Coordinating Massive Robot Swarms[*]

Bruce J. MacLennan

Department of Electrical Engineering & Computer Science
University of Tennessee, Knoxville, TN, USA
web.eecs.utk.edu/~mclennan

Abstract

This chapter addresses the problem of how to coordinate the behavior of very large numbers of micro-robots in order to assemble complex, hierarchically structured physical objects. The approach is patterned after morphogenetic processes during embryological development, in which masses of simple agents (cells) coordinate to produce complex three-dimensional structures. In order to ensure that the coordination mechanisms scale up to hundreds of thousands or millions of microrobots, the swarm is treated as a continuous mass using partial differential equations. The chapter presents algorithms and simulations for assembling segmented structures (artificial spines and legs) and for routing artificial neural fiber bundles.

Keywords: Artificial Morphogenesis, Microrobot, Morphogen, Morphogenesis, Neural Routing, Partial Differential Equation, Segmentation, Self-assembly, Substance

Goals

Although there has been considerable progress in the bulk assembly of nanostructured materials, many future applications of nanotechnology will require complex hierarchical systems, structured from the nanoscale up to the macroscale. Examples include future robots, computer systems, and peripheral devices. In some cases, technologies such as 3D printing will permit the fabrication of systems of moderate complexity. However, hierarchical systems that span the full range of scales, from nano to macro, will require self-assembly, at least at the small spatial scales.

The fabrication of biological-scale robots illustrates many of the problems: how can we assemble a brain-scale artificial nervous system, high-resolution sensors, effector systems with many degrees of freedom, etc.? Mammalian brains contain billions of neurons with trillions of interconnections, and it is plausible that artificial neural systems with similar capabilities will require comparable numbers of devices. Mammalian cortex is highly structured and functionally organized; how can we assemble comparable numbers of devices and interconnect them appropriately? For example, the human retina has perhaps 100 million receptors, which compress data into the approximately one million neurons of the optic nerve; we would like to be able to assemble sensors of similar complexity for future robots. Animals behave competently in the physical world by means of detailed proprioceptive, haptic, and other sensory information, which is used to control, in real time, a large number of muscle fibers to achieve fluent, finely controlled, and rapid movement. How can we assemble sensor and effector systems of comparable complexity?

We are investigating the use of swarms of microrobots to assemble such systems, but to do so we need techniques that will scale up to massive (biological) numbers. There is no specific goal number, of course, but we have at least hundreds of thousands in mind, and millions or billions may be required to assemble biological-scale robots. We cannot assume that coordination and communication strategies that work with dozens and hundreds — or even thousands — of microrobotic agents will scale up to biological numbers. As is explained in more detail later, we guarantee that our methods will scale by using the

---

same approach that biologists have used for describing the movement of massive numbers of cells: partial differential equations. In effect, we approximate massive numbers of agents by the continuum limit.

We do not know what sorts of microscopic agents will be used for the self-assembly of complex, hierarchical systems; possibilities include microrobots, nanobots, and genetically-engineered microorganisms (effectively organic microrobots). Since there are a number of possible technologies at various size scales, our goal is an abstract description, independent of the specifics of the agents. That is, we are developing abstract algorithms for self-organization that will produce the desired results so long as the agents have certain basic capabilities (which we are identifying; see **Assumed Capabilities** below).

Artificial Morphogenesis

Morphogenesis as a Model

We have an example of how hierarchical self-assembly can be accomplished in *embryological morphogenesis*, which coordinates billions or trillions of cells to assemble a complex, hierarchical body (Nüsslein-Volhard, 2008). Even a relatively simple animal has a large number of tissues, organs, vessels, nerves, etc. that are physically structured in a complex and functional organization. Moreover, multicellular organisms are hierarchically organized from the cellular (and indeed nano) level up to macroscopic level. Beginning from a single cell, the developing zygote begins to organize itself, establishing poles and layers, and the progressing organization governs future development, so that the microscopic agents (the cells) create the structure that governs their own future behavior. Cells migrate, following chemical gradients, and create forces and pressures that help to shape the tissues. Under the influence of structured signals, cells differentiate into functionally distinct tissues. Thus, the development of the embryo provides an inspiring example of how microscopic agents can coordinate their mutual behavior so as to self-organize into an immensely complicated structure. Our goal in *artificial morphogenesis* is to mimic these processes for the self-assembly of complex, hierarchical artificial systems by massive robot swarms.

Related Work

Artificial morphogenesis has some similarities to *amorphous computing* (Abelson, Allen, Coore, Hanson, Homsy, Knight Jr., Nagpal, Rauch, Sussman, & Weiss, 2000), especially in the earlier stages of morphogenesis, when the agents are in a relatively homogeneous mass. However, as morphogenesis proceeds, the agents arrange themselves into more organized structures, which leads to more structured signaling and interaction, which then leads to further organization. In other words, the goal of artificial morphogenesis is to transform an amorphous or homogeneous initial state into progressively more complex and specific structures. The physical arrangement of the microrobots and their emergent structure of communication and control reinforce each other in an ascending spiral.

Other researchers have recognized the value of morphogenesis as a model (Goldstein, Campbell & Mowry, 2005; Murata & Kurokawa, 2007; Nagpal, Kondacs, & Chang, 2003), but we believe they have not applied it systematically enough to scale up to biological numbers. Certainly, in any sort of bio-inspired computing, a crucial issue is how closely to mimic biological processes. This is also the crucial issue in modeling: What is the appropriate level for the model? For artificial morphogenesis, we believe that it is essential to adopt models that obviously apply to very large numbers of robotic agents, comparable to the number of cells in an embryo. This is also essential for the self-assembly of macroscopic structures organized from the nanoscale up.

There has been significant research modeling biological morphogenesis, but sometimes this is too specific to biological systems. For example, the COMPUCELL3D system models three-dimensional morphogenesis using a cellular Potts model to simulate changes in cell shape during morphogenesis (Cickovski, Huang, Chaturvedi, Glimm, Hentschel, Alber, Glazier, Newman, & Izaguirre, 2005). This is certainly an important issue in biological morphogenesis, but less so in artificial morphogenesis, in which

the robotic agents probably do not change shape or do so only in restricted ways. Therefore, we model morphogenesis at a higher, more abstract level, which is more tractable, analytically and computationally. We believe this is the optimal level to express biological-scale morphogenetic algorithms in a way that is relatively independent of specific implementation technology. Embryologists have found it useful for similar reasons (e.g., Forgacs & Newman, 2005; Meinhardt, 1982).

Morphogenetic Processes

Edelman (1988, p. 17) divides morphogenetic processes into (1) *driving forces* and (2) *regulatory mechanisms*. There are three driving forces: *cell proliferation*, *apoptosis* (programmed cell death), and *cell migration*. Cell proliferation is the mechanism by which embryos grow, but, in the absence of self-reproducing microrobots, it is likely to be less important in artificial morphogenesis. We can often achieve effects similar to cell proliferation either by providing an external supply of microrobots, which migrate to a growth site, or by having a fixed population of microrobots transport passive components to the growth site. One function of cell death in embryological development is to create cavities and passages in a tissue. In principle, we can implement the second driving force, apoptosis, in artificial morphogenesis by programming microrobots to disassemble themselves under appropriate circumstances, but it might be more practical to have them migrate away, or to sculpt cavities by having microrobots remove passive components. In artificial morphogenesis, as in embryogenesis, migration is often guided by chemical signals, for example, following the gradient of a *morphogen*. However, other controllable characteristics of the environment can also guide migration.

The other class of morphogenetic processes, the regulatory mechanisms, comprise *cell adhesion* and *cell differentiation*, both of which have direct analogs in artificial morphogenesis. Microrobots adhere to each other and to passive components in order to change their relative positions and to create permanent or temporary structures. Cells in embryos use adhesion molecules, and some microrobots will use molecular adhesion too, as well as electrostatic, magnetic, and mechanical adhesion (e.g., via latches). The second regulatory mechanism, cell differentiation, is the process by which cells assume different functions by enabling specific regulatory circuits to the exclusion of others. This is one of the central mechanisms of artificial morphogenesis, for different behavioral programs are enabled or disabled by variables within the microrobots. As will become apparent in our examples that follow, this permits identical microrobots to behave differently depending on the morphogenetic context in which they find themselves.

Biologists have identified about a dozen fundamental morphogenetic processes, which seem to be sufficient for biological development (Salazar-Ciudad, Jernvall, & Newman, 2003). Therefore they constitute an agenda for artificial morphogenesis (MacLennan, 2010, 2012b). The processes may be classified as (1) *cell autonomous mechanisms*, in which the cells do not interact, (2) *inductive mechanisms*, in which cells change state as a result of communication, and (3) *morphogenetic mechanisms*, which create patterns by rearranging cells without changing cell states.

The cell autonomous mechanisms are involved with cell division, and lead to daughter cells in different states or with different distributions of proteins and other substances. In the absence of self-reproduction, these mechanisms might not be used in robotic implementations of artificial morphogenesis, and they have a limited role even in natural morphogenesis (Salazar-Ciudad et al., 2003). The same effects can be accomplished by other means more suitable to artificial systems.

The inductive mechanisms are based on cell-to-cell communication, which can be implemented by means of diffusible chemicals, by membrane-bound chemicals, or by direct signaling through cell-to-cell junctions. All of these are potential inductive mechanisms in artificial morphogenesis as well. Inductive mechanisms are classified as *hierarchical*, in which signaling is primarily unidirectional, or *emergent*, in which cells signal each other reciprocally. The latter leads to more complex self-organization, as for example in reaction-diffusion systems (Meinhardt, 1982; Turing, 1952).

Morphogenetic processes, in the strict sense, rearrange cells without altering their state. One such process is *directed mitosis*, in which cells divide to yield daughter cells with specific orientations. In the absence of cell division, as is most likely the case in artificial morphogenesis, newly placed components will be oriented with respect to those already placed. Another important morphogenetic process is *differential growth*, which creates shape through the mechanical properties, such as the viscosity, elasticity, and cohesiveness, of cells and extracellular matrices (Salazar-Ciudad et al., 2003). *Apoptosis* — programmed cell death — is a third morphogenetic process, which can be used to create form by sculpting cavities. In artificial morphogenesis, microrobots might not "die," but they could migrate away or be disassembled, or microrobots might eliminate passive material.

Assumed Capabilities

To define the territory between biology and swarm robotics, we have a provisional list of assumed microrobot capabilities:

1. Components may be either active (microrobotic) or passive. Morphogenesis takes place by microrobots self-organizing themselves or organizing passive *material*. The extent to which the microrobots remain part of an assembled structure will depend on their cost and on whether the structure is intended to be able to reconfigure itself (i.e., *artificial metamorphosis*).

2. Microrobots have simple, local sensors (chemical, optical, electrical, mechanical, thermal, etc.). Typically these sensors are capable of responding to the intensity or concentration of a field or substance, or to its gradient.

3. Microrobots have simple effectors for local action (motion, shape, adhesion) and for signal generation (chemical, electrical, mechanical, etc.). Microrobots move primarily by simple mechanical interactions with their immediate surroundings (adhesion, traction, pressure, torque, etc.).

4. Long-range communication is accomplished in a variety of ways. One is direct robot-to-robot signaling, which allows a signal to propagate in a directed or undirected way through a sufficiently dense swarm of microrobots. Long-range signaling can also be accomplished by diffusion, which can be active or passive. Passive diffusion is a result of Brownian motion of microrobots or passive components (such as signaling molecules), and active diffusion is a result of microrobots moving under their own power, but randomly. Finally, long-range signaling can be implemented by means of more directed processes: microrobots can move from one place to another, and thereby represent information by their presence, or they can transport passive information-bearing objects.

5. The behavior of microrobots is governed by simple regulatory circuits, which need not be electrical. For example, biological microrobots may use genetic regulatory circuits. In general, regulation will be more like analog control than digital computation. Therefore, most often our programs for microrobot control will take the form of differential equations.

6. Microrobots might be self-reproducing or not, but, to ensure wider applicability, we focus on processes that can be accomplished without self-reproduction. For example, where biological morphogenesis uses cell proliferation, artificial morphogenesis might use microrobots to transport passive components to a growth area. (See the next section for more on this issue.)

7. Microrobots make use of ambient energy and/or distributed fuel. In general, they will not be able to store enough energy for prolonged operation, so they will make use of thermal, optical, chemical, electrical, mechanical, or other ambient and distributed energy sources. Metabolism may be used by organic microrobots (genetically engineered microorganisms).

Non-biological Implementations

Obviously some things must be done differently in non-biological systems from the way they are accomplished in embryos. For example, tissue growth and some aspects of biological morphogenesis are a result of cell proliferation, but in the likely absence of self-reproduction (cell division), artificial morphogenesis must accomplish them differently. One way is to have an external supply of microrobots migrate to the growth site, where they become a permanent part of the self-organizing system. Another is to have microrobots transport passive components to the growth site and to insert them in the tissue. Even with these alternative mechanisms, some biological morphogenetic processes will have to be adapted for artificial systems. For example, cell proliferation may take place in the interior of a cell mass, which is difficult to mimic when new components are provided externally. In these cases, we may have to arrange alternate ways of getting new components to the growth sites (e.g., through open passages).

Communication and Coordination Mechanisms

As previously remarked, the primary communication and coordination mechanisms are by means of contact, diffusion, and movement. Microrobots may affect and sense components (active or passive) in their immediate vicinity and thereby transfer information. For longer-distance "broadcasting" they may produce substances or disturbances in the medium, which disperse by diffusion or wave propagation. Most importantly, microrobots can move within the medium, thereby conveying both information and control by means of their presence or by means of the objects they move. In general, we do not assume any long-range communication or coordination mechanisms except what might be provided by macroscopic external fields (e.g., gravitational, electric, magnetic). Our approach to the coordination of microrobot swarms does not require establishment of any sort of global coordinate system or external patterning.

Morphogenetic Programming

Soft Matter

Embryological morphogenesis takes place in the regime of *soft matter*, or *viscoelastic material* (de Gennes, 1992). These are materials which stretch, bend, and fold when subjected to weaker forces, and flow viscously when subjected to stronger ones, which is the way tissues behave in the developing embryo. These viscoelastic properties are fundamental to the creation of embryonic form and structure. Similarly, application of these processes in artificial morphogenesis, involving very large numbers of microrobots and passive materials connected more or less strongly, suggests that these self-assembling structures be treated as soft matter ("tissues").

Continuum Mechanics Framework

The appropriate mathematical framework for describing viscoelastic materials is continuum mechanics expressed in partial differential equations (Beysens, Forgacs, & Glazier, 2000; Forgacs & Newman, 2005; Meinhardt, 1982; Taber, 2004) An artificial morphogenetic system is described as one or more three-dimensional *bodies* (which might be fluids) continuously evolving in three-dimensional space as a result of external and internal forces. Although the bodies may be composed of discrete elements (e.g., molecules, cells, or microrobots), we adopt a level of abstraction that permits bodies to be treated as continua. The dynamics of the system can be described from the perspective of an external observer (an *Eulerian* or *spatial* frame of reference) or from the perspective of an infinitesimal volume element of the material (a *Lagrangian* or *material* frame of reference). The latter is a more agent-oriented perspective, suitable for programming microrobots, which biologists have found valuable as well (Bonabeau, 1997). (Mathematical details can be found in prior reports: MacLennan, 2010, 2011.)

Continuum Programming Notation

We have developed an experimental programming notation for expressing morphogenetic algorithms at a level suitable for control of massive swarms of microscopic robots and for modeling their behavior (MacLennan, 2010, 2011, 2012b). Primarily, it is a notation for programming in stochastic partial differential equations, with suitable means for definition and initialization of variables. Our goal is that it be directly executable as a simulation language, but that it also serve as a specification language for designing microrobots and passive materials that can be applied to artificial morphogenesis.

Change Equations

Behavior is described by *change equations*, which can express continuous, discrete, and qualitative change. The notation $ÐX = F(X,Y,…)$ can be interpreted either as a PDE (partial differential equation), $\partial X / \partial t = F(X,Y,…)$, or as a temporal finite difference equation, $\Delta X / \Delta t = F(X,Y,…)$. This systematic ambiguity, which permits realization as either a discrete- or continuous-time system, is respected by the formal rules of manipulation for the notation. More precisely, change equations are treated as *dynamic equations on time scales* (Bohner & Peterson, 2001; Agarwal, Bohner, O'Regan & Peterson, 2002), in which in our case the time scale is either $\mathbb{R}$, the real numbers, or $(\Delta t)\mathbb{Z} = \{…,-2\Delta t,-\Delta t,0,\Delta t,2\Delta t,…\}$. Space can be similarly treated as ambiguously discrete or continuous by defining partial differential equations over time scales (Ahlbrandt & Morian, 2002), but treating it continuously helps to ensure that our coordination strategies scale to truly massive robot swarms.

Sometimes it is convenient to break a long change equation into parts, and so we allow it to be expressed in several *partial equations*, which might be textually separated in a program. For example, an equation of the form $ÐX = F(X,Y,…) + G(X,Y,…) - H(X,Y,…)$ could be broken into partial equations:

$$ÐX \mathrel{+}= F(X,Y,…),$$
$$ÐX \mathrel{+}= G(X,Y,…),$$
$$ÐX \mathrel{-}= H(X,Y,…).$$

This notation was first applied to morphogenesis by Kurt Fleischer (1995, p. 20).

Due to many stochastic factors, morphogenetic processes, both natural and artificial, must be robust. Therefore, in many cases, a precise functional dependence is not so important as whether one quantity tends to increase or decrease another. Embryologists express these relationships in *influence diagrams*, which show how one quantity promotes or represses another. Our morphogenetic programming notation has a similar concept, a change *regulation* (as opposed to a change *equation*). For example, $ÐX \sim X,-Y,Z$ means that the change of X is enhanced (positively regulated) by X and Z, but repressed (negatively regulated) by Y. More precisely, this regulation is interpreted as an equation $ÐX = F(X,-Y,Z)$ in which F is an unspecified function that is monotonically non-decreasing in each of its arguments. This allows us to specify algorithms that omit details that are irrelevant or that will be determined in a specific implementation.

For convenience we use a conditional notation to allow a threshold to gate the influence of a quantity; it is defined:

$$[x > \vartheta] = \begin{cases} 1, & \text{if } x > \vartheta, \\ 0, & \text{otherwise.} \end{cases}$$

More complex conditions have the obvious meaning; for example $[x > \vartheta \wedge y > \varphi] = [x > \vartheta] \times [y > \varphi]$. Conditions are one of the principal tools for enabling or disabling particular behavioral equations, which fulfills purposes similar to cell differentiation in biological morphogenesis.

Substances

Analogous to classes in object-oriented languages, we have *substances* with properties. For example, the following **substance** declaration defines "morphogen" to be a substance with several fixed properties (diffusion rate, decay rate, etc.) and a concentration that can vary over a region of space:

> **substance** morphogen:
>   **scalars**:
> D_a ‖ diffusion rate
> κ_a ‖ production rate
> τ_a ‖ decay time constant
> ϑ_A ‖ threshold
>   **scalar field** a ‖ concentration
>   **behavior**:
> $Ða = [A > \vartheta_A]\kappa_a S(1-a) + D_a \nabla^2 a - a/\tau_a$

The **behavior** part defines the dynamics of the substance's properties by means of formulas, primarily change equations and regulations.

Like classes in object-oriented languages, substances can be subclasses of other substances, and can inherit or override variables and behaviors defined in their superclasses. For example, a substance composed of a swarm of microrobots that follow the morphogen gradient could be defined:

> **substance** swarm **is** morphogen **with**:
>
>   **vector field v**   ‖ velocity field
>
>   **scalar field** M ‖ microrobot concentration
>
>   **scalar** μ ‖ microrobot mobility
>
>   **behavior:**
>
> $\mathbf{v} = \mu \nabla a$
> $ÐM = -\nabla \cdot M\mathbf{v}$

A substance can be a subclass of several different substances, thereby permitting it to inherit the properties of those substances. (Additional information can be found in prior publications: MacLennan, 2010, 2011, 2012b).

Bodies

Specific instances of substances are called *fields*, *tissues*, or *bodies* and are created by a **body** declaration, which gives initial values to all of its variables. For example, the following defines a spherical region of morphogen with all of the morphogen concentrated initially near the center:

> **body** MorphogenField(**p**) **of** morphogen:
>   **for** $\|\mathbf{p}\| \leq 1$:
> $D_a = 0.1$
> $\kappa_a = 0.2$
> $\tau_a = 0.01$

$$\vartheta_a = 0.6$$

$$\textbf{for } \|\mathbf{p}\| \le 0.001 : a = 1 \ \| \text{ concentration in center}$$

$$\textbf{for } \|\mathbf{p}\| > 0.001 : a = 0 \ \| \text{ concentration in remainder}$$

The vector variable **p** refers to an arbitrary location in the body. The **for**-declarations then define the spatially distributed initial values of a body of the specified substance. For example, D_a has a value of 0.1 throughout a sphere of radius 1 centered at the origin. The morphogen concentration is $a = 1$ within a smaller sphere (radius 0.001), but $a = 0$ outside of it. The intention is that the initial distributions in **body** declarations be simple so that they can be physically prepared.

Example: Segmentation

To further develop the technique and notation for morphogenetic programming, we have investigated a number of examples of morphogenesis. In each case we mimic processes that occur in biological morphogenesis, but we apply them in different contexts, motivated primarily by robotics.

Segmentation

In vertebrate animals, the number of vertebrae is characteristic of the species. How does an embryo generate such a precise number of segments (e.g., 33 for humans, 55 for chickens)? In vertebrate embryos, segmentation takes place through the clock-and-wavefront process (Cooke & Zeeman, 1976), and this process can be applied to similar ends in artificial morphogenesis. Therefore, our first example explores the assembly of segmented body parts, such as an artificial spine and legs. (We summarize the process here; additional detail can be found in prior publications: MacLennan, 2012a, 2012b.)

We begin with the spinal axis. The assembly is initialized with a small segment of *rostral tissue* ($S = 1$) at the head and a small segment of *caudal tissue* ($T = 1$) at the end (the *tail bud*). The growth duration t_G is regulated by a quantity G, which is initialized to G_0 and decays according to $\mathrm{D}G = -G / \tau_G$. Growth continues so long as G is greater than a threshold ϑ_G, and therefore $t_G = \tau_G \ln(G_0 / \vartheta_G)$. For convenience, we usually set $G_0 = 1$ and $\vartheta_G = 1/e$ so that $t_G = \tau_G$.

In an embryo this growth is a result of cell division; in artificial morphogenesis it is a result of external addition of microrobots, but the specific process is left open in this model. Therefore, we specify only that the microrobots in the tail bud are initialized in a direction **u** pointing away from the head and that they move in that direction at a rate r so long as growth continues. Therefore the effective velocity of these microrobots is $\mathbf{v} = [G > \vartheta_G]r\mathbf{u}$ and the resulting extension of the body is $r\tau_G \ln(G_0 / \vartheta_G)$, typically $r\tau_G$. The movement of the mass of microrobots constituting the tail bud is given by the negative divergence of the microrobot flux $T\mathbf{v}$:

$$\mathrm{D}T = -\nabla \cdot T\mathbf{v} = -T\nabla \cdot \mathbf{v} - \mathbf{v} \cdot \nabla T.$$

The spine grows linearly in the caudal direction through the addition of undifferentiated tissue ($M = 1$) between the head and tail buds, $\mathrm{D}M = rT / \lambda_{\text{TB}}$ where λ_{TB} is the length of the tail bud.

Two morphogens and the wavefront signal can be used to control the number and length of segments. A *caudal morphogen* diffuses from the tail bud, where it is produced up to saturation:

$$\mathrm{D}C = D_C \nabla^2 C - C / \tau_C + \kappa_C T(1 - C).$$

The *rostral morphogen* diffuses similarly from already differentiated tissue at the head end:

$$ÐR = D_R \nabla^2 R - R / \tau_R + \kappa_R S(1-R).$$

So long as growth continues ($G > \vartheta_G$), microrobots in the tail bud generates a pacemaker or clock signal:

$$ÐK = -\omega^2 L,$$
$$ÐL = [G > \vartheta_G] K.$$

When the clock is in the correct phase ($K > \vartheta_K$), the tail tissue generates a pulse of the *segmentation morphogen* α, which will propagate in a wave toward the head. This pulse is represented by $\psi = [G > \vartheta_G \wedge K > \vartheta_K] T$. The microrobot swarm constituting the undifferentiated tissue functions as an excitable medium, and segmentation morphogen concentration above a threshold ($\alpha > \vartheta_a$) causes it to generate a pulse of the morphogen, provided the tissue is not in its refractory period (represented by $\rho < \vartheta_\rho$). Therefore, this pulse of α is represented by $\phi = [\alpha > \vartheta_\alpha \wedge \rho < \vartheta_\rho] M$. The generation and propagation of the segmentation morphogen and the decay of the refractory factor are described by the change equations:

$$Ð\alpha = \phi + \psi + D_\alpha \nabla^2 \alpha - \alpha / \tau_\alpha,$$
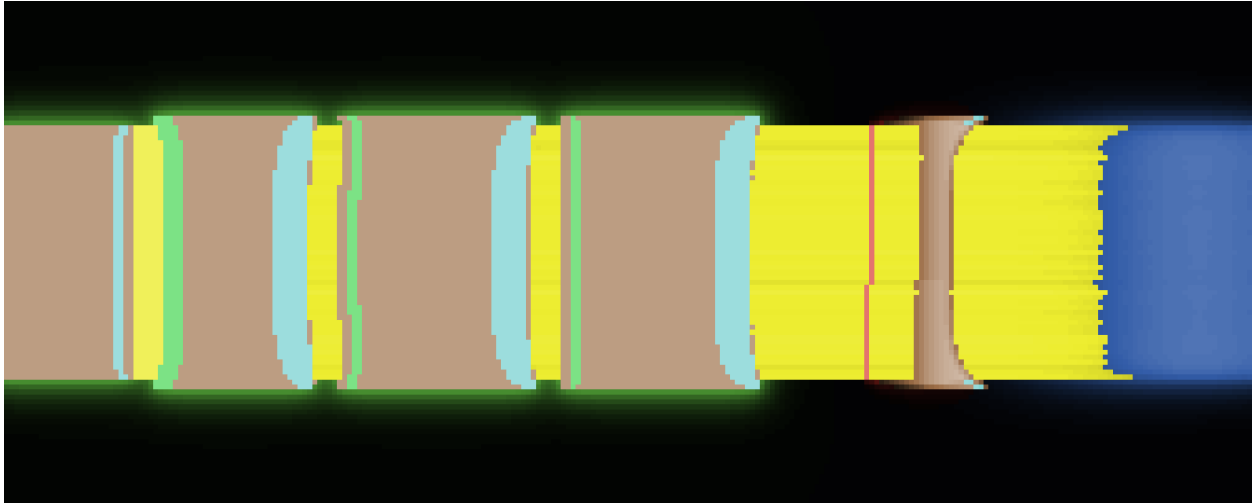$$Ð\rho = \phi - \rho / \tau_\rho.$$

Figure 1. Two-dimensional simulation of segmentation in progress. In addition to the initial head segment on the left, three complete segments have assembled and a fourth in the process of assembly. The thin line in this last segment is the leftward propagating pulse of the segmentation morphogen.

The new segment is formed between already segmented rostral tissue ($S \approx 1$) and the tail bud, a region identified by both rostral and caudal morphogens below threshold (R_{upb}, C_{upb}, respectively), as the wave of segmentation morphogen passes through in a rostral direction (Fig. 1):

$$ÐS = [\alpha > \alpha_{lwb} \wedge R < R_{upb} \wedge C < C_{upb}] k_S + \kappa_S S(1-S).$$

Segment differentiation is triggered by the first term, which is a transient pulse, and the second term continues its logistic growth to saturation (fully differentiated, $S = 1$).

The length λ of the segments is controlled by the ratio of the tail growth rate r and the clock frequency ω in cycles per unit time: $\lambda = 2\pi r / \omega$, since this is the distance the tail bud moves in one clock cycle. The number of segments n is given by the product of the frequency and duration of growth t_G: $n = \lfloor t_G \omega / 2\pi \rfloor$. Therefore, the number and length of the segments is controlled jointly by r, t_G, and ω.

We would like to be able to control the placement of legs within each segment, but this requires that we be able to distinguish the anterior and posterior ends of each segment. In particular, we want to polarize the segments by having the anterior ends of differentiated segment tissue ($S = 1$) further differentiate into anterior tissue ($A = 1$), and the posterior ends into posterior tissue ($P = 1$). This is accomplished by making the differentiated S tissue sensitive to appropriate values of the rostral and caudal morphogens when the segmentation wave passes through:

$$Ð A = \kappa_A S A (1 - A) + [c_A R_{\text{upb}} > R > c_a R_{\text{upb}} \wedge \alpha > \alpha_{\text{lwb}}] k_A - A / \tau_A,$$

$$Ð P = \kappa_P S P (1 - P) + [c_P C_{\text{upb}} > C > c_p C_{\text{upb}} \wedge \alpha > \alpha_{\text{lwb}}] k_P - P / \tau_P.$$

To determine the placement of the legs, we use anterior (a) and posterior (p) morphogens, which accumulate up to saturation in tissue where the anterior and posterior microrobots are sufficiently dense. The a and p morphogens then diffuse from A and P regions, respectively. This behavior is described as follows:

$$Ð a = [A > \vartheta_A] \kappa_a S (1 - a) + D_a \nabla^2 a - a / \tau_a,$$

$$Ð p = [P > \vartheta_P] \kappa_p S (1 - p) + D_p \nabla^2 p - p / \tau_p.$$

It is also easy for microrobots to determine whether they are on the outer surface of the spinal tissue ($E = 1$) as opposed to the interior ($E = 0$). For example, a microrobot can estimate the local population density in S by *quorum sensing*. The local density is near its maximum ($S = 1$ in our units) in the interior, and decreases nearer to the outer surface. The quorum sensing process can be expressed by a convolution, $K \otimes S$, where

$$(K \otimes S)(\mathbf{p}) = \int_{\mathbb{R}^3} K(\mathbf{p} - \mathbf{q}) S(\mathbf{q}) d\mathbf{q},$$

and K is a kernel representing a microrobot's perceptual range. For example, if the microrobots have a certain sensor range, then K will be a sphere of that radius. On the other hand, quorum sensing can be accomplished by having the microrobots emit a slowly diffusing, rapidly degrading signal, which can be sensed to estimate the local population density. In this case K will be a Gaussian kernel. Regardless of how quorum sensing is implemented, being on the surface of the spine is indicated by $E = 1$, where $E = [S_{\text{lwb}} < K \otimes S < S_{\text{upb}}]$.

The position of the "imaginal disks" ($I = 1$), where the legs will be assembled, is controlled by the a and p morphogens and the edge signal E. The morphogens must be in appropriate ranges to control the anterior-posterior position of legs within each spinal segment, and of course the disks form only on the surface:

$$Ð I = [a_{\text{upb}} > a > a_{\text{lwb}} \wedge p_{\text{upb}} > p > p_{\text{lwb}}] \kappa_I E S (1 - I).$$

The clock-and-wavefront process can also be applied to the generation of segmented legs; indeed the same morphogens can be used by changing the clock frequency and initial G value. To assemble appendages of length λ_A we use an initial G value $G_A = \vartheta_G \exp(\lambda_A / r \tau_G)$. To assemble n leg segments, we use a frequency $\omega_A = 2\pi n r / \lambda_A$. However, this process must be properly initialized by generating proper-

ly oriented terminal tissue analogous to the tail bud in spinal generation. This initialization can be triggered by rapid differentiation of the imaginal disk, $ÐI > \vartheta_{DI}$. I increases logistically with a maximum rate $\kappa_I / 2$. Therefore, the length of the initialization period is governed by $\kappa_I / 2 - \vartheta_{DI}$. During the initialization phase the microrobots in the imaginal disks differentiate into terminal tissue, $ÐT = [ÐI > \vartheta_{DI}]\kappa_{AT}I(1-T)$, and reorient themselves in an outward direction,

$$ Ð\mathbf{u} = [ÐI > \vartheta_{DI}]\kappa_A I \left(\frac{\nabla S}{\|\nabla S\|} - \mathbf{u} \right). $$

The same signal triggers rapid resetting of the growth and clock parameters:

$$ ÐG \mathrel{+}= [ÐI > \vartheta_{DI}]\kappa_{G0}I(G_A - G), $$
$$ Ð\omega \mathrel{+}= [ÐI > \vartheta_{DI}]\kappa_{\omega 0}I(\omega_A - \omega). $$

After this initialization, the equations defining the clock-and-wavefront process take over, but generating leg segments as determined by the G_A and ω_A parameters. See Fig. 2.
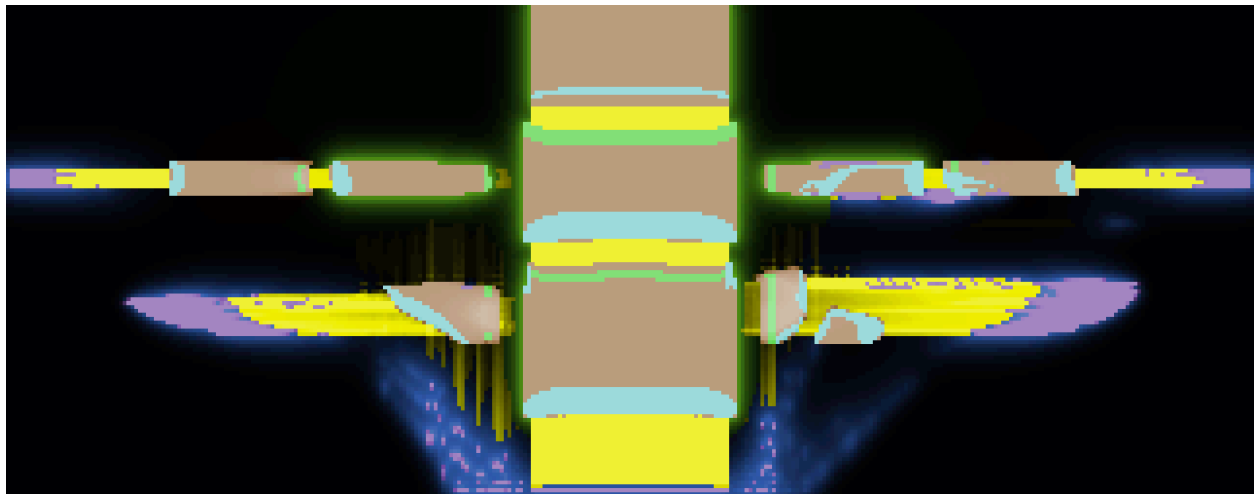

Figure 2. Two-dimensional simulation of growth and segmentation of first two pairs of legs.

Example: Neural Routing

As distant functional areas of the brain must be connected, so must the parts of an artificial nervous system. In embryos, a *growth cone* at the end of a growing axon follows chemical signals to its destination. We can use a similar approach, either by using multiple diffusible attractants to ensure that different simultaneously developing fibers find their destinations correctly, or by creating one fiber at a time, so there is no chance of interference. For the sake of this example, we use the latter, simpler approach. To illustrate the procedure for developing an algorithm for coordinating a massive robot swarm, we will solve this problem in three phases.

In the first phase a growth cone is represented by a single microrobot that generates a fiber from an origin to a destination (see MacLennan, 2012b for additional detail). To establish the pathway we have the goal region ($G = 1$) produce a diffusible attractant A, which also decays to avoid saturation of the growth space:

$$Đ A = D_A \nabla^2 A - A / \tau_A + \kappa_A G(1-A).$$

The growth-cone robot departs from the origin following the attractant gradient and creates the fiber (represented as a concentration of $P \approx 1$) in its wake. In the process of finding its way to its destination, we do not want the new fiber to collide with already established fibers. One way to accomplish this is to have existing fibers emit a repellant R to keep new fibers at a safe distance:

$$Đ R = D_R \nabla^2 R - R / \tau_R + \kappa_R P(1-R).$$

The growth cone then follows the gradient of the difference between the attractant and repellant. However, we want the growth cone to move at a constant rate in spite of the fact that the concentration of the gradient decreases exponentially with distance from the goal. Therefore we set the velocity $\mathbf{v}$ of the growth cone proportional to the versor of the gradient (i.e., to the normed gradient): $\mathbf{v} = r \nabla M / \lVert \nabla M \rVert$ where $M = A - R$. This velocity then determines the change in the growth cone's position, $Đ \mathbf{p} = \mathbf{v}$. (Every substance has an intrinsic vector field $\mathbf{p}$, the spatial location of every particle.) Therefore the growth cone substance is defined as a subclass inheriting the attractant and repellant substances:

> **substance** growth_cone **is** attractant **and** repellant **with**:
>   **scalar** r ‖ cone migration rate
>   **vector field v**   ‖ velocity field
>   **scalar field** M ‖ net attraction
>   **scalar field** P ‖ fiber material
>   **order**-2 **field** σ ‖ diffusion tensor
>   **order**-1 **random W**   ‖ random vector
>   **scalar** κ_P ‖ fiber deposition rate
>   **behavior:**
> $M = A - R$
> $\mathbf{v} = r \nabla M / \lVert \nabla M \rVert$
> $Đ \mathbf{p} = \mathbf{v} + \sigma Đ \mathbf{W}$
> $Đ P = \kappa_P (1 - P)$

Because $Đ \mathbf{W}$ is a normally-distributed random vector (explained in prior reports: MacLennan, 2011, 2012a, 2012b), the term $\sigma Đ \mathbf{W}$ generates a small random perturbation that keeps the growth cone from getting stuck. The fiber material P is deposited by the growth-cone microrobot: $Đ P = \kappa_P (1 - P)$.
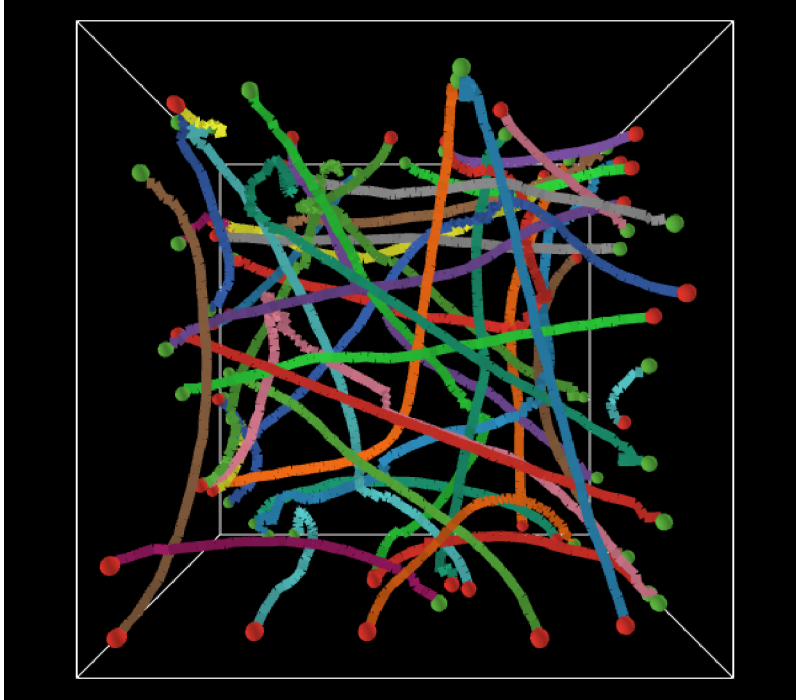
Figure 3. Simulation of path routing. Forty neural fibers are routed between randomly chosen origins and destinations on four surfaces.

Fig. 3 shows a typical result of simulating the sequential generation of 40 neural fibers between randomly generated origins and destinations. The fibers occupy about 4.5% of the interior space without collisions.

For the second phase of algorithm development, we extend this solution from a single agent generating a fiber to a small *cohort* of microrobots, which creates a fiber bundle by departing from the source region and following an attractant gradient to the destination. Each microrobot creates a nerve fiber in its wake. This creates a potential problem, since the streamlines defined by the gradient converge on the destination, so the diameter of the fiber bundle could decrease as it approaches its destination. We can keep its diameter constant by requiring the microrobots to keep within a certain distance range from each other. This suggests the use of a modified flocking algorithm to coordinate the movement of the microrobots (Reynolds, 1987; Spector, Klein, Perry & Feinstein, 2005).

As before, an attractant diffuses from the destination and decays, but instead of requiring a repellant, we assume that the microrobots can detect the proximity of an existing fiber bundle in some way. The acceleration of a microrobot is determined by the weighted sum of six "urges."

$$\mathbf{a} = w_d \mathbf{u}_d + w_a \mathbf{u}_a + w_c \mathbf{u}_c + w_s \mathbf{u}_s + w_v \mathbf{u}_v + w_w \mathbf{u}_w.$$

The *destination urge* $\mathbf{u}_d$ is proportional to the gradient of the attractant. The *avoidance urge* $\mathbf{u}_a$ steers away from already generated fibers within a specified sensor range, thus avoiding collisions between fibers. The *center urge* $\mathbf{u}_c$ is directed toward the centroid of the positions of the robot's "flock mates" (robots within a certain specified radius, if any); this tends to keep the cohort cohesive. The *spacing urge* $\mathbf{u}_s$ steers away from flock mates within a specified critical distance, and thus prevents the cohort from becoming too compact. It is the balance between the center and spacing urges that regulates the size of the cohort, and hence the diameter of the fiber bundle. The *velocity urge* $\mathbf{u}_v$ is the average velocity of

the flock mates, which encourages the robots to move in the same direction. The *wander urge* $\mathbf{u}_w$ is random, which helps prevent the cohort members from getting stuck. The resulting acceleration and velocity are both limited to be below specified maxima.
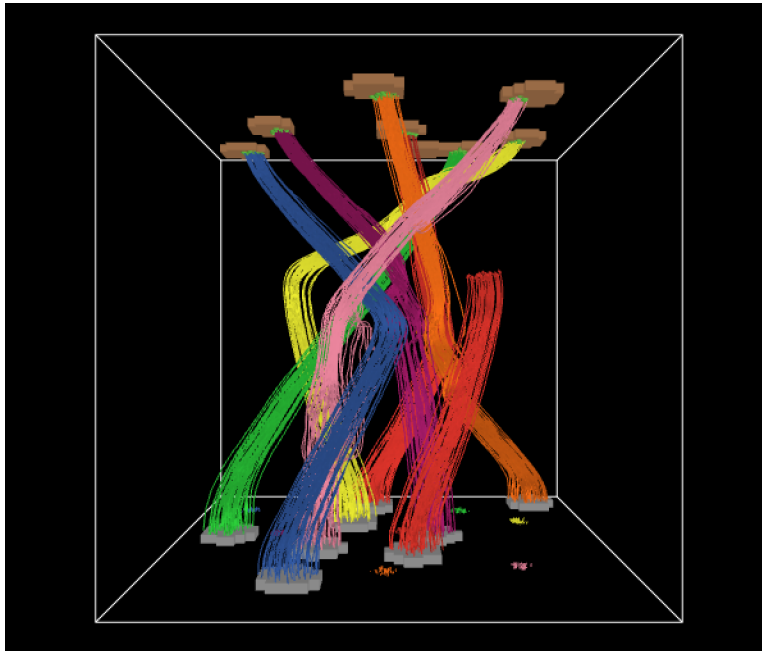

Figure 4. Simulation of the growth of six fiber bundles, with a seventh still growing, each comprising 60 fibers. Origins on the lower surface and destinations on the upper surface were chosen randomly.

Figure 4 shows a simulation of the growth of six fiber bundles, each comprising 60 fibers. Origins and destinations were randomly chosen on the lower and upper surfaces, respectively. The simulation was modified from a flocking model by Wilenski (2005) implemented in NetLogo 5.0.5 (Wilenski, 1999). We have observed that sometimes a microrobot gets separated from its cohort, or that a cohort might split in two to get around another fiber bundle, but the microrobots usually reach their destinations.

For the third phase of our algorithm development, we adopt the continuous approximation, treating an indefinitely large microrobot swarm as a continuous mass. We could have the microrobot mass move like a rigid body, with fixed spatial relations between the robots (with a result like Fig. 4), but such an inflexible approach is not required. Rather, we would like the mass to be able to deform and perhaps even to split in order to avoid obstacles, as a swarm of discrete agents can do. Therefore, we treat the microrobot mass as a nearly incompressible fluid, so that the mass maintains an approximately constant density of agents no matter how it moves. To accomplish this we use a spatially continuous flocking algorithm (cf. Carrillo, Fornasier, Toscani & Vecil, 2010; Chuang, D'Orsogna, Marthaler, Bertozzi & Chayes, 2007; Topaz, Bertozzi & Lewis, 2006).

Rather than using a repellant morphogen to avoid collision, the existing paths clamp the attractant to 0 (e.g., by quickly degrading it). Let A be the attractant concentration, G the goal-region density, and P the existing path density. Then the production, diffusion, decay, and clamping of the attractant is governed by the change equation:

$$Ð A = D_A \nabla^2 A - A / \tau_A + \kappa_G GA - PA / \tau_P.$$

Let C be the density of microrobots. We define a potential function $U(C)$ that is zero for acceptable robot densities and increases rapidly for densities outside this range (either too great or too small). Our

goal is for the robot mass to move up the *A* gradient without the density getting out of bounds, and so the desired velocity $\mathbf{v}'$ is proportional to the difference of the versor of the attractant gradient and the gradient of the density potential:

$$\mathbf{v}' = \frac{\kappa_A \nabla A}{\|\nabla A\|} - \lambda \nabla U(C).$$

where κ_A is the swarm speed and λ governs the importance of constant density. Therefore, we set the acceleration (in the material frame) to $\mathbf{a} = r(\mathbf{v}' - \mathbf{v})$, which steers the mass in the desired direction at a rate *r*. Subtracting the convection term gives the change in velocity in the spatial frame, $\mathbf{a} - \mathbf{v} \cdot \nabla \mathbf{v}$; therefore:

$$\mathrm{Ð}\mathbf{v} = r\left(\frac{\kappa_A \nabla A}{\|\nabla A\|} - \lambda \nabla U(C) - \mathbf{v} \right) - \mathbf{v} \cdot \nabla \mathbf{v}.$$

The microrobot flux is $C\mathbf{v}$, and therefore the changes in robot density and path concentration *P* are governed by:

$$\mathrm{Ð}C = -\nabla \cdot C\mathbf{v},$$
$$\mathrm{Ð}P = \kappa_P C(1 - P).$$

The equation for $\mathrm{Ð}C$ describes the motion of the robot mass in the spatial frame.

Figure 5 (nest page) shows a two-dimensional simulation of the system of PDEs in the Eulerian reference frame. It shows a single path deposited by a continuous mass of microrobots from an origin at the center of the top edge to a destination in the center of the bottom edge, which avoids several obstacles (e.g., previously generated paths).

Conclusions

We have described an approach to coordinating massive microrobot swarms that takes the number of agents to the continuum limit. This allows the use of partial differential equations and continuum mechanics, which guarantees that our coordination strategies scale up to very large number of microrobots (hundreds of thousands, millions, or more). To illustrate the technique, we have presented artificial morphogenesis as a promising approach to the coordination of microrobot swarms to assemble complex hierarchical structures. It is based on biological morphogenetic processes that are known to be effective and are commonly described in terms of partial differential equations. Therefore artificial morphogenesis provides a framework for coordination of massive swarms of microscopic robots for assembly of systems structured from the nanoscale up to the macroscale.
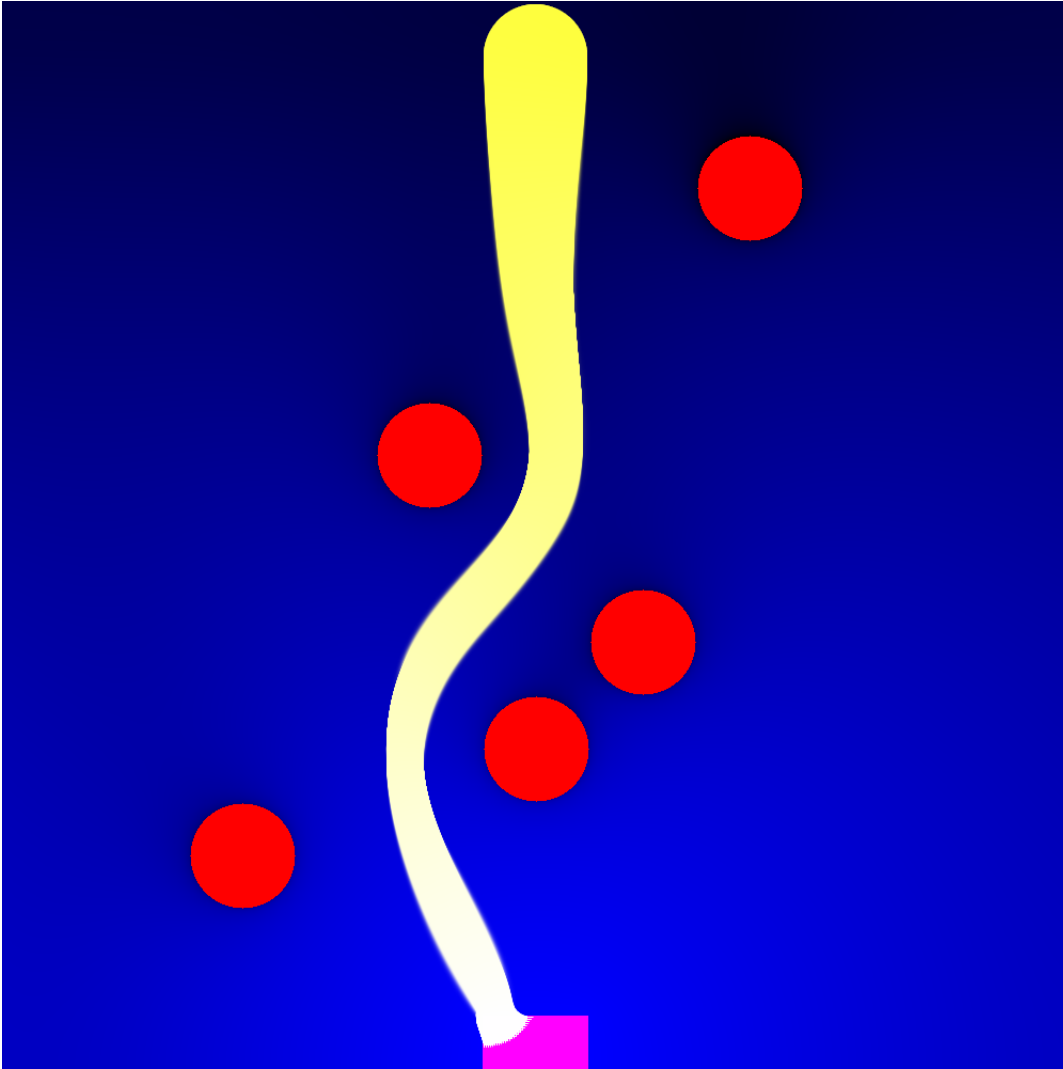
Acknowledgements

Figure 5. Two-dimensional simulation of path creation by massive robot swarm. The background color represents the concentration of an attractant diffusing from the destination at the bottom. The robot swarm emerges from the origin at the top and navigates its way between three obstacles representing previously grown paths in cross section.

References

Abelson, H., Allen, D. Coore, D., Hanson, C., Homsy, G., Knight, Jr., T. F., Nagpal, R., Rauch, E., Sussman, G. J., & Weiss, R. (2000). Amorphous computing. *Communications of the Association for Computing Machinery*, 43(5), 74–82.

Agarwal, R., Bohner, M., O'Regan, D., & Peterson, A. (2002). Dynamic equations on time scales: a survey. *Journal of Computational and Applied Mathematics*, 141, 1–26.

Ahlbrandt, C. D., & Morian, C. (2002). Partial differential equations on time scales. *Journal of Computational and Applied Mathematics*, 141, 35–55.

Beysens, D. A., Forgacs, G., & Glazier, J. A. (2000). Cell sorting is analogous to phase ordering in fluids. *Proceedings of the National Academy of Science USA*, 97, 9467–9471.

Bohner, M., & Peterson, A. (2001). *Dynamic Equations on Time Scales: An Introduction with Applications*. Boston: Birkäuser.

Bonabeau, E. (1997). From classical models of morphogenesis to agent-based models of pattern formation. *Artificial Life*, 3, 191–211.

Carrillo, J. A., Fornasier, M., Toscani, G., & Vecil, F. (2010). Particle, kinetic, and hydrodynamic models of swarming. In G. Naldi, L. Pareschi & G. Toscani (Eds.), *Mathematical Modeling of Collective Behavior in Socio-Economic and Life Sciences* (pp. 297–336). Boston: Birkäuser.

Chuang, Y.-L., D'Orsogna, M. R., Marthaler, D., Bertozzi, A. L., & Chayes, L. S. (2007). State transitions and the continuum limit for a 2D interacting, self-propelled particle system. *Physica D*, 232, 33–47

Cickovski, T. M., Huang, C., Chaturvedi, R., Glimm, T., Hentschel, H. G. E., Alber, M. S., Glazier, J. A., Newman, S. A., & Izaguirre, J. A. (2005). A framework for three-dimensional simulation of morphogenesis. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2(4), 273–278.

Cooke, J., & Zeeman, E. C. (1976). A clock and wavefront model for control of the number of repeated structures during animal morphogenesis. *Journal of Theoretical Biology*, 58, 455–476.

de Gennes, P. G. (1992). Soft matter. *Science*, 256, 495–497.

Edelman, G. M. (1988). *Topobiology: An Introduction to Molecular Embryology*. New York: Basic Books.

Fleischer, K. W. (1995). A multiple-mechanism developmental model for defining self-organizing geometric structures (Doctoral dissertation, California Institute of Technology, 1995). *Dissertation Abstracts International B*, 56/09, 4981, 1996.

Forgacs, G. & Newman, S. A. (2005). *Biological Physics of the Developing Embryo*. Cambridge, UK: Cambridge University Press.

Goldstein, S. C., Campbell, J. D., & Mowry, T. C. (2005). Programmable matter. *Computer*, 38(6), 99–101, June 2005.

MacLennan, B. J. (2010). Morphogenesis as a model for nano communication. *Nano Communication Networks*, 1(3), 199–208.

MacLennan, B. J. (2011). Artificial morphogenesis as an example of embodied computation. *International Journal of Unconventional Computing*, 7(1–2), 3–23.

MacLennan, B. J. (2012a). Embodied computation: Applying the physics of computation to artificial morphogenesis. *Parallel Processing Letters*, 22(3), 1240013.

MacLennan, B. J. (2012b). Molecular coordination of hierarchical self-assembly. *Nano Communication Networks*, 3(2), 116–128.

Meinhardt, H. (1982). *Models of Biological Pattern Formation*. London: Academic Press.

Murata, S., & Kurokawa, H. (2007). Self-reconfigurable robots: Shape-changing cellular robots can exceed conventional robot flexibility. *IEEE Robotics & Automation Magazine*, 71–78.

Nagpal, R., Kondacs, A., & Chang, C. (2003). Programming methodology for biologically-inspired self-assembling systems. In *AAAI Spring Symposium on Computational Synthesis: From Basic Building Blocks to High Level Functionality* (unnumbered). http://www.aaai.org/Papers/Symposia/Spring/2003/SS-03-02/SS03-02-024.pdf (retrieved January 13, 2015).

Nüsslein-Volhard, C. (2008). *Coming to Life: How Genes Drive Development*. Carlsbad, CA: Kales.

Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. *ACM SIGGRAPH Computer Graphics*, 21(4), 25–34.

Salazar-Ciudad, I,. Jernvall, J., & Newman, S. (2003). Mechanisms of pattern formation in development and evolution. *Development*, 130, 2027–2037.

Spector, L., Klein, J., Perry, C., & Feinstein, M. (2005). Emergence of collective behavior in evolving populations of flying agents. *Genetic Programming and Evolvable Machines*, 6(1), 111–125.

Taber, L. A. (2004). *Nonlinear Theory of Elasticity: Applications in Biomechanics*. Singapore: World Scientific.

Topaz, C. M., Bertozzi, A. L., & Lewis, M. A. (2006). A nonlocal continuum model for biological aggregation. *Bulletin of Mathematical Biology*, 68, 1601–1623.

Turing, A. M. (1952). The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, 237, 37–72.

Wilensky, U. (2005). NetLogo Flocking 3D Alternate model. http://ccl.northwestern.edu/netlogo/models/Flocking3DAlternate. Center for Connected Learning and Computer-Based Modeling, Northwestern University.

Wilensky, U. (1999). NetLogo. http://ccl.northwestern.edu/netlogo/. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University.

Key Terms and Definitions

Artificial Morphogenesis: A process for assembling or generating physical structures modeled more or less closely on biological *morphogenesis* (q.v.).

Body: In the context of *artificial morphogenesis* (q.v.), a specific instance of a *substance* (q.v.) occupying a region of space, which may vary as determined by the behavior of its constituent particles. A massive swarm of microrobots is treated as a body.

Caudal: Referring to the tail or posterior part of a structure.

Change Equation: An equation expressing the change in a quantity over time that can be interpreted either as a differential equation in continuous time or as a difference equation in discrete time. The change operator is written Đ.

Eulerian Frame of Reference: An approach to continuum mechanics in which spatially distributed variables are associated with fixed locations in space. Also called *spatial frame of reference*.

Lagrangian Frame of Reference: An approach to continuum mechanics in which spatially distributed variables are associated with fixed particles in a fluid or other substance, which may be moving through space. Also called *material frame of reference*.

Material Frame of Reference: Synonymous with *Lagrangian frame of reference* (q.v.).

Morphogen: A substance that diffuses and governs *morphogenesis* (q.v.) by means of variation in concentration.

Morphogenesis: The development of three-dimensional physical form during the development of an embryo.

Rostral: Referring to the head or anterior part of a structure.

Spatial Frame of Reference: Synonymous with *Eulerian frame of reference* (q.v.).

Substance: In the context of *artificial morphogenesis* (q.v.), a class of infinitesimal particles with particular defined properties and behaviors (analogous to a class in object-oriented programming). Substances are used to define the behavior of particular kinds of microrobots.