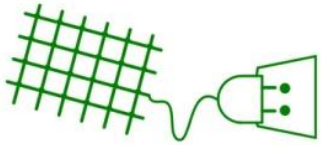

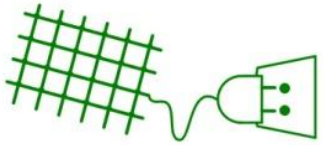

June 1, 2011

Data Center Energy-Efficient Network-Aware Scheduling

Dzmitry Kliazovich

University of Luxembourg

Why energy is important?

- Increased computing demand
  - Data centers are rapidly growing
  - Consume 10 to 100 times more energy per square foot than a typical office building

- Energy cost dynamics
  - Energy accounts for 10% of data center operational expenses (OPEX) and can rise to 50% in the next few years
  - Accompanying cooling system costs $2-$5 million per year

Distribution of data center energy consumption

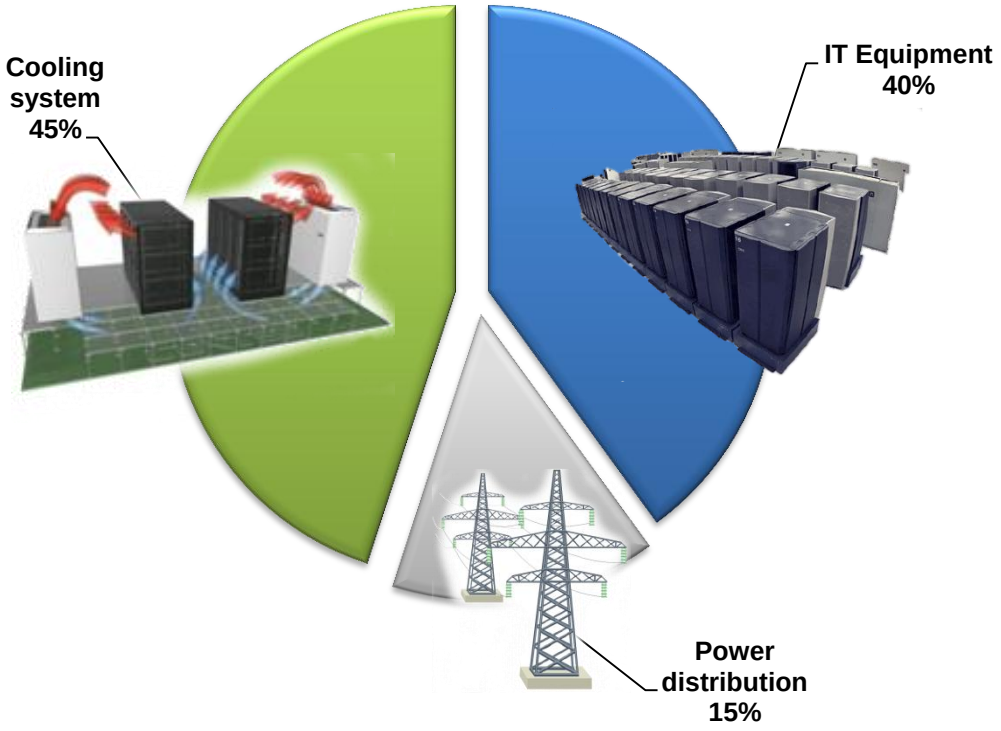

Data center architectures

- Past: Two-tier data center architecture
  - Access and Core layers
  - 1 GE and 10 GE links
  - Full mesh core network
  - Load balancing using ICMP

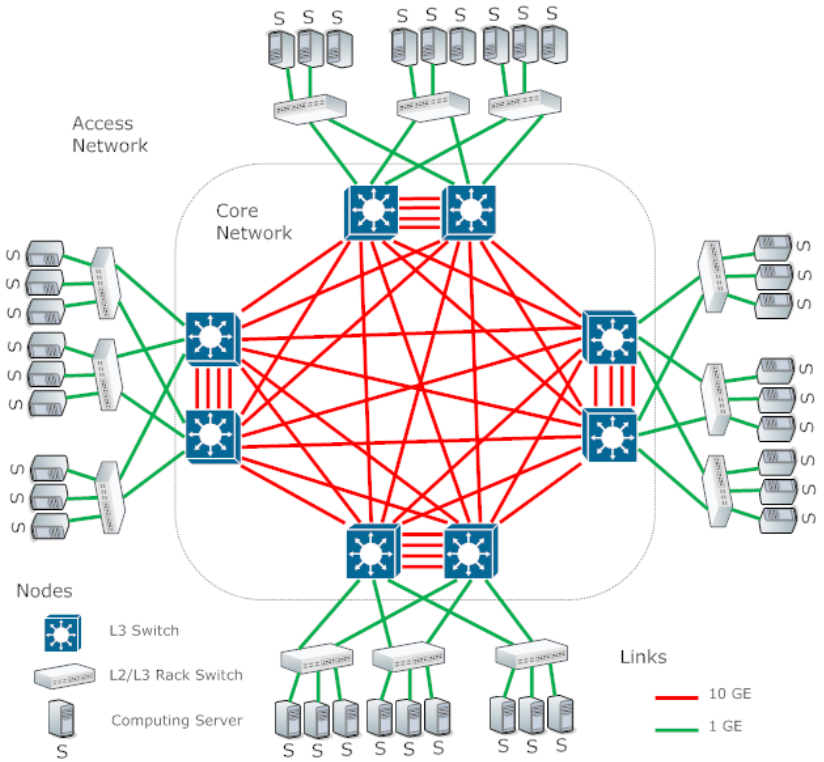

Data center architectures

- Present: Three-tier data center architecture
  - Most Widely Used Nowadays
  - Access, Aggregation, and Core layers
  - Scales to over 10,000 servers

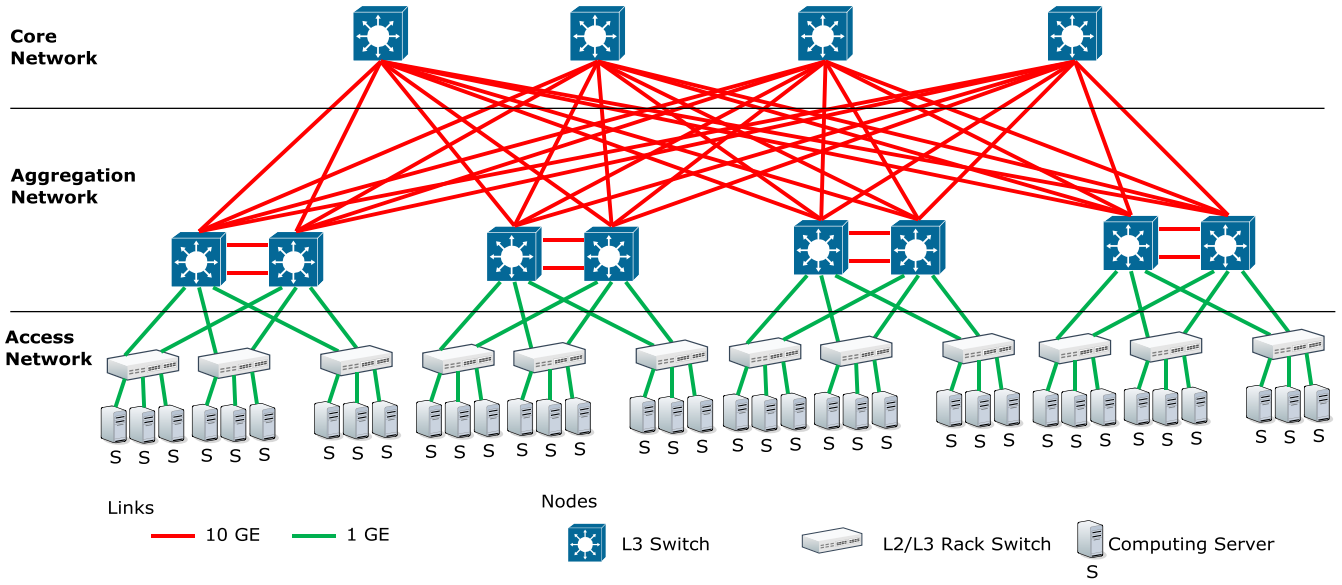

Data center architectures

- Present: Three-tier High-Speed architecture
  - Increased core network bandwidth
  - 2-way ECMP load balancing
  - 100 GE standard (IEEE 802.3ba) approved in June 2010

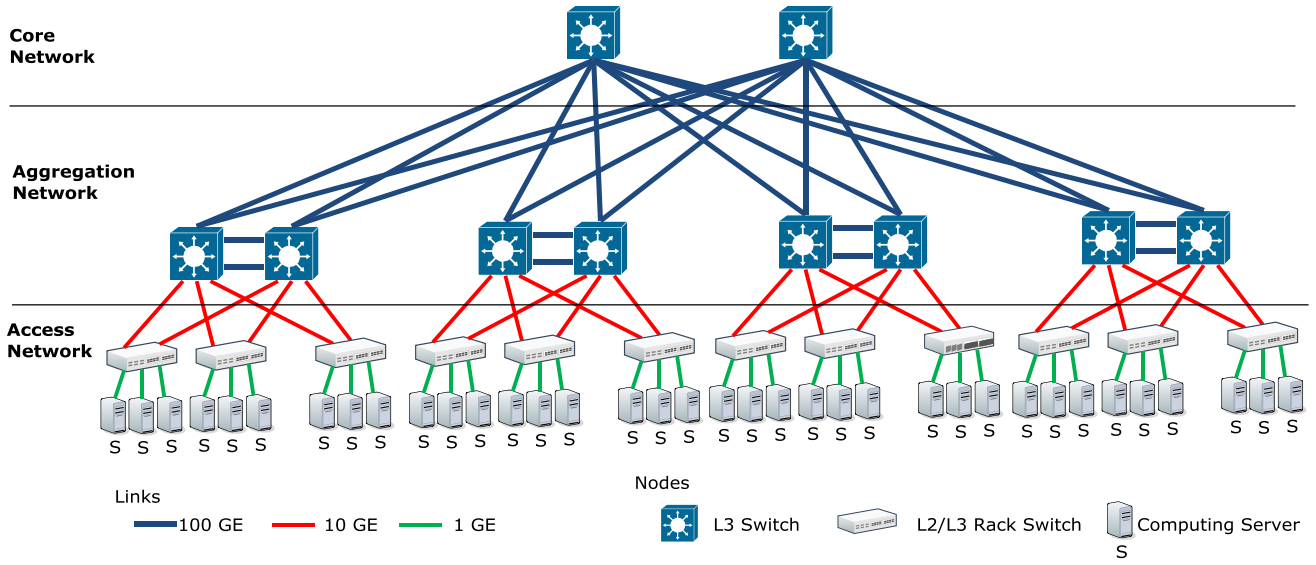

Data center components

- Servers' Energy Model

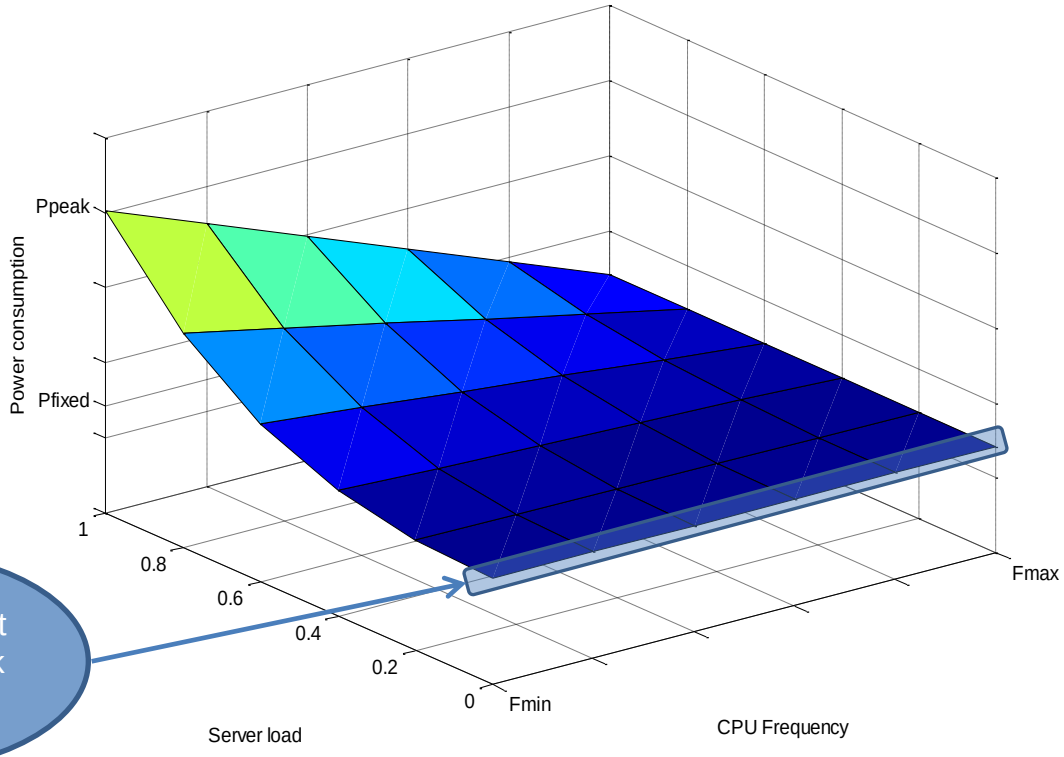

$$P = P_{fixed} + P_f * f^3$$

memory modules, disks, I/O resources

CPU

Idle server consumes about 66% of the peak load for all CPU frequencies

Data center components

- Switches
  - Most common Top-of-Rack (ToR) switches typically operate at Layer-2 interconnecting gigabit links in the access network
  - Aggregation and core networks host Layer-3 switches operating at 10 GE (or 100 GE)

- Links
  - Transceivers' power consumption depends on the quality of signal transmission in cables and is proportional to their cost
  - 1 GE links consume 0.4W for 100 meter transmissions over twisted pair
  - 10 GE links consume 1W for 300 meter transmission over optical fiber

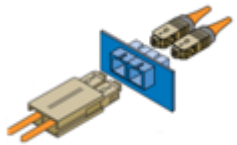

- Supported power management modes
  - Dynamic voltage scaling, dynamic shutdown, or both

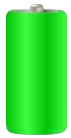

Simulator components

- Switches' Energy Model

$$P_{switch} = P_{chasis} + n_{linecards} \cdot P_{linecard} + \sum_{i=0}^{R} n_{ports.r} \cdot P_r$$

Chassis
~ 36%

Linecards
~ 53%

Port transceivers
~ 11%

P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, "A Power Benchmarking Framework for Network Devices," 8th international IFIP-TC 6 Networking Conference, Aachen, Germany, May 11 - 15, 2009.

DENS methodology

- DENS achieves balance between
  - Energy consumed by the data center
  - Individual job performances and their QoS requirements
  - Data center traffic demands

- DENS is architecture specific

$$M = \alpha \cdot f_s + \beta \cdot f_r + \gamma \cdot f_m$$

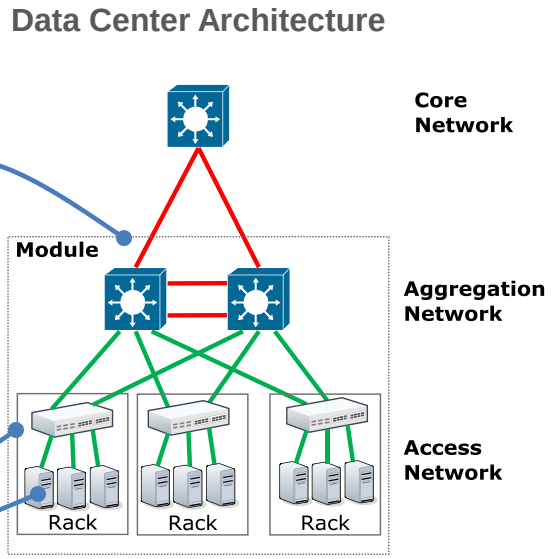

DENS methodology

- Computing server selection

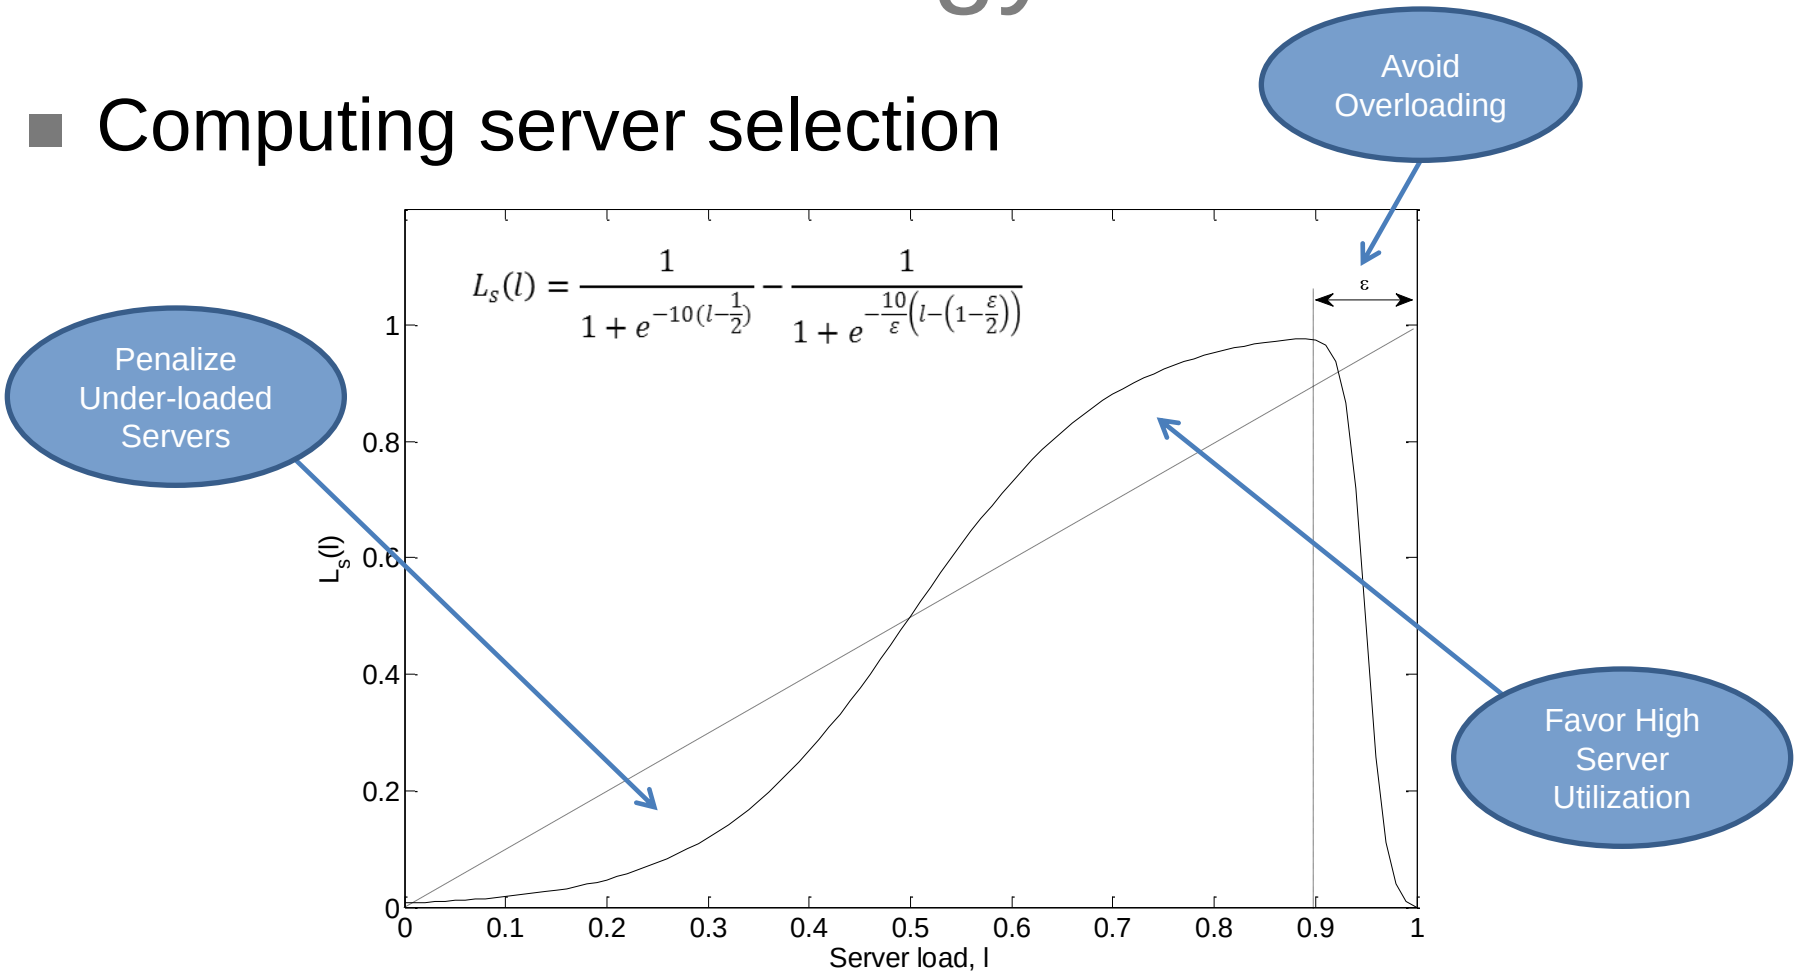

DENS methodology

- Computing server selection

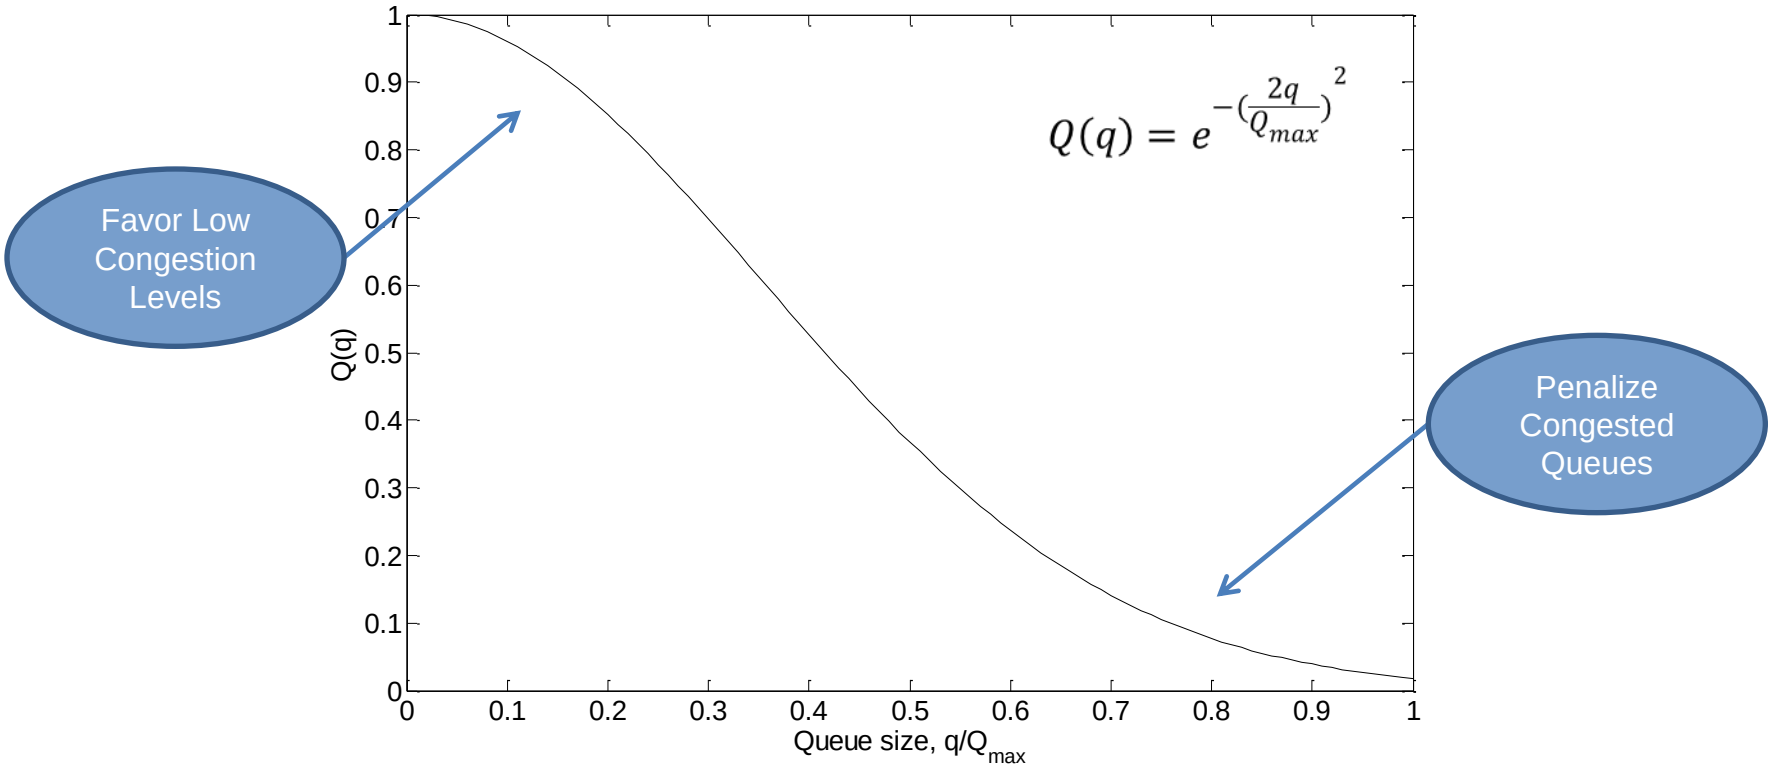

DENS methodology

- Computing server selection

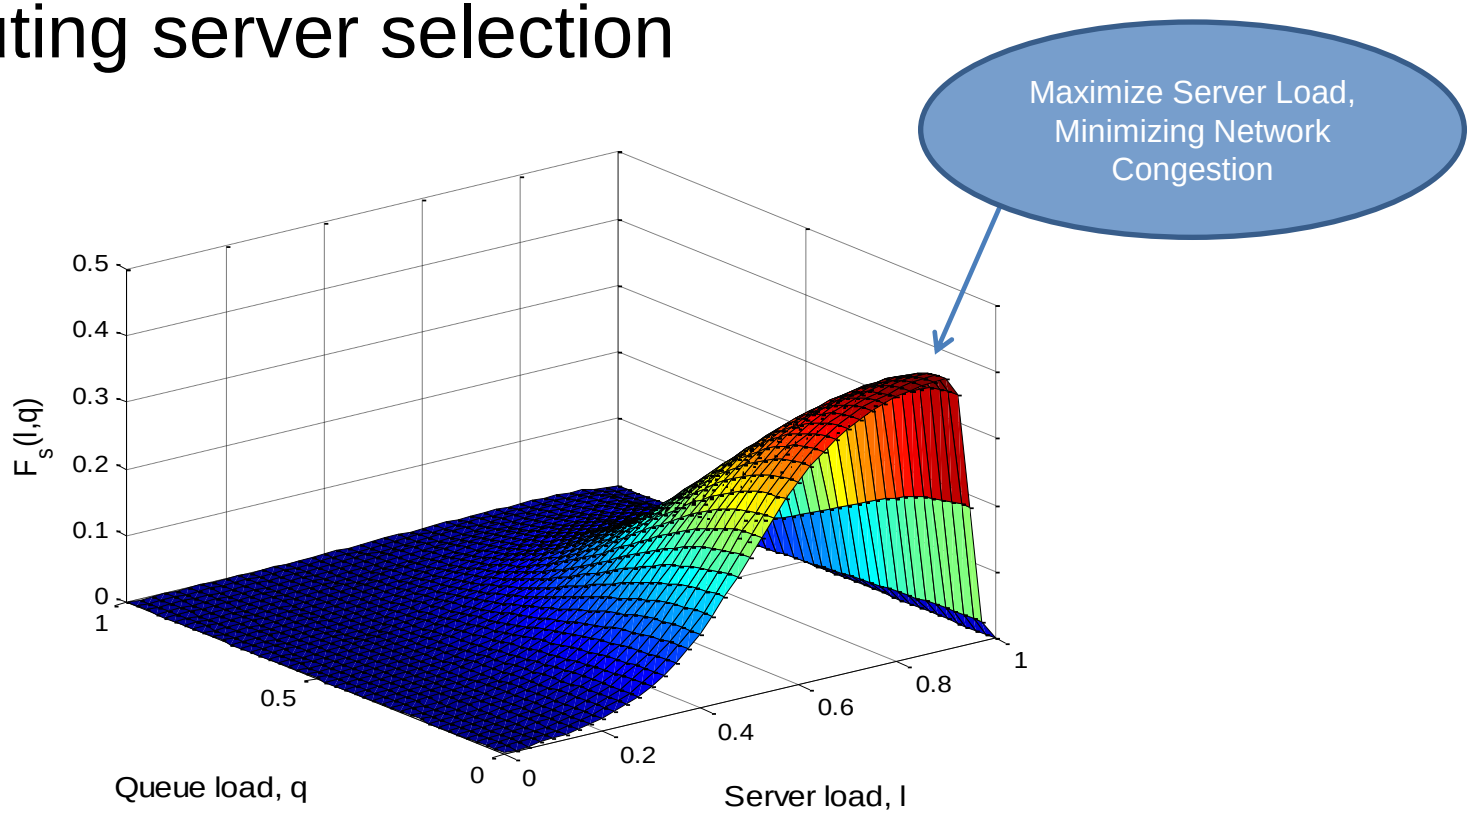

Performance Evaluation

- GreenCloud simulator is developed
- Three-tier data center topology
  - 1536 nodes, 32 racks, 4 core and 8 aggregation switches

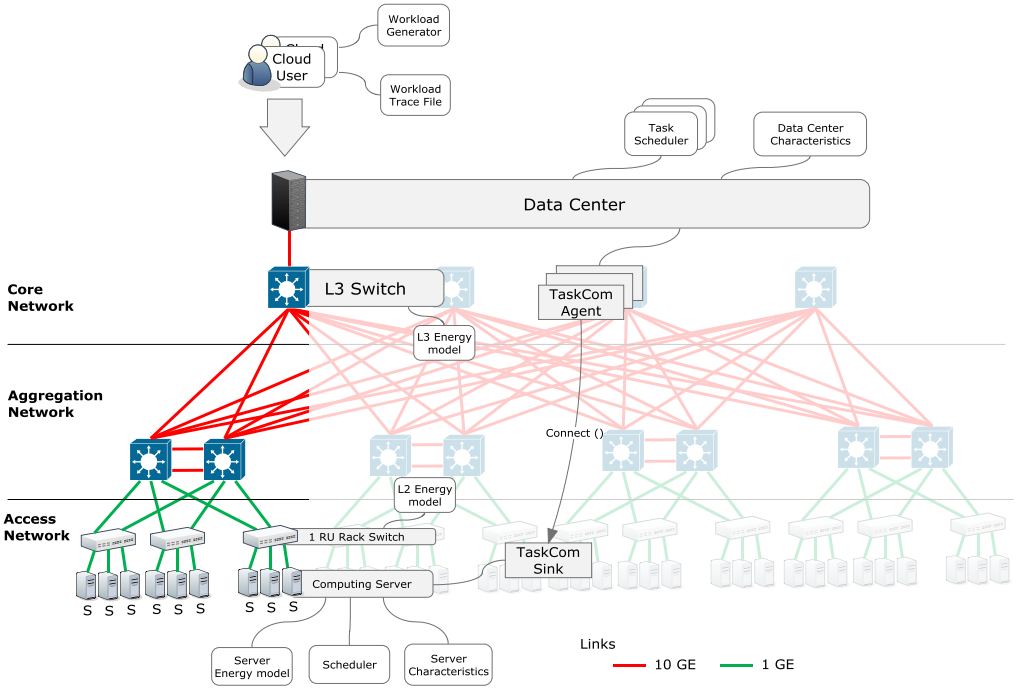

Performance Evaluation

- Server workload distribution

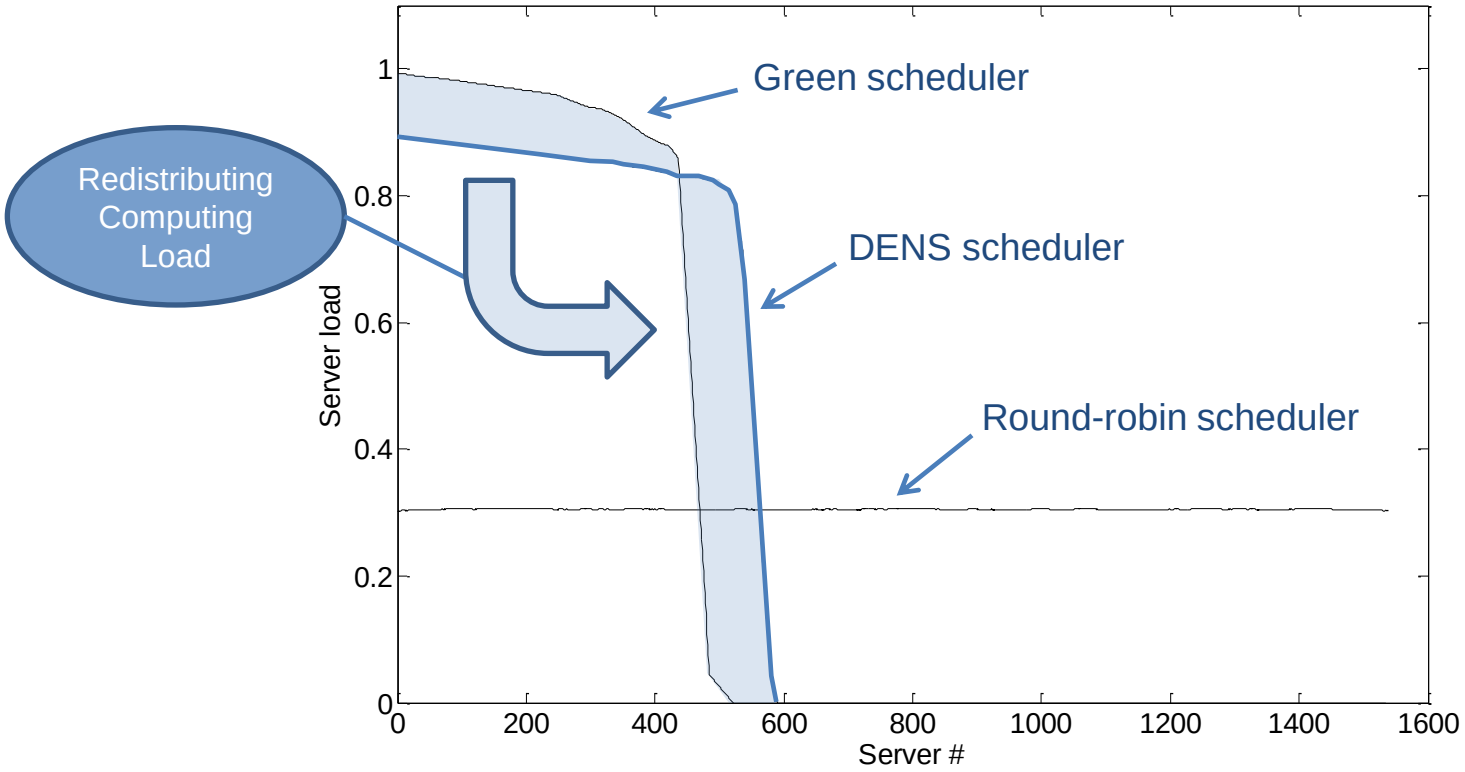

Performance Evaluation

- Network workload distribution

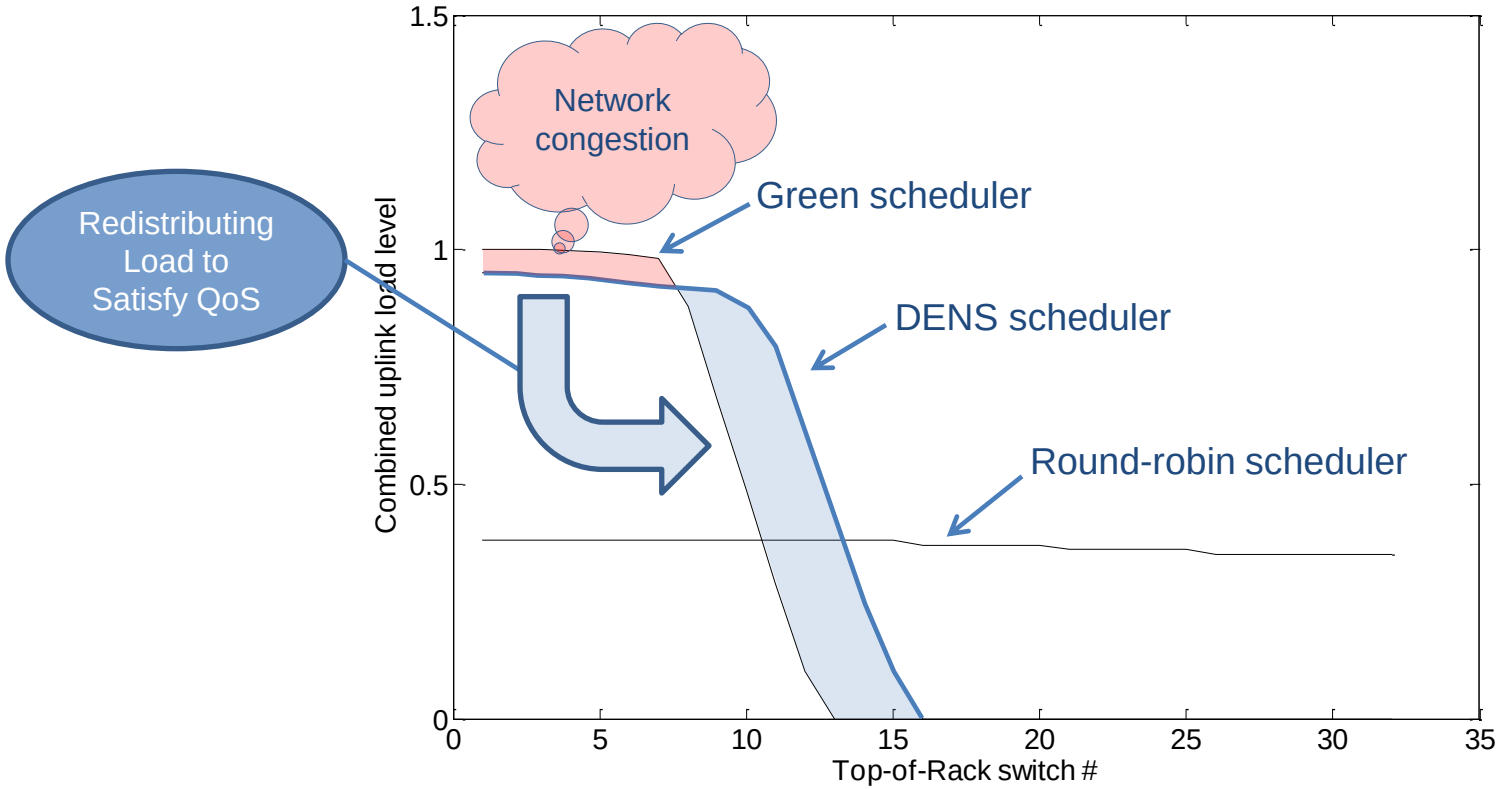

Performance Evaluation

- Top-of-Rack (ToR) switch load

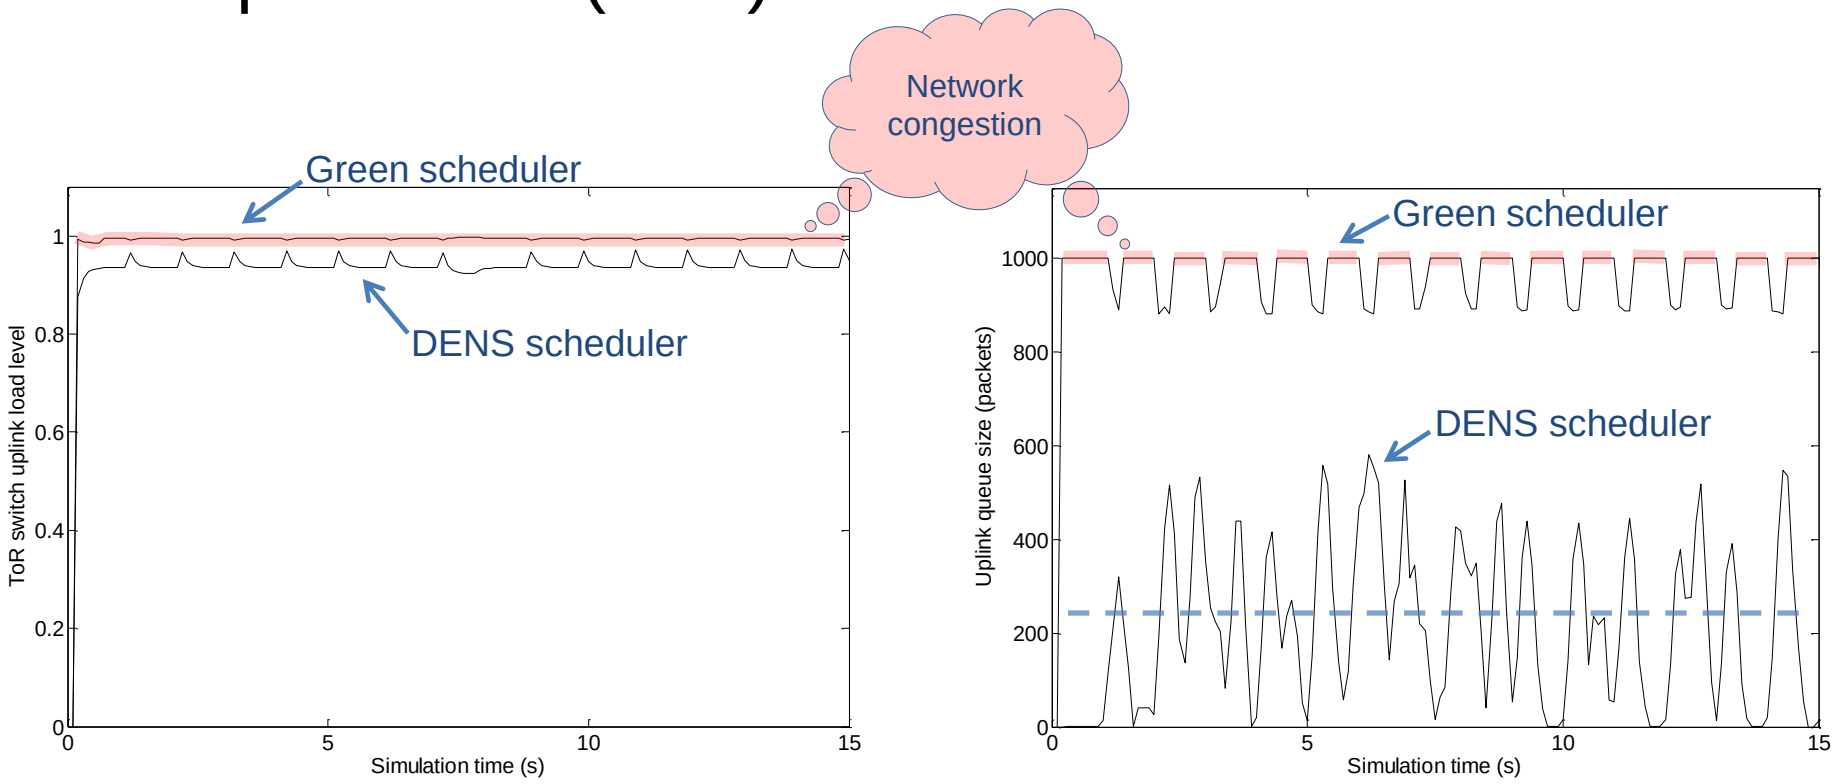

Performance Evaluation

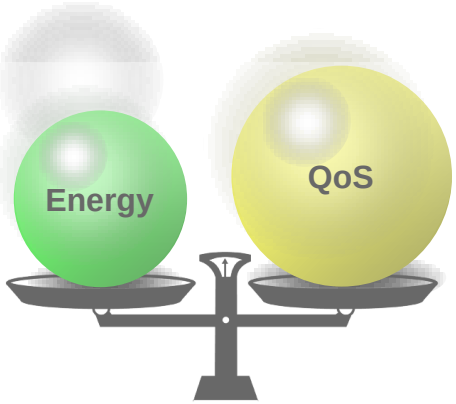

- Data center energy consumption

Parameter	Power Consumption (kW·h)		
	Round Robin scheduler	Green Scheduler	DENS scheduler
Data center	417.5K	203.3K (48%)	212.1K (50%)
Servers	353.7K	161.8K (45%)	168.2K (47%)
Network switches	63.8K	41.5K (65%)	43.9K (68%)

Conclusions

- We acknowledge

  - Funding form Luxembourg FNR in the framework of GreenIT project

  - Research fellowship provided by the European Research Consortium for Informatics and Mathematics (ERCIM)

References

- GreenCloud simulator
  - In Journal of Supercomputing, special issue on Green Networks
    "GreenCloud: A Packet-level Simulator of Energy-aware Cloud Computing Data Centers"
  - In IEEE Global Communications Conference (GLOBECOM)
    "GreenCloud: A Packet-level Simulator of Energy-aware Cloud Computing Data Centers"

GreenCloud
available at
http://greencloud.gforge.uni.lu

- Energy-Efficient Network-Aware Scheduling
  - In Cluster Computing, special issue on Green Networks
    "DENS: Data Center Energy-Efficient Network-Aware Scheduling"
  - In IEEE/ACM GreenCom [Best paper award]
    "DENS: Data Center Energy-Efficient Network-Aware Scheduling"

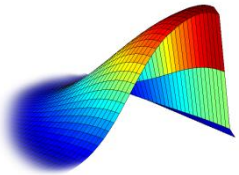

Thank you!