

A Visual Demonstration of Convergence Properties of Cooperative Coevolution

Liviu Panait, R. Paul Wiegand*, and Sean Luke

George Mason University, Fairfax, VA 22030

lpanait@cs.gmu.edu, paul@tesseract.org, and sean@cs.gmu.edu

Abstract. We introduce a model for cooperative coevolutionary algorithms (CCEAs) using partial mixing, which allows us to compute the expected long-run convergence of such algorithms when individuals' fitness is based on the maximum payoff of some N evaluations with partners chosen at random from the other population. Using this model, we devise novel visualization mechanisms to attempt to qualitatively explain a difficult-to-conceptualize pathology in CCEAs: the tendency for them to converge to suboptimal Nash equilibria. We further demonstrate visually how increasing the size of N , or biasing the fitness to include an ideal-collaboration factor, both improve the likelihood of optimal convergence, and under which initial population configurations they are not much help.

1 Introduction

Cooperative coevolutionary algorithms (CCEAs) are coevolutionary algorithms where individuals from different populations are evaluated based on how well they perform together as a team. Commonly all individuals in a team receive the same resulting fitness (payoff). Applications of this method include optimization of inventory control systems [1], learning constructive neural networks [2], rule learning [3], and multi-agent learning [4]. The presumed advantage of CCEAs is decomposition of the search space: each of the N populations learns a different aspect (projection) of the problem, instead of one single population having to learn the joint problem at one time. Unfortunately, though each CCEA population is searching its projection of the problem at a time, that projection is constantly changing. The result is that it is easy for the algorithm to get tricked by misleading information provided by poor samples of the projected space. This leads to algorithms that tend to prefer individuals in one population that will do well with *many* individuals in the other population(s) (*relative overgeneralization*), whether or not these combinations are globally optimal [5].

We have recently examined approaches to overcome relative overgeneralization. In [4, 6] we showed that this can be countered through judicious *biasing* of how individuals are evaluated. Specifically, the fitness of an individual is a

* R. Paul Wiegand currently holds a postdoctoral position with the American Society for Engineering Education and conducts research at the Naval Research Laboratory.

weighted sum of how well it performed when teamed with individuals from the other population(s), plus how well it performed with its *ideal collaborator* (or some approximation thereof). This biasing is quite effective at eliminating relative overgeneralization, at least in the simple problems we presented and analyzed formally.

Conceptualizing the phenomenon of overgeneralization, and the effects of ideal-collaboration, is quite difficult. A goal of this paper is to provide an intuitive understanding of the issues involved. The paper begins by extending the existing theoretical framework for analyzing cooperative coevolution by replacing the *complete mixing* model for fitness assessment with one using expected maximum payoff with multiple partners. This model can be used to determine the basins of attraction for the coevolutionary search. Then we will introduce a technique for visualizing these basins of attraction in a coevolutionary search space for small numbers of dimensions. As each “dimension” is a separate genome in the search space, obviously only trivial search spaces can be visualized in this way, but it will be more than sufficient for our goal. Last, we will use this new technique to show expected outcomes with and without biasing.

2 Evolutionary Game Theory

Our approach to coevolution follows Potter’s [7] model of cooperative coevolution: each population contains individuals that represent a particular component of the problem, so that one member from each population is needed in order to assemble a complete solution. Evaluation of an individual from a particular population is performed by assembling the individual with collaborating partners from other populations. To combat noise in the evaluation process due to choice of partners, multiple evaluations may be performed. Aside from evaluation, the populations are evolved independently.

An appealing abstract mathematical model for this system comes from the biology literature: evolutionary game theory [8, 9]. EGT provides a formalism based on traditional game theory and dynamical systems techniques to analyze the limiting behaviors of interacting populations under long-term evolution. For specifics about applying EGT to the analysis of multi-population cooperative coevolutionary algorithms, see [5].

In this paper, we consider only two-population models. Expressing the quality of complete solutions through a pair of *payoff matrices*, one for each population. For this paper, we assume a symmetric model where each payoff matrix is the transpose of the other. When individuals from the first population interact with individuals from the second, a payoff matrix A is used, while individuals from the second population receive payoffs defined by A^T . We will also use an *infinite population* model. A population can be thought of not as a set of individuals, but rather as a finite-length vector \mathbf{x} of proportions, where each element in the vector is the proportion of a given *genotype* in the population. As the proportions in a valid vector must sum to one, all legal vectors make up what is commonly known as the *unit simplex*, denoted Δ^n , n here is the number of distinct genotypes

possible, $\mathbf{x} \in \Delta^n : x_i \in [0, 1], \sum_{i=1}^n x_i = 1$. The joint, two-population space of a CCEA is the product simplex $\Delta^n \times \Delta^m$.

Previous approaches to modeling CCEAs via EGT [4–6, 10] assume that an individual’s fitness is assessed as the *average* of payoffs received when in combination with *every* member of the cooperating population; this is also known as *complete mixing*. Instead we will use a more realistic fitness assessment: the maximum payoff obtained when the individual is combined N times with individuals, chosen with replacement, from the other population. Maximum has been shown to produce superior results to average in recent studies [11].

Theorem 1. *Let the payoff for individual i when teamed with individual j be a_{ij} , and $(p_j)_{j \in 1..n}$ be the probability distribution for the individuals in the population of partners for i . If the a_{ij} values are sorted in increasing order ($a_{i1} \leq a_{i2} \leq \dots \leq a_{in}$), the expected maximum payoff of i over N pairwise combinations with random partners $j_1 \dots j_N$ from the other population is*

$$\sum_{j=1}^n a_{ij} \left(\left(\sum_{k=1}^j p_k \right)^N - \left(\sum_{k=1}^{j-1} p_k \right)^N \right)$$

Proof. The expected maximum payoff is a linear combination of the actual payoff a_{ij} times the probability that it is the maximum of pairwise combinations with N random partners. As the a_{ij} values are sorted, it follows that $\left(\sum_{k=1}^j p_k \right)^N$ represents the probability of choosing all N partners from the $1 \dots j$ set. Therefore, the probability that a_{ij} is the maximum of N combinations is equal to the probability of extracting all partners from the $1 \dots j$ set, minus the probability of choosing all such partners except j (that is, $\left(\sum_{k=1}^{j-1} p_k \right)^N$). Hence, the expected maximum is $\sum_{j=1}^n a_{ij} \left(\left(\sum_{k=1}^j p_k \right)^N - \left(\sum_{k=1}^{j-1} p_k \right)^N \right)$. ■

Interestingly, the extreme setting $N = 1$ results in same fitness per individual as in the complete mixing model. As we will see later on, our extension shows improvements as the number of partners is increased in the two populations.

3 Related Work

Traditionally, EC applications of evolutionary game theory have focussed almost entirely on modeling coevolutionary algorithms that use proportionate selection and employ complete mixing. There are several exceptions, however. Ficici [12] considers several alternative selection methods, including truncation selection, (μ, λ) -ES selection, linear ranking selection, and Boltzman selection; however, this work concerns single population, competitive models of coevolution. Wiegand [5] presents two partial mixing models for cooperative coevolution, one for a complete mixing but weighted with a scheme that prefers more fit individuals

Table 1. Examples of coordination games: (a) a simple 2x2 coordination game; (b) the Climb game; (c) the Penalty game.

| | |
|----|----|
| 20 | 0 |
| 10 | 15 |

(a)

| | | |
|----|----|----|
| 21 | 0 | 10 |
| 0 | 17 | 16 |
| 10 | 10 | 15 |

(b)

| | | |
|----|----|----|
| 20 | 10 | 0 |
| 10 | 15 | 10 |
| 0 | 10 | 20 |

(c)

(from previous generations) over less fit ones; and one that models fitness assessment as the average result of partnering with a finite number of collaborators. Finally, Ming [13] analyzes CEA behaviors on two-bit landscapes using various partnering strategies, indicating that even in such a simple setting such choices can have profound effects on runtime behavior.

Visualization of basins of attraction is a common approach for understanding behaviors in many kinds of dynamical systems [14]; however, they are often restricted to two dimensions and they are not commonly employed by the evolutionary computation community. Elsewhere, visualizing dynamical system properties have proved helpful in understanding aspects of other kinds of evolutionary systems. Visualizing reverse iterates of dynamical trajectories in genetic algorithms has revealed interesting properties of locations of fixed points outside the population space one normally considers [15]. In coevolutionary systems, the aforementioned paper by Ficici [12] attempts to illustrate chaotic and periodic behaviors using cobweb plots. Recent work in the analysis of cooperative coevolutionary algorithms has demonstrated the usefulness of plotting measures of relative sizes of basins of attraction using rain-gauge methods [4, 10]. Additionally, trajectory plots in the simplex and Cartesian-product simplex spaces of evolutionary and coevolutionary systems have revealed the at-times-counterintuitive dynamics of these systems [5, 10, 8]. Finally, visualization of transitional state information in finite populations using Markov models of evolution [16] and coevolution [17] have helped illustrate the differences between the long-term dynamical systems predictions and the true algorithm behaviors.

4 Visualizing Basins of Attraction

We employ the EGT model with expected maximum fitness as described in the previous section, using fitness proportionate selection and no variational operators. We iterate the model until the proportion for one of the genotypes in each population exceeds a threshold of 0.99995, or until 50000 generations. Given the initial configuration, EGT models the coevolutionary search as a deterministic process. That is, for each initial point in the search space, we can compute to which equilibrium point it converges. As [5] shows, the populations are expected to converge to Nash equilibrium points in the payoff matrix (elements that are maximal on their row and column).




Fig. 1. Basins of attraction for the simple 2x2 coordination game, when using different numbers of partners for each of the two populations. The proportion of 1s in the first population decreases from 1.0 (top) to 0.0 (bottom) in each squared image. Similarly, the proportion of 1s in the second population decreases from 1.0 (left) to 0.0 (right) in each squared image. The basin of attraction for the (1,1) point is colored white, and the one for (2,2) is colored gray.

4.1 2x2 Coordination Games

Imagine the payoff matrix (a) in Table 1, where each population contains two kinds of genotypes: “1”s and “2”s. The higher payoffs are achieved by either pairs (1,1) and (2,2). Both of these are Nash equilibria points. A payoff of 10 is received by the pair (2,1); this increases the difficulty of the search process.

As there are only two genotypes, the state of the coevolutionary search as modeled by EGT can be described at each instant by only two proportions — p_1 and p_2 (proportions of 2s in the first and in the second population). For visualization, we consider the cartesian product of the two spaces of initial configurations (ICs) for the two populations, which is a two dimensional square from (0.0,0.0) to (1.0,1.0). p_1 increases from top to bottom, and p_2 increases from left to right. The width and height of the square are divided into 100 segments, and the center of each segment is taken as the proportion of 2s in the initial configuration.

To visualize the basins of attraction, we mark with different colors the sets of initial configurations from which the EGT model converges to the specific equilibrium point. As we have only two Nash equilibria in our example, we use white for the (1,1) equilibrium and gray for (2,2).

Figure 1 show the basins of attraction when using different numbers of partners for each of the populations. As expected, if both populations start with very high proportions of 1s, the system converges to the (1,1) equilibrium. Similarly, if both populations start with very high proportions of 2s, the system converges to the (2,2) equilibrium. The visualization shows that using more partners for

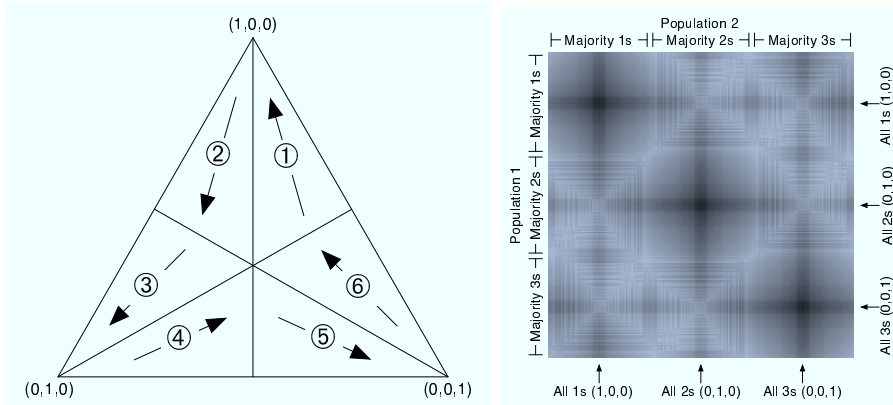


Fig. 2. (left): Our projection divides the simplex Δ^3 into six equal-area triangles; arrows shows the direction for sorting points in each area. (right): Visualization of the cartesian product of two simplices; see details in the text.

both populations leads to an increase in the size of the basin of attraction of the optimal Nash equilibrium point $(1,1)$. This result is in accordance with empirical findings reported in [11].

Additionally, gains in performance decrease with more partners. As more partners lead to longer runtimes, there may be a maximal justifiable number of partners. We plan to investigate this issue in the future.

Also, the visualization of the basins of attraction reveals an interesting artifact: the basin of attraction for the suboptimal peak tends to *grow* in the upper area of each image as we increase the number of collaborators for the second population. Thus, when the first population starts with a majority of 1s and the second population with a majority of 2s, the final results are *worse* if we provide the second population with more collaborators. One explanation: an increase in the number of partners for the second population significantly increases the expected fitness of the 2s, but does not dramatically affect the fitness of the 1s.

4.2 3x3 Coordination Games

Visualizing basins of attraction in 2x2 coordination games is relatively straightforward. In this section, we describe how the visualization can be extended to the more complicated case of 3x3 coordination games.

For a given population, possible proportions (points from the Δ^3 space) are projected onto the vertical or (for the second population) horizontal axis by first dividing the Δ^3 space into subregions, where all points in a given subregion have the same ordering of genotype proportion (1s more common than 2s, which are more common than 3s, for example). These subregions define the primary sorting for projecting points onto the axis. In this way we hope to see patterns emerging based on similarity in overall proportion.

Our projection of Δ^3 onto one dimension starts by dividing it into six equal-area triangles, as in Figure 2 (left). ICs in areas 1-2 have a majority of 1s in the population, and similarly areas 3-4 and 5-6 have majorities of 2s and 3s. The single axis is divided into six equal segments and points are assigned to the segments in the order of the triangular regions from which they come (e.g., points from area 1 are assigned to the first segment, etc.). Inside each area, ICs are ordered lexicographically in the direction of the arrow. More specifically, in regions 1-2, sorting is done primarily on p_1 , and secondarily on p_2 ; for 3-4, p_2 and p_3 ; for 5-6, p_3 and p_1 . Even-numbered regions are sorted ascending and odd-numbered regions are sorted descending. The objective of all these is to group together regions that we expect will converge to the same equilibrium. For this paper, we sample 216 ICs in the simplex: the six areas in Figure 2(left) are each divided into six triangles, and each of them is further divided into six more triangles. The center of each resulting triangle is an IC for our visualization. Our sampling does not cover ICs on the edges or vertexes of the simplex, but the probability that an EC algorithm starts from those ICs is 0 anyway.

The right image in Figure 2 is an example of the resulting projection of $(\Delta^3)^2$ onto 2-D. The sorting described above creates 2-D regions reflecting majority-1, majority-2, and majority-3 areas. Borders between those regions are the mixture of the two areas respectively. Dark lines in the figure show locations that are all 1s, 2s, and 3s in one or the other population (the vertices of the simplex).

In [4], we introduced a biased cooperative coevolution approach, and studied its application to two multi-agent coordination problems: Climb and Penalty. Both domains represent problems that associate penalties with miscoordinated actions, while providing suboptimal collaborations that avoid penalties. In this paper, we visualize basins of attraction for equilibria points in these two problems. The payoff matrices for Climb and Penalty are presented in Figure 1. Climb has two Nash equilibria, the optimum (1,1) and the suboptimum (2,2), and is strongly deceptive. Penalty has three Nash equilibria: both 1s, both 2s, and both 3s; both 2s is suboptimal, but is more forgiving if one or the other population deviates from the Nash equilibrium.

Figure 3(a) shows the basins of attraction¹ for the Climb coordination game when using different numbers of partners in each population. The images show that the “deceptiveness” of the problem domain decreases as the number of partners is increased. When using a single partner, it appears that the coevolutionary search will find the optima if at least one of the populations starts with a large number of 1s. However, as the number of partners is increased we observe that the basin of attraction for the suboptimal equilibria reduces to areas where at least one of the initial populations has a very large proportion of 2s or 3s: the more partners are used, the larger the proportion required to still converge to the sub-optimum.

¹ Some of the images, especially the one with 1 partner for each population, contain some gray dots. Those are visualization artifacts due to the threshold of 0.99995 we used for convergence.

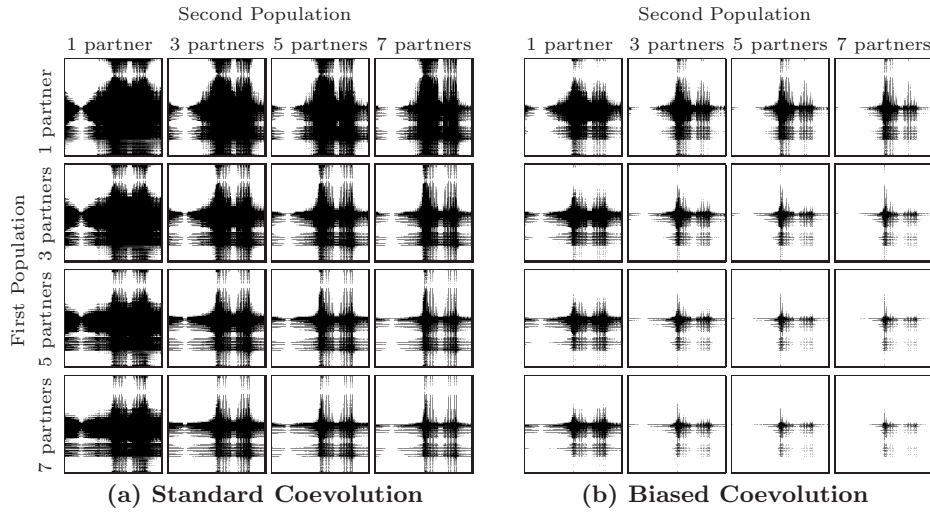


Fig. 3. Basins of attraction in Climb problem when using (a) standard and (b) biased coevolution, and 1, 3, 5 and 7 partners per population. White and black mark the basins of attraction for the (1,1) and (2,2) equilibria.

One interesting feature we note: if either population is dominated by 2s (as opposed to 3s), the system is less likely to converge to the optimum even if the other population is mostly 1s and 3s. This is due to the large, attractive basin of attraction for 2 is very large.

Figure 4(a) presents the basins of attraction for the Penalty game. We observe that the two global optima cover most of the space even when a single partner is used; the suboptimal equilibria covers mainly areas where at least one of the population started with a high percentage of 2s, and the other population has the 1s and 3s equally distributed — this increases the percentage of miscoordinations. As the number of partners is increased, the basin of attraction for the (2,2) point reduces to only areas where *both* populations start with almost only 2s. The visualization of the basins of attraction suggests that Penalty is a much easier coordination game than Climb.

We note a thin diagonal line in the 1-partner, 1-partner graph in Figure 4(a). Interestingly, this is due to the fact that if the proportion of 1s in one population is *equal* to the proportion of 3s in the other population, such combinations are heavily penalized and the system converges to the suboptimal (2,2) equilibrium.

Biased Coevolution As discussed in [4–6], much of the convergence to suboptimal solutions is due to relative overgeneralization, and one approach to dealing with this is to bias the fitness by basing it partly on the payoff of collaboration with the ideal partner for that individual (or an approximation thereof). Here, we compute the fitness of an individual as the average of the maximum payoff of N collaborations (as before), and the payoff of the individual with its ideal partner. We assume the ideal partner is known for each individual. Figures 3(b)

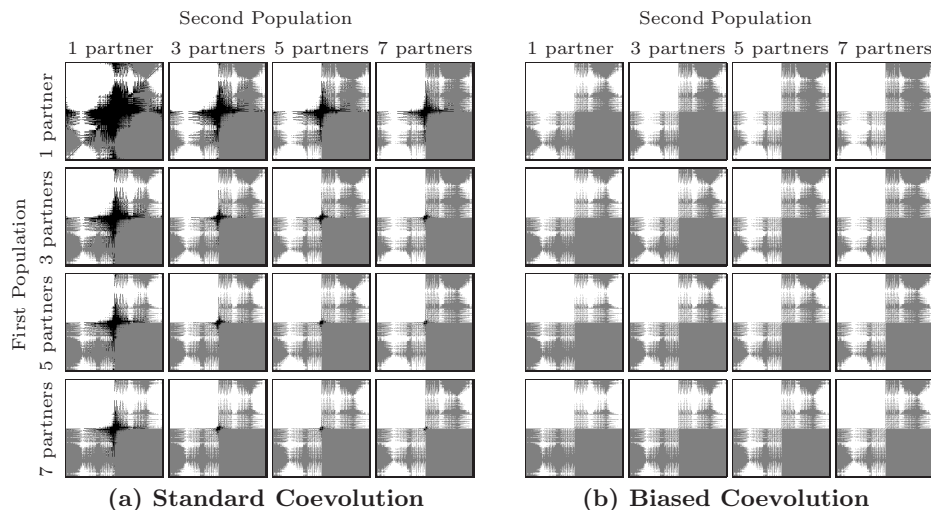


Fig. 4. Basins of attraction in Penalty problem when using (a) standard and (b) biased coevolution, and 1, 3, 5 and 7 partners per population. White, black and gray mark basins of attraction for the (1,1), (2,2) and (3,3) equilibria.

and 4(b) present the basins of attraction for the equilibrium points for the biased coevolutionary search. The images suggest that biasing further reduces the basins of attraction for suboptimal equilibria, and when biasing, increasing the number of partners helps even further. In fact in the Penalty domain, the basins of attraction for the two globally optimal equilibria cover the entire space, even with a single partner.

5 Conclusions and Future Work

In this paper we provided an evolutionary game theoretic formalism for computing the expected convergence when an individual is teamed with partners N times with replacement, and the maximum payoff is used as the individual's fitness. We then used this formalism to provide a visualization of convergence properties when teamed with multiple partners. The goal of the visualization was to demonstrate qualitatively how increases in number of partners affects the likely convergence to the globally optimal Nash equilibria, and how including collaboration with the ideal partner as part of the fitness function, as was done in [5, 4, 6], further reduces convergence to suboptima. This visualization was done straightforwardly for 2-genotype populations, and through a novel linearization of 3-D space, for 3-genotype populations.

Future research will investigate formal models for other complex collaboration schemes, and the visualization of basins of attraction for even more complicated problem domains. We believe that both the model and the visualization techniques improve our intuition of how CCEAs work, and how they can be better applied to optimization tasks.

References

1. Eriksson, R., Olsson, B.: Cooperative coevolution in inventory control optimisation. In Smith, G., Steele, N., Albrecht, R., eds.: Proceedings of the Third International Conference on Artificial Neural Networks and Genetic Algorithms, University of East Anglia, Norwich, UK, Springer (1997)
2. Potter, M., De Jong, K.: Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary Computation* **8** (2000) 1–29
3. Potter, M.A., De Jong, K.A., Grefenstette, J.J.: A coevolutionary approach to learning sequential decision rules. In: Proceedings from the Sixth International Conference on Genetic Algorithms, Morgan Kaufmann (1995) 366–372
4. Panait, L., Wiegand, R.P., Luke, S.: Improving coevolutionary search for optimal multiagent behaviors. In: Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI), Acapulco, Mexico, Morgan Kaufmann (2003) 653–658
5. Wiegand, R.P.: An Analysis of Cooperative Coevolutionary Algorithms. PhD thesis, George Mason University, Fairfax, Virginia (2004)
6. Panait, L., Wiegand, R.P., Luke, S.: A sensitivity analysis of a cooperative coevolutionary algorithm biased for optimization. In Poli, R., *et al*, eds.: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) 2004, Berlin, Germany, Springer (2004) (to appear)
7. Potter, M.: The Design and Analysis of a Computational Model of Cooperative CoEvolution. PhD thesis, George Mason University, Fairfax, Virginia (1997)
8. Maynard-Smith, J.: Evolution and the Theory of Games. Cambridge University Press (1982)
9. Hofbauer, J., Sigmund, K.: Evolutionary Games and Population Dynamics. Cambridge University Press (1998)
10. Wiegand, R.P., Liles, W., De Jong, K.: Modeling variation in cooperative coevolution using evolutionary game theory. In Poli, R., *et al*, eds.: Foundations of Genetic Algorithms (FOGA) VII, Morgan Kaufmann (2002) 231–248
11. Wiegand, R.P., Liles, W., De Jong, K.: An empirical analysis of collaboration methods in cooperative coevolutionary algorithms. In Spector, L., *et al*, eds.: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) 2001, Morgan Kaufmann (2001) 1235–1242
12. Ficici, S., Pollack, J.: Game-theoretic investigation of selection methods used in evolutionary algorithms. In Whitley, D., ed.: Proceedings of CEC 2000, IEEE Press (2000) 880–887
13. Chang, M., Ohkura, K., Ueda, K., Sugiyama, M.: Modeling coevolutionary genetic algorithms on two-bit landscapes: Partnering strategies. In Greenwood, G., *et al*, eds.: Proceedings of CEC 2004, IEEE Press (2004) (to appear)
14. Alligood, K., Sauer, T., Yorke, J.: Chaos: An Introduction to Dynamical Systems. Springer-Verlag (1996)
15. Juliany, J., Vose, M.: The genetic algorithm fractal. *Evolutionary Computation* **2** (1994) 165–180
16. Spears, W.M., De Jong, K.: Analyzing gas using markov models with semantically ordered and lumped states. In Belew, R., Vose, M., eds.: Foundations of Genetic Algorithms (FOGA) IV, Morgan Kaufmann (1996)
17. Liekens, A., Eikelder, H., Hilbers, P.: Predicting genetic drift in 2×2 games. In Poli, R., *et al*, eds.: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) 2004, Berlin, Germany, Springer (2004) (to appear)