Highly Efficient Data Migration and Backup for Big Data Applications in Elastic Optical Inter-Data-Center Networks

Ping Lu, Liang Zhang, Xiahe Liu, Jingjing Yao, and Zuqing Zhu

Abstract

This article discusses the technologies for realizing highly efficient data migration and backup for big data applications in elastic optical inter-data-center (inter-DC) networks. We first describe the impacts of big data applications on underlying network infrastructure and introduce the concept of flexible-grid elastic optical inter-DC networks. Then we model the data migration in such networks as dynamic anycast and propose several efficient algorithms. Joint resource defragmentation is also discussed to further improve network performance. For efficient data backup, we leverage a mutual backup model and investigate how to avoid the prolonged negative impacts on DCs' normal operation by minimizing the DC backup window.

Nowadays, the emergence of data-intensive applications, such as e-science, e-commerce, and teleconference, has brought us into the "big data" era. Different from the conventional ones, big data applications can generate huge volumes of data that the conventional systems can hardly capture, manage, store and analyze [1]. Recent advances on cloud computing have indicated that big data applications running on multi-data-center (multi-DC) systems can provide low-latency, high-quality, non-disruptive services to end users [2, 3]. Recently, a few big enterprises have built multi-DC systems and invested R&D efforts on the technologies for efficient data migration; for example, Google has deployed Effingo [4] for migrating large-scale data among its DCs. In [5], Laoutaris *et al.* developed NetStitcher, which used a relay-based data migration scheme that leverages in-network storage for buffering and forwards bulk data toward its destination when available bandwidth exists. However, the huge throughput associated with data migration and backup is still an open challenge for the underlying networks connecting multi-DC systems.

With agile spectrum management in the optical layer, flexible-grid elastic optical networks (EONs) allow the optical spectra to be allocated at the granularity of a few gigahertz or even smaller [6]. Consequently, the optical layer in EONs can react directly to variable bandwidth demands from the clients, and hence EONs have been considered as a promising underlying network infrastructure for supporting multi-DC systems and the big data applications running on them. This article discusses the technologies for realizing highly efficient data migration and backup for big data applications in elastic optical inter-DC networks. We first describe the characteristics of big data applications and cloud computing, and explain their impacts on the underlying inter-DC networks. Then we introduce the concept of flexible-grid elastic optical inter-DC networks and model the data migration in them as a dynamic anycast problem. Several efficient anycast algorithms are proposed, and joint resource defragmentation is also introduced to further improve network performance. For efficient data backup, we leverage a mutual backup model and investigate how to avoid the prolonged negative impacts on DCs' normal operation by minimizing the DC backup window. Finally, we discuss future challenges and summarize the article.

Ping Lu, Liang Zhang, Xiahe Liu, Jingjing Yao, and Zuqing Zhu are with the University of Science and Technology of China.

Big Data and Data Center Networks

Big Data

Big data applications can have different implementations, but they usually have common characteristics from some perspectives, for example, the volume, velocity and variety, the famous "3Vs" [1]. Here, volume describes the size of the data, and as the data volume becomes huge with the evolution of big data applications, one needs more computing/storage resources, as well as more efficient networking systems. Velocity reflects the frequency at which data is generated, processed, and transmitted. To adapt to the velocity requirement, large-scale and distributed service architecture and high-capacity underlying networks are needed. Variety refers to the formats and types of the data. Since both structured and unstructured data with various types need to be processed and transferred, big data applications require not only highly efficient data storage systems, but also flexible network infrastructure.

Cloud Computing

By leveraging resource virtualization, distributed processing, and other advanced technologies, cloud computing utilizes DCs and network equipment in the Internet to build an integrated platform and supports big data applications in a scalable and on-demand manner. Several programming models and file system architectures have been proposed for cloud computing. For instance, MapReduce, designed by Google, is a programming model to simplify the data processing in large DCs. Specifically, it maps massive data to hundreds or even thousands of servers or virtual machines (VMs) and executes data processing in a parallel way. To cooperate with MapReduce, the Google file system (GFS) is developed to manage the data storage for each server or VM.

Geo-Distributed Multi-DC Systems

With the rapid rise of big data and cloud computing, DC networks attract the interest of both industry and academia. Since cloud services can integrate the computing, storage, and bandwidth resources in inter- and intra-DC networks seamlessly and effectively, big data applications can leverage them and provide high-quality services to numerous end users. Recently, there is an increasing trend to deploy geographically distributed (geo-distributed) cloud systems with multiple DCs [7]. The rationales behind this trend are listed as follows:

- The service provider can run multiple DCs for different business or services, which are operated by geographically dispersed departments/branches.
- Deploying DCs close to end users improves user experience by minimizing network latency, and helps to promote the services at different geographical locations.
- Geo-distributed multi-DC systems provide additional redundancy, which is vital to ensure non-disruptive services when a natural disaster or human misconduct happens.

On the other hand, geo-distributed multi-DC systems also bring challenges to the underlying inter-DC networks. First of all, the applications may require a huge bulk of data that is located on different DCs, which causes massive data migration among DCs and consumes huge bandwidth. Also, periodic data backups need to accomplish bulk data transfers among DCs from time to time. Moreover, the data generated by the end users of big data applications can be huge, and how to deliver it to the DCs is an issue as well. In addition, some big data applications' features can make the traffic of data migration and backup exhibit high burstiness. To this end, we need to consider both the data migration from end users to DCs and that among DCs when designing the underlying inter-DC networks for geo-distributed multi-DC systems.

Inter-DC Networks

The geo-distributed multi-DC systems are interconnected with an inter-DC network. Generally, two types of traffic contribute to the data transfers in inter-DC networks. The first one is the user-to-DC (U2D) traffic, which is generated by the end users. For example, end users can submit data and tasks to the applications running on DCs. The other one is the DC-to-DC (D2D) traffic, which refers to the data migration and backup among DCs.

Big data applications result in some special characteristics of inter-DC networks. For instance, the applications can have different bandwidth and service time requirements, which calls for inter-DC networks that facilitate agile bandwidth management in both the time and spectrum domains with fine granularity. Also, unlike the conventional telecom demands that only consume bandwidth resources, the demands in inter-DC networks consume both the bandwidth resources in the network and the computing/storage resources in the DCs. That is to say, in addition to allocating necessary bandwidth resources to a demand, one also needs to reserve enough computing/storage resources in the DCs. Another interesting characteristic of the demands in inter-DC networks is that their destination DCs can be flexible, and hence the anycast scheme is commonly used. This is because multi-DC systems usually make multiple replications of the data or VMs available in different DCs.

Elastic Optical Networks

Issues with Conventional Fixed-Grid Optical Networks

In order to accommodate the overwhelming traffic growth from telecom and datacom applications, optical networks play an irreplaceable role in the Internet due to the tremendous bandwidth on optical fibers. Today's optical networks are implemented with dense wavelength-division multiplexing (DWDM) systems, which operate within wavelength channels/grids whose bandwidths are fixed. However, fixed-grid DWDM systems only provide limited scalability and flexibility in the optical layer, which makes the transport infrastructure too rigid to adapt to the uncertainty and heterogeneity of the traffic across inter-DC networks. For instance, the 50 GHz International Telecommunication Union (ITU) wavelength grid divides the optical spectrum into fixed 50 GHz frequency slots, but it will be difficult for bit rates at 400 Gb/s or higher to fit into this scheme [6].

Another issue with fixed-grid DWDM systems is the single-carrier scheme they use for data transmissions. Even if the ITU wavelength grid can be upgraded to a wider one, the corresponding high-speed (e.g, 400 Gb/s and beyond) data transmissions with a single carrier can hardly support long transmission distances all-optically due to the physical impairments. Hence, for carrying geo-distributed multi-DC systems, the underlying optical network needs repeated optical-to-electrical-to-optical (O/E/O) regenerations. Nevertheless, these O/E/O regenerations associate with relatively high capital expenditures (CAPEX) and operational expenditures (OPEX) due to the equipment cost and power consumption.

Last, and most important, it is difficult for fixed-grid DWDM networks to provide seamless and efficient support to big data applications with bandwidth requirements that are huge but can vary quickly at fine granularity. DWDM networks can only establish lightpaths and allocate bandwidth according to a coarse wavelength grid. Unfortunately, this scheme only provides low spectral efficiency when the carried traffic is highly dynamic.

Flexible-Grid Elastic Optical Networks

In order to properly address the issues with fixed-grid DWDM networks, one needs "elastic" optical networks equipped with bandwidth-variable (BV) transponders and switches that can allocate bandwidth at finer granularity and establish lightpath adaptively according to the actual traffic needs in inter-DC networks. Four elements are essential for EONs: BV transponders (BV-Ts), BV wavelength selective switches (BV-WSSs), flexible wavelength grid, and intelligent control plane.

Figure 1a shows the operation principles of BV-Ts and BV-WSSs in the process of data transmission. BV-Ts set up lightpaths for client traffic with just enough bandwidth resources, while BV-WSSs allow flexible spectra to be switched correctly from the input to the output ports. Figure 1b compares the existing fixed-grid wavelength plan with a flexible-grid one. Note that ITU Telecommunication Standardization Sector (ITU-T) has already started to define the flexible-grid wavelength plan in its Recommendation G.694.1 [8]. The flexible nature of EONs makes the intelligent control plane a "must have" for cost-effective resource management. For example, in EONs, the famous routing and wavelength assignment (RWA) problem in DWDM networks transforms into routing and spectrum assignment (RSA), which has to manipulate flexible spectra and hence requires more sophisticated algorithms [9].

The major difference between the flexible-grid EONs and the fixed-grid DWDM networks is that EONs can provision

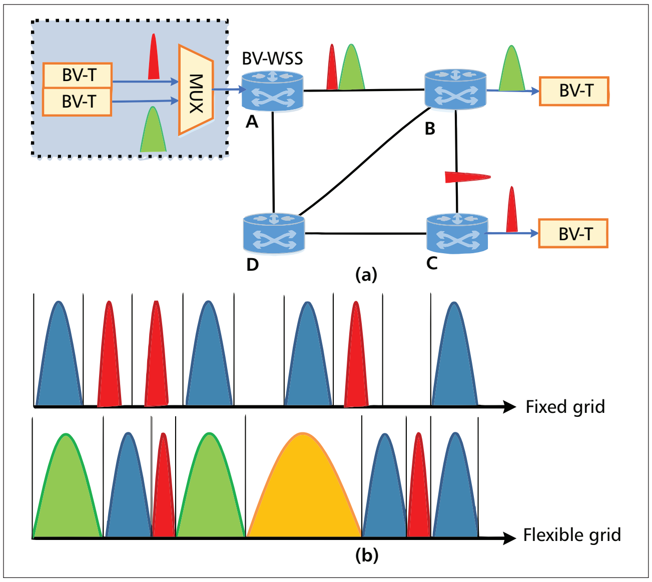

Figure 1. Flexible-grid elastic optical networks: a) operation principles of BV-Ts and BV-WSSs; b) fixed-grid and flexible-grid bandwidth allocation.

low and ultra-high bit rate demands efficiently with sub-wavelength frequency slots (FSs) and super-channels, respectively. Here, a sub-wavelength FS refers to spectrum occupation that is less than a full wavelength channel, while a supper-channel contains multiple closely packed FSs. Moreover, the BV-Ts can choose the modulation formats adaptively to accommodate various qualities of transmission [6], which brings another level of elasticity.

Enabling Techniques for Elastic Optical Networking

Multi-carrier transmission techniques, such as coherent optical orthogonal frequency-division multiplexing (CO-OFDM) [10] and Nyquist-WDM [11], have been proposed and demonstrated for implementing BV-Ts. These techniques enable the BV-Ts to groom the capacities of several spectrally contiguous FSs and achieve high-speed transmissions over them. Then the BV-Ts can change a lightpath's bandwidth allocation by adjusting the number of assigned FSs. Meanwhile, thanks to the technology advances on liquid crystal-on-silicon wavelength-selective switches (LCOS-WSSs), a BV-WSS can achieve switching granularity of 6.25 GHz or less [12].

Since EONs allocate variable-sized spectrum slices to lightpaths, there will be fragmentation in the spectrum over time for dynamic network operations. As shown in Fig. 2a, bandwidth fragmentation creates non-aligned, isolated, and small-sized unused spectrum slices in the optical spectrum, which is similar to the memory fragmentation in a computer hard disk [13]. Since these spectrum slices can hardly be used for future demands, fragmentation leads to low spectrum utilization in EONs. In order to relieve the fragmentation, we need a mechanism to periodically reconfigure the spectrum allocations in EONs (as in Fig. 2b) [13], that is, so-called defragmentation. Note that for minimizing traffic disruptions, defragmentation needs to precisely coordinate the spectrum reallocations at the BV-Ts and BV-WSSs along affected lightpaths. A hop-tuning-based spectrum retuning technique has been proposed and experimentally demonstrated in [14], which could support full-spectrum retuning within 1 μs.

Elastic Optical Inter-DC Networks

Figure 3 illustrates the architecture of an elastic optical inter-DC network. We have several DCs that are connected to switching nodes locally, while the switching nodes are interconnected with optical fibers. The underlying optical network uses flexible-grid elastic optical networking, and the bandwidth resources on each fiber are divided into fixed-sized FSs to provide the sub-wavelength bandwidth allocation granularity [9]. For each node that connects to a DC locally, we consider the computing/storage capacity of the DC. Then, to provision a U2D or D2D demand that requires certain bandwidth and computing/storage resources, we determine the destination DC, reserve sufficient computing/storage resources on it, and then solve the routing and spectrum assignment (RSA) problem to establish the lightpath.

For the RSA problem, we first find a routing path between the source and destination of a demand and then allocate sufficient spectrum resources on all the fiber links along the path to establish the lightpath. The fixed-sized FSs with a proper bandwidth (e.g., 12.5 GHz) are the smallest units for spectrum allocation, and the demand's bandwidth requirement determines how many FSs are allocated on the lightpath. The operation principle of the BV-Ts determines that we need to allocate spectrally contiguous FSs to a lightpath, satisfying the so-called spectrum-contiguous constraint [6]. The spectrum allocation also needs to satisfy the spectrum-continuity constraint, which is similar to the wavelength-continuity constraint for RWA and requires the lightpath to use the same spectrum on all the fiber links along the path.

Data Migration in Elastic Optical Inter-DC Networks

Data migration in elastic optical inter-DC networks can be formulated as an anycast problem. Based on the working principles of BV-Ts and BV-WSSs, we assume that there are B FSs on each fiber link. A subset of the switching nodes connect to DCs locally and are considered as the DC nodes. Each DC node associates with an attribute, which represents the available

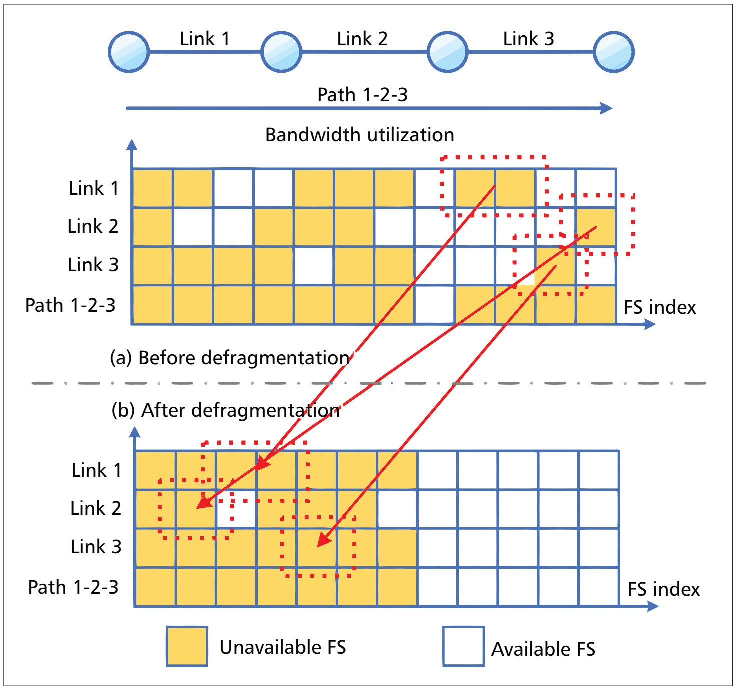

Figure 2. Examples of bandwidth fragmentation and defragmentation.

computing/storage capacity in its local DC. Note that a DC can have separate attributes to describe its computing and storage capacities, but for simplicity, we assume that the computing and storage capacities are correlated and hence can be modeled with one attribute. A U2D or D2D demand for data migration can be modeled as an anycast request $R(s, b, c)$, where s is the source node, b is the required bandwidth in FSs, and c is the computing/storage requirement.

Dynamic Anycast for Data Migration

In elastic optical inter-DC networks, we provision a data migration demand $R(s, b, c)$ with three steps:

1. Selecting DC node(s) as the destination(s)
2. Determining the amount of computing/storage capacities to be allocated on each destination DC to satisfy c
3. Performing RSA to establish the lightpath(s) from s to the destination DC(s) to satisfy b

Since the demands can come and leave on the fly, we have a dynamic anycast problem here. If sufficient resources cannot be allocated (i.e., either b or c cannot be satisfied) at the time when the demand arrives, it is blocked. We aim to minimize the blocking probability, since more blocked demands make the network less effective.

We first consider a simple scenario in which each demand is served with one DC. Then a greedy anycast algorithm with the shortest path routing is designed as follows:

- **Step 1:** Find the DC that has the largest available computing/storage capacity and select it as the demand's destination DC.
- **Step 2:** Calculate the shortest routing path from s to the destination DC.
- **Step 3:** Perform FS assignment on the routing path to satisfy b.

We denote this algorithm as greedy anycast with shortest path for single destination DC (*G-Anycast-SP-Single-DC*). *G-Anycast-SP-Single-DC* is straightforward, but it only considers the computing/storage resources for destination DC selection. Even though choosing the DC that has the largest available capacity in the first step can make the computing/storage loads be distributed evenly among DCs, a valid RSA solution may not be found in the consequent steps.

Therefore, it is desired that the anycast algorithm can consider computing/storage and bandwidth resources jointly when selecting the destination DC. This can be done by defining a metric as the product of the available capacity on the DC and the available bandwidth on the related routing path. Basically, for each source-destination pair in the network, we precalculate K shortest routing paths and then a balanced-load anycast algorithm with K-shortest path routing (*BL-Anycast-KSP-Single-DC*) is designed as follows:

- **Step 1:** Find all the DCs that have sufficient computing/storage capacities to accommodate c.
- **Step 2:** Enumerate all the feasible routing paths to the DCs and calculate the metric for each path-DC pair.
- **Step 3:** Select the path-DC pair that has the largest metric and perform FS assignment on the path to satisfy b.

Note that the *BL-Anycast-KSP-Single-DC* algorithm still may not fully utilize the bandwidth resources in the network, since only one destination DC can be used for each demand. Therefore, we allow a demand to be served by multiple DCs and modify *BL-Anycast-KSP-Single-DC* to a multi-DC-based algorithm, denoted as *BL-Anycast-KSP-Multi-DC*.

We evaluate the performance of the anycast algorithms with the network topology in Fig. 3. For each demand, the parameters s, b, and c are randomly generated. The demands arrive according to a Poisson traffic model with an average arrival rate of λ demands per time unit, and the service time

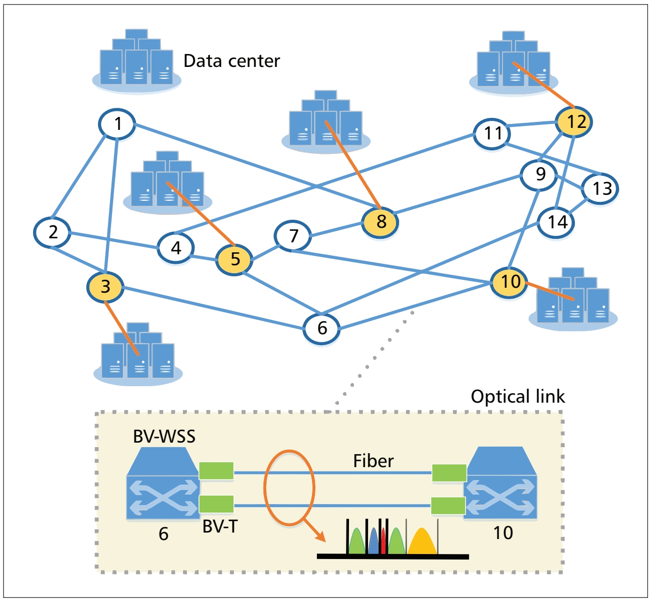

Figure 3. An example of elastic optical inter-DC networks.

of each demand follows the negative exponential distribution with an average of $1/\mu$ time units. Hence, the traffic load can be quantified with λ/μ in Erlangs. Figure 4a shows the results on blocking probability, which indicate that *G-Anycast-SP-Single-DC* provides the highest blocking probability, and the blocking performance of *BL-Anycast-KSP-Multi-DC* is the best.

Note that a data migration demand can be blocked in three cases:

1. The bandwidth resources on the path(s) are sufficient, but the computing/storage resources on the destination DC(s) are not (*DC blocking*).
2. The computing/storage resources are sufficient, but the bandwidth resources are not (*path blocking*).
3. Both resources are insufficient (*combinational blocking*).

We analyze the percentages of these three blocking cases for traffic loads at 170 and 240 Erlangs, and plot the results in Fig. 4b. We observe that the majority of the blocked demands from *G-Anycast-SP-Single-DC* are due to path blocking, which is a clear indication that this algorithm cannot utilize the bandwidth resources in the network wisely. The situation is improved with *BL-Anycast-KSP-Single-DC*, and for simulations with *BL-Anycast-KSP-Multi-DC*, the majority of the blocked demands are due to DC blocking. Therefore, the results confirm that *BL-Anycast-KSP-Multi-DC* can fully utilize the bandwidth resources and make the computing/storage resources become the bottleneck. Note that in multi-DC systems, upgrading DCs by putting in more servers is much easier and less expensive than upgrading the underlying optical network.

One issue with *BL-Anycast-KSP-Multi-DC* is that the demand can be split over many DCs, which makes the cost of BV-Ts at the source node too high and complicates the coordination among the DCs. Fortunately, this can be avoided by defining a path bandwidth granularity g [9]. Specifically, when a demand is provisioned over multiple DCs, the minimum number of FSs to be allocated on each routing path is g.

Resource Defragmentation for Data Migration

Even though the data migration algorithms can achieve high data transfer efficiency, the dynamic U2D or D2D demands from big data applications can fragment the resources in the network and cause a serious problem over time. Hence, we need to design defragmentation (DF) algorithms to alleviate

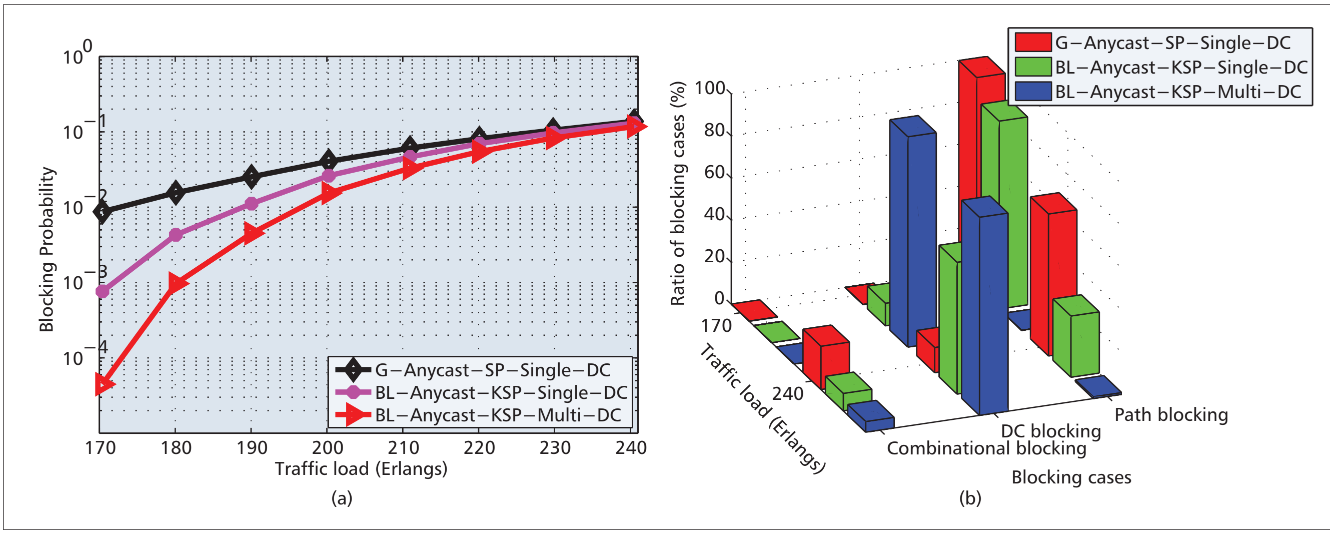

Figure 4. Results on data migration in elastic optical inter-DC networks: a) results on blocking probability; b) percentages of the three blocking cases.

this problem with network reconfiguration. Previous work has investigated bandwidth DF in EONs [13]. However, the DF algorithms were designed for telecom applications in which the demands' destinations are usually fixed. Applying those algorithms to elastic optical inter-DC networks only leads to re-optimization of the bandwidth allocations along the routing paths, while the computing/storage utilizations on the DCs stay the same, as shown in Fig. 5a. However, we need joint DF that can reallocate both the bandwidth and computing/storage utilizations as in Fig. 5b.

In order to control the operational cost of DF and avoid reconfiguring the network too frequently, we only invoke a DF operation when a demand would be blocked otherwise. In each DF operation, a network-wide partial reconfiguration is conducted to consolidate the utilizations of bandwidth and computing/storage resources. Here, "network-wide" reconfiguration ensures fairness in the network. Otherwise, if the DF operation only conducts local reconfiguration to accommodate a particular demand, that demand is treated with an unfairly high priority. "Partial" reconfiguration means that the DF operation only reallocates a portion of the in-service demands.

The joint DF consists of two steps:

1. Selecting the in-service demands to be reallocated
2. Re-optimizing the destination DCs and RSAs of the selected demands

For the first step, we design two selection strategies: one only considers the bandwidth utilization (*Bandwidth-View*), and the other takes both the bandwidth and computing/storage utilization into account (*Joint-View*). For the second step, we leverage the fragmentation-aware RSA [13] and design an anycast-based re-optimization procedure that can reduce bandwidth fragmentation on the fiber links as well as load distribution unevenness in the DCs. We also apply a constraint to restrict the service migration latency.

The DF algorithms are also evaluated with the network topology in Fig. 3, with the difference that this time, we assume each switching node connects to a DC locally. Figure 5c compares the results on blocking probability for three scenarios: without DF, Bandwidth-View DF, and Joint-View DF. It can be seen that compared to the one without DF, both DF algorithms reduce the blocking probability effectively. The Joint-View DF provides lower blocking probability than the Bandwidth-View DF, since it considers both the bandwidth and computing/storage utilization in the network.

Data Backup in Elastic Optical Inter-DC Networks

It is known that DCs are extremely valuable for big data applications, since they usually carry massive data and run thousands or even millions of applications. Nevertheless, this also makes DCs vulnerable to human misconduct and natural disasters. For instance, the 2008 Sichuan earthquake in China destroyed more than 60 enterprise DCs, and the 2011 Tōhoku earthquake and tsunami wiped out tens of DCs. Hence, data backup is not only important but also necessary for DCs. We consider a mutual DC backup model as shown in Fig. 6a, in which each DC in the elastic optical inter-DC networks can work as a backup DC for other DCs. When a DC failure happens, the services on it will be migrated to the corresponding backup DC(s), and hence service disruptions and data losses are minimized.

Minimizing the DC Backup Window

Note that the DC backup window, which is defined as the time required for backing up all the new data on DCs in an inter-DC network, is a key parameter for evaluating DC backup plans. Since DC backup usually needs to transfer huge bulks of data and hence consumes significant bandwidth, a prolonged DC backup window impacts normal operations of the DCs and causes long congestion in the network. Hence, we investigate how to reduce the DC backup window with joint optimization of the selection of backup DCs and the corresponding paths in elastic optical inter-DC networks.

We assume that the network operates on discrete time interval ΔT. At the beginning of each interval, each production DC chooses its backup DC and finds the backup path based on the network status; then the DC backup process will run for ΔT accordingly. Meanwhile, there are two constraints to apply:

1. Each production DC can only select one backup DC, while a backup DC can only receive data from one production DC too.
2. The production and backup DCs should be geographically dispersed and not fall into the same disaster zone.

The first constraint is for simplifying the data indexing during DC backup, while the second one ensures that a single disaster cannot destroy both DCs. The DC backup process ends when all the data on production DCs has been backed up; then the total backup duration is the DC backup window. To minimize the DC backup window, we optimize the backup process for each ΔT, that is, determining the production-back-

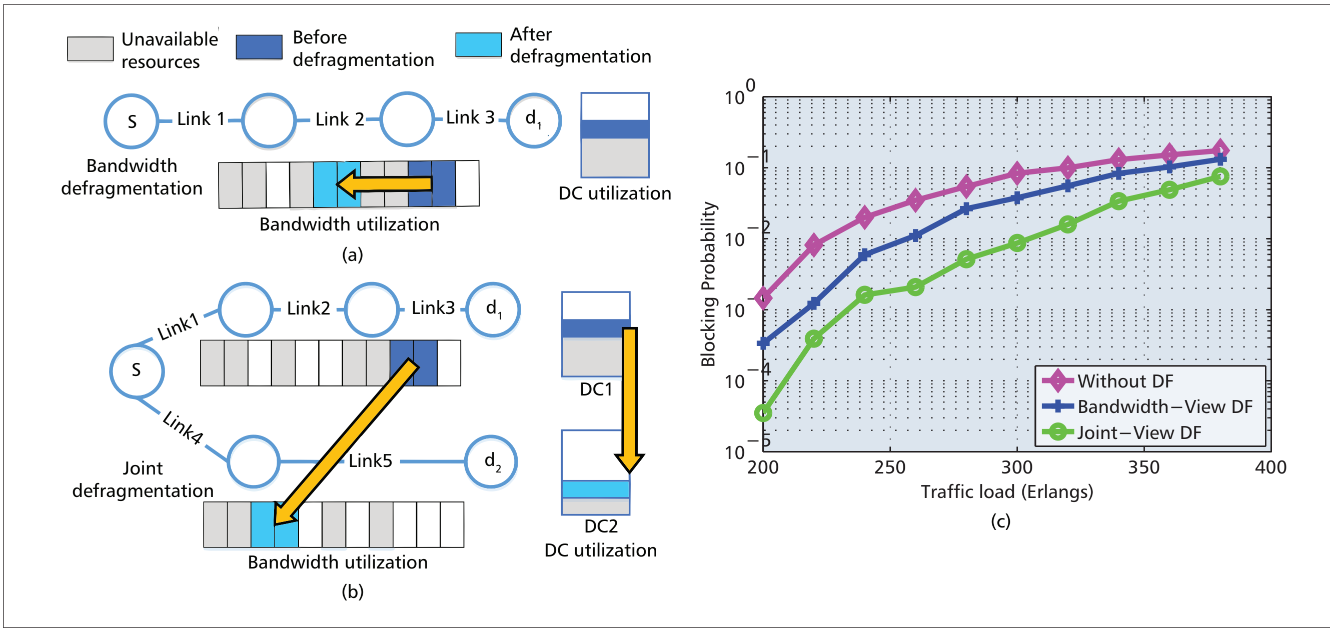

Figure 5. Results on resource defragmentation in elastic optical inter-DC networks: a) bandwidth defragmentation; b) joint defragmentation; and c) results on blocking probability.

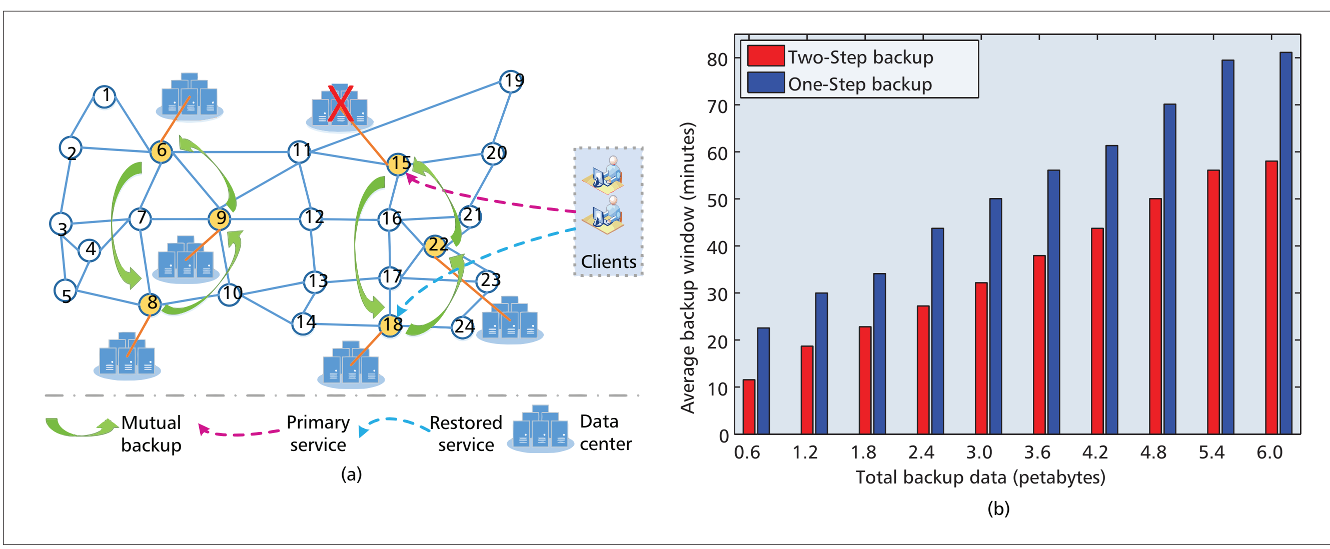

Figure 6. Data backup in elastic optical inter-DC networks: a) mutual backup in elastic optical inter-DC networks; and b) results on DC backup window.

up DC pairs and establishing lightpaths for each DC pair to transfer data. Note that in order to fully explore the network's throughput, we allow the production DC to set up multiple lightpaths to its backup DC simultaneously.

The backup optimization can be done either jointly (*One-Step*) or separately (*Two-Step*). In One-Step, we obtain the backup DC and the corresponding lightpaths for a production DC in one shot based on the global maximum flows in the network. Specifically, we first calculate the maximum flows for all the feasible production-backup DC pairs, and then choose the maximum flow that has the largest throughput (i.e., the global maximum flow) to set up the lightpaths for DC backup. The procedure is repeated until all the production DCs are served for the incoming time interval. In Two-Step, we sort the production DCs based on the amount of data for backup in descending order. For each production DC, we select the nearest feasible DC as the backup DC. Then we calculate the maximum flow for the production-backup DC pair and set up the lightpaths with it for backup for the incoming time interval. We evaluate the performance of these algorithms using the topology in Fig. 6a, where the DCs are located on nodes {6, 8, 9, 15, 18, 22}, and the available bandwidth on each link is uniformly distributed within [1,3] Tb/s. In each simulation, the total data volume for backup is fixed, while the actual data for backup on each production DC is randomly chosen. The time interval is $\Delta T = 10$ min, and the backup algorithm is performed at the beginning of each ΔT until all the data backup is done. The simulation results on the DC backup window are shown in Fig. 6b, where we observe that Two-Step provides a shorter DC backup window than One-Step. This is because Two-Step considers the distance between the DC pairs, and hence can avoid long routing paths and save network bandwidth. Consequently, since the bandwidth resources are used wisely, the DC backup window is reduced. Since

One-Step always uses the global maximum flow to set up the lightpaths, it only maximizes the backup throughput for the first DC pair. But if the distance between the DC pair is relatively long, this approach limits the throughput of the remaining DC pairs and prolongs the DC backup window.

Future Challenges

Despite the advantages of elastic optical inter-DC networks, there are also challenges for realizing highly efficient data migration and backup in them. For instance, the network control and management (NC&M) can be an issue. Almost all the algorithms discussed in this article are based on centralized NC&M, while the current NC&M in optical networks is usually distributed. Software-defined EONs (SD-EON) that adapts the concept of software-defined networking [15] can be a potential solution, as they separate the control and data planes of EONs, and apply centralized NC&M to manage network resource allocation.

Conclusion

This article discusses how to realize highly efficient data migration and backup for big data applications in elastic optical inter-DC networks. We first describe the impacts of big data applications on the underlying inter-DC networks and introduce the concept of flexible-grid elastic optical inter-DC networks that can properly address them. Then we model the data migration in these networks as a dynamic anycast problem and proposed several efficient algorithms. For data backup in these networks, we leverage a mutual backup model and discuss how to minimize the DC backup window.

Acknowledgments

This work was supported in part by the NCET Project NCET-11-0884, the NSFC Project 61371117, the Fundamental Research Funds for the Central Universities (WK2100060010), and the Strategic Priority Research Program of the Chinese Academy of Sciences (XDA06011202).

References

[1] http://en.wikipedia.org/wiki/Big_data
[2] W. Zhang *et al.*, "QoE-Driven Cache Management for HTTP Adaptive Bit Rate Streaming over Wireless Networks," *IEEE Trans. Multimedia*, vol. 15, Oct. 2013, pp. 1431–45.
[3] W. Zhang *et al.*, "Toward a Unified Elastic Computing Platform for Smartphones with Cloud Support," *IEEE Network*, vol. 27, Sept./Oct. 2013, pp. 34–40.
[4] D. Ziegler, "Distributed Peta-Scale Data Transfer," http://www.cs.huji.ac.il/simdhay/IND2011.html.
[5] N. Laoutaris *et al.*, "Inter-Datacenter Bulk Transfers with NetStitcher," *Proc. SIGCOMM 2011*, Aug. 2011, pp. 74–85.
[6] O. Gerstel *et al.*, "Elastic Optical Networking: A New Dawn for the Optical Layer?" *IEEE Commun. Mag.*, vol. 50, Apr. 2012, pp. S12–S20.
[7] A. Greenberg *et al.*, "The Cost of A Cloud: Research Problems in Data Center Networks," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 39, Jan. 2009, pp. 68–73.
[8] ITU-T Rec. G.694.1, "Spectral Grids for WDM Applications: DWDM Frequency Grid," Feb. 2012.
[9] Z. Zhu *et al.*, "Dynamic Service Provisioning in Elastic Optical Networks with Hybrid Single-/Multi-Path Routing," *J. Lightwave. Tech.*, vol. 31, Jan. 2013, pp. 15–22.
[10] J. Armstrong, "OFDM for Optical Communications," *J. Lightwave Tech.*, vol. 27, Mar. 2009, pp. 189–204.
[11] J. Cai *et al.*, "20 Tbit/s Transmission over 6860 km with Sub-Nyquist Channel Spacing," *J. Lightwave Tech.*, vol. 30, Feb. 2012, pp. 651–57.
[12] P. Colbourne and B. Collings, "ROADM Switching Technologies," *Proc. OFC 2011*, Mar. 2011, pp. 1–3.
[13] M. Zhang *et al.*, "Bandwidth Defragmentation in Dynamic Elastic Optical Networks with Minimum Traffic Disruptions," *Proc. ICC 2013*, June 2013, pp. 1–5.
[14] R. Proietti *et al.*, "Rapid and Complete Hitless Defragmentation Method Using a Coherent RX LO with Fast Wavelength Tracking in Elastic Optical Networks," *Opt. Express*, vol. 20, Nov. 2012, pp. 26,958–68.
[15] V. Gudla *et al.*, "Experimental Demonstration of OpenFlow Control of Packet and Circuit Switches," *Proc. OFC 2010*, Mar. 2010, pp. 1–3.

Biographies

PING LU received his B.S. degree from the Department of Optical Information Science and Technology, Xi'an University of Technology, China, in 2011. He is working toward his Ph.D. degree at the School of Information and Technology, University of Science and Technology of China (USTC). His research interests include cloud computing and datacenter networks.

LIANG ZHANG received his B.S. degree from the Department of Electronic Science Technology, Huazhong University of Science and Technology, China, in 2006. Now he is working toward his M.S. degree at the School of Information and Technology, USTC. His research interests include elastic optical networks and data center networks.

XIAHE LIU [S'13] received her B.S. degree from the School of Information Science and Technology, USTC, Hefei, China. She is working toward her M.S. degree at USTC. Her research interest is elastic optical networks.

JINGJING YAO received her B.S. degree from DaLian University of Technology, China, in 2013. She is working toward her M.S. degree at the School of Information and Technology, USTC. Her research interest is disaster-resilience data center networks.

ZUQING ZHU [M'07, SM'12] (zqzhu@ieee.org) received his Ph.D. degree from the Department of Electrical and Computer Engineering, University of California, Davis, in 2007. From July 2007 to January 2011, he worked in the Service Provider Technology Group of Cisco Systems, San Jose, California, as a senior R&D engineer. In January 2011, he joined USTC, where he is currently a full professor. He has published more than 140 papers in peer-reviewed journals and conferences. He has been on the Technical Program Committees (TPCs) of INFOCOM, ICC, GLOBECOM, and ICCCN, among others. He is also an Editorial Board member of *IEEE Communications Magazine*, *Elsevier Journal of Optical Switching and Networking*, *Springer Telecommunication Systems Journal*, *Wiley European Transactions on Emerging Telecommunications Technologies*, and others. He has received Best Paper Awards from IEEE ICC 2013, IEEE GLOBECOM 2013, IEEE ICNC 2014, and IEEE ICC 2015. He is a Senior Member of OSA.