Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Cost convexity controls the architecture of resistance-optimal networks

Daniel Pérez-Calixto

Daniel_Perez@ciencias.unam.mx

-instituto Nacional de }Medicina Genómica https://orcid.org/0009-0004-3988-1260

Research Article

Keywords: Optimal transport networks, Network topology, Constructal theory, Complex networks

Posted Date: September 29th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11174611/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Cost convexity controls the architecture of resistance-optimal networks

Hugo Tovar,[1,2] Genaro Vázquez-Victorio,[2] Enrique Hernández-Lemus,[1]
Alejandra Cervera,[1,*] and Daniel Pérez-Calixto[1,2,†]

[1] *Instituto Nacional de Medicina Genómica, Mexico City, Mexico*
[2] *Departamento de Física, Facultad de Ciencias, Universidad Nacional Autónoma de México, Mexico City, Mexico*

A conducting network improves access to its interior only by spending material, so architectures compare only at equal budget. We redistribute conductance under a power-law resource law $C_\alpha = \sum_e a_e w_e^\alpha$ and measure access by the reciprocal of the mean pairwise effective resistance. When the cost is strictly convex the allocation is unique and every admissible edge keeps positive conductance; at $\alpha = 1$ uniqueness survives but an edge can be priced out. The marginal worth of an edge is then its squared biharmonic distance, a Murray-type balance whose balanced quantity is global rather than junction-local. The same optimum scores any weighted network at its own budget, separating an allocation factor from a topological one. Under concave cost the marginal price of an incipient channel diverges, the feasible set is nonconvex, and conductance condenses onto sparse backbones. The uniform star and the uniform complete graph cross at $\alpha_c(n) = \log(n/2)/\log(n-1)$, near which intermediate states beat both. Up to $n = 64$ the best states found grow approximately linearly in n, so channels per node change much less than the active fraction does. Loops are governed instead by a collapse exponent $\alpha^*(n)$, resolved to 0.005 at every size: it rises from 0.555 at $n = 8$ to 0.675 at $n = 64$, and the collapse across it is abrupt, the support at $n = 64$ falling from 133 active edges to the 63 of a spanning tree. A fixed exponent therefore changes character with size by crossing this boundary. One prescribed source and sink instead gives a loopless optimum.

I. INTRODUCTION

Transport capacity and material cost are inseparable in a conducting network. Raising an edge conductance improves access to the interior, and it consumes resources that are then unavailable elsewhere. Any statement that one architecture transports better than another is therefore incomplete unless the two are held at the same budget.

Effective resistance is the natural measure of access when edge weights are conductances. It accounts for parallel routes instead of collapsing transport onto a single shortest path [1–3], and its sum over node pairs is the Kirchhoff index. Minimizing that sum under a linear budget is a classical convex design problem with a well-understood solution [4]. This leaves two questions. How does the optimal architecture change when the physical price of conductance is nonlinear? And how much of the available access does an arbitrary weighted network actually use?

Power-law prices arise whenever conductance is set by a channel geometry that itself costs volume, surface, or mass, as in vascular, porous, thermal, and electrical structures. The exponent is not a fitting convenience. Convex cost penalizes concentration and rewards spreading conductance thinly; concave cost does the reverse, and can drive the optimum onto a sparse support. Optimal transport networks have been shown to pass between loopless and redundant architectures as the cost law or the loading is changed [5–11]. Those studies prescribe sources and sinks, or minimize dissipated power. The loading here is different: every unordered node pair is weighted equally, so the objective is the mean effective resistance and the network is optimized for transport between any two sites rather than between selected terminals. The distinction is inert in the convex regime and decisive in the concave one, where the all-pairs objective can retain loops while the exact single-pair optimum cannot.

The two cost regimes also require different standards of evidence. In the convex case, the unique reference state makes the normalized efficiency well defined, and the gradient shows where additional conductance provides the largest gain in access. Below $\alpha = 1$, several supports may compete. Multiple local minima arise from the nonconvex formulation rather than from numerical noise, so the reference can only be reported as the best state found within a stated search. Over the sizes examined here, these states occupy a decreasing fraction of the candidate graph even though their absolute support continues to grow. We interpret this behavior through the constructal law, comparing each network with the best configuration allowed by the same resource model [12, 13].

II. ACCESS AT FIXED RESOURCE

A. Conductance network and global access

Consider a connected undirected graph $G = (V, E)$ with node set V, admissible edge set E, $n = |V|$, and $m = |E|$. The non-negative edge conductances are collected in $W = (w_e)_{e\in E}$. After fixing an arbitrary orientation for each edge, $b_e \in \mathbb{R}^n$ is its incidence vector, carrying $+1$

* acerverat@inmegen.gob.mx
† daniel_perez@ciencias.unam.mx

and -1 at the endpoints. The weighted Laplacian is

$$L(W) = \sum_{e \in E} w_e \, b_e b_e^{\mathsf{T}}. \tag{1}$$

For nodes s and t with standard basis vectors e_s and e_t, write $b_{st} = e_s - e_t$. The effective resistance between them is $r_{st} = b_{st}^{\mathsf{T}} L^{+} b_{st}$, with L^{+} the Moore–Penrose pseudoinverse of L. Averaging over the $n(n-1)/2$ unordered pairs gives

$$\overline{R}(W) = \frac{2}{n-1} \operatorname{tr}(L^{+}). \tag{2}$$

The prefactor $2/(n-1)$ converts the Kirchhoff index $n \operatorname{tr}(L^{+})$ into a pairwise mean, and it follows from $\sum_{i \neq j} r_{ij} = 2n \operatorname{tr}(L^{+})$ together with the fact that the rows of L^{+} sum to zero. Equation (2) represents a uniform ensemble of unit current injections and withdrawals. We take global access to be $A(W) = 1/\overline{R}(W)$, the reciprocal of the mean resistance and not the mean of the pair conductances.

B. Resource law

The price of conductance is

$$C_\alpha(W) = \sum_e a_e w_e^{\alpha}, \qquad a_e > 0, \quad \alpha > 0. \tag{3}$$

The coefficient a_e separates the candidate topology from the local price of an edge, and α fixes how that price grows with conductance. In the spatial examples below we use the length penalty $a_e = \ell_e$ with ℓ_e the Euclidean channel length.

Both α and the length dependence of a_e must be inherited from a physical model rather than chosen for convenience. If w_e is the end-to-end Hagen–Poiseuille conductance of a tube, then $w_e \propto r_e^4/\ell_e$ and the tube volume scales as $\ell_e^{3/2} w_e^{1/2}$, so a fixed-volume law corresponds to $\alpha = 1/2$ and $a_e \propto \ell_e^{3/2}$. Table I collects the pairs (α, a_e) that follow from the standard mechanisms in the same way. All of them place the physically motivated exponent at or below the convex boundary, which is why the concave regime is the relevant one rather than a mathematical curiosity. The simpler choice $a_e = \ell_e$ used in most of the spatial examples below is a controlled spatial penalty and is not offered as a universal tube-volume law; the fixed-volume case itself is computed in Sec. IX.

TABLE I. Cost exponents implied by standard channel mechanisms, obtained by expressing the conserved resource in terms of the end-to-end conductance w_e of a channel of length ℓ_e.

Channel	Conductance	Resource	$(\alpha,\, a_e)$
Viscous tube	$w \propto r^4/\ell$	volume $\propto r^2\ell$	$(1/2,\ \ell^{3/2})$
Viscous tube	$w \propto r^4/\ell$	wall area $\propto r\ell$	$(1/4,\ \ell^{5/4})$
Ohmic wire	$w \propto A/\ell$	volume $\propto A\ell$	$(1,\ \ell^{2})$

At budget $C_0 > 0$ the reference allocation solves

$$\min_{W} \overline{R}(W) \quad \text{subject to} \quad C_\alpha(W) \le C_0, \qquad W \ge 0, \tag{4}$$

with $W \ge 0$ understood componentwise. A weighting whose positive-conductance support is disconnected is assigned $\overline{R} = +\infty$. Because $\overline{R}$ strictly decreases when any edge conductance is raised, every finite optimum exhausts the budget, $C_\alpha(W) = C_0$.

A minimizer exists for every $\alpha > 0$, including the concave range where the feasible set is not convex. The constraint $C_\alpha(W) \le C_0$ with $a_e > 0$ bounds each coordinate by $w_e \le (C_0/a_e)^{1/\alpha}$, so the feasible set is compact; $\overline{R}$ is continuous on the subset of weightings whose support is connected and diverges as connectivity is lost, so the infimum is attained in the interior of that subset whenever the candidate graph is connected. Uniqueness is a separate question and is guaranteed only for $\alpha \ge 1$.

C. Scoring a network against its own optimum

The optimum of Eq. (4) is useful beyond its own value, because it converts a resistance into an efficiency:

$$\eta_\alpha(W \mid \Omega) = \frac{\overline{R}^{*}_{\Omega}(C_\alpha(W))}{\overline{R}(W)}. \tag{5}$$

Here Ω is the admissible edge set and $\overline{R}^{*}_{\Omega}(C)$ is the minimum of $\overline{R}$ over Ω at budget C. The set may contain only the edges of an observed network, or all unordered node pairs. For $\alpha \ge 1$ and a fixed candidate edge set, Eq. (4) is convex and $\overline{R}$ is strictly convex on connected weightings, so the optimal allocation is unique and $0 < \eta_\alpha(W \mid \Omega) \le 1$, with equality only at that allocation. We refer to this ratio as the constructal efficiency index.

For $0 < \alpha < 1$, the global reference generally cannot be obtained from an uncertified numerical minimum. We therefore retain η_α only where the global optimum is known. Elsewhere we report the resistance of the best state found, its support, and its dependence on initialization; using a local minimum as the denominator would give an apparently absolute score without the required guarantee.

III. THE MARGINAL WORTH OF AN EDGE

The effect of a single edge can be obtained without solving the optimization problem at all. Let $\mathbf{1}$ be the all-ones vector. As long as the conducting support stays connected, the nullspace of L is span$\{\mathbf{1}\}$ and $\delta L^{+} = -L^{+}(\delta L)L^{+}$ on the orthogonal subspace $\mathbf{1}^{\perp}$. Perturbing w_e then gives

$$\frac{\partial \overline{R}}{\partial w_e} = -\frac{2}{n-1} \, b_e^{\mathsf{T}} (L^{+})^2 b_e = -\frac{2}{n-1} \, d_B^2(e). \tag{6}$$

Equation (6) itself is not new: the derivative $\partial \operatorname{tr} L^{+}/\partial w_e = -\|L^{+}b_e\|_2^2$ is standard in the convex resistance-design literature [4], and its identification with a biharmonic distance follows Refs. [14, 15]. What is new here is Eq. (7) below, its division by the α-dependent marginal price, and the boundary criterion at $w_e = 0$ that this division produces.

The squared biharmonic distance is $d_B^2(e) = b_e^{\mathsf{T}}(L^{+})^2 b_e = \|L^{+}b_e\|_2^2$ [14, 15], the squared norm of the node-potential response to a unit current driven across e. It is not the square of the effective resistance of that edge. The two correlate strongly and differ at order one: on `gnp_random_graph(14, 0.35, seed=4)`, which has 36 edges, with $w_e \sim U(0.3, 1.0)$, they correlate at 0.990, yet substituting r_e^2 for $d_B^2(e)$ misstates the gradient by up to 89% and reorders the ranking of interventions. The substitution is literal, $-[2/(n-1)]\, r_e^2$ in place of $-[2/(n-1)]\, d_B^2(e)$: the two have different units, and a proxy rescaled to a common mean would give a smaller discrepancy.

Because $A = 1/\overline{R}$, the marginal access gain carries an additional factor $\overline{R}^{-2}$. Dividing by the marginal price of the same edge defines the gain per unit cost,

$$\widetilde{\kappa}_e \equiv \frac{\partial A/\partial w_e}{\partial C_\alpha/\partial w_e} = \frac{2\, d_B^2(e)}{(n-1)\,\overline{R}^{2}\,\alpha\, a_e\, w_e^{\alpha-1}}. \tag{7}$$

At a stationary point that exhausts the budget, the Karush–Kuhn–Tucker conditions equalize $\widetilde{\kappa}_e$ across all active edges. This network-level balance is a Murray-type statement, in that it equates a marginal transport benefit with a marginal material cost [16, 17], but the analogy becomes physical only once the dependence of w_e on channel geometry and dissipation is specified.

Durand reached a related result along a different route [8]: minimizing global flow resistance at fixed channel volume or surface constrains the cross sections of channels meeting at a junction, again generalizing Murray's law. Here the constraint acts on conductances rather than directly on channel geometry, so α belongs to the assumed material law and is not fixed by the flow regime. The marginal quantity in Eq. (7) is also different: it is a biharmonic distance that depends on the whole network, not only on one junction. A local balance is recovered only when the two descriptions coincide.

The boundary $w_e = 0$ is what separates the two regimes, and it does so through Eq. (7) rather than through any property of $\overline{R}$. For $\alpha > 1$ the marginal price $\alpha a_e w_e^{\alpha-1}$ vanishes at $w_e = 0$, so an admissible edge with finite d_B cannot stay absent at the optimum: its marginal opening cost vanishes while its transport benefit remains positive. For $0 < \alpha < 1$ the same marginal price diverges as $w_e \to 0$, which suppresses the opening of a new channel and makes sparse supports locally stable. At $\alpha = 1$ the marginal price is the constant a_e, neither vanishing nor divergent, and the split between zero and positive conductances is decided by the complementarity conditions of the convex problem. An edge stays closed whenever $d_B^2(e)/a_e$ falls below the common value that the active edges share, so the linear case admits sparse optima even though it is convex and its solution unique.[1]

IV. THE CONVEX OPTIMUM AS AN ABSOLUTE REFERENCE

When every unordered node pair is admissible and $a_e = 1$, the candidate graph is complete. Permutation symmetry together with strict convexity gives the uniform conductance

$$w_0 = \left[\frac{2C_0}{n(n-1)}\right]^{1/\alpha} \tag{8}$$

as the unique optimum for $\alpha \geq 1$, and substitution into Eq. (2) gives

$$\overline{R}^{*}_{\text{free}} = \frac{2}{n}\left[\frac{n(n-1)}{2C_0}\right]^{1/\alpha}. \tag{9}$$

A fixed candidate topology supplies a second reference. Writing $\overline{R}^{*}_{\text{fixed}}$ for its minimum resistance at the same budget, the efficiency relative to the free reference separates into two factors,

$$\eta_{\text{free}} = \frac{\overline{R}^{*}_{\text{free}}}{\overline{R}^{*}_{\text{fixed}}} \cdot \frac{\overline{R}^{*}_{\text{fixed}}}{\overline{R}} = \eta_{\text{top}}\, \eta_w, \tag{10}$$

which defines $\eta_{\text{top}} = \overline{R}^{*}_{\text{free}}/\overline{R}^{*}_{\text{fixed}}$ and $\eta_w = \overline{R}^{*}_{\text{fixed}}/\overline{R}$. The topological factor measures how much access the wiring pattern forfeits; the allocation factor measures how well the available conductance is distributed on that pattern. The factorization is exact because both numerator and denominator are evaluated at the same budget, and it is what makes the two questions separable at all.

Table II evaluates both factors for a set of standard graphs at common n, α, and budget, and the two columns move independently. The star, the ring, and the complete graph are edge-transitive, so their automorphism group acts transitively on the variables being optimized and convexity forces the optimum to be uniform, giving $\eta_w = 1$. Edge-transitivity, not vertex-transitivity, is the property this argument needs: the star has a hub and is not vertex-transitive. Their η_{top} nonetheless spans an order of magnitude. The barbell shows the converse and explains why an edge count is not a transport measure: it

[1] Three nodes suffice. Take the triangle with $a_{01} = a_{02} = 1$ and $a_{12} = a$, unit budget and $\alpha = 1$. For $a < 3$ all three edges carry conductance. At $a = 3$ the edge $(1, 2)$ closes, and for $a \geq 3$ the optimum is the path $1{-}0{-}2$ with $w_{01} = w_{02} = 1/2$ and $\overline{R} = 8/3$. The threshold is exact: at that state $d_B^2(1,2) = 8$ and $d_B^2(0,1) = d_B^2(0,2) = 8/3$, so the closure condition $d_B^2(1,2)/a \leq d_B^2(0,1)/1$ first binds at $a = 3$. The linear case is the classical problem of Ref. [4], where zero weights are a generic feature of the solution.

carries more edges than either random graph in the table, but most of them lie inside its two dense cliques, and all intergroup transport passes through a three-edge bridge. Its allocation factor is the lowest in the table, $\eta_w = 0.373$, because uniform weights waste budget on clique edges that cannot relieve the bridge.

V. WHERE CONCENTRATION BEGINS TO PAY

For $0 < \alpha < 1$ the complete graph remains a symmetric stationary state but need not be optimal, and the simplest way to see where it loses is to race it against the opposite extreme. At unit budget the uniform star and the uniform complete graph have resistance ratio

$$\frac{\overline{R}_{\mathrm{star}}}{\overline{R}_{\mathrm{complete}}} = (n-1)\left(\frac{2}{n}\right)^{1/\alpha}, \tag{11}$$

which equals one at

$$\alpha_c(n) = \frac{\log(n/2)}{\log(n-1)}, \tag{12}$$

giving $\alpha_c = 0.747$ at $n = 12$. This is a crossing between two symmetric ansätze and nothing more. Figure 1 shows what the crossing misses: near α_c the best state found by multistart search has lower resistance than either endpoint, so it is neither symmetric endpoint.

What the crossing does capture is a trend in system size. Since $\alpha_c(n)$ increases toward one, the interval in which the star beats the complete graph widens as the network grows: $\alpha_c = 0.712$, 0.747, 0.768, and 0.782 at $n = 8$, 12, 16, and 20, and 0.851 at $n = 100$. Figure 2 shows that the measured boundary moves in the same direction. Resolving the largest tree-like exponent on a grid of step 0.005 gives $\alpha^* = 0.555$, 0.585, 0.610, 0.630, 0.635, 0.645, 0.650, 0.660, 0.665, 0.670 and 0.675 at $n = 8$, 12, 16, 20, 24, 28, 32, 40, 48, 56 and 64. The boundary rises at every size, as $\alpha_c(n)$ does, but lies between 0.15 and 0.16 below it throughout: the two symmetric ansatze bracket the mechanism without locating it. Sparse backbones are therefore not an artifact of small systems. How far this trend goes is the subject of Sec. VIII.

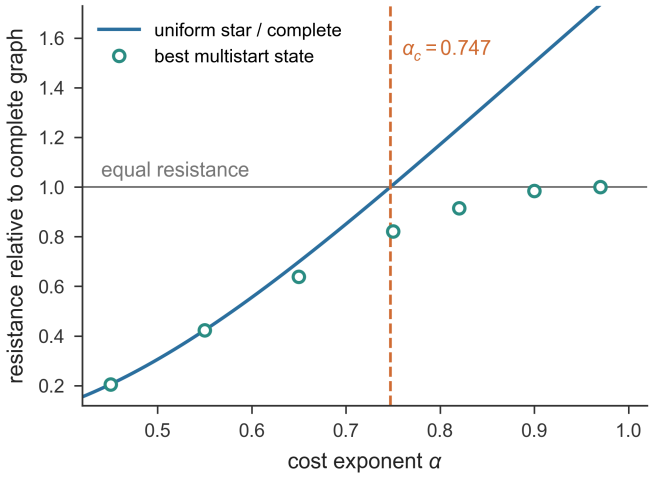


FIG. 1. Uniform star against uniform complete graph at $n = 12$ and $C_0 = 1$. Open circles give the lowest-resistance state obtained by multistart optimization, normalized by the complete-graph resistance at the same α. At the sampled values $\alpha \le 0.55$ the state found coincides with the star. Near the analytic crossing, intermediate supports beat both symmetric endpoints.

TABLE II. Convex reference factors for selected networks at $n = 24$ and $\alpha = 2$. η_w scores the allocation of conductance given the topology, η_{top} scores the topology against the complete graph, and $\eta_{\mathrm{free}} = \eta_{\mathrm{top}}\eta_w$ is the overall efficiency at the network's own budget. The budget is each network's own, $C_0 = \sum_e w_e^\alpha$ at unit weights, and $\overline{R}$ is compared at that budget using $\overline{R}(cW) = \overline{R}(W)/c$. The random graphs are `gnm_random_graph(24, 72, seed=1)`, `barabasi_albert_graph(24, 4, seed=1)` and `watts_strogatz_graph(24, 4, 0.5, seed=1)`; the remaining entries are `grid_2d_graph(6, 4)`, `star_graph(23)`, `barbell_graph(11, 2)`, `cycle_graph(24)` and `path_graph(24)`. Since $\alpha \ge 1$, no support truncation is applied.

Network	Edges	η_w	η_{top}	η_{free}
Complete	276	1.000	1.000	1.000
Erdős–Rényi	72	0.934	0.398	0.372
Barabási–Albert ($m = 4$)	80	0.878	0.397	0.348
Watts–Strogatz ($p = 0.5$)	48	0.979	0.294	0.288
Square lattice (6×4)	38	0.986	0.195	0.193
Star	23	1.000	0.151	0.151
Barbell	113	0.373	0.194	0.072
Ring	24	1.000	0.068	0.068
Path	23	0.968	0.036	0.035

VI. NUMERICAL PROCEDURE

Below $\alpha = 1$ it is convenient to work with the resource share of an edge,

$$p_e = \frac{a_e w_e^\alpha}{C_0}, \qquad p_e \ge 0, \qquad \sum_e p_e = 1, \tag{13}$$

from which the conductance is recovered as

$$w_e = \left[\frac{C_0\, p_e}{a_e}\right]^{1/\alpha}. \tag{14}$$

The budget constraint is then automatic and the search runs on the simplex.

The allocation search has two stages. The shares are first optimized in softmax coordinates with L-BFGS-B. Edges with $p_e \le \epsilon_{\mathrm{cut}}$, with $\epsilon_{\mathrm{cut}} = 10^{-5}$ throughout, are then set exactly to zero and the shares on each remaining connected support are refined with BFGS, which

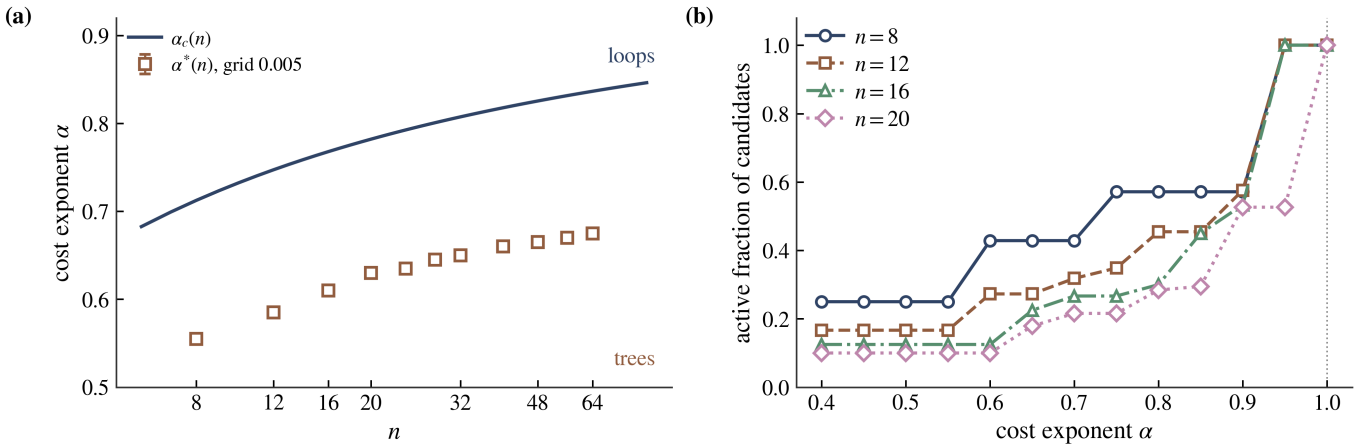


FIG. 2. Finite-size behavior of the sparse regime. (a) The analytic crossing $\alpha_c(n) = \log(n/2)/\log(n-1)$ (solid curve) against the loop-collapse exponent $\alpha^*(n)$, the largest exponent whose best state is a spanning tree, resolved at every tested size on a grid of step 0.005 by the downward sweep of Sec. VI. Error bars are one grid step. Both curves rise with n, so the tree regime widens with system size; α^* lies well below α_c, which compares only the two symmetric ansatze and therefore does not locate the boundary. (b) Fraction of candidate edges carrying conductance in the state found for the four sizes in the finer scan. Below $\alpha = 1$ the support occupies a smaller share of the candidate set as n grows.

removes any residual dependence on nearly vanishing conductances. The truncation stage is applied for $\alpha \leq 1$: only for $\alpha > 1$ does Sec. III give every admissible edge positive conductance at the optimum, so truncating there could only discard legitimately small weights. Feasibility of a support is decided topologically rather than from the algebraic connectivity, so that a connected support carrying uniformly small conductances is not misread as disconnected.

Each exponent is initialized from a uniform state, a star, a minimum spanning tree, a depth-first tree, and 60 seeded Dirichlet allocations, for 64 starts in total; the finite-size scan of Fig. 2 and the size scan of Sec. VIII use the same four structured states with 24 random starts, and the convex spatial cases use 16.

Equal numbers of starts at different sizes is not equal search quality. From $n = 8$ to $n = 64$ the number of allocation variables grows by a factor of 72, while the fraction of starts reaching the best value is already between one and five in 64 at $n = 12$. We therefore add a continuation stage: at fixed n the exponent is swept downward and the winning state at one exponent is supplied as an additional start at the next, and the better of the blind and continued results is kept. The outcome can only improve on blind multistart, since the blind starts are retained. Measured this way the blind protocol is the weaker of the two, and the states it misses differ structurally and not only in resistance. At $n = 32$ it returns $\beta = 42$ at $\alpha = 0.75$ and $\beta = 27$ at $\alpha = 0.66$, where continuation reaches $\beta = 53$ and $\beta = 33$ at resistances lower by 8.5×10^{-4} and 5.8×10^{-3} relative. The discrepancy in cycle rank is not one-signed: at $\alpha = 0.68$ the blind state carries one cycle *more* than the better state, $\beta = 34$ against 33, at a higher resistance. What the comparison establishes is therefore not a bias of fixed sign but that a cycle rank read off a blind search is a property of the state that search happened to find, and can differ from the best available state by as much as eleven independent loops. On the fine grid of Sec. VIII the same weakness appears as a systematic offset: at seven of the eleven sizes the blind protocol returns a spanning tree one grid step above α^*, where continuation still holds a looped state, so a collapse exponent read from blind search alone would sit one step too high.

Every state reported below, including the size scan of Table IV and Fig. 4, uses the continued protocol. Continuation improved on blind multistart at five of the 44 grid points: four at $\alpha = 0.75$, where it raises the cycle rank from 5 to 9 at $n = 8$, from 64 to 83 at $n = 48$, from 83 to 93 at $n = 56$ and from 88 to 114 at $n = 64$, lowering the resistance by 1.6, 5.6, 10.8 and 9.8×10^{-3} relative; and one at $\alpha = 0.65$ and $n = 24$, from 19 to 25 at a gain of 6.6×10^{-3}. It never improved on the blind result at $\alpha = 0.55$, where every state is a spanning tree, or at $\alpha = 0.90$.

We checked the implementation in both coordinate systems. Central differences in the conductances reproduce Eq. (6) with a maximum relative discrepancy of 2.5×10^{-8}. The optimizer works in logit coordinates, so we checked that composite gradient separately after projecting out the $\mathbf{1}$ direction left invariant by the softmax. Its maximum relative error is 9.0×10^{-8}. For every state reported below, the budget residual $|C_\alpha(W) - C_0|$ vanishes to machine precision. The largest infinity norm of the feasible projected gradient in the bounded logit coordinates is 3×10^{-6}.

The dispersion across starts is a practical measure of nonconvexity, and it is substantial. Counting a start as reaching the best state when its resistance lies within 10^{-6} relative to the smallest value found, between one

and five of 64 starts qualify at each exponent on the complete candidate graph. With $R_{\min}$ and $R_{\max}$ the smallest and largest final resistances, the relative spread $(R_{\max} - R_{\min})/R_{\min}$ is 3.87 at $\alpha = 0.45$, stays above one through $\alpha = 0.82$, and falls below one at $\alpha = 0.90$ and 0.97; the interquartile range divided by $R_{\min}$ falls from 0.092 to 0.003 across the same interval. Every support transition reported below therefore refers to the best state obtained at a stated search budget, not to a certified global minimum. The controls described below focus on the conclusions most sensitive to this distinction.

The calculations begin with the complete graph on 12 nodes and $a_e = 1$. Spatial effects are examined on a random geometric graph with 28 nodes, connection radius 0.38, seed 9, 95 edges, and $a_e = \ell_e$. The loading comparison of Sec. X uses a second geometric graph with 20 nodes, connection radius 0.42, seed 5, and 62 edges. These seeded graphs are controlled test beds and not empirical data sets.

VII. FROM BACKBONE TO MESH AT FIXED TOPOLOGY

To isolate the effect of cost convexity we hold the candidate topology, the budget, and the all-pairs loading fixed, and vary only α. The candidate graph is K_{12}, with 66 admissible edges and $a_e = 1$.

Table III reports the lowest-resistance allocation found at each exponent. The strongly concave states are sharply localized: at $\alpha = 0.45$ and 0.55 the best support has 11 active edges and agrees with the uniform star to numerical precision. It grows to 18 edges at $\alpha = 0.65$, 36 at $\alpha = 0.90$, and all 66 at $\alpha = 0.97$. The effective support tracks the active-edge count closely, which says the shares on the support are not far from uniform, while only one to five starts in 64 reach the best value at any exponent.

Figure 3 shows the same progression as a structural sequence. The active-edge count and the effective support sit at the tree value for $\alpha = 0.45$ and 0.55, then rise in steps as α approaches one, and near $\alpha = 0.97$ the search returns the complete graph. This approach to full support

TABLE III. Best multistart states on the complete candidate graph at $n = 12$, from 64 starts. Effective support is $1/\sum_e p_e^2$ and an edge counts as active when $p_e > 10^{-5}$. The last two columns give the fraction of starts reaching the best value and the relative range of final resistances across starts.

α	$\overline{R}_{\text{best}}/\overline{R}_{\text{complete}}$	Active edges	Eff. support	Starts at best	Spread
0.45	0.205	11	11.0	3.1%	3.87
0.55	0.423	11	11.0	3.1%	1.36
0.65	0.638	18	18.0	1.6%	1.48
0.75	0.821	23	22.8	1.6%	1.79
0.82	0.914	26	25.6	1.6%	2.08
0.90	0.984	36	35.4	1.6%	0.53
0.97	1.000	66	66.0	7.8%	0.73

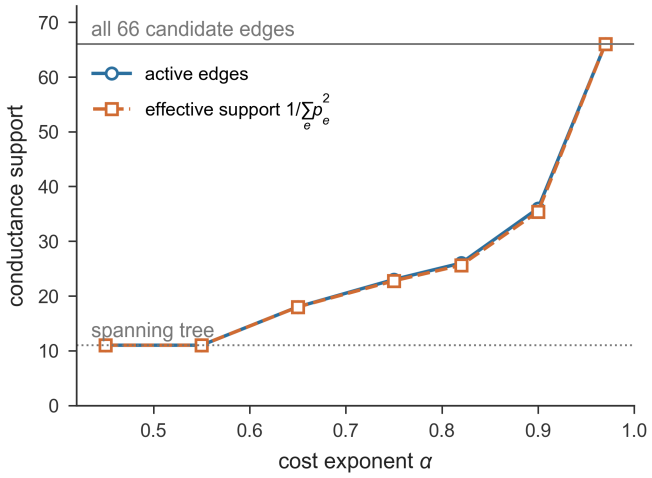


FIG. 3. Conductance support of the best multistart state on the complete candidate graph. The tree value $n - 1$ is marked by a dotted line. The support stays tree-like at small α and expands in steps as the cost becomes less concave.

is consistent with the convex limit at $\alpha = 1$, but the steps themselves should not be read as a thermodynamic critical sequence: the system is finite, the continuation in α is discrete, and the objective is nonconvex with genuinely competing minima. What is robust is the direction of travel, from a sparse backbone to a distributed support.

VIII. FINITE-SIZE GROWTH OF THE ACTIVE SUPPORT

The preceding results concern relatively small graphs. On a complete candidate graph the number of available edges grows as n^2, so the active fraction alone can give a misleading picture of sparsity. We instead examine the absolute support size. A dense state contains $O(n^2)$ active edges, whereas a backbone contains $O(n)$. Distinguishing these trends requires larger systems.

We fix four concave exponents and increase n to 64. At the largest size, the complete candidate graph has 2016 allocation variables. The search budget remains at 28 starts throughout the grid, keeping the comparison between sizes on equal footing. For a connected active support the cycle rank is $\beta = m_{\text{active}} - n + 1$, the number of independent loops. Table IV summarizes the behavior at each exponent and reports the largest-size values. Figure 4 shows the complete size dependence; all 44 numerical states are listed in Supplemental Table S1.

Over the range examined the states found grow approximately linearly in n. At $\alpha = 0.75$ the ratio active$/(n-1)$ fluctuates about 2.67 over the last four sizes; at $\alpha = 0.90$ the same ratio is near 4.69 over the last four. We quote these directly measured per-node ratios rather than a fitted slope. A straight-line fit over $n \geq 20$ returns a slope that need not agree with the local per-node ratio, because the fit also carries an intercept, and with only

TABLE IV. Summary of the size scan on the complete candidate graph, $a_e = 1$. Values in the last four columns correspond to the largest size, $n = 64$. All entries are best multistart states, not certified global optima. The complete grid of 44 (α, n) combinations is reported in Supplemental Table S1. Active edges are counted at $p_e > 10^{-5}$ and β is the cycle rank.

α	behavior over the tested sizes	Active	β	Active$/(n-1)$	β/n
0.55	tree at all 11 sizes	63	0	1.00	0.00
0.65	tree for $n \geq 32$	63	0	1.00	0.00
0.75	rising per-node ratio	177	114	2.81	1.78
0.90	looped; settles after $n = 40$	300	237	4.76	3.70

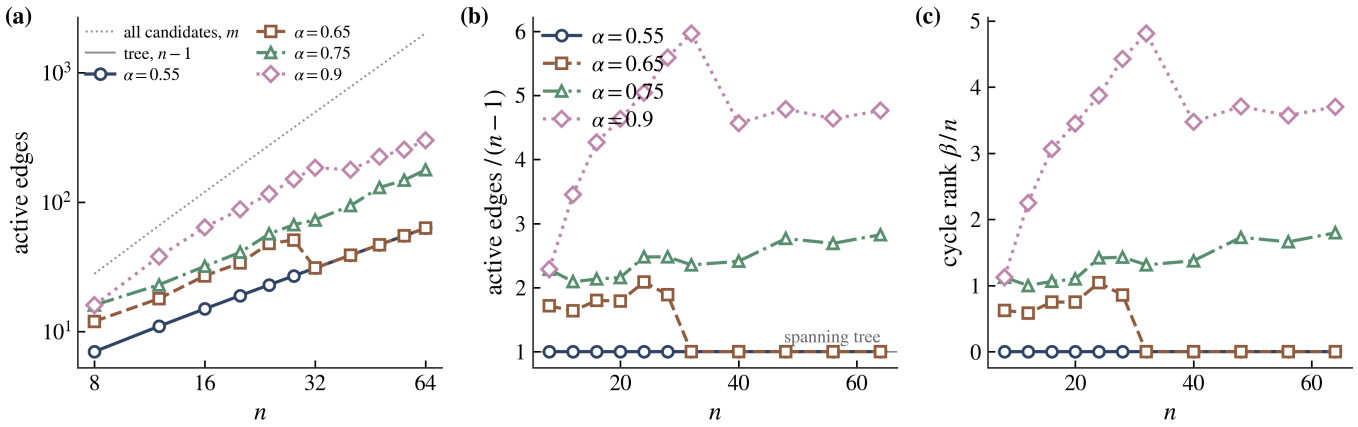


FIG. 4. Growth of the support in the concave states found. (a) Active edges against n on logarithmic axes, bracketed by the spanning-tree line $n-1$ and the full candidate set $m = n(n-1)/2$. At the larger sizes the trends are approximately linear in n, and every series stays far below the candidate line. (b) Active edges per node, which varies little at the larger sizes, with a level set by α; the $\alpha = 0.90$ series is still drifting below $n = 40$. (c) Cycle rank per node, which behaves in the same way, and drops to zero at $\alpha = 0.65$ between $n = 28$ and $n = 32$.

four points at the largest sizes the two estimators are not distinguishable in any case. The active fraction of the candidate set decreases roughly as $1/n$, which is not an independent observation but the algebraic consequence of linear growth with $m = n(n-1)/2$. At $\alpha = 0.75$, the product $n\,(\text{active}/m)$ stays between 4.7 and 5.6 from $n = 24$ to 64, while m grows by a factor of seven. At $\alpha = 0.55$ no fit is needed: the state found is a spanning tree, with $n-1$ edges and $\beta = 0$, at all eleven sizes. The normalized counts active$/(n-1)$ and β/n support a limited but direct conclusion. Over the accessible range, the number of channels per node changes much less than the fraction of the candidate graph that is used. Whether that per-node count remains bounded as $n \to \infty$ is not settled here, and the complete candidate graph does not test how it depends on the density of admissible edges.

These finite sizes do not establish an asymptotic growth law. Over $8 \leq n \leq 64$ a linear trend cannot be separated from $n^{1.08}$ or from $n \log n$: the effective power-law exponents at $\alpha = 0.75$ and 0.90 are 1.215 and 0.955, and with eleven sizes spanning less than a decade no model comparison among these candidates is meaningful. We therefore rely on the directly measured per-node ratios rather than assigning an asymptotic exponent to this finite range.

Every active-edge count and cycle rank in this work inherits the cut $p_e > 10^{-5}$, and it is worth being precise about what we have and have not tested there. Recounting the reported states at 10^{-4}, 10^{-5} and 10^{-6} leaves the active-edge count and the cycle rank identical at all three cuts for $(\alpha, n) = (0.55, 32)$, $(0.65, 28)$, $(0.65, 32)$, $(0.75, 32)$, $(0.75, 64)$, $(0.90, 32)$ and $(0.90, 64)$. Taken alone that invariance would be expected rather than informative, since the support-refinement step of Sec. VI sets the discarded shares to exactly zero and a recount afterwards cannot disagree with the cut that produced the support. We therefore ran the control the objection calls for. The search is repeated once with truncation disabled, and the refinement is rerun from that same pretruncation solution at each cut, so the active-edge count and the cycle rank are genuine functions of ϵ_{cut}. They do not move: at all three cuts the seven states give 31, 51, 31, 73, 151, 185 and 300 active edges with $\beta = 0$, 24, 0, 42, 88, 154 and 237, and the refined resistances agree to every printed digit. What also speaks to the question is the margin, since across these states the smallest retained share ranges from 2.0×10^{-3} to 3.2×10^{-2}, two to four orders of magnitude above the loosest cut, so nothing sits near 10^{-5} for a different choice to reclassify. The threshold-free effective support agrees closely as well, 31.0 against 31 active edges at $(0.55, 32)$ and 294.3 against 300 at $(0.90, 64)$.

Loops behave differently from the support that carries them, and the $\alpha = 0.65$ row of Table IV is where this shows. The best support found carries loops up to $n = 28$, with β/n between 0.58 and 1.04 and no clean trend, and is

an exact spanning tree at every one of the five larger sizes, $n = 32$, 40, 48, 56, and 64. Within this search budget the best states found therefore change character with system size at fixed cost exponent: looped states win at $n \leq 28$ and none was found that beats the tree at any larger size. This is an apparent numerical transition in the states we can locate, and it concerns more than Eq. (12), which compares only the two symmetric ansätze. Whether the true optimum changes character at some size between 28 and 32 is not established here.

Because the apparent loss of loops could result from incomplete sampling, we performed two additional checks. First, we raised the blind search budget from 28 to 64 and then 144 starts. The best resistance at $\alpha = 0.65$ remains unchanged in every printed digit, at 381.631590 for $n = 32$ and 546.797241 for $n = 40$, a relative change of exactly zero. Since blind Dirichlet sampling explores 500 to 800 dimensions poorly, the second check follows the looped state directly. We take the 51-edge, 24-cycle state found at $n = 28$, carry it into K_{32} and K_{40} by retaining its shares and attaching each new node to the busiest node of the old support, and use that as the initial condition. The seeded state relaxes onto the same spanning tree, with a resistance identical to the blind result. The looped state that wins at $n = 28$ is therefore not simply overlooked at $n = 32$ and 40.

Both checks vary the sampling. Neither reaches the move that actually opens a loop, and for $\alpha < 1$ that move is not infinitesimal: the marginal price $\alpha a_e w_e^{\alpha-1}$ diverges as $w_e \to 0$, so no gradient step ever creates a channel from nothing, and raising the number of blind starts cannot supply what the search does not contain. Two further controls address this directly.

The third carries a looped state *downward in* α at fixed n, the direction complementary to the seeded continuation above, on a grid of step 0.005. At $n = 32$ the continued state holds at 64 active edges with $\beta = 33$ at every exponent from $\alpha = 0.680$ down to $\alpha = 0.655$, and at $\alpha = 0.650$ it falls to 31 edges with $\beta = 0$ and $\overline{R} = 381.631590$, the value the blind search returns. The collapse is thus abrupt rather than gradual, and it is abrupt at every size: the transition is single at all eleven, the largest tree-like and smallest looped exponent differing by exactly one grid step, and the cycle rank immediately above it is 2, 7, 10, 18, 18, 30, 33, 41, 54, 62 and 70 from $n = 8$ to $n = 64$. At $n = 64$ the support falls from 133 active edges to the 63 of a spanning tree across a step of 0.005 in α. A discontinuous onset of loops has been reported for optimal supply networks under prescribed loading, where it arises through a saddle-node bifurcation as the cost of new edges is lowered [18]; the collapse seen here is the same phenomenon viewed along the cost exponent.

Resolving $\alpha^*(n)$ also reframes the $\alpha = 0.65$ row of Table IV. That row is not a place where loops stop being worth buying at large n; it is a horizontal line in the (α, n) plane crossing a boundary that rises with n. Since $\alpha^* = 0.635$ and 0.645 at $n = 24$ and 28 but 0.650 at $n = 32$, the fixed exponent 0.65 sits above the boundary at the two smaller sizes and at or below it from $n = 32$ on, which is exactly the change of character the row reports. Read that way the apparent size-driven loss of loops is a crossing rather than a mechanism, and it predicts where the same crossing occurs at other exponents: $\alpha = 0.60$ would cross between $n = 12$ and 16, and $\alpha = 0.68$ has not crossed by $n = 64$.

The fourth tests the finite move itself. Taking the tree found at $(\alpha, n) = (0.65, 32)$, each of the 465 closed edges in turn is assigned a finite share ε, the remaining shares are renormalized, and the state is refined. At $\varepsilon = 0.01$, 0.02 and 0.05 no insertion improves on the incumbent: all 465×3 trials relax back to the same tree, the best recovered resistance differing from it by less than 2×10^{-12} in relative terms. The tree is thus stable against opening any single channel at these scales. Varying the random start set at fixed $(\alpha, n) = (0.65, 32)$ over three independent seeds also returns the same state in every printed digit.

These are robustness checks and not a proof: they were run at these sizes only, they probe single-edge insertions rather than simultaneous ones, and they cannot exclude some other looped basin, at these or at larger n, that beats the tree. Certification of global optimality is not available for this problem.

One entry in Table IV does move with the search budget, and we report it rather than smooth it. At $\alpha = 0.90$ and $n = 40$, going from 28 to 144 starts lowers the best resistance by 1.4×10^{-4} in relative terms and shifts the support from 178 to 179 active edges. The $\alpha = 0.90$ series is the densest tested and sits closest to the convex boundary, where the number of competing supports is largest, and its per-node ratio rises monotonically to 5.97 at $n = 32$ before falling to the plateau near 4.7. We read the sizes below $n = 40$ there as pre-asymptotic rather than as evidence against saturation, and we report them so the reader can judge that call.

IX. A LENGTH-WEIGHTED SPATIAL NETWORK

The complete candidate graph has no geometry, so the same calculation is repeated with the length penalty $a_e = \ell_e$ on a fixed random geometric graph. Alongside the resource effective support $N_p = 1/\sum_e p_e^2$ we now also record the conductance effective support $N_w = 1/\sum_e q_e^2$ with $q_e = w_e/\sum_j w_j$. The two differ because Eq. (14) is nonlinear, so a concentrated budget can still produce a broad conductance profile, and the gap between them is itself informative.

The progression of Table V and Fig. 5 matches the dimensionless case. At $\alpha = 0.55$ the best state has 30 active edges and $\beta = 3$: a near-tree backbone carrying three independent cycles, on a 28-node graph whose spanning trees have 27 edges. Five cycles are present at $\alpha = 0.75$; by $\alpha = 0.95$, 60 edges are active and $\beta = 33$; and the two strictly convex cases activate every candidate edge, as Sec. III requires. This example uses one realization of the

geometric graph and is a controlled test bed rather than a sample from an ensemble, so it shows that geometry is compatible with the same progression, not how the transition depends on geometry.

The linear boundary is the sharpest test of Sec. III. At $\alpha = 1$ the problem is still convex and the optimum unique, but the marginal price a_e no longer vanishes at $w_e = 0$, so an edge whose biharmonic distance does not pay for its length can be priced out. With $a_e = \ell_e$ this is exactly what happens: 86 of the 95 candidate edges are active and nine are excluded, between the 60 of $\alpha = 0.95$ and the full support of $\alpha = 1.25$. Sparsity at the convex boundary is therefore not a numerical artifact of the concave regime; it is the complementarity predicted by Eq. (7), realized on a graph with unequal prices rather than on the three-node example of Sec. III.

The physically motivated case sits at the other end of the table. For a Hagen–Poiseuille tube at fixed volume, $\alpha = 1/2$ with $a_e \propto \ell_e^{3/2}$, the state found has 29 active edges and $\beta = 2$, a near-tree backbone on the same 28-node candidate graph, which is the point of contact between this formalism and vascular architecture. Its access gain, 15.27, is more than twice the largest value reached anywhere else in the table.

The two effective supports separate as α grows. They are close in the concave cases, 23.2 against 24.5 at $\alpha = 0.55$, and diverge strongly in the convex ones, 30.9 against 76.2 at $\alpha = 2$. Convex cost thus spreads the conductance far more evenly than it spreads the budget, because the budget-to-conductance map compresses differences when $\alpha > 1$. The map amplifies them when $\alpha < 1$, and the ordering does reverse, though only well below the boundary: at $\alpha = 1/2$ the conductance is more concentrated than the budget that buys it, $N_w = 15.8$ against $N_p = 18.4$. The access gain $A_{\mathrm{best}}/A_{\mathrm{equal}}$ in Fig. 6 runs in the opposite direction, from 15.27 at $\alpha = 1/2$ and 6.07 at $\alpha = 0.55$ to 1.48 at $\alpha = 2$: allocation is worth most precisely where the state found is sparsest and hardest to find, and therefore most in the physically motivated fixed-volume case.

TABLE V. Length-weighted spatial network, $n = 28$ and $m = 95$. Access gain is optimized access divided by equal-share access at the same α and budget; values at different α are not comparable in physical units. The first row is the Hagen–Poiseuille fixed-volume law of Sec. II B, $\alpha = 1/2$ with $a_e \propto \ell_e^{3/2}$; the remaining rows use $a_e = \ell_e$. All rows use the same random start set.

α	a_e	Active edges	Cycle rank	N_p	N_w	Access gain
0.50	$\ell_e^{3/2}$	29	2	18.4	15.8	15.27
0.55	ℓ_e	30	3	23.2	24.5	6.07
0.75	ℓ_e	32	5	22.6	23.5	2.89
0.95	ℓ_e	60	33	34.7	38.8	2.07
1.00	ℓ_e	86	59	38.4	44.8	1.99
1.25	ℓ_e	95	68	36.9	56.9	1.77
2.00	ℓ_e	95	68	30.9	76.2	1.48

X. THE LOAD DECIDES WHETHER LOOPS ARE WORTH BUYING

The cost law is not the only control of architecture. What the network is asked to do matters just as much, and the cleanest way to show this is to change the loading while holding everything else fixed. On the 20-node geometric graph we compare the all-pairs objective against a single source–sink pair separated by six hops.

For the single pair, let f_e be an oriented unit flow from s to t. Thomson's principle writes r_{st} as the minimum of $\sum_e f_e^2/w_e$ over such flows. Eliminating W at fixed f under Eq. (3) leaves the problem of minimizing

$$S(f) = \sum_e a_e^{1/(\alpha+1)} \, |f_e|^{2\alpha/(\alpha+1)}, \tag{15}$$

the remaining positive factors being independent of f. For $0 < \alpha < 1$ the exponent $q = 2\alpha/(\alpha+1)$ lies below one, so $|f_e|^q$ is concave in $|f_e|$ on the non-negative half-line. The function S is not concave on the full space of signed flows, and the argument requires the orientation to be fixed first. Any unit s–t flow decomposes into simple path flows plus circulations. Removing the circulations cannot raise the dissipation functional, because it is increasing in each $|f_e|$ for $q > 0$ and the decomposition lowers $|f_e|$ pointwise; what remains is an acyclic flow, and adopting the orientation that this acyclic flow induces on each edge makes all remaining coordinates non-negative. On that directed polytope S is concave, the polytope is the convex hull of the simple s–t path flows, and a concave function attains its minimum at an extreme point. Some minimizer is therefore supported on a single path and carries no cycles.

The structure of this argument is not new. Banavar and co-workers showed that minimizing a concave function of the link flows at prescribed sources and sinks selects a spanning tree, by exactly this route: concavity moves the minimum to an extreme point of the flow polytope [5], a result then developed for concave transport costs and loop analysis by Shao and Zhou [7]. Equation (15) is the same mechanism written for a resistance objective, and the single-pair result here should be read as recovering their conclusion in our variables rather than as an independent finding.

What differs is the loading. Their optimum is defined at a prescribed flow pattern, so the concave argument applies directly and looplessness follows for the whole network. Our objective averages over every node pair, which is not a single flow problem and admits no such reduction: the concave argument bites only in the rank-one case $Q = b_{st} b_{st}^{\mathsf{T}}$, and at finite n the all-pairs states in Table VI carry up to 32 independent cycles. Two further differences follow from working with $\overline{R}$ rather than with dissipation at fixed flow. The convex side of the same problem supplies the reference state that makes Eq. (5) an absolute score and splits it into η_{top} and η_w, which has no counterpart in the prescribed-flow formulation. And the all-pairs ensemble leaves room for the size dependence

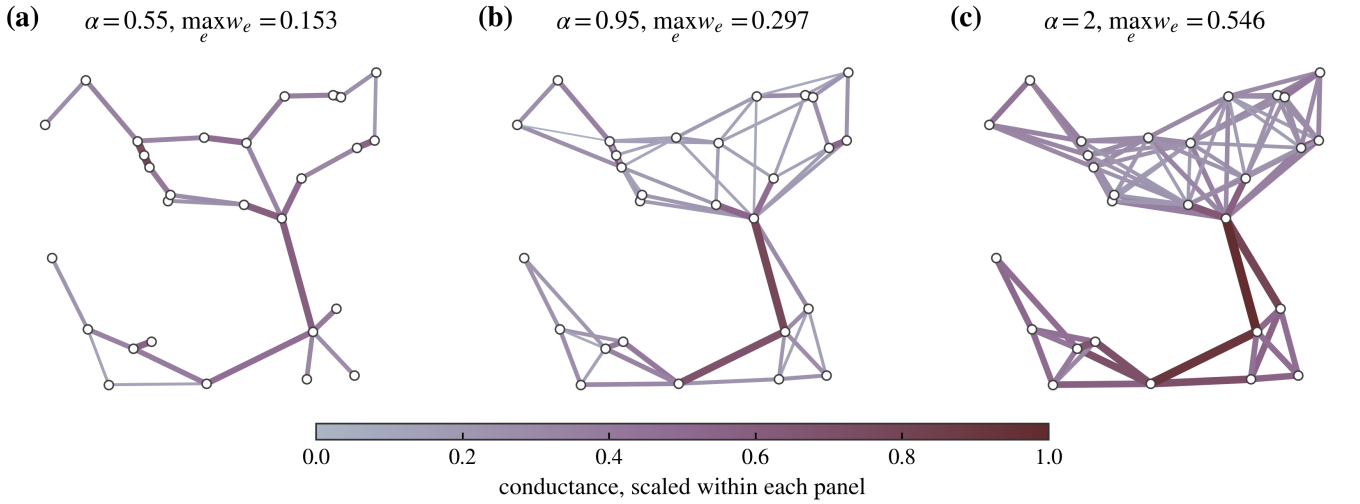


FIG. 5. Best multistart conductance states in the spatial candidate network at (a) $\alpha = 0.55$, (b) $\alpha = 0.95$, and (c) $\alpha = 2$. Line width and color encode conductance, scaled within each panel, with the panel maximum $\max_e w_e$ given in each title so that the three are comparable; edges with $p_e \leq 10^{-5}$ are omitted. The support changes from a sparse backbone to a loop-rich mesh as the cost law moves from the concave to the convex regime.

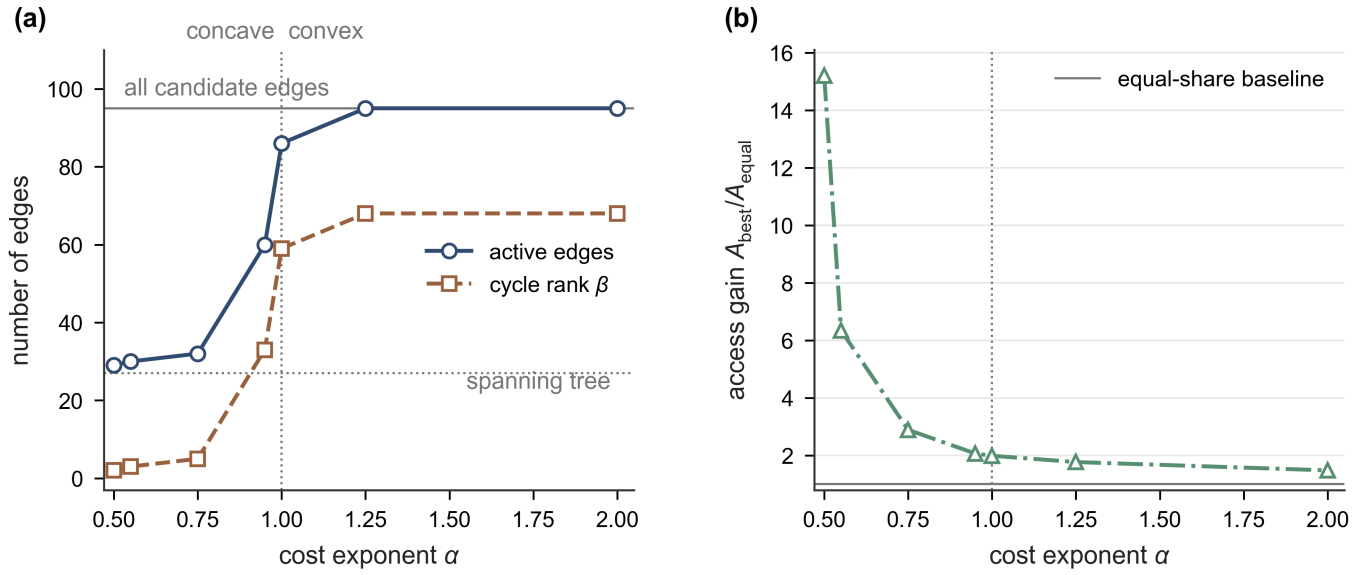


FIG. 6. Support topology and transport improvement in the spatial candidate graph. (a) Active edges and independent cycles against α; to the left of the dotted line at $\alpha = 1$ the cost is concave and a sparse backbone is stabilized, while convex cost recruits loop redundancy, with the spanning-tree and all-edges reference lines shown. (b) Access gain $A_{\text{best}}/A_{\text{equal}}$, the factor by which optimal allocation beats equal budget shares, largest in the concave regime where concentrating conductance on a few channels has the strongest effect.

of Sec. VIII, where whether loops survive is a question about n and not only about the cost exponent.

On a fixed path P, equalizing $\widetilde{\kappa}_e$ along it gives $w_e \propto a_e^{-1/(\alpha+1)}$ and an optimized resistance

$$r_{st}(P) = C_0^{-1/\alpha} \left[\sum_{e \in P} a_e^{1/(\alpha+1)} \right]^{1+1/\alpha}, \qquad (16)$$

so the optimal P is the shortest path under the transformed edge price $a_e^{1/(\alpha+1)}$. With equal prices this is the ordinary six-edge shortest route used in Table VI.

Table VI shows this contrast. Under all-pairs loading the effective support grows from 19.9 at $\alpha = 0.45$ to 33.5 at $\alpha = 0.97$ and the cycle rank from 2 to 32. The exact single-pair optimum has six active edges and no cycles at every exponent, and the looplessness holds for all $0 < \alpha < 1$, although which path is selected can depend on α when the coefficients a_e are unequal.

TABLE VI. All-pairs loading against the exact single-pair path on the same geometric candidate graph, $n = 20$ and $m = 62$. The path is loopless at every exponent shown. The all-pairs search uses the same two-stage optimizer and continuation protocol as the rest of the paper; relative to the earlier pipeline only the $\alpha = 0.65$ row moves, to a lower resistance with one additional active edge.

α	Eff. support (all pairs)	Active (all pairs)	Cycles (all pairs)	Edges in exact single-pair path
0.45	19.9	21	2	6
0.55	21.6	23	4	6
0.65	24.0	27	8	6
0.75	26.7	33	14	6
0.90	32.7	43	24	6
0.97	33.5	51	32	6

Prescribing a source and a sink removes the mechanism that pays for loops here. Redundancy is purchased when the network must serve pairs that do not share a route; for a single pair the concave flow cost selects one path exactly. Read together with Sec. VIII, the two results point the same way: loops in the best concave all-pairs states found are contingent, disappearing when the loading is narrowed to one pair and also, at $\alpha = 0.65$, when the system is simply made larger.

XI. READING THE MARGINAL LAW ON A NETWORK ONE DID NOT DESIGN

Equation (7) can be evaluated on a measured network without optimizing its topology at all, which is how it would be used in practice. Figure 7 does this for a barbell of two eight-node cliques joined by a three-edge bridge path through two intermediate nodes, with unit conductances, unit edge-price coefficients, and $\alpha = 2$. This is a smaller graph than the 24-node `barbell_graph(11, 2)` of Table II. The largest $\widetilde{\kappa}_e$ lie on the bridge, and $\widetilde{\kappa}_{\max}/\widetilde{\kappa}_{\min} = 144$, or 2.16 orders of magnitude. A uniform network is thus very far from equalizing marginal value, which is precisely the signal Eq. (7) is meant to carry.

Panel (b) reports the aggregate $\kappa_v = \sum_{e \ni v} \widetilde{\kappa}_e$. This is a diagnostic sum of incident edge scores, equivalent to assigning the same infinitesimal resource increment to every edge at v. It is not an intrinsic derivative of η with respect to a node, and it should not be quoted as one.

XII. DISCUSSION

The exponent α separates two mechanisms that are easy to conflate. For $\alpha > 1$ the marginal cost of activating an absent edge vanishes, so every admissible channel receives positive conductance at the optimum and the design problem is entirely one of allocation. For $\alpha < 1$ the first increment on a new channel has a divergent marginal cost, so access improves mainly by reinforcing an existing backbone, and topology becomes part of the answer rather than an input. Loops return only when the reduction they buy in the all-pairs effective resistance outweighs the price of opening additional channels.

Section VIII sharpens what "sparse" means here. Across the sizes we can reach, the concave states found use a decreasing fraction of the candidate set, but the number of channels per node changes much less. An active fraction measured at $n = 12$ or 20 will therefore not transfer directly to a larger graph. The behavior of loops is less simple. At $\alpha = 0.65$ the cycles present through $n = 28$ are gone at all five larger sizes, and seeded continuation at $n = 32$ and 40 does not recover them. Some of the redundancy seen in small concave states may therefore be a finite-size effect rather than a general consequence of the cost law. This does not happen over the accessible range at $\alpha = 0.75$, where β/n averages 1.64 over the last four sizes and reaches 1.78 at $n = 64$. If a size-driven crossover exists at that exponent, it lies beyond the present calculation.

The closest earlier comparison is the work of Shao and Zhou [7], who studied concave transport costs and the stability of loops for prescribed sources and sinks. Their objective is tied to that load and to the flows it produces. Here the mean effective resistance averages over every unordered node pair, so one added edge can improve many transport routes at once. The size dependence must therefore be read for this all-pairs ensemble. At $\alpha = 0.65$ the states found carry loops through $n = 28$ but become trees from $n = 32$ onward. At $\alpha = 0.75$, by contrast, β/n averages 1.64 over the last four sizes and is still rising at $n = 64$, where it reaches 1.78. Concave cost alone does not make loops inevitable; their value depends on the loading ensemble and on system size. Corson reached a complementary conclusion for prescribed loading [19]: with a single stationary source configuration the optimum is a tree, and loops appear once the load fluctuates, so redundancy is bought by variability in the demand rather than by the cost law. Our all-pairs objective is a stationary but non-degenerate load, and it sits between his two cases, which is consistent with loops that are present but contingent on α and n. Kaiser *et al.* showed that the onset of loops in optimal supply networks is discontinuous, arising through a saddle-node bifurcation [18]; the abrupt collapse localized in Sec. VIII has the same character along the cost exponent. This complements related studies of optimal transport architectures [6, 8, 10, 11].

The convex index is the part of this that is ready for use on data. Given an observed network and a candidate edge set, η_w scores the allocation of conductance and η_{top} scores the restriction imposed by the wiring, both at the network's own budget. The comparison is only meaningful when weight, cost exponent, geometric coefficient, node set, and transport ensemble carry the same interpretation in the two systems being compared. A correlation matrix, a hydraulic conductance, and an electrical admittance do not belong on one scale merely because each is called a weighted network.

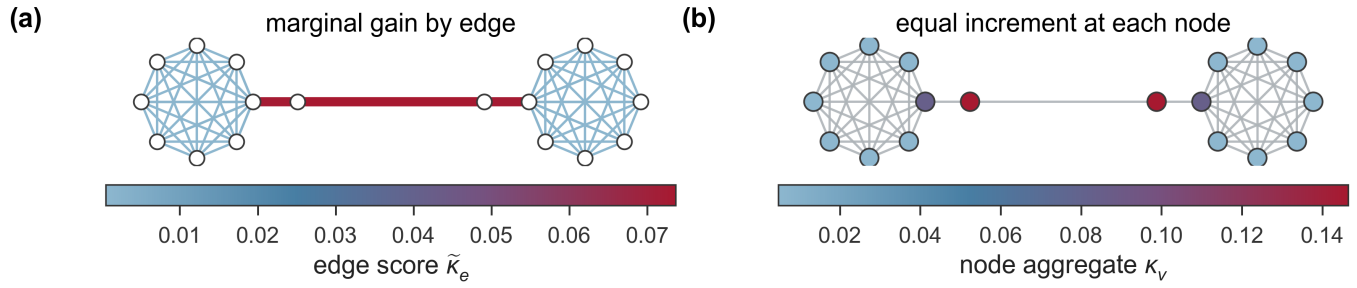


FIG. 7. Marginal access gain per unit cost in a barbell network with two eight-node cliques and two intermediate nodes, at $\alpha = 2$. Darker and thicker bridge edges carry the largest sensitivity; the maximum-to-minimum ratio is 144, or 2.16 orders of magnitude. Panel (a) shows the edge scores. Panel (b) sums the incident edge gains at each node for the equal-increment direction.

The spatial networks considered here are synthetic. Measured channel geometries will be needed to test the length penalty in a physical system, and the fixed-volume case of Sec. IX is the natural point of contact with measured vascular architectures [20]. More general loading is representable by a symmetric positive-semidefinite demand matrix Q with $Q\mathbf{1} = 0$, giving the objective $\mathrm{tr}(L^{+}Q)$; all-pairs loading is the isotropic choice on $\mathbf{1}^{\perp}$ and a single pair is the rank-one case $Q = b_{st}b_{st}^{\mathsf{T}}$. The perturbation argument of Sec. III carries over unchanged, and whether the per-node loop count is controlled by the spectrum of Q is a well-posed question that the two extreme cases here leave open. The concave all-pairs states remain numerical candidates throughout, and every support size, cycle rank, and crossover reported for $\alpha < 1$ is a property of the best state found at a stated search budget.

Certifying global optimality on general graphs is a separate problem, and the controls in Sec. VIII bound the risk without removing it. There is, however, a way to recover an absolute statement below $\alpha = 1$ without solving the nonconvex problem, and it seems worth stating as a target. Any lower bound $L \le \overline{R}^{*}$ obtained independently of the search, whether from a convex relaxation of Eq. (4) or from a combinatorial argument on the candidate set, immediately gives $L/\overline{R} \le \eta_\alpha \le U/\overline{R}$ with U the resistance of any feasible state, in particular the best one found. That is a certified interval rather than a point estimate, it narrows as either bound improves, and it would extend Eq. (5) to the concave regime where we currently decline to use it.

A useful relaxation is not obvious, since the difficulty sits in the concave cost rather than in $\overline{R}$, which is convex throughout; we have not pursued it here. The optimizer iterations should not be interpreted as a physical adaptation process. A biological or materials model would additionally need a mobility law, kinetic coefficients, and a mechanism linking local observations to conductance changes; adaptive dynamics of that kind, in which conductances respond to the flows they carry, are the subject of a separate line of work [21].

XIII. CONCLUSION

A finite resource budget gives a common physical basis for comparing weighted transport networks. Where the cost is convex, the mean effective resistance has a unique optimal allocation, which turns a normalized efficiency into a genuine measurement and factors it into an allocation part and a topological part. Its gradient is a squared biharmonic distance, giving the marginal worth of every channel and a marginal-value balance of Murray type in which the balanced quantity is a property of the whole network. Concave cost changes both the allocation and the active topology. Sparse supports and competing local minima appear, while the uniform star and complete graph provide only two analytic reference states. For $\alpha = 0.75$ and 0.90, the support found numerically grows approximately linearly with n over the sizes examined. The behavior of loops is less uniform. Resolving the loop-collapse exponent $\alpha^{*}(n)$ on a grid of step 0.005 shows it rising at every size from 0.555 at $n = 8$ to 0.675 at $n = 64$, between 0.15 and 0.16 below $\alpha_c(n)$ throughout, with a single and abrupt transition at every size. The change of character reported at the fixed exponent $\alpha = 0.65$ between $n = 28$ and $n = 32$ is this rising boundary being crossed, and at that point no looped state is recovered by seeded continuation in n, by continuation downward in α, or by opening any closed channel with a finite jump. Loading provides a separate control: for a single prescribed pair, Thomson's principle gives a loopless optimum throughout $0 < \alpha < 1$. Cost convexity and loading geometry together set both the access a network can reach and the architecture that delivers it.

ACKNOWLEDGMENTS

This work was supported by DGAPA-PAPIIT under Grant No. IN234029 and by SECIHTI under Grant No. CBF-2025-G-789.

DATA AVAILABILITY

All results reported here are generated by the executable notebook provided as Supplemental Material, which also documents the software environment used. Supplemental Table S1 lists the 44 states of the size scan.

Claude Opus 5 (Anthropic) was used for language editing, for assistance with code scripting, and to rerun part of the calculations in a second software environment, as recorded in the notebook. All results were verified by the authors, who take full responsibility for the content.

[1] D. J. Klein and M. Randić, Resistance distance, J. Math. Chem. **12**, 81 (1993).

[2] V. Latora and M. Marchiori, Efficient behavior of small-world networks, Phys. Rev. Lett. **87**, 198701 (2001).

[3] M. E. J. Newman, *Networks: An Introduction* (Oxford University Press, Oxford, 2010).

[4] A. Ghosh, S. Boyd, and A. Saberi, Minimizing effective resistance of a graph, SIAM Rev. **50**, 37 (2008).

[5] J. R. Banavar, F. Colaiori, A. Flammini, A. Maritan, and A. Rinaldo, Topology of the fittest transportation network, Phys. Rev. Lett. **84**, 4745 (2000).

[6] S. Bohn and M. O. Magnasco, Structure, scaling, and phase transition in the optimal transport network, Phys. Rev. Lett. **98**, 088702 (2007).

[7] Z. Shao and H. Zhou, Optimal transportation network with concave cost functions: Loop analysis and algorithms, Phys. Rev. E **75**, 066112 (2007).

[8] M. Durand, Architecture of optimal transport networks, Phys. Rev. E **73**, 016116 (2006).

[9] M. Durand, Structure of optimal transport networks subject to a global constraint, Phys. Rev. Lett. **98**, 088701 (2007).

[10] E. Katifori, G. J. Szöllősi, and M. O. Magnasco, Damage and fluctuations induce loops in optimal transport networks, Phys. Rev. Lett. **104**, 048704 (2010).

[11] H. Ronellenfitsch and E. Katifori, Global optimization, local adaptation, and the role of growth in distribution networks, Phys. Rev. Lett. **117**, 138301 (2016).

[12] A. Bejan and S. Lorente, The constructal law and the thermodynamics of flow systems with configuration, Int. J. Heat Mass Transfer **47**, 3203 (2004).

[13] A. Bejan and S. Lorente, The constructal law of design and evolution in nature, Philos. Trans. R. Soc. B **365**, 1335 (2010).

[14] Y. Lipman, R. M. Rustamov, and T. A. Funkhouser, Biharmonic distance, ACM Trans. Graph. **29**, 27 (2010).

[15] M. Black, L. Lin, W.-K. Wong, and A. Nayyeri, Biharmonic distance of graphs and its higher-order variants: Theoretical properties with applications to centrality and clustering, in *Proceedings of the 41st International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 235 (PMLR, 2024) pp. 4077–4102.

[16] C. D. Murray, The physiological principle of minimum work. I. The vascular system and the cost of blood volume, Proc. Natl. Acad. Sci. U.S.A. **12**, 207 (1926).

[17] A. Lonardi, E. Facca, M. Putti, and C. De Bacco, Designing optimal networks for multicommodity transport problem, Phys. Rev. Research **3**, 043010 (2021).

[18] F. Kaiser, H. Ronellenfitsch, and D. Witthaut, Discontinuous transition to loop formation in optimal supply networks, Nat. Commun. **11**, 5796 (2020).

[19] F. Corson, Fluctuations and redundancy in optimal transport networks, Phys. Rev. Lett. **104**, 048703 (2010).

[20] H. Ronellenfitsch and E. Katifori, Phenotypes of vascular flow networks, Phys. Rev. Lett. **123**, 248101 (2019).

[21] D. Hu and D. Cai, Adaptation and optimization of biological transport networks, Phys. Rev. Lett. **111**, 138701 (2013).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- costconvexity.ipynb.zip