# A Theory of Cloud Bandwidth Pricing for Video-on-Demand Providers

Di Niu, Chen Feng, Baochun Li
Department of Electrical and Computer Engineering
University of Toronto

***Abstract*—Current-generation cloud computing is offered with usage-based pricing, with no bandwidth capacity guarantees, which is however unappealing to bandwidth-intensive applications such as video-on-demand (VoD). We consider a new type of service where VoD providers, such as Netflix and Hulu, make reservations for bandwidth guarantees from the cloud at negotiable prices to support continuous media streaming. We argue that it is beneficial to multiplex such bandwidth reservations in the market using a profit-making *broker* while controlling the performance risks. We ask the question—in such a market, how much should each VoD provider pay for bandwidth reservation? We prove that the market has a unique Nash equilibrium where the bandwidth reservation price for a VoD provider critically depends on its demand correlation to the market. Real-world traces verify that our theory can significantly lower the market price for cloud bandwidth reservation.**

## I. Introduction

Cloud computing is changing the way that user-oriented multimedia applications such as video-on-demand (VoD) operate their businesses. Traditionally, VoD providers invest in commodity servers as business grows and acquire bandwidth deals from Internet Service Providers (ISPs). Nowadays, they can be freed from hardware maintenance and network administration by using computing and network resources in the public cloud. One particular example is that Netflix, a major VoD provider in North America, moved its streaming servers, encoding software, search engines, and huge data stores to Amazon Web Services (AWS) in 2010 [1].

Since video delivery has stringent rate requirements, a VoD provider wants to make sure that sufficient server bandwidth is provisioned to sustain continuous media delivery to end users. Thanks to recent advances in datacenter engineering, it is becoming technically feasible to offer bandwidth guarantees for egress traffic from virtual machines (VMs) [2], [3]. We believe that, when bandwidth-intensive applications (such as VoD and video gaming) are on board to use services from cloud providers (such as Amazon), there will be a *market* between the application providers and cloud providers. The commodities to be traded in such a market consist of *bandwidth reservations*, so that the application performance can be guaranteed.

The idea of reserving cloud bandwidth has posed several new challenges to resource management and pricing in the cloud. *First of all*, as demand varies dynamically, it is difficult for a VoD provider to estimate how much bandwidth to reserve at a certain point. Whereas under-provision causes performance issues, the common practice of over-provision incurs high service cost at the cloud and eventually leads to higher prices for VoD providers. *Second*, current cloud providers charge applications for bandwidth usage in a pay-as-you-go model based on the number of bytes transferred [4], which is however insufficient as a model to price *bandwidth guarantees*. Such a pricing problem is further complicated by the fact that cloud providers usually conduct resource over-subscription: the demands of multiple applications may be statistically multiplexed to save bandwidth reservation. With multiplexing and resource sharing among applications, it is even harder to fairly price bandwidth reservations for individual applications.

To address the above challenges, we propose an economically viable *cloud broker* that sells bandwidth guarantees to VoD providers *individually* under a certain pricing policy, while *jointly* booking bandwidth for them from the clouds to save reservation cost and maximize profit. With a reduced service cost, VoD providers may also expect a lower price for bandwidth reservation. Specifically, the broker operates in the following process. *First*, it predicts the demand statistics of VoD providers in the near future based on demand history available from cloud monitoring services, e.g., Amazon CloudWatch provides free resource monitoring to AWS customers at a 5-minute frequency [5] as of 2011. *Second*, the broker mixes demands of different VoD providers based on anti-correlation, and directs the mixed demands to multiple cloud providers for service. The broker repeats the above process periodically to adapt to demand changes.

Our study is inspired by the question—what the broker pricing policy should be and will be in such a broker-assisted market? The broker pricing policy is crucial because it not only determines the broker profit and bills for VoD providers, but also affects the way that the broker accommodates and directs demands, indirectly controlling resource allocation among multiple clouds. We start with observing the conflicting objectives of different entities in the market.

The first potential conflict lies between the broker and cloud providers. To enhance cloud resource efficiency, the load direction should be such that demands are optimally mixed with the minimum bandwidth resource reserved. However, depending on the pricing policy, a selfish broker may direct loads in a different way to maximize its profit, or even deny demand if serving it does not bring profit. To resolve this conflict, our first contribution is to characterize the entire region of good pricing policies under which a profit-driven broker behaves exactly like an altruistic social planner that directs loads for workload consolidation across multiple clouds. Good pricing can subsequently be enforced in a *controlled market* to promote market healthiness.

Another fundamental conflict lies between the broker profit and the cost of VoD providers. As the broker saves service cost through multiplexing, VoD providers may expect lower prices for bandwidth reservation. What is the maximum discount a VoD provider can enjoy? To answer this question, we study a *free market* where each VoD provider can bargain with the broker to negotiate the price it has to pay. However, a VoD provider cannot pay too little since in that case the broker will deny demand for profit concerns, hurting the VoD provider's utility. Using game theory, we prove that the free market will converge to a *unique Nash equilibrium*, where the equilibrium price for a VoD provider critically depends on its demand expectation, burstiness as well as correlation to the market. More interestingly, our study has discovered the *risk neutralizers* that may earn bonus for having demand negatively correlated to the market. We also point out the special meaning of equilibrium pricing in the good pricing region derived for the controlled market.

We have conducted online bandwidth reservation simulations driven by the workload traces of more than 200 video channels from an operational VoD system. The simulator incorporates demand prediction, resource reservation and pricing as interdependent components. We find that with the help of a broker, our theory can lower the market price for cloud bandwidth reservations by around 50% on average, and save cloud resource occupation by over 30%.

The remainder of this paper is organized as follows. We describe our system model in Sec. II and outline main contributions in Sec. III. We study bandwidth reservation pricing in controlled and free markets in Sec. IV and V, respectively, and present simulation results in Sec. VI. Sec. VII reviews related work, and Sec. VIII concludes the paper.

## II. System Model

We study a market of multiple public cloud providers, VoD providers as cloud *tenants*, and a broker. The broker sells *probabilistic bandwidth guarantees* to tenants while reserving the actual bandwidth from the clouds for the tenants. The system operates on a short-term basis. At the beginning of a short-term period, the broker analyzes each tenant's demand history, available from cloud monitoring services, to predict its expected demand as well as demand variation in the following period. Our previous work has shown that VoD demands are highly predictable even at a fine granularity of 10-minute intervals [6], [7]. While the broker sells guarantees to tenants individually, it jointly reserves bandwidth from multiple clouds for the mixed demand, exploiting statistical multiplexing to save reservation cost. As the aggregate service cost is reduced, the tenants may expect to pay a lower price. We start with describing different entities in the market in more detail.

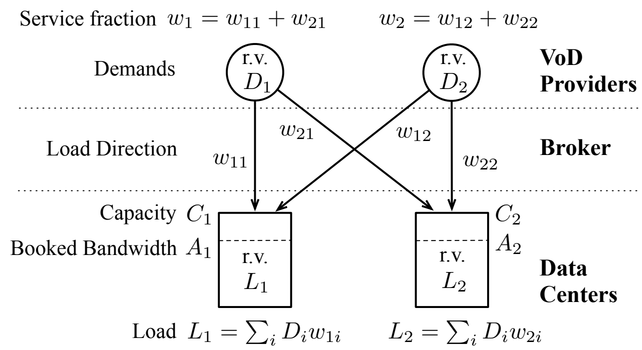


Fig. 1. A system of 2 cloud providers and 2 VoD providers. Random variables are labeled with "r.v."

### A. Cloud Providers, VoD Providers and the Broker

**Cloud Providers.** We assume there are $S$ cloud providers, each with an outgoing bandwidth capacity of $C_s$ available for reservation, $s = 1, \ldots, S$. Let $C_{\mathsf{sum}} := \sum_s C_s$ be the aggregate bandwidth capacity of all the cloud providers. Throughout this paper, we assume that $C_{\mathsf{sum}}$ is sufficiently large to accommodate the total demand in the market. This assumption is justified by the "illusion of infinite capacity" [4] and the fact that the cost of high-end routers is dropping more quickly than before.

**VoD Providers.** We assume that $N$ VoD providers are present as cloud tenants, hoping to receive bandwidth guarantees. Throughout the paper, we will use the terms *VoD provider* and *tenant* interchangeably. Bandwidth reservations need to be varied dynamically and periodically as demand changes. We consider one such *reservation period*. Suppose that in this period, tenant $i$'s bandwidth demand (due to video requests from VoD users) is a random variable $D_i$ with mean $\mu_i$ and variance $\sigma_i^2$.

Let $\boldsymbol{\mu} = [\mu_1, \ldots, \mu_N]^\mathsf{T}$ and $\boldsymbol{\sigma} = [\sigma_1, \ldots, \sigma_N]^\mathsf{T}$. Note that demands of different tenants may be intercorrelated due to the correlation between video genres, viewer preferences and video release times. Denote $\rho_{ij}$ the correlation coefficient of $D_i$ and $D_j$, with $\rho_{ii} \equiv 1$. Let $\mathbf{\Sigma} = [\sigma_{ij}]$ be the $N \times N$ symmetric demand covariance matrix, with $\sigma_{ii} = \sigma_i^2$ and $\sigma_{ij} = \rho_{ij}\sigma_i\sigma_j$ for $i \neq j$.

We define $\epsilon$ as the *risk factor*. To satisfy its demand $D_i$ with a high probability $1 - \epsilon$, tenant $i$ needs to reserve bandwidth $B_i := f_\epsilon(D_i)$, where $f_\epsilon(D_i)$ denotes $1 - \epsilon$ percentile of the distribution of $D_i$, so that $\Pr(D_i > B_i) = \epsilon$.

**The Broker** jointly books bandwidth for the tenants from multiple clouds based on demand estimates $\{\mu_i\}$, $\{\sigma_i^2\}$. The broker aims to mix demand based on anti-correlation to save bandwidth reservation. As shown in Fig. 1, such a decision essentially relates to a *load direction matrix* $\mathbf{W} = [w_{si}]_{S \times N}$, where $w_{si}$ represents the portion of tenant $i$'s demand $D_i$ directed to and served by cloud provider $s$, with $0 \leq w_{si} \leq 1$.

On the cloud side, given $[w_{si}]$, the bandwidth load imposed on cloud provider $s$ is a random variable $L_s = \sum_i w_{si} D_i$. Similar to $B_i := f_\epsilon(D_i)$, we define

$$A_s := f_\epsilon(L_s), \quad A_s \leq C_s, \tag{1}$$

as the bandwidth that the broker needs to reserve from cloud provider $s$ to satisfy load $L_s$ with a high probability $1-\epsilon$. Clearly, the reserved bandwidth $A_s$ must not exceed the capacity of cloud $s$, i.e., $A_s \le C_s$, or equivalently,

$$\Pr(L_s > C_s) \le \epsilon, \quad \forall s. \tag{2}$$

On the tenant side, given $[w_{si}]$, the total *guaranteed portion* of $D_i$ served by all the clouds is

$$w_i = \sum_s w_{si} \le 1. \tag{3}$$

The broker takes demand estimates $\boldsymbol{\mu}$, $\boldsymbol{\Sigma}$ and calculates a load direction matrix $\mathbf{W}^* = [w^*_{si}]$ to achieve a certain goal (e.g., optimal workload consolidation, or maximum broker profit), subject to cloud capacity constraints $\{C_s\}$. The output $\mathbf{W}^*$ corresponds to a way that $N$ inter-correlated random demands $D_1, \ldots, D_N$ are packed into $S$ clouds. Let $w^*_i = \sum_s w^*_{si}$. Note that depending on its objective, the broker does not always generate a guaranteed portion of $w^*_i = 100\%$ for each tenant.

### B. Pricing the Bandwidth Reservation

The broker charges tenant $i$ a fee of $P_i(w_i)$ for accommodating $w_i$ portion of its demand $D_i$ according to some pricing strategy $P_i(\cdot)$. Without loss of generality, we assume each cloud provider charges a fee of \$1 for every unit bandwidth reserved in each short-term period. Hence, the broker needs to pay $\$1 \cdot \sum_s A_s$ to the cloud providers for booking a total amount of $\sum_s A_s$ bandwidth. Therefore, under a given $\mathbf{W}$, the *broker profit* is

$$R(\mathbf{W}) = \sum_i P_i(w_i) - \sum_s A_s. \tag{4}$$

We formally define the pricing strategy and policy as follows:

**Definition 1:** A *bandwidth pricing strategy* $P_i(\cdot)$ relative to tenant $i$ is a concave function $P_i(w_i)$ of $w_i \in [0,1]$, with $P_i(0) = 0$. The collection of bandwidth pricing strategies $\{P_i(\cdot) : i = 1, \ldots, N\}$ forms a *pricing policy*.

These assumptions on $P_i(\cdot)$ are close to real-life situations: $P_i(\cdot)$ is concave because wholesale is cheaper than retail, i.e., the more a user buys some goods, the lower the unit price $P(w_i)/w_i$; $P_i(0) = 0$ because a user receiving no service pays nothing. In our model, $\{P_i(\cdot)\}$ can be authoritatively enforced in a controlled market, or negotiated between the tenants and broker in a free market.

## III. An Outline of Contributions

In our proposed market, different entities have different objectives. From the standpoint of cloud resource utilization, it is best to serve all the demands with the minimum amount of resource. In contrast, the broker wishes to turn a profit, while tenants favor a lower bandwidth reservation price. Our study is inspired by coordinating the conflicts between these objectives.

First of all, a basic goal of the cloud is to enhance resource efficiency, or to consolidate the workload. When the aggregate capacity is sufficient, this means to serve all the given demands $\{D_i\}$ by reserving as little bandwidth resource $\sum_s A_s$ as possible from the cloud. When multiple cloud providers coexist, a social planner can determine the load direction $\mathbf{W}^*$ that minimizes the aggregate bandwidth reservation by solving a *workload consolidation problem*:

$$\min_{\mathbf{W}} \sum_s A_s \tag{5}$$

$$\text{s.t.} \quad A_s \le C_s, \quad \forall s, \tag{6}$$

$$w_i = 1, \quad \forall i. \tag{7}$$

Note that different $\mathbf{W}$ can result into different $\sum_s A_s$ because $D_i$'s may be intercorrelated and *mixing anti-correlated demands can consolidate the workload*.

In reality, it is impractical to have a social planner perform optimal workload consolidation (5) across multiple clouds, which may not be cooperative to each other in the first place. Can we use a mechanism inherent in the market to optimally direct demands from tenants to different clouds? To achieve this, we resort to a *profit-driven* broker. We simply let the broker decide $\mathbf{W}^*$ by maximizing its profit $R(\mathbf{W})$ as follows:

$$\max_{\mathbf{W}} R(\mathbf{W}) = \sum_i P_i(w_i) - \sum_s A_s \tag{8}$$

$$\text{s.t.} \quad A_s \le C_s, \quad \forall s. \tag{9}$$

We observe that such a broker always has some incentive to operate, since its profit is non-negative for any pricing policy $\{P_i(\cdot)\}$, i.e., $R(\mathbf{W}^*) \ge R(\mathbf{W})\big|_{w_{si}=0, \forall s,i} = 0$.

However, a selfish broker has no obligation to accommodate all the demands for service: it may output $w^*_i < 1$ for some $i$. Neither will it necessarily optimize load direction for cloud workload consolidation. In other words, there is a fundamental conflict between workload consolidation across the clouds and profit maximization at the broker, since (8) is not equivalent to (5) in general. Such a conflict leads to the first question we ask:

(Q1) Under what pricing policy $\{P_i(\cdot)\}$ a profit-driven broker is also a social planner that optimizes workload consolidation across multiple clouds while accommodating all demands for service, i.e., problems (8) and (5) are equivalent?

An answer to this question enables us to enhance cloud resource efficiency simply by enforcing a good pricing policy at an economically viable broker, instead of resorting to impratical social planning across clouds. In Sec. IV, we characterize the entire region of all such good pricing policies. Furthermore, with consolidated workload and a reduced aggregate service cost $\sum_s A_s$, each tenant may indeed expect to pay a lower price $P_i(\cdot)$ for bandwidth reservation.

Now we look at the reservation price each tenant has to pay. First, a naive pricing policy $\{P^0_i(\cdot)\}$ is to charge \$1 for every unit bandwidth reserved for every tenant, where

$$P^0_i(w_i) = f_\epsilon(D_i w_i) \cdot 1, \text{ with } P^0_i(1) = f_\epsilon(D_i) = B_i, \ \forall i.$$

Since tenants wish to save money, the broker pricing policy $\{P_i(\cdot)\}$ must satisfy $P_i(w_i) \le P^0_i(w_i)$ for all $w_i \in [0,1]$ and for all $i$. Otherwise, tenants will reserve bandwidth from cloud providers directly and have no incentive to use the broker. What is the maximum and fair discount each tenant can enjoy? The pricing region characterization in Sec. IV also answers this question.

One may have noted that pricing policies found in (Q1) are enforced in the market by a supervisory agency other than the broker or VoD companies. However, in a free market, a selfish tenant may bargain with the broker to negotiate a reservation price. To closely model a free market, we assume that each tenant $i$ can submit to the broker any pricing strategy $P_i(\cdot)$ it prefers and accepts the service fraction $w_i^*$ returned by the broker. Based on $\{P_i(\cdot)\}$ collected from all the tenants, the broker determines load direction $\mathbf{W}^*$ and thus the service fraction $w_i^*$ for each $i$ *by maximizing its own profit*. We ask the question:

(Q2) In a free market where each selfish tenant (VoD provider) competes for service by submitting a pricing strategy $P_i(\cdot)$, what will $\{P_i(\cdot)\}$ eventually look like?

The free market can be modeled through a game played by all VoD providers as tenants, each with an independent strategy $P_i(\cdot)$, which is the pricing scheme it submits. In a system of sufficient capacity, we assume that a selfish tenant always 1) *expects to get fully served*, and 2) *tries to reduce the reservation price it has to pay if condition 1) is met*. To put the above formally, under a set of submitted strategies $\{P_i(\cdot)\}$, we define a utility function associated with each tenant $i$ as

$$U_i[P_1(\cdot), \ldots, P_N(\cdot)] = \begin{cases} -P_i(w_i^*), & \text{if } w_i^* = 1, \\ -\infty, & \text{if } w_i^* < 1, \end{cases} \tag{10}$$

where $w_i^*$ is determined by the broker via solving (22).

A utility function of the form (10) assumes that unlike the case of scarce metal (e.g., gold), it is not profitable for a VoD provider as a tenant to deliberately deny VoD user requests (making $w_i^* < 1$) in exchange for a lower bandwidth reservation cost. The reason is that in reality, popularity matters more to a VoD provider than instantaneous profit earned per video request from users. In a market of sufficient supply, if a VoD provider is found sacrificing some of its user requests to chase for cheaper bandwidth deals, it will lose reputation and revenue in the long run.

Intuitively, a VoD provider as a tenant cannot submit too low a $P_i(\cdot)$, since otherwise the broker will not serve $D_i$ completely out of profit concerns. Ironically, even if $P_i(\cdot)$ is high, $D_i$ may not get fully served either, depending on the prices submitted by other tenants. In Sec. V, we present a key finding of this paper:

**Theorem 1:** In a free market, $\{P_i(\cdot)\}$ will converge to a **unique Nash equilibrium** $\{P_i^*(\cdot)\}$, where $w_i^* = 1$ and

$$P_i^*(w_i^*) = \mu_i + \theta\sigma_i\rho_{iM}, \tag{11}$$

$\rho_{iM} \in [-1, 1]$ being the correlation coefficient between $D_i$ and $\sum_i D_i$.

It turns out that $\{P_i^*(\cdot)\}$ is also the lower border of the *good pricing region*, which means that without intervention, the free market competition will lead to cloud resource efficiency maximization, where each tenant receives the maximum price discount. In Sec. V and VI, we discuss the deeper connections between demand statistics $\mu_i, \sigma_i, \rho_{iM}$ and bandwidth reservation pricing, as revealed in Theorem 1.

## IV. Pricing Region in Controlled Markets

We address question (Q1) in this section. We first try to understand the structure of optimal load direction in the ideal case of social planning, and then characterize the pricing policies under which a profit-driven broker will exactly behave like a social planner and facilitate cloud workload consolidation. Such a pricing region also depicts the maximum price discount a tenant can enjoy from the broker.

Let us first introduce several frequently used notations. We define $\sum_i D_i$ as the *market demand* with its standard deviation denoted by $\sigma_M$. Denote $\sigma_{iM}$ the covariance between $D_i$ and the market demand $\sum_i D_i$. Clearly, we have

$$\sigma_{iM} = \mathbf{E}[(D_i - \mu_i)(\textstyle\sum_j D_j - \sum_j \mu_j)] = \sum_{j=1}^{N} \sigma_{ij}, \tag{12}$$

$$\sigma_M = \sqrt{\mathbf{Var}[\textstyle\sum_i D_i]} = \sqrt{\mathbf{1}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{1}} = \sqrt{\textstyle\sum_{i,j} \sigma_{ij}}. \tag{13}$$

Further denote $\rho_{iM} = \sigma_{iM}/\sigma_i\sigma_M \in [-1, 1]$ the correlation coefficient between $D_i$ and the market demand $\sum_i D_i$.

For convenience, let $\mathbf{L} = [L_1, \ldots, L_S]^\mathsf{T}$, $\mathbf{w}_s = [w_{s1}, \ldots, w_{sN}]^\mathsf{T}$, $\mathbf{w} = [w_1, \ldots, w_N]$, and $\mathbf{w}_s^* = [w_{s1}^*, \ldots, w_{sN}^*]$. Clearly, we have $\mathbf{w} = \sum_s \mathbf{w}_s$.

### A. The Optimal Load Direction of a Social Planner

We first consider the ideal case that there is a social planner that optimally directs load based on (5). For simplicity, we assume each random demand $D_i$ is Gaussian-distributed, which has been verified by real-world traces in [8] for large $D_i$. Recall that for a Gaussian random variable $X$, we have

$$f_\epsilon(X) = \mathbf{E}[X] + \theta\sqrt{\mathbf{Var}[X]}, \quad \theta = F^{-1}(1-\epsilon), \tag{14}$$

where $F(\cdot)$ is the CDF of normal distribution $\mathcal{N}(0, 1)$. For example, when $\epsilon = 2\%$, we have $\theta = 2.05$. Note that if each $D_i$ is Gaussian-distributed, so is $L_s = \sum_i w_{si} D_i$, and we have

$$\mathbf{E}[L_s] = \mu_1 w_{s1} + \ldots + \mu_N w_{sN} = \boldsymbol{\mu}^\mathsf{T}\mathbf{w}_s,$$
$$\mathbf{Var}[L_s] = \textstyle\sum_{i,j} \rho_{ij}\sigma_i\sigma_j w_{si} w_{sj} = \mathbf{w}_s^\mathsf{T}\boldsymbol{\Sigma}\mathbf{w}_s.$$

It follows that

$$B_i = f_\epsilon(D_i) = \mu_i + \theta\sigma_i,$$
$$A_s = f_\epsilon(L_s) = \boldsymbol{\mu}^\mathsf{T}\mathbf{w}_s + \theta\sqrt{\mathbf{w}_s^\mathsf{T}\boldsymbol{\Sigma}\mathbf{w}_s}.$$

Therefore, the cloud workload consolidation problem (5) has the following form under Gaussian demands:

$$\min_{\mathbf{W}} \textstyle\sum_s (\boldsymbol{\mu}^\mathsf{T}\mathbf{w}_s + \theta\sqrt{\mathbf{w}_s^\mathsf{T}\boldsymbol{\Sigma}\mathbf{w}_s}), \tag{15}$$

$$\text{s.t.} \quad \boldsymbol{\mu}^\mathsf{T}\mathbf{w}_s + \theta\sqrt{\mathbf{w}_s^\mathsf{T}\boldsymbol{\Sigma}\mathbf{w}_s} \le C_s, \quad \forall s, \tag{16}$$

$$\mathbf{w} = \textstyle\sum_s \mathbf{w}_s = \mathbf{1}, \tag{17}$$

$$\mathbf{0} \preceq \mathbf{w}_s \preceq \mathbf{1}, \quad \forall s, \tag{18}$$

where $\mathbf{1} = [1, \ldots, 1]^\mathsf{T}$ and $\mathbf{0} = [0, \ldots, 0]^\mathsf{T}$ are $N$-dimensional column vectors. Constraint (16) is equivalent to $A_s \le C_s$ and thus to (2).

Although problem (15) is convex optimization, it has coupled objectives and constraints, and takes numerical solvers a long time to converge for a large $S$. Our previous work [8] has given nearly closed-form solutions to problem (15) in the following theorem (please refer to [8] for the proof):

**Theorem 2:** When $C_{\mathsf{sum}} \geq \boldsymbol{\mu}^{\mathsf{T}}\mathbf{1} + \theta\sqrt{\mathbf{1}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{1}}$, an optimal solution $[w^*_{si}]$ to (15) is given by

$$w^*_{si} = \alpha_s, \quad \forall i, \quad s = 1, \ldots, S, \tag{19}$$

where $\alpha_1, \ldots, \alpha_S$ can be any solution to

$$\sum_s \alpha_s = 1,\ 0 \leq \alpha_s \leq \min\left\{1, \frac{C_s}{\boldsymbol{\mu}^{\mathsf{T}}\mathbf{1} + \theta\sqrt{\mathbf{1}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{1}}}\right\},\ \forall s. \tag{20}$$

When $C_{\mathsf{sum}} < \boldsymbol{\mu}^{\mathsf{T}}\mathbf{1} + \theta\sqrt{\mathbf{1}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{1}}$, there is no feasible solution that satisfies constraints (16) to (18).

Using Theorem 2, the broker can fast check if all demands can be served, simply by comparing the total cloud capacity $C_{\mathsf{sum}}$ with total bandwidth required for all demands combined:

$$f_\epsilon(\textstyle\sum_i D_i) = \boldsymbol{\mu}^{\mathsf{T}}\mathbf{1} + \theta\sqrt{\mathbf{1}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{1}}.$$

Hence, the assumption on sufficient capacity essentially means that we have assumed $\Sigma_s C_s \geq \boldsymbol{\mu}^{\mathsf{T}}\mathbf{1} + \theta\sqrt{\mathbf{1}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{1}}$. Theorem 2 implies that to optimally consolidate workload, the broker should direct the same portions $\alpha_1, \ldots, \alpha_S$ of $D_i$ to cloud provider $1, \ldots, S$ for service, respectively, no matter which tenant $i$ it is. The set of $\alpha_s$ can be found easily subject to the linear constraints (20).

The maximum bandwidth saving of joint bandwidth booking over individual booking for each tenant is

$$\begin{aligned}\Delta B(\mathbf{W}^*) &= \textstyle\sum_i B_i - \sum_s A_s \\ &= \textstyle\sum_i(\mu_i + \theta\sigma_i) - \sum_s(\boldsymbol{\mu}^{\mathsf{T}}\mathbf{w}^*{}_s + \theta\sqrt{\mathbf{w}^*{}_s^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{w}^*{}_s}) \\ &= \theta(\boldsymbol{\sigma}^{\mathsf{T}}\mathbf{1} - \sqrt{\mathbf{1}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{1}}) = \theta(\textstyle\sum_i \sigma_i - \sigma_M),\end{aligned} \tag{21}$$

which is $\theta$ times the gap between the sum of all demand standard deviations and the standard deviation of all demands combined. This confirms the belief that statistical multiplexing saves resource reservation.

### *B. The Good Pricing Region for the Broker*

However, social planning is not practical in the presence of rivalry clouds. Instead, in reality, we may only resort to a broker to optimally direct workloads. Recall that a selfish broker determines $\mathbf{W}^*$ by maximizing its own profit via (8). Similarly, under Gaussian demands, we can translate (8) into the following:

$$\max_{\mathbf{W}} \textstyle\sum_i P_i(w_i) - \sum_s(\boldsymbol{\mu}^{\mathsf{T}}\mathbf{w}_s + \theta\sqrt{\mathbf{w}_s^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{w}_s}), \tag{22}$$

$$\text{s.t.} \quad \boldsymbol{\mu}^{\mathsf{T}}\mathbf{w}_s + \theta\sqrt{\mathbf{w}_s^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{w}_s} \leq C_s, \quad \forall s, \tag{23}$$

$$\mathbf{w} = \textstyle\sum_s \mathbf{w}_s \preceq \mathbf{1}, \tag{24}$$

$$\mathbf{0} \preceq \mathbf{w}_s \preceq \mathbf{1}, \quad \forall s. \tag{25}$$

Note that a selfish broker not only has a different objective (22) than (15), but also has a different constraint (24) than (17), in that a profit-driven broker has *no obligation to accommodate all demands for service*, if doing so does not bring more profit. Instead, the broker decides the service fraction $w^*_i \geq 1$ for each tenant depending on the pricing policy $\{P_i(\cdot)\}$. We aim to find all good pricing policies such that the selfish broker behaves like an altruistic social planner, i.e., problem (22) is equivalent to (15).

The following theorem gives a necessary and sufficient condition for a good pricing policy:

**Theorem 3:** Broker profit maximization (22) and cloud workload consolidation (15) have a same optimal solution (19), if and only if

$$P'_i(1) \geq \mu_i + \theta \cdot \frac{\sigma_{iM}}{\sigma_M}, \quad \forall i, \tag{26}$$

where $\sigma_{iM}$ is the covariance between $D_i$ and $\sum_i D_i$ given by (12) and $\sigma_M$ is the standard deviation of $\sum_i D_i$ given by (13). Furthermore, if $P'_i(1) < \mu_i + \theta\sigma_{iM}/\sigma_M$ for some $i$, then $\mathbf{w}^* \neq \mathbf{1}$.

**Proof Sketch**: We first show that if $P'_i(1) \geq \mu_i + \theta\sigma_{iM}/\sigma_M$, $\forall i$, (19) is also an optimal solution to (22).

The function $f(\mathbf{w}_s) = \sqrt{\mathbf{w}_s^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{w}_s}$ is a convex cone. We have $f[(\mathbf{w}_1 + \mathbf{w}_2)/2] \leq [f(\mathbf{w}_1) + f(\mathbf{w}_2)]/2$. Applying this inequality iteratively, we can prove

$$\textstyle\sum_s \sqrt{\mathbf{w}_s^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{w}_s} \geq \sqrt{(\sum_s \mathbf{w}_s^{\mathsf{T}})\boldsymbol{\Sigma}(\sum_s \mathbf{w}_s)} = \sqrt{\mathbf{w}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{w}}. \tag{27}$$

Using the bound (27), problem (22) can be relaxed to the following concave problem:

$$\max_{\mathbf{w}} R_u(\mathbf{w}) = \textstyle\sum_i P_i(w_i) - (\boldsymbol{\mu}^{\mathsf{T}}\mathbf{w} + \theta\sqrt{\mathbf{w}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{w}}), \tag{28}$$

$$\text{s.t.} \quad \boldsymbol{\mu}^{\mathsf{T}}\mathbf{w} + \theta\sqrt{\mathbf{w}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{w}} \leq \textstyle\sum_s C_s, \tag{29}$$

$$\mathbf{w} \preceq \mathbf{1}, \tag{30}$$

with $R_u(\mathbf{w}) \geq R(\mathbf{W})$ being concave and (29) being a relaxation of constraints (23). Since $\boldsymbol{\mu}^{\mathsf{T}}\mathbf{1} + \theta\sqrt{\mathbf{1}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{1}} \leq \sum_s C_s$, $\mathbf{w} = \mathbf{1}$ is a feasible solution to (28).

$$\begin{aligned}\left.\frac{\partial R_u(\mathbf{w})}{\partial w_i}\right|_{\mathbf{w}=1} &= \left.P'_i(w_i)\right|_{w_i=1} - \mu_i - \left.\frac{\theta\sum_{j=1}^N \sigma_{ij}w_j}{\sqrt{\mathbf{w}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{w}}}\right|_{\mathbf{w}=1} \\ &= P'_i(1) - \mu_i - \frac{\theta\mathbf{Cov}[D_i, \sum_j D_j]}{\sqrt{\mathbf{1}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{1}}} \\ &= P'_i(1) - \mu_i - \frac{\sigma_{1M}}{\sigma_M}.\end{aligned} \tag{31}$$

If (26) is satisfied, i.e., $\partial R_u(\mathbf{w})/\partial w_i\big|_{\mathbf{w}=1} \geq 0$, $\forall i$, according to the gradient ascent algorithm for concave maximization, $\mathbf{w} = \mathbf{1}$ is the optimal solution to (28). If $w_{si} = \alpha_s$ given by (19), we have 1) all constraints (23)-(25) are satisfied, 2) $R_u(\mathbf{w}) = R(\mathbf{W})$ and 3) $R_u(\mathbf{w}) = R_u(\mathbf{1})$ in problem (28) reaches optimality. Thus, (19) is an optimal solution to (22).

Please refer to our technical report [9] for a proof of the converse. ■

It immediately follows that a good pricing policy must fall into a certain region, depicted by the following corollary:

**Corollary 1:** In a good pricing policy $\{P_i(\cdot) : i = 1, \ldots, N\}$, each $P_i(\cdot)$ must satisfy $\forall w_i \in [0, 1]$,

$$(\mu_i + \theta \cdot \frac{\sigma_{iM}}{\sigma_M})w_i \leq P_i(w_i) \leq (\mu_i + \theta\sigma_i)w_i. \tag{32}$$

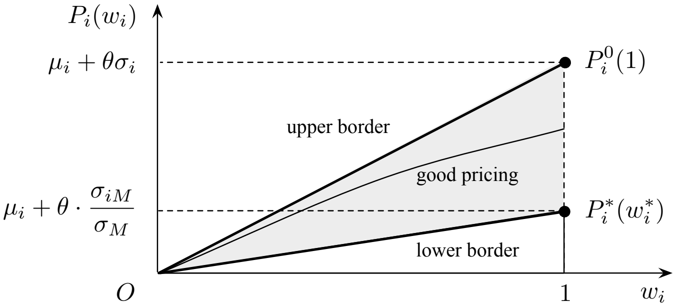


Fig. 2. The region of $P_i(\cdot)$ in a good pricing policy $\{P_i(\cdot)\}$. $P_i(\cdot)$ is between $P_i^*(\cdot)$ and $P_i^0(\cdot)$, and satisfies $P_i'(1) \geq P_i^{*\prime}(1)$.

**Proof**: First, $P_i(w_i)$ is upper-bounded by the cloud provider pricing function $P_i^0(w_i) = (\mu_i + \theta\sigma_i)w_i$, since otherwise the tenants would have booked bandwidth from cloud providers directly. Consider another pricing policy $\{P_i^*(\cdot)\}$, where

$$P_i^*(w_i) = (\mu_i + \theta \cdot \frac{\sigma_{iM}}{\sigma_M})w_i, \quad \forall i. \tag{33}$$

To prove that $P_i(\cdot)$ is lower-bounded by $P_i^*(\cdot)$, we need the following lemma (please refer to our technical report [9] for the proof):

**Lemma 1:** Let $f(x)$ be a concave function on $[0, 1]$ with $f(0) = 0$. Let $k$ be a (constant) real number. If $f'(1) \geq k$, then for each $x \in (0, 1]$, $f(x) \geq kx$. In this case, if $f(x') = kx'$ for some $x' \in (0, 1]$, then $f(x) \equiv kx$, $\forall x \in [0, 1]$.

By Theorem 3 and Lemma 1, in a good pricing policy, for each $i$, we have $P_i(w_i) \geq P_i^*(w_i)$, $\quad \forall w_i \in (0, 1]$, with equality achieved only if $P_i(\cdot) \equiv P_i^*(\cdot)$. ■

Fig. 2 illustrates the region of $P_i(\cdot)$ in a good pricing policy, where each $P_i(\cdot)$ is a concave function between $P_i^*(\cdot)$ and $P_i^0(\cdot)$ (inclusive) with $P_i'(1) \geq \mu_i + \theta\sigma_{iM}/\sigma_M$. In the good pricing region, a selfish broker behaves like a social planner facilitating workload consolidation without denying any demand. When a pricing policy falls out of this region, the broker will reject a part of demands to chase a profit. In other words, the broker will determine a $w_j^* < 1$ for some $j$, which is undesirable to tenants when supply is sufficient. Note that $P_i(\cdot)$ affects not only $w_i^*$ but also $w_j^*$ for any other tenant $j$. This phenomenon will be understood in Sec. V.

Now let us check the broker profit $R(\mathbf{W}^*)$. Under a good pricing policy, $w_i^* = 1$. Thus, the broker's total cost is

$$\begin{aligned}\textstyle\sum_s A_s &= \textstyle\sum_s(\boldsymbol{\mu}^\mathsf{T}\mathbf{w}_s^* + \theta\sqrt{\mathbf{w}_s^{*\mathsf{T}}\boldsymbol{\Sigma}\mathbf{w}_s^*}) = \boldsymbol{\mu}^\mathsf{T}\mathbf{1} + \theta\sqrt{\mathbf{1}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{1}} \\ &= \textstyle\sum_i \mu_i + \theta\sigma_M.\end{aligned}$$

Using Corollary 1, we obtain

$$\begin{aligned}R(\mathbf{W}^*) &\leq \textstyle\sum_i(\mu_i + \theta\sigma_i)w_i^* - \sum_s A_s = \theta(\sum_i \sigma_i - \sigma_M) \\ &= \Delta B(\mathbf{W}^*),\end{aligned} \tag{34}$$

with equality achieved when $P_i(\cdot) \equiv P_i^0(\cdot)$, and

$$\begin{aligned}R(\mathbf{W}^*) &\geq \textstyle\sum_i(\mu_i + \frac{\theta\sigma_{iM}}{\sigma_M})w_i^* - \sum_s A_s = \theta\frac{\sum_i \sigma_{iM}}{\sigma_M} - \theta\sigma_M \\ &= 0,\end{aligned} \tag{35}$$

with equality achieved when $P_i(\cdot) \equiv P_i^*(\cdot)$.

The above bounds show that the maximum broker profit is essentially the maximum achievable bandwidth saving $\Delta B(\mathbf{W}^*)$ from joint bandwidth booking given by (21). Moreover, the maximum profit is achieved when the broker adopts the same pricing policy $\{P_i^0(\cdot)\}$ as cloud providers do. On the other hand, when $P_i(\cdot) \equiv P_i^*(\cdot)$, the broker profit is zero, which means all the profit made from cloud bandwidth multiplexing has been rewarded to tenants as price discounts.

In a *controlled* market, the lowest price that tenant $i$ should pay for having all its demand served is

$$P_i^*(1) = \mu_i + \theta\sigma_{iM}/\sigma_M. \tag{36}$$

If any tenant $i$ pays less than $P_i^*(1)$, the broker cannot serve all the demands and not only $i$ but any tenant may be denied for service.

## V. Pricing in Free Markets

In Sec. IV, pricing is enforced as a policy by a supervisory agency other than the broker and tenants, which is still difficult to implement in reality. In this section, we consider a *free market*, where each tenant could bargain with the broker and submit its own pricing strategy $P_i(\cdot)$. Recall that we have defined utility $U_i$ in (10). All the tenants are essentially playing a game, each of whom aim to maximize its own utility $U_i$ which is a function of $P_1(\cdot), \ldots, P_N(\cdot)$. Intuitively, tenant $i$ cannot submit too low a $P_i(\cdot)$, beyond which the broker may deny a part of $D_i$ and return $w_i^* < 1$, leading to $U_i = -\infty$.

To find the equilibrium of a free market, we just need to find the Nash equilibrium in the above game, where no tenant $i$ can get a better utility by unilaterally changing $P_i(\cdot)$. From the results of Sec. IV, one may conjecture that $\{P_i^*(\cdot)\}$ is a Nash equilibrium. However, is this the unique Nash equilibrium? Could there be another equilibrium point where any tenant will submit a very low price and gets a utility of $-\infty$ without being able to better off by changing its pricing? In other words, can market collapse happen?

### A. Nash Equilibrium: Existence and Uniqueness

**Theorem 4:** If tenants have utility (10) and the broker decides $\mathbf{W}^*$ by maximizing its profit via (22), then $\{P_i(\cdot)\}$ will converge to a **unique Nash equilibrium** $\{P_i^*(\cdot)\}$, where

$$P_i^*(w_i) = (\mu_i + \theta \cdot \frac{\sigma_{iM}}{\sigma_M})w_i, \quad 0 \leq w_i \leq 1, \tag{37}$$

where $\sigma_{iM}$ and $\sigma_M$ are given by (12) and (13), respectively.

Let $P_{-i}(\cdot)$ represent the pricing strategies of all tenants except for tenant $i$. The proof of Theorem 4 relies on the following two lemmas and Theorem 3 that characterize the solution structure of (22) given $\{P_i(\cdot)\}$.

**Lemma 2:** If $P_i'(1) < \mu_i + \theta\sigma_{iM}/\sigma_M$, then $w_i^* < 1$ regardless of $P_{-i}(\cdot)$.

**Proof Sketch:** It suffices to show that

$$\left.\frac{\partial R_u(\mathbf{w})}{\partial w_i}\right|_{w_i=1} < 0 \text{ for all } \mathbf{w} \text{ with } w_i = 1,$$

since in this case we have $w_i^* = 1$ regardless of $\{w_j\}_{j \neq i}$ by the gradient ascent algorithm.

Recall that $\partial R_u(\mathbf{w})/\partial w_i$ has been given in (31). If $P_i'(w_i) < \mu_i + \theta\sigma_{iM}/\sigma_M$, then

$$\left.\frac{\partial R_u(\mathbf{w})}{\partial w_i}\right|_{w_i=1} < \theta\left(\frac{\sigma_{iM}}{\sigma_M} - \frac{\sum_{j=1}^N \sigma_{ij}w_j}{\sqrt{\mathbf{w}^\intercal\mathbf{\Sigma}\mathbf{w}}}\right)\Bigg|_{w_i=1}.$$

Now construct a function $g(w_1,\ldots,w_N)$ as

$$g(w_1,\ldots,w_N) = \sum_{j=1}^N \sigma_{ij}w_j/\sqrt{\mathbf{w}^\intercal\mathbf{\Sigma}\mathbf{w}}.$$

Clearly, $g(1,\ldots,1) = \sigma_{iM}/\sigma_M$. Hence, it suffices to show that

$$g(w_1,\ldots,w_N) \geq g(1,\ldots,1) \text{ for all } \mathbf{w} \text{ with } w_i = 1.$$

Without loss of generality, we only need to prove the case when $N = 2$ and $i = 1$, that is,

$$\frac{\sigma_1^2 + \rho_{12}\sigma_1\sigma_2 w_2}{\sqrt{\sigma_1^2 + \sigma_2^2 w_2^2 + 2\rho_{12}\sigma_1\sigma_2 w_2}} \geq \frac{\sigma_1(\sigma_1 + \rho_{12}\sigma_2)}{\sqrt{\sigma_1^2 + \sigma_2^2 + 2\rho_{12}\sigma_1\sigma_2}}.$$

It is easy to check that the above inequality indeed holds. Hence, $\partial R_u(\mathbf{w})/\partial w_i|_{w_i=1} < 0$ for all $\mathbf{w}$ with $w_i = 1$. ■

**Lemma 3:** If $P_i(w_i) = P_i^0(w_i) = (\mu_i + \theta\sigma_i)w_i$ for $w_i \in [0,1]$, then $w_i^* = 1$, regardless of $P_{-i}(\cdot)$.

**Proof Sketch:** It suffices to show that

$$\left.\frac{\partial R_u(\mathbf{w})}{\partial w_i}\right|_{w_i=1} \geq 0 \text{ for all } \mathbf{w} \text{ with } w_i = 1,$$

since in this case we have $w_i^* = 1$ regardless of $\{w_j\}_{j\neq i}$.

Recall that $\partial R_u(\mathbf{w})/\partial w_i$ has been given in (31). If $P_i(w_i) = (\mu_i + \theta\sigma_i)w_i$, then $P_i'(w_i) = \mu_i + \theta\sigma_i$. Thus,

$$\left.\frac{\partial R_u(\mathbf{w})}{\partial w_i}\right|_{w_i=1} = \theta\left(\sigma_i - \frac{\sum_{j=1}^N \sigma_{ij}w_j}{\sqrt{\mathbf{w}^\intercal\mathbf{\Sigma}\mathbf{w}}}\right)\Bigg|_{w_i=1}.$$

Now construct a random variable $X$ as $X = \sum_{j\neq i} w_j D_j + D_i$. By the Cauchy-Schwarz inequality, we have

$$\mathbf{E}^2[(D_i - \mu_i)(X - \mathbf{E}[X])] \leq \mathbf{E}[(D_i - \mu_i)^2]\mathbf{E}[(X - \mathbf{E}[X])^2],$$

which is $\left(\sigma_i^2 + \sum_{j\neq i}\sigma_{ij}w_j\right)^2 \leq \sigma_i^2\sqrt{\mathbf{w}^\intercal\mathbf{\Sigma}\mathbf{w}}|_{w_i=1}$. It follows immediately that

$$\sigma_i \geq \left.\frac{\sum_{j=1}^N \sigma_{ij}w_j}{\sqrt{\mathbf{w}^\intercal\mathbf{\Sigma}\mathbf{w}}}\right|_{w_i=1}.$$

Hence, $\partial R_u(\mathbf{w})/\partial w_i|_{w_i=1} \geq 0$ for all $\mathbf{w}$ with $w_i = 1$. ■

We are now ready to prove Theorem 4.

**Proof of Theorem 4**: We first show that $\{P_i^*(\cdot)\}$ is a Nash equilibrium. We need to show that at $\{P_i^*(\cdot)\}$, any tenant $i$ cannot increase its utility $U_i$ by unilaterally changing $P_i(\cdot)$, i.e.,

$$U_i[P_i^*(\cdot), P_{-i}^*(\cdot)] \geq U_i[P_i(\cdot), P_{-i}^*(\cdot)] \quad \forall i. \tag{38}$$

We exhaust the possibilities of $P_i(\cdot)$ by considering the range of $P_i'(1)$. First, if $P_i'(1) < \mu_i + \theta\sigma_{iM}/\sigma_M$, by Lemma 2, $w_i^* < 1$, and by (10),

$$U_i[P_i(\cdot), P_{-i}^*(\cdot)] = -\infty < -P_i^*(1) = U_i[P_i^*(\cdot), P_{-i}^*(\cdot)].$$

Second, if $P_i'(1) \geq \mu_i + \theta\sigma_{iM}/\sigma_M$, while $P_{-i}(\cdot) \equiv P_{-i}^*(\cdot)$, by Theorem 3, we have $w_i^* = \sum_s w_{si}^* = \sum_s \alpha_s = 1$. Hence,

$$U_i[P_i(\cdot), P_{-i}^*(\cdot)] = -P_i(1) \leq -P_i^*(1) = U_i[P_i^*(\cdot), P_{-i}^*(\cdot)].$$

The inequality is due to Lemma 1 and $P_i^*(w_i)$ is linear in $w_i$ and passes $(0,0)$. We have thus proved (38). Therefore, $\{P_i^*(\cdot)\}$ is indeed a Nash equilibrium.

Now we show $\{P_i^*(\cdot)\}$ is the unique Nash equilibrium, i.e., if $\{P_i(\cdot)\}$ is a Nash equilibrium, then $P_i(\cdot) \equiv P_i^*(\cdot)$. First of all, we must have

$$P_i'(1) \geq \mu_i + \theta\sigma_{iM}/\sigma_M, \quad \forall i. \tag{39}$$

We prove (39) by contradiction. Assume $P_i' < \mu_i + \theta\sigma_{iM}/\sigma_M$ for some $i$. By Lemma 2, $w_i^* < 1$ and thus $U_i = -\infty$. If tenant $i$ uses another strategy $P_i^0(w_i) = (\mu_i + \theta\sigma_i)w_i$, then by Lemma 3, $w_i^* = 1$, regardless of $P_{-i}(\cdot)$. Hence,

$$U_i[P_i^0(\cdot), P_{-i}(\cdot)] = -P_i^0(1) > -\infty = U_i[P_i(\cdot), P_{-i}(\cdot)],$$

contradicting with the definition of Nash equilibrium that unilaterally changing $P_i(\cdot)$ cannot increase $U_i$. Therefore, (39) must hold.

If (39) holds, by Theorem 3, $w_i^* = \sum_s w_{si}^* = \sum_s \alpha_s = 1$ for all $i$, and thus $U_i[P_i(\cdot), P_{-i}(\cdot)] = -P_i(1)$ for all $i$. If (39) holds, by Lemma 1, $P_i(1) \geq (\mu_i + \theta\sigma_{iM}/\sigma_M)\cdot 1 = P_i^*(1)$, with equality achieved *only if* $P_i(w_i) \equiv (\mu_i + \theta\sigma_{iM}/\sigma_M)w_i \equiv P_i^*(w_i)$ for $w_i \in [0,1]$. In other words, if $P_i(\cdot)$ is not $P_i^*(\cdot)$,

$$U_i[P_i(\cdot), P_{-i}(\cdot)] = -P_i(1) < -P_i^*(1) = U_i[P_i^*(\cdot), P_{-i}(\cdot)],$$

contradicting with the fact that $\{P_i(\cdot)\}$ is a Nash equilibrium. Therefore, $P_i(\cdot)$ must be $P_i^*(\cdot)$, which proves that $\{P_i^*(\cdot)\}$ is the unique Nash equilibrium. ■

### *B. Discussions*

Theorem 4 shows that even without a supervisory party, the selfishness of tenants will drive $\{P_i(\cdot)\}$ into an equilibrium point $\{P_i^*(\cdot)\}$, which is exactly the lower border of the good pricing region. Therefore, in a *free market*, a profit-driven broker is naturally a workload consolidator that optimizes the cloud resource efficiency.

In equilibrium market, we have $P_i(\cdot) \equiv P_i^*(\cdot)$ and broker profit $R(\mathbf{W}^*) = 0$ according to (35). This means the broker would be unable to exploit the *arbitrage* opportunity from bandwidth multiplexing in equilibrium. When *multiple brokers* exist, the unique Nash equilibrium still holds because the game is played by all the VoD providers, while each broker determines the service fraction $w_i^*$ based on the submitted $\{P_i(\cdot)\}$ in the same way. And the competition among brokers will necessarily lead to a zero broker profit. However, the broker can still earn income through other means such as agent fees, membership fees and commercials.

More importantly, let us take a closer look at the equilibrium bandwidth reservation price for each tenant. In equilibrium, we have $w_i^* = 1$ and $P_i^*(w_i^*) = P_i^*(1)$. Rewriting (36), we obtain

$$\$\ P_i^*(w_i^*) = \mu_i + \theta\sigma_i\rho_{iM} = \mu_i + [f_\epsilon(D_i) - \mu_i]\rho_{iM}, \tag{40}$$

proving Theorem 1, the key theorem of this paper.

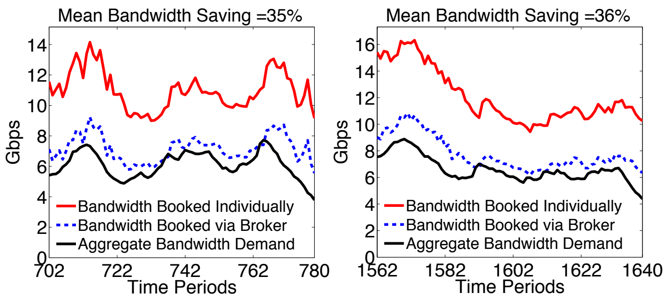


(a) $\epsilon = 5\%$, $\epsilon' = 5.06\%$ (b) $\epsilon = 1\%$, $\epsilon' = 1.27\%$

Fig. 3. The aggregate bandwidth $\sum_s A_s(t)$ booked by the broker, compared to the aggregate bandwidth $\sum_i B_i(t)$ needed if each channel books individually and the real aggregate demand $\sum_i D_i(t)$.

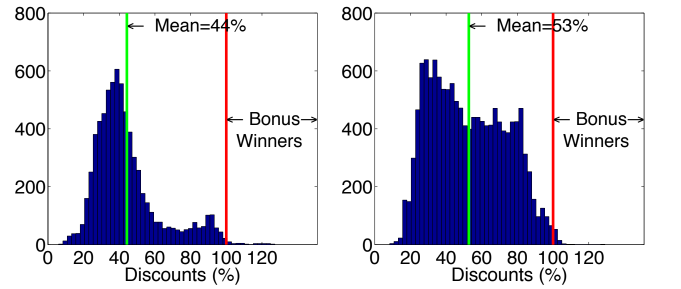


(a) $t$ =702—780, 91 channels (b) $t$ =1562—1640, 176 channels

Fig. 4. The histogram of payment discounts of all channels at all times in equilibrium, i.e., the histogram of $1 - P_i^*(1,t)/P_i^0(1,t)$ for all $i$ and $t$.

The free market automatically guarantees a fairness on the bandwidth costs of different tenants. Note that the tenant payment consists of two parts: \$1 per unit bandwidth for mean demand $\mu_i$, and \$$\rho_{iM}$ per unit bandwidth for reservation beyond the expected demand $\mu_i$. Apparently, the discount for tenant $i$ depends on its correlation $\rho_{iM}$ to the market demand. In the extreme case of $\rho_{iM} < 0$, $D_i$ is negatively correlated with the market demand. Surprisingly, except for paying for mean demand, tenant $i$ actually earns a *bonus* of $-\theta\sigma_i\rho_{iM} > 0$. The reason is that tenant $i$ serves as a risk neutralizer: whenever market demand has a random increase, $D_i$ will decrease to release occupied resources to accommodate the market surge. This helps the broker save the bandwidth reservation and hedge under-provision risks.

## VI. Trace-Driven Simulations

In this section, we simulate a broker-assisted cloud bandwidth trading system, and conduct trace-driven performance evaluation. We use workload traces collected from UUSee [10], an operational large-scale VoD service based in China. The dataset contains the bandwidth demand in UUSee video channels sampled every 10 minutes during 2008 Summer Olympics. We consider two 800-minute time spans in the traces, time periods 702—780 and 1562—1640, containing 91 and 176 concurrent video channels, respectively. We let each video channel $i$ represent the $i$th tenant of the cloud services, with the same demand $D_i$ as the channel's demand.

In our system, bandwidth reservation and trading are carried out online every $\Delta t = 10$ minutes. We believe a 10-minute frequency is close to the fastest frequency that the system can react to demand changes, since historical demand data is available for free from cloud monitoring services such as Amazon CloudWatch at a 5-minute frequency [5]. Before time $t$, the broker should have obtained estimates about expected demands $\boldsymbol{\mu}_t = [\mu_{1t}, \ldots, \mu_{Nt}]$ and the demand covariance matrix $\boldsymbol{\Sigma}_t = [\sigma_{ijt}]$ for all VoD providers in the coming period $[t, t + \Delta t)$. Such statistics can be predicted accurately based on historical demand data, since VoD demand follows repeated daily patterns and is highly predictable [6], [7], [11]. Some established time series forecasting tools in econometrics (e.g., [12], [13]) may be used for prediction. In particular, we implement an online version of seasonal ARIMA and GARCH models introduced in [6], [7] to predict $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}_t$. Once $\boldsymbol{\mu}_t$ and $\boldsymbol{\Sigma}_t$ are obtained, the broker maximizes its profit by solving (22), making load direction decision $\mathbf{W}^*$ and bandwidth reservation $A_s$ from each cloud provider $s$. The broker serves as a gateway that routes $w_{si}^*$ portion of VoD provider $i$'s incoming video requests to cloud provider $s$ for service during the period $[t, t + \Delta t)$. The above process is repeated for the next period $[t + \Delta t, t + 2\Delta t)$.

To evaluate the benefit of such a system, we quantify the equilibrium point in the market. Fig. 3 evaluates the saving on cloud bandwidth resources in equilibrium market. Due to multiplexing, the aggregate bandwidth $\sum_s A_{st}$ booked by the broker is less than the aggregate bandwidth $\sum_i B_{it}$ needed if each tenant books bandwidth individually. We set the risk factor as $\epsilon = 5\%$ and $\epsilon = 1\%$ for time periods 702—780 and 1562—1640, respectively, which results into an actual under-provision probability of $\epsilon' = 5.06\%$ and $\epsilon' = 1.27\%$ in Fig. 3(a) and Fig. 3(b), respectively. Although $\sum_i B_{it}$ always exceeds the real aggregate demand $\sum_i D_{it}$, it does not mean there is no outage when booking individually: each individual tenant is still subject to an under-provision probability $\epsilon$. In general, the broker can save bandwidth reservation by more than 30% on average with controllable quality risks.

More importantly, we further evaluate the price discount each tenant can enjoy. At time $t$, tenant $i$ pays \$ $P_{it}^0(1) = \mu_{it} + \theta\sigma_{it}$ if reserving bandwidth individually, but only pays \$ $P_{it}^*(1) = \mu_{it} + \theta\sigma_{it}\rho_{iMt}$ in the equilibrium market if it buys bandwidth guarantees from a broker. The discount it receives at time $t$ is $1 - P_{it}^*(1)/P_{it}^0(1)$. Fig. 4 plots the histogram of discounts $1 - P_{it}^*(1)/P_{it}^0(1)$ for all $i$ and over all time periods $t$. It shows that the mean price discount depends on demand statistics and even exceeds 50% during the second time span.

An interesting finding is that the discount is over 100% in some rare cases, which means that the payment $P_{it}^*(1) = \mu_{it} + \theta\sigma_{it}\rho_{iMt}$ of some tenant $i$ at some point $t$ is *negative*. We call such VoD providers "bonus winners," since the broker would even not mind paying to have them in the system. As pointed out in Sec. V, this is because "bonus winners" have demand negatively correlated to the market and thus are *risk neutralizers*: whenever there is an increase in market demand, the demand of "bonus winners" will drop to make resources available for other tenants in need of resources.

## VII. Related Work

Cloud bandwidth reservation is becoming technically feasible. There have been proposals on datacenter engineering to offer bandwidth guarantees for egress traffic from virtual machines (VMs) [2], or to connect the VMs of the same tenant in a virtual network with bandwidth guarantees [2], [3]. These advances have made the cloud more attractive to bandwidth-intensive applications such as VoD and video gaming. Virtualization techniques particularly for supporting cloud-based IPTV services are also being developed by major U.S. VoD providers such as AT&T [14]. Furthermore, video demand forecasting techniques have been proposed, such as the non-stationary time series models introduced in [6]–[8], and video access pattern extraction via principal component analysis in [11]. These prediction methods help to estimate the amount of resources to be reserved.

Cloud brokers, e.g., Zimory [15], have recently emerged as intermediators connecting buyers and sellers of computing resources. The engineering aspects of using brokerage to interconnect clouds into a global cloud market have been discussed in [16]. We propose a new type of cloud brokerage that multiplexes bandwidth reservations to save cost while providing quality guarantees to customers. Amazon Cluster Compute (as of 2011) [17] allows tenants to reserve, at a high cost, a dedicated 10 Gbps network with no multiplexing. Instead of over-provisioning a fixed amount of capacity, our proposed broker dynamically books resource in adaption to demand changes, exempting tenants from demand estimation, for which they have no expertise. We aim to find the pricing policies under which a selfish broker can also enhance cloud resource efficiency and save money for VoD providers. Different from usage-based pricing [4], our proposed pricing policy depends on demand statistics such as burtiness and correlation.

The idea of statistical multiplexing has been empirically evaluated for a shared hosting platform in [18]. VM consolidation with independent random bandwidth demands has also been considered in [19]. In contrast, our work exploits the unique characteristics of user-oriented applications. We leverage the fact that VoD demand is *fractionally splittable* into video requests, which can be optimally directed to different clouds and statistically mixed toward workload consolidation.

## VIII. Conclusions and Future Work

In this paper, we consider the scenario that multiple VoD providers make reservations for bandwidth guarantees from cloud service providers in order to support continuous media streaming. VoD providers attract inter-correlated random demands that can be directed to multiple cloud providers, subject to their bandwidth capacities. We introduce a profit-making *broker* that statistically mixes demands based on anti-correlation while controlling the quality risk. We study how each VoD provider should be charged in such a market.

The region of all good pricing policies is characterized, such that making a profit at the broker is equivalent to optimizing cloud resource efficiency. In a free market where VoD providers can negotiate the bandwidth price with the broker, we prove that the prices will converge to a unique Nash equilibrium, which forms the lower border of the good pricing region. Furthermore, the equilibrium bandwidth price of a VoD provider critically depends on its demand burstiness and correlation to the market demand. Trace-driven simulations verify that the presence of a broker can lower the market price for cloud bandwidth reservations by around 50% on average, and save cloud resources by over 30% given the same demand.

Interesting directions for future work include developing robust methods to deal with prediction inaccuracy and practical schemes to facilitate fast convergence to the market equilibrium, as well as extending this cloud pricing theory to other applications with predictable workload, such as video gaming.

## References


[1] "Four Reasons We Choose Amazon's Cloud as Our Computing Platform," *The Netflix "Tech" Blog*, December 14 2010.

[2] C. Guo, G. Lu, H. J. Wang, S. Yang, C. Kong, P. Sun, W. Wu, and Y. Zhang, "SecondNet: a Data Center Network Virtualization Architecture with Bandwidth Guarantees," in *Proc. of ACM International Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, 2010.

[3] H. Ballani, P. Costa, T. Karagiannis, and A. Rowstron, "Towards Predictable Datacenter Networks," in *Proc. of SIGCOMM'11*, Toronto, ON, Canada, 2011.

[4] M. Armbrust, A. Fox, R. Griffith, A. D. 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*, vol. 53, no. 4, pp. 50–58, 2010.

[5] "Amazon Web Services," http://aws.amazon.com/.

[6] D. Niu, Z. Liu, B. Li, and S. Zhao, "Demand Forecast and Performance Prediction in Peer-Assisted On-Demand Streaming Systems," in *Proc. of IEEE INFOCOM Mini-Conference*, 2011.

[7] D. Niu, B. Li, and S. Zhao, "Understanding Demand Volatility in Large VoD Systems," in *Proc. of the 21st International workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, 2011.

[8] D. Niu, H. Xu, B. Li, and S. Zhao, "Quality-Assured Cloud Bandwidth Auto-Scaling for Video-on-Demand Applications," in *Proc. of IEEE INFOCOM '12*, Orlando, Florida, 2012.

[9] D. Niu, C. Feng, and B. Li, "A Theory of Cloud Bandwidth Pricing for Video-on-Demand Providers," University of Toronto, Tech. Rep., 2011. [Online]. Available: http://iqua.ece.toronto.edu/~bli/papers/vPricing.pdf

[10] Z. Liu, C. Wu, B. Li, and S. Zhao, "UUSee: Large-Scale Operational On-Demand Streaming with Random Network Coding," in *Proc. of IEEE INFOCOM*, 2010.

[11] G. Gürsun, M. Crovella, and I. Matta, "Describing and Forecasting Video Access Patterns," in *Proc. of IEEE INFOCOM Mini-Conference*, 2011.

[12] G. E. P. Box, G. M. Jenkins, and G. C. Reinsel, *Time Series Analysis: Forecasting and Control*. Wiley, 2008.

[13] T. Bollerslev, "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, vol. 31, pp. 307–327, 1986.

[14] V. Aggarwal, X. Chen, V. Gopalakrishnan, R. Jana, K. K. Ramakrishnan, and V. A. Vaishampayan, "Exploiting Virtualization for Delivering Cloud-based IPTV Services," in *Proc. of IEEE INFOCOM Workshop on Cloud Computing*, 2011.

[15] "Zimory Cloud Computing," http://www.zimory.com/.

[16] 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. 599–616, 2009.

[17] "Amazon Cluster Compute, 2011," http://aws.amazon.com/ec2/hpc- applications/.

[18] B. Urgaonkar, P. Shenoy, and T. Roscoe, "Resource Overbooking and Application Profiling in Shared Hosting Platforms," in *Proc. of USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2002.

[19] M. Wang, X. Meng, and L. Zhang, "Consolidating Virtual Machines with Dynamic Bandwidth Demand in Data Centers," in *Proc. of IEEE INFOCOM 2011 Mini-Conference*, 2011.