Clonal Plasticity: An Autonomic Mechanism for Multi-Agent Systems to Self-Diversify

JAAMAS Track

Vivek Nallur
University College Dublin
Dublin, Ireland
vivek.nallur@ucd.ie

Siobhán Clarke
Trinity College Dublin
Dublin, Ireland
siobhan.clarke@scss.tcd.ie

ABSTRACT

Diversity has long been used as a design tactic in computer systems to achieve various properties. Multi-agent systems, in particular, have utilized diversity to achieve aggregate properties such as efficiency of resource allocations, and fairness in these allocations. However, diversity has usually been introduced manually by the system designer. This paper proposes a decentralized technique, clonal plasticity, that makes homogeneous agents self-diversify, in an autonomic way. We show that clonal plasticity is competitive with manual diversification, at achieving efficient resource allocations and fairness.

ACM Reference Format:
Vivek Nallur and Siobhán Clarke. 2018. Clonal Plasticity: An Autonomic Mechanism for Multi-Agent Systems to Self-Diversify. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. Andre, S. Koenig (eds.), Stockholm, Sweden, July 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

Socio-technical systems are all around us, and increasingly computers are an intrinsic part of these systems [3]. Smart-grids [11], vehicular ad-hoc networks [6, 9], smart buildings [17], e-procurement [16], cloud computing [12], healthcare [1], and transport [10] are some examples of domains where humans interact with computer systems to achieve a particular goal. Due to the speed, complexity, and frequency of decision-making involved in these domains, computers are often required to be autonomic and adaptive to the dynamic environment around them. Indeed this is a foundational reason why we create self-adaptive systems. Socio-technical systems were first defined by Emery and Trist [7] to be a complex interaction between humans, machines and their joint environment. This makes the aggregate outcome of individual agent actions dynamic and difficult to predict. An action taken by an agent at a particular point in time may have an outcome that is vastly different from the same action at a different point in time. In many such systems, agents make (or recommend) decisions that result in allocation of *common pool* resources [15]. That is, resources in a social context are divided amongst actors/entities in a manner that achieves some objective. In such systems, less-than-ideal outcomes are most clearly visible where some resource is being allocated, and there is no centralized mechanism to ensure that autonomous actions by each agent does not result in a disastrous allocation, at the aggregate level. This is specially so in cases where the agents are not, *apriori*, designed to be cooperative (*e.g.,* traffic jams, where each vehicle is exhibiting individually rational behaviour). In many domains, the natural model is a competitive set of agents where each agent attempts to fulfill its goals, without regard to the effect it has on the system-as-a-whole. Indeed, if the system is large enough, the agent cannot even view the system-as-a-whole. Examples of such domains include traffic, energy markets, financial markets, etc. In such domains, multi-agent system designers need a mechanism that can harness the multiplicity of agents to produce good outcomes, even where the agents are not explicitly designed to cooperate. In other words, resource allocation mechanisms in multi-agent systems must guard against some agents grabbing more than their fair share of resources without actually trampling on the autonomy of the agents themselves. As pointed out by *Chevaleyre et al.*, in some systems, it is impossible to allocate resources (roads, in this case) such that the result is both efficient and envy-free [5]. Furthermore, in certain simultaneous games, even solely envy-freeness is impossible to guarantee [2]. This paper is concerned with ameliorating the negative effect of the impossibility of envy-freeness in such games, by enabling long-term fairness through cumulative allocations. It is in this context that we introduce the idea of fairness. We distinguish fairness from envy-freeness, by observing that a series of envious allocations may, over time, turn out to be cummulatively fair. The contribution of this paper is to present a technique that allows us to achieve both efficiency and fairness. We borrow diversity as a metaphor from ecology, and model the multi-agent-system as an ecosystem, where each agent is concerned in a selfish way, only with its own survival. We had previously shown [14] that **diversity in an ecosystem of competing agents can drive the collective measure of efficiency and fairness higher**. However, the algorithmic diversity demonstrated was achieved via *deliberate human design*. This is not a scalable mechanism for long-lived, large multi-agent systems consisting of thousands of agents.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. Andre, S. Koenig (eds.), July 2018, Stockholm, Sweden.

In this paper, we present a mechanism (clonal plasticity) that allows agents to self-diversify in an autonomic fashion. We evaluate the diversity achieved via clonal plasticity, and reflect on its scalability.

2 ALGORITHMIC DIVERSITY

Decision-making and resource allocation mechanisms are all realized by various algorithms, and hence we view algorithmic diversity as an appropriate level of granularity for achieving our goals. We define *Algorithmic Diversity* as: a variability in the output produced by a set of algorithms, when faced with the same input/input-sequence. We do not consider differences in the internal data structures used (stacks, queues, trees, etc.) or control flow implemented (iteration, recursion, continuations, labelled jumps, etc.). This definition positions us in domains where there are multiple valid answers and no deterministic ways of calculating the optimal response in advance. Traffic flow management systems are a good example of such a domain. Each vehicle in traffic has (possibly) different source and destination and different strategies in how it gets there (*e.g.*, fastest route, shortest route, least polluted route, etc.). We use a self-organizing game called the Minority Game (MG) [4] as our exemplar Multi-Agent System in a traffic setting, primarily because it has been well-studied and well-understood. The Minority Game (MG) consists of an odd number (N) of agents, playing a game in an iterated fashion[1]. At each timestep, each agent chooses from one of two actions (A or B), and the group that is in the minority, wins. Since N is an odd integer, there is guaranteed to be a minority. Winners receive a reward, while losers receive nothing. After each round, all agents are told which action was in the minority. This feedback loop induces the agents to re-evaluate their action for the next iteration.

3 AUTOMATED GENERATION OF DIVERSITY

We introduce the main contribution of this paper: a self-adaptation mechanism called *clonal plasticity* which can be used by a group of agents to self-diversify. We use Clonal Plasticity, to generate diversity automatically amongst the algorithms implemented. We show that this generated diversity is competitive with the manually introduced diversity presented in the previous section. We had previously used clonal plasticity [13] to enable decentralized adaptation, based on localized environmental changes. Due to its decentralized nature, the adaptations performed by multiple agents in the system, results in variations being naturally produced at a system-wide level. Each minority game was played with a population size of 501 agents, through a simulation time period of 2000 steps. For each variation in the experimental setup, the data is reported as an average of 100 simulations[2].

[1]Note that the exact number of agents does not impact the results, as long as there is a well-defined minority at every stage.
[2]All code for this experiment can be found at:
https://bitbucket.org/viveknallur/clonal_plasticity_algo_diversity.git

Diversity	Reward	Gini
1.7	509	11.2
2.1	616	1.0
2.2	755	1.3
2.5	789	0.9
2.6	770	1.1
2.7	867	2.1

Table 1: Relationship between Diversity and Reward and Fairness

We measure the efficiency of the system through the median amount of rewards accumulated by the agent, after the simulation period. We measure the fairness of the system using the Gini index. The Gini index is used to calculate dispersion in the income distribution of a society [8]. In a society that is perfectly equal, the Gini index is equal to zero. Therefore, the closer the Gini index is to zero, the fairer the distribution of rewards.

3.1 Effectiveness of Diversity in Fairness and Efficiency

From Table 1 we see that the lowest level of diversity is also responsible for the highest level of the Gini index. Recall that the Gini index measures the inequality of the distribution of income. Thus, the closer the Gini index value is to zero, the more equitable the distribution is. This table shows that attaining both, efficiency of the mechanism (close to the theoretical optimum of 1000) and fairness (a theoretical optimum of zero) is a difficult problem. While increasing diversity directly increases the efficiency (867 is closer to 1000 than 509), it does not have the same clear impact on fairness. While the fairness achieved with the highest diversity (H-index of 2.7 results in Gini index of 2.1) is certainly better than the fairness achieved with low diversity (H-index of 1.7 results in Gini index of 11.2), it is interesting to see that a diversity level of 2.5 actually provides the highest amount of fairness (Gini index of 0.9). This indicates the system might have an optimal 'sweet spot' for diversity, and going beyond such a limit might not necessarily improve both efficiency and fairness.

4 CONCLUSION

In ecological sciences, as well as computer science, diversity has been shown to be an important property of an ecosystem. The concept of diversity has previously been used in computer security, search, machine-learning, as well as software engineering. However, instead of qualitative statements about diversity, we present a quantitative approach to measuring the effects of diversity, and an adaptation approach that increases the amount of diversity present in the system. This is especially important in socio-technical systems that involve human interactions. Further, we present an adaptation process (clonal plasticity) that can drive a system, from low levels of diversity to higher levels, depending on the number of plasticity points that individual agents possess. Being a completely decentralized mechanism, clonal plasticity is suitable for systems consisting of a large number of agents.

REFERENCES

[1] Fabrice Axisa, P. M. Schmitt, C. Gehin, G. Delhomme, E. McAdams, and A. Dittmar. 2005. Flexible technologies and smart clothing for citizen medicine, home healthcare, and disease prevention. *IEEE Transactions on Information Technology in Biomedicine* 9, 3 (Sept 2005), 325–336. https://doi.org/10.1109/TITB.2005.854505

[2] Eric Balkanski, David Kurokawa, Simina Brânzei, and Ariel D. Procaccia. 2014. Simultaneous Cake Cutting. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI'14)*. AAAI Press, 566–572. http://dl.acm.org/citation.cfm?id=2893873.2893962

[3] Gordon Baxter and Ian Sommerville. 2011. Socio-technical systems: From design methods to systems engineering. *Interacting with Computers* 23, 1 (2011), 4–17. https://doi.org/10.1016/j.intcom.2010.07.003 arXiv:http://iwc.oxfordjournals.org/content/23/1/4.full.pdf+html

[4] D. Challet and Y.C. Zhang. 1997. Emergence of cooperation and organization in an evolutionary game. *Physica A: Statistical Mechanics and its Applications* 246, 34 (1997), 407 – 418. https://doi.org/10.1016/S0378-4371(97)00419-6

[5] Yann Chevaleyre, Paul E. Dunne, Ulle Endriss, Jérôme Lang, J Lemaitre, Julian Maudet, Nicolas Padget, Steve Phelps, Juan A. Rodríguez-aguilar, and Sousa P. 2006. Issues in Multiagent Resource Allocation. *Informatica* 30, 2 (June 2006), 3–31. https://doi.org/10.1017/S0269888905000470

[6] M. Ding, X. Cheng, and G. Xue. 2003. Aggregation tree construction in sensor networks. In *Vehicular Technology Conference, 2003. VTC 2003-Fall. 2003 IEEE 58th*, Vol. 4. 2168–2172 Vol.4. https://doi.org/10.1109/VETECF.2003.1285913

[7] F.E. Emery and E.L. Trist. 1960. Socio-technical Systems. *Management Science, Models and Techniques* 2 (1960), 83–97.

[8] C Gini. 1936. On the Measure of Concentration with Special Reference to Income and Statistics. *Colorado College Publication* 208 (1936), 73–79.

[9] S. S. Manvi, M. S. Kakkasageri, and Jeremy Pitt. 2009. Multiagent Based Information Dissemination in Vehicular Ad Hoc Networks. *Mob. Inf. Syst.* 5, 4 (Dec. 2009), 363–389. https://doi.org/10.1155/2009/518042

[10] A. S. Masoum, S. Deilami, P. S. Moses, M. A. S. Masoum, and A. Abu-Siada. 2011. Smart load management of plug-in electric vehicles in distribution and residential networks with charging stations for peak shaving and loss minimisation considering voltage regulation. *IET Generation, Transmission Distribution* 5, 8 (August 2011), 877–888. https://doi.org/10.1049/iet-gtd.2010.0574

[11] A. H. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia. 2010. Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid. *IEEE Transactions on Smart Grid* 1, 3 (Dec 2010), 320–331. https://doi.org/10.1109/TSG.2010.2089069

[12] Vivek Nallur and Rami Bahsoon. 2010. Self-Adapting Applications Based on QA Requirements in the Cloud Using Market-Based Heuristics. In *Towards a Service-Based Internet (Lecture Notes in Computer Science)*, Elisabetta Di Nitto and Ramin Yahyapour (Eds.). Springer-Verlag, 51–62. https://doi.org/10.1007/978-3-642-17694-4_5

[13] Vivek Nallur, Nicolás Cardozo, and Siobhán Clarke. 2016. Clonal Plasticity: A Method for Decentralized Adaptation in Multi-agent Systems. In *Proceedings of the 11th International Workshop on Software Engineering for Adaptive and Self-Managing Systems (SEAMS '16)*. ACM, New York, NY, USA, 122–128. https://doi.org/10.1145/2897053.2897067

[14] Vivek Nallur, Eamonn O'Toole, Nicolás Cardozo, and Siobhán Clarke. 2016. Algorithm Diversity: A Mechanism for Distributive Justice in a Socio-Technical MAS. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems (AAMAS '16)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 420–428. http://dl.acm.org/citation.cfm?id=2936924.2936986

[15] Jeremy Pitt, Dídac Busquets, and Régis Riveret. 2015. The pursuit of computational justice in open systems. *AI & SOCIETY* 30, 3 (2015), 359–378. https://doi.org/10.1007/s00146-013-0531-6

[16] F. Wu, G. Zsidisin, and A. Ross. 2007. Antecedents and Outcomes of E-Procurement Adoption: An Integrative Model. *IEEE Transactions on Engineering Management* 54, 3 (Aug 2007), 576–587. https://doi.org/10.1109/TEM.2007.900786

[17] J. H. Yoon, R. Baldick, and A. Novoselac. 2014. Dynamic Demand Response Controller Based on Real-Time Retail Price for Residential Buildings. *IEEE Transactions on Smart Grid* 5, 1 (Jan 2014), 121–129. https://doi.org/10.1109/TSG.2013.2264970