Riemannian Optimization Applied to AC Optimal Power Flow

24PESGM114
8

Jonathan Maack[1], Devon Sigler[1], and Ariel Goodwin[2]
[1]National Renewable Energy Lab, [2]Cornell

Introduction

As transmission infrastructure ages and higher-penetrations of renewable generation become common, interest has increased in using alternating current (AC) physics in power systems design, analysis, and operation [1]. The core problem of many operations models is optimal power flow (OPF). Using AC physics with OPF results in a nonlinear, nonconvex optimization problem which is generally solved with interior point methods such as that implemented in IPOPT [2]. As the size of the power system increases, the ACOPF problem can be difficult to solve [3] due to the linear algebra problem at the core of interior point methods [4, 5].

The advent of optimization on smooth manifolds [6] suggests an alternative. The AC power flow equations form a smooth submanifold of Euclidean space called the power flow manifold [7]. We use the packages PowerModels.jl [8] and Manopt.jl [9] to demonstrate that Riemannian optimization techniques produce similar quality solutions as IPOPT as measured by the terminal objective value and maximum constraint violation.

Riemannian Manifolds

An **embedded submanifold** $\mathcal{M}$ of Euclidean space is given by

$$\mathcal{M} = \{x \in \mathbb{R}^n : q(x) = 0\}, \quad q \in C^\infty, \quad q : \mathbb{R}^n \to \mathbb{R}^m, \quad m < n$$

where Dq_x is the Jacobian of q at x and is of full rank for all $x \in \mathbb{R}^n$. We limit ourselves to manifolds of this form here.

Given a point $x \in \mathcal{M}$, the **tangent space** is defined by

$$T_x\mathcal{M} := \{v \in \mathbb{R}^n : v = \gamma'(0)\}, \quad \gamma : I \to \mathcal{M}, \quad I \subset \mathbb{R}, \quad \gamma(0) = x.$$

For a submanifold of Euclidean space, the tangent space coincides with the kernel of the Jacobian:

$$T_x\mathcal{M} = \{v \in \mathbb{R}^n : Dq_x v = 0\}.$$

That is, the tangent space is the tangent (hyper-)plane of the manifold at x.

The **tangent bundle** is the disjoint union of all tangent spaces

$$T\mathcal{M} = \{(x, v) : x \in \mathcal{M}, v \in T_x\mathcal{M}\}.$$

A **Riemannian manifold** is a smooth manifold paired with a Riemannian metric $g_x : T_x\mathcal{M} \times T_x\mathcal{M} \to \mathbb{R}$. The Riemannian metric generalizes the notion of the Euclidean inner product so that geometric notions such as angles can be given on smooth manifolds. We take our metric to be the standard Euclidean inner product

$$g_x(u, v) = \langle u, v \rangle_x := \sum_{i=1}^{n} u_i v_i, \quad u, v \in T_x\mathcal{M}.$$

The subscript emphasizes the fact that the vectors are restricted to the tangent space of the manifold at the point x.

For a function $f : \mathcal{M} \to \mathbb{R}$, the **Riemannian gradient** of f is defined as the unique vector field $\mathrm{grad} f$ on $\mathcal{M}$ such that for all $(x, v) \in T\mathcal{M}$, we have

$$Df(x)[v] = \langle \mathrm{grad} f(x), v \rangle_x$$

where Df is the differential [10] of f. For our case, we have

$$\mathrm{grad} f(x) = P_x(\nabla \hat{f}(x)), \quad P_x = I - Q_x, \quad Q_x = Dq_x(Dq_x Dq_x^T)^{-1} Dq_x^T,$$

where $\hat{f}$ is any smooth extension of f to $\mathbb{R}^n$, and ∇ denotes the standard Euclidean gradient.

Riemannian Optimization

Consider the Riemannian optimization problem

$$\min_{x \in \mathcal{M}} f(x) \quad \text{s.t.} \quad h(x) \le 0.$$

To solve this problem, we need a few additional tools.

A **retraction** is a smooth map $R : T\mathcal{M} \to \mathcal{M} : (x, v) \to R_x(v)$ such that for each curve $\gamma(t) = R_x(tv)$, we have $\gamma(0) = x$ and $\gamma'(0) = v$. A retraction ensures that new iterates of an algorithm are on the manifold. The notion of a retraction is a generalization of the exponential map [6,10,11].

A **vector transport** is a smooth linear map $\mathcal{T} : T\mathcal{M} \oplus T\mathcal{M} \to T\mathcal{M} : (u, v) \to \mathcal{T}_u(v)$ such that, for all $x \in \mathcal{M}$ and for all $u, v \in T_x\mathcal{M}$, there exists a retraction R where $\mathcal{T}_u(v) \in T_{R_x(u)}\mathcal{M}$ and $\mathcal{T}_0(v) = v$ [6]. The notion of a vector transport is a generalization of parallel transport [6,11]. A vector transport allows us to move vectors from one tangent space to another tangent space.

A generic (first-order) Riemannian optimization algorithm iterates the following three steps:

1. Determine search direction s_k. This often uses a vector transport.
2. Determine step size α_k. Since $\phi(\alpha) = f(R_{x_k}(\alpha s_k))$ is a function from the reals to the reals, standard line search techniques are directly applicable.
3. Set $x_{k+1} = R_{x_k}(\alpha_k s_k)$. The retraction $R_x(u)$ guarantees that the new iterate is on the manifold.

This generic procedure is depicted in Figure 1.

Test Setup

For each method, we require a retraction and potentially a vector transport. Our retraction is taken from [12]. The idea is to search for the manifold in a direction normal to the tangent space. This is often called the orthographic retraction in the literature. For a point $x \in \mathcal{M}$, a search direction s_k, and step size α_k, we perform the iteration

$$y_k^{\ell+1} = y_k^\ell - Dq_{x_k}^T (Dq_{x_k} Dq_{x_k}^T)^{-1} q(y_k^\ell), \quad y_k^0 = x_k + \alpha_k s_k$$

until $\|q(y_k^\ell)\|$ is sufficiently small at which point we set $x_{k+1} = y_k^\ell$.

For the embedded submanifold case, an obvious vector transport is given by taking the orthogonal projection of a vector into the necessary tangent space. To be exact, for any retraction R , we define $\mathcal{T}_u(v) = P_{R_x(u)} v$ where $x \in \mathcal{M},\ u, v \in T_x\mathcal{M}$, and $P_{R_x(u)}$ is the previously defined orthogonal projector to the tangent space.

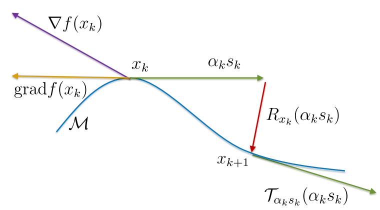

Fig. 1. Depiction of a single iteration of a first-order Riemannian optimization algorithm.

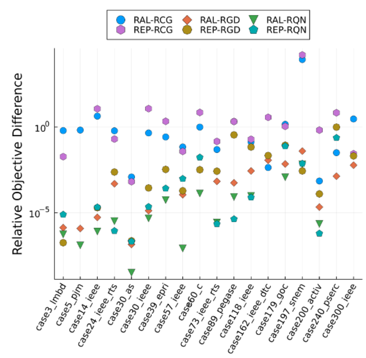

Fig. 2. Relative difference in objective function compared to IPOPT value. Values are omitted when the algorithm failed to terminate.

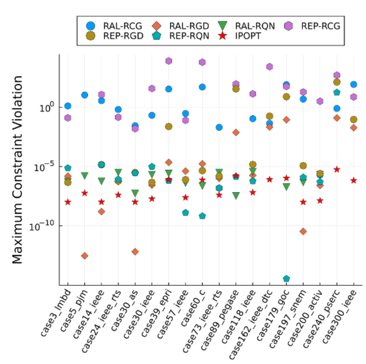

Fig. 3. Maximum constraint violation. Values are omitted when the algorithm failed to terminate.

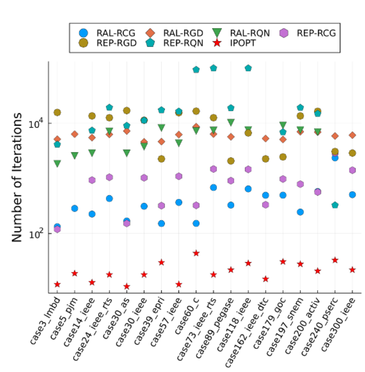

Fig. 4. Number of total inner loop iterations before termination. Values are omitted when the algorithm failed to terminate.

To handle the inequality constraints, we use the Riemannian augmented Lagrangian (RAL) method and the Riemannian exact penalty (REP) method [13]. To solve the unconstrained subproblems, we test out several algorithms: Riemannian gradient descent (RGD), Riemannian conjugate gradient descent (RCG), and Riemannian quasi-Newton method (RQN). All subsolvers use the default line-search and implementations from Manopt.jl version v0.4.41 [9].

We briefly discuss some references on the algorithms used here (see the Manopt documentation for a more thorough discussion). The generalization of gradient descent to the Riemannian case is discussed in [6]. For RCG, we used the conjugate descent (CD) coefficient update generalized to the Riemannian setting as described in [14]. The RQN algorithm is a Riemannian generalization of the BFGS method as proposed and analyzed in [15].

The AC optimal power flow model is constructed by version v0.19.9 of PowerModels.jl [8]. The detailed formulation of the problem is given in the PowerModels documentation. All equality constraints imposed by this model form our manifold. Our test problems are drawn from version 23.07 of the pglib-opf repository [16]. We chose cases from the repository with fewer than 300 buses so that the optimization completed in a reasonable amount of time.

Results

Fig. 2 gives the relative difference in terminal objective value for each tested algorithm. This value is computed as

$$\left| \frac{f(x_{ro}) - f(x_{ipopt})}{f(x_{ipopt})} \right|,$$

where the *ro* subscript indicates a Riemannian optimization solution and the *ipopt* subscript indicates the IPOPT solution. Fig. 3 gives the maximum constraint violation for the tested algorithms and IPOPT. Fig. 4 gives the number of iterations performed prior to termination. For IPOPT, this is the number of iterations reported by the optimizer output. For the Riemannian algorithms, this is the number of iterations performed by the unconstrained subproblem optimizer. For all three figures, failure to terminate is indicated by omitted values.

The RAL and REP with RGD subsolver perform reasonably well with RAL-RGD outperforming REP-RGD. In all but one case, RAL-RGD has a relative objective difference within 1% of the IPOPT solution whereas REP-RGD differences tend to be higher. The constraint violations are small being on the order of 10^{-5} in most tests. These are larger than the IPOPT values (between 10^{-8} and 10^{-5}). REP-RGD does produce solutions with significant constraint violation.

Both RCG variants perform poorly. While all cases terminated, the solutions are of poor quality. The relative objective function differences are significant (on the order of 1). The constraint violation is frequently large and rarely less than 10^{-2}. The poor performance seems to be the result of early termination of the RAL-RCG and REP-RCG algorithms.

The RQN based algorithms clearly performed the best with RAL-RQN slightly outperforming REP-RQN. These two methods also failed to terminate more than the others. This is likely due to un-optimized implementations of the retraction and vector transport rather than a failing in the algorithm. Both algorithms consistently produce solutions comparable to IPOPT in terms of objective function value and constraint violation.

References

[1] D. K. Molzahn, I. A. Hiskens et al., "A survey of relaxations and approximations of the power flow equations," Foundations and Trends® in Electric Energy Systems, vol. 4, no. 1-2, pp. 1–221, 2019.
[2] A. Wachter and L.T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," Mathematical programming, vol. 106, pp. 25–57, 2006.
[3] F. Capitanescu, "Critical review of recent advances and further developments needed in AC optimal power flow," Electric Power Systems Research, vol. 136, pp. 57–68, 2016.
[4] J. Maack, "Optimal power flow derived sparse linear solver benchmarks," National Renewable Energy Laboratory (NREL), Golden, CO (United States), Tech. Rep., 2023.
[5] M. Benzi, G. H. Golub, and J. Liesen, "Numerical solution of saddle point problems," Acta numerica, vol. 14, pp. 1–137, 2005.
[6] N. Boumal, An introduction to optimization on smooth manifolds. Cambridge University Press, 2023.
[7] S. Bolognani and F. D¨orfler, "Fast power system analysis via implicit linearization of the power flow manifold," in 2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton). IEEE, 2015, pp. 402–409.
[8] C. Coffrin, R. Bent, K. Sundar, Y. Ng, and M. Lubin, "PowerModels.jl: An open-source framework for exploring power flow formulations," in 2018 Power Systems Computation Conference (PSCC), June 2018, pp. 1–8.
[9] R. Bergmann, "Manopt.jl: Optimization on manifolds in Julia," Journal of Open Source Software, vol. 7, no. 70, p. 3866, 2022.
[10] J. M. Lee, Smooth manifolds. Springer, 2012.
[11] ——, Introduction to Riemannian manifolds. Springer, 2018, vol. 2.
[12] D. Luenberger and Y. Ye, Linear and Nonlinear Programming, ser. International Series in Operations Research & Management Science. Springer International Publishing, 2015.
[13] C. Liu and N. Boumal, "Simple algorithms for optimization on Riemannian manifolds with constraints," Applied Mathematics & Optimization, vol. 82, pp. 949–981, 2020.
[14] H. Sato, "Riemannian conjugate gradient methods: General framework and specific algorithms with convergence analyses," SIAM Journal on Optimization, vol. 32, no. 4, pp. 2690–2717, 2022.
[15] W. Huang, P.-A. Absil, and K. A. Gallivan, "A Riemannian BFGS method without differentiated retraction for nonconvex optimization problems," SIAM Journal on Optimization, vol. 28, no. 1, pp. 470–495, 2018.
[16] S. Babaeinejadsarookolaee, A. Birchfield, R. D. Christie, C. Coffrin, C. DeMarco, R. Diao, M. Ferris, S. Fliscounakis, S. Greene, R. Huang et al., "The power grid library for benchmarking AC optimal power flow algorithms," arXiv preprint arXiv:1908.02788, 2019.

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. This work was supported by the Laboratory Directed Research and Development (LDRD) Program at NREL. This research was performed using computational resources sponsored by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government.

2024 IEEE PES GENERAL MEETING
Seattle, Washington
July 21-25, 2024
NREL/PO-2C00-90522