

# Dynamic Cloud Resource Reservation via Cloud Brokerage

Wei Wang\*, Di Niu<sup>†</sup>, Baochun Li\*, Ben Liang\*

\* Department of Electrical and Computer Engineering, University of Toronto

<sup>†</sup> Department of Electrical and Computer Engineering, University of Alberta

Email: weiwang@eecg.toronto.edu, dnu@ualberta.ca, bli@eecg.toronto.edu, liang@utoronto.ca

**Abstract**—Infrastructure-as-a-Service clouds offer diverse pricing options, including on-demand and reserved instances with various discounts to attract different cloud users. A practical problem facing cloud users is how to minimize their costs by choosing among different pricing options based on their own demands. In this paper, we propose a new *cloud brokerage service* that reserves a large pool of instances from cloud providers and serves users with price discounts. The broker optimally exploits both pricing benefits of long-term instance reservations and multiplexing gains. We propose dynamic strategies for the broker to make instance reservations with the objective of minimizing its service cost. These strategies leverage dynamic programming and approximate algorithms to rapidly handle large volumes of demand. Our extensive simulations driven by large-scale Google cluster-usage traces have shown that significant price discounts can be realized via the broker.

## I. INTRODUCTION

Infrastructure-as-a-Service (IaaS) cloud enables IT services to elastically scale *computing instances* to match their time-varying computational demands. Thanks to the economies of scale, an IaaS cloud is capable of offering such on-demand computational services at a low cost [1]. Cloud users usually pay for the usage (counted by the number of instance-hours incurred) in a pay-as-you-go model, and are therefore freed from the prohibitive upfront investment on infrastructure, which is usually over-provisioned to accommodate peak demands.

A cloud provider prefers users with predictable and steady demands, which are more friendly to capacity planning. In fact, most cloud providers offer an additional pricing option, referred to as the *reservation option*, to harvest long-term risk-free income. Specifically, this option allows the user to prepay a one-time reservation fee and then to reserve a computing instance for a long period (usually in the order of weeks, months or years), during which the usage is either free or charged under a significant discount [2], [3], [4]. If fully utilized, such a *reserved instance* can easily save its user more than 50% of the expense.

However, whether and how much a user can benefit from the reservation option critically depends on its demand pattern. Due to the prepayment of reservation fees, the cost saving of a reserved instance is realized only when the accumulated instance usage during the reservation period exceeds a certain threshold (varied from 30% to 50% of the reservation period). Unless heavily utilized, the achieved saving is not significant. For this reason, users with sporadic and bursty demands only launch instances on demand.




Fig. 1. The proposed cloud broker. Solid arrows show the direction of instance provisioning; dashed arrows show the direction of money flow.

Unfortunately, on-demand instances are economically inefficient to users, not only because of the higher rates, but also because there is a fundamental limit on how small the billing cycle can be made. For example, Amazon Elastic Compute Cloud (EC2) charges on-demand instances based on running hours. In this case, an instance running for only 10 minutes is billed *as if* it were running for a full hour [2]. Similar limitations are also imposed by other cloud providers, *e.g.*, [3], [4]. Such billing inefficiency becomes more salient when longer billing cycles are employed (*e.g.*, in VPS.NET [5], even a single hour is charged at a daily rate) and user demands are sporadic, with a substantial amount of partial usage.

In general, to what extent a cloud user can enjoy cost savings due to reservation, while avoiding its inefficiency due to coarse-grained billing cycles, is limited by its own demand pattern. A natural question arises: Can we go beyond this limitation to further lower the cost for all cloud users? In particular, can users with arbitrary demand pattern benefit from reservation options while reducing the costs of instance-hours that are not fully utilized?

In this paper, we propose a *cloud brokerage service* to address these challenges. As illustrated in Fig. 1, instead of trading directly with cloud providers, a user will purchase instances from the *cloud broker*, who has reserved a large pool of instances from the cloud providers to serve a major part of the incoming user demand, while accommodating request bursts by launching on-demand instances. From the perspective of users, their behavior resembles launching instances “on demand” provided by the broker, yet at a lower price. Intuitively, the cloud broker leverages the “wholesale” model and the pricing gap between reserved and on-demand instances to reduce the expenses for all the users. More importantly, the broker can optimally coordinate different users to achieve additional cost savings, with the following benefits.




Fig. 2. The broker can time-multiplex partial usage from different users in the same instance-hour. In this case, serving two users only takes one instance-hour, instead of two.

**Better exploiting reservation options:** The broker aggregates the demand from a large number of users for service, smoothing out individual bursts in the aggregated demand curve, which is more stable and suitable for service through reservation. In contrast, individual users usually have bursty and sporadic demands, which are not friendly to the reservation option.

**Reducing wasted cost due to partial usage:** Partial usage of a billing cycle always incurs a full-cycle charge, making users pay for more than what they use. As illustrated in Fig. 2, without the broker, Users 1 and 2 each has to purchase one instance-hour, and pays the hourly rate even if it only uses the hour partially. In contrast, the broker can use a single instance-hour to serve both users by time-multiplexing their usage, reducing the total service cost by one half. Such a benefit can be realized at the broker by scheduling the aggregated user demands to the pooled instances.

**Enjoying volume discounts:** Most IaaS clouds offer significant volume discounts to those who have purchased a large number of instances. For example, Amazon provides 20% or even higher volume discounts in EC2 [2]. Due to the sheer volume of the aggregated demand, the cloud broker can easily qualify for such discounts, which further reduces the cost of serving all the users.

It is through all the mechanisms above that the broker reduces the expenses for cloud users, while turning profits for itself. However, a major challenge in operating such a broker is the decision on how many instances the broker should reserve, how many instances it should launch on demand, and when to reserve, as the demands change dynamically over time. To solve this challenge, we formulate the problem of dynamic instance reservation given user demand data, and derive the optimal reservation strategy via dynamic programming. Unfortunately, such dynamic programming is computationally prohibitive. We propose two efficient approximate algorithms that we prove to offer worst-case cost guarantees. We also propose an effective online algorithm that makes reservation decisions without having access to future demand information.

We conduct large-scale simulations driven by 180 GB of Google cluster usage traces [6] involving over 900 cloud users' workload in a recent month. We empirically evaluate the aggregate and individual cost savings brought forth by the broker, under the proposed reservation strategies. Our results suggest that the broker is the most beneficial for users with medium demand fluctuations, reducing their total expenses by more than 40%. As for general users, 70% of

them receive discounts more than 25%. Such cost savings are more significant in IaaS clouds employing longer reservation periods or longer billing cycles (e.g., [5]).

## II. DYNAMIC RESOURCE RESERVATION

In this section, we briefly review the prevalent cloud pricing schemes and formulate the broker's optimal instance reservation problem to satisfy given demands, with an objective of cost minimization.

### A. Prevalent Cloud Pricing Schemes

Most IaaS clouds provide users with multiple purchasing options, including on-demand instances, reserved instances, and other instance types [2], [3], [4]. *On-demand instances* allow users to pay a fixed rate in every *billing cycle* (e.g., an hour) with no commitment. For example, if the hourly rate of an on-demand instance is  $p$ , an instance that has run for  $n$  hours is charged  $n \cdot p$ . As another purchasing option, a *reserved instance* allows a user to pay a one-time fee to reserve an instance for a certain amount of time, with reservation pricing policies subtly different across cloud providers. In most cases, the cost of a reserved instance is *fixed*. For example, in [3], [4], [5], the cost of a reserved instance is equal to the reservation fee. As another example, in Amazon EC2 [2], the cost of a Heavy Utilization Reserved Instance is billed as a reservation fee plus a heavily discounted hourly rate charged over the entire reservation period, no matter whether the instance is used or not. EC2 also offers other reservation options (e.g., Light/Medium Utilization Reserved Instances), with cost dependent on the actual usage time of the reserved instance. Throughout the paper, we limit our discussions to reservations with fixed costs, which represent the most common cases in IaaS clouds.

### B. The Instance Reservation Problem

The broker asks cloud users to submit their demand estimates over a certain horizon, based on which dynamic reservation decisions are made. Note that even if a user trades directly with cloud providers, it needs to estimate its future demand to decide how many instances to reserve at a particular time. In the case where users are unable to estimate demand at all, we propose an online reservation strategy in Sec. IV-C to make decisions based on history.

Suppose cloud users submit to the broker their demand estimates up to time  $T$  into the future (in terms of billing cycles). The broker aggregates all the demands. Suppose it requires  $d_t$  instances in total to accommodate all the requests at time  $t$ ,  $t = 1, 2, \dots, T$ . The broker makes a decision to reserve  $r_t$  instances at time  $t$ , with  $r_t \geq 0$ . Each reserved instance will be effective from  $t$  to  $t + \tau - 1$ , with  $\tau$  being the *reservation period*.

At time  $t$ , denote the number of reserved instances that remain effective by  $n_t = \sum_{i=t-\tau+1}^t r_i$ . Note that these  $n_t$  reserved instances may not be sufficient to accommodate the aggregate demand  $d_t$ . The broker thus needs to launch




Fig. 3. State illustration. The reservation period is  $\tau = 4$ . All four reservations are highlighted as the shaded area. At stage 2,  $\mathbf{s}_2 = (2, 2, 2, 1)$ .

$(d_t - n_t)^+$  additional on-demand instances at time  $t$ , where we denote by  $X^+ := \max\{0, X\}$ .

Let  $\gamma$  denote the one-time reservation fee for each reserved instance, and  $p$  denote the price of running an on-demand instance per billing cycle. Then, the total cost to accommodate all the demands  $d_1, \dots, d_T$  is

$$\sum_{t=1}^T r_t \gamma + \sum_{t=1}^T (d_t - n_t)^+ p, \quad (1)$$

where the first term is the total cost of reservations and the second is the cost of all on-demand instances. The broker's problem is to make dynamic reservation decisions  $r_1, \dots, r_T$  to minimize its total cost, i.e.,

$$\begin{aligned} \min_{\{r_1, \dots, r_T\}} \quad & \text{cost} = \sum_{t=1}^T r_t \gamma + \sum_{t=1}^T (d_t - n_t)^+ p, \\ \text{s.t.} \quad & n_t = \sum_{i=t-\tau+1}^t r_i, \forall t = 1, \dots, T. \end{aligned} \quad (2)$$

Problem (2) is integer programming. In general, complex combinatorial methods are needed to solve it.

### III. DYNAMIC PROGRAMMING: OPTIMALITY AND LIMITATIONS

In this section, we resort to dynamic programming to characterize the optimal solution to problem (2). Using a set of recursive Bellman equations, the original combinatorial optimization problem can be decomposed into a number of subproblems, each of which can be solved efficiently. However, we also point out that computing such a dynamic programming is practically infeasible, and is highly inefficient to handle a large amount of data.

#### A. Dynamic Programming Formulation

We start by defining stages and states. The decision problem (2) consists of  $T$  stages, each representing a billing cycle. A state at stage  $t$  is denoted by a  $\tau$ -tuple  $\mathbf{s}_t := (x_1^t, \dots, x_{\tau-1}^t)$ , where  $x_i^t$  denotes the number of instances that are reserved no later than  $t$  and remain effective at stage  $t+i$ , for  $i = 1, \dots, \tau-1$ . Here, we use a  $\tau$ -tuple to represent a state because no instance reserved before or at stage  $t$  will remain effective after stage  $t+\tau-1$ . And it is easy to check that  $x_1^t \geq \dots \geq x_{\tau-1}^t$ . For example, in Fig. 3, three instances are reserved at stage 1, 2, and 3, respectively, with a reservation period  $\tau = 4$ . We see that at stage 2,  $\mathbf{s}_2 = (2, 2, 1)$ , where the first element is 2 because two instances are reserved no later than stage 2 and remain effective at stage 3.

We note that given state  $\mathbf{s}_{t-1} := (x_1^{t-1}, \dots, x_{\tau-1}^{t-1})$  at the previous stage, the current state  $\mathbf{s}_t := (x_1^t, \dots, x_{\tau-1}^t)$  is




Fig. 4. Illustration of a transition into state  $\mathbf{s}_t = (x_1^t, \dots, x_{\tau-1}^t)$  from state  $\mathbf{s}_{t-1} = (x_1^{t-1}, \dots, x_{\tau-1}^{t-1})$  according to (3). The shaded area stands for the  $x_{\tau-1}^{t-1}$  reservations made in hour  $t$ .

independent of the states  $\mathbf{s}_{t-2}, \mathbf{s}_{t-3}, \dots$ . In fact,  $\mathbf{s}_t$  can be characterized by  $\mathbf{s}_{t-1}$ , with state transition equations

$$x_i^t = x_{i+1}^{t-1} + x_{\tau-1}^{t-1}, \quad i = 1, \dots, \tau-2. \quad (3)$$

To see the rationale behind (3), let us consider a state  $\mathbf{s}_t$  in Fig. 4. At stage  $t+\tau-1$ , there are  $x_{\tau-1}^t$  reservations that remain effective. Clearly, all these reservations are made at stage  $t$  (because the reservations made before stage  $t$  have all expired at  $t+\tau-1$ ), i.e.,  $r_t = x_{\tau-1}^t$ . These  $x_{\tau-1}^t$  reserved instances (the shaded area in Fig. 4) will add to the effective reservations starting from stage  $t$ . Since  $\mathbf{s}_{t-1} := (x_1^{t-1}, \dots, x_{\tau-1}^{t-1})$ , we know that at stage  $t+i$ , there remain  $x_{i+1}^{t-1}$  effective reservations made before  $t$ ,  $i = 1, \dots, \tau-2$ . Adding the  $x_{\tau-1}^t$  instances reserved at stage  $t$  leads to  $x_i^t = x_{i+1}^{t-1} + x_{\tau-1}^t$  instances that are reserved no later than stage  $t$  and remain effective at stage  $t+1$ .

Define  $V(\mathbf{s}_t)$  as the minimum cost of serving the demands  $d_1, \dots, d_t$  up to stage  $t$ , conditioned on that state  $\mathbf{s}_t$  is reached at stage  $t$ . Then we have the following recursive Bellman equations:

$$V(\mathbf{s}_t) = \min_{\mathbf{s}_{t-1}} \{V(\mathbf{s}_{t-1}) + c(\mathbf{s}_{t-1}, \mathbf{s}_t)\}, \quad t > 0, \quad (4)$$

where the minimization is over all states  $\mathbf{s}_{t-1}$  that can transit to  $\mathbf{s}_t$  following (3). In (4) the minimum cost of reaching  $\mathbf{s}_t$  is given by the minimum cost of reaching a previous state  $\mathbf{s}_{t-1}$  plus a transition cost  $c(\mathbf{s}_{t-1}, \mathbf{s}_t)$ , minimized over all possible previous states  $\mathbf{s}_{t-1}$ . The transition cost is defined as

$$c(\mathbf{s}_{t-1}, \mathbf{s}_t) := \gamma x_{\tau-1}^t + p(d_t - x_1^{t-1} - x_{\tau-1}^{t-1})^+, \quad (5)$$

where the transition from  $\mathbf{s}_{t-1}$  to  $\mathbf{s}_t$  follows (3).

The rationale of (5) is straightforward. As has been noted,  $x_{\tau-1}^t$  instances are reserved at stage  $t$ , and  $x_1^{t-1}$  instances are reserved before  $t$  and remain effective at  $t$ . To accommodate the demand  $d_t$  at stage  $t$ , the broker needs to launch  $(d_t - x_1^{t-1} - x_{\tau-1}^{t-1})^+$  on-demand instances. As a result, the broker pays  $\gamma x_{\tau-1}^t$  to reserve the  $x_{\tau-1}^t$  instances, and  $p(d_t - x_1^{t-1} - x_{\tau-1}^{t-1})^+$  for on-demand instances, leading to a transition cost of (5).

The boundary conditions of (4) are given by

$$V(\mathbf{s}_0) := \gamma x_1^0, \quad \text{for all } \mathbf{s}_0 := (x_1^0, \dots, x_{\tau-1}^0), \quad (6)$$

since an initial state  $\mathbf{s}_0 := (x_1^0, \dots, x_{\tau-1}^0)$  indicates that the broker has already reserved  $x_1^0$  instances at the beginning (time 1) and paid  $\gamma x_1^0$ .




Fig. 5. The Periodic Decisions algorithm, with  $\gamma = \$2.5$  and  $p = \$1$ . The algorithm is (a) optimal when  $T \leq \tau$ , (b) not optimal when  $T > \tau$ .

Through the analysis above, we have converted problem (2) into an equivalent dynamic programming problem defined by (3), (4), (5), and (6).

### B. The Curse of Dimensionality

Although the dynamic programming presented above is optimal, it is computationally infeasible for large data. This is because to derive the minimum cost  $V(s_T)$ , one has to compute  $V(s_t)$  for all states  $s_t$ . Since each state  $s_t$  is defined as a high dimensional tuple  $(x_1^t, \dots, x_\tau^t)$ , there exist exponentially many such states. Therefore, going through all of them results in exponential time complexity. Also, since the computed  $V(s_t)$  has to be stored for every state  $s_t$ , the space complexity is exponential as well. This is known as the *curse of dimensionality* suffered by high-dimensional dynamic programming [7].

A classical method to handle the curse of dimensionality is to use *Approximate Dynamic Programming* (ADP) [7]. ADP estimates the minimum cost at each state first and refines such estimates in an iterative fashion. It is known that ADP converges to the optimal solution if the initial estimates are “optimistic”, i.e., the estimates are even better than the actual optimal solution [7]. However, the more optimistic the initial estimates are, the slower the convergence will be. Through extensive simulations, we find that even by setting some “smart” initial estimations, the convergence speed of ADP is still not satisfactory to handle the large amount of demand data in our problem. Refer to our technical report [8] for a more detailed discussion and simulation results.

## IV. APPROXIMATE ALGORITHMS

To overcome the prohibitive complexity of dynamic programming, in this section, we develop approximate algorithms to solve (2). These algorithms are highly efficient and are proved to have worst-case performance guarantees. Furthermore, we also propose an online reservation strategy which can be applied when future demand data is not available.

### A. A 2-Competitive Heuristic with Short-Term Predictions

First, we present a simple heuristic that relies on short-term demand estimations and in the worst case, incurs twice the minimum cost, given arbitrary demands. We start off by dividing the demands into *levels*. Let  $\bar{d} := \max_t d_t$  be the peak demand. As shown in Fig. 5a, the total demands are divided

### Algorithm 1 Heuristic: Periodic Decisions

1. Segment  $T$  into intervals  $\{I_i\}$ , each with length  $\tau$ .
2. **for all** intervals  $I_i$  **do**
3. Reserve  $l$  instances at the beginning of this interval, such that  $u_l^i \geq \gamma/p > u_{l+1}^i$ , where  $u_l^i := \sum_{t \in I_i} d_t^l$  is the utilization of level  $l$  in interval  $i$ .
4. **end for**

into  $\bar{d} = 5$  levels, with level 1 being the bottom (labeled as “L1” in Fig. 5a) and level  $\bar{d}$  being the top. Define  $d_t^l$  as the demand at time  $t$  in level  $l$ , such that  $d_t^l = 1$  if  $d_t \geq l$ , and  $d_t^l = 0$  otherwise. For example, in Fig. 5a, level 3 has demands only at time 3 and 5 (i.e.,  $d_3^3 = d_5^3 = 1$ ).

We now consider a special case, when all given demands are within a single reservation period, i.e.,  $T \leq \tau$ . In this case, it is sufficient to *make all the reservations at time 1*, since a reservation made anytime will remain effective for the entire horizon  $T$ . The question becomes how many instances to reserve at time 1.

Initially, we consider the first reserved instance that will be used to serve demands in level 1. Define *utilization*  $u_1$  as the number of billing cycles where this reserved instance will be used. It is easy to check  $u_1 = \sum_{t=1}^T d_t^1$ . The use of this reserved instance will be well justified if the reservation fee satisfies  $\gamma \leq pu_1$ ; otherwise, launching it on demand would be more cost efficient.

Next, suppose  $l - 1$  instances are already reserved in the bottom  $l - 1$  levels. We check if an instance should be reserved in level  $l$ . Define *utilization*  $u_l$  as the number of billing cycles where the  $l$ th reserved instance will be used, i.e.,

$$u_l := \sum_{t=1}^T d_t^l, \quad l > 0. \quad (7)$$

For convenience, we let  $u_0 := +\infty$  (for reasons that will be clear later). Again, the broker will adopt the  $l$ th reserved instance only if  $\gamma \leq pu_l$ . Noting that  $u_l$  is non-increasing in  $l$ , we obtain a very simple optimal algorithm: reserve  $l$  instances at time 1, such that  $u_l \geq \gamma/p > u_{l+1}$ .

Fig. 5a shows an example with  $\gamma = \$2.5$  and  $p = \$1$ . To run the algorithm, we first plot the demand curve  $d_t$ . We find  $u_l$  is the intersection area of a horizontal stripe in level  $l$  with the area below  $d_t$ , e.g.,  $u_3 = 2$ , as shown by the shaded area. In this case, the optimal strategy is to reserve 2 instances in the bottom 2 levels, as  $u_2 = 3 > 2.5 = \gamma/p$  while  $u_3 = 2 < \gamma/p$ .

When demands last for more than one reservation period, i.e.,  $T > \tau$ , a natural idea is to extend the algorithm above by letting the broker make periodic decisions. We segment the time axis into intervals, each with the same length as the reservation period  $\tau$ . The broker makes decisions for each  $\tau$ -interval separately, only at the beginning of that interval, by running the algorithm above. This leads to the Periodic Decisions algorithm described in Algorithm 1. Algorithm 1 is highly efficient, requiring only  $O(\bar{d}T)$  time and  $O(T)$  space. Besides, Algorithm 1 only requires demand estimations in the incoming interval (reservation period), and is suitable for users that can only make short-term predictions.

Algorithm 1 is not optimal, as reservations are placed *only* at the beginning of each interval. For example, in Fig. 5b, Algorithm 1 launches all instances on demand, incurring a cost of \$11. However, the optimal strategy is to reserve 2 instances at the beginning of time 6, with a lower total cost of \$10 =  $2.5 \times 2 + 5$ . The following theorem bounds the worst-case performance of Algorithm 1.

**Proposition 1:** Algorithm 1 is *2-competitive*. That is, for any demands, the cost incurred by Algorithm 1 is no more than twice the minimum cost obtained in Sec. III.

*Proof:* Divide time into non-overlapping intervals  $\{I_k\}$ , each of length  $\tau$ . Let  $I_k := [l_k, h_k]$  be the  $k$ th interval from  $l_k = (k-1)\tau + 1$  to  $h_k = k\tau$ . We call a reservation *interval-based* if instances are reserved only at the beginnings of intervals. Let  $r_t^I$  be the reservation decision of the interval-based strategy. Then  $r_t^I = 0, \forall t \neq l_k, k = 1, 2, \dots$

Given an arbitrary strategy, we construct the corresponding interval-based strategy as follows. Whenever the given strategy reserves  $r_t$  instances at time  $t$ , the constructed interval-based strategy also reserves  $r_t$  instances at the beginnings of all  $I_k$  such that  $I_k \cap R_t \neq \emptyset$ , where  $R_t = [t, t + \tau - 1]$  is the time span in which those  $r_t$  instances reserved by the given strategy are effective. For simplicity, let  $l_0 := 0$ . We then have

$$r_t^I = \begin{cases} \sum_{i=l_{k-1}+1}^{h_k} r_i, & t = l_k \text{ for some } k, \\ 0, & \text{o.w.} \end{cases}$$

Denote by  $n_t$  and  $n_t^I$  the numbers of reserved instances that remain effective at time  $t$  for the given strategy and the constructed interval-based strategy, respectively. Clearly,  $n_t^I \geq n_t$ . Let  $c$  and  $c^I$  be the costs of the given strategy and corresponding interval-based strategy, respectively. We have

$$\begin{aligned} c^I &= \sum_{t=1}^T r_t^I \gamma + \sum_{t=1}^T (d_t - n_t^I)^+ p \\ &= \sum_k \sum_{i=l_{k-1}+1}^{h_k} r_i \gamma + \sum_{t=1}^T (d_t - n_t^I)^+ p \\ &\leq 2 \sum_{t=1}^T r_t \gamma + \sum_{t=1}^T (d_t - n_t)^+ p \\ &\leq 2c, \end{aligned} \quad (8)$$

where the first inequality is due to  $n_t^I \geq n_t$  and the last inequality is due to the definition of the cost  $c$  (see (1)). Equation (8) indicates that the cost of the constructed interval-based strategy is at most twice the original strategy.

We make an important observation, that Algorithm 1 incurs the minimum cost among all interval-based strategies, *i.e.*,  $c_1 \leq c^I$  with  $c_1$  being the cost of Algorithm 1. By (8), we have  $c_1 \leq c^I \leq 2c$ , which suggests that Algorithm 1 incurs no more than twice the cost of any strategy. ■

### B. A Greedy Algorithm with Long-Term Predictions

It is possible to improve Algorithm 1 when users can make long-term predictions. One limitation of Algorithm 1 is that in each level, the reservations are made only at the beginnings of the  $\tau$ -intervals. A direct improvement of Algorithm 1 is to allow *arbitrary* reservation time in each level: we still consider from the bottom level up to the top, whereas in each level, we solve an optimal reservation problem using dynamic programming. Clearly, this strategy incurs less resource provisioning

### Algorithm 2 Greedy Reservation Strategy

1. **Initialization:**  $m_t^{\bar{d}} := 0$  for all  $t = 1, \dots, T$ .
2. **for**  $l = \bar{d}$  **down to** 1 **do**
3.   Make optimal reservations in level  $l$  via dynamic programming defined by (9), (10), and (11).
4.   Update  $m_t^{l-1}$ .
5. **end for**

cost than Algorithm 1 in each level due to the relaxation on reservation time. However, such a strategy remains inefficient, since it ignores the dependencies across different levels.

A simple fix can incorporate inter-level dependencies: instead of reserving bottom-up, we start to make reservations in the top level  $\bar{d}$  and proceed *top-down*. Every instance reserved in level  $l$  that is not used at time  $t$  will be passed over to the lower level  $l-1$ , so that it can be used to serve the demand at time  $t$  in that level. We then step to level  $l-1$  and make optimal reservations there, taking into account the “leftover” reserved instances passed over from the upper level. Undoubtedly, the algorithm becomes more efficient, since each level tries to utilize those “leftover” instances reserved in upper levels. Note that this is impossible for a bottom-up approach, where no “leftover” reserved instances can be passed from a lower level up.

In each level of the procedure above, optimal reservations can be efficiently made via dynamic programming. Suppose before we make reservations in level  $l$ ,  $m_t^l$  reserved instances are passed over from upper levels at time  $t$ , all of which can be utilized. Let  $V_l(t)$  be the minimum cost of serving demands  $d_1^l, \dots, d_t^l$  in level  $l$  up to time  $t$ . The Bellman equation is given by

$$V_l(t) = \min\{V_l(t-\tau) + \gamma, V_l(t-1) + c_l(t)\}, \quad (9)$$

which chooses the minimum between two options. The *first* is to serve the demand  $d_t^l$  with an instance reserved in the current level  $l$ . The best strategy is to optimally serve demands up to time  $t-\tau$  and reserve an instance at time  $t-\tau+1$ , incurring a cost of  $V_l(t-\tau) + \gamma$ . The *second* option is to optimally satisfy demands up to  $t-1$  and serve the demand at  $t$  using an on-demand instance, with a cost  $c_l(t)$  given by

$$c_l(t) = \begin{cases} p, & \text{if } d_t^l = 1 \text{ and } m_t^l = 0, \\ 0, & \text{otherwise,} \end{cases} \quad (10)$$

which means that we pay  $p$  to launch an on-demand instance only if there is demand at  $t$ , *i.e.*,  $d_t^l = 1$ , yet no reserved instance is leftover from upper levels to use at  $t$ , *i.e.*,  $m_t^l = 0$ . The boundary conditions are clearly

$$V_l(t) = 0, \quad t \leq 0. \quad (11)$$

After reservations have been made in level  $l$ , we update  $m_t^{l-1}$ , the number of reserved instances to be passed to level  $l-1$  at time  $t$  as follows:  $m_t^{l-1} := m_t^l + 1$ , if an instance is reserved in level  $l$  but is not used at time  $t$ ;  $m_t^{l-1} := m_t^l - 1$ , if demand  $d_t^l$  is served using a reserved instance passed over from upper levels; and  $m_t^{l-1} = m_t^l$ , otherwise.

---

**Algorithm 3** Online Reservation Made at Time  $t$ 

---

1. Let  $g_i = (d_i - n_i)^+$  for all  $i = t - \tau + 1, \dots, t$ .
  2. Run Algorithm 1 with  $g_{t-\tau+1}, \dots, g_t$  as the input demands. Let  $x$  be its output.
  3. Reserve  $r_t = x$  instances at time  $t$ .
  4. Update  $n_i = n_i + r_t$  for all  $i = t - \tau + 1, \dots, t + \tau - 1$ .
- 

Algorithm 2 summarizes the greedy reservation strategy above, which has time complexity  $O(dT)$  and space complexity  $O(T)$ . As has been analyzed earlier in this subsection, in each level, Algorithm 2 is more cost efficient than Algorithm 1, leading to the following proposition:

**Proposition 2:** Algorithm 2 incurs a cost no more than Algorithm 1, and is thus 2-competitive.

### C. An Online Reservation Strategy

Previous algorithms apply to users who predict their future demands. For those who cannot, we propose an *online strategy* that reserves instances based only on history, without accessing to any future information. Recall that  $n_t$  is the number of reserved instances that remain effective at time  $t$ . Our online strategy makes a reservation decision  $r_t$  at time  $t$  based on the historical information  $g_{t-\tau+1}, \dots, g_t$  in the past reservation period, where  $g_i := (d_i - n_i)^+$  for all  $i$ . We call  $g_i$  the *reservation gap* between demand  $d_i$  and the number of reservations  $n_i$  that remain effective at time  $i$ . Clearly, all these  $g_i$  reservation gaps have been filled by launching on-demand instances at time  $i$ .

Now at time  $t$ , we review the reservation decision made at time  $t - \tau + 1$  (we set  $r_i = 0$  for all  $i \leq 0$ ) and calculate how many *more* instances we should have reserved at time  $t - \tau + 1$ , if we knew that we would have to launch  $g_i$  instances on demand at time  $i = t - \tau + 1, \dots, t$ . These “should-have-reserved” instances can be computed by Algorithm 1, with outstanding gaps  $g_{t-\tau+1}, \dots, g_t$  as the input demands. Suppose the found value is  $x$ , we then reserve  $r_t = x$  instances at the current time  $t$ . In the meantime, we update the reservation history *as if* we had reserved  $x$  additional instances at time  $t - \tau + 1$ , by setting  $n_i := n_i + x$  for all  $i = t - \tau + 1, \dots, t$ , which will be used in computing the next decision  $r_{t+1}$ . Algorithm 3 details this process. Note that initially (at time 0), we set  $d_i = n_i = 0$  for all  $i$ .

The computational complexity of the online reservation strategy above is the same as Algorithm 1 at every time  $t$ .

## V. PERFORMANCE EVALUATION

In this section, we conduct simulations driven by a large volume of real-world traces to evaluate the performance of the proposed brokerage service and reservation strategies, with an extensive range of scenarios.

### A. Dataset Description and Preprocessing

Workload traces in public clouds are often confidential: no IaaS cloud has released its usage data so far. For this reason, we use Google cluster-usage traces that were recently released




Fig. 6. The demand curves of three typical users.




Fig. 7. Demand statistics and the division of users into 3 groups according to demand fluctuation level.

[6] in our evaluation. Although Google cluster is not a public IaaS cloud, its usage traces reflect the computing demands of Google engineers and services, which can represent demands of public cloud users to some degree. The dataset contains 180 GB of resource demand/usage information of 933 users over 29 days in May 2011, on a cluster of 12,583 physical machines. In the Google traces, a user submits work in the form of *jobs*. A job consists of several *tasks*, each of which has a set of resource requirements on CPU, disk, memory, etc.

**Instance Scheduling.** We take such a dataset as input, and ask the question: How many computing instances would each user require if she were to run the same workload in a public IaaS cloud? It is worth noting that in Google cluster, tasks of different users may be scheduled onto the same machine, whereas in IaaS clouds each user will run tasks only on her own computing instances.

Therefore, we reschedule the tasks of each user onto instances that are exclusively used by this user. We set the instances to have the same computing capacity as Google cluster machines (most Google cluster machines are of the same computing capability, with 93% having the same CPU cycles), which enables us to accurately estimate the task run time by learning from the original traces.

For each user, we use a simple algorithm to schedule her tasks onto available instances that have sufficient resources to accommodate their resource requirements. Tasks that cannot



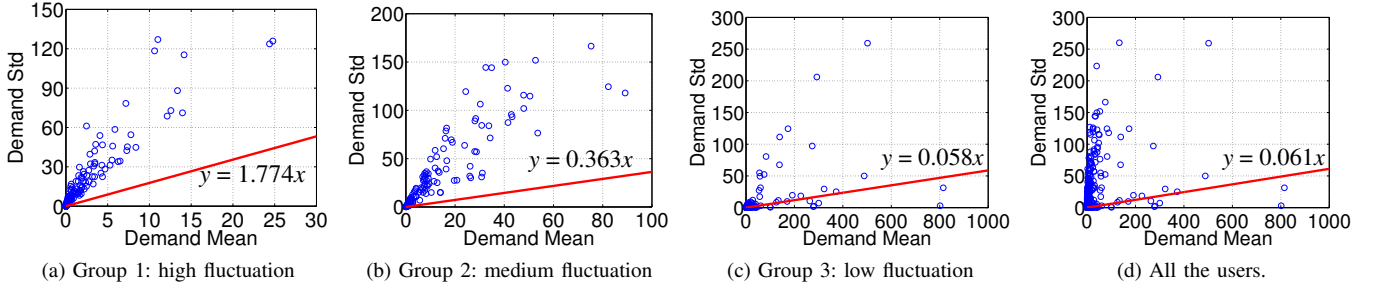


Fig. 8. Aggregation suppresses the demand fluctuation of individual users. Each circle represents a user. The line indicates the demand fluctuation level (the ratio between the demand standard deviation and mean) in the aggregate demand curve.




Fig. 9. Aggregation reduces the wasted instance-hours due to partial usage.

share the same machine (*e.g.*, tasks of MapReduce) are scheduled to different instances. (For simplicity, we ignore other complicated task placement constraints such as on OS versions and machine types.) Whenever the capacity of available instances is reached, a new instance will be launched. In the end, we obtain a demand curve for each user, indicating how many instances the user requires in each hour. Fig. 6 illustrates the demand curves of three typical users in the first 200 hours.

**Pricing.** Unless explicitly mentioned, we set the on-demand hourly rate to \$0.08, the same as Amazon EC2 small instances [2]. Since our data only spans one month, we assume each reservation is effective for one week, with a *full-usage discount* of 50%: the reservation fee is equal to running an on-demand instance for half a reservation period, which is a general pricing policy in most IaaS clouds [2], [3], [4], [5].

**Group Division.** To further understand the demand statistics of users, we calculate the demand mean and standard deviation for each user and plot the results in Fig. 7. As has been mentioned, to what extent a user can benefit from reservations critically depends on its demand pattern: the more fluctuated it is, the less is the benefit from using reserved instances. We hence classify all 933 users into the following three groups by the *demand fluctuation level* measured as the ratio between the demand standard deviation and mean:

**Group 1 (High Fluctuation):** Users in this group have a demand fluctuation level no smaller than 5. A typical user’s demand is shown in the top graph of Fig. 6. There are 271 users in this group, represented by “o” in Fig. 7. These users have small demands, with a mean less than 30 instances.

**Group 2 (Medium Fluctuation):** Users in this group have a demand fluctuation level between 1 and 5. A typical user’s demand is shown in the middle graph of Fig. 6. There are 286

users in this group, represented by “x” in Fig. 7. These users demand a medium amount of instances, with a mean less than 100.

**Group 3 (Low Fluctuation):** Users in this group have a demand fluctuation level less than 1, represented by “+” in Fig. 7. A typical user’s demand is shown in the bottom graph of Fig. 6. Almost all big users with a mean demand greater than 100 belong to this group.

Our evaluations are carried out for each group. We start to quantify to what extent the aggregation smooths out demand bursts of individual users. Fig. 8 presents the results, with “o” being the statistics of individual users and the line representing the fluctuation level of the aggregated demand. We see from Fig. 8a and 8b that aggregating bursty users (*i.e.*, users in Group 1 and 2) results in a steadier curve, with a fluctuation level much smaller than that of any individual user. For users that already have steady demands, aggregation does not reduce fluctuation too much (see Fig. 8c). In addition, Fig. 8d presents the result of aggregating all the users. In all cases, the aggregated demand is stabler and more suitable for service via reserved instances.

Another benefit of demand aggregation is that it reduces the wasted instance-hours incurred by partial usage. To see this, for each user, we count the total time during which it is billed but does not run any workload, when this user purchases directly from the cloud. In each group, we do the same count for the aggregate demand and compare it with the sum of the wasted instance-hours of all users in that group. Fig. 9 plots the results. As expected, we observe a reduction of wasted instance-hours in all four cases. Interestingly, the waste reduction is the most significant for users with medium fluctuation, instead of highly fluctuated users. This is because we do not have a sufficient amount of highly fluctuated demands to aggregate.

## B. Aggregate Cost Savings

We now evaluate the aggregate cost savings offered by the broker, under three different reservation strategies, namely, the Heuristic (Algorithm 1), Greedy (Algorithm 2) and Online (Algorithm 3). Assuming a specific strategy is adopted by *both users and the broker*, we compare the total service cost if users are using the broker with the sum of costs if users trade with the provider. Fig. 11 shows such comparisons in each user




Fig. 11. Aggregate service costs with and without broker in different user groups.




Fig. 10. Aggregate cost savings in different user groups due to the brokerage service.

group, while Fig. 10 shows the percentage of cost savings due to the use of a broker.

From Fig. 10, we see that the broker can bring a cost saving of close to 15% when it aggregates all the user demands. In terms of absolute values, the saving is over \$100K, as shown in Fig. 11d. However, the broker's benefit is different in different user groups: cost saving is the highest for users with medium demand fluctuation (40%), and the lowest for users with low demand fluctuation (5%). This is because when user demands are steady, they are heavily relying on reserved instances, regardless of whether they use the brokerage service or not. The broker thus brings little benefit, as shown in Fig. 11c. In contrast, for fluctuated demands, as shown in Fig. 11b, the broker can smooth out the demand curve through aggregation, better exploiting discounts of reserved instances. However, when users are *highly* fluctuated with bursty demands, as shown in Fig. 11a, even the aggregate demand curve is not smooth enough: these users can only leverage a limited amount of reserved instances, leading to less reservation benefit than for users with medium fluctuation. However, there is still 15 – 20% cost saving, partly due to aggregation and the reduction of partial usage.

We now compare the costs of different reservation strategies. Fig. 11 verifies that Greedy is the best strategy while Online is inferior due to the lack of future knowledge. However, as shown in Fig. 11a, the three strategies are similar for highly fluctuated users, because for bursty demands, on-demand instances are mainly used (especially without broker) and reservation strategies become less critical.




Fig. 12. CDF of price discounts for individual users due to the brokerage service, under different algorithms.

### C. Individual Cost Savings

We next evaluate the price discount each individual user can enjoy from the brokerage service. We consider a straight-forward usage-based pricing scheme adopted by the broker. That is, for each user, the broker calculates the area under its demand curve to find out the instance-hours it has used. The broker then lets users share the aggregate cost in proportion to their instance-hours. In Fig. 12, we plot the CDF of price discounts of individual users due to using the broker. In Fig. 13, we plot the costs with and without the broker for each individual user (represented by a circle), under Greedy strategy, where such costs are the same if the circle is on the straight line  $y = x$ . We do not plot for Group 3 (low fluctuation) because the benefit of broker is less significant. In this sense, users in Group 3 has less motivation to use the broker. Furthermore, we omit Group 1 (high fluctuation) because all their cost saving percentages are observed to be the same as the aggregate saving percentage. The reason is that with highly bursty demands, users in Group 1 will mainly use on-demand instances without the broker, leading to bills proportional to their usage. If these users choose to use the broker, their costs are also proportional to their usage. Therefore, the individual saving percentages are essentially the same as the aggregate saving percentage.

From Fig. 12a, we see that over 70% of users in Group 2 save more than 30%, while in Fig. 12b, we see that the broker can bring more than 25% price discounts to 70% of users if all users are aggregated. Several interesting phenomena are noted from Fig. 12 and Fig. 13. *First*, there is an upper limit on the price discount a user can get under Greedy, which is about 50%. *Second*, with Online, a majority (around 40 –






Fig. 13. Cost without the broker vs. with the broker for individual users, using Greedy strategy. Each circle is a user.

50%) of users receive a discount of around 30%. *Third*, when the broker charges users based on usage, very few users (less than 5%) do not receive discounts (with price discount below 0 or circles above the straight line in Fig. 13). Since these users only contribute to a very small portion of the entire demand (around 3%), the broker can easily guarantee to charge them at most the same price as charged by cloud providers, by compensating them with a portion of the profit gained from service cost savings.

It is worth noting that the usage-based billing above is only one of many possible pricing policies that the broker can use. We adopt it here because it is easy to implement and understand. Although it may cause the problem of compensating overcharged users as mentioned above, it is not typically an issue for the Google dataset. We note that more complicated pricing policies, such as charging based on users' Shapley value [9], can resolve this problem with guaranteed discounts for everyone. The discussion of these policies is out of the scope of this paper. As long as the cost saving is achieved by the broker, there are rich methods to effectively share the benefits among all participants (see Ch. 15 in [10]).

#### D. Reservation Periods and Billing Cycles

We now quantify the impact of other factors that may affect the broker's benefit. The first we consider is the reservation period  $\tau$ . In practice, different reservation periods are defined in different IaaS clouds, ranging from a month to years. To see how this affects the cost saving benefits, we fix the hourly on-demand rate, and try different reservation periods with 50% full-usage discount (*i.e.*, the reservation fee is equal to running on-demand instances for half of the reservation period). The results are plotted in Fig. 14. We observe that, in general, the longer the reservation period, the more significant the cost saving achieved by the broker. It is worth noticing that the broker offers very limited cost savings when there is no reserved instance offered in the IaaS cloud. In this case, the cost saving is only due to the reduction of partial usage.

The second factor that we take into account is how the length of billing cycle affects the cost saving. To see this, we change the billing cycle from an hour to a day, which is the case in VPS.NET [5]. We set the daily on-demand rate to 24 times the original hourly rate (*i.e.*,  $24 \times \$0.08 = \$1.92$ ). The full-usage reservation discount remains 50% (VPS.NET offers 41% full-usage reservation discount, though). Fig. 15a




Fig. 14. The aggregate cost saving as the reservation period varies in different user groups, using Greedy strategy, where "None" means no reservation is available.




Fig. 15. Cost savings with a daily billing cycle under the Greedy strategy.

and Fig. 15b present the simulation results using the Greedy strategy. As compared to the case of hourly billing cycle (Fig. 10 and 12b), we observe a significant cost saving improvement here. Intuitively, adopting a larger billing cycle results in more wasted partial usage, leading to more salient advantages of using the broker.

#### E. Discussions and Other Practical Issues

Let us further discuss several potential benefits of the broker that have not yet been investigated in the evaluation. First, by taking advantage of *volume discounts*, the cost of instance reservations would be significantly further reduced. As mentioned in the introduction, in practice, most IaaS clouds offer heavy volume discounts to large users. For example, in Amazon EC2, such volume discounts offer an additional 20% off on instance reservations [2]. Due to the sheer volume of the aggregated demand, the broker can easily qualify for these discounts.

Second, in addition to savings on the expenses of running instances, the broker can also help lower the costs of other cloud resources such as storage, data transfer, and bandwidth. Since all these prices are *sub-additive* [2], the cost of provisioning aggregated resources is much cheaper than the total cost of purchasing them individually from the cloud.

Third, thanks to the cost reduction, the broker lowers the bar of cloud adoption, which may attract more users and bring forth more demands to cloud providers. Since most of these demands are served with reserved instances, they are more friendly to capacity planning while generating long-term risk-free revenue for providers.

On the other hand, we also note several practical issues in

operating the brokerage service. First, the savings from partial usage reduction are conditioned on the pricing details of a specific cloud. It is worth noting that time-multiplexing users on an *on-demand instance* in EC2 will not lead to cost savings. This is because in EC2, stopping a user on an on-demand instance terminates a billing cycle, while loading a new user onto it opens a new one [2]. However, this is generally not an issue for other cloud providers such as ElasticHosts [3] or reserved instances with a fixed cost (e.g., EC2 Heavy Utilization Reserved Instances). In fact, even without time-multiplexing on-demand instances, our simulation indicates that the total cost savings will only decrease less than 10%.

Second, in reality a user may only have rough knowledge of its future demands, so the broker's demand estimate may not be accurate. In this case, they can still benefit from a broker that uses the online strategy, which does not rely on future information. After all, these users will face exactly the same situation when purchasing directly from the cloud.

Finally, in our simulation, we consider the case that the broker rewards all cost savings to users as price discounts. In reality, the broker can turn a profit by taking a portion of the savings as profit or through a commission.

## VI. RELATED WORK

Some existing works discussed how to leverage prevalent cloud pricing options to reduce instance running costs for an *individual user*. For example, Zhao *et al.* [11] proposed resource rental planning with EC2 Spot Price predictions to reduce the operational cost of cloud applications. Hong *et al.* [12] designed an instance purchasing strategy to reduce the "margin cost" of over-provisioning. They also presented a strategy to combine the use of on-demand and reserved instances [12], which is essentially a special case of Algorithm 1 when all demands are given in one reservation period. Vermeersch [13] implemented a software prototype that dynamically retrieves instances from Amazon EC2 based on the user workload. All these works make instance purchasing decisions for an individual user.

IaaS cloud brokers have recently emerged as *intermediators* connecting buyers and sellers of computing resources. For example, SpotCloud [14] offers a "clearinghouse" in which users can buy and sell unused cloud computing capacity. Buyya *et al.* [15] discussed the engineering aspects of using brokerage to interconnect clouds into a global cloud market. Song *et al.* [16] proposed a broker that predicts EC2 Spot Price, bids for Spot Instances, and uses them to serve cloud users. Unlike existing brokerage services that accommodate individual user requests *separately*, our broker serves the aggregated demands by leveraging instance multiplexing gains and instance reservations, and is a general framework not limited to a specific cloud.

The idea of resource multiplexing has also been extensively studied, though none of them relates to computing instance provisioning. Stanojevic *et al.* [17] made use of bandwidth burstable billing and proposed a cooperative framework in which multiple ISPs jointly purchase IP transit in bulk to

reduce individual costs. Urgaonkar *et al.* [18] empirically evaluated the idea of statistical multiplexing and resource overbooking in a shared hosting platform. Compared with these applications, exploiting multiplexing gains in cloud instance provisioning poses new challenges, mainly due to the newly emerged complex cloud pricing options. It remains nontrivial to design instance purchasing strategies that can optimally combine different pricing options to reduce cloud usage costs.

## VII. CONCLUDING REMARKS

In this paper, we propose a smart cloud brokerage service that serves cloud user demands with a large pool of computing instances that are either dynamically reserved or launched on demand from IaaS clouds. By taking advantage of instance multiplexing gains as well as the price gap between on-demand and reserved instances, the broker benefits cloud users with heavy discounts while gaining profits from the achieved cost savings. To optimally exploit the price benefits of reserved instances, we propose a set of dynamic strategies to decide when and how many instances to reserve, with provable performance guarantees. Large-scale simulations driven by real-world cloud usage traces quantitatively suggest that significant cost savings can be expected from using the proposed cloud brokerage service.

## REFERENCES

- [1] M. Armbrust, A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A view of cloud computing," *Communications of the ACM*, 2010.
- [2] Amazon EC2 Pricing, <http://aws.amazon.com/ec2/pricing/>.
- [3] ElasticHosts, <http://www.elastichosts.com/>.
- [4] GoGrid Cloud Hosting, <http://www.gogrid.com>.
- [5] VPS.NET, <http://vps.net>.
- [6] C. Reiss, J. Wilkes, and J. Hellerstein, "Google Cluster-Usage Traces," <http://code.google.com/p/googleclusterdata/>.
- [7] W. Powell, *Approximate Dynamic Programming: Solving the curses of dimensionality*. John Wiley and Sons, 2011.
- [8] W. Wang, D. Niu, B. Li, and B. Liang, "Dynamic cloud resource reservation via cloud brokerage," University of Toronto, Tech. Rep., 2012. [Online]. Available: <http://iqua.ece.toronto.edu/~weiwang/papers/broker.pdf>
- [9] A. E. Roth, Ed., *The Shapley Value, Essays in Honor of Lloyd S. Shapley*. Cambridge University Press, 1988.
- [10] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [11] H. Zhao, M. Pan, X. Liu, X. Li, and Y. Fang, "Optimal resource rental planning for elastic applications in cloud market," in *Proc. IEEE IPDPS*, 2012.
- [12] Y. Hong, M. Thottethodi, and J. Xue, "Dynamic server provisioning to minimize cost in an IaaS cloud," in *Proc. ACM SIGMETRICS*, 2011.
- [13] K. Vermeersch, "A broker for cost-efficient qos aware resource allocation in EC2," Master's thesis, University of Antwerp, 2011.
- [14] SpotCloud, <http://spotcloud.com/>.
- [15] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging it platforms: Vision, hype, and reality for delivering computing as the 5th utility," *Future Generation Computer Systems*, vol. 25, no. 6, pp. 599–619, 2009.
- [16] Y. Song, M. Zafer, and K.-W. Lee, "Optimal bidding in spot instance market," in *Proc. IEEE INFOCOM*, 2012.
- [17] R. Stanojevic, I. Castro, and S. Gorinsky, "CIPT: Using tuangou to reduce IP transit costs," in *Proc. ACM CoNEXT*, 2011.
- [18] B. Urgaonkar, P. Shenoy, and T. Roscoe, "Resource overbooking and application profiling in shared hosting platforms," in *Proc. USENIX OSDI*, 2002.