Error Bounds on the DC Power Flow Approximation: A Convex Relaxation Approach

Krishnamurthy Dvijotham[1], Daniel K. Molzahn[2]

***Abstract*— Power flow models are fundamental to power systems analyses ranging from short-term market clearing and voltage stability studies to long-term planning. Due to the nonlinear nature of the AC power flow equations and the associated computational challenges, linearized approximations (like the DC power flow) have been widely used to solve these problems in a computationally tractable manner. The linearized approximations have been justified using traditional engineering assumptions that under "normal" operating conditions, voltage magnitudes do not significantly deviate from nominal values and phase differences are "small". However, there is only limited work on rigorously quantifying when it is safe to use these linearized approximations. In this paper, we propose an algorithm capable of computing rigorous bounds on the approximation error in the DC power flow (and, in future extensions, more general linearized approximations) using convex relaxation techniques. Given a set of operational constraints (limits on the voltage magnitudes, phase angle differences, and power injections), the algorithm determines an upper bound on the difference in injections at each bus computed by the AC and DC power flow power flow models within this domain. We test our approach on several IEEE benchmark networks. Our experimental results show that the bounds are reasonably tight (i.e., there are points within the domain of interest that are close to achieving the bound) over a range of operating conditions.**

NOMENCLATURE

$\angle x$:	The phase of the complex number x ($\in [-\pi, \pi)$)
$M^{\mathcal{H}}$:	Hermitian conjugate (conjugate transpose) of a matrix M
$(x)^H$:	Conjugate of the complex number x
$\mathbf{j}$:	The imaginary unit, $\sqrt{-1}$
$\mathrm{Im}\,(x)$:	Imaginary part of the complex number x
$\mathrm{Re}\,(x)$:	Real part of the complex number x

I. INTRODUCTION

Power flow models are integral to the operation of electric power systems: some form of power flow analysis is used in all operational aspects of the power grid (state estimation, market clearing, transmission switching, etc.). In several applications, the nonlinear nature of the AC power flow equations presents computational and convergence challenges. Particularly in applications like transmission switching [1] or post-blackout restoration [2] (where a power flow model needs to be solved iteratively within a mixed-integer program) and market clearing (where it is imperative to have an algorithm that consistently converges to well-defined, unique Locational Marginal Prices), the AC power flow model is not acceptable. Although great advances have been made in convexifying the AC power flow model (e.g., [3]–[7]), these techniques have not yet achieved the degrees of reliability and scalability that are necessary for practical system operations.

Thus, for several applications, linearized power flow approximations have been more appropriate and seen widespread adoption. The DC power flow is the most popular approximation. This is a linearization of the active power flow equations that ignores reactive power flows and transmission losses and assumes zero deviations from nominal voltage magnitudes and small phase angle differences. Several studies have been done on the validity of the DC power flow approximation (and more accurate variants of it) for various applications [8]–[10]. The papers conclude that the accuracy of the DC power flow depends on the application setting and range of test cases considered. In fact, the authors of [10] note that "At no stage in the tests were we able to discern any statistical patterns in the dc-flow error scatters. This defeated all our attempts to find concise, meaningful indices with which to characterize and display dc-model accuracies." Thus, it is clear that characterizing DC power flow errors is an involved task and cannot be boiled down to simple rules of thumb or analytical measures.

Another motivation comes from changes in power system operations. As power systems undergo transformations with greater demand-side participation and the inclusion of distributed energy resources, the assumptions justifying linear approximations may break down more often. For example, it has been shown that under flow-reversal conditions, a power system can enter modes where it is stuck at a low-voltage solution [11]. Given the possibility for this and other non-standard behavior, it is likely that the assumptions justifying the use of linear power flow models will break down more often in future power grids with bidirectional flows, and early detection of such behavior is important.

Thus, irrespective of the details of the linear approximation used, it is of interest to quantify the error in the DC approximation and other linear approximations of the power flow equations. To date, to the best of our knowledge, there has not been much work that rigorously quantifies error bounds on linear power flow approximations. A notable exception is the work presented in [12], where the authors develop a new linear power flow approximation for distribution networks

[1]Krishnamurthy Dvijotham, is with Pacific Northwest National Laboratory, Richland, Washington, USA. dvij@cs.washington.edu

[2]Daniel Molzahn, is with the Energy Systems Division, Argonne National Laboratory, Argonne, IL 60439, USA. dmolzahn@anl.gov Support from the U.S. Department of Energy, Office of Electricity Delivery and Energy Reliability under contract DE-AC02-06CH11357.

Authors contributed equally.

and prove a bound on the approximation error under the assumption that the injections are contained within a ball centered at zero injections. Recently, researchers have studied alternative linearizations of the AC power flow equations and related error bounds, using both rectangular coordinates [13] and an implicit linearization scheme [14]. The error bound in the implicit linearization scheme proposed in [14] requires the computation of uniform bounds on the second derivatives of an appropriate representation of the power flow equations, which may itself be a computationally difficult problem. Additionally, [2] proposes another linear approximation of the AC power flow equations that is suitable for inclusion in mixed-integer problems (e.g., post-blackout restoration, transmission switching, etc.) but does not develop error bounds. Our proposed approach is applicable to bounding the error for these linear approximations and can also be combined with mixed-integer programming solvers to handle problems with discrete variables.

In this paper, we develop computational approaches based on convex optimization that provide upper bounds on the error of the linearized power flow equations relative to the nonlinear AC power flow equations. This paper focuses on the DC power flow approximation. Future extensions will consider a wider range of linearizations. The error bounds are valid over a range of operational conditions (limits on voltage magnitude, phase angle differences, and active/reactive injections). Numerical experiments show that we can obtain fairly tight bounds (with an gap of less than few MW between the upper bound from the convex relaxation and a lower bound from a feasible point) over a relatively wide range of operational conditions. We envision that this approach will provide an effective computational test of when it is safe to use linearized power flow equations as a substitute for the nonlinear AC equations, and hopefully also yield insights into designing better linear approximations tailored to particular operating conditions and system characteristics.

II. Power Flow Model and Approximations

We use the standard bus-injection AC power flow model. The network is described as an undirected graph with a set of vertices $\mathcal{N} = \{1, \ldots, n\}$ (vertices are also known as *buses*), and a set of edges $\mathcal{E} = \{(i,j)\}$ (edges are also known as *lines*). We write $i \sim k$ to denote that i and k are neighbors in the network.

The power network is characterized by a complex symmetric (but not necessarily Hermitian) *admittance matrix* $Y \in \mathbb{C}^{n \times n}$, where $\mathbb{C}$ denotes the set of complex numbers. The matrix Y retains the sparsity structure of the network: $Y_{ik} \neq 0$ if and only if $(i,k) \in \mathcal{E}$. The off-diagonal elements of Y are related to resistance and reactance of the transmission lines:

$$Y_{ik} = -\frac{1}{(r_{ik} + \mathbf{j}x_{ik})^H} = \frac{-r_{ik} + \mathbf{j}x_{ik}}{r_{ik}^2 + x_{ik}^2}$$

where $\mathbf{j} = \sqrt{-1}$ and the line connecting buses i and k has series impedance $r_{ik} + \mathbf{j}x_{ik}$.[1] The diagonal elements are $Y_{ii} = g_{sh,i} + \mathbf{j}b_{sh,i} - \sum_{k \sim i} Y_{ik}$, where $g_{sh,i} + \mathbf{j}b_{sh,i}$ denotes the shunt admittance at bus i.

At each bus in the network, there is a complex voltage phasor (V_i) that represents the steady-state sinusoidal voltage at that bus. Denote the magnitude and angle of V_i as $|V_i|$ and θ_i, respectively. Each bus also has a complex net power injection $p_i + \mathbf{j}q_i$. The AC power flow equations are

$$\sum_{k \in \mathcal{N}} (Y_{ik})^H V_i (V_k)^H = p_i + \mathbf{j}q_i \qquad \forall i \in \mathcal{N} \qquad (1)$$

We ignore the traditional distinctions between slack, PV, and PQ buses here because our application scenario involves bounding the error of a linearized approximation of the power flow equations over a wide range of operating conditions offline (at a time scale over which the voltage magnitude setpoints and possibly even the designation of the slack bus may change). However, it is not difficult to incorporate PV buses into our approach. In fact, this will only improve the quality of our error bounds since the voltage magnitude constraints significantly restrict the range of allowed operating conditions.

The equations (1) constitute a system of nonlinear (quadratic) equations in the variables $\{\mathrm{Re}\,(V_i), \mathrm{Im}\,(V_i)\}_{i=1}^n$. Due to the difficulties arising from the nonlinear nature of the AC power flow equations, a linearization of the AC power flow model is typically used as a proxy for the true nonlinear model. The DC power flow model is a commonly used approximation that employs the following assumptions:

a. Reactive power flows can be neglected.
b. The transmission lines are lossless (i.e., their resistance can be ignored), so
$$Y_{ik} = \mathbf{j}\mathrm{Im}\,(Y_{ik}) = -\mathbf{j}B_{ik} = -\frac{\mathbf{j}}{x_{ik}}.$$
Shunt elements are ignored as well, so $Y_{ii} = -\sum_{k \sim i} Y_{ik}$.
c. The voltage magnitudes at all buses are approximately equal, so we can assume $|V_i| = 1$ at all buses.
d. There are small angle differences between connected buses such that $\sin(\theta_i - \theta_k) \approx \theta_i - \theta_k$, $\forall\,(i,k) \in \mathcal{E}$.

Under these assumptions, the AC power flow equations (1) reduce to the DC power flow equations:

$$\sum_{k \sim i} B_{ik}(\theta_i - \theta_k) = p_i \qquad \forall i \in \mathcal{N} \qquad (2)$$

There are several variants of DC power flow:

a. Cold start (linearization about a fixed point like a flat voltage profile, as in (2)) or warm start (linearization about a specified AC power flow solution).
b. Augmentation with linear approximations of active power losses.
c. Accounting for transformers, shunt elements, etc.

[1] Our approach is easily extensible to more general line models which include shunt admittances, non-zero phase shifts, and off-nominal voltage ratios. The numerical experiments in Section IV are implemented using the line model in MATPOWER [15].

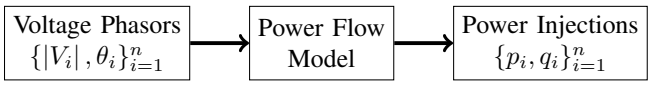

Fig. 1: A power flow model viewed as a mapping from voltages to power injections

These modifications are studied in detail in [10]. We do not get into the details of these more accurate approximations in this paper, but note that our techniques can easily be extended to analyze more sophisticated DC power flow models. Most transmission operators use one of these variants for the purposes of market clearing and only use the AC power flow model to check feasibility post-hoc. Hence, we will focus on quantifying the error bounds on the DC power flow model. Our technique also generalizes in principle to other linear approximations (e.g., the Linearized DistFlow [16] approximation that arises from a linearization of the Baran-Wu [17] model of distribution networks), but we leave concrete developments and applications for future work.

III. Main Technical Results

This section first provides a conceptual overview and then the technical details for our approach to bounding the DC power flow approximation error.

A. Overview of Our Approach

Given its widespread use and popularity, we focus on bounding the error in the DC power flow model relative to the AC power flow model. We view power flow models as mappings from the space of complex voltages to the space of power injections (see Fig. 1). In this view, the mappings are explicit: inputting the voltages in equations (1) and (2) yields power injections as outputs (either active and reactive power injections for the AC power flow model or just active power injections for the DC power flow model).

We envision our approach being applied to an operational power systems setting as follows. The system operator knows the current operating state of the network (complex voltages and power injections) and, from any appropriate load prediction algorithm, a range of expected operating conditions over the next few hours. The question of interest then is whether the DC power flow model will give satisfactory results over the range of expected future operating conditions. To give a quantitative answer to this question, we propose to solve the following optimization problem (stated in words here, with the exact mathematical formulation in Section III-B):

$$\max_i \quad \left| p_i^{AC}\left(|V|, \theta\right) - p_i^{DC}\left(\theta\right) \right|$$
$$\text{subject to operational bounds on } V, \theta, p^{AC}, \text{ and } q$$

where $p_i^{AC}\left(|V|, \theta\right)$ refers to the active power injection at bus i from the AC power flow model with inputs $|V|$ and θ, $p_i^{DC}\left(\theta\right)$ refers to the active injection at bus i from the DC power flow model with input θ, and $|\cdot|$ is the absolute value operator. The global solution to this optimization problem yields a guaranteed bound on the difference between the outputs (i.e., power injections) of the AC and DC power flow models with the same inputs (i.e., angles θ), provided that the system remains within the operational bounds on voltage magnitudes, phase angle differences, and power injections predicted by the system operator. If the maximum difference in outputs is small enough, one can confidently use the DC power flow model as a proxy for the AC power flow model. If not, the approximation may need to be refined, and the optimization algorithm can provide guidance on the refinement by suggesting parameter values (voltage magnitudes/angles/injections) at which the difference between the DC and AC power flow models can be large.

A similar approach could also be used in longer-term planning studies, where the computational burden of solving many scenarios and the possible formulation of bilevel/multi-stage optimization problems necessitates the use of linear power flow approximations. The proposed approach can quantify when DC power flow models can safely be used. If the DC power flow model is not shown to be sufficiently accurate, the proposed approach can suggest regions of the parameter space for refinement of the approximation.

We note that the viewpoint of power flow models as nonlinear mappings from voltages to power injections is the converse of the way power flow models are often used in practice. System operators typically specify power injections and solve the power flow equations to determine the corresponding voltages and flows. While it would be desirable to work with the reverse map directly (i.e., constrain injections to be the same in both the AC and DC power flow models and bound the differences in the voltage magnitudes and phase angles), the existence of multiple power flow solutions in the AC model complicates matters. When we maximize the difference between the outputs of the AC and DC power flow models, an optimization algorithm may choose a low-voltage, unstable solution of the AC power flow equations to maximize the difference from the DC power flow model. Since the system will be operated in a state corresponding to a high-voltage, stable power flow solution, selection of a low-voltage, unstable solution in the error bounding algorithm would result in an overestimate of the worst-case error between the AC and DC power flow models. One approach to dealing with this issue is to impose operational constraints on the AC model so that it is forced to pick the high-voltage solution (assuming that there indeed exists a unique high-voltage power flow solution). It is possible to extend our approach this setting, but we leave this for future work.

B. Mathematical Formulation of Error Quantification

We now formalize our approach. Consider the set of operational constraints on voltage magnitudes, phase angle differences, and active and reactive power injections:

$$\left(\underline{v}_i\right)^2 \leq |V_i|^2 \leq \left(\overline{v}_i\right)^2 \qquad \forall i \in \mathcal{N} \qquad (3a)$$
$$\underline{\theta}_{ik} \leq \theta_i - \theta_k \leq \overline{\theta}_{ik} \qquad \forall (i,k) \in \mathcal{E} \qquad (3b)$$
$$\underline{p}_i \leq p_i \leq \overline{p}_i \qquad \forall i \in \mathcal{N} \qquad (3c)$$
$$\underline{q}_i \leq q_i \leq \overline{q}_i \qquad \forall i \in \mathcal{N} \qquad (3d)$$

where $\overline{(\cdot)}$ and $\underline{(\cdot)}$ denote specified upper and lower limits, respectively, on the corresponding quantities. Note that the

active and reactive power injection limits in (3) are not necessarily generation capacity limits. Typically, we consider a range of injections around a nominal AC power flow solution, with the range determined by a forecast of how much load demands and generator outputs are expected to change over the time period of interest. The angle difference and voltage magnitude constraints typically correspond to network-defined limits. In practice, operator actions and automatic controllers will work to ensure that voltage magnitudes and angle differences are within reasonable limits.

Our formulation is intended to quantify the error from the DC approximation over a range of operational conditions, which may include variations in the voltage setpoints of generator buses. Thus, the active and reactive power injections and voltage magnitudes are all allowed to change.

The optimal objective value of the following problem gives an error bound on the active power injection at bus i:

$$\max_{V,\theta,p^{AC},q,p^{DC}} \quad |p_i^{AC} - p_i^{DC}| \qquad \text{subject to} \tag{4a}$$

$$\sum_{k\in\mathcal{N}} (Y_{ik})^H V_i (V_k)^H = p_i^{AC} + \mathbf{j} q_i \qquad \forall i \in \mathcal{N} \tag{4b}$$

$$\sum_{k\sim i} -\mathrm{Im}\,(Y_{ik})\,(\theta_i - \theta_k) = p_i^{DC} \qquad \forall i \in \mathcal{N} \tag{4c}$$

$$\angle V_i = \theta_i \qquad \forall i \in \mathcal{N} \tag{4d}$$

$$(\underline{v}_i)^2 \le |V_i|^2 \le (\overline{v}_i)^2 \qquad \forall i \in \mathcal{N} \tag{4e}$$

$$\underline{\theta}_{ik} \le \theta_i - \theta_k \le \overline{\theta}_{ik} \qquad \forall\,(i,k) \in \mathcal{E} \tag{4f}$$

$$\underline{p}_i \le p_i^{AC} \le \overline{p}_i \qquad \forall i \in \mathcal{N} \tag{4g}$$

$$\underline{q}_i \le q_i \le \overline{q}_i \qquad \forall i \in \mathcal{N} \tag{4h}$$

$$\theta_1 = 0 \tag{4i}$$

Constraint (4b) enforces the AC power flow equations, (4c) enforces the DC power flow equations, and (4d) enforces consistency between the AC power flow variables V and the DC power flow variables θ, where $\angle(\cdot)$ denotes the phase angle of the corresponding complex variable. Constraints (4e)–(4h) correspond to (3). Constraint (4i) sets the angle reference at bus 1.

While (4) characterizes the DC power flow accuracy in terms of errors in active power injections, note that our approach can easily be extended to studying error bounds on flows or other quantities of interest by modifying the objective function (4a).

Unfortunately, due to the nonlinear AC power flow constraints (4b) and the consistency constraints (4d), (4) is nonconvex and finding the global solution is challenging. Hence, we propose to use convex relaxation techniques to upper bound the global solution of (4).

The quadratic nonlinearities in (4b) can be handled using a semidefinite programming (SDP) relaxation [4], and the connection between the angles θ and the voltages V can be handled using the QC relaxation [18]. Applying these relaxations enlarges the feasible space, therefore providing an upper bound on the globally optimal objective value of (4). We also obtain lower bounds by using a nonlinear programming algorithm to obtain a local optimum of (4). The combination of these upper and lower bounds characterizes the worst-case error in the DC power flow approximation.

C. Convex Relaxation of (4)

We now describe the convex relaxation of (4). We define the Hermitian matrix $W = VV^{\mathcal{H}} \in \mathbb{C}^{n\times n}, W = W^{\mathcal{H}}$. Further, we introduce the auxiliary variables $v_i = |V_i| \in \mathbb{R}^n$. Both the power injection constraints (4b) and the voltage magnitude constraints (4e) are linear in terms of the entries of the matrix W. The entries of W can be expressed in terms of v_i and θ_i:

$$W_{ik} = v_i v_k \left(\cos(\theta_i - \theta_k) + \mathbf{j}\sin(\theta_i - \theta_k)\right), \tag{5a}$$

$$W_{ii} = v_i^2. \tag{5b}$$

The QC relaxation constructs convex envelopes for these nonlinear equality constraints. The right-hand sides of the conditions in (5) are replaced by a convex set that depends on the values of v and θ, such that the overall condition becomes a convex constraint on v, θ, and W. The QC relaxation is built using a set of templates that are described below.

Given variables $x \in [\underline{x}, \overline{x}]$ and $y \in [\underline{y}, \underline{y}]$, define

$$\langle x \rangle^T = \left\{ t : (\underline{x} + \overline{x})\, x - \underline{x}\overline{x} \ge t \ge x^2 \right\}, \tag{6a}$$

$$\langle x, y \rangle^M = \left\{ t : \begin{cases} t \ge y\underline{x} + x\underline{y} - \underline{x}\underline{y} \\ t \ge y\overline{x} + x\overline{y} - \overline{x}\overline{y} \\ t \le y\underline{x} + x\overline{x} - \underline{x}\overline{y} \\ t \le y\overline{x} + x\underline{y} - \overline{x}\underline{y} \end{cases} \right\}. \tag{6b}$$

Note that $\langle x \rangle^T$ and $\langle x, y \rangle^M$ are *set-valued functions* of x and y, with the sets depending on the bounds $\underline{x}$, $\overline{x}$, $\underline{y}$, and $\overline{y}$. The QC relaxation also defines convex envelopes of the sine and cosine functions. Define $x^m = \max(|\underline{x}|, |\overline{x}|)$ and

$$\langle x \rangle^S = \left\{ s : \begin{cases} s \le \cos\left(\frac{x^m}{2}\right)\left(x - \frac{x^m}{2}\right) + \sin\left(\frac{x^m}{2}\right) \\ s \ge \cos\left(\frac{x^m}{2}\right)\left(x + \frac{x^m}{2}\right) - \sin\left(\frac{x^m}{2}\right) \end{cases} \right\}, \tag{7a}$$

$$\langle x \rangle^C = \left\{ c : \begin{cases} c \le 1 - \frac{1-\cos(x^m)}{(x^m)^2} x^2 \\ c \ge \frac{\cos(\underline{x}) - \cos(\overline{x})}{\underline{x} - \overline{x}} (x - \underline{x}) + \cos(\underline{x}) \end{cases} \right\}. \tag{7b}$$

The convex envelopes $\langle x \rangle^T$, $\langle x, y \rangle^M$, $\langle x \rangle^S$, and $\langle x \rangle^C$ enclose the square, bilinear product, sine, and cosine functions, respectively.

The accuracy of the QC relaxation depends on the tightness of bounds on the voltage magnitudes and angle differences. The bounds on the voltage magnitudes and angle differences specified in the operational constraints (3) may not actually be achievable due to the limitations imposed by other constraints. To infer tighter bounds derived from the potentially weaker bounds specified in the operational constraints, we use a *bound tightening* algorithm [19]. (See [20], [21] for other bound tightening approaches.) Bound tightening algorithms identify situations where the combination of the AC power flow equations with the operational constraints imply tighter bounds on voltage magnitudes and angle differences than those originally specified in the operational

constraints. In order to perform this inference, we solve a set of optimization problems:

$$\min / \max \quad f \quad \text{subject to (1) and (3)} \tag{8}$$

where f is a placeholder for $|V_i|$ at each bus $i \in \mathcal{N}$ as well as the phase angle differences $\theta_i - \theta_k$ for each line $(i,k) \in \mathcal{E}$. We can use any relaxation technique to handle the non-convexity in (8) introduced by the power flow equations (1) and voltage magnitude limits (3a). Solving a relaxation of (8) provides a bound on the corresponding quantity f that must hold at any solution of the AC power flow equations within the operational constraints (3). If the bound resulting from (8) is tighter than the previously specified bound, the specified bound is replaced by the solution to (8).[2] Tighter bounds on certain quantities can improve the bounds for other quantities, so the bound tightening is performed iteratively until reaching a fixed point where no further bounds can be tightened. See [19]–[21] for details regarding efficient implementations of bound tightening algorithms.[3]

Using the convex envelopes defined in (6) and (7), we formulate the following convex relaxation of (4) for $\sigma \in \{-1,+1\}$:[4]

$$\max_{W,V^{\text{prod}},v,c,s,\theta,p^{AC},q,p^{DC}} \sigma\left(p_i^{AC} - p_i^{DC}\right) \quad \text{subject to} \tag{9a}$$

$$\sum_{k\in\mathcal{N}} (Y_{ik})^H W_{ik} = p_i^{AC} + \mathbf{j} q_i \qquad \forall i \in \mathcal{N} \tag{9b}$$

$$W \succeq 0,\ W = W^{\mathcal{H}} \tag{9c}$$

$$\sum_{k\sim i} -\text{Im}\left(Y_{ik}\right)\left(\theta_i - \theta_k\right) = p_i^{DC} \qquad \forall i \in \mathcal{N} \tag{9d}$$

$$W_{ii} \in \langle v_i \rangle^T \qquad \forall i \in \mathcal{N} \tag{9e}$$

$$\text{Re}\left(W_{ik}\right) \in \left\langle V_{ik}^{\text{prod}}, c_{ik} \right\rangle^M \qquad \forall (i,k) \in \mathcal{E} \tag{9f}$$

$$\text{Im}\left(W_{ik}\right) \in \left\langle V_{ik}^{\text{prod}}, s_{ik} \right\rangle^M \qquad \forall (i,k) \in \mathcal{E} \tag{9g}$$

$$V_{ik}^{\text{prod}} \in \langle v_i, v_k \rangle^M \qquad (i,k) \in \mathcal{E} \tag{9h}$$

$$V_{ii}^{\text{prod}} = W_{ii} \qquad \forall i \in \mathcal{N} \tag{9i}$$

$$c_{ik} \in \langle \theta_i - \theta_k \rangle^C \qquad \forall (i,k) \in \mathcal{E} \tag{9j}$$

$$s_{ik} \in \langle \theta_i - \theta_k \rangle^S \qquad \forall (i,k) \in \mathcal{E} \tag{9k}$$

$$\left(v_i^{\min}\right)^2 \le W_{ii} \le \left(v_i^{\max}\right)^2 \qquad \forall i \in \mathcal{N} \tag{9l}$$

$$v_i^{\min} \le v_i \le v_i^{\max} \qquad \forall i \in \mathcal{N} \tag{9m}$$

$$\tan\left(\underline{\theta}_{ik}\right)\text{Im}\left(W_{ik}\right) \le \text{Re}\left(W_{ik}\right) \le \tan\left(\overline{\theta}_{ik}\right)\text{Im}\left(W_{ik}\right) \qquad \forall (i,k) \in \mathcal{E} \tag{9n}$$

$$\underline{\theta}_{ik} \le \theta_i - \theta_k \le \overline{\theta}_{ik} \qquad \forall (i,k) \in \mathcal{E} \tag{9o}$$

$$p_i^{\min} \le p_i^{AC} \le p_i^{\max} \qquad \forall i \in \mathcal{N} \tag{9p}$$

$$q_i^{\min} \le q_i \le q_i^{\max} \qquad \forall i \in \mathcal{N} \tag{9q}$$

$$\theta_1 = 0 \tag{9r}$$

$$\text{Im}\left(W_{1,1}\right) = 0 \tag{9s}$$

The maximum of the solutions' objective values for $\sigma = \{-1,+1\}$ gives an upper bound on $|p_i^{AC} - p_i^{DC}|$. Constraints (9b) and (9c) form the SDP relaxation of the quadratic constraint (4b) (i.e., $W = VV^{\mathcal{H}}$ is relaxed to $W \succeq 0$). (We also exploit network sparsity using the approach in [23], [24].) Constraint (9d) formulates the DC power flow model. Constraints (9e)–(9k) form the QC relaxation with variables V^{prod}, v, c, s, and θ and construct the links to the SDP relaxation's variable W. Constraints (9l)–(9q) enforce the operational constraints (3). Constraints (9r) and (9s) set the angle reference.

We need to solve two convex optimization problems per bus, which could potentially be computationally intensive. However, the different problems are independent and can thus be solved in parallel. Further, the applications discussed previously only require solving the problem offline once every few hours.

D. Exactness of Relaxation

The SDP relaxation employed in (9) is *exact* when its solution satisfies the condition $\text{rank}\,(W) = 1$ [4]. For the purposes of obtaining an upper bound on the worst-case error in the DC power flow approximation, exactness of the relaxation is not necessary: we only use the bound given by the relaxation's objective value, not the decision variables. However, a solution for which $\text{rank}\,(W) = 1$ can also provide a *lower bound* on the worst-case error in addition to the upper bound available from the solution's objective value. Specifically, in this case, let η be a unit-length eigenvector corresponding to the non-zero eigenvalue λ of W. The voltage vector $V_{SDP} = \sqrt{\lambda}\,\eta$ is a feasible solution to (4), and thus the corresponding objective value lower bounds the worst-case error in the DC power flow approximation. Note that the potential for differences between the voltages in (9) associated with the QC and SDP relaxations (i.e., $v\angle\theta$ for the QC relaxation and the vector V_{SDP} implied by a rank-one W matrix for the SDP relaxation) mean that the objective value associated with V_{SDP} is not necessarily an upper bound on the worst-case error. In other words, exactness of (9) requires both the SDP and QC relaxations to be exact, which is generally not the case.[5]

IV. Numerical Results

In order to numerically evaluate our approach, we computed a nominal operating point by solving an AC Optimal Power Flow (OPF) problem for the IEEE test cases in the MATPOWER package [15]. This section compares

[2]Tighter in this context means that maximization of (8) yields a value that is less than the specified upper bound or that minimization of (8) yields a value that is greater than the specified lower bound.

[3]The numerical results in Section IV use the bound tightening algorithm described in [19], including constraints from both the QC relaxation [18] and the semidefinite programming relaxation [4].

[4]Note that the numerical results in Section IV also strengthen (9) using "Lifted Nonlinear Cuts" and slightly stronger relaxations of the sine and cosine functions for cases where the upper and lower angle difference limits $\overline{\theta}_{ik}$ and $\underline{\theta}_{ik}$ are either both non-negative or both non-positive for some line $(i,k) \in \mathcal{E}$. See [22] for the formulations of these constraints.

[5]Satisfaction of the rank condition $\text{rank}\,(W) = 1$ indicates that tightening the constraints corresponding to W (using, e.g., higher-order moment relaxations [5]) will not improve the optimal objective value of (9). However, even if $\text{rank}\,(W) = 1$, it may still be possible to tighten the constraints in (9) corresponding to the QC relaxation variables V^{prod}, v, c, s, and θ.

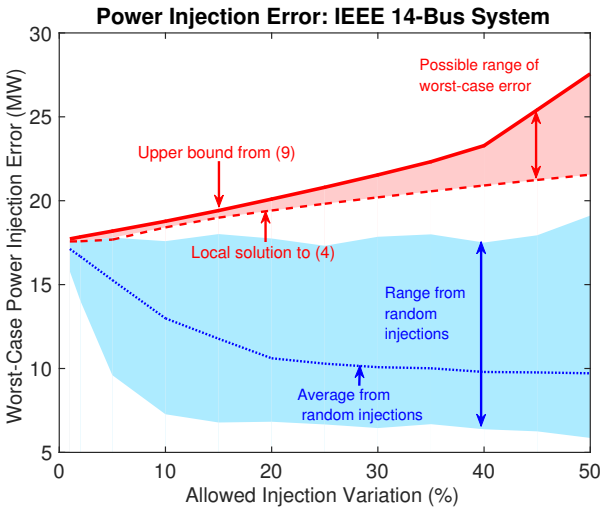

Fig. 2: Comparison of upper and lower bound on the worst-case error in active power injections from applying the DC power flow approximation to the IEEE 14-bus system. The proposed algorithm is able to estimate the error bound up to 6 MW while allowing a 50% variation in injections around the nominal operating point for this system. The blue region shows the range of worst-case errors from 1000 sets of active and reactive power injections (in steps of 5% variation) chosen from a uniform random distribution.

a lower bound on the worst-case error in the DC power flow approximation (obtained from a local optimum of (4) calculated using a modification of MATPOWER's OPF solver) with the upper bound (obtained from solving (9) for $\sigma = \{1, -1\}$ and taking the maximum of the two optimal values) for a range of injections around this nominal point. The computations were performed using MATLAB 2013a, YALMIP 2015.06.26 [25], and Mosek 7.1.0.28 on a computer with a quad-core 2.70 GHz processor and 16 GB of RAM.

The values $p_i^{\min}$, $p_i^{\max}$, $q_i^{\min}$, and $q_i^{\max}$ were obtained by allowing a fixed percentage variation around the nominal (OPF) solution (that is, the active and reactive injections are allowed to vary a fixed percentage from their nominal values, while also remaining within the specified upper and lower generation limits for buses with generators). The voltage magnitudes are constrained to the values specified in the test cases and a maximum angle difference of $\pm 30^\circ$ was enforced between connected buses.[6]

For the IEEE 14-bus system, Fig. 2 shows the upper and lower bounds on the worst-case error in the active power injections for the DC power flow approximation as a function of the allowed injection variation. The results show that even allowing 50% variation in injections, the upper and lower bounds are within 6 MW of each other. Given the actual variation in loads on real power systems, this should allow us to capture the variability over several hours.

Fig. 2 also shows the results of a computational experiment. Specifically, for each level of allowed injection variation (in steps of 5%), we considered 1000 sets of active and reactive power injections that satisfied the operational constraints (3), randomly selected from a uniform distribution. The blue region in Fig. 2 illustrates the range of worst-case error between the DC and AC power flow models among these power injections. The dotted blue line shows the average value of the worst-case error among the random injections. With values less than the upper bound from (9), these results numerically validate the proposed approach. The results also show that specifically chosen power injections may induce significantly larger worst-case errors for the DC power flow as compared to randomly chosen injections.

[6]The bound tightening approach [19] reduced the ranges of angle differences and voltage magnitudes in the QC relaxation's constraints in (9).

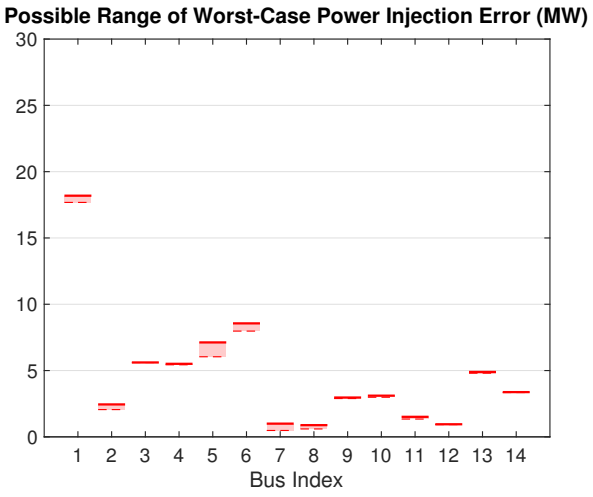

(a) 5% Allowed Power Injection Variation

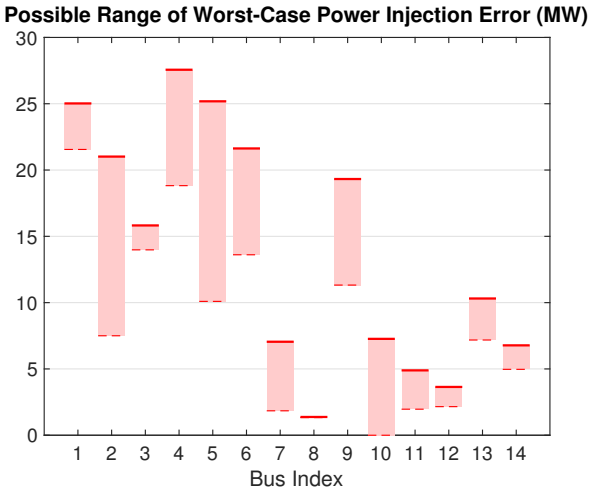

(b) 50% Allowed Power Injection Variation

Fig. 3: Comparison of upper and lower bounds on the worst-case error in active power injections at each bus from applying the DC power flow approximation to the IEEE 14-bus system for 5% and 50% allowed injection variation.

Plots showing upper and lower error bounds at each bus in the IEEE 14-bus system are given in Fig. 3a (for 5% allowed variation) and Fig. 3b (for 50% allowed variation). This information is valuable because it indicates the buses at which the DC power flow approximation can have the largest

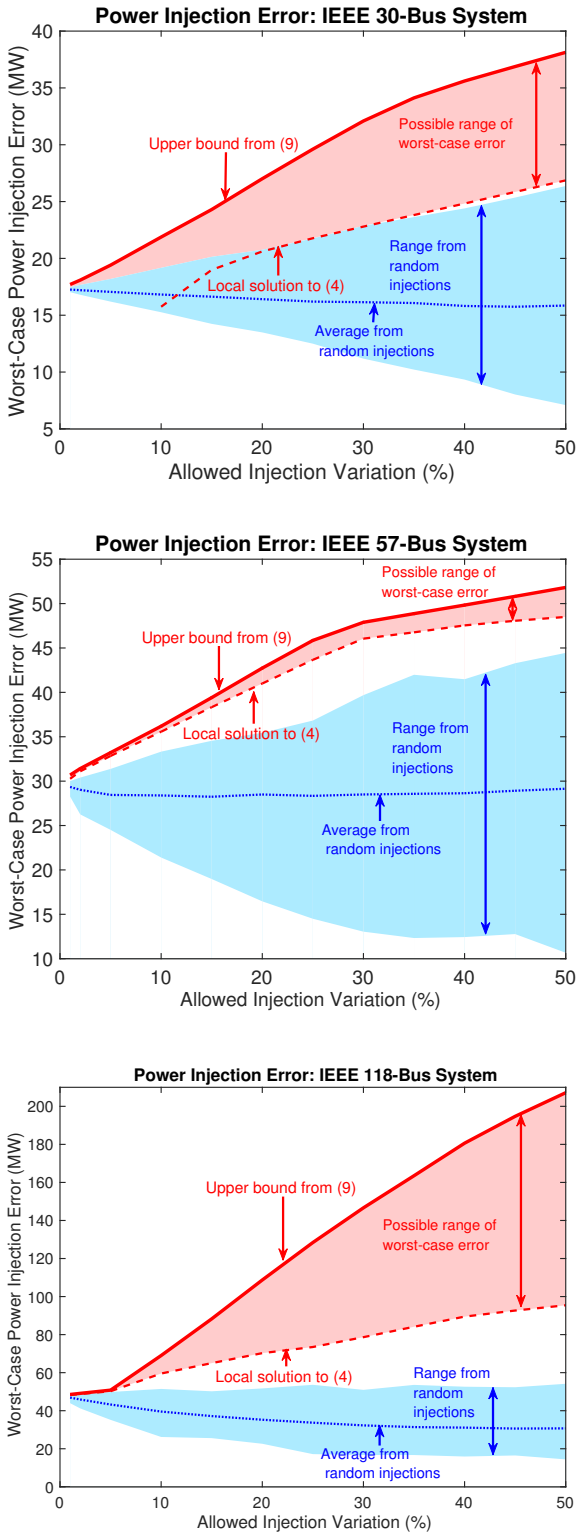

Fig. 4: Comparison of upper and lower bounds on the worst-case error in active power injections from applying the DC power flow approximation to several IEEE test cases. The blue region shows the range of worst-case errors from 1000 sets of active and reactive power injections (in steps of 5% variation) from a uniform random distribution.

error. We observe that the worst-case error varies by bus, which motivates future work to investigate specific power system characteristics that are associated with large potential error in the DC power flow approximation.

Fig. 4 shows the upper and lower bounds for the IEEE 30-, 57-, and 118-bus systems. Results for the IEEE 30- and 57-bus systems are similar to those for the IEEE 14-bus system in that the upper and lower bounds are relatively close (less than 10 MW) over a wide range of allowed injection variation. For the IEEE 118-bus system, the gap between the bounds is relatively small (less than 10 MW) when the allowed injection variation is less than 10% but increases quickly afterwards. As shown in blue, the results for the computational experiment with random power injections is repeated for each system, with qualitatively similar results to those in for the IEEE 14-bus system.

These results motivate future work to improve both the upper and lower bounds. Tightening the upper bounds may be possible by enhancing the convex relaxations, using techniques such as those in [5], [21], [22], [26]. We anticipate benefits from tightening the constraints associated with both the SDP relaxation (using, e.g., Lasserre's *moment relaxation hierarchy* [5], [27], which generalizes the SDP relaxation) and the QC relaxation (using, e.g., the stronger constraints on the angle variables in [21] and [26]).

Improving the lower bounds may be possible through better local solution techniques (e.g., initializing the local solver with the solution to the relaxation). We note that local solutions to (4) have been observed in the numerical experiments. Local solutions are evident for the IEEE 30-bus system in Fig. 4, where the DC power flow approximation had larger worst-case errors for some sets of random power injections than injections obtained via local solution of (4).

Computations for the IEEE 14-, 30-, 57-, and 118-bus systems required an average of approximately 1.0, 2.8, 13.3, and 75.0 minutes, respectively, in total for all $2n$ solutions of (9) to calculate an upper bound. The lower bounds had average solution times of 0.5, 2.4, 5.9, and 39 seconds, respectively, for the four systems. The vast majority (approximately 98% for the IEEE 118-bus system) of the solution time for (9) was spent performing the bound tightening. Alternative bound-tightening approaches (e.g., not including the SDP constraints in the bound tightening procedure, adopting from the approach in [20], etc.) thus have the potential to significantly reduce the computational burden of the proposed algorithm. Since all of the bound-tightening computations and the $2n$ solutions of (9) can be performed in parallel, we anticipate that the proposed approach can be scaled to practical power systems.

V. Conclusions and Future Work

We have presented a novel approach to quantifying errors in linearized power flow approximations. Preliminary results on IEEE test cases indicate that our technique is able to characterize relatively tight bounds on the worst-case error over a range of operating conditions. Given the motivations outlined in Section I, we envision that this will eventually

develop into a useful tool for detecting situations where linearized models cannot be trusted and one needs to resort to nonlinear power flow models to obtain sufficient accuracy. However, we acknowledge that the results, while promising, are preliminary and need to be tested carefully across a wider range of test cases and error metrics. The following are a few concrete directions for future work:

1. *Alternative power flow linearizations and test cases*: This paper used a basic cold-start DC power flow approximation. As noted in [10], more sophisticated variants of the DC power flow approximation often have better accuracy, and other power flow linearizations have been developed for particular applications (e.g., the Linearized DistFlow model [16], [28] and the linearizations in [2], [12]–[14]). Our approach extends in a straightforward manner to such linearizations, and we will investigate our results in this more general context. Also, we will study the robustness of the results across a broader set of power system test cases and perform comparisons to other error bounds [12]–[14].
2. *Tighter relaxations*: Recent literature [5], [21], [22], [26] has proposed enhancements to the QC and SDP relaxations used in this work. Using these enhanced relaxations should allow us to obtain tighter error bounds over wider ranges of operational conditions.
3. *Error characterization*: The proposed approach characterizes the worst-case linearization error for each bus, and natural extensions to other metrics (e.g., errors in line flows) can be similarly localized. This capability suggests the potential for identifying specific aspects of power system models (e.g., topology characteristics, electrical parameter values, etc.) that result in good or poor performance of the power flow approximations.
4. *Designing new linearizations*: Given a method for obtaining upper bounds on the error of linear power flow approximations, a natural next step is trying to design "optimal linearizations" that are tailored to a range of operational conditions. One potential approach may be to iterate between bound computation (finding a voltage profile that maximizes the error bound) and optimizing coefficients of a linear approximation to minimize the error at the previous worse-case voltage profile.

ACKNOWLEDGEMENT

Guidance from N.-J. Simon regarding the display of the figures is gratefully acknowledged.

REFERENCES

[1] K. Hedman, R. O'Neill, E. Fisher, and S. Oren, "Optimal Transmission Switching–Sensitivity Analysis and Extensions," *IEEE Trans. Power Syst.*, vol. 23, no. 3, pp. 1469–1479, 2008.

[2] C. Coffrin and P. Van Hentenryck, "A Linear-Programming Approximation of AC Power Flows," *INFORMS J. Comput.*, vol. 26, no. 4, pp. 718–734, 2014.

[3] R. Jabr, "Optimal Power Flow Using an Extended Conic Quadratic Formulation," *IEEE Trans. Power Syst.*, vol. 23, no. 3, pp. 1000–1008, Aug. 2008.

[4] J. Lavaei and S. Low, "Zero Duality Gap in Optimal Power Flow Problem," *IEEE Trans. Power Syst.*, vol. 27, no. 1, pp. 92–107, Feb. 2012.

[5] D. Molzahn and I. Hiskens, "Sparsity-Exploiting Moment-Based Relaxations of the Optimal Power Flow Problem," *IEEE Trans. Power Syst.*, vol. 30, no. 6, pp. 3168–3180, Nov. 2015.

[6] S. Low, "Convex Relaxation of Optimal Power Flow–Parts I: Formulations and Equivalence," *IEEE Trans. Control Network Syst.*, vol. 1, no. 1, pp. 15–27, Mar. 2014.

[7] ——, "Convex Relaxation of Optimal Power Flow–Part II: Exactness," *IEEE Trans. Control Network Syst.*, vol. 1, no. 2, pp. 177–189, Jun. 2014.

[8] T. Overbye, X. Cheng, and Y. Sun, "A Comparison of the AC and DC Power Flow Models for LMP Calculations," in *37th Hawaii Int. Conf. Syst. Sci. (HICSS)*, Jan. 2004.

[9] K. Purchala, L. Meeus, D. Van Dommelen, and R. Belmans, "Usefulness of DC Power Flow for Active Power Flow Analysis," in *IEEE PES General Meeting*, June 2005, pp. 454–459.

[10] B. Stott, J. Jardim, and O. Alsaç, "DC Power Flow Revisited," *IEEE Trans. Power Syst.*, vol. 24, no. 3, pp. 1290–1300, Aug. 2009.

[11] H. Nguyen and K. Turitsyn, "Voltage Multistability and Pulse Emergency Control for Distribution System With Power Flow Reversal," *IEEE Trans. Smart Grid*, vol. 6, no. 6, pp. 2985–2996, Nov. 2015.

[12] S. Bolognani and S. Zampieri, "On the Existence and Linear Approximation of the Power Flow Solution in Power Distribution Networks," *IEEE Trans. Power Syst.*, vol. 31, no. 1, pp. 163–172, Jan. 2016.

[13] S. Dhople, S. Guggilam, and Y. Chen, "Linear Approximations to AC Power Flow in Rectangular Coordinates," in *53rd Annu. Allerton Conf. Commun., Control, and Comput.*, Sept. 2015, pp. 211–217.

[14] S. Bolognani and F. Dörfler, "Fast Power System Analysis via Implicit Linearization of the Power Flow Manifold," in *53rd Annu. Allerton Conf. Commun., Control, and Comput.*, Sept. 2015, pp. 402–409.

[15] R. Zimmerman, C. Murillo-Sánchez, and R. Thomas, "MATPOWER: Steady-State Operations, Planning, and Analysis Tools for Power Systems Research and Education," *IEEE Trans. Power Syst.*, vol. 26, no. 1, pp. 12–19, 2011.

[16] M. Farivar, L. Chen, and S. Low, "Equilibrium and Dynamics of Local Voltage Control in Distribution Systems," in *IEEE 52nd Annu. Conf. Decis. Control (CDC)*, Dec. 2013, pp. 4329–4334.

[17] M. Baran and F. Wu, "Optimal Capacitor Placement on Radial Distribution Systems," *IEEE Trans. Power Del.*, vol. 4, no. 1, pp. 725–734, Jan. 1989.

[18] C. Coffrin, H. Hijazi, and P. Van Hentenryck, "The QC Relaxation: A Theoretical and Computational Study on Optimal Power Flow," *IEEE Trans. Power Syst.*, vol. 31, no. 4, pp. 3008–3018, July 2016.

[19] ——, "Strengthening Convex Relaxations with Bound Tightening for Power Network Optimization," in *Principles and Practice of Constraint Programming*, ser. Lecture Notes in Computer Science, G. Pesant, Ed. Springer, 2015, vol. 9255, pp. 39–57.

[20] C. Chen, A. Atamtürk, and S. S. Oren, "Bound Tightening for the Alternating Current Optimal Power Flow Problem," *IEEE Trans. Power Syst.*, vol. 31, no. 5, pp. 3729–3736, Sept. 2016.

[21] B. Kocuk, S. Dey, and A. Sun, To appear in *Oper. Res.,* preprint available at *arXiv:1504.06770*.

[22] C. Coffrin, H. Hijazi, and P. Van Hentenryck, "Strengthening the SDP Relaxation of AC Power Flows with Convex Envelopes, Bound Tightening, and Lifted Nonlinear Cuts," *arXiv:1512.04644*, Jan. 2016.

[23] R. Jabr, "Exploiting Sparsity in SDP Relaxations of the OPF Problem," *IEEE Trans. Power Syst.*, vol. 27, no. 2, pp. 1138–1139, May 2012.

[24] D. Molzahn, J. Holzer, B. Lesieutre, and C. DeMarco, "Implementation of a Large-Scale Optimal Power Flow Solver Based on Semidefinite Programming," *IEEE Trans. Power Syst.*, vol. 28, no. 4, pp. 3987–3998, 2013.

[25] J. Löfberg, "YALMIP: A Toolbox for Modeling and Optimization in MATLAB," in *IEEE Int. Symp. Compu. Aided Control Syst. Des.*, 2004, pp. 284–289.

[26] K. Bestuzheva, H. Hijazi, and C. Coffrin, "Convex Relaxations for Quadratic On/Off Constraints and Applications to Optimal Transmission Switching," Preprint: *http://www.optimization-online.org/DB_FILE/2016/07/5565.pdf*, 2016.

[27] J. Lasserre, *Moments, Positive Polynomials and Their Applications*. Imperial College Press, 2010, vol. 1.

[28] L. Gan and S. Low, "Convex Relaxations and Linear Approximation for Optimal Power Flow in Multiphase Radial Networks," *18th Power Syst. Comput. Conf. (PSCC)*, Aug. 2014.