Accurate Linear Cutting-Plane Relaxations for ACOPF

Daniel Bienstock · Matías Villagra

Abstract We present a pure linear cutting-plane relaxation approach for rapidly proving tight and accurate lower bounds for the Alternating Current Optimal Power Flow Problem (ACOPF) and its multi-period extension with ramping constraints. Our method leverages outer-envelope linear cuts for well-known second-order cone relaxations for ACOPF together with modern cut management techniques and reformulations to attain numerical stability. These techniques prove effective on a broad family of ACOPF instances, including the largest ones publicly available, quickly and robustly yielding tight bounds. Additionally, we consider the (frequent) case where an ACOPF instance is handled following a small or moderate change in problem data, e.g., load changes and generator or branch shut-offs. We provide significant computational evidence, on single and multi-period ACOPF instances, that the cuts computed on the prior instance provide very good lower bounds when warm-starting our algorithm on the perturbed instance.

Keywords Linear Relaxations · Outer-Approximation · Cutting-planes · QCQP · SOCP · Electricity Grid Optimization

Mathematics Subject Classification (2020) 90C25 · 90C05 · 90C06

1 Introduction

The Alternating-Current Optimal Power Flow (ACOPF) problem is a well-known challenging computational task. It is nonlinear, non-convex and with feasible region that may be disconnected; see [44], [22] and Proposition 8. In [71,18] it is shown that the feasibility version of the problem is strongly

This work was supported by an ARPA-E GO award.

Daniel Bienstock · Matías Villagra
500 W 120th St #315, New York, NY 10027
Tel.: +1 (212) 854-2942
E-mail: dano@columbia.edu, mjv2153@columbia.edu

NP-hard, and [53] proved that this decision problem is weakly NP-hard on star-networks.

In the current state-of-the-art, special-purpose interior point methods are empirically successful at computing very good solutions but cannot provide any bounds on solution quality. Strong lower bounds are available through second-order cone (SOC) relaxations [46,50]; however all solvers struggle when handling such relaxations for large or even medium cases (see [31]; we will provide additional evidence for this point). Other techniques, such as spatial-branch-and-bound methods applied to McCormick (linear) relaxations of quadratically-constrained formulations for ACOPF, tend to yield poor performance unless augmented by said SOC inequalities – interior point methods are still needed for upper bounds.

In this paper we present a fast (linear) cutting-plane method used to obtain tight relaxations for ACOPF, and its multi-period extension with ramping constraints, by appropriately handling the SOC relaxations. Our approach scales well to the largest ACOPF instances currently available. Our research thrust is motivated by three key observations:

(1) Linearly constrained convex quadratic optimization technology is, at this point, very mature – many solvers are able to handle massively large instances quickly and robustly; these attributes extend to the case where formulations are dynamically constructed and updated, as would be the case with a cutting-plane algorithm.
(2) Whereas the strength of the SOC relaxations has been known for some time, an adequate theoretical understanding for this behavior was not available.
(3) In power engineering practice, it is often the case that a power flow problem (either in the AC or DC version) is solved on data that reflects a recent, and likely limited, update on a case that was previously handled. In power engineering language, a 'prior solution' was computed, and the problem is not solved 'from scratch.' In the context of our type of algorithm, this feature opens the door for the use of *warm-started formulations*, i.e., the application of a cutting plane procedure that leverages previously computed cuts to obtain sharp bounds more rapidly than from scratch.

As an additional attribute arising from our work, our formulations are, of course, linear. This feature paves the way for effective *pricing* schemes, i.e., extensions of the locational marginal pricing setup currently used for real-time and day-ahead energy markets [61,40,12].

Our contributions

- We describe very tight and numerically stable linearly-constrained relaxations for ACOPF. The relaxations can be constructed and solved robustly and quickly via a cutting-plane algorithm that relies on proper cut management. On medium to (very) large instances our algorithm is competitive with or better than what was previously possible using nonlinear relaxations, both in terms of bound quality and solution speed. The robustness of our method is especially prominent in the multi-period setting, where our method is able to provide tight lower bounds with high accuracy, while the nonlinear relaxations are simply out of reach for nonlinear solvers.
- We provide a theoretical justification for the tightness of the SOC relaxation for ACOPF as well as for the use of our linear relaxations. As we argue, the SOC constraints (or tight relaxations thereof, or equivalent formulations) are *necessary* in order to achieve a tight relaxation. More precisely, we point out that the (linear) active-power loss inequalities (introduced in [15]) are outer-cone envelope approximations for the SOC constraints. Moreover, the loss inequalities provide a fairly tight relaxation because they imply that every unit of demand and loss is matched by a corresponding unit of generation (a fact not otherwise guaranteed) via a flow-decomposition argument [13].
- We demonstrate, through extensive numerical testing, that the warm-start feature for our cutting-plane algorithm is indeed effective. It is worth noting that this capability stands in contrast to what is possible using nonlinear (convex) solvers, in particular interior point methods. We also show numerical convergence of the dual variables of active-power balance constraints (which are used as input for computing locational marginal prices).
- The warm-start capability yields practicable and tight relaxations for *multi-period* formulations of ACOPF. The tightness of our formulations is certified through a heuristic for such problems.
- We bring forth to the ACOPF literature a discussion on approximate solutions to convex relaxations and their accuracy as lower bounds. We argue that linearly-constrained relaxations with convex quadratic objective possess robust theoretical bounding guarantees, in sharp contrast to what nonlinear relaxations such as SOCPs can offer.

The structure of this paper is as follows: in Section 2 we formally define ACOPF and its multi-period extension; in Section 3 we review the literature on convex relaxations for ACOPF with emphasis on linear relaxations; in Section 4 we describe two second-order cone relaxations, provide theoretical justification for the tightness of these relaxations, and introduce a new numerically better-behaved conic relaxation; in Section 5 we describe our cutting-plane framework, as well as the cut separation and cut management heuristics deployed in our algorithm; in Section 6 we provide a discussion on the accuracy of lower bounds of convex relaxations for ACOPF; in Section 7 we present significant computational experiments on the performance of SOCPs versus our algorithm and warm-started formulations for the single and multi-period

settings; and finally in Section 8 we conclude and outline future research directions. Our proofs can be found in Appendix 8.

This work significantly extends a much shorter conference version of this paper where the single-period setting is addressed [19]. The relaxation in Section 4.5 is new, as well as the multi-period computational results and the sections on accuracy and behavior of the dual variables of active-power balance constraints.

Fig. 1: ACOPF Formulation using polar coordinates

$$[\text{ACOPF}]: \quad \min \quad \sum_{k\in\mathcal{G}} F_k(P_k^g) \tag{1a}$$

subject to:

$\forall$ bus $k \in \mathcal{B}$:

$$\sum_{\{k,m\}\in\delta(k)} P_{km} = \sum_{\ell\in\mathcal{G}_k} P_\ell^g - P_k^d \tag{1b}$$

$$\sum_{\{k,m\}\in\delta(k)} Q_{km} = \sum_{\ell\in\mathcal{G}_k} Q_\ell^g - Q_k^d \tag{1c}$$

$\forall$ branch $\{k,m\} \in \mathcal{E}$:

$$P_{km} = G_{kk}|V_k|^2 + |V_k||V_m|(G_{km}\cos(\theta_k-\theta_m) + B_{km}\sin(\theta_k-\theta_m)) \tag{1d}$$

$$P_{mk} = G_{mm}|V_m|^2 + |V_k||V_m|(G_{mk}\cos(\theta_k-\theta_m) - B_{mk}\sin(\theta_k-\theta_m)) \tag{1e}$$

$$Q_{km} = -B_{kk}|V_k|^2 - |V_k||V_m|(B_{km}\cos(\theta_k-\theta_m) - G_{km}\sin(\theta_k-\theta_m)) \tag{1f}$$

$$Q_{mk} = -B_{mm}|V_m|^2 - |V_k||V_m|(B_{mk}\cos(\theta_k-\theta_m) + G_{mk}\sin(\theta_k-\theta_m)) \tag{1g}$$

$\forall$ generator $k \in \mathcal{G}$:

$$P_k^{\min} \le P_k^g \le P_k^{\max} \tag{1h}$$

$$Q_k^{\min} \le Q_k^g \le Q_k^{\max} \tag{1i}$$

$\forall$ bus $k \in \mathcal{B}$:

$$(V_k^{\min})^2 \le |V_k|^2 \le (V_k^{\max})^2 \tag{1j}$$

$\forall$ branch $\{k,m\} \in \mathcal{E}$:

$$\max\left\{P_{km}^2 + Q_{km}^2, P_{mk}^2 + Q_{mk}^2\right\} \le U_{km} \tag{1k}$$

$$|\theta_k - \theta_m| \le \bar{\Delta}_{km} \tag{1l}$$

2 Formal Problem Statement

In standard (single-period) ACOPF we are given, as input, a *power system* consisting of an undirected network endowed with numerical parameters describing physical attributes; a set of *generators*, and a set of complex-valued *loads* (or demands). The objective of the problem is to set complex voltages at the buses, and generator outputs, so as to satisfy loads at minimum cost – cost is incurred by the generators – while transmitting power following laws of physics and equipment constraints. A mathematical description is given above in (1) using the so-called *polar representation.*

We denote by $\mathcal{N} := (\mathcal{B}, \mathcal{E})$ the network, where $\mathcal{B}$ denotes the set of nodes, which (following power engineering conventions) we will refer to as *buses*, and $\mathcal{E}$ denotes the set of edges, which we will refer to as transmission *lines* or *transformers*, or generically, *branches*. We denote by $\mathcal{G}$ the set of generators of the grid, each of which is located at some bus; for each bus $k \in \mathcal{B}$, we denote by $\mathcal{G}_k \subseteq \mathcal{G}$ the generators at bus k.

Each bus k has a fixed load $P_k^d + jQ_k^d$, where $P_k^d \geq 0$ is termed active power load, and $-\infty < Q_k^d < +\infty$ is reactive power load; lower $V_k^{\min} \geq 0$ and upper $V_k^{\max} \geq 0$ voltage limits. For each branch $\{k, m\}$ we are given thermal limit $0 \leq U_{km} \leq +\infty$, and maximum angle-difference $|\Delta_{km}| \leq \pi$. Thus, the goal is to find voltage magnitude $|V_k|$ and phase angle θ_k at each bus k, active P^g and reactive Q^g power generation for every generator g, so that power is transmitted by the network so as to satisfy active P^d and reactive Q^d power demands at minimum cost.

In the above formulation (1), the physical parameters of each line $\{k, m\} \in \mathcal{E}$ are described by

$$Y_{km} := \begin{pmatrix} G_{kk} + jB_{kk} & G_{km} + jB_{km} \\ G_{mk} + jB_{mk} & G_{mm} + jB_{mk}, \end{pmatrix}$$

which is the (complex) admittance matrix of the transmission line $\{k, m\}$ (see [11] for background on transmission line modeling). These parameters model, in (1d)-(1g), active and reactive power flows. Moreover, inequalities (1k)-(1l) correspond to flow capacity constraints, and inequalities (1h)-(1j) impose operational limits on power generation and voltages. Constraints (1b)-(1c) impose active and reactive power balance; the left-hand side represents power injection at bus $k \in \mathcal{B}$, while the right-hand side represents net power generation (generation minus demand) at bus k. Finally, for each generator $k \in \mathcal{G}$, it is customary to assume the functions $F_k : \mathbb{R} \to \mathbb{R}$ in the objective (1a) are convex piecewise-linear or convex quadratic.

We remark that, often, constraint (1l) is not present, and, when explicitly given, concerns angle limits $\bar{\Delta}_{km}$ that are *small* (smaller than $\pi/2$). Under such assumptions there are equivalent ways to restate (1l) involving the arctangent function and other variables present in the formulation (the same applies to convex relaxations).

Please refer to the surveys [60], [14] for alternative, but equivalent, ACOPF formulations.

2.1 Multi-Period Formulation

Our techniques apply to a multi-period setting for ACOPF; this extension is important because it corresponds to the way that energy markets are operated. We will rely on a simple multi-period setting for ACOPF, where in each period t there is a standard ACOPF formulation. The variables for this period are linked to those other periods $t' \neq t$[1] via generator ramping constraints (defined

[1] As a simple case, $t' = t - 1$.

below). Other variants of this problem include "uptime/downtime" costs for generators, energy storage reserve requirements, and network topology control, among others (see e.g. [26,42,45]).

The multi-period problem can be described as follows: given a time-horizon of $T \in \mathbb{N}$ periods, a time-invariant undirected network endowed with numerical parameters describing physical attributes, and complex network loads $P^d_{i,t} + jQ^d_{i,t}$ for $i \in \mathcal{B}$ and $t \in \{0,\ldots,T-1\}$, the goal is to set complex voltages $V_{i,t}$ at the buses, and generator outputs $P^g_{k,t}+jQ^g_{k,t}$ for $k \in \mathcal{G}$ and $t \in \{0,\ldots,T-1\}$, so as to satisfy loads at minimum cost while transmitting power following laws of physics and equipment costs, and satisfying generator ramping constraints. For every generator $k \in \mathcal{G}$, and any time-period $t \in \{0,\ldots,T-1\}$, the associated ramping up and down constraints are given by[2]

$$(1 + r^u_{k,t})P^g_{k,t} - P^g_{k+1,t} \geq 0 \tag{2a}$$

$$(1 - r^d_{k,t})P^g_{k,t} - P^g_{k+1,t} \leq 0. \tag{2b}$$

where $r^u_{k,t}, r^d_{k,t} \in (0,1)$ are the ramping up and down rates, respectively. In summary, multi-period ACOPF for a time-invariant network $(\mathcal{B},\mathcal{E})$ with rational data $(P^d_{i,t}, Q^d_{i,t})_{i\in\mathcal{B},t\in\{0,\ldots,T-1\}}$ and $(Y_{km}, U_{km})_{\{k,m\}\in\mathcal{E}}$ can be defined by (1) and (2) on variables

$$(V_{i,t},\theta_{i,t})_{i\in\mathcal{B}},\ (P_{km,t},P_{mk,t},Q_{km,t},Q_{mk,t})_{\{k,m\}\in\mathcal{E}},\ (P^g_{k,t},Q^d_{k,t})_{k\in\mathcal{G}}$$

for each $t \in \{0,\ldots,T-1\}$.

3 Literature Review

First, we will briefly review the extensive literature on convex relaxations for ACOPF. An extensive review is provided in [60].

The simplest relaxations use, a starting point, a rectangular formulation of the ACOPF problem (rather than the polar setup described above) yielding a QCQP (quadratically constrained quadratic program) and rely on the well-known McCormick [57] reformulation to linearize bilinear expressions. This straightforward relaxation has long been known to provide very weak bounds.

The SOC relaxation introduced in [46], which is widely known as the Jabr relaxation (see Section 4.1), has had significant impact due to its effectiveness as a lower bounding technique. Briefly, it linearizes the power flow definitions (1d)-(1g) using $|\mathcal{B}| + 2|\mathcal{E}|$ additional variables and adds $|\mathcal{E}|$ valid SOC inequalities (see Section 4.1 for a complete derivation). Moreover, in [46] it is shown that for radial transmission networks, i.e., networks that have a tree structure, the load flow problem can be solved via an SOCP whose conic constraints correspond to the latter SOC inequalities. While on the one hand the SOC relaxation is strong, it also yields formulations that, in the case of large ACOPF instances, are very challenging even for the best solvers.

[2] Since P^g_k can potentially be negative for some generator $g \in \mathcal{G}$, then the latter constraints need to be formulated in terms of the absolute value of P^g_k.

A wide variety of techniques have been proposed to strengthen the Jabr relaxation. In [50] *arctangent constraints* are associated with cycles, with the goal of capturing the relationship between the additional variables in the Jabr relaxation and phase angles – these are formulated as bilinear constraints, and then linearized via McCormick inequalities. The two other strengthened SOC formulations proposed in [50] add polyhedral envelopes for *arctangent functions*, and dynamically generate semi-definite cuts for cycles in the network. [51] proposes a *minor-based* formulation for ACOPF (which is a reformulation of the rank-one constraints in the QCQP formulation for ACOPF [60]), which is relaxed to generate cutting-planes improving on the tightness of the Jabr relaxation. [31] introduced the Quadratic Convex (QC) relaxation (an SOC relaxation) which is related to the Jabr relaxation (Section 4.2), strengthened with polyhedral envelopes for sine, cosine and bilinear terms appearing in the power flow definitions (1d)-(1g). A novel machine learning-based approach is followed in [62] to obtain fast dual bounds for the Jabr relaxation for cases with up to 2869 buses.

A semidefinite programming relaxation based on the *Shor relaxation* [66] is presented in [7] followed by [52]. This formulation is at least as tight as the Jabr relaxation at the expense of higher computational cost [51]. In [52] it is proven that the SDP relaxation achieves zero duality gap under appropriate network assumptions. Experiments with SDP relaxations have been constrained by current SDP technology capabilities.

Next we review linear relaxations for ACOPF. [15,16] (also see [60]) introduce the so-called active-power loss linear inequalities which impose the fact that on any branch the active power loss is nonnegative. The resulting relaxation, which we term the linear-loss-relaxation, is shown to yield good lower bounds. In a similar vein, [30] propose the network flow and the copper plate relaxations. The network flow relaxation amounts to the linear loss-relaxation with additional sparse linear inequalities that lower bound net reactive power losses in appropriate cases[3]. Moreover, the "copper plate" relaxation is obtained from the network flow relaxation by neglecting the power flow equations entirely via aggregation of all active and reactive power injections in the network. Along these lines [69] provides a relaxation which enforces a (valid) linear relationship between active and reactive power losses by relaxing linear combinations of (1d)-(1g).

The technique in [15], inspired by Glover [39], ϵ-approximates the products of continuous variables (arising from the rectangular formulation of ACOPF [60]), to arbitrary precision, using binary expansions and McCormick inequalities. This process yields a mixed integer linear ϵ-approximation for ACOPF. Another linear ϵ-approximation, which is based on the Jabr relaxation, is used in [59]. Their main contribution is using the SOC linear approximation developed in [10] which requires $O(k_\ell \log(1/\epsilon))$ linear constraints and additional variables to ϵ-approximate a conic constraint of row size $k_\ell \in \mathbb{N}$. One of the al-

[3] Note that it is physically possible to have reactive power gains, i.e., negative reactive losses, see [11]).

gorithms in [58] is a successive linear programming (SLP) method – originally introduced by Kelley in [48] to address convex problems – which is applied to the Jabr relaxation. This process yields a linear relaxation for ACOPF, which was [58] on instances with up to 3375 buses.

The well-known, and widely used in practice, Direct Current Optimal Power Flow (DCOPF) model is a linear approximation for ACOPF [67,24]. It is based on a number of simplifications that are approximately valid under normal system operations; under this formulation active power losses are zero and reactive power is completely ignored [11]. A main drawback of this formulation is that AC feasible solutions might not even be feasible for DCOPF [9] – losses are the main cause of infeasibility – in contrast to the previously mentioned linear ϵ-approximations. In any case, this linear model is the canonical approximation of ACOPF for many extensions and applications such as welfare maximization and pricing in energy markets, day-ahead security-constrained unit commitment (SCUC), real-time security-constrained dispatch (SCD) and transmission switching (TS) among many others.

We refer the reader to the surveys [60,14,76] for additional material on convex relaxations for single-period ACOPF.

Single-period convex relaxations can naturally be extended to multi-period ACOPF. To the best of our knowledge, all of the computational experiments address small to medium-sized cases – which is not surprising given that the single-period (nonlinear) convex relaxations are already challenging enough (see e.g. Section 7.1.2). It is worth mentioning that most of the research in the multi-period setting has focused on obtaining AC primal bounds for numerous multi-period ACOPF features. Multi-period Jabr SOCPs are solved in: [56] to find approximate solutions to a robust multi-period ACOPF; [33] and [8] as subproblems of Benders decomposition algorithms for AC-constrained UC; [63] for SOC-constrained UC; and a multi-period mixed integer Jabr SOCP is solved [55] for AC-constrained UC.

Moreover, [77] solves a sequence of Jabr SOCP relaxations using penalty terms to enforce AC feasibility. On the other hand, notable linear approximations for multi-period ACOPF are the copper-plate linear, DCOPF, and strengthened versions of DCOPF [74] which account for reactive power and voltage magnitudes. These approximations are key ingredients for solving scalable UC or Security Constrained ACOPF (SC-ACOPF). In [3] a linearized multi-period Security Constrained Stochastic ACOPF based on the strengthened DC approximation [74] is presented.

Next, we succinctly review literature on heuristics – mostly based on convex relaxations or linear approximations – used to find AC feasible solutions. In the single-period setting, the linear programming Hot-Start and Warm-Start approximation models are proposed in [32] for finding locally optimal single-period AC solutions. The Hot-Start model assumes there is a solved AC base-point solution available, and it uses to fix voltage magnitudes in the power AC power flow definitions. Morevoer, they approximate $\sin(\theta)$ by

θ and use a linear convex approximation of $\cos(\theta)$ (under the assumption $\theta \in (-\pi/2, \pi/2)$). Since fixing the voltages might be too restrictive in some cases, their Warm-Start model chooses takes a subset of the buses as target buses, and linearizes deviations from the targeted voltage magnitudes. In [26] and [58] SLP algorithms are proposed for finding locally optimal single-period AC solutions.

With respect to the multi-period setting, [25] presents an outer-approximation iterative algorithm for the UC problem with AC network constraints (i.e, an multi-period ACOPF is being solved) – this method solves a finite sequence of MILP master problems and nonlinear subproblems (ACOPF with fixed binary variables). The nonlinear subproblems are tackled using the SLP presented in [27]. Moreover, in [33] an SLP is used as a penalty method to recover a UC problem with AC-constraints feasible solution from the multi-period Jabr relaxation. This is different from the technique used in [77] where a multiobjective SOCP relaxation is used to obtain an AC feasible solution. In [2] a Benders decomposition algorithm is used for AC-constrained UC – the subproblem is a multi-period ACOPF on a small network. Note that all of the previously mentioned references on convex relaxations of some multi-period ACOPF variant are used for AC feasibility.

Two general purpose nonlinear solvers which can be used to find solutions for single and multi-period ACOPF are Knitro [23] and IPOPT [72].

4 Convex Relaxations and Power Losses

4.1 Jabr relaxation

A well-known convex relaxation of ACOPF is the *Jabr relaxation* [46]. It linearizes the power flow definitions (1d)-(1g) using $|\mathcal{B}|+2|\mathcal{E}|$ additional variables and adds $|\mathcal{E}|$ rotated-cone inequalities. A simple derivation is as follows: For any line $\{k, m\} \in \mathcal{E}$, we define

$$v_k^{(2)} := |V_k|^2,\ c_{km} := |V_k||V_m|\cos(\theta_k - \theta_m),\ s_{km} := |V_k||V_m|\sin(\theta_k - \theta_m). \quad (3)$$

Clearly we have the following valid non-convex quadratic relation

$$c_{km}^2 + s_{km}^2 = v_k^{(2)} v_m^{(2)}, \quad (4)$$

which in Jabr [46] is relaxed into the (convex) inequality

$$c_{km}^2 + s_{km}^2 \leq v_k^{(2)} v_m^{(2)}. \quad (5)$$

This is a rotated-cone inequality hence it can be represented as a second-order cone constraint. Note that (3) can be used to represent the power flow

equations in (1d)-(1g) as, $\forall\{k,m\} \in \mathcal{E}$,

$$P_{km} = G_{kk} v_k^{(2)} + G_{km} c_{km} + B_{km} s_{km} \tag{6a}$$
$$P_{mk} = G_{mm} v_m^{(2)} + G_{mk} c_{km} - B_{mk} s_{km} \tag{6b}$$
$$Q_{km} = -B_{kk} v_k^{(2)} + B_{km} c_{km} - G_{km} s_{km} \tag{6c}$$
$$Q_{mk} = -B_{mm} v_m^{(2)} + B_{mk} c_{km} + G_{mk} s_{km}. \tag{6d}$$

In summary, the Jabr relaxation can be obtained from the formulation (1) by: (i) adding the $c_{km}, s_{km}, v_k^{(2)}$ variables; (ii) replacing (1d)-(1e) with (6); and (iii) adding the constraints (5).[4]

4.2 An alternative SOCP

Recall that complex power injected into branch $\{k,m\} \in \mathcal{E}$ at bus $k \in \mathcal{B}$ is defined by $S_{km} := V_k I_{km}^*$, hence

$$|S_{km}|^2 = |V_k|^2 |I_{km}|^2 \tag{7}$$

holds. Moreover, since complex power can be decomposed into active and reactive power as $S_{km} = P_{km} + jQ_{km}$, and recalling that $v_k^{(2)} := |V_k|^2$ while denoting $i_{km}^{(2)} := |I_{km}|^2$, we have

$$P_{km}^2 + Q_{km}^2 = v_k^{(2)} i_{km}^{(2)}. \tag{8}$$

By relaxing the equality (8) we obtain the *i2 rotated-cone inequality* [31,34]

$$P_{km}^2 + Q_{km}^2 \leq v_k^{(2)} i_{km}^{(2)}. \tag{9}$$

Since the variable $i_{km}^{(2)}$ can be defined by

$$i_{km}^{(2)} = \alpha_{km} v_k^{(2)} + \beta_{km} v_m^{(2)} + \gamma_{km} c_{km} + \zeta_{km} s_{km} \tag{10}$$

for appropriate constants $\alpha_{km}, \beta_{km}, \gamma_{km}$ and ζ_{km} (Appendix B), we obtain an alternative SOC relaxation. This formulation, which we call the *i2 relaxation*, is comprised by (1a)-(1c), (1h)-(1k), the definition of $i_{km}^{(2)}$ (10), and the rotated-cone inequalities (9).

It is known [68,31,29] that for each branch $\{k,m\}$ the system defined by the linearized power flows (6) plus the Jabr inequality (5), and the system defined by the linearized power flows (6), the rotated-cone inequality (9) and the linear definition of $i_{km}^{(2)}$ (29), are equivalent. It is to be noted that by this we mean that for each feasible solution to one system there is a feasible solution to the other one. However, equivalence may fail to hold if $i^{(2)}$ is upper bounded.

[4] We stress that (3) is *not* added.

Proposition 1 *The Jabr and the i2 relaxations are equivalent if the $i^{(2)}$ variables are not upper bounded, and otherwise the i2 relaxation can be strictly stronger.*

Proof Sufficiency was proven in [31]. For an example where the i2 relaxation is strictly stronger than the Jabr relaxation see Appendix A.

Our computational experiments corroborate this fact; we have found that linear outer-approximation cuts for the rotated-cone inequalities (5) and (9) have significantly different impact in lower bounding ACOPF (Section 7).

4.3 Losses

Transmission line losses are a structural feature of power systems [11]. It is a physical phenomenon in which the (complex) power sent from bus k to bus m along branch km will not necessarily be equal to the power received at bus m. This feature of power grids is captured (in steady state) by the power flow equations (1d)-(1g). Indeed, active-power loss on line $\{k, m\}$ can be defined as the quantity

$$\ell_{km} := P_{km} + P_{mk}. \tag{11}$$

First, we will focus on active-power losses. It is well-known (see, e.g., [34, 15]) that for a simple transmission line, i.e., a non-transformer branch without shunts, active-power loss equals

$$\ell_{km} = g_{km}|V_k - V_m|^2 \tag{12}$$

where g_{km} denotes line conductance (for simple lines we have $G_{kk} = G_{mm} =: g_{km}$). Thus, if $g_{km} \geq 0$ (i.e., the standard case [11,75]), the linear inequality $P_{km} + P_{mk} \geq 0$ is valid for ACOPF. This inequality still holds under (traditional) branch shunts and transformers provided that $g_{km} \geq 0$. In an arbitrary relaxation to ACOPF we might have $\ell_{km} < 0$, and, as we demonstrate below; this feature will make the relaxation weak. Moreover, critically, Proposition 2 below shows that the Jabr inequalities imply nonnegative active power losses as outer-envelope inequalities.

Additionally, the following theorem in [13] sheds light on the relationship between active-power losses and (global) active power balance. Consider an ACOPF instance on an undirected network $\mathcal{N} = (\mathcal{B}, \mathcal{E})$. Let us subdivide each branch $\{k, m\}$ by introducing a new *node*, denoted n_{km}. To properly place this result, we consider a solution to any relaxation for ACOPF. *Any* relaxation is of interest (even an exact relaxation) so long as variables P_{km} and P_{mk} are present, and active power flow balance constraints (1b) are given (or implied). Now, given a feasible solution to the relaxation we assign real flow values to the edges of the subdivided network, and orient those edges, as follows. Consider an arbitrary branch $\{k, m\}$.

(1) If $P_{km} \geq 0$ the edge between k and n_{km} is oriented from k to n_{km} and assigned flow P_{km}. If $P_{km} < 0$ the edge between k and n_{km} is oriented from n_{km} to k and assigned flow $-P_{km}$.
(2) Similarly, if $P_{mk} \geq 0$ the edge between m and n_{km} is oriented from m to n_{km} and assigned flow P_{mk}. If $P_{mk} < 0$ the edge between m and n_{km} is oriented from n_{km} to m and assigned flow $-P_{mk}$.

Thus, each edge in the subdivided network is oriented in the direction of the flow value that was assigned. Using the well-known network flow-decomposition theorem [1] yields that the flows in the subdivided network can be decomposed into a sum of *flow-carrying paths*, i.e., directed paths where

- Each path starts from a *source* (a bus with positive net generation or a node on a branch with negative loss) and ends at a *sink* (a bus with negative net generation or a node on a branch with positive loss),
- Each path carries a fixed positive flow amount.

In summary,

Theorem 1 (Theorem 1.2.11 [13]) *Consider any relaxation to an ACOPF instance that includes, or implies, active power flow balance constraints (1b). Then*

(a) The sum of active-power generation minus negative active power losses amounts equals the sum of active-power loads plus positive active-power losses. Furthermore, there exists a family of flow-carrying paths that accounts for all the flows in the subdivided network.
(b) There is a similar statement regarding reactive power flows (paraphrased here for brevity). The sum of reactive-power generation amounts plus the sum of reactive-power generated by line shunt elements, equals the sum of reactive-power loads and losses. Furthermore, there exist a family of paths that account for all reactive-power generation, loads, and losses. Each path has as origin a generator (with positive reactive-power generation) or a shunt element, and as destination a bus with a positive reactive-power load or a line loss. If all shunt susceptances are zero, then each path has a bus with positive reactive-power generation as origin.

The relevance of (a) is that when negative losses are present, then total generation may be smaller than total loads – effectively, the negative losses amount to a source of free generation and directly contribute to a lower objective value for ACOPF than AC feasible. While this point is also made in [31] regarding a particular relaxation, the above theorem provides a clear explanation for this fact that applies to *any* relaxation.

As an experiment, suppose we remove, from the Jabr relaxation, just one of the SOC constraints, say for branch $\{\hat{k}, \hat{m}\}$, thus obtaining a weaker relaxation; and suppose that, furthermore, the branch limits U_{km} are large (the usual case). Then, as we substantiate next, it is quite likely that the loss on $\{\hat{k}, \hat{m}\}$ will be negative in the weakened relaxation and the value of the relaxation will

be (much) weaker – precisely because that phenomenon would allow a solution with less generation, and thus, lower cost, than otherwise possible.

Table 1 provides empirical verification; it reports on 100 experiments in which exactly one Jabr constraint is removed from the formulation. "Avg Loss" represents the average loss across all branches; "Avg (br)" and "Min (br)" are the average and minimum (resp.) loss at the branch whose SOC constraint was bypassed; and "Jabr" denotes the value of the Jabr relaxation, while "weak Jabr" is the average value of the 100 weakened Jabr SOCs. For reference, "AC-L" represents the total loss of an AC locally optimal solution, while "SOC-L" is the total loss of an optimal solution to the Jabr SOC.

Table 1: Average and minimum losses 100 repetitions of removing a randomly selected SOC constraint.

Case	Avg Loss	Avg (br)	Min (br)	Jabr	weak Jabr	AC-L	SOC-L
case14	-0.3808	-0.4906	-1.7443	8075.12	6292.78	0.0929	0.0918
case118	0.1084	-0.7046	-5.1803	129340.00	126982.72	0.7740	0.7125
case300	1.8485	-1.1652	-6.1421	718654.00	714858.26	3.0274	2.8064

In Table 2 we provide additional evidence in the same direction. We report on 100 experiments, for larger cases, where we remove 100 Jabr constraints randomly selected using the uniform distribution. To highlight the relative orders of magnitude of the losses, we exhibit the total load for each case. The column "Sample" is an average over all the samples of the 100 repetitions.

Table 2: Average and minimum losses 100 repetitions of removing 100 randomly selected SOC constraints.

		Weak Jabr				Jabr		AC
		Avg Losses		Min Losses				
Case	Load	Total	Sample	Sample	Objective	Objective	Losses	Losses
1354pegase	730.60	-150.38	-379.60	-764.52	58021.55	74008.58	9.49	10.09
2869pegase	1324.37	-197.83	-402.59	-1119.85	112663.19	133872.69	14.24	15.51
3375wp	483.63	-65.50	-117.11	-343.72	6467352.82	7385372.70	6.90	8.30

The following example sheds light on the importance of imposing the non-negative active-power loss requirement.

Remark 1 Solving the McCormick relaxation of the rectangular QCQP formulation [60] of case3 from MATPOWER, we obtain an optimal solution with zero active-power generation. In particular[5], the active-power losses are

$$\ell_{12} = 0.0789, \quad \ell_{13} = -1.1941, \quad \ell_{23} = 0.0331,$$

[5] Since total active-power demand is 1.00 and there is a bus shunt at bus 2 with shunt conductance 0.1, and the magnitude of voltage at bus 2 squared is 0.8211 in our solution, we have that generation equals losses plus the sum of the loads.

i.e., total active-power losses equals -1.0821.

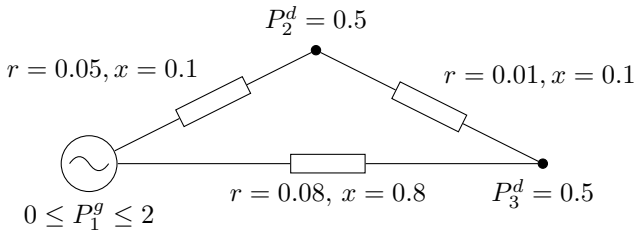

Fig. 2: Circuit representation of case3 from MATPOWER.

Given that we obtain active power loss as a linear combination of the power flow equations it is natural to wonder if the SOC relaxations to ACOPF (or outer approximations thereof) accurately approximate active-power losses. In Table 3 we provide empirical evidence that the quality of the approximations to total losses provided by the SOC relaxations is still very weak – even though the objective lower bounds proved by the SOC relaxations are quite close to ACOPF primal bounds.

Table 3: Active-power losses and optimality gaps.

	AC		Jabr SOC			
Case	Objective	Losses	Objective	Losses	Opt Gap	Loss Gap
1354pegase	74069.35	10.09	74008.58	9.48	0.08%	6.04%
ACTIVSg10k	2485898.75	23.01	2466666.10	13.88	0.77%	39.68%
ACTIVSg70k	16439499.83	156.83	16217263.66	121.31	1.35%	22.65%

It must be noted, however, that from a theoretical perspective the existence of active-power losses is not what makes ACOPF intractable; the NP-hardness proof in [71,18] relies on a lossless system.

4.4 Losses and cutting-planes

Next we provide theoretical justification for using an outer-approximation cutting-plane framework on the Jabr and $i2$ relaxations. We show that for transmission lines with $G_{kk} > 0 > G_{km} = G_{mk} \ge -G_{kk}$ and $B_{km} = B_{mk}$, in particular lines with no transformer or shunt elements, active-power loss inequalities are implied by the Jabr inequalities, and also by the definition of the $i^{(2)}$ variable. We begin with two simple technical observations.

First, consider a (generic) rotated cone inequality

$$x^2 + y^2 \le wz \tag{13}$$

where $x, y \in \mathbb{R}$ and $w, z \in \mathbb{R}_+$. Then

$$x^2 + y^2 \leq wz \iff (2x)^2 + (2y)^2 \leq (w+z)^2 - (w-z)^2 \tag{14a}$$

$$\iff ||(2x, 2y, w-z)^\top||_2 \leq w + z. \tag{14b}$$

Next, let $\lambda \in \mathbb{R}^3$ satisfy $||\lambda||_2 = 1$. Then, by (14b),

$$2\lambda_1 x + 2\lambda_2 y + \lambda_3(w-z) \leq ||\lambda||_2 \, ||(2x, 2y, w-z)^\top||_2 \leq w + z. \tag{15}$$

Inequality (15) provides a generic recipe to obtain outer-envelope inequalities for the rotated cone (13). As a result of these developments, we have:

Proposition 2 *For a transmission line $\{k, m\} \in \mathcal{E}$ with $G_{kk} > 0 > G_{km} = G_{mk} \geq -G_{kk}$ and $B_{km} = B_{mk}$, the Jabr inequality $c_{km}^2 + s_{km}^2 \leq v_k^{(2)} v_m^{(2)}$ implies, as an outer envelope approximation inequality, that $\ell_{km} \geq 0$.*

Proof See Appendix C.

Moreover, for simple transmission lines the definition of the variable $i^{(2)}$ implies the active-power loss inequalities [68,28]. Actually, if a line $\{k, m\}$ has no shunts, then it is folklore (and follows from basic physics) that $i_{km}^{(2)} \geq 0$ implies $\ell_{km} \geq 0$.

Proposition 3 *For any line transmission line $\{k, m\} \in \mathcal{E}$ with $y_{km}^{sh} = 0$, the definition of the variable $i_{km}^{(2)}$ implies the active-power loss inequality associated with that line.*

Proof See Appendix D.

4.5 A numerically better-behaved i2 relaxation

The i2 relaxation is at least as strong as the Jabr relaxation, and in congested networks it can be strictly stronger (Proposition 1). Thus, the i2 relaxation is a natural candidate for outer-approximation via a cutting-plane framework (Section 5); we numerically verify the strength of a linear outer-approximation of the i2 relaxation in Section 7.1.1. The strength of this relaxation comes at a price: numerical instability. To be precise, the potentially large coefficients in the definition of the $i^{(2)}$ variables can result in a ill-conditioned SOC instance (see, e.g., [49]). Next we illustrate this issue with a simple case to show how the Jabr outer-envelope inequalities complement the strength of the $i^{(2)}$ definitions. We conclude this subsection with a new, numerically stable modified i2 relaxation, which we term i2+(ρ), where $\rho > 0$.

Let $\{k, m\}$ be a transmission line with a limit $U_{km} < +\infty$, and without any transformers or shunt elements. By (29) we have

$$i_{km}^{(2)} = \left(\frac{1}{r_{km}^2 + x_{km}^2}\right)\left(v_k^{(2)} + v_m^{(2)} - 2c_{km}\right) \tag{16}$$

where r_{km} and x_{km} denote the line's resistance and reactance. Moreover, equation (7) implies that $i^{(2)}_{km}$ can be upper-bounded by the constant $H_{km} := U_{km}/(V^{\min}_k)^2$, where $V^{\min}_k$ corresponds to the lower limit for voltage magnitude at bus k (both voltage limits are close to 1 under standard data conditions).

Suppose that this simple line $\{k, m\}$ has small resistance, e.g., on the order of 10^{-5}. Then $1/(r^2_{km} + x^2_{km})$ can be on the order of 10^8. In addition, assume that the limit U_{km} is *small*. Thus, $(r^2_{km} + x^2_{km})H_{km}$ can be fairly small, which yields

$$v^{(2)}_k + v^{(2)}_m - 2c_{km} \leq (r^2_{km} + x^2_{km})H_{km} \approx 0. \tag{17}$$

Since $v^{(2)}_k + v^{(2)}_m - 2c_{km} \geq 0$ is a Jabr outer-envelope cut (proof of Proposition 2), inequality (17) is forcing our solutions to lie near the surface of the rotated-cone $c^2_{km} + s^2_{km} \leq v^{(2)}_k v^{(2)}_m$, i.e., it is cutting-off solutions which are in the interior of the Jabr cone for line $\{k, m\}$. Notice that the more capacitated the transmission line is, the stronger the effect of the $i^{(2)}$ variables (Figure 3).

This key observation motivates a numerically more stable relaxation. Consider the definition of $i^{(2)}_{km}$ in Appendix B for appropriate constants $\alpha_{km}, \beta_{km}, \gamma_{km}$ and ζ_{km}:

$$i^{(2)}_{km} = \alpha_{km} v^{(2)}_k + \beta_{km} v^{(2)}_m + \gamma_{km} c_{km} + \zeta_{km} s_{km}. \tag{18}$$

Let $\rho > 0$ be some fixed parameter, and consider the following heuristic: for every transmission line $\{k, m\} \in \mathcal{E}$,

(i) if $\alpha_{km} > \rho$, we say that the branch induces a *bad-i2*, and we define the following inequalities (which are valid by the previous discussion)

$$v^{(2)}_k + \frac{\beta_{km}}{\alpha_{km}} v^{(2)}_m + \frac{\gamma_{km}}{\alpha_{km}} c_{km} + \frac{\zeta_{km}}{\alpha_{km}} s_{km} \geq 0, \tag{19a}$$

$$v^{(2)}_k + \frac{\beta_{km}}{\alpha_{km}} v^{(2)}_m + \frac{\gamma_{km}}{\alpha_{km}} c_{km} + \frac{\zeta_{km}}{\alpha_{km}} s_{km} \leq \frac{H_{km}}{\alpha_{km}}; \tag{19b}$$

(ii) otherwise, we say that the branch induces a *good-i2*, and we define $i^{(2)}_{km} = \alpha_{km} v^{(2)}_k + \beta_{km} v^{(2)}_m + \gamma_{km} c_{km} + \zeta_{km} s_{km}$.

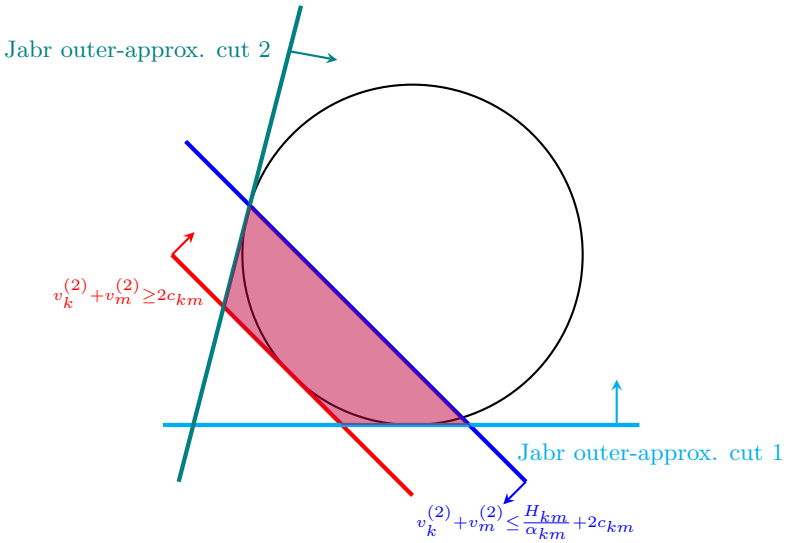

Fig. 3: The black circle represents a horizontal cross-section of a projected Jabr cone (5) onto the $(v_k^{(2)}, v_m^{(2)}, c_{km})$-space. The red and blue inequalities represent (19) for a simple transmission line, while the teal and cyan inequalities represent two Jabr outer-approximation cuts.

Empirically, we see in Table 4 that the coefficients of the inequalities (19) are relatively small; clearly closer to the range of coefficients of the linear constraints of the Jabr SOCP, and to the coefficients[6] of the Jabr outer-envelope cuts which will be deployed in our cutting-plane algorithm (refer to the linearly constrained relaxation M in Section 5.2.2). We observe that roughly around 90% of the branches correspond to *bad-i2*s.

Table 4: Statistics of the coefficient α in (18) and of the coefficients in (19).

		α		$\frac{\beta}{\alpha}$		$\frac{\gamma}{\alpha}$		$\frac{\zeta}{\alpha}$		$\frac{H}{\alpha}$		
Case	%bad-i2s	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Min	Avg
ACTIVSg10k	91.33	1.00e10	1.09e06	1.47	1.01	-1.80	-2.01	0.88	0.00	18.52	2.57e-09	3.06
10000goc-sad	88.50	2.23e06	2.41e05	1.21	1.00	-1.90	-2.00	0.01	0.00	1.89	4.64e-06	0.48
30000goc-api	96.13	2.41e07	3.66e05	1.31	1.00	-1.94	-2.00	0.01	5.90e-06	1.89	2.16e-06	0.71

By making reasonable assumptions about transmission line parameters we can derive the following upper bounds for the coefficients in inequalities (19).

Proposition 4 *Let $\{k,m\} \in \mathcal{E}$ be a transmission line that satisfies $|b| > b^{sh}$, $g^{sh} = 0$ and $|b| > g$. Then, the coefficients $\frac{\beta_{km}}{\alpha_{km}}$, $\frac{\gamma_{km}}{\alpha_{km}}$, and $\frac{\zeta_{km}}{\alpha_{km}}$ in (19) are*

[6] The size of coefficients of these cuts are linear in the values of variables c, s and v in any solution to a linear relaxation of the Jabr SOCP; see Proposition 5.

bounded by $\max\{\tau_{km}^2, 3\tau_{km}\}$*, where* τ_{km} *is the transformer tap ratio located at bus* k*.*[7]

Proof See Appendix G.

In consequence, given $\rho > 0$, we define our numerically better-behaved i2+(ρ) relaxation by (1a)-(1c), (1h)-(1k), the definition of $i_{km}^{(2)}$ (29) and the rotated-cone inequalities (9) for every good-i2 branch $\{k, m\}$, and inequalities (19) for every bad-i2 branch $\{k, m\}$. We use the i2+(ρ) relaxation for our multi-period experiments (see discussion and results in Section 7.3).

5 Cutting-plane Framework: Cut Separation and Management

In this paper we use a dynamically generated linearly-constrained relaxation as a lower bounding procedure for ACOPF; relying on a perspective sometimes associated with integer programming.

Given a set X in $\mathbb{R}^n$, we say that a convex inequality $g(x) \leq d$ is *valid* for X if $g(x) \leq d$ holds for all $x \in X$. Given X and $\overline{x} \notin X$, then we say that $c^\top x \leq d$ is a (linear) *cut* for X if the inequality is valid for X but it is not satisfied by $\overline{x}$. A (linear) cutting-plane algorithm [48,64] for a set X is an iterative procedure in which, starting from an initial (linear) relaxation, in every round (linear) cuts are added to approximate X. Typically, these cuts are computed in iterative fashion; at each round an optimal solution $\overline{x} \notin X$ to the current relaxation is separated using a computed cut. We note that in addition to Kelley's cutting-plane algorithm [48] other related algorithms have been proposed in the literature. See [65] for background.

In the case of single-period experiments our target set X is the i2 relaxation of ACOPF, while we use the i2+(ρ) relaxation for the multi-period setting. The reason is simple: the multi-period i2 relaxation for medium to large networks is not numerically tractable.

In support of this statement, we note that direct solution of the Jabr and i2 relaxations of ACOPF, for large instances, is computationally prohibitive and often results in non-convergence (see tables 6, 7, 8, 9, 12, 16 and 17). Our empirical evidence further shows that outer-approximation of the rotated-cone inequalities (in either case) requires a large number of cuts in order to achieve a tight relaxation value. Moreover, employing such large families of cuts yields a relaxation that, while linearly constrained, still proves challenging – both from the perspective of running time and numerical tractability. Nonetheless, a characteristic feature of our iterative procedure is its robustness to potentially suboptimal termination of the oracle used to solve the LPs or convex QPs; *independent* of the quality of the primal solution obtained, our linear cuts will always be valid.

However, as we show, adequate cut management proves successful, yielding a procedure that is (a) rapid, (b) numerically stable, and (c) provides a very

[7] If the line does not have a transformer, then $\tau_{km} = 1$; generally, tap ratios do not take extreme values.

tight relaxation (Tables 8 and 9). The critical ingredients in this procedure are: (1) fast cut separation; (2) appropriate violated cut selection; and (3) effective dynamic cut management, including rejection of *nearly-parallel* cuts and removal of *expired* cuts, i.e., previously added cuts that are slack (see Section 5.2.1).

Our procedure possesses efficient warm-starting capabilities – this is a central goal of our work. Previously computed cuts, for some given instance, can be re utilized and loaded into new runs of *related* instances, hence leveraging previous computational effort. It is worth noting that this reoptimization feature stands in sharp contrast to what is possible using nonlinear (convex) solvers. In Section 7.2 we justify the validity of this feature and Tables 8 and 9 summarize extensive numerical evidence on its performance relative to solving the SOCPs 'from scratch'. We remark that adequate cut management is what enables this capability in the case of large single and multi-period ACOPF instances.

5.1 Two Simple Cut Procedures

The following (routine) results provide a fast procedure for separating over the rotated-cone inequalities

$$c_{km}^2 + s_{km}^2 \leq v_k^{(2)} v_m^{(2)}, \qquad P_{km}^2 + Q_{km}^2 \leq v_k^{(2)} i_{km}^{(2)}. \tag{20}$$

Proposition 5 *Consider the second-order cone* $C := \{(x, s) \in \mathbb{R}^n \times \mathbb{R}_+ : ||x||_2 \leq s\}$. *Suppose* $(x', s') \notin C$ *with* $s' \geq 0$. *Then the projection of* (x', s') *onto* C *is given by*

$$x_0 := s_0 \frac{x'}{||x'||}, \quad s_0 := \frac{||x'|| + s'}{2}.$$

Moreover, the hyperplane which achieves the maximum distance from (x', s') *to any hyperplane separating* C *and* (x', s') *is given by*

$$(x')^t x \leq s||x'||.$$

Proof See Appendix E.

Corollary 1 *Let* $C := \{(x, y, w, z) \in \mathbb{R}^2 \times \mathbb{R}^2_+ : x^2 + y^2 \leq wz\} \subseteq \mathbb{R}^4$ *and suppose that* $(x', y', w', z') \notin C$ *where* $w' + z' > 0$. *Then the hyperplane separating* C *from* (x', y', w', z') *and at maximum distance from this point is given by*

$$(4x')x + (4y')y + ((w' - z') - n_0)w + (-(w' - z') - n_0)z \leq 0, \tag{21}$$

where $n_0 := ||(2x', 2y', w' - z')^\top||_2$.

Proof Rewriting the rotated-cone inequality as (14b) and a direct application of Proposition 5 gives us the desired separating hyperplane.

The following proposition gives us a simple procedure for computing linear cuts for violated limit inequalities

$$P_{km}^2 + Q_{km}^2 \leq U_{km}. \tag{22}$$

Proposition 6 *Consider the Euclidean ball in* $\mathbb{R}^2$ *of radius* r, $S_r := \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq r^2\}$, *and let* $(x', y') \notin S_r$. *Then the separating hyperplane that is at maximum distance from* (x', y') *is given by*

$$(x')x + (y')y \leq r||(x', y')||_2. \tag{23}$$

Proof See Appendix F.

5.2 Basic Cutting-Plane Algorithm

5.2.1 Single-period

In what follows we describe our cutting-plane algorithm for the single-period setting. First, we define a linearly constrained base model M_0 as follows:

$$[M_0]: \qquad \min \qquad \sum_{k \in \mathcal{G}} F_k(P_k^g) \tag{24a}$$

subject to:

$$\text{constraints (1b), (1c), (6), (1h) – (1i), (1j)} \tag{24b}$$

In other words, we consider the linearized power flow equations of the Jabr SOCP and all the linear constraints in (1). Recall that for single-period ACOPF, our target set X is the i2 relaxation.

In every round of our iterative procedure, linear constraints will be added to and removed from the current relaxation, which is initialized as M_0. We will denote by M the relaxation at a generic iteration of our cutting-plane algorithm.

Given a feasible solution $\bar{x}$ to M, and letting $f_{km}(x) \leq 0$ be some valid convex inequality (20) or (22), our measure of *cut-quality* is $\max\{f_{km}(\overline{x}), 0\}$, i.e., the amount by which $\overline{x}$ violates the inequality. Let $\epsilon > 0$. For each type $\tau \in \{\text{Jabr}, \text{i2}, \text{limit}\}$ of inequality we select the top p_τ percent –a fixed parameter– most highly violated inequalities of type τ whose violation is greater than ϵ – these are the τ-candidates (Line 7 of the pseudocode).

For each list of τ-candidates, we compute cuts for the corresponding branches using the cut procedures described in Section 5.1. Candidate cuts will be rejected if they are *nearly parallel* to incumbent cuts in M – near-linear dependency is a common source of ill-conditioning in optimization models [43,49]. To be precise, given $\epsilon_{par} > 0$, we say that two linear inequalities $c^t x \leq 0$ and $d^t x \leq 0$ are ϵ_{par}*-parallel* if the cosine of the angle between their normal vectors c and d is strictly more that $1 - \epsilon_{par}$.

Finally, we describe a heuristic for *cleaning-up* our formulation. For each added cut, we keep track of its *cut-age*, i.e., the count of rounds since it was

added. Then, in every iteration, if a cut $c^\top x \leq d$ has age greater or equal than a fixed parameter $T_{age} \in \mathbb{N}$, and it is ϵ-*slack*, i.e., $d - c^\top \overline{x} > \epsilon$, then it is dropped from M.

Algorithm 1 Cutting-Plane Algorithm

```
procedure CUTPLANE
    Initialize r ← 0, M ← M_0, z_0 ← +∞
    while t < T and r < T_ftol do
        z ← min M and x̄ ← argmin M
        Check for violated inequalities by solution x̄.
        Sort inequalities by violation.
        Compute cuts for the most violated inequalities.
        Add computed cuts if they are not ϵ-parallel to existing cuts in M.
        Drop cuts of age ≥ T_age whose slack is ≥ ϵ_j.
        if z − z_0 < z_0 · ϵ_ftol then
            r ← r + 1
        else
            r ← 0
        end if
        z_0 ← z
    end while
end procedure
```

In addition to M_0 and the parameters $p_\tau, \epsilon, \epsilon_{par}, T_{age}$, other inputs for our procedure are: a time limit $T > 0$; the number of admissible iterations without sufficient objective improvement $T_{ftol} \in \mathbb{N}$; and a threshold for objective relative improvement $\epsilon_{ftol} > 0$.

5.2.2 Multi-period

We outer-approximate the i2+ relaxation in the multi-period setting. Thus, the linearly constrained base model M_0^T is given by

$$[M_0^T]: \qquad \min \quad \sum_{t=0}^{T-1} \sum_{k \in \mathcal{G}} F_k(P_{k,t}^g) \tag{25a}$$

subject to:

$$\text{constraints (1b), (1c), (6), (1h) – (1i), (1j)} \tag{25b}$$

$$\text{for every good-i2 branch } \{k,m\},\ (29),\ (9) \tag{25c}$$

$$\text{for every bad-i2 branch } \{k,m\},\ (19). \tag{25d}$$

In other words, we consider all the linear constraints of our i2+ relaxation.

For each time-period t, we sort violated inequalities $\tau \in \{\text{Jabr}, \text{i2}, \text{limit}\}$ by violation, as in the single-period setting, and pick as (τ, t)-candidates branches the top p_τ percentage of the most violated branches per period t. The rest follows as in Algorithm 1.

Our experiments show that our cut selection is robust, and that Gurobi version 10.0.1 [41] is able to handle surprisingly well large volume of cuts, see tables 10 11, 12.

6 On the Accuracy of Lower Bounds

In this section we address an issue of fundamental importance in numerical optimization, and which, in our opinion, has not received adequate attention in the ACOPF literature.

Most optimization solvers, commercial and academic, work with finite precision, floating-point arithmetic (some notable exceptions are the LP solvers [4] and [38]) and are subject to roundoff errors, usually small. To be precise, whenever an optimization instance is "solved" and declared "optimal" by a solver, it is highly likely that the provided solution will have small infeasibilities, i.e., the solution will exhibit constraint violations up to some tolerance. Usually, this *infeasibility tolerance* can be controlled by the user.

A natural question then is: given an ϵ-feasible $\tilde{x}$ solution to a convex optimization problem $[P]$ with optimal solution $\overline{x}$ (likely unknown by the user), how *superoptimal* can the approximate solution be? In other words, are there any general guarantees which permit us bound the superoptimality of approximate solutions? For example, if f denotes the objective function for $[P]$, we seek a guarantee of the form

$$f(\tilde{x}) \geq f(\overline{x}) - h_P(\epsilon)$$

for a certain function h_P which only depends on the instance's data. Ideally, we would like for h_P to be monotonic on ϵ and polynomial on the size (bit-length) of the problem $[P]$ data. If $[P]$ is a linear program then h_P is linear function whose slope coefficient has bit-length polynomial on the problem's bit encoding [64]. Moreover, if $[P]$ is a convex QP we also have a good guarantee as in the LP case. Our proof leverages an argument used to show that QP is in NP [70].

Proposition 7 *Let $[P]$ be a convex QP with objective given by $f(x) := x^\top Hx + c^\top x$, where $H \in \mathbb{Q}^{n\times n}$ is a positive definite matrix and $c \in \mathbb{Q}^n$, and feasible set $\mathcal{X} := \{x \in \mathbb{R}^n : Ax \geq b\}$, where $A \in \mathbb{Q}^{m\times n}$ and $b \in \mathbb{Q}^m$. Let $\overline{x} \in \mathbb{Q}^n$ be an optimal solution to this convex QP. Suppose there is a point $\tilde{x} \in \mathbb{Q}^n$ that is ϵ-feasible for $[P]$. Then,*

$$f(\tilde{x}) \geq f(\overline{x}) - ||\overline{\lambda}||_1 \epsilon. \tag{26}$$

where $\overline{\lambda}$ is an optimal solution to the dual of $[P]$.

Moreover, $||\overline{\lambda}||_1$ can be bounded by a constant $g(A, b, H, c)$ whose bit-length is polynomial on the bit-complexity of the input (A, b, H, c).

Proof See Appendix I.

Remark 2 This proposition can be readily applied to the dual of the convex QP which is again a convex QP (the quadratic term in the objective is negative definite). Indeed, if we denote by q the dual objective and by $\tilde{\lambda}$ an ϵ-feasible dual solution, we can apply our result to the convex QP with objective $q' := -q$ which yields

$$q'(\tilde{\lambda}) \geq q'(\overline{\lambda}) - ||\overline{x}||_1 \epsilon \iff q(\tilde{\lambda}) \leq q(\overline{\lambda}) + ||\overline{x}||_1 \epsilon.$$

Next, we elaborate on the relevance of Proposition 7. Suppose we are given:

- A (not necessarily convex) optimization problem $[Z]$ whose optimal objective value is denoted by $\overline{z}$;
- A convex relaxation $[P]$ of $[Z]$, with objective function f, and its dual $[D]$, with objective function q;
- An ϵ-feasible primal-dual pair $(\tilde{x}, \tilde{\lambda})$ and an optimal primal-dual pair $(\overline{x}, \overline{\lambda})$ for $[P]$-$[D]$.

Assume that strong duality holds[8] for the pair of convex problems $[P]$-$[D]$. Figure 4 depicts two inherent risks when dealing with ϵ-feasible solutions: incurring in (1) invalid lower bounds, and (2) poor primal bounds for $[Z]$. Indeed, (1) an *infeasible* (even if ϵ-feasible) dual solution could provide an invalid lower bound for $[Z]$, i.e., $\overline{z} < q(\tilde{\lambda})$. On the other hand, (2) if a primal ϵ-solution $\tilde{x}$ to the relaxation $[P]$ is used as an approximate or even an exact solution[9] to $[Z]$ and we do not have guarantees such as those provided by Proposition 7, then the value $f(\tilde{x})$ may be significantly off (far from $\overline{z}$).

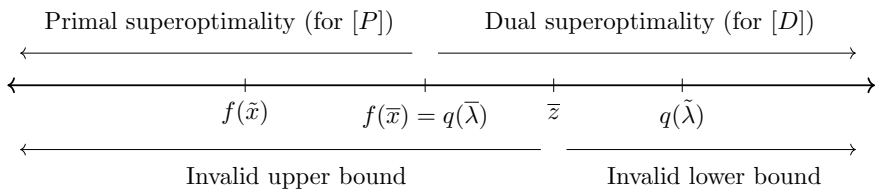

Fig. 4: Validity of bounds and superoptimality of ϵ-feasible solutions.

Results such as inequality (26) are appealing since the bound has a linear dependence on ϵ. Unfortunately, small dual violations are not the only ingredient for accurate bounds; problem structure is key. In Appendix J we show that if $[P]$ and $[D]$ are a primal-dual pair of SOCPs for which strong duality holds, then any ϵ-feasible solution $\tilde{x}$ satisfies $c^\top \tilde{x} \geq c^\top \overline{x} - (||\overline{\lambda}||_1 + ||\overline{u}||_1)\epsilon$ where c denotes the primal objective and $(\overline{x}, (\overline{\lambda}, \overline{u}))$ a primal-dual optimal pair. Thus, a natural question is: can we always provide a 'reasonable' bound on $||\overline{\lambda}||_1 + ||\overline{u}||_1$? From a fundamental perspective, such a bound should be

[8] An analogous analysis can be done is there exists a duality gap; we assumed strong duality for clarity of explanation.

[9] For instance, for radial networks the Jabr SOCP relaxation is exact.

polynomial on the bit-length of the problem data – as in the case for LPs or convex QPs[10]. In what follows we provide a partially negative answer to the latter question even for the case of the Jabr SOCP for the simplest possible network – two buses and one transmission line.

Proposition 8 *There exists an ACOPF instance on a two-bus network with standard line parameters and rational data, where the only two AC feasible solutions are irrational. Additionally, the unique minimizer of the Jabr SOCP in this instance is also irrational. Furthermore, the dual problem to this Jabr SOCP has a unique maximizer, which is irrational as well.*

Proof See Appendix H.

Remark 3 This is an example of an ACOPF instance (in the simplest possible non-trivial setting) whose feasible region is disconnected while its convex hull is described by the feasible region of the Jabr SOCP (since the network is radial [46]). Also see the classical example in [44]. Moreover, this example shows that to answer whether AC-Feasibility [53,18] is in NP or not one cannot rely on an AC rational feasible solution as a polynomial-size certificate.

In summary, LP and convex QP relaxations possess robust theoretical guarantees for bounding challenging optimization problems, in sharp contrast to what nonlinear relaxations such as SOCPs can offer. Given that this paper focuses on lower bounds for ACOPF, we are particularly interested in dual infeasibility of primal-dual solutions provided by the solvers (the primal being a convex relaxation of ACOPF). In Section 7.3 we report on dual infeasibilities attained by solvers run on our linearly constrained models.

7 Computational Experiments

We ran all of our experiments on an Intel(R) Xeon(R) 64-bit Linux machine, with 2 E5-2687W v3 3.10GHz CPUs, 20 physical cores, 40 logical processors, and 256 GB RAM. We used three commercial solvers: Gurobi version 10.0.1 [41], Artelys Knitro versions 13.2.0 and 14.1.0 beta [23], and Mosek 10.0.43 [5]. As a common interface for all SOCP and ACOPF models we used AMPL [37] via Python. We note that unlike Gurobi and Knitro, Mosek does not accept a constraint of the form $x^2 + y^2 \leq z^2$ or $x^2 + y^2 \leq wz$ as a conic constraint; thus we reformulated the SOCPs using a format that Mosek-AMPL was able to read.

We report extensive numerical experiments on instances from the following data sets: medium and large-sized standard IEEE instances available from MATPOWER [75], the Pan European Grid Advanced Simulation and State Estimation (PEGASE) project [47] [36], ACTIVSg synthetic cases developed as part of the US ARPA-E GRID DATA research project [20], [21], and the

[10] The existence of SOCPs whose unique solution is irrational are folklore results, see [17] for examples.

largest instances from the Power Grid Library for Benchmarking AC Optimal Power Flow Algorithms [6] (PGLIB). Our main focus are cases with 9000 buses or more.

Our cutting-plane algorithm is implemented in Python 3 and uses Gurobi 10.0.1 to solve LPs and convex QPs. All of our reported single-period experiments were obtained with the following parameter setup (Section 5.2): barrier convergence tolerance and absolute feasibility and optimality tolerances equal to 10^{-6}, $\epsilon = 10^{-5}$, $p_{Jabr} = 0.55$, $p_{i2} = 0.15$, $p_{limit} = 1$, $T_{age} = 5$, $\epsilon_{par} = 10^{-5}/2$, $\epsilon_{ftol} = 10^{-5}$, and $T_{ftol} = 5$. For our multi-period experiments we used the numerically better-behaved i2+(ρ) relaxation with $\rho = 10^2$, and, as a result, we were able to increase solution accuracy: we set absolute feasibility and optimality tolerances equal to 10^{-8}, and we relaxed T_{ftol} from 5 to 3 for $T = 12$ and to 2 for $T = 24$. All of our code, AMPL model files, and solution files can be downloaded from www.github.com/matias-vm.

Next we describe the symbols used in our tables. The character " – " denotes that the solver did not converge, while "TLim" means that the solver did not converge within our time limit. By convergence we mean that the solver *declares* to have obtained an *optimal* solution, within the previously defined tolerances. Further, "INF" means that the instance was declared infeasible by the solver, while "LOC INF", used by KNITRO, means that the solver converged to an infeasible point. Moreover, if Gurobi declares *numerical trouble* while solving our LPs or convex QPs at some iteration of our cutting-plane algorithm, we report the objective value of the previous iteration followed by the character " $*$ ". The objective values and running times are reported with 2 decimal places.

We remark that, to the best of our knowledge, this is the first computational study which compares the performance of three leading commercial solvers on the Jabr SOCP using a common framework (AMPL). We evaluate the *solvers* on the Jabr SOCP, and compare their performance with that of our warm-started i2+ formulations. For the solvers we use Jabr instead of the i2 SOCP because the former is numerically better behaved from the solvers' perspective (Section 4.5). We report on these numerical issues in Section 7.1.2; also compare tables 11 and 12.

7.1 Single-Period Experiments

SOCPs: Solver parameters We set a time limit of 1,000 seconds for all of our SOCP experiments[11]. Next we describe additional parameter specifications for each solver.

- We use Gurobi's default homogeneous self-dual embedding interior-point algorithm (barrier method without *Crossover*, and *Bar Homogeneous* set to 1), and we set the parameter *Numeric Focus* equal to 1. Barrier convergence tolerance and absolute feasibility and optimality tolerances were set to 10^{-6}. Gurobi was allowed 20 threads.

[11] One iteration of our algorithm requires approximately $60 - 100$ seconds.

- We use Knitro's default Interior-Point/Barrier Direct Algorithm, with absolute feasibility and optimality tolerances equal to 10^{-6}. We used the HSL MA57 sparse symmetric indefinite linear solver, and the Intel Math Kernel Library (MKL) functions for Basic Linear Algebra Subroutines (BLAS), i.e., for basic vector and matrix computations. Moreover, 20 threads were used with Knitro. SOCPs were explicitly input to Knitro as convex problems. When computing primal bounds, we employed the linear solver HSL MA97 whenever Knitro under MA57 was not converging.
- We use Mosek's default homogeneous and self-dual interior-point algorithm for conic optimization. We set the relative termination tolerance, as well as primal and dual feasibility tolerances to 10^{-6}. We used 20 threads with Mosek.

Remark 4 We did not set a time limit for computing single-period ACOPF primal bounds.

7.1.1 Cut Computations

Table 5 summarizes outcomes on cut computations for a substantial number of instances from the libraries described above. Under "Cutting-Plane," "Objective" reports the objective of the last iteration of our algorithm, "Time" reports the total running time (in seconds) of our method; "Computed" reports the number of cuts computed throughout the entire procedure; "Added" exhibits the total number of cuts in our linearly constrained relaxation at the last round; and "Rnd" is the number of rounds or our algorithm. Finally, "Primal bound" reports the objective value of a feasible solution to ACOPF and the amount of time (in seconds) required by Knitro for this purpose.

Overall, our cut management heuristics yield very tight linearly constrained relaxations using a relatively small number of cuts - note that we could potentially add $3|\mathcal{E}|$ cuts per round (for each branch $\{k, m\}$ there are three inequalities (20) and (22) that might be violated). For instance, case ACTIVSg70k has 88207 branches and after 10 rounds of cuts we only keep 123431 out of the 350572 linear cuts computed throughout the course of our algorithm. Thus, fewer than 1.5 linear cuts per branch give us a relaxation with *optimality gap*[12] at most 0.69%.

We remark that for some instances the objective value of our procedure can be higher than the objective value of the Jabr SOCP since our algorithm is outer-approximating the feasible region of the i2 SOCP (Proposition 1).

We also note that 30000goc-sad was the only instance for which Gurobi experienced numerical difficulties while solving the linearly constrained relaxation (indicated by the character "$*$" next to the objective value). In this case the reported objective value corresponds to the previous iteration – setting a

[12] Given a primal bound of a minimization problem, we define the optimality gap of a relaxation of the given problem as $\frac{z_p - z_r}{z_p}$, where z_p denotes the objective value of the primal bound and z_r denotes the objective value of the relaxation.

Table 5: Performance of Cutting-Plane Algorithm (Not Warm-Started)

	Cutting-Plane Algorithm					Primal bound	
Case	Objective	Time	Computed	Added	Rounds	Objective	Time
9241pegase	309221.81	378.82	135599	29875	23	315911.56	96.74
9241pegase-api	6924650.57	277.32	128316	30230	21	7068721.98	73.85
9241pegase-sad	6141202.28	386.51	113686	27273	21	6318468.57	33.92
9591goc-api	1346373.10	187.26	87812	22469	22	1570263.74	42.85
9591goc-sad	1055493.25	246.87	90153	20514	27	1167400.79	28.15
ACTIVSg10k	2476851.62	132.16	60803	18183	19	2485898.75	76.71
10000goc-api	2502026.03	147.12	73084	19666	24	2678659.51	23.46
10000goc-sad	1387303.02	114.97	58984	18528	17	1490209.66	103.06
10192epigrids-api	1849488.30	152.87	97921	24882	22	1977687.11	117.15
10192epigrids-sad	1672819.53	185.02	95740	23726	23	1720194.13	23.74
10480goc-api	2708819.18	200.48	114967	29805	21	2863484.4	38.71
10480goc-sad	2287314.69	270.38	118122	28004	24	2314712.14	27.93
13659pegase	379084.55	841.83	176962	37297	22	386108.81	1184.15
13659pegase-api	9270988.77	326.57	147479	34390	19	9385711.45	44.43
13659pegase-sad	8868216.24	301.87	130682	32662	19	9042198.49	42.08
19402goc-api	2448812.41	440.67	213564	52388	22	2583627.35	87.33
19402goc-sad	1954047.79	488.33	218291	49749	25	1983807.59	64.01
20758epigrids-api	3042956.88	464.17	189436	46124	25	3126508.30	61.39
20758epigrids-sad	2612551.03	379.36	180790	44624	24	2638200.23	58.11
24464goc-api	2560407.12	471.14	226595	57162	22	2683961.9	533.03
24464goc-sad	2605128.51	506.39	222908	55242	23	2653957.66	73.87
ACTIVSg25k	5993266.85	592.39	156285	43851	28	6017830.61	56.69
30000goc-api	1531110.84	464.16	142385	41840	24	1777930.63	134.71
30000goc-sad	1130733.51*	147.74	76546	76546	6	1317280.55	565.05
ACTIVSg70k	16326225.66	1065.76	350572	123431	13	16439499.83	240.55
78484epigrids-api	15877674.54	1007.99	556893	240576	10	16140427.68	1079.03
78484epigrids-sad	15175077.19	1062.55	501202	313587	8	15315885.86	343.45

more aggressive cut management heuristic, for instance decreasing T_{age} from 5 to 4, gave us numerically more stable cuts and a better bound.

Finally, we obtained our primal bounds by running Knitro with a *flat-start*, i.e., we provided as initial point voltage magnitudes set to 1 and all $\theta_{km} = 0$.

7.1.2 Performance on SOCPs

In Table 6, we observe that for the cases in which at least two solvers converge, the reported bounds for the Jabr SOCP agree on the first 2 to 3 most significant digits. These differences in bounds reflect how numerically challenging the given instances are. We remind the readers of the parameter choices that we made in order for the solvers to achieve termination – the solvers otherwise would often fail.

As we mentioned at the beginning of this section, the i2 SOCP is numerically even more challenging for the solvers than the Jabr SOCP. Indeed, in Table 7 we can see that the solvers do struggle (see also Table 12). We studied in detail some cases where Gurobi AMPL declared optimality –for example case ACTIVSg70k– while reporting variable bound max scaled violation equal to 8.43 as well as large primal and dual residuals (0.0128 and

Table 6: Solvers' Performance on Jabr SOCP

	Objective			Time		
Case	Gurobi	Knitro	Mosek	Gurobi	Knitro	Mosek
9241pegase	-	309234.16	-	82.11	34.68	31.11
9241pegase-api	-	6840612.84	-	116.32	23.39	72.29
9241pegase-sad	-	6083747.85	-	111.05	26.01	75.99
9591goc-api	1346480.71	1348107.89	1345869.72	38.25	23.74	36.60
9591goc-sad	1055698.54	1058606.56	1054379.58	49.29	32.83	37.61
ACTIVSg10k	-	2468172.93	2466666.10	40.18	21.48	26.08
10000goc-api	-	2507034.94	2498948.00	48.63	35.19	30.13
10000goc-sad	1387288.49	1388679.63	1386041.07	23.58	26.27	23.68
10192epigrids-api	-	1849684.14	1848873.47	75.82	42.69	29.09
10192epigrids-sad	-	1672989.96	1672534.72	83.85	28.33	28.63
10480goc-api	-	2708973.58	2707828.26	75.94	27.21	56.82
10480goc-sad	-	2286454.3	2285547.23	149.93	38.17	59.48
13659pegase	379135.73	379144.11	-	33.61	43.26	34.92
13659pegase-api	-	9198542.14	-	162.21	30.64	105.11
13659pegase-sad	8826902.31	8826958.23	8787429.86	83.75	31.84	108.74
19402goc-api	-	2449020.25	2447799.72	158.12	152.89	103.04
19402goc-sad	-	1954331.70	1952550.06	203.56	155.89	104.88
20758epigrids-api	-	-	3040421.02	143.99	TLim	93.46
20758epigrids-sad	-	-	2610196.94	98.30	TLim	75.88
24464goc-api	2548335.96	-	2558631.63	603.95	TLim	129.90
24464goc-sad	-	-	2603525.46	333.50	TLim	128.50
ACTIVSg25k	5956787.54	5964417.54	5955368.56	169.66	87.14	87.18
30000goc-api	-	1531256.65	1529197.81	207.60	118.80	123.38
30000goc-sad	-	-	1130868.71	191.22	TLim	84.90
ACTIVSg70k	-	16221577.73	16217263.66	553.26	320.98	232.47
78484epigrids-api	-	-	-	756.00	TLim	637.48
78484epigrids-sad	15180775.21	-	15169401.54	463.17	TLim	601.04

3.25, respectively). Moreover, we noticed inconsistent termination status for cases 10192epigrids-sad, 10480goc-api, 20758epigrids-sad, and 30000goc-sad on Gurobi and Gurobi through AMPL (Gurobi-AMPL) using the same model; Gurobi AMPL declares optimal termination for these instances while Gurobi does not. Cases for which we were able to identify *low quality* solutions or inconsistencies have been denoted with the character "†" next to their reported objective value in Table 7.

7.2 Warm-Started Single-Period Instances

As described in the introduction, in power engineering practice it is of interest to address problem instances where a limited change in data occurs after solving a previous problem. In the context of this paper, we can therefore assume that we have a *warm-started formulation*, i.e., one where we leverage previously computed cuts. Here we present this warm-starting feature of our algorithm; we justify its validity and show via numerical experiments its appealing lower bounding capabilities.

Table 7: Solvers' Performance on i2 SOCP

	Objective			Time		
Case	Gurobi	Knitro	Mosek	Gurobi	Knitro	Mosek
10192epigrids-api	1849683.44	1849684.2	-	37.44	19.4	30.74
10192epigrids-sad	1672998.72†	1672998.73	-	23.36	21.48	24.53
10480goc-api	2709110.52†	2709110.71	-	41.44	27.78	35.28
10480goc-sad	2287736.73	2287715.33	-	41.11	28.46	28.48
13659pegase	379142.67	-	-	52.12	TLim	36.49
13659pegase-api	9287242.7	9287244.72	-	66.39	236.65	30.01
13659pegase-sad	8878803.69	-	-	63.11	TLim	30.48
19402goc-api	2449100.15†	2449102.05	-	79.38	55.6	54.18
19402goc-sad	1954367.11†	1954367.2	-	180.97	59.46	82.55
20758epigrids-api	-	3043275.95	-	79.03	64.17	56.02
20758epigrids-sad	2612841.71†	2612841.8	-	48.32	84.07	58.87
24464goc-api	2560829.65†	-	-	132.34	TLim	88.79
24464goc-sad	2605532.65†	-	-	74.43	916.1	65.41
ACTIVSg25k	5994727.45	-	-	70.61	TLim	52.38
30000goc-api	1531320.78	1531322.2	-	96.91	593.73	71.38
30000goc-sad	1132242.88†	1132256.94	-	78.0	325.11	74.61
ACTIVSg70k	16333807.38†	-	-	300.98	TLim	209.3
78484epigrids-api	15882668.49	15882668.46	15882654.42	216.15	315.31	203.81
78484epigrids-sad	15180792.15	15180792.0	15180763.6	250.43	376.82	222.17

We note that the convex inequalities (20), which we use to develop cuts, do not depend on input data such as loads or operational limits. Any such cut remains valid and can be used if the associated branch remains operational. This will be our strategy, below.

We created two kinds of perturbed instances: **a)** Instances were the load of each bus was perturbed by a Gaussian $(\mu, \sigma) = (0.01 \cdot P_d, 0.01 \cdot P_d)$, where P_d denotes the original load, subject to the newly perturbed load being non negative; and **b)** instances were the transmission line which carries the largest amount of active power in an ACOPF solution is turned off. With regards to cases of type a), we focus on relatively small individual load changes to model re-optimization after a relatively small time change. In the multi-period setting, below, we consider much larger load changes, to capture hourly models. Cases of type b) do change the structure of the network and in their warm-start we removed any cuts associated with the switched-off branch.

Tables 8 and 9 summarize our warm-started experiments on perturbed instances from our data set in Table 5 and compare our algorithm and solvers' performance on the Jabr SOCP. "First Round" reports the objective value and running time of the relaxation M_0 loaded with the cuts computed in Table 5, i.e., our warm-started relaxation. Moreover, under "Cutting-plane", we report on the objective value (at termination) and the total running time of our cutting-plane procedure (on the warm-started relaxation). "Jabr SOCP" and "Primal bound" report, respectively, on the objective value and running time of the Jabr SOCP for the three solvers, and ACOPF primal solutions.

We stress that after just *one* round, the bound provided by our algorithm is already quite good – and also point out the comparison between the running time for our first round, and the solvers' running time.

Loads perturbed – Gaussian deltas. For most of these cases, our procedure proves very tight lower bounds in less than 25 seconds ("First Round" column). Judging by the time required by Knitro and by the number of cases in which the solvers converge running the SOCPs, it appears that these instances are, overall, more challenging than their unperturbed counterparts.

Our procedure also stands out in quickly lower bounding the largest cases. For instance, a very sharp bound for case ACTIVSg70k is obtained in 102.25 seconds, taking less than half of the time it takes the fastest SOCP solver to converge. Similar performance is achieved on the largest epigrids cases where our method is 3 to 5 times faster.

An interesting empirical fact is that our cuts are robust with respect to load perturbations. Indeed, our evidence shows that there is not a considerable improvement from the "First Round" to the "Last Round" objectives. Thus, the previously pre-computed cuts loaded to M_0 in the first iteration are, already, accurately outer-approximating the SOC relaxations.

Our linearly constrained relaxations are able to prove infeasibility for the modified case 9241pegase-api in 23.10 seconds while none of the three solvers were able to provide a certificate of infeasibility for the Jabr SOCP. Knitro required 1845.42 seconds to declare convergence to a locally infeasible solution. Similar results are obtained for case 24464ogc-api.

The only case where our method fails to provide a valid lower bound is case 30000goc-sad – our minimization oracle reports numerical trouble and fails to provide a solution to our warm-started relaxation. This is not surprising since difficult numerical behavior was noticed when computing cuts for this case.

Transmission line with largest flow switched off. Overall, our method achieves a similar performance on this set of perturbed instances as in a); sharp lower bounds are obtained in about 25 seconds for most of the cases.

For this data set, our method and all of the SOCP solvers are able to prove infeasibility relatively quickly. On the other hand, our method proves a lower bound for ACTIVSg70k relatively quickly in the first round, but fails to converge in the next round due to numerical trouble caused by the newly added cuts.

As when perturbing loads, our warm-started formulation achieves a good performance on the largest epigrid cases – bounds are sharp with respect to the SOC relaxations and it is at least 3x faster.

7.3 Multi-Period Results

Data. At present, there is a dearth of publicly-available, realistic multi-period AC load and generation data. One partial source is the set of cases developed

Table 8: Warm-Started Relaxations, Loads perturbed Gaussian $(\mu, \sigma) = (0.01 \cdot P_d, 0.01 \cdot P_d)$

	Cutting-Plane				Jabr SOCP							
	First Round				Objective			Time			Primal bound	
Case	Objective	Time	Objective	Time	Gurobi	Knitro	Mosek	Gurobi	Knitro	Mosek	Objective	Time
9241pegase	309288.32	13.78	309299.97	160.28	-	309302.67	-	73.12	32.21	36.04	315979.53	101.48
9241pegase-api	INF	23.10	INF	23.10	-	-	-	134.53	TLim	72.96	LOC INF	1845.92
9241pegase-sad	6153913.91	16.18	6154117.59	136.78	-	6096743.03	-	97.51	26.07	83.43	6333763.92	43.71
9591goc-api	1343642.47	11.06	1343670.62	56.36	1343767.43	1345384.57	1343190.29	39.36	25.36	35.30	1571582.59	54.16
9591goc-sad	1058124.48	12.62	1058157.44	65.37	1058337.76	1061275.83	1057323.31	51.85	34.04	37.52	1178895.53	29.53
ACTIVSg10k	2475041.43	9.52	2475078.69	50.51	-	2466383.20	-	42.31	21.75	29.33	2484093.15	57.24
10000goc-api	2502049.28	8.51	2502098.01	36.48	2501946.30	2507074.78	2499373.75	31.91	43.44	32.33	LOC INF	1677.21
10000goc-sad	1388833.86	8.70	1388859.09	44.50	1388824.91	1390230.41	1387588.17	25.96	29.31	23.67	1493481.44	93.72
10192epigrids-api	1848085.36	10.27	1848133.48	45.84	-	1848285.26	1847120.93	65.38	41.17	25.99	LOC INF	1458.35
10192epigrids-sad	1672358.89	10.33	1672398.61	53.37	-	1672533.02	1671364.67	73.64	28.61	35.66	1717429.36	23.89
10480goc-api	2704157.29	12.43	2704252.95	58.45	-	2704373.73	2703432.85	197.17	27.57	55.92	2868495.28	36.89
10480goc-sad	2294908.37	12.81	2294990.69	70.93	-	2294080.35	2292830.56	185.22	35.90	58.31	2322198.81	27.34
13659pegase	379742.62	60.74	379794.51	426.88	379799.37	379804.43	-	34.21	43.17	32.75	386765.25	370.23
13659pegase-api	9253539.07	21.25	9253773.43	109.20	9181205.93	9181269.20	-	97.11	30.41	118.31	9368277.57	62.20
13659pegase-sad	8865733.59	21.28	8865892.49	113.04	8824442.20	8824486.03	-	86.49	33.19	102.59	9039904.52	40.02
19402goc-api	2452185.69	23.55	2452270.83	120.10	-	2452448.33	2451708.50	146.87	120.39	103.32	LOC INF	4440.99
19402goc-sad	1956255.19	23.28	1956313.91	113.89	-	1956570.60	1955018.07	231.90	172.82	102.19	1986936.95	66.02
20758epigrids-api	3043006.76	22.34	3043076.56	104.06	-	-	3032919.24	134.60	TLim	78.32	LOC INF	12425.89
20758epigrids-sad	2610197.53	20.46	2610261.88	93.09	-	-	2608090.26	143.69	TLim	72.19	2635892.81	49.25
24464goc-api	2561680.14	26.28	INF	50.38	-	LOC INF	-	223.07	573.37	118.6	-	19444.54
24464goc-sad	2606391.76	26.78	2606473.78	133.54	-	-	2604708.86	423.12	TLim	128.84	2655942.01	72.48
ACTIVSg25k	5988886.18	28.24	5989016.75	198.58	5952404.50	5960068.30	5949381.04	138.01	73.75	109.39	6013477.05	57.87
30000goc-api	1527412.96	25.35	1527487.45	151.75	-	1528338.73	1525625.64	243.61	369.83	119.92	LOC INF	3407.47
30000goc-sad	-	46.33	-	46.33	-	-	1132715.53	257.94	TLim	75.20	1318389.55	620.27
ACTIVSg70k	16316572.42	102.25	16317886.35	536.51	-	16210682.53	16206290.43	498.80	309.56	229.07	16428367.50	243.84
78484epigrids-api	15862318.24	115.76	15865624.98	883.93	-	-	15859950.52	757.64	TLim	642.24	-	8113.53
78484epigrids-sad	15176866.00	151.77	15180592.27	1118.02	15182602.75	-	15174716.43	420.56	TLim	589.46	15316872.94	353.13

Table 9: Warm-Started Relaxations, Transmission Line with largest flow turned off

	Cutting-Plane				Jabr SOCP							
	First Round				Objective			Time			Primal bound	
Case	Objective	Time	Objective	Time	Gurobi	Knitro	Mosek	Gurobi	Knitro	Mosek	Objective	Time
9241pegase	INF	10.86	INF	10.86	INF	INF	INF	8.36	7.55	10.22	INF	8.37
9241pegase-api	INF	7.55	INF	7.55	INF	INF	INF	7.92	8.02	10.22	INF	8.23
9241pegase-sad	INF	7.33	INF	7.33	INF	INF	INF	8.14	8.10	10.31	INF	8.46
9591goc-api	1346470.95	10.42	1346859.06	60.76	1346969.75	1348591.44	1346437.99	39.30	17.98	36.89	1395829.51	28.08
9591goc-sad	1055823.53	11.51	1056267.57	101.64	1056447.48	1059382.10	1055501.31	45.09	35.41	37.18	1199276.44	29.90
ACTIVSg10k	2477043.05	9.94	2477537.79	75.85	-	2468821.96	2466981.35	44.81	21.60	17.52	LOC INF	7092.45
10000goc-api	2506671.15	8.06	2509971.69	46.10	2509846.00	2514991.16	2506236.75	31.03	33.95	32.15	2692320.35	23.28
10000goc-sad	1387382.65	8.76	1387515.89	66.14	1387480.33	1388870.68	1386283.75	26.53	34.24	24.14	1506187.88	108.19
10192epigrids-api	1849901.82	9.47	1850621.81	68.73	-	1850788.76	1849821.44	69.81	38.01	25.30	2021493.05	117.18
10192epigrids-sad	1673575.50	11.08	1674274.99	74.49	-	1674417.21	1673564.57	69.91	43.91	28.54	1734014.50	24.11
10480goc-api	2710040.46	11.33	2711100.23	73.85	-	2711224.27	2710520.15	95.40	27.99	56.58	2862699.50	225.06
10480goc-sad	2288069.64	13.47	2288969.47	98.88	-	2288069.23	2286864.08	106.54	37.65	59.43	2318279.76	26.13
13659pegase	379102.13	53.29	379163.58	199.96	379177.99	379182.14	-	33.18	345.83	31.90	386126.93	394.93
13659pegase-api	INF	10.81	INF	10.81	INF	INF	INF	10.94	8.90	13.79	INF	11.55
13659pegase-sad	INF	10.63	INF	10.63	INF	INF	INF	11.02	10.80	13.76	INF	11.73
19402goc-api	2450110.09	23.93	2451621.60	171.31	-	2451793.39	2450488.01	154.42	132.03	104.58	2587915.50	403.20
19402goc-sad	1954365.39	23.70	1954881.06	191.52	-	1955116.35	1953676.05	258.93	156.10	102.63	1985954.83	63.38
20758epigrids-api	3043482.21	20.44	3044690.46	133.92	-	-	3041974.74	112.78	TLim	96.19	3132571.31	52.82
20758epigrids-sad	2612646.70	20.62	2612786.37	115.89	-	-	2610315.41	169.67	TLim	73.02	2638560.64	47.87
24464goc-api	2560669.11	25.94	2561110.03	161.80	2550118.22	-	2559240.55	440.81	TLim	118.05	2684708.93	1663.63
24464goc-sad	2605179.75	26.98	2605369.03	166.81	-	2605474.23	2603609.34	564.27	74.69	124.19	2654344.45	76.39
ACTIVSg25k	6045885.88	27.52	6048122.86	238.67	6009656.52	6018875.03	6009500.57	144.28	65.42	81.41	LOC INF	1634.38
30000goc-api	1531110.55	25.02	1531159.95	130.30	-	1532013.41	1529195.12	195.74	135.71	119.77	LOC INF	3203.26
30000goc-sad	-	45.60	-	45.60	-	-	1130917.78	218.12	TLim	74.27	1324622.71	186.43
ACTIVSg70k	16426522.74	98.51	16426522.74*	98.51	-	-	-	150.77	TLim	129.60	LOC INF	3160.81
78484epigrids-api	15888353.48	104.32	15892229.11	¡916.88	15894055.55	-	15880422.78	322.22	TLim	625.74	16169740.92	2328.79
78484epigrids-sad	15179882.22	149.65	15185980.69	1151.33	15188085.99	-	15182701.65	437.59	TLim	594.84	15330674.69	272.41

for the GO3 Competition [45][13]. Another partial source is the data set from the optimization package "UnitCommitment.jl" for the SCUC problem [73]; however, the largest instance contains less than 14,000 nodes, which is not among the largest publicly available instances. We are also aware that creating *good* ACOPF data is a non-trivial task [54], [20], but given that there were no publicly available multi-period load time-series to match the standard largest AC cases we developed our own data sets[14].

We created and worked with three data sets: (1) $T = 4$ with increasing[15] (on expectation) active power loads, to be precise, at each time-period $t \in \{1, 2, 3\}$, $P^d_{k,t}$ is set to $(1 + 0.01 \cdot t + u_{k,t}) \cdot P^d_k$, where $u_{k,t}$ is a uniformly distributed random variable with range $[0, 0.025]$; $T = 12$ where a half-a-day active power load curved is simulated, see Figure 5; and (3) $T = 24$ where a 24 hr active power load curve is simulated, see Figure 6. Moreover, we assumed 50% generator ramp up r^u and down r^d rates for consecutive time-periods (see Section 2.1) in all of the three settings.

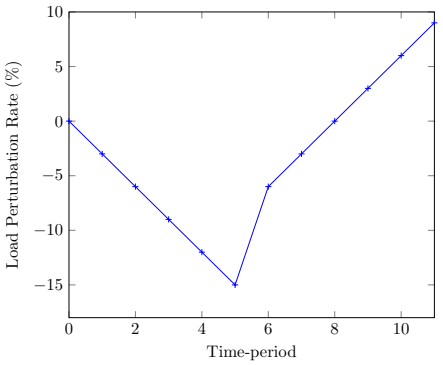

Fig. 5: 12hr load curve.

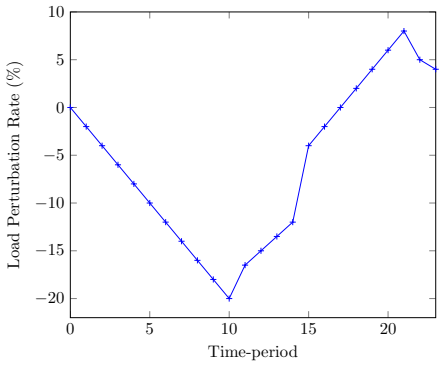

Fig. 6: 24hr load curve.

Upper bound heuristics. In order to evaluate how tight a given lower bound is, an upper bound is needed. We relied on two simple heuristics to compute upper bounds. The first heuristic, which we term H1, consists in providing to Knitro the complete multi-period formulation. The parameter setup for H1 is: absolute feasibility tolerance equal to 10^{-6}, while we set relative optimality tolerance to 10^{-3}; simple variables bounds are enforced throughout the optimization; and $T_{ftol} = 3$; and time limit of 2400 seconds for $T = 4$, and $7,200$ seconds for $T = 12, 24$. If H1 fails, then we proceed with heuristic H2 which solves

[13] The GO3 Competition does not provide explicit load data, since a welfare maximization problem needs to be solved to determine the active and reactive power consumption. We purposely avoided this dependence on the solution of another challenging optimization problem.

[14] These data sets can be downloaded from www.github.com/matias-vm.

[15] Given that the PGLIB library [6] contains congested instances, we applied decreasing (on expectation) active power loads perturbations to the PGLIB instances.

one single-period ACOPF instance at a time, and upon convergence, fixes the upper and lower limits of active power generation of the next instance as a function of generator outputs in the current period. To be precise, for every generator $k \in \mathcal{G}$ and any time-period $t \in \{0, \ldots, T-2\}$, and provided that the nonlinear solver converges for the single-period instance at time t, then the output of generator k for the time-period $t+1$ is bounded by $P^{min}_{k,t+1} \leq P^{g}_{k,t+1} \leq P^{max}_{k,t+1}$ where $P^{min}_{k,t+1} := (1 - r^{d}_{k,t+1})P^{g}_{k,t}$ and $P^{max}_{k,t+1} := (1 + r^{u}_{k,t+1})P^{g}_{k,t}$. If the solver fails to find a solution for some time-period, then we declare failure of H2. Overall, H2 succeeded in most of the instances for which we were able to obtain a primal bound. The parameter setup for H2 was the same as for H1 except that we set the time limit for single-period ACOPFs to $1,200$ seconds.

SOCPs: Solver parameters for multi-period instances. With respect to our SOCP experiments, we set a time limit of $2,400$ seconds for $T = 4$ instances, $4,800$ seconds for $T = 12$, while we set $8,000$ seconds for $T = 24$ instances. In what follows we describe the parameter specifications for each solver for the SOCP runs.

- We use Gurobi's default homogeneous self-dual embedding interior-point algorithm (barrier method without *Crossover*, and *Bar Homogeneous* set to 1), and we set the parameter *Numeric Focus* equal to 1. Barrier convergence tolerance and absolute feasibility and optimality tolerances were set to 10^{-6}. Gurobi was allowed 20 threads.
- We use Knitro's default Interior-Point/Barrier Direct Algorithm, with absolute feasibility and optimality tolerances equal to 10^{-6}. We used the Intel MKL PARDISO sparse symmetric indefinite solver, and the Intel Math Kernel Library (MKL) functions for Basic Linear Algebra Subroutines (BLAS), i.e., for basic vector and matrix computations. Moreover, we gave Knitro 20 threads to use for parallel computing features (10 threads were given to the linear solver). When solving the SOCPs, we explicitly instructed Knitro to handle the problem as convex, and to enforce simple variable bounds throughout the optimization.
- We use Mosek's default homogeneous and self-dual interior-point algorithm for conic optimization. We set the relative termination tolerance, as well as primal and dual absolute feasibility tolerances to 10^{-6}. We assigned 20 threads to Mosek.

We recall the reader that we use the i2+(ρ) relaxation, with $\rho = 100$, as target set for our cutting-plane multi-period experiments. For our cutting-plane algorithm, when $T = 4$ a $1,200$ seconds time limit was enforced prior to *starting* a new round of cutting; the time limit was $3,600$ seconds for $T = 12$, and $6,000$ seconds for $T = 24$.

7.3.1 Not Warm-Started Cutting-Plane Algorithm: $T = 4$

Table 10 summarizes our not warm-started cutting-plane experiments on multi-period instances with $T = 4$. Columns "# Vars", "# Cons", and "FTime"

denote the number of variables, number of constraints and formulation time of the base model M_0^T. Under "First Iteration" we report on the objective value, running time, and maximum dual infeasibility error (see Section 6) of the relaxation M_0^T, i.e., without any cuts. On the other hand, "Last Iteration" presents the objective value, running time, total number of cuts held by the model M_{last}^T on the last iteration of the algorithm, and maximum dual infeasibility error. "Total" presents the total running time of our algorithm and the total number of rounds of cuts, and "ACOPF" describes the number of variables, number of constraints, and the objective value and running time of each ACOPF primal solution.

Overall, we can see that lower bounds are fairly tight; even our base model M_0^T without any cuts, whose objective value is shown under "First Iteration", attains quick and sharp lower bounds for these multi-period instances. For instance, an optimality gap of 6.91% is achieved by M_0^T for case ACTIVSg70k, and, after three rounds of cuts, nearly 800,000 cuts drive the gap down to 1.75%.

Our cut management heuristics prove to be empirically successful: dual infeasibility reported by Gurobi is either maintained or slightly worsened (compare "DInfs" of "First Iteration" versus that of "Last Iteration") after adding a significant number of cuts. This means that our cuts are keeping our linearly constrained relaxation numerically stable and that our reported lower bounds are reasonably accurate (Section 6). Moreover, we observe that instances which require a significant number of rounds of cuts, say 10 or more, under the given time limit, notoriously benefit from our cut clean-up heuristic. We see that a small number of cuts (roughly 10% of the total number of constraints) is good enough to outer-approximate the rotated-cones and attain sharp ACOPF lower bounds (as in the single-period case; Section 7.1.1).

7.3.2 Warm-Started Cutting-Plane Algorithm: T = 4, 12, 24

We use a simple methodology to warm-start the algorithm described in Section 5.2.2: we propagate cuts for the original instances to every time-period. To be precise, we create a pool that includes, for each branch $\{k, m\} \in \mathcal{E}$, all the associated single-period cuts – which are assumed to have been previously computed (Section 7.1.1). The cuts in the pool are propagated to every time-period $t \in \{0, \ldots, T-1\}$ and are added to the base model M_0^T.

In Tables 11 and 12 we present the results of our warm-started experiments with $T = 4$. In Table 11 we compare our lower bounds (obtained by warm-starting our base model with precomputed cuts) against the lower bounds obtained by solving the Jabr SOCPs with each of the three commercial solvers. In Tables 12, 16, and 17 we compare our bounds against the performance of the solvers on the i2+(ρ) relaxation plus Jabr inequalities for every branch whose i2 inequality is bad; this relaxation is numerically better behaved than the i2 SOCP (see discussion in Section 4.5).

In both tables the columns "# Vars", "# Cons", and "FTime" denote the number of variables, number of constraints and formulation time of the

Table 10: Cutting-Plane Not Warm-Started, $T = 4$ and loads perturbed $\text{Unif}(0, 0.025) + 0.01 \cdot t$ for $t \in \{1, 2, 3\}$, ramping rates 50%

	Cutting-Plane												ACOPF			
				First Iteration			Last Iteration				Total				PBound	
Case	#Vars	#Cons	FTime	Obj	Time	DInfs	Obj	Time	#Cuts	DInfs	Time	Rnds	#Vars	#Cons	Obj	Time
1354pegase	66897	72564	1.35	299460.87	2.33	5.80e-10	303439.90	2.31	11043	3.32e-09	37.32	12	53772	70220	303727.82	7.60
2869pegase	150431	161568	3.31	543415.33	6.16	1.23e-07	549381.91	7.83	26938	6.00e-10	154.55	14	120712	158388	549988.85	92.99
3375wp	148217	159692	3.18	29810390.36	6.79	5.13e-08	30318794.35	10.92	16457	1.67e-07	171.70	16	117364	151844	30410712.08	63.05
6468rte	319517	323396	6.57	349845.41	13.77	1.29e-04	355995.14	11.83	46902	2.72e-06	354.29	15	247296	321888	356969.89	2596.76
9241pegase	524028	538308	10.93	1250007.83	27.06	1.92e-07	1269609.76	258.55	318602	5.63e-05	1403.50	7	412248	543530	1297205.96	172.18
ACTIVSg10k	456684	490360	9.98	10035162.94	24.82	7.07e-09	10222977.88	27.19	52082	2.10e-08	399.43	13	363940	470558	10265810.55	228.28
10000goc-api	465676	495636	10.04	8916264.44	24.76	1.50e-09	9495292.75	20.41	61212	5.28e-07	492.16	19	368928	478650	10019708.68	1049.69
10000goc-sad	465676	495636	10.03	5373568.66	28.48	8.39e-07	5456260.77	29.84	61388	1.12e-06	403.37	11	368928	478650	5866821.40	241.57
13659pegase	723245	735740	15.46	1528552.09	31.02	1.29e-08	1551553.61	422.84	469991	9.30e-07	1380.08	6	567716	739636	1584844.81	3083.66
13659pegase-api	723245	735740	15.71	34787242.54	39.17	2.84e-05	35703095.45	207.32	316521	2.24e-08	1207.94	9	567716	739636	36163923.41	1288.60
13659pegase-sad	723245	735740	15.7	33913060.97	40.32	1.46e-08	34547415.27	97.98	243868	9.98e-08	1218.18	10	567716	739636	35199898.78	545.24
20758epigrids-api	1106271	1127168	23.52	11540550.41	175.25	3.38e-08	11834277.30	373.71	453981	2.33e-06	1294.19	5	859924	1131168	12088587.29	1018.17
20758epigrids-sad	1106271	1127168	24.30	10060842.01	86.20	4.97e-08	10219196.77	306.29	410645	1.22e-07	1294.96	7	859924	1131168	10317105.84	709.85
ACTIVSg25k	1145783	1212392	24.89	24097056.99	80.20	1.93e-08	24710933.54	120.53	239603	9.12e-08	1228.25	8	902588	1170088	24835784.33	451.45
30000goc-api	1253699	1366264	27.91	4916231.92	73.67	5.00e-07	5705500.71	82.73	112230	1.31e-06	1238.19	12	990172	1280368	6563873.25	1669.13
30000goc-sad	1253699	1366264	27.62	4151693.26	87.16	7.53e-06	4340911.74	84.31	104180	7.58e-07	1213.47	11	990172	1280368	-	3666.85
ACTIVSg70k	3115935	3316972	68.93	65302843.92	269.38	1.85e-08	68917427.01	389.02	799462	2.03e-08	1268.46	3	2448820	3175256	70148300.24	2034.37
78484epigrids-api	4179320	4225032	89.55	60235793.42	347.62	1.49e-04	61087514.34	677.62	1496304	1.51e-04	1816.21	3	3230648	4252514	62011780.96	4831.70
78484epigrids-sad	4179320	4225032	89.43	57816998.73	387.48	1.23e-04	58689459.71	571.38	1081287	1.18e-04	1361.42	2	3230648	4252514	-	-

base model M_0^T (before adding the previously computed cuts). 'First Iteration" reports on the objective value, running time, the number of cuts with which the instance was warm-started, and the maximum dual infeasibility error (Section 6) of the warm-started relaxation M_T^0. Multi-column "Last Iteration" presents the objective value, running time, total number of cuts held by the model M_{last}^T on the last iteration of the algorithm, and maximum dual infeasibility error of model M_{last}^T. "Total" presents the total running time of our algorithm and the total number of rounds of cuts. Under the header "Jabr SOCP", columns "#Vars" and "#Cons" denote the number of variables, number of constraints of the Jabr SOCP (which are the same for the three solvers given that we used a common AMPL model file), and "Obj" and "Time" show the objective value and running time of the Jabr SOCP for each solver. Finally, under column "ACOPF" we report objective value of an ACOPF primal solution.

As in the single-period experiments (Tables 8 and 9), "First Iteration" shows that our warm-started relaxation, i.e., model M_0^T loaded with the propagated precomputed cuts, attains sharp bounds fairly quickly. A notable example, in which at least 2 solvers converge in Table 11, is case ACTIVSg25k for which Gurobi solves the convex QP in 94 seconds resulting in 2.1% optimality gap. Moreover, after 9 rounds of cuts our algorithm is able to drive the down the optimality gap to 0.5% at very high accuracy (Figure 8).

Dual infeasibility in (11), as in the not warm-started experiments, is not significantly worsened by our cuts, i.e., they do not seem to ill-condition our instances. On the contrary, we see that for some instances, our cut management heuristics are able to improve on this metric. In Figure 7 we illustrates the evolution of the number of outer-envelope cuts with respect to iterations of our cutting-plane algorithm for two $T = 4$ warm-started cases. For instance, ACTIVSg10k was warm-started with 28,184 cuts, and peaked 111,323 cuts at iteration 6 which corresponds to the chosen threshold T_{age} (Section 5.2.1). From the 7th iteration onward, there is a steady cut-cleansing ending up with $51,188$ cuts. As expected, both instances end up with more cuts than were present when warm-started.

With respect to the solvers' performance on the SOCPs, we see that all of the solvers struggle. In Table 11 we see differences up to the second digit in the objective values of instances on which at least 2 solvers converge, which corroborates that these are numerically challenging instances. Table 12 exhibits even more striking results; here we ran the solvers on the i2+(ρ) relaxation (the relaxation that we are outer-approximating with our cutting-plane algorithm).

In Appendix K, we present tables 16 and 17 which show computational results for our $T = 12$ and $T = 24$ instances. We see that Gurobi is able to robustly handle our very large LPs and convex QPs, such as case ACTIVSg70k, with almost 19 million variables and 20 million constraints. For this instance, we load 1.5 million precomputed cuts, and after two rounds of cuts we end up with 3.35 million cuts, and the solver outputs a lower bound with very small dual infeasibilities. For these large instances, we were only able to compute an

AC primal bound for case 1354pegase $T = 12$; our heuristics failed in all the other instances with $T = 12$ and $T = 24$.

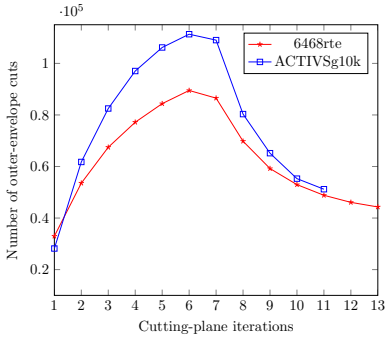

Fig. 7: Cut management for $T = 4$ warm-started instances.

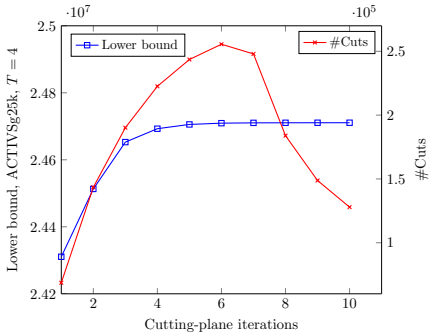

Fig. 8: Lower bounds for case ACTIVSg25k and number of cuts.

On tables 13 and 14 we compare our linearly-constrained warm-started relaxation against multi-period DCOPF[16]. We report on the number of variables and constraints for each instance, as well as the objective, total running time, and dual infeasibility[17]. We remind the reader that DCOPF is a linearly-constrained approximation of ACOPF, not a relaxation [9]; hence, DCOPF infeasibility does not necessarily imply that the associated ACOPF is infeasible, in sharp contrast to what any relaxation for ACOPF guarantees.

Our ACOPF relaxations quadruple the size (in terms of number of variables and constraints) of the corresponding multi-period DCOPF instances. Gurobi clearly runs faster on multi-period DCOPF. Interestingly, there are some DCOPF instances for which Gurobi does not converge[18] – either because of suboptimal termination or because it runs into numerical trouble – while it does converge for our ACOPF relaxations attaining small dual infeasibilities.

7.4 Stability of Dual Variables for Active-Power Balance Constraints

In current energy pricing rules [61,40], dual variables associated with active-power balance constraints (1b) of a welfare maximization problem constitute a key input for computing Locational Marginal Prices (LMPs). In general, linear

[16] Multi-period DCOPF is solved with Gurobi using the same parameter setup as our cutting-plane algorithm.

[17] For our relaxation, we report on the last iteration of our cutting-plane algorithm. Moreover, the number of constraints includes the total number of cuts kept during the last iteration.

[18] We remind the reader that there are alternative formulations of DCOPF. Rather than directly solving an LP or convex QP, such alternative approaches could significantly enhance computational efficiency by exploiting problem structure (deploying a delayed constrained generation algorithm where the decision variables are nodal injections).

Table 11: Warm-Started Relaxation versus Jabr SOCP, $T = 4$ and loads perturbed $\text{Unif}(0, 0.025) + 0.01 \cdot t$ for $t \in \{1, 2, 3\}$, ramping rates 50%

	Cutting-Plane												Jabr SOCP								ACOPF
				First Iteration				Last Iteration				Total			Obj			Time			
Case	#Vars	#Cons	FTime	Obj	Time	#Cuts	DInfs	Obj	Time	#Cuts	DInfs	Time	Rnds	#Vars	#Cons	Gurobi	Knitro	Mosek	Gurobi	Knitro	Mosek	PBound
1354pegase	66897	72564	1.41	301493.85	3.42	6000	2.55e-08	303438.51	1.93	10345	2.95e-08	36.04	12	55280	66149	-	303420.78	303415.82	15.34	11.65	14.79	303727.82
3375wp	148217	159692	3.17	30132477.23	8.76	10448	1.42e-06	30318337.40	8.59	15786	4.44e-07	124.03	13	118128	142915	-	-	30233269.55	60.00	TLim	36.2	30410712.08
6468rte	319517	323396	6.81	352645.45	18.24	32960	2.68e-06	355986.22	11.76	42631	2.72e-06	205.53	13	252240	302520	356099.86	-	355981.27	66.75	TLim	57.88	356969.89
9241pegase	524028	538308	11.36	1258422.42	48.99	68472	3.61e-06	1270138.75	100.47	189019	5.18e-08	1268.90	7	433700	515921	1262979.52	-	-	182.32	TLim	89.64	1297205.96
ACTIVSg10k	456684	490360	10.27	10117471.23	30.65	28184	5.72e-09	10222978.04	23.94	51188	5.01e-08	306.72	10	364824	437972	-	9894331.31	9876639.32	243.69	357.69	98.62	10265810.55
10000goc-api	465676	495636	9.88	9129427.52	28.01	35856	6.70e-09	9495973.72	20.46	59211	3.92e-07	301.67	11	383978	467735	-	-	9479414.94	304.47	TLim	132.55	10019708.68
10000goc-sad	465676	495636	10.21	5414194.06	30.94	35224	1.29e-06	5456194.0	26.47	53520	1.31e-06	336.21	10	383978	467735	-	-	5448060.51	316.29	TLim	127.01	5866821.40
13659pegase	723245	735740	15.67	-	365.52	97036	-	-	365.52	97036	-	384.05	0	578580	686433	1488477.49	-	-	308.57	TLim	100.23	1584844.81
13659pegase-api	723245	735740	15.28	35162915.42	74.46	85584	2.42e-06	35703716.41	178.18	218790	4.06e-07	1348.20	8	594950	719171	35070776.73	-	-	717.60	TLim	479.87	36163923.41
13659pegase-sad	723245	735740	15.60	34093210.18	52.16	78612	2.89e-08	34547590.80	98.42	136983	5.16e-07	1273.65	12	594950	719171	34196027.45	-	33517709.2	523.11	TLim	466.32	35199898.78
20758epigrids-api	1106271	1127168	23.34	-	195.65	80904	5.54	11833435.09	658.47	260055	9.29e-05	1575.77	3	918078	1105639	-	-	11830133.74	1093.22	TLim	322.49	12088587.29
20758epigrids-sad	1106271	1127168	23.64	10138699.01	92.13	77228	2.76e-08	10219195.88	181.38	185780	8.84e-08	1374.00	8	918078	1105639	-	-	10211969.97	398.31	TLim	497.46	10317105.84
ACTIVSg25k	1145783	1212392	25.46	24310520.83	93.95	68676	1.01e-07	24710662.67	92.3	127912	1.61e-08	1206.14	9	912168	1099187	23293596.46	-	23640456.52	768.78	TLim	395.86	24835784.33
30000goc-api	1253699	1366264	30.24	5185673.79	97.36	59732	3.54e-07	5706550.06	80.15	105806	9.79e-07	1146.25	12	1013234	1246465	-	-	5637560.70	TLim	TLim	644.65	6563873.25
30000goc-sad	1253699	1366264	28.03	4316717.55	138.08	179704	7.84e-07	4341119.26	77.37	91071	1.06e-06	994.70	9	1013234	1246465	-	-	4325675.85	985.49	TLim	569.21	-
ACTIVSg70k	3115935	3316972	68.7	67721489.41	350.58	258228	3.04e-08	69422696.57	369.05	597271	1.11e-07	1301.25	3	2480088	3003929	-	-	65293153.18	TLim	TLim	925.00	70148300.24
78484epigrids-api	4179320	4225032	96.42	60947456.94	494.72	625780	1.53e-04	61271585.98	563.69	1047809	1.52e-04	1523.35	2	3420774	4126501	-	-	-	TLim	TLim	1109.78	62011780.96
78484epigrids-sad	4179320	4225032	97.61	58930587.37	743.6	888916	1.20e-04	59051305.24	845.85	1133203	1.20e-04	1569.46	1	3420774	4126501	59124882.16	-	-	1586.33	TLim	1103.22	-

Table 12: Warm-Started Relaxation versus i2 SOCP+, $T = 4$ and loads perturbed $\text{Unif}(0, 0.025) + 0.01 \cdot t$ for $t \in \{1, 2, 3\}$, ramping rates 50%

	Cutting-Plane												i2 SOCP+								ACOPF
				First Iteration				Last Iteration				Total			Obj			Time			
Case	#Vars	#Cons	FTime	Obj	Time	#Cuts	DInfs	Obj	Time	#Cuts	DInfs	Time	Rnds	#Vars	#Cons	Gurobi	Knitro	Mosek	Gurobi	Knitro	Mosek	PBound
1354pegase	66897	72564	1.41	301493.85	3.42	6000	2.55e-08	303438.51	1.93	10345	2.95e-08	36.04	12	56326	84155	303424.53	-	-	41.94	TLim	7.86	303727.82
3375wp	148217	159692	3.17	30132477.23	8.76	10448	1.42e-06	30318337.40	8.59	15786	4.44e-07	124.03	13	121822	179665	-	-	-	60.00	TLim	15.28	30410712.08
6468rte	319517	323396	6.81	352645.45	18.24	32960	2.68e-06	355986.22	11.76	42631	2.72e-06	205.53	13	269070	373246	355828.72	-	-	468.73	TLim	41.89	356969.89
9241pegase	524028	538308	11.36	1258422.42	48.99	68472	3.61e-06	1270138.75	100.47	189019	5.18e-08	1268.90	7	451546	643811	-	-	-	411.55	TLim	92.33	1297205.96
ACTIVSg10k	456684	490360	10.27	10117471.23	30.65	28184	5.72e-09	10222978.04	23.94	51188	5.01e-08	306.72	10	364824	437972	-	-	-	243.69	TLim	79.58	10265810.55
10000goc-api	465676	495636	9.88	9129427.52	28.01	35856	6.70e-09	9495973.72	20.46	59211	3.92e-07	301.67	11	390042	567215	-	-	-	559.76	TLim	148.77	10019708.68
10000goc-sad	465676	495636	10.21	5414194.06	30.94	35224	1.29e-06	5456194.0	26.47	53520	1.31e-06	336.21	10	390042	567215	-	-	-	419.79	TLim	104.86	5866821.40
13659pegase	723245	735740	15.67	-	365.52	97036	-	-	365.52	97036	-	384.05	0	618066	859791	-	-	-	457.45	TLim	117.18	1584844.81
13659pegase-api	723245	735740	15.28	35162915.42	74.46	85584	2.42e-06	35703716.41	178.18	218790	4.06e-07	1348.20	8	618066	859791	-	-	-	1093.04	TLim	155.84	36163923.41
13659pegase-sad	723245	735740	15.60	34093210.18	52.16	78612	2.89e-08	34547590.80	98.42	136983	5.16e-07	1273.65	12	618066	859791	-	-	-	891.76	TLim	150.4	35199898.78
20758epigrids-api	1106271	1127168	23.34	-	195.65	80904	5.54	11833435.09	658.47	260055	9.29e-05	1575.77	3	950270	1340191	-	-	-	1529.93	TLim	356.82	12088587.29
20758epigrids-sad	1106271	1127168	23.64	10138699.01	92.13	77228	2.76e-08	10219195.88	181.38	185780	8.84e-08	1374.00	8	950270	1340191	-	-	-	1397.93	TLim	247.43	10317105.84
ACTIVSg25k	1145783	1212392	25.46	24310520.83	93.95	68676	1.01e-07	24710662.67	92.3	127912	1.61e-08	1206.14	9	961854	1387817	-	-	-	589.52	TLim	200.22	24835784.33
30000goc-api	1253699	1366264	30.24	5185673.79	97.36	59732	3.54e-07	5706550.06	80.15	105806	9.79e-07	1146.25	12	1018714	1524129	-	-	-	1761.14	TLim	286.2	6563873.25
30000goc-sad	1253699	1366264	28.03	4316717.55	138.08	179704	7.84e-07	4341119.26	77.37	91071	1.06e-06	994.70	9	1018714	1524129	-	-	-	727.32	TLim	247.6	-
ACTIVSg70k	3115935	3316972	68.7	67721489.41	350.58	258228	3.04e-08	69422696.57	369.05	597271	1.11e-07	1301.25	3	2584042	3765747	-	-	-	1312.16	TLim	405.88	70148300.24
78484epigrids-api	4179320	4225032	96.42	60947456.94	494.72	625780	1.53e-04	61271585.98	563.69	1047809	1.52e-04	1523.35	2	3558322	4997073	-	-	-	TLim	TLim	1117.23	62011780.96
78484epigrids-sad	4179320	4225032	97.61	58930587.37	743.6	888916	1.20e-04	59051305.24	845.85	1133203	1.20e-04	1569.46	1	3558322	4997073	-	-	-	TLim	TLim	1117.35	-

Table 13: Warm-Started Relaxation versus DCOPF, $T = 12$

	Warm-Started Relaxation					DCOPF				
Case	#Vars	#Cons	Obj	Time	DInfs	#Vars	#Cons	Obj	Time	Dinfs
1354pegase	201209	250008	860664.27	92.63	1.06e-08	62369	67828	850023.48	2.44	1.72e-10
9241pegase	1574972	1980693	3565252.49	3811.42	4.32e-04	447608	477952	3627501.88	18.06	3.76e-10
ACTIVSg10k	1375020	1639853	28690135.28	1308.80	3.38e-07	449628	501812	28187885.77	19.50	2.34e-07
10000goc-api	1401204	1686928	28459989.11	1066.39	2.40e-08	446364	490232	INF	20.15	-
10000goc-sad	1401204	1647684	16474320.51	1450.12	3.36e-06	446364	490232	16114986.72	19.96	2.21e-07
ACTIVSg25k	3447015	4472632	69534077.06	4073.97	9.75e-07	1097943	1199456	67992943.20	48.82	7.88e-07
30000goc-api	3768147	4484763	17079448.81	3667.40	1.06e-05	1225815	1299860	-	75.87	-
30000goc-sad	3768147	4373646	13190466.44	3901.76	2.45e-06	1225815	1299860	12846216.92	53.33	7.65e-08
ACTIVSg70k	9368583	11913750	188736994.5	4702.30	3.66e-08	2977455	3195644	-	175.35	-
78484epigrids-api	12551704	14607420	INF	447.39	-	3555448	3699780	INF	4.19	-
78484epigrids-sad	12551704	16185978	177335668.05	5465.67	1.36e-04	3555448	3699780	-	175.91	-

Table 14: Warm-Started Relaxation versus DCOPF, $T = 24$

	Warm-Started Relaxation					DCOPF				
Case	#Vars	#Cons	Obj	Time	DInfs	#Vars	#Cons	Obj	Time	Dinfs
1354pegase	402677	503633	1666609.97	180.61	2.97e-08	124997	136696	1646661.83	4.61	2.15e-10
9241pegase	3151388	4613025	6925029.02	8419.66	1.25e-03	896660	961684	7013508.62	36.02	5.25e-10
ACTIVSg10k	2752524	3290672	55184650.09	2869.25	7.18e-08	901740	1013564	54241448.25	37.77	5.44e-08
10000goc-api	2804496	3403822	53800966.46	2273.58	3.18e-07	894816	988820	INF	40.25	-
10000goc-sad	2804496	3262316	32678842.78	3099.33	7.97e-06	894816	988820	32120062.03	41.02	1.97e-06
ACTIVSg25k	6898863	8756843	134142620.00	6991.16	2.08e-06	2200719	2418248	131325272.36	95.18	5.69e-09
30000goc-api	7539819	9332691	31616476.13	6799.92	2.06e-05	2455155	2613824	-	171.42	-
30000goc-sad	7539819	9661317	25811075.96	6249.53	1.23e-05	2455155	2613824	25279287.29	108.86	2.22e-08
ACTIVSg70k	18747555	23469443	357083446.24	8496.71	2.98e-08	5965299	6432848	-	829.66	-
78484epigrids-api	25110280	24855336	-	2138.78	-	7117768	7427052	INF	380.67	-
78484epigrids-sad	25110280	30188832	344805634.00	5399.62	1.34e-04	7117768	7427052	-	513.08	-

models such as DCOPF are used for pricing because of their computational tractability. The trade-offs of using such approximations are well-known. As stated in [12], "If DCOPF does not adequately reflect the physics of the power grid, the prices retrieved from the DCOPF solution will not be accurate. Thus, tighter relaxations of ACOPF could lead to prices that better reflect scarcity in the physical network." Among their findings, the authors empirically show that nonlinear tighter relaxations such as the Jabr SOCP or the QC relaxation mitigate biasedness in price signals, in particular in congested networks, and lead to higher welfare.

Given the theoretical and empirical evidence presented in this paper, we believe that our linearly constrained relaxation is well-placed as a potential substitute for DCOPF in electricity market pricing. A first step in evaluating its potential is to understand whether the dual variables associated to active-power balance constraints of our dynamically built linearly constrained relaxations numerically converge to some vector of duals. In Table 15 and Figure 9 below we provide preliminary evidence of numerical convergence for medium-sized $T = 4$ instances. At each iteration $k > 1$ of our cutting-plane algorithm, we obtain the corresponding vector of duals λ^k and compute its Euclidean distance to its predecessor λ^{k-1}.

8 Conclusions and Future Work

In this paper we present a fast linear cutting-plane method used to obtain tight relaxations for even the largest single and multi-period ACOPF instances, by appropriately outer-approximating the SOC relaxations. Our relaxations

Table 15: On the convergence of dual variables of active-power balance constraints $T = 4$

	$\|\|\lambda^k - \lambda^{k-1}\|\|_2$								
Case	$k=2$	$k=4$	$k=6$	$k=8$	$k=10$	$k=12$	$k=14$	$k=16$	$k=18$
1354pegase	42.5161	9.2024	2.9110	2.0099	1.0506	0.4326	0.189	0.1038	0.0494
ACTIVSg20000	2590.985	481.1567	133.0116	67.344	38.3696	13.2228	7.9999	2.8572	2.5617
2869pegase	84.6365	14.806	5.7585	3.4687	2.4777	0.9821	0.8868	0.4770	1.8608
6468rte	232.8488	42.6396	14.7121	7.3318	4.1112	1.5908	1.2435	0.3554	0.6279
ACTIVSg10k	6394.2841	1712.8186	475.1414	104.5385	54.3611	17.6938	8.4489	7.7985	4.2688

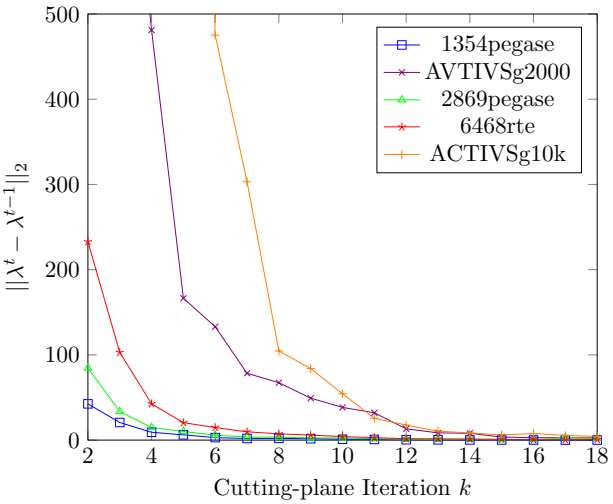

Fig. 9: Distance between dual variables of active-power balance constraints of consecutive cutting-plane iterations.

can be constructed and solved robustly and quickly via a cutting-plane algorithm that relies on proper cut management and leverages mature linear programming technology. Moreover, we provide a theoretical justification for the tightness of two SOC relaxations for ACOPF as well as for the use of our linearly-constrained relaxations.

The central focus on this paper concerns *reoptimization*. We show that our procedure possesses efficient warm-starting capabilities – previously computed cuts, for some given instance, can be re-utilized and loaded into new runs of *related* instances, hence leveraging previous computational effort. As a main contribution we demonstrate, through extensive numerical testing in medium to (very) large multi-period instances, that the warm-start feature for our cutting-plane algorithm yields tight and accurate bounds far faster than otherwise possible. It is worth noting that this capability stands in contrast to what is possible using nonlinear (convex) solvers.

We bring forth to the ACOPF literature a discussion on approximate solutions to convex relaxations to ACOPF and their accuracy as lower bounds. We show that linearly-constrained relaxations with convex quadratic objective

possess robust theoretical guarantees for bounding challenging optimization problems, in sharp contrast to what nonlinear relaxations such as SOCPs can offer.

We believe our work paves the way for promising new research directions. For instance, since our relaxations are linear they could be deployed for practical pricing schemes which could increase welfare and mitigate biasedness in price signals [12]. Moreover, we believe our relaxation can be used for harder multi-period ACOPF features such as unit-commitment or security constraints, hence it would be interesting to evaluate its performance on these challenging problems. We are also interested in developing new heuristics for finding multi-period ($T = 12$ and 24 time-periods) ACOPF solutions for large-scale networks exploiting solutions to our linearly-constrained relaxations.

Acknowledgements

We would like to thank Erling Andersen, Bob Fourer, Ed Klotz and Richard Waltz.

Sun.Jan.26.094450.2025@freestyle

References

1. Ahuja, R., Magnanti, T., Orlin, J.: Network Flows: Theory, Algorithms, and Applications. Pearson (1993)
2. Alguacil, N., Conejo, A.: Multiperiod optimal power flow using benders decomposition. IEEE Transactions on Power Systems **15**, 196–201 (2000)
3. Alizadeh, M., Capitanescu, F.: A tractable linearization-based approximated solution methodology to stochastic multi-period ac security-constrained optimal power flow. IEEE Transactions on Power Systems **38**, 5896–5908 (2023)
4. Applegate, D.L., Cook, W., Dash, S., Espinoza, D.: Exact Solutions to Linear Programming problems. Operations Research Letters **35**, 693–699 (2007)
5. ApS, M.: Mosek for AMPL User's Guide (2024)
6. Babaeinejadsarookolaee, S., Birchfield, A., Christie, R., Coffrin, C., DeMarco, C., Diao, R., Ferris, M., Fliscounakis, S., Greene, S., Huang, R., Josz, C., Korab, R., Lesieutre, B., Maeght, J., Mak, T., Molzahn, D., Overbye, T., Panciatici, P., Park, B., Snodgrass, J., Tbaileh, A., Van Hentenryck, P., Zimmerman, R.: The power grid library for benchmarking ac optimal power flow algorithms. arXiv:1908.02788v2 (2021)
7. Bai, X., Wei, H., Fujisawa, K., Wang, Y.: Semidefinite Programming for Optimal Power Flow problems. International Journal of Electrical Power and Energy Systems **30**, 383–392 (2008)
8. Bai, Y., Zhong, H., Xia, Q., Kang, C., Xie, L.: A decomposition method for network-constrained unit commitment with ac power flow constraints. Energy **88** (2015)
9. Baker, K.: Solutions of DC OPF Are Never AC Feasible. In: Proceedings of the Twelfth ACM International Conference on Future Energy Systems, p. 264–268. Association for Computing Machinery (2021)
10. Ben-Tal, A., Nemirovski, A.: On Polyhedral Approximations of the Second-Order Cone. Mathematics of Operations Research **26**(2), 193–205 (2001)
11. Bergen, A.R., Vittal, V.: Power Systems Analysis. Prentice-Hall (1999)
12. Bichler, M., Knörr, J.: Getting Prices right on Electricity Spot Markets: On the Economic Impact of Advanced Power Flow models. Energy Economics **126** (2023)

13. Bienstock, D.: Electrical Transmission System Cascades and Vulnerability, an Operations Research viewpoint. Society for Industrial and Applied Mathematics (2015)
14. Bienstock, D., Escobar, M., Gentile, C., Liberti, L.: Mathematical Programming Formulations for the Alternating Current Optimal Power Flow problem. 4OR **18**(3), 249–292 (2020)
15. Bienstock, D., Munoz, G.: On Linear Relaxations of OPF Problems. arXiv:1411.1120 pp. 1–15 (2014)
16. Bienstock, D., Munoz, G.: Approximate method for AC transmission switching based on a simple relaxation for ACOPF problems. In: 2015 IEEE Power & Energy Society General Meeting (PESGM), pp. 1–5. IEEE (2015)
17. Bienstock, D., Pia, A., Hildebrand, R.: Complexity, Exactness, and Rationality in Polynomial Optimization. Mathematical Programming **197** (2023)
18. Bienstock, D., Verma, A.: Strong NP-hardness of AC power flows feasibility. Operations Research Letters **47**(6), 494–501 (2019)
19. Bienstock, D., Villagra, M.: Accurate and Warm-Startable Linear Cutting-Plane Relaxations for ACOPF. https://optimization-online.org/?p=25584 (2024)
20. Birchfield, A., Xu, T., Gegner, K., Shetye, K., Overbye, T.: Grid structural characteristics as validation criteria for synthetic networks. IEEE Transactions on Power Systems **32**, 3258–3265 (2017)
21. Birchfield, A., Xu, T., Overbye, T.: Power flow convergence and reactive power planning in the creation of large synthetic grids. IEEE Transactions on Power Systems **33**, 6667–6674 (2018)
22. Bukhsh, W., Grothey, A., McKinnon, K., Trodden, P.: Local Solutions of the Optimal Power Flow Problem. IEEE Transactions on Power Systems **28**(4), 4780–4788 (2013)
23. Byrd, R., Nocedal, J., Waltz, R.: Knitro: An Integrated Package for Nonlinear Optimization. In: Large-Scale Nonlinear Optimization, vol. 83, pp. 35–59. Springer (2006)
24. Cain, M.B., O'Neill, R.P., Castillo, A.: History of Optimal Power Flow and Formulations. Federal Energy Regulatory Commission (2012)
25. Castillo, A., Laird, C., Silva-Monroy C.A. Watson, J., O'Neill, R.: The unit commitment problem with ac optimal power flow constraints. IEEE Transactions on Power Systems **31** (2016)
26. Castillo, A., Lipka, P., Watson, J., Oren, S., O'Neill, R.: A Successive Linear Programming approach to solving the IV-ACOPF. IEEE Transactions on Power Systems **31**(4), 2752–2763 (2016)
27. Castillo, A., Lipka, P., Watson, J., Oren, S., O'Neill, R.: A successive linear programming approach to solving the iv-acopf. IEEE Transactions on Power Systems **31** (2016)
28. Coffrin, C., Hijazi, H., Van Hentenryck, P.: Distflow Extensions for AC Transmission Systems. arXiv:1411.1120 (2015)
29. Coffrin, C., Hijazi, H., Van Hentenryck, P.: Network Flow and Copper Plate Relaxations for AC Transmission Systems. arXiv:1506.05202 (2015)
30. Coffrin, C., Hijazi, H., Van Hentenryck, P.: Network Flow and Copper Plate Relaxations for AC Transmission Systems. In: Power Systems Computation Conference (PSCC), pp. 1–8 (2016)
31. Coffrin, C., Hijazi, H.L., Van Hentenryck, P.: The qc relaxation: A theoretical and computational study on optimal power flow. IEEE Transactions on Power Systems **31**(4), 3008–3018 (2016)
32. Coffrin, C., Van Hentenryck, P.: A linear-programming approximation of ac power flows. INFORMS Journal on Computing **26**, 718–734 (2014)
33. Constante-Flores, G.E., Conejo, A.J., Qiu, F.: Ac network-constrained unit commitment via relaxation and decomposition. IEEE Transactions on Power Systems **37**, 2187–2196 (2022)
34. Farivar, M., Clarke, C.R., Low, S.H., Chandy, K.M.: Inverter VAR Control for Distribution Systems with Renewables. IEEE International Conference on Smart Grid Communications pp. 457–462 (2011)
35. Farivar, M., Low, S.: Branch Flow Model: Relaxations and Convexification—Part I. IEEE Transactions on Power Systems **28**, 2554–2564 (2013)
36. Fliscounakis, S., Panciatici, P., Capitanescu, F., Wehenkel, L.: Contingency ranking with respect to overloads in very large power systems taking into account uncertainty, preventive and corrective actions. IEEE Trans. on Power Systems pp. 4909–4917 (2013)

37. Fourer, R., Gay, D., Kernighan, B.: AMPL: A Mathematical Programing Language. Algorithms and Model Formulations in Mathematical Programming **51** (1989)
38. Gleixner, A., Steffy, D., Wolter, K.: Iterative refinement for linear programming. INFORMS Journal on Computing **28**, 449–464 (2016)
39. Glover, F.: Improved Linear Integer Programming Formulations of Nonlinear Integer problems. Mgt. Science **22**, 455–460 (1975)
40. Gribik, P., Hogan, W., Pope, S.: Market-Clearing Electricity Prices and Energy Uplift. Harvard Electricity Policy Group (2007)
41. Gurobi Optimization, L.: Gurobi Optimizer Reference Manual (2023)
42. Hedman, K., Oren, S., O'Neill, R.: A Review of Transmission Switching and Network Topology Optimization. IEEE Power and Energy Society General Meeting pp. 1–7 (2011)
43. Higham, N.: Accuracy and Stability of Numerical Algorithms. SIAM (1996)
44. Hiskens, I., Davy, R.: Exploring the Power Flow Solution Space Boundary. IEEE Transactions on Power Systems **16**(3), 389–395 (2001)
45. Holzer, J., Coffrin, C., DeMarco, C., Duthu, R., Elbert, S., Eldridge, B., Elgindy, T., Garcia, M., Greene, S., Guo, N., Hale, E., Lesieutre, B., Mak, T., McMillan, C., Mittelman, H., O'Neill, R., Overbye, T., Palmintier, B., Parker, R., Safdarian, F., Tbaileh, A., Van Hentenryck, P., Veeramany, A., Wangen, S., Wert, J.: Grid Optimization Competition Challenge 3 Problem Formulation. Tech. rep., Pacific Northwest National Laboratory (2023)
46. Jabr, R.: Radial distribution Load Flow using Conic programming. IEEE Transactions on Power Systems **21**(3), 1458–1459 (2006)
47. Josz, C., Fliscounakis, S., Maeght, J., Panciatici, P.: Ac power flow data in matpower and qcqp format: itesla, rte snapshots, and pegase. arXiv:1603.01533 (2016)
48. Kelley, J.E.J.: The Cutting-Plane Method for Solving Convex Programs. Journal of the Society for Industrial and Applied Mathematics pp. 703–712 (1960)
49. Klotz, E.: Identification, assessment, and correction of ill-conditioning and numerical instability in linear and integer programs. Bridging Data and Decisions pp. 54–108 (2014)
50. Kocuk, B., Dey, S., Sun, X.: Strong SOCP Relaxations for the Optimal Power Flow Problem. Operations Research **64**(6), 1177–1196 (2016)
51. Kocuk, B., Dey, S.S., Sun, X.A.: Matrix Minor Reformulation and SOCP-based Spatial Branch-and-Cut method for the AC Optimal Power Flow problem. Mathematical Programming Computation **10**(4), 557–596 (2018)
52. Lavaei, L., Low, S.H.: Zero Duality Gap in Optimal Power Flow problem. IEEE Transactions on Power Systems **27**, 92–107 (2012)
53. Lehmann, K., Grastien, A., Van Hentenryck, P.: AC-Feasibility on Tree Networks is NP-Hard. IEEE Transactions on Power Systems **31**(1), 798–801 (2016)
54. Li, H., Yeo, J., Bornsheuer, A., Overbye, T.: The Creation and Validation of Load Time Series for Synthetic Electric Power Systems. IEEE Transactions on Power Systems **36**, 961–969 (2021)
55. Liu, J., Laird, C., Scott, J., Watson, J., Castillo, A.: Global solution strategies for the network-constrained unit commitment problem with ac transmission constraints. IEEE Transactions on Power Systems **34** (2018)
56. Lorca, A., Sun, X.: The adaptive robust multi-period alternating current optimal power flow problem. IEEE Transactions on Power Systems **33**, 1993–2003 (2017)
57. McCormick, G.: Computability of Global Solutions to Factorable Nonconvex programs: Part I—Convex underestimating problems. Math. Program. **10**, 147–175 (1976)
58. Mhanna, S., Mancarella, P.: An Exact Sequential Linear Programming Algorithm for the Optimal Power Flow Problem. IEEE Transactions on Power Systems **37**(1), 666–679 (2022)
59. Mhanna, S., Verbic, G., Chapman, A.C.: Tight LP Approximations for the Optimal Power Flow problem. In: 2016 Power Systems Computation Conference (PSCC), pp. 1–7. IEEE (2016)
60. Molzahn, D., Hiskens, I.: A Survey of Relaxations and Approximations of the Power Flow Equations. Foundations and Trends in Electric Energy Systems **4**, 1–221 (2019)

61. O'Neill, R., Sotkiewicz P, M., Hobbs, B., Rothkopf, M., Stewart, W.R.: Efficient Market-Clearing Prices in Markets with nonconvexities. European Journal of Operational Research **164**(1), 269–285 (2005)
62. Qiu, G., Tanneau, M., Van Hentenryck, P.: Dual conic proxies for ac optimal power flow. Power Systems Computation Conference (PSCC) (2024)
63. Quan, R., Jian, J., Mu, Y.: Tighter relaxation method for unit commitment based on second-order cone programming and valid inequalities. Electrical Power and Energy Systems **55** (2014)
64. Schrijver, A.: Theory of Linear and Integer Programming. Wiley (1998)
65. Serrano, F., Schwarz, R., Gleixner, A.: On the relation between the extended supporting hyperplane algorithm and Kelley's cutting plane algorithm. J. Global Optimization pp. 161–179 (2020)
66. Shor, N.: Quadratic Optimization Problems. Soviet Journal of Computer and Systems Sciences **25**, 1–11 (1987)
67. Stott, B., Jardim, J., Alsac, O.: Dc power flow revisited. IEEE Transactions on Power Systems **24** (2009)
68. Subhonmesh, B., Low, S., Chandy, K.: Equivalence of Branch Flow and Bus Injection models. In: 2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pp. 1893–1899 (2012)
69. Taylor, J.A., Hover, F.S.: Linear Relaxations for Transmission System Planning. IEEE Transactions on Power Systems **26**(4), 2533–2538 (2011). DOI 10.1109/TPWRS.2011.2145395
70. Vavasis, S.: Quadratic Programming is in NP. Information Processing Letters **36**, 73–77 (1990)
71. Verma, A.: Power Grid Security Analysis : An Optimization Approach. Ph.D. Thesis, Columbia University (2009)
72. Wächter, A., Biegler, L.T.: On the Implementation of a Primal-Dual Interior Point Filter Line Search Algorithm for Large-Scale Nonlinear Programming. Mathematical Programming **106**, 25–57 (2006)
73. Xavier, A.S., Kazachkov, A.M., Yurdakul, O., He, J., Qiu, F.: UnitCommitment.jl: A Julia/JuMP Optimization Package for Security-Constrained Unit Commitment (Version 0.4). Zenodo (2024)
74. Yang, Z., Zhong, H., Bose, A., Zheng, T., Xia, Q., Kang, C.: A Linearize OPF Model with Reactive Power and Voltage Magnitude: A Pathway to Improve the MW-Only DC OPF. IEEE Transactions on Power Systems **33**, 1734–1745 (2018)
75. Zimmerman, R., Murillo-Sanchez, C., Gan, D.: Matpower, A MATLAB Power System Simulation Package. IEE Trans. Power Sys. **26**, 12–19 (2011)
76. Zohrizadeh, F., Josz, C., Jin, M., Madani, R., Lavaei, J., Sojoudi, S.: A Survey on Conic Relaxations of Optimal Power Flow problem. European Journal of Operational Research **287**(2), 391–409 (2020)
77. Zohrizadeh, F., Kheirandishfard, M., Nasir, A., Madani, R.: Sequential relaxation of unit commitment with ac transmission constraints. IEEE Conference on Decision and Control pp. 2408–2413 (2018)

A i2 SOCP strictly stronger than Jabr SOCP

We numerically show that for case 1354pegase its i2 SOCP relaxation is strictly stronger than its Jabr SOCP relaxation. Knitro attains an optimal solution to the Jabr SOCP of value 74009.28 while an optimal solution to the i2 SOCP has value 74013.68. As a sanity check, we fixed, within tolerance $\pm 10^{-5}$, the solution to the Jabr SOCP in the i2 SOCP formulation; Gurobi declared the resulting SOCP infeasible and provided the following Irreducible Inconsistent Subsystem (IIS):

$$
\begin{aligned}
i^{(2)}_{549,5002} &= 11822.45384038167\, v^{(2)}_{549} \\
&\quad + 11822.45384038167\, v^{(2)}_{5002} \\
&\quad - 23644.8888107824\, c_{549,5002} \\
&\quad - 29.87235441454166\, s_{549,5002} \\
v^{(2)}_{549} &\geq 1.209989999283128 \\
v^{(2)}_{5002} &\geq 1.19745626014781 \\
c_{549,5002} &\leq 1.195643087927643 \\
s_{549,5002} &\leq 0.0246578041355137 \\
i^{(2)}_{549,5002} &\leq 39.69.
\end{aligned}
$$

B Definition of the i2 variable

Let the admittance matrix for line $\{km\}$ be

$$
Y := \begin{pmatrix} \left(y + \frac{y^{sh}}{2}\right)\frac{1}{\tau^2} & -y\frac{1}{\tau e^{-j\sigma}} \\ -y\frac{1}{\tau e^{j\sigma}} & y + \frac{y^{sh}}{2} \end{pmatrix} = G + jB
$$

where $y = g + jb$ denotes it's series admittance, shunt admittance is denoted by $y^{sh} = g^{sh} + jb^{sh}$, and $N := \tau e^{j\sigma}$ denotes the transformer ratio of magnitude $\tau > 0$ and phase shift angle σ.

Recall that Ohm's Law states that the current flowing from bus k to m via transmission line $\{km\}$ is given by $\begin{pmatrix} I_{km} \\ I_{mk} \end{pmatrix} = Y \begin{pmatrix} V_k \\ V_m \end{pmatrix}$, hence $I_{km} = \frac{y}{\tau}\left(\frac{1}{\tau}V_k - e^{j\sigma}V_m\right) + V_k\frac{y^{sh}}{2\tau^2}$. Therefore, the modulus of I_{km} squared is

$$
\begin{aligned}
|I_{km}|^2 &= I_{km}I^*_{km} \\
&= \left(\frac{y}{\tau}\left(\frac{1}{\tau}V_k - e^{j\sigma}V_m\right) + V_k\frac{y^{sh}}{2\tau^2}\right)\left(\frac{y^*}{\tau}\left(\frac{1}{\tau}V^*_k - e^{-j\sigma}V^*_m\right) + V^*_k\frac{y^{sh*}}{2\tau^2}\right) \\
&= \frac{|y|^2}{\tau^2}\left|\frac{1}{\tau}V_k - e^{j\sigma}V_m\right|^2 + 2\mathrm{Re}\left[\frac{y}{\tau}\left(\frac{1}{\tau}V_k - e^{j\sigma}V_m\right)V^*_k\frac{y^{sh*}}{2\tau^2}\right] \\
&\quad + \frac{1}{4\tau^4}|y^{sh}|^2|V_k|^2. \qquad (27)
\end{aligned}
$$

Notice that

$$
\begin{aligned}
\mathrm{Re}\left[\frac{y}{\tau}\left(\frac{1}{\tau}V_k - e^{j\sigma}V_m\right)V^*_k\frac{y^{sh*}}{\tau^2}\right] &= \frac{1}{\tau^3}\mathrm{Re}\left[yy^{sh*}\left(\frac{1}{\tau}V_k - e^{j\sigma}V_m\right)V^*_k\right] \\
&= \frac{1}{\tau^3}\mathrm{Re}\left[yy^{sh*}\left(\frac{1}{\tau}|V_k|^2 - e^{j\sigma}V_mV^*_k\right)\right] \\
&= \frac{1}{\tau^3}[(gg^{sh} + bb^{sh})\left((1/\tau)|V_k|^2\right. \\
&\quad -|V_k||V_m|\cos(\theta_{km} - \sigma)) \\
&\quad \left. -\ (bg^{sh} - gb^{sh})|V_k||V_m|\sin(\theta_{km} - \sigma)\right]
\end{aligned}
$$

since $yy^{sh*} = (gg^{sh} + bb^{sh}) + j(bg^{sh} - gb^{sh})$. Moreover, the first term of the RHS of (27) can be written as

$$
\frac{|y|^2}{\tau^2}\left|\frac{1}{\tau}V_k - e^{j\sigma}V_m\right|^2 = \frac{|y|^2}{\tau^2}\left(\frac{1}{\tau^2}|V_k|^2 + |V_m|^2 - \frac{2}{\tau}|V_k||V_m|\cos(\theta_{km} - \sigma)\right)
$$

Putting everything together yields

$$
\begin{aligned}
|I_{km}|^2 = & \frac{(g^2+b^2)}{\tau^2}\left(\frac{1}{\tau^2}|V_k|^2 + |V_m|^2 - \frac{2}{\tau}|V_k||V_m|\cos(\theta_{km}-\sigma)\right) \\
& + \frac{1}{\tau^3}\left[(gg^{sh}+bb^{sh})\left(\frac{1}{\tau}|V_k|^2 - |V_k||V_m|\cos(\theta_{km}-\sigma)\right)\right. \\
& \left. - (bg^{sh}-gb^{sh})|V_k||V_m|\sin(\theta_{km}-\sigma)\right] \\
& + \frac{1}{4\tau^4}(g^{sh2}+b^{sh2})|V_k|^2.
\end{aligned}
$$

Since $|V_k||V_m|\cos(\theta_{km}-\sigma) = c_{km}\cos(\sigma)+s_{km}\sin(\sigma)$ and $|V_k||V_m|\sin(\theta_{km}-\sigma) = s_{km}\cos(\sigma)-c_{km}\sin(\sigma)$ we can represent $i^{(2)}_{km} := |I_{km}|^2$, linearly, in terms of the fundamental variables $v^{(2)}_k$,$v^{(2)}_m$, c_{km}, and s_{km} as

$$
\begin{aligned}
i^2_{km} = & \frac{(g^2+b^2)}{\tau^2}\left(\frac{1}{\tau^2}v^{(2)}_k + v^{(2)}_m - \frac{2}{\tau}(c_{km}\cos(\sigma)+s_{km}\sin(\sigma))\right) \\
& + \frac{1}{\tau^3}\left[(gg^{sh}+bb^{sh})\left(\frac{1}{\tau}v^{(2)}_k - (c_{km}\cos(\sigma)+s_{km}\sin(\sigma))\right)\right. \\
& \left. - (bg^{sh}-gb^{sh})(s_{km}\cos(\sigma)-c_{km}\sin(\sigma))\right] \\
& + \frac{1}{4\tau^4}(g^{sh2}+b^{sh2})v^{(2)}_k \qquad (28) \\
= & \alpha_{km}v^{(2)}_k + \beta_{km}v^{(2)}_m + \gamma_{km}c_{km} + \zeta_{km}s_{km}, \qquad (29)
\end{aligned}
$$

where

$$
\begin{aligned}
\alpha_{km} :&= \frac{1}{\tau^4}\left((g^2+b^2) + (gg^{sh}+bb^{sh}) + \frac{(g^{sh2}+b^{sh2})}{4}\right) \\
\beta_{km} :&= \frac{(g^2+b^2)}{\tau^2} \\
\gamma_{km} :&= \frac{1}{\tau^3}\left(\cos(\sigma)(-2(g^2+b^2) - (gg^{sh}+bb^{sh})) + \sin(\sigma)(bg^{sh}-gb^{sh})\right) \\
\zeta_{km} :&= \frac{1}{\tau^3}\left(\sin(\sigma)(-2(g^2+b^2) - (gg^{sh}+bb^{sh})) - \cos(\sigma)(bg^{sh}-gb^{sh})\right).
\end{aligned}
$$

C Proof of Proposition 2

By (14b) we have that the Jabr inequality $c^2_{km}+s^2_{km} \le v^{(2)}_k v^{(2)}_m$ can be written as $||(2c_{km}, 2s_{km}, v^{(2)}_k - v^{(2)}_m)||_2 \le v^{(2)}_k + v^{(2)}_m$. Hence, taking $\lambda = (1,0,0)^\top$, and using (15) we have

$$
v^{(2)}_k + v^{(2)}_m - 2c_{km} \ge 0. \qquad (30)
$$

On the other hand, by summing up equations (6b) and (6a) we have

$$
\begin{aligned}
\ell_{km} &= G_{kk}v^{(2)}_k + G_{mm}v^{(2)}_m - 2G_{km}c_{km} \\
&\ge \min\{G_{kk}, G_{mm}\}(v^{(2)}_k + v^{(2)}_m - 2c_{km})
\end{aligned}
$$

given our assumptions on $G_{kk}, G_{mm}, G_{km}, G_{mk}$. Therefore, (30) implies $\ell_{km} \ge 0$.

D Proof of Proposition 3

First, we note that if $y^{sh}_{km} = 0$, a simple calculation yields

$$\ell_{km} = g_{km} \left| \frac{1}{\tau} V_k - e^{j\sigma} V_m \right|^2,$$
$$I_{km} = \frac{y}{\tau} \left(\frac{1}{\tau} V_k - e^{j\sigma} V_m \right)$$

see B. Therefore $i^{(2)}_{km} \geq 0$ clearly implies $\ell_{km} \geq 0$.

E Proof of Proposition 5

First, we show that for every $(x, s) \in C$ the inequality $(x')^\top x \leq ||x'||s$ holds. Indeed, by Cauchy-Schwartz inequality we have that $(x')^\top x \leq ||x'||||x||$, and since $||x|| \leq s$ the inequality follows. Next, we show that the point (x_0, s_0) defined by $s_0 := (||x'|| + s')/2$ and $x_0 := (s_0/||x'||)x'$ lies on $(x')^\top x = ||x'||s$. Certainly, $(x')^\top x_0 = (s_0/||x'||)||x'||^2 = s_0||x'||$. Moreover, (x_0, s_0) lies on C since $||x_0|| = ||(s_0/||x'||)x'|| = s_0$. Finally, we show that the vector $(x_0 - x', s_0 - s')$ is orthogonal to the plane (it suffices to take the vector (x_0, s_0) since $(0, 0)$ lies on C). It can be readily checked that $||(x_0, s_0)||^2 - (x')^\top x_0 - s's_0 = 0$.

F Proof of Proposition 6

Since $(x')^2 + (y')^2 > r^2$, there exists some $0 < t_0 < 1$ such that $(t_0x')^2 + (t_0y')^2 = r^2$. It can be readily checked that $t_0(x', y')$ is the projection of (x', y') onto S_r. Therefore, the normal to the separating hyperplane is $(1 - t_0)(x', y')$ and the RHS is $(1 - t_0)t_0||(x', y')||_2^2$, in other words, $(x')x + (y')y \leq r||(x', y')||_2$ is the desired valid inequality since $0 < t_0 < 1$.

G Proof of Proposition 4

Let $\{k, m\} \in \mathcal{E}$ be an arbitrary transmission line, and denote by $A := g^2 + b^2$, $B := gg^{sh} + bb^{sh}$ and $C := (g^{sh2} + b^{sh2})/4$, see (29). We drop the subscripts for clarity of exposition. Then, the coefficients in (19) can be written as

$$\frac{\beta}{\alpha} = \tau^2 \frac{A}{A + B + C}$$
$$\frac{\gamma}{\alpha} = \tau \frac{-\cos(\sigma)(2A + B) + \sin(\sigma)(bg^{sh} - gb^{sh})}{A + B + C}$$
$$\frac{\zeta}{\alpha} = -\tau \frac{\sin(\sigma)(2A + B) + \cos(\sigma)(bg^{sh} - gb^{sh})}{A + B + C}$$

Since we assumed $|b| > b^{sh}$, $g^{sh} = 0$, and $|b| > g$ we have that $A + B > 0$, hence the ratios $\frac{\gamma}{\alpha}$ and $\frac{\zeta}{\alpha}$ can be upper-bounded by

$$\tau \left(\frac{2A + |bb^{sh}|}{A + B + C} + \frac{|gb^{sh}|}{A + B + C} \right).$$

In consequence, we have the following upper-bounds

$$\frac{\beta}{\alpha} \leq \tau^2, \quad \frac{\gamma}{\alpha} \leq 3\tau, \quad \frac{\zeta}{\alpha} \leq 3\tau.$$

H Proof of Proposition 8

Consider a network consisting of two buses and a line joining them. There is one generator located at bus 1 and a load located at bus 2. The transmission line is simple (i.e., no shunts or transformers) and has series admittance $y = g + jb$ with $g > 0$, $b < 0$ and $-b > g$. Assume that the voltages at both buses are fixed to 1, that reactive power generation is unconstrained and that the line limit is very large. Finally, suppose that the objective consists in minimizing active generation. Additionally, there is no explicit reactive power load. We can make the following observations:

1. Active power balance constraints: $P_{1,2} = P_1^g$ and $P_{2,1} = -P_2^d$;
2. Power flows: $P_{1,2} = g - gc_{1,2} - bs_{1,2}$ and $P_{2,1} = g - gc_{1,2} + bs_{1,2}$;
3. Jabr inequality: $c_{1,2}^2 + s_{1,2}^2 \leq 1$.

Since reactive power generation is unconstrained and the line limit is very large, we do not have to write reactive power flows. Therefore, we can write the Jabr SOCP for this AC network as

$$\begin{aligned} OPT := \min \quad & g - gc_{1,2} - bs_{1,2} \\ \text{subject to:} \quad & \\ & g - gc_{1,2} + bs_{1,2} = -P_2^d && \text{(31a)} \\ & c_{1,2}^2 + s_{1,2}^2 \leq 1 && \text{(31b)} \end{aligned}$$

Since the network is radial, we know that imposing (31b) as an *equation* yields an exact ACOPF formulation [46,35]. Given that (31a) is linear in $c_{1,2}$, $s_{1,2}$ and constraint (31b) represents the unit circle, for a particular choice of parameters g, b, and P_2^d, the line represented by (31a) will intersect the boundary of the unit circle at exactly two points. These points will be the unique two AC feasible solutions for this instance, while their convex hull represents the feasible region of the SOCP. Finally unless the objective is orthogonal to the hyperplane defined by constraint (31a) the SOCP optimum will be exactly one of the two AC feasible points. We derive analytically these two AC points. Indeed, by substituting the expression for $s_{1,2}$ in (31a) into (31b) we obtain

$$\begin{aligned} \left(\frac{g^2+b^2}{b^2}\right) c_{1,2}^2 - 2\left(\frac{g}{b}\right)\gamma c + (\gamma^2 - 1) = 0, \quad \gamma := \frac{P^d + g}{b} \\ \implies c_{1,2} = \frac{(g/b)\gamma \pm \sqrt{\alpha}}{(g^2+b^2)/b^2}, \quad \alpha := 1 + \left(\frac{g}{b}\right)^2 - \gamma^2. \end{aligned}$$

Moreover, solving for P^d in terms of α we have

$$P^d = -g \pm \sqrt{g^2 - b^2(\alpha - 1)}.$$

Letting $g = 3$, $b = -8$, $\alpha = 3/4$, gives us $P^d = 2$ and $\gamma = -5/8$. This choice of parameters implies that the $s_{1,2}$-intercept in the $c_{1,2} - s_{1,2}$ plane of the line (31a) is strictly between $(-1, 1)$, hence there are exactly two AC feasible solutions and moreover they are irrational

$$\begin{aligned} \tilde{c}_{1,2} &= \frac{(15/64) + (\sqrt{3}/2)}{(9/64) + 1}, \ \tilde{s}_{1,2} = -(3/8)\left(\frac{(15/64) + (\sqrt{3}/2)}{(9/64) + 1}\right) + (5/8) \\ \hat{c}_{1,2} &= \frac{(15/64) - (\sqrt{3}/2)}{(9/64) + 1}, \ \hat{s}_{1,2} = -(3/8)\left(\frac{(15/64) - (\sqrt{3}/2)}{(9/64) + 1}\right) + (5/8) \end{aligned}$$

Given that the objective is not orthogonal to the linear constraint, and that it lies in the second quadrant, we conclude that the optimal solution to the Jabr SOCP is $(\tilde{c}_{1,2}, \tilde{s}_{1,2}) \approx (0.9647, 0.2632)$.

We conclude by showing that the dual of (31) has a unique maximizer which is irrational as well. Indeed, let $\lambda \in \mathbb{R}$ and $\delta \in \mathbb{R}_+$ be the dual variables associated to the constraints (31a) and (31b), respectively, and consider the following Lagrangean dual:

$$\begin{aligned} L(c_{1,2}, s_{1,2}; \lambda, \delta) : = g - gc_{1,2} - bs_{1,2} + \lambda(g - gc_{1,2} + bs_{1,2} + P_2^d) \\ + \delta(c_{1,2}^2 + s_{1,2}^2 - 1) \end{aligned}$$

The function $L(c_{1,2}, s_{1,2}; \lambda, \delta)$ is convex quadratic in the variables $c_{1,2}, s_{1,2}$. Given that Slater's condition holds for (31), we have that an optimal primal-dual pair must satisfy the KKT condition (Lagrangean optimality)

$$OPT = \min_{c_{1,2}, s_{1,2}} \quad L(c_{1,2}, s_{1,2}; \lambda^*, \delta^*) \tag{32}$$

$$= L(c_{1,2}^*, s_{1,2}^*; \lambda^*, \delta^*) \tag{33}$$

where $(c_{1,2}^*, s_{1,2}^*, \lambda^*, \delta^*)$ is an optimal primal-dual pair. It can be readily checked that the first-order stationarity condition implies that

$$2\delta^* c_{1,2}^* - g(1 + \lambda^*) = 0 \tag{34}$$

$$2\delta^* s_{1,2}^* + b(\lambda^* - 1) = 0. \tag{35}$$

Since we previously showed that the unique minimizer of (31) is the irrational solution $(\tilde{c}_{1,2}, \tilde{s}_{1,2})$, we have that $c_{1,2}^* = \tilde{c}_{1,2}$ and $s_{1,2}^* = \tilde{s}_{1,2}$. Therefore, equations (34) and (35) imply that (λ^*, δ^*) must also be irrational.

I Proof of Proposition 7

Let $\tilde{x} \in \mathbb{R}^n$ such that $A\tilde{x} \geq b - \epsilon e$, where e denotes the all-ones vector. Let $[D]$ denote the Lagrangean Dual problem of $[P]$ on variables $\lambda \in \mathbb{R}^m$. Consider a primal-dual optimal pair $(\overline{x}, \overline{\lambda}) \in \mathbb{R}^n \times \mathbb{R}^m$ such that $\overline{\lambda}$ is a rational vector that satisfies $||\overline{\lambda}||_1 \leq g(A, b, H, c)$, where g is a polynomial in the length of the input A, b, H, c. Such a solution $\overline{\lambda}$ exists by strong duality for convex QPs, we assumed $[P]$ is feasible and bounded, and by Section 2 in [70]). Strong duality also implies that $(\overline{x}, \overline{\lambda})$ satisfies KKT conditions (i) Lagrangean optimality $\nabla f(\overline{x}) = A^\top \overline{\lambda}$ and (ii) complementary slackness $\overline{\lambda}^\top (b - A\overline{x}) = 0$. In addition, given that f is convex and differentiable we have that $f(\overline{x}) + \nabla f(\overline{x})^\top (x - \overline{x}) \leq f(x)$ for every $x \in \mathbb{R}^n$. In consequence,

$$\begin{aligned} f(\tilde{x}) &\geq f(\overline{x}) + \nabla f(\overline{x})^\top (\tilde{x} - \overline{x}) \\ &= f(\overline{x}) + (A^\top \overline{\lambda})^\top (\tilde{x} - \overline{x}) \\ &= f(\overline{x}) + \overline{\lambda}^\top (A\tilde{x}) - \overline{\lambda}^\top (A\overline{x}) \\ &\geq f(\overline{x}) + \overline{\lambda}^\top (b - A\overline{x}) - \epsilon ||\overline{\lambda}||_1 \\ &= f(\overline{x}) - \epsilon ||\overline{\lambda}||_1 \\ &\geq f(\overline{x}) - \epsilon g(A, b, H, c). \end{aligned}$$

J Superoptimality and SOCPs

Consider the following SOCP:

$$
\begin{aligned}
[1] \qquad \overline{p} := \quad & \min \quad c^\top x \\
& \text{subject to:} \\
& \qquad ||y_i||_2 \le t_i \quad i \in [m] \\
& \quad y_i = A_i x + b_i, \quad t_i = c_i^\top x + d_i \quad i \in [m] \\
& \qquad Ax \ge b
\end{aligned}
$$

Its dual is given by:

$$
\begin{aligned}
[2] \qquad \overline{d} := \quad & \min \quad \lambda^\top b - \sum_{i=1}^{m} (u_i d_i + v_i^\top b_i) \\
& \text{subject to:} \\
& \qquad ||v_i||_2 \le u_i \quad i \in [m] \\
& \quad c^\top = \lambda^\top A + \sum_{i=1}^{m} (v_i^\top A_i + u_i c_i^\top) \\
& \qquad \lambda, u_i \ge 0
\end{aligned}
$$

Lemma 1 *Suppose that [1] is strictly feasible and $\overline{p} > -\infty$. Then strong duality holds, i.e., $\overline{p} = \overline{d}$, and there exists an optimal primal-dual pair $(\overline{x}, (\overline{y}_i, \overline{t}_i)_{i\in[m]})$, $(\overline{\lambda}, (\overline{v}_i, \overline{u}_i)_{i\in[m]})$. Let $(\tilde{x}, (\tilde{y}_i, \tilde{t}_i)_{i\in[m]})$ be an ϵ-feasible solution to the primal problem [1], i.e., it satisfies*

$$
\epsilon := \max \left\{ \max\{b_i - A_i^\top \tilde{x} \,:\, i \in [m]\}, \max\{||\tilde{y}_i||_2 - \tilde{t}_i \,:\, i \in [m]\} \right\}.
$$

Then,

$$
c^\top \tilde{x} \ge \overline{p} - \epsilon(||\overline{\lambda}||_1 + ||\overline{u}||_1). \tag{36}
$$

Proof Let e_m the all-ones vector in $\mathbb{R}^m$. Then,

$$
\begin{aligned}
c^\top \tilde{x} &= \overline{\lambda}^\top A\tilde{x} + \sum_{i=1}^{m} (\overline{v}_i^\top A_i \tilde{x} + \overline{u}_i c_i^\top \tilde{x}) \\
&\ge \overline{\lambda}^\top (b - \epsilon e_m) + \sum_{i=1}^{m} (\overline{v}_i^\top (\tilde{y}_i - b_i) + \overline{u}_i (\tilde{t}_i - d_i)) \\
&= \overline{\lambda}^\top b - \epsilon||\overline{\lambda}||_1 + \sum_{i=1}^{m} (\overline{v}_i^\top \tilde{y}_i + \overline{u}_i \tilde{t}_i) - \sum_{i=1}^{m} (\overline{v}_i^\top b_i + \overline{u}_i d_i) \\
&\ge \overline{p} - \epsilon||\overline{\lambda}||_1 + \sum_{i=1}^{m} (\overline{v}_i^\top \tilde{y}_i + \overline{u}_i (||\tilde{y}_i||_2 - \epsilon)) \\
&= \overline{p} - \epsilon||\overline{\lambda}||_1 - \epsilon||\overline{u}||_1 + \sum_{i=1}^{m} (\overline{v}_i^\top \tilde{y}_i + \overline{u}_i ||\tilde{y}_i||_2) \\
&\ge \overline{p} - \epsilon(||\overline{\lambda}||_1 + ||\overline{u}||_1),
\end{aligned}
$$

where the last inequality follows since if $||\overline{v}_i||_2 \le \overline{u}_i$ then $-\overline{v}_i^\top \tilde{y}_i \le ||\overline{v}_i||_2 ||\tilde{y}_i||_2 \le \overline{u}_i ||\tilde{y}_i||_2$ for all $i \in [m]$.

K Warm-Started Relaxation $T = 12, 24$

Table 16: Warm-Started Relaxation versus i2 SOCP+, $T = 12$ hr active power load curve, ramping rates 50%

	Cutting-Plane												i2 SOCP+								ACOPF
				First Iteration				Last Iteration				Total			Obj			Time			
Case	#Vars	#Cons	FTime	Obj	Time	#Cuts	DInfs	Obj	Time	#Cuts	DInfs	Time	Rnds	#Vars	#Cons	Gurobi	Knitro	Mosek	Gurobi	Knitro	Mosek	PBound
1354pegase	201209	219772	4.25	855287.79	9.75	18000	1.70e-08	860664.27	6.16	30236	1.06e-08	92.63	12	168974	257483	-	-	-	245.39	177.08	35.72	861473.50
9241pegase	1574972	1626484	34.46	3554525.89	163.27	205416	5.04e-07	3565252.49	271.12	354209	4.32e-04	3811.42	8	1354634	1969307	-	-	-	1439.33	TLim	352.61	-
ACTIVSg10k	1375020	1490960	30.21	28360735.35	106.03	84552	7.50e-09	28690135.28	94.69	148893	3.38e-07	1308.80	11	1151066	1714202	-	-	-	1543.11	TLim	190.39	-
10000goc-api	1401204	1503620	32.03	27353308.25	96.60	107568	3.27e-07	28459989.11	74.63	183308	2.40e-08	1066.39	11	1170122	1736383	-	-	-	TLim	TLim	256.71	-
10000goc-sad	1401204	1503620	30.42	16363341.55	133.30	105672	2.59e-06	16474320.51	105.74	144064	3.36e-06	1450.12	10	1170122	1736383	-	-	-	1918.80	TLim	247.71	-
ACTIVSg25k	3447015	3675848	74.34	68451480.40	357.88	206028	2.68e-07	69534077.06	972.97	796784	9.75e-07	4073.97	5	2885558	4247241	-	-	-	3539.20	TLim	523.18	-
30000goc-api	3768147	4127000	83.88	15513235.10	367.19	179196	7.00e-07	17079448.81	291.10	357763	1.06e-05	3667.40	9	3056138	4657273	-	-	-	TLim	TLim	756.77	-
30000goc-sad	3768147	4127000	86.66	13129167.74	521.03	539112	1.68e-06	13190466.44	329.66	246646	2.45e-06	3901.76	8	3056138	4657273	-	-	-	1694.63	TLim	888.11	-
ACTIVSg70k	9368583	10034036	207.22	184410148.60	1199.80	774684	4.07e-08	188736994.5	1276.64	1879714	3.66e-08	4702.30	3	7752122	11515211	-	-	-	3596.20	TLim	1124.19	-
78484epigrids-api	12551704	12730080	263.06	INF	137.79	1877340	-	INF	137.79	1877340	-	447.39	0	10674890	15270737	-	-	-	2875.97	TLim	1254.95	-
78484epigrids-sad	12551704	12730080	273.21	176961641.05	2697.25	2666748	1.36e-04	177335668.05	2828.46	3455898	1.36e-04	5465.67	1	10674890	15270737	-	-	-	TLim	TLim	1241.86	-

Table 17: Warm-Started Relaxation versus i2 SOCP+, $T = 24$ hr active power load curve, ramping rates 50%

	Cutting-Plane												i2 SOCP+								ACOPF
				First Iteration				Last Iteration				Total			Obj			Time			
Case	#Vars	#Cons	FTime	Obj	Time	#Cuts	DInfs	Obj	Time	#Cuts	DInfs	Time	Rnds	#Vars	#Cons	Gurobi	Knitro	Mosek	Gurobi	Knitro	Mosek	PBound
1354pegase	402677	440584	8.66	1656321.34	20.38	36000	2.46e-08	1666609.97	12.74	63049	2.97e-08	180.61	11	337946	517475	-	-	-	345.44	TLim	53.32	-
9241pegase	3151388	3258748	66.05	6865741.86	368.23	410832	4.33e-08	6925029.02	2427.02	1354277	1.25e-03	8419.66	4	2709266	3957551	-	-	-	1972.91	TLim	525.04	-
ACTIVSg10k	2752524	2991860	66.13	54565119.63	226.82	169104	7.5e-09	55184650.09	211.53	298812	7.18e-08	2869.25	11	2302130	3446078	-	-	-	3918.14	TLim	396.91	-
10000goc-api	2804496	3015596	62.15	51667906.99	208.71	215136	5.822e-07	53800966.46	175.46	388226	3.18e-07	2273.58	10	2340242	3490135	-	-	-	1332.84	TLim	542.88	-
10000goc-sad	2804496	3015596	60.63	32512183.58	284.34	211344	3.46e-06	32678842.78	276.70	246720	7.97e-06	3099.33	9	2340242	3490135	-	-	-	2994.53	TLim	714.14	-
ACTIVSg25k	6898863	7371032	150.16	132201219.46	816.18	412056	4.74e-07	134142620.00	2042.59	1385811	2.08e-06	6991.16	3	5771114	8536377	-	-	-	4451.39	TLim	963.91	-
30000goc-api	7539819	8268104	167.43	28586461.36	755.74	358392	1.7313e-06	31616476.13	884.66	1064587	2.06e-05	6799.92	7	6112274	9356989	-	-	-	TLim	TLim	1139.21	-
30000goc-sad	7539819	8268104	169.75	25717079.58	1091.27	1078224	2.52e-06	25811075.96	1021.24	1393213	1.23e-05	6249.53	5	6112274	9356989	-	-	-	3213.23	TLim	1125.14	-
ACTIVSg70k	18747555	20109632	410.49	350310736.78	2641.18	1549368	2.407e-07	357083446.24	2687.31	3359811	2.98e-08	8496.71	2	1550424	23139407	-	-	-	5061.29	TLim	1331.11	-
78484epigrids-api	25110280	24855336	533.07	-	1521.64	3754680	-	-	1521.64	3754680	-	2138.78	0	21349922	30681233	-	-	-	TLim	287.62	1434.04	-
78484epigrids-sad	25110280	24855336	532.93	344805634.00	4749.24	5333496	1.34e-04	-	TLim	6915253	-	11561.02	1	21349922	30681233	-	-	-	TLim	284.28	1448.87	-