Towards the Coevolution of Cellular Automata Controllers for Chemical Computing with the B-Z Reaction

Christopher Stone, Rita Toth, Andrew Adamatzky, Ben de Lacy Costello & Larry Bull
Unconventional Computing Group
University of the West of England
Bristol BS16 1QY, U.K.
+44 (0)117 3283161
{Christopher3.Stone, Rita.Toth, Andrew.Adamatzky, Ben.Delacycostello, Larry.Bull}@uwe.ac.uk

ABSTRACT
We propose that the behaviour of non-linear media can be controlled automatically through coevolutionary systems. By extension, forms of unconventional computing, i.e., massively parallel non-linear computers, can be realised by such an approach. In this study a light-sensitive sub-excitable Belousov-Zhabotinsky reaction in which a checkerboard image comprised of varying light intensity cells projected onto the surface of a catalyst loaded gel is controlled using a heterogeneous cellular automaton. Pulses of wave fragments are injected onto the gel resulting in rich spatio-temporal behaviour and a coevolved cellular automaton is shown able to either increase or decrease the chemical activity through dynamic control of the light intensity within each cell in both simulated and real chemical systems.

Categories and Subject Descriptors
I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods and Search – *backtracking, control theory, dynamic programming, graph and tree search strategies, heuristic methods, plan execution formation and execution, scheduling.*

General Terms
Algorithms, Experimentation.

Keywords
Cellular Automata, Coevolution, Hill-climber, Non-linear media, Unconventional computing.

1. INTRODUCTION
There is growing interest in research into the development of hybrid wetware-silicon devices focused on exploiting their potential for 'non-linear computing.' The aim is to harness the as yet only partially understood intricate dynamics of non-linear media to perform complex 'computations' more effectively than with traditional architectures and to further the understanding of how such systems function. The area provides the prospect of radically new forms of machines. We are developing an approach by which networks of non-linear media — reaction-diffusion systems — can be produced to achieve a user-defined computation in a way that allows direct control of the media. Coevolutionary algorithms are used to design the appropriate networks by searching a defined behavioural space to create a computing resource capable of satisfying a given objective(s). In this paper we examine a Belousov-Zhabotinsky (BZ) [20] reaction-diffusion system in which the networks are created via light and we present initial results from the general control/programming scenario.

Excitable and oscillating chemical systems have been used to solve a number of computational tasks (see [4] for a recent review). In this paper we adapt a system described by Wang *et al*. [19] and explore the computational potential based on the movement and control of wave fragments. In the system they describe, the application of Gaussian noise (where the mean light level is fixed at the subexcitable threshold of the reaction) in the form of light projected onto a light sensitive analogue of the BZ reaction was observed to induce wave formation and subsequently "avalanche behaviour" whereby a proliferation of open ended excitation wave fragments were formed.

Evolutionary Algorithms (EA) (e.g., [6]) are being increasingly used in the design and analysis of complex systems. Example applications include data mining, time series analysis, scheduling, process control, robotics and electronic circuit design. Such techniques can be used for the design of computational resources in a way that offers substantial promise for application in non-linear media computing since the algorithms are almost independent of the medium in which the computation occurs. This is important in order to achieve effective non-linear media computing since an EA does not need to directly manipulate the material to facilitate learning and the task itself can be defined in a fairly unsupervised manner. In contrast, most traditional learning algorithms use techniques that require detailed knowledge of and control over the computing substrate involved. Indeed, Harding and Miller [10] have recently described the use of an EA to design a computational system using liquid crystals. In this paper we control a BZ network via an approach which uses coevolutionary computing to create heterogeneous Cellular Automata (CAs).

GECCO'07, July 7–11, 2007, London, England, United Kingdom.

2. MODEL

The features of the chemical system are simulated using a two-variable Oregonator model modified to account for the photochemistry [7, 12, 15]:

$$\frac{\partial u}{\partial t} = \frac{1}{\varepsilon}\left(u - u^2 - (fv + \Phi)\frac{u - q}{u + q} \right) + D_u \nabla^2 u$$

$$\frac{\partial v}{\partial t} = u - v \,.$$

The variables u and v represent the instantaneous local concentrations of the bromous acid autocatalyst and the oxidized form of the catalyst, $HBrO_2$ and tris (bipyridyl) Ru (III), respectively, scaled to dimensionless quantities. The ratio of the time scales of the two variables, u and v, is represented by ε, which depends on the rate constants and reagent concentration; f is a stoichiometric coefficient. The rate of the photo-induced bromide production is designated by Φ, which also denotes the excitability of the system in which low light intensities facilitate excitation while high intensities result in the production of bromide that inhibits the process, experimentally verified by Kádár et al. [12]. The scaling parameter, q, depends on reaction rates only. The system was integrated using the Euler method with a five-node Laplacian operator, time step Δt=0.001 and grid point spacing Δx=0.62. The diffusion coefficient, D_u, of species u was unity, while that of species v was set to zero as the catalyst is immobilized in the gel. The parameters were ε = 0.11, f = 1.1 and q = 0.0002. The system was oscillatory at the dark state which made it possible to initiate waves in a cell by setting its light intensity to zero. The system was subexcitable at Φ=0.04.

3. COEVOLUTIONARY CONTROL

The characteristics of the chosen chemical system are very much akin to those of two-dimensional cellular automata, such as the Game of Life [8]. That is, fragments of excitation travel across the surface of the gel, often colliding to form other fragments or self-extinguishing, as do the gliders in "Life." Further, the light projections which cause such behaviour are arranged in a regular grid of cells over the gel surface. We are therefore interested in using cellular automata to control the behaviour of the fragments to implement computation, particularly forms of collision-based computing (e.g., [1]).

Previously, several good results from the evolution of cellular-automaton rules to perform useful tasks have been published. Mitchell *et al*. (e.g., [16][17][5]) have investigated the use of EAs to learn the rules of uniform one-dimensional, binary CAs. There a Genetic Algorithm (GA) [11] produces the entries in the update table used by each cell, candidate solutions being evaluated with regard to their degree of success for the given task — density and synchronization. Andre et al. [2] repeated Mitchell *et al*.'s work, using Genetic Programming [2] to evolve update rules. They report similar results.

More related to our approach, Sipper [18] has presented a non-uniform, or heterogeneous, approach to evolving CAs. Each cell of a one- or two-dimensional CA is also viewed as a GA population member, mating only with its lattice neighbours and receiving an individual fitness. He shows an increase in performance over Mitchell et al.'s work, exploiting the potential for spatial heterogeneity in the tasks. However the reliance upon each cell having access to its own fitness means it is not applicable in the majority of chemical computing scenarios we envisage; like Mitchell *et al*.'s work, fitness will be based on emergent global phenomena in our approach. Also, we expect a high degree of spatial heterogeneity will typically be required rendering the use of recombination between cell controllers somewhat superfluous. Thus, following [13], we begin here by using a coevolutionary approach wherein each cell of a two-dimensional CA controller is developed via a genetic hill-climber.

For a given experiment, a random set of CA controllers is created for a two-dimensional array of size 10-by-10, i.e., 100 cells. The grid edges are not connected and the neighbourhood size of each cell is of radius 1; cells consider neighbourhoods of varying size depending upon their spatial position, varying from three in the corners, to five for the other edge cells, and eight everywhere else. In the model each of the 100 cells consists of 400 (20-by-20) simulation points for the reaction, as defined above.

Waves were initiated by setting the light intensity to zero for all grid points within one cell at the bottom. These waves were broken up into 12 fragments by choosing an appropriate light pattern as shown in Figure 2(a). The black area represents the excitable medium whilst the white area is non excitable. After the initiation three light levels were used: one is sufficiently high to inhibit the reaction; one at the sub-excitable threshold such that excitation just manages to propagate; and, the other low enough to fully enable it. The modeled chemical system is then run for 10 seconds of simulated real time. A 100-bit description of the grid is then passed to the CA. Each bit corresponds to a cell and it is set to true if the average level of activity within the given cell is greater than a pre-determined threshold of 10%. The CA returns a 100-digit trinary action string, each digit of which indicates whether the high (Φ=0.093023), threshold (Φ=0.04) or low (Φ=0.000876) intensity light should be projected onto the given cell, ascertained by each CA cell controller considering its own state and that of its neighbours in the traditional way. Another 10 seconds of real time are then simulated with those light-levels projected, etc. until 25 iterations have passed.

After 25 iterations, the fitness of the emergent behaviour is calculated. We here consider two cases: to increase the amount of excitation on the gel surface, i.e., the CA controller must create new fragments; and, to decrease the amount of excitation. The global fitness is assigned to each CA cell. Some proportion, here 1000 randomly chosen genes, of the CAs are then mutated. That is, the transition value for one of the possible states it considers is changed to one of the other possible light intensities. The simulation is reset and repeated as described.

The mutation algorithm keeps track of which CA states are visited since mutation. On the next fitness evaluation (at the end of a further 25 iterations) mutations in states that were not visited are discarded on the grounds that they have not contributed to the global fitness value and are thus unreliable estimators of worth. Fitness values are calculated as a running average across learning problems using the Widrow-Hoff delta rule with learning rate β=0.2: fitness ← fitness + β[new fitness – fitness]

4. RESULTS

Figure 1(a) shows the fitness of our coevolutionary approach averaged over ten runs for the inhibition task wherein fitness is calculated as the number of cells in the grid that have an activity level less than the 10% threshold. As can be seen, the amount of excitation decreases during learning. Figure 2 shows snapshots of the spatio-temporal behaviour of a typical solution produced.

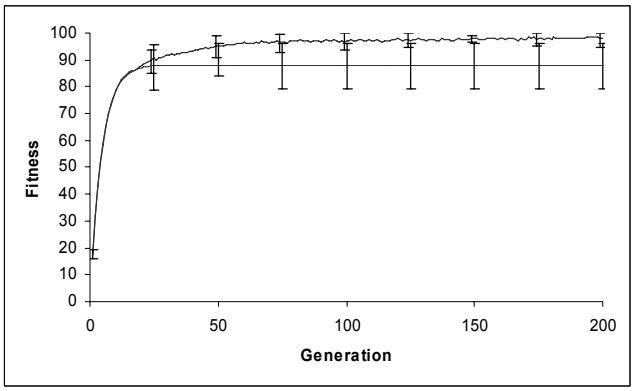

(a)

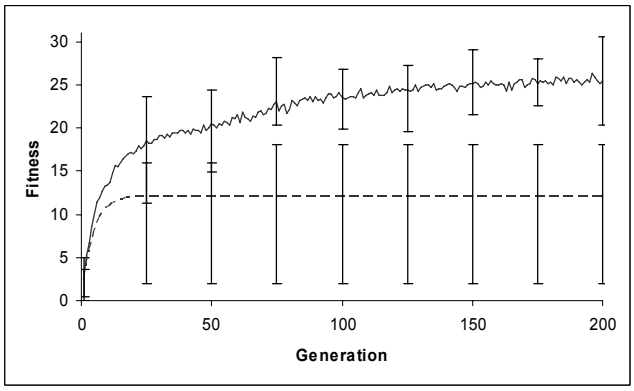

(b)

Figure 1. Showing the typical fitness over time for both the inhibition (a) and excitation tasks (b) with the model. Solid lines: breeders1000, dashed lines: breeders 0.

Figure 1(b) shows the fitness of our coevolutionary approach averaged over ten runs for the excitation task wherein fitness is calculated as the number of cells in the grid that have an activity level greater than or equal to the 10% threshold. As can be seen, the amount of excitation increases during learning. Figure 3 shows snapshots of the spatio-temporal behaviour of a typical solution produced. If the task were to have been tackled by a person with prior knowledge of the BZ reaction the number of fragments would be increased simply via the predominant projection of the lowest light level. However, it appears from experimental observation that the number of fragments and thus the total excitation level is increased via the projection of appropriate high and low light intensity cells placed by the CA so as to manoeuvre and split existing fragments. However, the exact mechanism of this process is yet to be fully understood.

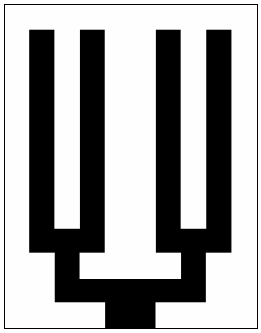

(a)

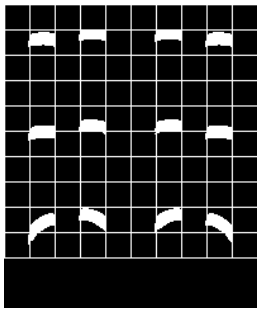

(b)

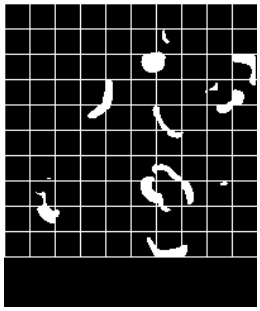

(c)

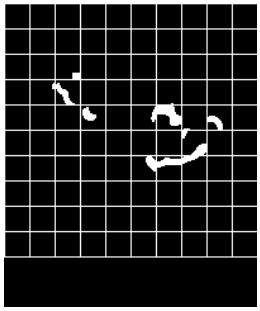

(d)

Figure 2. Showing initialization pattern (a) and example solutions to the modelled inhibition task after initialization (b) and coevolved at generations 30 (c) and 200 (d). State on last of the 25 control cycles shown.

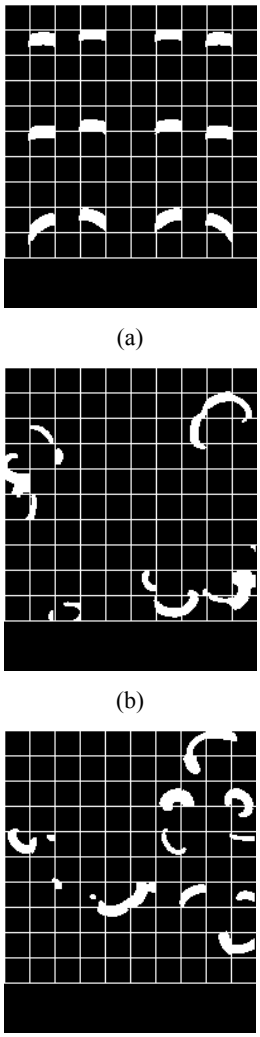

(a)

(b)

(c)

Figure 3. Showing example solutions to the modelled excitation task after initialization (a) and coevolved at generations 30 (b) and 200 (c). State on last of the 25 control cycles shown.

As noted above, we have made two modifications to the standard genetic hill-climber approach in that we introduce a running average of successive fitness evaluations and remove non-utilized mutations. Figure 4 shows the fitness of our coevolutionary cellular automata approach averaged over ten runs of the excitation task with the running fitness average mechanism disabled. As can be seen, little or no learning is observed over the same period that learning is clearly seen in Figure 1(b).

We suggest that the high dependence between the CA cells and the non-linear nature of the chemistry means that the standard approach is unable to cope with the constant fitness landscape distortions experienced by the set of CA controllers every time a mutation is introduced.

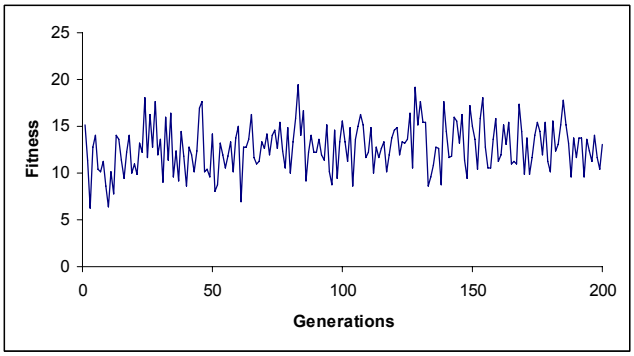

Figure 4. Showing performance on the excitation task without a running fitness average

5. EXPERIMENTAL IMPLEMENTATION

Given the success of our approach using the simulated chemistry, we have attempted to implement the same two tasks using a real chemical reaction constructed using the following methodology.

Sodium bromate, sodium bromide, malonic acid, sulphuric acid, tris(bipyridyl) ruthenium (II) chloride, 27% sodium silicate solution stabilized in 4.9 M sodium hydroxide were purchased from Aldrich and used as received unless stated otherwise.

To create the gels a stock solution of the sodium silicate solution was prepared by mixing 222 mL of the purchased sodium silicate solution with 57 mL of 2 M sulphuric acid and 187 mL of deionised water [19]. $Ru(bpy)_3SO_4$ was recrystallised from the chloride salt with sulphuric acid [9]. Pre-cured solutions for making gels were prepared by mixing 2.5 mL of the acidified silicate solution with 0.6 mL of 0.025 M $Ru(bpy)_3SO_4$ and 0.65 mL of 1.0 M sulphuric acid solution. Using capillary action, portions of this solution were quickly transferred into a custom-designed 25 cm long 0.3 mm deep perspex mould covered with microscope slides. The solutions were left for 3 hours to permit complete gellation. After gellation the adherence to the Perspex mould is negligible leaving a thin gel layer on the glass slide. After 3 hours the slides were carefully removed from the mould and the gels on the slides were washed in deionised water at least five times to remove by products. The gels were 26 mm by 26 mm, with a wet thickness of approximately 300 μm. The gels were stored under water and rinsed right before use.

The catalyst-free reaction mixture was freshly prepared in a 30 mL continuously-fed stirred tank reactor (CSTR), which involved the *in situ* synthesis of stoichiometric bromomalonic acid from malonic acid and bromine generated from the partial reduction of sodium bromate. This CSTR in turn continuously fed a thermostatted open reactor with fresh catalyst-free BZ solution in order to maintain a nonequilibrium state. The final composition of the catalyst-free reaction solution in the reactor was: 0.42 M sodium bromate, 0.19 M malonic acid, 0.64 M sulphuric acid and 0.11 M bromide. The residence time was 30 minutes.

An InFocus Model Projector was used to illuminate the computer-controlled image. Images were captured using a Lumenera Infinity2 USB 2.0 scientific digital camera. The open reactor was

surrounded by a water jacket thermostatted at 22 ^{o}C. Peristaltic pumps were used to pump the reaction solution into the reactor and remove the effluent. A diagrammatic representation of the experimental setup is shown in Figure 5.

The spatially distributed excitable field on the surface of the gel was made possible by the projection of a 10-by-10 cell checkerboard grid pattern generated using a computer. The checkerboard image comprised of light levels from a low of 0.35 mW cm^{-2} to a high of 3.5 mW cm^{-2} intensity cells, in three equal bins, representing excitable, the threshold, and non-excitable domains respectively.

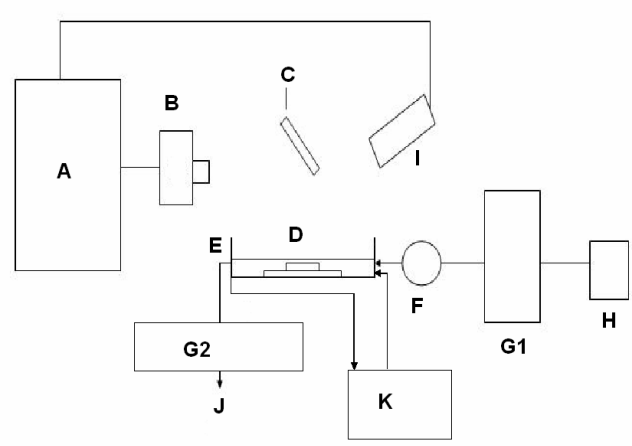

Figure 5: A block diagram of the experimental setup where A: computer, B: projector, C: mirror, D: microscope slide with the catalyst-laden gel, E: thermostatted Petri dish, F: CSTR, G1 and G2: pumps, H: stock solutions, I: camera, J: effluent flow, K: thermostatted water bath.

The checkerboard grid pattern was projected onto the catalyst-laden gel through a 455nm narrow bandpass interference filter and 100/100mm focal length lens pair and mirror assembly. The size of the projected grid was approximately 20mm square. Every 10 seconds, the checkerboard pattern was replaced with a uniform grey level of 3.5 mW cm^{-2} for 10 ms during which time an image of the BZ fragments on the gel was captured. The purpose of removing the grid pattern during this period was to allow activity on the gel to be more visible to the camera and assist in subsequent image processing of chemical activity.

Captured images were processed to identify chemical wave activity. This was done by differencing successive images on a pixel by pixel basis to create a black and white thresholded image. Each pixel in the black and white image was set to white, corresponding to chemical activity; if the intensity of the red or blue channels differed in successive images by more than X out of Y pixels (1.95%). Pixels at locations not meeting this criterion were set to black. The images were cropped to the grid location and the grid superimposed on the thresholded images to aid analysis of the results.

6. EXPERIMENTAL RESULTS

Figures 6 and 7 show examples of the spatio-temporal dynamics exhibited by the real chemical system, averaged over three runs each. In both cases very similar behaviour is seen as was observed in simulation. Moreover, our coevolutionary approach appears able to control the chemical system to achieve the desired goal to a similar degree of accuracy, i.e., fitness, as was seen in the simulations, as shown in Figures 8(a) and (b) (compare with Figures 1(a) and (b) over the same period).

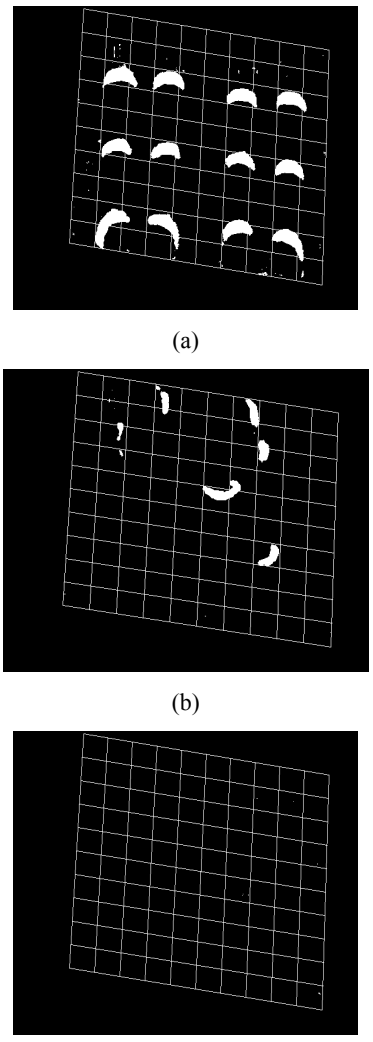

(a)

(b)

(c)

Figure 6. Showing an example solution to the inhibition task run in the real chemistry after initiation (a), generation 12 (b) and 30 (c). State on the last of the 25 control cycles is shown.

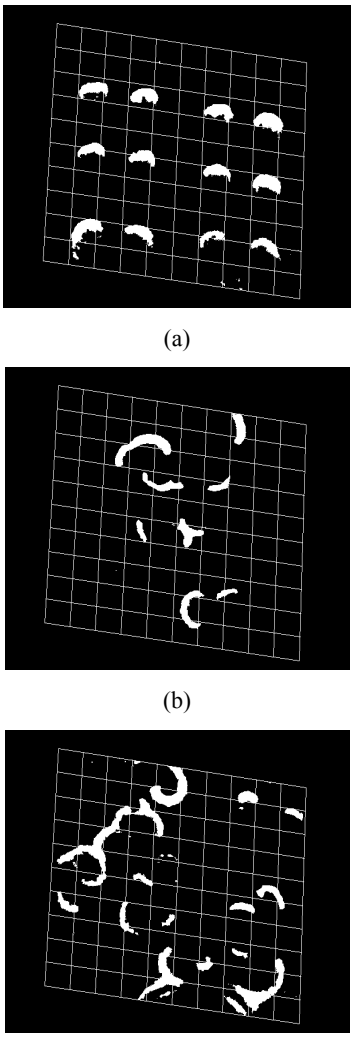

(a) (b) (c)

Figure 7. Showing an example solution to the excitation task run in the real chemistry after initiation (a), generation 12 (b) and 40 (c). State on the last of the 25 control cycles is shown.

If, for example, the excitation task were to have been tackled by a person with prior knowledge of the BZ reaction, the number of fragments would be increased simply via the predominant projection of the lowest light level. However, it appears from experimental observation that the number of fragments and thus the total excitation level is increased *also* via the projection of appropriate high and low light intensity cells placed by the CA so as to manoeuvre and split existing fragments. The exact mechanism of this process is yet to be fully understood.

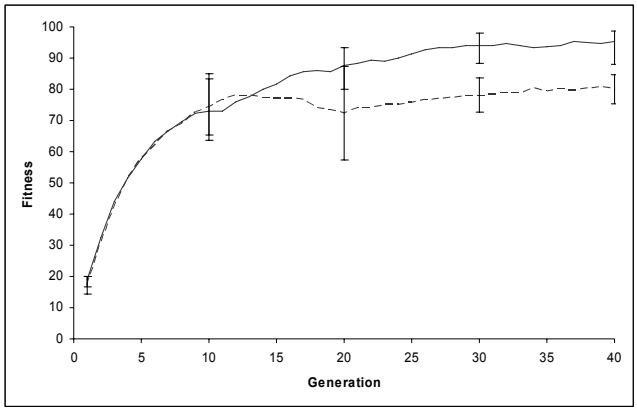

(a)

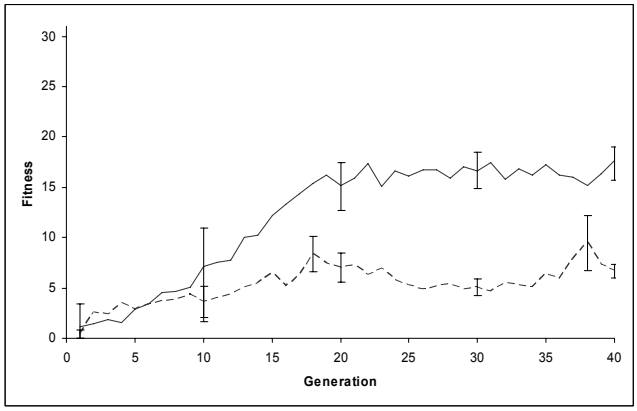

(b)

Figure 8. Showing the typical fitness over time for both the inhibition (a) and excitation tasks (b) on the real chemical system. Solid lines: breeders1000, dashed lines: breeders 0.

7. CONCLUSIONS

Excitable and oscillating chemical systems have previously been used to solve a number of computational tasks. However we suggest that the lack of compartmentalization in the majority of the previous systems limits the domain of solvable tasks thus making it difficult to realize general-purpose computing. We propose that utilizing networks of coupled oscillating chemical reactions will open chemical computing to wider domains and are interested in using light to create multiple compartments — here cells in a grid. In this paper we have presented initial results from a methodology by which to achieve the complex task of designing such systems — through the use of coevolutionary learning techniques. We have shown using both simulated and real systems that it is possible to control the behaviour of a light-sensitive BZ reaction in this way, showing fundamental control by increasing or decreasing the amount of excitation.

Current work is exploring the development of collision-based computing systems within our approach. As noted in the introduction, the fragments of excitation can, in principle, be made to behave like the moving structures — gliders — in Conway's Game of Life [8]. It has been shown that the Game of Life can support universal computation through the construction

of AND, OR and NOT gates [3] wherein streams of gliders representing binary inputs are collided together and the outcome — a glider or annihilation — interpreted as the result. We are examining the coevolution of such structures in the continuous, non-linear 2D media described above.

The use of larger gels, making use of solutions learned in simulation to seed initial populations for use on the real chemical system, and the use of populations to evolve individual cell controllers are also being considered.

8. ACKNOWLEDGMENTS

This work was support under EPSRC Grant No. GR/T11029.

9. REFERENCES

[1] Adamatzky, A. (2002)(ed) *Collision-based Computing*. Springer.

[2] Andre, D., Koza, J.R., Bennett, F.H. & Keane, M. (1999) *Genetic Programming III*. MIT Press.

[3] Berlekamp, E., Conway, J. & Guy, R. (1982) *Winning Ways for Your Mathematical Plays*. Academic Press.

[4] Budd, A., Stone, C., Masere, J., Adamatzky, A., DeLacyCostello, B. & Bull, L. (2006) Towards Machine Learning Control of Chemical Computers. In A. Adamatzky & C. Teuscher (eds) *From Utopian to Genuine Unconventional Computers*. Luniver Press, pp17-36

[5] Das, R., J, Mitchell M, & Hanson J. (1995) Evolving Globally Synchronized Cellular Automata. In L. Eshelman (ed) *Proceedings of the 6th International Conference on Genetic Algorithms*. Morgan Kaufmann, pp336-343.

[6] DeJong, K. (2005) *Evolutionary Computation*. MIT Press.

[7] Field, R. J. & Noyes, R. M. (1973) Oscillations in chemical systems. IV. Limit cycle behavior in a model of a real chemical reaction. *J. Chem. Phys*. 60: 1877-1884.

[8] Gardner, M. (1970) The Fantastic Combinations of John Conway's New Solitaire Game 'Life' *Scientific American* 223(4): 120-123.

[9] Gao, Y., Försterling, H.-D. (1995) Oscillations in the Bromomalonic Acid/Bromate System Catalyzed by $[Ru(bipy)_3]^{2+}$ *J. Phys.Chem.* 99: 8638-8644.

[10] Harding, S. & Miller, J. (2004) A Tone Discriminator in Liquid Crystal. In *Proceedings of the IEEE Congress on Evolutionary Computation*. IEEE, pp1800-1807.

[11] Holland, J.H. (1975) *Adaptation in Natural and Artificial Systems.* Ann Arbor Press.

[12] Kádár, S., Amemiya, T. & Showalter, K. (1997) Reaction Mechanism for Light Sensitivity of the $Ru(bpy)_3^{2+}$-Catalyzed Belousov-Zhabotinsky Reaction. *J. Phys.Chem.* 101: 8200.

[13] Kauffman, S.A. (1993) *On the Origins of Order*. Oxford.

[14] Koza, J.R. (1992) *Genetic Programming*. MIT Press.

[15] Krug, H.-J., Pohlmann, L. & Kuhnert, L. (1990) Analysis of the modified complete Oregonator accounting for oxygen sensitivity and photosensitivity of Belousov-Zhabotinsky systems*, J. Phys. Chem*. 94: 4862.-4866.

[16] Mitchell,M., Hraber P, & Crutchfield J. (1993) Revisiting the Edge of Chaos: Evolving Cellular Automata to Perform Computations. *Complex Systems* 7: 83-130.

[17] Mitchell, M., Crutchfield, J. & Hraber, P. (1994) Evolving Cellular Automata to Perform Computations: Mechanisms and Impediments. *Physica D* 75: 361-391

[18] Sipper, M. (1997) *Evolution of Parallel Cellular Machines*. Springer.

[19] Wang, J., Kádár, S., Jung, P. & Showalter, K. (1999) Noise driven Avalanche behavior in subexcitable media. *Physical Review Letters* 82: 855-858.

[20] Zaikin, A. N. & Zhabotinsky, A. M. (1970) Concentration wave propagation in two-dimensional liquid-phase self-oscillating system. *Nature* 225: 535-537.