Chapter 14
Morphological Adaptation Approaches to Path Planning

"The lights outside tonight are far from home."

(Kurt Wagner, 2000.)

14.1 Path Planning: Unconventional Computing Approaches

Path planning (or motion planning) is a common application of computer science and robotics where a path has to be found between points (typically two points, source and destination point) within an arena. The representation of the arena may already be known or may be discovered by localisation and mapping methods (in this chapter we consider examples where the arena layout is known in advance). The resultant path should be short, minimising distance between the points. Other constraints may also apply, such as requiring paths of sufficient width, avoiding walls, avoiding obstacles, or minimising the number of turns.

Unconventional computing seeks to utilise the computing potential of natural physical systems to solve useful problems. Since these systems are localised in space, they typically use different mechanisms to classical approaches. In recent years physical propagation through space in chemical substrates has been used as a search strategy. Babloyantz first suggested that travelling wave-fronts from chemical reactions in excitable media could be used to approximate spatial problems [262]. Wave propagation in the Belousov-Zhabotinsky (BZ) chemical reaction was subsequently used to discover the path through a maze [72]. In this research a trigger wave was initiated at the bottom left corner of a maze and its propagating wave front recorded by time-lapse photography. Direction of wave propagation was calculated from the collective time-lapse information to give vectors which indicated the direction of the travelling wave. The path from any point on the maze to the

J. Jones, *From Pattern Formation to Material Computation,*
Emergence, Complexity and Computation 15, DOI: 10.1007/978-3-319-16823-4_14

exit (the source of the diffusion) was followed by tracking backwards (using the vector information) to the source.

Wave-front propagation generates a solution from any (and indeed *every*) point in the arena. Branching paths (for example around obstacles) are searched in parallel and the solution time is dependent on the spatial size (in terms of maximum path length) of the arena and the wave-front propagation speed. Although computationally efficient, a direct spatial encoding of the problem (arena, desired start and end points) must be stored, as opposed to a more compact graph or grid encoding in classical approaches.

Reading the output of the parallel calculations is not a simple approach using chemical substrates. Although the propagating wave solves the shortest path for all points in the arena, finding and tracking the desired path from start to end point requires separate processes. Different approaches have been attempted including image processing [263], using two wave-fronts in both directions [264], and hybrid chemical and cellular automata approaches [171]. More recently, a direct visual solution to path planning was devised in which an oil droplet (exploiting convection currents and surface tension effects) migrated along a pH gradient formed within a maze to track the shortest path through the maze [71].

In this chapter we continue the exploration of material computation by morphological adaptation seen in [216] and [243] and examine its application to path planning. Taking inspiration from the behaviour of slime mould, we use a large sheet, or 'blob' of virtual slime mould which is located within an arena in which a path between two points (represented by attractants) must be found. By shrinking this blob over time, it withdraws from the confines of the arena boundary and adapts its shape to connect the start and end points of the path. Examples of the shrinkage method are given in Section 14.2, along with more challenging additions to the problem such as multiple-path options, collision-free paths and obstacle avoidance. We conclude in Section 14.3 by summarising the approach and its contribution to unconventional computing methods of path planning in terms of its simplicity.

14.2 Morphological Adaptation in the Model Slime Mould

We placed a large population of particles within the confines of a 2D arena (Fig. 14.1a), so that the virtual plasmodium completely filled the arena (Fig. 14.1b). Start and end points of the path were represented by projection of attractant into the arena at their respective locations. The virtual plasmodium was attracted to these start and end points. The population size was reduced by adjusting the parameters governing the growth and shrinkage in favour of shrinkage. The collective 'blob' began to shrink and, as it did so, adapted its shape to maintain connectivity to the start and end points and conform to to the borders of the arena (Fig. 14.1c-e). Any extraneous

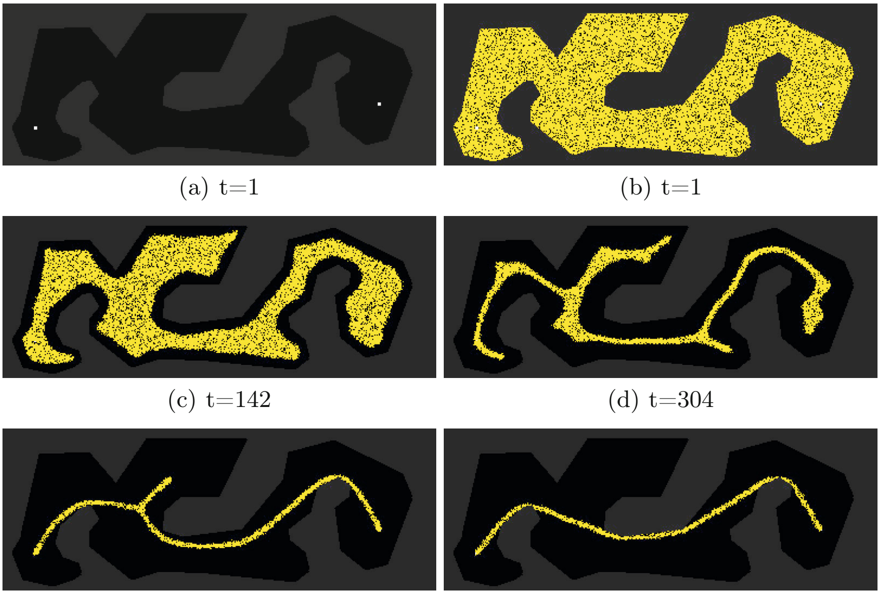

(a) t=1 (b) t=1

(c) t=142 (d) t=304

(e) t=1534 (f) t=4709

Fig. 14.1 Approximation of shortest path between two points by morphological adaptation in a virtual slime mould. a) 2D Arena defined by borders (grey), habitable regions (black) and start and end points (white), b) initialisation of virtual plasmodium in habitable region, c-f) shrinkage of blob causes adaptation of shape and attraction to points, ultimately forming the shortest path between the points.

pseudopodium-like appendages were withdrawn until only a single path connected the start and end points, forming the shortest path between the two points (Fig. 14.1f). This path was composed of a thin band of particles.

14.2.1 Multiple Paths and Many-to-One Path Selection

Most applications for path planning involve finding a path between two points. However it may be necessary to calculate a connecting path between multiple points. This can be achieved in the morphological adaptation approach by having multiple nutrient attractant sources. When a large blob is placed within a polygon arena containing multiple sources, it retains its connection to all points as it shrinks (Fig. 14.2). The final network path (Fig. 14.2d) indirectly connects all source points. Note that the connections between the sources are not simple edges to and from each point. This is not possible because straight edges would pass through the protruding arena boundaries. Instead there is a core curved path running between the outermost points (1 and 4) with pseudopodium-like extensions protruding from this path to connect the inner points (2 and 3). Note that this core path

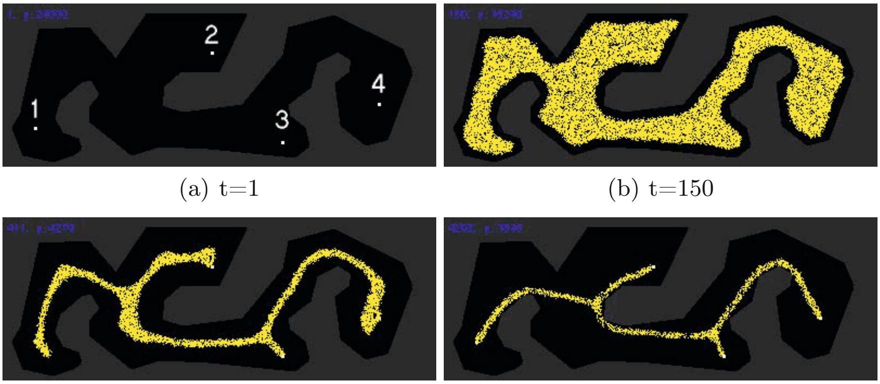

(a) t=1 (b) t=150

(c) t=411 (d) t=4292

Fig. 14.2 Multiple source path planning by morphological adaptation. a) arena with four locations (numbered white points), b-d) shrinkage of virtual plasmodium yields a network with a core path extending from the outermost points with extensions connecting the inner points.

passing between four points differs from that of Fig. 14.1f which passes between only two outer points. The connection between the core path and the two inner points (2 and 3) distorts the core path in the direction of the inner points.

Given this path connecting all four points, how does the morphological adaptation method respond to the removal of certain path options? In Figs. 14.3 and 14.4 we demonstrate the effect of removing different path options. Beginning with the path connecting four points (Fig. 14.3a) we remove points 2 and 3 by deleting the attractant sources at these locations from the lattice. The pseudopodium-like projections connecting these branches to the core path both retract into the core path due to the lack of attractant from these sources (Fig. 14.3b,c). The retracting projections merge with the main flow of particles in the core path. When retraction of these branches is complete the core path connecting points 1 and 4 continues its adaptation to adopt a minimal path between the outer points following the contours of the arena boundary (Fig. 14.3d).

If we remove different source points from the same starting configuration, the adaptation takes a different course. Fig. 14.4 shows the results of a different experiment with the same four points initially connected (Fig. 14.4a). When source points 2 and 4 are removed from the lattice the pseudopodium withdrawal again commences (Fig. 14.4b,c) and the final remaining path adopts a minimal connection between points 1 and 3 (Fig. 14.4d).

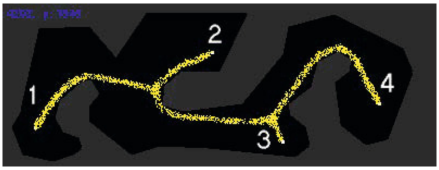

(a) t=4292

(b) t=4415

(c) t=4795

(d) t=7872

Fig. 14.3 From multiple paths to single path by morphological adaptation. a) multiple paths connecting four attractants (numbered), b-d) removal of attractant source from nodes 2 and 3 causes withdrawal of pseudopodia from previous sources and ultimately a single path formed between nodes 1 and 4.

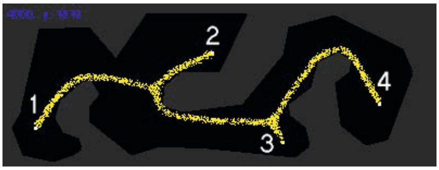

(a) t=4000

(b) t=4550

(c) t=5941

(d) t=10019

Fig. 14.4 From multiple paths to single path by morphological adaptation. a) multiple paths connecting four attractants (numbered), b-d) removal of attractant source from nodes 2 and 4 causes withdrawal of pseudopodia from previous sources and ultimately a single path formed between nodes 1 and 3.

14.2.2 Collision-Free Paths

In many applications a collision-free path may be required, for example if the desired path has to avoid close proximity to walls. To achieve this method by morphological adaptation we represented the walls of the arena as repellent sources (repellent sources project negatively weighted values into the diffusive lattice). We used the same arena as in earlier experiments, but with different start and end points (Fig. 14.5a,b). As the blob shrunk it formed the shortest path (following the walls) when repellent diffusion was not activated (Fig. 14.5c-f). When repellent diffusion from the arena walls was activated the virtual plasmodium still maintained its connectivity to the start and end points but also avoided the diffusing repellent values projecting from the walls of the arena (Fig. 14.5g). Further increasing the concentration of the repellent source increased the distance of the path from the walls (Fig. 14.5h).

14.2.3 Avoiding Obstacles and Preventing Multiple Paths

The response of the morphological adaptation in the presence of obstacles is shown in Fig. 14.6. Again the model is initialised within the habitable region of the arena. When the virtual plasmodium shrinks and adapts its shape in the presence of obstacles we see that multiple paths are formed around the obstacles which connect the start and end points (Fig. 14.6f). Even if we projected repellents from these obstacles, the effect would only be to widen the distance of the multiple paths from these obstacles, and not to form a single path.

The final shrunken multiple paths connecting the nutrient sources are observed because they actually already existed upon initialisation, as the blob was initialised around the obstacles (see Fig. 14.6b). To ensure only a single path is generated we devised a two-part repulsion mechanism. The first part of the mechanism occurs by initialising the blob to cover the *entire* arena (including the obstacles). This part in isolation would not solve the problem of multiple paths, however: If the obstacles repelled the blob immediately then the virtual plasmodium would simply flee the obstacle regions from all directions and multiple paths would still be retained. The second part of the mechanism ensures that only a single path is retained. The shrinkage process is performed more slowly and we generate repellent fields only from obstacles (more specifically, exposed *fragments* of large obstacles) that have been partially uncovered by the shrinkage of the blob.

The mass of the blob is thus shifted away from obstacles by their emergent repulsion field. Because the blob shrinks slowly inwards from the outside of the arena obstacles are slowly uncovered and the repulsion field further pushes the blob inwards until a single path connecting the source attractants is formed. The shrinkage and repulsion mechanism is illustrated in Fig. 14.7 where the arena (including obstacles) is completely covered by a large mass

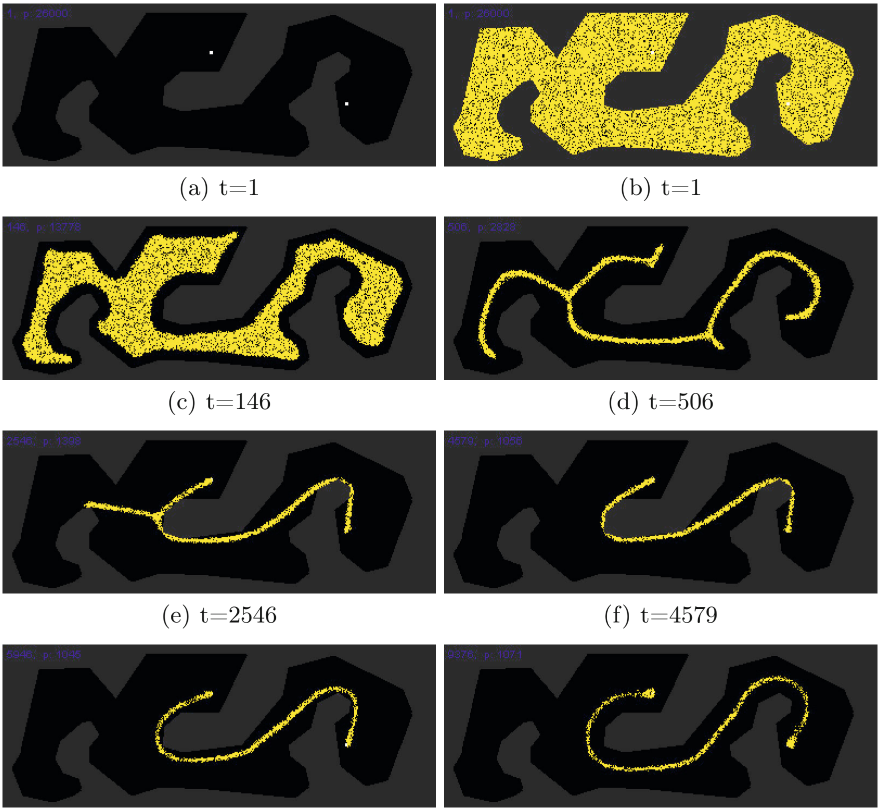

(a) t=1 (b) t=1

(c) t=146 (d) t=506

(e) t=2546 (f) t=4579

(g) t=5946 (h) t=9376

Fig. 14.5 Approximation of collision-free shortest path by morphological adaptation and repulsion. a) 2D Arena defined by borders (grey), habitable regions (black) and start and end points (white), b) initialisation of virtual plasmodium in habitable region, c-f) shrinkage of blob causes adaptation of shape and attraction to points, forming the shortest path between the points, g) repulsion field emitted from arena walls causes virtual plasmodium to avoid wall regions, forming a collision-free path, h) increasing concentration of repulsion field causes further adaptation of the virtual plasmodium away from walls.

of particles comprising the virtual plasmodium (Fig. 14.7b). The blob shrinks inwards as the periphery of the blob is drawn inwards (Fig. 14.7c). When an obstacle is partially uncovered repellent is projected into the diffusive lattice at exposed obstacle fragments (Fig. 14.7d, arrowed). The blob at these regions is repelled and moves away from the exposed obstacle fragment. The shrinkage process continues and when a larger obstacle is partially exposed the repellent projected into the lattice again causes the blob to move away

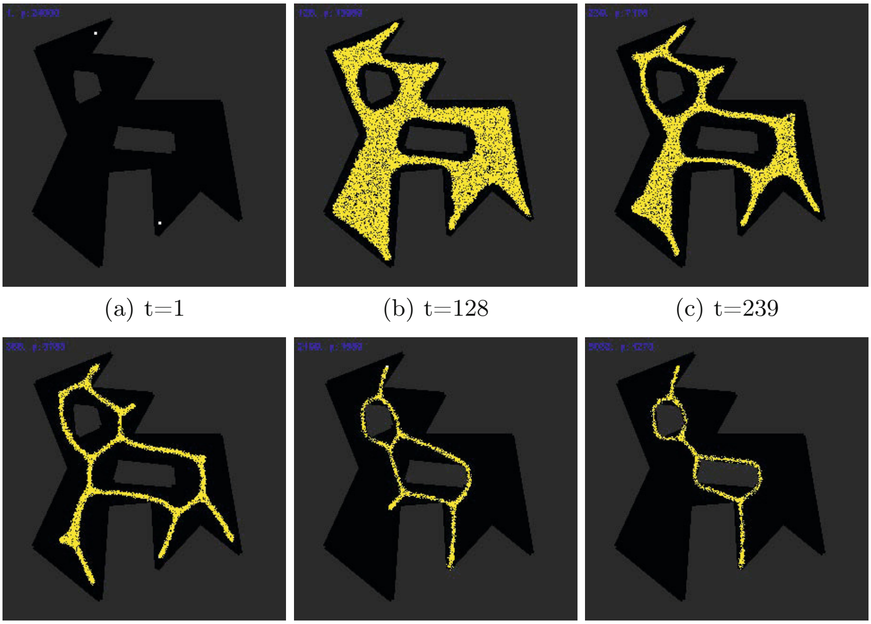

(a) t=1 (b) t=128 (c) t=239

(d) t=368 (e) t=2199 (f) t=5032

Fig. 14.6 Addition of obstacles results in multiple paths connecting start and end points. a) 2D Arena defined by borders and obstacles (grey), habitable regions (black) and start and end points (white), b-e) shrinkage of virtual plasmodium within habitable region, f) shrinkage of virtual plasmodium around obstacles causes multiple paths around obstacles on the path between start and end points.

from this region (Fig. 14.7f, arrowed). Further exposure of this large lower obstacle causes the blob to continue to be repelled away (Fig. 14.7g) until eventually only a single path remains which connects the source attractants and threads between the obstacles (Fig. 14.7h).

In the presence of a large number of obstacles, the repulsion field emanating from newly-exposed obstacles acts to deform the shrinkage of the virtual plasmodium. The mass of particles is deformed both by the attractant stimuli from the start and end points of the path (Fig. 14.8d) and the gradual exposure of the obstacles as the blob shrinks. Fig. 14.8 shows the deformation of the blob and also the changing concentration gradient field as the shrinkage continues, until only a collision-free path between the obstacles remains.

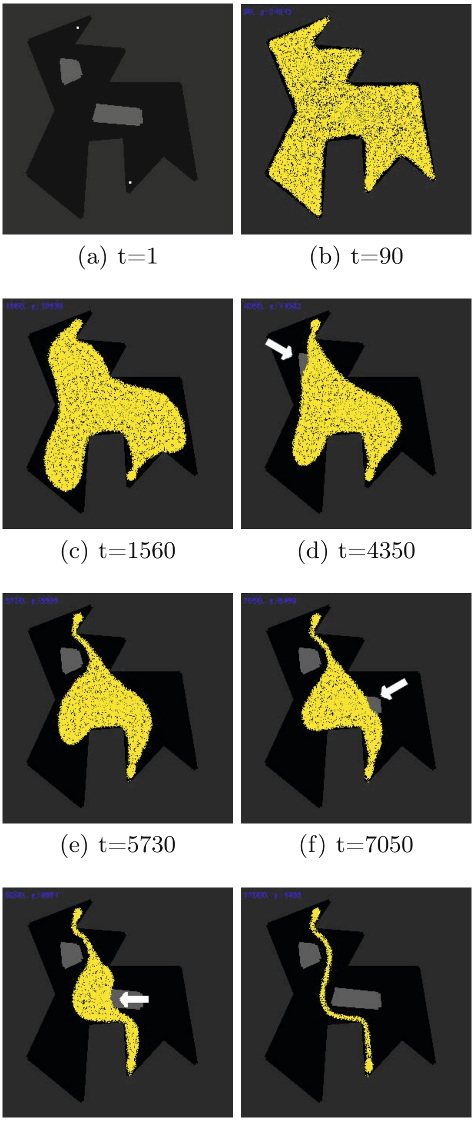

(a) t=1 (b) t=90

(c) t=1560 (d) t=4350

(e) t=5730 (f) t=7050

(g) t=8240 (h) t=17000

Fig. 14.7 Mechanism of shrinkage combined with repulsion at exposed obstacle fragments generates a single path. a) arena with habitable areas (black), inhabitable areas (dark grey), obstacles light grey and path source locations (white). b) blob initialised on entire arena, including obstacles, c) gradual shrinkage of blob, d) exposure of obstacle fragment generates repellent field at exposed areas (arrow), e) blob moves away from repellent field of obstacle, f) lower obstacle is exposed causing repellent field at these locations (arrow), g) further exposure causes migration of blob away from these regions (arrow), h) final single path connects source points whilst avoiding obstacles.

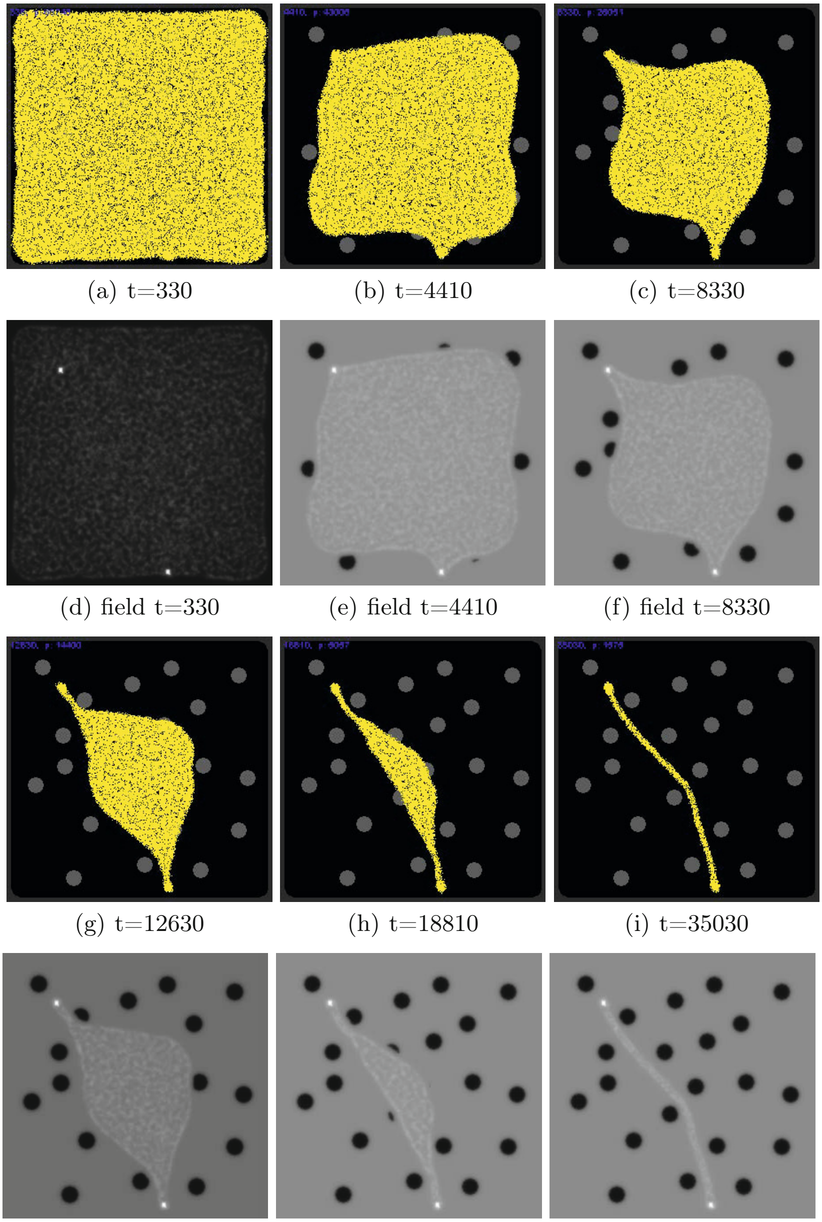

(a) t=330 (b) t=4410 (c) t=8330

(d) field t=330 (e) field t=4410 (f) field t=8330

(g) t=12630 (h) t=18810 (i) t=35030

(j) field t=12630 (k) field t=18810 (l) field t=35030

Fig. 14.8 Shrinkage and exposed repulsion method in complex obstacle field. (a-c) and (g-i) uniform shrinkage of blob is distorted by attraction to start and end stimuli and repulsion from exposed obstacles, (d-f) and (j-l) visualisation of gradient field showing attractants at start and end stimuli (bright spots), blob (mid-grey mass) and repulsion field from exposed obstacles (dark circles). Greyscale gradient field images transformed by gamma correction ($\gamma = 0.6$) to improve clarity.

14.2.4 Model Parameters and Problem Representation

Path start and end points were represented by projection of chemo-attractant to the diffusive lattice at their positions indicated by white pixels on the source configuration images. Lattice size varied with each particular arena but varied from 150 pixels minimum size to 416 pixels maximum size. The attractant projection concentration was 6.375 units per scheduler step. Projection of repellent sources from arena boundaries (used in collision-free planning experiments) was implemented by negatively valued projection (−6.375 units) into the lattice at arena boundary locations causing repulsion of the blob from these regions. Uncovering of obstacles by the shrinking blob acted to project repellent into the lattice (−6.375 units) at exposed areas of obstacles, causing the blob to be repelled from these regions. Each pixel comprising the obstacle was deemed to be uncovered if there were any agent particles within an 11×11 window of the pixel. If an obstacle was covered by particles the repellent projection was suppressed by reducing the projection of repellent to −0.006375 units. This ensured that repulsion from obstacles only occurred when a significant part of the obstacle had been uncovered by the shrinking blob. Diffusion in the lattice was implemented at each scheduler step and at every site in the lattice via a simple mean filter of kernel size 3×3. Damping of the diffusion distance, which limits the distance of chemo-attractant gradient diffusion, was achieved by multiplying the mean kernel value by 0.9 per scheduler step.

The blob was initialised by creating a population of particles and inoculating the population within the habitable regions of the arena The exact population size differed depending on the size of the arena (and habitable area) initial blob was typically composed of between 24000 and 70000 particles. Particles were given random initial positions within the habitable area and also random initial orientations. Particle sensor offset distance (SO) was 7 pixels. Angle of rotation (RA) was set to 45° and sensor angle (SA) was set to 90°. Agent forward displacement was 1 pixel per scheduler step and particles moving forwards successfully deposited 5 units of chemo-attractant into the diffusive lattice. This value is slightly less than the attractant projection value, causing the particles to be anchored to projection sites and ultimately constraining the shape of the shrinking blob. Both data projection stimuli and agent particle trails were represented by the same chemo-attractant ensuring that the particles were attracted to both data stimuli and other agents' trails. The collective behaviour of the particle population was cohesion, minimisation, and morphological adaptation to the configuration of stimuli.

14.2.5 Shrinkage Mechanism

Adaptation of the blob size was implemented via tests at regular intervals. Growth of the population was implemented as follows: If there were between 1 and 10 particles in a 9×9 neighbourhood of a particle, and the particle had moved forward successfully, the particle attempted to divide into two if there was a space available at a randomly selected empty location in the immediate 3×3 neighbourhood surrounding the particle. Shrinkage of the population was implemented as follows: If there were between 0 to 79 particles in a 9×9 neighbourhood of a particle the particle survived, otherwise it was deleted. Deletion of a particle left a vacant space at this location which was filled by nearby particles (due to the emergent cohesion effects), thus causing the blob to shrink slightly. As the process continued the blob shrunk further and adapted its shape to the stimuli provided by the configuration of path source points, arena boundaries and repellent obstacles. The frequency at which the growth and shrinkage of the population was executed determined the turnover rate for the population. The frequency of testing for particle division was every 10 scheduler steps and the frequency for testing for particle removal was every 2 scheduler steps. Since the shrinking blob method is only concerned with the reduction in size of the population it might be asked as to why there were tests for particle division at all. The particle division mechanism was present to ensure that the adaptation of the blob was uniform across the sheet to prevent 'tears' or holes forming within the blob sheet, particularly at the start of an experiment before flux within the blob was initially stabilised.

14.3 Summary: Shaping a Material to Find a Path

We have demonstrated how unconventional computation of path planning problems may be performed directly in 2D space by morphological adaptation in a virtual material inspired by the adaptation of slime mould *Physarum polycephalum*. Unlike previous implementations of path planning problems in chemical substrates the method does not rely on a two-stage computation (one stage to perform the computation, another stage to highlight the path). The method computes a simple path with only two attractant sources. Multiple paths were represented by having more than two attractant sources and a single path was selected between two of these sources by removal of redundant sources. Collision-free paths were discovered by the simultaneous addition of repellent sources at arena boundaries. Obstacle avoiding paths were discovered using a mechanism whereby obstacles were represented by a gradual exposure of repellent sources. The contribution of this method is in the simplicity of the approach: the behaviour of the shrinking blob is distributed within the material itself and emerges from the simple and local interactions between the particles which comprise the blob. The path finding process is governed, to a large extent, by the spatial configuration of the arena and the obstacles within the arena. Since the blob initially occupies all of the space within the arena the path finding method may be described as subtractive — all redundant or inefficient paths are removed during the shrinkage process. This is achieved by withdrawal of pseudopodia (for example from dead-ends in the arena) and also by displacement of the blob by the repellent fields emitted from the gradually exposed obstacles. Unlike chemical-based approaches the method is not initiated at either the path start or end points but is initialised by shrinkage from the arena boundary. Diffusion from the source points still occurs but is merely used to anchor the blob material at these points and does not require propagation of the diffusion front throughout the entire arena. Likewise, repellent diffusion occurs from the boundary (for collision-free paths) and from obstacles (for obstacle-avoiding paths) but this diffusion also is only local and does not require propagation throughout the entire arena.

In this and previous chapters we have examined the network adaptation behaviour of the multi-agent model. The relative predictability of the model adaptation is due to the smooth particle flow that occurs when particle collisions occur under the default motor behaviour. In the *Physarum* plasmodium the protoplasmic flux is not as uniform or predictable. Differences in protoplasmic transport cause temporary blockages, changes in pressure, surging transport of material and oscillating thickness of the plasmodium, and network adaptation is typically more complex and unpredictable. Chapters 15 and 16 demonstrate how such oscillatory phenomena emerge (and can be usefully utilised) in the virtual plasmodium when smooth particle flux is de-activated.