

A Diversity Dilemma in Evolutionary Markets

Peter R. Lewis
School of Computer Science
University of Birmingham
United Kingdom
p.r.lewis@cs.bham.ac.uk

Paul Marrow
BT Group plc
Adastral Park, Ipswich
United Kingdom
paul.marrow@acm.org

Xin Yao
School of Computer Science
University of Birmingham
United Kingdom
x.yao@cs.bham.ac.uk

ABSTRACT

Markets are useful mechanisms for performing resource allocation in fully decentralised computational and other systems, since they can possess a range of desirable properties, such as efficiency, decentralisation, robustness and scalability. In this paper we investigate the behaviour of co-evolving *evolutionary market agents* as adaptive offer generators for sellers in a multi-attribute posted-offer market. We demonstrate that the evolutionary approach enables sellers to automatically position themselves in market niches, created by heterogeneous buyers. We find that a trade-off exists for the evolutionary sellers between maintaining high population diversity to facilitate movement between niches and low diversity to exploit the current niche and maximise cumulative payoff. We characterise the trade-off from the perspective of the system as a whole, and subsequently from that of an individual seller. Our results highlight a decision on risk aversion for resource providers, but crucially we show that rational self-interested sellers would not adopt the behaviour likely to lead to the ideal result from the system point of view.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Intelligent Agents, Multiagent Systems*;
J.4 [Computer Applications]: Social and Behavioural Sciences—*Economics*

General Terms

Algorithms, Economics.

Keywords

Diversity, Co-evolution, Online evolution, Market-based control, Self-interested agents.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICEC '11 Liverpool, UK
Copyright 2011 ACM ...\$10.00.

1. INTRODUCTION

Markets are now widely regarded [6, 7, 12, 19] to be useful mechanisms for performing resource allocation in fully decentralised computational and other systems. They can be efficient, dynamic and decentralised, leading to desirable properties such as robustness and scalability. Individual resource providers and users are typically represented by autonomous agents, acting as buyers or sellers in a market for resources. The agents follow some predefined or adaptive strategy, in order to act in the best interests of their host. In both strategy design and modelling, evolutionary algorithms are often used, due to an analogy between fitness-seeking evolution and rational profit-seeking self-interested economic agents [1, 5, 11, 19, 21, 22].

In this paper, we build upon a previous retail-inspired approach [19], in which *evolutionary market agents* decide the prices of resources in a posted-offer market mechanism. In the original model, self-interested competitive behaviour between agents led to evenly balanced allocation of resources when buyers valued resources equally, and stable uneven allocations reflective of buyers' valuations when they valued the resources differently. These are useful outcomes from the perspective of system designers harnessing economic principles to build decentralised computing systems.

Our new extension in this paper allows for resources to be described over additional quality attributes, rather than simply their price. In this extended model, resource users may have different preferences and constraints over multiple attributes, and as such form niches in the market into which sellers may segment. The key contribution of this paper is a deeper understanding of the complex interactions between the sellers' strategies in such markets, with a particular emphasis on those which lend themselves to an evolutionary interpretation. This insight will enable the better design of such strategies for posted-offer and other markets in the future. Throughout, we are concerned with the use of artificial markets as a control method for engineered systems. We therefore contrast the behaviour of the system from the global perspective of the system designer with that of the individual agents.

Specifically, we show that service providing nodes using *evolutionary market agents* [19] are able to position themselves in niches created through heterogeneous buyers automatically, providing levels of a quality attribute appropriate to them. However, we also show that algorithmic artefacts may distort this. Sellers are faced with a classic trade-off between exploring the attribute space for more profitable regions and exploiting a known good market position. In

evolutionary algorithms higher levels of exploration are a feature of greater population diversity. Evolutionary sellers are therefore faced with the decision of either maintaining a high population diversity to perpetuate search or maintaining a low diversity to exploit a market position and hence gain a higher immediate payoff. This trade-off is illustrated in the context of *evolutionary market agents* using a hypermutation operator.

Of course, the trade-off between exploration and exploitation is not new, and is well known particularly in the reinforcement learning community (see e.g. [23]). Indeed, it is often characterised in terms of a family of *multi-armed bandit* problems [2]. The problem is also well known in the evolutionary computation community [10] and according to Holland [13] a classic problem affecting all adaptive and learning systems.

However, what is investigated here is different from those evolutionary cases previously studied, due to the online nature of the fitness evaluation and the dynamics brought about by co-evolution. Particularly, in classical evolutionary search, the goal is that one individual within the population finds the optimum. In the case studied here, every fitness evaluation performed counts towards the performance of the algorithm. The goal in optimising the trade-off presented in this paper is to maximise total payoff over all fitness evaluations across the lifetime of the algorithm. This is similar to the performance metrics often used in dynamic optimisation (e.g. [20, 26]), though these cases do not consider co-evolution. Furthermore, in this paper we consider the competing outcomes of the achievement of the local objectives of the individual agents, and the global objective of the system designer.

Additionally, the idea of niching has been well used in the design of evolutionary algorithms [9]. However, the niching described in this paper is concerned with the proximity of offers, each of which is the result of an *independent* evolutionary algorithm rather than an individual within a single population. Since each evolutionary algorithm acts independently of the others as it converges to a particular point in the attribute space, no landscape information is shared between algorithms. Any such sharing between algorithms would require private market information to be shared between independent sellers, something which is ruled out by our model and often illegal in real markets.

The contributions of this paper therefore fall into two main areas. Firstly we show how the market-based control model previously proposed by us [19] can be extended to include multiple attributes of resources, not simply price. This allows for agents' preferences and constraints to be considered and leads to the emergence of niches in the market, into which *evolutionary market agents* can position themselves automatically, in order to provide attributes appropriate to the users. Secondly, the diversity trade-off simultaneously presents each individual evolutionary algorithm with a dilemma between high expected payoff and risk aversion. This diversity dilemma leads to sellers often being unlikely to find the optimal offers, instead becoming stuck in a less profitable niche. We characterise the trade-off using a randomly triggered hypermutation operator, the use of which may help *evolutionary market agents* to avoid this problem. However, subsequent experimental results show that self-interested agents facing this dilemma are unlikely to use the required level of hypermutation in their evolution-

ary algorithms, since this leads to a loss in expected payoff and a loss in confidence of achieving the expectation.

The remainder of this paper is structured as follows. In section 2 we describe our multi-attribute extension to the market model presented in [19], including describing agent behaviour and how this may induce niches in the market. We then demonstrate the emergence of selling agents into these niches in simulation in section 3. In section 4 we examine the number of sellers finding each niche, and demonstrate that the standard *evolutionary market agent* algorithm can lead to uneven segment sizes. Section 5 describes the explore - exploit trade-off and subsequent population diversity dilemma, which explains why the uneven segments occur and sellers are unable to obtain the optimal payoff. The dilemma is investigated in section 6 by means of a parametrised hypermutation operator, and section 7 examines what behaviour might be expected of self-interested sellers in the context of the diversity dilemma. Finally, sections 8 and 9 conclude the paper and look toward future directions respectively.

2. THE MARKET MODEL

The model which forms the basis for the investigation in this paper is the retail-inspired posted-offer market, which we previously studied in the context of computational resource allocation scenarios [19]. This approach has been shown to be capable of achieving a balanced load across providers [17] as well as other stable resource allocations [18] in decentralised computational systems. In the original model, resources are described purely in terms of their price, and are assumed otherwise to be perfectly substitutable.

Using the original notation, a resource π , is offered by a seller s_i for the price $p_{s_i}^\pi$. If π may instead be described over a number of attributes $\pi_1, \pi_2, \dots, \pi_n$, then we replace the offered price $p_{s_i}^\pi$ with a vector of offered attribute values, $\vec{x}_{s_i}^\pi$:

$$\vec{x}_{s_i}^\pi = \langle x_{s_i}^{\pi_1}, x_{s_i}^{\pi_2}, \dots, x_{s_i}^{\pi_n} \rangle \quad (1)$$

As described in section 1, the market being investigated is intended for use as a resource allocation tool for a distributed computing system. However at this stage, the effect of the embodiment of the approach in any particular application or network environment is not considered, since this would make it unnecessarily difficult to analyse and understand the underlying behaviour of the model. With this in mind, a number of simplifying assumptions are taken:

1. That the system proceeds synchronously in discrete time-steps,
2. that each buyer desires exactly one unit of π per time-step,
3. that the actual provision of π may be regarded as instantaneous, such that it does not interfere with the mechanism,
4. that each seller has sufficient quantity of π available to satisfy all the buyers, should it be so requested, and
5. that network connectivity is uniform.

The first two assumptions are present at this abstract stage in order to aid the analysis of the system, and though

their presence has not been fully investigated, there appears to be no obvious intuition as to why they should alter the underlying behaviour being demonstrated. Assumptions 3 and 4 may not be appropriate for certain embodiments of the approach, however are representative in theory of the provision of information-based services such as HTTP requests, and are present in other related work such as Kuwabara et al.'s [16] model. Finally, assumption 5 replicates the network conditions found in Wolski et al.'s [25] G-Commerce system.

2.1 Agent Behaviour

As in the original model, buying agents are faced with a decision problem: given the offers available, which one(s) should be chosen? Sellers are faced with a similar decision: what offer will provide the maximum payoff given the buyers and competitors? In order to decide this, the agent is endowed with a utility model which reflects the decision making priorities of its owner: the user or supplier of the resource. To a certain extent, it does not matter what this utility model is; the agent may treat its utility function as a black box, into which it places an offer, and from which it receives a utility value for that offer. In this way, offers may be compared against each other both in terms of preference and the magnitude of the preference.

However, in order to aid analysis we make use of the common utility models of multi-attribute utility theory (MAUT) [15], a branch of decision theory concerned with multiple criteria. Its aim is to analyse and compare the attractiveness of a set of alternatives, by calculating a single conjoint utility value for each alternative. This approach has long been useful in real-world decision making [3, 14, 24], where an expert's time is valuable or limited, and an analytical model may be used as an aid to the decision process. Similar reasoning applies here, where for example, the human user of a resource may not be willing or able to assess each and every offer for its attractiveness themselves, and an agent in possession of a model may save considerable time and effort.

The translation of human or business needs and preferences into utility functions and weights for their respective agents is a large and complex task, and we do not pretend to contribute to that field here. It is important however to note the assumptions relied upon by MAUT, particularly that of *utility independence* of attributes. Briefly, if the utility function of attribute x does not shift strategically as attribute y varies, then x is *utility independent* of y . In other words, for x to be *utility independent* of y , all the utility functions for x , defined as y varies, must be positive linear transformations of each other [15].

2.2 Additive Utility Model

As discussed, agents' utility functions may be constructed in a variety of ways in order to represent the underlying desirability of different combinations of attributes. One of the simplest utility models in MAUT is the additive utility function [15]. The additive utility model assumes a refinement of utility independence: *additive independence*. For attribute x to be *additively independent* of attribute y , the utility contribution of x must not be dependent on the value of y . If this assumption is true for all attributes, then the additive model may be used. For a more formal definition of *additive independence*, see Keeney and Raiffa [15].

In the additive model the decision maker calculates a

weighted sum of the utility associated with each attribute. An attribute's weight represents the importance of that attribute to the decision maker and gives its contribution to the overall utility of the alternative. The outcome of the function is a conjoint measure of each alternative, which may be used to determine the overall relative preference. Describing a buyer's consideration of an offer by seller s_i , the buyer's utility would be as follows:

$$u(\vec{x}_{s_i}^\pi) = \sum_{j=1}^n k_j u_j(x_{s_i}^{\pi_j}) \quad (2)$$

The attribute's utility function u_j determines the unweighted contribution of the j th attribute, the attribute's weight or scaling constant k_j represents the importance of the attribute to the decision maker. The scaling constants are normalised, such that $\forall k_j, 0 \leq k_j \leq 1$ and $\sum_{j=1}^n k_j = 1$.

2.3 A Simple Example

The decision process may be illustrated by means of a simple example. Consider a buyer faced with the following offers from sellers s_1 and s_2 :

$$\vec{x}_{s_1}^\pi = \langle 600, 300 \rangle$$

$$\vec{x}_{s_2}^\pi = \langle 800, 500 \rangle$$

where a high value for the first attribute is undesirable to the buyer, but a high value for the second attribute is desirable. The buyer agent's attribute utility functions will of course determine the final utility value, but in this example assume that they are linear functions which normalise from between some minimum value 0 and maximum value 1000, and in the case of the first attribute, obtains the inverse. This may result in a conjoint utility of, for example,

$$u(\vec{x}_{s_1}^\pi) = -0.6k_1 + 0.3k_2$$

$$u(\vec{x}_{s_2}^\pi) = -0.8k_1 + 0.5k_2$$

Since $u(\vec{x}_{s_1}^\pi) > 0$ if and only if $k_2 > 2k_1$, the offer from s_1 is only acceptable to a buyer for whom the second attribute is twice as important as the first. For s_2 's offer, $u(\vec{x}_{s_2}^\pi) > 0$ if and only if $k_2 > 1.6k_1$, which is somewhat less demanding.

Note that though this example is described in terms of a buyer agent, the same approach is applied to a seller, when comparing the utilities of two or more offers which it is considering making.

2.4 Buyer and Seller Strategies

For consistency with previous work [19], buyers then make use of a given fixed strategy in order to decide amongst offers, given their expected utility values. In this paper, we use *spread buyers*, who purchase a proportion of the resource from each acceptable seller, relative to the expected utility gain from purchasing from that seller. *Spread buyers* have been shown to lead to stable desirable outcome resource allocations due to the smooth nature of their payoff functions. Further details and analysis of these strategies can be found in [18] and [19]. Since we fix the buyers' strategies, in this paper our investigations are restricted to those cases where adaptation is driven by sellers. An important extension to this work will be the consideration of co-adaptive systems.

In the context of the *evolutionary market agent* algorithm, which we recall from [19] in algorithm 1, the conjoint utility

measure provided by the multi-attribute utility function provides the payoff associated with a certain offer, and hence a fitness measure for each agent’s evolutionary algorithm. This is the most straightforward approach taken in extending the single attribute case, is sufficient for our purposes here, and therefore that taken in this paper.

Algorithm 1: The evolutionary market agent algorithm [19]

1. Decide upon the design parameters to be used: initial price range $[p_{min}, p_{max}]$, population size and mutation factor, α . In the simulations described, $p_{min} = 0$, and $p_{max} = 500$. A population size of 20 was used, with a mutation factor, $\alpha = 3.0$.
 2. Generate an initial population, Pop , and set $k = 1$. Each individual in Pop is a real value, drawn from the uniform random distribution $[p_{min}, p_{max}]$.
 3. *Initial fitness testing*
 - (a) Set the seller’s offer to the value of the first individual in Pop , and enter the market for one market time-step. Record the seller’s payoff, P_{s_i} as that individual’s fitness.
 - (b) Repeat for the next individual in Pop , until all initial individuals have been evaluated in the market.
 4. *Probabilistic tournament selection*
 - (a) Select four individuals, x_1, x_2, x_3 and x_4 from Pop , at random, such that $x_1 \neq x_2 \neq x_3 \neq x_4$.
 - (b) Let champion c_1 be either x_1 or x_2 , the fitness of whichever is greater with probability 0.9, the fitness of whichever is less otherwise.
 - (c) Let champion c_2 be either x_3 or x_4 , the fitness of whichever is greater with probability 0.9, the fitness of whichever is less otherwise.
 5. Let the offspring, o , be a new individual with its price equal to the mid-point of c_1 and c_2 .
 6. Mutate o , by perturbing its value by a random number drawn from a normal distribution with mean zero and standard deviation α .
 7. Select the individual in $\{x_1, x_2, x_3, x_4\}$ with the lowest fitness value, remove it from Pop , and insert o into Pop .
 8. Set the seller’s offer to the value encoded in o , and enter the market for one market time-step. Record the seller’s payoff, P_{s_i} as o ’s fitness.
 9. Repeat from step 4.
-

2.5 Budget Constraints

The utility functions above describe how an agent may order and compare offers described over multiple attributes. However, they do not consider that an offer, though attractive, may be unattainable. A simple example of this is an

offer which, from a buyer’s perspective, provides a highly attractive resource at a reasonable but unaffordable cost. In this scenario, an undesirable attribute such as price, may be beyond the buyer’s budget constraint.

Similar constraints may exist on desirable attributes; a buying agent may be instructed not to accept an offer where, for example, a security attribute is below a certain threshold. Similarly, a seller agent may be physically unable to provide a resource with a latency below a certain limit value. Regardless of its utility, we consider that an offer is unacceptable if any attribute is beyond the budget constraint of the decision maker. Such offers are therefore rejected.

The presence of heterogeneous hard constraints in a population of buyers leads to the presence of market niches even when in all other respects the buyers are homogeneous. We call this *constraint-induced niching*. It is important to be aware that there are other forms of market niching, for example those induced through a trade-off between multiple desirable or undesirable attributes, especially in the presence of more complex non-linear utility functions and constraints. However, in this paper we investigate the performance of the evolutionary market algorithm in the presence of such *constraint-induced niching*, though we see no reason why the results presented here would not apply to any other form of market niching.

3. EMERGENCE OF MARKET NICHES

We firstly illustrate, by means of simulation, that *evolutionary market agents*, as introduced in [17] and without modification are able to position sellers in market niches, by considering a conjoint utility value for each offer as that offers fitness value in the evolutionary algorithm.

Our experimental set up consisted of 20 seller agents and 1000 buyer agents trading a resource with two attributes, π_1 and π_2 , notionally price and quality respectively. Both buyers and sellers had homogeneous linear attribute utility functions, such that attribute π_1 was desirable to the sellers but undesirable to the buyers, while attribute π_2 was undesirable to the sellers but desirable to the buyers. For the additive utility functions, sellers were given the weights 0.7, 0.3, while buyers were given the weights 0.3, 0.7. The exact values of these weights, and their symmetry, proved not to be important, however they were found to provide a good range within the attribute space for feasible offers.

In order to generate the niches, each buyer had a hard budget constraint in terms of attribute π_1 , notionally price, such that any offer for which x^{π_1} was greater than this constraint is rejected as unacceptable by that buyer. The value of each buyer’s budget constraint was drawn with equal probability from one of two Gaussian distributions, with a relatively low standard deviation when compared with the offer space as a whole. The means of these distributions were 100 and 200.

Figure 1 shows the position of each of the 20 sellers’ offers in the attribute space, at four occasions during the simulation. As can be clearly seen, from an initial random distribution, sellers’ offers quickly found the feasible region between the two diagonal lines. The offers then converged to two niches, forming market segments, one for each subset of the buyer population. The system then remains stable, other than for ongoing small mutations about the centre of each niche. In summary, we observe the seller agents self-organising to equilibrium values for their offers’ attributes.




Figure 1: Evolution of the offers in the attribute space through time for a typical simulation run. Each point represents an individual seller's offer. The space between the diagonal lines is the feasible region in which offers are acceptable to both buyers and sellers. The two vertical lines indicate the means of the Gaussian distributions from which the buyers' budget constraints are drawn. Evolutionary market agents lead the sellers to find the niches created by the heterogeneous budget constraints.




Figure 2: Segment A, the lower niche, repeatedly contains fewer sellers than segment B, the higher niche.

4. QUANTIFYING MARKET NICHEs

We know therefore that the sellers use of *evolutionary market agents* allows them to find and position their offers in otherwise unknown market niches. However, it is also clear from the results in figure 1 that the sellers are not evenly split between the niches. Figure 2 shows the number of sellers positioned within a radius of 30 from the centre of each niche, as the simulation progresses. Mean and standard deviation are shown for 30 independent runs.

We can verify that this is not optimal seller behaviour by looking at the payoffs the sellers obtain in each niche. Figure 3 shows the mean payoffs, along with their standard deviation, for the sellers in each segment, over time.

The sellers in segment A generally obtain a higher payoff than those in segment B, however their payoff is highly unpredictable, indicated by the very high standard deviation across the independent runs. Conversely, the sellers in segment B are obtaining a lower payoff, though this is highly stable.

The question then arises of why the sellers in segment B do not move their offers to the niche occupied by segment A, given the increased likelihood of obtaining a higher payoff for doing so. Indeed, beyond iteration 50, there is relatively little change in either segment size or payoff. This is due to the lack of diversity in the sellers' populations of offers after this stage in the simulation. Since the populations are highly converged, making a jump to the more attractive niche is very unlikely.

5. THE DIVERSITY DILEMMA

Since each seller makes use of its own independent evolutionary algorithm to generate competitive offers, the performance of each generated offer at a given time is dependent in part on the current offers from the other sellers. Of vital importance therefore when considering the behaviour of evolutionary sellers is that since the seller has no model of the market, in order to obtain a fitness value for an offer in its population, it must test this offer in the live market. This is due to information being private and therefore




Figure 3: Mean and standard deviation seller payoffs for each segment. The payoff for sellers in segment A, the lower niche, is greater though highly unpredictable. The payoff for segment B is lower though more stable.

a seller having no knowledge of the buyers or behaviour of its competitors, other than the payoff returned from its actions. Performing a fitness evaluation on a given offer in the population therefore counts towards the total payoff which is obtained by the seller over time.

Since every fitness evaluation counts, once a seller can be sure it has found the best position for its offer, it can maximise its payoff over time by always taking that position. In terms of the evolutionary algorithm, this means the seller's population of offers converging to that point, losing all diversity.

However, as is usual with evolutionary algorithms, a seller can never be sure it has found this position, since if it is at a local optimum, a better position may exist elsewhere in the attribute space. This is what we observe occurring in the simulation results described in section 4. In order for a seller to find a position in the attribute space which would lead to a higher payoff, or indeed to gain confidence that there does not exist such a position, the seller must perform a search. In the context of the evolutionary algorithm, this search requires population diversity. It is here that the diversity dilemma presented is different from those previously studied; the algorithm must balance the diversity required to increase the likelihood of finding better solutions against the loss of cumulative fitness that will occur when an evaluation of a poor individual is made (due to higher diversity). Since every fitness evaluation must be made in the live market, the real problem instance, the algorithm cannot achieve both simultaneously. Additionally, since each seller has its own independent evolutionary algorithm, which all competitively co-evolve and do not share information, previously accumulated landscape knowledge and knowledge of good search behaviour (with respect to the diversity dilemma) will have only a limited use.

Since the *evolutionary market agent* algorithm has been hitherto tested primarily in single attribute environments with single equilibria, it is no surprise that the algorithm favours an initial search period with high population diver-

sity, in which the equilibrium is found, followed by a loss of diversity in the population, as the discovered position may be exploited to obtain maximum payoff. In scenarios such as that investigated in this paper, where, from the perspective of the algorithm, more than one optimum exists, a different approach to the diversity dilemma may be needed.

6. INVESTIGATING THE DIVERSITY DILEMMA

Our focus in this section is to explore the diversity dilemma described in section 5, and attempt to mitigate its effect on the overall payoff which a seller achieves. Increased population diversity is a characteristic which is often found to be desirable in evolutionary algorithms applied to dynamic optimisation problems [4]. Methods of diversity generation and preservation include immigrant schemes, random restarts and hypermutation [8]. In this section we use hypermutation in order to characterise the diversity dilemma, as it has been shown to generate and preserve diversity without discarding as much information as alternative schemes such as random restarts [8]. Hypermutation is typically applied periodically in order to inject a greater amount of diversity into the population, and involves the mutation factor applied being substantially higher than the base mutation factor, but for just one or a small number of mutations. For the vast majority of mutations, the normal mutation factor is used.

However, our departure from hypermutation as originally presented is that rather than rely upon a trigger, such as a drop in evolutionary fitness, to initiate the increase in mutation factor, we instead trigger the hypermutation randomly. In our scenario, it is not clear what might constitute an effective trigger, since it is not a drop in fitness or a lack of diversity *per se* which indicates the need for an increased search. Instead, as described in section 5, any seller could at any time be in a suboptimal market position without knowledge of this. Therefore, we instead perform a hypermutation instead of a regular mutation with a very small probability, each time a mutation is made.

In the first experiment in this section, each time a mutation was due to be performed, a hypermutation was instead made with probability 0.01. The occurrence of a hypermutation event caused the standard deviation of the Gaussian distribution from which the mutation factor was drawn to increase from 3.0 to 50.0 for that mutation only. All else remained the same as the simulation described in sections 3 and 4. Figure 4 shows that despite the initially converged populations being similar to in the version without hypermutation, the addition of random hypermutation allows sellers in segment A to move occasionally to segment B in order to take advantage of the higher payoff available. This plot may be compared with that in figure 2.

As the number of sellers in segment A increases, so the market share served by each seller in that segment decreases, and conversely so for segment B. This is reflected in the payoff obtained by the sellers, illustrated in figure 5, which may be compared with figure 3. It is clear that the ability of even a small number of sellers to move to the more profitable segment A leads to a somewhat more equitable payoff between the segments. However, more importantly, the addition of the hypermutation operator allows the sellers to continue to move between niches throughout the simulation, with-




Figure 4: The addition of hypermutation allows sellers to continue to move between niches throughout the simulation, without requiring a higher base mutation factor.




Figure 5: The addition of hypermutation leads to both more equitable and stable payoffs between the two segments.




Figure 6: Mean sellers' payoff after convergence, by segment, calculated over thirty independent runs, as the number of hypermutations, n , increases. As n grows, the segments' payoffs become more equitable.

out requiring a higher base mutation rate and despite their converged populations.

This result in itself is useful, however even given hypermutation, we still do not observe the segments' respective payoffs equalising within any reasonable time, rather drifting together at a slow pace. The rate of equalisation may be increased by increasing the level of hypermutation. In the following set of experiments, the hypermutation operator was modified such that, when triggered, the subsequent n mutations were hypermutations. The probability of a hypermutation occurring remained the same.

Here we observe the mean payoff of the sellers in each segment, calculated between iterations 4000 to 4999, late enough such that the sellers' populations have completed their initial convergence period, and any observed population diversity and subsequent movement is due to mutation or hypermutation. Figure 6 shows the mean and standard deviation of these values, calculated over for thirty independent simulation runs, plotted as the number of hypermutations per hypermutation event, n varies. Similarly, figure 7 shows the mean of the standard deviations between the sellers' payoffs in each run, again calculated over the same thirty independent runs.

As figure 6 shows, increasing the number of hypermutations performed, n , leads to increasingly equitable payoffs between the sellers, as an increasing amount of movement is facilitated between the niches, and they become less uneven in size. However, there is a price to pay for this equity; this increase in diversity brought about by more hypermutation leads also to an increased standard deviation between the sellers' payoffs, as shown in figure 7. This is due to the sellers' populations becoming more diverse, and the evaluation of each individual contributing towards the total payoff.

7. EXPECTED SELLER BEHAVIOUR

In the model investigated, each seller makes use of its own independent evolutionary algorithm. The algorithms competitively co-evolve in the live market, do not share in-

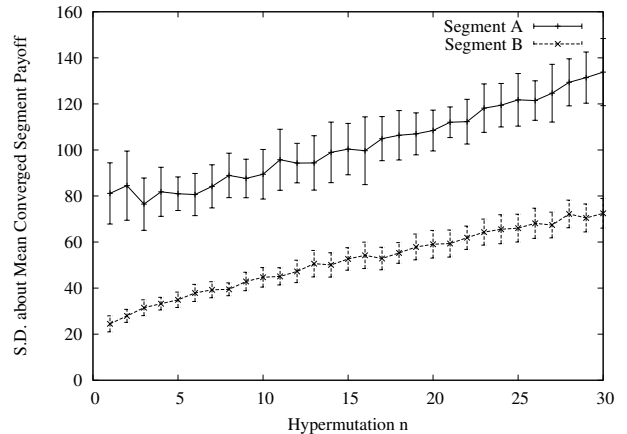


Figure 7: Standard deviation between the sellers' payoffs after convergence, by segment, calculated over thirty independent runs, as the number of hypermutations, n , increases. As n grows, the standard deviation between the sellers' payoffs within the segments also increases.

formation or possess more than local knowledge. However, the results and discussion presented in section 6 are largely from the perspective of the system as a whole, from a *bird's eye view*, where global knowledge, full information about all sellers is available. This information is of course not available to the individual sellers, who rely upon private information only. What behaviour then should we expect of self-interested sellers' *evolutionary market agents*, given the information they have available, with regard to the diversity of their populations?

Figure 8 shows the mean payoff and standard deviation over all sellers in the model, during the converged phase between iterations 4000 and 4999, regardless of their position or payoff. This may be interpreted as an expectation of a seller's payoff during this phase, given that which segment they find first is not yet determined. Again, we vary the number of hypermutations, n , on the horizontal axis.

Firstly, we observe that the highest expected payoff may be achieved by using no hypermutation at all, when n is zero. Secondly, this comes at a high price in terms of the certainty of the expectation, indicated by the high accompanying standard deviation amongst the sellers. This is due to there being a chance of the seller finding a niche which it shares with relatively few competitors. However, this is not guaranteed, and the payoff it receives will be much lower if it finds only the more populated niche. Figures 2 and 3 illustrated this.

Thirdly, by adopting a small amount of hypermutation (e.g. $n = 1$), a seller can drastically increase its chance over time of finding the more profitable niche and hence the certainty of the expectation. This comes at a small but significant cost to its expected payoff.

The highest degree of certainty was achieved by sellers in this simulation by performing a small number of hypermutations. Values for n between 2 and 4 give rise to the lowest standard deviations, again with a small penalty in the expected payoff.

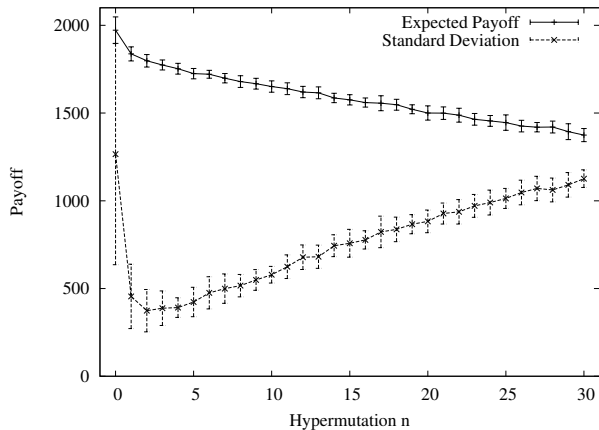


Figure 8: Mean of and standard deviation between the sellers’ payoffs after convergence, for all sellers, calculated over thirty independent runs, as the number of hypermutations, n , increases. Rational sellers face a choice between a high expected payoff with a low certainty, and a lower expected payoff with a greater certainty. High levels of hypermutation do not appear beneficial.

Larger number of hypermutations however led to a decrease in the level of certainty with which a seller would obtain the expected payoff, and a decrease in the expected payoff itself. Despite that from the perspective of the system as a whole, this leads to more equitable payoffs and evenly sized segments, such behaviour is hard to motivate from a seller’s perspective. Rather, we expect that sellers might choose between maximising their expected payoff, at a cost of certainty, or else sacrifice a small amount of expected payoff in order to buy a greater amount of certainty over their payoff. This is how the diversity dilemma may be characterised from a seller’s perspective, and either option leads to the greatest likelihood being an uneven balance of sellers between the segments.

8. CONCLUSIONS

In this paper we have extended significantly our previous evolutionary market-based approach to resource allocation [19], in order that resources may be described over multiple attributes in addition to price only. Resource users may have different preferences and constraints over the attributes, and as such form niches in the market, into which sellers may segment. Resource providers’ pricing strategies are modelled using using *evolutionary market agents*, which decide the offers for their resources, in a posted-offer market mechanism.

We showed, through simulation, that *evolutionary market agents* are able to position themselves automatically in niches created by heterogeneous buyers, providing levels of a quality attribute appropriate to them. However, we highlight a dilemma faced by evolutionary sellers, between a maintaining high population diversity to facilitate search and a low diversity to exploit a market position and hence gain a higher payoff. This dilemma arises from an instance

of the explore - exploit trade-off, which was illustrated and characterised using a parametrised hypermutation operator.

We found that the standard *evolutionary market agent* algorithm can lead to uneven segment sizes, since populations of offers with low diversity and hence less exploratory behaviour are unlikely to move between niches. As a result, sellers are unable to obtain the optimal payoff. We further found that by introducing a randomly triggered hypermutation operator, sellers become more likely to move between niches in order to obtain a higher payoff, making the segment sizes more even. This increase in exploratory behaviour leads to an increased certainty of achieving the expected payoff. Certainty comes at a cost however, as increased exploration also means that sellers must accept a slightly lower overall expected payoff. This presents the seller with a decision, which will be made based on their attitude to risk.

However, the levels of exploration (through evolutionary diversity) required to achieve truly even segment sizes and fluid movement between segments are unlikely to be realised by self-interested sellers, since when this is the case both the expected payoff and the confidence of achieving the expectation are reduced. Importantly, the exploration by use of hypermutation has highlighted the need to consider exploration - exploitation trade-offs in the design of adaptive pricing algorithms for market-based control systems.

9. FUTURE WORK

Since nothing in the implementation of the agents or evolutionary algorithm is specific to the particular way in which the niches in the market were generated in our experiments, we expect that the results obtained here will apply equally to other forms of market niching. Nevertheless, it will be useful to confirm this experimentally in more complex markets, for example through the study of niches induced not just by agents’ constraints, but also by their preferences over a greater number of attributes.

It is important to note that in this paper we are not arguing that hypermutation should be a preferred way of tackling the sellers’ diversity dilemma; though we demonstrate that it provides one such method, we use it primarily as an exploratory tool. Therefore, though the trade-offs faced by sellers appear to be inherent to the system itself, other diversity mechanisms could lead to potentially more attractive positions within the trade-off. This will of course depend on the objectives of a particular resource provider, and there is therefore significant scope for numerous other diversity mechanisms to be explored in this context.

Finally, it will be interesting to validate these results in real-world applications which make use of market-based approaches to resource allocation. We have begun work to apply this approach to cloud computing, which provides a large scale test bed for such multi-attribute market-based control methods.

10. ACKNOWLEDGEMENTS

The research leading to these results has received funding from the European Union Seventh Framework Programme under grant agreement n° 257906.

11. REFERENCES

- [1] E. J. Anderson and T. D. H. Cau. Modeling implicit collusion using coevolution. *Operations Research*, 57(2):439–455, 2009.
- [2] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning*, 47:235–256, 2002.
- [3] M. Barbuceanu and W.-K. Lo. Multi-attribute utility theoretic negotiation for electronic commerce. In *Agent-Mediated Electronic Commerce III, Current Issues in Agent-Based Electronic Commerce Systems*, pages 15–30, London, UK, 2001. Springer-Verlag.
- [4] J. Branke. *Evolutionary Optimization in Dynamic Environments*. Kluwer, 2001.
- [5] A. Byde. Applying evolutionary game theory to auction mechanism design. In *IEEE International Conference on E-Commerce*, pages 347–354, 2003.
- [6] S. H. Clearwater, editor. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, Singapore, 1996.
- [7] D. Cliff and J. Bruten. Simple bargaining agents for decentralized market-based control. Technical Report HPL-98-17, HP Laboratories, Bristol, UK, 1998.
- [8] H. G. Cobb and J. J. Grefenstette. Genetic algorithms for tracking changing environments. In *Proceedings of the 5th international conference on genetic algorithms*, pages 523–530, 1993.
- [9] P. Darwen and X. Yao. Every niching method has its niche: Fitness sharing and implicit sharing compared. In *Parallel Problem Solving from Nature—PPSN IV*, page 398–407. Springer, 1996.
- [10] A. E. Eiben and C. A. Schippers. On evolutionary exploration and exploitation. *Fundamenta Informaticae*, 35:1–16, 1998.
- [11] T. Eymann. Co-evolution of bargaining strategies in a decentralized multi-agent system. In *AAAI Fall 2001 Symposium on Negotiation Methods for Autonomous Cooperative Systems*, pages 126–134, 2001.
- [12] T. Eymann, M. Reinicke, O. Ardaiz, P. Artigas, L. Diaz de Cerio, F. Freitag, R. Messeguer, L. Navarro, D. Royo, and K. Sanjeevan. Decentralized vs. centralized economic coordination of resource allocation in grids. In *Lecture Notes in Computer Science*, volume 2970, pages 9–16. Springer, 2004.
- [13] J. H. Holland. Genetic algorithms. *Scientific American*, 267(1):66–72, 1992.
- [14] G. P. Huber. Multi-attribute utility models: A review of field and field-like studies. *Management Science*, 20(10):1393–1402, 1974.
- [15] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley and Sons, New York, 1976.
- [16] K. Kuwabara, T. Ishida, Y. Nishibe, and T. Suda. An equilibratory market-based approach for distributed resource allocation and its applications to communication network control. In S. H. Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, pages 53–73. World Scientific, Singapore, 1996.
- [17] P. R. Lewis, P. Marrow, and X. Yao. Evolutionary market agents for resource allocation in decentralised systems. In *Parallel Problem Solving From Nature - PPSN X*, volume 5199 of *Lecture Notes in Computer Science*, pages 1071–1080. Springer, 2008.
- [18] P. R. Lewis, P. Marrow, and X. Yao. Evolutionary market agents and heterogeneous service providers: Achieving desired resource allocations. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 904–910, 2009.
- [19] P. R. Lewis, P. Marrow, and X. Yao. Resource allocation in decentralised computational systems: An evolutionary market based approach. *Autonomous Agents and Multi-Agent Systems*, 21(2):143–171, 2010.
- [20] N. Mori and H. Kita. Genetic Algorithms for Adaptation to Dynamic Environments - A Survey. In *IECON 2000: 26th Annual Conference of the IEEE Industrial Electronics Society, Vol 4*, pages 2947–2952, Nagoya, Japan, 2000.
- [21] S. Phelps, P. McBurney, S. Parsons, and E. Sklar. Co-evolutionary auction mechanism design: A preliminary report. In *Lecture Notes in Computer Science*, volume 2531, pages 123–142. Springer, 2002.
- [22] T. C. Price. Using co-evolutionary programming to simulate strategic behaviour in markets. *Journal of Evolutionary Economics*, 7(3):219–254, 1997.
- [23] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, 1998.
- [24] G. W. Torrance, M. H. Boyle, and S. P. Horwood. Application of multi-attribute utility theory to measure social preferences for health states. *Operations Research*, 30(6):1043–1069, 1982.
- [25] R. Wolski, J. Plank, J. Brevik, and T. Bryan. Analyzing market-based resource allocation strategies for the computational grid. *International Journal of High Performance Computing Applications*, 15(3):258, 2001.
- [26] X. Yu, Y. Jin, K. Tang, and X. Yao. Robust Optimization over Time – A New Perspective on Dynamic Optimization Problems. In *WCCI 2010 IEEE World Congress on Computational Intelligence*, pages 3998–4003, 2010.