Coordinated Colocation Datacenters for Economic Demand Response *

Nguyen H. Tran, Cuong T. Do, Choong Seon Hong
Kyung Hee University
{nguyenth, dtcuong, cshong}@khu.ac.kr

Shaolei Ren
Florida International University, USA
sren@cis.fiu.edu

Zhu Han
University of Houston, U.S.A
zhan2@uh.edu

ABSTRACT

Demand response of datacenters recently has received increasing efforts due to huge demands and flexible power control knobs. However, very few works focus on a critical segment of datacenter business: multi-tenant colocation. In colocation datacenters, while there exist multiple tenants who manage their own servers, the colocation operator only provides other facilities such as cooling, reliable power, and network connectivity. Therefore, colocation has a unique challenge for the demand response: uncoordinated power management among tenants. To tackle this challenge, we study how to coordinate tenants for economic demand response. We show that there is an interaction between the operator and tenants' strategies, where each side maximizes its own benefit. Hence, we apply a two-stage Stackelberg game to analyze this scenario and derive this game's equilibria. Finally, trace-based simulations are also provided to illustrate the efficacy of our proposed incentive schemes.

1. INTRODUCTION

Datacenters have been consuming billions of kilowatt-hour energy each year to support the explosive cloud services. While their huge energy demands have certainly raised concerns with environmental impacts and operating costs, datacenter's operation is highly-automated (through, e.g., workload and resource management), leading to a great flexibility in energy demand which has been increasingly recognized as a valuable demand response resource that helps to balance grid's power supply and demand. Recent studies have shown that a 20 MW datacenter can be worth up to 5 million dollars of energy storage for power grid, and the whole datacenter demand response market in the U.S. can be billions of dollars [7].

While the existing research focuses on Google-type data centers [5–7, 10], we consider colocation datacenters (colos), which consume nearly 40% of all data center energy [8]. There are many reasons to advocate more research efforts on colos. First, colos' customers diversely include many popular Internet websites (such as Twitter and Wikipedia) and various cloud-computing services (such as Salesforce and Box) [9]. Second, the growth of colos continues increasing sharply: currently there are more than 1200 colos in the U.S. alone [3], and the colos market is expected to grow from current $25 billions to $43 billion in the next five years [1]. Finally, colos are ideal contributors to the demand response programs: (i) colos also have extreme power demands, e.g. colos' demands in New York exceeds 400MW [3]; (ii) colos are often located in urban areas, e.g. Los Angeles.

Instead of fully controlling all facilities, a colo is a shared multi-tenant datacenter where multiple tenants house and fully control their servers while the colo's operator is mainly responsible for facility support such as power, cooling, and network access. Thus, there exists a split-incentive hindrance for colos' demand response: the operator may need reducing energy usage upon the request of LSE in order to receive financial reimbursement, while tenants have little intention to reduce power demand because tenants' billings are based mainly on peak-power subscription with fixed rates, which is independent of their actual usage [2]. The first study that attempts to break the split-incentive issue of colos is [9], though its mechanism is simple and relies on the tenants' best-effort, which can unfulfill reduction targets as well as truthfulness of strategic tenants. A recent work [11] has overcome these issues with a randomized auction mechanism that can guarantee a 2-approximation in social welfare cost. However, both [9] and [11] are based on reverse auction, in which tenants must *voluntarily* submit bids. Since tenants are naturally not concerned in load reduction, treating tenant biddings as voluntary tasks can lead to pessimistic results on the number of participants.

In this work, we study how to coordinate tenants to perform economic demand response. In our proposal, the operator will reward tenants with monetary incentives to perform demand response to a level that can maximize the operator's profit, which can be the financial compensation from LSE or achieving green certificates. Consequently, upon receiving the announced reward from operator, self-optimized tenants will maximize their net utility individually. We model this mechanism as a Stackelberg game and analyze its equilibria. We also propose an algorithm to obtain the optimal solution of the operator's mixed-boolean nonlinear problems. Our key contributions are not only reflected in the efficient performance guarantee but also validated by trace-based simulations. A wide range of numerical case studies demonstrate that our linear-complexity scheme can achieve the same performance as the exhaustive search method for

*This work was supported by the ICT R&D program of MSIP/IITP. [R0126-15-1009, Development of Smart Mediator for Mashup Service and Information Sharing among ICBMS Platform]. Dr. CS Hong is the corresponding author.

Copyright is held by author/owner(s).

the mixed-boolean programming problem.

2. SYSTEM MODEL

We consider a colo-datacenter in which a set of $\mathcal{I} = \{1, \ldots, I\}$ tenants house their servers. Tenant i has M_i homogeneous servers. A tenant with heterogeneous servers can be viewed as multiple virtual tenants, each having homogeneous servers. We consider a one-period demand response, as in [6, 9], where its duration T is controlled by an LSE, e.g. 15 minutes or 1 hour. During a period, the workload arrival rate to tenant i is denoted by λ_i.

If tenant i has no intention to participate in demand response, all of its servers are active and the workload will be distributed to all servers evenly to optimize performance [5]; hence, the energy consumption of this case is $e_i = M_i(p_{i,s} + p_{i,d}\frac{\lambda_i}{M_i\mu_i})T$ [9], where $p_{i,s}$ and $p_{i,a}$ are the static and active powers of each server, respectively, μ_i is a server's service rate measured in terms of the amount of workloads processed per unit time, and $\frac{\lambda_i}{M_i\mu_i}$ is the server utilization with M_i active servers. In contrast, when doing demand response by turning off m_i servers, the energy consumption of tenant i is $e_i' = (M_i - m_i)\big(p_{i,s} + p_{i,d}\frac{\lambda_i}{(M_i-m_i)\mu_i}\big)T$. Therefore, IT-only (e.g., not including cooling) energy reduction by tenant i is

$$\Delta e_i = e_i - e_i' = m_i\frac{p_{i,s}\cdot T}{PUE}, \tag{1}$$

where PUE is the power usage effectiveness measuring the energy efficiency of the colo. In the sequel, we assume $\frac{p_{i,s}\cdot T}{PUE} = 1$ without loss of generality (w.l.o.g.); hence, we will use Δe_i and m_i interchangeably.

Turning servers off can more or less have negative effects to tenants' performance, inducing tenants' costs. We rely on two typical costs that are widely used for tenants: the wear-and-tear cost and Service Level Agreement (SLA) cost [5,9].

Wear-and-tear cost: This cost, which occurs when tenants switch/toggle servers between active and idle states in every period, can be modeled as $\omega_{i,1}\cdot m_i$, where $\omega_{i,1}$ is a monetary weight (i.e., \$/server.)

SLA cost: Since many Internet services hosted in datacenters are sensitive with response/delay time, the SLA cost can be viewed proportionally to tenants' average response time. By using the M/M/1 queue model, the average response time of each tenant i's workload is $\frac{1}{\mu_i - \frac{\lambda_i}{M_i-m_i}}$. The total SLA cost of a tenant can be modeled as $\omega_{i,2}\cdot d_i(m_i)$ where $d_i(m_i) = \frac{\lambda_i}{\mu_i - \frac{\lambda_i}{M_i-m_i}}$, and $\omega_{i,2}$ is a monetary weight (i.e., \$/delay.)

Therefore, tenant i's total cost when turning m_i servers off is

$$C_i(m_i) = \omega_{i,1}\cdot m_i + \omega_{i,2}\cdot d_i(m_i). \tag{2}$$

3. INCENTIVE MECHANISM FOR COLOS' ECONOMIC DEMAND RESPONSES

In this section, we first introduce the economic demand response of colos. We then study this scenario using the Stackelberg game.

3.1 Economic Demand Response: a Two-stage Stackelberg Game Approach

Economic demand response programs generally indicate how customers can actively respond to price signals [4]. For example, during peak times with high wholesale prices, the customers (i.e. colos), who receive signaling from the LSE, can reduce their consumption to receive some economic benefits corresponding to the amount of energy reduction. Since the reduction volume is not necessarily fixed, many customers find this program appealing due to its flexibility.

In this scenario, even though a colo can freely determine a desired reduction volume, its operator cannot directly control the tenants' servers to proceed the demand response. Therefore, the operator's purpose is to incentivize tenants to reduce their energy up to a level that can maximize the operator's benefit. Consequently, upon receiving the announced reward from the operator, rational tenants will maximize their own profits individually. Observing this hierarchical structure between the operator and tenants, we tackle this economic demand response for colos by using a Stackelberg game approach. The strategies of players in each stage of this game will be presented sequentially.

Tenants (Stage II). Since the operator is the leader that has a first-move advantage, it will first announce a reward rate r (e.g., \$/kWh) that it is willing to pay tenants for turning off their servers. Given r, at Stage II, each rational tenant i's strategy is to choose a number of turned-off servers m_i that will maximize its net utility as follows

$$\underset{m_i}{\text{maximize}} \qquad u_i(m_i, r) = rm_i - C_i(m_i) \tag{3}$$

$$\text{s.t.} \qquad m_i \geq 0. \tag{4}$$

Since the number of servers can be very large, e.g. thousands, we can relax m_i as a continuous variable [6]. We have

$$C_i^{''}(m_i) = \frac{2\lambda_i^2\mu_i\omega_{i,2}}{\big((M_i - m_i)\mu_i - \lambda_i\big)^3}, \tag{5}$$

which means $C_i(m_i)$ is a strictly convex function when tenant i's workload is less than its service rate, i.e., $C_i^{''}(m_i) > 0$ when $\frac{\lambda_i}{M_i-m_i} < \mu_i$. We further relax the feasible constraint $0 \leq m_i \leq M_i$ to (4), which has no effect to problem (3) since its feasible solutions are always strictly less than M_i (i.e., $C_i(m_i') = \infty$, $m_i' \geq M_i$). Then, since $u_i(m_i)$ is strictly concave, there exist a unique solution $m_i^*(r)$, $\forall i$, for a given r in Stage II.

Operator (Stage I). Knowing that each tenant i's strategy will be $m_i^*(r)$, the operator's strategy is to choose the optimal $r^\star$ of the following profit maximization problem

$$\max_{r\geq 0}. \quad \mathcal{U}(r, \{m_i^*\}) = U\Big(\textstyle\sum_{i\in\mathcal{I}} m_i^*(r)\Big) - r\textstyle\sum_{i\in\mathcal{I}} m_i^*(r), \tag{6}$$

where $U(\cdot)$ is the colo-datacenter's utility, which can represent a financial compensation from LSE or a green certificate achieved with respect to its energy reduction, balanced with the cost spent for incentiving tenants $r\sum_{i\in\mathcal{I}} m_i^*(r)$. Even though we have no assumption on a specific utility function for our proposal, some typical candidates are provided for case studies in Section 4.

Stackelberg Equilibrium. Denote a solution of the operator's profit maximization by $r^\star$, we have the following definition

DEFINITION 1. *$(r^\star, \{m_i^\star\})$ is a Stackelberg equilibrium if it satisfies the following conditions for any values of r and*

$\{m_i\}$

$$\mathcal{U}(r^\star, \{m_i^\star\}) \geq \mathcal{U}(r, \{m_i^\star\}), \quad (7)$$

$$u_i(m_i^\star, r^\star) \geq u_i(m_i, r^\star), \forall i. \quad (8)$$

Next, we use the backward-induction method to analyze the Stackelberg equilibria.

3.2 Stackelberg Equilibrium: Analysis and Algorithm

By first-order condition $\frac{\partial u_i}{\partial m_i} = r - C_i'(m_i) = 0$, we have the unique solution m_i^* of tenant i for a given r as follows

$$m_i^*(r) = [f_i(r)]^+ := \left[M_i - \rho_i \left(1 + \sqrt{\frac{\omega_{i,2}}{r - \omega_{i,1}}} \right) \right]^+, \forall i, \quad (9)$$

where $[x]^+ = \max\{x, 0\}$, and $\rho_i := \frac{\lambda_i}{\mu_i}$.

Then, by substituting (9) into (6), the operator's problem is formulated as follows

$$\begin{aligned} \underset{r}{\text{maximize}} \quad & U\left(\sum\nolimits_{i\in\mathcal{I}} [f_i(r)]^+\right) - r\sum\nolimits_{i\in\mathcal{I}} [f_i(r)]^+ \quad (10) \\ \text{s.t.} \quad & r \geq 0. \end{aligned}$$

We see that due to the operator $[.]^+$, problem (10) is non-convex. Specifically, if we define a new variable

$$z_i = \begin{cases} 1, & r > \kappa_i; \\ 0, & \text{otherwise}, \end{cases} \quad (11)$$

where

$$\kappa_i := \omega_{i,1} + \frac{\omega_{i,2}\rho_i^2}{(M_i - \rho_i)^2}, \quad (12)$$

then we see that $m_i^*(r) > 0$ when $z_i = 1$ and $m_i^*(r) = 0$ when $z_i = 0$. Therefore, problem (10) is equivalent to

$$\begin{aligned} \underset{r, \{z_i\}_{i\in\mathcal{I}}}{\text{maximize}} \quad & U\left(\sum\nolimits_{i\in\mathcal{I}} z_i \cdot f_i(r)\right) - r\sum\nolimits_{i\in\mathcal{I}} z_i \cdot f_i(r) \quad (13) \\ \text{s.t.} \quad & r \geq 0, \\ & z_i \in \{0, 1\}, \forall i. \end{aligned}$$

We see that problem (13) is a mixed-boolean programming, which we may acquire an exponential-complexity effort (i.e. 2^I configurations of $\{z_i\}_{i\in\mathcal{I}}$) to solve by the exhaustive search. However, by unveiling its special structure, we propose an algorithm, namely Algorithm 1, that can find the solutions of problem (13) with linear complexity as follows.

Proposition 1. *Algorithm 1 can solve the Stage-I's equivalent problem* (13) *with linear complexity.*

Proof. Since the tenants are sorted according to increasing κ_i (line 1), when the sufficient condition $r > \kappa_i$ satisfies, we have $z_j = 1$, $\forall j \leq i$. In this case, the operator's problem (13) becomes (14), which is a single-variable and continuous problem and can be solved efficiently using any numerical methods (e.g., bisection, Newton, etc.) (lines 1-4). Therefore, we assume that (14) is available first, then find its solutions and only keep those satisfying the sufficient condition (line 5). By successively solving (14) and checking the sufficient condition (lines 5-8), we cover all possible cases of the equivalence between problem (13) and (14). Finally, we just compare and pick solutions that result in the highest operator's profit (line 10). Clearly, with a single loop, Algorithm 1 has the complexity $O(cI)$, where c is the complexity to solve problem (14). □

Algorithm 1 Operator's Revenue Maximizer

1: Sort tenants according to $\kappa_1 < \kappa_2 < \ldots < \kappa_I$.
2: $\mathcal{A} = \{\}$, $\mathcal{B} = \mathcal{I}$, $j = I$;
3: **while** $j > 0$ **do**
4: Find the solutions r_j of the following problem

$$\max_{r \geq \kappa_1} U\left(\sum\nolimits_{i\in\mathcal{B}} f_i(r)\right) - r\sum\nolimits_{i\in\mathcal{B}} f_i(r) \quad (14)$$

5: **if** $r_j > \kappa_j$ **then** $\mathcal{A} = \mathcal{A} \cup \{r_j\}$;
6: **end if**
7: $\mathcal{B} = \mathcal{B} \setminus j$;
8: $j = j - 1$;
9: **end while**
10: Return $r_j \in \mathcal{A}$ with highest optimal values of (14).

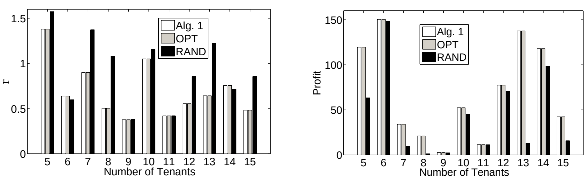

(a) (b)

Figure 1: Comparison of three schemes in economic demand response with utility U is a log function: a) Reward rates, b) Operator's profit.

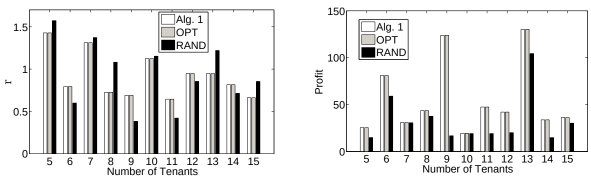

(a) (b)

Figure 2: Comparison of three schemes in economic demand response with utility U is a linear function: a) Reward rates, b) Operator's profit.

Denoting the Algorithm 1' outputs $r^\star$ (which can be multiple values) and $m_i^\star = m_i^*(r^\star)$, we have the straightforward result, whose proof is omitted due to limited space.

Theorem 1. *The Stackelberg equilibria of colos' economic demand response are the set of pairs* $(r^\star, \{m_i^\star\})$.

Based on these equilibria analysis, we next examine how to implement the Stackelberg game-based mechanism.

3.3 Implementation Operations

The main operation of colos' economic demand response can be implemented in the following order. First, each self-optimized tenant submits its best response (9) to the operator. Then, after collecting all of these best responses, the operator solves its profit maximization (6) using Algorithm 1 to achieve $r^\star$ and broadcast this $r^\star$ to all tenants. Finally, based on this $r^\star$, each tenant will correspondingly turn $m_i^\star$ servers off.

4. SIMULATION RESUTLS

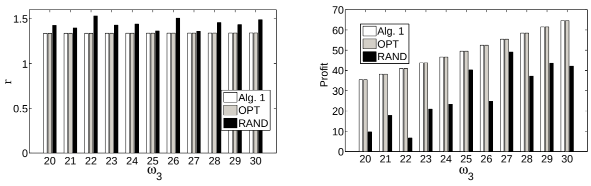

(a) (b)

Figure 3: Comparison of three schemes in economic demand response vs. monetary weight ω_3: a) Reward rates, b) Operator's profit.

We consider a colo-datacenter with varying number of tenants for performance evaluation, and each tenant i has a number of maximum servers M_i that varies uniformly from 3,000 to 10,000 representing for heterogeneous tenant's business. The wear-and-tear and delay cost weights, $\omega_{i,1}$ and $\omega_{i,2}$, respectively, also are uniformly distributed on $[0.1, 3]$, which captures a wide range of tenants' cost sensitivity. We uses two basic traces "MSR" and "FIU" [10] to generate synthetic workloads for tenants. Each tenant's workload is normalized with respect to its service rate μ_i, which is set to 1000 jobs/s.

We compare the performance of Algorithm 1 (Alg. 1) with two baselines. The first baseline, named OPT, is the optimal solutions of problem (10) using exhaustive search. The second baseline, called RAND, is a random price ν^{rand} uniformly distributed in $[\min_i\{C_i'(0)\}, \max_i\{C_i'(0)\}]$ to enable feasible solutions, which represents a simple but inefficient schemes. When $U = \omega_3 \log\left(1 + \sum_{i\in I} m_i^*(r)\right)$, where ω_3 is uniformly distributed on $[0.2, 50]$ and log term reflects the diminishing return on the amount of reduced load, we show the values of the reward rates of different schemes and the corresponding operator's profit in Figs. 1a and 1b, respectively. When $U = \omega_4\left(\sum_{i\in I} m_i^*(r)\right) + \omega_5$, where ω_4 and ω_5 are uniformly distributed on $[1, 2]$ and $[5, 10]$, respectively, we show the operator's reward rate and profit of three schemes in Figs. 2a and 2b, respectively. Since the operator can have a wide range of possible utility values depending on many factors such as LSE's reimbursement, peak or non-peak demand response period, and colo's characteristics, we have the freedom to choose the weight parameters such that feasible solutions exist. In all scenarios, we see that while Alg. 1 and OPT achieve the same performance in all figures, the scheme RAND is not as efficient as the others.

We also examine the effect of ω_3 in Fig. 3. We can see that ω_3 has an impact to the operator's profit. Specifically, the optimal operator's profit increases linearly when ω_3 increases, while the optimal reward rates are unchanged.

5. CONCLUSIONS

In this paper, we addressed the demand response of crucial but less-studied segment of datacenter market: colocation datacenters (colos). We tackled the split-incentive hindrance between colo's tenants and operator in the context of economic demand response, which is based on top of the two-stage Stackelberg game. Then the operator is the leader who can set its incentive reward rate, and the tenants are the followers who decides how much energy to reduce given the operator's reward. We first analyzed this hierarchical game structure using the backward induction method and then proposed a linear time complexity to find its equilibria. Finally, the trace-based simulation results validated efficacy of our proposed scheme.

6. REFERENCES

[1] Colocation Market - Worldwide Market Forecast and Analysis (2013 - 2018).

[2] E. Consulting. Pricing data center colocation services, 2009.

[3] DatacenterMap. Colocation USA.

[4] EnerNOC. Demand Response: A Multi-Purpose Resource for Utilities and Grid Operations.

[5] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska. Dynamic right-sizing for power-proportional data centers. In *Proc. IEEE INFOCOM*, Shanghai, China, Apr. 2011.

[6] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. H. Andrew. Greening geographical load balancing. In *Proc. ACM SIGMETRICS*, San Jose, CA, 2011.

[7] Z. Liu, I. Liu, S. Low, and A. Wierman. Pricing data center demand response. *ACM SIGMETRICS Perform. Eval. Rev.*, 42(1):111–123, 2014.

[8] NRDC. Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers. In *Issue Pap.*, 2014.

[9] S. Ren and M. A. Islam. Colocation Demand Response : Why Do I Turn Off My Servers? In *USENIX ICAC*, Philadelphia, PA, 2014.

[10] N. H. Tran, S. Ren, Z. Han, S. man Jang, S. I. Moon, and C. S. Hong. Demand Response of Data Centers: A Real-time Pricing Game between Utilities in Smart Grid. In *USENIX Feed. Comput. Work.*, Philadelphia, PA, 2014.

[11] L. Zhang, S. Ren, C. Wu, and Z. Li. A Truthful Incentive Mechanism for Emergency Demand Response in Colocation Data Centers. In *IEEE INFOCOM*, Hong Kong, China, 2015.