

New Bandwidth Sharing and Pricing Policies to Achieve A Win-Win Situation for Cloud Provider and Tenants

Haiying Shen and Zhuozhao Li
Department of Electrical and Computer Engineering
Clemson University, Clemson, SC 29631
Email: {shenh, zhuozhl}@clemson.edu

Abstract—For predictable application performance or fairness in network sharing in clouds, many bandwidth allocation policies have been proposed. However, with these policies, tenants are not incentivized to use idle bandwidth or prevent link congestion, and may even take advantage of the policies to gain unfair bandwidth allocation. Increasing network utilization while avoiding congestion not only benefits cloud provider but also the tenants by improving application performance. In this paper, we propose a new pricing model that sets different unit prices for reserved bandwidth, the bandwidth on congested links and on uncongested links, and makes the unit price for congested links proportional to their congestion degrees. We use game theory model to analyze tenants' behaviors in our model and the current pricing models, which shows the effectiveness of our model in providing the incentives. With the pricing model, we propose a network sharing policy to achieve both min-guarantee and proportionality, while prevent tenants from earning unfair bandwidth. We further propose methods for each virtual machine to arrange its traffic to maximize its utility. As a result, our solution creates a win-win situation, where tenants strive to increase their benefits in bandwidth sharing, which also concurrently increases the utilities of cloud provider and other tenants. Our simulation and trace-driven experimental results show the effectiveness of our solution in creating the win-win situation.

I. INTRODUCTION

Cloud computing attracts many enterprises (e.g., Dropbox, Facebook video storage) to migrate their business or services to the clouds without the need to build their own datacenters. Cloud provider (provider in short) multiplexes computation, storage and network resources among different tenants, enabling them to independently run their own jobs on the cloud. Nowadays, on the Infrastructure as a Service (IaaS) (e.g., Amazon EC2), the resources are charged based on the renting time period of virtual machines (VMs) and VM types (with different CPU and memory storage). Though the CPU and memory storage of a VM are dedicated resources to a tenant, each network link is shared among tenants, which makes it non-trivial to guarantee the provision of a certain bandwidth to a tenant. Current best-effort bandwidth provision is insufficient to guarantee the quality-of-service to tenants (i.e., satisfy Service Level Agreement (SLA)). Congested links lead to slow traffic rate, which not only degrades the performance of tenants' applications but also increases their cost due to longer VM usage.

Previous research studied the problem of bandwidth allocation among different tenants. Popa *et al.* [1] indicated that a

desirable allocation solution should meet three requirements: *min-guarantee*, *high utilization* and *network proportionality*, which however are difficult to achieve simultaneously due to their tradeoffs. *Min-guarantee* means guaranteeing the minimum bandwidth that tenants expect for each VM, irrespective of the network utilization of other tenants. It is essential for predictable network performance [2], [3] and a lack of it would impede cloud adoption by applications (e.g., transaction processing web applications [4] and video-on-demand (e.g., YouTube)). *High utilization* means maximizing network utilization in the presence of unsatisfied demands. This means an application can use the idle bandwidth, which shortens job completion time (that benefits tenants) and enables more jobs to be deployed in the infrastructure (that increases the provider's revenue). *Network proportionality* means that network resources allocated to tenants are proportional to their payments, which aims to achieve fairness between tenants.

Many bandwidth allocation policies [1], [5]–[10] have been proposed to achieve min-guarantee or network proportionality. However, they cannot achieve high utilization to benefit both the provider and the tenants; tenants would try to gain more benefits at the cost of the provider or other tenants. For example, a tenant tries to compete bandwidth in a more congested link even though it can use an idle link; it may also purposely change its actual bandwidth demand to receive more bandwidth allocation, which reduces network utilization [1]. Thus, a significant problem is how to achieve a win-win situation, where tenants strive to increase their utility in bandwidth sharing, which also concurrently increases the network utilization, profit and SLA conformance of the provider. However, no previous research has studied this problem.

To address this problem, we propose a new bandwidth pricing model in this paper. Unlike the previous works that allocate bandwidth based on tenant payment, our model determines each tenant's payment based on allocated bandwidth. Thus, network proportionality is achieved since the allocated bandwidths of tenants are always proportional to their payments. In the current flat-rate per VM payment model, tenants compete for bandwidth since the consumed bandwidth does not affect payment. In our pricing model, the consumed bandwidth determines the payment, which encourages tenants to be cooperative in bandwidth sharing to reduce their payment.

Our pricing model considers three parts in determining the payment of a tenant (P_{t_i}): min-guarantee bandwidth (M_{t_i}),

consumed congested bandwidth ($B_{t_i}^c$) and consumed uncongested bandwidth ($B_{t_i}^u$) of all VMs of the tenant; $P_{t_i} = \alpha M_{t_i} + \beta B_{t_i}^c + \gamma B_{t_i}^u$ ($\alpha > \beta > \gamma$), where α , β and γ are unit prices and β is proportional to link congestion degree. Therefore, to reduce payment, a tenant will buy the minimum bandwidth on a VM based on its real minimum demand, which reduces the provider's reserved but unused resources and increases network utilization. Also, a tenant will try to use idle bandwidth and avoid more congested bandwidth, which increases network utilization and decreases SLA violations. High network utilization in turn increases the performance of applications and hence benefits the tenants.

Our bandwidth allocation strategy first satisfies the min-guarantee, and then achieves proportionality (network, congestion or link proportionality [1]) on the residual bandwidth. With our pricing model, tenants are disincentivized to take advantage of the allocation policies (or even cheat) for more bandwidth which would otherwise lead to low network utilization [1]. As a result, our solution simultaneously achieves the above-stated three requirements – an unsolved problem in previous research. We also propose methods for each VM to arrange its traffic flows to maximize its utility.

The contributions of our paper are summarized as follows:

- We use the game theory model to analyze the behaviors of tenants in the current pricing models and allocation policies. We find that tenants may try to gain more benefits at the cost of the provider and other tenants.
- We propose a pricing model to create a win-win situation, where tenants try to gain more utility which also concurrently increases the benefits of other tenants and the provider. Our analysis on the tenant behaviors confirms the advantages of our pricing model.
- We propose a network sharing policy to achieve both min-guarantee and different types of proportionality, while preventing tenants from earning unfair bandwidth.
- We propose a traffic flow arrangement policy for each VM to determine the links to traverse its traffic flows to their destinations, and the destination VMs for flows without fixed destinations in order to maximize its utility.

Consequently, with our solution, the competitive cloud environment is transformed to a cooperative environment, which increases the benefits of both the provider and tenants, and helps create a harmonious ecosystem. Our experimental results verify the advantages of our solution. The rest of this paper is structured as follows. Section II presents a concise review of related work. Section III analyzes the behaviors of tenants in current bandwidth allocation and pricing model and shows that competitive bandwidth sharing does not benefit either tenants or the provider. Section IV presents our proposed policies, and analyzes their effectiveness in increasing the benefits of both sides. Section V presents the performance of our proposed policies in comparison to previous bandwidth allocation strategies. Finally, section VI concludes this paper with remarks on future work.

II. RELATED WORK

Recently, several bandwidth allocation mechanisms have been proposed that assign weights to VMs or tenants/services for bandwidth competition in clouds. Some works [8], [9] provide proportional network sharing based on VM weight (or payment), while other works [5]–[7] provide minimum bandwidth guarantee by reserving bandwidth.

Seawall [8] is a hypervisor-based mechanism to enforce the bandwidth allocation in each congested link based on the weights of the VMs which are communicating along that link. Netshare [9] is a statical multiplexing mechanism that enables tenants to receive constant proportionality throughout the cloud. Popa *et al.* [1] proposed PS-L and PS-N to achieve proportionality. PS-L achieves *link proportionality*, in which the allocated bandwidth in a congested link is proportional to the sum of the weights of a tenant's VMs that communicate through the link. PS-N achieves *congestion proportionality*, in which the total allocated bandwidth on congested links of a tenant is proportional to the sum of the weights of a tenant's VMs. Although these policies can achieve proportionality, they cannot provide min-guarantee for predictable performance.

Popa *et al.* [1] also proposed PS-P to support minimum bandwidth guarantees by assigning the weight of on link between a VM-pair based on the weight of the VM closer to the link. Oktopus [5] and SecondNet [6] use static reservations in the network to achieve minimum bandwidth guarantees. Gatekeeper [7] is a per-VM hose model with work conservation. Guo *et al.* [10] proposed to achieve min-guarantee and then share the residual bandwidth among VM-pairs for link proportionality. However, this policy does not support network proportionality or congestion proportionality.

In all the above works, since bandwidth is allocated based on weight determined by flat-rate payment, all tenants will try to compete for bandwidth, which reduces network utilization and increases SLA violations. Different from these policies, our solution provides utilization incentive, and simultaneously achieves the three aforementioned requirements.

Niu *et al.* [11] proposed a pricing model for pricing cloud bandwidth reservation in order to maximize social welfare. Feng *et al.* [12] utilized the bargaining game to maximize the resource utilization in video streaming datacenters. Wilson *et al.* [13] proposed a congestion control protocol to allocate bandwidth according to flow deadlines, and charge bandwidth usage. Different from these pricing models, our pricing model aims to provide incentives to tenants to use uncongested links to increase network utilization, and prevent congestion to reduce SLA violations, which helps create a win-win situation for both the provider and tenants.

III. COMPETITIVE BANDWIDTH SHARING IN CURRENT POLICIES

A. Problems in Bandwidth Allocation and Our Solutions

We argue that the ultimate objective of the three desired requirements in bandwidth allocation (i.e., min-guarantee, high utilization and network proportionality) is to maximize the




Fig. 1: Failure to achieve high utilization in bandwidth allocation [1].

benefits of the provider and each tenant; that is, increasing the profit of the provider and the performance of tenants' applications based on their payments. Therefore, we should not simply aim to develop a bandwidth allocation policy that can meet a part of the three requirements. Rather, we should develop a policy that can flexibly meet the three requirements and achieve the ultimate goal. Below, we present the unsolved problems in bandwidth allocation indicated in [1] that prevent us from simultaneously achieving these requirements and briefly explain our solutions.

a) Tradeoff Between Min-guarantee and Network Proportionality: Suppose tenant A employs 2 VMs and tenant B employs 10 VMs. We assume the weights of VMs are the same for simplicity. VMs A_1 and B_1 are hosted on the same physical machine (PM) that communicate with other VMs that belong to the same tenant. According to the network proportionality, A_1 receives $2/12$ of the access link, while B_1 receives $10/12$. A_1 's allocation may be lower than its minimum guarantee, failing to satisfy its min-guarantee. Also, tenant B can buy many VMs for B_1 to communicate in order to dominate the link, which would degrade A 's application performance. To address this tradeoff problem, we first satisfy the min-guarantee of each VM and then follow the network proportionality in allocating the residual bandwidth. We also set the highest unit price for the min-guarantee bandwidth, so that tenants will try to limit the minimum bandwidth to their exact needs, which prevents the domination situation to a certain degree.

b) Tradeoff Between High Utilization and Network Proportionality: Consider two tenants A and B , each employing 4 VMs with the same weight. Their flows traverse the same congested link l with capacity C as shown in Figure 1(a). Based on the network proportionality, each tenant receives $C/2$ bandwidth. Now assume VMs A_1 and A_3 start communicating along an uncongested path l_3 (Figure 1(b)). In order to maintain network proportionality, tenant A 's allocation is decreased along link l . If A 's traffic along l is more important than that along path l_3 , A is disincentivized to use path l_3 , which degrades network utilization and also increases the probability of link congestion.

To address this problem, we assign lower unit price to uncongested links than congested links and make the unit price for congested links proportional to the congestion degree. In this way, tenants are incentivized to use uncongested links, and avoid competing for bandwidth in the congested links. The higher the congestion of a link, the lower probability for a tenant to compete for bandwidth on the link. As a result, the network utilization is increased and the congestions are prevented or mitigated, which enhances application perfor-

mance and also increases the provider's profit and reduces SLA violations.

Popa *et al.* [1] suggested that congestion proportionality can achieve utilization incentives but tenants may cheat to gain more bandwidth which reduces network utilization. Since the uncongested links are not considered in bandwidth allocation, tenants are incentivized to use uncongested links. However, a tenant can reduce its demand on purpose to change a congested link to an uncongested link in order to increase its own allocation and reduce others' allocation, which decreases network utilization, as illustrated in the example below.

Assume ϵ is a very small number. In Figure 1(c), if the demand of $B_3 \rightarrow B_4 = \epsilon$, the allocation $A_3 \rightarrow A_4 = C - \epsilon$, and then $B_1 \rightarrow B_2 = C - \epsilon$ and $A_1 \rightarrow A_2 = \epsilon$. Tenant A can purposely change its demands on l_2 to $C - 2\epsilon$. Then, l_2 becomes uncongested and is not considered in congestion proportionality. Finally, tenant A receives $3C/2 - 2\epsilon$ and tenant B receives $C/2 + \epsilon$. The network utilization is decreased from $2C$ to $2C - \epsilon$.

Suppose D_l and C_l denote the total bandwidth demand and capacity on link l , we argue that *congested links* should be defined as the links with $D_l > C_l$ rather than $D_l \geq C_l$ as in [1] and *uncongested links* should be defined as the links with $D_l \leq C_l$. Because when $D_l = C_l$, the link can exactly satisfy the tenants' demands and there is no need for them to compete for bandwidth. With this new definition, a tenant only has incentives to purposely reduce its demand when $D_l > C_l$ to make it $D_l = C_l$, and when $D_l = C_l$ (the link is fully utilized), the tenants have no incentives to reduce their demand. In a congested link, each tenant will check its gain and cost to decide if it should reduce demand to make it $D_l = C_l$. The gain includes more allocation in other congested links and lower payment in our pricing model. Note that instead of preventing tenants from reducing their demands when $D_l > C_l$, we encourage such behavior, because it will not reduce network utilization. In addition, such behavior avoids link congestion and hence increases application performance for tenants and reduces SLA violations of the provider. Though finally tenant A may receive more bandwidth in another congested link, it still needs to pay for this bandwidth in our pricing model, which achieves proportionality.

B. Game Theory Based Analysis on Current Pricing Models

We analyze the behaviors of tenants and the provider using the non-cooperative game theory [14], in which each game player tries to maximize its payoff. We first analyze the current price model in Amazon EC2, where tenants pay a fixed flat-rate per VM for each type of VMs. When a link is congested, a previously proposed bandwidth allocation strategy (min-guarantee, network proportional, congestion proportionality or link proportionality) is used. Currently, the provider supplies bandwidth in the best-effort provision manner. Therefore, we assume that without min-guarantee requirement, the bandwidth provision does not affect the SLA violations, and with this requirement, failures of providing the min-guarantee bandwidth lead to SLA violations.

The utility of the provider (i.e., cloud profit) is the difference between its total revenue and total cost, which includes the cost for consumed bandwidth and for SLA violations. We use N_{V_i} ($1 \leq i \leq m$) to represent the total number of sold type- i VMs, use m to represent the number of VM types in the system and use p_i to denote the payment of a type- i VM. As in [11], we assume tier-1 ISPs charge the provider b for each unit bandwidth actually used. B^a denotes the allocated bandwidth of all tenants and $B_{t_i}^a$ denotes the allocated bandwidth of tenant t_i . M_{v_i} denotes the min-guarantee bandwidth for VM v_i . The min-guarantee bandwidth for tenant t_i (M_{t_i}) is the sum of the min-guarantee bandwidths of t_i 's VMs: $M_{t_i} = \sum_{v_k} M_{t_i, v_k}$. We use $H_{t_i} = M_{t_i} - B_{t_i}^a$ to denote the unsatisfied bandwidth for t_i to meet the min-guarantee requirement. It leads to $F_c(H_{t_i})$ utility loss of the provider caused by the reputation degradation and potential revenue loss. We use $F_{t_i}(H_{t_i})$ to denote the utility loss of tenant t_i due to unfilled demands from clients. With the min-guarantee requirement, reserving bandwidth capacity K will incur a reservation cost of cK [11]. Then, the provider's utility can be represented by:

$$U_c = \begin{cases} \sum_i p_i N_{V_i} - bB^a, & \text{w/o min-g} \\ \sum_i p_i N_{V_i} - bB^a - \sum_{t_i} F_c(H_{t_i}) - cK, & \text{w/ min-g,} \end{cases} \quad (1)$$

in which "min-g" denotes min-guarantee requirement. A tenant's utility can be represented by:

$$U_{t_i} = g_{t_i} B_{t_i}^a - \sum_k p_k N_{V_{k t_i}} - F_{t_i}(H_{t_i}), \quad (2)$$

where g_{t_i} represents the earned utility of each used bandwidth unit and $N_{V_{k t_i}}$ denote the number of type- k VMs bought by tenant t_i .

Based on Equation (1), for the provider, in order to maximize its utility, it needs to increase the number of sold VMs (N_{V_i}), reduce the total used bandwidth (B^a). With min-guarantee, the provider also needs to reduce provision failure on reserved bandwidth (reduce congestion) and reduce reserved bandwidth. Given a certain number of PMs, to increase N_{V_i} , the provider can place many VMs on one PM. To reduce B^a , the provider can employ strategies such as placing the VMs of the same tenant in the same or nearby PMs (which is out of the scope of this paper). Given a certain VM placement, the provider supplies bandwidth in the best-effort manner, and it has no control over B^a . Consequently, it tries to maximize the number of VMs placed in a PM while guarantee the minimum bandwidth for VM and reduce link congestion. Though the provider can use bandwidth allocation policies to achieve different proportionality, it has no control on tenants' bandwidth demand to reduce the link congestion situation. Thus, the provider needs an additional policy for this purpose to increase cloud profit.

Based on Equation (2), in order to increase utility, a tenant tries to receive more $B_{t_i}^a$, buy fewer and less-expensive VMs and also reduce the unsatisfied demand. As a result, tenants will try to be economical when buying VMs and compete for more bandwidth. As explained in Section III-A, in the network proportionality or congestion proportionality policy, the

competition leads to low network utilization, which reduces the utility of the provider and other tenants.

We then analyze the recently proposed pricing model in [11]. Each tenant pays p for every unit bandwidth consumed and pays $k_{t_i} w_{t_i}$ for having w_{t_i} portion of its demand guaranteed. Then, the utilities of the provider and tenant are:

$$U_c = \sum_{t_i} (p B_{t_i}^a + k_{t_i} w_{t_i}) - bB^a - \sum_{t_i} F_c(w_{t_i} D_{t_i} - B_{t_i}^a) - cK, \quad (3)$$

$$U_{t_i} = g_{t_i} B_{t_i}^a - (p B_{t_i}^a + k_{t_i} w_{t_i}) - F_{t_i}(w_{t_i} D_{t_i} - B_{t_i}^a), \quad (4)$$

where $p, g_{t_i} > b$.

Equation (3) indicates that to increase utility, the provider wishes to increase network utilization (B^a) and reduce unsatisfied demands. However, it has no control on bandwidth demands from tenants. Equation (4) shows that to maximize its utility, given a reserved portion, a tenant tends to compete for usage bandwidth in demand. Since the unit price for used bandwidth is the same regardless of the congestion degree of links, tenants tend to compete for more important bandwidth to them, as explained in Figure 1(b).

Both pricing models lead to bandwidth competition among tenants. As explained in Section III-A, though different allocation policies can be used in bandwidth competition, the competition still can lead to low network utilization and reduce the benefits of other tenants and the provider. That is, the pursuit of higher utility of a tenant decreases the utility of the other tenants and the provider. We need a policy to create a harmonious environment where all tenants cooperate to increase their utilities and also concurrently increase the system utility and reduce unsatisfied demands, which not only benefits all tenants but also the provider. To achieve this goal, we propose our pricing model and network sharing policy in the next section and use game theory to analyze their effectiveness.

IV. PROPOSED POLICIES FOR COOPERATIVE BANDWIDTH SHARING

In this section, we present our pricing model that can help achieve high network utilization and also avoid congested links, thus increase application performance and reduce SLA violations. More importantly, this pricing model transforms the competitive environment to a cooperative environment, in which a tenant can receive more benefits by being cooperative than by being non-cooperative.

We assume a multi-path or multi-tree topology [15]–[18], where each VM has multiple links to connect to other VMs. In Figure 2, we only drew the multiple links for A_1 and A_7 as an example for easy readability. As in [1], [10], we consider a hose model [19], where each VM is connected to non-blocking switches by dedicated connection.

A. A New Bandwidth Pricing Model

When a tenant buys VMs, it can specify the min-guarantee of each VM. We use *congested bandwidth* ($B_{t_i}^c$) and *uncongested bandwidth* ($B_{t_i}^u$) to represent tenant t_i 's consumed bandwidth on congested links and on uncongested links, respectively. Then, t_i 's total allocated bandwidth $B_{t_i}^a = B_{t_i}^c +$




Fig. 2: An example of multi-tree topology [16].

$B_{t_i}^c$. We use M_{t_i, v_j} , B_{t_i, v_j}^c and B_{t_i, v_j}^u to represent the minimum guaranteed bandwidth, the congested and uncongested bandwidth of VM v_j of tenant t_i ; $B_{t_i, v_j}^a = B_{t_i, v_j}^u + B_{t_i, v_j}^c$.

We use α , β and γ to denote the unit price of minimum guaranteed bandwidth, congested bandwidth and uncongested bandwidth and $\alpha > \beta > \gamma$. Then, each tenant's payment consists of three parts:

$$P_{t_i} = \alpha M_{t_i} + \beta B_{t_i}^c + \gamma B_{t_i}^u \\ = \alpha \sum_{v_j} M_{t_i, v_j} + \beta \sum_{v_j} B_{t_i, v_j}^c + \gamma \sum_{v_j} B_{t_i, v_j}^u. \quad (5)$$

For tenants, the reserved bandwidth is more valuable than non-reserved bandwidth, because a tenant is guaranteed to receive the reserved bandwidth. Therefore, it should pay more for reserved bandwidth. If its price is low, each tenant would try to buy more minimum bandwidth, which would generate much reserved but unused bandwidths and hence reduce the cloud profit. Reserved bandwidth (K) incurs additional cost of cK to the provider. On the other hand, it reduces the utility loss due to poor performance of applications. Then, to increase profit, the provider should encourage tenants to reserve no more bandwidth than their exact need, which also increases network utilization. Thus, we set α to the highest value among the unit prices, i.e., $\alpha > \beta, \gamma$.

In the ideal situation, each link achieves $D_l = C_l$; i.e., the network is fully utilized and all bandwidth demands are satisfied. Then, both the provider and tenants earn the maximum profit and experience the least utility loss due to unfulfilled demands. To make the system approach the ideal situation, we need to encourage tenants to use uncongested links and avoid using congested links. Accordingly, the unit price (β) of congested bandwidth should be higher than the unit price (γ) of uncongested bandwidth. To tenants, congested bandwidth is more valuable than uncongested bandwidth as they must compete for it. With $\beta > \gamma$, tenants are incentivized to use uncongested links and avoid using congested links to reduce payment.

We define a link's *congestion degree* as $\frac{D_l}{C_l}$. To avoid exacerbating the congestion situation, the tenants should be more strongly disincentivized to use more congested links. Thus, we set a congested link's β to be proportional to its congestion degree: $\beta = \gamma(\min\{\frac{D_l}{C_l}, \delta\})$ ($\frac{D_l}{C_l} > 1$). $\delta > 1$ is used to limit the infinite increase of β .

B. Network Bandwidth Sharing

To consider both min-guarantee and proportionality in a congested link, each VM first receives its min-guarantee, and then receives its share on the residual bandwidth based on




Fig. 3: Communication between VMs. Fig. 4: Bandwidth sharing in a link.

the proportionality allocation policy, which can be network proportionality, congestion proportionality or link proportionality. Let D_{v_i} denote the total demand of VM v_i . Then, we have $D_{v_i} = \sum_{v_k} D_{v_i, v_k}$, where v_k denotes each VM v_i communicates with and D_{v_i, v_k} denotes the traffic demand between VM v_i and v_k . The total bandwidth allocated to VM v_i is denoted by $B_{v_i}^a = \sum_{v_k} B_{v_i, v_k}^a$. Below, we first introduce a method to calculate the min-guarantee bandwidth for a pair of VMs to ensure that the min-guarantee of each VM is guaranteed. Then, we introduce how to calculate the weight of a pair of VMs in bandwidth allocation. Finally, we introduce the entire process of bandwidth requesting and allocation.

VM v_i may communicate with other VMs through a link, as shown in Figure 3. To ensure that $B_{v_i}^a$ satisfies M_{v_i} , v_i 's min-guarantee should be distributed among these VMs and v_j should receive its portion equals to M_{v_j} over the sum of the min-guarantee of all of these VMs, i.e., $M_{v_i} \frac{M_{v_j}}{\sum_{D_{v_i, v_k} \neq 0} M_{v_k}}$.

Similarity, v_i should receive $M_{v_j} \frac{M_{v_i}}{\sum_{D_{v_k, v_j} \neq 0} M_{v_k}}$. Then, we define the min-guarantee of a pair of VM v_i and v_j over a link as:

$$M_{v_i, v_j} = \rho M_{v_i} \frac{M_{v_j}}{\sum_{D_{v_i, v_k} \neq 0} M_{v_k}} + (1 - \rho) M_{v_j} \frac{M_{v_i}}{\sum_{D_{v_k, v_j} \neq 0} M_{v_k}}, \quad (6)$$

where $\rho = 1$ for all links in the tree topology that are closer to v_i than v_j , and $\rho = 0$ for all links closer to v_j than v_i .

Suppose VM v_i demands bandwidth D_{v_i, v_j} to VM v_j on a link. We define: $L_{v_i, v_j} = \min\{D_{v_i, v_j}, M_{v_i, v_j}\}$. If the link has residual bandwidth no less than L_{v_i, v_j} , v_i receives $L_{i, j}$ and there is no competition on the link. Otherwise, each pair of communicating VMs v_i' and v_j' on the link receive their $L_{i', j'}$, and then the residual bandwidth is allocated among the pair of VMs that have unsatisfied demands based on proportionality.

We directly use the min-guarantees of VMs as the weights of VMs in bandwidth allocation. The cloud can also specify different levels of competition ability for the tenants to purchase as the weights of VMs in bandwidth competition. The weight of a pair of VMs v_i and v_j on a link equals:

$$W_{v_i, v_j} = M_{v_i} \frac{M_{v_j}}{\sum_{D_{v_i, v_k} \neq 0} M_{v_k}} + M_{v_j} \frac{M_{v_i}}{\sum_{D_{v_k, v_j} \neq 0} M_{v_k}}, \quad (7)$$

As shown in Figure 3, $\sum_{D_{v_i, v_k} \neq 0} M_{v_k}$ means the sum of the min-guarantees of all VMs that v_i communicates with through this link, across the entire network, and in all congested links in the link proportionality, network proportionality and congestion proportionality policy, respectively.

In the following, we explain the process of bandwidth requesting and allocation with our pricing model. The implementation of the policy can rely on switch support or hypervisors as explained in [1]. When VM v_i declares its bandwidth demand to v_j on a link, if the residual bandwidth is no less than the demand (i.e., accepting v_i 's demand will not congest the link), v_i receives its demanded bandwidth. Otherwise, the link will be congested and the unit price for the bandwidth on this link increases. In this case, v_i can consider if it can reduce its demand to make $D_l = C_l$ based on the traffic's delay tolerance. Recall we assume a multi-path or multi-tree topology. v_i can also seek other alternative uncongested links. If it must make a demand that leads to $D_l > C_l$, the VMs on the link are notified the possible congestion. Since the congestion leads to higher unit price for all VMs on the link, the VMs will try to constrain the link congestion degree. Since some applications are delay-tolerant (e.g., high-throughput computing task) while others are delay-sensitive (e.g., VoD applications), the VMs of delay-tolerant applications can reduce their bandwidth demand if its performance degradation is tolerable. The notified VMs also seek other alternative uncongested links to transmit data. Then, if the link still will become congested, as shown in Figure 4, the L_{v_i, v_j} of each VM should be first satisfied, and the residual bandwidth will be allocated among VMs with $B_{v_k}^a < D_{v_k}$ using our allocation policy. Higher unit price for higher congestion links incentivizes VMs to cooperatively try to reduce demand to reduce congestion in order to reduce their payment. As each tenant tries to avoid congested links and use uncongested links, and also constrain the congestion degrees of congested links, the network utilization is increased and the SLA violations are reduced, which benefits the provider and also the tenants. To encourage tenants to reduce their unimportant bandwidth demands, we can also employ reward policies, which we leave as our future work.

C. Analysis of Our Pricing Model

We use R_c to denote the provider's revenue. In our pricing model, the utility of the provider equals:

$$\begin{aligned} U_c &= R_c - bB^a - \sum_{t_i} F_c(H_{t_i}) - cK \\ &= \sum_{t_i} \{(\alpha M_{t_i} + \beta B_{t_i}^c + \gamma B_{t_i}^u) - b(B_{t_i}^c + B_{t_i}^u) - F_c(H_{t_i}) - cM_{t_i}\} \\ &\geq (\alpha - c) \sum_{t_i} M_{t_i} + (\beta - b)B^a - \sum_{t_i} F_c(H_{t_i}) \end{aligned} \quad (8)$$

G_{t_i} denotes the gain of tenant t_i from receiving bandwidth, P_i denotes the money payment and O_i denotes the utility loss due to unsatisfied demand. The utility of a tenant equals:

$$\begin{aligned} U_{t_i} &= G_{t_i} - P_{t_i} - O_{t_i} \\ &= g_{t_i}(B_{t_i}^c + B_{t_i}^u) - (\alpha M_{t_i} + \beta B_{t_i}^c + \gamma B_{t_i}^u) - F_{t_i}(H_{t_i}). \end{aligned} \quad (9)$$

$(g_{t_i} - \gamma)B_{t_i}^a - \alpha M_{t_i} - F_{t_i}(H_{t_i}) \leq U_{t_i} \leq (g_{t_i} - \beta)B_{t_i}^a - \alpha M_{t_i} - F_{t_i}(H_{t_i})$. Based on Equation (8), for the provider, in order to increase utility, it needs to increase B^a , i.e., increase the network utilization, sell more reserved bandwidth, and decrease unsatisfied demands. As indicated in Section III-B, in the current pricing models, the provider has no control on how much and in which

links the tenants demand bandwidth. Only when a link is congested, the provider allocates the link bandwidth among the tenants based on min-guarantee or proportionality. Therefore, the provider cannot actively try to increase its utility. Using our proposed pricing model, the provider can guide how much and in which links that tenants demand bandwidth to increase their utility, which also increases the provider's utility.

Equation (9) shows that in order to increase utility, a tenant needs to gain more allocated bandwidth, reduce min-guarantee M_{t_i} and reduce unsatisfied demands H_{t_i} . Reducing min-guarantee also reduces the tenant's bandwidth competing ability and hence increases H_{t_i} , resulting in utility decrease. Though increasing M_{t_i} strengthens a tenant's competing ability, it generates a much higher additional payment cost in our pricing model. Therefore, tenants are incentivized to limit their min-guarantee bandwidth to their exact needs. Bandwidth demand prediction [11], [20] can help tenants to estimate their demands. For a given demand $B_{t_i}^a = (B_{t_i}^c + B_{t_i}^u)$, the payment cost is $\beta B_{t_i}^c + \gamma B_{t_i}^u$ ($\beta > \gamma$); β is proportional to link congestion degree. Then, tenants are incentivized to use uncongested links instead of competing on congested links, to use less congested links and constrain link congestion. Consequently, with our network sharing policy, delay-tolerant applications may reduce unimportant demand or use less-important links to avoid bandwidth competition and congested links in order to pay less. The applications that compete for bandwidth are delay-sensitive applications, which however must pay high prices for their competed bandwidth. As a result, the bandwidth is allocated among applications based on their delay tolerance degree; more delay-sensitive applications have higher priority to receive bandwidth and also pay more for this priority, and vice versa. Then, the cloud achieves high overall performance for different delay-tolerant applications. These incentivized tenant behaviors benefit all tenants, increase network utilization and decrease unsatisfied demands, which increases the provider's utility.

In Section III, we presented problems in the previous allocation policies: i) nodes are disincentivized to use uncongested links, and ii) nodes may cheat to gain more bandwidth allocation, both of which decrease network utilization. With our pricing model, tenants are incentivized to use uncongested links because they are cheaper than congested links; so problem i) is resolved. We then see if problem ii) is resolved. First, our definition of uncongestion is $L_l/C_l \leq 1$. If a link satisfies $L_l/C_l = 1$ (i.e., fully utilized), it is not congested, so it will not be considered in congestion proportionality. Thus, tenants on the links with $L_l/C_l = 1$ have no intention to reduce demands as it will not increase their allocation. If a link satisfies $L_l/C_l > 1$, it is congested and will be considered in congestion proportionality. Then, tenants are incentivized to reduce their demands to make the link satisfy $L_l/C_l = 1$ because of the cheaper unit price for uncongested links. This increases the utility of not only tenants but also the provider by reducing unsatisfied demands. Even though the tenant can gain more allocation, it still needs to pay for its gained additional bandwidth, which keeps proportionality.




Fig. 5: Link mapping.




Fig. 6: Destination VM selection.

D. Traffic Flow Arrangement Policy

We use $\mathcal{PM} = \{p_1, p_2, \dots\}$ to denote the set of all PMs in the cloud. Suppose tenant t_i has N_{t_i} VMs. The objective of a tenant is to maximize its utility U_{t_i} . To achieve this objective, we let each VM distributively determine the links for its traffic flows. Each row in the matrix means v_i sends data to each v_j ($1 \leq j \leq N_{t_i}$). For a particular v_j , v_i can have multiple paths to send data to v_j [15]–[18]. We classify the flows that v_i attempts to send out into two types: *destined flow* and *non-destined flow*. A *destined flow* must traverse to a specified VM, while a *non-destined flow* can change its destination. For example, the data that is needed by a task executed in VM v_i is destined flow to v_i . The data of a computing task (e.g., WordCount) that can be assigned to any VM that has enough capacity to handle the task is non-destined flow.

Each path of v_i in multiple paths has an average price for the bandwidth usage. Recall that communicating along uncongested links are more expensive than communicating along congested links, while higher congested links are more expensive than less congested links. Therefore, v_i tries to choose the cheapest link (i.e., least congested) to traverse the flows. Always choosing the least congested link for each flow may not maximize $\sum_{1 \leq j \leq N_{t_i}} U_{v_i v_j}$ globally because the residual bandwidth in the least congested link may be fragmented, which otherwise can support a high-demand flow. Failing to find a link to support a high-demand flow leads to competition. To handle this problem, we propose a link mapping algorithm as shown in Figure 5. v_i orders all flows based on bandwidth demand in descending order and orders the links based on residual bandwidth in ascending order. For each flow, v_i checks the link list in sequence until it finds one that has residual bandwidth no less than the flow’s demand, and assigns this flow to this link. If a flow fails to find such a link, it is assigned to the last link with the maximum residual bandwidth, which can minimize the congestion degree. After each assignment, the two lists are updated. Using this way, the flows are assigned to links that have sufficient bandwidth to support the flow first or that lead to the least unsatisfied demand, thus increasing the utility of both the tenant and the provider. In the latter case, the bandwidth allocation should be conducted. If the flow is non-destined flow, v_i can assign it to any v_k ($1 \leq k \leq N_{t_i}$) that has enough capacity (i.e., CPU and storage) to handle the task of the flow. We introduce the destination VM selection policy to help v_i gain more bandwidth. As shown in Figure 6, v_i can choose a VM that leads to the highest allocation based on Equations (6) and (7): $M_{v_i, v_j} + R \frac{W_{v_i, v_j}}{\sum_{v_m, v_n} W_{v_m, v_n}}$, where R is the residual bandwidth and v_m and v_n are the VM-pairs

that are using the link’s bandwidth. Consequently, each VM selects links and destinations for its flows to use uncongested links and constrain the congestion degree, which increases network utilization and reduce unfilled demands. For each flow transmission, the policy in Section IV-B is used to prevent the occurrence of congestion, and allocate bandwidth based on min-guarantee and proportionality in congested links. The link mapping and destination VM selection policies help better arrange a VM’s multiple flows to increase the utilities of both the tenants and the provider.

V. PERFORMANCE EVALUATION

We use simulation and trace-driven experiments to evaluate the performance of our proposed policies in comparison with the previous bandwidth allocation policy. Specifically, we use PS-P [1] as baseline. As it achieves minimum allocation without our proposed pricing model, we use *min-w/o* to denote it. We use link proportionality as an example in our allocation policy though it can support different proportionalities. In order to see the contributions of our different policies, we use *min-P-w/o* to denote our min-guarantee plus proportionality allocation policy without our pricing model. We use *min-P-w/* to denote our allocation policy with our pricing model, where tenants are only incentivized to use the least congested links, and use *min-P-w/V* to denote the case when tenants further are incentivized to volunteer to reduce unimportant demands.

We use a tree topology as shown in Figure 2 [1], [16] in our experiments. It has 16 servers and 2 tenants A and B. Each tenant has one VM in each of the servers. Each server connects to its switch (named as local link) and three other switches so that each VM has three links (named as foreign links) connecting to other VMs not in the same server. We assume tenant A’s VMs communicate with other tenant A’s VMs using a one-to-one communication pattern (i.e., $A_i \leftrightarrow A_{i+8}$, where $i = 1, 2, \dots, 8$), while tenant B’s VMs communicate with all other tenant B’s VMs (i.e., $B_i \leftrightarrow B_j$, where $i \neq j$).

Tenant B has two sets of VMs: B_i and B'_i ($1 \leq i \leq 16$). Each B_i has 40Mbps minimum-guarantee and has already been allocated with 80Mbps bandwidth on each local link. Each B'_i has minimum bandwidth of 20Mbps and has been allocated with bandwidth randomly chosen from $(0, 100)$ Mbps on A_i ’s selected foreign link. Each of tenant A’s VMs will make requests of bandwidth randomly selected from $[60, 70)$ Mbps and their minimum-guarantees are randomly selected from $[30, 40)$ Mbps. We set $\alpha = 1$ and $\gamma = 0.3$. The changes to these parameters will not affect the relative performance differences between different policies.

For the trace-driven experiments, we deployed Hadoop on a cluster running the WordCount benchmark job and then collected the transmitted and received bytes of each VM every second for 100 seconds. We use this trace in the experiments with the same settings as above. In the experiment, each VM made request for bandwidth based on the trace. We measured the metric each second for 100 seconds and present the median, the 95th and 5th percentiles of the metric results.

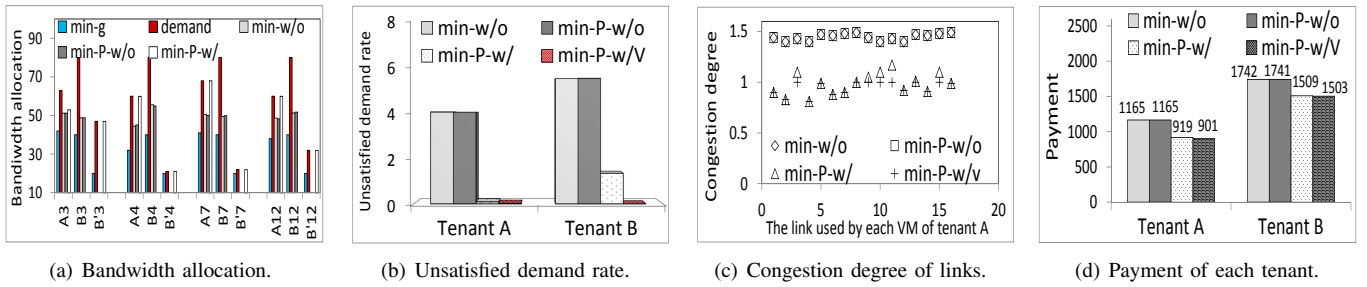


Fig. 7: Experimental results in simulation.




Fig. 8: Trace-driven experimental results.

A. Effectiveness of Pricing and Network Sharing Policies

Figures 7(a) and 8(a) show the min-guarantees and demands of VMs, and their allocated bandwidths in different policies in simulation and trace-driven experiments, respectively. We show the results of 4 pairs of VMs rather than all 16 pairs to make the figure easy to read. We see that in *min-w/o*, tenant A's VMs with larger min-guarantees receive more bandwidth and vice versa. For example, in Figure 7(a), A_4 has a smaller min-guarantee than A_3 , so A_4 receives more bandwidth than A_3 . This is because the min-guarantees of B's VMs on the local links are fixed. Then, tenant A's VM with a higher min-guarantee has a higher weight, so it receives more bandwidth and hence tenant B's VM receives less bandwidth. We see that in *min-w/o*, the VMs of tenant A and tenant B always cannot receive their demanded bandwidth. With the pricing model, in *min-P-w*, tenant A's VMs avoid using congested local links and are incentivized to use the least congested foreign links. From the figure, we see that only B_3 receives bandwidth less than its demand. This is because the least congested foreign link also becomes congested. Then, our allocation policy is employed to allocate bandwidth, which ensures VM's min-guarantee first and then allocates the bandwidth based on the proportionality. These experimental results show the advantage of our price model to incentivize tenants to avoid bandwidth competition and fully utilize bandwidth resources, while ensuring min-guarantee.

Our simulation results also show that when A_{15} does not reduce its unimportant demand to make the link uncongested, both A_{15} and B_{15} cannot receive their demanded bandwidths. When A_{15} does this, both receive their demanded bandwidth. Due to space limit, we do not show this result in a figure. This result shows that our pricing model is incentivizing tenants to reduce their unimportant demands to make links uncongested.

Unsatisfied demand rate for a tenant is defined as the sum of unsatisfied demand percentage of each VM of the tenant;

$\sum_{v_i \in V_t} \frac{D_{v_i} - B_{v_i}^a}{D_{v_i}}$. Figures 7(b) and 8(b) show the unsatisfied demand rate in each method in simulation and trace-driven experiments, respectively. They indicate that without our pricing, *min-w/o* and *min-P-w/o* only achieve different fairness in allocation but cannot prevent bandwidth competition. With our pricing model, *min-P-w* reduces the unsatisfied demand for both tenants A and B because tenants are incentivized to select the least congested links in order to reduce payment. We see that *min-P-w/V* further reduces the unsatisfied demand rate for both tenants A and B. Tenant A volunteers to reduce its unimportant demand, which reduces the unit price for bandwidth consumption and SLA violations.

Figure 7(c) shows the congestion degree of the link used by each VM of tenant A in simulation. Figure 8(c) shows the average congestion degree of links used by tenant A in the trace-driven experiments. We see that without our pricing model (*min-w/o* and *min-P-w/o*), all links are congested. With our pricing model (*min-P-w* and *min-P-w/V*), the congestion degree stays around 1. Since the unit price for uncongested links is lower than that of congested links, tenant A is incentivized to use uncongested links, leading to low link congestion degrees. *min-P-w/V* further reduces the congestion degree of the link used by A_{15} from 1.1 in *min-P-w* to 1, which indicates its effectiveness in maintaining the uncongestion situation by encouraging tenants to reduce unimportant demand.

Figures 7(d) and 8(d) show the total payment of tenant A and tenant B (including VMs B_i and B'_i) in simulation and trace-driven experiments, respectively. We also use our pricing model policy to measure the payment in *min-w/o* and *min-P-w/o* to show the incentives. The figure indicates that if tenants use the less congested links, they can pay less. We also see that *min-P-w/V* produces slightly less payment than *min-P-w* for both tenants because A_{15} reduces its unimportant demand.

B. Effectiveness of Traffic Flow Arrangement Policy

Consider that a VM has three available links (l_1 , l_2 and l_3) with capacities equal to $10Mbps$, $40Mbps$, $100Mbps$, respectively. The VM needs to send data to three other VMs (VM_1 , VM_2 and VM_3) with demands of $10Mbps$, $40Mbps$ and $100Mbps$, respectively. We assume that without our link mapping policy, VM1 will be allocated in priority, then VM2 and VM3. Without the link mapping policy, the allocation is $VM_1 \rightarrow l_3$, $VM_2 \rightarrow l_3$, and $VM_3 \rightarrow l_3$ because l_3 always has the most available bandwidth. With this policy, the allocation is $VM_3 \rightarrow l_3$, $VM_2 \rightarrow l_2$, and $VM_1 \rightarrow l_1$.

TABLE I: Bandwidth allocation with and without the link mapping policy.

	Min-g	Demand	W/o mapping	W/ mapping
VM1	5	10	6.7	10
VM2	20	40	26.7	40
VM3	50	100	66.7	100

TABLE II: Performance with and without the link mapping policy.

	Unsatisfied demand rate	Cong. degree	Payment	Total # of cong. links
W/o mapping	0.3, 0.3, 0.3	1, 1, 1.5	10, 41, 103	1
W/ mapping	0, 0, 0	1, 1, 1	8, 32, 80	0

Table I and Table II show the different metrics. From Table I, we see that the bandwidth demand for all three destination VMs are satisfied with the policy, but are not satisfied without this policy. From Table II, we see that this mapping policy reduces the unsatisfied demand rate, congestion degree, the payment for bandwidth usage, and the number of congested linked. The mapping policy globally considers the bandwidth demands and tries to satisfy each demand while avoids link congestion. More importantly, its payment reduction can incentivize tenants to carefully arrange their flows to different available links, which benefits both the provider and the tenants. Figure 9 shows the unsatisfied bandwidth rate with and without our destination VM selection policy. We see that this policy is effective in reducing the unsatisfied demand.




Fig. 9: Destination VM selection.

VI. CONCLUSIONS

In this paper, we analyzed the behaviors of tenants in the current pricing models and previously proposed bandwidth allocation policies in clouds. These policies incentivize tenants to compete for bandwidth and even gain unfair allocation, which leads to low network utilization and degrades the benefits of both the cloud provider and other tenants. We propose bandwidth sharing and pricing policies to transform the competitive environment to a win-win cooperative environment, where tenants strive to increase their utility, which also concurrently increases the utilities of the cloud provider and other tenants. Specifically, we propose a new bandwidth pricing model, a network bandwidth sharing policy and flow arrangement policies. These policies incentivize tenants to use uncongested links and constrain congestion,

which increases network utilization and reduces unfulfilled bandwidth demands. The bandwidth allocation on congested links also meets the three desired requirements (min-guarantee, high utilization, and network proportionality) – an unsolved problem in previous research. Our experimental results show the effectiveness of our proposed policies. In our future work, we will consider rewarding tenants for reducing demand to maintain the uncongested link states.

ACKNOWLEDGEMENTS

This research was supported in part by U.S. NSF grants IIS-1354123, CNS-1254006, CNS-1249603, CNS-1049947, CNS-0917056 and CNS-1025652, Microsoft Research Faculty Fellowship 8300751.

REFERENCES

- [1] L. Popa, G. Kumar, M. Chowdhury, A. Krishnamurthy, S. Ratnasamy, and I. Stoica. Faircloud: sharing the network in cloud computing. In *Proc. of SIGCOMM*, 2012.
- [2] 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*, 2010.
- [3] J. Schad, J. Dittrich, and J. Quiané-Ruiz. Runtime measurements in the cloud: observing, analyzing, and reducing variance. In *Proc. of VLDB*, 2010.
- [4] D. Kossmann, T. Kraska, and S. Loesing. An evaluation of alternative architectures for transaction processing in the cloud. In *Proc. of SIGMOD*, 2010.
- [5] H. Ballani, P. Costa, T. Karagiannis, and A. Rowstron. Towards predictable datacenter networks. In *Proc. of the SIGCOMM*, 2011.
- [6] 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 CoNEXT*, 2010.
- [7] H. Rodrigues, J. R. Santos, Y. Turner, P. Soares, and D. Guedes. Gatekeeper: supporting bandwidth guarantees for multi-tenant datacenter networks. In *Proc. of WIOV*, 2011.
- [8] A. Shieh, S. Kandula, A. Greenberg, C. Kim, and B. Saha. Sharing the data center network. In *Proc. of NSDI*, 2011.
- [9] T. Lam, S. Radhakrishnan, A. Vahdat, and G. Varghese. Netshare: Virtualizing data center networks across services. Technical report, 2010.
- [10] J. Guo, F. Liu, D. Zeng, J. C.S. Lui, and H. Jin. A cooperative game based allocation for sharing data center networks. In *Proc. of Infocom*, 2013.
- [11] D. Niu, C. Feng, and B. Li. Pricing cloud bandwidth reservations under demand uncertainty. In *Proc. of SIGMETRICS*, 2012.
- [12] Y. Feng, B. Li, and B. Li. Bargaining towards maximized resource utilization in video streaming datacenters. In *Proc. of INFOCOM*, 2012.
- [13] C. Wilson, H. Ballani, T. Karagiannis, and A. Rowstron. Better never than late: meeting deadlines in datacenter networks. In *Proc. of SIGCOMM*, 2011.
- [14] M. J. Osborne and A. Rubinstein. *A course in game theory*. The MIT Press, July 1994.
- [15] C. Raiciu, S. Barre, C. Pluntke, A. Greenhalgh, D. Wischik, and M. Handley. Improving datacenter performance and robustness with multipath tcp. In *Proc. of SIGCOMM*, 2011.
- [16] A. Greenberg, J. R. Hamilton, N. Jain, S. Kandula, C. Kim, P. Lahiri, D. A. Maltz, P. Patel, and S. Sengupta. V12: a scalable and flexible data center network. In *Proc. of SIGCOMM*, 2009.
- [17] C. Guo, G. Lu, D. Li, H. Wu, X. Zhang, Y. Shi, C. Tian, Y. Zhang, and S. Lu. Bcube: A high performance, server-centric network architecture for modular data centers. In *Proc. of SIGCOMM*, 2009.
- [18] J. Mudigonda, P. Yalagandula, M. Al-Fares, and J. C. Mogul. Spain: Cots data-center ethernet for multipathing over arbitrary topologies. In *Proc. of NSDI*, 2010.
- [19] N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. E. van der Merive. A flexible model for resource management in virtual private networks. In *Proc. of SIGCOMM*, 1999.
- [20] D. Xie, N. Ding, Y. C. Hu, and R. R. Kompella. The only constant is change: incorporating time-varying network reservations in data centers. In *Proc. of SIGCOMM*, 2012.