

A Tight Characterization of Fast Failover Routing: Resiliency to Two Link Failures is Possible

Wenkai Dai (University of Vienna, Austria), Klaus-Tycho Foerster (TU Dortmund, Germany), and Stefan Schmid (TU Berlin, Germany)





Motivation: Control-Plane Is Relatively Slow

- Traditionally, control-plane triggers recomputation of routing tables to tackle link-failures
 E.g., OSPF and IS-IS
 - \circ Very slow (convergence time $\approx 100 ms$), unacceptable for critical applications





Fast Rerouting (FRR) on Data-Plane

- Modern communication networks support fast reroute: local failover without invoking control plane, by pre-installing alternative paths
- Challenge: conditional forwarding rules can only depend on local information







- Network is a connected undirected graph G = (V, E)
- Routing function at node v may only match on:
 - 1. Destination *t*





- Network is a connected undirected graph G = (V, E)
- Routing function at node v may only match on:
 - 1. Destination *t*
 - 2. Incoming port from $E(v) \cup \{\bot\}$





- Network is a connected undirected graph G = (V, E)
- Routing function at node v may only match on:
 - 1. Destination *t*
 - 2. Incoming port from $E(v) \cup \{\bot\}$
 - 3. Incident edge failures $F \cap E(v)$





- Network is a connected undirected graph G = (V, E)
- Routing function at node v may only match on:
 - 1. Destination *t*
 - 2. Incoming port from $E(v) \cup \{\bot\}$
 - 3. Incident edge failures $F \cap E(v)$

To determine an outgoing link in E(v) - F





- Network is a connected undirected graph G = (V, E)
- Routing function at node v may only match on:
 - 1. Destination *t*
 - 2. Incoming port from $E(v) \cup \{\bot\}$
 - 3. Incident edge failures $F \cap E(v)$
 - 4. Optional: Source s (our model)

To determine an outgoing link in E(v) - F





- Network is a connected undirected graph G = (V, E)
- Routing function at node v may only match on:
 - 1. Destination *t*
 - 2. Incoming port from $E(v) \cup \{\bot\}$
 - 3. Incident edge failures $F \cap E(v)$
 - 4. Optional: Source s (our model)
 - To determine an outgoing link in E(v) F
- Static routing tables, deterministic behavior, no rewriting bits in packets



Goal: Computing a *k***-Resilient Routing Scheme**

- Given an undirected graph G = (V, E) and a destination $t \in V$
- A routing scheme is a set of routing functions defined on V
- Let $F \subseteq E(G)$ denote k arbitrary failed edges in G
- A routing scheme is called k-Resilient $(1 \le k \le |E|)$:
 - any $v \in V$ can reach a destination $t \in V$
 - as long as v t is still connected in G F
- Goal: Compute a k-resilient routing scheme for given $G = (V, E), t \in V$ and k



Related Works

- Fast failover is a widely studied area, see recent survey
 - Chiesa et al. (2021): A Survey of Fast Recovery Mechanisms in Packet-Switched Networks
- Randomized resilient routing has also been studied (not our model here)





Related Works (Without Matching Source)

Per- destination	Per source	Incoming port	Incident links	<i>k</i> -edge- connected graph	Packet rewriting bits	Resiliency	Ref.
X			X			No	[Kwong et al. 2011]
X		X	X			1	[Feigenbaum et al. 2012]
X		X	X			No 2-resilient	[Chiesa et al. 2016]
X		X	X	$k \leq 5$		(k - 1)	[Chiesa et al. 2016]
X		X	X	$k \to \infty$		[<i>k</i> /2]	[Chiesa et al. 2016]
X		X	X	k > 5		?	Open
X		X	X	<i>k</i> > 5	3 bits	(k - 1)	[Chiesa et al. 2016]
X		X	X	<i>k</i> > 5	log k bits	(k - 1)	[Chiesa et al. 2016]



Related Works (Matching Source)

Per- destination	Per source	Incoming port	Incident links	<i>k</i> -edge- connected graph	Packet rewriting bits	Resiliency	Ref.
X	X	X	X	$k \to \infty$		(k - 1)	[Foerster et al. 2019]
X	X	X	X			No ∞-resiliency	[Foerster et al. 2021]
X	X	X	X			1	[Feigenbaum et al. 2012]
X	X	X	X			? = 2, 3,	Open
X	Х	Х	X			2	[This paper]
Х	Х	X	X			No 3-resiliency	[This paper]





Summary of Our Results

• A 2-resilient source-matched routing scheme is computable within $O(n \cdot m)$, where m = |E| and n = |V|



 No ≥ 3-resilient source-matched routing scheme in general graphs





3-Resilient Source-Matched Routing Is Impossible

 In the following network, the source s cannot reach t after three link failures even if s - t is connected







 From now on, k-connected always means k-edge-connected

Computing a 2-Resilient Source-Matched Routing Scheme

- 2-resiliency is obvious if s t is at least 3-connected in G
- Find 3 edge-disjoint s t paths: {P₁, P₂, P₃}
- Routing on $\{P_1 > P_2 > P_3\}$ sequentially against 2 failures



 P_1



Reduced to sub-instances: Source and destination are exactly 2-connected

Reduction to Instances Where s - t Is Exactly 2-Connected

- If s t is 1-connected in G, by removing all bridges, s t impact connected component is at least 2-connected
- Obvious when s t is 3-connected





 P_2

Still Challenging When s - t Is Exactly 2-Connected in G

- Two edge-disjoint s t paths P_1 and P_2 are not enough for two failures P_1
- A path avoiding failures might exist
- Hard to determine useful structure in general graphs
- Further reduction:

• Assume two edge-disjoint s - t paths P_1 and P_2 in G are also node-disjoint (splitting articulation points)



Obtain Connected Components of *G* **after Two Failures**

- For any $\{e_1, e_2\} \subset E$, if s t disconnected in $G \{e_1, e_2\}$, let $E^c = \bigcup \{e_1, e_2\}$
- Obtain a graph $G E^{c}$ containing a set of connected components \mathbb{C}





Finding Gadgets for Each Connected Component $C_i \in \mathbb{C}$

- For each $C_i \in \mathbb{C}$ in G, our algorithm can find a gadget \mathcal{G}_i of C_i :
 - A gadget \mathcal{G}_i is a **subgraph** of $\mathcal{C}_i \in \mathbb{C}$, i.e., $\mathcal{G}_i \subseteq \mathcal{C}_i$
 - ∘ By replacing C_i with G_i in G, the s t connectivity remains unchanged against two failures $\{e_1, e_2\} \subset E$





Representing Gadgets for Components

• Gadgets for all components $C_i \in \mathbb{C}$ can be represented by the following simple patterns





Representing Gadgets for Components (2)

• If $C_i \in \mathbb{C}$ contains s or t, its gadget can be represented by these patterns





Obtain a Kernel Graph Based on Gadgets

- In a graph G, replacing each component $C_i \in \mathbb{C}$ by its gadget \mathcal{G}_i
- Obtain a *kernel graph* $G \coloneqq \bigcup_{c_i \in \mathbb{C}} g_i \cup E^c$, where $G \subseteq G$
- Fact: The *s t* connectivity remains the same in *G* as *G* after any two failures





Determine 2-Resilient Source-Matched Routing in Kernel Graph

- If a gadget has *crossing paths*, simplify it by ignoring one path
- Now, simplified kernel graph is a planar graph, where paths $Q_1 \cup Q_2$ are the **boundary** of the **outer face** on the drawing.







Determine 2-Resilient Source-Matched Routing in Kernel Graph

- If a gadget has *crossing paths*, simplify it by ignoring one path
- Now, simplified kernel graph is a planar graph, where paths $Q_1 \cup Q_2$ are the **boundary** of the **outer face** on the drawing.









Summary

• A 2-resilient source-matched routing scheme is computable within $O(n \cdot m)$, where m = |E| and n = |V|



 No ≥ 3-resilient source-matched routing scheme in general graphs









A Tight Characterization of Fast Failover Routing: Resiliency to Two Link Failures is Possible

Wenkai Dai (University of Vienna, Austria), Klaus-Tycho Foerster (TU Dortmund, Germany), and Stefan Schmid (TU Berlin, Germany)



Slides by Wenkai Dai (University of Vienna)



Intuitions for Devising 2-Resilient Routing

- E.g., define routing function at v_0 as follows:
 - Then it cannot be 2-resilient since a routing loop occurs





Intuitions for Devising 2-Resilient Routing (2)

• Another routing function can easily avoid the loop







Correctness of 2-Resilient Routing in Simplified Kernel Graph

• Proof of Correctness:

- Starting at s, always walking on the outer face $s \rightarrow t$ of the residual graph up to 2 failures
- Each failure makes the detouring paths being new outer face of residual graph





Problem of Simplified Kernel Graph



- Drawback of simplification:
 - $\circ s t$ is still connected in the example, but routing cannot reach t anymore





Additional Routing Rule for Simplified Gadgets

- First, recover the ignored crossing paths
- Additional Rule: add one directed path as shown in the example (green path)
- Correctness: only used after seeing two failures, and directed path cannot introduce a loop



Recover