Interdependent networks: The fragility of control

2 authors, including:

Marc Barthelemy
Atomic Energy and Alternative Energies Commission
239 PUBLICATIONS   **28,310** CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:

Project   Book: Morphogenesis of spatial networks (Springer) View project

Project   GeoHistoricalData View project

Interdependent networks: the fragility of control

Richard G. Morris & Marc Barthelemy

Institut de Physique Théorique, CEA, CNRS-URA 2306, F-91191, Gif-sur-Yvette, France.

Recent work in the area of interdependent networks has focused on interactions between two systems of the same type. However, an important and ubiquitous class of systems are those involving monitoring and control, an example of interdependence between processes that are very different. In this Article, we introduce a framework for modelling 'distributed supervisory control' in the guise of an electrical network supervised by a distributed system of control devices. The system is characterised by degrees of freedom salient to real-world systems— namely, the number of control devices, their inherent reliability, and the topology of the control network. Surprisingly, the behavior of the system depends crucially on the reliability of control devices. When devices are completely reliable, cascade sizes are percolation controlled; the number of devices being the relevant parameter. For unreliable devices, the topology of the control network is important and can dramatically reduce the resilience of the system.

The study of interdependent, networked, systems is an area that has recently received a lot of attention[1–11], where the majority of work has so far focussed on the interactions between different 'critical infrastructures'[12–16]. We argue that critical infrastructures should themselves be viewed as a special class of interdependent systems, due to the presence of in-built monitoring and control mechanisms[12,17,18]. The type of control most prevalent in such systems is so-called 'supervisory' control— as distinguished from, say, controllability[19] — which typically involves monitoring an underlying process, with the option of a pre-defined intervention once a critical state is reached. Here, in keeping with the picture of interdependent networks, both monitoring and intervention are local processes, associated with specific points on the underlying network. Furthermore, we are interested in the case when the control is 'distributed', that is the local interventions are somehow coordinated via communications between sensors. At the most general level, we are interested in building a physics-like model of such systems: that is, complicated enough to encompass any interesting behaviour, but sufficiently idealized that the mechanisms at play can be easily identified and understood.

Our ideas are based on the supervisory control and data acquisition (SCADA) concept, ubiquitous in real-world monitoring of industrial manufacturing, power generation, and distribution processes (*e.g.*, electricity, gas, and water)[20]. To this end, our model comprises an underlying system, here, an electrical network, where a simple control device is placed on each transmission line with a probability p (see Fig. 1). The device monitors the load of that line and, if it is overloaded, then the device can dissipate the excess load with a probability of success q, and prevent the failure. In the opposite case, the line fails and the load is redistributed. The redistribution of loads may then lead to the overloading and failure of further power lines, and so on, potentially resulting in large system-wide outages[21]. If, at any stage during this process, more than one line becomes overloaded, then it is assumed that the next line to fail will be the one with the largest excess load. In the case where these lines are supervised, it therefore helps if the control devices respond in a coordinated way— always dissipating the excess load on the line under the greatest threat of failure. We therefore stipulate that for a control device to be operational, it must be in contact with a central processing unit (CPU). We envisage a communication network composed of ICT-like links connecting the devices and the CPU where, in keeping with a distributed SCADA picture, each device can also act as a signal relay— so called 'daisy chaining'. Crucially, this means that when a control device fails, it can disconnect many other devices from the CPU, rendering them useless— and dramatically increasing the fragility of the system. The structure of the supervisory network is therefore very important, and we consider two extremes. On one hand, a Euclidean minimum spanning tree (EMST) minimises the total length of the control network— and hence the cost— but typically sacrifices direct connectivity to the CPU. On the other hand, a mono-centric network maximises direct connectivity to the CPU, but can be very costly in terms of the total length of network needed. We interpolate between these two configurations by using a simple rewiring process: for each node in a EMST, replace with probability μ, the edge connected to the neighbor closest to the CPU (along the network), by an edge that connects directly to the CPU. The result is that the topology of the supervisory network
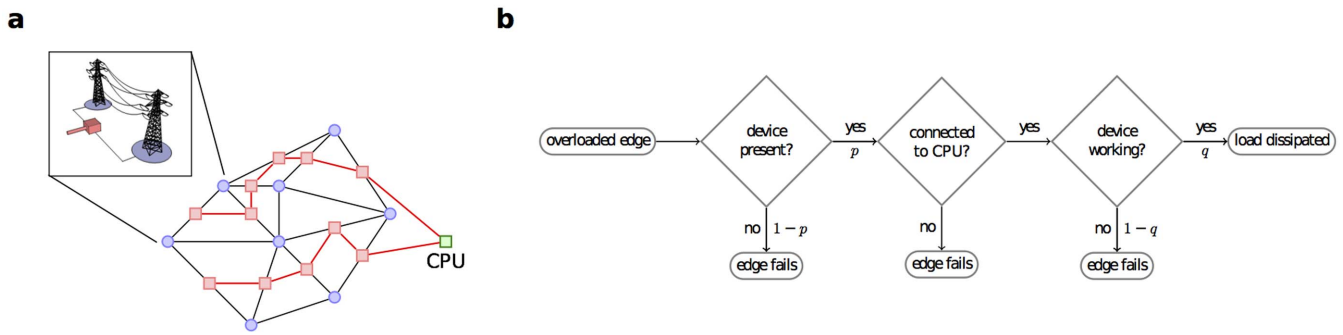
Figure 1 | Supervisory control. (a): The underlying power grid is represented by blue nodes and black edges. The load carried by power-lines (edges) is supervised by control devices, shown by red squares. The control devices act as signal relays and form a supervisory network (red edges) with the CPU. (b): For a given overloaded edge, there is a probability p that a control device is present. If this device is connected to the CPU, then it can be determined if the edge is carrying the largest excess load in the system. If so, the device will attempt to dissipate the excess load, with a success rate q.

relies on one continuous parameter $\mu \in [0, 1]$, such that $\mu = 0$ and $\mu = 1$ correspond to EMST and mono-centric networks respectively.

For modelling the electrical network we adopt a straightforward approach which has been proposed and analysed elsewhere[22]. The idea assumes a set of producers and consumers linked by power lines, where the resulting load carried by each line, or edge, may be represented by a random variable drawn from a uniform distribution U. Since U is properly normalized, the upper and lower bounds of the distribution are related to the average load $\bar{l}$, such that

$$U(\bar{l}) = \begin{cases} 1/2\bar{l} & \forall\, l \in [0, 2\bar{l}] & \text{if} \quad \bar{l} \leq 1/2 \\ 1/2(1-\bar{l}) & \forall\, l \in [2\bar{l}-1, 1] & \text{if} \quad \bar{l} > 1/2 \end{cases}. \quad (1)$$

In keeping with the above, it is also assumed that the transmission lines have an intrinsic carrying capacity (assumed here, without loss of generality, to be one) which, if exceeded, causes the line to fail and the load to be redistributed evenly amongst its nearest neighbours[23]. The crucial departure from Ref. 22, is in our choice of network topology. Since many critical infrastructures are, to a good approximation, planar subdivisions[24], we use the well known Delaunay triangulation[25], which is a simple, reasonable model for planar networks such as power grids.

Results

We test the vulnerability of our model against failure cascades by using computer simulations (see Methods section for details). For given values of the parameters p, q, μ, and $\bar{l}$, we repeatedly generate instances of the ensemble, each time initiating a cascade according to a 'fallen tree' approach— that is, an unspecified external event removes an edge and, if it is supervised, the associated control device. Following each cascade, N_{lcc}, the size of the remaining largest connected component of the underlying electricity network, is recorded. We assume that administrators/designers of real systems are interested in ensuring that cascades are bounded by a certain size. To this end, we consider

$$P_\varepsilon = P(1 - N_{lcc}/N \leq \varepsilon), \quad (2)$$

the probability that, following a cascade, the number of nodes disconnected from the largest connected component— the effective cascade 'size': $1 - N_{lcc}/N$— is less than a fraction $\varepsilon \in (0, 1]$ of the original nodes.

In general, as one would expect, the larger the average load carried by the system, the smaller the probability that the cascade size is bounded (see Fig. 2a). However, we also observe another feature of this type of cascading model, first identified in Ref. 22: for each value of p, there is a non-zero critical value

$$\bar{l}_\varepsilon = \sup\{\bar{l} \in (0,\, 1) : P_\varepsilon(\bar{l}) = 1\}, \quad (3)$$

that corresponds to the maximum average load below which cascade sizes are bounded with probability one (within a given accuracy, here 1 part in 5×10^3). Plotting the values of $\bar{l}_{1/2}$ against p, a sharp transition can be observed at some point p^* (see Fig. 2b). Above this value, the fraction of disconnected nodes is always bounded by $\varepsilon = 1/2$, regardless of how much load the system is carrying. In the completely reliable case ($q = 1$), p^* just corresponds to the percolation threshold p_c (~ 0.33 for Delaunay triangulations[26]). The cascades are then 'percolation controlled' due to the formation of a giant component connected by supervised edges, coined here the giant supervised component (GSC). The upper bound on cascade size that is enforced by the GSC can be lowered by employing more control devices— i.e., increasing p (see Fig. 2c). For $p \gtrsim 1 - p_c$, most nodes are connected by supervised edges and cascades cannot disconnect nodes from the giant component.

Whilst $q = 1$, the only impact of decreasing μ is to increase the number of devices disconnected by the initial external shock. Disregarding the correlation induced by starting the cascade at the point of disconnection, this effect corresponds to a small shift

$$\delta \sim \langle s \rangle/N, \quad (4)$$

in the positive x-direction of Figs. 2b and 2c. Here, $\langle s \rangle$ is the average sub-tree size associated with a randomly chosen node (see Fig. 2c inset). Figure 2d shows the effects of this shift when $p > p^*$, for both large and small ε. Here, it is natural to characterize changes in μ by a normalized cost

$$C = L(\mu)/L(1), \quad (5)$$

where $L(\mu)$ is the total length of the supervisory network. The message of Fig. 2d is that: increasing the number of direct CPU connections at the cost of increased network length, is only beneficial if the suppression of small cascades is desired.

If, in contrast to above, the control devices have an inherent rate-of-failure ($q < 1$), then a GSC may be either disconnected or reduced in size as control devices fail. In the best case scenario, when the supervising network is mono-centric and q is close to one, the picture is one of 'effective percolation' with (see Methods)

$$p^* = p_c q^{-\alpha}, \quad (6)$$

where α is determined by the topology of the underlying network (~ 2.4 for a Delaunay triangulation, see Methods section for details). This simple form shows good agreement with direct estimates of the value of p^* (see Fig. 3b and Methods for details). For lower values of q, percolation-like descriptions are no longer appropriate: regardless of the number of control devices, it is not possible to bound cascade sizes in a way that is independent of the average load carried by the system. Indeed, if control devices are both unreliable ($q < 1$) and the
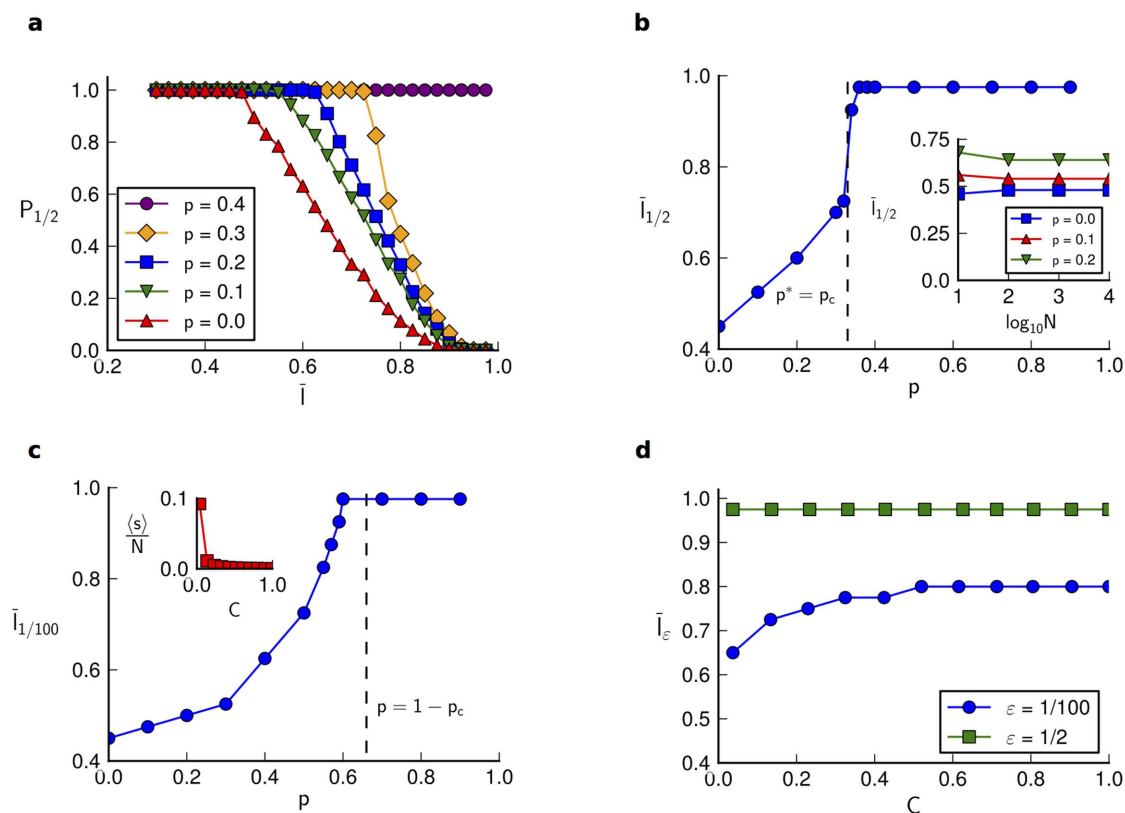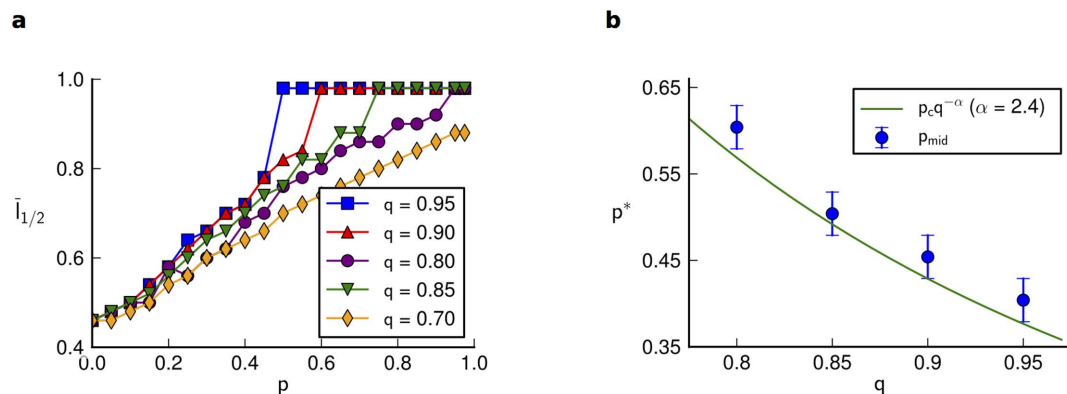
Figure 2 | The effects of reliable control devices ($q = 1$). (a): The probability that, following a cascade, the remaining largest connected component of the underlying grid contains more than half of the nodes $P_{1/2}$, is dependent on the average load carried by the system $\bar{l}$, and the number of control devices present p. For each value of p, the system is characterized by a critical average load $\bar{l}_{1/2}$. Below this critical value, cascades never disconnect more than half of the system ($P_{1/2} = 1$), whilst above it, there is always a finite chance that this will happen ($P_{1/2} < 1$). (b): As the bond-percolation threshold $p^* = p_c \sim 0.33$ is approached, the critical value $\bar{l}_{1/2}$ rises sharply to one due to the formation of a giant supervised component (GSC). Inset: Results are unchanged by increases in system size. (c): The bound on cascade size can be lowered by increasing $p > p_c$ and therefore the size of the GSC. For $p \geq 1 - p_c$ most nodes are connected by supervised edges and therefore cascades cannot disconnect any nodes completely. Inset: For values of the cost (estimated as the total length of the supervisory network) above $C \sim 1/2$ the average sub-tree size $\langle s \rangle$ of the control network— and therefore the average number of devices disconnected at cascade initiation— is less than 1 and negligible as a fraction of the system size. (d): Critical value $\bar{l}_\varepsilon$ for $p > p^*$. In this case, increasing the cost of the supervisory network only increases the critical load associated with bounding small cascades, and not those of the order of the system size.

control network is tree-like ($\mu < 1$), the system is very susceptible to large failure cascades, with little impact made by increasing p (see Fig. 4). In this case, we see that for both large and small cascades, the topology of the control network is very relevant and can induce extreme fragility in the control system (see Fig. 5).

Discussion

In conclusion, we have introduced a minimal model which incorporates the salient features of many real-world control systems. Firstly, the control devices are simple: they only have so-called 'supervisory' functions of monitoring and intervention. Secondly, the system is

Figure 3 | The effects of control device failure ($q \neq 1$) when every device is connected directly to the CPU ($\mu = 1$). (a): As the reliability of the control devices decreases, more devices are needed to maintain the same critical load. (b): Agreement between the numerical value p_{mid} obtained for the $\bar{l}_{1/2}$ transition— and the theoretical form $p_c q^{-\alpha}$ motivated by simple arguments (see main text and Methods Section).
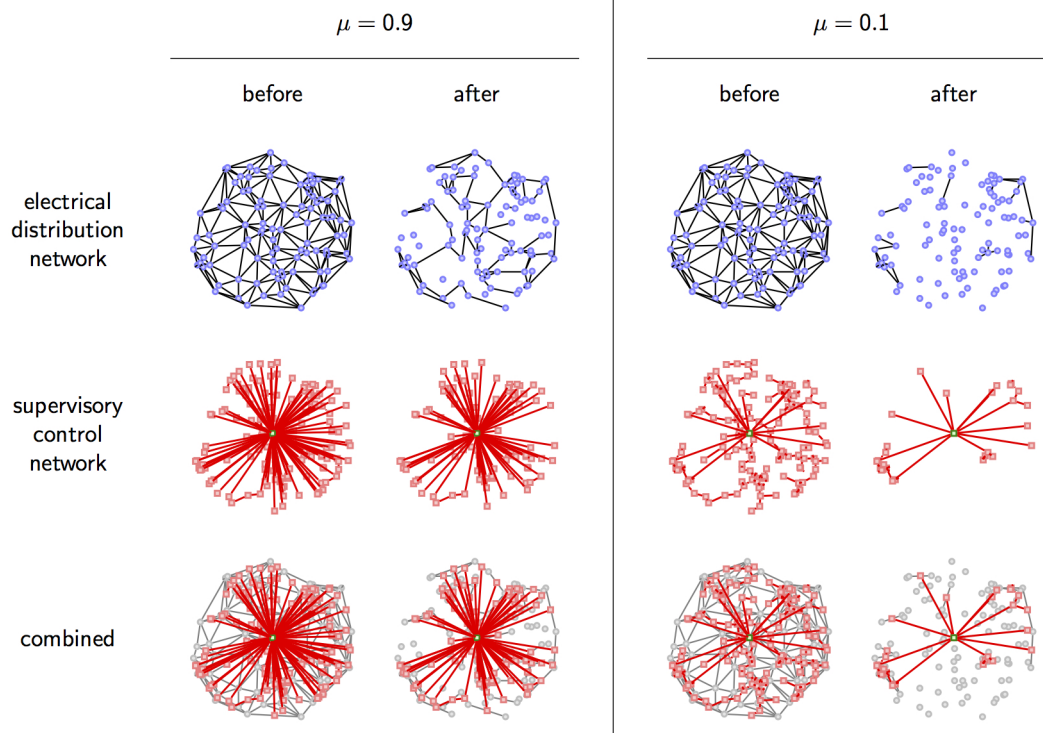
Figure 4 | The effect of μ ($p = 0.5$, $q = 0.9$). When the supervising network is almost mono-centric ($\mu = 0.9$) very few control devices fail and therefore the remaining largest connected component connects 85% of the nodes in the system. If the supervising network is almost a tree ($\mu = 0.1$) then even though the inherent failure rate is low, many devices become disconnected from the CPU and therefore only 10% of nodes are left connected following a cascade.

'distributed', that is, not only are the devices positioned in space but they require coordination— in this case, by connection to a CPU. Thirdly, we also incorporate the effects of devices having an inherent rate-of-failure. With only these simple characteristics, the resulting behaviour is very rich. The primary feature concerns the fragility of such control systems: a small reduction of control device reliability leads to a regime where the ability to suppress cascades is dramatically affected by the topology of the control network. Our results suggest that it is much more cost-effective to try to improve the reliability of control devices rather than working on the stability of the supervisory control network. We believe that these results make a first step in understanding distributed supervisory control, whilst also providing helpful guidelines to designers and administrators of real systems. We welcome further work in the area.
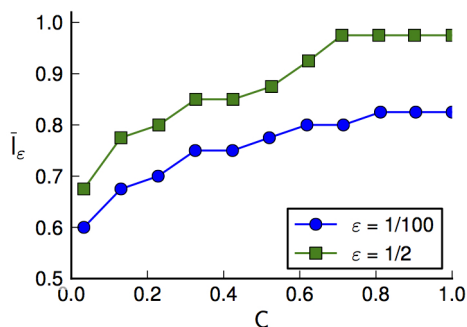
Methods

Simulations. To simulate the system, N nodes are placed in the plane at random, the Delaunay triangulation is then formed, and loads are allocated to the resulting edges according to $U(\bar{l})$. The supervisory network is incorporated by first adding a control device to each edge with probability p, then forming the network according to the rewiring procedure described in the main text (dependent on parameter μ). Cascades are initiated by assuming an external event that causes an edge to be removed at random and its load is redistributed amongst its nearest neighbors. If the failing edge was supervised, then the control device is also removed. During the ensuing cascade, we stipulate that for a control device to work, it must be connected to the CPU, a special node that cannot be removed. If a control device is unconnected, then it cannot work and is of no use. However, if a control device is connected, and it is supervising an edge that is about to fail— *i.e.*, it is carrying the largest excess load in the system— then there is a probability q that the excess load is dissipated and the load of the edge is reset to $\bar{l}$. The quantity q can be thought of as the inherent reliability of a device.

Simulations were written in C++ and implemented using the Boost Graph library[27] where possible. Delaunay triangulations were produced using an iterative algorithm[25].

Results are presented for systems of size $N = 500$ ($\sim 3 \times 10^3$ edges) and statistics are calculated over 5×10^3 instances of each ensemble (defined by parameters p, q, μ, and $\bar{l}$). Critical values $\bar{l}_\varepsilon$ and p^* are accurate up to an error of approximately ± 0.02, since they are identified by varying the underlying parameter by finite increments. In Figs. 4c and 5, $\bar{l}_\varepsilon$ corresponds to $P_\varepsilon > 0.99$ in order to accommodate the noise associated with different control network structures.

Figure 5 | The effect of topology in control networks with unreliable devices ($q = 0.9$). For all cascade sizes (ie. regardless of ε), the critical load depends strongly the structure of the supervisory network, in contrast with the completely reliable case. In particular, even when there are already many direct-CPU links ($C > 1/2$), the critical load that the system can carry is drastically reduced by introducing more dependency into the supervisory network.

Formation of an effective GSC. Labelling each supervised edge by $i = 1, 2, \ldots, E_s$, the probability that a supervised edge survives a cascade is q^{n_i}, where n_i is the number of times a device is solicited— *i.e.*, it tries to dissipate its excess load with probability q. Here, for large enough systems the number of supervised edges is given by $E_s = pE$. (Since the average degree of a Delaunay triangulation is peaked around six, the total number of edges E is well approximated by $E \sim 3N$.) Using a bar to denote system average $\overline{q^n} = 1/E_s \sum_i q^{n_i}$, we know that if Var $[n]$ is small, then $\overline{q^n} \sim q^{\bar{n}}$. Approximating a large system average with an ensemble average $\langle \ldots \rangle$ over many smaller systems, the results are given in Table 1. Here it is clear that the average $\langle n \rangle$ is well approximated by the value 2.4, regardless of p and q, and that the variance is

4

Table 1 | For different values of p and q, $\langle n \rangle \sim 2.4$, and Var $[n] \ll \langle n \rangle$

	$\langle n \rangle$	Var $[n]$
$q = 0.95, p = 0.4$	2.36	0.0004
$q = 0.9, p = 0.45$	2.40	0.0007
$q = 0.85, p = 0.5$	2.40	0.001
$q = 0.8, p = 0.6$	2.44	0.001

always very small compared to the average. We can then write the effective probability that a generic edge resists failure as

$$p_{\text{eff}} = p \cdot q^{\alpha}, \qquad (7)$$

with $\alpha = \langle n \rangle \simeq 2.4$. The system will then be resilient if $p_{\text{eff}} = p_c$, which implies Eq. (6).

Equation (6) may be contrasted with a direct approximation of when an effective GSC forms. From simulation results, we associate each transition with the value p_{mid}, defined as halfway between p_c and the lowest value of p for which $\bar{l}_\varepsilon$ is maximal (*i.e.*, the midpoint of the transition).

1. Buldyrev, S. V., Parshani, R., Paul, G. & Stanley, H. E. Catastrophic failures in interdependent networks. *Nature (London)* **464**, 1025–1028 (2010).
2. Parshani, R., Buldyrev, S. & Havlin, S. Interdependent Networks: Reducing the Coupling Strength Leads to a Change from a First to Second Order Percolation Transition. *Phys. Rev. Lett.* **105**, 048701 (2010).
3. Huang, X., Gao, S., Buldyrev, S. V., Havlin, S. & Stanley, H. E. Robustness of interdependent networks under targeted attack. *Phys. Rev. E* **83**, 065101(R) (2011).
4. Gu, C.-G. *et al.* Onset of cooperation between layered networks. *Phys. Rev. E* **84**, 026101 (2011).
5. Gao, J., Buldyrev, S. V., Stanley, H. E. & Havlin, S. Networks formed from interdependent networks. *Nature Phys.* **8**, 40–48 (2011).
6. Brummitt, C. D., D'Souza, R. M. & Leicht, E. A. Suppressing cascades of load in interdependent networks. *Proc. Natl. Acad. Sci. USA* **109**, 680–689 (2012).
7. Brummitt, C. D., Lee, K.-M. & Goh, K.-I. Multiplexity-facilitated cascades in networks. *Phys. Rev. E* **85**, 045102(R) (2012).
8. Saumell-Mendiola, A., Serrano, M. A. & Boguna, M. Epidemic spreading on interconnected networks. *Phys. Rev. E* **86**, 026106 (2012).
9. Li, W., Bashan, A., Buldyrev, S. V., Stanley, H. E. & Havlin, S. Cascading Failures in Interdependent Lattice Networks: The Critical Role of the Length of Dependency Links. *Phys. Rev. Lett.* **108**, 228702 (2012).
10. Bashan, A., Berezin, Y., Buldyrev, S. V. & Havlin, S. The extreme vulnerability of interdependent spatially embedded networks. ArXiv:1206.2062 (2012).
11. Morris, R. G. & Barthelemy, M. Transport in spatial networks. *Phys. Rev. Lett.* **109**, 128703 (2012).
12. Rinaldi, S., Peerenboom, J. & Kelly, T. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control. Syst. Magn.* **21**, 11–25 (2001).
13. Heller, M. Interdependencies in Civil Infrastructure Systems The Bridge. *Natl. Acad. Eng.* **31**, 9–15 (2001).
14. Carreras, B. A., Newman, D. E., Gradney, P., Lynch, V. E. & Dobson, I. Interdependent Risk in Interacting Infrastructure Systems in Proceedings of the 40th Annual Hawaii International Conference on System Sciences. (2007).
15. Rosato, V. *et al.* Modelling interdependent infrastructures using interacting dynamical models. *Int. J. Crit. Infrastruct.* **4**, 63–79 (2008).
16. Vyatkin, V. *et al.* Standards-enabled Smart Grid for the future Energy Web 2010 Innovative Smart Grid Technologies (ISGT). IEEE. New York, 2010, pp. 1–9.
17. Shea, D. A. Critical Infrastructure: Control Systems and the Terrorist Threat CRS Report for Congress. RL31534 (2004).
18. Bessani, A. N., Sousa, P., Correia, M., Neves, N. F. & Verissimo, P. The CRUTIAL Way of Critical Infrastructure Protection. *Security & Privacy, IEEE* **6**, 44–51 (2008).
19. Liu, Y.-Y., Slotine, J.-J. & Barabasi, A.-L. Controllability of complex networks. *Nature (London)* **473**, 167–173 (2011).
20. Communication Technologies Inc., NCS Technical Information Bulletin 04-01. (2004).
21. Baldick, R. *et al.* 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. *Power and Energy Society General Meeting-Conversion and Delivery of Electrical Energy in the 21st Century, IEEE* (pp. 1–8). IEEE (2008).
22. Moreno, Y., Pastor-Satorras, R., Vázquez, A. & Vespignani, A. Critical load and congestion instabilities in scale-free networks. *Europhys. Lett.* **62**, 292–298 (2003).
23. Duxbury, P. M., Leath, P. L. & Beale, P. D. Breakdown properties of quenched random systems: The random-fuse network. *Phys. Rev. B* **36**, 367–380 (1987).
24. Barthelemy, M. Spatial Networks. *Phys. Rep.* **499**, 1–101 (2011).
25. Guibas, L. & Stolfi, J. Primitives for the manipulation of general subdivisions and the computation of Voronoi. *ACM T. Graphic.* **4**, 74–123 (1985).
26. Becker, A. & Ziff, R. Percolation thresholds on two-dimensional Voronoi networks and Delaunay triangulations. *Phys. Rev. E* **80**, 041101 (2009).
27. Siek, J. G., Lee, L.-Q. & Lumsdaine, A. Boost Graph Library: User Guide and Reference Manual (Addison-Wesley, Boston, 2002).

Acknowledgements

Author contributions

Together, RGM and MB conceived of, and analysed the model, with RGM producing simulations. The manuscript was written collaboratively between both authors.

Additional information

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Morris, R.G. & Barthelemy, M. Interdependent networks: the fragility of control. *Sci. Rep.* **3**, 2764; DOI:10.1038/srep02764 (2013).