Noname manuscript No.
(will be inserted by the editor)

A Relaxation-based Probabilistic Approach for PDE-constrained Optimization under Uncertainty with Pointwise State Constraints

Drew P. Kouri · Mathias Staudigl · Thomas M. Surowiec

Received: date / Accepted: date

Abstract We consider a class of convex risk-neutral PDE-constrained optimization problems subject to pointwise control and state constraints. Due to the many challenges associated with almost sure constraints on pointwise evaluations of the state, we suggest a relaxation via a smooth functional bound with similar properties to well-known probability constraints. First, we introduce and analyze the relaxed problem, discuss its asymptotic properties, and derive formulae for the gradient the adjoint calculus. We then build on the theoretical results by extending a recently published online convex optimization algorithm (OSA) to the infinite-dimensional setting. Similar to the regret-

DPK was sponsored by the Department of Energy Office of Science under the Early Career Research Program. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.
MST and TMS acknowledge the support of the COST Action CA16228 "European Network for Game Theory" for support, which made the initial discussions for this topic possible.

D. P. Kouri
Sandia National Laboratories
P.O. Box 5800
MS-1320, Albuquerque, NM, USA
E-mail: dpkouri@sandia.gov

M. Staudigl
Department of Advanced Computing Sciences (DACS)
Maastricht University, The Netherlands
E-mail: m.staudigl@maastrichtuniversity.nl

T. M. Surowiec
Department of Numerical Analysis and Scientific Computing
Simula Research Laboratory
Kristian Augusts gate 23, 0164 Oslo, Norway
E-mail: thomasms@simula.no

based analysis of time-varying stochastic optimization problems, we enhance the method further by allowing for periodic restarts at pre-defined epochs. Not only does this allow for larger step sizes, it also proves to be an essential factor in obtaining high-quality solutions in practice. The behavior of the algorithm is demonstrated in a numerical example involving a linear advection-diffusion equation with random inputs. In order to judge the quality of the solution, the results are compared to those arising from a sample average approximation (SAA). This is done first by comparing the resulting cumulative distributions of the objectives at the optimal solution as a function of step numbers and epoch lengths. In addition, we conduct statistical tests to further analyze the behavior of the online algorithm and the quality of its solutions. For a sufficiently large number of steps, the solutions from OSA and SAA lead to random integrands for the objective and penalty functions that appear to be drawn from similar distributions.

Keywords Optimization under Uncertainty · PDE-Constrained Optimization · State Constraints · Probability Constraints · Expectation Constraints · First-Order Methods · Stochastic Approximation

Mathematics Subject Classification (2010) 49M20, 49M41, 65K05, 65K10, 90C15, 90C25

1 Introduction

In this paper, we propose a comprehensive relaxation-based approach for the numerical solution of a risk-neutral PDE-constrained optimization problem subject to control and pointwise state constraints. For our algorithm develop-opent, we assume that the objective function has the general form

$$j(z) \triangleq \mathbb{E}_{\mathbb{P}}[J(z,\boldsymbol{\xi})], \tag{1.1}$$

where $J(z,\boldsymbol{\xi})$ is convex in the control variables z almost surely (a.s.) and $\boldsymbol{\xi}$ denotes the random inputs. In the context of PDE-constrained optimization, we will assume that the solution operator $z \mapsto u_{\boldsymbol{\xi}}(z)$ of the PDE (as an implicit function of the controls z) is included in the definition of $J(z,\boldsymbol{\xi})$ so that

$$J(z,\boldsymbol{\xi}) \triangleq \widehat{J}(u_{\boldsymbol{\xi}}(z), z, \boldsymbol{\xi}).$$

One common example of this is the standard tracking-type function

$$\widehat{J}(u_{\boldsymbol{\xi}}(z), z, \boldsymbol{\xi}) \triangleq \frac{1}{2}\|u_{\boldsymbol{\xi}}(z) - u_d\|_U^2 + \frac{\alpha}{2}\|z\|_Z^2 \tag{1.2}$$

in which u_d is a deterministic target state, $\|\cdot\|_U$ and $\|\cdot\|_Z$ are appropriate Hilbert space norms, and $\alpha > 0$.

In many applications, the state variable $u_{\boldsymbol{\xi}}(z)$ is restricted by certain pointwise bounds, which may be in the form of a static obstacle, a minimum temperature, or a maximum allowable concentration. Consequently, we consider the situation where the PDE solution $u_{\boldsymbol{\xi}}(z)$ must satisfy the pointwise constraint

$$u_{\boldsymbol{\xi}}(z) \geq \psi \quad \text{a.e./a.s.}, \tag{1.3}$$

where ψ is a prescribed function.

Minimizing the convex objective function (1.1) while enforcing the state constraint (1.3) represents significant challenges. In particular, any algorithmic approach necessarily relies on a tractable formulation of the uncertain constraints. One of the main contributions of this paper is the derivation of such a tractable formulation that follows a convex relaxation approach. To illustrate our approach, consider the relaxation of the a.s. constraints into probability constraints. Here, we only require the bound to hold with high probability. In doing so, we fix a pre-specified confidence level $p \in (0,1)$ and require that

$$\mathbb{P}\left(u_{\boldsymbol{\xi}}(z) \geq \psi \quad \text{a.e.}\right) \geq p. \tag{1.4}$$

Joint chance constraints of this form are popular in many engineering problems such as hydro reservoir control and mechanics [2, 8]. However, optimization problems with probability constraints are difficult already in finite-dimensions, where there exist many structural results related to differentiability and convexity [10, 16, 26]. See [30] or [32, Chap. 4] for additional details. In this work, we make further progress on this relevant question. Our basic ideas can be described as follows:

1. Replace the almost sure state constraint (1.3) by a degenerate *global* expectation constraint of the type

$$\Phi(z) = 0, \tag{1.5}$$

   for some suitably defined expectation functional $\Phi(z)$. This relaxation approach is inspired in part by the notion of "integrated chance constraints," defined in [5, 6, 15].
2. Relax this constraint by introducing a small slack $\varepsilon > 0$, and impose the weaker restriction

$$\Phi(z) \leq \varepsilon. \tag{1.6}$$

This relaxation strategy allows us to cast the PDE-constrained optimization problem as a stochastic optimization problem with expectation constraints of the form

$$\min_{z \in \mathcal{Z}_{\mathrm{ad}}} j(z) \quad \text{subject to} \quad \Phi(z) \leq \varepsilon, \tag{1.7}$$

where $\mathcal{Z}_{\mathrm{ad}}$ is the set of admissible controls. This reformulation allows for global violations of the state constraint similar to the chance constraint (1.4), but with the decisive advantage of guaranteed convexity and smoothness, independent of the nature of the contaminating noise. This offers an alternative to probability constraints that is numerically tractable and exhibits similar properties (see Section 3). The price we pay for this regularity is that both the objective function, as well as the constraints appear in terms of an expectation. Hence, to solve such an optimization problem numerically, we need to resort to sampling and simulation-based techniques. Accordingly, we develop a new online stochastic approximation (OSA) algorithm for the relaxed formulation (1.7), which is designed to solve an infinite-dimensional stochastic optimization problem with expectations appearing in the objective function and the constraints.

Recently, the PDE-constrained optimization community has devoted a significant amount of interest in developing numerical schemes for control problems with randomly perturbed coefficients. Many approaches employ an empirical approximation for the random integrands using either Monte-Carlo [22], Quasi-Monte Carlo [14], Multilevel Monte Carlo [35] or adaptive sparse grids [19, 20] to obtain a deterministic PDE-constrained problem. The deterministic solvers employed in these (and related) papers are typically inexact Newton approaches, which allow for massive parallelization for the gradient and Hessian-vector products and avoid expensive matrix computations. Recently, stochastic approximation methods in the spirit of [29] have been adapted to stochastic PDE-constrained optimization problems in [12, 13] and [28]. Variance reduction ideas, originating from machine learning, have witnessed applications in this field [27] as well.

From the broader perspective of stochastic optimization, not many numerical schemes for solving convex stochastic optimization problems subject to expectation constraints are known. The only existing alternative to the scheme developed here is essentially the recent work by [23]. Here, the authors adapt the proximal gradient method to functional constrained optimization problems, where the constraint needs to be sampled at the current position of the algorithm. Their scheme is very flexible, and extends to the Bregman setup easily. Additionally, given a desired solution accuracy $\epsilon_{\text{sol}} > 0$, they state an $O(\epsilon_{\text{sol}}^{-2})$ iteration complexity result in terms of the ergodic average. This gives an upper bound on the number of iterations needed to arrive at a solution of accuracy $\epsilon_{\text{sol}} > 0$ on the order of $\epsilon_{\text{sol}}^{-2}$. However, their analysis is restricted to finite-dimensional optimization problems. In addition, it seems that their analysis relies on an a priori decomposition of the set of iterations, which appears to be challenging, at least, to verify in practice.

In contrast to the method in [23], we instead follow a machine learning inspired approach and extend a recent online convex optimization algorithm with time-varying constraints, due to [37]. We enrich their method to allow for periodic restarts at pre-defined epochs, similar to the regret-based analysis of time-varying stochastic optimization problems laid out in [3, 9]. The restart-based algorithm does not improve the theoretical iteration complexity, but allows us to use larger, epoch-dependent step sizes, which can be a crucial factor in practice.

Recently, a conditional gradient framework for stochastically constrained convex programming problems has been introduced in [36]. We believe that this approach can be applied to the current situation as well, once the technical details associated with the infinite-dimensionality stemming from the PDE-constraint are resolved. However, the iteration complexity of that scheme is $O(\epsilon_{\text{sol}}^{-6})$, which is much slower than the present $O(\epsilon_{\text{sol}}^{-2})$. At least for PDE-constrained optimization, it is important to note that each stochastic gradient (or subgradient) evaluation requires the solution of two partial differential equations. Therefore, even a rough solution accuracy of $\epsilon_{\text{sol}} = 10^{-3}$ indicates that the conditional gradient approach would require approximately 10^{18} PDE solves, which is clearly not acceptable from a computational standpoint. In contrast, the Newton-based solvers mentioned above typically require several million PDE solves, which is comparable to the 10^6 PDE solves for an $O(\epsilon_{\text{sol}}^{-2})$ algorithm.

The rest of this paper is organized as follows. Section 2 introduces the notation used in this paper and describes the PDE model in detail. Section 3 describes the stochastic optimization problem subject to the PDE constraint and the obstacle. In that section the penalty reformulation is also explained in detail. Section 4 presents the algorithm we use to solve the relaxed problem (1.7) numerically. Finally, Section 5 provides a numerical example with a rigorous statistical analysis comparing the proposed OSA algorithm and a sample average approximation (SAA) approach.

2 Notation and PDE Description

Throughout, $(\Omega, \mathcal{F}, \mathbb{P})$ is a complete probability space and $D \subset \mathbb{R}^n$ is an open and bounded set. Given a Banach space U, we denote the associated norm by $\|\cdot\|_U$. If V is another Banach space, then we denote the space of continuous linear operators mapping U into V by $\mathrm{Lin}(U, V)$. The topological dual space of V is accordingly $V^* \triangleq \mathrm{Lin}(V, \mathbb{R})$. Given a functional $h : V \to \mathbb{R}$, we denote the Fréchet derivative of h at $v \in V$ by $h'(v) \in V^*$. When V is a Hilbert space, we denote its inner product by $(\cdot, \cdot)_V$ and assume that the associated norm is given by $\|f\|_V^2 \triangleq (f, f)_V$. When the context is clear, we will omit the space from the norm and inner product. If V is a Hilbert space and $h : V \to \mathbb{R}$ is Fréchet differentiable, we denote the gradient associated with $h'(v)$ by $\nabla h(v) \in V$. Finally, we denote strong (norm) convergence by '$\to$' and weak convergence by '$\rightharpoonup$.'

We consider PDEs with random inputs, the solutions of which are random fields. Consequently, it is necessary to introduce various function spaces for the solution variables. We denote the Lebesgue space of (equivalence classes of) square-integrable functions from D to $\mathbb{R}$ by $L^2(D)$. We further denote the Sobolev space of $L^2(D)$ functions with $L^2(D)$ weak derivatives by $H^1(D)$ and the closed subspace of $H^1(D)$ functions with zero boundary trace by $H_0^1(D)$. We denote the topological dual space of $H_0^1(D)$ by $H^{-1}(D)$. For more on Lebesgue and Sobolev spaces, see e.g., [1]. The Sobolev $H_0^1(D)$ is a common solution space for deterministic linear elliptic PDEs. However, since we consider PDEs with random inputs, the solution is a random field, which belongs to the Bochner space

$$\mathcal{U} \triangleq L^2(\Omega, \mathcal{F}, \mathbb{P}; H_0^1(D))$$

of $\mathcal{F}$-strongly measurable square $\mathbb{P}$-integrable mappings $v : \Omega \to H_0^1(D)$, endowed with the natural norm

$$\|v\|_{\mathcal{U}}^2 \triangleq \mathbb{E}_{\mathbb{P}}\left[\|v\|_{H_0^1(D)}^2\right].$$

Other Bochner spaces, e.g., $L^\infty(\Omega, \mathcal{F}, \mathbb{P}; H^1(D))$, are defined analogously [17]. For $v \in \mathcal{U}$, we recall that $v(\omega) \in H_0^1(D)$ a.s. When needed, we denote the evaluation of $v(\omega)$ at $x \in D$ by $v(\omega, x)$.

We refer to the optimization variables $z \in \mathcal{Z} \triangleq L^2(D)$ as the "controls" and the PDE solutions $u \in \mathcal{U}$ as the "states." Let (Ξ, Σ) be a measurable space. The measurable mapping $\boldsymbol{\xi} : \Omega \to \Xi$ is a random element that parametrizes the PDE coefficients. Without loss of generality, we assume that $\boldsymbol{\xi}(\Omega) = \Xi$. We distinguish between the random element $\boldsymbol{\xi}$ and its possible values $\xi \in \Xi$ using bold text. Given a control $z \in \mathcal{Z}$, the state solves the random PDE: Find $u \in \mathcal{U}$ that satisfies

$$\begin{aligned} &\mathbb{E}_{\mathbb{P}}\left[\int_D \kappa(x, \boldsymbol{\xi}) \nabla_x u(\cdot, x) \cdot \nabla_x v(\cdot, x)\, dx\right] \\ &\quad = \mathbb{E}_{\mathbb{P}}\left[\int_D ((B(\boldsymbol{\xi})z)(x) + f(x, \boldsymbol{\xi})) v(\cdot, x)\, dx\right] \quad \forall\, v \in \mathcal{U}. \end{aligned} \tag{2.1}$$

For the development of numerical methods, it is often convenient to consider the equivalent "parametric" weak form of (2.1): For fixed $z \in \mathcal{Z}$ and $\xi \in \Xi$, find $u_\xi \in H_0^1(D)$ that satisfies

$$\begin{aligned} &\int_D \kappa(x, \xi) \nabla u_\xi(x) \cdot \nabla \phi(x)\, dx \\ &\quad = \int_D ((B(\xi)z)(x) + f(x, \xi)) \phi(x)\, dx \quad \forall\, \phi \in H_0^1(D). \end{aligned} \tag{2.2}$$

The equivalence to (2.1) in one direction can be seen by choosing test functions of the type $v(\omega, x) = \chi_A(\omega)\varphi(x)$ such that $A \in \mathcal{F}$ and $\varphi \in H_0^1(D)$ and then substituting these into (2.1). Since (2.1) holds for all $A \in \mathcal{F}$, we obtain (2.2). The reverse direction from (2.2) to (2.1) is a special case of the nonlinear elliptic setting considered in [22], where it is shown that measurability follows from a measurable selection theorem and integrability from standard a priori estimates for linear elliptic PDE.

We now postulate several basic assumptions on the data for (2.1).

Assumption 1 *We assume that*

1. $f(\cdot, \boldsymbol{\xi}) \in L^\infty(\Omega, \mathcal{F}, \mathbb{P}; \mathcal{Z})$;
2. *There exist positive constants* $0 < \kappa_0 \le \kappa_1 < +\infty$ *such that*

$$\kappa_0 \le \kappa(\cdot, \xi) \le \kappa_1 \quad a.e. \quad \forall\, \xi \in \Xi;$$

3. *The control operator* $\omega \mapsto B(\boldsymbol{\xi}(\omega)) : \Omega \to \mathrm{Lin}(\mathcal{Z}, H^{-1}(D))$ *is uniformly measurable, essentially bounded and completely continuous:*

$$z_n \rightharpoonup z \text{ in } \mathcal{Z} \quad \Rightarrow \quad B(\boldsymbol{\xi})z_n \to B(\boldsymbol{\xi})z \text{ in } H^{-1}(D) \quad a.s.$$

Under Assumption 1, the Lax-Milgram Lemma applies to show that a solution to (2.1) exists and is unique. Owing to the linearity of the PDE, the solution can be written as $S(z) + u_f$, where $S(z)$ is the unique solution of (2.1) obtained by setting $f \equiv 0$ and u_f is the unique solution of (2.1) obtained by setting $z \equiv 0$. Note that S is a bounded linear operator from $\mathcal{Z}$ into $\mathcal{U}$. Using the equivalence between (2.1) and (2.2), we may also use the ξ-dependent solution operator $S_\xi : \mathcal{Z} \to H_0^1(D)$ such that the solution to (2.2) is $S_\xi(z) + u_{f(\cdot,\xi)}$ for $\xi \in \Xi$.

Remark 1 As shown in [22], a much larger class of semilinear elliptic PDEs can be analyzed in an optimization context. However, significant difficulty arises from the state constraint. For this reason, we have chosen to develop the theory and algorithm in this paper for the linear elliptic case. One of the main difficulties in the semilinear case is a full convergence proof of the algorithm, since the nonlinearity renders the optimization problems nonconvex.

3 The Optimization Problem

3.1 The Objective Function, Constraints, and Further Data Assumptions

We consider optimal control problems with the convex objective function (1.1) and the control constraints

$$\mathcal{Z}_{\mathrm{ad}} \triangleq \{v \in \mathcal{Z} \mid a \le v \le b \text{ a.e.}\}. \tag{3.1}$$

In addition, we impose a unilateral state constraint on the solution operator $z \mapsto S(z)$ of the form:

$$S_{\boldsymbol{\xi}}(z)[x] \ge \psi(x, \boldsymbol{\xi}) - u_{f(\cdot,\boldsymbol{\xi})}(x) \quad \text{for a.a. } x \in D \quad \text{a.s.} \tag{3.2}$$

In order to prove existence of solutions and analyze the algorithm below, we require the following mild regularity assumptions.

Assumption 2 *We assume that*

1. $D \subset \mathbb{R}^n$ *is an open and bounded set with Lipschitz boundary* $\Gamma \subset \mathbb{R}^{n-1}$*;*
2. $a, b \in L^2(D)$ *with* $a < b$ *a.e.;*
3. $\psi : \overline{D} \times \Xi \to \mathbb{R}$ *is continuous, satisfies* $\psi(x, \xi) \le 0$ *for all* $(x, \xi) \in \Gamma \times \Xi$*, and* $\psi(\cdot, \xi) \in H^1(D)$ *for all* $\xi \in \Xi$*;*
4. $j(\cdot)$ *is proper, weakly lower-semicontinuous, and convex.*

For readability, we define the state-constrained feasible set by

$$\mathsf{C} \triangleq \left\{ z \in \mathcal{Z} \,|\, \mathbb{P}\left(S_{\boldsymbol{\xi}}(z)[x] \geq \psi(x, \boldsymbol{\xi}) - u_{f(\cdot,\boldsymbol{\xi})}(x) \quad \text{for a.a. } x \in D \right) = 1 \right\}. \quad (3.3)$$

The next assumption ensures that the "original" optimization problem admits a solution. Moreover, it is necessary for the asymptotic statements of the relaxation approach.

Assumption 3 *The feasible set is nonempty, i.e.,* $\mathsf{C} \cap \mathcal{Z}_{\text{ad}} \neq \varnothing$.

We will later discuss a weaker assumption in the context of the relaxed problems below for cases in which it is unclear whether Assumption 3 holds.

As an operator from $\mathcal{Z}$ into $\mathcal{U}$, S is bounded and linear. From this, one readily shows that (under Assumption 3) $\mathsf{C} \cap \mathcal{Z}_{\text{ad}}$ is a nonempty, closed, bounded, and convex set in $\mathcal{Z}$ and therefore weakly compact. Consequently, the convex infinite-dimensional stochastic optimization problem

$$\inf_{z \in \mathcal{Z}} \{ j(z) \,|\, z \in \mathsf{C} \cap \mathcal{Z}_{\text{ad}} \} \quad \text{(P)}$$

admits a solution $z^\star$, which is unique if j is strictly convex as in (1.2). Under further regularity conditions on D, e.g., if D is of type $C^{1,1}$ or a convex polyhedron, it is possible to show that $u(z) = S(z) + u_f \in L^\infty(\Omega, \mathcal{F}, \mathbb{P}; H^2(D) \cap H_0^1(D))$. This fact was discussed in detail in the recent paper [11]. It provides sufficient regularity to argue for the existence of a Slater point for the state constraint and derive optimality conditions for (P) using standard Lagrangian duality as in [4, Chap. 3]. However, for general domains D, it is unclear whether multiplier-based optimality conditions exist.

3.2 A Relaxation Approach

Our relaxation strategy employs a fairly broad class of penalty functions, defined by the following properties.

Definition 1 (Regular Penalty) A function $\varphi : \mathbb{R} \to \mathbb{R}_+$ is a *regular penalty* if the following conditions hold:

C.1 φ is a continuously differentiable convex function;
C.2 $\varphi(r) = 0$ for all $r \leq 0$;
C.3 $\varphi(r) > 0$ for all $r > 0$;
C.4 φ has a Lipschitz continuous gradient with modulus L_φ:

$$(\forall t, s \in \mathbb{R}) : |\varphi'(t) - \varphi'(s)| \leq L_\varphi |t - s|. \quad (3.4)$$

Assumption C.4 implies that for all $t, s \in \mathbb{R}$

$$|\varphi(t) - \varphi(s) - \varphi'(s)(t - s)| \leq \frac{L_\varphi}{2} |t - s|^2. \quad (3.5)$$

In addition, since $\varphi \geq 0$, $\varphi(0) = 0$, and φ' is Lipschitz, we have $|\varphi'(t)| \leq L_\varphi |t|$. A concrete example for a regular penalty function $\varphi : \mathbb{R} \to \mathbb{R}_+$ is as follows.

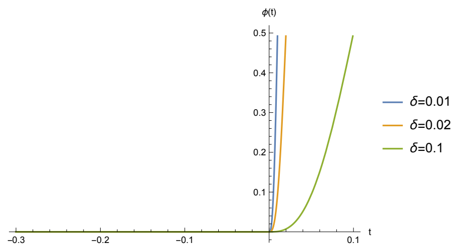

Fig. 1: The regular penalty function (3.6) for various values of δ.

Example 1 Consider the function $r : \mathbb{R} \to \mathbb{R}$ given by

$$r(t) \triangleq \begin{cases} t - \frac{1}{2} & \text{if } t \geq 1, \\ t^3 - \frac{t^4}{2} & \text{if } t \in (0,1), \\ 0 & \text{else.} \end{cases}$$

It can be easily checked that $r(\cdot)$ satisfies conditions C.1-C.4. In particular, r is globally Lipschitz smooth with $|r'(t)| \leq 1$ for all $t \in \mathbb{R}$. Let $\delta > 0$ and define

$$\varphi(t) \triangleq r(\delta^{-1}t). \tag{3.6}$$

Then, $|\varphi'(t)| \leq \frac{1}{\delta} \equiv L_\varphi$ for all $t \in \mathbb{R}$. See Figure 1 for an illustration. We see that as $\delta \to 0^+$ the map $t \mapsto r(t/\delta)$ approximates the indicator function equal to 0 on $(-\infty, 0]$ and $+\infty$ in $(0, \infty)$.

We employ regular penalty functions to convert the pointwise constraints (3.2) to expectation constraints. Let $\theta(z, \xi) \in H^1(D)$ for $z \in \mathcal{Z}$ be defined pointwise by

$$\theta(z, \xi) \triangleq \psi(\cdot, \xi) - (u_{f(\cdot,\xi)} + S_\xi(z)). \tag{3.7}$$

The evaluation of $\theta(z, \xi)$ at $x \in D$ is denoted by $\theta(z, \xi)[x]$. Observe that $\theta(z, \boldsymbol{\xi}) \leq 0$ a.s. whenever $z \in \mathsf{C} \cap \mathcal{Z}_{\mathrm{ad}}$. Next, consider the function $\Phi : \mathcal{Z} \to \mathbb{R}_+$, defined by

$$\Phi(z) \triangleq \mathbb{E}_{\mathbb{P}}[F(\theta(z, \boldsymbol{\xi}))] \qquad \forall z \in \mathcal{Z}, \tag{3.8}$$

where for measurable functions $v : D \to \mathbb{R}$, F is defined by

$$F(v) \triangleq \int_D \varphi(v(x))\, dx. \tag{3.9}$$

Using Φ, we arrive at the family of relaxed optimization problems

$$\min_{z \in \mathcal{Z}_{\mathrm{ad}}} \{ j(z) \,|\, \Phi(z) \leq \varepsilon \}, \tag{P$_\varepsilon$}$$

where $\varepsilon > 0$ is a given tolerance for constraint violation. The next lemma shows that our relaxation approach has similar implications in terms of constraint violation like standard chance constraints: Making Φ small guarantees constraint satisfaction with high probability.

Lemma 1 *Let $\varphi : \mathbb{R} \to \mathbb{R}_+$ be a regular penalty function. For a fixed $z \in \mathcal{Z}$, we have*

$$\theta(z, \boldsymbol{\xi}) \leq 0 \quad \text{a.e./a.s.} \qquad \Leftrightarrow \qquad \Phi(z) = 0.$$

Proof Fix $z \in \mathcal{Z}$. Clearly, if $\theta(z, \boldsymbol{\xi}) \leq 0$ a.e./a.s., then $\varphi(\theta(z, \boldsymbol{\xi})) = 0$ a.e./a.s. Consequently, $\Phi(z) = 0$. Conversely, suppose $\Phi(z) = 0$, and let $M \subset D \times \Omega$ be a set of positive measure on which $\theta(z, \boldsymbol{\xi}) > 0$. Without loss of generality, assume that $\theta(z, \boldsymbol{\xi}) \leq 0$ holds on $(D \times \Omega) \setminus M$. Then, by the properties of φ, we would have $\varphi(\theta(z, \boldsymbol{\xi})) > 0$ a.e. on M. It follows that

$$0 < \int_M \varphi(\theta(z, \boldsymbol{\xi}(\omega))[x]) \, dx \, d\mathbb{P}(\omega) \leq \mathbb{E}_{\mathbb{P}} \left[\int_D \varphi(\theta(z, \boldsymbol{\xi})[x]) \, dx \right] = \Phi(z) = 0,$$

which is a contradiction. ∎

Remark 2 The use of Φ is linked to probability constraints and the original local setting. As Lemma 1 shows, $\Phi(z) \leq 0$ is equivalent to the original constraint (3.2). By Markov's inequality, we have for any $c > 0$ that

$$\mathbb{P}(F(\theta(z, \boldsymbol{\xi})) \geq c) \leq \frac{1}{c} \Phi(z).$$

Choosing $\varepsilon = c^2$ in (1.6), any point $z \in \mathcal{Z}_{\text{ad}}$ that satisfies (1.6) would then fulfill

$$\mathbb{P}(F(\theta(z, \boldsymbol{\xi})) < c) \geq 1 - c.$$

Thus, although the relaxation allows for violations of the local state constraint, the probability of such events can be tamed by choosing moderate values for c; e.g., $c = 10^{-2}$ and $\varepsilon = 10^{-4}$.

3.3 Analysis of Φ

We now derive the regularity properties of the function Φ that are needed for the development of optimization algorithms. We start by providing technical lemmata on integral functionals that make the verifying the properties of Φ more readable.

Lemma 2 *Let $\varphi : \mathbb{R} \to \mathbb{R}_+$ be a regular penalty function. Define the functional $F : H_0^1(D) \to \mathbb{R}_+$ by (3.9). Then, F is convex, globally Lipschitz continuous with modulus L_F, and Fréchet differentiable.*

Proof Convexity is straightforward. For global Lipschitz continuity, we have for any u, $v \in H_0^1(D)$ the inequality

$$|F(v) - F(u)| \leq \int_D |\varphi(v(x)) - \varphi(u(x))| \, dx \leq L_\varphi \|v - u\|_{L^1(D)}.$$

Since D is bounded, we also have Lipschitz continuity in $L^2(D)$ with modulus $\hat{L}_F \triangleq L_\varphi \mathrm{Vol}(D)^{1/2}$. Furthermore, letting c_{emb} be the constant from Poincaré's inequality, we can set $L_F \triangleq c_{\mathrm{emb}} \hat{L}_F$ and obtain

$$|F(v) - F(u)| \le L_F \|v - u\|_{H_0^1(D)}.$$

Next, we prove differentiability. Fix arbitrary $u, v \in H_0^1(D)$ and $t > 0$. From (3.5), we have for almost all $x \in D$:

$$-\frac{t^2 L_\varphi}{2}|u(x)|^2 \le \varphi(v(x) + tu(x)) - \varphi(v(x)) - t\varphi'(v(x))u(x) \le \frac{t^2 L_\varphi}{2}|u(x)|^2.$$

By dividing both sides by $t > 0$, integrating over x, and letting $t \to 0^+$, the Lebesgue dominated convergence theorem implies that

$$F'(v; u) = \lim_{t \to 0^+} \frac{F(v + tu) - F(v)}{t} = \int_D \varphi'(v) u \, dx.$$

Since $|\varphi'(v)| \le L_\varphi |v|$ a.e. and $|v| \in H_0^1(D)$, $(\varphi'(u), \cdot)_{L^2(D)}$ defines a bounded linear functional on $H_0^1(D)$. Hence, $F'(v; \cdot)$ is linear and continuous on $H_0^1(D)$ and therefore, F is Gâteaux differentiable at v. Since F is Lipschitz, Gâteaux and Hadamard differentiability coincide, see e.g., [4, Prop. 2.49].

To prove that F is in fact Fréchet differentiable, let $u, v \in H_0^1(D)$. For almost every $x \in D$, we have

$$|\varphi(v(x) + u(x)) - \varphi(v(x)) - \varphi'(v(x))u(x)| \le \frac{L_\varphi}{2}|u(x)|^2.$$

Integrating over $x \in D$ and using Poincaré's inequality, we have

$$\frac{|F(v + u) - F(v) - F'(v)u|}{\|u\|_{H_0^1(D)}} \le \frac{L_\varphi c_{\mathrm{emb}}}{2} \|u\|_{H_0^1(D)}.$$

Passing to the limit as $\|u\|_{H_0^1(D)} \to 0$ we obtain the assertion. ■

Proposition 1 *The function $\Phi : \mathcal{Z} \to \mathbb{R}_+$ defined in (3.8) is convex, globally Lipschitz continuous, and continuously differentiable with derivative*

$$\Phi'(z)h = \mathbb{E}_{\mathbb{P}}\left[(\varphi'(\theta(z, \boldsymbol{\xi})), -S(h))_{L^2(D)}\right] \quad \forall\, h \in \mathcal{Z} \tag{3.10}$$

and Lipschitz continuous gradient,

$$\nabla \Phi(z) = -\mathbb{E}_{\mathbb{P}}[\eta_{\boldsymbol{\xi}}] \tag{3.11}$$

where $\eta = \eta_\xi(z) \in H_0^1(D)$, for fixed $\xi \in \Xi$, fulfills

$$\int_D \kappa(x, \xi) \nabla \eta(x) \cdot \nabla v(x) \, dx = \int_D \varphi'(\theta(z, \xi)) v(x) \, dx \quad \forall\, v \in H_0^1(D). \tag{3.12}$$

Proof The convexity of Φ is a result of the linearity of S and the convexity of the regular penalty function φ. To see that Φ is globally Lipschitz, we appeal to the properties of F. The Lipschitz continuity of F established in Lemma 2 immediately gives for all $u_1, u_2 : \Omega \to H_0^1(D)$,

$$|F(u_1) - F(u_2)| \leq L_F \|u_1 - u_2\|_{H_0^1(D)} \quad \text{a.s.}$$

Squaring both sides followed by taking the expectation yields

$$\mathbb{E}_{\mathbb{P}}\left[|F(u_1) - F(u_2)|^2\right] \leq L_F \mathbb{E}_{\mathbb{P}}\left[\|u_1 - u_2\|^2_{H_0^1(D)}\right].$$

Applying Jensen's inequality to the left-hand side and taking the square root of both sides yields

$$|\mathbb{E}_{\mathbb{P}}[F(u_1)] - \mathbb{E}_{\mathbb{P}}[F(u_2)]| \leq \sqrt{L_F \mathbb{E}_{\mathbb{P}}[\|u_1 - u_2\|^2_{H_0^1(D)}]}.$$

It follows that

$$\mathcal{U} \ni u \mapsto \widehat{\Phi}(u) \triangleq \mathbb{E}_{\mathbb{P}}[F(u)]$$

is globally Lipschitz. Clearly, $\widehat{\Phi}(\theta(z, \boldsymbol{\xi})) = \Phi(z)$. Using to the linearity of S, we deduce that

$$\begin{aligned} |\Phi(z_1) - \Phi(z_2)| &= |\mathbb{E}_{\mathbb{P}}[F(\theta(z_1, \boldsymbol{\xi}))] - \mathbb{E}_{\mathbb{P}}[F(\theta(z_2, \boldsymbol{\xi}))]| \\ &\leq \sqrt{L_F \mathbb{E}_{\mathbb{P}}[\|S(z_1) - S(z_2)\|^2_{H_0^1(D)}]} \\ &\leq M \|z_1 - z_2\|_{L^2(D)} \end{aligned}$$

for some $M > 0$. Notice that the constant terms u_f and ψ disappear in the Lipschitz bound on F before estimating from above by the $H_0^1(D)$-norm.

Next, we show that Φ is continuously Fréchet differentiable with Lipschitz derivative. Since $S(z) + u_f$ is a continuous affine mapping from $\mathcal{Z}$ into $\mathcal{U}$, we need only consider $\widehat{\Phi}(\cdot)$. Using the chain rule [34, Th. 20.9], we would then obtain the formula in (3.10). For any $u, h \in \mathcal{U}$, we will demonstrate that

$$\widehat{\Phi}'(u)h = \mathbb{E}_{\mathbb{P}}\left[\int_D \varphi'(u)h \, dx\right]. \tag{3.13}$$

Using the same arguments as in Lemma 2, the functional on the righthand side in (3.13) is clearly bounded and linear on $\mathcal{U}$.

By the assumptions on the regular penalty function φ, we have for all $u, h \in \mathcal{U}$,

$$\begin{aligned} -\frac{L_\varphi}{2}|h(\omega,x)|^2 &\leq [\varphi(u(\omega,x) + h(\omega,x)) - \varphi(u(\omega,x))] - \varphi'(u(\omega,x))h(\omega,x) \\ &\leq \frac{L_\varphi}{2}|h(\omega,x)|^2. \end{aligned}$$

for almost all $(x,\omega) \in D \times \Omega$. Integrating over $D \times \Omega$ yields

$$
\begin{aligned}
-\frac{L_\varphi}{2}\mathbb{E}_{\mathbb{P}}\left[\int_D |h(\cdot,x)|^2\,dx\right] &\le \mathbb{E}_{\mathbb{P}}\left[\int_D (\varphi(u+h)-\varphi(u))\,\,dx\right] - \mathbb{E}_{\mathbb{P}}\left[\int_D \varphi'(u)h\,dx\right]\\
&= \left(\widehat{\Phi}(u+h)-\widehat{\Phi}(u)\right) - \mathbb{E}_{\mathbb{P}}\left[\int_D \varphi'(u)h\,dx\right]\\
&\le \frac{L_\varphi}{2}\mathbb{E}_{\mathbb{P}}\left[\int_D |h(\cdot,x)|^2\,dx\right].
\end{aligned}
$$

Since $\mathcal{U}$ is continuously embedded into $L^2(\Omega,\mathcal{F},\mathbb{P};\mathcal{Z})$, there exists a constant $c>0$ such that

$$
\left|\left(\widehat{\Phi}(u+h)-\widehat{\Phi}(u)\right) - \mathbb{E}_{\mathbb{P}}\left[\int_D \varphi'(u)h\,dx\right]\right| \le c\|h\|_{\mathcal{U}}^2
$$

and it follows that $\widehat{\Phi}$ is Fréchet differentiable.

In addition, there exists a constant $c'>0$ for any $u,v \in \mathcal{U}$ and $h \in \mathcal{U}$ such that for almost all $\omega \in \Omega$ we have

$$
\begin{aligned}
(F'(u(\omega)) - F'(v(\omega)))h(\omega) &\le L_\varphi \int_D |u(\omega)-v(\omega)||h(\omega)|\,dx\\
&\le c'\|u(\omega)-v(\omega)\|_{H_0^1(D)}\|h(\omega)\|_{H_0^1(D)}.
\end{aligned}
\tag{3.14}
$$

Taking the expectation on both sides and applying Hölder's inequality yields

$$
\widehat{\Phi}'(u)h - \widehat{\Phi}'(v)h \le c'\|u-v\|_{\mathcal{U}}\|h\|_{\mathcal{U}}.
$$

Replacing h by $-h$ yields the same inequality with

$$
|\widehat{\Phi}'(u)h - \widehat{\Phi}'(v)h| \le c'\|u-v\|_{\mathcal{U}}\|h\|_{\mathcal{U}}.
$$

Then taking the supremum over all $h \in \mathcal{U}$ with $\|h\|_{\mathcal{U}} = 1$ proves that $\widehat{\Phi}' : \mathcal{U} \to \mathcal{U}^*$ is Lipschitz.

It remains to verify (3.11). For any $h \in \mathcal{Z}$, it holds that

$$
(\varphi'(\theta(z,\boldsymbol{\xi})), -S(h))_{\mathcal{Z}} = (-S^*\varphi'(\theta(z,\boldsymbol{\xi})), h)_{\mathcal{Z}} \quad \text{a.s.},
$$

where the canonical embedding $\iota_{H_0^1 \hookrightarrow L^2}$ in front of $S(h)$ and its adjoint $\iota_{L^2 \hookrightarrow H^{-1}}$ in front of the φ'-term have been suppressed in the notation above. Taking the expectation of both sides and applying Fubini's theorem, we then have

$$
\Phi'(z)h = (\mathbb{E}_{\mathbb{P}}[-S^*\varphi'(\theta(z,\boldsymbol{\xi}))], h)_{\mathcal{Z}}
$$

for any pair $z,h \in \mathcal{Z}$. Finally, by recognizing that $\eta_\xi(z) = -S_\xi^*\varphi'(\theta(z,\xi))$ solves the adjoint equation (3.12), we deduce the final assertion. ■

3.4 Pointwise and Uniform Bounds on the Objective and Reduced Gradients

The numerical method considered in the next section requires a number of bounds on the objective functional, the constraint functional, and the gradients of the integrands in j and Φ. Based on the structural assumptions, these can be easily verified using standard a priori bounds in the analysis of linear elliptic PDEs. We require the following additional properties on j.

Assumption 4 *We assume that $j : \mathcal{Z} \to \mathbb{R}$ has the form*

$$j(z) = \mathbb{E}_{\mathbb{P}}[J_0(u_{\boldsymbol{\xi}}(z))] + J_1(z),$$

where $J_0, J_1 : \mathcal{Z} \to \mathbb{R}$ are convex and continuously differentiable and, given $z, h \in \mathcal{Z}$, we have

$$j'(z)h = \mathbb{E}_{\mathbb{P}}[J_0'(u_{\boldsymbol{\xi}}(z))u_{\boldsymbol{\xi}}'(z)h] + J_1'(z)h$$

Moreover, we assume that J_0' and J_1' map bounded sets into bounded sets.

These assumptions are directly inspired by the model problem (1.2). We gather several usual bounds that arise from the assumptions, which are necessary for the convergence analysis in the following sections.

Proposition 2 *Let Assumptions 1, 2, and 4 hold.*

1. *There exist constants M_{obj} and M_{bd} such that*

$$|j(z)| \leq M_{\mathrm{obj}} \quad \text{and} \quad |\Phi(z)| \leq M_{\mathrm{bd}} \quad \forall z \in \mathcal{Z}_{\mathrm{ad}}. \tag{3.15}$$

2. *The gradient mapping $\nabla_z J(z, \boldsymbol{\xi})$ has the form*

$$\nabla_z J(z, \boldsymbol{\xi}) = B^*(\boldsymbol{\xi})\lambda_{\boldsymbol{\xi}} + \nabla J_1(z),$$

where for fixed $\xi \in \Xi$, λ_ξ solves

$$\int_D \kappa(x, \xi)\nabla\lambda(x)\cdot\nabla v(x)dx = -\int_D J_0'(u_\xi(z))v(x)dx \quad \forall\, v \in H_0^1(D). \tag{3.16}$$

3. *There exists a constant M_{adj} such that*

$$\|\nabla_z J(z, \boldsymbol{\xi})\|_{\mathcal{Z}} \leq M_{\mathrm{adj}} \quad \forall z \in \mathcal{Z}_{\mathrm{ad}}, \ \text{a.s.} \tag{3.17}$$

4. *There exists a constant $M_{\mathrm{ctr}} > 0$ such that*

$$\|\nabla_z (F \circ \theta)(z, \boldsymbol{\xi})\|_{\mathcal{Z}} \leq M_{\mathrm{ctr}} \quad \forall z \in \mathcal{Z}_{\mathrm{ad}}, \ \text{a.s.} \tag{3.18}$$

Proof By boundedness of $\mathcal{Z}_{\mathrm{ad}}$ and the continuity of j and Φ there exist uniform bounds M_{obj} and M_{bd} such that (3.15) in 1. holds.

Next, we derive a general bound on the solution operator $u_{\boldsymbol{\xi}}$. Recalling that $u_\xi(z) = S_\xi(z) + u_{f(\cdot,\xi)}$ solves (2.2), we can use $\phi = u_\xi$ in (2.2) as a test function. If follows from Assumption 1 that

$$\kappa_0\|u_\xi(z)\|_{H_0^1(D)} \leq c_{\mathrm{emb}}(\|B(\xi)\|_{\mathrm{op}}\|z\|_{\mathcal{Z}} + \|f(\xi)\|_{\mathcal{Z}}) \tag{3.19}$$

holds for all $\xi \in \Xi$. By assumption, $\|B(\boldsymbol{\xi})\|_{\text{op}}, \|f(\boldsymbol{\xi})\|_{\mathcal{Z}} \in L^\infty(\Omega, \mathcal{F}, \mathbb{P})$ and $\mathcal{Z}_{\text{ad}}$ is bounded. Therefore, (3.19) implies the existence of a constant M_{st} such that

$$\|u_{\boldsymbol{\xi}}(z)\|_{H_0^1(D)} \leq M_{\text{st}} \quad \forall z \in \mathcal{Z}_{\text{ad}} \quad \text{a.s.} \tag{3.20}$$

Similarly, using the standard rules of adjoint calculus, see e.g., [25, Chap. 2], [33], or [18], we derive a bound for the adjoint equation (3.16) associated with $\nabla_z J(z, \boldsymbol{\xi})$. Denoting the adjoint state by $\lambda_{\boldsymbol{\xi}}$ we have

$$\kappa_0 \|\lambda_{\boldsymbol{\xi}}\|_{H_0^1(D)} \leq c_{\text{emb}} \|J_0'(u_{\boldsymbol{\xi}}(z))\|_{\mathcal{Z}^*} \quad \text{a.s.}$$

By the properties of J_0' and (3.20), there exists a constant $M'_{\text{adj}} > 0$ such that

$$\|\lambda_{\boldsymbol{\xi}}\|_{H_0^1(D)} \leq M'_{\text{adj}} \quad \forall z \in \mathcal{Z}_{\text{ad}} \quad \text{a.s.}$$

Consequently, the following bound holds independently of z a.s.:

$$\|\nabla_z J(z, \boldsymbol{\xi})\|_{\mathcal{Z}} \leq M_{\text{adj}} \quad \text{a.s.}$$

where $M_{\text{adj}} = M'_{\text{adj}} \|B(\boldsymbol{\xi})\|_{\text{op}} + \sup_{z \in \mathcal{Z}_{\text{ad}}} \|\nabla J_1(z)\|_{\mathcal{Z}}$.

We can proceed analogously for the constraint mapping by exploiting the statements in Proposition 1 and using the associated adjoint equation (3.12). To this end, fix an arbitrary $z \in \mathcal{Z}_{\text{ad}}$. Then the stochastic gradients associated with the integrand $F(\theta(z, \boldsymbol{\xi}))$ of Φ satisfy (3.12) (excluding the embedding into $\mathcal{Z}$). Using $\eta_{\boldsymbol{\xi}} = v$ as a test function and continuing as above for J we have

$$\kappa_0 \|\eta_{\boldsymbol{\xi}}\|_{H_0^1(D)} \leq c_{\text{emb}} \|\varphi'(\theta(z, \boldsymbol{\xi}))\|_{\mathcal{Z}} \quad \text{a.s.}$$

Condition C.4 of Definition 1 implies that $\eta_{\boldsymbol{\xi}} \in L^\infty(\Omega, \mathcal{F}, \mathbb{P}; H_0^1(D))$. Indeed, this condition readily gives

$$|\varphi'(\theta(z, \boldsymbol{\xi}))| = |\varphi'(\theta(z, \boldsymbol{\xi})) - \varphi'(0)| \leq L_\varphi |\theta(z, \boldsymbol{\xi})| \quad \text{a.s.}$$

Combining this with (3.19), there exists a constant $M_{\text{ctr}} > 0$ (in ω and $z \in \mathcal{Z}_{\text{ad}}$) such that $\|\eta_{\boldsymbol{\xi}}\|_{H_0^1(D)} \leq M_{\text{ctr}}$ a.s. We then have for all $z \in \mathcal{Z}_{\text{ad}}$

$$\|\nabla_z (F \circ \theta)(z, \boldsymbol{\xi})\|_{\mathcal{Z}} \leq M_{\text{ctr}} \quad \text{a.s.} \tag{3.21}$$

■

3.5 Asymptotic Considerations

Using the favorable properties of the penalty functional Φ established above, we now investigate the asymptotic properties of the problem (P_ε) as $\varepsilon \downarrow 0$. For convenience, we recall (P) and (P_ε) here:

$$j_{\text{opt}} \triangleq \inf_z \{ j(z) \,|\, z \in \mathsf{C} \cap \mathcal{Z}_{\text{ad}} \}, \qquad j_{\text{opt}}(\varepsilon) \triangleq \min_{z \in \mathcal{Z}_{\text{ad}}} \{ j(z) \,|\, \Phi(z) \leq \varepsilon \}$$

Let $z^\star$ and $z^\star_\varepsilon$ denote controls satisfying $j(z^\star) = j_{\text{opt}}$ and $j(z^\star_\varepsilon) = j_{\text{opt}}(\varepsilon)$, respectively.

Proposition 3 *Suppose Assumptions 1, 2, 3, and 4 are fulfilled. Furthermore, assume that j is strongly convex. Then, for all sequences $\varepsilon_k \downarrow 0$, the sequence $\{z_k^\star\}_{k\in\mathbb{N}}$ with $z_k^\star = z_{\varepsilon_k}^\star$ converges weakly in $\mathcal{Z}$ to $z^\star$.*

Proof As discussed above, the assumptions and strong convexity ensure that both problems have unique optimal solutions $z^\star$ and $z_\varepsilon^\star$. Next, we note that by Assumption 2 the path of solutions $\{z_\varepsilon^\star\}_{\varepsilon>0}$ is uniformly bounded in $\mathcal{Z}$. Therefore, for any sequence $\varepsilon_k \downarrow 0$, the sequence of unique solutions $\{z_k^\star\}$ admits a weakly convergent subsequence $\{z_{k_l}^\star\}$ with limit point $\bar{z}$. By Proposition 1, Φ is weakly lower-semicontinuous. It follows from Lemma 1 and

$$\Phi(\bar{z}) \leq \liminf_{l\to+\infty} \Phi(z_{\varepsilon_{k_l}}^\star) \leq \liminf_{l\to+\infty} \varepsilon_{k_l} = 0,$$

that $\bar{z}$ is feasible for (P).

Next, by Assumption 3, $z \in \mathsf{C} \cap \mathcal{Z}_{\mathrm{ad}}$ is feasible for (P_ε) for all $\varepsilon > 0$. Therefore, we have that $j(z_\varepsilon^\star) \leq j(z) \quad \forall z \in \mathsf{C} \cap \mathcal{Z}_{\mathrm{ad}}$. Since j is convex and continuous, it is weakly lower-semicontinuous. Hence, for any $z \in \mathsf{C} \cap \mathcal{Z}_{\mathrm{ad}}$, we have

$$j(\bar{z}) \leq \liminf_{l\to+\infty} j(z_{k_l}^\star) \leq j(z),$$

including $z = z^\star$. Since $z^\star$ is unique and $\bar{z}$ is feasible for (P), $\bar{z} = z^\star$. Note that this same argument would hold for every weakly convergent subsequence of $\{z_k^\star\}$. It follows by the Urysohn property that the entire sequence $\{z_k^\star\}$ converges weakly to $z^\star$.

■

Remark 3 Most of the arguments in the proof of Proposition 3 can be relaxed to the convex (not strongly convex) case. The statement would then read: for all sequences $\varepsilon_k \downarrow 0$, there exist a subsequence of solutions $\{z_{k_l}^\star\}$ with $z_{k_l}^\star = z_{\varepsilon_{k_l}}^\star$ that converges weakly to a solution of (P).

Remark 4 The arguments in the proof of Proposition 3 could also be used for the case in which Assumption 3 is relaxed to require that

$$\{z \in \mathcal{Z}_{\mathrm{ad}} \mid \Phi(z) \leq \varepsilon_{\min}\} \neq \varnothing$$

for some minimal, but positive $\varepsilon_{\min}$. Therefore, if Assumption 3 cannot be verified for a given $\varepsilon > 0$, a path-following argument for the relaxed problems is still available. Either way, we have shown that the mapping

$$[0, +\infty) \ni \varepsilon \mapsto z_\varepsilon^\star \in \mathcal{Z}$$

is weakly continuous and for the case when Assumption 3 does hold:

$$\Phi(z_\varepsilon^\star) = O(\varepsilon) \text{ and } \Phi(z_\varepsilon^\star) = o(\varepsilon^q)$$

for any $q \in (0, 1)$. Furthermore, when the integrand has the form (1.2), it is clear that $j_{\mathrm{opt}} \geq j_{\mathrm{opt}}(\varepsilon) \geq 0$ for all ε, even if Assumption 3 does not hold, in which case $j_{\mathrm{opt}} = +\infty$. However, under the assumptions of Proposition 3, this special case of a strongly convex objective can be used to show that $[0, +\infty) \ni \varepsilon \mapsto z^\star_\varepsilon \in \mathcal{Z}$ is in fact strongly continuous, due to the presence of the squared-Z-norm in the objective. From this the results of Proposition 3 demonstrate that $j_{\text{opt}} : [0, \infty) \to \mathbb{R}$ is continuous everywhere on $(0, \infty)$, right continuous at 0, and decreasing as ε increases. In fact, $j_{\text{opt}}(\varepsilon)$ would decrease monotonically as ε increases.

4 The Algorithm

4.1 Stochastic Approximation

To approximate the value function $j_{\text{opt}}(\varepsilon)$, we introduce an online stochastic approximation (SA) approach in which data is drawn anew each iteration to evaluate the objective and constraint functions as well as their derivatives. Our algorithm is applicable to any convex stochastic optimization problem with expectation constraints, and hence is not specific to the PDE-constrained optimization application, although it does apply to it. As such, we assume to have access to a first-order *stochastic oracle* (SO) with the following properties. In the subsequent definition, we use the letter J to refer to quantities related to the objective function and G for the expectation constraint. Throughout this section, $(\cdot, \cdot)$ denotes the inner product on $\mathcal{Z}$ and $\| \cdot \| = \sqrt{(\cdot, \cdot)}$.

Definition 2 (Stochastic Oracle) Let $(\Omega, \mathcal{F}, (\mathcal{F}_k)_k, \mathbb{P})$ be a filtered probability space. Given a control $z \in \mathcal{Z}_{\text{ad}}$ and an iteration k, a *stochastic oracle* (SO) is a black-box device whose output is a set of random elements $J_k(z)$, $J'_k(z)$, $G_k(z)$, and $G'_k(z)$, with the following properties:

1. $J_k(z)$, $J'_k(z)$, $G_k(z)$, and $G'_k(z)$ are unbiased estimators of $j(z)$, $\nabla j(z)$, $\Phi(z) - \varepsilon$, and $\nabla \Phi(z)$, respectively, in the sense that
$$\begin{aligned} \mathbb{E}[J_k(z)|\mathcal{F}_k] &= j(z), \quad \mathbb{E}[G_k(z)|\mathcal{F}_k] = \Phi(z) - \varepsilon, \\ \mathbb{E}[J'_k(z)|\mathcal{F}_k] &= \nabla j(z), \quad \mathbb{E}[G'_k(z)|\mathcal{F}_k] = \nabla \Phi(z) \end{aligned} \tag{4.1}$$
holds a.s.
2. There exists D_1 and $D_2 > 0$ independent of $z \in \mathcal{Z}_{\text{ad}}$ and $k \geq 1$:
$$\|J'_k(z)\| \leq D_1 \text{ and } \|G'_k(z)\| \leq D_2. \tag{4.2}$$
3. There exists $M > 0$ independent of $z \in \mathcal{Z}_{\text{ad}}$ such that $|G_k(z)| \leq M$.

We recall that Proposition 2 provides the required bounds on the objective function gradient for our target PDE-constrained application. To derive the bound on G_k, Lemma 2 and (3.20) ensure that

$$|F(\theta(z, \boldsymbol{\xi})) - F(\theta(z', \boldsymbol{\xi}))| \leq \|u_{\boldsymbol{\xi}}(z) - u_{\boldsymbol{\xi}}(z')\|_{H^1_0(D)} \leq 2M_{\text{st}} \text{ a.s. } \forall\, z,\, z' \in \mathcal{Z}_{\text{ad}}.$$

Taking $z' \in \mathcal{Z}_{\text{ad}}$ such that $\theta(z', \boldsymbol{\xi}) \leq 0$, Lemma 1 indicates that $F(\theta(z', \boldsymbol{\xi})) = 0$ a.s., which yields

$$|F(\theta(z, \boldsymbol{\xi}))| \leq 2M_{\text{st}} \quad \forall z \in \mathcal{Z}_{\text{ad}} \text{ a.s.}$$

Hence, the single samples $J_k(z) = J(z, \xi_k)$ and $G_k(z) = F(\theta(z, \xi_k))$, where ξ_k is an independent and identically distributed (i.i.d.) copy of $\boldsymbol{\xi}$, produce a SO when endowed with the filtration $\mathcal{F}_0 = \{\varnothing, \Omega\}$ and $\mathcal{F}_k = \sigma(\{\xi_1, \dots, \xi_k\})$. In the following example, we see that these observations can be easily extended to (mini) batches of samples.

Example 2 A common example for an SO is to construct Monte-Carlo estimators of the random data involved in the stochastic optimization problem. To obtain such an SO, let $m \geq 1$ be a given integer (the sample budget) and assume it is easy to simulate m i.i.d. copies of the random element $\boldsymbol{\xi}$. Let $\xi_k = (\xi_k^{(1)}, \dots, \xi_k^{(m)})$ denote the sample at iteration k and set

$$J_k(z) \triangleq \frac{1}{m}\sum_{t=1}^{m} J(z, \xi_k^{(t)}), \quad J'_k(z) \triangleq \frac{1}{m}\sum_{t=1}^{m} \nabla_z J(z, \xi_k^{(t)}),$$
$$G_k(z) \triangleq \frac{1}{m}\sum_{t=1}^{m} F(\theta(z, \xi_k^{(t)})), \quad G'_k(z) \triangleq \frac{1}{m}\sum_{t=1}^{m} \nabla_z (F \circ \theta)(z, \xi_k^{(t)})$$

for all $z \in \mathcal{Z}_{\text{ad}}$. It follows directly from the above discussion that this mechanism gives rise to an admissible SO. Note that the generation of these estimators requires solving a sequence of PDEs, one for each random variable $\xi_k^{(t)}$. Hence, the computational complexity of this SO at each iteration k is $m \times C$, where C is an upper bound on the cost of evaluating the objective function, the constraint, and their derivatives.

4.2 A penalty-based first-order algorithm

Our algorithmic strategy begins with the construction of a suitable penalty function with adaptive weights. Given $z \in \mathcal{Z}_{\text{ad}}$, consider the function

$$\mathcal{L}_{\gamma,k}(z, w) \triangleq J_k(z) + \frac{w}{\gamma} G_k(z) \tag{4.3}$$

where $w \geq 0$ is a penalty parameter (chosen by the algorithm), $\gamma > 0$ is a user-specified parameter, and $k \geq 1$ is an iteration counter. The ratio w/γ measures the importance of the constraint violation over reducing the objective function value while executing the optimization algorithm, and thus we can interpret (4.3) as a penalty formulation of the original stochastic optimization problem. Querying the SO at the pair (z, w) allows us to evaluate the function $\mathcal{L}_{\gamma,k}(z, w)$, as well as

$$\mathcal{L}'_{\gamma,k}(z, w) = J'_k(z) + \frac{w}{\gamma} G'_k(z). \tag{4.4}$$

Our numerical treatment of the stochastic optimization problem builds on successive restarts of a master algorithm, to be denoted $\mathbb{X}_{\mathbb{T}}^{(z,w)}(\alpha, \gamma)$, which is formally described in Algorithm 1. This procedure takes as inputs an initial guess $(z, w) \in \mathcal{Z}_{\text{ad}} \times [0, \infty)$, an index set of iteration counters $\mathbb{T} \subseteq \mathbb{N} \cup \{0, +\infty\}$,

Algorithm 1: Master Algorithm $\mathbb{X}_{\mathbb{T}}^{(z,w)}(\alpha, \gamma)$

Input: $\mathbb{T} \subseteq \mathbb{N}$ iteration counter set, Initial condition $(z, w) \in \mathcal{Z}_{\text{ad}} \times [0, \infty)$; Parameters $\alpha, \gamma \in (0, \infty)$.
Output: Sequence $\{(\mathsf{z}^t, \mathsf{w}^t), \min \mathbb{T} \leq t \leq \sup \mathbb{T} + 1\}$.
Set $\mathsf{z}_{\min \mathbb{T}} = z$ and $\mathsf{w}_{\min \mathbb{T}} = w$;
for $t = \min \mathbb{T}, \dots, \sup \mathbb{T}$ **do**
    Compute $\mathsf{z}^{t+1} = P_{\mathcal{Z}_{\text{ad}}}(\mathsf{z}^t - \frac{\gamma}{\alpha}\mathcal{L}'_{\gamma,t}(\mathsf{z}^t, \mathsf{w}^t))$;
    Compute $\mathsf{w}^{t+1} = \max\{0, \mathsf{w}^t + G_t(\mathsf{z}^t) + (G'_t(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)\}$;
end for

as well as the user-specified parameters α and γ, whose role will be explained below. To simplify the presentation, we assume that the batch $\mathbb{T}$ is a set of adjacent integers, with smallest element $\min \mathbb{T} \in \mathbb{N} \cup \{0\}$ and largest element $\sup \mathbb{T} \in \mathbb{N} \cup \{0, +\infty\}$, i.e $\mathbb{T} = \{\min \mathbb{T}, \min \mathbb{T} + 1, \dots, \sup \mathbb{T}\}$. If $\sup \mathbb{T} = \infty$, we define $\sup \mathbb{T} + 1 \triangleq \infty$. The master process generates a sequence $\{(\mathsf{z}^t, \mathsf{w}^t); \min \mathbb{T} \leq t \leq \sup \mathbb{T} + 1\}$ via the updates

$$\mathsf{z}^{t+1} = P_{\mathcal{Z}_{\text{ad}}}(\mathsf{z}^t - \frac{\gamma}{\alpha}\mathcal{L}'_{\gamma,t}(\mathsf{z}^t, \mathsf{w}^t)),$$
$$\mathsf{w}^{t+1} = \max\{0, \mathsf{w}^t + G_t(\mathsf{z}^t) + (G'_t(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)\},$$

where $P_{\mathcal{Z}_{\text{ad}}}(z) \triangleq \operatorname{argmin}_{z' \in \mathcal{Z}_{\text{ad}}} \frac{1}{2}\|z' - z\|^2$ is the orthogonal projection onto $\mathcal{Z}_{\text{ad}}$. The first updating equation is just a projected gradient descent step, using the sampled data embodied in the random variable $\mathcal{L}'_{\gamma,t}(\mathsf{z}^t, \mathsf{w}^t)$ as first-order feedback. The second step updates the penalty parameter via a first-order approximation of the sampled constraint function. If the local linearized model has a positive value, then the iterates are moving away from the feasible set $\mathcal{Z}_{\text{ad}}$. The algorithm reacts to this by increasing the weight w^{t+1}.

The master algorithm is the basic pillar in our restart-based optimization strategy, culminating in Algorithm 2. This scheme takes as inputs a sequence of time iteration counters $\mathbb{T}_1, \dots, \mathbb{T}_s$ (the "batches"), and a corresponding sequence of optimization parameters $(\alpha_j, \gamma_j), 1 \leq j \leq s$, as well as suitably chosen initial conditions $(\mathsf{z}_j^1, \mathsf{w}_j^1), 1 \leq j \leq s$. In practical implementations, we choose the batches of nearly equal size. Specifically, given the predefined total number of iterations N, we let the user define another input parameter $\Delta_N \in \{1, \dots, N\}$, which defines the length of batches. To be precise, $s \triangleq \lceil \frac{N}{\Delta_N} \rceil$ is the number of restarts (meaning the number of calls of the master algorithm). We then set $|\mathbb{T}_j| = \Delta_N$ for $1 \leq j < s$, and $|\mathbb{T}_s| = N - (s-1)\Delta_N$. In other words, all batches except the last one have the same size Δ_N.

Once the batches have been defined in this way, we construct a sequence $\{(\mathsf{z}_j^t, \mathsf{w}_j^t); 1 \leq t \leq \sup \mathbb{T}_j + 1\}$, by calling the master algorithm $\mathbb{X}_{\mathbb{T}_j}^{(\mathsf{z}_j^1, \mathsf{w}_j^1)}(\alpha_j, \gamma_j)$. We warm-start each call to the master algorithm by setting

$$(\mathsf{z}_j^1, \mathsf{w}_j^1) = (\mathsf{z}_{j-1}^{\sup \mathbb{T}_{j-1}+1}, \mathsf{w}_{j-1}^{\sup \mathbb{T}_{j-1}+1})$$

for $j = 1, \dots, s$. By default, we set $\mathbb{T}_0 = \{0\}$, and provide inputs $z_1 = \mathsf{z}_0^1, w_1 = \mathsf{w}_0^1$ to the algorithm. In a final post-processing step, we concatenate these

Algorithm 2: Epoch-dependent online SA (OSA)

Input: $1 \leq \Delta_N \leq N$. Initial condition $z_0 \in \mathcal{Z}_{\text{ad}}$ and $w_0 = 0$;
Epoch-specific parameters $\{\alpha_j\}_{j=1}^s$ and $\{\gamma_j\}_{j=1}^s$;
Set $\mathsf{z}_0^1 = z_0$ and $\mathsf{w}_0^1 = w_0$;
Set $s \equiv \lceil \frac{N}{\Delta_N} \rceil$;
Construct batches $\mathbb{T}_0, \mathbb{T}_1, \dots, \mathbb{T}_s$ with $\mathbb{T}_0 = \{0\}$, and

$$\mathbb{T}_j = \{1, \dots, \Delta_N\} \quad \text{for } 1 \leq j < s,$$
$$\mathbb{T}_s = \{1, \dots, N - (s-1)\Delta_N\}.$$

for $j = 1, \dots, s$ **do**
    Set $\mathsf{z}_j^1 = \mathsf{z}_{j-1}^{\sup \mathbb{T}_{j-1}+1}$ and $\mathsf{w}_j^1 = \mathsf{w}_{j-1}^{\sup \mathbb{T}_{j-1}+1}$;
    Compute $\{(\mathsf{z}_j^t, \mathsf{w}_j^t); 1 \leq t \leq \sup \mathbb{T}_j + 1\}$ by calling master algorithm $\mathbb{X}_{\mathbb{T}_j}^{(\mathsf{z}_j^1, \mathsf{w}_j^1)}(\alpha_j, \gamma_j)$;
end for
For $k \in \{1, \dots, N\}$, set $z_k = \mathsf{z}_j^t$ and $w_k = \mathsf{w}_j^t$ for $k - t = (j-1)\Delta_N, t \in \mathbb{T}_j$;
Report $\bar{z}_N = \frac{1}{N} \sum_{k=1}^N z_k$.

trajectories to obtain a sequence $\{(z_k, w_k)\}_{k=1}^N$, and its ergodic average $\bar{z}_N \triangleq \frac{1}{N} \sum_{k=1}^N z_k$. The specific concatenation procedure is given by

$$z_k = \mathsf{z}_j^t \text{ and } w_k = \mathsf{w}_j^t \text{ for } k = (j-1)\Delta_N + t, 1 \leq j \leq s, t \in \mathbb{T}_j. \tag{4.5}$$

We describe this procedure in Algorithm 2.

Remark 5 Note that Algorithm 2 is still not an executable scheme since we have not specified a strategy to choose the epoch-dependent parameters α_j and γ_j. In the convergence analysis (Section 4.4), we will pin down a simple rule determining these parameters.

4.3 Preparatory Estimates for the master algorithm

Algorithm 2 is defined by s consecutive restarts of the master, where s is a user-defined parameter. Each restart differs only in the initial guess and the parameters (α_j, γ_j). We begin the analysis of the complexity of Algorithm 2 by analysing the master algorithm $\mathbb{X}_{\mathbb{T}}^{(z,w)}(\alpha, \gamma)$ for given inputs $(\mathbb{T}, (z, w), \alpha, \gamma)$. Since the restarts differ only in these input parameters, all estimates derived in this section are valid when applied to the analysis of Algorithm 2.

Let $\mathcal{A}_0 \triangleq \{\varnothing, \Omega\}$ and $\mathcal{A}_t \triangleq \sigma(\mathsf{z}^\tau; \tau \leq t)$ be the natural filtration associated with the process defining $\mathbb{X}_{\mathbb{T}}^{(z,w)}(\alpha, \gamma)$. We proceed with our analysis in several steps.

By definition,

$$\frac{1}{2}(\mathsf{w}^{t+1})^2 \leq \frac{1}{2}\left[\mathsf{w}^t + G_t(\mathsf{z}^t) + (G_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)\right]^2$$
$$= \frac{1}{2}(\mathsf{w}^t)^2 + \mathsf{w}^t\left[G_t(\mathsf{z}^t) + (G_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)\right]$$
$$+ \frac{1}{2}\left[G_t(\mathsf{z}^t) + (G_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)\right]^2.$$

Using the triangle inequality, it follows that

$$|G_t(\mathsf{z}^t) + (G_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)| \leq |G_t(\mathsf{z}^t)| + |(G_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)|$$
$$\leq M + |s_k| \leq M + D_2 R,$$

where R denotes the diameter of the feasible set $\mathcal{Z}_{\text{ad}}$. This in turn yields

$$\frac{1}{2}(\mathsf{w}^{t+1})^2 \leq \frac{1}{2}(\mathsf{w}^t)^2 + \mathsf{w}^t\left[G_t(\mathsf{z}^t) + (G_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)\right] + \frac{1}{2}(M + D_2 R)^2. \quad (4.6)$$

Lemma 3 *Consider the master algorithm* $\mathbb{X}_{\mathbb{T}}^{(z,w)}(\alpha, \gamma)$ *with pseudo-code given in Algorithm 1. Then, for all* $\min \mathbb{T} \leq t_1 \leq t_2 \leq \sup \mathbb{T}$*, we have*

$$\sum_{t=t_1}^{t_2} G_t(\mathsf{z}^t) \leq \mathsf{w}^{t_2+1} - \mathsf{w}^{t_1} + D_2 \sum_{t=t_1}^{t_2} \|\mathsf{z}^{t+1} - \mathsf{z}^t\|.$$

Proof For all t, it holds true that

$$\mathsf{w}^{t+1} = \max\{0, \mathsf{w}^t + G_t(\mathsf{z}^t) + (G_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)\}$$
$$\geq \mathsf{w}^t + G_t(\mathsf{z}^t) + (G_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)$$
$$\geq \mathsf{w}^t + G_t(\mathsf{z}^t) - \|G_t'(\mathsf{z}^t)\|\|\mathsf{z}^{t+1} - \mathsf{z}^t\|$$
$$\geq \mathsf{w}^t + G_t(\mathsf{z}^t) - D_2\|\mathsf{z}^{t+1} - \mathsf{z}^t\|.$$

Rearranging, this yields

$$G_t(\mathsf{z}^t) \leq \mathsf{w}^{t+1} - \mathsf{w}^t + D_2\|\mathsf{z}^{t+1} - \mathsf{z}^t\|.$$

Summing over $t = t_1, \dots, t_2$ verifies the claim. ■

In the following we need the Pythagorean identity

$$2\langle w - v, u - v\rangle = \|w - v\|^2 - \|w - u\|^2 + \|u - v\|^2. \quad (4.7)$$

Recall that our primal update z^t is a forward step involving the gradient estimator (4.4). The optimality condition for the update z^{t+1} therefore reads as

$$(\mathsf{z}^{t+1} - \mathsf{z}^t + \frac{\gamma}{\alpha}\mathcal{L}_{\gamma,t}'(\mathsf{z}^t, \mathsf{w}^t), z - \mathsf{z}^{t+1}) \geq 0 \quad \forall\, z \in \mathcal{Z}_{\text{ad}}.$$

This implies

$$(\gamma J_t'(\mathsf{z}^t), z - \mathsf{z}^{t+1}) + \mathsf{w}^t (G_t'(\mathsf{z}^t), z - \mathsf{z}^{t+1}) \\ \geq \alpha(\mathsf{z}^t - \mathsf{z}^{t+1}, z - \mathsf{z}^{t+1}) \quad \forall\, z \in \mathcal{Z}_{\mathrm{ad}}. \tag{4.8}$$

Convexity and (4.7) gives then the a.s. inequality

$$\mathcal{L}_{\gamma,t}(z, \mathsf{w}^t) \geq \mathcal{L}_{\gamma,t}(\mathsf{z}^t, \mathsf{w}^t) + (\mathcal{L}'_{\gamma,t}(\mathsf{z}^t, \mathsf{w}^t), \mathsf{z}^{t+1} - \mathsf{z}^t) \\ + \frac{\alpha}{\gamma}\left(\frac{1}{2}\|\mathsf{z}^{t+1} - z\|^2 - \frac{1}{2}\|\mathsf{z}^t - z\|^2 + \frac{1}{2}\|\mathsf{z}^{t+1} - \mathsf{z}^t\|^2\right)$$

for all $z \in \mathcal{Z}_{\mathrm{ad}}$. Splitting up terms, this reads explicitly as

$$\gamma J_t(z) + \mathsf{w}^t G_t(z) + \frac{\alpha}{2}\left(\|\mathsf{z}^t - z\|^2 - \|\mathsf{z}^{t+1} - z\|^2\right) \\ \geq \gamma J_t(\mathsf{z}^t) + (\gamma J_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t) + \frac{\alpha}{2}\|\mathsf{z}^{t+1} - \mathsf{z}^t\|^2 \\ + \mathsf{w}^t\left(G_t(\mathsf{z}_t) + (G_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)\right) \tag{4.9}$$

for all $z \in \mathcal{Z}_{\mathrm{ad}}$.

Lemma 4 *Consider the master algorithm $\mathbb{X}_{\mathbb{T}}^{(z,w)}(\alpha, \gamma)$ with pseudo-code given in Algorithm 1. For all $t \in \mathbb{T}$, we have*

$$\|\mathsf{z}^{t+1} - \mathsf{z}^t\| \leq \frac{\gamma}{\alpha} D_1 + \frac{\mathsf{w}^t}{\alpha} D_2. \tag{4.10}$$

Proof By choosing $z = \mathsf{z}^t \in \mathcal{Z}_{\mathrm{ad}}$ in (4.8), we readily obtain

$$\alpha\|\mathsf{z}^{t+1} - \mathsf{z}^t\|^2 \leq \gamma(J_t'(\mathsf{z}^t), \mathsf{z}^t - \mathsf{z}^{t+1}) + \mathsf{w}^t(G_t'(\mathsf{z}^t), \mathsf{z}^t - \mathsf{z}^{t+1}) \\ \overset{(4.2)}{\leq} \left(\gamma D_1 + \mathsf{w}^t D_2\right)\|\mathsf{z}^{t+1} - \mathsf{z}^t\|$$

and rearranging yields (4.10). ■

For the following lemma, we recall that $\mathcal{A}_t = \sigma(\{\mathsf{z}^\tau; \tau \leq t\})$ encapsulates the information generated by the stochastic process up to iteration t.

Lemma 5 *Let $\hat{z} \in \mathcal{Z}_{\mathrm{ad}}$ be such that $\Phi(\hat{z}) = 0$. Then, for all $t \in \mathbb{T}$,*

$$\mathbb{E}[\mathsf{w}^t G_t(\hat{z})|\mathcal{A}_t] = -\varepsilon\mathbb{E}(\mathsf{w}^t|\mathcal{A}_t). \tag{4.11}$$

Proof Via the Tower property of conditional expectations and (4.1), we immediately deduce that

$$\begin{aligned} \mathbb{E}[\mathsf{w}^t G_t(\hat{z})|\mathcal{A}_t] &= \mathbb{E}[\mathbb{E}(\mathsf{w}^t G_t(\hat{z})|\mathcal{A}_{t-1})|\mathcal{A}_t] \\ &= \mathbb{E}[\mathsf{w}^t \mathbb{E}(G_t(\hat{z})|\mathcal{A}_{t-1})|\mathcal{A}_t] \\ &= \mathbb{E}[\mathsf{w}^t(\Phi(\hat{z}) - \varepsilon)|\mathcal{A}_t] \\ &= -\varepsilon\mathbb{E}[\mathsf{w}^t|\mathcal{A}_t]. \end{aligned}$$

■

Lemma 6 *For all $t \in \mathbb{T}$, we have*

$$\mathsf{w}^t - D_2 R \leq \mathsf{w}^{t+1} \leq \mathsf{w}^t + M \tag{4.12}$$

Proof We start with establishing the upper bound. From convexity, we get

$$G_t(\mathsf{z}^t) + (G'_t(\mathsf{z}^t), z - \mathsf{z}^t) \leq G_t(z). \tag{4.13}$$

Hence, by definition of the updating we have

$$(\mathsf{w}^{t+1})^2 \leq \left(\mathsf{w}^t + G_t(\mathsf{z}^t) + (G'_t(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)\right)^2 \leq \left(\mathsf{w}^t + G_t(\mathsf{z}^{t+1})\right)^2.$$

Hence, by the triangle inequality, and part 3 of Definition 2, we conclude

$$\begin{aligned} |\mathsf{w}^{t+1}| &\leq |\mathsf{w}^t + G_t(\mathsf{z}^{t+1})| \leq |\mathsf{w}^t| + |G_t(\mathsf{z}^{t+1})| \\ &\leq \mathsf{w}^t + M. \end{aligned}$$

For the lower bound, we use part 2 of Definition 2 to observe that

$$\begin{aligned} |\mathsf{w}^{t+1} - \mathsf{w}^t| &\leq |G_t(\mathsf{z}^t) + (G'_t(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)| \\ &\leq |G_t(\mathsf{z}^t)| + \|G'_t(\mathsf{z}^t)\| \cdot \|\mathsf{z}^{t+1} - \mathsf{z}^t\| \\ &\overset{(4.2)}{\leq} M + D_2 R. \end{aligned}$$

■

We remark that Lemma 6 implies that for all $t_1, t_2 \in \mathbb{T}$,

$$\mathsf{w}^{t_1+t_2} \geq \mathsf{w}^{t_1} - t_2(M + D_2 R), \tag{4.14}$$

Set $L_t \triangleq \frac{1}{2}[(\mathsf{w}^{t+1})^2 - (\mathsf{w}^t)^2]$.

Lemma 7 *For all $t \in \mathbb{T}$,*

$$L_t \leq \mathsf{w}^t[G_t(\mathsf{z}^t) + (G'_t(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)] + \frac{1}{2}(M + D_2 R)^2. \tag{4.15}$$

Proof This is (4.6). ■

The next result is fundamental to our approach. It gives a drift lemma for the penalty process $\{\mathsf{w}^t\}_{t\in\mathbb{T}}$, in the spirit of [37]. This result will be instrumental to prove an $L^1(\Omega, \mathcal{A}, \mathbb{P})$ bound on the penalty process. The next Lemma shows that the penalty process $\{\mathsf{w}^t\}_t$ satisfies the conditions in [37, L. 5].

Lemma 8 *Let $\mathbb{T} = \mathbb{N}$ and consider the master algorithm $\mathbb{X}^{(z,w)}_{\mathbb{T}}(\alpha, \gamma)$. Let $\varepsilon < 2(M + D_2 R)$ and n an arbitrary integer. Then, for each $t \in \mathbb{N}$, we have*

$$|\mathsf{w}^{t+1} - \mathsf{w}^t| \leq M + D_2 R \quad \text{and} \tag{4.16}$$

$$\mathbb{E}[\mathsf{w}^{t+n} - \mathsf{w}^t | \mathcal{A}_{t-1}] \leq \begin{cases} -\frac{\varepsilon}{2} n & \text{if } \mathsf{w}^t \geq \delta(\varepsilon, n) \\ n(M + D_2 R) & \text{if } \mathsf{w}^t < \delta(\varepsilon, n), \end{cases} \tag{4.17}$$

where

$$\delta(\varepsilon, n) \triangleq \frac{\varepsilon n}{2} + \frac{\alpha}{n\varepsilon} R^2 + \frac{2}{\varepsilon}\left[\gamma D_1 R + \frac{1}{2}(M + D_2 R)^2\right] + n(M + D_2 R). \tag{4.18}$$

Proof Condition (4.16) is just a restatement of Lemma 6. To verify (4.17), recall that $L_t = \frac{1}{2}(\mathsf{w}^{t+1})^2 - \frac{1}{2}(\mathsf{w}^t)^2$. Condition (4.8) is equivalent to

$$\begin{aligned}
&\gamma(J_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t) + \mathsf{w}^t(G_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t) + \frac{\alpha}{2}\|\mathsf{z}^{t+1} - \mathsf{z}^t\|^2 \\
&\quad \leq \gamma(J_t'(\mathsf{z}^t), z - \mathsf{z}^t) + \mathsf{w}^t(G_t'(\mathsf{z}^t), z - \mathsf{z}^t) + \frac{\alpha}{2}\|\mathsf{z}^t - z\|^2 - \frac{\alpha}{2}\|\mathsf{z}^{t+1} - z\|^2
\end{aligned}$$

for all $z \in \mathcal{Z}_{\text{ad}}$. Adding $\mathsf{w}^t G_t(\mathsf{z}^t)$ to both sides and using (4.13), it follows that

$$\begin{aligned}
&\mathsf{w}^t\left[G_t(\mathsf{z}^t) + (G_t'(\mathsf{z}^t), \mathsf{z}^{t+1} - \mathsf{z}^t)\right] \\
&\quad \leq \gamma(J_t'(\mathsf{z}^t), z - \mathsf{z}^{t+1}) + \mathsf{w}^t G_t(z) \\
&\quad + \frac{\alpha}{2}\left(\|z - \mathsf{z}^t\|^2 - \|z - \mathsf{z}^{t+1}\|^2 - \|\mathsf{z}^{t+1} - \mathsf{z}^t\|^2\right) \\
&\quad \leq \gamma\|J_t'(\mathsf{z}^t)\| \cdot \|z - \mathsf{z}^{t+1}\| + \mathsf{w}^t G_t(z) + \frac{\alpha}{2}\left(\|z - \mathsf{z}^t\|^2 - \|z - \mathsf{z}^{t+1}\|^2\right) \\
&\quad \overset{(4.2)}{\leq} \gamma D_1 R + \mathsf{w}^t G_t(z) + \frac{\alpha}{2}\left(\|z - \mathsf{z}^t\|^2 - \|z - \mathsf{z}^{t+1}\|^2\right).
\end{aligned} \tag{4.19}$$

By (4.15) and (4.19), we readily obtain

$$L_t \leq \gamma D_1 R + \mathsf{w}^t G_t(\hat{z}) + \frac{\alpha}{2}\left(\|\hat{z} - \mathsf{z}^t\|^2 - \|\hat{z} - \mathsf{z}^{t+1}\|^2\right) + \frac{1}{2}(M + D_2 R)^2.$$

for all $t \geq 0$ and $\hat{z} \in \mathcal{Z}_{\text{ad}}$. Therefore, for all $t \geq 1$ and $n \geq 1$ we get

$$\begin{aligned}
\frac{1}{2}(\mathsf{w}^{t+n})^2 - \frac{1}{2}(\mathsf{w}^t)^2 &= \sum_{\tau=t}^{t+n-1} L_\tau \\
&\leq n\left[\gamma D_1 R + \frac{1}{2}(M + D_2 R)^2\right] + \frac{\alpha}{2}R^2 + \sum_{\tau=t}^{t+n-1} \mathsf{w}^\tau G_\tau(\hat{z}).
\end{aligned}$$

Whence,

$$(\mathsf{w}^{t+n})^2 \leq (\mathsf{w}^t)^2 + 2n\left[\gamma D_1 R + \frac{1}{2}(M + D_2 R)^2\right] + \alpha R^2 + 2\sum_{\tau=t}^{t+n-1} \mathsf{w}^\tau G_\tau(\hat{z}).$$

Let us pick a point $\hat{z} \in \mathcal{Z}_{\text{ad}}$ for which $\mathbb{E}[G_t(\hat{z})|\mathcal{A}_t] = -\varepsilon$, i.e. a feasible control satisfying $\Phi(\hat{z}) = 0$. Lemma 5 and the law of iterated expectations shows that for $\tau \geq t$, we have

$$\begin{aligned}
\mathbb{E}\left(\mathsf{w}^\tau G_\tau(\hat{z})|\mathcal{A}_{t-1}\right) &= \mathbb{E}\left[\mathbb{E}(\mathsf{w}^t G_\tau(\hat{z})|\mathcal{A}_\tau)|\mathcal{A}_{t-1}\right] = -\varepsilon\mathbb{E}\left[\mathbb{E}(\mathsf{w}^\tau|\mathcal{A}_\tau)|\mathcal{A}_{t-1}\right] \\
&= -\varepsilon\mathbb{E}[\mathsf{w}^\tau|\mathcal{A}_{t-1}].
\end{aligned}$$

Then, taking $\mathcal{A}_{t-1}$-conditional expectations on both sides gives

$$
\begin{aligned}
&\mathbb{E}[(\mathsf{w}^{t+n})^2|\mathcal{A}_{t-1}] \\
&\quad\le (\mathsf{w}^t)^2 + 2n\left[\gamma D_1 R + \tfrac{1}{2}(M+D_2R)^2\right] + \alpha R^2 - 2\varepsilon\sum_{\tau=t}^{t+n-1}\mathbb{E}[\mathsf{w}^\tau|\mathcal{A}_{t-1}] \\
&\quad\le (\mathsf{w}^t)^2 + 2n\left[\gamma D_1 R + \tfrac{1}{2}(M+D_2R)^2\right] + \alpha R^2 - 2\varepsilon\sum_{j=1}^{n-1}\mathbb{E}[\mathsf{w}^{t+j}|\mathcal{A}_{t-1}] \\
&\quad\overset{(4.14)}{\le} (\mathsf{w}^t)^2 + 2n\left[\gamma D_1 R + \tfrac{1}{2}(M+D_2R)^2\right] + \alpha R^2 - 2\varepsilon\sum_{j=0}^{n-1}\left(\mathsf{w}^t - j(M+D_2R)\right) \\
&\quad\le (\mathsf{w}^t)^2 + 2n\left[\gamma D_1 R + \tfrac{1}{2}(M+D_2R)^2\right] + \alpha R^2 - 2\varepsilon n\mathsf{w}^t + \varepsilon n^2(M+D_2R) \\
&\quad= (\mathsf{w}^t)^2 - \varepsilon n\mathsf{w}^t + n\left[2\left(\gamma D_1 R + \tfrac{1}{2}(M+D_2R)^2\right) + \varepsilon n(M+D_2R) - \varepsilon\mathsf{w}^t + \tfrac{\alpha}{n}R^2\right],
\end{aligned}
$$

where we have used $\sum_{j=1}^{n-1} j = \frac{n(n-1)}{2} \le \frac{n^2}{2}$ in the fourth inequality. From this we deduce that if $\mathsf{w}^t \ge \delta(\varepsilon, n)$, then

$$\mathbb{E}[(\mathsf{w}^{t+n})^2|\mathcal{A}_{t-1}] \le (\mathsf{w}^t)^2 - \varepsilon n\mathsf{w}^t - \frac{\varepsilon^2 n^2}{2} \le (\mathsf{w}^t - \frac{\varepsilon}{2}n)^2.$$

This, finally, leads to the bound

$$\mathbb{E}[\mathsf{w}^{t+n}|\mathcal{A}_{t-1}] \le \sqrt{\mathbb{E}[(\mathsf{w}^{t+n})^2|\mathcal{A}_{t-1}]} \le \mathsf{w}^t - \frac{\varepsilon}{2}n$$

provided that $\mathsf{w}^t \ge \delta(\varepsilon, n)$. Conversely, if $\mathsf{w}^t < \delta(\varepsilon, n)$, then we can use (4.16) to obtain $\mathsf{w}^{t+n} - \mathsf{w}^t \le n(M + D_2R)$. ■

Corollary 1 *Under the same assumptions as in Lemma 8, we have for all* $t \ge 0$*,*

$$\mathbb{E}[\mathsf{w}^t] \le \delta(\varepsilon, n) + \frac{8n(M+D_2R)^2}{\varepsilon}\log\left(1 + \frac{32(M+D_2R)^2}{\varepsilon^2}e^{\frac{\varepsilon}{8(M+D_2R)}}\right).$$

Proof This follows from Lemma 8 and Part 1) of [37, L. 5]. ■

We remark that the constant

$$\mathsf{C} \triangleq \frac{8(M+D_2R)^2}{\varepsilon}\log\left(1 + \frac{32(M+D_2R)^2}{\varepsilon^2}e^{\frac{\varepsilon}{8(M+D_2R)}}\right) \tag{4.20}$$

is in fact an absolute constant, independent of algorithm parameters. We further remark, that if we choose $n = \lceil\sqrt{\Delta}\rceil, \alpha = \lceil\Delta\rceil$ and $\gamma = \sqrt{\Delta}$ for some $\Delta > 0$, then

$$\mathbb{E}[\mathsf{w}^t] \le \delta(\varepsilon, \lceil\sqrt{\Delta}\rceil) + \lceil\sqrt{\Delta}\rceil\mathsf{C} = \mathcal{O}(\sqrt{\Delta}). \tag{4.21}$$

Corollary 2 *Suppose* $\sup \mathbb{T} < \infty$. *Then, for any* $\min \mathbb{T} \leq t_1 < t_2 \leq \sup \mathbb{T}$, *we have*

$$\sum_{t=t_1}^{t_2} G_t(\mathsf{z}^t) \leq \mathsf{w}^{t_2+1} + D_1 D_2 (t_2 - t_1 + 1)\frac{\gamma}{\alpha} + \frac{D_2^2}{\alpha}\sum_{t=t_1}^{t_2} \mathsf{w}^t. \tag{4.22}$$

In particular, if $\min \mathbb{T} = t_0 + 1, \sup \mathbb{T} = t_0 + \Delta$ *for some* $\Delta \in \mathbb{N}$, *and* $\alpha = \Delta, \gamma = \sqrt{\Delta}$, *then*

$$\mathbb{E}\left[\sum_{t\in\mathbb{T}} G_t(\mathsf{z}^t)\right] \leq \mathcal{O}(\sqrt{\Delta}) \tag{4.23}$$

Proof Combining Lemma 3 with (4.10) yields immediately (4.22). Now, using Corollary 1, together with (4.21), it follows that

$$\mathbb{E}\left[\sum_{t=t_0+1}^{t_0+\Delta} G_t(\mathsf{z}^t)\right] \leq \mathcal{O}(\sqrt{\Delta}) + D_1 D_2 \sqrt{\Delta} + \frac{D_2^2}{\Delta}\mathcal{O}(\Delta\sqrt{\Delta}) = \mathcal{O}(\sqrt{\Delta}).$$

■

4.4 Main Convergence Argument

In this section, we give detailed proofs on the convergence properties of Algorithm 2. Recall, that we construct the sequences $\{(z_k, w_k)\}_{k=1}^N$ by concatenating the trajectories produced by the master algorithm on the batches $\mathbb{T}_1, \ldots, \mathbb{T}_s$, as described in (4.5). We let $\mathcal{F}_0 \triangleq \{\varnothing, \Omega\}$ and $\mathcal{F}_k \triangleq \sigma(z_1, \ldots, z_k)$, denote the natural filtration induced by the so-constructed process. To emphasize that the batches are computed using i.i.d. information, we let $G_{j,t}(z)$ and $J_{j,t}(z)$ represent the random estimators reported by the SO in epoch $j \in \{1, \ldots, s\}$ and inner iteration $t \in \mathbb{T}_j$ at position $z \in \mathcal{Z}_{\text{ad}}$, and let $\{(\mathsf{z}_j^t, \mathsf{w}_j^t), 1 \leq t \leq \sup \mathbb{T}_j + 1\}$ denote the subsequence computed by the master algorithm $\mathbb{X}_{\mathbb{T}_j}^{(\mathsf{z}_j^1, \mathsf{w}_j^1)}(\alpha_j, \gamma_j)$ in the j-th restart. The filtration used to measure the concatenated process $\{(z_k, w_k)\}_{k=1}^N$ is intrinsically related to the filtration induced by the master process $\{\mathcal{A}_{j,t}\}_{t=0}^{\sup \mathbb{T}_j}$. Specifically, if $k = (j-1)\Delta_N + t$ for some $t \in \mathbb{T}_j$, we have $\mathcal{F}_k = \sigma\left(\bigcup_{i=1}^{j-1}\{\mathcal{A}_{i,t}\}_{t=0}^{\sup \mathbb{T}_i} \cup \mathcal{A}_{j,1} \cup \ldots \cup \mathcal{A}_{j,t}\right)$.

Our first result is a bound on the expected constraint violation in terms of the ergodic average.

Proposition 4 *Consider Algorithm 2 with epochs* $j \in \{1, \ldots, s\}$ *and epoch-specific step sizes* $\gamma_j = \sqrt{\Delta_N}$ *and* $\alpha_j = \Delta_N$, *where* $\Delta_N \triangleq \lfloor N^a \rfloor$ *for some* $a \in (0, 1]$. *Then,*

$$\mathbb{E}[\Phi(\bar{z}_N)] \leq \mathcal{O}(N^{-a/2})(1 + \mathcal{O}(N^{a-1})) + \varepsilon, \tag{4.24}$$

where $\varepsilon > 0$ *is the a-priori fixed relaxation parameter.*

Proof Corollary 2 gives for each $j = 1, 2, \ldots, s-1$,

$$\begin{aligned}\sum_{k=(j-1)\Delta_N+1}^{j\Delta_N} G_k(z_k) &= \sum_{t=1}^{\Delta_N} G_{j,t}(\mathsf{z}_j^t)\\ &\le \mathsf{w}_j^{\Delta_N+1} - \mathsf{w}_j^1 + \frac{\gamma_j}{\alpha_j}\Delta_N D_1 D_2 + \frac{D_2^2}{\alpha_j}\sum_{t=1}^{\Delta_N}\mathsf{w}_j^t\\ &= \mathsf{w}_{j+1}^1 - \mathsf{w}_j^1 + \frac{\gamma_j}{\alpha_j}\Delta_N D_1 D_2 + \frac{D_2^2}{\alpha_j}\sum_{t=1}^{\Delta_N}\mathsf{w}_j^t,\end{aligned}$$

where the last equality uses the definition $\mathsf{w}_{j+1}^1 = \mathsf{w}_j^{\Delta_N+1}$. Furthermore, for $j = s$, Corollary 2 yields

$$\begin{aligned}\sum_{k=(s-1)\Delta_N+1}^{N} G_k(z_k) &= \sum_{t=1}^{N-(s-1)\Delta_N} G_{s,t}(\mathsf{z}_s^t)\\ &\le \mathsf{w}_s^{N-(s-1)\Delta_N+1} - \mathsf{w}_s^1 + \frac{\gamma_s}{\alpha_s}\Delta_N D_1 D_2 + \frac{D_2^2}{\alpha_j}\sum_{t=1}^{N-(s-1)\Delta_N}\mathsf{w}_j^t,\end{aligned}$$

Therefore,

$$\begin{aligned}\sum_{k=1}^{N} G_k(z_k) &= \sum_{j=1}^{s-1}\sum_{k=(j-1)\Delta_N+1}^{j\Delta_N} G_k(z_k) + \sum_{k=(s-1)\Delta_N+1}^{N} G_k(z_k)\\ &= \sum_{j=1}^{s-1}\sum_{t=1}^{\Delta_N} G_{j,t}(\mathsf{z}_j^t) + \sum_{t=1}^{N-(s-1)\Delta_N} G_{s,t}(\mathsf{z}_s^t)\\ &\le \Delta_N D_1 D_2 \sum_{j=1}^{s}\frac{\gamma_j}{\alpha_j} + \sum_{j=1}^{s}\frac{D_2^2}{\alpha_j}\sum_{t\in\mathbb{T}_j}\mathsf{w}_j^t\\ &+ (\mathsf{w}_2^1 - \mathsf{w}_1^1) + (\mathsf{w}_3^1 - \mathsf{w}_2^1) + \ldots + (\mathsf{w}_s^1 - \mathsf{w}_{s-1}^1) + (\mathsf{w}_s^{N-(s-1)\Delta_N+1} - \mathsf{w}_s^1)\\ &\le \mathsf{w}_s^{N-(s-1)\Delta_N+1} + \Delta_N D_1 D_2 \sum_{j=1}^{s}\frac{\gamma_j}{\alpha_j} + \sum_{j=1}^{s}\frac{D_2^2}{\alpha_j}\sum_{t\in\mathbb{T}_j}\mathsf{w}_j^t.\end{aligned}$$

Corollary 1 implies

$$\begin{aligned}&\mathbb{E}\left[\mathsf{w}_s^{N+1-(s-1)\Delta_N}\right]\\ &\le \delta(\varepsilon, n) + \frac{8n(M+D_2R)^2}{\varepsilon}\log\left(1 + \frac{32(M+D_2R)^2}{\varepsilon^2}e^{\frac{\varepsilon}{8(M+D_2R)}}\right)\end{aligned}$$

and

$$\mathbb{E}\left[\mathsf{w}_j^t\right] \le \delta(\varepsilon, n) + \frac{8n(M+D_2R)^2}{\varepsilon}\log\left(1 + \frac{32(M+D_2R)^2}{\varepsilon^2}e^{\frac{\varepsilon}{8(M+D_2R)}}\right),$$

where n and ε are to be chosen. In particular, setting $n = \lceil \sqrt{\Delta_N} \rceil$, it follows from (4.21) that $\mathbb{E}[\mathsf{w}_s^{N+1-(s-1)\Delta_N}] \leq \mathcal{O}(\sqrt{\Delta_N})$, and $\mathbb{E}[\mathsf{w}_j^t] \leq \mathcal{O}(\sqrt{\Delta_N})$. Thus, taking expectations on both sides of the penultimate display, and using the specification $\gamma_j = \sqrt{\Delta_N}, \alpha_j = \Delta_N$ for all $j \in \{1, 2, \dots, s\}$, as well as $s = \lceil N/\Delta_N \rceil$, we arrive at following full sequence counterpart to (4.23):

$$\mathbb{E}\left[\sum_{k=1}^{N} G_k(z_k)\right] \leq \mathcal{O}(\sqrt{\Delta_N}) + D_1 D_2 s\sqrt{\Delta_N} + D_2^2 \mathcal{O}(s\sqrt{\Delta_N})$$
$$= \mathcal{O}(s\sqrt{\Delta_N})\,(1 + \mathcal{O}(\Delta_N/N))\,.$$

By choosing $\Delta_N = \lfloor N^a \rfloor$ for $a \in (0, 1]$, it follows $\mathcal{O}(s\sqrt{\Delta_N}) = \mathcal{O}(N^{1-a/2})$. Hence,

$$\mathbb{E}\left[\sum_{k=1}^{N} G_k(z_k)\right] \leq \mathcal{O}(N^{1-a/2})(1 + \mathcal{O}(N^{a-1}))$$

The law of iterated expectations and Jensen's inequality implies

$$\frac{1}{N}\mathbb{E}\left[\sum_{k=1}^{N} G_k(z_k)\right] = \frac{1}{N}\mathbb{E}\left[\sum_{k=1}^{N} (\Phi(z_k) - \varepsilon)\right] \geq \mathbb{E}[\Phi(\bar{z}_N)] - \varepsilon,$$

where we have used the convexity of the penalty function Φ, established in Proposition 1. Whence, we arrive at the expected constrained violation bound

$$\mathbb{E}[\Phi(\bar{z}_N)] \leq \mathcal{O}(N^{-a/2})(1 + \mathcal{O}(N^{a-1})) + \varepsilon. \tag{4.25}$$

■

Remark 6 If we choose $a = 1$, then no restart effectively takes place and we recover the standard $\mathcal{O}(N^{-1/2})$ constraint violation bound from [37]. Restart leads to slightly worse constraint violation bounds, but comes with the decisive advantage that it allows us to choose larger, epoch-dependent, step sizes. This has potentially significant impacts on the practical performance of the algorithm, as we will demonstrate in Section 5.

We conclude the analysis of Algorithm 2 by proving a convergence rate in terms of objective function values. Following the notation of Proposition 3, we fix $\varepsilon > 0$ and denote any corresponding solution to (P_ε) by $z_\varepsilon^\star$.

Theorem 1 *Consider Algorithm 2 with epochs $j \in \{1, \dots, s\}$ and epoch-specific step sizes $\gamma_j = \sqrt{\Delta_N}$ and $\alpha_j = \Delta_N$. If $\Delta_N = \lfloor N^a \rfloor$ for $a \in (0, 1]$, then*

$$\mathbb{E}[j(\bar{z}_N) - j(z_\varepsilon^\star)] \leq \mathcal{O}(N^{-a/2}). \tag{4.26}$$

Proof Let $j \in \{1,\ldots,s\}$ and $t \in \mathbb{T}_j$ be arbitrary. Choosing $z = z_\varepsilon^\star$ in (4.9), it follows that

$$\begin{aligned}\gamma_j\left[J_{j,t}(\mathsf{z}_j^t) - J_{j,t}(z_\varepsilon^\star)\right] \leq\ & \mathsf{w}_j^t G_{j,t}(z_\varepsilon^\star) - \gamma_j(J'_{j,t}(\mathsf{z}_j^t), \mathsf{z}_j^{t+1} - \mathsf{z}_j^t) \\ &+ \frac{\alpha_j}{2}\left(\|\mathsf{z}_j^t - z_\varepsilon^\star\|^2 - \|\mathsf{z}_j^{t+1} - z_\varepsilon^\star\|^2 - \|\mathsf{z}_j^{t+1} - \mathsf{z}_j^t\|^2\right) \\ &- \mathsf{w}_j^t\left[G_{j,t}(\mathsf{z}_j^t) + (G'_{j,t}(\mathsf{z}_j^t), \mathsf{z}_j^{t+1} - \mathsf{z}_j^t)\right]\end{aligned}$$

Using (4.15) with $L_{j,t} \triangleq \frac{1}{2}(\mathsf{w}_j^{t+1})^2 - \frac{1}{2}(\mathsf{w}_j^t)^2$, we arrive at

$$\begin{aligned}\gamma_j\left[J_{j,t}(\mathsf{z}_j^t) - J_{j,t}(z_\varepsilon^\star)\right] \leq\ & \mathsf{w}_j^t G_{j,t}(z_\varepsilon^\star) - \gamma_j(J'_{j,t}(\mathsf{z}_j^t), \mathsf{z}_j^{t+1} - \mathsf{z}_j^t) - L_{j,t} \\ &+ \frac{\alpha_j}{2}\left(\|\mathsf{z}_j^t - z_\varepsilon^\star\|^2 - \|\mathsf{z}_j^{t+1} - z_\varepsilon^\star\|^2 - \|\mathsf{z}_j^{t+1} - \mathsf{z}_j^t\|^2\right) \\ &+ \frac{1}{2}(M + D_2R)^2.\end{aligned}$$

Next, by Cauchy-Schwarz and the definition of the SO, we observe

$$\begin{aligned}-(J'_{j,t}(\mathsf{z}_j^t), \mathsf{z}_j^{t+1} - \mathsf{z}_j^t) - \frac{\alpha_j}{2\gamma_j}\|\mathsf{z}_j^{t+1} - \mathsf{z}_j^t\|^2 &\leq D_1\|\mathsf{z}_j^{t+1} - \mathsf{z}_j^t\| - \frac{\alpha_j}{2\gamma_j}\|\mathsf{z}_j^{t+1} - \mathsf{z}_j^t\|^2 \\ &= \frac{\gamma_j D_1^2}{2\alpha_j} - \frac{\alpha_j}{2\gamma_j}\left(\|\mathsf{z}_j^{t+1} - \mathsf{z}_j^t\| - \frac{D_1\gamma_j}{\alpha_j}\right)^2 \\ &\leq \frac{\gamma_j D_1^2}{2\alpha_j}.\end{aligned}$$

When combined with the previous display, this yields

$$\begin{aligned}\gamma_j\left[J_{j,t}(\mathsf{z}_j^t) - J_{j,t}(z_\varepsilon^\star)\right] \leq\ & \mathsf{w}_j^t G_{j,t}(z_\varepsilon^\star) - L_{j,t} + \frac{\gamma_j^2 D_1^2}{2\alpha_j} \\ &+ \frac{\alpha_j}{2}\left(\|\mathsf{z}_j^t - z_\varepsilon^\star\|^2 - \|\mathsf{z}_j^{t+1} - z_\varepsilon^\star\|^2\right) + \frac{1}{2}(M + D_2R)^2.\end{aligned}$$

Aggregating these estimates for $t \in \mathbb{T}_j$ gives

$$\begin{aligned}\sum_{t\in\mathbb{T}_j}\gamma_j\left[J_{j,t}(\mathsf{z}_j^t) - J_{j,t}(z_\varepsilon^\star)\right] \leq\ & \sum_{t\in\mathbb{T}_j}\mathsf{w}_j^t G_{j,t}(z_\varepsilon^\star) + \frac{1}{2}(\mathsf{w}_j^1)^2 - \frac{1}{2}(\mathsf{w}_j^{\sup\mathbb{T}_j+1})^2 + \Delta_N\frac{\gamma_j^2 D_1^2}{2\alpha_j} \\ &+ \frac{\alpha_j}{2}\left(\|\mathsf{z}_j^1 - z_\varepsilon^\star\|^2 - \|\mathsf{z}_j^{\sup\mathbb{T}_j+1} - z_\varepsilon^\star\|^2\right) + \frac{\Delta_N}{2}(M + D_2R)^2.\end{aligned}$$

Next, we sum over all epochs $j \in \{1,\ldots,s\}$ and set $\gamma_j = \sqrt{\Delta_N}, \alpha_j = \Delta_N$. Using the warm start condition $\mathsf{w}_j^{\sup\mathbb{T}_j+1} = \mathsf{w}_{j+1}^1$ for $j \in \{1,\ldots,s-1\}$, we obtain

$$\begin{aligned}\sum_{k=1}^N[J_k(z_k) - J_k(z_\varepsilon^\star)] &= \sum_{j=1}^s\sum_{t\in\mathbb{T}_j}\left[J_{j,t}(\mathsf{z}_j^t) - J_{j,t}(z_\varepsilon^\star)\right] \\ &\leq \frac{1}{\sqrt{\Delta_N}}\sum_{k=1}^N w_k G_k(z_\varepsilon^\star) + \frac{1}{2\sqrt{\Delta_N}}(\mathsf{w}_1^1)^2 - \frac{1}{2\sqrt{\Delta_N}}(\mathsf{w}_s^{\sup\mathbb{T}_s+1})^2 \\ &+ \frac{s\sqrt{\Delta_N}}{2}(M + D_2R)^2 + \frac{D_1^2 s\sqrt{\Delta_N}}{2} + \frac{R^2\sqrt{\Delta_N}}{2}.\end{aligned}$$

By definition, $\Phi(z_\varepsilon^\star) \le \varepsilon$, so that $\mathbb{E}[G_k(z_\varepsilon^*)] \le 0$. Via the law of iterated expectations, this implies

$$\begin{aligned}\mathbb{E}[w_k G_k(z_\varepsilon^*)] &= \mathbb{E}\left[\mathbb{E}[w_k G_k(z_\varepsilon^\star)|\mathcal{F}_k]\right] = \mathbb{E}\left[w_k \mathbb{E}[G_k(z_\epsilon^\star)|\mathcal{F}_k]\right] \\ &= \mathbb{E}\left[w_k(\Phi(z_\varepsilon^\star) - \varepsilon)\right] \le 0.\end{aligned}$$

Using this, and the fact that $\mathsf{w}_1^1 = w_1 = 0$, we can take expectations on both sides of the penultimate display, and finally arrive at

$$\mathbb{E}\left[\sum_{k=1}^N (J_k(z_k) - J_k(z_\varepsilon^\star))\right] \le \frac{s\sqrt{\Delta_N}}{2}(M + D_2 R)^2 + \frac{D_1^2 s\sqrt{\Delta_N}}{2} + \frac{R^2\sqrt{\Delta_N}}{2}. \tag{4.27}$$

This readily implies

$$\mathbb{E}[j(\bar{z}_N) - j(z_\varepsilon^\star)] \le \frac{s\sqrt{\Delta_N}}{2N}(M + D_2 R)^2 + \frac{D_1^2 s\sqrt{\Delta_N}}{2N} + \frac{R^2\sqrt{\Delta_N}}{2N}.$$

Since $s = \lceil \frac{N}{\Delta_N} \rceil$, it follows $\frac{s\sqrt{\Delta_N}}{N} = \mathcal{O}(1/\sqrt{\Delta_N})$. Hence, for $\Delta_N = \lfloor N^a \rfloor$ for $a \in (0, 1]$, we get the bound

$$\mathbb{E}[j(\bar{z}_N) - j(z_\varepsilon^\star)] \le \frac{s\sqrt{\Delta_N}}{2N}\left[(M + D_2 R)^2 + D_1^2 + \frac{R^2}{s}\right] \le \mathcal{O}(N^{-a/2}).$$

■

Remark 7 Stochastic approximation algorithms often allow mean convergence statements for the (ergodic) trajectory in the presence of strong convexity. However, due to the fact that our problem formulation includes a stochastic approximation of the functional bound constraint, it is rather difficult, if not impossible, to extend the standard arguments. If it could be guaranteed that $\Phi(z_\varepsilon^\star) < \varepsilon$, then a convergence rate for the trajectory can be readily derived using the fact that j is smooth and strongly convex.

5 Implementation and Numerical Experiments

In this section we test the OSA Algorithm 2 on a strongly convex problem arising from the optimal control of a linear elliptic PDE with uncertain coefficients. Our test problem is motivated by the example in [21] and has the form

$$\min_{z \in L^2(D)} \frac{\alpha}{2}\mathbb{E}\left[\int_D ([u_{\boldsymbol{\xi}}(z)](x) - w(x))^2 \,\mathrm{d}x\right] + \frac{1}{2}\int_D z(x)^2 \,\mathrm{d}x \tag{5.1a}$$

$$\text{subject to} \quad -10 \le z \le 10 \quad \text{a.e.}, \qquad u_{\boldsymbol{\xi}}(z) \ge \psi \quad \text{a.e./a.s.}, \tag{5.1b}$$

where $u = u_\xi(z) \in H^1(D)$ is a weak solution to

$$-\nabla \cdot (\kappa(x,\xi)\nabla u(x)) + v(x,\xi) \cdot \nabla u(x) = f(x,\xi) + z(x) \quad \text{for } x \in D \tag{5.1c}$$

$$\kappa(x,\xi)\nabla u(x) \cdot n = 0 \quad \text{for } x \in \Gamma_n \tag{5.1d}$$

$$u(x) = 0 \quad \text{for } x \in \Gamma_d \tag{5.1e}$$

Here, $D \triangleq (0,1)^2$, $\alpha = 10^4$, $w(x) \triangleq (x - 0.5)^\top (x - 0.5)$,

$$\psi(x) \triangleq \begin{cases} \frac{1}{4} & \text{if } \|x - (\frac{1}{2}, \frac{1}{2})^\top\|_2 \le \frac{1}{4} \\ 0 & \text{otherwise} \end{cases},$$

$\Gamma_d \triangleq \{0\} \times (0,1)$, $\Gamma_n \triangleq \partial D \setminus \Gamma_d$,

$$\kappa(x,\xi) \triangleq 0.5 + c\exp(\beta(x,\xi)), \quad v(x,\xi) = \begin{pmatrix} b(\xi) - a(\xi)x_1 \\ a(\xi)x_2 \end{pmatrix},$$

and f is the sum of five Gaussian sources whose locations, widths and magnitudes are random. The explicit form for β is described in [21, §4] (where it is denoted by δ). The random inputs $\boldsymbol{\xi}$ are uniformly distributed on $[-1,1]^{37}$. For our results, we replace the state constraint in (5.1a) with the smooth penalty constraint as in (P_ε) with φ as in Example 1. We chose the state constraint penalty parameters to be $\varepsilon = 10^{-2}$ and $\delta = 10^{-2}$.

Since an exact solution is unknown, we solved the problem using sample-average approximation (SAA) for the expectations. We employ the augmented Lagrangian (AL) algorithm described in [7] to solve the resulting deterministic problem. At each AL iteration, we solve the bound-constrained subproblem using the projected trust-region Newton method described in [24]. Our implementation of this algorithm is available in the Rapid Optimization Library [31]. Since the AL algorithm uses second-order information for the subproblem solves, it is reasonable to assume that it makes more progress each iteration than the OSA algorithm, with more computational effort. We treat the resulting solution as the “true” solution and empirically study the performance of the OSA algorithm. In contrast to the OSA algorithm, for which we have a full convergence proof in Hilbert space, the SAA approach would require, amongst other things, a (statistical) consistency result to be fully justified. For this we would need to investigate the asymptotic behavior of the random set-valued mappings

$$\Upsilon_\varepsilon(\mathbb{P}_N) := \{ z \in \mathcal{Z}_{\text{ad}} \,|\, \mathbb{E}_{\mathbb{P}_N}[F(\theta(z,\boldsymbol{\xi}))] \le \varepsilon \},$$

which goes beyond the scope of this paper.

We discretized (5.1c) using continuous piecewise linear finite elements on a uniform 64×64 quadrilateral mesh. To obtain our reference solution, we applied the aforementioned SAA approach with 10^3 samples. We stopped AL when the optimality and feasibility criteria were smaller than 10^{-8}, which required 7 iterations (12 subproblem iterations). The final values for the optimality and feasibility criteria were 5.33×10^{-9} and 2.84×10^{-11}, respectively. For reference, AL required 20 function and gradient evaluations as well as 145 applications of the Hessian to a vector, resulting in a total of 495,000 deterministic PDE solves. Since Algorithm 2 requires 3 deterministic PDE solves per iteration, the cost of AL is comparable to running Algorithm 2 with $N = 165,000$. We also solved (5.1a) using Algorithm 2 with $N \in \{10^3, 10^4, 10^5\}$ iterations and epoch lengths $\varDelta_N = N$ and $\varDelta_N = 500$. We ran each of these cases five times. This provided us with five controls per configuration of N and $\varDelta_N$ for comparison in the statistical tests detailed below.

Figure 2 depicts the empirical distribution of the objective function for each of these runs using 10^4 samples, which were chosen to be different than the samples used for SAA and Algorithm 2. The five runs with 10^3 (red), 10^4 (blue) and 10^5 (green) iterations produced similar distributions for their respective settings. Therefore, for a given value N, there is significant overlap of the distributions; Figure 2. Although the SAA solution (black) generally produced smaller objective function values, the OSA distributions appear to be converging to that of the SAA solution. This is especially apparent for $N = 10^5$ and $\varDelta_N = 500$.

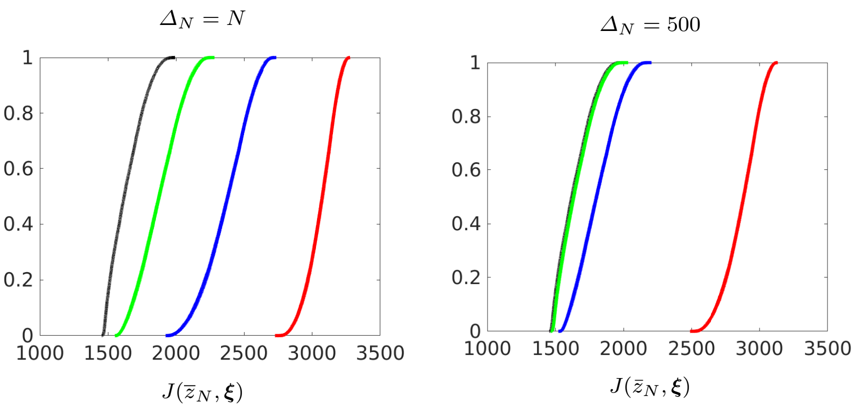

Fig. 2: Objective function empirical distribution using 10^4 samples for the SAA optimal control (black) and the OSA controls computed with 10^3 (red), 10^4 (blue), and 10^5 (green) iterations for epoch lenghts $\varDelta_N = N$ (left) and $\varDelta_N = 500$ (right).

Section 4.4 provides the mean convergence statements for feasibility (Proposition 4) and the optimal values (Theorem 1). However, since we generate our solutions via simulations, i.e., sample paths of a rather complicated stochastic process, we believe it is also of interest to conduct a further statistical test to judge the reliability of a single reported solution. This kind of post-optimality test is rarely included in the literature. Our goal is to compare the quality of solutions via hypothesis testing of the empirical distribution functions for the objective J and penalty function F out of sample.

We performed a two-sample Kolmogorov-Smirnov (KS) test using the empirical distributions (cdfs) generated by $m = 10^4$ samples of the objective and penalty functions for two computed controls z_1 and z_2 either from SAA or from OSA with fixed N. The KS statistics are defined by

$$D^J(z_1, z_2) := \sup_{t\in\mathbb{R}} |\widehat{F}_{J(z_1,\boldsymbol{\xi})}(t) - \widehat{F}_{J(z_2,\boldsymbol{\xi})}(t)|$$
$$D^F(z_1, z_2) := \sup_{t\in\mathbb{R}} |\widehat{F}_{F(\theta(z_2,\boldsymbol{\xi}))}(t) - \widehat{F}_{F(\theta(z_1,\boldsymbol{\xi}))}(t)|,$$

where $\widehat{F}_{J(z,\boldsymbol{\xi})}(\cdot)$ is the empirical cdf for the objective and $\widehat{F}_{F(\theta(z,\boldsymbol{\xi}))}(\cdot)$ is the empirical cdf for the penalty. We recall that the KS test is nonparametric and makes no assumptions about the form of the true distributions. It merely reports the maximum difference between the two cdfs. The null hypothesis for the KS test is that the random variables $J(z_1, \boldsymbol{\xi})$ and $J(z_2, \boldsymbol{\xi})$ are sampled from populations with identical distributions. Typically the null hypothesis will be rejected if the KS statistic is larger than a certain critical threshold, which depends on the number of the samples used to generate the empirical cdfs.

For our problem set up, the null hypothesis is rejected at level α if

$$D^J(z_1, z_2) > \frac{\sqrt{-0.5\ln\left(\frac{\alpha}{2}\right)}}{\sqrt{N}}$$

and similarly for $D^F(z_1, z_2)$. Typical values for α range from 0.2 to 0.001, but there are technically no restrictions.

We list the computed KS statistics in Table 1. The upper triangle of each table corresponds to the objective function and the lower triangle corresponds to the penalty. The first column compares the penalty distribution for the SAA control and the OSA controls computed using 10^3, 10^4, and 10^5 iterations. Similarly, the first row compares the objective distribution for the SAA control and the OSA controls. The (i, j)-entry for $j > i$ lists the KS statistic for the objective functions computed using the ith and jth controls. The (i, j)-entry for $j < i$ lists the KS statistic for the corresponding penalty functions. As Figure 2 already suggests, the OSA and SAA solutions generate random variables that generally appear to be from different distributions. This observation is confirmed by the KS test. A simple computation shows that the null hypothesis is rejected for the objective function in all cases with the largest α being $\alpha \approx 10^{-20}$ (i.e., $N = 10^5$ and $\Delta_N = 500$). However, the OSA objective function distributions do appear to converge to the SAA distribution as N increases (cf., Figure 2). In contrast, the null hypothesis is accepted for the objective function distributions for the different OSA controls $z^\star_{N,i}$ for $i = 1, \ldots, 5$ for any α below a minimum level of $\alpha = 0.17$, which provides confidence that the OSA controls, computed using the same N and Δ_N, generate objective function values that are drawn from the same distribution.

6 Conclusion and Outlook

PDE-constrained optimization is an important class of infinite-dimensional optimization problems. Motivated by applications in engineering and physics,

N	$\Delta_N = N$						$\Delta_N = 500$					
	0	1.0000	1.0000	1.0000	1.0000	1.0000	0	1.0000	1.0000	1.0000	1.0000	1.0000
	1.0000	0	0.0127	0.0121	0.0098	0.0130	1.0000	0	0.0108	0.0109	0.0095	0.0120
	1.0000	0.0153	0	0.0109	0.0093	0.0097	1.0000	0.0141	0	0.0116	0.0074	0.0104
10^3	1.0000	0.0100	0.0148	0	0.0111	0.0145	1.0000	0.0093	0.0156	0	0.0131	0.0145
	1.0000	0.0123	0.0086	0.0129	0	0.0080	1.0000	0.0108	0.0072	0.0137	0	0.0068
	1.0000	0.0121	0.0105	0.0145	0.0069	0	1.0000	0.0117	0.0097	0.0160	0.0070	0
	0	0.9986	0.9980	0.9983	0.9982	0.9986	0	0.4250	0.4270	0.4192	0.4301	0.4318
	0.8090	0	0.0114	0.0114	0.0101	0.0152	0.1219	0	0.0108	0.0104	0.0093	0.0150
	0.8132	0.0173	0	0.0103	0.0059	0.0115	0.1271	0.0065	0	0.0110	0.0060	0.0106
10^4	0.8078	0.0096	0.0182	0	0.0130	0.0157	0.1284	0.0096	0.0079	0	0.0130	0.0149
	0.8148	0.0140	0.0092	0.0151	0	0.0089	0.1241	0.0056	0.0040	0.0077	0	0.0086
	0.8246	0.0171	0.0116	0.0205	0.0104	0	0.1276	0.0073	0.0052	0.0060	0.0038	0
	0	0.5665	0.5691	0.5658	0.5706	0.5778	0	0.0609	0.0655	0.0629	0.0635	0.0672
	0.2055	0	0.0104	0.0105	0.0094	0.0143	0.0032	0	0.0115	0.0102	0.0110	0.0160
	0.2094	0.0067	0	0.0107	0.0065	0.0110	0.0028	0.0008	0	0.0107	0.0053	0.0131
10^5	0.2111	0.0091	0.0078	0	0.0136	0.0152	0.0042	0.0021	0.0017	0	0.0130	0.0156
	0.2156	0.0106	0.0070	0.0082	0	0.0086	0.0025	0.0013	0.0009	0.0020	0	0.0116
	0.2142	0.0125	0.0069	0.0060	0.0053	0	0.0024	0.0012	0.0008	0.0023	0.0007	0

Table 1: Kolmogorov-Smirnov (KS) statistic of the objective, D^J, and penalty functions, D^F, quantifying the discrepancy between the empirical distributions for the controls computed by SAA and OSA. Smaller values indicate that the two sets of samples come from the same distribution. The upper triangle of each table corresponds to the objective and the lower corresponds to the penalty. The first column compares the penalty distribution for the SAA optimal control and the OSA controls computed using 10^3, 10^4, and 10^5 iterations. Similarly, the first row compares the objective distribution for the SAA optimal control and the OSA controls. The (i, j)-entry for $j > i$ lists the KS statistic for the objective functions computed using the ith and jth controls. The (i, j)-entry for $j < i$ lists the KS statistic for the corresponding penalty functions.

we considered a class of convex stochastic optimization problems in which the solution of the PDE needs to satisfy a pointwise constraint in an a.s. sense. We proposed a penalty-based relaxation approach that transforms this challenging problem into a numerically tractable form. We then developed a tailor-made online stochastic approximation scheme to effectively solve the resulting convex optimization problem.

We provide a full convergence analysis in infinite-dimensional Hilbert space. Assuming that we employ a conforming spatial discretization, there is sufficient stability near the fully continuous solution, and the numerical bias can be controlled as a function of N, then the convergence statements should carry over to discretization refinements of the fully discrete problem. A deeper analysis of this, as in [13, 27], will be the subject of future research.

Data availability statement

Data sharing not applicable to this article as no datasets were generated or analysed during the current study. However, the relevant code used in the numerical experiments will be made available in future versions of the Rapid Optimization Library (ROL), which is available here: `https://trilinos.github.io/rol.html` as part of the Trilinos Project: `https://github.com/trilinos/Trilinos`.

Declarations

DPK was sponsored by the Department of Energy Office of Science under the Early Career Research Program. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. MST and TMS acknowledge the support of the COST Action CA16228 "European Network for Game Theory" for support, which made the initial discussions for this topic possible. The authors have no financial or proprietary interests in any material discussed in this article.

References

1. Adams, R.A., Fournier, J.J.F.: Sobolev spaces, *Pure and Applied Mathematics (Amsterdam)*, vol. 140, second edn. Elsevier/Academic Press, Amsterdam (2003)
2. Andrieu, L., Henrion, R., Römisch, W.: A model for dynamic chance constraints in hydro power reservoir management. European J. Oper. Res. **207**(2), 579–589 (2010). DOI 10.1016/j.ejor.2010.05.013. URL https://doi.org/10.1016/j.ejor.2010.05.013
3. Besbes, O., Gur, Y., Zeevi, A.: Non-stationary stochastic optimization. Operations Research **63**(5), 1227–1244 (2015). URL http://pubsonline.informs.org/doi/abs/10.1287/opre.2015.1408
4. Bonnans, J.F., Shapiro, A.: Perturbation analysis of optimization problems. Springer Science & Business Media (2013)
5. Branda, M.: Stochastic programming problems with generalized integrated chance constraints. Optimization **61**(8), 949–968 (2012). DOI 10.1080/02331934.2011.587007. URL https://doi.org/10.1080/02331934.2011.587007
6. Branda, M., Dupacova, J.: Approximation and contamination bounds for probabilistic programs. Annals of Operations Research **193**(1), 3–19 (2012)
7. Conn, A.R., Gould, N.I.M., Toint, P.L.: A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds. SIAM Journal on Numerical Analysis **28**(2), 545–572 (1991). DOI 10.1137/0728030
8. Conti, S., Rumpf, M., Schultz, R., Tölkes, S.: Stochastic dominance constraints in elastic shape optimization. SIAM J. Control Optim. **56**(4), 3021–3034 (2018). DOI 10.1137/16M108313X. URL https://doi.org/10.1137/16M108313X
9. Duvocelle, B., Mertikopoulos, P., Staudigl, M., Vermeulen, D.: Multi-agent online learning in time-varying games. Forthcoming: Mathematics of Operations Research, arXiv preprint arXiv:1809.03066 (2022)
10. Farshbaf-Shaker, M.H., Henrion, R., Hömberg, D.: Properties of chance constraints in infinite dimensions with an application to PDE constrained optimization. Set-Valued and Variational Analysis **26**(4), 821–841 (2018). DOI 10.1007/s11228-017-0452-5. URL https://doi.org/10.1007/s11228-017-0452-5
11. Gahururu, D.B., Hintermüller, M., Surowiec, T.M.: Risk-neutral PDE-constrained generalized Nash equilibrium problems. Mathematical Programming (2022). DOI 10.1007/s10107-022-01800-z. URL https://doi.org/10.1007/s10107-022-01800-z

12. Geiersbach, C., Pflug, G.C.: Projected stochastic gradients for convex constrained problems in hilbert spaces. SIAM Journal on Optimization **29**(3), 2079–2099 (2019). DOI 10.1137/18M1200208. URL https://doi.org/10.1137/18M1200208
13. Geiersbach, C., Wollner, W.: A stochastic gradient method with mesh refinement for PDE-constrained optimization under uncertainty. SIAM Journal on Scientific Computing **42**(5), A2750–A2772 (2020). DOI 10.1137/19M1263297. URL https://doi.org/10.1137/19M1263297
14. Guth, P.A., Kaarnioja, V., Kuo, F.Y., Schillings, C., Sloan, I.H.: A quasi-Monte Carlo method for optimal control under uncertainty. SIAM/ASA Journal on Uncertainty Quantification **9**(2), 354–383 (2021). DOI 10.1137/19M1294952. URL https://doi.org/10.1137/19M1294952
15. Haneveld, W.K.K., Van Der Vlerk, M.H.: Integrated chance constraints: reduced forms and an algorithm. Computational Management Science **3**(4), 245–269 (2006)
16. Henrion, R., Möller, A.: A gradient formula for linear chance constraints under Gaussian distribution. Mathematics of Operations Research **37**(3), 475–488 (2012). DOI 10.1287/moor.1120.0544. URL https://doi.org/10.1287/moor.1120.0544
17. Hille, E., Phillips, R.S.: Functional analysis and semi-groups. American Mathematical Society Colloquium Publications, vol. 31. American Mathematical Society, Providence, R. I. (1957). Rev. ed
18. Hinze, M., Pinnau, R., Ulbrich, M., Ulbrich, S.: Optimization with PDE constraints, *Mathematical Modelling: Theory and Applications*, vol. 23. Springer, New York (2009)
19. Kouri, D.P., Heinkenschloss, M., Ridzal, D., van Bloemen Waanders, B.G.: A trust-region algorithm with adaptive stochastic collocation for PDE optimization under uncertainty. SIAM Journal on Scientific Computing **35**(4), A1847–A1879 (2013). DOI 10.1137/120892362. URL https://doi.org/10.1137/120892362
20. Kouri, D.P., Heinkenschloss, M., Ridzal, D., van Bloemen Waanders, B.G.: Inexact objective function evaluations in a trust-region algorithm for PDE-constrained optimization under uncertainty. SIAM J. Sci. Comput. **36**(6), A3011–A3029 (2014). DOI 10.1137/140955665. URL https://doi.org/10.1137/140955665
21. Kouri, D.P., Surowiec, T.M.: Existence and optimality conditions for risk-averse PDE-constrained optimization. SIAM/ASA Journal on Uncertainty Quantification **6**(2), 787–815 (2018)
22. Kouri, D.P., Surowiec, T.M.: Risk-averse optimal control of semilinear elliptic PDEs. ESAIM Control Optim. Calc. Var. **26**, Paper No. 53, 19 (2020). DOI 10.1051/cocv/2019061. URL https://doi.org/10.1051/cocv/2019061
23. Lan, G., Zhou, Z.: Algorithms for stochastic optimization with function or expectation constraints. Computational Optimization and Applications **76**(2), 461–498 (2020). DOI 10.1007/s10589-020-00179-x. URL https://doi.org/10.1007/s10589-020-00179-x
24. Lin, C.J., Moré, J.J.: Newton's method for large bound-constrained optimization problems. SIAM Journal on Optimization **9**(4), 1100–1127 (1999)
25. Lions, J.L.: Optimal Control of Systems Governed by Partial Differential Equations, *Grundlehren der mathematischen Wissenschaften*, vol. 170. Springer-Verlag Berlin Heidelberg (1971)
26. Marti, K.: Differentiation of probability functions: The transformation method. Computers & Mathematics with Applications **30**(3), 361–382 (1995). DOI https://doi.org/10.1016/0898-1221(95)00113-1. URL https://www.sciencedirect.com/science/article/pii/0898122195001131
27. Martin, M., Nobile, F.: PDE-constrained optimal control problems with uncertain parameters using SAGA. SIAM/ASA Journal on Uncertainty Quantification **9**(3), 979–1012 (2021). DOI 10.1137/18M1224076. URL https://doi.org/10.1137/18M1224076
28. Martin, Matthieu, Krumscheid, Sebastian, Nobile, Fabio: Complexity analysis of stochastic gradient methods for PDE-constrained optimal control problems with uncertain parameters. ESAIM: M2AN **55**(4), 1599–1633 (2021). DOI 10.1051/m2an/2021025. URL https://doi.org/10.1051/m2an/2021025

29. Nemirovski, A., Juditsky, A., Lan, G., Shapiro, A.: Robust stochastic approximation approach to stochastic programming. SIAM Journal on optimization **19**(4), 1574–1609 (2009)
30. Prékopa, A.: Stochastic programming, *Mathematics and its Applications*, vol. 324. Kluwer Academic Publishers Group, Dordrecht (1995). DOI 10.1007/978-94-017-3087-7. URL http://dx.doi.org/10.1007/978-94-017-3087-7
31. Ridzal, D., Kouri, D.P., von Winckel, G.J.: Rapid optimization library. Tech. rep., Sandia National Lab.(SNL-NM), Albuquerque, NM (United States) (2017)
32. Shapiro, A., Dentcheva, D., Ruszczyński, A.: Lectures on Stochastic Programming. Society for Industrial and Applied Mathematics (2009). DOI doi:10.1137/1.9780898718751. URL https://doi.org/10.1137/1.9780898718751
33. Tröltzsch, F.: Optimal Control of Partial Differential Equations. Graduate Studies in Mathematics. American Mathematical Society (2010)
34. Vaart, A.W.v.d.: Asymptotic Statistics. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press (1998). DOI 10.1017/CBO9780511802256
35. Van Barel, A., Vandewalle, S.: MG/OPT and multilevel Monte Carlo for robust optimization of PDEs. SIAM Journal on Optimization **31**(3), 1850–1876 (2021). DOI 10.1137/20M1347164. URL https://doi.org/10.1137/20M1347164
36. Vladarean, M.L., Alacaoglu, A., Hsieh, Y.P., Cevher, V.: Conditional gradient methods for stochastically constrained convex minimization. In: H.D. III, A. Singh (eds.) Proceedings of the 37th International Conference on Machine Learning, *Proceedings of Machine Learning Research*, vol. 119, pp. 9775–9785. PMLR (2020). URL https://proceedings.mlr.press/v119/vladarean20a.html
37. Yu, H., Neely, M., Wei, X.: Online convex optimization with stochastic constraints. Advances in Neural Information Processing Systems pp. 1428–1438 (2017)