The Minimum Cost D-Geodiverse Anycast Routing with Optimal Selection of Anycast Nodes

Amaro de Sousa[1,2], Dorabella Santos[3]

[1] Instituto de Telecomunicações
Portugal

[2] Universidade de Aveiro
Portugal

[3] INESC Coimbra
Portugal

Outline

- Motivation
- Definition of D-geodiverse anycast routing
- The minimum cost D-geodiverse anycast problem
- Computational results
- Conclusions

Motivation

- Disaster based failures can seriously disrupt a telecommunications network, making its services unavailable
- It is important not only to quickly recover the network in the disaster area (post-disaster problem) but also to minimize the disaster impact between network nodes outside the disaster area (pre-disaster problem)
- Improving the preparedness of telecommunication networks to disasters is becoming a key issue
- To enhance the preparedness of networks to disasters, one approach is path geo-diversification:
  - *to take into consideration the geographical diversity of the network topology when making routing decisions*

Motivation

- In anycast, network nodes are partitioned into two sets:
  - the source nodes
  - the anycast (destination) nodes
- The traffic of each source node is routed towards the anycast node providing the minimum cost routing path.
- Content Delivery Networking (CDN) - content replicated over multiple data centers (DCs) and users retrieve content from the closest DC.
- Software Defined Networking (SDN) - control plane separated from the data plane and based on a set of physically distributed SDN controllers; switches query the closest (primary) controller for routing decisions.

Definition of D-geodiverse anycast routing

- Consider a given geographical network G=(N,A) and a distance parameter D.
- A D-geodiverse anycast routing solution guarantees for each source node that:
  - there are two routing paths,
  - each one towards a different anycast node, such that
  - the geographical distance between the two paths is at least D.
- A disaster with a coverage diameter below D (involving neither the source node nor its entire set of links):
  - cannot affect both paths simultaneously,
  - thus, enhancing the network robustness to natural disasters.

Definition of D-geodiverse anycast routing

Source node: 1
Anycast nodes: 3, 5 and 8

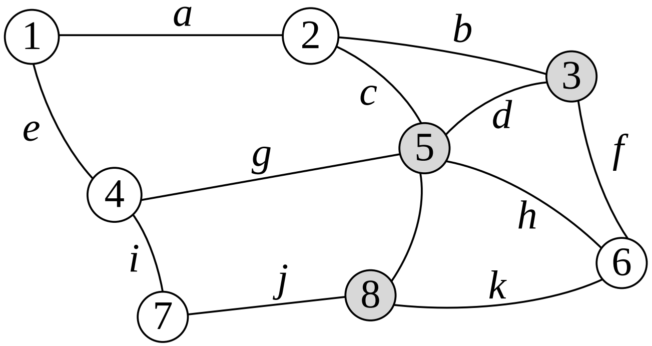

Assume: $D_1 < D_2 < D_3$

If D < D_1

- all solutions are feasible

If D > D_3

- none of these solutions is feasible
- either another set of anycast nodes can provide a feasible solution
- or node 1 must be a anycast node

Solution 1:

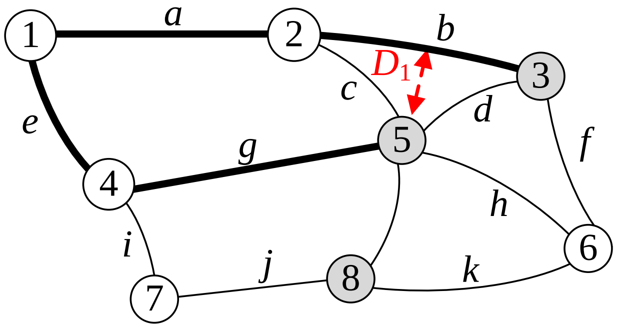

Solution 2:

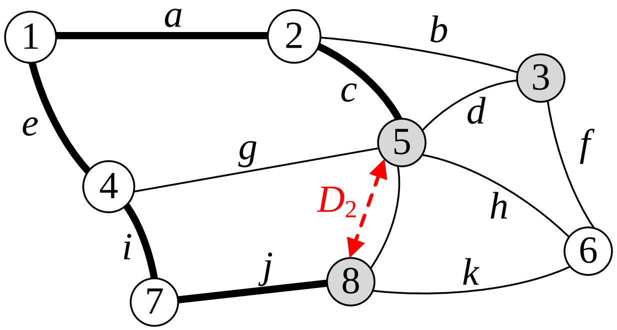

Solution 3:

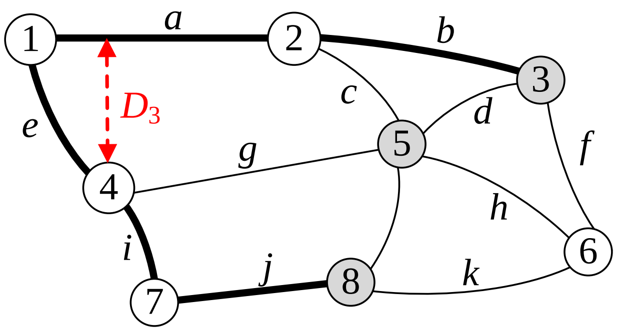

Minimum cost D-geodiverse anycast problem

- The selection of the anycast nodes has an impact both on the feasibility and cost of a D-geodiverse anycast routing solution.
- Given:
  - a geographical network G=(N,A) with a routing cost c_{ij} associated to each arc $(i,j) \in A$
  - a geographical distance parameter D
  - a number of anycast nodes R
- The minimum cost D-geodiverse anycast problem (MCD-GAP) aims to select a set of R anycast nodes that obtain a minimum cost routing solution.

Minimum cost D-geodiverse anycast problem

- MCD-GAP is modelled and solved by Integer Linear Programming (ILP).
- Pre-processing:
  - Based on the geographical information of the network, we compute $\delta(a, b)$ as the shortest distance between any point of link a and any point of link b.
  - Then, for each node $s \in N$, we compute the set of link pairs P_s whose minimum distance between them is lower than D.
  - Special care is taken for link pairs sharing the source node $s \in N$ as two of such links must exist whatever D is defined (details in paper).

Minimum cost D-geodiverse anycast problem

- Then, MCD-GAP is formulation as an ILP model guaranteeing that for each node $s \in N$:
  - if s is not an anycast node, at most one of each pair of links P_s is in the two routing paths from s to two different anycast nodes
  - if s is an anycast node, no routing paths need to exist in the solution starting from s
- The objective function is the sum of the costs of all pairs of paths from each source node
- We mitigate the symmetry problem with a variable elimination rule (details in paper)

Minimum cost D-geodiverse anycast problem with vulnerability regions

- MCD-GAP assumes that a disaster can happen at any region of the network
- In practice, the probability of natural hazards is not uniform in the geographical area of a given network
- Network operators might want to tailor the network robustness to the different hazard types and regions, which are referred as vulnerability regions
- Moreover, the network operator might consider different geographical parameters D for each vulnerability region, depending on its hazard type

Minimum cost D-geodiverse anycast problem <u>with vulnerability regions</u>

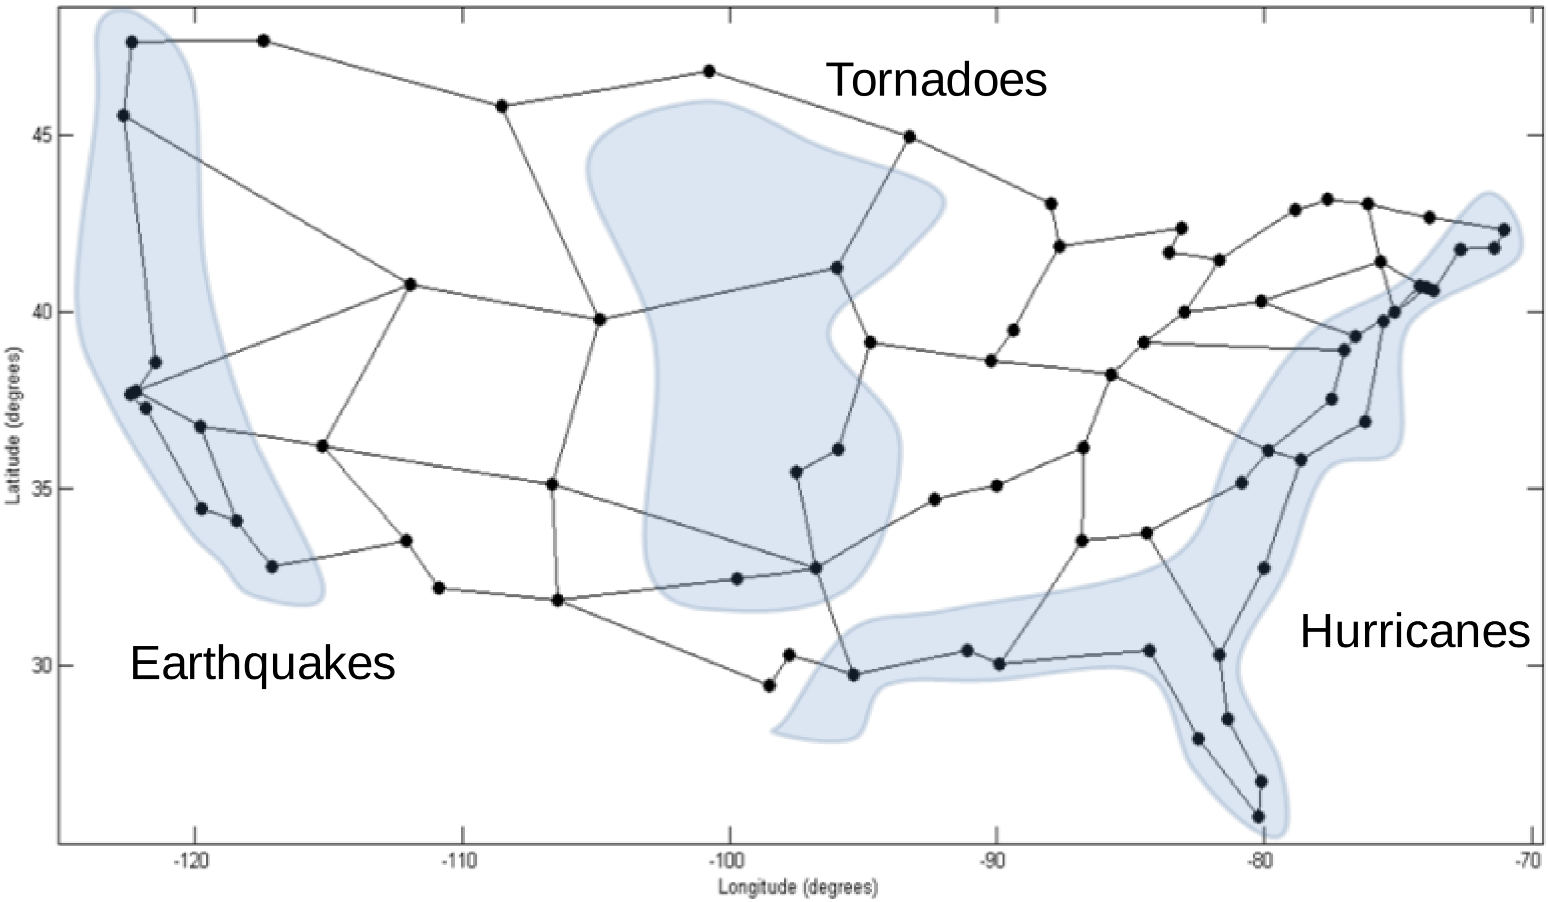

Minimum cost D-geodiverse anycast problem with vulnerability regions

- Consider:
  - a set of V vulnerability regions and
  - a distance D_v associated to each region $v = 1 \ldots V$
- The aim is that each pair of paths is:
  - node disjoint outside regions and
  - D_v-geodiverse inside vulnerability region $v = 1 \ldots V$
- A pair of links is in set P_s only if both links belong to a region v and if their minimum distance is lower than D_v
- If a pair of links belongs to different regions, the largest value D_v among all involved regions is used

Computational results

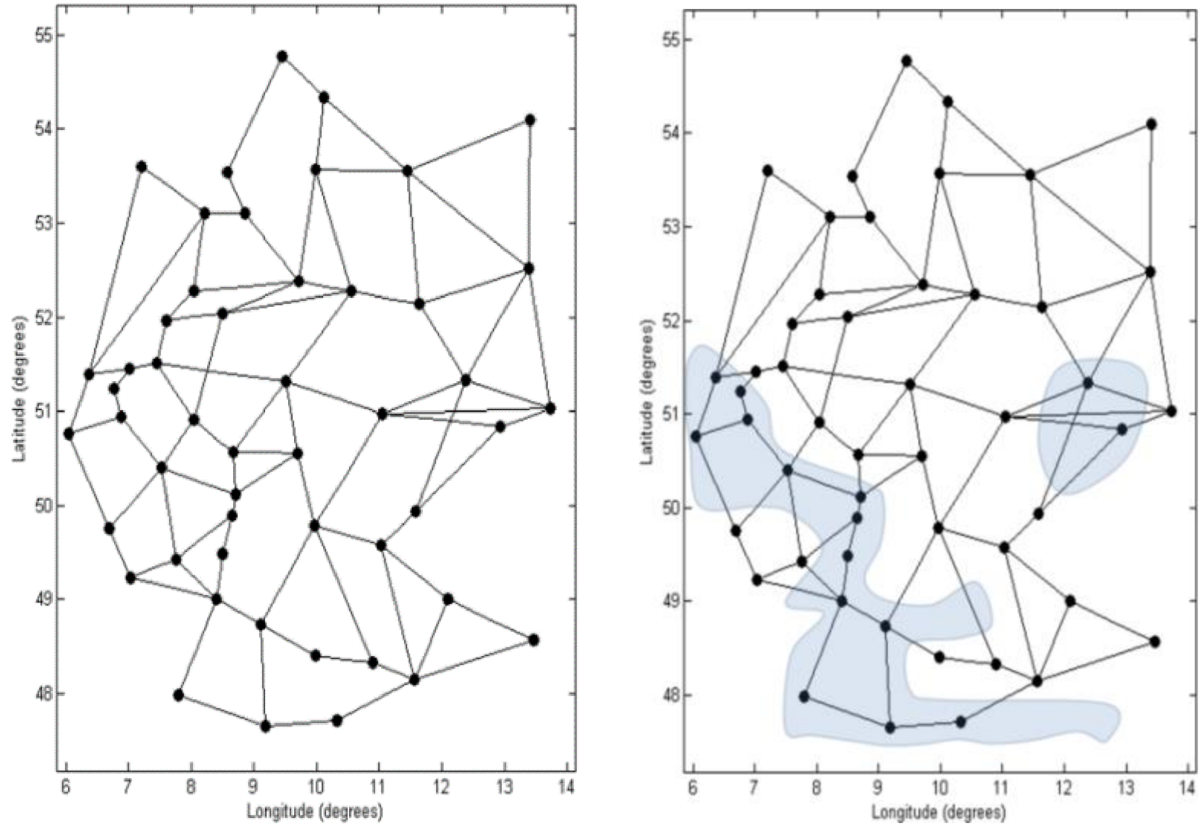

- Germany50: 50 nodes, 88 links, average node degree 3.52 (http://sndlib.zib.de)
- Germany seismic hazard map from:
  M. Müller et al., CEDIM Risk Explorer – a map server solution in the project Risk Map Germany, Natural Hazards and Earth System Sciences, vol. 6, pp. 711-720, 2006

Computational results

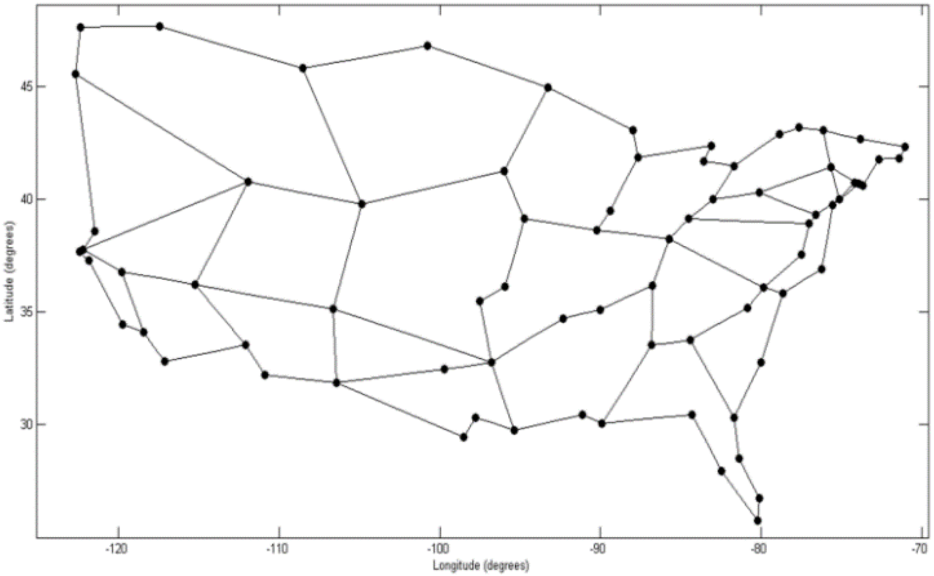

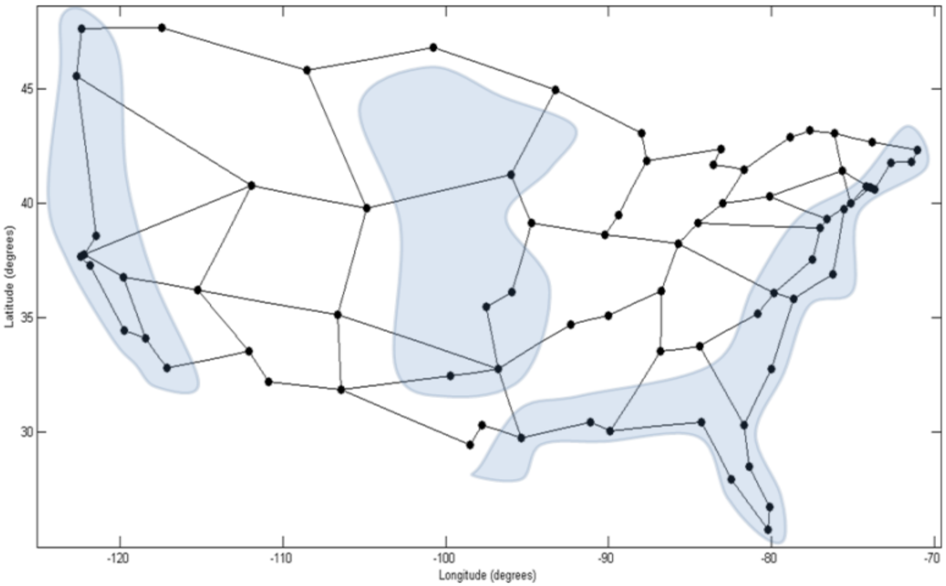

- CORONET CONUS: 75 nodes, 99 links, average node degree 2.64 (http://monarchna.com/topology.html)
- USA natural hazard risk map from: http://alertsystemsgroup.com/earthquake-early-warning/informative-maps

Computational results

Germany 50 results *without vulnerability regions*

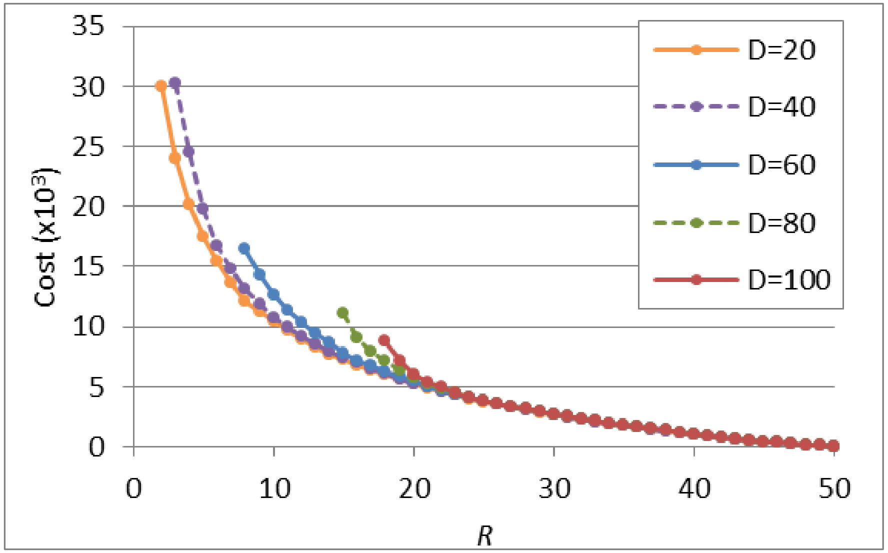

- ❖ The cost decreases with higher number of anycast nodes R
- ❖ For the same R, cost increases with larger geographical distances D
- ❖ The minimum required number of anycast nodes R increases with larger values of D

Computational results

Germany 50 results <u>with and without vulnerability regions</u>

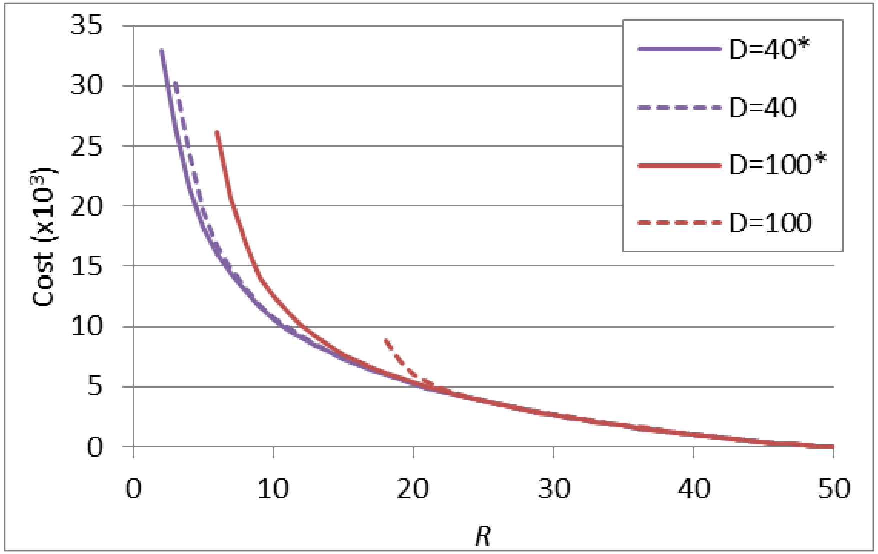

- For the pairs of values D and R such that both cases are feasible, there are cost gains in considering vulnerability regions
  - These cost gains are higher for higher values of D
- The minimum R is lower when vulnerability regions are considered
  - This reduction is higher for higher values of D

Computational results

CORONET CONUS results <u>without vulnerability regions</u>

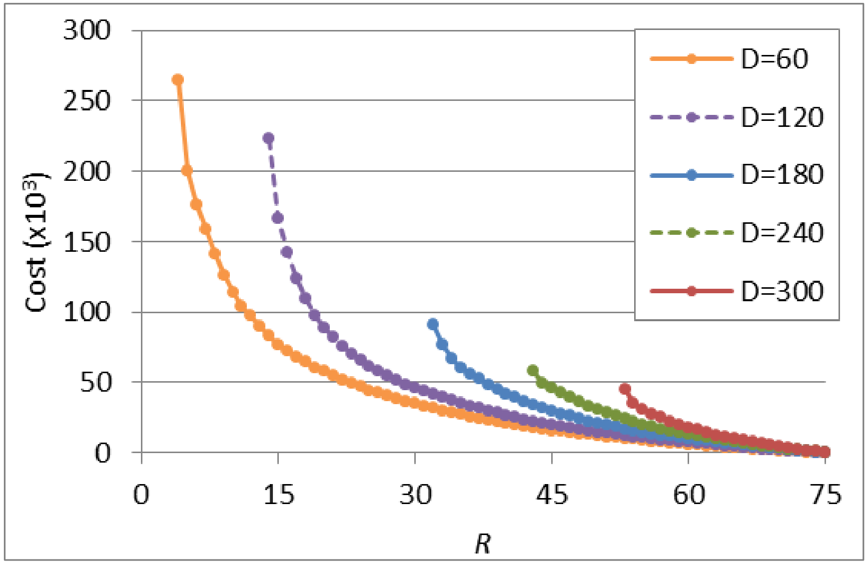

- The cost decreases with higher number of anycast nodes R
- For the same R, cost increases with larger geographical distances D
- The minimum required number of anycast nodes R increases with larger values of D

Computational results

CORONET CONUS results with and without vulnerability regions

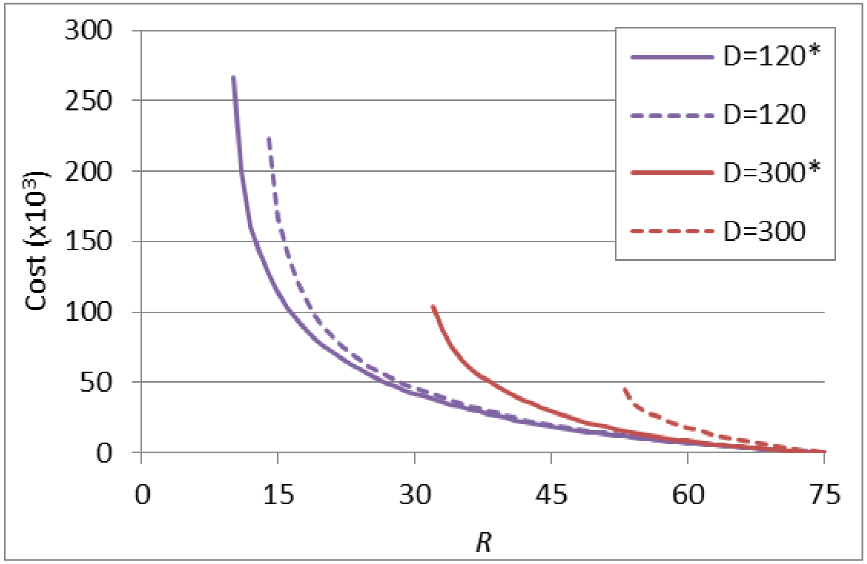

- For the pairs of values D and R such that both cases are feasible, there are cost gains in considering vulnerability regions
  - These cost gains are higher for higher values of D
- The minimum R is lower when vulnerability regions are considered
  - This reduction is higher for higher values of D

Computational results

Anycast nodes of CORONET CONUS for D = 120 Km and R = 14 anycast nodes

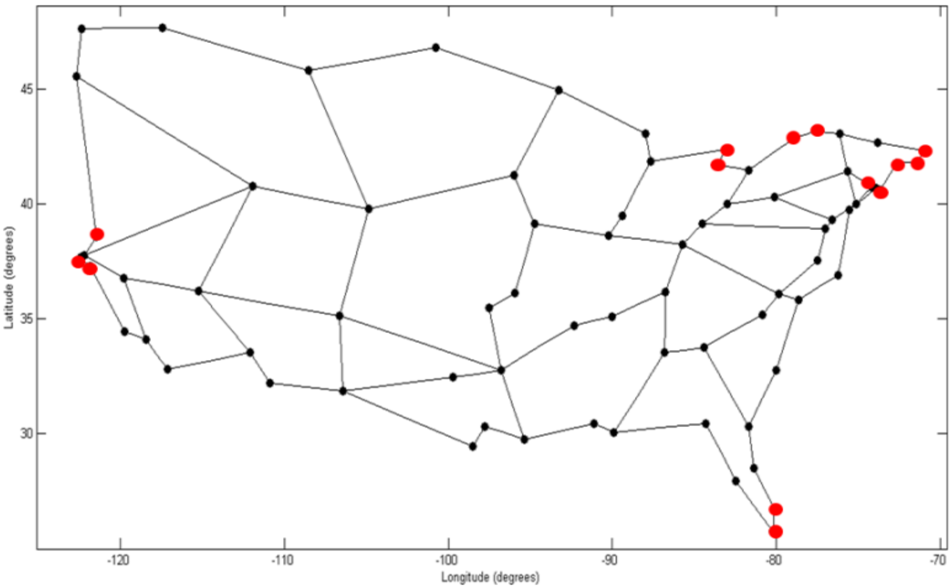

(a) Cost = 223008

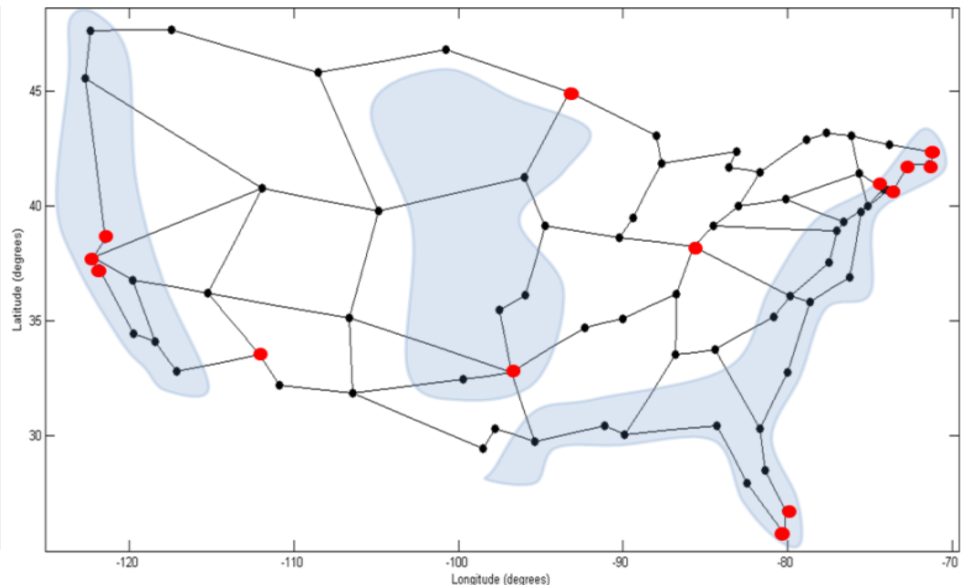

(b) Cost = 126362

- Anycast nodes are mainly selected on the network parts with closer nodes and shorter links in the whole network (a) or only inside the vulnerability regions (b).
- With vulnerability regions, anycast nodes are selected more uniformly throughout the network obtaining in this way an huge cost reduction of 43.3%.

Computational results

Germany 50 running times

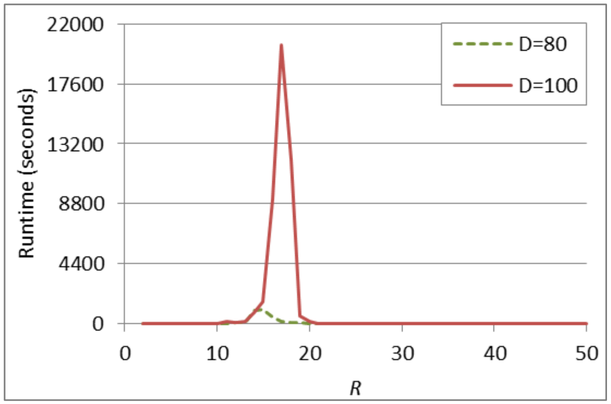

(a)

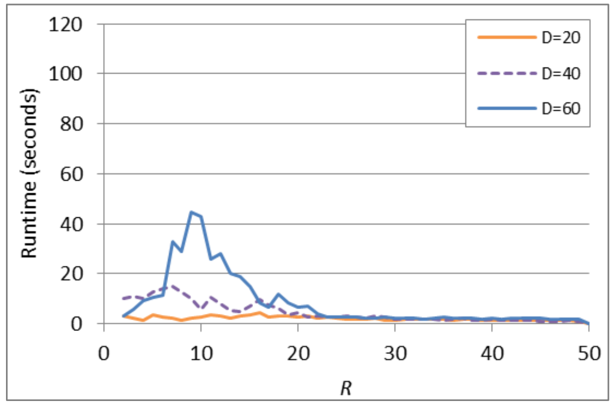

(b)

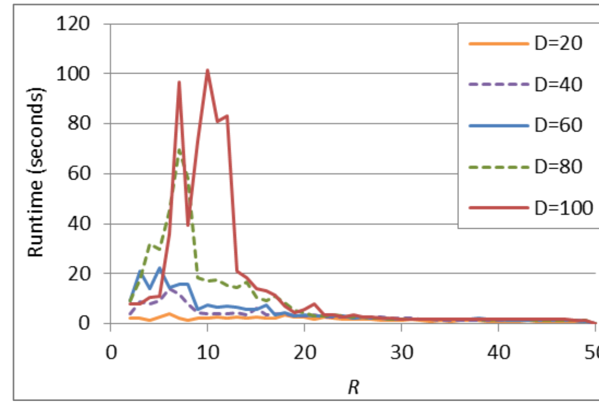

(c)

- CPLEX runtime without (a-b) and with (c) survivability regions
- Instances with larger values of D are harder to be solved
- Without vulnerability regions, the instances are harder to be solved than with vulnerability regions
- The worst runtime becomes 1030 seconds for D = 80 Km and almost 20500 seconds (around 5 hours and 40 minutes) for D = 100 Km

Computational results

CORONET CONUS running times

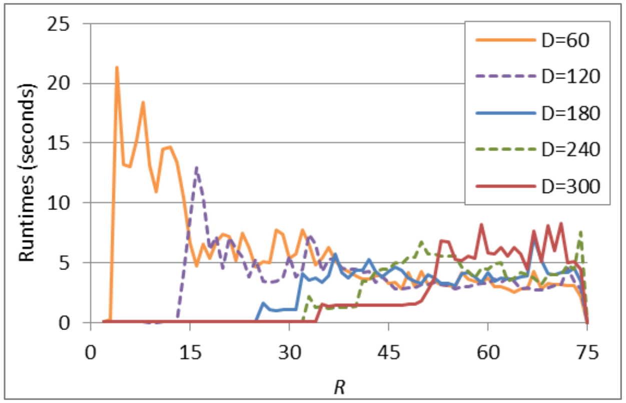

(a)

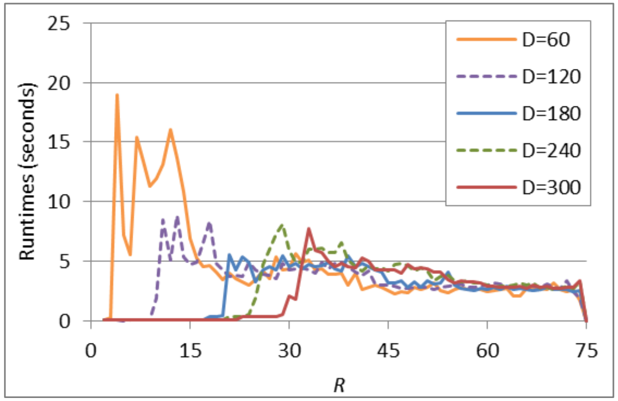

(b)

- CPLEX runtime without (a) and with (b) survivability regions
- These instances are much easier to solve than the Germany50 ones
- Recall that the average node degree of CORONET CONUS is much lower than the one of Germany50
- So, the number of paths between pairs of nodes is smaller which, in turn, makes the problems easier to solve

Conclusions

- We have exploited path geodiversity in anycast communications to enhance the network robustness against natural disasters.
- We have defined and solved the minimum cost D-geodiverse anycast routing problem with optimal selection of anycast nodes.
- We have extended it to consider the existence of vulnerability regions.
- We have presented computational results based on two well-known network topologies using real information of their hazard regions.

Conclusions

- We were able to compute the optimal solutions for all cases of interest.
- The results showed that, in general, improving the robustness to natural disasters:
  - increases the routing costs and
  - requires a higher minimum number of anycast nodes.
- A careful characterization of the vulnerability regions allows the operator to achieve:
  - either improved robustness with the same cost and number of anycast nodes
  - or reduced number of anycast nodes and routing costs for the same robustness

Thank you for your attention!

cost
EUROPEAN COOPERATION IN SCIENCE AND TECHNOLOGY

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

FCT
Fundação para a Ciência e a Tecnologia
MINISTÉRIO DA CIÊNCIA, TECNOLOGIA E ENSINO SUPERIOR

CENTRO-01-0145-FEDER-029312 - ResNeD
UID/EEA/50008/2013