A Fundamental Tradeoff Between Total and Brown Power Consumption in Geographically Dispersed Data Centers

---

Citation:
A. Kiani and N. Ansari, "A fundamental tradeoff between total and brown power consumption in geographically dispersed data centers," *IEEE Communications Letters*, vol. 20, no. 10, pp. 1955, 2016.

URL:
http://ieeexplore.ieee.org/document/7536170/

A Fundamental Tradeoff Between Total and Brown Power Consumption in Geographically Dispersed Data Centers

Abbas Kiani *Student Member, IEEE*, Nirwan Ansari *Fellow, IEEE*

***Abstract*—This paper aims at deriving a fundamental tradeoff between the total and brown power consumption associated with geographical dispersed data centers, where utilizing more green energy mostly happens at the cost of increasing the total power consumption. To this end, we define a new service efficiency parameter for data centers in satisfying the QoS requirements based on the queueing analysis. More importantly, we propose the idea of modeling geo-dispersed data centers with an information flow graph to capture a total-brown power consumption tradeoff region. Accordingly, we characterize the achievable tradeoff between total and brown power consumption.**

***Keywords*—Data centers, Information flow graph, Green energy**

I. Introduction

THE significant growing demand for online services has led to a multitude of challenges in provisioning Data Center Networks (DCNs) from DCN architecture design, congestion notification, TCP Incast, virtual machine migration, to routing in DCNs [1]. Most importantly, in recent years, preparing DCNs as a scalable computing infrastructure with hundreds of thousands of servers has witnessed a significant surge in the electric power usage.

Among the studies that have focused on reducing the data centers' power consumption, a small and cohesive body of work has investigated workload distribution across multiple data centers mainly to utilize the diversity of electricity price and renewable energy generation, and a variety of policies and algorithms have been proposed. In one of the earliest papers published on the subject, two request distribution policies were proposed to enable the requests to manage their energy consumption and cost while ensuring Service Level Agreements (SLAs) [2]. A similar policy that can act in accordance with the caps on the brown energy consumption is presented in [3]. The social impacts of geographical load balancing is explored in [4] and two distributed algorithms are provided to compute the optimal routing as well as capacity provisioning decisions for Internet-scale systems. Recently, Zhao *et al.* [5] took into consideration of dynamic Virtual Machine (VM) pricing and designed a new algorithm to maximize the long-term cloud provider's profit. Also, Kiani and Ansari [6] proposed a workload distribution strategy based on the notion of green workload and green service rate versus brown workload and brown service rate, respectively, and real-time monitoring of the queue lengths. Moreover, the possibility of purchasing power both at the Day-Ahead Market (DAM) and the Real-Time Market (RTM) has been proposed in [7].

Authors are with the Advanced Networking Lab., Department of Electrical and Computer Engineering, New Jersey Institute of Technology, Newark, NJ, 07102 USA. (e-mail:abbas.kiani@njit.edu and nirwan.ansari@njit.edu)

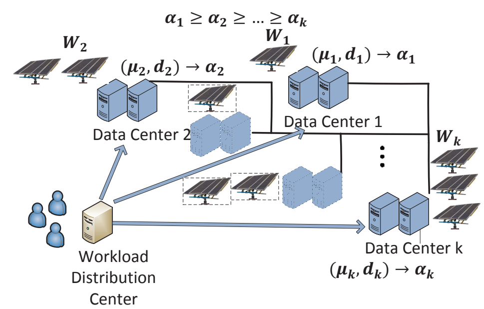

Fig. 1. System model.

Most of the proposed workload distribution strategies aim at reducing the energy cost and brown energy consumption via utilizing more renewable energy. Brown energy here refers to the energy which is produced from polluting sources and incurs environmental impacts. However, such a strategy may increase the total power consumption due to the fact that different data centers have different servers with different service capabilities, and also a request sent to different data centers experiences different network delays. In other words, the idea of sending a request to another data center with higher network delay or less service capability only in order to utilize more renewable energy may lead to a significant increase in the total power consumption. It is worth mentioning that the extra green energy at a data center can be injected into the power grid, and the data center can receive compensation for the injected power. In other words, the more green energy utilization at the cost of increasing the total energy consumption is not necessarily the best option. While, to our best knowledge, no prior work has investigated this tradeoff problem, we propose a new scheme to derive a fundamental tradeoff between the total and brown power consumption. To this end,

- We define a new service efficiency parameter for geo-dispersed data centers based on an M/GI/1 Processor Sharing (PS) queue analysis by taking into consideration of the network delay.
- We develop a new information flow graph based model for geo-dispersed data centers to capture the tradeoff between the total and brown power consumption.
- Based on the developed model, we characterize the achievable tradeoff between total and brown power consumption.

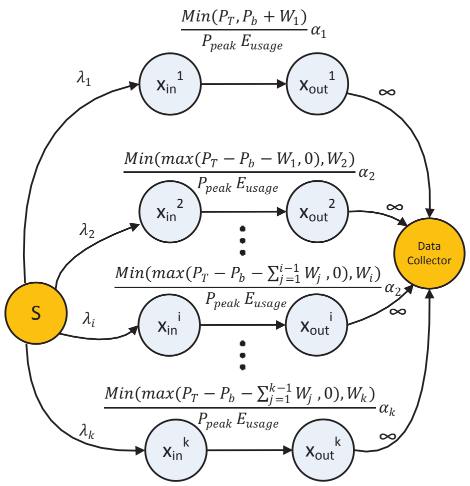

Fig. 2. Information flow graph of the system model.

The rest of paper is organized as follows. Section II describes the system model and problem formulation. In Section III, we propose our workload distribution strategy. Finally, Sections IV and V present numerical results and conclude the paper, respectively.

II. System Model and Problem Formulation

Fig. 1 shows the proposed system model in which we consider a group of N data centers dispersed at different regions. The service requests are initiated by users and arrive at a Workload Distribution Center (WDC). One or a group of servers can serve as the workload distribution center [8]. These servers can be treated as the front-end devices that exist in multi-data center Internet services like Google and Itunes [2]. The distribution center facilitates workload flexibility at the demand side. In other words, this center inspects the arriving requests from all users and manages the distribution of the incoming workload to the geo-dispersed data centers.

We divide the runtime of the data centers into a sequence of time slots at equal length, T, e.g., a few minutes. Our goal is to capture a fundamental tradeoff between the total and brown power consumption. To this end, we propose an optimization problem to be solved at the beginning of each time slot in which we update the number of the allocated requests to each data center. Note that for the analysis, we consider a single time slot, e.g., Δ as the time slot of interest, and omit the explicit time dependence in the notations.

The data centers are supplied by both on-grid and renewable types of power. The main power supply of each data center is on-grid or brown energy. To capitalize on the environmental and sustainability advantages of green energy, we also assume that each data center either is equipped with a renewable power source or has access to a nearby renewable energy source such as solar panels or a wind farm. Let W_i be the total available renewable power at data center i at the beginning of the time slot.

The allocated requests to a data center are first placed in a queue before they can be processed by any available server. We model each queue as an M/GI/1 PS queue which has been commonly adopted in modeling the waiting time of the requests at a data center in many studies like [4]. Therefore, the queuing delay at data center i can be computed as $\frac{1}{\mu_i - \frac{\lambda_i}{m_i}}$, where λ_i and μ_i are the allocated requests to data center i and the service rate of a single server at data center i, respectively. Also, m_i represents the total number of servers at data center i. The total number of servers that are turned on and run at full utilization can be computed as $m_i = \frac{P_i}{P_{peak} E_{usage}}$, where P_i is the power consumption of data center i. P_{peak} also indicates the average peak power of a turned on server in handling a service request. Moreover, E_{usage} is the Power Usage Effectiveness (PUE) of a data center and is defined as the ratio of the data center's total power consumption to the power consumption of the servers [6], [7].

To satisfy the QoS requirements, the queueing delay for each service request should be limited by a given deadline determined by the Service Level Agreement (SLA) between the data centers and clients. Let D be the SLA deadline. Therefore, according to our queuing delay, the allocated rate to each data center is upper bounded by

$$\lambda_i \leq \frac{P_i}{P_{peak} E_{usage}} (\mu_i - \frac{1}{D - d_i}), \tag{1}$$

where d_i denotes the network delay from the workload distribution center to data center i. The workload distribution center sorts N data centers based on $\alpha_i \triangleq \mu_i - \frac{1}{D-d_i}$ such that $\alpha_{i-1} \geq \alpha_i$.

Denote λ_T as the total number of requests arrived at the workload distribution center at the beginning of the time slot. To capture the tradeoff between the total and brown power consumption, we model geo-dispersed data centers with an information flow graph. The information flow graph is a directed acyclic graph which includes three types of nodes: (i) a single source node (S), (ii) some intermediate nodes, and (iii) data collector nodes [9], [10]. As depicted in Figure 2, the workload distribution center can be thought as the source node which is the source of original requests (WDC node). Also, the intermediate nodes are data centers, and data collector node can correspond to the users that receive processed requests.

The information flow graph, which models the geo-dispersed data centers, varies across time. At any given time, each node in the graph is either active or inactive. At the initial time of each time slot, the WDC node as the only active node contacts all N data center nodes and sorts them based on the service efficiency parameter, i.e., α_i. Then, it connects to a set of the first k data center nodes, i.e., $i = 1, ..., k$, with capacities of the edges equal to the allocated workloads to these nodes. It is assumed the total service provided by all the available renewable energy at these k data centers is not more than the required service to serve all the arriving requests. In fact, brown energy consumption is also required to serve all the requests and satisfy the QoS requirements. As the first data center has the highest service efficiency parameter and is assumed to have enough resources to satisfy the QoS requirements, it is more efficient to consume the brown energy only at this data center. Therefore, we have $P_1 = min(P_T, P_b + W_1)$ and $P_i = min(max(P_T - P_b - \sum_{j=1}^{i-1} W_j, 0), W_i)$ for $i = 2, ..., k$, where P_b is the brown power consumption, and P_T is assumed to be the total power consumption of all k data centers. Note that the brown power consumption depends on the number of connected data center nodes to the WDC node, i.e., k. In other words, in our model, connecting to different number of data centers will result in different amount of brown power consumption, and accordingly total power consumption. From this point onwards, WDC becomes and remains inactive, and selected data center nodes become active. Note that each data center node is represented by a pair of incoming and outgoing nodes connected by a directional edge whose capacity is the maximum number of requests that the data center can handle by the deadline. Finally, when the deadline comes, the data collector node becomes active and connects to the data center nodes to receive the processed requests. The edges that connect from the data center nodes to the data collector node are assumed to have infinite capacity, i.e., users have access to all the processed requests. In the next section, we will show how this model can capture the whole trade-off region between the total and brown power consumption.

III. Total-Brown Power Consumption Trade-off

In this section, we will characterize the optimal total-brown power consumption tradeoff region. As mentioned earlier, our workload allocation strategy needs k active data center nodes to connect to, and has to be designed such that the WDC node allocates $\lambda_1 = \frac{min(P_T, P_b+W_1)\alpha_1}{P_{peak}E_{usage}}$ requests to the first node and $\lambda_i = \frac{min(max(P_T-P_b-\sum_{j=1}^{i-1} W_i, 0), W_i)\alpha_i}{P_{peak}E_{usage}}$ requests to nodes $i = 2, ..., k$.

Theorem 1 *For some given (k, P_T), there exists $P_b^*(k, P_T)$ such that if $P_b \geq P_b^*(k, P_T)$, the points (k, P_b, P_T) are feasible, i.e., $P_b - P_T$ tradeoff is achievable. If $P_b \leq P_b^*(k, P_T)$, it is information theoretically impossible to serve all the arriving requests by the deadline. The threshold function $P_b^*(k, P_T)$ is,*

$$\overset{\star}{P}_b(k, P_T) =$$
$$\begin{cases} \frac{\lambda_T P_{peak}E_{usage} - \sum_{j=1}^{k} W_j\alpha_j}{\alpha_1}, & P_T \in \big[f(k), \infty\big) \\ \frac{\lambda_T P_{peak}E_{usage} - \sum_{j=1}^{i-1} W_j(\alpha_j-\alpha_i) - P_T\alpha_i}{\alpha_1-\alpha_i}, & P_T \in \big[f(i-1), f(i)\big), \end{cases} \tag{2}$$

where
$$f(i) \triangleq \frac{\lambda_T P_{peak}E_{usage} - \sum_{j=1}^{i} W_j(\alpha_j - \alpha_1)}{\alpha_1}, \tag{3}$$
and $i = 2, ..., k$.

Note that the tradeoff region which is verified in (2) has two extremal points corresponding to the minimum P_T and the minimum P_b, respectively. The point that minimizes P_T is always achieved when we send all the requests to the first data center. In (2), this point can be verified by letting $i = 2$, i.e., $(P_b, P_T) = (\frac{\lambda_T P_{peak}E_{usage} - \sum_{j=1}^{2-1} W_j(\alpha_j-\alpha_2) - P_T\alpha_2}{\alpha_1-\alpha_2}, f(1))$. On the other hand, the point that minimizes P_b is achieved when $P_T = f(k)$, i.e., when we send the requests to all available data centers.

Proof: Consider a given information flow graph. The minimum cut is a cut between the source node (WDC node) and the data collector node in which its total sum of the edge capacities is the smallest. According to Fig. 2, the capacity of the WDC-data collector minimum cut can be computed as

$$C = min(P_T, P_b + W_1)\frac{\alpha_1}{P_{peak}E_{usage}} +$$
$$\sum_{i=2}^{k} min(max(P_T - P_b - \sum_{j=1}^{i-1} W_i, 0), W_i)\frac{\alpha_i}{P_{peak}E_{usage}}. \tag{4}$$

If C is larger than or equal to the total number of requests (λ_T), the data collector node can receive all the processed requests by the deadline, and so the workload distribution strategy can meet the SLA requirements. To derive the optimal tradeoff between P_b and P_T, one can fix P_T and k (to some integer values) and then find the minimum value of P_b that satisfies $C \geq \lambda_T$. To this end, we define $\overset{\star}{P}_b(k, P_T)$ as follows:

$$\overset{\star}{P}_b(k, P_T) \triangleq \text{min } P_b$$
$$\text{subject to}: C \ \geq \ \lambda_T. \tag{5}$$

Note that C is a function of P_b. Therefore, $C(P_b)$ can be computed by considering the possible intervals of P_b.

$$C(P_b)P_{peak}E_{usage} =$$
$$\begin{cases} P_b\alpha_1 + \sum_{j=1}^{k} W_j\alpha_j, & P_b \in (0, P_T - \sum_{j=1}^{k} W_j] \\ P_b(\alpha_1 - \alpha_k) + P_T\alpha_k + \sum_{j=1}^{k-1} W_j(\alpha_j - \alpha_k), & P_b \in (P_T - \sum_{j=1}^{k} W_j, P_T - \sum_{j=1}^{k-1} W_j] \\ \vdots \\ P_b(\alpha_1 - \alpha_i) + P_T\alpha_i + \sum_{j=1}^{i-1} W_j(\alpha_j - \alpha_i), & P_b \in (P_T - \sum_{j=1}^{i} W_j, P_T - \sum_{j=1}^{i-1} W_j] \\ \vdots \\ P_b(\alpha_1 - \alpha_2) + P_T\alpha_2 + W_1(\alpha_1 - \alpha_2), & P_b \in (P_T - \sum_{j=1}^{2} W_j, P_T - W_1]. \end{cases}$$

As a result by noting $C \geq \lambda_T$ and letting $\overset{\star}{P}_b(k, P_T) = C^{-1}(\lambda_T)$, we have

$$\overset{\star}{P}_b(k, P_T) =$$
$$\begin{cases} \frac{\lambda_T P_{peak}E_{usage} - \sum_{j=1}^{k} W_j\alpha_j}{\alpha_1}, & \lambda_T P_{peak}E_{usage} \in A \\ \frac{\lambda_T P_{peak}E_{usage} - \sum_{j=1}^{i-1} W_j(\alpha_j-\alpha_i) - P_T\alpha_i}{\alpha_1-\alpha_i}, & \lambda_T P_{peak}E_{usage} \in B, \end{cases}$$

where $A \triangleq (\sum_{j=1}^{k} W_j\alpha_j, P_T\alpha_1 + \sum_{j=1}^{k} W_j(\alpha_j - \alpha_1)]$ and $B \triangleq (P_T\alpha_1 + \sum_{j=1}^{i} W_j(\alpha_j - \alpha_1), P_T\alpha_1 + \sum_{j=1}^{i-1} W_j(\alpha_j - \alpha_1)]$. By changing the conditions in the above expression from $\lambda_T P_{peak}E_{usage}$ to P_T, our tradeoff region, i.e., (2), is derived. ■

IV. Numerical Results

We consider $k = 6$ data centers, each integrated with a wind farm as a renewable power source. Our simulation data are based on the trends of wind power and the total workload shown in Figs. 3 and 4, respectively. Fig. 5 shows the tradeoff

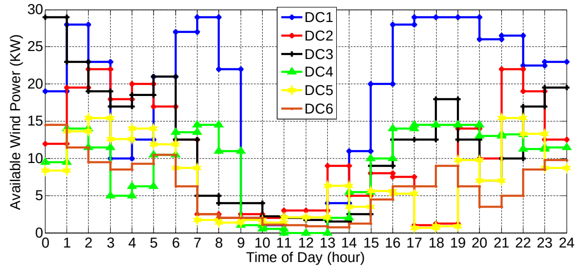

Fig. 3. Wind power generation.

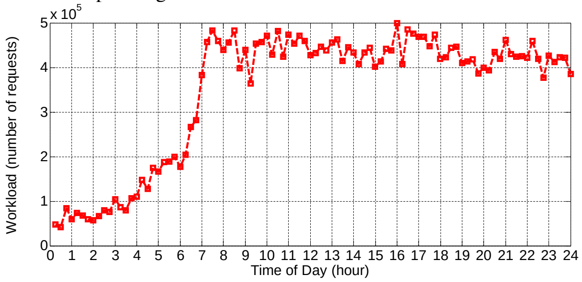

Fig. 4. The total incoming workload.

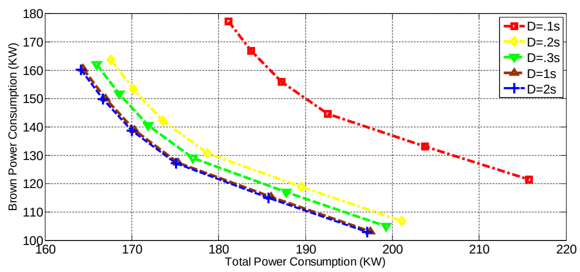

Fig. 5. Total-brown power consumption tradeoff curves for different values of D.

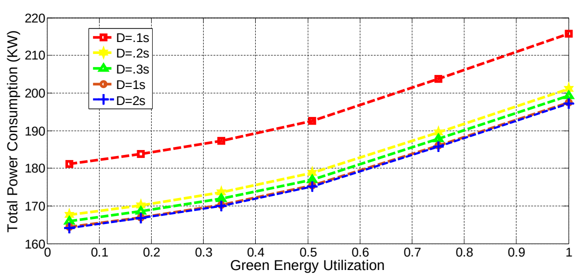

Fig. 6. Green power utilization-total power consumption tradeoff curves for different values of D.

curves between the total and brown power consumption for different values of D, which is the deadline to serve the requests. The tradeoff curves in this figure confirm that we can decrease brown power consumption by increasing the total power consumption. Also, the green power utilization-total power consumption tradeoff curves for different values of D are shown in Fig. 6. The green energy utilization is defined as the consumed wind power divided by the total available wind power at 6 data centers. Figs. 5 and 6 demonstrate that the curve corresponding to the highest deadline outperforms that of the curves with lower deadline values. Finally, Fig. 7 provides

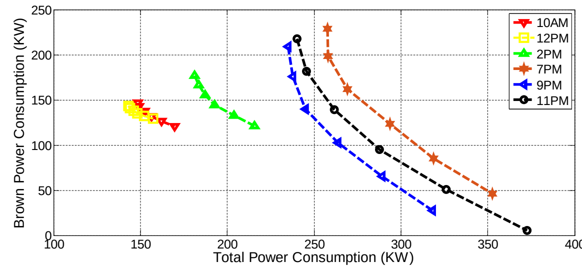

Fig. 7. Total-brown power consumption tradeoff curves at different hours of day.

the total-brown power consumption tradeoff at some sample hours of the day when $D = .1$. As shown in this figure, for example, the tradeoff curve at hour 12PM outperforms those of the other curves due to the less number of arrival requests.

V. Conclusion

In this paper, we have developed a new information flow graph based model for geo-dispersed data centers. Based on the developed model, we have derived a fundamental tradeoff between the total and brown power consumption. Furthermore, we have characterized the achievable points on this tradeoff in which one can know how much green energy is possibly utilized for a given amount of total power consumption budget.

References

[1] Y. Zhang and N. Ansari, "On architecture design, congestion notification, tcp incast and power consumption in data centers," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 1, pp. 39–64, 2013.

[2] K. Le, R. Bianchini, M. Martonosi, and T. D. Nguyen, "Cost-and energy-aware load distribution across data centers," in *Workshop on Power Aware Computing and Systems*. Big Sky, MT, October, 2009.

[3] K. Le, O. Bilgir, R. Bianchini, M. Martonosi, and T. D. Nguyen, "Capping the brown energy consumption of internet services at low cost," in *IEEE International Green Computing Conference*, August 2010.

[4] Z. Liu, M. Lin, A. Wierman, S. Low, and L.L. Andrew, "Greening geographical load balancing," *IEEE/ACM Transactions on Networking (TON)*, vol. 23, no. 2, pp. 657–671, 2015.

[5] J. Zhao, H. Li, C. Wu, Z. Li, Z. Zhang, and F. Lau, "Dynamic pricing and profit maximization for clouds with geo-distributed datacenters," *IEEE Infocom*, 2014.

[6] A. Kiani and N. Ansari, "Towards low-cost workload distribution for integrated green data centers," *IEEE Communications Letters*, vol. 19, no. 1, pp. 26–29, 2015.

[7] M. Ghamkhari, H. Mohsenian-Rad, and A. Wierman, "Optimal risk-aware power procurement for data centers in day-ahead and real-time electricity markets," in *IEEE Conference on Computer Communications (INFOCOM WKSHPS)*. Toronto, Canada, p. 2014.

[8] M. Ghamkhari and H. Mohsenian-Rad, "Optimal integration of renewable energy resources in data centers with behind-the-meter renewable generator," in *IEEE International Conference on Communications (ICC)*, 2012, pp. 3340–3344.

[9] A.G. Dimakis, P. Godfrey, Y. Wu, M.J. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *IEEE Transactions on Information Theory*, vol. 56, no. 9, pp. 4539–4551, 2010.

[10] A. Kiani and S. Akhlaghi, "Selective regenerating codes," *IEEE Communications Letters*, vol. 15, no. 8, pp. 854–856, 2011.