

Cost Minimization for Big Data Processing in Geo-Distributed Data Centers

Soniya Khan¹ Tanmay Naik² Ajay Temakar³ Sameer Mejari⁴ Prof. Madhuri Patil⁵

^{1,2,3,4,5}Department of Information Technology

^{1,2,3,4,5}MGM CET, India

Abstract— The explosive growth of demands on big data processing imposes a heavy burden on computation, storage, and communication in data centers, which hence incurs considerable operational expenditure to data center providers. Therefore, cost minimization has become an emergent issue for the upcoming big data era. Different from conventional cloud services, one of the main features of big data services is the tight coupling between data and computation as computation tasks can be conducted only when the corresponding data is available. As a result, three factors, i.e., task assignment, data placement and data movement, deeply influence the operational expenditure of data centers. In this paper, we are motivated to study the cost minimization problem via a joint optimization of these three factors for big data services in geo-distributed data centers. To describe the task completion time with the consideration of both data transmission and computation, we propose a two-dimensional Markov chain and derive the average task completion time in closed-form. Furthermore, we model the problem as a mixed-integer non-linear programming (MINLP) and propose an efficient solution to linearize it. The high efficiency of our proposal is validated by extensive simulation based studies.

Key words: Big Data, Data Flow, Data Placement, Distributed Data Centers, Cost Minimization, Task Assignment

I. INTRODUCTION

Data explosion in recent years leads to a rising demand for big data processing in modern data centers that are usually distributed at different geographic regions, e.g., Google's 13 data centers over 8 countries in 4 continents [1]. Big data analysis has shown its great potential in unearthing valuable insights of data to improve decision-making, minimize risk and develop new products and services. On the other hand, big data has already translated into big price due to its high demand on computation and communication resources [2]. Gartner predicts that by 2015, 71% of worldwide data center hardware spending will come from the big data processing, which will surpass \$126.2 billion. Therefore, it is imperative to study the cost minimization problem for big data processing in geo-distributed data centers.

Many efforts have been made to lower the computation or communication cost of data centers. Data center resizing (DCR) has been proposed to reduce the computation cost by adjusting the number of activated servers via task placement [3]. Based on DCR, some studies have explored the geographical distribution nature of data centers and electricity price heterogeneity to lower the electricity cost [4]–[6]. Big data service frameworks, e.g., [7], comprise a distributed file system underneath, which distributes data chunks and their replicas across the data centers for fine-grained load-balancing and high parallel

data access performance. To reduce the communication cost, a few recent studies make efforts to improve data locality by placing jobs on the servers where the input data reside to avoid remote data loading [7], [8].

Although the above solutions have obtained some positive results, they are far from achieving the cost-efficient big data processing because of the following weaknesses. First, data locality may result in a waste of resources. For example, most computation resource of a server with less popular data may stay idle. The low resource utility further causes more servers to be activated and hence higher operating cost.

Second, the links in networks vary on the transmission rates and costs according to their unique features [9], e.g., the distances and physical optical fiber facilities between data centers. However, the existing routing strategy among data centers fails to exploit the link diversity of data center networks. Due to the storage and computation capacity constraints, not all tasks can be placed onto the same server, on which their corresponding data reside. It is unavoidable that certain data must be downloaded from a remote server. In this case, routing strategy matters on the transmission cost. As indicated by Jin et al. [10], the transmission cost, e.g., energy, nearly proportional to the number of network link used. The more link used, the higher cost will be incurred. Therefore, it is essential to lower the number of links used while satisfying all the transmission requirements. Third, the Quality-of-Service (QoS) of big data tasks has not been considered in existing work. Similar to conventional cloud services, big data applications also exhibit Service-Level-Agreement (SLA) between a service provider and the requesters. To observe SLA, a certain level of QoS, usually in terms of task completion time, shall be guaranteed. The QoS of any cloud computing tasks is first determined by where they are placed and how many computation resources are allocated. Besides, the transmission rate is another influential factor since big data tasks are data-centric and the computation task cannot proceed until the corresponding data are available. Existing studies, e.g., [3], on general cloud computing tasks mainly focus on the computation capacity constraints, while ignoring the constraints of transmission rate.

To conquer above weaknesses, we study the cost minimization problem for big data processing via joint optimization of task assignment, data placement, and routing in geo-distributed data centers. Specifically, we consider the following issues in our joint optimization. Servers are equipped with limited storage and computation resources. Each data chunk has a storage requirement and will be required by big data tasks. The data placement and task assignment are transparent to the data users with guaranteed QoS. Our objective is to optimize the big data placement, task assignment, routing and DCR such that the overall

computation and communication cost is minimized. Our main contributions are summarized as follows:

To our best knowledge, we are the first to consider the cost minimization problem of big data processing with joint consideration of data placement, task assignment and data routing. To describe the rate-constrained computation and transmission in big data processing process, we propose a two-dimensional Markov chain and derive the expected task completion time in closed form.

Based on the closed-form expression, we formulate the cost minimization problem in a form of mixed-integer nonlinear programming (MINLP) to answer the following questions: 1) how to place these data chunks in the servers, 2) how to distribute tasks onto servers without violating the resource constraints, and 3) how to resize data centers to achieve the operation cost minimization goal.

To deal with the high computational complexity of solving MINLP, we linearize it as a mixed-integer linear programming (MILP) problem, which can be solved using commercial solver. Through extensive numerical studies, we show the high efficiency of our proposed joint-optimization based algorithm.

The rest of the paper is organized as follows. Section 2 summaries the related work. Section 3 introduces our system model. The cost optimization is formulated as an MINLP problem in Section 4 and then it is linearized in Section 5. The theoretical findings are verified by experiments in Section 6. Finally, Section 7 concludes our work.

II. RELATED WORK

A. Server Cost Minimization

Large-scale data centers have been deployed all over the world providing services to hundreds of thousands of users. According to [11], a data center may consist of large numbers of servers and consume megawatts of power. Millions of dollars on electricity cost have posed a heavy burden on the operating cost to data center providers. Therefore, reducing the electricity cost has received significant attention from both academia and industry [5], [11]–[13]. Among the mechanisms that have been proposed so far for data center energy management, the techniques that attract lots of attention are task placement and DCR.

DCR and task placement are usually jointly considered to match the computing requirement. Liu et al. [4] re-examine the same problem by taking network delay into consideration. Fan et al. [12] study power provisioning strategies on how much computing equipment can be safely and efficiently hosted within a given power budget. Rao et al. [3] investigate how to reduce electricity cost by routing user requests to geo-distributed data centers with accordingly updated sizes that match the requests. Recently, Gao et al. [14] propose the optimal workload control and balancing by taking account of latency, energy consumption and electricity prices. Liu et al. [15] reduce electricity cost and environmental impact using a holistic approach of workload balancing that integrates renewable supply, dynamic pricing, and cooling supply.

B. Big data Management

To tackle the challenges of effectively managing big data, many proposals have been proposed to improve the storage and computation process.

The key issue in big data management is reliable and effective data placement. To achieve this goal, Sathiamoorthy et al. [16] present a novel family of erasure codes that are efficiently repairable and offer higher reliability compared to Reed-Solomon codes. They also analytically show that their codes are optimal on an identified tradeoff between locality and minimum distance. Yazd et al. [8] make use of flexibility in the data block placement policy to increase energy efficiency in data centers and propose a scheduling algorithm, which takes into account energy efficiency in addition to fairness and data locality properties. Hu et al. [17] propose a mechanism allowing linked open data to take advantage of existing large-scale data stores to meet the requirements on distributed and parallel data processing.

Moreover, how to allocate the computation resources to tasks has also drawn much attention. Cohen et al. [18] present new design philosophy, techniques and experience providing a new magnetic, agile and deep data analytics for one of the world's largest advertising networks at Fox Audience Network, using the Greenplum parallel database system. Kaushik et al. [19] propose a novel, data-centric algorithm to reduce energy costs and with the guarantee of thermal-reliability of the servers. Chen et al. [20] consider the problem of jointly scheduling all three phases, i.e., map, shuffle and reduce, of the MapReduce process and propose a practical heuristic to combat the high scheduling complexity.

Constants	
\bar{J}_i	The set of servers in data center i
m_i	The switch in data center i
$w^{(u,v)}$	The weight of link (u, v)
φ_k	The size of chunk k
λ_k	The task arrival rate for data chunk k
P	The number of data chunk replicas
D	The maximum expected response time
P_j	The power consumption of server j
$\gamma^{(u,v)}$	The transmission rate of link (u, v)
Variables	
x_j	A binary variable indicating if server j is activated or not
y_{jk}	A binary variable indicating if chunk k is placed on server j or not
$z_{jk}^{(u,v)}$	A binary variable indicating if link (u, v) is used for flow for chunk k on server j
λ_{jk}	The request rate for chunk k on server j
θ_{jk}	The CPU usage of chunk k on server j
μ_{jk}	The CPU processing rate of chunk k on server j
$f_{jk}^{(u,v)}$	The flow for chunk k destined to server j through link (u, v)

Table.1: Notations

C. Data Placement

Shachnai et al. [21] investigate how to determine a placement of Video-on-Demand (VoD) file copies on the servers and the amount of load capacity assigned to each file copy so as to minimize the communication cost while ensuring the user experience. Agarwal et al. [22] propose an automated data placement mechanism Volley for geo-distributed cloud services with the consideration of WAN bandwidth cost, data center capacity limits, data inter-dependencies, etc. Cloud services make use of Volley by submitting logs of datacenter requests. Volley analyzes the logs using an iterative optimization algorithm based on data access patterns and client locations, and outputs migration recommendations back to the cloud service. Cidon et al. [23] invent MinCopysets, a data replication placement scheme that decouples data distribution and replication to improve the data durability properties in distributed data centers. Recently, Jin et al. [10] propose a joint optimization scheme that simultaneously optimizes virtual machine (VM) placement and network flow routing to maximize energy savings.

Existing work on data center cost optimization, big data management or data placement mainly focuses on one or two factors. To deal with big data processing in geo-distributed data centers, we argue that it is essential to jointly consider data placement, task assignment and data flow routing in a systematical way

III. SYSTEM MODEL

In this section, we introduce the system model. For the convenience of the readers, the major notations used in this paper are listed in Table 1.

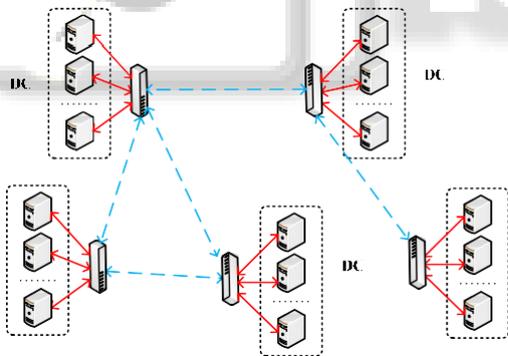


Fig.1: Data center topology

A. Network Model

We consider a geo-distributed data center topology as shown in Fig. 1, in which all servers of the same data center (DC) are connected to their local switch, while data centers are connected through switches. There are a set I of data centers, and each data center $i \in I$ consists of a set J_i of servers that are connected to a switch $m_i \in M$ with a local transmission cost of C_L . In general, the transmission cost C_R for inter-data center traffic is greater than C_L , i.e., $C_R > C_L$. Without loss of generality, all servers in the network have the same computation resource and storage capacity, both of which are normalized to one unit. We use J to denote the set of all servers i.e. $J = J_1, J_2 \dots J_{|I|}$. The whole system can be modeled as a directed graph $G = (N, E)$. The vertex set N

$= M \cup J$ includes the set M of all switches and the set J of all servers, and E is the directional edge set. All servers are connected to, and only to, their local switch via intra-data center links while the switches are connected via inter-data center links determined by their physical connection. The weight of each link (u, v) , representing the corresponding communication cost, can be defined as

$$w(u, v) = \begin{cases} C_R & \text{if } u, v \in M \\ C_L & \text{otherwise} \end{cases} \quad (1)$$

B. Task Model

We consider big data tasks targeting on data stored in a distributed file system that is built on geo-distributed data centers. The data are divided into a set K of chunks. Each chunk $k \in K$ has the size of ϕ_k ($\phi_k \leq 1$), which is normalized to the server storage capacity. P -way replica [19] is used in our model. That is, for each chunk, there are exactly P copies stored in the distributed file system for resiliency and fault-tolerance.

It has been widely agreed that the tasks arrival at data centers during a time period can be viewed as a Poisson process [9], [24]. In particular, let λ_k be the average task arrival rate requesting chunk k . Since these tasks will be distributed to servers with a fixed probability, the task arrival in each server can be also regarded as a Poisson process. We denote the average arrival rate of task for chunk k on server j as λ_{jk} ($\lambda_{jk} \leq 1$). When a task is distributed to a server where its requested data chunk does not reside, it needs to wait for the data chunk to be transferred. Each task should be responded in time D . Moreover, in practical data center management, many task prediction mechanisms based on the historical statistics have been developed and applied to the decision making in data centers [19]. To keep the data center settings up-to-date, data center operators may make adjustment according to the task prediction period by period [3], [14], [15]. This approach is also adopted in this paper.

IV. PROBLEM FORMULATION

In this section, we first present the constraints of data and task placement, remote data loading, and QoS. Then, we give the complete formulation of the cost minimization problem in a mixed-integer nonlinear programming form.

A. Constraints of Data and Task Placement

We define a binary variable y_{jk} to denote whether chunk k is placed on server j as follows,

$$y_{jk} = \begin{cases} 1 & \text{if chunk } k \text{ is placed on server } j \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

In the distributed file system, we maintain P copies for each chunk $k \in K$, which leads to the following constraint:

$$\sum_{j \in J} y_{jk} = P, \forall k \in K \quad (3)$$

Furthermore, the data stored in each server $j \in J$ cannot exceed its storage capacity, i.e.,

$$y_{jk} \cdot \phi_k \leq 1; k \in K, \forall j \in J \quad (4)$$

As for task distribution, the sum rates of task assigned to each server should be equal to the overall rate,

$$\lambda_k = \lambda_{jk} ; j \in J, \forall k \in K \quad (5)$$

Finally, we define a binary variable x_j to denote whether server j is activated, i.e.,

$$x_j = \begin{cases} 1 & \text{if this server is activated} \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

$$(y_{jk})_{k \in K} + (\lambda_{jk})_{k \in K} \leq x_j \leq (y_{jk})_{k \in K} + A ; k \in K, \forall j \in J$$

A server shall be activated if there are data chunks placed onto it or tasks assigned to it. Therefore, we have

$$\sum_{k \in K} y_{jk} + \sum_{k \in K} \lambda_{jk} \leq x_j \leq \sum_{k \in K} y_{jk} + A \sum_{k \in K} \lambda_{jk} ; \forall j \in J \quad (7)$$

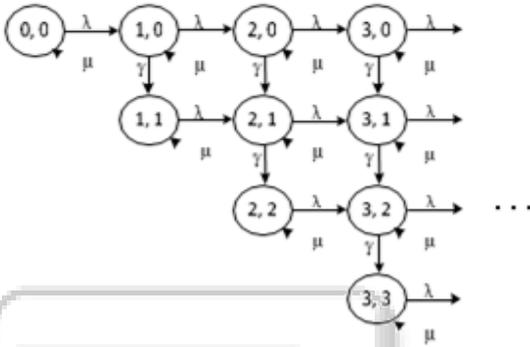


Fig. 2: Two-dimensional Markov Chain

B. Constraints of Data Loading

Note that when a data chunk k is required by a server j , it may cause internal and external data transmissions. This routing procedure can be formulated by a flow model. All the nodes N in graph tt , including the servers and switches, can be divided into three categories:

- Source nodes $u (u \in J)$. They are the servers with chunk k stored in it. In this case, the total outlet flows to destination server j for chunk k from all source nodes shall meet the total chunk requirement per time unit as $\lambda_{jk} \cdot \varphi_k$.
- Destination node $j (j \in J)$. When the required chunk is not stored in the destination node, i.e., $y_{jk} = 0$, it must receive the data flows of chunk k at a rate $\lambda_{jk} \cdot \varphi_k$.
- Let the flow over the link $(u, v) \in E$ carrying data of chunk $k \in K$ and destined to server $j \in J$. Then, the constraints on the above three categories of nodes can be expressed as follows respectively

$$f_{jk}^{(u,v)} \leq y_{uk} \cdot \lambda_k \cdot \varphi_k ; \forall (u, v) \in E, u \& j \in J, k \in K \quad (8)$$

Note that a non-zero flow $f(u, v)$ emitting from server u only if it keeps a copy of chunk k , i.e., $y_{uk} = 1$, as characterized in (8).

$$\sum_{(u,j) \in E} f_{jk}^{(u,j)} - \sum_{(v,w) \in E} f_{jk}^{(v,w)} = 0 ; \forall v \in N, j \in J, k \in K \quad (9)$$

The flow conservation is maintained on each switch as shown in (9). Finally, the destination receives all data $\lambda_k \varphi_k$ from others only when it does not hold a copy of chunk k , i.e., $y_{ik} = 0$. This is guaranteed by (10).

$$\sum_{(u,j) \in E} f_{jk}^{(u,j)} = (1 - y_{jk}) \cdot \lambda_{jk} \cdot \varphi_k ; \forall j \in J, k \in K \quad (10)$$

C. Constraints of QoS Satisfaction

Let μ_{jk} and γ_{jk} be the processing rate and loading rate for data chunk k on server j , respectively. The processing Markov chain as Fig. 2, where each state (p, q) represents p pending tasks and q available data chunks. We let θ_{jk} denote the amount of computation resource (e.g., CPU) that chunk k occupies. The processing rate of tasks is proportional to its computation resource usage, i.e.,

$$\mu_{jk} = \alpha_j \cdot \theta_{jk} ; j \in J, k \in K \quad (11)$$

where α_j is a constant relying on the speed of server j . Furthermore, the total computation resource allocated to all chunks on each server j shall not exceed its total computation resource, i.e.,

$$\sum_{k \in K} \theta_{jk} \leq 1 ; \forall j \in J \quad (12)$$

The loading rate of γ_{jk} is constrained by the rate on any link (u, v) denoted as $\gamma(u, v)$, if a non-zero flow f_{jk} goes through it. This condition can be described by a binary variable $z_{jk}^{(u,v)}$ as

$$f_{jk}^{(u,v)} \leq z_{jk}^{(u,v)} \leq A f_{jk}^{(u,v)} ; \forall (u, v) \in E, j \in J, k \in K \quad (13)$$

Finally, the constraints on γ_{jk} is given as

$$\gamma_{jk} \leq \gamma^{(u,v)} \cdot z_{jk}^{(u,v)} + 1 - z_{jk}^{(u,v)} ; \forall (u, v) \in E, j \in J, k \in K \quad (14)$$

Note that we consider sufficient bandwidth on each link such that $\gamma(u, v)$ can be handled as a constant number, which is mainly determined by I/O and switch latency [25].

By denoting $\pi_{jk}(p, q)$ as the steady state probability that the Markov chain stays at (p, q) , we can describe the transition process by a group of ODEs as follows. According to the transition characteristics, the whole figure can be divided into three regions.

Region-I: all states in the first line. In Region-I, except state $(0, 0)$, state $(p, 0)$ ($p > 1$) transits to two neighbouring states $(p + 1, 0)$ and $(p, 1)$. These can be described as:

$$\pi'_{jk}(0,0) = -\lambda_{jk}\pi_{jk}(0,0) + \mu_{jk}\pi_{jk}(1,1); \forall j \in J, k \in K \quad (15)$$

$$\pi'_{jk}(p,0) = -\lambda_{jk}(\pi_{jk}(p,0) - \pi_{jk}(p-1,0)) + \mu_{jk}\pi_{jk}(p+1,1) - \gamma\pi_{jk}(p,0); \forall j \in J, k \in K \quad (16)$$

Region II: all states in the diagonal line except $(0, 0)$. In this region, all the pending tasks have already obtained their needed data chunk to proceed. Therefore, each state (p, q) in Region-II will transit to $(p-1, q-1)$ after processing one data chunk. Then, we have:

$$\pi'_{jk}(p,p) = -\lambda_{jk}\pi_{jk}(p,p) + \mu_{jk}(\pi_{jk}(p+1,p+1) - \pi_{jk}(p,p)) + \gamma\pi_{jk}(p,p-1); \forall j \in J, k \in K \quad (17)$$

Region-III: all remaining states in the central region. Each state (p, q) in Region-III relies on its three neighbouring states and also will transit to the other three neighbouring

states. As shown in Fig. 2, the transition relationship can be written as:

$$\begin{aligned} \pi'_{jk}(p, q) = & -\lambda_{jk} \left(\pi_{jk}(p, q) - \pi_{jk}(p-1, q-1) \right) \\ & + \mu_{jk} \left(\pi_{jk}(p+1, q+1) - \pi_{jk}(p, q) \right) \\ & - \gamma \left(\pi_{jk}(p, q) - \pi_{jk}(p-1, q-1) \right); \\ & \forall j \in J, k \in K \end{aligned} \quad (18)$$

By solving the above ODEs, we can derive the state probability $\pi_{jk}(p, q)$ as:

$$\pi_{jk}(p, q) = \frac{(\lambda_{jk})^p (\mu_{jk})^{B-q} (\gamma_{jk})^{B-p+q}}{\sum_{q=0}^B \sum_{p=0}^B (\lambda_{jk})^p (\mu_{jk})^{B-q} (\gamma_{jk})^{B-p+q}}; \quad \forall j \in J, k \in K \quad (19)$$

where B is the task buffer size on each server. When B goes to infinity, the mean number of tasks for chunk k on server j T_{jk} is

$$T_{jk} = \lim_{B \rightarrow \infty} \frac{\sum_{q=0}^B \sum_{p=0}^B p (\lambda_{jk})^p (\mu_{jk})^{B-q} (\gamma_{jk})^{B-p+q}}{\sum_{q=0}^B \sum_{p=0}^B (\lambda_{jk})^p (\mu_{jk})^{B-q} (\gamma_{jk})^{B-p+q}} \quad \forall j \in J, k \in K \quad (20)$$

By applying the multivariate l'Hospital's rule, (20) can be simplified to

$$T_{jk} = \frac{\lambda_{jk}}{\mu_{jk} \gamma_{jk} - \lambda_{jk}} \quad \forall j \in J, k \in K \quad (21)$$

According to the Little's law, the expected delay d_{jk} of user requests for chunk k on server j is

$$d_{jk} = \frac{T_{jk}}{\lambda_{jk}} = \frac{1}{\mu_{jk} \gamma_{jk} - \lambda_{jk}}; \quad \forall j \in J, k \in K \quad (22)$$

According to the QoS requirement, i.e., $d_{jk} \leq D$, we have

$$\mu_{jk} \cdot \gamma_{jk} - \lambda_{jk} \geq \frac{\mu_{jk}}{D}; \quad \forall j \in J, k \in K \quad (23)$$

Where

$$\mu_{jk} = \begin{cases} 1 & \text{if } \lambda_{jk} \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad (24)$$

The binary variable of u_{jk} can be described by constraints:

$$\lambda_{jk} \leq \mu_{jk} \leq A \lambda_{jk}; \quad \forall j \in J, k \in K \quad (25)$$

where A is an arbitrary large number, because of $0 < \lambda_{jk} < 1$ and $u_{jk} \in \{0, 1\}$.

D. An MINLP Formulation

The total energy cost then can be calculated by summing up the cost on each server across all the geo-distributed data centers and the communication cost, i.e.,

$$C_{total} = \sum_{j \in J} x_j \cdot P_j + \sum_{j \in J} \sum_{k \in K} \sum_{(u,v) \in E} f_{jk}^{(u,v)} \cdot w^{(u,v)} \quad (26)$$

where P_j is the cost of each activated server j .

Our goal is to minimize the total cost by choosing the best settings of x_j , y_{jk} , $z_{jk}^{(u,v)}$, θ_{jk} , λ_{jk} and $f^{(u,v)}$. By summarizing all constraints discussed above, we can formulate this cost minimization as a mixed-integer non-linear programming (MINLP) problem as:

MINLP:

min : (26),

s.t. : (3) - (5), (7) - (14), (23), (25),

$x_j, y_{jk}, z_{jk}, u_{jk} \in \{0, 1\}, \forall j \in J, k \in K$

Note that the above formulation is based on the setting that the number of replicas for each data chunk is a predetermined constant. If it is a part of the optimization, denoted by an integer variable p , the total cost can be further minimized by the formulation below.

MINLP-2:

min : (26),

s. t. : $\sum_{j \in J} y_{jk} = p, \forall k \in K, p > 1, (4), (5), (7) - (14), (23), (25),$

$x_j, y_{jk}, z_{jk}, u_{jk} \in \{0, 1\}, \forall j \in J, k \in K$

V. LINEARIZATION

We observe that the constraints (8) and (10) are nonlinear due to the products of two variables. To linearize these constraints, we define a new variable δ_{jk} as follows:

$$\delta_{jk} = y_{jk} \lambda_{jk}; \quad \forall j \in J, \forall k \in K \quad (27)$$

which can be equivalently replaced by the following linear constraints:

$$0 \leq \delta_{jk} \leq \lambda_{jk}; \quad \forall j \in J, \forall k \in K \quad (28)$$

$$\lambda_{jk} + y_{jk} - 1 \leq \delta_{jk} \leq y_{jk} \nu; \quad \forall j \in J, \forall k \in K \quad (29)$$

The constraints (8) and (10) can be written in a linear form as:

$$\delta_{jk}^{(u,v)} \leq \delta_{uk} \varphi_k; \quad \forall (u, v) \in E, \forall j \in J, \forall k \in K \quad (30)$$

$$\sum_{(u,j) \in E} f_{jk}^{(u,j)} = (\lambda_{jk} - \delta_{jk}) \cdot \varphi_k \quad \forall j \in J, k \in K \quad (31)$$

We then consider the remaining nonlinear constraints (14) and (23) that can be equivalently written as:

$$\begin{aligned} \gamma^{(u,v)} \cdot \mu_{jk} \cdot z_{jk}^{(u,v)} + 1 - \mu_{jk} \cdot z_{jk}^{(u,v)} - \lambda_{jk} & \geq \frac{\mu_{jk}}{D} \\ \forall (u, v) \in E, \forall j \in J, \forall k \in K \end{aligned} \quad (32)$$

In a similar way, we define a new variable ϵ_{jk} as:

$$\epsilon_{jk}^{(u,v)} = \mu_{jk} \cdot z_{jk}^{(u,v)} \quad (33)$$

such that constraint (32) can be written as:

$$\begin{aligned} \gamma^{(u,v)} \cdot \epsilon_{jk}^{(u,v)} + 1 - \epsilon_{jk}^{(u,v)} - \lambda_{jk} & \geq \frac{\mu_{jk}}{D} \\ \forall (u, v) \in E, \forall j \in J, \forall k \in K \end{aligned} \quad (34)$$

The constraint (33) can be linearized by:

$$0 \leq \epsilon_{jk}^{(u,v)} \leq \mu_{jk}; \quad \forall (u, v) \in E, \forall j \in J, \forall k \in K \quad (35)$$

$$\mu_{jk} + z_{jk}^{(u,v)} - 1 \leq \epsilon_{jk}^{(u,v)} \leq z_{jk}^{(u,v)}; \quad \forall j \in J, k \in K \quad (36)$$

Now, we can linearize the **MINLP** problem into a mixed-integer linear programming (**MILP**) as

MILP:

min : (26),

s. t. : (3) – (5), (7), (9), (11) – (13),
 (25), (28) – (31), (34) – (36),
 $x_j, y_{jk}, z_{jk}, u_{jk} \in \{0, 1\}, \forall j \in J, k \in K.$

VI. PERFORMANCE EVALUATION

In this section, we present the performance results of our joint-optimization algorithm (“Joint”) using the MILP formulation. We also compare it against a separate optimization scheme algorithm (“Non-joint”), which first finds a minimum number of servers to be activated and the traffic routing scheme using the network flow model as described in Section 4.(B).

In our experiments, we consider $|J| = 3$ data centers, each of which is with the same number of servers. The intra- and inter-data center link communication cost are set as $C_L = 1$ and $C_R = 4$, respectively. The cost P_j on each activated server j is set to 1. The data size, storage requirement, and task arrival rate are all randomly generated. To solve the MILP problem, commercial solver Gurobi [26] is used.

The default settings in our experiments are as follows: each data center with a size 20, the number of data chunks $|K| = 10$, the task arrival rates $\lambda_k \in [0.01, 5], k \in K$, the number of replicas $P = 3$, the data chunk size $\varphi_k \in [0.01, 1], k \in K$, and $D = 100$. We investigate how various parameters affect the overall computation, communication and overall cost by varying one parameter in each experiment group.

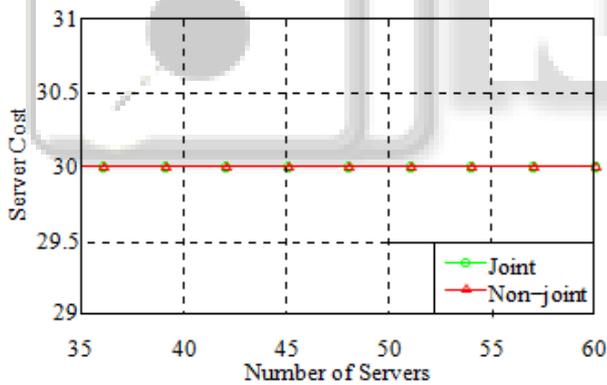


Fig. 3(a): Effect of Server cost on no. of servers

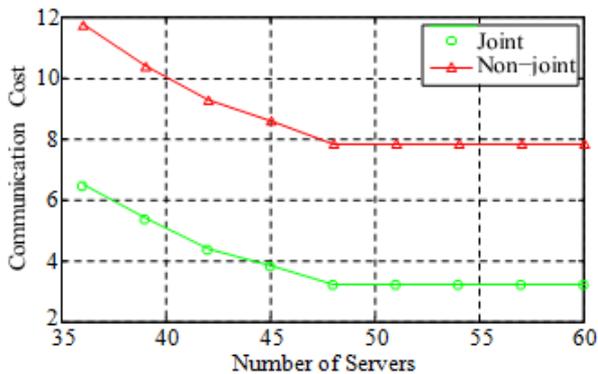


Fig. 3(b): Effect of Communication cost on no. of servers

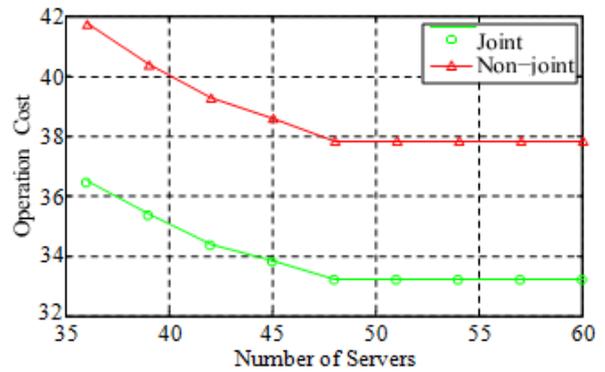


Fig. 3(c): Effect of Overall cost on no. of servers

Fig. 3 shows the server cost, communication cost and overall cost under different total server numbers varying from 36 to 60. As shown in Fig. 3(a), we can see that the server cost always keep constant on any data center size. As observed from Fig. 3(b), when the total number of servers increases from 36 to 48, the communication costs of both algorithms decrease significantly. This is because more tasks and data chunks can be placed in the same data center when more servers are provided in each data center. Hence, the communication cost is greatly reduced. However, after the number of server reaching 48, the communication costs of both algorithms converge. The reason is that most tasks and their corresponding data chunks can be placed in the same data center, or even in the same server. Further increasing the number of servers will not affect the distributions of tasks or data chunks any more. Similar results are observed in Fig. 3(c).

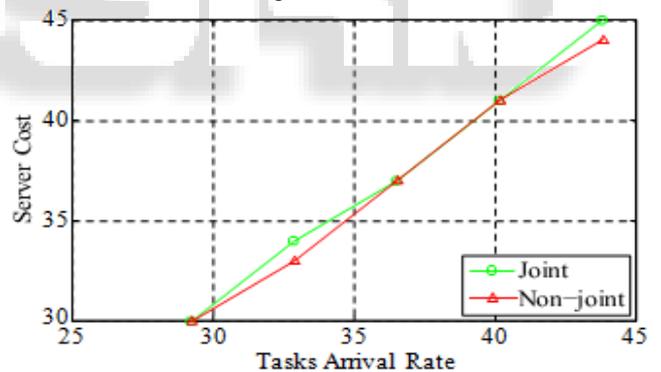


Fig. 4(a) : Effect of Server Cost on Task Arrival Rate

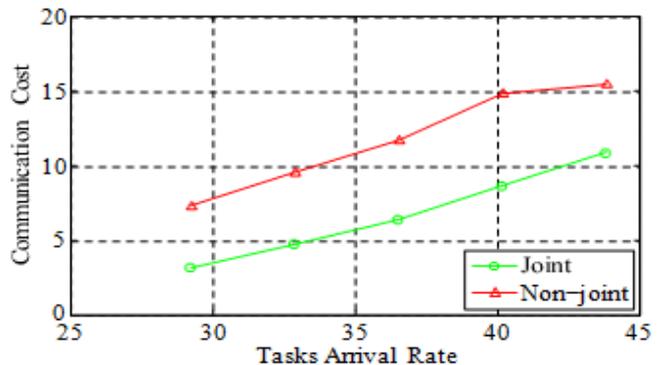


Fig. 4(b): Effect of Communication Cost on Task Arrival Rate

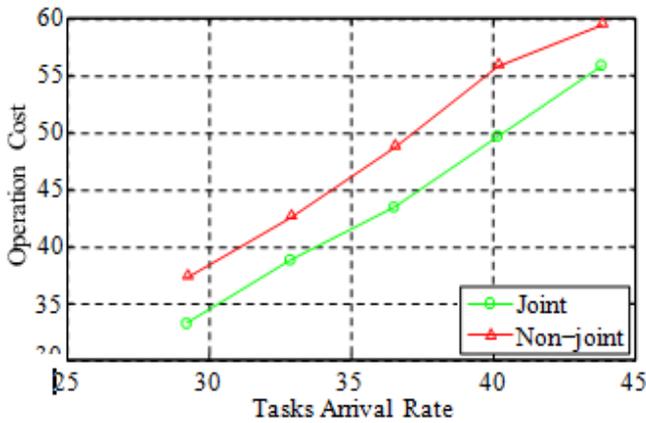


Fig. 4(c): Effect of Overall cost on Task Arrival Rate

Then, we investigate how the task arrival rate affects the cost via varying its value from 29.2 to 43.8. The evaluation results are shown in Fig. 4. We first notice that the total cost shows as an increasing function of the task arrival rates in both algorithms. This is because, to process more requests with the guaranteed QoS, more computation resources are needed. This leads to an increasing number of activated servers and hence higher server cost, as shown in Fig. 4(a). An interesting fact noticed from Fig. 4(a) is that “Joint” algorithm requires sometimes higher server cost than “Non-joint”. This is because the first phase of the “Non-joint” algorithm greedily tries to lower the server cost. However, “Joint” algorithm balances the tradeoff between server cost and communication cost such that it incurs much lower communication cost and thus better results on the overall cost, compared to the “Non-joint” algorithm, as shown in Fig. 4(b) and Fig. 4(c), respectively.

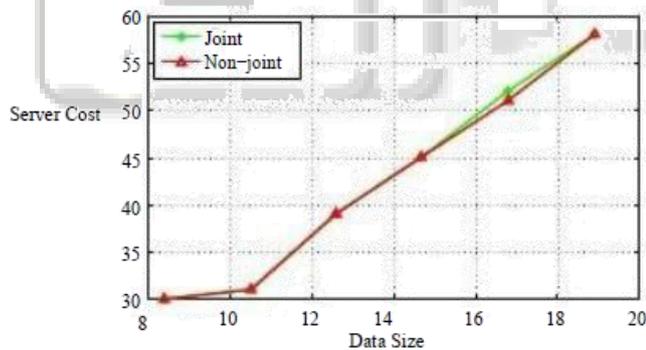


Fig. 5(a): Effect of Server cost on data size

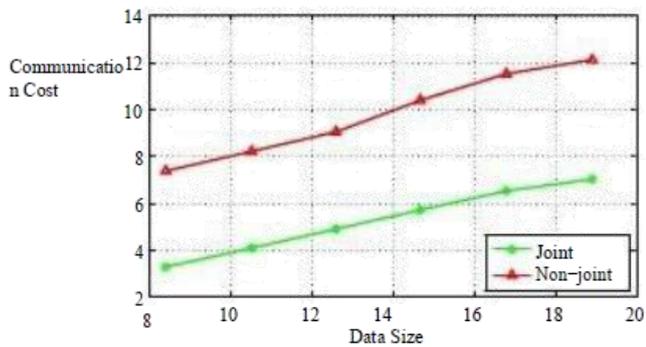


Fig. 5(b): Effect of Communication cost on data size

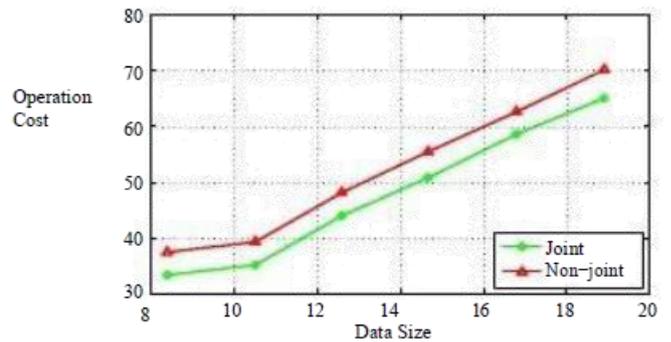


Fig. 5(c): Effect of Overall cost on data size

Fig. 5 illustrates the cost as a function of the total data chunk size from 8.4 to 19. Larger chunk size leads to activating more servers with increased server cost as shown in Fig. 5(a). At the same time, more resulting traffic over the links creates higher communication cost as shown in Fig. 5(b). Finally, Fig. 5(c) illustrates the overall cost as an increasing function of the total data size and shows that our proposal outperforms “Non-joint” under all settings.

Next we show in Fig. 6 the results when the expected maximum response time D increases from 20 to 100. From Fig. 6(a), we can see that the server cost is a non-increasing function of D . The reason is that when the delay requirement is very small, more servers will be activated to guarantee the QoS. Therefore, the server costs of both algorithms decrease as the delay constraint increases. A looser QoS requirement also helps find cost-efficient routing strategies as illustrated in Fig. 6(b). Moreover, the advantage of our “Joint” over “Non-joint” can be always observed in Fig. 6(c).

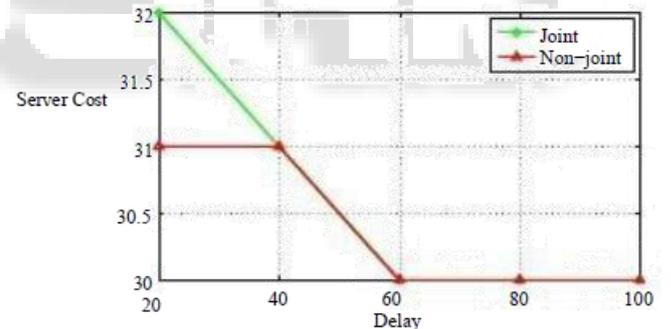


Fig. 6(a): Effect of Server cost on expected task completion delay

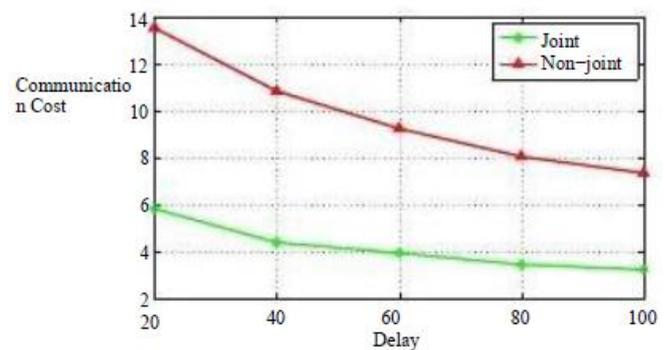


Fig. 6(b): Effect of Communication cost on expected task completion delay

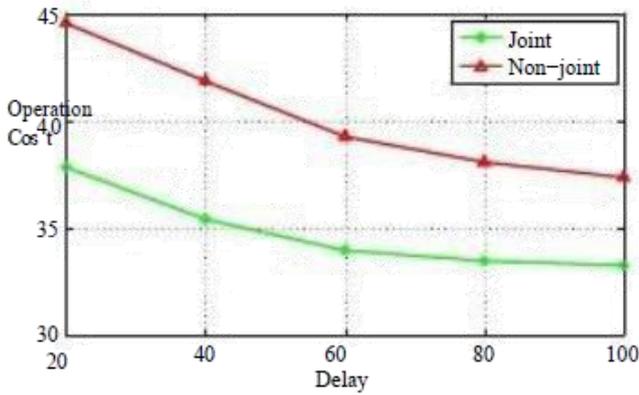


Fig. 6(c): Effect of Overall cost on expected task completion delay

Finally, Fig. 7 investigates the effect of the number of replicas for each data chunk, which is set from 1 to 6. An interesting observation from Fig. 7(c) is that the total cost first decreases and then increases with the increasing number of replicas. Initially, when the replica number increases from 1 to 4, a limited number of activated servers are always enough for task processing, as shown in Fig. 7(a). Meanwhile, it improves the possibility that task and its required data chunk are placed on the same server. This will reduce the communication cost, as shown in Fig. 7(b). When the replica number becomes large, no further benefits to communication cost will be obtained while more servers must be activated only for the purpose of providing enough storage resources. In this case, the server and hence the overall costs shall be increased, as shown in Fig. 7(a) and Fig. 7(c), respectively.

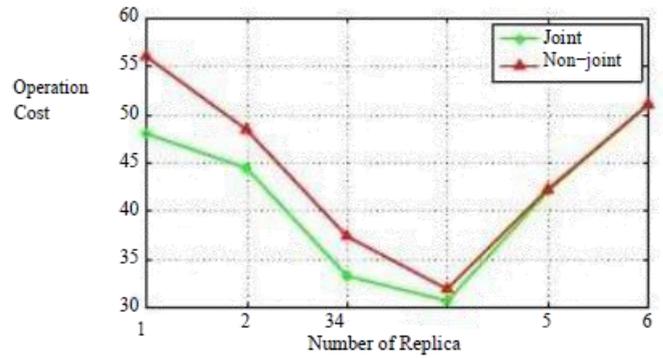


Fig. 7(c): Effect of Overall cost on no. of replica

Our discovery that the optimal number of chunk replicas is equal to 4 under the network setting above is verified one more time by solving the formulation given in Section 4(D) that is to minimize the number of replicas with the minimum total cost. Additional results are given under different settings via varying the task arrival rate and chunk size in the ranges of $[0.1, \lambda_U]$ and $[\phi_L, 1.0]$, respectively, where a number of combinations of (λ_U, ϕ_L) are shown in Fig. 8. We observe that the optimal number of replica is a non-decreasing function of the task arrival rate under the same chunk size while a non-increasing function of the data chunk size under the same task arrival rate.

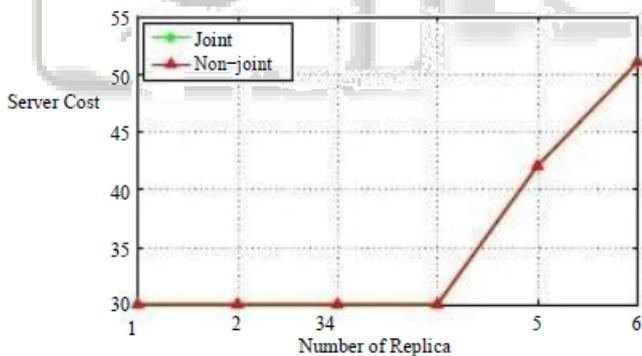


Fig. 7(a): Effect of Server cost on no. of replica

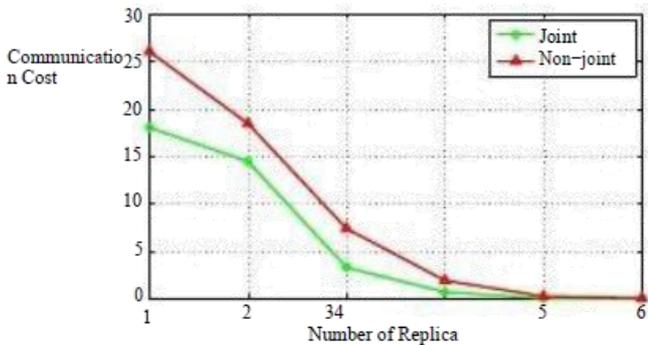


Fig. 7(b): Effect of Communication cost on no. of replica

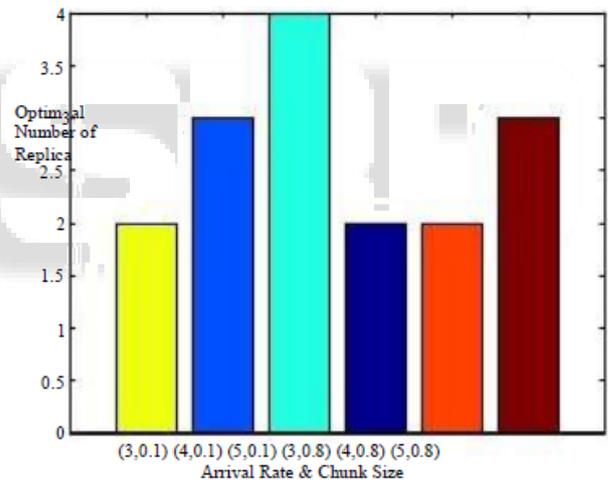


Fig. 8: Optimal number of replica

VII. CONCLUSION

In this paper, we jointly study the data placement, task assignment, data center resizing and routing to minimize the overall operational cost in large-scale geo-distributed data centers for big data applications. We first characterize the data processing process using a two-dimensional Markov chain and derive the expected completion time in closed-form, based on which the joint optimization is formulated as an MINLP problem. To tackle the high computational complexity of solving our MINLP, we linearize it into an MILP problem. Through extensive experiments, we show that our joint-optimization solution has substantial advantage over the approach by two-step separate optimization. Several interesting phenomena are also observed from the experimental results.

REFERENCES

- [1] "Data Center Locations," [http:// www.google.com/about/data-centers/inside/locations/index.html](http://www.google.com/about/data-centers/inside/locations/index.html).
- [2] R. Raghavendra, P. Ranganathan, V. Talwar, Z. Wang, and X. Zhu, "No "Power" Struggles: Coordinated Multi-level Power Management for the Data Center," in *Proceedings of the 13th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. ACM, 2008, pp. 48–59.
- [3] L. Rao, X. Liu, L. Xie, and W. Liu, "Minimizing Electricity Cost: Optimization of Distributed Internet Data Centers in a Multi- Electricity-Market Environment," in *Proceedings of the 29th International Conference on Computer Communications (INFOCOM)*. IEEE, 2010, pp. 1–9.
- [4] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. Andrew, "Greening Geographical Load Balancing," in *Proceedings of International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*. ACM, 2011, pp. 233–244.
- [5] R. Urgaonkar, B. Urgaonkar, M. J. Neely, and A. Sivasubramaniam, "Optimal Power Cost Management Using Stored Energy in Data Centers," in *Proceedings of International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*. ACM, 2011, pp. 221–232.
- [6] B. L. Hong Xu, Chen Feng, "Temperature Aware Workload Management in Geo-distributed Data centers," in *Proceedings of International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*. ACM, 2013, pp. 33–36.
- [7] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," *Communications of the ACM*, vol. 51, no. 1, pp. 107–113, 2008.
- [8] S. A. Yazd, S. Venkatesan, and N. Mittal, "Boosting energy efficiency with mirrored data block replication policy and energy scheduler," *SIGOPS Oper. Syst. Rev.*, vol. 47, no. 2, pp. 33–40, 2013.
- [9] Marshall and C. Roadknight, "Linking cache performance to user behaviour," *Computer Networks and ISDN Systems*, vol. 30, no. 223, pp. 2123 – 2130, 1998.
- [10] H. Jin, T. Cheoherngnarn, D. Levy, A. Smith, D. Pan, J. Liu, and N. Pissinou, "Joint Host-Network Optimization for Energy-Efficient Data Center Networking," in *Proceedings of the 27th International Symposium on Parallel Distributed Processing (IPDPS)*, 2013, pp. 623–634.
- [11] Qureshi, R. Weber, H. Balakrishnan, J. Gutttag, and B. Maggs, "Cutting the Electric Bill for Internet-scale Systems," in *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM)*. ACM, 2009, pp. 123–134.
- [12] X. Fan, W.-D. Weber, and L. A. Barroso, "Power Provisioning for A Warehouse-sized Computer," in *Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA)*. ACM, 2007, pp. 13–23.
- [13] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar, "Benefits and Limitations of Tapping Into Stored Energy for Datacenters," in *Proceedings of the 38th Annual International Symposium on Computer Architecture (ISCA)*. ACM, 2011, pp. 341–352.
- [14] P. X. Gao, A. R. Curtis, B. Wong, and S. Keshav, "It's Not Easy Being Green," in *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM)*. ACM, 2012, pp. 211–222.
- [15] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, and C. Hyser, "Renewable and Cooling Aware Workload Management for Sustainable Data Centers," in *Proceedings of International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*. ACM, 2012, pp. 175–186.
- [16] M. Sathiamoorthy, M. Asteris, D. Papailiopoulos, A. G. Dimakis, R. Vadali, S. Chen, and D. Borthakur, "Xoring elephants: novel erasure codes for big data," in *Proceedings of the 39th international conference on Very Large Data Bases*, ser. PVLDB'13. VLDB Endowment, 2013, pp. 325–336.
- [17] Hu, N. Carvalho, L. Laera, and T. Matsutsuka, "Towards big linked data: a large-scale, distributed semantic data storage," in *Proceedings of the 14th International Conference on Information Integration and Web-based Applications & Services*, ser. IIWAS '12. ACM, 2012, pp. 167–176.
- [18] J. Cohen, B. Dolan, M. Dunlap, J. M. Hellerstein, and C. Welton, "Mad skills: new analysis practices for big data," *Proc. VLDB Endow.*, vol. 2, no. 2, pp. 1481–1492, 2009.
- [19] R. Kaushik and K. Nahrstedt, "T*: A data-centric cooling energy costs reduction approach for Big Data analytics cloud," in *2012 International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2012, pp. 1–11.
- [20] F. Chen, M. Kodialam, and T. V. Lakshman, "Joint scheduling of processing and shuffle phases in mapreduce systems," in *Proceedings of the 29th International Conference on Computer Communications (INFOCOM)*. IEEE, 2012, pp. 1143–1151.
- [21] H. Shachnai, G. Tamir, and T. Tamir, "Minimal cost reconfiguration of data placement in a storage area network," *Theoretical Computer Science*, vol. 460, pp. 42–53, 2012.
- [22] S. Agarwal, J. Dunagan, N. Jain, S. Saroiu, A. Wolman, and H. B. hogan, "Volley: Automated Data Placement for Geo-Distributed Cloud Services," in *The 7th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2010, pp. 17–32.
- [23] Cidon, R. Stutsman, S. Rumble, S. Katti, J. Ousterhout, and M. Rosenblum, "MinCopssets: Derandomizing Replication In Cloud Storage," in *The 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2013.
- [24] S. Gunduz and M. Ozsu, "A poisson model for user accesses to web pages," in *Computer and Information Sciences - ISCIS 2003*, ser. Lecture Notes in Computer

Science. Springer Berlin Heidelberg, 2003, vol. 2869, pp. 332–339.

- [25] L. Kleinrock, “The latency/bandwidth tradeoff in gigabit networks,” *Communications Magazine, IEEE*, vol. 30, no. 4, pp. 36–40, 1992.
- [26] “Gurobi,” www.gurobi.com.

