Co-Evolution of the Dynamics in Population Games: the Case of Traffic Flow Assignment

Ana L. C. Bazzan
PPGC / UFRGS
Caixa Postal 15064, CEP 91501-970, Porto Alegre, RS, Brazil
bazzan@inf.ufrgs.br

ABSTRACT

In population games, one of the main interests is the evolution of the dynamics, i.e., how the distribution of individuals change along time. This is an abstract but elegant way to model population of drivers selecting routes. In this paper, a three-population asymmetric game is used to investigate the co-evolution of drivers' strategies. It is shown that the convergence to one of the Nash equilibria is achieved when the three populations co-evolve, under different rates of mutants in these populations.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Algorithms

Keywords

Game Theory, Multiagent Systems, Traffic Assignment

1. INTRODUCTION AND RELATED WORK

A population game models simultaneous interactions of a large number of simple agents distributed in a finite number of populations. In population games, typically, one is not only interested in constancy or equilibrium, but on changes. This is particularly the case when a game does not have an evolutionary stable strategy (ESS), or when it has more than one. In the present paper, game theory and evolutionary programming are combined in order to investigate the dynamics of demand in a traffic network in which the route choice of three populations of agents is modeled in a game-theoretic way. Because here three populations play a non-symmetric game with more than two actions, the replicator dynamics is not trivial to analyze and represent. Our aim is not the computation of the exact equilibrium, even because in the real-world this is probably a useless effort given that this equilibrium will not last long due to the dynamic nature of the environment. In particular, an environment is considered in which new drivers or agents replace existing ones in a way to reproduce the fact that in real-world networks, drivers unfamiliar with the network (e.g., non-commuters) also use it. This way, the goal is the investigation of the co-evolution of the three populations regarding the assignment of routes. The basic idea is that a population of strategies (for selecting a route) is reproduced, from generation to generation, proportionally to the fitness. Fitness is a function of the payoffs obtained by the driver agents after selecting and performing each joint strategy.

Game theoretic approaches have been used in traffic assignment. However, mostly, the focus is on the static game, as, e.g., studies involving the Braess paradox or the price of anarchy [4]. When the temporal aspects are considered, simplifications regarding the topology of the network are made, either by dealing with two-route scenarios ([3, 2]), or by considering a single origin-destination pair, or both ([1]).

GECCO'12 Companion, July 7–11, 2012, Philadelphia, PA, USA.
ACM 978-1-4503-1178-6/12/07.

2. METHODS AND RESULTS

A population game can be defined as follows:
populations $\mathcal{P} = \{1, ..., p\}$: society of $p \geq 1$ populations of agents where $|p|$ is the number of populations;
strategies $\mathcal{S}^p = \{s_1^p, ..., s_m^p\}$: set of strategies available to agents in population p;
payoff function $\pi(s_i^p, \mathbf{q^{-p}})$.

In this description, $|p|$ populations interact. Agents in population p have m^p possible strategies. For a (large) population of agents that can use a set of pure strategies $\mathcal{S}^p = \{s_1^p, ..., s_m^p\}$, a population profile is a vector σ that gives the probability $\sigma(s_i^p)$ with with strategy $s_i^p \in \mathcal{S}^p$ is played in the population p.

The previously mentioned idea that the composition of the population of agents (and hence of strategies) in the next generations changes with time (in this case generations) suggests that one can see these agents as replicators. In the replicator dynamics, it is assumed that members of each population p are programmed to adopt one pure strategy from a finite set $\mathcal{S}^p$.

Let n_i^p be the number of individuals using strategy $s_i^p \in \mathcal{S}^p$. Then, the fraction of agents using s_i^p is $x_i^p = \frac{n_i^p}{N^p}$, where N^p is the size of p. The interest here is on how the fraction of agents using each strategy changes with time, i.e., the derivative $\dot{x}_i^p$. Because payoffs represent reproductive fitness that is responsible for the number of successors using each strategy, one can write: $\dot{x}_i^p = (\pi(s_i^p, \mathbf{x^p}) - \bar{\pi}(\mathbf{x^p})) \times x_i^p$, with $\bar{\pi}(\mathbf{x^p})$ being the average payoff obtained by population p: $\bar{\pi}(\mathbf{x^p}) = \sum_{i=1}^{m} x_i^p \pi(s_i^p, \mathbf{x^p})$.

The previous description of population games can now be instantiate for our particular scenario. Formally, the set of **populations** is $\mathcal{P} = \{1, 2, 3\}$; the set of **strate-**

	G3		
	G2	S2	B2
G1	1/1/4	**5/6/7**	5/1/7
S1	3/4/6	4/6/8	4/1/8
B1	5/5/7	**5/6/8**	4/0/9

	T3		
	G2	S2	B2
G1	4/4/8	7/4/6	7/1/8
S1	4/6/8	5/4/6	5/1/8
B1	**5/7/8**	5/4/6	4/0/8

Table 1: Payoff matrices for the three-player traffic game; payoffs are for player 1 / player 2 / player 3 (boldface indicate the three Nash equilibria in pure strategies: σ_a, σ_b, and σ_c).

gies for each population $p \in \mathcal{P}$ is: $\mathcal{S}^1 = \{G1, S1, B1\}$, $\mathcal{S}^2 = \{G2, S2, B2\}$, and $\mathcal{S}^3 = \{G3, T3\}$; **payoff function** is as in Table 1 (explanation for these quantities omitted).

For the three-agent game whose payoffs are given in Table 1, there are five Nash equilibria, two of them in mixed strategies. Because in asymmetric games, all ESS are in pure strategies, only σ_a, σ_b, and σ_c are candidates for ESS. σ_a means agents in $p = 1$, $p = 2$, and $p = 3$ select G1, S2, and G3 respectively. In σ_b these agents select B1, S2, and G3. In σ_c agents select B1, G2, and T3. Clearly, among σ_a, σ_b, and σ_c, the first two are Pareto inefficient because σ_c is an outcome that make all agents better off.

The dynamics of the process is then accomplished with a genetic algorithm. A population p is composed by N^p agents, each programmed to play a given strategy $s_i^p \in \mathcal{S}^p$. In each generation, each agent plays g games whose sum of payoffs is its fitness. After these g games are played, the populations of agents are reproduced according to their fitness. To reproduce the behavior of new drivers in the network, a mutation rate p_m is used: with probability p_m an agent in p is replaced by its mutated version, which means that its strategy is changed to another one randomly selected.

The focus of the investigation is on issues such as what happens if the populations start with each one using a given profile σ. For instance, if this profile is σ^*, under which conditions will it remain this way? How many mutants are necessary to shift this pattern? Also, if the population starts using any σ, what happens if it is close to σ^*? Will it tend to evolve towards σ^* or move away? If it reaches σ^*, how long has it taken? What happens if there are multiple equilibria?

By analytically checking which are the stable rest points of the static game, it was found that only σ_c is an ESS. When it comes to the dynamics, this investigation can only be done using numerical simulation. For instance, the issue about whether or not an ESS will establish depends on the mutation rate. If it is too high, then the populations never converge to the selection of any Nash equilibria, much less to the ESS, because perturbations happen too often. If it is too low, an initial condition may determine which Nash equilibria will establish, which may not be the ESS.

For the simulations of the evolutionary process, the values used for the main parameters of the model were: $\mathcal{P} = \{1, 2, 3\}$, $N^1 = N^2 = N^3 = 1000$, $g = 10,000$, $\Delta = 1000$ (number of generations), and p_m was varied.

Heatmaps were produced (not shown) to depict the intensity of the selection of each of the 18 joint actions that appear in Table 1. From these maps, it is possible to conclude that for high rates of mutation (e.g., $p_m = 10^{-1}$), either σ_a or σ_b appear more frequently. Performance is poor because there is a high rate of new strategies. When $p_m = 10^{-2}$ or $p_m = 10^{-3}$, the convergence pattern is clearer but still it is not possible to affirm that one profile has established. When this rate is decreased to $p_m = 10^{-4}$ or $p_m = 10^{-5}$, it is possible to observe that one of the two cases occur: either profile σ_c establishes right in the beginning, or there is a competition between σ_a and σ_b, with one or the other ending up establishing. With decrease in p_m, there is a decrease in the time needed to either σ_a or σ_b establish, if σ_c has not already set. Lower mutation rates follow the same trend.

3. CONCLUSION

In this paper, a three-population game was defined in order to model co-evolution of strategies in a scenario in which the payoffs of the populations are not symmetric, as it is common in the real world. Although the game considers three populations only, each having a few actions, this is not an unrealistic simplification. In fact, in the majority of the situations a traffic engineer has to deal with, there is a small number of commodities (origin-destination pairs) thus three populations is not far from reality. Regarding the number of actions, it is equally the case that in the majority of the real-world cases drivers do not have more than a handful of options to go from A to B.

The contribution of this paper is twofold: the modeling, whose analytical solution is not trivial given the number of variables involved; and the investigation of the dynamics of the co-evolution, by showing that the convergence to one of the Nash equilibria is achieved under given mutation rates only. The latter has the practical effect that, for networks where the number of newcomers (non-commuters) is high, and/or drivers tend to make experimentation (e.g., in response to information broadcast), there may be difficult to achieve some near-equilibrium conditions, meaning that the network may function in sub-optimal conditions.

Acknowledgment

The author and this work were partially supported by CNPq.

4. REFERENCES

[1] A. L. C. Bazzan and F. Klügl. Case studies on the Braess paradox: simulating route recommendation and learning in abstract and microscopic models. *Transportation Research C*, 13(4):299—319, August 2005.

[2] T. Chmura and T. Pitz. An extended reinforcement algorithm for estimation of human behavior in congestion games. *Journal of Artificial Societies and Social Simulation*, 10(2), 2007.

[3] F. Klügl and A. L. C. Bazzan. Route decision behaviour in a commuting scenario. *Journal of Artificial Societies and Social Simulation*, 7(1), 2004.

[4] T. Roughgarden and É. Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.