

Interpolation Constraints for Computing Worst-Case Bounds in Performance Estimation Problems

Anne Rubbens, Nizar Bousselmi, Sébastien Colla, and Julien M. Hendrickx

Abstract—The Performance Estimation Problem (PEP) approach consists in computing worst-case performance bounds on optimization algorithms by solving an optimization problem: one maximizes an error criterion over all initial conditions allowed and all functions in a given class of interest. The maximal value is then a worst-case bound, and the maximizer provides an example reaching that worst case. This approach was introduced for optimization algorithms but could in principle be applied to many other contexts involving worst-case bounds. The key challenge is the representation of infinite-dimensional objects involved in these optimization problems such as functions, and complex or non-convex objects as linear operators and their powers, networks in decentralized optimization etc. This challenge can be resolved by interpolation constraints, which allow representing the effect of these objects on vectors of interest, rather than the whole object, leading to tractable finite dimensional problems. We review several recent interpolation results and their implications in obtaining of worst-case bounds via PEP.

I. INTRODUCTION

A. Worst-case bounds

A classical way of characterizing the efficiency of optimization is through the derivation of *worst-case bounds*: guarantees on the evolution of a quantity of interest valid for all functions in a given class \mathcal{F} . As a pedagogical example, consider the gradient descent algorithm with a constant step-size $\alpha > 0$

$$x_{i+1} = x_i - \alpha \nabla f(x_i), \quad i = 0, \dots, N-1, \quad (1)$$

and let \mathcal{F}_L be the set of L -smooth convex functions defined on \mathbb{R}^d , i.e. functions for which for all x, y there holds

$$\begin{cases} f(\lambda x + (1-\lambda)y) \leq \lambda f(x) + (1-\lambda)f(y) \quad \forall \lambda \in [0, 1], \\ \|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\|. \end{cases} \quad (2)$$

One can prove, e.g. from [1, Theorem 2.1.13], that for every function $f \in \mathcal{F}_L$, every N and every $0 < \alpha \leq \frac{2}{L}$ there holds

$$f(x_N) - f(x^*) \leq \frac{2L\|x_0 - x^*\|^2}{4 + NL\alpha(2 - L\alpha)}. \quad (3)$$

Such bounds can be used not only for the intrinsic guarantees they provide, but also to compare the performances of algorithms, and tune their parameters, e.g. α in (1). One would indeed assume that an algorithm with a better bound is

The authors are with ICTEAM Institute, UCLouvain, Belgium. firstname.name@uclouvain.be Their work is supported by FRIA, FNRS, and a L'Oréal-UNESCO For Women in Science fellowship. This work is also supported by the "RevealFlight" ARC at UCLouvain, by the Incentive Grant for Scientific Research (MIS) "Learning from Pairwise Data" and by the KORNET project from F.R.S.-FNRS

more efficient, and typically select the algorithm parameters leading to the best bound. However, this seemingly intuitive approach may lead to sub-optimal or misguided choices when the bounds are too conservative: an algorithm with a stronger bound is not necessarily more efficient, and may just be more amenable to the analysis technique that was used. Similarly, the parameters minimizing the bound do not necessarily optimize the actual performance, and may just be those for which the proof technique introduces the smallest amount of conservatism. Far from being theoretical possibilities, these phenomena have been observed even in very simple settings. Fig. 1 compares the bound (3) with the exact worst-case bound in this setting, obtained via the method described in Section II. One can see that the classical bound (3) is minimized by $\alpha = L^{-1}$, while the actual performance is minimized by $\alpha \approx 1.834L^{-1}$ [2]. This step-size improves the bound by a factor 2 compared to $\alpha = L^{-1}$, and can be shown to actually divide the number of iterations required to achieve a given accuracy by 2. Moreover, this phenomenon is accentuated as N increases: the optimal step-sizes approach $2L^{-1}$, a range of values that would seem absurd based on the theoretical bound [2].




Fig. 1: Evolution with the adimensional step-size $h = L\alpha$ of two worst-case guarantees on $f(x_N) - f(x^*)$ for L -smooth convex functions f , where the number of iterations N is set to 10. The plot shows (i) the theoretical bound (3) and (ii) the exact bound for this setting. This tight bound allows improving the step-size selection, since the choice $h = 1$ minimizing the bound (3) requires doubling the number of steps to obtain the same performance as the true optimal step-size. This highlights the benefits of relying on tight bounds for parameter selection and algorithm comparisons.

This demonstrates the importance of bounds that are tight, or as tight as possible, to not only understand the performances of algorithms but also to compare them on a sound base in order to select or tune them, and to guide further progress in algorithmic development.

B. Sources of Conservatism

At an abstract level, almost all proofs of worst-case bounds consist in (i) translating conceptual assumptions on the function classes into algebraic relations between quantities playing a role in the algorithm, often the iterates x_i and the gradient $\nabla f(x_i)$ and/or function values $f(x_i)$ at these iterates, and (ii) combining these relations with the algorithm description to obtain a bound on a quantity of interest, see e.g. [1]. Conservatism can then come from two possible sources: an insufficiently effective translation of the conceptual assumptions, and/or a sub-optimal combination of the algebraic relations obtained with the algorithm description.

We defer the discussion of the second aspect to Section II and focus here on the first one, coming back to our example set \mathcal{F}_L of convex L -smooth functions defined in (2). Due to the use of quantifier $\forall \lambda$, (2) cannot be easily directly manipulated for finite set of points. However, an equivalent definition of \mathcal{F}_L , involving only easily discretizable constraints is the following [1]:

$$\begin{cases} f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle \\ \|\nabla f(x) - \nabla f(y)\| \leq L\|x - y\| \end{cases} \quad \forall x, y. \quad (4)$$

An alternative definition, if less intuitive one, is

$$f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle + \frac{1}{2L} \|\nabla f(x) - \nabla f(y)\|^2, \quad (5)$$

$\forall x, y$. One can show that constraints (4) and (5) are both completely equivalent to the definition (2) of \mathcal{F}_L when they are imposed for every couple $(x, y) \in \mathbb{R}^d$ [1]. However, most proofs will involve a finite number of equations, meaning that the constraints would be imposed on a finite number of pairs (x, y) , corresponding typically to the iterates of the algorithm and to remarkable points of the functions. And the equivalence between the constraints no longer holds when they are imposed on only a finite number of pairs: Constraint (4) is indeed weaker than (5), and sometimes significantly weaker as demonstrated in Fig. 2. Consequently, for most problems, the best bound that can be obtained using (4) will be weaker than that obtainable from (5), as further illustrated in Fig. 3: one can observe that the tightest bound obtained based on constraint (4) is minimized by a step-size of $\alpha \approx 0.694L^{-1}$, which yields an even worse actual performance than the classical choice of $\alpha = L^{-1}$, and that the bound explodes as α approaches $2L^{-1}$. In other words, relying on (4) to derive performance guarantees directly introduces more conservatism than using (5).

C. Interpolation Constraints

The simple example in Section I-B highlights the importance of working with algebraic characterizations of functions classes that introduce the smallest amount of conservatism, or none at all if possible. This is in fact the case of




Fig. 2: Comparison between the admissible values for f_2 and g_2 under respectively the discretized versions of (4) (in grey) and (5) (in black), assuming $x_2 = 0, x_1 = 1, g_1 = 1, f_1 = 0$. This shows that, despite the equivalence between these constraints when imposed on all pairs (x, y) , (4) is significantly weaker than (5). In particular, the points in the grey (but not black) area do not correspond to any actual function $f \in \mathcal{F}_L$.




Fig. 3: Evolution with the adimensional step-size $h = L\alpha$ of three worst-case guarantees on $f(x_{10}) - f(x^*)$ for L -smooth convex functions f . The plot shows (i) the theoretical bound (3), (ii) the best possible bound (i.e. PEP-based, see Section II), relying on the non-tight representation of \mathcal{F}_L (4) (see Section III), and (iii) the tight worst-case bound, i.e. the best possible bound relying on interpolation constraints (5). As already observed in Fig. 1, one can see that the step-sizes minimizing non-tight bounds lead to poorer actual performances, even if the bound is PEP-based, that is the best possible bound given a class representation, when the representation is non-tight. This highlights the necessity of relying on tight descriptions of classes to efficiently tune methods.

(5): it has been shown that for any arbitrary set of N triples $\{(x_i, g_i, f_i)\}_{i \in I}$ satisfying its discretized version

$$f_j \geq f_i + \langle g_i, x_i - x_j \rangle + \frac{1}{2L} \|g_i - g_j\|^2 \quad \forall i, j, \quad (6)$$

there exists an actual function $f \in \mathcal{F}_L$ such that $f(x_i) = f_i$ and $\nabla f(x_i) = g_i$ for every i . Replacing the assumption $f \in \mathcal{F}_L$ by (6) being satisfied for all pairs of points appearing in the problem does not introduce any conservatism, since there is always a function in \mathcal{F}_L that would correspond to these values. We say that such a function *interpolates* the $\{(x_i, g_i, f_i)\}$ and that (6) is an *interpolation constraint* since it guarantees the existence of an interpolating function. By contrast, the weaker constraint (4) is not an interpolation constraint for \mathcal{F}_L , as sets of triples $\{(x_i, g_i, f_i)\}$ satisfying it do not necessarily correspond to an actual function $f \in \mathcal{F}_L$, see Fig. 2. Hence any bound established using (4) is valid for a class of unspecified “objects” larger than the functions in \mathcal{F}_L and will thus in most cases be conservative for \mathcal{F}_L .

D. Paper Organization

In this tutorial paper we will first describe how interpolation constraints can be used to automatically compute worst-case bounds in Section II, and then review interpolation constraints for several classes of objects: functions and general operators in Section III, linear operators in Section IV, and doubly-stochastic matrices related to networks in decentralized optimization in Section V.

II. AUTOMATIC COMPUTATION OF WORST-CASE BOUNDS

We come back to the second sources of conservatism identified in Section I-B, i.e. the combination of algebraic translations of the conceptual assumptions and the algorithm description to obtain a bound. Combining such algebraic relations efficiently could in principle be solved by skilled manipulations, but finding the best way of manipulating them can in many cases be very challenging. Fortunately, this process can be automated, thanks among others to the Performance Estimation Problem (PEP) approach, which was the initial motivation for the formalization of the notion of interpolation constraints in this context [2]–[5] and which we will demonstrate here on our pedagogical example of the gradient descent algorithm (1).

Suppose we want to obtain the best possible bound of the form (3). The idea of PEP is to compute the maximal value of $f(x_N) - f(x^*)$ over all possible functions $f \in \mathcal{F}_L$ and starting points x_0 for which $\|x_0 - x^*\|^2 \leq 1$ (We can then obtain a bound proportional to $\|x_0 - x^*\|^2$ using a simple scaling argument). This a priori infinite-dimensional problem can be re-expressed in a finite manner: rather than considering the function f as a whole, we will only consider a discrete set $\{(x_i, g_i, f_i)\}_{i \in I = \{0, \dots, N, *\}}$ of iterates and optimal point x_i together with their gradient-vectors g_i and function values f_i , *under the constraint that the set is consistent with an actual function $f \in \mathcal{F}_L$* , i.e. that there exists a function $f \in \mathcal{F}_L$ interpolating it. The optimality of

x^* and the fact that the iterates are obtained by the algorithm (3) can then also be expressed as constraints. The PEP takes thus the form:

$$\begin{aligned} & \max_{\{(x_i, g_i, f_i)\}_{i \in I}} && f_N - f^* && (7) \\ \text{s.t.} &&& && \\ \text{algorithm:} &&& x_{i+1} = x_i - \alpha g_i \text{ for } i = 0, \dots, N-1. && \\ \text{optimality:} &&& g^* = 0 && \\ \text{initial cond.:} &&& \|x_0 - x^*\|^2 \leq 1 && \\ \text{consistency of} &&& \{(x_i, g_i, f_i)\}_{i \in I} \text{ with some } f \in \mathcal{F}_L. && \end{aligned}$$

The optimum of this problem gives by definition a tight worst-case performance bound for the gradient descent with L -smooth convex functions. Moreover, the optimal solution corresponds to an instance of function and initial condition actually reaching this upper bound, which can provide very relevant information on the bottlenecks faced by the algorithm. In many cases, these worst-case instances have actually a surprisingly simple structure [3].

In (7), the consistency of $\{(x_i, g_i, f_i)\}_{i \in I}$ with some $f \in \mathcal{F}_L$ is implicit, but can be directly replaced by an interpolation constraint when one is available, as (6) in the case of \mathcal{F}_L . When no interpolation constraint is known, one can always relax the consistency constraint into a *necessary constraint for interpolability*, as e.g. (4). The solution to the problem obtained will then be a valid worst-case upper bound, and in fact the best upper bound that could possibly be obtained using the necessary constraint used, but in general not a tight one, as illustrated in Fig. 3.

Observe that in both the cases of (6) and (4), the resulting problem is quadratic and potentially non-convex in the iterates and gradients vectors, but it is linear in the scalar products of these and in the function values. Hence, letting the matrix of these scalar products and the vector of function values be the decision variables of the PEP, one can reformulate it as a semi-definite program (SDP) and solve it efficiently, see e.g. [2, Section 3.2] for more details. We refer to any constraint linear in these scalar products and function values as *PEP-representable*.

The PEP framework allows computing the exact performance of a wide range of first-order methods on various classes of functions and objects for which PEP-representable interpolation constraints are known. In particular, any method, whether implicit or explicit, consisting in (sub)gradient queries, linear operators, and with PEP-representable constraints can be straightforwardly analyzed. This includes e.g the gradient method, the proximal gradient method [6], the coordinate descent [7], inexact gradient methods [8], Bregman gradient methods [9], variational problems and splitting methods [10], [11], and their combinations. In general, the analysis requires pre-determined step-sizes, possibly time-varying but independent of the function values or gradients. But some efforts have also been made in adaptive first-order methods [12], [13]. The formulation and resolution of Performance Estimation Problems is im-

plemented in Matlab and Python toolboxes: PESTO [14] and PEPit [15]. These toolboxes allow for a simple and human readable description of the algorithms and objects, close to their mathematical expression.

Note that interpolation constraints are also used to obtain tight automated convergence guarantees via another approach, relying on tools from robust control [16]–[18]. The idea is to automatically derive a quadratic Lyapunov function that serves as certificate for the linear convergence of various first-order methods, by analyzing a single step of the method. Provided that the method is analysed over a class of functions for which interpolation constraints are available, the Lyapunov function obtained is tight in the sense that it guarantees the best decrease rate possible for any *single* (or fixed number of) iteration of the algorithm and for a quadratic Lyapunov function. Compared to the PEP framework, one of the assets of this approach is its tractability, since the SDP to solve is of small size. One the other hand, it cannot deal with time-varying steps or sublinear rates, and the PEP framework can achieve better rates by analyzing several steps.

III. FUNCTIONS AND OPERATORS INTERPOLATION

In this section, we review the functions classes \mathcal{F} for which the following question has been answered: what necessary and sufficient constraints, i.e. interpolation constraints, must a set of data (points, functions values and (sub)gradients) satisfy to ensure its \mathcal{F} -interpolability:

Definition 1 (\mathcal{F} -interpolability): Let I be a finite set of indices. A set of triples $\{(x_i, g_i, f_i)\}_{i \in I}$ is \mathcal{F} -interpolable if

$$\exists F \in \mathcal{F} : \begin{cases} f_i = F(x_i) \\ g_i = \nabla F(x_i) \end{cases} \quad \forall i \in I.$$

The concept of interpolation constraint applied to optimization was introduced in [2], to tightly analyze first-order methods on the class of smooth strongly convex functions $\mathcal{F}_{\mu,L}$, classically defined by:

$$\begin{cases} f(x) \geq f(y) + \langle \nabla f(y), x - y \rangle + \frac{\mu}{2} \|x - y\|^2 \\ \|\nabla f(x) - \nabla f(y)\| \leq L \|x - y\| \end{cases} \quad \forall x, y. \quad (8)$$

When $\mu = 0$, definition (8) amounts to definition (4) of \mathcal{F}_L . Separately, the discretized versions of the two constraints in (8) are interpolation constraints: for instance, a data set $\{(x_i, g_i, f_i)\}_{i \in I}$ is interpolable by a strongly convex function if and only if it satisfies

$$f_i \geq f_j + \langle g_j, x_i - x_j \rangle + \frac{\mu}{2} \|x_i - x_j\|^2 \quad \forall i, j \in I.$$

However, it can be shown that the juxtaposition of these two interpolation constraints, i.e. the discretized version of Definition (8), is not tight, while an interpolation constraint for $\mathcal{F}_{\mu,L}$ is given by :

Theorem 1 ($\mathcal{F}_{\mu,L}$ -interpolation constraint [2], [19]): A set of triples $\{(x_i, g_i, f_i)\}_{i \in I}$ is $\mathcal{F}_{\mu,L}$ -interpolable if

$$\begin{aligned} f_i &\geq f_j + \langle g_j, x_i - x_j \rangle + \frac{\|g_i - g_j\|^2}{2(L - \mu)} \\ &+ \frac{\mu L \|x_i - x_j\|^2}{2(L - \mu)} - \frac{\mu}{L - \mu} \langle g_i - g_j, x_i - x_j \rangle \quad \forall i, j, \end{aligned}$$

when $\mu \neq L$, and

$$\begin{aligned} f_i &\geq f_j + \frac{1}{2} \langle g_i + g_j, x_i - x_j \rangle \\ &+ \frac{\|g_i - g_j - L(x_i - x_j)\|^2}{L} \quad \forall i, j, \end{aligned}$$

otherwise.

Hence, an analysis relying on these interpolation constraints is a priori tighter than any analysis relying on different definitions of this function class, as shown in Fig. 3 for $\mu = 0$. It was later shown in [20] that this result also holds for negative values of μ as well, that is for smooth *weakly convex* functions, i.e. convex up to the addition of a quadratic term. Interpolation constraints for $\mathcal{F}_{\mu,L}$ that do not explicitly imply function values, in case one would not want to use these as variables, also exist, based on Rockafellar's cyclic monotonicity conditions [3], [21]:

Theorem 2 ($\mathcal{F}_{\mu,L}$ -interpolation without function values):

A set of pairs $\{(x_i, g_i)\}_{i \in I}$ is $\mathcal{F}_{\mu,L}$ -interpolable, in the sense that $\exists F \in \mathcal{F}_{\mu,L} : g_i = \nabla F(x_i) \quad \forall i \in I$, if for any cyclic sequence $(x_{i_1}, g_{i_1}), \dots, (x_{i_N}, g_{i_N}), (x_{i_1}, g_{i_1})$,

$$\begin{aligned} &\sum_{k=1}^N \langle g_{i_k}, x_{i_k} - x_{i_{k+1}} \rangle - \frac{\langle g_{i_k}, g_{i_k} - g_{i_{k+1}} \rangle}{L} \\ &- \mu \langle x_{i_k}, x_{i_k} - x_{i_{k+1}} \rangle - \frac{\mu \langle x_{i_k}, g_{i_k} - g_{i_{k+1}} \rangle}{L} \geq 0. \end{aligned}$$

However, checking these constraints rapidly becomes intractable when the data set size increases, so that when possible, introducing function values as variables is more efficient.

In [2], interpolation constraints are also given for smooth functions with bounded subgradient $\mathcal{C}_{L,M}$ and strongly convex functions with bounded domain $\mathcal{S}_{\mu,D}$. While, in general, the juxtaposition of two interpolation constraints does not remain an interpolation constraint, it is the case in these particular cases:

Theorem 3 ($\mathcal{C}_{L,M}$ -interpolation constraints [3]): A set of triples $\{(x_i, g_i, f_i)\}_{i \in I}$ is $\mathcal{C}_{L,M}$ -interpolable if

$$\begin{cases} f_i \leq f_j + \langle g_j, x_i - x_j \rangle + \frac{L}{2} \|x_i - x_j\|^2 \\ \|g_i\| \leq M. \end{cases}$$

On the other hand, one can show that this property does not hold anymore for the class of weakly convex functions with bounded subgradient, for which the juxtaposition of the two constraints is not interpolable [19]. Finally, interpolation constraints are also available for indicator functions [3]:

$$i_C(x) = \begin{cases} 0 & \text{if } x \in C \\ 1 & \text{otherwise.} \end{cases}$$

for C a closed convex set. Let \mathcal{I}_M be the class of indicator functions bounded in radius by some M , i.e. $\|x\| \leq M \quad \forall x \in C$. These functions are of especially high importance for projections in constrained optimization, and their interpolation constraints are given by:

Theorem 4 (\mathcal{I}_M -interpolation constraints [3]): A set of triples $\{(x_i, g_i, f_i)\}_{i \in I}$ is \mathcal{I}_M -interpolable if

$$\begin{cases} f_i = 0 \\ \langle g_j, x_i - x_j \rangle \leq 0 \\ \|x_i\| \leq M. \end{cases}$$

Aside from these interpolation constraints for function classes, we can adapt Definition 1 to consider operator classes \mathcal{Q} . In this work, we focus on real operators which are defined as multi-dimensional mappings $Q : \mathbb{R}^d \rightarrow \mathbb{R}^m$.

Definition 2 (\mathcal{Q} -interpolability): Let I be a finite set of indices. A set of pairs $\{(x_i, q_i)\}_{i \in I}$ is \mathcal{Q} -interpolable if

$$\exists Q \in \mathcal{Q} : q_i = Q(x_i) \quad \forall i \in I.$$

Note that operator interpolation can be viewed as zeroth order function interpolation for multi-valued functions. Moreover, the concept of interpolation without function values, e.g Theorem 2, can be seen as operator interpolation for a subclass of operators, coherent as gradients of a function. Interpolation constraints exist for (strongly) monotone \mathcal{M}_μ , cocoercive \mathcal{C}_β and Lipschitz operators \mathcal{N}_L [10], [22], [23]:

Theorem 5 ($\mathcal{M}_\mu, \mathcal{C}_\beta, \mathcal{N}_L$ -interpolation [10]): A set of pairs $\{(x_i, q_i)\}_{i \in I}$ is

$$\mathcal{M}_\mu\text{-interpolable if } \forall i, j, \langle q_i - q_j, x_i - x_j \rangle \geq \mu \|x_i - x_j\|^2,$$

$$\mathcal{C}_\beta\text{-interpolable if } \forall i, j, \langle q_i - q_j, x_i - x_j \rangle \geq \beta \|q_i - q_j\|^2,$$

$$\mathcal{N}_L\text{-interpolable if } \forall i, j, \|q_i - q_j\| \leq L \|x_i - x_j\|.$$

However, no interpolation result is known for operators combining several of these properties.

As explained in Section II, combining the PEP framework with any of these interpolation constraints, when PEP-representable (i.e. linear in the function values and scalar products of points and (sub)gradients), allows tightly analyzing a wide variety of first-order methods on these function and operator classes. Even when interpolation constraints are not PEP-representable, involving e.g. absolute values or exponentials, it holds that any theoretical analysis based on them will be a priori tighter than an analysis based on non interpolable constraints, which explains why deriving interpolation constraints has an interest far beyond the PEP framework.

IV. LINEAR OPERATOR INTERPOLATION

A. Motivations

Many optimization problems involve a particular subclass of operators that has not yet been described in Section III, that is linear operators $\mathcal{Q}(x) = Mx$. These problems, including a large number of classical optimization problems as total variation deblurring, basis pursuit, resource allocation, etc., contain the linear operator in their objective function, e.g.,

$$\min_x f(x) + h(Mx),$$

or in their constraints

$$\min_x f(x) \text{ s.t. } Mx = b,$$

possibly with more terms in the objective function or constraint (e.g. $f(x) + g(M_1x) + h(M_2x)$ or $M_1x + M_2y = b$).

Consequently, many methods have been designed to solve such structured problems, e.g. Chambolle-Pock method [24], primal-dual fixed point method [25] or alternating direction method of multipliers [26]. They typically combine gradient and linear operator queries. Therefore, tightly analyzing these methods in PEP requires obtaining interpolation constraints representing a linear operator or, more precisely, representing its effect.

In this section, we consider finite-dimensional linear operators, i.e. matrices. The matrices considered are not necessarily symmetric and have their maximal singular value bounded by some $L \geq 0$. In other words, we work with the following set of linear operators:

$$\mathcal{L}_L = \{M : \sigma_{\max}(M) \leq L\}.$$

B. Interpolation Constraints for Linear Operators

To allow tightly analyzing methods involving linear operators, we now explain how to represent the application of a linear operator in PEP.

Let us consider a function $F(x) = h(Mx)$, where $h \in \mathcal{F}$, and its (sub)-gradient $\nabla F(x_i) = M^T \nabla h(Mx_i)$ at x_i . Introducing intermediate variables y_i , u_i and v_i allows decomposing the gradient as

$$\begin{aligned} y_i &= Mx_i, \\ u_i &= \nabla h(y_i), \\ v_i &= M^T u_i = \nabla F(x_i). \end{aligned}$$

If interpolation constraints are known for \mathcal{F} , the existence of some $h \in \mathcal{F}$ such that $u_i = \nabla h(y_i)$ is guaranteed by imposing the corresponding interpolation constraints on the set $\{(y_i, u_i)\}_{i \in I}$, involving potentially unused function values h_i if necessary, see Section III for more details. To impose constraints $y_i = Mx_i$ and $v_i = M^T u_i$, one could consider M as a variable of the PEP and try to impose conditions on its singular values. However, the problem would then become non convex since involving the multiplication of variables. We rather take the view of interpolation constraints for operators and impose convex constraints on the sets $\{(x_i, y_i)\}$ and $\{(u_i, v_i)\}$, guaranteeing the existence of a matrix $M \in \mathcal{L}_L$ interpolating them.

Definition 3 (\mathcal{L}_L -interpolability): Let $I = \{1, \dots, N\}$ a set of indices. Sets of pairs $\{(x_i, y_i)\}_{i \in I}$ and $\{(u_i, v_i)\}_{i \in I}$ are \mathcal{L}_L -interpolable if

$$\exists M \in \mathcal{L}_L : \begin{cases} y_i = Mx_i, & \forall i \in I, \\ v_i = M^T u_i, & \forall i \in I. \end{cases}$$

In [27], interpolation constraints for the class \mathcal{L}_L are derived:

Theorem 6 (\mathcal{L}_L -interpolation constraints): Sets of pairs $\{(x_i, y_i)\}_{i \in I}$ and $\{(u_i, v_i)\}_{i \in I}$ are \mathcal{L}_L -interpolable if

$$\begin{cases} X^T V = Y^T U, \\ Y^T Y \preceq L^2 X^T X, \\ V^T V \preceq L^2 U^T U, \end{cases}$$

where $X = (x_1 \dots x_N)$, $Y = (y_1 \dots y_N)$, $U = (u_1 \dots u_N)$ and $V = (v_1 \dots v_N)$.

Observe that these constraints are convex semidefinite constraints on the scalar products of the points and therefore PEP-representable. They can hence be used to analyze worst-case performance of any method involving linear operators via PEP. In [27], the authors exploit this result to analyze, e.g., the Chambolle-Pock method [24] or the gradient method on a structured function $h(Mx)$, where $h \in F$ and \mathcal{F} is a function class for which interpolation constraints are known. Note that Theorem 6 can be extended to symmetric linear operators with bounded eigenvalues, see [27].

V. NETWORK MATRICES INTERPOLATION

In this section, we focus on the interpolation of a specific set of finite-dimensional linear operators: the set of symmetric doubly-stochastic matrices, which have a given range of eigenvalues. These matrices are often used to represent consensus steps over a network of agents, which are a building block of most distributed optimization methods. Therefore, having such interpolation constraints allows building a PEP computing worst-case guarantees for distributed optimization algorithms. The main difference with the set of linear operator \mathcal{L}_L characterized in Section IV is the stochasticity of the matrix, required for having an averaging consensus.

A. Distributed Optimization

Distributed optimization methods aim at minimizing the average of local functions that are distributed across a network of agents $\{1, \dots, V\}$. They are used, for example, in large-scale machine learning or sensor networks for various reasons, including data privacy, easy scaling to large size problems, robustness to failure and possible speed up. Formally, the goal of decentralized optimization algorithms is to collaboratively solve the following optimization problem:

$$\underset{x \in \mathbb{R}^d}{\text{minimize}} \quad f(x) = \frac{1}{V} \sum_{a=1}^V f_a(x),$$

where $f_a : \mathbb{R}^d \rightarrow \mathbb{R}$ is the private function locally held by agent a . Each agent a holds its own version $x_a \in \mathbb{R}^d$ of the decision variable x , performs local computations and exchanges local information with its neighbors to come to an agreement on the minimizer x^* of the global function f . In this work, we focus on decentralized algorithms where the exchanges of information take the form of an average consensus on some quantity, e.g., on the x_a . One of the first and simplest algorithm of this form is the *distributed gradient descent* (DGD) [28], [29] for which iteration i can be described as

$$\begin{aligned} y_{a,i} &= \sum_{b=1}^V w_{ab} x_{b,i}, & \text{for } a = 1, \dots, V. & \quad (9) \\ x_{a,i+1} &= y_{a,i} - \alpha_i \nabla f_a(x_{a,i}), & \text{for } a = 1, \dots, V. & \end{aligned}$$

The vectors $y_{a,i} \in \mathbb{R}^d$ represent the result of the consensus step. The matrix $W \in \mathbb{R}^{V \times V}$ contains all the averaging

weights w_{ab} and is called the network or averaging matrix. Consensus step (9) can be written as a matrix multiplication where the variables of all the agents are stacked in single columns $\mathbf{x}_i, \mathbf{y}_i \in \mathbb{R}^{Vd}$:

$$\mathbf{y}_i = (W \otimes I_d) \mathbf{x}_i, \quad \text{with} \quad \mathbf{y}_i = \begin{bmatrix} y_{1,i} \\ \vdots \\ y_{V,i} \end{bmatrix} \quad \text{and} \quad \mathbf{x}_i = \begin{bmatrix} x_{1,i} \\ \vdots \\ x_{V,i} \end{bmatrix},$$

where \otimes denotes the Kronecker product and I_d the identity matrix of size d . This corresponds to applying a linear operator with a particular matrix shape and properties. In what follows, we consider the following set of network matrices, for a fixed number of agents $V > 1$ and a fixed bound $\lambda \in [0, 1)$ on the non-principal eigenvalues:

$$\mathcal{W}_\lambda = \left\{ W \in \mathbb{R}^{V \times V} : \begin{array}{l} W \text{ is symmetric and } \lambda_1(W) = 1, \\ -\lambda \leq \lambda_V(W) \leq \dots \leq \lambda_2(W) \leq \lambda \end{array} \right\}.$$

The performance guarantees from the literature, see [29] for a survey, usually add a non-negativity assumption for the network matrices, by qualifying them as *doubly-stochastic*, but these guarantees are in fact often valid for all the matrices from the set \mathcal{W}_λ . Representing \mathcal{W}_λ in PEP through interpolation constraints is therefore important to obtain general and comparable bounds for DGD or any other decentralized algorithm involving consensus steps.

B. Interpolation Constraints for Consensus Steps

We now present interpolation constraints to describe the effect of consensus steps with an averaging matrix from \mathcal{W}_λ . This allows to formulate the performance estimation problem for any decentralized algorithm using averaging consensus. To analyze such algorithms with the PEP framework, we need to find the worst-case for the local function and the sequence of local iterates of each agent. The same techniques as in Section II can be applied to discretize the problem, using proper interpolation constraints on each local function. In this section, we focus on the representation of the consensus steps that are part of these algorithms:

$$\mathbf{y}_i = (W \otimes I_d) \mathbf{x}_i \quad \text{with } W \in \mathcal{W}_\lambda.$$

As in Section IV, we cannot add matrix W as variable of the PEP problem because it makes it intractable but we rather impose convex constraints on the (scalar products of the) iterates \mathbf{y}_i and \mathbf{x}_i that are necessary for the existence of the desired network matrix W . The idea is thus to obtain a relaxed representation of the network matrix by formulating necessary constraints on the discretized points. As explained in Section II, sufficiency of the formulation can be verified a posteriori and will be discussed below.

Definition 4 (\mathcal{W}_λ -interpolability): Let I be a set of indices of consensus steps. A set of pairs $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i \in I}$ is \mathcal{W}_λ -interpolable if,

$$\exists W \in \mathcal{W}_\lambda : \mathbf{y}_i = (W \otimes I_d) \mathbf{x}_i \quad \text{for } i \in I.$$

While we do not have necessary and sufficient interpolability constraints for \mathcal{W}_λ , Theorem 7 below provides necessary conditions that have been shown to lead to good

results in practice. It is expressed in terms of the matrices $X = (\mathbf{x}_1 \dots \mathbf{x}_N)$, $Y = (\mathbf{y}_1 \dots \mathbf{y}_N)$, as well as their decomposition in two terms, the average and centered parts:

$$X = (\bar{X} \otimes \mathbf{1}_V) + X_\perp, \quad Y = (\bar{Y} \otimes \mathbf{1}_V) + Y_\perp,$$

where \bar{X} and \bar{Y} are agents average vectors in $\mathbb{R}^{d \times N}$, defined as $\bar{X}_{\cdot i} = \frac{1}{V} \sum_{a=1}^V x_{a,i}$, $\bar{Y}_{\cdot i} = \frac{1}{V} \sum_{a=1}^V y_{a,i}$ for $i = 1, \dots, N$ and $\mathbf{1}_V = [1 \dots 1]^T \in \mathbb{R}^V$.

Theorem 7 (\mathcal{W}_λ -interpolation necessary constraints [30]): If a set of pairs $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i \in I}$ is \mathcal{W}_λ -interpolable, then it holds that

$$\text{(Average preservation)} \quad \bar{X} = \bar{Y}, \quad (10)$$

$$\text{(Symmetry)} \quad X_\perp^T Y_\perp = Y_\perp^T X_\perp, \quad (11)$$

$$\text{(Variance reduction)} \quad Y_\perp^T Y_\perp \preceq \lambda^2 X_\perp^T X_\perp. \quad (12)$$

Constraints (10) of average preserving come from the double-stochasticity of the network matrix W ; the symmetry constraint (11) directly reflects the symmetry of W ; and the constraint (12) of variance reducing is a consequence of the bound λ on the eigenvalues of $W \in \mathcal{W}_\lambda$. All these constraints can be expressed as linear constraints in terms of scalar products of the iterates and are therefore PEP-representable, allowing to automatically compute performance guarantees of a given distributed optimization methods. The guarantees obtained are valid for any averaging matrix $W \in \mathcal{W}_\lambda$. This theorem can be extended to general range of non-principal eigenvalues $[\lambda^-, \lambda^+] \neq [-\lambda, \lambda]$ with $\lambda^- \leq \lambda^+ \in (-1, 1)$, see [30, Theorem 1].

In contrast with previous sections, the constraints from Theorem 7 are necessary, but not sufficient to the \mathcal{W}_λ -interpolability, which remains an open question. This means that these constraints describe a feasible set containing at least all the sets of iterates X, Y that are \mathcal{W}_λ -interpolable, but maybe also some others, for which there exists no matrix $W \in \mathcal{W}_\lambda$ such that $Y = (W \otimes I_d)X$. Therefore, using these constraints in a PEP provides a valid upper bound on the worst-case performance of the algorithm we analyze. If the worst-case solution obtained with PEP corresponds to iterates for which there is a matrix $W^* \in \mathcal{W}_\lambda$ such that $Y = (W^* \otimes I_d)X$, then we know that the worst-case bound is exact. Moreover, based on the results of Section IV, one can show that we can always recover a matrix $M \in \mathbb{R}^{Vd \times Vd}$ which is symmetric and with $\lambda_1(M) = 1$, $-\lambda \leq \lambda_V(M) \leq \dots \leq \lambda_2(M) \leq \lambda$ but for which we cannot guarantee the shape $M = W \otimes I_d$.

Another way of verifying the exactness of the PEP bound is to obtain the same worst-case value using an exact PEP problem in which the network matrix W is fixed a priori, with some guess of worst-case network matrix W^* . This technique allowed us to verified exactness of the PEP bounds for different algorithms in our different experiments [30], [31] and [32]. This is shown, for example, for 10 iterations of DGD in Fig. 4. The figure also shows the large improvement achieved by the PEP bound with respect to the theoretical bound from the literature for constant step-size [29], especially for large values of λ , corresponding to situations where the network is poorly connected.




Fig. 4: This plot is inspired from [30, Fig. 2] and shows the evolution with λ of the worst-case performance of $N = 10$ iterations of DGD with $A = 3$ agents and $\alpha = \frac{1}{\sqrt{N}}$. The plot shows (i) the theoretical bound from [29] (in pink), largely above (ii) the worst-case performance obtained with PEP using constraints from Theorem 7 (in blue) and (iii) the exact worst-case performance for the averaging matrix $W^{(1)}$ (in green). The blue and green curves are matching, which indicates the tightness of the spectral PEP bound for DGD.

VI. CONCLUSIONS

We have shown how choices made when translating conceptual assumptions on functions, operators etc. into algebraic relations between quantities involved in an algorithm could introduce significant conservatism in the derivation of worst-case bounds on optimization algorithm, before even beginning to combine these relations. This motivates the study of interpolation constraints: algebraic relations that do not introduce any degree of conservatism because they are necessary and sufficient for the existence of a function or operator consistent with the quantities analyzed.

We have reviewed the interpolation constraints for different classes of functions and general operators, but also for a spectral class of linear operators, that are used in many methods designed to solve structured problems, and for a spectral class of symmetric and stochastic matrices, that are typically used in distributed optimization algorithms.

By definition, interpolation constraints allow in principle obtaining tight bounds on the performance of any optimization algorithm provided one can combine them in an appropriate, or optimal, manner. This is the goal of performance estimation problems (PEP), which obtain bounds by solving an optimization problem of which the algebraic descriptions of the objects at stakes and the algorithmic description are constraints, and which lead to (numerical) tight bounds if interpolation constraints are used. PEPs motivated the study of interpolation constraints, but we stress that their benefit are also valid in classical analysis, for instance when analyzing function classes or methods that are not PEP representable, and advise therefore always using interpolation constraints when they are available.

There remain, however, several challenges. First, interpolation constraints are only available for a subset of the interesting classes of functions and operators. To the best of our knowledge, no interpolation constraint is known yet for e.g. blockwise smooth functions [7], Holder smooth functions [33], smooth functions satisfying a Polyak-Łojasiewicz (PL) condition [34], or even functions that are smooth and convex but only defined on a subset of \mathbb{R}^d , and for which condition (6) has been shown to be too strict [35]. Available interpolation constraints also concern mostly first-order descriptions; there are no tractable interpolation constraints combining information about first and second-order derivatives, preventing their use in the analysis of second-order methods. Second, we have no principled way of combining interpolation constraints to tightly describe function classes defined by several properties. Indeed, it is known that simply juxtaposing the interpolation constraints for two classes of functions, e.g. smooth and convex functions, does in general not provide an interpolation constraint for the intersection of these classes, e.g. smooth convex functions. Finally, understanding the actual impact of using approximate (or only necessary) interpolation constraints also remains an open question.

REFERENCES

- [1] Y. Nesterov, *Introductory lectures on convex optimization: A basic course*. Springer Science & Business Media, 2003, vol. 87.
- [2] 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.
- [3] 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.
- [4] A. B. Taylor, J. M. Hendrickx, and F. Glineur, “Exact worst-case performance of first-order methods for composite convex optimization,” *SIAM Journal on Optimization*, vol. 27, 12 2015.
- [5] Y. Drori and M. Teboulle, “Performance of first-order methods for smooth convex minimization: A novel approach,” *Mathematical Programming*, vol. 145, 06 2012.
- [6] 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, pp. 455–476, 2018.
- [7] Y. Kamri, J. M. Hendrickx, and F. Glineur, “On the worst-case analysis of cyclic coordinate-wise algorithms on smooth convex functions,” *arXiv preprint arXiv:2211.17018*, 2022.
- [8] A. Taylor and F. Bach, “Stochastic first-order methods: non-asymptotic and computer-aided analyses via potential functions,” in *Conference on Learning Theory*. PMLR, 2019, pp. 2934–2992.
- [9] R.-A. Dragomir, A. B. Taylor, A. d’Aspremont, and J. Bolte, “Optimal complexity and certification of bregman first-order methods,” *Mathematical Programming*, pp. 1–43, 2021.
- [10] E. K. Ryu, A. B. Taylor, C. Bergeling, and P. Giselsson, “Operator splitting performance estimation: Tight contraction factors and optimal parameter selection,” *SIAM Journal on Optimization*, vol. 30, no. 3, pp. 2251–2271, 2020.
- [11] J. Lee, S. Yi, and E. K. Ryu, “Convergence analyses of davis-yin splitting via scaled relative graphs,” *arXiv preprint arXiv:2207.04015*, 2022.
- [12] E. De Klerk, F. Glineur, and A. B. Taylor, “On the worst-case complexity of the gradient method with exact line search for smooth strongly convex functions,” *Optimization Letters*, vol. 11, pp. 1185–1199, 2017.
- [13] M. Barré, A. Taylor, and A. d’Aspremont, “Complexity guarantees for polyak steps with momentum,” in *Conference on Learning Theory*. PMLR, 2020, pp. 452–478.
- [14] A. B. Taylor, J. M. Hendrickx, and F. Glineur, “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.
- [15] 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.
- [16] L. Lessard, B. Recht, and A. Packard, “Analysis and design of optimization algorithms via integral quadratic constraints,” *SIAM Journal on Optimization*, vol. 26, no. 1, pp. 57–95, 2016.
- [17] A. Taylor, B. Van Scoy, and L. Lessard, “Lyapunov functions for first-order methods: Tight automated convergence guarantees,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 4897–4906.
- [18] M. Upadhyaya, S. Banert, A. B. Taylor, and P. Giselsson, “Automated tight lyapunov analysis for first-order methods,” 2023.
- [19] A. Rubbens and J. M. Hendrickx, “A constraint-based approach to function interpolation, with application to performance estimation for weakly convex optimisation,” in preparation.
- [20] T. Rotaru, F. Glineur, and P. Panagiotis, “Tight convergence rates of the gradient method on hypoconvex functions,” *arXiv preprint arXiv:2203.00775*, 2022.
- [21] R. T. Rockafellar, *Convex analysis*. Princeton university press, 1997, vol. 11.
- [22] H. Whitney, “Analytic extensions of differentiable functions defined in closed sets,” *Hassler Whitney Collected Papers*, pp. 228–254, 1934.
- [23] H. H. Bauschke, P. L. Combettes et al., *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2011, vol. 408.
- [24] A. Chambolle and T. Pock, “A first-order primal-dual algorithm for convex problems with applications to imaging,” *Journal of mathematical imaging and vision*, vol. 40, pp. 120–145, 2011.
- [25] P. Chen, J. Huang, and X. Zhang, “A primal-dual fixed point algorithm for minimization of the sum of three convex separable functions,” *Fixed Point Theory and Applications*, vol. 2016, no. 1, pp. 1–18, 2016.
- [26] D. Gabay and B. Mercier, “A dual algorithm for the solution of nonlinear variational problems via finite element approximation,” *Computers & mathematics with applications*, vol. 2, no. 1, pp. 17–40, 1976.
- [27] N. Bousselmi, J. M. Hendrickx, and F. Glineur, “Interpolation conditions for linear operators and applications to performance estimation problems,” *arXiv preprint arXiv:2302.08781*, 2023.
- [28] K. Yuan, Q. Ling, and W. Yin, “On the convergence of decentralized gradient descent,” *SIAM Journal on Optimization*, vol. 26, 10 2013.
- [29] 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.
- [30] S. Colla and J. M. Hendrickx, “Automatic performance estimation for decentralized optimization,” *IEEE Transactions on Automatic Control*, pp. 1–15, 2023.
- [31] —, “Automated worst-case performance analysis of decentralized gradient descent,” in *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021, pp. 2627–2633.
- [32] —, “Automated performance estimation for decentralized optimization via network size independent problems,” in *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022, pp. 5192–5199.
- [33] R. Bonic and J. Frampton, “Smooth functions on banach manifolds,” *Journal of Mathematics and Mechanics*, vol. 15, no. 5, pp. 877–898, 1966.
- [34] H. Karimi, J. Nutini, and M. Schmidt, “Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition,” in *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2016, Riva del Garda, Italy, September 19-23, 2016, Proceedings, Part I 16*. Springer, 2016, pp. 795–811.
- [35] Y. Drori, “On the properties of convex functions over open sets,” *arXiv preprint arXiv:1812.02419*, 2018.