Research Repository UCD

Title	Clonal Plasticity: An Autonomic Mechanism for Multi-Agent Systems to Self-Diversify
Authors(s)	Nallur, Vivek, Clarke, Siobhán
Publication date	2017-12-07
Publication information	Nallur, Vivek, and Siobhán Clarke. "Clonal Plasticity: An Autonomic Mechanism for Multi-Agent Systems to Self-Diversify" 32, no. 2 (December 7, 2017).
Publisher	Springer
Item record/more information	http://hdl.handle.net/10197/9934
Publisher's statement	The final publication is available at Springer via http://dx.doi.org/10.1007/s10458-017-9380-x .
Publisher's version (DOI)	10.1007/s10458-017-9380-x

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

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.

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

Vivek Nallur · Siobhán Clarke

the date of receipt and acceptance should be inserted later

1 Introduction

Socio-technical systems are all around us, and increasingly computers are an intrinsic part of these systems [1]. Smart-grids [2], vehicular ad-hoc networks [3,4], smart buildings [5], e-procurement [6], cloud computing [7], healthcare [8], and transport [9] 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 [10] to be a complex interaction between humans, machines and their joint environment. This means that design choices made for the construction/selection of computer programs would also significantly affect the humans in the system.

A commonly accepted definition of an agent is given by Woolridge [11] - "An agent is a computer system that is situated in some environment, and that is capable of autonomous action in this environment in order to meet its design objectives". Computer systems that function in socio-technical systems are increasingly agents that are, by definition, autonomous. However, in a socio-technical system filled with such agents, each with its own set of constraints, goals and strategies, the resulting collective system may not be predictable. Rather, the multiplicity of agents, their possible actions, and corresponding feedback loops could result in a complex system [12]. This makes the aggregate outcome of individual 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 *com-*

mon pool resources [13]. 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. Typically, resource allocation in multi-agent systems are analyzed for two types of aggregate outcomes: (a) efficiency, and (b) social welfare. The issue of efficiency is quantified as *Pareto-Efficiency*: where the total allocation is such that it is not possible to improve the allocation for some agents without making it worse for some other agents. A pareto-efficient allocation, therefore, represents a situation where the solution cannot be bettered, without hurting at least one other agent. Social welfare is usually codified as *envy-freeness*: no agent would rather hold the resource bundle held by another agent [14]. For example, consider a route-choice problem through a city *i.e.*, pick road A or road B to reach a destination. This is a familiar situation in many cities, where multiple humans, utilizing commonly known (albeit coarse-grained) information about traffic flows, along with their own private heuristics (*e.g.*, road A is generally congested from 7:30 - 8:30 am), make route choices to reach their destination. These agents may use additional machine intelligence in the form of real-time traffic updates through GPS devices/smartphones, and shortest-path algorithms, while making their choices. The reward for this choice is determined dynamically, based on the aggregate number of agents that make a particular choice without prior coordination, *i.e.*, the agent that chooses the ‘road less travelled’ would experience lower congestion. There is no way to force an agent to make a particular choice, or to choose a particular strategy. However, regardless of the strategy used to pick a route, it must yield a satisfactory outcome for all agents, *at least some of the time*, else some agents would be envious of others. Envious allocations are not a good outcome in socio-technical systems, in contrast to efficiency-based problem formulations, where the only concern is to optimize the traffic flow through that route. 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 [15]. Furthermore, in certain simultaneous games, even solely envy-freeness is impossible to guarantee [16].

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 al-

locations. 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 cumulatively fair. We invoke Rescher’s notion of fairness through distributive justice [17] over the total resources allocated amongst all agents. All agents have a legitimate claim on the road, and distributive justice consists of reconciling these claims in the face of plurality of allocation choices. Rescher says, “... a claim represents what an individual ought *ideally* to be given in the light of the principles of justice ...” [18]. In the case of the route-choice problem, this would be the evenness of the rewards (experience lower congestion) that accumulate to drivers, over multiple days. That is, even though it is impossible to ensure that all users in a road network get the same rewards on any given day, it is **sufficient** to ensure that, over time all users get **approximately** the same cumulative reward. The lack of envy-freeness on any given day, is compensated by the distribution of envious allocations over time and across agents.

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 [19] 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 (or indeed, a sustainable mechanism) for long-lived, large multi-agent systems consisting of thousands of agents. 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.

The rest of this paper is structured as follows: First, we review resource allocation in multi-agent systems and the various ways in which diversity has been used in software systems (section 2). We also review various mechanisms of introducing and sustaining diversity used in other domains of computer science (section 3). Next, we make explicit what we mean by diversity in multi-agent systems, and how the diversity present in a multi-agent system can be quantified (section 4). We consider a multi-agent system, where multiple agents act autonomously to secure resources (section 5) and show that increasing the amount of diversity results in an increased distributive justice amongst the agents (section 6). We then introduce a self-tuning mechanism (clonal plasticity) that measurably increases the diversity amongst the agents (section 7) and formalize the expected fairness from clonal plasticity (section 8). We show experimentally that clonal plasticity does increase diversity, compare its effectiveness against manual diversification, and discuss the scalability of clonal plasticity (section 9). We also contrast it against other phenotypic-diversity generation mechanisms (section 10) as well as non-phenotypic diversity generation mechanisms (section 11). We reflect on the limitations and weaknesses of clonal plasticity (section 12) and finally, we conclude the paper (section 13) with a few thoughts on open questions and potential areas of further research.

2 Related Work

We first look at the concept of diversity, as transferred from the ecological sciences, and how it has been implemented in computer systems. Then we review the work on decentralized resource allocation in multi-agent systems.

2.1 Diversity in Software Systems

In the natural world, diversity has long been considered a contributor to stability of natural ecosystems [20, 21]. While not sufficient to guarantee the stability of an ecosystem, it is nevertheless considered critically important, due to the *Insurance Hypothesis* [22–24]. Briefly, the Insurance Hypothesis relies on the observation that “differential response by species to a perturbation or variation in environment causes a dampening of the effect of the perturbation”. Simply put, an increase in diversity increases the probability that there will exist some species that is able to respond differently to environmental perturbations. The presence of diversity is therefore hypothesized to be an insurance against an ecosystem collapsing due to perturbations.

2.1.1 Security

Researchers have applied the notion of diversity in order to protect systems from attacks and intrusions. The principal use of diversity in these approaches has been in the form of creating a *moving target* for defense. The idea is to reduce the predictability in the structural layout of programs in memory, thus making it more expensive for attackers to create an attack that can work across a range of systems. Multiple stages in the lifecycle of software development such as, *creation* [25], *compilation* [26–29], *deployment of software* [30–33] have benefited from the introduction of variability.

2.1.2 Search and Machine Learning

Differential response has also been used to explore different parts of a search space, in order to aid search techniques such as genetic algorithms [34] and particle swarm optimization [35, 36]. The presence of diversity has been used to create ensemble learners that learn faster than other learners [37]. There has been work in the area of support vector machines that incorporate the notion of diversity to create more accurate active-learning [38, 39]. Other computational intelligence techniques such as Artificial Immune Systems have also recognized diversity [40, 41] to be integral to its effectiveness. Almost all techniques in the field of search and computational intelligence now explicitly incorporate mechanisms to generate and maintain diversity.

2.1.3 Software Engineering

Apart from resilience in the presence of attackers, diversity as a concept has also been investigated for software product lines [42] in the form of design diversity, management of functional diversity using plugins [43] and survivability [44]. Diversity, as a

software engineering phenomenon has also been extensively studied in various surveys [45–47].

All of these surveys and papers have considered "differential response by a species" to be the critical element of diversity. However, to the best of our knowledge, we are not aware of any prior work that proposes a systematic technique for diversification of multi-agent systems (MAS). This is critically important since agents in a MAS are being used to control and coordinate complex pieces of software, such as smart-grids [2], vehicular ad-hoc networks [3, 4], smart buildings [5], healthcare [8], and transport [9] systems.

2.2 Resource Allocation in Multi-Agent Systems

The allocation of resources within a society of agents such that each individual agent's preferences are met forms a multi-disciplinary research area, with inputs from Computer Science and Economics. There are many approaches to realizing an allocation, and each approach has its own properties with regard to three fundamental axes: **(a)** what kind of resources (divisible / indivisible) are being allocated? **(b)** what is the allocation procedure to be followed? and **(c)** what constitutes a good allocation? Each of these axes can be sub-divided further: **(i)** how are resources and preferences represented? **(j)** how do agents negotiate amongst themselves? **(k)** what is the computational complexity of the entire allocation process, and so on. Depending on the domain where multi-agent resource allocation (MARA) occurs, different combinations of these properties will take priority over others. The breadth of domains for which MARA has been used is broad enough (network routing [48], air traffic control [49], manufacturing [50, 51], transport and logistics [52], e-procurement [53], and even sharing satellite resources [54]) that these properties become the critical factors on which the choice of allocation approach in a particular software system hinges.

2.2.1 Types of Resources

Different types of resources may be in contention for allocation. Some resources are divisible (*e.g.*, network bandwidth [48]) whereas others are indivisible (*e.g.*, slot for landing aircraft [49]). Some resources are shareable (*e.g.*, images from a satellite [54]), while others are non-shareable (*e.g.*, resources in a factory [50]). Some domains may even have multiple types of resources, but may have business constraints that only allow certain *bundles* of resources to be allocated (*e.g.*, multi-unit e-procurement [53]).

2.2.2 Allocation Approaches

Fundamentally, all allocation procedures can be classified as being *centralized* or *decentralized*. Centralized approaches are mostly in the form of combinatorial auctions [55], which may themselves use multiple strategies for trading (*e.g.*, machine-learning [7], game-theory [56]). Decentralized approaches include negotiation protocols such as *Contract-Net* [57] and its variants (Leveled Commitments [58], Constraint Propagation [59]). Centralized approaches have the advantage of being easy

to understand and implement, however they suffer from their inability to scale. For a large system, consisting of hundreds or thousands of agents, scalability is a critical issue. Decentralized approaches, on the other hand, scale well to large numbers of agents, which are typical of large socio-technical systems such as traffic on city roads. There have also been attempts to extract the best of both worlds by performing resource allocation in a hybrid manner, with a mix of centralized and decentralized decision making. [60], [61], [62] are all approaches that utilize particle-swarm optimization as a middle ground, with each individual particle behaving in a decentralized manner, along with neighbourhood rules that result in the creation of small hierarchies (partial centralization).

2.2.3 Social Welfare

One of the objectives of MARA is the achievement of allocations that reflect the desired preferences of the individual agents. This introduces an axis, where the aggregate social distribution of resources is as important as the efficiency and ease of allocation. Concepts from social welfare and social choice theory [63] influence the allocation procedures, and the desired ordering of resource allocations. Bentham's utilitarianism [64] is concerned with maximizing the total utility in the system. Other theories such as Rawls's theory of justice [65] advocates an egalitarian allocation which is measured by the utility gained by the agent that is *worst-off*. Other notions of welfare include *elitist social welfare* (measure the utility of the agent that is best-off) and *maximin ordering* (ranks allocations such that the worst-off agent has the highest possible utility) which can be used to evaluate the allocations generated by various schemes.

2.3 Distributive Justice

The notion of distributive justice has been dealt with in multiple works in the field of social choice theory. However, in the field of computer science, we are aware of only one thread of research [66,67], that explicitly attempts to ensure distributive justice in multi-agent systems. We view our work as complementary, as we approach the same research theme, but from two different ends of the spectrum. Pitt et al. develop a formal framework for operationalizing distributive justice rules in self-organizing systems [66]. In this work, we present an empirical approach that measures the amount of distributive justice under certain conditions and present a mechanism to autonomically increase this justice in homogeneous systems.

3 Diversity Generation Mechanisms in Literature

As is evident from section 2, diversity has been well-recognized as a mechanism for positively influencing system properties, in many domains. However, there hasn't been a similar wealth of work in actually generating diversity, in a homogeneous system. In the field of security, diversity through randomization [27–33] has been the

most dominant tactic. Most work in automated diversity generation has been in the domain of evolutionary and swarm-based algorithms, where populations (and therefore diversity) are a natural solution concept.

3.1 Evolutionary Computing

Search techniques such as Evolutionary Algorithms, Genetic Algorithms, Genetic Programming, etc. all rely on population-based approaches along with fitness functions that aim to mimic the natural process of evolution [68]. In all of these techniques, diversity is a key factor in whether the population is able to converge to a good solution. Note that diversity, *per se*, is not a requirement, rather it is a strong observational recommendation. This is echoed by McPhee and Hopper in [69], "Progress in evolution depends fundamentally on the existence of variation of population. Unfortunately, a key problem in many Evolutionary Computation (EC) systems is the loss of diversity through premature convergence". Diversity, hence, is generated or preserved through various approaches. Cobb and Grefenstette introduced a random immigrants approach [70] where random individuals are inserted into the population in every generation. Introduction of specialized selection operators (such as *incest prevention* [71] and *mating restrictions* [72]) also encourage diversity in the population by disallowing genetically similar individuals to mate. Črepinšek et al. [73] classify different strategies for diversity maintenance and diversity control in a population and observe that most methods simply expect an increase in diversity and hence, a better balance between exploration and exploitation. They however caution that not all diversification is useful. Depending on whether more exploration is required or more exploitation is necessary, the levels of useful diversity will vary. That is, increasing diversity by itself does not necessarily mean that the system will perform better.

3.2 Swarm-Based Optimization Algorithms

Swarm-based algorithms have been used for several combinatorial and optimization problems. Algorithms in this category view swarms as any restrained collection of interacting agents or individuals [74]. Particle Swarm Optimization (PSO) was introduced by Eberhart and Kennedy [75] as a mechanism for optimization of non-linear functions. In a similar vein, swarm-based algorithms such as Ant Colony Optimization [76] and Bee Colony algorithms [74, 77] have been proposed as meta-heuristics for solving combinatorial problems in traffic and transportation [78]. The central problem in all of these optimization algorithms is how to avoid local minima in complex search spaces, and achieve global optima quickly. Diversity is usually achieved in these methods using concepts such as 'critical thresholds' [79], 'electrostatic charges' [80], etc. The fundamental idea in these methods is to probabilistically avoid dense neighbourhoods by making particles move away from each other upon triggering a certain event/threshold.

3.3 Machine Learning

In the field of Machine Learning, generation of diversity has been explicitly considered in ensemble-based methods. Diversity of neural network architectures has been investigated through the process of speciation [81]. In [81], the problem of combining multiple artificial neural networks is tackled via a Darwinian evolutionary process along with *fitness sharing* as the key ingredient that promotes speciation, and thus the creation of diverse neighbourhoods. Their mechanism specifies a population-wide calculation of fitness, and subsequent adjustment of individual-fitness based on the sharing radius. Fitness sharing penalizes individuals in dense neighbourhoods, while individuals in sparse neighbourhoods are promoted, thus increasing the probability that diverse individuals are carried forward into the next generation.

Concept Drift is an important topic in online learning mechanisms, where classifiers that learn on a particular dataset, are subsequently required to classify using variables that have changed with time. Minku et al. ([82]) had shown that while ensembles with low diversity showed lower test errors, they were unable to cope with a dynamic environment and consequent concept drift. Ensembles with higher diversity were able to cope better, regardless of the type of drift encountered. Subsequently, they proposed a method that used four ensembles with two levels of diversity [37] to successfully handle environments, both without drift as well as with concept drift. The levels of diversity are generated using different values for presenting a training example. These different values (drawn from a poisson distribution) result in different ensembles learning different classifications.

4 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.) for two reasons:

1. Detecting whether two algorithms are the same even when they are exceedingly simple in their construction is so hard that for practical purposes it is infeasible [83]
2. More philosophically, from the perspective of the environment (or other agents), if two systems produce exactly the same output (behaviour), regardless of the input sequence, then they are identical, regardless of their internal construction.

This definition positions us in domains where there are multiple valid answers and no deterministic ways of calculating the optimal response in advance. Autonomous systems fall squarely into this category since the environment in which they operate is dynamic in nature. Traffic flow management systems are a good example of dynamically changing environments. 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.). It is in these scenarios that aggregate outcomes such as efficiency and fairness are important. While autonomous agents may negotiate amongst themselves to reach an efficient allocation of road resources, it is not necessary that this allocation is fair. It is in this context that we hypothesize that diversification of algorithms leads to a fairer distribution of rewards amongst the population, regardless of the actual algorithms in use. This hypothesis was first formulated and tested in [19]. In this work, we extend the results as well as introduce an autonomic mechanism to allow agents to self-diversify. We also compare the effectiveness of the autonomic mechanism against the manual efforts earlier.

We use a self-organizing game called the Minority Game (MG) [84] as our exemplar Multi-Agent System in a traffic setting, primarily because it has been well-studied and well-understood. The Minority Game (MG), introduced by Challet and Zhang, consists of an odd number (N) of agents, playing a game in an iterated fashion¹. 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.

5 Minority Game

While simple in its conception, MG has been used in many fields like econophysics [85–87], multi-agent resource allocation [88, 89], emergence of cooperation [90], and heterogeneity [91, 92]. The Minority Game (MG) is intuitively applicable to many domains, such as traffic (the driver that chooses the ‘less-travelled’ road experiences less congestion), packet traffic in networks, ecologies of foraging animals, etc. MG has been acknowledged to be an example of a complex adaptive system, where the individual behaviour is simple to understand, yet the aggregate dynamics are complex and have served as a proxy in multiple domains such as financial markets [93], distributed decision-making in wireless networks [94], and smart-buildings [95].

5.1 Description of the Game

There is a set of $N = \{1, \dots, n\}$ agents ($N = \text{odd}$), each implementing an algorithm $A = \{a_1, \dots, a_n\}$. At every iteration of the game G_i , each agent chooses from a binary set $\{A, B\}$. Thus, at each iteration i of the game (G):

$$G_i(a_j) \rightarrow A|B$$

refers to agent j making a binary choice from A or B . If the subset of agents that chose A is smaller than the subset of agents that chose B , then the agents that chose A are said to be in the minority ($\text{min}G_i$) at iteration i . And vice versa. The agents that

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.

are in $\min G$ for any particular iteration, earn a reward of one unit. The agents that are not in $\min G$ receive no reward. If R_i is the reward earned by an agent for a particular iteration, then the total reward earned by an agent after T iterations is given by:

$$\sum_{i=1}^T R_i(a)$$

Clearly, the allocation of rewards is most efficient when the size of $\min G = N/2$. While the allocation is least efficient when the size of $\min G = 1$ (one agent chooses A, and all the rest choose B).

5.2 Measuring Algorithm Diversity

In MG, each agent shares the same functional objective, to continue to be on the side of the minority, and thereby accumulate a reward. The difference in each agent is the algorithm used to compute the best action (A or B). Depending on the type of algorithm used, differences in initialization parameters could lead to different actions chosen (*e.g.*, in evolutionary or machine learning algorithms). Thus, any algorithm that exhibits a different output-sequence is categorized as a different variety. The calculation of diversity in the game depends only on the variety and balance in the population of agents (implementing those algorithms). This calculation is most famously done by the *Shannon Index* (eq. 1). The Shannon Index measures the entropy (or uncertainty) in picking an individual of a particular type, from a dataset [96]. Thus, for instance, if all the agents in a MAS implement exactly the same algorithm, the entropy (or uncertainty) will be zero. The higher the diversity in the population, the higher is the Shannon Index. In eq. 1, R is the number of types in the population, and p is the proportion of individuals in the population with type i .

$$H' = - \sum_{i=1}^R p_i \ln p_i \quad (1)$$

The Shannon Index is a measure of both, the *richness of species* (number of different algorithms) and *evenness of species* (abundance of agents implementing an algorithm). We use the Shannon Index to measure how much algorithms differ, in their key parameters. Thus, if a particular initialization parameter (that has an effect on the output-sequence³) can take multiple values, the Shannon Index will measure how many such differentiated types exist and how many individuals implement that type, as a proportion of the total population.

6 Effect of Algorithmic Diversity

In this section, we re-cap the important aspects of the results from [19]. Diversification of algorithms was performed on two levels: (a) Parameter Diversification, and (b) Strategy Diversification². Three algorithms were chosen as the candidate algorithms:

² All code and data for this experiment is available at <https://bitbucket.org/viveknallur/aamas2016.git>

1. **BestPlay**: Introduced by the creators of the MinorityGame, this is a simple strategy that nevertheless exhibits interesting dynamics at the aggregate level [84]. In this strategy, all players had access to a bounded collective memory (m) of the game, along with a pool of strategies(s). Collective memory is implemented in the form of a history of which action (+1 or -1) was in the minority, at each timestep. Each player drew s strategies randomly from a strategy space, that indicated which action to play next, based on the history of the game. Each player allocated virtual points amongst all its strategies, based on the continually changing history of the game. Strategies that correctly predicted the correct output, regardless of whether they were used or not, received virtual points. Each player, for the next round, used the top-ranking strategy amongst his pool of strategies. Table 1 shows a sample strategy pool used by a player. The size of each strategy depends on m , the memory of the game, while the total number of strategies in the pool depends on the computational power the agent possesses. Table 1 is therefore taken from a game with a memory of size 3, while the agent has a pool size (k) of 4. Typically, each agent in the entire game has the same pool size (k). While it is intuitively obvious that the greater the pool size, the greater the probability of an agent drawing good strategies, it is impossible for an agent to have a pool size that exhaustively covers every possible strategy, since the strategy space grows hyper-exponentially(2^{2^m}) with increase in m .

Strategy	Output
[+1 -1 +1]	-1
[-1 -1 -1]	-1
[-1 -1 +1]	+1
[+1 +1 +1]	-1

Table 1: Sample Strategy Pool for a Player

2. **Evolutionary**: Evolutionary algorithms have been used quite successfully in multiple domains [97, 98]. As pointed out in [99], evolutionary algorithms are now being compared to human beings, in their ability to solve hard problems. We have implemented an evolutionary algorithm that uses *BestPlay*'s strategy representation as a genome and performs crossover and mutation on it. The strategy is as follows: At every reproduction cycle, the ten most poorly performing agents get rid of their strategies and adopt a combination of strategies from two parents randomly selected from the ten best performing agents. Effectively, two of the ten best performing agents have reproduced to create an offspring that replaces one of the ten worst agents. The offspring mutates the strategies obtained from its parents, by some small mutation probability. This allows the offspring to get better, as well as explore more of the strategy space. The reproduction cycle refers to the rounds after which evolution occurs. A frequency of 20 cycles means that after every 20 rounds of the game, the worst 10 agents discard their strategies

and inherit good strategies from 2 of the top 10 agents. These are then mutated according to the agents' own mutation probability.

3. **Roth-Erev:** We implement a Reinforcement-Learning algorithm, called Roth-Erev [100]. Reinforcement-Learning(RL) is a class of machine-learning algorithms where an agent tries to learn the optimal action to take, in a certain environment, based on a succession of *action-reward* pairs. Given a set of actions (a) that an agent can take, the Roth-Erev learning algorithm performs the following steps:
 - (a) Initialize initial propensities (q) amongst all actions (N)
 - (b) Initialize strength (s) of initial propensities
 - (c) Initialize recency (ϕ) as the 'forgetting' parameter
 - (d) Initialize epsilon (ϵ) as the exploration parameter
 - (e) Choose action (a_k) at time (t), based on q
 - (f) Based on reward(r_k) for a_k , update q using the following formula:

$$q_j(t+1) = [1 - \phi]q_j(t) + E_j(\epsilon, N, k, t)$$

$$E_j(\epsilon, N, k, t) = \begin{cases} r_k(t)[1 - \epsilon] & \text{if } j = k \\ r_k(t) \frac{\epsilon}{N-1} & \text{otherwise} \end{cases}$$

- (g) Probability of choosing action j at time (t) is given by:

$$P_j(t) = \frac{q_j(t)}{\sum_{n=1}^N [q_n(t)]}$$

- (h) Repeat from step 3e

6.1 Parameter Diversity

Parameter diversity refers to diversity in the initialization parameters of the algorithm being implemented by the agents. That is, for any game, all agents still implement the same code, but some parameter in the algorithm is diversified. For the Evolutionary and Roth-Erev algorithms, even minute differences within the initialization parameters can cause the aggregate average payoff to vary. For each of the three families, the following parameters were diversified within the algorithm:

1. Best Play algorithm: k – which is the pool of strategies, a player has (strategies-per-agent).
2. Evolutionary algorithm: The mutation probability that allows an agent to explore the strategy space, by mutating known good strategies.
3. Roth-Erev algorithm: The parameters of *recency* (ϕ), and *exploration* (ϵ)

6.2 Strategy Diversity

We mix agents implementing different algorithms (strategies) into one population. The diversification here takes place on two levels, parameter-diversification (as before), as well as proportion of population implementing a particular strategy. The total population of agents in all scenarios is kept constant, but the proportion of agents playing a particular strategy is varied. This leads to different levels of diversification, for each combination of strategies. The parameters that were varied were as follows:

1. **Number of Algorithms:** This reflects the effective number of families (or *genera*) of species that are present in the population.
2. **Proportion of population implementing an algorithm:** Depending on the relative number of each specie, the aggregate reward in the game changes.
3. **Parameters within algorithms that were diversified:** Variation in critical parameters for an algorithm results in effectively creating a new specie, since the functional pathway (and resultant output) begins to change depending on the environment.

Table 2 shows the constant factors in each experiment. Each minority game was played with a population size of 501 (N) agents, through a simulation time period of 2000 steps (T). For each variation in the experimental setup, the data is reported as an average of 100 simulations.

Variable	Value
Population Size	501
Simulation Period (rounds)	2000
Repetition of Variation	100

Table 2: Experimental constants for Minority Game

6.3 Results

We combine the results from all of the diversification to illustrate the comparative rewards achieved at different levels of diversity. The graphs show the distribution of the population achieving a particular level of average reward. The ‘fatter’ the graph, the more evenly rewards are distributed, and vice-versa. Figure 1 clearly shows that the higher the diversity level, the fairer the average reward. Not only is the reward fairer, it is also higher, which is an important consideration for resource allocation strategies. Since it is the same three algorithms that are mixed in different proportions to achieve different levels of diversity, it is clear that diversity-level is the distinguishing factor in the level of rewards.

As was pointed out earlier in this section, MG represents a class of games where envy-freeness is impossible to achieve. After any given round, a majority of the agents

Fig. 1: Payoff across all diversification

are envious of the allocations secured by the minority. Algorithmic Diversification offers an alternate resolution, in such games, where multiple iterations of the game result in approximate envy-freeness. That is, on multiple iterations of an envious game, the introduction of algorithmic diversity, moves all the agents towards receiving approximately the same amount of reward. As seen in Figure 1, the increase in diversity results in increased fairness, which can be seen as a proxy for envy-freeness. This occurs due to a 'rotation' of agents that receive a reward, through the multiple iterations. The faster this rotation, in allowing agents that have previously not been rewarded to receive rewards, the greater the likelihood of envy-freeness.

7 Automated Generation of Diversity

In this section, 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 [101] 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. The process of clonal plasticity is briefly explained in the following sub-section.

7.1 Clonal Plasticity Process

Clonal Plasticity is an adaptation strategy, formed through the combination of two processes present in plants: (a) *clonal reproduction*, and (b) *phenotypic plasticity*. *Clonal Reproduction* refers to "... sequential reiteration of a basic structural unit or module consisting of the leaf with its associated auxiliary bud and a segment of stem that connects it to other units" [102]. The fundamental idea in this form of reproduction is that the daughter plant is genetically identical to the parent plant, and hence the term *clone*. *Phenotypic Plasticity* refers to the ability of an organism to vary its structure (height of stem, width of leaves, etc.) or behaviour (connect to its clonal sibling or disconnect) according to environmental influences. Importantly, this adaptation does not include variation in an organism due to genetic factors. We combine these two mechanisms to propose a self-adaptation mechanism called *clonal plasticity*. Clonal Plasticity is distinguished from other evolutionary mechanisms such as *genetic algorithms*, in that the genome of the organism is **not** modified at all. From a software system's perspective, this means that a clonally plastic software will deviate in some aspects of its behaviour, but **will functionally remain the same**.

The process of plastic reproduction, depicted in Figure 2, is given by the following steps:

1. **Identify Plasticity Points:** At a given moment in time, t , of the system's lifetime, each individual agent in the system (depicted a circle in Figure 2) identifies all its parts that can take multiple representations, values, or behaviour.
2. **Evaluate Environmental Input:** Simultaneously, the fitness of each individual agent is self-evaluated in the following way: for each plastic point, environmental feedback is evaluated to check whether it impedes or favours that point. Feedback from the environment may come from different sensors or neighbouring agents, or the agent's own reproduction memory (m in Figure 2).
3. **Plasticity Memory:** All agents that have reproduced before, keep a record of the previous Plasticity Range chosen during the last reproduction cycle (variable m in Figure 2). This allows an individual agent to repeat a good adaptation strategy, if it previously had a positive reward, or alternatively abandon a poor adaptation strategy.
4. **Choose Plasticity Range:** During the following time-step ($t+1$), each agent chooses a plastic response, based on the evaluated environmental input. The available responses may be one of:
 - (a) Exact Clone,
 - (b) Low Plasticity, and
 - (c) High Plasticity.

5. **Clone and Modify Plastic Points:** During the same time-step, each agent makes a clone of itself. The individual agent's plasticity points may be adapted according to the chosen strategy. The entire process starts again with a new time-step, t' , on all clones.

Fig. 2: Clonal plasticity process

7.1.1 Cloning Frequency:

Every adaptation action, in any system, incurs a cost, either computational or communicative. It must be possible to evaluate the effect of the cloning process, before creating the next plastic clone. $t + 1$ refers to the next time-step when it is feasible to perform the next cloning process. In the agent that implements plasticity, the granularity of time when the next iteration of the cloning process is done, is domain dependent. For instance, domains such as traffic and smart-grids are different in how dynamically their respective environments change. To allow for each adaptation to take effect, we introduce the notion of *cloning frequency*. Cloning Frequency refers to the time-period for which an agent is evaluated before the next iteration of the cloning process starts.

7.1.2 Choosing the Plasticity Range:

The process of choosing the plasticity range is the critical part in generating a diverse population. This process of modifying the plastic points considers both positive and negative feedback from the environment. For a given time instant t , the adaptation decision depends on two factors:

```

Data: PlasticSet, Memory, LastReward
Result: AlgorithmClone
1 Algorithm CreateClone(PlasticSet, Memory)
2   foreach Agent a do
3      $LastReward_a \leftarrow EvalEnviron(a, PlasticSet_a)$ 
4      $WhichRange_a \leftarrow ChoosePlasticityRange>LastReward_a, Memory_a)$ 
5     switch  $WhichRange_a$  do
6       case Exact-Clone do
7          $clone \leftarrow a$ 
8         return clone
9       case Low-Plastic-Clone do
10         $clone \leftarrow a$ 
11         $plasticity - point, direction \leftarrow Memory_a$ 
12         $ModifyAgent(clone, plasticity - point, direction)$ 
13        return clone
14       case High-Plastic-Clone do
15         $clone \leftarrow a$ 
16         $last-plasticity-point, last-direction \leftarrow Memory_a$ 
17         $plasticity-point \leftarrow chooseDifferent(PlasticSet_a, last-plasticity-point)$ 
18         $direction \leftarrow chooseDifferent(last - direction)$ 
19         $ModifyAgent(clone, plasticity - point, direction)$ 
20        return clone
21 Algorithm ChoosePlasticityRange(LastReward, Memory)
22   if  $LastReward < 0$  then
23     return Low-Plastic-Clone
24   else
25     if  $Memory > 0$  then
26       return Exact-Clone
27     else
28       return High-Plastic-Clone
29 Algorithm EvalEnviron(Agent a)
30 Algorithm ModifyAgent(Agent a,  $plasticity - point, direction$ )

```

Fig. 3: Create a plastic clone of an algorithm

1. Feedback from environment at time t , whether it is positive or negative.
2. Memory of the modification action, taken at time $t-1$.

If the feedback from the environment is positive, and memory of the previous action is also positive, then the Plasticity Range is chosen to be *Exact Clone*. If the feedback from the environment is positive, but the memory of the previous action is negative, then the Plasticity Range is chosen to be *Low Plasticity*. If the feedback from the environment is negative, then the Plasticity Range is always *High Plasticity*.

7.1.3 Exact Clone

In this case, the individual simply makes a copy of itself, with no change at all.

7.1.4 Low Plasticity

Modification with Low Plasticity implies that the individual chooses the same plasticity point, as chosen during the previous cloning process, and moves in the same direction as before.

7.1.5 High Plasticity

Modification with High Plasticity implies that the individual chooses a different plasticity point, than was chosen during the previous cloning process, and moves in a random direction.

7.2 Clonal Plasticity for MG

Here, we modify the agents, such that each agent is able to reproduce in a clonally plastic way. Cloning, as opposed to other evolution-based mechanisms, is completely decentralized as there is no universal fitness function that sorts the population in to *fit* and *unfit* individuals.

Making an algorithm plastic requires following the process of plastic reproduction

1. **Identification of Plasticity Points:** BestPlay has only one variable that influences its decision-making, *i.e.*, memory. Hence, there is only one plasticity point for agents implementing BestPlay. Similarly, for the Evolutionary strategy, there is one variable that influences how the decision making varies within the population, *i.e.*, mutation. Roth-Erev has two variables that can be plasticized, *i.e.*, recency and epsilon
2. **Evaluate Environmental Input:** The environment allows each agent to know how well its algorithm is doing. That is, depending on the performance of the agent in successfully choosing the minority group, the agent would have either get positive or negative feedback. In the context of MG, cloning frequency refers to the number of times an agent plays the game, before its performance is evaluated. Thus, a cloning frequency of 5 means that an agent's performance is evaluated after every 5 rounds. Clearly, the lower this number, the faster the agent reacts to the environmental feedback. Agents playing BestPlay or Roth-Erev were evaluated with cloning frequency ranging from 5 to 105 rounds. Agents playing the Evolutionary strategy were evaluated with cloning frequency ranging from 20 to 100. This difference is due to the fact that preliminary experimentation revealed that the Evolutionary algorithm required 20 rounds to produce a noticeable effect in an agent's performance. The effect of cloning frequency on diversity generation is discussed in section 9.2. The memory (m) of the previous algorithm parameters allows an agent to return to a previously known good configuration, if the current modification resulted in losses.
3. **Choose Plasticity Range:** The plasticity range chosen by a particular agent defines how the agent decides to adapt. Each of the algorithms (BestPlay, Evolutionary, Roth-Erev), depending on its choice of parameters, decides how the agent

Feedback	Plasticity Memory	Plasticity Range Chosen
Negative	(Disregarded)	High Plasticity
Positive	wins-this-cycle $\leq$ wins-last-cycle	Low Plasticity
Positive	wins-this-cycle $>$ wins-last-cycle	Exact Clone

Table 3: Criteria for choosing plasticity range

would play upcoming rounds of MG. Any change of these parameters could cause the agent to change its decision, and thereby affect the rewards that it accumulates.

7.3 Worked Out Example

Now we perform the clonal plasticity process step-by-step for BestPlay. For details on the clonal plasticity process for the Evolutionary and RothErev algorithms, see appendices A and B (the process differs only in the plasticity points identified for each algorithm).

1. **Identify Plasticity Points:** The BestPlay algorithm critically depends on the size of the memory, that it keeps of the game (see Table 1). Two agents that diversify into different directions (one increasing the memory size, and the other decreasing) could end up with drastically different rewards. Here, m is identified as a plasticity point of the algorithm. Another plasticity point is the number of strategies (k) that the agent possesses.
2. **Evaluate Environmental Input:** For any given time instant, the algorithm has a 50% chance of picking the right group (*i.e.*, being in the minority). Hence, for any given period of cloning frequency, if the agent has accumulated (wins-this-cycle) more than 50% of the rewards during the period, the feedback is said to be positive. If the agent accumulates less than 50% of the rewards, then the feedback is negative.
3. **Plasticity Memory:** The agent keeps track of the rewards that it had accumulated during the previous adaptation (wins-last-cycle). Along with the feedback, this helps it decide what plasticity range to choose next (see Table 3).
4. **Clone and Modify Plastic Points:** In the event of *Exact Clone*, the agent does not change any of its parameters. In the event of *Low Plasticity*, the agent picks one of its plastic points and modifies it in the same direction, as the plasticity memory from the previous generation. In the event of *High Plasticity*, the agent modifies *all* its plasticity points to return to the values that it had from the previous generation.

8 Requirements For Distributive Justice

In section 5.1, we saw that efficiency in allocations would only occur, if $\min G = N/2$. However, it is entirely possible that in repeated iterations of the game, the same sub-

set of N agents continue to be a part of $\text{min}G$. This would mean that the allocations while efficient, are unfair, leading to envy amongst the agents. Ideally, we would like *envy-freeness* to be a property of clonal plasticity. While *envy-freeness* is an important concern in many domains, there are some areas, specifically in simultaneous games, where envy-freeness is impossible [16]. Balkanski *et al.* [16] show that for any resource allocation protocol which involves simultaneous choice by more than 3 agents, it is impossible to construct an envy-free protocol. This impossibility is intuitive in the case of the Minority Game where, by construction, during every round, a majority of the agents are denied a reward and are envious of the minority.

In view of the impossibility of an envy-free protocol ever being created, we focus on an alternate form of fairness, *i.e.*, *distributive justice*.

Distributive justice, in the context of MG, implies the following condition:

$$\sum_{t=1}^T R_t(a_j) \approx \sum_{t=1}^T R_t(a_k) \forall (a_k \neq a_j) \quad (2)$$

Equation 2 simply states that, over a period T , every agent (a_j) in MG must get approximately the same reward ($R(a_j)$), as any other agent (a_k). In the ideal case, every agent would get **exactly** the same reward as every other agent. However, this can be trivially shown to be impossible under the following conditions of agent abilities:

1. *Autonomy*: Each agent is autonomous and cannot be coerced by any other agent. The agent's choice is dependent solely on its own algorithm.
2. *Simultaneity*: Each agent makes its choice in a simultaneous fashion, and no agent has access to any other agent's decision before it makes its own.

Proof:

Assume a set of algorithms A that are implemented in the agents playing MG. If this set A meets two conditions, then MG would be both, efficient and fair:

1. The size of the minority is the largest possible, while still being a minority, and
2. The members of the minority group, in any iteration of MG, have never been in a minority before

then, after two iterations, there will be only one agent that has never been in a minority before. Therefore, in iteration 3, the sole agent participates in the minority group, and thus meets the game is both, efficient and fair.

For any mechanism, meeting condition 1 requires that the size of the $\text{min}G$ be $N/2$ for iteration 1 **and** iteration 2. Meeting condition 2 requires that the list of agents that have already participated in $\text{min}G$ be known at all times, and any agent that has previously been a part of $\text{min}G$ be restricted by some mechanism. Meeting condition 2 violates the environmental assumption of *Autonomy*, while meeting condition 1 violates our environmental assumption of *Simultaneity*. Thus, there is no way for any mechanism to assure equality of rewards and be efficient, without violating our operating conditions.

8.1 Desirable Properties of Clonal Plasticity:

It is evident that equation 2 would be impossible to achieve with MG, if it were played in a one-shot manner. That is, for a single round of play, regardless of the algorithm chosen, the agents that are in the minority will get a reward, while the agents in the majority will not. This means that equation 2 will never hold. However, if the same agents play MG in a repeated manner, it is probable that different agents are in a minority in other rounds, such that their rewards accumulated over time would be approximately equal. That is:

$$\lim_{T \rightarrow \infty} \sum_{t=1}^T R_t(a_j) \approx \sum_{t=1}^T R_t(a_k) \forall (a_j \neq a_k) \quad (3)$$

Recall, that we are not interested solely in distributive justice, but also that the mechanism should be efficient. In the case of MG, the system would be efficient if allocation of rewards is as high as possible in any given round (see subsection 5.1). At its most efficient, the game provides the highest amount of reward (when the size of *minG* is $N/2$):

$$R_1 = \frac{N}{2} \quad (4)$$

This implies that after T rounds of highest possible efficiency, the total rewards in the system must be:

$$\sum_{t=1}^T R_t = \frac{N}{2} * T \quad (5)$$

From equations 4 and 5, it is clear that at optimal efficiency, the expected rewards accumulated by any single agent is given by:

$$\begin{aligned} \sum_{t=1}^T R_t(a_j) &= \frac{N/2 * T}{N} \\ &= \frac{T}{2} \end{aligned} \quad (6)$$

Thus, we would like a mechanism to achieve **both** equation 3 and 6.

9 Experimental Setup and Results

Table 4 shows the constant factors in each experiment³. 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. We measure the efficiency of the system through the median amount of rewards accumulated by the agent, after the simulation period. We

³ All code for this experiment can be found at:
https://bitbucket.org/viveknallur/clonal_plasticity_algo_diversity.git

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 [103]. 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.

Variable	Value
Population Size	501
Simulation Period (rounds)	2000
Repetition of Variation	100

Table 4: Experimental constants - Same as manual process of diversification

Figures 4 and 5 show the results of diversification introduced by clonal plasticity. As is immediately obvious from the x-axis (that shows diversity generated), manual diversification is able to generate a higher level of diversity in the populations of agents. Clonal Plasticity is able to reach a diversity level (also called H-index) of upto 2.7, whereas manual diversification (see Figure 1) is able to reach an H-index of upto 5.8. This is due to the fact that a human (knowing the functional pathways of a particular algorithm) is able to specify diverse individuals that will generate diverse outputs. However, in terms of efficiency and equity, clonal plasticity confirms the results found by *Nallur et al.* [19]. Note that as the diversity goes up, the median reward accumulated by the population also increases (Figure 4). At the lower end of the diversity scale ($H = 1.7$), the median reward is 509, while at the higher end of the scale ($H = 2.7$), the median reward is 867. Increasing the diversity of the population increases the reward achieved by each agent. Figure 5 shows the effect on equity of rewards achieved. The y-axis shows the Gini index, which is a quantitative measure of equality of reward. In Figure 5, the increase in diversity results in a decrease in the Gini index (and therefore, a fairer distribution). *Note:* Both Figures 4 and 5 are shown with error bars that indicate the 95% confidence intervals. The medians were established using 100 experimental trials for each variation (7 variations in algorithms [see Table 6]), and the confidence intervals for each diversity value across the 700 trials were established by bootstrapping at 100 iterations.

9.1 Effectiveness of Diversity in Fairness and Efficiency

Recall from Section 8, the theoretical optimum efficiency that we would like is $T/2$. That is, the accumulated reward at the end of the simulation should be equal to half the number of rounds of simulation. In all our experiments, the T is kept as a constant 2000, hence the optimum efficiency is 1000. We see from Figure 4, that increasing diversity results in a increased median reward. At the highest diversity level ($H = 2.7$), the accumulated reward is 867, while at the lowest level of ($H = 1.7$), it is 509.

Note also from Table 5, 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

Fig. 4: Impact of diversity on Reward

Fig. 5: Impact of diversity on Fairness

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 5: Relationship between Diversity and Reward and Fairness

Algorithm	Diversity	Num_Species
BestPlay	1.7	60
Evo	2.7	501
RothErev	2.7	486
BestPlay-Evo	2.1	271
RothErev-BestPlay	2.2	303
RothErev-Evo	2.7	493
Evo-RothErev-BestPlay	2.5	419

Table 6: Diversity and Number of Species

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) 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 (Gini index of 2.7) is certainly better than the fairness achieved with low diversity (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. Automatically deriving the pareto curve of diversity, for a particular system, would be an interesting avenue of further research.

9.2 Evaluating Clonal Plasticity in Generation of Diversity

We now look at Clonal Plasticity itself and evaluate how the cloning frequency influences the amount of diversity achieved within the population. Cloning frequency refers to how many rounds (of time) elapse before the cloning process is run. Intuitively, the fewer the number of rounds between each cloning process, the higher the amount of diversity that will be achieved. This is due to the fact that each run of the cloning process has the potential to create a new specie. This is borne out by the data in Table 6 which illustrates that an increase in mean diversity has a direct correspondence with an increase in the mean number of species generated by the cloning process. Figure 6 shows a population-wide distribution of the same analysis

In Figure 6, we see the distribution of diversity across the entire population as bubbles (y -axis), with change in cloning frequency (x -axis). The size of the bubbles represents the number of species. As we go up the y -axis, notice that the size of the bubbles increases. This indicates that the greater the diversity, the greater the number of species generated. Also, the lower the number of rounds between each cloning process (*i.e.*, the lower the frequency), the higher the mean amount of diversity produced. Table 9 in Appendix C shows the full breakdown of the number of species produced, as the cloning frequency varies, per algorithm. It is interesting to note that

Fig. 6: Diversity, cloning frequency and number of species together

the cloning frequency does not, by and large, affect the diversity achieved *within* an algorithmic group.

9.3 Effect of Plasticity on Different Algorithms

Figures 7 and 8 show an algorithm-wise breakup of differences in reward achieved as well as the Gini index. Read in conjunction with Figure 9, which shows that three algorithms (Evo, RothErev and RothErev-Evo) achieved the highest diversity, it also shows that high levels of diversity are linked to high levels of reward. Table 9, from Appendix C shows how this can be further broken down based on how rapidly the cloning process operates.

Consider Figure 7 and 8 along with Table 5 and 6. Both BestPlay-Evo and RothErev-BestPlay achieve a mean Gini-index very close to zero (1.0 and 1.3), however it is the combination of Evo-RothErev-BestPlay that achieves the best Gini-index (0.9) and as well as a very high level of mean reward (780). Evo-RothErev-BestPlay performs very well across both indicators of reward and Gini. However, if one considers only rewards, it is RothErev with the highest diversity level of 2.7, that has the highest reward level. In contrast, Evo is another algorithm that reached a high level of diversity; almost all its individuals consistently formed separate species (see Figure 9) and a mean Gini-index very close to zero. However, an examination of Figure 8 shows that the distribution of Gini values fluctuated quite a bit (see Table 9 from Appendix C for a cloning-frequency-wise breakdown), indicating that it is sensitive to cloning-frequency values. Taken together, the results show that raising diversity

Fig. 7: Algorithm-wise breakup of Reward achieved

Fig. 8: Algorithm-wise breakup of Gini Index

levels indiscriminately is not a silver bullet for achieving fairness, however reward levels rise in congruence with diversity.

9.4 Scalability of Clonal Plasticity

Although the clonal plasticity process is effective at generating diversity, its scalability is an important factor in deciding whether it is feasible for implementation in real systems.

Before the cloning process can begin, the identification of the plasticity points must be done. This is a one-time function and needs to be done only at the initializa-

Fig. 9: Algorithm-wise breakup of number of species at the end of cloning

tion of the plasticity process. For any algorithm a , with k interface variables, the first step is to identify whether those variables modify the output of the algorithm. This is clearly $O(k)$ in complexity. The second major step is to identify the plasticity range for each plasticity point. This involves identifying the valid set of values a plasticity point may take, as well as the constraints on its value. Assuming a time complexity of r per variable, the first phase of clonal plasticity is $O(r * k)$ in complexity.

Recall that the `ChoosePlasticityRange` and `CreateClone` functions form the core of the plasticity process (see Algorithm 3). `ChoosePlasticityRange` is a linear function ($O(c)$) since it involves merely decision-making, based on the `LastReward` and `Memory` variables. `CreateClone` involves looking at the value of each plasticity point, and assuming a time complexity of p per plasticity point, its complexity can be calculated as $O(p * k)$. The computational complexity of the clonal plasticity process is therefore: $O(r * k) + O(c) + O(p * k)$. This is important because, if the adaptation process is computationally expensive, then its benefits are limited to systems with few plasticity points. The linear complexity of clonal plasticity implies that the mechanism is scalable to large-scale multi-agent systems, with a high number of plasticity points.

The `EvalEnviron` is a domain-specific function to evaluate the current performance of the agent, and is therefore left unspecified. Any adaptation mechanism would require an `EvalEnviron` function and as such, it is a price that all adaptation functions must pay for sensible self-modification.

Crucially, it is important to note that the entire process works on a single agent and is independent of other agents' performance, in any direct fashion. This allows the plasticity mechanism to work in a completely decentralized manner. Of course, the presence of other agents affects the evaluation of the environment (in an indirect fashion via their actions), and thus provides local competitive pressure that moves the agent to a better configuration.

10 Differences between Clonal Plasticity and Phenotypic Diversity-Based Mechanisms

In the presence of large search spaces for adaptation, there have been a number of *intelligent trial-and-error* based algorithms that have been proposed. These differ from the more commonly known EA/GA/GP methods in that there is an explicit absence of an optimizing function. Notable examples include techniques such as Novelty Search [104], Functional Blueprints [105] and more recently, MAP-Elites [106], SHINE [107], etc. Among these techniques, there is an acknowledgement of the fundamental importance of phenotypic diversity, similar to clonal plasticity. The key difference between these methods and clonal plasticity is *pre-computation* of phenotypically diverse behaviours. For instance, Multi-dimensional Archive of Phenotypic Elites (MAP-Elites) [106] creates a number of discrete behaviour bins, which contain the fittest individual mapped to that bin so far. The bins are then sampled at random to produce an offspring and mapped to a bin depending on its behaviour characterization. The individual is then saved or culled depending on its relative fitness to other individuals in the bin. In this way, due to uniform sampling, the algorithm passively acquires diversity [108]. SHINE [107] uses the same fundamental idea, but uses a tree-structure for maintenance and selection of potential individuals. This, along with a technique of penalising crowded areas, allows the algorithm to explore the phenotypic landscape more rapidly than MAP-Elites.

Real-world domains that have been explored by these techniques include bi-pedal walking by a robot, walking by hexapod robot and damaged robotic arms. A peculiarity of these domains is that the adaptation required cannot wait for *trial-and-error* to conclude, and hence pre-computation is a necessity for such domains. Pre-computation requires simulation and a thorough exploration of the behaviour space, before the final solutions are transferred to the hardware device. Clonal Plasticity does not do any pre-computation and hence is unsuitable for domains that require a rapid response.

11 Differences between Clonal Plasticity and Non-Phenotypic Diversity-Based Mechanisms

Clonal Plasticity is a meta-heuristic, in the same vein as the other methods in evolutionary computation. However, it differs in significant ways from the diversity generation approaches presented before, in section 3. These axes affect how suitable an approach is, for a specific domain:

1. **Representation**(Neutral or Dependent): If the diversification relies on the distance-measure adopted, then the representation of the individual becomes a critical factor in calculating diversity.
2. **Temporal Dependence** (Yes or No): Does the diversification depend on *when* the differences between individuals is calculated. For example, methods such as *incest prevention* and *mating restrictions* require that the diversity between individuals be calculated **before** every new generation. Similarly, the ensemble mech-

anism also requires that different values for ensemble training are chosen (from a poisson distribution) before the training starts.

3. **Locus of Control** (Yes or No): Does the mechanism involve calculating the diversity of the entire population in a centralized fashion or can different subsets diversify at different rates?

11.1 Representation

Evolutionary Algorithms (EAs), Genetic Algorithms (GAs), Genetic Programming (GP) typically use a distance measure in the genotype of the individual to force diversity into the population. For example, Sareni and Krähenbühl utilize the notion of fitness sharing and creation of niches to ensure that the population retains diversity, after multiple generations [109]. However, this mechanism quickly becomes computationally intractable when using genotype representations such as Artificial Neural Networks (ANNs) [110]. Behavioural diversity based approaches such as novelty search [104] avoid the problems of genotype-based representations.

The traditional particle-swarm [75] uses a tuple of D-dimensional vectors (where D is the dimensionality of the search space), position and fitness of best point, as well as index of best particle to represent a particle in the swarm. Diversity generation mechanisms (*e.g.*, [79, 80]) continue to use the tuple as a standard representation of a particle.

Machine Learning approaches such as ensembles are independent of any particular representation.

Clonal Plasticity does not have (or require) a particular representation of an agent. The agent must merely be able to reflect on, and tune its own parameters.

Representation is an important aspect during development, since many application domains where multi-agent systems are used (modelling electricity grids, transportation systems, etc.) do not lend themselves to arbitrary representations. Therefore, for agents that need to model domain-specific behaviour it would be more convenient to use employ either ensemble-based diversity generation or clonal plasticity.

11.2 Temporal Dependence

EA/GA/GP have a high dependence on the temporality of the diversity calculation. Due to their inherent selection-pressure based nature, diversity mechanisms must be deployed at every generation, to ensure that the selection pressure does not lead to a homogeneous/convergent population.

Swarm-based methods are also dependent on temporality of diversity calculation, since every particle's next-step will influence whether the resulting swarm consists of

a dense neighbourhood(s) or not.

Machine Learning mechanisms such as ensemble-learning are not dependent on generations, however diversity is induced during the training phase by using variation in the examples [37]. However, other mechanisms such as evolutionary modification of neural networks [81] require a generation-by-generation diversity calculation.

Clonal Plasticity belongs to a class of methods which do not explicitly force selective pressure towards diversity [111], rather diversity is emergent from the self-adaptation behaviour.

The need to ensure diversity at every generation makes EA/GA/GP/swarm-based approaches unsuitable for a dynamic environment that requires some action from the system-under-consideration. Unlike optimization problems where the final calculated solution is the only determiner of quality, other domains (such as human-facing systems) will often need software to respond, even if sub-optimally. In such situations, clonal plasticity is the best way to ensure that the system continues to function, while moving towards more diversity.

11.3 Locus of Control

EA/GA/GP use methods that calculate genotype-distance require the entire population to be available for diversity calculations. Even the behaviour-based novelty search creates a distance-measure among behaviours that have been observed in the population, and hence require computation over the entire population.

Swarm-based methods operate on each particle at a time, however, neighbourhood measurements of density require computations over subsets of the population, which make diversification more of a centralized process than a decentralized one.

Machine Learning techniques all explicitly aim for diversity during their training phase. This requires that after a training datum, the feedback and subsequent training take the current homogeneity of all available learners into account.

Since clonal plasticity does not explicitly aim towards diversity, the diversity seen in the population is an emergent phenomenon from the actions of the individual agents. Each agent does, however, perform a neighbourhood-based calculation of its relative fitness. This makes clonal plasticity not a fully decentralized process, rather a distributed process.

No approach that we know of is able to achieve a fully decentralized operation to ensure diversity. The best achieved (so far) reduces computation to small subsets, instead of the entire population. Decentralization is required where the system-under-consideration is a large-scale system, and using centralized operations are infeasible.

	EA / GA / GP	Swarm-Based Approaches	Machine Learning	Clonal Plasticity
Representation	Distance measure based on genotype	Tuple-based representation of particle	Independent	Independent
Temporal Dependence	Diversity is calculated after every generation	Diversity is calculated after every generation	Ensemble-learning dependent on training phase	Not dependent
Locus of Control	Centralized	Centralized	Centralized	Distributed

Table 7: Contrast between diversity generation approaches

ble. Examples of such systems include vehicular ad-hoc networks [3, 4], smart buildings [5], e-procurement [6], cloud computing [7], healthcare systems [8].

12 Weaknesses of Clonal Plasticity

Domain Limitations: Clonal Plasticity is an intelligent trial-and-error mechanism that is able to dynamically adjust to a changing environment by exhibiting phenotypic diversity. However, in certain domains such as robotics or autonomous vehicles the rate of response required from the system is quite fast. Due to real-time response requirements, other phenotypic diversity mechanisms such as MAP-Elites [106] use *pre-computation* to explore the possible behaviour space before deployment. Clonal Plasticity does not use pre-computation and hence is slow to jump from one known behaviour to another. Domains requiring such fast adaptation responses medical devices, transportation, and industrial robotics are not suitable for clonal plasticity.

Algorithmic Limitations: Clonal plasticity is a meta-algorithm that achieves diversity regardless of the underlying problem-specific algorithm being implemented. However, the diversity-seeking behaviour makes it unlikely that clonal plasticity can be used to achieve optimality in any search space. This diversity, unlike others such as Novelty Search [104] and SHINE [107], is an emergent characteristic of the interaction between plasticity and the environment. This implies that if by some chance, clonal plasticity hits on the optimal behaviour, it is likely to stay at that point. However, it does not actively seek out optimal behaviour. By contrast, EA/GA/GP have a well-established record of optimization with both, theoretical analysis [112] as well as in practical domains [113, 114]. Hence, clonal plasticity is unsuited for situations that require optimization.

13 Conclusion

In ecological sciences, as well as computer science, diversity has been shown to be an important property of an ecosystem. As pointed out in Section 2, the concept of diversity has 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. The *Minority Game* has been shown to be a simple, yet good proxy for socio-technical systems such as traffic [115], financial markets [93], distributed decision-making in wireless networks [94], smart-buildings [95], etc. Results shown in the *Minority Game* setting indicate that diversity can play an important role in affecting system properties such as efficiency and fairness. We do not wish to suggest that diversity is a silver bullet for achieving optimal levels of both, but rather that it is an important mechanism in a software engineer's toolkit while engineering socio-technical systems.

13.1 Open Research Questions

Although clonal plasticity has been shown to lead to successful diversification, it would be interesting to know if there are limits to how quickly such diversification can occur. An investigation into the formal properties of the clonal plasticity process would be invaluable in determining its suitability for time-critical systems. Are there properties of the constituent algorithms or the number of plasticity points that can predict the speed of, or magnitude of adaptive movement? Can we automatically identify the pareto space of the system, where increasing diversity does not necessarily result in better values of the desired system properties? All of these are subjects of future research that we think will contribute greatly to a more detailed and nuanced understanding of this adaptive strategy.

13.2 Concluding Remarks

Multi-agent systems are increasingly being used to model complex socio-technical systems that involve human interactions. In such systems, it is difficult to simultaneously guarantee properties of efficiency and fairness. This paper presents a measurable system property (diversity) that can be engineered to produce high levels of both. Further, it presents 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. G. Baxter and I. Sommerville, "Socio-technical systems: From design methods to systems engineering," *Interacting with Computers*, vol. 23, no. 1, pp. 4–17, 2011. [Online]. Available: <http://iwc.oxfordjournals.org/content/23/1/4.abstract>
2. A. H. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, "Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid," *IEEE Transactions on Smart Grid*, vol. 1, no. 3, pp. 320–331, Dec 2010.
3. M. Ding, X. Cheng, and G. Xue, "Aggregation tree construction in sensor networks," in *Vehicular Technology Conference, 2003. VTC 2003-Fall. 2003 IEEE 58th*, vol. 4, Oct 2003, pp. 2168–2172 Vol.4.

4. S. S. Manvi, M. S. Kakkasageri, and J. Pitt, "Multiagent based information dissemination in vehicular ad hoc networks," *Mob. Inf. Syst.*, vol. 5, no. 4, pp. 363–389, Dec. 2009. [Online]. Available: <http://dx.doi.org/10.1155/2009/518042>
5. J. H. Yoon, R. Baldick, and A. Novoselac, "Dynamic demand response controller based on real-time retail price for residential buildings," *IEEE Transactions on Smart Grid*, vol. 5, no. 1, pp. 121–129, Jan 2014.
6. F. Wu, G. Zsidisin, and A. Ross, "Antecedents and outcomes of e-procurement adoption: An integrative model," *IEEE Transactions on Engineering Management*, vol. 54, no. 3, pp. 576–587, Aug 2007.
7. V. Nallur and R. Bahsoon, "Self-adapting applications based on qd requirements in the cloud using market-based heuristics," in *Towards a Service-Based Internet*, ser. Lecture Notes in Computer Science, E. Di Nitto and R. Yahyapour, Eds. Springer-Verlag, December 2010, pp. 51–62.
8. F. Axisa, P. M. Schmitt, C. Gehin, G. Delhomme, E. McAdams, and A. Dittmar, "Flexible technologies and smart clothing for citizen medicine, home healthcare, and disease prevention," *IEEE Transactions on Information Technology in Biomedicine*, vol. 9, no. 3, pp. 325–336, Sept 2005.
9. A. S. Masoum, S. Deilami, P. S. Moses, M. A. S. Masoum, and A. Abu-Siada, "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*, vol. 5, no. 8, pp. 877–888, August 2011.
10. F. Emery and E. Trist, "Socio-technical systems," *Management Science, Models and Techniques*, vol. 2, pp. 83–97, 1960.
11. M. Wooldridge, "Intelligent agents: The key concepts," in *ECCAI Advanced Course on Artificial Intelligence*. Springer, 2001, pp. 3–43.
12. J. Ladyman, J. Lambert, and K. Wiesner, "What is a complex system?" *European Journal for Philosophy of Science*, vol. 3, no. 1, pp. 33–67, 2013. [Online]. Available: <http://dx.doi.org/10.1007/s13194-012-0056-8>
13. J. Pitt, D. Busquets, and R. Riveret, "The pursuit of computational justice in open systems," *AI & SOCIETY*, vol. 30, no. 3, pp. 359–378, 2015. [Online]. Available: <http://dx.doi.org/10.1007/s00146-013-0531-6>
14. Y. Chevalere, U. Endriss, J. Lang, and N. Maudet, *A Short Introduction to Computational Social Choice*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 51–69. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-69507-3_4
15. Y. Chevalere, P. E. Dunne, U. Endriss, J. Lang, J. Lemaitre, J. Maudet, Nicolas Padget, S. Phelps, J. A. Rodríguez-aguilar, and S. P., "Issues in multiagent resource allocation," *Informatica*, vol. 30, no. 2, pp. 3–31, Jun. 2006. [Online]. Available: <http://dx.doi.org/10.1017/S0269888905000470>
16. E. Balkanski, D. Kurokawa, S. Brânzei, and A. D. Procaccia, "Simultaneous cake cutting," in *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, ser. AAAI'14. AAAI Press, 2014, pp. 566–572. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2893873.2893962>
17. N. Rescher, *Distributive Justice: A Constructive Critique of the Utilitarian Theory of Distribution*. New York: The Bobbs-Merrill Co. Inc., 1966.
18. N. Rescher, *Fairness: Theory and Practice of Distributive Justice*, N. Rescher, Ed. New Brunswick, USA: Transaction Publishers, 2002.
19. V. Nallur, E. O'Toole, N. Cardozo, and S. Clarke, "Algorithm Diversity: A Mechanism for Distributive Justice in a Socio-Technical MAS," in *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, ser. AAMAS '16. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2016, pp. 420–428. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2936924.2936986>
20. E. P. Odum, *Fundamentals of Ecology*, 5th ed. Cengage Learning, 2004.
21. R. MacArthur, "Fluctuations of animal populations and a measure of community stability," *Ecology*, vol. 36, pp. 533–536, 1955.
22. J. Lawton and V. Brown, *Redundancy in Ecosystems*, E. Schulze and H. Mooney, Eds. Springer, 1993.
23. S. Naeem and S. Li, "Biodiversity enhances ecosystem reliability," *Nature*, vol. 390, pp. 507–509, September 1997.
24. S. Naeem, "Species redundancy and ecosystem reliability," *Conservation Biology*, vol. 12, no. 1, pp. 39–45, 1998.
25. A. Avizienis, "The methodology of n-version programming," *Software fault tolerance*, vol. 3, pp. 23–46, 1995.

26. F. Cohen, "Operating system protection through program evolution," *computers and security*, vol. 12, no. 6, pp. 565–584, October 1993.
27. G. S. Kc, A. D. Keromytis, and V. Prevelakis, "Countering code-injection attacks with instruction-set randomization," in *Proceedings of the 10th ACM conference on Computer and communications security*. ACM, 2003, pp. 272–280.
28. C. Giuffrida, A. Kuijsten, and A. S. Tanenbaum, "Enhanced operating system security through efficient and fine-grained address space randomization," in *Proceedings of the 21st USENIX Conference on Security Symposium*, ser. Security'12. Berkeley, CA, USA: USENIX Association, 2012, pp. 40–40. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2362793.2362833>
29. K. Z. Snow, F. Monrose, L. Davi, A. Dmitrienko, C. Liebchen, and A.-R. Sadeghi, "Just-in-time code reuse: On the effectiveness of fine-grained address space layout randomization," in *Proceedings of the 2013 IEEE Symposium on Security and Privacy*, ser. SP '13. Washington, DC, USA: IEEE Computer Society, 2013, pp. 574–588. [Online]. Available: <http://dx.doi.org/10.1109/SP.2013.45>
30. E. G. Barrantes, D. H. Ackley, T. S. Palmer, D. Stefanovic, and D. D. Zovi, "Randomized instruction set emulation to disrupt binary code injection attacks," in *Proceedings of the 10th ACM Conference on Computer and Communications Security*, ser. CCS '03. New York, NY, USA: ACM, 2003, pp. 281–289.
31. E. Shioji, Y. Kawakoya, M. Iwamura, and T. Hariu, "Code shredding: Byte-granular randomization of program layout for detecting code-reuse attacks," in *Proceedings of the 28th Annual Computer Security Applications Conference*, ser. ACSAC '12. New York, NY, USA: ACM, 2012, pp. 309–318. [Online]. Available: <http://doi.acm.org/10.1145/2420950.2420996>
32. L. V. Davi, A. Dmitrienko, S. Nürnberger, and A.-R. Sadeghi, "Gadge me if you can: Secure and efficient ad-hoc instruction-level randomization for x86 and arm," in *Proceedings of the 8th ACM SIGSAC Symposium on Information, Computer and Communications Security*, ser. ASIA CCS '13. New York, NY, USA: ACM, 2013, pp. 299–310. [Online]. Available: <http://doi.acm.org/10.1145/2484313.2484351>
33. L. Szekeres, M. Payer, T. Wei, and D. Song, "Sok: Eternal war in memory," in *Proceedings of the 2013 IEEE Symposium on Security and Privacy*, ser. SP '13. Washington, DC, USA: IEEE Computer Society, 2013, pp. 48–62. [Online]. Available: <http://dx.doi.org/10.1109/SP.2013.13>
34. M. Lozano, F. Herrera, and J. R. Cano, "Replacement strategies to preserve useful diversity in steady-state genetic algorithms," *Information Sciences*, vol. 178, no. 23, pp. 4421 – 4433, 2008, including Special Section: Genetic and Evolutionary Computing. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0020025508002867>
35. F. Han and Q. Liu, "A diversity-guided hybrid particle swarm optimization based on gradient search," *Neurocomputing*, vol. 137, pp. 234 – 240, 2014, advanced Intelligent Computing Theories and Methodologies Selected papers from the 2012 Eighth International Conference on Intelligent Computing (ICIC 2012). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0925231214002665>
36. L. Vitorino, S. Ribeiro, and C. Bastos-Filho, "A mechanism based on artificial bee colony to generate diversity in particle swarm optimization," *Neurocomputing*, vol. 148, pp. 39 – 45, 2015. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0925231214009412>
37. L. Minku and X. Yao, "Ddd: A new ensemble approach for dealing with concept drift," *Knowledge and Data Engineering, IEEE Transactions on*, vol. 24, no. 4, pp. 619–633, April 2012.
38. K. Brinker, "Incorporating diversity in active learning with support vector machines," in *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, 2003, pp. 59–66.
39. L. I. Kuncheva and C. J. Whitaker, "Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy," *Machine learning*, vol. 51, no. 2, pp. 181–207, 2003.
40. S. A. Hofmeyr and S. Forrest, "Immunity by design: An artificial immune system," in *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation-Volume 2*. Morgan Kaufmann Publishers Inc., 1999, pp. 1289–1296.
41. D. Dasgupta, "Advances in artificial immune systems," *IEEE computational intelligence magazine*, vol. 1, no. 4, pp. 40–49, 2006.
42. K. Pohl, G. Böckle, and F. J. v. d. Linden, *Software Product Line Engineering: Foundations, Principles and Techniques*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2005.
43. J. G. Wijnstra, "Supporting diversity with component frameworks as architectural elements," in *Software Engineering, 2000. Proceedings of the 2000 International Conference on*, 2000, pp. 51–60.
44. M. A. Hiltunen, R. D. Schlichting, C. A. Ugarte, and G. T. Wong, "Survivability through customization and adaptability: the cactus approach," in *DARPA Information Survivability Conference and Exposition, 2000. DISCEX '00. Proceedings*, vol. 1, 2000, pp. 294–307 vol.1.

45. J. C. Knight, "Dependable and historic computing," C. B. Jones and J. L. Lloyd, Eds. Berlin, Heidelberg: Springer-Verlag, 2011, ch. Diversity, pp. 298–312. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2184121.2184147>
46. I. Schaefer, R. Rabiser, D. Clarke, L. Bettini, D. Benavides, G. Botterweck, A. Pathak, S. Trujillo, and K. Villela, "Software diversity: state of the art and perspectives," *International Journal on Software Tools for Technology Transfer*, vol. 14, no. 5, pp. 477–495, 2012. [Online]. Available: <http://dx.doi.org/10.1007/s10009-012-0253-y>
47. B. Baudry and M. Monperrus, "The multiple facets of software diversity: Recent developments in year 2000 and beyond," *ACM Comput. Surv.*, vol. 48, no. 1, pp. 16:1–16:26, Sep. 2015. [Online]. Available: <http://doi.acm.org/10.1145/2807593>
48. R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode, *Selfish Routing in Non-cooperative Networks: A Survey*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 21–45. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-45138-9_2
49. L. Castelli, R. Pesenti, and A. Ranieri, "The design of a market mechanism to allocate air traffic flow management slots," *Transportation Research Part C: Emerging Technologies*, vol. 19, no. 5, pp. 931 – 943, 2011, freight Transportation and Logistics (selected papers from {ODYSSEUS} 2009 - the 4th International Workshop on Freight Transportation and Logistics). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0968090X10001087>
50. K. P. Sycara, S. P. Roth, N. Sadeh, and M. S. Fox, "Resource allocation in distributed factory scheduling," *IEEE Expert: Intelligent Systems and Their Applications*, vol. 6, no. 1, pp. 29–40, Feb. 1991. [Online]. Available: <http://dx.doi.org/10.1109/64.73815>
51. E. Kutanoglu and S. D. Wu, "On combinatorial auction and lagrangean relaxation for distributed resource scheduling," *IIE Transactions*, vol. 31, no. 9, pp. 813–826, 1999. [Online]. Available: <http://dx.doi.org/10.1023/A:1007666414678>
52. P. J. 't Hoen and J. A. La Poutré, *A Decommitment Strategy in a Competitive Multi-agent Transportation Setting*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 56–72. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-25947-3_4
53. A. Reyes-Moro and J. A. Rodríguez-Aguilar, *iAuctionMaker: A Decision Support Tool for Mixed Bundling*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 202–214. [Online]. Available: http://dx.doi.org/10.1007/11575726_15
54. M. Lemaître, G. Verfaillie, F. Jouhaud, J.-M. Lachiver, and N. Bataille, "Selecting and scheduling observations of agile satellites," *Aerospace Science and Technology*, vol. 6, no. 5, pp. 367 – 381, 2002. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1270963802011732>
55. P. Cramton, Y. Shoham, and R. Steinberg, Eds., *Combinatorial Auctions*, 1st ed. MIT Press, December 2005.
56. J. Bredin, R. T. Maheswaran, c. Imer, T. Başar, D. Kotz, and D. Rus, "A game-theoretic formulation of multi-agent resource allocation," in *Proceedings of the Fourth International Conference on Autonomous Agents*, ser. AGENTS '00. New York, NY, USA: ACM, 2000, pp. 349–356. [Online]. Available: <http://doi.acm.org/10.1145/336595.337525>
57. R. G. Smith, "The contract net protocol: High-level communication and control in a distributed problem solver," *IEEE Trans. Comput.*, vol. 29, no. 12, pp. 1104–1113, Dec. 1980. [Online]. Available: <http://dx.doi.org/10.1109/TC.1980.1675516>
58. T. W. Sandholm and V. R. Lesser, "Leveled commitment contracts and strategic breach," *Games and Economic Behavior*, vol. 35, pp. 212–270, 2001.
59. P. Sousa, C. Ramos, and J. Neves, "The fabricare scheduling prototype suite: Agent interaction and knowledge base," *Journal of Intelligent Manufacturing*, vol. 14, no. 5, pp. 441–455, 2003. [Online]. Available: <http://dx.doi.org/10.1023/A:1025749208437>
60. J. B. Cruz Jr, G. Chen, D. Li, and X. Wang, "Particle swarm optimization for resource allocation in uav cooperative control," in *AIAA Guidance, Navigation, and Control Conference and Exhibit*. Providence, USA, 2004, pp. 1–11.
61. P.-Y. Yin and J.-Y. Wang, "A particle swarm optimization approach to the nonlinear resource allocation problem," *Applied mathematics and computation*, vol. 183, no. 1, pp. 232–242, 2006.
62. F. Tao, D. Zhao, Y. Hu, and Z. Zhou, "Resource service composition and its optimal-selection based on particle swarm optimization in manufacturing grid system," *IEEE Transactions on industrial informatics*, vol. 4, no. 4, pp. 315–327, 2008.
63. A. Sen, *Collective Choice and Social Welfare*. Elsevier B.V., 1970.
64. J. Bentham, *An Introduction to the Principles of Morals and Legislation*. Library of Economics and Liberty [Online]: Oxford Clarendon Press, 1907. [Online]. Available: <http://www.econlib.org/library/Bentham/bnthPML.html>

65. J. Rawls, *A Theory of Justice*. Harvard University Press, 1971.
66. J. Pitt, J. Schaumeier, D. Busquets, and S. Macbeth, "Self-organising common-pool resource allocation and canons of distributive justice," in *Self-Adaptive and Self-Organizing Systems (SASO), 2012 IEEE Sixth International Conference on*, Sept 2012, pp. 119–128.
67. J. Pitt, D. Busquets, and S. Macbeth, "Distributive justice for self-organised common-pool resource management," *ACM Trans. Auton. Adapt. Syst.*, vol. 9, no. 3, pp. 14:1–14:39, Oct. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2629567>
68. K. A. De Jong, *Evolutionary computation: a unified approach*. MIT press, 2006.
69. N. F. McPhee and N. J. Hopper, "Analysis of genetic diversity through population history," in *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation - Volume 2*, ser. GECCO'99. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1999, pp. 1112–1120. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2934046.2934071>
70. H. G. Cobb and J. J. Grefenstette, "Genetic algorithms for tracking changing environments." DTIC Document, Tech. Rep., 1993.
71. L. J. Eshelman and J. D. Schaffer, "Preventing premature convergence in genetic algorithms by preventing incest," in *ICGA*, vol. 91, 1991, pp. 115–122.
72. K. Deb and D. E. Goldberg, "An investigation of niche and species formation in genetic function optimization," in *Proceedings of the 3rd international conference on genetic algorithms*. Morgan Kaufmann Publishers Inc., 1989, pp. 42–50.
73. M. Črepinšek, S.-H. Liu, and M. Mernik, "Exploration and exploitation in evolutionary algorithms: A survey," *ACM Computing Surveys (CSUR)*, vol. 45, no. 3, pp. 1–33, 2013.
74. D. Karaboga and B. Basturk, "A powerful and efficient algorithm for numerical function optimization: artificial bee colony (abc) algorithm," *Journal of global optimization*, vol. 39, no. 3, pp. 459–471, 2007.
75. J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proceedings of IEEE International Conference on Neural Networks*, vol. 4, 1995, pp. 1942–1948.
76. M. Dorigo and G. Di Caro, "Ant colony optimization: a new meta-heuristic," in *Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on*, vol. 2. IEEE, 1999, pp. 1470–1477.
77. V. Tereshko and A. Loengarov, "Collective decision making in honey-bee foraging dynamics," *Computing and Information Systems*, vol. 9, no. 3, p. 1, 2005.
78. D. Teodorovic and M. Dell'Orco, "Bee colony optimization—a cooperative learning approach to complex transportation problems," *Advanced OR and AI methods in transportation*, pp. 51–60, 2005.
79. M. Lovbjerg and T. Krink, "Extending particle swarm optimisers with self-organized criticality," in *Evolutionary Computation, 2002. CEC'02. Proceedings of the 2002 Congress on*, vol. 2, IEEE. Institute of Electrical and Electronics Engineers (IEEE), 2002, pp. 1588–1593.
80. T. M. Blackwell, "Particle swarms and population diversity," *Soft Computing*, vol. 9, no. 11, pp. 793–802, 2005.
81. K. Kyung-Joong and C. Sung-Bae, "Evolutionary ensemble of diverse artificial neural networks using speciation," *Progress in Modeling, Theory and Application of Computational Intelligence - 15th European Symposium on Artificial Neural Networks*, vol. 71, no. 7-9, p. 15, 2007.
82. L. L. Minku, A. P. White, and X. Yao, "The impact of diversity on online ensemble learning in the presence of concept drift," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 5, pp. 730–742, May 2010.
83. A. Blass, N. Dershowitz, and Y. Gurevich, "When are two algorithms the same?" *CoRR*, vol. abs/0811.0811, 2008. [Online]. Available: <http://arxiv.org/abs/0811.0811>
84. D. Challet and Y. Zhang, "Emergence of cooperation and organization in an evolutionary game," *Physica A: Statistical Mechanics and its Applications*, vol. 246, no. 34, pp. 407 – 418, 1997. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0378437197004196>
85. D. Challet, "Inter-pattern speculation: Beyond minority, majority and \$-games," *Journal of Economic Dynamics and Control*, vol. 32, no. 1, pp. 85–100, Jan. 2008. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0165188907000413>
86. G. Bianconi, T. Galla, M. Marsili, and P. Pin, "Effects of Tobin taxes in minority game markets," *Journal of Economic Behavior & Organization*, vol. 70, no. 1-2, pp. 231–240, May 2009. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167268108002084>
87. C. A. Zapart, "On entropy, financial markets and minority games," *Physica A: Statistical Mechanics and its Applications*, vol. 388, no. 7, pp. 1157–1172, Apr. 2009. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0378437108009801>

88. R. Metzler and C. Horn, "Evolutionary minority games: the benefits of imitation," *Physica A: Statistical Mechanics and its Applications*, vol. 329, no. 3-4, pp. 484–498, Nov. 2003. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0378437103006265>
89. Y. Li and R. Savit, "Toward a theory of local resource competition: the minority game with private information," *Physica A: Statistical Mechanics and its Applications*, vol. 335, no. 1-2, pp. 217–239, Apr. 2004. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0378437103011348>
90. D. Dhar, V. Sasidevan, and B. K. Chakrabarti, "Emergent cooperation amongst competing agents in minority games," *Physica A: Statistical Mechanics and its Applications*, vol. 390, no. 20, pp. 3477–3485, Oct. 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0378437111003876>
91. G. W. Greenwood, "Deceptive strategies for the evolutionary minority game," in *5th Intl Conf on Computational Intelligence and Games*, 2009, pp. 25–31. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1719293.1719308>
92. B. A. Mello and D. O. Cajueiro, "Minority games, diversity, cooperativity and the concept of intelligence," *Physica A: Statistical Mechanics and its Applications*, vol. 387, no. 2-3, pp. 557–566, Jan. 2008. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0378437107009995>
93. T. Galla, G. Mosetti, and Y.-C. Zhang, "Anomalous fluctuations in minority games and related multi-agent models of financial markets."
94. S. Ranadheera, S. Maghsudi, and E. Hossain, "Minority games with applications to distributed decision making and control in wireless networks," *IEEE Wireless Communications*, vol. PP, no. 99, pp. 2–10, 2017.
95. C. Zhang, W. Wu, H. Huang, and H. Yu, "Fair energy resource allocation by minority game algorithm for smart buildings," in *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2012*. IEEE, 2012, pp. 63–68.
96. C. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, pp. 379–423, 1948.
97. D. Ashlock, *Evolutionary Computation for Modeling and Optimization*. Springer-Verlag New York, 2006.
98. K. De Jong, *Evolutionary Computation: A Unified Approach*. MIT Press, 2006.
99. A. E. Eiben and J. Smith, "From evolutionary computation to the evolution of things," *Nature*, vol. 521, pp. 476–482, 2015.
100. A. E. Roth and I. Erev, "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term," *Games and Economic Behavior*, vol. 8, no. 1, pp. 164–212, 1995. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S089982560580020X>
101. V. Nallur, N. Cardozo, and S. Clarke, "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*, ser. SEAMS '16. New York, NY, USA: ACM, 2016, pp. 122–128. [Online]. Available: <http://doi.acm.org/10.1145/2897053.2897067>
102. J. L. Harper, "Plant demography and ecological theory," *Oikos*, vol. 35, no. 2, pp. 244–253, 1980.
103. C. Gini, "On the measure of concentration with special reference to income and statistics," *Colorado College Publication*, vol. 208, pp. 73–79, 1936.
104. J. Lehman and K. O. Stanley, "Improving evolvability through novelty search and self-adaptation," in *2011 IEEE Congress of Evolutionary Computation (CEC)*. Institute of Electrical and Electronics Engineers (IEEE), jun 2011.
105. J. Beal, "Functional blueprints: An approach to modularity in grown systems," *Swarm Intelligence*, vol. 5, no. 3, pp. 257–281, 2011.
106. A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, "Robots that can adapt like animals," *Nature*, vol. 521, pp. 503–507, 2015.
107. D. Smith, L. Tokarchuk, and G. Wiggins, "Rapid phenotypic landscape exploration through hierarchical spatial partitioning," in *International Conference on Parallel Problem Solving from Nature*. Springer, 2016, pp. 911–920.
108. J. K. Pugh, L. B. Soros, and K. O. Stanley, "Quality diversity: A new frontier for evolutionary computation," *Frontiers in Robotics and AI*, vol. 3, p. 40, 2016.
109. B. Sareni and L. Krähenbühl, "Fitness sharing and niching methods revisited," *IEEE Trans. Evol. Comput.*, vol. 2, no. 3, pp. 97–106, 1998.

110. J. B. Mouret and S. Doncieux, "Encouraging behavioral diversity in evolutionary robotics: An empirical study," *Evol. Comput.*, vol. 20, no. 1, pp. 91–133, Mar. 2012. [Online]. Available: http://dx.doi.org/10.1162/EVCO_a_00048
111. H. Hamann, "Lessons from speciation dynamics: How to generate selective pressure towards diversity," *Artificial Life*, vol. 21, no. 4, pp. 464 – 480, 2015. [Online]. Available: <http://search.ebscohost.com/login.aspx?direct=true&db=a9h&AN=110916959&site=ehost-live>
112. K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan, "A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii," in *International Conference on Parallel Problem Solving From Nature*. Springer, 2000, pp. 849–858.
113. M. Stillger and M. Spiliopoulou, "Genetic programming in database query optimization," in *Proceedings of the 1st Annual Conference on Genetic Programming*. Cambridge, MA, USA: MIT Press, 1996, pp. 388–393. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1595536.1595591>
114. P. Barmapalexis, K. Kachrimanis, A. Tsakonas, and E. Georgarakis, "Symbolic regression via genetic programming in the optimization of a controlled release pharmaceutical formulation," vol. 107, no. 1, pp. 75–82, May 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0169743911000153>
115. A. L. C. Bazzan, R. H. Bordini, G. K. Andriotti, R. M. Vicari, and J. Wahle, "Wayward agents in a commuting scenario (personalities in the minority game)," in *Proceedings Fourth International Conference on MultiAgent Systems*, 2000, pp. 55–62.

A

The clonal plasticity process step-by-step for *Evo* is as follows:

1. **Identify Plasticity Points:** The *Evo* algorithm diversifies itself based on the mutation probability of an individual. Two agents that diversify into different directions (one increasing the mutation probability, and the other decreasing) could end up with drastically different strategies after their evolutionary reproduction. Here, *ownMutationRate* is identified as a plasticity point of the algorithm.
2. **Evaluate Environmental Input:** For any given time instant, the algorithm has a 50% chance of picking the right group (*i.e.*, being in the minority). Hence, for any given period of cloning frequency, if the agent has accumulated (*wins-this-cycle*) more than 50% of the rewards during the period, the feedback is said to be positive. If the agent accumulates less than 50% of the rewards, then the feedback is negative.
3. **Plasticity Memory:** The agent keeps track of the rewards that it had accumulated during the previous adaptation (*wins-last-cycle*). Along with the feedback, this helps it decide what plasticity range to choose next (see Table 8).
4. **Clone and Modify Plastic Points:** In the event of *Exact Clone*, the agent does not change its plasticity point. In the event of *Low Plasticity*, the agent modifies its *ownMutationRate* in the same direction, as the plasticity memory from the previous generation. In the event of *High Plasticity*, the agent modifies its *ownMutationRate* to return to the value that it had from the previous generation.

Feedback	Plasticity Memory	Plasticity Range Chosen
Negative	(Disregarded)	High Plasticity
Positive	<i>wins-this-cycle</i> <= <i>wins-last-cycle</i>	Low Plasticity
Positive	<i>wins-this-cycle</i> > <i>wins-last-cycle</i>	Exact Clone

Table 8: Criteria for choosing plasticity range

B

The clonal plasticity process step-by-step for *RothErev* is as follows:

1. **Identify Plasticity Points:** The *RothErev* algorithm depends on two variables, *Recency* and *Epsilon*. These affect how much weight the algorithm gives to the recent past, and how much exploration it does, respectively. Any difference in the values of these variables can cause agents that start out the same, to diverge in their strategies.
2. **Evaluate Environmental Input:** For any given time instant, the algorithm has a 50% chance of picking the right group (*i.e.*, being in the minority). Hence, for any

given period of cloning frequency, if the agent has accumulated (*wins-this-cycle*) more than 50% of the rewards during the period, the feedback is said to be positive. If the agent accumulates less than 50% of the rewards, then the feedback is negative.

3. **Plasticity Memory:** The agent keeps track of the rewards that it had accumulated during the previous adaptation (*wins-last-cycle*). Along with the feedback, this helps it decide what plasticity range to choose next (see Table 8).
4. **Clone and Modify Plastic Points:** In the event of *Exact Clone*, the agent does not change any of its parameters. In the event of *Low Plasticity*, the agent picks one of its plastic points and modifies it in the same direction, as the plasticity memory from the previous generation. In the event of *High Plasticity*, the agent modifies *all* its plasticity points to return to the values that it had from the previous generation.

Note that in this paper, we have differentiated the algorithms based only on their plasticity points. It is also possible to modify the rules by which the environmental input is evaluated and how the plasticity range is chosen, for each individual algorithm.

C

In this section, we show detailed information about how cloning frequency affects both, the diversity in the population as well as the *Gini* index. Recall that the Gini index measures the equity of reward in a society. The closer the absolute value of the index is to zero, the more equitable a society is. Conversely, the higher the Gini index, the less equitable distribution of rewards in that society. Table 9 shows the mean diversity, rewards and gini index for each cloning frequency, for each mix of algorithm. The final column (*std dev*) records the standard deviation for the rewards achieved. Note that each population with a single algorithm (*BestPlay*, *RothErev* and *Evo*) has a different cloning frequency, while each population with a mix of algorithms has a common frequency (60, 120, 180, 240, 300). This is due to the fact that *Evo* requires at least 20 rounds for speciation to show a difference in behaviour.

	cloning frequency	diversity	reward	gini	std dev
BestPlay	5	1.7	605	2.2	3.830861
	15	1.7	525	7.3	3.140808
	25	1.7	509	18.5	2.813394
	35	1.7	503	23.7	2.711638
	45	1.7	499	24.9	2.802296
	55	1.7	496	25.1	2.857674
	65	1.7	494	24.4	2.970102
	75	1.7	493	22.8	3.053186
	85	1.7	492	21.7	3.142969
	95	1.7	491	19.3	3.249604
	105	1.7	490	18.1	3.321539
BestPlay-Evo	60	2.1	598	0.4	90.388678
	120	2.1	596	0.9	102.229262
	180	2.1	595	1.3	108.460674
	240	2.1	595	1.5	112.134200
	300	2.1	595	1.5	113.784304
Evo	20	2.7	849	-7.3	22.793937
	40	2.7	702	4.3	25.876079
	60	2.7	574	-5.5	35.548009
	80	2.7	516	5.9	32.871732
	100	2.7	504	7.4	26.535574
Evo-RothErev-BestPlay	60	2.6	766	1.0	116.125807
	120	2.5	785	0.9	115.274892
	180	2.5	796	0.8	114.048802
	240	2.5	803	0.7	112.573399
	300	2.5	808	0.8	111.829513
RothErev	5	2.7	931	15.8	2.643791
	15	2.7	952	7.5	5.226056
	25	2.5	953	9.4	7.318583
	35	2.5	950	5.1	10.374762
	45	2.6	948	3.1	13.096147
	55	2.6	946	3.8	19.383790
	65	2.7	946	3.9	22.110698
	75	2.7	946	3.5	28.256625
	85	2.7	946	3.2	31.609064
	95	2.7	945	2.8	36.288345
	105	2.7	945	2.9	37.789350
RothErev-BestPlay	60	2.2	910	1.0	217.470952
	120	2.2	916	1.6	221.771987
	180	2.2	922	1.6	224.047112
	240	2.2	927	2.0	225.288587
	300	2.2	927	2.5	226.053332
RothErev-Evo	60	2.7	925	4.0	123.979601
	120	2.7	931	3.4	121.939317
	180	2.7	935	3.0	120.620793
	240	2.7	937	2.7	119.954384
	300	2.7	938	2.7	119.438351

Table 9: All mixes of algorithms and cloning frequencies and their associated diversity, gini, reward values