Designing resilient optical grids/clouds

Chris Develder*, Jens Buysse*, Marc De Leenheer*, Bart Dhoedt*, Brigitte Jaumard†

*: Ghent University – IBBT, INTEC – IBCN, Ghent, Belgium

†: CSE Dept., Concordia University, Montreal, Canada

Agenda

1. Introduction
2. Problem statement, Model & ILP solutions
3. Case study
4. Conclusions

Why optical grids/clouds? (1)

- eScience:
  - By 2015 it is estimated that **particle physicists** will require exabytes (10^{18}) of storage and **petaflops** (10^{15}) per second of computation
  - CERN's LHC Computing Grid (LGC), when fully operational will generate **15 petabytes** annually (that's ~2Gbit/s)

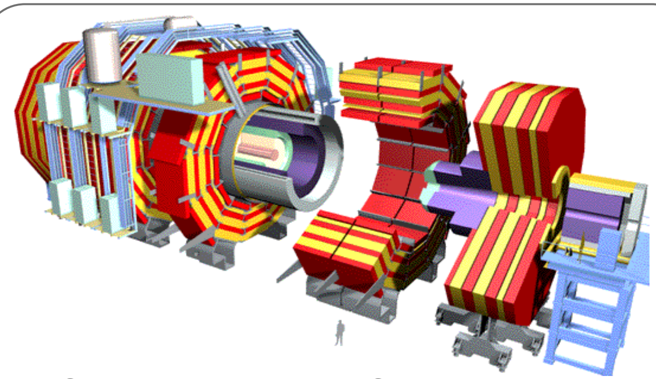

LHC = Large Hadron Collidor: particle accellerator

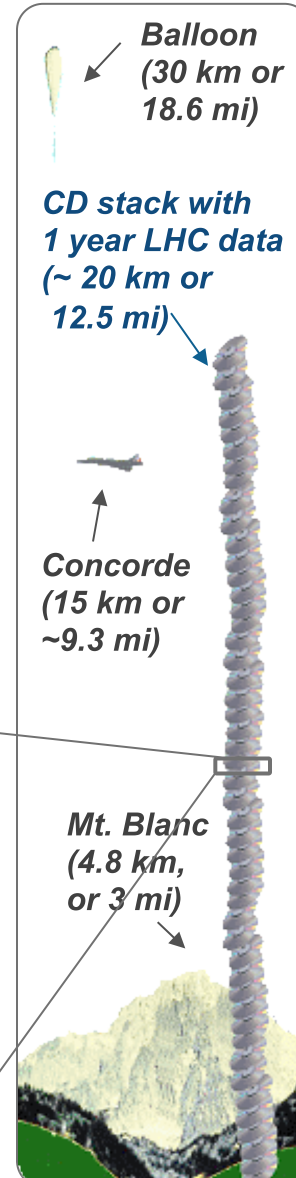

Why optical grids/clouds? (2)

- Consumer service:
  - Eg. **video editing**: 2Mpx/frame for HDTV, suppose effect requires 10 flops/px/frame, then evaluating 10 options for 10s clip is **50 Gflops** (today's high performance PC: <5 Gflops/s)

Online gaming:
e.g. Final Fantasy XI: 1.500.000 gamers

*Virtual reality: rendering of 3*10^8 polygons/s → 10^4 GFlops*

Multimedia editing

Optical grids/clouds

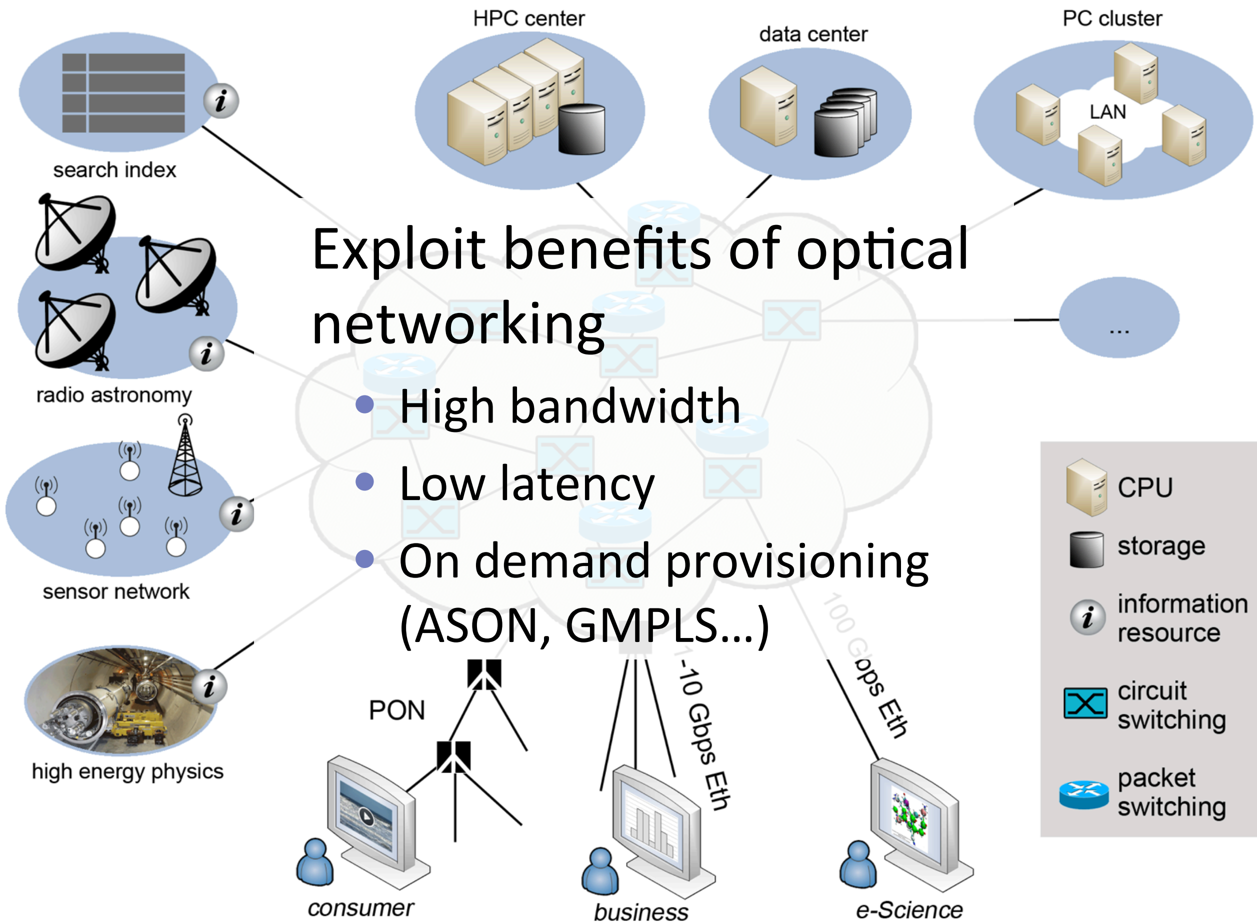

Exploit benefits of optical networking

- High bandwidth
- Low latency
- On demand provisioning (ASON, GMPLS…)

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

Anycast

- Users do (in general) NOT care on what location their applications are being served
  - E.g., virtual machines in IaaS can be instantiated anywhere
  - E.g., bag-of-tasks grid jobs can be run at any server

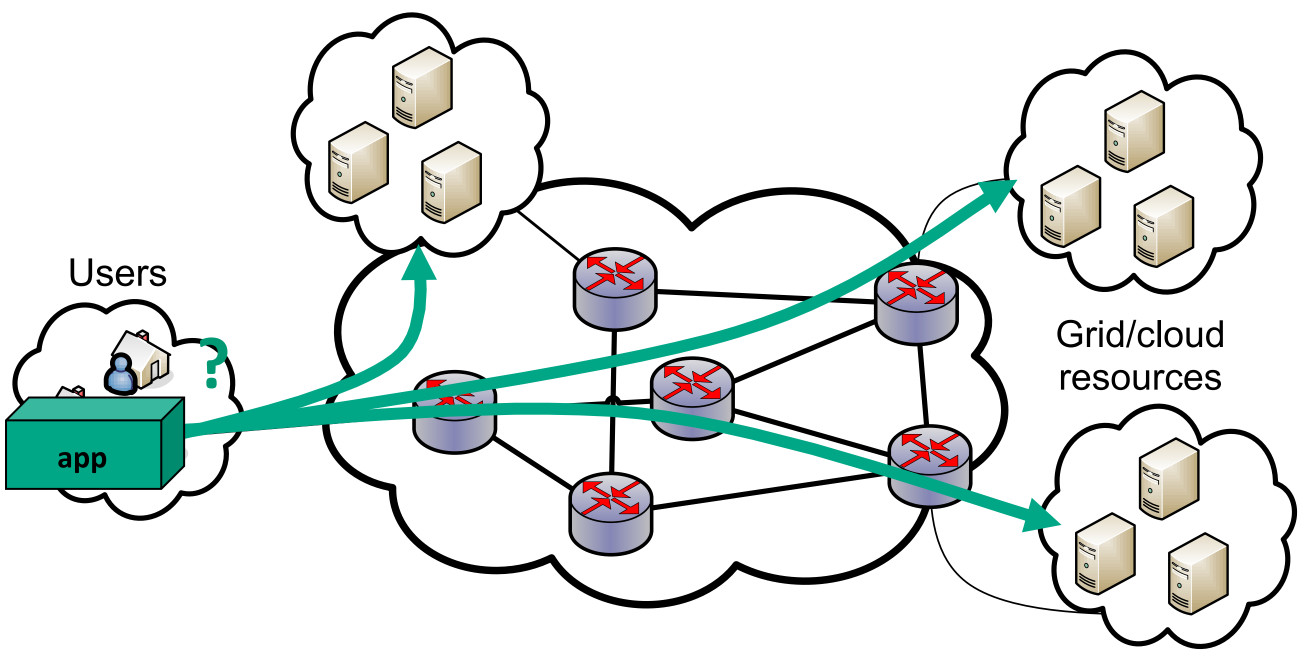

C. Develder, et al., "Designing resilient optical grids/clouds", PS 2012

Exploiting relocation

- Dimension optical grid/cloud so that it is resilient against failures
- Exploit anycast principle: allow rerouting jobs to other destinations

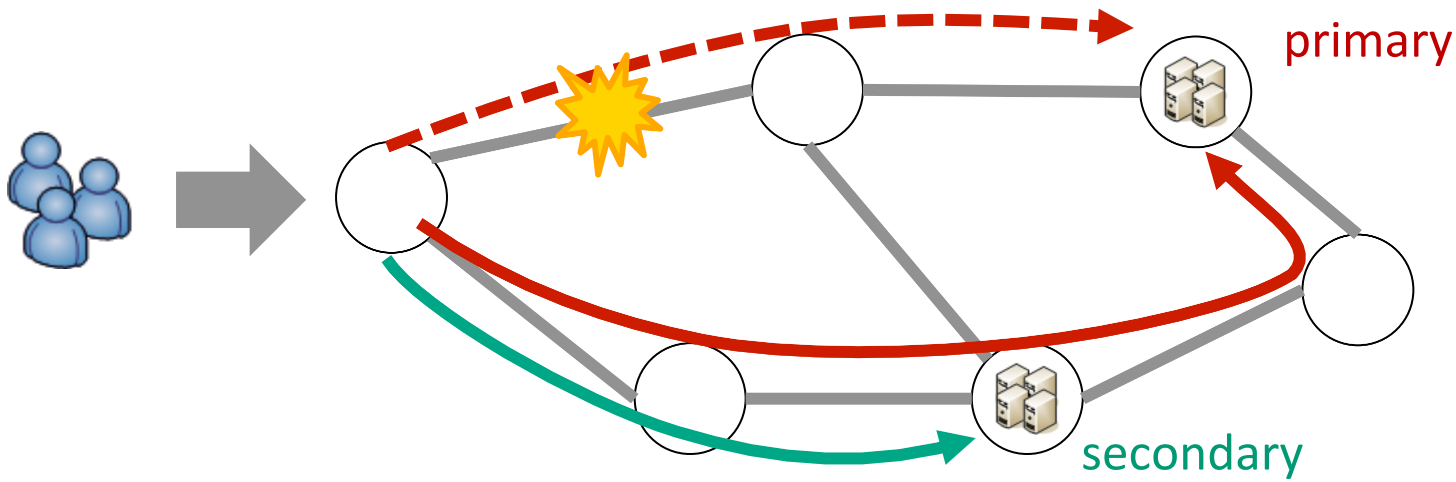

J. Buysse, M. De Leenheer, B. Dhoedt and C. Develder, *"Providing resiliency for optical grids by exploiting relocation: A dimensioning study based on ILP"*, Comput. Commun., Vol. 34, No. 12, Aug. 2011.

A. Shaikh, J. Buysse, B. Jaumard and C. Develder, *"Anycast routing for survivable optical grids: scalable solution methods and the impact of relocation"*, IEEE/OSA J. Opt. Commun. Netw., Vol. 3, No. 9, Sep. 2011.

Agenda

1. Introduction
2. Problem statement, Model & ILP solutions
3. Case study
4. Conclusions

Problem statement

Given

- Topology (sources, candidate data center locations, OXCs)
- Demand (for given sources)
- Survivability requirements (e.g. link and/or node failures)

Shared protection

Find

- K locations (chosen from candidate data center locations)
- Destination sites and routes
- Network and server capacity

Such that

- Network and server resources are minimized

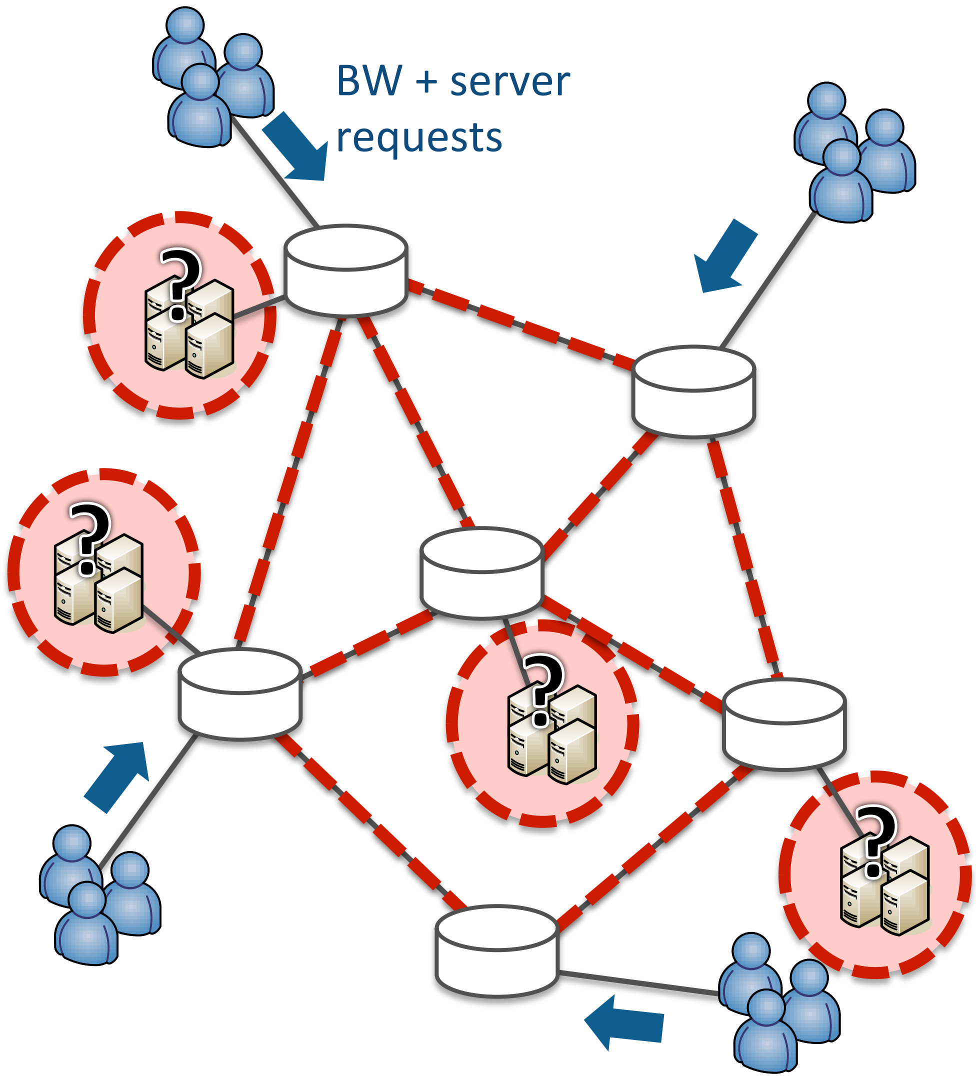

Solution approach

Step 1: Find the K best data center locations

Step 2: Find the primary/secondary destinations + paths towards them

Failure-Independent (FID) rerouting
=> Column generation

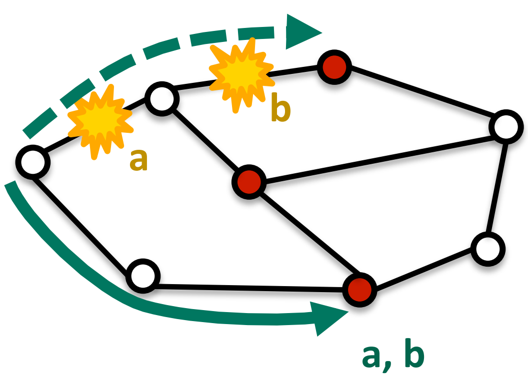

Failure-Dependent (FD) rerouting
=> Single ILP

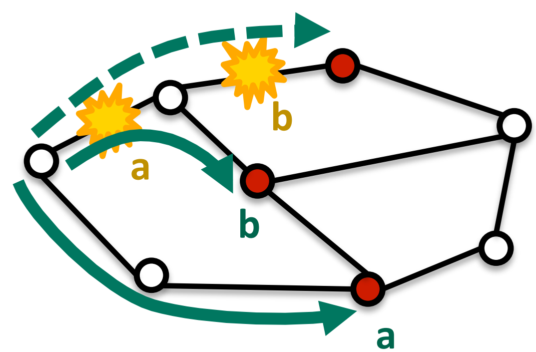

C. Develder, et al., "Designing resilient optical grids/clouds", PS 2012

Step 1: Finding the K "best" locations

- Binary variables:
  - t_v = 1 iff site v is server location
  - $f_{vv'}$ = 1 iff request from source v is directed to v'

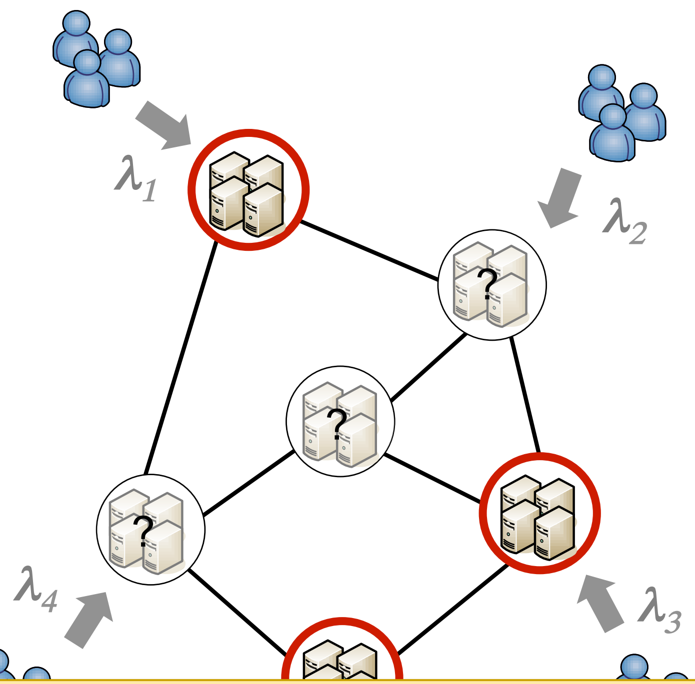

- Constants:
  - $h_{vv'}$ = cost for sending 1 unit request from source v to server site v'
  - Δ_v = number of unit requests from source v

$$\min \sum_{v} \sum_{v'} \Delta_v \cdot h_{vv'} \cdot f_{vv'}$$

subject to

$$\begin{cases} \sum_{v} t_v = K \\ \sum_{v'} f_{vv'} = 1 \quad \forall v \\ f_{vv'} \leq t_v \quad \forall v, v' \end{cases}$$

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

Step 2: Find destinations and routes towards them

- Network model:

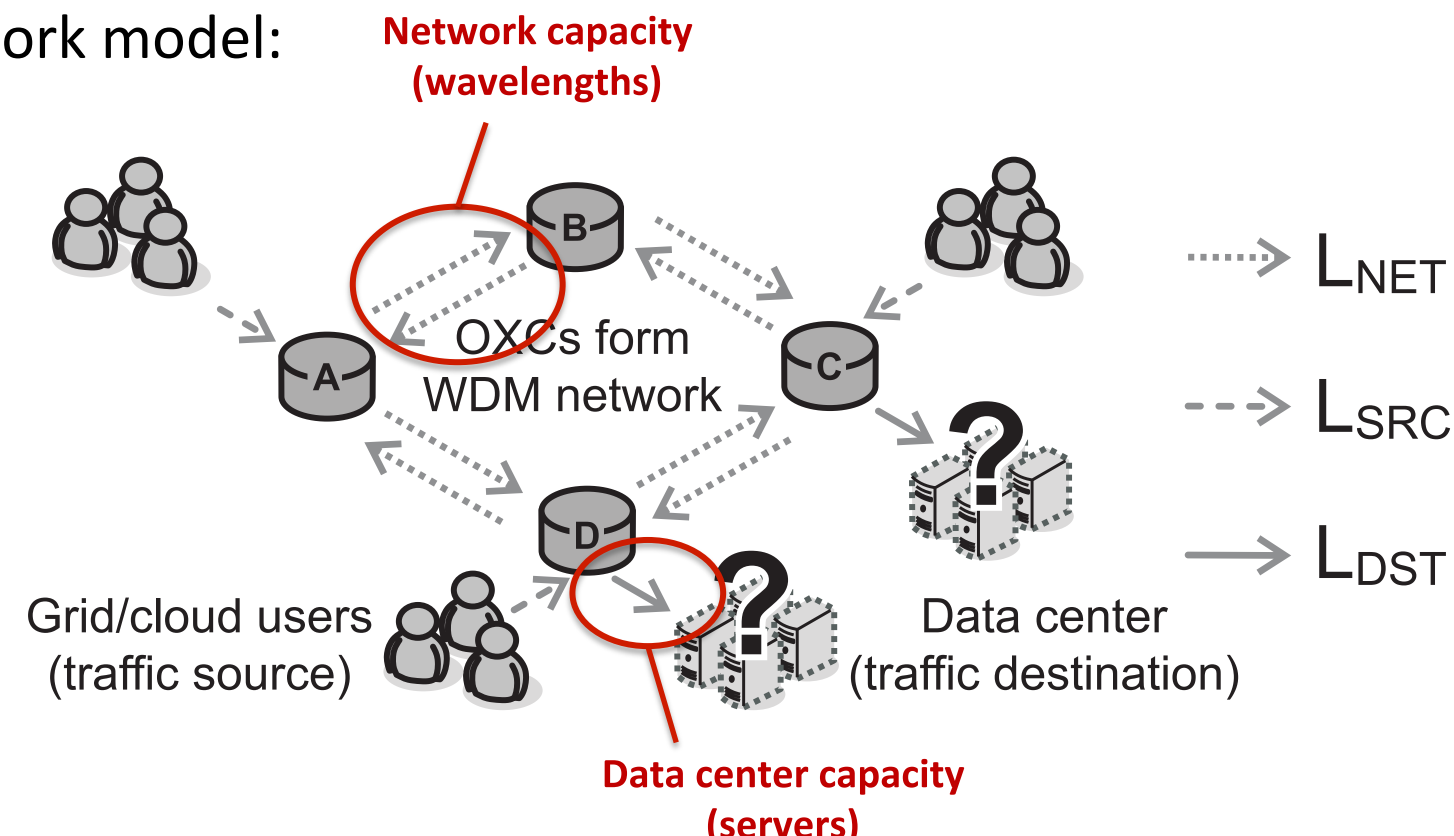

- w_l = "capacity" of link l
- Capacity = wavelengths for NET links, server for DST links!

Step 2: Find destinations and routes towards them

- Failure: modeled as SRLG = set of links that simultaneously fail
- Single link failure:

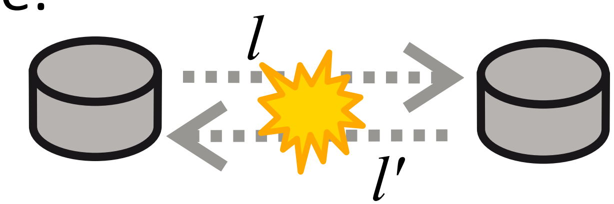

$$s = \{l, l'\}$$

- Single server failure: *1:N* protection [= add 1 for case single one out of N fails]
  - No relocation:
    - Let x = number of servers under working conditions
    - Then we need $\lceil (1+1/N)\cdot x \rceil$

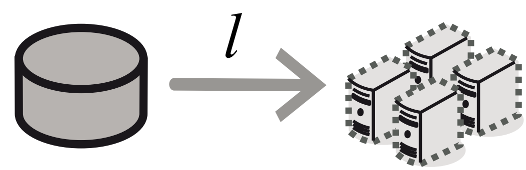

$$w_l \geq \rho_l \cdot x$$
$$\rho_l = 1 + 1/N$$

  - Relocation: consider $(1+N)$ parallel links, at most 1 fails

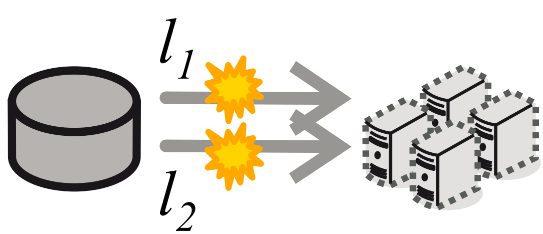

$$s_1 = \{l_1\}$$
$$s_2 = \{l_2\}$$

Step 2: Find destinations and routes towards them

- **Failure dependent (FD) rerouting** => Single ILP
- Variables:
  - p_{vls} : number of unit demands with source v that cross link l under failure s
  - w_l : capacity on link l
- Objective:

Network capacity (wavelengths) **Data center capacity (servers)**

$$\min \left(\sum_{\ell \in L_{\mathrm{NET}}} w_\ell + \alpha \sum_{\ell \in L_{\mathrm{DST}}} w_\ell \right)$$

- Constraints:
  - p_{vls} : flow constraints + don't use failing links when protecting against s
  - w_l : count capacity

1 for network link
1+1/N for server link, in case of no relocation (NR)

$$w_\ell \geq \rho_\ell \cdot \sum_{v \in V_{\mathrm{SRC}}} p_{v\ell s} \qquad \forall s \in S$$

Step 2: Find destinations and routes towards them

- **Failure-independent (FID) rerouting** => Column generation:
  - Assume: given "configurations" = combination of working and backup paths
  - Restricted Master Problem (RMP) finds best combination of configurations
  - Pricing Problem (PP) finds new configuration that can reduce cost

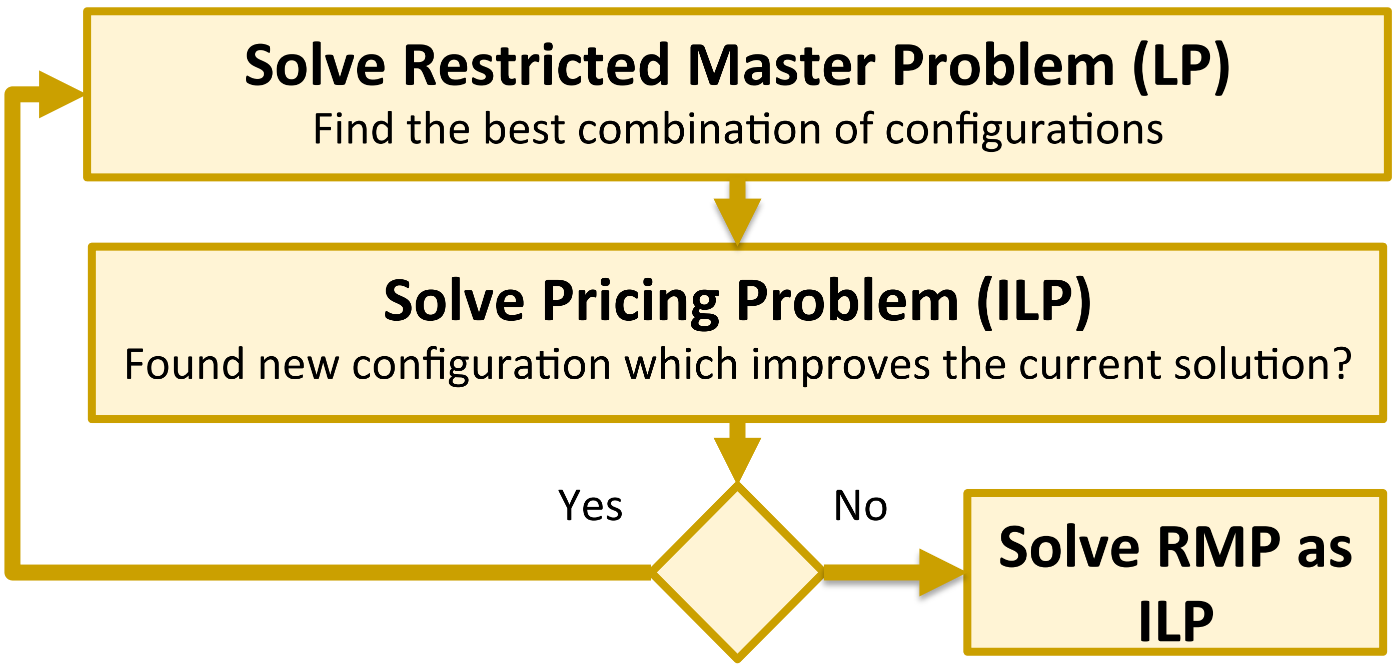

Agenda

1. Introduction
2. Problem statement, Model & ILP solutions
3. Case study
4. Conclusions

Case study set-up

- Topology
  - European network
  - 28 nodes and 41 bidir links

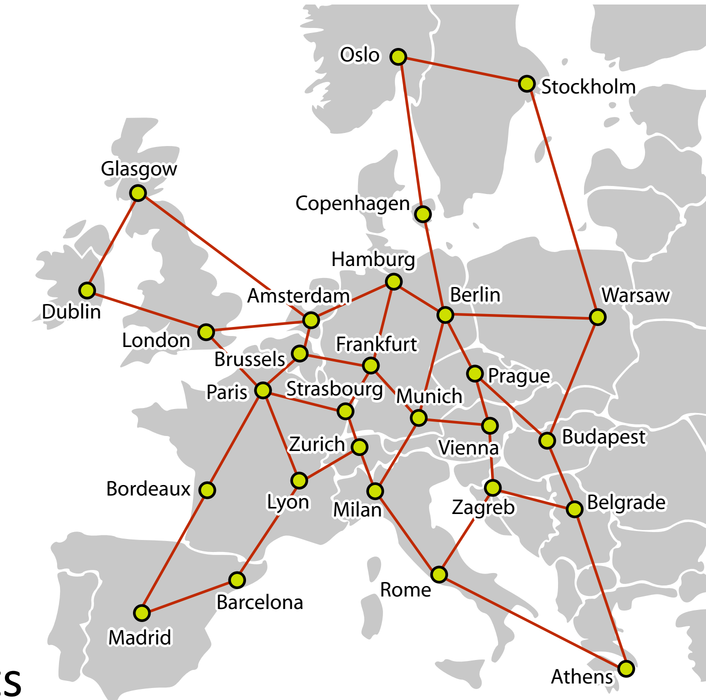

- Demand
  - Randomly generated requests (10-350)
  - 10 instances for each number of requests

- **Four scenario's:**

	No relocation		Exploiting relocation
Single link failures:	*1L, NoReloc*	⇔	*1L, Reloc*
Single failures of either link or server:	*1LSN, NoReloc*	⇔	*1LS, Reloc*

The impact of relocation

- 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

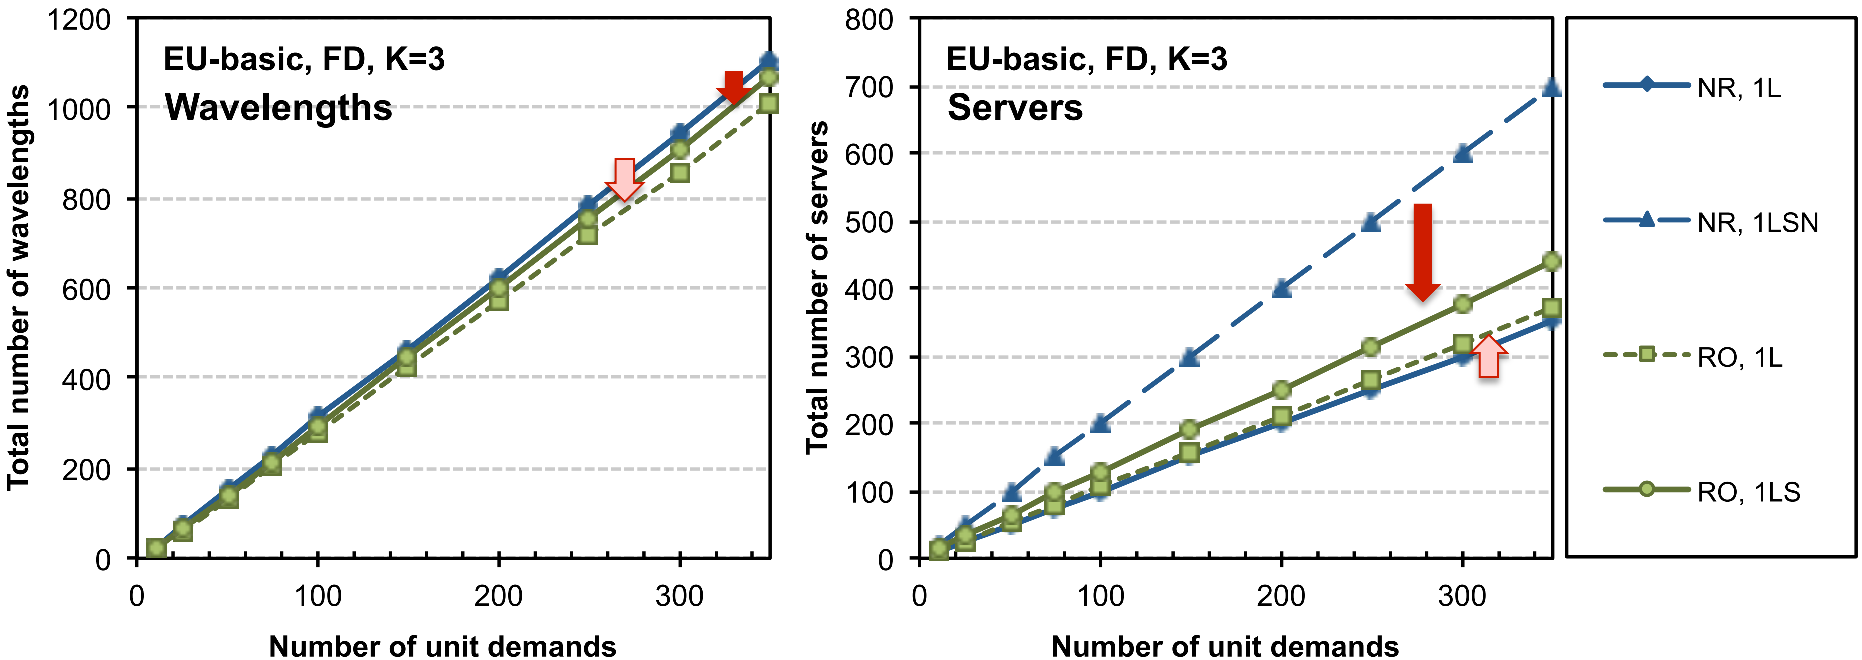

The impact of relocation

- 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

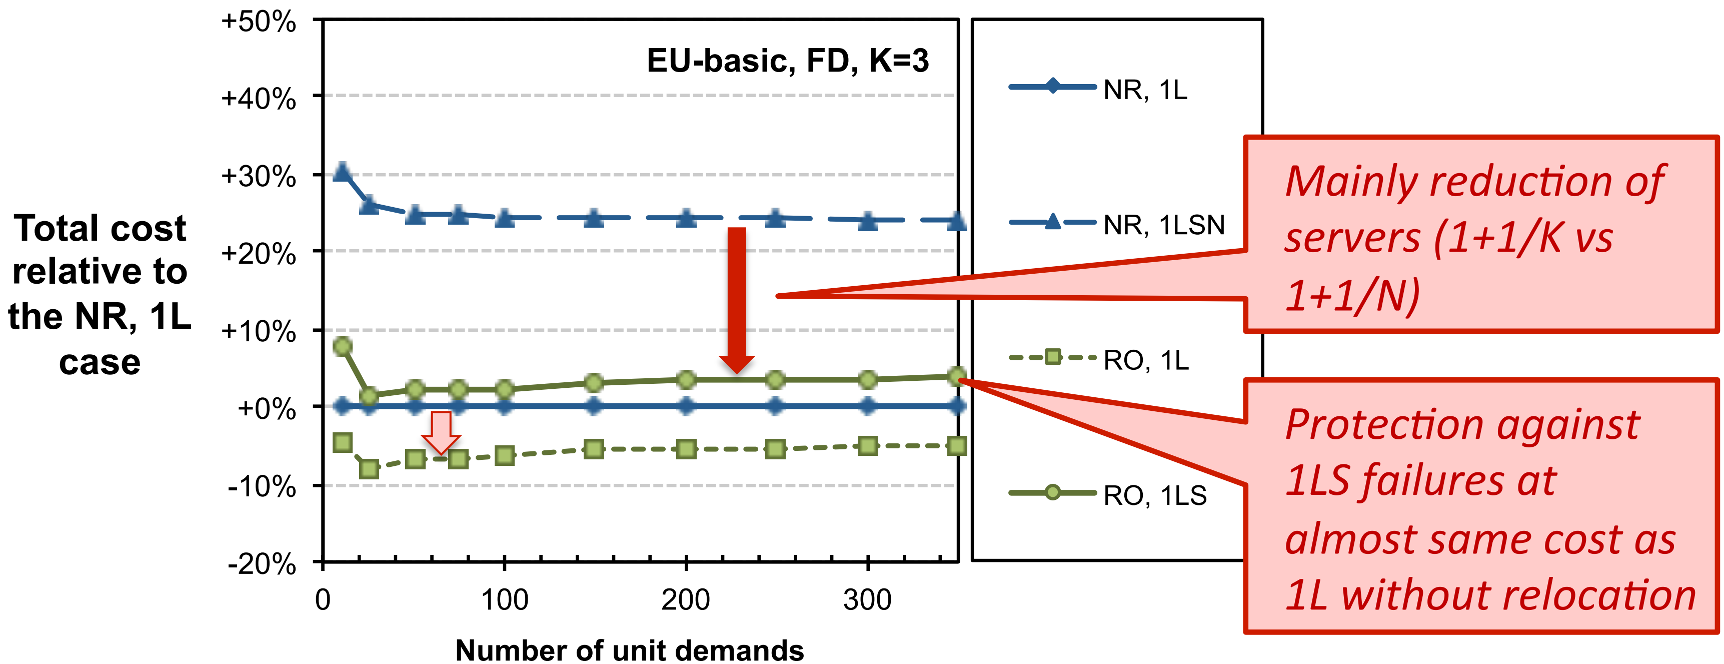

Failure dependent rerouting? (FD vs FID)

Failure-Independent (FID) rerouting

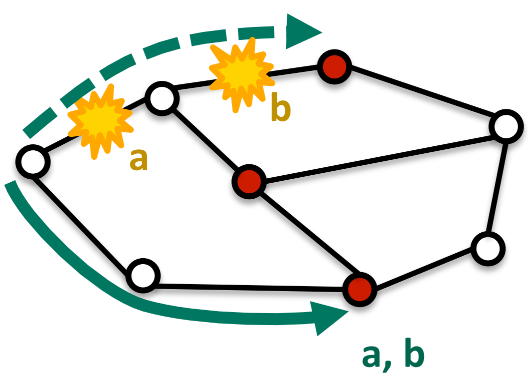

Failure-Dependent (FD) rerouting

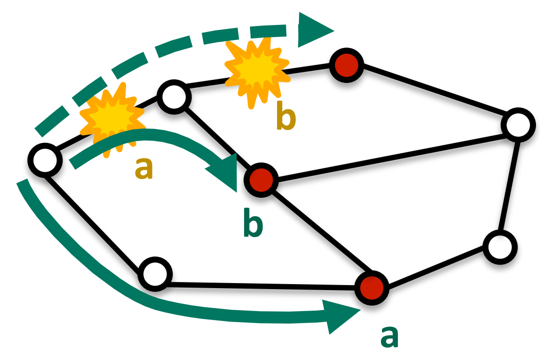

Failure dependent rerouting? (FD vs FID)

- FD is best, obviously
- Yet, difference is limited (few %)
  – at least for small K (= number of server sites)

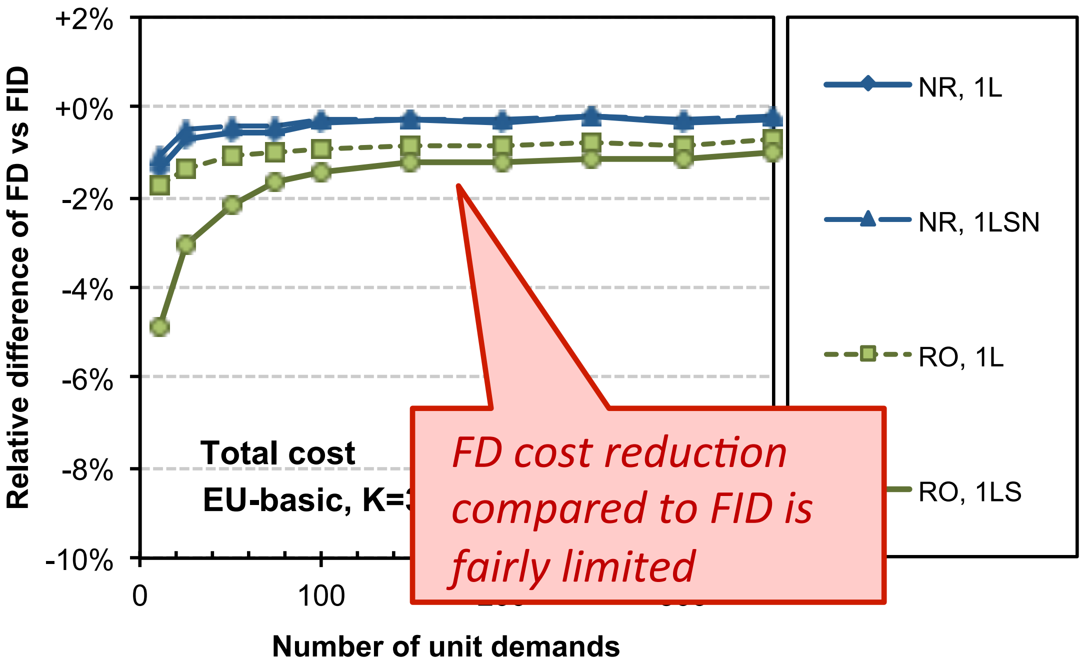

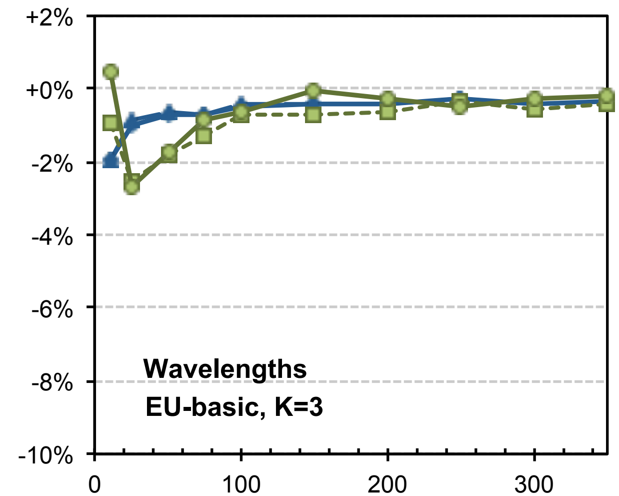

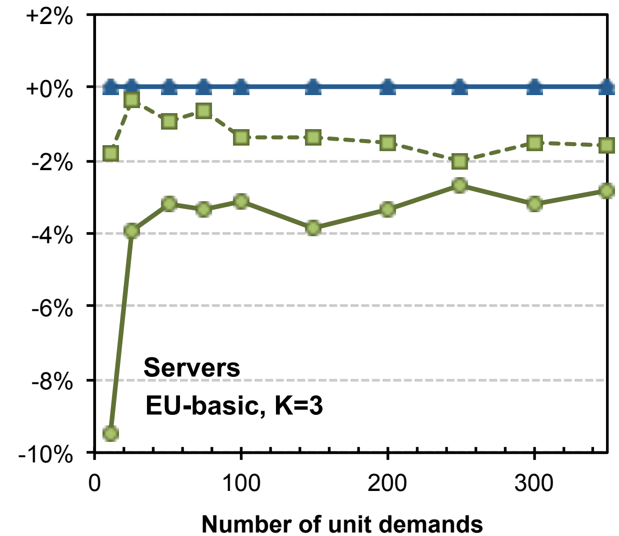

Failure dependent rerouting? (FD vs FID)

- FD is best, obviously
- Yet, difference is limited (few %)
  – at least for small K (= number of server sites)

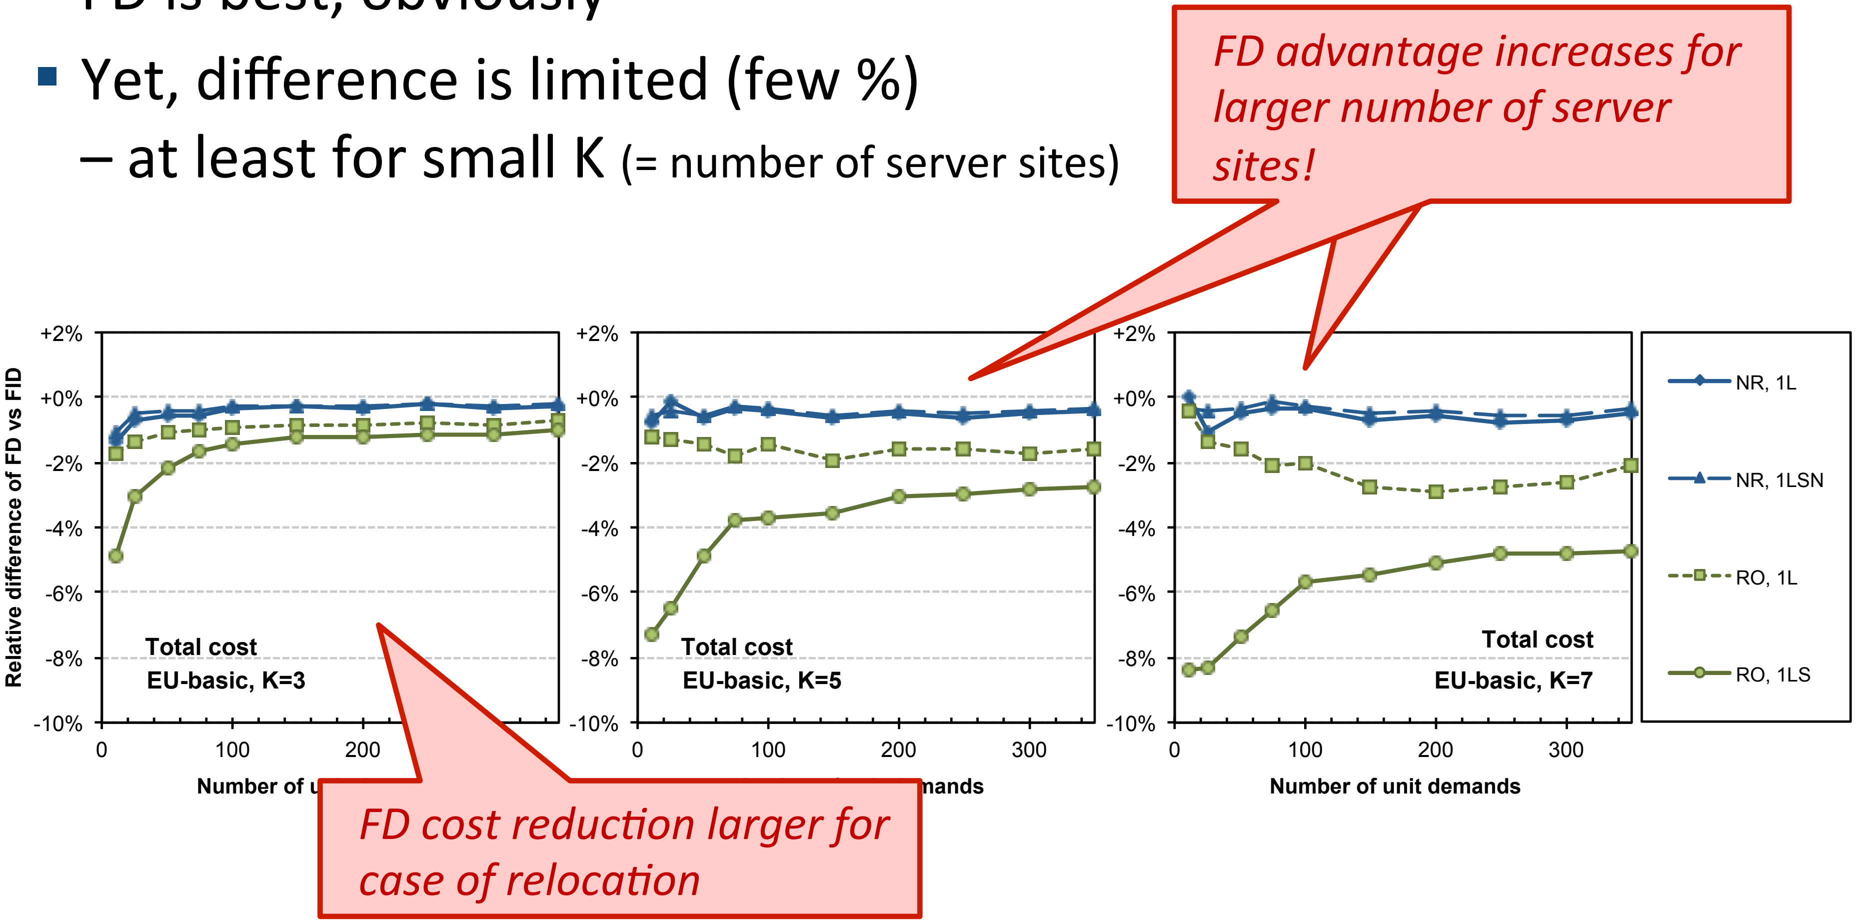

Influence of K?

- K ➚
  - Wavelength reduction more pronounced
  - Lower extra cost to provide single server failure protection

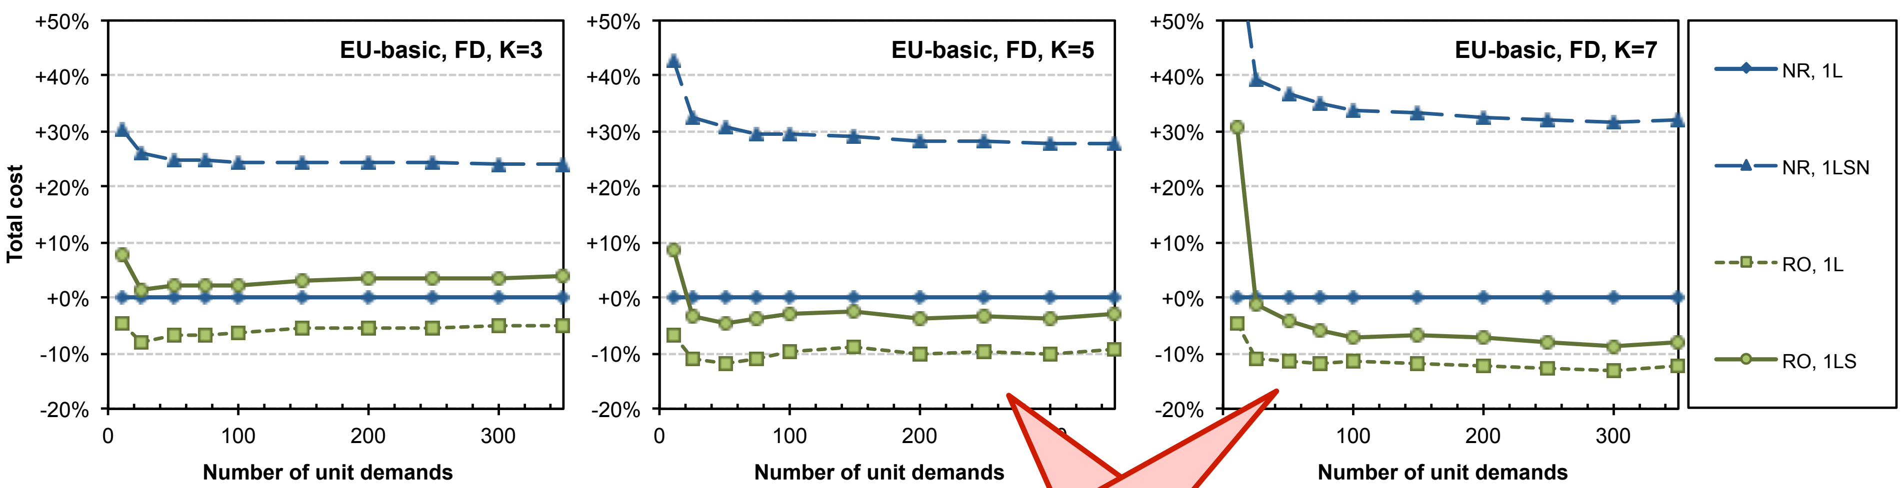

Relocation advantage increases for larger number of server sites!

Agenda

1. Introduction
2. Problem statement, Model & ILP solutions
3. Case study
4. Conclusions

Conclusions

- Dimensioning algorithm for resilient optical grids
  - Exploit relocation for resiliency
  - Small ILP for finding K best locations
  - ILP (Column generation) for server & network dimensions
  - Generic model based on SRLG concept

- Case study on EU network topology [10-350 unit demands]
  - Relocation offers cost adantage of up to 10% to protect against single network link failures
  - Total cost to protect against 1LS with relocation
    ~ Cost to protect against 1L only, without relocation
  - Relocation advantage more substantial for larger number of server sites
  - Failure-dependent rerouting advantage if we use relocation (couple of %)

Future/ongoing work

- Impact of topology
- Refinement of K-location choosing ILP
- Incorporating choice of K sites into CG-ILP approach
- …

B. Jaumard, A. Shaikh and C. Develder, *"Selecting the best locations for data centers in resilient optical grid/cloud dimensioning"*, in Proc. 14th Int. Conf. Transparent Optical Netw., Coventry, UK, 2-5 Jul. 2012

Thank you … any questions?

Prof. Chris Develder

chris.develder@intec.ugent.be

Ghent University – IBBT

Acknowledgments

GEYSERS: Generalised Architecture for Dynamic Infrastructure Services
www.geysers.eu

Research Foundation – Flanders (FWO)
Post-doctoral fellowships (M. De Leenheer, C. Develder)
www.fwo.be

IWT: Agency for Innovation by Science and Technology
PhD scholarship (J. Buysse)
www.iwt.be

Ghent University
Stevin Supercomputer Infrastructure
www.ugent.be