Comparing dynamics of cascading failures between network-centric and power flow models

Valentina Cupac,[1] Joseph T. Lizier,[2,1,*] and Mikhail Prokopenko[1]

[1] *CSIRO Information and Communication Technologies Centre, PO Box 76, Epping, NSW 1710, Australia*

[2] *Max Planck Institute for Mathematics in the Sciences, Inselstraße 22, 04103 Leipzig, Germany*

(Dated: November 27, 2012)

Small initial failures in power grids can lead to large cascading failures, which have significant economic and social costs, and it is imperative to understand the dynamics of these systems and to mitigate risks of such failures. We directly compare two prominent models for cascading failures in power grids: a power flow model — ORNL-PSerc-Alaska (OPA) and a complex networks model — Crucitti-Latora-Marchiori (CLM). Quantitative comparison of these two models is not trivial and has not been previously performed, so we present a method to quantitatively compare the two using transmission capacity as a common variable. Primarily, we find that the two models exhibit similar phase transitions in average network damage (load shed/demand in OPA, path damage in CLM) with respect to transmission capacity, sharing a common critical region and similar transitions in probability distributions of network damage size. Furthermore, we find that both OPA and CLM reveal similar impacts of network topology, size and heterogeneity of transmission capacity, with respect to vulnerability to large cascades. Thus, our analysis indicates that the CLM model, despite neglecting realistic power flow assumptions and exhibiting differences in behaviour at the local scale, nonetheless exhibits ensemble properties which are consistent with the more realistic OPA fast-scale model. Given the advantages of simplicity and scalability of CLM, these results provide impetus for the use of complex networks models to study ensemble properties of cascading failures in larger power grid networks.

Keywords: cascading failures, vulnerability of infrastructure, power grids, OPA model, complex networks

I. INTRODUCTION

Complex networks are susceptible to cascading failures, whereby small initial disturbances may propagate through the system causing significant damage [1–4]. Cascading failures pose a significant risk to power transmission systems. When a line outages, power flow becomes shifted onto the remaining lines. Consequently, the lines carrying extra flow may become overloaded themselves, forcing further redistribution to other lines in turn. The spread of the cascading failure may cause large blackouts across the power grid [2]. Since blackouts have significant social and economic costs, it is imperative to model the complex dynamics of cascading failures in power grids, and to undertake preventive or corrective measures in mitigating the risk of large blackouts [5–7].

A major challenge lies in developing models that are (i) *realistic*, adhering to basic power flow and network constraints, (ii) *scalable*, mapping to large topologies of real power grids, and (iii) *computationally efficient*, permitting real-time applications of the model. These trade-offs are apparent when considering two prominent classes of cascading failure models: *power flow* models, such as ORNL-PSerc-Alaska (OPA) [8–10], and *complex network theory* models, such as Crucitti-Latora-Marchiori (CLM) model [1, 11]. These models are described in detail in Section II.

* joseph.lizier@csiro.au; phone: +61 (0) 2 9372 4711

OPA offers a more faithful representation of power systems as it explicitly incorporates the DC power flow equations and minimisation of generation cost and load shed [9]. However, this also results in notable computational limitations, whereby use of OPA is limited to networks with relatively smaller number of nodes and requires considerable time for solving constrained linear optimisation functions [2]. Furthermore, OPA requires a more involved process in setting parameters and has more parameters compared to CLM.

CLM is a simpler and more abstract model, designed to study common features of flows in many systems [1], e.g. transportation flows, internet traffic flows, and power flows. It offers advantages in modelling cascading failures with fewer parameters (simplicity) and can be more readily applied to larger networks (scalability). On the other hand, CLM neglects basic power flow constraints and is unable to directly represent generator capacity and load demand. Furthermore, it does not provide direct physical measures of blackout size (load shed), but rather resorts to abstract measures (path damage).

Thus, OPA is able to offer a more realistic model at the expense of scalability, whilst CLM is able to offer a more scalable and simpler abstract model at the expense of neglecting basic power flow network constraints.

Whilst CLM provides theoretically interesting applications of graph theory in modelling cascading failures, it is not immediately clear whether or not it has adequate practical significance in specifically modelling power transmission systems due to its abstract nature.

A critical issue lies in model validation, the extent to which abstract variables within CLM are meaningful and translatable to more realistic measures used within OPA. While CLM is certainly a scalable and computationally efficient model, it is necessary to ascertain whether it provides valid measures of damage resulting from cascading failures.

Most previous research in this area typically focused on single models [1, 10–12], with a lack of cross-validation and in-depth comparative analysis. Some studies [2, 5] have provided *qualitative* comparison of these models — identifying similarities and differences, and evaluating advantages and disadvantages. We provide a detailed qualitative comparison in Section III. Crucially though, there has been no rigorous attempt to directly *quantitatively* compare the behaviour of CLM with respect to well-established power flow models such as OPA, using a set of common network topologies and comparable transmission capacities.

Consequently, our primary motivation is to directly compare the critical behaviour and resultant distributions of power grid damage in CLM versus OPA, and to evaluate the extent to which abstract variables and measures within CLM are meaningful in the context of more well-understood measures within OPA.

Such a comparison is not straightforward due to differences in representing system capacity, differences in iterative algorithms they rely on, and differences in measures of system damage. In Section IV, we describe how we will set parameters and measure responses to damage in each model to ensure they are comparable.

The scope of this paper is restricted to studying the OPA fast dynamics model which is analogous in temporal horizons to the CLM model, and ignore the OPA slow dynamics model with does not have a CLM counterpart. Furthermore, we study transmission capacity which is defined for both CLM and OPA, and ignore generation capacity which is defined within OPA but not within CLM.

Certainly, OPA and CLM are inconsistent at the local scale, and we show that load shed/damage in OPA is poorly correlated with path damage in CLM when triggered by the same simulated individual failure (see Appendix A). Nonetheless, we seek to examine whether there are similarities in global or ensemble behaviour at the network level (e.g. susceptibility to large cascades as a function of transmission capacity). Specifically, we aim to evaluate the extent to which CLM is consistent with OPA regarding whether CLM and OPA exhibit similar:

1. phase transitions in average network damage and probability distributions of network damage, with respect to transmission capacity;

2. ranking of grid topologies and grid sizes, based on their susceptibility to large cascading failures; and

3. shifts in probability distributions of failure, with respect to heterogeneity in transmission capacity across the network.

The analysis of probability distributions of resultant failures and damage is crucial to our paper. In the case of OPA, we examine the probability distribution of load shed/demand. In the case of CLM, we examine the probability distribution of efficiency loss (i.e. damage). We then analyse in Section V the extent to which the two models exhibit similarities in their distributions, and the extent to which they exhibit similarities in changes within distributions arising from changes in transmission capacity. This is motivated by the hypothesis that even though CLM and OPA may model power flow dynamics very differently, they may agree in terms of the resultant high-level ensemble properties and probability distributions of cascading failures.

II. MODELS OF CASCADING FAILURES

Modelling cascading failures poses a number of challenges due to the diversity of mechanisms which are involved in initial failures and in subsequent interactions within the power system [5]. These challenges have motivated the development of various cascading failure models. These models can be classified into several broad categories [2, 13]: artificial *power flow*, conventional reliability, and *complex network theory* approaches.

Artificial power flow approaches seek to provide realistic power flow models, and examine the changes in flows as a result of initial failures. OPA, which as described in Section II A is the most prominent of these, considers the dynamics of the flows in time on a given underlying network structure. Another power flow model is the CASCADE model [12], though it is considered "too simple" in that it "disregards the system structure, neglects the times between adjacent failures and generation adaptation during failure" [13].

Conventional reliability approaches, including the FTA model [14], provide simple insights to end-users, but typically require models specialized to individual networks and events, and complex analysis which is difficult to verify and construct [13].

Complex network theory approaches including the CLM model (as described in Section II B) formulate more abstract models, using the tools of graph theory to obtain general insights incorporating network structure in an efficient manner. [1]

In seeking to evaluate how accurately the CLM model performs on a given network structure, we compare it to the OPA model, since this is the most prominent of the power flow models, and is specifically designed to consider the dynamics of flows in time and the underlying

[1] Separate to considerations of *modelling* cascading failures, other complex systems approaches to studying cascading failures include evaluation of the goodness of fit of failure event sizes to power-laws [15], and analysing the dynamics of cascading failure events using state transition graphs [16].

network structure. We describe these models in detail in the following section.

A. OPA model

ORNL-PSerc-Alaska (OPA) is a model proposed by researchers at Oak Ridge National Laboratory (ORNL), Power System Engineering Research Center of Wisconsin University (PSerc), and Alaska University (Alaska) [8–10]. The power grid is represented as a set of nodes with power injections (P_i) which are positive for generators (G) and negative for loads (L). Generators have limited power supply (P_i^{max}) and loads have fixed demand (P_i^{dem}). A transmission line connecting nodes i and j has a power flow F_{ij}, which is limited by the maximum power flow F_{ij}^{max} that it can carry. A transmission line connecting nodes i and j has susceptance b_{ij}, impedance z_{ij}, and resistance is neglected.

The OPA model considers dynamics in terms of two interdependent time scales. The *slow* time scale describes how the system evolves as demand grows over longer timeframes, and the subsequent network upgrades which occur in response to demand growth and blackouts. The *fast* time scale, on the other hand, describes cascading failures of transmission lines at a fixed point during the slow dynamics. Our comparative study focuses on the fast dynamics, which are described below.

Power dispatch is formulated as a linear optimisation problem, where the goal is to minimise a cost function[2] which gives preference to generation shift whilst assigning a high cost to ($W = 100$) to load shedding (reduction of power supply to consumers). It is assumed that all generators operate at the same cost and that all loads are served with equal priority:

$$Cost = \sum_{i \in G} P_i(t) + W \sum_{j \in L} P_j(t). \tag{1}$$

The optimisation is subject to the following constraints. Generators always have positive power injections and cannot exceed generator capacity limits:

$$0 \leq P_i \leq P_i^{max}, \quad i \in G. \tag{2}$$

Loads always have negative power injections (cannot generate power)[3]:

$$P_j^{dem} \leq P_j \leq 0, \quad j \in L. \tag{3}$$

The absolute power flow along edges is subject to power flow limits:

$$-F_{ij}^{max} \leq F_{ij} \leq F_{ij}^{max}, \quad i, j \in N. \tag{4}$$

Total power generation and consumption remain balanced:

$$\sum_{i \in G \cup L} P_i = 0. \tag{5}$$

Based on standard DC power flow analysis, power flows are expressed as:

$$F = AP, \tag{6}$$

where F is a vector of power flows along the edges, P is a vector of power flows injected at each node P_i, and A is a matrix that depends on the network structure and impedances (see [8] for details about the computation of A).

After solving the linear optimisation problem, we examine which lines are overloaded. A line is considered to be *overloaded* if F_{ij} is within 1% of F_{ij}^{max}. Each overloaded line may *outage* with probability p_1. When line outage occurs, its power flow limit F_{ij}^{max} is divided by a very large number (κ_2) to ensure that practically no power may flow through the line. Furthermore, to avoid a matrix singularity resulting from the line outage, line impedance is multiplied by a large number (κ_1), resulting in changes to the network matrix A. The updated set of line flow limits and the updated network matrix lead to the recomputation of the optimisated state and, once again, some lines may outage, leading to further iterations. A cascade is represented as a series of probabilistic line failures and the subsequent power flow redistribution. The process terminates when a solution is found with no line outages.

During the cascade process, load shed occurs when lines are outaged due to an imbalance between supply and demand. Load shed (S_i) for an individual node i is defined as the difference between its power injection and demand, and is always non-negative:

$$S_i = |P_i^{dem} - P_i|. \tag{7}$$

Consequently, total load shed (S) for the system is:

$$S = \sum_{i \in L} S_i. \tag{8}$$

Finally, total load shed is normalised with respect to total demand D, and used as a measure of damage to the system resulting from the cascade:

$$S/D = \frac{\sum_{i \in L} S_i}{\sum_{i \in L} P_i^{dem}}. \tag{9}$$

A crucial result arising from OPA simulations is the presence of scale-free (power-law) regions in the distribution of blackout size, consistent with historical blackout data for real power systems [10]. In addition, the observed power law distribution of blackout size provided evidence for self-organised criticality in power systems.

OPA is able to incorporate basic elements of power system components, power flows, and the relations between

[2] There was a typographical error in the optimisation function written in [9, 10], where there was an incorrect minus sign in the optimisation function, which should have been a plus sign instead [17].

[3] Papers including [9, 10] omitted the lower bound for Eq. 3 [17].

supply and demand. However, the applications of OPA have generally been limited to networks with a relatively small number of nodes (typically hundreds), compared to real power grids [2].

B. CLM model

Complex network theory has been used to formulate more abstract models, where the power grid is represented as an undirected graph with nodes (substations) and edges (transmission lines). A prominent approach is the CLM model [1], including extensions to differentiate generators and loads [11]. The power grid is represented as a set of N nodes representing N_G generators and N_L loads representing distribution substations, interconnected by a set of edges representing transmission lines.

An edge connecting nodes i and j is characterised by edge efficiency $e_{ij}(t)$, representing relative link capacity. Initially we assume that edges are perfectly efficient $e_{ij}(0) = 1$, although the efficiency is subject to change over consequent iterations. At any time step, edge weight is computed as the inverse of edge efficiency. It is assumed that electricity flowing between any two nodes i and j takes the most efficient path connecting those two nodes. The most efficient path is the shortest path between i and j, where the length of a given path is defined as the sum of edge weights on that path in the graph between i and j. The efficiency $\epsilon_{ij}(t)$ of the most efficient path from i to j is the inverse of the shortest path length [18].

Each node is characterised by a *loading* $L_i(t)$[4], which varies over time and a fixed limited capacity C_i. The loading of each node with transmitting capabilities is defined as the number of most efficient paths from generators to distribution substations that pass through that node [19]. (In the original model [1] which considered paths between *all node pairs*, the loading was equivalent to *betweenness centrality* [19].) Consequently, the more shortest paths that pass through a node, the higher the loading it has. The loading at a given node amounts to the total flow passing through that node when all generator/distribution substation pairs transmit energy between them [19].

Node capacity C_i represents the maximum loading that it can handle without performance degradation. The capacity is assumed to be directly proportional to the initial loading of the node [20]. The fixed network tolerance parameter $\alpha \geq 1$ represents the ability of nodes to handle increased loading thereby resisting perturbations:

$$C_i = \alpha L_i(0). \tag{10}$$

An initial breakdown of edges surrounding a node causes power to be redistributed in the network, reflected by changes in the most efficient paths and, consequently, changes in the loading at each node. Some nodes are then forced to operate above capacity (being overloaded), represented by decreases in efficiency of the edges of that node (slightly altering the definition of [1], as per [21]):

$$e_{ij}(t+1) = \begin{cases} e_{ij}(0)\min(\frac{C_i}{L_i(t)}, \frac{C_j}{L_j(t)}) & \text{if } L_i(t) > C_i \\ & \text{or } L_j(t) > C_j \\ e_{ij}(0) & \text{otherwise.} \end{cases} \tag{11}$$

The changes in the efficiency of these edges may change the most efficient paths between node pairs, and indeed changes in most efficient paths alter the loading across the network. Convergence in this iterative process occurs when the state dynamics become stable or periodic.

Network efficiency is computed as the average path efficiency of the most efficient path from each generator to each load [11][5]:

$$E(t) = \frac{1}{N_G N_L} \sum_{i \in G} \sum_{j \in L} \epsilon_{ij}(t). \tag{12}$$

Network damage D caused by the cascade is defined as the normalised loss of average path efficiency [11]:

$$D(t) = \frac{E(0) - E(t)}{E(0)}. \tag{13}$$

CLM can specifically examine the effect of network structure on cascades and their consequences (for example, random graphs appear to be more resistant to cascading failures than scale-free graphs [1]). Results based on the North American power network indicate that the critical tolerance levels for loading-based node removal were substantially higher than the critical tolerance levels for random node removals [11]. Furthermore, it was shown that nodes can be divided into different classes: the removal of some nodes causes very little damage for any tolerance levels, whilst other nodes' removal leads to an efficiency damage which is a function of their tolerance level.

A key advantage of complex network models, such as CLM, is that they enable us to utilise techniques from graph theory and perform simulations on large networks due to their scalability; for example, Kinney et al. [11] have demonstrated the applicability of CLM to *at least* tens of thousands of nodes (an increase of two orders of

[4] CLM papers typically use the term *load* to qualitatively describe the amount of power flowing through a node. At the same time, it can be used to signify node type, whereby a *load* node represents a distribution substation. Thus, to avoid potential ambiguity, we use the term *loading* for the former case.

[5] Note that the original formulation in [1] considered most efficient paths between all node pairs, failing to differentiate between generators and loads.

magnitude over the capability of OPA [2]), though there are challenges in reflecting particular features of power grid operations. Nonetheless, such an approach provides important insights and has gained increasing applications in modelling cascading failures in power grids.

III. QUALITATIVE COMPARISON

Since OPA and CLM rely on different variables and algorithms outlined in the previous section, we present here a general qualitative comparison between these two models, followed by a quantitative comparative analysis in the next section.

A. Demand and generation

The two models differ in the extent to which they are able to represent demand and generator capacity. OPA provides direct representation of load demand and generator capacity. For example, when mapping any network, it enables us to specify the load demand of individual distribution substations, and to specify the capacity of individual generation substations. On the other hand, CLM does not quantitatively represent demand and generation capacity (although it does quantitatively represent transmission capacity). OPA allows variations in the total demand / generator capacity ratios. CLM does not allow such variation, total demand and total generator capacity are implicitly equal since every generator is connected to every load and vice versa.

Furthermore, CLM assumes that every distribution substation is connected to every generator, whereby there is only one shortest path from any distribution substation and every generator. This implies that every distribution substation attempts to extract an equal amount of power from every generator. OPA does not have such a simplistic restriction, a distribution substation may be receiving power from just *some* generators (rather than from *all* generators, as is the case in CLM) and furthermore, it may disproportionately receive more power from some generators than others (rather than attempting to connect equally to all generators, as is the case in CLM). In CLM, every generator is used at all times, whilst in OPA, some generators may be fully utilised, some partially utilised, and some might not provide any power at all (see Fig. 4. in [9]).

Therefore, OPA is able to represent transitions in relative generation capacity of the power grid, whilst CLM is unable to do so.

B. Transmission

The input variable that is directly represented in both models is transmission capacity, the capacity to carry power. OPA represents transmission capacity as the maximum line flow, which is associated with *edges* in the model. On the other hand, CLM represents transmission capacity as the maximum number of shortest paths, which is associated with *nodes* in the model. In both cases, transmission capacity varies across the system, i.e. in OPA, different lines may have different limits, and in CLM, different nodes may have different limits.

In both models, transmission capacity is the crucial factor in the cascade process, described below. In OPA, *edges* fail when transmission capacity is reached — an edge is overloaded when its power flow reaches its power flow limit. In CLM, *nodes* fail when transmission capacity is exceeded — a node is overloaded when the number of shortest paths exceeds the maximum number of shortest paths that can pass through the node.

Therefore, both OPA and CLM are able to represent transitions in relative transmission capacity of the power grid.

C. Initial failure

We mentioned that edges fail in OPA, whilst nodes fail in CLM, in alignment with the terminology used in the paper which describe the CLM model. It should be noted that in CLM, when we speak of nodes failing, it does not imply that actual substations fail, but rather that the transmission lines immediately attached to those nodes would fail. Consequently, in the implementation of CLM, node failure is actually represented as the failure of its adjacent edges.

D. Transmission paths

A major difference is that in OPA, power flows through multiple paths from one node to another, whilst in CLM, power is considered to flow only through a single (shortest) path. That said, the two models still exhibit some similarities in determining the paths through which power will flow. Edge *impedance* in OPA is related to edge *efficiency* in CLM, and edge *susceptance* in OPA is related to edge *weight* in CLM. When an initial failure occurs, the susceptance and efficiency of affected edges is decreased in OPA and CLM, respectively. This leads to a corresponding increase of those edges' impedances and weights in OPA and CLM, respectively.

Both algorithms iteratively consider updates to the power flows based on changes in impedance and weights, though they do this in different ways. We have the following important differences regarding how the models respond to overloading: (i) magnitude of the change and (ii) permanence of the change. Firstly, in OPA, the changes are more extreme: edge failure leads to setting the susceptance to almost zero and the impedance to a very large number, and the result is that there is theoretically no flow through that edge. On the other hand, in

CLM, edge efficiency is decreased depending on the extent of overloading, hence there may be smaller or larger decreases (unlike OPA, where the susceptance is guaranteed to go to almost zero). Similar principles apply to weights in CLM: they are not necessarily increased to a very large number. Secondly, in OPA, the changes are permanent for the duration of the cascade, whilst in CLM the changes are temporary (efficiency may recover after it is decreased, during a single cascade). It should also be noted that whilst failure is deterministic in CLM (an overloaded node fails with certainty), it is probabilistic in OPA (an overloaded edge fails with probability p_1).

E. Measurement of cascade damage

At the completion of the cascade, both models measure the extent of damage caused by the cascade. OPA primarily measures blackout size using load shed, representing the difference between total power demanded and total power supplied to the distribution substations. Thus, it measures the extent to which consumers failed to receive the power they demanded which occurred either due to line outages or inadequate generator capacity.

CLM is unable to represent load shed because, as mentioned before, it quantifies neither power demand nor power injections which are necessary for the computation of load shed. What it does do is to represent damage by computing the average efficiency reduction. These are the only two measures that we can use to evaluate the extent of damage in both models.

F. Power-law distribution of cascade damage

When analysing cascading failures, an important consideration is the probability distribution of extent of damage, and in particular the tail behaviour. In a normal distribution, failures are symmetrically distributed about a mean, and extreme failures are relatively rare. In a power law distribution, the occurrence of extreme failures is much more probable than in a normal distribution.

Analysis of North American blackouts, based on 15 years of normalised NERC data, indicates the presence of power law tails in the distribution of blackout sizes [9], demonstrating that the probability of large blackouts is higher than what would be expected if blackouts were normally distributed. It is important for models to be able to capture this tail behaviour to ensure that real power systems are adequately modelled.

Studies performed with OPA have shown that the probability distribution of normalised load shed has power law tails — in a remarkable agreement with the North American historical blackout data, indicating that the North American power systems have operated close to a critical point [9]. Furthermore, it was shown with an analytically tractable loading-dependent model of cascading failure that, at a critical loading, a power-law region emerges in the distribution of number of components failed [12].

A study performed using the CLM model applied to the North American power grid found that the cumulative damage distributions with $\alpha = 1.1$ and $\alpha = 1.2$ were in approximate agreement with the power-law distributions found previously using OPA [11].

IV. METHODOLOGY

We seek to compare the OPA and CLM models by simulating cascading failure experiments, defined by the following parameters:

- A model for cascading failures (OPA, CLM). Ensuring comprability between the two models necessitated some adjustments as described in the remainder of this section;
- A network specifying the grid topology (several alternatives are described in subsection IV A);
- A transmission capacity factor (α), representing the average or global transmission capacity for the system; and
- A heterogeneity factor (r), representing local fluctuations in α across the network.

For each set of parameters considered, we ran experiments as follows:

1. Initialise network capacity, which incorporates α and r, as described in Section IV B.
2. Initiate random failure and model the iterative cascade process, as described in Section IV C.
3. Measure the resulting damage caused by the cascading failure, as described in Section IV D.

We studied network damage for 21 levels α ($\alpha = 0.90$ to $\alpha = 1.2$ with increments of 0.02, and $\alpha = 1.2$ to $\alpha = 1.40$ with increments of 0.04). For each of these levels, we ran 1,000 experiments — each with a different randomly-selected point of the initial failure. Furthermore, for any given level, even though the global transmission capacity would be fixed for all experiments performed at that level, local variations would result from the renewed application of the heterogeneity factor r for each experiment, as explained in Section IV B.

We also examined probability distributions for a subset of levels α ($\alpha = 0.80, 1.00, 1.10, 1.20, 1.60, 2.00$). For each of these subset levels, we ran 10,000 experiments to gather enough results to show the probability distribution of network damage.[6]

[6] The number of experiments was selected by qualitatively analyzing a. stability to increase in number of experiments, and b.

A. Network topology

The graph is specified by a set of nodes (a node is either a load or generator) connected by edges (representing transmission lines). In both OPA and CLM case, we constructed a network with nodes which were either generator nodes (representing net generation substations) or load nodes (representing net distribution substations). In our OPA implementation, to facilitate comparability with CLM, all the generators were assumed to have equal capacity, and all the loads were assumed to have equal demand. Furthermore, all edge impedances were set to 1 in OPA, to align with the uniform edge weight initialisation in CLM.

We considered different types of networks, comprising of the standard 300-node IEEE300 bus test case [22] as well as a randomly-generated 300-node scale free network (SCALE300) and a 300-node Erdös-Renyi network (ER300). In constructing IEEE300, nodes with generator buses were classed as generator nodes and all other nodes were classed as load nodes. When constructing SCALE300 and ER300 we ensured that they had the same percentage of generator nodes as IEEE300 (approximately 23%) — for each node, the probability of it being selected as a generator was 0.23 and the probability of being selected as a load was 0.77. The average degree d was relatively similar between these networks ($d = 2.73$ for IEEE300, $d = 2.94$ for SCALE, and $d = 3.23$ for ER300).

We also considered two larger networks, IEEE418 and SCALE400. We constructed "IEEE418" [23] by combining the standard IEEE300 and IEEE118 bus test cases [22]. These two networks were combined by inserting 30 additional edges to connect the two networks, resulting in a combined network with average path efficiency which approximated the weighted average path efficiency of the networks considered separately. We constructed SCALE400 using the same method as generating SCALE300, and ensured that the proportion of generators in SCALE400 was same as in IEEE418.

B. Initialisation of transmission capacity

We used α as a common measure of relative transmission capacity for both networks. The CLM model assumes that $\alpha \geq 1$, representing excess transmission capacity. In our study, we use a more relaxed approach $\alpha \geq 0$ which enables us to consider cases with deficient transmission capacity. This ensured comparability with studies performed in OPA which considered both under- and over-capacity system.

For both models, we computed initial transmission (initial loading in CLM, initial flows in OPA), and applied the α multiplier to compute transmission capacity, which remained fixed for the duration of the experiment.

In CLM, for each node i, initial loading $L_i(0)$ was computed as described in Section II B. To incorporate heterogeneity r, the capacity for node i is initialised as follows, where q is a uniformly distributed random variable:

$$C_i = \alpha q L(0), q \in [1 - r, 1 + r]. \tag{14}$$

Regarding OPA, prior studies set the initial limits (demand, generator capacity, line flow limits) by evolving the network using combined fast-slow dynamics until the network reaches a steady state (stabilisation of average fractional overload) [17]. In order to ensure comparability with CLM, and taking into consideration that we limited the scope of our comparsion to fast dynamics, we used a simpler initialisation strategy which did not require the use of the slow model.

To do this, we make certain assumptions in order to estimate the initial flows, and then set the flow limits based on these estimates. Demand for all load nodes was fixed to a constant amount (we used 100), and total generation capacity was set to be equal to total demand, and equally divided among the generators. Then, we estimated power flows along the lines, on the assumption that every load node would obtain an equal amount of power from every generator. These are aligned with assumptions in the CLM model regarding generation and demand (see Section III A). The initial flow estimations were implemented by selecting a generator (one at a time), setting all other generator capacities to zero (to prevent them from being used by the optimiser) and then computing power flows to each load node. When the power flow magnitudes were summed (resulting in $F^*_{ij}(0)$), they would reflect equal contribution from every generator, analogous to the initialisation process in CLM. The power flow along each line would be multiplied by α and, incorporating heterogeneity r, the maximum capacity for an edge connecting nodes i and j is initialised as follows, where q is a uniformly distributed random variable:

$$F^{max}_{ij} = \alpha q |F^*_{ij}(0)|, q \in [1 - r, 1 + r]. \tag{15}$$

The $F^*_{ij}(0)$ were used only for the purposes of setting the flow limits, and it is important to note that they are neither computed using the full OPA algorithm nor do they feed into the dynamics of the algorithm after the initialisation. They are only used to select the initial flow limits in a manner that is comparable to the limits used by CLM.

We attempt to experimentally verify whether the transmission capacity initialisation schemes applied to CLM and OPA were indeed comparable. The transmission capacity for a node i was computed directly in CLM (C_i). In OPA, there is no direct transmission capacity

smoothness, of the probability distribution function. In particular, we examined these qualities in the tail behaviour. We found that 5,000 samples was suitable, but doubled this to 10,000 as a buffer. For the investigations of average behaviour, we found that 500 samples provided stable results, but again double this to 1,000 as a buffer.

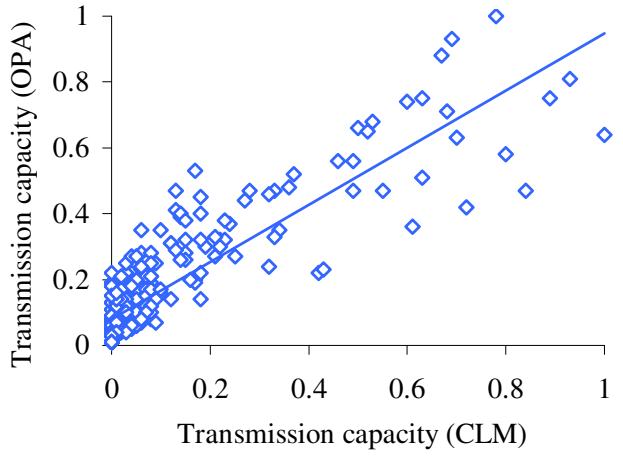

FIG. 1. Scatterplot of normalised node transmission capacities in OPA versus CLM, both using the IEEE300 network, with homogeneous transmission capacity ($r = 0$). Node transmission capacity initialisation in CLM is highly correlated with transmission capacity in OPA ($r_{OPA,CLM} = 0.88$). The line of best fit was added to the plot.

measure for each node i, so we modelled it based on the transmission capacity F_{ij}^{max} of lines connected to it as follows: the transmission capacity for node i was modelled as the sum of maximum power flows of adajcent edges $\sum_{j\in J} F_{ij}^{max}$ where J is the set of nodes directly connected to i via an edge.

In Fig. 1, we plot the transmission capacities determined for each node under both models (with $r = 0$). This verified that transmission capacity in CLM was highly correlated with transmission capacity in OPA (with correlation coefficient 0.88). That is, nodes with low transmission capacity in CLM also tended to have low transmission capacity in OPA, and nodes with high transmission capacity in CLM tended to have high transmission capacity in OPA. This indicates that our initialisation scheme was consistent for CLM and OPA.

For both models, we ran the bulk of our experiments with relatively low heterogeneity factor ($r = 0.1$) which enabled us to ensure that the capacity of individual components did not substantially deviate from the global capacity factor α, whilst being able to represent diversity in cascade events. However, in experiments where we specifically studied the impact of heterogeneity, we studied the effect the behaviour of the network using a higher heterogeneity factor ($r = 0.5$).

C. Cascade algorithm

To ensure comparability with CLM, we ignored the effects of network upgrades over time, and thus applied fast dynamics only.

The cascade was initiated by the failure of a single node in each model. The single node is randomly selected from the set of nodes in the network, where each node is equally likely to be selected. In CLM, the initial single-node failure is represented by the removal of the edges connected to it. In OPA, the initial single-node failure is represented by outaging the edges connected to it.

Then, we applied the algorithms which were defined in Section II. The algorithms ran over several iterations until they either converged or exceeded the maximum number of steps (we used maximum 15 iterations for CLM and OPA). Normal convergence in CLM would be when the state dynamics become stable or periodic, whilst normal convergence in OPA would be when no edge fails in the prior iteration, or if there is no solution to the linear optimisation (which is treated as a blackout [8]).

In OPA, we set the probability of an overloaded edge to outage to $p_1 = 1$, to ensure comparability with CLM, where an overloaded node failed with certainty.

D. Cascade damage

We examined the extent of blackout, or damage to the network that were caused by the different types of node removals. In CLM, we computed path damage, i.e. normalised efficiency loss (previously defined by [11]), which represents damage to the network. In OPA, we computed the load shed/demand ratio, representing extent of blackout. The measurements used in our results are made at the end of the cascade (i.e. at the completion of the simulation algorithm, as described above).

In an abstract sense, we will use the term *network damage* to refer to both cases.

V. RESULTS

It is evident that OPA and CLM provide different results at the local scale (see Appendix A), however we evaluated the extent to which the two models are consistent at the network level. Firstly, we examined whether CLM and OPA exhibit similar characteristics in phase transitions of network damage with respect to transmission capacity (see Section V A). Secondly, we investigated whether they provide a consistent ranking of grid topologies, based on their tolerance to large cascading failures (see Section V B). Thirdly, we verified whether CLM and OPA show consistent behaviour as network size increases (see Section V C). Fourthly, we examined whether network susceptibility to large failure changes consistently with variations in heterogeneity of transmission capacity across the system (see Section V D).

We note again that it is the better computational efficiency and scalability of CLM which motivates us to examine its consistency with OPA at the network level. While we do not specifically investigate scalability here (see [2, 11]), we report that our CLM simulations in this section were generally 10-20 times faster than OPA for analysis of corresponding networks here. If CLM is found to produce consistent results, this will prompt further use

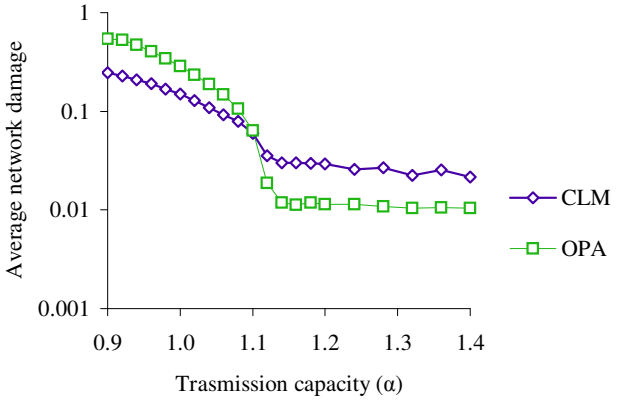

FIG. 2. Phase transition in average network damage (path damage for CLM, load shed/demand for OPA) for a range of α values. Error bars (standard error of the mean) were too small to be visible on this plot.

of CLM to study ensemble properties of cascading failures in larger networks.

A. Phase transitions

The network damage resulting from cascading failures is dependent on the relative capacity of the system. When relative capacity is quite high, a system has high tolerance to failure, whereby a small disturbance would tend to cause relatively low damage on average, with low probability of very high damage. As relative capacity *decreases* and passes a critical capacity level, it increases the probability of large failure. We focus on how the probability distribution of average network damage for IEEE300 changes over different transmission capacity levels α.

We first investigated the change in *average* network damage as a function of relative capacity α. This is the typical analysis used in CLM (e.g. see Fig. 4 [1]) and is akin to the identification of critical demand levels in OPA (e.g. See Fig. 5 in [9]). This comparison between the two models is only possible with the adjustments we introduced. Figure 2 shows that both models exhibit a strongly stable regime for large $\alpha \geq 1.14$ where network damage is small and remains relatively constant as α changes. Additionally, both models exhibit an unstable regime for small $\alpha < 1.10$ where small disturbances can cause large amounts of network damage. Most importantly, this figure shows that both models exhibit a sharp increase in network damage at the same critial α region ($1.10 \leq \alpha \leq 1.12$). That is to say, we have identified very similar phase transitions in the space of α (control parameter), with respect to average path damage or average load shed/demand (order parameter).

The presence of the critical α region is characterised by a rapid increase in load shed/demand as transmission capacity is reduced. This is consistent with prior OPA studies which have identified a second order transition point occurring at a particular power demand, when power flows in transmission lines reach limits set by the transmission capacity of the grid (e.g. see Fig. 5 [9]). The critical phase transition shown for CLM is also consistent with prior studies (e.g. see Fig. 2(b) in [1]).

We then examined the *probability distributions* of magnitude of network damage for particular values of α.[7] This is the typical analysis conducted in OPA studies (e.g. see Fig. 8 in [9]), and while a similar examination has been presented for CLM (e.g. see Fig. 3 in [11]) there has been no quantitative comparison on the same network with a common measure of transmission capacity. Figures 3(a) and 3(b) plot these probability distributions for several α (of course, the expectation values of each probability distribution correspond to the single averages plotted for each α on Fig. 2). The most important result is that both OPA and CLM show the same trends in these probability distributions. These figures indicate that when the relative capacity is in the stable regime ($\alpha \geq 1.2$ for this figure), probability distributions do not change with α and are skewed towards low network damage. This indicates that both models show similar susceptibility to failure when transmission capacity is sufficiently above the critical value and that, in those circumstances, increasing transmission capacity has no impact on overall failure reduction. However, as relative capacity enters the critical region (see $\alpha = 1.1$ here), the probability distribution shifts to the right and the likelihood of large failures increases significantly in both models. The distribution continues shifting to the right as relative capacity is reduced even further (see $\alpha = 0.8, 1.0$ here).

Certainly, the exact averages and distributions of network damage may not be identical between the models: CLM shows larger network damage than OPA in the stable regime and lower network damage in the unstable region, and appears slightly less sensitive to early critical loading effects. Indeed, the two models may not provide the same results on a local level (see Appendix A), so one cannot use the resulting measure of path damage (CLM) from a *specific* individual failure to infer the measure of load shed/demand (OPA) for the same specific failure. However, what is important is that both models exhibit the same network-wide or *ensemble* behaviour with respect to transmission capacity α. Crucially, they identify a common critical region, and probability distributions undergo similar shifts as relative transmission capacity α changes.

[7] Note that the magnitude or "scale" of network damage here is a very different concept to our use of the term "scale" to refer to micro or local scale features of the cascading failures (i.e. specifics of power flows on specific edges due to specific failures) versus macro or network-wide scale features (i.e. probability distributions of failure event sizes with respect to capacity or structure type).

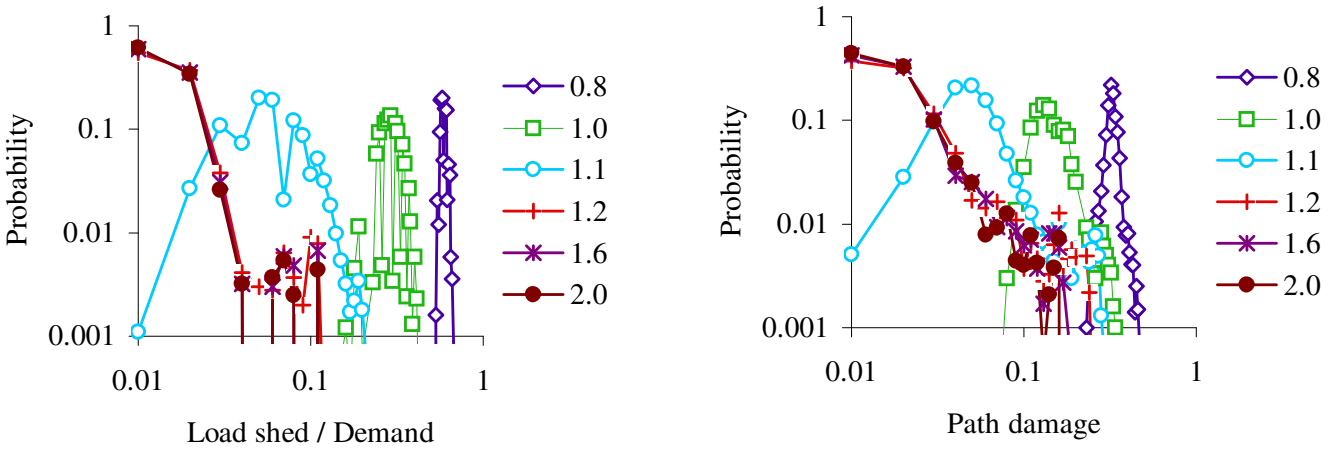

FIG. 3. Probability distribution functions (PDFs) for network damage for IEEE300: (a) PDFs for load shed/demand in OPA; (b) PDFs for path damage in CLM. Different curves correspond to different levels α.

B. Impact of network topology

We examined how network topology affects susceptibility to *large* cascading failures in CLM and OPA. We studied three networks with identical node count, comprising the standard IEEE300 network, a Erdös-Renyi 300-node network (ER300) and a 300-node scale-free network (SCALE300).

We found similar results from OPA and CLM by computing the *maximum* damage caused by the cascade over a range of relative capacity (α) levels. Figures 4(a) and 4(b) illustrate that rankings of topology-dependent cases of high damage were generally preserved over a range of transmission capacity levels, whereby the topology-based differences in ranking were most evident in high capacity ($\alpha \geq 2$) cases. At critical and lower capacities ($\alpha \leq 1.1$), topology appeared to have a less significant impact: CLM showed convergence in maximum path damage across all three networks (see Fig. 4(b)) and OPA showed overlapping of two networks (see Fig. 4(a)).

Figures 5(a) and 5(b) show that network structure *does* impact tail behaviour, at relatively high capacity levels ($\alpha = 2.0$). Both models produced consistent rankings of high-cascade risk associated with network topology, evident in tail behaviour for load shed/demand in OPA, and path damage in CLM.

Thus, our results indicate that network structure has an impact on susceptibility to large cascading failures in both models, being of particular importance for higher transmission capacity levels. Analysis of tail behaviour for OPA and CLM showed that, for an identical node count, Erdös-Renyi graphs are the most tolerant to large failure, and scale-free graphs were most susceptible to large failure. This is because scale-free graphs have large hubs and we know that the network is prone to large cascades when these hubs are affected [24]. Importantly, the standard IEEE300 network showed intermediate tolerance. This is consistent with prior research which shown that the power grid topology is an intermediate between Erdös-Reyni graphs and scale-free graphs [19, 25] (as cited in [11]).

We can confirm that OPA and CLM reveal similar phase transitions in *average* damage with respect to capacity for scale-free graphs and Erdös-Renyi graphs (as we observed for the IEEE300 in Fig. 3), though the transition revealed by CLM for Erdös-Renyi was significantly more broad. Importantly, the critical regions identified by the two models overlapped in each case. We should also note that the correspondence of ranking of different topologies with respect to damage size by OPA and CLM only holds for maximal damage events, not the average damage events.

C. Impact of network size

We investigated the impact of increased network size on susceptibility to large cascades. Firstly, we studied the probability distributions for IEEE300 versus IEEE418. Figures 6(a) and 6(b) show that OPA and CLM generate similar distributions as network size changes from IEEE300 to IEEE418. The larger network exhibited earlier and sharper tail cut-offs, indicating that increasing network size increases network tolerance to large failure.

Similar results were found when comparing the profiles of SCALE300 and SCALE400 (not shown) where the larger network again exhibited sharper tail cut-offs.

Thus, our results indicate that increasing network size (but keeping size of initial failure the same) increases tolerance to large cascades, and importantly, this is detected by both OPA and CLM models. This may be because the larger networks offer more redundant transmission paths which can mitigate failures.

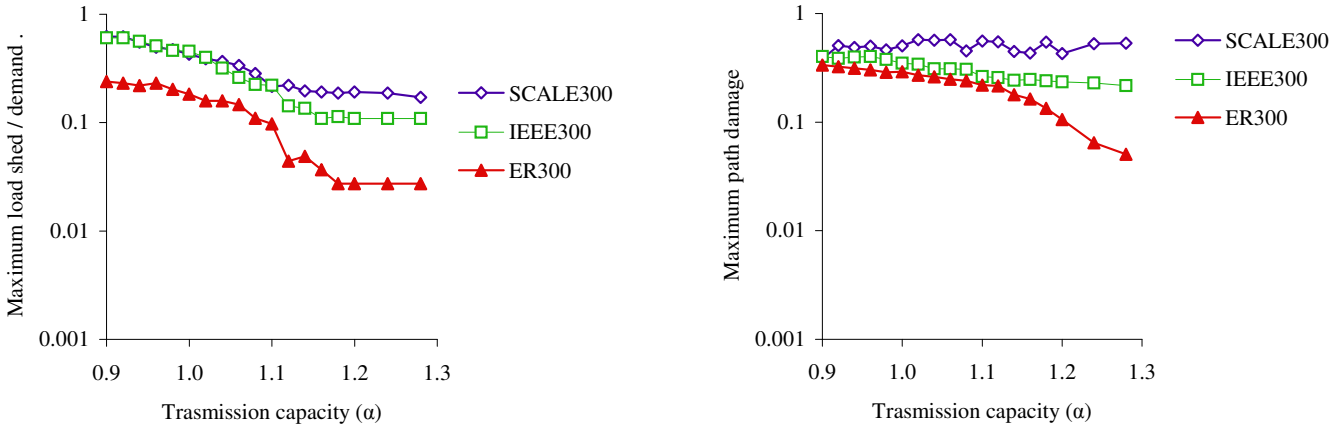

FIG. 4. Maximum observed network damage for SCALE300, IEEE300, ER300: (a) maximum observed load shed/demand in OPA; (b) maximum observed path amage in CLM. Results are shown for a range of levels α.

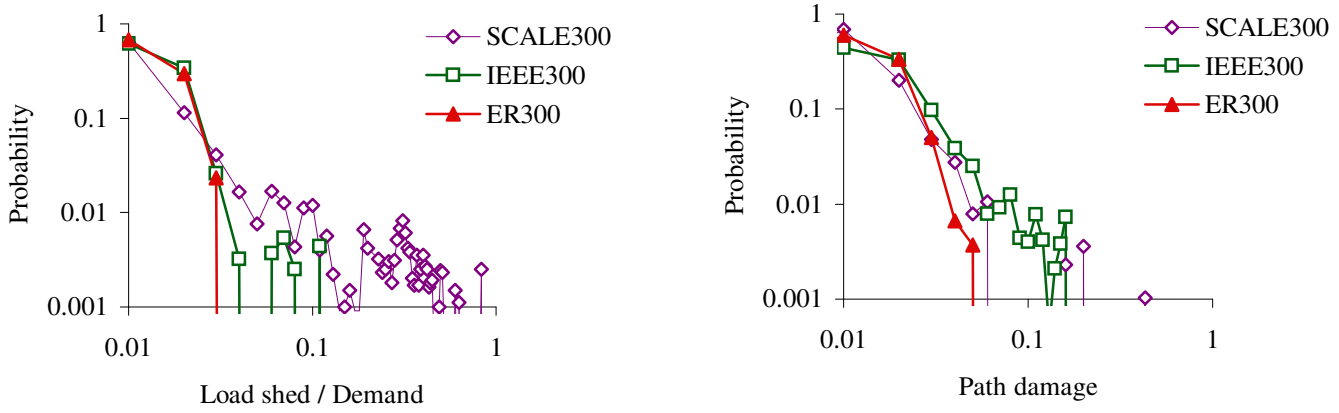

FIG. 5. Probability distribution of network damage for SCALE300, IEEE300, ER300 for $\alpha = 2.0$: (a) probability distribution of load shed/demand in OPA; (b) probability distribution of path damage in CLM.

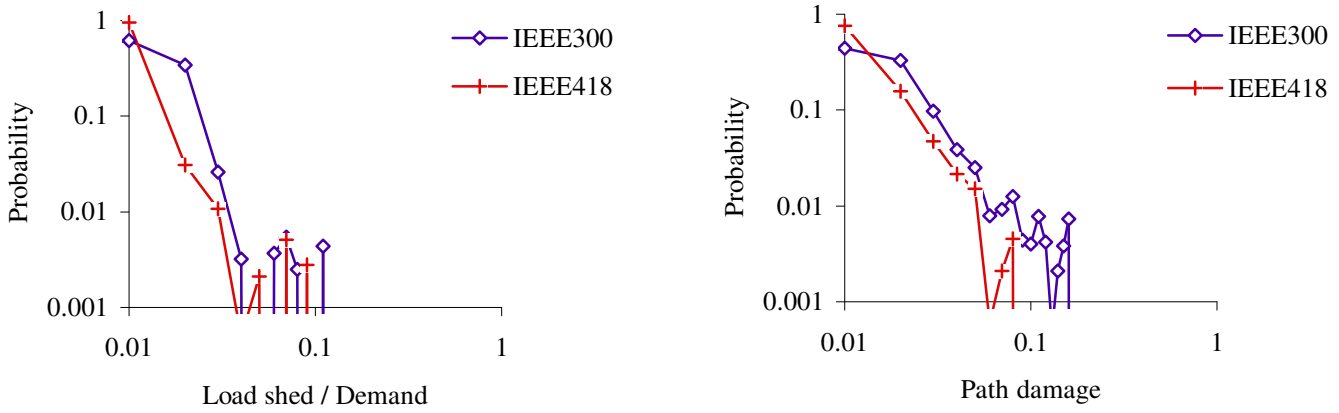

FIG. 6. Probability distribution of network damage for IEEE300 and IEEE418: (a) probability distribution of load shed/demand in OPA; (b) probability distribution of path damage in CLM

D. Impact of heterogeneity in transmission capacity

We examined the extent to which OPA and CLM exhibit consistent behaviour as heterogeneity in transmission capacity increases across the system. Our previous results were computed using $r = 0.1$ as described in Section IV B to represent random local fluctuations in relative transmission capacity. Here we compare those results to an experiment using a larger level of heterogeneity ($r = 0.5$), for a sample transmission capacity $\alpha = 1.2$.

Figures 7(a) and 7(b) indicate that increasing the heterogeneity factor r shifts the probability distribution to larger network damage values, i.e. increasing the probability of larger cascades. We find that certain α values (e.g. $\alpha = 1.2$ here) which are stable for low heteregeneity levels ($r = 0.1$ here, see Section V A) become unstable for higher heteregeneity levels ($r = 0.5$ here). Both OPA and CLM reveal this behaviour with increasing heteregeneity.

The tendency for heterogeneity to increase the probability of larger cascades, thereby shifting the critical region to a higher transmission capacity α level, is consistent with prior OPA studies which examined the effect of random demand fluctuations and found that the level of fluctuations causes the network exhibit criticality even when total demand is below total generator capacity limits (see Section V in [9]).

VI. DISCUSSION AND CONCLUSION

We investigated the behaviour of two approaches to modelling cascading failures: a power flow model — ORNL-PSerc-Alaska (OPA), and a complex networks model — Crucitti-Latora-Marchiori (CLM) model. Our analysis indicates that the CLM model, despite neglecting realistic network constraints and being inconsistent with OPA at the local scale, nonetheless exhibits ensemble or network-level properties which are consistent with the more realistic OPA fast-scale model with respect to transmission capacity α.

Primarly, we found that the OPA and CLM exhibited similar phase transitions in network damage with respect to transmission capacity α: both models exhibited a strongly stable regime for larger α and an unstable regime for smaller α. Crucially, they shared a common critical α region. Furthermore, we found both models exhibited similar probability distribution of network damage, in response to changes in α.

Additionally, we found that OPA and CLM provided consistent rankings of network topologies with respect to susceptibility to large failures (i.e. tail behaviour): both models showed that scale-free networks were most susceptible to large cascades, whilst Erdös-Renyi networks were most resilient to large cascades. Importantly, IEEE300 (a more realistic network) showed intermediate susceptibility to large cascades relative to these reference structures. We also found both models showed that larger networks were more resilient than smaller networks. Furthermore, OPA and CLM exhibited similar transitions in the probability distribution of network damage with respect to changes in heterogeneity in transmission capacity α across the network: increasing heterogeneity increased susceptibility to large cascades, and led to instability even for higher α.

The results from this comparative paper provide an important contribution regarding the usefulness of an abstract model (CLM). Even though CLM is inconsistent with OPA at the local scale, it is nonetheless useful in studying properties of a power system at a global scale and profiling the risk of cascading failures in various network topologies, with the advantages of simplicity and scalability compared to a detailed and more realistic power flow model (OPA). That is to say, despite the modelling power of OPA, it is much less computationally efficient and there are limits to the size of networks it can study [2] in comparison to CLM; as such, our finding that CLM exhibits ensemble properties consistent with OPA is particularly important for the study of larger networks.

Our study provides a starting point for more extensive quantitative validation of abstract models against more realistic ones, and to identify conditions under which they exhibit comparable behaviour. It is important to note that the current comparative study was performed under a specific set of restrictions, constraints and simplifications. Firstly, we introduced α to OPA to enable us to directly compare the behaviour of the two models with respect to a common measure of transmission capacity α. Secondly, we outlined an ad hoc method of using α to initialise transmission capacity of edges in OPA, with the assumption that demand and generation are equal and (only for the purposes of capacity initialisation) that each distribution substation attempts to equally extract power from each generator. We showed that this initialisation scheme ensured that the transmission capacity was comparable between OPA and CLM. Nonetheless, it should be noted that typically OPA would be initialised by evolving the network with both slow- and fast-scale dynamics, and only the latter were considered in our study. Thirdly, the following simplifications were applied to OPA to ensure comparability with CLM for the purposes of this study: nodes were treated as either purely generators or purely loads, load nodes had uniform demand, generator nodes had uniform generation capacity, edges had uniform impedance, overloaded edges failed with certainty.

Future studies may consider extensions to CLM which could enable more comprehensive comparative studies and broaden the scope of applications of abstract models of power systems. Firstly, it may be useful to model not only variations in transmission capacity, but also variations in generation capacity, and to consider situations where generation capacity and demand are not equally distributed, which is aligned with more realistic cases of power grids. One could for example consider supply

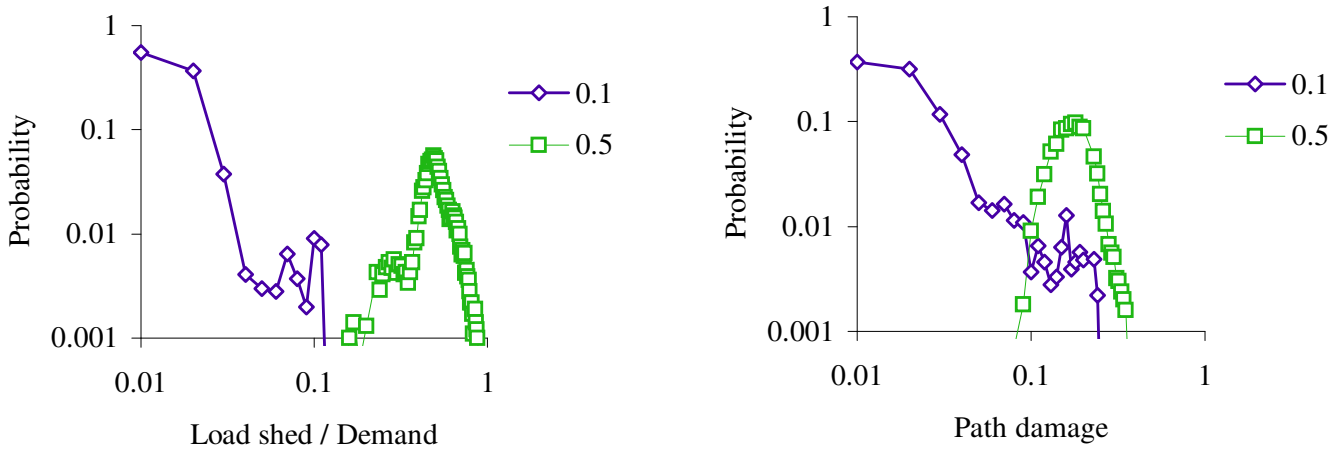

FIG. 7. Probability distribution of network damage for IEEE300 for $\alpha = 1.2$ with $r = 0.1$ and $r = 0.5$: (a) probability distribution of load shed/demand in OPA; (b) probability distribution of path damage in CLM.

from a given generator to only a portion of, rather than all, distributed substations. Secondly, it may be useful to consider cases with non-uniform initial edge weights, corresponding to non-uniform edge impedances in reality which are an important element in dictating power flows. Thirdly, slight modifications to betweenness centrality may be applied so that loading and transmission capacity are computed for edges rather than nodes. While we have considered heterogeneous transmission capacities across the network here, using random uniformly distributed fluctuations around α, it may be worthwhile to consider direct individual settings for α (with fluctuations as here) in different parts of the network. These extensions would enable us to map real power grid data to CLM, and enable more comprehensive studies and evaluations of the extent to which CLM remains useful in these more general cases.

ACKNOWLEDGMENTS

We thank B.A. Carreras [17] for providing valuable clarification regarding the implementation of the OPA model and providing the IEEE118 data used in prior OPA papers. We also thank the High Performance Computing and Communications Centre (http://www.hpccc.gov.au/) for the use of their supercomputer clusters in performing the experiments for this paper.

[1] Paolo Crucitti, Vito Latora, and Massimo Marchiori, "Model for cascading failures in complex networks," Physical Review E **69**, 045104 (2004).

[2] Ke Sun and Zhen-Xiang Han, "Analysis and comparison on several kinds of models of cascading failure in power system," in *Transmission and Distribution Conference and Exhibition: Asia and Pacific, 2005 IEEE/PES* (2005) pp. 1 – 7.

[3] Adilson E. Motter and Ying-Cheng Lai, "Cascade-based attacks on complex networks," Physical Review E **66**, 065102 (2002).

[4] Charles D. Brummitt, Raissa M. D'Souza, and E. A. Leicht, "Suppressing cascades of load in interdependent networks," Proceedings of the National Academy of Sciences **109**, E680–E689 (2012).

[5] R. Baldick, B. Chowdhury, I. Dobson, Zhaoyang Dong, Bei Gou, D. Hawkins, H. Huang, M. Joung, D. Kirschen, Fangxing Li, Juan Li, Zuyi Li, Chen-Ching Liu, L. Mili, S. Miller, R. Podmore, K. Schneider, Kai Sun, D. Wang, Zhigang Wu, Pei Zhang, Wenjie Zhang, and Xiaoping Zhang, "Initial review of methods for cascading failure analysis in electric power transmission systems IEEE PES CAMS task force on understanding, prediction, mitigation and restoration of cascading failures," in *IEEE Power and Energy Society General Meeting – Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE* (2008) pp. 1–8.

[6] R. Hooshmand and M. Moazzami, "Optimal design of adaptive under frequency load shedding using artificial neural networks in isolated power system," International Journal of Electrical Power & Energy Systems **42**, 220–228 (2012).

[7] P. A. Trodden, W. A. Bukhsh, A. Grothey, and K. I. M. McKinnon, "MILP formulation for controlled islanding of power networks," International Journal of Electrical Power & Energy Systems **45**, 501–508 (2013).

[8] I. Dobson, B.A. Carreras, V.E. Lynch, and D.E. Newman, "An initial model for complex dynamics in electric power system blackouts," in *Proceedings of the 34th Annual Hawaii International Conference on System Sciences, 2001* (2001) pp. 710 – 718.

[9] B.A. Carreras, V.E. Lynch, I. Dobson, and D.E. New-

man, "Critical points and transitions in an electric power transmission model for cascading failure blackouts," Chaos **12**, 985 – 994 (2002).
[10] B.A. Carreras, D.E. Newman, I. Dobson, and A.B. Poole, "Evidence for self-organized criticality in a time series of electric power system blackouts," IEEE Transactions on Circuits and Systems I: Regular Papers **51**, 1733 – 1740 (2004).
[11] R. Kinney, P. Crucitti, R. Albert, and V. Latora, "Modeling cascading failures in the North American power grid," European Physical Journal B **46**, 101 – 107 (2005).
[12] I. Dobson, B.A. Carreras, and D.E. Newman, "A loading-dependent model of probabilistic cascading failure," Probability in the Engineering and Informational Sciences **19**, 15 – 32 (2005).
[13] R. Leelaruji and V. Knazkins, "Modeling adequacy for cascading failure analysis," in *Australasian Universities Power Engineering Conference (AUPEC '08)* (IEEE, 2008) pp. 1–6.
[14] M. Rausand and A. Hoyland, *System Reliability Theory Models, Statis- tical Methods, and Application* (John Wiley and Sons, Hoboken, NJ, 2004).
[15] Martí Rosas-Casals and Ricard Solé, "Analysis of major failures in europe's power grid," International Journal of Electrical Power & Energy Systems **33**, 805–808 (2011).
[16] Liang Chang and Zhigang Wu, "Performance and reliability of electrical power grids under cascading failures," International Journal of Electrical Power & Energy Systems **33**, 1410–1419 (2011).
[17] B.A. Carreras, personal communication (2012).
[18] Vito Latora and Massimo Marchiori, "Efficient behavior of small-world networks," Physical Review Letters **87**, 198701 (2001).
[19] K. I. Goh, B. Kahng, and D. Kim, "Universal behavior of load distribution in scale-free networks," Physical Review Letters **87**, 278701 (2001).
[20] Adilson E. Motter and Ying-Cheng Lai, "Cascade-based attacks on complex networks," Physical Review E **66**, 065102 (2002).
[21] Joseph T. Lizier, David J. Cornforth, and Mikhail Prokopenko, "The information dynamics of cascading failures in energy networks," in *Proceedings of European Conference on Complex Systems, Coventry, UK, Sept 2009* (2009) p. 54.
[22] "IEEE300 and IEEE118 standard bus case," `http://www.ee.washington.edu/research/pstca/`.
[23] Valentina Cupac, Joseph T. Lizier, and Mikhail Prokopenko, "IEEE418," `http://prokopenko.net/ieee418.html`.
[24] Yang Nan, Liu Wenying, and Guo Wei, "Study on scale-free characteristic on propagation of cascading failures in power grid," in *IEEE Energytech 2011* (2011) pp. 1 –5.
[25] Réka Albert and Albert-László Barabási, "Statistical mechanics of complex networks," Reviews of Modern Physics **74**, 47–97 (2002).

Appendix A: Local behaviour in response to specific failures

It can be illustrated that CLM and OPA models exhibit differences in local behaviour in response to *specific* failures. We construct a simple network (see Fig. 8) with nodes $N_i, i \in 1...8$ and edges $E_{ij}, i, j \in 1...8$, and compare the two models' response to line failure, regarding: (i) the changes in state of network components, and (ii) the subsequent measure of network damage.

TABLE I. Network damage for each possible line failure

Edge removed ID	Load shed / demand (OPA)	Path damage (CLM)	Absolute difference
E_{12}	0.17	0.10	0.07
E_{23}	0.33	0.24	0.09
E_{34}	0.50	0.50	0.00
E_{45}	0.00	0.37	0.37
E_{56}	0.50	0.50	0.00
E_{67}	0.33	0.24	0.09
E_{78}	0.17	0.10	0.07

Figure 8(a) shows the impact of outaging edge E_{67} using OPA. After the outage, no power can flow through edges E_{67}, E_{78}. Thus load nodes N_7, N_8 receive no power, generator node N_5 supplies less power, and edge E_{56} transmits less power. We note that in the final state, the power output of generators is assymetrical, despite the symmetry of the network — this assymetry in generator output, due to the behaviour of the linear optimiser, has been previously documented (see Fig. 4 in [9]). It it important to note that, in theory, there may be several possible optimal solutions which satisfy power flow constraints in OPA, however the linear optimiser produces only a single optimal solution.

Figure 8(b) shows the impact of outaging edge E_{67} using CLM. The efficiency of edge E_{67} becomes zero, reducing the efficiency of any shortest paths passing through that edge to zero. Consequently, less paths are offered, reducing the transmission loading at nodes $N_5 - N_7$.

Therefore, OPA and CLM exhibit some differences their network components responded to the outage of edge E_{67}. In the case of OPA, edges E_{56}, E_{67}, E_{78} and nodes N_5, N_7, N_8 were affected. In the case of CLM, edges E_{67} and nodes $N_5 - N_7$ were affected by the outage.

Table I shows the load shed/demand (S/D) and path damage (D) resulting from all possible edge failures (one failure at a time). In general, the measures of load shed/demand and path damage differed, with a prominent difference observed for edge E_{45}. The ratios in network damage also differed, e.g. the impact of outaging edge E_{23} relative to E_{12} was $0.33/0.17 = 1.94$ in OPA and $0.24/0.10 = 2.4$ in CLM.

Based on individual failures shown in Table I, we computed the linear correlation between the measures of load shed/demand in OPA and path damage in CLM, finding some weak correlation ($R^2 = 0.35$). Thus, there was no strong correlation between these two measures even on a very simple network.

Furthermore, we examined the correlation of load shed/demand and path damage for individual node removals for each specific α value, using IEEE300 —

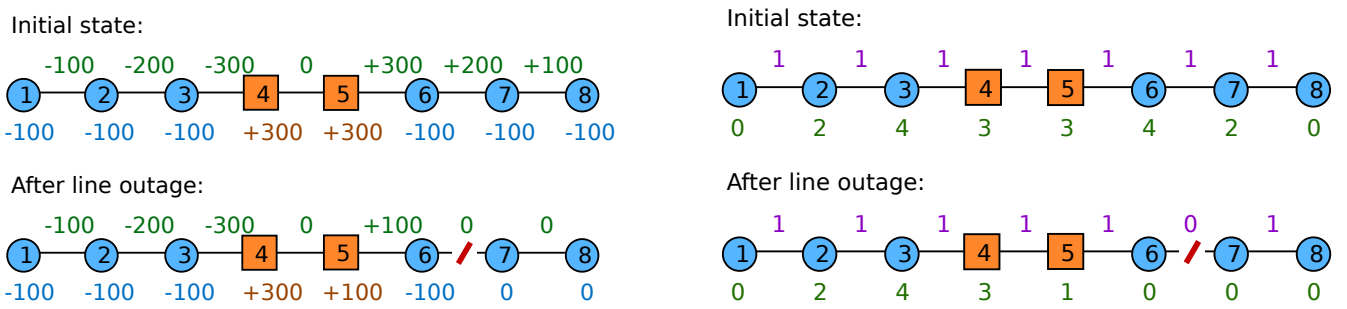

FIG. 8. Simple network with load nodes (blue circles) and generator nodes (orange squares), numbered to indicate position: (a) using OPA, indicating power flows F_{ij} shown above edges, power injections P_i shown below nodes (positive for generator nodes and negative for load nodes); (b) using CLM, indicating edge efficiencies e_{ij} shown above edges, and transmission loading L_i shown below nodes.

analagous to the experiments on the network level in Section V A, however with homogenous transmission capacity (i.e. $r = 0$). We found low levels of correlation: $R^2 = 0.27$ when $\alpha = 1.50$, $R^2 = 0.04$ when $\alpha = 1.10$, $R^2 = 0.03$ for $\alpha = 1.05$, $R^2 = 0.002$ for $\alpha = 1.01$. This indicates that, when we simulate the failure of one specific node, the resulting measure of load shed/demand (OPA) and path damage (CLM) are poorly correlated. Thus, OPA and CLM are not consistent at the local scale.