Resource Allocation in Decentralised Computational Systems: An Evolutionary Market-Based Approach

Lewis, Peter; Marrow, Paul; Yao, Xin

DOI:

[10.1007/s10458-009-9113-x](https://doi.org/10.1007/s10458-009-9113-x)

Citation for published version (Harvard):

Lewis, P, Marrow, P & Yao, X 2010, 'Resource Allocation in Decentralised Computational Systems: An Evolutionary Market-Based Approach', *Autonomous Agents and Multi-Agent Systems*, vol. 21, no. 2, pp. 143-171. <https://doi.org/10.1007/s10458-009-9113-x>

[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.

Resource Allocation in Decentralised Computational Systems: An Evolutionary Market-Based Approach

Peter R. Lewis · Paul Marrow · Xin Yao

2010

Abstract We present a novel market-based method, inspired by retail markets, for resource allocation in fully decentralised systems where agents are self-interested. Our market mechanism requires no coordinating node or complex negotiation. The stability of outcome allocations, those at equilibrium, is analysed and compared for three buyer behaviour models. In order to capture the interaction between self-interested agents, we propose the use of competitive coevolution. Our approach is both highly scalable and may be tuned to achieve specified outcome resource allocations. We demonstrate the behaviour of our approach in simulation, where *evolutionary market agents* act on behalf of service providing nodes to adaptively price their resources over time, in response to market conditions. We show that this leads the system to the predicted outcome resource allocation. Furthermore, the system remains stable in the presence of small changes in price, when buyers' decision functions degrade gracefully.

Keywords decentralised systems · market-based control · coevolution · load balancing · resource allocation · self-interested agents

A preliminary version of this paper [36] was presented at *Parallel Problem Solving From Nature X*.

Peter R. Lewis
School of Computer Science,
University of Birmingham, UK
Tel.: +44 121 414 8201
Fax: +44 121 414 5370
E-mail: p.r.lewis@cs.bham.ac.uk

Paul Marrow
BT plc,
Adastral Park,
Ipswich, UK
E-mail: paul.marrow@bt.com

Xin Yao
School of Computer Science,
University of Birmingham, UK
E-mail: x.yao@cs.bham.ac.uk

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 [50]. In such systems, both service providing and service using nodes may arrive, organise and dissipate, as computational capabilities are formed and reformed as needed, without reference to a central authority or coordinator.

As the Internet matures, it is predicted that the majority of the interactions it carries will be carried out by autonomous agents on behalf of their owners [29]. In such distributed systems, which we consider in the broadest sense, where there exists a distribution of work or resource about a network of nodes, 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 [40]. As part of this, there remains the problem of how to allocate such resource amongst the nodes [35].

From an engineering perspective, any resource allocation problem will have an objective: a desired allocation or outcome. A common example of this is a balanced load, where that the task of providing a resource is shared evenly between a group of nodes. More complex objectives may involve stable, uneven allocations, where account is taken of other factors. Such factors might include users' preferences over quality of service issues, underlying costs to the service provider, or differences in the ability of nodes to provide an equivalent resource. The ability to allocate resources in a desired configuration, in a scalable and robust manner, is essential.

In many systems, nodes are either assumed or designed to be *self-interested*, each wishing, for example, to maximise its allocation. In these cases, the total demand for a resource may exceed its total supply. The allocation of resources to individuals in such scarcity is the classic problem studied by economics. In the human world, building on the simple idea of bilateral exchange, it is often approached using the structure and rules of a market, and its tool, price.

Externality pricing has been proposed for the provision of computational services, which are otherwise virtually zero cost per-use. Gupta et al. [27] argue that this approach, where service users self-select their quantity based on price, is preferable to the alternative of provider or regulator enforced quantity limits: service rationing.

Markets can also be used in settings where it is impractical or unwanted to associate monetary payment with the resource allocation process. Indeed, market-based control is a broad approach to resource allocation in a wide range of real-world applications [14]. In these scenarios, artificially created markets and pricing can enable stable, efficient, self-organising resource allocation to be achieved. For example, Wolski et al. [56] describe how market economies might be used to allocate resources in the computational grid. In this model, applications treat resources as interchangeable commodities, which may be provided by anyone. They argue that in this scenario, resource decisions must be self-interested.

However, the market mechanisms employed in these systems to date, though operating in a distributed manner, often rely upon some form of centralised coordination or control, such as a central auctioneer, specialist, or some set of super-nodes [16]. We believe that this lack of full decentralisation leads inevitably to unfulfilled potential in terms of both scalability and robustness to failure.

In this paper, our contribution is twofold. Firstly, we demonstrate the ability of a posted-offer market to achieve a range of desired resource allocations in a decentralised computational system where nodes are self-interested, outlining a game theoretic method for outcome allocations to be predicted in a given scenario. Secondly, we use the coevolution of

sellers' offers to demonstrate the achievement of the predicted outcomes in simulation. We propose this as a potential approach to realising self-interested adaptive pricing under the assumption of private information present in the posted-offer model.

This paper extends our previous work in a number of ways: providing a theoretical underpinning for the experimental results presented in [36], and extending the work on uneven resource allocations in [37] to many-node systems. Throughout, we consider a variety of buyer behaviours, including mixed populations, and their effect on the outcomes. Finally, we explore the scalability of our approach.

The remainder of this paper is structured as follows. In section 2 we discuss a number of existing classical and market-based approaches to resource allocation, highlighting why we believe a new approach is required in systems with no central control or coordination. In section 3 we formulate a problem for study, based on the above scenario. We then describe our fully decentralised market-based approach and introduce the buyer behaviours we investigate. In section 4 we analyse the expected behaviour of the system in a variety of instances of the problem, with each buyer type, and in section 5 we describe our *evolutionary market agents*. In section 6 the behaviour is demonstrated experimentally in simulation, where we also investigate mixed buyer populations and scalability. We conclude the paper, and discuss future prospects and research directions in section 7.

2 Related Work

2.1 Classical Approaches to Resource Allocation

Classically, resource allocation objectives are achieved in a centralised manner, often relying on a single node responsible for, say, load balancing [5]. A balanced load, though by no means the only interesting outcome, is often used as an example of a desired resource allocation. Load balancing is in itself interesting, since it is useful in numerous real-world scenarios, including telecommunications networks, road networks and electricity and water distribution networks. In many of these domains, even in very large scale systems, centralisation is the usual approach taken [5].

Resource allocation techniques can be divided into two groups, stateless and state-based [10]. Perhaps the most widely known and easily understood stateless approach, used to balance the load on web servers, is round-robin DNS. A more complex example is proportional share scheduling [35], in which resources are allocated to jobs according to a set of pre-determined weights. However, stateless approaches such as this are unable to take account of current server load or availability, leading to no guarantee that the desired outcome is achieved. Simple state-based extensions permit the usage of information about the resources being managed, and enable the proximity to the desired allocation to be measured. Examples of state-based resource allocation approaches include those which make use of geographical information and previous usage levels in order to determine an appropriate allocation of resource. A useful review and comparison of these approaches in the web server domain may be found in [10].

2.2 Decentralised Approaches

Centralised resource allocation methods do however have a number of drawbacks [23]. These include:

- the requirement that the environment remain static while the central coordinator is calculating the optimal resource allocation,
- that the coordinator has global knowledge of the system and all nodes within it,
- that all coordination messages must route through the central point, both counteracting the benefit from having resources distributed about the network, reducing scalability [1] and creating a fundamentally brittle system [3].

This brittleness may be mitigated against to a certain degree, with backup coordinator nodes, however even in these cases the wider system is reliant upon the existence and performance of a small number of key nodes. Failure at these key points in the network may well cripple wider functionality, at best [45]. These drawbacks lead to the need for a truly decentralised approach to the allocation of resources that does not rely on a central coordinator [23].

In the field of grid computing, examples include the hierarchical agent-based approach found in [9], and TURBO [1]. In the latter, allocations are achieved through the reliance on altruistic behaviour between cooperating peers, which collaborate in order to reach a global objective.

Balanced overlay networks are another effective and generic technique for balancing the load across a decentralised network [6]. In this approach, service providing nodes present an estimation of their availability to other local nodes to which they are connected. Newly arriving jobs take a random walk through the network and select the providing node with the highest availability. Upon accepting and completing a job, a provider node updates its availability estimate. Here, service users are self-interested within the bounds of the providers observed within their random walk, though the providers themselves are relied upon both to provide an honest and accurate account of their availability and to facilitate the random walk by exposing their local connections. In the case where such cooperation may not be relied upon, it is likely that the system's performance would deteriorate significantly.

In decentralised peer-to-peer storage systems where both cooperation and global knowledge of the system may be assumed, Surana et al.'s [52] approach may be used. Here the case is considered when moving loads around the network also uses bandwidth. Their objective is therefore a balance between achieving an even load and minimising the amount of load moved. Their fully decentralised approach is, in effect, tantamount to performing a centralised calculation at each node, periodically requiring cooperative reassignment of a load, based on global knowledge of the system.

A non-cooperative, decentralised approach exists in the domain of downloading replicated files. Dynamic parallel access schemes [48,12] make use of self-interested smart clients to increase the speed of file downloads. It is not clear however, how this approach might be generalised to other service-based systems.

2.3 Existing Market-Based Approaches

Making the case for economics-inspired approaches to resource allocation, Buyya et al. [7] argue that classical approaches either rely on central coordination and complete knowledge, or else cooperation between nodes. Khan and Ahmad [32] show that in any cooperative approach, global optima can only be achieved when all the nodes cooperate. In a discussion of lessons learnt from experience with load management in giant-scale web services, Brewer [5] proposes the idea of incorporating, into a request for a resource, a notion of its value or cost. This, along with the use of smart clients, would allow for responsive adaptation in

the presence of changes to the network, as well as graceful degradation. Similarly, Gupta et al. [27] argue that in the provision of virtually zero cost per-use computational services, a mechanism involving pricing and user self-selection is preferable to the alternative of provider or regulator enforced limits or rationing.

The application of economic ideas to resource allocation problems in computational systems is approached in the field of market-based control, an introduction to which may be found in [14]. Using the terminology of Casavant and Kuhl's [11] taxonomy of scheduling in distributed computing systems, this is a family of distributed mechanisms for dynamic global resource allocation in a network, relying on the non-cooperative and self-interested behaviour of nodes. Fundamentally, these approaches rely on the rational behaviour of self-interested peers, represented by agents attempting to maximise their payoff from interactions or exchanges within some market environment. The theories of microeconomics predict that, through such repeated exchanges, efficient resource allocations may be achieved. A brief and accessible introduction to some relevant microeconomics may be found in [17].

It is important to note that the non-cooperation of agents does not imply self-interest. Indeed, in Khan and Ahmad's [32] study of various games-based resource allocation methods, non-cooperative agents bid for jobs based on an honest estimation of the estimated time to complete a job, without consideration of the benefit they expect to derive from the bid. This is clearly not the self-interested behaviour relied upon by economics. Self-interest, at least in a boundedly rational sense, is a key assumption of economics inspired computation, and without it, the game theoretic approach is not useful.

Agents in a market-based system may interact through any of a number of different mechanisms. Common examples include English, Dutch and Vickrey auctions, as well as double auctions such as the Continuous Double Auction. Research in the field of automated mechanism design also suggests that other less obvious auction mechanisms may lead to more efficient outcomes [15,55].

However, both Cliff and Bruten [16] and Eymann et al. [23] note that due to the mechanisms employed, a large proportion of market-based control systems are not truly decentralised, relying on a centralised price-fixing process rather than the participants between them determining prices. This is true of Wolski et al.'s [56] G-Commerce model, which relies upon a central market-maker. Cliff and Bruten [16] argue that the presence of such a centralised process or component removes the primary advantage of using a market-based system: its robust, self-organising properties.

A number of distributed auction mechanisms have also been proposed [19,28,33], which do not rely on one central coordinating node. These approaches reduce the fragility associated with reliance upon a single point, provide more scalability and allow for dynamic composition of auctions. Typically, either the central auctioneer is replaced by a number of local ones, which communicate through some secure means, or else the auctioneer role is fulfilled by a spare, disinterested node. Double auctions, for example, though relying on a specialist to match bids and asks [42], may be decentralised by the presence of multiple specialists between which the participants may choose [39]. These techniques do reduce bottlenecks at certain points within the network and the removal of a single node cannot lead to system-wide failure. However, similarly to the replicated round-robin DNS approaches discussed in section 2.1 above, the system is still largely reliant on a small subset of its nodes.

However, it may be possible in systems such as this to scale up the number of auctioneers or specialists, in order to achieve a suitable degree of redundancy and decentralisation. This issue is worth investigating further, though intuition suggests that a system which relies upon a set of super-nodes can never provide the level of robustness of a system without such a need, even if the super-nodes were present in abundance. Approaches such as this also

raise questions of incentive compatibility for those acting as super-nodes. Therefore, if an approach exists without the need for such complexity, it should be preferred.

2.4 Decentralised Market Mechanisms

Cliff and Bruten [16] conclude from their critique that, rather than depend upon a central node such as an auctioneer, market mechanisms should instead rely on the ability of intelligent agents to bargain between themselves in order to arrive at acceptable prices. This approach is taken in the AVALANCHE [21], and CATNET [22, 24, 23, 3] systems. These take inspiration from Agent-based Computational Economics (ACE) [53] in attempting to replicate the dynamics of human markets with complex cognitive agents. These approaches use highly developed strategies, as agents negotiate bilaterally in order to determine the provision of a resource. Furthermore, in the approach taken in CATNET, resources are relied upon to forward requests to neighbouring hosts, without any consideration of the effect of this on its own interests [23]. This appears to be at odds with the self-interested nature of the agents.

A somewhat simpler application of market principles to computational resource allocation can be seen in Spawn [54]. In this approach, consuming agents bid in sealed-bid auctions hosted by providing agents, for their resources. This again requires a high level of strategic ability on the part of consumers, as they must decide in which auctions to participate. Of course, consumers may win multiple auctions, and questions then arise of how to handle these situations.

Though not discussed in detail in their paper, Cliff and Bruten [16] also briefly mention retail markets as an alternative to auctions and bilateral negotiation. The mechanism used in modern retail markets is usually referred to as the *posted-price* or *posted-offer* model [43, 31], though in online content delivery it is sometimes referred to as the *quoted-price* model [30]. It is a fully decentralised approach to the determination of price [43] without the need for complex bilateral negotiation, and we believe that this provides a potentially simpler alternative.

Chavez et al. [13] use an approach of this type in Challenger, where offers are broadcast to the nodes in a network, though instead of using price, bids contain an honest reporting of a job's priority. Similarly to the example in [32] discussed in section 2.3 this not self-interested behaviour. Xiao et al. [57] describe their system GridIS, in which buyers broadcast job requests and sellers reply by posting offers to perform them at a price. However, the sellers used require the knowledge of certain global information in determining their price, both in the form of the latest accepted market price, which we consider to be private information, and also the level of aggregate supply of all the providers in the network. Our assumption of private information forbids this also.

Kuwabara et al. [34] propose an application of a posted-offer market to decentralised computational systems, observing the quantities provided at the equilibria at which the markets arrive. No central component, such as an auctioneer or specialist is used; prices are determined privately by the sellers and then posted via a broadcast mechanism. The buyers then decide the quantity to purchase from each seller.

2.5 Adaptive and Evolutionary Pricing

As in human retail markets, key to the efficiency of the posted-offer market in [34] is the assumption of iterated transactions. For example, where a large number of buyers arrive over a period of time, each seller is able to adapt its price such that the payoff from its transactions is maximised. Sellers therefore compete on price, over time. From a seller's perspective, we are therefore faced with the questions of finding the best price, and how this price should be adapted in response to market conditions. Of course, other sellers will also respond, as prices co-adapt. We call this the *adaptive pricing game*.

Critically however, though the sellers in Kuwabara et al.'s [34] system are non-cooperative, they are not self-interested, since they do not consider their payoff when determining their price. Instead, a seller node's previous usage is encoded directly as the seller's next price.

One effective way to model self-interested competitive behaviour computationally is with coevolution [18], an extension of evolutionary computation in which the fitness function is in part dependent upon the actions of others. Co-evolution has been used to learn strategies in repeated games such as the Iterated Prisoner's Dilemma [4] and also to find Nash Equilibria in games with continuous strategy spaces [51]. In the study of markets, coevolution has also been used in the optimisation of parametrised bidding and bargaining strategies [8,21,41,2]. Of course, in many cases, the optimisation is dynamic, as it is performed against a moving optimum, as the competitors also update their strategies.

In [44], Price demonstrates that, rather than optimising the parameters of a particular bidding strategy, certain classic competitive behaviour can be achieved by coevolving prices directly. By drawing an analogy between payoff and evolutionary fitness, coevolution is used to drive competition in the market. In a sense, from a seller's perspective, evolution itself is the strategy. Amongst other examples, Bertrand competition is demonstrated by competitively coevolving two sellers' prices.

In the approach described in this paper, we demonstrate the ability of a posted-offer market to achieve a range of desired resource allocations in a decentralised computational system where nodes are self-interested. We then outline a game theoretic method for outcome allocations to be predicted in a given scenario. We use the direct coevolution of sellers' offers to demonstrate the achievement of the predicted outcomes in simulation, and propose this as a potential approach to realising self-interested adaptive pricing under the assumption of private information.

As such, this work extends preliminary experimental explorations [36] in which we demonstrated that sellers using *evolutionary market agents* in a posted-offer market may converge to an equilibrium at which the load is balanced evenly. In a two-node scenario, we further explored *evolutionary market agents* in the presence of service providing nodes with heterogeneous abilities to provide the resource. We showed in [37] that taking account of this heterogeneity through differences in the sellers' cost and valuations of the resource, enables us to achieve desired uneven resource allocations.

3 Problem Formulation

3.1 Scenario and Objectives

Our objective differs from that of many market-based resource allocation mechanisms, in that we are interested in achieving a particular outcome resource allocation in a given scenario. Our approach begins from the starting point of a desired allocation of resources which

the system designer or owner wishes to achieve. An artificial market is then created in order to bring this allocation into effect, under the assumptions of decentralisation and self-interest.

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 the members of S are self-interested. Nodes in S may vary in their ability to provide π , for example, one node may be able to perform twice as much work as another.

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. If s_i is a node in S and b_j is a node in B , we define q_{ij} to be the quantity of the resource π provided by s_i to b_j .

The total quantity of π provided by s_i at a given instant, its *load*, l_{s_i} , is therefore:

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

Our objective is to achieve a stable, desired configuration for the provision of π by the nodes in S . Such a configuration may be expressed by the vector $\mathbf{L}_S = \langle l_{s_1}, l_{s_2}, \dots, l_{s_n} \rangle$, where $n = |S|$. For convenience and ease of comparison between scenarios, we often normalise this vector by the total resource being provided.

In this way, any desired outcome allocation may be described. An evenly balanced load, for example, may be written as $\langle \frac{1}{n}, \frac{1}{n}, \dots, \frac{1}{n} \rangle$. Similarly, an allocation between six providing nodes, where the even numbered nodes provide exactly twice the resource as odd numbered nodes would be $\langle \frac{1}{9}, \frac{2}{9}, \frac{1}{9}, \frac{2}{9}, \frac{1}{9}, \frac{2}{9} \rangle$. Though this is a trivial problem when cooperation may be assumed, we wish to achieve this using only self-interest, in a fully decentralised manner with no central or regional control, and with only private information available.

It is worth noting that alternative approaches exist to the measurement of load. For example, we may wish to adopt a longer-term view, as in the *long-term proportional share* approach described by Lai [35]. Here, load is averaged over a realistically set sliding time window. Casavant and Kuhl [11] also discuss the need to determine an appropriate measure of the load for the particular system. In this paper however, we only consider the instantaneous load as described above.

3.2 Mechanism and Assumptions

At a given instant, a service provider, $s_i \in S$, advertises π at the price $p_{s_i}^\pi$ per unit via a broadcast mechanism. Each service user, a buyer in this case, then has the option of purchasing 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 independently adapt their prices to the market conditions over time.

At this stage, we are not considering the effect of the embodiment of our approach in any particular application or network environment, since this would make it unnecessarily difficult to analyse and understand the underlying behaviour of the model. This is of course an obvious area for future work. With this in mind, we make a number of simplifying assumptions:

1. That the system proceeds synchronously in discrete time-steps,
2. that each buyer $b_j \in B$ 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 $s_i \in S$ has sufficient quantity of π available to satisfy all the buyers in B 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 we do not believe that they 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 [34]. In our future work we will consider scenarios where there exist hard limits on supply. Finally, assumption 5 replicates the network conditions found in Wolski et al.'s G-Commerce [56].

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 purchase nothing. These constraints mean therefore that $\sum_{i=1}^{|S|} q_{ij} \in \{0, 1\}$ for all $b_j \in B$.

3.3 Buyer Behaviour

Both buyers and sellers accrue a payoff, or utility gain, from their interactions in the marketplace. For buyers, this is the value they associate with the price paid subtracted from the value they associate with the purchased resource. If we denote buyer b_j 's unit valuation of π by $v_{b_j}^\pi$, then its payoff from a unit transaction with s_i will be $v_{b_j}^\pi - p_{s_i}^\pi$. Since any buyer accepting a price above $v_{b_j}^\pi$ would lead to a negative payoff, this is its reserve or limit price. From a buyer's perspective, if a seller's price would not lead to a negative payoff for the buyer, then we say that the price is *acceptable*. We use S_{b_j} to denote the subset of S which contains exactly those sellers in S whose price is *acceptable* to buyer b_j . When buyers are homogeneous in so far as they have the same reserve prices, such that $v_{b_j}^\pi = v^\pi, \forall b_j \in B$, we are able to refer to a single set of acceptable sellers, $S_B \subseteq S$, such that $S_B = \{s_i : s_i \in S, p_{s_i}^\pi \leq v^\pi\}$.

We assume that buyers are self-interested and loosely rational, such that they would prefer higher payoffs. This is manifested through the following of some strategy, incorporating a decision function, which describes the quantity (which may be zero) to buy from each seller. This is a similar approach to that taken in [26], in which buyers may be either hyper-rational *bargain hunters*, seeking out the best possible price, or else *time savers* who will purchase from any *acceptable* seller, chosen at random. The possibility of complex buyer decision functions means that there may not be a straightforward mapping between seller abilities and valuations of the resource, and the subsequent outcome allocation.

Here, we consider the above problem in the presence of three buyer types: hyper-rational *bargain hunters* and *time savers* from [26] and the risk-averse *spread buyers* as described in [36].

3.3.1 Bargain Hunters

Bargain hunters always attempt to maximise their instantaneous payoff. In each iteration, they check the prices of all the sellers, selecting the one seller which provides the most attractive offer (i.e. the lowest price). If this price is acceptable, then the buyer purchases its

entire unit of π from that seller. In the event that more than one seller provides an equally attractive and acceptable offer, the buyer purchases an even proportion of π from each such seller.

3.3.2 Time Savers

Time savers do not check the price of every seller in the system when deciding from whom to buy. Instead, they select a seller at random, and if its price is acceptable, then they purchase the entire unit of π from that seller. If it is not, then they continue selecting previously unchecked random sellers until they find an acceptable price. If no seller has an acceptable price, then they purchase nothing.

3.3.3 Spread Buyers

Spread buyers are rudimentarily risk-averse, preferring to spread their purchases across a number of sellers. At each time-step, the buyer looks at all the available offers, and purchases a proportion of π from each seller with a price below $v_{b_j}^\pi$, relative to the expected utility gain from purchasing from that seller. Specifically, the quantity purchased by buyer j from seller i , with an acceptable price, is determined according to the following calculation:

$$q_{ij} = \frac{(v_{b_j}^\pi - p_{s_i}^\pi)}{(nv_{b_j}^\pi - \sum_{k=1}^n p_{s_k}^\pi)}. \quad (2)$$

3.3.4 Alternative Buyer Behaviours

It is worth noting that although we consider these three buyer strategies here, many other behaviours could be considered. For *spread buyers*, for example, an alternative might be to motivate risk-aversion through the mechanism itself. However we prefer not to complicate the model at this stage, instead favouring the clarity gained by the assumption of risk averse behaviour.

A behaviour with a high level of practical relevance is that of *sticky* buyers. Sticky buyers prefer to continue using service providers which they have previously used. They might switch to another seller if the long-term gain were greater than some threshold, and this might be motivated by an explicit cost to switching. However, at this stage we leave this as an item for future research.

3.4 Seller Behaviour

Sellers also receive a payoff, defined by their payoff function. We denote seller s_i 's payoff as P_{s_i} . In its simplest form, this is its revenue from the sale of π :

$$P_{s_i} = \sum_{j=1}^{|B|} p_{s_i}^\pi q_{ij}, \quad (3)$$

or indeed

$$P_{s_i} = p_{s_i}^\pi \times l_{s_i}. \quad (4)$$

However, such a payoff function takes no account of the heterogeneity of service provider abilities, as in [37]. One service providing node, for example, may be able to perform more work than another, due to an increased ability to provide the resource.

Such heterogeneities of service providing nodes' abilities may be represented by their selling agents using a notional cost of provision, along with different relative valuations of π when compared with the notion of money [37]. Since equivalent quantities of work are substitutable irrespective of which node performed them, we would expect that a node able to provide more work for the same cost would place a lower valuation on performing that quantity of work.

This valuation may be built in to our existing model in the form of a sellers' payoff function which takes account of the cost of the provision of π , $c_{s_i}^\pi$, and a preference weight, w_{s_i} , on the price:

$$P_{s_i} = \sum_{j=1}^{|B|} q_{ij} (w_{s_i} p_{s_i}^\pi - c_{s_i}^\pi), \quad (5)$$

or

$$P_{s_i} = (w_{s_i} p_{s_i}^\pi - c_{s_i}^\pi) \times l_{s_i}. \quad (6)$$

Clearly, a seller wishing to maximise its revenue would aim to increase both its price and its market share. However as we have seen from the buyers' behaviour, the market share will depend upon the relationship between its price and those of its competitors.

4 Predicted Outcomes

One motivation for employing an artificial market is that competition between self-interested sellers drives the system towards equilibrium. It is at this equilibrium that the system is stable, and thus we describe the allocation of resources in this stable state as the outcome resource allocation. Of course, in realistic scenarios, continuing changes to the system may lead to the existence of attractors rather than equilibria in the classic sense, though we do not consider this case at this stage.

The model we describe here is, in essence, a generalised version of the Bertrand game [38]. The classical Bertrand game consists of two sellers, both of whom offer to sell a certain good to a population of buyers. Each seller must decide what price to charge for the good, and then supply the quantity subsequently demanded by the buyers. The buyers in the classic Bertrand game behave hyperrationally, as with the *bargain hunters* described above, always buying from the seller with the lowest price, or half from each seller if the prices are identical.

In this game either seller can take the entire market by offering a price only fractionally lower than its competitor. However, since this applies to both sellers, the non-cooperative Nash equilibrium for the game is for both sellers to charge a zero-profit price. If each seller's costs are equal, then the equilibrium price for each seller will also be equal. This leads to the sellers sharing the market equally at equilibrium. As we demonstrated experimentally in [36], it is this basic idea which provides us with a balanced load in the simplest case.

However, in the more general case, where buyers may follow any of a number of strategies, calculating the expected outcome resource allocation may be a more complex task. Here we follow the methodology of determining the sellers' best response at each iteration, by solving payoff equations constructed from the given buyer behaviour. This enables us to

identify the Nash equilibrium position, where each and every seller's best response is equal to its previous position.

In the following discussions, we assume that the buyers do have an identical reserve price, $v^\pi = 300$, and therefore that we have a single acceptable set of sellers, S_B . Any seller in S but not in S_B will of course attract no buyers at all, and will hence receive no payoff and have a load of zero. For the sake of clarity, in the remainder of this section, we consider only those sellers in S_B .

4.1 Bargain Hunters

4.1.1 Simple Case

Let us first consider a scenario with two identical service providing nodes, such that $S = \{s_1, s_2\}$, each with equivalent abilities and costs of zero. In this case the payoff given in equation 4 will be sufficient for both sellers, so that

$$P_{s_1} = p_{s_1}^\pi \times l_{s_1}. \quad (7)$$

and

$$P_{s_2} = p_{s_2}^\pi \times l_{s_2}. \quad (8)$$

As in Bertrand competition, B is a large population of hyperrational buyers, *bargain hunters*, as described above. Recalling the decision function for these buyers, and our assumption that each buyer wishes to purchase exactly one unit of π , we may therefore say that

$$P_{s_1} = \begin{cases} |B| \times p_{s_1}^\pi & \text{if } p_{s_1} < p_{s_2}; \\ 0.5 \times |B| \times p_{s_1}^\pi & \text{if } p_{s_1} = p_{s_2}; \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

and the equivalent for s_2 respectively.

From a game theoretic perspective, given an observed value for their competitor's price, both s_1 and s_2 will wish to respond with the best response. In this case, this will be to undercut the competitor's price, if possible, in order to receive the payoff given by the first case in equation 9. The competing seller will of course act similarly, leading to a price war where each undercuts the other until their zero-payoff price is reached. Assuming that a seller would rather not participate than receive a negative payoff, once $p_{s_1} = p_{s_2} = 0$, the rational course of action is to maintain a price of 0, accepting the second case.

Recalling that the current load on a service providing node is given by equation 1 above, we therefore have that at equilibrium,

$$l_{s_1} = 0.5 \times |B|, \quad (10)$$

and

$$l_{s_2} = 0.5 \times |B|. \quad (11)$$

This is indeed an evenly balanced load, i.e.

$$\mathbf{Ls} = \left\langle \frac{1}{2}, \frac{1}{2} \right\rangle. \quad (12)$$

The theory of Bertrand competition (which is described more fully in [38]) demonstrates that when competing on price alone, two sellers are enough for the perfectly competitive outcome described here. Since the same logic applies to larger number of sellers, this evenly balanced outcome also holds for larger systems under the same assumptions.

4.1.2 Complex Case

Now let us consider the more complex case, where service providers may vary in their ability to provide the resource π . As described above, this may be represented in the seller's payoff function, by a heterogeneity in the cost of provision, $c_{s_i}^\pi$, and preference weight, w_{s_i} , as in equation 6. Competition on price will still occur as in the more trivial case. However, the lower limit price for each seller will depend upon its cost and preference weight. As we might expect, that seller (or those sellers) which are able to offer the lowest price will take the entire market, leaving those unable to compete with nothing.

We may illustrate this by means of a three-node example, such that $S = \{s_1, s_2, s_3\}$, where $c_{s_1}^\pi = c_{s_2}^\pi = 100$, $c_{s_3}^\pi = 150$ and $w_{s_1} = w_{s_2} = w_{s_3} = 1$. We therefore have the following payoff functions:

$$P_{s_1} = (p_{s_1}^\pi - 100) \times l_{s_1}, \quad (13)$$

$$P_{s_2} = (p_{s_2}^\pi - 100) \times l_{s_2}, \quad (14)$$

and

$$P_{s_3} = (p_{s_3}^\pi - 150) \times l_{s_3}. \quad (15)$$

Considering first s_1 , clearly if its offer is to be accepted by any buyer, indicated by a positive load, $l_{s_1} > 0$, a non-negative payoff is obtained if and only if $p_{s_1}^\pi \geq 100$. The same is true respectively of s_2 . However, for s_3 , this is the case if and only if $p_{s_3}^\pi \geq 150$. s_3 is therefore unable to compete with s_1 and s_2 and will discontinue participation in the price war once the best price dropped below 150. s_1 and s_2 however, will continue undercutting each other until their prices reach 100, at which point any further price-cut would leave them with a negative payoff themselves. However, since they both share the same limit price, they will remain at equilibrium with $p_{s_1} = p_{s_2} = 100$, and therefore $\mathbf{L}_S = \langle \frac{1}{2}, \frac{1}{2}, 0 \rangle$.

Now let us add a fourth node to S , s_4 , such that $c_{s_4}^\pi = 150$ and $w_{s_4} = 2.5$. s_4 has the following payoff function:

$$P_{s_4} = (2.5 \times p_{s_4}^\pi - 150) \times l_{s_4}. \quad (16)$$

Now, s_4 's cost is also 150, as with s_3 , however since its preference weight is 2.5, its payoff will not drop below zero for positive loads so long as $p_{s_4} \geq 60$. This is clearly a more competitive position than s_1 and s_2 are able to take. The expected allocation at equilibrium for this four-node system is therefore $\mathbf{L}_S = \langle 0, 0, 0, 1 \rangle$.

More generally, for purchases made by hyperrational *bargain hunters*, the seller or sellers able to offer the lowest price will share the load evenly between them. Any unable to offer this price will have a load of zero. The limit price for a seller s_i is $\frac{c_{s_i}^\pi}{w_{s_i}}$.

This example enables us to observe the *all or nothing* nature of a population of buyers made up entirely of *bargain hunters*, and it is here that the limit of Bertrand competition in its classic sense is reached. In order to elicit more complex resource allocations, we must turn to other buyer behaviours.

4.2 Time Savers

Intuitively, a population of *time savers* will possess less of the *all or nothing* nature of *bargain hunters*, as each will prefer potentially any seller whose price is acceptable.

Considering the trivial two-node example described above, what outcome do we expect with a population of *time savers*? Recalling that we only consider those sellers in S_B at present, we expect the payoff for s_1 and s_2 to be

$$P_{s_1} = \frac{p_{s_1}^\pi}{|S_B|} \quad (17)$$

$$P_{s_2} = \frac{p_{s_2}^\pi}{|S_B|} \quad (18)$$

Here, unlike with *bargain hunters*, there is no advantage to undercutting the price of a competing seller, since this will only serve to reduce the payoff. Instead, the dominant position is to charge the highest possible price, whilst still remaining in S_B ; the equilibrium is at $p_{s_1} = p_{s_2} = v^\pi$.

Similarly to *bargain hunters* however, since $p_{s_1} = p_{s_2}$, then $\mathbf{L}_S \approx \langle \frac{1}{2}, \frac{1}{2} \rangle$. Note that due to the probabilistic nature of the buyers' decision function, the load will tend towards this as the probabilities average out.

Indeed, the same is true of more complex cases with heterogeneous costs and preference weights. Since each seller's market share is independent of their price, given a seller's membership of S_B , any operation on the price will not affect the load at equilibrium.

4.3 Spread Buyers

4.3.1 Simple Case

For a population of spread buyers, as described above, the sellers' payoff functions for the simple two-node case are

$$P_{s_1} = \sum_{j=1}^{|B|} \frac{v^\pi - p_{s_1}^\pi}{2v^\pi - (p_{s_1}^\pi + p_{s_2}^\pi)} \times p_{s_1}^\pi, \quad (19)$$

and

$$P_{s_2} = \sum_{j=1}^{|B|} \frac{v^\pi - p_{s_2}^\pi}{2v^\pi - (p_{s_2}^\pi + p_{s_1}^\pi)} \times p_{s_2}^\pi. \quad (20)$$

Sellers s_1 and s_2 will each then attempt to maximise their respective payoff function as before. The outcome resource allocation occurs when the system is at equilibrium. Figure 1a illustrates the payoff function for s_1 , when $v^\pi = 300$ and $p_{s_2}^\pi = 250$.

Clearly, the best response price for s_1 is less than $p_{s_2}^\pi$; in fact in this instance it is 217.71. However, given this value as $p_{s_1}^\pi$, s_2 is then faced with the payoff function illustrated in figure 1b. Of course, s_1 will respond to this new value for $p_{s_2}^\pi$. Its best response is in this case 204.92. By iteratively calculating each seller's best response, we observe that this particular system is at equilibrium when $p_{s_1}^\pi = p_{s_2}^\pi = 200$.

Clearly at this point the market share, and hence load, of each seller is also equal: $\mathbf{L}_S = \langle \frac{1}{2}, \frac{1}{2} \rangle$.

Fig. 1 (a) Seller s_1 's payoff function with one competitor, and (b) s_2 's subsequent payoff function from s_1 's best response.

4.3.2 Complex Case

The same methodology may be employed in more complex cases to determine equilibria, and hence outcome resource allocations. For the three-node system discussed in section 4.1.2, the sellers' payoff functions are

$$P_{s_1} = \sum_{j=1}^{|B|} \frac{v^\pi - p_{s_1}^\pi}{3v^\pi - (p_{s_1}^\pi + p_{s_2}^\pi + p_{s_3}^\pi)} (w_{s_1} p_{s_1}^\pi - c_{s_1}^\pi), \quad (21)$$

$$P_{s_2} = \sum_{j=1}^{|B|} \frac{v^\pi - p_{s_2}^\pi}{3v^\pi - (p_{s_1}^\pi + p_{s_2}^\pi + p_{s_3}^\pi)} (w_{s_2} p_{s_2}^\pi - c_{s_2}^\pi), \quad (22)$$

and

$$P_{s_3} = \sum_{j=1}^{|B|} \frac{v^\pi - p_{s_3}^\pi}{3v^\pi - (p_{s_1}^\pi + p_{s_2}^\pi + p_{s_3}^\pi)} (w_{s_3} p_{s_3}^\pi - c_{s_3}^\pi). \quad (23)$$

Following the above method, we calculate that this system is at equilibrium when $p_{s_1}^\pi = p_{s_2}^\pi = 221.75$ and $p_{s_3}^\pi = 237.49$. At this point, none of the sellers may increase their respective payoff by unilaterally choosing a different price.

The allocation at these prices, therefore, is our expected outcome allocation. This may be calculated for each service provider using equation 6:

$$l_{s_i} = \sum_{j=1}^{|B|} \frac{v^\pi - p_{s_i}^\pi}{3v^\pi - (p_{s_1}^\pi + p_{s_2}^\pi + p_{s_3}^\pi)}. \quad (24)$$

In the example discussed here, we therefore have an outcome resource allocation of $\mathbf{L}_S = \langle 0.3575, 0.3575, 0.2850 \rangle$.

In other words, given the above valuation and cost parameters, we expect s_1 and s_2 to take 35.75% of the load each, and s_3 to take 28.5%.

Adding in the fourth seller at this point, we repeat the calculations. In the four-node example discussed in section 4.1.2, the prices at equilibrium are $p_{s_1}^\pi = p_{s_2}^\pi = 214.53$, $p_{s_3}^\pi = 233.25$ and $p_{s_4}^\pi = 200.73$. The outcome resource allocation is therefore $\mathbf{L}_S = \langle 0.2536, 0.2536, 0.1981, 0.2946 \rangle$.

4.4 Achieving Desired Outcome Allocations

We have so far described how to predict an outcome allocation, the allocation at equilibrium, given an initial configuration of nodes. How then, should the the sellers' valuation and cost values be set in order for the system to stabilise at a particular desired outcome?

The method for achieving this is essentially the reverse of the above process. Firstly, using the desired load vector as a starting point, we calculate the prices at equilibrium which give rise to the desired loads. Secondly, suitable values for the sellers' preference weight and cost parameters are chosen to give rise to equilibrium at the required prices. This method was originally described in [37], and will now be illustrated through the use of examples.

4.4.1 A Trivial Example

Suppose initially that we wish to achieve an evenly balanced load between two providers, in the presence of a population of homogeneous spread buyers. Our desired outcome allocation is $\mathbf{L}_S = \langle \frac{1}{2}, \frac{1}{2} \rangle$. Recall that these values are normalised by $|B|$. However, since the population of buyers is homogeneous, for each seller, $s_i \in S$ it must be that

$$\frac{\sum_{j=1}^B q_{ij}}{|B|} = q_{ik}, \forall b_k \in B, \quad (25)$$

where q_{ik} is the quantity of π bought by buyer b_k from seller s_i . Therefore, for any seller s_i , the non-normalised load $l_{s_i} = 0.5 \times |B|$ if and only if $q_{ij} = 0.5$, for all $b_j \in B$. In the case of homogeneous buyers, we may therefore calculate the required equilibrium prices as if there were a single buyer, which for the sake of consistency we call b_j .

Now, given the buyer decision function for spread buyers (from equation 2 above), we know that we require

$$\frac{(v_{b_j}^\pi - p_{s_1}^\pi)}{(v_{b_j}^\pi - p_{s_1}^\pi) + (v_{b_j}^\pi - p_{s_2}^\pi)} = 0.5, \quad (26)$$

which simplifies to

$$p_{s_1}^\pi = p_{s_2}^\pi. \quad (27)$$

In other words, an evenly balanced load will be the outcome allocation when both sellers quote the same price for π at equilibrium. In this trivial case, we have already seen in section 4.3.1 that this equilibrium may be achieved with zero cost values, $c_{s_i}^\pi$, and when $w_{s_1} = w_{s_2} = 1$.

4.4.2 A More Complex Example

Now let us consider a more complex desired outcome allocation, $\mathbf{L}_S = \langle \frac{2}{3}, \frac{1}{3} \rangle$. We wish for s_1 to provide twice the load of s_2 .

Following the same method, for a given homogeneous buyer valuation, $v_{b_j}^\pi$, we may calculate the required relationship between $p_{s_1}^\pi$ and $p_{s_2}^\pi$ at equilibrium to achieve our desired outcome is

$$\frac{(v_{b_j}^\pi - p_{s_1}^\pi)}{(v_{b_j}^\pi - p_{s_1}^\pi) + (v_{b_j}^\pi - p_{s_2}^\pi)} = \frac{2}{3}, \quad (28)$$

which simplifies to

$$p_{s_2}^\pi = \frac{v_{b_j}^\pi + p_{s_1}^\pi}{2}. \quad (29)$$

In order to achieve our desired outcome allocation, the prices $p_{s_1}^\pi$ and $p_{s_2}^\pi$ must conform to this relationship at equilibrium.

The question then arises of what preference weight and cost values, w_{s_1} , w_{s_2} , $c_{s_1}^\pi$ and $c_{s_2}^\pi$, can be chosen in order to satisfy this constraint. In our example, when $v_{b_j}^\pi = 300, \forall b_j \in B$, we therefore require that

$$p_{s_2}^\pi = 150 + \frac{p_{s_1}^\pi}{2}. \quad (30)$$

An appropriate method may then be used to find suitable parameters to satisfy this relationship at equilibrium. As is shown in [37], one approach is to fix the cost values $c_{s_1}^\pi$ and $c_{s_2}^\pi$, say at 1.0, in which case the constraint is satisfied when $w_{s_1} = 1.0$ and $w_{s_2} = 0.42105$. Therefore, by using the correct cost and weight values, our desired outcome allocation of $\mathbf{L_S} = \langle \frac{2}{3}, \frac{1}{3} \rangle$ may be achieved. A similar approach may be taken with other buyer decision functions.

5 Evolutionary Market Agents

We outlined in section 2.5 that the key problem to be addressed by a seller is that of how to adapt its price over time, such as to maximise its payoff through transactions in the market. We have also seen how buyers might make decisions about from whom to purchase the resource. We now turn our attention to the sellers' pricing agents, proposing the use of autonomous *evolutionary market agents*.

The calculation of a best response by a seller, as described in the previous section, makes two unrealistic assumptions: firstly, that each seller possesses current knowledge of all its competitors' prices, and secondly that it knows the decision function of every buyer. The first assumption may not be a problem in small systems, but in larger systems consisting of tens of thousands of nodes, this kind of global knowledge may not be available. However, the absence of knowledge of buyers' decision functions is a classic assumption in economic games, especially in the presence of a heterogeneous population.

Without knowledge of the buyers' decision functions, sellers cannot calculate a best response price, and must instead take an exploratory approach to finding it. Evolutionary computation provides an effective method of achieving this, given the *black box* nature of the problem. Indeed, a coevolutionary approach has long been used to model competing players in repeated games such as the Iterated Prisoner's Dilemma [4]. With the evolutionary algorithm outlined in this section, sellers expect to find good responses to the market. This approach allows us to achieve an approximation to the game theoretic outcome above, under the assumption of private information.

5.1 The Evolutionary Market Agent Algorithm

An *evolutionary market agent* operates on behalf of a particular service providing node in order to adaptively price its resource over time. Using evolutionary computation techniques, the agent evolves the market position of its host over time, in response to current market

conditions. A population of prices is evolved on-line, with payoff information from the market being used as fitness values, and the self-interested objective of payoff maximisation [36]. Competition between sellers is therefore driven by the coevolution of their respective *evolutionary market agents*.

As in our previous work [36], in this model a market position consists simply of price. Therefore each member of the population, an *individual* in evolutionary computation terminology, represents a real-valued price. For each interaction in the market, an individual's price is adopted, and the resulting payoff provides its fitness.

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

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, $\alpha = 0.1$.
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.

5.2 Predicted Outcomes with Evolutionary Market Agents

Most fundamentally to note about our *evolutionary market agents* is that they, as expected, do not have knowledge of the buyers' decision functions, and hence do not calculate a best response to the current state of the market. Their behaviour is instead exploratory and myopic. This has important implications for the outcome of their competition in the presence of *bargain hunters* and *time savers*, since these buyer types have a step in their decision functions.

By contrast, *spread buyers* have the benefit of giving rise to a seller market share function, and hence payoff function, which degrades gracefully as a seller moves away from the optimal price. Our *evolutionary market agents* lend themselves better to this behaviour. As

an illustration, consider the scenario when a seller s_1 is competing with s_2 , where $v^\pi = 300$ and $p_{s_2}^\pi = 250$. Figure 2 illustrates the payoff s_1 can expect to receive as a function of its own next price, $p_{s_1}^\pi$, for the three buyer types we discussed in section 3.3.

Fig. 2 Seller s_1 's payoff function with one competitor, for each of the buyer types.

The step function is clearly visible here, for sellers dealing with either *bargain hunters* or *time savers*. We may expect therefore, that in the absence of full knowledge of the buyers' decision functions (i.e. the curves shown here), and relying upon a myopic exploratory heuristic, that a seller will oscillate between receiving a near-optimal payoff, and one of zero.

Such an oscillation between high and low payoffs is not conducive of stable outcomes. Indeed, for *bargain hunters*, the position of the step for a given seller will depend upon its competitors' prices. Since the same is also true of the other competitors, the exact equilibrium position of the step, given by the sellers' prices, will be fairly arbitrary.

For *spread buyers*, however, a small move away from the optimal price leads to a small change in payoff for the seller. The resulting impact upon the other sellers' payoff functions is also minimal. This incremental characteristic of a market of *spread buyers* leads to a less brittle, more predictable system.

6 Simulation Results

By using this coevolutionary approach we expect the system to reach an approximation to the predicted outcome. However, given the *evolutionary market agents'* stochastic, exploratory nature, and the uncertainty introduced by their myopic behaviour, statistical results gained experimentally are useful.

We now explore the scenarios described in section 4, and some larger, more complex examples, in simulation.

6.1 A Baseline Scenario

Figure 3 illustrates the evolution of this behaviour over time in a scenario with two service providing nodes, such that $S = \{s_1, s_2\}$. Both s_1 and s_2 are represented by an *evolutionary market agent*, as described above, each with a population size of 20 and a mutation factor of 0.1. B consists of ten homogeneous *spread buyers*, such that $v^\pi = 300$.

Fig. 3 Evolution of price (above) and load (below) of two service providing nodes over time, with a population of 10 *spread buyers*. Stable even load balancing emerges. Mean and standard deviation over 30 independent runs.

Figure 3 clearly shows the ability of our approach to achieve a roughly even load between the two service providers in a short time. This is due to the evolutionary agents' competitively coevolving their prices to within close proximity of each other quickly, resulting in roughly even shares of the market. Following these exploratory fluctuations, the loads stabilise as the prices converge to the predicted equilibrium, here 200. At this point, the loads are highly equal. Due to diverse populations within each agent's population however, their prices, and hence the allocation of resources, continue to vary to a small degree.

6.2 Comparing Buyer Types

We now turn our attention to the behaviour of the system in the presence of *bargain hunters* and *time savers*. Intuitively, from section 5.2 above, we expect oscillatory behaviour when

these buyers' decision functions, with their step characteristics, are present. When observing the load on service providing nodes in simulation, this unreliable behaviour, along with a high standard deviation, is indeed manifested. Figures 4 and 5 illustrate this with two sellers and a population of 10 *bargain hunters*, and *time savers* respectively. The behaviour illustrated above in figure 3, for *spread buyers*, is clearly more desirable.

Fig. 4 Evolution of load of two service providing nodes over time, with a population of 10 *bargain hunters*. The step in the buyers' decision function leads to oscillations and unpredictability. Mean and standard deviation over 30 independent runs.

The presence of the step in the decision function for both *bargain hunters* and *time savers* clearly leads to oscillations in the load allocated between the service providing nodes; the allocations do not remain stable over time. The graceful degradation of the *spread buyers*' decision function, however, leads to a highly stable and evenly balanced load.

6.3 Mixed Buyer Populations

As we have shown, one benefit of *spread buyers*' behaviour is the graceful degradation of a seller's market share as its price moves away from the optimum (or the optimum moves away as competitors update their prices). This smooth curve allows the sellers' evolutionary algorithm to easily find the optimum price. By contrast, as we saw in section 6.2, buyer behaviours with a step in their decision function lead to more erratic, less stable equilibria.

However, how disruptive is the presence of such buyer behaviour in a population otherwise consisting of *spread buyers*? The experimental result in figure 6 indicates a graceful degradation of performance in the presence of an increasing proportion of *bargain hunters* in the population.

Fig. 5 Evolution of load of two service providing nodes over time, with a population of 10 *time savers*. The step in the buyers' decision function has less of an oscillatory effect. Mean and standard deviation over 30 independent runs.

6.4 Scalability

Due to the distributed, decentralised nature of our approach, it scales well. This scalability leads to the behaviour demonstrated in the above simple scenario also being observed in much larger scenarios. Figure 7 illustrates the evolution of price and load for 1,000 sellers and 10,000 buyers in a typical run of the simulation. Figure 8 shows the mean and standard deviation of the load variance for 30 independent runs. It is clear that in this respect the approach is highly scalable.

Here results are of a similar form to the smaller simulation. Figure 9 shows the time taken for the system to reach within 1% of the predicted outcome allocation, from the initial allocation, as the number of sellers increases.

A further important question concerning the scalability of the approach, and one that is shared with many other market-based mechanisms such as CATNET [3], is that of how knowledge of offers is disseminated around the system. In some embodiments of the model, such as networks making use of wireless technology, broadcasting may indeed be achievable. In peer-to-peer networks, the broadcast facility may be emulated through algorithms such as flooding [47], distributed hash tables [46] or epidemic algorithms [20]. However, Ardaiz et al. [3] argue quite correctly that decentralised and self-interested resource allocation approaches should prefer decentralised, incentive-compatible discovery mechanisms.

Fig. 6 The proximity of the outcome resource allocation to the predicted outcome degrades gracefully with respect to the proportion of *bargain hunters* in a population otherwise consisting of *spread buyers*. Each data point is from 30 independent runs of a simulation with two sellers, each with identical preferences and costs, and 100 buyers, which vary along the x axis.

6.5 Uneven Resource Allocations

As described in section 4 above, by incorporating a notion of preference and cost into sellers' payoff functions, we may achieve stable, desired uneven resource allocations. We now illustrate this in simulation, considering the four-node case described in section 4.1.2, in the presence of 100 *spread buyers*. Table 6.5 compares the predicted and experimental outcomes for this example. The experimental results show the state of the system at iteration 1,000, for 30 independent runs.

Seller	Predicted Price	Experimental Price	Standard Deviation
s_1	214.53	214.79	6.41
s_2	214.53	217.68	6.10
s_3	233.25	233.26	5.98
s_4	200.73	201.62	10.34
Seller	Predicted Load	Experimental Load	Standard Deviation
s_1	0.2536	0.2503	0.0161
s_2	0.2536	0.2535	0.0156
s_3	0.1981	0.1999	0.0133
s_4	0.2946	0.2967	0.0220

Table 1 Comparison of predicted and experimental outcome allocations: uneven allocation example. Experimental results calculated over 30 independent runs.

Fig. 7 Evolution of price and load variance between 1,000 service providing nodes over time, with a population of 10,000 *spread buyers*. Results from a typical run.

Fig. 8 Load variance between 1,000 service providing nodes over time, with a population of 10,000 *spread buyers*. The approach is clearly highly scalable. Mean and standard deviation over 30 independent runs.

Fig. 9 Time taken to reach close proximity to the predicted outcome allocation (within 1%). Each data point represents the mean and standard deviation over 30 independent runs.

Clearly, the coevolutionary approach is able to achieve results very close to the predicted outcome, with a high degree of reliability. It is however, worth noting that the unusually high standard deviation of s_4 is due to an outlier in the data obtained. Since the data represents a snapshot of the system at iteration 1,000, and given the continued Gaussian mutation of prices about the equilibrium, occasional, temporary outliers are to be expected.

6.6 Design of the Evolutionary Algorithm

The evolutionary algorithm used here is not the result of extended experimental tuning or design, and as such we fully expect that it may be improved upon. We do not at this stage claim that our *evolutionary market agents* are superior to other already existing adaptive pricing algorithms, however, this is not our primary concern at this stage. Rather, we are keen to emphasise the sufficiency of the coevolutionary approach to drive competition between sellers and achieve the effects described. A useful extension to this work would be to perform a comparison with other suitable strategies, such as the Gjerstad-Dickhaut [25] and Roth-Erev [49] algorithms.

However, during the course of the experimentation, some results concerning the evolutionary algorithm itself were obtained which we now briefly report. Particularly, the decision to use either stochastic or deterministic selection and replacement operators affects the algorithm's performance. Two variants of each operator were tested. For selection, the variants were the stochastic selection method described in section 5 and a more deterministic version where the winner of each tournament was always chosen to be a parent. For replacement, the variants were the deterministic operator described in section 5, where the worst individual out of those in the tournament was always replaced, and an alternative where with a

probability of 0.1 the individual replaced was instead selected at random from those in the tournament. Combining these, the four variants of the algorithm tested were as follows:

- Both deterministic selection and replacement operators,
- Deterministic selection but probabilistic replacement,
- Probabilistic selection and probabilistic replacement, and
- Probabilistic selection but deterministic replacement.

In the experiments, all variants of the algorithm were able to find the equilibrium price, but differed on their behaviour once at the optimum. Figure 10 shows the standard deviation of each algorithm's price about the mean for a seller over 30 independent runs.

Fig. 10 Probabilistic selection, combined with the deterministic replacement operator enables the algorithm to remain closer to the equilibrium price.

As can be seen, the use of both deterministic operators led to a high and highly volatile standard deviation, indicating frequent significant fluctuations in price. The two variants with probabilistic replacement operators performed very similarly, with lower and significantly more stable standard deviations. The lowest standard deviation was achieved by the variant with the probabilistic selection but deterministic replacement operators. This is the algorithm described in section 5 and it is for this reason that it was chosen for the majority of our experiments.

We have yet to carry out a detailed analysis of the reasons for this behaviour, however intuitively we expected the probabilistic operators would improve performance. This is since, in the early part of the simulation, a price higher than the equilibrium might yield a high payoff for a seller, but once the other competitors had responded to this, it would no longer be possible for any seller to reach such a high payoff. This price would remain in the population, and frequently win tournaments, despite its adoption at a later iteration being unlikely

to reproduce the high payoff. The market would have moved on and its fitness would be out of date. Probabilistic selection ensures that other individuals, with lower fitnesses but perhaps more suited to the current market conditions, can get selected instead. It is interesting then that probabilistic replacement did not bring about a similar performance gain. However, though it might be useful to occasionally select individuals with only an average fitness in future iterations, it is likely that individuals with very low fitness would perform poorly in any scenario, and there would therefore be little to gain by retaining them. A more formal analysis of the populations over time will prove illuminating.

7 Conclusions and Future Work

7.1 Conclusions

We have described a resource allocation problem, motivated by an emerging computational paradigm: dynamic, decentralised, service-based systems. Making use of the retail-inspired posted-offer market model, we proposed a fully decentralised, evolutionary market-based solution, which uses competitive coevolution between self-interested sellers to achieve a desired resource allocation. We illustrated this by using the approach to achieve a balanced load, and a number of other stable outcome allocations.

We have analysed the predicted equilibria in the system, and hence demonstrated a methodology for, given an initial configuration, determining the outcome resource allocation in the presence of various buyer behaviours. We applied this methodology in order to validate both previous [36, 37] and new simulation results.

Furthermore, we explored the relationship between experimental performance, in simulation, and such predictions for three buyer behaviour models. We concluded that behaviours with a step in their decision function lead to unstable, erratic outcome allocations. Those with smooth decision functions however, such as *spread buyers*, lead to highly accurate and stable allocations, as well as being robust to small changes in price. We further showed that this performance degrades gracefully as the proportion of buyers with step functions in the population increases.

Crucially, our approach assumes not only a lack of cooperation between agents but also self-interest. We believe that it is more suited to this scenario than other decentralised approaches, since it accounts for such self-interested utility maximising behaviour. Unlike many market-based systems, our approach requires no central or regional point of control or coordination, such as an auctioneer or specialist. Only private information is available to the agents, as sellers have no knowledge of the size of the marketplace, the number of competitors or any history. Our approach is highly scalable. We showed that the time taken for the system to converge to close proximity of the predicted outcome allocation appears to grow asymptotically to the number of service providing nodes. Additionally, the system has no point which is weaker than any other, and is hence robust to failure.

7.2 Future Work

The future directions identified for this work fall into a number of distinct areas. Firstly, we aim to consider higher level strategies for both buyers and sellers, for example those able to make use of historical information, as well as others found in the literature such as the Gjerstad-Dickhaut [25] and Roth-Erev [49] algorithms. The outcomes from this will be

compared with existing results. Crucially, the strategies presented here act as price-takers, despite the effect they have on the market. Higher level strategies may be able to exploit this. As discussed in section 3.3.4, sticky buyers, who prefer to continue using service providers which they have previously used unless there is significant reason to switch, provide a model of potentially highly realistic behaviour. Furthermore, the degree of stickiness may be parametrised, creating space of buyer strategies to be explored.

Secondly, it is highly likely that more realistic scenarios will be dynamic, where service providers may be added to or removed from the system during its operation. In addition, the population of buyers may change over time, and there may also be external disturbances. It will be desirable for the system to automatically adjust to such changes, and also for us to be able to predict how quickly this is achieved. A future, more realistic model should include further issues with which to describe the service, such as a measure of quality of service. Sellers could then achieve product differentiation, able to attract buyers with different profiles. Outcomes in this case are likely to be more complex than in the model investigated here. A further consideration in a more realistic model might be a potential performance hit as a result of breaking a task up into several sub-tasks, as is done here. If this were to be the case, then this should be quantified and its effect built into the model.

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. An adaptive mutation factor, in order to allow sellers to explore the market widely when necessary, but to compete without reckless price changes, may be a useful starting point.

Acknowledgements This work is partially supported by a UK EPSRC/BT CASE studentship project on "Autonomic Solutions for Virtualised ICT Systems" and an EPSRC project (No. GR/T10671/01) on "Market Based Control of Complex Computational Systems".

References

1. 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), pp. 377–382 (2006)
2. Anderson, E.J., Cau, T.D.H.: Modeling implicit collusion using coevolution. *Operations Research* **57**(2), 439–455 (2009)
3. Ardaiz, O., Artigas, P., Eymann, T., Freitag, F., Navarro, L., Reinicke, M.: The catallaxy approach for decentralized economic-based allocation in grid resource and service markets. *Applied Intelligence* **25**(2), 131–145 (2006)
4. Axelrod, R.: The evolution of strategies in the iterated prisoners dilemma. In: L. Davis (ed.) *Genetic algorithms and simulated annealing*, pp. 32–41. Pittman, London (1987)
5. Brewer, E.A.: Lessons from giant-scale services. *IEEE Internet Computing* **5**(4), 46–55 (2001)
6. Bridgewater, J.S.A., Boykin, P.O., Roychowdhury, V.P.: Balanced overlay networks (BON): An overlay technology for decentralized load balancing. *IEEE Transactions on Parallel and Distributed Systems* pp. 1122–1133 (2007)
7. Buyya, R., Abramson, D., Giddy, J., Stockinger, H.: Economic models for resource management and scheduling in grid computing. *Concurrency and computation: practice and experience* **14**(13–15), 1507–1542 (2002)
8. Byde, A.: Applying evolutionary game theory to auction mechanism design. In: *IEEE International Conference on E-Commerce*, pp. 347–354 (2003)
9. Cao, J., Spooner, D.P., Jarvis, S.A., Nudd, G.R.: Grid load balancing using intelligent agents. *Future Generation Computer Systems* **21**(1), 135–149 (2005)
10. Cardellini, V., Colajanni, M., Yu, P.S.: Dynamic load balancing on web server systems. *IEEE Internet Computing* **3**(3), 28–39 (1999)

11. Casavant, T.L., Kuhl, J.G.: A taxonomy of scheduling in general-purpose distributed computing systems. *IEEE Transactions on Software Engineering* **14**(2), 141–154 (1988)
12. Chang, R., Guo, M., Lin, H.: A multiple parallel download scheme with server throughput and client bandwidth considerations for data grids. *Future Generation Computer Systems* **24**(8), 798–805 (2008)
13. Chavez, A., Moukas, A., Maes, P.: Challenger: A multi-agent system for distributed resource allocation. In: *Proceedings of the First International Conference on Autonomous Agents* (1997)
14. Clearwater, S.H. (ed.): *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, Singapore (1996)
15. Cliff, D.: Explorations in evolutionary design of online auction market mechanisms. *Journal of Electronic Commerce Research and Applications* **2**(2), 162–175 (2003)
16. Cliff, D., Bruten, J.: Simple bargaining agents for decentralized market-based control. Tech. Rep. HPL-98-17, HP Laboratories, Bristol, UK (1998)
17. Cliff, D., Bruten, J.: Shop til you drop I: Market trading interactions as adaptive behavior. Tech. Rep. HPL-98-58, HP Labs, Bristol, UK (1998)
18. Eiben, A.E., Smith, J.E.: *Introduction to Evolutionary Computing*. Springer Verlag, Heidelberg (2003)
19. Esteva, M., Padget, J.: Auctions without auctioneers: Distributed auction protocols. In: *Lecture Notes in Artificial Intelligence*, vol. 1788, chap. II, pp. 20–28. Springer-Verlag, Berlin, Germany (2000)
20. Eugster, P., Guerraoui, R., Kermarrec, A., Massoulie, L.: Epidemic information dissemination in distributed systems. *Computer* **37**(5), 60–67 (2004)
21. Eymann, T.: Co-evolution of bargaining strategies in a decentralized multi-agent system. In: *AAAI Fall 2001 Symposium on Negotiation Methods for Autonomous Cooperative Systems* (2001)
22. Eymann, T., Padovan, B., Schoder, D.: The catallaxy as a new paradigm for the design of information systems. In: *Proceedings of the 16th World Computer Congress of the International Federation for Information Processing* (2000)
23. Eymann, T., Reinicke, M., Ardaiz, O., Artigas, P., de Cerio F., L.D., Freitag, Messeguer, R., Navarro, L., Royo, D., Sanjeevan, K.: Decentralized vs. centralized economic coordination of resource allocation in grids. *Lecture notes in computer science* pp. 9–16 (2004)
24. Eymann, T., Reinicke, M., Ardaiz, O., Artigas, P., Freitag, F., Navarro, L.: Self-organizing resource allocation for autonomic networks. In: *Database and Expert Systems Applications, 2003. Proceedings. 14th International Workshop on*, pp. 656–660 (2003)
25. Gjerstad, S., Dickhaut, J.: Price formation in double auctions. *Games and Economic Behavior* **22**(1), 1–29 (1998)
26. Greenwalt, A.R., Kephart, J.O.: Shopbots and pricebots. *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence* **1**, 506–511 (1999)
27. Gupta, A., Stahl, D.O., Whinston, A.B.: The economics of network management. *Communications of the ACM* **42**(9), 57–63 (1999)
28. 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*, vol. 3, pp. 1583–1587 (2005)
29. He, M., Jennings, N.R., Leung, H.F.: On agent-mediated electronic commerce. *IEEE Transactions on Knowledge and Data Engineering* **15**(4), 985–1003 (2003)
30. Jagannathan, S., Almeroth, K.C.: Price issues in delivering e-content on-demand. *ACM SIGecom Exchanges* **3**(2), 18–27 (2002)
31. Ketcham, J., Smith, V.L., Williams, A.W.: A comparison of posted-offer and double-auction pricing institutions. *The Review of Economic Studies* **51**(4), 595–614 (1984)
32. Khan, S.U., Ahmad, I.: Non-cooperative, semi-cooperative, and cooperative games-based grid resource allocation. In: *20th International Parallel and Distributed Processing Symposium, 2006*, p. 10 (2006)
33. Kikuchi, H., Hotta, S., Abe, K., Nakanishi, S.: Distributed auction servers resolving winner and winning bid without revealing privacy of bids. In: *Proceedings of the Seventh International Conference on Parallel and Distributed Systems: Workshops*, p. 307. IEEE Computer Society, Washington, DC, USA (2000)
34. 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: S.H. Clearwater (ed.) *Market-Based Control: A Paradigm for Distributed Resource Allocation*, pp. 53–73. World Scientific, Singapore (1996)
35. Lai, K.: Markets are dead, long live markets. *ACM SIGecom Exchanges* **5**(4), 1–10 (2005)
36. Lewis, P.R., Marrow, P., Yao, X.: Evolutionary market agents for resource allocation in decentralised systems. In: *Parallel Problem Solving From Nature - PPSN X, Lecture Notes in Computer Science*, vol. 5199, pp. 1071–1080. Springer (2008)
37. Lewis, P.R., Marrow, P., Yao, X.: Evolutionary market agents and heterogeneous service providers: Achieving desired resource allocations. In: *Proceedings of the IEEE Congress on Evolutionary Computation*, pp. 904–910 (2009)

38. Mas-Colell, A., Whinston, M.D., Green, J.R.: *Micro-Economic Theory*. Oxford University Press, Oxford (1995)
39. Niu, J., Cai, K., Parsons, S., Gerding, E., McBurney, P., Moyaux, T., Phelps, S., Shield, D.: JCAT: a platform for the TAC market design competition. In: *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: Demo papers*, pp. 1649–1650. International Foundation for Autonomous Agents and Multiagent Systems (2008)
40. Papazoglou, M.P., Traverso, P., Dustdar, S., Leymann, F.: *Service-oriented computing: Research roadmap* (2006)
41. Phelps, S., McBurney, P., Parsons, S., Sklar, E.: Co-evolutionary auction mechanism design: A preliminary report. *Lecture Notes in Computer Science* pp. 123–142 (2002)
42. Phelps, S., Parsons, S.D., McBurney, P.: An evolutionary game-theoretic comparison of two double-auction market designs. In: *Proceedings of the 6th Workshop on Agent Mediated Electronic Commerce*. New York, NY, USA (2004)
43. Plott, C.R., V. L. Smith, V.: An experimental examination of two exchange institutions. *The Review of economic studies* pp. 133–153 (1978)
44. Price, T.C.: Using co-evolutionary programming to simulate strategic behaviour in markets. *Journal of Evolutionary Economics* **7**(3), 219–254 (1997)
45. Ramasubramanian, V., Sirer, E.G.: Perils of transitive trust in the domain name system. Tech. Rep. TR2005-1994, Cornell University, Ithaca, New York, USA (2005)
46. Ratnasamy, S., Francis, P., Handley, M., Karp, R., Schenker, S.: A scalable content-addressable network. In: *Proceedings of the 2001 SIGCOMM conference*, vol. 31, pp. 161–172 (2001)
47. Ripeanu, M., Foster, I.: Peer-to-peer architecture case study: Gnutella network. In: *Proceedings of International Conference on Peer-to-peer Computing*, vol. 101 (2001)
48. Rodriguez, P., Biersack, E.W.: Dynamic parallel access to replicated content in the internet. *IEEE/ACM Transactions on Networking* **10**(4), 455–465 (2002)
49. Roth, A.E., Erev, I.: Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and economic behavior* **8**(1), 164–212 (1995)
50. Singh, M.P., Huhns, M.N.: *Service-Oriented Computing: Semantics, Processes, Agents*. John Wiley and Sons, Chichester, West Sussex (2005)
51. Son, Y.S., Baldick, R.: Hybrid coevolutionary programming for nash equilibrium search in games with local optima. *IEEE Transactions on Evolutionary Computation* **8**(4), 305–315 (2004)
52. Surana, S., Godfrey, B., Lakshminarayanan, K., Karp, R., Stoica, I.: Load balancing in dynamic structured peer-to-peer systems. *Performance Evaluation* **63**(3), 217–240 (2006)
53. Tesfatsion, L., Judd, K.: *Handbook of computational economics*. Elsevier (2006)
54. Waldspurger, C., Hogg, T., Huberman, B., Kephart, J., Stornetta, W.: Spawn: A distributed computational economy. *IEEE Transactions on Software Engineering* **18**(2), 103–117 (1992)
55. Walia, V., Byde, A., Cliff, D.: Evolving market design in zero-intelligence trader markets. In: *Proceedings of the IEEE International Conference on E-Commerce*, pp. 157–163. IEEE Computer Society (2003)
56. Wolski, R., Plank, J., Brevik, J., Bryan, T.: Analyzing market-based resource allocation strategies for the computational grid. *International Journal of High Performance Computing Applications* **15**(3), 258 (2001)
57. Xiao, L., Zhu, Y., Ni, L.M., Xu, Z.: GridIS: An incentive-based Grid Scheduling. In: *Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium*, p. 65b (2005)