Restoration in Optical Cloud Networks With Relocation and Services Differentiation[1]

Presenter: Yongcheng (Jeremy) Li

PhD student, School of Electronic and Information Engineering,

Soochow University, China

Email: liyongcheng621@163.com

Group Meeting, Friday, October 14, 2016

[1] C. N. da Silva, L. Wosinska, S. Spadaro *et al.*, "Restoration in optical cloud networks with relocation and services differentiation," *Journal of Optical Communications and Networking*, vol. 8, no. 2, pp. 100-111, Feb. 2016.

UCDAVIS

Outline

1. Background
2. Path restoration with service relocation and differentiation (PR-SRD) problem
3. ILP for Relocation With Priorities Model
4. Heuristic for Relocation With Priorities
5. Performance evaluation
6. Conclusions

Background

- With an ever-growing demand for cloud-based services, optical cloud concept is becoming increasingly popular
- Client nodes require access to *information technology (IT) resources* offered at selected data center (DC) nodes located in different geographical areas
  - Storage and computation services
- DC nodes are interconnected and made accessible to client nodes via *optical transport resources*
  - high-speed optical wavelength division multiplexing (WDM) fiber links
- *Resiliency* is a crucial aspect to consider
  - The availability of both transport and IT resources
  - At the occurrence of a failure, we need to guarantee the required availability levels for the cloud services already provisioned
    - Proactive protection-based techniques
    - Reactive restoration-based techniques

Background (contd.)

- Protection-based techniques
  - Pre-reserve backup resources during the provisioning phase
  - Guarantee 100% survivability against the failure scenarios
  - Translate into additional transport and IT resources
  - Affect the blocking probability for service requests
- Restoration-based techniques
  - Without pre-reserved backup resources and more flexible
  - Allow for lower deployment costs and/or better network blocking probability performance
  - Less than 100% survivability against failures
  - The restorability performance can be improved by leveraging on the anycast nature of cloud services
    - The portion of disrupted cloud services that can be successfully restored

UC DAVIS
UNIVERSITY OF CALIFORNIA

Background (contd.)

- Cloud service relocation
  - If the original DC node where the cloud service is running is not reachable anymore after a failure, the cloud service can be migrated to a different DC
  - Migration process is not instantaneous and needs to be used carefully to not significantly affect the cloud service downtime
  - Not be able to recover all the cloud services disrupted by a failure because of a lack of transport resources
- Cloud service differentiation
  - A selection of cloud services that should be given priority when using the existing spare wavelengths needs to be made.
    - Differentiate the cloud services based on their survivability requirements
- Investigating the benefits of applying both the service relocation and the service differentiation concepts while restoring optical cloud services

Path restoration with service relocation and differentiation (PR-SRD) problem

- Problem statement
  - Objective
    - Minimize the *average downtime* of the cloud services
      - The portion of the service holding time during which a cloud service is not available
  - Assumptions
    - Dynamic provisioning scenario requires continuous bandwidth
    - A single fiber link failure
    - Cloud services are divided into traffic classes, each one with a different priority value in order to reflect their importance
      - When competing for the same spare resources, cloud services belonging to a traffic class with a high priority value are given precedence over the ones with a low priority
    - If a cloud service cannot be successfully restored upon the occurrence of a failure, the cloud service is dropped

PR-SRD problem (contd.)

- *G(N,E)* represents optical transport network after a fiber link failure
- N_{DC} is the set of data center (DC) nodes, $N_{DC} \in N$
- $\mathrm{DC}_k \in N_{DC}, \mathrm{DC}_k^{st}$ available storage units, DC_k^{pu} available processing units
- Q is the set of cloud services disrupted by a failure that needs to be restored
- $Q_i \in Q$ requires Q_i^{st} storage, and Q_i^{pu} processing units, with Q_i^{st}, $Q_i^{pu} \in Z_+^*$
- The arrival and holding time values of each cloud service are represented by Q_i^{at} and Q_i^{ht}, respectively, with Q_i^{at}, $Q_i^{ht} \in R_+^*$
- The source node of Q_i is Q_i^{src}
- The DC node that was serving Q_i before the failure is Q_i^{dst}
- Let $\alpha_i = 1, 2, 3 \dots M$ with $M \in Z_+$ is the set of priority values (one for each traffic class) that can be assigned to the cloud service Q_i
  - The higher the value of α_i, the higher the importance of the cloud service

UC DAVIS
UNIVERSITY OF CALIFORNIA

PR-SRD problem (contd.)

- At the occurrence of a failure, for each $Q_i \in Q$ we define a quantity called remaining service time $Q_i^{rt} = Q_i^{ht} - (CT - Q_i^{at})$, Where CT represents the time at which the failure occurs
- Q_i^{rt} is normalized as follows:

RT = 10

$$t_i = \left\lceil 100 \times \frac{Q_i^{rt}}{\mathrm{RT}} \right\rceil, \qquad \mathrm{RT} = \max\left(Q_i^{rt}\right), Q_i \in Q$$

- The remaining service has a value t_i that is always within the [1,100] interval

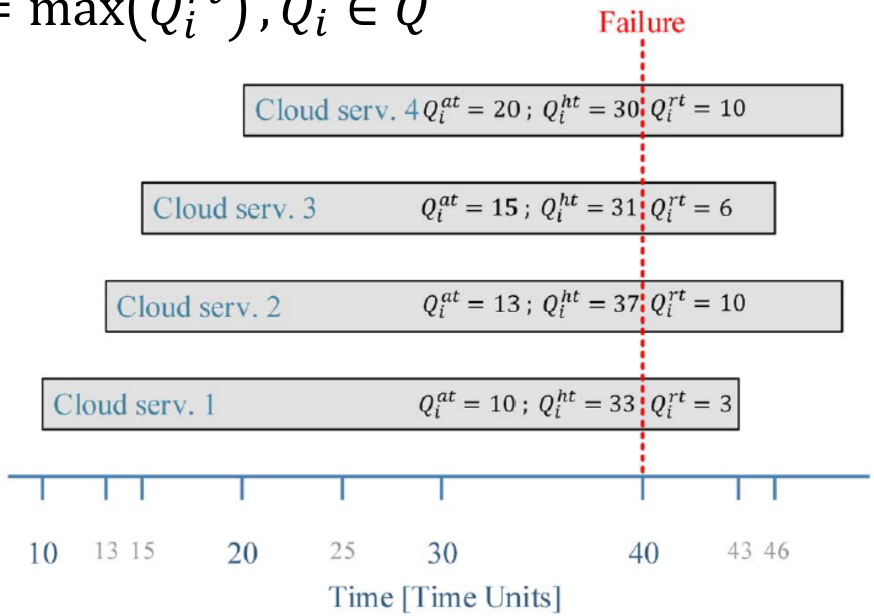

PR-SRD problem (contd.)

- $Q_{i,k}^{rd}$ is the relocation downtime for a given cloud service i relocated to DC k (i.e., $Q_{i,k}$)
  - Q_i^{st} is the number of storage units to be relocated
  - Δ is the rate at which storage units are transmitted
  - $d_{Q_i^{dst},k}$ is the distance between the DC node, where Q_i was being served before the failure and
  - Θ is the propagation time

$$Q_{i,k}^{\mathrm{rd}} = \frac{Q_i^{\mathrm{st}}}{\Delta} + \frac{d_{Q_i^{\mathrm{dst}},k}}{\Theta}, \quad \forall\, Q_i \in Q, \quad k \in N_{\mathrm{DC}}. \qquad (3)$$

- When a cloud service is relocated, the value of t_i is modified as follows

$$t_{i,k} = \begin{cases} \left\lceil 100 \times \frac{(Q_i^{\mathrm{rt}} - Q_{i,k}^{\mathrm{rd}})}{\mathrm{RT}} \right\rceil, & \text{if } Q_i^{\mathrm{rt}} - Q_{i,k}^{\mathrm{rd}} > 0 \\ 0, & \text{otherwise} \end{cases},$$

$$\mathrm{RT} = \max(Q_i^{\mathrm{rt}}), \quad \forall\, Q_i \in Q, \qquad k \in N_{\mathrm{DC}}. \qquad (4)$$

UC DAVIS
UNIVERSITY OF CALIFORNIA

PR-SRD problem (contd.)

- Inputs

- W_{xy}: the number of free wavelengths on fiber link $(x,y) \in E$;
- $\alpha_i \in A$: the set of priority values of $Q_i \in Q$;
- t_i: normalized value of Q_i^{rt}, when $Q_i \in Q$ is restored but not relocated, calculated according to Eq. (2);
- $t_{i,k}$: normalized value of Q_i^{rt}, when $Q_i \in Q$ is relocated to DC $k \in N_{\mathrm{DC}}$, calculated according to Eq. (4);
- Q_i^{dst}: the node serving $Q_i \in Q$ before the failure;
- Q_i^{st}: storage units required by $Q_i \in Q$;
- Q_i^{pu}: processing units required by $Q_i \in Q$.

- Variables:

- wl_{xy}: the total number of wavelengths used by the restoration paths on fiber link $(x,y) \in E$;
- $wl_{xy}^i \in \{0,1\}$: equal to 1 if the restoration path of cloud service Q_i traverses fiber link (x,y), 0 otherwise;
- $A_i \in \{0,1\}$: equal to 1 if the service Q_i is successfully restored, 0 otherwise;
- $A_{i,k} \in \{0,1\}$: equal to 1 if cloud service Q_i is successfully restored using the DC at node $k \in N_{\mathrm{DC}}$.

PR-SRD problem (contd.)

$$\min \sum_{Q_i \in Q} \left[\alpha_i \left(t_i - \sum_{k \in N_{\mathrm{DC}}} t_{i,k} \times A_{i,k} \right) \right] + \beta \sum_{Q_i \in Q} \sum_{k \in N_{\mathrm{DC}} | k \neq Q_i^{\mathrm{dst}}} A_{i,k} + \gamma \sum_{(x,y)} wl_{x,y}. \quad (5)$$

The first one is the sum of the downtime of all the cloud services that cannot be restored

The second term counts the number of cloud services that needed relocation while being restored

The third term accounts for the number of wavelength links used by all the restoration paths

$$\sum_{\forall n \in N} wl^i_{nj} - \sum_{\forall m \in N} wl^i_{jm} = \begin{cases} -A_i, & \text{if } j = Q_i^{\mathrm{src}} \\ A_{i,j}, & \text{if } j \in N_{\mathrm{DC}} \\ 0, & \text{otherwise} \end{cases}, \quad \forall Q_i \in Q, \quad \forall j \in N, \quad (6)$$

Constraint (6) guarantees the flow conservation of each restored cloud service

$$wl_{xy} = \sum_{\forall Q_i} wl^i_{xy}, \quad \forall (x,y) \in E, \quad (7)$$

Constraint (7) computes the total number of wavelengths used on each fiber link for restoration purposes.

$$wl_{xy} \leq W_{xy}, \quad \forall (x,y) \in E, \quad (8)$$

Constraint (8) ensures that the number of wavelengths used on each fiber link does not exceed the actual number of available wavelengths

$$\sum_{\forall k} A_{i,k} \leq 1, \quad \forall Q_i \in Q, \quad (9)$$

Constraint (9) checks that each relocated cloud service uses at most one DC

$$A_{i,k} = 0, \quad \forall k \in N_{\mathrm{DC}}, \quad \forall Q_i \in Q | k \neq Q_i^{\mathrm{dst}} \wedge Q_i^{\mathrm{rt}} \leq Q_i^{\mathrm{rd}}, \quad (10)$$

Constraint (10) ensures that a cloud service cannot be relocated to DC node k if the relocation downtime to that DC node is larger than the remaining service time

PR-SRD problem (contd.)

$$\sum_{\forall Q_i \in Q} (Q_i^{\mathrm{st}} \times A_{i,k}) \leq \mathrm{DC}_k^{\mathrm{st}}, \quad \forall\, k \in N_{\mathrm{DC}}, \qquad (11)$$

$$\sum_{\forall Q_i \in Q} (Q_i^{\mathrm{pu}} \times A_{i,k}) \leq \mathrm{DC}_k^{\mathrm{pu}}, \quad \forall\, k \in N_{\mathrm{DC}}, \qquad (12)$$

Constraints (11) and (12) ensure that a DC node cannot be used to relocate a cloud service if it does not have enough IT resources (either storage or processing units) to accommodate it

$$t_{i,k} = \begin{cases} t_i, & \text{if } k = Q_i^{\mathrm{dst}} \\ \left\lceil 100 \times \frac{(Q_i^{\mathrm{rt}} - Q_{i,k}^{\mathrm{rd}})}{\mathrm{RT}} \right\rceil, & \text{if } Q_i^{\mathrm{rt}} - Q_{i,k}^{\mathrm{rd}} > 0, \\ 0, & \text{otherwise} \end{cases}$$

$$\mathrm{RT} = \max(Q_i^{\mathrm{rt}}), \quad \forall\, Q_i \in Q, \qquad k \in N_{\mathrm{DC}}. \qquad (13)$$

Constraint (13) computes the normalized value of the remaining service time of $Q_i \in Q$

UC DAVIS
UNIVERSITY OF CALIFORNIA

Heuristic for Relocation With Priorities

Algorithm 1 HRP

$Q' = Q$ sorted by Q_i^{wt}
for all $Q_i \in Q'$ **do**
  curRoute = selRoute = selDC = NULL
  if $shortestPath\ (Q_i^{src}, Q_i^{dst}) \neq NULL$ **then**
    $selRoute = shortestPath(\ Q_i^{src}, Q_i^{dst})$
    $restorePath\ (Q_i, selRoute)$
  else
    for all $DC_k \in N_{DC} | Q_i^{rt} > Q_{i,k}^{rd} \wedge DC_k^{st} \geq Q_i^{st} \wedge C_k^{pu} \geq Q_i^{pu}$ **do**
      curRoute = $shortestPath\ (Q_i^{src}, Q_i^{dst})$
      if $hopCount$(curRoute) < $hopCount$(selRoute) **then**
        selRoute = curRoute
      end if
    end for
    if selRoute ≠ NULL **then**
      $relocateAndRestorePath(\ Q_i, selRoute)$
    else
      $dropService(Q_i)$
    end if
  end if
end for

The objective of the heuristic is the same as that of the IRP,

average downtime value of all unrestored cloud services

number of successfully restored cloud services that needed to be relocated

number of wavelength links used in the restoration process.

Performance evaluation

- Topologies

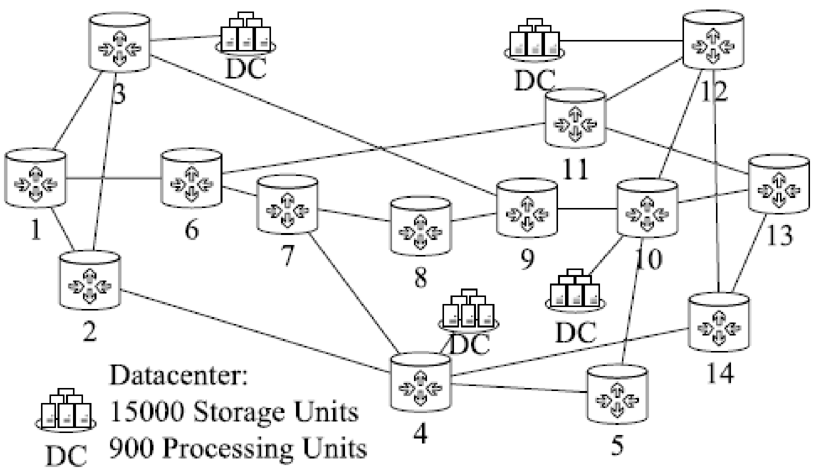

NSFNET

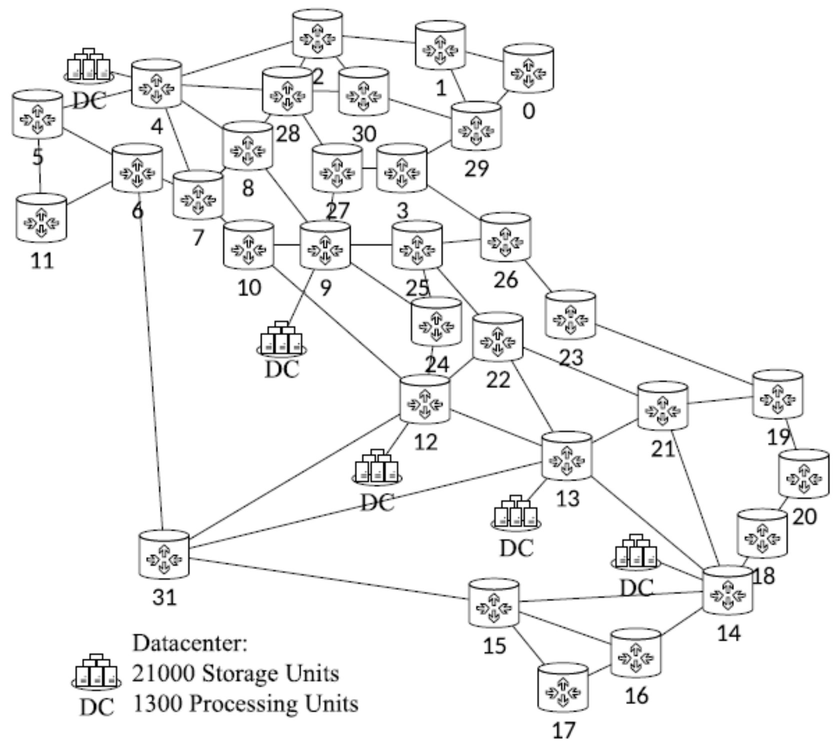

Italian network topology

Performance evaluation

- Simulation Scenarios
  - Scenario A
    - For NSFNET
    - All cloud services belong to the same traffic class
  - Scenario B and C
    - Two different traffic classes
      - Up to 20% of the cloud services in each experiment have high priority
      - The remaining part has low priority
    - B use NSFNET, and C uses Italian network topology
  - All fiber links in the network are bidirectional, with 80 wavelengths in each direction
  - DCs are assumed to be co-located with the network nodes to which they are connected
  - Network nodes have full wavelength conversion capability

Performance evaluation (contd.)

- Each simulation experiment consists of establishing 1 million cloud services
- Each service from a client to a DC node has enough storage and processing resources to accommodate cloud service requirements.
- The amounts of storage and processing units required by each cloud service are chosen uniformly in the intervals [1,100] and [1,5]
- Connecting a client node to a DC node requires the establishment of a lightpath with a capacity equal to the capacity of one wavelength channel
- The holding time of each cloud service is exponentially distributed with an average value of 60 time units
- The arrival rate of the cloud services follows a Poisson distribution,

Performance evaluation (contd.)

- The client node at which a cloud service originates is uniformly selected among all non-DC network nodes
- Transponders in the optical networks work at 100 Gbps, with storage units of approximately 1.3 Gb in size
- Δ is equal to 100 [storage units/s]
- Θ is assumed to be equal to 2×10^5
- All the results are the average of 100 different experiments

Performance evaluation (contd.)

- Blocking probability
  - The ratio between the number of cloud services that could not be successfully provisioned in the network and the total number of service requests
- average availability
  - The ratio between the sum of the uptime of all the provisioned cloud services and the sum of their service holding time values
- Average restorability
  - The ratio between the number of cloud services that were successfully restored and the number of cloud services disrupted by a failure
- Average relocations
  - The ratio between the number of restored cloud services that required relocation and the number of cloud services that were successfully restored

Performance evaluation (contd.)

- Scenario A

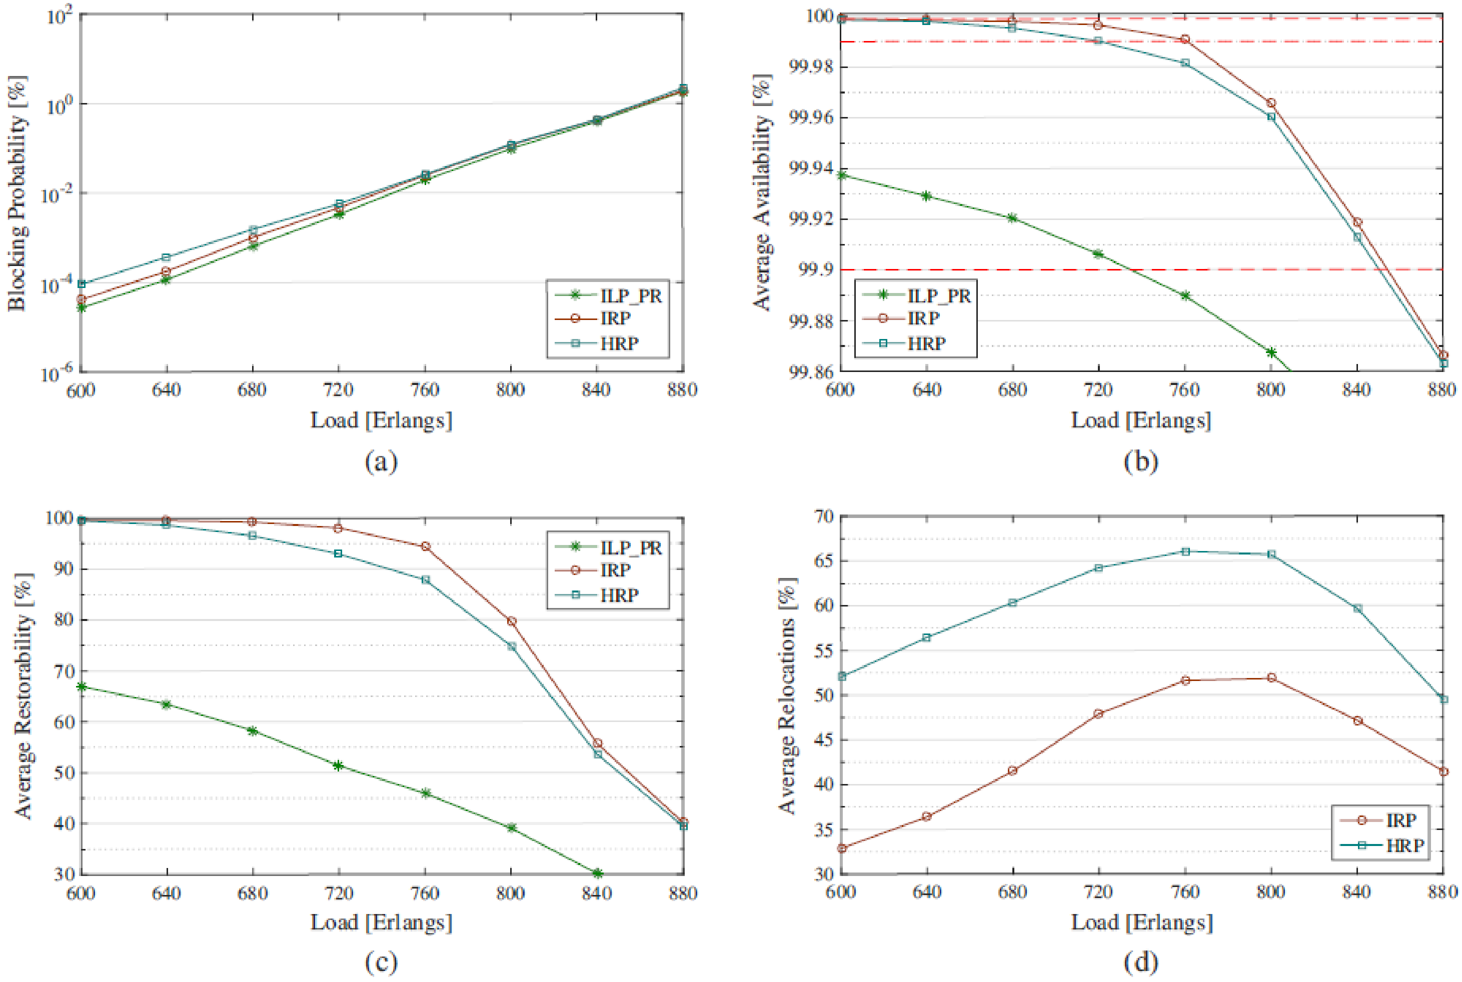

UCDAVIS UNIVERSITY OF CALIFORNIA

Performance evaluation (contd.)

- Scenario B

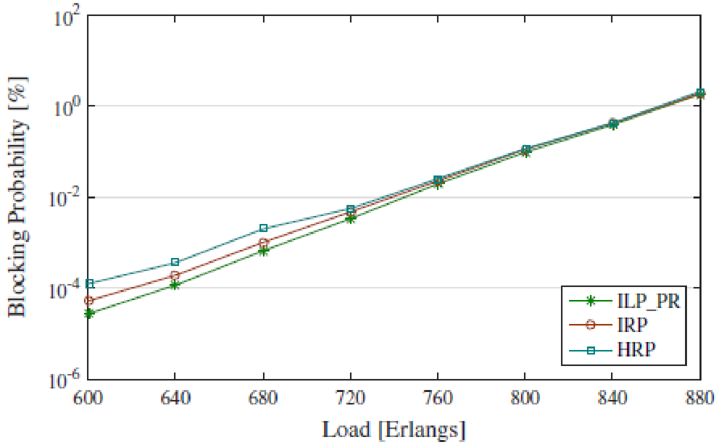

(a)

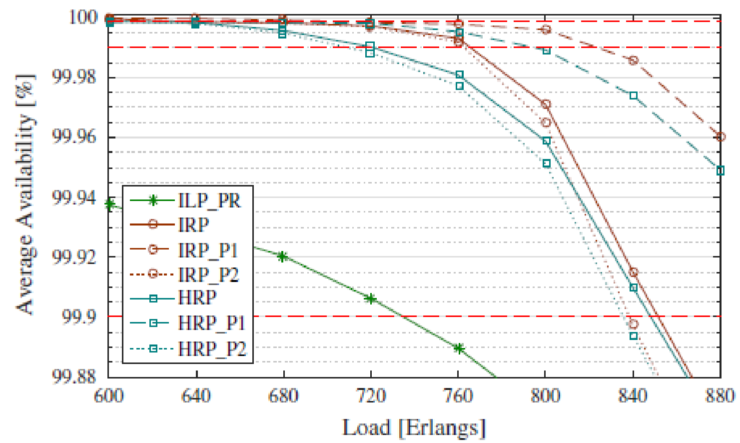

(b)

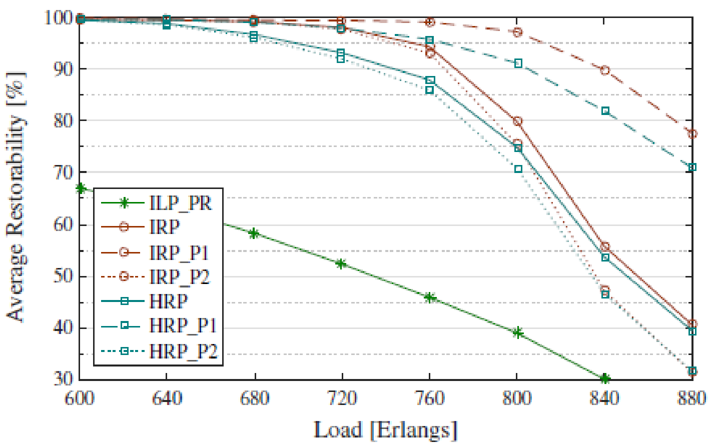

(c)

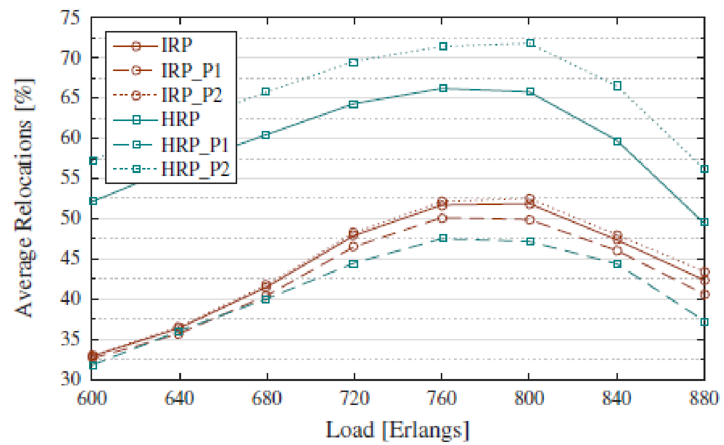

(d)

Performance evaluation (contd.)

- Scenario C

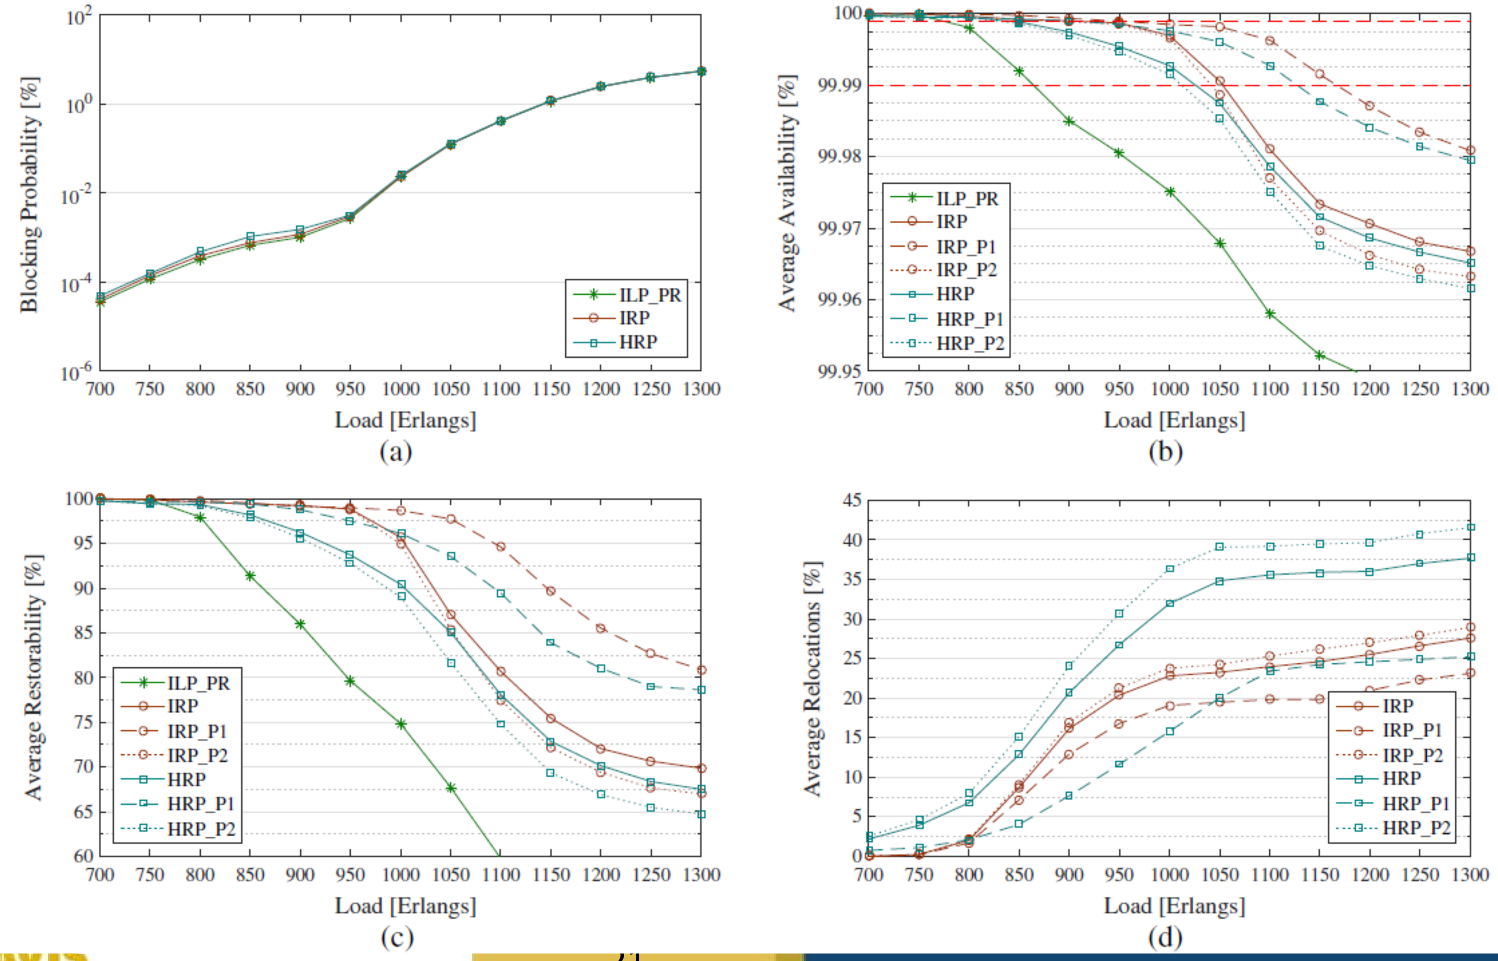

UC DAVIS
UNIVERSITY OF CALIFORNIA

Conclusion

- Proposes a restoration-based survivability strategy to recover cloud services disrupted by fiber link failures
- Proposed an ILP formulation model (IRP) and a heuristic algorithm (HRP)
- HRP provides results very close to IRP with a significantly lower processing time
- Both IRP and HRP are able to improve the average service availability and restorability performance with a limited number of cloud service relocations when compared to conventional restoration based techniques
- The availability and restorability performance of critical cloud services are very close to those achievable with a protection-based strategy, but with the inherent benefits in terms of efficient resource usage deriving from a restoration-based approach

Thank you for your attention!

Presenter: Yongcheng (Jeremy) Li

PhD student, School of Electronic and Information Engineering,

Soochow University, China

Email: liyongcheng621@163.com

Group Meeting, Friday, October 14, 2016

UCDAVIS
UNIVERSITY OF CALIFORNIA