A designed tree model with a sharp congestion–fragmentation transition

Alec Kriebel
Independent researcher
ORCID: 0009-0001-9320-500X

21 October 2026 (version 1.0; unrefereed preprint)

Abstract

We construct a bounded-degree physical tree model in which the normalized marginal admitted demand drops from one to zero while the spare-capacity graph changes from the entire tree to isolated vertices. Capacities arise from the unique cheapest fixed-topology enlargement supporting doubled initial traffic. Each of the k^2 cross-root streams has an independent, fixed-law exponential growth rate. An all-cuts estimate confines both changes, with probability tending to one, to an interval of width $16\sqrt{\log k/k}$ around mean doubled demand. Both outer-phase conclusions hold for every optimizer. This is a specialized existence response to Aldous's conceptual congestion toy-model request. The width is an upper bound; no critical profile or universal traffic law is asserted.

1 The request and the model

Aldous's problem page [1] asks for a tractable model connecting a rapid decline in the marginal amount of extra traffic admitted with fragmentation of the graph of edges having spare capacity. One suggested design mechanism minimizes the cost of increasing capacities to support doubled current demand. The accompanying index classifies the request as conceptual, rather than as one uniquely formulated conjecture. We implement that mechanism in an explicit family. Its topology, traffic pattern and growth process are part of the model.

Let T_H be the complete rooted binary tree of height $H \geq 2$, with $k = 2^{H-1}$ leaves in each root subtree and $N = 4k - 1$ vertices. Write L, R for its left and right leaf sets. There is one commodity from each $i \in L$ to each $j \in R$, initially of size one; all other demands vanish. Each commodity uses its unique simple physical path. For an edge e, let k_e count its descendant leaves. Its initial load and capacity are kk_e. Give every capacity increment a strictly positive linear price and require all doubled demands to fit. The unique cheapest resulting capacities are

$$C_e = 2kk_e. \tag{1}$$

Indeed, the paths force these lower bounds separately, and the bounds are simultaneously feasible. Positive prices force the componentwise minimum. Terminal capacities grow as $2k$; the two root-edge capacities are $2k^2$. Each finite network is then held fixed.

Sample independent $X_{ij} \sim \mathrm{Exp}(1)$, and for $t \geq 1$ put

$$d_{ij}(t) = 1 + (t-1)X_{ij}, \qquad S_k = \sum_{i,j} X_{ij} > 0 \quad \text{a.s.}$$

These are quenched random slopes, not fresh samples at successive times. Each stream has mean t and variance $(t-1)^2$, independently of k. Thus this is a specialization different from the linear-variance example on the source page.

Let $F_k(t)$ be the maximum of $\sum_{ij} y_{ij}$, with $0 \leq y_{ij} \leq d_{ij}(t)$, subject to physical capacities. Define

$$r_k(t) = F'_{k,+}(t)/S_k.$$

The denominator is the realized offered-demand slope, not initial demand k^2. For each optimizing admitted matrix let $E_k(t)$ contain exactly the edges with loads strictly below capacity, on the full vertex set of T_H; write $M_k(t)$ for its largest component size. Inside the transition this graph may depend on the optimizer. The two phase conclusions below quantify over every optimizer. Component-size convergence applies to any measurable selection of an optimizer; the outer-phase event itself is uniform over all optimizers.

2 The joint transition

Theorem 1. *Put*

$$\delta_k = 8\sqrt{\frac{\log k}{k}}, \qquad \epsilon_k = 2k^{-31} + \frac{2k^{-2}}{(1-k^{-2})^2} + \left(\frac{4}{k^2}\right)^k.$$

For all binary k large enough that $\delta_k < 1$, with probability at least $1-\epsilon_k$, both statements hold simultaneously:

(i) For every $1 \le t \le 2-\delta_k$, all demand is admitted, $r_k(t) = 1$, and every optimizer has $E_k(t) = T_H$, hence $M_k(t) = N$.

(ii) For every $t \ge 2+\delta_k$, $F_k(t) = 2k^2$, $r_k(t) = 0$, and every optimizer has edgeless $E_k(t)$, hence $M_k(t) = 1$.

Consequently, for fixed $1 \le t < 2$, $r_k(t) \to 1$ and $M_k(t)/N \to 1$ in probability; for fixed $t > 2$ both tend to zero. For every fixed $t \ge 1$, including $t = 2$,

$$F_k(t)/k^2 \longrightarrow \min(t, 2) \quad \textit{in probability.}$$

No limit of $r_k(2)$, critical component-size profile, or optimal transition-width order is asserted.

2.1 Exact deterministic reduction

A left leaf's load is its row sum $\sum_j y_{ij}$; a right leaf's load is its column sum. Each other load is the sum of the corresponding descendant terminal loads. Thus (1) is equivalent exactly to

$$\sum_j y_{ij} \le 2k, \qquad \sum_i y_{ij} \le 2k, \qquad 0 \le y_{ij} \le d_{ij}(t). \tag{2}$$

Internal constraints introduce no further restriction.

The source–L–R–sink network has capacities $2k$, $d_{ij}(t)$, $2k$ on its three arc layers. Max-flow/min-cut [2] gives the finite formula

$$F_k(t) = \min_{\substack{A \subseteq L \\ B \subseteq R}} \left[2k(k - |A| + |B|) + |A|(k - |B|) + (t-1) \sum_{i \in A, j \notin B} X_{ij} \right]. \tag{3}$$

It is a finite minimum of nondecreasing affine functions, hence is continuous, concave and piecewise affine. Its right derivative is the smallest active slope. With $\mathcal{I}_k(t)$ the active cuts, including all ties,

$$r_k(t) = \min_{(A,B) \in \mathcal{I}_k(t)} \frac{\sum_{i \in A, j \notin B} X_{ij}}{S_k} \in [0, 1]. \tag{4}$$

These formulas do not assert efficient exhaustive enumeration or an elementary finite-k expectation.

If all offered row and column sums are strictly less than $2k$, full admission is feasible and is the unique admitted matrix attaining the offered total: every entry is individually upper bounded. Every physical edge then has strict slack. Conversely, $F_k(t) \le 2k^2$ always. Equality forces every row and column of every optimizer to equal $2k$, saturating every physical edge. These observations establish the optimizer quantifiers.

2.2 Exponential tails and the lower phase

For a sum Z of n independent Exp(1) variables, $\mathbb{E}e^{\lambda Z} = (1-\lambda)^{-n}$ for $\lambda < 1$. Exponential Markov with $\lambda = 1 - 1/a$ gives

$$\mathbb{P}(Z \geq na) \leq e^{-nI(a)} \ (a > 1), \qquad \mathbb{P}(Z \leq na) \leq e^{-nI(a)} \ (0 < a < 1), \quad I(a) = a - 1 - \log a.$$

For $0 < \delta < 1$,

$$I((1-\delta)^{-1}) \geq \delta^2/2, \qquad I((1+\delta)^{-1}) \geq \delta^2/8. \tag{5}$$

Both sides vanish at zero. The left-hand derivatives in δ are $\delta/(1-\delta)^2 \geq \delta$ and $\delta/(1+\delta)^2 \geq \delta/4$, proving (5) by integration.

At $t = 2 - \delta$, an offered row sum is $k + (1-\delta)\sum_j X_{ij}$. It reaches $2k$ only if that exponential sum reaches $k/(1-\delta)$; the same holds for a column. A union bound over $2k$ events, without joint independence, gives

$$\mathbb{P}(\text{some offered row or column reaches } 2k) \leq 2ke^{-k\delta^2/2}.$$

At $\delta = \delta_k$ this is $2k^{-31}$. Otherwise all rows and columns have strict slack at the endpoint and at all earlier $t \geq 1$. The endpoint slack also persists in a right neighborhood. The deterministic observations prove (i), including its endpoint right derivative.

2.3 All upper-phase cuts

At $t = 2 + \delta$, all cuts in (3) are at least $2k^2$ exactly when

$$\sum_{i \in A, j \notin B} d_{ij}(t) \geq 2k(|A| - |B|).$$

Cuts with $|A| \leq |B|$ are automatic. For the others put $D = R \setminus B$, $a = |A|$, $d = |D|$, so $a + d > k$. Since

$$ad - k(a + d - k) = (k - a)(k - d) \geq 0,$$

failure implies

$$ad + (1+\delta)\sum_{A \times D} X_{ij} < 2k(a + d - k) \leq 2ad, \qquad \sum_{A \times D} X_{ij} < ad/(1+\delta).$$

Its probability is at most e^{-cad}, where $c = I((1+\delta)^{-1}) \geq \delta^2/8$. Different cuts need not be independent.

Put $q = \min(a, d)$, $\ell = k - \max(a, d)$. Then $0 \leq \ell < q$ and $ad = q(k - \ell)$. If $q \leq \lfloor k/2 \rfloor$, the number of choices for fixed q, ℓ, allowing interchange, is at most $2\binom{k}{q}\binom{k}{\ell}$. Summing over $\ell < q$ bounds it by $2qk^{2q}$, whereas $ad \geq qk/2$. For $q > k/2$, $ad > k^2/4$ and there are at most 4^k choices in total. The probability that any cut fails is therefore at most

$$\sum_{q=1}^{\lfloor k/2 \rfloor} 2qk^{2q}e^{-cqk/2} + 4^k e^{-ck^2/4}.$$

At δ_k, $c \geq 8\log(k)/k$, bounding this by

$$\sum_{q \geq 1} 2qk^{-2q} + \left(\frac{4}{k^2}\right)^k = \frac{2k^{-2}}{(1-k^{-2})^2} + \left(\frac{4}{k^2}\right)^k.$$

When all cuts pass, $F_k(2 + \delta_k) = 2k^2$. Increasing demands retain a maximizing admitted flow for all later times. The upper bound $2k^2$ then makes throughput constant, its right derivative zero, and every edge saturated for every optimizer. Combining the two failure bounds proves (i)–(ii).

2.4 Limits

We have $\mathbb{E}(S_k/k^2) = 1$ and $\mathrm{Var}(S_k/k^2) = k^{-2}$. The fixed-$t$ conclusions away from two follow from (i)–(ii), $\epsilon_k \to 0$, and $\delta_k \to 0$. At two, on the same good event,

$$k^2 + (1 - \delta_k)S_k = F_k(2 - \delta_k) \le F_k(2) \le 2k^2.$$

This gives the remaining throughput limit without exchanging limits and derivatives. The failure bounds are summable along binary k; hence the outer-phase events also hold eventually almost surely under any common coupling, by the first Borel–Cantelli lemma.

3 Context, scope and reproducibility

Capacitated random transportation and flow-cut analysis are classical. Hassin–Zemel [3] and Karp–Motwani–Nisan [4, 5] are essential antecedents, alongside max-flow/min-cut and elementary concentration. The highly optimized tolerance paradigm is prior work [6], as are congestion-induced fragmentation and shortest-path percolation [7]. The latter sequentially consumes edges for individual requests, rather than optimizing simultaneous admitted demand under fixed capacities. The contribution here is the explicit cheapest-design tree construction and its simultaneous normalized-marginal and every-optimizer spare-connectivity theorem, not a new general transport-feasibility principle.

The sufficient exponent separation in [3, Theorem 1] does not cover the present balanced near-capacity regime. The directed theorem in the actual 1988 report [5, Section 4] fixes bounded integer capacities and a positive constant mean surplus; Section 5 separately proposes a vanishing-surplus extension. That report has not been equated with an inspected 1993 final edition: the final journal body [4] was unavailable. These comparisons do not exclude every possible derivation from prior work.

The physical graph is sparse but traffic is dense, capacities increase with size, and internal constraints reduce to terminal constraints. Only the fixed-topology design alternative is implemented. No adaptive capacity growth, generic cyclic-network theorem, linear-variance process, critical profile or universal critical law is supplied. Throughput saturates; the marginal drops. The width is a sufficient bound without a matching lower bound. Aldous's 2005 discussion [8, slide 5] already contrasts easy but unrealistic tree analysis with preferred mean-field models. Our tree has no rerouting and does not remove that explanatory limitation. The suitability of one example as the uniquely *right* conceptual model is not a mathematical conclusion.

A dated bounded audit inspected 16 primary source bodies and recorded 52 literature searches, locating no earlier exact construction and theorem in that corpus. The final-edition gap above, unpublished work and incomplete indexing remain limitations; no first-discovery or global novelty certificate is asserted. The source and theorem comparisons are in the supplement.

The proof is analytic; finite controls do not establish rare-event asymptotics. The original 24-file two-turn packet is preserved unchanged. The supplement's new replay derivative verifies those frozen files and requires exact active second-turn output bytes. It admits bounded absolute drift only at 40 specified older floating diagnostics, together with explicit runtime and public-input projection rules; all other structure and controls remain exact. Python 3.9.6, NumPy 2.0.2 and SciPy 1.13.1 were tested. The strict historical wrapper fails at roundoff in that environment and is not advertised as portable. Exact policies, guard tests and limitations are provided; other platform/library combinations are not guaranteed.

AI involvement and review status. AI tools were extensively involved in mathematical research, model and proof development, computational verification, literature search, manuscript preparation and adversarial auditing. This is an unrefereed preprint. Automated independent-agent reviews are not independent human peer review; the manuscript has not undergone formal independent human peer review.

References

[1] D. Aldous, *Congestion in networks: SOC, HOT and the percolation-fragmentation phase transition*, problem page posted May 2003, update marked February 2018, accessed 21 October 2026. Primary problem; conceptual problem index.

[2] L. R. Ford, Jr. and D. R. Fulkerson, Maximal flow through a network, *Canadian Journal of Mathematics* **8** (1956), 399–404. doi:10.4153/CJM-1956-045-5.

[3] R. Hassin and E. Zemel, Probabilistic analysis of the capacitated transportation problem, *Mathematics of Operations Research* **13**(1) (1988), 80–89. doi:10.1287/moor.13.1.80.

[4] R. M. Karp, R. Motwani and N. Nisan, Probabilistic analysis of network flow algorithms, *Mathematics of Operations Research* **18**(1) (1993), 71–97. doi:10.1287/moor.18.1.71.

[5] R. M. Karp, R. Motwani and N. Nisan, *Probabilistic analysis of network flow algorithms*, Berkeley technical report CSD-88-392, December 1988. Actual report.

[6] J. M. Carlson and J. Doyle, Complexity and robustness, *Proceedings of the National Academy of Sciences* **99**(Suppl. 1) (2002), 2538–2545. doi:10.1073/pnas.012582499.

[7] M. Kim and F. Radicchi, Shortest-path percolation on random networks, *Physical Review Letters* **133** (2024), 047402. doi:10.1103/PhysRevLett.133.047402.

[8] D. Aldous, *Flows through random networks*, Ottawa talk, June 2005. Primary slides.