Les Cahiers du GERAD ISSN: 0711–2440

Acceleration of umbrella constraint discovery in generation scheduling problems

A. Jahanbani Ardakani
F. Bouffard

G–2014–39

June 2014

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Acceleration of umbrella constraint discovery in generation scheduling problems

Ali Jahanbani Ardakani

François Bouffard

GERAD & Department of Electrical and Computer Engineering, McGill University, Montréal (Québec) Canada, H3A 0E9

ali.jahanbaniardakani@mail.mcgill.ca
francois.bouffard@mcgill.ca

June 2014

Les Cahiers du GERAD

G–2014–39

Abstract: Security-constrained optimal power flow (SCOPF) and security-constrained unit commitment (SCUC) problems are necessary tools for system operators for operational planning and near to real-time operation. The solution time of these problems are challenging mainly due to their inherent large size. Previous studies have shown that relatively few of those problems' constraints serve to enclose their feasible set of solutions. Therefore, the constraints that do not contribute to the feasible set of solutions could be discarded to decrease the size of these problems and their associated solution times. Umbrella constraint discovery (UCD) has been proposed to identify and rule out redundant constraints in dc-SCOPF problems. In this paper, we propose an improvement over the original UCD formulation that exploits the structure of its parent SCOPF problem. This new *partial UCD* approach can lead to significant speed-ups in terms of UCD solution time and size. Based on the encouraging results for partial UCD on SCOPF, we apply the technique on SCUC. Alike in SCOPF, partial UCD can efficiently strip out non-umbrella constraints off SCUC. We find, however, that because of its structure, SCUC has a much lower proportion of non-umbrella constraints.

Key Words: Computational complexity, convex optimization, generation dispatch, linear programming, preventive control, security-constrained optimal power flow, umbrella constraint, unit commitment.

Résumé: Les problèmes d'écoulement de puissance optimaux sous contraintes de sécurité (EPOCS) et les problèmes d'ordonnancement des groupes de production sous contraintes de sécurité (OGCS) sont essentiels à la bonne conduite et la minimisation des coûts de production dans les grands réseaux électriques. Les temps de calculs pour ces problèmes sont généralement importants étant donné la taille des grands réseaux électriques modernes. Des études antérieures ont démontré que relativement peu des contraintes de ces problèmes sont nécessaires afin de borner l'ensemble des solutions possibles. Il est donc envisageable de se départir d'un grand nombre de contraintes ne participant pas à la définition de l'ensemble des solutions possibles de ces problèmes. On peut ainsi réduire grandement la taille de ces problèmes et leurs temps de calcul. Le problème d'identification des contraintes parapluies (ICP) fut proposé pour éliminer les contraintes redondantes dans le contexte des problèmes d'EPOCS linéaires. Dans cet article, nous proposons une amélioration du problème ICP exploitant la structure des problèmes d'EPOCS. Cette nouvelle approche "d'ICP partielle" permet de réduire la taille et le temps de calcul du problème d'ICP de manière significative. Basé sur ces résultats encourageants, nous appliquons cette technique sur le problème d'OGCS. Tout comme dans le cas du problème d'EPOCS, l'ICP partielle permet de réduire le nombre de contraintes de l'OGCS. Étant donné sa structure, le problème d'OGCS ne peut bénéficier de l'apport de l'ICP partielle comme l'EPOCS.

Acknowledgments: This work was supported by the Fonds de recherche du Québec – Nature et technologies, Québec (Québec) Canada, and benefited from the IBM Academic Initiative through a free license for the ILOG CPLEX solver library.

1 Introduction

The uncertainties in the present large-scale interconnected electricity networks is on a rapid rise. With higher levels of renewable resources integration over the coming years and the continuous aging of network infrastructure, the uncertainties will even increase more. These uncertainties should be appropriately represented in the tools available for system operation and planning to study the impact of this wide spectrum of uncertainties. Security-constrained optimal power flow (SCOPF) and security-constrained unit commitment (SCUC) problems are the tools that system operators and planners use to deal with these uncertainties in scheduling the generation for the next day, the next hour and well into real time [1]. For instance, most of the large US ISOs solve some version of a SCUC on a daily basis to clear their day-ahead energy market [2]. In both SCOPF and SCUC problems the dimensions of the resulting problem, arising from the large scale of the networks themselves and of the large number of possible contingencies, often prevents the operators to take full advantage of these tools.

Previous work [3–5] and operator experience have shown that in such network-constrained generation scheduling problems only a very small proportion of constraints are indeed binding. That stems from the fact that such problems have a very high level of redundancy in their constraint set. Otherwise put, there are very few *umbrella constraints* which are both necessary and sufficient to describe their feasible set [6]. The remaining *non-umbrella constraints*, despite not having the potential to ever be binding, make use of computing resources during the solution process of these problems. Therefore, by identifying and removing these unnecessary constraints from such problems, they run faster.

This observation led to the proposal of the *umbrella constraint discovery* (UCD) problem in [6]. UCD is an optimization-based methodology which guarantees the identification of *all* constraints which form the feasible set of solutions (so-called umbrella constraints) of a given convex optimization problem. The method's efficacy has been proven mathematically and by applying the tool on SCOPF problems for different standard test networks. For example, in [6] a SCOPF implementing $N-1$ transmission security on the IEEE 118-bus system [7] was found to have 99.3% of its constraints to be non-umbrella. Removing those non-umbrellas allows the reduced-size SCOPF to run in one hundredth of its original (full model) time. Other authors have also proposed similar approaches, but generally focusing on potentially-binding constraints only. This is the case of Zhai *et al.* [5] which proposes a relaxation approach for the identification of inactive line constraints in a network-constrained unit commitment. However, in [5] the approach does not consider line constraints arising from operation under contingencies. In [8] and [9], the authors propose a classification of contingencies (dominated and dominating) to identify which constraints should be considered as part of a way to speed up a SCOPF decomposition solution approach.

As was evidenced and discussed in [6], the UCD problem itself is not easy as its own number of constraints increases with the number of constraints in the original problem. To overcome this issue, a divide and conquer decomposition technique was proposed to solve UCD. The success of the decomposition approach relied upon the inherent structure of the original SCOPF. In addition, empirical evidence generated in [6] showed that the membership in the umbrella set is relatively insensitive in light of nodal demand changes. In fact, the union of all yearly umbrella constraints can be used as a quasi umbrella set which contains all constraints necessary to bound network operation over an entire year. In the case of the IEEE Reliability Test System – 1996 [10], this union consists of 95 constraints only (instead of 2864 constraints for the full model in every hour).

In spite of the progress made by the use of decomposition, UCD is hard and requires significant solution time. In this paper, based on further observations of the SCOPF structure, we propose a new formulation of UCD which allows significantly faster non-umbrella constraint identification. This formulation, called *partial UCD*, can be used in the early stages of UCD decomposition to quickly maximize the pre-identification of non-umbrella constraints before solving a full UCD problem. With this added step, the resulting UCD problems are much smaller and run significantly faster.

In this paper, unlike in [6], we not only address the application of UCD on SCOPF problems, but we also consider the more general SCUC given the promising results for SCOPF. The goal here, in terms of the application of UCD to SCUC, is to generate simplifications via the uncovering of non-umbrella constraints in the unit commitment-based constraints of the SCUC which would lead to strong branching and time decoupling. This should be on top of constraint reductions associated with identifying network security-based non-umbrella constraints (as in the case of SCOPF).

2 Generation scheduling problem statement

The generation scheduling problem formulations considered in this paper for SCOPF and SCUC are presented in detail in Appendices A and B respectively. For the general case and where both continuous and binary decision variables are present in the generation scheduling problem, a compact mixed-integer problem (MIP) formulation takes the form

$$\min f(p,u) \tag{1}$$

subject to

$$[A|B]\begin{bmatrix} p \\ - \\ u \end{bmatrix} \leq \begin{bmatrix} b \\ - \\ c \end{bmatrix} \tag{2}$$

$$u \in \{0,1\}^M \tag{3}$$

where p is a vector in $\mathbb{R}^N$ lumping all of the network's continuous decision variables, u is a vector representing all (M) binary decision variables of the network, $f(p,u)$ is some objective function, A and B are constant matrices and b and c are constant vectors with appropriate dimensions.

3 Umbrella constraint discovery

For the benefit of the reader, the fundamentals and concepts behind the umbrella constraint discovery problem proposed in [6] are repeated here.

3.1 Definitions

Definition 1 (Umbrella constraint) *Let $\mathcal{J}$ be the set of indices corresponding to all the constraints (rows) of an optimization problem formulated as (1)–(3), and let $j \in \mathcal{J}$. Constraint (row) j is an umbrella constraint of $\mathcal{J}$ if and only if removing it from $\mathcal{J}$ alters set of feasible solutions of the original optimization problem.*

Definition 2 (Umbrella set) *The umbrella set $\mathcal{U} \subseteq \mathcal{J}$ of an optimization problem is the set containing the* minimum *number of constraints (rows) necessary to form its set of feasible solutions. Removing any member of the umbrella set $\mathcal{U}$ would alter the set of feasible solutions of the original optimization problem, while adding any of the constraints in the non-umbrella set, $\overline{\mathcal{U}} \subseteq \mathcal{J}$, for which $\mathcal{U} \cap \overline{\mathcal{U}} = \emptyset$ and $\mathcal{U} \cup \overline{\mathcal{U}} = \mathcal{J}$, would not change the set of feasible solutions.*

3.2 Problem formulation

In the following sections, unless stated otherwise, we consider the linear programming (LP) relaxation of the mixed-integer problem in (1)–(3), *i.e.* (3) is replaced by $0 \leq u_i \leq 1$ for all $i \in M$. This is possible because the structure of the constraint set in (2) is unaffected by the nature of the decision variables. It is not hard to prove that if a constraint in (2) is umbrella (non-umbrella) for the LP relaxation, it is umbrella (non-umbrella) for the MIP.

For an optimization problem defined by the LP relaxation of (1)–(3), the goal of UCD is to identify all umbrella and non-umbrella constraints in the set $\mathcal{J}$. The basic observation that led to the proposal of UCD is that *there is always at least* one point *lying on an umbrella constraint that is feasible with respect to all the other constraints in a convex optimization problem.* In the case of a non-umbrella constraint, *there is* no single point *available on the constraint that can be feasible with respect to all other constraints of the problem* [6].

Now consider a_j, which denotes the vector corresponding to the j^{th} row of matrix $[A|B]$ and b_j the j^{th} element of the $[b|c]^T$ vector of the system of inequalities in (2). Let us also define an arbitrary vector $w_j \in \mathbb{R}^I$ where I corresponds to the total number of decision variables (p and u). We claim that constraint j is umbrella if and only if there exists a point w_j, which lies on the hyperplane $a_j^T w_j = b_j$, while the chosen point w_j is consistent with respect to all other constraints $j' = 1, \ldots, J$, *i.e.* w_j is such that it satisfies $a_{j'}^T w_j \le b_{j'}$ for all j'.

In the event where there is no feasible w_j (indicating that constraint j is non-umbrella), we add a slack variable $s_j \ge 0$ such that $a_j^T w_j + s_j \ge b_j$ can be satisfied while all other constraints on w_j are met. This addition of a non-zero slack ensures the consistency of non-umbrella constraint j with respect to all other constraints j'. Here s_j can be interpreted as the "distance traveled" by constraint j to enter the umbrella set. In the event where constraint j is umbrella, setting $s_j = 0$ is possible. This allows us to define the identification rule

$$s_j = \begin{cases} 0, & \text{if constraint } j \text{ is umbrella} \\ > 0, & \text{otherwise.} \end{cases} \tag{4}$$

By finding the values of s_j for all j, we can identify the membership of the umbrella set of (2) (as well as possibly the relaxed constraints corresponding to (3)). This is done by solving the following linear program (LP)

$$\min_{w, s \ge 0} \sum_{j=1}^{J} s_j \tag{5}$$

subject to, for all $j = 1, \ldots, J$,

$$a_{j'}^T w_j \le b_{j'} \qquad j' = 1, \ldots, J \tag{6}$$

$$a_j^T w_j + s_j \ge b_j \tag{7}$$

As UCD is an LP, it can be solved in polynomial time in theory [11]. However, UCD is much larger than the original SCOPF and 0-1 relaxed SCUC problems in (1)–(3), (which are also LPs). In the case of SCOPF, UCD has $J(I+1)$ variables and $J(J+2)$ constraints, as opposed to I variables and J constraints. Therefore, even for small power systems, the dimensions of UCD become unmanageable. For example, for the small-sized IEEE 118-bus standard test system [7]—with $I = 118$ and $J = 65\,730$ implementing $N-1$ transmission security—cannot be solved directly. In order to overcome this shortcoming, a decomposition approach has been proposed in [6] to break UCD problem in tractable sub-problems. As for the SCUC 0-1 LP relaxation, the application of UCD on its constraint set suffers also from the same dimensionality problem, albeit at a wider scale for similar-sized systems.

3.3 UCD decomposition approach

The UCD problem decomposition approach relies on the following principle, which is proved in [6]:

Lemma 1 (Non-Umbrella Constraint Lemma) *If constraint $j \in \tilde{\mathcal{J}}$, where $\tilde{\mathcal{J}} \subseteq \mathcal{J}$, is found to be non-umbrella when solving UCD for the set of constraints $\tilde{\mathcal{J}}$ then constraint j is also found to be non-umbrella when solving UCD for the full set of constraints $\mathcal{J}$.*

The Non-Umbrella Constraint Lemma provides a simple way to flag and remove constraints which will automatically be non-umbrella as new constraints are brought in for comparison. This motivates the decomposition algorithm from [6]:

Step 1: Partition the full set of constraints $\mathcal{J} = \{\mathcal{J}_1, \ldots, \mathcal{J}_M\}$ into computationally-manageable constraint blocks and solve UCD for each of those. This step can be executed by parallel processors.

Step 2: Recombine the umbrella constraints generated by the solution of UCD in Step 1 into computationally-manageable constraint blocks, and solve UCD for each of those (again this can be done over parallel processors).

Step 3: Repeat Step 2 until the solution of a single UCD problem is obtained or when a required constraint number reduction goal is attained.

We note here that the outcome of this decomposition process does not lead to the exact computation of the objective function in (5). It merely serves as a mechanism to progressively flag out non-umbrella constraints, *i.e.* once a non-zero s_j value is obtained for constraint j in either Steps 1 or 2, it is flagged as non-umbrella, and it is not considered in subsequent iterations.

3.4 Computational inefficiencies in UCD

The vast majority of UCD constraints is coming from the set of constraints in (6). Throughout the solution process, the LP solver has to look for w_j values for each constraint in (7) while ensuring its feasibility against all other constraints in (6). In the event where such a w_j cannot be found, s_j must be increased until feasibility is re-established. The observation one makes here is that the moment $s_j > 0$, there is no longer a need to keep on searching for a feasible w_j as constraint j is proved to be non-umbrella (as per the Non-Umbrella Constraint Lemma). Therefore, it is pointless to keep on comparing other constraints to a proven non-umbrella constraint. It is also equally true that it is wasteful to compare non-umbrella constraints against each other. For instance, in [6] when considering the IEEE 118-bus test system, we found that only 451 constraints were indeed umbrella. Therefore, throughout the UCD solution process 65 279 SCOPF non-umbrella constraints were being compared against each other without providing any more information about the actual set of umbrella constraints. Thus, as a rule of thumb, we should seek to maximize the comparison of potentially non-umbrella constraints only with those constraints which would be likely umbrellas. Such likely umbrella constraints could be the generation limits and the power balance which, in the case of the IEEE RTS, are members of the umbrella set throughout the year [6].

4 Partial umbrella constraint discovery

As mentioned above, in UCD a constraint j is flagged as umbrella if the LP solver is able to find a point w_j on that constraint that satisfies all other constraints. This "proof" requires that the constraint be compared against *all* other constraints. On the other hand, to flag a constraint j as non-umbrella, a single cross-constraint comparison could be enough if that comparison is able to show that there is no point on constraint j that can satisfy that other constraint. We recall also that if a constraint is flagged as non-umbrella at any point during the solution process, that constraint *is non-umbrella for the entire set of constraints*.

Obviously, it is not practical to interrupt the LP solution process every time an instance of $s_j > 0$ is being found. Instead, we propose here to solve an approximation of UCD as a pre-processing step to its full solution. This pre-processing step, which we call *partial UCD* (P-UCD), serves to rapidly identify non-umbrella constraints with the objective of reducing the number of necessary constraint comparisons in the solution of the full UCD.

In P-UCD, the SCOPF constraint set is partitioned into two subsets: potential umbrella ($\mathcal{J}_\nu$) and potential non-umbrella ($\mathcal{J}_\eta$) constraints such that $\mathcal{J} = \mathcal{J}_\nu \cup \mathcal{J}_\eta$ and $\mathcal{J}_\nu \cap \mathcal{J}_\eta = \emptyset$. The membership in the potential umbrella and non-umbrella constraint sets are determined heuristically. The heuristic here should be based essentially on operator experience and/or on past UCD results. For instance, if the results of a prior UCD run are available, it is reasonable to assume that the current UCD results will be similar given similar operating conditions, as seen in [6]. Thus, in that case, $\mathcal{J}_\nu$ would contain the umbrellas from the past UCD run and $\mathcal{J}_\eta$ would contain its non-umbrellas. Another valuable partition would be to have $\mathcal{J}_\nu$ include the economic dispatch constraints (*i.e.* the generation limits and the network-wide power balance), while $\mathcal{J}_\eta$ regroups all line flow limits. This is also in line with the evidence found in [6], which shows that the vast majority of SCOPF umbrella constraints come from the generation side.

4.1 Formulation of partial UCD

Let us pick a constraint $j \in \mathcal{J}_\nu$. Also, let a_j denote the vector corresponding to the j^{th} row of matrix $[A|B]$ and b_j the j^{th} element of the $[b|c]^T$ vector of the system of inequalities in (2). Similarly, for $j' \in \mathcal{J}_\eta$, let $a_{j'}$

denote the vector corresponding to the j'^{th} row of matrix $[A|B]$ and $b_{j'}$ the j'^{th} element of the $[b|c]^T$ vector of (2). Let us also define an arbitrary vector $w_j \in \mathbb{R}^I$, where $I = M + N$ (the number of variables in the SCOPF). We claim that by solving the following optimization problem, the non-umbrella constraints in the set of $\mathcal{J}_\eta$ with respect to the constraints in $\mathcal{J}_\nu$ are identified.

Proposition 1 (Partial UCD (P-UCD))

$$\min_{w,s\geq 0} \sum_{j\in\mathcal{J}_\eta} s_j \tag{8}$$

subject to,

$$a_{j'}^T w_j \leq b_{j'} \qquad j \in \mathcal{J}_\eta,\ j' \in \mathcal{J}_\nu \tag{9}$$

$$a_j^T w_j + s_j \geq b_j \qquad j \in \mathcal{J}_\eta \tag{10}$$

The interpretation of partial UCD is very similar to that of UCD. If constraint $j \in \mathcal{J}_\eta$ is indeed umbrella, there exists a w_j that can lie on the hyperplane $a_j^T w_j = b_j$ (as dictated by (10)) while still remaining feasible with respect to all constraints in subset $\mathcal{J}_\nu$, (9). In the case where constraint $j \in \mathcal{J}_\eta$ is not umbrella, to maintain the feasibility of P-UCD, s_j should be increased until one finds a w_j which can lie on the hyperplane $a_j^T w_j = b_j - s_j$ while simultaneously satisfying all constraints in subset $\mathcal{J}_\nu$. This is just like UCD where the slack variables s_j serve to bring all non-umbrella constraints to the border of the feasibility region to ensure the feasibility of the whole problem. In the case of the partial UCD problem, nonetheless, the feasibility region is defined by the constraints in subset $\mathcal{J}_\nu \subseteq \mathcal{J}$ which are predetermined based on user input. This is unlike UCD in (6) and (7) where its feasibility region is defined by all constraints of the SCOPF.

Figure 1 shows an example of P-UCD. Here, $\mathcal{J}_\nu$ comprises of all constraints except 1–4 and 9 and 10. Clearly, constraints 1–4 are umbrella. The only non-umbrella constraints identified through P-UCD are constraints 9 and 10. Partial UCD would find all non-umbrellas if $\mathcal{J}_\nu$ contained only constraints 1–4. Otherwise, some true non-umbrellas would not be identified as such. In all cases, it is necessary to follow up with a UCD run to formally identify all umbrella constraints.

P-UCD allows for s_j and w_j variables to be only allocated to potentially non-umbrella constraints and, therefore, it can have a significantly lower dimensionality in comparison to UCD. It also has a much lower constraint count. As can be seen in the P-UCD formulation, to rule out the non-umbrellas in the subset

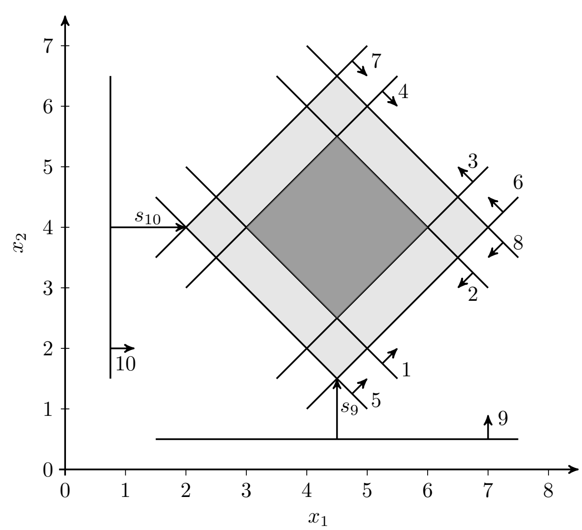

Figure 1: Illustration of partial UCD: we assume here that constraints 5–8 are potential umbrellas (members of $\mathcal{J}_\nu$). The remaining constraints (1–4 and 9 and 10) are potential non-umbrella constraints (members of $\mathcal{J}_\eta$). Solving partial UCD will identify constraints 9 and 10 as non-umbrella. The subsequent UCD run on constraints 1–8 would identify constraints 1–4 as umbrella.

$\mathcal{J}_\eta$, its elements are only compared against constraints in the subset $\mathcal{J}_\nu$. The number of constraints in a partial UCD problem is $|\mathcal{J}_\eta| \cdot |\mathcal{J}_\nu|$ which can be significantly less than the number of the corresponding full UCD problem, that is $J(J+2)$. The number of variables in P-UCD is $|\mathcal{J}_\eta|(I+1)$ as opposed to the umber of variables in the corresponding UCD problem which is $J(I+1)$. For instance, in the case of the IEEE RTS [10], if we let $\mathcal{J}_\nu$ contain all 50 economic dispatch constraints and $\mathcal{J}_\eta$ hold the remaining 2814 line flow limits, P-UCD has 140 700 constraints and 67 536 variables. This is in stark contrast with the full UCD which would have over eight million constraints and 71 600 variables.

Obviously, the price to pay with the solution of P-UCD is that it *only identifies the constraints in* $\mathcal{J}_\eta$ *which are non-umbrella*. It cannot identify any umbrella constraints unless $\mathcal{J}_\nu$ is the umbrella set. As $\mathcal{J}_\nu$ is only a good guess of what the umbrella set is, $\mathcal{J}_\nu$ may contain some non-umbrellas and $\mathcal{J}_\eta$ may have umbrellas within itself. What P-UCD provides is a way to fast-track the identification of non-umbrellas so that a full UCD can be performed only on those constraints in $\mathcal{J}_\nu$ and those in $\mathcal{J}_\eta$ not ruled out as non-umbrellas by P-UCD. This way UCD can be run on a much reduced set of constraints which will be identified as umbrella or not—providing the proof of umbrella set membership.

It is important to note that P-UCD cannot leave out a *true* umbrella constraint by identifying it as a non-umbrella constraint. If a constraint $j \in \mathcal{J}_\eta$ (*i.e.* in the set of potential non-umbrellas) is indeed umbrella, one can find a w_j satisfying all constraints in (9). Otherwise, it is non-umbrella by the Non-Umbrella Constraint Lemma. This property ensures that the feasible set of solutions of the original SCOPF problem will never be altered by P-UCD.

As mentioned before, the closer the set $\mathcal{J}_\nu$ is to $\mathcal{U}$, the more non-umbrella constraints in $\mathcal{J}_\eta$ get identified by P-UCD. Therefore, a good estimation of the potential umbrella constraint set is essential for efficient function of P-UCD. In the next Section, we provide empirical evidence of the efficacy of P-UCD as a pre-processing step to a full UCD run.

5 Partial UCD test results

In this Section, we apply partial UCD approach on two SCOPF problems which were previously investigated in [6]. We highlight the importance of potential umbrella ($\mathcal{J}_\nu$) and non-umbrella ($\mathcal{J}_\eta$) constraints set selection and how it affects the performance of partial UCD.

5.1 Application on IEEE-RTS

We use the single-area version of the IEEE Reliability Test System–1996 [10] and form the SCOPF problem considering only line contingencies in the first hour of the RTS year. The corresponding SCOPF problem has a constraint set of size 2864 constraints (set $\mathcal{J}$), where 2814 constraints are arising from line constraints, 48 from generation limits and two constraints from load balance constraints. The solution time of the corresponding SCOPF problem subject to all constraints is 0.015 seconds. By running UCD on the SCOPF constraint set, 87 umbrella constraints are identified (3% of constraints are umbrella). The run time for UCD (ran without decomposition) is about 16 minutes, while using the line-based decomposition technique proposed in [6], the UCD run time can be decreased to about 4.5 minutes.

To implement partial UCD, an estimate of actual umbrella constraints is necessary. The experience of running UCD on this network (see [6]) suggests that most of the generation limits are part of the final umbrella set along with the load balance for any network. Therefore, a good estimation of the umbrella set $\mathcal{J}_\nu$, as suggested in the previous section, would be the set of generation limits and the load balances constraints. The remaining constraints, which are all arising from line operations under no-contingency and contingency states, form the set $\mathcal{J}_\eta$.

These potentially umbrella and non-umbrella constraint sets are applied to (9) and (10), and the corresponding solution time for partial UCD is 0.484 s, while 2776 out of 2814 constraints in $\mathcal{J}_\eta$ are flagged as non-umbrella. Only 38 constraints in set $\mathcal{J}_\eta$ are flagged as potential umbrellas.

In the next stage, the 38 potential umbrellas from $\mathcal{J}_\eta$ are merged with the 50 constraints from $\mathcal{J}_\nu$. UCD is run with those 88 remaining constraints, and it takes 0.016 s to run. The result of UCD is that only one generation limit from the set $\mathcal{J}_\nu$ is finally identified as non-umbrella. Therefore, the final umbrella set has 87 constraints of which 38 are coming from line constraints and 49 are coming from generation limits and load balance constraints (as before). The total time for running P-UCD followed by UCD takes 0.500 s. Following this, if the minimum cost SCOPF is solved subjected to the umbrella set only the solution time is virtually 0 s.

As it can be seen from the results, the solution time of UCD can decrease significantly if the constraints are preprocessed using P-UCD; here we see a reduction from 4.5 minutes (270 s) down to 0.5 s (0.19% of the UCD original time with line-based decomposition). As was discussed before, the efficacy of P-UCD is strongly dependent upon the accuracy of the potential umbrella set $\mathcal{J}_\nu$. The estimation of the potential umbrella set is not always a simple task. As will be shown in the next subsection, it is not always enough to include generation limits and load balance constraints in $\mathcal{J}_\nu$. A poor estimation of the potential umbrella set may result in the identification of only a few non-umbrella constraints in $\mathcal{J}_\eta$ thus making the P-UCD step counterproductive.

5.2 Application on IEEE 118-bus system

The corresponding SCOPF problem of the IEEE 118-bus system[1] has 65 730 constraints. Generation limits count for 236 constraints, and the load balance constraints contribute two more inequality constraints. Only 176 line contingencies are considered, and the constraints arising from the operation of the lines under these contingencies account for the remaining 65 492 constraints. The SCOPF problem takes 2.808 seconds to be solved. The corresponding UCD problem of this SCOPF is too large to be solved directly as the one formulated with for the IEEE RTS. Line-based decomposition, introduced in [6] and explained briefly in Section 3.3, is used to partition UCD problems in smaller sub-problems. These sub-problems are independent from each other and can be run using parallel computing. The total run time of the UCD using line-based decomposition takes over 3700 seconds to complete.

Partial UCD is used here to preprocess the set of constraints and rule out some non-umbrella constraints in order to reduce the size and solution time of the full UCD. To set up P-UCD, an estimation of potential umbrella constraints is necessary. We shall assume the same $\mathcal{J}_\nu - \mathcal{J}_\eta$ partition used with the IEEE RTS above. Therefore, $\mathcal{J}_\nu$ has 238 members and the rest of the constraints (line constraints) belong to $\mathcal{J}_\eta$ which has 65 492 constraints. P-UCD has about 15.5 million constraints, which is too large to be solved directly. The solution is to break down the set $\mathcal{J}_\eta$ into smaller subsets (100 here) and run partial UCD on each of them. These subsets can be formed using different approaches such as line-based or contingency-based decomposition. These sub-problems can be solved independently using parallel computing.

Each of the sub-problems are solved using partial UCD with $\mathcal{J}_\nu$ as the potential umbrella constraint set. P-UCD takes 52.673 s to run in total, with 0.0804 s on average for each subset, and it identifies 39 806 constraints from the set $\mathcal{J}_\eta$ as non-umbrellas. This leaves us with 25 686 constraints as potential umbrellas in $\mathcal{J}_\eta$. The large number of remaining constraints in $\mathcal{J}_\eta$ could be because some of the generation limits in $\mathcal{J}_\nu$ are not as stringent as in the RTS case, while many of the line constraints in $\mathcal{J}_\eta$ are more demanding and cannot be flagged as non-umbrella by the generation and load balance constraints alone.

The remaining number of constraints in $\mathcal{J}_\eta$ is still too large to be considered for a single run of a UCD to identify all umbrella constraints. This is mostly because the estimation of the umbrella set was not accurate enough. To address this issue, there are two options. Firstly, the problem can be solved using decomposed UCD which will identify all umbrella constraints at the expense of more computation time. Otherwise, we can use a better estimation of the potential umbrella constraints in the first place. An example of such a more exact estimation of the umbrella constraints could be the umbrella set of the previous operating hour.

We consider now the case where the set $\mathcal{J}_\nu$ is set to $\mathcal{U}$ from the previous hour of operation which typically has similar loading patterns to that of the current hour. The set $\mathcal{J}_\nu$ has 451 constraints, while the set $\mathcal{J}_\eta$

[1] The complete data set of 118-bus system which was used in this study can be found at http://cl.ly/3x0L3L393r24.

Table 1: Summary of UCD solution times with and without the application of P-UCD

Test Network	UCD	Partial UCD
IEEE RTS	240 s	0.5 s
IEEE 118-bus	3700 s	701.5 s

includes 26 513 constraints. As before, the set of potential non-umbrella constraints, $\mathcal{J}_\eta$, is divided into smaller subsets, each with 100 constraints here, and partial UCD is run for each of the subsets. Each subset takes 2.48 s to run on average and the cumulative run time for all subsets is 633.6 s. In this stage, only 34 constraints remain in $\mathcal{J}_\eta$ which shows the efficiency of the alternative potential umbrella set $\mathcal{J}_\nu$. Including $\mathcal{J}_\nu$, there are now 485 constraints which have the potential to be in the final umbrella set. The number of remaining umbrella constraints is small enough to run a full UCD on them. UCD takes 15.23 s to run and identifies 37 constraints as non-umbrella. The final umbrella set contains 448 constraints of which 213 constraints come from line constraints (about 0.3% of all line constraints in the original set of constraints) and 235 constraints come from generation limits and load balance constraints. The total run time for this method, including partial UCD followed by UCD, is 701.5 s which is less than 20% of the best UCD run time using decomposition alone.

Table 1 summarizes the progress made through the use of P-UCD to streamline the solution of UCD. In the case of the IEEE RTS especially, the results are impressive. However, the observations here may be artifacts of the structure of this particular test network which is known to be *transmission strong* and *generation weak*. The fact that most of the generation limits are umbrella indicates that the generation system restricts the feasible space of the SCOPF the most, while transmission constraints do not contribute much to the restriction of operations. In the case of the IEEE 118-bus system the situation is not as clearcut, but the results show a significant improvement over the best case without P-UCD. Given those improvements in UCD performance when combined with P-UCD, one can consider its application to the security-constrained unit commitment, as we do next.

6 Security-constrained unit commitment

UCD has proven to be a very effective tool in identifying non-umbrella constraints in SCOPF problems. One of the most important tools for ISOs and RTOs in day-ahead market clearing is security-constrained unit commitment (SCUC), which is solved over and over. By solving SCUC problems, operators aim to minimize the operation cost of the system or the total bid cost while satisfying a subset of technical constraints of the power system. Technical constraints arise from operation of the system in pre- and post-contingency states and include system load demand, line limits and generation units constraints, such as ramping limits, minimum up/down times, *etc.* A generic SCUC problem formulation is presented in Appendix B. SCUC problems are more challenging than SCOPF problems mainly due to the presence of binary variables and because of time couplings. Thus, it could be beneficial to apply UCD as a pre-processing step to SCUC because of the relative stability of the power system feasibility region over time spans of 12 to 24 hours.

What we investigate next is the application of UCD and P-UCD to the 0-1 LP relaxation of a SCUC formulation. We will also discuss some non-umbrella constraint identification which could have significant computational benefits for SCUC. Here we make use of the IEEE RTS to illustrate the potential benefits and drawbacks of UCD and P-UCD on SCUC solution.

6.1 Application of UCD on static unit commitment

In this section, we implement UCD on the corresponding single-period (static) unit commitment problem of IEEE-RTS without considering the network for illustrative purposes only. In order to form a UCD problem for a static unit commitment problem (UC), we need to consider its 0-1 linear programming relaxation, wherein all binary variables are converted to continuous variables and a lower (0) and upper (1) bounds are added to the set of constraints for each variable).

The IEEE-RTS has 32 generating units; therefore, in the corresponding UC, the generation limits contribute 64 constraints—one for the lower and one the for upper generation limits. The load balance equality constraint counts contributes two equivalent inequality constraints. When forming the corresponding UCD problem, 64 constraints are added for upper and lower limits on the relaxed binary variables. Therefore, in total, the set of constraints has 130 elements.

The system load has been varied and the corresponding UC and UCD problems are solved. For different ranges of loads, constraints join or retire from the umbrella set of the problem as expected. All of these constraints are umbrella for loads below 3005 MW. The interesting observation is that if a lower limit constraint of a binary variable is non-umbrella, this indicates that the generating unit corresponding to that binary variable would have to be "on" in that range of load. For the IEEE-RTS this situation happens for loads over 3005 MW where, for example at 3006 MW, lower limit constraints of units 22 and 23 become non-umbrella, signalling that these two generating units must remain online for any load above that threshold. Also, if the load is increased even more, we observe that the minimum generation limits of those units will also leave the umbrella set. The same phenomenon repeats itself with other units if we keep increasing the load. This observation continues until the load reaches a point where the problem becomes infeasible. Thus, if the lower bound of a relaxed binary variable is flagged as non-umbrella, the corresponding generating unit has to be online. A similar argument can be made in the case of upper bounds on relaxed binary variables which are non-umbrella. In that case, the corresponding generator has to be offline.

Knowing ahead of time that some of the binary variables should be fixed prior to the UC solution process is quite advantageous as it would contribute to decrease the computational cost of the corresponding branch-and-cut algorithm. This type of information is only available when the system is either highly or lightly loaded, however. This type of information is not usually available at mid-range loads of the system since all binary variable limits are umbrella constraints.

6.2 Application of UCD on dynamic unit commitment

In this section, we apply UCD and P-UCD on multi-period (dynamic) UC problems with pre-contingency network constraints. The test network is still the IEEE-RTS and the UC problem is modeled as given in Appendix B.

For the purpose of this study, a planning horizon comprising the first four hours of the IEEE-RTS year is considered. This system has 32 generating units and 38 transmission lines. The corresponding SCUC has 1 532 constraints of which 304 come from line limits, while the rest are associated to generation-related inequalities (see Appendix B). The solution time for the SCUC problem here is 0.063 s.

Now, if we consider the corresponding UCD problem, it has an extra 128 constraints associated with the 0-1 relaxation of the binary variables. Therefore, UCD has 1 660 constraints in total here. By running UCD, which takes 69.701 s to solve, only 491 constraints are identified as non-umbrella (29.6%). From the constraints corresponding to the line limits, all of them are non-umbrella for the whole time horizon, except for upper flow limit on line 11 (between buses 7 and 8). This result was expected as the net generation at bus 7 is over the upper flow limit of line 11. From generation-related constraints, all system spinning reserve requirements are non-umbrella as well as most of ramp limits. The interesting result is than most of the ramp limits are not the restricting limits when it comes to the inter-period couplings. The minimum up/down time restrictions are those playing this role as they are never flagged as non-umbrella. The compact size original problem takes 0.047 s to solve (25% faster than with all constraints). Table 2 summarizes the results of application of UCD on static and dynamic unit commitment problems associated with the RTS.

A scenario where both ramp limits and minimum up/down time constraints flagged as non-umbrella would be very interesting from a SCUC computational point of view, as this indicates that two successive time periods have no coupling. If this is the case for all generating units simultaneously, it would be possible to split the SCUC between the periods where those non-umbrella constraints arise and solve them independently. This is obviously desirable property; however, it is very unlikely to happen in practice.

Table 2: Application of UCD on static and dynamic unit commitment

	Static UC	Dynamic UC
Horizon	1 hr	4 hr
# of constraints	66	15 322
UC run time	0 s	0.063 s
# of UCD constraints	130	1660
UCD run time	0.109 s	69.701 s
# of non-umbrella	0	491
Compact UC run time	0 s	0.047 s

For time horizons over four hours in the case of the IEEE RTS, UCD becomes intractable and cannot be solved directly. In this case, we can use P-UCD to pre-identify some of the non-umbrella constraints. The results of P-UCD when $\mathcal{J}_\eta$ contains all line constraints and $\mathcal{J}_\nu$ holds all generation limits match the results presented in [5].

6.3 Discussion

As it can be seen from the results of the application of UCD on dynamic SCUC problems, not many constraints are proven non-umbrellas. Most of the non-umbrella constraints, as expected, are line flow constraints and most of constraints arising from generation limits are umbrella. One observation here is that what makes line flow limits good non-umbrella constraint candidates is that they couple a large number of the problem's decision variables [see (14) which couples all p_i variables]. Geometrically speaking, these constraints affect multiple dimensions of the problem simultaneously. On the other hand, generation-based constraints contribute to restrict at most four variables simultaneously ($u_{i,t}$, $u_{i,t-1}$, $p_{i,t}$ and $p_{i,t-1}$). Therefore, whether a constraint pertaining to a specific generating unit is flagged as umbrella or not is essentially independent from the constraints on other generating units (unless the system load is very high or very low). In addition, with the introduction of a new time period or a new generating unit, new dimensions are added to the original feasible space. For a constraint which was non-umbrella before the addition of a new variable, this constraint will become umbrella of the new feasible space unless it is explicitly coupled with the added dimension.

The effective conclusion here is that one should use UCD and P-UCD to reduce the number of transmission constraints (as with SCOPF). However, it is clear that the benefit of applying UCD and P-UCD on unit commitment constraints is very slim given their computational cost.

7 Conclusion

UCD is a very effective tool in identifying non-umbrella constraints. The complexity and computational burden associated to UCD can be problematic. In this paper we proposed partial UCD as a fast pre-processing step to UCD. Partial UCD can rapidly identify non-umbrella constraints in the set of constraints and reduce the number of constraints that UCD has to consider. The efficacy of partial UCD is dependent upon an *a priori* estimation of the final umbrella set. Good set estimates include the results of previous runs of UCD on the same problem and/or user experience.

UCD and partial UCD can reduce the complexity of transmission security-constrained optimization problems in power systems (SCOPF and SCUC) by identifying and removing non-umbrella constraints prior to their solution. Given the structure of SCOPF and SCUC problems, typical non-umbrella constraints are those coupling several dimensions of problems as transmission line limits do. Constraints which apply to individual generating units only are typically umbrella constraints since they do not couple multiple dimensions within power system optimization problems. Given the much larger number of generation-based constraints in SCUC in comparison to SCOPF, this is why UCD is not as effective on SCUC as it is on SCOPF. Nevertheless, UCD and P-UCD can contribute substantially to reduce the dimensionality of the transmission constraint set for both SCOPF and SCUC.

In ongoing work, the authors investigate how UCD and its variations can help identify the sources of infeasibility in SCOPF and SCUC problems as well as the application of UCD on ac-SCOPF formulations.

A Security-constrained optimal power flow

The goal of the preventive SCOPF problem is to find the dispatch levels for each generating unit such that the operational cost of a network is minimized while all relevant technical constraints corresponding to pre- and post-contingency states are satisfied simultaneously. Similar to the model used in [6], only line contingencies are considered. In preventive control, the optimal values of the control variables, generation levels here, are to remain constant following the occurrence of any credible line or transformer contingency. In most instances of preventive SCOPF problems, the network technical constraints include: (i) the system power balance; (ii) generation upper and lower production limits; and, (iii) line flow limits enforced on the basis of the dc power flow both in pre- and post-contingency states.

The SCOPF objective function, shown in (11), generally minimizes the cost of operation

$$\min \sum_{i=1}^{I} C_i(p_i) \tag{11}$$

where $C_i(\cdot)$ is the generation cost function at bus i, p_i the generation level at bus i, and I is the number of network buses. The above objective function is subjected to the system-wide load balance, for which we ignore losses here,

$$\sum_{i=1}^{I} p_i = \sum_{i=1}^{I} d_i \tag{12}$$

and where d_i is the load at bus i. Generation levels have to be bounded at each node $i = 1, \ldots, I$ as well

$$\underline{p}_i \leq p_i \leq \bar{p}_i \tag{13}$$

Here $\underline{p}_i$ and $\bar{p}_i$ are respectively the lower and upper dispatch limits of generation at bus i. In addition, there are two constraints associated with each transmission line; one associated with its upper limit and the other with its lower limit. Each line limit constraint pair has to be included to bound the line usage in all pre- and post-contingency states. Hence, for all lines in the network, $\ell = 1, \ldots, L$, and for all line contingency states, $k = 0, \ldots, K$, the SCOPF solution has to satisfy

$$-\bar{f}_\ell(k) \leq \sum_{i=1}^{I} h_{\ell i}(k)(p_i - d_i) \leq \bar{f}_\ell(k) \tag{14}$$

where L is the number of lines of the network, $\bar{f}_\ell(k)$ is the upper limit of line ℓ in contingency state k, and $h_{\ell i}(k)$ is the linear generation shift factor (*a.k.a.* power transmission distribution factor (PTDF)) relating power injection at bus i to the flow in line ℓ while contingency state k is occurring [12, 13]. We point out that contingency state $k = 0$ corresponds to the pre-contingency, *i.e.* normal, operating condition. For ease of exposition here, we shall consider contingency states k where a single transmission component is down ($N-1$ security); however, this can be generalized to multiple component contingencies ($N-m$ security).

B Security-constrained unit commitment

The goal of the SCUC problem is to find the dispatch level and on-off status of each generating unit with the objective of minimizing the operational cost of a network, while all technical constraints are satisfied. These technical constraints include system-wide power balance, system-wide reserve requirements, line flow limits, generation lower and upper bounds, ramping limits and minimum up and down times of generating units in pre- and post-contingency states.

The objective of SCUC, shown in (15), minimizes the cost of operation

$$\min \sum_{i=1}^{I} \sum_{t=1}^{T} [C_i(p_{i,t}) + S_i(u_{i,t})] \tag{15}$$

where $C_i(\cdot)$ is the generation cost function at generating unit i, $p_{i,t}$ is the generation level of generating unit i at time t and I is the number of generating units. Binary variables $u_{i,t}$ denote the respective on ($u_{i,t} = 1$) or off ($u_{i,t} = 0$) status of unit i in period t. The operator of the system also attempts to minimize the costs of starting up units through the $S_i(u_{i,t})$ terms. This is subjected to the system-wide load balance as in (12) which now applies for each time period $t = 1, \ldots, T$, while line flows are restricted as in (14) also for all time periods.

The total system spinning reserve should also be met

$$\sum_{i=1}^{I} r_{i,t} \geq R_t \tag{16}$$

where R_t is the system spinning reserve requirement at time t and $r_{i,t}$ is the contribution of generating unit i to the total reserve in period t. Generation levels have to be bounded for each generating unit and for each time period

$$u_{i,t}\underline{p}_i \leq p_{i,t} \leq u_{i,t}\bar{p}_i \tag{17}$$

The ramp up and down rates for each unit and time period are limited

$$|p_{i,t} - p_{i,t-1}| \leq \Delta_i \tag{18}$$

and individual reserve contributions $r_{i,t}$ also bounded

$$r_{i,t} = u_{i,t} \min\{\bar{r}_i, \bar{p}_i - p_{i,t}\} \tag{19}$$

The above constraints for ramp up and down of each unit should be active only if the unit remains on at time $t-1$ and t. $p_{i,0}$ corresponds to the initial condition of the unit, which is usually determined by previous runs of unit commitment problem. Minimum up and down time for each unit is modeled using the formulation presented in [14] which we omit here for the sake of brevity.

References

[1] F. Capitanescu, J.M. Ramos, P. Panciatici, D. Kirschen, A.M. Marcolini, L. Platbrood, and L. Wehenkel, State-of-the-art, challenges, and future trends in security constrained optimal power flow, Elect. Power Syst. Res., 81(8), 1731–1741, Aug. 2011.

[2] R.P. O'Neill, T. Dautel, and E. Krall, Recent ISO software enhancements and future software and modeling plans, Federal Energy Regulatory Commission, Washington, DC, Tech. Rep., 2011.

[3] F. Bouffard, F.D. Galiana, and A.J. Conejo, Market-clearing with stochastic security—Part II: Case studies, IEEE Trans. Power Syst., 20(4), 1827–1835, Nov. 2005.

[4] F. Bouffard, F.D. Galiana, and J.M. Arroyo, Umbrella contingencies in security-constrained optimal power flow, in Proc. 15th Power Systems Computation Conf., Liège, Belgium, 2005.

[5] Q. Zhai, X. Guan, J. Cheng, and H.Wu, Fast identification of inactive security constraints in SCUC problems, IEEE Trans. Power Syst., 25(4), 1946–1954, Nov. 2010.

[6] A. Ardakani and F. Bouffard, Identification of umbrella constraints in dc-based security-constrained optimal power flow, IEEE Trans. Power Syst., 28(4), 3924–3934, 2013.

[7] IEEE 118 bus power flow test case. [Online]. Available: www.ee.washington.edu/research/pstca/pf118/ieee118psp.txt.

[8] F. Capitanescu, M. Glavic, D. Ernst, and L. Wehenkel, Contingency filtering techniques for preventive security-constrained optimal power flow, IEEE Trans. Power Syst., 22(4), 1690–1697, Nov. 2007.

[9] F. Capitanescu and L. Wehenkel, A new iterative approach to the corrective security-constrained optimal power flow problem, IEEE Trans. Power Syst., 23(4), 1533–1541, Nov. 2008.

[10] Reliability Test System Task Force, The IEEE reliability test system—1996, IEEE Trans. Power Syst., 14(3), 1010–1020, Aug. 1999.

[11] D. Bertsimas and J.N. Tsitsiklis, Introduction to Linear Optimization. Belmont, MA: Athena Scientific, 1997.

[12] A.J. Wood and B.F. Wollenberg, Power Generation Operation and Control, 2nd ed. New York, NY: Wiley-Interscience, 1996.

[13] J. Zhu, Optimization of Power System Operation. Hoboken, NJ: Wiley, 2008, ch. 3, pp. 43–84.

[14] M. Carrion and J. Arroyo, A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem, IEEE Trans. Power Syst., 21(3), 1371–1378, 2006.