



A social cost gradient centrality for transportation networks

Jonas Wassmer^{1,3,a} , Nils Antary⁴, Carsten Hartmann⁵, and Norbert Marwan^{1,2,3}

¹ Potsdam Institute for Climate Impact Research (PIK), Member of the Leibniz Association, Potsdam, Germany

² Institute of Physics and Astronomy, University of Potsdam, Potsdam, Germany

³ Institute of Environmental Science and Geography, University of Potsdam, Potsdam, Germany

⁴ Meteorology, Department of Earth System Sciences, University of Hamburg, Hamburg, Germany

⁵ Institute of Energy and Climate Research – Energy System Engineering (ICE-1), Forschungszentrum Jülich, Jülich, Germany

Received 20 February 2026 / Accepted 17 March 2026 / Published online 12 April 2026
© The Author(s) 2026

Abstract Urban road networks are prototypical complex systems in which large numbers of individual agents interact through shared infrastructure, giving rise to collective traffic equilibria. These equilibria emerge from decentralized route choice decisions and exhibit a strong sensitivity to local changes in network properties. Even small modifications to the capacity or free-flow travel time of a single link can trigger large-scale reconfigurations of traffic flows, as illustrated by the Braess paradox. Conventional measures of link importance often overlook this systemic sensitivity and its implications for network efficiency. In this study, we introduce a social cost-based centrality measure that quantifies the marginal impact of link-level free-flow travel time perturbations on total social cost under Wardrop equilibrium. The measure is derived analytically from the linear formulation of the traffic assignment problem, enabling efficient and interpretable computation of local sensitivities on fixed active-support subgraphs (and thus piecewise linear behavior across regime changes). We demonstrate the approach in synthetic and real-world urban networks, revealing structurally critical links that exert disproportionate influence on travel costs. Beyond improving network efficiency, this framework supports targeted interventions that can reduce unnecessary travel, lower emissions, and contribute to the transition toward more sustainable and livable cities.

1 Introduction

Traffic patterns in cities emerge from many local route choices made on shared road infrastructure rather than from centralized coordination. This decentralized interaction between heterogeneous drivers produces macroscopic states such as congestion patterns and traffic equilibria, making road systems a paradigmatic setting for self-organization in complex systems [1–5].

The behavior of such systems is governed by the emergence of macroscopic order from microscopic interactions. Through cooperative behavior, simple components can give rise to pattern formation, phase transitions, and system-wide stability [6, 7]. In traffic networks, local route choices aggregate into collective flow patterns and systemic vulnerabilities. The importance of individual roads is therefore not determined by their physical properties alone, but by their role in shaping the global organization of flows. Identifying which local components exert the greatest influence on system-wide behavior is essential for understanding, predicting, and managing these self-organizing dynamics.

A defining feature of these networks is their sensitivity to local perturbations, consistent with the principle of synergetic systems near instability points as described by Haken [6]. Even small interventions, such as modifying the capacity or free-flow travel time of a single link, can trigger large-scale reconfigurations of traffic flows. The Braess paradox illustrates this sensitivity: adding a road can paradoxically increase total travel times by shifting route choices [8–10]. This demonstrates that transport network performance must be understood as a systemic property, rather than a simple sum of link-level characteristics.

^a e-mail: jonas.wassmer@pik-potsdam.de (corresponding author)

Wardrop [11] formulated two fundamental principles of traffic assignment. The *user equilibrium* describes the state in which no driver can reduce their travel time by unilaterally changing routes, while the *system optimum* minimizes total travel time across all users. The gap between these two states reflects the inefficiency of decentralized decision-making. Understanding how local infrastructure changes influence these efficiency gaps remains a central challenge in network science and transport research [9, 12, 13].

Previous work has addressed traffic equilibria [14–17], equilibrium algorithms [18–20], link removal and capacity management [21–23], and paradoxical network effects [8, 10, 24, 25]. Yet, systematic and scalable methods for quantifying how marginal changes to individual links affect overall network performance are still limited. Identifying critical and paradoxical edges through edge-impact analysis can provide valuable guidance for infrastructure management and policy [22, 23].

1.1 Motivation and contribution

Many cities face the dual challenge of mitigating congestion while reducing emissions and reclaiming public space from car traffic [26, 27]. Expanding road capacity has long been a standard response, but this strategy often induces additional traffic and fails to improve long-term performance [28, 29]. The Braess paradox provides a theoretical basis for why more capacity does not always lead to better outcomes. Addressing this challenge requires tools that can identify the network elements with the greatest systemic impact and thereby support targeted interventions such as congestion pricing, access restrictions, or adaptive signal control [30–32].

To this end, we introduce *social cost gradient centrality* (SCGC), a novel edge importance measure that quantifies how marginal changes in free-flow travel times affect total social cost. SCGC highlights both critical links that strongly influence system efficiency and Braessian links whose improvement would paradoxically worsen it.

The main contributions of this paper are

1. a closed-form matrix formulation of static traffic assignment with linear cost functions, showing that total social cost is piecewise linear in free-flow travel times,
2. the definition of SCGC, unifying the detection of critical and paradoxical edges within a single framework,
3. applications to synthetic and real-world urban networks, demonstrating how self-organized traffic patterns emerge from the interplay of network topology, user behavior, and link characteristics.

1.2 Structure of the paper

The remainder of this paper is structured as follows. Section 2 reviews the traffic assignment model and the underlying cost function. We also revisit the Braess paradox and present an extended formulation. Section 3 derives the SCGC measure analytically. Section 4 applies the method to synthetic and real-world networks. Section 5 discusses policy implications, emphasizing the identification of structurally important and paradoxical links, and concludes with an outlook on future research directions.

2 Theory: traffic modeling

2.1 Cost function

Transportation networks are naturally modeled as directed graphs $G = (V, E)$, where edges $e \in E \subseteq V \times V$ represent road segments and nodes $n \in V$ represent intersections. If a road between two intersections n and m has lanes in opposite directions, both edges $e = (n, m)$ and $e = (m, n)$ belong to the edge set E .

To analyze the traffic flow $f_e \geq 0$ on a given road segment e , we define a cost function $t_e(f_e)$ that represents travel time as a function of both traffic volume and road capacity. Under the assumption that drivers follow basic traffic rules (e.g., matching the speed of the vehicle ahead, obeying speed limits, and maintaining a minimum speed), travel time is often modeled using the ‘Bureau of Public Roads’ (BPR) formulation [33]:

$$t_e^{(k)}(f_e) = \beta_e(1 + \alpha_e f_e^k), \quad (1)$$

where β_e is the free-flow travel time, α_e is a parameter that captures how sensitively travel time reacts to congestion, and k determines the degree of nonlinearity in the relationship. Throughout this paper, the term *cost function* refers to the edge-level travel time function $t_e(f_e)$. The network-level quantity $SC(f_e)$ introduced below is the corresponding *social cost*, i.e., the flow-weighted sum of travel times over all edges.




Fig. 1 Traffic flows f_e in a simple network where 10 drivers travel from node A to node D under two routing regimes: **a** user equilibrium, where each driver minimizes their individual travel time, and **b** social optimum, where the total system travel time is minimized. The congestion parameter is set to $\alpha_e = 1$ for all edges. The free-flow travel times are given by $\beta_{AB} = 0$, $\beta_{AC} = 10$, $\beta_{BD} = 10$, $\beta_{BC} = 1$, and $\beta_{CD} = 0$

In this work, we focus on the linear approximation of Eq. (1) (i.e., $k = 1$), which simplifies the travel time function to

$$t_e(f_e) = \alpha_e f_e + \beta_e, \quad t_e(f_e) \in [t_e^0, t_e^{\max}]. \quad (2)$$

The lower bound t_e^0 corresponds to the free-flow travel time determined by the road speed limit, whereas the upper bound $t_e^{\max}$ is set by walking speed travel time under extreme congestion.

2.2 Traffic equilibria and social cost

We analyze traffic dynamics by modeling each driver as a rational agent seeking to minimize travel time between a given origin and destination. We further assume that drivers have perfect, real-time knowledge of traffic conditions, effectively implying universal GPS usage. These assumptions lead to the concept of a user equilibrium: a state in which no driver can reduce travel time by unilaterally changing routes. This principle, also known as Wardrop's first principle [11], implies that all used routes between an origin–destination pair have equal and minimal travel times, while unused routes have higher travel times. The resulting traffic pattern satisfies the conditions of a Nash equilibrium in a network game among travelers [9, 34].

While the user equilibrium represents an individually optimal state, it does not necessarily minimize overall congestion. In contrast, if routing is coordinated to minimize total travel time across the entire network,

$$SC(f_e) = \sum_{e \in E} f_e t_e(f_e), \quad (3)$$

the system reaches a system optimum—a state in which traffic is distributed to maximize overall efficiency. This objective is given by the social cost in Eq. (3). This principle, also known as Wardrop's second principle [11], yields an allocation in which some drivers may experience longer travel times than in user equilibrium, but total congestion and aggregate delay are minimized.

The discrepancy between social cost at user equilibrium and at system optimum highlights inefficiencies introduced by selfish routing. This inefficiency is captured by the price of anarchy [34], which quantifies how strongly uncoordinated behavior degrades overall performance. In Fig. 1, we compare traffic flows f_e in a simple network where 10 drivers travel from node A to node D under (a) user equilibrium and (b) system optimum. In user equilibrium, each driver chooses routes to minimize individual travel time. In system optimum, routing is coordinated to minimize total travel time across all drivers. The total social cost is higher in user equilibrium ($SC(f_e) = 200$) than in system optimum ($SC(f_e) = 190$), illustrating how decentralized decisions can produce inefficient system-level outcomes. In both panels, the edge flows f_e shown are the resulting flow allocations for 10 drivers traveling from A to D under the stated link cost functions: one under user equilibrium and one under system optimum. The reported SC values are the corresponding total social costs obtained from these flows via Eq. (3).

2.3 The Braess paradox

The Braess paradox describes the counterintuitive phenomenon in which adding a link to a transportation network can increase overall congestion and travel times [8]. This paradox arises under user equilibrium because drivers




Fig. 2 The Braess paradox in a simple four-node network. **a** Total social cost of the user equilibrium as a function of the free-flow travel time β_{CB} . **b** Network topology with 40 drivers traveling from origin A (red) to destination D (blue), along with the cost functions of all edges. The red cross indicates the case where edge (C, B) is removed ($\beta_{CB} \rightarrow \infty$)

minimizing their individual travel times can inadvertently create a globally suboptimal state. Similar to extensions of this concept to supply networks [35, 36], we define the Braess paradox more broadly: a link is considered Braessian if adding capacity to it increases social cost in user equilibrium. This naturally extends the classical definition, which corresponds to the case where the link initially has zero capacity. For the following argument, consider the linear cost approximation in Eq. 2. Since β_e represents free-flow travel time on edge e , we use it to identify Braessian behavior: if total social cost decreases as β_e increases, then edge e is Braessian.

Definition II.1 (*Braessian edge*) An edge e in a transportation network is called *Braessian* if an increase in its capacity (or equivalently, a decrease in its free-flow travel time β_e) leads to an increase in the total social cost of the system. Formally, an edge e is Braessian if

$$\frac{\partial SC(f_e)}{\partial \beta_e} < 0. \quad (4)$$

In Fig. 2, we illustrate the Braess paradox for a simple network topology (panel b) with the corresponding edge cost functions. In this example, 40 drivers travel from node A (red) to node D (blue). Panel a shows the total social cost in user equilibrium (see Eq. (3)) as a function of free-flow travel time β_{CB} . As expected, increasing β_{CB} from zero to approximately one initially raises social cost. However, as β_{CB} increases further (up to about 2.5), social cost declines and eventually stabilizes at a value lower than in the initial case of zero free-flow travel time. The red cross marks the social cost when edge (C, B) is removed entirely, corresponding to $\beta_{CB} \rightarrow \infty$.

This behavior highlights the paradox. When edge (C, B) is highly efficient (small β_{CB}), it attracts additional flow, so many drivers choose the apparently fastest route. The resulting overuse can increase congestion. Initially, increasing free-flow travel time on this edge raises overall travel time, as expected. Beyond a threshold, however, a paradoxical effect appears: as β_{CB} grows further, drivers redistribute to alternative, less congested routes, reducing total social cost below the initial configuration. As the edge becomes less attractive, inefficient routing is mitigated and the system approaches a more efficient user equilibrium.

2.4 The traffic assignment problem

Having established the concepts of user equilibrium and system optimum, we now turn to the formal mathematical formulation of traffic assignment. The traffic assignment problem (TAP) seeks to determine how trips are distributed across a network for a given demand and network topology. This problem can be formulated in various ways, including link-based, path-based, and cycle-based representations [37], which are equivalent under appropriate conditions (see Appendix 1 for details). Variational-inequality formulations provide an equivalent equilibrium characterization and theoretical basis for these models [38].

In this work, we employ a node-link formulation of the TAP [39]. To formally describe the network structure, we define the node-edge incidence matrix E_{ne} , which encodes the connectivity between nodes and edges:

$$E_{ne} = \begin{cases} 1, & \text{if edge } e \text{ originates at node } n, \\ -1, & \text{if edge } e \text{ terminates at node } n, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

The TAP is then stated as the optimization problem

$$\begin{aligned}
 & \min_{f_e} \quad \sum_{e \in E} \mathcal{F}(f_e), \\
 & \text{subject to} \quad \sum_{e \in E} E_{ne} f_e^w = p_n^w, \quad \forall n \in V, \quad \forall w \in W, \\
 & \quad f_e^w \geq 0, \quad \forall e \in E, \quad \forall w \in W, \\
 & \quad f_e = \sum_{w \in W} f_e^w, \quad \forall e \in E.
 \end{aligned} \tag{6}$$

Here, f_e^w denotes the flow on edge e associated with the origin–destination (OD) tuple w , while f_e represents the total flow on edge e , obtained by summing over all OD tuples. For the analytical derivations below, each OD tuple is taken in the form $w = (o, d_1, d_2, \dots, d_n)$, i.e., one origin node o supplying multiple destinations. This single-origin convention matches the computational setup and provides one gauge condition ($\lambda_o^w = 0$) per tuple-specific support graph. The opposite convention (multiple origins, one destination) can be treated analogously by reversing edge directions. The demand corresponding to an OD tuple w is represented by the OD matrix whose components are defined as

$$p_n^w = \tag{7}$$

$$\begin{cases} y^o, & \text{if node } n \text{ is an origin in } w, \\ y^d, & \text{if node } n \text{ is a destination in } w, \\ 0, & \text{otherwise.} \end{cases} \tag{8}$$

Here, $y^o > 0$ denotes the total demand injected at an origin node, and $y^d < 0$ the total demand withdrawn at a destination node. Additionally, the demand vector $\mathbf{p}^w$ is constructed to satisfy flow conservation: $\sum_{n \in V} p_n^w = 0$.

The objective function $\mathcal{F}(f_e)$ for the user equilibrium of the TAP is given by the cumulative travel cost experienced by all drivers on each edge, which is represented by the integration of the cost function [40]

$$\mathcal{F}_{\text{ue}}(f_e) = \int_0^{f_e} t_e^{(k)}(u) du. \tag{9}$$

In the linear approximation of the travel time function (see Eq. (2)), this becomes

$$\mathcal{F}_{\text{ue}}(f_e) = \frac{1}{2} \alpha_e f_e^2 + \beta_e f_e. \tag{10}$$

The system optimum, in contrast, is obtained by minimizing the total social cost across all edges (see Eq. 3). Under the same linear approximation, the corresponding objective function is

$$\mathcal{F}_{\text{so}}(f_e) = f_e t_e(f_e) = \alpha_e f_e^2 + \beta_e f_e. \tag{11}$$

Solving the optimization problem in Eq. (6) allows us to determine the equilibrium flow configuration that corresponds to either the user equilibrium or the system optimum, depending on the form of the objective function. This formulation therefore provides the analytical foundation for evaluating how changes in network structure or edge costs affect overall traffic patterns and system efficiency.

It is worth noting that, under the linear travel time approximation, the user equilibrium and system optimum formulations differ only by a factor of $\frac{1}{2}$ in the quadratic term of the objective function. Consequently, both problems share the same feasible region but yield distinct flow allocations due to the different weighting of congestion effects.

3 Methods: solving the optimization problems

3.1 Algebraic solution

In this section, we derive an analytical solution for the user equilibrium of the TAP using the linear cost function specified in Eq. 2. To obtain a closed form, we first consider the equality-constrained system on a fixed active set (edges with positive flow). In practice, this active set is identified once from the full constrained problem, and the algebraic derivation below is then carried out on that fixed support. For consistent terminology, we define

$$\mathcal{A} := \{(w, e) \in W \times E : f_e^w > 0\} \quad (\text{stacked active set}), \quad (12)$$

$$E^w := \{e \in E : f_e^w > 0\} \quad (\text{active-edge set indexed by } w), \quad (13)$$

$$G^w := (V, E^w) \quad (\text{active-support subgraph}). \quad (14)$$

Thus, fixing the active set $\mathcal{A}$ is equivalent to fixing all active-support subgraphs G^w . The corresponding Lagrangian function is formulated as follows:

$$\begin{aligned} \mathcal{L}(f_e^w, \lambda_n) = & \sum_{e \in E} \left(\frac{1}{2} \alpha_e f_e^2 + \beta_e f_e \right) \\ & + \sum_{n \in V} \sum_{w \in W} \lambda_n^w \left(p_n^w - \sum_{e \in E} E_{ne} f_e^w \right), \end{aligned} \quad (15)$$

where λ_n^w denotes the Lagrange multiplier associated with node n and OD tuple w .

To compute the optimal flows, we take the derivative of the Lagrangian with respect to f_e^w and set it to zero:

$$\frac{\partial \mathcal{L}}{\partial f_e^w} = \alpha_e \left(\sum_{w \in W} f_e^w \right) + \beta_e - \sum_n \lambda_n^w E_{ne} \stackrel{!}{=} 0, \quad (16)$$

where we have used the total flow condition $f_e = \sum_{w \in W} f_e^w$ from the optimization problem Eq. 6.

We can solve this equation to obtain the flows directly. However, because we have not yet enforced the non-negativity constraint $f_e^w \geq 0$, the resulting flows may contain negative values, which are physically meaningless for vehicle flow. For an edge $e = (n, m)$ in the active set of tuple w , Eq. 16 implies

$$f_e = \frac{1}{\alpha_e} \left(\sum_{n \in V} \lambda_n^w E_{ne} - \beta_e \right) = \frac{1}{\alpha_e} (\lambda_n^w - \lambda_m^w - \beta_e) \geq 0. \quad (17)$$

From this, we obtain the following necessary condition for active edges:

$$\lambda_n^w \geq \lambda_m^w + \beta_e, \quad (18)$$

which implies that, for each OD tuple w , the multiplier at any node n must be at least as large as the multiplier at its neighboring node m , plus the connecting edge length β_e .




Fig. 3 The directional-consistency condition imposed by Eq. (18). Along each directed edge carrying positive flow, the multiplier at the tail node exceeds the multiplier at the head node by at least β_e . Consequently, the active support cannot contain directed loops when all edge lengths are strictly positive




Fig. 4 Decomposition of flows in an example graph under a positive-flow constraint. **a** Total flow f_e , composed of the individual flow components f_e^w shown in **b–e**. Each active-support subgraph $G^w(V, E^w)$ also displays its respective OD matrix components p_n^w color-coded on the nodes

The implications of this condition are illustrated in Fig. 3. For a directed edge $n \rightarrow m$ in the active set, Eq. (18) gives $\lambda_n \geq \lambda_m + \beta_{nm}$. Summing this inequality along a directed cycle C would yield

$$\lambda_n \geq \lambda_n + \sum_{e \in C} \beta_e, \quad (19)$$

which is impossible if all $\beta_e > 0$. Hence, under this positivity assumption, each active-support subgraph is acyclic. This statement concerns the active-support subgraphs used in the algebraic solution, not necessarily the full physical network.

We can circumvent this limitation by decomposing the original graph into a set of active-support subgraphs. Consider an origin–destination tuple $w = (o, d_1, d_2, \dots, d_{|V|})$ and define the corresponding active-edge set by $E^w = \{e \in E : f_e^w > 0\}$. Under Eq. (18) with strictly positive β_e , the active-support subgraph $G^w(V, E^w)$ is acyclic.

Repeating this construction for multiple origin–destination pairs, where each node serves as an origin exactly once, naturally leads to a decomposition of the overall flow into active-support subgraphs $G^w(V, E^w)$. Each active-support subgraph shares the same node set V but includes only edges with positive-flow values (so that Eq. (18) is satisfied), i.e., $E^w \subseteq E$. Consequently, each index w defines one active-support subgraph, ensuring that all flow values remain non-negative. The effective edge incidence matrix for each active-support subgraph can be expressed as

$$E_{ne}^w = \begin{cases} 1 & \text{if edge } e \text{ originates at node } n \text{ and } f_e^w > 0, \\ -1 & \text{if edge } e \text{ terminates at node } n \text{ and } f_e^w > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (20)$$

By considering these active-support subgraphs, we can reformulate the system using a generalized node–edge incidence matrix, where columns corresponding to edges that do not exist in active-support subgraph w are omitted:

$$E = \begin{bmatrix} E^1 & 0 & \dots & 0 \\ 0 & E^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & E^W \end{bmatrix} \in \mathbb{R}^{|V| \times (|E^1| + |E^2| + \dots + |E^W|)}. \quad (21)$$

The total numbers of edges and nodes in all subflows are

$$N_{t,e} = \sum_{w \in W} |E^w| \leq |E| \cdot |W|, \quad (22)$$

$$N_{t,n} = \sum_{w \in W} |V^w| = |V| \cdot |W|. \quad (23)$$

For convenient notation, we introduce the function $e(i)$, which returns the index of the edge in the original graph that corresponds to the edge in the active-support subgraph referred to by i in the matrix $\mathbf{E}$.

This provides a structured way to represent the disaggregated flow components. If the active-support subgraph structure is known, the problem can be solved immediately using Eq. 16. The active-support subgraph structure itself can be determined numerically by solving the problem once, after which the analytical formulation can be applied directly. This decomposition is particularly useful because it enables an analytical solution while preserving the non-negativity of flows. Thus, we can adjust the parameters α_e and β_e of the system to generate similar solutions. The advantage of this approach is that it allows us to obtain multiple solutions while requiring a numerical solver only once. The approach relies on the assumption that small variations in these parameters do not alter the active-support subgraph structure, ensuring the validity of the analytical solution.

In Fig. 4, we illustrate the traffic flow on a small graph under a positive-flow constraint. Panel a depicts the total flow f_e , while panels b–e show the corresponding flow decompositions for each source w . To determine the individual active-support subgraph structures $G^w(V, E^w)$, we first solve the full system numerically using the MOSEK or OSQP API for Python [41, 42] to obtain the respective Lagrange multipliers λ_n^w . In this example, we consider four sources ($W = 4$), leading to four active-support subgraphs, which are shown in panels b–e. Combining these four active-support subgraphs reconstructs the total flow f_e shown in panel a.

3.2 Matrix formulation of the TAP

In order to prove in later sections that the social cost at user equilibrium is linear in the free-flow travel time β_e , we first reformulate the optimality conditions of Eq. (16) in matrix form. This allows us to express the TAP as a linear system of equations, which is convenient for subsequent analytical steps.

We start by incorporating the flow conservation constraint of the TAP,

$$\sum_{e \in E} E_{ne} f_e^w = p_n^w, \quad (24)$$

and combine it with Eq. (16) to obtain

$$M \begin{pmatrix} \mathbf{f} \\ \boldsymbol{\lambda} \end{pmatrix} = \begin{pmatrix} -\boldsymbol{\beta} \\ \mathbf{p} \end{pmatrix}. \quad (25)$$

where the vector $\mathbf{f}$ collects the flow variables f_e^w and the vector $\boldsymbol{\lambda}$ collects the Lagrange multipliers λ_n^w for all nodes and OD-tuples w . The flow vector $\mathbf{f}$ is defined as

$$\mathbf{f} = (\mathbf{f}^1 \mathbf{f}^2 \dots \mathbf{f}^{|W|})^\top \in \mathbb{R}^{N_{t,e}}. \quad (26)$$

The entries of $\mathbf{f}$ are ordered by stacking the flows for each OD tuple w consecutively; the first $|E^1|$ elements correspond to active-support subgraph $G^1(V, E^1)$, the next $|E^2|$ to $G^2(V, E^2)$, and so on (compare with Fig. 4, panels b–e).

The vector of Lagrange multipliers $\boldsymbol{\lambda}$ is defined as

$$\boldsymbol{\lambda} = (\boldsymbol{\lambda}^1 \boldsymbol{\lambda}^2 \dots \boldsymbol{\lambda}^{|W|})^\top \in \mathbb{R}^{N_{t,n}}, \quad (27)$$

where $\boldsymbol{\lambda}^w$ contains the multipliers for active-support subgraph $G^w(V, E^w)$. The ordering follows the same stacking convention as for $\mathbf{f}$.

Analogously, we define the right-hand-side vectors: $\boldsymbol{\beta} \in \mathbb{R}^{N_{t,e}}$ contains the edge cost coefficients $\beta_i = \beta_{e(i)}$, and $\mathbf{p} \in \mathbb{R}^{N_{t,n}}$ contains the nodal demands p_n^w .

Finally, the coefficient matrix $\mathbf{M}$ has the block structure

$$\mathbf{M} = \begin{pmatrix} \mathbf{K} & \mathbf{E}^\top \\ \mathbf{E} & \mathbf{0} \end{pmatrix} \in \mathbb{R}^{[N_{t,e}+N_{t,n}] \times [N_{t,e}+N_{t,n}]}, \quad (28)$$

where $\mathbf{E}$ is the stacked edge–node incidence matrix introduced in Eq. (21). Furthermore, the matrix $\mathbf{K} \in \mathbb{R}^{N_{t,e} \times N_{t,e}}$ is defined as

$$\mathbf{K} = \begin{cases} \alpha_{e(j)} & \text{if } e(l) = e(j), \\ 0 & \text{otherwise.} \end{cases} \quad (29)$$

and $\mathbf{0}$ denotes the zero matrix.

For many configurations, this problem is underdetermined, because only the total flows on the edges matter, which allows for multiple configurations of how the subflows create these optimal total flows. Furthermore, the solutions depend only on differences in λ values, so one reference λ must be fixed for each active-support subgraph. Finally, the condition that every active-support subgraph must have the same inflow and outflow is also fixed by the structure of the equations. In order to obtain an invertible problem, each of these issues must be addressed.

To fix how the subflows combine into the total flows, the most straightforward way is to minimize the square of the individual flows; this favors an even share between different flows on the individual edges. This is done by adding a small value δ to the diagonal of $\mathbf{K}$. This δ should be much smaller than all α , so it only resolves the problem of multiple solutions without changing the total flows.

The λ value of the single source for each origin–destination tuple is set to zero, and the continuity equation for that same node is removed in each active-support subgraph because it is redundant.

This gives the resulting matrix,

$$\tilde{\mathbf{M}} = \begin{pmatrix} \tilde{\mathbf{K}} & \tilde{\mathbf{E}}^\top \\ \tilde{\mathbf{E}} & \mathbf{0} \end{pmatrix} \in \mathbb{R}^{[N_{t,e}+N_{t,n}-|W|] \times [N_{t,e}+N_{t,n}-|W|]}, \quad (30)$$

$$\tilde{\mathbf{K}} = \mathbf{K} + \delta. \quad (31)$$

This matrix is invertible and the solutions for the flows and λ 's can therefore be written as

$$\begin{pmatrix} \mathbf{f} \\ \tilde{\lambda} \end{pmatrix} = \tilde{\mathbf{M}}^{-1} \begin{pmatrix} -\beta \\ \tilde{\mathbf{p}} \end{pmatrix}. \quad (32)$$

To simplify the notation, the tildes are omitted from here on. Using the Schur complement,

$$\mathbf{S} = -\mathbf{E}\mathbf{K}^{-1}\mathbf{E}^\top, \quad (33)$$

which corresponds to the negative of the Laplacian for the system of active-support subgraphs. Then, we can write

$$\mathbf{M}^{-1} = \begin{pmatrix} \mathbf{K}^{-1} + \mathbf{K}^{-1}\mathbf{E}^\top\mathbf{S}^{-1}\mathbf{E}\mathbf{K}^{-1} & -\mathbf{K}^{-1}\mathbf{E}^\top\mathbf{S}^{-1} \\ -\mathbf{S}^{-1}\mathbf{E}\mathbf{K}^{-1} & \mathbf{S}^{-1} \end{pmatrix}. \quad (34)$$

Using this matrix formulation, we provide a compact algebraic representation of the TAP solution, which will be particularly useful for deriving analytical properties of the equilibrium solution in subsequent sections.

3.3 Linearity of social cost

Lemma 1 (Linearity of social cost in user equilibrium) *For a fixed active-support subgraph in the traffic assignment problem, the social cost at user equilibrium is linear in the free-flow travel times β_e ; across active-support subgraph changes, it is piecewise linear.*

Proof First, we note that the total social cost is given by

$$SC = \sum_e f_e t_e \quad (35)$$

$$= \sum_e f_e (\alpha_e f_e + \beta_e) \quad (36)$$

$$= \sum_w \sum_e f_e^w (\alpha_e f_e + \beta_e). \quad (37)$$

From Eq. (16), we know that this is equal to

$$SC = \sum_w \sum_e f_e^w (\lambda_{o(e)}^w - \lambda_{d(e)}^w). \quad (38)$$

where $\lambda_{o(e)}^w$ and $\lambda_{d(e)}^w$ denote the OD-specific multipliers at the origin and destination nodes of edge e . By noting that the vector $\mathbf{p}$ contains the flow starting at the unique origin node of active-support subgraph w , $o(w)$, and ending at node n , and that the total travel time is independent of the path taken, we have

$$\begin{aligned} t_{\text{total}, n}^w &= \sum_{e \in \text{path}(o(w), n)} t_e \\ &= \sum_{e \in \text{path}(o(w), n)} (\lambda_{o(e)}^w - \lambda_{d(e)}^w) \\ &= \lambda_{o(w)}^w - \lambda_n^w. \end{aligned} \quad (39)$$

Therefore,

$$SC = - \sum_w \sum_n p_n^w t_{\text{total}, n}^w = - \sum_w \sum_n p_n^w (\lambda_{o(w)}^w - \lambda_n^w). \quad (40)$$

Here, $\mathbf{p}$ is the reduced stacked vector that contains destination-withdrawal magnitudes only. Using the stacked notation from before and the fact that we set $\lambda_{o(w)}^w = 0$, we obtain

$$SC = - \sum_i p_i \lambda_i. \quad (41)$$

With Eq. (32), this can be rewritten as

$$SC = - \sum_i p_i (\mathbf{S}^{-1} \mathbf{E} \mathbf{K}^{-1} \boldsymbol{\beta} + \mathbf{S}^{-1} \mathbf{p})_i. \quad (42)$$

Hence, SC is linear in every β_k . The derivative of SC with respect to β_k is

$$\frac{\partial SC}{\partial \beta_k} = -p_q \mathbf{S}_{qp}^{-1} \mathbf{E}_{pl} \mathbf{K}_{lj}^{-1} \delta_{e(j), k}, \quad (43)$$

where summation over repeated indices is implied. The analytical expression for $\mathbf{K}_{lj}^{-1}$ is

$$\mathbf{K}_{lj}^{-1} = \frac{1}{\delta} \begin{cases} 1 - \frac{\alpha_{e(j)}}{\delta + \alpha_{e(j)} n(e(j))}, & l = j, \\ -\frac{\alpha_{e(j)}}{\delta + \alpha_{e(j)} n(e(j))}, & e(l) = e(j) \text{ and } l \neq j, \\ 0, & \text{otherwise.} \end{cases} \quad (44)$$

where $n(e(i))$ is the number of subflows that include edge $e(i)$. Defining $\mathbf{L}_{lj} = \delta \mathbf{K}_{lj}^{-1}$, Eq. (43) becomes

$$\frac{\partial SC}{\partial \beta_k} = p_q (\mathbf{E} \mathbf{L} \mathbf{E}^\top)^{-1}_{qp} \mathbf{E}_{pl} \mathbf{L}_{lj} \delta_{e(j), k}. \quad (45)$$

For a fixed active-support subgraph, this expression is independent of β_k (indeed of $\boldsymbol{\beta}$), so the derivative is constant in that regime. Writing it this way removes the prefactor $1/\delta$, which makes the step-wise calculation on a computer more accurate, even for very small δ . $\square$

In this proof, we used the algebraic solution of user equilibrium flows (given by Eq. (32) for the optimization problem in Eq. (6)) to compute the derivative of social cost with respect to an edge's free-flow travel time β_e . We showed that this derivative is constant with respect to β_e . This property allows us to determine whether social cost increases or decreases as β_e changes, thereby assessing whether individual edges contribute to system efficiency or inefficiency.

However, there are important caveats. First, we must solve the system numerically once to identify active-support subgraphs $G^w(V, E^w)$ where algebraic solutions exist. Second, linearity of social cost can only be established within a fixed active-support subgraph. If β_e changes substantially, the active-support subgraph may change, and the algebraic solution may no longer apply. Consequently, this approach characterizes the system under infinitesimal changes in β_e ; equivalently, it establishes piecewise linearity (see Appendix 2 for a quantitative stability scan and Appendix 3 for the local first-breakpoint predictor).

3.4 Social cost gradient centrality

We refer to the newly introduced measure defined in Eq. (43) as *social cost gradient centrality* (SCGC). It quantifies the sensitivity of the total social cost to incremental variations in the free-flow travel time parameter β_e . By explicitly capturing how marginal adjustments to edge characteristics influence network-wide performance, SCGC ranks edges according to their functional significance in maintaining, improving, or impairing the system's efficiency.

Edges with high-SCGC values are critical for overall network performance, as increases in their free-flow travel time would lead to substantial rises in total social cost. Such edges form the structural backbone of the network and require careful operational management. In contrast, edges with low SCGC values have little influence on aggregate efficiency, making them flexible candidates for reallocation, traffic calming, or other interventions. Finally, edges with negative SCGC values correspond to *Braess edges*: decreasing their free-flow travel time would paradoxically increase total system cost. This offers a direct and quantitative criterion for detecting Braess-like behavior in large-scale networks without relying on stylized examples or exhaustive scenario testing.

Classical formulations of the Braess paradox [8] identify paradoxical behavior through specific network topologies or through numerical experiments that remove edges and recompute equilibria. SCGC generalizes this concept: rather than requiring $|E|+1$ equilibrium computations, we obtain the gradient information analytically after solving the traffic assignment problem once. This makes the method computationally tractable even for large real-world networks.

Leveraging Lemma 1, the SCGC can be computed efficiently, as it only requires the pseudo-inverse of the Schur complement in Eq. (33), obtained from the active-set linear system for Eq. (6). The measure provides immediate insight into whether small increases in edge cost parameters improve or worsen system efficiency. It is important to note that the method assumes the active-support subgraph $G^w(V, E^w)$ remains unchanged under small variations in β_e . While valid for marginal perturbations, this assumption may not hold for larger interventions such as complete edge removal. Therefore, SCGC provides a reliable local indicator for identifying Braess edges and other critical links, while larger structural changes may require additional equilibrium computations. In practice, SCGC can also serve as a warm-start for brute-force Braess-edge discovery by restricting expensive equilibrium resolves to a top-ranked candidate subset.

4 Results: applications to real-world road networks

In the following, we apply the SCGC framework to both synthetic and real-world urban road networks to demonstrate its computational efficiency and interpretive value.

4.1 Validating social cost gradient centrality

To validate our approach, Fig. 5 illustrates the derivative of social cost with respect to β_e for each edge in a small example graph. Panel a shows the analytical solution (see Eq. (43)). Panels b–f display numerical evaluations obtained through multiple numerical solutions, specifically, solving the system 25 times per edge, amounting to 125 solutions in total for this small example. The analytical method thus significantly reduces computational effort.

Edges highlighted in red exhibit positive derivatives, indicating that increasing their costs leads to higher social costs. In contrast, edges highlighted in blue show negative derivatives, implying that raising their costs reduces total social cost, thereby classifying them as Braessian.

Panels b–f further show that analytical and numerical solutions coincide in a neighborhood of the initial β_e values (highlighted in matching colors), where the relationship between social cost and β_e is piecewise linear. As β_e deviates further from its baseline, discrepancies emerge due to structural changes in the active-support subgraph,




Fig. 5 Derivative of the social cost with respect to β_e for each edge in a small example graph. **a** Analytical solution compared to **b–f** numerical approximations. Positive values (red) indicate that increasing β_e raises the social cost, while negative values (blue) suggest a decrease, marking the edge as Braessian




Fig. 6 Visualization of **a** traffic flows and **b** SCGC on a synthetic planar graph $G(V, E)$ with $|V|=50$ nodes and $|E|=270$ edges. **c** Scatter plot comparing traffic flows and SCGC values

underscoring that the analytical approximation is most reliable for small perturbations of β_e (see Appendices 2 and 3).

Additionally, in panels b–e, we demonstrate that the SCGC is positive, meaning that increasing the corresponding β_e values results in higher overall social cost. In contrast, panel f presents an edge with a negative SCGC, indicating that an increase in β_e for this edge would reduce social cost, revealing a Braessian effect.

4.2 Ranking edges by social cost gradient centrality

The SCGC quantifies the sensitivity of the total social cost with respect to changes in an edge's free-flow travel time β_e (e.g., due to speed limit adjustments). High-SCGC values indicate edges where small increases in travel time result in disproportionately large increases in total social cost, highlighting their systemic importance.

To illustrate this concept, Fig. 6 presents traffic flows and SCGC values for a synthetic planar graph $G(V, E)$ with $|V| = 50$ nodes and $|E| = 270$ edges, representing a stylized urban road network. The cost function parameter α_e (see Eq. (2)) was sampled uniformly from the interval $[0.1, 1]$. The free-flow travel time parameter β_e was determined from the geodesic length of each edge, assuming a uniform speed limit of 50, km/h.

Next, we assigned node populations according to a power-law distribution with an exponent of -1 , resulting in a few highly populated nodes and many nodes with relatively small populations, reflecting typical real-world urban network structures. Based on these populations, we constructed an origin–destination (OD) matrix in which a fraction of each node's population travels to other nodes.

In this formulation, we assume that the likelihood of drivers originating from or traveling to a node is determined solely by its population. As a result, the probability of traveling between any two nodes is independent of the distance separating them. While this assumption simplifies the model, a more realistic representation could incorporate distance effects, for example by using a gravity model [43]. We therefore interpret the synthetic benchmark as a controlled stress test of the method rather than as a calibrated demand model.

Panel a of Fig. 6 shows the resulting traffic flows, while panel b presents the corresponding SCGC values for each edge. Centrally located edges tend to exhibit both higher traffic flows and larger SCGC values. This is because more shortest paths pass through these edges, making them critical for network-wide accessibility. Consequently, even small perturbations to their free-flow travel times β_e can strongly affect total social cost.

Interestingly, five edges display negative SCGC values, indicating that increasing their free-flow travel times β_e would *decrease* the total social cost and thus identifying them as Braessian edges.

Panel c presents a scatter plot comparing traffic flows with SCGC values. While a positive correlation is evident, the relationship is not strictly linear or one-to-one. This demonstrates that SCGC captures additional structural information beyond simple traffic volumes, offering a more nuanced perspective on network vulnerability and intervention points.

4.3 Event-driven traffic in Cologne

To illustrate the practical relevance of our framework, we analyze the traffic impact of a major football event in Cologne, Germany. We construct the city's street network from OpenStreetMap (OSM) following the procedure in Appendix 5, resulting in a directed graph $G = (V, E)$ with $|V| = 835$ nodes and $|E| = 2438$ edges. Each road segment is assigned a linear cost function as described in Appendix 6, and travel demand is generated from population data according to Appendix 7. This yields a realistic yet computationally tractable representation of the city's road infrastructure and demand during large-scale events.

The primary source of traffic demand is the set of parking lots surrounding the football stadium. In the baseline calibration, this corresponds to $n_{\text{vehicles}} = 15,000$ event-related trips. The parking locations are geocoded, and the nearest nodes in the network are defined as destination nodes W . We then solve the traffic assignment problem as described in Sect. 3.1, yielding user equilibrium flows f_e on each edge. As shown in Fig. 7a, the resulting flow pattern exhibits strong spatial concentration along the city's main arterial roads leading to the stadium, which form the backbone of the event-related traffic system.

To assess the network's vulnerability to inefficiencies, we compute the derivative of the social cost with respect to β_e , $\partial SC / \partial \beta_e$ (SCGC), for each edge as outlined in Sect. 3.4. Edges with negative SCGC values indicate Braess edges, links where a decrease in capacity, for example by lowering speed limits (i.e., increasing free-flow travel time), could paradoxically reduce overall travel costs. Figure 7b highlights these edges in blue, revealing structurally important corridors and potentially counterproductive links whose management can substantially influence system performance.

Across the entire network of 2438 edges (with a total length of 2105 km), we identify 32 Braess edges in the baseline event scenario (about 1.31% of all links), with a combined length of 24.1 km (1.14% of total network length). The minimum SCGC value is $\min(\partial SC / \partial \beta_e) \approx -70.1$, while the largest positive value reaches $\max(\partial SC / \partial \beta_e) \approx 2846.7$, concentrated along a few high-capacity arterial corridors. The median SCGC is relatively low (2.66), reflecting that only a limited number of edges exert a dominant influence on network efficiency.

To test robustness, we scale the same OD matrix by factors 0.8, 1.0, and 1.2. The number of Braess edges increases from 18 (0.74%) at $0.8\times$ demand to 32 (1.31%) at baseline and 66 (2.71%) at $1.2\times$ demand; the corresponding Braess-edge length rises from 14.2 to 24.1 km and 60.4 km. Hence, paradoxical corridors become more prevalent as event demand intensifies.

We additionally perform a quantitative stress test of edge ranking methods (SCGC vs flow vs betweenness vs random) by applying $+1\%$ perturbations to selected free-flow times and recomputing equilibrium social cost.




Fig. 7 Visualization of vehicle flows and derivative social costs (SCGC) during a football event in Cologne, Germany. Color coding in **a** shows vehicle flows f_e across the road network. Color coding in **b** depicts the derivative of the social cost with respect to β_e , highlighting Braess edges in blue

The chosen perturbation magnitude is deliberately local: it lies below the median active-set stability radius of about 3.5% reported in Appendix 2, so first-order rankings are evaluated primarily within the locally stable regime. Intuitively, this asks a simple question: if we slightly “stress” the roads that each ranking declares most important, which ranking causes the largest system-wide damage under the same intervention budget? In this sense, a better ranking should produce a larger realized increase in social cost because it identifies truly critical links. SCGC-linear ranking dominates all comparators across budgets $k = 1, \dots, 8$ for all three demand levels (win share 100%). At baseline demand, the mean realized increase in social cost across budgets is 0.01485% for SCGC, versus 0.01098% (flow), 0.00882% (betweenness), and 0.00488% (random). For single-edge perturbations, the rank correlation between predicted and observed damage is $\rho \approx 0.9997$ for SCGC, compared with $\rho \approx 0.5457$ for flow and $\rho \approx 0.1029$ for betweenness. This confirms that SCGC is not only interpretable but also predictive for local intervention impact in this real-world case (see Appendix 4).

Major arterial roads exhibit the highest positive SCGC values, reflecting their structural role in maintaining efficient access to the stadium. These roads are critical for event-related traffic and should remain as free of disturbances as possible. Measures such as restricting through traffic, imposing parking bans, or providing dedicated lanes could help safeguard their capacity and ensure smooth flows.

In contrast, the identified Braess edges are primarily located on secondary or connector roads. Although they carry less flow, they have a disproportionately negative effect on network-wide performance. Reducing their attractiveness, e.g., through temporary speed limit reductions, dynamic access restrictions, or partial closures, can encourage more efficient routing and alleviate congestion elsewhere. Because these links are not central access routes to the stadium, such interventions can be implemented with minimal disruption to event traffic.

A promising traffic management strategy would be to prioritize access to high-SCGC (red) corridors for event traffic while diverting background traffic to lower-SCGC (blue or neutral) roads. This separation of flow classes can help prevent capacity bottlenecks and stabilize travel times. More generally, sensitivity-based indicators such as SCGC provide a systematic basis for identifying and managing structurally important or counterproductive links in urban road networks, enabling targeted, adaptive, and event-specific traffic control strategies.

Overall, this case study demonstrates how real-world data, traffic assignment modeling, and sensitivity analysis can be combined to produce actionable insights for urban mobility management. Rather than expanding capacity

uniformly, targeted interventions on a small number of critical and Braess edges offer a cost-effective means of improving network performance during major events.

4.4 Daily urban traffic analysis in Potsdam

To demonstrate the applicability of our framework to everyday mobility patterns, we analyze the daily traffic structure of Potsdam, Germany. The city's street network is constructed from OpenStreetMap using the methodology described in Appendix 5, resulting in a directed graph $G = (V, E)$ with $|V| = 112$ nodes and $|E| = 296$ edges. Each edge is assigned a linear cost function following Appendix 6, and travel demand is generated from population data as outlined in Appendix 7. We choose the OD matrix such that the resulting flows represent a moderate traffic scenario, which is representative of typical conditions in many medium-sized European cities. In contrast to the Cologne scenario, where demand is concentrated near a stadium, here it is distributed evenly across the network to reflect typical daily mobility.

We begin by analyzing edge utilization, defined as

$$u_e = 1 - \frac{\beta_e}{\alpha_e f_e + \beta_e}.$$

This measure reflects the share of available capacity that is effectively used on a given edge: values close to one indicate heavily used links operating near capacity, while values close to zero correspond to underutilized links with little traffic relative to their capacity. Figure 8a visualizes the resulting utilization pattern across the network. The average utilization is 0.46, with a median of 0.46 and a maximum of 0.94, indicating a clear differentiation between heavily used corridors and underutilized links. We identify 26 edges (8.8% of the network) with utilization above 0.8, corresponding to a total length of 16.4 km. These high-utilization corridors likely coincide with the city's primary arterials, concentrating most of the daily traffic flow and playing a crucial role in maintaining network accessibility and efficiency.

At the same time, 48 edges (16.2% of the network), covering 79.4 km of roadway, exhibit utilization values below 0.2. This underutilization highlights a large share of the network where capacity remains unused. However, low




Fig. 8 Visualization of vehicle flows and derivative social costs (SCGC) in Potsdam, Germany. Color coding in **a** shows the edge utilization $1 - \beta_e / (\alpha_e f_e + \beta_e)$ across the road network, while color coding in **b** depicts the derivative of the social cost with respect to β_e , highlighting Braess edges in blue

utilization alone does not necessarily imply low importance for overall network performance, since some low-flow links may still be structurally relevant. To determine their actual impact on efficiency, we turn to the SCGC analysis.

We next examine the derivative of the social cost with respect to β_e , $\partial SC/\partial \beta_e$ (SCGC), as described in Sect. 3.4. Figure 8b shows the spatial distribution of SCGC values, which range from $\min(\partial SC/\partial \beta_e) \approx -16.4$ to $\max(\partial SC/\partial \beta_e) \approx 1017.8$, with a mean of 200.4 and a median of 143.3. Most edges exhibit positive SCGC values, indicating that increasing travel times on these links would negatively affect overall network performance. These edges largely overlap with the high-utilization corridors identified earlier, underscoring their central role in supporting efficient mobility.

In contrast, we identify only 2 Braess edges, with a total length of 1.5 km. These links have negative SCGC values, meaning that increasing their free-flow travel times could paradoxically improve total network performance.

In contrast to the Cologne case (see Sect. 4.3), where demand and supply are highly asymmetric, a more balanced spatial distribution of origins and destinations substantially reduces the likelihood of Braess edges emerging. When flows are dispersed more evenly across the network, the structural conditions that give rise to paradoxical effects become less pronounced.

While many of these edges exhibit relatively low utilization, it is their low or even negative SCGC values, rather than low flow alone, that identify them as non-critical for overall system efficiency. In other words, their marginal contribution to total social cost is small or beneficial when reduced. Such links therefore provide strategic flexibility for targeted interventions, including speed limit reductions, lane reallocation to cycling or bus infrastructure, or temporary closures during peak periods, without compromising aggregate network performance.

This case study illustrates how combining utilization and SCGC analysis provides a more complete picture of network structure and function. High-SCGC and high-utilization corridors form the operational backbone of the system and should be protected to maintain stable traffic flow. Conversely, edges with low or negative SCGC values represent promising candidates for mode shift or other interventions that can promote more sustainable mobility without sacrificing efficiency. Potsdam's network structure thus reveals both clear operational priorities and opportunities for targeted, low-impact interventions. The cross-city intervention-ranking validation is provided in Appendix 4.

5 Discussion and conclusion

This study proposed *social cost gradient centrality* (SCGC), a novel sensitivity-based graph centrality for identifying critical and paradoxical edges in transportation networks. By linking marginal changes in link-level free-flow travel times to total social cost, SCGC provides a direct measure of how local infrastructure characteristics influence global network performance. Unlike traditional topological centrality measures, SCGC is grounded in the equilibrium flow structure of the network and reflects the emergent properties of self-organized traffic states.

5.1 Interpretation and implications

The results demonstrate that not all links contribute equally to system performance. A relatively small subset of edges exerts a disproportionately large influence on the overall social cost, highlighting their structural and functional centrality. Positive SCGC values indicate critical links whose degradation would significantly increase travel times, whereas negative values reveal Braessian edges where capacity expansion or travel time reductions would paradoxically worsen overall network efficiency. This provides a powerful diagnostic lens for understanding how local perturbations propagate through a self-organizing system.

From a conceptual standpoint, these findings align with the view of transport networks as complex adaptive systems and with Haken's principles of synergetics, where global order emerges from the interplay of a few dominant structures and many interacting components [44]. In this perspective, edges with high-SCGC values can be interpreted as effective 'order parameters' that shape macroscopic flow patterns. SCGC captures this local–global coupling by quantifying how infinitesimal link-level interventions affect system-wide performance.

A key implication is that SCGC extends conventional edge ranking beyond purely structural notions of importance. Metrics based only on topology or observed flow volumes may identify heavily used corridors, but they do not directly quantify the marginal system-level consequence of changing a link attribute. SCGC closes this gap by linking each edge to a directional sensitivity: it indicates not only whether an edge matters, but also whether local improvement is expected to help or harm global efficiency under equilibrium behavior.

At the same time, SCGC should be interpreted as a local response measure around the current operating point. Its value is strongest when used as a differential indicator that complements scenario-based planning, rather than replacing it. In practical terms, SCGC can prioritize where detailed simulations or intervention studies should be concentrated, thereby reducing the search space for policy design in large urban networks. This empirical usefulness is supported by the cross-city intervention validation (Appendix 4): across 42 tested budget points, SCGC ranking

is best in 100% of cases and attains substantially higher one-edge rank correlation with realized impact than flow or betweenness baselines. Put simply, when we perturb the top-ranked roads by the same small amount, the SCGC-ranked set consistently produces the largest realized deterioration in total network performance, confirming that SCGC better captures true marginal criticality.

5.2 Policy relevance

For urban mobility planning, the implications are twofold. First, the identification of critical edges enables targeted investment in infrastructure maintenance or enhancement where it has the greatest systemic benefit. Second, the detection of Braessian edges provides an analytical foundation for demand management, road space reallocation, or capacity reduction strategies that can improve overall efficiency despite appearing counterintuitive from a local perspective. These insights support the shift away from car-centric expansion strategies toward more intelligent network management and demand-side regulation.

Moreover, SCGC offers a quantitative tool for prioritizing measures such as congestion pricing, access restrictions, or adaptive signal control. By focusing on links with the strongest impact on total social cost, policymakers can implement targeted and cost-effective interventions, avoiding broad and inefficient measures that overlook the network's self-organizing structure.

From an implementation perspective, SCGC can support a staged decision workflow: first, compute a citywide sensitivity map to identify candidate links; second, perform robustness checks under alternative demand snapshots and operational assumptions; third, test selected interventions in pilot corridors before network-wide rollout. This progression combines computational tractability with policy caution and helps avoid expensive interventions on links with low systemic leverage.

Importantly, the presence of Braessian links does not imply that every capacity reduction is beneficial. Rather, it highlights where local changes warrant closer evaluation because network feedbacks are non-intuitive. In this sense, SCGC contributes to evidence-based governance by making such feedbacks measurable and transparent.

5.3 Limitations and future work

While the current formulation provides a clear analytical foundation, several important limitations remain. A key restriction arises from the fact that SCGC exploits the piecewise linearity of total social cost with respect to free-flow travel times, which is inherent to the underlying linear link cost functions and static equilibrium formulation. This allows for an elegant and computationally efficient derivation of marginal sensitivities, but it also means that SCGC is strictly valid only within a fixed active-support subgraph (equilibrium) regime. When marginal changes in link attributes lead to discontinuous shifts in route choice patterns or the activation of alternative paths, the derivative structure itself changes. Consequently, SCGC captures local sensitivity around the current equilibrium but does not provide a global picture of how larger interventions might reshape flow distributions and system performance (Appendices 2 and 3).

This limitation is particularly relevant in networks with multiple near-optimal equilibria or where congestion dynamics are dominated by threshold effects. In such settings, even small infrastructure or demand changes can trigger qualitative shifts in traffic patterns that fall outside the linear response captured by SCGC. Future research should address this by extending the method to nonlinear cost functions or by embedding SCGC within dynamic traffic assignment frameworks that explicitly model time-dependent and path-switching behavior.

A second promising direction lies in broadening the scope of SCGC beyond single-layer, single-mode networks. Integrating multimodal and multilayer structures would enable the assessment of resilience and efficiency across more realistic urban mobility systems, where capacity interventions or disruptions can have cascading effects across modes.

Finally, coupling SCGC with control or optimization approaches represents a powerful avenue for policy applications. By identifying the links with the highest marginal impact, SCGC could inform targeted road pricing, capacity management, or resilience strategies. Extending the concept to capture higher-order interactions between edges, such as pairwise or groupwise sensitivities, would further enhance its ability to reveal the structural mechanisms through which local interventions propagate through complex transport networks.

Further limitations concern data and model calibration. In real-world deployments, uncertainties in speed limits, lane counts, demand estimation, and network simplification steps can propagate into α_e , β_e , and OD matrices, thereby affecting SCGC rankings. Systematic uncertainty quantification, including confidence intervals for edge sensitivities, is therefore an important extension for operational use.

A related direction is temporal and behavioral heterogeneity. The present static equilibrium perspective does not explicitly resolve peak spreading, departure-time choice, incident response, or adaptive traveler learning. Incorporating these effects in dynamic or stochastic formulations would improve realism and clarify how stable SCGC rankings remain across daily and seasonal variability. Related work on scalable data-driven macroscopic traffic modeling has also explored network-level formulations that do not rely on fixed OD demand assumptions [45]. A

further step is to evaluate SCGC under richer demand constructions, including empirically calibrated OD matrices and other more realistic mobility-demand estimates.

5.4 Conclusion

This work advances the understanding of how local infrastructure characteristics shape global network behavior in self-organizing transport systems. SCGC provides a simple yet powerful tool to identify edges with disproportionate systemic importance, encompassing both critical and paradoxical effects. By grounding edge importance in network flows and social costs, this measure offers actionable insights for infrastructure planning, demand management, and policy design.

Methodologically, the contribution is twofold: a tractable analytical sensitivity framework for equilibrium traffic assignment under linear costs, and a centrality interpretation that connects marginal link perturbations to system-wide efficiency outcomes. Empirically, applications to synthetic and real-world networks show that SCGC can separate structurally important corridors from links that are operationally flexible, including rare but policy-relevant Braessian edges.

For practitioners, SCGC is most valuable as a prioritization layer: it identifies where limited planning and investment capacity is likely to yield the largest system-level effects, and where conventional intuition about ‘improvements’ may fail. For research, it opens a path toward sensitivity-based control of mobility systems that combines equilibrium theory, network science, and decision support.

SCGC operationalizes Haken’s insight that global order in complex systems arises from local interactions and a few dominant components. More broadly, it demonstrates how sensitivity analysis can bridge microscopic infrastructure attributes and macroscopic system performance, contributing to the broader study of self-organization and control in complex networks.

Funding Open Access funding enabled and organized by Projekt DEAL. Funding was provided by GRK 2043: Natural Hazards and Risks in a Changing World (NatRisk-Change) (Grant No. 251036843).

Data availability Input network data are obtained from OpenStreetMap via OSMnx, and population data are obtained from GHSL 2023 [50]. The preprocessing and calibration steps used to construct the final analysis datasets are implemented in the repository scripts and can be rerun end-to-end from raw sources.

Code availability All code required to reproduce the results, figures, and manuscript in this study is available at <https://github.com/jwassmer/tap-flows>. Core model and solver implementations are located in `src/`, experiment and figure scripts in `analysis/`. Generated figures and cached experiment tables are stored in `figs/` and `cache/`, respectively.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix 1: Link–path to node–link formulation

The traffic assignment problem (TAP) captures this behavior. Following [38, 40, 46], it can be formalized as an optimization problem

$$\begin{aligned}
 & \min_{f_e} \sum_{e \in E} \mathcal{F}(f_e), \\
 & \text{subject to } \sum_{r \in R^w} f_r^w = d^w, \quad \forall w \in W, \\
 & \quad f_r^w \geq 0, \quad \forall r \in R^w, w \in W, \\
 & \quad f_e = \sum_{w \in W} \sum_{r \in R^w} \Delta_{e,r}^w f_r^w,
 \end{aligned} \tag{46}$$

where f_e describes the flow on edge e , while f_r^w describes the flow on path $r \in R$ between origin–destination tuple (OD) $w \in W$, such that demand d^w is satisfied. Moreover, we define the link–path incidence matrix $\Delta_{e,r}^w$, which

is equal to one if edge e lies on path r for OD tuple w , and zero otherwise:

$$\Delta_{e,r}^w = \begin{cases} 1, & \text{if edge } e \text{ is on path } r \in R^w \text{ of OD tuple } w, \\ 0, & \text{otherwise.} \end{cases} \quad (47)$$

For the linear cost model used in this paper (see Eq. 2), the user equilibrium objective is

$$\mathcal{F}_{\text{ue}}(f_e) = \frac{1}{2}\alpha_e f_e^2 + \beta_e f_e. \quad (48)$$

The objective function for the system optimum of the TAP is given by the total travel time experienced by all drivers (i.e., the social cost), given by

$$\mathcal{F}_{\text{so}}(f_e) = f_e t_e. \quad (49)$$

Both objective functions are convex on the domain $[0, f_e^{\max}]$. Hence, the optimization problem (46) can be solved in principle in polynomial time [47]. This formulation of the TAP is referred to as the *link-path* formulation, as it requires the calculation of all paths from all origin–destination pairs. However, the number of possible paths between an origin and a destination generally increases rapidly (often exponentially) with graph size, making this formulation impractical for large urban networks.

To overcome the issue of having to compute all possible paths for a origin–destination pair, an alternative approach is to reformulate the TAP in a *node-link* formulation. This can be achieved by defining the node-edge incidence matrix of a graph as

$$E_{ne} = \begin{cases} 1 & \text{if edge } e \text{ originates at node } n, \\ -1 & \text{if edge } e \text{ terminates at node } n, \\ 0 & \text{otherwise.} \end{cases} \quad (50)$$

Then, we can write down the TAP in the following way:

$$\begin{aligned} & \min_{f_e} \sum_{e \in E} \mathcal{F}(f_e), \\ & \text{subject to } \sum_{e \in E} E_{ne} f_e^w = p_n^w, \quad \forall n \in V, w \in W \\ & \quad f_e^w \geq 0, \quad \forall e \in E, w \in W, \\ & \quad f_e = \sum_{w \in W} f_e^w \quad \forall e \in E. \end{aligned} \quad (51)$$

with

$$\mathcal{F}(f_e) = \frac{1}{2}\alpha_e f_e^2 + \beta_e f_e. \quad (52)$$

where

$$p_n^w = \begin{cases} d^w, & \text{if node } n \text{ is the origin in a OD tuple } w, \\ -d^w, & \text{if node } n \text{ is the destination in a OD tuple } w, \\ 0, & \text{otherwise.} \end{cases} \quad (53)$$

We will refer to the vector $\mathbf{p}^w$ with the elements p_i^w , $i \in V$ as source-sink vector.

Lemma 2 (Equivalence of link-path and node-link constraints) *The constraints of the link-path formulation of the Traffic Assignment Problem (TAP) in Eq. 46 can be equivalently transformed into the constraints of the node-link formulation of TAP in Eq. 51.*

Proof For one OD pair $w = (o, d)$, define edge flows induced by path flows as

$$f_e^w = \sum_{r \in R^w} \Delta_{e,r}^w f_r^w, \quad f_e = \sum_{w \in W} f_e^w. \quad (54)$$

Applying the node-edge incidence matrix to f^w gives

$$\sum_{e \in E} E_{ne} f_e^w = \sum_{r \in R^w} \left(\sum_{e \in E} E_{ne} \Delta_{e,r}^w \right) f_r^w. \quad (55)$$

For any path $r \in R^w$, the inner sum equals +1 at the origin node o , −1 at the destination node d , and 0 elsewhere. Therefore,

$$\sum_{e \in E} E_{ne} f_e^w = \begin{cases} \sum_{r \in R^w} f_r^w, & n = o, \\ -\sum_{r \in R^w} f_r^w, & n = d, \\ 0, & \text{otherwise.} \end{cases} \quad (56)$$

Using the link-path demand constraint $\sum_{r \in R^w} f_r^w = d^w$, we obtain

$$\sum_{e \in E} E_{ne} f_e^w = \begin{cases} d^w, & n = o, \\ -d^w, & n = d, \\ 0, & \text{otherwise,} \end{cases} = p_n^w, \quad (57)$$

which is exactly the node-link flow conservation constraint. The non-negativity constraints are preserved directly, since $f_r^w \geq 0$ implies $f_e^w \geq 0$ by construction. $\square$

Lemma 3 (Flow decomposition theorem [37, 48]) *Every route flow f_r has a unique representation as non-negative link flows f_e . Conversely, every non-negative link flow can be represented by route flows, though this representation need not be unique.*

Appendix 2: Active-set stability under local β -perturbations

To quantify the practical validity range of the fixed active-support subgraph assumption used by SCGC, we perform an active-set stability scan on a synthetic planar network. Starting from a baseline user equilibrium, we perturb one-edge free-flow parameter β_k at a time and recompute the multicommodity equilibrium. We then compare the resulting active set to the baseline active set.

In the experiment shown in Fig. 9, we use a synthetic network with 50 nodes and evaluate perturbations in the range $\Delta\beta_k/\beta_k \in [-5\%, 5\%]$ for a selected set of high-flow edges. Panel a reports, for each perturbed edge and perturbation level, the number of active-set entries that differ from baseline. Panel b summarizes two stability indicators over all tested edges: the exact stable fraction (black) and the mean Jaccard overlap of active sets (blue dashed).


Let $\mathcal{A}_0$ denote the baseline active set and $\mathcal{A}_{k,\delta}$ the active set after perturbing edge k by relative amount δ . The Jaccard similarity is computed as

$$J(\mathcal{A}_0, \mathcal{A}_{k,\delta}) = \frac{|\mathcal{A}_0 \cap \mathcal{A}_{k,\delta}|}{|\mathcal{A}_0 \cup \mathcal{A}_{k,\delta}|}. \quad (58)$$

Thus, $J = 1$ indicates exact equality of active sets, while $J < 1$ indicates structural changes.

The results show that active-set stability is highest close to $\Delta\beta_k = 0$, where the exact stable fraction is 1 and the Jaccard similarity is exactly 1 for all tested edges. At the perturbation bounds, the mean Jaccard values remain high ($\bar{J} \approx 0.9978$ at −5% and $\bar{J} \approx 0.9983$ at +5%), even though exact stability drops (stable fraction ≈ 0.55 at −5%). The global minimum over all tested edge–perturbation pairs is $J_{\min} \approx 0.9855$, observed for edge (13, 46) at −5%, corresponding to 6 changed active-set entries out of 414 baseline active entries. The global maximum is $J_{\max} = 1$, attained at $\delta = 0$ (and also at several nonzero perturbations for robust edges). For this benchmark, the median local stability radius is approximately 3.5%, supporting the interpretation of SCGC as a local sensitivity measure that is reliable for small perturbations and should be complemented by full re-equilibration for larger interventions.

Fig. 9 Active-set stability under edge-wise perturbations of free-flow travel time. **a** Number of active-set entries that differ from the baseline active set when perturbing one-edge parameter β_k at a time. **b** Aggregate stability metrics across tested edges: exact stable fraction (black) and mean Jaccard similarity (blue dashed). The annotation reports the median local perturbation radius around zero for which exact active-set identity is preserved



Appendix 3: Local first-breakpoint predictor

This appendix formalizes the local predictor used to estimate when a perturbation of one-edge parameter β_k first changes the active set. We denote by f_i the commodity-edge flow entries (stacked over tuples and edges), and by

$$s_i = \alpha_{e(i)} f_{e(i)} + \beta_{e(i)} - (\mathcal{E}^\top \lambda)_i \quad (59)$$

the corresponding complementarity slack, where $f_{e(i)} = \sum_w f_{e(i)}^w$ is total flow on physical edge $e(i)$. At baseline, active entries satisfy $f_i^0 > \varepsilon_{\text{act}}$ and (up to tolerance) $s_i^0 = 0$, while inactive entries satisfy $f_i^0 \approx 0$ and $s_i^0 > 0$.

For a perturbation on one edge, $\beta_k = \beta_k^0 + \eta$, and fixed active set, linearizing the active-set optimality system yields first-order responses

$$f_i(\eta) = f_i^0 + \dot{f}_i \eta + o(\eta), \quad s_i(\eta) = s_i^0 + \dot{s}_i \eta + o(\eta). \quad (60)$$

Candidate breakpoints are where either an active flow hits the activity threshold or an inactive slack hits zero. For the positive direction ($\eta > 0$), these candidates are

$$\eta_i^{+, \text{leave}} = \frac{\varepsilon_{\text{act}} - f_i^0}{\dot{f}_i} \quad \text{for active } i \text{ with } \dot{f}_i < 0, \quad (61)$$


$$\eta_j^{+, \text{enter}} = -\frac{s_j^0}{\dot{s}_j} \quad \text{for inactive } j \text{ with } \dot{s}_j < 0. \quad (62)$$

Analogously, for the negative direction ($\eta < 0$), candidate magnitudes are

$$\eta_i^-, \text{leave} = -\frac{\varepsilon_{\text{act}} - f_i^0}{\dot{f}_i} \quad \text{for active } i \text{ with } \dot{f}_i > 0, \quad (63)$$

$$\eta_j^-, \text{enter} = \frac{s_j^0}{\dot{s}_j} \quad \text{for inactive } j \text{ with } \dot{s}_j > 0. \quad (64)$$

Fig. 10 Transfer test of the raw local linear first-breakpoint predictor. Each point corresponds to one edge-direction sample; color indicates graph family and marker indicates demand/beta regime. The dashed line marks perfect agreement ($\hat{r} = r^{\text{obs}}$). Deviations from the diagonal quantify mismatch between local linear prediction and observed active-set stability radius



The predicted directional first breakpoints are the smallest positive candidates,

$$\hat{\eta}_k^+ = \min(\eta_i^{+, \text{leave}}, \eta_j^{+, \text{enter}}), \quad (65)$$

$$\hat{\eta}_k^- = \min(\eta_i^{-, \text{leave}}, \eta_j^{-, \text{enter}}), \quad (66)$$

and the predicted symmetric local radius is

$$\hat{r}_k = \frac{1}{\beta_k^0} \min(\hat{\eta}_k^+, \hat{\eta}_k^-). \quad (67)$$

The observed directional radii are obtained by repeated full equilibrium solves and active-set comparison:

$$r_{k,\pm}^{\text{obs}} = \sup\{r \geq 0 : \mathcal{A}_{k,\pm}(\rho) = \mathcal{A}_0 \ \forall \rho \in [0, r]\}, \quad (68)$$

with $\beta_k = \beta_k^0(1 \pm \rho)$.

Figure 10 compares predicted and observed radii across graph families and demand/beta regimes. The concentration around the diagonal in some regimes indicates local usefulness, while the broad scatter away from the diagonal in others confirms that this first-order predictor should be interpreted as a local heuristic, not a globally reliable extrapolator.

Appendix 4: Intervention-ranking validation across real-world cases

To validate whether SCGC-based ranking predicts realized intervention impact beyond a single network, we run the same stress test protocol on two distinct real-world cases: event-driven traffic in Cologne and daily urban traffic in Potsdam, each evaluated under three demand multipliers. Starting from each calibrated baseline equilibrium, we define a candidate set C as the top-flow edges (Cologne: up to 60 candidates, Potsdam: up to 25 candidates). For each ranking method and budget k , we select the top- k edges from C and apply a positive perturbation to free-flow travel times:

$$\beta'_e = \beta_e(1 + \rho), \quad \rho = 0.01. \quad (69)$$

Table 1 Intervention-validation summary pooled across Cologne and Potsdam over all tested demand multipliers

Method	Mean realized ΔSC (%)	Mean Spearman ρ
SCGC-linear	0.0219	0.995
Flow	0.0164	0.578
Betweenness	0.0092	0.041
Random	0.0091	–

For each method $m \in \{\text{SCGC}, \text{flow}, \text{betweenness}, \text{random}\}$ and budget k , we choose a set $S_k^{(m)}$ of k edges according to the method ranking and recompute the full multicommodity equilibrium. The realized impact is measured as

$$\Delta SC(S_k^{(m)}) = 100 \cdot \frac{SC(\beta') - SC(\beta)}{SC(\beta)}. \quad (70)$$

Higher values indicate that the selected edges are more critical (larger system damage under the same perturbation budget). For the random baseline, we average over repeated random top- k selections (20 repeats in Cologne, 16 in Potsdam).

For SCGC, the local predicted single-edge impact is

$$\widehat{\Delta SC}_e = \frac{\partial SC}{\partial \beta_e} \Delta \beta_e, \quad (71)$$

with $\Delta \beta_e = \rho \beta_e$ in the stress test setting. We then compare this first-order ranking to realized one-edge impacts and to multi-edge top- k outcomes.

Across all tested real-world budgets and demand levels (42 budget points in total, i.e., $3 \times (8 + 6)$ from Cologne budgets $k = 1, \dots, 8$ and Potsdam budgets $k = 1, \dots, 6$), SCGC attains a strict 100% budget-wise win rate over flow, betweenness, and random baselines. Pooling both cities and all demand multipliers, the mean realized social cost increase is 0.0219% for SCGC (95% bootstrap CI 0.0180–0.0259), compared with 0.0164% for flow (0.0133–0.0199), 0.0092% for betweenness (0.0078–0.0106), and 0.0091% for random (0.0070–0.0113). The mean SCGC–flow gap is 0.0054 percentage points (95% bootstrap CI 0.0046–0.0063), i.e., a relative gain of about 33% over flow. For single-edge perturbations, pooled rank correlation between predicted and observed impact is $\rho \approx 0.995$ for SCGC (95% bootstrap CI 0.989–0.999), versus $\rho \approx 0.578$ for flow (0.543–0.609) and $\rho \approx 0.041$ for betweenness (–0.011–0.088). These results support the claim that SCGC captures marginal system sensitivity more faithfully than topology-only or flow-only baselines across heterogeneous urban operating conditions (Table 1; Figs. 11, 12).

Fig. 11 Case-level intervention curves for Cologne: realized social cost change ΔSC under +1% perturbations of selected β_e values as a function of intervention budget k . Curves show median performance across demand multipliers; shaded bands indicate interquartile ranges





Fig. 12 Case-level intervention curves for Potsdam, using the same protocol as in Cologne. SCGC remains the best-performing ranking across all tested budgets and demand multipliers



Appendix 5: Constructing real-world networks

We construct directed graphs using road network data from the OpenStreetMap (OSM) project, leveraging the Python package OSMnx [49]. OSM provides detailed geographical and infrastructural information about roads in a specified region, including their location, length, speed limit, number of lanes, and other metadata.

Real-world graph networks can become very large. For example, consider Berlin, the largest city in Germany: its road network comprises $N = 5711$ nodes and $M = 11,882$ edges. If every node were allowed to act as an origin in the traffic assignment problem (TAP), i.e., $|W| = 5711$, the resulting system would involve $5711 \times 11,882$ flow variables and 5711×5711 constraints. Solving such a system directly is computationally infeasible.

To address this challenge, we apply a series of spatial and topological simplification steps that preserve the essential structure of the network while making the problem tractable.

First, we exclude minor roads, in particular residential streets, which contribute little to large-scale traffic flow but significantly increase network size and complexity. This step effectively filters the graph to include only major road classes such as motorways, trunk roads, and primary or secondary streets.

Second, we use the `osmnx.simplification consolidate_intersections` function to merge nearby nodes representing the same physical intersection. Divided roads, roundabouts, or complex intersections are often represented in OSM by multiple closely spaced nodes, which unnecessarily inflate the number of variables and constraints in the TAP formulation. Consolidating these nodes reduces network complexity while preserving its essential connectivity and flow structure. The tolerance parameter for node consolidation is set to 100 m, reflecting typical urban intersection spacing.

Finally, we restrict the set of destination nodes W to a strategically selected subset that captures the essential travel demand structure. For example, we may consider centroids of administrative districts, major transportation hubs, or points of interest such as football stadiums. By focusing on a representative set of destinations rather than the full node set, we drastically reduce the dimensionality of the problem while retaining meaningful large-scale flow patterns.

These preprocessing steps yield a simplified yet topologically representative network, enabling efficient computation of user equilibrium flows on real-world road networks without sacrificing key structural characteristics.

Appendix 6: Accessing the cost function

To compute the cost function (see Eq. (2)), we infer the parameters α_e and β_e for each road segment $e \in E$. These parameters are derived directly from the physical and infrastructural characteristics of the road network obtained from OpenStreetMap.

The free-flow travel time β_e is determined from the length l_e of the road segment and its speed limit v_e , ensuring that baseline travel times reflect realistic traffic conditions under uncongested flow:

$$\beta_e = \frac{l_e}{v_e}.$$

To represent congestion effects, we compute the congestion factor α_e using the road length, speed limit, and number of lanes m_e . Additionally, we define the edge capacity $x_e^{\max}$ and the maximum travel time $t_e^{\max}$ based on a reference walking speed, which acts as an upper bound in extreme congestion scenarios.

The congestion factor α_e is given by

$$\alpha_e = \frac{t_e^{\max} - \beta_e}{5 x_e^{\max}},$$

where $t_e^{\max} = \frac{l_e}{v_{\text{walk}}}$ represents the travel time at walking speed and $x_e^{\max}$ is the effective capacity of the road segment. The scaling factor 5 controls the slope of the cost function and can be tuned to adjust how quickly travel times increase with flow.

Appendix 7: Constructing the OD matrix

To generate the OD matrix p_n^w , we use the population distribution of the region under study. Specifically, we rely on the Global Human Settlement Layer (GHSL) [50] to assign a population count to each node in the network. This is achieved by computing Voronoi polygons around the nodes and intersecting them with the GHSL raster, allowing us to allocate population values proportionally to their spatial influence area (see [43, 51] for a detailed explanation of this process).

Based on these node-level population values, we assume that the likelihood of trips originating from or traveling to a node is solely determined by its population. As a result, the probability of traveling from one node to another is independent of the distance between them. While more sophisticated approaches could incorporate distance-based effects, such as a gravity model, this simplification is justified by the typical spatial scale of our networks (urban areas), where trip length distributions are relatively homogeneous and distance effects are less pronounced.

The OD matrix is generated by distributing a fixed fraction γ of each node's population across a selected set of destination nodes in proportion to their population. This results in

$$p_n^w = \gamma \frac{P_w}{\sum_{w' \in W} P_{w'}} P_n,$$

where P_n denotes the population at origin node n and P_w the population at destination node w .

References

1. D. Helbing, Traffic and related self-driven many-particle systems. *Rev. Mod. Phys.* **73**, 1067–1141 (2001). <https://doi.org/10.1103/RevModPhys.73.1067>
2. K. Nagel, M. Schreckenberg, A cellular automaton model for freeway traffic. *J. Phys. I* **2**, 2221–2229 (1992). <https://doi.org/10.1051/jp1:1992277>
3. B.S. Kerner, *The Physics of Traffic* (Springer, Heidelberg, Berlin, 2004)
4. D. Chowdhury, L. Santen, A. Schadschneider, Statistical physics of vehicular traffic and some related systems. *Phys. Rep.* **329**, 199–329 (2000). [https://doi.org/10.1016/S0370-1573\(99\)00117-9](https://doi.org/10.1016/S0370-1573(99)00117-9)
5. D. Helbing, L. Buzna, A. Johansson, T. Werner, Self-organized pedestrian crowd dynamics: experiments, simulations, and design solutions. *Transp. Sci.* **39**(1), 1–24 (2005). <https://doi.org/10.1287/trsc.1040.0108>
6. H. Haken, Synergetics: an overview. *Rep. Prog. Phys.* **52**(5), 515 (1989). <https://doi.org/10.1088/0034-4885/52/5/001>
7. S. Camazine, J.-L. Deneubourg, N.R. Franks, J. Sneyd, G. Theraulaz, E. Bonabeau, *Self-Organization in Biological Systems* (Princeton University Press, Princeton, 2001)
8. D. Braess, Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* **12**, 258–268 (1968). <https://doi.org/10.1007/BF01918335>
9. T. Roughgarden, É. Tardos, How bad is selfish routing? *J. ACM* **49**, 236–259 (2002). <https://doi.org/10.1145/506147.506153>
10. R. Steinberg, W.I. Zangwill, The prevalence of Braess' paradox. *Transp. Sci.* **17**(3), 301–318 (1983). <https://doi.org/10.1287/trsc.17.3.301>
11. J.G. Wardrop, Road paper. Some theoretical aspects of road traffic research. *Proc. Inst. Civ. Eng.* **1**(3), 325–362 (1952). <https://doi.org/10.1680/ipeds.1952.11259>
12. J.R. Correa, A.S. Schulz, N.E. Stier-Moses, A geometric approach to the price of anarchy in nonatomic congestion games. *Games Econ. Behav.* **64**(2), 457–469 (2008). <https://doi.org/10.1016/j.geb.2008.01.001>

13. H. Youn, M.T. Gastner, H. Jeong, Price of anarchy in transportation networks: efficiency and optimality control. *Phys. Rev. Lett.* **101**(12), 128701 (2008). <https://doi.org/10.1103/PhysRevLett.101.128701>
14. Y. Sheffi, *Urban Transportation Networks* (Prentice Hall, Englewood Cliffs, NJ, 1985)
15. M. Patriksson, *The Traffic Assignment Problem* (VSP, Utrecht, 1994)
16. M. Florian, D. Hearn, Network equilibrium models and algorithms. In *Handbook of Transportation Science*, edited by R. W. Hall (Kluwer Academic Publishers, Boston, 1995).
17. M.J. Smith, The existence, uniqueness and stability of traffic equilibria. *Transp. Res. Part B: Methodol.* **13**(4), 295–304 (1979). [https://doi.org/10.1016/0191-2615\(79\)90022-5](https://doi.org/10.1016/0191-2615(79)90022-5)
18. L.J. LeBlanc, E.K. Morlok, W.P. Pierskalla, An efficient approach to solving the road network equilibrium traffic assignment problem. *Transp. Res.* **9**(5), 309–318 (1975). [https://doi.org/10.1016/0041-1647\(75\)90030-1](https://doi.org/10.1016/0041-1647(75)90030-1)
19. H. Bar-Gera, Origin-based algorithm for the traffic assignment problem. *Transp. Sci.* **36**(4), 398–417 (2002). <https://doi.org/10.1287/trsc.36.4.398.549>
20. R.B. Dial, A path-based user-equilibrium traffic assignment algorithm that obviates path storage and enumeration. *Transp. Res. Part B: Methodol.* **40**(10), 917–936 (2006). <https://doi.org/10.1016/j.trb.2006.02.008>
21. A. Nagurney, Congested urban transportation networks and emission paradoxes. *Transp. Res. Part D* **5**(2), 145–151 (2000). [https://doi.org/10.1016/S1361-9209\(99\)00031-0](https://doi.org/10.1016/S1361-9209(99)00031-0)
22. A. Nagurney, Q. Qiang, A network efficiency measure for congested networks. *Europhys. Lett. (EPL)* **79**(3), 38005 (2007). <https://doi.org/10.1209/0295-5075/79/38005>
23. E. Jenelius, T. Petersen, L.-G. Mattsson, Importance and exposure in road network vulnerability analysis. *Transp. Res. Part A: Policy Pract.* **40**(7), 537–560 (2006). <https://doi.org/10.1016/j.tra.2005.11.003>
24. E.I. Pas, S.L. Principio, Braess' paradox: some new insights. *Transp. Res. Part B: Methodol.* **31**(3), 265–276 (1997). [https://doi.org/10.1016/S0191-2615\(96\)00024-0](https://doi.org/10.1016/S0191-2615(96)00024-0)
25. M. Passacantando, G. Gnecco, Y. Hadas, M. Sanguineti, Braess' paradox: a cooperative game-theoretic point of view. *Networks* **78**(3), 264–283 (2021). <https://doi.org/10.1002/net.22018>
26. P. Newman, J. Kenworthy, *The End of Automobile Dependence* (Island Press, Washington, DC, 2015)
27. F. Creutzig et al., Transport: A roadblock to climate change mitigation? *Science* **350**(6263), 911–912 (2015). <https://doi.org/10.1126/science.aac8033>
28. A. Downs, The law of peak-hour expressway congestion. *Traffic Q.* **16**, 393–409 (1962)
29. G. Duranton, M.A. Turner, The fundamental law of road congestion: evidence from US cities. *Am. Econ. Rev.* **101**(6), 2616–2652 (2011). <https://doi.org/10.1257/AER.101.6.2616>
30. H. Yang, H.-J. Huang, *Mathematical and Economic Theory of Road Pricing* (Elsevier, Amsterdam, 2005)
31. K.A. Small, E.T. Verhoef, *The Economics of Urban Transportation* (Routledge, London, 2007)
32. A. de Palma, R. Lindsey, Traffic congestion pricing methodologies and technologies. *Transp. Res. Part C* **19**(6), 1377–1399 (2011). <https://doi.org/10.1016/j.trc.2011.02.010>
33. US Bureau of Public Roads. Office of Planning. Urban Planning Division, *Traffic Assignment Manual for Application with a Large, High Speed Computer* (US Department of Commerce, Washington, DC, 1964)
34. E. Koutsoupias, C. Papadimitriou. Worst-case equilibria. In *Annual Symposium on Theoretical Aspects of Computer Science*, edited by C. Meinel and S. Tison (Springer, 1999), pp. 404–413. https://doi.org/10.1007/3-540-49116-3_38
35. B. Schäfer, T. Pesch, D. Manik, J. Gollenstede, G. Lin, H.-P. Beck, D. Witthaut, M. Timme, Understanding Braess' paradox in power grids. *Nat. Commun.* **13**(1), 5396 (2022). <https://doi.org/10.1038/s41467-022-32917-6>
36. D. Manik, D. Witthaut, M. Timme. Predicting Braess' paradox in supply and transport networks. *arXiv preprint. arXiv: 2205.14685* (2022)
37. R.K. Ahuja, T.L. Magnanti, J.B. Orlin, *Network Flows: Theory, Algorithms and Applications* (Prentice Hall, Englewood Cliffs, NJ, 1995)
38. S. Dafermos, Traffic equilibrium and variational inequalities. *Transp. Sci.* **14**(1), 42–54 (1980). <https://doi.org/10.1287/trsc.14.1.42>
39. S. Nguyen, An algorithm for the traffic assignment problem. *Transp. Sci.* **8**(3), 203–216 (1974). <https://doi.org/10.1287/trsc.8.3.203>
40. M. Beckmann, C.B. McGuire, C.B. Winsten, Studies in the economics of transportation. Technical report (1956)
41. MOSEK ApS (2024), *MOSEK Optimizer API for Python 11.0.4*. URL <https://docs.mosek.com/latest/pythonapi/index.html>
42. B. Stellato, G. Banjac, P. Goulart, A. Bemporad, S. Boyd, OSQP: an operator splitting solver for quadratic programs. *Math. Program. Comput.* **12**(4), 637–672 (2020). <https://doi.org/10.1007/s12532-020-00179-2>
43. J. Wassmer, B. Merz, N. Marwan, Resilience of transportation infrastructure networks to road failures. *Chaos: Interdiscip. J. Nonlinear Sci.* (2024). <https://doi.org/10.1063/5.0165839>
44. H. Haken, *Synergetics, An Introduction*. Springer Series in Synergetics (1983). <https://doi.org/10.1007/978-3-642-88338-5>
45. T. Firat, D. Eroglu, Data-driven modeling of traffic flow in macroscopic network systems. *Chaos* **35**(9), 093134 (2025). <https://doi.org/10.1063/5.0285930>
46. S.C. Dafermos, F.T. Sparrow, Traffic assignment problem for a general network. *J. Res. Natl Bur. Stand. Sect. B: Math. Sci.* **73B**(2), 91 (1969). <https://doi.org/10.6028/jres.073B.010>
47. M.K. Kozlov, S.P. Tarasov, L.G. Khachiyan, Polynomial solvability of convex quadratic programming. *Doklady Akademii Nauk*, **248**, 1049–1051 (1979)

48. A. Krylatov, V. Zakharov, T. Tuovinen, *Optimization Models and Methods for Equilibrium Traffic Assignment*, (Springer, Cham, 2020)
49. G. Boeing, Modeling and analyzing urban networks and amenities with OSMnx. *Geogr. Anal.* **57**(4), 567–577 (2025). <https://doi.org/10.1111/gean.70009>
50. M. Schiavina, M. Melchiorri, M. Pesaresi, P. Politis, S. M. Carneiro Freire, L. Maffenini, P. Florio, D. Ehrlich, K. Goch, A. Carioli, J. Uhl, P. Tommasi, T. Kemper, GHSL data package 2023 (GHS P2023) (2023). <https://publications.jrc.ec.europa.eu/repository/handle/JRC133256>
51. J. Wassmer, S. Bryant, P. Schimansky, L.T. Keegan, M. Pregnolato, J. Kurths, N. Marwan, B. Merz, Unveiling hidden risks in healthcare from flood-induced transportation disruption in Germany. *Commun. Earth Environ.* **6**(1), 676 (2025). <https://doi.org/10.1038/s43247-025-02645-y>