A Riemannian Subgradient Algorithm for Economic Dispatch with Valve-Point Effect*

Pierre B. Borckmans^{\dagger ‡} S. Easter Selvan^{\dagger}

Nicolas Boumal^{†§}

P.-A. Absil[†]

July 13, 2013

Abstract

The economic load dispatch problem (ELDP) is a classical problem in the power systems community. It consists in the optimal scheduling of the output of power generating units to meet the required load demand subject to unit and system inequality and equality constraints. This optimization problem is challenging on three different levels: the geometry of its feasible set, the non-differentiability of its cost function and the multimodal aspect of its landscape. For this reason, ELDP has received much attention in the past few years and numerous derivative-free techniques have been proposed to tackle its multimodal and nondifferentiable characteristics. In this work, we propose a different approach exploiting the rich geometrical structure of the problem. We show that the (nonlinear) equality constraint can be handled in the framework of Riemannian manifolds and we develop a feasible (all iterates satisfy the constraints) subgradient descent algorithm to provide fast convergence to local minima. To this end, we show that Clarke's calculus can be used to compute a deterministic admissible descent direction by solving a simple, low-dimensional quadratic program. We test our approach on four real data sets. The proposed method provides fast local convergence and scales well with respect to the problem dimension. Finally, we show that the proposed algorithm, being a local optimization method, can be incorporated in existing heuristic techniques to provide better exploration of the search space.

Key words: economic load dispatch; nonsmooth optimization; differential geometry

1 Introduction

The economic load dispatch problem (ELDP) is the optimal scheduling of the output of power generating units to meet the required load demand subject to unit and system equality and inequality constraints [SCC03]. In the traditional ELDP, the cost function for each generator is modeled by a single quadratic function. Nevertheless, in practice, one has to take into account highly nonlinear input–output characteristics arising due to valve-point loadings or generating unit ramp rate limits. As a consequence, we end up with a nonsmooth, equality- and inequality-constrained optimization problem, (5), which is in general multimodal (it presents several local optima) and for which classical smooth optimization techniques are thus not suitable.

For this reason, the ELDP has received much attention in the past few years and numerous derivative-free techniques have been proposed to tackle its multimodal and nondifferentiable aspects. Popular techniques include genetic algorithms (GA) [WS93], evolutionary programming [YYH96], particle swarm optimisation (PSO) [PLSL05], and differential evolution (DE) [NI08].

^{*}This paper presents research results of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme initiated by the Belgian Science Policy Office. This work was financially supported by the Belgian FRFC (Fonds de la Recherche Fondamentale Collective).

[†]Department of Mathematical Engineering, ICTEAM Institute, Université catholique de Louvain, B-1348 Louvain-la-Neuve, Belgium (http://sites.uclouvain.be/absil/).

[‡]Supported by a FRIA (Fonds pour la formation à la Recherche dans l'Industrie et dans l'Agriculture) fellowship. [§]Supported by a FNRS (Fonds National de la Recherche Scientifique) fellowship.

More recently, "global-local" hybrid methods have appeared that combine a "global" method having good global searching abilities and a "local" method with better fine-tuning abilities. Even though there may not be a clear frontier between the two groups of methods, it is quite evident that methods such as GA, PSO and DE belong to the "global" group whereas a method such as sequential quadratic programming (SQP) belongs to the "local" group. Desirable properties of a local method include the following. (i) The method should be *feasible*, i.e., all the iterates should satisfy the equality and inequality constraints. Indeed, a drawback of infeasible methods is that even if the outcome is close to satisfying the constraints, the cost function value at the outcome may be quite different from its value at the nearest point satisfying the constraints. Another downside of infeasibility is that the intermediate results are not directly usable since they do not respect the balance constraint. (ii) The iterates should converge to a local minimizer of the cost function. (iii) There is a well-known trade-off between low numerical cost of the iteration and fast convergence of the iterates; as a simplistic illustration, concatenating a few steps of a given method improves the speed of convergence per iterate but degrades the numerical cost per iterate. If the sought accuracy is sufficiently low—which is arguably the case of the ELDP where the coefficient values and even the model itself are debatable—then it is preferable to run a lot of cheaper iterations, as this makes it possible to check a stopping criterion more frequently and hence to avoid the unnecessary computational effort inherent to overconvergence. (iv) Methods that exploit the very particular structure of the ELDP should be preferred over "out of the box" numerical algorithms.

SQP is a popular choice for the local method in global-local ELDP algorithms; see, e.g., [VJ04, AKTH02, CLL⁺12, ASAO10]. Whereas SQP is a welcome complement to global search methods, it does not satisfy any of the properties mentioned above. Specifically, it is an infeasible method, since the equality constraint is only satisfied in the limit. Convergence, while empirically observed, is not guaranteed because the classical convergence theory of SQP assumes, among other things, that the cost function and the constraint functions have continuous first derivatives [NW06, theorem 18.3], which is not the case in the ELDP (5). As a quadratically convergent method [NW06, theorem 18.4], it can be seen as favoring fast convergence of the iterates at the expense of a higher numerical cost per iterate. Finally, it is a general-purpose algorithm, not specifically tailored to the ELDP.

Other local methods have been proposed in combination with global methods. The Nelder-Mead (NM) method has been shown to be an effective local method in combination with PSO [Nik10], but it does not exploit the readily accessible first-order information on the cost function. The modified subgradient (MSG) method used in [YÖ11] exploits first-order information on the cost function, but feasibility of the iterates is achieved only after a certain number of steps in view of the *sharp* augmented Lagrangian approach. Shor's r-algorithm, combined with an improved differential evolution (IDE) method in [YWZY09], is also of the subgradient type, but the iterates are not feasible.

In this paper, we develop a novel method that satisfies properties (i), (iii), and (iv). It is also built to satisfy (ii) and does it empirically, however, as a consequence of (iv), its convergence does not immediately follow from an existing result, and a detailed convergence analysis is beyond the scope of this paper.

Property (i), feasibility of the iterates, is enforced using the framework of Riemannian optimization. The adequacy of this framework stems from the fact that the equality constraint (3) defines an ellipsoid, which admits a natural structure of a Riemannian manifold. Riemannian optimization, also called optimization on manifolds, is nowadays a vibrant area of research, whose foundations can be found, e.g., in [HM94, EAS98, AMS08]. Since this framework is new in the ELDP context, we dedicate a significant part of this paper to laying out the necessary background.

Property (iii) comes by preferring a steepest-descent approach over a second-order approach. However, in view of the non-smoothness due to the valve-point effect, the steepest-descent approach does not rely on gradient techniques but rather on subgradient techniques, using the framework of Clarke's generalized calculus [Cla75].

In view of the above, the proposed method fits in the framework of Riemannian subgradient descent. With respect to the general-purpose Riemannian subgradient descent method of Dirr et

al. [DHL07], a contribution of our development is to incorporate bound constraints in order to handle the generator capacity constraints (2) present in the ELDP. Another contribution of this work is that, whereas many heuristic algorithms for the ELDP consist of (a combination of) existing black-box optimization techniques, the proposed method strives to exploit as much as possible the very particular structure of the ELDP. This allows notably for an efficient representation of the generalized gradient, which enables fast computation of a descent direction by solving a low-dimensional quadratic program.

In summary, the proposed technique provides fast convergence to a nearby local minimum of the ELDP (5), while satisfying the power balance (3) and capacity constraints (2) throughout the optimization process. Therefore, the aforementioned heuristics largely explored in the literature and the proposed subgradient descent algorithm present very complementary properties: the multimodal aspect of the ELDP can be addressed using any global feasible exploration tool while the local refinement of a potential solution is efficiently provided by the proposed approach, including a check for the stationarity of the final iterate.

The remainder of the paper is organized as follows. In Section 2, the ELDP with the valvepoint effect is briefly presented, followed by a detailed treatment of the underlying geometry of the optimization problem and an introduction to the necessary differential geometry tools in Section 3. The subgradient descent algorithm is presented in Section 4. Subsequently, its formulation on the Riemannian manifold and its specialization for the ELDP are discussed. Section 5 presents the implementation details. In Section 6, numerical results are presented and we show how our algorithm can be hybridized with a global scheme to deal with the multimodal aspect of the ELDP. Finally, conclusions are drawn in Section 7.

2 ELDP Considering Valve-point Effect

2.1 Problem Statement

Traditionally, the generating-unit cost functions are considered to be convex with the heat rate curves exhibiting monotonically increasing characteristics. However, in reality, the steam admission valves in the large steam turbines cause discontinuities in the incremental heat rate curves. Therefore, to accurately model the ELDP, the valve-point loadings in the n generating units are to be incorporated, leading to nonconvex input-output characteristics of the generating units [VJ05]; the cost function is then stated as

$$f_{\rm T}(\boldsymbol{p}) = \sum_{i=1}^{n} f^{i}(p_{i}), \qquad \boldsymbol{p} = [p_{1}, p_{2}, \dots, p_{n}]^{\rm T}$$
$$= \sum_{i=1}^{n} a_{i} p_{i}^{2} + b_{i} p_{i} + c_{i} + \left| d_{i} \sin \left[e_{i} \left(p_{i}^{\min} - p_{i} \right) \right] \right|.$$
(1)

Here $f_{\rm T}(\mathbf{p})$ is the total production cost (\$/h) pertaining to the *n*-dimensional output power vector \mathbf{p} and $f^i(p_i)$ is the incremental fuel cost function (\$/h) pertaining to the real power output of the *i*th unit, p_i . For the *i*th generating unit, the cost coefficients are denoted by a_i, b_i, c_i , and the constants from the valve-point effect by d_i, e_i . All the coefficients are positive. In matrix notation, we have

$$f_{\mathrm{T}}(\boldsymbol{p}) = \boldsymbol{p}^{\mathsf{T}}\mathrm{Diag}(\boldsymbol{a})\boldsymbol{p} + \boldsymbol{b}^{\mathsf{T}}\boldsymbol{p} + \boldsymbol{c}^{\mathsf{T}}\boldsymbol{1} + \boldsymbol{d}^{\mathsf{T}}\left|\sin\left[\boldsymbol{e}\circ(\boldsymbol{p}^{\min}-\boldsymbol{p})
ight]
ight|,$$

where \circ denotes the component-wise product, $\text{Diag}(\cdot)$ denotes the diagonal matrix obtained from the entries of its vector argument, and the sine term and the absolute value are taken component-wise.

The cost function $f_{\rm T}(\mathbf{p})$ is to be minimized subject to the following inequality and equality constraints:

(a) generator capacity constraints

$$p_i^{\min} \le p_i \le p_i^{\max}, \quad i = 1, \dots, n \tag{2}$$

where p_i^{\min} and p_i^{\max} are the lower and upper power generating limits of the *i*th unit (MW); (b) real power balance constraint

$$\sum_{i=1}^{n} p_i = p_{\rm D} + p_{\rm L}(\boldsymbol{p}), \tag{3}$$

where $p_{\rm D}$ is the power demand (MW) and $p_{\rm L}(p)$ stands for the power loss (MW) expressed as

$$p_{\rm L}(\boldsymbol{p}) = \sum_{i=1}^{n} \sum_{j=1}^{n} p_i B_{ij} p_j + \sum_{i=1}^{n} b_i^0 p_i + b^{00}, \qquad (4)$$

or in matrix notation $p_{\rm L}(\boldsymbol{p}) = \boldsymbol{p}^{\mathsf{T}} \boldsymbol{B} \boldsymbol{p} + \boldsymbol{p}^{\mathsf{T}} \boldsymbol{b}^0 + b^{00}$. Coefficients B_{ij}, b_i^0, b^{00} are the transmission loss coefficients (B-coefficients) given by the elements of the square matrix \boldsymbol{B} of size $n \times n$, the vector \boldsymbol{b}^0 of length n, and the constant b^{00} , respectively. The matrix \boldsymbol{B} is symmetric positive-definite, hence $p_{\rm L}(\boldsymbol{p})$ is a convex quadratic function of \boldsymbol{p} .

To summarize, we consider the following optimization problem, that we will refer to as the ELDP:

$$\min_{\boldsymbol{p}\in\mathbb{R}^n} f_{\mathrm{T}}(\boldsymbol{p}) \quad (1)$$
subject to (2) and (3).
(5)

2.2 Geometry of the Feasible Set

In this subsection, we explore the geometrical interpretation and consequences of constraints (2) and (3).

The bound constraints (2) force p to lie inside a box whose faces are parallel to the reference frame. The equality constraint (3) imposes that p lies on a quadric surface, specifically on an ellipsoid since B is positive definite. The feasible set of (5), denoted by Ω , is thus composed of the points on the ellipsoid that are also inside the box:

$$\Omega := \{ \boldsymbol{p} \in \mathbb{R}^n : \boldsymbol{p} \text{ satisfies } (2) \text{ and } (3) \}$$

$$= \{ \boldsymbol{p} \in \mathbb{R}^n : \boldsymbol{p}^{\min} \le \boldsymbol{p} \le \boldsymbol{p}^{\max}, \ \boldsymbol{p}^{\top} \boldsymbol{B}^{\text{ell}} \boldsymbol{p} + \boldsymbol{p}^{\top} \boldsymbol{b}^{\text{ell}} + c^{\text{ell}} = 0 \},$$
(6)

where **1** is the vector of all ones, $\mathbf{B}^{\text{ell}} = \mathbf{B}$, $\mathbf{b}^{\text{ell}} = \mathbf{b}^{\mathbf{0}} - \mathbf{1}$, and $c^{\text{ell}} = b^{00} + p_d$. In general, Ω may be composed of more than one connected component, but we did not observe this in practical ELDP instances. The feasible set Ω is depicted in Figure 1 for the cases n = 2 and n = 3.



Figure 1: Representation of the **p**-space along with the constraints (inside the box and on the ellipsoid) and the feasible set Ω (darker blue). Notice that Ω has measure zero in \mathbb{R}^n .

2.3 Structure of the Optimization Landscape

In this subsection, we provide some insight about the cost function $f_{\mathrm{T}}(\mathbf{p})$ defined in (1). A first observation is that this function is separable, i.e., it consists in a sum of components f^i , each depending only on p_i , the *i*th component of \mathbf{p} . However, problem (5) does not decompose into independent one-dimensional problems because the components of \mathbf{p} are coupled through the equality constraint (3). The cost function $f_{\mathrm{T}}(\mathbf{p})$ is continuous but not everywhere differentiable. Indeed, at any point $\mathbf{q} \in \mathbb{R}^n$ for which one or more components q_i cancel the corresponding sine term in f^i , the gradient of f_{T} does not exist because of the absolute value. These points, termed *nondifferentiable points* or *kink points*, are given by the following expression:

$$\boldsymbol{q} = (q_1, \dots, q_n), \; \exists i \text{ s.t. } q_i = p_i^{\min} + \frac{k\pi}{e_i}, k \in \mathbb{N}.$$
(7)

Nevertheless, at all other points p in \mathbb{R}^n , the gradient of f_T can be computed:

$$\nabla f_{\rm T}(\boldsymbol{p}) = 2\text{Diag}(\boldsymbol{a})\boldsymbol{p} + \boldsymbol{b} - \boldsymbol{d} \circ \boldsymbol{e} \circ \cos \boldsymbol{\theta} \circ \text{sign}(\sin \boldsymbol{\theta}), \tag{8}$$

where θ stands for $[e \circ (p^{\min} - p)]$. The function $f_{\rm T}(p)$ is thus piecewise smooth (as defined in [Roc03]) and one can use the canvas of Clarke's generalized calculus [Cla75] to compute its generalized gradient at the nondifferentiable points. This analysis will be carried out in Section 4.

Now, recall that problem (5) consists of minimizing $f_{\rm T}(\mathbf{p})$ over the feasible set Ω . As depicted in Figure 2 for an illustrative 2-dimensional case, the restriction of $f_{\rm T}(\mathbf{p})$ to Ω presents a partially smooth and multimodal landscape. If the portion Ω of the ellipsoid is not too large, a similar



Figure 2: Landscape of the cost function $f_{\rm T}(\mathbf{p})$ for an illustrative 2-dimensional case. (a) Level curves of $f_{\rm T}(\mathbf{p})$ and the ellipsoid constraint. (b) Graph of the cost function $f_{\rm T}(\mathbf{p})$. (c) Graphs of the restriction of $f_{\rm T}(\mathbf{p})$ to Ω . The cost function is multimodal and piecewise smooth.

representation can be obtained for a 3-dimensional case, by projecting the feasible set Ω onto the nearest 2-dimensional plane; see Figure 3, where the observations made for the 2-dimensional case can also be noted.

2.4 Summary of the Optimization Challenges

From the previous sections, it appears that the ELDP is a challenging optimization problem on three different levels: the geometry of its feasible set, the non-differentiability of its cost function, and the multimodal aspect of its landscape. The rest of this paper presents our approach to deal with these difficulties.



Figure 3: Landscape of the cost function $f_{\rm T}(\mathbf{p})$ for an illustrative 3-dimensional case. This representation is made possible by projecting the 3 dimensional feasible set Ω onto the nearest plane. (a) Level curves of the restriction of $f_{\rm T}(\mathbf{p})$ to the projection $\hat{\Omega}$ of Ω . (b) Graph of the restriction of $f_{\rm T}(\mathbf{p})$ to $\hat{\Omega}$. The cost function is multimodal and piecewise smooth.

3 Optimization Exploiting the Geometry of Ω

The proposed local method can be viewed as an adaptation of the classical line-search scheme

$$\boldsymbol{p}^{k+1} = \boldsymbol{p}^k + \alpha^k \boldsymbol{d}^k, \tag{9}$$

where \mathbf{p}^k and \mathbf{p}^{k+1} denote the current and next iterates, \mathbf{d}^k is the search direction, and α^k is the step length. In this section, with property (i)—feasibility—in mind, we generalize the "+" operation by means of the concept of retraction and we introduce "admissibility" conditions on \mathbf{d}^k in order to guarantee that \mathbf{p}^{k+1} remains in Ω for all α^k sufficiently small. Then, in Section 4, we will show how to choose \mathbf{d}^k as the *steepest* admissible direction. Finally, the selection of α^k , as well as other implementation details, will be addressed in Section 5.

3.1 The Ellipsoid Manifold

As explained in Section 2.2, the feasible set Ω of the ELDP (5) has a rich geometrical structure that we would like to exploit. The set of points that satisfy the equality constraint (3) is an ellipsoid centered at

$$\boldsymbol{a}^{\text{ell}} = -\frac{1}{2} (\boldsymbol{B}^{\text{ell}})^{-1} \boldsymbol{b}^{\text{ell}}.$$
 (10)

This smooth surface has a natural structure of a manifold, specifically of an (n-1)-dimensional submanifold of \mathbb{R}^n . We will call it the ellipsoid manifold:

$$\mathcal{E}^{n-1} := \{ \boldsymbol{p} \in \mathbb{R}^n : \boldsymbol{p} \text{ satisfies (3)} \}$$

$$= \{ \boldsymbol{p} \in \mathbb{R}^n : \boldsymbol{p}^\top \boldsymbol{B}^{\text{ell}} \boldsymbol{p} + \boldsymbol{p}^\top \boldsymbol{b}^{\text{ell}} + c^{\text{ell}} = 0 \}.$$
(11)

The gist of the proposed Riemannian optimization approach is to restrict the optimization domain \mathbb{R}^n to the constraint manifold, in this case with the resulting advantage that every iterate belongs to \mathcal{E}^{n-1} , i.e., satisfies the equality constraint (3). The ELDP (5) then becomes:

$$\min_{oldsymbol{p}\in\mathcal{E}^{n-1}} f_{\mathrm{T}}(oldsymbol{p}) \quad \mathrm{subject \ to} \quad oldsymbol{p}^{\min} \leq oldsymbol{p} \leq oldsymbol{p}^{\max}$$

We mention that an alternative way, used by several ELDP heuristics, of respecting the equality constraint (3) is to resort to a slack variable, which amounts to performing the optimization on (n-1) variables while computing the last variable explicitly. A difficulty with this approach is that

the description of the cost function $f_{\rm T}$, of the feasible set Ω , and of the nondifferentiable points in terms of the (n-1) remaining variables becomes more intricate. This difficulty is avoided in the Riemannian approach.

Riemannian Optimization Ingredients 3.2

The classical line-search scheme (9) is not meant to produce iterates that remain on a submanifold such as \mathcal{E}^{n-1} . In this subsection, we provide the needed ingredients of differential geometry (the branch of mathematics that studies manifolds) in order to produce iterates on an abstract submanifold \mathcal{M} of \mathbb{R}^n . Then, in Section 3.3, we will specialize the ingredients to the manifold \mathcal{E}^{n-1} , and finally, in Section 3.5, we will reintroduce the bound constraints (2).

Instrumental to the definition of an admissible direction d^k is the notion of tangent space to \mathcal{M} at a point $p \in \mathcal{M}$, denoted by $T_{p}\mathcal{M}$, and defined as the following vector subspace of \mathbb{R}^{n} :

$$T_{\boldsymbol{p}}\mathcal{M} = \{\boldsymbol{\xi} \in \mathbb{R}^n : \exists c : \mathbb{R} \to \mathcal{M} \text{ with } c(0) = \boldsymbol{p}, \ c'(0) = \boldsymbol{\xi}\},\$$

where c'(0) is the usual derivative at 0 of the curve c (assumed to exist). From this definition, we see that $T_{\boldsymbol{p}}\mathcal{M}$ is the set of vectors that are tangent to the manifold at \boldsymbol{p} . Geometrically, this notion coincides with the concept of tangent plane to a smooth surface, as depicted in Figure 4. The tangent bundle T \mathcal{M} is the collection of the tangent spaces at all $p \in \mathcal{M}$.

The notion of *steepest* admissible direction will require a norm on $T_{\mathbf{p}}\mathcal{M}$. Each tangent space $T_{p}\mathcal{M}$ is a vector space and as such can be endowed with an inner product $\langle \cdot, \cdot \rangle_{p}$. The natural way of doing this is by restricting the canonical inner product of \mathbb{R}^n to $T_p\mathcal{M}$, i.e.,

$$\langle .,. \rangle_{\boldsymbol{p}} : \mathrm{T}_{\boldsymbol{p}}\mathcal{M} \times \mathrm{T}_{\boldsymbol{p}}\mathcal{M} \to \mathbb{R} : (\boldsymbol{\xi}, \boldsymbol{\zeta}) \mapsto \langle \boldsymbol{\xi}, \boldsymbol{\zeta} \rangle_{\boldsymbol{p}} = \langle \boldsymbol{\xi}, \boldsymbol{\zeta} \rangle = \boldsymbol{\xi}^{\mathsf{T}} \boldsymbol{\zeta}.$$

We then say that \mathcal{M} is a *Riemannian submanifold* of \mathbb{R}^n . The inner product induces a notion of norm: $\|\boldsymbol{\xi}\|_{\boldsymbol{p}} = \langle \boldsymbol{\xi}, \boldsymbol{\xi} \rangle_{\boldsymbol{p}}^{1/2}$. The normal space to \mathcal{M} at a point $\boldsymbol{p} \in \mathcal{M}$, denoted as $N_{\boldsymbol{p}}\mathcal{M}$, is the orthogonal complement

of $T_{\boldsymbol{p}}\mathcal{M}$ in \mathbb{R}^n :

$$N_{\boldsymbol{p}}\mathcal{M} = (T_{\boldsymbol{p}}\mathcal{M})^{\perp}.$$

One can then compute the projection $P_{\mathbf{p}}(\mathbf{v})$ of a vector $\mathbf{v} \in \mathbb{R}^n$ onto $T_{\mathbf{p}}\mathcal{M}$ by removing the normal component of \boldsymbol{v} .

We now turn to the generalization of the "+" operation of (9). A retraction on \mathcal{M} [ADM⁺02, AMS08] is a smooth mapping R from the tangent bundle T \mathcal{M} onto \mathcal{M} that satisfies $R(0_p) = p$ for all \boldsymbol{p} (where $0_{\boldsymbol{p}}$ denotes the origin of $T_{\boldsymbol{p}}\mathcal{M}$) and $\frac{d}{dt}R(t\boldsymbol{\xi}_{\boldsymbol{p}})|_{t=0} = \boldsymbol{\xi}_{\boldsymbol{p}}$ for all $\boldsymbol{\xi}_{\boldsymbol{p}} \in T_{\boldsymbol{p}}\mathcal{M}$. The restriction of R to $T_{\boldsymbol{p}}\mathcal{M}$ is denoted by $R_{\boldsymbol{p}}$. Observe that R given by $R_{\boldsymbol{p}}(\alpha d) = \boldsymbol{p} + \alpha d$ is a valid retraction in \mathbb{R}^n , hence the generalization.

Using these tools, one can adapt the iterative process (9) as follows: assuming that $p^k \in \mathcal{M}$, construct $\boldsymbol{p}^{k+1} \in \mathcal{M}$ according to

$$\boldsymbol{p}^{k+1} = R_{\boldsymbol{p}^k} \left(\alpha^k P_{\boldsymbol{p}}^k(\boldsymbol{d}^k) \right), \ \alpha^k \in \mathbb{R}, \ \boldsymbol{d}^k \in \mathbb{R}^n.$$

Optimization Ingredients on \mathcal{E}^{n-1} 3.3

We now specifically provide the aforementioned tools for the ellipsoid manifold \mathcal{E}^{n-1} . A schematic diagram is given in Figure 4.

Let c(t) be a curve on \mathcal{E}^{n-1} parametrized by $t \in \mathbb{R}$:

$$c: \mathbb{R} \to \mathcal{E}^{n-1}: t \to c(t), \text{ s.t. } c(0) = \boldsymbol{p}, \ \dot{c}(0) = \boldsymbol{\xi}$$

For all $t \in \mathbb{R}$, the ellipsoid equation gives $c(t)^{\top} B^{\text{ell}} c(t) + c(t)^{\top} b^{\text{ell}} + c^{\text{ell}} = 0$. Taking the derivative with respect to t and evaluating at t = 0 (substituting for **p** and $\boldsymbol{\xi}$), one obtains:

$$\boldsymbol{\xi}^{\mathsf{T}} \boldsymbol{B}^{\mathrm{ell}} \boldsymbol{p} + \boldsymbol{p}^{\mathsf{T}} \boldsymbol{B}^{\mathrm{ell}} \boldsymbol{\xi} + \boldsymbol{\xi}^{\mathsf{T}} \boldsymbol{b}^{\mathrm{ell}} = 0.$$

Since B^{ell} is a symmetric matrix, one can rewrite:

$$\boldsymbol{\xi}^T (2\boldsymbol{B}^{\text{ell}}\boldsymbol{p} + \boldsymbol{b}^{\text{ell}}) = 0.$$

The *tangent space* is thus defined as follows:

$$T_{\boldsymbol{p}} \mathcal{E}^{n-1} = \{ \boldsymbol{\xi} \in \mathbb{R}^n : \boldsymbol{\xi}^{\mathsf{T}} (2\boldsymbol{B}^{\mathrm{ell}} \boldsymbol{p} + \boldsymbol{b}^{\mathrm{ell}}) = 0 \}.$$
(12)

The *normal space* is then obtained by:

$$N_{\boldsymbol{p}} \mathcal{E}^{n-1} = \{ \boldsymbol{\nu} \in \mathbb{R}^{n} : \boldsymbol{\xi}^{T} \boldsymbol{\nu} = 0, \ \forall \boldsymbol{\xi} \in T_{\boldsymbol{p}} \mathcal{E}^{n-1} \}$$

= $\{ \tau \left(2\boldsymbol{B}^{\text{ell}} \boldsymbol{p} + \boldsymbol{b}^{\text{ell}} \right), \ \tau \in \mathbb{R} \}.$ (13)

The projection $P_{p}(v)$ of a vector $v \in \mathbb{R}^{n}$ onto $T_{p}\mathcal{E}^{n-1}$ can then be constructed so as to remove the normal component of v:

$$P_{\boldsymbol{p}}(v) = \boldsymbol{v} - \tau (2\boldsymbol{B}^{\mathrm{ell}}\boldsymbol{p} + \boldsymbol{b}^{\mathrm{ell}})$$

where the value of $\tau \in \mathbb{R}$ must be determined to ensure that $P_{p}(v)$ belongs to $T_{p}\mathcal{E}^{n-1}$:

$$(\boldsymbol{v} - \tau (2\boldsymbol{B}^{\text{ell}}\boldsymbol{p} + \boldsymbol{b}^{\text{ell}}))^{\mathsf{T}} (2\boldsymbol{B}^{\text{ell}}\boldsymbol{p} + \boldsymbol{b}^{\text{ell}}) = 0,$$

which yields

$$\tau = \frac{\boldsymbol{v}^{\mathsf{T}}(2\boldsymbol{B}^{\mathrm{ell}}\boldsymbol{p} + \boldsymbol{b}^{\mathrm{ell}})}{\left\| (2\boldsymbol{B}^{\mathrm{ell}}\boldsymbol{p} + \boldsymbol{b}^{\mathrm{ell}}) \right\|^2}.$$

Defining $n_p = \frac{(2B^{\text{ell}}p + b^{\text{ell}})}{\|(2B^{\text{ell}}p + b^{\text{ell}})\|}$ as the unit normal vector to the ellipsoid at p, the projection becomes:

$$P_{\boldsymbol{p}}(\boldsymbol{v}) = \boldsymbol{v} - \boldsymbol{n}_{\boldsymbol{p}} \boldsymbol{n}_{\boldsymbol{p}}^{\top} \boldsymbol{v} = (I - \boldsymbol{n}_{\boldsymbol{p}} \boldsymbol{n}_{\boldsymbol{p}}^{\top}) \boldsymbol{v}.$$
(14)

It remains to choose a retraction R on \mathcal{E}^{n-1} . On a Riemannian manifold, the exponential map, based on geodesics, is a possible choice for R [AMS08, §5.4], but computing geodesics on the ellipsoid manifold \mathcal{E}^{n-1} is a research area on its own; see, e.g., [Per02, Kar13] and references therein. Instead, we use a simpler choice that consists, for $\boldsymbol{p} \in \mathcal{E}^{n-1}$ and $\boldsymbol{\xi} \in T_{\boldsymbol{p}} \mathcal{E}^{n-1}$, in defining $R_{\boldsymbol{p}}(\boldsymbol{\xi})$ to be the intersection with \mathcal{E}^{n-1} of the line segment between $\boldsymbol{p} + \boldsymbol{\xi}$ and the center $\boldsymbol{a}^{\text{ell}}$ of the ellipsoid. Specifically,

$$R_{\boldsymbol{p}}(\boldsymbol{\xi}) = \boldsymbol{a}^{\text{ell}} + \beta \left(\boldsymbol{p} + \boldsymbol{\xi} - \boldsymbol{a}^{\text{ell}} \right)$$
$$= \boldsymbol{a}^{\text{ell}} + \beta \boldsymbol{w} = \boldsymbol{q}, \tag{15}$$

where $\boldsymbol{a}^{\text{ell}}$ is given by (10) and $\beta \in \mathbb{R}$ must be chosen so as to satisfy $\boldsymbol{q} \in \mathcal{E}^{n-1}$:

$$\boldsymbol{q}^{\mathsf{T}}\boldsymbol{B}^{\mathrm{ell}}\boldsymbol{q} + \boldsymbol{q}^{\mathsf{T}}\boldsymbol{b}^{\mathrm{ell}} + c^{\mathrm{ell}} = 0, \tag{16}$$

which yields:

$$\beta^{2} \left(\boldsymbol{w}^{\mathsf{T}} \boldsymbol{B}^{\text{ell}} \boldsymbol{w} \right) + \beta \left(\boldsymbol{w}^{\mathsf{T}} (2\boldsymbol{B}^{\text{ell}} \boldsymbol{a}^{\text{ell}} + \boldsymbol{b}^{\text{ell}}) \right) + \left(\boldsymbol{a}^{\text{ell}} (\boldsymbol{B}^{\text{ell}} \boldsymbol{a}^{\text{ell}} + \boldsymbol{b}^{\text{ell}}) + c^{\text{ell}} \right) = 0 \quad (17)$$

Equation (17) defines a parabola in β . The roots of this parabola correspond to the two intersections with the ellipsoid of the line going trough a^{ell} and $(p + \xi)$ (see q and q' in Figure 4). Among these two points, the nearest to $(p + \xi)$ corresponds to the closest root to 1, which is therefore chosen for β . The fact that R_p is a retraction follows from [AM12, theorem 15].



Figure 4: Illustration of the optimization tools on the ellipsoid manifold \mathcal{E}^{n-1} . The vector $\boldsymbol{v} \in \mathbb{R}^n$ is projected onto the tangent space $T_p \mathcal{E}^{n-1}$: $\boldsymbol{\xi} = P_p(\boldsymbol{v})$. The retraction of the tangent vector $\boldsymbol{\xi}$ is then computed: $\boldsymbol{q} = R_p(\boldsymbol{\xi})$.

3.4 Restriction to Sub-Ellipsoids

As will be discussed in Section 4.3, it will prove useful at times to consider the restriction of the ellipsoid \mathcal{E}^{n-1} obtained by fixing some coordinates in $\boldsymbol{p} = (p_1, \ldots, p_n)^{\mathsf{T}}$. Let $\mathcal{C} \subset \{1, \ldots, n\}$ be the indices of these constant coordinates with $|\mathcal{C}| = n_c < n$, and let $\boldsymbol{P} = (P_1, \ldots, P_{n_c})^{\mathsf{T}}$ be the corresponding constants. The following notation is introduced: given a matrix \boldsymbol{M} , a vector \boldsymbol{v} and two sets of indices \mathcal{I}_1 and \mathcal{I}_2 , the sub-matrix $[M_{i,j}]_{i\in\mathcal{I}_1,j\in\mathcal{I}_2}$ is denoted as $\boldsymbol{M}_{\mathcal{I}_1,\mathcal{I}_2}$ and the sub-vector $[v_i]_{i\in\mathcal{I}_1}$ is denoted as $\boldsymbol{v}_{\mathcal{I}_1}$.

We are interested in the following set:

$$\mathcal{E}^{n-1}_{\mathcal{C},\boldsymbol{P}} := \{ \boldsymbol{p} \in \mathcal{E}^{n-1} : \boldsymbol{p}_{\mathcal{C}} = \boldsymbol{P} \}.$$

This set describes a Riemannian manifold which is the intersection of the ellipsoid \mathcal{E}^{n-1} with the intersection of the n_c axis-aligned hyperplanes defined by $p_{\mathcal{C}} = P$. This manifold is in fact again an ellipsoid of dimension $(n - n_c - 1)$, as depicted in Figure 5.



Figure 5: Illustration of a sub-ellipsoid for n = 3, with $C = \{1\}$ and $\mathcal{F} = \{2, 3\}$.

Letting $\mathcal{F} = \{1, ..., n\} \setminus \mathcal{C}$ be the set of free coordinates, this sub-ellipsoid is given by the following equation:

$$\mathcal{E}_{\mathcal{C},\mathbf{P}}^{n-1} := \left\{ \boldsymbol{p} \in \mathbb{R}^n : \boldsymbol{p}_{\mathcal{F}}^\top \widehat{\boldsymbol{B}^{\text{ell}}} \boldsymbol{p}_{\mathcal{F}} + \boldsymbol{p}_{\mathcal{F}}^\top \widehat{\boldsymbol{b}^{\text{ell}}} + \widehat{\boldsymbol{c}^{\text{ell}}} = 0, \ \boldsymbol{p}_{\mathcal{C}} = \boldsymbol{P} \right\},$$
(18)

where

$$\begin{split} \mathbf{B}^{\text{ell}} &= \mathbf{B}^{\text{ell}}_{\mathcal{F},\mathcal{F}}, \\ \widehat{\mathbf{b}^{\text{ell}}} &= \left(\mathbf{b}^{\text{ell}}_{\mathcal{F}} + 2\mathbf{B}^{\text{ell}}_{\mathcal{F},\mathcal{C}}\mathbf{p}_{\mathcal{C}}\right), \\ \widehat{c^{\text{ell}}} &= \mathbf{p}^{\top}_{\mathcal{C}} \left(\mathbf{B}^{\text{ell}}_{\mathcal{C},\mathcal{C}}\mathbf{p}_{\mathcal{C}} + \mathbf{b}^{\text{ell}}_{\mathcal{C}}\right) + c^{\text{ell}} \end{split}$$

The center of this new ellipsoid is then given by $\hat{a}_{\mathcal{F}}^{\text{ell}} = -\frac{1}{2} (\widehat{\boldsymbol{B}^{\text{ell}}})^{-1} \widehat{\boldsymbol{b}^{\text{ell}}}$ and $\hat{a}_{\mathcal{C}} = \boldsymbol{P}$.

The ingredients presented in the previous section can now be adapted to the sub-ellipsoid when needed. In particular, when computing the retraction $\boldsymbol{q} = R_{\boldsymbol{p}}(\boldsymbol{\xi})$ of a vector $\boldsymbol{\xi} \in T_{\boldsymbol{p}} \mathcal{E}_{\mathcal{C},\boldsymbol{pc}}^{n-1}$, one expects \boldsymbol{q} to remain on $\mathcal{E}_{\mathcal{C},\boldsymbol{pc}}^{n-1}$. With the definition (15) of R, this is not the case since the scaling factor β affects all of the coordinates of \boldsymbol{q} . However, the following retraction \widehat{R} (where \mathcal{C} is omitted in the notation as it can be deduced from the argument $\boldsymbol{\xi}$, see (44) for the robustified version) can be used to ensure the desired property:

$$\widehat{R}_{\boldsymbol{p}}(\boldsymbol{\xi}) = \boldsymbol{q}, \text{ with } \begin{cases} \boldsymbol{q}_{\mathcal{C}} = \boldsymbol{p}_{\mathcal{C}}, \\ \boldsymbol{q}_{\mathcal{F}} = \hat{\boldsymbol{a}} + \beta \left(\boldsymbol{p}_{\mathcal{F}} + \boldsymbol{\xi}_{\mathcal{F}} - \hat{\boldsymbol{a}} \right), \end{cases}$$
(19)

where β is computed using (17) with the corresponding $\widehat{B^{\text{ell}}}$, $\widehat{b^{\text{ell}}}$, $\widehat{c^{\text{ell}}}$ and $\widehat{a^{\text{ell}}}$. The difference between the two retractions R and \hat{R} is illustrated in Figure 6.



Figure 6: Retraction of the tangent vector $\boldsymbol{\xi}$ at the point \boldsymbol{p} . When using the retraction R given by (15), the point $\boldsymbol{q'}$ is obtained, outside of the subellipsoid going through \boldsymbol{p} . When using the retraction \hat{R} (19) relative to the center of this subellipsoid, the point \boldsymbol{q} is obtained, belonging to the subellipsoid.

Noting that $\hat{R}_{p}(\boldsymbol{\xi}) = R_{p}(\boldsymbol{\xi})$ when $\mathcal{C} = \emptyset$, equation (19) will be used whenever a retraction is needed in the remainder of this paper.

3.5 Respecting the Bound Constraints

So far in this section, we have considered performing optimization on \mathcal{E}^{n-1} . However, as presented in Section 2.2, the feasible set Ω is only a portion of \mathcal{E}^{n-1} , delimited by the bound constraints (2). In view of the sought property (i) mentioned in the introduction, we want to produce iterates that remain within these bounds at all times. To this end, a possible approach would be to penalize the cost function $f_{\rm T}$ with a log barrier term:

$$\widehat{f_{\mathrm{T}}}(\boldsymbol{p}) = \begin{cases} f_{\mathrm{T}}(\boldsymbol{p}) + \mu_k \phi(\boldsymbol{p}) & \text{if } \boldsymbol{p}^{\min} \leq \boldsymbol{p} \leq \boldsymbol{p}^{\max}, \\ +\infty & \text{otherwise,} \end{cases}$$
with $\phi(\boldsymbol{p}) = \ln(\boldsymbol{p} - \boldsymbol{p}^{\min}) + \ln(\boldsymbol{p}^{\max} - \boldsymbol{p}),$

where $\mu_k \downarrow 0$ when $k \to \infty$. This technique aims at repelling iterates from the boundaries of Ω to avoid premature convergence to nonstationary points. Other techniques to handle the bound

constrains include penalty terms and augmented Lagrangian (see, e.g., chapter 17 in [NW06]), but they do not ensure feasibility of the iterates.

We choose to adopt an alternative to these approaches. The principle consists of computing the steepest admissible direction d^k at the current iterate p^k , and then setting $p^{k+1} = \hat{R}_p(\alpha^k d^k)$ where α^k is selected by means of a line-search technique. In this context, by admissible direction, we mean the following.

Definition 3.1 An admissible direction at $\mathbf{p} \in \Omega$ is a vector $\mathbf{d} \in T_{\mathbf{p}} \mathcal{E}^{n-1}$ for which there exists $\epsilon > 0$ such that $\widehat{R}_{\mathbf{p}}(\alpha \mathbf{d})$ belongs to Ω for all $\alpha \in [0, \epsilon]$.

The computation of the steepest admissible direction will be presented in Section 4 and the linesearch technique will be discussed in Section 5.3

We already mention that, in order to be functional in combination with the chosen basic Armijotype line-search technique, the proposed approach requires a robustification mechanism, described in Section 5.2, both for the capacity constraints (2) and for the nondifferentiable points (7). Otherwise, since the nondifferentiable points and boundary points form a zero-measure set, all the iterates would be likely to be considered differentiable points that strictly satisfy the bound constraints (2), with the consequence that the computed steepest admissible direction would merely be the projected gradient (20) at all times, yielding unsatisfactory convergence as illustrated by the squares in Figure 9.

4 Subgradient Descent for the ELDP

This section chiefly concerns the computation of the steepest admissible direction for the ELDP (5). We will proceed gradually. The concept of subgradient is first recalled in an unconstrained setting in Section 4.1, then it is generalized to the Riemannian setting in Section 4.2 with a view towards considering the ELDP cost function on the ellipsoid manifold \mathcal{E}^{n-1} , which is done in Section 4.3, before reintroducing the bound constraints in Section 4.4.

As stated in Section 2.3, the cost function $f_{\rm T}(\mathbf{p})$ is piecewise smooth, i.e., differentiable almost everywhere. Let S, resp. D, denote the set of nondifferentiable points, resp. differentiable points, of $f_{\rm T}$ in Ω :

$$S = \left\{ \boldsymbol{p} \in \Omega \ : \ \exists i \text{ s.t. } p_i = p_i^{min} + \frac{k\pi}{e_i}, k \in \mathbb{N} \right\}, \quad D = \Omega \setminus S.$$

We will also consider the set S^o , resp. D^o , of points of S, resp. D, that are not on the boundary of the box (2). Note that in ELDP instances, one must expect that both S^o and $S \setminus S^o$ are nonempty, and likewise for D^o and $D \setminus D^o$.

For $p \in D^o$, the steepest admissible direction $\operatorname{grad} f_{\mathrm{T}}(p)$ is simply the projection of the gradient of $f_{\mathrm{T}}(p)$ onto $T_p \mathcal{E}^{n-1}$:

$$\operatorname{grad} f_{\mathrm{T}}(\boldsymbol{p}) = P_{\boldsymbol{p}}\left(\nabla f_{\mathrm{T}}(\boldsymbol{p})\right),\tag{20}$$

where $P_{\mathbf{p}}$ is given by (14) and $\nabla f_{\mathrm{T}}(\mathbf{p})$ by (8). For all other points $\mathbf{q} \in \Omega \setminus D^{o}$, describing the steepest admissible direction requires more sophisticated developments that we present in the rest of this section.

4.1 Subgradient Descent

Given a nonsmooth, nonconvex function $f : \mathbb{R}^n \to \mathbb{R}$, Clarke's generalized directional derivative is given by:

$$f^{\circ}(\boldsymbol{x}; \boldsymbol{d}) = \limsup_{\boldsymbol{y} \to \boldsymbol{x}, \ t \downarrow 0} \frac{f(\boldsymbol{y} + t\boldsymbol{d}) - f(\boldsymbol{y})}{t}.$$

The generalized gradient of f, noted $\partial f(\mathbf{x})$, is then defined as the set of the subgradients \mathbf{s} of $f^{\circ}(\mathbf{x}; \cdot)$:

$$\partial f(\boldsymbol{x}) = \left\{ \boldsymbol{s} \in \mathbb{R}^n : f^{\circ}(\boldsymbol{x}; \boldsymbol{v}) \geq \boldsymbol{v}^{\top} \boldsymbol{s}, \forall \boldsymbol{v} \in \mathbb{R}^n \right\}.$$

For a general nonsmooth function, this set can be difficult to describe in practice. However, if the function is locally Lipschitz, which is the case of the ELDP cost function (1), then the generalized gradient at \boldsymbol{x} is the convex hull of all points \boldsymbol{s} of the form

$$\boldsymbol{s} = \lim_{i \to \infty} \nabla f(\boldsymbol{x}_i),$$

where $\{x_i\}$ is a sequence converging to x such that f is differentiable at each x_i (see [Cla76, proposition 5]). If moreover, as is the case of the ELDP cost function (1), $f : \mathbb{R}^n \to \mathbb{R}$ can be expressed as the pointwise maximum of m smooth functions $f_j : \mathbb{R}^n \to \mathbb{R}$, i.e.,

$$f(\boldsymbol{x}) = \max_{j=1,\dots,m} f_j(\boldsymbol{x}),\tag{21}$$

then the generalized gradient can be simply described as

$$\partial f(\boldsymbol{x}) = \operatorname{co} \left\{ \nabla f_j(\boldsymbol{x}) \mid j \in \mathcal{I}_f(\boldsymbol{x}) \right\},$$
(22)

where $co\{\cdot\}$ denotes the convex hull and $\mathcal{I}_f(\boldsymbol{x})$ denotes the set of indices for which the maximum in (21) is attained, i.e.,

$$\mathcal{I}_f(\boldsymbol{x}) = \{ j \in \{1, \dots, m\} \mid f(\boldsymbol{x}) = f_j(\boldsymbol{x}) \}.$$
(23)

Using this framework, and without considering any constraint for the time being, the steepest descent direction d^k is given by the opposite of the shortest vector in $\partial f(x^k)$ [BLO05, lemma 2.1], which can be obtained by solving the following quadratic programming problem:

$$\min_{\substack{\lambda_j \ge 0\\ \sum \lambda_j = 1}} \left\| \sum_{j \in \mathcal{I}_f(\boldsymbol{x}^k)} \lambda_j \nabla f_j(\boldsymbol{x}^k) \right\|^2.$$
(24)

In practice, when the set $\partial f(\boldsymbol{x}^k)$ is composed of a single element $\nabla f_j(\boldsymbol{x}^k)$, the subproblem becomes trivial as the descent direction is simply given by $\boldsymbol{d}^k = \nabla f_j(\boldsymbol{x}^k)$. When $\partial f(\boldsymbol{x}^k)$ contains at least two elements, solving the subproblem allows to compute a descent direction even though the function is not differentiable.

4.2 Riemannian subgradient descent

In order to extend the subgradient descent to a Riemannian manifold setting, one has to redefine the needed ingredients. Let \mathcal{M} be a submanifold of \mathbb{R}^n and $f: \mathbb{R}^n \to \mathbb{R}$ a piecewise smooth function, defined as the maximum of m smooth functions f_j . Using the tools presented in Section 3.2, the projected gradient of each function f_j at a point $\boldsymbol{x} \in \mathcal{M}$ is defined as follows:

grad
$$f_j(\boldsymbol{x}) = P_{\boldsymbol{x}}(\nabla f_j(\boldsymbol{x})) \in T_{\boldsymbol{x}}\mathcal{M}.$$

The generalized projected gradient is then given by:

grad
$$f(\boldsymbol{x}) = \operatorname{co} \left\{ \operatorname{grad} f_j(\boldsymbol{x}) \mid j \in \mathcal{I}_f(\boldsymbol{x}) \right\}$$

= $\operatorname{co} \left\{ P_{\boldsymbol{x}}(\nabla f_j(\boldsymbol{x})) \mid j \in \mathcal{I}_f(\boldsymbol{x}) \right\}.$ (25)

The steepest admissible direction d^k is obtained by computing the shortest vector in grad $f(x^k)$, solving the following quadratic program:

$$\min_{\substack{\lambda_j \ge 0\\ \sum \lambda_j = 1}} \left\| \sum_{j \in \mathcal{I}(\boldsymbol{x}^k)} \lambda_j \operatorname{grad} f_j(\boldsymbol{x}^k) \right\|^2$$

$$= \min_{\substack{\lambda_j \ge 0\\ \sum \lambda_j = 1}} \left\| \sum_{j \in \mathcal{I}(\boldsymbol{x}^k)} \lambda_j P_{\boldsymbol{x}^k}(\nabla f_j(\boldsymbol{x}^k)) \right\|^2$$

$$= \min_{\substack{\lambda_j \ge 0\\ \sum \lambda_j = 1}} \left\| P_{\boldsymbol{x}^k} \left(\sum_{j \in \mathcal{I}_f(\boldsymbol{x}^k)} \lambda_j \nabla f_j(\boldsymbol{x}^k) \right) \right\|^2.$$
(26)

The steepest-descent admissible direction is computed as follows:

$$\boldsymbol{d}^{k} = -P_{\boldsymbol{x}^{k}}\left(\sum_{j\in\mathcal{I}_{f}(\boldsymbol{x}^{k})}\lambda_{j}^{*}\nabla f_{j}(\boldsymbol{x}^{k})\right),$$

where λ^* is the solution of (26).

4.3 Application to the ELDP cost function on \mathcal{E}^{n-1}

In this subsection, we adapt the subgradient descent scheme to the ELDP, temporarily ignoring the bound constraints (2) until Section 4.4. We first show that the ELDP cost function $f_{\rm T}$ (1) can be expressed as the pointwise maximum of smooth functions, then we compute the associated generalized gradient and we formulate the search for a descent direction as a quadratic programming problem.

Recalling that the ELDP cost function $f_{\rm T}(\mathbf{p})$ is the sum of n functions $f^i(p_i)$ (i = 1, ..., n), let $\hat{f}^i(\mathbf{p}) \doteq f^i(p_i)$. Each of these functions can be decomposed into two components, $\hat{f}^i_{\rm q}(\mathbf{p}) = a_i p_i^2 + b_i p_i + c_i$ (the quadratic term) and $\hat{f}^i_{\rm s}(\mathbf{p}) = d_i \sin \left[e_i \left(p_i^{\rm min} - p_i\right)\right]$ (the sinusoidal term). The ELDP cost function can then be written as follows:

$$f_{\rm T}(\boldsymbol{p}) = \sum_{i=1}^{n} \hat{f}_{\rm q}^{i}(\boldsymbol{p}) + \left| \hat{f}_{\rm s}^{i}(\boldsymbol{p}) \right|$$

= $\sum_{i=1}^{n} \max\left(\hat{f}_{\rm q}^{i}(\boldsymbol{p}) - \hat{f}_{\rm s}^{i}(\boldsymbol{p}), \ \hat{f}_{\rm q}^{i}(\boldsymbol{p}) + \hat{f}_{\rm s}^{i}(\boldsymbol{p}) \right).$ (27)

Clearly, $f_{\rm T}$ can be expressed as the maximum of 2^n functions $f_{{\rm T},j}$, taking all the possible combinations of the maximum arguments in (27):

$$\begin{split} f_{\mathrm{T}}(\boldsymbol{p}) &= \max_{j=1,...,2^{n}} f_{\mathrm{T},j}(\boldsymbol{p}) \\ &= \max_{j=1,...,2^{n}} \left(\sum_{i=1}^{n} \hat{f}_{\mathrm{q}}^{i}(\boldsymbol{p}) + \sum_{i=1}^{n} (-1)^{\left\lfloor \frac{j\%2^{i}}{2^{i-1}} \right\rfloor} \hat{f}_{\mathrm{s}}^{i}(\boldsymbol{p}) \right) \\ &= f_{\mathrm{Q}}(\boldsymbol{p}) + \max_{j=1,...,2^{n}} f_{\mathrm{S},j}(\boldsymbol{p}), \end{split}$$

where a%b denotes the the remainder of the division a/b.

The set of indices $\mathcal{I}_{f_{\mathrm{T}}}(\mathbf{p}) = \{j \in \{1, \ldots, 2^n\} \mid f_{\mathrm{T}}(\mathbf{p}) = f_{\mathrm{T},j}(\mathbf{p})\}$ can now be described. Let $\mathcal{S}(\mathbf{p})$ be the set of indices of the entries p_i in \mathbf{p} that cancel the sine term, and let $\mathcal{F}(\mathbf{p})$ be the other indices of \mathbf{p} :

$$\mathcal{S}(\boldsymbol{p}) = \{s_1, \dots, s_{n_s}\}$$
$$= \left\{ i \in \{1, \dots, n\} \mid p_i = p_i^{min} + \frac{k\pi}{e_i}, \ k \in \mathbb{N} \right\},$$
$$\mathcal{F}(\boldsymbol{p}) = \{f_1, \dots, f_{n_s}\}$$
(28)

$$(\mathbf{p}) = \{J_1, \dots, J_{n_f}\}$$

= $\{1, \dots, n\} \setminus \mathcal{S}(\mathbf{p}).$ (29)

As depicted in Figure 7, the cardinality of $I_{f_{\mathrm{T}}}(\boldsymbol{p})$ is then given by $|I_{f_{\mathrm{T}}}(\boldsymbol{p})| = 2^{n_s}$. If $\mathcal{S}(\boldsymbol{p})$ is empty, only one function $f_{\mathrm{T},j}$ has to be considered, thus $|I_{f_{\mathrm{T}}}(\boldsymbol{p})| = 2^0 = 1$.

As per (25), the generalized projected gradient of $f_{\rm T}$ at p is now defined as follows:

grad
$$f_{\mathrm{T}}(\boldsymbol{p}) = \operatorname{co} \left\{ \operatorname{grad} f_{\mathrm{T},j}(\boldsymbol{p}) \mid j \in I_{f_{\mathrm{T}}}(\boldsymbol{p}) \right\}$$

= $P_{\boldsymbol{p}}(\nabla f_{\mathrm{Q}}(\boldsymbol{p})) + \operatorname{co} \left\{ P_{\boldsymbol{p}}(\nabla f_{\mathrm{S},j}(\boldsymbol{p})) \mid j \in I_{f_{\mathrm{T}}}(\boldsymbol{p}) \right\}.$ (30)



Figure 7: The ELDP cost function $f_{T}(\mathbf{p})$ as the pointwise maximum of 2^{n} functions $f_{T,j}(\mathbf{p})$.

We now consider the quadratic subproblem (26) for the ELDP to determine the descent direction d at a point $p \in \mathcal{E}^{n-1}$:

$$\min_{\substack{\lambda_j \ge 0\\ \sum \lambda_j = 1}} \left\| P_{\boldsymbol{p}} \left(\sum_{j \in \mathcal{I}_{f_{\mathrm{T}}}(\boldsymbol{p})} \lambda_j \nabla f_{\mathrm{T},j}(\boldsymbol{p}) \right) \right\|^2.$$
(31)

This is a quadratic program with 2^{n_s} variables λ_j . However, as illustrated in Figure 8, the number of variables can be drastically reduced since the generalized projected gradient grad $f_{\rm T}(\mathbf{p})$ can actually be described with only n_s vectors. Indeed, noting that $\left|\hat{f}_{\rm s}^i(\mathbf{p})\right|$ is differentiable for $i \in \mathcal{F}(\mathbf{p})$ and recalling that $f_{\rm T}$ is a separable function, grad $f_{\rm T}(\mathbf{p})$ can be advantageously rewritten in the following way:

$$\operatorname{grad} f_{\mathrm{T}}(\boldsymbol{p}) = P_{\boldsymbol{p}}(\nabla f_{\mathrm{Q}}(\boldsymbol{p})) + \sum_{i \in \mathcal{F}(\boldsymbol{p})} P_{\boldsymbol{p}}(\nabla \left| \hat{f}_{\mathrm{s}}^{i}(\boldsymbol{p}) \right|) + \left\{ \sum_{i \in \mathcal{S}(\boldsymbol{p})} \lambda_{i} P_{\boldsymbol{p}}(\nabla \hat{f}_{\mathrm{s}}^{i}(\boldsymbol{p})) \mid \lambda_{i} \in [-1, 1] \right\}. \quad (32)$$

The quadratic subproblem (31) hence becomes:

$$\min_{\lambda_i \in [-1,1]} \left\| P_{\boldsymbol{p}} \left(\nabla f_{\mathbf{Q}}(\boldsymbol{p}) + \sum_{i \in \mathcal{F}(\boldsymbol{p})} \nabla \left| \hat{f}_s^i(\boldsymbol{p}) \right| + \sum_{i \in \mathcal{S}(\boldsymbol{p})} \lambda_j \nabla \hat{f}_s^i(\boldsymbol{p}) \right) \right\|^2.$$
(33)

Introducing the following notation:

$$\boldsymbol{S} = \left[P_{\boldsymbol{p}}(\nabla \hat{f}_{\mathrm{s}}^{s_{1}}(\boldsymbol{p})) \cdots P_{\boldsymbol{p}}(\nabla \hat{f}_{\mathrm{s}}^{s_{n_{s}}}(\boldsymbol{p})) \right], \qquad (34)$$

$$\boldsymbol{g} = P_{\boldsymbol{p}}\left(\nabla f_{\mathrm{Q}}(\boldsymbol{p}) + \sum_{i \in \mathcal{F}(\boldsymbol{p})} \nabla \left| \hat{f}_{s}^{i}(\boldsymbol{p}) \right| \right),$$
(35)

the quadratic subproblem can be rewritten as a general bound-constrained convex quadratic pro-



Figure 8: Illustration of the generalized gradient (dashed area). The description of grad $f_{\rm T}(\mathbf{p})$ (30) using 2^{n_s} vectors and the description (32) using n_s vectors are equivalent.

gram of the form:

$$\lambda^* = \arg\min_{\lambda} \|\boldsymbol{g} + \boldsymbol{S}\boldsymbol{\lambda}\|^2$$

$$= \arg\min_{\lambda} \boldsymbol{\lambda}^{\mathsf{T}} \boldsymbol{S}^{\mathsf{T}} \boldsymbol{S}\boldsymbol{\lambda} + 2\boldsymbol{\lambda}^{\mathsf{T}} \boldsymbol{S}^{\mathsf{T}} \boldsymbol{g}$$
subject to $-1 \leq \boldsymbol{\lambda} \leq 1$, (36)

where $S^{\top}S$ is positive semidefinite (positive definite if S has full rank). The descent direction d is then computed as $d = -(g + S\lambda^*)$.

4.4 Including the Bound Constraints

Solving the subproblem (36) does not always provide an admissible descent direction for points on the boundary of the feasible set Ω . Indeed, if the resulting direction d points outside Ω , no progress can be made by following it. However, the bound constraints can be incorporated into the quadratic problem in order to ensure an admissible descent direction for all points in Ω . Let $l_i(\mathbf{p})$ and $u_i(\mathbf{p})$, $i = 1, \ldots, n$, denote the lower and upper bound constraints:

$$l_i(\boldsymbol{p}) = -p_i + p_i^{\min} \le 0$$
$$u_i(\boldsymbol{p}) = p_i - p_i^{\max} \le 0.$$

Let $\mathcal{L}(p)$ and $\mathcal{U}(p)$ be the indices of the active constraints, and let $\mathcal{B}(p)$ be their union:

$$\mathcal{L}(\boldsymbol{p}) = \{l_1, \dots, l_{n_l}\} \\ = \{i \in \{1, \dots, n\} \mid l_i(\boldsymbol{p}) = 0\}, \\ \mathcal{U}(\boldsymbol{p}) = \{u_1, \dots, u_{n_u}\} \\ = \{i \in \{1, \dots, n\} \mid u_i(\boldsymbol{p}) = 0\}, \\ \mathcal{B}(\boldsymbol{p}) = \mathcal{L}(\boldsymbol{p}) \big| \quad \big| \mathcal{U}(\boldsymbol{p}).$$
(37)

Noting that $\nabla l_i(\mathbf{p}) = -\mathbf{e}_i$ and $\nabla u_i(\mathbf{p}) = \mathbf{e}_i$, we introduce the following notation:

$$\boldsymbol{B} = \left[P_{\boldsymbol{p}}(-\boldsymbol{e}_{l_1}) \cdots P_{\boldsymbol{p}}(-\boldsymbol{e}_{l_{n_l}}) \mid P_{\boldsymbol{p}}(\boldsymbol{e}_{u_1}) \cdots P_{\boldsymbol{p}}(\boldsymbol{e}_{u_{n_u}}) \right].$$
(38)

The search for an admissible descent direction can now include the projected gradient of the active constraints an can thus be expressed as the following quadratic program:

$$(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) = \arg\min_{\boldsymbol{\lambda}, \boldsymbol{\mu}} \|\boldsymbol{g} + \boldsymbol{S}\boldsymbol{\lambda} + \boldsymbol{B}\boldsymbol{\mu}\|^2$$
subject to
$$\begin{cases}
-1 \le \boldsymbol{\lambda} \le 1, \\ \boldsymbol{\mu} \ge 0.
\end{cases}$$
(39)

This is again a convex quadratic problem of dimension $|\mathcal{S}(p) \bigcup \mathcal{B}(p)| \leq n$. The steepest-descent admissible direction is finally computed as follows:

$$\boldsymbol{d} = -(\boldsymbol{g} + \boldsymbol{S}\boldsymbol{\lambda}^* + \boldsymbol{B}\boldsymbol{\mu}^*). \tag{40}$$

4.5 First-Order Stationarity Condition

Let $f : \mathbb{R}^n \to \mathbb{R}$ be a nonsmooth (Lipschitz) function to be minimized subject to a set of constraints $c_i(\boldsymbol{x}) \leq 0$ (i = 1, ..., m). The first order necessary stationarity conditions for this constrained nonsmooth problem can be stated as follows (see *e.g.* [Cla76], Section 3):

Theorem 4.1 If x^* minimizes f(x) locally, then there exist numbers μ_i (i = 1, ..., m) such that:

- (a) $\mu_i \ge 0$,
- (b) $\mu_i c_i(x^*) = 0$,
- (c) $\mathbf{0} \in \partial f(\mathbf{x}^*) + \sum_{i=1}^m \mu_i \nabla c_i(\mathbf{x}^*).$

Applying this theorem to the ELDP, we obtain the following characterization of its local minimizers:

Theorem 4.2 If p^* minimizes $f_T(p)$ locally in Ω , then there exist numbers μ_i^- , μ_i^+ (i = 1, ..., n) such that:

(a)
$$\mu_i^- \ge 0, \ \mu_i^+ \ge 0$$

(b)
$$\mu_i^- l_i(\mathbf{p}^*) = 0, \ \mu_i^+ u_i(\mathbf{p}^*) = 0,$$

(c) $\mathbf{0} \in \operatorname{grad} f_{\mathrm{T}}(\boldsymbol{p}^*)$ + $\sum_{i=1}^n \left(\mu_i^- \operatorname{grad}(l_i(\boldsymbol{p}^*)) + \mu_i^+ \operatorname{grad}(u_i(\boldsymbol{p}^*)) \right).$

The technique proposed in the previous section to compute admissible descent directions offers an important advantage over the classical alternatives (barrier term, augmented Lagrangian, etc.). Computing the direction $d = -(g + S\lambda^* + B\mu^*)$ at a local minimizer point p^* yields the zero vector. Therefore, the norm of the direction d can be used as a stopping criterion for the descent algorithm, ensuring that it will stop at a (numerically) stationary point.

5 Implementation

In the previous sections, all the ingredients needed to perform local optimization of $f_{\rm T}(\mathbf{p})$ in Ω were introduced. In this section, we present how, given a feasible initial iterate $\mathbf{p}^0 \in \Omega$, the Riemannian subgradient descent method is applied to produce a sequence of iterates $\mathbf{p}^k \in \Omega$. The implementation details are now given.

5.1 Generating a feasible iterate

Recalling the geometrical considerations presented in Section 2.2, the feasible set Ω is the intersection of an axis-aligned box and an ellipsoid surface. From there, a simple approach to generate a feasible point $\mathbf{p} \in \Omega$ is given by Algorithm 1. The loop is needed because the scaling performed by R may output a point outside the bounds constraints. Note that Algorithm 1 becomes superfluous if the proposed local method is combined with a feasible global method (i.e., with iterates on Ω).

Algorithm 1 Feasible point generation Output: $p \in \Omega$, a feasible point repeat Draw x from the uniform distribution on $[p_1^{\min}, p_1^{\max}] \times \cdots \times [p_n^{\min}, p_n^{\max}]$ $p \leftarrow R_x(\mathbf{0})$, where R is given by formula (15) until $p \in \Omega$ return p

5.2 Computing the descent direction

The steepest-descent admissible direction at a point \mathbf{p}^k is described by (40). It can be computed, provided that the sets $\mathcal{S}(\mathbf{p}^k)$ and $\mathcal{B}(\mathbf{p}^k)$ are available. However, in practice, it is numerically unlikely for an iterate \mathbf{p}^k to ever cancel exactly any of the sine terms in $f_{\mathrm{T}}(\mathbf{p}^k)$. Therefore, $\mathcal{S}(\mathbf{p}^k)$ and its complement $\mathcal{F}(\mathbf{p}^k)$ are approximated in the following way:

$$\mathcal{S}_{\epsilon}(\boldsymbol{p}) = \left\{ i \in \{1, \dots, n\} : \left| p_i - \left(p_i^{min} + \frac{k\pi}{e_i} \right) \right| \le \epsilon \right\},\tag{41}$$

$$\mathcal{F}_{\epsilon}(\boldsymbol{p}) = \{1, \dots, n\} \setminus \mathcal{S}_{\epsilon}(\boldsymbol{p}).$$
(42)

Similarly, the bound constraints indices are approximated as follows:

$$\mathcal{B}_{\epsilon}(\boldsymbol{p}) = \left\{ i \in \{1, \dots, n\} \mid l_i(\boldsymbol{p}) \ge -\epsilon \text{ or } u_i(\boldsymbol{p}) \ge -\epsilon \right\}.$$
(43)

An " ϵ -steepest-descent ϵ -admissible direction" is then computed at each iteration, using Algorithm 2. In its nondifferentiable-related aspect, this approach is reminiscent of Goldstein's ϵ -generalized gradient [Gol77]. In its bound-related aspect, it has an active-set flavor.

Algorithm 2 DESCENTDIRECTION $(p, \mathcal{S}_{\epsilon}(p), \mathcal{F}_{\epsilon}(p), \mathcal{B}_{\epsilon}(p))$ - Descent direction

Input: $p \in \Omega$, $S_{\epsilon}(p)$, $\mathcal{F}_{\epsilon}(p)$, $\mathcal{B}_{\epsilon}(p)$ Output: d, an admissible descent direction Compute S, g and B at p according to (34), (35) and (38), in which S(p), $\mathcal{F}(p)$, $\mathcal{B}(p)$ are replaced by $S_{\epsilon}(p)$, $\mathcal{F}_{\epsilon}(p)$, $\mathcal{B}_{\epsilon}(p)$ Compute (λ^*, μ^*) solving the quadratic programming sub-problem (39) $d \leftarrow -(g + S\lambda^* + B\mu^*)$ return d

5.3 Computing the step size

Once a descent admissible direction has been computed, the step size must be determined. We propose to adopt Armijo's rule with two minor modifications. First, the initial step size is retained from the previous step size computation. Second, before performing classical backtracking, a few trials are performed by expanding the previous step size. These two modifications provide possible faster convergence of the descent algorithm by allowing bigger step sizes. The line search uses the retraction \hat{R} defined by (19). In practice, given a vector $\boldsymbol{\xi}$, the corresponding set \mathcal{C} involved in (19) is approximated as follows:

$$\mathcal{C}_{\epsilon} = \{ i \in \{1, \dots, n\} \mid |\boldsymbol{\xi}_i| \le \epsilon \}.$$
(44)

The computation of the step size is described in Algorithm 3.

Algorithm 3 STEPSIZE $(p, g, \alpha_0, \beta, \gamma, n_f, n_b)$ - Line search using Armijo's rule

Input: $p \in \Omega$, g search direction, α_0 initial step size, β sufficient decrease constant, γ step size shrinking factor, $n_{\rm f}$ number of forward steps, $n_{\rm b}$ number of backtracking steps

Output: α , a step size ensuring sufficient decrease

```
Consider \tilde{f}_{T}(\boldsymbol{p}) = \begin{cases} f_{T}(\boldsymbol{p}) & \text{if } (2), \\ +\infty & \text{otherwise.} \end{cases}

Step 1 (Forward exploration):

for s = n_{f} \rightarrow 1 do

if \tilde{f}_{T}(\hat{R}_{\boldsymbol{p}}(\gamma^{-s} \frac{\boldsymbol{g}}{\|\boldsymbol{g}\|})) < f_{T}(\boldsymbol{p}) - \beta \alpha_{0} \gamma^{-s} \|\boldsymbol{g}\| then

return \alpha_{0} \gamma^{-s}

end if

end for

Step 2 (Backtracking):

for s = 0 \rightarrow n_{b} do

if \tilde{f}_{T}(\hat{R}_{\boldsymbol{p}}(\gamma^{s} \frac{\boldsymbol{g}}{\|\boldsymbol{g}\|})) < f_{T}(\boldsymbol{p}) - \beta \alpha_{0} \gamma^{s} \|\boldsymbol{g}\| then

return \alpha_{0} \gamma^{s}

end if

end for

return 0
```

5.4 Subgradient Descent for the ELDP

The subgradient descent algorithm can now be completely described and is given by Algorithm 4. In the taxonomy of optimization methods [NW06], the proposed approached can be considered an active-set–like feasible ϵ -subgradient–descent method for nonsmooth nonconvex problems with both equality and bound constraints.

Algorithm 4 Riemannian Subgradient Descent for the ELDP

```
Input: p^0 \in \Omega, \epsilon singular point detection tolerance, \delta descent direction norm tolerance, \beta suffi-
   cient decrease constant, \gamma step size shrinking factor, n_{\rm f} number of forward steps, n_{\rm b} number of
   backtracking steps
Output: p^* \in \Omega, a generalized stationary point
   Set k = 0, \ \alpha^0 = 1
   Step 1 (Compute the admissible descent direction)
         Compute \mathcal{S}_{\epsilon}(\mathbf{p}^k), \mathcal{F}_{\epsilon}(\mathbf{p}^k) and \mathcal{B}_{\epsilon}(\mathbf{p}^k) (41)–(43)
         d^k \leftarrow \text{DescentDirection}(p^k, \mathcal{S}_{\epsilon}(p^k), \mathcal{F}_{\epsilon}(p^k), \mathcal{B}_{\epsilon}(p^k))
         if \|\boldsymbol{d}^k\| < \delta then
               return p^k
         end if
         oldsymbol{g}^k \leftarrow oldsymbol{d}^k / oldsymbol{\|} oldsymbol{d}^k oldsymbol{\|}
   Step 2 (Compute the step size)
         \alpha^k \leftarrow \text{STEPSIZE}(\boldsymbol{p}^k, \boldsymbol{g}^k, \alpha^{k-1}, \beta, \gamma, n_{\text{f}}, n_{\text{b}})
   Step 3 (Compute the new iterate)
         if \alpha^k > 0 then
              \pmb{p}^{k+1} \leftarrow \hat{R}_{\pmb{p}^k}(\alpha^k \pmb{g}^k)
               k \leftarrow k+1
               Go to Step 1
         else
               return p^k
         end if
```

6 Numerical Experiments

The Riemannian subgradient descent algorithm presented in the previous section was implemented in MATLAB. The data from four instances of ELDP with valve-point effect were collected, for dimensions n = 3, 5, 6 and 15 (sources: [Gai03], [MuAWA10], [Gai04], [SHP+10]). (We are not aware of publicly-available ELDP instances with valve-point effect and a quadratic loss function for which n is greater than 15.) These data sets contain the following information for each instance:

• the coefficients of $f_{\rm T}$ for the quadratic term:

$$a_i [\$/(h.MW^2)], b_i [\$/(h.MW)], c_i [\$/h],$$

• the coefficients of $f_{\rm T}$ for the sine term:

$$d_i [\$/h], e_i [1/MW],$$

- the bound constraints: p_i^{\min} , p_i^{\max} [MW],
- the power demand: p_d [MW],
- the transmission loss coefficients:

$$B [1/MW], B^0 [-], B^{00} [MW].$$

An important remark about these data sets is that for the cases n = 5 and n = 6, only the first two dimensions present the valve-point effect ($d_i = e_i = 0, i \ge 2$). The consequence is that the multimodal aspect for these two data sets is much less pronounced compared to the the two other data sets, for which the valve-point effect is present in all dimensions.

The MATLAB code and data sets are available online at http://www.inma.ucl.ac.be/~borckmans/ ELDP. In the following experiments, the default parameters were chosen as follows: $\epsilon = 10^{-8}$, $\beta = 10^{-4}$, $\gamma = 0.5$, $n_{\rm f} = 3$, $n_{\rm b} = 50$, $\delta = 10^{-12}$.



(a) Local minima for n = 3

Figure 9: Local convergence illustration for n = 3. The trajectories from 9 starting points (stars) are depicted. The final iterates obtained using a naive steepest-descent algorithm (squares) and the proposed Riemannian subgradient algorithm (dots) are represented. The dashed blue lines correspond to the nondifferentiable sub-ellipsoids.

6.1 Local Convergence

The trajectories generated by Algorithm 4 for different starting points are illustrated in Figure 9 for the case n = 3. Whenever the sequence of iterates approaches a nondifferentiable point, the projected subgradients are computed, and the descent direction resulting from the quadratic subproblem follows the corresponding sub-ellipsoid. These sub-ellipsoids are represented as dashed lines in Figure 9. For comparison purpose, applying a simple Riemannian gradient descent (ignoring the nondifferentiable aspect of $f_{\rm T}$) from the same starting points yields convergence to nonstationary points.

Algorithm 4 was tested by performing a large number of trials for each ELDP instance, starting from random points. The final iterates of the trials were tested for the necessary stationarity condition; the norm of the final descent direction, as computed by Algorithm 2 was smaller than 10^{-12} for all the trials. Considering two points to be different if their corresponding costs are at least 10^{-8} apart from each other, the proposed algorithm was able to consistently converge to a small set of potential local minima for the cases n = 3, 5 and 6 (9 minima for n = 3, 7 minima for n = 5, and 7 minima for n = 6). As mentioned earlier, the small number of local minima for the cases n = 5 and n = 6 is due to the limited valve-point effect. On the other hand, for the case n = 15, the number of local minima is much higher since the valve-point effect affects all the dimensions. After 10^6 trials, the proposed approach consistently identified solutions in a set of around 2×10^5 points.



Figure 10: Local convergence properties for n = 3.



Figure 11: Influence of ϵ (n = 15)

6.2 Parameter Influence

As depicted in Figure 10(a) and 10(b), the convergence speed is affected by the line search scaling factor γ and is enhanced when reusing the previous successful step size and allowing forward exploration ($n_{\rm f}$ steps).

Another important parameter of the proposed algorithm is ϵ , the tolerance for detecting the singular entries in (41). Smaller values for ϵ lead to better precision but at the cost of taking more function evaluations to detect singularities. This can be observed in Figure 11, where the plateaus correspond to iterates approaching singular sub-ellipsoids of Ω .



Figure 12: Using the proposed subgradient descent (SD) algorithm as a post processing tool. Stopping the DE algorithm progression at some point and refining the output using the proposed SD algorithm (red curve) provides faster convergence than DE alone (blue curve).

6.3 Global Exploration

The Riemannian subgradient descent presented in this paper is inherently a local optimizer. It is very efficient at computing a local minimizer in the vicinity of a given starting point in Ω , but it is not meant to roam around the cost function landscape in search of the global minimum. On the other hand, the global heuristics that are largely explored in the literature have very complementary properties: they often provide good exploration of the search space but they tend to struggle to refine the final iterate and they do not attempt to check stationarity (and even feasibility in some cases) of their output. This calls for hybrid algorithms that combine the best of both approaches.

In order to illustrate this idea, we implemented the Differential Evolution (DE) algorithm [SP97] as proposed in [BC10]. This population-based heuristic is known to provide reasonably good exploration of the search space, relying quite heavily on random updates, combining selected individuals. As in many heuristics proposed to solve the ELDP, the authors of [BC10] use the concept of a slack variable to respect the equality constraint (3).

A first approach to hybridize DE with our algorithm, known as sequential approach [Sel11], is to stop the global exploration at some point before applying the subgradient descent (Algorithm 4). This post-processing step avoids the slow refinement of the solution and produces in practice a final iterate that satisfies the stationary conditions from Theorem 4.2. The typical behavior of such an approach is depicted in Figure 12.

Another option is to integrate the subgradient descent within the global technique, following the cyclical approach mentioned in [Sel11]. The idea is to refine the best individual produced by the DE algorithm at each iteration using the projected subgradient descent. The benefit of this approach is illustrated in Figure 13.



Figure 13: Benefit of hybridizing DE with the subgradient descent (SD) algorithm. In this cyclical hybrid approach, DE is used to escape local minima (green curve), then the proposed SD algorithm is applied to refine the best individual (red curve). The outcome is a faster convergence than when the DE technique is applied alone (blue curve).

7 Conclusion

In this paper, we showed how the geometrical structure of the ELDP (5) can be exploited to perform local optimization of its associated cost function. The canvas of optimization on Riemannian manifolds was used to maintain feasible iterates at all times. The nondifferentiable nature of the cost function was handled using Clarke's generalized calculus. A simple procedure was presented to compute an admissible descent direction and to perform a curvilinear search along that direction while satisfying the constraints. The resulting local method is unique in that it possesses the combination of desirable properties mentioned in the introduction. Finally, we showed that the multimodal aspect of the ELDP can be addressed by hybridizing the proposed local method with a global optimization technique.

Acknowledgements

The authors would like to thank Arnaud Browet and Samuel Melchior for their valuable comments and suggestions that came out of many insightful conversations.

References

- [ADM⁺02] Roy L. Adler, Jean-Pierre Dedieu, Joseph Y. Margulies, Marco Martens, and Mike Shub. Newton's method on Riemannian manifolds and a geometric model for the human spine. *IMA J. Numer. Anal.*, 22(3):359–390, July 2002. doi: 10.1093/imanum/22.3.359.
- [AKTH02] P. Attaviriyanupap, H. Kita, E. Tanaka, and J. Hasegawa. A hybrid EP and SQP for dynamic economic dispatch with nonsmooth fuel cost function. *Power Engineering Review, IEEE,* 22(4):77, april 2002. doi:10.1109/MPER.2002.4312139.
- [AM12] P.-A. Absil and Jérôme Malick. Projection-like retractions on matrix manifolds. SIAM Journal on Optimization, 22(1):135–158, 2012. doi:10.1137/100802529.
- [AMS08] P.-A. Absil, R. Mahony, and R. Sepulchre. Optimization Algorithms on Matrix Manifolds. Princeton University Press, Princeton, NJ, 2008. URL: http: //www.inma.ucl.ac.be/~absil/amsbook/.
- [ASAO10] J. S. Alsumait, J. K. Sykulski, and A. K. Al-Othman. A hybrid GA-PS-SQP method to solve power system valve-point economic dispatch problems. *Applied Energy*, 87(5):1773 - 1781, 2010. doi:10.1016/j.apenergy.2009.10.007.
- [BC10] A. Bhattacharya and P.K. Chattopadhyay. Hybrid differential evolution with biogeography-based optimization for solution of economic load dispatch. *IEEE Trans*actions on Power Systems, 25(4):1955-1964, nov. 2010. doi:10.1109/TPWRS.2010. 2043270.
- [BLO05] James V. Burke, Adrian S. Lewis, and Michael L. Overton. A robust gradient sampling algorithm for nonsmooth, nonconvex optimization. SIAM J. Optim., 15(3):751– 779, 2005. doi:10.1137/030601296.
- [Cla75] Frank H. Clarke. Generalized gradients and applications. Transactions of the American Mathematical Society, 205:pp. 247-262, 1975. URL: http://www.jstor.org/ stable/1997202.
- [Cla76] F. H. Clarke. A new approach to Lagrange multipliers. Mathematics of Operations Research, 2:165–174, 1976.
- [CLL⁺12] Jiejin Cai, Qiong Li, Lixiang Li, Haipeng Peng, and Yixian Yang. A hybrid CPSO–SQP method for economic dispatch considering the valve-point effects. *Energy Conversion and Management*, 53(1):175 181, 2012. doi:10.1016/j.enconman.2011. 08.023.
- [DHL07] Gunther Dirr, Uwe Helmke, and Christian Lageman. Nonsmooth Riemannian optimization with applications to sphere packing and grasping. In F. Allgwer, P. Fleming, P. Kokotovic, A.B. Kurzhanski, H. Kwakernaak, A. Rantzer, J.N. Tsitsiklis, Francesco Bullo, and Kenji Fujimoto, editors, Lagrangian and Hamiltonian Methods for Nonlinear Control 2006, volume 366 of Lecture Notes in Control and Information Sciences, pages 29–45. Springer Berlin / Heidelberg, 2007. doi:10.1007/978-3-540-73890-9_2.
- [EAS98] Alan Edelman, Tomás A. Arias, and Steven T. Smith. The geometry of algorithms with orthogonality constraints. SIAM J. Matrix Anal. Appl., 20(2):303–353, 1998. doi:10.1137/S0895479895290954.
- [Gai03] Z. L. Gaing. Particle swarm optimization to solving the economic dispatch considering the generator constraints. *IEEE Transactions on Power Systems*, 18(3):1187– 1195, AUG 2003. doi:10.1109/TPWRS.2003.814889.

- [Gai04] Z. L. Gaing. Closure to discussion of particle swarm optimization to solving the economic dispatch considering the generator constraints. *IEEE Transactions on Power Systems*, 19(4):2122–2123, 2004. doi:10.1109/TPWRS.2004.831708.
- [Gol77] A. A. Goldstein. Optimization of Lipschitz continuous functions. Math. Programming, 13(1):14-22, 1977. doi:10.1007/BF01584320.
- [HM94] Uwe Helmke and John B. Moore. Optimization and Dynamical Systems. Communications and Control Engineering Series. Springer-Verlag London Ltd., London, 1994. With a foreword by R. Brockett.
- [Kar13] Charles F. F. Karney. Algorithms for geodesics. Journal of Geodesy, 87:43–55, 2013.
 doi:10.1007/s00190-012-0578-z.
- [MuAWA10] Tahir Nadeem Malik, Azzam ul Asar, Mudasser F. Wyne, and Shakil Akhtar. A new hybrid approach for the solution of nonconvex economic dispatch problem with valve-point effects. *Electric Power Systems Research*, 80(9):1128 – 1136, 2010. doi: 10.1016/j.epsr.2010.03.004.
- [NI08] Nasimul Noman and Hitoshi Iba. Differential evolution for economic load dispatch problems. *Electric Power Systems Research*, 78(8):1322 – 1331, 2008. doi:10.1016/ j.epsr.2007.11.007.
- [Nik10] Taher Niknam. A new fuzzy adaptive hybrid particle swarm optimization algorithm for non-linear, non-smooth and non-convex economic dispatch problem. *Applied Energy*, 87(1):327 - 339, 2010. doi:10.1016/j.apenergy.2009.05.016.
- [NW06] Jorge Nocedal and Stephen J. Wright. *Numerical optimization*. Springer Series in Operations Research and Financial Engineering. Springer, New York, second edition, 2006.
- [Per02] A. M. Perelomov. A note on geodesics on ellipsoid, 2002. arXiv:math-ph/0203032.
- [PLSL05] Jong-Bae Park, Ki-Song Lee, Joong-Rin Shin, and K.Y. Lee. A particle swarm optimization for economic dispatch with nonsmooth cost functions. *Power Systems*, *IEEE Transactions on*, 20(1):34 42, feb. 2005. doi:10.1109/TPWRS.2004.831275.
- [Roc03] R. T. Rockafellar. A property of piecewise smooth functions. *Comput. Optim. Appl.*, 25(1-3):247–250, March 2003. doi:10.1023/A:1022921624832.
- [SCC03] N. Sinha, R. Chakrabarti, and R. K. Chattopadhyay. Evolutionary programming techniques for economic load dispatch. *IEEE Transactions on Evolutionary Computation*, 7(1):83–94, feb 2003. doi:10.1109/TEVC.2002.806788.
- [Sel11] A. Immanuel Selvakumar. Enhanced cross-entropy method for dynamic economic dispatch with valve-point effects. *International Journal of Electrical Power & Energy* Systems, 33(3):783 – 790, 2011. doi:10.1016/j.ijepes.2011.01.001.
- [SHP⁺10] Nidul Sinha, Kaustabh Moni Hazarika, Shantanu Paul, Himanshu Shekhar, and Amrita Karmakar. Improved real quantum evolutionary algorithm for optimum economic load dispatch with non-convex loads. In Bijaya Panigrahi, Swagatam Das, Ponnuthurai Suganthan, and Subhransu Dash, editors, Swarm, Evolutionary, and Memetic Computing, volume 6466 of Lecture Notes in Computer Science, pages 689–700. Springer Berlin / Heidelberg, 2010. doi:10.1007/978-3-642-17563-3_81.
- [SP97] Rainer Storn and Kenneth Price. Differential evolution a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11:341–359, 1997. doi:10.1023/A:1008202821328.

- [VJ04] TAA Victoire and AE Jeyakumar. Hybrid PSO-SQP for economic dispatch with valve-point effect. *Electric Power Systems Research*, 71(1):51–59, SEP 2004. doi: 10.1016/j.epsr.2003.12.017.
- [VJ05] T. A. A. Victoire and A. E. Jeyakumar. Deterministically guided PSO for dynamic dispatch considering valve-point effect. *Electric Power Systems Research*, 73(3):313– 322, MAR 2005. doi:10.1016/j.epsr.2004.07.005.
- [WS93] D. C. Walters and G. B. Sheble. Genetic algorithm solution of economic dispatch with valve point loading. *Power Systems, IEEE Transactions on*, 8(3):1325-1332, aug 1993. doi:10.1109/59.260861.
- [YÖ11] Celal Yaşar and Serdar Özyön. A new hybrid approach for nonconvex economic dispatch problem with valve-point effect. *Energy*, 36(10):5838 5845, 2011. doi: 10.1016/j.energy.2011.08.041.
- [YWZY09] Xiaohui Yuan, Liang Wang, Yongchuan Zhang, and Yanbin Yuan. A hybrid differential evolution method for dynamic economic dispatch with valve-point effects. Expert Systems with Applications, 36(2, Part 2):4042 – 4048, 2009. doi: 10.1016/j.eswa.2008.03.006.
- [YYH96] Hong-Tzer Yang, Pai-Chuan Yang, and Ching-Lien Huang. Evolutionary programming based economic dispatch for units with non-smooth fuel cost functions. *Power* Systems, IEEE Transactions on, 11(1):112-118, feb 1996. doi:10.1109/59.485992.