Hedging Cloud Energy Costs via Risk-free Provision Point Contracts

Owen Rogers and Dave Cliff
Department of Computer Science, University of Bristol, Bristol, U.K.

Keywords: Pricing, Derivatives, Energy Market, Electricity Futures.

Abstract: The cost of electricity is a major concern to public providers of cloud computing services. On-demand pricing, common amongst cloud providers, does not aid the provider in planning future demand and therefore purchasing energy at discounted rates. In this paper, we describe a number of advance pricing schemes for cloud computing resources based on *provision-point* contracts, commonly used by deal-of-the-day websites such as Groupon. We propose three models – Group Provision Points, Contributory Provision Points, and Variable Reward Forwards – that each reward consumers with reduced prices for advance reservations, while allowing providers to make accurate forecasts of energy usage. Furthermore, we show how the schemes are risk-free for the provider, guaranteeing to be at least as profitable as on-demand schemes. We present results from a simulation of the schemes, and compare the results to our analytically derived predictions.

1 INTRODUCTION

Consumers of cloud computing resources typically pay a single price to access a virtual machine for a specified period of time. This single price covers the virtual machine's fraction of the cost of the physical server itself, maintenance and repairs, the physical datacentre space, the electricity needed to power it, and the cost of air conditioning to cool the datacentre.

In on-demand pricing, consumers gain access to the resource immediately and are charged for the amount of time they use the resource.

In forward pricing, consumers gain access to the resource at a specified time in the future, and have access for a pre-agreed duration.

Air conditioning and datacentre space are generally fixed costs. Regardless of how many servers are placed in the datacentre, these costs will essentially be the same.

Electricity costs for powering servers are variable costs. The total electricity required by the provider is proportional to the amount of virtual machines demanded by the provider's customers.

Energy costs are a significant cost for providers of public cloud computing resources.

Estimates for the contribution of server electricity to the total cost of ownership (TCO) of a physical server vary between 3% and 15% (Barroso and Hölzle 2009; Berl et al., 2009). Volume servers account for 34% of datacentre electricity usage (Brown, 2008). A full review of datacentre costs can be found in (Patel and Shah, 2005).

This cost therefore impacts the price paid by consumers to access virtual machines, and the profit achieved by the provider. In a competitive marketplace, keeping prices as low as possible is critical for commercial success.

Currently, research is being focussed on reducing the power consumption of computing technology (Barroso and Holzle 2007; Lee and Zomaya, 2010). The primary focus of this research is reducing carbon footprint, but reducing expenditure is an important factor too.

Typically, a cloud provider would purchase electricity on-demand for a fixed price to power its datacentre. This could be directly from an energy supplier, or from a broker who hedges market-traded instruments to offer fixed prices to its clients.

Larger cloud providers might purchase electricity directly from the spot-market, where prices vary over time to match supply with demand. These larger providers may also generate their own electricity and be able to contribute energy to the grid as well as consuming it through bilateral agreements.

Some research has been directed at moving virtual machines between datacentres with the aim

Rogers O. and Cliff D..
Hedging Cloud Energy Costs via Risk-free Provision Point Contracts.
DOI: 10.5220/0004374502440252
In *Proceedings of the 3rd International Conference on Cloud Computing and Services Science* (CLOSER-2013), pages 244-252
ISBN: 978-989-8565-52-5

of finding the cheapest spot-price.

Qureshi *et al.* were the first to propose dynamically assigning computational workloads in distributed systems to locations where electricity may be cheaper. They found savings of millions of dollars could be achieved through a simulation (Qureshi et al., 2009).

A similar method was suggested by Rao *et al.*, but the dynamic allocation also takes into account the latency between different locations, so that QoS metrics would be met while electricity cost reduced (Rao et al., 2010).

Buchbinder *et al.* extended these methods so that only batch applications would be migrated to cheaper markets (Buchbinder et al., 2011). In this way, applications that could tolerate a delay would use the cheapest electricity, and interactive applications would not cause poor user-experience as a result of the overhead involved in migrating the application

Ding *et al.* also proposed that virtual machines could be moved between datacentres based on electricity prices (Guo et al., 2011).

However, little research has been conducted on if providers can use *derivative contracts* to purchase electricity in advance for a discount.

The cloud provider could potentially decrease its costs by purchasing electricity futures directly (Hull, 2008). *Futures contracts* are a type of derivative that give buyers guaranteed access to the resource in advance of when it is delivered: the user is *obliged* to take ownership of the resource on the delivery date that the contract specified. The provider could then engage a broker to provide fixed-price electricity to top up its pre-bought electricity capacity

A futures contract typically details the size of the commodity being purchased. In electricity futures, the commodity is a quantity of electricity delivered for a fixed period of time, typically a month or a quarter.

However, the use of electricity futures can have significant associated risks. If the provider invests in a future which is subsequently not fully utilised by customers, then it is possible it will not cover the investment. Electricity delivered to the cloud provider cannot be stored; if it is not used as it is delivered, then it is wasted.

Considering an electricity future for one months delivery of 1MW costs over $35,000, this risk can be sizeable[1].

[1] ICE UK Base Electricity Futures, November 2012

In this paper, we propose three pricing schemes that allow the provider to purchase electricity futures with no-risk that they will subsequently fail to utilise their investment effectively. The provider is guaranteed to be at least as profitable as using a traditional on-demand pricing scheme.

Our schemes are based on *provision-point contracts* (also known as *assurance contracts*). In a provision-point mechanism, members of a group pledge to contribute to an action if a threshold of some order is met. If this threshold is met, the action is taken and the public goods are provided; otherwise no party is bound to carry out the action and money paid is refunded (Bagnolli and Lipman 1989).

Such a mechanism is used by deal-of-the-day website Groupon[2]. Users make requests for special offers by purchasing a coupon. When a threshold is reached, the deal is profitable to the provider and the offer is confirmed.

In previous work, we showed how provision-point contracts can be used to schedule virtual machines more effectively on a large-scale cloud infrastructure (Rogers and Cliff, 2012; Rogers and Cliff, 2012).

In this paper, we amend traditional provision points by changing the beneficiaries of the contract and the value of the offer to create a number of new pricing schemes.

Consumers of cloud computing resources can purchase these in advance for discount, while retaining the ability to purchase additional resources on-demand. The cloud provider subsequently uses this information to purchase electricity futures.

We show how Group Provision Points, Contributory Provision Points, and Variable Reward Forwards allow providers to make accurate forecasts of energy usage and therefore reduce their costs through the purchase of electricity.

We present results from a simulation of the schemes, and show that our schemes have benefits for both provider and consumer compared to traditional on-demand and forward pricing.

2 PRICING SCHEMES

2.1 On-demand Pricing

In standard on-demand pricing there is a period of duration N intervals, where resources are purchased and then immediately available.

[2] www.groupon.com

The provider charges customers a cost C_o to use the computing resource for an interval i. The total demand experienced for resources in time interval i is t_i. In this case, the total revenue (REV) achieved by the provider over the period is the total demand experienced at the on-demand price:

$$REV = C_o \sum_{i=0}^{N} t_i$$

The provider will be required to pay for electricity for the duration of the interval that the virtual machine is running at a cost E_o from the energy supplier or broker. The electricity required per virtual machine for the interval is β. In this case, the cost of electricity (COE) to the provider is the total demand experienced, at the cost of on-demand electricity per virtual machine:

$$COE = \beta E_o \sum_{i=0}^{N} t_i$$

Therefore, the provider's profit using an on-demand model is:

$$P_{od} = (C_o - \beta E_o) \sum_{i=0}^{N} t_i$$

2.2 Forward Contracts

Consider a pricing model for cloud computing which uses two periods, each period consisting of N time intervals.

In the first period, 'the reservation period', consumers purchase advance reservations (or forwards) at a cost C_r, which allows them to use a resource at a specific interval i during the next period. The total number of resources reserved in a time interval i is r_i.

In the second period, "the execution period", consumers gain access to their reservations at the specified time interval. Consumers may also purchase access to a resource for the duration of an interval at a cost C_o. The total demand experienced for resources in time interval i is t_i

In this case, the revenue achieved over the period is the sum of reserved resources bought at the reserve price, plus the additional resources bought on-demand at the on-demand cost:

$$REV = \sum_{i=0}^{N} [r_i C_r + (t_i - r_i) C_o]$$

As the provider has committed to deliver a number of resources through the sale of forward contracts on computing resources, she can use this information to purchase forward contracts on electricity to obtain a saving on consumption. The provider can choose to buy θ forward electricity contracts, where each contract entitles them to use I units of electricity for a period of N time intervals at a cost E_r per time interval.

The cost over the period is the cost of purchasing reserved electricity across the entire period, plus the sum of the cost of purchasing on-demand electricity required in addition to the reserved electricity.

$$COE = \theta N E_r + \left[E_o \beta \sum_{i=0}^{N} (t_i) - \theta N E_o I \right]$$

Therefore the profit obtained via hedging electricity consumption through the use of forward contracts on electricity is:

$$P_{res} = (C_r - C_o) \sum_{i=0}^{N} r_i + (C_o - E_o \beta) \sum_{i=0}^{N} t_i + N\theta(IE_o - E_r),$$

For the model to be worth implementing for the provider, it must offer a greater profit than using an on-demand model:

$$P_{res} > P_{od}$$

$$C_o < C_r + \frac{\theta N(IE_o - E_r)}{\sum_0^N r_i}$$

However, for the model to be beneficial to the user, the user must be incentivised to provide a forecast. Therefore, the cost of reserving a resource must be less than the cost of buying a resource on-demand:

$$C_o > C_r$$

So our conditions for the model to be beneficial to all parties are:

$$C_o > C_r \quad (1)$$

$$C_r > C_o - \frac{\theta N(IE_o - E_r)}{\sum_0^N r_i} \quad (2)$$

With forward pricing on computing resources, the provider might choose to fix C_o and C_r so that customers are fully aware of the pricing they will be charged. In this case condition (1) is satisfied, and consumers will use the service.

However, as condition (2) is dependent on $\sum_0^N r_i$, the provider is not aware of if the scheme will be more profitable than on-demand pricing until all users have purchased forward contracts and the provider must deliver the resource.

The provider must provide users with access to their reserved instances for smaller cost (and therefore less revenue), but may not benefit from cheaper electricity costs in all cases.

2.3 Group Provision Points (GPP)

This issue can be circumvented with a provision point contract. We now introduce an additional, intermediate phase – the 'confirmation phase':

1. **Reservation Phase:** Users request resources to be consumed in the execution phase
2. **Confirmation Phase:** If the provider finds that they will benefit as a result of the model by conditions (1) and (2) being met, they will confirm user's requests and the contracts are confirmed. If either condition is not met, all contracts are cancelled.
3. **Execution Phase:** Users gain access to their confirmed resources, and may also buy additional on-demand resources.

If the requirements of the user population are found not to produce an increase in profit, the provider cancels all contracts and no revenue is lost as a result. If the scheme is profitable, all contracts are confirmed. This is equivalent to a traditional provision-point contract used by deal-of-the-day websites such as Groupon.

2.4 Contributor Provision Points (CPP)

The forward and GPP schemes are extremes. In the forward scheme, all users who submit a reservation benefit from reduced prices, in spite of it sometimes not benefitting the provider. In the GPP scheme either all, or no, users benefit from reduced prices depending on whether an advantage is gained by the provider or not.

A compromise might be to only confirm contract requests to the consumers that contribute to the purchase of advanced electricity during the confirmation phase. This could be based on the earliest consumers who request a reservation. Customers who submitted a late reservation would have their contract cancelled, as their discount would not contribute to cheaper electricity.

The provider would typically determine how much advance electricity θ to purchase based on some function of the profile of the reserved resources over the month.

$$\theta = f([r_o \cdots r_N])$$

If the provider chooses to purchase θ forward contracts on electricity, this will provide the provider with $I\theta$ units of electricity each interval for N intervals. Therefore, the total electricity available to the provider over the period is $I\theta N$. This will support q contracts:

$$q = \frac{1}{\beta} I\theta N$$

If we confirm only q contracts, and cancel all others:

$$q = \sum_0^N r_i$$

Substituting into (2):

$$C_r + \frac{\theta N(IE_o - E_r)}{\frac{1}{\beta} I\theta N} > C_o$$

$$C_r + \frac{\beta(IE_o - E_r)}{I} > C_o \qquad (3)$$

The vulnerability of the forwards has now been removed, as the conditions for profitability no longer depend on the uncontrollable number of reservations.

As long as prior to implementing the scheme conditions (1) and (3) are met and E_r is set to be the maximum likely cost of an electricity future, the scheme will generate a profit over on-demand pricing.

This scheme also protects the provider against changes in the cost of electricity forwards. If the cost of a forward does not satisfy the following, the provider should cancel all contracts:

$$E_r < IE_o - \frac{I}{\beta}(C_o - C_r)$$

2.5 Variable Reward Forwards (VR)

In the variable reward model, consumers are given the guarantee that when purchasing a forward in the reservation period, the price payable for the forward will be the same, or less, than the cost of an on-demand resource. The exact value of C_r is not known until the execution period and is determined on a profit-sharing basis, where μ is the share desired by the provider.

$$\delta = \frac{\theta N(IE_o - E_r)}{\sum_0^N r_i}$$
$$C_{r_{min}} = (C_o - \delta)$$

Users pay the minimum C_r to be profitable, plus a share of the saving achieved:

$$C_r = C_{r_{min}} + \mu(C_o - C_{r_{min}})$$
$$C_r = C_o + \delta(\mu - 1)$$

If $C_r > C_o$ then on-demand instances are cheaper than reserved and the model will fail. In this case, no discount is to be offered and $C_r = C_o$. The condition for this is:

$$(C_o - \delta) > C_o$$
$$\delta < 0$$
$$C_r = \begin{cases} C_o, & \delta < 0 \\ C_o + \delta(\mu - 1), & \delta \geq 0 \end{cases}$$

This will always be as least as profitable as on-demand instances as users pay the on-demand price if no saving can be made.

3 SIMULATION

3.1 Setup

A simulation was written in Python, the primary aims being to verify that the models outperform conventional on-demand and forward pricing schemes when applied to practical applications, and that users can make savings using a rational approach to forecasting. Furthermore, a simulation will aid comparing models where the cost of electricity futures varies over time.

In our first simulation, we wish to determine which contract model generates most profit in a monopoly market where the broker is the only (or at least the preferred) provider of cloud resources. Our objective is to understand the profitability implications for the provider of such schemes, and the cost implications for the consumer.

For electricity, we assume that the broker may purchase electricity futures for a period of a calendar month, which supplies 1MWh of electricity per hour. We obtain prices of ICE UK Base Electricity Futures over a 39 month period from March 2012 (Figure 1). The cost of electricity on-demand from the grid is £0.01/kWh, based on (Barroso and Hölzle 2009) which is the most reliable source of this information in current academic research.

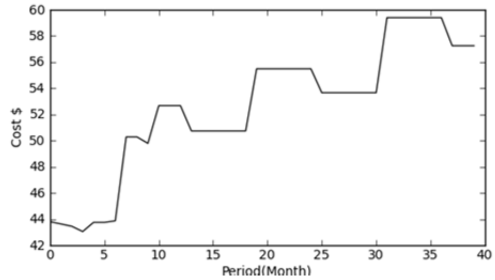

Figure 1: Cost over electricity futures over time.

Group provision points (GPP), contributory provision points (CPP) and standard forwards all have the same prices for reservations and on-demand instances. An on-demand instance is set at £0.01/Computational-Unit/hour, which is a reasonable figure in the current market.

The contributory provision point provider believes that the maximum she will need to pay for electricity in the foreseeable future is £55, Therefore,

$$C_r > C_o - \frac{\beta(IE_o - E_r)}{I}$$
$$0.007863 < C_r < C_o$$

The CPP provider sets C_r = £0.007864. This will guarantee her a benefit over on-demand pricing as long as electricity does not go higher than £55.

The forward and GPP providers set C_r to be the same. The variable reward forward (VR) provider sets $C_o = 0.01$, but has no basis for determining a reserved price. She decides that she requires 50% of any saving used by the scheme to be retained as profit, and the other 50% to be split to consumers who reserved.

We simulated a demand curve varying over time using a combination of 5 types of users:

- *Flat profile* represents where demand is constant, and hence trivially easy to predict;
- *Random profile* represents stochastically unpredictable demand, chosen randomly from a normal distribution;
- *Sine profiles* (with period of 24 hours) are an approximation to daily rhythms, where demand varies sinusoidally, peaking in the middle of the day and at a minimum in the middle of the night. More precisely, in our simulations this sinusoidal demand pattern peaks around mid-day, and demand can never be negative, so a function of the form $1+\cos(2\pi h/24)$ is used, where h is the hour-number in the day. We have explored three variations of these sinusoid patterns:

a. *Flat Sine* represents constant a constant baseline of demand with periodic variations across each day;
b. *Growing Sine* represents daily periodic demand, with the baseline increasing steadily across the month;
c. *Shrinking Sine* represents daily periodic demand, shrinking through the month.

We create a demand curve by combining different quantities of these users such that demand is generally growing over time so that the benefit of purchasing additional electricity futures can be seen in our results. There are 2000 users in total using the simulation.

The aim of the demand curve is to determine if the scheme can be profitable in a heterogeneous market of different users with different demands, which follows an increasing trend. We are not aware of any real-world data on public cloud demand which we could use over such timescales, so this is a suitable approximation in this preliminary study.

We assume the provider has servers that can support 8 virtual machines, and each server uses 380W.

3.2 Results

3.2.1 Provider Cost Reduction

For clarity, figures 3-8 show data points averaged over the last 2 months with the corresponding standard deviation shown in error bars.

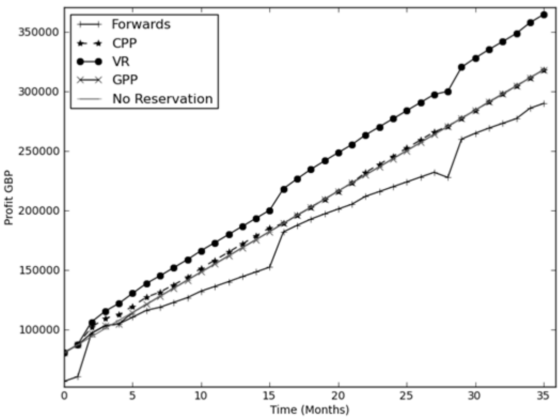

Figure 2: Profit achieved using pricing models.

Figure 2 shows that the CPP model does generate more profit for the provider than on-demand pricing alone. Initially this is around a 10% increase, but this decreases as forward electricity prices increases. In month 17 (hour 11424), electricity futures rise above the point where our reserved pricing is profitable, and so all contracts are cancelled. This can be seen in Figure 3 where no energy resources are reserved as the cost of electricity forwards goes higher than our reserved pricing threshold.

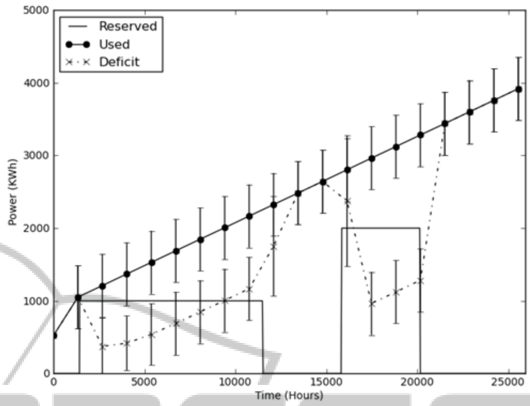

Figure 3: Power used by CPP model.

In this case, the profit achieved is the same as on-demand pricing. In Figure 4 it can be seen at this point that no contracts are confirmed, and all users must purchase on-demand resources.

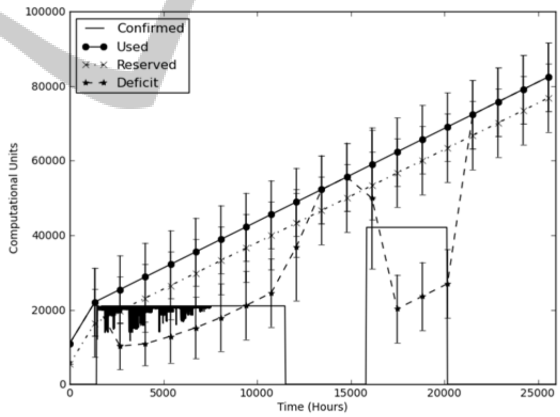

Figure 4: Resource allocation in CPP model.

The VR model is the most profitable for the provider, being up to 16% more profitable than on-demand pricing. This is because the provider is not committed to giving a specific discount to the consumer. The fact that any benefit obtained through advance reservations is shared means that the provider gains when big savings are achieved, and doesn't lose out when a loss is likely. The VR model is also not negatively impacted as a result of changes in the price of electricity futures and, unlike the CPP model, generates more profit than on-demand during these price hikes. Figure 5 shows the purchase of electricity by the provider. Figure 6 shows the purchase of resources by consumers.

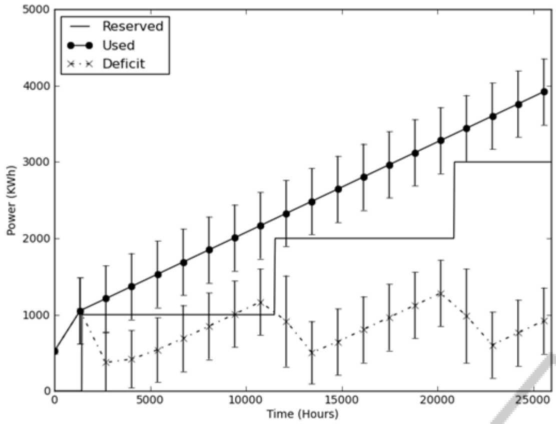

Figure 5: Power used by forwards and VR model.

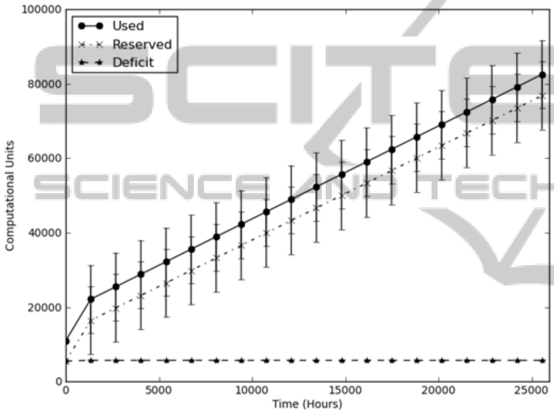

Figure 6: Resource allocation using forwards and VR.

Forwards are generally less profitable than on-demand resources, by quite a large margin. Clearly, the pricing is too low for purchasing advance reservations to be profitable, but determining this price is not easy as the number of reservations is not known until they have all been requested.

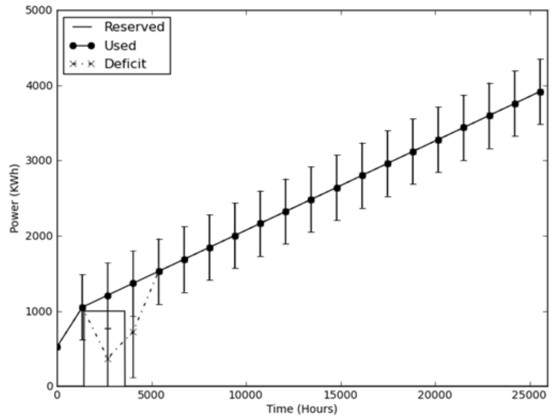

Figure 7: Power used by GPP model.

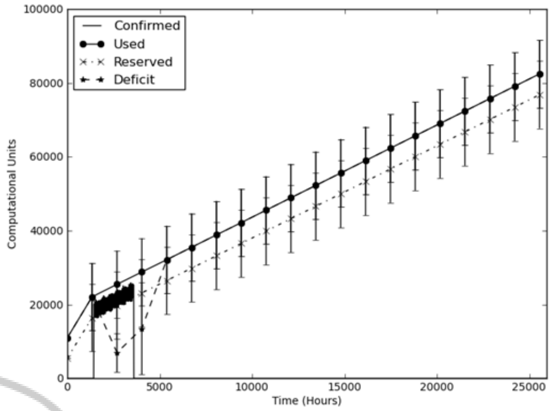

Figure 8: Resource allocation using GPP model.

The GPP model also fails to deliver significant gains in profit. The model protects losses as a result of not giving discounts when forwards are less profitable than on-demand, but it doesn't achieve high profits when forwards are more profitable as everyone receives the discount (figures 7 and 8).

3.2.2 Provider Cost Reduction

Forwards are generally the most beneficial to the consumer achieving a mean saving of around 20% the cost of an on-demand instance, and reducing costs for all market demand profiles (Figure 9). This is because the consumer always gains access to the resource, and thus their net costs are reduced. The mean price does not equal the cost of the reserved resource in all situations because sometimes a user purchases a resource that subsequently she does not require, but which she has already agreed to pay for.

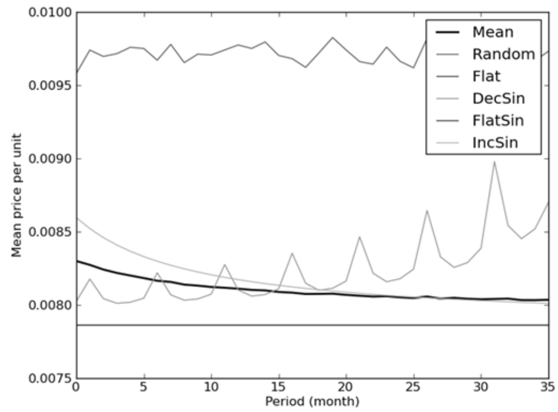

Figure 9: Mean price per computational unit using forwards.

The CPP model rewards consumers with around a 15-10% saving when a cost saving is achieved, with this rising to no discount when electricity prices increase (Figure 10). All consumers make a saving using the CPP model.

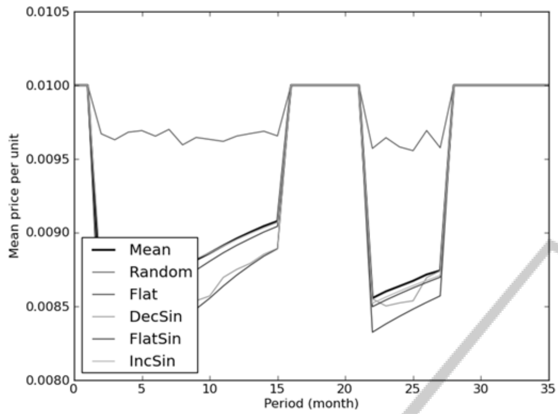

Figure 10: Mean price per computational unit using CPP.

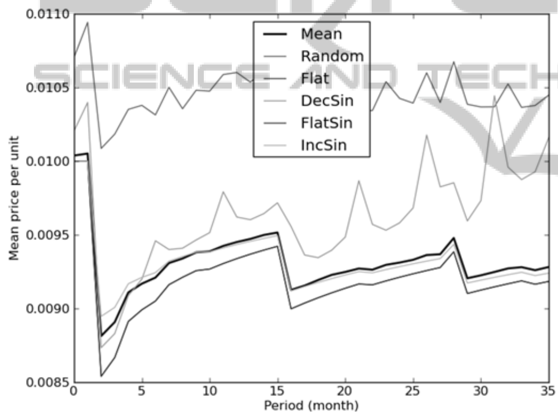

Figure 11: Mean price per computational unit using VR.

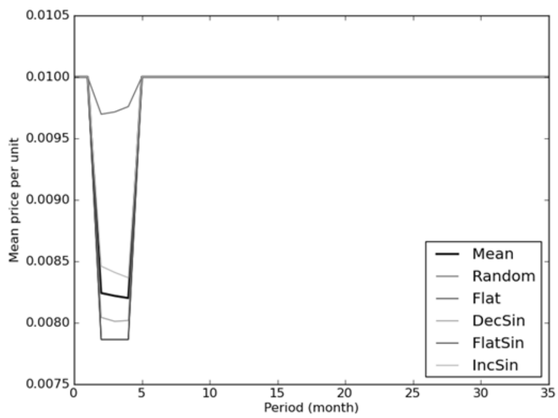

Figure 12: Mean price per computational unit using GPP.

The VR model provides a discount of around 5% on average (Figure 11). However, users with random demand profiles spend more using the scheme than if using just on-demand resources. This is as a result of poor predictability, which results in the purchase of advance resources which are subsequently not used.

The GPP model is generally unattractive to consumers (Figure 12). Only occasionally is a discount awarded, and it is unlikely this would not occur enough to be of interest.

4 CONCLUSIONS

In this paper, we have introduced and analysed a number of novel pricing schemes for cloud computing which we have shown to offer opportunities for increasing profits by reducing the cost of purchasing electricity.

Group Provision Points are unlikely to be implemented in a commercial offering, as the scheme does not take full advantage of information acquired through the sale from consumers. Consumers do not receive regular enough discounts to make forecasting worthwhile, nor does the provider benefit from reduced electricity costs.

We believe Contributory Provision Points and Variable Reward Forwards are the most attractive of the schemes discussed. Contributory Provision Points will favour those who can predict their future demands earlier. Variable Reward Forwards gives everyone who contributed to a reduced cost with a share of the saving. It is likely Variable Reward Forwards would be seen as fairer by the user-base as everyone is rewarded; not just those who contribute to the discount, which cannot be established beforehand.

Both of these schemes can be configured to outperform on-demand pricing by setting reserved pricing appropriately.

However, because of the size of the electricity futures involved, only larger providers would be able to take advantage of the schemes.

Further investigation should be conducted on how these schemes can be used in bilateral arrangements, where datacentres may produce their own electricity which may be ploughed back into the electricity grid. Furthermore, can these schemes be enhanced through the use of cloud spot-markets, or reserved instances?

In this work, we assumed air conditioning was a fixed cost which doesn't change with increasing number of servers. However, a gradual increase in air conditioning energy is likely to be seen as a result of the increased heat generated by servers. In future work, including air conditioning costs in the model could further reduce expenditure.

A significant amount of work is still required to determine if these schemes can be implemented commercially. In future work, we plan to create a simulation of a competitive market of providers utilising the scheme. Our objective is to see if one scheme becomes dominant in the marketplace. We also wish to investigate if the providers can change pricing with a view to acquire more business. This could eventually lead to a market for provision point contracts in cloud computing.

ACKNOWLEDGEMENTS

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

REFERENCES

Bagnolli, M, and B Lipman. 1989. "Provision of Public Goods: Fully Implementing the Core Through Private Contributions." *Review of Economics Studies* 56: 583–601.

Barroso, L. A., and U. Holzle. 2007. "The case for energy-proportional computing." *Computer* 40(12): 33–37. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4404806 (November 30, 2011).

Barroso, Luiz André, and Urs Hölzle. 2009. "The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines." *Synthesis Lectures on Computer Architecture* 4(1): 1–108. http://www.morganclaypool.com/doi/abs/10.2200/S00193ED1V01Y200905CAC006.

Berl, a., E. Gelenbe, M. Di Girolamo, G. Giuliani, H. De Meer, M. Q. Dang, and K. Pentikousis. 2009. "Energy-Efficient Cloud Computing." *The Computer Journal* 53(7): 1045–1051. http://comjnl.oxfordjournals.org/cgi/doi/10.1093/comjnl/bxp080 (November 7, 2012).

Brown, R. 2008. "Report to congress on server and data center energy efficiency: Public law 109-431." http://escholarship.org/uc/item/74g2r0vg.pdf (January 21, 2013).

Buchbinder, N., N. Jain, and I Menache. 2011. "Online job-migration for reducing the electricity bill in the cloud." *NETWORKING 2011*. http://www.springerlink.com/index/86060362741H6X88.pdf (January 23, 2013).

Guo, Yuanxiong, Zongrui Ding, Yuguang Fang, and Dapeng Wu. 2011. "Cutting Down Electricity Cost in Internet Data Centers by Using Energy Storage." In *2011 IEEE Global Telecommunications Conference - GLOBECOM 2011*, IEEE, p. 1–5. http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6134209 (January 23, 2013).

Hull, John C. 2008. *Fundamentals of Futures and Options Markets*. 6th ed. Prentice Hall.

Lee, Young Choon, and Albert Y. Zomaya. 2010. "Energy efficient utilization of resources in cloud computing systems." *The Journal of Supercomputing*. http://www.springerlink.com/index/10.1007/s11227-010-0421-3 (July 18, 2011).

Patel, C. D., and A. J. Shah. 2005. "Cost model for planning, development and operation of a data center." *Development* 107: 1–36. http://www.hpl.hp.com/techreports/2005/HPL-2005-107R1.pdf (November 30, 2011).

Qureshi, Asfandyar, Rick Weber, Hari Balakrishnan, John V. Guttag, and Bruce Maggs. 2009. "Cutting the Electric Bill for Internet-Scale Systems." http://dspace.mit.edu/handle/1721.1/62585 (January 21, 2013).

Rao, Lei, Xue Liu, Le Xie, and Wenyu Liu. 2010. "Minimizing Electricity Cost: Optimization of Distributed Internet Data Centers in a Multi-Electricity-Market Environment." In *2010 Proceedings IEEE INFOCOM*, IEEE, p. 1–9. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5461933' escapeXml='false'/> (January 7, 2013).

Rogers, O., and D. Cliff. 2012. "Options, forwards and provision-point contracts in improving cloud infrastructure utilisation." *Journal of Cloud Computing: Advances, Systems and Applications* 1(21).

Rogers, Owen, and Dave Cliff. 2012. "The Use of Provision Point Contracts for Improving Cloud Infrastructure Utilisation." In *Grid Economics and Business Models*,.