21, rue d'Artois, F-75008 PARIS
http : //www.cigre.org

CIGRE US National Committee
2014 Grid of the Future Symposium

Implications of Cyber Attacks on Distributed Power System Operations

J. ZHANG*, L. SANKAR, K. HEDMAN
Arizona State University
USA

SUMMARY

The electric grid is a complex physically distributed and inter-connected network managed by a large number of entities (e.g., systems operators, utilities) to ensure reliable transmission, generation, and distribution of power. Sustained and reliable operation with dynamic situational awareness in the grid requires continued data sharing amongst the grid entities. Lack of automated communications and coordination between distributed operators in the grid contributes significantly to the lack of global situational awareness occasionally with serious consequences of runaway cascading failures. While wide-area monitoring and information sharing has been proposed by the Federal Energy Regulatory Commission, real-time data sharing in the grid is still done in an ad hoc manner between connected areas. Furthermore, the mode, amount, and granularity of data shared are not standardized. A smart adversary can design man-in-the middle attacks that limit the information shared between adjacent areas, thereby threatening the reliable operation of the system.

In this paper, we focus on a class of topology-targeted man-in-the-middle (MitM) communication attacks aimed at limiting information sharing between adjacent areas, particularly when one or both areas experience topology changes (e.g., line outages). To understand the broader consequences of such attacks on actual power systems operation, we develop a tractable temporal model for energy management system (EMS) operations that allows studying the time progression of the cyber-attack introduced in one area and its effect on both areas. The aim of this work is two-fold: (i) understanding the physical consequences of a class of cyber-attacks; and (ii) mimicking data sharing conditions that in practice led to blackouts when local outages were not shared in real-time between connected areas (e.g., Northeast black 2003).

Our results demonstrate that such an MitM communication attack in a distributed power network leads to a range of consequences, some more severe than others: these include relatively benign oscillations in the power flow solutions between the two areas that eventually fix themselves (infrequent) to more complex situations (more likely) over time including power flow overload violations caused by thermal limit relaxations, progressively severe lack of convergence of OPF in both areas, as well as actual physical line overflows that are not observable from the cyber solution but can eventually cause line overheating and cascading outages. Based on these observations, in addition to the traditional countermeasure of human operator-based data sharing (which have been shown to be error-prone and delayed too), it is essential to have more resiliency via automated data sharing mechanisms. To this end, we propose an interactive distributed data processing platform. This could help both areas become aware of inconsistencies over faster time-scales including: (a) enable local topology processing to include interactive updating; (b) enable real-time coordination of dispatch between the two areas; and (c) create and share a list of external contingencies caused to other areas by an internal component outage.

KEYWORDS

Cyber-attack, topology, information sharing, distributed power system operation.

jzhan188@asu.edu

I. Introduction

The electric grid is a complex physically distributed and inter-connected network managed by a large number of entities (e.g., systems operators, utilities) to ensure reliable transmission, generation, and distribution of power. Sustained and reliable operation with dynamic situational awareness in the grid requires continued data sharing amongst the grid entities. The grid is fast converging towards a Smart Grid characterized by (a) vastly expanded data acquisition, (b) highly variable environments due to integration of renewables, and (c) distributed processing and control. In this new paradigm, timely and controlled information exchange is critical not only to ensure reliability and stability but also to thwart cyber attacks that could potentially bring down the entire grid with one or more local outages.

In this paper, we focus on a class of topology-targeted man-in-the-middle (MitM) communication attacks aimed at limiting information sharing between adjacent areas, particularly when one or both areas experience topology changes (e.g., line outages). While wide-area monitoring and information sharing has been proposed by the Federal Energy Regulatory Commission based on observations that lack of seamless data sharing is an important factor in cascading failures, real-time data sharing in the grid is still done in an ad hoc manner between connected areas. For example, in the Northeast blackout of 2003 [1], [2], a line out in one area (Ohio) was not conveyed for a sufficient period of time to neighboring regions leading to convergence failure of the state estimator and other cascading problems. Furthermore, the mode, amount, and granularity of data shared is not standardized; for example, two connected areas may only share limited topology information such as low granularity network equivalent models which in turn are insufficient to capture the complexity of the electric grid and ensure wide-area reliability (e.g., the Yuma-Southern California outage of 2011 [3]). In fact, changes in the grid topology are often communicated via human operators and not in an automated manner which adds to communication delays and errors. In the light of such limitations, a smart adversary can limit information sharing in a number of ways. We seek to understand the effects of such limited data sharing scenarios (both adversarial and otherwise) on the electric power system real-time operations.

We introduce a class of distributed communication attacks wherein an attack on the Energy Management System (EMS) of one area prevents the sharing of topology changing information with the other area (in automated systems where topology may be shared real-time or frequently, this can be achieved via man-in-the-middle attacks). We assume that the attacker is either involved in bringing down a line remotely (breakers can be remotely tripped in some cases) or is aware of a line out (again possible via presence of software trojans in the EMS). The attacker, therefore, is assumed to have some knowledge of the network topology.

There has been much recent interest on cyber attacks on the grid, in particular false data injection (or integrity) attacks, where as the name suggest false data is introduced in specific measurement and computing units of the EMS such as state estimation (e.g., [4], [5]), automatic generation control (e.g., [6]), generator frequency control (e.g., [7]), topology processing (e.g., [8]), as well as attack consequences on markets (e.g., [9], [10]). However, the consequences of such cyber-attacks on system operations are yet to be demonstrated. An important question that remains to be addressed is whether serious damage such as instability, cascading failures, and potential blackouts, which can cripple society and the economy, can be caused by cyber attacks on the grid.

To understand the broader consequences of (unobservable) attacks on measurements or shared data, we develop a layered systems model that enables the modeling of the time progression of attacks. In [11], Liang et al. introduced a time progression based system model for EMS functionalities and used it to demonstrate how an unobservable false data injection attack on AC state estimation (SE), by a sophisticated attacker, can lead to a physical generation dispatch when none was needed. In this paper, we focus on a distributed two-area (managed by two operators and EMSs) setting to demonstrate the consequences of limited information sharing. Specifically, we focus on attacks that create or exploit outages in one area and limit information sharing via a communication attack thereby affecting the optimal power flow solutions and dispatch in a connected area that has incorrect topology information. Our results demonstrate that such an attack in a distributed power network leads to a range of possibilities; these include relatively benign oscillations in the power flow solutions between the two areas that eventually fix themselves (infrequent) to more complex situations (more likely) over time including power flow overload violation caused by thermal limit relaxation, progressively severe lack of convergence of OPF in both areas, as well as actual physical line overflows that are not observable from the cyber solution but can eventually cause line overheating and cascading outages. Our time-progression based system model allows us to capture the major computational components of EMSs including AC state estimation and optimal power flow (OPF) including generation dispatch and understand the temporal consequences of attacks. Based on our observations, we also present countermeasures for such attacks.

The paper is organized as follows. In Section II, we introduce the layered system model for the two area network including the data sharing, and computational models. In Section III, we introduce the attacker model. In Section IV, we illustrate an attack for an RTS 24-bus system appropriately modeled as a two-area network. We conclude in Section V.

II. System Model

We consider a two-area network model in which each area uses its measurements to evaluate the state of the system, compute the optimal power flow, and determine generation dispatch. It is assumed that the computations are performed at a local control center as shown in Fig. 1, and henceforth, when we refer to the two areas sharing information, it implies that information is exchanged between the control centers. We make the following assumptions about the information shared between the two areas.

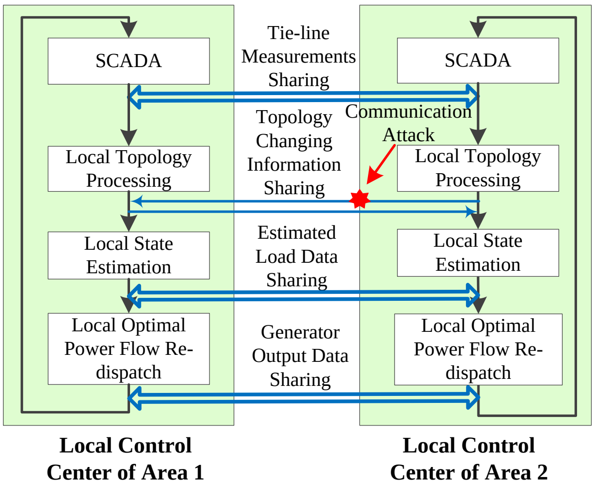
SCADA
Tie-line Measurements Sharing
SCADA
Topology Changing Information Sharing
Communication Attack
Local Topology Processing
Local Topology Processing
Local State Estimation
Estimated Load Data Sharing
Local State Estimation
Local Optimal Power Flow Re-dispatch
Generator Output Data Sharing
Local Optimal Power Flow Re-dispatch
Local Control Center of Area 1
Local Control Center of Area 2

Fig. 1 Computational units and data interactions between the two areas of the network.

A. Information sharing model

To illustrate the distributed effects of a communication attack, we assume that the two areas share as much information as relevant and based on current practices. The primary assumption is that each area performs its own computations with some data (depending on the computation block) obtained from the other side. Our assumptions are as follows:

- Static topology information: The static topology information is shared among all areas of the interconnected power system.
- Dynamic topology information: Each area is assumed to communicate the topology changing information among the whole system in real-time. Thus, once a topology error is found, the local operator should send this information to other areas immediately, which allows them to update the whole system topology information in time.
- Generation: The generation schedule of each unit is shared among areas in real-time.
- Measurements: The tie-line measurements are shared between adjacent areas in real-time. In general, more measurements can be shared but we assume that each area does its own local state estimation (as is often the case in practice).
- Estimated load: The estimated load data is shared among areas in real-time.
- Network models for power flow: Each area computes its own AC OPF. In practice, each area uses a network equivalent model of its connected areas to simplify the OPF computation. However, since we seek to understand the effect of a communication cyber-attack on dispatch and power flows (line

overloads often contribute to outages), we choose the best case network model, i.e., we assume each area uses the complete network model of the other side in computing its OPF. However, each area can only dispatch its own generators, and thus, computes the OPF while keeping the dispatch for the other area fixed according to the generation data sharing model.

B. Computational models

We briefly outline the mathematical model for each of the computational units we consider here. The different computational units and their interactions across the two areas are illustrated in Fig. 1.

1) State Estimation: Each area applies a weighted least-squares (WLS) AC state estimation to calculate its system state (complex voltages) using the measurements from meters in its area as well as tie-line measurements. [12]

2) Optimal Power Flow: Assuming perfect network equivalent models (i.e., complete sharing of neighboring network graphs for OPF), area *i*, $i = 1, 2$, computes the OPF for the entire two-area network while allowing changes in dispatch only for its own area; in other words, area *i* runs its OPF with the dispatch for area *j*, $j = 1, 2, j \neq i$ fixed around values that were shared from the previous time period that area *j* ran its own OPF. The resulting OPF problem can be viewed as each area performing a centralized power flow problem but with the capability to only dispatch local units.

Let *B* and *Br* denote the set of buses and branches in the entire two-area network, and B_i and B_j denote the set of buses in area *i*, $i = 1, 2$ and area *j*, $j = 1, 2, j \neq i$, respectively. Further, let G_n denote the set of generators at bus *n*, $\{G_n\}_{n \in B_i}$ denote the set of generators in area *i*, $i = 1, 2$. We henceforth, use *i*, *i*=1, 2, to denote the area under study and *j*, $j = 1, 2, j \neq i$, to denote the connected area. Let $c_g(.)$ denote the cost function for generator *g*. The OPF for area *i*, $i = 1, 2$ can be formulated as the following optimization problem:

$$\text{Min} \sum_{g \in \{G_n\}_{n \in B_i}} c_g(P_g) \tag{1}$$

$$\text{s.t.} \sum_{g \in G_n} P_g + \sum_{\forall k(n,;)} P_k - \sum_{\forall k(;,n)} P_k = P_{dn} \qquad \forall n \in B \tag{2}$$

$$\sum_{g \in G_n} Q_g + \sum_{\forall k(n,;)} Q_k - \sum_{\forall k(;,n)} Q_k = Q_{dn} \qquad \forall n \in B \tag{3}$$

$$P_k = V_n^2(g_{sn} + g_{nm}) - V_n V_m (g_{nm} \cos(\theta_n - \theta_m) + b_{nm} \sin(\theta_n - \theta_m)) \qquad \forall k \in Br \tag{4}$$

$$Q_k = -V_n^2(b_{sn} + b_{nm}) - V_n V_m (g_{nm} \sin(\theta_n - \theta_m) - b_{nm} \cos(\theta_n - \theta_m)) \qquad \forall k \in Br \tag{5}$$

$$\sqrt{P_k^2 + Q_k^2} \leq S_k^{\max} \qquad \forall k \in Br \tag{6}$$

$$P_g^{\min} \leq P_g \leq P_g^{\max} \qquad \forall g \in \{G_n\}_{n \in B_i} \tag{7}$$

$$Q_g^{\min} \leq Q_g \leq Q_g^{\max} \qquad \forall g \in \{G_n\}_{n \in B_i} \tag{8}$$

$$V^{\min} \leq V_n \leq V^{\max} \qquad \forall n \in B \tag{9}$$

$$\hat{P}_g^{fix} - \Delta\bar{P} \leq P_g \leq \hat{P}_g^{fix} + \Delta\bar{P} \qquad \forall g \in \{G_n\}_{n \in B_j} \tag{10}$$

$$\hat{Q}_g^{fix} - \Delta\bar{Q} \leq Q_g \leq \hat{Q}_g^{fix} + \Delta\bar{Q} \qquad \forall g \in \{G_n\}_{n \in B_j} \tag{11}$$

where P_g is the active power output of generator *g* with maximum and minimum limit $P_g^{\max}$ and $P_g^{\min}$, Q_g is the reactive power output of generator *g* with maximum and minimum limit $Q_g^{\max}$ and $Q_g^{\min}$, b_{nm} and g_{nm} are the susceptance and conductance, respectively, of line *k* from bus *n* to bus *m*, b_{sn} and g_{sn} are the shunt branch susceptance and conductance, respectively, of bus *n*, *k(n,;)* is the set of lines *k* with bus *n* as its receiving bus and *k(;,n)* is the set of lines *k* with bus n as its sending bus, $\hat{P}_g^{fix}$ and

$\hat{Q}_g^{fix}$ are the fixed active and reactive power outputs with $\Delta\overline{P}$ and $\Delta\overline{Q}$ deviation for generator *g* in area *j*, respectively, P_k and Q_k are the active and reactive power flows, respectively, on line *k* with line capacity limit $S_k^{\max}$, P_{dn} and Q_{dn} are the active and reactive power demands, respectively, at bus n, θ_n is the voltage angle for bus *n*, V_n is the voltage magnitude for bus *n* with maximum and minimum limits $V_n^{\max}$ and $V_n^{\min}$, respectively.

The objective in (1) is to minimize the total active power generation cost of area *i*, $i=1,2$. Constraints (2) and (3) represents the active and reactive power balance constraints for each bus in the centralized system (two-area network). The constraints in (4) and (5) are the active and reactive transmission line power flow constraints for the whole system while (6) is the thermal limit for each transmission line. Constraints (7) and (8) are the local (for area *i* only) unit active and reactive power output limits while (9) defines the voltage magnitude limits for each bus in the whole system. Finally, (10) and (11) incorporate the unit active and reactive power output limits for area *j*, $j\neq i$, i.e., the power output of generation units external to area *i* are fixed around the values shared by the other areas.

When no feasible solution (i.e., a solution which satisfies (2)-(11)) can be found, the distributed OPF program fails to converge. In practice, to find a feasible solution, system operators often relax the constraints. In this paper, the thermal limit constraint on the congested line is the first constraint to be relaxed. Multiple iterations of relaxing the line limits may be needed to obtain a feasible solution; to this end, we model the relaxed limits as follows:

$$\sqrt{P_k^2+Q_k^2}\leq S_k^{\max}+u\Delta\overline{S} \qquad (12)$$

where line *k* is the congested line, $\Delta\overline{S}$ is the incremental value by which the line limit is relaxed in each iteration, and $u \leq u_{max}$ is the iteration number. In each iteration, the thermal limit is relaxed by increasing the rating of line *k* by $\Delta\overline{S}$, and the OPF program is executed to check whether it converges. This process is repeated until the OPF program converges or the relaxation time reaches its maximum value. Following this, other important lines (such as those with high reactive power flow) will be relaxed using the same procedure. If both methods fail to work, then we consider the test case as a *not converge case*.

III. Attacker Model

We assume that the attacker has access to the data being shared between areas and can corrupt the data. Examples abound of such data corruption attacks including the oft cited Stuxnet virus attack. The attacker is assumed to either participate in creating a line outage in one area or be aware of such an outage and then act to corrupt the topology information shared with the other area. Our attack model also captures simple human errors in information sharing between connected areas, including delays and miscommunications. In the interest of understanding worst case attacks and data sharing limitations, the area with the outage is assumed to be aware of the outage shortly after. This assumption is based on frequently seen patterns of limited data sharing that precede (and are a cause of) large blackouts.

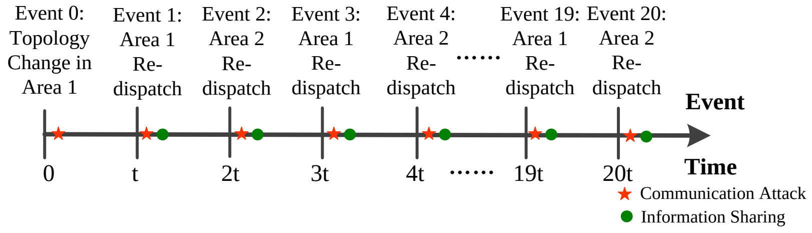
Event 0: Topology Change in Area 1
Event 1: Area 1 Re-dispatch
Event 2: Area 2 Re-dispatch
Event 3: Area 1 Re-dispatch
Event 4: Area 2 Re-dispatch
……
Event 19: Area 1 Re-dispatch
Event 20: Area 2 Re-dispatch
Event
0
t
2t
3t
4t
……
19t
20t
Time
★ Communication Attack
● Information Sharing

Fig. 2 Time sequences of events at the two areas at the time of and following an attack in one area.

In order to understand the effect of such an attack, we study the time progression of the attack. We consider the following time-progression of the attack and system behavior includes the following steps:

1) Event 0: System/Attack: Outage occurs in Area i, *i*=1,2. Area *i* becomes aware of outage and updates its topology in the next computation time period (Event 1) to perform SE, OPF and dispatch.

Attacker replaces updated topology information shared with area j, $j = 1, 2$, $j \neq i$, with the previous static topology information.

2) Event 1: Area i: Area i uses measurements with updated topology to compute SE, OPF, and dispatch. Shares dispatch status with Area j. Attacker sustains attack.

3) Event 2: Area j: Area j uses measurements with updated topology to compute SE, OPF, and dispatch. Shares dispatch status with Area i.

4) Events repeated back and forth until alarms are set off either due to repeated lack of convergence or physical line overloads. All the while it is assumed that the attacker sustains the attack.

We illustrate this time sequence in Fig. 2 for the case in which Area 1 experiences a line outage while Area 2 does not have the real-time topology information following the outage due to a communication attack.

IV. Illustration of Results

In this section, we illustrate our distributed communication attack and its consequences. The test system is an IEEE 24-bus reliable test system (RTS) as shown in Fig. 3, which is decomposed into two areas that are connected by four tie lines. Each area is assumed to have its own local control center that performs local state estimation with local measurements and tie-line power flow measurements shared from adjacent areas, following which it shares its estimated load information with the other area. This is followed by an OPF re-dispatch keeping the generator outputs external of the other area fixed (due to the reactive power supply, the reactive power generator on bus 14 is assumed to be dispatched by both sides). If the OPF re-dispatch fails to converge, constraints on congested lines and other important lines (i.e. #10 in the test system) are relaxed, one at a time as described earlier, with a 2MVA incremental relaxation in rating and a maximum 20 iterations. This process alternates between the two areas every t time units (see Fig. 2).

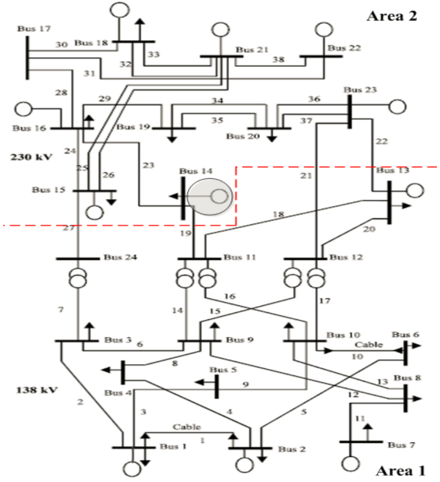

Fig. 3 An IEEE RTS 24-bus divided into two areas (separated by red dashed line).

The attack is modeled as a line outage in one area. In order to understand the worst-case effect of the attack, the area without knowledge of the outage is assumed to have a congested line prior to the attack. The attacker, aware of this outage in one area, compromises the topology changing communication signals such that the same static topology prior to the attack is shared. All possible choices of line outages in one area and congested lines in the other are considered exhaustively to demonstrate the effect of the possible attack cases. The system behavior is followed over 20t time units following the outage and over this time the two areas perform SE, OPF, and dispatch is a round-robin (i.e., one after the other continuously) fashion. The area with correct topology information is assumed to re-dispatch first after topology changing happened. The events sequence when Area 1 has an outage and Area 2 is affected by the communication attack is shown in Fig. 2. The time immediately after topology changing is assumed as Event 0 (denoted E0). After t time units, the area (area i) with the correct topology information dispatches following SE and OPF -- this is Event 1. Event 2 follows time units later when the area (area j) with the false topology information re-dispatches following its own SE and OPF. The two areas continue re-dispatching alternately in the simulation time period to obtain Event *Et* in time *t*.

To include all possible attacks, two directions of attacks are studied. Thus, directions 1 and 2 result when the line outage occurs in Area 1 and the congested line lies in Area 2 and vice-versa, respectively. Since our focus is on worst case attacks, we assume that the area without real-time topology information has some lines congested. This is achieved in simulation by reducing the line rating to 90% of the base case power flow to create congestion. We first document our results in tables and then provide the detailed analysis and plots.

Table I. Post E0 system behavior with sustained attack for both cases of overload and no overload following attack following E0

Total Feasible Case	Overload following E0				No Overload following E0		
	PF Overload Oscillate	PF Overload Violation	PF Not Converge	PF Reduction	No Violation	PF Overload Violation	PF Not Converge
540	35.93%	8.15%	3.70%	3.52%	44.07%	2.96%	1.67%

Table I shows the numbers in percentage of the seven possible long term (20 events) outcomes of an attack after E0, of which 4 outcomes correspond to cases in which there is an overload immediately after E0 and 3 outcomes correspond to the no overload cases after E0. These attack consequences are quantified by comparing the cyber power flow and physical power flow in the area without the real-time topology over the entire attack time duration. The cyber power flow is the OPF solution calculated by the control center with fixed external generation. The physical power flow, on the other hand, is the real power flow values of the system after dispatching with the most recent OPF dispatch solution with the true topology information. Therefore, for the area with false topology information, the cyber power flow values will be different from the physical power flow values. Four kinds of disparities are observed between the cyber and physical power flows if there is an overload in the area with wrong topology information following E0; we name them *PF Overload Oscillating*, *PF Overload Violation*, *PF Not Converge*, and *PF Reduction*. On the other hand, if no overload is observed following E0, then the resulting three disparities are referred to as *PF No Overload*, *PF Overload Violation* and *PF Not Converge*. We describe these disparities in detail below.

E0 Overload & PF Overload Oscillating cases: In these cases, the overload problem on the previous congested line is corrected each time the area with correct topology information re-dispatches its generator outputs. However the overload problem reappears when the other area uses wrong topology information to dispatch at the next event. The power flow on the previous congested line oscillates in the simulation time period. A typical power flow overload oscillating case plot is shown in Fig. 4(a). For such cases, the target lines can get heated repeatedly due to the dispatch of the area with false topology information. The accumulation may eventually cause the line to overheat and trip offline. Thus these cases are assumed to be *successful attack outcome*.

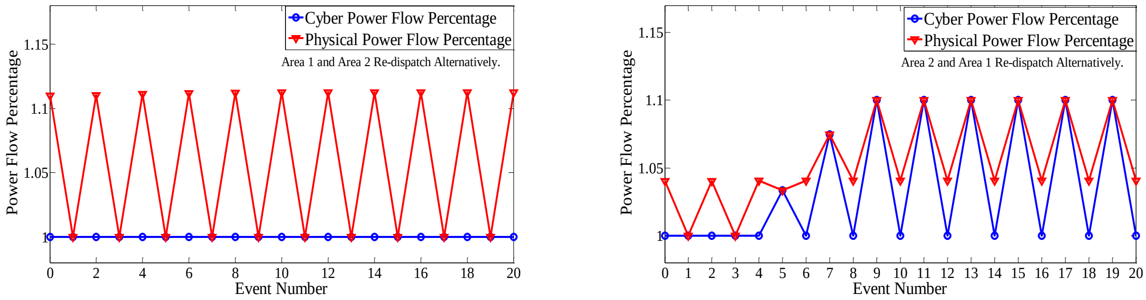

(a) E0 Overload & PF Oscillate: Power flow on line #24 (area 2) when line #3 (in area 1) is tripped.

(b) E0 Overload & PF Overload Violation case: Power flow for Line #2 (area 1) when line #30 (area 2) is tripped.

Fig. 4 Typical power flow variation on congested line for 2 successful attacks.

E0 Overload & PF Overload Violation cases: A typical PF Overload Violation case is shown in Fig. 4(b). In these cases, the power flow oscillate in the first few events following an overload after E0, then no local dispatch plan for the area with correct topology information that can satisfy all the constraints of the system can be found. To obtain a feasible re-dispatch, thermal limit constraints of lines have to be relaxed in the rest events. Thus, after several re-dispatch process, the congested line will keep overloading due to relaxed solution. The heat accumulation may eventually cause the line to overheat and trip offline. Therefore, these cases can be viewed as *successful attack outcomes*.

E0 Overload & Not Converge cases: In these cases, critical topology change happened. To get a feasible solution requires the jointly re-dispatch of both areas since a large amount of active and reactive power output needs to be re-dispatched on both areas. Since overload problem can't be solved by local control center, the worse operation states will continue until more serious consequences happened. Therefore, these cases can be viewed as *successful attack outcomes.*

E0 Overload & PF Reduction cases: For these cases, a line overloaded after E0 can finally reduce below 100% of the rating in the simulation time period. Though the re-dispatch plan of the area with wrong topology still give a wrong calculation values of the system, no further problem caused by the wrong plan. We, therefore, view this attack leading to such cases as an *unsuccessful attack*.

E0 No Overload & No Violation cases: For the cases in which no overload happens immediately after E0, it is possible for the system to continue without thermal limit violations during the entire simulation time period. Therefore, an attack leading to these cases is an *unsuccessful attack*.

E0 No Overload & PF Overload Violation cases: In these cases, although no overload happens immediately after E0 and in the first few events, thermal limit constraints have to be relaxed in the rest events to get a feasible solution. Thus, after several re-dispatch process, the congested line will keep overloading due to relaxed solution. Therefore, this attack can be viewed as a *successful attack*.

E0 No Overload & Not Converge cases: Same as the Not Converge cases of PF overload following E0, the topology change in such cases is critical. Despite no overload following E0, to obtain a feasible solution requires centralized re-dispatch. This attack can also be viewed as a *successful attack.*

We observe a total of 283 successful attack cases, i.e., 52.41% of the total attack cases. We define the subclass of successful attacks for which the power flow of 105% relative to the flow following Event 0 as critical (successful) attacks, and note that the total number of critical attacks for the RTS system is 63, which is 11.67% of the total attack cases. These results demonstrate the potential vulnerability of a topology-based communication attack.

V. COUNTERMEASURES AND CONCLUDING REMARKS

We have introduced a new class of distributed communication (man-in-the-middle) attacks specifically targeting the topology sharing units between connected areas in the electric grid. We have demonstrated the time consequences of such attacks and have shown that such attacks can often lead to serious consequences if active intervention is not present. In this context, we observe that in addition to the traditional countermeasure of human operator-based data sharing (which have been shown to be error-prone and delayed too), it is essential to have more resiliency via automated data sharing mechanisms. Our attack is successful because the two areas process data largely independently except for data sharing and do not employ a more interactive distributed processing platform. This could help both areas become aware of inconsistencies over faster time-scales including: (a) enable local topology processing to include interactive updating; (b) enable real-time coordination of dispatch between the two areas; and (c) create and share a list of external contingencies caused to other areas by an internal component outage. It is worth noting that, while some of these mechanisms are being considered or even used currently in the grid, it is not done in a uniform manner and this work highlights the limitations of not doing so.

BIBLIOGRAPHY

[1] "Federal Energy Regulatory Commission (FERC): Final report on the August 14th blackout in the United States and Canada: Causes and recommendations," http://www.ferc.gov/industries/electric/indusact/reliability/blackout/ch1-3.pdf, April 2004.

[2] "Federal Energy Regulatory Commission (FERC): Mandatory reliability standards for interconnection reliability operating limits," http://www.ferc.gov/whats-new/commmeet/2011/031711/E-8.pdf, March 2011.

[3] "Federal Energy Regulatory Commission (FERC) and the North American Reliability Corporation (NERC): Arizona-Southern California outages on September 8, 2011," http://www.nerc.com/files/AZOutage-Report-01MAY12.pdf, April 2012.

[4] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," in Proceedings of the 16th ACM Conference on Computer and Communications Security, ser. CCS '09, Chicago, Illinois, USA, 2009, pp. 21–32.

[5] O. Kosut, L. Jia, R. J. Thomas, and L. Tong, "Malicious data attacks on the smart grid," IEEE Transactions on Smart Grid, vol. 2, no. 4, pp. 645–658, 2011.

[6] S. Sridhar and M. Govindarasu, "Model-based attack detection and mitigation for automatic generation control," IEEE Transactions on Smart Grid, vol. 5, no. 2, pp. 580–591, 2014.

[7] J. Wei, D. Kundur, T. Zourntos, and K. Butler-Purry, "A flocking-based dynamical systems paradigm for smart power system analysis," in Power and Energy Society General Meeting, 2012 IEEE, 2012, pp.1–8.

[8] J. Kim and L. Tong, "On topology attack of a smart grid: Undetectable attacks and countermeasures," IEEE JSAC, vol. 31, no. 7, pp. 1294–1305, 2013.

[9] L. Jia, J. Kim, R. J. Thomas, and L. Tong, "Impact of data quality on real-time locational marginal price," IEEE Trans. Power Systems, vol. 29, no. 2, pp. 627–636, 2014.

[10] L. Xie, Y. Mo, and B. Sinopoli, "Integrity data attacks in power market operations," IEEE Transactions on Smart Grid, vol. 2, no. 4, pp. 659–666, 2011.

[11] J. Liang, O. Kosut, and L. Sankar, "Cyber-attacks on ac state estimation: Unobservability and physical consequences," in IEEE PES General Meeting, Washington, DC, July 2014.

[12] A. Abur and A. G. Exposito, Power System State Estimation: Theory and Implementation. New York: CRC Press, 2004.