Fast Control-Path Recovery After Failures In SDN

Sedef Savas

Motivation

- Even a single physical failure affects several connections between controllers and switches (***control paths***).

- When switches lose their connection to controller, they become "**uncontrolled**", no longer receives new instructions, hence although the switches are physically operational, they can not be used to establish new data traffic or reroute the disrupted data traffic.

- ***Control-Path Routing*** (path establishment for all control paths) affects the restoration time of disrupted data traffic after failures.

Multi-Stage Control-Path Recovery

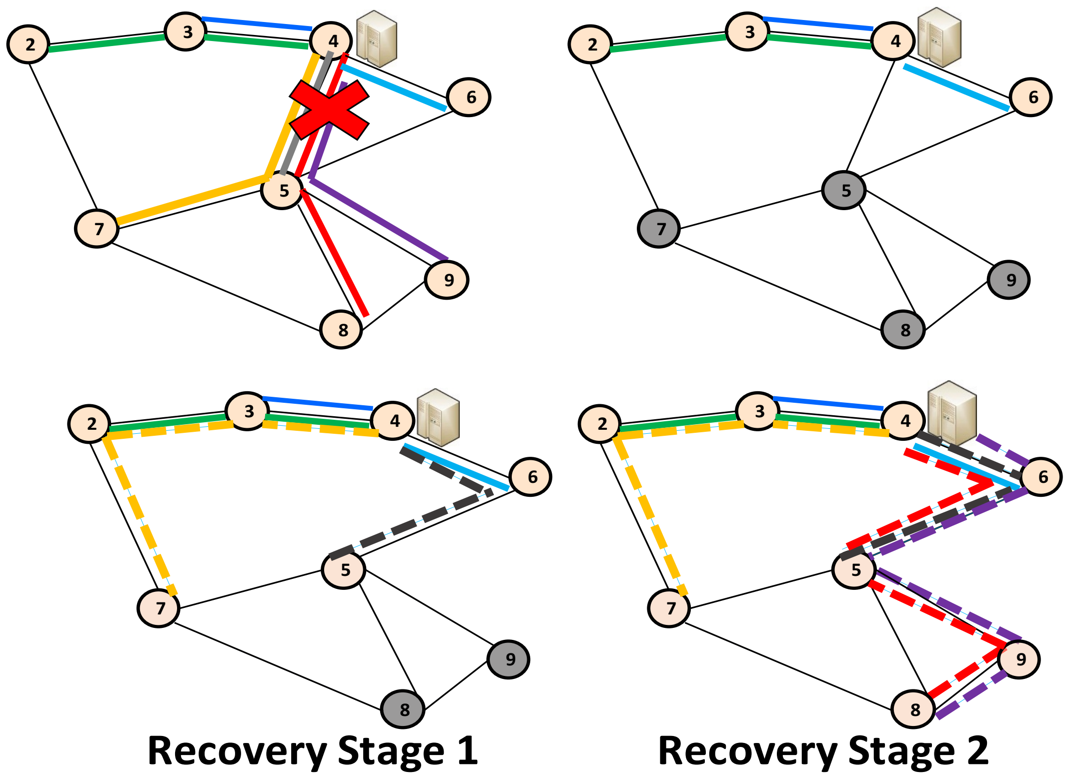

2 consequences of multi-stage control-path recovery

1) Longer data-path recovery times
Disrupted data paths that are destined to node 9 has to wait until stage 2 finishes.

2) Longer data paths
Data paths that will be recovered before all control paths are restored takes longer routes as not all the physically functional nodes will not be available until the end of last stage.

Multiple nodes can be re-connected to their controller in a single stage.

Related Works

- There are control-path routing studies, but they do not consider the multi-stage recovery of control paths. Mostly, shortest path routing.

- Hence, they don't show the effect of control-path routing on data-path recovery times.

Goal

Fast data-path recovery requires fast control-path recovery.

After a failure, to reach uncontrolled nodes, controller needs a **fully-controlled path**, where all nodes of that path is controlled, to establish flow.

Multi-stage recovery is required because after a failure, some uncontrolled nodes do not have a fully-controlled path from the controller.

Route control paths such that after any failure considered, maximum number of stages required to restore all control paths are minimum.

Stage 3

Link 3-5 fails.
Nodes 10, 11, 1, 3, and 4 is uncontrolled.

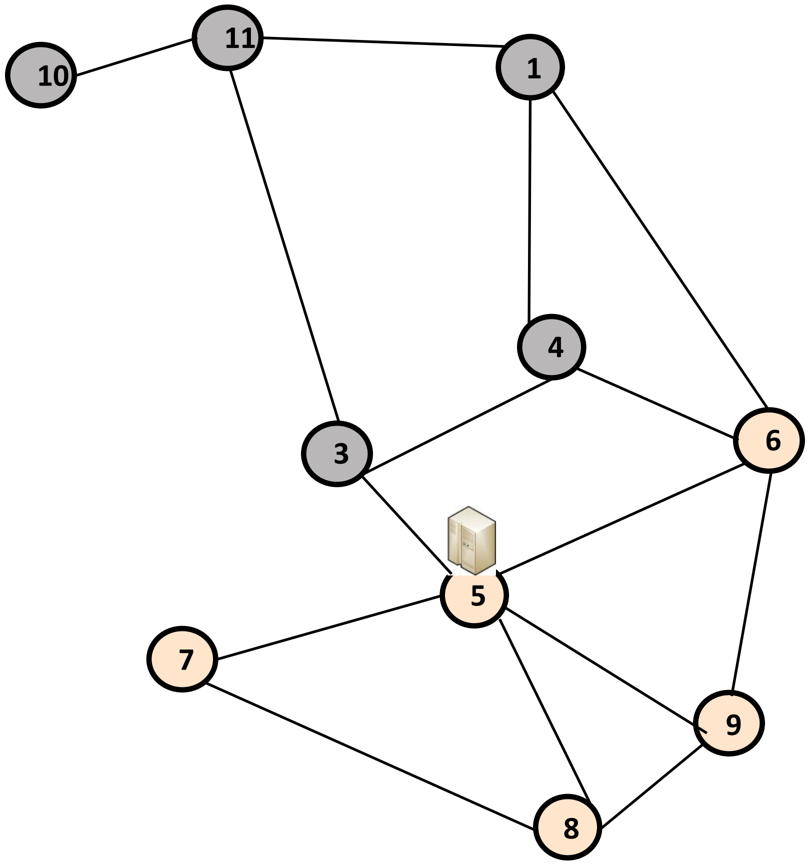

What is the minimum number of stages required to restore node 10's control path?

Requirement: To restore a control path, we need **a fully-controlled path** between node and controller.

Look all paths from node 10 to controller (node 5). Find the path that will be recovered soonest.

Remark 1: After a failure, number of stages required to recover a disrupted control path is equal to the # of uncontrolled nodes on a path, which has the minimum uncontrolled nodes on.

Why not reaching to any controlled node is not enough to recover a control path?

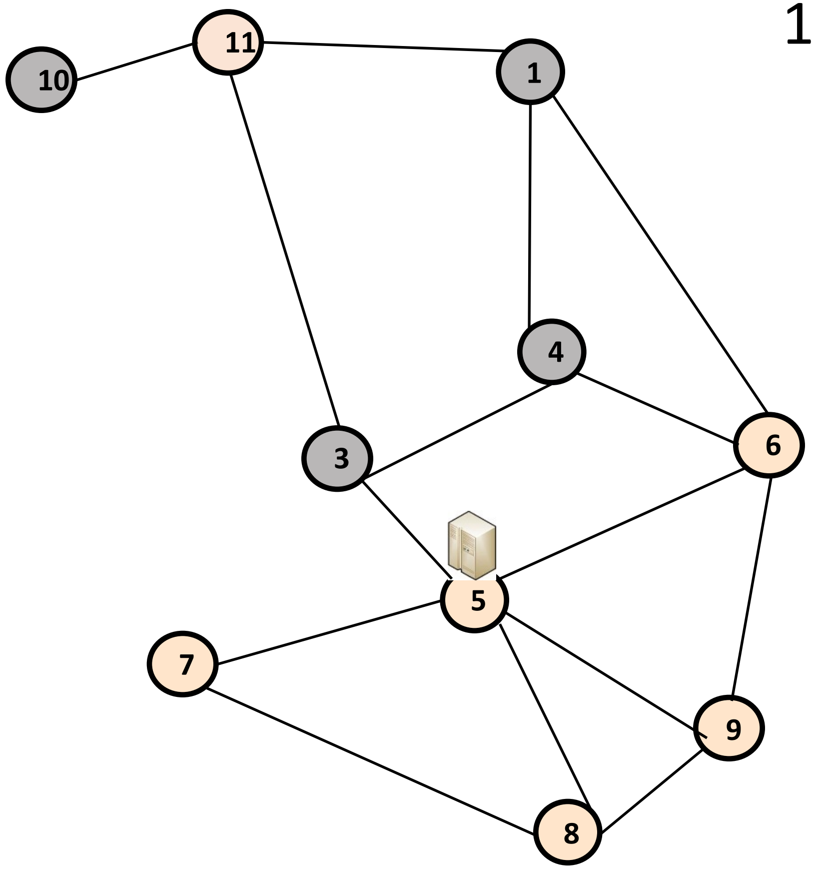

Link 5-6 fails.

10, 1, 3, and 4 is uncontrolled.

Max # of stages required is SPANNING TREE but that tree is changing with the failures.

Does min # of uncontrolled node on a restoration path give the minimum number of stages?
How does multiple control-path recovery in a single stage affect this?

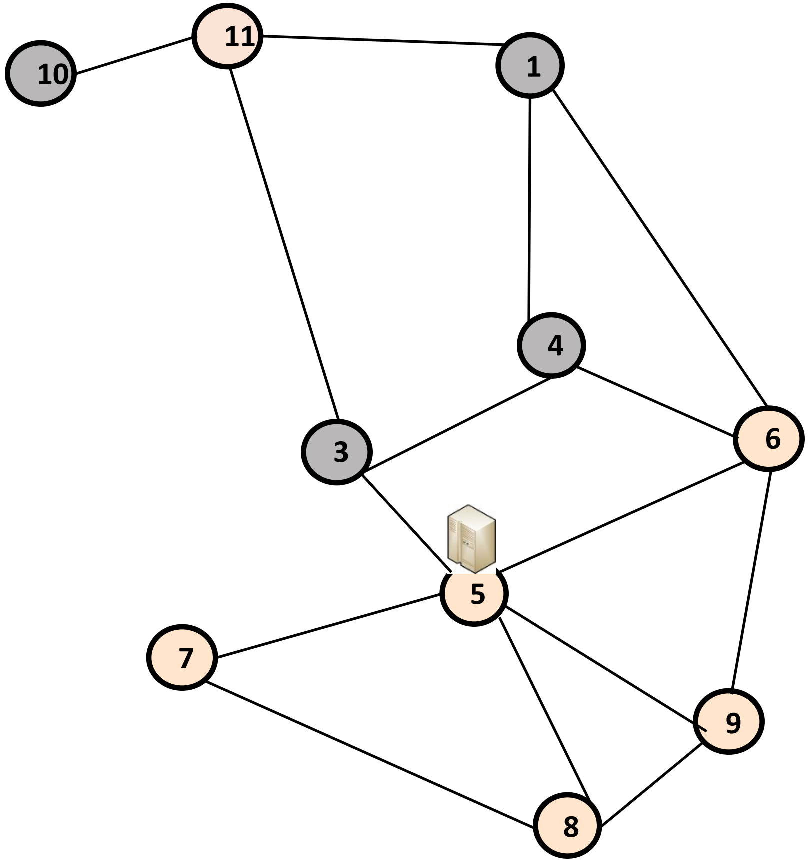

Yes, starting from the uncontrolled nodes closest to the controller, select the best path to restore that node. All the best paths on the restoration path of node 10 is selectable from node 10.

Remark 2: Number of stages required to restore all control paths after any failure considered is equal to the control path's recovery stage # that has the highest # of uncontrolled node on its restoration path.

# of recovery stages for node v after y =
(min # of uncontrolled nodes on restoration path p of node v after y)

Max # of recovery stages for node v after any failure =
max (# of recovery stages for node v after y) $\forall y \in Y$

# of recovery stages for any nodes after any failure =
Max(Max # of recovery stages for node v after any failure) $\forall v \in V$

Problem Statement

Given the network topology, single controller's location, and set of failures, find a *control path routing* that minimizes the maximum number of stages required to recover any control path after any failure considered.

To achieve this objective:

1. Find a *control-path routing* with the objective of minimum resource consumption such that after any failures in given failure set, any control path can be recovered within K stages, where K ∈ { 1, 2, ... (N - # of controllers)}, where N is number of nodes.
2. Starting from K = 1 and increasing it at each try, solve the previous problem until a solution is found. Find the smallest K that has a solution.

As a result, we find a routing which requires minimum number of stages after any failure and consumes minimum bandwidth resources.

Mathematical Formulation

Objective:

Minimize total resource usage of control paths.

Subject to:

After any failure y, there must exist at least 1 path from all uncontrolled nodes to a controller, where that path has at most K uncontrolled nodes on it.

Abilene network (11 node)
Controller is on node 7

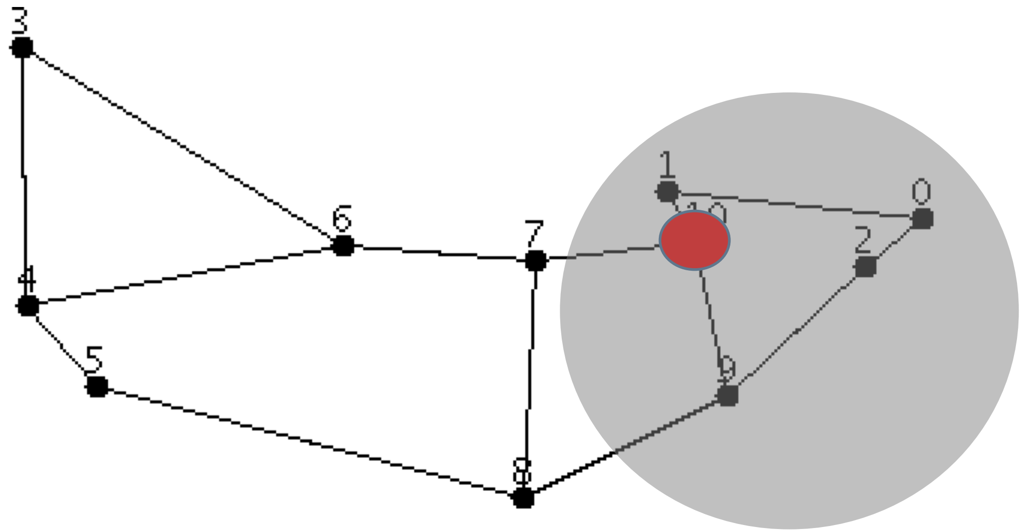

Control Path Routing
0 - 1 - 10 – 7
1 - 10 – 7
2 - 9 - 10 – 7
3 - 6 – 7
4 - 6 – 7
5 - 8 – 7
6 – 7
8 – 7
9 - 10 – 7
10 - 7

Node 1-0-2-9 lost control paths.

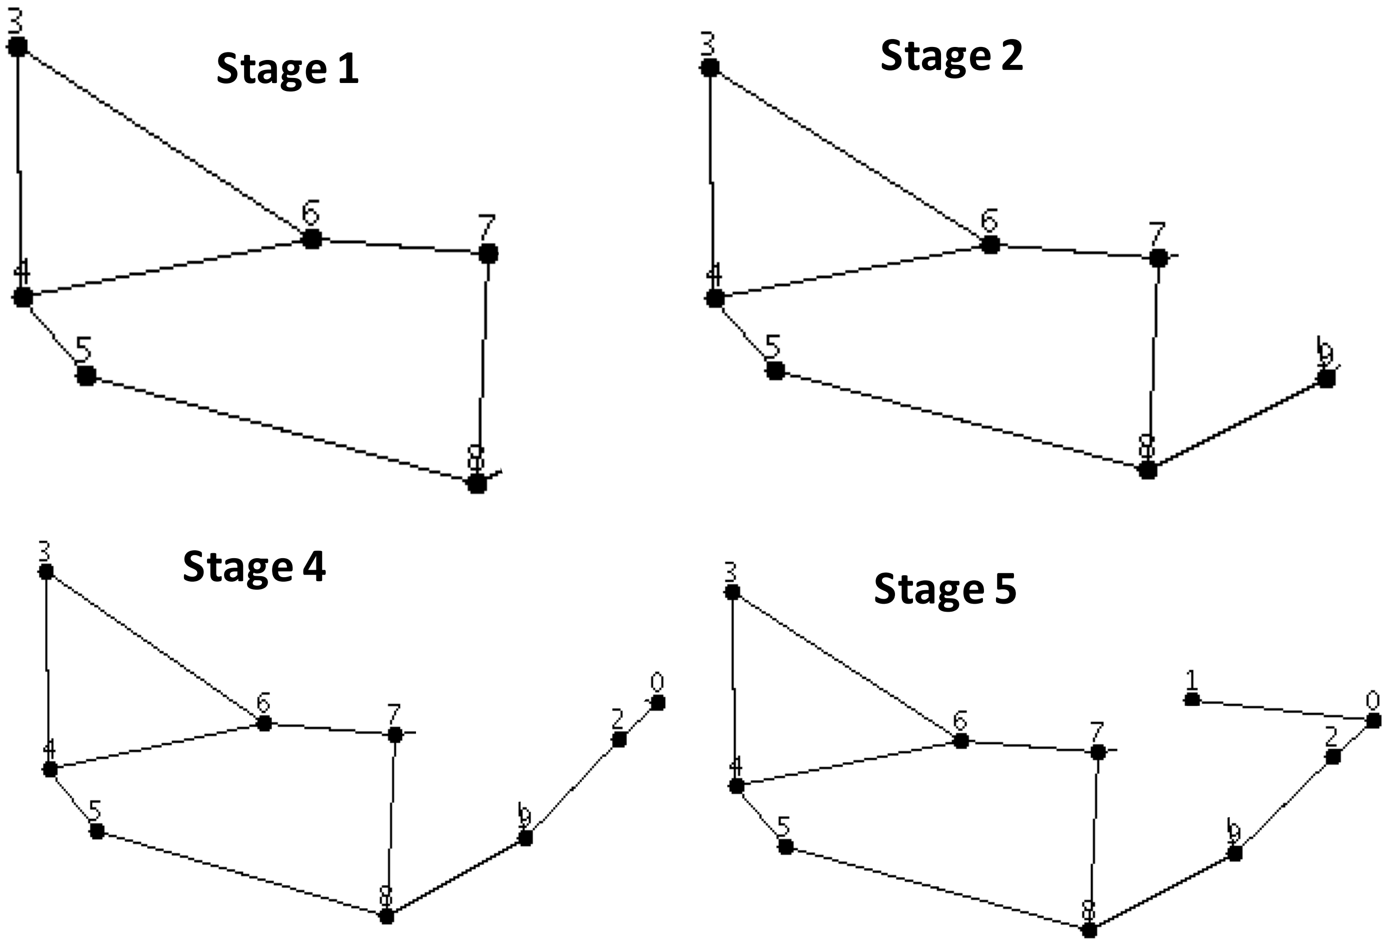

	Stage 1	Stage 2	Stage 3	Stage 4	Stage 5
Restored Control Paths		Node 9	Node 2	Node 0	Node 1
Recovered Path Number	0	6	6	8	14

K =2 after any node or link failure

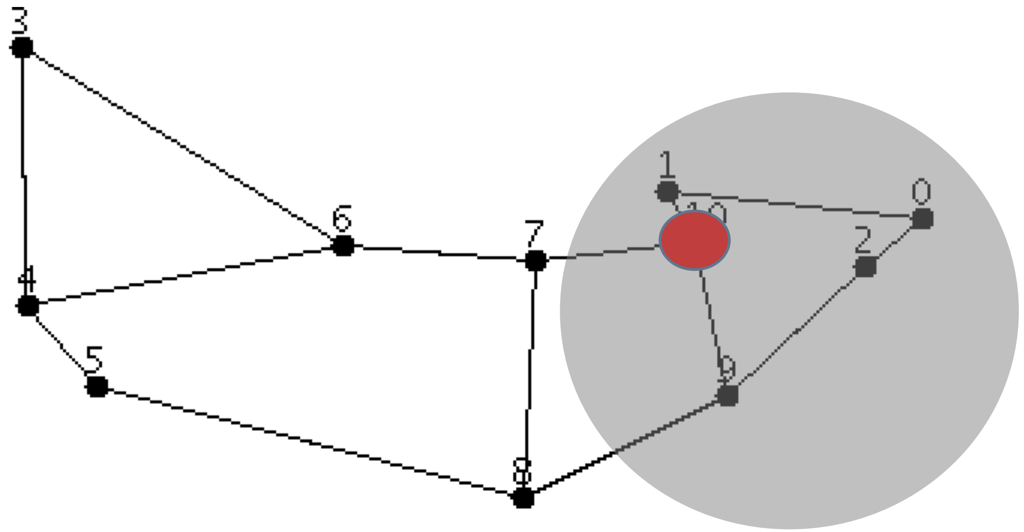

Node 1-0-2-9 lost control paths.

Shortest path: Control Path Routing
0 - 1 - 10 – 7
1 - 10 – 7
2 - 9 - 10 – 7
3 - 6 – 7
4 - 6 – 7
5 - 8 – 7
6 – 7
8 – 7
9 - 10 – 7
10 - 7

NEW Control Path Routing
0 - 1 - 10 - 7
1 - 10 - 7
2 - 9 - 8 - 7
3 - 6 - 7
4 - 6 - 7
5 - 8 - 7
6 - 7
8 - 7
9 - 8 - 7
10 - 7

Minimum # of controller for K=1 (Link Failure)

CONTROLLER LOCATION

- 1, 6
- 0, 6
- 2,6

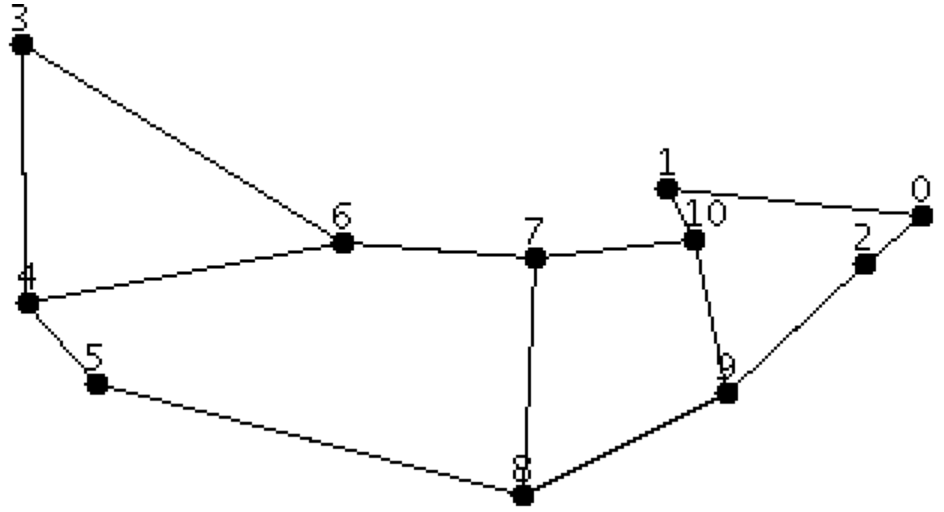

- Node 0 controlpath: 0 - 1
- Node 2 controlpath: 2 - 9 - 10 - 1
- Node 3 controlpath: 3 - 6
- Node 4 controlpath: 4 - 6
- Node 5 controlpath: 5 - 4 - 6
- Node 7 controlpath: 7 - 6
- Node 8 controlpath: 8 - 7 - 6
- Node 9 controlpath: 9 - 10 - 1
- Node 10 controlpath: 10 – 1

K = 2 , 2 controllers are needed located 4 and 7.
14 nodes are connected to controller 7 and 24 nodes are controlled by controller 4.

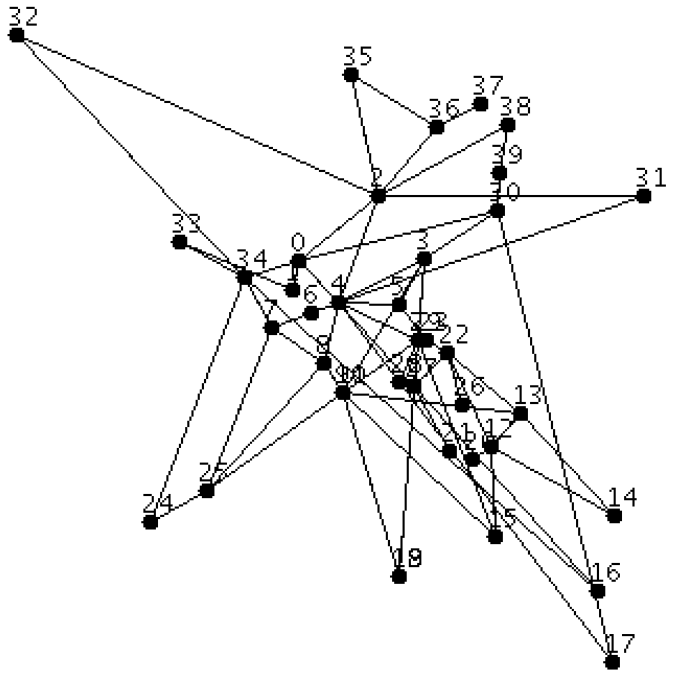

Node 0 controlpath: 0 - 34 - 7
Node 1 controlpath: 1 - 0 - 4
Node 2 controlpath: 2 - 4
Node 3 controlpath: 3 - 4
Node 5 controlpath: 5 - 4
Node 6 controlpath: 6 - 4
Node 8 controlpath: 8 - 7
Node 9 controlpath: 9 - 8 - 7
Node 10 controlpath: 10 - 3 - 4
Node 11 controlpath: 11 - 13 - 22 - 23 - 29 - 4
Node 12 controlpath: 12 - 15 - 9 - 8 - 7
Node 13 controlpath: 13 - 22 - 23 - 29 - 4
Node 14 controlpath: 14 - 13 - 22 - 23 - 29 - 4
Node 15 controlpath: 15 - 9 - 8 - 7
Node 16 controlpath: 16 - 4
Node 17 controlpath: 17 - 4
Node 18 controlpath: 18 - 9 - 8 - 4
Node 19 controlpath: 19 - 3 - 4
Node 20 controlpath: 20 - 12 - 22 - 23 - 29 - 4
Node 21 controlpath: 21 - 27 - 28 - 29 - 4
Node 22 controlpath: 22 - 23 - 5 - 4
Node 23 controlpath: 23 - 5 - 4
Node 24 controlpath: 24 - 25 - 7
Node 25 controlpath: 25 - 7
Node 26 controlpath: 26 - 22 - 23 - 29 - 4
Node 27 controlpath: 27 - 28 - 29 - 4
Node 28 controlpath: 28 - 29 - 9 - 8 - 7
Node 29 controlpath: 29 - 9 - 8 - 7
Node 30 controlpath: 30 - 0 - 34 - 7
Node 31 controlpath: 31 - 4
Node 32 controlpath: 32 - 34 - 7
Node 33 controlpath: 33 - 34 - 7
Node 34 controlpath: 34 - 7
Node 35 controlpath: 35 - 2 - 4
Node 36 controlpath: 36 - 2 - 0 - 34 - 7
Node 37 controlpath: 37 - 36 - 2 - 4
Node 38 controlpath: 38 - 2 - 4
Node 39 controlpath: 39 - 30 - 3 - 4

Shortest path (K = 3) vs our approach (K = 2).

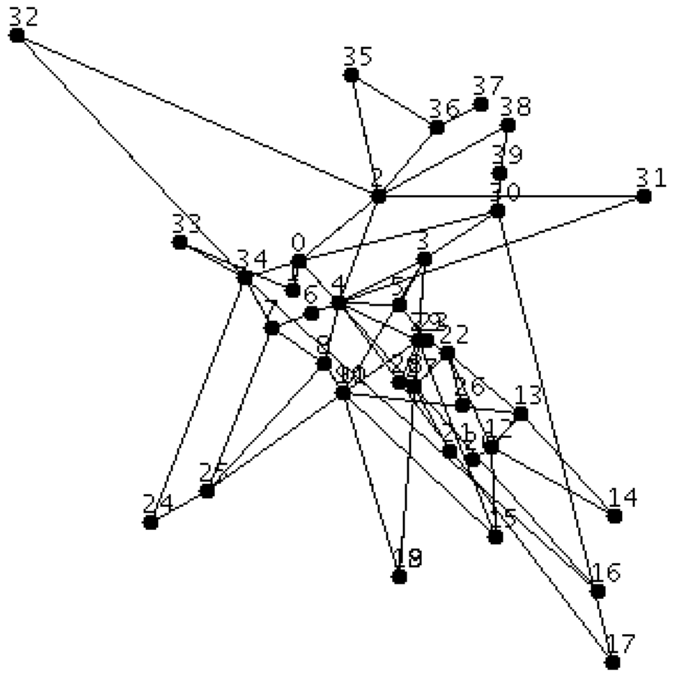

Node 0 controlpath: 0 - 4
Node 1 controlpath: 1 - 0 - 4
Node 2 controlpath: 2 - 4
Node 3 controlpath: 3 - 4
Node 5 controlpath: 5 - 4
Node 6 controlpath: 6 - 4
Node 8 controlpath: 8 - 4
Node 9 controlpath: 9 - 8 - 4
Node 10 controlpath: 10 - 3 - 4
Node 11 controlpath: 11 - 13 - 22 - 23 - 29 - 4
Node 12 controlpath: 12 - 22 - 23 - 29 - 4
Node 13 controlpath: 13 - 22 - 23 - 29 - 4
Node 14 controlpath: 14 - 13 - 22 - 23 - 29 - 4
Node 15 controlpath: 15 - 29 - 4
Node 16 controlpath: 16 - 4
Node 17 controlpath: 17 - 4
Node 18 controlpath: 18 - 9 - 8 - 4
Node 19 controlpath: 19 - 3 - 4
Node 20 controlpath: 20 - 12 - 22 - 23 - 29 - 4

Node 0 controlpath: 0 - 34 - 7
Node 1 controlpath: 1 - 0 - 4
Node 2 controlpath: 2 - 4
Node 3 controlpath: 3 - 4
Node 5 controlpath: 5 - 4
Node 6 controlpath: 6 - 4
Node 8 controlpath: 8 - 7
Node 9 controlpath: 9 - 8 - 7
Node 10 controlpath: 10 - 3 - 4
Node 11 controlpath: 11 - 13 - 22 - 23 - 29 - 4
Node 12 controlpath: 12 - 15 - 9 - 8 - 7
Node 13 controlpath: 13 - 22 - 23 - 29 - 4
Node 14 controlpath: 14 - 13 - 22 - 23 - 29 - 4
Node 15 controlpath: 15 - 9 - 8 - 7
Node 16 controlpath: 16 - 4
Node 17 controlpath: 17 - 4
Node 18 controlpath: 18 - 9 - 8 - 4
Node 19 controlpath: 19 - 3 - 4
Node 20 controlpath: 20 - 12 - 22 - 23 - 29 - 4

Shortest path (K = 3) vs our approach (K = 2).

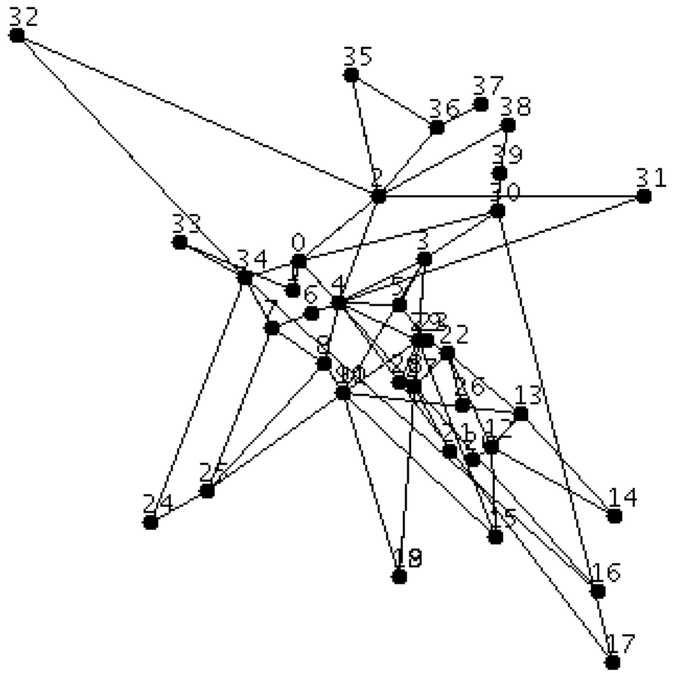

Node 21 controlpath: 21 - 27 - 28 - 29 - 4
Node 22 controlpath: 22 - 23 - 29 - 4
Node 23 controlpath: 23 - 29 - 4
Node 24 controlpath: 24 - 25 - 7
Node 25 controlpath: 25 - 7
Node 26 controlpath: 26 - 22 - 23 - 29 - 4
Node 27 controlpath: 27 - 28 - 29 - 4
Node 28 controlpath: 28 - 29 - 4
Node 29 controlpath: 29 - 4
Node 30 controlpath: 30 - 3 - 4
Node 31 controlpath: 31 - 4
Node 32 controlpath: 32 - 34 - 7
Node 33 controlpath: 33 - 34 - 7
Node 34 controlpath: 34 - 7
Node 35 controlpath: 35 - 2 - 4
Node 36 controlpath: 36 - 2 - 4
Node 37 controlpath: 37 - 36 - 2 - 4
Node 38 controlpath: 38 - 2 - 4
Node 39 controlpath: 39 - 30 - 3 - 4

→

Node 21 controlpath: 21 - 27 - 28 - 29 - 4
Node 22 controlpath: 22 - 23 - 5 - 4
Node 23 controlpath: 23 - 5 - 4
Node 24 controlpath: 24 - 25 - 7
Node 25 controlpath: 25 - 7
Node 26 controlpath: 26 - 22 - 23 - 29 - 4
Node 27 controlpath: 27 - 28 - 29 - 4
Node 28 controlpath: 28 - 29 - 9 - 8 - 7
Node 29 controlpath: 29 - 9 - 8 - 7
Node 30 controlpath: 30 - 0 - 34 - 7
Node 31 controlpath: 31 - 4
Node 32 controlpath: 32 - 34 - 7
Node 33 controlpath: 33 - 34 - 7
Node 34 controlpath: 34 - 7
Node 35 controlpath: 35 - 2 - 4
Node 36 controlpath: 36 - 2 - 0 - 34 - 7
Node 37 controlpath: 37 - 36 - 2 - 4
Node 38 controlpath: 38 - 2 - 4
Node 39 controlpath: 39 - 30 - 3 - 4

Controller is on node 0 (has average node degree)

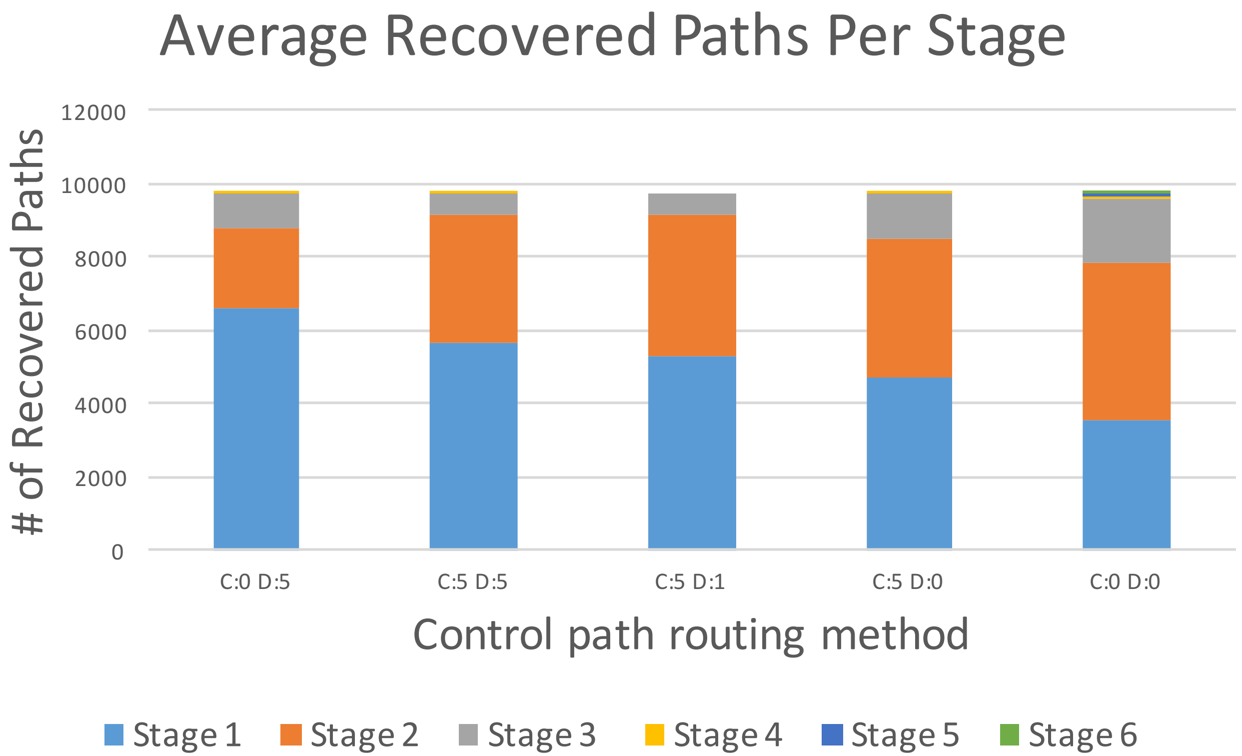

	Stage 1	Stage 2	Stage 3	Stage 4	Stage 5	Stage 6
C:0 D:5	6569	2247	886	58		
C:5 D:5	5640	3546	522	52		
C:5 D:1	5322	3808	630			
C:5 D:0	4680	3791	1263	26		
C:0 D:0	3536	4310	1726	66	74	48

Which questions are important?

- What is the minimum # of stages achievable when controllers and set of failures are given?

- To restore control paths in K stages after any failure in the failure set given, how many controllers are required and where to place them.