are able to rotate continuously at a maximum rate of 90° per second and can be used as wheels, allowing individual modules to move by differential drive. These faces have thin rubber tires to enable driving on a variety of surfaces. The circular top face is also able to rotate continuously at a maximum rate of 30° per second, and is referred to as the *pan* joint. A central bending joint (referred to as the *tilt* joint) has a 180° range of motion, allowing the top face to bend forward or backward until it is perpendicular to the bottom face. Each of the joints is equipped with a custom potentiometer for position sensing [28], and modules perform feedback control for all joints at a rate of 20hz.

The motions that a SMORES-EP cluster can perform are limited by the strength of the motors and connectors. A pair of connected EP-Faces can sustain a maximum bending load of 1.8Nm, equivalent to supporting 3.1 modules in cantilevered horizontally against gravity [27]. If necessary, this limitation can be alleviated by installing connector plates that screw into the faces of two modules to rigidly connect them. In this case, four modules can be supported in cantilever before the motor torque limits are exceeded. When two modules are attached using connector plates, they lose the ability to disconnect and self-reconfigure.

The SMORES-EP system also includes passive cubes that can act as lightweight structural elements in SMORES-EP robot configurations. These plastic cubes have the same 80mm form factor as modules, and have an array of 8 permanent magnets on each face, allowing them to make a strong connection to modules.

Each module has its own 600mAh LiPo battery, microcontroller (STM32F303), and 802.11b WiFi module (TI CC3000), allowing it to move and operate independently or as part of a cluster. Battery life is typically about one hour, assuming typical usage of the motors and wireless communication. The EP-magnets require very little energy, and usage does not significantly affect battery life.

In this work, a cluster of modules is controlled by a central computer running a Python program that sends wireless messages (UDP packets) to the movement and magnets of each module. Wireless networking is provided by a standard off-the-shelf router, with a range of about 100 feet. Because individual SMORES-EP modules do not have any way of sensing their environment, localization is provided by AprilTag markers [18] mounted to modules and objects of interest, tracked by an overhead camera. The AprilTag tracker, high-level planner, and module control software run with a control loop time of about 4hz on a laptop a 2.4Ghz processor and 4Gb of RAM.

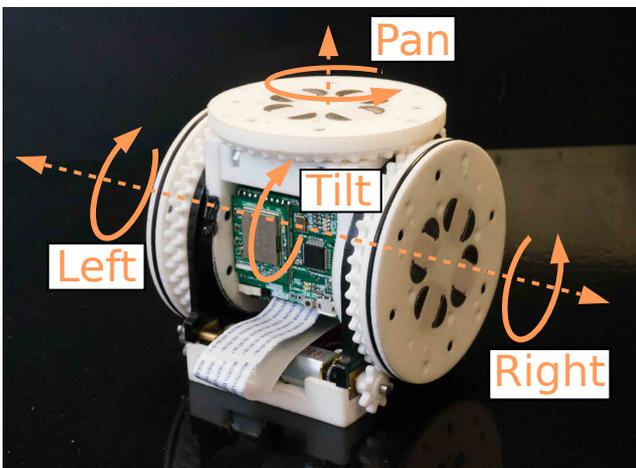


Fig. 3: SMORES-EP module

4.2 Design and Simulation Tool: VSPARC

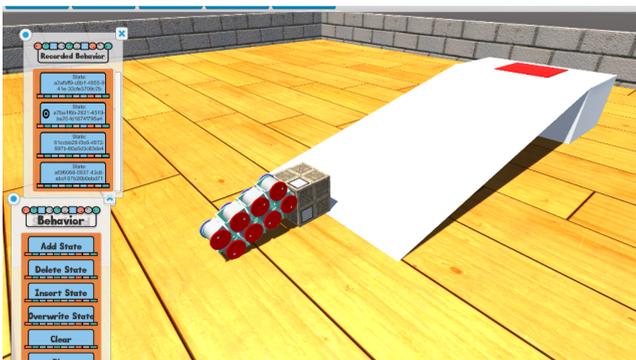


Fig. 4: VSPARC user interface

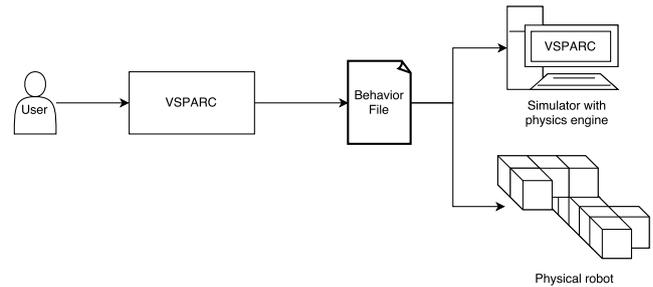


Fig. 5: The same behavior file can be used by both the simulator and the physical robot.

VSPARC, which stands for **V**erification, **S**imulation, **P**rogramming And **R**obot **C**onstruction, is our interactive design tool that allows users to design configurations and behaviors for SMORES-EP robots, and simulate them with a real-time physics engine. As shown in Figure 4, the graphical user interface, powered by the Unity3D Engine [1], allows users with little background in robotics to design and test different robot configurations and behaviors. The ability to control each joint of each module grants more experienced users the possibility to create complex designs.

VSPARC provides realistic physical modelling of SMORES-EP, taking into consideration factors such as the connector and actuator force limits. This allows users to test and verify behaviors before running them with physical modules and receive early warning if, for example, their behavior would likely cause the connection between two modules to break.

VSPARC is available for free online at www.vsparc.org, and enables users to save and share their designs to a central server, allowing a large number of users to contribute to our design library. VSPARC's main features are listed below:

- Design configurations with unlimited number of modules and visualize the design in a 3D environment.
- Command positions or velocities for each joint of all modules.
- Design behaviors for any configuration by creating a sequence of joint commands.
- Simulate the performance of any behavior in a physics engine.
- Create and share designs online. Test and improve other users' designs.

As shown in Figure 5, behaviors designed in VSPARC can be exported as XML files and then run on SMORES-EP modules, providing seamless translation of behaviors from the simulator to physical robots.

4.3 Design Library

In this section, we introduce a library-driven framework to organize configurations and behaviors created in VS-PARC. We introduce the notion of *properties*, which specify the functionality and constraints of behaviors, and the *robot design library*, which can be searched to find configurations and behaviors with desired properties.

Definition 5 (Property) Properties provide high-level descriptions of the intended effects of a behavior, as well as the environment in which the behavior is appropriate. We define a property as $p = (p_n, \Omega)$, where p_n is the name of the property (i.e. a description title, in English) and Ω is the set of values of the property. For example, a behavior with the property $p = (\text{Action}, \{\text{Move}, \text{Push}\})$ can perform both *Move* and *Push* actions. Properties are also used to describe the environmental conditions required for the behavior to run as expected. For example, the property $p = (\text{ObjectWeight}, [2, 5])$ indicates that the behavior can appropriately interact with an object if its weight is between 2 and 5 module-weights. In this case, the property is a quantitative description of the environment. We say a property $p_1 = (p_{n_1}, \Omega_1)$ *satisfies* a property $p_2 = (p_{n_2}, \Omega_2)$ if and only if $p_{n_1} = p_{n_2}$ and $\Omega_1 \subseteq \Omega_2$.

Properties connect tasks with behaviors that are appropriate to address them. In Section 4.4.1, we discuss how correct behaviors for a task can be automatically selected based on requirements over property values. Table 1 lists some examples of environment and behavior properties that might be used for common robot tasks.

Table 1: Examples of property names

Properties for Robot Behavior	Properties for Environment
Speed	Box_Mass
Width	Stair_Height
Height	Ground_Roughness
Action	Tunnel_Height

Definition 6 (Robot Design Library) The design library is a collection of modular robot configurations and behaviors labeled with environment and robot behavior properties. The library \mathcal{L} consists of a set of library entries, $\mathcal{L} = \{l_1, l_2, \dots\}$. Each library entry is defined as $l = (C, B_C, P_e, P_r)$, where C is the configuration and B_C is a behavior associated with C . P_e and P_r are sets of properties that describe the environment conditions and robot behavior functionality, respectively.

As an example, the library entry:

$$l = (C = \text{snake}, B_C = \text{climb}, P_e, P_r)$$

$$\text{where : } P_e = \{(\text{Ledge_Height}, [2, 3])\}$$

$$\text{and : } P_r = \{(\text{Action}, [\text{Climb}]), (\text{Speed}, [1])\}$$

represents a **snake** shape configuration with a **climb** behavior that can climb a ledge with a height of two to three module-lengths, with the speed of 1 module-length per second. Moreover, we say a library entry l *satisfies* a property p if there exist a property $p' \in P_e \cup P_r$ such that p' satisfies p .

To populate the library with different configurations and behaviors designs, we made our design tool available online at www.vsparc.org and distributed the tool to undergraduate and graduate student volunteers, hosting three hackathons in which participants created designs for various robot tasks. Currently, the library includes 52 configurations and 97 behaviors contributed by 20 volunteers. Since the full library is too large to list in this paper, we provide a representative sampling of configurations, behaviors, and properties in Table 2. The unit for length is the side length of a single SMORES-EP module. The unit for mass is the mass of a single SMORES-EP module.

4.3.1 Environmentally Adaptive Parametric Behaviors

In this section, we introduce an extension to the formalism we presented in [9]. As explained in Section 3, standard behaviors are defined as a series of joint angles or joint velocities for the modular robot cluster. These are discrete, open-loop actions can be sequenced by the high-level mission planner to complete tasks. Here, we present *environmentally-adaptive parametric behaviors* (EAP behaviors) which provide additional functionality, allowing low-level behaviors to directly respond to sensed conditions in sophisticated ways. These behaviors are *parametric* because they take input arguments, called parameters, which allow them to produce a continuous range of motions. They are *environmentally adaptive* because their parameters are intelligently assigned as a function of the state of the robot and environment.

Definition 7 (Environmentally Adaptive Parametric Behavior)

We define an Environmentally-Adaptive Parametric Behavior as $B_C^{EAP} = (\{b_1, b_2, \dots, b_n\}, \mathbf{p}, f)$, where $\{b_1, b_2, \dots, b_n\}$ is a sequence of behavior states, $\mathbf{p} \in \mathbb{R}^m$ is a vector of parameters, and $f : \mathbb{R}^n \rightarrow \mathbf{p}$ is the controller function.

Configuration Name	Single module	Rolling Loop	DoubleDriver	Stair Climber
Number of modules	1	8	7	4
				
Locomotion				
Max robot height	1	2.5	1.5	2
Max robot width	1	1	3	1
Max robot length	1	5	3	3.5
Terrain - Smooth	✓	✓	✓	✓
Terrain - Rough		✓		✓
Terrain - Sloped		✓	✓	✓
Driving - Straight	✓	✓	✓	✓
Driving - Differential drive	✓		✓	
Driving - Holonomic				
Ledge ascent - Max height		0.25		0.75
Ledge ascent - # modules lifted		all		all
Ledge descent - Max height		1		1.5
Ledge descent - # modules lowered		all		all
Manipulation				
Attachment - Push	✓	✓	✓	✓
Attachment - Magnetic	✓		✓	✓
Attachment - Carry				
Workspace size	$X : [-inf, inf]$ $Y : [-inf, inf]$ $Z : [0, 1]$	$X : [-inf, inf]$ $Y : [0, 1]$ $Z : [0, 2.5]$	$X : [-inf, inf]$ $Y : [-inf, inf]$ $Z : [0, 1.5]$	$X : [-inf, inf]$ $Y : [-inf, inf]$ $Z : [0, 2]$
Payload mass	1	2	4	2
Configuration Name				
	Swerve Lifter	backhoe	snake7	
Number of modules	9	9	7	
				
Locomotion				
Max robot height	2	4	4	
Max robot width	4	4	1	
Max robot length	3	7	7	
Terrain - Smooth	✓		✓	
Terrain - Rough				
Terrain - Sloped			✓	
Driving - Straight	✓		✓	
Driving - Differential drive	✓			
Driving - Holonomic				
Ledge ascent - Max height			3	
Ledge ascent - # modules lifted			4	
Ledge descent - Max height			3	
Ledge descent - # modules lowered			all	
Manipulation				
Attachment - Push	✓	✓	✓	
Attachment - Magnetic		✓	✓	
Attachment - Carry	✓			
Workspace size	$X : [-inf, inf]$ $Y : [0, 1]$ $Z : [0, 2]$	$X : [-inf, inf]$ $Y : [-3, 3]$ $Z : [0, 4]$	$X : [-3, 3]$ $Y : [-3, 3]$ $Z : [0, 4]$	
Payload mass	3	1	1	

Table 2: Matrix of designs and properties. Length and mass units are module-lengths and module-masses.

Like standard behaviors, EAP behaviors consist of a sequence of behavior states $\{b_1, b_2, \dots, b_n\}$. However, some of the joint commands of these states are *parametric*: instead of encoding fixed joint angles or velocities, they introduce a variable (called a *parameter* of the behavior) that can be assigned their value whenever the behavior is called. Additionally, we associate with each EAP behavior a *controller function* $f : \mathbb{R}^n \rightarrow \mathbf{p}$, which takes as input information about the robot and environment and produces as output the parameters of the behavior. This function is a feedback controller which lets the behavior adapt to environment conditions.

EAP behaviors expand the capabilities of our system. For example, consider a single SMORES-EP module, which can drive on smooth terrain using its two wheels. Using VSPARC, we can create a parametric **Drive** behavior that commands it to turn its wheels, assigning the wheel velocities to two parameters, e.g. $\mathbf{p} = \{V_{\text{left}}, V_{\text{right}}\}$. Using Python, we can now write a controller function for path following, taking as input the current location of the module and producing as output appropriate parameter values (wheel velocities) to drive the module along the path. In Section 5.2, we demonstrate how a similar **Drive** behavior and a path planner are used to direct a module to explore different regions on a tabletop.

As another example, consider the **Backhoe** configuration in Table 2. Using VSPARC, we can create a behavior that assigns parameters to the angles of all pan and tilt joints of the arm, providing access to the 7-DOF forward kinematics of the robot. For the controller function, we can write code that takes the position of an object as input and solves an inverse kinematics problem, providing output joint angles that cause the arm to touch an object.

As the above examples imply, EAP behaviors have a two-step design process. First, a parametric behavior is created using VSPARC, which we have extended to allow users to assign any joint value to a parameter rather than a fixed value. This process is no more difficult than creating a non-parametric behavior. Next, a controller function is written, to provide the mapping from sensor data to parameter values. Controller functions can be quite sophisticated (examples include motion planners and feedback controllers) and are typically written in Python by an expert user. However, if existing controller functions are available, novice users can re-use them to produce new EAP behaviors. For example, the path-following controller developed for a single module could be re-used by a novice user to create a similar behavior for the **DoubleDriver** configuration (Table 2), which is also capable of differential drive.

4.4 Reactive Controller Synthesis and Execution with the Library

In this section, we describe how our high-level mission planner synthesizes and executes controllers capable of accomplishing tasks using configurations and behaviors from the design library. This process has three parts: (1) matching library entries with boolean variables, (2) generating additional LTL constraints imposed by mapping, and (3) executing the controller. This framework is illustrated in Figure 6, and described in the following subsections.

4.4.1 Matching library entries with boolean variables

In our high-level mission planner, users specify tasks using robot and environment variables that abstract robot actions and environmental conditions, as discussed in Section 3.2. They label each variable with behavior and environment properties, to encode the desired functionality and constraints. Our system searches the design library for a set of library entries that satisfy the properties, and maps them to the corresponding boolean variable. Consider an example robot task specification:

if the robot senses **Cup** then do **Push**.

The robot variable **Push** might be described with:

$$P = \{ \{(\text{Cup_Mass}, [1, 3])\}, \\ \{(\text{Action}, [\text{Drive}]), (\text{Speed}, [1])\} \}$$

indicating that robot needs to be able to drive with speed of 1 with a cup that weights 1 to 3 module-weights. With this specification, we can search through the robot design library to find a set of library entries $L_y = \{l_1, \dots, l_k\}$ that satisfies all properties in the set P .

4.4.2 Generating additional LTL formulas imposed by matching

During the matching process, additional necessary LTL constraints are automatically created among the robot variables. Consider a set of robot boolean variables \mathcal{Y} used in a task specification. We define a mapping relation $\lambda : \mathcal{Y} \rightarrow 2^{\mathcal{L}}$ that maps each variable $y \in \mathcal{Y}$ to a set of library entries L_y that satisfies the user specified set of properties P for y . We say a library entry l can *implement* a variable y if $l \in \lambda(y)$. For any $y \in \mathcal{Y}$, if $\lambda(y) = \emptyset$, we need to make sure variable y is never **True**, because no library entry can implement y . For any $y, y' \in \mathcal{Y}$, if $\lambda(y) \cap \lambda(y') = \emptyset$, we need to make sure variable y and y' can never be **True** at the same time,

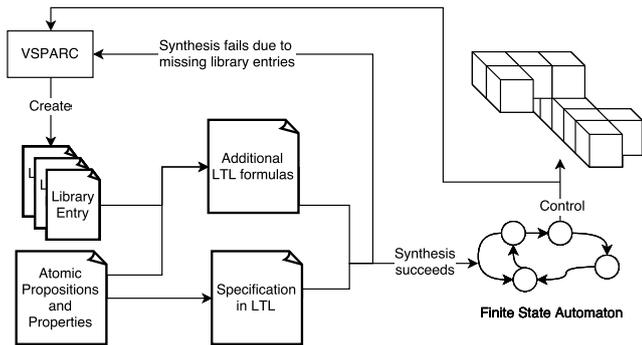


Fig. 6: Controller synthesis and execution

because there does not exist a library entry that can implement both y and y' . To encode the mutual exclusion between robot variables into the task specification, we specify them in the form of LTL formulas that are used together with the original task specification to generate robot controllers during synthesis.

4.4.3 Controller Execution

The synthesized finite-state automaton can be used to control simulated or the physical robots. If synthesis fails, possibly due to lack of library entries that implement some robot variables, LTLMoP will notify the user, who can then design suitable configurations and behaviors with VSPARC.

A synthesized controller is executed by running behaviors based on the value of each robot action variable. If a variable maps to a non-parametric behavior, the behavior is simply executed when the variable becomes **True**. A behavior is stopped when the corresponding variable becomes **False**.

To execute an environmentally-adaptive parametric behavior, the values of all parametric joint commands are decided during execution by calling the controller function each time the behavior is executed. For example, if the robot variable **Explore** matches with the **EAP Drive** behavior of the **Single Module** configuration, the behavior will be executed whenever **Explore** is **True**. A path planner function computes values of parameters in **Drive** in order to control the robot as a two-wheel differential-drive car.

If two consecutive behaviors must be satisfied by two different configurations, reconfiguration is required. To reduce overall mission time, when multiple behaviors match with a robot boolean variable, we avoid unnecessary reconfiguration by biasing towards the behavior that requires no reconfiguration.

5 Experimental Results

We validate the capabilities of our system through experiments in simulation and hardware, illustrated in Figures 9, 10, and 11, as well as the attached video. Faced with various task requirements, the system responds by synthesizing appropriate solutions. The simulation experiments demonstrate how the high-level mission planner can automatically synthesize and execute solutions to tasks using configurations and behaviors from the library. The hardware experiments validate that the system is capable of accomplishing complex physical tasks, such as carrying objects and climbing ledges.

5.1 Simulated Task Scenarios

We present two simulated task scenarios. A straightforward task is matched with a simple solution that uses one configuration, while a more complex task is addressed by reconfiguring between three different configurations, to leverage their wide-ranging capabilities.

5.1.1 Scenario 1

In Scenario 1, our system must solve a multi-part task in the environment shown at the top of Figure 7. The environment includes a button, a lightweight block, a gap in the ground, and a ramp, all in a straight line. Pressing the button causes the block to drop to the ground, where it can be pushed into the gap, forming a bridge between the flat region and ramp. When the task begins, the robot is initially positioned in front of the button. The objective is to reach a goal area at the top of the ramp. The high-level action definitions for this task are provided in Table 3.

After searching the library, the high-level mission planner discovers that the **rollingLoop** configuration has behaviors that satisfy the requirements of all three actions needed for this task (See Table 2). To complete the task, the mission planner synthesizes a controller that commands the loop to press the button, push the block into the gap, and ascend the ramp, as shown in Figure 7.

In response to this straightforward task, our system produces a simple solution. As discussed in Section 4.4.3, the system attempts to minimize reconfiguration when completing a task, and so will opt to solve the entire task with a single configuration whenever possible.

Action Definition	Properties
pushButton:	type = Manipulation_Push height = 1.5
pushBox:	type = Manipulation_Push payload = 2 distance_x = 3
climb:	type = Locomotion drive = Straight terrain = Sloped

Table 3: High-level Action Definitions for Scenario 1

5.1.2 Scenario 2

Like Scenario 1, Scenario 2 requires the robot to move from a starting position to a goal position. However, several small changes have been made to the environment that makes the task more difficult. The button has been moved to the side of the map, and floats at a height of 4 module-lengths above the ground. The box is twice as heavy, weighing 4 module-weights rather than 2. The ramp has been replaced with stairs with a step height of 0.75 module-lengths. Table 4 provides the high-level action definitions for this scenario.

These changes make it impossible for the `rollingLoop` to complete the task - it can't reach the button, it's not strong enough to push the block, and it can't ascend steps more than 0.25 module-lengths high. Instead, the high-level planner compiles a more complicated controller that uses behaviors from three different configurations in the library, shown in Figure 7. To push the button, the planner selects the `backhoe`, because it is the only configuration with a large enough vertical workspace. To push the block into the gap, the robot reconfigures into the `doubleDriver`, which is capable of driving, turning, and pushing objects as heavy as 5 module-weights. To climb the stairs, the robot reconfigures into the `stairClimber`, which can easily ascend 0.75 module-length steps.

This scenario demonstrates how our system leverages the flexibility of modular robots. This challenging task requires the diverse capabilities provided by all three configurations, and could not be accomplished by any one of them alone. Note that for the purposes of this work we do not provide strategies to autonomously perform self-reconfiguration. Instead, we assume that the robot can self-reconfigure between any two configurations as long as the initial configuration has an equal or greater number of modules than the final configuration. This does not fundamentally limit the power of our system: techniques for autonomous self-reconfiguration with SMORES-EP have been recently developed, and are being published in parallel.

Action Definition	Properties
pushButton:	type = Manipulation_Push height = 4
pushBox:	type = Manipulation_Push payload = 4 distance_x = 3
climb:	type = Locomotion drive = Straight ledge height = 0.75

Table 4: High-level Action Definitions for Scenario 2

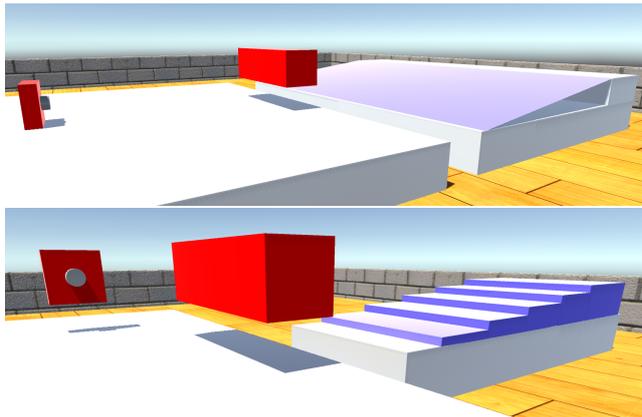


Fig. 7: Environments for Scenarios 1 (top) and 2 (bottom) in the simulator.

5.2 Hardware Experiments

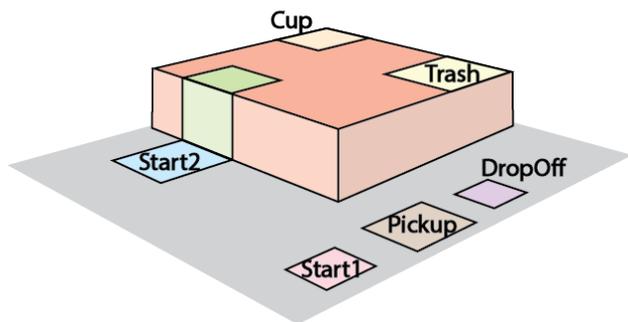


Fig. 8: Map of the hardware demo

Our hardware experiments demonstrate that our system can accomplish a complex physical task using physical SMORES-EP robot modules. The robot is required to clean the top of a table, operating in the environment shown in Figure 8. To do so, the robot must first move a waste bin from its initial location (labelled "Pickup") to a target location next to the table (la-

belled “Dropoff”). Then, the robot must climb to the top of the table and explore the surface. Whenever it encounters an object, it must react appropriately: if it is garbage, it should push it off the table and into the waste bin, and if not, it should notify a human to remove it.

This experiment showcases the seamless translation of behaviors from the VSPARC simulator to hardware, and the ability to use the LTLMoP high-level planner to create mission plans that can be directly executed by the modules. AprilTags tracked by an overhead camera provide information about the position of modules and objects in the environment, serving as sensory feedback for the high-level planner.

This experiment also demonstrates how the design library is continually expanded as users develop designs to address new tasks. While the library encompasses a wide range of functionality, it is by no means complete: when a high-level specification was first created for this experiment, the mission planner reported that it could not be satisfied using existing elements in the library. Consequently, two new configurations (the `swerveLifter` and `snake7` configurations) were created, and low-level behaviors were iteratively developed to fulfill the needs of each component of the task. Once these configurations and behaviors were made available in the library, the high-level planner was able to successfully synthesize and execute controllers to accomplish the tasks.

5.2.1 Moving the Waste Bin

The robot begins its task in region `Start1`, and must move the waste bin from `Pickup` to `Dropoff`, a distance of 10 module lengths. Once the waste bin is in place beside the table, the robot must travel to the edge of the table (`Start2`), where it can begin the next phase of the task (exploring the tabletop).

The waste bin is a box supported by four legs, making it impossible for any design less than two module-heights tall to push it. This constraint rules out most car-like configurations in the library. The 10-module distance over which the bin must be transported imposes a workspace requirement that rules out all stationary manipulators. Consequently, the `swerveLifter` configuration was designed to meet all the criteria. The `swerveLifter` uses four SMORES-EP modules as powered caster wheels, allowing omnidirectional movement (sometimes called *swerve drive*). It can also raise and lower, enabling it to lift and carry objects by driving underneath them.

The high-level description of this phase of the task is shown in Specification 1, and Figure 10 shows how

the robot completes it. The task is reactive: the robot waits until it senses the waste bin before beginning the `pickup` action (Line 3 of Specification 1). Once the waste bin appears (i.e. the AprilTag marking it comes the camera view), the robot lowers itself, drives beneath the waste bin, and carries it to the `Dropoff` region. It then moves back out from beneath the waste bin, and executes a series of omnidirectional driving behaviors to travel to the edge of the table.

Specification 1 Moving the Wastebin

1. `carry` is set on `pickup` and reset on false
 2. `dropped` is set on `drop` and reset on false
 3. do `pickup` if and only if you were sensing `wasteBin` and you are not activating `carry`
 4. do `goToTable` if and only if you are activating `dropped`
 5. do `drop` if and only if you were activating `carry` and you are not activating `dropped`
-

5.2.2 Table Exploration

With the waste bin in place, the robot begins the second phase of the task: cleaning the top of the table. The robot needs to climb to the tabletop, explore, and react to what it finds. The `snake7` configuration was designed to be capable to do this. As shown in Table 2, the `snake7` configuration can use its `climbup` and `climbdown` behaviors to ascend and descend ledges up to 3 module-heights tall. However, it is unable to lift its entire body up to the tabletop, and even if it could, it would be too large to effectively explore. Instead, the robot reconfigures, detaching the front module of the snake to act as a `module1` configuration that can use its EAP behavior `differentialDrive` to explore the tabletop, and its `spin`, and `push` behaviors to clean.

Specification 2 provides the high-level task description, and Figure 11 shows the robot completing the task. The robot begins in the `snake7` configuration, positioned at the edge of the table in the `ground` region. An AprilTag is fixed to the front module of the snake, allowing the mission planner to determine its location at all times. Sensing that it is in the `ground` region, the `snake7` executes `climbup` (line 8 of Specification 2). After climbing, the mission planner senses that the head of the snake has reached the `dock` region at the edge of the tabletop, and executes the `undock` behavior to detach the head module from the snake (line 6), allowing it to operate on its own as a `module1`.

The module then uses `differentialDrive` to visit two regions of interest on the tabletop (`loc1` and `loc2`). `differentialDrive` is an EAP behavior that allows

the robot to explore its environment in a continuous fashion. The driving behavior and its parameters are the same as the driving behavior presented as an example in Section 4.3.1: two parameters specify the left and right wheel velocities. The controller function is a potential field path planner that maps the robot’s current position (sensed via AprilTag) to a desired linear and angular velocity, which are converted to wheel velocities.

When it reaches **loc1**, the robot senses a coffee mug (marked with an AprilTag), and responds by executing a **spin** behavior to notify a nearby human that it should be removed (line 1). When it reaches **loc2**, it senses a piece of trash, and it correctly responds by performing a **push** to move it off the table and into the waste bin. Having fully explored the table, the module returns to the dock point and re-attaches to the body of the snake (line 5). The snake then executes **climbdown** to descend back to the floor, completing its mission.

Specification 2 Cleaning the Tabletop

1. if you are sensing **mug** then do **spin**
 2. if you are sensing **trash** then do **push**
 3. **loc1visited** is set on **loc1** and reset on false
 4. **loc2visited** is set on **loc2** and reset on false
 5. do **docking** if and only if you were in **dock** and you are activating (**loc1visited** and **loc2visited**)
 6. do **undock** if and only if you were in **dock** and you are not activating (**loc1visited** or **loc2visited**)
 7. do **climbdown** if and only if you were in **dock** and you activated (**loc1visited** and **loc2visited**)
 8. do **climbup** if and only if you were in **ground** and you are not activating (**loc1visited** or **loc2visited**)
 9. infinitely often do **docking**
-

5.2.3 Challenges

In general, the hardware experiment was successful, with the high-level planner successfully executing library behaviors to complete this task. While running the experiment, several notable challenges were encountered. During the first phase (moving the waste bin), achieving accurate steering with the **swerveLifter** proved difficult. The **swerveLifter** steers by aligning four caster wheels in the same direction, a process that is sensitive to encoder calibration errors across modules. Recently, more sophisticated calibration procedures for the SMORES-EP encoders have been developed, and encoder performance has been improved [28].

During the second phase of the experiment (exploring the tabletop), careful initial positioning was required for the open-loop **climbUp** behavior to succeed - in sev-

eral trials, the snake was started too close to the ledge, causing it to collide with the corner of the table and break. This problem could be alleviated by developing an EAP behavior allowing the robot to autonomously drive to the appropriate distance before beginning to climb.

In both phases of the experiment, limited magnetic connector strength between modules presented a significant challenge. The **swerveLifter** configuration had to be constructed with a passive cube in its center in order to perform its raising and lowering behaviors without breaking. During descent from the table, bending forces experienced at the center of the **snake7** configuration would sometimes cause connections between modules to break.

The limited strength of the magnetic connectors can be viewed as a trade-off for ease of reconfiguration. Connection and disconnection between the head and body of the snake takes very little time, and the forgiving area-of-acceptance of the connector [27] makes it possible to dock the head of the snake to the body even though the exact position of the body is not known (only the head module had an AprilTag). Autonomous docking succeeded about 25% of the time. This performance could be improved by applying more recently developed techniques for autonomous self-reconfiguration with SMORES-EP.

6 Discussion and Future Work

6.1 Simulator-to-hardware translation

Translation of behaviors from VSPARC to the hardware was largely successful, and the ability to prototype designs and behaviors in a simulator resulted in significant time savings over prototyping in hardware. Disparities between performance in the simulator and hardware tended to arise from real-world phenomena the simulator did not model accurately. For example, variability in magnetic connector strength (which differs from module to module [27]) sometimes resulted in connections breaking unexpectedly, and encoder calibration errors could cause behaviors requiring very precise position control to perform poorly.

Incorporation of on-board sensing will allow our system to operate autonomously in unknown environments. At the time of writing, a “brain module” has been developed that allows SMORES-EP clusters to carry an RGB-D camera and computer unit. In the future, VSPARC could be expanded to include simulated sensing capabilities, making it easier to develop closed-loop EAP behaviors.

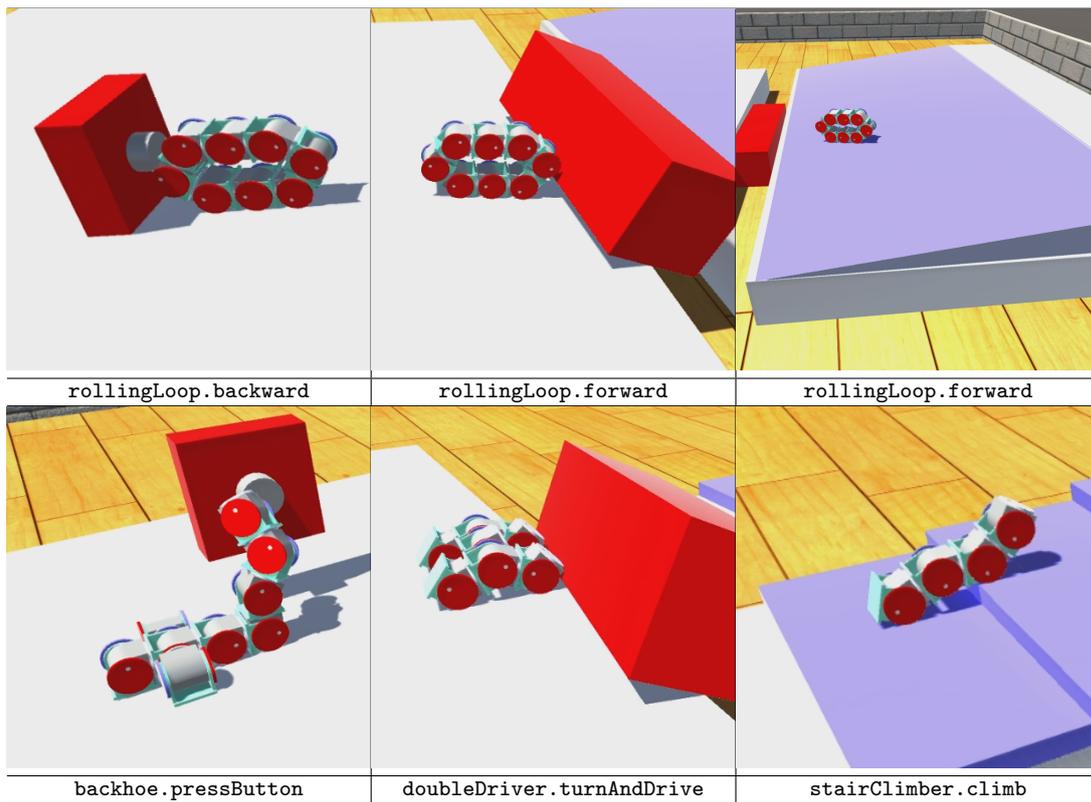


Fig. 9: Simulated Demo

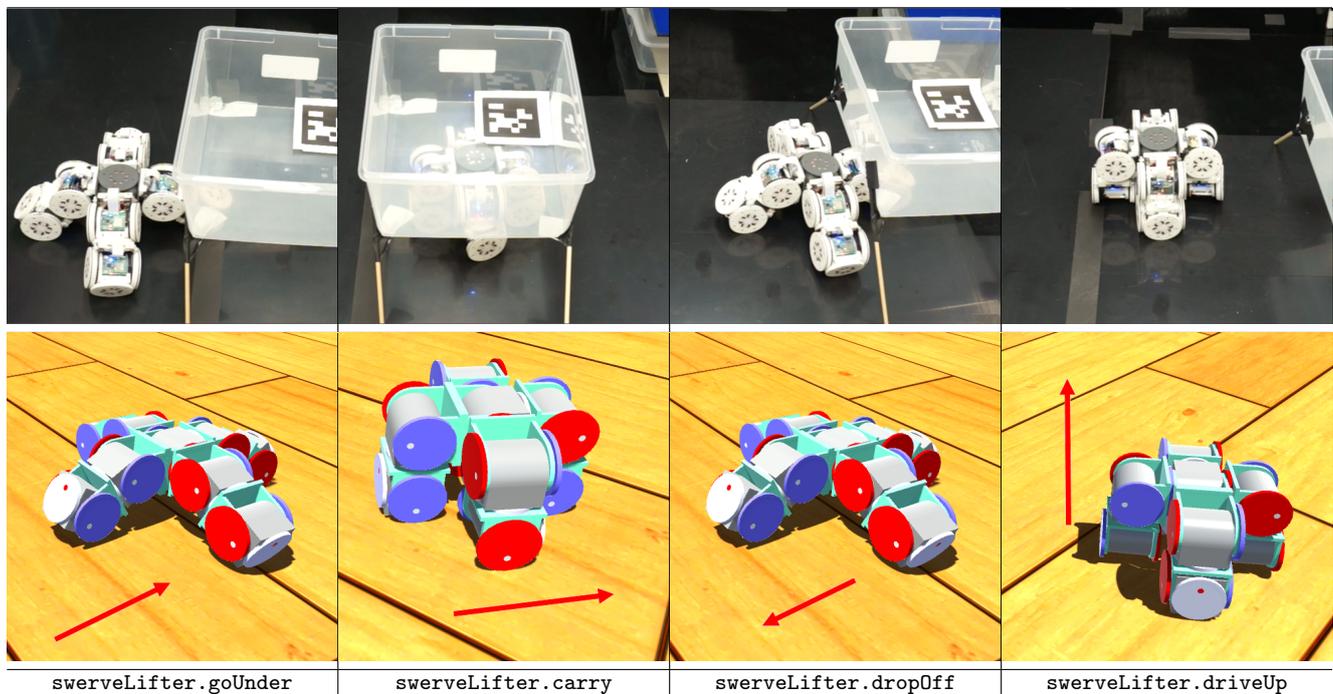


Fig. 10: Moving the Waste Bin

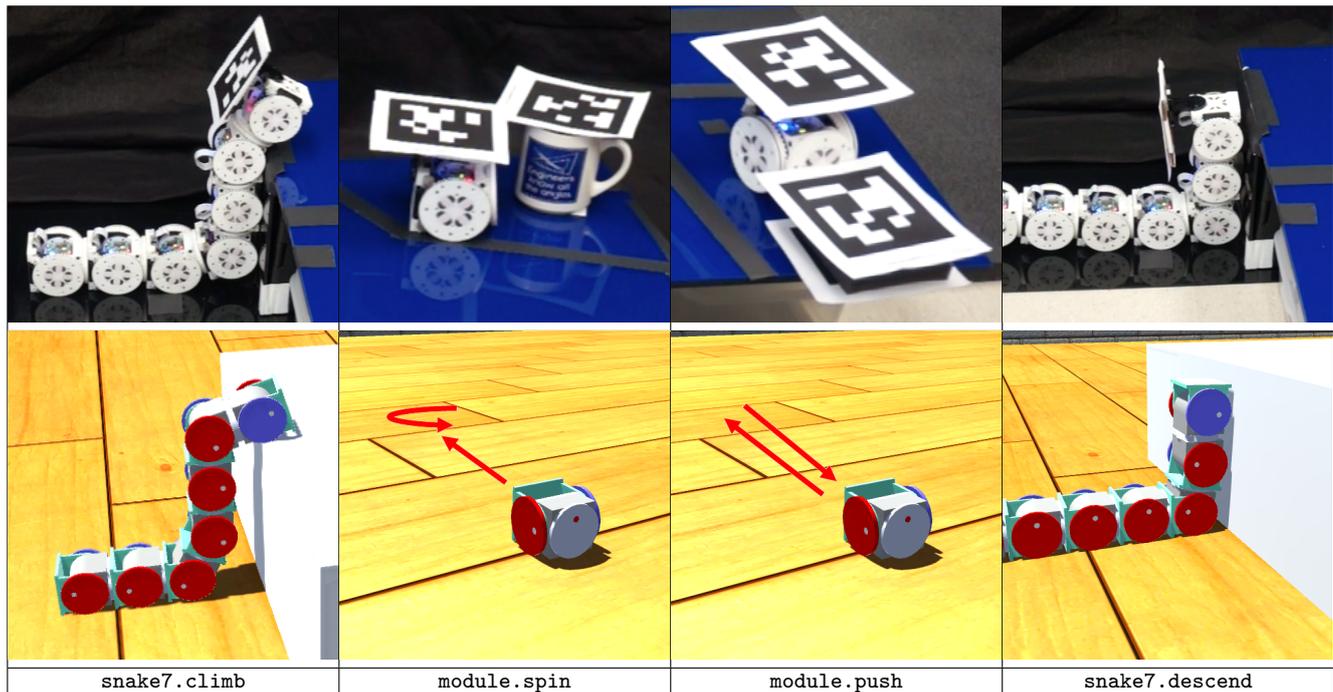


Fig. 11: Cleaning the Table

6.2 Library Creation: Lessons Learned

Early on in the development of this system, we intended to populate our design library through crowdsourcing, using a system such as Amazon Mechanical Turk where a large number of online users could create configurations and behaviors using VSPARC. We quickly realized that this strategy would not produce high-quality designs: developing sophisticated designs and behaviors in the simulator requires skill and experience. Holding hackathons with undergraduate engineers proved to be a much more effective strategy, because participants would become significantly more adept at creating designs and behaviors through hours of practice. Newcomers would typically spend about an hour creating a useful behavior, where well-practiced users would spend about twenty minutes.

Interestingly, users spent significantly more time creating behaviors than configurations. Most users required only a few minutes to build a new configuration and conceive of the fundamental motions they wanted it to perform. The majority of the design time was spent coding joint trajectories to achieve the desired motion while maintaining balance and avoiding connector strength overload. In the future, we hope to incorporate motion planning tools within VSPARC that will allow novice users to specify desired motions without explicitly cod-

ing joint angles. Evolutionary techniques will also be explored to generate behaviors automatically.

An existing algorithm for modular robot design embedding detection could also be used to automatically generate behaviors for new configurations [15]. This algorithm can automatically detect when one a subset of the joints of one configuration can be used to replicate the kinematics of another (a condition known as *embedding*), and generates a mapping that can be used to transfer behaviors originally developed for one configuration to any other configuration that embeds it. This could also allow behaviors developed for SMORES-EP to be ported to other modular robot systems, or vice-versa.

6.3 Composing Library Elements to Complete Missions

Environment and behavior properties provide an expressive way for the user to specify the requirements of a task. However, the fact that a behavior is labeled with a specific property does not guarantee it will perform as intended in all circumstances. Adapting behaviors to environments different from the one in which they were designed can cause them to fail, as evidenced by the problems in establishing proper initial robot position for the `climbUp` behavior in the table cleaning

scenario. Development of more closed-loop parametric behaviors will help address this issue.

Methods for automatically analyzing tasks and environments are actively being researched [26]. Determining optimal sets of environment factors and integrating methods for automatic task analysis and would be an interesting avenue for future work.

It's worth noting that some behaviors are much more tolerant to varying environments than others. In our hardware experiments with the `stairClimber` configuration, we found that a single open-loop gait was able to climb steps of several varying sizes with no problems. Establishing confidence bounds on behavior success as a function of environment parameters and including this information in the library is future work.

7 Conclusion

We presented a system for addressing high-level tasks with modular self-reconfigurable robots. We demonstrated how our physics-based simulator allows SMORES-EP configurations and behaviors to be easily created and stored in the design library, and how our a framework for labeling each entry in the library with descriptive properties allows them to be organized by functionality. Integration with a high-level mission planner allowed users to provide high-level task specifications, which were used to synthesize reactive controllers that use configurations and behaviors from the library. The capabilities of our system are validated through experiments in simulation and with physical modular robots. Building beyond our earlier work, we also expanded the system by introducing environmentally-adaptive parametric behaviors, which allowed sophisticated motion planners and feedback controllers to be used within our framework.

This system is among the first to address complex, reactive, high-level tasks with modular self-reconfigurable robots. By providing this framework and demonstrating its success in the lab, we hope to lay the foundation for future modular robot systems to address tasks in the real world.

Acknowledgments

This work was funded by NSF grant numbers CNS-1329620 and CNS-1329692.

References

1. Unity3d. <http://unity3d.com/>. Accessed: 2015-04-35

2. Belta, C., Bicchi, A., Egerstedt, M., Frazzoli, E., Klavins, E., Pappas, G.: Symbolic planning and control of robot motion [grand challenges of robotics]. *Robotics Automation Magazine, IEEE* **14**(1), 61–70 (2007)
3. Bhatia, A., Kaviraki, L.E., Vardi, M.Y.: Sampling-based motion planning with temporal goals. In: *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 2689–2696. IEEE (2010)
4. Castro, S., Koehler, S., Kress-Gazit, H.: High-level control of modular robots. In: *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pp. 3120–3125. IEEE (2011)
5. Davey, J., Kwok, N., Yim, M.: Emulating self-reconfigurable robots-design of the smores system. In: *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pp. 4464–4469. IEEE (2012)
6. Finucane, C., Jing, G., Kress-Gazit, H.: Ltllmop: Experimenting with language, temporal logic and robot control. In: *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pp. 1988–1993 (2010)
7. Fukuda, T., Kawauchi, Y.: Cellular robotic system (cebot) as one of the realization of self-organizing intelligent universal manipulator. In: *ICRA (1990)*
8. Hornby, G.S., Lipson, H., Pollack, J.B.: Generative representations for the automated design of modular physical robots. *ICRA (2003)*
9. Jing, G., Tosun, T., Yim, M., Kress-Gazit, H.: An end-to-end system for accomplishing tasks with modular robots. In: *Robotics and Systems (2016)*
10. Kloetzer, M., Belta, C.: A fully automated framework for control of linear systems from temporal logic specifications. *Automatic Control, IEEE Transactions on* **53**(1), 287–297 (2008)
11. Kress-Gazit, H., Fainekos, G.E., Pappas, G.J.: Temporal logic based reactive mission and motion planning. *IEEE Transactions on Robotics* **25**(6), 1370–1381 (2009)
12. Kuindersma, S., Deits, R., Fallon, M., Valenzuela, A., Dai, H., Permenter, F., Koolen, T., Marion, P., Tedrake, R.: Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot. *Autonomous Robots* **40**(3), 429–455 (2016)
13. Kurokawa, H., et al.: Distributed self-reconfiguration of m-tran iii modular robotic system. *IJRR (2008)*
14. Lipson, H., Pollack, J.B.: Towards continuously reconfigurable self-designing robotics. In: *ICRA (2000)*
15. Mantzouratos, Y., Tosun, T., Khanna, S., Yim, M.: On embeddability of modular robot designs. In: *IEEE International Conference on Robotics and Automation (2015)*
16. Mehta, A., Bezzo, N., An, B., Gebhard, P., Kumar, V., Lee, I., Rus, D.: A design environment for the rapid specification and fabrication of printable robots. In: *14th International Symposium on Experimental Robotics (ISER14), Marrakech/Essaouira, Morocco, June*, pp. 15–18 (2014)
17. Mehta, A.M., DelPreto, J., Shaya, B., Rus, D.: Cogeneration of mechanical, electrical, and software designs for printable robots from structural specifications. In: *Intelligent Robots and Systems (IROS 2014)*, pp. 2892–2897. IEEE (2014)
18. Olson, E.: Apriltag: A robust and flexible visual fiducial system. In: *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pp. 3400–3407. IEEE (2011)
19. Østergaard, E.H., Kassow, K., Beck, R., Lund, H.H.: Design of the atron lattice-based self-reconfigurable robot. *Autonomous Robots* **21**(2), 165–183 (2006)

20. Raman, V., Donzé, A., Sadigh, D., Murray, R.M., Seshia, S.A.: Reactive synthesis from signal temporal logic specifications. In: Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control, HSCC '15, pp. 239–248. New York, NY, USA (2015)
21. Salemi, B., Shen, W.M., Will, P.: Hormone-controlled metamorphic robots. In: ICRA (2001)
22. Schulz, A., Sung, C., Spielberg, A., Zhao, W., Cheng, Y., Mehta, A., Grinspun, E., Rus, D., Matusik, W.: Interactive robogami: data-driven design for 3d print and fold robots with ground locomotion. In: SIGGRAPH 2015: Studio, p. 1. ACM (2015)
23. Sproewitz, A., Moeckel, R., Maye, J., Ijspeert, A.J.: Learning to move in modular robots using central pattern generators and online optimization. *The International Journal of Robotics Research* **27**(3-4), 423–443 (2008)
24. Stoy, K., Shen, W.M., Will, P.M.: Using role-based control to produce locomotion in chain-type self-reconfigurable robots. *Mechatronics, IEEE/ASME Trans. on* (2002)
25. Sung, C., Bern, J., Romanishin, J., Rus, D.: Reconfiguration planning for pivoting cube modular robots. In: Robotics and Automation (ICRA), 2015 IEEE International Conference on, pp. 1933–1940. IEEE (2015)
26. Sung, J., Jin, S.H., Lenz, I., Saxena, A.: Robobarista: Learning to manipulate novel objects via deep multimodal embedding. arXiv preprint arXiv:1601.02705 (2016)
27. Tosun, T., Davey, J., Liu, C., Yim, M.: Design and characterization of the ep-face connector (2016). DOI 10.1109/IROS.2016.7759033
28. Tosun, T., Edgar, D., Liu, C., Tsabedze, T., Yim, M.: Paintpots: Low cost, accurate, highly customizable potentiometers for position sensing (2017)
29. Tosun, T., Jing, G., Kress-Gazit, H., Yim, M.: Computer-aided compositional design and verification for modular robots. In: International Symposium on Robotics Research (2015)
30. Walter, J.E., Tsai, E.M., Amato, N.M.: Choosing good paths for fast distributed reconfiguration of hexagonal metamorphic robots. In: ICRA (2002)
31. Wongpiromsarn, T., Topcu, U., Murray, R.M.: Receding horizon control for temporal logic specifications. In: Proceedings of the 13th ACM International Conference on Hybrid Systems: Computation and Control, HSCC '10, pp. 101–110. New York, NY, USA (2010)
32. Yim, M.: Locomotion with a unit-modular reconfigurable robot. Ph.D. thesis, Stanford University (1994)
33. Yim, M., Shen, W.M., Salemi, B., Rus, D., Moll, M., Lipson, H., Klavins, E., Chirikjian, G.S.: Modular self-reconfigurable robot systems [grand challenges of robotics]. *IEEE Robotics & Automation Magazine* **14**(1), 43–52 (2007)
34. Yim, M., Shirmohammadi, B., Sastra, J., Park, M., Dugan, M., Taylor, C.J.: Towards robotic self-reassembly after explosion. In: International Conference on Intelligent Robots and Systems, 2007, pp. 2767–2772. IEEE (2007)
35. Zhang, Y., et al.: Phase automata: a programming model of locomotion gaits for scalable chain-type modular robots. In: IROS (2003)