Introduction – Chris Develder

- PhD, Ghent University, 2003
  - "Design and analysis of optical packet switching networks"
- Professor at Ghent University since Oct. 2007
  - *Research Interests*: dimensioning, modeling and optimizing **optical** (grid/cloud) networks; **smart grids**; multimedia and home networks; **information retrieval**
  - Visiting researcher at UC Davis, CA, USA, Jul-Oct. 2007 (optical grids)
  - Visiting researcher at Columbia Univ., NY, USA, 2013-15 (IR/IE)
- Industry Experience: network planning/design tools
  - OPNET Technologies (now part of Riverbed), 2004-05
- More info: http://users.atlantis.ugent.be/cdvelder

UNIVERSITEIT GENT

FACULTY OF ENGINEERING AND ARCHITECTURE

Dimensioning (optical) networks for cloud computing

Chris Develder, *et al.*

Ghent University – iMinds
Dept. of Information Technology – IBCN

Networking for big data applications

Optical networks crucial for increasingly demanding cloud services, e.g.,

- Computing:
  - High energy physics
  - Amazon EC2, Microsoft Azure
- Online storage:
  - Dropbox, Google Drive, etc.
- Collaboration tools:
  - MSOffice 365, Google docs
- Video streaming:
  - Netflix, YouTube

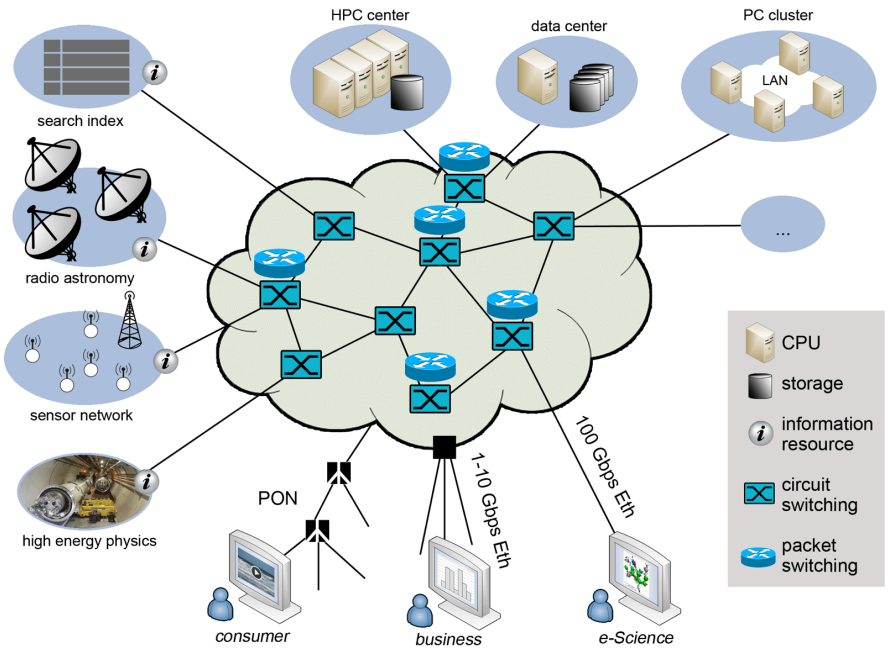

C. Develder, et al., *"Optical networks for grid and cloud computing applications"*, Proc. IEEE, Vol. 100, No. 5, May 2012, pp. 1149-1167.

Networking for big data applications

Optical networks crucial for increasingly demanding cloud services, e.g.,

- Computing:
  - High energy physics
  - Amazon EC2, Microsoft Azure
- Online storage:
  - Dropbox, Google Drive, etc.
- Collaboration tools:
  - MSOffice 365, Google docs
- Video streaming:
  - Netflix, YouTube

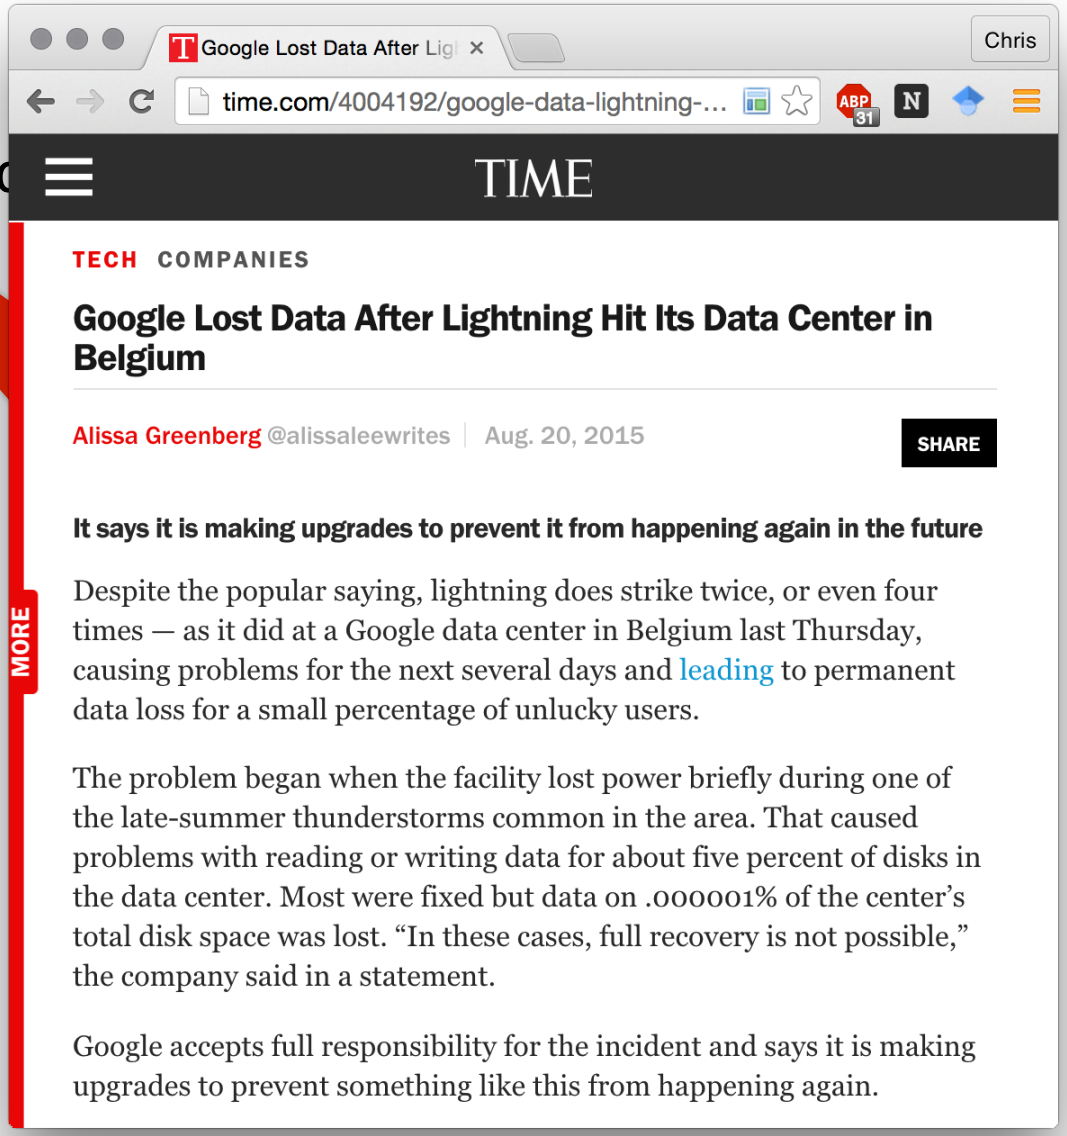

Dimensioning networks for multi-site Data Centers

Given:

- Cloud service requests
  (bandwidth + server capacity)
- Network topology
  (w/ candidate DC locations)

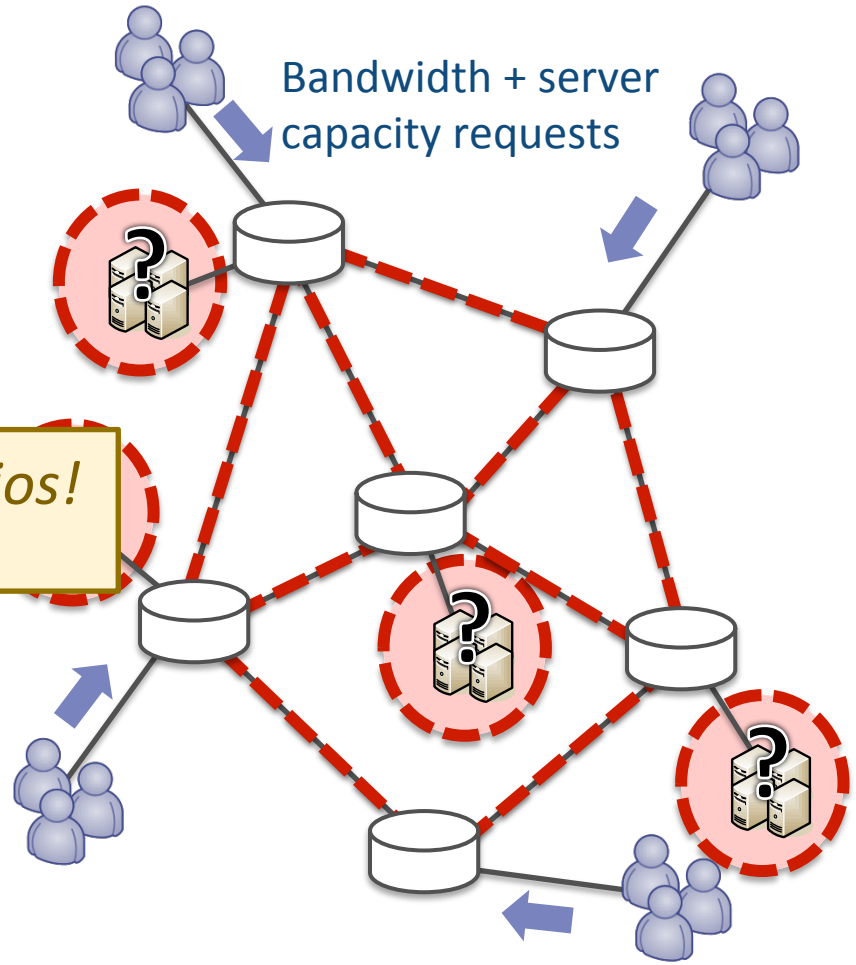

Also under failure scenarios!
(We'll assume shared protection)

Find:

- Minimal resource capacity to satisfy requests?
- Routes to follow for each request?
- How many DCs and where?

Dimensioning for clouds: What's different?

Anycast

Users do (in general) **NOT** care where applications are served

- E.g., virtual machines in IaaS can be instantiated anywhere
- E.g., bag-of-tasks computational jobs can be run at any server

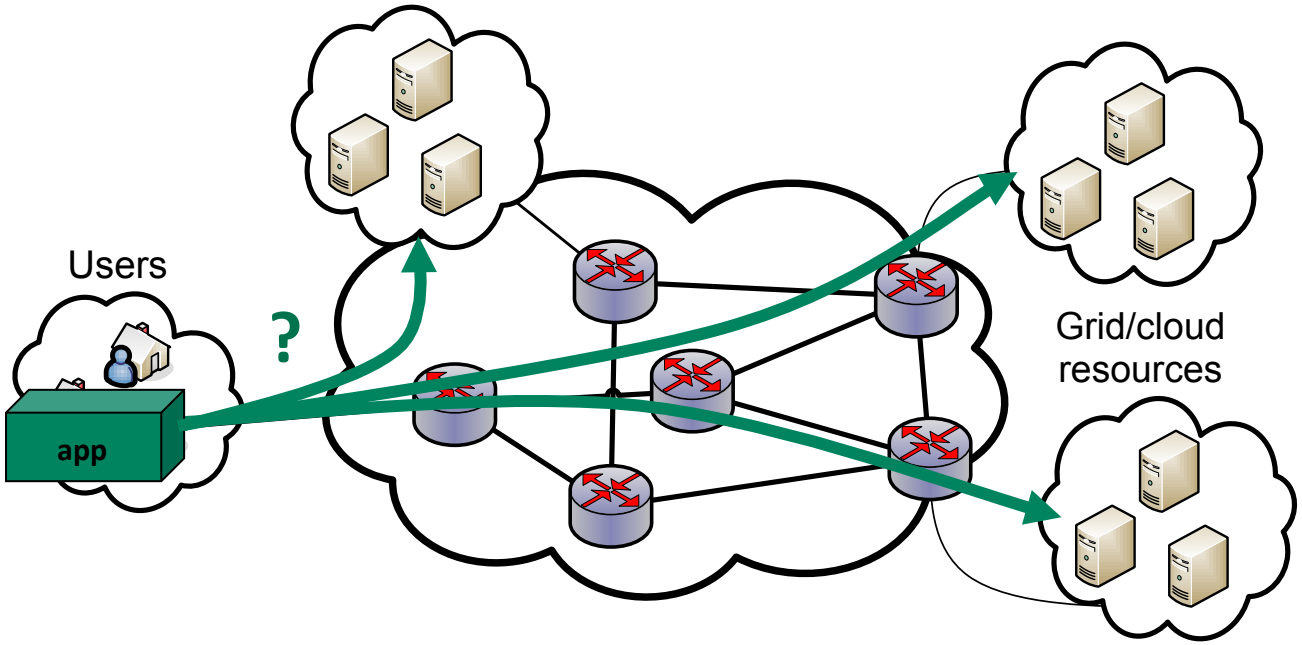

Network virtualization

Physical network is logically partitioned in isolated virtual networks

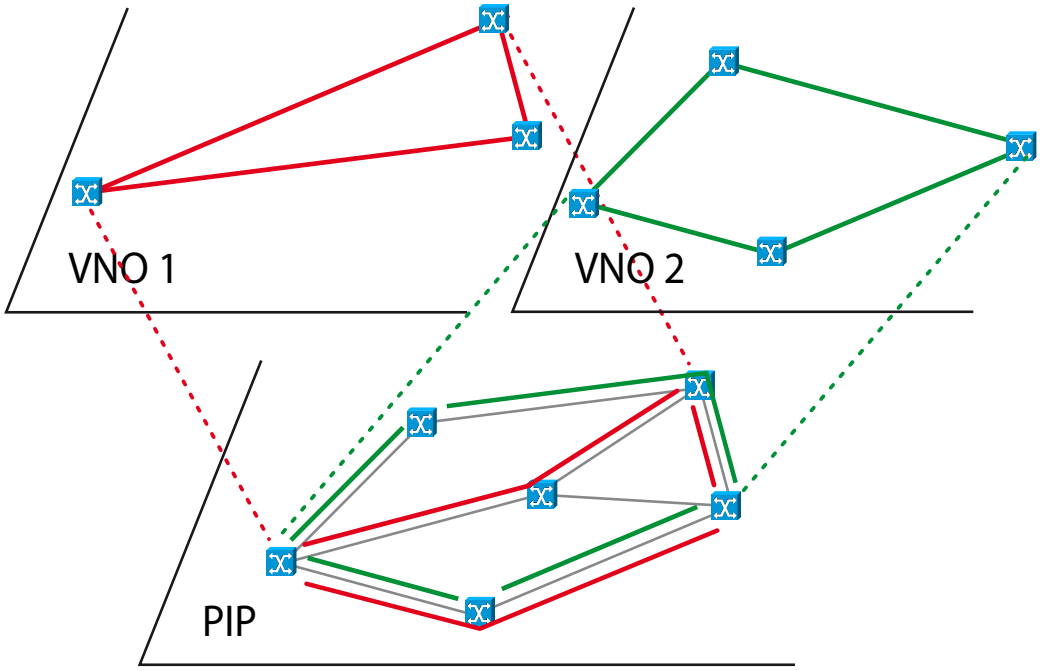

- **Virtual Network Operators (VNO)** operate logically separated networks
- **Physical Infrastructure Providers (PIP)** have full control over infrastructure (fibers, OXCs)

J.A. García-Espín, et al., *"Logical Infrastructure Composition Layer: the GEYSERS holistic approach for infrastructure virtualisation"*, in Proc. TERENA Networking Conference (TNC 2012), Reykjavík, Iceland, 21-24 May 2012.

Key questions?

Exploiting anycast to minimize capacity?

1. Does choice of anycast algorithm highly impact network bandwidth requirements?

2. What is benefit of relocating to alternate DC for resilience?

3. Under time-varying traffic, can changing (backup) routes save bandwidth?

(1) Impact of anycast routing on bandwidth req.

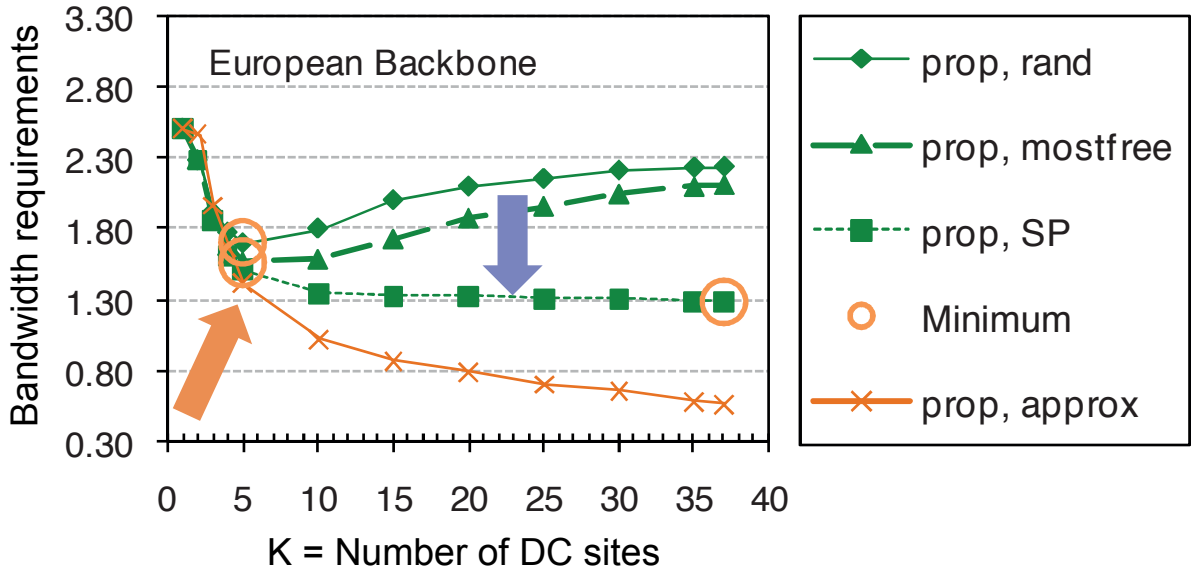

Impact of *# DC sites*:

Optimal value with minimal bandwidth, depends on the scheduling algorithm & server distribution

Impact of *scheduling*: (rand vs mostfree vs SP)

Nearest free server (SP) scheduling
→ min. bandwidth

Impact of *server capacity distribution*: (unif vs prop)

Smart, non-uniform server distribution (prop)
→ bandwidth reduction (compared to e.g., uniform)

(2) Relocation to maximally share resources

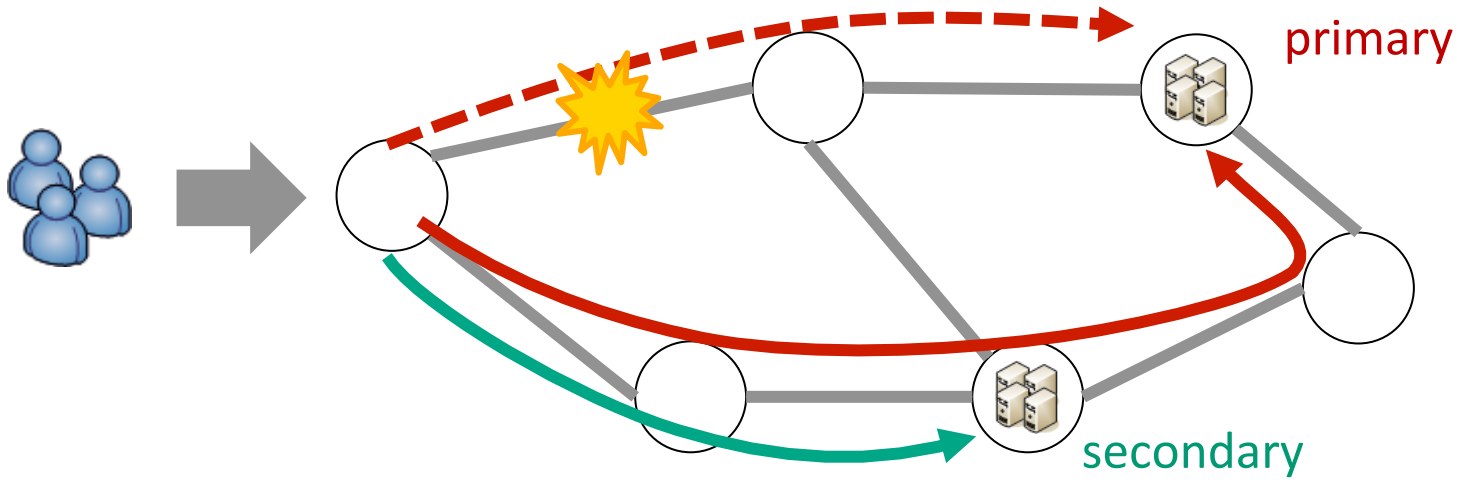

Intuition: save bandwidth by **relocating** to alternate DC for resilience

(2) Relocation to maximally share resources

Single **link** failures (**1L**):

- Reduction of backup wavelengths
- Slight increase in server capacity

Single **link/server** failure (**1LS**)

- Reduction of backup wavelengths
- Fewer servers than 1:N server protection (N=1)

Total cost relative to the NR, 1L case

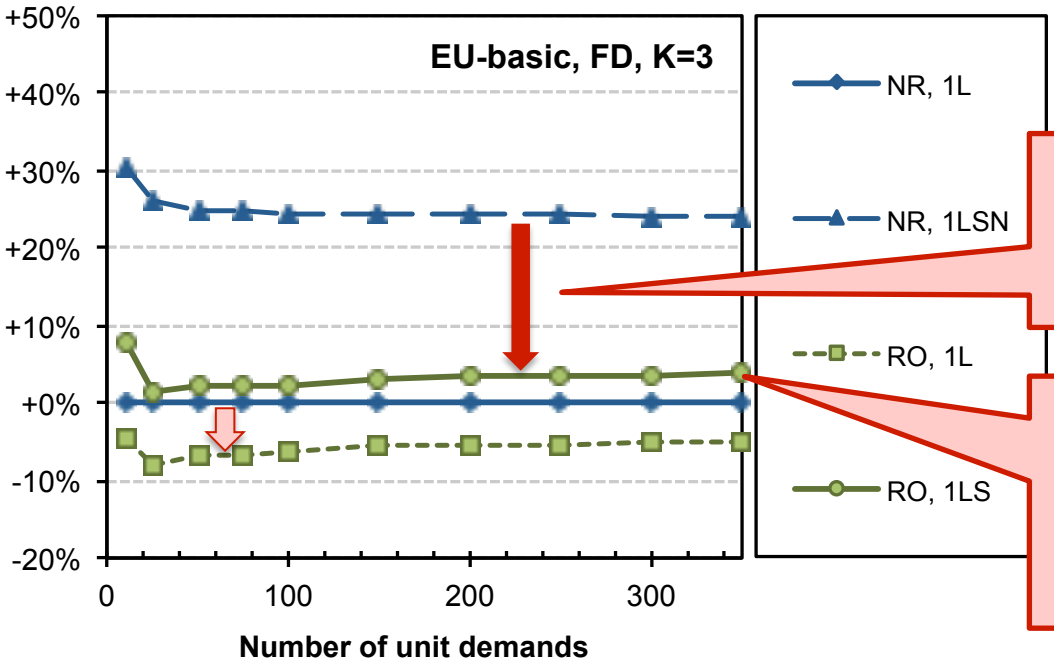

Mainly reduction of servers (1+1/K vs 1+1/N)

Protection against 1LS failures at almost same cost as 1L without relocation

(3) Changing routes for time-varying traffic

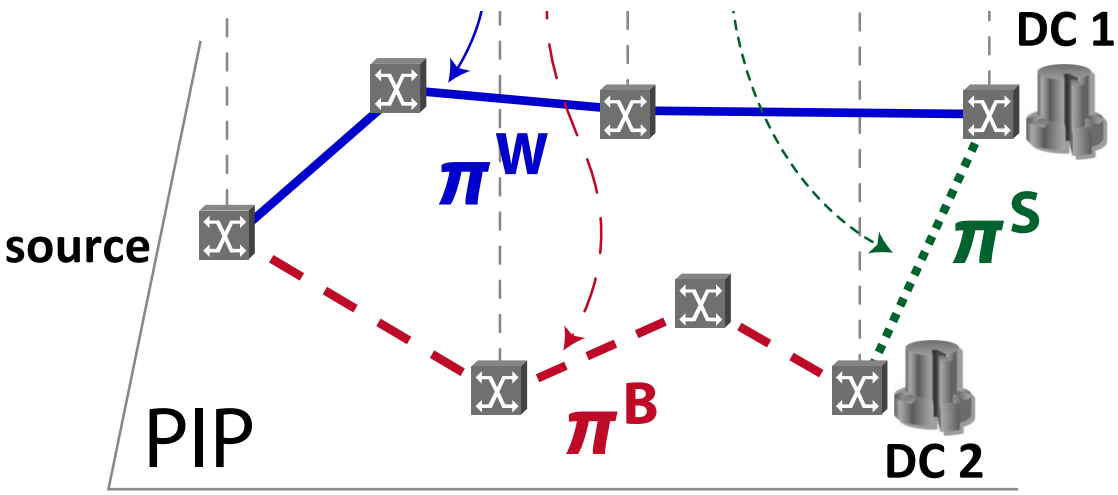

W = working path (primary)
B = backup path (secondary)
S = synchronization path (sync)

Resilience scenarios:

- **Scenario I**: Do NOT change
- **Scenario II**: May change backup & synchronization paths
- **Scenario III**: May change all

Intuition: bandwidth saving mainly by changing secondary DC and thus backup & sync. paths

(3) Changing routes for time-varying traffic

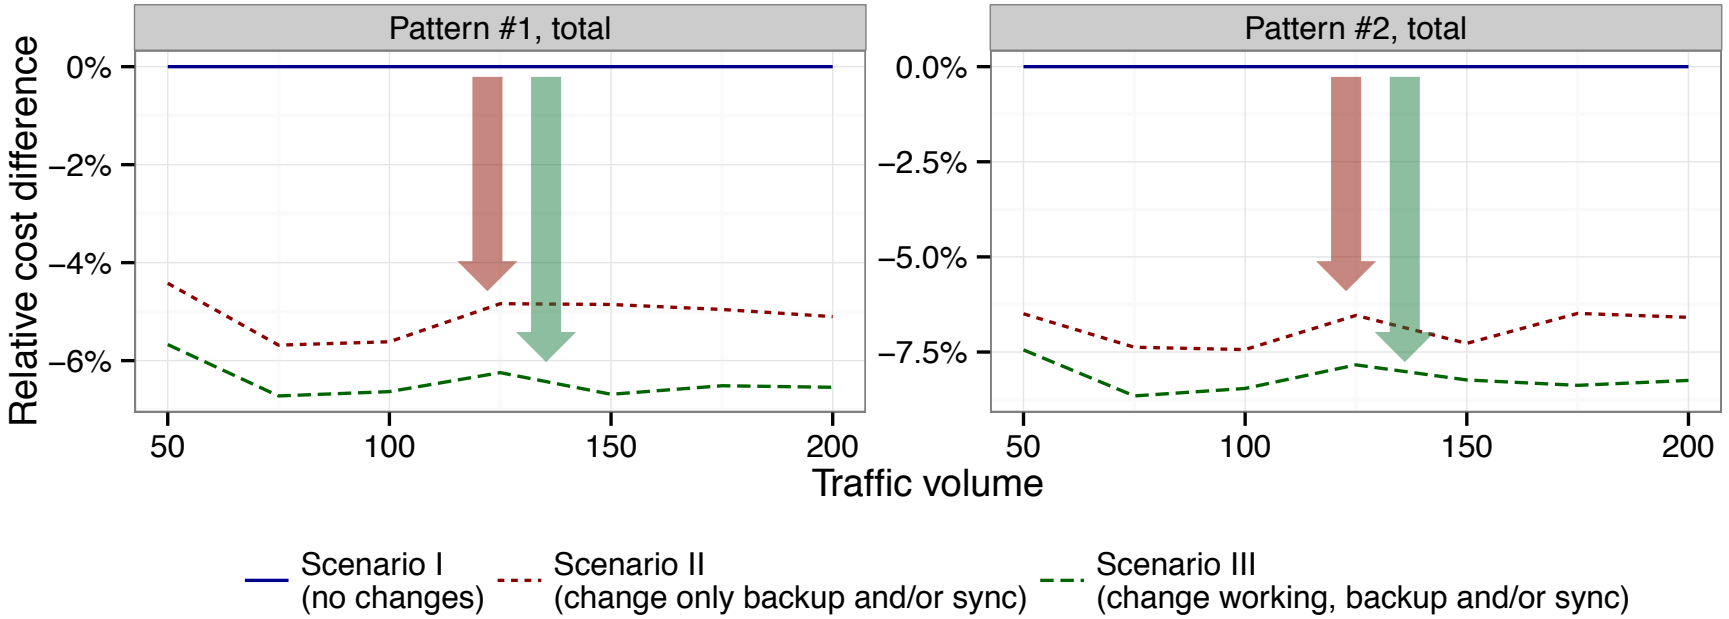

- Total cost savings up to almost 8% (pattern #2, i.e., more multi-period traffic)
- Savings mainly by **sharing of backup** (backup savings up to 14%)
- Saving by only changing backup/synchronization (Scenario II) almost as good as when also changing working (Scenario III)

Wrap-up

- Cloud computing: anycast routing is key difference
- Bandwidth requirements can be minimized by exploiting freedom offered by anycast:
  1. Choice of destination will impact network capacity requirements
  2. Relocation to alternate DC for resilience allows overall bandwidth savings
  3. Changing backup DC for time-varying traffic allows bandwidth savings
- Future work: truly scalable algorithms, implementation through software defined networking, …?

C. Develder, B. Mukherjee, B. Dhoedt and P. Demeester, ***"On dimensioning optical grids and the impact of scheduling"***, Photonic Netw. Commun., Vol. 17, No. 3, Jun. 2009, pp. 255-265.

C. Develder, J. Buysse, B. Dhoedt and B. Jaumard, ***"Joint dimensioning of server and network infrastructure for resilient optical grids/clouds"***, IEEE/ACM Trans. Netw., Vol. 22, No. 5, Oct. 2014, pp. 1591-1606.

T. Wang, B. Jaumard, C. Develder, ***"Network mapping for resilient multi-site data centers"***, Proc. IEEE Int. Conf. Advanced Netw. and Telecommun. Sys. (ANTS 2015), Kolkata, India, 15-18 Dec. 2015.

Thank you … any questions?

Chris Develder

chris.develder@intec.ugent.be

Ghent University – iMinds