UNIVERSITEIT GENT

FACULTY OF ENGINEERING AND ARCHITECTURE

Time-varying resilient virtual network mapping for multi-location cloud data centers

Minh Bui[1], Ting Wang[1], Brigitte Jaumard[1], Deep Medhi[2] and Chris Develder[3]

1: CSE, Concordia University, Montreal, Canada

2: CSEE, University of Missouri, Kansas City, MO, USA

3: INTEC – IBCN, Ghent University – iMinds, Ghent, Belgium

iMinds FUTURE INTERNET DEPT.

INTEC

Optical clouds

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

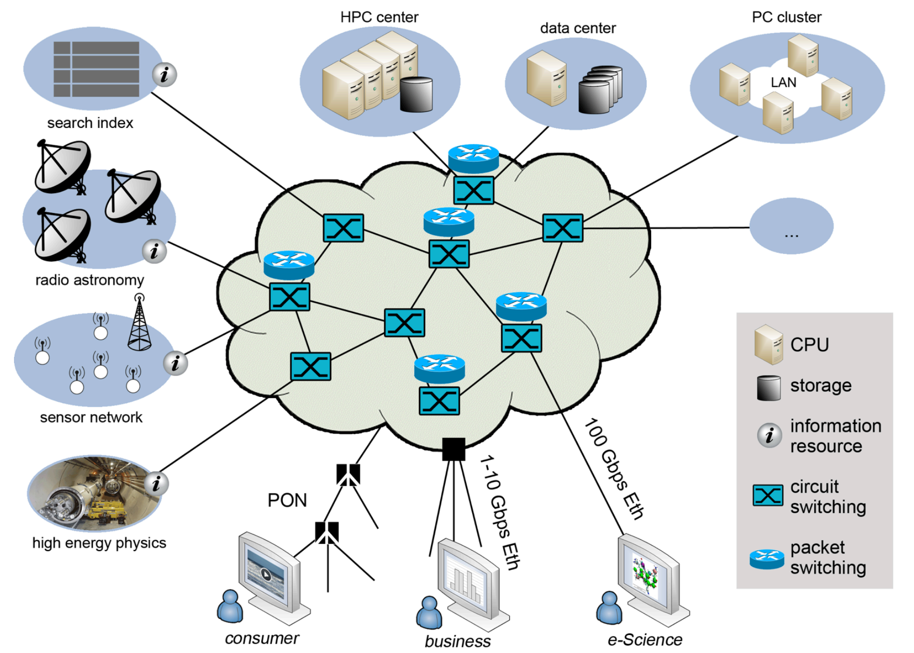

C. Develder, et al., "*Optical networks for grid and cloud computing applications*", Proc. IEEE, Vol. 100, No. 5, May 2012, pp. 1149-1167.

M. Bui, et al., "Time-varying resilient virtual network mapping for multi-location cloud data networks", ICTON 2014

Network virtualization

Physical network is logically partitioned in isolated virtual networks

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

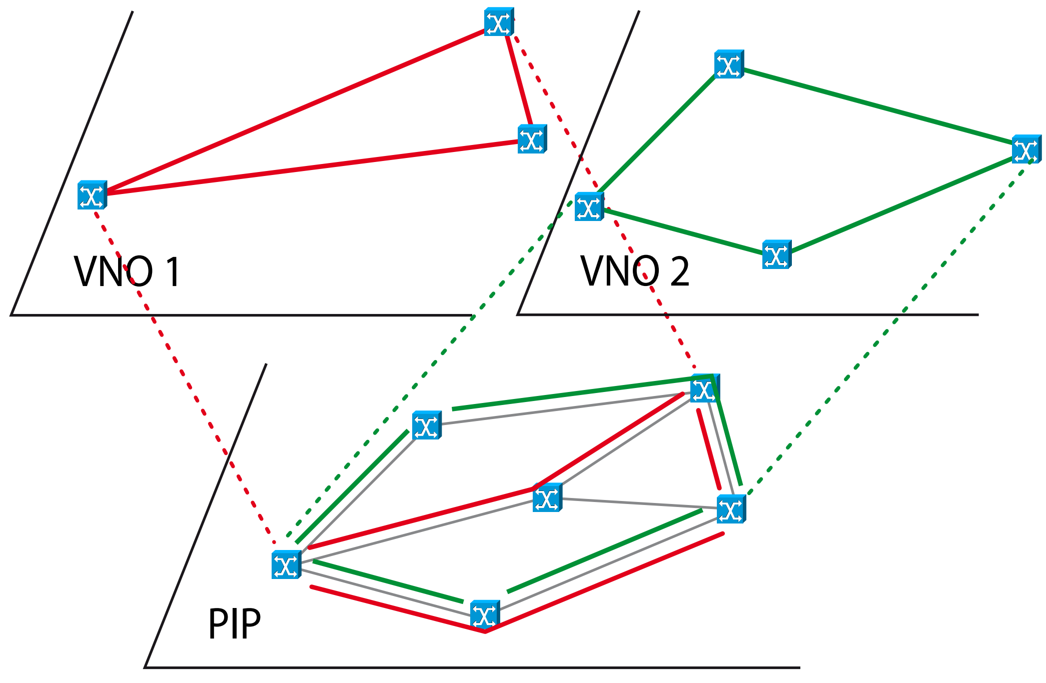

J.A. García-Espín, et al., *"Logical Infrastructure Composition Layer: the GEYSERS holistic approach for infrastructure virtualisation"*, in Proc. TERENA Networking Conference (TNC 2012), Reykjavík, Iceland, 21-24 May 2012.

Overview

Overview

Resiliently provisioning virtual cloud networks

How to choose **the virtual to physical mapping**, such that

- Services remain available in case of network failures
- Bandwidth for providing services is minimal

?

Note:

- Anycast: requests coming from users can be served by any server
- Cloud services offered by VNO
- Cloud services run on top of PIP

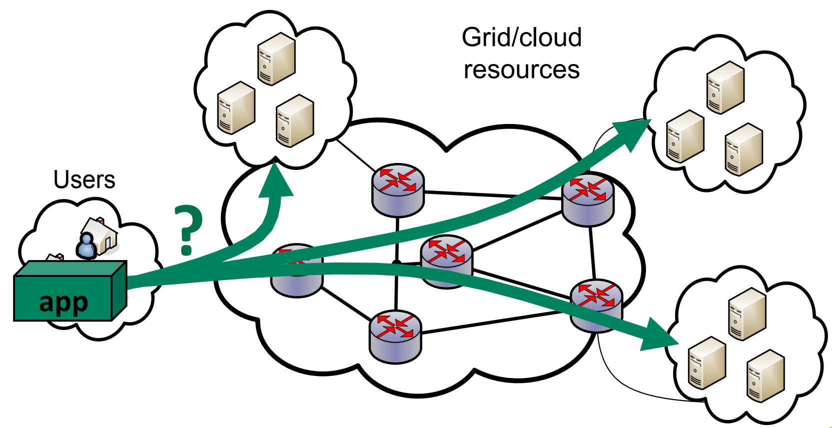

B. Jaumard, A. Shaikh and C. Develder, *"Selecting the best locations for data centers in resilient optical grid/cloud dimensioning (Invited Paper)"*, in Proc. 14th Int. Conf. Transparent Optical Netw. (ICTON 2012), Coventry, UK, 2-5 Jul. 2012.

M. Bui, et al., "Time-varying resilient virtual network mapping for multi-location cloud data networks", ICTON 2014

Two proposed protection schemes:

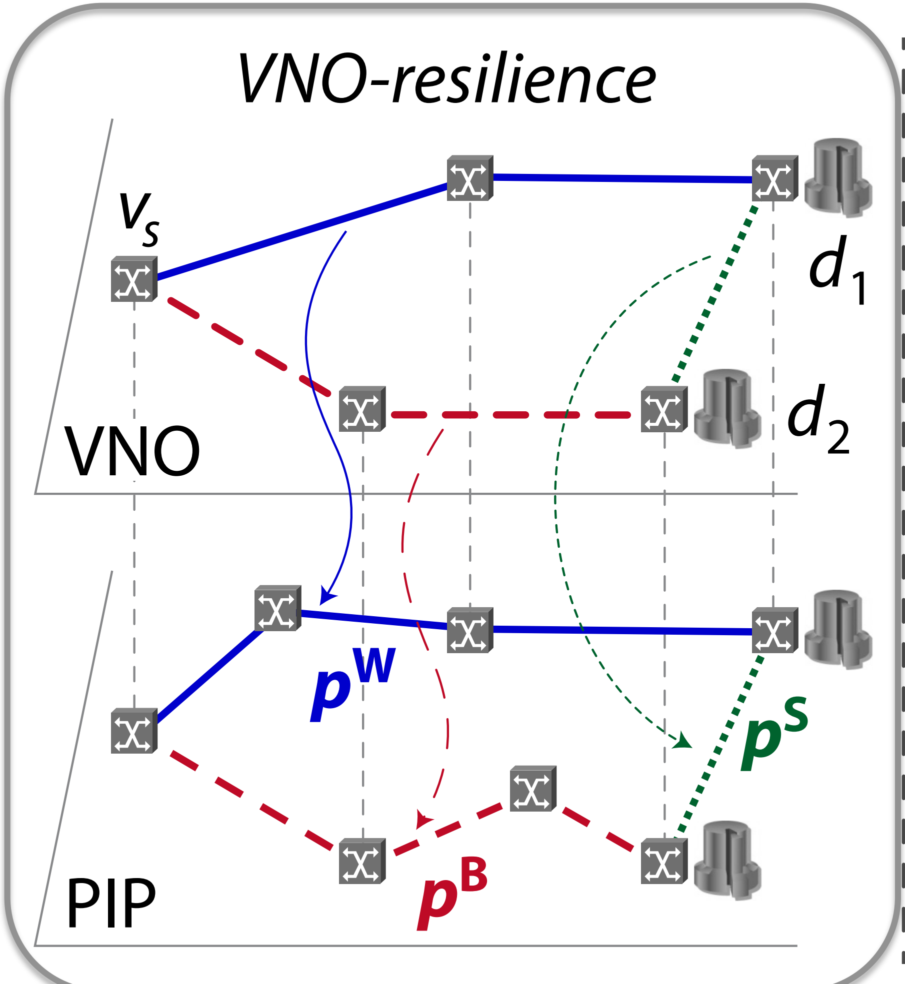

This paper

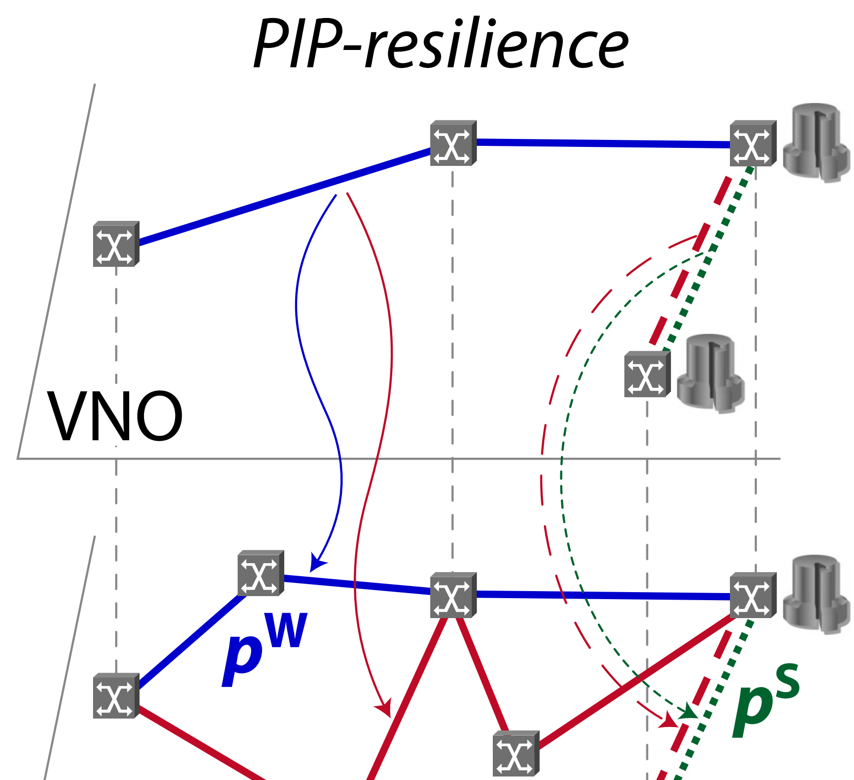

M. Bui, B. Jaumard, and C. Develder, *"Anycast end-to-end resilience for cloud services over virtual optical networks"* (Invited Paper), in Proc. 15th Int. Conf. Transparent Optical Netw. (ICTON 2013), Cartagena, Spain, 23-27 Jun. 2013.

M. Bui, et al., "Time-varying resilient virtual network mapping for multi-location cloud data networks", ICTON 2014

Related work: Static traffic scenarios

- **Traditional dimensioning (no virtualisation, no resilience)**:
  - Develder *et al*. 2009: Anycast, flexibility in choosing data center
- **Resilient** dimensioning problem:
  - Shaikh *et al.* 2011, Develder *et al.* 2013: scalable method, no synchronization between working and backup DCs
- Routing cloud service requests and **mapping a VNet** to the physical infrastructure separately:
  - Lee *et al.* 2009, Yu *et al.* 2010: Survivable VNet embedding, but *assume VNet is given*
  - Jiang *et al.* 2012, Alicherry *et al.* 2012: Optimal server selection and routing of anycast services in the physical layer for intra- and inter-DC networks but *no resilient network design in the virtual layer*
- **VNet planning** problem:
  - Barla *et al.* 2012, Barla *et al.* 2013: using mixed integer linear programming, but *no synchronization between working and backup DCs*
  - Bui *et al.* 2013 (ICTON): first model that incorporates synchronisation path, but *still static traffic!*

Problem statement

- Study **time-varying traffic**:
  Traffic pattern changes from one period (t) to the next (t+1)

- Key research question:
  Benefit (in network resource usage) of **changing routes** for legacy traffic, i.e., that continues from t to t+1?
  - Does it help to only change backup paths?
  - Or do we need to change working as well?
  - For *all* legacy traffic?

M. Bui, et al., "Time-varying resilient virtual network mapping for multi-location cloud data networks", ICTON 2014

Problem statement

Given:

- Cloud network topology: $G = (V, L)$, with V = nodes, L = links
- Locations of the data centers, $V_D \subseteq V$
- Set of service requests, K partitioned into
  - K^{LEG}: requests in period t, that continue into the next t+1
  - K^{ADD}: requests in subsequent period t+1

  and characterized by
  - v_k: source of service
  - Δ_k: bandwidth requirement in period t
  - Services originating from the same source are aggregated
- Routing of the requests in period t: K^{LEG}

Find:

- Choice of primary and backup **DC locations** for each service in period t+1
- Primary, backup and synchronization **paths** in period t+1

Such that:

- Total network **bandwidth utilization is minimized**
- K^{LEG} are (i) unchanged, (ii) only changed for backup, (iii) freely changed

Overview

1. Introduction
2. Problem statement
3. Model & solution approach
4. Case study
5. Conclusions

Overview

1. Introduction
2. Problem statement
3. Model & solution approach
4. Case study
5. Conclusions

Solution: Column generation model

- Column generation idea:
  - Many different “configurations”
  - Start from a restricted subset of such “configurations”
  - Iteratively find additional configurations that reduce the cost:
    (1) Restricted Master Problem (RMP)
    (2) Pricing Problem (PP) to find new configs
- A configuration =
  - **Working** path
  - **Backup** path
  - **Sync** path between the primary & backup DCs

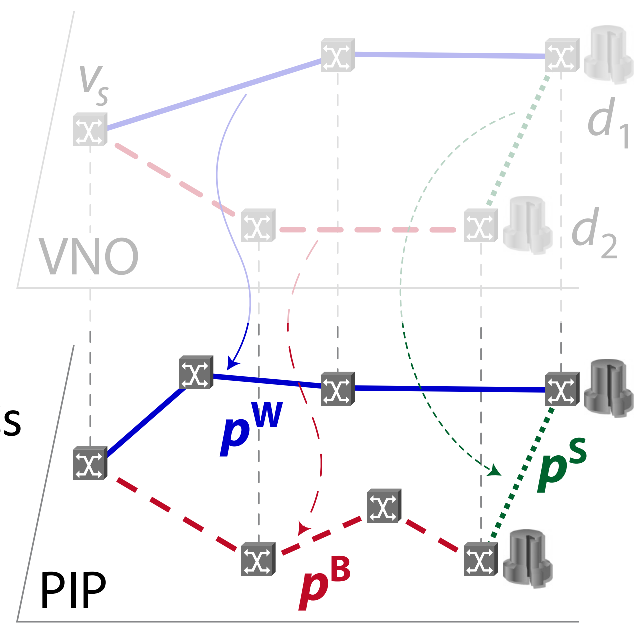

Column generation solution algorithm

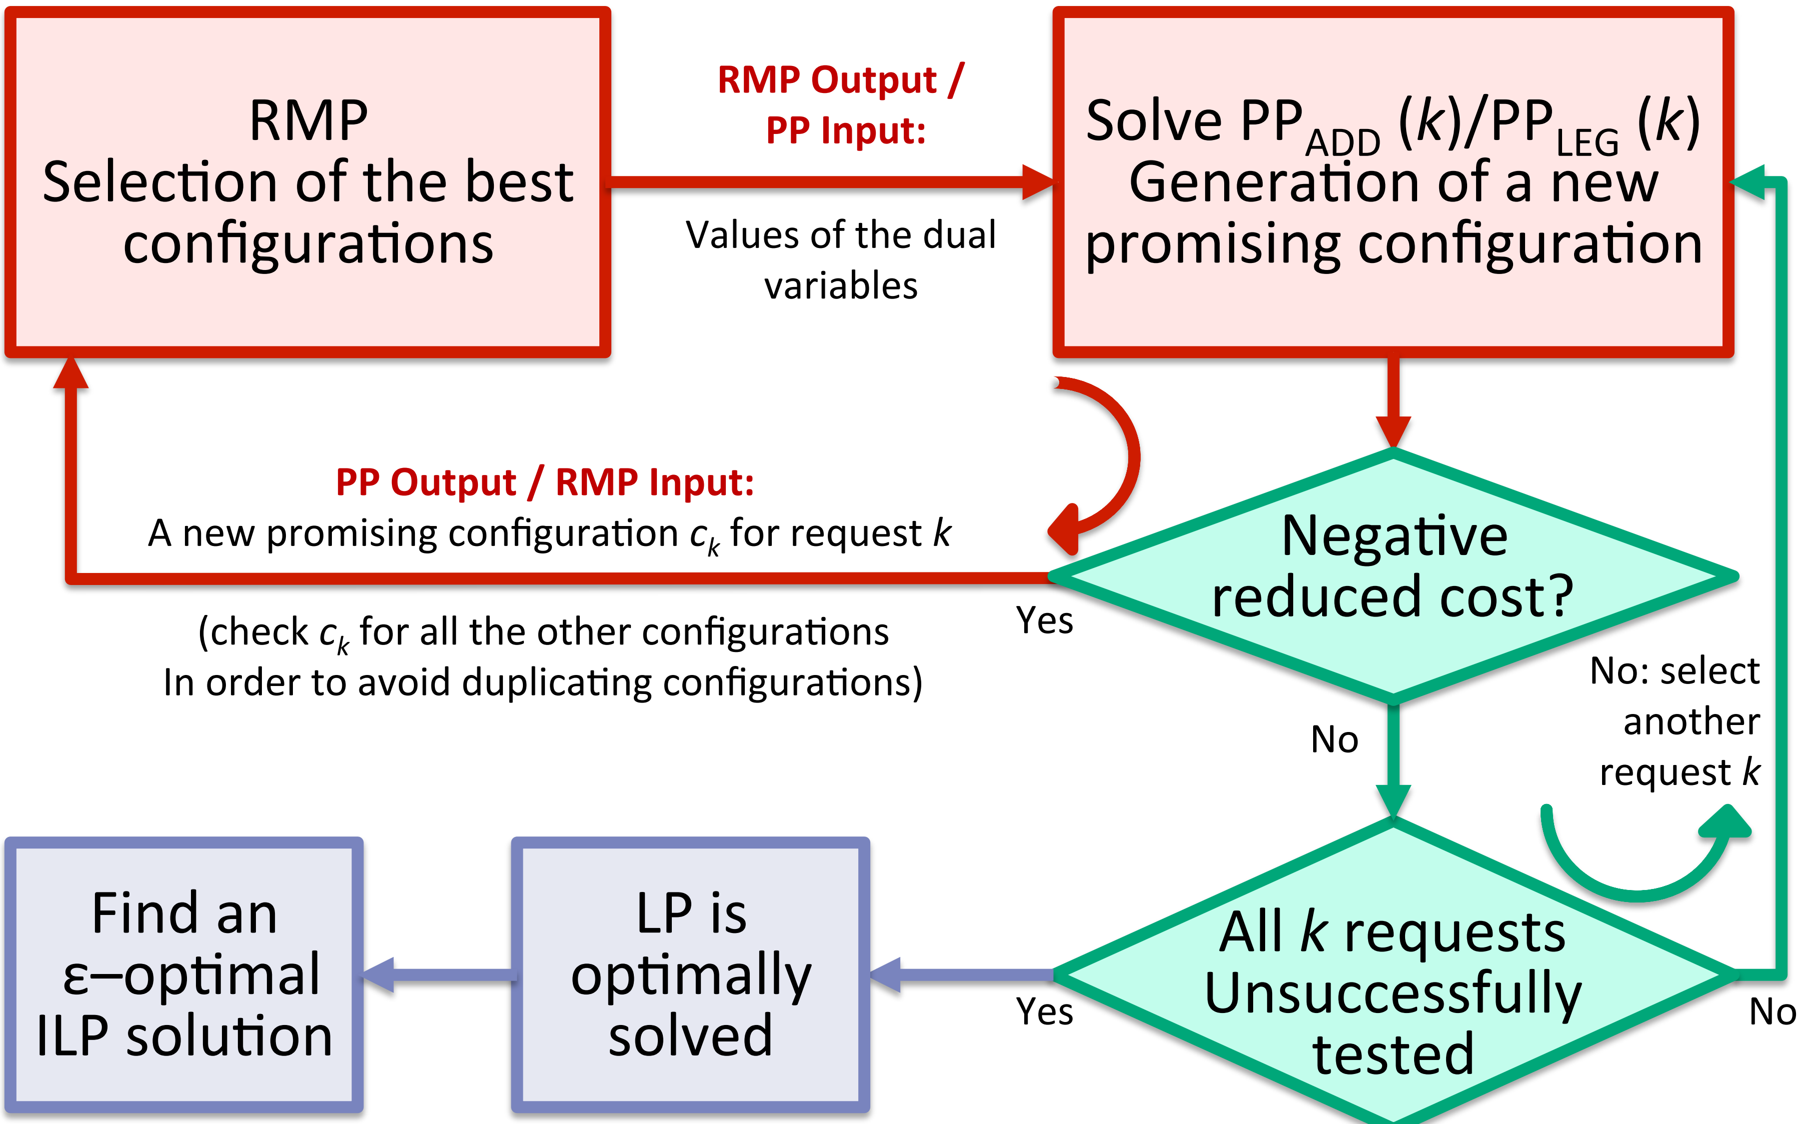

M. Bui, et al., "Time-varying resilient virtual network mapping for multi-location cloud data networks", ICTON 2014

Restricted Master Problem (RMP)

$$\min \sum_{\ell \in L} \underbrace{\left(\beta_\ell^{\mathrm{W}} + \beta_\ell^{\mathrm{B}} + \beta_\ell^{\mathrm{S}}\right)}_{\mathrm{BW}_\ell} \cdot \|\ell\|$$

$$+ \, \mathrm{PENAL}^{\mathrm{DISRUPT\_B}} \sum_{k \in K^{\mathrm{LEG}}} x_k^{\mathrm{LEG\_BS}}$$

$$+ \, \mathrm{PENAL}^{\mathrm{DISRUPT\_W}} \sum_{k \in K^{\mathrm{LEG}}} x_k^{\mathrm{LEG\_W}}$$

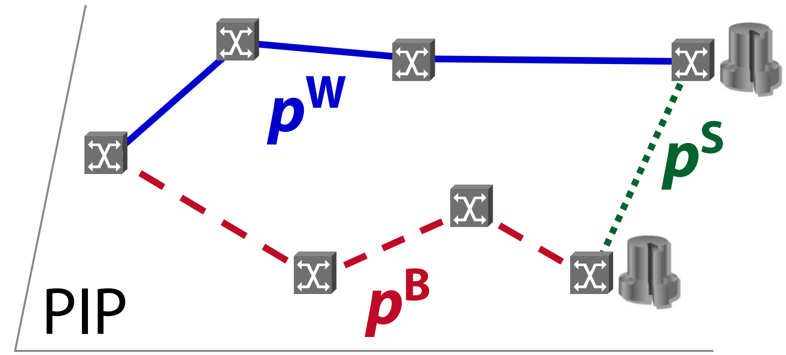

Case (ii):
penalty for disrupting B/S
path of legacy

Case (iii):
penalty for disrupting W
path of legacy

Constraints:

- Assure all requests are granted
- Count legacy changes x^{LEG_BS}, x^{LEG_W}
- Compute W, B, S bandwidths
- Check capacity constraints on data centers

Overview

1. Introduction
2. Problem statement
3. Model & solution approach
4. Case study
5. Conclusions

Overview

Case study

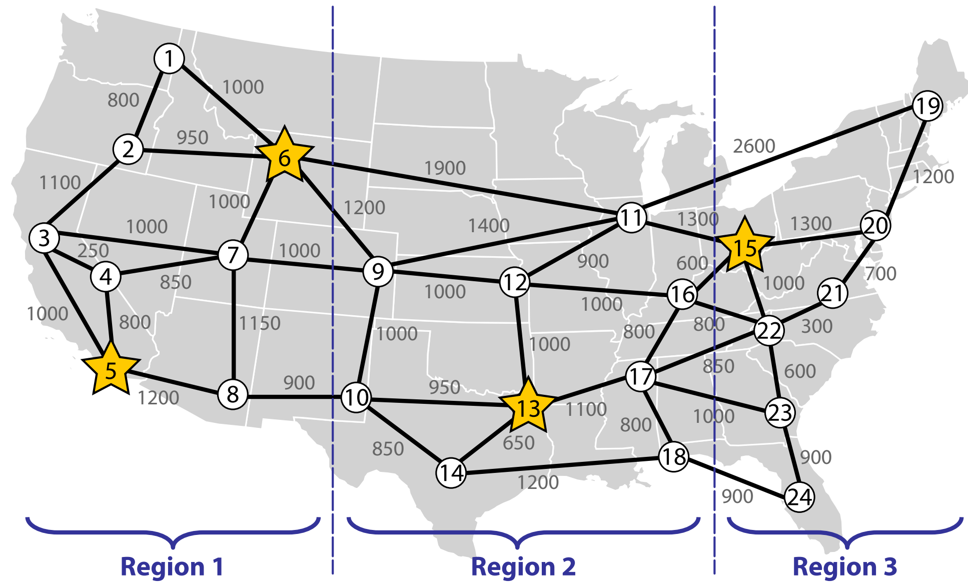

- Topology:
  - 24 nodes, 43 links
  - Data centers in ☆: CA, WY, TX, OH
- Traffic:
  - **Total** of 20…80 requests (same for *t* and *t*+1) with Δ_k in [0,1] units; synchronisation fraction $\delta_k = 0.1$
  - **Period *t***: 30% region 1, 50% region 2, 20% region 3
  - **Period *t*+1**:

Scenario	Region 1	Region 2	Region 3
40% legacy	20% drop 20% add	30% drop 10% add	10% drop 30% add
60% legacy	10% drop 10% add	20% drop 10% add	10% drop 20% add
80% legacy	10% drop 10% add	10% drop -	- 10% add

M. Bui, et al., "Time-varying resilient virtual network mapping for multi-location cloud data networks", ICTON 2014

Results: Net total bandwidth savings?

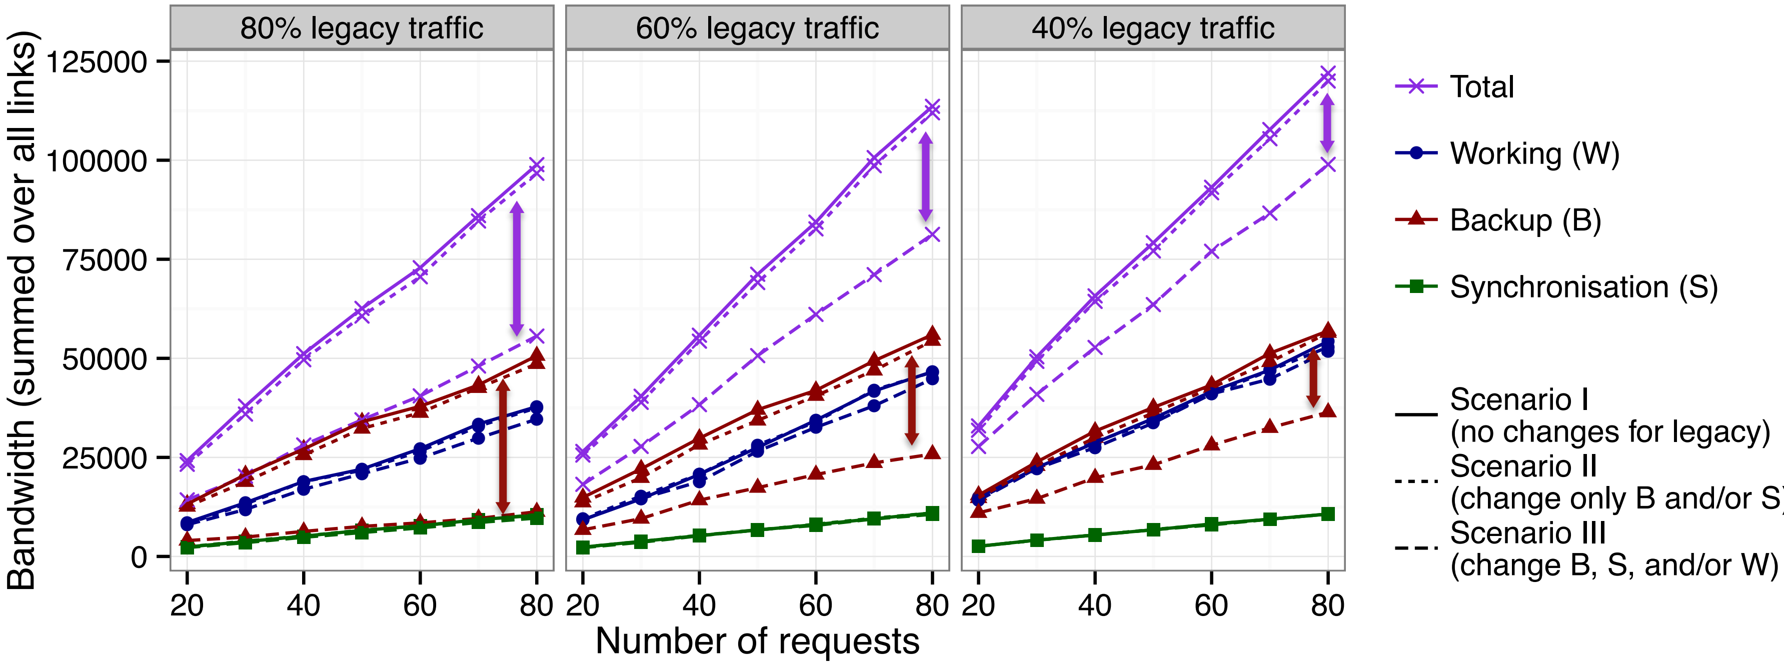

1. Only changing backup (II,----) does not **save** much; changing also working (III, – – -) does
2. Capacity savings are realized through **sharing of backup**
3. Savings obviously diminish with decreasing legacy fraction (L→R)

M. Bui, et al., "Time-varying resilient virtual network mapping for multi-location cloud data networks", ICTON 2014

Results: Change *all* legacy routes?

Of the legacy requests, we typically change >50%, but only about **20%** need to **reroute the working path**

Legacy traffic	Path changes	Total demand (number of requests) 20	30	40	50	60	70	80	α_{LEG} over all demands
80%	Scen. II – Changed B and/or S	9.6	12.8	16.6	17.2	20.6	22.6	22.8	46.73%
	Scen. III – Fixed W, changed B/S	6.2	8.2	12.4	16.8	16.2	18.0	25.6	37.08%
	Scen. III – Changed W	3.2	3.8	6.8	5.6	10.0	12.8	7.8	18.14%
	Scen. III – Total changes	9.4	12.0	19.2	22.4	26.2	30.8	33.4	55.22%
60%	Scen. II – Changed B and/or S	8.0	10.0	12.2	15.6	17.4	19.6	15.2	50.25%
	Scen. III – Fixed W, changed B/S	4.8	6.2	8.8	12.2	11.4	11.8	16.0	34.98%
	Scen. III – Changed W	3.0	3.0	5.8	6.0	9.4	11.4	6.6	21.83%
	Scen. III – Total changes	7.8	9.2	14.6	18.2	20.8	23.2	22.6	56.82%
40%	Scen. II – Changed B and/or S	5.8	7.0	9.4	12.2	10.6	13.4	11.6	54.12%
	Scen. III – Fixed W, changed B/S	3.0	3.6	6.8	7.4	7.2	7.6	11.2	34.16%
	Scen. III – Changed W	1.8	2.4	4.2	3.6	5.8	6.0	4.4	20.87%
	Scen. III – Total changes	4.8	6.0	11.0	11.0	13.0	13.6	15.6	55.03%

M. Bui, et al., "Time-varying resilient virtual network mapping for multi-location cloud data networks", ICTON 2014

Overview

Overview

1. Introduction
2. Problem statement
3. Model & solution approach
4. Case study
5. Conclusions

Conclusions

- Scalable column-generation based method for resilient VNet planning of time-varying traffic
- Our case study shows that:
  - Changing legacy traffic from one period to the next only really saves if we allow changing the working paths
  - … but we need only to change around 20% of them
- Future work:
  - Optimization of choice of DC locations (e.g., ‘scattered’ vs ‘paired’, see ICTON 2013)
  - Extensive studies of different traffic patterns, over multiple periods

M. Bui, et al., "Time-varying resilient virtual network mapping for multi-location cloud data networks", ICTON 2014

Thank you … any questions?

Prof. Chris Develder

chris.develder@intec.ugent.be

Ghent University – iMinds

UNIVERSITEIT GENT

iMinds FUTURE INTERNET DEPT.

M. Bui, et al., "Time-varying resilient virtual network mapping for multi-location cloud data networks", ICTON 2014

INTEC