Protection Plan Design for Cloud Tenants with Bandwidth Guarantees

Hyame Assem Alameddine, Sara Ayoubi, Chadi Assi
Concordia University, Montreal, Canada

DRCN 2016 – Paris

Presented by: Hyame Assem Alameddine

Outline

- Introduction
- Bandwidth allocation in data centers
- Motivation and challenges
- Protection plan design
- Numerical results
- Conclusion
- Future work
- References
- Questions

Introduction

Introduction (1)

Cloud computing overview

Cloud Computing

Sharing of computing and network resources in a pay-as-you-use basis.

Supports heterogeneous applications [14]:

- Online banking (dedicated bandwidth and high survivability)
- Audio/video streaming (dedicated bandwidth and bounded delay)
- E-mail and web browsing (best-effort delivery)

⬇

Network Virtualization

Introduction (2)

Problem 1: Bandwidth guarantee in the cloud

- Why cloud providers do not guarantee bandwidth for their clients?
  - Shared network links among tenants
  - Transmission control protocol (TCP) congestion control establishes flow fairness but not tenant fairness
- What are the implications of the lack of bandwidth guarantee in the cloud?
  - Unpredictable applications performance
  - Bounding cost for running applications in the cloud with the current pricing policy (pay-as-you-go)
  - Lost of potential clients who require network performance for their critical applications
- How bandwidth is guaranteed in the cloud?
  - Through the use of abstraction models (Pipe model [1],[8], hose model [1][8][6], TAG model [4][5], etc.)

Introduction (3)

Problem 2: Applications reliability in the cloud

- Why applications reliability in the cloud is important (some numbers) ? [3]
  - Cost of one hour downtime of critical applications (banking, retail systems, etc.) varies between $25,000 and $150,000
  - One of 10 companies requires more than 99.999% availability
- How reliability can be guaranteed in the cloud ?
  - Provisioning of additional computing and network resources (backup virtual machines (VMs), backup bandwidth)
  - Providing worst case survival (WCS) [4][5][10]

Introduction (4)

Our Contribution

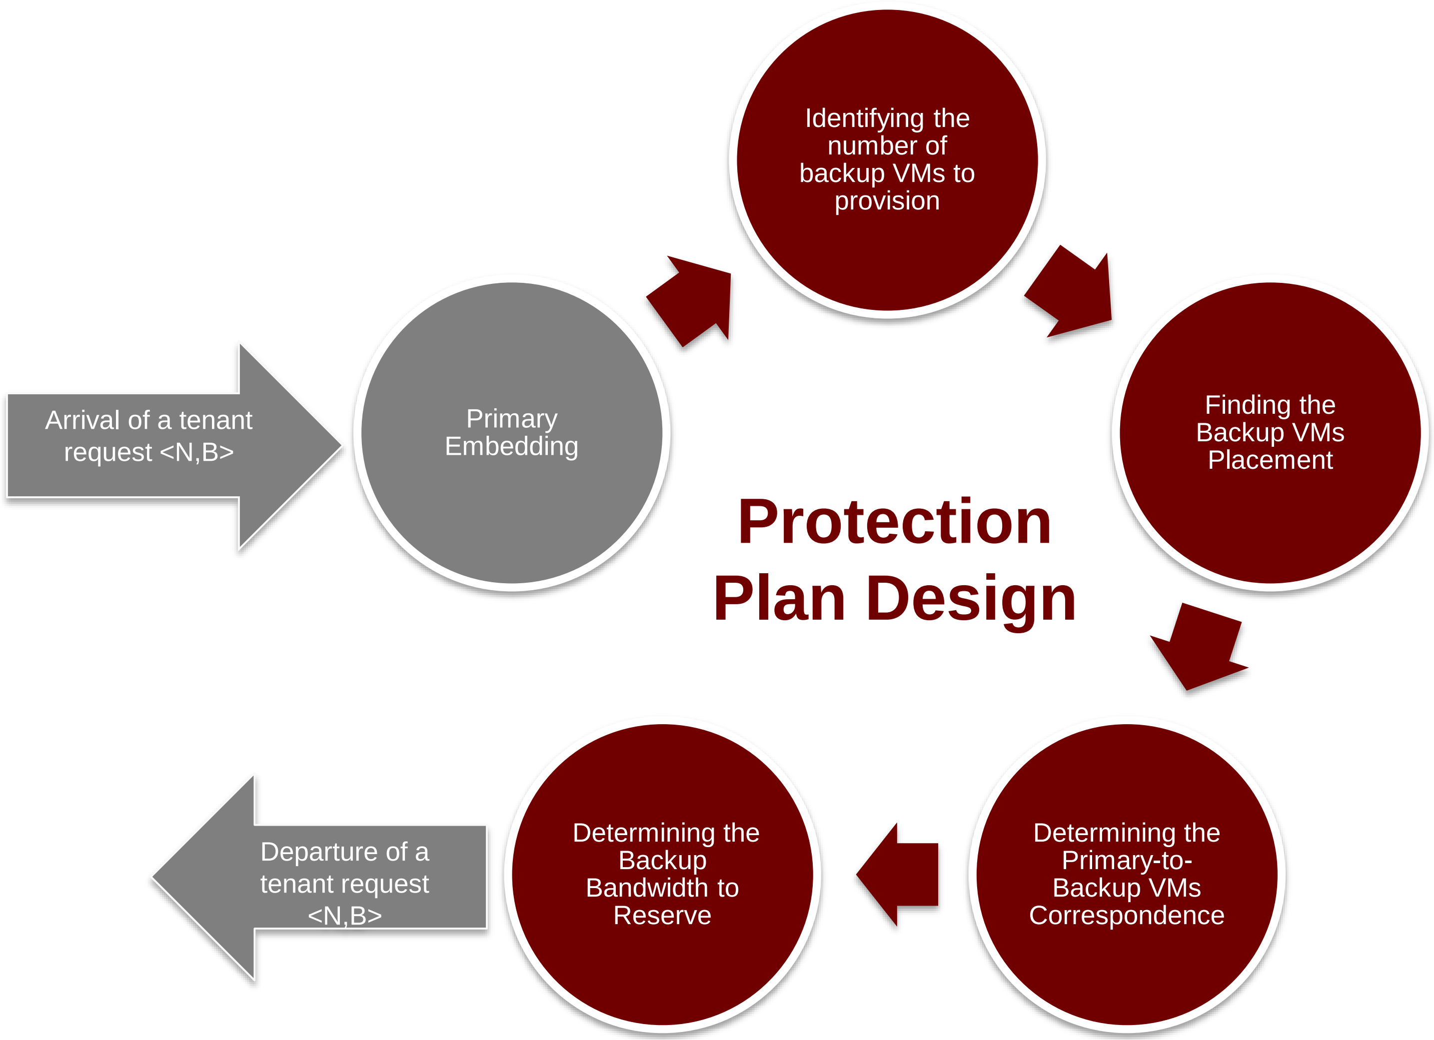

Bandwidth Allocation in Data Centers

Allocating Bandwidth in data centers

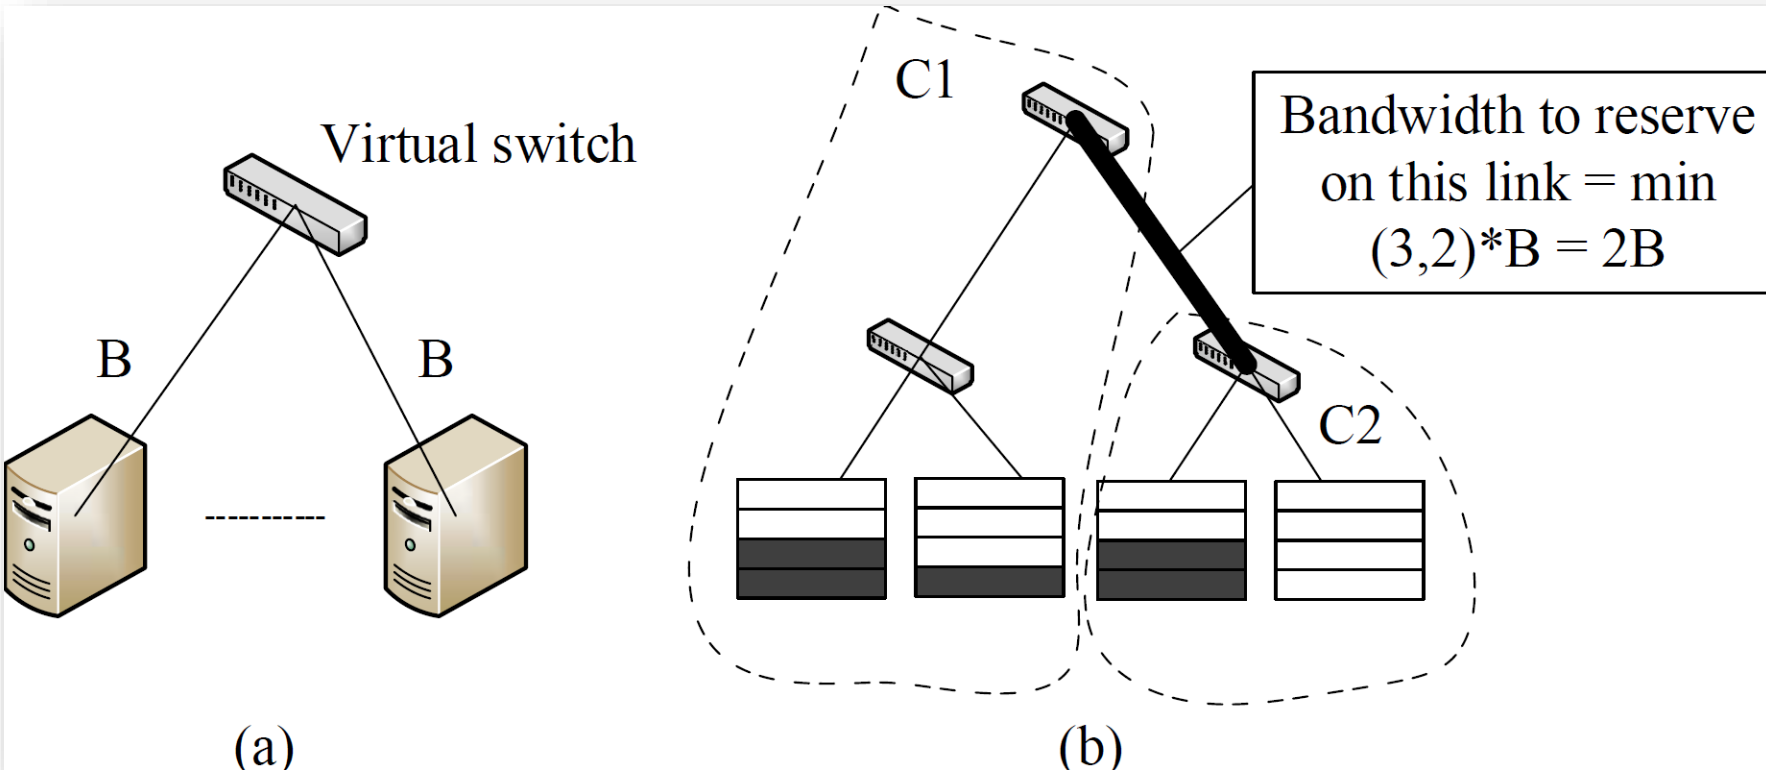

Fig.1: Bandwidth allocation based on the hose model

Motivation and Challenges

Motivation and Challenges (1)

Step 1: Identifying the number of backup virtual machines to provision

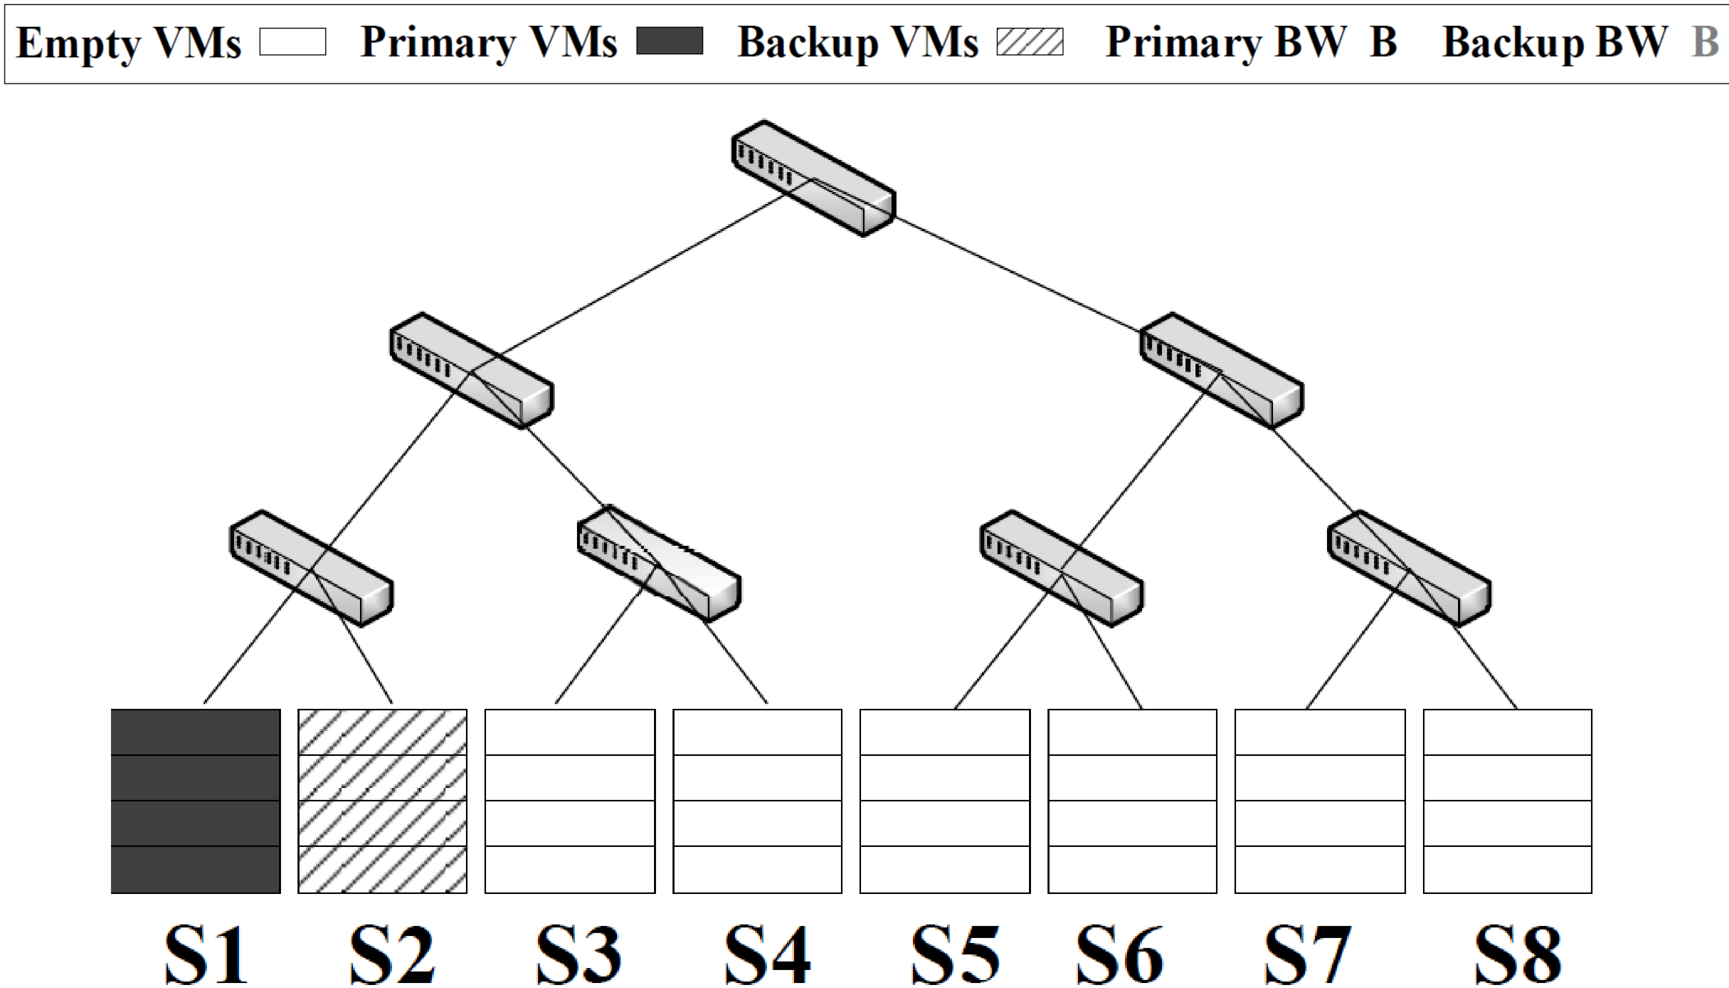

(a) Total bandwidth to reserve = 0B

Motivation and Challenges (1)

Step 1: Identifying the number of backup virtual machines to provision

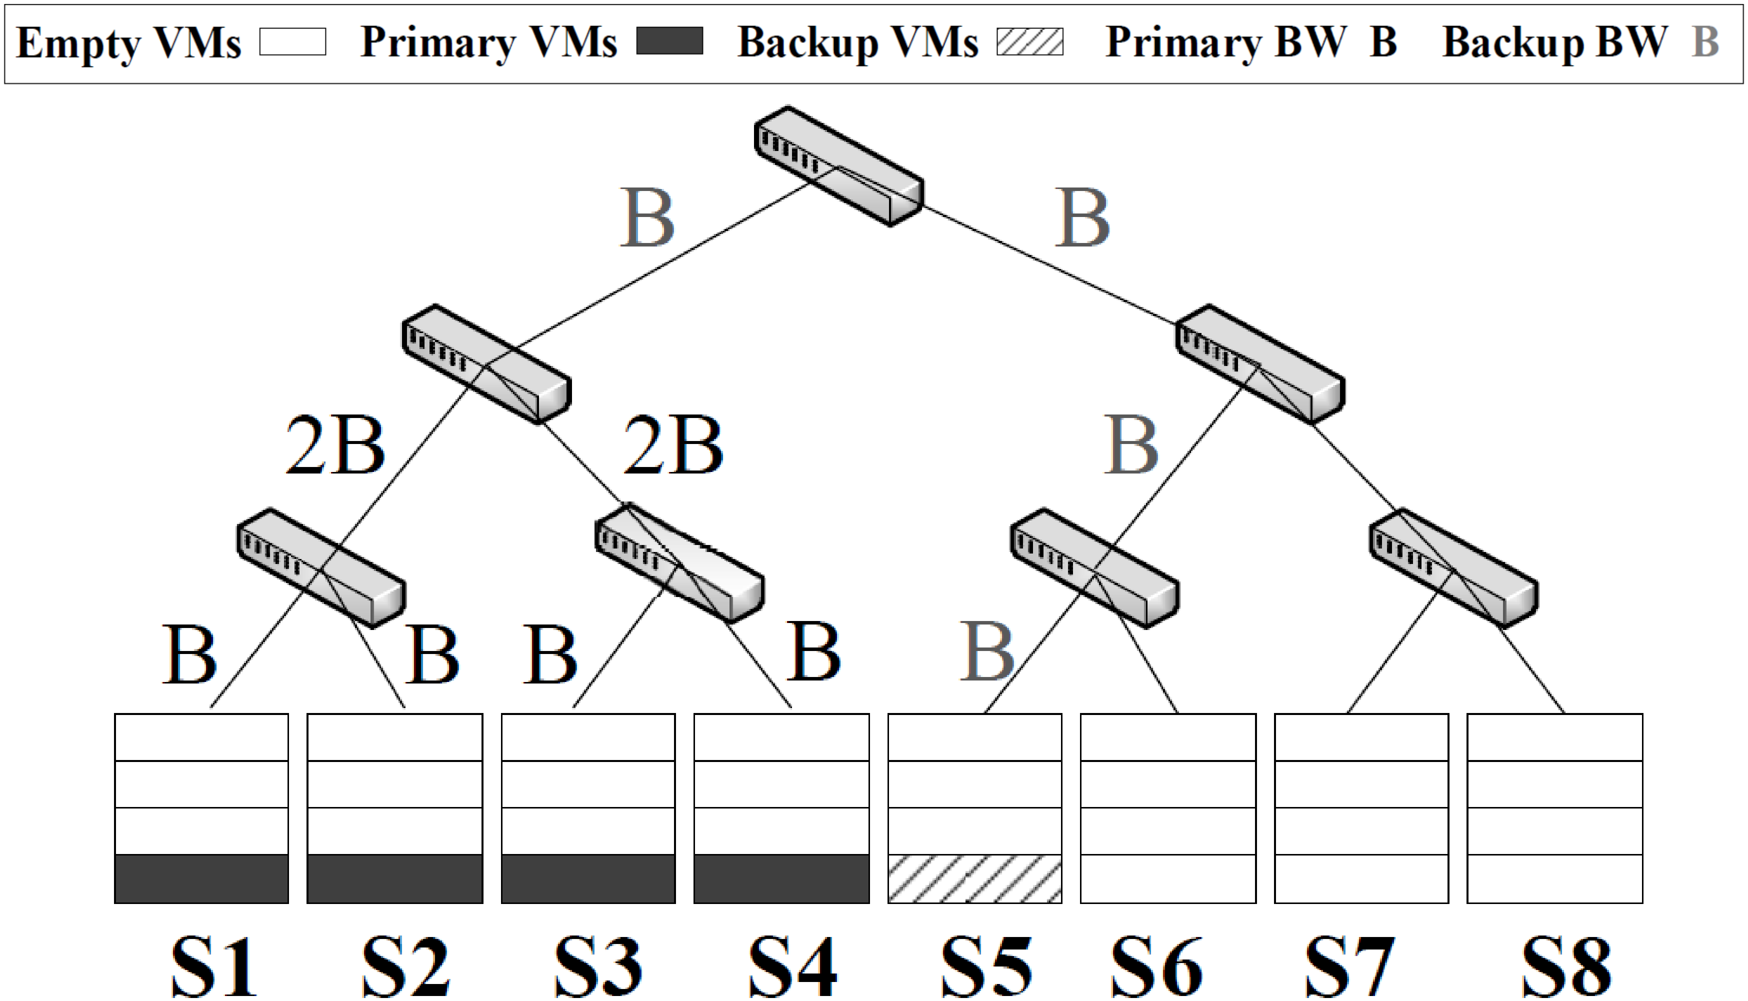

(b) Total bandwidth to reserve = 8B + 4B = 12B

Motivation and Challenges (1)

Step 1: Identifying the number of backup virtual machines to provision

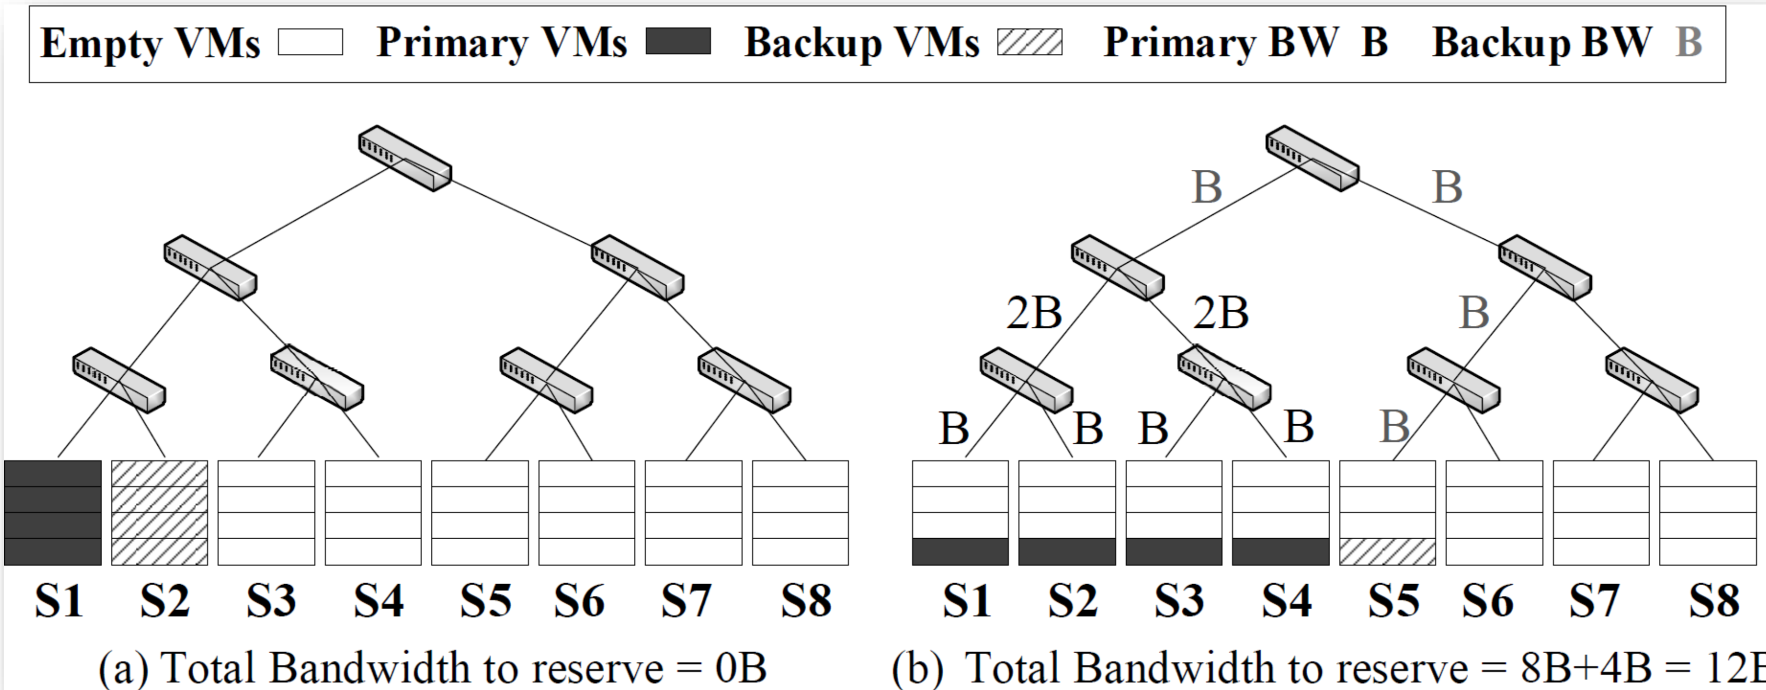

Fig.2: Trade-off between the primary embedding solution and the incurred backup footprint

Observations:

- Collocation : **+** reduces the bandwidth to reserve
  - increases the number of backup VMs
- In this work, we assume that primary VMs is performed based on collocating VMs under the smallest sub-tree.

Motivation and Challenges (2)

Step 2: Finding the backup virtual machines placement

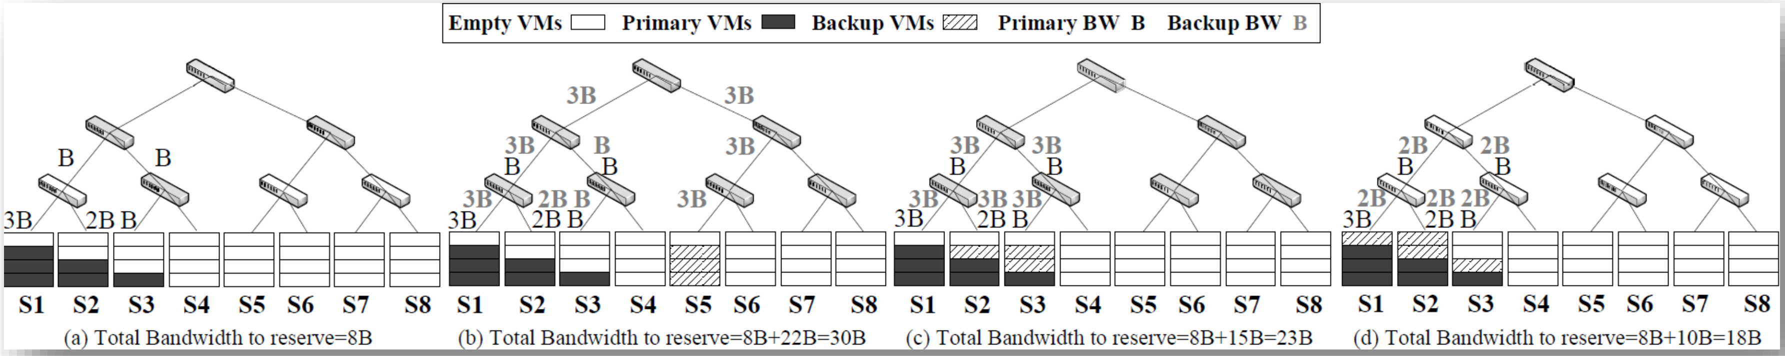

Fig.3: Tenant of 6 VMs requirement with its needed backup VMs for 100% availability

Observations:

- The placement of backup VMs has a reciprocal impact on their number, in addition to the incurred backup footprint.
- Minimum number of backup VMs needed = maximum number of primary VMs hosted on the same physical server.

Motivation and Challenges (3)

Step 3: Determining the Primary-to-Backup virtual machines correspondence

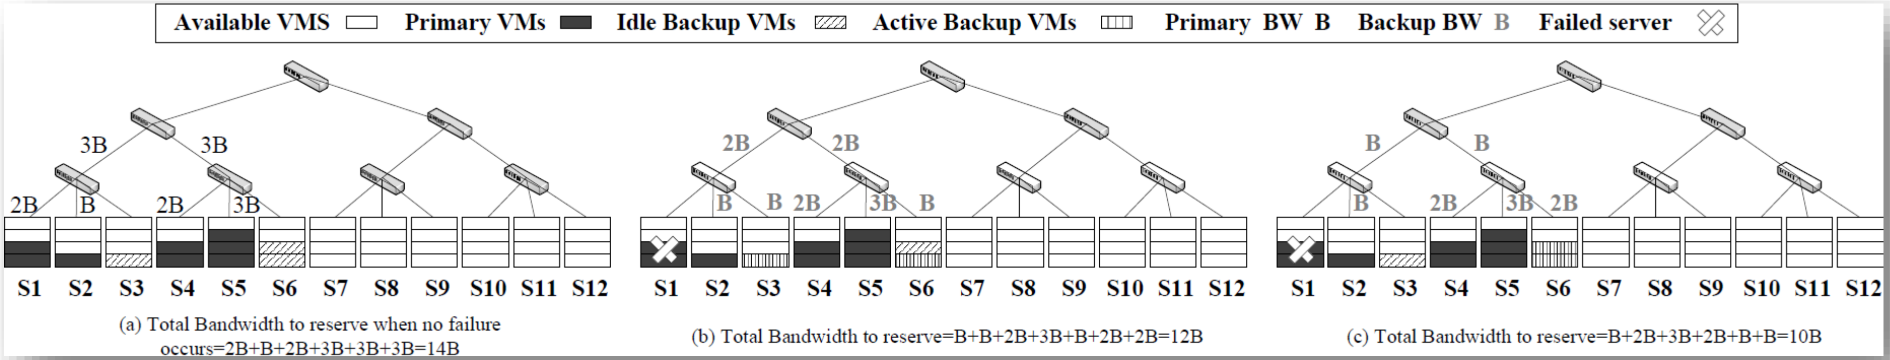

Fig.4: Different protection plans for the same server failure results in different bandwidth consumption

Observations:

- The amount of backup bandwidth to reserve is affected by the primary-to-backup VMs correspondence.
- The best primary-to-backup VMs correspondence is the one that consumes the least bandwidth.

Motivation and Challenges (4)

Step 4: Determining the backup bandwidth to reserve

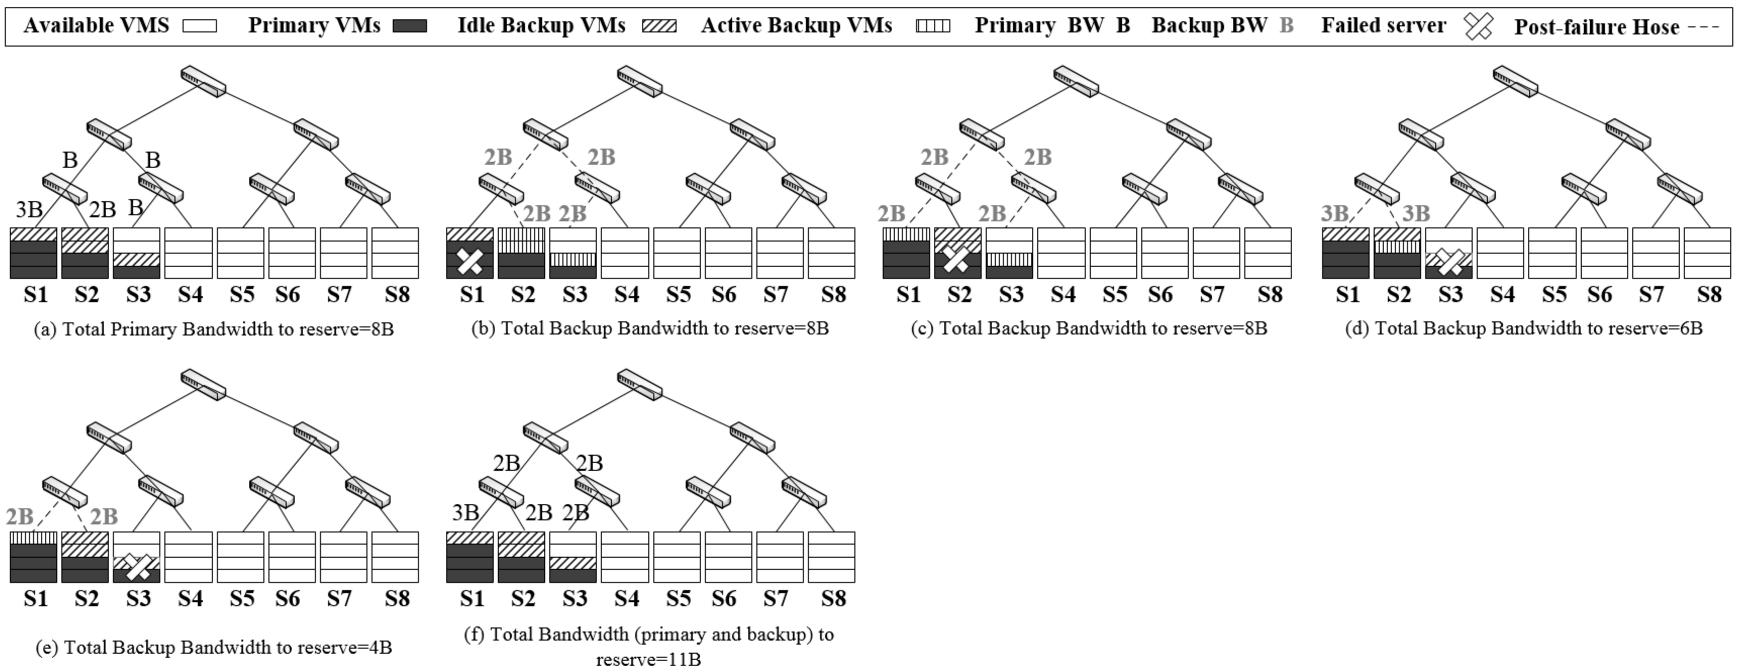

Fig.5: Methodology to determine the backup bandwidth to reserve while considering bandwidth reuse

Observations:

- The backup bandwidth to reserve on a link e is calculated as the maximum bandwidth to reserve on this link by considering the failure of each server hosting the primary VMs of the tenant.
- By considering the reuse of primary bandwidth as backup bandwidth upon a failure, the backup bandwidth to reserve is decreased by the primary bandwidth.

Protection Plan Design

Protection Plan Design (1)

Survivable Virtual Machines Placement with Bandwidth Guarantees (SVMP-BG)-Model

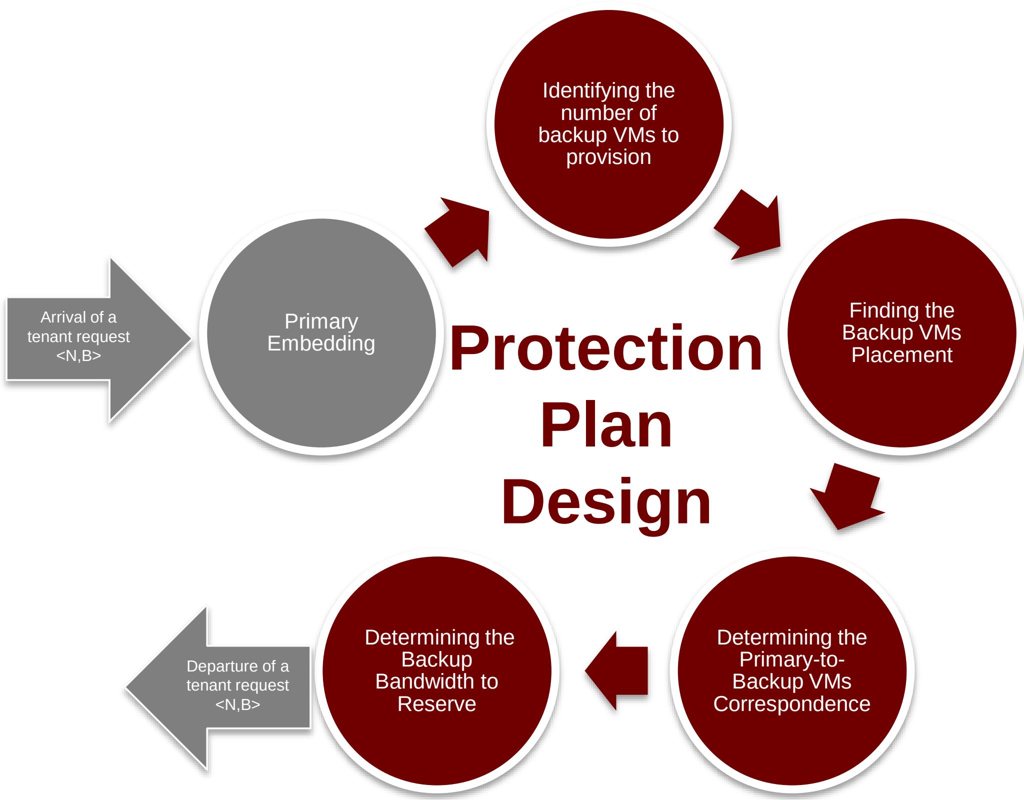

SVMP-BG (Model)

Given:

- Substrate network (single path tree topology)
- Tenant request <N,B>
- Primary embedding of the request

Find:

- Lowest cost protection plan design
- Objective function:

$$\text{Minimize} \quad \alpha\left(\sum_{p=1}^{P}\sum_{k=1}^{K} y_{kp}\right) + \frac{(1-\alpha)}{B}\left(\sum_{ij \in E} \hat{t}_{ij}\right)$$

Protection Plan Design (2)

Survivable Virtual Machines Placement with Bandwidth Guarantees (SVMP-BG)-Heuristic

SVMP-BG can be decomposed into two sub-problems:

The backup placement problem

- Decides on the placement of backup VMs
- Number of backup VMs to embed = maximum number of primary VMs hosted on a server
- Motivates two types of searches for backup VMs placement:
  1. Search with collocation
  2. Search without collocation

Protection plan design (PPD) problem

A relaxation of the SVMP-BG model

Given:

- Substrate network (single path tree topology)
- Tenant request <N,B>
- Primary and backup VMs embedding of the request

Find:

- Lowest cost (backup bandwidth) correspondence between primary and backup VMs
- Objective function:

$$\text{Minimize} \quad \sum_{ij \in E} \hat{t}_{ij}$$

Numerical Results

Université Concordia University

Numerical results (1)

SVMP-BG model vs heuristic

	SVMP-BG Model				SVMP-BG Heuristic			
α	Rejection Rate (%)	Exec. Time (ms)	Tot. Res. Bandwidth (MBps)	Tot. Res. Backup VMs	Rejection Rate (%)	Exec. Time (ms)	Tot. Res. Bandwidth (MBps)	Tot. Res. Backup VMs
0	20	486 094	18230	35	20	330	18846	26
0.25	0	152 068	13740	39	0	389	13740	40
0.5	10	3 718	18054	30	10	405	18358	30
0.75	10	27 829	19156	30	10	470	21612	31
1	20	4 531	12332	27	30	330	9414	27

Table 1: SVMP-BG Model and heuristic comparison over a small network of 12 physical servers

SVMP-BG objective function: $$\text{Minimize} \quad \alpha\left(\sum_{p=1}^{P}\sum_{k=1}^{K} y_{kp}\right) + \frac{(1-\alpha)}{B}\left(\sum_{ij\in E} \hat{t}_{ij}\right)$$

Observations:

SVMP-BG heuristic is:

- More scalable than the SVMP-BG model.
- Able to perform the balance between backup VMs consumption and the backup bandwidth to reserve.

Numerical results (2)

SVMP-BG heuristic vs PPDR algorithm

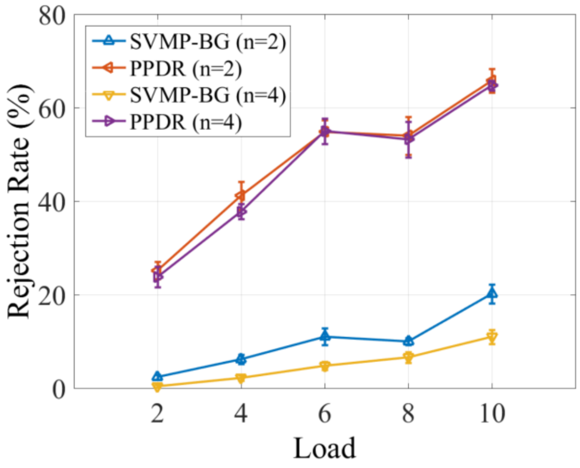

Fig.6: Rejection rate over load

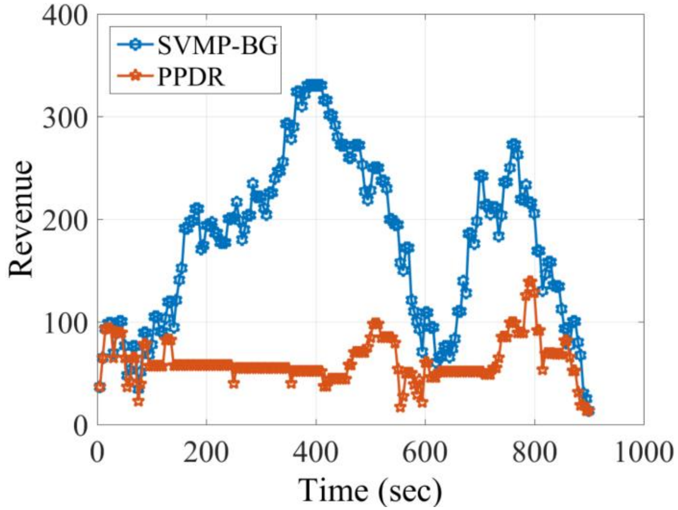

Fig.7: Revenue over time

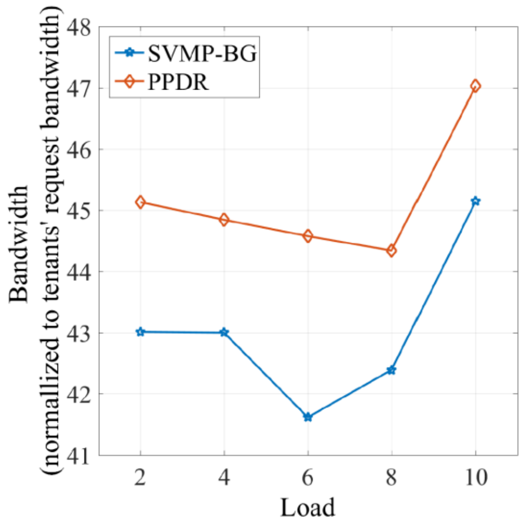

Fig.8: Average bandwidth over load

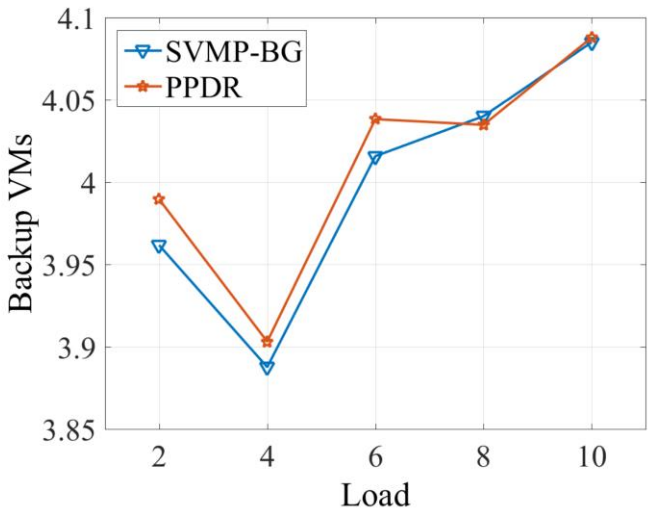

Fig.9: Average backup VMs over load

Conclusion

Conclusion

- Designing a protection plan for a cloud tenant able to provide both bandwidth guarantees and reliability is a complex problem
- Such protection plan can be achieved by the SVMP-BG model that leaves the choice for the cloud operator to realize the balance between bandwidth use and VMs consumption vs prioritizing one for another
- Given that the SVMP-BG model is NP-complete, we develop an SVMP-BG heuristic that balances the use of bandwidth and VMs
- SVMP-BG heuristic is proved to be much more scalable than the SVMP-BG model and outperforms the PPDR benchmark algorithm

Future Work

Future Work

- Explore the bandwidth required to synchronize the primary and backup VMs
- Studying our protection plan design using different primary embedding solutions (not based on collocation)
- Exploring further bandwidth saving opportunities

References

References

- [1] L. Chen et al. Allocating bandwidth in datacenter networks: A survey.Journal of Computer Science and Technology, 29(5):910–917, 2014.
- [2] Md G. Rabbani et al. On achieving high survivability in virtualized data centers. IEICE Transactions on Communications, 97(1):10–18, 2014.
- [3] M. Perlin. Downtime, outages and failures-understanding their true costs. Retrieved November, 25:2012, 2012.
- [4] J. Lee et al. Cloudmirror: Application-aware bandwidth reservations in the cloud. USENIX HotCloud, 20, 2013.
- [5] J. Lee et al. Application-driven bandwidth guarantees in datacenters. In SIGCOMM, pages 467–478. ACM, 2014.
- [6] H. Ballani et al. Towards predictable datacenter networks. In SIGCOMM, volume 41, pages 242–253. ACM, 2011.
- [7] H. Ballani et al. Chatty tenants and the cloud network sharing problem. In NSDI, pages 171–184, 2013.
- [8] J.C. Mogul et al. What we talk about when we talk about cloud network performance. SIGCOMM, 42(5):44–48, 2012.
- [9] Q. Zhang et al. Venice: Reliable virtual data center embedding in clouds. In INFOCOM, pages 289–297. 2014 Proceedings IEEE, 2014.
- [10] P. Bod´ık et al. Surviving failures in bandwidth-constrained datacenters. In SIGCOMM, pages 431–442. ACM, 2012.
- [11] C. Develder et al. Dimensioning backbone networks for multi-site data centers: exploiting anycast routing for resilience. In RNDM, pages 34 – 40. 2015 Proceedings IEEE, 2015.
- [12] P. Gill et al. Understanding network failures in data centers: measurement, analysis, and implications. In SIGCOMM, volume 41, pages 350–361. ACM, 2011.
- [13] H. Yu et al. Cost efficient design of survivable virtual infrastructure to recover from facility node failures. In ICC, pages 1–6. IEEE, 2011.
- [14] M. M. A. Khan, N. Shahriar, R. Ahmed, and R. Boutaba. SiMPLE: Survivability in Multi-Path Link Embedding. IEEE/ACM/IFIP International Conference on Network and Service Management (CNSM), Barcelona (Spain), November 9-13, 2015.

Questions?

THANK YOU!

UNIVERSITÉ

Concordia

UNIVERSITY

www.concordia.ca