

A new approach to Laplacian solvers and flow problems

Patrick Rebeschini

PATRICK.REBESCHINI@YALE.EDU

*Department of Electrical Engineering
Yale University, New Haven,
New Haven, CT 06511, USA*

Sekhar Tatikonda

SEKHAR.TATIKONDA@YALE.EDU

*Department of Electrical Engineering
Yale University, New Haven,
New Haven, CT 06511, USA*

Abstract

This paper investigates message-passing algorithms for solving systems of linear equations in the Laplacian matrices of graphs and to compute electric flows. These two problems are fundamental primitives that arise in several domains such as computer science, electrical engineering, operations research, and machine learning. Despite the extensive literature on approximately solving these problems in quasi-linear time, the algorithms that have been proposed are typically centralized and involve multiple graph theoretic constructions or sampling mechanisms that make them difficult to implement and analyze. On the other hand, message-passing routines are distributed, simple, and easy to implement. In this paper we establish a framework to analyze message-passing algorithms to solve voltage and flow problems. We characterize the error committed by the algorithms in d -regular graphs with equal weights. We show that the convergence of the algorithms is controlled by the total variation distance between the distributions of non-backtracking random walks that start from neighbor nodes. More broadly, our analysis of message-passing introduces new insights to address generic optimization problems with constraints.

Keywords: Laplacian solver, flow problem, message-passing, min-sum, distributed algorithms.

1. Introduction

Voltage and flow problems. Laplacian matrices of graphs are fundamental. Solving linear systems of equations in these matrices constitutes the back-bone of many algorithms with applications to a wide variety of fields, such as computer science, machine learning, electrical engineering, operations research, computer vision, network analysis, and computational biology. Laplacian solvers capable of providing a ε -approximation in the Laplacian-modified ℓ_2 norm with a quasi-linear running time $\tilde{O}(m \log^c n \log 1/\varepsilon)$ — where n and m are the number of vertices and edges in the graph, respectively, c is a constant greater than 0, and where the \tilde{O} notation hides $\log \log n$ factors — have recently been used to break longstanding barriers in a wide list of fundamental graphs problems. See Spielman (2011); Vishnoi (2013) for surveys on these solvers and their applications.

When a weighted graph is viewed as an electric circuit where each weighted edge represents a resistor and where each vertex represents a node to which some external current is

applied, solving linear systems of equation in the Laplacian matrix of the graph corresponds to computing the voltages induced to each node by the external currents. For this reason, we refer to the Laplacian problem as the “voltage problem.” Alongside with the problem of computing the voltages is the problem of computing the current induced across each resistor, which we refer to as the “flow problem.” The two problems are related by Ohm’s law — which says that electric flows are induced by voltage differences across the nodes — and they both can be casted in the form of quadratic optimization problems. The voltage problem is an unconstrained problem to find the voltage assignment that maximizes the quadratic Laplacian form. The flow problem is a constrained optimization problem to find the flow assignment that minimizes the circuit energy while satisfying Kirchhoff’s conservation law. The voltage problem is the dual of the flow problem. Recently, the voltage problem has received more attention than the flow problem, despite the fact that voltages are defined up to differences — Laplacian solvers aims at computing the *unique* set of voltages that sums to zero — while flows are uniquely defined.

Previous literature on Laplacian solvers. The extensive literature on approximately solving the voltage and flow problems in quasi-linear time typically involves algorithms that rely on multiple graph theoretic constructions or sampling mechanisms which make them difficult to implement and analyze. The first quasi-linear time solver was given in Spielman and Teng (2004, 2014), and it achieves a running time of $\tilde{O}(m \log^c n \log 1/\varepsilon)$, where c is a constant close to 100. Since then, many algorithms have steadily decreased the exponent c , with the work in Koutis et al. (2011) that has reduced the constant c to 1. Currently, the best running time is in Cohen et al. (2014) where $c = 1/2$. These algorithms all rely on the same general architecture, which combines the Chebyshev method or the conjugate gradient algorithm with the use of recursive preconditioning, spectral sparsifiers, low-stretch spanning trees, and expander graphs. Despite their theoretical guarantees, these algorithms present noticeable difficulties to practitioners due to the involved machinery they rely upon. Lately, new algorithms have been proposed that depart from this framework and use less graph theoretic constructions. The method in Kelner et al. (2013), for instance, which operates in the flow domain and yields a fast solver with $c = 2$, using very little of the previous machinery — although it still requires low-stretch spanning trees. More recently, the work in Kyng and Sachdeva (2016), which relies on sparse Gaussian elimination for Laplacian matrices to yield a fast solver with $c = 3$, without the use of any graph theoretic construction. While being more appealing, these new types of algorithms still rely on random sampling so the convergence analysis is only performed in expectation, or with high probability, and their pseudo-code is still not as neat as one could hope for. Moreover, all these algorithms are centralized. There has also been some work in parallelizing these solvers, with the first parallel solver with near linear work and polylogarithmic depth presented in Peng and Spielman (2014). However, these parallel solvers work in the shared memory model, and are not completely distributed.

Message-passing algorithms. The aforementioned limitations prompted us to investigate the behavior of a popular class of distributed algorithms called message-passing to solve both the voltage and the flow problem. Message-passing algorithms are iterative, distributed, deterministic (i.e., not randomized), general-purpose routines for a wide class of optimization and statistical inference problems that have a graph structure. In optimization, message-passing algorithms are used to optimize objective functions that are sums

of component functions supported on a given graph. These algorithms have emerged as canonical procedures to address large scale problems in a variety of domains, including signal processing, statistics, and machine learning. Their success comes from the fact that the algorithms are simple and easy to implement, require minimal data structure and low communication per iteration, and in many applications they have been shown to converge to the solution of the problem. Interest in message-passing have been triggered by their success in solving certain classes of challenging combinatorial optimization problems, such as decoding low-density parity-check codes, see Berrou et al. (1993); Richardson and Urbanke (2001) for instance, and solving certain type of satisfiability problems, see Mezard et al. (2002); Braunstein et al. (2005), for instance.

Despite their successes, the convergence analysis of message-passing algorithms remains limited. The canonical message-passing algorithm in optimization takes the name of “min-sum,” as it consecutively involves minimizations and summations of real-valued functions, also called “messages.” For quadratic optimization problems and, more generally, convex optimization problems, the min-sum algorithm has been shown to converge to the problem solution for any graph topology under the assumption that the objective function is scaled diagonal dominant (Moallemi and Van Roy, 2010). This condition is restrictive. In particular, it does not apply to quadratic objective functions in Laplacians matrices, hence to the voltage problem, and it does not apply to constrained optimization problems, hence to the flow problem. While some work has been done to study the behavior of min-sum in constrained optimization, this only seems to apply to linear objective functions. See Gamarnik et al. (2012) and references therein.

Results. To the best of our knowledge, this paper presents the first results on the convergence behavior of the min-sum algorithm applied to the voltage and flow problems. More broadly, we establish a general framework to analyze the convergence of the min-sum paradigm, which goes beyond the typical assumption of scaled-diagonal dominance. The algorithms that we analyze — Algorithm 3 for the voltage problem and Algorithm 5 for the flow problem — are easy to implement, with a pseudo-code that fits two lines. We study the behavior of these algorithms in d -regular graphs. In particular, for graphs with equal weights, when $d \geq 3$, we show that the error the algorithms commit is *characterized* by quantities involving the difference of the distributions of non-backtracking random walks that start from neighbor nodes. We show that the convergence of the algorithms (we perform the analysis with respect to various norms) can be controlled by the total variation distance between these distributions. For some classes of graphs, such as $d/2$ -connected cycles and $d/2$ -dimensional tori, where d is even, we present empirical evidence that the total variation distances of interest decay like $O(1/\sqrt{t})$, where t is the iteration time, independently of the graph dimensions n and m . This empirical evidence allows us to show that in these graphs, potentially under some additional assumptions on the problem inputs (depending on the norm used for the analysis), the min-sum algorithm can yield ε -approximate solutions for both the voltage and the flow problem with running time $O(m/\varepsilon^2)$, which is *linear* in m .

Novelty in the analysis of message-passing. Our analysis of the min-sum algorithm presents two keys contributions that go beyond the specifics of the problems here addressed.

1. **Beyond contraction arguments.** Typically, the convergence analysis of message-passing algorithms is based on exploiting a decay of correlation property in the computation tree that is obtained by unraveling the algorithm operations with time. This

decay of correlation property is established by imposing strong, *local* assumptions on the original problem, such as the assumption of scaled diagonal dominance for the min-sum algorithm. These conditions essentially yield a bound on the sum of the local interactions in the computation tree by a quantity strictly less than one. In this case, a contraction argument can be performed at each level of the computation tree, and an exponential decay from root to leaves can be established, which eventually results in the exponential convergence of the algorithm. On the other hand, the analysis of message-passing algorithms that we give in this paper is based on a *global* analysis of the problem supported on the computation tree, which can not be captured by local contraction arguments based on worst-case bounds for the local interactions. We show that what drives the convergence of the min-sum algorithm in our applications is not a decay of correlation *in magnitude* from root to leaves, but a decay of correlation that keeps track of the positive and negative signs. Ultimately, the convergence of the algorithms that we analyze is driven by the cancellations in the summation over the leaves of the computation tree, as it is particularly evident in the flow problem.

2. **Beyond unconstrained optimization problems.** Typically, the convergence analysis of message-passing algorithms relies on finding the fix point(s) of the algorithms. For unconstrained optimization problems, the key insight to find fix point(s) is that the first order optimality conditions for the original problem are respected by the problem supported on the computation tree. See Moallemi and Van Roy (2010), for instance. On the other hand, the convergence analysis that we present for the flow problem — which is a constrained problem — is based on the realization that also the KKT optimality conditions for constrained problems are respected on the computation tree. This insight allows to extend the analysis of the min-sum algorithm to more general constrained optimization problems.

Structure of the paper. The structure of the paper is as follows. In Section 2 we define the voltage and flow problems as optimization procedures. In Section 3 we introduce the general-purpose min-sum algorithm, which we then specialize to the voltage and flow problem, respectively. In Section 4 we present the error characterization for the min-sum algorithms for d -regular graphs, and we discuss their convergence behaviors. In Section 5 we develop the main framework to investigate the min-sum algorithm, and we present the proofs of the results given in the previous section. In Appendix A we establish a general connection between the inverse of restricted Laplacian matrices and hitting times of (regular, i.e., walks that *can* backtrack) random walks on graphs — which is the main computational tool used in our analysis. In Appendix B we append some technical results.

Remark 1 (Notation) *Throughout, for a given real-valued matrix M , we denote by M^T its transpose, by M^{-1} its inverse, and by M^+ its Moore-Penrose pseudoinverse. Given a set \mathcal{I} and $i \in \mathcal{I}$, we denote by i both the element $i \in \mathcal{I}$ and the subset $\{i\} \subseteq \mathcal{I}$. For instance, we write $\mathcal{I} \setminus i$ to mean $\mathcal{I} \setminus \{i\}$. Given a function $f : \mathbb{R}^{\mathcal{I}} \rightarrow \mathbb{R}$, $i \in \mathcal{I}$, we write $f(x_{\mathcal{I} \setminus i} \cdot)$ to mean the function $z \in \mathbb{R} \rightarrow f(x_{\mathcal{I} \setminus i} z)$.*

2. Voltage and flow problems

Throughout, let us consider a simple (i.e., no self-loops, and no multiple edges), undirected, weighted graph G given by the triple (V, E, W) , where V is a set of n vertices, E is a set of m edges, and $W \in \mathbb{R}^{V \times V}$ is a symmetric matrix that assigns a positive weight to every edge, namely, $W_{vw} > 0$ if $\{v, w\} \in E$, and $W_{vw} = 0$ if $\{v, w\} \notin E$. The matrix W is called the *adjacency matrix* of the weighed graph G . Define the weighted degree of a vertex v by $d_v := \sum_{w \in V} W_{vw}$, and let $D \in \mathbb{R}^{V \times V}$ be the diagonal matrix defined by $D_{vv} := d_v$. The Laplacian L of the graph G is the matrix defined as $L := D - W$. Without loss of generality, in what follows we assume that the graph G is connected, otherwise we can treat each connected component on its own. It is easy to verify that in this case the null space of the Laplacian matrix L is spanned by the all-ones vector $\mathbf{1} \in \mathbb{R}^V$, so that the range of L is spanned by the vectors orthogonal to $\mathbf{1}$.

Voltage problem. We are interested in solving linear equations in the Laplacian matrix L . That is, given b in the range of L , i.e., $\mathbf{1}^T b = 0$, our goal is to find ν that satisfies

$$L\nu = b. \quad (1)$$

As it can immediately be checked from first order optimality conditions, the set of solutions to (1) coincides with the set of solutions to the “voltage problem,” namely:

$$\text{maximize} \quad -\frac{1}{2}\nu^T L\nu + b^T \nu. \quad (2)$$

We are interested in the unique solution that is orthogonal to the all one vector, namely, $\nu^* = L^+ b$, where L^+ is the Moore-Penrose pseudoinverse of L . Given an accuracy parameter $\varepsilon > 0$, the quasi-linear algorithms mentioned in the introduction achieve a quasi-linear running time $\tilde{O}(m \log^c n \log 1/\varepsilon)$ to return an estimate $\hat{\nu}$ that satisfies $\|\nu^* - \hat{\nu}\|_L \leq \varepsilon \|\nu^*\|_L$, where $\|\nu\|_L := \sqrt{\nu^T L \nu}$ is the Laplacian-modified ℓ_2 norm, also known as L -norm. It can be easily seen that $\|\cdot\|_L$ is a pseudo-norm, as $\|\nu\|_L = 0$ for any vector ν orthogonal to the all one vector. In many applications where Laplacian solvers are needed, the L -norm is the natural norm to run the analysis. See Vishnoi (2013).

Flow problem. To introduce the flow problem, let us fix an arbitrary orientation of the edges in E . Denote this set of directed edges by \vec{E} , and let $\vec{G} = (V, \vec{E})$ be the corresponding directed graph. Let $A \in \mathbb{R}^{V \times \vec{E}}$ be the *signed vertex-edge incidence* matrix of the graph \vec{G} , defined as $A_{ve} := 1$ if edge e leaves node v , $A_{ve} := -1$ if edge e enters node v , $A_{ve} := 0$ otherwise. To each directed edge $e = (v, w)$ let $R \in \mathbb{R}^{\vec{E} \times \vec{E}}$ be the diagonal matrix defined by $R_{ee} := 1/W_{vw}$. It is easy to verify that the Laplacian L of the undirected graph G can be expressed as $L = AR^{-1}A^T$. Note that while A depends on the choice of the direction of the edges in \vec{E} , L does not. The “flow problem” is defined as follows:

$$\begin{aligned} &\text{minimize} && \frac{1}{2}x^T R x \\ &\text{subject to} && Ax = b. \end{aligned} \quad (3)$$

This problem can be interpreted as a minimal-energy electrical network problem. To each directed edge $e \in \vec{E}$ is associated a resistor of value R_{ee} (analogously, a conductor of value $1/R_{ee}$) through which a current x_e flows, with the convention that $x_e > 0$ if the current

is in the direction of the edge, $x_e < 0$ if the current is in the direction opposite the edge. The energy of having current x_e on edge e is given by the $\frac{1}{2}R_{ee}x_e^2$. For each node $v \in V$, b_v represents a given external current: $b_v > 0$ represents a source where the current enters node v , whereas $b_v < 0$ represents a sink where the current exits node v . The constraint equations $Ax = b$ embody Kirchoff's conservation law, by which at each vertex the net sum of the incoming and outgoing internal flows equals the external flow.

Connection. From the KKT optimality conditions, it is easy to verify (see Proposition 23 in Appendix B) that the unique solution of problem (3) reads $x^* = R^{-1}A^TL + b = R^{-1}A^T\nu^*$. That is, if $e = (v, w) \in \vec{E}$, we have $x_e^* = W_{vw}(\nu_v^* - \nu_w^*)$ which represents Ohm's law. Finally, notice that problem (2) is the dual of (3). Recall that the Lagrangian of (3) is the function \mathcal{L} from $\mathbb{R}^{\vec{E}} \times \mathbb{R}^V$ to \mathbb{R} defined as

$$\mathcal{L}(x, \nu) := \frac{1}{2}x^TRx + \nu^T(b - Ax). \quad (4)$$

The dual function is $q(\nu) = \min_{x \in \mathbb{R}^{\vec{E}}} \mathcal{L}(x, \nu) = -\frac{1}{2}\nu^TL\nu + b^T\nu$, and the dual problem reads $\max_{\nu \in \mathbb{R}^V} q(\nu)$, which is problem (2).

3. Min-sum algorithm

The min-sum algorithm is a popular distributed routine to optimize a cost function that has a graph structure. Let \mathcal{V} and \mathcal{F} be two finite sets, whose elements are respectively referred to as *variables* and *factors*, to be interpreted as two groups of vertices in a certain undirected bipartite graph with edge set \mathcal{E} , where each edge in \mathcal{E} connects an element of \mathcal{V} with an element of \mathcal{F} . For any variable node $i \in \mathcal{V}$, we use the notation ∂i to denote the set of neighbors of vertex i in the factor graph. Analogously, for any factor node $a \in \mathcal{F}$, we use the notation ∂a to denote the set of neighbors of vertex a in the factor graph. Consider the following optimization problem over $x \in \mathbb{R}^{\mathcal{V}}$:

$$\text{minimize} \quad g(x) := \sum_{i \in \mathcal{V}} g_i(x_i) + \sum_{a \in \mathcal{F}} g_a(x_{\partial a}), \quad (5)$$

for a given set of functions $g_i : \mathbb{R} \rightarrow \mathbb{R}$, for each $i \in \mathcal{V}$, and $g_a : \mathbb{R}^{\partial a} \rightarrow \mathbb{R}$, for each $a \in \mathcal{F}$. The min-sum algorithm is an iterative algorithm that proceeds by updating a set of *messages* associated to each edge in the underlying factor graph. Messages are real functions on the optimization variables x_i , for $i \in \mathcal{V}$, and each edge in the factor graph is associated with two messages at any given iteration of the algorithm. For any time $s \geq 0$, and $i \in \mathcal{V}, a \in \mathcal{F}$ so that $\{i, a\} \in \mathcal{E}$, $\mu_{i \rightarrow a}^s : \mathbb{R} \rightarrow \mathbb{R}$ represents a message from variable to factor, and $\mu_{a \rightarrow i}^s : \mathbb{R} \rightarrow \mathbb{R}$ represents a message from factor to variable.

In this work we consider the *synchronous* implementation of the min-sum algorithm, where at each iteration all messages are updated. For a given iteration step $t \geq 1$, the algorithm reads as in Algorithm 1. The iterations of the min-sum algorithm correspond to the dynamic programming iterations to solve (5) when the underlying graph is a tree. In general, the min-sum algorithm is not guaranteed to converge, and even when it converges it is not guaranteed that it converges to an optimal point of problem (5). The min-sum algorithm is suited for cooperative distributed multi-agents optimization. In this framework, to each variable $i \in \mathcal{V}$ is associated an agent who has access only to the local objective

Algorithm 1: Min-sum

Input: Initial messages $\{\mu_{a \rightarrow i}^0\}$, for each $\{i, a\} \in \mathcal{E}$;
for $s \in \{1, \dots, t\}$ **do**
 Compute, for each $\{i, a\} \in \mathcal{E}$,
 $\mu_{i \rightarrow a}^s = g_i + \sum_{b \in \partial i \setminus a} \mu_{b \rightarrow i}^{s-1}$, $\mu_{a \rightarrow i}^s = \min_{x \in \mathbb{R}^V} \{g_a(x_{\partial a \setminus i \cdot}) + \sum_{j \in \partial a \setminus i} \mu_{j \rightarrow a}^s(x_j)\}$;
Compute $\mu_i^t = g_i + \sum_{a \in \partial i} \mu_{a \rightarrow i}^t$, for each $i \in \mathcal{V}$;
Output: $\hat{x}_i^t := \arg \min_{z \in \mathbb{R}} \mu_i^t(z)$, for each $i \in \mathcal{V}$.

functions g_i and g_a , for each $a \in \partial i$. At each time step, each agent can only exchange information with his near-neighbors in the factor graph, i.e., the agents who share at least one factor with him. Agents' goal is to cooperate over time to solve problem (5). The min-sum algorithm is a mean to achieve this goal: iteratively, each agent i updates a set of messages $\{\mu_{i \rightarrow a}^s\}$, $a \in \partial i$, using the messages that he has received from his neighbors at the previous time step, i.e., $\{\mu_{j \rightarrow a}^{s-1}\}$, $a \in \partial j$, for each neighbor j of i , and using the locally known functions g_i and g_a , for each $a \in \partial i$.

We now specialize the min-sum algorithm to the voltage and flow problem, respectively.

3.1 Min-sum for the voltage problem

In order to apply the min-sum algorithm to solve the voltage problem (2), throughout the rest of this paper we consider the equivalent minimization form: minimize $\frac{1}{2} \nu^T L \nu - b^T \nu$. We rewrite this problem in the unconstrained form (5).¹ To this end, let us identify $\mathcal{V} \leftrightarrow V$ and $\mathcal{F} \leftrightarrow E$. In this context, the notation ∂v indicates the set of edges in E are connected to $v \in V$; analogously, the notation ∂e indicates the pair of vertices that are connected by edge e . For each $v \in V$, define the function $g_v(\nu_v) = -b_v \nu_v$, and for each $e \in E$ define the function $g_e(\nu_{\partial e}) := \frac{1}{2} W_{vw} (\nu_v - \nu_w)^2$, where $\partial e = \{v, w\}$. Recalling that $\nu^T L \nu = \sum_{\{v, w\} \in E} W_{vw} (\nu_v - \nu_w)^2$, the problem above is equivalent to the problem of minimizing the function $\sum_{v \in V} g_v(\nu_v) + \sum_{e \in E} g_e(\nu_{\partial e})$, which is of the form (5). As in the factor graph interpretation each factor node (i.e., edge in E) is adjacent to exactly two variable nodes (i.e., vertices in V), it is convenient to rewrite the min-sum algorithm in terms of updates of messages from factors to variables only. At each time step $t \geq 1$, to each edge $e = \{v, w\} \in E$ are associated two messages: $\mu_{e \rightarrow v}^t$ and $\mu_{e \rightarrow w}^t$. It is easy to verify that in the present setting Algorithm 1 reads as follows.

Algorithm 2 applies for any choice of the initial messages $\{\mu_{e \rightarrow v}^0\}$, for $e \in E, v \in \partial e$. If we choose initial messages that are quadratic functions, then also the sequent messages computed by the algorithm remain quadratic. The following Lemma presents the quadratic updates, not stating formulas for the constant terms as these terms do not influence the output of the min-sum algorithm, as we will show next.

Lemma 2 *For a given $s \geq 1$, for each $e \in E, v \in \partial e$ let $\mu_{e \rightarrow v}^{s-1}(z) = \frac{1}{2} W_{e \rightarrow v}^{s-1} z^2 + w_{e \rightarrow v}^{s-1} z + c_{e \rightarrow v}^{s-1}$. Then, if the graph G has no leaves, for each $e \in E, v \in \partial e$, we have $\mu_{e \rightarrow v}^s(z) =$*

1. It is well known that different ways of mapping the original optimization problem to the form (5) yield different convergence results. See Moallemi and Roy (2009), for instance.

Algorithm 2: Min-sum, voltage problem

Input: Initial messages $\{\mu_{e \rightarrow v}^0\}$, for $e \in E, v \in \partial e$;
for $s \in \{1, \dots, t\}$ **do**
 Compute, for each $e \in E, v \in \partial e$, with $w = \partial e \setminus v$
 $\mu_{e \rightarrow v}^s = \min_{\nu_w \in \mathbb{R}} \{g_w(\nu_w) + g_e(\nu_w \cdot) + \sum_{f \in \partial w \setminus e} \mu_{f \rightarrow w}^{s-1}(\nu_w)\}$;
 Compute $\mu_v^t = g_v + \sum_{e \in \partial v} \mu_{e \rightarrow v}^t$, for each $v \in V$;
Output: $\hat{\nu}_v^t := \arg \min_{z \in \mathbb{R}} \mu_v^t(z)$, for each $v \in V$.

$\frac{1}{2}W_{e \rightarrow v}^s z^2 + w_{e \rightarrow v}^s z + c_{e \rightarrow v}^s$, where $\{c_{e \rightarrow v}^s\}$ is a certain set of constants, and (here $w = \partial e \setminus v$)

$$W_{e \rightarrow v}^s = W_{vw} \frac{\sum_{f \in \partial w \setminus e} W_{f \rightarrow w}^{s-1}}{W_{vw} + \sum_{f \in \partial w \setminus e} W_{f \rightarrow w}^{s-1}}, \quad w_{e \rightarrow v}^s = W_{vw} \frac{\sum_{f \in \partial w \setminus e} w_{f \rightarrow w}^{s-1} - b_w}{W_{vw} + \sum_{f \in \partial w \setminus e} W_{f \rightarrow w}^{s-1}}.$$

Proof Fix $s \geq 1$, $e \in E$, $v \in \partial e$, and let $w = \partial e \setminus v$. The min-sum update rule reads $\mu_{e \rightarrow v}^s(\nu_v) = \min_{z \in \mathbb{R}} \{-b_w z + \frac{1}{2}W_{vw}(\nu_v - z)^2 + \frac{1}{2}z^2 \sum_{f \in \partial w \setminus e} W_{f \rightarrow w}^{s-1} + z \sum_{f \in \partial w \setminus e} w_{f \rightarrow w}^{s-1} + \sum_{f \in \partial w \setminus e} c_{f \rightarrow w}^{s-1}\}$. The solution of the minimization reads $z^* = \frac{W_{vw}\nu_v + b_w - \sum_{f \in \partial w \setminus e} w_{f \rightarrow w}^{s-1}}{W_{vw} + \sum_{f \in \partial w \setminus e} W_{f \rightarrow w}^{s-1}}$, and the proof is immediately concluded by substitution. \blacksquare

From Lemma 2 it follows that if we initialize Algorithm 2 with quadratic messages of the form $\mu_{e \rightarrow v}^0(z) = \frac{1}{2}W_{e \rightarrow v}^0 z^2 + w_{e \rightarrow v}^0 z$, then, modulo the constant terms, we can track the evolution of the quadratic messages computed by the algorithm by evaluating recursively the set of parameters $\{W_{e \rightarrow v}^s\}, \{w_{e \rightarrow v}^s\}$ as prescribed by Lemma 2. The belief function associated to vertex v at time t then reads $\mu_v^t(z) = -b_v z + \sum_{e \in \partial v} (\frac{1}{2}W_{e \rightarrow v}^t z^2 + w_{e \rightarrow v}^t z + c_{e \rightarrow v}^t)$, so that $\hat{\nu}_v^t = \arg \min_{z \in \mathbb{R}} \mu_v^t(z) = \frac{b_v - \sum_{e \in \partial v} w_{e \rightarrow v}^t}{\sum_{e \in \partial v} W_{e \rightarrow v}^t}$, independent of $\{c_{e \rightarrow v}^t\}$. The final procedure so obtained, in the case of quadratic initial messages is given by Algorithm 3.

Algorithm 3: Min-sum, voltage problem, quadratic messages

Input: Initial messages $\{W_{e \rightarrow v}^0\}, \{w_{e \rightarrow v}^0\}$, $e \in E, v \in \partial e$;
for $s \in \{1, \dots, t\}$ **do**
 Compute, for each $e \in E, v \in \partial e$, with $w = \partial e \setminus v$
 $W_{e \rightarrow v}^s = W_{vw} \frac{\sum_{f \in \partial w \setminus e} W_{f \rightarrow w}^{s-1}}{W_{vw} + \sum_{f \in \partial w \setminus e} W_{f \rightarrow w}^{s-1}}, \quad w_{e \rightarrow v}^s = W_{vw} \frac{\sum_{f \in \partial w \setminus e} w_{f \rightarrow w}^{s-1} - b_w}{W_{vw} + \sum_{f \in \partial w \setminus e} W_{f \rightarrow w}^{s-1}};$
Output: $\hat{\nu}_v^t = \frac{b_v - \sum_{e \in \partial v} w_{e \rightarrow v}^t}{\sum_{e \in \partial v} W_{e \rightarrow v}^t}$, for each $v \in V$.

Algorithm 3, with the choice of initial conditions $W_{e \rightarrow v}^0 = W_{vw}$ and $w_{e \rightarrow v}^0 = 0$, for each $e = \{v, w\} \in E, v \in \partial e$, will be the focus of our investigation for the voltage problem — the convergence analysis will rely on the general formulation of Algorithm 2.

3.2 Min-sum for the flow problem

To apply the min-sum algorithm to solve the flow problem (3), we first rewrite this problem in the unconstrained form (5). To this end, let us identify $\mathcal{V} \leftrightarrow \vec{E}$ and $\mathcal{F} \leftrightarrow V$. In this

context, the notation ∂v indicates the set of edges in \vec{E} that either enter or leave node $v \in V$; analogously, the notation ∂e indicates the pair of vertices that are connected by edge e . For each $e \in \vec{E}$, define the function $g_e(x_e) = \frac{1}{2}R_{ee}x_e^2$, and for each $v \in V$ define the function $g_v(x_{\partial v}) := 0$ if $\sum_{e \in \vec{E}} A_{ve}x_e = b_v$, and $g_v(x_{\partial v}) := \infty$ otherwise. Clearly, problem (3) is equivalent to the problem of minimizing the function $\sum_{e \in \vec{E}} g_e(x_e) + \sum_{v \in V} g_v(x_{\partial v})$, which is of the form (5). As in the factor graph interpretation each variable node (i.e., edge in \vec{E}) is adjacent to exactly two factor nodes (i.e., vertices in V), it is convenient to rewrite the min-sum algorithm in terms of updates of messages from variables to factors only. At each time step $t \geq 1$, to each edge $e \in \vec{E}$ such that $\partial e = \{v, w\}$ (that is, either $e = (v, w)$ or $e = (w, v)$) are associated two messages: $\mu_{e \rightarrow v}^t$ and $\mu_{e \rightarrow w}^t$. It is easy to verify that in the present setting Algorithm 1 reads as follows.²

Algorithm 4: Min-sum, flow problem

Input: Initial messages $\{\mu_{e \rightarrow v}^0\}$, for $e \in \vec{E}, v \in \partial e$;
for $s \in \{1, \dots, t\}$ **do**
 Compute, for each $e \in \vec{E}, v \in \partial e$, with $w = \partial e \setminus v$
 $\mu_{e \rightarrow v}^s = g_e + \min_{x \in \mathbb{R}} \left\{ g_w(x_{\partial w \setminus e} \cdot) + \sum_{f \in \partial w \setminus e} \mu_{f \rightarrow w}^{s-1}(x_f) \right\};$
 Compute $\mu_e^t = \mu_{e \rightarrow v}^t + \mu_{e \rightarrow w}^t - g_e$, for $e = (v, w) \in \vec{E}$;
Output: $\hat{x}_e^t := \arg \min_{z \in \mathbb{R}} \mu_e^t(z)$, for each $e \in \vec{E}$.

Note that if $\partial w \setminus e = \emptyset$, then the message $\mu_{e \rightarrow v}^s$ in Algorithm 4 reads $\mu_{e \rightarrow v}^s(z) = \frac{1}{2}R_{ee}b_w^2$ if $z = A_{we}b_w$, and $\mu_{e \rightarrow v}^s(z) = \infty$ otherwise, for any $s \geq 1$, irrespective of the choice of the initial messages. This is a consequence of the fact that if vertex $w \in V$ is a leaf of G , then by Kirchoff's law we know that the value of the flow on the unique edge e that is associated to w has to be $x_e = A_{we}b_w$. Without loss of generality, in what follows we assume that the graph G has no leaves. Otherwise, by removing one leaf at a time as described above, we can reduce the original problem to a new problem supported on a graph that has no leaves (if G is a tree we can exactly compute all the components of the solution x^* in this way).

Algorithm 4 applies for any choice of the initial messages $\{\mu_{e \rightarrow v}^0\}$, for $e \in \vec{E}, v \in \partial e$. If we choose initial messages that are quadratic functions, then also the sequent messages computed by the algorithm remain quadratic. The following Lemma presents the quadratic updates, not stating formulas for the constant terms as these terms do not influence the output of the min-sum algorithm, as we will show next.

Lemma 3 *For a given $s \geq 1$, for each $e \in \vec{E}, v \in \partial e$ let $\mu_{e \rightarrow v}^{s-1}(z) = \frac{1}{2}R_{e \rightarrow v}^{s-1}z^2 + r_{e \rightarrow v}^{s-1}z + c_{e \rightarrow v}^{s-1}$. Then, if the graph G has no leaves, for each $e \in \vec{E}, v \in \partial e$, we have $\mu_{e \rightarrow v}^s(z) = \frac{1}{2}R_{e \rightarrow v}^s z^2 + r_{e \rightarrow v}^s z + c_{e \rightarrow v}^s$, where $\{c_{e \rightarrow v}^s\}$ is a certain set of constants, and (here $w \in \partial e \setminus v$)*

$$R_{e \rightarrow v}^s = R_{ee} + \frac{1}{\sum_{f \in \partial w \setminus e} 1/R_{f \rightarrow w}^{s-1}}, \quad r_{e \rightarrow v}^s = -A_{we} \frac{\sum_{f \in \partial w \setminus e} A_{wf} r_{f \rightarrow w}^{s-1}/R_{f \rightarrow w}^{s-1} + b_w}{\sum_{f \in \partial w \setminus e} 1/R_{f \rightarrow w}^{s-1}}.$$

2. For convenience of notation, we present Algorithm 4 with initial messages indexed by time 0. Hence, messages are indexed by a time parameter that is shifted by one unit compared to Algorithm 1.

Proof Fix $s \geq 1$, $e \in \vec{E}$, $v \in \partial e$, and let $w \in \partial e \setminus v$. As we assume that the graph G has no leaves, then $\partial w \setminus e \neq \emptyset$ and the min-sum update rule reads

$$\mu_{e \rightarrow v}^s(z) = g_e(z) + \min_{x \in \mathbb{R}^{\partial w \setminus e}: \tilde{A}x = b_w - A_{we}z} \sum_{f \in \partial w \setminus e} \mu_{f \rightarrow w}^{s-1}(x_f),$$

where $\tilde{A} := A_{w, \partial w \setminus e}$ is the submatrix of A given by the w -th row and the columns indexed by the elements in $\partial w \setminus e$. The minimization in the definition of $\mu_{e \rightarrow v}^s$ can be written as

$$\begin{aligned} & \text{minimize} && \frac{1}{2} x^T \tilde{R} x + \tilde{r}^T x + 1^T \tilde{c} \\ & \text{subject to} && \tilde{A} x = b_w - A_{we} z \end{aligned} \tag{6}$$

over $\tilde{x} \in \mathbb{R}^{\partial w \setminus e}$, with $\tilde{R} \in \mathbb{R}^{(\partial w \setminus e) \times (\partial w \setminus e)}$ diagonal with $\tilde{R}_{ff} := R_{f \rightarrow w}^{s-1}$, $\tilde{h} \in \mathbb{R}^{\partial w \setminus e}$ with $\tilde{h}_f := r_{f \rightarrow w}^{s-1}$, and $\tilde{c} \in \mathbb{R}^{\partial w \setminus e}$ with $\tilde{c}_f := c_{f \rightarrow w}^{s-1}$. By Proposition 23 in Appendix B the solution of (6) reads $\tilde{x}^* = \tilde{r} - A_{we} z \tilde{y}$, with $\tilde{r} := b_w \tilde{y} - (I - \tilde{y} \tilde{A}) \tilde{R}^{-1} \tilde{h}$, $\tilde{y} := \tilde{R}^{-1} \tilde{A}^T \tilde{L}^{-1}$, and $\tilde{L} := \tilde{A} \tilde{R}^{-1} \tilde{A}^T \in \mathbb{R}$. Substituting \tilde{x}^* in (6) we get $\mu_{e \rightarrow v}^s(z) = \frac{1}{2} R_{e \rightarrow v}^s z^2 + r_{e \rightarrow v}^s z + c_{e \rightarrow v}^s$, with $R_{e \rightarrow v}^s := R_{ee} + \tilde{y}^T \tilde{R} \tilde{y}$, and $r_{e \rightarrow v}^s := -A_{we}(\tilde{y}^T \tilde{R} \tilde{r} + \tilde{r}^T \tilde{y})$. The proof follows as $\tilde{L} = \sum_{f \in \partial w \setminus e} 1/R_{f \rightarrow w}^{s-1}$, $\tilde{y}^T \tilde{R} \tilde{y} = \tilde{L}^{-1}$, $\tilde{y}^T \tilde{R} \tilde{r} = \tilde{L}^{-1} b_w$, $\tilde{r}^T \tilde{y} = \tilde{L}^{-1} \sum_{f \in \partial w \setminus e} A_{wf} r_{f \rightarrow w}^{s-1} / R_{f \rightarrow w}^{s-1}$. ■

From Lemma 3 it follows that if we initialize Algorithm 4 with quadratic messages $\mu_{e \rightarrow v}^0(z) = \frac{1}{2} R_{e \rightarrow v}^0 z^2 + r_{e \rightarrow v}^0 z$, then, modulo the constant terms, we can track the evolution of the quadratic messages computed by the algorithm by evaluating recursively the set of parameters $\{R_{e \rightarrow v}^s\}, \{r_{e \rightarrow v}^s\}$ as prescribed by Lemma 3. The belief function associated to edge e at time t then reads $\mu_e^t(z) = \frac{1}{2} (R_{e \rightarrow v}^t + R_{e \rightarrow w}^t - R_{ee}) z^2 + (r_{e \rightarrow v}^t + r_{e \rightarrow w}^t) z + c_e^t$, for a certain constant c_e^t , so that $\hat{x}_e^t = \arg \min_{z \in \mathbb{R}} \mu_e^t(z) = -\frac{r_{e \rightarrow v}^t + r_{e \rightarrow w}^t}{R_{e \rightarrow v}^t + R_{e \rightarrow w}^t - R_{ee}}$, independent of c_e^t . The final procedure so obtained, in the case of quadratic initial messages, for a graph G with no leaves, is given by Algorithm 5.

Algorithm 5: Min-sum, flow problem, quadratic messages, no leaves

Input: Initial messages $\{R_{e \rightarrow v}^0\}, \{r_{e \rightarrow v}^0\}$, $e \in \vec{E}, v \in \partial e$;

for $s \in \{1, \dots, t\}$ **do**

Compute, for each $e \in \vec{E}$, $v \in \partial e$, with $w = \partial e \setminus v$

$$\left[\begin{aligned} R_{e \rightarrow v}^s &= R_{ee} + \frac{1}{\sum_{f \in \partial w \setminus e} 1/R_{f \rightarrow w}^{s-1}}, & r_{e \rightarrow v}^s &= -A_{we} \frac{\sum_{f \in \partial w \setminus e} A_{wf} r_{f \rightarrow w}^{s-1} / R_{f \rightarrow w}^{s-1} + b_w}{\sum_{f \in \partial w \setminus e} 1/R_{f \rightarrow w}^{s-1}}; \end{aligned} \right.$$

Output: $\hat{x}_e^t = -\frac{r_{e \rightarrow v}^t + r_{e \rightarrow w}^t}{R_{e \rightarrow v}^t + R_{e \rightarrow w}^t - R_{ee}}$, for $e = (v, w) \in \vec{E}$.

Algorithm 5, with the choice of initial conditions $R_{e \rightarrow v}^0 = R_{ee}$, $r_{e \rightarrow v}^0 = 0$, for each $e \in \vec{E}, v \in \partial e$, will be the focus of our investigation for the flow problem — the convergence analysis will rely on the general formulation of Algorithm 4.

4. Results

In this section we characterize the error committed by the min-sum algorithms to solve the voltage and flow problem when applied to connected d -regular graphs, and we present some

convergence results. Henceforth, let $\hat{\nu}^t$ be the outcome of Algorithm 3 at time t with the choice of initial conditions $W_{e \rightarrow v}^0 = W_{vw}$ and $w_{e \rightarrow v}^0 = 0$, for each $e = \{v, w\} \in E, v \in \partial e$. Analogously, let \hat{x}^t be the outcome of Algorithm 5 at time t with the choice of initial conditions $R_{e \rightarrow v}^0 = R_{ee}$, $r_{e \rightarrow v}^0 = 0$, for each $e \in \vec{E}, v \in \partial e$. The proofs of the main results here presented — Lemma 4 and Lemma 6 — are given in Section 5, where we develop the general machinery to analyze the behavior of the min-sum algorithm to solve constrained optimization problems. The proof of the other technical results are given in Appendix B.

Henceforth, given a weighted directed graph $\vec{G} = (V, \vec{E}, W)$, we let $G = (V, E, W)$ be the weighted undirected graph naturally associated to \vec{G} , i.e., E is obtained by removing the orientation from the edges in \vec{E} .

4.1 Error characterization

The case $d = 2$, i.e., a cycle graph, allows a characterization of the error committed by the min-sum algorithms for weighted graphs in terms of the voltage solution $\nu^* = L^+b$.

Lemma 4 (Regular graphs, $d = 2$) *Let $\vec{G} = (V = \{0, \dots, n-1\}, \vec{E}, W)$ be a weighted directed cycle, and let $G = (V, E, W)$ be the corresponding undirected graph. Define the function $i \in \mathbb{N} \rightarrow \rho(i) := i \bmod n$. For $v \in V$, $e = \{\rho(v), \rho(v+1)\} \in E$, $t \geq 2$, we have*

$$\begin{aligned} \nu_v^* - \hat{\nu}_v^t &= \alpha_{v,t} \nu_{\rho(v-t-1)}^* + (1 - \alpha_{v,t}) \nu_{\rho(v+t+1)}^*, & \alpha_{v,t} &:= \frac{\sum_{k=v}^{v+t} 1/W_{\rho(k)\rho(k+1)}}{\sum_{k=v-t-1}^{v+t} 1/W_{\rho(k)\rho(k+1)}}, \\ x_e^* - \hat{x}_e^t &= A_{ve} \beta_{e,t} (\nu_{\rho(v-t)}^* - \nu_{\rho(v+t+1)}^*), & \beta_{e,t} &:= \frac{1}{\sum_{k=v-t}^{v+t} 1/W_{\rho(k)\rho(k+1)}}. \end{aligned}$$

If each edge has the same weight ω , then $\alpha_{v,t} = 1/2$ and $\beta_{e,t} = \omega/(2t+1)$.

For $d \geq 3$ we can provide a characterization of the error committed by the min-sum algorithms in the case of graphs with equal weights. Before stating these results, we define and bound some quantities of interest. The proof of the following proposition is given in Appendix B.

Proposition 5 *Fix $d \geq 3$. For $s \geq 1$ let $h_s := \frac{1}{(d-1)^s - 1}$, and let $\delta_0 := \frac{1}{d}$, $\delta_1 := \frac{1}{d-1} + d - 1$, and $\delta_s := \frac{1}{d-1} (2 + \frac{(d-2)^2}{d-1} (1 + h_{s+2})) \delta_{s-1} - \frac{1}{(d-1)^2} \delta_{s-2}$ for $s \geq 2$. For any $t \geq 3$ define*

$$\begin{aligned} b_{d,t} &:= \frac{1}{(d-1)^2} \left(1 + (d-2)(1 + h_{t+1}) \right) \delta_{t-2} - \frac{(d-2)(1 + h_{t+1})}{(d-1)^3} \delta_{t-3}, \\ c_{d,t} &:= \frac{1}{d-1} \left(1 + \frac{1}{(d-2)(1 + h_t)} \right) \delta_{t-2} - \frac{1}{(d-1)^2} \delta_{t-3}. \end{aligned}$$

Then we have $\frac{1}{2} \leq \frac{d-2}{d-1} \leq b_{d,t} \leq c_{d,t} \leq 1 + \varepsilon_d < 4$ and $c_{d,t} \geq 1$, where ε_d is a positive decreasing function of $d \geq 3$ so that $\varepsilon_d \rightarrow 0$ as $d \rightarrow \infty$ and $\varepsilon_3 < 3$.

For each $v \in V$ we let \mathbf{P}_v be the law of a time-homogeneous *non-backtracking* random walk Y_0, Y_1, Y_2, \dots on the vertices of G with initial state $Y_0 = v$, which is defined by

$$\mathbf{P}_v(Y_0 = w) = \begin{cases} 1 & \text{if } w = v, \\ 0 & \text{otherwise,} \end{cases} \quad \mathbf{P}_v(Y_1 = w) = \begin{cases} 1/d & \text{if } \{w, v\} \in E, \\ 0 & \text{otherwise,} \end{cases}$$

and, for any $\{z, z'\} \in E$, by

$$\mathbf{P}_v(Y_{t+1} = w | Y_t = z, Y_{t-1} = z') = \begin{cases} \frac{1}{d-1} & \text{if } \{w, z\} \in E, w \neq z', \\ 0 & \text{otherwise.} \end{cases}$$

For ease of notation, for $t \geq 0$, $w \in V$, define the row-stochastic matrices $P^{(t)}$ and $P^{(t,w)}$ as

$$P_{vz}^{(t)} := \mathbf{P}_v(Y_t = z), \quad P_{vz}^{(t,w)} := \mathbf{P}_v(Y_t = z | Y_1 \neq w), \quad v, z \in V.$$

We can now state the main results for d -regular graphs with $d \geq 3$ and equal weights, namely, $W_{vw} = \omega$ for any $\{v, w\} \in E$, for some $\omega > 0$.

Lemma 6 (Regular graphs, $d \geq 3$, equal weights) *Let $\vec{G} = (V, \vec{E}, W)$ be a connected directed d -regular graph, with $d \geq 3$, where each edge has the same weight $\omega > 0$, and let $G = (V, E, W)$ be the corresponding undirected graph. For $v \in V$, $e = (v, w) \in \vec{E}$, $t \geq 3$,*

$$\nu_v^* - \hat{\nu}_v^t = \frac{1}{b_{d,t}} (P^{(t+1)} \nu^*)_v, \quad x_e^* - \hat{x}_e^t = \frac{\omega}{c_{d,t}} \frac{d-1}{d} ((P^{(t,w)} \nu^*)_v - (P^{(t,v)} \nu^*)_w),$$

where $b_{d,t}$ and $c_{d,t}$ are defined and bounded as in Proposition 5.

Given the error characterizations provided by Lemma 4 and Lemma 6, we will now investigate the convergence properties of the min-sum algorithms.

4.2 Convergence results for the voltage problem

We first analyze the convergence behavior of the min-sum algorithm when applied to the voltage problem: Algorithm 3. Following the literature on Laplacian solvers, as already noticed in Section 2, we perform the analysis in the L -norm, which is defined as $\|\nu\|_L := \sqrt{\nu^T L \nu}$ for any $\nu \in \mathbb{R}^V$, where $\nu^T L \nu = \sum_{\{v,w\} \in E} W_{vw} (\nu_v - \nu_w)^2$.

In the case $d = 2$, the results of Lemma 4 attest that the output $\hat{\nu}^t$ of Algorithm 3 does *not* converge, as it keeps oscillating as a function of t . This is particularly evident in the case of equal weights, as the following corollary attests.

Corollary 7 (Regular graphs, $d = 2$) *Consider the setting of Lemma 4, in the case when each edge has the same weight ω . For $t \geq 2$ we have*

$$\|\nu^* - \hat{\nu}^t\|_L^2 = \frac{1}{2} \|\nu^*\|_L^2 + \frac{\omega}{2} \sum_{v=0}^{n-1} \nu_v^* (2\nu_{\rho(v+2t+2)}^* - \nu_{\rho(v+2t+3)}^* - \nu_{\rho(v+2t+1)}^*).$$

Proof The proof follows immediately from Lemma 4 noticing that $\|\nu^* - \hat{\nu}^t\|_L^2$ equals

$$\omega \sum_{v=0}^{n-1} (\nu_{v+1}^* - \hat{\nu}_{v+1}^t - (\nu_v^* - \hat{\nu}_v^t))^2 = \frac{\omega}{4} \sum_{v=0}^{n-1} (\nu_{\rho(v-t)}^* - \nu_{\rho(v-t-1)}^* + \nu_{\rho(v+t+2)}^* - \nu_{\rho(v+t+1)}^*)^2,$$

which reads $\frac{1}{2} \|\nu^*\|_L^2 + \frac{\omega}{2} \sum_{v=0}^{n-1} (\nu_{\rho(v-t)}^* - \nu_{\rho(v-t-1)}^*)(\nu_{\rho(v+t+2)}^* - \nu_{\rho(v+t+1)}^*)$. The result follows by grouping and relabeling the terms in the sum. \blacksquare

The oscillatory behavior of the min-sum algorithm on cycle graphs in the voltage problem marks a clear difference with what happens in the flow problem where, as we will see in Corollary 9 below, the algorithm *does* converge to the problem solution.

Let us now consider the case $d \geq 3$. In this case, as Corollary 8 below attests, the convergence of the min-sum algorithm in the L -norm can be controlled by the total variation distance between the probability distributions of non-backtracking random walks that emanate from neighbor nodes. As in bipartite graphs random walks that start at two neighbor locations never overlap, we also consider an averaged version of the min-sum algorithm, where the output at two consecutive iterations is properly averaged.

In this framework, a natural quantity to consider is the total variation normalized version of the L -norm. Given a row-stochastic matrix $M \in \mathbb{R}^{V \times V}$ and $\nu \in \mathbb{R}^V$, define

$$\|\nu\|_{L,M}^2 := \sum_{\{v,w\} \in E} \frac{W_{vw}}{(\frac{1}{2} \sum_{z \in V} |M_{vz} - M_{wz}|)^2} ((M\nu)_v - (M\nu)_w)^2.$$

The quantity $\frac{1}{2} \sum_{z \in V} |M_{vz} - M_{wz}|$ corresponds to the total variation distance between the probability mass functions defined by the v -th and w -th rows of the matrix M , respectively. Clearly, $\|\nu\|_{L,M} = \|M\nu\|_{L'}$, where L' is the Laplacian of the weighted graph (V, E, W') , with $W'_{vw} := W_{vw}/(\frac{1}{2} \sum_{z \in V} |M_{vz} - M_{wz}|)^2$. Furthermore, it is immediate to show that

$$\|M\nu\|_L \leq \frac{1}{2} \left(\max_{\{v,w\}} \sum_{z \in V} |M_{vz} - M_{wz}| \right) \|\nu\|_{L,M}. \quad (7)$$

We now state a convergence result for the min-sum algorithm for the voltage problem, in the case of d -regular graphs, $d \geq 3$, with the same edge weights, with respect to the total variation normalized L -norm. For ease of notation, for any $t \geq 0$, define $\Delta^{(t)} \in \mathbb{R}^{\vec{E} \times V}$ by

$$\Delta_{ez}^{(t)} := P_{vz}^{(t)} - P_{wz}^{(t)}, \quad e = (v, w) \in \vec{E}, z \in V.$$

Corollary 8 (Regular graphs, $d \geq 3$, equal weights) *Consider the setting of Lemma 6. For $t \geq 4$ let $\hat{\nu}_{\text{ave}}^t := \frac{b_{d,t-1}}{b_{d,t-1} + b_{d,t}} \hat{\nu}^{t-1} + \frac{b_{d,t}}{b_{d,t-1} + b_{d,t}} \hat{\nu}^t$ and $Q^{(t)} := (P^{(t-1)} + P^{(t)})/2$. Then,*

$$\|\nu^\star - \hat{\nu}^t\|_L \leq \|\Delta^{(t+1)}\|_\infty \|\nu^\star\|_{L, P^{(t+1)}}, \quad \|\nu^\star - \hat{\nu}_{\text{ave}}^t\|_L \leq \frac{1}{2} \|\Delta^{(t)} + \Delta^{(t+1)}\|_\infty \|\nu^\star\|_{L, Q^{(t+1)}}.$$

Proof From Lemma 6, using (7) with $M = P^{(t+1)}$, we get

$$\|\nu^\star - \hat{\nu}^t\|_L^2 = \frac{1}{b_{d,t}^2} \|P^{(t+1)}\nu^\star\|_L^2 \leq \|\Delta^{(t+1)}\|_\infty^2 \|\nu^\star\|_{L, P^{(t+1)}}^2,$$

as $b_{d,t} \geq 1/2$ by Proposition 5. From Lemma 6 $\nu^\star - \hat{\nu}^t = P^{(t+1)}\nu^\star/b_{d,t}$ so that for $t \geq 4$

$$(b_{d,t-1} + b_{d,t})(\nu^\star - \hat{\nu}_{\text{ave}}^t) = b_{d,t-1}(\nu^\star - \hat{\nu}^{t-1}) + b_{d,t}(\nu^\star - \hat{\nu}^t) = (P^{(t)} + P^{(t+1)})\nu^\star.$$

Proceeding as above we get

$$\|\nu^\star - \hat{\nu}_{\text{ave}}^t\|_L^2 = \frac{4}{(b_{d,t-1} + b_{d,t})^2} \|Q^{(t+1)}\nu^\star\|_L^2 \leq \left(\frac{1}{2} \|\Delta^{(t)} + \Delta^{(t+1)}\|_\infty \right)^2 \|\nu^\star\|_{L, Q^{(t+1)}}^2. \quad \blacksquare$$

Corollary 8 shows that the convergence of the min-sum estimate $\hat{\nu}^t$ and its averaged version $\hat{\nu}_{\text{ave}}^t$ in the L -norm is controlled by the behavior with time t of the ℓ_∞ norm of the matrices $\Delta^{(t+1)}$ and $\Delta^{(t)} + \Delta^{(t+1)}$, respectively. These quantities yield a uniform control on the total variation distances (modulo $1/2$ -factors) between the probability distributions of non-backtracking random walks that emanate from neighbor nodes. For some choices of graphs these quantities are seen to exhibit decay with time *independently* of the size of the graph. This is the case, for instance, for the $d/2$ -connected cycle and the $d/2$ -dimensional torus, with $d \geq 4$ even. See Figure 1 and Figure 2.




Figure 1: A 2-connected cycle (left). Plot in logarithmic scale of $\|\Delta^{(t)}\|_\infty$ versus time t for various $d/2$ -connected cycles, where d is the graph degree (right).




Figure 2: A 2-dimensional torus — periodic boundary conditions are not drawn (left). Plot in logarithmic scale of $\|\Delta^{(t)} + \Delta^{(t+1)}\|_\infty$ versus time t for various $d/2$ -dimensional tori, where d is the graph degree (right).

In both these two classes of graphs the respective quantities of interest are empirically seen to decay like $O(1/\sqrt{t})$, independently³ of the graph dimensions $n = |V|$ and $m = |E|$.

3. It is implied that the decay with time of these quantities is independent of n and m as long as time t satisfies $t \lesssim D/d$, where D is the diameter of the graph. Once the random walks reach the periodic boundary conditions the decay is even faster due to the extra overlaps and cancellations of the distributions.

Each iteration of the min-sum algorithm requires a complexity that is linear in the number of edges, i.e., $O(m)$. Hence, in problem instances where one can show that $\|\nu^*\|_{L,P(t)} \lesssim \|\nu^*\|$ or $\|\nu^*\|_{L,Q(t)} \lesssim \|\nu^*\|$ for any t (potentially under some assumptions on the problem input b), in graphs where $\|\Delta^{(t+1)}\|_\infty$ or $\|\Delta^{(t)} + \Delta^{(t+1)}\|_\infty$ decays like $O(1/\sqrt{t})$, for instance, the min-sum estimate is seen to attain a ε relative error in the L -norm with a running time $O(m/\varepsilon^2)$, *linear* in the number of edges m .

4.3 Convergence results for the flow problem

We now analyze the convergence behavior of the min-sum algorithm when applied to the flow problem: Algorithm 5. We perform the analysis in the ℓ_∞ norm.

For the case $d = 2$, Lemma 4 immediately yields the following result.

Corollary 9 (Regular graphs, $d = 2$) *Consider the setting of Lemma 4. Let $\omega_M := \max_{\{v,w\} \in E} W_{vw}$. Then,*

$$\|x^* - \hat{x}^t\|_\infty \leq \frac{2\omega_M}{2t+1} \|\nu\|_\infty.$$

Proof From Lemma 4, by the triangle inequality, we have $|x_e^* - \hat{x}_e^t| \leq \frac{\omega_M}{2t+1} (|\nu_{\rho(v-t)}^*| + |\nu_{\rho(v+t+1)}^*|)$. The proof follows immediately by the triangle inequality for norms. \blacksquare

As each iteration of the min-sum algorithm requires a complexity $O(m)$, the min-sum estimate for the cycle graph is guaranteed to get a ε relative error in the ℓ_∞ norm with a running time $O(m/\varepsilon)$, linear in m .

Let us now consider the case $d \geq 3$. For any $t \geq 1$, define the matrix $\tilde{\Delta}^{(t)} \in \mathbb{R}^{\vec{E} \times V}$ as:

$$\tilde{\Delta}_{ez}^{(t)} := P_{vz}^{(t,w)} - P_{wz}^{(t,v)}, \quad e = (v, w) \in \vec{E}, z \in V. \quad (8)$$

Lemma 6 yields $x^* - \hat{x}^t = \omega(d-1)\tilde{\Delta}^{(t)}\nu^*/(dc_{d,t})$, which allows for an immediate error analysis of the min-sum algorithm in terms of the behavior with time t of the ℓ_∞ norm of $\tilde{\Delta}^{(t)}$. As in bipartite graphs random walks that start at two neighbor locations never overlap, we also look at an averaged version of the min-sum algorithm, where the output at two consecutive iterations is properly averaged.

Corollary 10 (Regular graphs, $d \geq 3$, equal weights) *Consider the setting of Lemma 6. For $t \geq 4$ define $\hat{x}_{\text{ave}}^t := \frac{c_{d,t-1}}{c_{d,t-1}+c_{d,t}}\hat{x}^{t-1} + \frac{c_{d,t}}{c_{d,t-1}+c_{d,t}}\hat{x}^t$. Then,*

$$\|x^* - \hat{x}^t\|_\infty \leq \omega \frac{d-1}{d} \|\tilde{\Delta}^{(t)}\|_\infty \|\nu^*\|_\infty, \quad \|x^* - \hat{x}_{\text{ave}}^t\|_\infty \leq \frac{\omega}{2} \frac{d-1}{d} \|\tilde{\Delta}^{(t-1)} + \tilde{\Delta}^{(t)}\|_\infty \|\nu^*\|_\infty.$$

Proof From Lemma 6 we have $x^* - \hat{x}^t = \frac{\omega}{c_{d,t}} \frac{d-1}{d} \tilde{\Delta}^{(t)} \nu^*$ so that for any $t \geq 4$ we have

$$(c_{d,t-1} + c_{d,t})(x^* - \hat{x}_{\text{ave}}^t) = c_{d,t-1}(x^* - \hat{x}^{t-1}) + c_{d,t}(x^* - \hat{x}^t) = \omega \frac{d-1}{d} (\tilde{\Delta}^{(t-1)} + \tilde{\Delta}^{(t)}) \nu^*.$$

The proof follows immediately using the sub-multiplicative property of induced norms, and using that $c_{d,t} \geq 1$ from Proposition 5. \blacksquare

Corollary 10 shows that the convergence of the min-sum estimate \hat{x}^t and its averaged version \hat{x}_{ave}^t in the ℓ_∞ norm is controlled by the behavior with time t of the ℓ_∞ norm of the matrices $\tilde{\Delta}^{(t)}$ and $\tilde{\Delta}^{(t-1)} + \tilde{\Delta}^{(t)}$, respectively. These quantities yield a uniform control on the total variation distances (modulo $1/2$ -factors) between the probability distributions of non-backtracking random walks that emanate from neighbor nodes, conditioned on the event that the first step of the two walks avoids each other's locations. Proposition 24 in Appendix B yields a recursive formula for computing the vectors $\tilde{\Delta}_e^{(t)}$'s in terms of the adjacency matrix of the undirected unweighted graph $G = (V, E)$. Analogously to the voltage case previously discussed, for some choices of graphs these quantities are seen to exhibit decay with time *independently* of the size of the graph. This is the case, once again, for the $d/2$ -connected cycle and the $d/2$ -dimensional torus, with $d \geq 4$ is even, where the respective quantities of interest are seen to decay like $O(1/\sqrt{t})$. For these graphs the min-sum estimate is seen to attain a ε relative error in the ℓ_∞ norm with a running time $O(m/\varepsilon^2)$, linear in m .

Remark 11 (Analysis in ℓ_2 norm) *For d -regular graphs with $d \geq 3$, the analysis in the ℓ_2 norm for the flow problem follows the same steps as the analysis in the L -norm for the voltage problem. In fact, $\|x^* - \hat{x}^t\|_2 = \frac{\omega}{c_{d,t}} \frac{d-1}{d} \|\tilde{\Delta}^{(t)} \nu^*\|_2$, while $\|\nu^* - \hat{\nu}^t\|_L = \frac{1}{b_{d,t}} \|\Delta^{(t)} \nu^*\|_2$.*

5. Proofs

This section contains the proofs of the main results in Section 4, namely, Lemma 4 and Lemma 6. While the main architecture of the proofs for the results concerning the min-sum algorithm applied respectively to the voltage problem (Algorithm 3) and to the flow problem (Algorithm 5) is similar, there are some key differences that need to be taken into account. First, as the voltage algorithm updates functions (i.e., messages) on the vertices while the flow algorithm updates functions on the edges, the two algorithms give rise to different topologies for the computation tree that supports their operations with time. Second, as the voltage problem is an unconstrained optimization problem while the flow problem is an optimization problem with constraints, a different mechanism to find the fix point(s) of the algorithm is needed. As in the literature, to the best of our knowledge, previous analyses of the min-sum algorithm for quadratic and, more generally, convex problems only focused on the unconstrained case, we first give the full details of the proofs for the flow problem. Then, we give an overview of the proofs for the voltage problem, outlining in details only the parts where the analysis differs significantly from the one done for the flow problem.

5.1 Proof for the flow problem, Algorithm 3

To begin with, we follow the general scheme presented in Moallemi and Van Roy (2010) for unconstrained optimization problems, and investigate the evolution of the min-sum algorithm under a set of linear perturbations. Henceforth, we consider Algorithm 5 when the initial messages are parametrized by $\{R_{e \rightarrow v}^0 = R_{ee}\}, \{r_{e \rightarrow v}^0 = p_{e \rightarrow v}\}$, for a certain set of real numbers $p = \{p_{e \rightarrow v}\}$. In fact, as far as the analysis is concerned, it is convenient (and more amenable to generalizations) to consider the form of Algorithm 4. For a given choice

of perturbation parameters $p = \{p_{e \rightarrow v}\}$, define the initial messages as

$$\mu_{e \rightarrow v}^0(\cdot, p) : z \in \mathbb{R} \longrightarrow \mu_{e \rightarrow v}^0(z, p) = \frac{1}{2} R_{ee} z^2 + p_{e \rightarrow v} z,$$

and denote by $\{\mu_{e \rightarrow v}^s(\cdot, p)\}$, $s \geq 1$, the corresponding sequence of messages generated by the min-sum algorithm. Denote the estimates at time t as $\hat{x}^t(p)_e := \arg \min_{z \in \mathbb{R}} \mu_e^t(z, p)$, where $\mu_e^t(\cdot, p)$ is the corresponding belief function. We also introduce the notation $\mathcal{N}(v) := \{w \in V : \{v, w\} \in E\}$ to denote the neighborhood of vertex v in G .

5.1.1 COMPUTATION TREE (FLOW PROBLEM)

To investigate the convergence behavior of the min-sum algorithm, we recall the concept of the computation tree. Given an edge $\tilde{e} = (\tilde{v}, \tilde{w}) \in \vec{E}$, the computation tree rooted at \tilde{e} of depth t supports the optimization problem that is obtained by unfolding the computations involved in the min-sum estimate $\hat{x}^t(p)_{\tilde{e}}$. Formally, the computation tree is a directed tree $\vec{\mathbb{T}} = (\mathbb{V}, \vec{\mathbb{E}})$ where each vertex in \mathbb{V} is mapped to a vertex in V through a map $\sigma : \mathbb{V} \rightarrow V$ that preserves the edge structure, namely, if $\mathfrak{e} = (\mathfrak{v}, \mathfrak{w}) \in \vec{\mathbb{E}}$ then $\sigma(\mathfrak{e}) := (\sigma(\mathfrak{v}), \sigma(\mathfrak{w})) \in \vec{E}$. The tree $\vec{\mathbb{T}}$ is defined iteratively. Initially, at level 0, the tree corresponds to a single root edge $\tilde{\mathfrak{e}} = (\tilde{\mathfrak{v}}, \tilde{\mathfrak{w}}) \in \vec{\mathbb{E}}$ corresponding to \tilde{e} , i.e., $\sigma(\tilde{\mathfrak{e}}) = \tilde{e}$. At this stage, $\tilde{\mathfrak{v}}$ and $\tilde{\mathfrak{w}}$ are the leaves of the tree. For the remaining t levels, the following procedure is repeated. The leaf in the tree are examined. Given a leaf \mathfrak{v} with $\sigma(\mathfrak{v}) = v$ that is connected to a vertex \mathfrak{w} with $\sigma(\mathfrak{w}) = w$, for any $z \in \mathcal{N}(v) \setminus w$, a vertex \mathfrak{z} with $\sigma(\mathfrak{z}) = z$ and a directed edge $(\mathfrak{v}, \mathfrak{z})$ (resp. $(\mathfrak{z}, \mathfrak{v})$) are added to the next level of the tree if $(v, z) \in \vec{E}$ (resp. $(z, v) \in \vec{E}$). Figure 3 gives an example. We denote the set of vertices and edges in the k -th level of the tree respectively by $\mathbb{V}^k \subset \mathbb{V}$ and $\mathbb{E}^k \subset \mathbb{E}$. In what follows we also extend the usual neighborhood notation to vertices and edges in the graph \mathbb{T} , namely, $\partial \mathfrak{v}$ is the set of edges in \mathbb{T} that are connected to node \mathfrak{v} , and $\partial \mathfrak{e}$ is the set of vertices in \mathbb{T} that are connected to edge \mathfrak{e} .




Figure 3: Graph $\vec{G} = (V = \{1, 2, 3, 4\}, \vec{E})$ (left) and computation tree $\vec{\mathbb{T}} = (\mathbb{V}, \vec{\mathbb{E}})$ of depth $t = 3$ with root edge corresponding to $(1, 2) \in \vec{E}$ (right). For the sake of simplicity, vertices in $\vec{\mathbb{T}}$ are labeled by the corresponding vertices in \vec{G} .

The t -th (last) level of the tree $\vec{\mathbb{T}}$ is of particular relevance, as it supports the initial conditions of the min-sum algorithm, as we are now about to describe. Let $\mathbb{R} \in \mathbb{R}^{\vec{\mathbb{E}} \times \vec{\mathbb{E}}}$ be the diagonal matrix defined by $\mathbb{R}_{\mathfrak{e}\mathfrak{e}} := R_{\sigma(\mathfrak{e})\sigma(\mathfrak{e})}$. For a given choice of the perturbation parameters $p = \{p_{e \rightarrow v}\}$, define $\mathbb{p}(p) \in \mathbb{R}^{\vec{\mathbb{E}}}$ by

$$\mathbb{p}(p)_{\mathfrak{e}} := \begin{cases} p_{\sigma(\mathfrak{e}) \rightarrow \sigma(\mathfrak{v})} & \text{if } \mathfrak{e} \in \vec{\mathbb{E}}^t, \mathfrak{v} = \partial \mathfrak{e} \cap \mathbb{V}^{t-1} \\ 0 & \text{otherwise.} \end{cases}$$

Let $\bar{\mathbb{V}} := \mathbb{V} \setminus \mathbb{V}^t$. Define the matrix $\bar{\mathbb{A}} \in \mathbb{R}^{\bar{\mathbb{V}} \times \bar{\mathbb{E}}}$ as $\bar{A}_{\mathbf{ve}} := A_{\sigma(\mathbf{v})\sigma(\mathbf{e})}$ if $\mathbf{e} \in \partial \mathbf{v}$ and $\bar{A}_{\mathbf{ve}} := 0$ otherwise. Define the vector $\bar{\mathbf{b}} \in \mathbb{R}^{\bar{\mathbb{V}}}$ as $\bar{b}_{\mathbf{v}} := b_{\sigma(\mathbf{v})}$. Consider the following problem supported on the tree $\bar{\mathbb{T}}$, over $\mathbf{x} \in \mathbb{R}^{\bar{\mathbb{E}}}$:

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \mathbf{x}^T \mathbb{R} \mathbf{x} + \mathbb{p}(p)^T \mathbf{x} \\ & \text{subject to} && \bar{\mathbb{A}} \mathbf{x} = \bar{\mathbf{b}}. \end{aligned} \tag{9}$$

Let $\mathbf{x}^*(p)$ denote the unique optimal solution of this problem, as a function of the perturbation p . The relationship between the min-sum algorithm and the computation tree is made explicit by the following lemma, which can be easily established by inductively examining the operations of Algorithm 4.

Lemma 12 *Given the setting of this section, $\hat{x}^t(p)_{\bar{\mathbf{e}}} = \mathbf{x}^*(p)_{\bar{\mathbf{e}}}$.*

We now introduce a few quantities related to the computation tree $\bar{\mathbb{T}}$ and the problem (9) that will play a key role in our analysis. The signed vertex-edge incidence matrix $\mathbb{A} \in \mathbb{R}^{\mathbb{V} \times \bar{\mathbb{E}}}$ associated to the directed graph $\bar{\mathbb{T}}$ is defined for each $\mathbf{v} \in \mathbb{V}$, $\mathbf{e} \in \bar{\mathbb{E}}$ as $A_{\mathbf{ve}} := 1$ if edge \mathbf{e} leaves node \mathbf{v} , $A_{\mathbf{ve}} := -1$ if edge \mathbf{e} enters node \mathbf{v} , and $A_{\mathbf{ve}} := 0$ otherwise. It is easy to verify that $\bar{\mathbb{A}}$ is the submatrix of \mathbb{A} that corresponds to the rows associated to $\bar{\mathbb{V}}$. Let $\mathbb{W} \in \mathbb{R}^{\mathbb{V} \times \mathbb{V}}$ be a symmetric matrix that assigns a positive weight to every edge in $\bar{\mathbb{E}}$, defined as $W_{\mathbf{vw}} := W_{\sigma(\mathbf{v})\sigma(\mathbf{w})}$ if either $(\mathbf{v}, \mathbf{w}) \in \bar{\mathbb{E}}$ or $(\mathbf{w}, \mathbf{v}) \in \bar{\mathbb{E}}$, and $W_{\mathbf{vw}} := 0$ otherwise. Henceforth, we think of $\bar{\mathbb{T}} = (\mathbb{V}, \bar{\mathbb{E}}, \mathbb{W})$ as the corresponding weighted graph, and $\mathbb{T} = (\mathbb{V}, \mathbb{E}, \mathbb{W})$ as the undirected graph naturally associated to $\bar{\mathbb{T}}$, where the edge set \mathbb{E} is obtained by removing the orientation from the edges in $\bar{\mathbb{E}}$. As \mathbb{T} is connected by assumption, it is easy to show that the range of \mathbb{A} is spanned by the vectors orthogonal to the all-ones vector. It follows that the matrix $\bar{\mathbb{A}}$ has full row rank. Define the weighted degree of a vertex \mathbf{v} by $d_{\mathbf{v}} := \sum_{\mathbf{w} \in \mathbb{V}} W_{\mathbf{vw}}$, and let $\mathbb{D} \in \mathbb{R}^{\mathbb{V} \times \mathbb{V}}$ be the diagonal matrix defined by $D_{\mathbf{vv}} := d_{\mathbf{v}}$. The Laplacian \mathbb{L} of the graph $\bar{\mathbb{T}}$ is the matrix defined as $\mathbb{L} := \mathbb{D} - \mathbb{W}$. Let $\bar{\mathbb{L}} \in \mathbb{R}^{\bar{\mathbb{V}} \times \bar{\mathbb{V}}}$ denote the submatrix of \mathbb{L} that corresponds to the rows and columns associated to $\bar{\mathbb{V}}$. Note that $R_{\mathbf{ee}} := 1/W_{\mathbf{vw}}$. It is easy to verify that the following equalities holds: $\mathbb{L} = \mathbb{A} \mathbb{R}^{-1} \mathbb{A}^T$ and $\bar{\mathbb{L}} = \bar{\mathbb{A}} \mathbb{R}^{-1} \bar{\mathbb{A}}^T$. We also define the transition matrix of the ordinary diffusion random walk on \mathbb{T} as the matrix $\mathbb{P} := \mathbb{D}^{-1} \mathbb{W}$.

5.1.2 FIX POINT (FLOW PROBLEM)

Clearly, the choice $p = 0$ for the linear perturbation parameters yields the original (unperturbed) algorithm. The next theorem below shows that for a particular choice of p the min-sum algorithm yields the optimal solution at any time $t \geq 1$. The optimal choice of p is a function of the optimal Lagrangian multipliers for problem (3). Recall from (4) that the Lagrangian of this optimization problem is given by $\mathcal{L}(x, \nu) := \frac{1}{2} x^T R x + \nu^T (b - Ax)$, where $\nu = (\nu_v)_{v \in V}$ is the vector formed by the Lagrangian multipliers. Define the function Ψ from $\mathbb{R}^{\bar{\mathbb{E}}} \times \mathbb{R}^V$ to $\mathbb{R}^{\bar{\mathbb{E}}} \times \mathbb{R}^V$ as $\Psi(x, \nu) := (\nabla_x \mathcal{L}(x, \nu), Ax - b) = (Rx - A^T \nu, Ax - b)$. Recall that x^* denotes the unique optimal solution of problem (3). As the constraints are linear, by the Lagrange multiplier theorem there exists $\nu^* \in \mathbb{R}^V$ so that $\Psi(x^*, \nu^*) = 0$, i.e.,

$$\begin{cases} R_{ee} x_e^* - (A^T \nu^*)_e = 0 \text{ for } e \in \bar{\mathbb{E}}, \\ \sum_{e \in \bar{\mathbb{E}}} A_{ve} x_e^* = b_v \text{ for } v \in V. \end{cases} \tag{10}$$

Note that the first line in (10) represents Ohm's law, and the second line represents Kirchhoff's law. Given the pair (x^*, ν^*) , we are now ready to define the set of parameters $p^* = \{p_{e \rightarrow v}^*\}$ so that the min-sum algorithm yields the optimal solution at any time step.

Theorem 13 *For each $e \in \vec{E}$, $v \in \partial e$, let $w = \partial e \setminus v$ and define $p_{e \rightarrow v}^* := -A_{we}\nu_w^*$. Then, for any $\tilde{e} \in \vec{E}$, $t \geq 1$, we have $\hat{x}^t(p^*)_{\tilde{e}} = x_{\tilde{e}}^*$.*

Proof Fix $\tilde{e} \in \vec{E}$, $t \geq 1$. Consider the setting in Section 5.1.1, and let \vec{T} be the computation tree of depth t rooted at \tilde{e} . The Lagrangian of the optimization problem (9) corresponding to the choice $p = p^*$ is the function \mathbb{L} from $\mathbb{R}^{\vec{E}} \times \mathbb{R}^{\vec{V}}$ to \mathbb{R} :

$$\mathbb{L}(x, \nu) := \frac{1}{2}x^T R x + p(p^*)^T x + \nu^T (\bar{b} - \bar{A}x),$$

where $\nu = (\nu_v)_{v \in \vec{V}}$. Let Ψ from $\mathbb{R}^{\vec{E}} \times \mathbb{R}^{\vec{V}}$ to $\mathbb{R}^{\vec{E}} \times \mathbb{R}^{\vec{V}}$:

$$\Psi(x, \nu) := (\nabla_x \mathbb{L}(x, \nu), \bar{A}x - \bar{b}) = (Rx + p(p^*) - \bar{A}^T \nu, \bar{A}x - \bar{b}).$$

The Lagrangian multiplier theorem says for the unique minimizer $x^* \equiv x^*(p^*) \in \mathbb{R}^{\vec{E}}$ of (9) there exists $\nu^* \in \mathbb{R}^{\vec{V}}$ so that $\Psi(x^*, \nu^*) = 0$, namely,

$$\begin{cases} R_{ee}x_e^* - (\bar{A}^T \nu^*)_e = 0, e \in \vec{E} \setminus \vec{E}^t, \\ R_{ee}x_e^* + p_{\sigma(e) \rightarrow \sigma(v)}^* - (\bar{A}^T \nu^*)_e = 0, e \in \vec{E}^t, v = \partial e \cap \mathbb{V}^{t-1}, \\ \sum_{e \in \vec{E}} \bar{A}_{ve}x_e^* = \bar{b}_v, v \in \vec{V}. \end{cases}$$

Using (10) it is easy to check that the choice $x_e^* = x_{\sigma(e)}^*$ and $\nu_v^* = \nu_{\sigma(v)}^*$ satisfies the above system of equations. The proof is concluded by Lemma 12. \blacksquare

5.1.3 ERROR CHARACTERIZATION (FLOW PROBLEM)

The next lemma characterizes the sensitivity of the \tilde{e} -th component (root edge of \vec{T}) of the optimal solution of the optimization problem (9) with respect to perturbations of the parameters $p_{e \rightarrow v}$'s, which are supported on \vec{E}^t .

Lemma 14 *Consider the setting in Section 5.1.1. Fix $\tilde{e} = (\tilde{v}, \tilde{w}) \in \vec{E}$, $t \geq 1$, and let \vec{T} be the computation tree of depth t rooted at \tilde{e} , with root edge $\tilde{e} = (\tilde{v}, \tilde{w}) \in \vec{E}$. For any $p = \{p_{e \rightarrow v}\}$, $e \in \vec{E}$ and $v \in \partial e$, we have*

$$\frac{\partial}{\partial p_{e \rightarrow v}} x^*(p)_{\tilde{e}} = \frac{1}{R_{\tilde{e}\tilde{e}}} \sum_{v \in \mathbb{V}^{t-1}} (\bar{\mathbb{L}}_{\tilde{v}v}^{-1} - \bar{\mathbb{L}}_{\tilde{w}v}^{-1})(y_{e \rightarrow v})_v,$$

where $y_{e \rightarrow v} \in \mathbb{R}^{\mathbb{V}^{t-1}}$ is defined, for each $v \in \mathbb{V}^{t-1}$, as $(y_{e \rightarrow v})_v := \frac{A_{ve}}{R_{ee}} \mathbf{1}_{\sigma(v)=v} \mathbf{1}_{\sigma(\partial v \cap \vec{E}^t)=e}$.

Proof By Proposition 23 in Appendix B, the solution of problem (9) reads $x^*(p) = R^{-1} \bar{A}^T \bar{\mathbb{L}}^{-1} \bar{b} + (R^{-1} \bar{A}^T \bar{\mathbb{L}}^{-1} \bar{A} - I) R^{-1} p(p)$, where $I \in \mathbb{R}^{\vec{E} \times \vec{E}}$ is the identity matrix. As $p(p)$ is supported on the edges of \mathbb{T} that are connected to its leaves and \tilde{e} is the root edge, we have

$$\frac{\partial}{\partial p_{e \rightarrow v}} x^*(p)_{\tilde{e}} = \left(R^{-1} \bar{A}^T \bar{\mathbb{L}}^{-1} \bar{A} R^{-1} \frac{\partial}{\partial p_{e \rightarrow v}} p(p) \right)_{\tilde{e}} = \frac{1}{R_{\tilde{e}\tilde{e}}} \sum_{v \in \mathbb{V}^{t-1}} (\bar{\mathbb{L}}_{\tilde{v}v}^{-1} - \bar{\mathbb{L}}_{\tilde{w}v}^{-1})(y_{e \rightarrow v})_v,$$

where $(y_{e \rightarrow v})_{\mathbf{v}} = (\bar{\mathbb{A}}\mathbb{R}^{-1} \frac{\partial}{\partial p_{e \rightarrow v}} \mathbb{P}(p))_{\mathbf{v}} = \frac{A_{ve}}{R_{ee}} \mathbf{1}_{\sigma(\mathbf{v})=v} \mathbf{1}_{\sigma(\partial \mathbf{v} \cap \vec{\mathbb{E}}^t)=e}$ for each $\mathbf{v} \in \mathbb{V}^{t-1}$. \blacksquare

As a consequence of Lemma 14, we immediately have the following characterization of the error committed by Algorithm 5 with initial conditions $\{R_{e \rightarrow v}^0 = R_{ee}\}$, $\{r_{e \rightarrow v}^0 = 0\}$.

Theorem 15 *Consider the setting in Section 5.1.1. Fix $\tilde{e} = (\tilde{v}, \tilde{w}) \in \vec{E}$, $t \geq 1$, and let $\vec{\mathbb{T}}$ be the computation tree of depth t rooted at \tilde{e} , with root edge $\tilde{e} = (\tilde{\mathbf{v}}, \tilde{\mathbf{w}}) \in \vec{\mathbb{E}}$. We have*

$$x_{\tilde{e}}^* - \hat{x}_{\tilde{e}}^t = \mathbb{W}_{\tilde{\mathbf{v}}\tilde{\mathbf{w}}} \sum_{\mathbf{v} \in \mathbb{V}^{t-1}} (\bar{\mathbb{L}}_{\tilde{\mathbf{v}}\mathbf{v}}^{-1} - \bar{\mathbb{L}}_{\tilde{\mathbf{w}}\mathbf{v}}^{-1}) \sum_{\mathbf{w} \in \mathbb{V}^t} \mathbb{W}_{\mathbf{v}\mathbf{w}} \nu_{\sigma(\mathbf{w})}^*.$$

Proof From Theorem 13 and Lemma 14, by the fundamental theorem of calculus and the chain rule of differentiation, we get

$$\begin{aligned} x_{\tilde{e}}^* - \hat{x}_{\tilde{e}}^t &= \hat{x}^t(p^*)_{\tilde{e}} - \hat{x}^t(0)_{\tilde{e}} = \int_0^1 d\theta \frac{d}{d\theta} x^*(\theta p^*)_{\tilde{e}} = \int_0^1 d\theta \sum_{e \in \vec{E}, v \in \partial e} \frac{\partial}{\partial p_{e \rightarrow v}} x^*(\theta p^*)_{\tilde{e}} p_{e \rightarrow v}^* \\ &= \mathbb{W}_{\tilde{\mathbf{v}}\tilde{\mathbf{w}}} \sum_{\mathbf{v} \in \mathbb{V}^{t-1}} (\bar{\mathbb{L}}_{\tilde{\mathbf{v}}\mathbf{v}}^{-1} - \bar{\mathbb{L}}_{\tilde{\mathbf{w}}\mathbf{v}}^{-1}) \sum_{e \in \vec{E}, v \in \partial e} (y_{e \rightarrow v})_{\mathbf{v}} p_{e \rightarrow v}^*. \end{aligned}$$

By Theorem 13 $p_{e \rightarrow v}^* = -A_{we} \nu_w^*$, $w = \partial e \setminus v$, and $\sum_{\substack{e \in \vec{E} \\ v \in \partial e}} (y_{e \rightarrow v})_{\mathbf{v}} p_{e \rightarrow v}^* = \sum_{\mathbf{w} \in \mathbb{V}^t} \mathbb{W}_{\mathbf{v}\mathbf{w}} \nu_{\sigma(\mathbf{w})}^*$. \blacksquare

We now use the error characterization provided by Theorem 15 and the general connection between restricted Laplacians and random walks developed in Appendix A to prove the results in Section 4. Henceforth, for each $\mathbf{v} \in \mathbb{V}$, let $\mathbf{P}_{\mathbf{v}}$ be the law of a time homogeneous Markov chain X_0, X_1, X_2, \dots on \mathbb{V} with transition matrix $\mathbb{P} := \mathbb{D}^{-1}\mathbb{W}$ and initial condition $X_0 = \mathbf{v}$. Analogously, denote by $\mathbf{E}_{\mathbf{v}}$ the expectation with respect to this law. The hitting time to a set $\mathbb{Z} \subseteq \mathbb{V}$ is defined as $T_{\mathbb{Z}} := \inf\{k \geq 0 : X_k \in \mathbb{Z}\}$. We also extend the notation $\mathcal{N}(\mathbf{v}) := \{\mathbf{w} \in \mathbb{V} : \{\mathbf{v}, \mathbf{w}\} \in \mathbb{E}\}$ to denote the vertex-neighborhood of vertex \mathbf{v} in \mathbb{T} .

The next lemma, which builds on the results developed in Appendix A, gives an expression for the restricted Laplacian in terms of hitting probabilities of random walks. This result represents the computational device that we will use repeatedly in what follows.

Lemma 16 *Consider the setting in Section 5.1.1. Fix $\tilde{e} = (\tilde{v}, \tilde{w}) \in \vec{E}$, $t \geq 1$, and let $\vec{\mathbb{T}}$ be the computation tree of depth t rooted at \tilde{e} , with root edge $\tilde{e} = (\tilde{\mathbf{v}}, \tilde{\mathbf{w}}) \in \vec{\mathbb{E}}$. For each $\mathbf{v} \in \bar{\mathbb{V}}$, $\mathbf{w} \in \mathbb{V}^{t-1}$, with $\mathbf{z} = \mathcal{N}(\mathbf{w}) \setminus \mathbb{V}^t$, we have*

$$\bar{\mathbb{L}}_{\mathbf{v}\mathbf{w}}^{-1} = \frac{\mathbf{P}_{\mathbf{v}}(T_{\mathbf{w}} < T_{\mathbb{V}^t})}{[1 - \mathbf{P}_{\mathbf{z}}(T_{\mathbf{w}} < T_{\mathbb{V}^t}) \mathbf{P}_{\mathbf{w}}(X_1 = \mathbf{z})] d_{\mathbf{w}}}. \quad (11)$$

Proof Fix $\mathbf{v} \in \bar{\mathbb{V}}$, $\mathbf{w} \in \mathbb{V}^{t-1}$, and let $\mathbf{z} = \mathcal{N}(\mathbf{w}) \setminus \mathbb{V}^t$. Define the random variable $A := \sum_{k=0}^{T_{\mathbb{V}^t}} \mathbf{1}_{X_k = \mathbf{w}}$. From Proposition 22 we have $\bar{\mathbb{L}}_{\mathbf{v}\mathbf{w}}^{-1} = \frac{1}{d_{\mathbf{w}}} \mathbf{E}_{\mathbf{w}}[A] \mathbf{P}_{\mathbf{v}}(T_{\mathbf{w}} < T_{\mathbb{V}^t})$. By the Markov property, performing a first step analysis for the random walk, we get

$$\mathbf{E}_{\mathbf{w}}[A] = 1 + \sum_{\tilde{\mathbf{z}} \in \mathcal{N}(\mathbf{w})} \mathbf{E}_{\tilde{\mathbf{z}}}[A] \mathbf{P}_{\mathbf{w}}(X_1 = \tilde{\mathbf{z}}) = 1 + \mathbf{E}_{\mathbf{z}}[A] \mathbf{P}_{\mathbf{w}}(X_1 = \mathbf{z}),$$

where we used that $\mathbf{E}_{\tilde{z}}[A] = 0$ for any $\tilde{z} \in \mathbb{V}^t$. Conditioning on the event $\{T_{\mathbf{w}} < T_{\mathbb{V}^t}\}$ and its complementary, using that $\mathbf{E}_{\mathbf{z}}[A|T_{\mathbf{w}} > T_{\mathbb{V}^t}] = 0$, we also get $\mathbf{E}_{\mathbf{z}}[A] = \mathbf{E}_{\mathbf{z}}[A|T_{\mathbf{w}} < T_{\mathbb{V}^t}] \mathbf{P}_{\mathbf{z}}(T_{\mathbf{w}} < T_{\mathbb{V}^t})$. By the strong Markov property we have $\mathbf{E}_{\mathbf{z}}[A|T_{\mathbf{w}} < T_{\mathbb{V}^t}] = \mathbf{E}_{\mathbf{w}}[A]$, so that we obtain the equation $\mathbf{E}_{\mathbf{w}}[A] = 1 + \mathbf{E}_{\mathbf{w}}[A] \mathbf{P}_{\mathbf{z}}(T_{\mathbf{w}} < T_{\mathbb{V}^t}) \mathbf{P}_{\mathbf{w}}(X_1 = \mathbf{z})$, whose solution reads $\mathbf{E}_{\mathbf{w}}[A] = 1/(1 - \mathbf{P}_{\mathbf{z}}(T_{\mathbf{w}} < T_{\mathbb{V}^t}) \mathbf{P}_{\mathbf{w}}(X_1 = \mathbf{z}))$. \blacksquare

We can now present the proofs of the results concerning the flow problem in Lemma 4 and Lemma 6 in Section 4.

Proof [Proof of Lemma 4, flow problem] Let $\vec{G} = (V = \{0, \dots, n-1\}, \vec{E}, W)$ be a weighted directed cycle. Consider the setting in Section 5.1.1. Fix $v \in V$ and assume that $\tilde{e} = (\tilde{v} = \rho(v), \tilde{w} = \rho(v+1)) \in \vec{E}$ (the case $e = (\tilde{w} = \rho(v+1), \tilde{v} = \rho(v))$ yields to the same result with a minus sign in front). Fix $t \geq 2$, and let $\vec{\mathbb{T}}$ be the computation tree of depth t rooted at \tilde{e} , with root edge $\tilde{e} = (\tilde{\mathbf{v}}, \tilde{\mathbf{w}}) \in \vec{\mathbb{E}}$. Label the tree vertices as $\mathbb{V} = \{0, 1, \dots, 2t+1\}$, as in Figure 4. Clearly, $\tilde{\mathbf{v}} = t, \tilde{\mathbf{w}} = t+1$, $\mathbb{V}^{t-1} = \{1, 2t\}$, and $\mathbb{V}^t = \{0, 2t+1\}$.



Figure 4: Cycle $G = (V, E)$ (left) and computation tree $\mathbb{T} = (\mathbb{V}, \mathbb{E})$ of depth t with root edge $\{\tilde{\mathbf{v}} = t, \tilde{\mathbf{w}} = t+1\} \in \mathbb{E}$ (right).

We will apply Lemma 16 to compute the quantities $\bar{\mathbb{L}}_{\tilde{\mathbf{v}}\mathbf{v}}^{-1} - \bar{\mathbb{L}}_{\tilde{\mathbf{w}}\mathbf{v}}^{-1}$, for each $\mathbf{v} \in \mathbb{V}^{t-1}$. We write the computations for $\mathbf{v} = 1$, as the case $\mathbf{v} = 2t$ follows immediately by symmetry. For convenience of notation, we write \mathbb{W}_s to indicate $\mathbb{W}_{s,s+1}$. For each $s \in \mathbb{V}$, define $f_s := \mathbf{P}_s(T_1 < T_{\mathbb{V}^t})$, $p_s := \mathbf{P}_s(X_1 = s+1) = \mathbb{W}_s/\mathbb{d}_s$, $\mathbb{d}_s := \mathbb{W}_{s-1} + \mathbb{W}_s$, and $q_s := \mathbf{P}_s(X_1 = s-1) = 1 - p_s$. By a first step analysis of the random walk and the Markov property, we have

$$f_s = f_{s-1}q_s + f_{s+1}p_s, \quad s \in \{2, \dots, 2t\}, \quad (12)$$

with boundary conditions $f_1 = 1, f_{2t+1} = 0$. The above system of equations corresponds to the classical Gambler Ruin's problem, with final states 1 and $2t+1$, see Figure 5.




Figure 5: Gambler's Ruin.

To solve these equations, using that $1 = p_s + q_s$, we first rewrite them as $f_{s+1} - f_s = \frac{q_s}{p_s}(f_s - f_{s-1})$, for $s \in \{2, \dots, 2t\}$. Using that $q_s/p_s = \mathbb{W}_{s-1}/\mathbb{W}_s$ we get $f_{s+1} - f_s = \frac{\mathbb{W}_1}{\mathbb{W}_s}(f_2 - f_1)$.

Summing these equations side by side, using that $f_1 = 1$, we get $f_s = 1 + (f_2 - 1)\mathbb{W}_1(1/\mathbb{W}_1 + 1/\mathbb{W}_2 + \dots + 1/\mathbb{W}_{s-1})$. Using that $f_{2t+1} = 0$ we can solve for f_2 and finally obtain $f_s = \frac{1/\mathbb{W}_s + 1/\mathbb{W}_{s+1} + \dots + 1/\mathbb{W}_{2t}}{1/\mathbb{W}_1 + 1/\mathbb{W}_2 + \dots + 1/\mathbb{W}_{2t}}$ for any $s \in \{1, \dots, 2t\}$. From Lemma 16, noticing that $z = 2$ in this case, we have $\bar{\mathbb{L}}_{\tilde{\mathbf{v}}1}^{-1} - \bar{\mathbb{L}}_{\tilde{\mathbf{w}}1}^{-1} = (f_t - f_{t+1})/((1 - f_2 p_1) d_1)$, which equals $\bar{\mathbb{L}}_{\tilde{\mathbf{v}}1}^{-1} - \bar{\mathbb{L}}_{\tilde{\mathbf{w}}1}^{-1} = \frac{1/\mathbb{W}_t}{\mathbb{W}_0(1/\mathbb{W}_0 + \dots + 1/\mathbb{W}_{2t})}$. By symmetry, we have $\bar{\mathbb{L}}_{\tilde{\mathbf{v}},2t}^{-1} - \bar{\mathbb{L}}_{\tilde{\mathbf{w}},2t}^{-1} = \frac{-1/\mathbb{W}_t}{\mathbb{W}_{2t}(1/\mathbb{W}_0 + \dots + 1/\mathbb{W}_{2t})}$. By Theorem 15, as $\mathbb{W}_s = W_{\sigma(s), \sigma(s+1)}$, we finally get $x_{\tilde{\mathbf{e}}}^* - \hat{x}_{\tilde{\mathbf{e}}}^t = \frac{\nu_{\sigma(0)}^* - \nu_{\sigma(2t+1)}^*}{1/W_{\sigma(0)\sigma(1)} + \dots + 1/W_{\sigma(2t)\sigma(2t+1)}}$. The proof follows as the map $\sigma : \mathbb{V} \rightarrow V$ reads $\sigma(s) = \rho(s - t + v)$. ■

Proof [Proof of Lemma 6, flow problem] Let $\vec{G} = (V, \vec{E}, W)$ be a connected d -regular graph, with $d \geq 3$, where each edge has the same weight ω . Consider the setting in Section 5.1.1. Fix $\tilde{\mathbf{e}} = (\tilde{\mathbf{v}}, \tilde{\mathbf{w}}) \in \vec{E}$, $t \geq 3$, and let $\vec{\mathbb{T}}$ be the computation tree of depth t rooted at $\tilde{\mathbf{e}}$, with root edge $\tilde{\mathbf{e}} = (\tilde{\mathbf{v}}, \tilde{\mathbf{w}}) \in \vec{\mathbb{E}}$. Fix $\mathbf{v} \in \mathbb{V}^{t-1}$, and assume that \mathbf{v} is contained in the subtree with root vertex $\tilde{\mathbf{v}}$. For what follows, it is convenient to label some of the vertices in \mathbb{V} . Let $\mathbf{v} = 0, 1 = \mathcal{N}(\mathbf{v}) \setminus \mathbb{V}^t, 2 = \mathcal{N}(1) \setminus \mathbb{V}^{t-1}, \dots, k+1 = \mathcal{N}(k) \setminus \mathbb{V}^{t-k}$, until we reach $\tilde{\mathbf{v}} = t-1$. Label $\tilde{\mathbf{w}} = t$. See Figure 6. We will apply Lemma 16 to compute the quantity $\bar{\mathbb{L}}_{\tilde{\mathbf{v}}\mathbf{v}}^{-1} - \bar{\mathbb{L}}_{\tilde{\mathbf{w}}\mathbf{v}}^{-1}$. To this




Figure 6: Computation tree $\mathbb{T} = (\mathbb{V}, \mathbb{E})$ of depth $t = 3$ associated to any 3-regular graph.

end, we need to compute the hitting probabilities $f_{\mathbf{w}} := \mathbf{P}_{\mathbf{w}}(T_{\mathbf{v}} < T_{\mathbb{V}^t})$ for $\mathbf{w} \in \{1, t-1, t\}$. To do this, it is convenient to first reduce the network topology by invoking the equivalence between random walks and electrical circuits. See Peter G. Doyle (1984). The electrical circuit associated to the random walk X on \mathbb{T} is constructed by assigning to each edge $\mathbf{e} = \{\mathbf{v}, \mathbf{w}\} \in \mathbb{T}$ a conductor ω , which corresponds to a resistor $1/\omega$. Then, the quantities $f_{\mathbf{w}}$, $\mathbf{w} \in \mathbb{V}$, correspond to the voltages induced on $\mathbb{V} \setminus \{\mathbf{v} \cup \mathbb{V}^t\}$ when a voltage of intensity 1 is applied to \mathbf{v} ($f_{\mathbf{v}} = 1$), and a voltage of intensity 0 is applied to the sites in \mathbb{V}^t ($f_{\mathbf{w}} = 0$, $\mathbf{w} \in \mathbb{V}^t$). By repeatedly applying the standard network reduction operations for conductors in parallel and in series, it is easy to verify that, as long as the quantities f_1, f_2, \dots, f_t are concerned, the original network is equivalent to the one represented in Figure 7.




Figure 7: Reduced network.

In the reduced network, the conductance between the pairs of nodes s and $s + 1$, for $s \in \{0, \dots, t - 1\}$, remain ω as in the unreduced network, i.e., the computation tree. The effect of all the nodes different from $0, \dots, t$ in the original graph is summarized in the reduced network by node $t + 1$. For any $s \in \{0, \dots, t\}$, between node s and node $t + 1$ there is a new conductor \mathbb{C}_s , which is obtained by reducing the subtree having root vertex s in the original graph. It is easy to verify that

$$\mathbb{C}_0 = \omega(d - 1), \quad \mathbb{C}_s = \omega \frac{(d - 2)^2}{d - 1} (1 + h_{s+1}), s \in \{1, \dots, t - 1\}, \quad \mathbb{C}_t = \omega(d - 2)(1 + h_t),$$

with $h_s := \frac{1}{(d-1)^{s-1}}$. Let $p_0 := \frac{\omega}{\omega + \mathbb{C}_0}$, $q_0 := 1 - p_0$, $p_s = q_s := \frac{\omega}{2\omega + \mathbb{C}_s}$ for $s \in \{1, \dots, t - 1\}$, $p_t := \frac{\mathbb{C}_t}{\omega + \mathbb{C}_t}$, $q_t := 1 - p_t$. By the Markov property we have

$$f_s = f_{s-1}q_s + f_{s+1}p_s, \quad s \in \{1, \dots, t\}, \quad (13)$$

with $f_0 = 1, f_{t+1} = 0$. This system of equations corresponds to a modification of the classical Gambler Ruin's problem, where states 0 and $t + 1$ are the final states, and each intermediate state s has a probability $r_s := 1 - p_s - q_s$ to reach state $t + 1$. See Figure 8.




Figure 8: Gambler's Ruin on the reduced network.

Without loss of generality, we set $r_t = 0$. We can easily write the solution of this system of equations in a compact form. Define $\alpha_1 = \alpha_2 := 1, \alpha_s := \alpha_{s-1} - p_{s-2}q_{s-1}\alpha_{s-2}$ for $s \geq 3$. Define $\beta_1 = \beta_2 := 1, \beta_s := \beta_{s-1} - p_{s-1}q_s\beta_{s-2}$ for $s \geq 3$. Then, as a function of f_1 , we can write the solution of (13) as $f_s = \frac{1}{p_1 \cdots p_{s-1}} (\alpha_s f_1 - \beta_{s-1} q_1)$, for $s \in \{2, \dots, t + 1\}$. In fact, it is easy to verify that this expression satisfies system (13). Using $f_{t+1} = 0$ we find $f_1 = q_1 \beta_t / \alpha_{t+1}$ and $f_s = \frac{p_0 q_1}{p_0 \cdots p_{s-1}} (\frac{\alpha_s \beta_t}{\alpha_{t+1}} - \beta_{s-1})$ for $s \in \{1, \dots, t + 1\}$, with $\beta_0 := 0$. Lemma 16 yields the following, using that $\mathbb{C}_v = \omega d$,

$$\begin{aligned} \bar{\mathbb{L}}_{\tilde{v}v}^{-1} &= \frac{f_{t-1}}{(1 - f_1 p_0) \mathbb{C}_v} = \frac{1}{\omega d} \frac{q_1}{p_1 \cdots p_{t-2} (\alpha_{t+1} - p_0 q_1 \beta_t)} (\alpha_{t-1} \beta_t - \alpha_{t+1} \beta_{t-2}), \\ \bar{\mathbb{L}}_{\tilde{v}v}^{-1} &= \frac{f_t}{(1 - f_1 p_0) \mathbb{C}_v} = \frac{1}{\omega d} \frac{q_1}{p_1 \cdots p_{t-2} (\alpha_{t+1} - p_0 q_1 \beta_t)} \frac{1}{p_{t-1}} (\alpha_t \beta_t - \alpha_{t+1} \beta_{t-1}). \end{aligned}$$

Using the recursive formulas for α_{t+1} and β_t , and the fact that $p_t + q_t = 1$, we get

$$\bar{\mathbb{L}}_{\tilde{v}v}^{-1} - \bar{\mathbb{L}}_{\tilde{v}v}^{-1} = \frac{1}{\omega d} \frac{q_1}{p_1 \cdots p_{t-2} (\alpha_{t+1} - p_0 q_1 \beta_t)} p_t (\alpha_{t-1} \beta_{t-1} - \alpha_t \beta_{t-2}),$$

which, as $\alpha_s \beta_s - \alpha_{s+1} \beta_{s-1} = p_{s-1} q_s (\alpha_{s-1} \beta_{s-1} - \alpha_s \beta_{s-2})$, becomes

$$\bar{\mathbb{L}}_{\tilde{v}v}^{-1} - \bar{\mathbb{L}}_{\tilde{w}v}^{-1} = \frac{p_t}{q_t} \bar{\mathbb{L}}_{\tilde{w}v}^{-1} = \frac{1}{\omega d(d-1)^{t-1}} \frac{1}{\xi_{t-1}}, \quad (14)$$

where $\xi_{t-1} := \frac{1}{(d-1)^{t-1}} \frac{\alpha_{t+1} - p_0 q_1 \beta_t}{q_1 \cdots q_{t-1} p_t}$ is defined by $\xi_s = \frac{p_s}{q_s p_{s+1} (d-1)} \xi_{s-1} - \frac{p_{s-1} p_s q_{s+1}}{q_{s-1} q_s p_{s+1} (d-1)^2} \xi_{s-2}$ for $s \in \{2, \dots, t-1\}$, with $\xi_0 := \frac{p_0 q_1}{p_1}$ and $\xi_1 = \frac{1}{d-1} \frac{1 - p_0 q_1 - p_1 q_2}{q_1 p_2}$. Plugging in the definitions of the transition probabilities, recalling the quantities defined in Proposition 5, it is immediate to verify that $\delta_s \equiv \xi_s$ for $s \in \{0, \dots, t-2\}$ and $c_{d,t} \equiv \xi_{t-1}$. By symmetry, it is clear that $\bar{\mathbb{L}}_{\tilde{v}v}^{-1} - \bar{\mathbb{L}}_{\tilde{w}v}^{-1} = \frac{1}{\omega d(d-1)^{t-1}} \frac{1}{c_{d,t}}$ holds for *any* $v \in \mathbb{V}^{t-1}$ that is contained in the subtree of \mathbb{T} with root vertex \tilde{v} . On the other hand, again by symmetry, for any $v \in \mathbb{V}^{t-1}$ contained in the subtree of \mathbb{T} with root vertex \tilde{w} we have $\bar{\mathbb{L}}_{\tilde{v}v}^{-1} - \bar{\mathbb{L}}_{\tilde{w}v}^{-1} = -\frac{1}{\omega d(d-1)^{t-1}} \frac{1}{c_{d,t}}$. Theorem 15 yields

$$x_{\tilde{e}}^* - \hat{x}_{\tilde{e}}^t = \omega^2 \sum_{v \in \mathbb{V}^{t-1}} (\bar{\mathbb{L}}_{\tilde{v}v}^{-1} - \bar{\mathbb{L}}_{\tilde{w}v}^{-1}) \sum_{w \in \mathbb{V}^t} \nu_{\sigma(w)}^* = \frac{\omega}{d(d-1)^{t-1}} \frac{1}{c_{d,t}} \sum_{v \in V} \nu_v^* (\#_{\tilde{v} \rightarrow v, \tilde{w}}^t - \#_{\tilde{w} \rightarrow v, \tilde{v}}^t),$$

where $\#_{\tilde{v} \rightarrow v, \tilde{w}}^t$ is the number of non-backtracking paths in the original undirected graph $G = (V, E)$ that connect node \tilde{v} to node v in t time steps, where the first step is different from \tilde{w} . The proof is easily concluded noticing that the total number of non-backtracking paths after t steps is $d(d-1)^{t-1}$, so that $\mathbf{P}_{\tilde{v}}(Y_1 \neq \tilde{w}, Y_t = v) = \#_{\tilde{v} \rightarrow v, \tilde{w}}^t / (d(d-1)^{t-1})$. ■

5.2 Proof for the voltage problem, Algorithm 5

We now give an overview of how the proofs for Lemma 4 and Lemma 6 work for the results concerning the voltage problem. The main framework for these proofs is the same as the one used in the previous section to investigate the behavior of the min-sum algorithm applied to the flow problem. For this reason, as outlined at the beginning of Section 5, here we only discuss in details the parts of the proofs that present major differences compared to the flow problem. The reader should not confuse the new quantities that we are about to introduce with the analogous quantities that were previously defined in Section 5.1. Even if we use much of the same notation, the new quantities are now supported on a different graph, i.e., on the computation tree for the voltage problem.

Henceforth we consider Algorithm 3 when the initial messages are parametrized by $\{W_{e \rightarrow v}^0 = W_{vw}, e = \{v, w\}\}, \{h_{e \rightarrow v}^0 = p_{e \rightarrow v}\}$, for a certain set of real numbers $p = \{p_{e \rightarrow v}\}$. Also in this case, as far as the analysis is concerned, it is convenient to consider the form of Algorithm 2. For a given choice of perturbation parameters $p = \{p_{e \rightarrow v}\}$, define the initial messages as

$$\mu_{e \rightarrow v}^0(\cdot, p) : z \in \mathbb{R} \longrightarrow \mu_{e \rightarrow v}^0(z, p) = \frac{1}{2} W_{vw} z^2 + p_{e \rightarrow v} z,$$

where $e = \{v, w\}$, and denote by $\{\mu_{e \rightarrow v}^s(\cdot, p)\}$, $s \geq 1$, the corresponding sequence of messages generated by the min-sum algorithm. Denote the estimates at time t as $\hat{\nu}_v^t = \arg \min_{z \in \mathbb{R}} \mu_v^t(z, p)$, where $\mu_v^t(\cdot, p)$ is the corresponding belief function.

5.2.1 COMPUTATION TREE (VOLTAGE PROBLEM)

The computation tree for the voltage problem differs from the one in the flow problem due to the fact that the min-sum algorithms for the voltage problem propagates messages that are supported on vertices, not on edges. For this reason, the computation tree that we will now define is rooted at a vertex, not at an edge as the one defined in Section 5.1.1.

Given a vertex $\tilde{v} \in V$, the computation tree rooted at \tilde{v} of depth t supports the optimization problem that is obtained by unfolding the computations involved in the min-sum estimate $\hat{\nu}^t(p)_{\tilde{v}}$. Formally, the computation tree is an undirected tree $\mathbb{T} = (\mathbb{V}, \mathbb{E})$ where each vertex in \mathbb{V} is mapped to a vertex in V through a map $\sigma : \mathbb{V} \rightarrow V$ that preserves the edge structure, namely, if $\mathfrak{e} = \{\mathfrak{v}, \mathfrak{w}\} \in \mathbb{E}$ then $\sigma(\mathfrak{e}) := \{\sigma(\mathfrak{v}), \sigma(\mathfrak{w})\} \in E$. The tree \mathbb{T} is defined iteratively. Initially, at level -1 , the tree is made by a single root vertex $\tilde{\mathfrak{v}} \in \mathbb{V}$ corresponding to \tilde{v} , i.e., $\sigma(\tilde{\mathfrak{v}}) = \tilde{v}$. Consecutively, for any $z \in \mathcal{N}(\tilde{v})$, a vertex \mathfrak{z} with $\sigma(\mathfrak{z}) = z$ and an edge $\{\tilde{\mathfrak{v}}, \mathfrak{z}\}$ are added to the level 0 of the tree. For the remaining $t - 1$ levels, the following procedure is repeated. The leaf in the tree are examined. Given a leaf \mathfrak{v} with $\sigma(\mathfrak{v}) = v$ that is connected to a vertex \mathfrak{w} with $\sigma(\mathfrak{w}) = w$, for any $z \in \mathcal{N}(v) \setminus w$, a vertex \mathfrak{z} with $\sigma(\mathfrak{z}) = z$ and an edge $\{\mathfrak{v}, \mathfrak{z}\}$ are added to the next level of the tree. Figure 9 gives an example. We denote the set of vertices and edges in the k -th level of the tree respectively by $\mathbb{V}^k \subset \mathbb{V}$ and $\mathbb{E}^k \subset \mathbb{E}$. In what follows we also extend the usual neighborhood notation to vertices and edges in the graph \mathbb{T} , namely, $\partial \mathfrak{v}$ is the set of edges in \mathbb{T} that are connected to node \mathfrak{v} , and $\partial \mathfrak{e}$ is the set of vertices in \mathbb{T} that are connected to edge \mathfrak{e} .




Figure 9: Graph $G = (V = \{1, 2, 3, 4\}, E)$ (left) and computation tree $\mathbb{T} = (\mathbb{V}, \mathbb{E})$ of depth $t = 2$ with root edge corresponding to $1 \in V$ (right). For the sake of simplicity, vertices in \mathbb{T} are labeled by the corresponding vertices in G .

The t -th (last) level of the tree is of particular relevance as it supports the initial conditions of the min-sum algorithm, as we are now about to describe. Let $\mathbb{W} \in \mathbb{R}^{\mathbb{V} \times \mathbb{V}}$ be the diagonal matrix defined by $\mathbb{W}_{\mathfrak{vw}} := W_{\sigma(\mathfrak{v})\sigma(\mathfrak{w})}$ if $\{\mathfrak{v}, \mathfrak{w}\} \in \mathbb{E}$ and $\mathbb{W}_{\mathfrak{vw}} := 0$ otherwise. Let $\bar{\mathbb{V}} := \mathbb{V} \setminus \mathbb{V}^t$. Define the weighted degree of a vertex \mathfrak{v} by $\mathfrak{d}_{\mathfrak{v}} := \sum_{\mathfrak{w} \in \mathbb{V}} \mathbb{W}_{\mathfrak{vw}}$, and let $\mathbb{D} \in \mathbb{R}^{\mathbb{V} \times \mathbb{V}}$ be the diagonal matrix defined by $\mathbb{D}_{\mathfrak{vv}} := \mathfrak{d}_{\mathfrak{v}}$. The Laplacian \mathbb{L} of the graph \mathbb{T} is the matrix defined as $\mathbb{L} := \mathbb{D} - \mathbb{W}$. Let $\bar{\mathbb{L}} \in \mathbb{R}^{\bar{\mathbb{V}} \times \bar{\mathbb{V}}}$ denote the submatrix of \mathbb{L} that corresponds to the rows and columns associated to $\bar{\mathbb{V}}$. For a given choice of the perturbation parameters $p = \{p_{e \rightarrow v}\}$, define $\bar{\mathfrak{p}}(p) \in \mathbb{R}^{\bar{\mathbb{V}}}$ by

$$\bar{\mathfrak{p}}(p)_{\mathfrak{v}} := \begin{cases} \sum_{\mathfrak{f} \in \partial \mathfrak{v} \setminus \mathfrak{e}} p_{\sigma(\mathfrak{f}) \rightarrow \sigma(\mathfrak{v})} & \text{if } \mathfrak{v} \in \mathbb{V}^{t-1}, \mathfrak{e} = \partial \mathfrak{v} \setminus \mathbb{E}^t, \\ 0 & \text{otherwise,} \end{cases} \quad (15)$$

and let $\bar{\mathbb{b}} \in \mathbb{R}^{\bar{\mathbb{V}}}$ defined by $\bar{\mathbb{b}}_{\mathfrak{v}} := b_{\sigma(\mathfrak{v})}$ for each $\mathfrak{v} \in \bar{\mathbb{V}}$.

Consider the following problem supported on the computation tree \mathbb{T} , over $\bar{\nu} \in \mathbb{R}^{\bar{\mathbb{V}}}$:

$$\text{minimize} \quad \frac{1}{2} \bar{\nu}^T \bar{\mathbb{L}} \bar{\nu} + (\bar{\rho}(p) - \bar{\mathbb{b}})^T \bar{\nu}. \quad (16)$$

It is a known fact that while the Laplacian \mathbb{L} is positive semi-definite, the restricted Laplacian $\bar{\mathbb{L}}$ is positive definite. Hence, for any choice of the perturbation p the objective function in problem (16) is strictly convex and it admits a unique solution, which we denote by $\bar{\nu}^*(p)$. The relationship between the min-sum algorithm and the computation tree is made explicit by the following lemma, which can be easily established by inductively examining the operations of Algorithm 2.

Lemma 17 *Given the setting of this section, $\hat{\nu}^t(p)_{\tilde{v}} = \bar{\nu}^*(p)_{\tilde{v}}$.*

5.2.2 FIX POINT (VOLTAGE PROBLEM)

The choice $p = 0$ for the linear perturbation parameters yields the original (unperturbed) algorithm. The next theorem below shows that for a particular choice of p the min-sum algorithm yields the optimal solution at any time $t \geq 1$. The optimal choice of p is a function of the optimal solution ν^* for problem (2). Recall that, by definition, ν^* is the unique vector that satisfy $L\nu^* = b$ and $1^T \nu^* = 0$. Hence, component-wise, for each $v \in V$,

$$d_v \nu_v^* - \sum_{w \in \partial v} W_{vw} \nu_w^* - b_v = 0, \quad (17)$$

where recall that $d_v := \sum_{w \in \partial v} W_{vw}$. Given ν^* , we are now ready to define the set of parameters $p^* = \{p_{e \rightarrow v}^*\}$ so that the min-sum algorithm yields the optimal solution at any time step.

Theorem 18 *For each $e \in E$, $v \in \partial e$, let $w = \partial e \setminus v$ and define*

$$p_{e \rightarrow v}^* := -W_{vw} \nu_w^*.$$

Then, for any $\tilde{v} \in V$, $t \geq 1$, we have $\hat{\nu}^t(p^)_{\tilde{v}} = \nu_{\tilde{v}}^*$.*

Proof Fix $\tilde{v} \in V$, $t \geq 1$. Consider the setting in Section 5.2.1, and let \mathbb{T} be the computation tree of depth t rooted at \tilde{v} . From the first order optimality conditions for problem (16) corresponding to the choice $p = p^*$, we know that $\bar{\nu}^* \equiv \bar{\nu}^*(p^*)$ is the unique vector so that

$$\begin{cases} d_{\mathbb{v}} \bar{\nu}_{\mathbb{v}}^* - \sum_{\mathbb{w} \in \partial \mathbb{v}} W_{\mathbb{v}\mathbb{w}} \bar{\nu}_{\mathbb{w}}^* - \bar{\mathbb{b}}_{\mathbb{v}} = 0, \mathbb{v} \in \bar{\mathbb{V}} \setminus \mathbb{V}^{t-1}, \\ d_{\mathbb{v}} \bar{\nu}_{\mathbb{v}}^* - W_{\mathbb{v}\mathbb{w}} \bar{\nu}_{\mathbb{w}}^* + \bar{\rho}(p^*)_{\mathbb{v}} - \bar{\mathbb{b}}_{\mathbb{v}} = 0, \mathbb{v} \in \mathbb{V}^{t-1}, \{\mathbb{v}, \mathbb{w}\} = \partial \mathbb{v} \setminus \mathbb{E}^t. \end{cases}$$

Using (17) it is easy to check that the choice $\bar{\nu}_{\mathbb{v}}^* = \nu_{\sigma(\mathbb{v})}^*$ satisfies the above system of equations. The proof is concluded by Lemma 17. \blacksquare

5.2.3 ERROR CHARACTERIZATION (VOLTAGE PROBLEM)

The next lemma characterizes the sensitivity of the \tilde{v} -th component (root vertex of \mathbb{T}) of the optimal solution of the optimization problem (16) with respect to perturbations of the parameters $p_{e \rightarrow v}$'s, which are supported on \mathbb{V}^{t-1} .

Lemma 19 *Consider the setting in Section 5.2.1. Fix $\tilde{v} \in V$, $t \geq 1$, and let \mathbb{T} be the computation tree of depth t rooted at \tilde{v} , with root edge \tilde{v} . For any $p = \{p_{e \rightarrow v}\}$, $e \in E$ and $v \in \partial e$, we have*

$$\frac{\partial}{\partial p_{e \rightarrow v}} \bar{v}^*(p)_{\tilde{v}} = - \sum_{\mathbf{v} \in \mathbb{V}^{t-1}} \bar{\mathbb{L}}_{\tilde{v}\mathbf{v}}^{-1} (\mathbf{y}_{e \rightarrow v})_{\mathbf{v}},$$

where $\mathbf{y}_{e \rightarrow v} \in \mathbb{R}^{\mathbb{V}^{t-1}}$ is defined, for each $\mathbf{v} \in \mathbb{V}^{t-1}$, as $(\mathbf{y}_{e \rightarrow v})_{\mathbf{v}} := \mathbf{1}_{\sigma(\mathbf{v})=v} \mathbf{1}_{\sigma(\partial \mathbf{v} \cap \tilde{\mathbb{E}}^t)=e}$.

Proof From the first order optimality conditions, we know that the unique solution of problem (16) reads $\bar{v}^*(p) = \bar{\mathbb{L}}^{-1}(\bar{\mathbb{b}} - \bar{\mathbb{p}}(p))$ so that $\frac{\partial}{\partial p_{e \rightarrow v}} \bar{v}^*(p) = -\bar{\mathbb{L}}^{-1} \frac{\partial}{\partial p_{e \rightarrow v}} \bar{\mathbb{p}}(p)$. Using (15) it is easy to check that $(\mathbf{y}_{e \rightarrow v})_{\mathbf{v}} = (\frac{\partial}{\partial p_{e \rightarrow v}} \bar{\mathbb{p}}(p))_{\mathbf{v}}$. \blacksquare

As a consequence of Lemma 19, we immediately have the following characterization of the error committed by Algorithm 3 with $\{W_{e \rightarrow v}^0 = W_{vw}, e = \{v, w\}\}$, $\{h_{e \rightarrow v}^0 = 0\}$.

Theorem 20 *Consider the setting in Section 5.2.1. Fix $\tilde{v} \in V$, $t \geq 1$, and let \mathbb{T} be the computation tree of depth t rooted at \tilde{v} , with root vertex $\tilde{v} \in \mathbb{V}$. We have*

$$\nu_{\tilde{v}}^* - \hat{\nu}_{\tilde{v}}^t = \sum_{\mathbf{v} \in \mathbb{V}^{t-1}} \bar{\mathbb{L}}_{\tilde{v}\mathbf{v}}^{-1} \sum_{\mathbf{w} \in \mathbb{V}^t} \mathbb{W}_{\mathbf{vw}} \nu_{\sigma(\mathbf{w})}^*.$$

Proof From Theorem 18 and Lemma 19, by the fundamental theorem of calculus and the chain rule of differentiation,

$$\begin{aligned} \nu_{\tilde{v}}^* - \hat{\nu}_{\tilde{v}}^t &= \hat{\nu}_{\tilde{v}}^t(p^*)_{\tilde{v}} - \hat{\nu}_{\tilde{v}}^t(0)_{\tilde{v}} = \int_0^1 d\theta \frac{d}{d\theta} \bar{v}^*(\theta p^*)_{\tilde{v}} = \int_0^1 d\theta \sum_{e \in E, v \in \partial e} \frac{\partial}{\partial p_{e \rightarrow v}} \bar{v}^*(\theta p^*)_{\tilde{v}} p_{e \rightarrow v}^* \\ &= - \sum_{\mathbf{v} \in \mathbb{V}^{t-1}} \bar{\mathbb{L}}_{\tilde{v}\mathbf{v}}^{-1} \sum_{e \in E, v \in \partial e} (\mathbf{y}_{e \rightarrow v})_{\mathbf{v}} p_{e \rightarrow v}^*. \end{aligned}$$

From Theorem 18 we get $p_{e \rightarrow v}^* := -W_{vw} \nu_w^*$ for $e = \{v, w\}$, and it is easy to verify that $\sum_{e \in E, v \in \partial e} (\mathbf{y}_{e \rightarrow v})_{\mathbf{v}} p_{e \rightarrow v}^* = - \sum_{\mathbf{w} \in \mathbb{V}^t} \mathbb{W}_{\mathbf{vw}} \nu_{\sigma(\mathbf{w})}^*$. \blacksquare

Lemma 16 for the flow problem holds also in the present case, upon referring the notation used (i.e., $\mathbf{P}_{\mathbf{v}}$, $\mathbf{E}_{\mathbf{v}}$, etc.) to the Markov chain defined on the computation tree for the voltage problem. We can now present the outline of the proofs for the results concerting the voltage problem in Lemma 4 and Lemma 6 in Section 4. Here we assume that the reader is already familiar with the proofs for the flow case.

Proof [Proof of Lemma 4, voltage problem] Let $G = (V = \{0, \dots, n-1\}, E, W)$ be a weighted cycle. Consider the setting in Section 5.2.1. Fix $v \in V$, $t \geq 2$, and let \mathbb{T} be

the computation tree of depth t rooted at \tilde{v} , with root vertex \tilde{v} , i.e., $\sigma(\tilde{v}) = \tilde{v}$. Label the vertices of the computation tree as $\mathbb{V} = \{0, 1, \dots, 2t+2\}$, as in Figure 10. Clearly, $\tilde{v} = t+1$, $\mathbb{V}^{t-1} = \{1, 2t+1\}$, and $\mathbb{V}^t = \{0, 2t+2\}$.



Figure 10: Ring $G = (V, E)$ (left) and computation tree $\mathbb{T} = (\mathbb{V}, \mathbb{E})$ of depth t with root vertex $\tilde{v} = t+1 \in \mathbb{V}$ (right).

Given this set up for the computation tree, the remaining of the proof follows exactly the steps of the proof in the flow case. Using the analogue of Lemma 16 for the voltage case, we can compute the quantities $\bar{\mathbb{L}}_{\tilde{v}\mathbf{v}}^{-1}$, for each $\mathbf{v} \in \mathbb{V}^{t-1}$, by solving a system of equations like (12). If we write \mathbb{W}_s to indicate $\mathbb{W}_{s,s+1}$, we find that $\bar{\mathbb{L}}_{\tilde{v}1}^{-1} = \frac{1/\mathbb{W}_{t+1} + \dots + 1/\mathbb{W}_{2t+1}}{\mathbb{W}_0(1/\mathbb{W}_0 + \dots + 1/\mathbb{W}_{2t+1})}$, and, by symmetry, $\bar{\mathbb{L}}_{\tilde{v},2t+1}^{-1} = \frac{1/\mathbb{W}_0 + \dots + 1/\mathbb{W}_t}{\mathbb{W}_{2t+1}(1/\mathbb{W}_0 + \dots + 1/\mathbb{W}_{2t+1})}$. By Theorem 20 we get $\nu_v^* - \hat{\nu}_v^t = \sum_{\mathbf{v} \in \mathbb{V}^{t-1}} \bar{\mathbb{L}}_{\tilde{v}\mathbf{v}}^{-1} \sum_{\mathbf{w} \in \mathbb{V}^t} \mathbb{W}_{\mathbf{v}\mathbf{w}} \nu_{\sigma(\mathbf{w})}^* = \frac{1/\mathbb{W}_{t+1} + \dots + 1/\mathbb{W}_{2t+1}}{1/\mathbb{W}_0 + \dots + 1/\mathbb{W}_{2t+1}} \nu_{\sigma(0)}^* + \frac{1/\mathbb{W}_0 + \dots + 1/\mathbb{W}_t}{1/\mathbb{W}_0 + \dots + 1/\mathbb{W}_{2t+1}} \nu_{\sigma(2t+2)}^*$, and the proof follows using that $\mathbb{W}_s = W_{\sigma(s), \sigma(s+1)}$ and noticing that the mapping $\sigma : \mathbb{V} \rightarrow V$ now reads $\sigma(v) = \rho(s - t - 1 + \tilde{v})$. \blacksquare

Proof [Proof of Lemma 6, voltage problem] Let $G = (V, E, W)$ be a connected d -regular graph, with $d \geq 3$, where each edge has the same weight ω . Consider the setting in Section 5.2.1. Fix $\tilde{w} \in V$, $t \geq 3$, and let \mathbb{T} be the computation tree of depth t rooted at \tilde{w} , with root vertex $\tilde{w} \in \mathbb{V}$. Fix $\mathbf{v} \in \mathbb{V}^{t-1}$. For what follows, it is convenient to label some of the vertices in \mathbb{V} . Let $\mathbf{v} = 0, 1 = \mathcal{N}(\mathbf{v}) \setminus \mathbb{V}^t, 2 = \mathcal{N}(1) \setminus \mathbb{V}^{t-1}, \dots, k+1 = \mathcal{N}(k) \setminus \mathbb{V}^{t-k}$, until we reach $\tilde{w} = t$. See Figure 11. Given this set up for the computation tree, the remaining




Figure 11: Computation tree $\mathbb{T} = (\mathbb{V}, \mathbb{E})$ of depth $t = 2$ associated to any given d -regular graph, with $d = 3$.

of the proof follows exactly the steps of the proof in the flow case. Using the analogue of Lemma 16 for the voltage case, we can compute $\bar{\mathbb{L}}_{\tilde{w}\mathbf{v}}^{-1}$ by solving a system of equations of the form (13) supported on the a reduced network that is exactly like the one in Figure 7. The

only difference is that now $\mathbb{C}_t = \omega(d-2)(1+h_{t+1})$, while all other effective conductances \mathbb{C}_s , for $s \in \{0, \dots, t-1\}$, read exactly as for the flow problem. Rearranging (14) we get

$$\bar{\mathbb{L}}_{\tilde{\mathbf{w}}\tilde{\mathbf{v}}}^{-1} = \frac{1}{\omega d(d-1)^{t-1}} \frac{q_t}{p_t} \frac{1}{\xi_{t-1}},$$

where $\xi_{t-1} := \frac{1}{(d-1)^{t-1}} \frac{\alpha_{t+1}-p_0q_1\beta_t}{q_1 \dots q_{t-1}p_t}$ is defined as in the flow case by $\xi_s = \frac{p_s}{q_s p_{s+1}(d-1)} \xi_{s-1} - \frac{p_{s-1}p_s q_{s+1}}{q_{s-1}q_s p_{s+1}(d-1)^2} \xi_{s-2}$ for $s \in \{2, \dots, t-1\}$, with $\xi_0 := \frac{p_0q_1}{p_1}$ and $\xi_1 = \frac{1}{d-1} \frac{1-p_0q_1-p_1q_2}{q_1p_2}$. Plugging in the definitions of the transition probabilities, noticing that p_t and q_t differ from the analogous quantities defined in the flow case due to the new definition of \mathbb{C}_t , recalling Proposition 5, it is immediate to verify that $\delta_s \equiv \xi_s$ for $s \in \{0, \dots, t-2\}$ and $\frac{d-1}{(d-2)(1+h_{t+1})} b_{d,t} \equiv \xi_{t-1}$, so that $\bar{\mathbb{L}}_{\tilde{\mathbf{w}}\tilde{\mathbf{v}}}^{-1} = \frac{1}{\omega d(d-1)^t} \frac{1}{b_{d,t}}$. Theorem 20 yields

$$\nu_{\tilde{w}}^* - \hat{\nu}_{\tilde{w}}^t = \sum_{\mathbf{v} \in \mathbb{V}^{t-1}} \bar{\mathbb{L}}_{\tilde{\mathbf{w}}\tilde{\mathbf{v}}}^{-1} \sum_{\mathbf{w} \in \mathbb{V}^t} \mathbb{W}_{\tilde{\mathbf{w}}\mathbf{w}} \nu_{\sigma(\mathbf{w})}^* = \frac{1}{d(d-1)^t} \frac{1}{b_{d,t}} \sum_{v \in V} \nu_v^* \#_{\tilde{w} \rightarrow v}^{t+1}.$$

where $\#_{\tilde{w} \rightarrow v}^{t+1}$ is the number of non-backtracking paths in the original undirected graph $G = (V, E)$ that connect node \tilde{w} to node v in $t+1$ time steps. The proof is easily concluded noticing that the total number of non-backtracking paths after $t+1$ steps is $d(d-1)^t$, so that $\mathbf{P}_{\tilde{w}}(Y_{t+1} = v) = \#_{\tilde{w} \rightarrow v}^{t+1} / (d(d-1)^t)$. \blacksquare

Appendix A. Inverse of restricted Laplacians and hitting times

In this appendix we establish a standalone general connection between the inverse of restricted Laplacian matrices — which are obtained by removing certain rows and columns from Laplacians matrices — and hitting times of random walks on graphs. This connection (Proposition 22 below) is used repeatedly in Section 5 to prove the error characterization of the min-sum algorithm.

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 positive 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 $v \in V$. Let $L := D - W$ be the Laplacian matrix for G . 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 $P = D^{-1}W$ and initial condition $X_0 = v$. Analogously, denote by \mathbf{E}_v the expectation with respect to this law. The hitting time to a set $Z \subseteq V$ is defined as $T_Z := \inf\{k \geq 0 : X_k \in Z\}$. Let $Z \subseteq V$ be fixed, and define \bar{W} and \bar{D} as the matrix obtained by removing the rows and columns associated to Z from W and D , respectively. Let $\bar{V} := V \setminus Z$, $\bar{E} := E \setminus \{\{u, v\} \in E : u \in Z \text{ or } v \in Z\}$ and consider the graph $\bar{G} := (\bar{V}, \bar{E})$. The quantity $\bar{L} := \bar{D} - \bar{W}$ represents a restricted Laplacian — sometimes also called grounded Laplacian. Let $\bar{P} := \bar{D}^{-1}\bar{W}$ be the transition matrix of the transient part of the killed random walk that is obtained from X by adding cemeteries at the sites in Z . Creating cemeteries at the sites in Z means modifying the walk X so that each site in Z becomes a recurrent state, i.e., once the walk is in state $\bar{z} \in Z$ 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.

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.

Proposition 21 *For any $v, w \in \bar{V}$, $k \geq 0$, we have $\bar{P}_{vw}^k = \mathbf{P}_v(X_k = w, T_Z > k)$.*

Proof Proof is by induction. Clearly, for any $v, w \in \bar{V}$, we have $\mathbf{P}_v(X_0 = w, T_Z > 0) = \mathbf{P}_v(X_0 = w) = \mathbf{1}_{v=w} = \bar{P}_{vw}^0$, which proves the statement for $k = 0$ ($\mathbf{1}_{v=w}$ is the indicator function). Assume that the statement holds for any time $i \geq 0$ up to $k > 0$. By the properties of conditional expectation, as $\{T_Z > k + 1\} = \{X_0 \notin Z, \dots, X_{k+1} \notin Z\}$, we have

$$\begin{aligned} \mathbf{P}_v(X_{k+1} = w, T_Z > k + 1) &= \mathbf{E}_v[\mathbf{P}_v(X_{k+1} = w, T_Z > k + 1 | X_0, \dots, X_k)] = \\ &= \mathbf{E}_v[\mathbf{1}_{\{X_0 \notin Z, \dots, X_k \notin Z\}} \mathbf{P}_v(X_{k+1} = w, X_{k+1} \notin Z | X_0, \dots, X_k)]. \end{aligned}$$

for any $v, w \in \bar{V}$. By the Markov property, on the event $\{X_k \notin Z\}$, we have

$$\mathbf{P}_v(X_{k+1} = w, X_{k+1} \notin Z | X_0, \dots, X_k) = \mathbf{P}_{X_k}(X_1 = w, X_1 \notin Z) = \bar{P}_{X_k w},$$

so that by the induction hypothesis we have

$$\mathbf{P}_v(X_{k+1} = w, T_Z > k + 1) = \mathbf{E}_v[\mathbf{1}_{\{T_Z > k\}} \bar{P}_{X_k w}] = \sum_{u \in V \setminus Z} \mathbf{P}_v(X_k = u, T_Z > k) \bar{P}_{uw} = \bar{P}_{vw}^{k+1},$$

which proves the statement for $k + 1$. ■

We can now relate the inverse of the reduced Laplacian \bar{L} with the hitting times of the original random walk X . The following result represents the main computational tool used in Section 5 to establish decay of correlation in the computation tree that supports the operations of the min-sum algorithm over time.

Proposition 22 *For each $v, w \in \bar{V}$, we have*

$$\bar{L}_{vw}^{-1} = \bar{L}_{ww}^{-1} \mathbf{P}_v(T_w < T_Z), \quad \bar{L}_{ww}^{-1} = \frac{1}{d_w} \mathbf{E}_w \left[\sum_{k=0}^{T_Z} \mathbf{1}_{X_k=w} \right].$$

Proof Let us first assume that \bar{G} is connected. The matrix \bar{P} is clearly sub-stochastic. 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 Horn and Johnson (1986), for instance), so that the Neumann series $\sum_{k=0}^{\infty} \bar{P}^k$ converges. The Neumann series expansion for \bar{L}^{-1} yields $\bar{L}^{-1} = \sum_{k=0}^{\infty} (I - \bar{D}^{-1} \bar{L})^k \bar{D}^{-1} = \sum_{k=0}^{\infty} \bar{P}^k \bar{D}^{-1}$, or, $\bar{L}_{vw}^{-1} = \frac{1}{d_w} \sum_{k=0}^{\infty} \bar{P}_{vw}^k$. As $\bar{P}_{vw}^k = \mathbf{P}_v(X_k = w, T_Z > k)$ by Proposition 21, by the Monotone convergence theorem we can take the summation inside the expectation and get $\sum_{k=0}^{\infty} \bar{P}_{vw}^k = \sum_{k=0}^{\infty} \mathbf{E}_v[\mathbf{1}_{X_k=w} \mathbf{1}_{T_Z > k}] = \mathbf{E}_v[\sum_{k=0}^{T_Z-1} \mathbf{1}_{X_k=w}]$, which equals $\mathbf{E}_v[\sum_{k=0}^{T_Z} \mathbf{1}_{X_k=w}]$ as $X_{T_Z} \in Z$, and $w \notin Z$. 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 the 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_Z are two stopping times for X , and so is their minimum $S := T_w \wedge T_Z$. As either $X_S = w$ or $X_S \in Z$, we have $\mathbf{E}_v[\sum_{k=0}^{T_Z} \mathbf{1}_{X_k=w}] = \mathbf{E}_v[\sum_{k=0}^{T_Z} \mathbf{1}_{X_k=w} | X_S = w] \mathbf{P}_v(X_S = w)$, where we used that, conditionally on $\{X_S \in Z\} = \{T_w > T_Z\}$, clearly $\sum_{k=0}^{T_Z} \mathbf{1}_{X_k=w} = 0$. Conditionally on $\{X_S = w\} = \{T_w < T_Z\} = \{S = T_w\}$, we have $T_Z = S + \inf\{k \geq 0 : X_{S+k} \in Z\}$, 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 by assumption so that the Markov chain will almost surely eventually hit either w or a site in Z)

$$\mathbf{E}_v \left[\sum_{k=0}^{T_Z} \mathbf{1}_{X_k=w} \middle| X_S = w \right] = \mathbf{E}_v \left[\sum_{k=0}^{\inf\{k \geq 0 : X_{S+k} \in Z\}} \mathbf{1}_{X_{S+k}=w} \middle| X_S = w \right] = \mathbf{E}_w \left[\sum_{k=0}^{T_Z} \mathbf{1}_{X_k=w} \right].$$

Putting everything together we have $\bar{L}_{vw}^{-1} = \frac{1}{d_w} \mathbf{E}_w[\sum_{k=0}^{T_Z} \mathbf{1}_{X_k=w}] \mathbf{P}_v(T_w < T_Z)$. As $\mathbf{P}_w(T_w < T_Z) = 1$, clearly $\bar{L}_{vw}^{-1} = \frac{1}{d_w} \mathbf{E}_w[\sum_{k=0}^{T_Z} \mathbf{1}_{X_k=w}]$ so that $\bar{L}_{vw}^{-1} = \bar{L}_{ww}^{-1} \mathbf{P}_v(T_w < T_Z)$. The argument just presented extends easily to the case when \bar{G} is not connected. 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 the argument above can be applied to each block. ■

Appendix B. Technical results

This appendix contains several technical results that are used throughout in the paper. The following result in quadratic optimization has been used repeatedly in the main text.

Proposition 23 *Given $R \in \mathbb{R}^{m \times m}$ symmetric positive definite, $h \in \mathbb{R}^m$, $A \in \mathbb{R}^{n \times m}$, $b \in \mathbb{R}^n$ in the range of A , the unique solution of the optimization problem*

$$\begin{aligned} & \text{minimize} && \frac{1}{2} x^T R x + h^T x \\ & \text{subject to} && A x = b. \end{aligned}$$

over $x \in \mathbb{R}^m$ reads $x^* = R^{-1} A^T L^+ b + (R^{-1} A^T L^+ A - I) R^{-1} h$ with $L := A R^{-1} A^T$.

Proof The KKT equations for this problem read

$$\begin{pmatrix} R & -A^T \\ A & 0 \end{pmatrix} \begin{pmatrix} x \\ \nu \end{pmatrix} = \begin{pmatrix} -h \\ b \end{pmatrix}.$$

Solving the first equation for x we have $x = R^{-1}(A^T \nu - h)$, which, upon substitution into the second equation $Ax = b$ yields a solution $\nu = L^+(A R^{-1} h + b)$, and the proof follows. ■

We now present the proof of Proposition 5 in Section 4.

Proof [Proof of Proposition 5] Fix $d \geq 3$, $t \geq 3$. As $s \rightarrow h_s$ monotonically converges to 0 exponentially fast, we expect that the function $s \rightarrow \delta_s$ — and hence the values of

both $b_{d,t}$ and $c_{d,t}$ — can be controlled by the system that is obtained by replacing h_s with 0 for any $s \geq 0$. Thus, we define $\hat{\delta}_0 := \delta_0$, $\hat{\delta}_1 := \delta_1$, and $\hat{\delta}_s := \frac{1}{d-1}(2 + \frac{(d-2)^2}{d-1})\hat{\delta}_{s-1} - \frac{1}{(d-1)^2}\hat{\delta}_{s-2}$, $s \geq 2$. The characteristic equation of the dynamical system $\hat{\delta}_s$, $s \geq 0$, reads $(d-1)^2\lambda^2 - (d^2 - 2d + 2)\lambda + 1 = 0$. This equation has two distinct real solutions: $\lambda_+ = 1$ and $\lambda_- = 1/(d-1)^2 < 1$. Imposing the initial conditions for $\hat{\delta}_0$ and $\hat{\delta}_1$, we find $\hat{\delta}_s = c_+ + c_- \lambda_-^s \geq 0$ with $c_+ = \frac{d(d-1)((d-1)^2+1)-1}{d((d-1)^2-1)} > 0$ and $c_- = 1/d - c_+ < 0$.

We begin by analyzing $c_{d,t}$. To this end, define $\hat{c}_{d,t} := \frac{1}{d-1}(1 + \frac{1}{(d-2)(1+h_t)})\hat{\delta}_{t-2} - \frac{1}{(d-1)^2}\hat{\delta}_{t-3}$. First, we prove that $c_{d,t} \geq 1$. By induction it is easy to show that $(d-1)(\delta_s - \hat{\delta}_s) \geq \delta_{s-1} - \hat{\delta}_{s-1} \geq 0$ for any $s \geq 1$, so we have $c_{d,t} \geq \hat{c}_{d,t}$. As $\hat{\delta}_s \geq c_+ + c_- \lambda_-$ for any $s \geq 1$ and $\hat{\delta}_s \leq c_+$ for any $s \geq 0$, we get $\hat{c}_{d,t} \geq \frac{1}{d-1}(1 + \frac{1}{(d-2)(1+h_1)})(c_+ + c_- \lambda_-) - \frac{1}{(d-1)^2}c_+ \geq 1$. Second, we prove that $c_{d,t} \leq 1 + \varepsilon_d$. Given $\alpha := 1 + \frac{1}{(d-1)^2}$, we first show that $(d-1)(\alpha\hat{\delta}_s - \delta_s) \geq \alpha\hat{\delta}_{s-1} - \delta_{s-1} \geq 0$ for any $s \geq 1$, so that $c_{d,t} \leq \alpha\hat{c}_{d,t}$. Define $\xi_s := \alpha\hat{\delta}_s - \delta_s$. We proceed by induction. Note that $(d-1)\xi_1 \geq \xi_0 \geq 0$. For a given $s \geq 3$, assume that $(d-1)\xi_r \geq \xi_{r-1}$ for all $r \leq s-1$. We will prove that $(d-1)\xi_s \geq \xi_{s-1}$. We have $(d-1)\xi_s - \xi_{s-1} = \frac{(d-1)\xi_{s-1} - \xi_{s-2}}{d-1} + \frac{(d-2)^2}{d-1}((1+h_{s+2})\xi_{s-1} - \alpha h_{s+2}\hat{\delta}_{s-1})$. As $\hat{\delta}_{s-1} \leq c_+$ and $(d-1)^{s-1}h_{s+2} \leq h_3$, with $c_+h_3 < 1$, by the induction hypothesis we have $\xi_{s-1} \geq \frac{\xi_1}{(d-1)^{s-2}} = \frac{(\alpha-1)\delta_1}{(d-1)^{s-2}}$ and $(d-1)\xi_s - \xi_{s-1} \geq \frac{(d-2)^2}{(d-1)^s}((d-1)(\alpha-1)\delta_1 - \alpha) = 0$, which proves the induction step. As $\frac{1}{d} = c_+ + c_- \leq \hat{\delta}_s \leq c_+$, we get $c_{d,t} \leq \alpha\hat{c}_{d,t} \leq \bar{\alpha} := \alpha(\frac{c_+}{d-1}(1 + \frac{1}{d-2}) - \frac{1}{d(d-1)^2})$. The function $d \geq 3 \rightarrow \varepsilon_d := \bar{\alpha} - 1$ is decreasing, with $\varepsilon_d \rightarrow 0$ as $d \rightarrow \infty$, and $\varepsilon_3 < 3$.

By induction it is immediate to show that $(d-1)\delta_s \geq \delta_{s-1} \geq 0$ for any $s \geq 1$, so that a direct comparison yields $b_{d,t} \leq c_{d,t}$. Finally, we prove that $b_{d,t} \geq \frac{d-2}{d-1}$. Define $\hat{b}_{d,t} := \frac{1}{(d-1)^2}(1 + (d-2)(1+h_{t+1}))\hat{\delta}_{t-2} - \frac{(d-2)(1+h_{t+1})}{(d-1)^3}\hat{\delta}_{t-3}$. As $(d-1)(\delta_s - \hat{\delta}_s) \geq \delta_{s-1} - \hat{\delta}_{s-1} \geq 0$ for $s \geq 1$, we have $b_{d,t} \geq \hat{b}_{d,t}$. As $\hat{\delta}_s \geq \hat{\delta}_3 = c_+ + c_- \lambda_-^3$ for $s \geq 3$ and $\hat{\delta}_s \leq c_+$ for $s \geq 0$, by the monotonicity of h_s we get $\hat{b}_{d,t} \geq \frac{1}{(d-1)^2}(1 + (d-2))(c_+ + c_- \lambda_-^3) - \frac{(d-2)(1+h_1)}{(d-1)^3}c_+ \geq \frac{d-2}{d-1}$. ■

The next proposition yields a recursive formula to compute the quantities $\tilde{\Delta}_e^{(t)}$'s in (8).

Proposition 24 (Non-backtracking random walks) *Let $\vec{G} = (V, \vec{E})$ be a connected d -regular graph with $d \geq 2$, and let $B \in \mathbb{R}^{V \times V}$ be the adjacency matrix for the corresponding undirected graph $G = (V, E)$, that is, the symmetric matrix defined by $B_{vw} := 1$ if $\{v, w\} \in E$, $B_{vw} := 0$ otherwise. For $e = (v, w) \in \vec{E}$, the quantity $\tilde{\Delta}_e^{(t)}$ in (8) can be computed as*

$$\tilde{\Delta}_e^{(1)} = \frac{B(\mathbf{1}_v - \mathbf{1}_w)}{d-1} + \frac{\mathbf{1}_v - \mathbf{1}_w}{d-1}, \quad \tilde{\Delta}_e^{(2)} = \frac{B\tilde{\Delta}_e^{(1)}}{d-1} + \frac{\mathbf{1}_w - \mathbf{1}_v}{d-1}, \quad \tilde{\Delta}_e^{(t)} = \frac{B\tilde{\Delta}_e^{(t-1)} - \tilde{\Delta}_e^{(t-2)}}{d-1}, \quad t \geq 3.$$

Proof Fix $v, w \in V$, and let $\#_z^t$ be the number of non-backtracking paths that connect node v to node z in t time steps, where the first step is different from w . By definition B_{vz} equals the number of paths that connect node v to node z in a single step. Clearly, $\#_z^1 = B_{vz} - (\mathbf{1}_w)_z$, which, by the symmetry of B , can be written in vector form as $\#^1 = B\mathbf{1}_v - \mathbf{1}_w$. It is immediate to verify that $\#^2 = B\#^1 - (d-1)\mathbf{1}_v$, as there are $(d-1)$ paths that can backtrack to v at time $t = 2$. For $t \geq 3$ we have $\#^t = B\#^{t-1} - (d-1)\#^{t-2}$.

In fact, $(B\#^{t-1})_z$ equals the number of paths that go from v to z in t steps, where the first move is different from w , such that the first $t-1$ steps are non-backtracking. To count only the paths that are non-backtracking for all t steps, we need to subtract from $(B\#^{t-1})_z$ the number of paths that do backtrack at the t -th step, which are the ones that are at z both at step $t-2$ and t . Clearly, there are $(d-1)\#_z^{t-2}$ of these paths, as vertex z has d neighbors but only $d-1$ moves are allowed at the $(t-1)$ -th step due to the non-backtracking nature of the paths up to step $t-1$. The proof follows noticing that the total number of non-backtracking after t steps is $d(d-1)^{t-1}$, so that $\mathbf{P}_v(Y_1 \neq w, Y_t = z) = \#_z^t / (d(d-1)^{t-1})$ and $P_{vz}^{(t,w)} = \mathbf{P}_v(Y_t = z | Y_1 \neq w) = \mathbf{P}_v(Y_1 \neq w, Y_t = z) / \mathbf{P}_v(Y_1 \neq w)$, $\mathbf{P}_v(Y_1 \neq w) = (d-1)/d$. ■

References

- C. Berrou, A. Glavieux, and P. Thitimajshima. Near shannon limit error-correcting coding and decoding: Turbo-codes. 1. In *Communications, 1993. ICC '93 Geneva. Technical Program, Conference Record, IEEE International Conference on*, volume 2, pages 1064–1070 vol.2, 1993.
- A. Braunstein, M. Mézard, and R. Zecchina. Survey propagation: An algorithm for satisfiability. *Random Struct. Algorithms*, 27(2):201–226, 2005.
- Michael B. Cohen, Rasmus Kyng, Gary L. Miller, Jakub W. Pachocki, Richard Peng, Anup B. Rao, and Shen Chen Xu. Solving sdd linear systems in nearly $m \log 1/2n$ time. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing, STOC '14*, pages 343–352. ACM, 2014.
- David Gamarnik, Devavrat Shah, and Yehua Wei. Belief propagation for min-cost network flow: Convergence and correctness. *Operations Research*, 60(2):410–428, 2012.
- Roger A. Horn and Charles R. Johnson, editors. *Matrix Analysis*. Cambridge University Press, New York, NY, USA, 1986.
- Jonathan A. Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. A simple, combinatorial algorithm for solving sdd systems in nearly-linear time. *CoRR*, abs/1301.6628, 2013.
- I. Koutis, G. L. Miller, and R. Peng. A nearly- $m \log n$ time solver for sdd linear systems. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 590–598, 2011.
- Rasmus Kyng and Sushant Sachdeva. Approximate gaussian elimination for laplacians: Fast, sparse, and simple. *CoRR*, abs/1605.02353, 2016.
- M. Mezard, G. Parisi, and Riccardo Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297:812 – 815, 2002.
- C. C. Moallemi and B. Van Roy. Convergence of min-sum message passing for quadratic optimization. *IEEE Transactions on Information Theory*, 55(5):2413–2423, 2009.

- Ciamac C. Moallemi and Benjamin Van Roy. Convergence of min-sum message-passing for convex optimization. *Information Theory, IEEE Transactions on*, 56(4):2041–2050, 2010.
- Richard Peng and Daniel A. Spielman. An efficient parallel solver for sdd linear systems. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pages 333–342, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2710-7. doi: 10.1145/2591796.2591832. URL <http://doi.acm.org/10.1145/2591796.2591832>.
- J. Laurie Snell Peter G. Doyle. *Random Walks and Electric Networks*, volume 22. Mathematical Association of America, 1 edition, 1984.
- T. J. Richardson and R. L. Urbanke. The capacity of low-density parity-check codes under message-passing decoding. *IEEE Transactions on Information Theory*, 47(2):599–618, 2001.
- Daniel A. Spielman. *Algorithms, graph theory, and linear equations in Laplacian matrices.*, pages 2698–2722. Hackensack, NJ: World Scientific; New Delhi: Hindustan Book Agency, 2011.
- Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing*, STOC '04, pages 81–90. ACM, 2004.
- Daniel A. Spielman and Shang-Hua Teng. Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM Journal on Matrix Analysis and Applications*, 35(3):835–885, 2014.
- Nisheeth K. Vishnoi. $Lx = b$ Laplacian Solvers and Their Algorithmic Applications. *Foundations and Trends® in Theoretical Computer Science*, 8(1–2):1–141, 2013.