Content Accessibility in Optical Cloud Networks Under Targeted Link Cuts

Carlos Natalino[1], Aysegul Yayimli[2], Lena Wosinska[1], Marija Furdek[1]

[1] KTH Royal Institute of Technology, Stockholm, Sweden
[2] Istanbul Technical University, Turkey

marifur@kth.se

RECODIS
Resilient communication services protecting end-user applications from disaster-based failures

ONDM 2017, Budapest, Hungary, May 15-17, 2017

COST Action 15127

Resilient Communication Services Protecting End-User Applications from Disaster-Based Failures

RECODIS

Resilient communication services protecting end-user applications from disaster-based failures

- WG 1: Large-scale natural disasters
- WG 2: Weather-based disruptions
- WG 3: Technology-related disruptions
- WG 4: Malicious human activities
    - How to quantify network vulnerability to attacks?
    - How to measure the level of difficulty for an attacker to affect the network?

Outline

- Introduction
- Content Delivery Networks
- Gauging CDN Robustness
  - Average Two-Terminal Reliability
  - Average Content Accessibility
- Simulation results
- Conclusions

Introduction

- Immense growth of the amount and variety of network traffic[1]
- Intensive growth of data center traffic and cloud computing[2]
  - Annual global data center traffic will reach 10.4 zettabytes by 2019
  - More than 86% of workload will be processed by cloud data centers

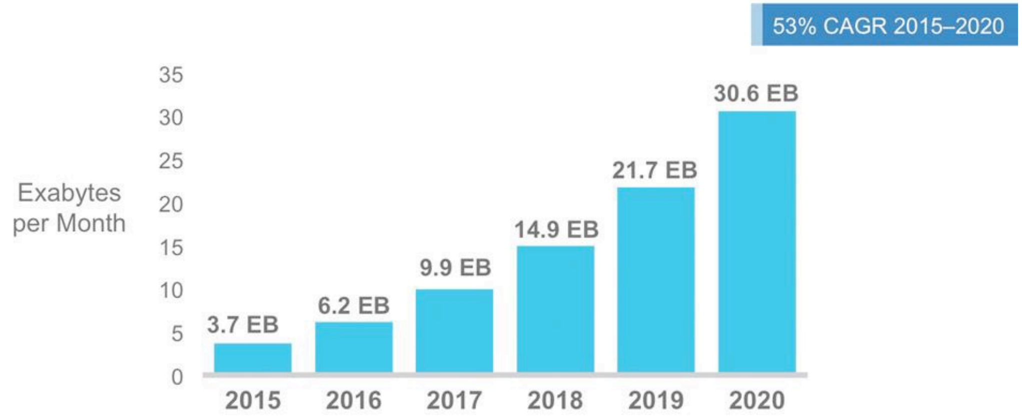

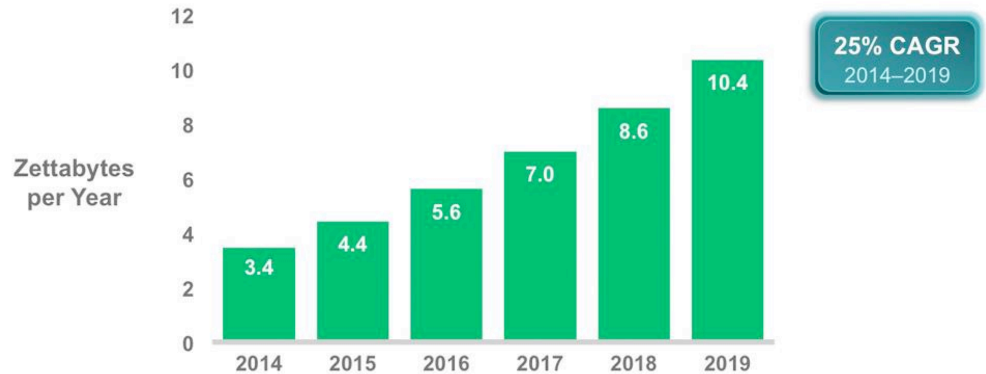

[1] Cisco Visual Networking Index: Global mobile data traffic forecast update, 2015-2020
[2] Cisco Global Cloud Index: Forecast and methodology, 2014-2019

Content Delivery Networks (CDNs)

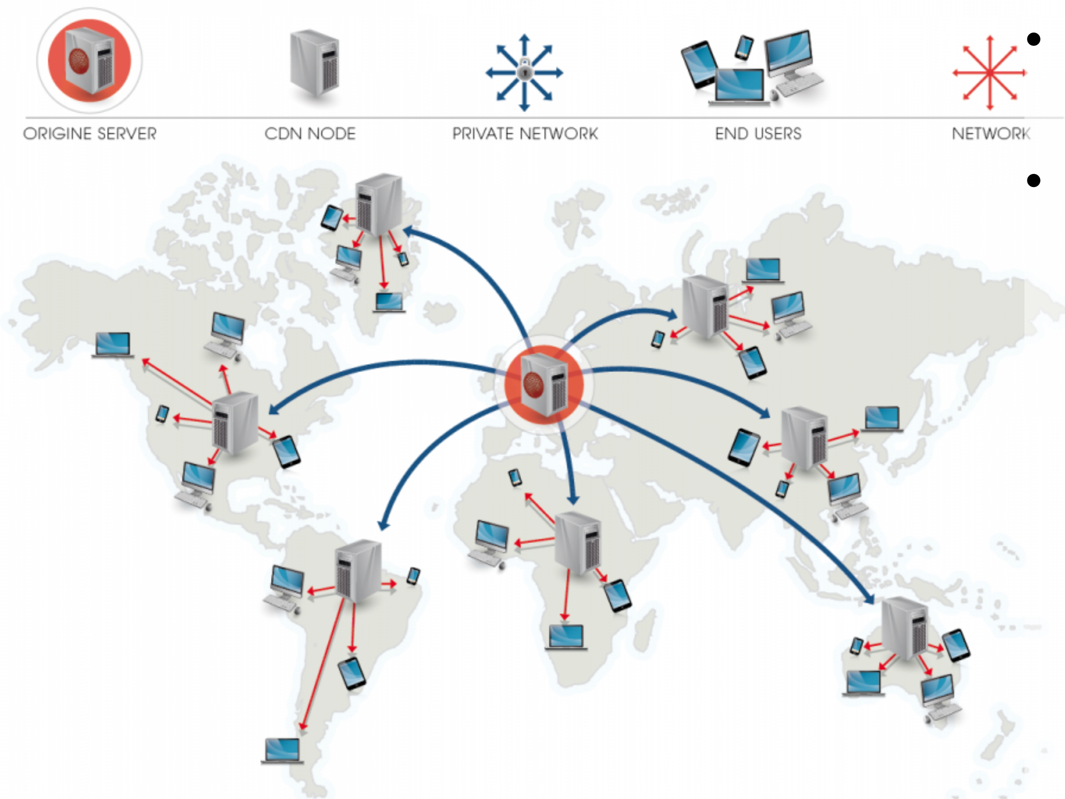

- Content is replicated over a set of data centers
- Users can connect to any replica (**anycast**)
    - Lower latency
    - More efficient network resource usage
    - Higher availability and accessibility
    - Inherently higher robustness

Robustness of CDNs

- CDNs are vulnerable to a wide range of physical-layer attacks aimed at service degradation
- **Link cut attacks**
  - Relatively low level of sophistication
  - Can cause outright service interruption
  - Efficiency of attacks is boosted by targeting the most critical links

A series of attacks on fiber network in San Francisco area
- Investigated by FBI
- $250.000 award offered by AT&T for information

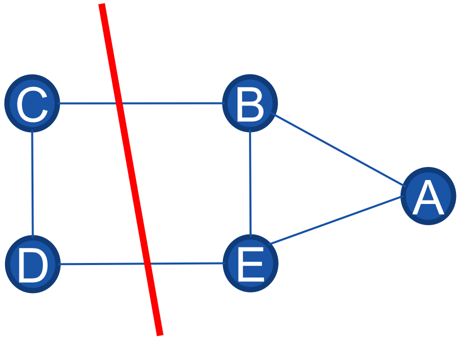

By cutting only 2 links, the network is partitioned

Robustness of CDNs

- How to model the effects of link cut attacks in the anycast traffic scenarios?
- **Content accessibility**: the ability of a region in the network topology (e.g., a set of users connected to an aggregation node) to access a particular content that is replicated over a number of nodes
  - Depends on the replica placement and the link cut set

Degree centrality: determined by the nodal degree

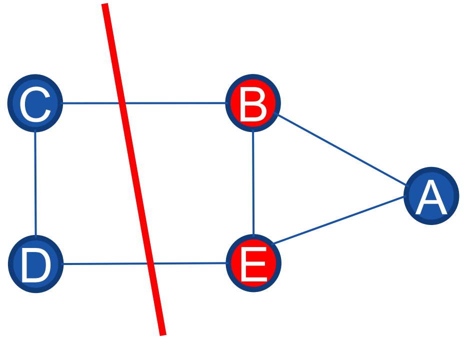

No content accessibility for nodes C and D

Robustness of CDNs

- How to model the effects of link cut attacks in the anycast traffic scenarios?
- **Content accessibility**: the ability of a region in the network topology (e.g., a set of users connected to an aggregation node) to access a particular content that is replicated over a number of nodes
  - Depends on the replica placement and the link cut set

Betweeness centrality: equal to the number of shortest paths traversing the element

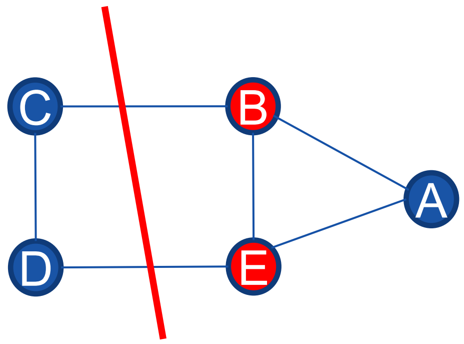

No content accessibility for nodes C and D

Robustness of CDNs

- How to model the effects of link cut attacks in the anycast traffic scenarios?
- **Content accessibility**: the ability of a region in the network topology (e.g., a set of users connected to an aggregation node) to access a particular content that is replicated over a number of nodes
  - Depends on the replica placement and the link cut set

Closeness centrality:
Based on the average distance to all other nodes

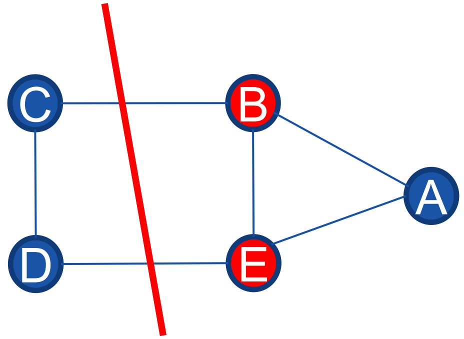

No content accessibility for nodes C and D

Robustness of CDNs

- How to model the effects of link cut attacks in the anycast traffic scenarios?
- **Content accessibility**: the ability of a region in the network topology (e.g., a set of users connected to an aggregation node) to access a particular content that is replicated over a number of nodes
  - Depends on the replica placement and the link cut set

Clustering-based placement:
Nodes are clustered and the content is placed at the cluster centroids

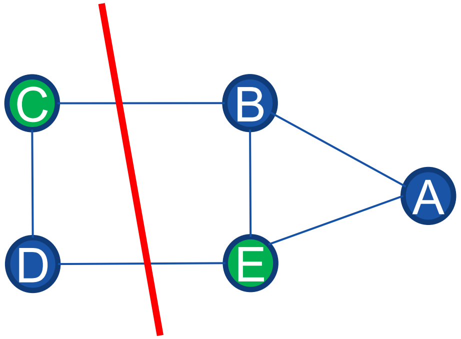

All nodes can access content

Average 2 Terminal Reliability (A2TR)

- A well-known connectivity measure under link cuts from the literature
- Defined as the probability that a randomly chosen pair of nodes is connected[1,2].

  A2TR = 1 → graph fully connected

  A2TR = 0 → graph completely disconnected

- Parameters:
  - Graph G(V,E)
  - Set of subgraphs C

$$A2TR = \frac{\sum_{i=1}^{|C|} |C_i| \times (|C_i| - 1)}{|V| \times (|V| - 1)}$$

[1] S. Rai and D. P. Agrawal, Distributed Computing Network Reliability. Los Alamitos, CA, USA: IEEE Computer Society Press, 1990.
[2] S. Neumayer and E. Modiano, "Network Reliability With Geographically Correlated Failures," 2010 Proceedings IEEE INFOCOM, San Diego, CA, 2010.

A2TR Example

- Fully connected network

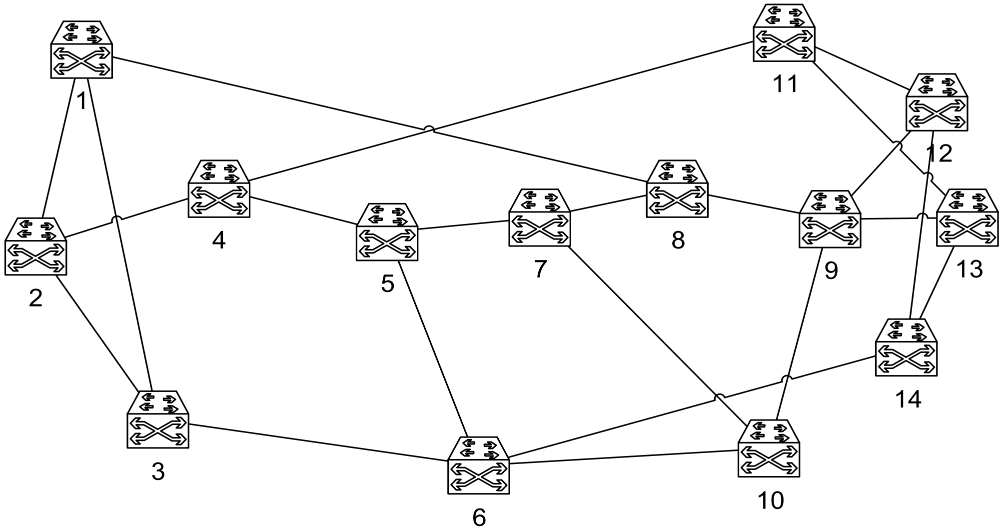

$$A2TR = \frac{\sum_{i=1}^{|C|} |C_i| \times (|C_i| - 1)}{|V| \times (|V| - 1)} = \frac{14 \times (14 - 1)}{14 \times (14 - 1)} = 1$$

A2TR Example

- Completely disconnected network

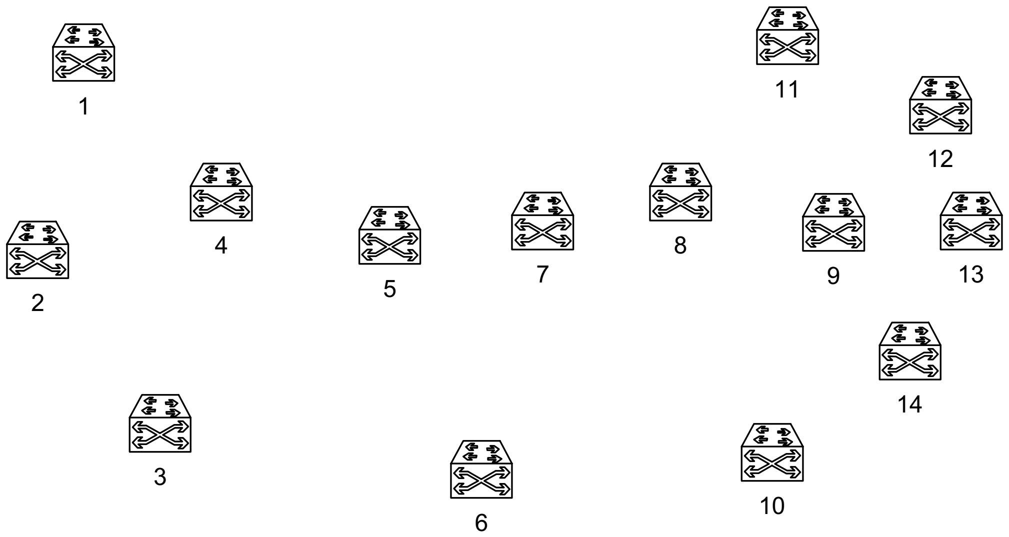

$$A2TR = \frac{\sum_{i=1}^{|C|} |C_i| \times (|C_i| - 1)}{|V| \times (|V| - 1)} = \frac{1 \times (1-1) + 1 \times (1-1) + \ ... + 1 \times (1-1)}{14 \times (14-1)} =$$

$$\frac{0}{14 \times (14-1)} = 0$$

A2TR Example

- A random cut

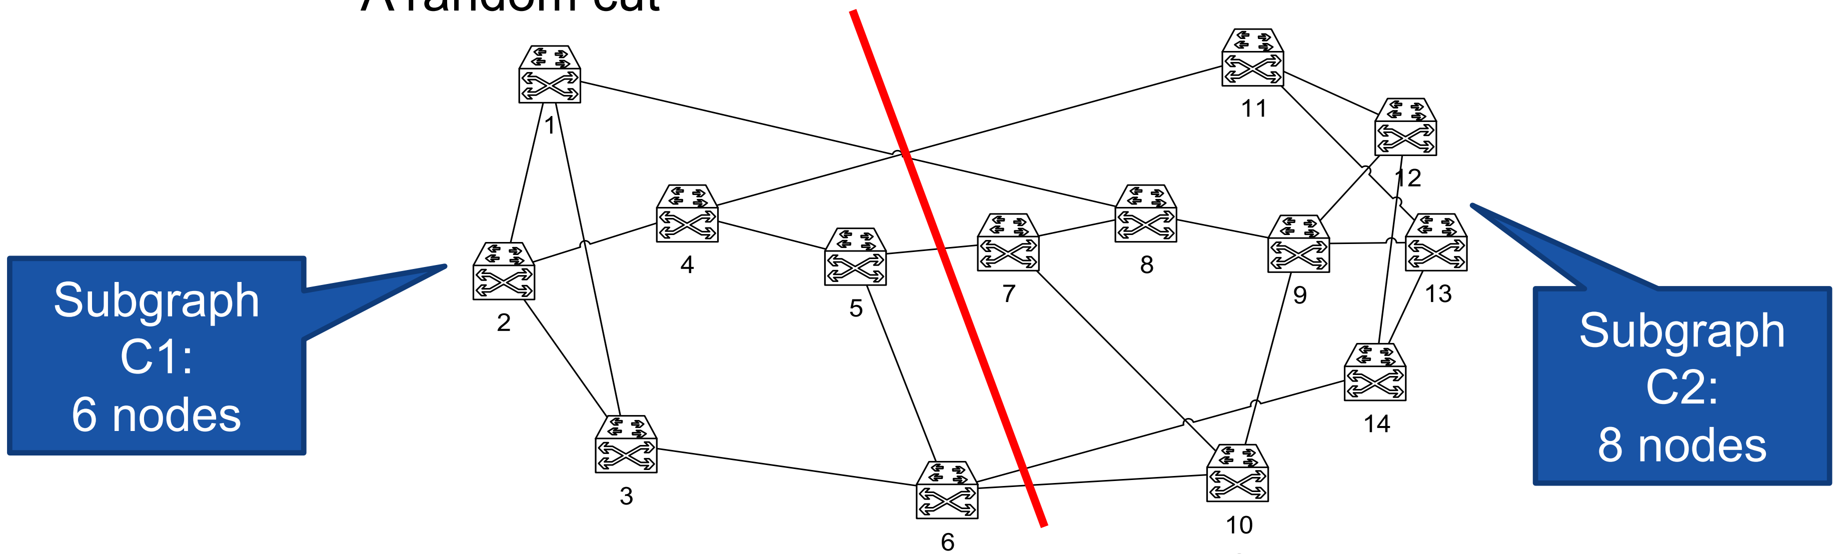

$$A2TR = \frac{\sum_{i=1}^{|C|}|C_i|\times(|C_i|-1)}{|V|\times(|V|-1)} = \frac{\overbrace{6\times(6-1)}^{C1}+\overbrace{8\times(8-1)}^{C2}}{14\times(14-1)} = \frac{86}{182} = 0.4725$$

- A randomly selected pair of nodes can be connected in 47% of cases

Content Accessibility in CDNs

- How to quantify the content accessibility on the example below?
- 2 replicas
- Best Case Scenario
- Worst Case Scenario
- Real Case Scenario

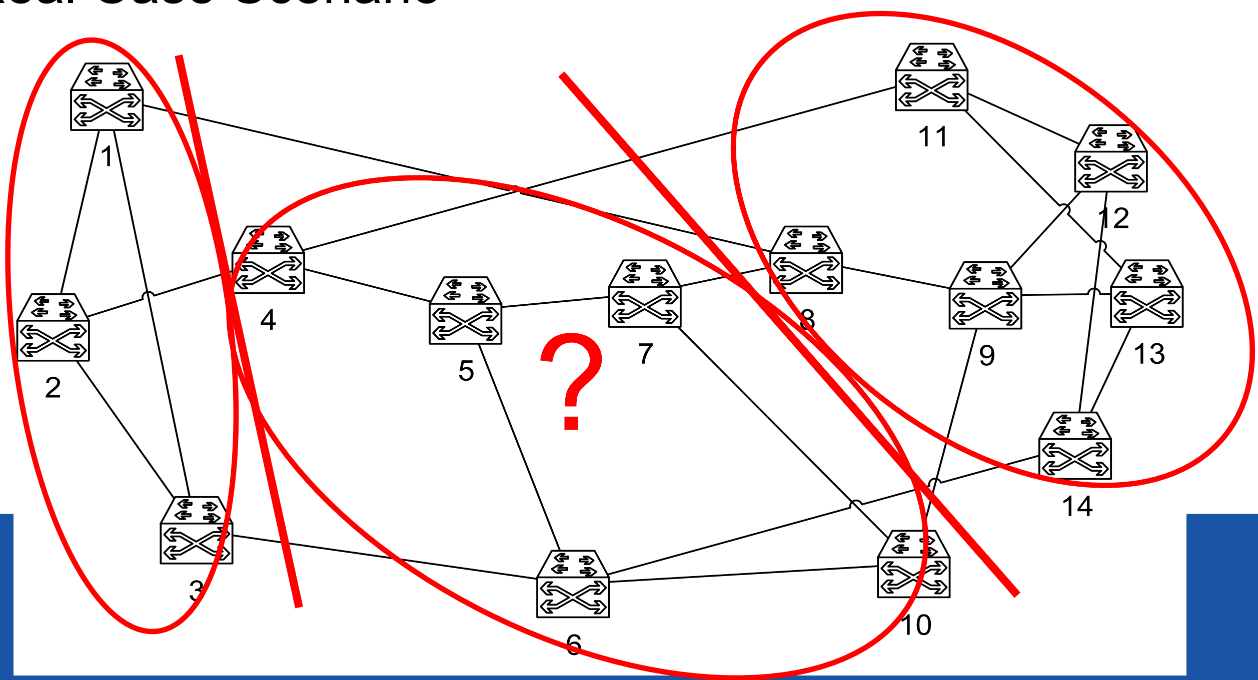

Average Content Accessibility (ACA)

- Measures the portion of nodes that are still able to connect to a replica for a given portion of cut links

Symbol	Description
$G(V,E)$	Network graph with nodes and links
r	Number of replicas
C	Set of connected components
C_i	A particular connected component with $\lvert C_i \rvert$ nodes
x_i	1 if there is a replica in connected component C_i, 0 otherwise

- **Best Case Scenario** – Replicas are spread across the largest connected components
- **Worst Case Scenario** – Replicas are confined in the smallest connected components
- **Real Case Scenario** – Replica placement is given

ACA in the Best Case Scenario (ACA-BCS)

- Content replicas are spread across the largest connected components
- Gives an upper bound on the ACA for a given number of replicas

$$\boxed{ACA_{bcs}(r) = \frac{\sum_{i=1}^{r} |C_i^{desc}|}{|V|}}$$

$$ACA_{bcs}(1) = \frac{\sum_{i=1}^{r} |C_i^{desc}|}{|V|} = \frac{6}{14} = 0.42$$

$$ACA_{bcs}(2) = \frac{\sum_{i=1}^{r} |C_i^{desc}|}{|V|} = \frac{6+5}{14} = 0.78$$

$$ACA_{bcs}(3) = \frac{\sum_{i=1}^{r} |C_i^{desc}|}{|V|} = \frac{6+5+3}{14} = 1$$

$|V| = 14 ; |E| = 22$

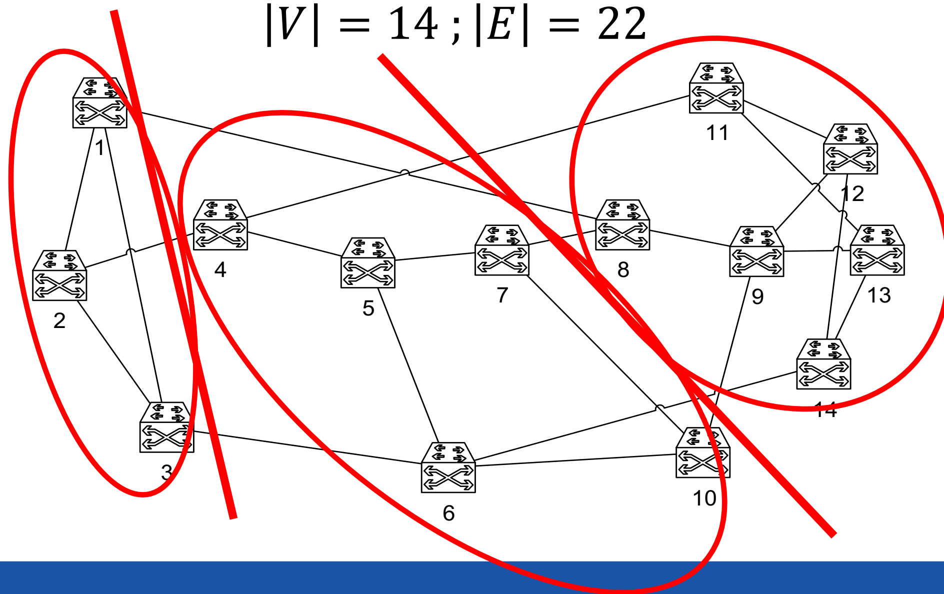

ACA in the Worst Case Scenario (ACA-WCS)

- Gives a lower bound on ACA
- Replicas are confined in the smallest connected components
- **Exact fit:**
  - The replicas are confined in connected components whose size is equal to the number of replicas
- **Best fit:**
  - The replicas are located in connected components whose size is the closest to the number of replicas

Algorithm 1: Algorithm for the ACA_{wcs}

Data: $G(V, E)$, r, C

Result: $ACA_{wcs}(r)$

Exact fit

for *combination in binary* $0..2^{|C|} - 1$ **do**

    $sum \leftarrow \sum_{i=1}^{|C|} |C_i| \times combination_i$;

    if $sum = r$ **then**

        return $\frac{sum}{|V|}$;

$\bar{r} \leftarrow r$; $CP \leftarrow C$; $sum \leftarrow 0$;

Best fit

while $\bar{r} > 0$ **do**

    if $\exists_i$ *such that* $|CP_i| > \bar{r}$ **then**

        $C_{BF} \leftarrow min_i(|CP_i| - \bar{r})$;

        $\bar{r} \leftarrow \bar{r} - |C_{BF}|$;

        $sum \leftarrow sum + |C_{BF}|$;

        $CP \leftarrow CP \backslash C_{BF}$;

    else

        $C_{BF} \leftarrow min_i(\bar{r} - |CP_i|)$;

        $\bar{r} \leftarrow \bar{r} - |C_{BF}|$;

        $sum \leftarrow sum + |C_{BF}|$;

        $CP \leftarrow CP \backslash C_{BF}$;

return $\frac{sum}{|V|}$;

ACA in the Worst Case Scenario (ACA-WCS)

$$ACA_{wcs}(1) = \frac{3}{14} = 0.21$$

$$ACA_{wcs}(2) = \frac{3}{14} = 0.21$$

$$ACA_{wcs}(3) = \frac{3}{14} = 0.21$$

$$ACA_{wcs}(4) = \frac{5}{14} = 0.36$$

$|V| = 14\,;|E| = 22$

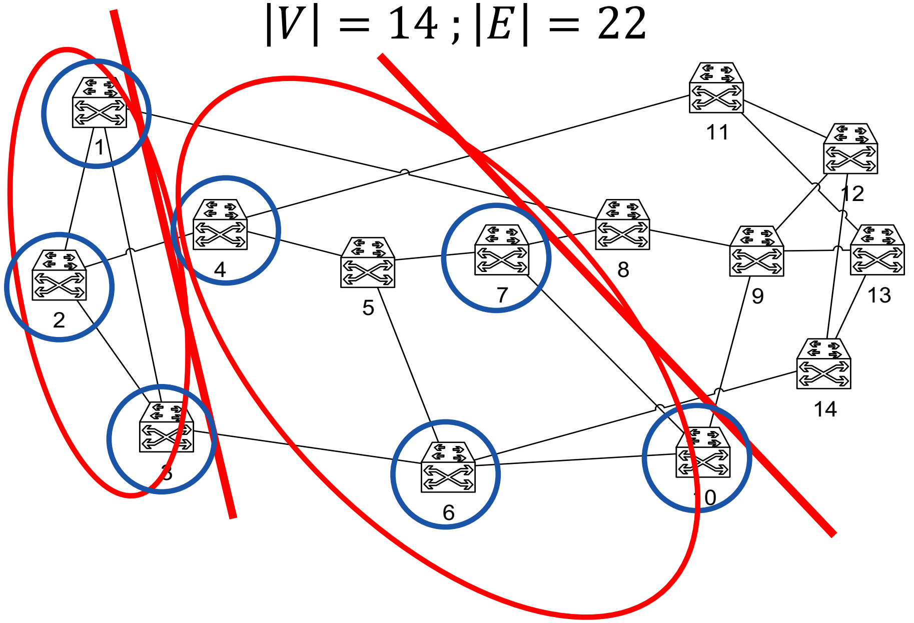

ACA in the Real Case Scenario (ACA-RCS)

- Content replica placement is given beforehand

$$\boxed{ACA_{rcs}(r) = \frac{\sum_{i=1}^{|C|} |C_i| \times x_i}{|V|}}$$

$|V| = 14\,; |E| = 22$

$$ACA_{rcs}(1) = \frac{\sum_{i=1}^{|C|} |C_i| \times x_i}{|V|} = \frac{5}{14} = 0.35$$

$$ACA_{rcs}(2) = \frac{\sum_{i=1}^{|C|} |C_i| \times x_i}{|V|} = \frac{5+3}{14} = 0.57$$

$$ACA_{rcs}(2) = \frac{\sum_{i=1}^{|C|} |C_i| \times x_i}{|V|} = \frac{5+3}{14} = 0.57$$

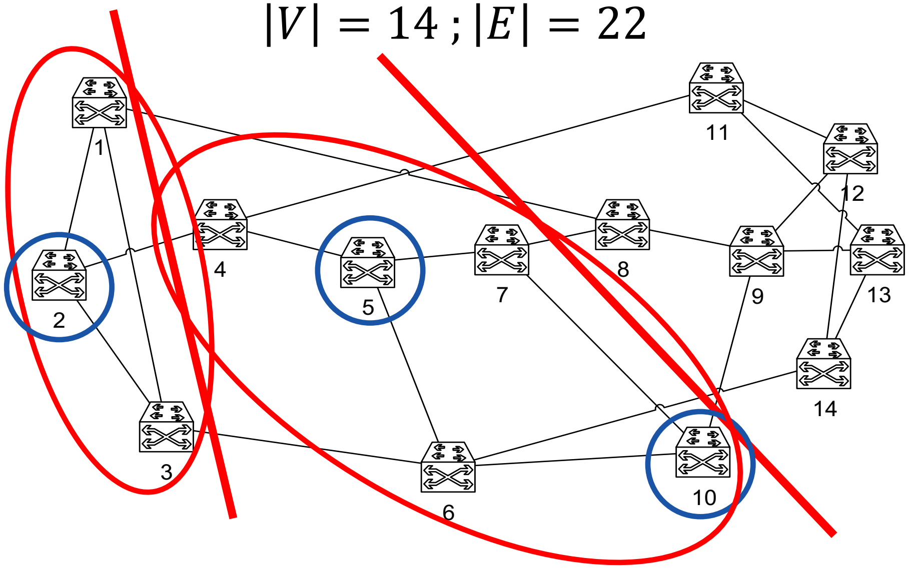

Simulation setup

- Scenarios:
  - 3 network topologies[1]
  - Replica placement strategies:
    - Degree centrality
    - Betweeness centrality
    - Closeness centrality
    - Clustering with K-Means
  - Simultaneous and sequential link cut attacks
    - Based on link betweeness

Topology	n	m	$\mathbf{k} \pm \sigma$	D
Sprint	11	18	3.27$\pm$1.42	4
Géant	40	61	3.05$\pm$1.92	8
Garr	61	75	2.45$\pm$2.58	8

[1] S. Knight, H. X. Nguyen, N. Falkner, R. Bowden, and M. Roughan, “The internet topology zoo,” IEEE J. Sel. Areas Commun., vol. 29, no. 9, pp. 1765–1775, Oct 2011.

Discrepancies between A2TR and ACA

- Sprint network (11 nodes, 18 links)

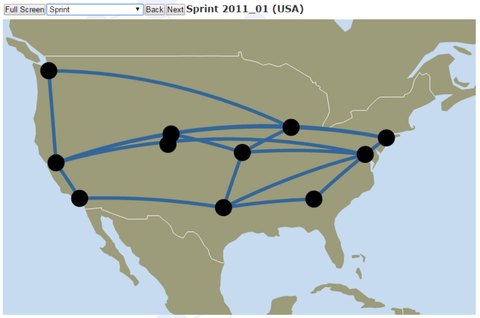

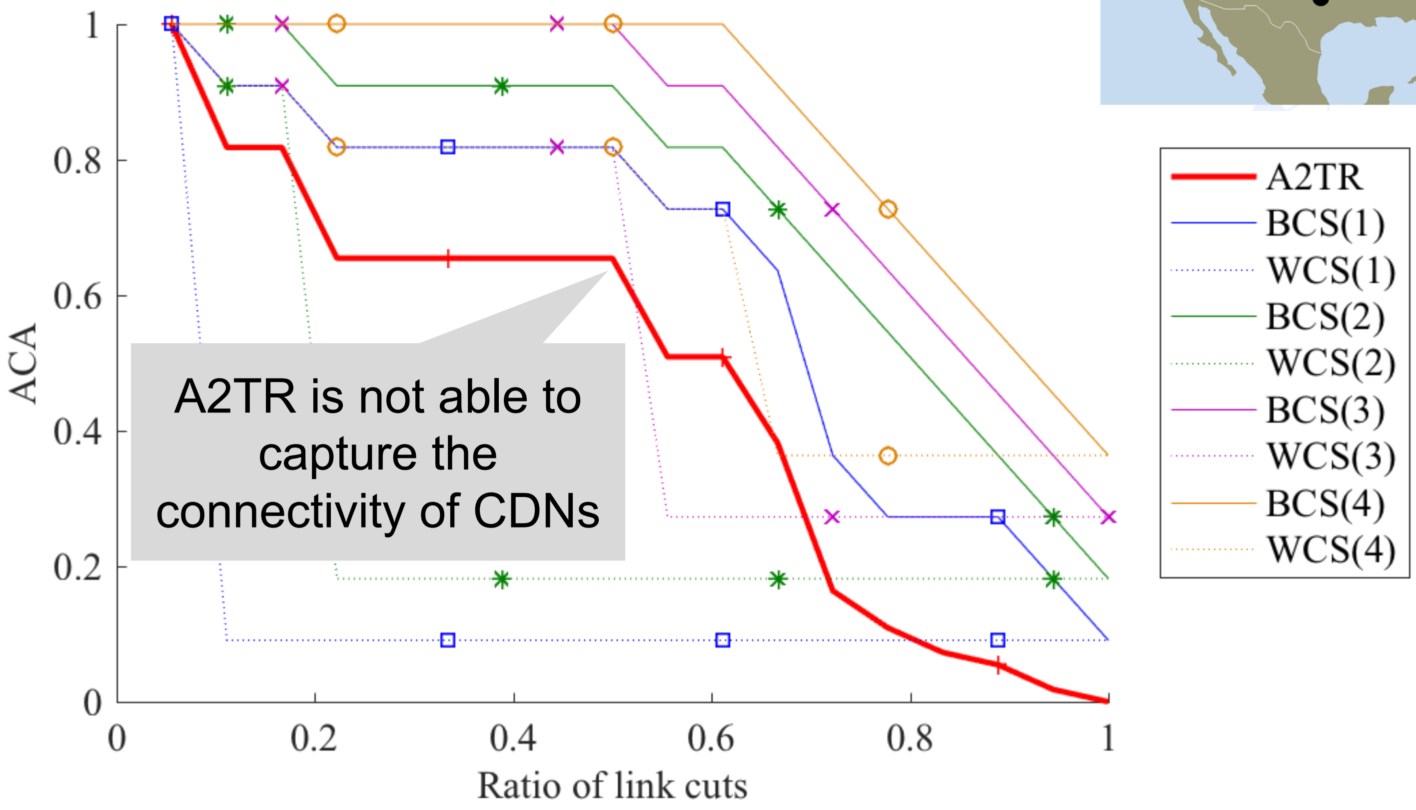

ACA-BCS vs. ACA-WCS

- Sprint network (11 nodes, 18 links)

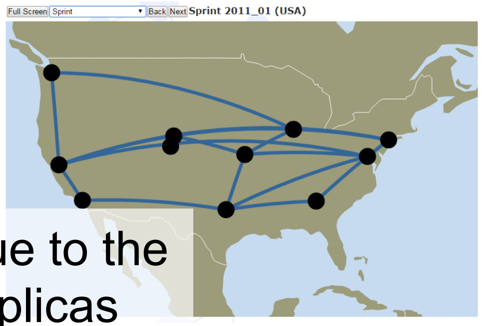

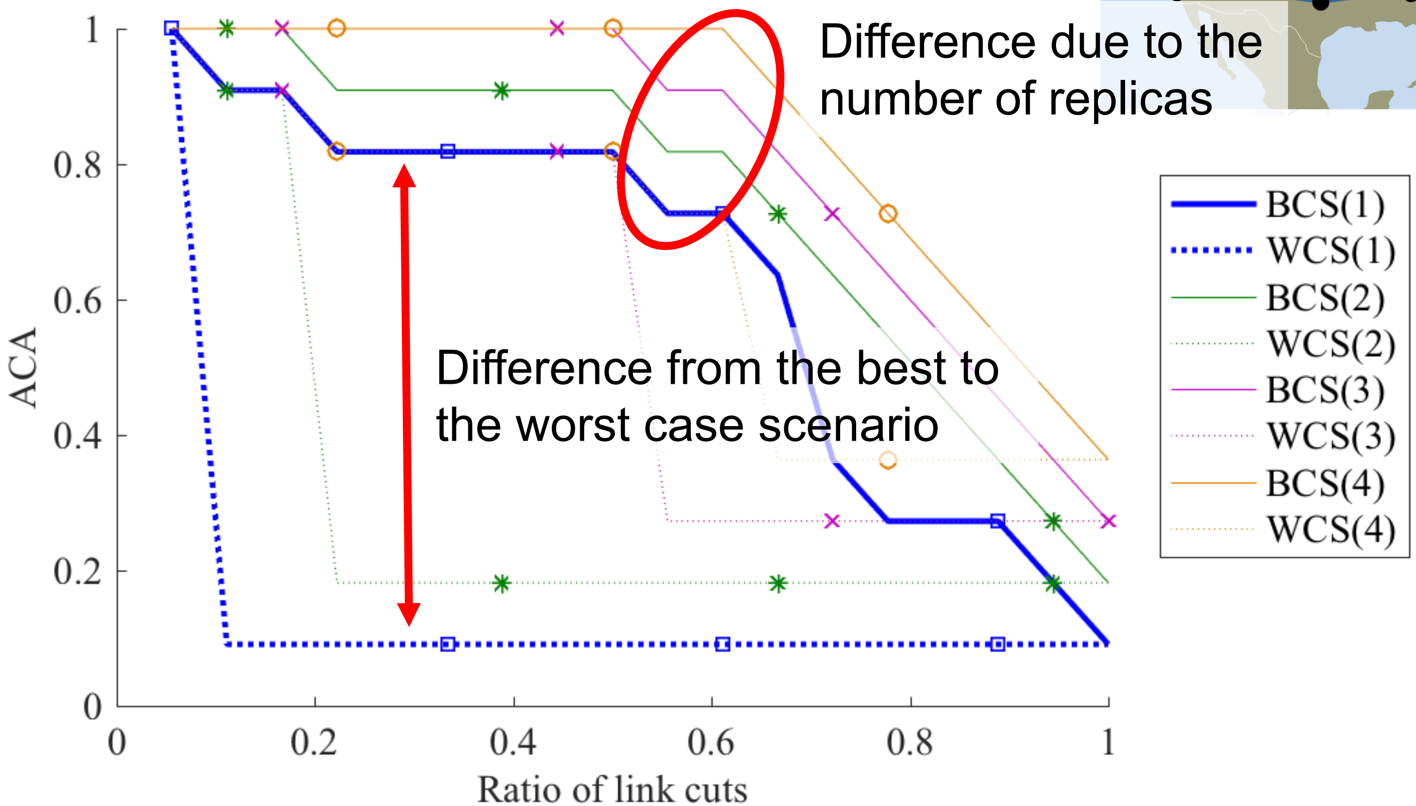

Impact of the number of replicas

- How does the increase in the number of replicas change ACA?

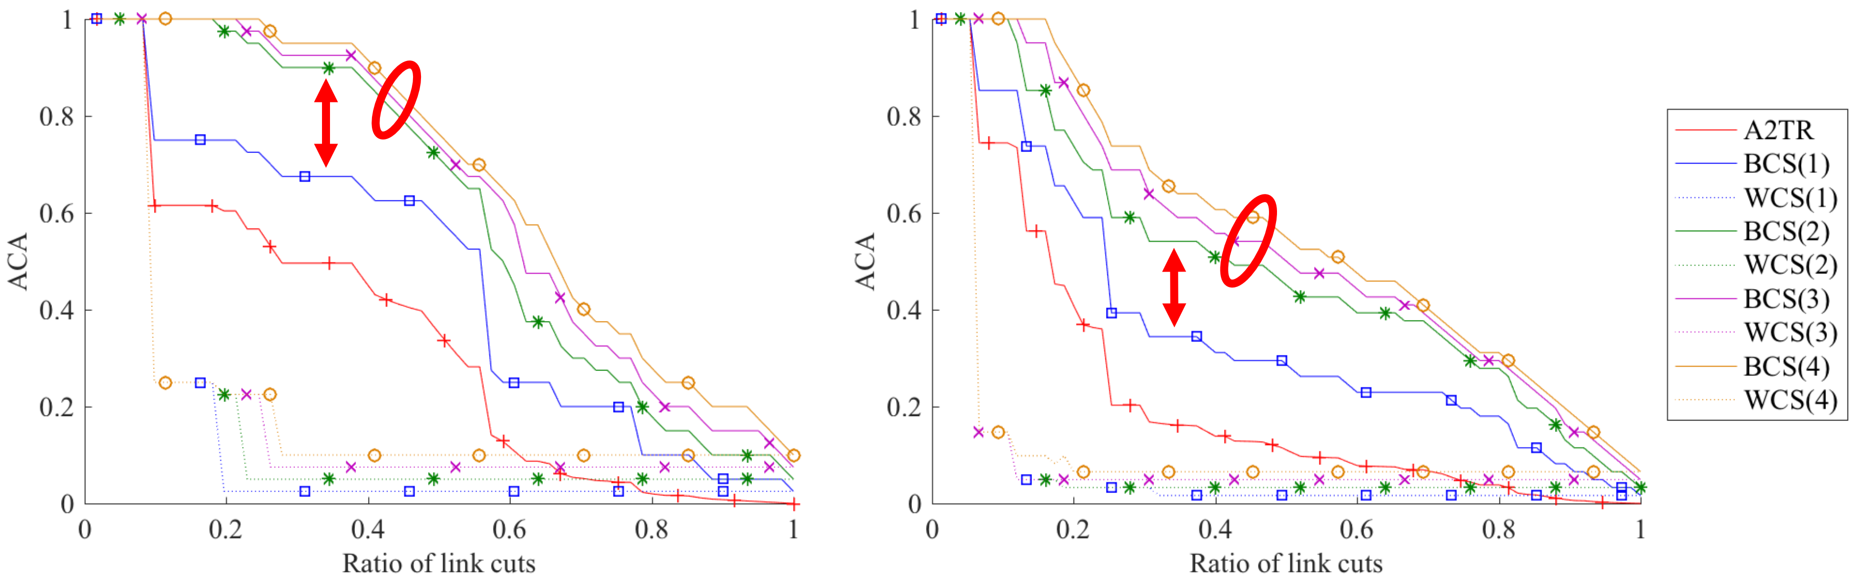

Géant

Garr

Impact of the replica placement on ACA-RCS

- Géant network (40 nodes, 61 links, 2 replicas)

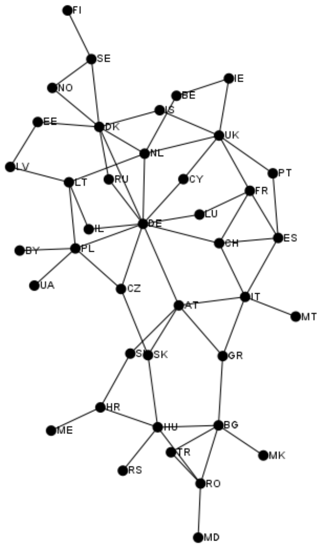

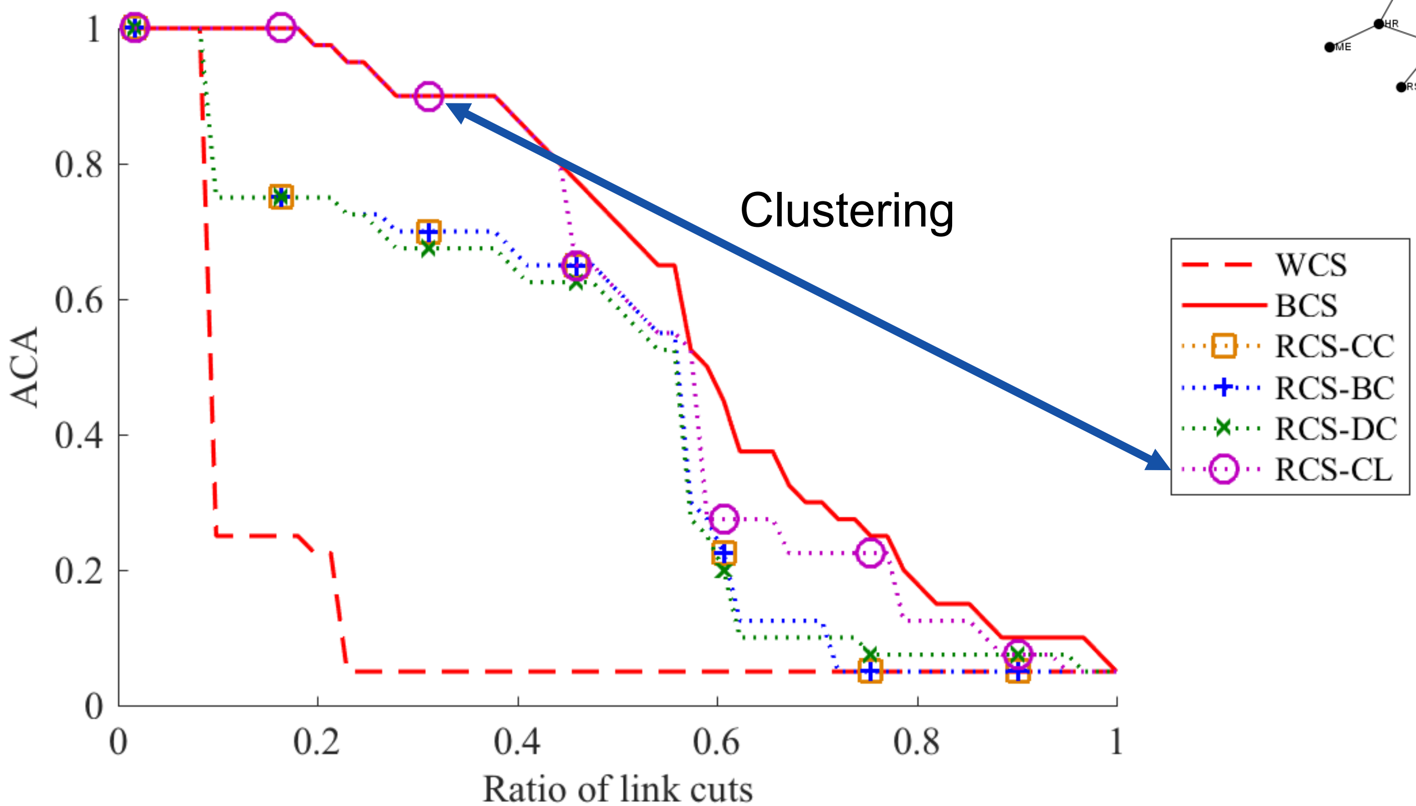

Impact of the replica placement on ACA-RCS

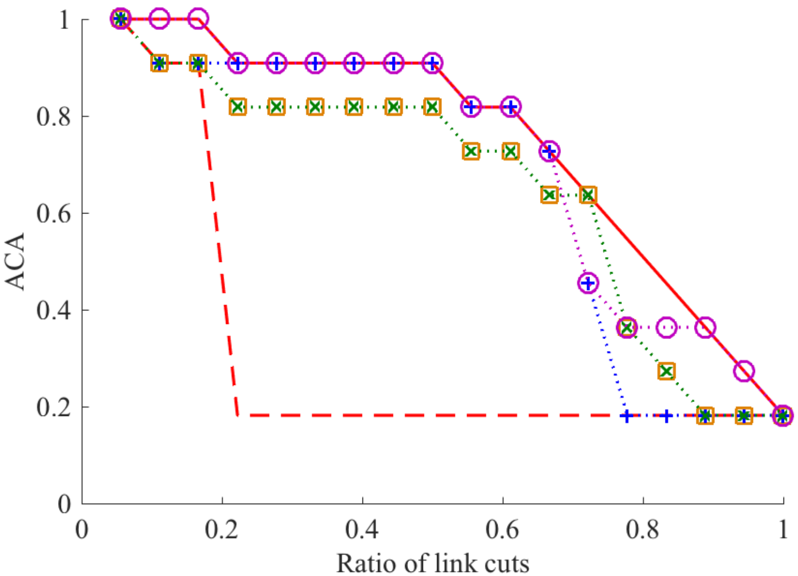

Sprint

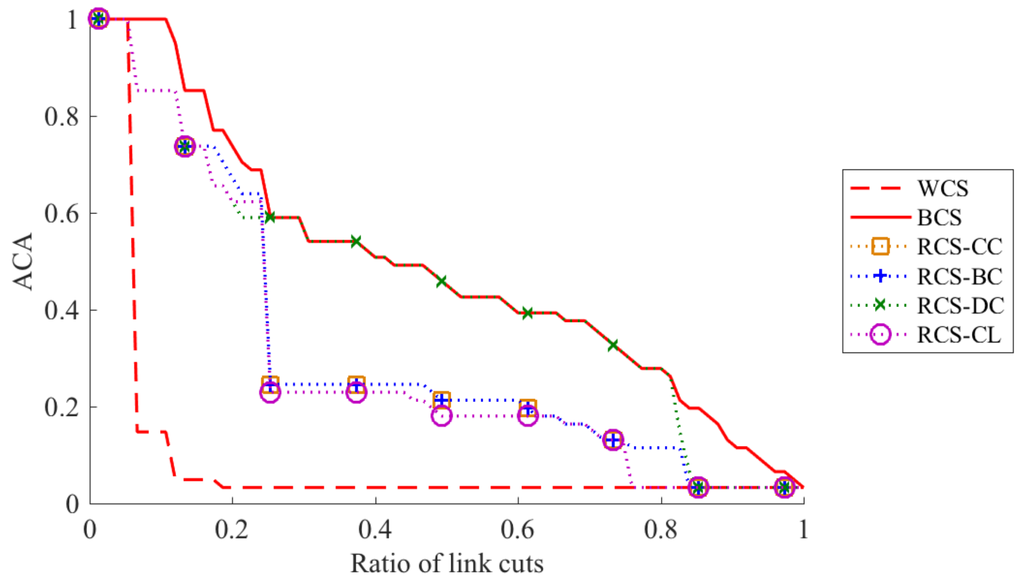

Garr

Impact of the type of attack on ACA

- Sprint network (11 nodes, 18 links, 2 replicas)
- Simultaneous cuts: link criticality evaluated once
- Sequential cuts: link criticality re-evaluated in the modified topology upon each cut

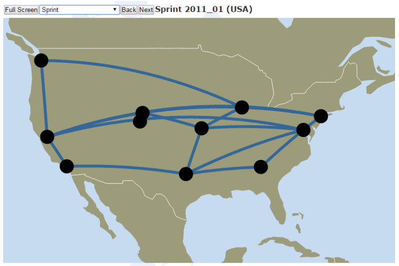

BCS difference between simultaneous and sequential cuts

WCS difference between simultaneous and sequential cuts

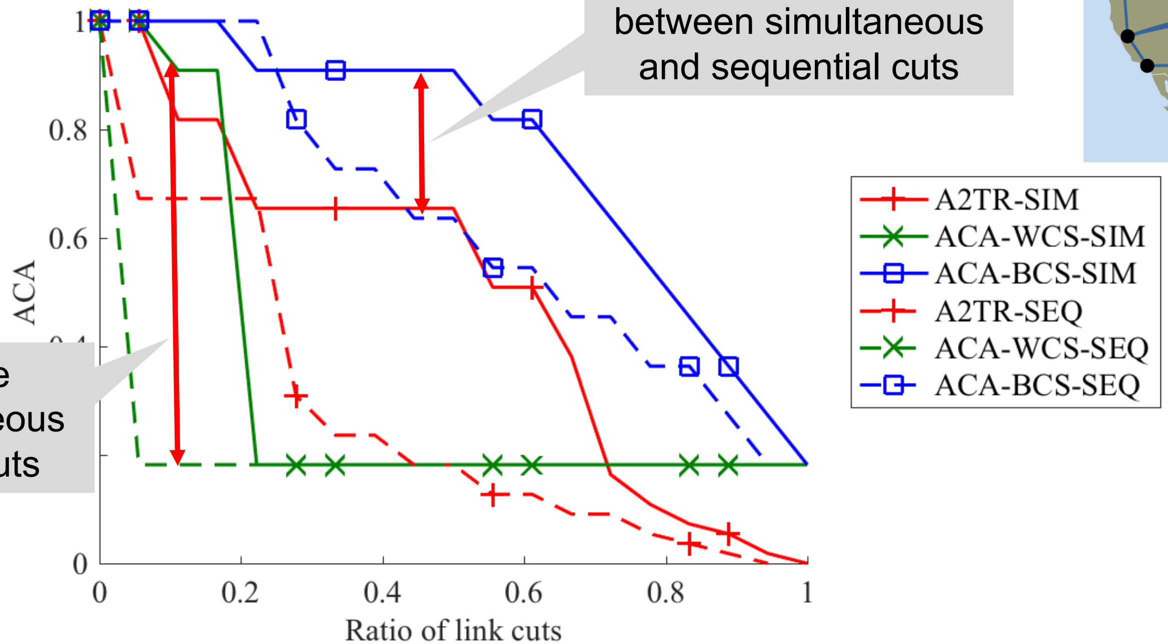

Impact of the type of attack on ACA

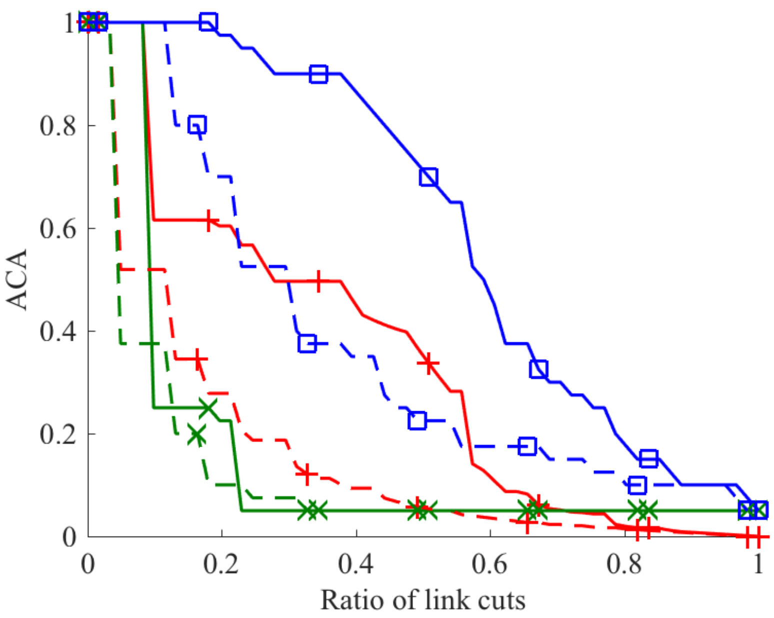

Géant

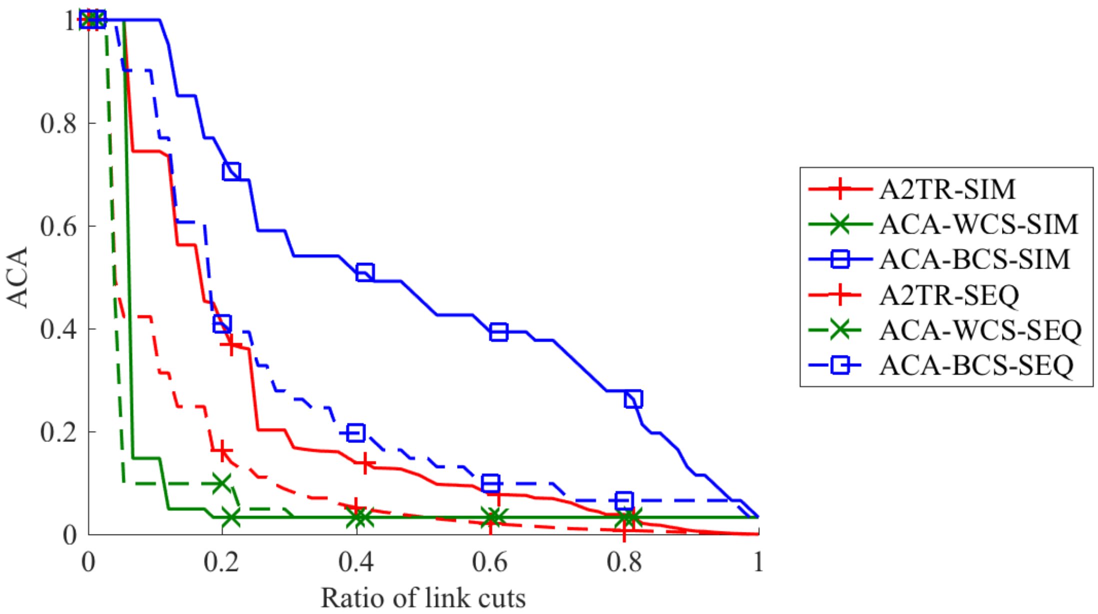

Garr

Conclusions and next steps

- State-of-the-art (A2TR) strategies are not applicable to gauge CDN robustness to link cuts
- The proposed Average Content Accessibility (ACA) measure can capture CDN robustness in the worst, the best and realistic case
- Adding replicas does not always significantly increase content accessibility
- Content placement strategies greatly impact content accessibility
- Simultaneous and sequential attacks (link cuts) affect the content accessibility in different ways
- Next steps:
  - Consider the impact of link cuts to other parameters, e.g., latency and network resource usage
  - Analyze/propose content placement strategies considering content accessibility
  - Find the right number of replicas to support a required robustness level
  - Develop network topology update/enhancement approaches to improve content accessibility in CDNs

Questions?

Thank you for your attention!
☺

Content Accessibility in Optical Cloud Networks Under Targeted Link Cuts

Carlos Natalino, Aysegul Yayimli, Lena Wosinska, Marija Furdek
marifur@kth.se