

Programming Modular Robots with Locally Distributed Predicates

Michael De Rosa, Padmanabhan Pillai,
Jason Campbell, Seth Copen Goldstein, Peter Lee

Abstract— We present a high-level language for programming modular robotic systems, based on *locally distributed predicates* (LDP), which are distributed conditions that hold for a connected subensemble of the robotic system. An LDP program is a collection of LDPs with associated actions which are triggered on any subensemble that matches the predicate. The result is a reactive programming language which efficiently and concisely supports ensemble-level programming. We demonstrate the utility of LDP by implementing three common, but diverse, modular robotic tasks.

I. INTRODUCTION

There are a significant number of challenges to programming modular robots. These challenges can be broadly be divided into two areas: managing the ensemble and controlling the individual modules. In this paper we present an approach to programming the ensemble based on *locally distributed predicates* (LDP). LDP lets a programmer specify how the entire ensemble should behave by breaking the problem down into how small groups of robots should interact. In this manner LDP significantly reduces the disparate problems of inter-robot timing and concurrency, resource management, and the lack of global knowledge at any individual robot.

Traditional imperative programming languages, such as C/C++, Java, do little to address the ensemble-level issues involved in programming modular robots. These languages are inherently oriented towards a single processing node, and require significant additional effort when used in a distributed setting. In addition to creating a representation of the data needed for an algorithm, the programmer must determine what information is available locally and what must be obtained from remote nodes, the messages and protocol used to transfer this data, mechanisms to route or propagate information through multiple hops as needed, and a means to ensure the consistency of this data. Furthermore, in algorithms to control ensembles, it is often necessary to express and test conditions that span multiple modules. Languages that constrain the programmer to the perspective of a single node make such algorithms difficult to implement.

Prior work in this field can be roughly divided into three categories: logical declarative languages for programming distributed systems, reactive programming techniques for robots, and functional approaches with roots in sensor network research. In the first category we have such tools as P2 [1] and Meld [2], which provide a logic programming facility for distributed systems (and modular robots in particular, in the case of Meld). These tools are powerful, in that programs written in them have certain provable properties, but this provability limits the expressive range

of the languages. Subsumption architectures [3] and reactive programming languages [4] provide other expressive formats for programming single robots, but are not specialized for ensemble-level multi-robot programming. Functional approaches include such languages as Regiment [5] and Proto [6]. These approaches also raise the abstraction level, but are less concerned with geometry (i.e., the neighbor relationships) and actuation.

We extend prior work on distributed watchpoints [7] to a reactive programming approach aimed at modular robots. The language, LDP, is used to specify lists of actions that are predicated on local distributed conditions. The language can express conditions that span multiple modules, can incorporate both temporal and spatial relations, and use various computational and logical operations. This mechanism breaks free from the node-centric paradigm enforced by C-like languages, and moves towards implementing algorithms from an ensemble perspective. Furthermore, LDP relieves the programmer from having to coordinate data distribution, as the system automates the distribution and coordination of all data needed to satisfy the conditions.

In the next section of the paper we give an overview of LDP. We follow this description with detailed examples of how LDP can be used to efficiently implement problems from three different domains. In Section III we demonstrate how LDP can implement a snake-gait in a chain-style modular robot. Section IV describes a typical data aggregation example by creating a spanning tree and then combining values from each robot's sensors. Finally, in Section V we demonstrate how LDP can be used in a metamodule system, implementing a complete planner in 9 lines of LDP. Overall, we show that LDP is effective for concisely expressing real-world modular robotic programs, can greatly reduce programmer effort in implementing complex distributed algorithms, and can do so efficiently.

II. LOCALLY DISTRIBUTED PREDICATES

Locally distributed predicates are useful for describing and detecting distributed state configurations in subsets of an ensemble. In contrast to classical global predicate evaluation [8], [9], which attempts to detect conditions over the entire ensemble, LDP operates on fixed-size, connected subgroups of modules. Thus, LDP is a more natural way to program an ensemble and can be used to detect conditions that occur in small groups of robots. Furthermore, LDP can detect such conditions efficiently, as searching for a small, connected subgroup is much less costly than inspecting the entire ensemble.

TABLE I
LDP OPERATORS AND PRIMITIVES

Boolean	&	!
Mathematical	+	- * / %
Comparison	>	< >= <= == !=
Temporal	prev()	next()
Topological	neighbors()	
Set	union	size any ...

A. LDP Syntax

An LDP program consists of data declarations and a series of statements, each of which has a predicate clause and a collection of action clauses. When a predicate matches on a particular sub-ensemble, the actions are carried out on that sub-ensemble. Notice that LDP has no explicit control structures, such as looping or function calls.

Each predicate begins with a declaration for each module involved in the statement. These modules are searched for in the order listed and, most importantly, there must be a path between all modules in a matching subensemble. The condition itself is composed of numeric state variables (expressed as `module.variableName`), temporal offsets (the operators `prev()` and `next()`), and topology restrictions (via the `neighbor` relation `neighbors(moduleA,moduleB)`). These primitives can be linked together with the mathematical, boolean, and comparison operators summarized in Table I.

The core language of LDP extends the condition grammar for distributed watchpoints with the addition of set variables (variables prefixed with a `$` are set variables, as in `moduleName.$setVar`) and the requisite operators for manipulating these variables. Specifically, we have implemented `intersection()`, `union()`, `size()`, `any()`, `add()`, and `remove()`.

B. Distributed Predicate Detection

The core of the LDP execution model is the *PatternMatcher*. A *PatternMatcher* contains all the information needed to determine if the predicate for its associated statement should be triggered. This object migrates around the sub-ensemble until either it fails to match or it matches. In addition to the active *PatternMatchers*, each robot has a collection of continually running threads (one for each statement in the program) which creates new *PatternMatchers* at every event of importance, e.g., a clock tick, a new sensor reading, etc. In its simplest form, a new *PatternMatcher* is created for each statement on each robot at every time tick.

The *PatternMatcher* is an object which encapsulates a search attempt for a particular predicate. Every *PatternMatcher* contains an expression tree, which encodes the boolean condition that the LDP is attempting to match. This expression tree contains storage for state variable values, to allow for comparison of state between multiple modules.

When a *PatternMatcher* is created the current module id is bound to the first named module in the statement and the local values of its state variables are filled in. The

expression tree is then examined for success or failure of the boolean predicate. If the expression tree is successful, then the action clauses of the statement are executed. If the tree is unsuccessful, the *PatternMatcher* is discarded. If no determination can be made, the *PatternMatcher* is forwarded to all of the module's neighbors, where the above process is repeated.

PatternMatchers provide numerous opportunities for optimization, allowing for boolean short-circuiting, as well as more intelligent search strategies than spreading to all neighbors. Additionally, *PatternMatchers* allow for backtracking in search paths, allowing for the detection of nonlinear configurations of matching modules. These extensions, as well as a full description of the distributed predicate detection algorithm, are presented in detail in [10].

C. Triggering Actions

By themselves, distributed watchpoints were insufficient to serve as a programming language, as they could not trigger any action on predicate matches. For LDP, we add a final clause to the predicate—the trigger. We define three types of triggers: (1) setting a state variable to a value, (2) changing the topology of the system, and (3) calling an arbitrary function implemented by the robot's runtime. Any predicate may have more than one trigger action. However, we require that all the actions must be executed on the same module. This eliminates the need for locking or synchronization across multiple actions and/or modules. Notice that the trigger actions may be executed on any module in the matching subensemble. Thus, we use the route gathered during the initial search for the subensemble to notify the acting module of the predicate's match.

D. Implementing LDP

Using LDP in any give modular robotic system is straightforward. The system must call an initialization step that sets up the LDP data structures, e.g., the *PatternMatcher* and state history queues.¹ The runtime requires the invocation of three basic routines which (1) enumerate a module's current neighbors, (2) transmit *PatternMatchers* between neighboring robots, and (3) invoke the statement threads at appropriate intervals, e.g., the `tick()` function. Finally, the system must ensure that incoming LDP messages trigger the appropriate callback.

Each application that uses LDP must additionally implement variable initialization, access, and modification for any state variables used in the program. The programmer must also implement any custom library functions that will be called from LDP actions.

III. EXAMPLE PROGRAM: PHASE AUTOMATA FOR SNAKE-STYLE GAIT

Phase automata [11] are a technique for scalably describing cyclic gaits in chain-style modular robots, such

¹More detail on implementing LDP can be found in [10]. Note to reviewers: a draft copy of this paper can be found at <http://www.cs.cmu.edu/~claytronics/LDP/Watchpoints-IJRR.pdf>




Fig. 1. a) Snake gait phase automaton, b) Joint angle vs. time graph, and c) Resulting snake gait in chain-style modules. Black modules are actuating negative joint angle, white modules are actuating a positive joint angle.

```
// per-module state variables
int id; // the id number of module
int parent = -1; // the id number of the previous module in the chain
float time; // the current time at the module
float offset = 0; // the module's phase offset
float angle = 0; // the joint angle of the module's central joint. Changing this value actuates the module's motor.

// set up parents and phase offset
1 forall (a) where (a.id == 1) do a.parent = a.id;
2 forall (a,b) where (a.prev(1).parent != a.parent) & (a.id < b.id)
   do b.parent = a.id, b.phase = a.phase + 0.1;
// perform movement
3 forall (a) where ((a.time + a.phase) % 1.0 == 0.5) & (a.parent != -1) do a.angle = 15.0;
4 forall (a) where ((a.time + a.phase) % 1.0 == 0.0) & (a.parent != -1) do a.angle = -15.0;
```

Fig. 2. Complete Source Code for Snake Gait Example, $\Delta\phi = 0.1, \alpha = 15.0, T = 1.0$

as Polypod [12] and Superbot [13]. A phase automaton consists of a set of multiple states with associated actions, whose transitions are governed either by external events or an internal globally-synchronized clock. A phase automaton additionally possesses an initial time offset ϕ , which can vary from module to module.

A simple phase automaton for a snake-like robot is shown in Figure 1a. In this automaton, the joint angle of a particular module is set to either $+\alpha$ or $-\alpha$ in a cyclic manner, with period T . The initial phase offset ϕ is determined by a module's position in the chain, and increases by a constant $\Delta\phi$ at each module. This automaton produces the gait shown in Fig. 1c.

To implement this automaton in a modular robotic system, there are two fundamental tasks: distributing the correct phase offset to each module, and setting the joint angle to the correct value based on the current time and offset. Figure 2 shows the complete code for these two steps. This example program assumes that the modules have unique id numbers that are ordered in increasing fashion from the “head” of the chain, which has id 1.

The first statement sets the `parent` of the head module to be its own id. The second statement triggers only when a module has changed its parent variable, and traverses one link of the chain at a time, setting the parent and phase variables of successive modules.

The next two lines implement the actual snake gait of the phase automaton. For each module, these statements check to see if the current time, modified by the phase ϕ and period T , corresponds to one of the transition points of the automaton. If so, the joint angle of the module is set appropriately. It is interesting to note that the first two statements are broadly applicable to any chain-style robot, and that they may be

reused for different gaits.

There are several interesting features of the phase automaton code which bear closer examination. As all of the statements in the program run simultaneously and concurrently, it is necessary to enable and disable the various steps of the algorithm by the use of *gating subpredicates*. These are predicates which match only once, or only for a certain period of time. There are two such subpredicates in Figure 2. The first, `a.prev(1).parent != a.parent` in line 2, ensures that the predicate matches only on the tick after the `parent` variable has changed. This prevents the continual (and unnecessary) reassignment of parents to the modules in the chain. The second subpredicate `a.parent != -1`, in lines 3 and 4, prevents the predicate from matching (and motion from occurring) until a parent and phase offset have been assigned to the module.

We evaluated the snake gait program on chains of 5 to 20 modules, and found that the resulting gait appeared visually similar to that presented in the original paper. The additional gaits (rolling and centipede) described in [11] could also be implemented using similar LDP programs.

IV. EXAMPLE PROGRAM: DATA AGGREGATION

A common task in modular robots, and distributed systems in general, is the aggregation of a distributed set of values at a central point. In this example program, we implement distributed averaging of a scalar variable over the entire ensemble. This is useful for such tasks as distributed sensing, localization, and center of mass estimation.

To obtain the average of a variable over all modules, we use a technique popular in sensor networks. We begin by designating one module as the root of a spanning tree, and having all modules transmit their value up the hierarchy of

```

int isSeed; // set to 1 on the spanning tree's root, 0 otherwise
int id; // the id number of module
int parent = -1; // the id number of the parent module in the tree
set<int> $notChildren = {}; // the set of module ids of neighbors who are not children
set<int> $children = {}; // the set of module ids of neighbors who are children, and have provided data
set<int> $neighbors; // the set of a module's neighbors' ids. Updated when accessed
int isLeaf = 0; // set to 1 if a module is a leaf
int sensor; // the variable to average over
int sum = 0; // sum of all sensor values received
int count = 0; // number of modules that have transmitted data
int average = 0; // average of the sensor value over all modules

// build spanning tree
1 forall (a) where (a.isSeed == 1) do a.parent = a.id;
2 forall (a,b) where (a.parent != -1) & (a.parent != a.prev(1).parent) & (b.parent == -1)
   do b.parent = a.id;
// build notChildren sets
3 forall (a,b) where (a.parent != -1) & (a.parent != a.prev(1).parent) & (a.parent != b.id)
   do b.$notChildren.add(a.id);
// start propagation at leaves
4 forall (a) where (size(a.$neighbors) == size(a.$children) + size(a.$notChildren))
   do a.isLeaf = 1;
// propagate data up
5 forall (a,b) where (a.isLeaf != a.prev(1).isLeaf) & (b.id == a.parent)
   do b.sum = b.sum + a.sum + a.sensor,
      b.count = a.count + b.count + 1,
      b.$children.add(a.id);
// get average
6 forall (a) where (a.count > 0) do a.average = a.sum / a.count;

```

Fig. 4. Complete Source Code for Data Aggregation Example




Fig. 3. Data Aggregation Algorithm: a) Available communications links b) Establishment of spanning tree rooted at R c) Leaves propagating data upwards (shaded circles) d-e) Additional levels of propagation f) Data aggregation complete

the tree to the root, where it is accumulated (Fig.3). The naïve implementation of such an algorithm would be for each module to transmit its variable's value, and for that value to be propagated all the way to the root of the tree, where the root module would add it to a running total.

Propagating each value independently is clearly inefficient, and so instead we implement summing and averaging at each level of the tree, so that only one data value must be passed up to a module's parent. The difficulty with this technique lies in knowing when all of a module's children have sent it information, so that the module can propagate the sum to a higher level of the tree. To solve this, we have each module maintain two set variables. One tracks immediate

neighbors that are known not to be its children. The other tracks immediate neighbors that are its children and have already provided it with data. When the size of these two sets sums to the total number of neighbors that a module has, it can transmit its own information up the tree. This algorithm is not the most efficient or robust choice [14], [15], but it serves to clearly illustrate how one might implement such a task.

The code in Figure 4 is an implementation of this averaging algorithm. It has 6 statements, spread over 5 different phases. The first two statements establish a spanning tree rooted at the designated module. The next statement adds all of module a's neighbors who are in the tree but not children of a to the set a.\$notChildren. The next statement begins propagation at each level by setting the isLeaf variable to 1 once all children have provided data. If a module becomes a leaf, the fifth statement adds its running count and total to that of its parent, and adds it to the parent's \$children set. Finally, the sixth statement continually sets the known average to be the total of all reported values divided by the count of reporting modules. The code as presented computes the average only once, but the addition of a "reset" mechanism is a simple change.

We evaluated the performance of the data aggregation algorithm on a simulated robot ensemble. The modules were arranged in a square lattice, in flat planes with sizes ranging from 5 by 5 (25 modules) to 20 by 20 (400 modules). In all cases, the number of messages required was exactly three




Fig. 5. Metamodule-based Shape Planner. Grey metamodules at top of structure are being created, while those at the bottom are generating deletion trees (red arrows) and being destroyed.

```

int isSeed;           // set to 1 on the metamodule which initiates the motion planner, 0 otherwise
int id;               // the id number of metamodule
int parent = -1;      // the id number of the parent metamodule in the deletion tree
int inside;           // set to 1 if the metamodule is inside the target shape, 0 otherwise
set<int> $notChildren = {}; // the set of metamodule ids of neighbors who are not children
set<int> $spaces;      // the set of adjacent free locations where an additional metamodule should be created,
                      // updated when accessed
set<int> $neighbors;   // the set of a metamodule's neighbors' ids. Updated when accessed

function create(int); // creates a new metamodule adjacent to the calling one at a given space
function destroy();   // destroys the current metamodule by dispersing it into adjoining metamodules

// propagate FINAL state from seed outward
1 forall (a) where (a.isSeed == 1) do a.state = FINAL;
2 forall (a,b) where (a.state == FINAL) & (b.inside == 1) do b.state = FINAL;
// create new metamodules at edges of start shape
3 forall (a) where (a.state == FINAL) & (size(a.$spaces) > 0) do a.create(a.$spaces.any());
// propagate PATH state to all metamodules outside start shape
4 forall (a,b) where (a.state == FINAL) & (b.inside == 0) &
  (b.state == NEUTRAL) do b.state = PATH;
5 forall (a,b) where (a.state == PATH) & (b.state == NEUTRAL) do b.state = PATH;
// build deletion trees from FINAL out through PATH metamodules
6 forall (a,b) where (b.state == PATH) & (b.parent == -1) &
  ((a.parent != -1) | (a.state == FINAL)) do b.parent = a.id;
// build notChildren sets
7 forall (a,b) where (a.parent != -1) & (a.parent != b.id); b.$notChildren.add(a.id);
8 forall (a,b) where (a.state == FINAL) do b.$notChildren.add(a.id);
// delete PATH metamodules with no children
9 forall (a) where (a.state == PATH) &
  (size(a.$neighbors) == size(intersect(a.$notChildren,a.$neighbors))) do a.destroy();

```

Fig. 6. Complete Source Code for the Basic Metamodule Planner Example

times the number of discrete communications links between the modules. Each phase of the algorithm (spanning tree construction, not-child set construction, and data aggregation) required that each adjacent pair of modules exchange one message. In terms of time complexity, the algorithm required time linear in the depth of the spanning tree to complete.

V. EXAMPLE PROGRAM: METAMODULE PLANNER

As our final example, we explore the problem of distributed shape planning for an ensemble of lattice-style modular robots. We use an extension of the shape change algorithm described in [16]. The algorithm produces a distributed asynchronous plan for a group of modules to transform from a feasible start state to a feasible goal state, while maintaining global connectivity throughout the execution of the plan. Furthermore, the algorithm provides provable guarantees of completeness: if there exists a globally connected path, it will be found. A film strip of the planner in action is shown in Fig.5.

A. The planning algorithm

The basic planner in [16] finds a sequence of rearrangements to go from a starting configuration to reach a target shape while maintaining global connectivity. The algorithm operates on metamodules, i.e., particular structures of modules, which are assumed to provide an abstraction where one metamodule can spawn a new metamodule in an adjoining empty spot, or absorb an adjacent metamodule to create an empty spot². Initially, the ensemble of metamodules is in the start shape. During execution, metamodules are created and destroyed to reach the target shape.

The plan starts with a seed metamodule, which is in the

²In the full version of this algorithm, the modules that form new metamodules and are absorbed from deleted ones are treated as resources, and are passed around by the metamodules to where they are needed. For simplicity and space considerations, in this paper we show the basic version of the algorithm that does not consider resources. An LDP implementation of the full algorithm is available at <http://www.cs.cmu.edu/~claytronics/LDP/mm-planner-full.txt>.

intersection of the start and goal shapes. The seed metamodule recruits every neighboring metamodule in the goal shape to enter the *final* state. It marks every neutral neighbor not in the goal shape as a candidate for removal, with state *path*. To avoid breaking global connectivity when removing metamodules, the planner uses path state metamodules to grow trees out from the portions of the ensemble to be preserved into the regions to be deleted, and only deletes metamodules at the leaves. Every path metamodule has a link from its parent, and as long as the link is not broken, the module will remain connected to the goal shape. Eventually, the path-trees will have no further space to expand, at which point, the leaves can be trimmed without loss of connectivity. In Fig.5, the start shape is indicated by the lighter colored metamodules, the goal shape by the darker colored ones, and trees are indicated by red arrows.

B. Implementation

The implementation of the motion planner runs at the *metamodule* level, on structured subgroups of modules. This allows for the creation and destruction of metamodules, as their constituent modules can be absorbed or provided by other nearby metamodules. To implement this application, we must port the LDP runtime to the metamodule level, which requires implementing communication and state variable storage across multiple modules. The code for the planning algorithm is shown in Fig.6. The state variables `isSeed`, `inside`, and `$spaces` are dependent on the start and goal shapes of the specific plan, and are initialized and managed by the low-level support code.

The first two lines spread the *final* state to the seed, and then to every contiguous metamodule which is inside the target shape. The third line causes new metamodules to be created at *final* locations with free spaces for creation. The fourth and fifth lines propagate the *path* state outwards from the edges of the *final* region to all metamodules that are outside the goal shape. The sixth line creates a forest of trees rooted at the edges of the *final* region, and spreading into the *path* deletion region. Statements seven and eight create `notChildren` sets, in a similar fashion to the data aggregation of Section IV. Finally, the ninth statement deletes *path* metamodules with no remaining children.

VI. DISCUSSION & CONCLUSIONS

We have demonstrated the utility of LDP by implementing three common classes of modular robotic algorithm. LDP provides a concise abstraction of distributed state, and helps to separate the actions of the algorithm from the support code necessary in traditional imperative languages.

As with any language, there are certain tasks which are more or less difficult to express in LDP. In particular, distributed state comparison and simple temporal relationships are quite naturally written in LDP. The lack of any explicit control structures or ordering makes certain other tasks more difficult. In particular, in order to impose an execution order of different “phases” on the program, we must use gating subpredicates to bound the times when certain statements

can be active. This is especially important for statements which we want executed exactly once on each module.

Finally, we note that the design of the LDP language and runtime is deliberately amenable to extension with new primitives and operators. The addition of new primitives (such as per-edge variables or subexpression quantification) would allow LDP to address more specialized application domains. As LDP is derived from the distributed watchpoint language, it is also possible to debug LDP programs in a distributed fashion within the language.

REFERENCES

- [1] B. Loo, T. Condie, J. Hellerstein, P. Maniatis, T. Roscoe, and I. Stoica, “Implementing declarative overlays,” in *Proceedings of ACM Symposium on Operating System Principles (SOSP)*, 2005.
- [2] M. Ashley-Rollman, S. Goldstein, P. Lee, T. Mowry, and P. Pillai, “Meld: A declarative approach to programming ensembles,” in *Proceedings of the IEEE International Conference on Robots and Systems IROS '07*, 2007.
- [3] R. Brooks, “A robust layered control system for a mobile robot,” *IEEE Journal Of Robotics And Automation*, RA-2, pp. 14–23, April 1986.
- [4] G. Berry, “The esterel v5 language primer,” Centre de Mathematiques Appliquees, Ecole des Mines and INRIA, Tech. Rep., 1999. [Online]. Available: <http://www-sop.inria.fr/meije/esterel/esterel-eng.html>
- [5] R. Newton, G. Morrisett, and M. Welsh, “The regiment macroprogramming system,” in *IPSN '07: Proceedings of the 6th international conference on Information processing in sensor networks*. New York, NY, USA: ACM Press, 2007, pp. 489–498.
- [6] J. Beal and J. Bachrach, “Infrastructure for engineered emergence on sensor/actuator networks,” *IEEE Intelligent Systems*, vol. 21, no. 2, pp. 10–19, 2006.
- [7] M. DeRosa, S. Goldstein, P. Lee, J. Campbell, and P. Pillai, “Distributed watchpoints: Debugging large multi-robot systems,” in *Proceedings of the IEEE International Conference on Robotics and Automation ICRA '07*, 2007.
- [8] C. M. Chase and V. K. Garg, “Detection of global predicates: Techniques and their limitations,” *Distributed Computing*, vol. 11, no. 4, pp. 191–201, 1998.
- [9] R. Cooper and K. Marzullo, “Consistent detection of global predicates,” in *Proceedings of the ACM/ONR Workshop on Parallel and Distributed Debugging*, published in *ACM SIGPLAN Notices*, vol. 26, 1991, pp. 167–174.
- [10] M. De Rosa, S. C. Goldstein, P. Lee, J. Campbell, and P. Pillai, “Distributed watchpoints: Debugging large modular robotic systems,” *International Journal of Robotics Research (special issue, to appear)*, 2007.
- [11] Z. Ying, M. Yim, C. Eldershaw, D. Duff, and K. Roufas, “Phase automata: a programming model of locomotion gaits for scalable chain-type modular robots,” in *Proceedings of 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems, (IROS 2003)*, 2003.
- [12] D. Duff, M. Yim, and K. Roufas, “Evolution of polybot: A modular reconfigurable robot,” in *Proc. of COE/Super-Mechano-Systems Workshop*, 2001.
- [13] B. Salemi, M. Moll, and W.-M. Shen, “SUPERBOT: A deployable, multi-functional, and modular self-reconfigurable robotic system,” in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems IROS '06*, 2006.
- [14] A. Manjhi, S. Nath, and P. B. Gibbons, “Tributaries and deltas: efficient and robust aggregation in sensor network streams,” in *SIGMOD '05: Proceedings of the 2005 ACM SIGMOD international conference on Management of data*. New York, NY, USA: ACM Press, 2005, pp. 287–298.
- [15] S. Nath, P. B. Gibbons, S. Seshan, and Z. R. Anderson, “Synopsis diffusion for robust aggregation in sensor networks,” in *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*. New York, NY, USA: ACM Press, 2004, pp. 250–262.
- [16] D. Dewey and S. Srinivasa, “A metamodule shape planner for modular robots,” in *Proceedings of the IEEE International Conference on Robotics and Automation ICRA '08 (in submission)*, 2008.