

On PMU Location Selection for Line Outage Detection in Wide-area Transmission Networks

Yue Zhao, *Member, IEEE*, Andrea Goldsmith, *Fellow, IEEE*, and H. Vincent Poor, *Fellow, IEEE*

Abstract—The optimal PMU locations to collect voltage phase angle measurements for detecting line outages in wide-area transmission networks are investigated. The problem is established as one of maximizing the minimum distance among the voltage phase angle signatures of the outages, which can be equivalently formulated as an integer programming problem. Based on a greedy heuristic and a linear programming relaxation, a branch and bound algorithm is proposed to find the globally optimal PMU locations. Using this algorithm, the optimal trade-off between the number of PMUs and the outage detection performance is characterized for IEEE 14, 24 and 30 bus systems. The algorithm is shown to find the globally optimal PMU locations in a small number of iterations. It is observed that it is sufficient to have roughly one third of the buses providing PMU measurements in order to achieve the same outage detection performance as with all the buses providing PMU measurements.

Index Terms—Phasor measurement unit, location selection, outage detection, transmission networks, branch and bound

I. INTRODUCTION

In high voltage transmission networks, lack of wide-area situational awareness has been one of the major causes of large-scale blackouts [1]. One of the reasons for the lack of situational awareness has been the limitations of the conventional sensors and SCADA (supervisory control and data acquisition) systems. Phasor measurement units (PMUs), as compared to conventional sensors, are able to provide GPS-synchronized, more accurate and temporally much denser measurements of voltage and current phasors [2]. Thus, the deployment of PMUs for grid sensing is widely considered to be a major driving force for improving the reliability of transmission networks.

There has been considerable research investigating how PMU measurements can be exploited in various tasks for achieving a reliable grid, including outage detection [3], [4], [5], state estimation [6], [7], [8], stability analysis [9], etc. As the cost of installing and networking PMUs is relatively high, two of the major questions are i) what is the necessary or sufficient number of PMUs to use for achieving good performance in all these tasks, and ii) where are the optimal locations to collect PMU measurements? For achieving full network observability, it was estimated that about one third of

the buses need to provide PMU measurements [7]. However, it is unclear whether this estimate applies to *other* kinds of tasks including outage detection. Moreover, for different tasks, the objectives for which the PMU measurements are used differ considerably, and the corresponding optimal PMU locations can vary greatly.

In this paper, we focus on the problem of using PMU measurements of *voltage phase angles* for real time detection of *line outages* over wide areas (including the ones that occur at places where no PMU measurements are available at the control center). We assume that network state estimation during normal conditions is available over relatively slow timescales, (e.g., using NERC System Data Exchange (SDX) [10]), which provides the pre-outage base case system parameters and states. In [3], [4], and [5], it was demonstrated that many line outages can be detected with only a *subset* of the buses providing PMU measurements. However, a comprehensive understanding of the outage detection performance given arbitrary constraints on the number of PMUs to use is left open.

We address this open question of characterizing the optimal trade-off between the number of PMUs used and the outage detection performance. The central problem is finding the *optimal locations* at which to collect PMU measurements, which is an NP hard combinatorial optimization. We consider *both non-adaptive and adaptive* PMU location selections.

We first define a set of voltage phase angle signatures associated with the potential outages. Intuitively, the outage detection performance depends on the degrees of separation among the outage signatures. We employ the minimum distance among the signatures as an indicator of their degrees of separation. Accordingly, we introduce the max-min distance criterion for optimizing the PMU locations given any constraint on the number of PMUs. A simple greedy heuristic is developed that provides lower bounds on the global optimum. Next, we show that finding the optimal PMU locations can be equivalently formulated as an integer programming problem (IP), which allows a relaxation as a linear program (LP) that provides upper bounds on the global optimum. We then develop a branch and bound algorithm for PMU location selection: at each iteration, we strategically fix a bus to either provide or not provide PMU measurements, and compute new lower and upper bounds using the greedy heuristic and LP relaxation.

To understand the optimal trade-off between the number of PMUs and the minimum distance among the outage signatures, we apply the proposed branch and bound algorithm on detecting all single line outages in IEEE 14, 24 and 30 bus systems [11], and show that the globally optimal PMU locations can be found in a small number of iterations. The

This work was supported in part by the NSF under Grant CNS-09-05086, and in part by the DTRA under Grant HDTRA1-08-1-0010.

Y. Zhao is with the Department of Electrical Engineering, Princeton University, Princeton, NJ, 08544 USA, and with the Department of Electrical Engineering, Stanford University, Stanford, CA, 94305 USA (e-mail: yuez@princeton.edu).

A. Goldsmith is with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305 USA (e-mail: andrea@stanford.edu)

H. V. Poor is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: poor@princeton.edu).

optimal trade-offs show that having roughly *one third* of the buses to provide PMU measurements is sufficient for achieving the same outage detection performance as with all the buses providing PMU measurements.

II. PROBLEM FORMULATION

We consider a power transmission network with N buses and L transmission lines. We denote the set of all the buses by $\mathcal{N} = \{1, 2, \dots, N\}$, and the set of all the lines by $\mathcal{L} = \{1, 2, \dots, L\}$. We study the problem of *line outage detection*, where an outage event corresponds to the loss of a subset of lines in \mathcal{L} . Clearly, there are in total $2^L - 1$ different possible outage events, which correspond to all the non-empty subsets of \mathcal{L} . In practice, however, there are several reasons for us to consider a much smaller subset of the outage events:

- In case a single line outage results in the overloading and overheating of another line, it often takes minutes after the initial single line outage for the overheated line to trip. Thus, within a few tens of seconds, it is very unlikely to see many line outages happen together due to their probabilistic independence.
- Not all the outage events have the same level of impact on the grid. For example, some line outages do not lead to the overheating of other lines, while others do. The latter have a greater impact on the grid as they may cause cascading failures.
- As we would like to detect and locate outages in real time with good detection performance, the limitations of computational power may forbid us to consider too many outage events.

Based on the above considerations and the real world situation, a set of *outage events of interest* shall be selected, which we denote by $\mathcal{E} = \{E_1, E_2, \dots, E_K\}$. For example, in [3], [4], \mathcal{E} contains all single and double line outages, respectively.

Given an outage event set \mathcal{E} , we augment it by the *non-outage event* (i.e., normal condition) E_0 , and denote the augmented set by $\tilde{\mathcal{E}} = \{E_0, E_1, \dots, E_K\}$. We investigate the following problem:

- The grid starts in normal condition E_0 .
- Either the grid stays in normal condition $E = E_0$, or an arbitrary outage event $E = E_k \in \mathcal{E}$ occurs.
- We make a detection decision $\hat{E} \in \tilde{\mathcal{E}}$. The detection is successful if $\hat{E} = E$, otherwise not.

We detect outages by observing the *voltage phase angles at a subset of the buses*, where the phasor measurements are provided by PMUs. A system diagram for this problem is depicted in Figure 1. As outages shall be detected as quickly as possible to prevent cascading effects, we aim at detecting an outage within about ten seconds after its occurrence. We note that the phase angles typically stabilize within a few seconds after a line outage occurs [3], and we assume that the power injections of the network remain the same within these few seconds. This is a reasonable assumption as power injections change relatively slowly. We make use of the stabilized phase angles before and after an outage to make detection decisions.

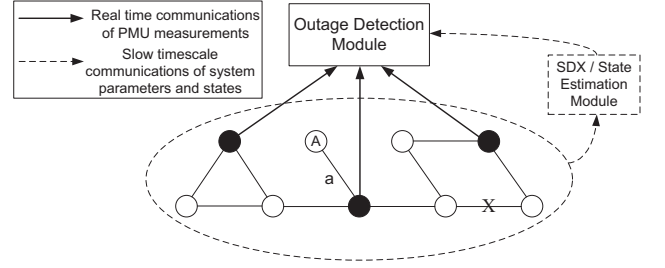


Fig. 1. System diagram for the problem of line outage detection using PMU measurements of voltage phase angles. Buses and transmission lines are represented by circles and connecting segments. The solid circles are the buses where PMU measurements are communicated to the outage detection module in the control center. The solid links represent real time communications of the measurements. The dashed SDX / state estimation module and links provide the slow timescale updates of the system parameters and states to the outage detection module. The loss of line X is an example of single line outage.

A. Pre-computing the Outage Signatures

For each event $E_k \in \mathcal{E}$, we denote the corresponding stabilized phase angles at all the buses by $\theta^{(k)} \in \mathbb{R}^{N \times 1}$. We employ the DC power flow model to compute $\theta^{(k)}$ as follows:

$$\mathbf{P} = \mathbf{B}^{(k)} \theta^{(k)}, k = 0, 1, \dots, K, \quad (1)$$

where $\mathbf{P} \in \mathbb{R}^{N \times 1}$ denotes the power injections, and $\mathbf{B}^{(k)} \in \mathbb{R}^{N \times N}$ is determined by the line status information after event E_k occurs [12]. We note the following properties of $\mathbf{B}^{(k)}$:

- $\mathbf{B}^{(k)} \mathbf{1} = \mathbf{0}$, (2)
and we can arbitrarily add the same constant phase angle to every entry of θ , without having any change of \mathbf{P} .
- $\text{rank}(\mathbf{B}^{(k)}) = N - 1$ if and only if the grid is connected. In the case where the grid contains islands, $\text{rank}(\mathbf{B}^{(k)}) = N - C$ where C is the number of connected components (islands) of the grid.

We note that if an outage creates islands in the grid, the balance of the power injections may not be satisfied by the same \mathbf{P} , and (1) may not hold anymore. For example, in Figure 1, the loss of line a will create a single bus island A which may not be self-balanced (i.e., the pre-outage power injection at A may not be zero). We restrict our consideration to the cases where the post-outage grid remains connected, and leave the islanding cases as future work.

In practice, *noisy* versions of the pre-outage and post-outage phase angle vectors are observed:

$$\theta = \theta^{(0)} + \mathbf{z}^0, \tilde{\theta} = \theta^{(\kappa)} + \mathbf{z}^1, \quad (3)$$

where $E_\kappa, \kappa \in 0, 1, \dots, K$ is the actual outage (or non-outage if $\kappa = 0$) that occurred, and $\mathbf{z}^0, \mathbf{z}^1$ are observation noise vectors that account for the errors in the measurement data and the system parameter data. In the following sections, we assume that the noises at all the buses are i.i.d.. We will generalize this assumption in Section V-A. The task of line outage detection can then be formulated as the following hypothesis testing problem:

From observing the pre-outage and post-outage phase angle vectors $\theta, \tilde{\theta}$ (3), identify which event $E_k \in \tilde{\mathcal{E}}$ has occurred.

For this hypothesis testing problem, $\theta^{(0)}, \theta^{(1)}, \dots, \theta^{(K)}$ can be viewed as the *signatures* of the events E_0, E_1, \dots, E_K .

To identify any event $E_k \in \mathcal{E}$, all the signatures $\{\theta^{(k)}, k = 0, 1, \dots, K\}$ shall be collected before an outage occurs.

From (1), we pre-compute the signatures by

$$\theta^{(k)} = \mathbf{B}^{(k)+} \mathbf{P}, k = 0, 1, \dots, K, \quad (4)$$

where $\mathbf{B}^{(k)+}$ is the pseudoinverse of $\mathbf{B}^{(k)}$. We assume that the knowledge of the power injections \mathbf{P} and the normal condition $\mathbf{B}^{(0)}$ matrix are available from system-wide data sources (e.g., the NERC SDX [10], or other state estimation mechanisms that operate over relatively slow timescales, cf. Figure 1.) Note that, with the knowledge of the open line indices for event E_k , the post-outage $\mathbf{B}^{(k)}$ matrices ($k = 1, 2, \dots, K$) can be derived from the normal condition $\mathbf{B}^{(0)}$ matrix. As a result, while the grid is working under normal conditions, the signatures $\{\theta^{(k)}, k = 0, 1, \dots, K\}$ can be pre-computed by (4) in preparation for detecting the potential line outages in \mathcal{E} . Intuitively, *the more separated the signatures are from each other, the better detection performance can be achieved*. In later sections, we characterize the separation by the minimum distance among the signature set.

From (2), a phase angle vector θ remains functionally equivalent by adding the same constant to all its entries. Without loss of generality (WLOG), we can choose any one of the N buses as the *reference bus*, and subtract the phase angle at the reference bus from the phase angles at all the N buses, such that the phase angle at the reference bus is always kept at zero. In principle, it does not matter which bus among the N buses is chosen to be the reference bus. However, it does matter in the case when only a *subset* of the buses' phase angles are available. This is because once we choose a bus as the reference bus, we implicitly assume that the phase angle at this bus is available.

B. PMU Location Selection

We now establish the main problem of this paper, namely, PMU location selection. In (3), synchronous and timely updates of the complete phase angle vector require *all* the buses to have PMU measurements, and moreover low latency communication links that convey the PMU measurements from all the buses to the control center that performs timely outage detection.

In practice, however, it may not be economically desirable that every bus has a PMU and constantly communicates the PMU measurements to the control center, *particularly when there is very little performance loss with only a fraction of the buses providing PMU measurements*. We motivate the selection of a *subset* of the buses to provide PMU measurements in the following two application scenarios:

- 1) As the cost of PMU installation in high voltage transmission networks is relatively high, we may want to install PMUs only at a subset of the buses to reduce cost.
- 2) Suppose over time the cost of PMUs drops and all (or many) buses have PMUs installed. Due to the *information redundancy* in the PMU measurements, it may not be effective to communicate the data from all the buses to the control center, consuming an unnecessary amount of the expensive low-latency communication

link capacity, and also incurring a longer data collection delay. Thus, the control center may select only a subset of the PMUs to provide measurements.

In the following sections, we address the second application scenario, in which the PMU location selection can adapt to the changes of the power injections \mathbf{P} and the network \mathbf{B} matrix over time. We show later in section V-B that the results can be directly extended to the first application scenario where a non-adaptive PMU location selection is required.

The problem of PMU location selection entails the following two questions:

- *Given that we want to choose M ($2 \leq M \leq N$) of the N buses to provide PMU measurements, which M buses should we choose, and how can we characterize the corresponding outage detection performance?*
- *What is the optimal tradeoff between the number of PMUs M and the outage detection performance?*

As mentioned in the last section, it makes no sense to have a bus chosen as the reference bus without having a PMU measuring its phase angle. Therefore, to find the optimal PMU locations, we cannot simply choose a reference bus arbitrarily, as it does lose generality. Instead, we *traverse all the N cases of choosing each bus as the reference bus*.

From now on, we assume that *we have chosen a reference bus r during this outer traversal, and optimize the location selection of the other $M - 1$ PMUs*. We assume that all the phase angle vectors have been adjusted so that their r^{th} entries always equal 0:

$$\theta_r^{(k)} = 0, \forall k = 0, 1, \dots, K.$$

Accordingly, we denote by \mathcal{M} the subset of buses with PMU measurements in addition to r , $\mathcal{M} \subseteq \{1, 2, \dots, N\}$, $|\mathcal{M}| = M - 1$. As a result, for any two phase angle vectors, we cannot distinguish their entries at the other $N - M + 1$ buses in $\mathcal{N} \setminus \mathcal{M}$. Therefore, for any phase angle vector θ , we project it into a sub-vector by extracting the $M - 1$ entries of θ whose indices are in \mathcal{M} (it is not necessarily an orthogonal projection) and denote the projected vector by $\theta_{\mathcal{M}} \in \mathbb{R}^{(M-1) \times 1}$. Accordingly, every set \mathcal{M} leads to a set of projected signatures $\{\theta_{\mathcal{M}}^{(k)}, k = 0, 1, \dots, K\}$.

Clearly, with different choices of \mathcal{M} , the *degrees of separation* among the projected signatures can differ considerably. Intuitively we want to optimize \mathcal{M} to get better separation among the projected signatures, and thus better detection performance.

We model the separation among the projected signatures (and hence the outage detection performance) by the *minimum distance* in p -norm among them:

$$d_{\min}(\mathcal{M}) = \min_{0 \leq i < j \leq K} \|\theta_{\mathcal{M}}^{(i)} - \theta_{\mathcal{M}}^{(j)}\|_p, \quad (5)$$

where p is a parameter to choose, and $p = 2$ corresponds to the Euclidian distance. Given M as the total number of PMUs to use, \mathcal{M} can then be optimized under the following *Max-Min Distance Criterion*:

$$\max_{\mathcal{M}, |\mathcal{M}|=M-1, r \notin \mathcal{M}} d_{\min}(\mathcal{M}). \quad (6)$$

Clearly, (6) is an NP hard combinatorial optimization of the set \mathcal{M} .

C. An Integer Programming Formulation

Note that $d_{\min}(\mathcal{M})$ can be re-written as

$$d_{\min}(\mathcal{M}) = (\min(\mathbf{w}^T \Theta))^{\frac{1}{p}}, \quad (7)$$

where

$$\bullet \quad \mathbf{w} \in \mathbb{R}^{N \times 1}, \text{ and } w_i = \begin{cases} 1, & \text{if } i \in \mathcal{M} \text{ or } i = r, \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

$$\bullet \quad \Theta \in \mathbb{R}^{N \times \binom{K+1}{2}}, \text{ and its columns are constructed by collecting the following } \binom{K+1}{2} N \times 1 \text{ vectors:}$$

$$\forall 0 \leq i < j \leq K,$$

$$|\theta^{(i)} - \theta^{(j)}|^p, \quad (9)$$

where the $|\cdot|^p$ operation is applied *elementwise*.

In other words, the indices of the non-zero entries of \mathbf{w} denote the M buses that are chosen to provide PMU measurements (including the reference bus r). We name the binary vector \mathbf{w} the *bus selection indicator vector*. Consequently, (6) is equivalent to the following integer programming problem:

$$\max_{\mathbf{w}} \min(\mathbf{w}^T \Theta) \quad (10)$$

$$\text{s.t. } w_i \in \{0, 1\}, i = 1, \dots, N, \quad (11)$$

$$\sum_{i=1}^N w_i = M, \quad w_r = 1.$$

Clearly, finding the global optimum of (10) requires a worst-case computational complexity of $\binom{N-1}{M-1}$.

III. PMU LOCATION SELECTION WITH MAX-MIN DISTANCE CRITERION

In this section, we provide two algorithms for solving the combinatorial optimization of PMU location selection (6).

A. A Greedy Heuristic

We develop a greedy heuristic that generates a series of PMU location selection solutions $\mathcal{M}_2, \mathcal{M}_3, \dots, \mathcal{M}_N$ for $M = 2, 3, \dots, N$ respectively, that satisfy the following *consistency property*:

$$\mathcal{M}_2 \subset \mathcal{M}_3 \subset \dots \subset \mathcal{M}_N. \quad (12)$$

We present the greedy algorithm as follows in a slightly more general form with an arbitrary initial set of selected buses \mathcal{M}_{ini} :

Algorithm 1: Greedy PMU Location Selection

Given M and an initial set of buses \mathcal{M}_{ini} :

$m = |\mathcal{M}_{ini}|, \mathcal{M}_m = \mathcal{M}_{ini}.$

Repeat

$m \leftarrow m + 1,$

$$\mathcal{M}_m = \mathcal{M}_{m-1} \cup \left\{ \underset{n \in \mathcal{N} \setminus \mathcal{M}_{m-1}}{\operatorname{argmax}} d_{\min}(\mathcal{M}_{m-1} \cup \{n\}) \right\} \quad (13)$$

Until $m = M.$

In other words, given the total number of PMUs M and

a starting set of chosen buses \mathcal{M}_{ini} , we choose another $M - |\mathcal{M}_{ini}|$ buses *one by one*: At each step, we keep the already chosen buses; from the remaining buses, we choose the one that *maximizes the current step's minimum distance*, and include it in the set of the chosen buses. When the only prior knowledge of the bus selection is the chosen reference bus r , $\mathcal{M}_{ini} = \{r\}$, and Algorithm 1 generates the set of greedy consistent solutions (12) as M increases from 2 to N .

B. A Branch and Bound Algorithm

First, we note that the integer programming formulation (10) has a concave objective function, and hence has the following *relaxation* as a convex optimization:

$$\max_{\mathbf{w}} \min(\mathbf{w}^T \Theta) \quad (14)$$

$$\text{s.t. } 0 \leq w_i \leq 1, i = 1, \dots, N, \quad (15)$$

$$\sum_{i=1}^N w_i = M, \quad w_r = 1.$$

In fact, this convex optimization can be equivalently cast as a linear program [13]. Accordingly, the optimal value of (14) serves as an *upper bound*, denoted by U_1 , on the global optimum of (10). Meanwhile, Algorithm 1 with $\mathcal{M}_{ini} = \{r\}$ provides a *lower bound*, denoted by L_1 .

Remark 1 (A note on the rounding heuristic): Another lower bounding heuristic is to find an integral solution by *rounding* the fractional solution obtained from the relaxed problem (14). In particular, we consider the heuristic by rounding *the M largest* fractional entries to 1, and the others to 0. In all scenarios that we simulated, we compared this rounding heuristic with the greedy heuristic: somewhat to our surprise, *the greedy heuristic uniformly outperforms the rounding heuristic*. The reason is that, in the relaxed fractional solution, it is often some *very small* non-zero entries that are *critical* in the sense that losing them will drastically reduce the minimum distance. Consequently, for lower bounding the global optimum, we propose the greedy heuristic instead of the rounding heuristic, as the former is both much cheaper computationally and much better in performance. Thus, we use the relaxation technique *not* to provide a fractional solution to round, but to *upper bound* the global optimum and to develop a *branch and bound* algorithm that can significantly improve the greedy solutions in a few iterations.

For any bus $n \neq r$, (10) can be split into two sub-problems by fixing w_n to be either 0 or 1:

$$\max_{\mathbf{w}} \min(\mathbf{w}^T \Theta) \quad (16)$$

$$\text{s.t. } w_i \in \{0, 1\}, i = 1, \dots, N, \quad (17)$$

$$\sum_{i=1}^N w_i = M, \quad w_n = 0, w_r = 1.$$

and

$$\max_{\mathbf{w}} \min(\mathbf{w}^T \Theta) \quad (18)$$

$$\text{s.t. } w_i \in \{0, 1\}, i = 1, \dots, N, \quad (19)$$

$$\sum_{i=1}^N w_i = M, \quad w_n = 1, w_r = 1.$$

Similarly to (14), relaxations of these two sub-problems can be formed by replacing (17) and (19) with (15), and they provide upper bounds, denoted by $u_2^{(0)}$ and $u_2^{(1)}$, on the global optimum of (16) and (18) respectively. Meanwhile, applying the greedy heuristic under the constraint $w_n = 0$ or $w_n = 1$ provides lower bounds, denoted by $l_2^{(0)}$ and $l_2^{(1)}$, on these sub-problems' global optima. Define

$$U_2 \triangleq \max\{u_2^{(0)}, u_2^{(1)}\}, \text{ and } L_2 \triangleq \max\{l_2^{(0)}, l_2^{(1)}\}. \quad (20)$$

Then, U_2 and L_2 are new upper and lower bounds on the original global optimum (10) [14].

More generally, the above splitting procedure with relaxations and greedy heuristics can be applied on the sub-problems themselves to form more children sub-problems with upper and lower bounds. For example, for any bus $s, (s \neq n, r,)$ (16) can be further split into two sub-problems by adding yet another constraint $w_s = 0$ or $w_s = 1$ respectively.

We define the following upper and lower bounding oracles, as well as an oracle that returns the next bus to split:

Definition 1: Oracle $UB(\mathcal{C})$ takes a *constraint set* \mathcal{C} as input, where \mathcal{C} specifies a set of buses whose selection indicator variables are pre-determined to be either 0 or 1. An IP under the constraints \mathcal{C} is formed, a relaxation is solved, and the optimum of this relaxation is output by $UB(\mathcal{C})$ as an upper bound on the optimum of the constrained IP.

For example, in (16) and (18), the constraint sets are $\mathcal{C}^{(0)} = \{w_n = 0, w_r = 1\}$ and $\mathcal{C}^{(1)} = \{w_n = 1, w_r = 1\}$, respectively.

Definition 2: Oracle $LB(\mathcal{C})$ takes a constraint set \mathcal{C} as input. An IP under the constraints \mathcal{C} is formed, a greedy solution is found by Algorithm 1, and the achieved objective value is output by $LB(\mathcal{C})$ as a lower bound on the optimum of the constrained IP.

Definition 3: Based on the *order* of the buses chosen by Algorithm 1, Oracle $next(\mathcal{C})$ outputs *the first* bus that is chosen by this heuristic.

When a sub-problem with constraints \mathcal{C} needs to be split further, $next(\mathcal{C})$ is the bus we choose to perform the splitting by fixing $w_{next(\mathcal{C})}$ to be either 0 or 1.

We now provide a branch and bound algorithm as in Algorithm 2 where i_{\max} is the maximum number of iterations allowed. As the algorithm progresses, a *binary tree is developed where each node represents a constraint set*. The *leaf nodes* are kept in \mathcal{S} . The tree starts with a single node, $\{w_r = 1\}$, corresponding to the prior knowledge that the reference bus is r who uses a PMU. When a sub-problem corresponding to a leaf node \mathcal{C}^* is split into two new sub-problems, the two new constraint sets $\mathcal{C}^{(0)}$ and $\mathcal{C}^{(1)}$ become the children of the parent constraint set \mathcal{C}^* .

In Algorithm 2, (22) is a generalization of (20). It means that the current global upper bound equals the *highest* upper bound among all the *leaf node* constraint sets. This is true because all the leaf nodes \mathcal{S} represent a *complete partition* of the original parameter space [14]. At the beginning of every iteration, in choosing which leaf node to split (21), we select the one that gives the *highest* upper bound (i.e., the current *global* upper bound). It is a heuristic based on the reasoning that, by further splitting this critical leaf node, a lower global upper bound

may be obtained, (whereas splitting any other node will leave the global upper bound unchanged.) At iteration i , the current upper and lower bounds on the global optimum are available as U_i and L_i . When these two bounds meet, i.e., $U_i - L_i < \epsilon$, the solution that achieves the current lower bound is guaranteed to be globally optimal.

Algorithm 2:

PMU Location Selection using Branch and Bound

Initial step: $i = 1$,

the initial constraint set: $\mathcal{C}_1 = \{w_r = 1\}$,

the initial set of leaves of the tree of constraint sets (initially a single node): $\mathcal{S} = \{\mathcal{C}_1\}$.

Compute $U_1 = UB(\mathcal{C}_1)$, $L_1 = LB(\mathcal{C}_1)$.

While $U_i - L_i > \epsilon$ or $i < i_{\max}$, repeat

Choose which leaf node constraint set to split:

$$\mathcal{C}^* = \operatorname{argmax}_{\mathcal{C} \in \mathcal{S}} \{UB(\mathcal{C})\}. \quad (21)$$

Choose the next bus to split, $n = next(\mathcal{C}^*)$,

Form two new constraint sets,

$$\mathcal{C}_{i+1}^{(0)} = \mathcal{C}^* \cup \{w_n = 0\}, \mathcal{C}_{i+1}^{(1)} = \mathcal{C}^* \cup \{w_n = 1\}.$$

In the set of leaves \mathcal{S} , replace the parent constraint set \mathcal{C}^* with the two children $\mathcal{C}_{i+1}^{(0)}$ and $\mathcal{C}_{i+1}^{(1)}$:

$$\mathcal{S} \leftarrow (\mathcal{S} \setminus \{\mathcal{C}^*\}) \cup \{\mathcal{C}_{i+1}^{(0)}\} \cup \{\mathcal{C}_{i+1}^{(1)}\}.$$

Compute new upper and lower bounds for the two new constrained IP:

$$UB(\mathcal{C}_{i+1}^{(0)}), UB(\mathcal{C}_{i+1}^{(1)}), LB(\mathcal{C}_{i+1}^{(0)}), LB(\mathcal{C}_{i+1}^{(1)}),$$

Update the global upper and lower bounds,

$$U_{i+1} = \max_{\mathcal{C} \in \mathcal{S}} \{UB(\mathcal{C})\}, L_{i+1} = \max_{\mathcal{C} \in \mathcal{S}} \{LB(\mathcal{C})\}. \quad (22)$$

$i \leftarrow i + 1$.

Choose the best achieved solution so far:

$$\hat{\mathcal{C}} = \operatorname{argmax}_{\mathcal{C} \in \mathcal{S}} LB(\mathcal{C}).$$

Return the greedy solution under the constraint set $\hat{\mathcal{C}}$.

We note that Algorithm 1 is a degraded version of Algorithm 2 with just one iteration. As the total number of possible constraint sets is 2^N (corresponding to the 2^N bus selection indicator vectors \mathbf{w}), Algorithm 2 is guaranteed to converge in 2^N iterations (and in practice much less as will be shown later.) To limit the algorithm's run time, a maximum number of iterations i_{\max} can be enforced as in Algorithm 2.

Finally, we define i_{achieve} to be the number of iterations used to *achieve* the globally optimal solution, and i_{prove} the number of iterations used to *prove* its global optimality. In other words, it takes i_{achieve} iterations for the *lower bound* to reach the global optimum, while it takes i_{prove} iterations for *both* the upper and lower bounds to reach the global optimum. As will be shown next, typically we have $i_{\text{achieve}} \ll i_{\text{prove}}$.

IV. PERFORMANCE EVALUATION AND THE OPTIMAL M - d_{\min} TRADEOFF

In this section, we simulate the proposed algorithms in IEEE 14, 24, and 30 bus systems using the software toolbox MATPOWER [11]. We set the outage event set \mathcal{E} to be

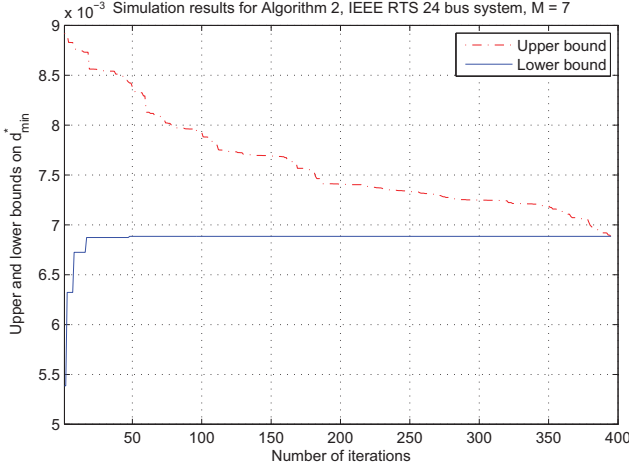


Fig. 5. The instant upper and lower bounds on the maximum achievable minimum distance as Algorithm 2 iterates; IEEE RTS 24 bus system, $M = 7$.

- In these three systems, Algorithm 2 always finds the globally optimal solution in *less than 19 iterations*. For solving the $M = 4$ case in the 30 bus system, $i_{\text{achieve}} = 19$. For most other cases, i_{achieve} is much smaller.
- When M is relatively small, it takes a much larger number of iterations to *prove* that the solutions achieved within 19 iterations are indeed globally optimal. The maximum i_{prove} is 395 for solving the $M = 7$ case in the 24 bus system. For this particular case, the upper and lower bounds achieved as Algorithm 2 progresses are plotted in Figure 5.
- i_{achieve} and i_{prove} are much smaller than the *combinatorial complexity* of this NP hard problem. E.g., for solving the case of $M = 7$ in the 24 bus system, an exhaustive search has to traverse $\binom{24}{7} = 346104$ PMU location selections, whereas Algorithm 2 finds the optimal solution in $i_{\text{achieve}} = 17$ iterations, and verifies its optimality in $i_{\text{prove}} = 395$ iterations.

The fact that i_{achieve} is typically small demonstrates that Algorithm 2 is very efficient in finding the globally optimal PMU locations. We note that, in practice, the sometimes larger i_{prove} does not matter at all. This is because we always use the best solution found within a predetermined run time, and as long as i_{achieve} is sufficiently small, we will find the optimal solution (albeit without proving its optimality).

Finally, as an example of the optimal PMU location selection, we depict the buses that are chosen for the 30 bus system with $M = 10$ (cf. Figure 6). In this case, bus 1, 5, 8, 9, 14, 21, 22, 24, 26, 29 are chosen to provide PMU measurements, achieving the same minimum distance among the outage signatures as with all the 30 buses chosen. As shown in Table 3, since $i_{\text{achieve}} = 1$ for $M = 10$, this globally optimal solution can in fact be found by the simple greedy method of Algorithm 1.

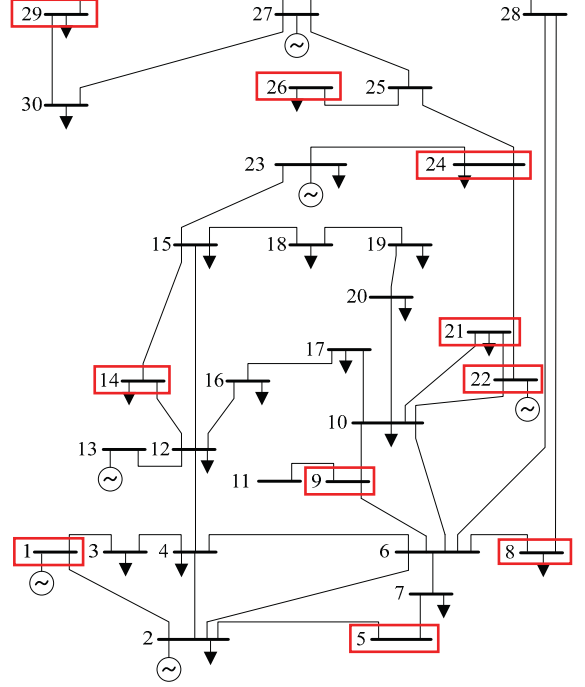


Fig. 6. One line diagram of the IEEE 30 bus system, and the optimal PMU location selection with $M = 10$. The buses enclosed by rectangles are the optimally chosen buses.

V. DISCUSSION

A. Normalizing Distances by Noise Variances

The max-min distance criterion we have used is for countering the effect of the noise (3) by separating the outage signatures as much as possible. In general, the variances of the noise in the PMU measurements *may differ at different buses*. This shall be taken into account when defining the *effective distances* among the outage signatures.

For two outage signatures $\theta^{(i)}$ and $\theta^{(j)}$, denote their difference in absolute value by $\Delta\theta = |\theta^{(i)} - \theta^{(j)}|$, where $|\cdot|$ is applied elementwise. The n^{th} entry $\Delta\theta_n$ is the distance between the two signatures *at the n^{th} bus*.

Suppose now that the noises are independently Gaussian with variances $\sigma_1^2, \sigma_2^2, \dots, \sigma_N^2$ at the N buses. We can then *normalize* the distance at the n^{th} bus by a factor of $\frac{1}{\sigma_n}$:

$$\Delta\tilde{\theta} = \Sigma^{-1}\Delta\theta, \quad (24)$$

where $\Sigma = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_N)$. After this normalization, it is now appropriate to *add up* all the buses' contributions in separating the signatures by computing the p -norm of $\Delta\tilde{\theta}$:

$$\|\Delta\tilde{\theta}\|_p = \left(\sum_{n=1}^N \Delta\tilde{\theta}_n^p \right)^{\frac{1}{p}}. \quad (25)$$

Accordingly, the definition of the minimum distance among the signatures (5) is generalized as follows:

$$\tilde{d}_{\min}(\mathcal{M}) = \min_{0 \leq i < j \leq K} \left\| \left(\Sigma^{-1} \left(\theta^{(i)} - \theta^{(j)} \right) \right) \right\|_{\mathcal{M}} \Big|_p$$

For the IP formulation (10), it was assumed that $\Sigma = \sigma \mathbf{I}$ for some σ in the previous definition of Θ (9). With an arbitrary Σ , the matrix of the distances among the outage signatures is then generalized by the following normalization:

$$\tilde{\Theta} = \Sigma^{-p} \Theta. \quad (26)$$

B. Extension to Non-Adaptive PMU Location Selection

We have provided two application scenarios in Section II-B that motivate the problem of PMU location selection. In the previous sections, we have focused on the second scenario in which the outage signatures $\{\theta^{(k)}, k = 0, 1, \dots, K\}$, and hence the optimal PMU location selection, depend on the current network setting \mathbf{P} and \mathbf{B} .

In the first scenario, however, location selection for *PMU installation* is a planning problem whose solution must work *non-adaptively* in the sense that the installed PMUs are not transferable to other locations. As a result, the PMU locations must be chosen to accommodate, if not all, *the most typical* network settings in terms of \mathbf{P} and \mathbf{B} . As we consider wide-area transmission networks, this requirement is addressed as follows:

- The network topology \mathbf{B} is in general slowly and slightly changing, and an approximate estimate of it would suffice for computing typical sets of outage signatures.
- Typical power injections \mathbf{P} are available from grid statistics for high voltage transmission networks: collecting the typical data for *different seasons* and during *daytime and nighttime* would suffice for computing typical sets of outage signatures.

With a set of typical power injections $\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_T$, and an estimated \mathbf{B} , there are T typical sets of outage signatures:

$$\{\theta^{(k,t)}, k = 0, 1, \dots, K, t = 1, 2, \dots, T. \quad (27)$$

The minimum distance (5) can be generalized as

$$d_{\min}(\mathcal{M}) = \min_{1 \leq t \leq T} \min_{0 \leq i < j \leq K} \|\theta_{\mathcal{M}}^{(i,t)} - \theta_{\mathcal{M}}^{(j,t)}\|_p. \quad (28)$$

Similarly to (9), we construct $\Theta \in \mathbb{R}^{N \times T \cdot \binom{K+1}{2}}$, and its columns are constructed by collecting the following $T \cdot \binom{K+1}{2}$ $N \times 1$ vectors: $\forall 1 \leq t \leq T, \forall 0 \leq i < j \leq K$,

$$|\theta^{(i,t)} - \theta^{(j,t)}|^p, \quad (29)$$

where the $|\cdot|^p$ operation is applied *elementwise*.

The problem of non-adaptive PMU location selection with (28) can then be formulated as the same integer programming problem (10), for which Algorithm 1 and 2 can be applied.

VI. CONCLUDING REMARKS

We have studied the problem of PMU location selection for detecting line outages in wide-area transmission networks. We have first introduced the voltage phase angle signatures of the outages. We have then established the PMU location selection problem as maximizing the minimum distance among the set of outage signatures, and have shown that it can be equivalently formulated as an integer programming problem. For this NP hard problem, we have developed a greedy

heuristic and a linear programming relaxation, providing a lower and an upper bound on the global optimum respectively. Based on this heuristic and relaxation, we have proposed a branch and bound algorithm to find the globally optimal PMU locations. Using this algorithm, we have characterized the optimal trade-offs between the number of PMUs and the minimum distance among the outage signatures in IEEE 14, 24 and 30 bus systems. For all the simulated cases, the optimal PMU locations are found in at most 19 (and in most cases a much less number of) iterations. From the optimal trade-offs, we have observed that it is sufficient to have roughly one third of the buses to provide PMU measurements in order to achieve the same outage detection performance as with all the buses providing PMU measurements.

For future work, while we have employed the max-min distance criterion among the outage signatures, it is interesting to investigate the use of other criteria in characterizing outage detection performance. We would also like to examine the optimal PMU location selections with outage event sets more general than line outages. Finally, it remains an interesting open question how to take into account the needs from *multiple* tasks simultaneously (e.g., outage detection *and* state estimation) while selecting PMU locations.

REFERENCES

- [1] *US-Canada Power System Outage Task Force, Final Report on August 14, 2003 the Blackout in the United States and Canada*, 2004.
- [2] A. G. Phadke, "Synchronized phasor measurements in power systems," *IEEE Comput. Appl. Power*, vol. vol.6, no.2, pp. 10 – 15, Apr. 1993.
- [3] J.E. Tate and T.J. Overbye, "Line outage detection using phasor angle measurements," *IEEE Transactions on Power Systems*, vol. vol.23, no.4, pp. 1644 – 1652, Nov. 2008.
- [4] J.E. Tate and T.J. Overbye, "Double line outage detection using phasor angle measurements," *Proc. of IEEE Power and Energy Society General Meeting, Calgary, Canada*, July 2009.
- [5] H. Zhu and G.B. Giannakis, "Lassoing line outages in the smart power grid," *Proc. of 2nd Intl. Conf. on Smart Grid Communications, Brussels, Belgium*, Oct. 2011.
- [6] L. Zhao and A. Abur, "Multi area state estimation using synchronized phasor measurements," *IEEE Transactions on Power Systems*, vol. vol.20, no.2, pp. 611 – 617, May 2005.
- [7] T.L. Baldwin, L. Mili, Jr. M.B. Boisen, and R. Adapa, "Power system observability with minimal phasor measurement placement," *IEEE Transactions on Power Systems*, vol. vol.8, no.2, pp. 707 – 715, May 1993.
- [8] B. Xu and A. Abur, "Observability analysis and measurement placement for systems with PMUs," *IEEE PES Power Systems Conference and Exposition, New York, USA*, vol. vol.2, pp. 943 – 946, Oct. 2004.
- [9] K. Sun, S. Likhate, V. Vittal, V.S. Kolluri, and S. Mandal, "An online dynamic security assessment scheme using phasor measurements and decision trees," *IEEE Transactions on Power Systems*, vol. vol.22, no.4, pp. 1935 – 1943, Nov. 2007.
- [10] *US-Canada Power System Outage Task Force, Final Report on the Implementation of the Task Force Recommendations*, 2006.
- [11] R. D. Zimmerman, C. E. Murillo-Sanchez, and R. J. Thomas, "MAT-POWER steady-state operations, planning and analysis tools for power systems research and education," *IEEE Transactions on Power Systems*, vol. vol.26, no.1, pp. 12 – 19, Feb. 2011.
- [12] A. J. Wood and B. F. Wollenberg, *Power Generation, Operation, and Control*, John Wiley and Sons, New York, NY, USA, 2nd edition, 1996.
- [13] S.P. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, Cambridge, UK, 2004.
- [14] V. Balakrishnan, S. Boyd, and S. Baleni, "Branch and bound algorithm for computing the minimum stability degree of parameter-dependent linear systems," *Int. J. of Robust and Nonlinear Control*, vol. 1(4), pp. 295 – 317, October - December 1991.