

Protoswarm: A Language for Programming Multi-Robot Systems Using the Amorphous Medium Abstraction

(Short Paper)

Jonathan Bachrach
MIT CSAIL
Cambridge, MA 02139 USA
jrb@csail.mit.edu

James McLurkin
MIT CSAIL
Cambridge, MA 02139 USA
jamesm@csail.mit.edu

Anthony Grue
Microsoft
Redmond, WA 98052 USA
tonyg@alum.mit.edu

ABSTRACT

Multi-robot systems are becoming increasingly prevalent, but programmability is a major barrier to their deployment. Present systems force programmers to think in terms of individual agents. Application code becomes entangled with details of coordination and robustness and often does not compose well or translate to other domains. We offer an alternate approach whereby the programmer controls a single virtual *spatial computer* which fills the environment space. The computations on this spatial computer are actually performed by a large number of locally-interacting individual agents. This abstracts the actual computational hardware behind the spatial computer interface, and allows the programmer to focus on a single model of global computation. We achieve this abstraction with two components: a language that embodies continuous space and time semantics and a runtime library that implements these semantics approximately. We demonstrate the efficacy of our approach with multi-agent algorithms in both simulation and on a group of 40 robots.

Categories and Subject Descriptors

D.1.3 [Software]: Programming Techniques—*Concurrent Programming*; D.3 [Software]: Programming Languages—*Data-flow languages, concurrent programming structures, Compilers, Optimization, Run-time environments*; I.2.9 [Computing Methodologies]: Artificial Intelligence—*Robotics, Autonomous vehicles*

General Terms

Algorithms, Design, Languages

Keywords

Programming Languages, Spatial Computing, and Swarm Robotics

1. INTRODUCTION

Multi-robot applications involve coordinating the movement of robots in space over time. However, programming

Cite as: Protoswarm: A Language for Programming Multi-Robot Systems Using the Amorphous Medium Abstraction (Short Paper), Jonathan Bachrach, James McLurkin and Anthony Grue, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 1175-1178.

Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

multi-robot applications requires the user to write software for individual robots and then imagine how these robots will interact to produce the final application. The mapping from robot actions to group actions is often complex and difficult to invert, making programming these systems challenging.

The dream is that by using a high level language to program a multi-robot application, we would empower a programmer to succinctly implement low-level repair and group functionalities and to quickly compose new programs out of these components. Programs would be developed with a more intentional macro perspective, and modularity would be strongly promoted.

In this paper, we present a new language, called Protoswarm, that is inspired by the continuous space-time model of Proto and extends this type of model to program swarms of robots. In particular, we present a virtual *spatial computer*, built out of a continuous *Amorphous Medium*, which fills the environment space. Programmers develop code for this medium without considering the details of the individual agents. The computations on this spatial computer are actually performed by a population of locally-interacting agents. The agents approximate the virtual computer presented to the user. We achieve this Amorphous Medium Abstraction [1] using two mechanisms: a language, called Protoswarm, which provides continuous space and time semantics, and a runtime library which approximates the semantics on the given hardware.

2. RELATED WORK

There are many domain-specific programming models for spatial computers, Swarm [9], TinyOS [4], and Paintable Computing [3], but they all involve programming the behavior of the devices, rather than the behavior of the aggregate. A notable exception is CMost, the operating system for the CM-5 [11], which allows operations on fields of devices, but assumes a fixed population of devices arranged in a grid.

In related swarm languages, programmers are similarly forced to program and manage individual robots. Mataric [7] introduced the notion of basis behaviors and group computing, but the basis behaviors are more challenging to combine than in Protoswarm. More recently, works by Klavins [5] and Kloetzer [6] have promoted the idea of high-level descriptors for swarm flocking, and the ability to compile out rules. The high-level is more akin to what we will show. However these systems mainly produce motion control laws on more capable robots – with GPS, and global clocks – where interactions are less critical in determining robot behavior. Furthermore, the languages are focussed on motion, and do not provide

very expressive means of distributed sensing and distributed state, that one might like to do with a robot swarm application.

In contrast, sensor networks have focussed almost exclusively on data collection. They have focussed on using well-known complete languages such as SQL or functional languages, which come with strong guarantees about what can be computed and many algorithmic tools to support the language implementation. For example, the Regiment [10] programming language operates on geometric regions of space, but is targeted towards sensor-network data-gathering and only distributes some operations across space.

Our goal is to combine these two points of view – both sensing and motion control are fundamental parts of programming robot swarms. We would like to take advantage of these types of languages in robot swarms, since many times mobile sensor networks are essentially robot swarms and vice versa. The Proto language is described in [2] and the amorphous medium abstraction was first proposed in [1].

3. MULTI-AGENT PROGRAMMING

In this section, we introduce the Protoswarm language, and build up facilities that support high-level modular multi-agent programming. In Protoswarm, the computational model is based on manifolds of space that execute code, called the Amorphous Medium. The medium has computational state and physical extent, both of which evolve over time. We assume that the medium is populated by an infinite number of agents, and each agent can only communicate with neighbors within a fixed distance. Programs for continuous regions are then run approximately on a discrete set of agents. Each agent runs identical code but their execution diverges due to differing local state and environment and interactions with neighbors.

Protoswarm programs are written as expressions over fields, where fields are mappings from manifolds to values. Expressions are executed repeatedly, producing streams of fields. Behaviors are produced from vector fields by points in space moving in the direction of the vectors.

Protoswarm is inspired by the programming model of Proto [2]. We treat the world as fields/streams and the computing constructs compute on these streams. Also we treat the system as a spatial computer, so all computing constructs work on neighborhoods without reference to exact neighbors. Unlike Proto however agents can move in space, and this is described by adding movement actuation. Consult [2] for more information on the Proto programming language on which the Protoswarm language is based. Now we describe some constructs one can build in this language to program at the group level.

In order to build basic behaviors, we assume we have a number of basic spatial and temporal functions that are useful in programming the Amorphous Medium developed in [1]. In particular, we assume functions to measure distance between any point and a given source region (`distance-to`), to elect a leader (`elect-leader`), to designate subregions (`dilate`), and to provide a global clock (`time`).

3.1 Building Basic Behaviors

In this section, we build upon our basic spatial and temporal functions to produce simple motion primitives for wandering, clustering, and dispersion. We move regions by defining a vector field over a region and using this vector field to




Figure 1: Cluster-to and contour-field vector field examples run on 500 simulated robots.

move the points in the region. For example, a random vector field is produced by:

```
(def brownian (s)
  (tup (rnd (- s) s) (rnd (- s) s)))
```

This produces a tuple at each point in the region that represents a random change in that point's current goal position. Note that this is simply a field; we have not produced any motion yet. We can move each point to its goal position with (`mov (brownian)`), which produces the desired behavior.

Regions can be clustered into a set of smaller regions by moving each point towards the average of the positions of all neighboring points:

```
(def cluster () (int-hood (nbr-vec)))
```

As this code executes, points near the boundary of the region move towards the center of the region. No attempt is made to keep the region coherent as it clusters. Eventually, the region will contract to a set of points.

Conversely, regions can be dispersed by creating virtual springs between points with a resting length of d . The following fragment:

```
(def disperse (d)
  (int-hood
   (* (- 1 (/ d (nbr-range))) (nbr-vec))))
```

moves points to minimize the spring energy between neighbors. This eventually results a uniform dispersion [12].

In order to perform clustering and dispersion directed towards (or away from) a region, we need a way to determine the direction to a region. We can interpret a field of scalars in a region as the z-values of a topographic terrain. We can then compute the gradient at any point in the region by finding directions of maximal increase in height:

```
(def grad (field)
  (int-hood
   (* (/ (- (nbr field) field) (nbr-range))
      (nbr-vec))))
```

The third line computes a vector towards each neighbor, with magnitude equal to the gradient of the `field` towards each neighbor. The `int-hood` operator integrates all the gradient vectors within a local region around each point, called a *neighborhood*. The neighborhood of a point is the circle of radius `comm-range` centered at that point. Essentially, this function computes the average gradient vector towards the source from all the points in the neighborhood.

Distance-based dispersion and clustering can be defined by moving towards or away from this gradient vector:

```
(def disperse-from (src)
  (grad (distance-to src)))

(def cluster-to (src)
  (* -1 (grad (distance-to src))))
```

The left picture of Figure 1 shows the vector field produced by `(cluster-to (is-light))`, where the vectors points towards lit regions.

Finally, we can follow a contour line in the topography of the field. Our approach is to create another field that has a stable limit cycle along the contour at a given level. We can generate this field by summing vectors pointed both towards and tangential to the desired topographic line:

```
(def contour-field (field level)
  (let* ((vec (grad field))
        (+ (* c (- level field) vec)
           (rotate pi/2 vec))))
```

where `c` is a feedback constant less than one. The following example, produces a vector field causing points to orbit at 0.5 meter around the lit region:

```
(contour-field (distance-to (is-light)) 0.5)
```

as shown in the right picture of Figure 1.

3.2 Behavioral Combinators

In order to construct more complicated behaviors, we need a method for behavioral composition. The first mechanism spatially composes behaviors. For example, we can create a behavior that makes agents disperse and remain somewhat stationary over certain areas, while wandering everywhere else. The code `(mov (cover-light))` creates a dispersal field in lit areas, and a brownian field in other regions:

```
(def cover-light ()
  (where (is-light) (disperse) (brownian)))
```

The second mechanism composes behaviors over time, sequencing behaviors according to events. For example, we can sequence dispersion for 2 seconds followed by wandering for 3 seconds:

```
(loop (while (wait 2) (disperse))
      (while (wait 3) (brownian)))
```

or wander until coming in contact with an object, then pushing it for 5 seconds:

```
(loop (while (not (is-near-object)) (brownian))
      (while (wait 5) (push-object)))
```

In general, we can sequence arbitrary behaviors by introducing the notion of finite streams, which are truncated by some event. The `while` function creates a stream of fields while the predicate is true. Finite streams are represented as a tuple of value fields and a boolean *active field* that is true when the stream is active. The `loop` function transitions from finite stream to stream based on the active field.

4. IMPLEMENTATION

The Protoswarm implementation addresses three separate challenges: (a) how to implement the primitives in a fault-tolerant manner in the face of agent movement, (b) how to translate swarm programs onto actual robots in an efficient




Figure 2: Distance-to data plotting actual versus computed distances on 28 robots.

and portable manner, and (c) how to support program development.

Neighborhood communications and localization are supported by a best-effort communication scheme. The most recent information on neighbor’s relative positions and shared variables are stored in a table [3]. Neighborhood operations then access the table, combining the most recent values into an approximate summary value. The virtual machine maintains the table by gathering shared values during each round of execution. These are then transmitted each round while receipt of packets proceeds in the background.

The simulator permits the running of much larger networks (over 10,000 agents), larger applications, flexible visualization, and friendlier code development and debugging. As in the robot port, only a small amount of platform specific code is necessary. The bulk of the simulator code facilitates visualization, code development, and debugging.

We implemented Protoswarm on a group of 40 autonomous mobile robots designed by iRobot. Each “SwarmBot” is autonomous and is equipped with bump sensors, light sensors, and an infra-red inter-robot communication and localization system [8]. The inter-robot localization system enables each robot to determine the positions of its neighbors relative to its own local coordinate system. The infra-red communication system is used to maintain the neighborhood table.

5. EXPERIMENTS

We tested `elect-leader`, `distance-to`, `cluster-to`, and `dilate` on the robots. Data was collected from the robots using a ceiling-mounted vision tracking system that recorded the positions of each robot over time. Telemetry from each robot was recorded to monitor each robot’s internal state.

The `distance-to` function measures the distance between any point in the medium space and a source region. We ran the following code: `(distance-to (elect-leader (id)))`, which elects a robot to be the source region, and then measures the distance to that robot from all other robots. Figure 2 compares the estimated distance to the source region to the actual distance. The distance estimate is accurate over the entire workspace. Because the paths for messages are constrained to only travel over the communication graph, the distance estimate will be an overestimate of the actual path. The longest path through the network was four communications hops.

The `dilate` function uses the `distance-to` function to defined a region around a source. We tested:

```
(dilate (elect-leader (id)) 0.8)
```

to produce a region of 0.8 meters around the leader. The picture in Figure 3 shows a snapshot of `dilate` running on robots. The graph shows the extent of the dilation region




Figure 3: Dilate results on robots. The lefthand picture shows a swarm of 40 robots running dilate. The righthand graph shows the actual dilated region versus the ideal region.




Figure 4: Cluster-to results on robots. The lefthand plot shows five paths towards the source the bottom left, and the righthand plot shows the efficiency of these paths, where the efficiency is the ration of shortest possible path length to actual path length.

around the leader robot as it is driven around using radio control. The black line shows the probability of a neighboring robot considering itself part of the dilation region. The transition point is shifted to smaller radii because `distance-to` is an overestimate of the actual distance. We suspect that the slope of the transition is caused by the voids in the network and the convergence speed of `distance-to` relative to the speed of the robot.

The `cluster-to` produced a vector field which is used to drive regions towards source regions. We tested the following code:

```
(mov (mux (elect-leader (- (id)))
          (cluster-to (elect-leader (id))))
      (tup 0 0))
```

which drives an “anti-leader” robot to a leader robot. The left plot in Figure 4 shows five paths from various starting positions. The right plot shows the path efficiency for each of these paths, where path efficiency is the ratio of shortest possible path length to actual path length. In twenty runs, the robot always converged to the source.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we introduce a continuous spatial computer abstraction to programming multi-agent behaviors. Our approach is built upon an Amorphous Medium Abstraction which frees the programmer from needing to consider individual robots. The Protoswarm language uses this abstraction to provide the user with a high-level programming model. We describe several core algorithms written in Protoswarm useful for constructing larger applications. We tested program fragments on 40 to 10000 in simulation and on a physical swarm of 40 robots. In all cases, the programs behaved as expected and the resulting behaviors were robust and scalable.

The power of our approach is that we can write scalable applications once and deploy them approximately on a number of multi-robot platforms with each platform incurring a certain approximation error. We think it is important to characterize this error, but at this time we are unable to make formal or statistical bounds on it or guarantees on correctness of high level programs. We have been working steadily on this challenge and hope to have results soon.

Although the demonstrated examples are limited, the programming model is a promising tool for multi-robot systems. In the future, we hope to expand the list of group level behaviors and applications and deploy the model on a wider range of multi-robot systems.

7. ACKNOWLEDGMENTS

The authors would like to thank all our sponsors. In particular, this work was funded under NSF grant CCF-0621897. We would like to acknowledge Jacob Beal’s large contribution to this work both in helping flush out the robotics approach and for his key role in the development of the Amorphous Medium Abstraction and the Proto language.

8. REFERENCES

- [1] J. Beal. Programming an amorphous computational medium. In *Unconventional Programming Paradigms International Workshop*, September 2004.
- [2] J. Beal and J. Bachrach. Infrastructure for engineered emergence in sensor/actuator networks. *IEEE Intelligent Systems*, pages 10–19, March/April 2006.
- [3] W. Butera. *Programming a Paintable Computer*. PhD thesis, MIT, 2002.
- [4] J. Hill, R. Szewczyk, A. Woo, D. Culler, S. Hollar, and K. Pister. System architecture directions for networked sensors. In *ASPLOS*, November 2000.
- [5] E. Klavins. A language for modeling and programming cooperative control systems. In *Proceedings of the International Conference on Robotics and Automation*, 2004.
- [6] M. Kloetzer and C. Belta. Hierarchical abstractions for robotic swarms. In *IEEE International Conference on Robotics and Automation*, 2006.
- [7] M. Mataric and M. Marjanovic. Synthesizing complex behaviors by composing simple primitives. In *Proceedings, Self Organization and Life, From Simple Rules to Global Complexity, European Conference on Artificial Life (ECAL-93)*, pages 698–707, May 1993.
- [8] J. McLurkin. Stupid robot tricks: A behavior-based distributed algorithm library for programming swarms of robots. Master’s thesis, MIT, 2004.
- [9] N. Minar, R. Burkhart, C. Langton, and M. Askenazi. The swarm simulation system, a toolkit for building multi-agent simulations. Technical Report Working Paper 96-06-042, Santa Fe Institute, 1996.
- [10] R. Newton and M. Welsh. Region streams: Functional macroprogramming for sensor networks. In *First International Workshop on Data Management for Sensor Networks (DMSN)*, Aug. 2004.
- [11] J. Palmer and J. G.L. Steele. Connection machine model cm-5 system overview. In *Fourth Symposium on the Frontiers of Massively Parallel Computation*, 1992.
- [12] W. Spears, D. Spears, J. Hamann, and R. Heil. Distributed, physics-based control of swarms of vehicles. *Autonomous Robots*, 17(2-3), August 2004.