Control of Communications-Dependent Cascading Failures in Power Grids

Jose Cordova-Garcia, *Member, IEEE,* Xin Wang, *Member, IEEE,* Dong-Liang Xie, *Member, IEEE,* Yue Zhao, *Member, IEEE,* and Lei Zuo, *Member, IEEE,*

***Abstract*—The most severe failures in power grids are often characterized as cascading failures, where an initial event triggers consequent failures all along the grid often leading to blackouts. Upon identifying a failure and its cascade potential, timely control actions should be performed by the grid operators to mitigate the effect of the cascade. These actions have to be delivered to one or more control devices, creating a dependency between the power grid and its control network. In this paper we study the dependency of the operation of the power grid on the control network. Different from literature studies on failure control, our dependency model captures the impact of networking parameters. We formulate an algorithmic model that describes the impact of this dependency on cascade control. Based on this model, we propose an efficient cascade control algorithm for power grids. Finally, we evaluate the impact of the power-communication network dependency with uncontrolled grids, ideal/simple control grids and our proposed control scheme, and our results demonstrate that our proposed algorithm can significantly reduce the failure of power lines while sustaining larger power demand for users.**

***Index Terms*—Cascading failure, delay effects, fault propagation, load shedding, networked control, wide-area control.**

NOMENCLATURE

α	Heat parameter
Δ	Expected delays of all nodes
$\boldsymbol{\Theta}$	Nodes voltage angles vector
$\mathbf{B}$	Nodal admittance matrix
$\mathbf{c}, c_i$	Controllability vector, and controllability of node i
$\mathbf{P}$	Nodes injections vector
$\mathcal{P}$	Goal powers determined by control scheme
E	Set of edges (power lines)
F	Record of failed lines
N	Set of nodes (substations)
T	Power grid graph
$\delta(\pi_i)$	End-to-end delay from control center to node i
μ_{ij}	Thermal capacity limit of line (i, j)
π_i	Path from control center to node i
θ_i	Voltage angle of node i
$\tilde{f}_{ij}$	Heat condition of line (i, j)
f_{ij}	Power flow of line (i, j)
$P_i, \hat{P}_i$	Power injection and goal power of node i
r_n	Grid state of uncontrolled cascade
s_m^c	Grid state due to changes induced by control packets
x_{ij}	Reactance of power line (i, j)

Manuscript received June 29, 2017; revised February 11, 2018.

J. Cordova-Garcia is with the Department of Electrical and Computer Engineering, Escuela Superior Politecnica del Litoral (ESPOL), Guayaquil, Ecuador email: jecordov@espol.edu.ec

X. Wang and Y. Zhao are with the Department of Electrical and Computer Engineering, Stony Brook University, Stony Brook, NY, 11794 USA e-mail: {x.wang, yue.zhao.2}@stonybrook.edu.

DL. Xie is with the Beijing University of Posts and Telecommunications, Beijing, China email: xiedl@bupt.edu.cn

L. Zuo is with the Department of Mechanical Engineering, Virginia Tech, Blacksburg, VA 24061, USA email: leizuo@vt.edu

I. INTRODUCTION

A CASCADING failure of the transmission network of a power grid has a great negative impact on the operation of the network. While unusual, it can cause blackouts and considerable losses, both operational and economical. Such events are triggered by failures in either the power lines or the substations. The vegetation, climate or accidents make power lines more susceptible to failure than substations. Both failures can also occur due to attacks.

Different from data flows, power lines carry electric flows, which cannot be intentionally managed but follow the laws of physics. Once a power line fails, the lost power flow is automatically carried over the remaining lines, and these additional flows may make one or more of the operating lines exceed the capacity. Power lines can carry excessive flows for some period of time before they are heated up to a certain level and become inoperable. A cascading failure can be initiated once one or more power lines fail.

Due to the importance of power lines in the transmission network, they are constantly monitored, and the states of the lines are sent to a control center (CC) through a communication network. Based on the monitoring data, operators at the CC can take preventive and corrective actions, often at the substations of the power grid. Example actions include the manual tripping of lines, load shedding and varying the output of power generators.

During a cascading failure, several *rounds* of changes in power levels and topologies can take place in the power network until no further failures occur and the power grid regains normal conditions. It has been shown [1] that failures under a cascade process can take place in the whole extension of the power grid, even in locations far away from the initial failure.

Existing control used in the power grid is based on offline rules to set up local protections that will be automatically triggered upon failures. In this approach, overloaded power lines are tripped without considering that the consequent power flow redistribution can overload the remaining lines. Thus, such an automatic response to power line failures is ineffective and can initiate or extend cascading failures [1].

A situation-specific control strategy can be planned by the CC and transmitted through the communication network with the goal of mitigating cascades. However, as the cascade is a temporal process, the delay in delivering control packets plays a crucial role on the effectiveness of the control strategy. A delayed control action not only leads to ineffective stopping of cascading failures, but could also trigger other failures in the grid which aggravates the problem.

To summarize, this paper makes the following contributions:

- Presenting the problems associated with the conventional control of cascading failures in power grids.
- Proposing a communication-dependent cascading failure model to capture the failure behaviors.
- Designing a simple control policy to reduce the extension of the cascading failures under the communication-dependent cascading failure model.
- Evaluating our proposed control policy and comparing its performance with a representative group of policies from peer work. The results demonstrate the effectiveness of our control algorithm in different scenarios while maintaining at least the same level of power supply to customers as that of peer schemes.

The rest of the paper is organized as follows. We present the related work in Section II. We introduce the system model that motivates our work in Section III. In Sections IV and V, we introduce and illustrate our problem, and present our proposed control mechanism. Simulation results are given in Section VI. We conclude the work in Section VII.

II. Related Work

Although the power grid has been widely deployed for many years, it has been only in the past few years that the Smart Grid (SG) efforts have been initiated to enhance its performance. Computational and communications advances can provide tools for great improvements in power management.

Due to its importance and applicability, simulation models for cascading failures have become the focus of recent SG literature. In [2], authors consider two levels of system dynamics, short-term fast dynamics and long-term slow dynamics. The authors consider the failure of lines due to heating as part of the fast dynamic processes. A non-linear temperature model, as a function of power flow, is used to describe the heating status of lines. Our focus in this paper is on the power flow changes induced by failure control that can cause the tripping of overheated lines, and the heating model study is complementary to our work.

In [3], cascading processes of different time scales are represented through a quasi-dynamic framework, which avoids incurring in a large computational cost to simulate a highly dynamic system. Line outages due to overheating are simulated by randomly selecting lines in the grid to disconnect. Differently, in this paper we enable the operator to evaluate the effect of failure control on line overheating by using the low complexity heat model based on flow moving average. Such a model allows us to evaluate different scenarios of line overheating after a control is applied and when the cascade ends. Moreover, in [1] a thorough study of cascading failures using grid data from the western US interconnect system is presented. The study verifies the validity of using cascading failure models [4] based on linear moving averages. Also, a load shedding algorithm is presented to control failures. Different from [1], we do not consider the communication of control packets ideal. Instead, we incorporate packet delays into the cascade model and control formulation.

Recent literature has started to highlight the importance of considering communication delays when designing efficient system-wide control for power grids. The communication delays in the delivery of control messages from the control center to power nodes are examined in [5], [6], where the popular control delay model in [7] is validated through simulations and network measurements [8], [9], [10], [11]. Instead of studying delay characteristics in the SG environment, we evaluate their effect on the cascade and on control effectiveness. Efforts in [12], [13] propose modifying the communication network to guarantee low delays in the delivery of wide-area control commands. Rather than changing the control network structure to lower delays, in this work we propose that power operators incorporate delay parameters into the control program in order to improve its effectiveness under adverse delays conditions. In [14], authors characterize the development of cascading failures due to delays in the transmission of load shedding messages. However, the delays that affect control in [14] are much longer than the ones present in commonly deployed communication networks. Differently, the goal of this work is to improve the effectiveness of cascade control when different load shedding packets experience different delays of any length.

Moreover, our proposed scheme can be useful to complement recent advances on the risk study of cascades [15], which introduces a risk assessment approach based on Markovian tree searches to find risky states a power grid can reach during cascading failures. Our control methodology can be incorporated into such an analysis to find risky states where delayed packets make control completely ineffective.

Other works as in [16], [17] study the problem of vulnerability and robustness of power grids. Different from [1], these works introduce models of interdependency between the power grid and the communication network that controls it. The basic idea behind these models is that: a power node depends on a control node for proper operation, while the control node needs power to perform control. This interdependency model relies on the power conditions only without considering the state of the communication network.

As a direct extension of the dependency models above, the control mechanisms in [18], [19] simply remove nodes that cannot be controlled or powered and attempt to mitigate the cascading effects with the remaining nodes. We compare our scheme with such kind of control mechanism in section VI. Furthermore, different from the aforementioned schemes, we assess the effect of control on the development of a cascade failure considering the delay of the communication network, and propose a control mechanism that takes into account the delays of all nodes in the network to reduce the size of consequent failures.

III. System Models

In this section, we describe the network models used for the power transmissions and control communications, as well as a cascading failure model.

A. Transmission Network of Power Grids

A line failure in the transmission network of the power grid can interrupt power delivery from generators to distribution substations that serve thousands of customers or other substations. We focus only on the transmission network of the power grid and use *transmission network* and *power grid/network* interchangeable.

A transmission network is represented by the graph $T(N, E)$, where the node set N corresponds to substations and the edge set E corresponds to the power lines connecting substations. Each node $i \in N$ is associated with a value of power, in the unit of MW, P_i. A supply node i corresponding to a generator substation has a positive P_i, while a demand node (i.e. a load) such as a distribution substation has a negative P_i. A node can also have $P_i = 0$, which means it is a neutral node that does not consume or generate any power. The transmission network T delivers power through the lines E covering thousands of kilometers.

For its computational tractability, the DC model is commonly used as a relaxed version of its AC counterpart when analyzing transmission networks [20]. We thus use the DC power flow model to characterize the behavior of the power grid, which is expressed in its matrix form as:

$$\mathbf{P} = \mathbf{B} \cdot \mathbf{\Theta} \tag{1}$$

where $\mathbf{B}$ is the nodal admittance matrix [21]. In this model, a power flow f_{ij} depends on the power line characteristics, namely reactance x_{ij}, and voltage angles θ_i as $\theta_i - \theta_j = x_{ij} f_{ij}$. The active powers are represented as a vector $\mathbf{P} = [P_1, \ldots P_i, \ldots P_N]^T$ where each P_i in (1) corresponds to $\sum_{j \in N(i)} f_{ij} = P_i, \forall i, j \in N, (i, j) \in E$ and $N(i)$ is the set of neighbors of i.

B. Communication Network for Power Grid Control

The proper operation of the power transmission network is controlled using a dedicated communication network. Although these networks have been used for a while (e.g. the SCADA system), different communication technologies, both wired and wireless, are still under deployment or modernization to provide enhanced performance through timely control, a goal for future smart grids.

Consider a graph $C(N_c, E_c)$, where N_c is the set of nodes in the communication network that provide control functions to the power grid, and E_c the set of communication links that connect control nodes between themselves and connect control nodes with the control center.

The number of nodes in N_c to control the power grid is restricted by its extension and topology. The minimum control network for a Smart Grid requires C to have at least one $n_i \in N_c$ to control each $i \in N$. A one-to-one interdependency model makes $|N_c| = |N|$.

Due to the topology restriction of power grid, commonly the topology of the control communication network follows the topology of the power network. Control packets are originated at a control center (CC) and sent to a control node $n_i \in N_c$. Upon a power network failure, the states of the links $e_{(n_i, n_j)} \in E_c$ determine the performance of the corrective actions taken at CC. Intuitively loss and delay of control packets can affect the expected results of a corrective action.

More specifically, after a decision is made at the control center, certain number of control packets traverse C through paths of the form $\pi_i = \{CC, h_1, h_2, \ldots, n_i\}$ in order to control the power grid node i. Traversing each communication relay node h_k involves the transmission delay and the propagation delay. Moreover, the delay can dramatically increase under high data traffic, equipment failure or cyber attacks.

The delays experienced in paths to controllable nodes can further extend the level of interdependency of these networks, and impact the effectiveness of a control action taken at a controllable node.

The CC determines the shortest path π_i to node n_i and the estimated end-to-end delay can be calculated as [7]:

$$\delta(\pi_i) = \sum_{h_j \in \pi_i} \frac{P_s}{R_{h_j \to h_{j+1}}} + \sum_{h_j \in \pi_i} \frac{d_{h_j \to h_{j+1}}}{v_{h_j \to h_{j+1}}} + \mathrm{E}[T_{br}] \tag{2}$$

where P_s is the packet size and $R_{a \to b}$ corresponds to the link rate between nodes a and b. T_{br} is the service latency induced by all the hops of π_i which depends on the state of the control network, thus we use its expected value. The second term in (2) represents the propagation delay that depends on the medium used for each communication link in E_c, and the extension of the network. $d_{a \to b}$ represents the distance between nodes a and b and $v_{a \to b}$ the corresponding propagation speed, set according to the technology used for each communication link.

C. Cascading Failures of Power Grids

A failure in the power grid corresponds to the removal of a single or a group of nodes/edges from the graph $T(N, E)$. A node removal represents the shutdown of a substation, while an edge removal represents a power line failure. In this paper we focus on power line failures.

Natural disasters, equipment malfunction, and intentional attacks are among the causes of power line failures. Thus, the initial failure that triggers a cascade can be located at any grid location and affect any number of lines. In the graph model, after the initially failed lines are removed, the power flows are automatically (and uncontrollably) reassigned in the remaining graph(s). A cascading failure model can be summarized in the following steps:

1) Remove failed lines(s) from the grid graph.
2) Modify power injections such that: $\sum_{P_i \in \mathbf{P}} P_i = 0$.
3) Solve the model in (1).
4) Determine lines that exceed their capacity limits, mark them as failed lines.

5) If there are new lines marked as failed line(s), repeat, else, stop the cascade.

In Step 1), following the initial failure and without control, overloaded lines will be tripped by local protections, which extends the initial failure thus causes a new "failure" in the system. Step 1) could disconnect the original graph into several components, i.e. islands. In this case the model is applied for each component. Note that this cascade model provides no control, and the changes in generators and loads in 2) are only necessary to make 3) feasible. Also, in Step 4, the model inspects each flow $f_{i,j}, \forall(i,j) \in E$ obtained in Step 3 and determines if the line has exceeded its capacity $u_{i,j}$, if so, it is considered as a new failure. Note that a simple control under this model would affect the power injections in Step 2 as determined in the CC, however it does not consider the effect of the communication network on the cascade development.

IV. Problem Description

In this section we illustrate the problem of controlling power grid failures under a communication-dependent cascade model. First, we review a naive approach to control cascading failures, the ideal control. Then, we raise two issues that build up the mentioned problem under the interdependency between the power and communication networks introduced in the previous section.

A. Ideal Control of Cascading Failures

Once a failure is identified at the control center, the corrective action is taken to modify the power injections of nodes in order to maintain all power flows within their range of operation: $f_{i,j} \leq \mu_{i,j}\ ,\ \ e_{i,j} \in E$, where $u_{i,j}$ defines the capacity of the power flow on line $e_{i,j}$. Such an action can require several or even all power nodes to perform a modification to stop the cascade. However, changing the power injection at loads directly affects the number of customers served and power generation cannot be abruptly changed, but only slowed down. Thus, this strategy should also attempt to make these changes as small as possible.

Such strategy can be formulated through a linear program. Using the model introduced in the last section, the result of such program, i.e. the $\hat{P}_i$ values associated with this control policy, can be used in Step 2 of the cascade model to mitigate the cascade failure. Such corrections $\hat{P}_i$ are transmitted as control messages, and the model in III-C implies that all such messages are delivered simultaneously and with no delay as soon as a control decision is made at CC, which is not realistic. A control policy should consider the delays of the links of controllable nodes and moreover the model used for evaluating the cascade behavior should also take the state of the communication network into account.

B. Communications-Dependent Control

For illustration, consider the cascade evolution diagrams presented in Figure 1. Labels F represent the time instant a new failure occurs, i.e. lines are disconnected. Dashed lines represent the power flow changes induced either by failures or the modification of power at controlled nodes. Following power flow changes, the moment that overloaded power lines exceed their capacity limit is marked by black solid circles on the time line. Local protections disconnect the lines that exceed their limits after some time (set by the operator) originating a new cascade round, marked with solid red arrows. In the point of view of the cascade evolution, the disconnected lines are newly failed lines that contribute to the duration of the cascade over time. In figure 1a we illustrate the evolution of a cascade over time when no control is applied. Failed lines in F_0 initate the cascade, the power flow redistribution in the remaining lines leads to some lines becoming overloaded, marked with dashed orange arrows and labeled with $f^t_{j,k} > \mu_{j,k}$ in the figure. Eventually, overloaded lines trigger their protection equipment disconnecting them from the grid, F_1, and creating a new cascade round r_1. This process is repeated until no more lines are overloaded and the grid reaches a balanced status, usually resulting in a large blackout.

The case of ideal control of cascading failures is shown in figure 1b. Blue arrows represent the time of arrival of a control packet to its associated power node, $\hat{P}_i$. The label s^c represents the state the power grid reaches after a control message is received and executed at a power node. Ideally, all N control messages arrive simultaneously represented by the single blue arrow $s^c_{1\ldots N}$. The new $\hat{P}_i$ associated to $s^c_{1\ldots N}$ causes flow redistribution that should lead the grid to reach a state where $f^t_{j,k} < \mu_{j,k}$ for all power lines. Thus, overloaded lines regain their normal operation condition and the next cascade round is not realized (shown with the blue "x" marks).

There exist two major implicit assumptions in such an ideal control. First, in order to stop the cascade, the ideal control should drive all power lines to work under their capacity limits. This requires an accurate model of the operation of power lines to be used as a power flow restriction. The assumption that the control center can accurately describe the temporal evolution of every power line and its relationship with its thermal limit is overly optimistic.

Second, the communication network used to transmit control packets is ideal, i.e. no communication delays occur. In 1c we show a more realistic scenario, where control packets arriving at different time instants due to communication delays. Every time a control message arrives at a node i, a change in the corresponding P_i occurs. Hence, this results in new values for power flows, not only on the lines directly connected to the node i but also anywhere in the network according to (1). This is illustrated with dashed blue arrows.

In terms of the operating condition of the power line, the modification of power flow results in heating cost, even for lines that are not overloaded. Thus, consecutive modifications of power flow represent higher heating possibility, hence a higher chance that the grid is driven to a state where power lines exceed their thermal limits. While under the ideal control the operator can expect the power lines to reach an operating condition that is under the thermal limit, under realistic communication delays such an assumption compromises the effectiveness of control. As shown in figure 1c, after all delayed control messages arrive, the grid reaches a state where lines exceed their thermal limits, hence eventually those lines

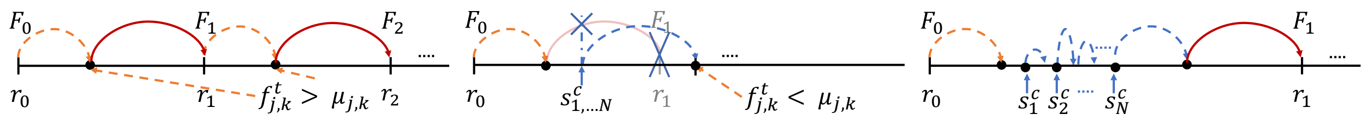

(a) No control cascade evolution (b) Ideal control cascade evolution (c) Delayed control cascade evolution

Fig. 1. Effect of Control Policy packets delay on cascading failures.

in F_1 are disconnected and the cascade is not stopped.

These two issues bring up the following question: *Can we find a control policy that minimizes the adverse effect resulting from the different arrival time between control actions taken at different parts of the power grid?*

The problem we are interested in this paper follows the insight of the previous discussion: Given a communication network that controls a power grid, and the need of proper operation of the power grid with balanced P_is and flows $f_{i,j}$ within capacity, upon a failure, find a control policy that:

1) Minimizes the change of power injections P_i while regaining balance with $f_{i,j}$ under limits of operation.
2) Reduces the number of new tripped lines caused by link delays in the arrival of control packets.
3) Reduces the thermal effect of flow variations due to the inevitably different arrival time of control commands to different control points.

Note that the last two goals imply the need of using the power grid and communication network together in order to plan the control policy. Moreover, a new model of cascading failures is also needed in order to evaluate the effects described in this section.

V. Control of Communications-Dependent Cascading Failures

To more effectively control the cascading failures, we first extend the simple cascading failure model in III-C to account for the effect of grid control in the presence of delays in the communication network. Then, taking into account the communication delay and different arrival time of control commands, we propose a control policy to reduce the damage of the cascading failure to the power grid.

A. Communications-Dependent Cascading Failure Model

Following the discussions in IV-B, $\delta(\pi_i)$ defines another level of interdependence between the power and communication networks, as the delay of control packets affects the consequent behavior of the power grid upon their arrival. Our model of communication-dependent cascading failures based on $\delta(\pi_i)$ is shown in Algorithm 1.

Two different "round" indexes are defined in our model, r_n and s^c_m. The index r_n corresponds to the time when new failures appear in the power grid under the uncontrolled cascade evolution. The cascading failure model captures this behavior through the iteration at the Step 5 of the simple cascading failure model in Section III-C.

The second round index s^c_m depends on the communication network C and represents any change of grid status due to the control, even if such a change is not a response to a line disconnection. The execution order of a control action m depends on the delay that the corresponding control message experiences. Then, s^c_{max} corresponds to the expected status reached by the grid once all control messages have been delivered.

A control scheme determines how the powers $\{P_i \mid i \in N\}$ should be re-dispatched, i.e. through load shedding, in order to mitigate the cascade. Let $\mathcal{P}$ collect the new values of power injections after shedding as determined by a control scheme. Once a control packet is transmitted to its destined node i, it will modify the corresponding P_i to a new value $\hat{P}_i$, which causes changes of power line flows $\tilde{f}^{r_{n+1}}_{j,k}$ and consequently results in a new state of the power grid. s^c_m accounts for changes induced by a control message m that, if delivered, is executed before the next failure. The model described by Algorithm 1 uses $\mathcal{P}$ to account for the effect of control on the cascade development. Lines 10-20 show that the occurrence of r_{n+1} may be altered through the control-induced "rounds".

As a result of the delay of the communication network, each control packet defined in $\mathcal{P}$ may arrive at a different time instant $\delta(\pi_i)$. Thus, with $m = 1$, s^c_1 corresponds to the round at min$\{\delta(\pi_i)\}$ when the first control packet arrives at a power node. Every time a control action is executed, there is one less message in $\mathcal{P}$ to impact the cascade evolution (line 19). Ideally, once all control packets have arrived, the final status of the grid at s^c_{max} should not result in any overload, hence F should be empty and there will not be another round at r_{n+1}. However, as discussed before, power lines after all control actions will only work under capacity if the thermal capacity model is completely accurate, which is an overly optimistic assumption. Moreover, power flows at each s^c_m, caused by the difference in control time at different nodes, modify the operating condition of the power lines and contribute to their heating-up. Thus, the occurrence of r_{n+1} is subject to a physical model of the line condition as determined by the flows at every intermediate s^c_m instead of solely by r_n. A new cascade round r_{n+1} will occur if at least one line exceeds its capacity (line 17) at any time during or after the control. While unlikely, our model can evaluate failures occurring before all control messages have been delivered (line 7). Also, our model can evaluate cascade behavior when control is not available. This model will be used to evaluate the cascading failure control proposed bellow.

B. Communication Delay Dependent Control

To reduce the negative effects of cascading failures induced by delayed control packets, we propose a novel control policy to reduce the size of F. When making a simple schedule for networked control, a control message is generated for every

Algorithm 1 Communications-Dependent Cascade Model

Input: Grid $T(N, E)$, initial failure F_0, control policy $\mathcal{P}$

$F = F_0$, $n = 0$, $m = 0$, $F^{r_n} = F_0$, $F^{s^c_m} = \emptyset$
$\mathcal{P}$ contains all power modifications $\hat{P}_i$ specified in control

while $F^{r_n} \cup F^{s^c_m}$ is not empty OR $\mathcal{P}$ is not empty **do**

  Remove failed lines: $E = E \backslash F^{r_n}$, and set $F^{r_n} = \emptyset$

  Remove failed lines: $E = E \backslash F^{s^c_m}$, and set $F^{s^c_m} = \emptyset$

  Balance supply and demand in the grid

  Compute flows $f^{r_{n+1}}_{j,k}, \forall (j, k) \in E$ using (1)

  Compute heating effect, due to uncontrolled grid:
  $\tilde{f}^{r_{n+1}}_{j,k} = \alpha |f^{r_{n+1}}_{j,k}| + (1 - \alpha) \tilde{f}^{r_n}_{j,k}, \forall (j, k) \in E$

  Add to $F^{r_{n+1}}$ all lines $\tilde{f}^{r_{n+1}}_{j,k} > \mu_{j,k}$ that exceed their protection deadline

  Update total record $F = F \cup F^{r_{n+1}}$

  if $\mathcal{P}$ is not empty **then**

    Initialize $\tilde{f}^{s^c_m}_{j,k} = |f^{r_n}_{j,k}|$

    Estimate delays of $\hat{P}_i$s not delivered, $\delta(\pi_i)$

    Sort $\mathcal{P}$ to follow an ascending order specified by δs

    Modify P_i according to $\hat{P}_i, i$ corresponds to the first element of $\mathcal{P}$

    Compute flows $f^{s^c_{m+1}}_{j,k}$ using (1)

    Compute heating effect, due to control:
    $\tilde{f}^{s^c_{m+1}}_{j,k} = \alpha |f^{s^c_{m+1}}_{j,k}| + (1 - \alpha) \tilde{f}^{s^c_m}_{j,k}, \forall (j, k) \in E_n$

    Add to $F^{s^c_{m+1}}$ all lines $\tilde{f}^{s^c_{m+1}}_{j,k} > \mu_{j,k}$ that exceed their protection deadline

    Update total record $F = F \cup F^{s^c_{m+1}}$,
    and $\tilde{f}^{r_{n+1}}_{j,k} = \tilde{f}^{s^c_{m+1}}_{j,k}$

    $m = m + 1$, Remove the first element of $\mathcal{P}$

  end if

  $n = n + 1$

end while

node in the grid with different amount of power shedding, thus $|\mathcal{P}| = |N|$. As described in Algorithm 1, the grid state resulted from controls for different nodes determines the duration of the cascade. The assumption of the ideal scenario where all control messages arrive at the same time results in a single grid state s^c_{sync} where the subscript $sync$ indicates that all control packets get delivered and all required changes take place synchronously. Thus, the heating effect of such changes is minimized. However, control messages may experience different delay in a realistic communication network. As a result of the difference in the time of controlling different nodes, the already stressed grid may undergo further large changes in power flow patterns that accelerate the line overloading.

To reduce the large power flow fluctuations and thus the duration of the cascade, the delayed control packets should have a smaller effect on the load change of a line. Thus the power grid can more quickly reach normal status. Following this control policy, we can drive the cumulative power line operating condition with the control amount decreasing with the packet delay.

To realize this control policy, we first define a *controllability* vector $\mathbf{c}$. For a node i, the controllability factor c_i is set based on the difference between the estimated transmission delay $\delta(\pi_i)$ of a control packet to the node and a parameter Δ. During a favorable network status, there is low delay variation and Δ can be set to the expected average delay of all nodes in the grid. When nodes experience longer than expected delays the operator can set Δ accordingly using the average of the measured past transmission delays.

We introduce this controllability factor to determine how much load shedding should be performed for a specific node in order to address a failure. More specifically, we would like the control policy to set low power changes, and consequently create less effect on the grid, to nodes whose delays exceed Δ. On the other hand, nodes experiencing low delays can perform more shedding so that larger power flow changes occur as fast as possible. To induce such a shedding behavior, we define c_i as an exponential function of the difference between the expected delay and the delay experienced by the node to be controlled as follows:

$$c_i = e^{\delta(\pi_i) - \Delta}. \tag{3}$$

The exponential function allows us to rapidly reduce the shedding at nodes whose delay values grow larger than the expected. Using (3), we can derive a control policy to minimize the power changes in the grid, while considering the effect that the communication network has on the development of cascading failures as a result of these changes. We build control policies that achieve this goal with the following linear program:

$$\underset{\mathcal{P}}{\text{minimize}} \quad \mathbf{c} \cdot |\mathcal{P} - \mathbf{P}^{r_n}| + \beta ||\mathbf{P}^{s^c_{max}} - \mathbf{P}^{r_n}|| \tag{4a}$$

$$\text{subject to} \quad \mathcal{P} = \mathbf{B}^{r_n} \cdot \mathbf{\Theta}\,, \tag{4b}$$

$$|\hat{P}_i| \leq |P^{r_n}_i|, \forall i \in N\,. \tag{4c}$$

$$f^{s^c_{max}}_{j,k} < \mu_{j,k}, \forall e_{j,k} \in E \tag{4d}$$

where $\mathcal{P}$ contains the target set of power $\hat{P}_i$s that make up the control scheme, and r_n is the cascade round where control is applied. We consider applying control as soon as a cascade is detected at r_0. The changes required to achieve the target set of $\hat{P}_i$s correspond to shedding power at load nodes as well as the possible ramping down of generation units. As the time it takes to ramp down a generator can be longer than the communication delay, the grid may have already been driven to the desired operation values when the generator is reduced to a lower value. In order for power lines to reach the desired operating points, the formulation in (4) can be modified to restrict the usage of generator nodes in the control policy to alleviate the side effect of ramping delay. However, our paper provides no node restrictions with the intention of evaluating a larger solution space with all nodes available for control. The exponential function in (3) enables the control scheme to assign much lower load shedding as the delay grows longer than Δ. In (4a) we have included a regularization term with the parameter β that allows the control center to minimize the number of large changes required for control. Thus, the operator can choose a large beta, usually larger than 100 in our simulations, to induce large load shedding to be performed only by a few nodes. A reduced number of control actions

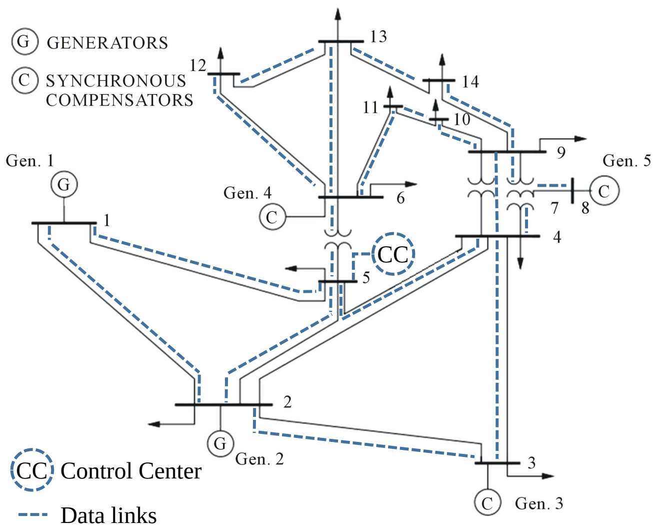

Fig. 2. IEEE 14 bus power grid and its control communication network.

that require large changes can minimize the stress imposed on the grid with the control of all nodes, particularly for larger systems. Furthermore, the power operator can introduce security mechanisms, only for the nodes required to perform large changes, without the otherwise large overhead incurred when protecting communications of every node in the grid.

The objective function in (4a) can also be seen as a cost-based relationship between changes in the power grid and the delay in the communication network. Constraint 4b models the DC power flow model behavior as explained before, and 4c limits loads and generators to be within their original normal operation values.

Also, (4a) aims to:

1) Use all controllable nodes in order to create a policy that mitigates the cascade by the time the most delayed packet arrives, corresponding to the index r^c_{max}
2) Change power injections considering the tradeoff between the change of power of the nodes being controlled and the arrival time of the packets that command the node to change power.

Our proposed policy does not explicitly exclude nodes with very high message delay from control, but instead makes their power changes smaller to reduce the effect of changes on the development of cascading failures. With our policy, a node i with a higher controllability will have a larger modification of its power P_i in the optimization equation, thus reducing the amount of power to change. This helps to increase the resiliency of control against adverse network delays thus further reducing the chance of extending the cascade due to post-control line overloads.

VI. Performance Evaluation

In this section, we evaluate the performance of our proposed control scheme under the model introduced in section V-A. We also introduce the metrics used for evaluation and the control strategies of peer works.

We will first introduce the simulation setup with the parameters used for performance studies. We will then show the results obtained for scenarios based on network sizes, different adverse delay scenarios, different behaviors of power lines due to flow dynamics and new failure protection deadlines.

A. Metrics and Evaluation

Cascading failures occur when the size of an initial failure extends with the appearance of new power line failures all around the power grid. Under the communications-dependent cascade model and due to the effects of dynamic changes in the power line flows, a number of $|F|$ lines will be overloaded and eventually disconnected after a control policy is applied. We use the number of "hot" lines after control to evaluate the performance of our proposed scheme. Also, we show the demand supplied to the customers at the end of the cascade. That is, we are interested in the demand left for customers when the cascade stops and no further line disconnections occur, after control.

Furthermore, we will compare our proposed control mechanisms with some peer work described in section II. As discussed before, such schemes use only a power criteria in the objective function and are either not clear how to establish a delay criteria [19] or delays are ignored [18]. Given that this type of strategies basically reduces to the simple power-based control, we refer to this type as "Simple control".

B. Simulation Setup

The communication network C used to transmit control messages follows the topology of the power grid network. The variations of the topology of C, albeit possible, are restricted by the large extension of the power grid. Power lines define paths connecting substations, thus communication links follow these paths. To account for topology variations, we set the path lengths to be uniformly distributed within [500, 4000] km. A substation in the power grid is randomly selected as the control center, and communication delays are estimated based on the paths to each controlled power node. The link data rate is determined by the technology (e.g., wireless, optical) used for the communication network, and can be estimated based on the link capacity. To estimate the service delay T_{br} in (2), we set the link rate to 2Mbps. The incoming and outgoing transmission rates of routers at each power node are set to 100Mbps and 50Mbps respectively, which are common parameters used in wide area control [7], [8], [9], [10], [11]. The power grids tested are the IEEE 14, 30, 57 and 118 bus systems with 20, 41, 80, 186 power lines. The IEEE 14 bus power grid is presented in Figure 2 for illustration. In the figure, node 5 has been randomly selected as the control center. Control packets generated at node 5 will be delivered to the grid nodes following the data paths depicted by the dashed lines. Also, the control computation time observed in our evaluations was in the millisecond range, which is also commonly used in the literature [8], [9], [10], [11]. Although power lines can be heated up during the control computation, the computation time is generally much smaller than the time it takes for a line to exceed its thermal limit.

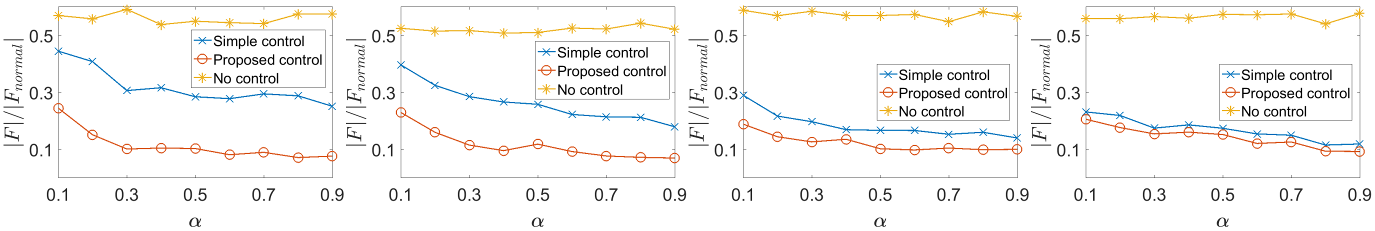

(a) IEEE 14 (b) IEEE 30 (c) IEEE 57 (d) IEEE 118

Fig. 3. Effectiveness of the proposed control in terms of the ratio of lines that can create a new cascade round after the control is applied, $|F|/|F_{normal}|$. α defines different power line heat conditions.

We are interested in evaluating the performance of control during and after a cascading failure. Once initiated, the cascade follows the round-based behavior as presented in sections III-C and V-A, regardless of the size of the initial failure. Such models represent discrete realizations of the temporal cascading process. Thus, the initial failure that triggers the cascade is set to be a random selection of 10% of the lines in the grid. This choice of the initial failure size is large enough to trigger a cascade while avoiding unusual scenarios where a large number of lines fail simultaneously. Moreover, from the system-wide operation point of view, it is important to evaluate the status of the grid when control is applied and once the cascade stops. This allows the power operator to be aware of control effectiveness and plan further actions accordingly. Therefore, we will evaluate the proposed control using the discrete round-based model of section V-A rather than examining the precise occurrence time of the cascade events. Furthermore, the operator can interpret scenarios defined by high and low values of the parameter α in Algorithm 1 as cascade rounds occurring at fast and slow time scales. In addition, we specify a “protection deadline” as the time it takes an overloaded line to be disconnected by its protection device. The deadline setup for the power line protection devices is performed offline by power operators. Thus, α determines how overloaded lines are heated up to exceed their capacity and trigger the running of their protection devices to disconnect them after a predefined time duration. Common values for the protection deadline setting are in the range of milliseconds [6], [9]. Then, when control is available, line disconnections depend on the different control delays, heating parameters α and the protection deadlines. In order to highlight the effect of control packet delays, we consider all grid lines to have similar heating characteristics, i.e. αs are set to be the same, and all protection devices share a common setting.

The default parameters, unless otherwise specified, are $\alpha = 0.5$, $\Delta = 2 \cdot E[\delta]$ with $\delta = [\delta(\pi_1), \ldots, \delta(\pi_N)]$, and the protection deadline is set to 100ms. An α value of 0.5 indicates that the status of a line depends equally on its previous condition and the new one induced by control, i.e. line status. A protection deadline of 100ms is large enough to expect all control actions to be effective as long as their delays do not exceed the deadline. With Δ set to be double the expected delay among all nodes, our proposed control is driven to only avoid large shedding of load in nodes experiencing very long delays.

Different values of α, Δ, and protection deadlines will be tested to evaluate the effectiveness of the proposed control mechanism. These parameter evaluations represent different timescales of cascades. Furthermore, as described in sections III-B and IV-B, communication delay depends only on the status of the control network and its timescale should be in the millisecond range under normal delay conditions. We are also interested in evaluating adverse communication delay conditions where the timescale difference between a line disconnection and the data delay varies greatly, i.e. delayed packets could arrive after the line disconnection. For example, using the default parameters, the expected average communication delays will be well under 100ms [7], [8], [10], [11]. With $\alpha = 0.5$, and following Algorithm 1, lines do not become instantly overheated. Instead, with such an α the operator can expect overheating to take place in the seconds to minutes range [2], [3], [1]. Additionally, the protection deadline adds more time before lines are disconnected [6], [9] . Below we test different parameter settings in order to represent different cascading scenarios and timescales.

C. Line status after control

In Figure 3, we test the scalability of the proposed control using the IEEE test cases. We evaluate how grid size and topology affect the control performance, and use the ratio between the number of “hot” lines in F and the number of lines before the initial failure, F_{normal}.

We evaluate the effect of α which in the cascade model described represents, in a simplified way, the physical properties of the power lines. In practice, along with the conductor material, weather can affect greatly the physical properties of a power line. This effect can be captured by α in the model of Algorithm 1. Thus, α directly affects how the power lines physically respond, in terms of heat, to the variation of power flows induced by failures and controls. Such a response determines the operating condition of the power lines and whether a line becomes “hot”. The proposed control does not assume that the operator knows the exact thermal operating points of the power lines. Thus, the moving average used in Algorithm 1 allows us to evaluate the effectiveness of control under different thermal operating conditions without the physical details of the power lines. From the scenario where the power line operation point accumulates a large amount of heat due to the control actions taken (low α) to the scenario where control actions quickly modify the operation

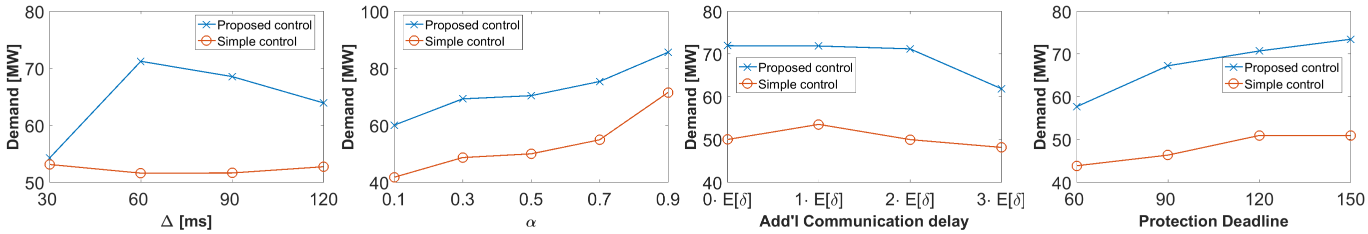

(a) Effect of Δ (b) Effect of α (c) Effect of delays (d) Effect of protection

Fig. 4. Total demand supplied at the end of cascade.

point towards the desired line status (high α). The study of monitoring line temperatures and weather conditions that influence α is out of the scope of this paper. However, if such parameters are available at the control center, the operator can process that information to obtain an equivalent α and incorporate it into the power line constraint (4d).

Without any control in place, $|F|$ consistently grows with the grid size and is nearly independent of α. In this case, the cascade evolves freely and there are no control actions to further modify the status according to the α characteristic of the power lines. For the controlled cases, $|F|$ also grows with the grid size but at a much smaller scale. In figures 3a-3c, our scheme can provide a consistent reduction of the number of power line failures at the end of control by about a half compared to that of simple control. This demonstrates that our communication-aware control scheme can more effectively mitigate the cascade by reducing the number of lines working over their thermal limits. For the larger grid in figure 3d, however, our control scheme can provide only about 15% reduction of the size of $|F|$ compared to the simple control. Thus the effectiveness of control reduces when more nodes are affected by large and variable delays. This suggests that operators should consider applying networked control of power line failures in a distributed fashion with a large grid divided into different power management domains as often done in practical systems.

D. Grid status at the end of cascade

We further test the effect of several parameters on the final state the grid reaches after the cascading failure stops. Following the previous results, we consider that the power operator performs failure control over a small grid areas where delays are shorter and less variable. In this case, the control is expected to be more effective and our algorithm plays less role, so our evaluation is more conservative. Thus, for this section tests are conducted over the 30-node IEEE grid and we evaluate the amount of demand left to support the customers at the end of cascade.

1) Effect of Δ: The parameter Δ in (3) sets the delay "threshold" that determines the amount of power a node has to change for failure control. In our policy, nodes that experience delays beyond Δ will perform smaller changes. In this test, the delays just vary naturally according to the grid size and topology of the communication network. The average communication delay is 30ms.

As shown in Figure 4a, setting Δ to double the average delay provides the best performance. To recover the grid from failure, our policy controls nodes whose delays may drift away from the average but are also not extremely large. Without considering the status of the communication network, the grid performance is significantly compromised under the simple control. It may make large power changes to nodes with delays exceeding Δ, which aggravates the cascade failure. We set $\Delta = 2 \cdot E[\delta]$ in remaining tests.

2) Effect of α: In Figure 4b, as expected, the total demand increases as the heat effects diminish, i.e. with a large α. By considering the effect of delays and consequently the effect of the changes posed by control, our proposed scheme can provide about 50% more demand compared to the simple control.

3) Effect of communication delays: In previous evaluations, communication delays are only related to the topology and size of the grid. To test the robustness of control to adverse communication scenarios, we introduce additional delays to 10 randomly selected nodes. Their delays can be as large as $3 \cdot E[\delta]$, which increases the delay variability. This can happen due to events or cyber attacks to targeted nodes.

Figure 4c shows that our proposed control consistently provides 50% more demand to customers at the end of the cascade compared to simple control strategies. Simple control does not take into account delay conditions and the figure shows that its performance starts to degrade quickly after the additional delay is larger than $E[\delta]$. On average, such additional delays start to approach 100ms, the default protection deadline. Hence, failures of overloaded lines occur before all control packets are delivered to their destinations. When a delayed control message reaches its target node, the grid has already undergone line disconnections and the control is no longer valid, as it was not planned for such a changed grid. Our proposed control does not make the overly optimistic assumption that the control center knows beforehand the exact delay that each control packet will experience. Instead, it introduces an estimation of the delays in the objective function to reduce the number of nodes required to perform large shedding. In this way, our control can avoid performing large (critical) shedding from nodes experiencing large delays.

4) Effect of protection delays: In Figure 4d we evaluate the effect of varying the protection deadline, beyond which overloaded power lines are disconnected. It is usually set by the operator. The control center usually plans the control and transmits control packets after the initial failure. Thus, the

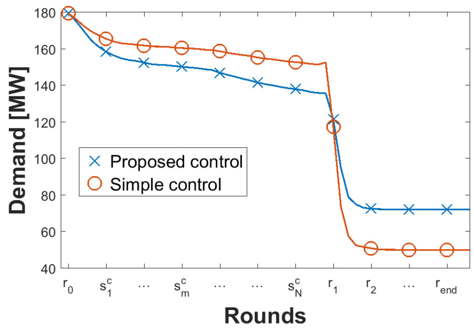

Fig. 5. Cascade evolution over time in terms of demand supplied to customers.

control is planned to be executed according to the nodes and lines that are active following the initial failure. The control center can be unaware of overloaded power lines that get disconnected after the protection deadline expires.

Ideally, all control messages should be transmitted and executed before the protections are triggered to disconnect the corresponding power lines. Delayed control packets delivered after lines are disconnected by protections will be executed in a grid that has changed already, leading the grid to an unplanned state that potentially worsens the original problem.

As it can be seen in the figure, when the protection deadline is as short as 60ms, the control performance is compromised. However, our proposed control still supplies about 30% more total demand at the end of the cascade than simple control. This is because our proposed control is aware of the communication delays and performs most control actions using as many nodes with low delay as possible. The figure also shows that for the case of more relaxed protection deadlines, i.e. $>$ 90ms, the simple control can improve its performance as delayed packets have more time to arrive before the occurrence of new failures. However, such an improvement is not comparable with that of our control, which provides the usual 50% larger demand to customers.

E. Cascade evolution over time

Finally, Figure 5 shows the temporal evolution of the cascade using the default parameters specified in the simulation setup. From the initial failure round until the cascade stops, the figure provides an insight on how the proposed algorithm mitigates the cascade in terms of the total demand supplied to customers. The x-axis corresponds to the cascade rounds and it has been quantized using equal spacing between rounds for illustration purposes. Following the cascade model of Algorithm 1, the grid states corresponding to the arrival of delayed packets are marked with s^c_m, and r_i marks the states reached by the grid after all control actions have been applied, i.e. all r_n in Algorithm 1 after s^c_{max}. The figure shows that our proposed control performs larger power changes (i.e. load shedding) faster and over nodes with low delays, which causes the reduction of the demand available for customers initially. By the end of the first round of control, the amount of load shedding performed by our scheme is about 15% larger than a simple control, providing less available demand. After all control packets have been received at s^c_N, the power operator takes no further actions and expects the shedding performed to be effective in stopping the cascade. However, without further planned actions, the overloaded lines at s^c_N fail and automatic load shedding takes place aiming to regain power balance by r_1. Following the simple cascade model, new rounds past r_1 occur with further shedding being performed automatically, as illustrated by the steep decay of demand in the figure, until no line is overloaded. Given that our control provides a reduced number of line overloads, the grid status following r_1 is less likely to perform large amounts of automatic shedding while the grid following the simple control eventually provides less demand when the cascade stops. Thus, our control scheme requires slightly more load shedding early on the evolution of the cascade in order to make the grid more quickly recover to the balance state to support more demand at the end.

VII. Conclusion

In this paper, we study the effect of communication delay on the effectiveness of controlling cascading failures. We illustrate the problems associated with communication-dependent control and the issues of peer schemes (referred as "simple control") that ignore such dependency. We propose a new failure model to capture the impact of communication delay on cascade failures in power grid. Built upon the model, we further propose a control mechanism to effectively mitigate the cascade failure using information from the states of both the power grid and the communication network. Our simulation results show that the communication delay has different effects on the development of cascading failures in different scenarios and power grid conditions. Specifically, we investigate the effect of control on the power lines operation point, the effect of adverse communication delays, and different protection trigger time, as well as the effects of different network sizes and topologies. Our evaluation results demonstrate that our algorithm can consistently achieve much better performance compared to that of peer works in mitigating cascading failures.

References

[1] A. Bernstein, D. Bienstock, D. Hay, M. Uzunoglu, and G. Zussman, "Power grid vulnerability to geographically correlated failures; analysis and control implications," in *INFOCOM, 2014 Proceedings IEEE*, April 2014, pp. 2634–2642.

[2] J. Qi, S. Mei, and F. Liu, "Blackout model considering slow process," *IEEE Transactions on Power Systems*, vol. 28, no. 3, pp. 3274–3282, Aug 2013.

[3] R. Yao, S. Huang, K. Sun, F. Liu, X. Zhang, and S. Mei, "A multi-timescale quasi-dynamic model for simulation of cascading outages," *IEEE Transactions on Power Systems*, vol. 31, no. 4, pp. 3189–3201, July 2016.

[4] D. Bienstock, "Optimal control of cascading power grid failures," in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, Dec 2011, pp. 2166–2173.

[5] M. Chenine and L. Nordstrom, "Modeling and simulation of wide-area communication for centralized pmu-based applications," *IEEE Transactions on Power Delivery*, vol. 26, no. 3, pp. 1372–1380, July 2011.

[6] Y. Deng, H. Lin, A. G. Phadke, S. Shukla, J. S. Thorp, and L. Mili, "Communication network modeling and simulation for wide area measurement applications," in *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*, Jan 2012, pp. 1–6.

[7] J. Stahlhut, T. Browne, G. Heydt, and V. Vittal, "Latency viewed as a stochastic process and its impact on wide area power system control signals," *Power Systems, IEEE Transactions on*, vol. 23, no. 1, pp. 84–91, Feb 2008.

[8] K. Zhu, L. Nordstrm, and A. T. Al-Hammouri, "Examination of data delay and packet loss for wide-area monitoring and control systems," in *2012 IEEE International Energy Conference and Exhibition (ENERGYCON)*, Sept 2012, pp. 927–934.

[9] C. Huang, F. Li, T. Ding, Y. Jiang, J. Guo, and Y. Liu, "A bounded model of the communication delay for system integrity protection schemes," *IEEE Transactions on Power Delivery*, vol. 31, no. 4, pp. 1921–1933, Aug 2016.

[10] F. Zhang, L. Cheng, X. Li, M. Zeng, Y. Sun, and J. H. Chow, "Estimation and measurement of closed-loop delays in the actual wacs of guizhou power grid," in *2016 IEEE Power and Energy Society General Meeting (PESGM)*, July 2016, pp. 1–5.

[11] F. Zhang, Y. Sun, L. Cheng, X. Li, J. H. Chow, and W. Zhao, "Measurement and modeling of delays in wide-area closed-loop control systems," *IEEE Transactions on Power Systems*, vol. 30, no. 5, pp. 2426–2433, Sept 2015.

[12] J. Zare, F. Aminifar, and M. Sanaye-Pasand, "Communication-constrained regionalization of power systems for synchrophasor-based wide-area backup protection scheme," *IEEE Transactions on Smart Grid*, vol. 6, no. 3, pp. 1530–1538, May 2015.

[13] Y. Wang, P. Yemula, and A. Bose, "Decentralized communication and control systems for power system operation," *IEEE Transactions on Smart Grid*, vol. 6, no. 2, pp. 885–893, March 2015.

[14] M. Wei, Z. Lu, and W. Wang, "Dominoes with communications: On characterizing the progress of cascading failures in smart grid," in *2016 IEEE International Conference on Communications (ICC)*, May 2016, pp. 1–6.

[15] R. Yao, S. Huang, K. Sun, F. Liu, X. Zhang, S. Mei, W. Wei, and L. Ding, "Risk assessment of multi-timescale cascading outages based on markovian tree search," *IEEE Transactions on Power Systems*, vol. 32, no. 4, pp. 2887–2900, July 2017.

[16] A. Sen, A. Mazumder, J. Banerjee, A. Das, and R. Compton, "Identification of $\cal K$ Most Vulnerable Nodes in Multi-layered Network Using a New Model of Interdependency," *ArXiv e-prints*, Jan. 2014.

[17] S. Neumayer and E. Modiano, "Assessing the effect of geographically correlated failures on interconnected power-communication networks," in *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*, Oct 2013, pp. 366–371.

[18] E. M. M. Parandehgheibi and D. Hay, "Mitigating cascading failures in interdependent power grids and communication networks," in *SmartGridComm, Communications and Networks Symposium. IEEE*, November 2014.

[19] M. Rahnamay-Naeini and M. Hayat, "On the role of power-grid and communication-system interdependencies on cascading failures," in *Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE*, Dec 2013, pp. 527–530.

[20] K. Purchala, L. Meeus, D. Van Dommelen, and R. Belmans, "Usefulness of dc power flow for active power flow analysis," in *Power Engineering Society General Meeting, 2005. IEEE*, June 2005, pp. 454–459 Vol. 1.

[21] A. R. Bergen, *Power systems analysis*, ser. Prentice-Hall series in electrical and computer engineering. Englewood Cliffs (N.J.): Prentice Hall, 1986. [Online]. Available: http://opac.inria.fr/record=b1100177