

RESEARCH

Open Access

Contributory provision point contracts – a risk-free mechanism for hedging cloud energy costs

Owen Rogers* and Dave Cliff

* Correspondence:
csorr@bristol.ac.uk
Department of Computer Science,
University of Bristol, Merchant
Venturers Building, Bristol, UK

Abstract

Cloud computing services rely on electricity to power compute-servers, network equipment, cooling systems, and other supporting infrastructure. As such, energy costs are a substantial outgoing to public providers of cloud computing services. On-demand pricing, where consumers are not required to give advance notice of requirements, does not aid the provider in planning future demand, and therefore makes it more difficult to purchase energy at discounted rates. In this paper, we propose an advance pricing mechanism for cloud computing resources based on *provision-point* contracts, commonly used by deal-of-the-day websites such as Groupon. We show how our Contributory Provision Point (CPP) contracts reward consumers with reduced prices for advance reservations, while allowing providers to make accurate forecasts of energy usage. We show how CPP contracts are risk-free for the provider, guaranteeing to be at least as profitable as on-demand mechanisms where electricity is purchased ad-hoc by the provider. Through a computer simulation, we demonstrate that CPP contracts can be more profitable for the provider compared to a traditional method of hedging electricity futures using a popular forecasting algorithm. Furthermore, we show that CPP contracts encourage consumers to forecast honestly by rewarding them with discounted rates, while remaining profitable for the provider, even when forecasts are not completely accurate.

Keywords: Utility computing, Market-orientated computing, Assurance contracts, Hedging, Energy futures, Forward contracts, Derivatives

Introduction

Consumers of infrastructure-as-a-service resources typically pay a single price to access a virtual machine for a specified period of time. This price covers the costs of the virtual machine's fraction of the physical server, any maintenance and repairs, the physical datacentre space, the energy needed to power it, the cost of air conditioning to cool the datacentre, and its fraction of other costs incurred in creating the cloud service, such as administration and advertising.

One of the primary attractive features of cloud computing is on *on-demand* pricing, where consumers gain access to the resource immediately and are charged for the amount of time they use the resource [1]. An alternative is advance or *forward* pricing, where consumers gain access to the resource at a specified time in the future, and have access for a pre-agreed duration.

The provider has both fixed and variable costs. Air-conditioning and datacentre space are examples of fixed costs; regardless of how many servers are placed in the datacentre, these costs will essentially be the same.

The provider also has variable costs, where the cost is proportional to the amount of virtual machines demanded by the provider's customers. Electricity costs for powering servers are variable costs.

The public cloud provider relies on electricity to successfully provide consumers with a service. As such, energy costs can be a significant outgoing 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% [2-4]. This cost impacts the price paid by consumers to access virtual machines, and the profit achieved by the provider. In a competitive marketplace, keeping prices low is critical for commercial success.

Currently, research effort is being directed at reducing the power consumption of computing technology [5,6]. The primary objective of this research is to reduce carbon footprint, but reducing expenditure is also an important factor.

Many datacentres purchase electricity as per most household consumers; on-demand for a fixed cost per unit consumed. This could be directly from an energy supplier, or from an energy broker who hedges market-traded instruments to offer fixed prices to its clients. The broker is typically an intermediary party who sits between the electricity providers and cloud providers, negotiating between both parties. Agreements will exist between sellers, buyers and brokers to ensure that only purchased electricity is consumed, and penalties will be payable in the event of breaking these agreements.

However, some larger data centres may be built with their own power source such as hydro-electric generators or solar cells. In this case, the datacentre can also become a provider of electricity, pumping unused electricity into the electricity grid and subsequently receiving a payment. Such a system is governed through the use of a *bilateral agreement*, which details costs and conditions for both selling and purchasing electricity resources.

Some research has been conducted on optimising the placement of virtual machines such that the cheapest electricity available is used by the datacentre.

Qureshi *et al.* were the first to propose dynamically assigning computational workloads in distributed systems to locations where electricity may be cheaper. Their simulation studies indicated that savings of millions of dollars could be achieved using their method [7].

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 [8].

Buchbinder *et al.* extended these methods so that only batch applications would be migrated to cheaper markets [9]. 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

A currently under-explored area of research is if cloud providers can use derivative contracts in energy to decrease its costs [10]. *Futures contracts* are a type of financial 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 specifies. The provider could then engage a broker to provide fixed-price electricity to top up its pre-bought electricity capacity.

A futures contract specifies the size of the commodity being purchased. In electricity futures, the commodity is a quantity of electrical power (typically a MW) delivered for a fixed period of time, typically a month or a quarter-year.

However, the use of electricity futures can have significant associated risks. If the provider buys a future contract for electricity which is subsequently not fully utilised by customers, then it is possible that the provider will make a loss. Electricity delivered to the cloud provider cannot be stored; if is not used as it is delivered, then it is wasted. Given that an electricity future for one month's delivery of 1 MW costs over £35,000, this risk can be sizeable.^a

In this paper, we propose a pricing mechanism that allows cloud providers to purchase electricity futures with no risk that they will subsequently fail to utilise their investment effectively; the mechanisms guarantee that the provider will make at least as much profit as offering on-demand resources and purchasing electricity ad-hoc.

The novel mechanisms we explore in this paper are based on *provision-point contracts* (also known as *assurance contracts*). In such a mechanism, members of a group pledge to contribute to an action if a threshold of some type 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 the monies paid are then refunded [11].

Such a mechanism is used by deal-of-the-day website *Groupon.com*. Users make requests for special offers by purchasing an offer. When a threshold number of users is reached, the offer is profitable to the provider and the offer is confirmed. In this case, both the provider and consumer benefit; if the threshold is not reached, those users who did purchase the offer have their costs refunded.

In previous work [12], we showed how provision-point contracts can be used to schedule virtual machines more effectively on a large-scale cloud infrastructure.

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

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.

In this paper we show how our *Contributory Provision Points* contracts allow providers to make accurate forecasts of energy usage and therefore reduce their costs through the purchase of electricity.

Furthermore, we present results from a computer-based sensitivity analysis of our model. We show that CPP contracts can achieve larger profits for the cloud provider compared to traditional hedging of electricity resources using a popular accurate forecasting algorithm. We also show that CPP contracts are more profitable for the provider than hedging energy using standard forward contracts. Furthermore, our results show that our provision-point schemes encourage honesty in the user population, but can tolerate small deviations in forecasting accuracy to provide both consumer and provider with financial benefits.

Pricing mechanisms

On-demand pricing

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

The provider charges customers a cost C_o to use the computing resource for an interval i . The interval i is chosen by the provider, and represents the minimum period (typically an hour) for which a resource may be used, and for which the consumer will be charged. 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 E_o for electricity for the duration of the interval that the virtual machine is running; we assume this cost is constant, and doesn't vary with time of day or quantity. 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:

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

Forward Contracts

Now, 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 allow 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 total resources reserved at the reserve price, plus the additional resources bought on-demand at the on-demand price.

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

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 electricity futures contracts, where each contract entitles her 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 total cost of purchasing on-demand electricity required in addition to the reserved electricity.

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

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

$$P_{res} = REV - COE$$

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

For the mechanism 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 consumer, the consumer must be incentivised to provide a forecast. Therefore, reserving a resource must be cheaper than buying a resource on-demand.

$$C_o > C_r$$

Therefore, the conditions for the mechanism to be beneficial to all parties are:

$$C_o > C_r \tag{1}$$

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

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 are incentivised to use the service.

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

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

Provision-point contracts

This issue can be circumvented by introducing a provision-point contract into the mechanism. An intermediate phase – the ‘confirmation phase’ – is now introduced:

1. Reservation Phase: Consumers 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 mechanism by conditions (1) and (2) being met, they will confirm consumers' requests and the contracts are confirmed. If either condition is not met, all contracts are cancelled.
3. Execution Phase: Consumers gain access to their confirmed resources, and may also buy additional on-demand resources.

If the requirements of the consumer 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 mechanism 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. This format of provision-point contract is hereby referred to as a Group Provision Point (GPP) contract.

Contributor Provision Points (CPP)

The forward and group provision-point mechanisms are extremes. In the forward mechanism, all consumers who submit a reservation benefit from reduced prices, in spite of it sometimes not benefitting the provider. In the group provision-point mechanism either all, or no, consumers 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. 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.

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

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 only q contracts are confirmed, and all others are cancelled:

$$q = \sum_0^N r_i$$

Substituting into (2):

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

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

The vulnerability of the forward contracts 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 mechanism conditions (1) and (3) are met and E_r is set to be the maximum likely cost of an electricity future, the mechanism will generate a profit over on-demand pricing.

This mechanism 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:

$$E_r < IE_o - \frac{1}{\beta}(C_o - C_r)$$

Simulation

In Pricing mechanisms section, we showed that CPP contracts are guaranteed to be at least as profitable to the provider as traditional on-demand contracts, where the provider purchases electricity *ad hoc* to meet demand.

However, a more likely situation is that the provider would independently predict future demand based on past performance and would purchase electricity futures based on this forecast. In this scenario, the provider could reduce energy costs by purchasing in advance, without asking consumers for advance reservations or giving them a discount. However, this does involve risk for the provider – if the forecast is wrong, the provider may invest in energy which is subsequently not used.

To determine if CPP contracts could generate more profit for the provider than forecasting, a simulation was coded in *Python* [13], which used the *RPy*^b package to interface with the statistical programming language *R* [14].

It is likely that determining in exactly what situations when our models will outperform provider forecasting will be an impossible task. Different providers will use different methods for forecasting, will understand their respective user-base to different degrees, and will take different approaches to risk. Furthermore, different providers will have different target markets which may have differing levels of volatility.

We wish to compare the performance of the provision-point mechanism to that of a provider who has a very accurate view of future demand. We assume that in a real-world scenario, providers will submit a less accurate forecast than our ideal simulated scenario. In this manner we can determine the minimum benefit achieved by the CPP mechanism compared to forecasting.

In these simulations, we give the provider the ability to forecast accurately by choosing a simple and predictable waveform to represent market demand (a sinusoid with a small amount of noise) and a simple and powerful forecasting algorithm (Holt-Winters) set with ideal parameter-values (derived from preliminary tests).

Our objective is to determine if there is a potential for providers to benefit from our models. In Discussion section, we give further discussion of deploying these models commercially.

We use two outcome variables to measure performance:

- The increase in the provider's profitability over the simulated period using our model compared to a traditional on-demand model where the provider forecasts energy needs.

- The decrease in the consumers' mean cost per resource over the simulated period using our model compared to a traditional on-demand model where the provider forecasts energy needs.

For the model to be useful the provider must increase profitability and the consumer must reduce expenditure.

We manipulate three experimental variables:

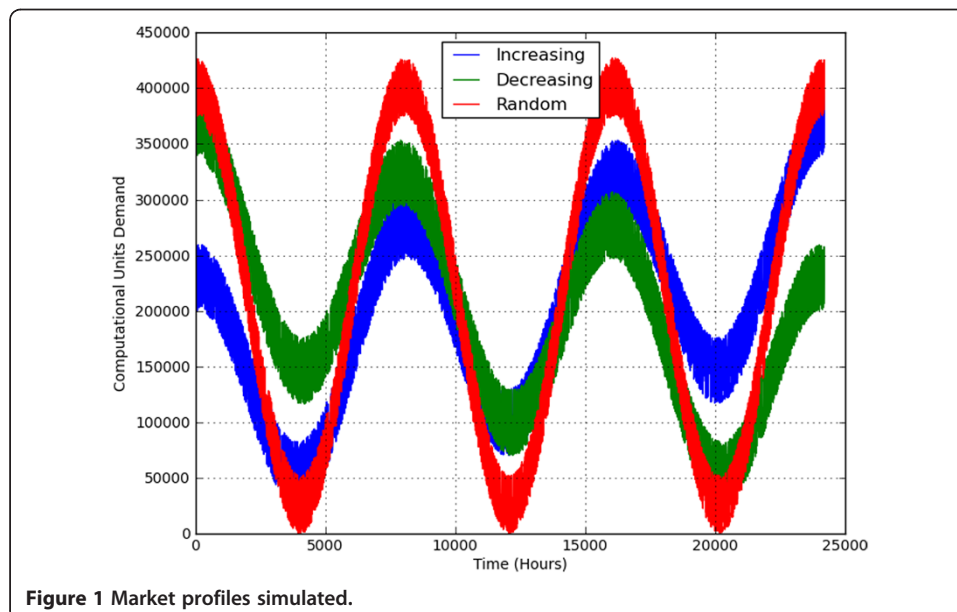
- The cost of an advance resource.
- The accuracy of consumers' predictions, expressed as a percentage overestimate.
- The demand for resources over the simulation

Resource costs

We choose the cost of a resource to be £0.006, which is approximately the Pound Sterling value of an Amazon Web Services micro-instance hour in December 2012. Our objective is to determine relative savings, rather than absolute benefits. Note that the cost of an advance resource is different to the decrease in consumers' mean cost per resource, as the consumer may purchase resources that are subsequently not used. In this case, the mean cost per resource may increase but the cost of a resource will remain the same.

Demand profiles and forecast accuracy

We simulate a population of 3000 users, across three market profiles. A market profile defines the change in demand experienced by the provider over the simulation period. We use three market profiles, as shown in Figure 1. These market profiles show increasing, decreasing, and flat trends, with random fluctuations and a sine periodicity.



Eighty percent of consumers in the simulation experience demand as a function of the form $1 + \cos(2\pi h/24)$, where h is the hour-number in the day. Twenty percent of consumers experience demand chosen randomly from a normal distribution. This ratio was chosen so that the random component would still play an appreciable role in determining overall demand but would not swamp the predictability of the sinusoid profile.

Consumers follow a highly conservative approach of only reserving a resource if it has been used at the same time for the last three periods.

A cosine was chosen as it reflects a seasonal aggregate demand where there are periods of high demand (eg. at Christmas) followed by a lull in sales. A cloud provider whose customers may be retailers might experience such seasonal cycles. Most importantly, the cosine is representative of a trend that should be relatively easy to predict by the provider using a forecasting algorithm.

The forecast accuracy describes the percentage amount by which consumers overestimate usage, and thereby purchase advance resources. A negative value represents an underestimate.

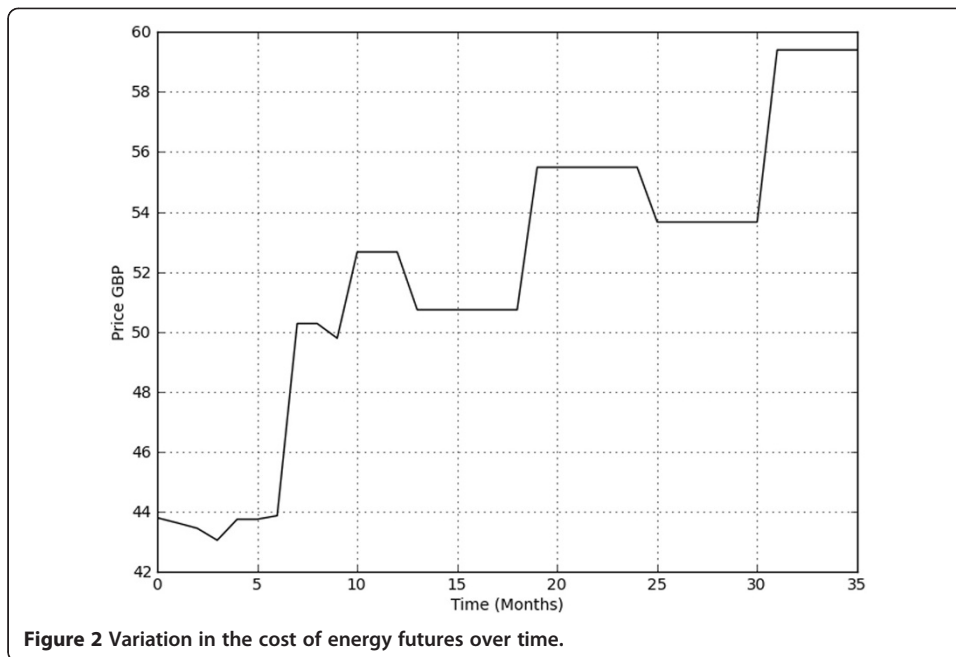
Forecast accuracy is an important characteristic of consumers for two reasons. Firstly, consumers may try to obtain a greater discount by submitting dishonest forecasts to the provider. So it is important the mechanism must not reward consumers for dishonest behaviour, which will impact the ability of the provider to use the forecast. Secondly, consumers may purchase advanced resources based on what they believe to be an honest forecast, but where the forecast may become inaccurate as a result of changing circumstances. In this case, the mechanism must allow some margin of error in consumer forecasts, as consumers are unlikely to be completely accurate all of the time.

Provider forecasting

For electricity, we assume that the provider may purchase electricity futures for a period of four weeks, 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 2). The cost of electricity on-demand from the grid is £0.01/kWh, based on [2].

The provider will purchase an electricity future if its estimated expected usage, based on consumers' advance reservations, is such that it will be cheaper than subsequently buying the same amount of energy on-demand.

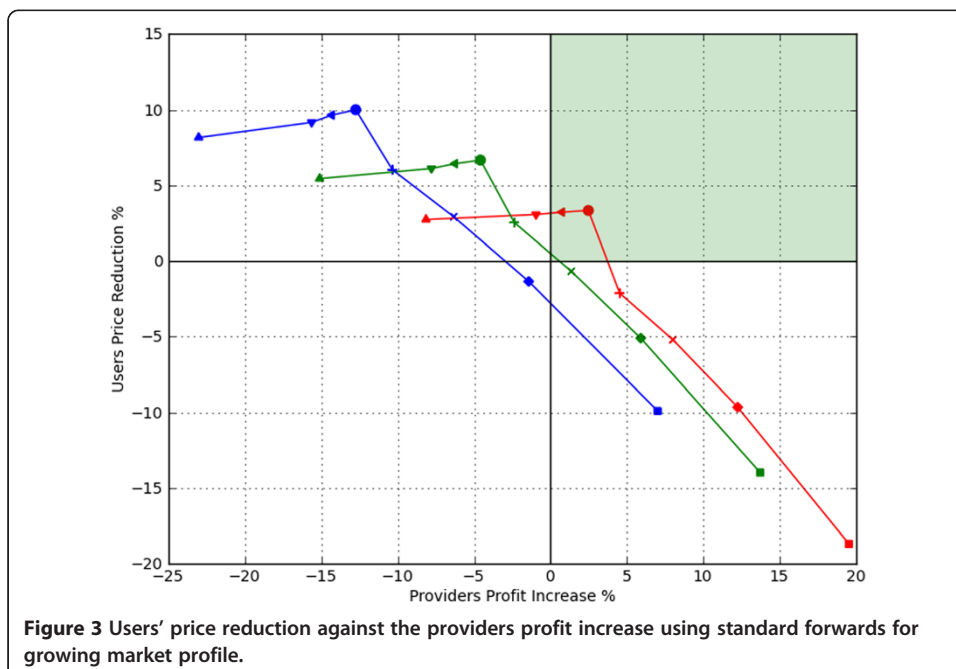
As a first approximation to real-world forecasting approaches, in the simulations reported here the provider uses the Holt-Winter exponential smoothing algorithm to forecast energy usage [15]. This algorithm is easily implementable and has successfully been used previously to forecast energy usage so is likely to be an option for a provider who wishes to forecast energy usage [16]. The algorithm allows the specification of up to three parameters, α , β and γ , which define the amount of smoothing, trend and seasonal effect respectively to be applied where forecasting. During preliminary studies on our demand profiles, we found that the specifying values of β and γ caused large fluctuations in forecasting demand. We decided to use the "single" version of the algorithm (where β and γ are set to 0) as this was the most stable method of forecasting for the provider using our test market profiles.



The provider computes a value of α , the smoothing factor, from historical data of market performance once every six months. During preliminary investigation, it was found that this process proved more accurate than the case where the provider was supplied with α calculated from the entire demand profile in advance.

Results

Figures 3, 4, and 5 shows the users' price reduction against the providers profit increase for a range of resource prices and accuracies, when using standard forward contracts. Figures 6, 7, and 8 shows the same measures using CPP contracts. In these figures, colours represent



the resource price, and different markers represent forecast accuracy as shown in the legend in Figure 9.

For a mechanism to be profitable, it must increase the provider's profits while reducing users' costs. In Figures 3, 4, 5, 6, 7, and 8, this occurs in the green shaded quadrant. In all market profiles, there is at least one resource price which is found to be both beneficial to the consumer and provider for both standard forwards and CPP contracts.

In both standard forwards, and CPP contracts, the user benefits financially most when they are honest. A reliable forecast will give them the most cost saving. Underestimating usage results in a smaller saving for the consumer as they fail to maximise cost savings through the purchase of cheaper resources. However, overestimating often results in a smaller saving, or even an overspend compared to on-demand purchasing, as a result of purchasing resources that are subsequently not needed. Overspending is most likely to be seen where the advance resource cost is close to the on-demand cost. This implies consumers are more likely to underestimate to protect their risk.

In our simulation studies, CPP contracts are always more profitable to the provider than standard forwards. In some cases, this means the provider can make more of a profit for the same advance resource price. For example, in the flat market, an advance price of £0.0056 produces only a 1% increase in profit for the provider. In the same market when CPP contracts are offered, an increase of 4% is achieved. The sacrifice is consumers' mean cost, because in CPP contracts the consumers are not guaranteed to receive the discount. However, in terms of an attractive product, being able to offer computational units with a lower list price is likely to encourage consumers. It is possible consumers would make decisions based on published costs rather than their mean cost per resource.

Setting the price of CPP contracts is essentially a business decision, based on how much of a saving consumers would typically expect for requesting resources (especially

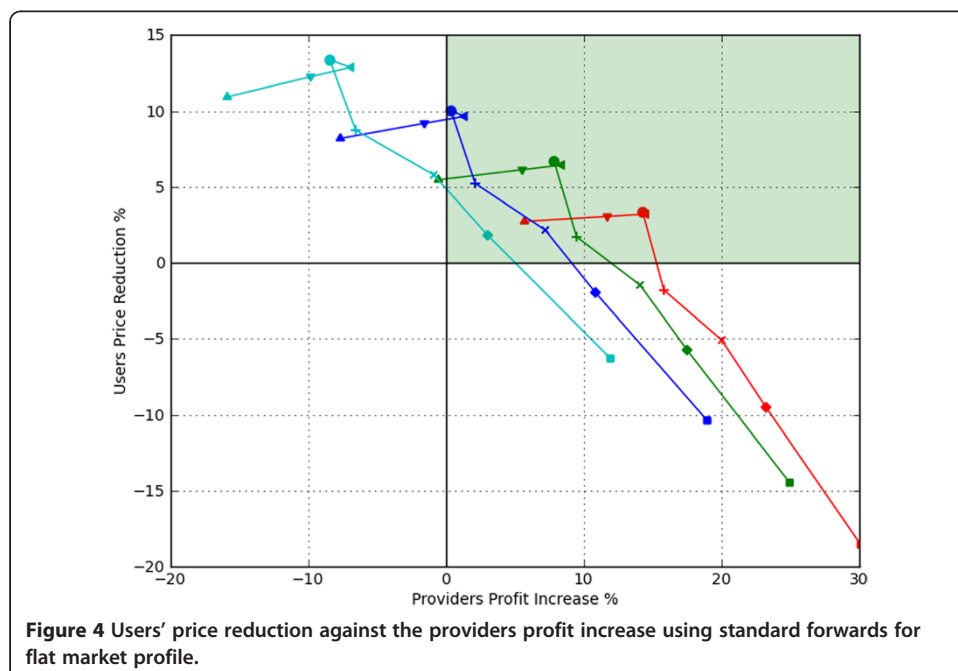



Figure 4 Users' price reduction against the providers profit increase using standard forwards for flat market profile.



considering that subsequent supply is not guaranteed) against how much of a profit increase is required by the provider for the mechanism to be worthwhile.

In the flat market, there are a number of CPP contracts prices which are beneficial to both parties (Figure 7). In a competitive market, the provider would be better off offering a cheaper list price to entice customers, even if this results in a smaller profit increase as a result of reduced energy costs.

A cost of £0.0058 suits everyone in these simulations. Customers make a minimum saving of 3%, the provider makes a saving of between 3 and 15% and the user can underestimate their usage by 10% with both parties still benefitting. However, the consumer must avoid overestimating as this can quickly result in a loss for the consumer.


Discussion

Contributory provision points in a cloud environment

We believe that provision-point contracts can be used in combination with on-demand pricing to give consumers the flexibility of being able to start and stop computing resources as and when required, while benefitting from discounted resources.

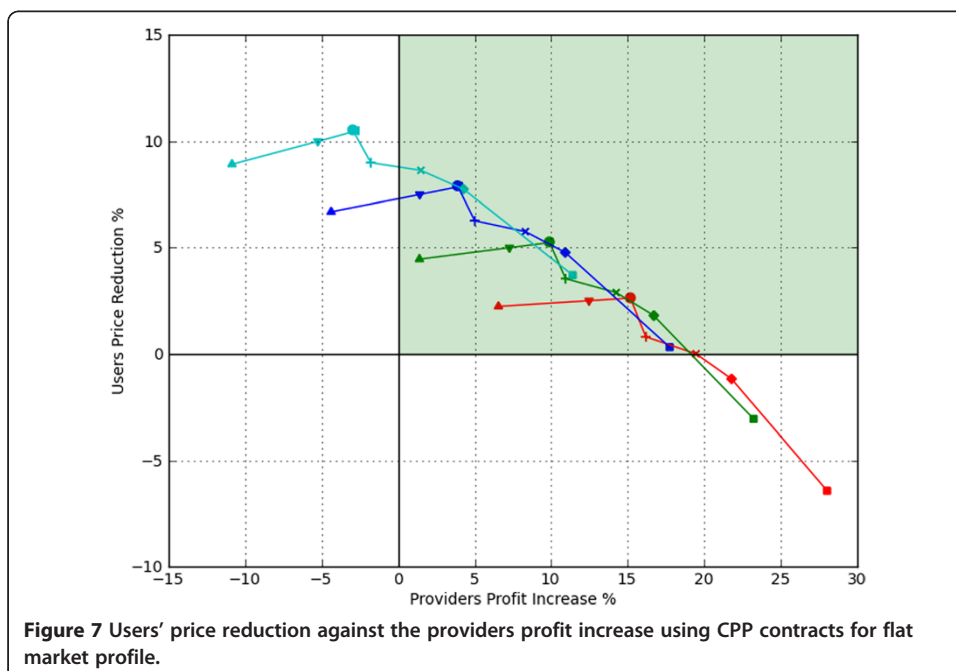
Most applications deployed on a cloud infrastructure require a minimum number of virtual machines to operate around the clock. For example, a typical web application would always require a web server, an application server and a database server. The owner of this web application can seek a discount by requesting three CPP contracts for three virtual machines on which are run a web, application and database server. The consumer knows that these three virtual machines will be required in the next period, so has nothing to lose by making a request for them.


Consider a consumer which has gained access to these virtual machines through CPP contracts and has therefore obtained a discount. If there is sudden surge in demand, and the application is struggling to handle web requests, the consumer can still



purchase an on-demand virtual machine to provide the application with an additional web server. The only pre-requisite is that the application has been coded to integrate with resources started on-demand.

In this scenario, the consumer has obtained a discount by purchasing CPP contracts for the baseline application servers, but can still grow capability to meet unexpected demand by purchasing additional on-demand virtual machines.





The same consumer might notice that at a particular time of day, the application server struggles to handle requests perhaps as a result of an update process. The consumer could request an additional CPP contract for a virtual machine that can be used to host an additional application server while this update takes place. Again, the consumer has reduced her expenditure by requesting a reservation.

Consumer attractiveness against on-demand pricing

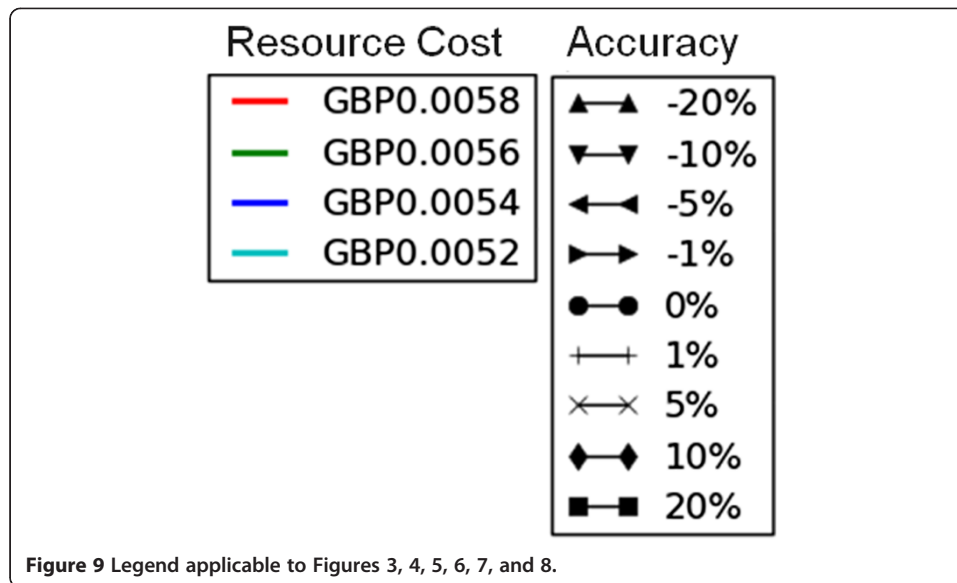
One obvious question prompted by our work is this: will consumers reserve resources if there is no guarantee on the provider's part that they will be delivered (prior to the confirmation phase)?

There are already a large number of deal-of-the-day websites which rely on consumers purchasing offers with no guarantee that the offer will be confirmed: *Groupon*; *LivingSocial*; *TravelZoo*; and *Wowcher* are just a selection. For the consumer, they have nothing to lose financially – if the deal is not confirmed, they are refunded whatever was paid. If the deal is confirmed, the consumer benefits with a discount.

However, these websites are aimed at consumers – would businesses that rely on cloud infrastructure for critical business operations use CPP contracts?

Amazon Web Services already offer non-guaranteed reservation of computing resources through its Spot Market [17]. In this model, a consumer places a bid for a quantity of resources. If the current spot price is below the bid, the consumer receives access to the resource at a cost lower than the on-demand price. If the spot price rises above the bid, the consumer loses access to the resource.

In this model, the consumer has no guarantee of receiving the resource. They request a resource through their bid. The consumer does not know the request is confirmed, until the spot price changes in their favour and they are given access.



If the consumer is confident they will need the resource, they have nothing to lose by placing a bid. If the bid is not successful, they will purchase resources at the on-demand price. If the bid is successful, they will have access to resources for a lower cost.

The success of AWS Spot Instances implies that CPP contracts will be desired by the consumer. Many businesses will have an understanding of their future demands, and so have nothing to lose by making a request for discounted future resources. These consumers have the opportunity to make savings by purchasing CPP contracts. This saving is not guaranteed, but the consumer would have needed to buy the resource anyway.

Consumer attractiveness against forward pricing

In a competitive market, would consumers prefer to use standard forwards or provision-point contracts? If consumers are savvy then they would choose to assess benefit based on the mean cost per resource. In this case, the consumer would prefer to use standard forwards as this delivers a better saving.

However, consumers are not necessarily rational. The provider will be able to offer a cheaper list price for CPPs than standard forwards while receiving the same profit. Consumers are likely to be enticed (at least initially) by the cheaper list price of the resource.

For example, consider the shrinking profile where the provider is happy for an advance model to be as profitable as an on-demand model, perhaps as a means of attracting new customers through discounted pricing.

In this scenario, CPP contracts breakeven when the resource price is £0.0056. However, standard forwards breakeven when the Resource Price is approximately £0.00564. So to the consumer, the most attractive model would appear to be the CPP mechanism.

Performance and implementation

It is difficult to compare the performance of CPP contracts against forecasting using other techniques. The ability for a cloud provider to predict its energy usage depends on a large

number of things, and this is likely to vary widely between providers [18]. For example, the following may affect the provider's ability to forecast:

- Understanding of its market, and cycles or trends;
- Forecasting algorithm chosen;
- Period in existence (for historical data);
- Accuracy of smoothing parameters;
- Customer churn or growth, marketing campaigns;
- Effective CRM management and sales process.

The CPP mechanism relies on consumers to forecast usage. Again, the ability of the consumer to predict usage relies on a number of factors, including those specified for the provider. However, it is a reasonable assumption that a consumer of resources is likely to be able to predict future demands to a higher degree of accuracy than the provider. This is simply because the consumer has a more detailed understanding of its target market, customer cycles and trends.

In the simulations reported here, simple waveforms were chosen that should be easy to predict by the provider. Furthermore, the provider was configured such that it would learn the best smoothing-factor time value. Even in these unrealistically ideal circumstances (i.e., forecasting for real-world providers is a much more challenging issue, and hence likely to be significantly less accurate), it was found that CPP contracts still outperformed forecasting and forward contracts. We believe that if our simulations has used more challenging circumstances, making the provider's forecasts more error-prone, it is clear that CPP contracts would perform even better.

A provider considering offering such a mechanism should run side-by-side comparison trials to determine if CPP contracts perform better than their current methods of forecasting demand.

Conclusions

In this paper, we have proposed a novel pricing model for cloud services based on provision-point contracts. We have shown that Contributory Provision Point contracts are guaranteed to be as least as profitable as on-demand pricing mechanisms in the case where the provider purchases electricity on an *ad hoc* basis. We have also shown how CPP contracts are more profitable for the provider than standard forwards.

Through simulation studies, we have shown how CPP contracts can outperform on-demand pricing where energy is purchased based on the Holt-Winters forecasting algorithm, both by maximising profit and reducing consumer expenditure. Furthermore, our results show that our provision-point mechanisms encourage honesty in the user population, but can tolerate small deviations in forecasting accuracy to provide both consumer and provider with financial benefits.

The mechanism is not guaranteed to perform better than a cloud provider that offers on-demand pricing and hedges energy costs with a high degree of accuracy. However, the model is guaranteed to perform better than a cloud provider who offers on-demand pricing and purchases energy in an on-demand manner. This means the provider is guaranteed to receive a minimum gross margin whether forecasting or not; a

purchased energy future will always be utilised enough that a loss cannot be made. As such, the mechanism carries no risk to the provider that investments in electricity will not be worthwhile.

CPP contracts could potentially be offered by any provider of a service, where the providers' variable costs can be reduced by reserving in advance.

We believe that CPP contracts could be implemented on a commercial basis by providers of cloud computing services. The existence of numerous deal-of-the-day websites and AWS Spot Prices implies consumers are willing to request discounted resources, even if they are not guaranteed to gain access to them. However, we believe it is difficult to compare our method for hedging energy costs against forecasting based on historical data, simply because of the number of variables that will differ between providers. As such, we suggest providers considering offering CPP contracts to consumers should trial the model next to their existing hedging mechanism to see if financial gains can be realised.

The mechanism described here relies on a fixed cost for on-demand electricity to be available, and agreed in advance by an electricity broker or supplier. It is likely this could be achieved commercially, but as the broker is absorbing risk on behalf of the provider, the consumer will probably pay more than if they purchase electricity from a spot market. Further work should be conducted on if the provision-point mechanism can be used to hedge electricity where the provider purchases electricity on-demand from a spot market. In such a mechanism there will always be an element of risk, as the spot-price might rise such that the energy future is no longer financially beneficial. However, provision-points could potentially lower this risk by giving the provider the opportunity to cancel contracts if it sees a trend that indicates a likely future loss.

Provision-point contracts can play a part in reducing costs where cloud providers with their own electricity generating equipment participate in bilateral agreements with energy providers: PPCs could give providers an indication of future demand, which could then potentially be used to sell or buy electricity at a better price. Exploring this is another area of future research.

In this paper, air conditioning has been considered a fixed cost for simplicity. However, this is not actually the case in a large scale datacentre; more servers will require more energy to be fed into coolers to counteract the extra hot air produced by those servers, although the rate of change may be very small. It is likely to be the case that improved server consolidation will result in less consumption of energy for air conditioning. Furthermore, energy futures could be purchased such that air conditioning cooling costs will be reduced as well as direct server costs.

Our intention is that this paper provides a basis for further investigation into the use of derivative contracts for reducing cloud datacentre costs: clearly there are many more issues that remain to be explored.

Endnotes

^aICE UK Base Electricity Futures, November 2012.

^bRpy.sourceforge.net.

Competing interests

Owen Rogers' PhD studentship is part-funded by Hewlett-Packard Research Labs. Primary funding for both Rogers' and Cliff's work on this project was provided by a research grant from the UK Engineering and Physical Sciences Research

Council. The authors list these sources of funding for completeness, and as a mark of acknowledgement and grateful thanks. Neither author considers these sources of funding to constitute competing or conflicting interests in the work.

Authors' contributions

OR designed and implemented the simulations, analysed the results, and wrote the first complete draft of this paper. The novel aspects of the CPP contract model described in this paper are entirely OR's contribution. DC provided direction and advice, proposed research areas and reviewed and edited earlier drafts of this paper. Both authors read and approved the final manuscript.

Authors' information

Owen Rogers is Senior Analyst for Digital Economics at 451 Research. Owen graduated from Cardiff University with a MEng in Computer Systems Engineering in 2005. Following this, he joined telecommunications firm Cable & Wireless, first as a Graduate Engineer, and then as Product Development Manager for Managed Security Services. In 2009, Owen joined managed services provider Claranet as Product Portfolio Manager, before returning to academia in 2010. He completed his PhD in early 2013 at the University of Bristol. Owen is a Chartered Engineer through the British Computer Society.

Dave Cliff is a professor of computer science at the University of Bristol. He has previously held faculty positions at the universities of Sussex and Southampton in the UK, and at the MIT Artificial Intelligence Lab, USA. He also spent seven years working in industrial technology research, first for Hewlett-Packard Labs and then for Deutsche Bank's Foreign Exchange Complex Risk Group. Since 2005, Dave has been Director of the UK Large-Scale Complex IT Systems Initiative (www.lscits.org). He is a Fellow of the British Computer Society.

Acknowledgements

We thank the UK Engineering and Physical Sciences Research Council (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.

Received: 8 March 2013 Accepted: 26 April 2013

Published: 30 April 2013

References

1. Weinman J Time is money: the value of "on-demand". http://joeweinman.com/Resources/Joe_Weinman_Time_Is_Money.pdf, Accessed 25th April 2013 (2011)
2. Barroso LA, Hölzle U (2009) The datacenter as a computer: an introduction to the design of warehouse-scale machines. *Synthesis Lect Comp Ar* 4:1–108
3. Berl A, Gelenbe E, Di Girolamo M, Giuliani G, De Meer H, Dang MQ, Pentikousis K (2009) Energy-efficient cloud computing. *Comp J* 53:1045–1051
4. Patel CD, Shah AJ (2005) Cost model for planning, development and operation of a data center. *Development* 107:1–36
5. Lee YC, Zomaya AY (2012) Energy efficient utilization of resources in cloud computing systems. *J Supercomput* 60(2):268–280
6. Barroso LA, Holzle U (2007) The case for energy-proportional computing. *Computer* 40:33–37
7. Qureshi A, Weber R, Balakrishnan H, Guttag JV, Maggs B (2009) Cutting the electric bill for internet-scale systems. <http://dspace.mit.edu/handle/1721.1/62585>, Accessed 25th Apr 2013 (2009)
8. Rao L, Liu X, Xie L, Liu W (2010) Minimizing electricity cost: optimization of distributed internet data centers in a multi-electricity-market environment. 2010 Proceedings IEEE INFOCOM:1–9, IEEE
9. Buchbinder N, Jain N, Menache I (2011) Online job-migration for reducing the electricity bill in the cloud. *Networking* 1:172–185, 2011
10. Hull JC (2008) *Fundamentals of Futures and Options Markets*. Prentice Hall, London
11. Bagnolli M, Lipman B (1989) Provision of public goods: fully implementing the core through private contributions. *Rev Econ Stud* 56:583–601
12. Rogers O, Cliff D (2012) Options, forwards and provision-point contracts in improving cloud infrastructure utilisation. *J Cloud Comput: Advances, Systems and Applications* 1:1–22
13. Kiusalaas J (2010) *Numerical Methods in Engineering with Python* (Google eBook). Cambridge University Press, Cambridge
14. Crawley MJ (2007) *The R Book*. Wiley, Oxford
15. Holt CC (2004) Forecasting seasonals and trends by exponentially weighted moving averages. *Int J Forecasting* 20:5–10
16. Taylor JW (2010) Triple seasonal methods for short-term electricity demand forecasting. *Eur J Oper Res* 204:139–152
17. Ben-Yehuda OA, Ben-Yehuda M, Schuster A, Tsafir D (2011) Deconstructing Amazon EC2 spot instance pricing. 2011 IEEE Third Int Conf Cloud Comput Tech Sci:304–311, IEEE
18. Brockwell P, Davis R (2002) *Introduction to time series and forecasting*. Springer, Heidelberg

doi:10.1186/2192-113X-2-10

Cite this article as: Rogers and Cliff: Contributory provision point contracts – a risk-free mechanism for hedging cloud energy costs. *Journal of Cloud Computing: Advances, Systems and Applications* 2013 **2**:10.