# The Effects of Truthfulness on a Computing Resource Options Market

Owen Rogers and Dave Cliff
csorr@bristol.ac.uk csdtc@bristol.ac.uk
Department of Computer Science
University of Bristol
Bristol, UK

***Abstract*—Grid, cluster and cloud computing provide the opportunity for resources to be traded as a commodity in a open marketplace. An options market for computing resources would allow users to reserve a resource for a fee, and then pay an additional fee later should they actually need to use it. However, a major issue is ensuring that users do not falsify their likely requirements with the objective of reducing costs while keeping their right to use the resource. This paper describes a computer simulation of a two-period model that was proposed by Wu, Zhang and Huberman in 2008, which they claimed promoted truth-telling among the resource buyers population. Wu et al. provided a theoretical description and analysis of their model, but presented no empirical results. Our work, reported in this paper, allows for detailed exploration of less abstract, and more heterogeneous, instantiations of their model. We explore the Wu et al. model via techniques similar to replicator dynamics used in studies of evolutionary processes. A discussion on our empirical findings, and how they can be utilised in a commercial offering of the model, is included.**

***Index Terms*—market-orientated computing, resource reservation, risk assessment, utility computing**

## I. Introduction

Grid, cluster and, most recently, cloud computing have all promised to transform computing resources into a commodity, that can be delivered in a manner similar to that of existing utilities, such as electricity, gas, water and telephone services [2]. Cloud computing in particular is primed to deliver a new level of freedom to the consumer, allowing different levels of service and quality to be delivered on an as-needed basis without the need for capital investment.

This *utility based model* provides users with the ability to purchase computing resources as if they were any other commodity such as coal or steel. By providing a suitable mechanism for buying and selling, this *market orientated* computing opens up a wide range of trading possibilities - CPU cycles, storage capacity, or memory allocations can be bought and sold, for current or future use. This is already happening to some extent in the market place, and a wide range of economic and resource sharing models for grids, clusters and clouds are publicly accessible [10], [6].

However, the variable nature of IT usage means that pricing the service so that competitiveness and profitability are balanced has an element of risk. For the enterprise, determining and hedging their future demand for a resource is not an easy task [8].

Currently, users purchase capability from the provider directly - the use of a centralised marketplace or an intermediary party is still an area of research opposed to a reality.

Such a centralised mechanism would enable a true Service Orientated Architecture where customer needs are matched to the most suitable computing resources using brokers or coordinators. This would be controlled through Service Level Agreements (SLA) which would define the metrics that must be achieved (e.g. uptime, latency) and the compensation that would be due to the customer should the metric not be achieved [3], [5].

To take account of future requirements for resource, users could reserve resources through a derivatives market involving futures and/or options. A futures contract is an agreement to buy or sell an asset for a certain price at a certain time in the future. An options contract gives the contract holder the right to buy, or sell, an asset by a certain date for a certain price, without obligation [7].

It has been proposed that *swing options*, originally developed for trading electrical power, can be used to price a future reservation of computing resources [4]. Analogous to electricity, computing resources are non-storable with volatile usage patterns, so such a model would provide customers with flexibility in terms of amount and duration of resource requirement, and resource providers with an estimate of demand.

This presents two problems. Firstly, how can users accurately predict their future resource requirement. Secondly, how can the user be trusted to submit a true representation of their likely resource requirements.

The first issue can be solved using a forecasting tool, such as that proposed in [4]. For the second issue, Wu et al. [9] proposed a reservation model which was shown to lead to a truthful reservation on the user's part.

This paper will focus on simulating the reservation model proposed by Wu et al, in a multiple user, heterogeneous, variable market. Wu et al's model involved a number of users who require the resource, plus a central authority ("the coordinator") responsible for receiving and resolving resource requests. The first section of our paper concentrates on determining if the model favours those who are honest. The second section focuses on determining if the model encourages its users to become more honest. The third section determines if a more honest user-base is beneficial to the coordinator. Finally, a discussion on real world applications of the model and simulation is included.

## II. Previous Work

Wu et al. proposed a two-period model for resource reservation in which at the first period the user knows her probability of using the resource at the second period, and purchases a reservation whose price depends on that probability.

Consider N users who live for two discrete periods. Each user can purchase a unit of resource to use in the second period, either at a discounted rate of 1 in Period 1, or at higher price C, where C > 1, in Period 2. In Period 1, each user only knows the probability that they will need the resource in Period 2 - it is not known for certain until the next period.

A third agent, the coordinator, is introduced who makes a profit by aggregating the user's probabilities and absorbing risk through a two period game described below:

1) Period 1: Each user $i$ submits to the coordinator a probability, $q_i$, which does not have to be the real probability, $p_i$, that they will require a unit of resource in Period 2.
2) Period 1: The coordinator reserves $\sum q_i n_i$ units of resource from the resource provider at the discount price for use in Period 2, where $n_i$ is the number of units of resource required by each user. For simplicity in this simulation, $n_i = 1$ for all users.
3) Period 2: The coordinator delivers the reserved resources to users who claim them. If the amount reserved by the coordinator is not enough to cover the demand, more is purchased from the resource provider at the higher unit price C.
4) Period 2: User *i* pays

$$\begin{cases} f(q_i) & \text{if resource is required} \\ g(q_i) & \text{if resource is not required} \end{cases}$$

where $f, g : [0, 1] \to \mathcal{R}^+$

The contract can be regarded as an option if $g(q_i)$ is paid in Period 1 (i.e. as a premium), and $f(q_i) - g(q_i)$ (i.e. as a price) is paid in Period 2 should the resource be required. In Period 1, the resource is reserved, but the user is not under any obligation to purchase.

Wu et al. showed that if the following conditions could be met, the coordinator would make a profit:

- Condition A: The coordinator can make a profit by providing the service.
- Condition B: Each user prefers to use the service provided by the coordinator, rather than to deal with the resource provider.

The following truth telling conditions are not completely necessary, but are useful, for conditions A and B to hold:

- Condition T1 (truth telling): Each user submits his true probability in Period 1 so that he expects to pay the least later.
- Condition T2 (truth telling): When a user does not need a resource in Period 2, it is reported to the coordinator.

The following specific case was proved to be valid, where k is a constant chosen to alter the price paid by the customer:

$$f(q_i) = 1 + \frac{k}{2} - kp_i + \frac{kp_i^2}{2}$$

$$g(q_i) = \frac{kp_i^2}{2}$$

This paper describes a computer simulation of the model, using the specific case analysed in [9] where C = 2, and k = 1.5, to see if these conditions are met in a real-world, multi-user scenario, where users submit different resource probabilities using different levels of honesty. It is this heterogeneity of probabilities and levels of honesty that makes our simulation an extension of the theoretical model and supporting analysis provided by Wu et al. in 2008.

## III. Methodology

A computer simulation was implemented in Python to replicate the model as an options contract[1]. The algorithm performs the following steps:

1) Each user *i* in the range 1 to N is assigned an "honesty", $H_i$, chosen randomly from a uniform distribution over [0,1] which describes the accuracy with which a probability of future resource requirement is provided to the coordinator. An honesty of 1 means a user will always exercise their right to purchase as per their forecast probability. A honesty of 0 means a user will never exercise their right to purchase.
2) A *replicator dynamics* approach is adopted, whereby for every two units of total time T, a user is randomly chosen to undergo a *mutation*, and this user is given a new honesty [1]. The user tries bidding and executing as per the following steps for a sample size S, using the new honesty.
   a) For each user, a random resource probability, $p_i$ is assigned. A probability of 0 means that a resource will definitely not be required in the next time period. A resource probability of 1 means a resource will definitely be required in the next time period
   b) Each user is given the opportunity to request a resource to be utilised in the next time period. The user submits a resource probability to the coordinator.
   c) If the user is underestimating its future requirement, it will submit the following resource probability to the coordinator:
   $$q_i = H_i p_i$$
   d) If the user is overestimating its future requirement, it will submit the following resource probability to the coordinator:
   $$q_i = (1 - H_i)(1 - p_i) + p_i$$
   e) The user is charged a premium $\dfrac{kq_i^2}{2}$ to reserve the resource. The premium and a fixed transaction

[1]Copies of the Python source code are available on request.

fee are removed from the user's bank balance, and added to the coordinator's balance.

f) The coordinator purchases units from the resource provider at a cost of 1 per unit. This is removed from the coordinator's balance and added to the resource provider's balance. As an example, a user *i* with $H_i = 1$ who anticipates the future requirement with a probability of 0.8, will submit a probability of 0.8 to the coordinator as $q_i = H_i p_i = 1 \times 0.8 = 0.8$. Therefore, this user is considered "honest" in the language of Wu et al.

g) A user *i* with $H_i < 1$ will not always exercise their right as per their submitted probability. For instance, a user with honesty 0.7 who anticipates the future resource requirement with a probability of 0.8, will submit a probability of future resource requirement with $q_i = H_i p_i = 0.7 \times 0.8 = 0.56$ when underestimating. Therefore, this user is classed as "dishonest".

h) Each user is now given the option of exercising their right to use their resource.

i) A user will exercise their right where $p_i \geq A$ where A is the availability of the resource. An availability of 0 means that there is no surplus of the resource, and all users will exercise their right. An availability of 1 means that that the resource is abundant, and no users will exercise their right. For the simulation, assume the availability varies with uniform randomness.

j) If a user wishes to use the resource, they are charged a price, $1 + \frac{k}{2} - kq_i$ which is removed from their balance and added to the coordinator's balance.

k) If the coordinator has not previously purchased enough resource from the resource provider, they will purchase the deficit at a cost of C per unit.

l) This cost is removed from the coordinator's balance and added to the resource providers.

3) Steps a-l are repeated S times, to ensure each user provides a range of probabilities to the coordinator using the same honesty.
4) If the user finds that the mean cost per required resource over the sample S is lower using the mutated honesty, the honesty remains and the new behaviour is adopted by the user. If not, the honesty returns to its previous value and the old behaviour continues.
5) A new mutation is determined as per Step 2, and this process continues until T iterations have passed

## IV. Does Telling the Truth Benefit the User?

### A. Setup

The aim of our first simulation reported here is to determine if a user who submits an untrue probability to the coordinator benefits more than a user who submits a true probability. The most likely scenario is where a user submits a probability of less than they predict they will require, such that they pay a lower premium but still have the right to purchase a resource should they require it. The situation where a user submits a higher probability than they required is also simulated, thereby providing a complete view as to whether the the system can be played. The simulation was setup with the parameters shown in Table 1.

This simulation was constructed such that no mutations would occur, instead, a wide range of fixed honesties submitting a range of probabilities were simulated.

| Parameter | Description | Value |
|---|---|---|
| C | Cost to Coordinator per Unit Resource | 2 |
| k | Cost Multiplier to User | 1.5 |
| F | Fixed Transaction Fee | 0.01 |
| N | Number of Bidders | 1000 |
| S | Number of Samples | 1000 |
| T | Number of Mutation Repetitions | 1 |

Table I
Parameters for Comparison Simulation

### B. Results

The simulation was run, and the honesty of the user plotted against their mean cost per resource. This is shown in Figure 1.

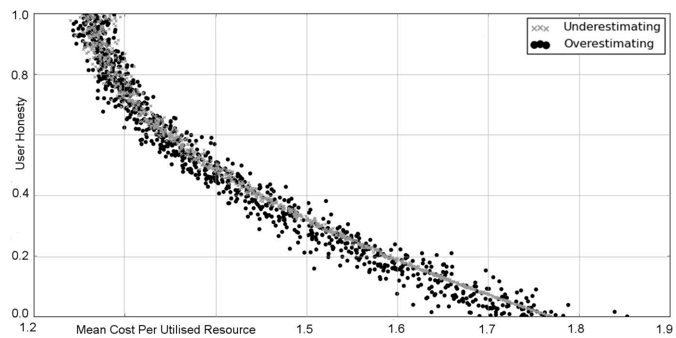


Figure 1. Plot of user honesty against mean cost per utilised resource

### C. Discussion

As can be seen from Figure 1, users with a higher honesty generally pay a smaller price per unit resource than those who lie. This value depends on the exact probabilities submitted, $p_i$, which varies due to the randomness in the simulation.

Over infinite samples, all values of $p_i$ between 0 and 1 will be simulated. When $H_i = 1$, $q_i = p_i$, so the mean cost per resource is the mean premium and strike price of all of the values of $p_i$ between 0 and 1. Thus,

$$\int_0^1 1 + \frac{k}{2} - kp + \frac{kp^2}{2}\mathrm{dp} = \left[p + \frac{kp}{2} - \frac{kp^2}{2} + \frac{kp^3}{6}\right]_0^1 = 1.25$$

When $H_i = 0$, $q_i = 0$. Therefore, the mean cost per utilised resource is:

$$1 + k/2 = 1.75$$

The determined mean cost per utilised resource of 1.25 when $H_i = 1$ and 1.75 when $H_i = 0$ match that obtained through simulation.

It is interesting to note that the mean cost per utilised resource when using the coordinator's service is always cheaper than when purchased directly from the resource provider - the cost direct from the resource provider is 2. Therefore, in this case, Condition B is proved as the user has an incentive to use the coordinator's service as it provides a mean saving compared to using the resource provider's service directly.

## V. Does the model cause a user to change its behaviour?

### A. Setup

The aim of our second simulation is to determine if the model encourages a user to change their honesty.

For this, only the situation where a user underestimates their usage is considered, as this is most likely to occur in a real-life scenario (as a user may decide to provide an underestimate of their probability, to obtain a lower premium but ensure the right to purchase a resource).

The parameters for the simulation are shown in Table 2.

The simulation alternates between mutating $H_i$ such that user *i* becomes more honest, $H_{mutation >} H_i$, and mutating $H_i$ such that user *i* becomes less honest, $H_{mutation <} H_i$. This simulates a real world scenario whereby a user may spontaneously decide to become more or less honest to see if they benefit financially as a result. As a mutation only occurs every two iterations, a sample size of 100 has been chosen so that it will be possible to see if there is a change in the average honesty of the user-base during a realistic computational time frame. Initially, all users have an honesty of 0.

| Parameter | Description | Value |
|---|---|---|
| C | Cost to Coordinator per Unit Resource | 2 |
| k | Cost Multiplier to User | 1.5 |
| F | Fixed Transaction Fee | 0.01 |
| N | Number of Bidders | 100 |
| S | Number of Samples | 100 |
| T | Number of Mutation Repetitions | 2000 |

Table II
PARAMETERS FOR MUTATION SIMULATION.

### B. Results

The simulation was run, and the status of each user's honesty per iteration was recorded. This was plotted as a chart showing the honesty of each user during the duration of the simulation in Figure 2, and as a mean honesty in Figure 3.

### C. Discussion

Figures 2 and 3 show that user base as a whole becomes more honest as more mutations occur. Initially, the mean honesty is 0 as all users are set to have a honesty of 0. As the simulation continues, mutations in the favour of honesty are adopted by the users, and the mean honesty of the user base

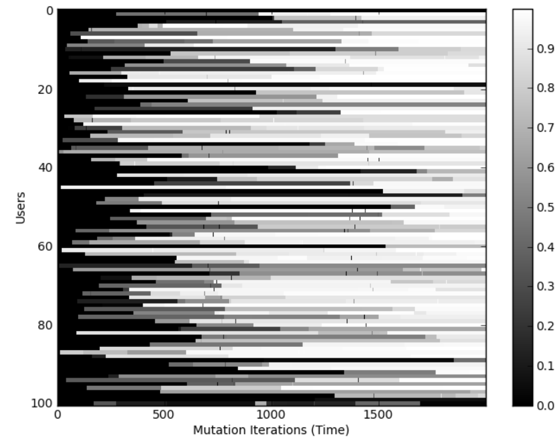


Figure 2. Plot showing the state of all users honesty at each iteration where the initial honesty of all users is 0, with increasing honesty being shown by increasing brightness as per the colour bar.

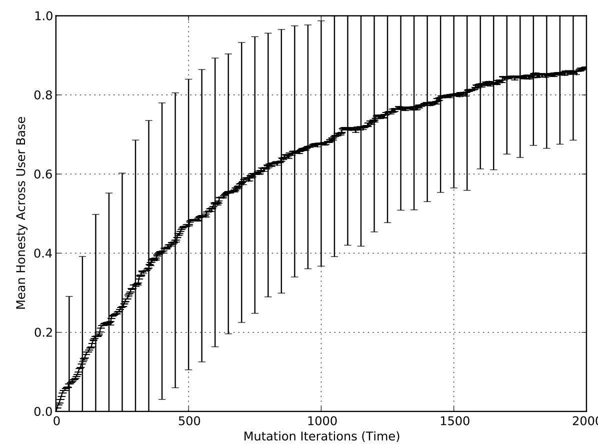


Figure 3. Plot showing changes in the honesty of the user-base at each iteration, where the initial honesty of all users is 0. Error bars show +/- one standard deviation.

increases as shown in Figure 3. Figure 2 shows how increases in honesty are frequently adopted by the user (shown as an increase in brightness), while decreases in honesty frequently result in the old behaviour being continued (shown as very small dark patches).

The data in Figure 3 bears some further analysis. While it is clear that the mean level of honesty in the population is increasing over time, it is not clear what the ultimate long-run trend is: does Figure 3 show a steady progression to a state where all users attain 100% honesty, or is it instead an indication of an asymptotic approach, where 100% honesty is never reached because once honesty is sufficiently high across the general user population, that population can then tolerate or support a small number of dishonest "free-riders"? To explore this question, we ran another simulation experiment where the initial user-base were all assigned an honesty of 1, and then subjected to mutation as per our previous experiment. The results are shown in Figure 4.

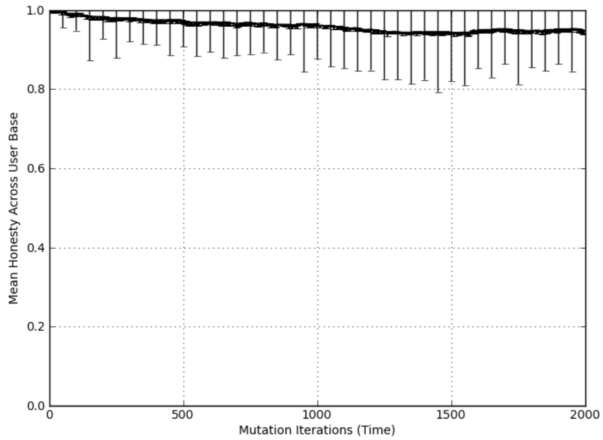


Figure 4. Plot showing changes in the honesty of the user-base at each iteration, where the initial honesty of all users is 1. Error bars show +/- one standard deviation.

As can be seen from Figure 4, the level of honesty in the population of initially-honest users remains stable near 100%, with a small deviation being caused by a mutation occurring in a user at each iteration, and so we conclude that the user population in Figure 3 is converging toward a stable state of 100% honesty.

As users are compelled to submit their true probabilities, hence Condition T1 is proved.

# VI. Do changes in user behaviour benefit the coordinator?

## A. Setup

The simulation was setup to record the coordinator's profit for each iteration, as the mean honesty of the user-base increases, as per the parameters in Table 2.

## B. Results

Figure 5 shows a plot of mean honesty of the user-base against the coordinator's profit.

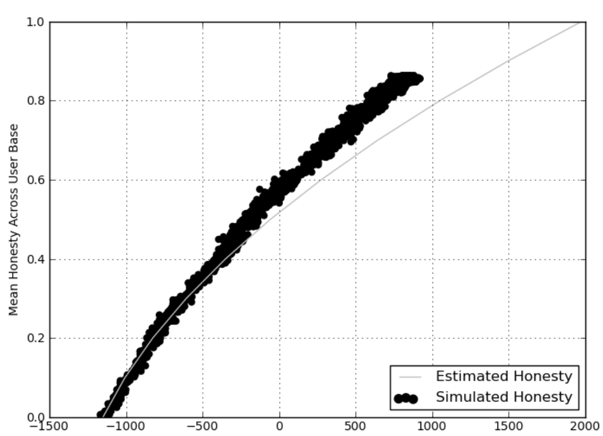


Figure 5. Plot of mean honesty of user-base against coordinator's profit.

## C. Discussion

From the methodology described earlier, it can be seen that the coordinator's profit per user per sample is:

*Profit (Period 1) = (Premium + Fee - Discount Resource Fee)* for all users who make a reservation

*Profit (Period 2) = (Price - Standard Resource Fee)* for only those users who exercise their option

Substituting equations for the premium, price, discount resource fee and standard resource fee, and considering the mean values of $p_i$ and $q_i$ across the user base across *S* samples and *N* users, we can derive the profit, or earnings, $E$:

$$E = NS\left(\left[\frac{k\bar{q}^2}{2} + F - \bar{q}\right] + \left[(1 + \frac{k}{2} - k\bar{q})\bar{p} - C(\bar{p} - \bar{q})\right]\right)$$

Figure 5 shows a plot of this estimated profit, where $\bar{q} = \bar{p}\bar{H}$, and a plot of the coordinators profit against honesty from the simulation.

It can be seen that the estimated honesty closely matches the results generated from the simulation. There is a discrepancy however at higher mean values of $H_i$, which are most likely due to noise in the simulation, specifically the use of pseudo-random number generators for $H_i$, $p_i$ and $A$, and lost resolution due to averaging.

Figure 5 shows that as the honesty of the user-base increases, so too does the coordinator's profit.

Therefore, Condition A is proved - the coordinator can make a profit from the model.

# VII. Further work

## A. Market Modelling

The simulator provides an opportunity to measure the performance of the model under a variety of market conditions.

In the simulations discussed here, the market has varied uniformly over the duration so that option execution is performed as per the user's actual probability, $p_i$. The model can be extended such that the likelihood of exercising an option changes with changes in resource availability and therefore demand. A market profile would define the volatility of the availability of the resource. Commercially useful market profiles would include a steady market with low availability, a steady market with high availability, a highly volatile market, and a steady market with an infrequent crash.

In the simulations discussed here, the number of resource units, $n_i = 1$. It would be interesting to simulate the situation where $n_i > 1$ as, in reality, it is likely a user would require a larger number of resource units than 1. Some users may have a typical resource requirement from period to period, for example, a managed hosting provider may have a good view of the maximum number of resources required for a period. Other users may change their resource requirement regularly, for example, a provider of complex IT will vary sporadically depending on likely new implementations during a period.

### B. Risk Offsetting

The coordinator can protect her risk by anticipating user demand, and charging an appropriate transaction fee to offset this risk. Figure 6 shows how profit is affected by mean honesty for a small sample of transaction fees.

Further research may establish that a particular customer base is more likely to be dishonest. For example, one geographical region may be less likely to be honest to a coordinator in a different geographical region due to previous existing social, economic, political or cultural issues, which causes an inherent lack of trust.

The term “honesty” can be here reasonably interchanged with *reliability*. It may be that a particular customer base has the best intentions, but regularly reserve resources with an incorrect probability. For example, a user who deals with implementing complex systems may find it more difficult to predict future usage accurately due to the longer sales, implementation and acceptance cycles brought about by determining complicated design requirements. On the other hand, a particular customer base may have a very accurate view of future requirements, such as a website that has a fixed number of users.

If market research is used to determine the mean honesty of a particular demographic, this information can be used to provide the most competitive price to that specific segment without exposing the coordinator to unnecessary risk. In Figure 6, for example, if it is found that 40% of users are dishonest in a particular market segment, a transaction fee of 0.05 would mean that the coordinator would generally break even. If the coordinator is very risk averse, they could set the transaction fee to 0.14 which should prevent any loss.

Further work should be undertaken to determine if such inherent honesties/reliabilities exist in the addressable market, and to determine a transaction fee for each market segment such that risk and competitive pricing are balanced. Combined with market modelling previously discussed, this segment specific pricing could be a powerful differentiator.

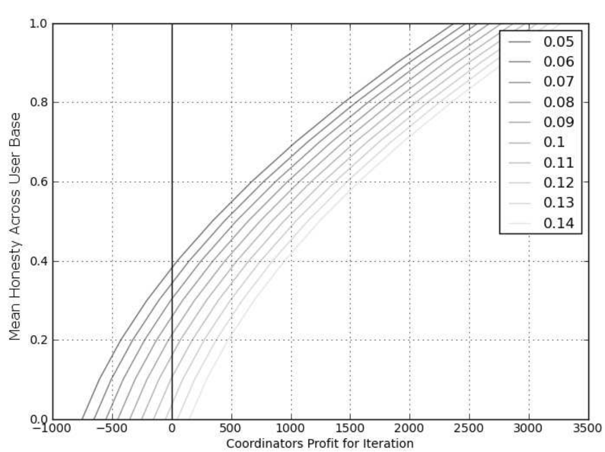


Figure 6. Plot showing relationship between honesty of user-base and coordinator’s profit for different transaction fees.

## VIII. Conclusion

This paper has shown how a truth telling reservation model for computing resources described by Wu et al. can provide the basis for an options market, which is both price-competitive for the user and profitable for the coordinator. The model was simulated with multiple heterogeneous users, submitting a wide range of probabilities over a long term. It was found to benefit users who are honest with a reduced cost, and to encourage those who are dishonest to become more truthful. Furthermore, the results show that the coordinator benefits financially as the honesty of the user-base increases.

The simulation has shown that the reservation model may be suitable for real-world application. The model provides a platform for further risk assessment work to be undertaken and, as discussed, the simulator can be further extended to simulate a variety of market conditions, or specific user demands.

It has been shown that an appropriate transaction fee can offset risk caused by different levels of honesty, or reliability, in a specific market segment. This inherent honesty could be determined through market research, and integrated into a future simulation so that a transaction fee can be determined under different market profiles, and user demands.

By taking the results from this paper and extending them with future research into the performance of the model under different conditions and inherent honesties, in different segments, a commercial offering that is profitable to the coordinator, beneficial to the user, and with a calculated level of risk looks likely to be achievable.

## IX. Acknowledgements

We thank the UK EPSRC for funding the Large-Scale Complex IT Systems Initiative (www.lscits.org) as well as HP Labs Adaptive Infrastructure Lab for providing additional financial support.

## References


[1] T Borgers. Learning Through Reinforcement and Replicator Dynamics. *Journal of Economic Theory*, 77(1):1–14, November 1997.

[2] Rajkumar Buyya, Chee Shin Yeo, Srikumar Venugopal, James Broberg, and Ivona Brandic. Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Computer Systems*, 25(6):599–616, June 2009.

[3] Andrew Byde, Mathias Sallé, Claudio Bartolini, and Mathias Sall. Market-Based Resource Allocation for Utility Data Centers Market-Based Resource Allocation for Utility Data Centers. *Work*, 2003.

[4] Scott H Clearwater and Bernando Huberman. Swing Options : A Mechanism for Pricing IT Peak Demand. *Proceedings of 11th International Conference on Computing in Economics*, 2005.

[5] L Grit. Broker Architectures for Service-oriented Systems. 2005.

[6] David Hilley. Cloud Computing : A Taxonomy of Platform and Infrastructure-level Offerings Cloud Computing : A Taxonomy of Platform and Infrastructure-level Offerings. *Technology*, (April), 2009.

[7] John C Hull. *Fundamentals of Futures and Options Markets*. 2005.

[8] A Khajeh-Hosseini, I Sommerville, and I Sriram. Research Challenges for Enterprise Cloud Computing. *Arxiv preprint arXiv: ...*, 2010.

[9] F Wu, L Zhang, and BA Huberman. Truth-telling reservations. *Proceedings of 11th International Conference on Computing in Economics*, 2008.

[10] Chee Shin Yeo and Rajkumar Buyya. A taxonomy of market-based resource management systems for utility-driven cluster. *Cluster Computing*, (June):1381–1419, 2006.