Implementing Reaction-Diffusion on an Amorphous Computer

Daniel Coore, Radhika Nagpal *
{newts, radhi}@martigny.ai.mit.edu
MIT Artificial Intelligence Laboratory

Abstract

Traditionally reaction-diffusion simulations have been performed on regularly discretised spaces. In an Amorphous Computer [1] the processors are distributed densely but randomly. By implementing the Gray-Scott model on an amorphous computer simulation, we show that the patterns obtained by uniform discretisation can also be realised qualitatively on irregular, amorphous discretisations if diffusion is modeled correctly.

1 Introduction

The Gray-Scott model of reaction-diffusion is based upon the following chemical system involving chemical species U, V and P:

$$U + 2V \rightarrow 3V$$

$$V \rightarrow P$$

The reaction kinetics of this system give rise to non-linear partial differential equations for u and v, the respective concentrations of U and V.

$$\frac{\partial u}{\partial t} = r_u \nabla^2 u - uv^2 + f(1-u) \quad (1)$$

$$\frac{\partial v}{\partial t} = r_v \nabla^2 v + uv^2 - (f+k)v \quad (2)$$

*Coore and Nagpal are advised by Profs. G.J. Sussman and H. Abelson and are supported in part by DARPA under contract number N00014-96-1-1228 administered by ONR. Nagpal is also supported by AT&T GRPW Program. Address: Room 434, 545 Technology Square, Cambridge MA 02139.

Pearson[2] has catalogued the patterns obtained at twelve points in f, k parameter space. In those simulations, space was discretised into 256×256 squares, each of which acted as a reaction locus. Our goal was to try to obtain qualitatively equivalent patterns at similar locations in parameter space, despite the irregular spatial discretisation.

2 Modeling Reaction-Diffusion

In an amorphous computer, the processors have no precision interconnect, and are placed randomly, often in a way that tries to keep the variance in density small. Processors communicate locally with all processors within some small but fixed radius of communication. Each processor represents a reaction site, and maintains information on the concentrations of the different chemical species at that locus. Reaction is modeled independently of the layout of the loci. Modeling diffusion however, is far more interesting because it involves neighbour interaction, and is therefore critically dependent on the layout of the loci.

The finite difference model of the diffusion process on a regular grid typically computes the mean neighbour value of u to represent the term $\nabla^2 u$. Extending this idea to the amorphous layout leads to the following approximation:

$$\nabla^2 u(x) \approx \sum_{x_n \in N(x)} \frac{u(x_n) - u(x)}{|N(x)|}$$

$N(x)$ represents the set of neighbours of the processor at position x. However, this approximation does not conserve the total amount of U in the system since the value of $|N(x)|$ is no longer the same at every position x.

If we assume that for each processor the density of its neighbours is the same in all radially outward directions, then we obtain:

$$\nabla^2 u(x) \approx \sum_{x_n \in N(x)} (u(x_n) - u(x)).$$

This approximation solves the conservation problem but has stability problems. In a forward-Euler integration of the system, the single step decrease in u at a locus can sometimes be larger than the current value of u (since r_u and Δt are set by the user). Indeed we observed negative values at some loci, which also caused extremely large values at other sites (since the total quantity of U is conserved). To remedy this problem, we can divide the sum by a constant N, chosen so that $N \geq \max_x |N(x)|$.

There is yet another model for diffusion. It is based on the observation that the ∇^2-operator is a continuous description of an essentially discrete process at the molecular level. Since our simulation domain is discrete to begin with, we can model diffusion at the micro scale by performing exchanges of U quantities between neighbouring processors as if actual molecules were being exchanged. In this model, the amount transferred from a processor at x to a neighbour at y is proportional to $u(x) - u(y)$. Transfers occur atomically, so there is never a danger of producing negative quantities of U at any site. This is the diffusion model that we used to simulate the Gray-Scott reactions.

2.1 Implementation

There were two implementations of the Gray-Scott simulator. In the first simulator each processor was chosen in turn to undergo the transfer process with all its neighbours. In the second simulator, processors were chosen at random to transfer molecules to all their neighbours. We found that both simulators produced qualitatively similar patterns, showing that reaction-diffusion can be modeled asynchronously.

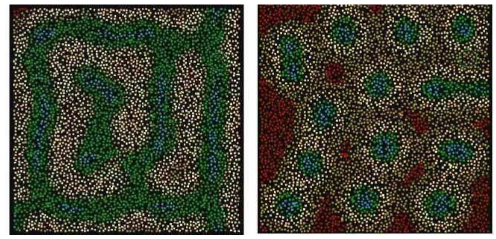

Figure 1: Patterns (left to right) obtained when $f = .04$, $k = .06$ and $f = .035$, $k = .065$ (labelled θ and λ by Pearson)

3 Results

Figure 1 shows the patterns obtained at the indicated parameter values with the synchronous simulator, written in Java[1]. These figures are qualitatively similar to those produced by Pearson [2], indicating that regular discretisation is not an essential feature for interesting pattern formation in the Gray-Scott model of reaction-diffusion.

References

[1] H. Abelson, T. F. Knight, and G. J. Sussman. Amorphous computing. *White paper*, October 1995.

[2] J. E. Pearson. Complex patterns in a simple system. *Science*, 261:189–192, July 1993.

[1] http://www-swiss.ai.mit.edu/ switz/amorphous/GrayScott