European Conference on Complex Systems 2009 manuscript No.
(will be inserted by the editor)

The information dynamics of cascading failures in energy networks

Joseph T. Lizier · Mikhail Prokopenko ·
David J. Cornforth

Received: 31 March 2009

Abstract Small failures in electrical energy networks can lead to cascading failures that cause large and sustained power blackouts. These can disrupt important services and cost millions of dollars. It is important to understand these events so that they may be avoided. We use an existing model for cascading failures to study the information dynamics in these events, where the network is collectively computing a new stable distribution of flows. In particular, information transfer and storage across the network are shown to exhibit sensitivity to reduced network capacity earlier than network efficiency does, and so could be a useful indicator of critical loading. We also show that the local information dynamics at each node reveals interesting relationships between local topological features and computational traits. Finally, we demonstrate a peak in local information transfer in time coinciding with the height of the cascade's spread.

Keywords information theory · information transfer · cascading failures

1 Introduction

Modern energy, communication and transportation networks have evolved in various complex ways to satisfy the demands of their end-users under various resource constraints, but of some concern is that they are all subject to *cascading failure events* [1, 8]: local failures that trigger avalanche mechanisms with large effects over the whole network. Our focus is energy networks, where usage has increased faster than investment in infrastructure (much of which is reaching the end of its useful life) and the grid has become critically loaded. As such, catastrophic blackouts and outages are occurring more frequently (see examples in [1]) and with more adverse impact.

Joseph T. Lizier · Mikhail Prokopenko
CSIRO Information and Communications Technology Centre, Locked Bag 17, North Ryde, NSW 1670, Australia

Joseph T. Lizier
School of Information Technologies, The University of Sydney, NSW 2006, Australia
E-mail: jlizier@it.usyd.edu.au

David J. Cornforth
CSIRO Energy Technology, PO Box 330, Newcastle, NSW 2300, Australia

It is likely that in the future energy networks will act as "smart grids", with intelligent devices at key points in the grid sensing the network's performance and collaborating to manage it. But in order to see such technology succeed, we first need a proper understanding of how the network operates during disturbances. Our approach is to examine the way the network *intrinsically* processes information during these extreme events. Some of the motivation here arises from recent advances in the field of complex systems science and artificial life, which has provided many examples of information processing in the natural world that can be applied in the manufactured environment, e.g. studies of regulatory networks (see [9,6]). Make no mistake: during cascading failures, the network is in fact *computing* it's new stable state (attractor), so understanding this computation can help understand the dynamics here.

Here, we examine an existing model of cascading failures in power networks [1] in order to measure the *information storage* and *transfer* taking place in the network during these events. These measurements are made using a framework for the information dynamics of distributed computation [5,7] that has previously been applied to a model of biological regulatory networks [6]. Importantly the framework is built on the *application-independence* of information theory, allowing an interchangeability of learnings with other complex systems. Our results here demonstrate maximisations in these computational capabilities near a phase transition in network efficiency as a function of spare capacity in the network. We also show interesting relationships between the information dynamics at each node in the network and the topological properties of the nodes. Finally, we find that the information transfer in the network is maximised as the cascade is growing at its fastest rate, demonstrating an important duality between the physical event and its informational interpretation.

2 Cascading failures model

Cascading failures in energy networks are studied here using the model described in [1] (which is also applicable to communication or transport networks, e.g. the Internet). The network is constructed as a weighted, undirected graph of N nodes (representing substations) and K edges (representing the transmission lines). The model describes the weights of each edge and loads of each node, and how these interdependently evolve after the breakdown of a node.

Each edge has an efficiency $e_{ij}(n)$ (akin to relative capacity in the energy networks analogy) that changes with time n; if there is an edge from node i to j then $e_{ij}(n) \in (0, 1]$ and is initialised to $e_{ij}(0) = 1$ (if there is no edge $e_{ij}(n) = 0$). Edge efficiency is the inverse of edge weight, while the efficiency $\epsilon_{ij}(n)$ of the most efficient *path* from i to j is the inverse of the shortest path length (which is of course the same path) [4].

Each node i has a load $L_i(n)$, being the total number of most efficient paths passing through it at time n (i.e. the load is the *betweenness centrality* of the node) [3]. Each node is also assigned a capacity C_i, being the maximum load it can handle without performance degradation. The capacity is assumed to be proportional to the initial load of the node: $C_i = \alpha L_i(0)$ [8], with $\alpha \geq 1$ being the *fixed* network tolerance.

The edge efficiencies become sub-optimal if (slightly altering the definition in [1]) *either* end-point node is operating above capacity:

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

Changes induced in edge efficiencies by excess loads cause changes in most efficient paths and therefore the distribution of loads at the next time step. While the initial state of the network is stable (since $\alpha \geq 1.0$), the removal of a node (simulating the breakdown of a substation) triggers a dynamical process where the network loads are redistributed. This process can cause other nodes to overload, shunting their loads onwards to other nodes who then overload etc, stimulating a cascading failure in time.

The performance of the network as a whole during these events is tracked using the average efficiency between all node pairs, $E(n) = \langle \epsilon_{ij}(n) \rangle$. Typically, for large α the network efficiency is relatively unaffected by node removal, however as α becomes closer to 1 the lower tolerance means that node removal can have dramatic cascading effects. Indeed the average network efficiency after node removals falls very quickly with α as $\alpha \rightarrow 1$, moving from a stable (efficient) state to an unstable one.

3 Information dynamics

We have proposed a framework for the local information dynamics of distributed computation in [5]. The framework describes computation in terms of information storage, transfer and modification at each spatiotemporal point in a complex system. Built on information theory, it provides a common, non-linear, application-independent language in which to analyse and design complex systems.

Information storage refers to the amount of information in the past of a node that is relevant to predicting its future. Memory has previously been studied in similar bursty or cascade-style systems in [2] using correlation coefficients of inter-event times, though information theory allows one to capture non-linear effects and is application independent; indeed, memory is about information. As such here we use the *active information storage* as the stored information that is currently in use in computing the next state of the node [5]. Specifically, the *local active information storage* for node X is the local (or unaveraged) mutual information between its semi-infinite past $x_n^{(k)}$ (as $k \rightarrow \infty$) and its next state x_{n+1} at time step $n+1$:

$$a_X(n+1,k) = \log_2 \frac{p(x_n^{(k)}, x_{n+1})}{p(x_n^{(k)})p(x_{n+1})}. \tag{2}$$

The *active information* is the average over time: $A_X(k) = \langle a_X(n,k) \rangle$.

Information transfer is formulated by the *transfer entropy* [10] as the information provided by a source node about a destination's next state that was not contained in the past of the destination. Specifically, the *local transfer entropy* [7] from a source node Y to a destination X is the local mutual information between the previous state of the source y_n and the next state of the destination x_{n+1}, *conditioned* on the semi-infinite past of the destination $x_n^{(k)}$ (as $k \rightarrow \infty$):

$$t_{Y \rightarrow X}(n+1,k) = \log_2 \frac{p(x_{n+1} | x_n^{(k)}, y_n)}{p(x_{n+1} | x_n^{(k)})}. \tag{3}$$

Again, the *transfer entropy* is the (time) average: $T_{Y \rightarrow X}(k) = \langle t_{Y \rightarrow X}(n,k) \rangle$. It can be measured for any two time series, but only represents information transfer when measured on a causal link. Importantly, the transfer entropy properly measures a directed, dynamic flow of information, unlike previous investigations of information transfer in networks with mutual information [9, 11] which measure correlations only.

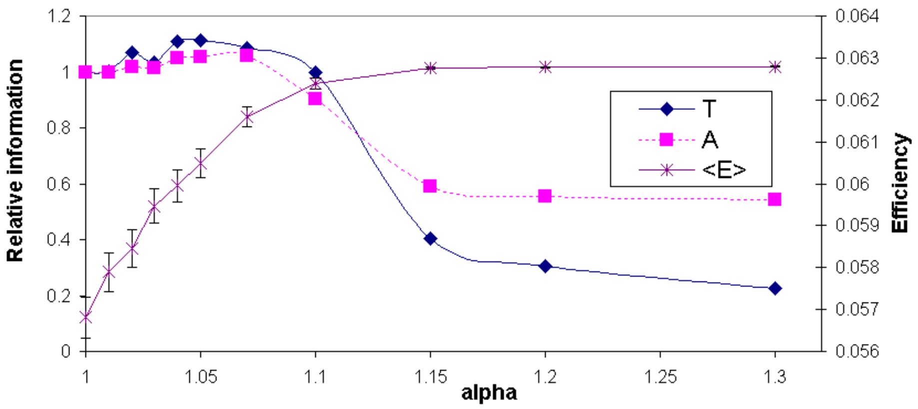

Fig. 1 Information dynamics versus network tolerance α, plotted relative to their values for $\alpha = 1.00$. Transfer entropy T is averaged over all time steps for each link, then averaged over all links, with $T(\alpha = 1.00) = 0.0155$ bits. Similarly, active information storage A is averaged over all nodes, with $A(\alpha = 1.00) = 1.10$ bits. Also plotted is the final network efficiency $\langle E \rangle$ averaged over each simulated node removal.

4 Results and discussion

To measure the information dynamics of cascading failure events in power grids, we use the model in Sect. 2 with the topology of the electrical power grid of the western United States [12] ($N = 4941$ and $K = 6594$). For $1.0 \leq \alpha \leq 1.3$, we simulate the breakdown of substations by removing 149 randomly selected nodes (one at a time) and then measure the resulting time series of node loads and edge efficiencies until the network reaches a (possibly periodically) stable state.

For each α, the active information storage A_i is measured from the loads for every node i, and the transfer entropies $T_{i \to j}$ and $T_{j \to i}$ are measured in each direction for each causal link ij between their time series of loads. We use a history length $k = 3$. Probability distribution functions are estimated from the set of time series of observations of loads obtained from all the separate node knockouts, using kernel estimation (e.g. see [10]) with a kernel width of 0.5 standard deviations of each variable.

First, we examine the **average information transfer** (across all links, $T = \langle T_{i \to j} \rangle$) and **average information storage** (across all nodes, $A = \langle A_i \rangle$) in the network **as a function of tolerance** α. Fig. 1 shows that both reflect the phase transition in average network efficiency, being maximised in the vicinity of this transition. Indeed, as α is reduced the information dynamics change more rapidly than network efficiency (see $1.1 \leq \alpha \leq 1.15$) and *could be a useful early indicator of critical loading.* The maximisations in these measures occur *before* the system is pushed into a chaotic state at $\alpha = 1.0$, in alignment with results from random boolean networks [6]. This is because the increased activity in the network as $\alpha \to 1.0$ begins to obscure the individual contribution of nodes and so measurement of information transfer falls.

We then examine the **information transfer values *locally* at each link**, $T_{i \to j}$: Fig. 2 shows a strong correlation of the transfer entropy (averaged over all applicable links) to source node degree. Correlations are large ($r \geq 0.80$) and significant ($p < 0.01$) for α in the vicinity of the phase transition, becoming somewhat weaker in the stable region. Similar results are found for active information storage and node degree. These effects are explained in that the more neighbours a node has, the more diversity

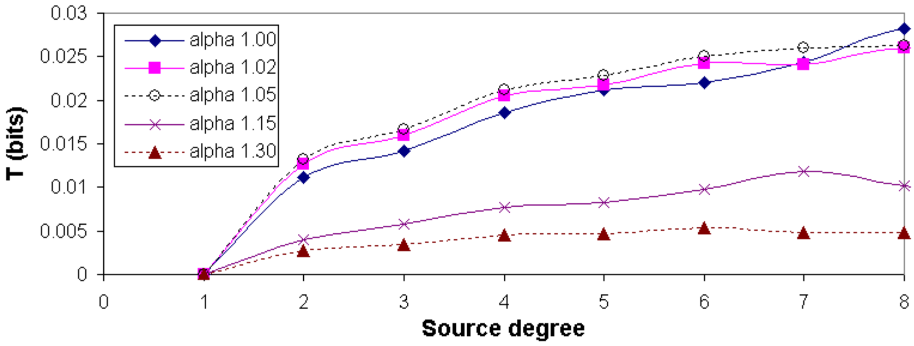

Fig. 2 Average transfer entropy T versus degree of the source node for various α. The small number of nodes with degree larger than eight introduces large errors beyond this point.

in its activity: this intrinsic variation in a node is a dominant factor in how much information it can store and transfer. Indeed, this is why nodes with degree one have zero information: with zero betweenness centrality their load never changes, so there is no diversity they can transfer and no influence on them to measure. Note however, that there is no corresponding correlation of the transfer entropy to the degree of the destination node. Certainly, the increased diversity in the destination provides greater scope for information transfer to it, but the more difficult it becomes to identify the coherent effect of the source. We also measure small but significant (at $p < 0.001$) correlation coefficients between the initial load of a node (i.e. its betweenness centrality) and the information it stores and transfers to other nodes during the cascade (e.g. for $\alpha = 1.05$ these are 0.24 and 0.13 respectively). In a similar fashion to degree, as the betweenness centrality of a node increases so to does its propensity to be influenced by other nodes, and therefore the diversity of its activity increases.

Finally, we examine **how the local information dynamics evolve in time** as the cascade unfolds. Fig. 3 shows that for $\alpha = 1.05$ the local information transfer values exhibit a strong peak around 12 time steps after the cascading failure (similar plots are observed for other values of α). This peak lags the initial failure event but coincides with the steepest drop in network efficiency. The cascade takes some time to build up in size, but causes the steepest change in network efficiency and largest information transfer when it is spreading at its most rapid. The relation between the *change* in network efficiency and local information transfer (in time) is quite strong: they have a correlation score of -0.94. The local information transfer could therefore be a useful indicator of the peak of the cascade, as well as its subsequent passing. Indeed, were the local values to be examined in space (i.e. on particular links) as well as time, they could be used as indicators of the direction of spread of the cascade. Interestingly though, the local values in space and time do not exhibit a strong correlation to local loads - it appears the utility of the local information transfer in time discussed above is an self-organised emergent effect which can be seen only at the macroscopic level.

5 Conclusion

We have characterised the information dynamics in cascading failure events in energy networks, as the network computes its new stable state. We have demonstrated maximisations in information dynamics across the network near the phase transition in

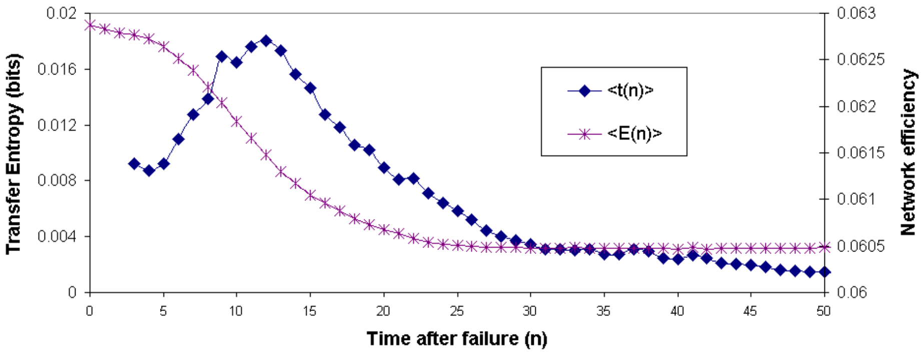

Fig. 3 Local transfer entropies in time (averaged across all links for the given time step n: $\langle t_{i\to j}(n)\rangle_{ij}$) and network efficiency $E(n)$ versus time after initial node removal. Both measures are averaged across all simulated node knockouts with $\alpha = 1.05$.

network efficiency, with the information dynamics initially changing faster than network efficiency as the tolerance drops. We have also shown that the local information dynamics (information stored, and transferred from) at each node are correlated with node degree. Finally, we have demonstrated a strong relationship between the spreading of the cascade (in an application-dependent level) and information transfer across the network (in an application-independent manner). In future work, we plan to examine the effect of altering the topology of the grid (e.g. with mini-grid topologies), and whether the insights gained here can be used to control such cascade events.

References

1. Crucitti, P., Latora, V., Marchiori, M.: Model for cascading failures in complex networks. Phys. Rev. E **69**(4), 045104 (2004)
2. Goh, K.I., Barabási, A.L.: Burstiness and memory in complex systems. Europhysics Letters **81**(4), 48002 (2008)
3. Goh, K.I., Kahng, B., Kim, D.: Universal behavior of load distribution in scale-free networks. Phys. Rev. Lett. **87**(27), 278701 (2001)
4. Latora, V., Marchiori, M.: Efficient behavior of small-world networks. Phys. Rev. Lett. **87**(19), 198701 (2001)
5. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: A framework for the local information dynamics of distributed computation in complex systems (2008). URL http://arxiv.org/abs/0811.2690. arXiv:0811.2690
6. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: The information dynamics of phase transitions in random boolean networks. In: S. Bullock, J. Noble, R. Watson, M.A. Bedau (eds.) Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems (ALife XI), Winchester, UK, pp. 374–381. MIT Press, Cambridge, MA (2008)
7. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Local information transfer as a spatiotemporal filter for complex systems. Phys. Rev. E **77**(2), 026110 (2008)
8. Motter, A.E., Lai, Y.C.: Cascade-based attacks on complex networks. Phys. Rev. E **66**(6), 065102 (2002)
9. Ribeiro, A.S., Kauffman, S.A., Lloyd-Price, J., Samuelsson, B., Socolar, J.E.S.: Mutual information in random boolean models of regulatory networks. Phys. Rev. E **77**(1), 011901 (2008)
10. Schreiber, T.: Measuring information transfer. Phys. Rev. Lett. **85**(2), 461–464 (2000)
11. Solé, R.V., Valverde, S.: Information transfer and phase transitions in a model of internet traffic. Physica A **289**(3-4), 595–605 (2001)
12. Watts, D.J., Strogatz, S.: Collective dynamics of small-world networks. Nature **393**, 440–442 (1998)