

Optimal Boundary Control of Diffusion on Graphs via Linear Programming ^{*}

Harbir Antil^{†1}, Rainald Löhner^{‡2}, and Felipe Pérez^{§3}

^{1,3}Department of Mathematical Sciences and the Center for Mathematics and Artificial Intelligence (CMAI), George Mason University, Fairfax, VA 22030, USA.

²Department of Physics and the Computational Fluid Dynamics Center, George Mason University, Fairfax, VA 22030, USA.

Abstract

We propose a linear programming (LP) framework for steady-state diffusion and flux optimization on geometric networks. The state variable satisfies a discrete diffusion law on a weighted, oriented graph, where conductances are scaled by edge lengths to preserve geometric fidelity. Boundary potentials act as controls that drive interior fluxes according to a linear network Laplacian. The optimization problem enforces physically meaningful sign and flux-cap constraints at all boundary edges, derived directly from a gradient bound $|\nabla u| \leq \phi_{\max}$. This yields a finite-dimensional LP whose feasible set is polyhedral, and whose boundedness and solvability follow from simple geometric or algebraic conditions on the network data.

We prove that under the absence of negative recession directions—automatically satisfied in the presence of finite box bounds, flux caps, or sign restrictions—the LP admits a global minimizer. Several sufficient conditions guaranteeing boundedness of the feasible region are identified, covering both full-rank and rank-deficient flux maps. The analysis connects classical results such as the Minkowski–Weyl decomposition, Hoffman’s bound, and the fundamental theorem of linear programming with modern network-based diffusion modeling.

Two large-scale examples illustrate the framework: (i) A typical large stadium in a major modern city, which forms a single connected component with relatively uniform corridor widths, and a (ii) A complex street network emanating from a large, historical city center, which forms a multi-component system.

In both cases, the LP formulation yields stable and physically consistent flux fields that satisfy sign and conservation constraints to within double-precision tolerance. The approach generalizes naturally to other diffusion-controlled phenomena such as thermal transport, chemical dispersion, or pedestrian flow in complex geometric environments.

1 Introduction

Many planning, transport, and diffusion processes on geometric networks can be formulated as steady-state boundary-control problems. The goal is to prescribe or optimize boundary potentials (Dirichlet data) so that the resulting flux field achieves a desired steady-state configuration inside the network. Examples include heat transport through composite materials, contaminant propagation in pipe systems, and pedestrian or vehicular flow in constrained infrastructures such as stadiums,

^{*}This work is partially supported by the Office of Naval Research (ONR) under Award NO: N00014-24-1-2147, NSF grant DMS-2408877, and the Air Force Office of Scientific Research (AFOSR) under Award NO: FA9550-25-1-0231.

[†]Email: hantil@gmu.edu

[‡]Email: rlohner@gmu.edu

[§]Email: fperezsi@gmu.edu

terminals, or urban layouts. In all such systems, the governing equations are linear, yet the control variables—boundary conditions—are subject to physical and geometric constraints that must be respected. Optimizing these boundary controls thus leads naturally to *PDE-constrained optimization* problems posed on graphs.

This work develops a general *linear programming (LP)* formulation for such diffusion-control problems on weighted, oriented graphs. Each edge $e = (t(e), h(e))$ (tail and head, respectively) is assigned a length L_e and a cross-sectional weight k_e , defining a conductance

$$c_e = k_e/L_e$$

consistent with Fick’s or Fourier’s law. The discrete flux $q_e = -c_e(u_{h(e)} - u_{t(e)})$ links vertex potentials $u_{t(e)}$ and $u_{h(e)}$, and vertex-level balance enforces discrete conservation. Boundary vertices are divided into inflow and outflow sets $(\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}})$, and the boundary potentials $\mathbf{g}$ serve as controls. The interior state is determined uniquely by solving a discrete Laplace equation, producing an *affine control-to-state mapping*

$$(\mathbf{u}, \mathbf{q}, \Phi) = (\mathbf{u}_0, \mathbf{q}_0, \Phi_0) + (\mathbf{U}, \mathbf{Q}, \mathbf{P}) \mathbf{g}.$$

This affine reduction converts the continuous diffusion control problem into a finite-dimensional LP with explicit geometric and physical inequality constraints.

Modeling. Three structural elements distinguish the present framework:

- (i) **Edge-level sign constraints.** Backflow is eliminated by enforcing orientation-aware sign restrictions on boundary fluxes, ensuring outwardness and physical admissibility.
- (ii) **Flux caps from gradient bounds.** A maximum gradient constraint $|\partial_s u| \leq \phi_{\max}$ implies per-edge flux caps $|q_e| \leq k_e \phi_{\max}$, yielding linear inequalities in $\mathbf{g}$.
- (iii) **Affine elimination of the state.** The linear diffusion operator permits exact elimination of interior variables, resulting in a polyhedral feasible set and a globally convex optimization problem.

Mathematical framework. The feasible region of the LP

$$P = \{\mathbf{g} \in \mathbb{R}^n : A\mathbf{g} \leq b, \mathbf{g}_{\min} \leq \mathbf{g} \leq \mathbf{g}_{\max}\}$$

is a convex polyhedron defined by affine flux relations and capacity bounds. We show that the objective functional, being linear in $\mathbf{g}$, attains a global minimum whenever no negative recession direction $\mathbf{d} \in \text{rec}(P)$ satisfies $\mathbf{c}^\top \mathbf{d} < 0$ (see (9) for definition of $\mathbf{c}$). This condition, verified using the Minkowski–Weyl decomposition of polyhedra, guarantees existence of an optimal solution without assuming boundedness *a priori*. Practical sufficient conditions for boundedness arise naturally: finite box bounds, full-rank flux maps with two-sided caps, or rank-deficient flux maps regularized by sign constraints. These results connect geometric arguments from convex analysis and linear programming—such as Hoffman’s error bound [7], Rockafellar’s theory of convex sets [11], and Schrijver’s foundational work on polyhedral theory [13]—to diffusion control on networks.

Algorithmic formulation. After eliminating the interior state, the optimization problem reduces to

$$\min_{\mathbf{g}} -(\mathbf{c}^*)^\top \mathbf{g} \quad \text{s.t.} \quad A_{\text{cap}} \mathbf{g} \leq b_{\text{cap}}, A_{\text{edge}} \mathbf{g} \leq b_{\text{edge}}, \mathbf{g}_{\min} \leq \mathbf{g} \leq \mathbf{g}_{\max},$$

where $(A_{\text{edge}}, b_{\text{edge}})$ encode per-edge sign constraints with a small slack $\varepsilon \geq 0$, and $(A_{\text{cap}}, b_{\text{cap}})$ enforce $|q_e| \leq \phi_{\max} k_e$. This LP can be solved directly using standard solvers, and the affine control–state map allows post-solution reconstruction of all nodal potentials and fluxes. The LP structure ensures global optimality, simple feasibility checks, and numerical robustness without requiring penalty parameters or nonlinear relaxations.

Connections to prior work. Diffusion and flow on graphs are classical subjects [4, 5, 3], and linear programming has long been used for network-flow optimization [1, 2, 14]. However, traditional approaches typically optimize edge flows or capacities directly. Our formulation differs in two key respects: (i) fluxes arise from a diffusive (Ohmic) relation rather than from independent capacity constraints; and (ii) the optimization acts on *boundary potentials*, not on edge injections, while enforcing sign constraints that preserve physical outwardness. On the analysis side, solvability and boundedness follow from convex-analytic principles (recession cones, Hoffman bounds) rather than combinatorial flow arguments. In the context of evacuation modeling, this framework is consistent with continuum potential-flow theories (e.g., [9]) but operates natively on networks, which is essential when only a line-graph abstraction is available. We also refer to [12], and references therein, for gas transportation on networks using PDE constrained optimization frameworks.

Contributions.

- A unified LP formulation for boundary-control diffusion problems with edge-level sign and flux-cap constraints.
- A geometric proof of existence and boundedness using recession cones and Minkowski–Weyl decomposition.
- Simple sufficient conditions for automatic boundedness under physical modeling assumptions.
- Numerical validation on two real-world networks confirming conservation, sign correctness, and geometric consistency.

2 Model on a metric graph

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a finite connected graph. Each edge $e \in \mathcal{E}$ is a 1D segment of length $L_e > 0$ with constant conductivity $k_e > 0$, oriented from its tail $t(e)$ to head $h(e)$. We are given disjoint boundary sets: inflows $\mathcal{V}_{\text{in}}$ and outflows $\mathcal{V}_{\text{out}}$. The user chooses a fixed subset $\mathcal{V}_{\text{fix}} \subseteq \mathcal{V}_{\text{in}}$ with prescribed temperatures u_{fix} . All remaining boundary nodes are controls:

$$\mathcal{V}_{\text{ctrl}} := (\mathcal{V}_{\text{in}} \cup \mathcal{V}_{\text{out}}) \setminus \mathcal{V}_{\text{fix}}, \quad \mathbf{g} := \mathbf{u}|_{\mathcal{V}_{\text{ctrl}}} \in \mathbb{R}^{|\mathcal{V}_{\text{ctrl}}|}.$$

Special cases: $\mathcal{V}_{\text{fix}} = \emptyset$ (all boundary are controls) or $\mathcal{V}_{\text{ctrl}} = \emptyset$ (no optimization). We denote nodal temperatures by u_v for $v \in \mathcal{V}$ and edgewise temperatures by $u_e(x)$ for $x \in (0, L_e)$ with $x = 0$ at $t(e)$ and $x = L_e$ at $h(e)$.

On edge e , the steady diffusion equation holds

$$-\frac{d}{dx} \left(k_e \frac{du}{dx} \right) = 0$$

The solution is affine: $u_e(x) = u_{t(e)} + \frac{x}{L_e} (u_{h(e)} - u_{t(e)})$. The (signed) constant flux along e (positive in the $t(e) \rightarrow h(e)$ direction) is

$$q_e := -k_e \frac{du_e}{dx} = -\frac{k_e}{L_e} (u_{h(e)} - u_{t(e)}) = -c_e (u_{h(e)} - u_{t(e)}),$$

with edge conductance $c_e := k_e/L_e$.

At node $v \in \mathcal{V}$, the nodal flux balance (inflow minus outflow) is

$$\Phi_v := \sum_{e \in \mathcal{E}: h(e)=v} q_e - \sum_{e \in \mathcal{E}: t(e)=v} q_e, \quad \sum_{v \in \mathcal{V}} \Phi_v = 0.$$

Interior nodes enforce conservation: $\Phi_v = 0$ for all $v \notin \mathcal{V}_{\text{in}} \cup \mathcal{V}_{\text{out}}$.

3 Discrete operators and shapes

Index nodes by $1, \dots, n_V$ and orient each edge with a tail and head. We use an *edge-by-node* incidence matrix $\mathbf{B} \in \mathbb{R}^{n_E \times n_V}$:

$$(\mathbf{B})_{e,v} = \begin{cases} -1, & v = t(e), \\ +1, & v = h(e), \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$




Figure 1: Orientation and signs. Incidence $\mathbf{B}$ is edge-by-node with -1 at tail and $+1$ at head. Then $\mathbf{q} = -\mathbf{CB}\mathbf{u}$ and $\Phi = \mathbf{B}^\top \mathbf{q}$ implement the definitions compactly.

Let $\mathbf{C} = \text{diag}(c_e) \in \mathbb{R}^{n_E \times n_E}$. For a nodal vector $\mathbf{u} \in \mathbb{R}^{n_V}$,

$$\begin{aligned} \text{edge fluxes: } \mathbf{q} &= -\mathbf{CB}\mathbf{u} \in \mathbb{R}^{n_E}, \\ \text{nodal balances: } \Phi &= \mathbf{B}^\top \mathbf{q} = -\mathbf{B}^\top \mathbf{CB}\mathbf{u} =: -\mathbf{L}\mathbf{u}, \end{aligned} \quad (2)$$

where $\mathbf{L} := \mathbf{B}^\top \mathbf{CB} \in \mathbb{R}^{n_V \times n_V}$ is the weighted graph Laplacian. To summarize, we have: $\mathbf{B} \in \mathbb{R}^{n_E \times n_V}$, $\mathbf{C} \in \mathbb{R}^{n_E \times n_E}$, $\mathbf{L} \in \mathbb{R}^{n_V \times n_V}$, $\mathbf{u} \in \mathbb{R}^{n_V \times 1}$, $\mathbf{q} \in \mathbb{R}^{n_E \times 1}$, $\Phi \in \mathbb{R}^{n_V \times 1}$.

Embedding matrices. For any node index set $S \subset \{1, \dots, n_V\}$, the embedding $\mathbf{E}_S \in \mathbb{R}^{n_V \times |S|}$ places a subvector at the global indices of S . We will use $\mathbf{E}_{\text{fix}}$ (for $\mathcal{V}_{\text{fix}}$), $\mathbf{E}_{\text{ctrl}}$ (for $\mathcal{V}_{\text{ctrl}}$), and $\mathbf{E}_{\text{int}}$ (for $\mathcal{V}_{\text{int}} := \{1, \dots, n_V\} \setminus (\mathcal{V}_{\text{fix}} \cup \mathcal{V}_{\text{ctrl}})$).

4 Mixed Dirichlet boundary and affine maps

Write the nodal state as

$$\mathbf{u} = \mathbf{E}_{\text{int}} \mathbf{u}_{\text{int}} + \mathbf{E}_{\text{fix}} \mathbf{u}_{\text{fix}} + \mathbf{E}_{\text{ctrl}} \mathbf{g}. \quad (3)$$

Interior conservation ($\Phi|_{\mathcal{V}_{\text{int}}} = 0$) gives

$$0 = -\mathbf{E}_{\text{int}}^\top \mathbf{L} \mathbf{u} = -\left(\underbrace{\mathbf{E}_{\text{int}}^\top \mathbf{L} \mathbf{E}_{\text{int}}}_{\mathbf{L}_{ii}} \mathbf{u}_{\text{int}} + \underbrace{\mathbf{E}_{\text{int}}^\top \mathbf{L} \mathbf{E}_{\text{fix}}}_{\mathbf{L}_{i,\text{fix}}} \mathbf{u}_{\text{fix}} + \underbrace{\mathbf{E}_{\text{int}}^\top \mathbf{L} \mathbf{E}_{\text{ctrl}}}_{\mathbf{L}_{i,\text{ctrl}}} \mathbf{g} \right). \quad (4)$$

Assuming each connected component has at least one Dirichlet node (fixed or controlled), $\mathbf{L}_{ii} \succ 0$ (see below) and

$$\mathbf{u}_{\text{int}}(\mathbf{g}) = \underbrace{-\mathbf{L}_{ii}^{-1}\mathbf{L}_{i,\text{fix}}\mathbf{u}_{\text{fix}}}_{A_0} + \underbrace{-\mathbf{L}_{ii}^{-1}\mathbf{L}_{i,\text{ctrl}}\mathbf{g}}_{A_1}. \quad (5)$$

Reconstruct the full state

$$\mathbf{u}(\mathbf{g}) = \mathbf{u}_0 + \mathbf{U}_g\mathbf{g}, \quad \mathbf{u}_0 := \mathbf{E}_{\text{int}}A_0 + \mathbf{E}_{\text{fix}}\mathbf{u}_{\text{fix}}, \quad \mathbf{U}_g := \mathbf{E}_{\text{int}}A_1 + \mathbf{E}_{\text{ctrl}}. \quad (6)$$

Then

$$\begin{aligned} \mathbf{q}(\mathbf{g}) &= \mathbf{q}_0 + \mathbf{Q}_g\mathbf{g} := -\mathbf{CB}\mathbf{u}_0 + (-\mathbf{CB}\mathbf{U}_g)\mathbf{g}, \\ \Phi(\mathbf{g}) &= \Phi_0 + \mathbf{P}_g\mathbf{g} := \mathbf{B}^\top\mathbf{q}_0 + (\mathbf{B}^\top\mathbf{Q}_g)\mathbf{g}. \end{aligned} \quad (7)$$

For boundary blocks, define

$$\begin{aligned} \Phi_{\text{in}}(\mathbf{g}) &= \Phi_0|_{\mathcal{V}_{\text{in}}} + K_{\text{in}}\mathbf{g}, \quad K_{\text{in}} := \mathbf{P}_g(\mathcal{V}_{\text{in}}, :), \\ \Phi_{\text{out}}(\mathbf{g}) &= \Phi_0|_{\mathcal{V}_{\text{out}}} + K_{\text{out}}\mathbf{g}, \quad K_{\text{out}} := \mathbf{P}_g(\mathcal{V}_{\text{out}}, :). \end{aligned} \quad (8)$$

We prove a few straightforward results which are critical to justify the above calculations, in-particular the invertability of $\mathbf{L}_{ii}$. Towards this end, let $\mathbf{1}_{\mathcal{G}_i} \in \mathbb{R}^{n_V}$ denotes the indicator vector of the i -th connected component $\mathcal{G}_i$ of the graph, defined by

$$(\mathbf{1}_{\mathcal{G}_i})_v = \begin{cases} 1, & \text{if vertex } v \in \mathcal{G}_i, \\ 0, & \text{otherwise.} \end{cases}$$

Each $\mathbf{1}_{\mathcal{G}_i}$ is constant on its component and vanishes elsewhere, so that the nullspace of the weighted Laplacian $\mathbf{L} = \mathbf{B}^\top\mathbf{CB}$ is precisely $\ker(\mathbf{L}) = \text{span}\{\mathbf{1}_{\mathcal{G}_1}, \dots, \mathbf{1}_{\mathcal{G}_K}\}$ which we characterize next.

Lemma 1 (Weighted Laplacian structure and nullspace). *Let $\mathbf{B} \in \mathbb{R}^{n_E \times n_V}$ be the edge-by-node incidence as given in (1). Let $\mathbf{C} = \text{diag}(c_e) \succ 0$ with edge conductances $c_e > 0$, and define $\mathbf{L}$ as in (2). Then:*

- (a) $\mathbf{L}$ is symmetric positive semidefinite (SPSD).
- (b) If the graph has K connected components $\mathcal{G}_1, \dots, \mathcal{G}_K$, then

$$\ker(\mathbf{L}) = \text{span}\{\mathbf{1}_{\mathcal{G}_1}, \dots, \mathbf{1}_{\mathcal{G}_K}\},$$

i.e., the kernel is spanned by the indicator vectors of the components (constant on each component, zero elsewhere).

Proof. (a) *SPSD.* Symmetry is immediate: $\mathbf{L}^\top = (\mathbf{B}^\top\mathbf{CB})^\top = \mathbf{B}^\top\mathbf{CB} = \mathbf{L}$ since $\mathbf{C}$ is diagonal. For any $\mathbf{u} \in \mathbb{R}^{n_V}$,

$$\mathbf{u}^\top\mathbf{L}\mathbf{u} = \mathbf{u}^\top\mathbf{B}^\top\mathbf{CB}\mathbf{u} = (\mathbf{B}\mathbf{u})^\top\mathbf{C}(\mathbf{B}\mathbf{u}) = \sum_{e \in \mathcal{E}} c_e ((\mathbf{B}\mathbf{u})_e)^2 \geq 0,$$

so $\mathbf{L}$ is positive semidefinite.

- (b) *Kernel characterization.* We have $\mathbf{u} \in \ker(\mathbf{L})$ iff $\mathbf{L}\mathbf{u} = \mathbf{0}$, i.e.,

$$0 = \mathbf{u}^\top\mathbf{L}\mathbf{u} = \sum_{e \in \mathcal{E}} c_e ((\mathbf{B}\mathbf{u})_e)^2,$$

which, since $c_e > 0$, is equivalent to $(\mathbf{B}\mathbf{u})_e = 0$ for all edges e . But $(\mathbf{B}\mathbf{u})_e = \mathbf{u}_{h(e)} - \mathbf{u}_{t(e)}$ by construction, so $\mathbf{u}_{h(e)} = \mathbf{u}_{t(e)}$ for every edge. Therefore $\mathbf{u}$ is *constant on each connected component*: along any path inside a component, adjacent nodes must have equal values. Conversely, any vector that is constant on each component satisfies $\mathbf{B}\mathbf{u} = \mathbf{0}$, hence $\mathbf{L}\mathbf{u} = \mathbf{0}$. The set of such vectors is precisely $\text{span}\{\mathbf{1}_{\mathcal{G}_1}, \dots, \mathbf{1}_{\mathcal{G}_K}\}$. $\square$

The next result discusses the invertability of $\mathbf{L}_{ii}$ and states the affine dependence of $\mathbf{u}$, $\mathbf{q}$, and Φ on $\mathbf{g}$.

Theorem 2 (Forward solve). *If each connected component of $\mathcal{G}$ contains at least one Dirichlet node (fixed or controlled), then the interior block*

$$\mathbf{L}_{ii} = \mathbf{E}_{\text{int}}^\top \mathbf{L} \mathbf{E}_{\text{int}}$$

is symmetric positive definite and the interior system has a unique solution. Consequently $\mathbf{u}(\mathbf{g})$, $\mathbf{q}(\mathbf{g})$, and $\Phi(\mathbf{g})$ depend affinely on $\mathbf{g}$.

Proof. By Lemma 1, $\ker(\mathbf{L})$ is spanned by componentwise constants. Consider $x \in \mathbb{R}^{|\mathcal{V}_{\text{int}}|}$ with $\mathbf{L}_{ii}x = \mathbf{0}$. Let $\tilde{x} := \mathbf{E}_{\text{int}}x$ be the vector on all nodes that equals x on $\mathcal{V}_{\text{int}}$ and zero on Dirichlet nodes $D := \mathcal{V}_{\text{fix}} \cup \mathcal{V}_{\text{ctrl}}$. Then

$$0 = x^\top \mathbf{L}_{ii}x = (\mathbf{E}_{\text{int}}x)^\top \mathbf{L}(\mathbf{E}_{\text{int}}x) = \sum_{e \in \mathcal{E}} c_e ((\mathbf{B} \mathbf{E}_{\text{int}}x)_e)^2,$$

so $\mathbf{B} \mathbf{E}_{\text{int}}x = \mathbf{0}$ and $\mathbf{E}_{\text{int}}x$ is constant on each connected component. But on any component that contains at least one Dirichlet node, $\mathbf{E}_{\text{int}}x$ is zero at that Dirichlet node by construction (since $\mathbf{E}_{\text{int}}$ places zeros on D). The only constant function that is zero at some node is the zero constant, hence $\mathbf{E}_{\text{int}}x = \mathbf{0}$ on that component. Because *every* component has at least one Dirichlet node by assumption, $\mathbf{E}_{\text{int}}x = \mathbf{0}$ globally, so $x = \mathbf{0}$. Thus $\mathbf{L}_{ii}$ is injective; being symmetric PSD, it is positive definite. Uniqueness and affinity of $(\mathbf{u}, \mathbf{q}, \Phi)$ in $\mathbf{g}$ follow (6), (7), and (8). $\square$

Corollary 3 (Necessity of anchoring each component). *If some connected component contains no Dirichlet node (neither fixed nor controlled), then the interior block $\mathbf{L}_{ii}$ is singular. The forward problem on that component is only determined up to an additive constant, and the overall state is not unique.*

Proof. If a component has no Dirichlet node, the corresponding constant vector lies in $\ker(\mathbf{L})$ and has support entirely on that component. Its restriction to $\mathcal{V}_{\text{int}}$ is nonzero and lies in $\ker(\mathbf{L}_{ii})$, so $\mathbf{L}_{ii}$ is not invertible. $\square$

Remark 4 (Practical interpretation). *Dirichlet nodes (fixed or controlled) anchor the potential on each connected component by removing the constant-mode nullspace of $\mathbf{L}$. Without at least one anchor per component, the “DC offset” is free (physically: only gradients are determined), making the linear system underdetermined and potentially causing numerical issues or unbounded directions in optimization.*

5 Objective: maximize outward boundary flux

Outward means $\Phi_{\text{in}} \leq 0$ at inflows (out of boundary into interior) and $\Phi_{\text{out}} \geq 0$ at outflows. We maximize the net outward boundary flux

$$J(\mathbf{g}) = -\mathbf{1}^\top \Phi_{\text{in}}(\mathbf{g}) + \mathbf{1}^\top \Phi_{\text{out}}(\mathbf{g}) = \text{const} + \underbrace{\left(\begin{bmatrix} K_{\text{in}} \\ K_{\text{out}} \end{bmatrix}^\top \begin{bmatrix} -1 \\ 1 \end{bmatrix} \right)}_{=: \mathbf{c}^\star} \mathbf{g}. \quad (9)$$

Since typical optimization solvers, e.g., `linprog` minimizes, we solve the following minimization problem:

$$\min_{\mathbf{g}} -(\mathbf{c}^\star)^\top \mathbf{g}.$$

We are still not done specifying the optimization problem fully.

6 No-backflow constraints

Edge-level (per-edge) sign constraints. For each boundary endpoint v and adjacent edge $e = (\text{tail} \rightarrow \text{head})$, outwardness implies:

$$\begin{aligned} v \in \mathcal{V}_{\text{in}} : \quad & t(e) = v \Rightarrow q_e \geq -\varepsilon, \quad h(e) = v \Rightarrow q_e \leq +\varepsilon, \\ v \in \mathcal{V}_{\text{out}} : \quad & h(e) = v \Rightarrow q_e \geq -\varepsilon, \quad t(e) = v \Rightarrow q_e \leq +\varepsilon. \end{aligned}$$

Collect these as $S\mathbf{q} \geq -\varepsilon$ with selector $S \in \mathbb{R}^{m \times n_E}$ (each row picks $\pm q_e$). With $\mathbf{q}(\mathbf{g}) = \mathbf{q}_0 + \mathbf{Q}_g \mathbf{g}$,

$$S(\mathbf{q}_0 + \mathbf{Q}_g \mathbf{g}) \geq -\varepsilon \iff (-S\mathbf{Q}_g) \mathbf{g} \leq S\mathbf{q}_0 + \varepsilon.$$

Here we denote

$$A_{\text{edge}} := -S\mathbf{Q}_g, \quad b_{\text{edge}} := S\mathbf{q}_0 + \varepsilon, \tag{10}$$

so that the edge sign constraints compactly read

$$A_{\text{edge}} \mathbf{g} \leq b_{\text{edge}}.$$

Each row of A_{edge} corresponds to one boundary edge condition, with the sign pattern determined by the selector matrix S . Edges between same-type boundary nodes (Vin–Vin or Vout–Vout) receive two opposite constraints, forcing $q_e \simeq 0$ (no boundary shunt).

The next result states the impact of slack ε on the total flux. The shunt case will be discussed in the result following this result.

Lemma 5 (Edge rules imply node sign constraints up to slack). *Let v be a boundary node with degree $\deg(v)$ (number of incident edges). Impose the edge-level outwardness rules with slack $\varepsilon \geq 0$:*

$$\begin{aligned} v \in \mathcal{V}_{\text{in}} : \quad & t(e) = v \Rightarrow q_e \geq -\varepsilon, \quad h(e) = v \Rightarrow q_e \leq +\varepsilon, \\ v \in \mathcal{V}_{\text{out}} : \quad & h(e) = v \Rightarrow q_e \geq -\varepsilon, \quad t(e) = v \Rightarrow q_e \leq +\varepsilon. \end{aligned}$$

Then the nodal flux balance satisfies

$$\begin{aligned} v \in \mathcal{V}_{\text{in}} : \quad & \Phi_v \leq \varepsilon \deg(v), \\ v \in \mathcal{V}_{\text{out}} : \quad & \Phi_v \geq -\varepsilon \deg(v). \end{aligned}$$

In particular, for $\varepsilon = 0$ the node-level sign constraints $\Phi_v \leq 0$ on $\mathcal{V}_{\text{in}}$ and $\Phi_v \geq 0$ on $\mathcal{V}_{\text{out}}$ follow from the edge rules.

Proof. Recall $\Phi_v = \sum_{h(e)=v} q_e - \sum_{t(e)=v} q_e$.

If $v \in \mathcal{V}_{\text{in}}$, each incident edge with $h(e) = v$ obeys $q_e \leq +\varepsilon$ and each with $t(e) = v$ obeys $q_e \geq -\varepsilon$, hence

$$\Phi_v = \sum_{h(e)=v} q_e - \sum_{t(e)=v} q_e \leq \#\{h(e)=v\} \varepsilon - (-\#\{t(e)=v\} \varepsilon) = \varepsilon \deg(v).$$

If $v \in \mathcal{V}_{\text{out}}$, then $h(e) = v \Rightarrow q_e \geq -\varepsilon$ and $t(e) = v \Rightarrow q_e \leq +\varepsilon$, so

$$\Phi_v \geq -\#\{h(e)=v\} \varepsilon - \#\{t(e)=v\} \varepsilon = -\varepsilon \deg(v).$$

Setting $\varepsilon = 0$ yields $\Phi_v \leq 0$ on $\mathcal{V}_{\text{in}}$ and $\Phi_v \geq 0$ on $\mathcal{V}_{\text{out}}$. □

Proposition 6 (Boundary–boundary edges are blocked). *If an edge e connects two boundary nodes of the same type, i.e., $t(e), h(e) \in \mathcal{V}_{\text{in}}$ or $t(e), h(e) \in \mathcal{V}_{\text{out}}$, then the edge rules imply*

$$\boxed{|q_e| \leq \varepsilon.}$$

Thus for $\varepsilon = 0$ such edges carry no flux (no boundary shunting).

Proof. For $t(e), h(e) \in \mathcal{V}_{\text{in}}$: the tail rule gives $q_e \geq -\varepsilon$ and the head rule gives $q_e \leq +\varepsilon$, hence $|q_e| \leq \varepsilon$. The Vout–Vout case is identical. $\square$

Remark 7 (Practical takeaway). *Lemma 5 shows that adding edge-level rules with a small slack ε automatically enforces the desired node-level sign constraints up to $O(\varepsilon)$, while Proposition 6 prevents artificial “short-circuits” along boundary–boundary edges. In implementation, these constraints appear as the linear inequality*

$$A_{\text{edge}} \mathbf{g} \leq b_{\text{edge}}, \quad A_{\text{edge}} = -S\mathbf{Q}_g, \quad b_{\text{edge}} = S\mathbf{q}_0 + \varepsilon,$$

where each row of S encodes the sign pattern for one boundary edge. Choosing ε at or slightly above the solver tolerance (e.g. 10^{-6} in physical units) effectively removes cosmetic backflow without altering the physical solution.

7 Flux caps from a gradient bound

A gradient cap $|u'| \leq \phi_{\text{max}}$ on edge e implies

$$|q_e| = k_e |u'| \leq \phi_{\text{max}} k_e = \phi_{\text{max}} c_e L_e.$$

If raw conductivities k_e are provided (as in data), take $q_{\text{max},e} = \phi_{\text{max}} k_e$ (equivalently $\phi_{\text{max}} c_e L_e$). In matrix form:

$$-\mathbf{q}_{\text{max}} \leq \mathbf{q}_0 + \mathbf{Q}_g \mathbf{g} \leq \mathbf{q}_{\text{max}} \iff \begin{bmatrix} \mathbf{Q}_g \\ -\mathbf{Q}_g \end{bmatrix} \mathbf{g} \leq \begin{bmatrix} \mathbf{q}_{\text{max}} - \mathbf{q}_0 \\ \mathbf{q}_{\text{max}} + \mathbf{q}_0 \end{bmatrix}.$$

For compactness, we define

$$A_{\text{cap}} := \begin{bmatrix} \mathbf{Q}_g \\ -\mathbf{Q}_g \end{bmatrix}, \quad b_{\text{cap}} := \begin{bmatrix} \mathbf{q}_{\text{max}} - \mathbf{q}_0 \\ \mathbf{q}_{\text{max}} + \mathbf{q}_0 \end{bmatrix},$$

so that the flux-cap inequalities take the standard form

$$A_{\text{cap}} \mathbf{g} \leq b_{\text{cap}}.$$

8 Linear program

Collecting all components, the optimization problem reads

$$\begin{aligned} \min_{\mathbf{g} \in \mathbb{R}^{|\mathcal{V}_{\text{ctrl}}|}} \quad & J(\mathbf{g}) := -(\mathbf{c}^*)^\top \mathbf{g}, \\ \text{s.t.} \quad & A_{\text{cap}} \mathbf{g} \leq b_{\text{cap}}, \\ & A_{\text{edge}} \mathbf{g} \leq b_{\text{edge}} \quad (\text{edge no-backflow with slack } \varepsilon), \\ & \mathbf{g}_{\text{min}} \leq \mathbf{g} \leq \mathbf{g}_{\text{max}} \quad (\text{optional box bounds}). \end{aligned} \tag{11}$$

Here all constraints are affine, so the feasible set is an intersection of finitely many halfspaces and slabs. To prove this result, we assume that the objective function is bounded below.

Theorem 8 (Convexity and existence of an optimum). *The feasible region of (11) is a convex polyhedron, and the objective is linear; therefore the problem is convex. If the feasible set is nonempty and the objective is bounded below, a global minimizer exists.*

Proof. Let the decision variable be $\mathbf{g} \in \mathbb{R}^n$ and write the constraints compactly as

$$A\mathbf{g} \leq b, \quad \mathbf{g}_{\min} \leq \mathbf{g} \leq \mathbf{g}_{\max}.$$

The feasible set

$$P := \{\mathbf{g} : A\mathbf{g} \leq b, \mathbf{g}_{\min} \leq \mathbf{g} \leq \mathbf{g}_{\max}\}$$

is an intersection of finitely many closed halfspaces (and possibly slabs, i.e., intersections of two halfspaces); hence P is a closed convex polyhedron [11]. Moreover, $J(\mathbf{g}) = \mathbf{c}^\top \mathbf{g}$ is linear and therefore convex.

Assume $P \neq \emptyset$ and $\inf_{\mathbf{g} \in P} \mathbf{c}^\top \mathbf{g} > -\infty$. Define the recession cone

$$\text{rec}(P) = \left\{ \mathbf{d} : A\mathbf{d} \leq 0, \begin{array}{l} d_i \geq 0 \text{ if } (g_{\min})_i > -\infty \text{ and } (g_{\max})_i = +\infty, \\ d_i \leq 0 \text{ if } (g_{\max})_i < +\infty \text{ and } (g_{\min})_i = -\infty, \\ d_i = 0 \text{ if both } (g_{\min})_i, (g_{\max})_i \text{ are finite} \end{array} \right\},$$

with no sign restriction when both bounds are infinite. Then the standard equivalence holds:

$$\mathbf{d} \in \text{rec}(P) \iff \mathbf{g} + \tau \mathbf{d} \in P \text{ for every } \mathbf{g} \in P \text{ and every } \tau \geq 0.$$

($\Rightarrow$) From $A\mathbf{d} \leq 0$ and $A\mathbf{g} \leq b$, $A(\mathbf{g} + \tau \mathbf{d}) \leq b$ for all $\tau \geq 0$. The coordinatewise sign/zero rules ensure the bounds are preserved. ($\Leftarrow$) If $\mathbf{g} + \tau \mathbf{d} \in P$ for all $\tau \geq 0$, then $A\mathbf{g} + \tau A\mathbf{d} \leq b$ implies $A\mathbf{d} \leq 0$ by dividing by τ and letting $\tau \rightarrow \infty$; the same “for all τ ” argument forces the stated sign/zero conditions on $\mathbf{d}$.

If there existed $\mathbf{d} \in \text{rec}(P)$ with $\mathbf{c}^\top \mathbf{d} < 0$, then $J(\mathbf{g} + \tau \mathbf{d}) = \mathbf{c}^\top \mathbf{g} + \tau \mathbf{c}^\top \mathbf{d} \rightarrow -\infty$, contradicting boundedness from below. Hence

$$\mathbf{c}^\top \mathbf{d} \geq 0 \quad \text{for all } \mathbf{d} \in \text{rec}(P).$$

By the Minkowski–Weyl representation (e.g., [2, Cor. 2.7]), there exist finite sets V (points) and R (rays) such that $P = \text{conv}(V) + \text{cone}(R)$ and $\text{rec}(P) = \text{cone}(R)$. Any $\mathbf{x} \in P$ can be written $\mathbf{x} = \sum_i \lambda_i v^i + \sum_j \mu_j r^j$ with $\lambda_i \geq 0$, $\sum_i \lambda_i = 1$, $\mu_j \geq 0$. Then

$$\mathbf{c}^\top \mathbf{x} = \sum_i \lambda_i \mathbf{c}^\top v^i + \sum_j \mu_j \mathbf{c}^\top r^j \geq \sum_i \lambda_i \mathbf{c}^\top v^i \geq \inf_{\mathbf{v} \in \text{conv}(V)} \mathbf{c}^\top \mathbf{v},$$

since $\mathbf{c}^\top r^j \geq 0$ for all rays. The reverse inequality holds by taking $\mu \equiv 0$, so

$$\inf_{\mathbf{x} \in P} \mathbf{c}^\top \mathbf{x} = \inf_{\mathbf{v} \in \text{conv}(V)} \mathbf{c}^\top \mathbf{v}.$$

If $V \neq \emptyset$, the feasible region P has at least one vertex. Since P is a closed polyhedron and the objective function $\mathbf{c}^\top \mathbf{g}$ is a continuous linear function, the minimum value is guaranteed to exist due to the boundedness-from-below assumption. The optimal value is attained at one of the vertices of the feasible region, a consequence of the fundamental theorem of linear programming. $\inf_{\mathbf{x} \in P} \mathbf{c}^\top \mathbf{x} = \min_{\mathbf{v} \in V} \mathbf{c}^\top \mathbf{v}$. The convex hull $\text{conv}(V)$ is a compact set (a bounded polytope) and the objective function is linear, so the minimum is attained at an extreme point (a vertex).

If $V = \emptyset$, the polyhedron P contains no vertices (extreme points). Since $P \neq \emptyset$, this implies that P is an unbounded polyhedron that contains a line. The existence of a finite optimal value, $\inf_{\mathbf{g} \in P} \mathbf{c}^\top \mathbf{g} = \alpha > -\infty$, guarantees that the optimal solution set is non-empty and closed. This optimal set is a face of the polyhedron P , defined by $P \cap \{\mathbf{g} : \mathbf{c}^\top \mathbf{g} = \alpha\}$. The minimum is therefore attained on this face, even though no vertices exist. $\square$

Next, we prove a general result, which shows that for $\mathbf{d} \in \text{rec}(P)$ with $\mathbf{c}^\top \mathbf{d} \geq 0$ then $J(\mathbf{g})$ is bounded below. In order to apply this result to our setting of Theorem 8, we will provide sufficient conditions in Proposition 12 which guarantees that $\mathbf{c}^\top \mathbf{d} \geq 0$.

We begin by stating the well-known Hoffman bound [7] and [8, pp. 175].

Lemma 9 (Hoffman bound for polyhedra). *Let $\mathcal{F} = \{x \in \mathbb{R}^n : Mx \leq q\}$ be nonempty, where $M \in \mathbb{R}^{p \times n}$, $q \in \mathbb{R}^p$. For any norm $\|\cdot\|$ on $\mathbb{R}^n$, there exists a constant $H = H(M, \|\cdot\|) \in (0, \infty)$ such that for every $x \in \mathbb{R}^n$,*

$$\text{dist}(x, \mathcal{F}) \leq H \|(Mx - q)_+\|_\infty,$$

where $(\cdot)_+$ is the componentwise positive part.

Lemma 10 (Linear growth of projection distances onto LP sublevel sets). *Let $P = \{g : Ag \leq b\}$ be nonempty and $c \in \mathbb{R}^n$. Fix $g^0 \in P$, and for each $k \in \mathbb{N}$ let*

$$P_{-k} := \{g \in P : c^\top g \leq -k\}$$

be nonempty. Choose $g_k \in \arg \min\{\|g - g^0\| : g \in P_{-k}\}$. Then there exists $H \in (0, \infty)$, depending only on (A, b, c, g^0) and the chosen norm, such that

$$\|g_k - g^0\| \leq H(1 + k) \quad \text{for all } k.$$

Proof. Write the augmented system for P_{-k} as

$$Mx \leq q_k \iff \begin{bmatrix} A \\ c^\top \end{bmatrix} x \leq \begin{bmatrix} b \\ -k \end{bmatrix}.$$

Let H be a Hoffman constant for the matrix $M = [A; c^\top]$ with respect to the chosen norm. Applying Hoffman's bound (Lemma 9) at $x = g^0$ gives

$$\text{dist}(g^0, P_{-k}) \leq H \|(Mg^0 - q_k)_+\|_\infty = H \left\| \left((Ag^0 - b)_+, (c^\top g^0 - (-k))_+ \right) \right\|_\infty.$$

Since $g^0 \in P$, we have $(Ag^0 - b)_+ = 0$, so

$$\text{dist}(g^0, P_{-k}) \leq H (c^\top g^0 + k)_+ = \begin{cases} 0, & c^\top g^0 \leq -k, \\ H(c^\top g^0 + k), & c^\top g^0 > -k. \end{cases}$$

By definition of g_k as a nearest point to P_{-k} ,

$$\|g_k - g^0\| = \text{dist}(g^0, P_{-k}) \leq H(c^\top g^0 + k)_+ \leq H(1 + |c^\top g^0|)(1 + k).$$

Renaming the constant $H \leftarrow H(1 + |c^\top g^0|)$ yields the claimed linear bound $\|g_k - g^0\| \leq H(1 + k)$ for all k . $\square$

Lemma 11 (Boundedness of a linear functional on a polyhedron). *Let $P := \{g \in \mathbb{R}^n : Ag \leq b\}$ be a nonempty polyhedron and $\text{rec}(P) := \{d \in \mathbb{R}^n : Ad \leq 0\}$ its recession cone. If every $d \in \text{rec}(P)$ satisfies $c^\top d \geq 0$, then the linear functional $J(g) = c^\top g$ is bounded below on P .*

Proof. Assume for contradiction that J is unbounded below on P . Then for each integer $k \geq 1$ the sublevel set $P_{-k} := \{g \in P : c^\top g \leq -k\}$ is nonempty. Fix $g^0 \in P$ and, for each k , choose $g_k \in P_{-k}$ minimizing $\|g - g^0\|$. Define the positive scalars

$$t_k := -(c^\top g_k - c^\top g^0) \xrightarrow{k \rightarrow \infty} +\infty, \quad \hat{d}_k := \frac{g_k - g^0}{t_k}.$$

By Lemma 10, $\|g_k - g^0\| \leq H(1 + k)$, hence

$$\|\hat{d}_k\| = \frac{\|g_k - g^0\|}{t_k} \leq \frac{H(1 + k)}{k - c^\top g^0} \leq 2H \quad \text{for all large } k,$$

so $(\hat{d}_k)$ is bounded. By Bolzano–Weierstrass, extract a convergent subsequence $\hat{d}_{k_j} \rightarrow \hat{d}$.

Since $Ag_k \leq b$ and $Ag^0 \leq b$,

$$A\hat{d}_k = \frac{Ag_k - Ag^0}{t_k} \leq \frac{b - Ag^0}{t_k} \xrightarrow{k \rightarrow \infty} 0,$$

which implies $A\hat{d} \leq 0$; hence $\hat{d} \in \text{rec}(P)$. Moreover,

$$c^\top \hat{d}_k = \frac{c^\top g_k - c^\top g^0}{t_k} = -1 \quad \Rightarrow \quad c^\top \hat{d} = -1 < 0.$$

Thus there exists a nonzero recession direction $\hat{d}$ with $c^\top \hat{d} < 0$, contradicting the hypothesis. Therefore J is bounded below on P . $\square$

Proposition 12 (When boundedness (or bounded-below) is automatic). *Let $P = \{g \in \mathbb{R}^n : Ag \leq b, g_{\min} \leq g \leq g_{\max}\}$, and suppose the flux map is affine: $q(g) = q_0 + Q_g g$. Then:*

- (i) **Explicit box bounds.** *If all components have finite bounds $g_{\min} \leq g \leq g_{\max}$, then P is a compact polytope (closed and bounded).*
- (ii) **Flux caps with full column rank.** *If two-sided flux caps $|q(g)| \leq q_{\max}$ hold and Q_g has full column rank on the control subspace, then $\|Q_g g\|_\infty \leq \|q_{\max} - q_0\|_\infty$ implies $\|g\|_2 \leq C$ for some $C < \infty$; hence P is bounded.*
- (iii) **Rank-deficient case with sign constraints.** *If Q_g is rank-deficient, but edge sign constraints eliminate every recession direction $d \in \text{rec}(P)$ with $c^\top d < 0$ (no “descent” rays), then the linear objective is bounded below on P (by Lemma 11). (Note: P itself can remain unbounded due to neutral rays with $c^\top d = 0$.)*

Proof. (i) The intersection with a finite hyperrectangle is closed and bounded, hence compact.

(ii) From $-q_{\max} \leq q_0 + Q_g g \leq q_{\max}$ we obtain $-q_{\max} - q_0 \leq Q_g g \leq q_{\max} - q_0$, i.e. $\|Q_g g\|_\infty \leq \|q_{\max} - q_0\|_\infty$. Let $Q_g \in \mathbb{R}^{m \times n}$ and assume it has full column rank, so $\sigma_{\min}(Q_g) > 0$. Then

$$\|g\|_2 \leq \frac{\|Q_g g\|_2}{\sigma_{\min}(Q_g)} \leq \frac{\sqrt{m} \|Q_g g\|_\infty}{\sigma_{\min}(Q_g)} \leq \frac{\sqrt{m} \|q_{\max} - q_0\|_\infty}{\sigma_{\min}(Q_g)}.$$

Thus g is norm-bounded and P is bounded.

(iii) Without flux caps, $\text{rec}(P)$ may be nontrivial. If the sign constraints (together with $Ag \leq b$ and the bound structure) exclude every $d \in \text{rec}(P)$ with $c^\top d < 0$, then by Lemma 11 the objective is bounded below on P . Neutral rays with $c^\top d = 0$ may remain, so P need not be bounded. $\square$

9 Numerical experiments

The main driver is Algorithm 1 with subroutines given in Algorithms 2–8. In all examples, we set $\phi_{\max} = 1$ and $\varepsilon = 0$.

We present two representative networks of contrasting structure and connectivity: a *typical large stadium in a major modern city*, which forms a single connected component with relatively uniform corridor widths, and a *complex street network emanating from a large, historical city center*, which decomposes into seven connected subgraphs with pronounced geometric heterogeneity. The same linear programming (LP) formulation applies to both, demonstrating that the proposed framework is robust to both dense single-component and fragmented multi-component geometries.

Gauge fixing for control invariance. When all Dirichlet values on a connected component are treated as free controls, adding a constant to every boundary value on that component leaves edge fluxes and nodal fluxes unchanged. The LP objective and constraints are therefore invariant to such rigid shifts, creating a null direction in control space. To remove this gauge freedom and improve conditioning, we impose one scalar gauge per connected component—either by fixing one Dirichlet node or enforcing a mean-zero constraint on its boundary controls. In the following examples, one node per component is fixed to a prescribed potential value.

Flux metrics. We visualize two complementary quantities:

- *Flux intensity*: the magnitude $|q_e|$, representing per-area flux density (diffusive strength per unit width);
- *Throughput*: the product $|q_e|A_e$, representing total transported quantity through an edge (aggregate capacity use).

Together, these measures distinguish between relative diffusion intensity and total corridor throughput.

The limit set for the flux corresponds to the empirically observed limit of pedestrian flux intensity, which is of the order of 1 person/meter/second for safe crowd walking conditions [6, 10]. Thus, the results obtained represent the upper limit of steady pedestrian flux that these networks are able to sustain in a safe manner.

9.1 Typical Large Stadium in a Modern City

We first demonstrate the LP framework on a realistic geometric network extracted via google maps for a large stadium located in a modern city. The network only considers a partial set of streets, no subway exits, and pedestrian motion only (no cars, buses or trains). The graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ encodes the concourse and corridor layout of the stadium and its immediate surroundings. Vertices correspond to corridor intersections, and edges to walkable links (avenues, streets, passages) with assigned cross-sectional areas. Conductances are scaled inversely with edge length to emulate diffusive pedestrian flow or steady-state heat conduction along corridors of varying width.

9.1.1 Geometry and cross-sectional area distribution

Figure 2 shows the reconstructed network geometry colored by effective edge area. Blue edges correspond to narrow passages, while yellow–red edges indicate wider corridors or open concourse sections. The lower panel zooms into the interior, revealing the main radial and circumferential corridors that govern overall circulation.




Figure 2: Stadium network colored by corridor cross-sectional area. Blue: narrow links; red: broad concourse regions. The zoomed view highlights the dense interior structure near the stadium core.

9.1.2 Temperature (potential) distribution

The steady-state potential field obtained from the optimized control is shown in Figure 3. Higher potentials (red–yellow) occur near entrance boundaries, while cooler regions (green–blue) correspond to exits or high-dissipation paths. We fix the potential at one interior node (113) to a reference value $u_{113} = 10$ to remove the additive nullspace of the steady-state diffusion operator. This normalization pins the global potential field without affecting gradients or flux balance.

9.1.3 Flux intensity and throughput

Figures 4–5 show the computed flux intensity $|q_e|$ and throughput $|q_e|A_e$ across the network. High-intensity fluxes (yellow–red) appear along the primary north–south and east–west corridors, indicating strong directional transport between major entry and exit points. Flux directions remain consistent with the stadium’s natural pedestrian flow topology. After multiplying by area, the



Figure 3: Optimized potential field u for the stadium network. Hotter colors (red–yellow) mark entrances or high-pressure zones; cooler tones (blue–green) mark exits and dissipative corridors. The reference potential was fixed at node 113 ($u_{113} = 10$).

throughput map emphasizes broad corridors that carry higher total flow.

9.1.4 Validation and observations

Figure 6 summarizes the overall network geometry. The LP formulation yields stable, physically consistent solutions even on complex real-world networks. No violation of the sign constraints was observed, and total inflow/outflow fluxes balanced to within 10^{-12} . The major flux paths align with the architectural concourse design, validating the model’s geometric sensitivity and predictive realism.

Quantity	Value
$\max(\Phi_{\text{in}})$ (should ≤ 0)	$-1.000 \times 10^{+1}$
$\min(\Phi_{\text{out}})$ (should ≥ 0)	-8.903×10^{-15}
Max boundary–edge sign violation (should ≤ 0)	8.90×10^{-15}
Global conservation $\sum_v \Phi_v$ (should ≈ 0)	4.44×10^{-14}
$\max_{v \in \mathcal{V}_{\text{int}}} \Phi_v $	3.936×10^{-13}
Amount entering at $\mathcal{V}_{\text{in}}$	$7.000000 \times 10^{+1}$
Amount leaving at $\mathcal{V}_{\text{out}}$	$7.000000 \times 10^{+1}$
In–out mismatch (should ≈ 0)	-1.236×10^{-12}
$\max_{\text{component}} \sum_{v \in \text{comp}} \Phi_v $	9.599×10^{-15}

Table 1: Stadium exit validation metrics. All diagnostics confirm sign correctness and flux conservation to within double-precision tolerance ($\sim 10^{-14}$).

Key remarks. The feasible set P is convex and bounded, ensuring LP solvability. The diffusion-like flux pattern shows that the optimized control redistributes flow to prevent local congestion. The




Figure 4: Flux intensity $|q_e|$ across the network. Top: global view of all corridors. Bottom: zoomed view near the stadium interior. Color represents per-area flow strength; arrows indicate flux direction.

area scaling reveals near-constant per-area flux intensity, confirming that the optimization respects both geometric and capacity constraints.

9.2 Complex Street Network Emanating from Large, Historical City Center

We next apply the same LP framework to a geometric network typical of those encountered around large, historical city centers. The graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ represents the pedestrian and structural corridors surrounding the historical center (whose network in turn may be very complex and chaotic). Unlike the previous stadium network, this geometry comprises seven disconnected components of varying scale. Each component was analyzed independently under identical flux-cap and sign constraints.




Figure 5: Flux throughput ($|q_e|A_e$) across the network. Top: global view. Bottom: zoom near the central concourse. Broader corridors support proportionally higher total flux.

9.2.1 Boundary and reference conditions

Dirichlet boundary conditions were imposed on inflow (entrance) and outflow (exit) nodes located at the periphery of each component. To remove the additive nullspace within each disconnected subgraph, one node per component was fixed to a prescribed potential:

$$\text{Nodes } \{10, 88, 437, 152, 441, 542, 550\}, \quad u_i = 10.$$

This ensures uniqueness of the steady-state potential on each connected component and eliminates the corresponding control-space gauge.

9.2.2 Temperature (potential) distribution

The resulting potential field is shown in Figure 7. High potentials (red–yellow) appear near inflow zones, while cooler regions (blue–green) correspond to outflow boundaries. Smooth gradients are observed across all seven components, reflecting physically consistent equilibrium states.




Figure 6: Overall stadium network geometry used in the simulation, overlaid on the planar embedding extracted




Figure 7: Optimized potential field u over the network. Each connected component exhibits smooth gradients between inflow and outflow boundaries. Fixed nodes ($u_i = 10$) provide a reference potential per component.

9.2.3 Flux intensity, direction, and throughput

Figure 8 shows the computed flux intensity $|q_e|$ over the full network, followed by directional and zoomed views (Figures 9–10). The LP enforcement of sign constraints yields consistent flux directions, while stronger magnitudes concentrate along primary corridors. Figure 11 presents the total throughput $|q_e|A_e$, highlighting wide corridors with greater transport capacity.




Figure 8: Flux intensity $|q_e|$ across the network. Color scale indicates per-area flux density (blue: weak, red: strong).

9.2.4 Validation and diagnostic metrics

All conservation and sign checks confirm the correctness of the LP solution. Table 2 summarizes the diagnostic results. Flux conservation holds to within 10^{-13} , and boundary sign violations are below 10^{-6} , demonstrating consistent performance across multiple components.

Quantity	Value
$\max(\Phi_{\text{in}})$ (should ≤ 0)	-1.998×10^{-15}
$\min(\Phi_{\text{out}})$ (should ≥ 0)	-3.492×10^{-6}
Max boundary-edge sign violation (should ≤ 0)	3.49×10^{-6}
Global conservation $\sum_v \Phi_v$ (should ≈ 0)	1.07×10^{-13}
$\max_{v \in \mathcal{V}_{\text{int}}} \Phi_v $	3.37×10^{-13}
Amount entering at $\mathcal{V}_{\text{in}}$	$2.729882 \times 10^{+2}$
Amount leaving at $\mathcal{V}_{\text{out}}$	$2.729882 \times 10^{+2}$
In-out mismatch (should ≈ 0)	4.547×10^{-13}
$\max_{\text{component}} \sum_{v \in \text{comp}} \Phi_v $	6.682×10^{-14}

Table 2: Network validation metrics. All diagnostics confirm flux conservation and sign correctness to within numerical tolerance.

Key remarks. The network decomposes into seven connected components, each internally flux-balanced. Independent gauge fixing ($u_i = 10$ per component) removes null directions in the control space, ensuring well-posedness. Flux balance and conservation hold to machine precision (10^{-13}), confirming numerical robustness.

These results demonstrate that the proposed framework extends naturally to multiply connected




Figure 9: Flux directions in two zoomed regions. Arrowheads indicate consistent directional flow along major corridors.

or fragmented networks while preserving physical and numerical consistency.

Conclusion and Outlook

We have presented a linear programming (LP) framework for steady-state diffusion and flux optimization on geometric networks. Starting from a discrete Laplacian model, we showed that boundary potentials can be treated as control variables that drive interior states and fluxes through an affine mapping. This reduction yields a finite-dimensional LP whose feasible set is polyhedral, whose convexity is guaranteed, and whose global solvability follows from geometric arguments on recession cones and the Minkowski–Weyl decomposition. Three classes of sufficient conditions—finite



Figure 10: Detailed flux orientation near junctions and intersections.




Figure 11: Flux throughput ($|q_e|A_e$) across the network. Wide corridors support proportionally higher total flux, consistent with capacity limits.

box bounds, two-sided flux caps with full-rank maps, and rank-deficient configurations stabilized by sign constraints—ensure boundedness automatically.

Numerical experiments on two real geometric networks, demonstrate the robustness and interpretability of the method. The LP formulation enforces flux conservation and sign correctness, handles disconnected components without additional regularization, and offers physically meaningful insights into crowd or transport behavior.

Beyond its immediate applications to diffusive transport, the proposed framework provides a foundational building block for *network-based digital twins*. The affine control-to-flux map enables:

- rapid re-optimization of boundary conditions under changing environments;
- direct integration into higher-level PDE-constrained or multi-scale optimization loops;
- uncertainty quantification through stochastic or robust LP formulations.

Several research directions follow naturally. A first extension concerns time-dependent diffusion or transient transport on networks, leading to dynamic LPs or model-predictive formulations. Second, coupling this graph-level LP with continuum PDEs at finer scales would enable hybrid *graph-PDE digital twins* for structures.

A Algorithms (pseudocode)

Algorithm 1 Mixed-boundary outward-flux maximization (high-level)

Require: Data file; user choice $(\mathcal{V}_{\text{fix}}, u_{\text{fix}})$; parameters $\phi_{\text{max}}, \varepsilon$

- 1: $(\mathbf{B}, \mathbf{C}, \mathbf{L}, \mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}}) \leftarrow \text{BUILDNETWORK}(\text{data})$ $\triangleright$ incidence, conductance, Laplacian
- 2: $(\mathcal{V}_{\text{ctrl}}, \mathcal{V}_{\text{int}}, \mathbf{E}_{\text{fix}}, \mathbf{E}_{\text{ctrl}}, \mathbf{E}_{\text{int}}) \leftarrow \text{PARTITIONANDEMBEDDINGS}(\mathcal{V}_{\text{fix}}, \mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}})$
- 3: $(\mathbf{u}_0, \mathbf{U}_g, \mathbf{q}_0, \mathbf{Q}_g, \mathbf{\Phi}_0, \mathbf{P}_g, K_{\text{in}}, K_{\text{out}}, \Phi_{0,\text{in}}, \Phi_{0,\text{out}}) \leftarrow \text{BUILDAFFINEMAPS}(\mathbf{L}, \mathbf{B}, \mathbf{C}, \mathbf{E}_{\text{fix}}, \mathbf{E}_{\text{ctrl}}, \mathbf{E}_{\text{int}}, u_{\text{fix}})$
- 4: $\mathbf{c}^* \leftarrow \text{BUILDOBJECTIVE}(K_{\text{in}}, K_{\text{out}})$ $\triangleright J = \text{const} + (\mathbf{c}^*)^\top \mathbf{g}$
- 5: $(A, b, \mathbf{g}_{\text{min}}, \mathbf{g}_{\text{max}}) \leftarrow \text{ASSEMBLECONSTRAINTS}(\mathbf{Q}_g, \mathbf{q}_0, \mathbf{P}_g, \mathbf{\Phi}_0, \mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}}, \phi_{\text{max}}, \varepsilon)$
- 6: $\mathbf{g}^* \leftarrow \text{SOLVELP}(-(\mathbf{c}^*), A, b, \mathbf{g}_{\text{min}}, \mathbf{g}_{\text{max}})$ $\triangleright$ linprog
- 7: $(\mathbf{u}^*, \mathbf{q}^*, \mathbf{\Phi}^*) \leftarrow \text{RECOVERSTATE}(\mathbf{u}_0, \mathbf{U}_g, \mathbf{q}_0, \mathbf{Q}_g, \mathbf{\Phi}_0, \mathbf{P}_g, \mathbf{g}^*)$
- 8: $\text{VERIFYANDREPORT}(\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}}, \varepsilon, \mathbf{q}^*, \mathbf{\Phi}^*)$ $\triangleright$ node signs, edge signs, conservation
- 9: **return** $(\mathbf{g}^*, \mathbf{u}^*, \mathbf{q}^*, \mathbf{\Phi}^*)$

Algorithm 2 BuildNetwork

Require: Data file with vertices, edges, raw k_e , lengths L_e , boundary sets $\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}}$

- 1: $c_e \leftarrow k_e/L_e$; $\mathbf{C} \leftarrow \text{diag}(c_e)$
- 2: Build edge-by-node incidence $\mathbf{B}$ with -1 at tail, $+1$ at head
- 3: $\mathbf{L} \leftarrow \mathbf{B}^\top \mathbf{C} \mathbf{B}$
- 4: **return** $(\mathbf{B}, \mathbf{C}, \mathbf{L}, \mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}})$

Algorithm 3 PartitionAndEmbeddings

Require: $\mathcal{V}_{\text{fix}}, \mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}}$; n_V (number of nodes)

- 1: $\mathcal{V}_{\text{ctrl}} \leftarrow (\mathcal{V}_{\text{in}} \cup \mathcal{V}_{\text{out}}) \setminus \mathcal{V}_{\text{fix}}$; $\mathcal{V}_{\text{int}} \leftarrow \{1:n_V\} \setminus (\mathcal{V}_{\text{fix}} \cup \mathcal{V}_{\text{ctrl}})$
- 2: Build embeddings $\mathbf{E}_{\text{fix}}, \mathbf{E}_{\text{ctrl}}, \mathbf{E}_{\text{int}}$ as sparse placement matrices
- 3: **return** $(\mathcal{V}_{\text{ctrl}}, \mathcal{V}_{\text{int}}, \mathbf{E}_{\text{fix}}, \mathbf{E}_{\text{ctrl}}, \mathbf{E}_{\text{int}})$

Algorithm 4 BuildAffineMaps

Require: $\mathbf{L}, \mathbf{B}, \mathbf{C}, \mathbf{E}_{\text{fix}}, \mathbf{E}_{\text{ctrl}}, \mathbf{E}_{\text{int}}, u_{\text{fix}}$

- 1: $\mathbf{L}_{ii} \leftarrow \mathbf{E}_{\text{int}}^\top \mathbf{L} \mathbf{E}_{\text{int}}; \mathbf{L}_{i,\text{fix}} \leftarrow \mathbf{E}_{\text{int}}^\top \mathbf{L} \mathbf{E}_{\text{fix}}; \mathbf{L}_{i,\text{ctrl}} \leftarrow \mathbf{E}_{\text{int}}^\top \mathbf{L} \mathbf{E}_{\text{ctrl}}$
 - 2: $A_0 \leftarrow -\mathbf{L}_{ii}^{-1} \mathbf{L}_{i,\text{fix}} u_{\text{fix}}; A_1 \leftarrow -\mathbf{L}_{ii}^{-1} \mathbf{L}_{i,\text{ctrl}}$
 - 3: $\mathbf{u}_0 \leftarrow \mathbf{E}_{\text{int}} A_0 + \mathbf{E}_{\text{fix}} u_{\text{fix}}; \mathbf{U}_g \leftarrow \mathbf{E}_{\text{int}} A_1 + \mathbf{E}_{\text{ctrl}}$
 - 4: $\mathbf{q}_0 \leftarrow -\mathbf{C} \mathbf{B} \mathbf{u}_0; \mathbf{Q}_g \leftarrow -\mathbf{C} \mathbf{B} \mathbf{U}_g$
 - 5: $\Phi_0 \leftarrow \mathbf{B}^\top \mathbf{q}_0; \mathbf{P}_g \leftarrow \mathbf{B}^\top \mathbf{Q}_g$
 - 6: $K_{\text{in}} \leftarrow \mathbf{P}_g(\mathcal{V}_{\text{in}}, :); K_{\text{out}} \leftarrow \mathbf{P}_g(\mathcal{V}_{\text{out}}, :)$
 - 7: $\Phi_{0,\text{in}} \leftarrow \Phi_0|_{\mathcal{V}_{\text{in}}}; \Phi_{0,\text{out}} \leftarrow \Phi_0|_{\mathcal{V}_{\text{out}}}$
 - 8: **return** $(\mathbf{u}_0, \mathbf{U}_g, \mathbf{q}_0, \mathbf{Q}_g, \Phi_0, \mathbf{P}_g, K_{\text{in}}, K_{\text{out}}, \Phi_{0,\text{in}}, \Phi_{0,\text{out}})$
-

Algorithm 5 BuildObjective

Require: $K_{\text{in}}, K_{\text{out}}$

- 1: $\mathbf{c}^* \leftarrow \begin{bmatrix} K_{\text{in}} \\ K_{\text{out}} \end{bmatrix}^\top \begin{bmatrix} -1 \\ 1 \end{bmatrix}$
 - 2: **return** $\mathbf{c}^*$
-

Algorithm 6 AssembleConstraints

Require: $\mathbf{Q}_g, \mathbf{q}_0, \mathbf{P}_g, \Phi_0, \mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}}, \phi_{\text{max}}, \varepsilon$

- 1: *Flux caps:* $\mathbf{q}_{\text{max}} \leftarrow \phi_{\text{max}} \vec{k}$
 - 2: $A_{\text{cap}} \leftarrow \begin{bmatrix} \mathbf{Q}_g \\ -\mathbf{Q}_g \end{bmatrix}; b_{\text{cap}} \leftarrow \begin{bmatrix} \mathbf{q}_{\text{max}} - \mathbf{q}_0 \\ \mathbf{q}_{\text{max}} + \mathbf{q}_0 \end{bmatrix}$
 - 3: *Edge no-backflow:* $(A_{\text{edge}}, b_{\text{edge}}) \leftarrow \text{EDGESIGNROWS}(\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}}, \mathbf{q}_0, \mathbf{Q}_g, \varepsilon)$
 - 4: $A \leftarrow \begin{bmatrix} A_{\text{cap}} \\ A_{\text{edge}} \end{bmatrix}; b \leftarrow \begin{bmatrix} b_{\text{cap}} \\ b_{\text{edge}} \end{bmatrix}$
 - 5: Choose bounds $\mathbf{g}_{\text{min}}, \mathbf{g}_{\text{max}}$ (optional)
 - 6: **return** $(A, b, \mathbf{g}_{\text{min}}, \mathbf{g}_{\text{max}})$
-

Algorithm 7 EdgeSignRows (build edge-level no-backflow)

Require: Graph orientation (tails/heads via $\mathbf{B}$ indices), sets $\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}}$, affine $(\mathbf{q}_0, \mathbf{Q}_g)$, slack ε

- 1: Build a sparse selector S with one row per boundary endpoint rule:
 - if $t \in \mathcal{V}_{\text{in}}$, add row picking $+q_e$; if $h \in \mathcal{V}_{\text{in}}$, add row picking $-q_e$;
 - if $h \in \mathcal{V}_{\text{out}}$, add row picking $+q_e$; if $t \in \mathcal{V}_{\text{out}}$, add row picking $-q_e$
 - 2: **return** $A_{\text{edge}} \leftarrow -S \mathbf{Q}_g; b_{\text{edge}} \leftarrow S \mathbf{q}_0 + \varepsilon$
-

Algorithm 8 RecoverState and VerifyAndReport

Require: $(\mathbf{u}_0, \mathbf{U}_g, \mathbf{q}_0, \mathbf{Q}_g, \Phi_0, \mathbf{P}_g, \mathbf{g}^*)$, sets $(\mathcal{V}_{\text{in}}, \mathcal{V}_{\text{out}})$, slack ε

- 1: **Recover:** $\mathbf{u}^* \leftarrow \mathbf{u}_0 + \mathbf{U}_g \mathbf{g}^*; \mathbf{q}^* \leftarrow \mathbf{q}_0 + \mathbf{Q}_g \mathbf{g}^*; \Phi^* \leftarrow \Phi_0 + \mathbf{P}_g \mathbf{g}^*$
 - 2: **Node signs:** ensure $\max(\Phi^*|_{\mathcal{V}_{\text{in}}}) \leq 0, \min(\Phi^*|_{\mathcal{V}_{\text{out}}}) \geq 0$
 - 3: **Edge signs:** rebuild S ; check $\min(S \mathbf{q}^*) \geq -\varepsilon$; list any violators
 - 4: **Conservation:** check $|\sum_v \Phi_v^*| \approx 0$
 - 5: **return** $(\mathbf{u}^*, \mathbf{q}^*, \Phi^*)$
-

References

- [1] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [2] Dimitris Bertsimas and John N Tsitsiklis. *Introduction to linear optimization*, volume 6. Athena scientific Belmont, MA, 1997.
- [3] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [4] Fan R. K. Chung. *Spectral Graph Theory*, volume 92 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, 1997.
- [5] Peter G. Doyle and J. Laurie Snell. *Random Walks and Electric Networks*. Mathematical Association of America, 2000. Original 1984 lecture notes.
- [6] John J Fruin. Pedestrian planning and design. Technical report, 1971.
- [7] Alan J. Hoffman. On approximate solutions of systems of linear inequalities. *J. Research Nat. Bur. Standards*, 49:263–265, 1952.
- [8] Mingyi Hong and Zhi-Quan Luo. On the linear convergence of the alternating direction method of multipliers. *Math. Program.*, 162(1-2):165–199, 2017.
- [9] Roger L. Hughes. A continuum theory for the flow of pedestrians. *Transportation Research Part B: Methodological*, 36(6):507–535, 2002.
- [10] Wsewolod M Predtetschenski and Anatolij I Milinski. Personenströme in gebäuden. 1971.
- [11] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [12] Martin Schmidt and Falk M Hante. Gas transport network optimization: Pde-constrained models. In *Encyclopedia of Optimization*, pages 1–7. Springer, 2023.
- [13] Alexander Schrijver. *Theory of linear and integer programming*. Wiley-Interscience Series in Discrete Mathematics. John Wiley & Sons, Ltd., Chichester, 1986. A Wiley-Interscience Publication.
- [14] Günter M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer, 1995.