About Omics Group

OMICS Group International through its Open Access Initiative is committed to make genuine and reliable contributions to the scientific community. OMICS Group hosts over 400 leading-edge peer reviewed Open Access Journals and organize over 300 International Conferences annually all over the world. OMICS Publishing Group journals have over 3 million readers and the fame and success of the same can be attributed to the strong editorial board which contains over 30000 eminent personalities that ensure a rapid, quality and quick review process.

About Omics Group conferences

- OMICS Group signed an agreement with more than 1000 International Societies to make healthcare information Open Access. OMICS Group Conferences make the perfect platform for global networking as it brings together renowned speakers and scientists across the globe to a most exciting and memorable scientific event filled with much enlightening interactive sessions, world class exhibitions and poster presentations
- Omics group has organised 500 conferences, workshops and national symposium across the major cities including SanFrancisco,Omaha,Orlado,Rayleigh,SantaClara,Chicago,Philadelphia,Unitedkingdom,Baltimore,SanAntanio,Dubai,Hyderabad,Bangaluru and Mumbai.

Networks for Large Data Flows - Revolution or Evolutions?

Weiqiang Sun and Weisheng Hu
Shanghai Jiao Tong University
{sunwq, wshu}@sjtu.edu.cn
sunwq@mit.edu

上海交通大學
SHANGHAI JIAO TONG UNIVERSITY

Optics 2014, Philadelphia, USA.
Sept. 8-10, 2014.

Outline

- The Increasing Challenges of Big Data on the network infrastructure
- The Characteristics of Data Flows
- Existing Efforts in Delivering Big Data
- Our Proposal - Integrated Data Flow Delivery with Built-in Mass Storage
- Some Results

The Ever-increasing Demand

- In 2013, data generated everyday exceeds 1EB (10^{18} bytes)
- In the coming 10 years, the data generated will increase by 50-folds
- The traffic between Data Centers will increase 34% every year

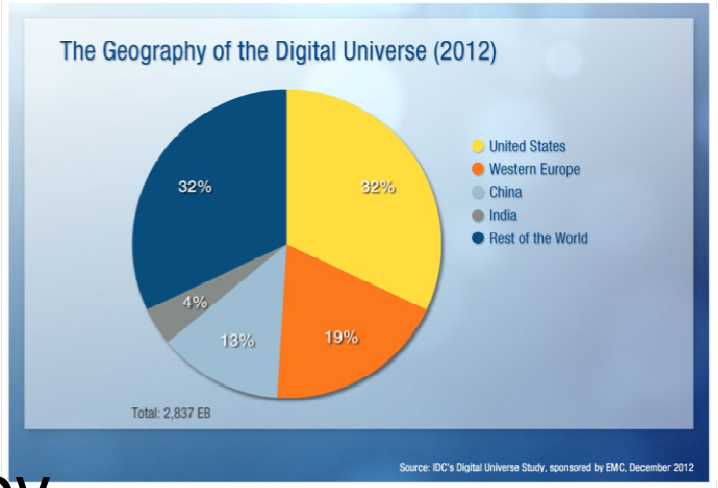

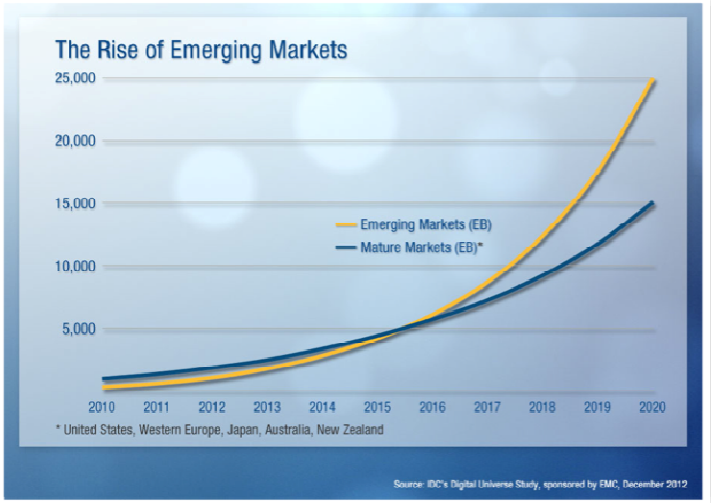

Big Data From the Science Domain

- The demand of moving large data set among research institutions has been immense
- Genomics
  - Data volume generated by HGP 10-folds every 12-18 Months
  - Still rely on courier mail of HDs or tapes (BGI) *
- High Energy Physics (LHC at CERN)
  - 10s - 100s TB, distributed on a daily basis
- Connectomics
  - 1 mm^3 of brain image could produce 1PB (10^{15} bytes)

Genomics

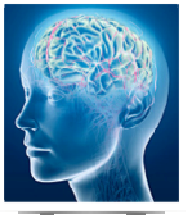
Connectomics

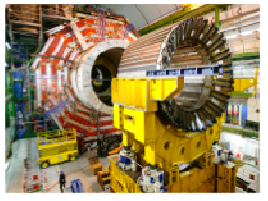
Physics


Meteorology

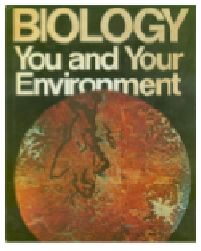

Biology


Environmental sciences

Big Data From the End Users

- Huge Email attachments
  - Gmail allows attachments up to 10GB from Nov. 2012. Before that, it was 25MB
  - outlook.com and hotmail allows 300GB (in OneDrive)
  - Yahoo mail allows unlimited attachments, if file is attached via dropbox
- This makes distributing and sharing of huge files a lot more easier than before
- Cloud storage plays an important role

sunwq@sjtu.edu.cn

Data Movement between DCs

- By 2015, the traffic betw. DCs will reach 1 ZB (10^{21} bytes)
- The cost of network infrastructure is dominated by the Inter-DC Net. [1]
  - Peering - where traffic are handled to ISP and then to users
  - The inter-DC links
  - Backhaul, metro-connectivity and others to reach the WAN sites

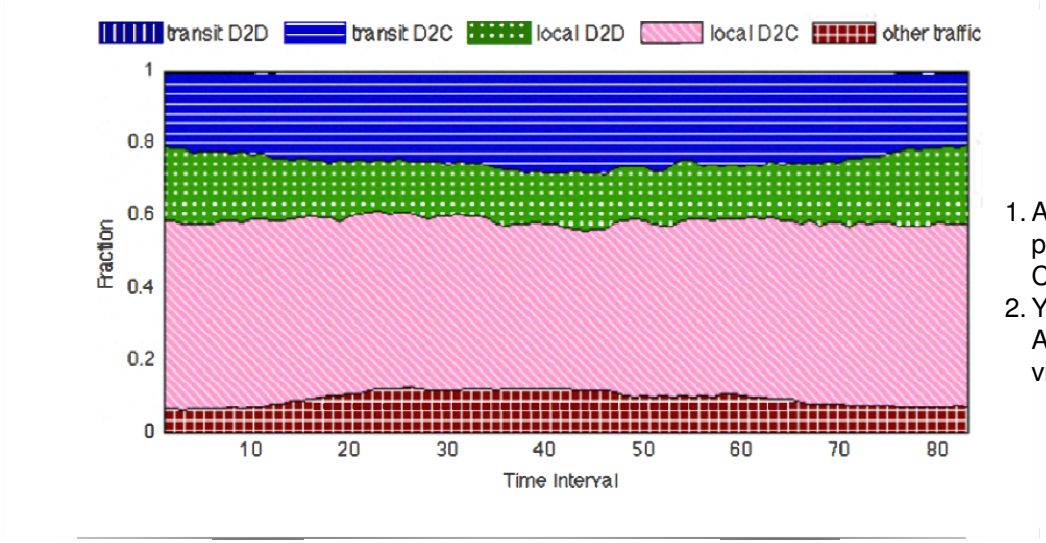

The fraction of different types of traffic at one data center [2]

1. A. Greenberg et al. "The cost of a cloud: research problems in data center networks." ACM SIGCOMM Computer Communication Review 39.1 (2008): 68-73.
2. Y. Chen, S. Jain, V. K. Adhikari, Z.-L. Zhang, and K. Xu. A First Look at Inter-Data Center Traffic Characteristics via Yahoo! Datasets. In IEEE INFOCOM'11.

What Are the Problems?

- Bandwidth increase can barely keep up with the pace of demand increase
- Bandwidth increased by 6-fold, but revenue and number of customers increase only by a single digit each year

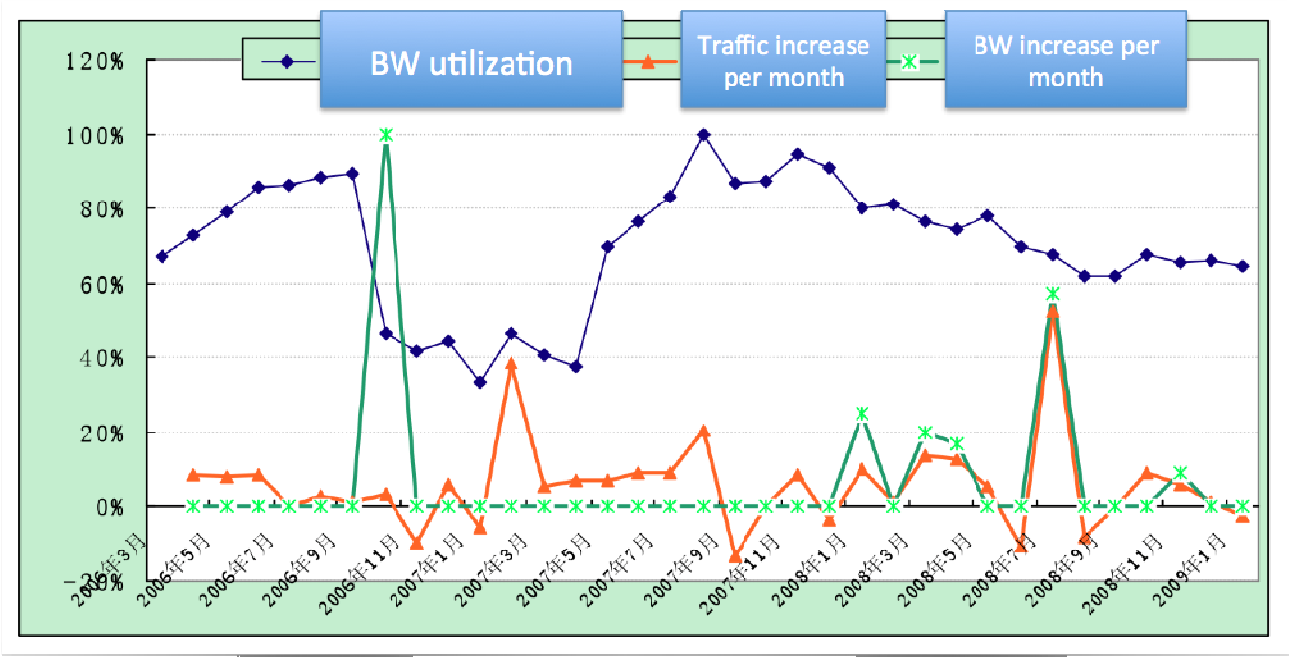

sunwq@sjtu.edu.cn

What Are the Problems? - cont.

- Elephant flows compete bandwidth with interactive but small flows
- Degrading the QoE of mice flows without bringing significant benefit to elephant flows
- Make resource sharing among large/small flow very difficult

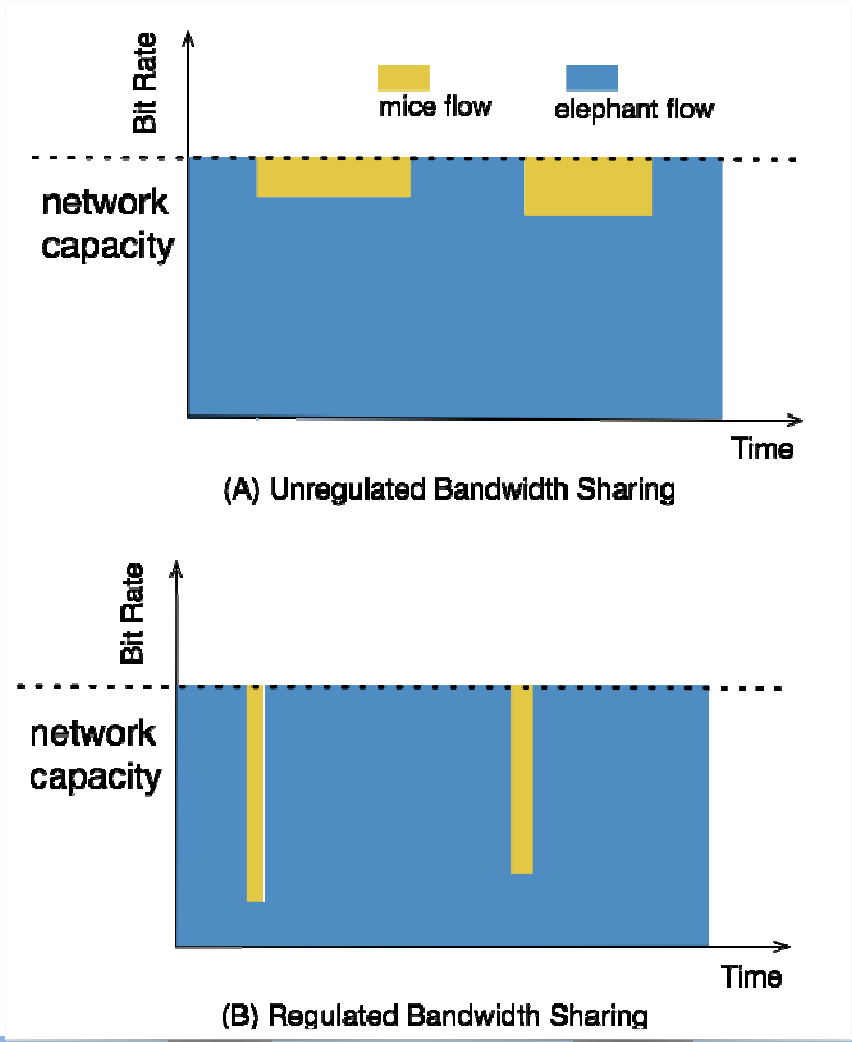

Outline

- The Increasing Challenges of Big Data on the network infrastructure
- The Characteristics of Data Flows
- Existing Efforts in Delivering Big Data
- Our Proposal - Integrated Data Flow Delivery with Built-in Mass Storage
- Some Results

Characteristics of Big Data Flows (From a transport network point of view)

Bulky

Delay Tolerant

Data Flows are Bulky

- The total traffic is dominated by bulk flows that is:
  - Very few in number: <1%
  - Big in size: between 100 - 1000MB
  - and most 100MB flows comes from larger flows
- Bulk flows occupy more than 90% of total bandwidth

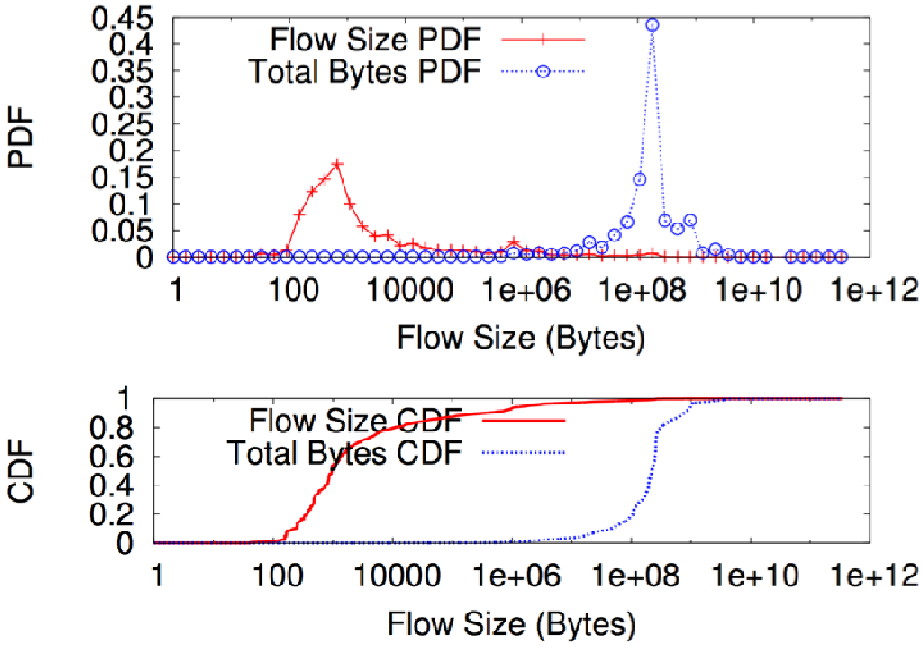

* A. Greenberg et al. "VL2: a scalable and flexible data center network." ACM SIGCOMM Computer Communication Review. Vol. 39. No. 4. ACM, 2009.

Data Flows are Delay Tolerant

- Data flows in E-Science is often delay tolerant
  - Genomic data
  - HEP data
- Data flows between DCs are dominated by background traffic *
  - Backups
  - Content distributions and so on

* Y. Chen, S. Jain, V. K. Adhikari, Z.-L. Zhang, and K. Xu. A First Look at Inter-Data Center Traffic Characteristics via Yahoo! Datasets. In IEEE INFOCOM'11.

Outline

- The Increasing Challenges of Big Data on the network infrastructure
- The Characteristics of Data Flows
- Existing Efforts in Delivering Big Data
- Our Proposal - Integrated Data Flow Delivery with Built-in Mass Storage
- Some Results

Big Data Movement - Existing Work

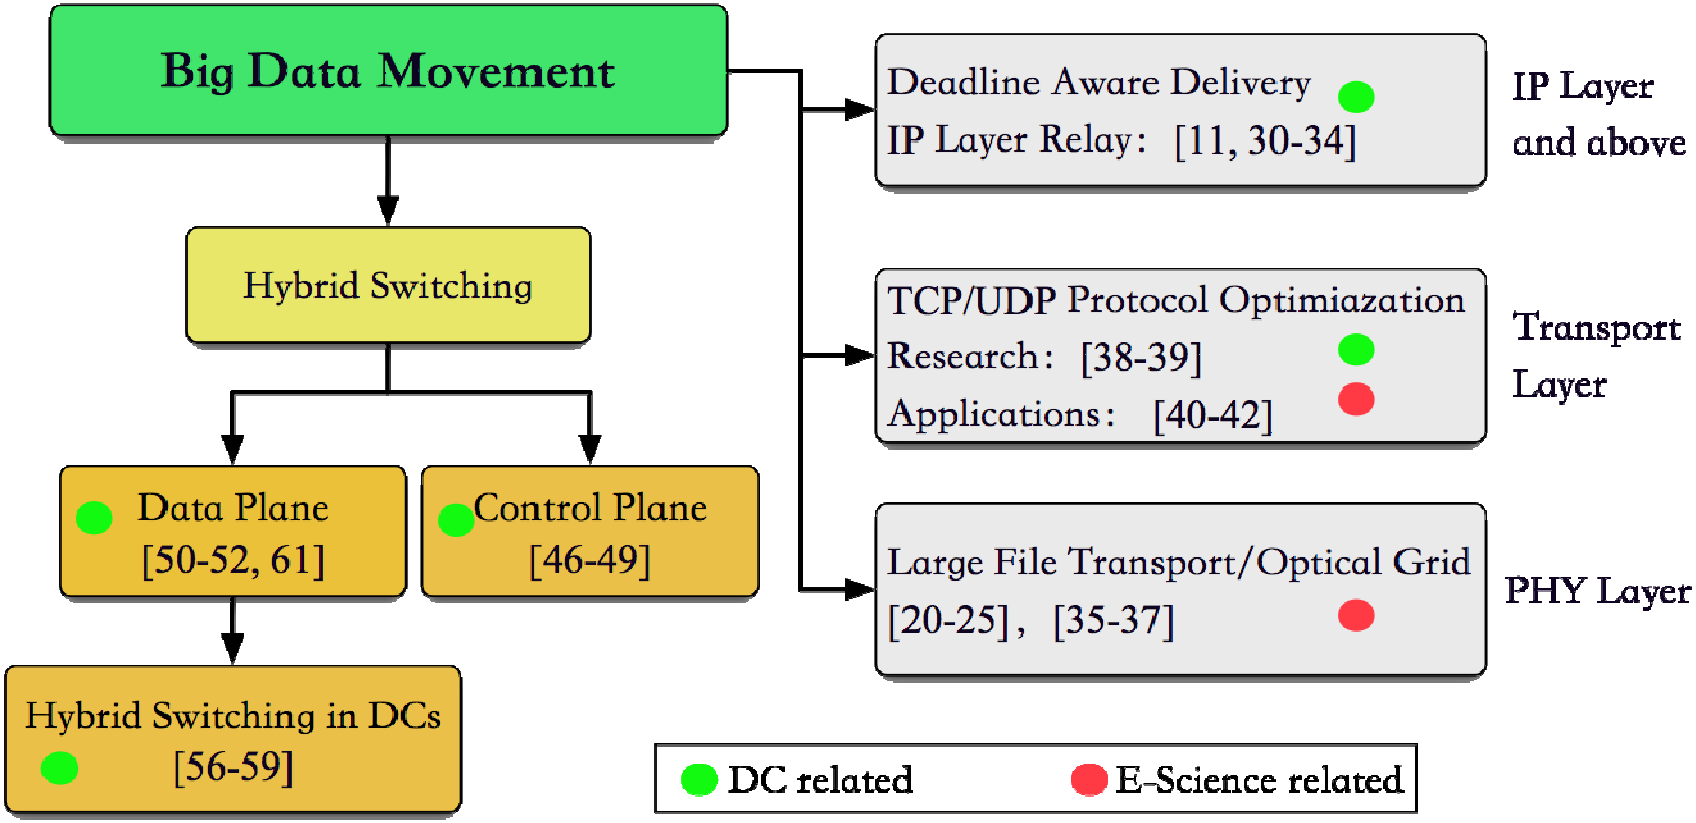

Moving bulk data with dedicated Optical Networks

- Can be dated back to early 2000s
- Use circuit switched optical networks
- A lot of interesting research and testbeds

	Networking Technology	Provisioning Method	Provisioned Channel capacity	Applications
USN	SONET/10 Gigabit Ethernet (core) and Multi-Service Provisioning Platform (edge)	Centralized scheduling and signaling within a VPN	10Gbps, 1Gbps and high-precision channels such as SONET OC-1	Large data transfer
CHEETAH	SONET (core) and Multi-Service Provisioning Platform (edge) with packet switching for dual homing	GMPLS with hardware-accelerated signaling	10Gbps and 1Gbps	Large file transfer and remote visualizations
UCLP/ CA*net4	Lightpath Switching	Centralized provisioning through Web Service	Wavelength (10Gbps) and sub-wavelength	High performance computing
DRAGON	DWDM switching (core) and Ethernet, TDM, IP (edge)	GMPLS with centralized broker	Wavelength, Ethernet and IP	e-VLBI
SURFnet6	Lambda switching and packet switching	User controlled provisioning	Wavelength	Large data transfer
OptiPuter	Lambda Switching with packet switching for dual homing	Client-provisioning through Web Service	Wavelength	Distributed Virtual Computer (DVC)
OMNInet	DWDM switching and L2/L3 devices	GMPLS with OIF UNI	10/100/1000Mbps	Grid applications
3TNet	SONET (core) and Multi-Service Provisioning Platform (edge)	GMPLS with OIF UNI	1000Mbps	Large data transfer

Resulted in good experience in building dedicated and small scale networks. But not intended for large scale deployment because of scalability issues.

sunwq@sjtu.edu.cn

Transport Protocol Optimizations

- Over High speed long distances networks
- Within Data Centers
- Over the public Internet
- Over dedicated high speed networks

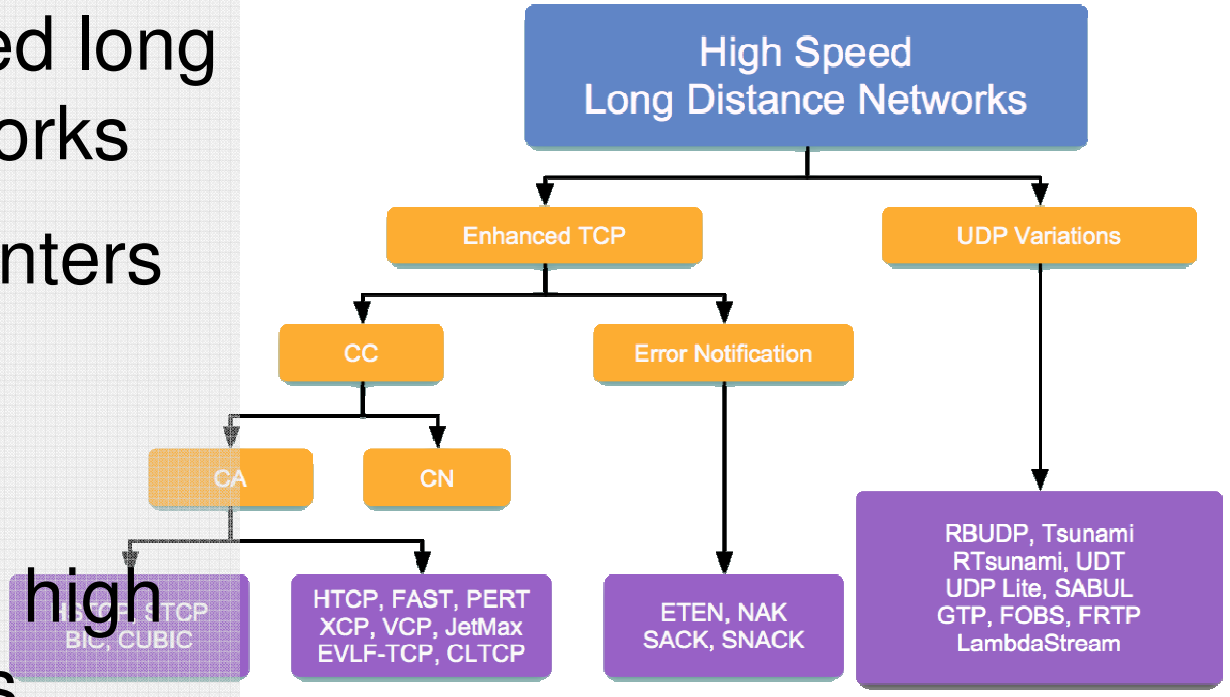

Necessary and important enhancements to existing protocols, but will not be able to address the scalability issues (capacity, power consumption and management).

Moving Bulk Data with the public Internet

- By optimizing the transport layer and application layer protocols
- By utilizing unused bandwidth

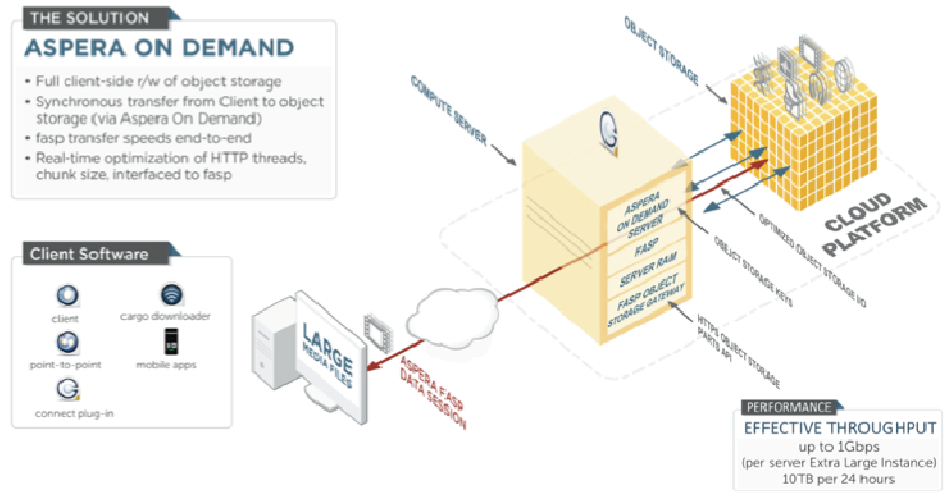

Make the best use of the current infrastructure and not considered to be a long term solution.

Hybrid Switching

- Try to leverage the advantages of both
  - Fine granular packet switching
  - Coarse granular, large capacity optical circuit switching
- Different modes
  - Parallel Mode
  - Client/Server Mode
  - Integrated Mode

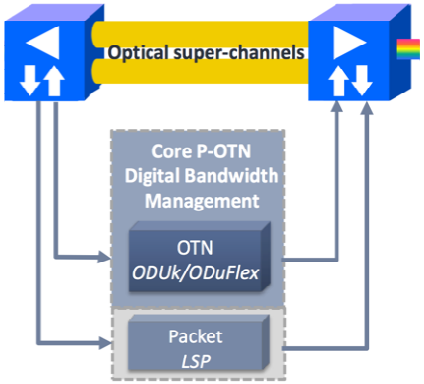

Infinera, ECOC 2013

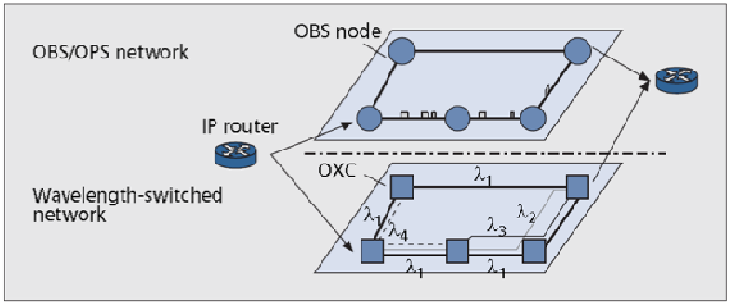

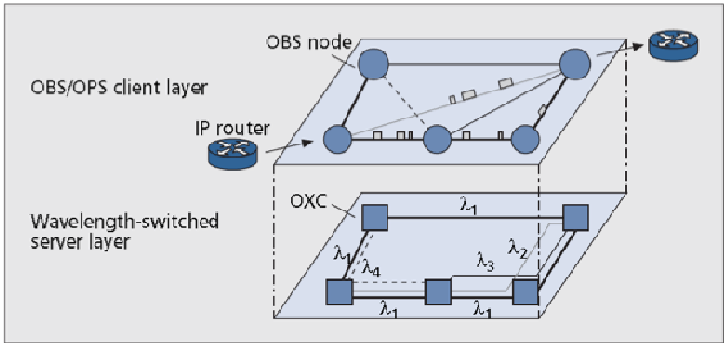

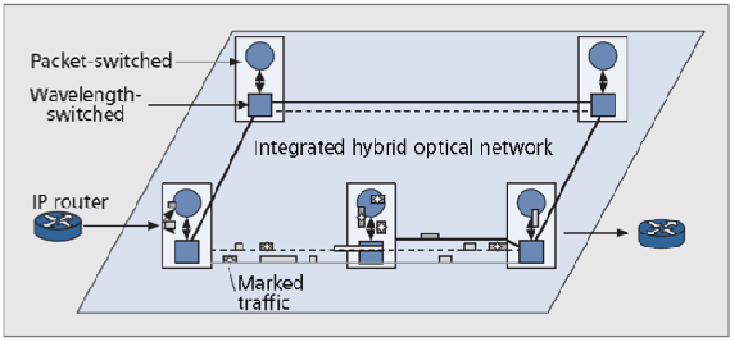

C. M. Gauger et al., COMMAG, 2006

Hybrid Switching in DCNs

- The packet-switched portion
  - all-to-all bandwidth for the bursty traffic
- The circuit-switched portion
  - baseline, slowly changing traffic

- Significant benefits
  - Up to a factor of 3 reduction in cost
  - A factor of 6 reduction in complexity
  - And a factor of 9 reduction in power consumption

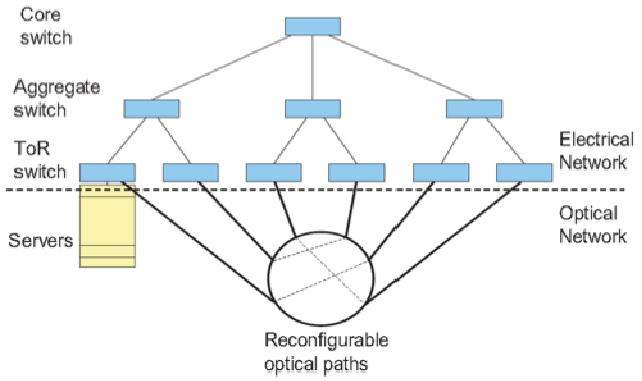

G. Wang et al., SIGCOMM 2010

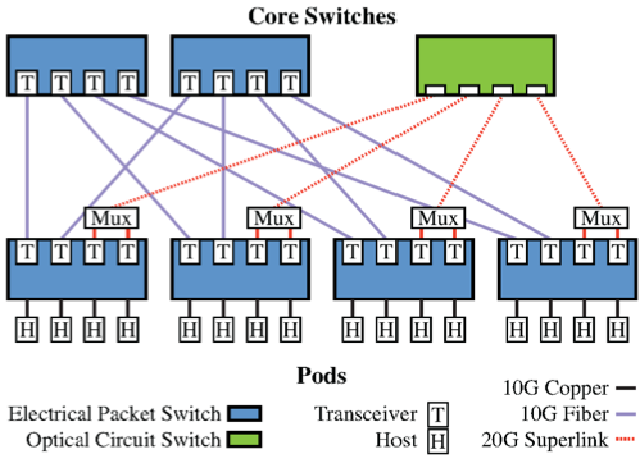

N. Farrington, et al., SIGCOMM 2010

SDN - the Control Plane-ng?

- A centralized way of controlling the network elements
- Separation of Data Plane and Control Plane
- Flow based management and control

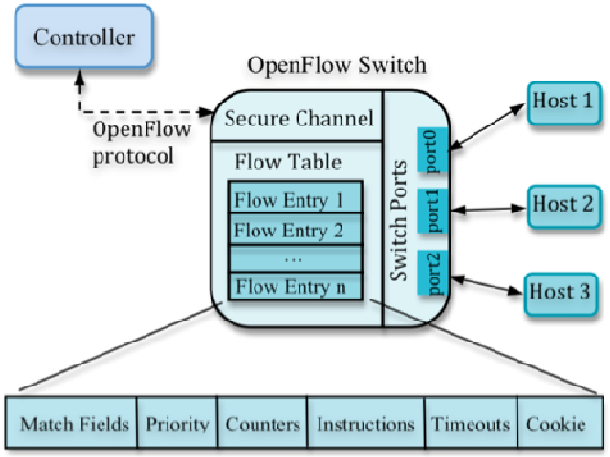

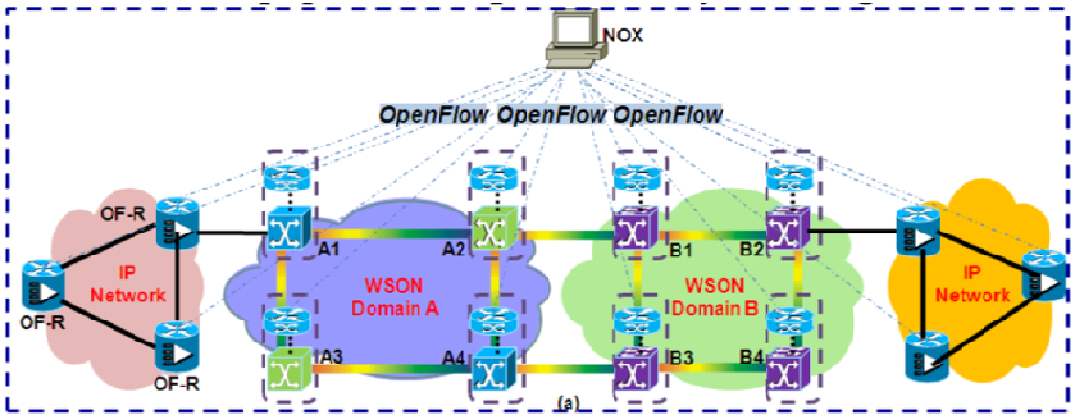

sunwq@sjtu.edu.cn

Big Data Movement - the Evolutions!

High Speed Transmission for E-Science

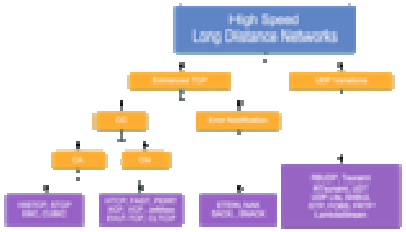

High Speed Transmission in LDN

Early 2000~

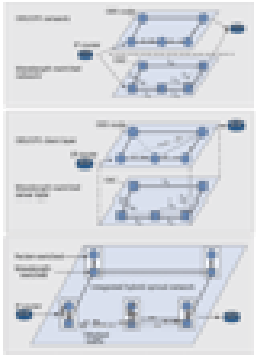

Hybrid Switching

2005~

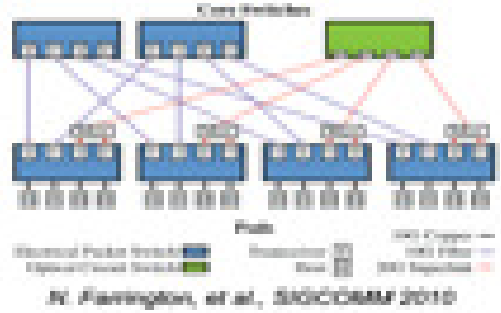

Hybrid Switching in DCN

2010~

Converging with Flow Switching

2012-

Outline

- The Increasing Challenges of Big Data on the network infrastructure
- The Characteristics of Data Flows
- Existing Efforts in Delivering Big Data
- Our Proposal - Integrated Data Flow Delivery with Built-in Mass Storage
- Some Results

SSS- Integrated Data Flow Delivery with Built-in Mass Storage

- High capacity optical switch for big data transfer and VT provisioning
- Low capacity Electronic Switch for fine granular packet switching
- In-network mess storage (for big data)

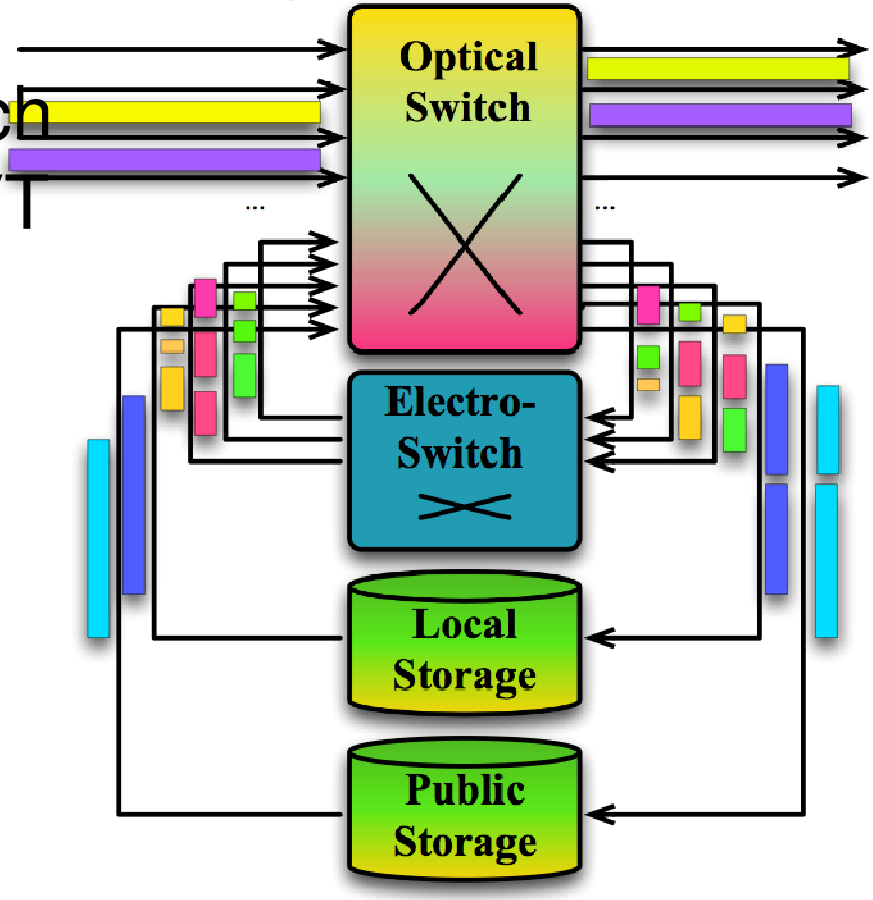

SSS - Integrated Data Flow Delivery with Built-in Mass Storage

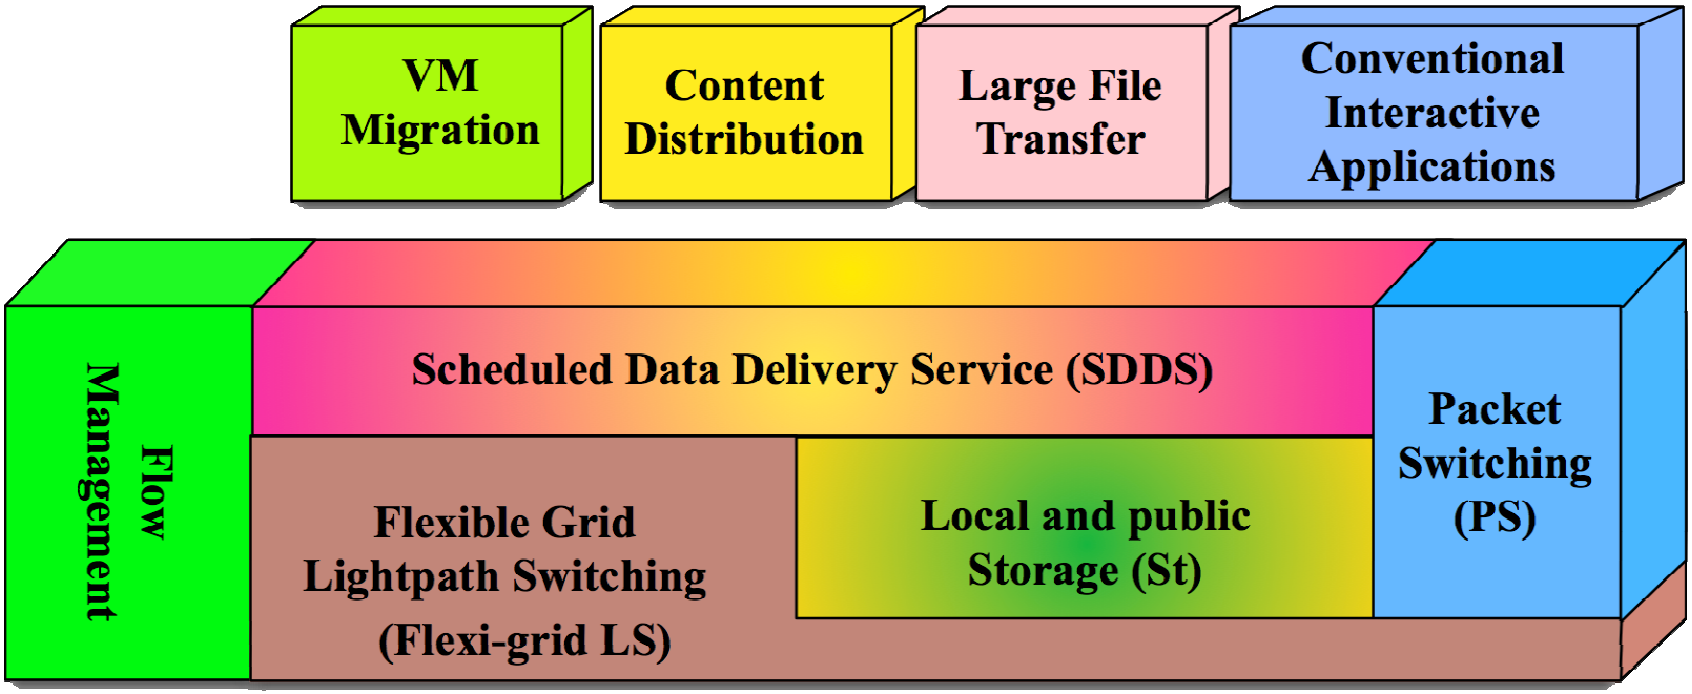

Prototype Implementation

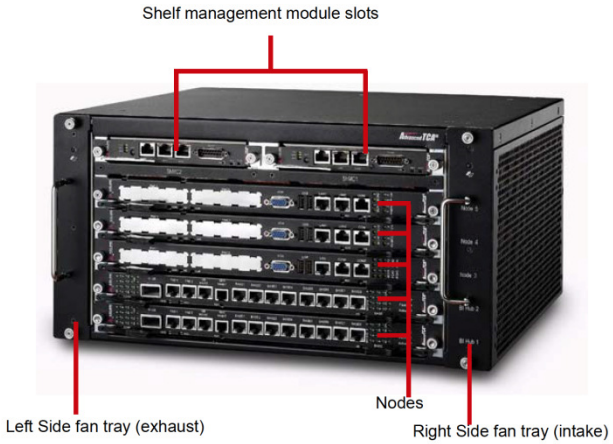

- Built on an ATCA chassis
  - AWGR+TWC for Optical Switching
  - E-Switching
  - Local Storage
- Network controller with OpenFlow

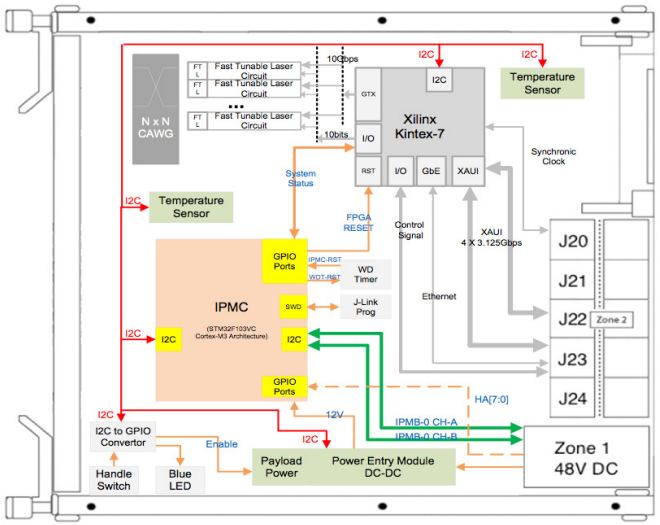

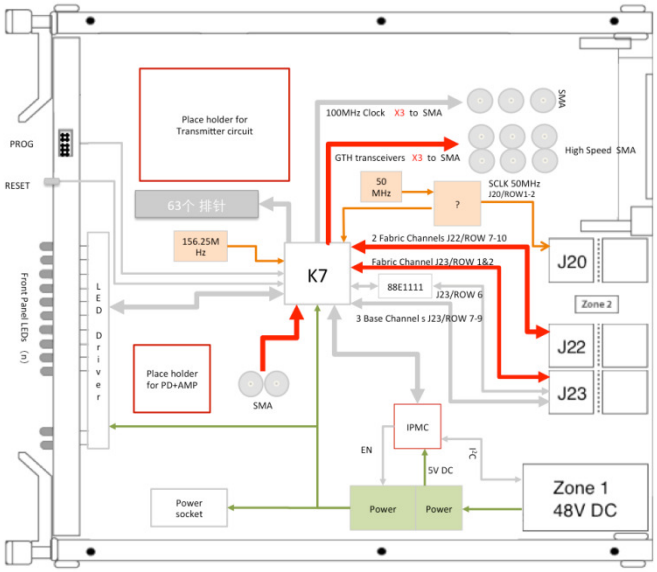

Outline

- The Increasing Challenges of Big Data on the network infrastructure
- The Characteristics of Data Flows
- Existing Efforts in Delivering Big Data
- Our Proposal - Integrated Data Flow Delivery with Built-in Mass Storage
- Some Results

Some Results

- Single queue with traffic aggregation
- Requests/clients has a deadline
- Use Earliest Deadline First (EDF) policy

Loss Rate vs. Load

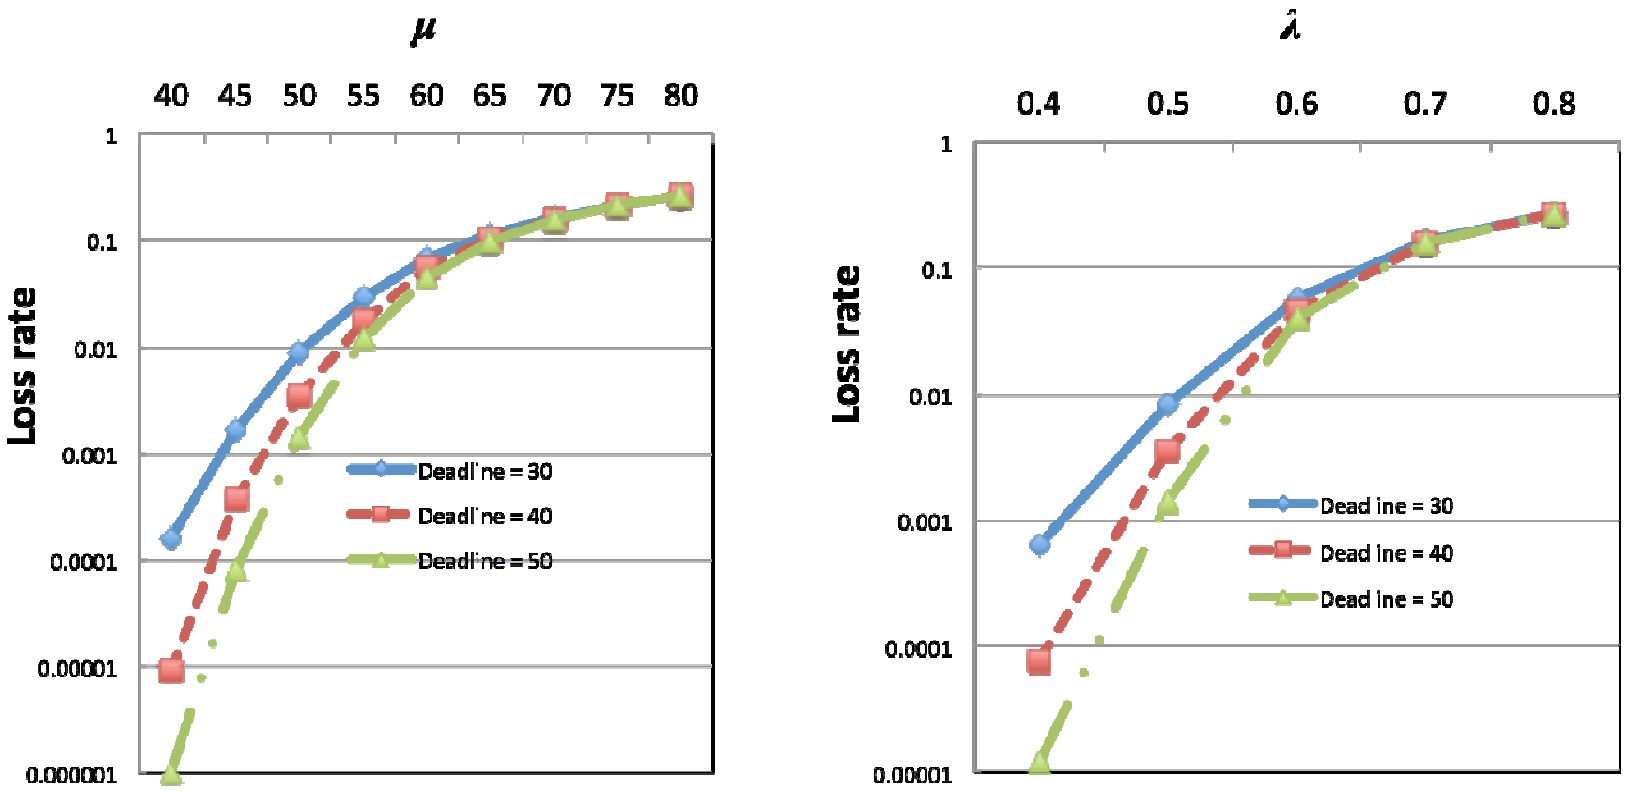

C = 30, # request in a batch = 2

Loss Rate vs. Deadline

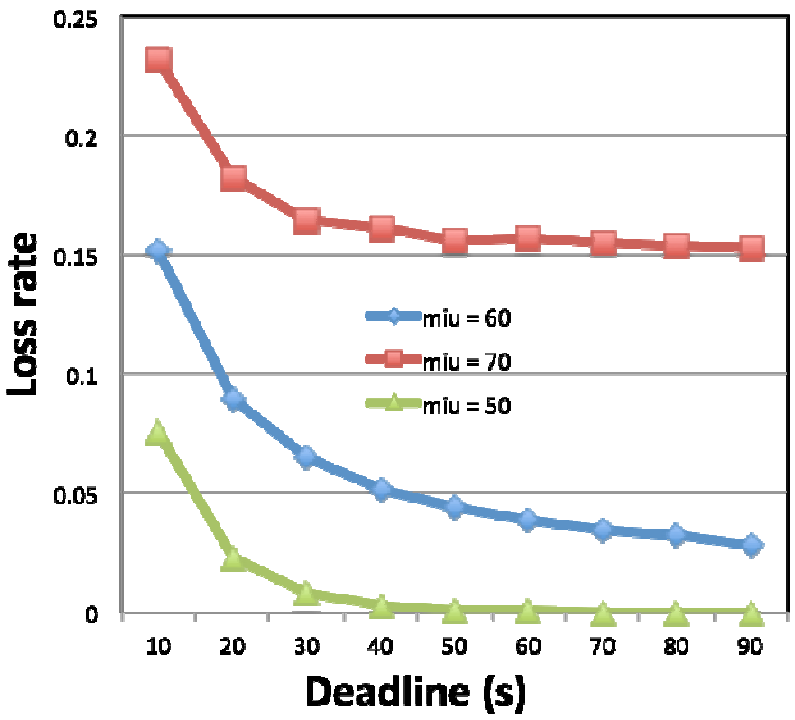

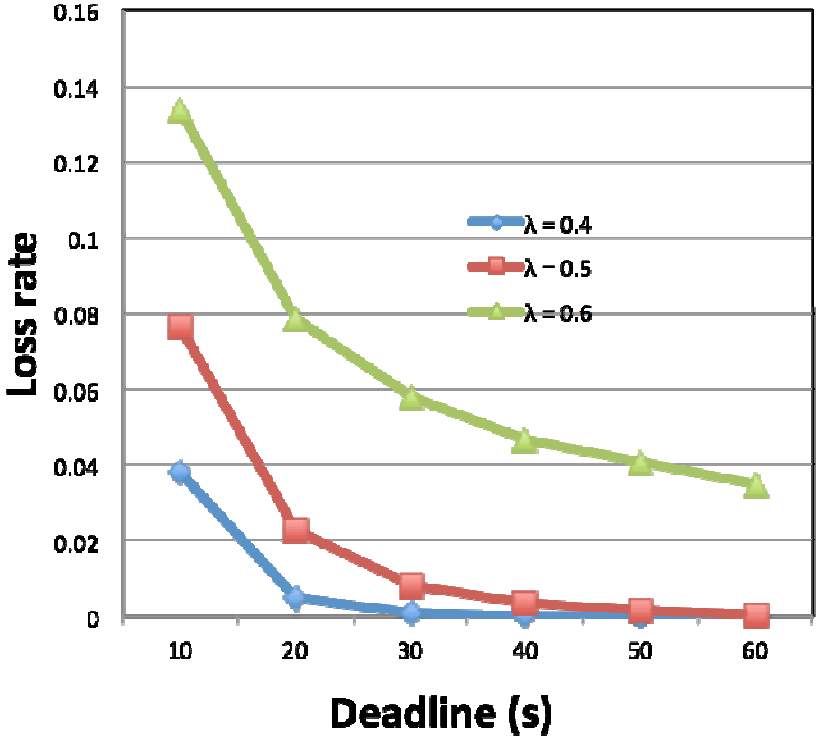

Link capacity = 30

sunwq@sjtu.edu.cn

Loss Rate vs. Batch Size

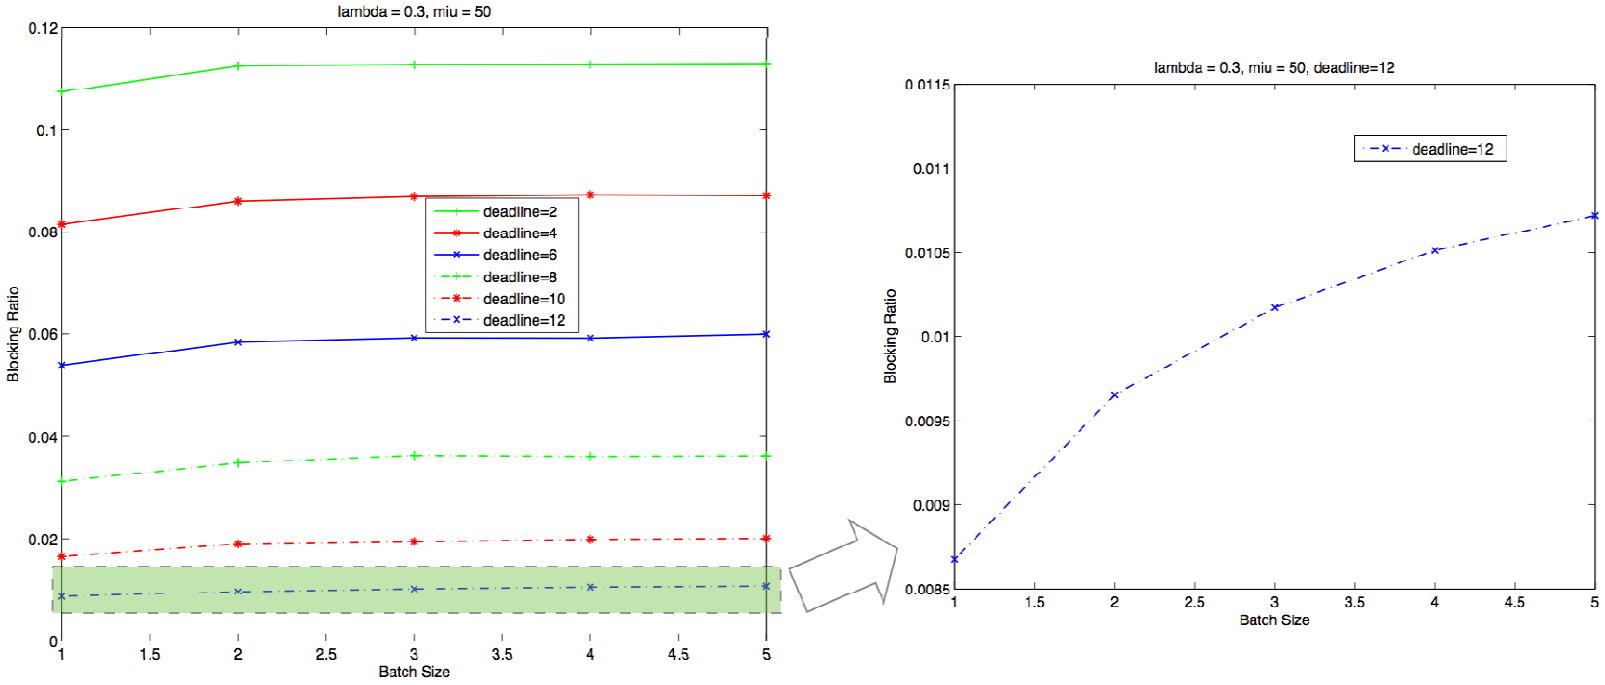

sunwq@sjtu.edu.cn

Average Delay vs. Load

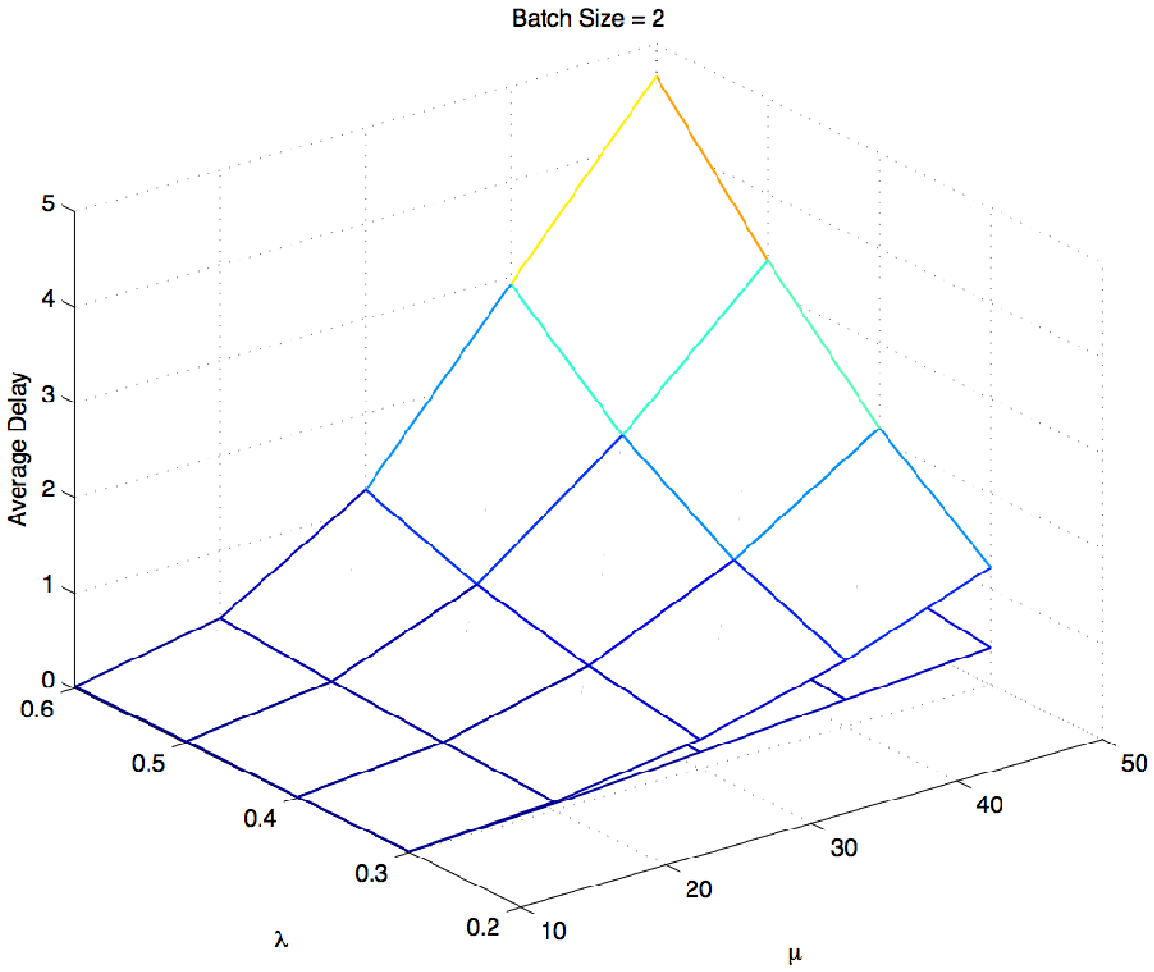

sunwq@sjtu.edu.cn

Conclusions

- Data increasing at an unprecedented pace
- Data flows are bulky and delay tolerant
- Innovations converge at the cloud age with flow switching and the SDN concept
- We propose to use hybrid switching with built-in storage to support big data delivery

Thank you!

Weiqiang Sun and Weisheng Hu
Shanghai Jiao Tong University
{sunwq, wshu}@sjtu.edu.cn

上海交通大學
SHANGHAI JIAO TONG UNIVERSITY

Optics 2014, Philadelphia, USA.
Sept. 8-10, 2014.

The Big Data Initiative by NSFC

- Date: 2015.1 - 2019.12
- Coverage: Storage, Computing, Processing, Transporting of Big Data
- Number of Projects: 12
- Support Level: 3.0 ~ 3.5 M RMB/project

More to come in 2015…

Let Us Meet Again

We welcome all to our future group conferences of Omics group international

Please visit:

www.omicsgroup.com

www.Conferenceseries.com

http://optics.conferenceseries.com/