ON RELIABILITY OF SIMULATIONS OF COMPLEX CO-EVOLUTIONARY PROCESSES

Peter Tiňo
School of Computer Science
The University of Birmingham
Edgbaston, Birmingham B15 2TT, UK
Email: P.Tino@cs.bham.ac.uk

Siang Yew Chong
School of Computer Science
The University of Nottingham, Malaysia Campus
Jalan Broga, Semenyih 43500, Malaysia
Email: Siang-Yew.Chong@nottingham.edu.my

Xin Yao
School of Computer Science
The University of Birmingham
Edgbaston, Birmingham B15 2TT, UK
Email: X.Yao@cs.bham.ac.uk

KEYWORDS

Evolutionary game theory, evolutionary computation, co-evolution, shadowing lemma.

ABSTRACT

Infinite population models of co-evolutionary dynamics are useful mathematical constructs hinting at the possibility of a wide variety of possible dynamical regimes - from simple attractive fixed point behavior, periodic orbits to complex chaotic dynamics. We propose to use the framework of shadowing lemma to link such mathematical constructs to large finite population computer simulations. We also investigate whether the imposition of finite precision computer arithmetic or the requirement that population ratios be rational numbers does not leave the infinite population constructs and theories irrelevant. We argue that if the co-evolutionary system possesses the shadowing property the infinite population constructs can still be relevant. We study two examples of hawk-dove game with Boltzmann and (μ, λ) selection. Whereas for Boltzmann selection there is a strong indication of the shadowing property, there is no shadowing in the case of (μ, λ) selection.

INTRODUCTION

Games have been used to model some important complex real-world problems in politics, economics, biology (Axelrod, 1984) and engineering (Nisan et al., 2007). They capture intrinsic properties of such problems through the specification of rules that constrain strategies to certain behaviors (legal moves), goals for strategies to meet (to win the game), and rewards under finite resources (payoffs). However, games that successfully abstract real-world problems are often analytically intractable. In such cases, alternative heuristic approaches, such as evolutionary algorithms (EAs), are used to solve these problems in practice.

Games are particularly important in the development of a class of EAs known as co-evolutionary algorithms (CEAs) (Chellapilla and Fogel, 1999). Unlike classical EAs that require an *absolute* quality measurement of solutions to guide the population-based, stochastic search process, in CEAs, the solution quality can be estimated only with respect to its performance against a (usually) small sample of test cases. In cases where an absolute quality measurement is not available, CEAs can still solve the problem by making use of some form of *strategic* interactions between competing solutions in the population to guide the search (Chong et al., 2008).

Despite early success of CEAs in solving games, there are well-documented failures leading to poor performance of CEAs under certain conditions. One example is the *overspecialization* of evolved strategies that performs well only against specific opponents rather than a large number of different opponents (Darwen and Yao, 2002; Chong et al., 2008). As such, there are strong motivations for more in-depth theoretical studies on CEAs to understand precisely how and why CEAs can solve the problem of games.

One theoretical framework that is naturally suited for analysis of CEAs is evolutionary game theory (EGT) (Smith, 1982). The classical game theory setting involves a rational individual (player) that has to choose between different strategies, one that maximizes its payoff when interacting against another player in a game (which in turn, also maximizes its own payoff). In contrast, the EGT setting involves an infinitely large population of players that are allowed to use a set of predefined strategies. These strategies are *inheritable* and all players compete for payoffs that decide their average reproductive success (Sigmund and Nowak, 1999). Different constructions (e.g., different games, different selection mechanisms etc.) will lead to different frequency-dependent population dynamics (Hofbauer and Sigmund, 2003a). As such, EGT provides the means with which one can study precisely the conditions that affect the outcome or success of some strategies over others in the population under evolutionary process.

However, there are few studies that apply EGT in the

Proceedings 24th European Conference on Modelling and Simulation ©ECMS Andrzej Bargiela, Sayed Azam Ali David Crowley, Eugène J.H. Kerckhoffs (Editors)
ISBN: 978-0-9564944-0-5 / ISBN: 978-0-9564944-1-2 (CD)

analysis of CEAs. For example, a simple EGT setting of the hawk-dove game that involves interactions between two strategies has been used to investigate the evolutionary process of CEAs under various conditions: the study in (Fogel et al., 1997) investigated the impact of finite population, while the study in (Ficici et al., 2005) investigated the impact of selection mechanisms.

While evolutionary dynamics typically involves rather simple dynamical scenarios, co-evolution can lead to a rich variety of dynamical behaviors, including chaos (see, e.g., (Ficici et al., 2005)). Such complex dynamics are usually (and conveniently) explored under the assumption of infinite populations, where the entities of interest are, e.g., the ratios of individuals adopting a particular strategy. Given that for some game settings, computer generated orbits of co-evolutionary dynamics, under the assumption of infinite populations, reveal intricate dynamical patterns (including chaotic dynamics), we ask two important questions regarding the information content in such simulations:

1. How informative are the observed chaotic trajectories, given that the computer precision is limited? In chaotic dynamics, nearby trajectories get locally exponentially separated, so round-off errors will inevitably lead to trajectories very different from the 'ideal' infinite population ones the equations are supposed to describe.

2. Often, we are ultimately interested in dynamical behaviors of potentially large, but finite populations, using infinite population formulations as convenient conceptual constructs. How informative are then the complex infinite population trajectories about the dynamics of large, but finite populations modeled on the computer?

We propose to address these questions in the context of shadowing lemma (see e.g. (Katok and Hasselblatt, 1995)). For simplicity of presentation, we only will consider co-evolution of games with two pure strategies. We consider in more detail some settings of the co-evolutionary dynamics of the two-strategy hawk-dove game with (μ, λ) and Boltzmann selection mechanisms. We will show that there is indeed a possibility of quite intricate infinite population co-evolutionary dynamics. We will also study whether such complex infinite population dynamical patterns can have direct relevance for large scale finite population computer simulations.

EVOLUTIONARY GAME THEORY

We first describe the standard EGT framework that makes the following assumptions:

1. An infinitely large population of players, each of which has a finite set of *pure strategies* to choose from.

2. *Complete mixing* - every player interacts with all players in the population. Each player accumulates payoff depending on the outcome of the games.

3. Players reproduce in proportion to their cumulative payoffs. Reproduction is asexual and without variation, i.e., players generate clones as their offspring.

In this paper, we consider a simple EGT setting similar to that of (Ficici et al., 2005) with a game involving two players (for illustration). Each player has a finite set of pure strategies to choose from, i.e., S_i for the first player and S_j for the second player. The game is *symmetric*, i.e., both players have the same set of strategies to choose from ($S_i = S_j$). The payoff (game outcome) for the first player is g_{ij} when the player chooses strategy $i \in S_i$ while the opponent chooses $j \in S_j$ (the payoff for the second player is g_{ji}).

For the case where there are only two *pure strategies* (i.e., $\{X, Y\} \in S$), the payoff matrix for the first player (row) in a game against the second player (column) can be constructed as follows:

	X	Y
X	a	b
Y	c	d

(1)

where each entry gives the respective payoff for the chosen pair of strategies. For example, the first player receives the payoff b when it chooses strategy X while its opponent chooses Y.

Each player in the population chooses only one of the two pure strategies. Let p be the proportion of players in the population choosing X. $1 - p$ is the proportion of players in the population choosing the other pure strategy Y. We can compute the cumulative payoffs for both strategies, w_X and w_Y, in the form of a pair of linear equations:

$$\begin{aligned} w_X &= ap + b(1-p) \\ w_Y &= cp + d(1-p). \end{aligned} \tag{2}$$

We consider games with the payoff structure satisfying $a < c$ and $b > d$. For any such game, there is one population state known as *polymorphic equilibrium*

$$p_{EQ} = \frac{d - b}{a - c + d - b} \tag{3}$$

in which the cumulative scores for both strategies are the same ($w_X = w_Y$) (Ficici et al., 2005). Interpreting the population as a mixture of strategies s, (i.e., uses pure strategies X and Y with probability $p_X = p$ and $p_Y = 1 - p$, respectively), the state p_{EQ} is a *Nash equilibrium*, whereby the mixture of strategies s is its own *best reply* (i.e., if a player uses s, the highest payoff that the opponent can obtain is when it also uses s). When both players use s, they are in Nash equilibrium since neither has the incentive to deviate unilaterally to use other strategies.

Hawk-Dove Game

Although there are many games that satisfy the constraints of $a < c$ and $b > d$, in this paper we consider the classical game setting of the hawk-dove game. The setting involves interactions of two distinct behaviors (pure strategies), *hawk* and *dove*, competing for gains (G) upon winning under the costs (C) of injury.

Hawks are aggressive and two hawks will fight until one retreats with an injury. Interactions between hawks would lead to the payoff expectation of $(G-C)/2$ given a probability of $1/2$ for injury. Doves, in contrast, will avoid a fight and perform threatening postures until both retreats without injury. In such a case, they share the gain with a payoff of $G/2$. Any interaction between a hawk and a dove will lead to the dove retreating immediately. The hawk will take the full gain (payoff of G) while the dove has zero gain (payoff is 0), with no cost on injury incurred to both parties.

The payoffs for hawk-dove interactions are summarized in the following payoff matrix:

	Hawk	Dove
Hawk	$(G-C)/2$	G
Dove	0	$G/2$

(4)

When the cost of injury is greater than the gain in winning the game (i.e., $G < C$), the hawk-dove game satisfy the constraints of $a < c$ and $b > d$. The polymorphic equilibrium can be easily obtained from Equation 3: $p_{EQ} = G/C$. Note that adding a constant value w to each entry in the payoff matrix (4) ensures that the cumulative payoffs for strategies are non-negative without changing the nature of the population dynamics (Ficici et al., 2005).

Replication and Selection Pressure

In the classical EGT, the population dynamics are described by the *replicator equation* that governs how the *frequency* (proportion) of strategies in the population changes with time t of the evolutionary process. For the case of the game specified by the payoff matrix (1), the replicator equation under a selection mechanism based on the proportion of cumulative payoffs is given by:

$$f(p) = \frac{pw_X}{pw_X + (1-p)w_Y}$$

where $f(p)$ is the frequency of strategy X in the population in the following generation $(t+1)$, given that its frequency in the current population is p.

In addition to the classical fitness-proportional selection, there is a variety of alternative selection mechanisms. In this paper, we consider two well known selection mechanisms: (μ, λ) and Boltzmann selection.

The (μ, λ) **selection** is usually associated with the selection operator that is used in a class of EAs known as *evolution strategies*. Each of μ parents generates k offspring, which results in $\lambda = k\mu$ offspring. In the context of an EGT setting of CEAs operating with infinite population, we are only concerned with the ratio $\gamma = \mu/\lambda$. The replicator equation can then be obtained as (Ficici et al., 2005):

$$f(p) = \begin{cases} 1 & \text{if} \quad p < p_{EQ} \text{ and } p \geq \gamma, \\ p/\gamma & \text{if} \quad p < p_{EQ} \text{ and } p < \gamma, \\ 1 + (p-1)/\gamma & \text{if} \quad p > p_{EQ} \text{ and } p > 1-\gamma, \\ 0 & \text{if} \quad p > p_{EQ} \text{ and } p \leq 1-\gamma, \\ p_{EQ} & \text{if} \quad p = p_{EQ}. \end{cases} \tag{5}$$

In **Boltzmann selection** the fitness that determines the selection of an individual is scaled as $w_{scaled} = e^{\beta w}$, where w is the original fitness. The *inverse temperature* parameter $\beta = 1/T$ determines the selection pressure. Selection of an individual is proportional to the scaled fitness w_{scaled}. For a CEA with an infinite population, the replicator equation with Boltzmann selection is given by (Hofbauer and Sigmund, 2003b; Ficici et al., 2005):

$$f(p) = \frac{pe^{\beta(pa-pb+b)}}{pe^{\beta(pa-pb+d)} - pe^{\beta(pc-pd+d)} + e^{\beta(pc-pd+d)}}. \tag{6}$$

THE SHADOWING PROPERTY

Given a game and a particular selection mechanism (that together imply the corresponding map f) and an initial condition p_0, the map f generates an orbit $p_n = f(p_{n-1})$, $n = 1, 2, ...$ If instead of the true iterands p_n we observed p_n corrupted by a bounded noise (for example due to rounding errors or finite population effects), but still used the dynamics f, we would obtain a *pseudo-trajectory* $\{\tilde{p}_n\}_{n \geq o}$,

$$|\tilde{p}_0 - p_0| < \delta, \;\; |f(\tilde{p}_{n-1}) - \tilde{p}_n| < \delta, \;\; n \geq 1,$$

where $\delta > 0$ is the range size of the bounded noise. Such a pseudo-trajectory is often referred to a δ-pseudo-trajectory.

Given an $\epsilon > 0$, we say that a trajectory $\{q_n\}_{n\geq 0}$ ϵ-shadows another trajectory $\{g_n\}_{n\geq 0}$, if $\{q_n\}_{n\geq 0}$ stays within the ϵ-tube around $\{g_n\}_{n\geq 0}$:

$$|g_n - q_n| < \epsilon, \quad n \geq 0.$$

The Shadowing lemma tells us that (remarkably) even for the most complex and locally exploding chaotic maps, under some circumstances, the corrupted pseudo-trajectories are informative: For *any* $\epsilon > 0$, there exists a $\delta > 0$, such that for *every* δ-pseudo-trajectory $\{\tilde{p}_n\}_{n\geq 0}$ there is a true (uncorrupted) trajectory $\{q_n\}_{n\geq 0}$ under f that ϵ-shadows the pseudo-trajectory $\{\tilde{p}_n\}_{n\geq 0}$:

$$|\tilde{p}_n - q_n| < \epsilon, \;\; q_{n+1} = f(q_n), \;\; n \geq 0.$$

Hence, even though one would be tempted to assume that, under chaos, trajectories that are disrupted at every point by a bounded noise cannot possibly represent anything real, in fact such trajectories can be closely shadowed by a true trajectory.

What are the conditions that guarantee the shadowing property? Virtually all studies of the shadowing property have been performed in the framework of continuous and smooth dynamical systems. If a system is (uniformly) hyperbolic on an invariant set Ω, i.e., (loosely speaking) at each point $p \in \Omega$ the (linearized) system has only local contracting and expanding subspaces that get consistently translated by f into the local contracting and expanding subspaces at $f(p)$, then the system will have the shadowing property (Hayesand and Jackson, 2005). In general, establishing that a dynamical system is hyperbolic can be rather difficult, but it can be shown that a smooth one-dimensional system acting on Ω is hyperbolic iff for all $p \in \Omega$, there exists an $n(p) > 0$, such that the derivative $(f^{n(p)})'(p)$ of $f^{n(p)}$ at p is (in absolute value) larger than 1. Setting $p_1 = p$ and $p_n = f(p_{n-1})$, $n = 2, 3, .., n(p) - 1$, (by chain rule) this translates into the requirement that

$$|\prod_{n=1}^{n(p)} f'(p_n)| > 1.$$

In what follows we concentrate on co-evolutionary dynamics of the hawk-dove game with Boltzmann and (μ, λ) selections. Through representative case studies, we first show that such infinite population systems are indeed inherently capable of generating very complex dynamical scenarios. We then discuss whether studying infinite population models of such co-evolutionary systems can be informative about the dynamics one would observe in large population size simulations.

In this contribution we use the framework of shadowing lemma for two purposes:

1. *To investigate whether the complex simulated co-evolutionary trajectories under the infinite population assumption represent anything real.* The computer arithmetic operates with finite precision and can only yield pseudo-trajectories that cannot be *a-priori* guaranteed to represent any true trajectory of the given complex system. If the co-evolutionary system has the shadowing property then one can be assured that the observed pseudo-trajectories are shadowed by true ones generated by the underlying system.

2. *To investigate whether the complex dynamical co-evolutionary patterns under infinite populations indicate complex dynamics in large finite population computer simulations.* For large population size, the effects of finite population size on the strategy ratios p can be considered bounded noise. Furthermore, the larger the population, the smaller the range size of the noise. In a system with shadowing property, finite population pseudo-trajectories are shadowed by the complex trajectories from the original infinite population model.

Shadowing under Boltzmann selection

As a case study, we use the following setting of the hawk-dove game and Boltzmann selection that can lead to complex chaotic population dynamics: $G = 7$, $C = 12$, $\beta > 3.8$ (Ficici et al., 2005).

To show that for a range of selection pressures $\beta = 1/T$, the hawk-dove game with Boltzmann selection leads to complex dynamics, we estimated the topological entropy h_0 of the system (Balmforth et al., 1994). Topological entropy quantifies the exponential growth rate in the number of periodic points of the system. In particular, for each continuous, piece-wise monotone map f (such as the replicator map under Boltzmann selection (6)), it holds $h_0(f) \leq \kappa(f)$, where $h_0(f)$ is the topological entropy of f and $\kappa(f)$ is the exponential growth rate of the number of periodic orbits of f (Katok and Mezhirov, 1998). Hence, for complex systems with positive topological entropy (calculated with log base 2), the number of periodic points of period m grows as $\approx 2^{h_0 m}$. For selection pressures given by $\beta = 3.8$, 5.0, 5.7 and 6.5 the estimated topological entropies were $h_0 = 0.281$, 0.842, 0.905 and 0.949, respectively, indicating intricate dynamical organization of the infinite population system at higher selection pressures. But are the computer generated trajectories informative about the true co-evolutionary dynamics of this game under the Boltzmann selection? Furthermore, do such trajectories tell us anything about the dynamical complexity one might expect when running simulations with large finite populations? To address these questions we turn to the shadowing lemma.

The system (6) is a smooth dynamical system over the unit interval $[0, 1]$. For higher selection pressures the map f has two critical points where the derivative vanishes:

$$c_i = \frac{1}{2} + (-1)^i \sqrt{\frac{1}{4} - \frac{2}{\beta C}}, \quad i = 1, 2.$$

Construct the set of all pre-images of the critical points under f;

$$\mathcal{E} = \{p \in [0, 1] \mid f^n(p) \in \{c_1, c_2\} \text{ for some } n \geq 0\}.$$

If the f-invariant set $\Omega = (0, 1) \setminus \mathcal{E}$ is a hyperbolic set under f, then the system (6) possesses the shadowing property on Ω. Showing analytically that Ω is hyperbolic is very difficult due to the (rather involved) nonlinearities in f. Nevertheless, we can obtain at least an indication of hyperbolicity of Ω by numerically checking whether for each p from a fine grid of points, there exists a natural number $n(p)$ such that $|(f^{n(p)})'(p)| > 1$. In Figure 1 we show the empirically determined fold numbers $n(p)$ for $G = 7$, $C = 12$ and $\beta = 5.0$. The figure indicates that f is hyperbolic (and so has the shadowing property) on Ω. Of course, such indication should be taken with a pinch of salt due to possible numerical inaccuracies in the neighborhood of critical points.

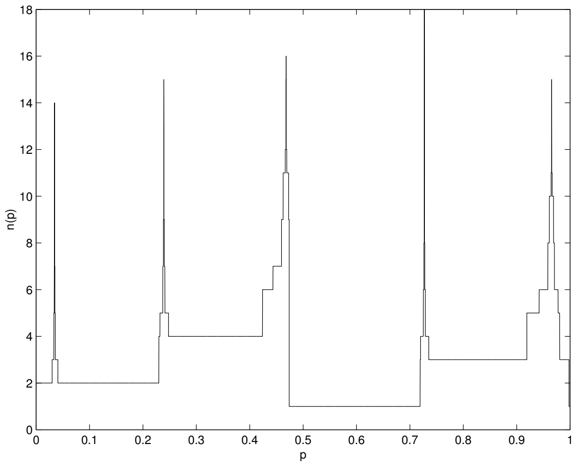

Figure 1: Hawk-dove game with $G = 7$, $C = 12$ and selection pressure $\beta = 5.0$. Shown is the fold number $n(p)$, such that $|(f^{n(p)})'(p)| > 1$, as a function of p.

Shadowing under (μ, λ) selection

Consider the hawk-dove game with the cost of injury twice the gain in winning, $C = 2G$, and a (μ, λ) selection with $\lambda = 2\mu$ (each of μ parents generates two offspring). The replicator equation (5) reads:

$$f(p) = \begin{cases} 2p & \text{if} \quad p \in [0, 1/2) \\ 1 + 2(p-1) & \text{if} \quad p \in (1/2, 1] \\ 1/2 & \text{if} \quad p = p_{EQ} = 1/2. \end{cases} \tag{7}$$

The map f acts as left-shift on binary representations of $p \in [0, 1/2) \cup (1/2, 1)$. As such, it can generate a wide variety of dynamical behaviors, dictated by the distribution of digits in the binary expansion $[p]_2$ of the initial condition p. The possible dynamics include periodic orbits of arbitrary periods, a-periodic and 'chaotic' orbits (arising from irrational initial conditions).

The map (7) is discontinuous. There is virtually no literature on shadowing property in discontinuous systems. Some discontinuous systems have the shadowing property, while others do not. It turns out that the difference between the well studied map $r(p) = 2p \bmod 1$ in chaotic dynamics and the map $f(p)$ in (7) - the existence of the additional equilibrium $p_{EQ} = 1/2$ for f - is crucial. Since $r(p)$ is a smooth expanding map on a smooth manifold (unit circle) (Katok and Mezhirov, 1998), it has the shadowing property. However, as we show below, this is not the case for the system (7).

For some small $\delta > 0$, consider a δ-pseudo-trajectory that gets within the δ neighborhood of $p_{EQ} = 1/2$. Then, for *arbitrary* number of time steps $m \geq 1$ the pseudo-trajectory can stay in p_{EQ}. After that, the trajectory can continue with a high, or low value of p, depending on to which side of $p_{EQ} = 1/2$ it shifts. More formally, consider a pseudo-trajectory $....\tilde{p}_{n-2}\tilde{p}_{n-1}\tilde{p}_n$ with

$$|f(\tilde{p}_n) - p_{EQ}| < \delta.$$

We can then set $\tilde{p}_{n+1} = p_{EQ}$ and let the pseudo-trajectory stay in p_{EQ} for m time steps: $\tilde{p}_{n+1} = \tilde{p}_{n+2} = ... = \tilde{p}_{n+m} = p_{EQ}$. The next element of the pseudo-trajectory will be

$$\tilde{p}_{n+m+1} = f(p_{EQ}) + \nu = p_{EQ} + \nu,$$

where $|\nu| < \delta$ and $\nu \neq 0$.

It is impossible to closely shadow such pseudo-trajectories with real trajectories of (7). Given a small $\epsilon > 0$, for a real trajectory to ϵ-shadow the equilibrium p_{EQ} for an arbitrarily large number of steps m, it must oscillate around p_{EQ} within the ϵ-neighborhood of p_{EQ}. Since the action of $f(p)$ is the left shift of the binary representation $[p]_2$ of p, this is not possible for arbitrarily small ϵ. In fact, this is only possible for $\epsilon \geq 1/6$, e.g., the orbit is given by the initial condition 0.101010... and oscillates between $[2/3]_2 = 0.101010...$ and $[1/3]_2 = 0.010101....$

It follows, that the general shadowing property does not hold for this setting of the game and (μ, λ) selection. However, on a reduced set of initial conditions, one can argue that a 'shadowing-like' property holds at least for a particular interpretation of the computer round-off error (Sharkovsky and Chua, 1993): If the computer can guarantee only M exact binary digits, then iterative application of f on $0.x_1x_2x_3...$ will lead to the pseudo-trajectory $\mathcal{O}$: $0.x_1x_2x_3...x_{M-1}x_M$, $0.x_2x_3...x_{M-1}x_My_1$, $0.x_3...x_{M-1}x_My_1y_2$, ... $0.y_1y_2...y_M$, $0.y_2...y_{M+1}$, ..., for some $y_j \in \{0, 1\}$, $j \geq 1$. Consider now the set $\mathcal{E}$ of all pre-images under f of $p_{EQ} = 1/2$,

$$\mathcal{E} = \{p \in [0, 1] \mid f^n(p) = p_{EQ} \text{ for some } n \geq 0\}.$$

The real trajectory of (7) starting in $0.x_2x_3...x_{M-1}x_My_1y_2...$ from the f-invariant set $\Omega = (0, 1) \setminus \mathcal{E}$ will ϵ-shadow the pseudo-trajectory $\mathcal{O}$ with $\epsilon = 2^{-M+1}$.

The set $\mathcal{E}$ contains all $p \in [0, 1]$ whose binary expansion $[p]_2$ contains any finite word over the alphabet $\{0, 1\}$ (including the empty word), followed by digit 1, followed by the right-infinite sequence of 0's:

$$\mathcal{E} = \{p \in [0, 1] \mid [p]_2 = .\{0, 1\}^*1000...\}.$$

The set $\mathcal{E}$ is dense in $[0, 1]$, since for any $p \in [0, 1]$ and arbitrarily small $\epsilon > 0$, there will be a $q \in [0, 1]$, such that $|p - q| < \epsilon$ and $[q]_2$ has an infinite tail of 0's. Analogously, it is easy to show that the set Ω is dense in $[0, 1]$ as well. However, Ω is much larger than $\mathcal{E}$ - in fact, while $\mathcal{E}$ is countable, Ω is uncountable since it contains infinite expansion rational numbers and all irrational numbers in $[0, 1]$. Under f, the set Ω contains seeds for a wide variety of dynamical regimes, including periodic and 'chaotic' orbits.

DISCUSSION AND CONCLUSION

Infinite population models of co-evolutionary dynamics are useful mathematical constructs hinting at the possibility of a wide variety of possible dynamical regimes

- from simple attractive fixed point behavior, periodic orbits to complex chaotic dynamics. We have used the framework of shadowing lemma (form the theory of complex dynamical systems) to link such mathematical constructs to large finite population computer simulations. It has also been investigated whether the imposition of finite precision computer arithmetic does not leave the infinite population constructs and theories irrelevant.

If the co-evolutionary system possesses the shadowing property, finite population pseudo-trajectories from large scale computer simulations are shadowed by trajectories from the original infinite population model. This can be used, e.g., to calibrate time consuming large population computer simulations into regimes of interesting complex dynamics revealed by the infinite population mathematical models. The effects of finite population size on co-evolutionary dynamics under stochastic replicator dynamics was studied in (Ficici and Pollack, 2007). The focus of our approach is different. We ask - can one expect that as the population size grows the complex dynamical features of infinite population co-evolutionary dynamics will shadow features observed in large scale finite population simulations?

As examples, we have dealt with some settings of the two-strategy hawk-dove game with (μ, λ) and Boltzmann selection mechanisms. It has been shown there is indeed a possibility of quite intricate infinite population co-evolutionary dynamics and that this dynamics can have direct relevance for large scale finite population computer simulations under the smooth dynamics driven by Boltzmann selection. However, the discontinuous nature of the transition map under the (μ, λ) selection can prevent the co-evolutionary dynamics to posses the shadowing property. Hence, one cannot expect the infinite population trajectories to have direct relevance for large-population simulations. On the other hand, it is not all bad news: the computer generated trajectories of *infinite* population models can be shadowed by the real *infinite* population dynamics.

To the best of our knowledge, this paper represents the first attempt to address the issues of relevance of infinite population co-evolutionary dynamics simulations within the context of shadowing properties of complex dynamical systems.

Acknowledgements

This work was partially supported by an EPSRC grant (No. GR/T10671/01) on "Market Based Control of Complex Computational Systems".

REFERENCES

Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books, New York.

Balmforth, N. J., Spiegel, E. A., and Tresser, C. (1994). Topological entropy of one-dimensional maps: Approximations and bounds. *Phys. Rev. Lett.*, 72(1):80–83.

Chellapilla, K. and Fogel, D. B. (1999). Evolution, neural networks, games, and intelligence. *Proceedings of the IEEE*, 87(9):1471–1496.

Chong, S. Y., Tiňo, P., and Yao, X. (2008). Measuring generalization performance in co-evolutionary learning. *IEEE Transactions on Evolutionary Computation*, 12(4):479–505.

Darwen, P. J. and Yao, X. (2002). Co-evolution in iterated prisoner's dilemma with intermediate levels of cooperation: Application to missile defense. *International Journal of Computational Intelligence and Applications*, 2(1):83–107.

Ficici, S. and Pollack, J. (2007). Evolutionary dynamics of finite populations in games with polymorphic fitness-equilibria. *Journal of Theoretical Biology*, 247(13):426–441.

Ficici, S. G., Melnik, O., and Pollack, J. B. (2005). A game-theoretic and dynamical-systems analysis of selection methods in coevolution. *IEEE Transactions on Evolutionary Computation*, 9(6):580–602.

Fogel, D. B., Fogel, G. B., and Andrews, P. C. (1997). On the instability of evolutionary stable strategies. *Biosystems*, 44:135–152.

Hayesand, W. and Jackson, K. (2005). A survey of shadowing methods for numerical solutions of ordinary differential equations. *Appl. Numer. Math.*, 53(2):299–321.

Hofbauer, J. and Sigmund, K. (2003a). Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519.

Hofbauer, J. and Sigmund, K. (2003b). Evolutionary game dynamics. *Bull. Amer. Math. Soc.*, 40:479–519.

Katok, A. and Hasselblatt, B. (1995). *Introduction to the modern theory of dynamical systems*. Cambridge University Press.

Katok, A. and Mezhirov, A. (1998). Entropy and growth of expanding periodic orbits for one-dimensional map. *Fundamenta Mathematicae*, 157:245–254.

Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V., editors (2007). *Algorithmic Game Theory*. Cambridge University Press, Cambridge, UK.

Sharkovsky, A. and Chua, L. (1993). Chaos in some 1-D discontinuous maps that appear in the analysis of electrical circuits,. *IEEE Transactions on Circuits and Systems I-Fundamental Theory and Applications*, 40(10):722–731.

Sigmund, K. and Nowak, M. A. (1999). Evolutionary game theory. *Current Biology*, 9:R503–R505.

Smith, J. M. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK.

AUTHOR BIOGRAPHIES

PETER TIŇO is Senior Lecturer with the School of Computer Science, the University of Birmingham, UK. His main research interests include probabilistic modeling and visualization of structured data, statistical pattern recognition, dynamical systems, evolutionary computation, and fractal analysis.

SIANG YEW CHONG is an Assistant Professor with the School of Computer Science, University of Nottingham, Malaysia Campus, and a member of the Automated Scheduling, Optimization and Planning (ASAP) Research Group, University of Nottingham, UK. His major research interests include evolutionary computation, machine learning, and evolutionary game theory.

XIN YAO is professor of computer science in the School of Computer Science, the University of Birmingham, and the Director of the Centre of Excellence for Research in Computational Intelligence and Applications (CERCIA). Prof. Yao is a Fellow of IEEE, a Distinguished Lecturer of the IEEE Computational Intelligence Society, and a Distinguished Visiting Professor at the Nature Inspired Computation and Applications Laboratory (NICAL) of University of Science and Technology of China, Hefei, China. His main interests include evolutionary computation (evolutionary optimization, evolutionary learning, evolutionary design), neural network ensembles and multiple classifiers, meta-heuristic algorithms, data mining and computational complexity of evolutionary algorithms.