1 Detection and Isolation of Power System Transmission Line Outages

1.1 Introduction

Many online decision-making tools used by the operators of a power transmission system, e.g., the state estimator, and contingency analysis and generation re-dispatch procedures, rely on a system model obtained offline assuming a certain system network topology. If the network topology changes due to a transmission line outage, the power system model needs to be updated; otherwise, the aforementioned decision-making tools may provide an inaccurate picture of the system state. For example, in the 2011 San Diego blackout , the system operators were not able to determine that certain transmission lines were overloaded because the system model was not updated following a transmission line outage [1]. The lack of situational awareness in this case prevented system operators from determining the next critical contingency and led to a cascading failure. Similarly, during the 2003 US Northeastern blackout, system operators were unaware of the loss of key transmission lines and as a results they did not initiate the necessary remedial actions schemes that could have prevented the cascading failure that ultimately led to the blackout [2].

In light of the discussion above, it is clear that there is a need for developing online techniques that enable detection and identification of transmission line outages in a timely fashion. In this chapter, we address this problem and discuss methods for online detection and identification of line outages in transmission networks recently developed by the authors, their students, and their postdoctoral research associates [3, 4, 5, 6, 7]; the foundations of the methods lie in the theory of quickest change detection (QCD) (see, e.g., [8, 9, 10]). In QCD, a decision maker observes a sequence of random variables, the probability distribution of which changes abruptly due to some event that occurs at an unknown time. The objective is then to detect this change in distribution as quickly as possible, subject to false alarm constraints.

Our line outage detection method relies on the fact that in transmission networks the incremental changes in the voltage phase angle at each bus of the network can be approximated by a linear combination of incremental changes in the power demand at all load buses. The coefficients of this linear combination can be easily obtained from the network admittance matrix (the sparsity pattern of which depends on the network topology) and some time-varying coefficients capturing the response of generators in the system to small changes in active power demand. An outage in any line will result in a change in the system admittance matrix. As such, the coefficients of the linear combination relating active power demands and voltage phase angles will change after a line outage. By assuming that the incremental changes in the power demand at load buses can be modeled as independent zero-mean Gaussian random variables, one can approximately obtain the probability distribution of the bus voltage angle incremental changes before the occurrence of any line outage, in a short transient period immediately after a line outage (governed by aforementioned coefficients capturing generator responses), and after all transient phenomena have died down and the system settles into a new steady state. Then, by using bus voltage phase angle measurements provided by phasor measurements units (PMUs), one can obtain a random sequence of incremental changes in the voltage phase angle measurements in real-time, and feed it to a QCD-based algorithm to detect a change in the probability distribution of this random sequence.

Earlier algorithms for topology change detection in the literature include those based on state estimation [11, 12], and rules that mimic system operator decisions [13]. More recent techniques have utilized measurements provided by PMUs for line outage detection [14, 15, 16]. All of these works rely on the voltage phase angle difference between two sets of measurements obtained before and after the outage occurrence, and attempt to identify the line outage by using hypothesis testing [14], sparse vector estimation [15], and mixed-integer nonlinear optimization [16]. However, these techniques are "one-shot" in the sense that they do not exploit the persistent nature of line outages and do not incorporate the transient behavior immediately after the outage occurrence. In addition, the work in [17] proposes a method to identify line outages using statistical classifiers where a maximum likelihood estimation is performed on the PMU data. This work considers the system transient response after the occurrence of a line outage by comparing data obtained offline via simulation against actual data. However, this method requires the exact time the line outage occurs to be known before executing the algorithm, i.e., it assumes line outage detection has already been accomplished by some other method, whereas the methods discussed in this chapter provide both detection and identification.

The remainder of this chapter is organized as follows. In Section 1.2, we provide some background on the theoretical foundations of QCD. In Section 1.3 we describe the power system model adopted in this work, and introduce the statistical model describing the voltage phase angle before, during, and after the occurrence of an outage. In Section 1.4, we describe several QCD-based line outage detection algorithms. Section 1.5 discusses how to identify line outages using QCD-based algorithms. In Section 1.6, we illustrate and compare the performance of the QCD-based line outage detection and identification algorithms discussed in Section 1.4 via numerical case studies on the IEEE 118-bus test system. Some concluding remarks are provided in Section 1.7.

1.2 Quickest Change Detection Background

In the basic quickest change detection (QCD) problem, we have a sequence of observations $\{X[k], k \geq 1\}$, which initially have certain distribution described by some probability density function (pdf), and at some point in time γ, due to some event, the distribution of the observations changes. We restrict our attention to models where the observations are independent and identically distributed (i.i.d.) in the pre- and post-change regimes since this is a good approximation for the line outage detection problem as we will see in Section 1.3. Thus

$$\begin{aligned} X[k] &\sim \text{i.i.d. with pdf } f_0 \text{ for } k < \gamma, \\ X[k] &\sim \text{i.i.d. with pdf } f_1 \text{ for } k \geq \gamma. \end{aligned} \tag{1.1}$$

We assume that f_0 and f_1 are known but that γ is unknown. We call γ the *change point*.

The goal in QCD is to detect the change in the distribution as quickly as possible (after it happens) subject to constraints on the rate of false alarms. A QCD procedure is constructed through a stopping time[1] τ on the observation sequence $\{X[k], k \geq 1\}$, with the change being declared at the stopping time. To have quick detection, we need to minimize some metric that captures the delay in detection $(\tau - \gamma)^+$ (where we use the notation $x^+ = \max\{x, 0\}$). Clearly, selecting $\tau = 1$ always minimizes this quantity, but if $\gamma > 1$, we end up with a false alarm event. The QCD problem is to find a stopping time τ that results in an optimal tradeoff between the detection delay and false alarm rate.

Before we discuss specific algorithms for the QCD problem, we introduce the Kullback-Leibler (KL) divergence, which is an information-theoretic measure of the discrepancy between two probability distributions, f and g, defined as:

$$D(f \| g) := \mathbb{E}_f \left[\log \frac{f(X)}{g(X)} \right] = \int f(x) \log \frac{f(x)}{g(x)} dx. \tag{1.2}$$

It is easy to show that $D(f\|g) \geq 0$, with equality if and only if $f = g$. The KL divergence plays an important role both in developing good QCD algorithms and in characterizing their performance.

We next introduce the likelihood ratio of the observations

$$L(X[k]) = \frac{f_1(X[k])}{f_0(X[k])}.$$

It is well known that the likelihood ratio plays a key role in the construction of good detection algorithms in general, not just for the QCD problem being discussed here [18].

A fundamental fact that is used in the construction of good QCD algorithms is that the mean of the log-likelihood ratio $\log L(X[k])$ in the pre-change regime, i.e., for $k < \gamma$ is given by

$$\mathbb{E}_{f_0} [\log L(X[k])] = -D(f_0 \| f_1) < 0$$

[1] See, e.g., [10] for the definition of a stopping time.

and in the post-change regime, $k \geq \gamma$, is given by

$$\mathbb{E}_{f_1}\left[\log L(X[k])\right] = -D(f_1 \| f_0) > 0.$$

1.2.1 Shewhart Test

A simple way to use the properties of the log-likelihood ratio $\log L(X[k])$ was first proposed by Shewhart [19], in which a statistic based on the current observation is compared with a threshold to make a decision about the change. That is, the Shewhart test is defined as

$$\tau_{\mathrm{S}} = \inf\{k \geq 1 : \log L(X[k]) > b\} \tag{1.3}$$

where b is chosen to meet a desired false alarm constraint. Shewhart's test is widely employed in practice due to its simplicity; however, significant gain in performance can be achieved by making use of past observations (in addition to the current observation) to make the decision about the change.

1.2.2 CuSum Test

In [20], Page proposed an algorithm that uses past observations, which he called the Cumulative Sum (CuSum) algorithm. The idea behind the CuSum algorithm is based on the behavior of the cumulative log-likelihood ratio sequence:

$$S[k] = \sum_{j=1}^{n} \log L(X[j]).$$

Before the change occurs, the statistic $S[k]$ has a negative "drift" due to the fact that $\mathbb{E}_0[\log L(X[j]) < 0$ and diverges towards $-\infty$. At the change point, the drift of the statistic $S[k]$ changes from negative to positive due to the fact that $\mathbb{E}_1[\log L(X[j]) > 0$, and beyond the change point $S[k]$ starts growing towards ∞. Therefore $S[k]$ roughly attains a minimum at the change point. The CuSum algorithm is constructed to detect this change in drift, and stop the first time the growth of $S[k]$ after change in the drift is large enough (to avoid false alarms).

Specifically, the stopping time for the CuSum algorithm is defined as

$$\tau_{\mathrm{C}} = \inf\left\{k \geq 1 : W[k] \geq b\right\} \tag{1.4}$$

where

$$\begin{aligned} W[k] &= S[k] - \min_{0 \leq j \leq k} S[j] \\ &= \max_{0 \leq j \leq k} S[k] - S[j] \\ &= \max_{0 \leq j \leq k} \sum_{\ell=j+1}^{k} \log L(X[\ell]) = \max_{1 \leq j \leq k+1} \sum_{\ell=j}^{k} \log L(X[\ell]) \end{aligned} \tag{1.5}$$

with the understanding that $\sum_{\ell=k+1}^{k} \log L(X[\ell]) = 0$. It is easily shown that $W[k]$ can be computed iteratively as follows:

$$W[k] = (W[k-1] + \log L(X[k]))^+, \quad W[0] = 0. \tag{1.6}$$

This recursion is useful from the viewpoint of implementing the CuSum test. Note that the threshold b in (1.4) is chosen to meet a desired false alarm constraint.

1.2.3 Optimization Criteria and Optimality of CuSum

Without any prior information about the change point, a reasonable measure of false alarms is the Mean Time to False Alarm (MTFA):

$$\text{MTFA}(\tau) = \mathbb{E}_\infty[\tau], \tag{1.7}$$

where $\mathbb{E}_\infty$ is the expectation with respect to the probability measure when the change never occurs. Finding a uniformly powerful test that minimizes the delay over all possible values of γ subject to a false alarm constraint, say $\text{MTFA}(\tau) \geq \beta$, is generally not possible. Therefore it is more appropriate to study the quickest change detection problem in what is known as a minimax setting in this case. There are two important minimax problem formulations, one due to Lorden [21] and the other due to Pollak [22].

In Lorden's formulation, the objective is to minimize the supremum of the average delay conditioned on the worst possible realizations, subject to a constraint on the false alarm rate. In particular, we define the Worst-case Average Detection Delay (WADD) as follows:

$$\text{WADD}(\tau) = \sup_{\gamma \geq 1} \text{ess sup } \mathbb{E}_\gamma \left[(\tau - \gamma + 1)^+ | X[1], \ldots, X[\gamma - 1] \right] \tag{1.8}$$

where $\mathbb{E}_\gamma$ denotes the expectation with respect to the probability measure when the change occurs at time γ. We have the following problem formulation.

Problem 1.1 (Lorden) Minimize $\text{WADD}(\tau)$ subject to $\text{MTFA}(\tau) \geq \beta$.

Lorden showed that the CuSum algorithm (1.4) is asymptotically optimal for Problem 1.1 as $\beta \to \infty$. It was later shown in [23] that an algorithm that is equivalent to the CuSum algorithm is actually exactly optimal for Problem 1.1. Although the CuSum algorithm enjoys such a strong optimality property under Lorden's formulation, it can be argued that WADD is a somewhat pessimistic measure of delay. A less pessimistic way to measure the delay, the Conditional Average Detection Delay (CADD), was suggested by Pollak [22]:

$$\text{CADD}(\tau) = \sup_{\gamma \geq 1} \mathbb{E}_\gamma[\tau - \gamma | \tau \geq \gamma] \tag{1.9}$$

for all stopping times τ for which the expectation is well-defined. It can be shown that

$$\text{CADD}(\tau) \leq \text{WADD}(\tau). \tag{1.10}$$

We then have the following problem formulation.

Problem 1.2 (Pollak) Minimize CADD(τ) subject to MTFA(τ) $\geq \beta$.

It can be shown that the CuSum algorithm ((1.4) is asymptotically optimal for Problem 1.2 as $\beta \to \infty$ (see, e.g., [10]).

1.2.4 Incompletely Specified Observation Models

Thus far we have assumed that pre- and post-change distributions, i.e., f_0 and f_1 are completely specified. For the line outage problem that is of interest in this chapter, while it reasonable to assume that the pre-change distribution is completely specified, the post-change distribution can only be assumed to come from a parametric family of distributions, with an unknown parameter a that depends on the line that goes into outage. The observation model of (1.1) becomes:

$$\begin{aligned} &X[k] \sim \text{i.i.d. with pdf } f_0 \text{ for } k < \gamma, \\ &X[k] \sim \text{i.i.d. with pdf } f_a \text{ for } k \geq \gamma, a \in \mathcal{A} \end{aligned} \tag{1.11}$$

where $\mathcal{A}$ is the set of all possible values that the parameter can take.

A powerful approach to constructing a good test for the observation model of (1.11) is the *generalized likelihood ratio* (GLR) approach [24, 25], where at any time step, all the past observations are used to first obtain a maximum likelihood estimate of the post-change parameter a, and then the post-change distribution corresponding to this estimate of a is used to compute the test statistic. In particular, the GLR-CuSum test can be constructed as:

$$\tau_{\mathrm{C}}^{\mathrm{GLR}} = \inf\{k \geq 1 : W_{\mathrm{G}}[k] \geq b\} \tag{1.12}$$

where

$$W_{\mathrm{G}}[k] = \max_{0 \leq j \leq k} \max_{a \in \mathcal{A}} \sum_{\ell=j+1}^{k} \log L_a(X[\ell]) \tag{1.13}$$

with

$$L_a(X[\ell]) = \frac{f_a(X[\ell])}{f_0(X[\ell])}.$$

Note that for general $\mathcal{A}$, the test statistic $W_{\mathrm{G}}[k]$ cannot be computed recursively as in (1.6). However, when the cardinality of $\mathcal{A}$ is finite, which will be the case for the line outage problem, we can implement $\tau_{\mathrm{C}}^{\mathrm{GLR}}$ via a recursive test statistic as follows. First, we swap the maxima in (1.13) to obtain:

$$W_{\mathrm{G}}[k] = \max_{a \in \mathcal{A}} \max_{0 \leq j \leq k} \sum_{\ell=j+1}^{k} \log L_a(X[\ell]) \tag{1.14}$$

Then if we define

$$W_a[k] = \max_{0 \leq j \leq k} \sum_{\ell=j+1}^{k} \log L_a(X[\ell]),$$

we can compute $W_a[k]$ recursively as:

$$W_a[k] = (W_a[k-1] + \log L_a(X[k]))^+, \quad W_a[0] = 0$$

and we have

$$\tau_{\mathrm{C}}^{\mathrm{GLR}} = \inf \left\{ k \geq 1 : \max_{a \in \mathcal{A}} W_a[k] \geq b \right\} \tag{1.15}$$

where b is chosen to meet a desired false alarm constraint. The GLR-CuSum test can be interpreted as running $|\mathcal{A}|$ CuSum tests in parallel, each corresponding to a different value of a, and stopping at the first time any one of the CuSum tests crosses the threshold b.

1.2.5 QCD Under Transient Dynamics

Another variant of the basic QCD problem that is relevant to the line outage detection problem is one where the change from the initial distribution to the final persistent distribution does not happen instantaneously, but after a series of transient phases. The observations within the different phases are generated by different distributions. The resulting observation model is then given by:

$$\begin{aligned} &X[k] \sim \text{i.i.d. with pdf } f_0 \text{ for } k < \gamma_0 \\ &X[k] \sim \text{i.i.d. with pdf } f_i \text{ for } \gamma_{i-1} \leq k < \gamma_i, \quad 1 \leq i \leq T-1 \\ &X[k] \sim \text{i.i.d. with pdf } f_T \text{ for } k \geq \gamma_{T-1} \end{aligned} \tag{1.16}$$

i.e., there are $T-1$ transient phases before the observations finally settle down to the persistent post-change distribution f_T. The change times γ_i, $0 \leq i \leq T-1$ are assumed be unknown, and the goal is to choose a stopping time τ is such a way as to the detect the change at γ_0 as quickly as possible subject to false alarm constraints.

A dynamic generalized likelihood approach can be used to construct a test for this problem [7, 26], and this approach leads to the following Dynamic CuSum (D-CuSum) Test:

$$\tau_{\mathrm{D}} = \inf \{ k \geq 1 : W_{\mathrm{D}}[k] \geq b \}. \tag{1.17}$$

Denoting

$$L_i(X[k]) = \frac{f_i(X[k])}{f_0(X[k])}$$

the statistic $W_{\mathrm{D}}[k]$ can be computed recursively through the following nested iterations. For $i = 1, 2, \ldots, T$, we first compute iteratively:

$$\Omega_i[k] = \max \{ \Omega_1[k-1], \Omega_2[k-1], \ldots, \Omega_i[k-1], 0 \} + \log L_i(X[k])$$

with $\Omega_i[0] = 0$, for $i = 1, 2, \ldots, T$. Then, we compute

$$W_{\mathrm{D}}[k] = \max \{ \Omega_1[k], \Omega_2[k], \ldots, \Omega_T[k], 0 \}. \tag{1.18}$$

It is established in [26] that the D-CuSum algorithm is asymptotically optimum as the mean time to false alarm goes to infinity. We will see in Section 1.4 that the line outage detection problem has both the transient and composite aspects, and therefore a combination of the GLR and dynamic CuSum algorithms will be used in constructing the most effective algorithm for line outage detection.

1.3 Power System Model

We consider a power system with N buses and L transmission lines, respectively indexed by the elements in $\mathcal{V} = \{1, \dots, N\}$ and $\mathcal{L} = \{1, \dots, L\}$. Let $V_i(t)$ and $\theta_i(t)$ respectively denote the voltage magnitude and phase angle of bus i at time t. Similarly, let $P_i(t)$ and $Q_i(t)$ respectively denote the net active and reactive power injection into bus i at time t. Then, the quasi steady-state behavior of the system can be described by the power flow equations (see e.g., [27]), which for bus i can be written as:

$$P_i(t) = \sum_{k=1}^{N} V_i(t)V_k(t)\Big(G_{ik}\cos\big(\theta_i(t) - \theta_k(t)\big) + B_{ik}\sin\big(\theta_i(t) - \theta_k(t)\big)\Big), \quad (1.19)$$

$$Q_i(t) = \sum_{k=1}^{N} V_i(t)V_k(t)\Big(G_{ik}\sin\big(\theta_i(t) - \theta_k(t)\big) - B_{ik}\cos\big(\theta_i(t) - \theta_k(t)\big)\Big) \quad (1.20)$$

where G_{ik} and B_{ik} are, respectively, the real and imaginary part of the (i, k) entry of the network admittance matrix. In the remainder we make the following two assumptions:

A1. Line outages are persistent, i.e., once a line (say line ℓ) is unintentionally tripped offline, its terminals will remain open until it is detected that the line is in such condition.

A2. Any single line outage will not cause the system to break into two electrically isolated islands, i.e., the underlying graph describing the power system network remains connected.

1.3.1 Pre-outage Model

Let Δt [s] denote the PMU sample time and define the k^{th} measurement sample of active and reactive power injections at bus i as

$$P_i[k] := P_i(k\Delta t)$$

and

$$Q_i[k] := Q_i(k\Delta t),$$

respectively. Similarly, define the k^{th} voltage magnitude and phase angle measurement sample at bus i as

$$V_i[k] := V_i(k\Delta t)$$

and

$$\theta_i[k] \coloneqq \theta_i(k\Delta t)$$

respectively. In addition, define the differences in bus i's voltage magnitudes and phase angles of between consecutive sampling times $k\Delta t$ and $(k+1)\Delta t$ as

$$\Delta V_i[k] \coloneqq V_i[k+1] - V_i[k]$$

and

$$\Delta\theta_i[k] \coloneqq \theta_i[k+1] - \theta_i[k]$$

respectively. Similarly, define the differences in the active and reactive power injections into bus i between two consecutive sampling times as

$$\Delta P_i[k] = P_i[k+1] - P_i[k]$$

and

$$\Delta Q_i[k] = Q_i[k+1] - Q_i[k]$$

respectively.

Next, we assume the DC power flow assumptions (see e.g., [27]) hold, namely

D1. flat voltage profile, i.e., $V_i(t) \approx 1$ p.u. for all $i \in \mathcal{V}$;
D2. negligible line resistances, i.e., $G_{ik} \approx 0$ for all $i, k \in \mathcal{V}$; and
D3. small phase angle differences, i.e.,

$$\cos\big(\theta_i(t) - \theta_k(t)\big) \approx 1$$

and

$$\sin\big(\theta_i(t) - \theta_k(t)\big) \approx \theta_i - \theta_k$$

for all i, k for which $B_{ik} > 0$.

Define $\Delta\tilde{P}[k] = \big[\Delta P_1[k], \ldots, \Delta P_N[k]\big]^\top$ and $\Delta\tilde{\theta}[k] = \big[\Delta\theta_1[k], \ldots, \Delta\theta_N[k]\big]^\top$. Then, by linearizing (1.19) around $(\theta_i[k], V_i[k], P_i[k], Q_i[k])$, $i = 1, \ldots, N$, we obtain

$$\Delta\tilde{P}[k] = \tilde{H}_0 \Delta\tilde{\theta}[k] \tag{1.21}$$

where $\tilde{H}_0 = [B_{ik}] \in \mathbb{R}^{N\times N}$. Without loss of generality, assume that bus 1 is the "slack" bus, i.e., it behaves as a voltage source with known magnitude and angle, which we take to be zero. Then, after omitting the equation corresponding to the slack bus, the relationship between variations in voltage phase angles and the variations in the active power injection is given by:

$$\Delta P[k] \approx H_0 \Delta\theta[k] \tag{1.22}$$

where $\Delta P[k]$ and $\Delta\theta[k] \in \mathbb{R}^{(N-1)}$ are respectively obtained by removing the first entry from $\Delta\tilde{P}[k]$ and $\Delta\tilde{\theta}[k]$, and $H_0 \in \mathbb{R}^{(N-1)\times(N-1)}$ is obtained by removing the first row and first column from the matrix $\tilde{H}_0$.

Next, assume that the buses are indexed in such a way that generator buses appear first in $\Delta P[k]$. Then, we can partition $\Delta P[k]$ as follows:

$$\Delta P[k] = \begin{bmatrix} \Delta P^g[k] \\ \Delta P^d[k] \end{bmatrix}$$

where $\Delta P^d[k] \in \mathbb{R}^{N_d}$ and $\Delta P^g[k] \in \mathbb{R}^{N_g}$, $N_d + N_g = N - 1$, respectively denote the changes in the load demand vector and the power generation vector at time instant k. Now, in order to capture the response of the generators in the system to random load demand fluctuations, we assume that

$$\Delta P^g[k] = B(t)\Delta P^d[k] \tag{1.23}$$

where $B(t) = [B_{ij}(t)]$ captures the response of the power injected at generation bus i to a change in demand at load bus j at instant k. The $B_{ij}(t)$'s depend on generator inertia coefficients and other control parameters, including droop coefficients and the automatic generation control (AGC) participation factors. In this work, we approximate $B(t)$ by assuming it can only take values in a finite set $\{B_0, B_1, \ldots, B_T\}$, where each B_i captures the generator response during some period of time $[t_i, t_{i+1})$. Let $B(t) = B_0$ during the pre-outage period and define $M_0 := H_0^{-1}$. Then, we can substitute (1.23) into (1.22) to obtain a pre-outage relation between the changes in the voltage angles and the active power demand at the load buses as follows:

$$\begin{aligned} \Delta\theta[k] &\approx M_0 \Delta P[k] \\ &= M_0 \begin{bmatrix} \Delta P^g[k] \\ \Delta P^d[k] \end{bmatrix} \\ &= [M_0^{(1)} \; M_0^{(2)}] \begin{bmatrix} B_0 \Delta P^d[k] \\ \Delta P^d[k] \end{bmatrix} \\ &= (M_0^{(1)} B_0 + M_0^{(2)}) \Delta P^d[k] \\ &= \tilde{M}_0 \Delta P^d[k] \end{aligned} \tag{1.24}$$

where $\tilde{M}_0 = M_0^{(1)} B_0 + M_0^{(2)}$.

1.3.2 Instantaneous Change During Outage

At the time an outage occurs, $t = t_f$, there is an instantaneous change in the mean of the voltage phase angle measurements that affects only one incremental sample, namely, $\Delta\theta[\gamma_0] = \theta[\gamma_0 + 1] - \theta[\gamma_0]$, with γ_0 such that $\Delta t \gamma_0 \leq t_f < \Delta t(\gamma_0 + 1)$, i.e., $\theta[\gamma_0]$ is obtained immediately prior to the outage, whereas $\theta[\gamma_0 + 1]$ is obtained immediately after the outage. Assume that the outaged line ℓ connects buses m and n. Then, the effect of an outage in line ℓ can be modeled with a power injection of $P_\ell[\gamma_0]$ at bus m and $-P_\ell[\gamma_0]$ at bus n, where $P_\ell[\gamma_0]$ is the pre-outage line flow across line ℓ from m to n. Following a similar approach as the one in [4], the relation between the voltage phase angle incremental change at the instant of outage, $\Delta\theta[\gamma_0]$, and the variations in the active power flow can be expressed as:

$$\Delta\theta[\gamma_0] \approx M_0 \Delta P[\gamma_0] - P_\ell[\gamma_0 + 1] M_0 r_\ell \tag{1.25}$$

where $r_\ell \in \mathbb{R}^{N-1}$ is a vector with the $(m-1)^{\text{th}}$ entry equal to 1, the $(n-1)^{\text{th}}$ entry equal to -1, and all other entries equal to 0. Then, by using (1.23) together with (1.25), we obtain

$$\Delta\theta[\gamma_0] \approx \tilde{M}_0 \Delta P^d[\gamma_0] - P_\ell[\gamma_0 + 1] M_0 r_\ell. \tag{1.26}$$

1.3.3 Post-Outage

Following a line outage, the power system undergoes a transient response governed by B_i, $i = 1, 2, \ldots, T-1$ until a new quasi steady state is reached, in which $B(t)$ settles to B_T. For example, immediately after the outage occurs, the power system is dominated by the inertial response of the generators, which is immediately followed by the governor response, and then followed by the response of the automatic generation control system response. As a result of the line outage, the system topology changes, which manifests itself in the matrix H_0. This change in the matrix H_0 resulting from the outage can be expressed as the sum of the pre-outage matrix and a rank-one perturbation matrix, ΔH_ℓ, i.e., $H_\ell = H_0 + \Delta H_\ell$. Then, by letting $M_\ell := H_\ell^{-1} = [M_\ell^{(1)} \; M_\ell^{(2)}]$, and proceeding in the same manner as the pre-outage model of (1.24), we obtain the post-outage relation between the changes in the voltage angles and the active power demand as:

$$\Delta\theta[k] \approx \tilde{M}_{\ell,i} \Delta P^d[k], \quad \gamma_{i-1} \leq k < \gamma_i \tag{1.27}$$

where $\tilde{M}_{\ell,i} = M_\ell^{(1)} B_i + M_\ell^{(2)}$, $i = 1, 2, \ldots, T$.

1.3.4 Measurement Model

We assume that the voltage phase angles are measured at only a subset of the load buses, and denote this reduced measurement set by $\hat{\theta}[k]$. Assume that we choose $p \leq N_d$ locations to deploy the PMUs. Then, there are $\binom{N_d}{p}$ possible locations to place the PMUs. Here, we assume that the PMU locations are fixed; in general, the problem of optimal PMU placement is NP-hard but one can resort to heuristics to obtain suboptimal placement policies (see, e.g., [6]).

Let

$$\tilde{M} = \begin{cases} \tilde{M}_0, \text{ if } 1 \leq k < \gamma_0, \\ \vdots \\ \tilde{M}_{\ell,T}, \text{ if } k \geq \gamma_T. \end{cases} \tag{1.28}$$

Then, the absence of a PMU at bus i corresponds to removing the i^{th} row of

$\tilde{M}$. Thus, let $\hat{M} \in \mathbb{R}^{p \times N_d}$ be the matrix obtained by removing $N-p-1$ rows from $\tilde{M}$. Therefore, we can relate $\hat{M}$ to $\tilde{M}$ in (1.28) as follows:

$$\hat{M} = C\tilde{M} \tag{1.29}$$

where $C \in \mathbb{R}^{p \times (N-1)}$ is a matrix of 1's and 0's that appropriately selects the rows of $\tilde{M}$. Accordingly, the increments in the phase angle can be expressed as follows:

$$\Delta\hat{\theta}[k] \approx \hat{M}\Delta P^d[k]. \tag{1.30}$$

Since small variations in the active power injections at the load buses, $\Delta P^d[k]$, can be attributed to random fluctuations in electricity consumption, we may model the $\Delta P^d[k]$'s as i.i.d. random vectors. Then, by using the Central Limit Theorem [28], we can argue that each $\Delta P^d[k]$ is a Gaussian vector, i.e., $\Delta P^d[k] \sim \mathcal{N}(0, \Lambda)$, where Λ is a covariance matrix (note that we make the additional assumption that the elements in $\Delta P^d[k]$ are roughly independent). Since $\Delta\hat{\theta}[k]$ depends on $\Delta P^d[k]$ through the linear relationship given in (1.30), we have that:

$$\Delta\hat{\theta}[k] \sim \begin{cases} f_0 := \mathcal{N}(0, \hat{M}_0\Lambda\hat{M}_0^\top), \text{ if } 1 \le k < \gamma_0, \\ f_\ell^{(0)} := \mathcal{N}(\mu_\ell, \hat{M}_0\Lambda\hat{M}_0^\top), \text{ if } k = \gamma_0, \\ f_\ell^{(1)} := \mathcal{N}(0, \hat{M}_{\ell,1}\Lambda\hat{M}_{\ell,1}^\top), \text{ if } \gamma_1 \le k < \gamma_2, \\ \vdots \\ f_\ell^{(i)} := \mathcal{N}(0, \hat{M}_{\ell,i}\Lambda\hat{M}_{\ell,i}^\top), \text{ if } \gamma_i \le k < \gamma_{i+1}, \\ \vdots \\ f_\ell^{(T)} := \mathcal{N}(0, \hat{M}_{\ell,T}\Lambda\hat{M}_{\ell,T}^\top), \text{ if } \gamma_T \le k \end{cases} \tag{1.31}$$

where $\mu_\ell := -P_\ell[\gamma+1]CM_0r_\ell$ is the instantaneous meanshift and $\gamma_1 = \gamma_0 + 1$. It is important to note that for $\mathcal{N}\left(0, \hat{M}\Lambda\hat{M}^\top\right)$ to be a nondegenerate pdf, its covariance matrix, $\hat{M}\Lambda\hat{M}^\top$, must be full rank. We enforce this by ensuring that the number of PMUs placed, p, is less than or equal to the number of load buses, N_d, and that they are deployed at nodes such that the measured voltage phase angles are independent. The matrices $\hat{M}$ are known based on the system topology following a line outage and Λ can be estimated from historical data.

1.4 Line Outage Detection Using QCD

In the line outage detection problem setting, the goal is to detect the outage in line ℓ as quickly as possible subject to false alarm constraints. The outage induces a change in the statistical characteristics of the observed sequence $\{\Delta\hat{\theta}[k]\}_{k\ge1}$. The aim is to design stopping rules that detect this change using the QCD methodology discussed in Section 1.2.

We begin by introducing the Meanshift test, a detection scheme that can be shown to be equivalent to that proposed in [14] (see also [16] and [15]) for detecting line outages.

1.4.1 Meanshift Test

The Meanshift test is a "oneshot" detection scheme like the Shewhart test of (1.3), i.e., only the most recent observation is used to calculate the test statistics. An additional simplification is that only a single log-likelihood ratio between the distribution of the observations at the changepoint and before the changepoint is used to construct the test. As seen in (1.31), there is a mean shift at exactly the time of outage γ_0, and only this mean shift is exploited in constructing the test statistic:

$$W_\ell^{\mathrm{M}}[k] = \log \frac{f_\ell^{(0)}(\Delta\hat{\theta}[k])}{f_0(\Delta\hat{\theta}[k])}. \tag{1.32}$$

Since an outage can occur at any line, we use a generalized test structure as in the GLR-CuSum test described in Section 1.2.4 , i.e., we take the maximum $W_\ell^{\mathrm{M}}[k]$ over ℓ to form the test statistic. Consequently, the stopping time for the Meanshift test is given by

$$\tau_{\mathrm{M}} = \inf \left\{ k \geq 1 : \max_{\ell \in \mathcal{L}} W_\ell^{\mathrm{M}}[k] > b \right\} \tag{1.33}$$

with $b > 0$ chosen to meet the false alarm constraint.

In addition to the fact that it is a "one-shot" detection scheme, the Meanshift test has the drawback that the details of the post-change statistics, the transient dynamics and the persistent distribution in (1.31), are not incorporated in the test statistic. As a result, the log-likelihood ratio used in the test statistic does not match the true distribution of the observations after the changepoint γ_0. For example, during the first transient period ($\gamma_0 < k \leq \gamma_1$), the expected value of the test statistic could be negative since

$$\mathbb{E}_\ell^{(1)} \left[\log \frac{f_\ell^{(0)}(\Delta\hat{\theta}[k])}{f_0(\Delta\hat{\theta}[k])} \right] = D(f_\ell^{(1)} \| f_0) - D(f_\ell^{(1)} \| \, f_\ell^{(0)}) \tag{1.34}$$

where $\mathbb{E}_\ell^{(1)}$ denotes the expectation under distribution $f_\ell^{(1)}$.

1.4.2 Generalized Shewhart Test

We can improve on the Meanshift test by incorporating all of the post-change statistics in (1.31) into a Shewhart type "one-shot" statistic. In particular, the test statistic corresponding to outage in line ℓ is given by:

$$W_\ell^{\mathrm{S}}[k] = \max_{i \in \{0,1,\ldots,T\}} \left\{ \log \frac{f_\ell^{(i)}(\Delta\hat{\theta}[k])}{f_0(\Delta\hat{\theta}[k])} \right\}. \tag{1.35}$$

and the corresponding generalized Shewhart test is given by:

$$\tau_{\mathrm{S}} = \inf\left\{k \geq 1 : \max_{\ell \in \mathcal{L}} W_\ell^{\mathrm{S}}[k] > b\right\} \tag{1.36}$$

with $b > 0$ chosen to meet the false alarm constraint. While we can expect the generalized Shewhart test to perform better than the Meanshift test, it is still a "one-shot" test that does not use the past information in the observations.

1.4.3 Generalized CuSum Test

We now explore the incorporation of past information in constructing the test statistic as in the CuSum test described in Section 1.2.2. We first make the approximation that the transition between pre- and post-outage periods is not characterized by any transient behavior other than the mean shift that occurs at the instant of outage. In particular, if line ℓ goes into outage, we make the approximation that the distribution of the observations goes from f_0 to $f_\ell^{(0)}$ at the time of outage, and then directly to $f_\ell^{(T)}$, the persistent change distribution for line ℓ. Then, we construct a slight generalization of the CuSum test, in which the test statistic corresponding to line ℓ being in outage incorporates the mean shift term in the following manner:

$$W_\ell^{\mathrm{C}}[k] = \max\left\{W_\ell^{\mathrm{C}}[k-1] + \log\frac{f_\ell^{(T)}(\Delta\hat{\theta}[k])}{f_0(\Delta\hat{\theta}[k])}, \log\frac{f_\ell^{(0)}(\Delta\hat{\theta}[k])}{f_0(\Delta\hat{\theta}[k])}, 0\right\} \tag{1.37}$$

with $W_\ell^{\mathrm{C}}[0] = 0$ for all $\ell \in \mathcal{L}$.

Although the above statistic does not take any transient dynamics into consideration, it is a good approximation when: (i) the transient distributions and the final post-change distribution are "similar", e.g., when the KL divergence between $f_\ell^{(i)}$ and $f_\ell^{(T)}$ is small, for $i = 1, 2, \ldots, T-1$, or (ii) when the expected detection delay is large so that the algorithm stops well into the persistent phase of the change.

Since the line that is in outage is not known *a priori*, we again use the GLR approach of Section 1.2.4 to construct the following generalized test, which we refer to as the G-CuSum test:

$$\tau_{\mathrm{C}} = \inf\left\{k \geq 1 : \max_{\ell \in \mathcal{L}} W_\ell^{\mathrm{C}}[k] > b\right\} \tag{1.38}$$

with $b > 0$ chosen to meet the false alarm constraint.

1.4.4 Generalized Dynamic CuSum Test

We now construct a test that takes into account all of the details of the statistical model of (1.31), the mean shift at the changepoint, the transient dynamics after the changepoint, and the final persistent change. We use the approach described in Section 1.2.5 to form the following test statistic for line ℓ:

$$W_\ell^{\mathrm{D}}[k] = \max\left\{\Omega_\ell^{(0)}[k], \ldots, \Omega_\ell^{(T)}[k], 0\right\} \tag{1.39}$$

where

$$\Omega_\ell^{(i)}[k] = \max\{\Omega_\ell^{(i)}[k-1], \Omega_\ell^{(i-1)}[k-1]\} + \log \frac{f_\ell^{(i)}(\Delta\hat{\theta}[k])}{f_0(\Delta\hat{\theta}[k])} \tag{1.40}$$

for $i \in \{1, \ldots, T\}$,

$$\Omega_\ell^{(0)}[k] := \log \frac{f_\ell^{(0)}(\Delta\hat{\theta}[k])}{f_0(\Delta\hat{\theta}[k])}$$

and $\Omega_\ell^{(i)}[0] := 0$, for all $\ell \in \mathcal{L}$ and all i. We refer to the corresponding generalized test as the Generalized Dynamic CuSum (G-D-CuSum) test, which is given by

$$\tau_{\mathrm{D}} = \inf\left\{k \geq 1 : \max_{\ell \in \mathcal{L}} W_\ell^{\mathrm{D}}[k] > b\right\} \tag{1.41}$$

with $b > 0$ chosen to meet the false alarm constraint.

1.5 Line Outage Identification

The line outage detection algorithms described in Section 1.4 can also be used to identify the line that has been affected by the outage. A natural strategy that exploits the generalized likelihood nature of the test statistics would be to declare the affected line as the one corresponding to the largest statistic, i.e.,

$$\hat{\ell} = \arg\max_{j \in \mathcal{L}} W_j[\tau] \tag{1.42}$$

where W_ℓ denotes the test statistic corresponding to line ℓ for the given test (Meanshift, Shewhart, G-CuSum, G-D-CuSum). However, algorithms based on maximum of line statistics that perform well for line outage detection may not perform well at line outage isolation. The reason is that the statistics for other lines may also increase following a line outage, even if the statistic corresponding to the line in outage has the largest drift. In particular, the drift (average increase) in the test statistic for line ℓ when line j is in outage is given by:

$$\mathbb{E}_{f_j}\left[\log \frac{f_\ell(\Delta\hat{\theta}[k])}{f_0(\Delta\hat{\theta}[k])}\right] = D(f_j \| f_0) - D(f_j \| f_\ell). \tag{1.43}$$

If $D(f_j \| f_\ell)$ is small compared to $D(f_j \| f_0)$, then the test statistic for line ℓ will have a positive drift when line j is in outage, and it is possible that the statistic for line ℓ is larger than that of line j when the threshold is crossed; see Fig. 1.2(b) for an example where line 36 is in outage, but the statistic for line 37 crosses the threshold first.

We could adopt a joint approach for detection and isolation, where we stop only when we are confident that the false isolation probability is small, as in the matrix CuSum test [9]. However, if the KL divergence between the post change distributions for some pairs of lines are small as in the example discussed in the previous paragraph, then good isolation will come at the cost of unacceptably long detection delays.

Another approach to obtaining good isolation error probabilities is to allow for the possibility that when an outage is declared by the detection algorithm, more than one line can be checked by the system operator. Then, we can create a pre-computed ranked list of candidate lines that should be checked when an outage in a particular line is declared by the algorithm. This is similar to the notion of list decoding in digital communications (see e.g. [29]). Then, false isolation occurs only if list of lines checked does not contain true outaged line.

We can create the ranked list based on the drifts of the line statistics given in (1.43). In particular, for the line ℓ statistic, we create the ranked list based on the smallest values of $D(f_j \| f_\ell)$, since the j's correspond to the lines that are most likely to the true outaged line if the line ℓ statistic crosses the threshold when the change is declared. If we constrain the ranked list to have r elements, then we cab define the ranked list corresponding to line ℓ as:

$$\mathcal{R}_\ell = \{\ell, j_2, \ldots, j_r\}. \tag{1.44}$$

Note that line ℓ has to be in $\mathcal{R}_\ell$ since $D(f_\ell \| f_\ell) = 0$. The remaining elements in $\mathcal{R}_\ell$, $j_2, \ldots, j_r$ are the indices of the lines with $D(f_j \| f_\ell)$ values that are successively larger.

To quantify the performance of our algorithm with respect to its ability to identify the outaged line accurately, we define the probability of false isolation (PFI). When line ℓ is in outage, a false isolation event occurs when line ℓ is not included in the ranked list $\mathcal{R}_{\hat{\ell}}$, with $\hat{\ell}$ chosen according to (1.42). Therefore we can define the PFI when line ℓ is outaged as:

$$\mathrm{PFI}_\ell(\tau) = \mathbb{P}\{\ell \notin \mathcal{R}_{\hat{\ell}} | \text{line } \ell \text{ in outage}\}. \tag{1.45}$$

The size r of the ranked list should be chosen to optimize the tradeoff between PFI and number of lines that need to be checked after an outage detection has occurred. In particular, larger ranked lists lead to lower PFI, but to a larger set of possibly outaged lines to check.

1.6 Numerical Results

In this section, we verify the effectiveness of the algorithm in (1.39)-(1.41) for detecting line outages in the IEEE 118-bus test system (see [30] for the model data). In order to simulate the dynamics of the system following a line outage, we utilize the Power System Toolbox (PST) [31], which includes detailed models of the generator dynamics and control. For simplicity, in (1.31) we only consider one transient period after the line outage occurrence, i.e., $T = 2$ with a duration

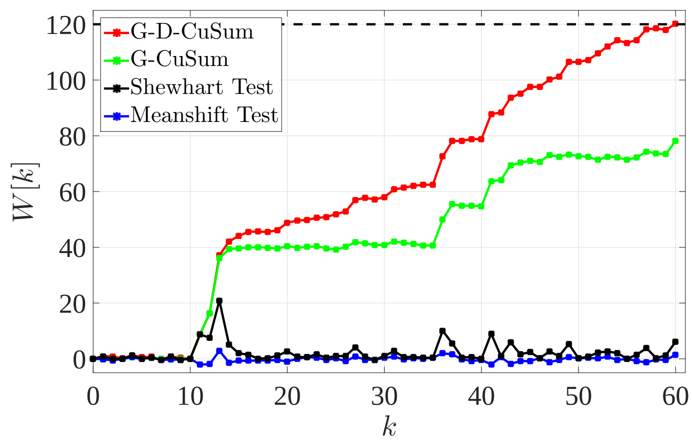

(a) Line 180 outage with b=120.

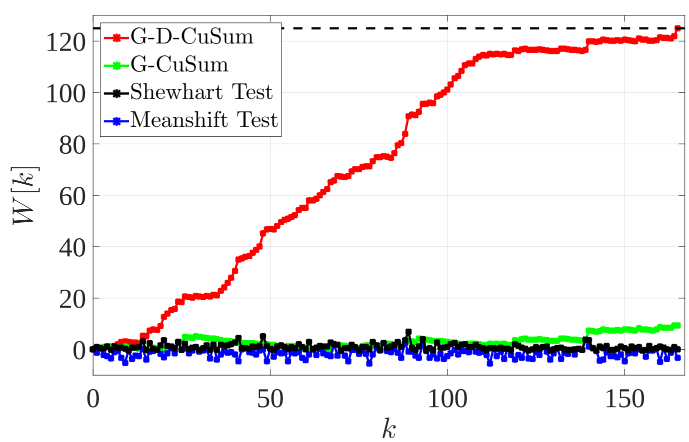

(b) Line 32 outage with b=125.

Figure 1.1 Sample paths of different algorithms for IEEE 118-bus test system.

of 100 samples. As discussed earlier, we assume the changes in load demands to be zero-mean independent Gaussian random variables with variance 0.03. We also assume that the PMUs provide 30 measurements per second and choose to place them at all load buses.

1.6.1 Line Statistic Evolution

We simulate two different line outages occurring at $k = 10$; the results are shown in Fig. 1.1. In one case, the detection takes place during the transient period, whereas in the other case the detection takes places after the transient period has ended. Figure 1.1(a) shows some typical progressions of $W_{180}[k]$ for the various line outage detection schemes discussed earlier with a detection threshold of

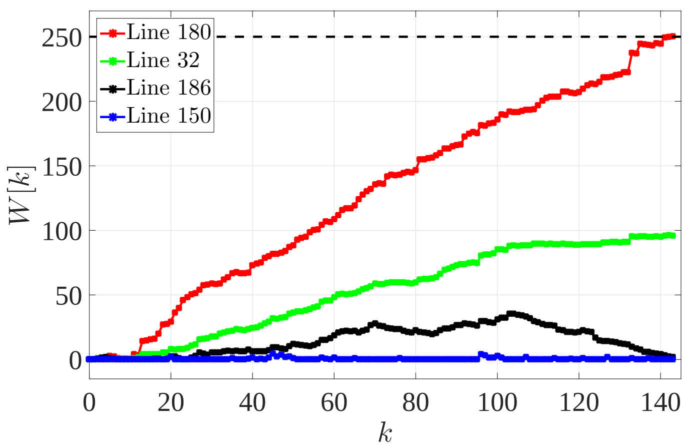

(a) Line 180 outage with b=250.

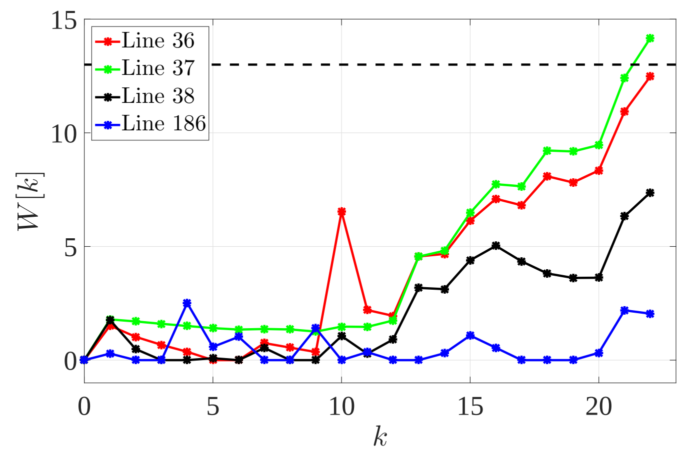

(b) Line 36 outage with b=13.

Figure 1.2 Sample paths of the G-D-CuSum algorithm for IEEE 118-bus test system.

$b = 120$. As shown, a line outage can be declared after 50 samples when the G-D-CuSum stream crosses the threshold. Also as shown in the figure, the other algorithms incur a much larger detection delay. The progress of $W_{32}[k]$ for the different algorithms for an outage in line 32 is shown in Fig. 1.1(b), where one can see that for the detection threshold, $b = 125$, the G-D-CuSum detects an outage 156 samples after its occurrence In this case, the detection occurs after the transient dynamics have ended. From the plots, we can conclude that the G-D-CuSum algorithm has a smaller detection delay than the G-CuSum algorithm even though the line outage is detected after the transient period has finished at approximately $k = 110$. Also note that the G-CuSum algorithm has a large delay because its statistic does not grow during the transient period.

Figure 1.2 shows the evolution of the G-D-CuSum statistic for different lines and two different outages in lines 180 and 36 occurring at time $k = 10$. For an outage in line 180, as shown in Fig. 1.2(a), $W_{180}[k]$ grows faster than that of other line statistics, and an an outage is declared after 135 samples, when $W_{180}[k]$ crosses the detection threshold, $b = 250$. Next, we illustrate the occurrence of a false isolation event. Specifically, we simulate an outage in line 36; however, in Fig. 1.2(b) we see that $W_{37}[k]$ crosses the detection threshold, $b = 13$, before $W_{36}[k]$ does.

1.6.2 Delay Performance

We performed Monte Carlo simulations for outages in lines 36, 180, and 104, and computed detection delay versus MTFA results for the G-D-CuSum, the Mean-shift test, the Shewhart test, and the G-CuSum algorithm; their performance for an outage in line 36—measured as the worst delay of all three outaged cases—is compared in Fig. 1.3. From the results in Fig. 1.3(a), we can conclude that the G-D-CuSum algorithm achieves the lowest detection delay among all algorithms for a given MTFA. In Fig. 1.3(b) we show the performance gain that is achieved when choosing the G-D-CuSum test over the Meanshift and Shewhart tests, and the G-CuSum test. As shown, for an outage in line 36 and a given false alarm rate, the delay achieved with the G-D-CuSum tests is an order of magnitude less than that of other tests.

1.6.3 Probability of False Isolation

Tables 1.1-1.3 contains the PFI versus MTFA for outages in lines 36, 104, and 180. We computed the PFI using the ranked list method discussed in Section 1.5 for a ranked list of fixed length 1, 3, and 5. Table 1.1 contains the PFI versus MTFA results for a ranked lists of length 1; this is equivalent to identifying the outaged line using (1.42), i.e., identifying the line with the highest statistic at the stopping time as outaged. Here it is important to note that while this technique might be efficient to handle some line outages (e.g., outage in line 180 and 104), some other line outages, e.g., an outage in line 36, may lead to large PFI values. As discussed in Section 1.5, this is due to the fact that many line statistics other than the one corresponding to the outaged line grow post-outage. Table 1.2 contains the PFI versus MTFA values for a ranked list of length 3; one can see that in this case the PFI is significantly reduced for line 36. Finally, Table 1.3 contains the PFI versus MTFA for a ranked list of length 5, where one case that in this case the PFI for line 36 is below 5%. Finally, note that the PFI decreases as the MTFA increases; this is due to the fact that larger MTFAs corresponds to larger thresholds, which result in smaller PFI values.

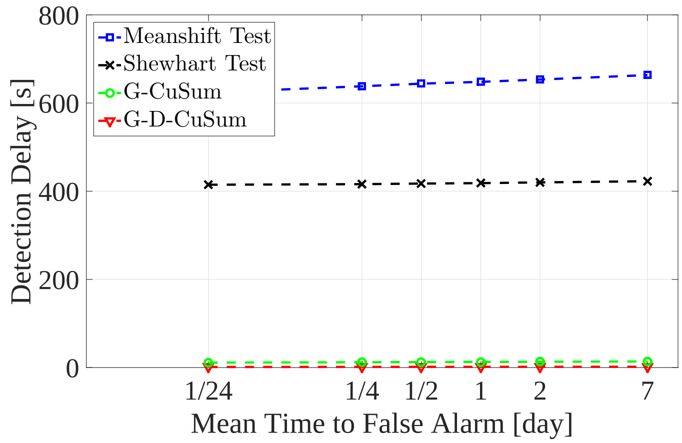

(a) Detection delay vs. mean time to false alarm for different algorithms.

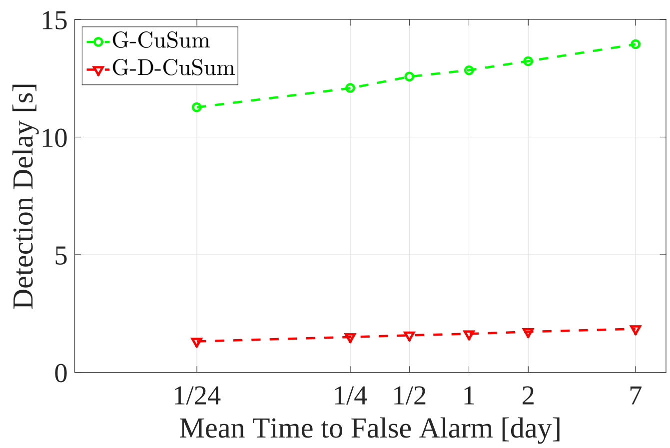

(b) Detection delay vs. mean time to false alarm for G-D-CuSum and G-CuSum algorithms.

Figure 1.3 Monte Carlo simulation results for an outage in line 36 for IEEE 118-bus test system.

1.7 Conclusion

In this chapter, we discussed algorithms based on the methodology of statistical quickest change detection for detecting and identifying line outages. The algorithms exploit the statistical properties of voltage phase angle measurements obtained from PMUs in real-time and feature a set of statistics that are used to capture each distribution shift that results from a line outage.

The most effective of these algorithms, the G-D-CuSum algorithm, takes a generalized likelihood ratio approach and incorporates the transients dynamics after the outage in the test statistic. We showed that the algorithm outperforms other

Table 1.1 Probability of false isolation for IEEE 118-bus test system simulated with a ranked list of length of 1.

$\mathbb{E}_\infty[\tau]$ [day]	1/24	1/4	1/2	1	2	7
Line 36	0.5999	0.5472	0.5266	0.5182	0.5102	0.4911
Line 180	0.0344	0.0237	0.0186	0.0160	0.0148	0.0091
Line 104	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$

Table 1.2 Probability of false isolation for IEEE 118-bus test system simulated with a ranked list of length of 3.

$\mathbb{E}_\infty[\tau]$ [day]	1/24	1/4	1/2	1	2	7
Line 36	0.1477	0.1392	0.1371	0.1303	0.1251	0.1219
Line 180	0.0197	0.0127	0.0112	0.0097	0.0062	0.0049
Line 104	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$

Table 1.3 Probability of false isolation for IEEE 118-bus test system simulated with a ranked list of length of 5.

$\mathbb{E}_\infty[\tau]$ [day]	1/24	1/4	1/2	1	2	7
Line 36	0.0295	0.0203	0.0158	0.0137	0.0108	0.0067
Line 180	0.0097	0.0074	0.0052	0.0041	0.0036	0.0027
Line 104	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$	$< 10^{-6}$

previously developed line outage detection methods. In particular, we compared the detection delay performance of this QCD-based algorithm against other line outage detection algorithms for line outages simulated on the IEEE 118-bus test system. We used this system to also simulate the delay performance of the G-D-CuSum for different line outage scenarios, as well as to illustrate the use of the algorithm for line outage identification. To achieve effective identification, we used the mechanism of ranked lists, i.e., sets of possibly outaged lines that are highly likely to contain the true outaged line.

The G-D-CuSum algorithm can be further improved by incorporating prior knowledge about the durations of the transients that follow the line outage, but our simulation studies for the IEEE 118-bus test system indicate that the improvement may be marginal. Another topic that is of interest to explore is the optimal placement of limited PMU resources in the network to maximize line outage detection performance. Some initial studies along these lines are presented in [6].

The techniques discussed in this chapter can be easily extended to double-line outage detection (and in general to multiple line outage detection), as in [4], by characterizing the distribution of the observations after a double-line outage. Then, the G-D-CuSum scheme can be used to detect such a change in the distribution. In particular, to handle double-line outages, a test statistic needs to be calculated for every possible pair of lines; however, this might be impractical in large systems due to combinatorial explosion. An alternative is just to just focus on the double-line outages that are more likely to occur. These can be characterized from historical data or from simulations by noting that a double-line outage is usually the result of a line tripping offline causing a second line to trip offline shortly after due to the particular loading conditions at the time the first line trips. Then, by taking into account the probability of each line tripping offline in the first place (which can also be determined from historical data for a particular system) and the loading conditions probability distribution, one can conduct Monte Carlo simulations to determine which lines are more likely to trip offline immediately after the occurrence of the first outage. As an alternative to Monte Carlo simulation, one can potentially use so-called line outage distribution factors (LODFs)—linear sensitivities that determine how the power flow on a particular line changes after the occurrence of an outage in another line—to determine which lines may become overloaded, and thus trip offline, following an outage in some particular line. This LODF-based method is less accurate that Monte Carlo simulation but also less computationally expensive.

References

[1] FERC and NERC. (2012, Apr.) Arizona-southern California outages on September 8, 2011: Causes and recommendations. [Online]. Available: http://www.ferc.gov.

[2] U.S.-Canada Power System Outage Task Force. (2004, Apr.) Final report on the August 14th blackout in the United States and Canada: causes and recommendations. [Online]. Available: http://energy.gov.

[3] T. Banerjee, Y. Chen, A. Domínguez-García, and V. Veeravalli, "Power system line outage detection and identification; a quickest change detection approach," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, May 2014, pp. 3450–3454.

[4] Y. C. Chen, T. Banerjee, A. D. Domínguez-García, and V. V. Veeravalli, "Quickest line outage detection and identification," *IEEE Trans. Power Syst.*, vol. 31, no. 1, pp. 749–758, Jan. 2016.

[5] G. Rovatsos, X. Jiang, A. D. Domínguez-García, and V. V. Veeravalli, "Comparison of statistical algorithms for power system line outage detection," in *Proc. of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Apr. 2016.

[6] X. Jiang, Y. C. Chen, V. V. Veeravalli, and A. D. Domínguez-García, "Quickest line outage detection and identification: Measurement placement and system partitioning," in *2017 North American Power Symposium (NAPS)*, Sep. 2017, pp. 1–6.

[7] G. Rovatsos, X. Jiang, A. D. Domínguez-García, and V. V. Veeravalli, "Statistical power system line outage detection under transient dynamics," *IEEE Transactions on Signal Processing*, vol. 65, no. 11, pp. 2787–2797, June 2017.

[8] H. V. Poor and O. Hadjiliadis, *Quickest Detection.* Cambridge University Press, 2009.

[9] A. G. Tartakovsky, I. V. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Change-Point Detection*, ser. Statistics. CRC Press, 2014.

[10] V. V. Veeravalli and T. Banerjee, *Quickest Change Detection.* Elsevier: E-reference Signal Processing, 2013.

[11] K. A. Clements and P. W. Davis, "Detection and identification of topology errors in electric power systems," *IEEE Trans. Power Syst.*, vol. 3, no. 4, pp. 1748–1753, Nov. 1988.

[12] F. F. Wu and W. E. Liu, "Detection of topology errors by state estimation," *IEEE Trans. Power Syst.*, vol. 4, no. 1, pp. 176–183, Feb. 1989.

[13] N. Singh and H. Glavitsch, "Detection and identification of topological errors in online power system analysis," *IEEE Trans. Power Syst.*, vol. 6, no. 1, pp. 324–331, Feb. 1991.

[14] J. E. Tate and T. J. Overbye, "Line outage detection using phasor angle measurements," *IEEE Trans. Power Syst.*, vol. 23, no. 4, pp. 1644–1652, Nov. 2008.

[15] H. Zhu and G. B. Giannakis, "Sparse overcomplete representations for efficient identification of power line outages," *IEEE Trans. Power Syst.*, vol. 27, no. 4, pp. 2215–2224, Nov. 2012.

[16] G. Feng and A. Abur, "Identification of faults using sparse optimization," in *Proc. of Communication, Control, and Computing (Allerton Conference)*, Sept 2014, pp. 1040–1045.

[17] M. Garcia, T. Catanach, S. V. Wiel, R. Bent, and E. Lawrence, "Line outage localization using phasor measurement data in transient state," *IEEE Trans. Power Syst.*, vol. 31, no. 4, pp. 3019–3027, Jul. 2016.

[18] P. Moulin and V. Veeravalli, *Statistical Inference for Engineers and Data Scientists*. Cambridge University Press, 2019.

[19] W. A. Shewhart, "The application of statistics as an aid in maintaining quality of a manufactured product," *J. Amer. Statist. Assoc.*, vol. 20, no. 152, pp. 546–548, Dec. 1925.

[20] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, no. 1/2, pp. 100–115, Jun. 1954.

[21] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Statist.*, vol. 42, no. 6, pp. 1897–1908, 12 1971. [Online]. Available: http://dx.doi.org/10.1214/aoms/1177693055

[22] M. Pollak, "Optimal detection of a change in distribution," *Ann. Statist.*, vol. 13, no. 1, pp. 206–227, Mar. 1985.

[23] G. V. Moustakides, "Optimal stopping times for detecting changes in distributions," *Ann. Statist.*, vol. 14, no. 4, pp. 1379–1387, Dec. 1986. [Online]. Available: http://dx.doi.org/10.1214/aos/1176350164

[24] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Trans. Inf. Theory*, vol. 44, no. 7, pp. 2917–2929, Nov. 1998.

[25] D. Siegmund and E. S. Venkatraman, "Using the generalized likelihood ratio statistic for sequential detection of a change-point," *Ann. Statist.*, vol. 23, no. 1, pp. 255–271, Feb. 1995.

[26] S. Zou, G. Fellouris, and V. Veeravalli, "Quickest change detection under transient dynamics: Theory and asymptotic analysis," *IEEE Transactions on Information Theory*, vol. 65, no. 3, March 2019.

[27] A. R. Bergen and V. Vittal, *Power Systems Analysis*. Prentice Hall, 2000.

[28] B. Hajek, *Random Processes for Engineers*. Cambridge university press, 2015.

[29] P. Elias, "Error-correcting codes for list decoding," *IEEE Trans. Inf. Theory*, vol. 37, no. 1, pp. 5–12, Jan. 1991.

[30] "Power system test case archive," Oct. 2012. [Online]. Available: http://www2.ee.washington.edu/research/pstca.

[31] J. Chow and K. Cheung, "A toolbox for power system dynamics and control engineering education and research," *IEEE Trans. Power Syst*, vol. 7, no. 4, pp. 1559–1564, Nov. 1992.

Index