Future Generation Computer Systems 105 (2020) 162–174

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

QoS provisioning for various types of deadline-constrained bulk data transfers between data centers☆

Aiqin Hou [a], Chase Q. Wu [b,*], Ruimin Qiao [a], Liudong Zuo [c], Michelle M. Zhu [d], Dingyi Fang [a], Weike Nie [a], Feng Chen [a]

[a] School of Information Science and Technology, Northwest University, Xi'an, Shaanxi 710127, China
[b] Department of Computer Science, New Jersey Institute of Technology, Newark, NJ 07102, USA
[c] Computer Science Department, California State University, Dominguez Hills, Carson, CA 90747, USA
[d] Department of Computer Science, Montclair State University, Montclair, NJ 07043, USA

ARTICLE INFO

Article history:
Received 21 February 2019
Received in revised form 27 September 2019
Accepted 28 November 2019
Available online 3 December 2019

Keywords:
Big data
Data center
High-performance networks
Software-defined networking
Bandwidth scheduling

ABSTRACT

An increasing number of applications in scientific and other domains have moved or are in active transition to clouds, and the demand for big data transfers between geographically distributed cloud-based data centers is rapidly growing. Many modern backbone networks leverage logically centralized controllers based on software-defined networking (SDN) to provide advance bandwidth reservation for data transfer requests. How to fully utilize the bandwidth resources of the links connecting data centers with guaranteed quality of service for each user request is an important problem for cloud service providers. Most existing work focuses on bandwidth scheduling for a single request for data transfer or multiple requests using the same service model. In this work, we construct rigorous cost models to quantify user satisfaction degree, and formulate a generic problem of bandwidth scheduling for multiple deadline-constrained data transfer requests of different types to maximize the request scheduling success ratio while minimizing the data transfer completion time of each request. We prove this problem to be not only NP-complete but also non-approximable, and hence design a heuristic algorithm. For performance evaluation, we establish a proof-of-concept emulated SDN testbed and also generate large-scale simulation networks. Both experimental and simulation results show that the proposed scheduling scheme significantly outperforms existing methods in terms of user satisfaction degree and scheduling success ratio.

1. Introduction

Similar to Internet Service Providers (ISPs), Cloud Service Providers (CSPs) aim to satisfy as many customers as possible with high throughput. Most of the leading CSPs deploy geographically-distributed data centers (DCs) to provide various types of services to their customers. These DCs are typically connected by links of high bandwidths across wide-area networks (WANs).

The network infrastructures that connect these geographically distributed DCs cost millions of dollars annually [2], but unfortunately have not been fully utilized. The average utilization of network resources on even those busy inter-DC links is 40%–60% [2], and is 30%–40% [3] on many others, partially due to the traditional best-effort transfer method on the Internet. As the number of cloud-based applications continues to increase, it has become a significant challenge to fully utilize the bandwidth resources of inter-DC network links to accommodate as many data transfer requests as possible and meanwhile maximize the throughput of the entire network system.

Nowadays, the backbones of many WANs employ new technologies to create high-performance networks (HPNs) (e.g., ESnet [4], Internet2 [5], etc.), which provide the capability of advance bandwidth reservation over dedicated channels provisioned by circuit-switching infrastructures or IP-based tunneling techniques for big data transfer. Particularly, the emerging software-defined networking (SDN) technologies greatly facilitate HPN deployments, and in fact, many HPNs have incorporated SDN capabilities into their network infrastructures to provide better Quality of Service (QoS). Such networks feature a virtual single-switch abstraction on top of data planes that employ both a bandwidth reservation system and SDN concepts [6,7]. Generally, bandwidth reservation allows a batch of user transfer requests accumulated over a period of time to be scheduled collectively in advance, and has proven to be an effective solution to providing

☆ Some preliminary results in this manuscript were published in INDIS'18 in conjunction with SC'18 Hou et al. (2018) [1].

* Corresponding author.
E-mail address: chase.wu@njit.edu (C.Q. Wu).

https://doi.org/10.1016/j.future.2019.11.039

QoS guaranteed transfer services and meanwhile achieving a high utilization of network resources.

There exist different types of inter-DC data transfers, among which, bulk data transfer requests (BDTRs) for large data volumes on the order of terabytes to petabytes with deadline constraints account for a major portion of traffic (e.g., 85%–95% in some WANs) [2,8–10].

However, most existing solutions for BDTRs are tailored for private cloud services, hence limiting their generalization and scope of application. For example, both Software-Driven WAN (SWAN) [2] and B4 [3] take traffic engineering approaches to improve the inter-DC WAN utilization by considering traffic characteristics and priorities (e.g., interactive > elastic > background). However, neither of them addresses the deadline constraint of BDTRs, one of the most common performance requirements from users [9]. As an increasing number of applications in scientific and many other domains have migrated from local computing and storage platforms to clouds, the demand for inter-DC data transfer with different types of BDTRs is rapidly growing, but the bandwidth scheduling problem in the emerging cloud environment still remains largely unexplored.

In this paper, we investigate a bandwidth scheduling problem for two types of BDTRs with fixed or variable bandwidth. Given multiple such BDTRs, we aim to fully utilize inter-DC link bandwidth resources and schedule as many BDTRs as possible while minimizing the earliest complete time (ECT) of each request. Specifically, we construct a rigorous cost model, define a new performance metric named user satisfaction degree, and then formulate a generic problem, Bandwidth Scheduling for Multiple Requests of Various Types, referred to as BS-MRVT. We prove BS-MRVT to be not only NP-complete but also non-approximable, and then propose an efficient heuristic scheduling algorithm. We conduct proof-of-concept experiments on a Mininet-based emulated testbed and also extensive simulations for bandwidth scheduling in both simulated and real-life networks. Both experimental and simulation results show that our proposed algorithm significantly outperforms existing methods in terms of user satisfaction degree and scheduling success ratio.

The rest of this paper is organized as follows. We conduct a survey of related work in Section 2. We construct network models and formulate BS-MRVT with complexity analysis in Section 3. We design the algorithm with a detailed illustration in Section 4. We conduct performance evaluation in Section 5 and conclude our work in Section 6.

2. Related work

There have been a number of successful research efforts in making full use of network resources for bulk data transfer between data centers using centralized traffic engineering techniques. For example, Jain et al. developed B4 [3], which is able to globally schedule massive bandwidth requirements at a modest number of sites. Hong et al. designed SWAN [2], a centrally controlling system that enables inter-DC WANs to carry more traffic. Kandula et al. developed TEMPUs [8], an online temporal planning scheme that packs long-running transfers across network paths and future time steps, while leaving enough capacity slack for future high-priority requests. However, the aforementioned work does not consider data transfer deadline. Moreover, most of these existing schemes are tailored for their private clouds and may not work as well for public clouds.

Zhang et al. designed Amoeba [9] for inter-DC WANs to ensure that data transfer be completed before a hard or soft deadline through centralized traffic engineering and bandwidth reservation scheme. Jin et al. presented Owan [10] to schedule bulk data transfer requests to meet their deadlines in a modern WAN infrastructure with an intelligent optical network-layer. Nandagopal et al. proposed GRESE, which attempts to minimize the overall bandwidth costs by leveraging the flexible nature of the deadlines of certain bulk data transfers [11]. Lin and Wu formulated and solved a class of transport-support workflow optimization problems for end-to-end data transfer path compositions and transport method selections to meet data transfer requests [12].

There also exists some work on routing policy for bulk data transfer. Yassine et al. designed a multi-rate bandwidth-on-demand broker [13], which employs a scheduling algorithm that considers both delay tolerant and intolerant multimedia data by using the concept of dense wavelength division multiplexing in backbone optical networks. Bandwidth-on-demand provides dynamic multi-rate for communication between geographically distributed cloud-based data centers to meet the requirements specified in the service level agreement. In [14], Shu and Wu investigated the problem of bandwidth scheduling for energy efficiency in high-performance networks.

Some efforts have been made to address the problem of scheduling multiple bandwidth reservation requests for bulk data transfer over a single path. Sharma et al. studied the problem of accommodating as many bandwidth reservation requests (BRRs) as possible while minimizing the total time to complete all data transfers on the same path. The problem was proved to be NP-hard and a heuristic algorithm was proposed. Zuo et al. investigated the problem of scheduling as many concurrent bandwidth reservation requests as possible on one dedicated channel in an HPN [15]. Wang et al. studied a periodic bandwidth scheduling problem to maximize the number of satisfied user requests for bandwidth reservation with deadline constraint on a fixed network path [16]. Note that these problems only consider a single network path, which is not common in real applications.

Some researchers investigated the problem of bandwidth scheduling for multiple user requests in HPNs. Zuo and Zhu studied the problem of scheduling as many concurrent bandwidth reservation requests as possible in an HPN while achieving the average earliest complete time (ECT) and the average shortest duration (SD) of scheduled BRRs [17]. Given multiple bandwidth reservation requests in a batch awaiting to be scheduled in a dedicated network, Zuo et al. studied two scheduling maximization problems: maximize the amount of data to be transferred and maximize the number of requests to be scheduled [18]. Srinivasan et al. proposed a bandwidth allocation scheme that flexibly and adaptively allocates bandwidth for big data transfer requests with the objective to maximize the acceptance ratio of the requests while satisfying the deadline constraints [19]. Wang et al. leveraged the temporal and spacial characteristics of inter-DC bulk data traffic and investigated the problem of scheduling multiple bulk data transfers to reduce network congestion [20].

Given one data transfer request aiming to achieve the ECT, Lin and Wu in [21] investigated bandwidth scheduling with an exhaustive combination of different path and bandwidth constraints: (i) fixed path with fixed bandwidth (FPFB), (ii) fixed path with variable bandwidth (FPVB), (iii) variable path with fixed bandwidth (VPFB), and (iv) variable path with variable bandwidth (VPVB). VPFB and VPVB are further divided into VPFB-0/1 and VPVB-0/1, respectively, depending on whether or not the path switch delay is negligible. They also analyzed the complexities of these problems and proposed optimal and heuristic algorithms.

The above survey of related work shows that although bandwidth scheduling has been studied extensively in various contexts in the literature, there are very limited efforts to address the problem of scheduling multiple concurrent bandwidth reservation requests of different types for big data transfer among data centers, which is the focus of our research.