International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Fault Tolerance of Mesh Interconnects with Sparse Shortcut Links: Minimum Cuts, Failure Probability and Deadlock-Free Rerouting

Roza Darwish*, Nasser Aba Rous**
* Holds a Master's degree in Electronics Technology Engineering, Department of Computer and Electronic Systems Engineering, Faculty of Information and Communications Technology Engineering, Tartous university, Syria.
** holder a Master's degree in Electronics, Faculty of Electrical Engineering, University of Aleppo.

Abstract - Do sparse shortcut (wrap-around) links make a mesh more reliable, and not only faster? We study two-dimensional k×k meshes with zero, two or all wrap links (mesh, edge pair, center pair, frame, torus) under random link and node failures. We prove that for all five designs every minimum edge cut isolates a single node, so the edge connectivity equals the minimum degree (2 for the mesh, edge pair and center pair; 3 for the frame; 4 for the torus) and the number of minimum cuts equals the number of minimum-degree nodes. Exhaustive enumeration at k = 4 and 6 confirms this. It gives a closed-form isolation estimator of the disconnection probability, which differs from Monte Carlo by at most 0.2 points for q ≤ 0.05 (k = 8, 16, 32), whereas the leading-order formula based on the minimum cuts alone is off by up to 2.2 points, because boundary nodes of degree 3 matter at scale. Two shortcut links do not improve reliability: the center pair behaves statistically like the mesh (k = 32, q = 0.05: P(connected) 0.971 vs 0.968), and the edge pair only helps by lifting one corner. Failed links barely lengthen the mean path (stretch ≤ 3.2% at q = 0.05), but deadlock-free up∗/down∗ rerouting on the damaged graph makes legal paths 6–29% longer at k = 16 with a central root (q = 0.05), and up to 38% with a corner root. After 8 of roughly 120 links fail at N = 64, the capacity under optimal routing drops by 15–23%. We report these findings even where the hybrids give no advantage, and release all code.

Keywords - interconnection networks, fault tolerance, edge connectivity, network reliability, up∗/down∗ routing, mesh, torus, HPC.

1. INTRODUCTION

Large HPC systems regularly lose links and nodes [1], so an interconnect must stay connected and usable in the presence of failures. Meshes are inexpensive but their boundary nodes have low degree; tori have uniform degree four but require $2k = 2\sqrt{N}$ wrap links, each reaching around the entire machine. A handful of shortcut links is known to capture most of the path-length savings of a torus, but whether such sparse hybrids also capture its fault tolerance is a different matter, and is the focus of this paper.

We want to know four things. First, what are the exact edge connectivity and minimum cuts of these graphs? Second, how likely is the network to disconnect under random failures, and can we find a closed form for it? Third, how much longer do the surviving paths become? Fourth, what does a deadlock-free routing function cost on the damaged graph, in path length and in capacity? Our contributions are:

1. Exact minimum cuts. For the mesh, edge pair, center pair, frame and torus (even $k \geq 4$) all minimum edge cuts are node-isolating (Proposition 1). We verified this exhaustively at k = 4 and 6, and by max-flow up to k = 16.
2. A predictive estimator, with its failure region stated. The standard leading-order estimate $n_\delta q^\delta$ [2] is accurate only while boundary nodes of degree 3 are insignificant. An isolation estimator over all nodes stays within 0.2 points for $q \leq 0.05$ at sizes up to k = 32.
3. Negative result. Two shortcut links do not improve connectivity over the plain mesh; only the frame (wraps on all four outer rows and columns) raises λ.
4. Cost of deadlock-free rerouting. We evaluate up∗/down∗ [3] on damaged graphs: path inflation, sensitivity to the root, and a capacity estimate, plus exact LP capacity after failures at N = 64.

2. RELATED WORK

Network reliability analysis is #P-hard in general [4], and small-failure expansions based on minimum cuts are the usual tool [2]; randomized min-cut algorithms [5] bound the number of near-minimum cuts. Deadlock-free routing on irregular or damaged topologies uses up∗/down∗ [3], the turn model [6], or acyclicity of the channel dependency graph [7], [8]. Failure statistics of real

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

systems are reported in [1]. Mesh/torus hybrids with shortcut links have been studied for latency and throughput [9], [10], [11], [12]; to our knowledge their fault tolerance has not been analyzed with exact cuts. Background can be found in [13], [14].

3. MODEL AND DESIGNS

Every design is an undirected graph on $N = k^2$ nodes. Mesh: all grid connections between neighbors. Torus: grid plus wrap links on every row and column. Edge pair: wrap links on row 0 and column 0 only. Centre pair: wrap links on row k/2 and column k/2. Frame: wrap links on rows 0 and k−1 and columns 0 and k−1. Link counts at k = 16 are 480, 482, 482, 484 and 512.

Failure models: (a) each link fails independently with probability q; (b) each node fails independently with probability q, which removes its incident links. Connectivity is measured on the surviving graph (surviving nodes for (b)). We report P(connected), the mean size of the largest component as a fraction of the survivors, and path statistics on connected samples.

4. EXACT EDGE CONNECTIVITY AND MINIMUM CUTS

Proposition 1. For the five designs and even $k \geq 4$, every minimum edge cut isolates a single node. Hence λ equals the minimum degree δ, and the number of minimum cuts equals the number n_δ of nodes of degree δ.

Evidence. The degree sequences are explicit: mesh corners have degree 2, boundary nodes 3, interior nodes 4; each wrap on a boundary row or column adds one to its two end nodes. We computed λ by max-flow for $k \leq 16$ and read λ, δ and n_δ off the degree sequence beyond that (Table 1). We then enumerated all λ-subsets of links at k = 4 and 6 (Table 2): in every case the number of disconnecting subsets equals n_δ and all of them isolate a node. We have no proof for general k, so the statement is a verified proposition, not a theorem.

Table 1. Edge connectivity λ, minimum degree δ and number of minimum-degree nodes n_δ.

k	Design	Links	δ	λ	n_δ
8	Mesh	112	2	2	4
8	Edge pair	114	2	2	1
8	Centre pair	114	2	2	4
8	Frame	116	3	3	24
8	Torus	128	4	4	64
16	Mesh	480	2	2	4
16	Edge pair	482	2	2	1
16	Centre pair	482	2	2	4
16	Frame	484	3	3	56
16	Torus	512	4	4	256
32	Mesh	1984	2	2	4
32	Edge pair	1986	2	2	1
32	Centre pair	1986	2	2	4
32	Frame	1988	3	3	120
32	Torus	2048	4	4	1024

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Table 2. Exhaustive enumeration of λ-link subsets: disconnecting subsets, node-isolating ones, and n_δ.

k	Design	λ	Disconnecting	Node-isolating	n_δ
4	Mesh	2	4	4	4
4	Edge pair	2	1	1	1
4	Centre pair	2	4	4	4
4	Frame	3	8	8	8
4	Torus	4	16	16	16
6	Mesh	2	4	4	4
6	Edge pair	2	1	1	1
6	Centre pair	2	4	4	4
6	Frame	3	16	16	16

The frame raises λ from 2 to 3 with only four wrap links, while the two-link designs keep λ = 2. The torus reaches λ = 4 with 2k wrap links (32 at k = 16).

5. Disconnection Probability

Leading-order formula. For small q, P(disconnected) ≈ $n_\delta q^\delta$ [2]. Isolation estimator. All other cut shapes are ignored and isolation events of different nodes are treated as independent, so P(connected) ≈ $\prod_v(1 - q^{\deg v})$. Both are compared with Monte Carlo (20,000 samples at k = 8 and 16; 6,000 at k = 32; 10,000 for node failures at k = 8).

Table 3. P(connected) under link failures: Monte Carlo / isolation estimator / λ-only formula, at q = 0.02, 0.05, 0.1.

k	Design	MC	Iso	λ	MC	Iso	λ	MC	Iso	Λ
16	Mesh	0.9977	0.9979	0.9984	0.9802	0.9819	0.9900	0.8764	0.8906	0.9600
16	Edge pair	0.9994	0.9991	0.9996	0.9885	0.9891	0.9975	0.9115	0.9160	0.9900
16	Centre pair	0.9978	0.9980	0.9984	0.9817	0.9824	0.9900	0.8790	0.8938	0.9600
16	Frame	0.9998	0.9995	0.9996	0.9911	0.9918	0.9930	0.9254	0.9268	0.9440
16	Torus	1.0000	1.0000	1.0000	0.9986	0.9984	0.9984	0.9750	0.9747	0.9744
32	Mesh	0.9975	0.9973	0.9984	0.9677	0.9698	0.9900	0.7705	0.7786	0.9600
32	Edge pair	0.9983	0.9985	0.9996	0.9753	0.9769	0.9975	0.7852	0.8007	0.9900
32	Centre pair	0.9975	0.9973	0.9984	0.9705	0.9703	0.9900	0.7563	0.7814	0.9600
32	Frame	0.9992	0.9989	0.9990	0.9792	0.9796	0.9850	0.8023	0.8102	0.8800
32	Torus	0.9998	0.9998	0.9998	0.9942	0.9936	0.9936	0.9020	0.9027	0.8976

Over all 45 settings with $q \le 0.05$, the worst absolute error of the isolation estimator is 0.22 points, against 2.23 points for the λ-only formula, which systematically over-estimates the reliability of mesh-like designs because it misses the 4(k−2) boundary nodes of degree 3. At q = 0.1 and k = 32 the estimators are less good (frame: MC 0.802, estimator 0.810), because failures of larger sets begin to matter; we regard $q \ge 0.1$ as outside their range. The Monte Carlo standard error is at most 0.007 at 6,000 samples.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

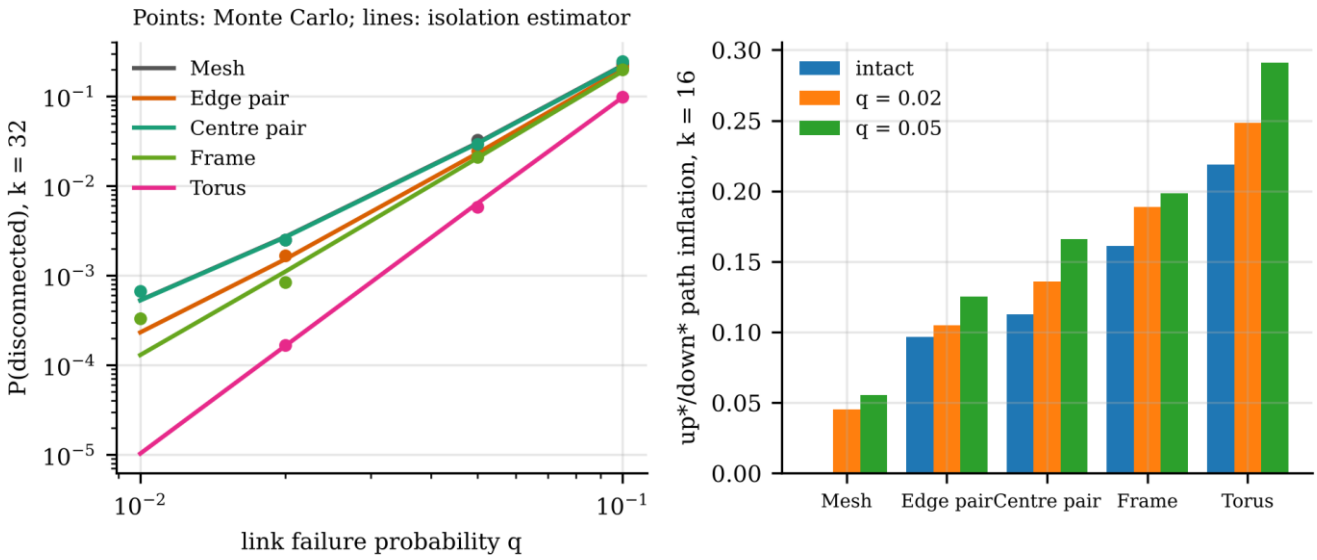


Fig. 1. Left: disconnection probability at k = 32 (points: Monte Carlo; lines: isolation estimator). Right: up∗/down∗ path inflation at k = 16 (root at the center).

The torus is clearly the most reliable at $q \geq 0.05$, with the frame second. The center pair is statistically equal to the mesh (k = 32, q = 0.05: 0.9705 vs 0.9677), while the edge pair is slightly better (0.9753) because it lifts one of the four corners from degree 2 to 3. Shortcuts that reduce latency do not improve reliability. One more result stands out: even when the graph disconnects, the largest component keeps at least 99.96% of the nodes at q = 0.05 for every design and size we ran, so most failures only cut off a corner node or two.

Table 4. k = 16: P(connected), link failures / node failures, at q = 0.02, 0.05, 0.1.

Design	Link 0.02	Node 0.02	Link 0.05	Node 0.05	Link 0.1	Node 0.1
Mesh	0.9977	0.9986	0.9802	0.9824	0.8764	0.8858
Edge pair	0.9994	0.9983	0.9885	0.9865	0.9115	0.9176
Centre pair	0.9978	0.9980	0.9817	0.9789	0.8790	0.8898
Frame	0.9998	0.9997	0.9911	0.9922	0.9254	0.9367
Torus	1.0000	1.0000	0.9986	0.9984	0.9750	0.9774

Node failures behave like link failures at the same q: the ranking and the effect of the boundary degree are the same.

6. PATH STRETCH AFTER FAILURES

Table 5. Connected samples: intact mean hops, mean-hop stretch at q = 0.05, mean diameter at q = 0.05, intact diameter.

k	Design	Mean hops	Stretch	Diameter (q=.05)	Diameter (intact)
8	Mesh	5.33	1.017	14.0	14
8	Edge pair	4.82	1.027	11.4	11
8	Centre pair	4.54	1.031	9.5	8
8	Frame	4.51	1.026	9.0	8
8	Torus	4.06	1.013	8.0	8
16	Mesh	10.67	1.009	30.0	30
16	Edge pair	9.69	1.020	23.9	23
16	Centre pair	9.02	1.022	18.6	16
16	Frame	9.09	1.020	17.9	16
16	Torus	8.03	1.008	16.1	16

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Mean path length is hardly changed (≤ 3.2% at q = 0.05). The diameter of the hybrids, however, grows more than the mesh diameter in relative terms, because their diametral pairs depend on the shortcut links; at k = 8 the mean diameter of the center pair rises from 8 to 9.5 at q = 0.05 (300 samples at k = 8, 120 at k = 16).

7. Deadlock-Free Rerouting with up∗/down∗

After a failure the routing must be recomputed without creating cyclic channel dependencies. We use up∗/down∗ [3]: nodes are ordered by BFS depth from a root (ties broken by identifier), a link is "up" toward the lower order, and a legal path is a sequence of up hops followed by down hops. Legal shortest paths are found by BFS over (node, phase) states. The scheme is deadlock-free because the channel dependency graph is acyclic [7]. We report mean legal hops against mean unrestricted shortest hops, and b_{est} = (N−1)/(maximum link load) when every ordered pair uses one deterministic legal path (a conservative estimate, not the optimum).

Table 6. k = 16, up∗/down∗ with center root unless stated: shortest and legal mean hops (intact), inflation (intact, corner root, q = 0.05 with center root; 6 failure samples), ρ_est intact and at q = 0.05.

Design	Short	Legal	Infl.	Infl. corner root	Infl. q=0.05	ρ_est	ρ_est q=.05
Mesh	10.67	10.67	0.0%	0.0%	5.5%	0.0279	0.0248
Edge pair	9.69	10.62	9.6%	37.7%	12.5%	0.0281	0.0253
Centre pair	9.02	10.04	11.3%	35.9%	16.6%	0.0347	0.0316
Frame	9.09	10.56	16.1%	33.0%	19.8%	0.0284	0.0246
Torus	8.03	9.79	21.9%	21.9%	29.1%	0.0347	0.0323

On the mesh, up∗/down∗ costs nothing with a central root. With shortcuts, legal paths are already 10–22% longer than shortest paths before any failure, the torus paying most (22%). The root makes a big difference: a corner root adds 22–38% to the hybrids, so it should be placed centrally. Failures add a further 3–7 points at q = 0.05, and ρ_{est} drops by 7–13%. Because ρ_{est} uses one fixed path per pair, it understates what a load-balancing implementation could achieve (Section 8). With 6 failure samples per cell at k = 16 and 25 at k = 8, differences of a point or two between cells are uncertain.

8. POST-FAILURE CAPACITY (N = 64)

We solve exact source-aggregated linear programs for the uniform-traffic capacity ρ under (a) minimal routing (any split over shortest paths of the damaged graph) and (b) arbitrary routing, after f random link failures, averaging 8 connected samples per cell (one for f = 0). Capacity is in flits/node/cycle, normalized so that the total demand per node is 1.

Table 7. N = 64: uniform-traffic capacity, minimal routing / optimal routing, after f random link failures.

Design	f=0 min	f=0 opt	f=2 min	f=2 opt	f=4 min	f=4 opt	f=8 min	f=8 opt
Mesh	0.492	0.492	0.457	0.457	0.435	0.435	0.403	0.409
Edge pair	0.321	0.554	0.324	0.516	0.330	0.496	0.356	0.429
Centre pair	0.260	0.554	0.262	0.500	0.261	0.492	0.263	0.463
Frame	0.432	0.615	0.385	0.582	0.373	0.547	0.357	0.526
Torus	0.984	0.984	0.915	0.915	0.890	0.891	0.815	0.816

Under optimal routing, 8 failed links (6–7% of the links) reduce capacity by 17% for the mesh, 23% for the edge pair, 16% for the center pair, 15% for the frame and 17% for the torus. Under minimal routing the two-link hybrids hardly change (plausibly because their bottleneck sits at the shortcut links, which random failures rarely hit; we did not test this), and the edge pair's minimal-routing capacity even rises slightly (0.321 to 0.356), because removing a link can remove a congested shortest path. This is a property of minimal routing, not a benefit of failures. The torus has the highest capacity at every f, and the gap between optimal and minimal routing in the hybrids persists after failures.

9. DISCUSSION AND LIMITATIONS

- Only independent random failures are modelled; correlated failures (a switch, a cable bundle, a power domain) can be far more damaging and are not included.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

- Proposition 1 is verified, not proved, for general k; for $k \geq 8$ the cut counts rely on max-flow for λ and on the degree sequence for n_δ.
- The isolation estimator ignores non-trivial cuts and assumes independence between nodes; it is reliable for $q \leq 0.05$ in our range and is not a bound.
- Monte Carlo sample sizes at k = 32 (6,000) give standard errors up to 0.007; stretch and up$*$/down$*$ use 6–300 samples, so differences of about a point are not significant.
- up$*$/down$*$ is only one of several deadlock-free schemes. Escape-channel adaptive routing [8] may recover some of the inflation but was not tested here. ρ_{est} uses a single path per pair.
- Capacity is evaluated by LP at N = 64 only, for uniform traffic, without cycle-accurate simulation.

In practice, when reliability matters, placing the shortcuts on the corners and boundary (the frame) is more important than the number of shortcuts; two center links give a latency benefit but do not affect connectivity.

10. CONCLUSION

For mesh interconnects with wrap links, the edge connectivity equals the minimum degree and all minimum cuts isolate a single node. This gives an accurate closed-form estimator of the disconnection probability for $q \leq 0.05$, shows that two shortcut links add no connectivity, and quantifies the cost of deadlock-free recovery: 6–29% longer paths (q = 0.05) and 15–23% lower optimal-routing capacity at N = 64 after 8 link failures. The torus is the most resilient; the frame is the best sparse compromise.

REPRODUCIBILITY

All code (ft.py, mc.py, updown.py, capfail.py, iso.py), all JSON outputs and the builder are in the accompanying archive; every number in the tables is generated from those JSON files. All random seeds are fixed. [Zenodo DOI to be added.]

REFERENCES

[1] B. Schroeder and G. A. Gibson, “A large-scale study of failures in high-performance computing systems,” IEEE Trans. Dependable Secure Comput., vol. 7, no. 4, pp. 337–350, Oct. 2010.

[2] C. J. Colbourn, The Combinatorics of Network Reliability. New York, NY, USA: Oxford Univ. Press, 1987.

[3] M. D. Schroeder et al., “Autonet: A high-speed, self-configuring local area network using point-to-point links,” IEEE J. Sel. Areas Commun., vol. 9, no. 8, pp. 1318–1335, Oct. 1991.

[4] M. O. Ball, “Computational complexity of network reliability analysis: An overview,” IEEE Trans. Reliab., vol. R-35, no. 3, pp. 230–239, Aug. 1986.

[5] D. R. Karger, “Minimum cuts in near-linear time,” J. ACM, vol. 47, no. 1, pp. 46–76, Jan. 2000.

[6] C. J. Glass and L. M. Ni, “The turn model for adaptive routing,” J. ACM, vol. 41, no. 5, pp. 874–902, Sep. 1994.

[7] W. J. Dally and C. L. Seitz, “Deadlock-free message routing in multiprocessor interconnection networks,” IEEE Trans. Comput., vol. C-36, no. 5, pp. 547–553, May 1987.

[8] J. Duato, “A new theory of deadlock-free adaptive routing in wormhole networks,” IEEE Trans. Parallel Distrib. Syst., vol. 4, no. 12, pp. 1320–1331, Dec. 1993.

[9] W. J. Dally, “Express cubes: Improving the performance of k-ary n-cube interconnection networks,” IEEE Trans. Comput., vol. 40, no. 9, pp. 1016–1023, Sep. 1991.

[10] U. Y. Ogras and R. Marculescu, “It’s a small world after all: NoC performance optimization via long-range link insertion,” IEEE Trans. VLSI Syst., vol. 14, no. 7, pp. 693–706, Jul. 2006.

[11] M. Koibuchi, H. Matsutani, H. Amano, D. F. Hsu, and H. Casanova, “A case for random shortcut topologies for HPC interconnects,” in Proc. ISCA, Portland, OR, USA, Jun. 2012, pp. 177–188.

[12] D. C. Jung, S. Davidson, C. Zhao, D. Richmond, and M. B. Taylor, “Ruche networks: Wire-maximal, no-fuss NoCs,” in Proc. IEEE/ACM NOCS, 2020.

[13] J. Duato, S. Yalamanchili, and L. Ni, Interconnection Networks: An Engineering Approach. San Francisco, CA, USA: Morgan Kaufmann, 2003.

[14] W. J. Dally and B. Towles, Principles and Practices of Interconnection Networks. San Francisco, CA, USA: Morgan Kaufmann, 2004.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)