

WHITE PAPER SERIES / WP-01 / SEPTEMBER 2026

Threshold blindness in degrading mesh networks

Averaged link metrics cannot predict when a network partitions,
because partition happens along the weakest path and the average
does not know where that path is

ACHYUT KUMAR / ABHINAV DUDDALA

DREITON DYNAMICS, ASHBURN, VIRGINIA

SUMMARY

A mesh network of drones, robots or vessels stays one network until the moment its links can no longer hold it together, and that moment is a threshold, not a slope. The metrics reported for such networks, average link availability, average signal strength, average neighbour count, move smoothly through it and take the same value on networks that are about to split and networks that are not. We call this threshold blindness.

In a simulated formation of 400 units with about 11 neighbours in radio range each, the average link availability was made to fall at a constant rate over a mission. For the first 71% of the reading's range nothing happened to connectivity. The network then went from whole, with nine tenths of its units in one group, to split, with half, while the reading moved from 29% to 19%, about a tenth of its range, over about a tenth of the mission.

Four networks with the same reading of 50% held 99%, 86%, 73% and 48% of their units in one group, depending only on where the failed links were. In the last case the formation was two networks of 48% and 39%.

Across 600 random scenarios, networks that had partitioned reported readings as high as 73% and networks that were whole reported readings as low as 24%. The best rule based on the average alone, chosen after the fact, called the outcome wrong 15% of the time. A reading that tracked the worst-connected tenth of the fleet was wrong 8% of the time, and a computation on the actual topology 6%. In predicting the size of the largest group, the average was off by 8.0 points on average and the topology computation by 3.5.

The reason is arithmetic. A network is connected only if every cut across it has working links crossing it, so partition is decided by the weakest cut, and an average has no term for where the weak links are. Whether a fleet is one network is a property of its topology and has to be computed from it.

1 WHY THIS QUESTION IS BEING ASKED NOW

The Secure World Foundation's 2026 assessment of counterspace capabilities, released in April, reported that jamming of GPS and of satellite communications is increasing both inside and outside armed conflicts, and that thirteen nations are now pursuing at least some capability to interfere with adversary satellites [1, 2]. Interference is no longer a condition that a fleet meets only in a war zone. It is a condition that a fleet planning for the next decade has to assume.

In the same month the Pentagon was moving in the other direction on autonomy. Defense One reported in May 2026 on new DARPA programmes for robots that form teams and carry out missions without a central node [3], and the Swarm Forge solicitation of March 2026 asked vendors for swarms that keep working while an adversary tries to jam, spoof, intercept or take over their control links, with decentralised control so that no single system is a point of failure [4].

Every one of those systems runs on a mesh of radio links between its units, and every one of them will be monitored, tested and accepted on link metrics. Every link metric in common use is an average. This paper is about what an average cannot say, and it says it with numbers rather than with a warning.

2 WHAT A DASHBOARD REPORTS AND WHAT A NETWORK NEEDS

A fleet's communications are usually described by a small set of quantities reported on a display or in a test report. The share of links that are up. The mean packet delivery ratio. The mean received signal strength. The mean number of neighbours each unit can hear. The mean latency. Each is an average over the fleet, or over a window of time, or both, and together they say whether the radios are working and whether the environment is getting worse.

What a mesh network needs is different. A mesh works while every unit can reach every other unit, directly or by relay through other units. A unit that cannot reach the rest is not part of the network, whatever its own radio reports. A fleet that has split into two connected groups is two networks. Whether the fleet is one network is a property of the whole arrangement of links, and two arrangements with the same average can differ in it.

We use the term threshold blindness for two properties that averaged link metrics have.

- The reading moves smoothly through a point at which the network changes abruptly: connectivity has a threshold, and the average passes through it at the same rate it passes through everything else.
- The reading takes the same value on networks that behave differently: the average of the link qualities does not depend on where the weak links are, and the network's connectivity does.

The rest of the paper shows both properties on a simulated formation, measures how much they cost in predictive power, and states what the number that does predict partition depends on.

THE SIMULATED FORMATION

All results come from a simulated flat formation of 400 units placed at random inside a square. Each unit is linked to every other unit inside a fixed radio range, set so that a unit in the interior of the formation has about 12 others in range; the average across the formation, including edge units, is about 11. Each link works with a probability set by the scenario, independently of the

other links. The dashboard reading is the share of in-range links that are working, which is what a display that polls each link and averages would show. A network is counted as whole if its largest connected group holds at least 90% of the units, and as partitioned otherwise. Section 7 lists what this leaves out.

3 A STEADY DECLINE IN THE READING, A STEP IN THE NETWORK

Figure 1 follows a formation through a mission in which the environment degrades at a constant rate. Every link's reliability falls from 100% at the start to zero at the end, so the dashboard reading falls in a straight line. The figure shows what the network does at the same time.




FIGURE 1 A formation of 400 units under link reliability that falls at a constant rate over a mission. The straight line is the dashboard reading. The other line is the share of units in the largest connected group. The shaded band is the interval during which the network goes from whole to split. Each curve averages 60 random formations.

For the first 71% of the mission, while the reading falls from 100% to 29%, the share of units in one group stays above 90%. Halfway through the mission the reading is 50% and 99% of the units are still in one group. Then the network splits. Between 71% and 81% of the mission, while the reading moves from 29% to 19%, the largest group falls from nine tenths of the fleet to half of it. Nearly three quarters of the reading's range corresponds to no change in connectivity, and about a tenth of it corresponds to all of the change.

The reading gives no indication of where in its range the step sits. It declines at the same rate before the step, through it and after it. An operator watching the reading fall from 60% to 50% sees the same thing as an operator watching it fall from 30% to 20%, and only the second is watching the network come apart. The threshold is a property of the formation, and a display that does not know the formation cannot mark it.

The step itself is not an artefact of the simulation. It is the giant-component transition of ran-

dom graph theory, first described for networks without geometry in 1960 [5] and established for networks in which units talk only to nearby units by the continuum percolation literature [6, 7]. Below the transition no large connected group exists; above it one group holds a fixed share of the units. For a fleet in which links fail at random, the location of the transition depends on the number of neighbours each unit has in range. A denser formation crosses it at a lower reading, and a sparser one at a higher reading. So a fleet that flies spread out and the same fleet that flies close have their thresholds at different points on the same dashboard scale, and nothing on the scale moves to show it.

4 SAME READING, DIFFERENT NETWORKS

The second property of averaged metrics is the more damaging one, because it holds even when nothing is degrading. Figure 2 shows four networks, each with a dashboard reading of about 50%, in which the failed links are arranged differently.



FIGURE 2 One realisation of each of four scenarios on the same formation, all with a dashboard reading near 50%. Units in the largest connected group are blue, units in a second large group are orange, and units in small fragments are grey. Working links are drawn in dark grey. From left to right, failures are spread uniformly, arranged in a gradient from right to left, concentrated in a footprint in the middle, and concentrated in a band down the middle.

- **Uniform:** every link works with probability 50%.
- **Gradient:** link reliability falls from high on one side of the formation to low on the other, as it would under an interference source, or a terrain mask, off one flank.
- **Footprint:** links inside a disc in the middle of the formation almost never work and links outside it work most of the time, as under an interference source with a footprint.
- **Band:** links inside a strip across the formation almost never work and links outside it work most of the time, as under a source whose footprint crosses the formation rather than sitting inside it.

In every case the reliabilities are chosen so that the share of working links is 50%. Table 1 gives the outcomes over 60 random formations of each.

ARRANGEMENT	READING	LARGEST GROUP	SECOND GROUP	PARTITIONED
Uniform	50%	99%	1%	2%
Gradient	50%	86%	2%	90%
Footprint	50%	73%	3%	100%
Band	50%	48%	39%	100%

TABLE 1 Four arrangements of failed links with the same dashboard reading, averaged over 60 random formations of 400 units. Partitioned means the largest group holds under 90% of the units.

The uniform network is whole in all but one of the sixty trials. The gradient cuts off the far flank and leaves 86% of the units in one group. The footprint cuts off the units inside it, which end up in small fragments rather than in a second group, and leaves 73%. The band splits the formation into two networks of 48% and 39%, and it does so in every trial. The dashboard reads 50% on all four.

The reason is a property of connected networks that has been known since 1927 [8]. Draw any line that separates the formation into two parts. The two parts are connected to each other only if at least one working link crosses the line. The network is connected only if that holds for every such line, so it is the line with the fewest working links across it that decides. In the band case, the links inside the band are 21% of all links and almost none of them work, while the rest work 62% of the time and pull the average up to 50%. The average is dominated by the links that do not matter. The cut is decided by the links that do.

Arrangements like the footprint and the band are not special cases. An interference source has a footprint. Terrain masks one part of a formation and not another. Range degrades links on the far side of a formation before the near side. The uniform case, in which every link fails with the same probability regardless of where it is, is the one arrangement in which the average carries the information, and it is the least likely arrangement to occur.

5 CAN THE AVERAGE PREDICT PARTITION AT ALL?

The two figures above are chosen cases. To measure how much predictive power the average has across cases that were not chosen, we generated 600 random scenarios, each a random formation with a random base level of link reliability and, in a random subset, a footprint or a band or both of random size, position and depth. For each scenario we recorded the dashboard reading and whether the network came out whole or partitioned. 31% of scenarios partitioned.




FIGURE 3 600 random scenarios. Each point is one network, placed by its dashboard reading and the share of its units in the largest group. Blue points are whole networks and orange points are partitioned ones. The shaded band is the range of readings over which both outcomes occur.

Figure 3 shows the result. Networks that partitioned reported readings as high as 73%. Networks that stayed whole reported readings as low as 24%. 79% of the scenarios had readings inside that overlap. Taking only the 91 scenarios whose reading was within five points of 50%, 23% were partitioned and the share of units in the largest group ranged from 38% to 100%.

We then asked how well three quantities predict the outcome, giving each the most favourable treatment we could.

- The dashboard reading: the best single threshold on the reading, chosen after the fact with full knowledge of the outcomes, calls the outcome wrong 15% of the time.
- The worst-connected tenth of the fleet: the number of working links a unit can expect, taken at the tenth percentile across units rather than at the mean, with its threshold chosen the same way, is wrong 8% of the time.
- A computation on the topology: running the same scenario ten times on the actual formation and link reliabilities and taking the majority outcome is wrong 6% of the time.

Predicting the size of the largest group rather than just whole or split separates the three further. A lookup table fitted to the dashboard reading is off by 8.0 points on average, one fitted to the worst tenth by 5.4 points, and the topology computation by 3.5 points.

Two things follow. The first is that the information needed to predict partition is present in the link qualities, but it sits in their distribution and their arrangement, not in their mean. A reading that keeps track of the worst-connected part of the fleet recovers more than half of what the average loses. The second is that no predictor reaches zero, because near the threshold the same formation with the same link reliabilities comes out whole on one draw and split on the next. The topology computation is at that floor. Its remaining error is chance, not blindness.

Measures that depend on arrangement rather than on the mean have been in the literature for a long time. The algebraic connectivity of a network, which is small when a network has a narrow cut and large when it does not, was introduced in 1973 [9]. What is new in the last two years is not the mathematics of the threshold but the number of fleets that will be operated against it.

6 WHAT THIS MEANS FOR A PROGRAMME OFFICE OR AN OPERATOR

A communications requirement stated as an average, or an acceptance test that reports one, is not a statement about whether the fleet will be one network in the environment it is built for. The following questions have computable answers for any proposed fleet.

- 1 Where does the threshold sit, in the units the dashboard reports, for this fleet at this spacing and this formation shape? The threshold moves with the neighbour count and the dashboard scale does not.
- 2 How much of the reading's range is above the threshold, and how far above it does the fleet sit in the interference environment it is expected to meet?
- 3 Which cut across the formation is the weakest, and what does the dashboard read when that cut has no working links across it?
- 4 What does the worst-connected tenth of the fleet look like at the moment the average still looks fine, and should the display show it?
- 5 At the rate the environment is expected to degrade, how much warning in time does the reading give before the fleet splits?

None of these is answered by an average, a claimed radio range, or a flight test in a clean environment. Each is a property of the fleet's topology together with the environment, and each can be computed before a flight test from published radio specifications, the formation and spacing, and a mission profile.

We compute them for specific fleets as a written assessment, the Connectivity Fragility Assessment. The method is the one in this paper, run on the actual formation geometry and radio characteristics rather than the random square used here. The assessment reports which links fail first as conditions degrade, when and where the fleet splits, how long each split lasts, and which units end up isolated.

7 LIMITS OF THIS ANALYSIS

The simulation is deliberately simple, and each simplification leans in a known direction.

- Link states are independent given the scenario: the footprint and band scenarios put correlation into the reliabilities, but within a scenario each link's state is drawn on its own. Correlated fading, in which nearby links go up and down together, makes the outcome more

variable at the same reading and widens the overlap in Figure 3. Correlated interference is the subject of a separate paper in preparation.

- **Snapshot in time:** units do not move and links do not recover. A moving fleet crosses in and out of footprints, and a unit cut off now may reconnect in seconds. This changes how long a split lasts and does not change where the threshold sits.
- **Links are either up or down:** a link with low but nonzero throughput counts as up. Modelling rate and latency lowers the effective reading at a given level of interference and moves the step earlier in Figure 1.
- **Random placement:** a structured formation with the same spacing has its threshold in nearly the same place and the same weakest-cut property.
- **No unit losses:** units are not removed. Losses thin the formation and lower the threshold in the reading; a later paper in this series treats them.
- **The predictors were fitted on the data they were scored on:** this is generous to the dashboard reading and to the worst-tenth measure, whose thresholds were chosen with the outcomes in hand. A rule fixed in advance would do worse.

The mathematics of what happens once the threshold is crossed, how a failure front moves through a connected group and how fast, is set out in [10, 11]. Those papers are mathematical and do not depend on the simplifications listed here.

8 METHOD NOTE

Units are placed uniformly at random in a unit square, and every pair within a radio range chosen so that an interior unit has 12 others in range is linked. For Figure 1, every link is kept with probability p , and p falls linearly from 1 to 0 across 101 steps; at each step the connected groups are found with a standard union-find pass, and the largest group's size is divided by the number of units. The dashboard reading at each step is the share of in-range links kept. The crossing points are read by linear interpolation. For Figure 2 and Table 1, each link's reliability is set from the position of its midpoint. The gradient sets it to 50% plus the midpoint's distance from the centre line, scaled so that it runs from near zero on one side to near one on the other. The footprint sets it to 3% inside a disc of radius 0.32 centred on the formation and to a single value outside chosen so that the mean over links is 50%. The band does the same with a strip of width 0.2 down the middle. Table 1 averages 60 formations of each. For Figure 3, each of 600 scenarios draws a base reliability uniformly between 20% and 90%, adds a disc with probability one half, with centre drawn uniformly in the middle 60% of the square, radius between 0.12 and 0.35 and reliability multiplier between 0 and 0.6, and adds a band with probability 0.4, with random orientation, centre drawn uniformly in the middle 40% of the square, width between 0.04 and 0.2 and multiplier between 0 and 0.6. The worst-tenth measure is the tenth percentile across units of each unit's expected number of working links. The topology predictor draws ten further realisations of the same scenario and takes the majority outcome, or the mean largest-group share for the size prediction. The best-threshold error rates search every threshold on the predictor and report the lowest error. The lookup-table errors fit 20 quantile bins on the predictor and report the mean absolute error of the bin means. Every number in the text is written from the simulation output by a script rather than transcribed by hand. The code is a few hundred lines of Python using numpy and scipy, runs from a fixed seed, and is available on request.

ABOUT THIS SERIES

This paper is part of a series from Dreiton Dynamics on threshold behaviour in connected systems. Each paper takes one question that programme offices, operators or planners are already asking, states the mathematics that governs it in plain terms, and shows what the numbers look like on a simulated example. The series does not describe any Dreiton product. WP-01 covers why averaged link metrics cannot predict when a network partitions. The papers can be read independently.

ABOUT DREITON DYNAMICS

Dreiton Dynamics is an independent research group in Ashburn, Virginia. Its work is the mathematics of threshold cascades on networks and the application of that mathematics to radio meshes, electric distribution and other connected systems. It offers two fixed-fee written assessments, the Connectivity Fragility Assessment for fleets of drones, robots and vessels, and the Cascading Outage Exposure Assessment for data center operators. Both are run on published specifications and public data unless the client chooses to share more.

AchyutKumar@dreitondynamics.com | abhinavduddala@dreitondynamics.com | dreitondynamics.com

REFERENCES

- [1] T. Hitchens, "As more nations seek counterspace chops, GPS jamming also rises: report," *Breaking Defense*, April 8, 2026.
- [2] Secure World Foundation, *Global Counterspace Capabilities: An Open Source Assessment*, 2026 edition, April 2026.
- [3] "Pentagon seeks smarter, self-organizing drones as autonomous-warfare budget is poised to skyrocket," *Defense One*, May 4, 2026.
- [4] J. Harper, "Pentagon preparing for drone swarm 'crucible'," *DefenseScoop*, March 31, 2026.
- [5] P. Erdős and A. Rényi, "On the evolution of random graphs," *Publ. Math. Inst. Hung. Acad. Sci.* 5, 17-61, 1960.
- [6] M. Penrose, *Random Geometric Graphs*, Oxford University Press, 2003.
- [7] P. Gupta and P. R. Kumar, "Critical power for asymptotic connectivity in wireless networks," in *Stochastic Analysis, Control, Optimization and Applications*, Birkhäuser, 1998.
- [8] K. Menger, "Zur allgemeinen Kurventheorie," *Fundamenta Mathematicae* 10, 96-115, 1927.
- [9] M. Fiedler, "Algebraic connectivity of graphs," *Czechoslovak Mathematical Journal* 23, 298-305, 1973.
- [10] A. Kumar, "Threshold cascades of interacting diffusions on unimodular random trees," arXiv:2608.15987, 2026.
- [11] A. Kumar and A. Duddala, "Two problems for threshold cascades of interacting diffusions on unimodular random trees: front propagation with a Bramson correction, and the continuous-type limit theory," arXiv:2608.21125, 2026.