



HAL
open science

Asynchrony and Acceleration in Gossip Algorithms

Hadrien Hendrikx, Laurent Massoulié, Mathieu Even

► **To cite this version:**

Hadrien Hendrikx, Laurent Massoulié, Mathieu Even. Asynchrony and Acceleration in Gossip Algorithms. 2021. hal-02989459v2

HAL Id: hal-02989459

<https://hal.science/hal-02989459v2>

Preprint submitted on 10 Feb 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Asynchrony and Acceleration in Gossip Algorithms

MATHIEU EVEN, Inria, ENS Paris, PSL Research University, France

HADRIEN HENDRIKX, Inria, ENS Paris, PSL Research University, France

LAURENT MASSOULIÉ, Inria, ENS Paris, PSL Research University, France

This paper considers the minimization of a sum of smooth and strongly convex functions dispatched over the nodes of a communication network. Previous works on the subject either focus on synchronous algorithms, which can be heavily slowed down by a few slow nodes (the *straggler problem*), or consider a model of asynchronous operation (Boyd et al. [5]) in which adjacent nodes communicate at the instants of *Poisson point processes*. We have two main contributions. 1) We propose *CACDM* (a Continuously Accelerated Coordinate Dual Method), and for the Poisson model of asynchronous operation, we prove *CACDM* to converge to optimality at an accelerated convergence rate in the sense of Nesterov and Stich [33]. In contrast, previously proposed asynchronous algorithms have not been proven to achieve such accelerated rate. While *CACDM* is based on discrete updates, the proof of its convergence crucially depends on a continuous time analysis. 2) We introduce a new communication scheme based on *Loss-Networks*, that is programmable in a fully asynchronous and decentralized way, unlike the Poisson model of asynchronous operation that does not capture essential aspects of asynchrony such as non-instantaneous communications and computations. Under this Loss-Network model of asynchrony, we establish for *CDM* (a Coordinate Dual Method) a rate of convergence in terms of the eigengap of the Laplacian of the graph weighted by local effective delays. We believe this eigengap to be a fundamental bottleneck for convergence rates of asynchronous optimization. Finally, we verify empirically that *CACDM* enjoys an accelerated convergence rate in the Loss-Network model of asynchrony.

Additional Key Words and Phrases: gossip algorithms, loss networks, distributed optimization, asynchrony, acceleration

1 INTRODUCTION

In this paper, we consider minimization of a function f given by a sum of local functions:

$$\min_{x \in \mathbb{R}^d} f(x) := \sum_{i=1}^n f_i(x). \quad (1)$$

A typical example is provided by Empirical Risk Minimization (ERM), in which the local functions f_i correspond to the empirical risk evaluated on subsets of the whole dataset. We further assume that there is an underlying communication network, and that each f_i , or gradients thereof, can only be computed at node i of this network. In the case of ERM, f_i represents the empirical risk for the dataset available at node i . We aim to solve Problem (1) in a decentralized fashion, where each node can only communicate with its neighbors in the graph.

Another important example is that of network averaging. It corresponds to $f_i(x) = \|x - c_i\|^2$ where c_i is a vector attached to node i . The solution of Problem (1) is then provided by $x^* = \frac{1}{n} \sum_{i=1}^n c_i$.

Typical decentralized approaches for this problem rely on gossip communications [39] and first order local gradient steps [6, 32, 38, 41, 43, 46]. Yet, these approaches often rely on global synchronous rounds, in which all nodes exchange with their neighbours at the same time. Such synchronous approaches are well suited to networks with homogeneous communication and computation delays. However the presence of a few slow links or nodes drastically degrades their performance. Our work targets asynchronous distributed algorithms, for which we aim to obtain

Authors' addresses: Mathieu Even, mathieu.even@inria.fr, Inria, ENS Paris, PSL Research University, France; Hadrien Hendrikx, hadrien.hendrikx@inria.fr, Inria, ENS Paris, PSL Research University, France; Laurent Massoulié, laurent.massoulie@inria.fr, Inria, ENS Paris, PSL Research University, France.

fast rates of convergence in networks with heterogeneous computation and communication delays, while being competitive with synchronous approaches in homogeneous environments.

1.1 Main Contributions

We consider local operations and communication schemes where each pair of neighbor nodes (i, j) can exchange local variables at *activation times* of the corresponding edge (ij) . We denote by $\mathcal{P}_{ij} \subset \mathbb{R}^+$ the *Point process* of the corresponding activation times. Upon activation of edge (ij) , nodes i and j can exchange local variables such as gradients of their local functions and update their local variables accordingly. We mainly study two models for the point processes \mathcal{P}_{ij} : i) the Poisson model of asynchrony popularized by Boyd et al. [5] where $(\mathcal{P}_{ij})_{(ij) \in E}$ are independent *Poisson point processes* of rates p_{ij} [17]. We refer to this model as the *Poisson point process model (P.p.p. model)*. ii) A more complex model, inspired by loss networks (Kelly [16]), that we call *Refined Loss Network Model (RLNM)*, designed to capture essential aspects of asynchronous communications and computations.

1.1.1 Randomized Gossip and P.p.p. model. We extend results obtained by [5] on gossip algorithms for network averaging to more general optimization problems of the form of Problem (1) through a dual formulation. We obtain a convergence rate that depends on both the condition number of the optimization problem and the Laplacian matrix of the graph, weighted by the rates of the *Poisson point processes* \mathcal{P}_{ij} . The proof relies on a continuous-time analysis, which paves the way for the introduction of an accelerated algorithm, *CACDM (Continuously Accelerated Coordinate Dual Method)*. *CACDM* can be interpreted as an accelerated coordinate gradient descent on the dual problem involving infinitesimal contractions. Using this interpretation we prove that *CACDM* converges at an accelerated rate in the sense of Nesterov and Stich [33]. To the best of our knowledge, this is the first asynchronous algorithm proven to achieve accelerated convergence rates in the *P.p.p. model*.

1.1.2 Refined Loss Network Model. Though the *P.p.p.* model is very convenient, it assumes that communications and computations are performed instantaneously. We thus modify the communication scheme in order to model communications in a more realistic way: busy nodes (*i.e.* computing or communicating nodes) are made unavailable for other nodes to communicate with. This model is directly inspired by Loss Networks, where *busy* nodes are locked away from the network, which we refine by adding a *busy-checking* operation. For this communication model, we derive a rate of convergence that depends on the Laplacian matrix of the graph weighted by local communication constraints. Thus, we are able to recover the robustness to stragglers that we had with the *P.p.p.* model, but with a theory that is more faithful to the implementation. The construction and analysis of this model enable us to identify key parameters of the communication network that condition achievable convergence rates for realistic asynchronous and distributed operation.

1.2 Related Work

1.2.1 Gossip Algorithms and Asynchrony. In gossip averaging algorithms [5, 10], nodes of the network communicate with their neighbors without any central coordinator in order to compute the global average of local vectors. These algorithms are particularly relevant since they can be generalized to address our distributed optimization problem with local functions f_i beyond the special case $f_i(x) = \|x - c_i\|^2$. Two types of gossip algorithms appear in the literature: synchronous ones, where all nodes communicate with each other simultaneously [4, 10, 38], and asynchronous ones also called randomized gossip [5, 14, 31], where at a defined time $t \geq 0$, only a pair of adjacent nodes can communicate. In the synchronous framework, the communication speed is limited by the slowest node (*straggler* problem).

Although qualified as asynchronous, the *P.p.p.* model cannot be programmed in a fully distributed and asynchronous structure: it assumes that communications and computations are instantaneous. Two different approaches can be considered to deal with the fact that communications and computations are in fact non-instantaneous: (i) when a node i receives information from a neighbor j at a time $t \geq 0$, account for the fact that this information is delayed, or (ii) forbid communications with a *busy* (i.e. communicating or computing) edge, thereby removing the need to handle delayed information. The first approach (i) is considered for asynchronous but centralized optimization by [19, 34], where delayed variables are modelled as so-called *perturbed iterates*. The second approach (ii) is reminiscent of *Loss-Networks*, initially considered for telecommunication networks [16], yet also adequate to reflect primitives in distributed computing such as *locks* and *atomic transactions*.

In the *perturbed iterate* modelling, a central unit delegates computations to workers. Asynchrony lies in the fact that these workers do not wait for the central unit to update their current version of the optimization variable x , but instead send gradients $\nabla f_i(x_i)$ whenever they can, even if based on outdated variable x_i . Thus, the parameter of the central unit is updated using perturbed (*delayed*) gradients [27]. Section 4 focuses on the second modelling: nodes behave as in the *P.p.p. model*, but are made *busy* and hence non-available for other nodes for a time $\tau_{ij} > 0$ after their activation. The system is asynchronous in the sense that communications are performed in a random pairwise fashion (instead of global synchronous rounds), and nodes do not wait for specific neighbours. Yet, received gradients are never out of date since nodes always finish their current operation (communicating or computing) before engaging in a new one.

1.2.2 Acceleration in an Asynchronous Setting. Acceleration means gaining order of magnitudes in terms of convergence speed, compared to classical algorithms. Accelerating gossip algorithms has been studied in previous works in the synchronous framework: *SSDA* [38], Chebyshev acceleration [30] Jacobi-Polynomial acceleration in the first iterations [4], or in the asynchronous *P.p.p. model*: Geographic Gossip [9], shift registers [23]. However, no algorithm in the *P.p.p. model* has been rigorously proven to achieve an accelerated rate for general graphs without additional synchronization between nodes. For instance, inspired by *ACDM* [33], [14] introduced *ESDACD*, where at each iteration, only a pair of adjacent nodes communicate, but all nodes need to make local contractions and thus need to know that an update is taking place somewhere else in the graph. This last requirement, also present in *Stochastic Heavy Balls* methods [26], is a departure from purely asynchronous operation, and thus a limitation of these methods. Section 3.3 presents a continuous alternative to *ACDM*, where the contractions previously cited are made continuously. Our algorithm (*CACDM*, for Continuously Accelerated Coordinate Descent Method) obtains in the *P.p.p. model* the same accelerated rate as [9, 14, 26] for any graph, without assuming access to any global iteration counter: it only needs local clock-synchronization between adjacent nodes. Although our analysis of *CACDM* does not extend to more general communication models such as those presented in Section 4, we observe empirically that *CACDM* enjoys accelerated rates in the Loss-Network model as well as in the *P.p.p. model*.

The detailed problem statement and notations are given in Section 2. Section 3 contains our results on asynchronous gossip in the *P.p.p. model*, first for a non-accelerated algorithm based on simple gradient descent steps, then for the accelerated algorithm *CACDM*. Section 4 finally presents our results for gossip algorithms in the *refined loss network model*.

2 PROBLEM FORMULATION AND NOTATIONS

2.1 Basic assumptions and notations

The communication network is represented by an undirected graph $G = (V, E)$ on the set of nodes $V = [n]$, and is assumed to be connected. Two nodes are said to be neighbors or adjacent in the graph, and we write $i \sim j$, if $(ij) \in E$. Two edges $(ij), (kl) \in E$ are adjacent in the graph if $(ij) = (kl)$ or if they share a node. Each node $i \in V$ has access to a local function f_i defined on \mathbb{R}^d , assumed to be L_i -smooth and σ_i -strongly convex [7], i.e. $\forall x, y \in \mathbb{R}^d$:

$$\begin{aligned} f_i(x) &\leq f_i(y) + \langle \nabla f_i(y), x - y \rangle + \frac{L_i}{2} \|x - y\|^2, \\ f_i(x) &\geq f_i(y) + \langle \nabla f_i(y), x - y \rangle + \frac{\sigma_i}{2} \|x - y\|^2. \end{aligned} \quad (2)$$

Let us denote $f(z) = \sum_{i \in [n]} f_i(z)$ for $z \in \mathbb{R}^d$ and $F(x) = \sum_{i \in [n]} f_i(x_i)$ for $x = (x_1^\top, \dots, x_n^\top) \in \mathbb{R}^{n \times d}$ where $x_i \in \mathbb{R}^d$ is attached to node $i \in [n]$. Let

$$L_{\max} := \max_i L_i \text{ and } \sigma_{\min} := \min_i \sigma_i \quad (3)$$

denote the global complexity numbers. Computing gradients and communicating them between two neighboring nodes $i \sim j$ is assumed to take time $\tau_{ij} > 0$. This constant takes into account both the communication and computation times, and should be understood as an upper-bound on the delays between nodes i and j .

In this decentralized setting, Problem (1) can be formulated as follows:

$$\min_{x \in \mathbb{R}^{n \times d}, x_1 = \dots = x_n} F(x), \quad (4)$$

where $x_1 = \dots = x_n$ enforces consensus on all the nodes. We add the following structural constraints:

- (1) *Local computations*: node i (and node i only) can compute first-order characteristics of f_i such as ∇f_i or ∇f_i^* ;
- (2) *Local communications*: node i can send information only to neighboring nodes $j \sim i$.

These operations may be performed asynchronously and in parallel, and each node possesses a local version $x_i \in \mathbb{R}^d$ of the global parameter x . The rate of convergence of our algorithms will be controlled by the smallest positive eigenvalue γ of the Laplacian of graph G [28], weighted by some constants v_{ij} that depend on the local communication and computation delays.

DEFINITION 1 (GRAPH LAPLACIAN). *Let $(v_{ij})_{(ij) \in E}$ be a set of non-negative real numbers. The Laplacian of the graph G weighted by the v_{ij} 's is the matrix with (i, j) entry equal to $-v_{ij}$ if $(ij) \in E$, $\sum_{k \sim i} v_{ik}$ if $j = i$, and 0 otherwise. In the sequel v_{ij} always refers to the weights of the Laplacian, and $\gamma(v_{ij})$ denotes this Laplacian's second smallest eigenvalue.*

For any function $g : \mathbb{R}^p \rightarrow \mathbb{R}$, g^* denotes its *Fenchel conjugate* on \mathbb{R}^p defined as

$$\forall y \in \mathbb{R}^p, g^*(y) = \sup_{x \in \mathbb{R}^p} \langle x, y \rangle - g(x) \in \mathbb{R} \cup \{+\infty\}.$$

Throughout the paper, \mathcal{F}_t for $t \in \mathbb{R}^+$ denotes the filtration of the point processes $\mathcal{P} = \bigcup_{(ij) \in E} \mathcal{P}_{ij}$ up to time t . If $t_k, k \in \mathbb{N}^*$ (and $t_0 = 0$) are the successive points in \mathcal{P} , we write if there is no ambiguity $\mathcal{F}_k = \mathcal{F}_{t_k}, k \in \mathbb{N}^*$.

2.2 Dual Formulation of the Problem

A standard way to deal with the constraint $x_1 = \dots = x_n$, is to use a dual formulation [14, 38, 45], by introducing a dual variable λ indexed by the edges. We first introduce a matrix $A \in \mathbb{R}^{n \times E}$ such that $\text{Ker}(A^\top) = \text{Vect}(\mathbb{1})$ where $\mathbb{1}$ is the constant vector $(1, \dots, 1)^\top$ of dimension n . A is chosen such that:

$$\forall (ij) \in E, Ae_{ij} = \mu_{ij}(e_i - e_j). \quad (5)$$

for some non-null constants μ_{ij} . We define $\mu_{ij} = -\mu_{ji}$ for this writing to be consistent. This matrix A is a square root of the laplacian of the graph weighted by $v_{ij} = \mu_{ij}^2$. The constraint $x_1 = \dots = x_n$ can then be written $A^\top x = 0$. The dual problem reads as follows:

$$\min_{x \in \mathbb{R}^{n \times d}, A^\top x = 0} \sum_{i=1}^n f_i(x_i) = \min_{x \in \mathbb{R}^{n \times d}} \max_{\lambda \in \mathbb{R}^E} \sum_{i=1}^n f_i(x_i) - \langle A^\top x, \lambda \rangle.$$

Let $F_A^*(\lambda) := F^*(A\lambda)$ for $\lambda \in \mathbb{R}^{E \times d}$ where F^* is the Fenchel conjugate of F . The dual problem reads

$$\min_{x \in \mathbb{R}^{n \times d}, x_1 = \dots = x_n} F(x) = \max_{\lambda \in \mathbb{R}^{E \times d}} -F_A^*(\lambda).$$

Thus $F_A^*(\lambda) = \sum_{i=1}^n f_i^*((A\lambda)_i)$ is to be minimized over the dual variable $\lambda \in \mathbb{R}^{E \times d}$.

We now make a parallel between pairwise operations between adjacent nodes in the network and coordinate gradient steps on F_A^* . As $F_A^*(\lambda) = \max_{x \in \mathbb{R}^{n \times d}} -F(x) + \langle A\lambda, x \rangle$, to any $\lambda \in \mathbb{R}^{E \times d}$ a primal variable $x \in \mathbb{R}^{n \times d}$ is uniquely associated through the formula $\nabla F(x) = A\lambda$. The partial derivative of F_A^* with respect to coordinate (ij) of λ reads :

$$\nabla_{ij} F_A^*(\lambda) = (Ae_{ij})^\top \nabla F^*(A\lambda) = \mu_{ij}(\nabla f_i^*((A\lambda)_i) - \nabla f_j^*((A\lambda)_j)).$$

Consider then the following step of coordinate gradient descent for F_A^* on coordinate (ij) of λ , performed when edge (ij) is activated at iteration k (corresponding to time t_k), and where $U_{ij} = e_{ij}e_{ij}^\top$:

$$\lambda_{t_{k+1}} = \lambda_{t_k} - \frac{1}{(\sigma_i^{-1} + \sigma_j^{-1})\mu_{ij}^2} U_{ij} \nabla_{ij} F_A^*(\lambda_{t_k}). \quad (6)$$

Denoting $v_k = A\lambda_{t_k} \in \mathbb{R}^{n \times d}$, we obtain the following formula for updating coordinates i, j of v when ij activated:

$$v_{k+1,i} = v_{k,i} - \frac{\nabla f_i^*(v_{k,i}) - \nabla f_j^*(v_{k,j})}{\sigma_i^{-1} + \sigma_j^{-1}}, \quad (7)$$

$$v_{k+1,j} = v_{k,j} + \frac{\nabla f_i^*(v_{k,i}) - \nabla f_j^*(v_{k,j})}{\sigma_i^{-1} + \sigma_j^{-1}}. \quad (8)$$

Such updates can be performed locally at nodes i and j after communication between the two nodes. We refer in the sequel to this scheme as the Coordinate Descent Method (CDM). While $\lambda \in \mathbb{R}^{E \times d}$ is a dual variable defined on the edges, $v \in \mathbb{R}^{n \times d}$ is also a dual variable, but defined on the nodes. The *primal surrogate* of v is defined as $x = \nabla F^*(v)$ i.e. $x_i = \nabla f_i^*(v_i)$ at node i . It can hence be computed with local updates on v ((7) and (8)). Thus CDM, based on coordinate gradient descent for the dual problem, translates into *local* updates for the primal variables x_i . Note that in order to perform CDM, an initialization $v(0) \in \text{Im}(A)$ at all nodes is required, to ensure the existence of $\lambda \in \mathbb{R}^{E \times d}$ such that $A\lambda(0) = v(0)$. We thus usually take $v_i(0) = 0$ for all nodes i .

REMARK 1. We hence have two notions of duality. For $x = (x_1, \dots, x_n) \in \mathbb{R}^{n \times d}$ the primal variables associated with the network nodes, $v = (v_1, \dots, v_n) \in \mathbb{R}^{n \times d}$ is its convex-dual conjugate with $v_i = \nabla f_i(x_i)$, while $\lambda \in \mathbb{R}^{E \times d}$ such that $A\lambda = v$ is its edge-dual conjugate.

REMARK 2. Matrix A is introduced only for the purpose of the analysis. Indeed, we analyze our algorithms through edge-dual formulations, with updates of the form (6) on these variables. However, we present the algorithm with the convex-dual variables, (7),(8), for which μ_{ij}^2 and hence the effect of matrix A disappears.

2.3 Gossip Averaging Problem

As previously mentioned, the initial problem (1) with functions $f_i(x) = \frac{1}{2}\|x - c_i\|^2$, $x \in \mathbb{R}^d$ for some vectors $c_1, \dots, c_n \in \mathbb{R}^d$ reduces to the gossip averaging problem that aims at computing in a decentralized way with local computations the value $\bar{c} = \frac{1}{n} \sum_{i=1}^n c_i$. We contrast in this particular framework the rates that can be obtained by synchronous and asynchronous methods. These rates are expressed in terms of the weighted graph Laplacian, where for synchronous updates the edge weights are tuned to the worst-case delay, whereas in the asynchronous case, the edge weights can be tuned to local delay. Thus the advantage of asynchronous methods over synchronous ones is captured by these different edge weights in the considered Laplacian.

Synchronous Communications: In Synchronous Gossip Algorithm iterations [10], all nodes update their values synchronously by taking a weighted average of the values of their neighbors (Appendix A.1 for more details). These algorithms converge linearly with a rate given by the smallest eigenvalue of the graph Laplacian weighted by weights $v_{ij} \leq 1$. Since every iteration takes a time τ_{max} , synchronous Gossip algorithms have a linear rate of convergence $\gamma_{synch} = \gamma(v_{ij})$ with weights $v_{ij} \leq \tau_{max}^{-1}$ for all $(ij) \in E$ (Definition 1). We rephrase this as the following

PROPOSITION 1 (SYNCHRONOUS GOSSIP). *Let $x(t) = (x_1(t), \dots, x_n(t))^T \in \mathbb{R}^{n \times d}$ be the matrix of vectors $x_i(t)$ attached to node i at time $t \geq 0$. For continuous time $t \geq 0$ and for synchronous gossip algorithms as in [10], we have:*

$$\|x(t) - \bar{c}\|^2 \leq \exp(-(t - \tau_{max})\gamma_{synch})\|x(0) - \bar{c}\|^2, \quad (9)$$

with γ_{synch} the second smallest eigenvalue of the graph Laplacian weighted by $v_{ij} \equiv \tau_{max}^{-1}$.

Asynchronous Communications in the P.p.p. model: This is the setting of randomized gossip as considered by [5], where point processes \mathcal{P}_{ij} are independent P.p.p. of rates $p_{ij} > 0$. When edge (ij) is activated, nodes i and j update their values by making a local averaging (Appendix A.2). We have the following convergence result.

PROPOSITION 2 (RANDOMIZED GOSSIP). *For randomized gossip as in [5], we have:*

$$\mathbb{E}[\|x(t) - \bar{c}\|^2] \leq \exp(-t\gamma_{asynch})\|x(0) - \bar{c}\|^2, \quad (10)$$

with γ_{asynch} the second smallest eigenvalue of the graph Laplacian weighted by $v_{ij} = p_{ij}$. Moreover, this rate is optimal in the sense that there exists $x(0) \in \mathbb{R}^{n \times d}$ such that (10) is an equality for all $t \geq 0$.

Proofs of (9) and (10) and details about synchronous and randomized gossip can be found in Appendix A. Equation (10) follows from derivations in [5], combined with a study of infinitesimal intervals of times $[t, t + dt]$. We generalize this result to the initial optimization problem (1) in next Section.

In the P.p.p., the terms $1/p_{ij}$ capture the average time between consecutive activations of edge (ij) and are thus naturally related to the delays τ_{ij} . This suggests that asynchrony brings about a speed-up reflected by the change in the Laplacian's spectral gap $\gamma(v_{ij})$ when the weights $v_{ij} \equiv \tau_{max}^{-1}$ are replaced by $v_{ij} = \tau_{ij}^{-1}$. The fact that γ_{asynch} is optimal leads us to believe that this quantity - the smallest non-null eigenvalue of the Laplacian with local weights - best describes the asynchronous speed-up.

The above argument identifying τ_{ij} with p_{ij}^{-1} is heuristic. Our analysis of the Loss-Network model will establish a more rigorous bridge between spectral gap of Laplacian with edge weights based on local delays and convergence speed of asynchronous schemes.

3 RANDOMIZED GOSSIP: THE P.P.P. MODEL

3.1 The P.p.p. Model and Randomized Gossip Algorithms

The P.p.p. model: Each edge $(ij) \in E$ has a clock that ticks at the instants of a *Poisson point process* \mathcal{P}_{ij} of intensity p_{ij} , where the \mathcal{P}_{ij} are mutually independent. At each tick of its clock, edge (ij) is activated and nodes i and j can communicate together. The process $\mathcal{P} = \bigcup_{(ij) \in E} \mathcal{P}_{ij}$, \mathcal{P} is again *P.p.p.* of intensity

$$I = \sum_{(ij) \in E} p_{ij}. \quad (11)$$

Randomized Gossip Algorithm: Each node i maintains a local variable $(x_i(t))_{t \geq 0}$. We denote $(v_i(t))_{t \geq 0}$ its local convex-dual conjugate and write $v(t) = (v_i(t))_i$. We initialize with $v_i(0) = 0$ at all nodes. Based on the dual problem formulation in Section 2.2, we consider CDM. Specifically, when clock (ij) ticks at time $t \geq 0$, perform the following update on variable $v(t)$:

$$\begin{aligned} v_i(t) &\stackrel{t}{\leftarrow} v_i(t) - \frac{\nabla f_i^*(v_i(t)) - \nabla f_j^*(v_j(t))}{\sigma_i^{-1} + \sigma_j^{-1}}, \\ v_j(t) &\stackrel{t}{\leftarrow} v_j(t) + \frac{\nabla f_i^*(v_i(t)) - \nabla f_j^*(v_j(t))}{\sigma_i^{-1} + \sigma_j^{-1}}. \end{aligned} \quad (12)$$

The desired output at node i and time t is then $x_i(t) = \nabla f_i^*(v_i(t))$. Note that as mentioned in Section 2.2, the outputs $v_i(t)$ and $x_i(t)$ at any node i and time t are all completely independent from the initial choice of matrix A , whose only use is for analysis. Observe that in the gossip averaging problem, $v_i(t) = x_i(t) - x_i(0)$, and Equation (12) simplifies to

$$x_i(t), x_j(t) \stackrel{t}{\leftarrow} \frac{x_i(t) + x_j(t)}{2}, \quad (13)$$

which coincides with classical randomized gossip updates for the averaging problem.

3.2 Continuous Time Convergence Analysis

The classical analysis of gossip algorithms [5] proceeds as follows: at every clock tick of \mathcal{P} , an edge (ij) is selected with probability $q_{ij} = \frac{p_{ij}}{I}$. A discrete time analysis of state variables at these ticking times is then performed. In order to derive bounds for continuous time t , we instead study infinitesimal intervals of time $[t, t + dt]$, giving us more degrees of freedom, as shown in Section 3.3.

THEOREM 1. *For the CDM updates (12), in the P.p.p. model with intensities p_{ij} , we have the following guarantees for all $t \geq 0$*

$$\mathbb{E}(F^*(v(t)) - F^*(v^*)) \leq (F^*(v(0)) - F^*(v^*)) \exp\left(-\frac{\sigma_{\min}}{2L_{\max}} \gamma_p t\right), \quad (14)$$

where $v^* = A\lambda^*$ is the minimizer of F^* on $\text{Im}(A)$, λ^* being a minimizer of F_A^* , $\gamma_p = \gamma(p_{ij})$ is the spectral gap of the graph Laplacian weighted by weights $v_{ij} = p_{ij}$ and σ_{\min}, L_{\max} are defined in (3).

Since $x(t) = \nabla F^*(v(t))$ and $x^* = \nabla F^*(v^*)$ where x^* is the minimizer of F under the consensus constraint, we have on primal variable $x(t)$ (Lemma 3):

$$\mathbb{E} \left[\|x_t - x^*\|^2 \right] \leq \frac{2L_{\max}}{\sigma_{\min}^2} (F^*(v(0)) - F^*(v^*)) \exp \left(-\frac{\sigma_{\min}}{2L_{\max}} \gamma_p t \right). \quad (15)$$

We thus obtain a factor γ_p in the rate of convergence that reflects communication speed, and $\frac{\sigma_{\min}}{L_{\max}}$ that is an upper-bound on the condition number of the objective function. The sketch of proof below relies on a classical analysis of coordinate descent algorithms adapted to continuous time. The technical details are deferred to Appendix B. We believe the proof technique to be of independent interest: it could be applied to analyze optimization methods such as gradient descent algorithms (stochastic, proximal or accelerated ones) with increments ruled by *Poisson point processes* with simple proofs based on establishment of differential inequalities.

PROOF. We prove Theorem 1 by considering edge-dual variables $\lambda_t \in \mathbb{R}^{E \times d}$ associated to $x(t)$ and $v(t)$, in particular with $A\lambda_t = v(t)$ and $A\lambda^* = v^*$. Since $v(0) = 0$, we take $\lambda_0 = 0$. We consider matrix A in (5) with $\mu_{ij}^2 = \frac{p_{ij}}{\sigma_i^{-1} + \sigma_j^{-1}}$. When clock (ij) ticks at time $t \geq 0$, the following update is performed on variable λ_t :

$$\lambda_t \xleftarrow{t} \lambda_t - \frac{1}{(\sigma_i^{-1} + \sigma_j^{-1})\mu_{ij}^2} U_{ij} \nabla_{ij} F_A^*(\lambda_t). \quad (16)$$

Furthermore, note that we have $F^*(v(t)) = F_A^*(\lambda_t)$. A key ingredient in the proof is the lemma below, which establishes a local smoothness property. Its proof is given in Appendix B,

LEMMA 1. For $\lambda \in \mathbb{R}^{E \times d}$ and $ij \in E$, we have:

$$F_A^* \left(\lambda - \frac{1}{\mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})} U_{ij} \nabla_{ij} F_A^*(\lambda) \right) - F_A^*(\lambda) \leq -\frac{1}{2\mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})} \|\nabla_{ij} F_A^*(\lambda)\|^2. \quad (17)$$

Then using this, for $t \geq 0$ and $dt > 0$:

$$\begin{aligned} \mathbb{E}^{\mathcal{F}_t} [F_A^*(\lambda_{t+dt}) - F_A^*(\lambda_t)] &= (1 - Idt) \mathbb{E}^{\mathcal{F}_t} [F_A^*(\lambda_{t+dt}) - F_A^*(\lambda_t) | \text{no activations in } [t, t + dt]] \\ &\quad + \sum_{(ij) \in E} p_{ij} dt \mathbb{E}^{\mathcal{F}_t} [F_A^*(\lambda_{t+dt}) - F_A^*(\lambda_t) | (ij) \text{ activated in } [t, t + dt]] \\ &= -dt \sum_{ij \in E} p_{ij} (F_A^*(\lambda_t) - F_A^*(\lambda_t - \frac{1}{(\sigma_i^{-1} + \sigma_j^{-1})\mu_{ij}^2} U_{ij} \nabla_{ij} F_A^*(\lambda_t))) + o(dt) \\ &\leq -dt \sum_{ij \in E} \frac{p_{ij}}{2(\sigma_i^{-1} + \sigma_j^{-1})\mu_{ij}^2} \|\nabla_{ij} F_A^*(\lambda_t)\|^2 + o(dt) \\ &= -\frac{dt}{2} \|\nabla F_A^*(\lambda_t)\|^2 + o(dt) \end{aligned}$$

Lemma 8 in the Appendix implies that $\|\nabla F_A^*(\lambda)\| \geq 2\sigma_A (F_A^*(\lambda) - F_A^*(\lambda^*))$, where σ_A is the strong convexity parameter of F_A^* with respect to the Euclidean norm on the orthogonal of $\text{Ker}(A)$. We thus have:

$$\mathbb{E}^{\mathcal{F}_t} [F_A^*(\lambda_{t+dt}) - F_A^*(\lambda_t)] \leq -dt \sigma_A (F_A^*(\lambda_t) - F_A^*(\lambda^*)) + o(dt).$$

Then, dividing by dt and taking $dt \rightarrow 0$ yields: $\frac{d}{dt} \mathbb{E} [F_A^*(\lambda_t) - F_A^*(\lambda^*)] \leq -\sigma_A \mathbb{E} [F_A^*(\lambda_t) - F_A^*(\lambda^*)]$. We then obtain an exponential rate of convergence σ_A by integrating. Finally, Lemma 5 in the Appendix gives $\sigma_A \geq \frac{\lambda_{\min}^+(AA^\top)}{L_{\max}}$ where $\lambda_{\min}^+(AA^\top)$ is the smallest non-null eigenvalue of AA^\top . As AA^\top is the

Laplacian of the graph with weights $v_{ij} = \mu_{ij}^2 = \frac{p_{ij}}{\sigma_i^{-1} + \sigma_j^{-1}}$ (Lemma 6), we have $\lambda_{\min}^+(AA^\top) \geq \sigma_{\min} \gamma_p / 2$ and (16) follows. \square

REMARK 3. *The above study of infinitesimal intervals of time directly leads to continuous-time bounds. These could also be derived from a discrete time analysis: Denote by $t_k \geq 0$ the time of k -th activation, $k \in \mathbb{N}^*$, and $t_0 = 0$. We can prove that:*

$$\mathbb{E}[F_A^*(\lambda_{t_k}) - F_A^*(\lambda^*)] \leq (1 - \sigma_A/I)^k (F_A^*(\lambda_0) - F_A^*(\lambda^*)), \quad (18)$$

where $I = \sum_{(ij) \in E} p_{ij}$. Then, we have in continuous time, for any $t \in \mathbb{R}^+$:

$$\begin{aligned} \mathbb{E}[F_A^*(\lambda_t) - F_A^*(\lambda^*)] &= \sum_{k \in \mathbb{N}} \frac{e^{-It} (It)^k}{k!} \mathbb{E}[F_A^*(\lambda_t) - F_A^*(\lambda^*) | k \text{ activations in } [0, t]] \\ &\leq \sum_{k \in \mathbb{N}} \frac{e^{-It} (It)^k}{k!} (1 - \sigma_A/I)^k (F_A^*(\lambda_0) - F_A^*(\lambda^*)) \\ &= e^{-\sigma_A t} (F_A^*(\lambda_0) - F_A^*(\lambda^*)), \end{aligned}$$

giving the same result. However in the next Section, we will see that the continuous time viewpoint is essential in the design of the CACDM algorithm, as well as for its analysis through consideration of infinitesimal intervals and differential calculus.

3.3 Accelerated Gossip in the *P.p.p.* model

Inspired by previous works [14, 33], we propose CACDM (Continuously Accelerated Coordinate Descent Method), a gossip algorithm that, for the *P.p.p.* model, provably obtains an accelerated rate of convergence in the sense of Nesterov and Stich [33] (Theorem 2).

3.3.1 *CACDM algorithm and convergence guarantees.* Similarly to other standard acceleration techniques, the algorithm requires maintaining two variables $x(t), y(t) \in \mathbb{R}^{n \times d}$, whose convex-dual conjugates are denoted $u(t), v(t)$. The variable $v(t)$ plays the role of a momentum. We initialize such that $u(0) = v(0) = 0$. As in last subsection, variables $u_i(t), v_i(t), x_i(t), y_i(t) \in \mathbb{R}^d$ for $i \in [n]$ are attached to node i . The algorithm involves two types of operations: continuous contractions, and pairwise updates along each edge (ij) when its Poisson clock ticks.

- (1) **Continuous Contractions:** For all times $t \in \mathbb{R}^+$ and node $i \in [n]$, for some fixed $\theta > 0$ to be specified, make the infinitesimal contraction

$$\begin{pmatrix} u_i(t+dt) \\ v_i(t+dt) \end{pmatrix} = \begin{pmatrix} 1 - dtI\theta & dtI\theta \\ dtI\theta & 1 - dtI\theta \end{pmatrix} \begin{pmatrix} u_i(t) \\ v_i(t) \end{pmatrix},$$

between times t and $t+dt$. Between times $s < t$, if there is no activation of i , it consists in performing the contraction:

$$\begin{pmatrix} u_i(t) \\ v_i(t) \end{pmatrix} = \exp\left((t-s)I \begin{pmatrix} -\theta & \theta \\ \theta & -\theta \end{pmatrix}\right) \begin{pmatrix} u_i(s) \\ v_i(s) \end{pmatrix} = \begin{pmatrix} \frac{1+e^{-2I\theta(t-s)}}{2} & \frac{1-e^{-2I\theta(t-s)}}{2} \\ \frac{1-e^{-2I\theta(t-s)}}{2} & \frac{1+e^{-2I\theta(t-s)}}{2} \end{pmatrix} \begin{pmatrix} u_i(s) \\ v_i(s) \end{pmatrix}. \quad (19)$$

- (2) **Local Updates:** Let γ_p be the smallest non-null eigenvalue of the Laplacian of the graph weighted by the local rates: $v_{ij} = p_{ij}$ (Definition 1), and L_{\max} defined in (3). When edge (ij)

is activated at time $t \geq 0$, perform the local update between nodes i and j :

$$u_i(t) \stackrel{t}{\leftarrow} u_i(t) - \frac{\nabla f_i^*(u_t(i)) - \nabla f_j^*(u_t(j))}{\sigma_i^{-1} + \sigma_j^{-1}}, \quad (20)$$

$$v_i(t) \stackrel{t}{\leftarrow} v_i(t) - \frac{\theta L_{\max}}{\gamma_p} \left(\nabla f_i^*(u_t(i)) - \nabla f_j^*(u_t(j)) \right), \quad (21)$$

and symmetrically at node j . The desired output at node i and at time t is then $x_i(t) = \nabla f_i^*(u_i(t))$ (Section 2.2).

This procedure can be performed asynchronously and at discrete times: the length $t - s$ between two activations of an edge that appears in the exponential contraction (19) is a local variable that can be computed from a local clock. More formally, the stochastic process defined above is the following, where $V_t = (u(t), v(t))^T$ and N_{ij} are independent *P.p.p.* of intensities p_{ij} :

$$dV_t = I \begin{pmatrix} -\theta & \theta \\ \theta & -\theta \end{pmatrix} V_t dt - \sum_{(ij) \in E} dN_{ij}(t) \begin{pmatrix} \frac{\nabla f_i^*(u_t(i)) - \nabla f_j^*(u_t(j))}{\sigma_i^{-1} + \sigma_j^{-1}} \\ \frac{\theta L_{\max}}{\gamma_p} \left(\nabla f_i^*(u_t(i)) - \nabla f_j^*(u_t(j)) \right) \end{pmatrix}.$$

Define the Lyapunov function

$$\mathcal{L}_t = \|v(t) - v^*\|_{(A^* \top A^*)^2}^2 + \frac{2\theta^2 S^2 L_{\max}^2}{\gamma_p^2} (F^*(u(t)) - F^*(v^*)), \quad (22)$$

where $v^* = A\lambda^*$ is the minimizer of F^* on $\text{Im}(A)$, λ^* being a minimizer of F_A^* , $\theta, S > 0$ to be defined, and A^* the pseudo-inverse of matrix A tuned with $\mu_{ij}^2 = p_{ij}$. Let γ_p be the smallest non-null eigenvalue of the Laplacian of the graph, with weights $v_{ij} = p_{ij}$ (Definition 1).

THEOREM 2. *For the CACDM algorithm defined by Equations (19), (20), (21) in the *P.p.p.* model, if $\theta = \sqrt{\frac{\gamma_p}{IS^2 L_{\max}}}$ for S verifying the inequality:*

$$S^2 \geq \sup_{(ij) \in E} \frac{(\sigma_i^{-1} + \sigma_j^{-1})}{2p_{ij}/I}, \quad (23)$$

where σ_i defined in (2), I in (11) and σ_{\min}, L_{\max} in (3), we have for all $t \in \mathbb{R}^+$:

$$\mathbb{E}[\mathcal{L}_t] \leq \mathcal{L}_0 e^{-I\theta t}.$$

where \mathcal{L}_t is defined in (22).

The proof of this theorem uses edge-dual variables and differential inequalities through the study of infinitesimal intervals $[t, t + dt]$ as in the proof of Theorem 1, further combined with inequalities in [33] for the study of accelerated coordinate descent. We first make a few comments on this theorem, and then proceed to its proof. Since $x(t) = \nabla F^*(u(t))$ and $v^* = \nabla F^*(x^*)$ where x^* is the minimizer of F under the consensus constraint, we have on primal variable $x(t)$:

$$\mathbb{E} \left[\|x_t - x^*\|^2 \right] \leq \frac{2L_{\max}}{\sigma_{\min}^2} \frac{\gamma_p}{2\theta^2 S^2 L_{\max}^2} \mathcal{L}_0 e^{-I\theta t}. \quad (24)$$

Remarks on the bound: (γ_p/I) is the normalized non-accelerated randomized gossip rate of convergence. It is divided by I so that the p_{ij} sum to 1. If there exists a constant $c > 0$ such that:

$$\forall (ij) \in E, \frac{p_{ij}}{I} \geq \frac{c}{|E|},$$

then $S^2 \geq \sigma_{\min}^{-1}|E|/c$, leading to the following rate of convergence:

$$I \times \sqrt{c \frac{\sigma_{\min}}{L_{\max}} \times \frac{Y_p}{I|E|}}.$$

Taking $I = 1$ (re-normalizing time) and the simple averaging problem leads to an improved rate of n^{-2} on the line graph instead of n^{-3} [28]. For the 2D-Grid, we have $n^{-3/2}$ instead of n^{-2} [28]. However, there is no improvement on the complete graph ($1/n$ in both cases). These rates are the same as [9, 14, 25]. Yet, our algorithm does not require to know the number of activations performed on the whole network, and only requires local clocks. Moreover, similarly to Hendrikx et al. [14], it works for any graph and for the more general problem of distributed optimization of smooth and strongly convex functions provided dual gradients of local functions are computable.

CACDM algorithm for the averaging problem: for the gossip averaging problem, we have $u(t) = x(t) - x(0)$, $v(t) = y(t) - y(0)$, and (20) and (21) read as:

$$\begin{aligned} x_i(t) &\stackrel{t}{\leftarrow} \frac{x_i(t) + x_j(t)}{2} \\ y_i(t) &\stackrel{t}{\leftarrow} y_i(t) - \frac{\theta}{Y_p} (x_t(i) - x_t(j)). \end{aligned}$$

The first variable thus performs classical local averagings while mixing continuously with the second one (the momentum).

3.3.2 Proof of Theorem 2. Let the two edge dual variables $\lambda, \omega \in \mathbb{R}^{E \times d}$ be the edge-dual conjugates of $x(t), y(t)$. Variable ω plays the role of the momentum. Since $u(0) = v(0) = 0$, we can take $\lambda_0 = \omega_0 = 0$. Operations (20) and (21) translate as follows on these variables when clock (ij) ticks. Let σ_A be the strong convexity parameter of F_A^* with respect to the Euclidean norm on the orthogonal of $\text{Ker}(A)$. In Appendix B, we prove that, if $\mu_{ij}^2 = p_{ij}$: $\sigma_A \leq \frac{Y_p}{L_{\max}}$.

While working with F_A^* and edge-dual variables, we use σ_A instead of $\frac{Y_p}{L_{\max}}$ as presented in the algorithm, in order to keep in mind its meaning for F_A^* . Define the coordinate gradient step:

$$\eta_{ij,t} = - \left(\begin{array}{c} \frac{1}{2\mu_{ij}^2(\sigma_i^{-1} + \sigma_j^{-1})} U_{ij} \nabla_{ij} F_A^*(\lambda_t) \\ \frac{\theta}{\sigma_A p_{ij}} U_{ij} \nabla_{ij} F_A^*(\lambda_t) \end{array} \right) \quad (25)$$

where $U_{ij} = e_{ij} e_{ij}^T$, and perform the gradient step:

$$\begin{pmatrix} \lambda_t \\ \omega_t \end{pmatrix} \stackrel{t}{\leftarrow} \begin{pmatrix} \lambda_t \\ \omega_t \end{pmatrix} + \eta_{ij,t} \quad (26)$$

Define:

$$L_t = \|\omega_t - \lambda^*\|^{*2} + \frac{2\theta^2 S^2}{\sigma_A^2} (F_A^*(\lambda_t) - F_A^*(\lambda^*)),$$

where $\|\cdot\|^*$ is the Euclidean norm on the orthogonal of $\text{Ker}(A)$, and λ^* is an optimizer of F_A^* . Note that we have $L_t = \mathcal{L}_t$ for all $t \geq 0$.

PROOF OF THEOREM 2. The proof closely follows the lines of Hendrikx et al. [14], Nesterov and Stich [33], adapted to fit our continuous time algorithm. Without loss of generality, we assume that $I = 1$ i.e. that the p_{ij} sum to 1 (by rescaling time with $t' = tI$). We denote $r_t = \|\omega_t - \lambda^*\|^*$, and $f_t = F_A^*(\lambda_t) - F_A^*(\lambda^*)$, such that $L_t = r_t^2 + \frac{2\theta^2 S^2}{\sigma_A^2} f_t$. Let $t \geq 0$ and $dt > 0$. The following equalities

and inequalities are true up to a $o(dt)$ approximation, which will disappear when we make $dt \rightarrow 0$. Let's start with the term r_t^2 :

$$\mathbb{E}^{\mathcal{F}_t} [r_{t+dt}^2] = (1 - dt)\mathbb{E}^{\mathcal{F}_t} [r_{t+dt}^2 | \text{no activations between } t \text{ and } t+dt] \quad (27)$$

$$+ dt\mathbb{E}^{\mathcal{F}_t} [r_{t+dt}^2 | 1 \text{ activation between } t \text{ and } t+dt] \quad (28)$$

For the first term, we get:

$$\begin{aligned} \mathbb{E}^{\mathcal{F}_t} [r_{t+dt}^2 | \text{no activation in } [t, t + dt]] &= \|(1 - \theta dt)\omega_t + \theta dt\lambda_t - \lambda^*\|^{*2} \\ &\leq (1 - \theta dt)r_t^2 + \theta dt\|\lambda_t - \lambda^*\|^{*2} \end{aligned}$$

where the inequality uses convexity of the squared function. For the other term, we decompose the event "1 activation between t and $t+dt$ " in the disjoint events " ij activated between t and $t+dt$ ", of probability $p_{ij}dt$, to get the following equation, which is true up to a $o(1)$ approximation (which is enough since we multiply by dt afterwards):

$$\begin{aligned} \mathbb{E}^{\mathcal{F}_t} [r_{t+dt}^2 | 1 \text{ activation between } t \text{ and } t+dt] &= \sum_{(ij) \in E} p_{ij} \left\| \omega_t - \frac{\theta}{p_{ij}\sigma_A} U_{ij} \nabla_{ij} F_A^*(\lambda_t) - \lambda^* \right\|^{*2} \\ &= \|\omega_t - \lambda^*\|^{*2} + \sum_{ij} p_{ij} \frac{\theta^2}{\sigma_A^2 p_{ij}^2} \|U_{ij} \nabla_{ij} F_A^*(\lambda_t)\|^{*2} - 2 \sum_{ij} p_{ij} \frac{\theta}{p_{ij}\sigma_A} \langle U_{ij} \nabla_{ij} F_A^*(\lambda_t), \omega_t - \lambda^* \rangle \quad (29) \end{aligned}$$

For the term $\sum_{ij} p_{ij} \frac{\theta^2}{\sigma_A^2 p_{ij}^2} \|U_{ij} \nabla_{ij} F_A^*(\lambda_t)\|^{*2}$, we get by definition of S^2 , and by a local smoothness inequality (namely, $\forall y, F_A^*(y) - F_A^*(y - \frac{1}{\mu_{ij}^2(\sigma_i^{-1} + \sigma_j^{-1})}) U_{ij} \nabla_{ij} F_A^*(y) \geq \frac{1}{2\mu_{ij}^2(\sigma_i^{-1} + \sigma_j^{-1})} \|\nabla_{ij} F_A^*(y)\|^2$ in Lemma 4):

$$\begin{aligned} \sum_{ij} p_{ij} \frac{\theta^2}{\sigma_A^2 p_{ij}^2} \|U_{ij} \nabla_{ij} F_A^*(\lambda_t)\|^{*2} &\leq \sum_{ij} p_{ij} \frac{2\theta^2 S^2}{\sigma_A^2 \mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})} \|U_{ij} \nabla_{ij} F_A^*(\lambda_t)\|^2 \\ &\leq \sum_{ij} p_{ij} \frac{2\theta^2 S^2}{\sigma_A^2} (F_A^*(\lambda_t) - F_A^*(\lambda_t - \frac{\theta}{\sigma_A p_{ij}} U_{ij} \nabla_{ij} F_A^*(\lambda_t))) \\ &= \frac{\theta^2 S^2}{\sigma_A^2} (F_A^*(\lambda_t) - \mathbb{E}^{\mathcal{F}_t} [F_A^*(\lambda_{t+dt}) | 1 \text{ activation in } [t, t+dt]]). \quad (30) \end{aligned}$$

For the term $-2 \sum_{ij} p_{ij} \frac{\theta}{p_{ij}\sigma_A} \langle U_{ij} \nabla_{ij} F_A^*(\lambda_t), \omega_t - \lambda^* \rangle$, we get, by adding and subtracting a λ_t in the bracket, and by convexity of $F_A^*(\sigma_A)$ is the strong convexity parameter of F_A^* :

$$\begin{aligned} -2dt \frac{\theta}{\sigma_A} \langle \nabla F_A^*(\lambda_t), \omega_t - \lambda^* \rangle &= -2dt \frac{\theta}{\sigma_A} \langle \nabla F_A^*(\lambda_t), \omega_t - \lambda_t \rangle - 2dt \frac{\theta}{\sigma_A} \langle \nabla F_A^*(\lambda_t), \lambda_t - \lambda^* \rangle \\ &\leq -2 \frac{1}{\sigma_A} \langle \nabla F_A^*(\lambda_t), \theta dt (\omega_t - \lambda_t) \rangle - 2dt \frac{\theta}{\sigma_A} (F_A^*(\lambda_t) - F_A^*(\lambda^*) + \sigma_A/2 \|\lambda_t - \lambda^*\|^{*2}) \end{aligned}$$

Then, let's define $\lambda'_{t+dt} = (1 - \theta dt)\lambda_t + \theta dt\omega_t = \mathbb{E}^{\mathcal{F}_t} [\lambda_{t+dt} | \text{no activations in } [t, t + dt]]$. By noticing that $\theta dt(\omega_t - \lambda_t) = \lambda'_{t+dt} - \lambda_t$, we get:

$$-2 \frac{1}{\sigma_A} \langle \nabla F_A^*(\lambda_t), \theta dt (\omega_t - \lambda_t) \rangle = -2 \frac{1}{\sigma_A} \langle \nabla F_A^*(\lambda_t), \lambda'_{t+dt} - \lambda_t \rangle \quad (31)$$

$$= -2 \frac{1}{\sigma_A} \langle \nabla F_A^*(\lambda'_{t+dt}), \lambda'_{t+dt} - \lambda_t \rangle \quad (32)$$

$$\leq -2 \frac{1}{\sigma_A} (F_A^*(\lambda'_{t+dt}) - F_A^*(\lambda_t)), \quad (33)$$

where from (31) to (32), the equality holds at $o(dt)$, as the left part of the bracket is true at $o(1)$ precision, and the right part of the bracket is a $O(dt)$. Then, we go from (32) to (33) using the convexity of F_A^* . By plugging (30) and (33) into (29), and rearranging the terms, we obtain:

$$\begin{aligned} \mathbb{E}^{\mathcal{F}_t} [r_{t+dt}^2] - r_t^2 &\leq -dt\theta r_t^2 + dt\theta \|\lambda_t - \lambda^*\|^2 \\ &\quad + dt \frac{2\theta^2 S^2}{\sigma_A^2} (F_A^*(\lambda_t) - \mathbb{E}^{\mathcal{F}_t} [F_A^*(\lambda_{t+dt})] | 1 \text{ activation in } [t, t+dt]) \\ &\quad - 2dt \frac{\theta}{\sigma_A} (F_A^*(\lambda_t) - F_A^*(\lambda^*) + \sigma_A/2 \|\lambda_t - \lambda^*\|^2) - 2 \frac{1}{\sigma_A} (F_A^*(\lambda'_{t+dt}) - F_A^*(\lambda_t)) \end{aligned}$$

Studying $\mathbb{E}^{\mathcal{F}_t} [F_A^*(\lambda_{t+dt})]$, we get:

$$\mathbb{E}^{\mathcal{F}_t} [F_A^*(\lambda_{t+dt})] = (1 - dt)F_A^*(\lambda'_{t+dt}) + dt\mathbb{E}^{\mathcal{F}_t} [F_A^*(\lambda_{t+dt}) - F_A^*(\lambda^*) | 1 \text{ activation in } [t, t+dt]] \quad (34)$$

Using $\theta^2 = \sigma_A/S^2$ (i.e $\theta^2 S^2/\sigma_A^2 = 1/\sigma_A$) and the above equality, Equation (29) become:

$$\begin{aligned} \mathbb{E}^{\mathcal{F}_t} [r_{t+dt}^2] - r_t^2 &\leq -dt\theta r_t^2 \\ &\quad + dt \frac{2}{\sigma_A} (F_A^*(\lambda_t) - \mathbb{E}^{\mathcal{F}_t} [F_A^*(\lambda_{t+dt})] | 1 \text{ activation in } [t, t+dt]) \\ &\quad - 2dt \frac{\theta}{\sigma_A} (F_A^*(\lambda_t) - F_A^*(\lambda^*)) - 2 \frac{1}{\sigma_A} (F_A^*(\lambda'_{t+dt}) - F_A^*(\lambda_t)) \\ &= -dt\theta r_t^2 - \frac{2}{\sigma_A} (\mathbb{E}^{\mathcal{F}_t} [F_A^*(\lambda_{t+dt}) - F_A^*(\lambda^*)] - F_A^*(\lambda_t) - F_A^*(\lambda^*)) \\ &\quad - 2dt \frac{\theta}{\sigma_A} (F_A^*(\lambda_t) - F_A^*(\lambda^*)) + 2dt \frac{\theta}{\sigma_A} (F_A^*(\lambda_t) - F_A^*(\lambda'_{t+dt})) \end{aligned}$$

Since the last term is a $o(dt)$, the previous equation simplifies to:

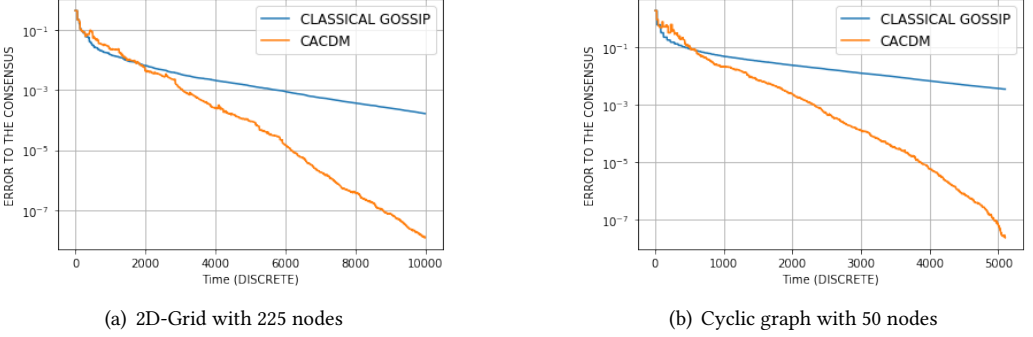
$$\mathbb{E}^{\mathcal{F}_t} [L_{t+dt}] - L_t \leq -\theta dt L_t$$

Finally, we take the expectation, divide by dt and make it tend to zero, which leads to $\frac{d}{dt}\mathbb{E}L_t \leq -\theta\mathbb{E}L_t$. Integrating this leads to the desired result, which writes:

$$\forall t \geq 0, \mathbb{E}L_t \leq \exp(-\theta t)L_0$$

□

Empirical Results: We consider the *P.p.p. model* on two graphs: the circle with 50 nodes and the 2D-Grid with 225 nodes. Our goal is to illustrate how the algorithms compare in a heterogeneous setting. Therefore, in both cases, 10% of the nodes (chosen uniformly at random) have a delay $\tau_i = 100$ time units, while the others have a delay equal to 1 time unit. The delay of an edge (ij) is then $\tau_{ij} = \max(\tau_i, \tau_j)$. Then, we take *Poisson* rates for the edges equal to the inverse of these delays: $p_{ij} = 1/\tau_{ij}$. The local functions for the gossip problems are chosen as $f_i(x) = \|x - c_i\|^2$, with $c_0 = 1$ and $c_i = 0$ otherwise, which is the worst case scenario in terms of mixing). Figure 1 shows the performances of classical (pairwise) gossip and *CACDM* in this setting. We see that *CACDM* is much faster than classical gossip, and that this is true in particular when the eigengap of the graph is small (of order $1/125000$ for our cyclic graph, compared to $1/50000$ for our grid), as predicted by Theorem 2.

Fig. 1. CACDM vs Randomized Gossip in the $P.p.p.$ model.

4 GOSSIP ON LOSS NETWORKS

4.1 Refined Loss-Network Communication Scheme and Detailed Algorithm

The $P.p.p.$ model is particularly amenable to analysis, and helps us understand quantitatively why asynchronous algorithms can outperform synchronous ones, but it assumes that communications and computations are done instantaneously. Thus, actual implementations differ from its underlying assumptions, unless further synchrony is assumed [14]. To alleviate this issue, with pairwise communications ruled by point processes as a baseline, we consider a protocol in which nodes are tagged as *busy* when they are already engaged in an update, and communications between busy nodes are forbidden. Our model is inspired from classical Loss Network models [16]. In our new model, edges are activated following the same procedure as in the $P.p.p.$ model, with a $P.p.p.$ of intensity p_{ij} . Note that we do not consider these intensities to be constraints of the problem, but rather parameters of the algorithm, that we tune below. Each node has an exponential clock of intensity $\frac{1}{2} \sum_{j \sim i} p_{ij}$. At each clock-ticking, if i is not busy, it selects a neighbor j with probability $p_{ij} / \sum_{k \sim i} p_{ik}$. Node i first checks if j is currently *busy*, an operation that takes time $\varepsilon \tau_{ij}$ for some small $\varepsilon > 0$ ($\varepsilon \ll 1$ if sending a simple request is much faster than sending a whole vector). If j is not *busy*, i and j compute and exchange information, becoming busy for a duration τ_{ij} . We can think of this procedure as classical gossip on an underlying random graph (Figure 2), that follows a Markov-Chain process if we extend the space of states with the inactivation time. We call our model the **Refined Loss Network Model of parameter ε (RLNM(ε))**. It is *refined* as the operation that consists in checking on its neighbors is not present in classical Loss Networks.

More precisely, asynchronous gossip on the Refined Loss-Network communication model runs as follows: each node has a local clock and a *Poisson Point Process* of intensity $\frac{1}{2} \sum_{j \sim i} p_{ij}$, where, with d_i the degree of node i and $\tau_{\max}(ij) = \max_{kl \sim ij} \tau_{kl}$:

$$p_{ij} = \min \left(\frac{1}{\tau_{\max}(ij)}, \frac{1}{2(\max(d_i, d_j) - 1)} \frac{1}{\tau_{ij}} \right). \quad (35)$$

Let $I = \sum_{ij \in E} p_{ij}$ the global activation intensity. Let node i 's local variable be $x_i(t)$ at time t , and let $v_i(t) = \nabla f_i(x_i(t))$ its convex-conjugate. Note λ_t the edge-dual variable at time t (we have $A\lambda_t = v(t)$). Initialize such that $v(0) = 0$ (and $\lambda_0 = 0$).

- (1) **"Busy-Checking" Operation:** when clock i rings at time t , select $j \sim i$ with probability $\frac{p_{ij}}{\sum_{k \sim i} p_{ki}}$ and check whether j is busy. This operation makes i busy for a timelapse of length $\varepsilon \tau_{ij}$.




Fig. 2. Underlying Markov Process for the Graph: edge (ij) activated at time t implies that while ij busy i.e. between times t and $t + \tau_{ij}$, all edges kl adjacent to ij are unavailable.

- (2) **Gradient Exchange:** if neighbor j (chosen at the previous step) is not busy, make both nodes busy for a time τ_{ij} , and i sends $\nabla f_i^*(v_i(t))$ to j (and reciprocally).
- (3) **Gradient Step:** when i receives gradient $\nabla f_j^*(v_j)$ from j , it updates its local value v_i using the following gradient step:

$$v_i(t) \xleftarrow{t} v_i(t) - \frac{\nabla f_i^*(v_i(t)) - \nabla f_j^*(v_j(t))}{\sigma_i^{-1} + \sigma_j^{-1}}. \tag{36}$$

The desired output at node i at time t is then $x_i(t) = \nabla f_i^*(v_i(t))$. Note that in the gossip averaging problem, these operations are equivalent to local averagings as shown in (13). Operations (2) and (3) both happen in the timelapse of length τ_{ij} , thus causing no asynchrony issues and avoiding the need to consider delayed gradients.

4.2 Convergence Results

Define the following constants, where p_{ij} is set as in (35) and I defined in (11):

$$\begin{cases} \tilde{\tau}_{ij} &= (1 + \epsilon)p_{ij}^{-1} \\ \tilde{\tau}_{\max} &= \max_{(ij) \in E} \tilde{\tau}_{ij} \\ T &= \frac{2 \log(6|E|)}{\log(1 - (1 - e^{-1})e^{-1})} I \tilde{\tau}_{\max} \end{cases}. \tag{37}$$

Define for $k \in \mathbb{N}$, $\mathcal{E}_k = F^*(v(t_k)) - F^*(v^*)$, where $v^* = A\lambda^*$ is a minimizer of F^* on $\text{Im}(A)$, λ^* being a minimizer of F_A^* and $t_k \in \mathbb{R}^+$ is the time of the k -th activation. For $k \in \mathbb{N}$, let \mathcal{L}_k be the following Lyapunov function:

$$\mathcal{L}_k = \frac{1}{T} \sum_{l=k}^{k+T-1} \mathcal{E}_l. \tag{38}$$

This choice of Lyapunov function is motivated by the fact that we want to take into account T successive values of \mathcal{E}_l (the dual error to the optimum), where T is the typical number of activations required to have all edges activated. Note that this Lyapunov function bears some resemblance with

Lyapunov-Krasovskii functionals (see e.g. [12]) used in the study of delayed differential systems, and which can be thought of as the continuous analog of \mathcal{L}_k , with an integral instead of a sum. We insist on the fact that considering this specific Lyapunov is a key step of our proof.

THEOREM 3 (DISCRETE-TIME RATE OF CONVERGENCE IN THE LOSS-NETWORK MODEL). *Consider the CDM algorithm (36), with node activations according to the RLNM with Poisson rates (35). Let $\Gamma_{RLNM} = \gamma(v_{ij})$ be the spectral gap of the weighed graph Laplacian with weights*

$$v_{ij} = \alpha \times \frac{\tilde{\tau}_{ij}^{-1} \min_{(kl) \sim (ij)} \frac{\tilde{\tau}_{ij}}{\tilde{\tau}_{kl}}}{d_{\max}^2 (\log(|E|) + \log(I\tilde{\tau}_{\max}))^2},$$

where $\alpha = \frac{32e^2}{\log(1-(1-e^{-1})e^{-1})^2}$ is a universal constant and d_{\max} is the maximal degree in the graph. Then, for all $k \in \mathbb{N}$:

$$\mathbb{E}[\mathcal{L}_k] \leq \left(\frac{1}{4} \left(1 - \frac{\sigma_{\min}}{L_{\max}} \Gamma_{RLNM} \right)^{T/3} + \frac{3}{4} \right)^{\lceil \frac{k}{2T} \rceil} \mathbb{E}[\mathcal{L}_0].$$

where Lyapunov function \mathcal{L}_k is defined in (38).

Theorem 3 gives precise results in a general setting, but it may be hard to parse. In order to present results in a more concise form, we introduce the simplifying Assumption 1, which in particular allows to obtain an asymptotic rate of convergence for \mathcal{E}_k .

ASSUMPTION 1 (DELAY CONSTRAINTS). *Let $\gamma_1 = \gamma(v_{ij})$ for $v_{ij} \equiv 1, (ij) \in E$ (Definition 1). Assume that:*

$$\frac{\tilde{\tau}_{\max}}{\tilde{\tau}_{\min}} \leq \frac{L_{\max}}{\sigma_{\min}} \times \frac{\alpha d_{\max}^2 \log(|E|)}{\gamma_1}. \quad (39)$$

Notice that the right-hand side of (39) reflects the complexity of the optimization problem through the first factor (generally referred to as the condition number of the optimization problem), and the topology of the graph (without the delays) through γ_1 . The more difficult the problem is, the bigger the right-hand side is. Assumption 1 will then be verified more easily for graphs with slow mixing times (γ_1^{-1} bigger) and less regular local functions. The order of magnitude of γ_1^{-1} is n^2 for the grid, and n for the line or the cyclic graph. More generally, the right-hand side of (39) is always of order bigger than n .

COROLLARY 1 (ASYMPTOTIC RATE). *Under Assumption 1, Theorem 3 gives:*

$$\limsup_{k \rightarrow \infty} \frac{1}{k} \log(\mathbb{E}[\mathcal{E}_k]) \leq -\frac{\sigma_{\min}}{L_{\max}} \times \frac{\Gamma_{RLNM}}{24e}.$$

Comments on the convergence rate: Theorem 3 and Corollary 1 are formulated in discrete time. The continuous exponential rate of convergence is obtained by multiplying by the global *P.p.p.* intensity I , up to a constant factor of order 1. The factor $\frac{1}{I}$ in the definition (3) of the weights v_{ij} is hence simply a normalization factor, due to a study in discrete time. As desired, the communication cost factor in the rate of convergence (Γ_{RLNM}) is captured by the Laplacian of the graph, weighted by *local* delays, instead of τ_{\max}^{-1} . We however observe slowdowns due to other factors.

- (1) Having $\tilde{\tau}_{ij}$ instead of τ_{ij} (as in the *P.p.p. model* (10)) means that the effective waiting time of edge ij between two activations is of order $\tilde{\tau}_{ij}$ (defined in (37)) and not τ_{ij} , which was expected since p_{ij} is tuned accordingly.
- (2) Adding the factor $\min_{(kl) \sim (ij)} \frac{\tilde{\tau}_{ij}}{\tilde{\tau}_{kl}}$ to the local weight in the Laplacian is a local slowdown: a node with a slow neighbor becomes less effective.

These first two remarks 1) and 2) suggest that by deleting some edges one could improve the rate of convergence. A similar phenomenon occurs in road-trafficking [3, 40], where deleting some roads can lead to reduced congestion (*Braess's paradox*).

- (3) The global factor $\frac{1}{d_{\max}}$ is not intuitive at first: the more connected the graph is, the higher the rate should be. We hence have a trade-off between $\frac{1}{d_{\max}}$ that decreases when adding edges, and the smallest eigenvalue of the Laplacian of the graph Γ that increases with connectivity. We believe that $\frac{1}{d_{\max}}$ is an artifact of the proof, but have not been able so far to remove it.
- (4) If some nodes are *stragglers* (i.e. with high delays compared to the others), the rate of convergence stated for RLNM improves over synchronous algorithms, as it takes into account *local* delays. If the delays are all of the same order of magnitude, a case favorable to synchrony, the rate obtained is the same as in synchronous algorithms, up to a factor of order $\frac{1}{d^2 \log(n)}$. The factor d^2 should not be of too much importance in d -regular graphs for $d \ll n$, such as grids or lines. The log factor comes from exponential tails of our random variables.

REMARK 4 (COMPARISON WITH A DELAYED INFORMATION APPROACH). *One may wonder how our model compares to a delayed information approach, in which nodes send gradients whenever they can. In the delayed information approach, delays increase the variance of the gradients, typically by a multiplicative factor τ equal to the discrete-time delay [13, 19] thus requiring step sizes to be scaled by the inverse of the delays. However, the few works done in this direction rely on a global upper-bound τ_{\max} on the delays, and as such provide slow rates in scenarios with heterogeneous local delays, compared to those achievable with our RLNM approach. Developing delayed information schemes that are competitive in heterogeneous scenarios is an open research direction.*

4.3 Sketch of Proof of Theorem 3

This proof follows three main steps: *i)* Deriving convergence results for more general communication schemes than RLNM, under deterministic assumptions on the delays. *ii)* Adapting Step *i)* to stochastic assumptions on the delays. *iii)* Deriving high-probability upper-bounds on the delays between two activations in RLNM in order to fall under the assumptions of Step *i)*.

As in the previous proofs, the analysis is done with edge-dual variable $\lambda_t \in \mathbb{R}^{E \times d}$, such that $A\lambda_t = v(t)$. Matrix A is tuned in the detailed proof (Appendix C). When nodes $i \sim j$ exchange gradients, it is equivalent to, on edge (ij) :

$$\lambda_t \leftarrow \lambda_t - \frac{1}{\mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})} \nabla_{ij} F_A^*(\lambda_t). \quad (40)$$

4.3.1 *Step 1: General Communication Schemes.* We consider general activation processes \mathcal{P}_{ij} . When edge (i, j) is activated, the update described in (36) is performed at nodes i and j . The delay of an edge is defined as its (random) waiting time between two activations. Two ergodicity-like conditions on the delays are needed: (i) edges activated regularly enough and (ii) incident edges must not be activated too many times. We now formally introduce these assumptions. We consider discrete time in this section: more precisely, $t \in \mathbb{N}$ stands for the t -th edge activation.

DEFINITION 2. *Consider a communication scheme with edge-activation point processes \mathcal{P}_{ij} . Let $t = 0, 1, 2, \dots$ index the consecutive edge activations. Let $s \in \mathbb{N}$, ij and $kl \in E$. Let $s_{ij} < t_{ij}$ such that $s_{ij} \leq s < t_{ij}$ be consecutive activation times (in discrete time) of (ij) . Denote $T_{ij}(s) = t_{ij} - s_{ij} - 1$ the total number of edge activations between the two consecutive activations of ij . Denote $N(kl, ij, s)$ the number of activations of edge kl in the activations $\{s_{ij}, s_{ij} + 1, \dots, t_{ij} - 1\}$.*

ASSUMPTION 2 (DELAY ASSUMPTIONS). *There exist $T \in \mathbb{N}^*$, $a, b > 0$, and $\ell_{ij} > 0$, $ij \in E$ such that, for the quantities and the communication scheme in Definition 2:*

- (1) For all $t \in \mathbb{N}$, all edges are activated between iterations t and $t + T - 1$.
- (2) $\forall s \geq 0, \forall (ij) \in E, T_{ij}(s) \leq a\ell_{ij}$: (ij) is activated at least every $a\ell_{ij}$ activations.
- (3) $\forall s \geq 0, \forall (ij), (kl) \in E$ such that $(kl) \sim (ij)$, $N(kl, ij, s) \leq \lceil \frac{b\ell_{ij}}{\ell_{kl}} \rceil$.

Assumption (1) is implied by Assumption (2) if $T = \max_{(ij)} \ell_{ij}$. Taking ℓ_{ij} as a deterministic upper-bound on the delays of edge (ij) between two activations in continuous time is sufficient to have Assumption (2) and (3), with some normalizing constant a , and b such that ℓ_{ij}/b is a lower-bound on these delays.

The main technical difficulty lies in the fact that at a defined activation time t , some nodes are not available: at any time $t \geq 0$, $\sum_{(ij) \in E \text{ not busy}} \nabla_{ij} F_A^*(\lambda_t)$ usually differs from $\nabla F_A^*(\lambda_t)$ as in *Markov-Chain Gradient Descent* [41], thus making an analysis such as in the *P.p.p. model* impossible. To alleviate this difficulty, in order to make sure that all edges are taken into account when performing the averaging, the Lyapunov function Λ_t that we study considers the value of the objective for T consecutive activation times. It is defined as follows on the dual variable: $\forall t \in \mathbb{N}, \Lambda_t = \frac{1}{T} \sum_{s=t}^{t+T-1} F_A^*(\lambda_s) - F_A^*(\lambda^*)$. Note that we have $\Lambda_t = \mathcal{L}_t$ for any $t \in \mathbb{N}$, \mathcal{L}_t as in (38): we simply changed notations as we work with edge-dual variables, and time is indexed in a different way. The first step of the proof of Theorem 3 consists in proving the following. A detailed proof of this can be found in Appendix C.1.

THEOREM 4. *Consider a general communication scheme as in Definition 2, that satisfies Assumption 2 for constants $\ell_{ij}, a, b > 0$. At every edge-activation of edge (ij) , update (40) is performed. Let γ be the smallest positive eigenvalue of the Laplacian of the graph with:*

$$v_{ij} = C\ell_{ij}^{-1} \min_{kl \sim ij} \frac{\ell_{kl}}{\ell_{ij}},$$

where $C = \frac{1}{2a+8d_{\max}^2 ab}$. Then, we have, for $t \in \mathbb{N}$:

$$\Lambda_t \leq \left(1 - \frac{\sigma_{\min}}{L_{\max}} \times \gamma\right)^t \Lambda_0.$$

4.3.2 Step 2: Introducing Stochasticity. Theorem 4 cannot be applied directly to *RLNM* since we have unbounded delays. Yet, Theorem 4 can be adapted to hold with relaxed assumptions: the conditions on the delays may only hold with some (not too low) probability instead of almost surely. More precisely, we prove the following in Appendix C.2.

PROPOSITION 3 (ADDING STOCHASTICITY). *Assume that, for all $t \in \mathbb{N}$, there exists a \mathcal{F}_{t+T-1} -measurable event A_t , such that $\mathbb{P}(A_t | \mathcal{F}_t) \geq \frac{1}{2}$ almost surely, and that under A_t , Assumption 2 holds for $t \leq s \leq t + T - 1$. Then, we have the following bound on L_t :*

$$\mathbb{E}[\Lambda_t] \leq \left(\frac{1}{4} \left(1 - \frac{\sigma_{\min}}{L_{\max}} \gamma\right)^{T/3} + \frac{3}{4}\right)^{\lceil \frac{t}{2T} \rceil} \mathbb{E}[\Lambda_0].$$

This proposition enables us to apply Theorem 4 to stochastic communication schemes that have unbounded yet stochastically controlled delays. This result and its proof are thus of independent interest: it encompasses more general communication schemes than *RLNM*. Furthermore, the methodology of this deterministic to stochastic conversion could be applied more generally to other problems.

4.3.3 Step 3: Controlling Inactivation Times in *RLNM*(ϵ). After studying general deterministic (Section 4.3.1) and then stochastic communication schemes (Section 4.3.2), we place ourselves back in the *RLNM*(ϵ) model. The following lemma controls how long a given edge can remain inactive in

our model, which is a key step of our analysis. Indeed, it allows us to specify the constants ℓ_{ij} , T , a , and b from Assumption 2 such that Proposition 3 can be applied.

LEMMA 2. For any $t_0 \geq 0$, $ij \in E$, if the Poisson intensities are such that $p_{ij} = \frac{1}{2^{\max(d_i, d_j)-1}}((1 + \varepsilon)\tau_{ij})^{-1}$ and $\tau_{\max}(ij) = \max_{kl \sim ij} \tau_{kl}$, let:

$$\ell_{ij} = \frac{\log(\delta^{-1})}{\log(1 - (1 - e^{-1})e^{-1})} (p_{ij}^{-1} + \tau_{\max}(ij))(1 + \varepsilon)$$

for any $\delta \in (0, 1)$. We have:

$$\mathbb{P}(ij \text{ not activated in } [t_0, t_0 + \ell_{ij}] | \mathcal{F}_{t_0}) \leq \delta. \quad (41)$$

PROOF OF LEMMA 2. Let $ij \in E$ and $t_0 \geq 0$ fixed. We use tools from queuing theory [44] ($M/M/\infty/\infty$ queues) in order to compute the probability that edge ij is activable at a time t or not. More formally, we define a process $N_{ij}(t)$ with values in \mathbb{N} , such that $N_{ij}(t_0) = 1$ if ij non-available at time t_0 and 0 otherwise. Then, when an edge kl such that $kl \sim ij$ is activated, we make an increment of 1 on $N_{ij}(t)$ (a customer arrives). This customer stays for a time $\tau_{kl}(1 + \varepsilon)$ and when he leaves, N_{ij} is decreased by 1. Thus $N_{ij} \geq 0$ a.s., and if $N_{ij} = 0$, then edge ij is available. For $t \geq \max_{kl \sim ij} \tau_{kl}(1 + \varepsilon) + t_0$, $N_{ij}(t)$ follows a Poisson law of parameter $\sum_{kl \sim ij} p_{kl} \tau_{kl}(1 + \varepsilon)$. For any $t \geq \max_{kl \sim ij} \tau_{kl}(1 + \varepsilon) + t_0$:

$$\mathbb{P}(ij \text{ available at time } t | \mathcal{F}_{t_0}) \geq \mathbb{P}(N_i(t) = 0) = \exp\left(-\sum_{kl \sim ij} p_{kl} \tau_{kl}(1 + \varepsilon)\right).$$

That leads to taking $p_{kl} = \frac{1}{2^{\max(d_k, d_l)-1}}((1 + \varepsilon)\tau_{kl})^{-1}$ for all edges, in order to have

$$\mathbb{P}(ij \text{ available at time } t | \mathcal{F}_{t_0}) \geq 1/e.$$

Then, $\mathbb{P}(ij \text{ rings in } [t, t + p_{ij}^{-1}]) = 1 - e^{-1}$, giving:

$$\begin{aligned} \mathbb{P}(ij \text{ activated in } [t_0, t_0 + (1 + \varepsilon)\tau_{\max}(ij) + p_{ij}^{-1}] | \mathcal{F}_{t_0}) &= \mathbb{P}(ij \text{ rings in } [t, t + p_{ij}^{-1}]) \\ &\times \mathbb{P}(ij \text{ available at time } t | \mathcal{F}_{t_0}, ij \text{ rings at a time } t \in [t_0 + (1 + \varepsilon)\tau_{\max}(ij), t_0 + (1 + \varepsilon)\tau_{\max}(ij) + p_{ij}^{-1}]) \\ &\geq (1 - e^{-1})e^{-1}, \end{aligned}$$

where we use the memoriless property of exponential random variables. Take $k \in \mathbb{N}$ such that $(1 - (1 - e^{-1})e^{-1})^k \leq \delta$, leading to $k = \log(6|E|)/\log(1 - (1 - e^{-1})e^{-1})$. Let

$$\ell_{ij} = k(p_{ij}^{-1} + \tau_{\max}(ij))(1 + \varepsilon).$$

Then we have a.s.:

$$\mathbb{P}(ij \text{ not activated in } [t_0, t_0 + \ell_{ij}] | \mathcal{F}_{t_0}) \leq \delta. \quad (42)$$


□

We then use this lemma in Appendix C.3 in order to tune the constants of Assumption 2 for *RLNM*.


4.4 Empirical Results

The results in Figure 3 correspond to the Loss-Network scheme on the same two heterogeneous graphs (50-node cycle and 225-node 2D-grid) as in Figure 1. We compare our algorithm on the Loss-Network to synchronous gossip. Time is indexed in a continuous way. Synchronous iterations are done every 100 units of time. The speed-up is significant when the fluctuation in term of delays in the graph is high, which illustrates the discussion at the end of Section 2.1.

Acceleration in *RLNM*(ε): The analysis of *CACDM* does not extend to more general models than the *P.p.p. model*. However, applying it to *RLNM* leads to an accelerated rate of convergence displayed in Figure 4, showing that our algorithm is quite robust to changes in edge activation statistics. In



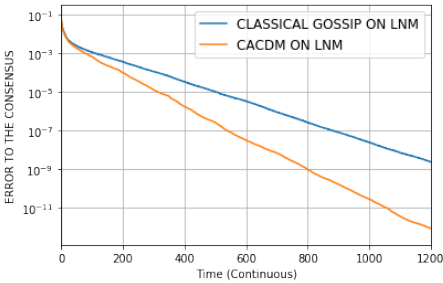
(a) 2D-Grid with 225 nodes




(b) Cyclic graph with 50 nodes

Fig. 3. Asynchronous Speed-Up: Classical synchronous gossip (Appendix A.1) VS Gossip on *RLNM*.

order to tune the algorithm, we take values p_{ij} as in (35). Time is indexed in a continuous way. 1000 units of time hence correspond to approximately $I \times 1000 \approx 10^5 - 10^6$ edge activations.



(a) 2D-Grid with 225 nodes



(b) Cyclic graph with 50 nodes

Fig. 4. *CACDM* vs Gossip in *RLNM*.

5 CONCLUSION

We studied asynchronous gossip algorithms in two frameworks: the popular *P.p.p. model* and the *refined loss network* model, a contribution of this paper. For the simple *P.p.p. model* of asynchronous operations we developed a novel analysis in continuous time of gradient descent which then enabled us to propose *CACDM*, a provably accelerated version of classical randomized gossip. *RLNM*, our refined model of asynchronous communications, provides a more realistic model of asynchrony than the *P.p.p. model*, as well as a framework that avoids the need to rely on delayed information. We obtained convergence rate guarantees for the *CDM* scheme under this model, that highlight the role of quantities such as local effective delays, local differences of delays, and node degrees. An interesting open question is whether our established rates of convergence enjoy some form of optimality, or how fundamental the local effective delays we identified, and the spectral gap of the associated weighted graph Laplacian, are intrinsic bottlenecks for the performance of asynchronous distributed optimization. We believe that both our main contributions (*CACDM* and *RLNM*) pave the way for fast asynchronous gossip algorithms with theoretical guarantees.

REFERENCES

- [1] Assran, M., Aytekin, A., Feyzmahdavian, H., Johansson, M., and Rabbat, M. (2020). Advances in asynchronous parallel and distributed optimization.
- [2] Aybat, N. and Gürbüzbalaban, M. (2017). Decentralized computation of effective resistances and acceleration of consensus algorithms.
- [3] Bean, N. G., Kelly, F. P., and Taylor, P. G. (1997). Braess’s paradox in a loss network. *Journal of Applied Probability*, 34(1):155–159.
- [4] Berthier, R., Bach, F., and Gaillard, P. (2018). Accelerated gossip in networks of given dimension using jacobi polynomial iterations.
- [5] Boyd, S., Ghosh, A., Prabhakar, B., and Shah, D. (2006). Randomized gossip algorithms.
- [6] Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3:1–122.
- [7] Bubeck, S. (2014). Convex optimization: Algorithms and complexity.
- [8] Colin, I., Bellet, A., Salmon, J., and Cléménçon, S. (2016). Gossip Dual Averaging for Decentralized Optimization of Pairwise Functions. In *International Conference on Machine Learning (ICML 2016)*, New York, United States.
- [9] Dimakis, A. D. G., Sarwate, A. D., and Wainwright, M. J. (2008). Geographic gossip: Efficient averaging for sensor networks. *IEEE Transactions on Signal Processing*, 56(3):1205–1216.
- [10] Dimakis, A. G., Kar, S., Moura, J. M. F., Rabbat, M. G., and Scaglione, A. (2010). Gossip algorithms for distributed signal processing. *Proceedings of the IEEE*, 98(11):1847–1864.
- [11] Duchi, J. C., Agarwal, A., and Wainwright, M. J. (2012). Dual averaging for distributed optimization: Convergence analysis and network scaling. *IEEE Transactions on Automatic Control*, 57(3):592–606.
- [12] Fridman, E. (2001). New lyapunov–krasovskii functionals for stability of linear retarded and neutral type systems. *Systems & Control Letters*, 43(4):309 – 319.
- [13] Hannah, R., Feng, F., and Yin, W. (2018). A2bcd: An asynchronous accelerated block coordinate descent algorithm with optimal complexity.
- [14] Hendrikx, H., Bach, F., and Massoulié, L. (2018). Accelerated decentralized optimization with local updates for smooth and strongly convex objectives.
- [15] Kakade, S., Shalev-Shwartz, S., and Tewari, A. (2009). On the duality of strong convexity and strong smoothness: Learning applications and matrix regularization. *Unpublished Manuscript*, <http://ttic.uchicago.edu/shai/papers/KakadeShalevTewari09.pdf>.
- [16] Kelly, F. P. (1991). Loss networks. *The Annals of Applied Probability*, 1(3):319–378.
- [17] Klenke, A. (2014). *The Poisson Point Process*, pages 543–561. Springer London, London.
- [18] Le Gall, J.-F. (2016). *Brownian Motion, Martingales, and Stochastic Calculus*, volume 274.
- [19] Leblond, R., Pedregosa, F., and Lacoste-Julien, S. (2016). Asaga: Asynchronous parallel saga.
- [20] Lee, S. and Nedich, A. (2013). Asynchronous gossip-based random projection algorithms over networks.
- [21] Lian, X., Zhang, C., Zhang, H., Hsieh, C.-J., Zhang, W., and Liu, J. (2017a). Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent.
- [22] Lian, X., Zhang, W., Zhang, C., and Liu, J. (2017b). Asynchronous decentralized parallel stochastic gradient descent.
- [23] Liu, J., Anderson, B. D., Cao, M., and Morse, A. S. (2013). Analysis of accelerated gossip algorithms. *Automatica*, 49(4):873 – 883.
- [24] Liu, J. and Wright, S. J. (2014). Asynchronous stochastic coordinate descent: Parallelism and convergence properties.
- [25] Loizou, N., Rabbat, M., and Richtárik, P. (2018). Provably accelerated randomized gossip algorithms.
- [26] Loizou, N. and Richtárik, P. (2018). Accelerated gossip via stochastic heavy ball method.
- [27] Mania, H., Pan, X., Papailiopoulos, D., Recht, B., Ramchandran, K., and Jordan, M. I. (2015). Perturbed iterate analysis for asynchronous stochastic optimization.
- [28] Mohar, Alavi, Y., Chartrand, G., and Oellermann, O. (1991). The laplacian spectrum of graphs. *Graph theory, combinatorics, and applications*.
- [29] Mokhtari, A. and Ribeiro, A. (2016). Dsa: Decentralized double stochastic averaging gradient algorithm. *Journal of Machine Learning Research*, 17(61):1–35.
- [30] Montijano, E., Montijano, J., and Sagues, C. (2011). Chebyshev polynomials in distributed consensus applications. *IEEE Transactions on Signal Processing*, 61.
- [31] Nedic, A. and Ozdaglar, A. (2009). Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61.
- [32] Nedich, A., Olshevsky, A., and Shi, W. (2016). Achieving geometric convergence for distributed optimization over time-varying graphs.
- [33] Nesterov, Y. and Stich, S. U. (2017). Efficiency of the accelerated coordinate descent method on structured optimization problems. *SIAM Journal on Optimization*, 27(1):110–123.

- [34] Niu, F., Recht, B., Re, C., and Wright, S. J. (2011). Hogwild!: A lock-free approach to parallelizing stochastic gradient descent.
- [35] Oreshkin, B. N., Coates, M. J., and Rabbat, M. G. (2010). Optimization and analysis of distributed averaging with short node memory. *IEEE Transactions on Signal Processing*, 58(5):2850–2865.
- [36] Pu, S., Shi, W., Xu, J., and Nedic, A. (2020). Push-pull gradient methods for distributed optimization in networks. *IEEE Transactions on Automatic Control*, pages 1–1.
- [37] Ram, S., Nedić, A., and Veeravalli, V. (2009). Asynchronous gossip algorithms for stochastic optimization. In *Proceedings of the 48th IEEE Conference on Decision and Control held jointly with 2009 28th Chinese Control Conference, CDC/CCC 2009*, Proceedings of the IEEE Conference on Decision and Control, pages 3581–3586. 48th IEEE Conference on Decision and Control held jointly with 2009 28th Chinese Control Conference, CDC/CCC 2009 ; Conference date: 15-12-2009 Through 18-12-2009.
- [38] Scaman, K., Bach, F., Bubeck, S., Lee, Y. T., and Massoulié, L. (2017). Optimal algorithms for smooth and strongly convex distributed optimization in networks.
- [39] Shah, D. (2009). Gossip algorithms. *Foundations and Trends® in Networking*, 3(1):1–125.
- [40] Steinberg, R. and Zangwill, W. I. (1983). The prevalence of braess’ paradox. *Transportation Science*, 17(3):301–318.
- [41] Sun, T., Sun, Y., and Yin, W. (2018). On markov chain gradient descent.
- [42] Sundhar Ram, S., Nedic, A., and Veeravalli, V. V. (2009). Asynchronous gossip algorithms for stochastic optimization. In *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pages 3581–3586.
- [43] Tang, H., Lian, X., Yan, M., Zhang, C., and Liu, J. (2018). D^2 : Decentralized training over decentralized data.
- [44] Tanner, M. (1995). *Practical queueing analysis*. IBM McGraw-Hill. McGraw-Hill, London.
- [45] Uribe, C. A., Lee, S., Gasnikov, A., and Nedić, A. (2020). A dual approach for optimal algorithms in distributed optimization over networks.
- [46] Zhang, Y., Duchi, J. C., and Wainwright, M. J. (2013). Communication-efficient algorithms for statistical optimization. *Journal of Machine Learning Research*, 14(68):3321–3363.
- [47] Zheng, S., Meng, Q., Wang, T., Chen, W., Yu, N., Ma, Z.-M., and Liu, T.-Y. (2016). Asynchronous stochastic gradient descent with delay compensation.

SUMMARY OF THE ARTICLE AND OF THE APPENDIX

Abstract	1
1 Introduction	1
1.1 Main Contributions	2
1.2 Related Work	2
2 Problem Formulation and Notations	4
2.1 Basic assumptions and notations	4
2.2 Dual Formulation of the Problem	5
2.3 Gossip Averaging Problem	6
3 Randomized Gossip: the <i>P.p.p. model</i>	7
3.1 The <i>P.p.p. Model</i> and Randomized Gossip Algorithms	7
3.2 Continuous Time Convergence Analysis	7
3.3 Accelerated Gossip in the <i>P.p.p. model</i>	9
4 Gossip on Loss Networks	14
4.1 Refined Loss-Network Communication Scheme and Detailed Algorithm	14
4.2 Convergence Results	15
4.3 Sketch of Proof of Theorem 3	17
4.4 Empirical Results	19
5 Conclusion	20
References	21
Summary of the Article and of the Appendix	22
A Gossip Algorithms: General Considerations on the Averaging Problem	23
A.1 Synchronous Gossip	23

A.2	Asynchronous Gossip	24
A.3	Laplacian Monotonicity	24
B	Preliminary Inequalities	25
C	Detailed Proof of Theorem 3	27
C.1	Proof Of Theorem 4	27
C.2	Proof Of Proposition 3: Adding Stochasticity	29
C.3	Study in the $RLNM(\varepsilon)$: Tuning the Parameters	30
C.4	Proof of Corollary 1	33

A GOSSIP ALGORITHMS: GENERAL CONSIDERATIONS ON THE AVERAGING PROBLEM

A.1 Synchronous Gossip

In the synchronous setting, all nodes are allowed to share a common clock, which enables them to perform operations synchronously. Formally, a *gossip matrix* is defined as follows:

DEFINITION 3 (GOSSIP MATRIX). *A gossip matrix is a matrix $W \in \mathbb{R}^{n \times n}$ such that:*

- $\forall (i, j) \in [n]^2, W_{i,j} > 0 \implies i \sim j \text{ or } i = j$ (supported by G),
- $\forall i \in [n], \sum_{j \sim i} W_{i,j} = 1$ (stochastic),
- $\forall (i, j) \in [n]^2, W_{i,j} = W_{j,i}$ (symmetric).

Iteratively, at times $t = 0, 1, 2, \dots$, if $x(t) = (x_i(t))_i \in \mathbb{R}^{n \times d}$ describes the information stacked locally at each node ($x_i(t)$ being the vector at node i), we perform the operation $x(t+1) = Wx(t)$. It is to be noted that, thanks to the sparsity of the gossip matrix, this operation is local: for all node i ,

$$x_i(t+1) = \sum_{j \sim i} W_{i,j} x_j(t), \quad (43)$$

where $i \sim j$ if they are neighbors or if $i = j$. The convergence bound will be stated below. Intuitively, at each iteration, each node i sends a proportion of its mass to each one of its neighbour, the condition $\sum_{j \sim i} W_{i,j} = 1$ being the mass conservation.

PROPOSITION 4 (SYNCHRONOUS GOSSIP). *Let γ_W be the eigengap of the laplacian of G weighted by $1 - W_{ij}$ at each edge. Then, for all $k = 0, 1, 2, \dots$:*

$$\|x(k) - \bar{c}\| \leq (1 - \gamma_W)^k \|c - \bar{c}\|, \quad (44)$$

where $x(0) = c$, and \bar{c} is when consensus is reached

PROOF. For $k \geq 0$,

$$\begin{aligned} x(k+1) - \bar{c} &= W(x(k) - \bar{c}) \\ \implies \|x(k+1) - \bar{c}\| &\leq \lambda_2(W) \|x(k) - \bar{c}\|, \end{aligned}$$

where λ_2 is the second largest eigenvalue of W , 1 being the largest (W is stochastic symmetric), and \bar{c} being in the corresponding eigenspace. We conclude by saying that $\lambda_2(W) = 1 - \gamma_W$ where γ_W is the smallest non null eigenvalue of $Id - W$. Notice that $Id - W$ is the laplacian of the graph weighted by $v_{ij} = 1 - W_{ij}$. \square

Then, since every iteration takes a time τ_{max} , denoting time in a continuous way by $t \in \mathbb{R}^+$, we have:

$$\|x(t) - \bar{c}\| \leq (1 - \gamma_W)^{t/\tau_{max}-1} \leq \exp\left(-\frac{\gamma_W}{\tau_{max}}(t - \tau_{max})\right), \quad (45)$$

and $\gamma_W/\tau_{max} \leq \gamma_{synch}$ where γ_{synch} is the smallest non-null eigenvalue of the laplacian of the graph with weights $v_{ij} = \tau_{max}$.

A.2 Asynchronous Gossip

Time is indexed in a continuous way, by \mathbb{R}^+ . For every edge $e = (ij) \in E$, let \mathcal{P}_{ij} be a Poisson point process (P.p.p.) of constant intensity $p_{ij} > 0$ that we will call "clocks", all independent from each other. Updates will be ruled by these processes: at every clock ticking of \mathcal{P}_{ij} , nodes i and j update the value they stack by the mean $\frac{x_i+x_j}{2}$. If we write $\mathcal{P} = \bigcup_{ij \in E} \mathcal{P}_{ij}$, \mathcal{P} is a P.p.p. of intensity $I := \sum_{ij \in E} p_{ij}$.

PROPOSITION 5 (ASYNCHRONOUS CONTINUOUS TIME BOUND). *Let $(x_t(i))_i$ be the vector stacked on the graph, and $\bar{c} = (\frac{1}{n} \sum_i c_i, \dots, \frac{1}{n} \sum_i c_i)^\top$ the consensus, where $c_i = x_i(0)$. Let σ_{asynch} be the smallest non null eigenvalue of the laplacian of the graph, weighted by the p_{ij} 's. For $t \geq 0$, we have:*

$$\mathbb{E}[\|x(t) - \bar{c}\|^2] \leq \exp(-t\sigma_{asynch})\|c - \bar{c}\|^2.$$

PROOF. First, it is to be noted that, if \mathcal{P} is a P.p.p. of intensity $\lambda > 0$, for all $t \in \mathbb{R}$ and $dt \rightarrow 0$:

$$\mathbb{P}([t, t + dt] \cap \mathcal{P} \neq \emptyset) = \lambda dt + o(dt). \quad (46)$$

When ij activated at time t , multiply $x(t)$ by $W_{ij} = I_n - \frac{{}^t(e_i - e_j)(e_i - e_j)}{2}$. By observing that $W_{ij}^2 = W_{ij}$ and that $\sum_{ij} p_{ij} W_{ij} = II_n - L$, where L is the laplacian of the graph weighted by the p_{ij} , we get that, with $R_t^2 = \|x(t) - \bar{c}\|^2$ the squared error to the consensus at time t , up to a $o(dt)$:

$$\begin{aligned} \mathbb{E}^{\mathcal{F}_t}[R_{t+dt}^2] &= (1 - Idt)\mathbb{E}^{\mathcal{F}_t}[R_{t+dt}^2 | \text{no activations in } [t, t + dt]] \\ &\quad + dt \sum_{ij} p_{ij} \mathbb{E}^{\mathcal{F}_t}[R_{t+dt}^2 | ij \text{ activated in } [t, t + dt]] + o(dt) \\ &= R_t^2 - dt(x(t) - \bar{c})^\top \sum_{ij} W_{ij}(x(t) - \bar{c}) \\ &\leq R_t^2 - dt\sigma_p R_t^2. \end{aligned}$$

Then, taking the mean, dividing by $dt \rightarrow 0$ and integrating concludes the proof. \square

A.3 Laplacian Monotonicity

We finish by proving the following intuitive result:

PROPOSITION 6 (MONOTONICITY OF THE LAPLACIAN). *Let $\Lambda(\lambda_{ij}, (ij) \in E)$ be the laplacian of the graph weighted by λ_{ij} . Then, its second smallest eigenvalue σ is a non decreasing function of each weight λ_{ij} .*

PROOF. First compute $\langle \Lambda u, u \rangle$, the weights λ_{ij} being fixed:

$$\begin{aligned} \langle \Lambda u, u \rangle &= \sum_i \sum_{j \sim i} u_i(u_i - u_j)\lambda_{ij} \\ &= \frac{1}{2} \sum_i \sum_{j \sim i} (u_i - u_j)^2 \lambda_{ij}. \end{aligned}$$

It appears that for any $u \in \mathbb{R}^n$, these are non decreasing quantities in each λ_{ij} . If we take Λ and Λ' two laplacians with weights $\lambda_{ij} \leq \lambda'_{ij}$, we get, for all $u \in \mathbb{R}^n$, $\langle \Lambda u, u \rangle \leq \langle \Lambda' u, u \rangle$. Then, using that $\sigma = \min_{\|u\|=1, \langle u, \mathbb{1} \rangle = 0} \langle \Lambda u, u \rangle$ (as $\mathbb{1}$ is a eigenvector associated to the eigenvalue 0), we have $\sigma' \leq \sigma$ the desired result. \square

B PRELIMINARY INEQUALITIES

We first present preliminary inequalities using properties on our function F_A^* . These properties were also proven in Hendrikx et al. [14] (except for Lemma 8) but we present them here for the paper to be self-contained.

LEMMA 3. *Let $x, v \in \mathbb{R}^{n \times d}$ such that $v = \nabla F(x)$ is the dual conjugate. Assume that there exists $\lambda \in \mathbb{R}^{E \times d}$ such that $A\lambda = v$. Let v^* be the minimizer of F^* on $\text{Im}(A) = \text{Vect}((1, \dots, 1)^\top)$, x^* the minimizer of F under consensus constraint and λ^* a minimizer of F_A^* . We have:*

$$\|x - x^*\|^2 \leq \frac{2L_{\max}}{\sigma_{\min}^2} (F^*(v) - F^*(v^*)). \quad (47)$$

PROOF.

$$\begin{aligned} \|x - x^*\|^2 &= \|\nabla F^*(v) - \nabla F^*(v^*)\|^2 \\ &\leq \frac{1}{\sigma_{\min}^2} \|v - v^*\|^2 \text{ (smoothness of } F^*) \\ &= \frac{1}{\sigma_{\min}^2} \|v - v^*\|_{\text{Im}(A)}^2 \\ &\leq \frac{2L_{\max}}{\sigma_{\min}^2} (F^*(v) - F^*(v^*)) \text{ (strong convexity of } F^*). \end{aligned}$$

\square

LEMMA 4. *For $\lambda \in \mathbb{R}^{E \times d}$ and $ij \in E$, we have:*

$$F_A^* \left(\lambda - \frac{1}{\mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})} U_{ij} \nabla_{ij} F_A^*(\lambda) \right) - F_A^*(\lambda) \leq -\frac{1}{2\mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})} \|\nabla_{ij} F_A^*(\lambda)\|^2. \quad (48)$$

PROOF. Let us define $h_{ij} = -\frac{1}{\mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})} U_{ij} \nabla_{ij} F_A^*(\lambda)$.

$$\begin{aligned} F_A^*(\lambda + h_{ij}) - F_A^*(\lambda) &= \sum_k f_k^*((A\lambda)_k + (Ah_{ij})_k) - f_k^*((A\lambda)_k) \\ &= f_i^*((A\lambda)_i + (Ah_{ij})_i) - f_i^*((A\lambda)_i) + f_j^*((A\lambda)_j + (Ah_{ij})_j) - f_j^*((A\lambda)_j), \end{aligned}$$

as (Ah_{ij}) is supported only by coordinates i and j . Moreover, as f_i^* is σ_i -smooth, we have:

$$f_i^*((A\lambda)_i + (Ah_{ij})_i) - f_i^*((A\lambda)_i) \leq \langle \nabla f_i^*((A\lambda)_i), (Ah_{ij})_i \rangle + \frac{\sigma_i^{-1}}{2} \|(Ah_{ij})_i\|^2,$$

and by summing for i and j and noticing that $(Ah_{ij})_i = \mu_{ij} \nabla_{ij} F_A^*(\lambda)$:

$$\begin{aligned} F_A^*(\lambda + h_{ij}) - F_A^*(\lambda) &\leq \langle \nabla_{ij} F_A^*(\lambda), h_{ij} \rangle + \frac{(\sigma_i^{-1} + \sigma_j^{-1}) \mu_{ij}^2}{2} \left(\frac{1}{\mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})} \right)^2 \|\nabla_{ij} F_A^*(\lambda)\|^2 \\ &= -\frac{1}{2\mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})} \|\nabla_{ij} F_A^*(\lambda)\|^2. \end{aligned}$$

\square

LEMMA 5. σ_A the strong convexity parameter of F_A^* on the orthogonal of $\text{Ker}(A)$ is lower bounded by $\lambda_{\min}^+(A^T A)/L_{\max}$, where $\lambda_{\min}^+(A^T A)$ is the smallest non null eigenvalue of $A^T A$.

PROOF. Let $\lambda, \lambda' \in \mathbb{R}^{E \times d}$. By L_i^{-1} and thus L_{\max}^{-1} -strong convexity of f_i^* :

$$f_i^*((A\lambda)_i) - f_i^*((A\lambda')_i) \geq \langle \nabla f_i^*((A\lambda')_i), (A(\lambda - \lambda'))_i \rangle - \frac{1}{2L_{\max}} \|(A(\lambda - \lambda'))_i\|^2$$

Summing over all $i \in [n]$ and using $\nabla F_A^*(\lambda') = \sum_i \nabla f_i^*((A\lambda')_i)$ leads to:

$$\begin{aligned} F_A^*(\lambda) - F_A^*(\lambda') &\geq \langle \nabla F_A^*(\lambda'), \lambda - \lambda' \rangle - \frac{1}{2L_{\max}} \|A(\lambda' - \lambda)\|^2 \\ &\geq \langle \nabla F_A^*(\lambda'), \lambda - \lambda' \rangle - \frac{\lambda_{\min}^+(A^T A)}{2L_{\max}} \|\lambda - \lambda'\|^{*2}. \end{aligned}$$

where $\|\cdot\|^*$ is the euclidian norm on the orthogonal of $\text{Ker}(A)$. □

LEMMA 6. AA^T is the laplacian of the graph G weighted by μ_{ij}^2 on the edges.

PROOF.

$$A^T e_i = \sum_{j \sim i} \mu_{ij} e_{ij}$$

For the diagonal, we have:

$$\begin{aligned} e_i AA^T e_i &= \sum_{k \sim i} \sum_{l \sim i} \mu_{ik} \mu_{il} \langle e_{ik}, e_{il} \rangle \\ &= \sum_{j \sim i} \mu_{ij}^2. \end{aligned}$$

Then, for $i \sim j, i \neq j$:

$$\begin{aligned} e_i AA^T e_j &= \sum_{k \sim i} \sum_{l \sim i} \mu_{ik} \mu_{jl} \langle e_{ik}, e_{jl} \rangle \\ &= \mu_{ij} \mu_{ji} \\ &= -\mu_{ij}^2. \end{aligned}$$

□

LEMMA 7. For $x, x' \in \mathbb{R}^{E \times d}$, and $ij \in E$, we have:

$$\|\nabla_{ij} F_A^*(x) - \nabla_{ij} F_A^*(x')\|^2 \leq 2(\sigma_i^{-1} + \sigma_j^{-1})^2 d_{ij} \mu_{ij}^2 \sum_{(kl) \sim (ij)} \mu_{kl}^2 \|x_{kl} - x'_{kl}\|^2. \quad (49)$$

PROOF. First, notice that $\nabla_{ij} F_A^*(x) = \mu_{ij} (\nabla f_i^*((Ax)_i) - \nabla f_j^*((Ax)_j))$. Then:

$$\begin{aligned} \|\nabla f_i^*((Ax)_i) - \nabla f_j^*((Ax)_j)\| &\leq \sigma_i^{-1} \|(A(x - x'))_i\| \text{ (smoothness)} \\ &\leq \sigma_i^{-1} \left\| \sum_{kl \sim ij} \mu_{kl} (x - x')_{kl} \right\| \\ &\leq \sigma_i^{-1} \sum_{kl \sim ij} \mu_{kl} \|x - x'\|_{kl} \end{aligned}$$

Conclude by taking the square and summing for i and j . □

LEMMA 8 (DISTANCE TO OPTIMUM). For any $\lambda \in \mathbb{R}^{E \times d}$ and for λ^* minimizing F_A^* , we have:

$$F_A^*(\lambda) - F_A^*(\lambda^*) \leq \frac{1}{2\sigma_A} \|\nabla F_A^*(\lambda)\|^2 \quad (50)$$

PROOF. We introduce Bregman divergences, which make the proof straightforward. For ϕ any real-valued function, differentiable, defined on an euclidian space \mathcal{V} , we define its Bregman divergence D_ϕ on \mathcal{V}^2 by:

$$D_\phi(x, y) = \phi(x) - \phi(y) - \langle \nabla \phi(y), x - y \rangle. \quad (51)$$

ϕ is thus L -smooth if and only if $D_\phi \leq LD_{\|\cdot\|^2/2}$. An important equality is the following, under convexity assumption for ϕ :

$$D_\phi(x, y) = D_{\phi^*}(\nabla \phi(y), \nabla \phi(x)). \quad (52)$$

Applying this to $\phi = F_A^*$, $x = \lambda$, $y = \lambda^*$, together with the fact that $(F_A^*)^*$ is σ_A^{-1} -smooth with respect to $\|\cdot\|^{*2}$ [15], the squared norm on the orthogonal of $\text{Ker}(A)$ leads to:

$$D_{F_A^*}(\lambda, \lambda^*) = D_{F_A^{**}}(\nabla F_A^*(\lambda^*), \nabla F_A^*(\lambda)) \leq \frac{1}{\sigma_A} D_{\|\cdot\|^{*2}/2}(\nabla F_A^*(\lambda^*), \nabla F_A^*(\lambda)),$$

and the result follows since $\nabla F_A^*(\lambda^*) = 0$ and $\|\nabla F_A^*(\lambda)\|^{*2} = \|\nabla F_A^*(\lambda)\|^2$. \square

C DETAILED PROOF OF THEOREM 3

C.1 Proof Of Theorem 4

To prove this intermediate theorem, we need to study every gradient step involved. At iteration s , not every coordinates is available, hence the need to study the impact of T gradient steps together. A gradient step alongside edge ij only involves edges in its neighborhood (thanks to the sparsity of the matrix A), a key element that will need to be explicitated. The proof involves three main steps.

Step 1: Applying Lemma 4 (local smoothness) gives, where ij is the t^{th} activated edge:

$$F_A^*(\lambda(t+1)) - F_A^*(\lambda(t)) \leq -\frac{1}{2(\sigma_i^{-1} + \sigma_j^{-1})\mu_{ij}^2} \|\nabla_{ij} F_A^*(\lambda(t))\|^2. \quad (53)$$

Hence, we get an inequality between L_t and L_{t+1} :

$$\Lambda_{t+1} = \frac{1}{T} \sum_{t \leq s < t+T} (F_A^*(\lambda(s+1)) - F_A^*(\lambda^*)) \leq \Lambda_t - \frac{1}{T} \sum_{t \leq s < t+T} \frac{1}{2(\sigma_i^{-1} + \sigma_j^{-1})\mu_{(ij)_s}^2} \|\nabla_{(ij)_s} F_A^*(\lambda(s))\|^2 \quad (54)$$

where $(ij)_s$ is the edge activated during activation s . Let's introduce the following quantity:

$$\frac{1}{T} \sum_{t \leq s < t+T} \sum_{ij \in E} \|\nabla_{ij} F_A^*(\lambda(s))\|^2 = \frac{1}{T} \sum_{t \leq s < t+T} \|\nabla F_A^*(\lambda(s))\|^2 \geq \sigma_A \Lambda_t \quad (55)$$

where where we used Lemma 8 (gradient domination), and σ_A is the strong convexity parameter of F_A^* (lower bounded by $\lambda_{\min}^+(A^T A)/L_{\max}$). Hence, if an inequality of the type

$$\frac{C}{T} \sum_{t \leq s < t+T} \sum_{ij \in E} \|\nabla_{ij} F_A^*(\lambda(s))\|^2 \leq \frac{1}{T} \sum_{t \leq s < t+T} \frac{1}{2(\sigma_i^{-1} + \sigma_j^{-1})\mu_{(ij)_s}^2} \|\nabla_{(ij)_s} F_A^*(\lambda(s))\|^2 \quad (56)$$

holds, we have (using (50)):

$$\Lambda_{t+1} \leq \Lambda_t - C \frac{1}{T} \sum_{t \leq s < t+T} \|\nabla F_A^*(\lambda(s))\|^2 \leq (1 - C\sigma_A)\Lambda_t. \quad (57)$$

We thus need to tune correctly the μ_{ij}^2 and C in order to have (56) verified.

Step 2: We are looking for necessary conditions for (56) to hold. In the left term, every coordinate is present at each time s . However, in the right hand side of the inequality, just the activated one is present. We will need to compensate this with a bigger factor in front of the gradients. In order to compare these quantities, we need to introduce upper bound inequalities on $\|\nabla_{ij}F_A^*(\lambda(s))\|^2$, that only make activated coordinates intervene. Let $s \in \{t, \dots, t+T-1\}$, and suppose that there exists $t \leq r \leq s < r+t_{ij} \leq t+T-1$ such that ij is activated at times r and $r+t_{ij}$. Thanks to the assumption on T , either one of these integers exists. If the other one doesn't, replace it with t for r , and by $t+T-1$ for $r+t_{ij}$. Thanks to our assumptions, we know that $t_{ij} \leq a\ell_{ij}$. We have the following basic inequalities:

$$\|\nabla_{ij}F_A^*(\lambda(s))\|^2 \leq (\|\nabla_{ij}F_A^*(\lambda(r))\| + \|\nabla_{ij}F_A^*(\lambda(s)) - \nabla_{ij}F_A^*(\lambda(r))\|)^2 \quad (58)$$

$$\leq 2(\|\nabla_{ij}F_A^*(\lambda(r))\|^2 + \|\nabla_{ij}F_A^*(\lambda(s)) - \nabla_{ij}F_A^*(\lambda(r))\|^2). \quad (59)$$

The quantity $\|\nabla_{ij}F_A^*(\lambda(s)) - \nabla_{ij}F_A^*(\lambda(r))\|^2$ then needs to be controlled. We know that thanks to (49), for $x, x' \in \mathbb{R}^{E \times d}$, we have

$$\|\nabla_{ij}F_A^*(x) - \nabla_{ij}F_A^*(x')\|^2 \leq 2(\sigma_i^{-1} + \sigma_j^{-1})^2 d_{ij} \mu_{ij}^2 \sum_{(kl) \sim (ij)} \mu_{kl}^2 \|x_{kl} - x'_{kl}\|^2. \quad (60)$$

Using this with

$$\|x_{kl} - x'_{kl}\|^2 = \left\| \sum_{r < u < s: (ij)_u = (kl)} \frac{1}{(\sigma_k^{-1} + \sigma_l^{-1}) \mu_{kl}^2} \nabla_{kl}F_A^*(\lambda(u)) \right\|^2 \quad (61)$$

$$\leq \sum_{r < u < r+t_{ij}: (ij)_u = (kl)} \left(\frac{1}{(\sigma_k^{-1} + \sigma_l^{-1}) \mu_{kl}^2} \right)^2 N(kl, ij, u) \|\nabla_{kl}F_A^*(\lambda(u))\|^2, \quad (62)$$

where we used (and will widely use again below) that $\|x_1 + \dots + x_n\|^2 \leq n(\|x_1\|^2 + \dots + \|x_n\|^2)$ (convexity of the squared norm), leads to:

$$\|\nabla_{ij}F_A^*(\lambda(s))\|^2 \leq 2\|\nabla_{ij}F_A^*(\lambda(r))\|^2 \quad (63)$$

$$+ 2d_{ij} \sum_{r < u < r+t_{ij}} N((ij)_u, ij, u) \frac{\mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})^2}{\mu_{(ij)_u}^2 (\sigma_{i_u}^{-1} + \sigma_{j_u}^{-1})^2} \|\nabla_{(ij)_u}F_A^*(\lambda(u))\|^2 \quad (64)$$

$$\leq 2\|\nabla_{ij}F_A^*(\lambda(r))\|^2 \quad (65)$$

$$+ 2d_{ij} \sum_{r < u < r+t_{ij}} \left[b \frac{\ell_{ij}}{L_{(ij)_u}} \right] \frac{\mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})^2}{\mu_{(ij)_u}^2 (\sigma_{i_u}^{-1} + \sigma_{j_u}^{-1})^2} \|\nabla_{(ij)_u}F_A^*(\lambda(u))\|^2 \quad (66)$$

The advantage of this last expression is that only activated quantities are present on the right hand side.

Step 3: The last step of the proof consists in summing the last inequality for $t \leq s < t+T$, $ij \in E$. When summing, each $\|\nabla_{(ij)_r}F_A^*(\lambda(r))\|^2$ appears on the right hand-side of the inequality, with a factor upper-bounded by $((ij)_r)$, noted (ij) :

$$2a\ell_{ij} + 2d_{ij} \sum_{kl \sim ij} a\ell_{kl} \left[\frac{b\ell_{kl}}{\ell_{ij}} \right] \frac{\mu_{kl}^2 (\sigma_k^{-1} + \sigma_l^{-1})^2}{\mu_{ij}^2 (\sigma_i^{-1} + \sigma_j^{-1})^2}. \quad (67)$$

We want the expression above multiplied by C defined in Step 1 to be upper-bounded by $\frac{1}{2(\sigma_i^{-1} + \sigma_j^{-1})\mu_{ij}^2}$, in order for (56) to be verified. This is possible if and only if:

$$C \left(2a\ell_{ij}\mu_{ij}^2(\sigma_i^{-1} + \sigma_j^{-1}) + 2d_{ij} \sum_{kl \sim ij} a \left\lceil \frac{b\ell_{kl}}{\ell_{ij}} \right\rceil \ell_{kl}\mu_{kl}^2 \frac{(\sigma_k^{-1} + \sigma_l^{-1})^2}{\sigma_i^{-1} + \sigma_j^{-1}} \right) \leq \frac{1}{2}, \quad (68)$$

where C is defined in step 1 of the proof. This is equivalent to:

$$C \left(2a\ell_{ij}\mu_{ij}^2(\sigma_i^{-1} + \sigma_j^{-1}) + 2d_{ij} \sum_{kl \sim ij} a \frac{b\ell_{kl}^2}{\ell_{ij}} \mu_{kl}^2 \frac{(\sigma_k^{-1} + \sigma_l^{-1})^2}{\sigma_i^{-1} + \sigma_j^{-1}} \right) \leq \frac{1}{4} \text{ if } \forall kl \sim ij, \ell_{ij} \leq b\ell_{kl}, \quad (69)$$

where we bounded $\left\lceil b \frac{\ell_{ij}}{\ell_{kl}} \right\rceil$ by $2 \frac{b\ell_{ij}}{\ell_{kl}}$ here. We here see that in this case, if

$$\mu_{ij}^2 = \frac{1}{\ell_{ij}(\sigma_i^{-1} + \sigma_j^{-1})} \times \min_{kl \sim ij} \frac{\ell_{kl}(\sigma_k^{-1} + \sigma_l^{-1})}{\ell_{ij}(\sigma_i^{-1} + \sigma_j^{-1})} \quad (70)$$

with $8a + 8d_{max}^2 b \leq C^{-1}$, our inequality holds. However, our inequality on the ceil operator seems not to work in the general case. Let's take kl a neighbor of ij such that $\ell_{ij} > b\ell_{kl}$. As $\ell_{ij} > b\ell_{kl}$, we have $\left\lceil \frac{b\ell_{kl}}{\ell_{ij}} \right\rceil = 1$, leading to $a \left\lceil \frac{b\ell_{kl}}{\ell_{ij}} \right\rceil \ell_{kl}\mu_{kl}^2 = a\ell_{kl}\mu_{kl}^2 \leq a \leq ab$. Hence, our result still holds.

Conclusion: We have our result for $C = \frac{1}{2a+8d_{max}^2 ab}$ and a laplacian weighted with local communication constraints: $\mu_{ij}^2 = \frac{1}{\ell_{ij}(\sigma_i^{-1} + \sigma_j^{-1})} \times \min_{kl \sim ij} \frac{\ell_{kl}(\sigma_k^{-1} + \sigma_l^{-1})}{\ell_{ij}(\sigma_i^{-1} + \sigma_j^{-1})}$. The final rate thus depends on the smallest eigenvalue of the laplacian weighted by:

$$\frac{1}{2a + 8d_{max}^2 ab} \frac{1}{L_{max}} \frac{1}{\ell_{ij}(\sigma_i^{-1} + \sigma_j^{-1})} \times \min_{kl \sim ij} \frac{\ell_{kl}(\sigma_k^{-1} + \sigma_l^{-1})}{\ell_{ij}(\sigma_i^{-1} + \sigma_j^{-1})}. \quad (71)$$

However, having local complexity constraints is not really of much interest to us, as the parameters σ_i entered in the algorithm are generally taken to be the same on all nodes. We thus formulate Theorem 2 with σ_{min} for simplicity (which is slightly weaker in general) which gives as final rate of convergence the smallest eigenvalue of the laplacian weighted by:

$$v_{ij} = \frac{1}{2a + 8d_{max}^2 ab} \frac{\sigma_{min}}{2L_{max}} \frac{1}{\ell_{ij}} \times \min_{kl \sim ij} \frac{\ell_{kl}}{\ell_{ij}}. \quad (72)$$

C.2 Proof Of Proposition 3: Adding Stochasticity

We now prove the other theorem, where we assume the existence of events A_t for $t \in \mathbb{N}$, under which the assumptions are true. Using the same arguments as in the proof of Theorem 2, we obtain:

$$\mathbb{E}[\Lambda_{t+1} - \Lambda_t | \mathcal{F}_t, A_t] \leq -\sigma \Lambda_t. \quad (73)$$

However, this is not enough to conclude. Under A_t^c , we only know that $\Lambda_{t+1} \leq \Lambda_t$ using Lemma 4 (our local gradient steps cannot increase distance to the optimum). Hence:

$$\mathbb{E}[\Lambda_{t+1} | \mathcal{F}_t] \leq (1 - \sigma \mathbb{I}_{A_t}) \Lambda_t. \quad (74)$$

And then, by induction:

$$\mathbb{E}[\Lambda_t] \leq \mathbb{E}[P_t \Lambda_0], \text{ where } P_t = \prod_{s=0}^{t-1} (1 - \sigma \mathbb{I}_{A_s}). \quad (75)$$

However, no direct bound on P_t exists. The interdependencies on the events A_t make it impossible for an induction to prove a bound of the form $\leq (1 - \sigma/2)^t$. However, the logarithm of the product seems easier to study:

$$\log(P_t) = \log(1 - \sigma) \sum_{s=0}^{t-1} \mathbb{I}_{A_s}, \quad (76)$$

giving us $\mathbb{E} \log(P_t) \leq \log(1 - \sigma)t/2$, as $\mathbb{P}(A_t) \geq 1/2$. We are thus going to make a study in probability. For $t \in \mathbb{N}$, let $X_t = \frac{1}{T} \sum_{s=t}^{t+T-1} \mathbb{I}_{A_s}$. Using Markov-type inequalities conditionnaly on \mathcal{F}_t gives:

$$\mathbb{P}(X_t \geq 1/3 | \mathcal{F}_t) + 1/3 \mathbb{P}(X_t \leq 1/3 | \mathcal{F}_t) \geq \mathbb{E}[X_t | \mathcal{F}_t] \geq 1/2 \implies \mathbb{P}(X_t \geq 1/3 | \mathcal{F}_t) \geq 1/4. \quad (77)$$

Thus, we have: $\mathbb{E}[\prod_{s=t}^{t+T-1} (1 - \mathbb{I}_{A_s} \sigma) | \mathcal{F}_t] \leq \frac{1}{4} (1 - \sigma)^{T/3} + \frac{3}{4}$. We then know how to control T consecutive factors of the product P_t . Skipping the next T terms, we have:

$$\mathbb{E} \left[\prod_{s=t}^{t+3T-1} (1 - \mathbb{I}_{A_s} \sigma) \right] = \mathbb{E} \left[\prod_{s=t}^{t+T-1} (1 - \mathbb{I}_{A_s} \sigma) \prod_{s=t+T}^{t+2T-1} (1 - \mathbb{I}_{A_s} \sigma) \prod_{s=t+2T}^{t+3T-1} (1 - \mathbb{I}_{A_s} \sigma) \right] \quad (78)$$

$$\leq \mathbb{E} \left[\prod_{s=t}^{t+T-1} (1 - \mathbb{I}_{A_s} \sigma) \prod_{s=t+2T}^{t+3T-1} (1 - \mathbb{I}_{A_s} \sigma) \right] \quad (79)$$

$$\leq \mathbb{E} \left[\prod_{s=t}^{t+T-1} (1 - \mathbb{I}_{A_s} \sigma) \mathbb{E}^{\mathcal{F}_{t+2T}} \left\{ \prod_{s=t+2T}^{t+3T-1} (1 - \mathbb{I}_{A_s} \sigma) \right\} \right] \quad (80)$$

as in the last right hand side, the first big product is \mathcal{F}_{t+2T} -measurable (our assumption on the A_s states that they are \mathcal{F}_{s+T-1} -measurable). Then, using inequality $\mathbb{E} \left[\prod_{s=t}^{t+T-1} (1 - \mathbb{I}_{A_s} \sigma) | \mathcal{F}_t \right] \leq \frac{1}{4} (1 - \sigma)^{T/3} + \frac{3}{4}$ twice, with t and $t + 2T$, we get:

$$\begin{aligned} \mathbb{E} \left[\prod_{s=t}^{t+3T-1} (1 - \mathbb{I}_{A_s} \sigma) \right] &\leq \mathbb{E} \left[\prod_{s=t}^{t+T-1} (1 - \mathbb{I}_{A_s} \sigma) \left(\frac{1}{4} (1 - \sigma)^{T/3} + \frac{3}{4} \right) \right] \\ &\leq \left(\frac{1}{4} (1 - \sigma)^{T/3} + \frac{3}{4} \right)^2. \end{aligned}$$

Proceeding the same way by induction leads us to:

$$\mathbb{E}[P_t] \leq \left(\frac{1}{4} (1 - \sigma)^{T/3} + \frac{3}{4} \right)^{\lfloor t/(2T) \rfloor}, \quad (81)$$

which is the desired bound. For the asymptotic one, $(1 - \sigma)^{T/3} \leq e^{-\sigma T/3}$. For σT small enough (less than $\log(2)$), we have $e^{-\sigma T/3} \leq 1 - \sigma T/3$, leading to $(\frac{1}{4} (1 - \sigma)^{T/3} + \frac{3}{4})^{\lfloor t/(2T) \rfloor} \leq (1 - T\sigma/12)^{\lfloor t/(2T) \rfloor} \leq e^{-(t+o(t))\sigma/24}$. The asymptotic rate of convergence thus holds if the assumption made in Corollary 1 holds.

C.3 Study in the $RLNM(\varepsilon)$: Tuning the Parameters

We first assume to be in the case $\varepsilon = 0$. We generalize to $\varepsilon > 0$ at the end. Let $t \in \mathbb{N}$ be fixed, and B_t be the event: "in the activations $t, t + 1, \dots, t + T - 1$, all edges are activated". Let then $C_t(ij, s)$ for $t \leq s < t + T$ be the event $\min(T_{ij}(s), t + T - s, s - t) \leq a \ell_{ij}$ and $D_t(kl, ij, s)$ be the event $N(kl, ij, s) \leq \lceil b \ell_{ij} / \ell_{kl} \rceil$, where $N(kl, ij, s)$ is the number of activations of kl between two activations of ij , around time s , where we only take into account the activations between times t and $t + T - 1$. Let then $A_t = B_t \cap (\cap_{kl, ij \in E, t \leq s < t+T} C_t(ij, s) \cap D_t(kl, ij, s))$. We want $\mathbb{P}(A_t) \geq 1/2$ for correct constants a, b, T and ℓ_{ij} (that can differ from τ_{ij}). Note that this event is \mathcal{F}_{t+T-1} -measurable,

as desired. We first study the length of time ℓ_{ij} edge ij must wait in order to be activated with high probability (*high* meaning more than $1 - \frac{1}{12|E|}$). This result is Lemma 2. Then, we use this length to determine the constants T, a, b, ℓ_{ij} needed.

LEMMA 9. For any $t_0 \geq 0$, $ij \in E$, if $p_{ij} = \frac{1}{2 \max(d_i, d_j) - 1} \tau_{ij}^{-1}$ and $\tau_{\max}(ij) = \max_{kl \sim ij} \tau_{kl}$, let $\ell_{ij} = \frac{\log(6|E|)}{\log(1 - (1 - e^{-1})e^{-1})} (p_{ij}^{-1} + \tau_{\max}(ij))$. We have:

$$\mathbb{P}(ij \text{ not activated in } [t_0, t_0 + \ell_{ij}] | \mathcal{F}_{t_0}) \leq \frac{1}{6|E|}. \quad (82)$$

PROOF OF LEMMA 2. Let $ij \in E$ and $t_0 \geq 0$ fixed. We use tools from queuing theory [44] ($M/M/\infty/\infty$ queues) in order to compute the probability that edge ij is activable at a time t or not. More formally, we define a process $N_{ij}(t)$ with values in \mathbb{N} , such that $N_{ij}(t_0) = 1$ if ij non-available at time t_0 and 0 otherwise. Then, when an edge $kl, kl \sim ij$ is activated, we make an increment of 1 on $N_{ij}(t)$ (a *customer* arrives). This customer stays for a time τ_{kl} and when he leaves we make N_{ij} decrease by 1. We have $N_{ij} \geq 0$ a.s., and if $N_{ij} = 0$, ij is available. For $t \geq \max_{kl \sim ij} \tau_{kl} + t_0$, $N_{ij}(t)$ follows a Poisson law of parameter $\sum_{kl \sim ij} p_{kl} \tau_{kl}$. For any $t \geq \max_{kl \sim ij} \tau_{kl} + t_0$:

$$\mathbb{P}(ij \text{ available at time } t | \mathcal{F}_{t_0}) \geq \mathbb{P}(N_{ij}(t) = 0) = \exp\left(-\sum_{kl \sim ij} p_{kl} \tau_{kl}\right). \quad (83)$$

That leads to taking $p_{kl} = \frac{1}{2 \max(d_k, d_l) - 1} \tau_{kl}^{-1}$ for all edges, in order to have $\mathbb{P}(ij \text{ available at time } t | \mathcal{F}_{t_0}) \geq 1/e$. Then, $\mathbb{P}(ij \text{ rings in } [t, t + p_{ij}^{-1}]) = 1 - e^{-1}$, giving:

$$\mathbb{P}(ij \text{ activated in } [t_0, t_0 + \tau_{\max}(ij) + p_{ij}^{-1}] | \mathcal{F}_{t_0}) = \mathbb{P}(ij \text{ rings in } [t, t + p_{ij}^{-1}]) \quad (84)$$

$$\times \mathbb{P}(ij \text{ available at time } t | \mathcal{F}_{t_0}, ij \text{ rings at a time } t) \quad (85)$$

$$t \in [t_0 + \tau_{\max}(ij), t_0 + \tau_{\max}(ij) + p_{ij}^{-1}) \quad (86)$$

$$\geq (1 - e^{-1})e^{-1}, \quad (87)$$

where we use the fact that exponential random variables have no memory. Take $k \in \mathbb{N}$ such that $(1 - (1 - e^{-1})e^{-1})^k \leq \frac{1}{6|E|}$, leading to $k \approx \log(6|E|)/\log(1 - (1 - e^{-1})e^{-1})$. Let $\ell_{ij} = k(p_{ij}^{-1} + \tau_{\max}(ij))$. Then we have a.s.:

$$\mathbb{P}(ij \text{ not activated in } [t_0, t_0 + \ell_{ij}] | \mathcal{F}_{t_0}) \leq \frac{1}{6|E|}. \quad (88)$$

□

Bounding T : A direct application of Lemma 2 leads, with $L = \max_{ij} \ell_{ij}$, to:

$$T = 2 \sum_{ij} \frac{L}{\tau_{ij}}. \quad (89)$$

Indeed, for all ij , not being activated in activations $t, t + 1, \dots, t + T - 1$ means not being activated for a continuous interval of time of length more than ℓ_{ij} . Hence:

$$\mathbb{P}(\exists (ij) \in E : (ij) \text{ not activated in } \{t, \dots, t + T - 1\} | \mathcal{F}_t) \quad (90)$$

$$\leq \sum_{ij \in E} \mathbb{P}((ij) \text{ not activated in } \{t, \dots, t + T - 1\} | \mathcal{F}_t) \quad (91)$$

$$\leq \sum_{ij \in E} \mathbb{P}((ij) \text{ not activated in } [t, t + \ell_{ij}] | \mathcal{F}_t) \quad (92)$$

$$\leq |E| \times \frac{1}{6|E|} \quad (93)$$

$$= 1/6. \quad (94)$$

Bounding T_{ij} : Applying Lemma 2 with $12|E|T$ instead of $6|E|$ leads to controlling all the inactivation lengths by a length ℓ'_{ij} , with a probability more than $1 - 1/(12|E|T)$. Let $ij \in E$ and $s \in \mathbb{N}$, $t \leq s < t + T$. Let $\alpha > 0$ to tune later. Denote by $\delta_{ij}(s)$ the (random) inactivation time of ij , around iteration s . Note that conditionnaly on the inactivation period $\delta_{ij}(s)$, $T_{ij}(s)$ is dominated in law by a Poisson variable of parameter $I\delta_{ij}(s)$, hence line (96):

$$\mathbb{P}(T_{ij}(s) \geq \alpha \ell'_{ij} | \mathcal{F}_t) \leq \mathbb{P}(T_{ij}(s) \geq \alpha \ell'_{ij} | \mathcal{F}_t, \delta_{ij} \leq \ell'_{ij}) \times \mathbb{P}(\delta_{ij} \leq \ell'_{ij}) + \mathbb{P}(\delta_{ij} \geq \ell'_{ij}) \quad (95)$$

$$\leq \mathbb{P}(\text{Poisson}(I\ell'_{ij}) \geq \alpha \ell'_{ij}) + \frac{1}{12|E|T} \quad (96)$$

$$\leq \frac{1}{12|E|T} + \frac{1}{12|E|T} \quad (97)$$

$$= \frac{1}{6|E|T}, \quad (98)$$

for some $\alpha > 0$ big enough, to determine with the following large deviation inequality:

LEMMA 10 (A LARGE DEVIATION INEQUALITY ON DISCRETE POISSON VARIABLES.). *Let $Z \sim \text{Poisson}(\lambda)$, for some $\lambda > 0$. Then, for all $u \geq 0$:*

$$\mathbb{P}(Z \geq u) \leq \exp(-u + \lambda(e - 1)). \quad (99)$$

This large deviation leads to taking $\alpha = 2eI$ for (97) to be true. Finally, we get:

$$\mathbb{P}(T_{ij}(s) \geq \alpha \ell'_{ij} | \mathcal{F}_t) \leq \frac{1}{6|E|T}. \quad (100)$$

Bounding $N(kl, ij, s)$: If $\delta_{ij}(s) \leq \ell'_{ij}$, this random variable is dominated by a Poisson variable of parameter $p_{kl}\ell'_{ij}$. Hence, still with Lemma 10, with probability more than $1 - \frac{1}{12|E|^2T}$, we can bound $N(kl, ij)$ by $e \log(12|E|^2T) + p_{kl}\ell_{ij}(e - 1) \leq 2ep_{kl}L_{ij}$.

Explicit writing of the union bound on A_t^C : $A_t^C = B_t^C \cup (\cup_{kl, ij \in E, t \leq s < t+T} C_t(ij, s)^C \cup D_t(kl, ij, s)^C) \in \mathcal{F}_{t+T-1}$. Thanks to the previous considerations, we have that $\mathbb{P}^{\mathcal{F}_t}(B_t^C) \leq 1/6$ with (94), $\mathbb{P}^{\mathcal{F}_t}(C_t(ij, s)^C) \leq \frac{1}{6|E|T}$ with (100) and $\mathbb{P}(D_t(kl, ij, s)^C | \mathcal{F}_t) \leq \frac{1}{6|E|^2T}$, for the following constants and weights:

- $\tilde{\tau}_{ij}^{-1} = p_{ij} = \min(\frac{1}{\tau_{\max}(ij)}, \frac{1}{2(\max(d_i, d_j) - 1)} \frac{1}{\tau_{ij}})$;
- $T = 2I \max_{ij \in E} \tilde{\tau}_{ij} \frac{\log(6|E|)}{\log(1 - (1 - e^{-1})e^{-1})}$;
- $a = 2eI \frac{\log(6|E|T)}{\log(1 - (1 - e^{-1})e^{-1})}$;
- $b = 2e \frac{\log(6|E|T)}{\log(1 - (1 - e^{-1})e^{-1})}$.

The union bound is the following:

$$\mathbb{P}^{\mathcal{F}_t}(A_t^C) \leq \mathbb{P}^{\mathcal{F}_t}(B_t^C) + \sum_{s,ij} \mathbb{P}^{\mathcal{F}_t}(C_t(ij,s)^C) + \sum_{s,ij} \mathbb{P}^{\mathcal{F}_t}(\cup_{kl} D_t(kl,ij,s)^C) \quad (101)$$

$$\leq 1/6 + |E|T/(6|E|T) \times 2 \quad (102)$$

$$\leq 1/2. \quad (103)$$

The rate of convergence ρ is then defined as the smallest non null eigenvalue of the laplacian of the graph, weighted by:

$$v_{ij} = \frac{\sigma_{\min}}{L_{\max}} \times \frac{\tilde{\tau}_{ij} \min_{kl \sim ij} \frac{\tau_{ij}}{\tau_{kl}}}{8a(1+d^2b)}. \quad (104)$$

Note that this analysis works for $\varepsilon = 0$, but also for **RLNM**($\varepsilon > 0$) by replacing τ_{ij} by $(1 + \varepsilon)\tau_{ij}$. Indeed, Lemma 2 still holds with $(1 + \varepsilon)\tau_{ij}$: the queuing construction still works.

C.4 Proof of Corollary 1

PROOF. First, notice that $\mathcal{E}_k \leq \mathcal{L}_k$ since the sequence $(\mathcal{E}_l)_l$ is non-increasing. Then:

$$\begin{aligned} \left(\frac{1}{4} \left(1 - \frac{\sigma_{\min}}{L_{\max}} \Gamma_{RLNM} \right)^{T/3} + \frac{3}{4} \right)^{\lceil \frac{k}{2T} \rceil} &\leq \left(\frac{1}{4} \exp\left(-\frac{\sigma_{\min}}{L_{\max}} \Gamma_{RLNM} \frac{T}{3}\right) + \frac{3}{4} \right)^{\lceil \frac{k}{2T} \rceil} \\ &\leq \left(1 - \frac{\sigma_{\min}}{L_{\max}} \Gamma_{RLNM} \frac{T}{12e} \right)^{\lceil \frac{k}{2T} \rceil} \quad \text{if } \frac{\sigma_{\min}}{L_{\max}} \Gamma_{RLNM} \frac{T}{12} \leq 1. \end{aligned}$$

That last condition is satisfied under Assumption 1 using monotonicity of the Laplacian. We thus have our result taking the logarithm and making $k \rightarrow \infty$. \square