Building a Network Highway for Big Data: Architecture and Challenges

Xiaomeng Yi, Fangming Liu, Jiangchuan Liu, and Hai Jin

Abstract

Big data, with their promise to discover valuable insights for better decision making, have recently attracted significant interest from both academia and industry. Voluminous data are generated from a variety of users and devices, and are to be stored and processed in powerful data centers. As such, there is a strong demand for building an unimpeded network infrastructure to gather geologically distributed and rapidly generated data, and move them to data centers for effective knowledge discovery. The express network should also be seamlessly extended to interconnect multiple data centers as well as interconnect the server nodes within a data center. In this article, we take a close look at the unique challenges in building such a network infrastructure for big data. Our study covers each and every segment in this network highway: the access networks that connect data sources, the Internet backbone that bridges them to remote data centers, as well as the dedicated network among data centers and within a data center. We also present two case studies of real-world big data applications that are empowered by networking, highlighting interesting and promising future research directions.

Our world is generating data at a speed faster than ever before. In 2010, 5 exabytes (1018 bytes, or 1 billion gigabytes) of data were created every two days, exceeding the total amount of information that was created by human beings from the dawn of civilization to 2003.[1] Till 2020, over 40 Zettabytes (10^{21} bytes) of data would be created, replicated, and consumed.[2] With the overwhelming amount of data pouring into our lives, from anywhere, anytime, and any device, we are undoubtedly entering the era of *big data*.

Big data brings big value. With advanced big data analyzing technologies, insights can be acquired to enable better decision making for critical development areas such as health care, economic productivity, energy, and natural disaster prediction, to name but a few. For example, by collecting and analyzing flu related keyword searches, Google has developed the *Flu Trends* service to detect regional flu outbreaks in near real time. Specifically, Google Flu Trends collected historical search frequency data of 50 million common keywords each week from 2003 to 2008. Then a linear model was used to compute the correlation coefficient between each keyword search history data and the actual influenza-like illness history data obtained from the Centers for Disease Control and Prevention (CDC) in the United States. After that, the keywords with the highest correlation coefficients were picked out, and their instant search frequencies were aggregated to predict future flu outbreaks in the United States. With big data in keyword searches, Google Flu Trends is able to detect flu outbreaks over a week earlier than CDC, which can significantly reduce the loss caused by the flu and even save lives. Another example comes from the United Parcel Service (UPS), who equips its vehicles with sensors to track their speed and location. With the sensed data, UPS has optimized its delivery routes and cut its fuel consumption by 8.4 million gallons in 2011.[3] It has been reported that big data analytics is among the top five catalysts that help increase U.S. productivity and raise the gross domestic product (GDP) in the coming years.[4]

In general, these big data are stored in data warehouses and processed in powerful data centers with massive interconnected server nodes. There have been significant studies on knowledge discovery and mining over the data and high-performance parallel computing tools for big data (e.g., MapReduce). But big data are not just born and reside there — they are transmitted from a variety of sources and are to be utilized at a variety of destinations for broad purposes. The life cycle of big data consists of multiple stages ranging from gen-

---

Xiaomeng Yi, Fangming Liu, and Hai Jin are with the Services Computing Technology and System Lab, Cluster and Grid Computing Lab in the School of Computer Science and Technology, Huazhong University of Science and Technology.

Jiangchuan Liu is with Simon Fraser University.

[1] *“Gold Rush: The Scramble to Claim and Protect Value in the Digital World,”* Deloitte Review, *2011*

[2] *“The Digital Universe in 2020: Big Data, Bigger Digital Shadows, and Biggest Growth in the Far East,”* IDC: Analyze the Future, *2012.*

[3] *“Big Data in Big Companies,” SAS Research Report, 2013.*

[4] *“Game Changers: Five Opportunities for US Growth and Renewal,” McKinsey Global Institute Report, 2013.*

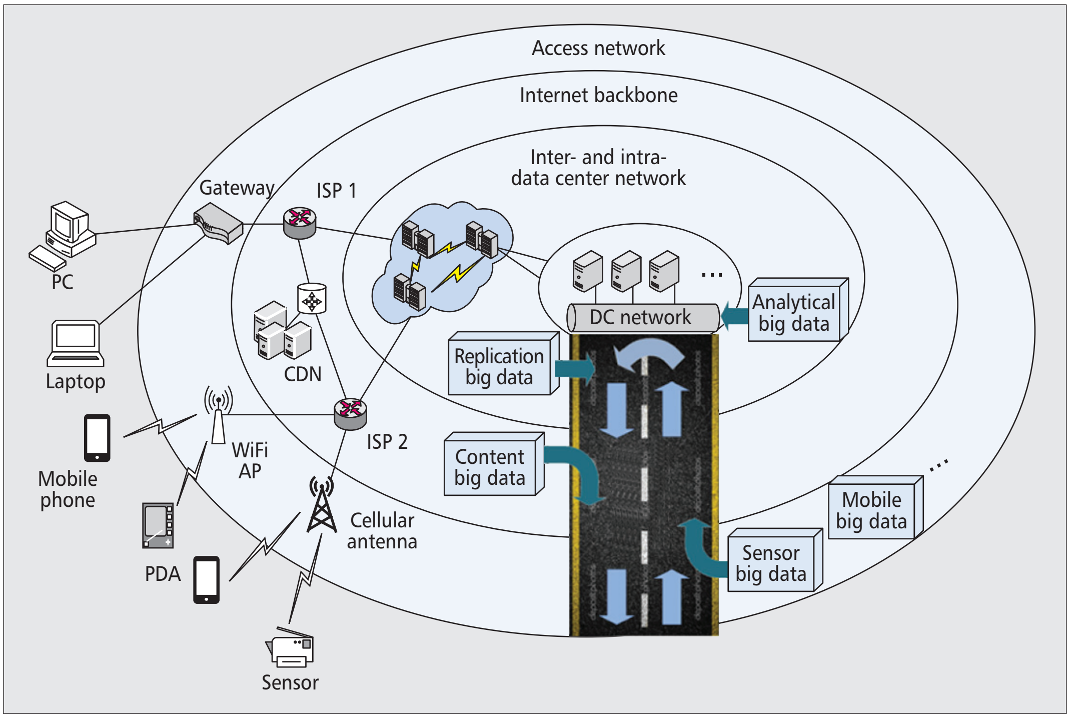

Figure 1. Three-layered network architecture from the perspective of big data applications.

eration, collection, aggregation, processing, and application delivery. While big data aggregation and processing occur mostly in data centers, the data are generated by and collected from geographically distributed devices, and the data knowledge/service are then distributed to interested users; the latter heavily depends on inter-data-center networks, access networks, and the Internet backbone, as depicted in Fig. 1. As such, networking plays a critical role as the digital highway that bridges data sources, data processing platforms, and data destinations. There is a strong demand to build an unimpeded network infrastructure to gather geologically distributed and rapidly generated data, and move them to data centers for effective knowledge discovery. The express network should also be extended to interconnect the server nodes within a data center and interconnect multiple data centers, thus collectively expanding their storage and computing capabilities.

Data explosion, which has been a continuous trend since the 1970s, is not news. The Internet has long grown together with the explosion, and has indeed greatly contributed to it. The three Vs (*volume*, *variety*, and *velocity*) from today's big data, however, are unprecedented. It remains largely unknown whether the Internet and related networks can keep up with the rapid growth of big data.

In this article, we examine the unique challenges when big data meet networks, and when networks meet big data. We check the state of the art for a series of critical questions: What do big data ask for from networks? Is today's network infrastructure ready to embrace the big data era? If not, where are the bottlenecks? And how could the bottlenecks be lifted to better serve big data applications? We take a close look at building an express network infrastructure for big data. Our study covers each and every segment in this network highway: the access networks that connect data sources, the Internet backbone that bridges them to remote data centers, and the dedicated network among data centers and within a data center. We also present two case studies of real-world big data applications that are empowered by networking, highlighting interesting and promising future research directions.

A Network Highway for Big Data: Why and Where?

With so much value hiding inside, big data have been regarded as the digital oil, and a number of government funded projects have been launched recently to build up big data analyzing systems, involving a number of critical areas ranging from healthcare to climate change and homeland security. A prominent example is the $200 million *National Big Data Research & Development Initiative* of the United States, announced by the White House in March 2012.[5] This initiative involves six departments and agencies in the United States, and aims to improve the tools and techniques needed to access, organize, and glean discoveries from big data. A list of representative projects worldwide is presented in Table 1.

In industry, there is also strong interest in exploiting big data to gain business profits. Advances in sensor networking, cyber-physical systems, and Internet of things have enabled financial service companies, retailers, and manufacturers to collect their own big data in their business processes. On the other hand, through such utility computing services as *cloud computing*, high-performance IT infrastructures and platforms that were previously unaffordable are now available to a broader market of medium and even small companies. As a result, more and more enterprises are seeking solutions to analyze the big data they generate. Gartner's survey on 720 companies in June 2013 shows that 64 percent of companies are investing or planning to invest in big data technologies.[6] A number of big data analytic platforms have already emerged in this competitive market. For example, Google launched

---

[5] *"Obama Administration Unveils 'Big Data' Initiative: Announces $200 Million in New R&D Investments," Office of Science and Technology Policy, The White House, 2012.*

[6] *"Survey Analysis: Big Data Adoption in 2013 Shows Substance Behind the Hype," Gartner, 2013.*

Project	Begin time	Department	Goal
1000 Genomes Project	1/2008	National Institutes of Health	To produce an extensive public catalog of human genetic variation, including SNPs and structural variants, and their haplotype contexts.
ARM Project	3/2012	Department of Energy	To collect and process climate data from all over the world to understand Earth's climate and come up with answers to climate change issues.
XDATA	3/2012	Defense Advanced Research Projects Agency (DARPA)	To develop new computational techniques and software programs that can analyze structured and unstructured big data sets faster and more efficiently.
BioSense 2.0	3/2012	Center for Disease Control and Prevention	To track public health problems and make data instantly accessible to end users across government departments.
The Open Science Grid	3/2012	National Science Foundation (NSF) and Department of Energy	To provide advanced fabric of services for data transfer and analysis to scientists worldwide for collaboration in science discovery.
Big Data for Earth System Science	3/2012	U.S. Geological Survey	To provide scientists with state-of-the-art computing capabilities and collaborative tools to make sense of huge data sets and better understand the earth.
Human Brain Project	2/2013	European Commission	To simulate the human brain and model everything that scientists know about the human mind using a supercomputer.
Unique Identification Authority	2/2009	The Indian Planning Commission	To create a biometric database of fingerprints, photographs, and iris scan images of all 1.2 billion people for efficient resident identification in welfare service delivery.

Table 1. Representative government-funded big data projects.

BigQuery, a big data service platform that enables customers to analyze their data by exploiting the elastic computing resources in Google's cloud; SAP also released its in-memory big data platform *HANA*, which is capable of processing large volumes of structured and unstructured data in real time.

In both government-funded research projects and business-oriented services, the life cycle of big data consists of multiple stages, as illustrated in Fig. 1. At first, the user data are generated from a variety of devices and locations, which are collected by wired and wireless networks. The data are then aggregated and delivered to data centers via the global Internet. In data centers, the big data are processed and analyzed. Finally, the results are delivered back to users or devices of interest and utilized.

Obviously, networks play a critical role in bridging the different stages, and there is a strong demand to create a fast and reliable interconnected network for the big data to flow freely on this digital highway. This network highway concerns not only just one segment of data delivery, but rather the whole series of segments for the life cycle of big data, from access networks to the Internet backbone, and to intra- and inter-data-center networks. For each layer of the network, the specific requirements big data transmission poses should be satisfied.

Access networks, which are directly connected to end devices such as personal computers, mobile devices, sensors, and radio frequency identification (RFID) devices, lie in the outer layer. On one hand, such raw data from fixed or mobile devices are transmitted into the network system; on the other hand, processed big data and analytics results are sent back to devices and their users. With the rapid development of wireless networks, more data are now collected from mobile devices that have limited battery capacity. As such, energy-efficient networking is expected to make batteries in mobile devices more durable. Wireless links also suffer from interference, which results in unstable bandwidth provisioning. Big data applications such as cinematic-quality video streaming require sustained performance over a long duration to guarantee the quality of user experience, which has become a critical challenge for wireless networking.

The Internet backbone serves as the intermediate layer that connects access networks and data center networks. Popular big data applications like photo sharing and video sharing allow users to upload multimedia contents to data centers and share them with their friends in real time. To enable good user experience, the Internet backbone needs to forward massive geologically distributed data to data centers with high throughput, and deliver processed data to users from data centers with low latency. As such, high-performance end-to-end links are required for uploading, and efficient content distribution networks (CDNs) are required for downloading.

Within a data center, big data are processed and analyzed with distributed computing tools such as MapReduce and Dryad, which involve intensive data shuffle among servers. A scalable, ultra-fast, and blocking-free network is thus needed to interconnect the server nodes. Multiple geologically distributed data centers can be exploited for load balancing and low-latency service provisioning, which calls for fast networking for data exchange, replication, and synchronization among the data centers. Moreover, inter-data-center links are leased from Internet service providers (ISPs) or dedicatedly deployed by cloud providers with nontrivial costs. Effective data transmission and traffic engineering schemes (e.g., using software

---

[7] *"The Zettabyte Era—Trends and Analysis," Cisco Visual Networking Index white paper, 2013.*

defined networking, SDN) that fully utilize the capacity of such links are therefore expected. For instance, Google and Microsoft both attempt to deploy SDN-based global data center WANs to achieve fault tolerance, utilization improvement, and policy control that can hardly be realized by traditional WAN architectures.

In summary, to attain full speed for big data transmission and processing, every segment of the network highway should be optimized and seamlessly concatenated. We proceed to investigate each part of the network system by identifying the unique challenges that big data applications pose and analyzing the state-of-the-art works that endeavor to build the network highway for big data.

Access Network: Linking Sources

With the fast progress of digitalization and the development of sensor networks, huge amounts of data are collected by all kinds of end devices like PCs, smartphones, sensors, and GPS devices. Meanwhile, applications like online social networks and video streaming push rich content big data into user devices. The access network plays a critical role here in gathering such distributed data and forwarding them through the Internet to data centers. While the *last mile* problem has been addressed well with today's high-speed home and office network connections, the wireless connection remains a severe bottleneck. Cisco predicts that traffic from wireless and mobile devices will exceed traffic from wired devices by 2016, which makes wireless network performance optimization of paramount importance.

Wireless broadband technologies have evolved significantly in recent years; but when facing data-intensive applications, they are still insufficient to satisfy the bandwidth requirements. Yoon *et al.* exploit the wireless broadcast advantage and use it to bridge the gap between wireless networking capacity and the bandwidth demands of video applications. MuVi [1], a multicast delivery scheme, is proposed to optimize video distribution. By prioritizing video frames according to their importance in video reconstruction and exploiting a resource allocation mechanism that maximizes the system utility, MuVi improves the overall video quality across all users in a multicast group.

Multiple-input multiple-output orthogonal frequency-division multiplexig (MIMO-OFDM) technologies, which can significantly increase wireless capacity, have become the default building blocks in the next generation of wireless networks. Liu *et al.* [2] observe that in wireless video streaming applications, both the video coding scheme and the MIMO-OFDM channel present non-uniform energy distribution among their corresponding components. Such non-uniform energy distribution in both the source and the channel can be exploited for fine-grained unequal error protection for video delivery in error-prone wireless networks. To this end, ParCast, an optimized scheme for video delivery in MIMO-OFDM channels, is proposed. It separates the video coding and wireless channel into independent components and allocates more important video components with higher-gain channel components. This leads to significantly improved quality for video over wireless.

In addition to the aforementioned works, there are several other promising approaches to improve the quality of wireless video streaming services. On the user's side, an application-aware MIMO video rate adaptation mechanism can be deployed. It detects changes in a MIMO channel and adaptively selects an appropriate transmission profile, thereby improving the quality of the delivered video. On the server's side, standard base station schedulers work on a fine-grained per-packet basis to decrease the delivery delay of single packets. However, it is insufficient to guarantee video watching experience at a coarse granularity, such as a fixed video bit rate over several seconds, which would typically consist of video content in thousands of packets. In response, a video management system, which schedules wireless video delivery at a granularity of seconds with knowledge of long-term channel states, has the potential to further improve user experience. Moreover, distortion and delay, which are two important user experience metrics, conflict with each other in wireless networks. The optimal trade-off between distortion and delay in wireless video delivery systems largely depends on the specific features of video flows. As such, a policy that smartly balances distortion and delay according to the features of video flows can also improve user experience.

Internet Backbone: From Local to Remote

Beyond the access network, the user or device generated data will be forwarded through the Internet backbone to data centers. For example, in mobile cloud computing services [3], where powerful cloud resources are exploited to enhance the performance of resource-constrained mobile devices, data from geologically distributed mobile devices are transmitted to the cloud for processing. Given that the local data come worldwide, the aggregated data toward a data center can be enormous, which creates significant challenges to the Internet backbone. Table 2 summarizes the Internet backbone solutions for big data.

End-to-End Transmission

With the growing capacity of access links, network bottlenecks are observed to be shifting from the network edges in access networks to the core links in the Internet backbone. To improve the throughput of end-to-end data transmission, path diversity should be explored, which utilizes multiple paths concurrently to avoid individual bottlenecks. A representative is mPath [4], which uses a large set of geographically distributed proxies to construct detour paths between end hosts. An additive increase and multiplicative decrease (AIMD) algorithm similar to TCP is used to deal with congested proxy paths to adaptively regulate the traffic over them, or even completely avoid them.

Besides uploading, the big data, after being processed by the data centers, also need to be downloaded by users to appreciate the inside value. Downloading, however, poses different demands. For applications like online social networks, it is critical to deliver user required contents with low latency while providing a consistent service to all users. Wittie *et al.* [5] reverse engineered Facebook, investigating the root causes of its poor performance when serving users outside of the United States. They suggest that this can be improved by exploring the locality of interest, which, with proxy and caching, can dramatically reduce the backbone traffic of such online social networks as well as its access delay.

Content Delivery Network

For geologically distributed data consumers, CDNs can be explored to serve them with higher throughput. High throughput is typically achieved in two ways: optimizing path selection to avoid network bottlenecks, and increasing the number of peering points. Yu *et al.* [6] introduce a simple model to illustrate and quantify the benefit of them. Using both synthetic and Internet network topologies, they show that increasing the number of peering points improves the

Internet backbone						
Approaches	Network infrastructure	Big data application	Goal	Technique	Evaluation method	Overhead
mPath [4]	Internet backbone	End-to-end transmission	Avoid bottleneck	AIMD algorithm	Implementation on PlanetLab	Proxy node deployment
Wittie *et al.* [5]	Internet backbone	Social networks	Reduce service latency	TCP proxy, caching	Trace-driven simulation	Cache and proxy deployment
Liu *et al.* [7]	CDN	Video streaming	Improve QoS	Bit rate adaption	Trace-driven simulation	Low scalability
Jiang *et al.* [8]	CDN	Video streaming	Reduce operational cost	CDN extension	Synthetic and trace-driven simulation	Tracker deployment
Intra- and inter-data-center networks						
Approaches	Network infrastructure	Big data application	Goal	Technique	Evaluation method	Overhead
Hedera [9]	Data center networks	Data processing	Optimize network utilization	Flow scheduling	Simulation, implementation on Portland testbed	Centralized control, low scalability
FlowComb [10]	Data center networks	MapReduce/ Dryad	Optimize network utilization	Flow scheduling	Implementation on Hadoop testbed	Monitor and transfer demand information
Orchestra [11]	Data center networks	MapReduce/ Dryad	Reduce job duration	Transfer scheduling	Implementation on Amazon EC2 and DETERlab	Modify the distributed framework
RoPE [12]	Data center networks	MapReduce/ Dryad	Reduce job duration	Execution plan optimization	Implementation on Bing's production cluster	Pre-run jobs to acquire job property
Camdoop [13]	Data center network topology	MapReduce/ Dryad	Decrease network traffic	Data aggregation	Implementation on CamCube	Special network topology deployment
Mordia [14]	OCS data center network	Data processing	Reduce switching delay	OCS, traffic matrix scheduling	Prototype implementation	Hardware/topology deployment
3D Beamforming [15]	Wireless data center network	Data processing	Flexible bandwidth provisioning	60 GHz wireless links	Local testbed, simulation	Physical antenna/ reflector deployment
NetStitcher [16]	Inter-data-center links	Data backup and migration	Improve network utilization	Store-and-forward algorithm	Emulation, live deployment	Periodical schedule recomputation
Jetway [17]	Inter-data-center links	Video delivery	Minimize link cost	Flow assignment algorithm	Simulation, implementation on Amazon EC2	Centralized controller, video flow tracking

Table 2. A taxonomy of Internet backbone, intra- and inter-data-center solutions for big data.

throughput the most, while optimal path selection has only limited contribution. Liu *et al.* [7] further find that video delivery optimized for low latency or high average throughput may not work well for high-quality video delivery that requires sustained performance over a long duration. This leads to an adaptive design with global knowledge of network and distribution of clients.

To further reduce the operational cost of big data traffic over CDN, Jiang *et al.* [8] suggest that the CDN infrastructures can be extended to the edges of networks, leveraging such devices such as set-top boxes or broadband gateways. Their resources can be utilized through peer-to-peer communications with smart content placement and routing to mitigate the cross-traffic among ISPs.

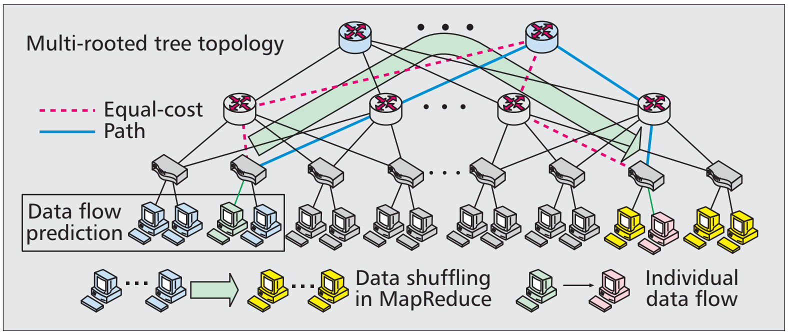

Figure 2. Data flow scheduling and data shuffling techniques in intra-data-center networks.

Data Center Networks: Where Big Data Are Stored and Processed

Big data collected from end devices are stored and processed in data centers. Big data applications such as data analysis and deep learning usually exploit distributed frameworks like MapReduce and Dryad to achieve inflexibility and scalability. Data processing in those distributed frameworks consists of multiple computational stages (e.g., map and reduce in the MapReduce framework). Between the stages, massive amounts of data need to be shuffled and transferred among servers. The servers usually communicate in an all-to-all manner, which requires high bisection bandwidth in data center networks. As such, data center networks often became a bottleneck for those applications, with data transfers accounting for more than 33 percent of the running time in typical workloads. In Table 2, we summarize the state-of-the-art solutions toward engineering better intra- and inter-data-center networks.

Dynamic Flow Scheduling

As shown in Fig. 2, there are multiple equal-cost paths between any pair of servers in a typical multi-rooted tree topology of a data center network. To better utilize the paths, Hedera [9] is designed to dynamically forward flows along these paths. It collects flow information from switches, computes non-conflicting paths for flows, and instructs switches to reroute traffic accordingly. With a global view of routing and traffic demands, Hedera is able to maximize the overall network utilization with only small impact on data flows. Through monitoring the flow information in switches, it schedules data flows only after they cause congestion at certain location. To actually avoid congestion, however, the data flows should be detected even before they occur. This is addressed in FlowComb [10], which predicts data flows effectively by monitoring MapReduce applications on servers through software agents. A centralized decision engine is designed to collect data from the agents and record the network information.

Transfer Optimization in MapReduce-Like Frameworks

In addition to scheduling individual flows, optimizing the entire data transfer has the potential to further reduce job completion time. To this end, Orchestra [11] optimizes common communication patterns like shuffle and broadcast by coordinating data flows. When there are concurrent transfers, Orchestra enforces simple yet effective transfer scheduling policies such as first-in first-out (FIFO) to reduce the average transfer time. This is further enhanced in RoPE [12] by optimizing job execution plans. An execution plan specifies the execution order of operations in a job, as well as the degree of parallelism in each operation. RoPE employs a composable statistics collection mechanism to acquire code and data properties in a distributed system, and then automatically generates optimized executions plans for jobs, which reduces the volume of data to be transferred.

Novel Topology and Hardware

Novel network topologies have also been proposed to improve the network performance in big data processing. Costa *et al.* [13] observe that in common workloads, data volumes are greatly reduced after the progress of data processing. With this insight, a data processing framework, Camdoop, is proposed. It is built on a novel network topology where the servers are directly connected, through which partial data can be aggregated on servers along the path to minimize data transmission.

Recently, optical circuit switching (OCS) has been suggested to accommodate the fast growing bandwidth demands in data center networks. The long circuit reconfiguration delay of OCSs, however, hinders their deployment in modern data centers. To address this issue, Porter *et al.* [14] propose novel traffic matrix scheduling (TMS), which leverages application information and short-term demand estimates to compute short-term circuit schedules, and proactively communicates circuit assignments to communicating entities. As such, it can support flow control in microseconds, pushing down the reconfiguration time by two to three orders of magnitude. On the other hand, wireless links in the 60 GHz band are attractive to relieve hotspots in oversubscribed data center networks. The 60 GHz wireless links require direct line of sight between sender and receiver, which limits the effective range of wireless links. Moreover, they can suffer from interference when nearby wireless links are working. To deal with these, Zhou *et al.* [15] propose a new wireless primitive, 3D beamforming, for data centers. By bouncing wireless signals off the ceiling, 3D beamforming avoids blocking obstacles, thus extending the range of each wireless link. Moreover, the signal interference range is significantly reduced, allowing nearby links to work concurrently.

Inter-Data-Center Links

Big-data-based services such as social networks usually exploit several geographically distributed data centers for replication and low-latency service provision. Those data centers are interconnected by high-capacity links leased from ISPs. Operations like data replication and synchronization require high bandwidth transformation between data centers. It is thus critical to improve the utilization or reduce the cost for such inter-data-center links.

Laoutaris *et al.* [16] observe a diurnal pattern of user demand on inter-data-center bandwidth, which results in low bandwidth utilization in off-peak hours. They propose NetStitcher to utilize the leftover bandwidth in off-peak hours for such non-real-time applications as backups and data migrations. As illustrated in Fig. 3, NetStitcher exploits a store-and-forward algorithm to transfer big data among data centers. The data are split into pieces and transferred to their destination along multiple paths, each of which consists of a series of intermediate data centers. A scheduling module decides when and where the data pieces should travel according to the available bandwidth.

For real-time big data applications like video streaming, Feng *et al.* [17] propose Jetway, which, based on a widely used *percentile charging* model, uses the qth largest traffic volume of all time intervals during the charging period as the charging volume. If the current time interval's traffic volume exceeds that of the qth percentile of previous time intervals, it will incur additional cost. Otherwise, the already paid for bandwidth will not be fully used. Jointly considering the link capac-

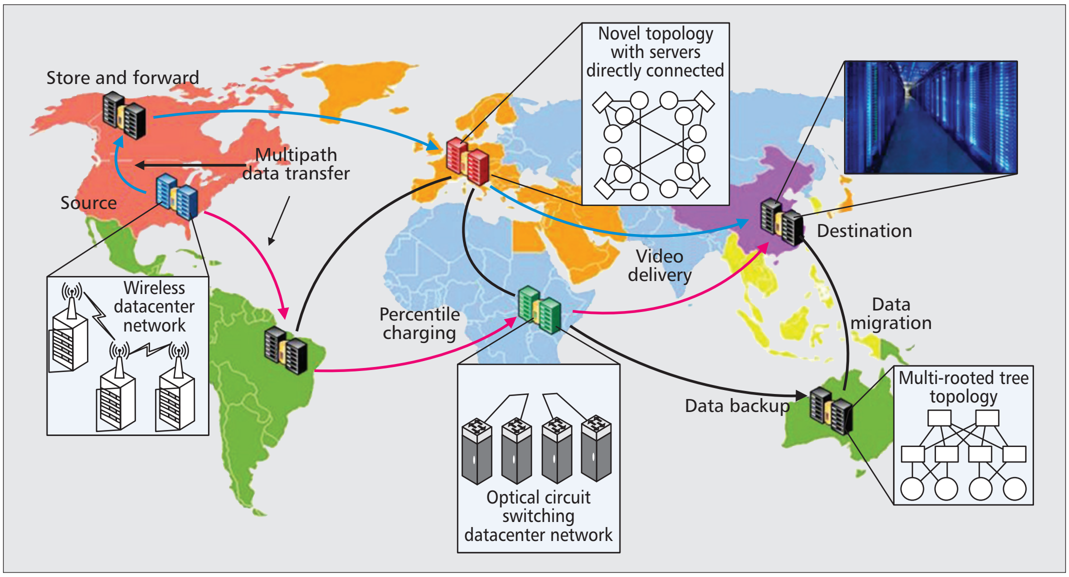

Figure 3. Emerging network architectures of different data centers, and multipath store-and-forward data transfer among geo-distributed data centers.

ity, bandwidth availability, delay-tolerant degree of flows, and previous traffic volume of links, Jetway can minimize the cost of inter-data-center links.

In Summary

The aforementioned research efforts propose solutions to serve big data applications at *different levels* ranging from the *macroscopic* multi-data-center level to the *microscopic* individual flows. Specifically, from the macroscopic level, inter-data-center links enable multiple geologically distributed data centers to provide such consistent reliable big data services as video streaming and social networking with low latency. At the single-data-center level, novel network topologies and hardware devices provide cost-effective solutions toward a high-performance network that interconnects all servers. At the microscopic level, coordinated all-to-all communications among a subset of servers in a data center mitigate network congestion and reduce the completion times of data processing jobs. In particular, for individual data flows between a specific pair of servers, dynamic flow scheduling techniques choose proper paths for data flows to avoid network resource competition. Nevertheless, there remain significant open questions to be addressed in the literature, such as how to schedule data flows in order to meet their specific completion deadlines, and how to provide guaranteed network performance to multiple concurrent data processing jobs in a data center. To achieve a comprehensive understanding of these problems, Xu *et al.* [18] have investigated the state-of-the-art research in the literature on providing guaranteed network performance for tenants in Internet as a service (IaaS) clouds.

Big-Data-Based Networking Applications

We generally classify big data applications into two categories, *Internet applications* and *mobile wireless network applications*, with regard to the networking infrastructure on which they work. For each category, we discuss the benefit and opportunity that big data brings, by analyzing representative applications depicted in Fig. 4.

Internet Applications

One of the prominent big data applications closely related to our daily lives is Netflix, which offers streaming video-on-demand services and now takes up a third of U.S. download Internet traffic during peak traffic hours. To support the combination of huge traffic and unpredictable demand bursts, Netflix has developed a global video distribution system using Amazon's cloud. Specifically, depending on customer demand, Netflix's front-end services are running on 500 to 1000 Linux-based Tomcat JavaServer and NGINX web servers. These are empowered by hundreds of other Amazon S3 and NoSQL Cassandra database servers using the Memcached high-performance distributed memory object caching system.

Netflix purchases master copies of digital films from movie studios and, using the powerful Amazon EC2 cloud machines, converts them to over 50 different versions with different video resolutions and audio quality, targeting a diverse array of client video players running on desktop computers, smartphones, and even DVD players or game consoles connected to television. The master copies and the many converted copies are stored in Amazon S3. In total, Netflix has over 1 petabyte of data stored on Amazon.

Thanks to insights from such big video data and the associated client behavior data as the programs they are watching, their demographics and preferences, Netflix made a big decision of bidding over $100 million for two seasons of the U.S. version of "House of Cards" in 2011, which turned out to be a great success. The "House of Cards" series, directed by David Fincher, starring Kevin Spacey, is based on a popular British series. Big data in Netflix show that its British version has been well watched, and the same subscribers who loved the original BBC production also like movies starring Kevin Spacey or directed by David Fincher. The analytical results of big data indicate that "House of Cards" would bring significant business value to Netflix. In fact, the series has brought Netflix over two million new subscribers in the United States and one million outside the United States. Netflix's purchase of "House of Cards" is largely regarded as a great success of big data.

Mobile Wireless Network Applications

Powered by advanced technologies in mobile networking, sensor networking, and the Internet of Things, mobile wireless big data applications are emerging in our lives. As a typical example of them, Nike+ provides improved service to Nike users with wireless connected sensing devices and smartphones. In particular, Nike enhances its products, such as shoes and wristbands, with built-in sensors that continuously

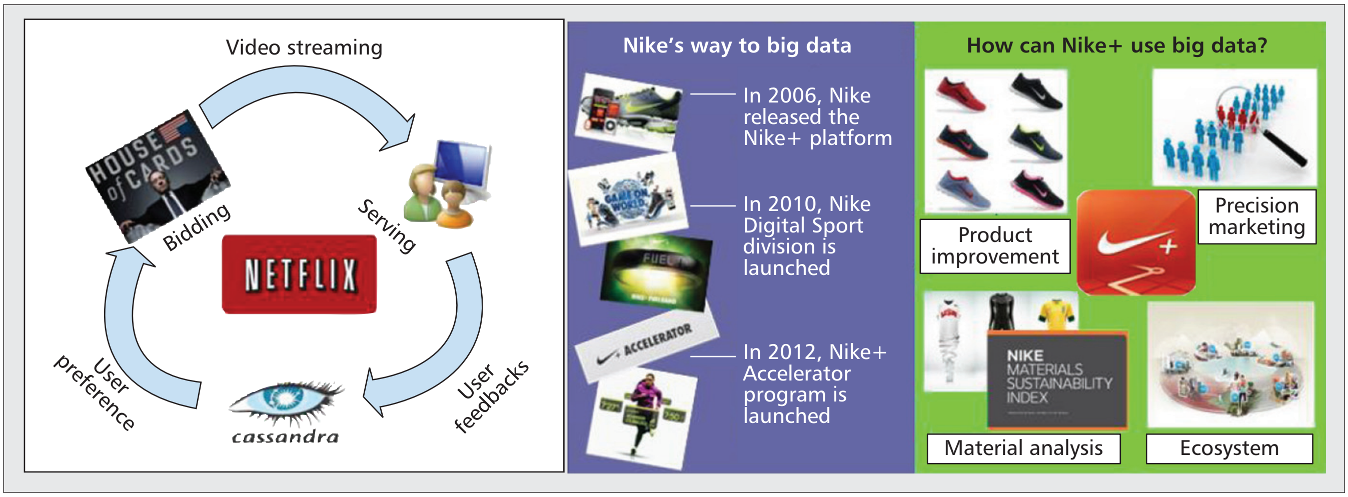

Figure 4. Representative big data applications in the Internet and wireless networks: left: Netflix's "House of Cards"; right: Nike+.

track the users' movement during their workouts. The users can install Nike+ apps in their smartphones, which collect data from the sensors through wireless connections. The collected data provide users with their instant exercise information such as their pace, GPS position, distance moved, and calories burned. Moreover, with mobile networks interconnecting users' smartphones, the Nike+ apps enable social interactions among users. They can share their progress and exercising experience to cheer each other on or create groups to go after a goal together. By now, Nike+ has become a big data platform that collects, stores, and processes data generated from more than 18 million users.

A series of issues in wireless networking, however, should be addressed when collecting and transferring user data in Nike+. To provide users with real-time feedback, both sensed data in wireless sensor networks and social interaction data in mobile cellular networks should be synchronized to users' smartphones frequently. Unfortunately, high-frequency data transmission in wireless networks would result in high energy consumption. Both wireless sensors and smartphones have limited battery capacity. Transferring data at too high a frequency would cause degradation in battery life. It is nontrivial for wireless networks to provide an energy-efficient solution for frequent data transmissions. In its latest product, Nike+ Fuelband, Nike+ adopts Bluetooth 4.0 wireless techniques for data transmission. By exploiting the novel Bluetooth LE (for low energy) protocol to synchronize sensed data to smartphones, Nike+ Fuelband has been made more durable.

Conclusion and Future Trends

So far we have reviewed the networking architecture and services for big data applications. We identify the major challenges big data applications bring to networking systems and discuss the state-of-the-art research efforts to meet the demand of big data over networking. With the rapid growth of big data applications, building the network system as a highway for big data transmission will continue to be a hot topic in both academia and industry. We now identify some notable future trends.

Standardization: Bridging the Fragmented World

The world of big data remains largely fragmented. Today's big data are generally stored and analyzed within a particular business entity or organization to obtain insights and guide decision making. There is a great demand to exchange these big data for better insight discovery. For example, manufacturing factories may require feedback data from retailers to discover user demands and help design new products. The retailers, in turn, would need product-related data to set proper prices for products and recommend them to target consumers. Standards are thus necessary to bridge the fragmented world. There has been active development on big data storage and exchange. For example, the National Institute of Standards and Technology (NIST) set up a big data working group on June 19, 2013, aimed at defining the requirements for interoperability, reusability, and extendability of big data analytic techniques and infrastructures.

From the perspective of networking, the standards are required to specify how to transfer big data between different platforms. The transferred big data may contain semi-structured or unstructured data, such as pictures, audios, videos, click streams, log files, and the output of sensors that measure geographic or environmental information. The standards should specify how these data should be encoded and transferred in network systems in order to facilitate network quality of service (QoS) management with low latency and high fidelity. Moreover, as novel technologies such as software defined networking (SDN) keep emerging in networking systems, corresponding standards are required to specify how these technologies can be exploited for efficient big data transmission. For example, the Open Network Foundation (ONF) has been releasing and managing the OpenFlow standard since 2011, which defines the communications interface between the control and forwarding layers in an SDN architecture.

Privacy and Security

With a variety of personal data such as buying preferences, healthcare records, and location-based information being collected by big data applications and transferred over networks, the public's concerns about data privacy and security naturally arise. While there have been significant studies on protecting data centers from being attacked, the privacy and security loopholes when moving crowsourced data to data centers remain to be addressed. There is an urgent demand on technologies that endeavor to enforce privacy and security in data transmission. Given the huge data volume and number of sources, this requires a new generation of encryption solutions (e.g., homomorphic encryption).

On the other hand, big data techniques can also be used to address the security challenges in networked systems. Network attacks and intrusions usually generate data traffic of specific patterns in networks. By analyzing the big data gathered by a network monitoring system, those misbehaviors can be identified proactively, thus greatly reducing the potential loss.

Empowered by real-time big data analyzing technologies, we expect that a series of advanced tools for identifying deep security loopholes in large-scale and complex system could be developed in the near future.

Acknowledgment

The corresponding author is Fangming Liu (fmliu@hust.edu.cn). The research was supported by a grant from the National Basic Research Program (973 program) under grant No. 2014CB347800.

References

[1] J. Yoon *et al.*, "MuVi: A Multicast Video Delivery Scheme for 4G Cellular Networks," *Proc. ACM Mobicom*, 2012.

[2] X. L. Liu *et al.*, "ParCast: Soft Video Delivery in MIMO-OFDM WLANs," *Proc. ACM Mobicom*, 2012.

[3] F. Liu *et al.*, "Gearing Resource-Poor Mobile Devices with Powerful Clouds: Architecture, Challenges and Applications," *IEEE Wireless Commun.*, Special Issue on Mobile Cloud Computing, vol. 20, no. 3, 2013.

[4] Y. Xu *et al.*, "mPath: High-Bandwidth Data Transfers with Massively-Multipath Source Routing," *IEEE Trans. Parallel and Distributed Systems*, vol. 24, issue 10, 2013.

[5] M. P. Wittie *et al.*, "Exploiting Locality of Interest in Online Social Networks," *Proc. ACM Eurosys*, 2010.

[6] M. Yu *et al.*, "Tradeoffs in CDN Designs for Throughput Oriented Traffic." *Proc. ACM Conext*, 2012.

[7] X. Liu *et al.*, "A Case for a Coordinated Internet Video Control Plane," *Proc. ACM Sigcomm*, 2012.

[8] W. Jiang *et al.*, "Orchestrating Massively Distributed CDNs," *Proc. ACM Conext*, 2012.

[9] M. Al-Fares *et al.*, "Hedera: Dynamic Flow Scheduling for Data Center Networks," *Proc. USENIX NSDI*, 2010.

[10] A. Das *et al.*, "Transparent and Flexible Network Management for Big Data Processing in the Cloud," *Proc. ACM Hotcloud*, 2013.

[11] M. Chowdhury *et al.*, "Managing Data Transfers in Computer Clusters with Orchestra," *Proc. ACM Sigcomm*, 2011.

[12] S. Agarwal *et al.*, "Re-optimizing Data-Parallel Computing," *Proc. USENIX NSDI*, 2012.

[13] P. Costa *et al.*, "Camdoop: Exploiting In-Network Aggregation for Big Data Applications," *Proc. USENIX NSDI*, 2012.

[14] G. Porter *et al.*, "Integrating Microsecond Circuit Switching into the Data Center," *Proc. ACM Sigcomm*, 2013.

[15] X. Zhou *et al.*, "Mirror Mirror on the Ceiling: Flexible Wireless Links for Data Centers," *Proc. ACM Sigcomm*, 2012.

[16] N. Laoutaris *et al.*, "Inter-Datacenter Bulk Transfers with NetStitcher," *Proc. ACM Sigcomm*, 2011.

[17] Y. Feng, B. Li, and B. Li, "Jetway: Minimizing Costs on Inter-Datacenter Video Traffic," *Proc. ACM MM*, 2012.

[18] F. Xu *et al.*, "Managing Performance Overhead of Virtual Machines in Cloud Computing: A Survey, State of Art and Future Directions," *Proc. IEEE*, vol. 102, no. 1, 2014.

Biographies

XIAOMENG YI (xiaomengyi@hust.edu.cn) is currently a Ph.D. student in the School of Computer Science and Technology, Huazhong University of Science and Technology, Wuhan, China. His current research interests focus on cloud computing, modeling and optimization.

FANGMING LIU [M] (fmliu@hust.edu.cn) is an associate professor in the School of Computer Science and Technology, Huazhong University of Science and Technology, and he is named the CHUTIAN Scholar of Hubei Province, China. He is the Youth Scientist of National 973 Basic Research Program Project on Software-Defined Networking (SDN)-Based Cloud Data Center Networks, which is one of the largest SDN projects in China. Since 2012, he has also been invited as a StarTrack Visiting Young Faculty at Microsoft Research Asia (MSRA), Beijing. He received his B.Eng. degree in 2005 from the Department of Computer Science and Technology, Tsinghua University, Beijing, and his Ph.D. degree in computer science and engineering from the Hong Kong University of Science and Technology in 2011. From 2009 to 2010, he was a visiting scholar at the Department of Electrical and Computer Engineering, University of Toronto, Canada. He was the recipient of Best Paper Awards from IEEE GLOBECOM 2011, IEEE IUCC 2012, and IEEE CloudCom 2012, respectively. His research interests include cloud computing and data center networking, mobile cloud, green computing and communications, software-defined networking and virtualization technology, large-scale Internet content distribution, and video streaming systems. He is a member of ACM, as well as a member of the China Computer Federation (CCF) Internet Technical Committee. He has been a Guest Editor of *IEEE Network* and *IEEE Systems Journal*, an Associate Editor of *Frontiers of Computer Science*, and served on the Technical Program Committees for IEEE INFOCOM 2013–2015, ICNP 2014, ACM Multimedia 2014, and IEEE GLOBECOM 2012–2014. He served as the IEEE LANMAN 2014 Publicity Chair and GPC 2014 Program Chair.

JIANGCHUAN LIU [SM] (jcliu@cs.sfu.ca) received his B.Eng. degree (cum laude) from Tsinghua University, Beijing, China, in 1999, and his Ph.D. degree from the Hong Kong University of Science and Technology in 2003, both in computer science. He is currently an associate professor in the School of Computing Science, Simon Fraser University, British Columbia, Canada, and was an assistant professor in the Department of Computer Science and Engineering at the Chinese University of Hong Kong from 2003 to 2004. He is a recipient of a Microsoft Research Fellowship (2000), Hong Kong Young Scientist Award (2003), and Canada NSERC DAS Award (2009). He is a co-recipient of the ACM Multimedia 2012 Best Paper Award, the IEEE GLOBECOM 2011 Best Paper Award, the IEEE Communications Society Best Paper Award on Multimedia Communications in 2009, as well as IEEE IWQoS '08 and IEEE/ACM IWQoS '12 Best Student Paper Awards. His research interests are in networking and multimedia, in particular, multimedia communications, peer-to-peer networking, cloud computing, social networking, and wireless sensor/mesh networking. He has served on the Editorial Boards of *IEEE Transactions on Multimedia*, *IEEE Communications Tutorial and Surveys*, *IEEE Internet of Things Journal*, *Elsevier Computer Communications*, and *Wiley Wireless Communications and Mobile Computing*. He is a TPC Co-Chair for IEEE/ACM IWQoS '14 in Hong Kong.

HAI JIN [SM] (hjin@hust.edu.cn) is a Cheung Kung Scholars Chair Professor of computer science at the Huazhong University of Science and Technology, China. He is now dean of the School of Computer Science and Technology at the university. He received his Ph.D. degree in computer engineering from the Huazhong University of Science and Technology in 1994. In 1996, he was awarded a German Academic Exchange Service fellowship to visit the Technical University of Chemnitz in Germany. He was awarded the Excellent Youth Award from the National Science Foundation of China in 2001. He is the chief scientist of ChinaGrid, the largest grid computing project in China, and chief scientist of the National 973 Basic Research Program Project of Virtualization Technology of Computing Systems. His research interests include computer architecture, virtualization technology, cloud computing and grid computing, peer-to-peer computing, network storage, and network security.