Optimal Intervention in Traffic Networks

(Article begins on next page)

Optimal intervention in traffic networks

Leonardo Cianfanelli[1], Giacomo Como[1,2], Asuman Ozdaglar[3], and
Francesca Parise[4]

[1]Dipartimento di Scienze Matematiche, Politecnico di Torino
[2]Department of Automatic Control, Lund University, BOX 118, SE-22100, Lund, Sweden
[3]Department of Electrical Engineering and Computer Science, MIT
[4]Department of Electrical and Computer Engineering, Cornell University

leonardo.cianfanelli@polito.it, giacomo.como@polito.it, asuman@mit.edu, fp264@cornell.edu
https://sites.google.com/view/leonardo-cianfanelli/, https://staff.polito.it/giacomo.como/,
https://asu.mit.edu/, https://sites.coecis.cornell.edu/parise/

Abstract

We present an efficient algorithm to identify which edge should be
improved in a traffic network to minimize the total travel time. Our
main result is to show that it is possible to approximate the variation of
total travel time obtained by changing the congestion coefficient of any
given edge, by performing only local computations, without the need of
recomputing the entire equilibrium flow. To obtain such a result, we re-
formulate our problem in terms of the effective resistance between two
adjacent nodes and suggest a new approach to approximate such effective
resistance. We then study the optimality of the proposed procedure for
recurrent networks, and provide simulations over synthetic and real trans-
portation networks.

Keywords: Wardrop equilibrium; transportation networks; effective resis-
tance; network design problem.

1 Introduction

Due to increasing populations living in urban areas, many cities are facing the
problem of traffic congestion, which leads to increasing levels of pollution and
massive waste of time and money (Schrank et al. [29]). The problem of mitigat-
ing congestion has been tackled in the literature from two main perspectives.
One approach is to indirectly influence the behaviour of the drivers, for in-
stance by road tolling (e.g., in Brown and Marden [7], Fleischer et al. [16],
Zhao and Kockelmann [39], Cole et al. [10]), information design (e.g., in Das
et al. [12], Meigs et al. [25], Wu et al. [35, 36]) or lottery rewards (Yue et al.
[38]), to minimize inefficiencies due to the autonomous uncoordinated decisions
of agents. A second approach is to intervene on the transportation network

directly, by building new roads or enlarging existing ones. The corresponding *network design problem* (i.e., the problem of optimizing the intervention on a transportation network subject to some budget constraints, see e.g. LeBlanc [22]) is very challenging because of its bi-level nature (Farahani et al. [15]), i.e., it involves a network intervention optimization problem given the flow distribution for that particular network. We assume that each edge of the network is endowed with a delay function and the flow distributes according to a Wardrop equilibrium, taking paths with minimum delay, defined as the sum of the delays of edges along a path (see Beckmann et al. [3] and Wardrop [34]). Characterization of Wardrop equilibrium is used to construct the lower level of the bilevel network design problem.

In this paper we study a special class of network design problem (NDP), where the planner can improve the delay function of a single edge. Our objective is to strike a balance between a model that is simple enough to guarantee tractable analysis, yet rich enough to allow insights for more general classes of NDPs. For this class of NDPs (with E denoting the number of edges), our first main result provides an analytical characterization of the cost variation corresponding to an intervention on a particular edge under a regularity assumption, which states that the edges that carry positive flow remain unchanged with an intervention. This assumption, tipically used in the traffic equilibrium literature (see Steinberg and Zangwill [31], Dafermos and Nagurney [11]) leads to characterization of Wardrop equilibria using a system of linear equations and enables representing edge interventions as rank-1 perturbations of the system. We show that this assumption is satisfied provided that the total incoming flow to the network is large enough, which may be of independent interest. We exploit the structure of our characterization and linearity of delay functions to express the cost variation using the effective resistance of an edge (i.e., between the enpoints of the edge), defined with respect to a related resistor network. Computing the effective resistance of a single edge requires the solution of a linear system with a matrix with dimension scaling with the size of the network (we indistinctly refer to the size of the network as the cardinality of the node and the edge sets, implicitly assuming that transportation networks are sparse in a such a way that the average degree of the nodes is independent of the number of nodes, inducing then a proportionality between the number of nodes and edges). Hence, solving the NDP requires the solution of E of these problems. Since this can be computationally untractable for large networks, our second main result proposes a method based on Rayleigh's monotonicity laws to approximate the effective resistance of each edge with a number of iterations independent of the network size, thus leading to a significant reduction of complexity. The key idea is that the effective resistance between two adjacent nodes i and j depends mainly on the local structure of the network around the two nodes (i.e., the set of nodes $\mathcal{N}_{\leq d}$ that are at distance no greater than a small given constant d from at least one of i and j), and may therefore be approximated by performing only local computations. Since for networks with bounded degrees (as typical in traffic networks, think for instance of the bidimensional square grid) the size of $\mathcal{N}_{\leq d}$

does not scale with the network size, we can guarantee that the approximation error and computational complexity of our method also do not scale. Finally, we give sufficient conditions on the network under which the approximation error vanishes asymptotically in the limit of infinite networks, and we show simulations on synthetic and real networks.

In our work we consider a special case of a NDP. These problems have been formalized in the last decades via many different formulations. Both *continuous* network design problems (e.g., Chiou [8], Li et al. [24], Wang et al. [32]), where the budget can be allocated continuously among the edges, and *discrete* formulations, in which the decision variables include which new roads to build (Gao et al. [19]), how many lanes to add to existing roads (Wang et al. [33]), or a mix of those two problems (Poorzahedy and Rouhani [26]), have been considered in the literature, together with *dynamical* formulations (Fontaine and Minner [17]), and formulations where the optimum is achieved by removing, instead of adding, edges, because of Braess paradox, as in Roughgarden [28] and Fotakis et al. [18]. For comprehensive surveys on the literature on NDP we refer to Yang and Bell [37] and Farahani et al. [15]. We stress that most of the literature focused on finding time polynomial algorithms to approximately solve NDPs in their most general form. As noted above, we instead consider a problem that can be solved with a polynomial algorithm by simply enumerating all the candidate edges and computing the cost corresponding to the intervention on each of those edges. Our main contribution is to define a simplified, more intuitive and tractable approach to solve such a design problem in quasilinear time instead of polynomial, as well as providing intuition and a complete new formulation. For the future we aim at extending our techniques to more general cases, like the multiple interventions case. Our work is related to Steinberg and Zangwill [31] and Dafermos and Nagurney [11], where the authors investigate the sign of total travel time variation when a new path is added to a two-terminal network, under similar assumptions to ours, providing sufficient conditions under which Braess paradox arises. In our work we instead suggest an efficient algorithm to select the best edge to improve. As mentioned, a key step of our approach is to reformulate the NDP in terms of a resistance problem. We do this in two steps: first we reformulate the NDP in terms of random walks over a network (following a similar approach as in Rebeschini and Tatikonda [27], where however the problem of finding conditions under which a perturbation of the incoming flow to a small part of the network leads to local perturbation of Wardrop equilibrium), then following standard literature we reformulate the random walk problem as an electrical one (see e.g. Doyle and Snell [13]). To summarize, the contribution of this paper is two-fold. From a methodological perspective, we provide a method to locally upper and lower bound the effective resistance between adjacent nodes, which may be of a separate interest beyond traffic applications. From the network design perspective, we provide a new formulation of the design problem in terms of resistor networks, and we exploit our methodological result to approximate in an efficient manner a simplified version of the design problem where a single edge can be improved. For the future

3

we aim at extending our methods to the case of multiple interventions. From a methodological perspective, it is worthwhile mentioning that the equivalence between Wardrop and electric flows has been already investigated in Klimm and Warode [21].

The remainder of the paper is organized as follows. In Section 2 we define the model and formulate the problem as a bi-level programming. In Section 3 we first translate the problem into a single-level program, and then rephrase it in terms of resistor networks. In Section 4 we provide our method to approximate effective resistance between neighbors. In Section 5 we analyze the asymptotic behaviour of the bounds in the limit of infinite networks. In Section 6 we show some simulations over relevant networks. Finally, in the conclusive section, we summarize the work and discuss future research lines.

1.1 Notation

$\mathcal{G} = (\mathcal{N}, \mathcal{E})$ denotes a two-terminal strongly connected directed network with origin o and destination d, where $\mathcal{N}$ and $\mathcal{E}$ are respectively the node and the edge sets. Let $\mathcal{P}$ denote the set of paths from o to d, and N, E and P denote the cardinality of $\mathcal{N}$, $\mathcal{E}$ and $\mathcal{P}$, respectively. We allow for multiple edges between the same pair of nodes. We do not allow for selfloops, since they are not relevant in traffic networks. δ_i, $\mathbb{1}$ and $\mathbb{I}$ denote the unitary vector with 1 in position i and 0 in all the other positions, the column vector of all ones, and the identity matrix, respectively, where the size of them may be deduced from the context. A^T and v^T denote the transpose of matrix A and vector v, respectively. For simplicity of notation we use A_{ij}^{-1} instead of $(A^{-1})_{ij}$.

2 Model and problem formulation

2.1 Model

The transportation network is modeled as a two-terminal strongly connected directed network $\mathcal{G} = (\mathcal{N}, \mathcal{E})$. Let $\tau > 0$ denote the throughput from the origin o to the destination d, and $\nu = \tau(\delta_o - \delta_d) \in \mathbb{R}^N$ denote the net inflow to the network. An admissible path flow is a vector $z \in \mathbb{R}^P$ satisfying the non-negativity and conservation of mass constraints

$$z \geq 0, \quad z^T \mathbb{1} = \tau. \tag{1}$$

Let $R \in \mathbb{R}^{E \times P}$ denote the edge-path incidence matrix, with entries $R_{lp} = 1$ if the edge l belongs to the path p or 0 otherwise. The path flow induces a unique edge flow $f \in \mathbb{R}^E$ via

$$f = Rz. \tag{2}$$

Let $B \in \mathbb{R}^{N \times E}$ denote the node-edge incidence matrix, with entries $B_{nl} = 1$ if the node n is the tail of the edge l, -1 if n is the head of l, or 0 otherwise. The

constraints can be reformulated in terms of edge flows as

$$f \geq 0, \quad Bf = \nu.$$

Every edge is endowed with a delay function, which is assumed affine, non-negative and strictly increasing,

$$d_l(f_l) = a_l f_l + b_l, \quad a_l > 0, \quad b_l \geq 0, \quad \forall l \in \mathcal{E}.$$

The cost of a path p, under flow distribution f, is

$$c_p(f) = \sum_{l \in \mathcal{E}} R_{lp} d_l(f_l), \tag{3}$$

which is the sum of the delays of the edges belonging to that path. We also define $A \in \mathbb{R}^{E \times E}$ and $b \in \mathbb{R}^E$ as

$$A := \begin{bmatrix} a_1 & 0 & \dots & 0 \\ 0 & a_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & a_E \end{bmatrix}, \quad b := \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_E \end{bmatrix}.$$

Definition 2.1 (Affine routing game). *An affine routing game is a quadruple* $(\mathcal{G}, A, b, \nu)$.

A Wardrop equilibrium is a flow distribution such that no one has incentive in changing his path.

Definition 2.2 (Wardrop equilibrium). *A path flow* z^*, *with associated edge flow* f^* *as defined in* (2), *is a Wardrop equilibrium if for every path* p

$$z_p^* > 0 \implies c_p(f^*) \leq c_q(f^*), \quad \forall q \in \mathcal{P}.$$

It is shown in Beckmann et al. [3] that an edge flow f^* is a Wardrop equilibrium of a routing game if and only if it solves the following minimization problem:

$$\begin{aligned} \underset{f}{\text{minimize}} \quad & \sum_{l \in \mathcal{E}} \int_0^{f_l} d_l(s) ds \\ \text{subject to} \quad & f \geq 0, Bf = \nu. \end{aligned} \tag{4}$$

Since the delay functions are assumed strictly increasing, the objective function is strictly convex and the Wardrop equilibrium f^* is unique.

We now define the social cost, which is the total travel time at the equilibrium.

Definition 2.3 (Social cost). *Let* f^* *be the unique Wardrop equilibrium of an affine routing game* $(\mathcal{G}, A, b, \nu)$. *The social cost is*

$$C(f^*) = \sum_{l \in \mathcal{E}} f_l^* d_l(f_l^*).$$

The social cost can be interpreted by a planner that aims at minimizing the overall congestion as a measure of performance of the transportation network.

2.2 Problem formulation

We consider a problem in which a planner can intervene on the network with the goal of minimizing the social cost. We propose as intervention to rescale the slope of one edge l by a scaling parameter $\kappa > 1$, so that the slope of the edge l gets reduced from a_l to $\tilde{a}_l = a_l/\kappa$. This intervention may correspond for instance to adding a new lane in a street. In fact, every intervention on a single edge may be seen as a rank-1 perturbation of the system and may be handled by our method (see §3.1). We aim at identifying which edge should be selected by the planner to minimize the social cost. Let $f^*(l)$ and $C(f^*(l))$ denote the Wardrop equilibrium when the slope of the edge l is rescaled, and the corresponding social cost, respectively. Hence, the problem can be expressed as follows.

Problem 1. *Let $(\mathcal{G}, A, b, \nu)$ be an affine routing game and $\kappa > 1$ be the scaling parameter. Find edge l^* such that*

$$l^* \in \underset{l \in \mathcal{E}}{\arg\min} \quad C(f^*(l)).$$

We stress the fact that Problem 1 is bi-level, in the sense that the planner optimizes the intervention over the edge set, but the cost function is a function of the Wardrop equilibrium f^*, which in turn is the solution of the optimization problem (4).

Problem 1 can be solved by a brute force approach, by enumerating all the edges and computing the corresponding equilibrium $f^*(l)$ by solving the convex program (4) with $d_l(f_l) = \tilde{a}_l f_l + b_l$ instead of $d_l(f_l)$. In this work we propose a method that, given f^* before the intervention (which is assumed to be observable and therefore known) and other electrical quantities computed on a resistance network related to the original unperturbed traffic network, provides an upper and lower bound to $C(f^*(l))$ with a computational complexity that does not scale with the size of the network. The main idea is that the effect of perturbing an edge may be well approximated by looking at a local portion of the network. Our method works under the assumption that the network is sparse in such a way that the average degree of the nodes does not depend on the size of the network, and under the assumption that the set of the used edges does not change after the intervention. The first assumption is suitable for traffic networks, and the second one is standard in the literature on intervention in traffic networks (see Steinberg and Zangwill [31] and Dafermos and Nagurney [11]). We provide a more detailed discussion on this assumption in §3.2.

3 Two equivalent reformulations

In this section we provide two equivalent formulations for Problem 1. The first formulation is based on the fact that modifying the slope of one edge is equivalent to introducing a rank-1 perturbation in the KKT conditions of (4) (for

an in depth discussion on KKT conditions we refer to Boyd and Vandenberghe [5]). The second formulation gives an interpretation using resistor networks.

3.1 KKT formulation

Let us introduce the dual variables $\lambda^*_{(i,j)}$ associated to $f^*_{(i,j)} \geq 0$ and γ^* associated to the constraint $Bf = \nu$. The KKT conditions of (4) are:

$$\begin{cases} a_{(i,j)} f^*_{(i,j)} + b_{(i,j)} + \gamma^*_j - \gamma^*_i - \lambda^*_{(i,j)} = 0 & \forall (i,j) \in \mathcal{E}, \\ \sum_{j:(j,i)\in\mathcal{E}} f^*_{(j,i)} - \sum_{j:(i,j)\in\mathcal{E}} f^*_{(i,j)} + \nu_i = 0 & \forall i \in \mathcal{N}, \\ \lambda^*_{(i,j)} f^*_{(i,j)} = 0 & \forall (i,j) \in \mathcal{E}, \\ \lambda^*_{(i,j)} \geq 0 & \forall (i,j) \in \mathcal{E}, \\ f^*_{(i,j)} \geq 0 & \forall (i,j) \in \mathcal{E}. \end{cases} \tag{5}$$

The third condition is known as complementary slackness, and implies that all the edges such that $\lambda^*_e > 0$ are not used at the equilibrium, i.e. $f^*_e = 0$. Let $\mathcal{E}_+$ denote the set of such edges. Thus, the edges in $\mathcal{E}_+$ and the last three conditions of (5) can be removed, without affecting the solution of (5). With a slight abuse of notation, from now on let $\mathcal{E}$ denote $\mathcal{E} \backslash \mathcal{E}_+$. Thus, the KKT conditions become:

$$\begin{cases} a_{(i,j)} f^*_{(i,j)} + b_{(i,j)} + \gamma^*_j - \gamma^*_i = 0 & \forall (i,j) \in \mathcal{E}, \\ \sum_{j:(j,i)\in\mathcal{E}} f^*_{(j,i)} - \sum_{j:(i,j)\in\mathcal{E}} f^*_{(i,j)} + \nu_i = 0 & \forall i \in \mathcal{N}, \end{cases} \tag{6}$$

where the constraint $f^*_{i,j} \geq 0$ can now be removed since the solution of (6) gives $f^*_{i,j} \geq 0$ for every edge $(i,j) \notin \mathcal{E}_+$. Without loss of generality, we order the nodes in such a way that the origin o and the destination d are the first and the last node respectively. Observe from (6) that the optimal flows $f^*_{(i,j)}$ depend on γ^* only via the difference $\gamma^*_i - \gamma^*_j$, so that γ^* remains a solution if a constant vector is added to it. This is due to the fact that the matrix B is not full rank. Removing the last row of B is therefore equivalent to setting $\gamma^*_d = 0$. Thus, we define $x \in \mathbb{R}^{N+E-1}$ and $y \in \mathbb{R}^{N+E-1}$ as

$$x := \left[\begin{array}{c} f \\ \gamma_- \end{array} \right], \quad y := -\left[\begin{array}{c} b \\ \nu_- \end{array} \right],$$

where γ_- and ν_- denote respectively γ and ν where the last element of both vectors is removed. Also, $B_- \in \mathbb{R}^{(N-1)\times E}$ denotes the node-edge incidence matrix where the last row is removed. Finally, we define $H \in \mathbb{R}^{(N+E-1)\times(N+E-1)}$ as

$$H := \left[\begin{array}{cc} A & -(B_-)^T \\ -B_- & \mathbb{0} \end{array} \right].$$

With this notation and assuming $\gamma^*_d = 0$, the KKT conditions (6) become:

$$Hx^* = y. \tag{7}$$

7

Because we assumed $\gamma_{\mathrm{d}}^* = 0$, x^* is unique, and

$$x^* = H^{-1}y. \tag{8}$$

The invertibility of H follows from the invertibility of A (the delays are strictly increasing) and from the invertibility of $Q := B_- A^{-1}(B_-)^T$ (see Horn and Johnson [20]), which we prove in the proof of Theorem 1. Let $A(l)$ and $H(l)$ denote the matrix A and H corresponding to the intervention on edge l. The optimal lagrangian multipliers γ^* have an useful interpretation, under the following assumption.

Assumption 1. *Let $\mathcal{E}_+(l)$ be the set of edges e for which $\lambda_e^*(l) > 0$ in the Wardrop equilibrium of $(\mathcal{G}, A(l), b, \nu)$. We assume that $\mathcal{E}_+(l) = \mathcal{E}_+$ for all $l \in \mathcal{E}$.*

The intuition is that under Assumption 1 the KKT conditions (7) before and after the intervention on the edge l involve the same set of edges and differ for the value of a_l only, allowing therefore to handle the intervention as rank-1 perturbation in H. A detailed discussion on such assumption is given in §3.2.

Proposition 1. *For any $l \in \mathcal{E}$ consider the modified game $(\mathcal{G}, A(l), b, \nu)$ obtained by changing the slope of edge l from a_l to $\tilde{a}_l = a_l/\kappa$ and construct the corresponding primal and dual solution $x^*(l)$ as in (8). Then,*

$$C(f^*) - C(f^*(l)) = \tau(\gamma_o^* - \gamma_o^*(l)),$$

where γ_o^* and $\gamma_o^*(l)$ are the $(\mathrm{E}+1)-$th component of x^* and $x^*(l)$ respectively.

Proof. See the Appendix. $\qquad\qquad\square$

Since τ is a given constant of the problem, Proposition 1 states that the goal of the planner should be to select the edge l minimizing $\gamma_o^*(l)$, that is, the optimal lagrangian multiplier of the origin after the intervention on the edge l. Observe that the brute force method requires the computation of the whole vector $x^*(l)$ for every edge l. A natural question is whether it is possible to evaluate $\gamma_o^*(l)$ for every edge l without computing the whole vector $x^*(l)$. We provide a positive answer under Assumption 1 in the next proposition, where the social cost variation is expressed in terms of the scaling parameter κ, the unperturbed equilibrium f^*, and selected elements of H^{-1}.

Proposition 2. *Let $(\mathcal{G}, A, b, \nu)$ be a routing game, $\kappa > 1$ be the scaling parameter and suppose that Assumption 1 holds. Then, the social cost variation corresponding to the intervention on edge l is equal to*

$$C(f^*) - C(f^*(l)) = \tau \frac{(\kappa - 1)a_l H_{(\mathrm{E}+1,l)}^{-1} f_l^*}{(\kappa - 1)a_l H_{ll}^{-1} - \kappa}. \tag{9}$$

Proof. See the Appendix. $\qquad\qquad\square$

In the next section we provide an interpretation to the required elements of H^{-1} in terms of electrical quantities. Before doing that we discuss Assumption 1 in detail.
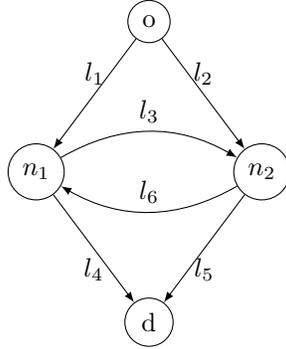
Figure 1: A directed network that is not series-parallel. On this network Assumption 1 is not guaranteed to hold.

3.2 On Assumption 1

In the following we show that Assumption 1 is without loss of generality on series-parallel networks, provided that the throughput is sufficiently high. First, we recall the definition of directed series-parallel networks, and then present the result in Proposition 3.

Definition 3.1 (Directed series-parallel networks). *A two-terminal directed network $\mathcal{G}$ is series-parallel if and only if (i) it is a single edge from the origin to the destination, or (ii) it is the result of connecting two directed series-parallel networks $\mathcal{G}_1$ and $\mathcal{G}_2$ in parallel, by merging o_1 with o_2 and d_1 with d_2, or (iii) it is the result of connecting two directed series-parallel networks $\mathcal{G}_1$ and $\mathcal{G}_2$ in series, by merging d_1 with o_2.*

Proposition 3. *Let $(\mathcal{G}, A, b, \nu)$ be a routing game. If $\mathcal{G}$ is a directed series-parallel network, it exists $\overline{\tau}$ such that for every $\tau \geq \overline{\tau}$, $\mathcal{E}_+ = \emptyset$. Furthermore, if $b = 0$, $\mathcal{E}_+ = \emptyset$ for every $\tau > 0$.*

Proof. See the Appendix. $\qquad \square$

Remark 1. *Proposition 3 immediately implies that Assumption 1 is without loss of generality on directed series-parallel networks provided that $\tau \geq \overline{\tau}$. However, $\overline{\tau}$ depends on A and thus may change after the intervention.*

The next example shows that without the assumption of series-parallel networks Proposition 3 may fail, even in case of linear delays.

Example 1. *Consider the network in Fig. 1, which is not series-parallel. Let us consider $\tau = 1$ and linear delay functions, with $a_1 = a_2 = a_3 = a_4 = a_6 = 1$ and $a_5 = 2$. By some computations,*

$$f_1^* = \frac{6}{11}, \quad f_2^* = \frac{5}{11}, \quad f_4^* = \frac{7}{11}, \quad f_5^* = \frac{4}{11}, \quad f_6^* = \frac{1}{11}, \quad \lambda_3^* = \frac{1}{11}, \quad f_3^* = \lambda_6^* = 0.$$

9

Turning a_5 from 2 to 1/2, we get:

$$f_1^* = \frac{6}{13}, \quad f_2^* = \frac{7}{13}, \quad f_3^* = \frac{1}{13}, \quad f_4^* = \frac{5}{13}, \quad f_5^* = \frac{8}{13}, \quad \lambda_6^* = \frac{1}{13}, \quad f_6^* = \lambda_3^* = 0.$$

Thus, $\mathcal{E}_+(a_5 = 2) = \{l_3\}$ and $\mathcal{E}_+(a_5 = 1/2) = \{l_6\}$. Since $\mathcal{E}_+(a_5 = 2) \neq \mathcal{E}_+(a_5 = 1/2)$, Assumption 1 is violated, even in case of linear delays.

3.3   Electrical formulation

In this section we explore the structure of H^{-1} to give an interpretation of it in terms of electrical quantities. To this end, by the well-known formula for the 2×2 block-matrices inversion (see Bernstein [4]), we get:

$$H^{-1} = \begin{bmatrix} A^{-1} - KQ^{-1}K^T & -KQ^{-1} \\ -Q^{-1}K^T & -Q^{-1} \end{bmatrix}, \tag{10}$$

where $K \in \mathbb{R}^{E \times (N-1)}$ and $Q \in \mathbb{R}^{(N-1) \times (N-1)}$ are

$$K := A^{-1}B_-^T, \qquad Q := B_- A^{-1} B_-^T.$$

From the definitions of B_- and A, it follows

$$K_{l:} = \frac{\delta_i^T - \delta_j^T}{a_l} \quad \forall l = (i,j) \in \mathcal{E}, \tag{11}$$

with the convention that $\delta_d = 0 \cdot \mathbb{1}$ (since we removed the destination), and

$$Q_{ij} = \begin{cases} -\sum_{l \in \{(i,j),(j,i)\}} \frac{1}{a_l} & \text{if } i \neq j \\ \sum_{l \in \partial i} \frac{1}{a_l} & \text{if } i = j. \end{cases} \quad \forall i,j \in \mathcal{N} \setminus d,$$

where ∂i denotes the in and out neighborhood edges of i, that is

$$\partial i := \{l \in \mathcal{E} : B_{il} \neq 0\}.$$

We remark that ∂i includes also edges pointing to the destination. The matrix Q allows for an interpretation in terms of electrical quantities. To this end, let us introduce the notion of resistor network and effective resistance between two nodes.

Definition 3.2. *A resistor network is an undirected weighted network, where the weight matrix W represents a conductance matrix, i.e., $W_{ij} = W_{ji}$ is the conductance between nodes i and j.*

Definition 3.3. *Let $\Delta V = V_i - V_j$ be a difference of potential that is set between nodes i and j on a resistor network. The effective resistance r_{ij} between i and j is*

$$r_{ij} = \frac{\Delta V}{i_i},$$

where i_i denotes the total current flowing from node i under such potential.

We define the resistor network $\mathcal{G}_R$ obtained by making every directed edge of the traffic network $\mathcal{G}$ undirected, with conductance matrix $W \in \mathbb{R}^{N \times N}$

$$W_{ij} := \begin{cases} \sum_{l \in \{(i,j),(j,i)\}} \frac{1}{a_l} & \text{if } i \neq j \\ 0 & \text{if } i = j. \end{cases}$$

Observe that W includes also the destination, is symmetric by construction, since $\mathcal{G}_R$ is undirected, and the coefficients a_l correspond to resistances. Finally, let $D \in \mathbb{R}^{N \times N}$ be the diagonal matrix of degrees in $\mathcal{G}_R$, i.e., $D = \text{diag}(W\mathbb{1})$, and $P = D^{-1}W \in \mathbb{R}^{N \times N}$ the normalized adjacency matrix. The matrix Q may be related to the truncated Laplacian of the resistor network $\mathcal{G}_R$. This is the key point to prove the next theorem. To this end, let us give the following definition.

Definition 3.4. *Let $V \in \mathbb{R}^N$ denote the potential over the nodes of $\mathcal{G}_R$ when the boundary conditions $V_o = 1$ and $V_d = 0$ are imposed. V is an harmonic function (see Levin and Peres [23]), i.e., it satisfies*

$$V_o = 1, \quad V_d = 0, \quad V_i = \sum_{j \in \mathcal{N}} P_{ij} V_j \quad \forall i \in \mathcal{N} \setminus o, d. \tag{12}$$

Theorem 1. *Let $(\mathcal{G}, A, b, \nu)$ be a routing game, $\kappa > 1$ be the scaling parameter, and suppose Assumption 1 holds. The social cost variation corresponding to the intervention on edge $l = (i, j)$ is*

$$\Delta C(l) := C(f^*) - C(f^*(l)) = \tilde{\tau} \frac{f_l^*(V_i - V_j)}{\frac{1}{\kappa - 1} + \frac{r_{ij}}{a_l}}, \tag{13}$$

where $\tilde{\tau}$ is a constant independent of l, and r_{ij} is the effective resistance between nodes i and j in $\mathcal{G}_R$.

Proof. See the Appendix. $\square$

Intuitively, Theorem 1 states the social cost variation after intervention on the edge $l = (i, j)$ depends:

- proportionally on $V_i - V_j$, which may be interpreted as a gradient of relative position from the tail to the head of l, since V_i may be seen as a measure of relative position with respect to the origin and the destination, with $V_o = 1$ at the origin, $V_d = 0$ at the destination, and every other node in between assuming an intermediate value that approaches 1 when the node is closer to the origin and far from the destination, and 0 in the opposite case;

- proportionally on the unpertured flow f_l^*, which is a measure of importance of the edge from the traffic perspective;

- inversely on r_{ij}/a_l, which is a non-negative quantity, no greater than 1; this term is maximum when the edge l is a bottleneck and decreases as the number of alternative paths from i to j increases.

11

In order to solve Problem 1 by the electrical formulation, we need to compute (13) for every edge l. The unperturbed equilibrium f^* is assumed to be observable and therefore given, and the potential V can be derived by solving the linear system (12). Observe that V has to be computed only once. However, the computation of r_{ij} involves the solution of a linear system, and is needed for every edge $l = (i, j)$, so that the solution of Problem 1 by the electrical formulation requires to solve E linear systems (see Aldous and Fill [1]), whereas by formulation (9) we need to compute a row and the diagonal of the inverse of H, which still is computationally onerous when the network is large. In the next section we propose a method to *approximate* the effective resistance between a pair of neighbors that, under a suitable assumption on the sparseness of the network, does not scale with the size of the network, allowing for a more efficient solution to Problem 1.

4 An approximate solution to Problem 1

As seen in the previous section, Problem 1 may be rephrased in terms of electrical quantities over a resistor network. However, even in this formulation the complexity of the problem scales badly because it requires to solve E linear systems whose size grows linearly with N. Since the computational bottleneck is represented by the effective resistance between every pair of adjacent nodes of the network, in the next subsection we propose a computationally cheaper method to approximate this quantity. The main idea of our method is that, even though the effective resistance depends on the entire network, when i and j are adjacent nodes, r_{ij} can be approximated by looking at a local portion of the network only. We then formulate an algorithm to approximately solve Problem 1 by exploiting the approximation of the effective resistance.

4.1 Approximating the effective resistance

Let us introduce the notion of cutting and shorting a network.

Definition 4.1 (Cutting at distance d). *A resistor network $\mathcal{G}_R$ is cut at distance d with respect to a pair of nodes (i, j) if every node at distance greater than d from both i and j is removed, and every edge having at least one end in the set of the removed nodes is removed. Let $\mathcal{G}_{ij}^{U_d}$ and $r_{ij}^{U_d}$ denote such a network and the effective resistance on it, respectively.*

Definition 4.2 (Shorting at distance d). *An resistor network $\mathcal{G}_R$ is shorted at distance d with respect to a pair of nodes (i, j) if all the nodes at distance greater than d from both i or j are shorted together, i.e., an infinite conductance is added between each pair of such nodes. Let $\mathcal{G}_{ij}^{L_d}$ and $r_{ij}^{L_d}$ denote such a network and the effective resistance on it, respectively.*

We refer to Fig. 2 for an example of these techniques applied to a regular grid. We next prove that $r_{ij}^{U_d}$ and $r_{ij}^{L_d}$ are respectively an upper and a lower
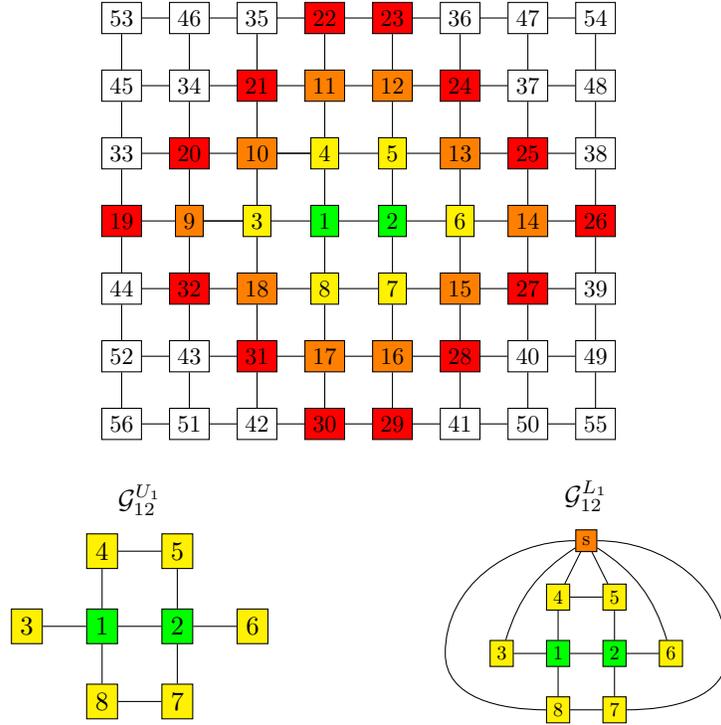
Figure 2: Square grid. Above: the yellow, orange and red nodes are at distance 1, 2 and 3, respectively from the green nodes. Bottom left: cut at distance 1. Bottom right: shorted at distance 1. We stress that in the bottom right network the edges connecting yellow nodes with node s do not have unitary weights.

bound for the effective resistance r_{ij} for every pair of adjacent nodes. To this end, let us introduce the Rayleigh's monotonicity laws.

Lemma 1 (Rayleigh's monotonicity laws (Levin and Peres [23])). *If the resistances of one or more edges are increased, the effective resistance r_{ij} between any two nodes i and j cannot decrease. If the resistances of one or more edges are decreased, r_{ij} cannot increase.*

Proposition 4. *Let $\mathcal{G}_R$ be a resistor network, D_{max} denote the maximal weighted degree of the network, and r_{ij} be the effective resistance between any two neighboring nodes i and j. Then,*

$$r_{ij}^{U_{d_1}} \geq r_{ij}^{U_{d_2}} \geq r_{ij} \geq r_{ij}^{L_{d_2}} \geq r_{ij}^{L_{d_1}}, \quad \forall d_2 > d_1 \geq 1.$$

Moreover,

$$1/D_{max} \leq r_{ij}^{L_d} \leq r_{ij}^{U_d} \leq 1/W_{ij}, \quad \forall d \geq 1. \tag{14}$$

13

Proof. Cutting a network at distance d is equivalent to setting to infinity the resistance of all the edges that have one node at distance greater than d. Shorting a network at distance d is equivalent to setting to zero the resistance between any pair of nodes at distance greater than d. Then, by Rayleigh's monotonicity laws, $r_{ij}^{U_d} \geq r_{ij} \geq r_{ij}^{L_d}$. Similar arguments may be used to show that, if $d_1 < d_2$, $r_{ij}^{U_{d_1}} \geq r_{ij}^{U_{d_2}}$ and $r_{ij}^{L_{d_1}} \leq r_{ij}^{L_{d_2}}$. The right inequality of (14) follows from noticing that, by Rayleigh's monotonicity laws, the effective resistance computed in the network with only nodes i and j, which is equal to $1/W_{ij}$, is an upper bound for $r_{ij}^{U_1}$. The left inequality follows from noticing that the effective resistance on the network in which every node except j is shorted with i, which results in a network with only two nodes and a conductance between i and j not greater than D_{max} (hence, resistance no less than $1/D_{max}$) is a lower bound of $r_{ij}^{L_1}$. $\square$

Proposition 4 states that cutting and shorting a network provides upper and lower bound for the effective resistance. Moreover, the tightness of the bounds is a monotone function of the distance d.

4.2 Our algorithm

Based on the method for approximating the effective resistance, we here propose an algorithm to approximately solve Problem. Our approach is detailed in Algorithm 1.

Algorithm 1:

Input: The resistor network $\mathcal{G}_R = (\mathcal{N}, \mathcal{E}_R, W)$, the rescale parameter κ and the distance $d \geq 1$ used to approximate the effective resistance.

Output: The optimal edge l^{*d} for the intervention.

Compute V by solving the sparse linear system

$$V_o = 1, \quad V_d = 0, \quad V_i = \sum_{j \in \mathcal{N}} P_{ij} V_j \quad \forall i \in \mathcal{N} \setminus o, d;$$

for *each* $l = (i, j) \in \mathcal{E}$ **do**
    Construct $\mathcal{G}_{ij}^{U_d}$ and $\mathcal{G}_{ij}^{L_d}$;
    Compute $r_{ij}^{U_d}$ and $r_{ij}^{L_d}$ on $\mathcal{G}_{ij}^{U_d}$ and $\mathcal{G}_{ij}^{L_d}$.
end

Select l^{*d} such that

$$l^{*d} \in \operatorname*{argmax}_{l=(i,j)\in\mathcal{E}} \quad \Delta C^d(l) := \frac{f_l^*(V_i - V_j)}{\frac{1}{\kappa-1} + \frac{r_{ij}^{U_d}+r_{ij}^{L_d}}{2a_l}}.$$

Note that the performance of Algorithm 1 depends on the choice of the parameter d. Specifically, the higher d is the better is the approximation of the

effective resistance and the closer is the output of Algorithm 1 to the achieving the minimum of Problem 1.

Theorem 2. *Let $\Delta C(l)$ be the cost variation corresponding to intervention on edge $l = (i,j) \in \mathcal{E}$ as given in Theorem 1, $\Delta C^d(l)$ be the cost variation estimated by Algorithm 1 for a given distance $d \geq 1$, and*

$$\epsilon_{ijd} := \frac{r_{ij}^{U_d} - r_{ij}^{L_d}}{a_l}.$$

Then,

$$\left| \frac{\Delta C(l) - \Delta C^d(l)}{\Delta C(l)} \right| \leq \frac{\epsilon_{ijd}}{2\left(\frac{1}{\kappa - 1} + \frac{r_{ij}^{U_d} + r_{ij}^{L_d}}{2a_l} \right)} \leq \frac{\epsilon_{ijd}}{2\left(\frac{1}{\kappa - 1} + \frac{1}{D_{max} \cdot a_l} \right)}$$

Furthermore,

$$\Delta C(l) \geq \tilde{\tau} \frac{f_l^*(V_i - V_j)}{\frac{1}{\kappa - 1} + \frac{r_{ij}^{U_d}}{a_l}}. \tag{15}$$

Proof. See the Appendix. $\qquad\qquad\square$

In the next section we provide conditions for ϵ_{ijd} to go to zero for large distance d in the limit of infinite networks. In the rest of this section we show that the tightness of the bounds (and therefore ϵ_{ijd}), and their computational complexity (for a fixed d) depend only on the local structure around the edge l, and do not scale with the size of the network, under a suitable assumption.

Assumption 2. *Let $\mathcal{G}_R$ be a resistor network, $l = (i,j)$ an arbitrary edge of the network, and $\mathcal{N}_{\leq d}$ denote the set of nodes that are at distance no greater than d from at least one of i and j. We assume that the network is sparse in such a way that the cardinality of $\mathcal{N}_{\leq d}$ does not depend on N for any d.*

Assumption 2 is suitable for transportation networks, because of physical constraints not allowing for the degree of the nodes to grow unlimitedly (think for instance of a square grid, where the degree of the nodes is 4 no matter what the size of the network is). Notice also that, under Assumption 2, N and E are proportional. Hence, from now on we refer indistinctly to N or E to denote the size of the network.

Proposition 5. *Let $\mathcal{G}_R = (\mathcal{N}, \mathcal{E}_R, W)$ be a resistor network, (i,j) be a pair of neighbors, $d \geq 1$. Then, the time complexity of the bounds and the tightness of the bounds are functions of the structure of $\mathcal{G}_R$ within distance $d+1$ from i and j only. Furthermore, under Assumption 2 they do not depend on the size of the network.*

Proof. See the Appendix. $\qquad\qquad\square$

Remark 2. *Proposition 5 states that under Assumption 2 the time complexity to approximate a single effective resistance does not scale with the size of the network for every distance d. Therefore, all the effective resistance may be approximated in linear time. V is computed via a diagonally dominant, symmetric and positive definite linear systems. The design of fast algorithms to solve this class of problem is an active field of research in the last years. To the best of our knowledge, the best algorithm has been provided by Cohen et al. [9] and has complexity $O(M \log^k N \log 1/\epsilon)$, where ϵ is the tolerance error, k is a constant, and M is the number of nonzero elements in the matrix of the linear system. Since in our case M scales with E, and since E scales with N under Assumption 2, Algorithm 1 is quasilinear in N.*

5 Bound analysis

In this section we provide a characterization of the tightness of the bounds of the effective resistance between neighbors in terms of random walks over the resistor networks $\mathcal{G}_R$, $\mathcal{G}_{ij}^{U_d}$ and $\mathcal{G}_{ij}^{L_d}$. We then use this characterization to provide a sufficient condition on the network under which the approximation error of the bounds vanishes asymptotically as the distance d grows. To this end, we introduce the following notation. Let

- T_S and T_S^+ denote the hitting time, (i.e., the first time $t \geq 0$ such that the random walk hits the set S), and return time (i.e., the first time $t > 0$ such that the random walk hits the set S), respectively.

- N_d denote the set of the nodes that are at distance d from either i or j and at distance greater or equal than d from the other node (we omit i and j for simplicity of notation).

- $p_k(X)$, $p_k^{U_d}(X)$ and $p_k^{L_d}(X)$, denote the probability that the event X occurs, given a random walk that starts in k at time 0 and evolves over the resistor networks $\mathcal{G}_R$, $\mathcal{G}_{ij}^{U_d}$ and $\mathcal{G}_{ij}^{L_d}$, respectively.

The next proposition provides a characterization for the distance between the upper and lower bound on r_{ij} in terms of probabilities of random walks over $\mathcal{G}_R$, $\mathcal{G}_{ij}^{U_d}$ and $\mathcal{G}_{ij}^{L_d}$.

Proposition 6. *Let $\mathcal{G}_R = (\mathcal{N}, \mathcal{E}_R, W)$ be a resistor network. Based on the random walk on the resistor network,*

$$r_{ij}^{U_d} - r_{ij}^{L_d} \leq \frac{D_{ii}}{(W_{ij})^2} \underbrace{p_i(T_{N_d} < T_j)}_{Term\ 1} \cdot \max_{g \in N_d} \underbrace{\left(p_g^{U_d}(T_i < T_j) - p_g^{L_d}(T_i < T_j)\right)}_{Term\ 2}. \quad (16)$$

Proof. See the Appendix. $\qquad\square$

In the next subsection we use this result to study the asymptotic behaviour of the error term $\epsilon_{ijd} = (r_{ij}^{U_d} - r_{ij}^{L_d})/a_l$. In §5.1 we show that this error goes

Table 1: All the four cases are possible, as shown in §5.2. Term $1 \to 0$ under the assumption that the network is recurrent, as proved in §5.1.

	Term $2 \to 0$	Term $2 \nrightarrow 0$
Term $1 \to 0$	2d grid	Ring
Term $1 \nrightarrow 0$	3d grid	Double tree

to zero for the class of recurrent networks. The core idea to prove this result is to show that Term 1 goes to zero. To generalize our analysis beyond recurrent networks, in §5.2 we study both Term 1 and 2 and provide examples showing that all combinations are possible (see Table 1). In particular, it is still possible that $\epsilon_{ijd} \to 0$ for non-recurrent networks (for which Term $1 \nrightarrow 0$, see Levin and Peres [23]) if Term $2 \to 0$.

5.1 Recurrent networks

In this section we show that a sufficient condition under which the distance between the upper and the lower bound vanishes as the distance d goes to infinity is that the network is recurrent. We start by introducing this class of networks.

Definition 5.1 (Recurrent random walk). *A random walk is recurrent if, for every starting point, it visits its starting node infinitely often with probability one (Levin and Peres [23]).*

Definition 5.2 (Recurrent network). *An infinite network $\mathcal{G} = (\mathcal{N}, \mathcal{E}, W)$ is recurrent if the random walk on the network is recurrent.*

The next theorem states that the distance between the upper and the lower bound on recurrent networks vanishes as d goes to infinity, provided that the degree of every node is finite.

Theorem 3. *Let $\mathcal{G}_R$ be an infinite recurrent resistor network, and let the maximal weighted degree D_{max} be finite. Then, for every edge $l = (i, j)$,*

$$\lim_{d \to +\infty} (r_{ij}^{U_d} - r_{ij}^{L_d}) = 0.$$

Proof. It is proved in Levin and Peres [23, Proposition 21.3] that a graph is recurrent if and only if

$$\lim_{d+ \to \infty} p_i(T_{N_d} < T_j) = 0 \quad \forall i, j \in \mathcal{N}. \tag{17}$$

Observe that, to hit any node at distance $d+1$, the random walk starting from i has to hit at least a node at distance d. Hence, the sequence $\left\{p_i(T_{N_d} < T_j)\right\}_{d=1}^{\infty}$ is non-increasing in d and the limit is well defined. Then, from (16), (17), from

the fact that $0 \leq p_g^{U_d}(T_i < T_j) - p_g^{L_d}(T_i < T_j) \leq 1$ for every node g, and from the assumptions $D_{max} < +\infty$ and $W_{ij} > 0$ (recall that i and j are adjacent nodes), it follows

$$\lim_{d \to +\infty} r_{ij}^{U_d} - r_{ij}^{L_d} \leq \frac{D_{max}}{(W_{ij})^2} \lim_{d \to +\infty} p_i(T_{N_d} < T_j) = 0,$$

which completes the proof. $\square$

Remark 3. *Theorem 3 implies that $\lim_{d \to +\infty} \epsilon_{ijd} = 0$ on recurrent networks for every neighboring nodes i and j. Hence, by Theorem 2, the cost variation corresponding to intervention on edge $l = (i,j)$ can be estimated with vanishing error. Observe that not only the error term ϵ_{ijd}, but also the relative error ϵ_{ijd}/r_{ij}, vanishes asymptotically, since $r_{ij} \geq 1/D_{max}$.*

Recurrence is a sufficient condition to guarantee $\lim_{d \to +\infty} \epsilon_{ijd} = 0$, but is not necessary, as discussed in the next subsection.

5.2 Beyond recurrence

We here provide examples of infinite networks for all of the cases in Table 1. Observe that, for every edge $l = (i,j) \in \mathcal{E}_R$, the network cut at distance d from l and the network shorted at distance d from l differ for a node only. Let s denote such node, which is the result of shorting all the nodes at distance greater than d from both i and j in a unique node. Intuitively speaking, our conjecture is that Term 2 in (16) is small when the network has many short paths. In fact, in this case, adding the node s does not affect too much the probability, starting from any node in N_d, of hitting i before j, thus making Term 2 small. This intuition can be made more clear by the next examples.

5.2.1 2d grid

Consider an infinite unweighted bidimensional grid as in Fig. 3. This network is relevant for the NDP since many transportation networks are very similar to grids. This network is recurrent (Levin and Peres [23]), hence Theorem 3 guarantees that Term 1 and thus ϵ_{ijd} go to 0 for large d. Our conjecture, confirmed by numerical simulations, is that, for every node $g \in N_d$,

$$\lim_{d \to +\infty} p_g^{U_d}(T_i < T_j) = 1/2, \quad \lim_{d \to +\infty} p_g^{L_d}(T_i < T_j) = 1/2.$$

Hence, this is recurrent network for which also Term 2 vanishes asymptotically.

5.2.2 3d grid

Consider an infinite unweighted tridimensional grid. This network is not recurrent (Levin and Peres [23]), therefore Term 1 does not go to 0 and we cannot
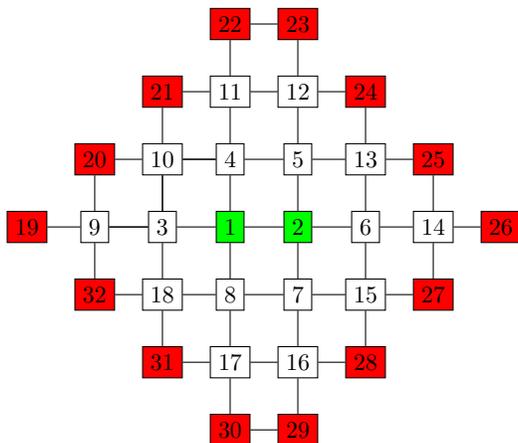
Figure 3: Bidimensional square grid, cut at distance $d = 3$. The red nodes belong to N_d. As d grows, $p_g(T_i < T_j)$ approaches $1/2$ for each $g \in N_d$, because there are many short paths.

conclude that $\epsilon_{ijd} \to 0$ from Theorem 3. Nonetheless, numerical simulations show that, similarly to the bidimensional grid, for every node $g \in N_d$,

$$\lim_{d \to +\infty} p_g^{U_d}(T_i < T_j) = 1/2, \quad \lim_{d \to +\infty} p_g^{L_d}(T_i < T_j) = 1/2.$$

Hence, this is a non-recurrent network for which Term 2 (and therefore ϵ_{ijd}) vanishes as the distance grows.

5.2.3 Ring

Consider an infinite unweighted ring as in Fig. 4. Consider nodes c and e as in Fig. 4. Then,

$$p_c^{U_d}(T_i < T_j) = 1, \quad p_e^{U_d}(T_i < T_j) = 0.$$

for each d (even $d \to +\infty$), whereas,

$$p_c^{L_d}(T_i < T_j) = \frac{d}{2d+1} \xrightarrow[d \to +\infty]{} \frac{1}{2}, \quad p_e^{L_d}(T_i < T_j) = \frac{d+1}{2d+1} \xrightarrow[d \to +\infty]{} \frac{1}{2},$$

since this case is equivalent to the gambler's ruin problem (Levin and Peres [23]). Hence, Term 2 does not vanish for the ring. This is due to the fact that, on the ring, all the paths from c to j not passing in i include the node s. Still, Term 1, and therefore ϵ_{ijd}, vanish asymptotically by Theorem 3, since this network is recurrent.

5.2.4 Double tree network

We finally propose an infinite network for which ϵ_{ijd} does not converge asymptotically. This network is not relevant for traffic applications, since it admits one
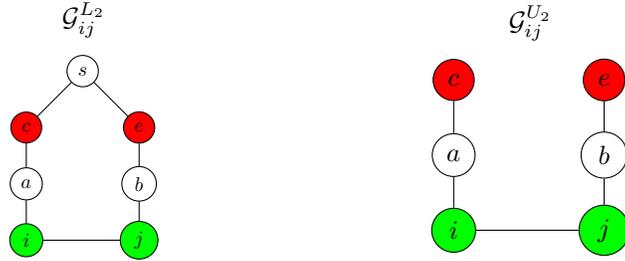
19

Figure 4: Left: shorted ring at distance $d = 2$. Right: cut ring at distance $d = 2$.
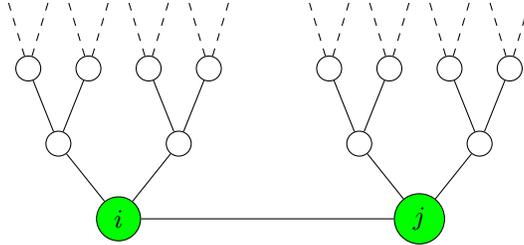

Figure 5: The double tree is an infinite non-recurrent network. On this network $\lim_{d \to \infty} \epsilon_{ijd} = 1/3$.

path only between every pair of nodes, but provides an interesting counterexample where the bounds do not converge asymptotically. The network is composed of two infinite regular trees, starting from node i and j respectively, linked by an edge $l = (i, j)$, as in Fig. 5, and it is assumed unweighted. It can be shown that on the double tree network the probability that the random walk, starting from i, returns on i is equal to the same quantity for a biased random walk over an infinite line (for more details we refer to the Supplementary Materials). Since the biased random walk on a line is not recurrent (see Levin and Peres [23]), this equivalence shows that the double tree network is non-recurrent, and Term 1 $\to 0$. Moreover, we show in the Supplementary Materials that

$$\lim_{d \to +\infty} r_{ij}^{U_d} - r_{ij}^{L_d} = \frac{1}{3},$$

thus implying that Term 2 $\nrightarrow 0$.

6 Simulations

6.1 Infinite grids

Infinite regular grids are useful to test the performance of the bounds. Indeed, despite having an infinite number of nodes, the effective resistance between
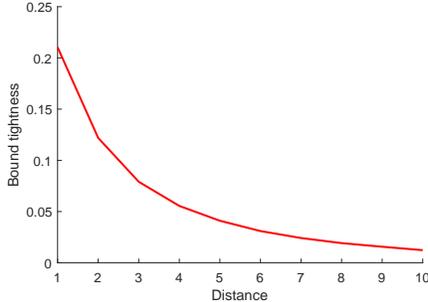
Figure 6: Average relative error of the bounds on Oldenburg network as a function of distance d.

adjacent nodes can be computed exploiting the symmetric structure of the grid. We focus on the square grid, but similar arguments can be applied to any regular infinite grid.

Lemma 2 (Bartis [2]). *Let $\mathcal{G}_R$ be an infinite square grid with unitary resistances. Then, the effective resistance between two neighboring nodes is $1/2$.*

In Table 2 the performances of the upper and lower bounds are shown. Numerical simulations show that for every edge $l = (i, j)$,

$$\frac{r_{ij}^{U_d} - r_{ij}}{r_{ij}} = \frac{r_{ij} - r_{ij}^{L_d}}{r_{ij}} = O(1/d^2).$$

We underline that the relative errors of the bounds are symmetric only in the square grid, but they scale similarly in all the regular bidimensional grids. Ob-

Table 2: Table of upper and lower bound in infinite square grid.

	$d=1$	$d=2$	$d=3$	$d=4$	$d=5$
$(r_{ij}^{U_d} - r_{ij})/r_{ij}$	$1/5$	0.0804	0.0426	0.0262	0.0178
$(r_{ij} - r_{ij}^{L_d})/r_{ij}$	$1/5$	0.0804	0.0426	0.0262	0.0178

serve that, despite the network being infinite, even at $d = 5$, the upper and the lower bounds give a good estimation of the true effective resistance.

6.2 Simulations on a real transportation network

In this section we present the performances of the cutting and shorting techniques on the traffic network of the city of Oldenburg (Brinkhoff [6]). The network is composed of 6105 nodes and 7035 edges, and its diameter is 104.

The network is assumed to be unweighted, with $a_l = 1$ for every edge l. The average relative error of the bounds, i.e.,

$$AT_d := \frac{1}{E} \sum_{(i,j) \in \mathcal{E}} \frac{r_{ij}^{U_d} - r_{ij}^{L_d}}{r_{ij}}$$

is shown in Table 3 and Fig. 6. Even for this network, the error of the bounds

Table 3: Table of the average relative error of the bounds at distance d.

	d=1	d=2	d=3	d=4	d=5	d=6	d=7	d=8	d=9	d=10
AT_d	0.21	0.12	0.079	0.056	0.041	0.031	0.024	0.019	0.016	0.012

decreases quickly, compared to the diameter of the network.

7 Conclusion

In this work we study a discrete network design problem, where a single edge can be improved. We reformulate the problem in terms of electrical quantities, in particular in terms of the effective resistance between the two nodes at the end of the edge. We then provide a method to approximate such effective resistance by performing only local computations. Both the tightness and the computational complexity of our bounds do not depend on the size of the network, but on the local structure only. Based on the electrical formulation and our approximation method for the effective resistance we propose an efficient algorithm to solve the original design problem.

An interesting direction for the future is a deeper analysis on tightness of the bounds on effective resistance as a function of the network and the distance d, since so far we have a result for the asymptotic behaviour only. Future research lines also include extending to the case of multiple interventions, dealing with different interventions (e.g building new roads), and the relaxation of some assumptions like the single origin and destination, and the assumption that the set of used edges is not affected by the intervention.

Acknowledgments.

References

[1] David Aldous and Jim Fill. Reversible markov chains and random walks on graphs. book in preparation. *URL for draft at http://www. stat. Berkeley. edu/users/aldous*, 2000.

[2] Francis J Bartis. Let's analyze the resistance lattice. *American Journal of Physics*, 35(4):354–355, 1967.

[3] Martin Beckmann, Charles B McGuire, and Christopher B Winsten. Studies in the economics of transportation. Technical report, 1956.

[4] Dennis S Bernstein. *Matrix mathematics: theory, facts, and formulas.* Princeton university press, 2009.

[5] Stephen Boyd and Lieven Vandenberghe. *Convex optimization.* Cambridge university press, 2004.

[6] Thomas Brinkhoff. A framework for generating network-based moving objects. *GeoInformatica*, 6(2):153–180, 2002.

[7] Philip N Brown and Jason R Marden. Studies on robust social influence mechanisms: Incentives for efficient network routing in uncertain settings. *IEEE Control Systems Magazine*, 37(1):98–115, 2017.

[8] Suh-Wen Chiou. Bilevel programming for the continuous transport network design problem. *Transportation Research Part B: Methodological*, 39(4):361–383, 2005.

[9] Michael B Cohen, Rasmus Kyng, Gary L Miller, Jakub W Pachocki, Richard Peng, Anup B Rao, and Shen Chen Xu. Solving sdd linear systems in nearly m log1/2 n time. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 343–352, 2014.

[10] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. How much can taxes help selfish routing? *Journal of Computer and System Sciences*, 72(3):444–467, 2006.

[11] Stella Dafermos and Anna Nagurney. On some traffic equilibrium theory paradoxes. *Transportation Research Part B: Methodological*, 18(2):101–110, 1984.

[12] Sanmay Das, Emir Kamenica, and Renee Mirka. Reducing congestion through information design. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1279–1284. IEEE, 2017.

[13] Peter G Doyle and J Laurie Snell. *Random walks and electric networks*, volume 22. American Mathematical Soc., 1984.

[14] Wendy Ellens, FM Spieksma, P Van Mieghem, A Jamakovic, and RE Kooij. Effective graph resistance. *Linear algebra and its applications*, 435(10):2491–2506, 2011.

[15] Reza Zanjirani Farahani, Elnaz Miandoabchi, Wai Yuen Szeto, and Hannaneh Rashidi. A review of urban transportation network design problems. *European Journal of Operational Research*, 229(2):281–302, 2013.

[16] Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 277–285. IEEE, 2004.

[17] Pirmin Fontaine and Stefan Minner. A dynamic discrete network design problem for maintenance planning in traffic networks. *Annals of Operations Research*, 253(2):757–772, 2017.

[18] Dimitris Fotakis, Alexis C Kaporis, and Paul G Spirakis. Efficient methods for selfish network design. *Theoretical Computer Science*, 448:9–20, 2012.

[19] Ziyou Gao, Jianjun Wu, and Huijun Sun. Solution algorithm for the bi-level discrete network design problem. *Transportation Research Part B: Methodological*, 39(6):479–495, 2005.

[20] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.

[21] Max Klimm and Philipp Warode. Computing all wardrop equilibria parametrized by the flow demand. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 917–934. SIAM, 2019.

[22] Larry J LeBlanc. An algorithm for the discrete network design problem. *Transportation Science*, 9(3):183–199, 1975.

[23] David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

[24] Changmin Li, Hai Yang, Daoli Zhu, and Qiang Meng. A global optimization method for continuous network design problems. *Transportation Research Part B: Methodological*, 46(9):1144–1158, 2012.

[25] Emily Meigs, Francesca Parise, Asuman Ozdaglar, and Daron Acemoglu. Optimal dynamic information provision in traffic routing. *arXiv preprint arXiv:2001.03232*, 2020.

[26] Hossain Poorzahedy and Omid M Rouhani. Hybrid meta-heuristic algorithms for solving network design problem. *European Journal of Operational Research*, 182(2):578–596, 2007.

[27] Patrick Rebeschini and Sekhar Tatikonda. Locality in network optimization. *IEEE Transactions on Control of Network Systems*, 6(2):487–500, 2018.

[28] Tim Roughgarden. On the severity of braess's paradox: designing networks for selfish users is hard. *Journal of Computer and System Sciences*, 72(5):922–953, 2006.

[29] David Schrank, Bill Eisele, Tim Lomax, and Jim Bak. 2015 urban mobility scorecard. 2015.

[30] Jack Sherman and Winifred J Morrison. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals of Mathematical Statistics*, 21(1):124–127, 1950.

[31] Richard Steinberg and Willard I Zangwill. The prevalence of braess' paradox. *Transportation Science*, 17(3):301–318, 1983.

[32] Guangmin Wang, Ziyou Gao, Meng Xu, and Huijun Sun. Models and a relaxation algorithm for continuous network design problem with a tradable credit scheme and equity constraints. *Computers & Operations Research*, 41:252–261, 2014.

[33] Shuaian Wang, Qiang Meng, and Hai Yang. Global optimization methods for the discrete network design problem. *Transportation Research Part B: Methodological*, 50:42–60, 2013.

[34] John Glen Wardrop. Road paper. some theoretical aspects of road traffic research. *Proceedings of the institution of civil engineers*, 1(3):325–362, 1952.

[35] Manxi Wu and Saurabh Amin. Information design for regulating traffic flows under uncertain network state. In *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 671–678. IEEE, 2019.

[36] Manxi Wu, Saurabh Amin, and Asuman E Ozdaglar. Value of information in bayesian routing games. *Operations Research*, 69(1):148–163, 2021.

[37] Hai Yang and Michael G H. Bell. Models and algorithms for road network design: a review and some new developments. *Transport Reviews*, 18(3):257–278, 1998.

[38] Jia Shuo Yue, Chinmoy V Mandayam, Deepak Merugu, Hossein Karkeh Abadi, and Balaji Prabhakar. Reducing road congestion through incentives: a case study. In *Transportation Research Board 94th Annual Meeting, Washington, DC*, 2015.

[39] Yong Zhao and Kara Maria Kockelman. On-line marginal-cost pricing across networks: Incorporating heterogeneous users and stochastic equilibria. *Transportation Research Part B: Methodological*, 40(5):424–435, 2006.

Appendix

Preliminaries on connection between Green's function, random walks and effective resistance

Let $\mathcal{G}_R = (\mathcal{N}, \mathcal{E}_R, W)$ denote a connected resistor network, $D = \mathrm{diag}W\mathbb{1} \in \mathbb{R}^{N \times N}$ denote the degree matrix, and $P = D^{-1}W \in \mathbb{R}^{N \times N}$ the associated random walk on the network. We define as $_kP \in \mathbb{R}^{(N-1) \times (N-1)}$ the matrix where the row and the column referring to the node k are deleted. $_kP$ can be thought of as the transition matrix of a killed random walk obtained by creating a cemetery in the node k. We then define the Green's function $_kG \in \mathbb{R}^{(N-1) \times (N-1)}$ as

$$_kG := \sum_{t=0}^{\infty} (_kP)^t = (\mathbb{I} - {}_kP)^{-1}, \tag{18}$$

where the last inequality follows from the connectedness of $\mathcal{G}_R$ and from the fact that $_kP$ is substochastic and irreducible. Hence, it has spectral radius $\rho < 1$ and $(\mathbb{I} - {}_kP)^{-1} = \sum_{t=0}^{\infty} (_kP)^t$ (see Horn and Johnson [20]). Since $((_kP)^t)_{ij}$ is the probability that the killed random walk starting from i is in j after t steps, $_kG_{ij}$ indicates the expected number of times that the killed random walk visits j starting from i before hitting k (Ellens and Spieksma [14]). While these results hold for any network, it is known (see Ellens and Spiesksma [14]) that the Green's function of the random walk on a resistor network can be related to electrical quantities. In particular, with the convention that

$$_kG_{ik} = {}_kG_{ki} = {}_kG_{kk} = 0 \quad \forall i \in \mathcal{N}, \tag{19}$$

it is known that for any node k,

$$\frac{_kG_{ii} - {}_kG_{ji}}{D_{ii}} + \frac{_kG_{jj} - {}_kG_{ij}}{D_{jj}} = \frac{_jG_{ii}}{D_{ii}} = \frac{1}{D_{ii}p_i(T_j < T_i^+)} = r_{ij}, \tag{20}$$

where $p_i(T_j < T_i^+)$ is as defined in Section 5, and r_{ij} is the effective resistance as defined in §3.3.

Proof of Proposition 1

From (6), for all the used edges,

$$\gamma_i^* - \gamma_j^* = a_{(i,j)} f_{(i,j)}^* + b_{(i,j)} \quad \forall (i,j) \in \mathcal{E}.$$

So, by (3), any used path $p = (\mathrm{o}, n_1, n_2, \cdots, n_p, \mathrm{d})$, namely a path containing only used edges, has cost at the equilibrium

$$
\begin{aligned}
c_p(f^*) &= a_{(\mathrm{o},n_1)} f_{(\mathrm{o},n_1)}^* + b_{(\mathrm{o},n_1)} + \cdots + a_{(n_p,\mathrm{d})} f_{(n_p,\mathrm{d})}^* + b_{(n_p,\mathrm{d})} \\
&= (\gamma_\mathrm{o}^* - \gamma_{n_1}^*) + (\gamma_{n_1}^* - \gamma_{n_2}^*) + \cdots + (\gamma_{n_p}^* - \gamma_\mathrm{d}^*) = \gamma_\mathrm{o}^* - \gamma_\mathrm{d}^* = \gamma_\mathrm{o}^*,
\end{aligned}
\tag{21}
$$

where the last equivalence follows from the assumption $\gamma_{\mathrm{d}}^* = 0$. Hence, all the used paths at the equilibrium have the same cost γ_{o}^*. Then, the social cost is

$$
\begin{aligned}
C(f^*) &= \sum_{l \in \mathcal{E}} f_l^* d_l(f_l^*) = \sum_{l \in \mathcal{E}} d_l(f_l^*) \sum_{p \in \mathcal{P}} R_{lp} z_p^* \\
&= \sum_{p \in \mathcal{P}} z_p^* \sum_{l \in \mathcal{E}} R_{lp} d_l(f_l^*) = \sum_{p \in \mathcal{P}} z_p^* c_p(f^*) = \gamma_{\mathrm{o}}^* \sum_{p \in \mathcal{P}} z_p^* = \gamma_{\mathrm{o}}^* \tau,
\end{aligned}
$$

where the second equivalence follows from (2), the forth from (3), the fifth from (21) and the last one from the mass constraint (1). The proof follows by applying similar arguments for $C(f^*(l))$.

Proof of Proposition 2

Let $x^*(l) = H(l)^{-1} y$ be the equilibrium after the intervention on edge l. Notice that

$$
H(l) = H - a_l \Big(1 - \frac{1}{\kappa}\Big) \delta_l \delta_l^T,
$$

Since only the element A_{ll} is perturbed, $H(l)$ is a rank-1 perturbation of H. Hence, its inverse can be computed by Sherman-Morrison formula (Sherman and Morrison [30]):

$$
H(l)^{-1} = H^{-1} - \frac{-a_l \big(1 - \frac{1}{\kappa}\big) H^{-1} \delta_l \delta_l^T H^{-1}}{1 - a_l \big(1 - \frac{1}{\kappa}\big) \delta_l^T H^{-1} \delta_l}.
$$

By right multiplying by y, and recalling that $x^* = H^{-1} y$, it follows

$$
x^* - x^*(l) = -\frac{a_l \big(1 - \frac{1}{\kappa}\big) H^{-1} \delta_l \delta_l^T x^*}{1 - a_l \big(1 - \frac{1}{\kappa}\big) \delta_l^T H^{-1} \delta_l}.
$$

Observe that $\gamma_{\mathrm{o}}^*(l)$ is $x_{E+1}^*(l)$ by construction. Then, by selecting the $(E-1)$-th element, we obtain

$$
\gamma_{\mathrm{o}}^* - \gamma_{\mathrm{o}}^*(l) = -\frac{a_l \big(1 - \frac{1}{\kappa}\big) H_{(E+1,l)}^{-1} x_l^*}{1 - a_l \big(1 - \frac{1}{\kappa}\big) H_{ll}^{-1}} = \frac{(\kappa - 1) a_l H_{(E+1,l)}^{-1} x_l^*}{(\kappa - 1) a_l H_{ll}^{-1} - \kappa}.
$$

The statement follows from Proposition 1 and from noticing that $x_l^* = f_l^*$ by construction.

Proof of Theorem 1

By letting $\tilde{W} \in \mathbb{R}^{(N-1) \times (N-1)}$ and $\tilde{D} \in \mathbb{R}^{(N-1) \times (N-1)}$ denote the restriction over $\mathcal{N} \setminus \mathrm{d}$, i.e., all the nodes except the destination, of W and D respectively, it is straightforward to check that

$$
Q = \tilde{D} - \tilde{W}.
$$

We notice that, $_\mathrm{d}P$, as defined in §7, is

$$_\mathrm{d}P = \tilde{D}^{-1}\tilde{W}.$$

$_\mathrm{d}P$ is substochastic, since the rows referring to nodes pointing to the destination sum to less than one. We now prove that Q is invertible. Indeed,

$$Q^{-1} = (\tilde{D}-\tilde{W})^{-1} = (\tilde{D}(\mathbb{I}-_\mathrm{d}P))^{-1} = (\mathbb{I}-_\mathrm{d}P)^{-1}\tilde{D}^{-1} = \sum_{t=0}^{\infty}(_\mathrm{d}P)^t\tilde{D}^{-1} = _\mathrm{d}G\tilde{D}^{-1}, \tag{22}$$

where the penultimate equivalence follows from strongly connectedness of $\mathcal{G}$ (and therefore, connectedness of $\mathcal{G}_R$) and (18).
Let $l = (i,j) \in \mathcal{E}$ an arbitrary edge. From (10) it follows:

$$H_{ll}^{-1} = \frac{1}{a_l} - K_{l,:}Q^{-1}K_{:,l}^T.$$

By (11),

$$H_{ll}^{-1} = \begin{cases} \frac{1}{a_l} - \frac{1}{a_l^2}(Q_{ii}^{-1} + Q_{jj}^{-1} - Q_{ij}^{-1} - Q_{ji}^{-1}) & \text{if } j \neq \mathrm{d} \\ \frac{1}{a_l} - \frac{1}{a_l^2}(Q_{ii}^{-1}) & \text{if } j = \mathrm{d}. \end{cases} \tag{23}$$

We now construct $\hat{Q}^{-1} \in \mathbb{R}^{N\times N}$ and $_\mathrm{d}\hat{G} \in \mathbb{R}^{N\times N}$ by adding a zero column and a zero row to Q^{-1} and $_\mathrm{d}G$ and $\hat{K} \in \mathbb{R}^{E\times N}$ by adding a zero column to K corresponding to the destination. Thus, (23) can be written as

$$H_{ll}^{-1} = \frac{1}{a_l} - \frac{1}{a_l^2}(\hat{Q}_{ii}^{-1} + \hat{Q}_{jj}^{-1} - \hat{Q}_{ij}^{-1} - \hat{Q}_{ji}^{-1}) \quad \forall l \in \mathcal{E},$$

and, by $\hat{Q}^{-1} = _\mathrm{d}\hat{G}D^{-1}$ (which follows from (22)),

$$H_{ll}^{-1} = \frac{1}{a_l} - \frac{1}{a_l^2}\left(\frac{_\mathrm{d}\hat{G}_{ii} - _\mathrm{d}\hat{G}_{ji}}{D_{ii}} + \frac{_\mathrm{d}\hat{G}_{jj} - _\mathrm{d}\hat{G}_{ij}}{D_{jj}}\right).$$

Finally, by noticing that the definition of $_\mathrm{d}\hat{G}$ is coherent with (19), and by (20), we get

$$H_{ll}^{-1} = \frac{1}{a_l} - \frac{r_{ij}}{a_l^2}. \tag{24}$$

Using the same notation with $\hat{Q}^{-1}$ and $\hat{K}^{-1}$ to handle also edges l pointing to the destination, from (10), (11), (22) and from symmetry of $\hat{Q}^{-1}$, it follows:

$$H_{\mathrm{E}+1,l}^{-1} = -(\hat{Q}^{-1}\hat{K}^T)_{1l} = -\frac{\hat{Q}_{1i}^{-1} - \hat{Q}_{1j}^{-1}}{a_l} = -\frac{\hat{Q}_{i1}^{-1} - \hat{Q}_{j1}^{-1}}{a_l} = -\frac{_\mathrm{d}\hat{G}_{io} - _\mathrm{d}\hat{G}_{jo}}{a_l D_{oo}}, \tag{25}$$

since the first node is the origin by construction.
We now prove that

$$V_i = \frac{_\mathrm{d}\hat{G}_{io}}{_\mathrm{d}\hat{G}_{oo}} \quad \forall i \in \mathcal{N}. \tag{26}$$

To this end, notice that both the potential and the Green's function are harmonic functions (see Levin and Peres [23]) satisfying same boundary conditions on the origin and the destination, i.e.,

$$V_{\mathrm{o}} = 1 = \frac{\mathrm{d}\hat{G}_{\mathrm{oo}}}{\mathrm{d}\hat{G}_{\mathrm{oo}}}, \quad V_{\mathrm{d}} = 0 = \frac{\mathrm{d}\hat{G}_{\mathrm{do}}}{\mathrm{d}\hat{G}_{\mathrm{oo}}},$$

since $_{\mathrm{d}}\hat{G}_{\mathrm{do}} = 0$ by construction, which implies (26) by Levin and Peres [23, Proposition 9.1]. Plugging (26) into (25), we get

$$H_{\mathrm{E}+1,l} = -\frac{\mathrm{d}\hat{G}_{\mathrm{oo}}}{a_l D_{\mathrm{oo}}}(V_i - V_j). \tag{27}$$

The statement follows from plugging (27) and (24) in (9) with the assignment $\tilde{\tau} = \tau \cdot {}_{\mathrm{d}}\hat{G}_{\mathrm{oo}}/D_{\mathrm{oo}} = \tau \cdot {}_{\mathrm{d}}G_{\mathrm{oo}}/D_{\mathrm{oo}}$.

Proof of Proposition 3

A sufficient condition under which $\mathcal{E}_+ = \emptyset$ is that the first E components of $x^* = H^{-1}b$, corresponding to equilibrium edge flows, are nonnegative. Indeed, since (4) is strictly convex, if the flows corresponding to $x^* = H^{-1}b$ satisfy the constraint $f^* \geq \mathbf{0}$, then f^* is feasible and is the unique Wardrop equilibrium, with $\lambda^* = 0$ because of the complementary slackness. Hence, we look for conditions satisfying $x_l^* \geq 0$ for every $l \in \{1, \cdots, \mathrm{E}\}$. Consider an arbitrary edge $l = (i, j)$. From (5) and (10), it follows:

$$x_l^* = -\frac{b_l}{a_l} + [KQ^{-1}K^T]_{l:}b + [KQ^{-1}]_{l:}(\nu_-).$$

With same arguments as in Proof of Theorem 1, we replace Q^{-1} with $\hat{Q}^{-1}$ and K with $\hat{K}$ to take into account edges pointing to the destination. Since $\nu_- = \tau\delta_{\mathrm{o}}$,

$$x_l^* = -\frac{b_l}{a_l} + [\hat{K}\hat{Q}^{-1}\hat{K}^T]_{l:}b + \tau[\hat{K}\hat{Q}^{-1}]_{l\mathrm{o}}$$

$$= -\frac{b_l}{a_l} + [\hat{K}\hat{Q}^{-1}\hat{K}^T]_{l:}b + \tau\frac{\mathrm{d}\hat{G}_{i\mathrm{o}} - {}_{\mathrm{d}}\hat{G}_{j\mathrm{o}}}{a_l D_{\mathrm{oo}}}$$

where the last equivalence follows from (25) and from the fact that the origin is the first node by construction. If $_{\mathrm{d}}\hat{G}_{i\mathrm{o}} - {}_{\mathrm{d}}\hat{G}_{j\mathrm{o}} > 0$, then, for any $\tau \geq \overline{\tau}_l$ with

$$\overline{\tau}_l = \frac{\frac{b_l}{a_l} - [K\hat{Q}^{-1}K^T]_{l:}b}{\frac{\mathrm{d}\hat{G}_{i\mathrm{o}} - {}_{\mathrm{d}}\hat{G}_{j\mathrm{o}}}{a_l D_{\mathrm{oo}}}}$$

it holds $x_l^* \geq 0$, which in turn implies that if $\tau \geq \overline{\tau} := \{\overline{\tau}_l\}_{l=1}^{\mathrm{L}}$, then $\mathcal{E}_+ = \emptyset$. Moreover, if the delays are linear, $_{\mathrm{d}}\hat{G}_{i\mathrm{o}} - {}_{\mathrm{d}}\hat{G}_{j\mathrm{o}} > 0$ implies $x_l^* \geq 0$ and $\mathcal{E}_+ = \emptyset$

for any value of τ, because $b = 0$. By (26) and Ohm's law (see [1]), $_\mathrm{d}\hat{G}_{io} - {}_\mathrm{d}\hat{G}_{jo}$ is proportional to $V_i - V_j = \mathrm{i}_l a_l$, where i_l denotes the current flowing through edge l. Then, it suffices to show that $\mathrm{i}_l > 0$. To this end, we observe that, by definition, if the network is series-parallel, it is a single edge (o, d) or it can obtained by connecting in series or in parallel two series-parallel networks. Thus, a series-parallel network can be reduced to a single edge by recursively i) merging two edges l_1 and l_2 connected in series into a single edge l_3, with $a_3 = a_1 + a_2$ (recall that the coefficients a_e correspond to resistances on the resistor network), and ii) merging two edges l_1 and l_2 connected in parallel into a single edge l_3, with $a_3 = a_1 a_2/(a_1 + a_2)$. Moreover, observe that in both cases $i_3 > 0$ if and only if $i_1 > 0$ and $i_2 > 0$. Indeed, in case i) $i_3 = i_1 = i_2$, and in case ii) $i_1 = i_2 a_2/a_1$ and $i_3 = i_1 + i_2$. Obviously, when the network is reduced to a single edge, the flow on the unique edge is positive because $\tau > 0$. Then, by applying those arguments recursively, for every edge $l = (i, j) \in \mathcal{E}$,

$$ i_l > 0 \implies V_i - V_j > 0, $$

implying that if $\tau \geq \bar{\tau}$ then $x_l^* \geq 0$ and $\mathcal{E}_+ = \emptyset$.

Proof of Theorem 2

Using the definitions,

$$
\begin{aligned}
|\Delta C(l) - \Delta C^d(l)| &= \left| \frac{\tilde{\tau} f_l^*(V_i - V_j)}{\frac{1}{\kappa-1} + \frac{r_{ij}}{a_l}} - \frac{\tilde{\tau} f_l^*(V_i - V_j)}{\frac{1}{\kappa-1} + \frac{r_{ij}^{U_d} + r_{ij}^{L_d}}{2a_l}} \right| \\
&= \left| \frac{\tilde{\tau} f_l^*(V_i - V_j)}{\frac{1}{\kappa-1} + \frac{r_{ij}}{a_l}} \right| \cdot \left| \frac{\frac{r_{ij}^{U_d} + r_{ij}^{L_d} - 2r_{ij}}{2a_l}}{\frac{1}{\kappa-1} + \frac{r_{ij}^{U_d} + r_{ij}^{L_d}}{2a_l}} \right|,
\end{aligned}
$$

Note also that

$$
\frac{|r_{ij}^{U_d} + r_{ij}^{L_d} - 2r_{ij}|}{a_l} \leq \frac{|r_{ij}^{U_d} - r_{ij}| + |r_{ij} - r_{ij}^{L_d}|}{a_l} = \frac{r_{ij}^{U_d} - r_{ij} + r_{ij} - r_{ij}^{L_d}}{a_l} = \frac{r_{ij}^{U_d} - r_{ij}^{L_d}}{a_l} = \epsilon_{ijd}.
$$

Putting those two together, and using (13), we get

$$
\left| \frac{\Delta C(l) - \Delta C^d(l)}{\Delta C(l)} \right| \leq \frac{\epsilon_{ijd}}{2\left(\frac{1}{\kappa-1} + \frac{r_{ij}^{U_d} + r_{ij}^{L_d}}{2a_l} \right)} \leq \frac{\epsilon_{ijd}}{2\left(\frac{1}{\kappa-1} + \frac{1}{D_{max} \cdot a_l} \right)},
$$

where the last inequality follows from (14). Finally, (15) follows from $r_{ij}^{U_d} \geq r_{ij}$.

Proof of Proposition 5

The cut and shorted networks are obtained by finding the neighbors within distance d and $d + 1$ from (i, j), respectively. The neighbors of a node i can

be found by checking the non-zero elements of $W(i,:)$. The neighbors within distance d can be found by iterating such operation d times. Hence, the time for building the cut and the shorted network depends on the local structure, which, under Assumption 2, does not depend on the size of the network. Since the bounds of the effective resistance are computed on these subnetwork, their time complexity and tightness depends on local structure, which, under Assumption 2, is independent of the size of the network.

Proof of Proposition 6

We introduce the following notation:

- The index U_d and L_d indicate that the random walk takes place over $\mathcal{G}_{ij}^{U_d}$ and $\mathcal{G}_{ij}^{L_d}$, respectively. So, for instance, $_k G_{ij}^{U_d}$ denotes the expected number of times that the random walk on the network $\mathcal{G}_{ij}^{U_d}$, starting from i, hits j before hitting k.

- $p_i(T_u < T_S)$, with $u \in S$, denotes the probability that the random walk starting from i hits the node $u \in S$ before hitting any other node in S.

By applying (20) to the effective resistance of edge $l = (i, j)$ in the shorted and the cut network, it follows

$$
r_{ij}^{U_d} = \frac{_j G_{ii}^{U_d}}{D_{ii}}, \qquad r_{ij}^{L_d} = \frac{_j G_{ii}^{L_d}}{D_{ii}},
$$

where we recall that $_j G_{ii}^{U_d}$ and $_j G_{ii}^{L_d}$ are the expected number of visits on i, before hitting j, starting from i, of the random walk defined on $\mathcal{G}_{ij}^{U_d}$ and $\mathcal{G}_{ij}^{L_d}$ respectively. The visits on i before hitting j can be divided in two disjoint sets: the visits before hitting j and before visiting any node in N_d, and the visits before hitting j but after at least a node in N_d has been visited. Let $G_{ii}^{<N_d}$ denote the expected number of visits to i, starting from i, before hitting any node in N_d and before hitting the absorbing node j (for simplicity of notation we omit the index j from now on). Observe that this term is the same for both $\mathcal{G}_{ij}^{U_d}$ and $\mathcal{G}_{ij}^{L_d}$, since these two networks differ only for the node s, which cannot be visited without visiting before at least a node in N_d. Moreover, let $G_{ii}^{U>N_d}$ and $G_{ii}^{L>N_d}$ denote the expected number of visits to i, starting from i, before hitting j but after at least one node in N_d has been visited in the network $\mathcal{G}_{ij}^{U_d}$ and $\mathcal{G}_{ij}^{L_d}$, respectively. Thus,

$$
G_{ii}^{U_d} = G_{ii}^{<N_d} + G_{ii}^{U>N_d},
$$
$$
G_{ii}^{L_d} = G_{ii}^{<N_d} + G_{ii}^{L>N_d}.
$$

This implies

$$
r_{ij}^{U_d} - r_{ij}^{L_d} = \frac{G_{ii}^{U>N_d} - G_{ii}^{L>N_d}}{D_{ii}}. \tag{28}
$$

Notice that $G_{ii}^{U>N_d}$ can be written as the sum over the nodes $g \in N_d$ of the probability, starting from i, of hitting g and going back to i without hitting j, multiplied by the expected number of visits on i starting from i, before hitting j, which is the derivative of a geometric sum. Therefore,

$$G_{ii}^{U>N_d} = \sum_{g \in N_d} \underbrace{p_i(T_g < T_{j \cup N_d})}_{(1)} \underbrace{p_g^{U_d}(T_i < T_j)}_{(2)} \sum_{k=1}^{\infty} k \underbrace{\left(p_i^{U_d}(T_i^+ < T_j)\right)^{k-1}}_{(3)} \underbrace{\left(1 - p_i^{U_d}(T_i^+ < T_j)\right)}_{(4)}$$

$$= \frac{\sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) p_g^{U_d}(T_i < T_j)}{1 - p_i^{U_d}(T_i^+ < T_j)},$$

where:

1. probability from i of hitting g before hitting j and any other node in N_d;

2. probability from g of hitting i before j;

3. probability from i of hitting i $k-1$ times before hitting j;

4. probability from i of hitting j before i.

Similarly,

$$G_{ii}^{L>N_d} = \sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) p_g^{L_d}(T_i < T_j) \sum_{k=1}^{\infty} k \left(p_i^{L_d}(T_i^+ < T_j)\right)^{k-1} \left(1 - p_i^{L_d}(T_i^+ < T_j)\right)$$

$$= \frac{\sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) p_g^{L_d}(T_i < T_j)}{1 - p_i^{L_d}(T_i^+ < T_j)}.$$

Substituting in (28) yields

$$r_{ij}^{U_d} - r_{ij}^{L_d} = \frac{1}{D_{ii}} \sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) \left(\frac{p_g^{U_d}(T_i < T_j)}{1 - p_i^{U_d}(T_i^+ < T_j)} - \frac{p_g^{L_d}(T_i < T_j)}{1 - p_i^{L_d}(T_i^+ < T_j)} \right).$$

From (20), it follows

$$r_{ij}^{U_d} = \frac{1}{D_{ii} p_i^{U_d}(T_j < T_i^+)} = \frac{1}{D_{ii}\left(1 - p_i^{U_d}(T_i^+ < T_j)\right)},$$

$$r_{ij}^{L_d} = \frac{1}{D_{ii} p_i^{L_d}(T_j < T_i^+)} = \frac{1}{D_{ii}\left(1 - p_i^{L_d}(T_i^+ < T_j)\right)}.$$
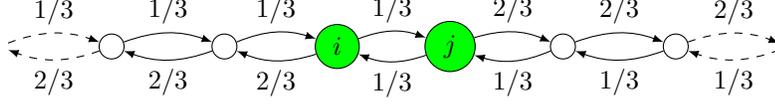
Figure 7: The double tree network is equivalent to a biased random walk like this.

Thus,

$$r_{ij}^{U_d} - r_{ij}^{L_d} = \sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) \big(p_g^{U_d}(T_i < T_j) r_{ij}^{U_d} - p_g^{L_d}(T_i < T_j) r_{ij}^{L_d} \big)$$

$$= \sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) \big(p_g^{U_d}(T_i < T_j) - p_g^{L_d}(T_i < T_j) \big) r_{ij}^{U_d} +$$

$$+ \sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) p_g^{L_d}(T_i < T_j) (r_{ij}^{U_d} - r_{ij}^{L_d})$$

$$\le \sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) \big(p_g^{U_d}(T_i < T_j) - p_g^{L_d}(T_i < T_j) \big) r_{ij}^{U_d} +$$

$$+ \sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) (r_{ij}^{U_d} - r_{ij}^{L_d})$$

$$= \sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) \big(p_g^{U_d}(T_i < T_j) - p_g^{L_d}(T_i < T_j) \big) r_{ij}^{U_d} +$$

$$+ p_i(T_{N_d} < T_j)(r_{ij}^{U_d} - r_{ij}^{L_d}),$$

where the last inequality follows from $p_g^L(T_i < T_j) \le 1$ and the last equality from the fact that $p_i(T_{N_d} < T_j) = \sum_{g \in N_d} p_i(T_g < T_{j \cup N_d})$. It follows

$$r_{ij}^{U_d} - r_{ij}^{L_d} \le \frac{\sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) \big(p_g^U(T_i < T_j) - p_g^L(T_i < T_j) \big) r_{ij}^{U_d}}{1 - p_i(T_{N_d} < T_j)}$$

$$\le \sum_{g \in N_d} p_i(T_g < T_{j \cup N_d}) \big(p_g^U(T_i < T_j) - p_g^L(T_i < T_j) \big) r_{ij}^{U_d} \frac{D_{ii}}{W_{ij}}$$

$$\le p_i(T_{N_d} < T_j) \cdot \max_{g \in N_d} \big(p_g^U(T_i < T_j) - p_g^L(T_i < T_j) \big) \frac{D_{ii}}{(W_{ij})^2}.$$

where the second inequality follows from $1 - p_i(T_{N_d} < T_j) = p_i(T_j < T_{N_d}) \ge P_{ij} = W_{ij}/D_{ii}$ and the last one from $r_{ij}^{U_d} \le 1/W_{ij}$ (as shown in (14)) and from $p_i(T_{N_d} < T_j) = \sum_{g \in N_d} p_i(T_g < T_{j \cup N_d})$.

More details on Section 5.2.4

We prove that the double tree network is not recurrent by showing that $p_i(T_i < T_{N_d})$ is the same as in a biased random walk. Indeed, from any d the probability of going from a node at distance d from i to a node at distance $d+1$ and $d-1$
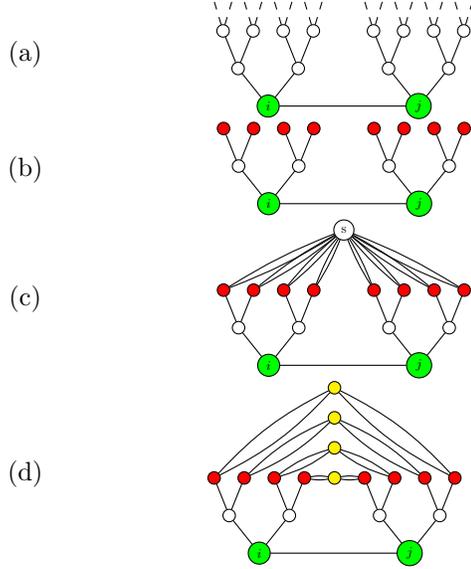
Figure 8: From above to below: (a) the double tree network; (b) the cut tree network at distance 2 from $l = (i,j)$; (c) the shorted tree network at distance 2 from $l = (i,j)$; (d) a network equivalent to the shorted one. In red, the nodes at distance 2.

are $2/3$ and $1/3$, respectively. Hence, the double tree is equivalent to a biased random walk on a line as in Fig. 7, which is not recurrent (see Levin and Peres [23]).

Since in the actual network and in the cut network there are no paths between i and j except the edge (i,j) (see Fig. 9 (a) and (b)),

$$r_{ij} = r_{ij}^{U_d} = 1.$$

Computing $r_{ij}^{L_d}$ is more involved. First, referring to Fig. 8, we note that, because of the symmetry of the network, the effective resistance between i and j in the shorted network (c), which is $r_{ij}^{L_d}$, is equivalent to the effective resistance in (d). Indeed, if we set potential $V_i = 1$ and $V_j = 0$, because of symmetry every yellow node has potential $1/2$. Thus, adding infinite conductance between all of them, i.e., shorting them, does not affect the current in the network (this procedure is also known in literature as *gluing*, see Levin and Peres [23]), and therefore the effective resistance. The network (d) is series-parallel, so that the effective resistance can be computed iteratively. Specifically, we refer to Fig. 9 to explain the recursion that leads to $r_{ij}^{L_d}$. From top to bottom, it is easy to see that the first network has effective resistance between the two blue nodes equal to 3. The second network is the parallel composition of two of these, in series with two single edges. This procedure is iteratively repeated $d-1$ times (in Fig. 9 only once, since $d = 2$), leading to a network that, composed in parallel with
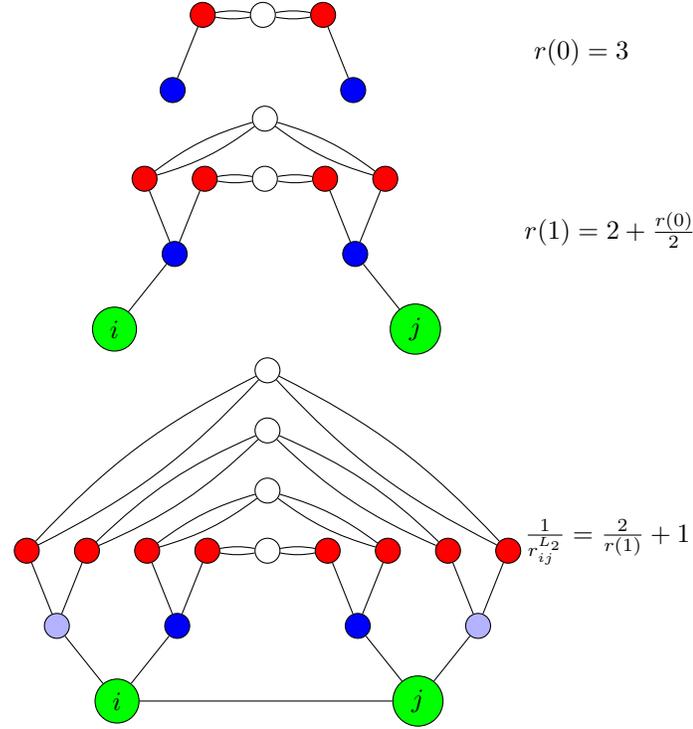
Figure 9: The network in Fig. 8(d) is series-parallel. Then, it can be obtained by recursively making parallel and series compositions of series-parallel networks.

a copy of itself and with a single edge, is $\mathcal{G}_{ij}^{L_d}$. Hence, $r_{ij}^{L_d}$ is the result of the following recursion.

$$\begin{cases} r(0) = 3, \\ r(n) = 2 + \frac{r(n-1)}{2}, \quad d > n \geq 1, \\ r_{ij}^{L_d} = (1 + \frac{2}{r(d-1)})^{-1}, \end{cases}$$

which has solution

$$\begin{cases} r(n) = (2^{d+2} - 1)/2^d, \quad d > n \geq 1, \\ r_{ij}^{L_d} = \frac{2^{d+1}-1}{2^{d+1}+2^d-1} \xrightarrow[d \to +\infty]{} \frac{2}{3}. \end{cases}$$