Evolutionary Market Agents for Resource Allocation in Decentralised Systems

Lewis, Peter; Marrow, Paul; Yao, Xin

DOI:

[10.1007/978-3-540-87700-4_106](https://doi.org/10.1007/978-3-540-87700-4_106)

Citation for published version (Harvard):

Lewis, P, Marrow, P & Yao, X 2008, 'Evolutionary Market Agents for Resource Allocation in Decentralised Systems', Paper presented at Parallel Problem Solving from Nature – PPSN X: Proceedings of the 10th International Conference, Berlin, Heidelberg, 1/09/08 pp. 1071-1080. https://doi.org/10.1007/978-3-540-87700-4_106

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Evolutionary Market Agents for Resource Allocation in Decentralised Systems

Peter R. Lewis¹, Paul Marrow^{1,2}, and Xin Yao¹

¹ University of Birmingham, Birmingham, UK.

² BT Group plc, Adastral Park, Ipswich, UK.

p.r.lewis@cs.bham.ac.uk, paul.marrow@bt.com, x.yao@cs.bham.ac.uk,

Abstract. We introduce self-interested evolutionary market agents, which act on behalf of service providers in a large decentralised system, to adaptively price their resources over time. Our agents competitively co-evolve in the live market, driving it towards the Bertrand equilibrium, the non-cooperative Nash equilibrium, at which all sellers charge their reserve price and share the market equally. We demonstrate that this outcome results in even load-balancing between the service providers.

Our contribution in this paper is twofold; the use of on-line competitive co-evolution of self-interested service providers to drive a decentralised market towards equilibrium, and a demonstration that load-balancing behaviour emerges under the assumptions we describe.

Unlike previous studies on this topic, all our agents are entirely self-interested; no cooperation is assumed. This makes our problem a non-trivial and more realistic one.

Key words: decentralised systems, market-based control, co-evolution, load-balancing, self-interested agents

1 Introduction

Emerging paradigms for the development and deployment of massively distributed computational systems allow resources to span many locations, organisations and platforms, connected through the Internet [1]. It has been predicted that the majority of transactions over the Internet will, in the future, be carried out by autonomous agents on behalf of their owners [2]. In this scenario, neither control nor even full knowledge of key resources may be assumed. There is therefore a need to find novel ways to understand and autonomically manage and control these large, decentralised and dynamic systems [3].

Gupta et al. [4] propose externality pricing for the provision of otherwise virtually zero cost per-use computational services. They argue that this approach, where service users self-select their quantity based on price, is a more preferable approach to the alternative of provider-side enforced quantity limits.

We propose the use of autonomous *evolutionary market agents* as an approach to achieving this. Evolutionary market agents operate on behalf of individual actors in a decentralised market-based system. Such decentralised markets have no auctioneers or market-makers. Instead, selling agents (the service

providers) advertise their services at a price, and buying agents (the service users) decide whether to buy, and from whom. Making use of only local information about the live market, the sellers adapt (or evolve) the offers of their host in order to maximise their payoff. When coupled with rational, self-interested buying agents, we demonstrate that this approach is able to provide load-balanced allocations across the system as a whole.

Market-based resource allocation and adaptive pricing are not new ideas (see [5] for an introduction) and our contribution in this paper is twofold. Firstly, we describe how decentralised computational markets can provide an emergent load-balancing behaviour between self-interested agents at equilibrium, and secondly, we demonstrate the use of competitive co-evolution to drive the market towards this equilibrium.

2 Related Work

Traditionally, load-balancing is done in a centralised manner [6]. Relying on a single node, such approaches have a central point of failure. Alfano and Di Caprio note that scalability is a critical factor in load-balancing systems [7]. They present a scalable, decentralised load-balancing mechanism, based upon cooperating peers. Our model, however, does not require cooperation. Indeed, its power lies in the self-interested competition of peers, over whom we may not have control.

As such, our approach falls into the broad category of market-based control, a methodology which has been applied to the allocation of resources in various real-world scenarios. Clearwater provides a useful introduction to the use of computational markets in scenarios such as bandwidth allocation and air-conditioning control [5]. A full review is beyond the scope of this paper, but may be found in [8].

Cliff and Bruten note that a large proportion of market-based control systems, however, either rely on a central auctioneer, or require human intervention [8]. Therefore, though much of the computation is done by individual agents and is distributed, these systems are often not decentralised. They argue that this leads to a brittleness of the system.

A number of distributed auction mechanisms have also been proposed [9–12], which do not rely on one central auctioneer. These reduce the fragility associated with reliance upon a single point, provide more scalability and allow for dynamic composition of auctions. Typically, the central auctioneer is replaced by a number of local ones, which communicate through some secure means. However, similarly to the vulnerabilities of the Internet’s domain name system [13], failure at certain points in the network may well cripple wider functionality, at best.

Kuwabara et al. propose what we believe to be the most decentralised market-based approach to the allocation of resources [14]. Here, no auctioneer, specialist or market-maker is used; prices are set solely by the sellers and advertised via a broadcast mechanism. Rational buyers then decide the quantity to purchase from each seller, in order to maximise their payoff. However, unlike our sellers’

strategies, those used by Kuwabara et al. are not driven by self-interest. The problem studied here is of the same general form as that in [14].

A variety of computational intelligence techniques have been used to attempt to replicate the behaviour of human marketing managers in the real world. Midgley et al. model brand managers in the retail coffee market as agents, which determine the occurrence and nature of various promotions and price discounts [15]. However, the difference between these cases and our problem is that we are not interested in modelling human behaviours, or in replicating them.

Applying evolutionary techniques to self-interested selling behaviour, Cheung et al. [16] used a simple model of the Australian consumer petroleum market to predict how sellers would modify their prices in a competitive environment. By giving sellers historical information about each other, they correctly replicated *implicit cartel* behaviour found in the real world, where despite the short-term rationality of cutting prices to increase market share, sellers in fact raised their prices in step. Their simulations resulted in sellers colluding in order to all charge the maximum price and share the market equally. If all other sellers cooperate, then this is the optimal strategy from a seller's point of view. However, only one seller not cooperating and undercutting the others, leads to it making a greater short-term payoff.

3 Problem Formulation

We consider a scenario consisting of a set of service providing nodes, S , each member of which provides an equivalent, quantitatively divisible service, the resource π , which may vary only in price. We assume that each member of S has an equal capacity for the provision of π , and that they cannot be relied upon to cooperate. We then imagine a large population of service users or buyers, B , each member of which aims to consume some of the resource π , at regular intervals. Our objective is to balance the load, such that all the service providers in S are providing an equal amount of π across the population of service users. Though this is a trivial problem when cooperation may be assumed, we wish to achieve this using self-interest, with no central control. Note that we are not modelling service providing nodes owned by competing businesses in the real-world, since then load-balancing would not be desirable; self-interested competition is instead artificially created in order to serve the purposes of the system owner.

At a given instant, a service provider, $s_i \in S$, advertises π at the price $p_{s_i}^\pi$ per unit. From a service user's point of view, this may be denoted as the offer X_{s_i} . Each service user, a buyer in this case, purchases some of the resource π should it be in their interest to do so at the price offered. The system iterates, with service providers able to adapt their prices to the market conditions over time. The actual provision of π may be regarded as instantaneous, such that it does not interfere with this mechanism.

For simplicity at this stage, we assume that the system proceeds in discrete time-steps, that each buyer $b_j \in B$ desires exactly one unit of π per time-step, and that each and every $s_i \in S$ has sufficient quantity of π available to satisfy all

the buyers in B should it be so requested. This final assumption is commonplace in the provision of information-based services, and is present in other related work such as [14]. We do not believe that its presence alters the underlying behaviour we demonstrate. We further simplify the model by stipulating that each service provider has no overhead cost associated with obtaining, producing or providing π . This assumption allows for buyers to purchase a tiny amount from each of a large number of sellers. Whether or not this is unrealistic will be determined by the application scenario.

Each time-step, each buyer, if it chooses to buy, may purchase any amount of π from any number of service providers in S , subject to the constraint that the total amount purchased per time-step is equal to exactly one unit. If no offer from any $s_i \in S$ is in its interest, the buyer may instead opt to purchase nothing. We therefore define q_{ij} to be the instantaneous quantity bought by buyer b_j from seller s_i . The constraints mean therefore that $\sum_{i=1}^{|S|} q_{ij} \in \{0, 1\}$ for all $b_j \in B$.

The quantity of π sold by a given seller s_i at a given time-step, its *load*, l_{s_i} , is therefore:

$$l_{s_i} = \sum_{j=1}^{|B|} q_{ij}. \quad (1)$$

Our stated objective is even load-balancing. In previous work [14], this is defined as being the minimisation of the variance between the service providers' loads, v_l , as shown in equation 2.

$$v_l = \frac{\sum_{i=1}^{|S|} (l_{s_i} - \mu)^2}{|S|}, \quad (2)$$

where

$$\mu = \frac{\sum_{i=1}^{|S|} l_{s_i}}{|S|}.$$

However, we prefer instead the measure, d_l , a normalised measure of mean absolute distance from the ideal load (here referred to as *NMA distance*), as described in equation 3. We find that this scales better with respect to $|S|$, making comparisons between simulations with different numbers of sellers simpler. Hence a high d_l indicates an uneven load, while a perfectly even load leads to a value of zero. We define μ as before.

$$d_l = \frac{\sum_{i=1}^{|S|} |l_{s_i} - \mu|}{|S|} \div \frac{2|S| - 2}{|S|^2}. \quad (3)$$

Both buyers and sellers accrue a payoff, or utility gain, from their interactions in the marketplace. For buyers, this will be the value they associate with the price paid subtracted from the value they associate with the purchased resource. We assume that buyers are self-interested, such that they attempt to maximise their utility. However, as in previous work [14], we do not assume that they are

hyperrational, behaving in an all-or-nothing manner in favour of the instantaneously most attractive option, since this behaviour may expose the buyer to a degree of risk.

Instead, we investigate buyers who are *risk-averse*, preferring to spread their purchases across a number of sellers. At each time-step, each buyer looks at the available offers, $X_{s_i} \forall s_i \in S$, and purchases a proportion of their total desired resource from each seller, relative to the expected utility gain from selecting the offer from that seller. An alternative would be to motivate risk-aversion through the game itself, however we prefer not to complicate the model, instead favouring the clarity gained by the assumption of risk averse behaviour.

Therefore firstly a buyer considers a unit transaction of π from each seller. The instantaneous expected utility, or payoff, for a buyer, b_j , in purchasing one unit of π from the seller s_i at the current offer,

$$E[U^{b_j}(X_{s_i})] = u^{b_j}(p_{s_i}^\pi, \pi), \quad (4)$$

where $E[U^{b_j}(X_{s_i})]$ represents buyer b_j 's expected payoff from accepting offer X_{s_i} , and $u^{b_j}(p, \pi)$ is the buyer's utility function over the goods: money and π .

The buyer then purchases a proportion of his total desired π from each seller. Any sellers which would provide a negative payoff are ignored. Recalling that q_{ij} is the quantity of π purchased by buyer b_j from seller s_i ,

$$q_{ij} = \frac{E[U^{b_j}(X_{s_i})]}{\sum_k E[U^{b_j}(X_{s_k})]}, \quad (5)$$

which ranges over k for which the expectation is positive.

The market model we have described here is an example of a Bertrand Game [17]. This is where two or more sellers compete by simultaneously setting prices for equivalent goods, and buyers then decide the quantity to purchase from each seller in a rational utility-maximising manner. The theoretical non-cooperative Nash equilibrium outcome is the Bertrand equilibrium, at which all sellers charge their reserve price and share the market equally. It is the equal sharing of the market at the Bertrand equilibrium which provides us with load-balancing.

However, Bertrand competition relies on the presence of a number of potentially unrealistic assumptions. Two of these are of particular interest. Firstly, Bertrand competition assumes no collusion between sellers. Cheung et al. showed that if sellers are able to reliably predict the behaviour of competitors, then they may implicitly collude in order to raise prices and hence their payoffs [16]. This behaviour is not observed in our model however, since the sellers do not retain historical information concerning each other. Such an ability might well be unfeasible in very large systems [16], though clearly a heterogeneous set of sellers, differentiated by strategic ability is likely to lead to an uneven market, and hence an uneven resource allocation.

Secondly, Bertrand competition assumes that sellers compete only on price, and are otherwise unable to differentiate their products in the market. This is unlikely, since more realistic service providing nodes' differing quality of service will provide product differentiation. Once competition exists other than on

price, Bertrand competition no longer applies, though other potentially useful equilibria will exist.

However, the introduction of issues other than price, and with it the heterogeneity of buyers, will also introduce the likelihood of ‘price-war’ behaviour. Kephart et al. showed that in certain circumstances the competitive behaviour of (non-colluding) sellers would lead to a never-ending series of price-wars [18]. This is shown to be the case when services are described over a number of issues, and a heterogeneous population of buyers have preferences over those issues such as to exist in different *niches* of the market. In these circumstances, sellers undercut each other in order to gain a greater market share and hence greater payoff. However, once the price has become sufficiently low, it becomes rational for the sellers to switch to providing for another niche of buyers, and the competition begins again. Once competition in that niche has driven the price down, the sellers will switch to a different niche, and so on. This result will be important to consider in future work.

4 Evolutionary Market Agents

Sellers also accrue a payoff, their *revenue*. In our model, this is income from the sale of π . At a given instant, the revenue of a seller, $s_i \in S$, is therefore:

$$r_{s_i} = \sum_{j=1}^{|B|} p_{s_i}^{\pi} q_{ij} . \quad (6)$$

Ideally, a seller will wish to maximise its revenue by increasing both its price and its market share, however as we have seen, the market share will depend upon the relationship between its price and those of its competitors.

An evolutionary market agent operates on behalf of each seller, with the self-interested objective of maximising its revenue. Using evolutionary computation techniques, the agent evolves the market position of its host over time. In this model, a market position consists simply of price, therefore each individual represents a real-valued price. For each interaction in the market, the price encoded by an individual is adopted, and the resulting payoff provides its fitness.

The evolutionary algorithm for seller s_i ’s agent proceeds as follows:

1. Decide upon the design parameters to be used: initial price range, population size and mutation factor. In the simulations described, an initial price range of 0 to 500 was chosen, along with a population size of 10 and mutation factor (α) of 0.1.
2. Generate an initial population, P , and set $k = 1$. Each individual in P is a real value, drawn from the uniform random distribution $[0, p_{max}]$.
3. *Initial fitness testing*
 - (a) Set the seller’s offer to the value of the first individual in P , and enter the market for one market time-step. Record the seller’s revenue, r_{s_i} as that individual’s fitness.

- (b) Repeat for the next individual in P , until all initial individuals have been fitness tested in the market.
- 4. *Probabilistic tournament selection*
 - (a) Select four individuals, x_1, x_2, x_3 and x_4 from P , 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 value equal to the mid-point of the values 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 P , and insert o into P .
- 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 revenue, r_{s_i} as o 's fitness.
- 9. Repeat from step 4.

5 Simulation Results

5.1 A Baseline Scenario

We firstly consider a small scenario with two service providers, such that $S = \{s_1, s_2\}$, each providing the resource π , at prices $p_{s_1}^\pi$ and $p_{s_2}^\pi$ respectively. Both s_1 and s_2 make use of an evolutionary market agent, as described in section 4 in order to determine these prices at each time-step.

We begin with 10 buyers (the service users), with identical linear utility functions,

$$u^{bj}(p_{s_i}^\pi, \pi) = 375 - p_{s_i}^\pi. \quad (7)$$

This represents a buyer being indifferent between a unit of π and its cost at a price of 375. This is an arbitrary positive value, and has little impact other than to provide a range of positive payoff values for a range of prices. An alternative approach would be to have the objective of minimising the buyer's spending, though this would remove the notion of a value placed upon π by the buyer. Actively considering negative buyer payoffs would introduce the question of buyer motivation, and considering a payoff able to range over both positive and negative values would slightly complicate the buyer's decision function for no gain. Linear utility functions are used to give an estimation of a service user's expected preferences, though the exact function will of course depend on the specific service and its users. The success of our approach with other forms of utility function remains a topic for research.

Figure 1 shows the normalised mean absolute distance from the ideal load, d_l between s_1 and s_2 , over time, for a typical run and across 30 independent runs.

Fig. 1. Load balancing performance over time: 2 sellers and 10 buyers. Typical run (left) and mean and standard deviation over 30 independent runs (right).

Here, the NMA distance drops quickly, indicating the approach’s ability to achieve a roughly even load between the two service providers in a short time. This is due to the evolutionary agents’ competitively co-evolving their prices to within close proximity of each other, resulting in roughly even shares of the market. Due to diverse populations within each agent however, their prices, and hence the allocation of resources, continue to vary.

Following these exploratory fluctuations, the NMA distance then stabilises as the populations converge. At this point, the load-balance is highly equal. It is important to note that Kephart’s price-wars [18], which would result in an unstable load allocation, do not occur here. This is the case since the sellers describe the service π over only a single issue, price, and hence our buyer population is homogeneous.

5.2 A More Complex Scenario

Due to the distributed, decentralised nature of our approach, it is highly scalable. The complexity of the agents’ evolutionary algorithms remains constant with respect to the number of buyers, whilst the complexity of the buyers’ algorithm grows only linearly with respect to the number of sellers.

Figure 2 shows both a typical run and mean and standard deviation calculated over 30 independent runs for $|S| = 1000$, $|B| = 10,000$.

Here results are of a similar form to the previous simulation. This shows the power of our approach to achieve load-balancing in a large, decentralised system where no individual has any desire in favour of this behaviour. Indeed to the contrary, though our outcome is similar to that in [14], our approach is novel in that it relies on self-interested behaviour. The presence of larger populations of agents appears to lead to more reliable results; the approach clearly scales well.

Fig. 2. Load balancing performance over time: 1,000 sellers and 10,000 buyers. Typical run (left) and mean and standard deviation over 30 independent runs (right).

6 Conclusions and Future Work

We have presented a resource allocation problem, motivated by an emerging computational paradigm; dynamic, decentralised, service-based systems. Based on the mechanism proposed by Kuwabara et al. [14], we have described a decentralised, evolutionary market-based approach, which makes use of Bertrand competition between self-interested sellers to achieve load-balancing. No cooperation is assumed. We believe that our approach is more suited to this scenario than other decentralised load-balancing mechanisms, since it accounts for self-interested utility maximising individuals.

Unlike the majority of market-based systems, our approach requires no central point of control or auctioneer. Agents have no knowledge of the size of the marketplace or any history. It has no point which is weaker than any other, and is hence both scalable and robust to failure. Sellers instead have the ability to advertise their prices through a broadcast mechanism.

A future, more realistic model may include a second issue with which to describe quality of service. Sellers could then achieve product differentiation. Outcomes in such a scenario are likely to be more complex than in the model investigated here, in which effectively sellers behave as Dutch auctioneers, since in general their prices only reduce over time.

It is also likely that more realistic scenarios will be dynamic, where service providers may be added to or removed from the system. In addition, the population of buyers may change over time, and there may also be external disturbances. It would be desirable for the system to autonomically achieve a new load-balance in the presence of such changes.

Finally, more advanced tuning of the evolutionary algorithm used in the sellers' evolutionary market agents should improve system performance, and analysis of the algorithm's properties, especially in dynamic environments, will be useful in achieving this. Adaptive mutation, in order to allow sellers to explore the market widely when necessary, but to compete without reckless price changes, will be one potentially useful technique.

References

1. Singh, M.P., Huhns, M.N.: *Service-Oriented Computing: Semantics, Processes, Agents*. John Wiley and Sons, Chichester, West Sussex (2005)
2. He, M., Jennings, N.R., Leung, H.: On agent-mediated electronic commerce. *IEEE Transactions on Knowledge and Data Engineering* **15**(4) (2003) 985–1003
3. Papazoglou, M.P., Traverso, P., Dustdar, S., Leymann, F.: *Service-oriented computing: Research roadmap*. (2006)
4. Gupta, A., Stahl, D.O., Whinston, A.B.: The economics of network management. *Communications of the ACM* **42**(9) (September 1999) 57–63
5. Clearwater, S.H., ed.: *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, Singapore (1996)
6. Ardellini, V.C., Olajanni, M.C., Yu, P.S.: Dynamic load balancing on web server systems. *IEEE Internet Computing* **3**(3) (May - June 1999) 28–39
7. Alfano, R., Caprio, G.D.: Turbo: an autonomous execution environment with scalability and load balancing features. In: *Proceedings of the IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and its Applications (DIS06)*. (2006) 377–382
8. Cliff, D., Bruten, J.: Simple bargaining agents for decentralized market-based control. Technical Report HPL-98-17, HP Laboratories, Bristol, UK (1998)
9. Esteva, M., Padget, J.: Auctions without auctioneers: Distributed auction protocols. In: *Lecture Notes in Artificial Intelligence*. Volume 1788. Springer-Verlag, Berlin, Germany (2000) 20–28
10. Kikuchi, H.: (m+1)st-price auction protocol. In: *Proceedings of the 5th International Conference on Financial Cryptography*, London, UK, Springer-Verlag (2002) 351–363
11. Hausheer, D., Stiller, B.: Peermart: The technology for a distributed auction-based market for peer-to-peer services. In: *Proceedings of the IEEE International Conference on Communications*. Volume 3. (2005) 1583–1587
12. Haque, N., Jennings, N.R., Moreau, L.: Scalability and robustness of a network resource allocation system using market-based agents. *Netnomics* **7**(2) (2005) 69–96
13. Ramasubramanian, V., Sirer, E.G.: Perils of transitive trust in the domain name system. Technical Report TR2005-1994, Cornell University, Ithaca, New York, USA (2005)
14. Kuwabara, K., Ishida, T., Nishibe, Y., Suda, T.: An equilibratory market-based approach for distributed resource allocation and its applications to communication network control. In Clearwater, S.H., ed.: *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, Singapore (1996) 53–73
15. Midgley, D.F., Marks, R.E., Cooper, L.G.: Breeding competitive strategies. *Management Science* **43**(3) (1997) 257–275
16. Cheung, Y., Bedingfield, S., Huxford, S.: Monitoring and interpreting evolved behaviours in an oligopoly. In: *Proceedings of the IEEE International Conference on Evolutionary Computation*. (1997) 697–701
17. Mas-Colell, A., Whinston, M.D., Green, J.R.: *Micro-Economic Theory*. Oxford University Press, Oxford (1995)
18. Kephart, J.O., Hanson, J.E., Sairamesh, J.: Price-war dynamics in a free-market economy of software agents. In: *Proceedings of the Sixth International Conference on Artificial Life*, Cambridge, MA, USA, MIT Press (1998) 53–62