CQNCR: Optimal VM Migration Planning in Cloud Data Centers

Presented By

Reaz Ahmed

David R. Cheriton School of Computer science
University of Waterloo

Joint work with Md. Faizul Bari, Mohamed Faten Zhani, Qi Zhang, and Raouf Boutaba

Outline

- Introduction
- Motivating Example
- Problem Formulation
- Proposed Heuristic
- Evaluation Results
- Conclusion & Future Work

VM Migration: Why do we need it?

- Virtualization has become a key enabling technology for cloud data centers
- Virtual Machine (VM) migration enables dynamic resource re-configuration for
  - Cost reduction
  - Maximizing utilization
  - Improving performance & reliability

Live Migration of VM

- Pre-copy live migration

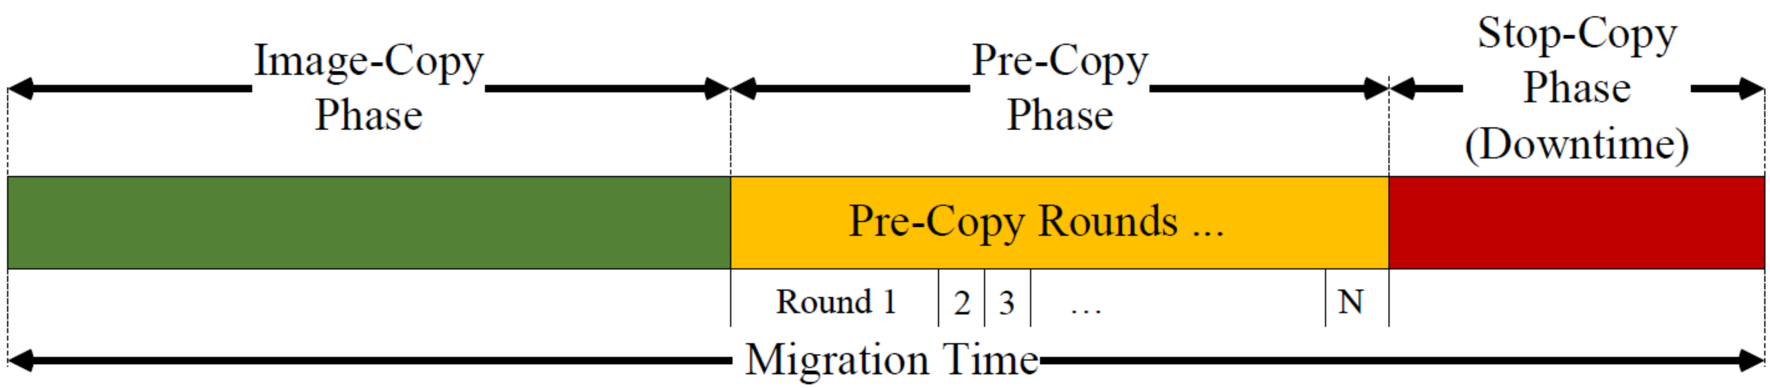

- **Image-Copy Phase:** the VM image is copied from the source to the destination
- **Pre-Copy Phase:** memory pages are copied periodically
- **Stop-Copy Phase:** the VM is paused and all remaining memory pages are transferred

VM Migration Sequencing Problem

- Depending on the objective, many VM migrations can be triggered at the same time
  - Service disruption
  - High resource consumption
  - Network congestion
  - Long provisioning time

- We need to find an optimal migration sequence
  - Reduce total migration time
  - Minimize service disruption time
  - Avoid server overload
  - Reduce network congestion

Outline

- Introduction
- Motivating Example
- Problem Formulation
- Proposed Heuristic
- Evaluation Results
- Conclusion & Future Work

Motivating Example

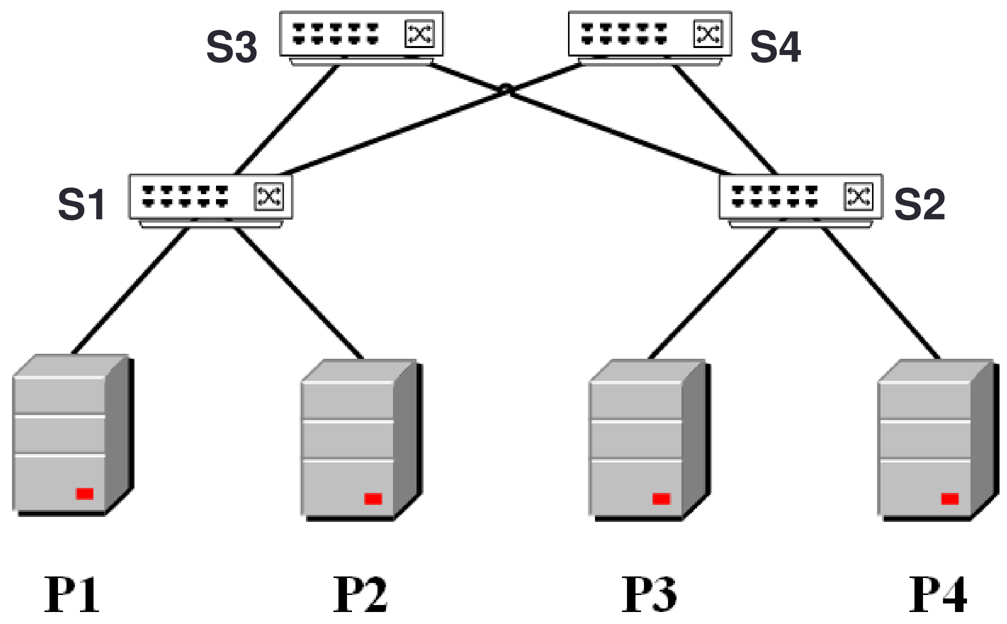

- Here we consider a network consisting of
  - 4 servers
  - 2 top-of-rack switches and
  - 2 aggregate switches

Motivating Example

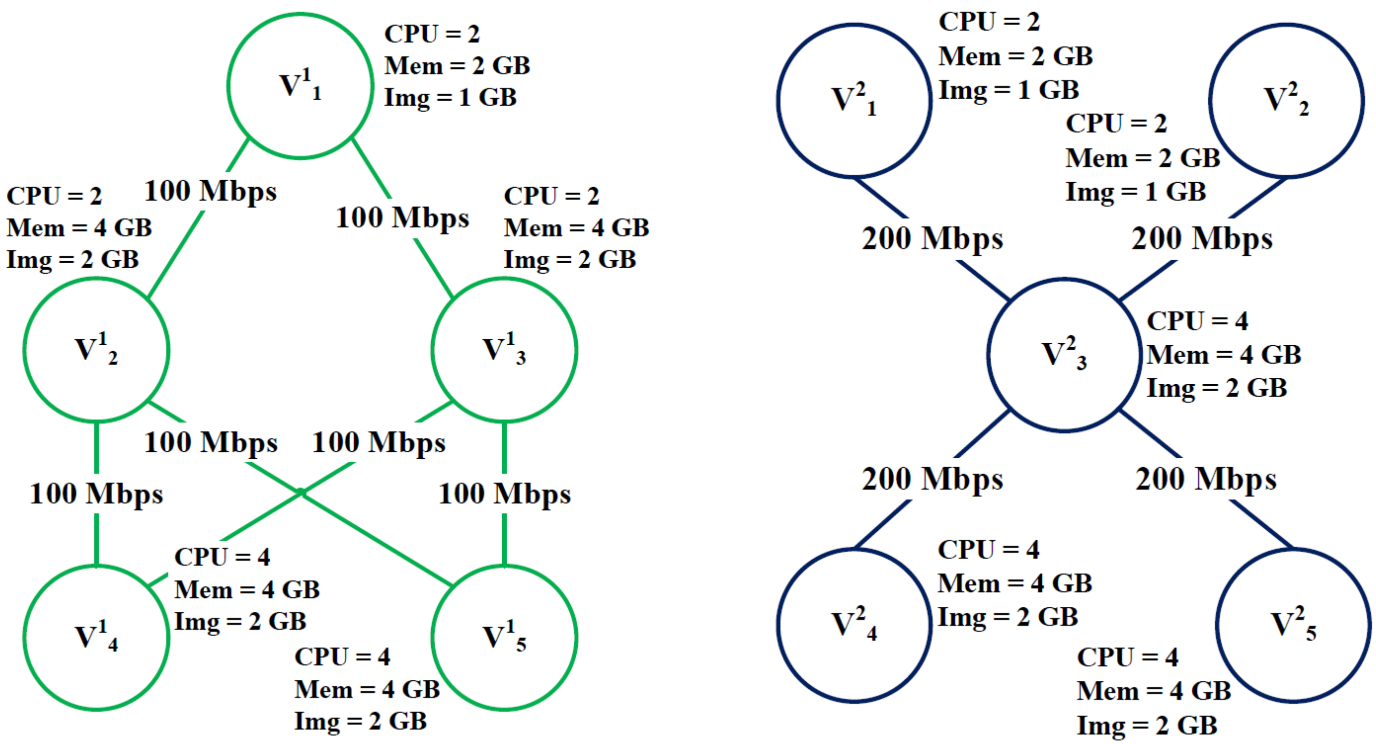

Example virtual networks to be deployed on the physical network

Motivating Example

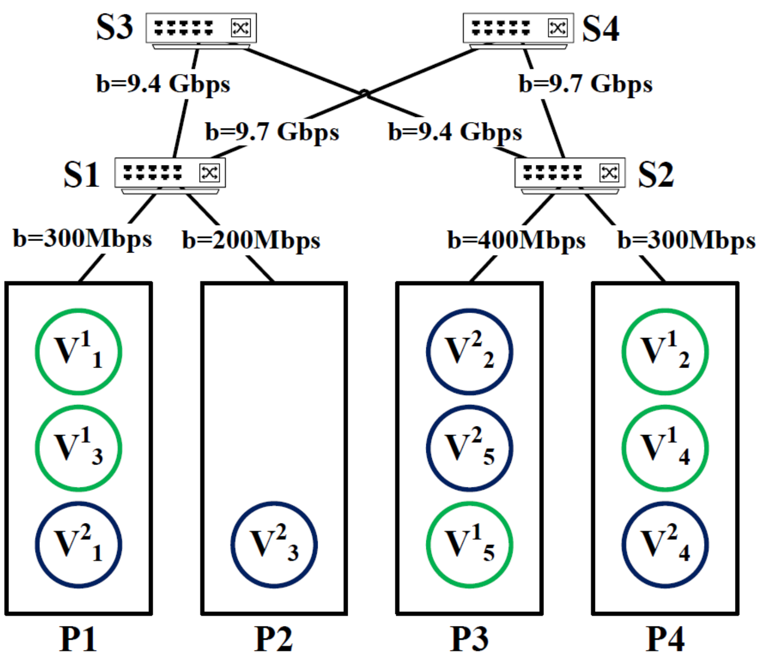

Initial mapping

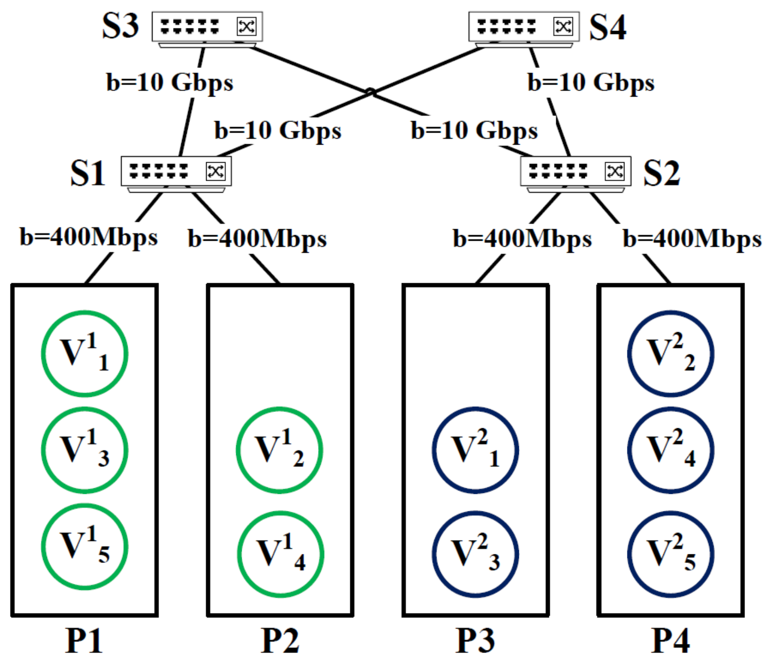

Final mapping

Motivating Example

Total migration time: 220 s
Average VM downtime: 4.14 s
Average VN downtime: 14.5 s

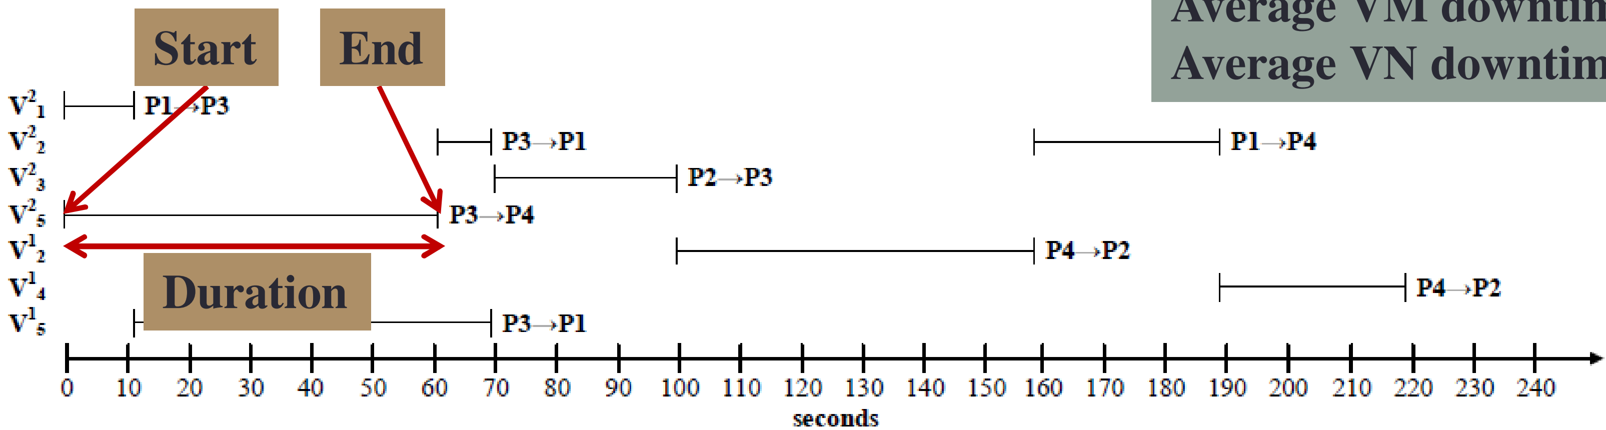

(a) Sequence 1

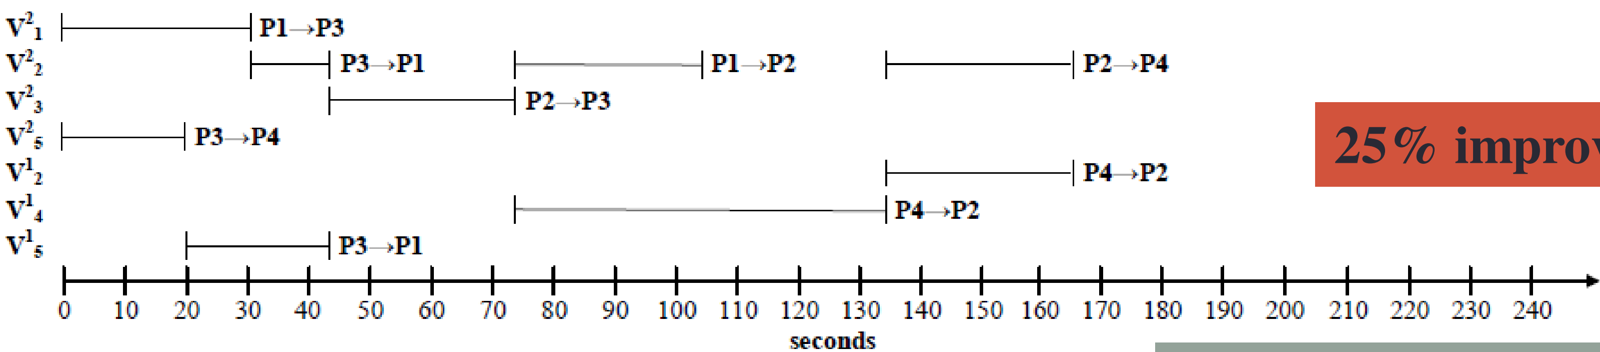

25% improvement

(b) Sequence 2

Total migration time: 165 s
Average VM downtime: 3.9 s
Average VN downtime: 13.7 s

Outline

- Introduction
- Motivating Example
- Problem Formulation
- Proposed Heuristic
- Evaluation Results
- Conclusion & Future Work

Problem Formulation

- The VM Migration Sequencing Problem can be formulated as an ILP
- Objective function

$$\min \left(\sum_{t=0}^{T} w^t + \delta \sum_{i=1}^{I} \sum_{n \in N^i} X_n^{it} P_n^i \right)$$

minimize **Migration time** **&** **Service down-time**

- Where
  - T is maximum allowable migration time
  - w^t represents whether migration occurs at time t
  - I is the number of VNs
  - N^i is the set of VMs present in VN i
  - X_n^{it} indicate whether n is under migration at time t
  - P_n^i represent the penalty (service downtime)

Problem Formulation

- Constraints
  - Only a single migration takes place for a single VM at any time instance
  - After a migration starts it cannot pause before completing
  - When a VM is under migration, it must be present in both the source and the destination
  - During the whole migration process capacity constraints for CPU, memory and disk must be satisfied.
- This problem generalizes the "Data Migration Problem" and hence it is *NP-Hard*.

Outline

- Introduction
- Motivating Example
- Problem Formulation
- Proposed Heuristic
- Evaluation Results
- Conclusion & Future Work

Proposed Heuristic

- Finding the migration sequence
  - Step 1: Find feasible migrations
  - Step 2: Group them in RIGs
  - Step 3: Compute migration cost
  - Step 4: Select the RIG to be migrated

Step 1: Find feasible migrations

- Take all the VMs that can be migrated at the current time

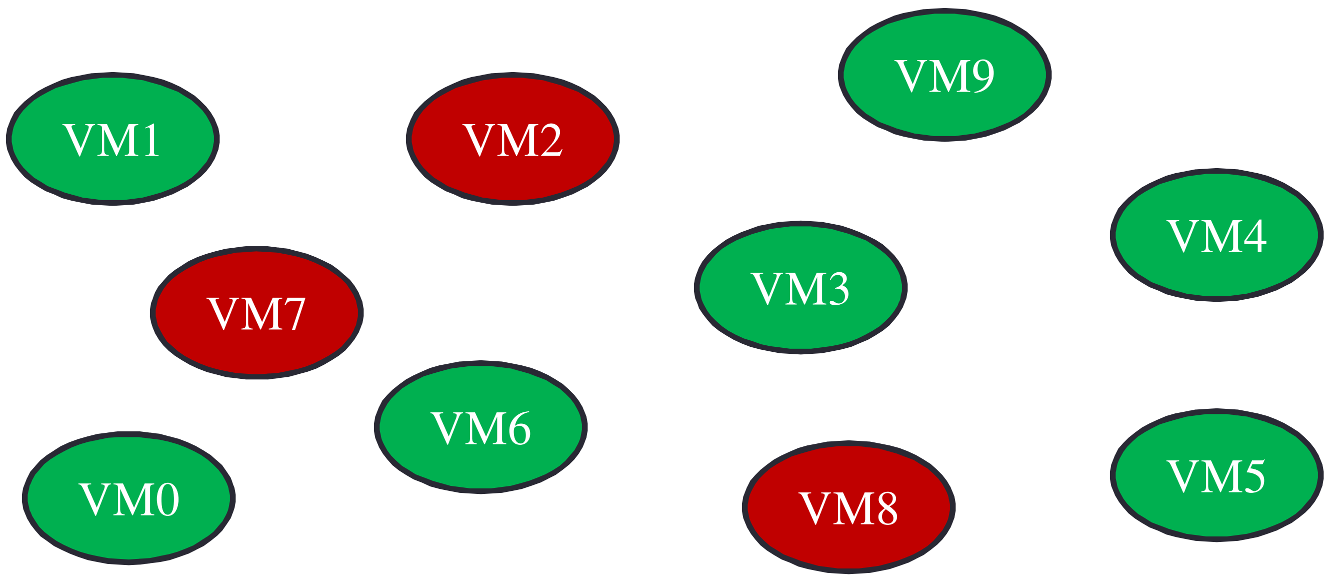

Green: feasible
Red: infeasible

Step 2: Group them in RIGs

- Group them into Resource Independent Groups (RIGs)
  - VMs in the same RIG
    - Can be migrated simultaneously using disjoin paths
    - Does not have the same source or destination

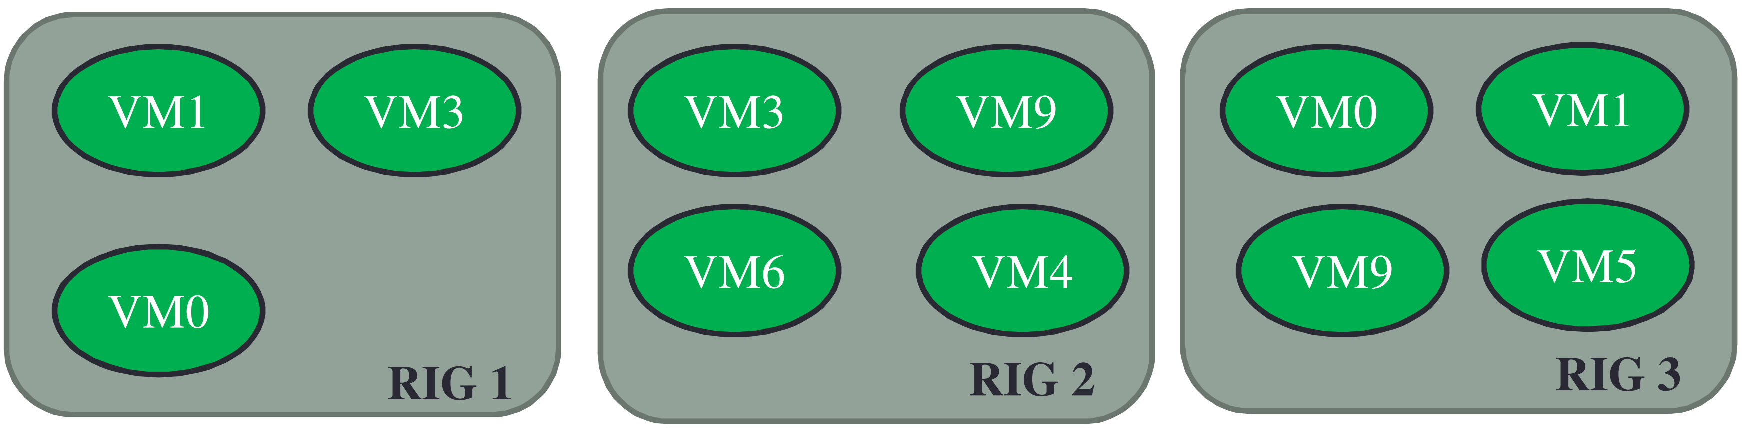

Step 3: Compute migration cost

- Compute migration time, wait time and impact-on-other-RIGs
  - Migration time (C_M): max {migration time of any VMs in a RIG}
  - Wait time (C_W): the time the scheduler has to wait before making the next decision
  - Impact on other RIGs (C_I): increase/decrease in migration time of other RIGs
- Total migration cost, $C_T = \alpha C_M + \beta C_W + \gamma C_I$

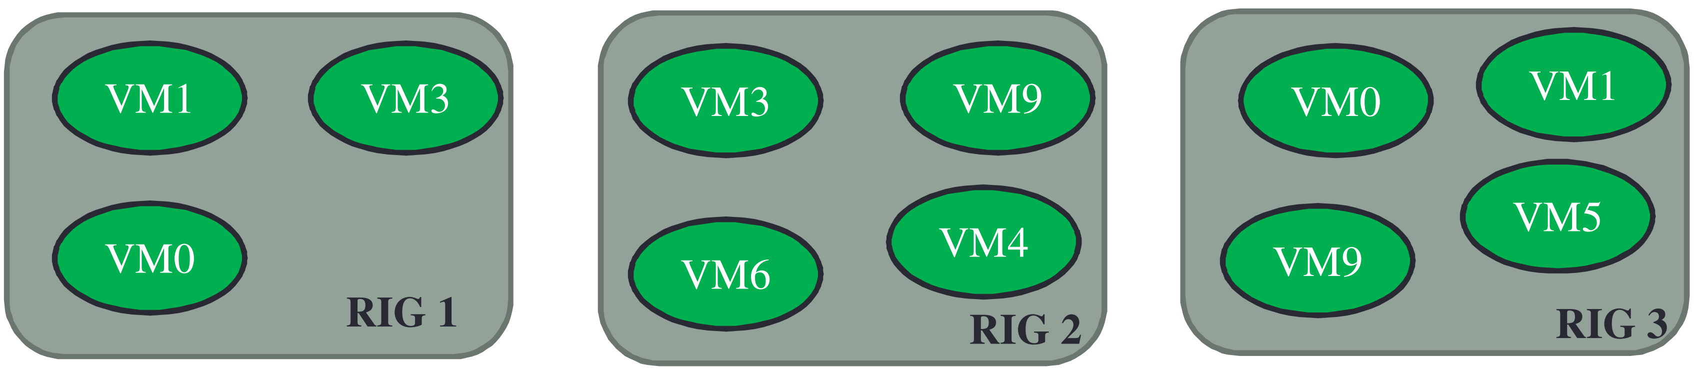

Step 4: Select the RIG to be migrated

- Select the RIG with the lowest C_T for migration at current time

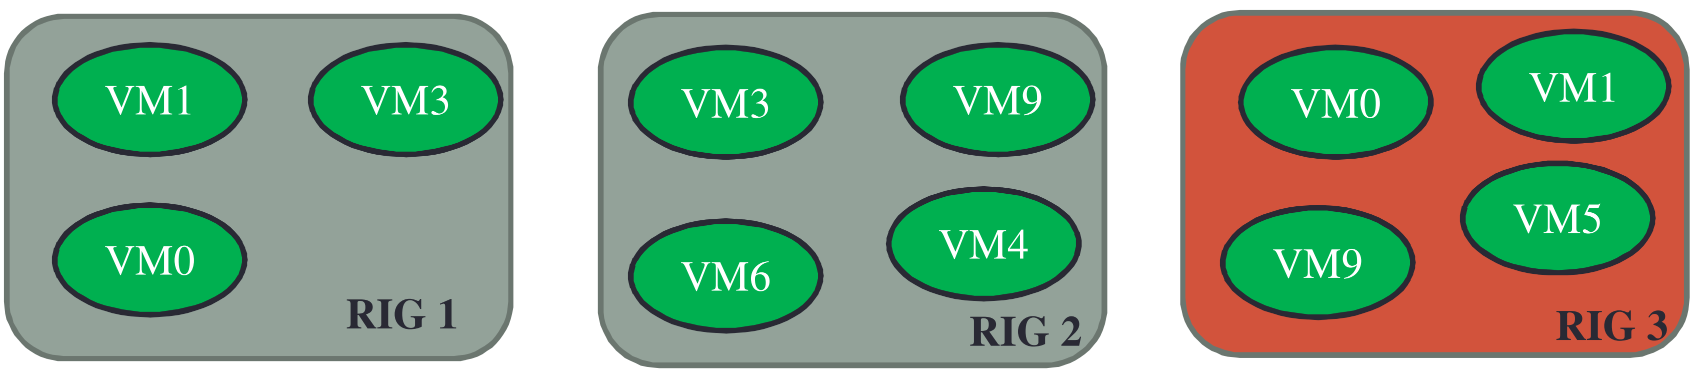

Outline

- Introduction
- Motivating Example
- Problem Formulation
- Proposed Heuristic
- Evaluation Results
- Conclusion & Future Work

Evaluation Results

- We generated 4 deployment scenarios

Scenarios	# of VNs	Initial mapping		Target mapping		# of Migrations
		Active Phy. Hosts	Active Phy. Links	Active Phy. Hosts	Active Phy. Links	
$S-1$	3	13	57	3	10	13
$S-2$	10	95	307	14	30	95
$S-3$	50	485	1483	75	194	486
$S-4$	100	935	2668	150	390	998

- VM specifications (CPU and memory) are taken from Amazon EC2 instances
- Bandwidth of virtual links are selected randomly between 250 and 750 Mbps

Evaluation Results

- We compare the performance of CQNCR with a baseline algorithm that
  - Chooses the VM with the shortest migration time first
  - Migrates multiple VMs simultaneously
  - Performs migrations using disjoint paths
- We compare the performance of CQNCR with the Baseline using 3 metrics
  - Total migration time
  - VM downtime
  - VN downtime (service disruption time)

Evaluation Results

- Total migration time

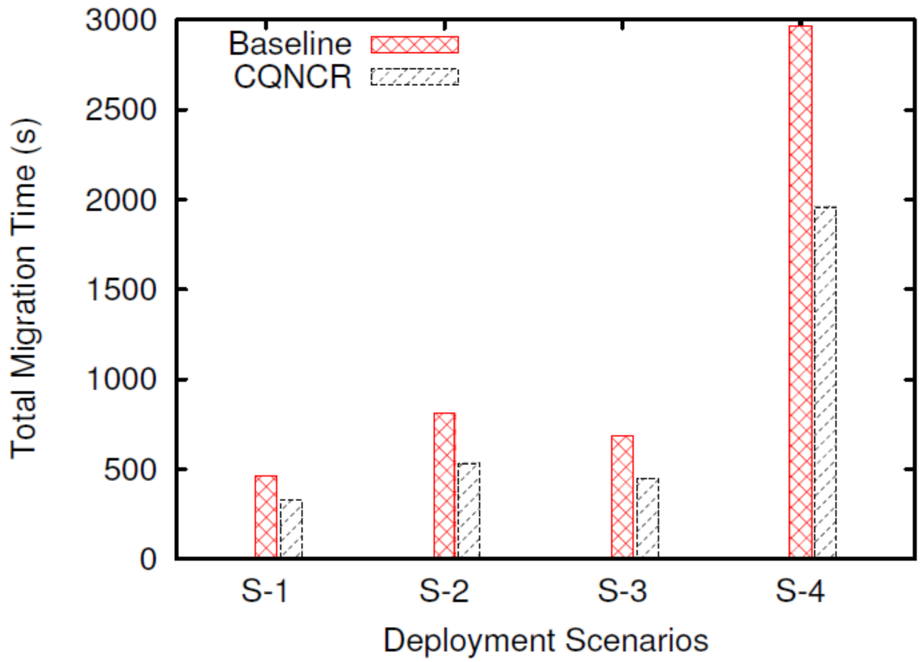

- CQNCR provides 35% improvement over the Baseline algorithm

Evaluation Results

- VM downtime

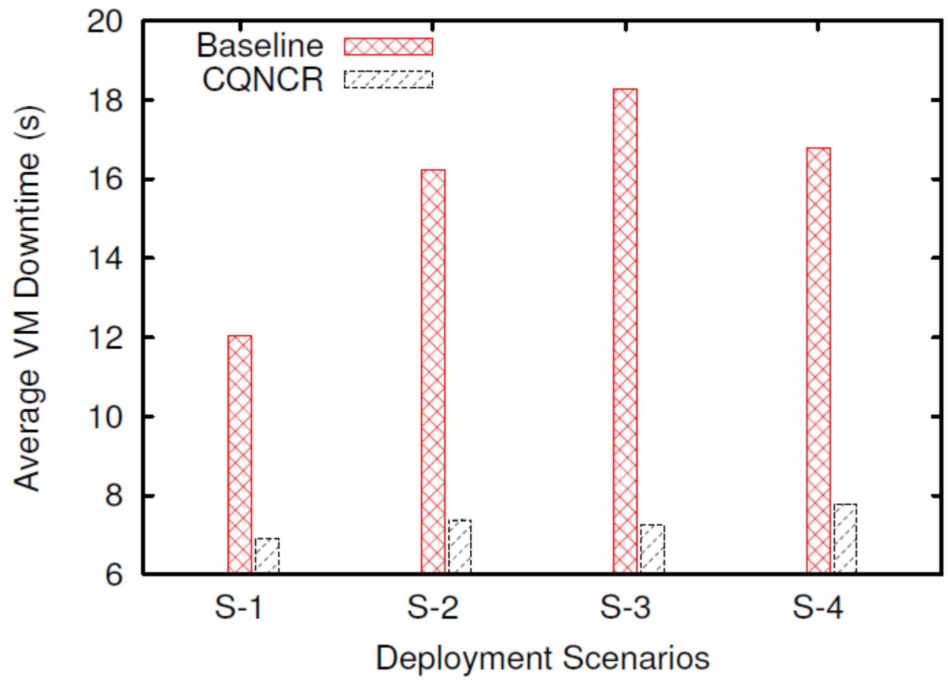

- CQNCR reduces VM downtime by up to 60%

Evaluation Results

- Virtual Network downtime

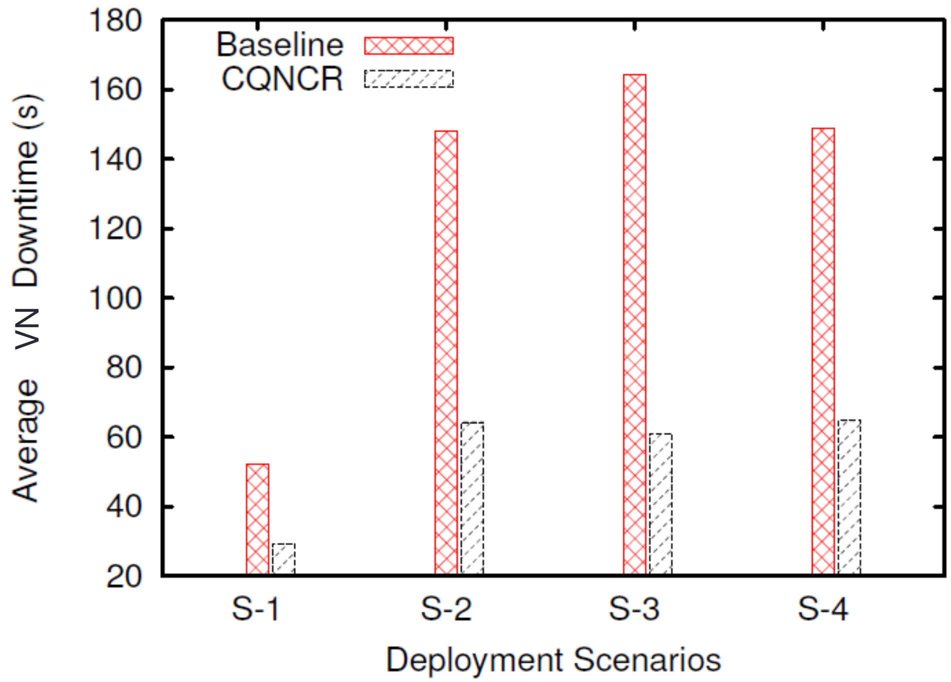

- CQNCR reduces VN downtime by up to 60%.

Outline

- Introduction
- Motivating Example
- Problem Formulation
- Proposed Heuristic
- Evaluation Results
- Conclusion & Future Work

Conclusion

- VM migration sequencing is an important problem as it has a direct impact on
  - Resource utilization
  - Service performance

- We proposed CQNCR for VM Migration Planning in Cloud data centers

- Benefits
  - Avoids network congestion
  - Reduce total migration time by up to 35%
  - Reduces VM/VN downtime by up to 60%

Future Work

- Assigning different priorities to different VNs
- Guarantee an upper bound on service downtime
- Consider the topology and workload of the VNs

Questions?

Backup Slides

Migration Graph & Infeasible Cycles

- Initial and final mappings

Virtual Machine	Physical Machine
VM1	P4
VM2	P3
VM3	P2
VM4	P2
VM5	P5
VM6	P1

Initial Mapping

Virtual Machine	Physical Machine
VM1	P5
VM2	P5
VM3	P5
VM4	P4
VM5	P3
VM6	P4

Final Mapping

Migration Graph & Infeasible Cycles

- Migration graph

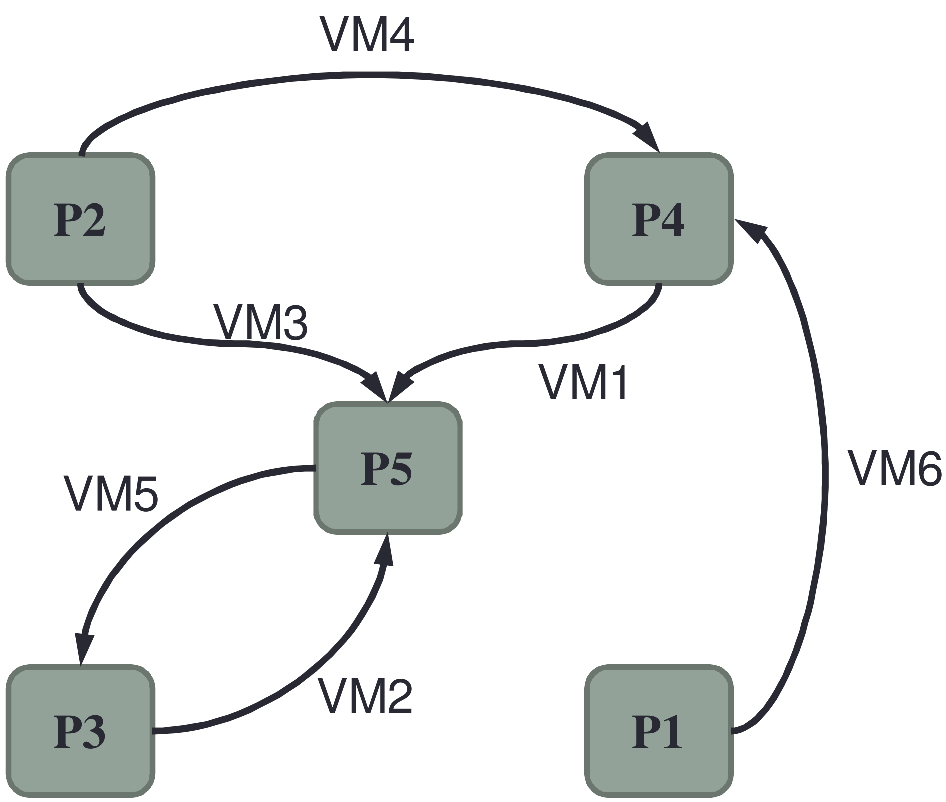

Migration Graph & Infeasible Cycles

- Infeasible cycle

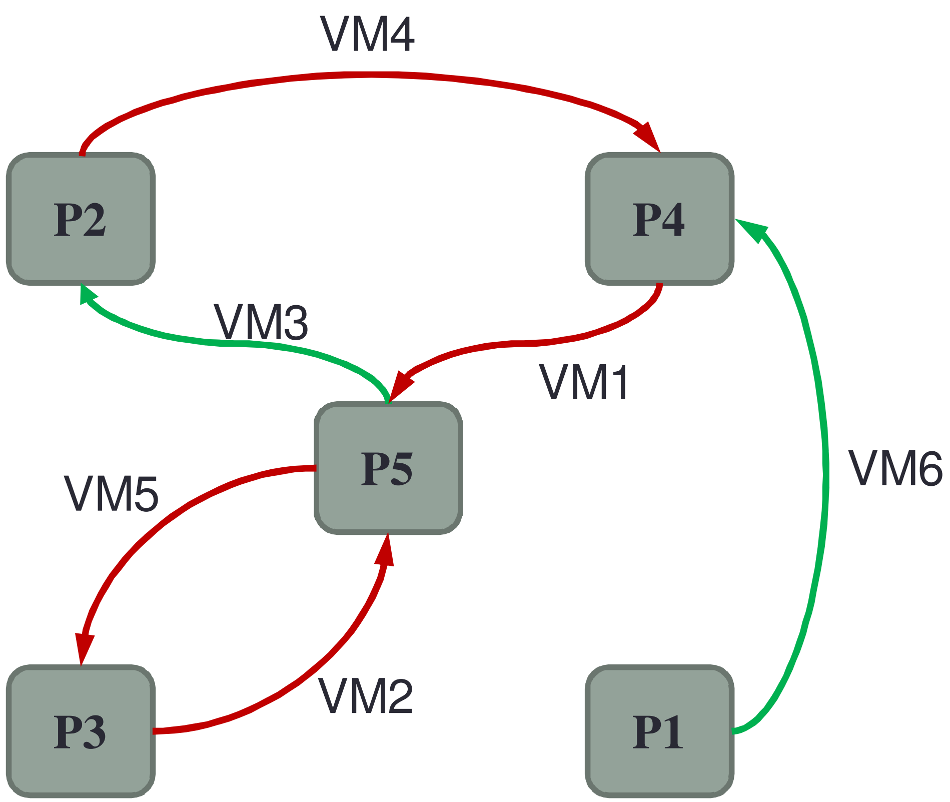

Migration Graph & Infeasible Cycles

- Breaking cycle with *pivot node*

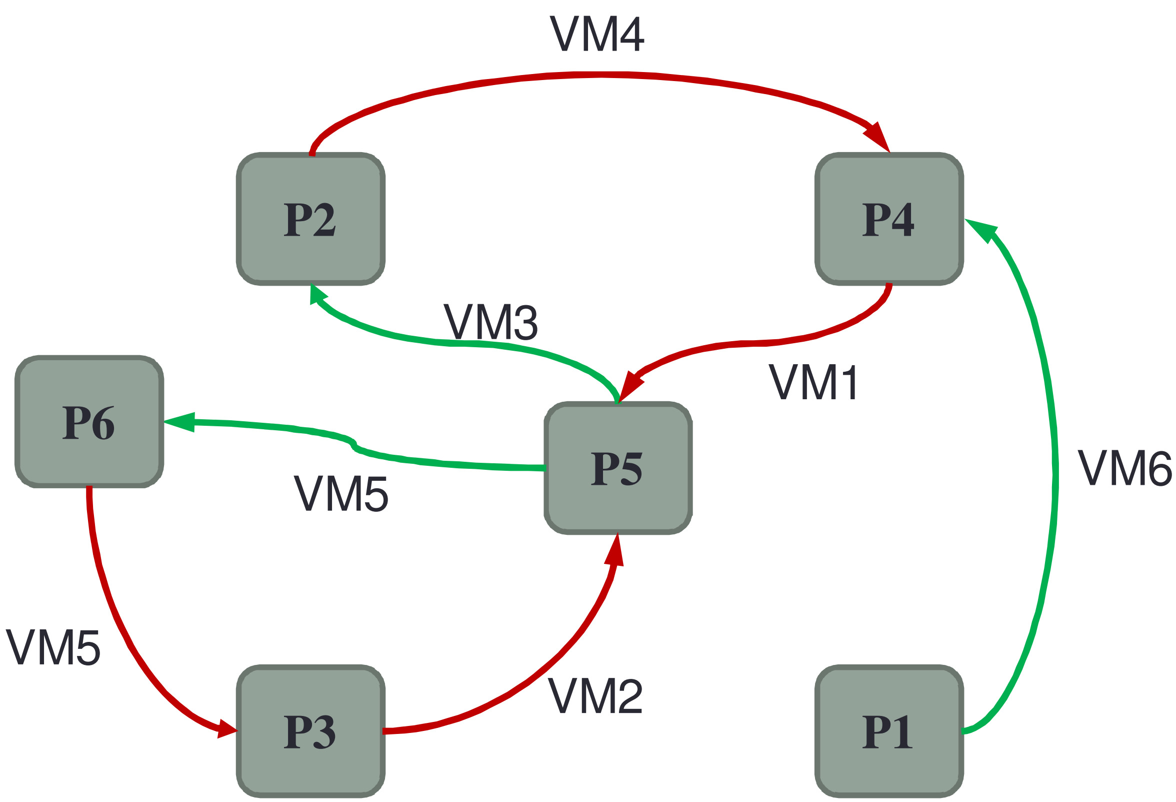

Evaluation Results

- VM specifications are taken from Amazon EC2 instances

EC2 Instance Type	vCPU	Memory (GB)
m3.xlarge	4	15
m3.2xlarge	8	30
m1.small	11	1.7
m1.medium	1	3.75
m1.large	2	7.5
m1.xlarge	4	15
c3.large	2	3.75
c3.xlarge	4	7
c3.2xlarge	8	15
c3.4xlarge	16	30
c1.medium	2	1.7
c1.xlarge	8	7
g2.2xlarge	8	15
cg1.4xlarge	16	22.5
m2.xlarge	2	17.1