

# Networks and/in data centers

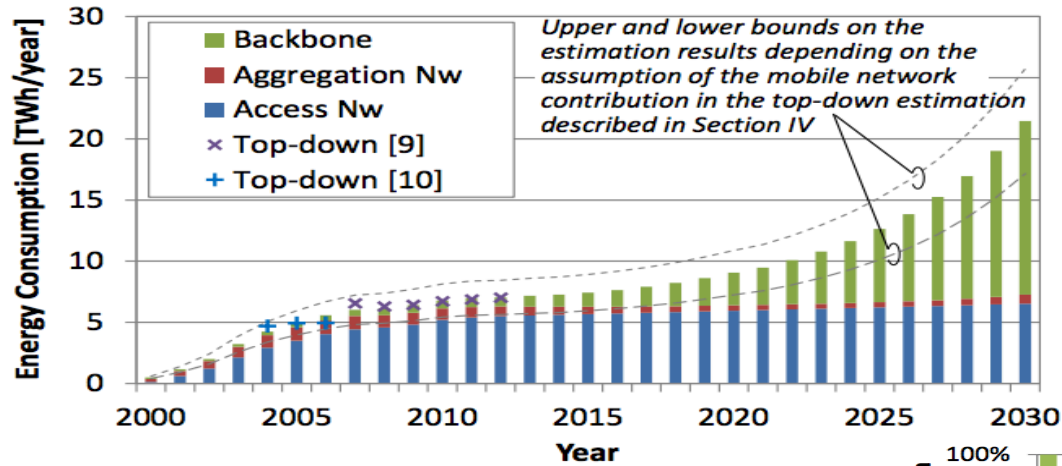
Dr. Paola Grosso

(with contributions from Fahimeh Alizadeh and Hao Zhu)

# Does the network matter?

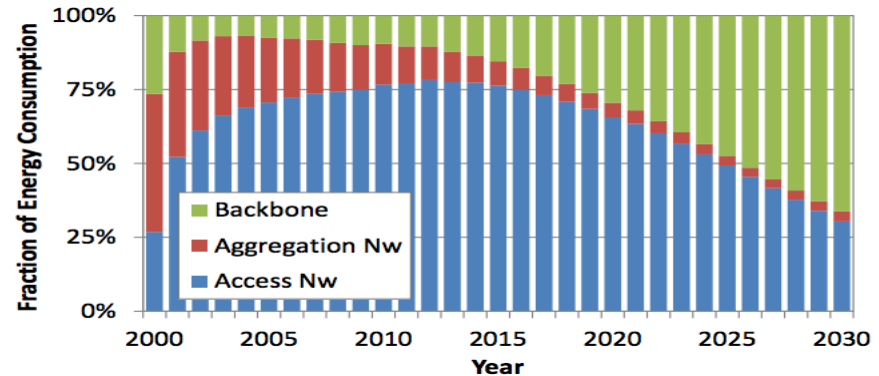
- The energy footprint of data center networks is 10-20% of total energy
- The proportion could be up to 50% if power management techniques are used on the server-side
- Global cloud IP traffic will 3.9-fold from 2013 to 2018.
- Networks suffer from inefficient power usage because of over-provisioning.

# Learned at ICT4S!



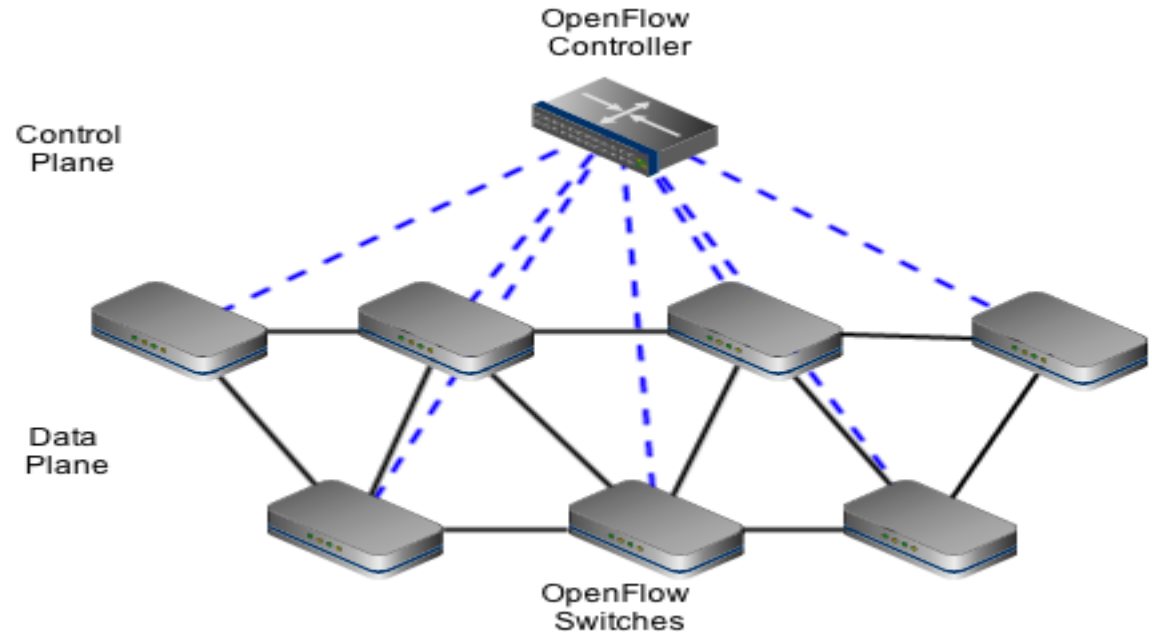
“Unifying Top-down and Bottom-up Approaches to Evaluate Network Energy Consumption”

Ishii et al. In: Lightwave Technology, Journal of , vol.PP, no.99, pp.1-1



# Whats new? Programmability


- 87% growth in SDN deployment in North American-based enterprises by 2016 [2]
- White-box networking
- Diverting the control to the software



[\*] Clifford Grossner, <http://www.infonetics.com/pr/2014/SDN-Strategies-Survey-Highlights.asp> (July 2014)]

# Emerging interest in SDN for energy efficiency

- Emerging studies improve on the energy consumption of servers by the VM migration.
- Some change the OpenFlow protocol to be energy-aware.
- All of them are implemented in intra-data center scale.
- All have a fixed initial traffic matrix.



Moghaddam, F. A., Lago, P., & Grosso, P. (2015). Energy-Efficient Networking Solutions in Cloud-Based Environments: A Systematic Literature Review. *ACM Computing Surveys (CSUR)*, 47(4), 64.

Yearly distribution of the OpenFlow technique adoption by decision frameworks from December 2008 to November 2013 [3]

# Green routing with SDN


Make a routing decision to aggregate traffic over a subset of links and devices in over-provision networks and switch off unused network components

Configured routes

Route: id0, Source/target: 192.168.122.111:192.168.121.204  
 Route: id1, Source/target: 192.168.122.111:192.168.121.205  
 Route: id2, Source/target: 192.168.122.111:192.168.121.203  
 Route: id0, Power/Cost/Emission: 374.0Watt, 0.27€/h, 0.01kg/h  
 Route: id1, Power/Cost/Emission: 445.0Watt, 0.30€/h, 0.01kg/h  
 Route: id2, Power/Cost/Emission: 378.0Watt, 0.28€/h, 0.01kg/h

Id	Src IP	Dst IP	Switch	Action
19	192.168.122.111	192.168.121.203	0	Delete
20	192.168.121.203	192.168.122.111	0	Delete
21	192.168.122.111	192.168.121.203	1	Delete
22	192.168.121.203	192.168.122.111	1	Delete
23	192.168.122.111	192.168.121.203	2	Delete
24	192.168.121.203	192.168.122.111	2	Delete
25	192.168.122.111	192.168.121.203	3	Delete
26	192.168.121.203	192.168.122.111	3	Delete

Remove this route



# Simulation

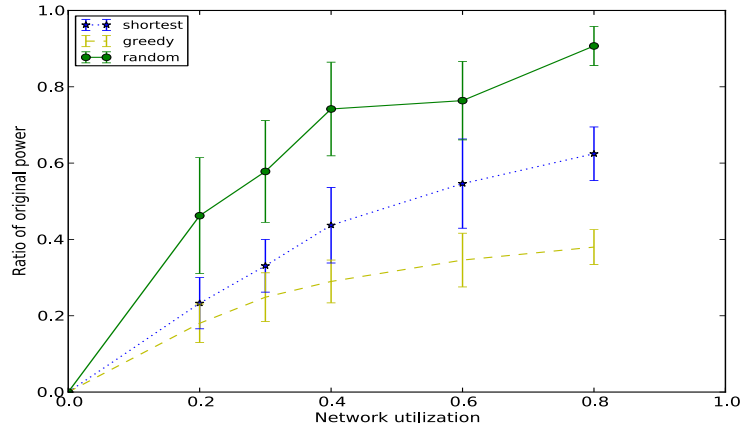
Topology: BCube(2,3) , 8 8-port switches in total.

Capacity of each link: 1Gbps;

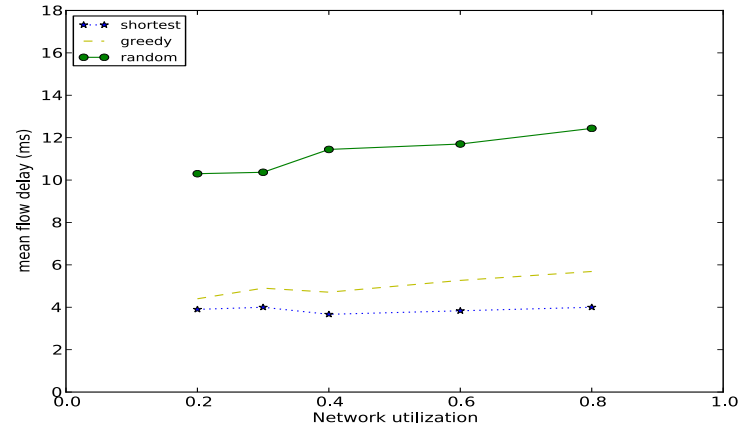
Delay of each link: 1 millisecond.

Mean power of switch: 120 watts.

Dynamic power/Overall power = 10%.



Network power saving against the network utilization in the BCube network



Mean flow delay time against the network utilization in the BCube network

# What next? Food for thoughts.

3 thoughts:

Moving from simulation to experimentation in real data center networks. *Will we see main stream use?*

Couple the programmability of the network with the application desired service. *How to manage the decentralization?*

The datacenter of the future is embedded in the fabric of the Cloud and driven by Big Data operations/analytics/computation.  
*How do we coordinate inter-data-center?*