

Automated Performance Estimation for Decentralized Optimization via Network Size Independent Problems

Sebastien Colla and Julien M. Hendrickx

Abstract—We develop a novel formulation of the Performance Estimation Problem (PEP) for decentralized optimization whose size is independent of the number of agents in the network. The PEP approach allows computing automatically the worst-case performance and worst-case instance of first-order optimization methods by solving an SDP. Unlike previous work, the size of our new PEP formulation is independent of the network size. For this purpose, we take a global view of the decentralized problem and we also decouple the consensus subspace and its orthogonal complement. We apply our methodology to different decentralized methods such as DGD, DIGing and EXTRA and obtain numerically tight performance guarantees that are valid for any network size.

I. INTRODUCTION

Decentralized optimization has received an increasing attention due to its useful applications in large-scale machine learning and sensor networks, see e.g. [1] for a survey. In decentralized methods for separable objective functions, we consider a set of agents $\{1, \dots, N\}$, working together to solve the following optimization problem:

$$\min_{x \in \mathbb{R}^d} f(x) = \frac{1}{N} \sum_{i=1}^N f_i(x), \quad (1)$$

where $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is the private function locally held by agent i . To achieve this goal, each agent i holds its own version x_i of the decision variable $x \in \mathbb{R}^d$. Agents perform local computations and exchange local information with their neighbors to come to an agreement on the minimizer x^* of the global function f . The problem (1) can thus be written using the separable function $F_s : \mathbb{R}^{Nd} \rightarrow \mathbb{R}$:

$$\min_{x_1, \dots, x_N \in \mathbb{R}^d} F_s(x_1, \dots, x_N) = \frac{1}{N} \sum_{i=1}^N f_i(x_i), \quad (2)$$

such that $x_1 = \dots = x_N$.

Exchanges of information often take the form of an average consensus step on some quantity, e.g., on the x_i . This corresponds to a multiplication by an averaging matrix $W \in \mathbb{R}^{N \times N}$, often assumed symmetric and *doubly stochastic*, i.e., a nonnegative matrix whose rows and columns sum to one. This matrix W indicates both the topology of the network of agents and the weights they are using during the average

consensus. Therefore we call it network or averaging matrix, without distinction.

One of the simplest decentralized optimization method is the distributed (sub)gradient descent (DGD) [2] where agents successively perform an average consensus step (3) and a local gradient step (4):

$$y_i^k = \sum_{j=1}^N w_{ij} x_j^k, \quad (3)$$

$$x_i^{k+1} = y_i^k - \alpha^k \nabla f_i(x_i^k), \quad (4)$$

for step-sizes $\alpha^k > 0$. We focus on decentralized methods where the only agent-specific part in their update is the consensus step, in which the agents use different averaging weights. This is the case of any method that combines consensus steps, e.g. (3), and calls to first-order oracles of local functions for implicit or explicit updates, e.g. (4), using identical step-sizes and parameters for all the agents. This class of decentralized methods, denoted \mathcal{M}_D , includes many algorithms such as DGD [2], EXTRA [3], DIGing [4], and NIDS [5], to name a few.

In general, the quality of an optimization method is evaluated *via* a worst-case guarantee. Having accurate performance bounds for a decentralized algorithm is important to correctly understand the impact of its parameters and the network topology on its performance. However, obtaining these guarantees can often be a challenging task, requiring combining the impact of the optimization component and of the interconnection network, sometimes resulting in bounds that are conservative or very complex.

A. Contributions and Previous works

In this work, we build on a computational approach that computes worst-case performance guarantee of an algorithm by solving an optimization problem, known as the performance estimation problem (PEP), see Section II for more details. The PEP approach has led to many results in centralized optimization, see e.g. [6], [7], [8], and we have recently proposed a formulation tailored for decentralized optimization [9], [10] using a spectral description for symmetric and *generalized*¹ doubly stochastic averaging matrices, leading to some interesting results and observations. This formulation considers each agent individually in the PEP, implying that its size grows with the number of agents N . However, we observed in all our experiments [10, Section IV] that the results are independent of N for classical

S. Colla and J. M. Hendrickx are with the ICTEAM institute, UCLouvain, Louvain-la-Neuve, Belgium. S. Colla is supported by the French Community of Belgium through a FRIA fellowship (F.R.S.-FNRS). J. M. Hendrickx is supported by the “RevealFlight” Concerted Research Action (ARC) of the Federation Wallonie-Bruxelles and by the F.R.S.-FNRS via its research project Kornet and its Incentive Grant for Scientific Research (MIS) “Learning from Pairwise Comparisons”. Email addresses: sebastien.colla@uclouvain.be, julien.hendrickx@uclouvain.be

¹A generalized doubly stochastic matrix has rows and columns that sum to one, without the need to be non-negative, see [10, Definition 2].

performance criteria. In this paper, we develop a new PEP formulation for decentralized optimization whose problem size is independent of N , though N could still appear as a scaling coefficient in some cases, see Section III. This PEP formulation takes a global view on the decentralized problem (2), by neglecting the separability of function F_s , which a priori corresponds to a relaxation. Then, we apply a change of variables to distinguish the consensus and orthogonal parts of a vector. In Section IV, we demonstrate on several examples (DGD, DIGing and EXTRA), and observe empirically the tightness of this new agent-independent PEP formulation with symmetric and generalized doubly stochastic averaging matrices. We also observe that EXTRA seems to perform better than DIGing for time-constant averaging matrices but is not robust to time-varying matrices.

B. Related work

An alternative approach with similar motivations for automated performance evaluation is proposed in [11] and is inspired by dynamical systems concepts. Integral quadratic constraints (IQC), usually used to obtain stability guarantees on complex dynamical systems, are adapted to obtain sufficient conditions for the convergence of optimization algorithms and deduce numerical bounds on the convergence rates. It has been applied to decentralized optimization in [12]. This methodology analyzes the convergence rate of a single iterate, which is beneficial for the problem size that remains small (and that is also independent of the number of agents in the decentralized case). However, this does not allow to deal with non-geometric convergences, e.g. on smooth convex functions, nor to analyze cases with properties that apply over several iterations, e.g. when averaging matrices are constant.

These are possible with the PEP approach which computes the worst-case performance on a given number K of iterations, but solving problems whose size grows with K .

II. THE GENERAL PEP APPROACH

In principle, a tight performance bound on an algorithm could be obtained by running it on every single instance - function and initial condition - allowed by the setting considered and selecting the worst performance obtained. This would also directly provide an example of “worst” instance if it exists. The performance estimation problem (PEP) formulates this abstract idea as a real optimization problem that maximizes the error measure of the algorithm result, over all possible functions and initial conditions allowed [13]. This optimization problem is inherently infinite-dimensional, as it contains a continuous function among its variables. Nevertheless, Taylor et al. have shown [6], [7] that PEP can be solved exactly using an SDP formulation, for a wide class of centralized first-order algorithms and different classes of functions.

The main ingredients of a PEP are: (i) a performance measure \mathcal{P} , e.g. $f(x^k) - f(x^*)$; (ii) a class of functions \mathcal{F} , e.g. the class $\mathcal{F}_{\mu,L}$ of μ -strongly convex and L -smooth functions; (iii) an optimization method \mathcal{M} ; (iv) a set of

optimal solutions \mathcal{X}^* ; (v) a set \mathcal{I}^0 of initial conditions, e.g. $\|x^0 - x^*\|^2 \leq 1$. These ingredients are organized into an optimization problem as

$$\begin{aligned} & \sup_{f, x^0, \dots, x^K, x^*} \mathcal{P}(f, x^0, \dots, x^K, x^*) & (5) \\ & \text{s.t. } f \in \mathcal{F}, \quad x^* \in \mathcal{X}^*, \\ & \quad x^k \text{ are iterates from method } \mathcal{M}, \\ & \quad \mathcal{I}^0 \text{ holds.} \end{aligned}$$

To overcome the infinite dimension of variable f , the problem can be discretized into $\{(x^k, g^k, f^k)\}_{k \in I}$, where g^k and f^k are respectively the (sub)gradient and the function value of f at point k and $I = \{0, \dots, K, *\}$. Then, the constraint $f \in \mathcal{F}$ is replaced by interpolation conditions ensuring that there is a function of class \mathcal{F} which interpolates those data points $\{(x^k, g^k, f^k)\}_{k \in I}$. Such constraints are provided for many different classes of functions in [7, Section 3]. They are generally quadratics and potentially non-convex in the iterates and the gradients vectors, but they are linear in the scalar products of these and in the function values. The same holds true for most classical performance criteria and initial conditions. We can then consider these scalar products directly as decision variables of the PEP. For this purpose, we define a Gram matrix G that contains scalar products between iterates $x^k \in \mathbb{R}^d$ and subgradients $g^k \in \mathbb{R}^d$.

$$G = P^T P, \text{ with } P = [g^0 \dots g^K g^* x^0 \dots x^K x^*].$$

By definition, G is symmetric and positive semidefinite. Therefore, under some conditions described below, a PEP in the form of (5) can be formulated as an equivalent positive semidefinite program (SDP) using the Gram matrix G and the vector of functional values $f_v = [f_i]_{i \in I}$ as variables. This SDP formulation is convenient because it can be solved numerically to global optimality and it provides the worst-case solution over all possible problem dimensions². See [7] for details about the SDP formulation of PEP, including ways of reducing the size of matrix G . However, the dimension of G always depends on the number of iterations K . From a solution G, f_v of the SDP formulation, we can construct a solution for the discretized variables $\{(x^k, g^k, f^k)\}_{k \in I}$, e.g. using the Cholesky decomposition. Since these points satisfy sufficient interpolation constraints, we can also construct a function from \mathcal{F} interpolating these points.

Proposition 1 states sufficient conditions under which a PEP in the form of (5) can be formulated as an SDP. The proposition uses the following definition.

Definition 1 (Gram-representable): Consider a Gram matrix G and a vector f_v , as defined above. We say that a constraint or an objective is linearly (resp. LMI) Gram-representable if it can be expressed using a finite set of linear (resp. LMI) constraints involving (part of) G and f_v .

Proposition 1 ([7, Proposition 2.6]): If the interpolation constraints of the class of functions \mathcal{F} , the satisfaction of the method \mathcal{M} , the performance measure \mathcal{P} and the set of constraints \mathcal{I} , which includes the initial conditions, are

²Indeed the dimension d of vectors x^k and g^k disappears when taking their scalar products

linearly (or LMI) Gram-representable, then, computing the worst-case for criterion \mathcal{P} of method \mathcal{M} after K iterations on objective functions in class \mathcal{F} with constraints \mathcal{I} can be formulated as an SDP, with $G \succeq 0$ and f_v as variables.

Remark: In [7], Proposition 1 was only formulated for linearly Gram-representable constraints, but its extension to LMI Gram-representable constraints is direct. Such constraints appear in the analysis of consensus steps with spectral classes of network matrices.

PEP techniques allowed answering several important questions in optimization, see e.g. the list in [8], and to make important progress in the tuning of certain algorithms including the well-known centralized gradient-descent. It was further exploited to design optimal first-order methods: OGM for smooth convex optimization [14] and its extension ITEM for smooth strongly convex optimization [15]. It can also be used to deduce proofs about the performance of the algorithms [16]. It has been made widely accessible *via* a Matlab [17] and Python [18] toolbox.

III. AGENT-INDEPENDENT PEP FORMULATION FOR DECENTRALIZED OPTIMIZATION

In previous work [10], we find worst-case performance of decentralized optimization algorithms by formulating a performance estimation problem (PEP) that considers individually each agent, with its own local function and local variables. Remembering that the size of the SDP PEP depends on the number of variables considered and not on their dimension, we will obtain a agent-independent representation by abstracting the variables and functions at the global level. For this, we consider a generalization of problem (2) in which each f_i is allowed to depend on all x_j . The problem becomes then one of optimizing a function $F : \mathbb{R}^{Nd} \rightarrow \mathbb{R}$ over $\mathbf{x} \in \mathbb{R}^{Nd}$, under the constraint that \mathbf{x} belongs to the consensus subspace (all $x_i \in \mathbb{R}^d$ are equal). This will admit a simple representation that only requires decomposing \mathbf{x} into two sets of components, respectively along the consensus subspace and along its orthogonal complement, as opposed to the N sets of components along each x_i in [10].

A. Class of algorithms and vector notations

We consider the class \mathcal{M}_D of all the decentralized methods whose update may combine the 3 following types of operations:

(i) *Gradient:* Each agent evaluates the (sub)gradient of its local function

$$g_i = \nabla f_i(x_i) \quad \text{for } i = 1, \dots, N.$$

Each agent can possibly hold different local functions, which can be involved in different gradient operations.

(ii) *Consensus:* All the agents perform a consensus step on any of their local variables

$$y_i = \sum_{j=1}^N w_{ij} x_j, \quad \text{for } i = 1, \dots, N. \quad (6)$$

with weights w_{ij} given by an averaging matrix W . Different consensus steps can possibly use the same averaging matrix.

(iii) *Combinations:* Each agent performs the same linear combination of its local variables.

These three operations also allow implicit updates, e.g. updates where the point at which the gradient is evaluated is not explicitly known. These operations can always be expressed using the vector notation that stacks vertically the agents local variables $x_i \in \mathbb{R}^d$ and their local (sub)gradients $\nabla f_i(x_i) \in \mathbb{R}^d$ into vectors \mathbf{x} , $\mathbf{g} \in \mathbb{R}^{Nd}$

$$\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_N \end{bmatrix}, \quad \mathbf{g} = \begin{bmatrix} \nabla f_1(x_1) \\ \vdots \\ \nabla f_N(x_N) \end{bmatrix}.$$

In particular, the consensus step (6) can be written for all agents at once as

$$\mathbf{y} = (W \otimes I_d)\mathbf{x}, \quad (7)$$

where $W \in \mathbb{R}^{N \times N}$ is the averaging matrix used in the consensus, which is assumed to be symmetric and generalized doubly stochastic, I_d denotes the identity matrix of size d and \otimes the Kronecker product. This notation means that we apply the same matrix W for the d dimensions of the agents variables.

B. Generalization of the decentralized problem

We first define the consensus subspace of \mathbb{R}^{Nd} .

Definition 2 (Consensus subspace): The consensus subspace of \mathbb{R}^{Nd} is denoted \mathcal{C} and is defined as

$$\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^{Nd} \mid x_1 = \dots = x_N \in \mathbb{R}^d\}.$$

An orthonormal basis of \mathcal{C} is given by $Q_{//} = \frac{1}{\sqrt{N}}(\mathbf{1}_N \otimes I_d)$. The dimension of \mathcal{C} is thus d . The orthogonal complement of \mathcal{C} is denoted \mathcal{C}^\perp and has dimension $(N-1)d$:

$$\mathcal{C}^\perp = \{\mathbf{x} \in \mathbb{R}^{Nd} \mid \mathbf{x}^T \mathbf{y} = 0, \text{ for all } \mathbf{y} \in \mathcal{C}\}.$$

Using this definition, as well as the vector notations, we consider the following optimization problem to generalize the decentralized problem (2)

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^{Nd}} \quad & F(\mathbf{x}), \\ \text{such that } \quad & \mathbf{x} \in \mathcal{C}, \end{aligned} \quad (8)$$

where $F : \mathbb{R}^{Nd} \rightarrow \mathbb{R}$ is a real function from a given class of functions \mathcal{F} . The class of problems of the form of (2) can be directly related to the class of problems of the form of (8) with the additional assumption that F is a separable function. Let us then define the two classes of functions we are using in this paper, and their corresponding class of separable functions. The results can be extended to the other classes of functions that are Gram-representable. In the following, we denote by \mathbb{E} a finite-dimensional real vector space.

Definition 3 (\mathcal{F}_R): A function $F : \mathbb{E} \rightarrow \mathbb{R}$ is in $\mathcal{F}_R(\mathbb{E})$ if F is convex and with (sub)gradient norm bounded as

$$\|\nabla F(\mathbf{x})\| \leq R \quad \text{for } R \geq 0 \text{ and for all } \mathbf{x} \in \mathbb{E}.$$

Definition 4 ($\mathcal{F}_{\mu,L}$): A function $F : \mathbb{E} \rightarrow \mathbb{R}$ is in $\mathcal{F}_{\mu,L}(\mathbb{E})$ if F is μ -strongly convex and L -smooth (with $0 \leq \mu \leq L$), i.e. if we have, for all $(\mathbf{x}_1, \mathbf{x}_2) \in \mathbb{E}^2$

$$\begin{aligned} F(\mathbf{x}_1) &\geq F(\mathbf{x}_2) + \langle \nabla F(\mathbf{x}_2), \mathbf{x}_1 - \mathbf{x}_2 \rangle + \frac{\mu}{2} \|\mathbf{x}_1 - \mathbf{x}_2\|^2, \\ \|\nabla F(\mathbf{x}_1) - \nabla F(\mathbf{x}_2)\| &\leq L \|\mathbf{x}_1 - \mathbf{x}_2\|. \end{aligned}$$

Definition 5 (\mathcal{F}^s): Let \mathcal{F} be a class of functions, e.g. \mathcal{F}_R or $\mathcal{F}_{\mu,L}$. Then $\mathcal{F}^s(\mathbb{E}^N)$ is the set of functions $F : \mathbb{E}^N \rightarrow \mathbb{R}$ that are separable such that

$$F(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N f_i(x_i), \quad \text{with } f_i \in \mathcal{F}(\mathbb{E}) \text{ for all } i$$

where $\mathbf{x} = (x_1, \dots, x_N) \in \mathbb{E}^N$.

We now argue that the worst-case performance of a decentralized algorithm solving problem (2) can be bounded by its performance on problem (8). The argument relies on the following proposition which follows from Definition 5 combined with Definitions 3 and 4.

Proposition 2 (Relaxation of separability): The sets of separable functions (Definition 5) are such that

$$\mathcal{F}_R^s(\mathbb{R}^{Nd}) \subseteq \mathcal{F}_{\frac{R}{\sqrt{N}}}(\mathbb{R}^{Nd}), \text{ and } \mathcal{F}_{\mu,L}^s(\mathbb{R}^{Nd}) \subseteq \mathcal{F}_{\frac{\mu}{N}, \frac{L}{N}}(\mathbb{R}^{Nd}).$$

Then, any optimization problem (e.g. a PEP) with constraint $F \in \mathcal{F}_{\frac{R}{\sqrt{N}}}$ (resp. $F \in \mathcal{F}_{\frac{\mu}{N}, \frac{L}{N}}$) is a relaxation of the one with constraint $F \in \mathcal{F}_R^s$ (resp. $F \in \mathcal{F}_{\mu,L}^s$). We can thus compute worst-case performance bounds of decentralized algorithms for solving (2) by analyzing their performance on the general problem (8). This consists, a priori, in a relaxation since we do no longer exploit the separability of the objective function.

To analyze the performance of an algorithm on problem (8) using PEP, we need to express its constraint $\mathbf{x} \in \mathcal{C}$ in a Gram-representable form. To do this, we consider the problem in another coordinates system, which relies on the decomposition of \mathbb{R}^{Nd} into subspaces \mathcal{C} and \mathcal{C}^\perp .

C. Change of variables and new coordinates system

By definition, we have that $\mathbb{R}^{Nd} = \mathcal{C} \oplus \mathcal{C}^\perp$, so that any vector $\mathbf{x} \in \mathbb{R}^{Nd}$ can be uniquely decomposed as

$$\mathbf{x} = \mathbf{x}_{\parallel} + \mathbf{x}_{\perp}, \quad \text{with } \mathbf{x}_{\parallel} \in \mathcal{C} \text{ and } \mathbf{x}_{\perp} \in \mathcal{C}^\perp.$$

Therefore, a vector $\mathbf{x} \in \mathcal{C}$ if and only if $\mathbf{x}_{\perp} = 0$. Moreover, as we will see, a consensus step with a generalized doubly stochastic averaging matrix does not impact \mathbf{x}_{\parallel} , but only contracts \mathbf{x}_{\perp} . These observations motivate the following change of variable.

Definition 6 (Change of variable ${}_b\mathbf{x}$):

Let $Q_{\parallel} = \frac{1}{\sqrt{N}}(\mathbf{1}_N \otimes I_d)$ be the orthonormal basis of \mathcal{C} and $Q_{\perp} \in \mathbb{R}^{Nd \times (N-1)d}$ the orthonormal basis of \mathcal{C}^\perp . For any $\mathbf{x} \in \mathbb{R}^{Nd}$, we define ${}_b\mathbf{x} \in \mathbb{R}^{Nd}$ as

$${}_b\mathbf{x} = \begin{bmatrix} \bar{x} \\ x_{\perp} \end{bmatrix} = \frac{1}{\sqrt{N}} \begin{bmatrix} Q_{\parallel}^T \\ Q_{\perp}^T \end{bmatrix} \mathbf{x},$$

where $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i \in \mathbb{R}^d$ and $x_{\perp} \in \mathbb{R}^{(N-1)d}$. By definition, the matrix $Q = [Q_{\parallel} \ Q_{\perp}]$ is orthogonal, and the change of variable can thus be inverted via

$$\mathbf{x} = \sqrt{N}Q {}_b\mathbf{x}. \quad (9)$$

This change of variable can be interpreted as a change of coordinate system: \mathbf{x} describes the coordinates in the standard basis and ${}_b\mathbf{x}$ its coordinates in a new basis b , with \bar{x} the part of the coordinates describing the position

along the consensus subspace \mathcal{C} , and x_{\perp} the part describing the position along the orthogonal complement \mathcal{C}^\perp . Using Definition 6, we can easily express the constraint $\mathbf{x} \in \mathcal{C}$.

Proposition 3: $\mathbf{x} \in \mathcal{C}$ iff $x_{\perp} = 0$.

Proof: We obtain this from (9),

$$\mathbf{x} = \sqrt{N}Q_{\parallel}\bar{x} + \sqrt{N}Q_{\perp}x_{\perp},$$

where the first term is in \mathcal{C} and the second is in \mathcal{C}^\perp . ■

We now review the impacts of this change of variable on the generalized problem, the formulation of a decentralized algorithm and the Performance Estimation Problem.

1) *Changes in the Generalized Decentralized Problem:*

We define a new function $\tilde{F}({}_b\mathbf{x}) : \mathbb{R}^{Nd} \rightarrow \mathbb{R}$ depending on the new variable ${}_b\mathbf{x}$ as

$$\tilde{F}({}_b\mathbf{x}) = F(\sqrt{N}Q {}_b\mathbf{x}) = F(\mathbf{x}). \quad (10)$$

Using our change of variable and function, as well as Proposition 3, problem (8) becomes

$$\begin{aligned} \min_{{}_b\mathbf{x} \in \mathbb{R}^{Nd}} \quad & \tilde{F}({}_b\mathbf{x}), \\ \text{such that } & x_{\perp} = 0. \end{aligned} \quad (11)$$

Let us analyze what are the properties of the function \tilde{F} and its gradient, in view of formulating the Performance Estimation Problem for algorithms solving such problem.

Proposition 4 (Gradients of \tilde{F}):

Let $\nabla \tilde{F}({}_b\mathbf{x})$ be the gradient of \tilde{F} at point ${}_b\mathbf{x}$ and $\nabla F(\mathbf{x})$ the gradient of F at point \mathbf{x} . There holds

$$\nabla \tilde{F}({}_b\mathbf{x}) = \sqrt{N}Q^T \nabla F(\mathbf{x}). \quad (12)$$

Proof: Using the chain rule, direct computation shows

$$\nabla_{{}_b\mathbf{x}} \tilde{F}({}_b\mathbf{x}) = \nabla_{{}_b\mathbf{x}} F(\sqrt{N}Q {}_b\mathbf{x}) = \sqrt{N}Q^T \nabla_{\mathbf{x}} F(\mathbf{x}),$$

where indices for ∇ indicate which variable is considered in the derivatives. ■

The gradient $\nabla \tilde{F}({}_b\mathbf{x})$ can be decomposed in two parts that respectively describe the variation of the value of $\tilde{F}({}_b\mathbf{x})$ with the consensus and orthogonal components:

$$\nabla \tilde{F}({}_b\mathbf{x}) = \begin{bmatrix} \nabla_{\parallel} \tilde{F}({}_b\mathbf{x}) \\ \nabla_{\perp} \tilde{F}({}_b\mathbf{x}) \end{bmatrix}.$$

Proposition 4 implies the following corollary.

Corollary 1: The norm of the gradient of \tilde{F} at ${}_b\mathbf{x}$ can be computed as

$$\|\nabla \tilde{F}({}_b\mathbf{x})\|_2 = \sqrt{N} \|\nabla F(\mathbf{x})\|_2.$$

Moreover, if $F \in \mathcal{F}^s$, so that it is separable as $F(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N f_i(x_i)$, then $\nabla \tilde{F}({}_b\mathbf{x})$ is the coordinate vector in basis b of the vector $\mathbf{g}^T = [\nabla f_1^T(x_1) \dots \nabla f_N^T(x_N)]$, i.e.

$$\nabla \tilde{F}({}_b\mathbf{x}) = \frac{1}{\sqrt{N}} Q^T \mathbf{g} = b\mathbf{g}.$$

In particular, we have that $\nabla_{\parallel} \tilde{F}({}_b\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N \nabla f_i(x_i)$.

Proposition 5 (Optimal Conditions of (11)): The vector ${}_b\mathbf{x}^*$ is an optimal solution for problem (11) if and only if

$$x_{\perp}^* = 0 \quad \text{and} \quad \nabla_{\parallel} \tilde{F}({}_b\mathbf{x}^*) = 0.$$

Proof: By definition of problem (11), $x_\perp = 0$ for any feasible solution and thus the problem is equivalent to

$$\min_{b\mathbf{x} \in \mathbb{R}^{Nd}} \tilde{F} \left(\begin{bmatrix} \bar{x} \\ 0 \end{bmatrix} \right),$$

from which the result follows by convexity. ■

Proposition 6 (Properties of \tilde{F}): In general,

- (i) If $F \in \mathcal{F}_{\frac{R}{\sqrt{N}}}$, then $\tilde{F} \in \mathcal{F}_R$;
- (ii) If $F \in \mathcal{F}_{\frac{\mu}{N}, \frac{L}{N}}$, then $\tilde{F} \in \mathcal{F}_{\mu, L}$.

In cases where F is separable, this becomes

- (iii) If $F \in \mathcal{F}_R^s$, then $\tilde{F} \in \mathcal{F}_R$;
- (iv) If $F \in \mathcal{F}_{\mu, L}^s$, then $\tilde{F} \in \mathcal{F}_{\mu, L}$.

Proof: (i) If $F \in \mathcal{F}_{\frac{R}{\sqrt{N}}}$, then $\|\nabla F(\mathbf{x})\|_2 \leq \frac{R}{\sqrt{N}}$ for all \mathbf{x} and using Corollary 1, we have $\|\nabla \tilde{F}(b\mathbf{x})\|_2 \leq R$. (ii) Starting from Definition 4 for $F \in \mathcal{F}_{\frac{\mu}{N}, \frac{L}{N}}$, we can use equation (9), (10) and (12) to express it in term of function \tilde{F} , its gradients and variables $b\mathbf{x}_1$ and $b\mathbf{x}_2$. Then, using orthogonality of Q , we recover conditions for $\tilde{F} \in \mathcal{F}_{\mu, L}$. Finally, (iii) and (iv) directly follows from (i) and (ii) combined with Proposition 2. ■

2) Changes in the Decentralized Algorithms:

After applying the changes of variables (9), an algorithm for solving problem (8) will be suited for solving problem (11). The impact of these changes of variables can be treated independently for each of the three possible operations of a decentralized algorithm from \mathcal{M}_D .

Gradient: Using Corollary 1, the computation of the vector g of local gradients of function F at point \mathbf{x} simply becomes the computation of $\nabla \tilde{F}(b\mathbf{x}) = b\mathbf{g}$.

Combinations: Since Q is orthogonal, the changes of variables does not modify the linear combination operations, e.g. applying (9) to the combination $\mathbf{z} = \alpha\mathbf{x} + \beta\mathbf{y}$ gives

$$\sqrt{N}Qb\mathbf{z} = \sqrt{N}Q(\alpha b\mathbf{x} + \beta b\mathbf{y}) \Leftrightarrow b\mathbf{z} = \alpha b\mathbf{x} + \beta b\mathbf{y}.$$

Consensus: Theorem 1 analyzes the effects of a consensus step on the new variables.

Theorem 1 (Consensus step in basis b): Let $W \in \mathbb{R}^{N \times N}$ be a symmetric generalized doubly stochastic matrix. The consensus step $\mathbf{y} = (W \otimes I_d)\mathbf{x}$, is expressed in basis b as

$$b\mathbf{y} = \begin{bmatrix} \bar{y} \\ y_\perp \end{bmatrix} = \begin{bmatrix} \bar{x} \\ \tilde{W}x_\perp \end{bmatrix}, \quad (13)$$

where $\tilde{W} \in \mathbb{R}^{(N-1)d \times (N-1)d}$ has the same eigenvalues as W with a multiplicity d times larger, except for $\lambda_1 = 1$ that is not an eigenvalue of \tilde{W} .

Proof: When applying changes of variables (9) in the standard consensus step (7), since Q is orthogonal, we have

$$b\mathbf{y} = {}_bW b\mathbf{x} \quad \text{with} \quad {}_bW = Q^T(W \otimes I_d)Q, \quad (14)$$

where ${}_bW \in \mathbb{R}^{Nd \times Nd}$ is the averaging matrix for the consensus expressed in the basis b and has the same eigenvalues as $(W \otimes I_d)$ by similarity. By definition, matrix $(W \otimes I_d)$ is symmetric, generalized doubly stochastic, has the same eigenvalues than W with a multiplicity d times larger for each of them, and it can thus be written using the eigen-decomposition:

$$W \otimes I_d = V\Lambda V^T = [V_\parallel \quad V_\perp] \begin{bmatrix} I_d & 0 \\ 0 & \Lambda_\perp \end{bmatrix} \begin{bmatrix} V_\parallel^T \\ V_\perp^T \end{bmatrix},$$

where $V_\parallel = Q_\parallel = \frac{1}{\sqrt{N}}(\mathbf{1}_N \times I_d)$ are the orthonormal eigenvectors corresponding to the d eigenvalues equal to 1 and V_\perp are the orthonormal eigenvectors corresponding to other eigenvalues Λ_\perp . All the eigenvectors are orthogonal to each other and therefore, since $V_\parallel = Q_\parallel$, the matrix ${}_bW$ from (14) is given by

$$\begin{bmatrix} Q_\parallel^T \\ Q_\perp^T \end{bmatrix} [V_\parallel \quad V_\perp] \begin{bmatrix} I_d & 0 \\ 0 & \Lambda_\perp \end{bmatrix} \begin{bmatrix} V_\parallel^T \\ V_\perp^T \end{bmatrix} [Q_\parallel \quad Q_\perp] = \begin{bmatrix} I_d & 0 \\ 0 & \tilde{W} \end{bmatrix},$$

where $\tilde{W} = Q_\perp^T V_\perp \Lambda_\perp V_\perp^T Q_\perp$ and is thus symmetric. Eigenvalues of \tilde{W} are the ones in Λ_\perp since ${}_bW$ and $(W \otimes I_d)$ have the same eigenvalues. Finally, we find (13) using this last expression for ${}_bW$ in the consensus step (14). ■

Theorem 1 formalizes the separate effects of a consensus step that we discussed at the beginning of Section III-C.

We now have access to the problem, its properties, and any decentralized algorithm $\mathcal{M} \in \mathcal{M}_D$ written using the new variables. There only remains to formulate the PEP that evaluates the worst-case performance of \mathcal{M} for solving problem (11) with a given class of functions.

3) *Performance Estimation Problem in basis b :* This subsection reviews the main ingredients of a PEP (5) after the changes of variables (9), namely, the ingredients of a PEP applied to a decentralized algorithm for solving problem (11). We explain how such PEPs can be formulated as SDPs that are independent of the dimension of vector $b\mathbf{x} \in \mathbb{R}^{Nd}$, and thus independent of N .

First, the optimal conditions for (11) are given in Proposition 5. We observe that they require the use of specific components x_\perp and $\nabla_\parallel \tilde{F}$, independently of the full vectors $b\mathbf{x}$ and $\nabla \tilde{F}$. This separation of consensus coordinates and orthogonal coordinates also occurs in the effect of consensus step in Theorem 1. To be able to use such components independently from each other, we split the vectors $b\mathbf{x}^k$ and $b\mathbf{g}^k = \nabla \tilde{F}(b\mathbf{x}^k)$ in two parts in the SDP PEP formulation and we consider the scalar product of each part separately. This results then in two Gram matrices:

$$G_\parallel = P_\parallel^T P_\parallel \quad \text{and} \quad G_\perp = P_\perp^T P_\perp,$$

where $P_\parallel \in \mathbb{R}^{d \times 2K+2}$ and $P_\perp \in \mathbb{R}^{(N-1)d \times 2K+2}$ contains respectively the consensus and the orthogonal components of vectors $b\mathbf{x}^k$ and $b\mathbf{g}^k$. Both Gram matrices have identical sizes, independent of the number of agents N . They can possibly be combined into a single matrix $G \succeq 0$ as

$$G = \begin{bmatrix} G_\parallel & 0 \\ 0 & G_\perp \end{bmatrix}.$$

This modification still allows writing all the constraints involving scalar products between the iterates $b\mathbf{x}$ or the gradient $b\mathbf{g}$, e.g. the usual interpolations conditions for imposing that $\tilde{F} \in \mathcal{F}$, using the fact that

$$(b\mathbf{x})^T (b\mathbf{x}) = \bar{x}^T \bar{x} + x_\perp^T x_\perp, \quad (b\mathbf{g})^T (b\mathbf{g}) = g_\parallel^T g_\parallel + g_\perp^T g_\perp, \\ \text{and} \quad b\mathbf{x}^T g = \bar{x}^T g_\parallel + x_\perp^T g_\perp.$$

The usual interpolation constraints are well linearly Gram-representable as they are expressed as linear combinations of scalar products between $\bar{x}^k, g_{\parallel}^k$ and between x_{\perp}^k, g_{\perp}^k , and they can thus be embedded in the SDP PEP formulation, see Proposition 1.

Concerning the algorithmic steps constraints, we need to decouple them into constraints for the consensus components and for the orthogonal components. Let us analyze separately each possible operation of a decentralized algorithm.

Gradient and Combinations: The gradient computation and the linear combinations can directly be decoupled for the consensus and the orthogonal components. These constraints are linear equalities and are thus linearly Gram-representable.

Consensus: For a consensus step with a symmetric and generalized doubly stochastic averaging matrix W , the update for the consensus and orthogonal parts is given in Theorem 1. For the consensus components, it also corresponds to linear equality that is linearly Gram-representable. For the orthogonal component, the linear equality involves a modified averaging matrix \tilde{W} that cannot be known exactly without distinguishing the different agents individually. Therefore, we consider a spectral description of the symmetric matrix \tilde{W} that specifies its range of eigenvalues: $\lambda(\tilde{W}) \in [\lambda^-, \lambda^+]$, with $\lambda^- \leq \lambda^+ \in (-1, 1)$, which is identical to the one of W , excluding $\lambda_1(W) = 1$ (Theorem 1). We stress that the range $[\lambda^-, \lambda^+]$ is known in our performance analysis, but does not necessarily need to be known for the execution of the algorithms. We now adapt previous work [10] to represent, in a Gram-representable manner, a set of K consensus steps using the same matrix \tilde{W} :

$$Y = \tilde{W}X \quad \text{with} \quad \tilde{W} \in \mathcal{W}_s(\lambda^-, \lambda^+), \quad (15)$$

where $\mathcal{W}_s(\lambda^-, \lambda^+)$ is the class of symmetric matrices of size $(N-1)d$ with eigenvalues between λ^- and λ^+ and $Y, X \in \mathbb{R}^{(N-1)d \times K}$ are the matrices whose k -th column is the vector involved in the k -th consensus step. Consensus steps (15) are replaced by necessary LMI Gram-representable constraints that are stated in Theorem 2. These constraints allow the PEP to find worst-case valid over an entire spectral class of averaging matrix.

Theorem 2 (Consensus Constraints from [10, Theorem 3]): If $Y = \tilde{W}X$ with $\tilde{W} \in \mathcal{W}_s(\lambda^-, \lambda^+)$, then the matrix $X^T Y$ is symmetric and the following LMI Gram-representable constraints are satisfied

$$\begin{aligned} \lambda^- X^T X &\preceq X^T Y \preceq \lambda^+ X^T X, \\ (Y - \lambda^- X)^T (Y - \lambda^+ X) &\preceq 0, \end{aligned}$$

where the notations \preceq denote negative semi-definiteness.

The proof of Theorem 2 is very similar to [10, Theorem 3] and is omitted.

The last important ingredients of PEP are the initial conditions and the performance criterion. Using $(\bar{x}^k, x_{\perp}^k, g_{\parallel}^k, g_{\perp}^k, \tilde{F}^k)_{k=0, \dots, K, *}$, as well as all their properties studied in previous subsections, we can formulate any metric that is not agent-specific. Propositions 7 and 8 gives two examples for initial conditions. Metrics used for initial conditions can also be used as performance criteria.

Proposition 7: The initialization of variables $x_i^0 \in \mathbb{R}^d$, for $i = 1, \dots, N$ such that

$$x_i^0 = x^0 \text{ for all } i, \text{ and } \|x^0 - x^*\|_2^2 \leq D^2,$$

where $x^* \in \mathbb{R}^d$ is an optimal solution of (1), is equivalent to the following linearly Gram-representable constraints

$$x_{\perp}^0 = 0, \quad \text{and} \quad \|\bar{x}^0 - \bar{x}^*\|_2^2 \leq D^2,$$

where ${}_b \mathbf{x}^* = \begin{bmatrix} \bar{x}^* \\ 0 \end{bmatrix} \in \mathbb{R}^{Nd}$ is an optimal solution of (11).

Proof: All the agents start with the same x_i^0 , which means that $\mathbf{x}^0 \in \mathcal{C}$ and thus $x_{\perp}^0 = 0$, see Proposition 3. We also have, by definition, that $\bar{x}^0 = x^0$ and $\bar{x}^* = x^*$. These conditions are linearly Gram-representable since they involve linear combinations of scalar products of variables. ■

Proposition 8: The initialization of variables $x_i^0 \in \mathbb{R}^d$, for $i = 1, \dots, N$ such that

$$\frac{1}{N} \sum_{i=1}^N \|x_i^0 - x^*\|_2^2 \leq D^2, \quad (16)$$

where $x^* \in \mathbb{R}^d$ is an optimal solution of (1), is equivalent to the following linearly Gram-representable constraints

$$\|{}_b \mathbf{x}^0 - {}_b \mathbf{x}^*\|_2^2 = \|\bar{x}^0 - \bar{x}^*\|_2^2 + \|x_{\perp}^0 - x_{\perp}^*\|_2^2 \leq D^2, \quad (17)$$

where ${}_b \mathbf{x}^* \in \mathbb{R}^{Nd}$ is an optimal solution of (11).

Proof: Using change of variables (9), the initial condition (16) becomes

$$\frac{1}{N} \sum_{i=1}^N \|x_i^0 - x^*\|_2^2 = \frac{1}{N} \|\mathbf{x}^0 - \mathbf{x}^*\|_2^2 = \|{}_b \mathbf{x}^0 - {}_b \mathbf{x}^*\|_2^2 \leq D^2.$$

We then obtain (17) by definition of the Euclidean norm. This condition is linearly Gram-representable since it involves a linear combination of scalar products of variables. ■

This last example can also apply for metric $\sum_{i=1}^N \|x_i^0 - x^*\|_2^2$, but in that case, N would appear as a scaling parameter for the formulation in basis b .

We now have all the ingredients to build a PEP formulation that can evaluate bounds on the worst-case performance of any decentralized algorithm under consideration (see Section III-A) to solve problem (2) for different classes of functions (\mathcal{F}_R and $\mathcal{F}_{\mu, L}$), and whose size is independent of the number of agents N . The bounds are valid for any averaging matrices that is symmetric, generalized doubly-stochastic and with the given range of eigenvalues. The obtained bounds may scale with N as a parameter but we can choose the initial conditions and the performance criterion to obtain bounds totally independent of N .

IV. NUMERICAL RESULTS

A. Comparison with the previous PEP formulation [10]

We implement this N -independent PEP formulation for analyzing the performance of K iterations of DGD and DIGing with constant step-size in the same settings as in [10, Section IV], to compare the results with both the previous N -dependent PEP formulation [10] and the theoretical guarantees (Fig. 1). We consider a class of averaging matrices that contains symmetric generalized doubly stochastic matrices, with a range of eigenvalues $[-\lambda, \lambda]$ for $\lambda \in [0, 1)$ that holds




Fig. 1: Evolution with λ of the worst-case performance of $K = 10$ iterations of DGD and DIGing. The agent-independent PEP bounds from our new formulation exactly match the PEP bounds for $N = 2$ from previous PEP formulation [10] and improve on the theoretical guarantees. (a) DGD with $f_i \in \mathcal{F}_R$ for all i ($R = 1$) and $\alpha = \frac{1}{\sqrt{K}}$; (b) DIGing with $f_i \in \mathcal{F}_{\mu,L}$ for all i ($\mu = 0.1$, $L = 1$) and optimized α .

for all eigenvalues, except for $\lambda_1 = 1$. Our framework allows for any performance criterion to be chosen, but we use the same as the theoretical guarantees we are comparing with.

For DGD, we analyze the worst-case of $f(x_{av}) - f(x^*)$, where $f(x) = \frac{1}{N} \sum_{i=1}^N f_i(x)$ is the function we minimize in the decentralized problem (1), x^* is an optimum of f , and $x_{av} = \frac{1}{N(K+1)} \sum_{i=1}^N \sum_{k=0}^K x_i^k$ is the average iterates over all the iterations and all the agents. This performance criterion has a theoretical guarantee [1] that is shown in Fig. 1a for comparison. This guarantee applies to a step-size $\alpha = \frac{1}{\sqrt{K}}$ and local functions $f_i \in \mathcal{F}_R$ for all i , i.e. convex functions with subgradient norm bounded by R . The initial conditions are the ones described in Proposition 7. In our experiments, we set $R = 1$ and $D = 1$, but results for general values of R and D can be recovered by scaling.

Fig. 1a shows the evolution as a function of λ of the results obtained for 10 iterations. We observe that the spectral bound obtained with our new agent-independent PEP formulation matches those obtained with the previous agent-dependent PEP formulation for different values of N [10]. The independence of N in this spectral bounds from [10] was observed but not guaranteed. However, the new bound is known to be valid for any value of N . This means that

- (i) The relaxation of the separability of the objective function that we have introduced with the generalized decentralized problem (Section III-B) has no impact on the PEP results for DGD in such settings.
- (ii) The agent-independent PEP formulation is numerically tight for DGD with symmetric generalized doubly stochastic averaging matrices since this is the case for the agent-dependent formulation [10].

The same observations also apply to DIGing whose results are shown in Fig. 1b, for a different setting. We consider metric from Proposition 8 as both initial condition and performance criterion, we suppose $f_i \in \mathcal{F}_{\mu,L}$ for all i , i.e. local functions are μ -strongly convex and L -smooth. In our experiments, we set $D = 1$, $\mu = 0.1$ and $L = 1$ but results

for general values of D and L can be recovered by scaling. Other values of μ lead to similar observations. Resulting PEP and theoretical bounds depend on the step-size α . For each bound and each λ , we separately choose the step-size leading to the smallest value. For the spectral PEP formulations, the best values of α for $\mu = 0.1$ are found empirically: $\alpha = 0.44(1 - \lambda)^2$. Fig. 1b uses a special log-scale for the y-axis, shifted by 1, the initial value of the performance criterion. We see that the theoretical guarantee from [4] for $N = 2$ remains close to 1 for all values of λ and this gets worst for larger N . This is not the case for our spectral agent-independent PEP bound which improves on the theoretical one.

B. Comparison of DIGing and EXTRA

Our framework enables easy comparison of algorithms based on their worst-case performance bounds. These upper bounds, and thus the resulting comparisons are valid over the entire given spectral class of networks, for all network sizes. For example, Fig. 2 compares DIGing [4] and EXTRA [3] performance after 10 iterations, in the same setting as in Fig. 1b: we consider metric from Proposition 8 both as initial condition and performance criterion and $f_i \in \mathcal{F}_{\mu,L}$ for all i . In our experiments, we set $D = 1$, $\mu = 0.1$ and $L = 1$ but results for general values of D and L can be recovered by scaling. We also consider the class of symmetric generalized doubly stochastic averaging matrices with range of eigenvalues $[-\lambda, \lambda]$. The algorithm EXTRA [3] uses a second averaging matrix W_2 which the authors recommend to set to $W_2 = (I + W)/2$. Consensus steps using W_2 are thus considered as $W_2 x = \frac{x}{2} + \frac{Wx}{2}$.

We compare the two algorithms in two different cases:

- **Constant averaging matrices.** It means that the same matrix is used for each consensus. In that case, the conditions from Theorem 2 are applied in PEP to all consensus steps at once.
- **Time-varying averaging matrices.** It means that each

consensus step can potentially use a different matrix from the same spectral class. In that case, the conditions from Theorem 2 are applied separately to each consensus steps in PEP. Therefore, for each of these applications of Theorem 2, X and Y have only one column and LMIs become scalar inequalities.

For each case, the step-size α is optimized for each value of λ . In Fig. 2, we observe that, for constant averaging matrices (plain lines), the worst-case performance of EXTRA after 10 iterations is better than the one DIGing for all the values of λ . However, EXTRA seems to be much less robust to time-varying matrices (dashed line). Indeed, in that case, DIGing has exactly the same performance as for constant matrices, while the performance of EXTRA deteriorates: it remains better than DIGing for small values of λ but explodes when λ increases. The robustness of DIGing to time-varying matrices was already known since its theoretical analysis [4], despite its non-tightness (see Fig. 1b), is valid for time-varying matrices. However, the worst-case behavior of EXTRA for time-varying matrices was unknown since its theoretical analysis [3] only focuses on constant matrices.

Note that these observations hold after 10 iterations but similar ones were made for larger number of iterations. However, the value of λ where dashed lines cross may slightly vary. Our observations are consistent with the fact that the alternative IQC methodology [12] does not provide any convergence guarantee for EXTRA with similar settings when $\lambda \geq 0.6$. Indeed, IQC evaluates the performance over one iteration of the algorithm and can therefore only consider the case where the averaging matrices are time-varying.

V. CONCLUSION

We develop a methodology to automatically compute numerical worst-case bounds of a broad class of decentralized optimization methods that are independent of the number of agents in the problem. This methodology also allows

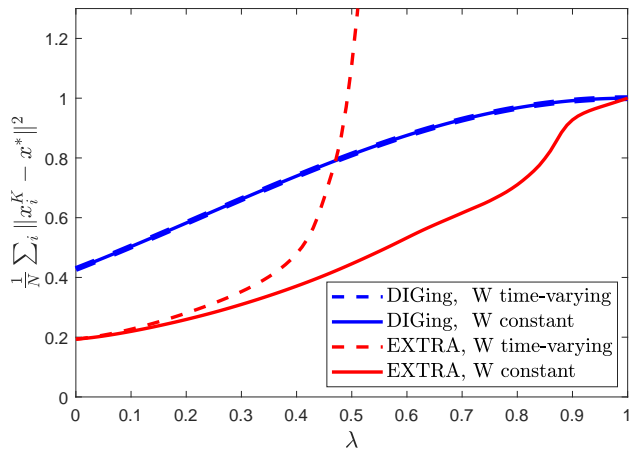


Fig. 2: Comparison between DIGing and EXTRA for 10 iterations with constant and time-varying averaging matrices. Evolution is shown with the range of eigenvalues $[-\lambda, \lambda]$. Optimized α and $f_i \in \mathcal{F}_{\mu, L}$ for all i ($\mu = 0.1$, $L = 1$).

considering time-dependent settings. This opens the way for computer-aided analysis of many decentralized algorithms, which could lead to improvements in their performance guarantees and parameter tuning. Moreover, the guarantees computed with our tool appear to be tight in many cases.

This methodology is based on the approach of performance estimation problem (PEP) and reuses results from our previous work to represent spectral classes of matrices. Although it formulates a relaxed problem, it appears to provide tight worst-case bounds for DGD and DIGing, which are valid for any number N of agents and improve on the theoretical guarantees.

REFERENCES

- [1] A. Nedić, A. Olshevsky, and M. G. Rabbat, “Network topology and communication-computation tradeoffs in decentralized optimization,” *Proceedings of the IEEE*, vol. 106, no. 5, pp. 953–976, 2018.
- [2] A. Nedic and A. Ozdaglar, “Distributed subgradient methods for multi-agent optimization,” *IEEE Transactions on Automatic Control*, vol. 54, pp. 48 – 61, 02 2009.
- [3] W. Shi, Q. Ling, G. Wu, and W. Yin, “Extra: An exact first-order algorithm for decentralized consensus optimization,” *SIAM Journal on Optimization*, vol. 25, 04 2014.
- [4] A. Nedic, A. Olshevsky, and W. Shi, “Achieving geometric convergence for distributed optimization over time-varying graphs,” *SIAM Journal on Optimization*, vol. 27, 07 2016.
- [5] Z. Li, W. Shi, and M. Yan, “A decentralized proximal-gradient method with network independent step-sizes and separated convergence rates,” *IEEE Transactions on Signal Processing*, vol. PP, 04 2017.
- [6] A. B. Taylor, J. M. Hendrickx, and F. Glineur, “Smooth strongly convex interpolation and exact worst-case performance of first-order methods,” *Mathematical Programming*, vol. 161, 02 2015.
- [7] —, “Exact worst-case performance of first-order methods for composite convex optimization,” *SIAM Journal on Optimization*, vol. 27, 12 2015.
- [8] A. Taylor, “Convex interpolation and performance estimation of first-order methods for convex optimization,” Ph.D. dissertation, UCLouvain, 2017, advisors : François Glineur and Julien M. Hendrickx.
- [9] S. Colla and J. M. Hendrickx, “Automated worst-case performance analysis of decentralized gradient descent,” in *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021, pp. 2627–2633.
- [10] —, “Automatic performance estimation for decentralized optimization,” *preprint*, 2022.
- [11] L. Lessard, B. Recht, and A. Packard, “Analysis and design of optimization algorithms via integral quadratic constraints,” *SIAM Journal on Optimization*, vol. 26, 08 2014.
- [12] A. Sundararajan, B. V. Scoy, and L. Lessard, “Analysis and design of first-order distributed optimization algorithms over time-varying graphs,” *IEEE Transactions on Control of Network Systems*, vol. 7, pp. 1597–1608, 2020.
- [13] Y. Drori and M. Teboulle, “Performance of first-order methods for smooth convex minimization: A novel approach,” *Mathematical Programming*, vol. 145, 06 2012.
- [14] D. Kim and J. Fessler, “Optimized first-order methods for smooth convex minimization,” *Mathematical Programming*, vol. 159, 06 2014.
- [15] A. B. Taylor and Y. Drori, “An optimal gradient method for smooth strongly convex minimization,” *Mathematical Programming*, pp. 1–38, 2022.
- [16] A. B. Taylor, J. M. Hendrickx, and F. Glineur, “Exact worst-case convergence rates of the proximal gradient method for composite convex minimization,” *Journal of Optimization Theory and Applications*, vol. 178, 08 2018.
- [17] —, “Performance estimation toolbox (PESTO): Automated worst-case analysis of first-order optimization methods,” in *IEEE 56th Annual Conference on Decision and Control (CDC)*, 2017, pp. 1278–1283.
- [18] B. Goujaud, C. Moucer, F. Glineur, J. Hendrickx, A. Taylor, and A. Dieuleveut, “PEPit: computer-assisted worst-case analyses of first-order optimization methods in Python,” *arXiv preprint arXiv:2201.04040*, 2022.