

Locality in Network Optimization*

Patrick Rebeschini[†] Sekhar Tatikonda[‡]

Abstract

In probability theory and statistics notions of correlation among random variables, decay of correlation, and bias-variance trade-off are fundamental. In this work we introduce analogous notions in optimization, and we show their usefulness in a concrete setting. We propose a general notion of correlation among variables in optimization procedures that is based on the sensitivity of optimal points upon (possibly finite) perturbations. We present a canonical instance in network optimization (the min-cost network flow problem) that exhibits locality, i.e., a setting where the correlation decays as a function of the graph-theoretical distance in the network. In the case of warm-start reoptimization, we develop a general approach to localize a given optimization routine in order to exploit locality. We show that the localization mechanism is responsible for introducing a bias in the original algorithm, and that the bias-variance trade-off that emerges can be exploited to minimize the computational complexity required to reach a prescribed level of error accuracy. We provide numerical evidence to support our claims.

Keywords: sensitivity of optimal points, decay of correlation, bias-variance, network flow, Laplacian, Green's function.

1 Introduction

Many problems in machine learning, networking, control, and statistics can be naturally posed in the framework of network-structured convex optimization. Given the huge problem size involved in modern applications, a lot of efforts have been devoted to the development and analysis of algorithms where the computation and communication are distributed over the network. A crucial challenge remains that of identifying the structural amount of information that each computation node needs to receive from the network to yield an approximate solution with a certain accuracy.

The literature on distributed algorithms in network optimization is gigantic, with much of the earlier seminal work contained in the textbook [6]. More recent work that explicitly relate the convergence behavior of the algorithms being considered to the network topology and size include — but it is certainly not limited to — [21, 29, 12] for first-order methods, and [36, 37] for second-order methods. Distributed algorithms are iterative in nature. In their synchronous implementations, at every iteration of the algorithm each computation

*Supported in part by the NSF Grant ECCS-1609484.

[†]P. Rebeschini was with the Department of Electrical Engineering, Yale University, New Haven, CT 06511, USA. He is now with the Department of Statistics, University of Oxford, Oxford, OX1 3LB, UK (e-mail: patrick.rebeschini@stats.ox.ac.uk).

[‡]S. Tatikonda is with the Department of Statistics and Data Science, Yale University, New Haven, CT 06511, USA (e-mail: sekhar.tatikonda@yale.edu).

node processes local information that come from its neighbors. Convergence analysis yields the number of iterations that guarantee these algorithms to meet a prescribed level of error accuracy. In turns, these results translate into bounds on the total amount of information processed by each node, as k iterations of the algorithm means that each node receives information coming from nodes that are at most k -hops away in the network. As different algorithms yield different rates of convergence, the bounds on the propagation of information that are so-derived are algorithm-dependent (and also depend on the error analysis being considered). As such, they do not capture structural properties of the optimization problem.

The main aim of this paper is to propose a general notion of “correlation” among variables in network optimization procedures that fundamentally characterizes the propagation of information across the network. This notion is based on the sensitivity of the optimal solution of the optimization problem upon perturbations of parameters locally supported on the network. This notion can be used to investigate “locality,” by which we mean problem instances where the correlation decays as a function of the natural distance in the network, so that, effectively, information only propagates across local portions of it. The phenomenon of locality characterizes situations where local perturbations are confined inside small regions of the network. How small the affected regions are, it depends on the particular problem instance at hand, and on the underlying graph topology. If locality is present, one could hope to exploit it algorithmically. In the case of localized perturbations, for instance, one could hope to localize known optimization routines so that only the affected regions are updated, hence yielding computational savings. While such a localization mechanism would introduce a structural “bias” with respect to the original algorithm — which updates every node in the network and not only the ones that are mostly affected by the perturbation — one could hope to exploit the “reduction in the variance” to achieve a prescribed level of error accuracy with a lower computational complexity. In the following we make all of this precise.

In this paper we define general notions of correlation, decay of correlation (locality), and bias-variance decomposition and trade-off in optimization, and we consider a concrete setting to illustrate their usefulness. The results that we present should be seen as an implementation of a general agenda that can be followed to establish and exploit locality in more general instances of network optimization.

This paper presents four main contributions discussed in separate sections, which we now summarize. Proofs are in the Appendices.

1) Sensitivity of optimal points. In Section 2 we present a theory on the sensitivity of optimal points for smooth convex problems with linear equality constraints upon (possibly finite) perturbations. We consider the problem of minimizing a smooth convex function $x \rightarrow f(x)$ subject to $Ax = b$, for a certain matrix A and vector $b \in \text{Im}(A)$, where $\text{Im}(A)$ denotes the image of A . We consider this problem as a function of the constraint vector b . We prove that if f is strongly convex, then the optimal point $b \rightarrow x^*(b)$ is continuously differentiable along $\text{Im}(A)$, and we explicitly characterize the effect that perturbations have on the optimal solution as a function of the objective function f , the constraint matrix A and vector b . Given a differentiable function $\varepsilon \in \mathbb{R} \rightarrow b(\varepsilon) \in \text{Im}(A)$, we show that the quantity $\frac{dx^*(b(\varepsilon))}{d\varepsilon}$ is a function of the Moore-Penrose pseudoinverse of the matrix $A\Sigma(b(\varepsilon))A^T$, where A^T is the transpose of A , and where $\Sigma(b)$ denotes the inverse of the Hessian of f evaluated at $x^*(b)$, namely, $\Sigma(b) := \nabla^2 f(x^*(b))^{-1}$. The literature on the sensitivity of optimal points (see Section 2 for a list of references) is typically only concerned with establishing infinites-

imal perturbations locally, i.e., on a neighborhood of a certain $b \in \text{Im}(A)$. On the other hand, the results that we present extend to *finite* perturbations as well, as we prove that the derivatives of the optimal point are continuous along $\text{Im}(A)$, and hence can be integrated to deal with finite perturbations. The workhorse behind our results is Hadamard’s global inverse function theorem [23], which in our setup yields necessary and sufficient conditions for the inverse of the KKT map to be continuously differentiable. Proofs are in Appendix A.

2) Notions of correlation in optimization. In Section 3 we provide an interpretation of the sensitivity theory previously developed in terms of notions of correlation among variables in optimization procedures, resembling analogous notions in probability theory. If the matrix A is full row rank, for instance, then the quantity $\frac{\partial x^*(b)_i}{\partial b_a}$ is well-defined and describes how much $x^*(b)_i$ — the i -th component of the optimal solution $x^*(b)$ — changes upon perturbation of b_a — the a -th component of the constraint vector b . We interpret $\frac{\partial x^*(b)_i}{\partial b_a}$ as a measure of the correlation between variables i and a in the optimization problem, and we are interested in understanding how this quantity behaves as a function of the geodesic distance between i and a . We motivate this terminology by establishing an analogy with the theory of correlations in probability, via the connection with Gaussian random variables (proofs are in Appendix B). We extend the notion of correlation beyond infinitesimal perturbations, and we show how our theory yields a first instance of comparison theorems for constrained optimization procedures, along the lines of the comparison theorems established in probability theory to capture stochastic decay of correlation and control the difference of high-dimensional distributions [11, 31].

In probability theory, decay of correlation characterizes the effective neighborhood dependency of random variables in a probabilistic network. Since the seminal work of Dobrushin [11], this concept has found many applications beyond statistical physics. Recently, it has been used to develop and prove convergence guarantees for fast distributed local algorithms for inference and decision problems on large networks in a wide variety of domains, for instance, probabilistic marginal inference [34], wireless communication [38], network learning [8], combinatorial optimization [16], and nonlinear filtering [32]. However, even in applications where the underlying problem does not involve randomness, decay of correlation is typically established upon endowing the model with a probabilistic structure, hence modifying the original problem formulation. Our results in network optimization, instead, show how locality can be described in a purely non-random setting, as a structural property of the original optimization problem. To the best of our knowledge, non-random notions of correlation in optimization have been previously considered only in [28], where the authors explicitly use the word “correlation” to denote the sensitivity of optimal points with respect to localized perturbations, and they analyze the correlation as a function of the natural distance in the graph. However, in their work correlation is regarded as a tool to prove convergence guarantees for the specific algorithm at hand (Min-Sum message-passing to solve *unconstrained* convex problems), and no general theory is built around it.

3) Locality: decay of correlation. As a paradigm for network optimization, in Section 4 we consider the problem of computing network flows. This is a fundamental problem that has been studied in various formulations by different communities. The min-cost network flow variant (where the objective function is typically chosen to be linear, although there are non-linear extensions, and the constraints also include inequalities) has been essential in the development of the theory of polynomial-times algorithms for optimizations (see [17] and references therein, or [1] for a book reference). In the case of quadratic

functions, the problem is equivalent to computing electrical flows, which is a fundamental primitive related to Laplacian solvers that has been extensively studied in computer science (see [35]) with also applications to learning problems ([5, 24], for instance). In the formulation we consider, a directed graph $\vec{G} = (V, \vec{E})$ is given (arrows indicate directed graphs/edges), with its structure encoded in the vertex-edge incidence matrix $A \in \mathbb{R}^{V \times \vec{E}}$. To each edge $e \in \vec{E}$ is associated a flow x_e with a cost $f_e(x_e)$, and to each vertex $v \in V$ is associated an external flow b_v . The network flow problem consists in finding the flow $x^*(b) \in \mathbb{R}^{\vec{E}}$ that minimizes the cost $f(x) := \sum_{e \in \vec{E}} f_e(x)$ and satisfies the conservation law $Ax = b$.

In this setting, the general sensitivity theory that we have previously developed allows to characterize the derivatives of the optimal flow in terms of graph Laplacians; in fact, in this case the matrix $A\Sigma(b)A^T$ corresponds to the Laplacian of an undirected weighted graph associated to \vec{G} , where each directed edge $e \in \vec{E}$ is given a weight $\Sigma(b)_{ee} > 0$. Exploiting a general connection between the Moore-Penrose pseudoinverse of graph Laplacians and the Green’s function of random walks on weighed graphs — which we present as standalone in Appendix C — we express the correlation term $\frac{dx^*(b(\varepsilon))_e}{d\varepsilon}$ in terms of *differences* of Green’s functions. Different graph topologies yield different decaying behaviors for this quantity, as a function of the geodesic distance $d(e, Z)$ between edge e and the set of vertices $Z \subseteq V$ where the perturbation is localized, namely, $Z := \{z \in V : \frac{db(\varepsilon)_z}{d\varepsilon} = 0\}$. In the case of expanders, we derive spectral bounds that show an exponential decay, with rate given by the second largest eigenvalue in magnitude of the diffusion random walk. In this case we establish various types of decay of correlation bounds, point-to-set and set-to-point. Appendix D contains the proofs of these results. For grid-like graphs, based on numerical simulations and on asymptotic results for the behavior of the Green’s function in infinite grids [26], we expect the correlation term to decay polynomially instead of exponentially.

4) Localized algorithms and bias-variance. In Section 5 we investigate applications of the framework that we propose to develop scalable computationally-efficient algorithms. To illustrate the main principle behind our reasoning, we consider the case when the solution $x^*(b)$ is given (also known as *warm-start* scenario) and we want to compute the solution $x^*(b+p)$ for the perturbed flow $b+p$, where p is a perturbation supported on a small localized subset $Z \subseteq V$. This setting belongs to the class of reoptimization problems typically studied in computer science and operations research (we could not to find previous literature on the setting that we consider; the case considered more frequently involves discrete problems and changes in the graph structure, as in [4]).

The decay of correlation property structurally exhibited by the min-cost network flow problem encodes the fact that when the external flow b is locally perturbed it suffices to recompute the solution only for the part of the network that is “mostly affected” by this perturbation, i.e., the set of edges that have a distance at most r from the perturbation set Z . Here, the radius r is tuned to meet the desired level of error tolerance given the size of the perturbation and given the graph topology being investigated. This allows us to show that it is possible to develop localized versions of canonical optimization algorithms that can exploit locality by only updating the edges in a subgraph of \vec{G} . We investigate regimes where localized algorithms can achieve computational savings against their global counterpart. The key behind these savings is the bias-variance decomposition that we give for the error of localized algorithms. This decomposition conveys the idea that by introducing some bias in a given optimization routine (in our setting, by truncating the problem size restricting the algorithm to a subgraph) one can diminish its variance and obtain

faster convergence rates. We illustrate this phenomenon theoretically and with numerical simulations for localized projected gradient descent. Proofs are in Appendix E.

Notation. For a given matrix M , let M^T be the transpose, M^{-1} be the inverse, and M^+ be the Moore-Penrose pseudoinverse. Let $\text{Ker}(M) := \{x : Mx = 0\}$ and $\text{Im}(M) := \{y : y = Mx \text{ for some } x\}$ be the kernel and image of M , respectively. Given an index set \mathcal{J} and subsets $K, L \subseteq \mathcal{J}$, if $M \in \mathbb{R}^{\mathcal{J} \times \mathcal{J}}$, let $M_{K,L} \in \mathbb{R}^{K \times L}$ denote the submatrix corresponding to the rows and columns of M indexed by K and L , respectively. Let I be the identity matrix, $\mathbb{1}$ the all-one vector (or matrix), and $\mathbb{0}$ the all-zero vector (or matrix), whose sizes are implied by the context. Given a vector $x \in \mathbb{R}^{\mathcal{J}}$, let $x_i \in \mathbb{R}$ be the component associated to $i \in \mathcal{J}$, and $x_K := (x_i)_{i \in K} \in \mathbb{R}^K$ the components associated to $K \subseteq \mathcal{J}$. Let $\|x\| := (\sum_{i \in \mathcal{J}} x_i^2)^{1/2}$ be the ℓ_2 -norm of x and $\|x\|_K := (\sum_{i \in K} x_i^2)^{1/2}$ be the localized ℓ_2 -norm on K . Clearly, $\|x\|_K = \|x_K\|$ and $\|x\|_{\mathcal{J}} = \|x\|$. We use the notation $|K|$ to denote the cardinality of K . If $\vec{G} = (V, \vec{E})$ is a directed graph with vertex set V and edge set \vec{E} , let $G = (V, E)$ be the undirected graph associated to \vec{G} , namely, $\{u, v\} \in E$ if and only if either $(u, v) \in \vec{E}$ or $(v, u) \in \vec{E}$.

2 Sensitivity of optimal points

Let \mathcal{V} be a finite set — to be referred to as the “variable set” — and let $f : \mathbb{R}^{\mathcal{V}} \rightarrow \mathbb{R}$ be a strictly convex function, twice continuously differentiable. Let \mathcal{F} be a finite set — to be referred to as the “factor set” — and let $A \in \mathbb{R}^{\mathcal{F} \times \mathcal{V}}$. Consider the following optimization problem over $x \in \mathbb{R}^{\mathcal{V}}$:

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && Ax = b, \end{aligned} \tag{1}$$

for $b \in \text{Im}(A) \subseteq \mathbb{R}^{\mathcal{F}}$, so that the feasible region is not empty. Throughout this paper we think of the function f and the matrix A as fixed, and we consider the solution of the optimization problem above as a function of the vector $b \in \text{Im}(A)$. By strict convexity, this problem has a unique optimal solution, that we denote by $x^*(b) := \arg\min\{f(x) : x \in \mathbb{R}^{\mathcal{V}}, Ax = b\}$.

Theorem 1 below provides a characterization of the way a perturbation of the constraint vector b on the subspace $\text{Im}(A)$ affects the optimal solution $x^*(b)$ in the case when the function f is strongly convex. The proof is given in Appendix A.

Theorem 1 (Sensitivity optimal point). *Let $f : \mathbb{R}^{\mathcal{V}} \rightarrow \mathbb{R}$ be a strongly convex function, twice continuously differentiable. Let $A \in \mathbb{R}^{\mathcal{F} \times \mathcal{V}}$. For $b \in \text{Im}(A)$, let $\Sigma(b) := \nabla^2 f(x^*(b))^{-1}$ and $D(b) := \Sigma(b)A^T(A\Sigma(b)A^T)^+$. Then, x^* is continuously differentiable along the subspace $\text{Im}(A)$, and given a differentiable function $\varepsilon \in \mathbb{R} \rightarrow b(\varepsilon) \in \text{Im}(A)$ we have*

$$\frac{dx^*(b(\varepsilon))}{d\varepsilon} = D(b(\varepsilon)) \frac{db(\varepsilon)}{d\varepsilon}.$$

Most of the literature on the sensitivity of optimal points for nonlinear programs investigates *local* results for infinitesimal perturbations, establishing the existence of a local neighborhood of the perturbed parameter(s) where the optimal point(s) has(have) some analytical properties such as continuity, differentiability, etc. Classical references are [19, 33, 13, 25]. Typically, the main tool used to establish this type of results is the

implicit function theorem applied to the first order optimality conditions (KKT map), which is a local statement (as such, it is flexible and it can be applied to a great variety of optimization problems). The sensitivity result that we present in Theorem 1 in the case of strongly convex functions, on the other hand, is based on Hadamard’s global inverse function theorem [23]. This is a *global* statement, as it shows that the optimal point x^* is continuously differentiable along the *entire* subspace $\text{Im}(A)$. This fact allows the use of the fundamental theorem of calculus to deal with finite perturbations, which is instrumental for the results developed in this paper (in particular, for the connection with comparison theorems in probability, Section 3.2, and for the results in Section 5). In the case of quadratic programs with linear constraints (the problem we consider can be thought of as an extension of this setting to strongly convex functions) there are many papers in parametric programming investigating the sensitivity of optimal points with respect to changes in the objective functions and constraints (see [30] and references therein). While most of these results are again local, some of them are closer to our approach and also address finite perturbations [7].

Theorem 1 characterizes the behavior of $x^*(b)$ upon perturbations of b along $\text{Im}(A) \subseteq \mathbb{R}^{\mathcal{F}}$. If the matrix A is full row rank, i.e., $\text{Im}(A) = \mathbb{R}^{\mathcal{F}}$, then the optimal point x^* is everywhere continuously differentiable, and we can compute its gradient. We have the following immediate corollary.

Corollary 1 (Sensitivity optimal point, full rank case). *Consider the setting of Theorem 1, with the matrix $A \in \mathbb{R}^{\mathcal{F} \times \mathcal{V}}$ having full row rank, i.e., $\text{Im}(A) = \mathbb{R}^{\mathcal{F}}$. Then, the function $b \in \mathbb{R}^{\mathcal{F}} \rightarrow x^*(b) \in \mathbb{R}^{\mathcal{V}}$ is continuously differentiable and*

$$\frac{dx^*(b)}{db} = D(b) = \Sigma(b)A^T(A\Sigma(b)A^T)^{-1}.$$

Remark 1. *Our goal is to present the simplest setting of interest where we can establish locality and illustrate the computational savings achieved by localized algorithms. The computational advantages provided by localization are already appreciable in the well-conditioned setting that we consider, as we explain in Section 5 below. For this reason, we are satisfied with the assumption of strong convexity in Theorem 1. The main tool behind Theorem 1 is Hadamard’s global inverse function theorem, which provides necessary and sufficient conditions for the inverse of a continuously differentiable function to be continuously differentiable. These conditions involve coercivity of the function of interest, along with a non-degeneracy condition on the determinant of the function. To relax the assumptions that we give in Theorem 1, one needs to consider a more refined application of Hadamard’s theorem to the KKT map.*

3 Notions of correlation in optimization

The sensitivity analysis presented in Section 2 suggests a natural notion of correlation between variables and factors in optimization, resembling notions of correlation among random variables in probability theory. If the matrix A is full row rank, then the quantity $\frac{\partial x^*(b)_i}{\partial b_a}$ is well-defined and it captures the interaction between variable $i \in \mathcal{V}$ and factor $a \in \mathcal{F}$ in the optimization procedure, and the quantity $D(b)_{ia}$ in Corollary 1 characterizes this correlation as a function of the constraint matrix A , the objective function f , and the optimal solution $x^*(b)$. Theorem 1 allows us to extend the notion of correlation between variables and factors to the more general case when the matrix A is not full rank. As an

example, let $b, p \in \text{Im}(A)$, and assume that p is supported on a subset $Z \subseteq \mathcal{F}$, namely, $p_a \neq 0$ if and only if $a \in Z$. Define $b(\varepsilon) := b + \varepsilon p$. Then, the quantity $\frac{dx^*(b(\varepsilon))_i}{d\varepsilon}$ measures how much a perturbation of the constraints in Z affects the optimal solution at $i \in \mathcal{V}$, hence it can be interpreted as a measure of the correlation between variable i and the factors in Z , which is characterized by the term $(D(b(\varepsilon))\frac{db(\varepsilon)}{d\varepsilon})_i = \sum_{a \in Z} D(b(\varepsilon))_{ia} p_a$ in Theorem 1.

We now discuss how these notions of correlation relate to analogous notions in probability.

3.1 Connection with Gaussian random variables

The resemblance between the notion of correlations in optimization and in probability is made explicit via the analogy to the theory of Gaussian random variables. Recall the following result (proofs of the results here discussed are given for completeness in Appendix B).

Proposition 1 (Conditional mean of Gaussian random variables). *Let \mathcal{V}, \mathcal{F} be two finite sets. Let $X \in \mathbb{R}^{\mathcal{V}}$ be a Gaussian random vector with mean $\mu \in \mathbb{R}^{\mathcal{V}}$ and covariance $\Sigma \in \mathbb{R}^{\mathcal{V} \times \mathcal{V}}$, possibly singular. Let $A \in \mathbb{R}^{\mathcal{F} \times \mathcal{V}}$ be given. Given a differentiable function $\varepsilon \in \mathbb{R} \rightarrow b(\varepsilon) \in \text{Im}(A)$, we have $\frac{d\mathbf{E}[X|AX=b(\varepsilon)]}{d\varepsilon} = \Sigma A^T (A \Sigma A^T)^{-1} \frac{db(\varepsilon)}{d\varepsilon}$. If Σ is invertible and A is full row rank, then for each $b \in \mathbb{R}^{\mathcal{F}}$ we have $\frac{d\mathbf{E}[X|AX=b]}{db} = \Sigma A^T (A \Sigma A^T)^{-1}$.*

Proposition 1 shows that for $i \in \mathcal{V}, a \in \mathcal{F}$, the quantity $\frac{\partial \mathbf{E}[X_i|AX=b]}{\partial b_a} = (\Sigma A^T (A \Sigma A^T)^{-1})_{ia}$ can be interpreted as a measure of correlation between the random variables X_i and $(AX)_a$, as it describes how much a perturbation of $(AX)_a$ impacts X_i , upon conditioning on AX . A similar interpretation can be given in optimization for the quantities in Corollary 1, with the difference that typically these quantities depend on $b \in \text{Im}(A)$, as they are functions of $x^*(b)$.

The sensitivity results that we derived in Section 2 also yield notions of correlation in optimization between variables. The following lemma, an immediate application of Corollary 1, shows that the local behavior of the optimal solution of the optimization problem (1) when we freeze some coordinates, upon perturbation of these coordinates, is analogous to the behavior of the conditional mean of a non-degenerate Gaussian random vector upon changing the coordinates we condition on. Recall (see the proof of Proposition 1) that if $X \in \mathbb{R}^{\mathcal{V}}$ is a Gaussian vector with mean $\mu \in \mathbb{R}^{\mathcal{V}}$ and positive definite covariance $\Sigma \in \mathbb{R}^{\mathcal{V} \times \mathcal{V}}$, then for $I \subseteq \mathcal{V}$ and $B := \mathcal{V} \setminus I$ $\frac{d\mathbf{E}[X_I|X_B=x_B]}{dx_B} = \Sigma_{I,B} (\Sigma_{B,B})^{-1}$.

Lemma 1 (Sensitivity with respect to boundary conditions). *Let $f : \mathbb{R}^{\mathcal{V}} \rightarrow \mathbb{R}$ be a strongly convex function, twice continuously differentiable. Let $I \subseteq \mathcal{V}$ be a nonempty set, and let $B := \mathcal{V} \setminus I$ not empty. Define the function $x_I^* : x_B \in \mathbb{R}^B \rightarrow x_I^*(x_B) := \text{argmin} \{f(x_I x_B) : x_I \in \mathbb{R}^I\}$. For $x_B \in \mathbb{R}^B$, let $H(x_B) := \nabla^2 f(x_I^*(x_B) x_B)$ and $\Sigma(x_B) := H(x_B)^{-1}$. Then, x_I^* is continuously differentiable and $\frac{dx_I^*(x_B)}{dx_B} = \Sigma(x_B)_{I,B} (\Sigma(x_B)_{B,B})^{-1} = - (H(x_B)_{I,I})^{-1} H(x_B)_{I,B}$.*

3.2 Comparison theorems

The connection between Theorem 1 and the theory of correlations in probability extends beyond infinitesimal perturbations. As previously discussed, Theorem 1 can be used to deal with finite perturbations, and so it can be interpreted as a comparison theorem to capture *uniform* correlations in optimization, along the lines of the comparison theorems

in probability theory to capture stochastic decay of correlation and control the difference of high-dimensional distributions (see the seminal work in [11], and [31] for generalizations).

To see this analogy, let us consider a simplified version of the Dobrushin comparison theorem that can be easily derived from the textbook version in [18], Theorem 8.20. Let I be a finite set, and let $\Omega := \prod_{i \in I} \Omega_i$ where Ω_i is a finite set for each $i \in I$. Define the projections $X_i : x \mapsto x_i$ for $x \in \Omega$ and $i \in I$. For any probability distribution μ on Ω , define the marginal $\mu_i(y) := \mu(X_i = y)$, and the conditional distribution $\mu_i^x(y) := \mu(X_i = y | X_{I \setminus \{i\}} = x_{I \setminus \{i\}})$. Define the total variation distance between two distributions ν and $\tilde{\nu}$ on Ω_i as $\|\nu - \tilde{\nu}\|_T := \frac{1}{2} \sum_{y \in \Omega_i} |\nu(y) - \tilde{\nu}(y)|$.

Theorem 2 (Dobrushin comparison theorem). *Let $\mu, \tilde{\mu}$ be probability distributions on Ω . For each $i, j \in I$, define $C_{ij} := \sup_{x, z \in \Omega: x_{I \setminus \{j\}} = z_{I \setminus \{j\}}} \|\mu_i^x - \mu_i^z\|_T$ and $\mathbf{b}_j := \sup_{x \in \Omega} \|\mu_j^x - \tilde{\mu}_j^x\|_T$, and assume that the Dobrushin condition holds: $\max_{i \in I} \sum_{j \in I} C_{ij} < 1$. Then the matrix sum $\mathbf{D} := \sum_{t \geq 0} \mathbf{C}^t$ is convergent, and for any $i \in I, y \in \Omega_i$, we have $\|\mu_i - \tilde{\mu}_i\|_T \leq \sum_{j \in I} \mathbf{D}_{ij} \mathbf{b}_j$.*

The Dobrushin coefficient C_{ij} is a (uniform) measure of the degree to which a perturbation of site j directly affects site i under the distribution μ . However, perturbing site j might also indirectly affect site i : it could affect another site k which in turn affects i , etc. The aggregate effect of a perturbation of site j on site i is captured by \mathbf{D}_{ij} . The quantity \mathbf{b}_j is a comparison term that measures the local difference at site j between μ and $\tilde{\mu}$ (in terms of the conditional distributions μ_j^x and $\tilde{\mu}_j^x$).

The formal analogy between the Dobrushin comparison theorem and the sensitivity results of Theorem 1 for the optimization problem (1) is made explicit by the fundamental theorem of calculus. This connection is easier to make if we assume that the matrix A has full row rank, and we consider the results in Corollary 1. In this setting, the optimal point x^* is everywhere continuously differentiable, and for each $i \in \mathcal{V}$, $b, \tilde{b} \in \mathbb{R}^{\mathcal{F}}$, $b \neq \tilde{b}$, we have $x^*(b)_i - x^*(\tilde{b})_i = \int_0^1 \frac{dx^*(\theta b + (1-\theta)\tilde{b})_i}{d\theta} d\theta = \sum_{a \in \mathcal{F}} D(b, \tilde{b})_{ia} (b_a - \tilde{b}_a)$, where $D(b, \tilde{b})_{ia} := \int_0^1 (\Sigma(b_\theta) A^T (A \Sigma(b_\theta) A^T)^{-1})_{ia} d\theta$ with $b_\theta := \theta b + (1-\theta)\tilde{b}$. If for each $i \in \mathcal{V}$ and $a \in \mathcal{F}$ we have $\sup_{b \in \mathbb{R}^{\mathcal{F}}} |(\Sigma(b) A^T (A \Sigma(b) A^T)^{-1})_{ia}| \leq D_{ia}$, then from the previous expression we find $|x^*(b)_i - x^*(\tilde{b})_i| \leq \sum_{a \in \mathcal{F}} D_{ia} |b_a - \tilde{b}_a|$, whose structure resembles the statement in Theorem 2. The quantity D_{ia} represents a uniform measure of the aggregate impact that a perturbation of the a -th component of the constraint vector b has to the i -th component of the optimal solution x^* , so the matrix D takes the analogous role of the matrix \mathbf{D} in Theorem 2 (and as we will see below in a concrete application, see Theorem 3, suitable series expansions of D yield the analogous of \mathbf{C}). The quantity $|b_a - \tilde{b}_a|$ is a comparison term that measures the local difference at factor a between b and \tilde{b} , resembling the role of \mathbf{b}_j in Theorem 2.

In the next section we investigate the notion of correlation just introduced in the context of network optimization, in a concrete instance when the constraints naturally reflect a graph structure, and we investigate the behavior of the correlations as a function of the natural distance in the graph.

4 Locality: decay of correlation

As a paradigm for network optimization, we consider the network flow problem that has been widely studied in various fields (see introduction). Consider a directed graph $\vec{G} := (V, \vec{E})$, with vertex set V and edge set \vec{E} , with no self-edges and no multiple edges.

Let $G = (V, E)$ be the undirected graph naturally associated with \vec{G} , that is, $\{u, v\} \in E$ if and only if either $(u, v) \in \vec{E}$ or $(v, u) \in \vec{E}$. Without loss of generality, assume that G is connected (otherwise we can treat each connected component on its own). For each $e \in \vec{E}$ let x_e denote the flow on edge e , with $x_e > 0$ if the flow is in the direction of the edge, $x_e < 0$ if the flow is in the direction opposite the edge. For each $v \in V$ let b_v be a given external flow on the vertex v : $b_v > 0$ represents a source where the flow enters the vertex, whereas $b_v < 0$ represents a sink where the flow leaves the vertex. Assume that the total of the source flows equals the total of the sink flows, that is, $\mathbb{1}^T b = \sum_{v \in V} b_v = 0$, where $b = (b_v)_{v \in V} \in \mathbb{R}^V$ is the flow vector. We assume that the flow satisfies a conservation equation so that at each vertex the total flow is zero. This conservation law can be expressed as $Ax = b$, where $A \in \mathbb{R}^{V \times \vec{E}}$ is the *vertex-edge incidence matrix* defined as $A_{ve} := 1$ if e leaves node v , $A_{ve} := -1$ if e enters node v , and $A_{ve} := 0$ otherwise.

For each edge $e \in \vec{E}$ let $f_e : \mathbb{R} \rightarrow \mathbb{R}$ be its associated cost function, assumed to be strongly convex and twice continuously differentiable. The network flow problem reads as problem (1) with $f(x) := \sum_{e \in \vec{E}} f_e(x_e)$. It is easy to see that since G is connected $\text{Im}(A)$ consists of all vectors orthogonal to the vector $\mathbb{1}$, i.e., $\text{Im}(A) = \{y \in \mathbb{R}^V : \mathbb{1}^T y = 0\}$. Henceforth, for each $b \in \mathbb{R}^V$ with $\mathbb{1}^T b = 0$, let $x^*(b)$ be the optimal flow.

We first apply the sensitivity theory developed in Section 2 to characterize the correlation between vertices (i.e., factors) and edges (i.e., variables) in the network flow problem. Then, we investigate the behavior of these correlations in terms of the natural distance on the graph G .

4.1 Correlation, graph Laplacians and Green's functions

In the setting of the network flow problem, Theorem 1 immediately allows us to characterize the derivatives of the optimal point x^* along the subspace $\text{Im}(A)$ as a function of the graph Laplacian [10]. For $b \in \mathbb{R}^V$ such that $\mathbb{1}^T b = 0$, let $\Sigma(b) := \nabla^2 f(x^*(b))^{-1} \in \mathbb{R}^{\vec{E} \times \vec{E}}$, which is a diagonal matrix with entries given by $\sigma(b)_e := \Sigma(b)_{ee} := (\frac{\partial^2 f_e(x^*(b)_e)}{\partial x_e^2})^{-1} > 0$. Each term $\sigma(b)_e$ is strictly positive as f_e is strongly convex by assumption. Let $W(b) \in \mathbb{R}^{V \times V}$ be the symmetric matrix defined, for each $u, v \in V$, as $W(b)_{uv} := \sigma(b)_e$ if $e = (u, v) \in \vec{E}$ or $e = (v, u) \in \vec{E}$, and $W(b)_{uv} := 0$ otherwise. Let $D(b) \in \mathbb{R}^{V \times V}$ be the diagonal matrix with entries given by $d(b)_v := D(b)_{vv} := \sum_{u \in V} W(b)_{vu}$, for $v \in V$. Let $L(b) := D(b) - W(b)$ be the graph Laplacian of the undirected weighted graph $(V, E, W(b))$, where to each edge $e = \{u, v\} \in E$ is associated the weight $W(b)_{uv}$. A direct application of Theorem 1 (upon choosing variable set $\mathcal{V} := \vec{E}$ and factor set $\mathcal{F} := V$, and noticing that $A\Sigma(b)A^T = L(b)$) shows that the derivatives of the optimal point x^* along $\text{Im}(A)$ can be expressed in terms of the Moore-Penrose pseudoinverse of $L(b)$. The connection between $L(b)^+$ and the Green's function of random walks with transition matrix $P(b) := D(b)^{-1}W(b)$ allows us to derive the following result (proofs are in Appendix C).

Theorem 3 (Sensitivity optimal flow). *Given $b \in \mathbb{R}^V$ with $\mathbb{1}^T b = 0$, let $D(b) := \Sigma(b)A^T L(b)^+$. The optimal network flow x^* is continuously differentiable along $\text{Im}(A)$, and given a differentiable function $\varepsilon \in \mathbb{R} \rightarrow b(\varepsilon) \in \text{Im}(A)$ we have $\frac{dx^*(b(\varepsilon))}{d\varepsilon} = D(b(\varepsilon)) \frac{db(\varepsilon)}{d\varepsilon}$. For $e = (u, v) \in \vec{E}$, we have*

$$\frac{dx^*(b(\varepsilon))_e}{d\varepsilon} = \sigma(b)_e \sum_{z \in V} \frac{1}{d(b)_z} \frac{db(\varepsilon)_z}{d\varepsilon} \sum_{t=0}^{\infty} (P(b)_{uz}^t - P(b)_{vz}^t).$$

Let $b, p \in \mathbb{R}^V$ such that $\mathbb{1}^T b = \mathbb{1}^T p = 0$, and assume that p is supported on a subset

$Z \subseteq V$, namely, $p_v \neq 0$ if and only if $v \in Z$. Define $b(\varepsilon) := b + \varepsilon p$. Then, as discussed in Section 3, the quantity $\frac{dx^*(b(\varepsilon))_e}{d\varepsilon}$ can be interpreted as a measure of the correlation between edge $e \in \vec{E}$ and the vertices in Z in the network flow problem. How does the correlation behave with respect to the graph distance between e and Z ? Theorem 3 shows that the correlation is controlled by the *difference* of the Green's function $\sum_{t=0}^{\infty} P(b)_{uz}^t$ with respect to two neighboring starting points u and v (note that the Green's function itself is infinite, as we are dealing with finite graphs). Different graph topologies yield different decaying behaviors for this quantity. In the case of expanders, we now derive spectral bounds that decay exponentially, with rate given by the second largest eigenvalue in magnitude of the diffusion random walk. For grid-like topologies, based on simulations and on asymptotic results for the behavior of the Green's function in infinite grids [26], we expect the correlation to decay polynomially rather than exponentially.

4.2 Decay of correlation for expanders

Let $n := |V|$ be the cardinality of V , and for each $b \in \text{Im}(A)$ let $-1 \leq \lambda_n(b) \leq \lambda_{n-1}(b) \leq \dots \leq \lambda_2(b) < \lambda_1(b) = 1$ be the real eigenvalues of $P(b)$.¹ Define $\lambda(b) := \max\{|\lambda_2(b)|, |\lambda_n(b)|\}$ and $\lambda := \sup_{b \in \text{Im}(A)} \lambda(b)$. For each $v \in V$, let $\mathcal{N}(v) := \{w \in V : \{v, w\} \in E\}$ be the set of node neighbors of v in the graph G . Let d be the graph-theoretical distance between vertices in the graph G , namely, $d(u, v)$ is the length of the shortest path between vertices $u, v \in V$. For subset of vertices $U, Z \subseteq V$, define $d(U, Z) := \min\{d(u, z) : u \in U, z \in Z\}$. For each subset of edges $\vec{F} \subseteq \vec{E}$, let $V_{\vec{F}} \subseteq V$ be the vertex set of the subgraph $(V_{\vec{F}}, \vec{F})$ of \vec{G} that is induced by the edges in \vec{F} . Recall the definition of the localized ℓ_2 -norm from Section 1. The following result attests that the correlation for the network flow problem is upper-bounded by a quantity that decays exponentially as a function of the distance in the graph, with rate given by λ . For graphs where λ does not depend on the dimension, i.e., expanders, Theorem 4 can be interpreted as a first manifestation of the decay of correlation principle (i.e., locality) in network optimization. The proof is given in Appendix D.

Theorem 4 (Decay of correlation for expanders). *Consider the setting defined above. Let $\varepsilon \in \mathbb{R} \rightarrow b(\varepsilon) \in \text{Im}(A)$ be a differentiable function such that for any $\varepsilon \in \mathbb{R}$ we have $\frac{db(\varepsilon)_v}{d\varepsilon} \neq 0$ if and only if $v \in Z$, for a given $Z \subseteq V$. Then, for any subset of edges $\vec{F} \subseteq \vec{E}$ and any $\varepsilon \in \mathbb{R}$, we have*

$$\left\| \frac{dx^*(b(\varepsilon))}{d\varepsilon} \right\|_{\vec{F}} \leq c \frac{\lambda^{d(V_{\vec{F}}, Z)}}{1 - \lambda} \left\| \frac{db(\varepsilon)}{d\varepsilon} \right\|_Z,$$

$$\text{with } c := \sup_{b \in \text{Im}(A)} \frac{\max_{v \in V_{\vec{F}}} \sqrt{2|\mathcal{N}(v) \cap V_{\vec{F}}|}}{\min_{v \in V_{\vec{F}}} d(b)_v} \max_{u, v \in V_{\vec{F}}} W(b)_{uv}.$$

Recall that $\left\| \frac{dx^*(b(\varepsilon))}{d\varepsilon} \right\|_{\vec{F}} \equiv (\sum_{e \in \vec{F}} (\frac{dx^*(b(\varepsilon))_e}{d\varepsilon})^2)^{1/2}$ and $\left\| \frac{db(\varepsilon)}{d\varepsilon} \right\|_Z \equiv (\sum_{v \in Z} (\frac{db(\varepsilon)_v}{d\varepsilon})^2)^{1/2}$. The bound in Theorem 4 controls the effect that a localized perturbation supported on a subset of vertices $Z \subseteq V$ has on a subset of edges $\vec{F} \subseteq \vec{E}$, as a function of the distance between \vec{F} and Z , i.e., $d(V_{\vec{F}}, Z)$ (note that we only defined the distance among vertices, not edges, and that this distance is with respect to the unweighted graph G). A key feature

¹This characterization of eigenvalues for random walks on connected weighted graphs follows from the Perron-Frobenius theory. See [27].

of Theorem 4 — which is essential for the results in Section 5 below — is that the bound presented does not depend on the cardinality of \vec{F} .

Theorem 4 controls the effect that a *single* localized perturbation (supported on multiple vertices, as it has to be that $|Z| \geq 2$ for the function $\varepsilon \in \mathbb{R} \rightarrow b(\varepsilon)$ to be on $\text{Im}(A)$) has on a *collection* of edges for the optimal solution, independently of the number of edges being considered. We refer to this type of decay of correlation as *set-to-point*. Analogously, it is possible to control the effect that *multiple* localized perturbations have on a *single* edge for the optimal solution, independently of the number of perturbations being considered. We refer to this type of decay of correlation as *point-to-set*. To illustrate in more detail these two types of decay of correlation, and for the sake of simplicity, we consider perturbations that are supported on exactly two vertices, corresponding to the endpoints of edges. Given $b \in \mathbb{R}^V$ such that $\mathbb{1}^T b = 0$, and $e = (u, v) \in \vec{E}$, we define the directional derivative of x^* along edge e evaluated at b as $\nabla_e x^*(b) := \frac{dx^*(b + \varepsilon(e_u - e_v))}{d\varepsilon} \Big|_{\varepsilon=0}$, where for each $v \in V$, $e_v \in \mathbb{R}^V$ is the vector defined as $(e_v)_w = 0$ if $w \neq v$ and $(e_v)_v = 1$. Then, we immediately have the following corollary of Theorem 4.

Corollary 2 (Set-to-point decay of correlation). *For $b \in \mathbb{R}^V$ with $\mathbb{1}^T b = 0$, $\vec{F} \subseteq \vec{E}$ and $e \in \vec{E}$, we have*

$$\|\nabla_e x^*(b)\|_{\vec{F}} \equiv \sqrt{\sum_{f \in \vec{F}} (\nabla_e x^*(b)_f)^2} \leq \sqrt{2c} \frac{\lambda^{d(V_{\vec{F}}, V_{\{e\}})}}{1 - \lambda},$$

where c is as in Theorem 4.

The key feature of the bound in Corollary 2 is that it does not depend on the cardinality of \vec{F} . Exploiting the symmetry of the identities involving the graph Laplacian, it is also easy to establish the following analogous result. The proof is given in Appendix D.

Lemma 2 (Point-to-set decay of correlation). *For $b \in \mathbb{R}^V$ with $\mathbb{1}^T b = 0$, $\vec{F} \subseteq \vec{E}$ and $f \in \vec{E}$, we have*

$$\sqrt{\sum_{e \in \vec{F}} (\nabla_e x^*(b)_f)^2} \leq \sqrt{2c'} \frac{\lambda^{d(V_{\vec{F}}, V_{\{f\}})}}{1 - \lambda},$$

with $c' = \sup_{b \in \text{Im}(A)} \frac{W(b)_{wz} \max_{v \in V_{\vec{F}}} \sqrt{2|\mathcal{N}(v) \cap V_{\vec{F}}|}}{\sqrt{\min\{d(b)_w, d(b)_z\}} \min_{v \in V_{\vec{F}}} \sqrt{d(b)_v}}$.

5 Localized algorithms and bias-variance

Let us consider the network flow problem defined in the previous section, for a certain external flow $b \in \mathbb{R}^V$ such that $\mathbb{1}^T b = 0$. Let $Z \subseteq V$ and choose $p \in \mathbb{R}^V$ supported on Z (i.e., $p_v \neq 0$ if and only if $v \in Z$) with $\mathbb{1}^T p = 0$. Assume that we perturb the external flow b by adding p . We want to address the following question: given knowledge of the solution $x^*(b)$ for the unperturbed problem, what is a computationally efficient algorithm to compute the solution $x^*(b + p)$ of the perturbed problem? The main idea that we want to exploit is that when locality holds and we can prove a decay of correlation property such as the one established in Theorem 4 for expander graphs, then a localized perturbation of the external flow will affect more the components of $x^*(b)$ that are close to the perturbed sites on Z . Hence, we expect that only a subset of the components of the solution around Z needs to be updated to meet a prescribed level of error tolerance, yielding savings on the computational complexity.

5.1 Local problem and localized algorithms

To formalize the argument given above, henceforth let $\vec{G}' = (V', \vec{E}')$ be a subgraph of $\vec{G} = (V, \vec{E})$ such that $Z \subseteq V'$. Let $G' = (V', E')$ be the undirected graph associated to \vec{G}' (see notation in Section 1), and assume that G' is connected. Define $V'^C := V \setminus V'$ and $\vec{E}'^C := \vec{E} \setminus \vec{E}'$. Let us consider the localized version of the network flow problem supported on the subgraph \vec{G}' . Let $A' := A_{V', \vec{E}'}$ be the submatrix of A corresponding to the rows indexed by V' and the columns indexed by \vec{E}' . For $b' \in \text{Im}(A') \subseteq \mathbb{R}^{V'}$, let $x'^*(b') \in \mathbb{R}^{\vec{E}'}$ denote the solution of the following problem over $x' \in \mathbb{R}^{\vec{E}'}$:

$$\begin{aligned} & \text{minimize} && f'(x') := \sum_{e \in \vec{E}'} f_e(x'_e) \\ & \text{subject to} && A'x' = b'. \end{aligned} \tag{2}$$

If locality holds and the subgraph \vec{G}' is large enough, we expect that the solution of the perturbed problem $x^*(b+p)$ will be close to the solution of the unperturbed problem $x^*(b)$ on \vec{E}'^C , and it will be substantially different on \vec{E}' . For this reason we investigate the performance of local iterative algorithms that operate only on the subgraph \vec{G}' , leaving the components supported on \vec{E}'^C unchanged.

Definition 1 (Local algorithm). *Given $b \in \mathbb{R}^V$ with $\mathbf{1}^T b = 0$, let $\mathcal{X}'_b := \{u \in \mathbb{R}^{\vec{E}} : (Au)_{V'^C} = b_{V'^C}\}$. A map $T'_b : \mathcal{X}'_b \rightarrow \mathcal{X}'_b$ defines a local algorithm on the subgraph $\vec{G}' = (V', \vec{E}')$ if the following two conditions hold for any choice of $x \in \mathcal{X}'_b$:*

- (i) $\lim_{t \rightarrow \infty} T_b^{t'}(x)_{\vec{E}'} = x'^*(b')$, $b' = b_{V'} - A_{V', \vec{E}'^C} x_{\vec{E}'^C}$;
- (ii) $T'_b(x)_{\vec{E}'^C} = x_{\vec{E}'^C}$.

This definition ensures that a local algorithm T'_b only updates the components of $x \in \mathcal{X}'_b$ supported on \vec{E}' and there converges to the solution of problem (2) with $b' = b_{V'} - A_{V', \vec{E}'^C} x_{\vec{E}'^C}$. The components that are left invariant on \vec{E}'^C play the role of boundary conditions. The algorithm that we propose to compute $x^*(b+p)$ given knowledge of $x^*(b)$ amounts to running a local algorithm on \vec{G}' for t iterations with boundary conditions $x^*(b)_{\vec{E}'^C}$, i.e., $T_{b+p}^{t'}(x^*(b))$. By definition

$$\lim_{t \rightarrow \infty} T_{b+p}^{t'}(x^*(b))_e = \begin{cases} x'^*(b_{V'} + p_{V'} - A_{V', \vec{E}'^C} x^*(b)_{\vec{E}'^C})_e & \text{if } e \in \vec{E}', \\ x^*(b)_e & \text{if } e \in \vec{E}'^C. \end{cases}$$

Designing a local algorithm to compute $x^*(b+p)$ is easy. Given any algorithmic procedure (for instance, a first-order method, a second-order method, or a primal-dual method), we can define its localized version by applying the algorithm to problem (2) with $b' = b_{V'} + p_{V'} - A_{V', \vec{E}'^C} x^*(b)_{\vec{E}'^C}$. We now provide a general analysis of the comparison between the performance of a given algorithm and its localized version.

5.2 Bias-variance decomposition and trade-off

The error committed by a chosen local algorithm T'_{b+p} on the subgraph \vec{G}' after $t \geq 1$ iterations is given by

$$\text{Error}(\vec{G}', t) := x^*(b+p) - T_{b+p}^{t'}(x^*(b)).$$

The analysis that we give is based on the error decomposition

$$\text{Error}(\vec{G}', t) = \text{Bias}(\vec{G}') + \text{Var}(\vec{G}', t),$$

with

$$\text{Bias}(\vec{G}') := x^*(b+p) - \left\{ \lim_{t \rightarrow \infty} T'_{b+p}{}^t(x^*(b)) \right\}, \quad (3)$$

$$\text{Var}(\vec{G}', t) := \left\{ \lim_{t \rightarrow \infty} T'_{b+p}{}^t(x^*(b)) \right\} - T'_{b+p}{}^t(x^*(b)). \quad (4)$$

This decomposition resembles the bias-variance decomposition in statistics, which motivates the choices of the terminology we use. In statistics, the term “bias” typically refers to the *approximation* error that is made from the simplifying assumptions built into the learning method, while the term “variance” refers to the *estimation* error that is made from the fluctuations of the learning method (trained on a given sample size) around its mean. Analogously, in our setting the bias term (3) represents the error that is made from the model restriction that we consider, namely, the localization of the chosen algorithmic procedure. The variance term (4) represents the error that is made by the deviation of the estimate given by the localized algorithm (at a given time t) from the optimal solution of the localized model. The bias term is algorithm-independent and it characterizes the error that we commit by localizing the optimization procedure per se, as a function of the subgraph \vec{G}' . The variance term depends on the algorithm that we run on \vec{G}' and on time.

We now show that in some regimes the bias introduced by localization can be exploited to lower the computational complexity associated to the variance term and yield savings for local algorithms. This bias-variance trade-off further motivates our analogy with statistics and the terminology we use.

Assume that we want to compare the computational complexity of a global algorithm versus its localized counterpart to achieve a prescribed error accuracy $\varepsilon > 0$. Let $t' := \min\{t > 0 : \|\text{Var}(\vec{G}', t)\| \leq \varepsilon\}$ be the minimal number of iterations that allows the localized algorithm $T'_{b+p}{}^t(x^*(b))$ to achieve a variance error (measured in the ℓ_2 -norm) less than ε . Let $\kappa(\vec{G}', 1/\varepsilon)$ be the computational cost needed to run the localized algorithm for t' iterations. Typically, $\kappa(\vec{G}', 1/\varepsilon)$ scales polynomially with $|V'|$ and $|\vec{E}'|$; it scales polynomially, logarithmically, or double-logarithmically with $1/\varepsilon$, depending on the regularity assumptions for the optimization problem and on the algorithmic procedure being used. For the sake of illustration, we assume that $\kappa(\vec{G}', 1/\varepsilon)$ is only a function of $|V'|$ and $1/\varepsilon$, so we write $\kappa(|V'|, 1/\varepsilon)$. The global algorithm ($\vec{G}' = \vec{G}$) has zero bias, so the computational cost $\kappa(|V|, 1/\varepsilon)$ guarantees to achieve $\|\text{Error}(\vec{G}, t)\| \leq \varepsilon$. The localized algorithm, on the other hand, has a non-zero bias due to localization. However, it will have a smaller computational cost for the variance term, as the algorithm runs on a subgraph of \vec{G} . To investigate the regime in which the added bias yields computational savings for the overall error term, we proceed as follows. By the triangle inequality, we have

$$\|\text{Error}(\vec{G}', t)\| \leq \|\text{Bias}(\vec{G}')\| + \|\text{Var}(\vec{G}', t)\|.$$

Assume that we can prove a bound of the following form:

$$\|\text{Bias}(\vec{G}')\| \leq \varphi \frac{1}{|V'|^\theta}, \quad (5)$$

for given universal constants $\varphi, \theta > 0$. Then, the choice $|V'| \geq (2\varphi/\varepsilon)^{1/\theta}$ guarantees that $\|\text{Bias}(\vec{G}')\| \leq \varepsilon/2$. By requiring $\|\text{Var}(\vec{G}', t)\| \leq \varepsilon/2$ we find that the computational cost $\kappa((2\varphi/\varepsilon)^{1/\theta}, 2/\varepsilon)$ will guarantee that the localized algorithm also achieves $\|\text{Error}(\vec{G}', t)\| \leq \varepsilon$. This argument shows that the localized algorithm is preferable when $\kappa((2\varphi/\varepsilon)^{1/\theta}, 2/\varepsilon) < \kappa(|V|, 1/\varepsilon)$, for instance. The regime where localized algorithms yield computational savings depends on a variety of factors: the regularity assumptions, the algorithm being chosen, the dimension of the original graph, and the required error tolerance. In the following we provide concrete settings where we can establish (5) both theoretically and numerically, and investigate the computational savings due to the bias-variance trade-off. The ideas here illustrated suggest a general framework to study the trade-off between accuracy and complexity for local algorithms in network optimization.

5.3 Well-conditioned setting and projected gradient descent

We now consider one of the simplest settings where we can establish locality and illustrate the computational savings that can be achieved by localizing a given optimization procedure. Henceforth, let each function f_e be α -strongly convex and β -smooth for some given parameters $\alpha, \beta \in (0, \infty)$, i.e., $0 < \alpha \leq \frac{d^2 f_e(x)}{dx^2} \leq \beta < \infty$ for any $x \in \mathbb{R}, e \in \vec{E}$. In this well-conditioned case it is well-known that gradient descent converges with a number of iterations that scales logarithmically with $1/\varepsilon$. Even in this favorable case, however, the computational savings achieved by a local algorithm can be considerable, as we now show.

For the sake of illustration, let us consider projected gradient descent. Let $\Pi_{\mathcal{X}}$ be the projection operator on a set \mathcal{X} , defined as $\Pi_{\mathcal{X}}(x) := \text{argmin}_{u \in \mathcal{X}} \|x - u\|$. The localized projected gradient descent on \vec{G}' is defined as follows.

Definition 2 (Localized projected gradient descent). *Given $x \in \mathcal{X}'_b$, define the set $\mathcal{X}'_b(x) := \{u \in \mathbb{R}^{\vec{E}'} : A_{V', \vec{E}'} u_{\vec{E}'} = b_{V'} - A_{V', \vec{E}'^C} x_{\vec{E}'^C}\}$. Localized projected gradient descent with step size $\eta > 0$ is the local algorithm defined by*

$$T'_b(x)_e = \begin{cases} \Pi_{\mathcal{X}'_b(x)}(x_{\vec{E}'} - \eta \nabla f'(x_{\vec{E}'}))_e & \text{if } e \in \vec{E}', \\ x_e & \text{if } e \in \vec{E}'^C. \end{cases}$$

When $\vec{G}' = \vec{G}$, this algorithm recovers the global algorithm applied to the whole graph \vec{G} . In the setting we consider, a classical result yields that projected gradient descent with step size $\eta = 1/\beta$ converges to the optimal solution for any starting point. This corresponds to condition (i) in Definition 1. The algorithm converges exponentially fast, with $\|T_b^{it}(x)_{\vec{E}'} - x'^*(b')_{\vec{E}'}\| \leq e^{-t/(2Q)} \|x_{\vec{E}'} - x'^*(b')_{\vec{E}'}\|$, where $Q = \beta/\alpha$ is the so-called *condition number* (see [9][Theorem 3.6], for instance). Under the (common) assumption that $\|x_{\vec{E}'} - x'^*(b')_{\vec{E}'}\| \leq R$, for a certain universal constant $R > 0$, the above convergence rate tells us that in order to reach a prescribed level of error accuracy ε , it is sufficient to run projected gradient descent for a number of iterations that does not depend on the dimension (nor on the topology) of the subgraph \vec{G}' it is applied to, and that only scales logarithmically in $1/\varepsilon$.

The bias-variance tradeoff exploits the fact that the computational cost per iteration *does* depend on the dimension, even if the number of iterations is dimension-free. In the case of projected gradient descent the cost per iteration is dominated by the cost of computing the projection step, and in general this cost scales polynomially with the graph

size. To be precise, it can be seen that

$$\begin{aligned} T'_{b+p}(x)_{\vec{E}'} &= (I - A'^T L'^+ A')(x_{\vec{E}'} - \eta \nabla f'(x_{\vec{E}'})) \\ &\quad + A'^T L'^+(b_{V'} + p_{V'} - A_{V', \vec{E}'^C} x_{\vec{E}'^C}), \end{aligned} \tag{6}$$

where $L' := A' A'^T$ is the graph Laplacian of the subgraph G' . Computing the pseudoinverse of L' exactly has a cost that scales like $O(|V'|^\omega)$, where $\omega > 2$ is the so-called matrix multiplication constant.² In this case we have that $\kappa(\vec{G}', 1/\varepsilon)$ scales like $O(|V'|^\omega + |\vec{E}'| \log(1/\varepsilon))$. If we consider graphs that have largest degree bounded above by a universal constant k , we have $|\vec{E}'| \leq (k/2)|V'|$ and we can use the argument given at the end of Section 5.2 to state the following result.

Proposition 2 (Computational cost, global versus localized). *Under the assumptions given in this section, assuming that (5) holds, projected gradient descent is guaranteed to compute a solution with accuracy ε with a cost that scales as:*

- *Global algorithm:* $O(|V|^\omega + |V| \log(1/\varepsilon))$;
- *Localized algorithm:* $O((1/\varepsilon)^{\omega/\theta} + (1/\varepsilon)^{1/\theta} \log(1/\varepsilon))$.

This result shows that the computational savings achieved by localized algorithms can be substantial even in a well-conditioned setting, in the case when $|V| \gg (1/\varepsilon)^{1/\theta}$, i.e., when the graph \vec{G} is very large compared to the inverse of the error tolerance $1/\varepsilon$. In the next two sections we investigate this regime both in theory (for expanders) and in simulations (for expanders and grids).

5.4 Expander graphs

In this section we assume that G is an expander graph. Using the decay of correlation property in Theorem 4, we provide upper bounds for the bias and variance terms defined in (3) and (4) as a function of the subgraph \vec{G}' and time t .

Let define the *inner boundary* of \vec{G}' as $\Delta(\vec{G}') := \{v \in V' : \mathcal{N}(v) \cap V'^C \neq \emptyset\}$, which represents the subset of vertices in V' that have at least one vertex neighbor outside V' (in the undirected graph G). Let $B \in \mathbb{R}^{V \times V}$ be the *vertex-vertex adjacency matrix* of the undirected graph $G = (V, E)$, which is the symmetric matrix defined as $B_{uv} := 1$ if $\{u, v\} \in E$, $B_{uv} := 0$ otherwise. Being real and symmetric, the matrix B has $n := |V|$ real eigenvalues which we denote by $\mu_n \leq \mu_{n-1} \leq \dots \leq \mu_2 \leq \mu_1$. Let $\mu := \max\{|\mu_2|, |\mu_n|\}$ be the second largest eigenvalue in magnitude of B . The next theorem yields bounds for the bias and variance error terms in the ℓ_2 -norm. The proof of this theorem is in Appendix E.

Theorem 5 (Error localized algorithm). *Let k_- and k_+ be, respectively, the minimum and*

²If we relax the requirement of performing an exact projection, we can consider efficient quasi-linear solvers to approximately compute the inverse of L' up to precision δ [22]. These solvers have a complexity that scales like $\tilde{O}(|\vec{E}'| \log |V'| \log(1/\delta))$. Even in this case, however, the savings of localized algorithms are still considerable in appropriate regimes (essentially, the same argument that we provide in the main text holds with $\omega \approx 1$).

maximum degree of G , $Q = \beta/\alpha$, and $\rho := \frac{Qk_+}{k_-} - 1 + \frac{Q}{k_-}\mu$. If $\rho < 1$, then

$$\begin{aligned} \|\text{Bias}(\vec{G}')\| &\leq \|p\| \gamma \frac{\rho^{d(\Delta(\vec{G}'), Z)}}{(1-\rho)^2} \mathbf{1}_{\vec{G}' \neq \vec{G}} \quad (\text{algorithm-free}) \\ \|\text{Var}(\vec{G}', t)\| &\leq \|p\| c \frac{e^{-t/(2Q)}}{1-\rho} \quad (\text{projected gradient descent}) \end{aligned}$$

where $\gamma := c(1 + c\sqrt{k_+ - 1})$ and $c := \frac{\sqrt{2k_+}}{k_-}$.

The bound for the bias decays exponentially with respect to the graph-theoretical distance (i.e., the distance in the undirected graph G) between the inner boundary of \vec{G}' , i.e., $\Delta(\vec{G}')$, and the region where the perturbation p is supported, i.e., $Z \subseteq V$. The rate is governed by the eigenvalue μ , the condition number Q , and the maximum/minimum degree of the graph. The bound for the variance decays exponentially with respect to the running time, with rate proportional to $1/Q$. We highlight that the constants appearing in the bounds in Theorem 5 do not depend on the choice of the subgraph \vec{G}' of \vec{G} , but depend only on μ , Q , k_+ , and k_- (as the proofs in Appendix E attest, a more refined analysis can yield better constants that do depend on the choice of \vec{G}'). In particular, the same constants apply for the analysis of the global algorithm. In this case, the bias term is zero, so that the error equals the variance component (see the indicator function in Theorem 5).

To investigate the computational savings achieved by localized gradient descent in the case of expanders, we will use Theorem 5 to derive a bound for the bias term as in (5) and then invoke Proposition 2. To this end, for the sake of simplicity, let G be a k -regular graph, where each vertex has k neighbors ($k_+ = k_- = k$). Let us introduce a collection of subgraphs that are centered on a given vertex and are parametrized by their radii. Namely, fix a vertex $v \in V$, let $V'_r := \{w \in V : d(v, w) \leq r\}$ denote the ball of radius $r > 0$ around vertex $v \in V$, and let $\vec{G}'_r := (V'_r, \vec{E}'_r)$ be the subgraph of \vec{G} that has vertex set V'_r and induced edge set \vec{E}'_r . Consider a perturbation vector $p \in \mathbb{R}^V$ that is supported on $Z := V'_z$, for a fixed $z \geq 1$. If we run the localized algorithm on \vec{G}'_r , with $r > z$, using the trivial bound $|V'_r| \leq k^r$, Theorem 5 yields bound (5) with $\varphi = \|p\| \gamma \frac{1}{(1-\rho)^2 \rho^z}$ and $\theta = \log(1/\rho)/\log k$. As $|E'_r| \leq (k/2)|V'_r|$, we have that the global projected gradient descent achieves error tolerance ε with a cost that scales like $O(|V|^\omega + (k/2)|V| \log(1/\varepsilon))$, while its localized counterpart achieves the same accuracy with a cost that scales like $O((2\varphi/\varepsilon)^\omega + (k/2)(2\varphi/\varepsilon)^{1/\theta} \log(2/\varepsilon))$.

Remark 2 (General graph topologies). *While the theoretical analysis that we present only holds for expanders, the main idea applies to any graph topology. In general, one needs to use Theorem 3 to establish locality results in the same spirit of Theorem 4 and Theorem 5, possibly establishing different rates of decay (i.e., not exponential) for different graphs, as discussed in Section 4.1 and as we will see in Section 5.5.*

5.5 Numerical Simulations

Let us consider the quadratic problem obtained with the choice $f_e(x) = \tau_e x^2/2$ for any $e \in \vec{E}$, where each τ_e is an independent sample from the uniform distribution in $[1, 2]$. The function $f = \sum_{e \in \vec{E}} f_e$ is α -strongly convex and β -smooth with $\alpha = 1$ and $\beta = 2$, as the Hessian $\nabla^2 f$ is diagonal with entries $\frac{d^2 f_e(x)}{dx^2} = \tau_e$. In this case the optimal solution of

the optimization problem can be computed analytically. If we define $\Sigma := (\nabla^2 f)^{-1}$ and let $\Sigma' := \Sigma_{V', V'}$ be the submatrix indexed by V' , we have $x^*(b) = \Sigma A^T (A \Sigma A^T)^+ b$ and $x'^*(b') = \Sigma' A'^T (A' \Sigma' A'^T)^+ b'$. Using these expressions we can compute the bias term (3). For an arbitrary $v \in V$, we take $Z = \{w \in V : d(v, w) \leq 1\}$. We draw the components of the vector b and the non-zero components of the vector p independently from the uniform distribution in $[-1, 1]$, imposing the conditions $\mathbf{1}^T b = 0$ and $\mathbf{1}^T p = 0$ (this is done by modifying an arbitrary component of the randomly-generated vectors to impose that the sum of the components equals zero). Given $r > 0$, let $\vec{G}'_r = (V'_r, \vec{E}'_r)$ be the subgraph of \vec{G} that includes all the vertices within a distance r from v .

We consider three graphs with 900 nodes each: a cycle, a two-dimensional square grid with periodic boundary conditions, and a 3-regular expander (uniformly sampled from the family of 3-regular graphs).

Figure 1 shows the behavior of the ℓ_2 -norm of the bias term as a function of the radius r and the size $|V'_r|$ (i.e., number of vertices) of the subgraph \vec{G}'_r for the three graph topologies of interest. Each topology gives rise to a different behavior. For the cycle graph, the bias error stays constant while $V'_r \neq V$ and drops to zero only when $V'_r = V$. This illustrates the fact that in the cycle there is no decay of correlation as perturbations do not dissipate. Borrowing again terminology from probability theory, we can say that the cycle represents a case of “long-range dependence.” For the two-dimensional grid, the bias decays polynomially with r while the subgraph \vec{G}'_r does not reach the boundaries of \vec{G} (as we conjectured, see Section 4.1), and then decays at a faster rate. For the expander, the bias decays exponentially with r , as we proved in Theorem 5. With respect to $|V'_r|$, the decay of the bias for both grids and expanders aligns with the polynomial ansatz in (5). This polynomial decay for expanders was proved in Section 5.4.




Figure 1: Typical realizations of $\|\text{Bias}(\vec{G}'_r)\|$ as a function of the radius r (left) and the size $|V'_r|$ (right) of the subgraph $\vec{G}'_r = (V'_r, \vec{E}'_r)$. The value of the bias for the expander is zero for any $r \geq 11$, as $V'_r = V$ in this case.

Figure 2 shows the behavior of the variance term for the localized projected gradient descent algorithm. As prescribed by classical results, and as we proved more precisely in the case of expanders (Theorem 5), the algorithm converges exponentially fast and the rate of convergence can be upper bounded by a quantity that does not depend on the graph topology nor on the graph size. As described in Section 5.3, the cost incurred by the algorithm to achieve a given accuracy for the variance is seen to increase polynomially with $|V'_r|$.




Figure 2: (Left) Typical realizations of $\log \|\text{Var}(\vec{G}'_r, t)\|$ for projected gradient descent as a function of the iteration step t , for different graph topologies and different choices of the radius r . We do not plot results for the cycle as in that case the algorithm converges in one iteration as there is a unique solution to $Ax = b$. We also plot the theoretical upper bound $-t/2Q$, with $Q = \beta/\alpha = 2$. (Right) Time of execution to run localized projected gradient descent to achieve $\|\text{Var}(\vec{G}'_r, t)\| \leq 10^{-7}$ as a function of $|V'_r|$. We use the `ginv` function from the MASS library in R to compute the pseudoinverse of the Laplacian matrix L' and run the algorithm as described in (6).

Finally, Figure 3 compares the computational cost of global and localized algorithms to achieve the same level of error accuracy. This plot aligns with the findings of Proposition 2 in the case of expanders. In particular, the plot shows that the computational savings achieved by localized algorithms are considerable in the regime $1/\varepsilon \ll |V|$.

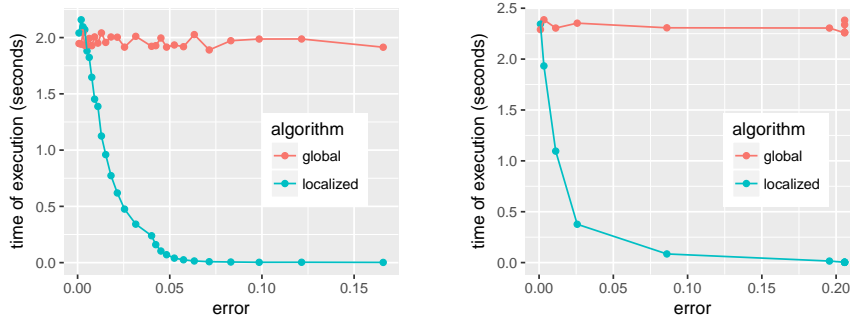


Figure 3: Time of execution to run projected gradient descent (global and localized) to achieve $\|\text{Error}(\vec{G}'_r, t)\| \leq \varepsilon + 10^{-7}$ as a function of the error tolerance parameter ε , for the grid (left) and the expander (right). As in the well-conditioned setting that we examine the algorithm drives the variance error to zero exponentially fast (in terms of number of iterations, see Figure 2), the error is dominated by the bias term. As a consequence, for the localized algorithm we adopt the strategy to choose the smallest radius r such that $\|\text{Bias}(\vec{G}'_r)\| \leq \varepsilon$ and then run the localized algorithms on \vec{G}'_r for the smallest number of iterations t such that $\|\text{Var}(\vec{G}'_r, t)\| \leq 10^{-7}$.

6 Conclusions

The main contribution of this paper is to derive a general analogy between natural concepts in probability and statistics (i.e., notions of correlation among random variables, decay of correlation, and bias-variance decomposition and trade-off) and similar notions that can be introduced in optimization. In this paper we have proposed notions of correlation that are based on the sensitivity of optimal points. We have illustrated how decay of correlation (locality) can be established in a canonical network optimization problem (min-cost network flow), and how it can be used to design local algorithms to compute the solution of a problem after localized perturbations are made to the system. In principle, the framework that we propose can be applied to *any* optimization problem. The key point is to establish decay of correlation for the problem at hand, deriving results that are analogous to the ones established in Section 4.2. In the case of the min-cost network flow problem, we showed that establishing locality reduces to bounding the discrete derivative of the Green's function of the diffusion random walk, as described in Theorem 3. We proved an exponential decay for the correlation in expander graphs (as a function of the distance from the perturbation), and we conjectured and provided numerical evidence for a polynomial decay in grid-like topologies. Once results on locality are established for the particular problem at hand, these results translate into a bound for the bias term of the error decomposition of localized algorithms, as the one that we give in Section 5.4. The analysis of the variance term, on the other hand, depends on the algorithm that one wants to localize. This part is not technically difficult, as it amounts to analyzing the performance of the chosen algorithm when applied to a subgraph with frozen boundary conditions. Establishing locality in more general settings and problems, and extending the ideas here presented to *cold-start* scenarios (where one wants to compute the optimal solution of an optimization problem starting from possibly any initial condition) remain open questions for future investigation.

References

- [1] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin, *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall, 1993.
- [2] Arthur E. Albert, *Regression and the moore-penrose pseudoinverse*, Mathematics in science and engineering, Academic Press, 1972.
- [3] David Aldous and James Allen Fill, *Reversible markov chains and random walks on graphs*, 2002, Unfinished monograph, recompiled 2014, available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- [4] Claudia Archetti, Luca Bertazzi, and M. Grazia Speranza, *Reoptimizing the traveling salesman problem*, Networks **42** (2003), no. 3, 154–159.
- [5] Mikhail Belkin and Partha Niyogi, *Semi-supervised learning on riemannian manifolds*, Machine Learning **56** (2004), no. 1, 209–239.
- [6] Dimitri P. Bertsekas and John N. Tsitsiklis, *Parallel and distributed computation: Numerical methods*, Athena Scientific, 1997.
- [7] J. C. G. Boot, *On sensitivity analysis in convex quadratic programming problems*, Operations Research **11** (1963), no. 5, 771–786.

- [8] Guy Bresler, Elchanan Mossel, and Allan Sly, *Reconstruction of markov random fields from samples: Some observations and algorithms*, Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques, Lecture Notes in Computer Science, vol. 5171, Springer Berlin Heidelberg, 2008, pp. 343–356.
- [9] Sébastien Bubeck, *Convex optimization: Algorithms and complexity*, Foundations and Trends® in Machine Learning **8** (2015), no. 3-4, 231–357.
- [10] F. R. K. Chung, *Spectral graph theory*, American Mathematical Society, 1997.
- [11] R. L. Dobrušin, *Definition of a system of random variables by means of conditional distributions*, Teor. Veroyatnost. i Primenen. **15** (1970), 469–497.
- [12] John C. Duchi, Alekh Agarwal, and Martin J. Wainwright, *Dual averaging for distributed optimization: Convergence analysis and network scaling.*, IEEE Trans. Automat. Contr. **57** (2012), no. 3, 592–606.
- [13] Anthony V. Fiacco, *Introduction to sensitivity and stability analysis in nonlinear programming / anthony v. fiacco*, Academic Press New York, 1983.
- [14] F. Fouss, A. Pirotte, J.-M. Renders, and M. Saerens, *Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation*, Knowledge and Data Engineering, IEEE Transactions on **19** (2007), no. 3, 355–369.
- [15] Tobias Friedrich and Thomas Sauerwald, *The cover time of deterministic random walks*, Computing and Combinatorics, vol. 6196, Springer Berlin Heidelberg, 2010, pp. 130–139.
- [16] David Gamarnik, David A. Goldberg, and Theophane Weber, *Correlation decay in random decision networks*, Mathematics of Operations Research **39** (2014), no. 2, 229–261.
- [17] David Gamarnik, Devavrat Shah, and Yehua Wei, *Belief propagation for min-cost network flow: Convergence and correctness*, Operations Research **60** (2012), no. 2, 410–428.
- [18] Hans-Otto Georgii, *Gibbs measures and phase transitions*, second ed., de Gruyter Studies in Mathematics, vol. 9, Walter de Gruyter & Co., Berlin, 2011.
- [19] Jürgen Guddat, *Stability in convex quadratic parametric programming*, Mathematische Operationsforschung und Statistik **7** (1976), no. 2, 223–245.
- [20] Roger A. Horn and Charles R. Johnson (eds.), *Matrix analysis*, Cambridge University Press, New York, NY, USA, 1986.
- [21] Björn Johansson, Maben Rabi, and Mikael Johansson, *A randomized incremental sub-gradient method for distributed optimization in networked systems*, SIAM J. on Optimization **20** (2009), no. 3, 1157–1170.
- [22] Ioannis Koutis, Gary L. Miller, and Richard Peng, *A nearly-m log n time solver for sdd linear systems*, FOCS, IEEE, 2011, pp. 590–598.

- [23] S.G. Krantz and H.R. Parks, *The implicit function theorem: History, theory, and applications*, The Implicit Function Theorem: History, Theory, and Applications, Birkhäuser, 2002.
- [24] Rasmus Kyng, Anup Rao, Sushant Sachdeva, and Daniel A. Spielman, *Algorithms for lipschitz learning on graphs.*, COLT, JMLR Workshop and Conference Proceedings, vol. 40, 2015, pp. 1190–1223.
- [25] Jerzy Kyparisis, *Uniqueness and differentiability of solutions of parametric nonlinear complementarity problems*, Mathematical Programming **36** (1986), no. 1, 105–113.
- [26] G.F. Lawler and V. Limic, *Random walk: A modern introduction*, Cambridge Studies in Advanced Mathematics, Cambridge University Press, 2010.
- [27] L. Lovász, *Random walks on graphs: A survey*, Combinatorics, Paul Erdos is Eighty **2** (1993), no. 1, 1–46.
- [28] Ciamac C. Moallemi and Benjamin Van Roy, *Convergence of min-sum message-passing for convex optimization*, Information Theory, IEEE Transactions on **56** (2010), no. 4, 2041–2050.
- [29] Damon Mosk-Aoyama, Tim Roughgarden, and Devavrat Shah, *Fully distributed algorithms for convex optimization problems*, SIAM Journal on Optimization **20** (2010), no. 6, 3260–3279.
- [30] H.X. Phu and N.D. Yen, *On the stability of solutions to quadratic programming problems*, Mathematical Programming **89** (2001), no. 3, 385–394.
- [31] Patrick Rebeschini and Ramon van Handel, *Comparison theorems for Gibbs measures*, Journal of Statistical Physics **157** (2014), no. 2, 234–281.
- [32] ———, *Can local particle filters beat the curse of dimensionality?*, Ann. Appl. Probab. **25** (2015), no. 5, 2809–2866.
- [33] Stephen M. Robinson, *Generalized equations and their solutions, Part I: Basic theory*, pp. 128–141, Springer Berlin Heidelberg, Berlin, Heidelberg, 1979.
- [34] Sekhar C. Tatikonda and Michael I. Jordan, *Loopy belief propagation and Gibbs measures*, Proc. UAI, vol. 18, 2002, pp. 493–500.
- [35] Nisheeth K. Vishnoi, *$Lx = b$* , Foundations and Trends® in Theoretical Computer Science **8** (2013), no. 1–2, 1–141.
- [36] E. Wei, A. Ozdaglar, and A. Jadbabaie, *A distributed Newton method for network utility maximization-I: Algorithm*, IEEE Transactions on Automatic Control **58** (2013), no. 9, 2162–2175.
- [37] ———, *A distributed Newton method for network utility maximization-II: Convergence*, IEEE Transactions on Automatic Control **58** (2013), no. 9, 2176–2188.
- [38] Dror Weitz, *Counting independent sets up to the tree threshold*, Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing, ACM, 2006, pp. 140–149.

- [39] F. Wu and C. Desoer, *Global inverse function theorem*, IEEE Transactions on Circuit Theory **19** (1972), 199–201.

A Hadamard's global inverse theorem

We prove Theorem 1. The proof relies on Hadamard's global inverse function theorem, which characterizes when a C^k function is a C^k diffeomorphism. Recall that a function from \mathbb{R}^m to \mathbb{R}^m is said to be C^k if it has continuous derivatives up to order k . A function is said to be a C^k *diffeomorphism* if it is C^k , bijective, and its inverse is also C^k .

The following important result characterizes when a C^k function is a C^k diffeomorphism.

Theorem 6 (Hadamard's global inverse function theorem). *Let Ψ be a C^k function from \mathbb{R}^m to \mathbb{R}^m . Then, f is a C^k diffeomorphism if and only if the following two conditions hold:*

1. *The determinant of the differential of Ψ is different from zero at any point, namely, $|\frac{d}{dz}\Psi(z)| \neq 0$ for any $z \in \mathbb{R}^m$.*
2. *The function Ψ is norm coercive, namely, for any sequence of points $z_1, z_2, \dots \in \mathbb{R}^m$ with $\|z_k\| \rightarrow \infty$ it holds $\|\Psi(z_k)\| \rightarrow \infty$ (for any choice of the norm $\|\cdot\|$, as norms are equivalent in finite dimension).*

Proof. See [39][Corollary of Lemma 2], for instance, or also [23] for a more general form of this inverse function to study diffeomorphisms on manifolds. \square

The following corollary is the backbone behind Theorem 1.

Lemma 3 (Diffeomorphism for Lagrangian multipliers map). *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a strongly convex function, twice continuously differentiable. Let $A \in \mathbb{R}^{m \times n}$ be a given matrix. Define the function Φ from $\mathbb{R}^n \times \mathbb{R}^m$ to $\mathbb{R}^n \times \mathbb{R}^m$ as*

$$\Phi(x, \nu) := \begin{pmatrix} \nabla f(x) + A^T \nu \\ Ax \end{pmatrix},$$

for any $x \in \mathbb{R}^n$, $\nu \in \mathbb{R}^m$. Then, the restriction of the function Φ to $\mathbb{R}^n \times \text{Im}(A)$ is a C^1 diffeomorphism.

Proof. Let us interpret Φ as the representation of a transformation \mathcal{T} in the standard basis of $\mathbb{R}^n \times \mathbb{R}^m$. Recall the orthogonal decomposition $\mathbb{R}^m = \text{Im}(A) \oplus \text{Ker}(A^T)$. Let the vectors $u_1, \dots, u_r \in \mathbb{R}^m$ form an orthogonal basis for $\text{Im}(A)$, where r is the rank of A , and let the vectors $v_1, \dots, v_{m-r} \in \mathbb{R}^m$ form an orthogonal basis for $\text{Ker}(A^T)$. Define the orthogonal matrix $Z = [u_1, \dots, u_r, z_1, \dots, z_{m-r}]$, which represents a change of basis in \mathbb{R}^m . As we have

$$\Phi(x, \nu) = \begin{pmatrix} \nabla f(x) + A^T Z Z^T \nu \\ Z Z^T Ax \end{pmatrix},$$

then the transformation \mathcal{T} is represented in the standard basis for \mathbb{R}^n and in the basis Z for \mathbb{R}^m by the following map

$$\tilde{\Phi}(x, \tilde{\nu}) := \begin{pmatrix} \nabla f(x) + \tilde{A}^T \tilde{\nu} \\ \tilde{A}x \end{pmatrix},$$

where $\tilde{A} := Z^T A$. In fact,

$$\tilde{\Phi}(x, Z^T \nu) = \begin{pmatrix} I & \mathbb{O} \\ \mathbb{O}^T & Z^T \end{pmatrix} \Phi(x, \nu),$$

where $I \in \mathbb{R}^{n \times n}$ is the identity matrix, and $\mathbb{O} \in \mathbb{R}^{n \times m}$ is the all-zero matrix. As $A^T Z = [A^T u_1, \dots, A^T u_r, \mathbb{O}_{n \times (m-r)}]$,

$$\tilde{A} = (A^T Z)^T = \begin{bmatrix} B \\ \mathbb{O}_{(m-r) \times n} \end{bmatrix},$$

where $B := [u_1, \dots, u_r]^T A \in \mathbb{R}^{r \times n}$. Therefore, the restriction of the transformation \mathcal{T} to the invariant subspace $\mathbb{R}^n \times \text{Im}(A)$ is represented in the standard basis for \mathbb{R}^n and in the basis $\{u_1, \dots, u_r\}$ for $\text{Im}(A)$ by the following map

$$\Psi(x, \xi) := \begin{pmatrix} \nabla f(x) + B^T \xi \\ Bx \end{pmatrix}.$$

As the function f is twice continuously differentiable, clearly the function Ψ is continuously differentiable, i.e., C^1 . We check that the two conditions of Theorem 6 are met.

The differential of Ψ evaluated at $(x, \xi) \in \mathbb{R}^n \times \mathbb{R}^r$ is

$$J(x, \xi) := \begin{pmatrix} \nabla^2 f(x) & B^T \\ B & \mathbb{O} \end{pmatrix}.$$

As f is strongly convex, $\nabla^2 f(x)$ is positive definite so invertible. Then, the determinant of the Jacobian can be expressed as $|J(x, \xi)| = |\nabla^2 f(x)| - B \nabla^2 f(x)^{-1} B^T$. As B has full row rank by definition, $B \nabla^2 f(x)^{-1} B^T$ is positive definite and we clearly have $|J(x, \xi)| \neq 0$.

To prove that the function Ψ is norm coercive, let us choose $\|\cdot\|$ to be the Euclidean norm and consider a sequence $(x_1, \xi_1), (x_2, \xi_2), \dots \in \mathbb{R}^n \times \mathbb{R}^r$ with $\|(x_k, \xi_k)\| \rightarrow \infty$. As for any $x \in \mathbb{R}^n, \xi \in \mathbb{R}^r$ we have $\|(x, \xi)\|^2 = \|x\|^2 + \|\xi\|^2$, clearly for the sequence to go to infinity one of the following two cases must happen:

- (a) $\|x_k\| \rightarrow \infty$;
- (b) $\|x_k\| \leq c$ for some $c < \infty$, $\|\xi_k\| \rightarrow \infty$.

Before we consider these two cases separately, let us note that, for any $x \in \mathbb{R}^n, \xi \in \mathbb{R}^r$,

$$\|\Psi(x, \xi)\|^2 = \|\nabla f(x) + B^T \xi\|^2 + \|Bx\|^2. \quad (7)$$

Let $\alpha > 0$ be the strong convexity parameter, and recall the following definition of strong convexity, for any $x, y \in \mathbb{R}^n$,

$$(\nabla f(x) - \nabla f(y))^T (x - y) \geq \alpha \|x - y\|^2. \quad (8)$$

- (a) Assume $\|x_k\| \rightarrow \infty$. Let P_{\parallel} be the projection operator on $\text{Im}(B^T)$, i.e., $P_{\parallel} := B^T (B B^T)^{-1} B$, and let $P_{\perp} = I - P_{\parallel}$ be the projection operator on $\text{Ker}(B)$, the orthogonal complement of $\text{Im}(B^T)$. As for any $x \in \mathbb{R}^n$ we have the decomposition $x = P_{\parallel} x + P_{\perp} x$ with $(P_{\parallel} x)^T P_{\perp} x = 0$, clearly $\|x\|^2 = \|P_{\parallel} x\|^2 + \|P_{\perp} x\|^2$. So, the condition $\|x_k\| \rightarrow \infty$ holds only if one of the two cases happens:

- (i) $\|P_{\parallel}x_k\| \rightarrow \infty$;
- (ii) $\|P_{\parallel}x_k\| \leq c$ for some $c < \infty$, $\|P_{\perp}x_k\| \rightarrow \infty$.

Consider (i) first. Let $x \in \mathbb{R}^n$ so that $P_{\parallel}x \neq 0$. As $BP_{\perp} = 0$, from (7) we have, for any $\xi \in \mathbb{R}^r$, $\|\Psi(x, \xi)\|^2 \geq \|Bx\|^2 = \|BP_{\parallel}x\|^2$, and $\min_{y \in \mathbb{R}^n: y \in \text{Im}(B^T), y \neq 0} \frac{y^T B^T B y}{\|y\|^2} \|P_{\parallel}x\|^2 = \lambda \|P_{\parallel}x\|^2$ is a lower bound, where λ is the minimum eigenvalue of $B^T B$ among those corresponding to the eigenvectors spanning the subspace $\text{Im}(B^T)$. Clearly, if $\lambda \neq 0$ (notice $\lambda \geq 0$ by definition) then the above yields that $\|\Psi(x_k, \xi_k)\| \rightarrow \infty$ whenever $\|P_{\parallel}x_k\| \rightarrow \infty$. To prove this, assume by contradiction that $\lambda = 0$. Then, there exists $y \in \mathbb{R}^n$ satisfying $y \in \text{Im}(B^T)$, $y \neq 0$, such that $B^T B y = \lambda y = 0$. As B^T has full column rank, the latter is equivalent to $B y = 0$ so that $P_{\perp}y = y \neq 0$, which contradicts the hypothesis that $y \in \text{Im}(B^T)$.

Consider now the case (ii). Decomposing the gradient on $\text{Im}(B^T)$ and its orthogonal subspace, from (7) we have, for any $x \in \mathbb{R}^n, \xi \in \mathbb{R}^r$, $\|\Psi(x, \xi)\|^2 \geq \|P_{\perp} \nabla f(x) + P_{\parallel} \nabla f(x) + B^T \xi\|^2 = \|P_{\perp} \nabla f(x)\|^2 + \|P_{\parallel} \nabla f(x) + B^T \xi\|^2$, so that $\|\Psi(x, \xi)\| \geq \|P_{\perp} \nabla f(x)\|$. Choosing $y = P_{\parallel}x$ in (8) we have

$$(P_{\perp} \nabla f(x) - P_{\perp} \nabla f(P_{\parallel}x))^T P_{\perp}x = (\nabla f(x) - \nabla f(P_{\parallel}x))^T P_{\perp}x \geq \alpha \|P_{\perp}x\|^2,$$

and applying Cauchy-Schwarz we get, for any x such that $P_{\perp}x \neq 0$, $\|P_{\perp} \nabla f(x)\| \geq \alpha \|P_{\perp}x\| - \|P_{\perp} \nabla f(P_{\parallel}x)\|$. By assumption f is twice continuously differentiable, so ∇f is continuous and it stays bounded on a bounded domain. Hence, we can conclude that $\|\Psi(x_k, \xi_k)\| \rightarrow \infty$ if $\|P_{\perp}x_k\| \rightarrow \infty$ with $(P_{\parallel}x_k)_{k \geq 1}$ bounded.

- (b) Assume $\|\xi_k\| \rightarrow \infty$ and $(x_k)_{k \geq 1}$ bounded. For any $\xi \in \mathbb{R}^r$, $\xi \neq 0$, we have

$$\|B^T \xi\|^2 = \frac{\xi^T B B^T \xi}{\|\xi\|^2} \|\xi\|^2 \geq \min_{y \in \mathbb{R}^r: y \neq 0} \frac{y^T B B^T y}{\|y\|^2} \|\xi\|^2 = \lambda_{\min} \|\xi\|^2,$$

where λ_{\min} is the minimum eigenvalue of $B B^T$, which is strictly positive as $B B^T$ is positive definite by the assumption that B has full row rank. From (7) we have $\|\Psi(x, \xi)\| \geq \|\nabla f(x) + B^T \xi\| \geq \|B^T \xi\| - \|\nabla f(x)\| \geq \sqrt{\lambda_{\min}} \|\xi\| - \|\nabla f(x)\|$, that, by continuity of ∇f , shows that $\|\Psi(x_k, \xi_k)\| \rightarrow \infty$ if $\|\xi_k\| \rightarrow \infty$ and $(x_k)_{k \geq 1}$ is bounded. □

We now present the proof of Theorem 1.

Proof of Theorem 1. The Lagrangian of the optimization problem is the function \mathcal{L} from $\mathbb{R}^V \times \mathbb{R}^{\mathcal{F}}$ to \mathbb{R} defined as $\mathcal{L}(x, \nu) := f(x) + \sum_{a \in \mathcal{F}} \nu_a (A_a^T x - b_a)$, where A_a^T is the a -th row of the matrix A and $\nu = (\nu_a)_{a \in \mathcal{F}}$ is the vector formed by the Lagrangian multipliers. Let us define the function Φ from $\mathbb{R}^V \times \mathbb{R}^{\mathcal{F}}$ to $\mathbb{R}^V \times \mathbb{R}^{\mathcal{F}}$ as

$$\Phi(x, \nu) := \begin{pmatrix} \nabla_x \mathcal{L}(x, \nu) \\ Ax \end{pmatrix} = \begin{pmatrix} \nabla f(x) + A^T \nu \\ Ax \end{pmatrix}.$$

For any fixed $\varepsilon \in \mathbb{R}$, as the constraints are linear, the Lagrange multiplier theorem says that for the unique minimizer $x^*(b(\varepsilon))$ there exists $\nu'(b(\varepsilon)) \in \mathbb{R}^{\mathcal{F}}$ so that

$$\Phi(x^*(b(\varepsilon)), \nu'(b(\varepsilon))) = \begin{pmatrix} 0 \\ b(\varepsilon) \end{pmatrix}. \quad (9)$$

As $A^T(\nu + \mu) = A^T\nu$ for each $\mu \in \text{Ker}(A^T)$, the set of Lagrangian multipliers $\nu^*(b(\varepsilon)) \in \mathbb{R}^{\mathcal{F}}$ that satisfies (9) is a translation of the null space of A^T . We denote the unique translation vector by $\nu^*(b(\varepsilon)) \in \text{Im}(A)$. By Hadamard's global inverse function theorem, as shown in Lemma 3, the restriction of the function Φ to $\mathbb{R}^{\mathcal{V}} \times \text{Im}(A)$ is a C^1 diffeomorphism, namely, it is continuously differentiable, bijective, and its inverse is also continuously differentiable. In particular, this means that the functions $x^* : b \in \text{Im}(A) \rightarrow x^*(b) \in \mathbb{R}^{\mathcal{V}}$ and $\nu^* : b \in \text{Im}(A) \rightarrow \nu^*(b) \in \text{Im}(A)$ are continuously differentiable along the subspace $\text{Im}(A)$. Differentiating both sides of (9) with respect to ε , we get

$$\begin{pmatrix} H & A^T \\ A & 0 \end{pmatrix} \begin{pmatrix} x' \\ \tilde{\nu} \end{pmatrix} = \begin{pmatrix} 0 \\ \frac{db(\varepsilon)}{d\varepsilon} \end{pmatrix},$$

where $H := \nabla^2 f(x^*(b(\varepsilon)))$, $x' := \frac{dx^*(b(\varepsilon))}{d\varepsilon}$, $\tilde{\nu} := \frac{d\nu^*(b(\varepsilon))}{d\varepsilon}$. As the function f is strongly convex, the Hessian $\nabla^2 f(x)$ is positive definite for every $x \in \mathbb{R}^{\mathcal{V}}$, hence it is invertible for every $x \in \mathbb{R}^{\mathcal{V}}$. Solving the linear system for x' first, from the first equation $Hx' + A^T\tilde{\nu} = 0$ we get $x' = -H^{-1}A^T\tilde{\nu}$. Substituting this expression in the second equation $Ax' = \frac{db(\varepsilon)}{d\varepsilon}$, we get $L\tilde{\nu} = -\frac{db(\varepsilon)}{d\varepsilon}$, where $L := AH^{-1}A^T$. The set of solutions to $L\tilde{\nu} = -\frac{db(\varepsilon)}{d\varepsilon}$ can be expressed in terms of the pseudoinverse of L as follows $\{\tilde{\nu} \in \mathbb{R}^{\mathcal{F}} : L\tilde{\nu} = -\frac{db(\varepsilon)}{d\varepsilon}\} = -L^+\frac{db(\varepsilon)}{d\varepsilon} + \text{Ker}(L)$. We show that $\text{Ker}(L) = \text{Ker}(A^T)$. We show that $L\nu = 0$ implies $A^T\nu = 0$, as the opposite direction trivially holds. In fact, let $A' := A\sqrt{H^{-1}}$, where $\sqrt{H^{-1}}$ is the positive definite matrix that satisfies $\sqrt{H^{-1}}\sqrt{H^{-1}} = H^{-1}$. The condition $L\nu = A'A^T\nu = 0$ is equivalent to $A'^T\nu \in \text{Ker}(A')$. At the same time, clearly, $A'^T\nu \in \text{Im}(A'^T)$. However, $\text{Ker}(A')$ is orthogonal to $\text{Im}(A'^T)$, so it must be $A'^T\nu = 0$ which implies $A^T\nu = 0$ as $\sqrt{H^{-1}}$ is positive definite. As $\text{Im}(L^+) = \text{Ker}(L)^\perp = \text{Ker}(A^T)^\perp = \text{Im}(A)$, we have that $\tilde{\nu} = -L^+\frac{db(\varepsilon)}{d\varepsilon}$ is the unique solution to $L\tilde{\nu} = -\frac{db(\varepsilon)}{d\varepsilon}$ that belongs to $\text{Im}(A)$. Substituting this expression into $x' = -H^{-1}A^T\tilde{\nu}$, we finally get $x' = H^{-1}A^TL^+\frac{db(\varepsilon)}{d\varepsilon}$. The proof follows as $\Sigma(b) = H^{-1}$. \square

B Correlation

In this section we provide the proofs of Proposition 1 and Lemma 1 in Section 3.

Proof of Proposition 1. Let \mathcal{J} be a finite set, and let $Z \in \mathbb{R}^{\mathcal{J}}$ be a Gaussian random vector with mean $\rho \in \mathbb{R}^{\mathcal{J}}$ and covariance $\Upsilon \in \mathbb{R}^{\mathcal{J} \times \mathcal{J}}$, not necessarily invertible. Let $\mathcal{V} \subseteq \mathcal{J}$ be not empty, and $\mathcal{F} := \mathcal{J} \setminus \mathcal{V}$ not empty. Recall that if $z_{\mathcal{F}} \in \text{Im}(\Upsilon_{\mathcal{F},\mathcal{F}})$, where $\Upsilon_{\mathcal{F},\mathcal{F}}$ is the submatrix of Υ indexed by the rows and columns referring to \mathcal{F} , then the conditional mean of $Z_{\mathcal{V}}$ given $Z_{\mathcal{F}} = z_{\mathcal{F}}$ is given by (see [2][Theorem 9.2.1], for instance): $\mathbf{E}[Z_{\mathcal{V}}|Z_{\mathcal{F}} = z_{\mathcal{F}}] = \rho_{\mathcal{V}} + \Upsilon_{\mathcal{V},\mathcal{F}}(\Upsilon_{\mathcal{F},\mathcal{F}})^+(z_{\mathcal{F}} - \rho_{\mathcal{F}})$. The statement of the proposition follows if we consider the Gaussian vector $Z := (X, AX)$ with $Z_{\mathcal{V}} := X$, $Z_{\mathcal{F}} := AX$, upon noticing that $\rho_{\mathcal{V}} = \mathbf{E}[X] = \mu$, $\rho_{\mathcal{F}} = \mathbf{E}[AX] = A\mu$, $\Upsilon_{\mathcal{V},\mathcal{F}} = \text{Cov}(X, AX) = \mathbf{E}[X(AX)^T] = \mathbf{E}[XX^T A^T] = \Sigma A^T$, and $\Upsilon_{\mathcal{F},\mathcal{F}} = \text{Cov}(AX, AX) = \mathbf{E}[AXX^T A^T] = A\Sigma A^T$. If Σ is invertible, i.e., positive definite, and A is full rank, then also $A\Sigma A^T$ is positive definite and invertible, so $(A\Sigma A^T)^+ = (A\Sigma A^T)^{-1}$. \square

Proof of Lemma 1. Given $b \in \mathbb{R}^B$, note that $x_I^*(b)$ corresponds to the components labeled by I of the solution $x^*(b)$ of the optimization problem (1), with $A := (0, I)$, where 0 is the all-zero matrix in $\mathbb{R}^{B \times I}$ and I is the identity matrix in $\mathbb{R}^{B \times B}$, and $x = (x_I, x_B)^T$. As the

matrix A is clearly full row rank, and $H = \nabla^2 f(x_I^*(b)) = \nabla^2 f(x^*(b))$, Corollary 1 yields

$$\frac{dx^*(b)}{db} = \Sigma A^T (A \Sigma A^T)^{-1} = \begin{pmatrix} \Sigma_{I,B} (\Sigma_{B,B})^{-1} \\ \mathbb{1} \end{pmatrix},$$

so that $\frac{dx_I^*(b)}{db} = \Sigma_{I,B} (\Sigma_{B,B})^{-1}$. The matrix identity $\Sigma_{I,B} (\Sigma_{B,B})^{-1} = -(H_{I,I})^{-1} H_{I,B}$ is standard textbook material. However, to show the statement involving the matrix H , we proceed from first principles, by applying the first order optimality condition to the restricted problem on I . Note that $x_I^*(x_B)$ is defined by the optimality conditions $\frac{df(x_I^*(x_B), x_B)}{dx_I} = 0$. Differentiating with respect to x_B , we get

$$\frac{d^2 f(x_I^*(x_B), x_B)}{dx_I^2} \frac{dx_I^*(x_B)}{dx_B} + \frac{d^2 f(x_I^*(x_B), x_B)}{dx_B dx_I} = 0,$$

or, equivalently, $H_{I,I} \frac{dx_I^*(x_B)}{dx_B} = -H_{I,B}$. The proof is concluded by inverting the matrix $H_{I,I}$, which is invertible as H is positive definite by assumption, so that any principal submatrix of it is also positive definite. \square

C Graph Laplacians and random walks

This section is self-contained and provides several connections between graph Laplacians and random walks on weighted graphs. In particular, the connection between pseudoinverses of graph Laplacians and Green's functions of random walks that we present in Lemma 4 below is the key that allows us to derive spectral bounds to prove the decay of correlation property in Section 4, as showed in Appendix D.

Throughout, let $G = (V, E, W)$ be a simple (i.e., no self-loops, and no multiple edges), connected, undirected, weighted graph, where to each edge $\{v, w\} \in E$ is associated a non-negative weight $W_{vw} = W_{wv} > 0$, and $W_{vw} = 0$ if $\{v, w\} \notin E$. Let D be a diagonal matrix with entries $d_v = D_{vv} = \sum_{w \in V} W_{vw}$ for each $v \in V$. For each vertex $v \in V$, let $\mathcal{N}(v) := \{w \in V : \{v, w\} \in E\}$ be the set of node neighbors of v . In this section we establish several connections between the graph Laplacian $L := D - W$ and the random walk $X := (X_t)_{t \geq 0}$ with transition matrix $P := D^{-1}W$. Henceforth, for each $v \in V$, let \mathbf{P}_v be the law of a time homogeneous Markov chain X_0, X_1, X_2, \dots on V with transition matrix P and initial condition $X_0 = v$. Analogously, denote by \mathbf{E}_v the expectation with respect to this law. The hitting time to the site $v \in V$ is defined as $T_v := \inf\{t \geq 0 : X_t = v\}$. Let π be the unique stationary distribution of the random walk, namely, $\pi^T P = \pi^T$. By substitution it is easy to check that $\pi_v := \frac{d_v}{\sum_{v \in V} d_v}$ for each $v \in V$. We adopt the notation $e_v \in \mathbb{R}^V$ to denote the vector whose only non-zero component equals 1 and corresponds to the entry associated to $v \in V$.

C.1 Restricted Laplacians and killed random walks

The connection between Laplacians and random walks that we present in Appendix C.2 below is established by investigating *restricted* Laplacians and *killed* random walks. Throughout this section, let $\bar{z} \in V$ be fixed. Let $\bar{V} := V \setminus \bar{z}$, $\bar{E} := E \setminus \{\{u, v\} \in E : u \in \bar{z} \text{ or } v \in \bar{z}\}$ and consider the graph $\bar{G} := (\bar{V}, \bar{E})$. Define \bar{W} and \bar{D} as the matrix obtained by removing the \bar{z} -th row and \bar{z} -th column from W and D , respectively. Then, $\bar{L} := \bar{D} - \bar{W} \in \mathbb{R}^{\bar{V} \times \bar{V}}$ represents the so-called restricted Laplacian that is obtained by removing the \bar{z} -th row and \bar{z} -th column from L . Define $\bar{P} := \bar{D}^{-1} \bar{W} \in \mathbb{R}^{\bar{V} \times \bar{V}}$. It is easy

to check that \bar{P} corresponds to the transition matrix of the transient part of the random walk that is obtained from X by adding a cemetery at site \bar{z} . Creating a cemetery at \bar{z} means modifying the walk X so that \bar{z} becomes a recurrent state, i.e., once the walk is in state \bar{z} it will go back to \bar{z} with probability 1. This is clearly done by replacing the \bar{z} -th row of P by a row with zeros everywhere but in the \bar{z} -th coordinate, where the entry is equal to 1. Note that \bar{L} does *not* correspond to the graph Laplacian of the graph $(\bar{V}, \bar{E}, \bar{W})$, as $\bar{D}_{vv} = \sum_{u \in \bar{V}} W_{vu}$ that does not equal $\sum_{u \in \bar{V}} \bar{W}_{vu}$ if $v \in \mathcal{N}(\bar{z})$. For the same reason, \bar{P} does *not* represent the ordinary random walk on $(\bar{V}, \bar{E}, \bar{W})$.

The relation between the transition matrix \bar{P} of the killed random walk and the law of the random walk X itself is made explicit in the next proposition. We omit the proof, which follows easily from the Markov property of X .

Proposition 3. *For $v, w \in \bar{V}$, $t \geq 0$, $\bar{P}_{vw}^t = \mathbf{P}_v(X_t = w, T_{\bar{z}} > t)$.*

The following proposition relates the inverse of the reduced Laplacian \bar{L} with the Green function of the killed random walk, namely, the function $(u, v) \in \bar{V} \times \bar{V} \rightarrow \sum_{t=0}^{\infty} \bar{P}_{uv}^t$, and with the hitting times of the original random walk X .

Proposition 4. *For each $v, w \in \bar{V}$, we have*

$$\bar{L}_{vw}^{-1} = \frac{1}{d_w} \sum_{t=0}^{\infty} \bar{P}_{vw}^t = \bar{L}_{wv}^{-1} \mathbf{P}_v(T_w < T_{\bar{z}}),$$

$$\text{and } \bar{L}_{wv}^{-1} = \frac{1}{d_w} \mathbf{E}_w \left[\sum_{t=0}^{T_{\bar{z}}} \mathbf{1}_{X_t=w} \right].$$

Proof. Let us first assume that \bar{G} is connected. The matrix \bar{P} is sub-stochastic as, clearly, if $v \notin \mathcal{N}(\bar{z})$ then $\sum_{w \in \bar{V}} \bar{P}_{vw} = 1$, while if $v \in \mathcal{N}(\bar{z})$ then $\sum_{w \in \bar{V}} \bar{P}_{vw} < 1$. Then \bar{P} is irreducible (in the sense of Markov chains, i.e., for each $v, w \in \bar{V}$ there exists t to that $\bar{P}_{vw}^t \neq 0$) and the spectral radius of \bar{P} is strictly less than 1 (see Corollary 6.2.28 in [20], for instance), so that the Neumann series $\sum_{t=0}^{\infty} \bar{P}^t$ converges. The Neumann series expansion for \bar{L}^{-1} yields $\bar{L}^{-1} = \sum_{t=0}^{\infty} (I - \bar{D}^{-1} \bar{L})^t \bar{D}^{-1} = \sum_{t=0}^{\infty} \bar{P}^t \bar{D}^{-1}$, or, entry-wise, $\bar{L}_{vw}^{-1} = \frac{1}{d_w} \sum_{t=0}^{\infty} \bar{P}_{vw}^t$. As $\bar{P}_{vw}^t = \mathbf{P}_v(X_t = w, T_{\bar{z}} > t)$ by Proposition 3, by the Monotone convergence theorem we can take the summation inside the expectation and get

$$\sum_{t=0}^{\infty} \bar{P}_{vw}^t = \sum_{t=0}^{\infty} \mathbf{E}_v[\mathbf{1}_{X_t=w} \mathbf{1}_{T_{\bar{z}} > t}] = \mathbf{E}_v \left[\sum_{t=0}^{T_{\bar{z}}} \mathbf{1}_{X_t=w} \right],$$

where in the last step we used that $X_{T_{\bar{z}}} = \bar{z}$ and $\bar{z} \neq w$. Recall that if S is a stopping time for the Markov chain $X := X_0, X_1, X_2, \dots$, then by the strong Markov property we know that, conditionally on $\{S < \infty\}$ and $\{X_S = w\}$, the chain $X_S, X_{S+1}, X_{S+2}, \dots$ has the same law as a time-homogeneous Markov chain $Y := Y_0, Y_1, Y_2, \dots$ with transition matrix P and initial condition $Y_0 = w$, and Y is independent of X_0, \dots, X_S . The hitting times T_w and $T_{\bar{z}}$ are two stopping times for X , and so is their minimum $S := \min\{T_w, T_{\bar{z}}\}$. As either $X_S = w$ or $X_S = \bar{z}$, we have

$$\mathbf{E}_v \left[\sum_{t=0}^{T_{\bar{z}}} \mathbf{1}_{X_t=w} \right] = \mathbf{E}_v \left[\sum_{t=0}^{T_{\bar{z}}} \mathbf{1}_{X_t=w} \middle| X_S = w \right] \mathbf{P}_v(X_S = w),$$

where we used that, conditionally on $\{X_S = \bar{z}\} = \{T_w > T_{\bar{z}}\}$, clearly $\sum_{t=0}^{T_{\bar{z}}} \mathbf{1}_{X_t=w} = 0$. Conditionally on $\{X_S = w\} = \{T_w < T_{\bar{z}}\} = \{S = T_w\}$, we have $T_{\bar{z}} = S + \inf\{t \geq 0 :$

$X_{S+t} = \bar{z}$ }, and the strong Markov property yields (note that the event $\{S < \infty\}$ has probability one from any starting point, as the graph G is connected and the chain will almost surely eventually hit either w or \bar{z}) that $\mathbf{E}_v[\sum_{t=0}^{T_{\bar{z}}} \mathbf{1}_{X_t=w} | X_S = w]$ can be written as $\mathbf{E}_v[\sum_{t=0}^{\inf\{t \geq 0: X_{S+t} = \bar{z}\}} \mathbf{1}_{X_{S+t}=w} | X_S = w] = \mathbf{E}_w[\sum_{t=0}^{T_{\bar{z}}} \mathbf{1}_{X_t=w}]$. Putting everything together we have $\bar{L}_{vw}^{-1} = \frac{1}{d_w} \mathbf{E}_w[\sum_{t=0}^{T_{\bar{z}}} \mathbf{1}_{X_t=w}] \mathbf{P}_v(T_w < T_{\bar{z}})$. As $\mathbf{P}_w(T_w < T_{\bar{z}}) = 1$, clearly $\bar{L}_{ww}^{-1} = \frac{1}{d_w} \mathbf{E}_w[\sum_{t=0}^{T_{\bar{z}}} \mathbf{1}_{X_t=w}]$ so that $\bar{L}_{vw}^{-1} = \bar{L}_{ww}^{-1} \mathbf{P}_v(T_w < T_{\bar{z}})$. The argument just presented extends easily to the case when \bar{G} is not connected. In fact, in this case the matrix \bar{P} has a block structure, where each block corresponds to a connected component and to a sub-stochastic submatrix, so that the argument above can be applied to each block separately. \square

The following result relates the inverse of the reduced Laplacian \bar{L} with the pseudoinverse of the Laplacian L , which we denote by L^+ . It is proved in [14][Appendix B].

Proposition 5. For $v, w \in \bar{V}$, $\bar{L}_{vw}^{-1} = (e_v - e_z)^T L^+ (e_w - e_z)$.

Proposition 4 and Proposition 5 allow us to relate the quantity L^+ to the *difference* of the Green's function of the random walk X , as we discuss next.

C.2 Laplacians and random walks

We now relate the Moore-Penrose pseudoinverse of the Laplacian $L := D - W$ with the Green's function $(u, v) \in V \times V \rightarrow \sum_{t=0}^{\infty} P_{uv}^t = \mathbf{E}_u[\sum_{t=0}^{\infty} \mathbf{1}_{X_t=v}]$ of the random walk, which represents the expected number of times the Markov chain X visits site v when it starts from site u . Notice that as the graph G is finite and connected, then the Markov chain X is recurrent and the Green's function itself equals infinity for any $u, v \in V$. In fact, the following result involves *differences* of the Green's function, not the Green's function itself.

Lemma 4. For any $u, v, w, z \in V$, we have

$$(e_u - e_v)^T L^+ (e_w - e_z) = \sum_{t=0}^{\infty} (e_u - e_v)^T P^t \left(\frac{e_w}{d_w} - \frac{e_z}{d_z} \right).$$

Proof. Using first Proposition 5 and then Proposition 4 we obtain, for any $u, v, w, z \in V$ (choose \bar{z} to be z in Appendix C.1), $(e_u - e_v)^T L^+ (e_w - e_z) = (e_u - e_z)^T L^+ (e_w - e_z) - (e_v - e_z)^T L^+ (e_w - e_z) = \bar{L}_{uw}^{-1} - \bar{L}_{vw}^{-1} = (e_w - e_z)^T L^+ (e_w - e_z) \{ \mathbf{P}_u(T_w < T_z) - \mathbf{P}_v(T_w < T_z) \}$. From (3.27) in the proof of Proposition 3.10 in Chapter 3 in [3], upon identifying $v \rightarrow u, x \rightarrow v, v_0 \rightarrow w, a \rightarrow z$, we immediately have the following relation between the difference of potentials and hitting times of the random walk X : $\mathbf{P}_u(T_w < T_z) - \mathbf{P}_v(T_w < T_z) = \pi_w \mathbf{P}_w(T_z < T_w^+) \{ \mathbf{E}_u T_z - \mathbf{E}_v T_z + \mathbf{E}_v T_w - \mathbf{E}_u T_w \}$, where $\pi_v := \frac{d_v}{\sum_{v \in V} d_v}$ is the v -th component of the stationary distribution of the random walk X , and $T_v^+ := \inf\{t \geq 1 : X_t = v\}$. From Corollary 8 in Chapter 2 in [3], we have

$$\pi_w \mathbf{P}_w(T_z < T_w^+) = \begin{cases} \frac{1}{\mathbf{E}_w T_z + \mathbf{E}_z T_w} & \text{if } w \neq z, \\ \pi_w & \text{if } w = z, \end{cases}$$

and we recall the connection between commute times and effective resistance (see, for example, Corollary 3.11 in [3]): $\mathbf{E}_w T_z + \mathbf{E}_z T_w = (e_w - e_z)^T L^+ (e_w - e_z) \sum_{v \in V} d_v$. Lemma 3.3 in [15] yields $\mathbf{E}_u T_z - \mathbf{E}_v T_z = \frac{1}{\pi_z} \sum_{t=0}^{\infty} (P_{vz}^t - P_{uz}^t)$, and $\mathbf{E}_u T_w - \mathbf{E}_v T_w = \frac{1}{\pi_w} \sum_{t=0}^{\infty} (P_{vw}^t - P_{uw}^t)$. The proof of the lemma follows by combining everything together. \square

We have the following corollary of Lemma 4, which immediately yields the proof of Theorem 3 in Section 4.

Corollary 3. For $u, v \in V$, $f = (f_z)_{z \in V} \in \mathbb{R}^V$ with $\mathbf{1}^T f = 0$,

$$(e_u - e_v)^T L^+ f = \sum_{z \in V} \sum_{t=0}^{\infty} (P_{uz}^t - P_{vz}^t) \frac{f_z}{d_z}.$$

Proof. From Lemma 4, by summing the quantity $(e_u - e_v)^T L^+(e_w - e_z)$ over $z \in V$, recalling that $\sum_{z \in V} e_z = \mathbf{1}$ and $L^+ \mathbf{1} = 0$ we have $(e_u - e_v)^T L^+ e_w = \sum_{t=0}^{\infty} (P_{uw}^t - P_{vw}^t) \frac{1}{d_w} - \frac{1}{|V|} \sum_{z \in V} \sum_{t=0}^{\infty} (P_{uz}^t - P_{vz}^t) \frac{1}{d_z}$. The identity in the statement of the Lemma follows easily as $f = \sum_{w \in V} f_w e_w$ and $\sum_{w \in V} f_w = 0$ by assumption. \square

D Decay of correlation

This appendix is devoted to the proofs of the decay of correlation properties stated in Section 4, namely, Theorem 4 (set-to-point) and Lemma 2 (point-to-set). These proofs rely on the sensitivity analysis for the network flow problem established in Theorem 3.

Henceforth, consider the general setting introduced in Appendix C, and let d denote the graph-theoretical distance on G : that is, $d(u, v)$ denotes the length of the shortest path between vertex u and vertex v . Note that $d(u, v) = \inf\{t \geq 0 : P_{uv}^t \neq 0\}$, as we assumed that to each edge $\{v, w\} \in E$ is associated a non-negative weight $W_{vw} = W_{wv} > 0$. Let $n := |V|$, and let $-1 \leq \lambda_n \leq \lambda_{n-1} \leq \dots \leq \lambda_2 < \lambda_1 = 1$ be the eigenvalues of P . Define $\lambda := \max\{|\lambda_2|, |\lambda_n|\}$.

The backbone behind the proof of Theorem 4 and Lemma 2 is given by the following lemma.

Lemma 5. For any $U, Z \subseteq V$ and $(f_z)_{z \in Z} \in \mathbb{R}^Z$ we have

$$\sqrt{\frac{1}{2} \sum_{\substack{u, v \in U: \\ \{u, v\} \in E}} \left(\sum_{z \in Z} \sum_{t=0}^{\infty} (P_{uz}^t - P_{vz}^t) f_z \right)^2} \leq \alpha \frac{\lambda^{d(U, Z)}}{1 - \lambda} \sqrt{\sum_{z \in Z} f_z^2 d_z},$$

$$\text{with } \alpha := \frac{\max_{v \in U} \sqrt{2|\mathcal{N}(v) \cap U|}}{\min_{v \in U} \sqrt{d_v}}.$$

Proof. Let $\Gamma := D^{1/2} P D^{-1/2} = D^{-1/2} W D^{-1/2}$. This matrix is clearly similar to P and symmetric. Let denote by ψ_n, \dots, ψ_1 the orthonormal eigenvectors of Γ corresponding to the eigenvalues $\lambda_n \leq \lambda_{n-1} \leq \dots \leq \lambda_2 \leq \lambda_1$, respectively. By substitution, it is easy to check that $\sqrt{\pi} \equiv (\sqrt{\pi_v})_{v \in V}$ is an eigenvector of Γ with eigenvalue equal to 1, where we recall that $\pi_v = d_v / \sum_{v \in V} d_v$. Since this eigenvector has positive entries, it follows by the Perron-Frobenius theory that $-1 \leq \lambda_n \leq \lambda_{n-1} \leq \dots \leq \lambda_2 < \lambda_1 = 1$ and that $\psi_1 = \sqrt{\pi}$. As $\Gamma = \sum_{k=1}^n \lambda_k \psi_k \psi_k^T$, by the orthonormality of the eigenvectors we have, for $t \geq 0$, $u, z \in V$,

$$P_{uz}^t = (D^{-1/2} \Gamma^t D^{1/2})_{uz} = \pi_z + \sum_{k=2}^n \lambda_k^t \psi_{ku} \psi_{kz} \sqrt{\frac{d_z}{d_u}},$$

where $\psi_{ku} \equiv (\psi_k)_u$ is the u -th component of ψ_k . As $P_{uz}^t = 0$ if $d(u, z) > t$, we have $P_{uz}^t - P_{vz}^t = \mathbf{1}_{d(U, Z) \leq t} (P_{uz}^t - P_{vz}^t)$ for any $u, v \in U, z \in Z$. Hence, for any $u, v \in U$, let

$g_{uv} := \sum_{z \in Z} \sum_{t=0}^{\infty} (P_{uz}^t - P_{vz}^t) f_z$, which equals

$$g_{uv} = \sum_{k=2}^n \left(\frac{\psi_{ku}}{\sqrt{d_u}} - \frac{\psi_{kv}}{\sqrt{d_v}} \right) \sum_{z \in Z} \psi_{kz} \sqrt{d_z} f_z \sum_{t=d(U,Z)}^{\infty} \lambda_k^t.$$

As $\lambda < 1$ by assumption, the geometric series converges for any $k \neq 1$. If we define $h_u := \sum_{k=2}^n \frac{\lambda_k^{d(U,Z)}}{1-\lambda_k} \frac{\psi_{ku}}{\sqrt{d_u}} \sum_{z \in Z} \psi_{kz} \sqrt{d_z} f_z$ for each $u \in V$, we have $g_{uv} = h_u - h_v$, and the triangle inequality for the ℓ_2 -norm yields $(\sum_{u,v \in U: \{u,v\} \in E} g_{uv}^2)^{1/2} \leq 2(\sum_{u,v \in U: \{u,v\} \in E} h_u^2)^{1/2}$ which is upper-bounded by $2\sqrt{\max_{u \in U} |\mathcal{N}(u) \cap U|} (\sum_{u \in U} h_u^2)^{1/2}$, where the factor 2 comes by the symmetry between u and v . Expanding the squares and using that $|\lambda_k| \leq \lambda$ for each $k \neq 1$, we get

$$\begin{aligned} d_u h_u^2 &\leq \frac{\lambda^{2d(U,Z)}}{(1-\lambda)^2} \sum_{k=1}^n \psi_{ku}^2 \left(\sum_{z \in Z} \psi_{kz}^2 d_z f_z^2 + \sum_{z \neq z'} \psi_{kz} \psi_{kz'} \sqrt{d_z d_{z'}} f_z f_{z'} \right) \\ &+ \sum_{k \neq k' \geq 2} \frac{\lambda_k^{d(U,Z)} \lambda_{k'}^{d(U,Z)}}{(1-\lambda_k)(1-\lambda_{k'})} \psi_{ku} \psi_{k'u} \sum_{z, z'} \psi_{kz} \psi_{k'z'} \sqrt{d_z d_{z'}} f_z f_{z'}, \end{aligned}$$

where we also used that $\sum_{k=2}^n x_k \leq \sum_{k=1}^n x_k$ if x_1, \dots, x_n are non-negative numbers. Let Ψ denote the matrix having the eigenvectors ψ_1, \dots, ψ_n in its columns, namely, $\Psi_{uk} = (\psi_k)_u = \psi_{ku}$. This is an orthonormal matrix, so both its columns and rows are orthonormal, namely, $\sum_{u=1}^n \psi_{ku} \psi_{k'u} = \mathbf{1}_{k=k'}$ and $\sum_{k=1}^n \psi_{ku} \psi_{kv} = \mathbf{1}_{u=v}$. Using this fact,

$$\sum_{u \in U} h_u^2 \leq \frac{\sum_{u \in V} d_u h_u^2}{\min_{u \in U} d_u} \leq \frac{1}{\min_{u \in U} d_u} \frac{\lambda^{2d(U,Z)}}{(1-\lambda)^2} \sum_{z \in Z} d_z f_z^2,$$

and the proof follows by putting all the pieces together, realizing that the upper bound in the statement of the lemma corresponds to $\frac{1}{\sqrt{2}} (\sum_{u,v \in U: \{u,v\} \in E} g_{uv}^2)^{1/2}$. \square

We are finally ready to present the proof of Theorem 4.

Proof of Theorem 4. Consider the setting developed in Section 4. Fix $\varepsilon \in \mathbb{R}$. From Theorem 3 we have $\frac{dx^*(b(\varepsilon))}{d\varepsilon} = \Sigma(b(\varepsilon)) A^T L(b(\varepsilon))^+ \frac{db(\varepsilon)}{d\varepsilon}$ or, entry-wise, for any $(u, v) \in \vec{E}$,

$$\frac{dx^*(b(\varepsilon))_{(u,v)}}{d\varepsilon} = W(b(\varepsilon))_{uv} (e_u - e_v)^T L(b(\varepsilon))^+ \frac{db(\varepsilon)}{d\varepsilon}.$$

Define $U := V_{\vec{F}}$, and let (U, F) be the undirected graph associated to (U, \vec{F}) (see Remark 1). Clearly, $(\sum_{e \in \vec{F}} (\frac{dx^*(b(\varepsilon))_e}{d\varepsilon})^2)^{1/2}$ is upper-bounded by

$$\gamma(b(\varepsilon)) \sqrt{\frac{1}{2} \sum_{u,v \in U: \{u,v\} \in E} \left((e_u - e_v)^T L(b(\varepsilon))^+ \frac{db(\varepsilon)}{d\varepsilon} \right)^2}, \quad (10)$$

where $\gamma(b) := \max_{u,v \in U} W(b)_{uv}$ for any $b \in \text{Im}(A)$. Corollary 3 yields (choosing $f = \frac{db(\varepsilon)}{d\varepsilon}$, using that $\mathbf{1}^T f = 0$ and $f_v \neq 0$ if and only if $v \in Z$) that $(e_u - e_v)^T L(b(\varepsilon))^+ \frac{db(\varepsilon)}{d\varepsilon}$ equals

$\sum_{z \in Z} \sum_{t=0}^{\infty} (P(b(\varepsilon))_{uz}^t - P(b(\varepsilon))_{vz}^t) \frac{1}{d(b(\varepsilon))_z} \frac{db(\varepsilon)_z}{d\varepsilon}$, so that (10) reads

$$\gamma(b(\varepsilon)) \sqrt{\frac{1}{2} \sum_{u,v \in U: \{u,v\} \in E} \left(\sum_{z \in Z} \sum_{t=0}^{\infty} (P(b(\varepsilon))_{uz}^t - P(b(\varepsilon))_{vz}^t) \frac{1}{d(b(\varepsilon))_z} \frac{db(\varepsilon)_z}{d\varepsilon} \right)^2}.$$

Lemma 5 yields (choosing $f_z = \frac{1}{d_z} \frac{db(\varepsilon)_z}{d\varepsilon}$) that the previous quantity is upper-bounded by

$$\gamma(b(\varepsilon)) \alpha(b(\varepsilon)) \frac{\lambda(b(\varepsilon))^{d(U,Z)}}{1 - \lambda(b(\varepsilon))} \sqrt{\sum_{z \in Z} \left(\frac{db(\varepsilon)_z}{d\varepsilon} \right)^2 \frac{1}{d(b(\varepsilon))_z}},$$

with $\alpha(b) := \frac{\max_{v \in U} \sqrt{2|\mathcal{N}(v) \cap U|}}{\min_{v \in U} \sqrt{d(b)_v}}$ for any $b \in \text{Im}(A)$. Combining everything together, we obtain

$$\left\| \frac{dx^*(b(\varepsilon))}{d\varepsilon} \right\|_{\bar{F}} \leq c(b(\varepsilon)) \frac{\lambda(b(\varepsilon))^{d(U,Z)}}{1 - \lambda(b(\varepsilon))} \left\| \frac{db(\varepsilon)}{d\varepsilon} \right\|_Z,$$

where $c(b) := \frac{\max_{v \in U} \sqrt{2|\mathcal{N}(v) \cap U|}}{\min_{v \in U} \sqrt{d(b)_v}} \gamma(b)$ for $b \in \text{Im}(A)$. The proof follows by taking the supremum over $b \in \text{Im}(A)$. \square

The proof of Lemma 2 follows analogously from the proof of Theorem 4, upon exploiting the symmetry of the pseudoinverse of the graph Laplacian.

E Localized algorithm

This section is devoted to the proof of Theorem 5, which states error bounds for the localized projected gradient descent algorithm. The proof relies on the decay of correlation property established in Theorem 4 for the network flow problem. Recall that the constants appearing in the bounds in Theorem 5 do not depend on the choice of the subgraph \bar{G}' of \bar{G} , but depend only on μ , Q , k_+ , and k_- . To prove this type of bounds, we first need to develop estimates to relate the eigenvalues of weighted subgraphs to the eigenvalues of the corresponding unweighted graph.

E.1 Eigenvalues interlacing

Let $G = (V, E)$ be a simple (i.e., no self-loops, and no multiple edges), connected, undirected graph, with vertex set V and edge set E . Let $B \in \mathbb{R}^{V \times V}$ be the vertex-vertex adjacency matrix of the graph, which is the symmetric matrix defined as $B_{uv} := 1$ if $\{u, v\} \in E$, $B_{uv} := 0$ otherwise. If $n := |V|$, denote by $\mu_n \leq \mu_{n-1} \leq \dots \leq \mu_2 \leq \mu_1$ the eigenvalues of B . Let $G' = (V', E')$ be a connected subgraph of G . Assume that to each edge $\{u, v\} \in E'$ is associated a non-negative weight $W_{uv} = W_{vu} > 0$, and let $W_{uv} = 0$ if $\{u, v\} \notin E'$. Let D' be a diagonal matrix with entries $D'_{vv} = \sum_{w \in V'} W'_{vw}$ for each $v \in V'$. Let $P' := D'^{-1}W'$. If $m := |V'|$, denote by $\lambda'_m \leq \lambda'_{m-1} \leq \dots \leq \lambda'_2 \leq \lambda'_1$ the eigenvalues of P' . The following proposition relates the eigenvalues of P' to the eigenvalues of B . In particular, we provide a bound for the second largest eigenvalue in magnitude of P' with respect to the second largest eigenvalue in magnitude of B , uniformly over the choice of G' .

Proposition 6 (Eigenvalues interlacing). *Let $w_- \leq W_{vw} \leq w_+$ for any $\{v, w\} \in E$, for some constants $w_-, w_+ > 0$. Let k_- and k_+ be, resp., the min and max degree of G . Then,*

$$1 - \frac{w_+ k_+}{w_- k_-} + \frac{w_+}{w_- k_-} \mu_{i+n-m} \leq \lambda'_i \leq 1 - \frac{w_- k_-}{w_+ k_+} + \frac{w_-}{w_+ k_+} \mu_i.$$

Therefore, if $\lambda' := \max\{|\lambda'_2|, |\lambda'_m|\}$ and $\mu := \max\{|\mu_2|, |\mu_n|\}$, we have $\lambda' \leq \frac{w_+ k_+}{w_- k_-} - 1 + \frac{w_+}{w_- k_-} \mu$.

Proof. Consider $\Gamma' := D^{1/2} P' D'^{-1/2} = D'^{-1/2} W' D'^{-1/2}$. As this matrix is similar to P' , it shares the same eigenvalues with P' . Let $L' := D' - W'$ be the Laplacian associated to G' . The Courant-Fischer Theorem yields

$$\lambda'_i = \max_{\substack{S \subseteq \mathbb{R}^m \\ \dim(S)=i}} \min_{x \in S} \frac{x^T \Gamma' x}{x^T x} = 1 + \max_{\substack{S \subseteq \mathbb{R}^m \\ \dim(S)=i}} \min_{y \in S} \frac{-y^T L' y}{y^T D' y},$$

where we used that $x^T \Gamma' x = x^T x - y^T L' y$ with $y := D'^{-1/2} x$, and that the change of variables $y = D'^{-1/2} x$ is non-singular (note that as G' is connected, then D' has non-zero entries on the diagonal). The Laplacian quadratic form yields $y^T L' y = \frac{1}{2} \sum_{u,v \in V'} W_{uv} (y_u - y_v)^2$, which is upper-bounded by $w_+ \frac{1}{2} \sum_{u,v \in V'} B_{uv} (y_u - y_v)^2 = w_+ y^T \mathcal{L}' y$, where \mathcal{L}' is the Laplacian of the *unweighted* graph $G' = (V', E') \equiv (V', E', B')$ with $B' := B_{V', V'}$. Note that we have $\mathcal{L}' = K' - B'$, where K' is diagonal and $K'_{vv} = \sum_{w \in V'} B'_{vw}$ is the degree of vertex $v \in V'$ in G' . As $y^T K' y = \sum_{v \in V'} K'_{vv} y_v^2 \leq k_+ y^T y$, we have $y^T \mathcal{L}' y \leq w_+ k_+ y^T y - w_+ y^T B' y$. At the same time, $y^T D' y \geq w_- k_- y^T y$. Therefore, $\lambda'_i \geq 1 - \frac{w_+ k_+}{w_- k_-} + \frac{w_+}{w_- k_-} \max_{\substack{S \subseteq \mathbb{R}^m \\ \dim(S)=i}} \min_{y \in S} \frac{y^T B' y}{y^T y}$, and by the Courant-Fischer Theorem the right-hand side equals $1 - \frac{w_+ k_+}{w_- k_-} + \frac{w_+}{w_- k_-} \mu'_i$, where $\mu'_n \leq \mu'_{n-1} \leq \dots \leq \mu'_2 \leq \mu'_1$ are the eigenvalues of B' . Analogously, it is easy to prove that $\lambda'_i \leq 1 - \frac{w_- k_-}{w_+ k_+} + \frac{w_-}{w_+ k_+} \mu'_i$. As B' is a principal submatrix of B , the eigenvalue interlacing theorem for symmetric matrices yields $\mu_{i+n-m} \leq \mu'_i \leq \mu_i$, and we have $\alpha + \beta \mu_{i+n-m} \leq \lambda'_i \leq \gamma + \delta \mu_i$, with $\alpha := 1 - \frac{w_+ k_+}{w_- k_-}$, $\beta := \frac{w_+}{w_- k_-}$, $\gamma := 1 - \frac{w_- k_-}{w_+ k_+}$, and $\delta := \frac{w_-}{w_+ k_+}$. Clearly, $|\lambda'_i| \leq -\alpha + \beta \max\{|\mu_{i+n-m}|, |\mu_i|\}$, so $\max\{|\lambda'_2|, |\lambda'_m|\} \leq -\alpha + \beta \max\{|\mu_{2+n-m}|, |\mu_2|, |\mu_n|, |\mu_m|\} = -\alpha + \beta \max\{|\mu_2|, |\mu_n|\}$. \square

E.2 Proof of Theorem 5

We now present the proof of Theorem 5. The proof relies on repeatedly applying Theorem 4 in Section 4 (which captures the decay of correlation for the network flow problem) and the fundamental theorem of calculus.

Proof of Theorem 5. Consider the setting of Section 5.

Bias term. Let us first bound the bias outside \vec{E}' . Let $n := |V|$, and for each $b \in \text{Im}(A)$ let $-1 \leq \lambda_n(b) \leq \lambda_{n-1}(b) \leq \dots \leq \lambda_2(b) < \lambda_1(b) = 1$ be the eigenvalues of $P(b)$. Let $\lambda(b) := \max\{|\lambda_2(b)|, |\lambda_n(b)|\}$ and $\lambda := \sup_{b \in \text{Im}(A)} \lambda(b)$. Define $b(\varepsilon) := b + \varepsilon p$, for any non-negative real number $\varepsilon \geq 0$. If $e \in \vec{E}'^C$, then $T'_{b(\varepsilon)}(x^*(b))_e = x^*(b)_e$ and

$$\text{Bias}(\vec{G}')_e = x^*(b(1))_e - x^*(b(0))_e = \int_0^1 d\varepsilon \frac{dx^*(b(\varepsilon))_e}{d\varepsilon}.$$

By the triangle inequality for the ℓ_2 -norm and Theorem 4,

$$\|\text{Bias}(\vec{G}')\|_{\vec{E}'C} \leq \sup_{\varepsilon \in \mathbb{R}} \left\| \frac{dx^*(b(\varepsilon))}{d\varepsilon} \right\|_{\vec{E}'C} \leq c\|p\| \frac{\lambda^{d(\Delta(\vec{G}'), Z)}}{1-\lambda},$$

where we used that $\sup_{\varepsilon \in \mathbb{R}} \left\| \frac{db(\varepsilon)}{d\varepsilon} \right\|_Z = \|p\|$, as $\frac{db(\varepsilon)_v}{d\varepsilon} = p_v$ for $v \in Z$ and $\frac{db(\varepsilon)_v}{d\varepsilon} = 0$ for $v \notin Z$, and $c := \sqrt{2k_+Q/k_-}$.

Let us now consider the bias inside \vec{E}' . Consider problem (2) in Section 5. For any $\varepsilon > 0, \theta > 0$, define

$$b'(\varepsilon, \theta) := b(\varepsilon)_{V'} - A_{V', \vec{E}'C} x^*(b(\theta))_{\vec{E}'C}. \quad (11)$$

Without loss of generality, we can index the elements of V' and \vec{E}' so that the matrix A has the following structure:

$$A = \begin{pmatrix} A_{V', \vec{E}'} & A_{V', \vec{E}'C} \\ A_{V'C, \vec{E}'} & A_{V'C, \vec{E}'C} \end{pmatrix} = \begin{pmatrix} A' & A_{V', \vec{E}'C} \\ \mathbb{0} & A_{V'C, \vec{E}'C} \end{pmatrix}.$$

For any x that satisfies the flow constraints on $\vec{E}'C$ with respect to $b(\varepsilon)$, $A_{V \setminus V', \vec{E}'C} x_{\vec{E}'C} = b(\varepsilon)_{V \setminus V'}$, we have

$$\left\{ \lim_{t \rightarrow \infty} T_{b+p}^{t'}(x) \right\}_{\vec{E}'} = x'^*(b(1)_{V'} - A_{V', \vec{E}'C} x_{\vec{E}'C}).$$

Clearly $x^*(b)$ satisfies the flow constraints on $\vec{E}'C$ with respect to $b(1)$, as p is supported on V' so that $b(\varepsilon)_{V'C} = b_{V'C}$. Recalling the definition of $b'(\varepsilon, \theta)$ in (11), we then have $(\lim_{t \rightarrow \infty} T_{b+p}^{t'}(x^*(b)))_{\vec{E}'} = x'^*(b'(1, 0))$. On the other hand, as $x^*(b(1))$ is a fixed point of the map $T_{b(1)}^{t'}$, we can characterize the components of $x^*(b(1))$ supported on \vec{E}' as

$$x^*(b(1))_{\vec{E}'} = \left\{ \lim_{t \rightarrow \infty} T_{b(1)}^{t'}(x^*(b(1))) \right\}_{\vec{E}'} = x'^*(b'(1, 1)).$$

It is easy to check that $b'(\varepsilon, \theta) \in \text{Im}(A')$ for each value of ε and θ . In fact, as \vec{G}' is connected by assumption, then $\text{Im}(A')$ corresponds to the subspace of $\mathbb{R}^{V'}$ orthogonal to the all-ones vector $\mathbb{1}$. We have $\mathbb{1}^T b'(\varepsilon, \theta) = \mathbb{1}^T b_{V'} + \varepsilon \mathbb{1}^T p_{V'} - \mathbb{1}^T A_{V', \vec{E}'C} x^*(b(\theta))_{\vec{E}'C}$. Note that $\mathbb{1}^T p_{V'} = 0$ by assumption. Also, $0 = \mathbb{1}^T b = \mathbb{1}^T b_{V'} + \mathbb{1}^T b_{V'C}$ so that $\mathbb{1}^T b_{V'} = -\mathbb{1}^T b_{V'C}$. Analogously, as $\mathbb{1}^T A = \mathbb{0}^T$, we have $\mathbb{1}^T A_{V', \vec{E}'C} = -\mathbb{1}^T A_{V'C, \vec{E}'C}$. Hence,

$$\mathbb{1}^T b'(\varepsilon, \theta) = -\mathbb{1}^T b_{V'C} + \mathbb{1}^T A_{V'C, \vec{E}'C} x^*(b(\theta))_{\vec{E}'C} = \mathbb{0}^T,$$

where the last equality follows as clearly $A_{V'C, \vec{E}'C} x^*(b(\theta))_{\vec{E}'C} = b_{V'C}$. Therefore, for $e \in \vec{E}'$,

$$\text{Bias}(\vec{G}')_e = \int_0^1 d\theta \frac{dx'^*(b'(1, \theta))_e}{d\theta}.$$

For each $b' \in \text{Im}(A')$, let $W'(b') \in \mathbb{R}^{V' \times V'}$ be a symmetric matrix defined as $W'(b')_{uv} = (\frac{\partial^2 f_e(x'^*(b')_e)}{\partial x_e^2})^{-1}$ if either $e = (u, w) \in \vec{E}$ or $e = (w, u) \in \vec{E}$, and $W'(b')_{uv} := 0$ otherwise. Let $D'(b') \in \mathbb{R}^{V' \times V'}$ be a diagonal matrix with entries $D'(b')_{vv} = \sum_{w \in V'} W'(b')_{vw}$.

Let $P'(b') := D'(b')^{-1}W'(b')$. If $m := |V|$, let $-1 \leq \lambda'_m(b') \leq \lambda'_{m-1}(b') \leq \dots \leq \lambda'_2(b') < \lambda'_1(b') = 1$ be the eigenvalues of $P'(b')$ (where this characterization holds as G' is connected by assumption). Define $\lambda'(b') := \max\{|\lambda'_2(b')|, |\lambda'_m(b')|\}$ and $\lambda' := \sup_{b' \in \text{Im}(A')} \lambda'(b')$. Proceeding as above, applying Theorem 4 to the optimization problem defined on \vec{G}' we get

$$\|\text{Bias}(\vec{G}')\|_{\vec{E}'} \leq \sup_{\theta \in \mathbb{R}} \left\| \frac{dx'^*(b'(1, \theta))}{d\theta} \right\|_{\vec{E}'},$$

which is upper-bounded by $c \frac{1}{1-\lambda'} \sup_{\theta \in \mathbb{R}} \left\| \frac{\partial b'(\varepsilon, \theta)_v}{\partial \theta} \right\|_{\Delta(\vec{G}'')}$, where we used that $\frac{\partial b'(\varepsilon, \theta)_v}{\partial \theta} = 0$ if $v \in V' \setminus \Delta(\vec{G}'')$, and clearly $d(V', \Delta(\vec{G}'')) = 0$ as $\Delta(\vec{G}'') \subseteq V'$. For $v \in \Delta(\vec{G}'')$ we have $\frac{\partial b'(\varepsilon, \theta)_v}{\partial \theta} = -\sum_{e \in \vec{E}'^C} A_{ve} \frac{dx^*(b(\theta))_e}{d\theta}$. If $\vec{F}'(v) := \{e \in \vec{E}'^C : e = (u, v) \text{ or } e = (v, u) \text{ for some } u \in V'^C\}$ denotes the set of edges that are connected to v but do not belong to \vec{E}' , we have $(\frac{\partial b'(\varepsilon, \theta)_v}{\partial \theta})^2 \leq (\sum_{e \in \vec{F}'(v)} |\frac{dx^*(b(\theta))_e}{d\theta}|)^2$, which by Jensen's inequality admits $|\vec{F}'(v)| \sum_{e \in \vec{F}'(v)} (\frac{dx^*(b(\theta))_e}{d\theta})^2$ as an upper bound. As $\max_{v \in \Delta(\vec{G}'')} |\vec{F}'(v)| \leq k_+ - 1$, applying Theorem 4 as above we get

$$\left\| \frac{\partial b'(1, \theta)}{\partial \theta} \right\|_{\Delta(\vec{G}'')} \leq \sqrt{k_+ - 1} \left\| \frac{dx^*(b(\theta))}{d\theta} \right\|_{\vec{E}'^C} \leq c \sqrt{k_+ - 1} \|p\| \frac{\lambda^{d(\Delta(\vec{G}''), Z)}}{1 - \lambda'}.$$

Therefore, $\|\text{Bias}(\vec{G}')\|_{\vec{E}'} \leq c^2 \sqrt{k_+ - 1} \|p\| \frac{\lambda^{d(\Delta(\vec{G}''), Z)}}{(1-\lambda')(1-\lambda)}$. By the triangle inequality for the ℓ_2 -norm, $\|\text{Bias}(\vec{G}')\| \leq \|\text{Bias}(\vec{G}')\|_{\vec{E}'} + \|\text{Bias}(\vec{G}')\|_{\vec{E}'^C}$, so we obtain

$$\|\text{Bias}(\vec{G}')\| \leq c(1 + c\sqrt{k_+ - 1}) \|p\| \frac{\lambda^{d(\Delta(\vec{G}''), Z)}}{(1-\lambda')(1-\lambda)}.$$

By Proposition 6 we have $\max\{\lambda, \lambda'\} \leq \frac{Qk_+}{k_-} - 1 + \frac{Q}{k_-} \mu$, and the bound for the bias term follows.

Variance term. As $(\lim_{t \rightarrow \infty} T'_{b+p}{}^t(x^*(b)))_{\vec{E}'} = x'^*(b'(1, 0))$, we get

$$\begin{aligned} \|\text{Var}(\vec{G}', t)\|_{\vec{E}'} &= \|x'^*(b'(1, 0)) - T'_{b+p}{}^t(x^*(b))_{\vec{E}'}\| \\ &\leq e^{-\frac{t}{2Q}} \|x'^*(b'(1, 0)) - x'^*(b'(0, 0))\|, \end{aligned}$$

where in the last inequality we used that $x^*(b)_{\vec{E}'} = x'^*(b'(0, 0))$. For each $e \in \vec{E}'$ we have

$$x'^*(b'(1, 0))_e - x'^*(b'(0, 0))_e = \int_0^1 d\varepsilon \frac{dx'^*(b'(\varepsilon, 0))_e}{d\varepsilon},$$

and using the triangle inequality for the ℓ_2 -norm, applying Theorem 4 to the optimization problem defined on \vec{G}' ,

$$e^{\frac{t}{2Q}} \|\text{Var}(\vec{G}', t)\|_{\vec{E}'} \leq \sup_{\varepsilon \in \mathbb{R}} \left\| \frac{dx'^*(b'(\varepsilon, 0))}{d\varepsilon} \right\|_{\vec{E}'} \leq c \|p\| \frac{1}{1 - \lambda'},$$

where we used that $\frac{\partial b'(\varepsilon, 0)_v}{\partial \varepsilon} = \frac{db(\varepsilon)_v}{d\varepsilon} = p_v$ for $v \in Z$ and $\frac{db(\varepsilon)_v}{d\varepsilon} = 0$ for $v \notin Z$, and that $d(V', Z) = 0$ as $Z \subseteq V'$. Clearly, $\text{Var}(\vec{G}', t)_e = 0$ for $e \in \vec{E}'^C$, as $T'_b{}^t(x^*(b))_e = x^*(b)_e$. Hence, $\|\text{Var}(\vec{G}', t)\| = \|\text{Var}(\vec{G}', t)\|_{\vec{E}'}$, and the proof is concluded as $\lambda' \leq \frac{Qk_+}{k_-} - 1 + \frac{Q}{k_-} \mu$ by Proposition 6. \square