CLOUD COVER

Challenges and Opportunities for Data-Intensive Computing in the Cloud

Eun-Sung Jung and Rajkumar Kettimuthu,
Argonne National Laboratory

Now running mostly on high-performance computers, data-intensive applications pose several important challenges as they move toward cloud deployment.

The era of big data is here. In science, business, industry, and IT management, the amount and rate of data generation is exploding. At Argonne National Laboratory's experimental Advanced Photon Source (APS) facility, for example, a tomography beamline, used for rendering objects at submicroscopic levels of detail, can produce 150 terabytes of data per day. Similarly, the University of Chicago's Data-Driven Urban Design and Analysis project (https://urbanccd.org/projects/urban-sciences-research-coordination-network-data-driven-urban-design-and-analysis) is constantly gathering massive amounts of data from street-level sensors all over the city to use for predictive analytics aimed at future urban design.

So far, to ensure timely computation, such data-intensive applications have had to run on dedicated high-performance computers (HPCs) available to very few users—and even then can scale up only as far as a particular facility's capacity allows. Cloud computing, on the other hand, can provide any user access to large computing resources for data-intensive applications—on-demand, with great flexibility, and at relatively low cost.

Successfully deploying data-intensive applications in clouds, however, requires seamlessly incorporating sophisticated big-data transfer and data-management frameworks into cloud computing infrastructures. Doing so poses some significant—though not insurmountable—challenges.

CLOUD LIMITATIONS

Data transfers into and out of the cloud, as well as among different clouds, typically must be carried out through system utilities such as secure copy (scp) and http, which are notoriously slow. In some cases, shipping disks using courier services like FedEx can actually be faster. Moreover, existing data-movement services' throughput—in the tens of megabytes per second—is an order of magnitude slower than the hundreds of megabytes per second that data-intensive applications require.

In addition, data-management platforms within datacenters

Published by the IEEE Computer Society

necessarily play key roles in cloud data transfer. However, these typically rely on popular cloud data-management software, such as Apache HBase, Google BigTable, and the open source product Memcached, which aren't directly applicable to data-intensive computing.

Current cloud computing paradigms simply do not satisfactorily meet data-intensive application requirements, due to limitations in both the underlying hardware infrastructure—specifically, network bandwidth shortfalls and high data-access latency—and unoptimized data-movement frameworks. Challenges in migrating data-intensive applications to the cloud from the perspective of federating diverse data-management and hardware systems are discussed elsewhere.[1] Here, we focus specifically on the data-movement framework.

What's needed are high-throughput and scalable data-management techniques, as well as higher-performance network and storage facilities coupled with stricter quality-of-service (QoS) requirements in the cloud.

CHALLENGES FOR DATA-INTENSIVE APPLICATIONS

Deploying data-intensive applications in the cloud faces several key challenges.

Distributed data sources

One key requirement for data-intensive computing in the cloud is the ability to efficiently move big data to clouds from increasingly varied sources: distributed sensors, genome-sequencing centers, and experimental facilities such as the APS, as well as observational sites such as Chile's Very Large Telescope. This requirement creates problems for several reasons:

- Simply put, more data must transfer to the cloud due to enhanced sensing capabilities and increasing global data sources.
- More distributed data leads to longer latency and inevitable bottlenecks in network links to a cloud, making single-cloud processing more difficult.
- Status changes in dynamic networks and the interplay of data movement and computation within clouds creates difficulties for sustaining data transfer rates.

Multicloud environments

In some cases, moving large datasets among multiple clouds might be necessary—for example, if a single cloud isn't sufficient for computation purposes or if a user or collaborators must exploit multiple cloud resources. While intercloud movement poses problems similar to those for source-to-cloud movement, two distinct questions arise:

- How many clouds will be required for processing the data, considering the costs involved in moving data among them?
- What resources (number of virtual machines, storage capacity, network bandwidth, and so forth) will each cloud require for optimal data movement?

Quality of service

Datacenter networks are critical to data-intensive application scalability and performance: optimizing QoS requires extensive communication among compute nodes. Most current datacenter cloud computing infrastructures simply can't provide scalable data-movement performance comparable to that of HPCs because they have insufficient network bandwidth and must base data movement on a best-effort policy.[2]

Predictable execution of data-intensive applications in these environments imposes two additional requirements:

- Data-intensive applications' data-movement performance should be determined separately from other applications that may be running.
- Provisioning resources for data-intensive applications should be autonomous, adapting to guarantee QoS despite dynamic network status changes.

Challenges like these are already being addressed, making prospects for cloud-operable data-intensive applications increasingly positive.

SOME WAYS FORWARD

Coping with big data from distributed sources and accommodating intercloud movement of large datasets at high transfer rates can be accomplished by implementing two measures: minimizing data-transfer overhead and better orchestration of multiple dataflows. For optimized intracloud movement within networks, emerging technologies such as QoS-aware resource management offer opportunities for fine-grained resource provisioning.

Minimizing overhead

As several research networks demonstrate (www.es.net/randd/100g-testbed), state of the art for high-speed network linkage is currently 100 Gbits per second. Proposed approaches for minimizing data-transfer overhead include

- eliminating redundant data copies using remote data-memory access;[3]
- harnessing data locality to avoid delayed data access time via application/CPU allocation-aware dataflow processing, which allocates application processing and protocol processing on the same core or on cores that share cache;[4]

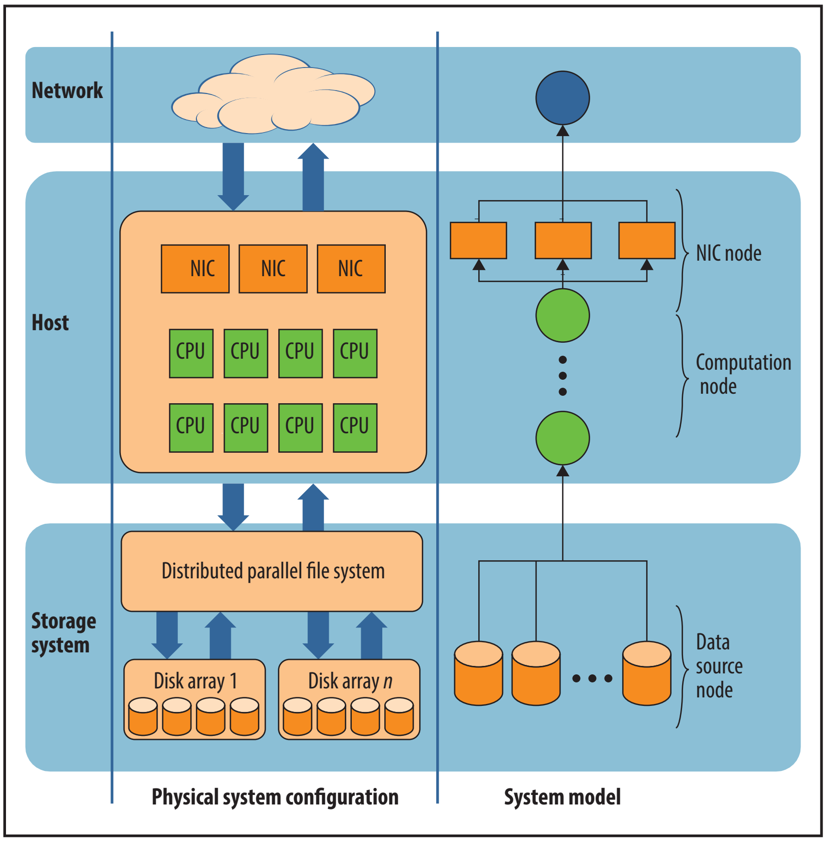

Figure 1. Graph-based modeling of a data-transfer system to optimize data transfer. In this approach, the optimization problem is formulated mathematically based on analytic models. NIC: network interface controller.

- reducing control-message wait time by, for example, using GridFTP pipelining for transferring many small files;[5] and
- isolating data transfer from computation through dedicated transfer nodes (DTNs) and a campus or lab "Science DMZ" (http://fasterdata.es.net/science-dmz) to create a stable data movement framework.

It should be noted that these approaches haven't yet been applied to cloud platforms, and can't be adopted for some cloud platforms directly. For example, for cloud architectures to incorporate DTNs and a Science DMZ, a dynamic resource manager is required to allocate DTNs in proper numbers and to reserve storage space on applications' data-transfer requests.

Orchestrating multiple dataflows

Multiple dataflows are already being used to achieve higher throughput. For example, multipath TCP (MPTCP) exploits multiple dataflows (www.multipath-tcp.org), even though optimized multipath establishment isn't an MPTCP role.

Such parallelism can also occur in the disk storage layer, the data-processing layer using multicores, the network-path layer, and the application layer using multiple nodes. One group of researchers has proposed an impedance-matching technique taking into account multiple layers involved in data movement.[6] Similarly, we earlier proposed graph-based modeling of data-transfer systems to optimize data transfer throughput,[7] as illustrated in Figure 1. Our approach can mathematically formulate the optimization problem based on analytic models of the components involved in data movement.

Nevertheless, improvement is clearly needed to accurately model time-varying components. To deal with the inherently dynamic status of wide area networks and unpredictable performance on cloud resources, and to guarantee sustainable data-transfer rates, researchers should focus on continuous monitoring and building predictable models on time-varying network status and data-movement performance in the cloud.

QoS-aware resource management

Intracloud movement can also be improved. Arguably, QoS-aware cloud resource management has been difficult because application performance on cloud resources is unpredictable and fine-grained resource provisioning isn't currently possible.

Nevertheless, emerging technologies such as network virtualization and software-defined networking have enabled researchers to do fine-grained traffic engineering and network path control, which could lead to better network utilization for datacenter networks. We believe cloud computing is moving toward advanced operation (or resource) management, where all resources are provisioned in fine-grained and predictable ways that weren't possible until now. Sophisticated QoS-aware scheduling algorithms should provide holistic resource provisioning while guaranteeing QoS.

Due to growing interest in embracing data-intensive applications, we envision a bright future for data-intensive computing on the cloud. Realizing this vision, however, requires enhanced data management over intercloud and intracloud networks, as well as adoption of emerging network and storage technologies, as outlined here.

Barriers to adopting these new infrastructures certainly exist. For example, the relatively high cost of advanced network and storage infrastructures such as the dragonfly network topology[8] might deter cloud service providers from deploying them. Nevertheless, increasing user demand for data-intensive applications in the cloud, and the potential benefits for cloud service providers such as reduced energy consumption through efficient data movement, will help accelerate adoption of new data-management infrastructures. C

References

1. G. Chen et al., "Federation in Cloud Data Management: Challenges and Opportunities," *IEEE Trans. Knowledge and Data Eng.*, vol. 26, no. 7, 2014, pp. 1670–1678.
2. A. Marathe et al. "A Comparative Study of High-Performance Computing on the Cloud," *Proc. 22nd Int'l Symp. High-Performance Parallel and Distributed Computing* (HPDC 13), 2013, pp. 239–250.
3. B. Tierney et al., "Efficient Data Transfer Protocols for Big Data," *Proc. 8th IEEE Int'l Conf. E-Science* (E-Science 12), 2012, pp. 1–9.
4. V. Ahuja, M. Farrens, and D. Ghosal, "Cache-Aware Affinitization on Commodity Multicores for High-Speed Network Flows," *Proc. 8th ACM/IEEE Symp. Architectures for Networking and Comm. Systems* (ANCS 12), 2012, pp. 39–48.
5. J. Bresnahan et al., "GridFTP Pipelining," TeraGrid 2007 Conf.; http://toolkit.globus.org/alliance/publications/papers/Pipelining.pdf.
6. B.W. Settlemyer et al., "A Technique for Moving Large Data Sets over High-Performance Long Distance Networks," *Proc. 27th IEEE Symp. Mass Storage Systems and Technologies* (MSST 11), 2011, pp. 1–6.
7. E.-S. Jung, R. Kettimuthu, and V. Vishwanath, "Toward Optimizing Disk-to-Disk Transfer on 100G Networks," *Proc. 2013 IEEE Int'l Conf. Advanced Networks and Telecommunications Systems* (ANTS 13), 2013, pp. 1–6.
8. J. Kim et al., "Technology-Driven, Highly-Scalable Dragonfly Technology," *Proc. 35th Int'l Symp. Computer Architecture* (ISCA 08), 2008, pp. 77–88.

***Eun-Sung Jung** is a postdoctoral researcher in the Mathematics and Computer Science Division at Argonne National Laboratory. Contact him at esjung@mcs.anl.gov.*

***Rajkumar Kettimuthu** is a project leader in the Mathematics and Computer Science Division at Argonne National Laboratory and a fellow at the University of Chicago's Computation Institute. He is a senior member of IEEE and ACM. Contact him at kettimut@mcs.anl.gov.*

Editor: San Murugesan, BRITE Professional Services, Sydney, Australia; san@computer.org