Dynamic Properties of Minimal Algorithms for Coevolution

Pablo Funes[1] and Enrique Pujals[2]

[1] Icosystem Corporation, 10 Fawcett St., Cambridge, MA, USA
pablo@icosystem.com
[2] Instituto de Matemática Pura e Aplicada (IMPA), Rio de Janeiro, RJ, Brasil
enrique@impa.br

Abstract. One characteristic that differentiates coevolution from regular evolution is the existence of intransitivities (rock-scissors-paper). It is only recently that abstract models have begun to be used to study coevolution. *Numbers Games* in particular, have been studied by several authors as minimal models of intransitivities.

Here we carry out an analytical study of the dynamics of basic coevolutionary algorithms in the presence of intransitivities, focusing on two-dimensional numbers games. Rather than testing out different algorithms, we focus on using formal proofs.

We show that depending on the nature of the problem, the coevolutionary hill-climber C(1+1) either makes progress with constant average speed or behaves as a random walk, wondering aimlessly. Also, depending on the space, optimal elements exist but are never reached. Larger populations make the analysis more difficult and do not bring qualitative changes into those dynamics, except on specific cases.

1 Coevolutionary Games

Coevolutionary algorithms share with evolutionary ones the same setup with finite population and mutation, recombination and selection stages. Selection however, is not defined by the usual fitness function $f : \mathcal{Y} \to \mathbb{R}$, where $\mathcal{Y}$ the genotype, or "state" space[1]. Instead, in coevolution a weaker *relative fitness* function

$$f : \mathcal{Y}^{\nu} \to \mathbb{R}^{\nu}$$

defines preferences within the members $(y_1, \ldots, y_\nu)$ of the current generation.

A frequent kind of coevolutionary setup takes place when one can compare pairs of individuals. This happens, for example, in two-player games. Players A and B challenge each other and the outcome is a winner and a loser — thus one has a fitness function which compares a pair of individuals and decides on a win, lose or tie:

$$f : \mathcal{Y} \times \mathcal{Y} \to \{-1, 0, 1\}$$

such that

$$f(A,B) = \begin{cases} 1 & \text{if } B \text{ wins} \\ -1 & \text{if } B \text{ loses} \\ 0 & \text{otherwise} \end{cases}$$

Definition 1. *We shall use the name* coevolutionary game *in a (genotype) state space* $\mathcal{Y}$ *to describe any function* $f : \mathcal{Y} \times \mathcal{Y} \to \{-1, 0, 1\}$.[3]

It is usually the case that when "*A* beats *B*" we can also say that "*B* loses against *A*". This is the symmetry property:

Definition 2. *A coevolutionary game* $f : \mathcal{Y} \times \mathcal{Y} \to -1, 0, 1$ *is* symmetric *if for every pair* $(y, y') \in \mathcal{Y}$,

$$f(y, y') = -f(y', y) \tag{1}$$

Remark 1. A symmetric coevolutionary game defines a partial order in the genotype space $\mathcal{Y}$:

$$y < y' \Leftrightarrow f(y, y') = 1$$

and symmetrically, any partial order $\leq$ on $\mathcal{Y}$ defines a symmetric coevolutionary game by

$$f(y, y') = \begin{cases} 1 & \text{if } y < y' \\ -1 & \text{if } y > y' \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

In the remainder of the paper we use the f and the $\leq$ notations interchangeably.

1.1 Intransitivity

One puzzling consequence of coevolution is the possibility of having internal cycles or "intransitivities", as in the well-known *rock, paper, scissors* game,

	rock	paper	scissors	B
rock	0	1	-1	
paper	-1	0	1	
scissors	1	-1	0	
A				$f(A, B)$

Intransitivity problems were found in early works in coevolution, giving rise to the name "Red Queen Effect": the fact that individuals in the current generation are able to beat those on the previous one does not mean that they are improving in general. The landscape is changing so it is conceivable that the evolutionary algorithm, like the Red Queen of Lewis Carrol, is going nowhere fast.

Definition 3. *A coevolutionary game is* transitive *if A beats B and B beats C implies that A shall beat C:*

$$y \geq y' \wedge y' \geq y'' \implies y \geq y'' \tag{3}$$

and intransitive *if it is not transitive.*

[3] In this paper we limit ourselves to deterministic games. In many cases however, pairs of strategies define instead just the probabilities of winning or losing, therefore adding an extra layer to the problem[12,10,8].

1.2 Coevolutionary games in $\mathbb{R}^2$

There has been a recent interest in studying coevolutionary "number games" in which the genotype space is $\mathbb{R}^2$ [13,9,3,11,4,7]. Watson [13, eq. 3] proposed the *Intransitive Numbers Game* (ING),

$$(x_1,y_1) \le (x_2,y_2) \iff \begin{cases} |x_1 - x_2| \le |y_1 - y_2| \text{ and } x_1 \le x_2 \\ \text{or} \\ |x_1 - x_2| \ge |y_1 - y_2| \text{ and } y_1 \le y_2 \end{cases} \tag{4}$$

ING was a good way to start a formalized discussion on the issues around coevolution and intransitivity. This game was designed to illustrate the fact that, even though there is a straightforward way to "improve" — namely, increasing the values of both x and y — there is a catch. Two points are compared only on the basis of their most similar dimension. If they are closer on x they will be compared with respect to x and if they are closer in y they will be compared in y. There is no confusion if one player beats the other in both dimensions, but if one pair is better than the other only in one dimension, then it is the most similar that prevails. This rule leads to intransitivities, for example: $(0,0) \ge (-1,2)$ and $(-1,2) \ge (2,1)$ but $(2,1) \ge (0,0)$.

Bucci [4] studies the intransitive numbers game as well as the *focusing game* he defines, which is not symmetric:

$$f((x_1,y_1),(x_2,y_2)) = 1 \iff \begin{cases} x_2 \ge x_1 \\ \text{or} \\ y_2 \ge y_1 \end{cases} \tag{5}$$

For contrast we can define the *sum* game, a game that is both transitive and symmetric[4],

$$(x_1,y_1) < (x_2,y_2) \iff x_1 + y_1 < x_2 + y_2 \tag{6}$$

1.3 Positive Region

Coevolutionary games like the ones above are characterized by their *positive regions*.

Definition 4. *The positive region of a fixed point $y_0 \in \mathcal{Y}$ is the set of all elements in the genotype space that are better,*

$$C^+(y_0) = \{y \in \mathcal{Y} : y_0 < y\} \tag{7}$$

and conversely, $C^-(y_0) = \{y \in \mathcal{Y} : y_0 > y\}$

Positive regions for the three games above (and two others) are shown in figure 1. All of them are examples of *pie games*, where winning regions are defined by angular relationships between y_0 and y.

[4] Also in [13, eq. 2]

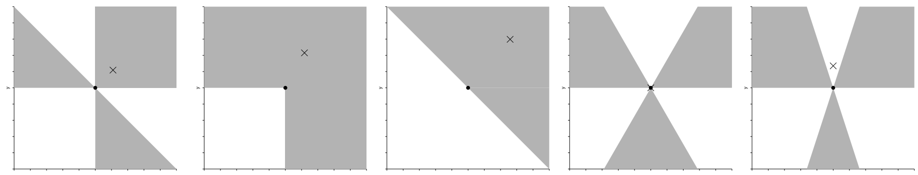

(a) (b) (c) (d) (e)

Fig. 1. Positive regions: (a) Watson's intransitive numbers game (eq. 4); (b) Bucci's *focusing game* (eq. 5); (c) s*um game*; (d) a balanced game and (e) *deceiving game* (e). A pair (x,y) at the center, is defeated by all strategies in the gray area (positive region), and in turn defeats the strategies in the white area (negative region). Game (b) is asymmetric – A beats B does not imply that B loses to A —because opposite areas do not have opposite colors. (c) is symmetric and transitive. (a) is symmetric but not transitive. The x marks the center of mass (see text).

2 C(1+1) Coevolutionary Algorithms

In order to do an analytical study of the dynamics of coevolution in the presence of intransitivities, we begin by investigating the simplest kind of coevolutionary algorithm we could think of :

Definition 5. *The* coevolutionary hillclimber $C(1+1)$ *is the following algorithm: For a coevolutionary game* $f : \mathcal{Y} \times \mathcal{Y} \to \{-1,0,1\}$

1. *Take an initial point* $y \in \mathcal{Y}$
2. *Repeat forever:*
   (a) *Generate a random mutation* y' *of* y
   (b) *Replace* y *with* y' *whenever* $y \leq y'$

We only need to consider steps in the positive regions of the game. That is, even though all mutations are equally likely to occur, only those which improve the fi tness are kept. Therefore it is useful to think only in terms of the steps taken,

Definition 6. $C^*(1+1)$ *is the following algorithm: Replace step 2(a) in def. 5 above with*
2 (a)' Generate a random mutation y' *of* y *such that* y' *beats* y

that is, in C^* we only increment the generation counter when the mutation has been successful.

All the sample games we have talked about, share the property that the positive regions for different points are congruent. This is the *uniformity* property.

Definition 7. *A coevolutionary numbers game* f *is* uniform *if for all y , y' and z,*

$$f(y,y') = f(y+z,y'+z) \tag{8}$$

Remark 2. A uniform coevolutionary game f can be defined by a one-variable fitness function $g : M_0 \to \{-1,0,1\}$.

$$g(y) = f(0,y) \tag{9}$$

where M_0 is the co-domain of the mutation operation. Throughout this paper we use random mutations in the unitary ball, so $M_0 = B_0(1) = \{(x,y) : ||(x,y)|| < 1\}$.

3 Evolution or Random Walk?

There are several ways we could try to characterize the evolutionary progress, or lack thereof, in coevolutionary algorithms.

- Victories: are the new generations increasingly likely to defeat the old ones?
- Exploration: will the evolving population wander aimlessly around the space or will it make progress in certain areas?

We thus introduce the following "undesirable properties":

Definition 8. ***Red Queen Property*** *(RQ) Let $y_0, y_1, \ldots, y_n, \ldots$ be the sequence of the (best) individuals from each generation in a coevolutionary run. The sequence has the* Red Queen Property *if an early generation k beats an infinite number of future generations:*

$$\forall n > k : P\{\exists m > n : y_k \text{ defeats } y_m\} = 1 \tag{10}$$

This definition agrees with the well-known notion, first proposed by Cliff and Miller [6], that coevolutionary progress can be measured by the increased frequency by which the best individual of the last generation beats the best individuals of earlier generations. RQ is the opposite property: no matter the number of generations, we keep coming back to a losing situation with nonvanishing frequency.

Definition 9. ***Random Walk*** *(RW) $y_0, y_1 \ldots, y_n, \ldots$ is a* random walk *if an early generation k is very close to an infinite number of future generations:*

$$\forall \varepsilon > 0 : P\{\exists m > ||y_k - y_m|| < \varepsilon\} = 1 \tag{11}$$

In fact, for all examples in this papers a stronger definition is useful[5],

Definition 10. ***Random Walk*** *(RW') $y_0, y_1 \ldots, y_n, \ldots$ is an* open random walk *when forever it keeps going to any arbitrary open set U:*

$$\forall n_0, U \, P\{\exists n \geq n_0 : y_n \in U\} = 1 \tag{12}$$

Remark 3. For uniform games in $\mathbb{R}^n$, each point starts an identical stochastic process, thus the RW and RQ properties above are valid for all $k \geq 0$ if and only if they are valid for a single k.

[5] We use the name "random walk" somewhat loosely here: random walks are usually restricted to lattices.

The number games have been defined geometrically. We intuitively consider that progress must mean movement in some direction, and this is in fact the case:

Theorem 1. *If a uniform game's negative regions are open, then* $RW \Rightarrow RQ$

Proof. If there are ever $n_1 < n_2$ such that $y_{n_1} > y_{n_2}$ then, given that y_{n_1} is inside a losing region of y_{n_2}, there is an ε such that $B_{y_{n_2}}(\varepsilon) \subset C^-(y_{n_1})$. This neighborhood of losers is visited an infinite number of times, because of the RW property.

Theorem 2. *Let* $g : B(0,1) \to \{-1,0,1\}$ *be a uniform coevolutionary game in* $\mathbb{R}^2$. *Then the algorithm C(1+1):*

1. Is RW if and only if the center of mass is zero (and RW' for games with open negative regions)
2. Moves in the direction of the center of mass with speed proportional to the norm of the center of mass .
3. Is RQ if and only if the center of mass lies outside C^+

Proof. Proof Let Y_n be a random variable representing the n-th generation element of $C*(1+1)$ and let $X_n = Y_{n+1} - Y_n$. Then $\{X_n\}_{n \in N}$ is a sequence of independent, identically distributed random variables. The average, or expected value $E(X)$ of the mutation operator is the center of mass,

$$E(X) = \int_{(x,y) \in C^+ \cap B_0(1)} (x,y)dxdy$$

By the law of large numbers we know that

$$\frac{1}{n} \sum_{i=1}^{n} (X_i - E(X_i)) \to N(0,1)$$

If $E(X) = 0$ then this means $\frac{1}{n} Y_n \to N(0,1)$, thus Y_n is not only centered around zero, but it "vibrates" around it with variance $\sqrt{n}$, thus is a random walk. If the negative regions are open, then the steps taken in fact "vibrate" in all directions, and thus the associated stochastic process visits every open set in the plane (see [5,2]). This proves part 1.

On the other hand, if $E(X)$ is nonzero, then it is $\frac{1}{n}(Y_n - nE(X)) \to N(0,1)$, the expected location after n steps is $nE(X)$ with a variance of $\sqrt{n}$ (point 2).

Finally, if the center of mass is in the interior of the winning region for (0,0), the distance between $nE(X)$ and the winning region grows linearly, but the variance only with the square root of n, therefore the probability of the process visiting the losing region approaches zero. Conversely, if the center of mass is in the losing region of zero, or the border, then the probability that Y_n is beaten by zero approaches 0 or $\frac{1}{2}$, respectively.

3.1 Comments

Theorem 2 is the main result of this paper. We have shown that the basic coevolutionary algorithm moves in a reasonable direction, with linear speed, avoiding the red queen effect, for many problems.

Well-behaved problems include the original ING and Bucci's focusing game and of course, all transitive problems. Looking back at figure 1, we can characterize the behavior of a pie game by looking at the position of x that indicates the center of mass. "Good problems" are those for which the circle is inside a gray region.

However, problems with center of mass at zero are hopeless random walks and the generations will drift aimlessly throughout the $\mathbb{R}^2$ universe. Fortunately, this can only happen in a perfectly balanced problem. This characteristic has a probability zero in the realm of all possible coevolutionary games. We shall thus say that *generically* the algorithm C(1+1) at least goes somewhere.

Problems with the center of mass in the wrong place (e.g. the "deception" problem of fig. 1e), confused by opposing intransitivities, go deep into the losing region. This is the worst kind of intransitivity, with the probability of losing vs. earlier generations approaching one.

Can we use larger populations to escape from RW/RQ behavior? Below we analyze two algorithms, C(1+N) — where instead of a random opponent it is the winner of a tournament of mutants that is chosen, and also $C_<(N+1)$ which tries to escape RQ by accepting new mutants only if they can beat all previous generations.

4 C(1+N)

The first improvement we can make to C(1+1) is to generate several mutants, instead of one, and choose the winner of a round-robin tournament as the next generation.

The result is that the new center of mass concentrates on the edge of the largest winning sector (if it exists), thus escaping the "deception" explained above, but not the RW when the center of mass is zero.

Note: the proofs in this and the next section are difficult. We have full proofs for three-sector, symmetric pie problems, but we can only fit a sketch of each in the available space.

Given a pie with sectors C_1, C_2, C_3 we define the vector $v_i = \int_{C_i}(x,y)dxdy$.

Definition 11. *We say that the property* **G** *(for* generic*) holds if there is i such that* $m(C_i) > m(C_j)$ *for any* $j \neq i$*. We named the associated vector* v_i *as* v_M*.*

$C(1+N)$ *is the following algorithm: given an initial point* Y*, repeat forever*

1. *let* $X_1,\dots,X_N$ *be N points chosen randomly in* B^+ *and take Z as the the best of* $X_1,\dots,X_N$*, meaning that* $Z = X_i$ *if and only if* $\sharp\{j : X_i > X_j\} \geq \sharp\{j : X_k > X_j\}$ *for any k and* $i < k$ *in case that* $\sharp\{j : X_i > X_j\} = \sharp\{j : X_k > X_j\}$
2. replace Y by $Y + Z$.

We define the measure m_N over as $m_N(U) = P(Z \in U)$.

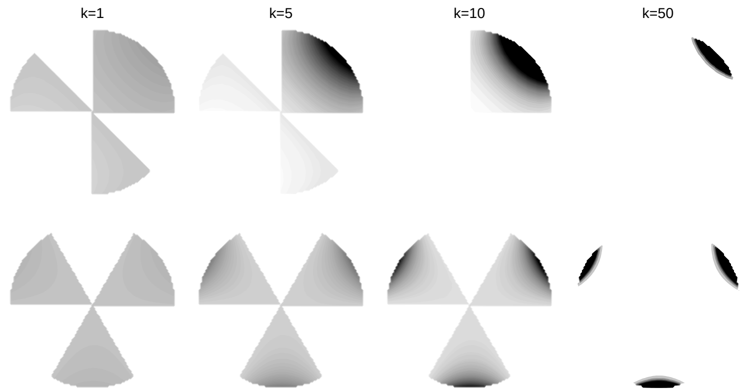

Fig. 2. Probability density of beating (0,0) plus k random opponents for increasing k in ING (A) and a balanced problem (B). Directionality concentrates for larger populations but cannot avoid the random walk problem in (B).

After a series of computations we find that,

$$m_N(U) = \Sigma_{k\geq N/2} D_k \mu(U)^k (1-\mu(U))^{N-1-k} \tag{13}$$

where $\mu(U) = \int_{(x,y)\in U} m((C^- + (x,y)) \cap B^+)dxdy$ and D_k is the number of all possible arrangement of events $\{X_1, ..., X_N\}$ with the property that the best scores exactly k points. Observe that the best has at least to score $N/2$ points (see fig. 2) . Then, we define the expectation vector related to this measure $E_N = \int_{B^+}(x,y)dm_N$.

Theorem 3. *The following results hold:*

1. *If* $C(1+1)$ *is a random walk then* $E_N = 0$ *for any* N
2. *If the property* **G** *holds then* $m_N \to \delta_{\frac{v_M}{|v_M|}}$ *and* $E_N \to \frac{v_M}{|v_M|}$, *where* δ *is the Dirac distribution concentrated on the single point* $\frac{v_M}{|v_M|}$.

Now we proceed to prove the Theorem. To get to point 1 we start with the following lemma.

Lemma 1. *Let* g *be a positive map over* $B_0(1)$ *such that* $g(v) = g(R^i_{\frac{2\pi}{3}}(v))$ *for any* $v \in B_1$*; where* $R^i_{\frac{2\pi}{3}}(v)$ *is the rotation of* v *by* $\frac{2\pi}{3}$*. Then* $\int_{B_0(1)}(x,y)g(x,y)dxdy = (0,0)$.

To finish the proof of the first part, we have to show that assuming $E_1 = (0,0)$, then density g_N associated to the measure m_N (i.e.: $m_N(U) = \int_U g_N dxdy$) verifies the hypothesis of the previous lemma. Recalling that m_N is the sum of different powers of μ and $1-\mu$ and the distribution of μ is given by $f(x,y) = m((C^- + (x,y)) \cap B^+)$, it is enough to show the hypothesis of the Lemma for $f(x,y)$. In fact, if $E_1 = (0,0)$ and having three sectors, it follows that if $(x,y) \in B^+$ then $R^i_{\frac{2\pi}{3}}(x,y) \in B^+$. Using this and the expression of f by straight calculation we conclude the proof.

For the second part of the theorem, roughly speaking, we want to show that the measure m_N is concentrated in the biggest sector. First we observe that if the property **G** holds, then the global maximum of $f(x,y) = m((C^- + (x,y)) \cap B^+)$ is equal to $\frac{v_M}{|v_M|}$ and so the measure μ^k is most concentrated over $\frac{v_M}{|v_M|}$. On the contrary, $(1-\mu(U))^{N-1-k}$ is most concentrated over the minimum of f. But since $k \geq N/2$ we get that $\mu(U)^k(1-\mu(U))^{N-1-k}$ is most concentrated over $\frac{v_M}{|v_M|}$. Moreover, as N increases, the number of terms that contribute to $\frac{v_M}{|v_M|}$ goes up. Again, the previous assertions follow from straight calculation.

5 $C_<(N+1)$

Definition 12. *$C_<(N+1)$ is the following algorithm: given a set of initial points $\{Y_1,...,Y_N\}$ such that $Y_k > Y_j$ for any $k > j$, repeat*

1. *Take $X_{N+1} \in B^+$ such that $Y_{N+1} = Y_N + X_{N+1}$ verifies that $Y_{N+1} > Y_k$ for any k;*
2. *Replace $\{Y_1,.....,Y_N\}$ with $\{Y_2,.....,Y_{N+1}\}$.*

We define the measure m_N over B^+ as $m_N(U) = P(X_{N+1} \in U)$. Moreover, observe that Y_{N+1} is obtained as $Y_1 + Z_N$ with $Z_N \in B^+(N)$ where $B^+(N) = B(N) \cap C_0^+$ and $B(N)$ is the ball of center 0 and radius N. Also, $Z_N = X_2 + X_3 + ... + X_{N+1}$ such that $X_i \in B^+$ and are chosen in B^+ with probability m_N. So, given m_N we also have a measure in $B^+(N)$ defined as $\bar{m}_N(V) = P(Z_N \in V)$

These probability measures m_N and $\bar{m}_N$ are defined by induction; assuming that we have defined m_{N-1} and $\bar{m}_{N-1}$ then

$$m_N(U) = P(X_{N+1} \in U) = \int_{z \in B^+(N-1)} m_{(N-1)}(U \cap (C_0^+ - z)) d\bar{m}_{(N-1)}(z)$$

Now, we define the expectation vector $E_N = \int_{B^+}(x,y)dm_N$.

Theorem 4. *The following results hold:*

1. *If $C(1+1)$ is a random walk then $E_N = (0,0)$ for any N;*
2. *If the property **G** holds then $m_N \to \delta_{\frac{v_M}{|v_M|}}$ and $E_N \to \frac{v_M}{|v_M|}$.*

From this theorem, we conclude that if $C(1+1)$ is a random walk, then $C_<(N+1)$ also behaves as a random walk. More precisely, since $E_N = (0,0)$ we have the same kind of Red Queen behavior. On the other hand, we know that (generically) the algorithm $C_<(N+1)$ explores the same sectors as $C(1+N)$.

To prove the theorem, we will use similar arguments as used in the case of $C(1+N)$. In fact, we use the lemma 1 after showing that the distribution associated to the measures m_N and $\bar{m}_N$ verify the hypothesis of the lemma. This is immediate for $m_2(U) = P(X_3 \in U) = \int_{(x,y) \in B^+} m_1[U \cap (C_0^+ - (x,y))]dx$. The rest follows by induction.

For the second part, again we want to show that the measures concentrate in the biggest C_j. In fact, the measures concentrated at $\frac{v_M}{v_M}$. For $N = 2$ we have to check the map $x \to m(U \cap (C_0^+ - x))$ and by symmetry it is the same situation for the measure associated with $C(1+2)$. The rest, holds by induction.

5.1 Comments

The previous theorems show that there is no chance of avoiding the random walk behavior using larger populations. However, for generic examples the previous algorithms can be extremely useful with large enough populations . Although the red queen is avoided generically, all we can achieve so far is to explores the same region. Mutations make progress only in the direction of the largest cone, leaving other secondary regions unexplored.

There is a correlation between the mass of a sector and the score that a typical point in it scores against a (large) set of random opponents. In a similar fashion as being a total winner becomes associated with being in the largest sector, a partial winner shall land in a minority region, if it exists.

We have shown that an algorithm $C(1+N,\lambda)$ that keeps players that achieve a certain ratio λ of wins vs. losses will land in a secondary area of the pie, provided there is one with corresponding mass. For instance, for $\lambda = 0$ we keep only loses and manage to move in the direction of the smallest slice of pie. This comes from the fact that each the measure associated is $D_0(1-\mu(U))^{N-1}$ (with the same symbols as eq. 13).

6 Reaching Global Optima in Closed Domains

So far we have considered coevolutionary games in unbounded domains. "Progress" is loosely defi ned in them; in ING we know intuitively that we should go up and to the right; for other games it can be more diffi cult to think of a defi nition of progress.

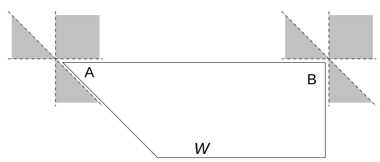

Fig. 3. Compact region W with maximal elements for ING. Corners A and B are optimal elements: a player in either of them beats everybody else. A coevolutionary hillclimber approaches B but never reaches it.

If we change the domain to a bounded, compact region, then global optima can exist. If they do, they are the target of any coevolutionary algorithm. In this section we examine the geometric conditions for the existence of maximal elements and their reachability.

Definition 13. *Consider a pie game in a compact, convex domain W. Given a point $p \in W$ we say that p is a maximal element if $(C_0^+ + p) \cap B_r(p) \cap W = \emptyset$ for some $r > 0$.*

The question we consider is: if there are maximal points, are they found by the algorithm $C(1+N)$? The problem is strongly dependent on the oddities of the boundary.

We are interested in those cases when corners are not too bizarre. In this regard we consider only maximal points (see fig. 3) that overlap the edge of the losing region associated with the corner.

Observe that a maximal element locally beats all other opponents; whether or not it is a global optimum for the problem depends on the shape of W. Moreover, not all compact domains have maximal elements; for example, if W is a disk then it has a single maximal element if and only if the problem is transitive.

We recall the concept of stationary distribution for a stochastic process. It describes the distribution of the accumulation points of sequences generated by the random process. If the stationary distribution is nonzero in all open sets, it follows that the accumulation points are spread everywhere. The condition is that there needs to be a transition probability between any pair of open sets U, V which depends only on the areas of U and V.

Theorem 5. *Let us consider the algorithm $C(1+N)$ for an intransitive (pie) game in the convex compact domain W. The stationary distribution of has support over open sets.*

The proof follows from the fact that we are in a compact set and that for any pair of open neighborhoods, there is a uniform positive probability to go from one to the other.

6.1 Comments

The above theorem implies that coevolutionary hillclimbers fail to reach the problem's optimal values, when those exist. This problem stems from the fact that, even though maximal points exist, and they are Nash equilibrium points, the main positive region disappears as we approach it. Transitive problems can only move toward the maximal, and thus reach it eventually, but intransitive ones can escape, as the alternative positive regions become proportionally larger.

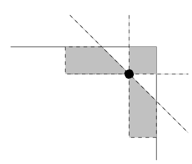

Fig. 4. Behavior of C(1+N) near a maximal point: As instances approach the maximal corner, the area of the main region disappears outside the edge and mutations escape along the secondary directions.

Mutation control strategies, like those found in Evolutionary Strategies and Simulated Annealing for example, might be able to avoid this problem by shrinking the mutation region as the border gets closer. Otherwise, we are condemned to stay around, but never reach, the maximal.

7 Conclusions and Future Work

The intransitive problems we have just analyzed might be stranger than originally thought. Watson [13] suggested that the intransitive numbers game was just a perturbation of the transitive case. We have seen here, however, that such perturbation complicates the very notion of optimization.

Although we have restricted the analysis to very particular intransitive problems ("pie problems"), this analysis is generic in the sense that any intransitive situation with real-valued parameters and continuous fi tness must end up looking like a pie — in a tiny neighborhood at least. This means that coevolution must deal with such intransitivities found along the way, by trying to get out of there in a constructive direction.

The simple C(1+1) coevolutionary hillclimber does have an answer, that works for many situations: it makes steady progress in some direction — not always the best one — increasing both x and y with constant average speed.

Using a larger population or keeping a fi nite number of previous generations around can fi x the focusing problem that arises when the solitary hillclimber ends up in the losing region, undecided between conflicting sub-goals.

The quest for an algorithm that could fi nd and explore the alternative directions of these problems is a harder one. Bucci [4, fi g. 3] proposes a form of a Pareto hillclimber that seems to spread out, exploring several interesting directions.We shall discuss Pareto optimization in an up-coming study.

Technically, in intransitive problems the Pareto front is everywhere, so the population could end up being a Brownian dust. An alternative could be the use of combinations of different winning rates: the highest one detects the main direction, and lower rates follow secondary ones.

References

1. H. G. Beyer. *The theory of evolution strategies*. Springer Verlag, 2001.
2. L. Breiman. *Probability*. Society for Industrial and Applied Mathematics, 1992.
3. A. Bucci and J.B. Pollack. A mathematical framework for the study of coevolution (prepublication). <I>Foundations of Genetic Algorithms 7</I>. Proceedings of FOGA VII, Torremolinos Spain, 4-6 September 2002 (prepublication version – to appear spring 2003), 2002.
4. Anthony Bucci and Jordan Pollack. Focusing versus intransitivity: geometrical aspects of coevolution. In et. al. Cantú-Paz, E., editor, *Genetic and Evolutionary Computation - GECCO 2003*, number 2724 in LNCS, Berlin, 2003. Springer-Verlag.
5. Chung and Fuchs. On the distribution of values of sums of random variables. *Mem. of the American Mathematical Society*, 6, 1951.
6. D. Cliff and G. Miller. Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In F. Morán, A. Moreno, J. J. Merelo, and P. Chacón, editors, *Advances in Artifi cial Life: Third European Conference on Artifi cial Life*, number 929 in Lecture Notes in Computer Science, pages 200–218, Berlin, New York, 1995. Springer.
7. Edwin D. De Jong and Jordan B. Pollack. Learning the ideal evaluation function. <i>Proceedings of GECCO 2003, pp. 277-288</i>, 2003.
8. Arpad E. Elo. *The Rating of Chessplayers, Past and Present*. Arco Pub., New York, 2nd edition, 1986.

9. Sevan G. Ficici and Jordan B. Pollack. Game theory and the simple coevolutionary algorithm: Some preliminary results on fitness sharing. <i>GECCO 2001 Workshop on Coevolution: Turning Adaptative Algorithms upon Themselves</i>, 2001.
10. Pablo Funes and Jordan B. Pollack. Measuring progress in coevolutionary competition. In J. Meyer, A. Berthoz, D. Floreano, H. Roitblat, and S. W. Wilson, editors, *From Animals to Animats 6*, pages 450–459. MIT Press, Cambridge Mass, London, 2000.
11. l. panait and s luke. A comparison of two competitive fitness functions. In GECCO-2002: Proceedings of the Genetic and Evolutionary Computation Conference. W. B. Langdon et al, eds. Morgan Kauffman. 503-511., 2002.
12. J. B. Pollack and A. D. Blair. Coevolution in the successful learning of backgammon strategy. *Machine Learning*, 32:225–240, 1998.
13. Richard Watson and Jordan B. Pollack. Coevolutionary dynamics in a minimal substrate. In . et al Spector, L, editor, *Proceedings of the 2001 Genetic and Evolutionary Computation Conference*. Morgan Kaufmann, 2001.