PHYSICAL REVIEW E **69**, 045104(R) (2004)

Model for cascading failures in complex networks

Paolo Crucitti,[1] Vito Latora,[2] and Massimo Marchiori[3]
[1]*Scuola Superiore di Catania, Via S. Paolo 73, 95123 Catania, Italy*
[2]*Dipartimento di Fisica e Astronomia, Università di Catania, and INFN, Sezione di Catania, Via S. Sofia 64, 95123 Catania, Italy*
[3]*W3C and Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

(Received 16 September 2003; published 29 April 2004)

Large but rare cascades triggered by small initial shocks are present in most of the infrastructure networks. Here we present a simple model for cascading failures based on the dynamical redistribution of the flow on the network. We show that the breakdown of a single node is sufficient to collapse the efficiency of the entire system if the node is among the ones with largest load. This is particularly important for real-world networks with a highly hetereogeneous distribution of loads as the Internet and electrical power grids.

DOI: 10.1103/PhysRevE.69.045104 PACS number(s): 89.75.Hc, 89.20.Hh, 89.30.−g, 89.75.Fb

Cascading failures are common in most of the complex communication and/or transportation networks [1,2] that are the basic components of our lives and industry. In fact, although most failures emerge and dissolve locally, largely unnoticed by the rest of the world, a few trigger avalanche mechanisms that can have large effects over the entire networks.

Cascading failures take place on the Internet, where traffic is rerouted to bypass malfunctioning routers, eventually leading to an avalanche of overloads on other routers that are not equipped to handle extra traffic. The redistribution of the traffic can result in a congestion regime with a large drop in the performance. For instance in October 1986, during the first documented Internet congestion collapse, the speed of the connection between the Lawrence Berkeley Laboratory and the University of California at Berkeley, two places separated only by 200 m, dropped by a factor 100 [3,4].

Cascading failures also take place in electrical power grids. In fact, when for any reason a line goes down, its power is automatically shifted to the neighboring lines, which in most of the cases are able to handle the extra load. Sometimes, however, these lines are also overloaded and must redistribute their increased load to their neighbors. This eventually leads to a cascade of failures: a large number of transmission lines are overloaded and malfunction at the same time. This is exactly what happened on 10 August 1996 [5,6] when a 1300-mw electrical line in southern Oregon sagged in the summer heat, initiating a chain reaction that cut power to more than 4 million people in 11 Western States. And probably this is also what happened on 14 August 2003 when an initial disturbance in Ohio [7] triggered the largest blackout in the U.S.'s history in which millions of people remained without electricity for as long as 15 h.

Large cascading failures are also present in social and economic systems [8].

How is it possible that a small initial shock, such as the breakdown of an Internet router (or of an electrical substation or line), can trigger avalanches mechanisms affecting a considerable fraction of the network and collapsing a system that in the past was proven to be stable with respect to similar shocks? In this paper we propose a simple model for cascading failures in complex networks. Resistance of networks to the removal of nodes or arcs, due either to random breakdowns or to intentional attacks, has been studied in Refs. [9–13]. Such studies have focused only on the *static properties* of the network showing that the removal of a group of nodes altogether can have important consequences. Here we show how the breakdown of a *single node* is sufficient to collapse the entire system simply because of the *dynamics of redistribution of flows* on the network. In our model each node is characterized by a given *capacity* to handle the traffic. Initially the network is in a stationary state in which the *load* at each node is smaller than its capacity. The breakdown (removal) of a node changes the balance of flows and leads to a redistribution of loads over other nodes. If the capacity of these nodes cannot handle the extra load this will be redistributed in turn, triggering a cascade of overload failures and eventually a large drop in the network performance such as those observed in real systems, like the Internet or the electrical power grids. The main differences with respect to previous models [14–16] are as follows.

(1) Overloaded nodes are not removed from the network. It is the communication passing through overloaded (congested) nodes that will get worse, so that eventually the information/energy will avoid congested nodes.

(2) The damage caused by a cascade is quantified in terms of the decrease in the network *efficiency*, a variable defined in Ref. [17].

First we introduce the model and then we show some applications to artificially created topologies, to the Internet, and to the electrical power grid of the western United States.

We represent a generic communication and/or transportation network as a valued (weighted) [18] undirected [19] graph **G**, with N nodes (the Internet routers or the substations of an electrical power grid) and K arcs (the transmission lines). **G** is described by the $N\times N$ adjacency matrix $\{e_{ij}\}$. If there is an arc between node i and node j, the entry e_{ij} is the value, a number in the range (0,1] attached to the arc; otherwise $e_{ij}=0$ [20]. Such a number is a measure of the efficiency in the communication along the arc. For instance, in the Internet, the smaller e_{ij} is, the longer it takes to exchange a unitary packet of information along the arc between i and j. Initially, at time $t=0$, we set $e_{ij}=1$ for all the existing arcs, meaning that all the transmission lines work perfectly and are equivalent. The model we will propose consists of a rule for the time evolution of $\{e_{ij}\}$ that mimics the dynamics of flow redistribution following the breakdown of a node. To define the network efficiency [17] we assume that the communication between a generic couple of nodes takes the most efficient path connecting them. The efficiency of a path is the so-called harmonic composition [21–23] of the efficiencies of the component arcs. By ϵ_{ij} we indicate the efficiency of the most efficient path between *i* and *j*. Matrix $\{\epsilon_{ij}\}$ is calculated by means of the algorithms used in Ref. [17]. Then the average efficiency of the network is

$$E(\mathbf{G}) = \frac{1}{N(N-1)} \sum_{i \neq j \in \mathbf{G}} \epsilon_{ij} \tag{1}$$

and is used as a measure of the performance of **G** at a given time.

The *load* $L_i(t)$ on node *i* at time *t* is the total number of most efficient paths passing through *i* at time *t* [24]. Each node is characterized by a *capacity* defined as the maximum load that node can handle. Following Ref. [14] we assume the capacity C_i of node *i* to be proportional to its initial load $L_i(0)$:

$$C_i = \alpha L_i(0), \quad i = 1,2, \ldots, N, \tag{2}$$

where $\alpha \geqslant 1$ is the tolerance parameter of the network [25]. This is a realistic assumption in the design of an infrastructure network, since the capacity cannot be infinitely large because it is limited by the cost. With such a definition of capacity, the network we have created is in a stationary state in which it operates with a certain efficiency *E*. The initial removal of a node [26], simulating the breakdown of an Internet router or of an electrical substation, starts the dynamics of redistribution of flows on the network. In fact the removal of a node changes the most efficient paths between nodes and consequently the distribution of the loads, creating overloads on some nodes. At each time *t* we adopt the following iterative rule:

$$e_{ij}(t+1) = \begin{cases} e_{ij}(0)\dfrac{C_i}{L_i(t)} & \text{if } L_i(t) > C_i \\ e_{ij}(0) & \text{if } L_i(t) \leqslant C_i, \end{cases} \tag{3}$$

where *j* extends to all the first neighbors of *i*. In this way if at time *t* a node *i* is congested, we reduce the efficiency of all the arcs passing through it, so that eventually the information/energy will take alternative paths (the new most efficient paths). This is a softer and, for some applications, a more realistic situation than the one considered in Ref. [14], in which the overloaded nodes are removed from the network. Rule (3) produces a decrease of the efficiency of the network *E* and, as we will show in the following, in some cases it can trigger an avalanche mechanism collapsing the whole system.

We illustrate how our model works in practice by considering two artificially created network topologies: (1) Erdös-Rényi (ER) random graphs [27]; (2) scale-free networks, i.e., graphs with an algebraic distribution of degree $P(k) \sim k^{-\gamma}$ with $\gamma = 3$ generated according to the Barabási-Albert (BA) model [28].

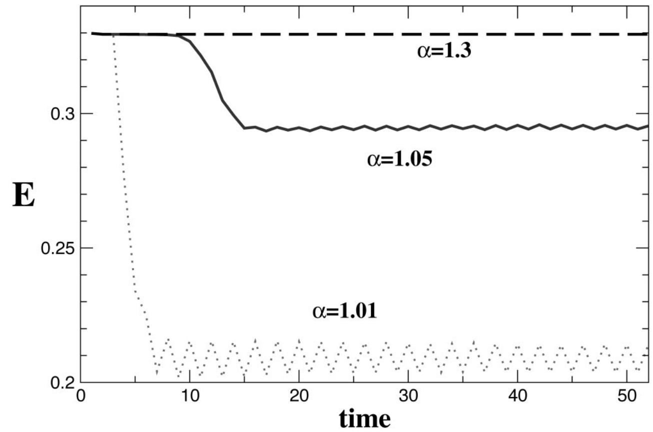

FIG. 1. Cascading failure in a BA scale-free network as triggered by the initial removal of a single node chosen at random. We plot the efficiency *E* of the network as a function of the time for three values of the tolerance parameter α. The curves correspond to an average over ten triggers.

In both cases we have constructed networks with $N = 2000$ and $K = 10\,000$. In Fig. 1 we report the typical time evolution of the network efficiency for the BA scale-free network. The dynamics of redistribution of flows is triggered by the removal at time $t = 0$ of a node chosen at random. We show the results for three values of the tolerance parameter, namely, $\alpha = 1.3$, 1.05, 1.01. In the first case the efficiency of the network is completely unaffected by the failure of the node. In the second case the network reaches a stationary state with an efficiency lower than the initial one. In the third case, because of the lower tolerance parameter, the cascading failures collapse the system: the network has lost 40% of the initial efficiency.

In Fig. 2 we report the final value of the efficiency, i.e., the efficiency after the system has relaxed to a stationary state, as a function of the tolerance parameter α. We consider both the ER random graph and the BA scale-free graph. Moreover, we adopt two different triggering strategies: *random removals* and *load-based removals*. In the first case (squares) the node removed initially is chosen at random: in this way we simulate the breakdown of the average node of the network. In the second case (full circles) the removed node is a very special one because it is the one with the largest load. Both for the random and for the scale-free network we observe a decrease of the efficiency for small values of the tolerance parameter α, and the collapse of the system for values smaller than a critical value α_c. ER random graphs appear to be more resistant to cascading failures than BA scale-free graphs (as also found in the model of Ref. [14]). In both cases the collapse transition is always sharper for load-based removals than for random removals, although the values of α_c can fluctuate for different realizations. For the ER random graphs considered we have obtained $\alpha_c = 1.02 \pm 0.002$ for random removals, and $\alpha_c = 1.06 \pm 0.005$ for load-based removals. For BA scale-free graphs $\alpha_c = 1.1 \pm 0.004$ for random removals, and $\alpha_c = 1.3 \pm 0.05$ for load-based removals [29]. The heterogeneity of the network plays an important role in the network stability. ER random graphs

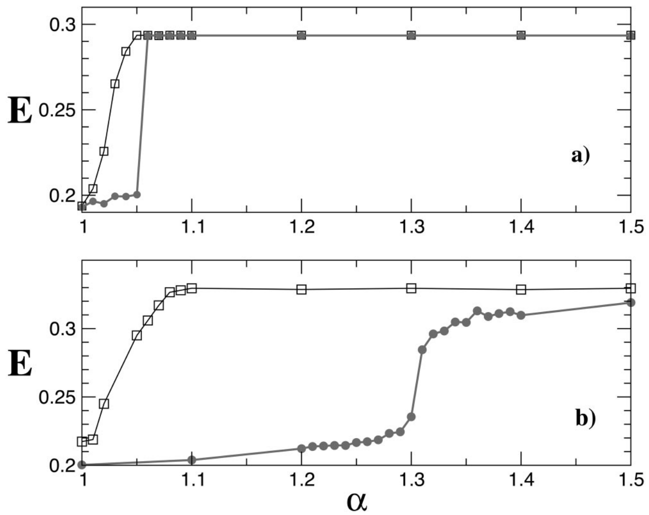

FIG. 2. Cascading failure in (a) ER random graphs and (b) BA scale-free networks as triggered by the removal of a node chosen at random (squares), or by the removal of the node with largest load (full circles). We report the final (after the cascade) efficiency E of the network as a function of the tolerance parameter α. Both the networks considered have $N=2000$ and $K=10\,000$. In the case triggered by the removal of a node chosen at random, the curve corresponds to an average over ten triggers.

have an exponential load distribution while BA networks exhibit a power-law distribution in the node load [24]. This makes a large difference between random removals and load-based removals in BA scale-free networks. In fact there are few nodes, the ones with extremely high initial load, that are far more likely than the other nodes (the most part of the nodes of network) to trigger cascades. Figure 2(b) shows the existence of a large region in the tolerance parameter, $1.1 \leqslant \alpha \leqslant 1.3$, where scale-free networks are stable with respect to random removals and are unstable with respect to load-based removals. If, for instance the nodes work with a tolerance of 30% above the standard load ($\alpha=1.3$), the network is in general very stable to an initial shock consisting in the breakdown of a node. This means that in most of the cases the failure is perfectly tolerated and reabsorbed by the system. However, there is always a finite, although very small, probability that the failure triggers an avalanche mechanism, collapsing the whole network.

As examples from the real world we study a network of the Internet (at the autonomous system level [1,30]) with $N=6474$ nodes and $K=12567$ arcs taken from Ref. [31], and the electrical power grid of the western United States from Ref. [32] having $N=4941$ and $K=6592$. Although the Internet exhibits a power law degree distribution (as for BA scale-free networks) while the electrical power grid has an exponential degree distribution (as for ER random graphs), we have checked that both the networks considered are very hetereogeneous from the point of view of the loads on nodes. In the insets of Fig. 3 and Fig. 4 we report $N(l)$, the number of nodes with a load larger than l, as a function of l: the straight lines indicate that the load distribution is consistent with a power-law with exponents, respectively, of 1.80 and 1.75. In the same figures we report the value of the efficiency

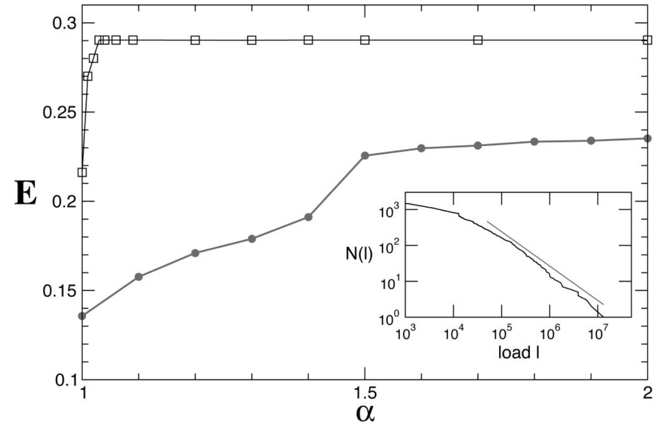

FIG. 3. Cascading failure in the Internet. The network considered is taken from Ref. [31]. For each value of α we report the efficiency E after the cascade triggered by the removal of a node chosen at random (squares), or by the removal of the node with the largest load (full circles). The curve reported for random removals is an average over ten different nodes. In the inset we plot the cumulative node load distribution.

after the cascade triggered by random failures and load-based failures. Due to the presence of a few nodes with an extremely high initial load, the figures show a large range of α where the network is stable against random failures and is vulnerable with respect to the breakdown of the most loaded nodes. Although the latter events have a very low probability, their occurrence may collapse the entire systems with a large effect on our life. These results are a possible explanation of the mechanism producing the experimentally observed Internet congestion collapses and the power blackouts. A small initial shock, such as the breakdown of an Internet router or of an electrical substation or line, may trigger avalanche mechanisms affecting a considerable fraction of a network that for years was proven to be stable with respect to similar shocks. As an example, if the electric power grid of the western United States of Fig. 4 works with a tolerance $\alpha=1.1$ ($\alpha=1.5$), a case in which the system is stable with respect to the failure of most of its nodes, the removal of a special node, the one with highest initial load, produces a drop of 30% (15%) of its efficiency.

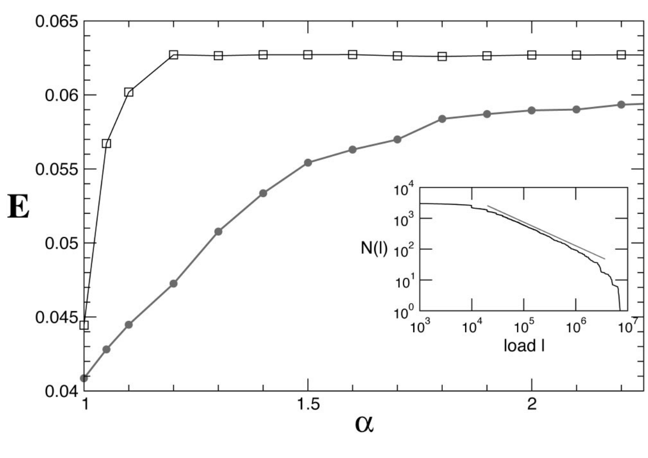

FIG. 4. Cascading failure in the electrical power grid of the western United States from Ref. [32]. Same plot as in Fig. 3.

Summing up, in this paper we have introduced a simple model to explain why large but rare cascade triggered by small initial shocks are present in most of the complex communication/transportation networks that are the basic components of our lives. The model is based on a dynamical redistribution of the flow triggered by the initial breakdown of a component of the system. The results show that the breakdown of a single node is sufficient to affect the efficiency of a network up to the collapse of the entire system if the node is among the ones with the largest load. This is particularly important for networks with a highly hetereogeneous distribution of node loads such as BA scale-free networks, but also real-world networks such as the Internet and electrical power grids. Our results show that it is only the breakdown of a selected minority of the nodes that can trigger the collapse of the system. It is also true that for the majority of the nodes nothing harmful happens, which leads us to the erroneous belief that our communication/transportation networks are safe. Therefore, it should be advisable to take into proper account, in the design of any complex network, the cascading failures effects analyzed here.

We thank D.J. Watts for the U.S. power-grid data from Ref. [32] and A. Rapisarda for useful comments.

[1] S.N. Dorogovtesev and J.F.F. Mendes, *Evolution of Networks* (Oxford University Press, Oxford, 2003).
[2] S.H. Strogatz, Nature (London) **410**, 268 (2001).
[3] V. Jacobson, Comput. Commun. Rev. **18**, 314 (1988).
[4] R. Guimerà, A. Arenas, A. Díaz-Guilera, and F. Giralt, Phys. Rev. E **66**, 026704 (2002).
[5] B.A. Carreras, D.E. Newman, I. Dolrou, and A.B. Poole, in Proceedings of Hawaii International Conference on System Sciences, Maui, Hawaii, 2000 (unpublished).
[6] M.L. Sachtjen, B.A. Carreras, and V.E. Lynch, Phys. Rev. E **61**, 4877 (2000).
[7] J. Glanz and R. Perez-Pena, ‘‘90 Seconds That Left Tens of Millions of People in the Dark,’’ New York Times, August 26, 2003.
[8] D.J. Watts, Proc. Natl. Acad. Sci. U.S.A. **99**, 5766 (2002).
[9] R. Albert, H. Jeong, and A.-L. Barabási, Nature (London) **406**, 378 (2000); **409**, 542(E) (2001).
[10] P. Holme, B.J. Kim, C.N. Yoon, and S.K. Han, Phys. Rev. E **65**, 056109 (2002).
[11] P. Crucitti, V. Latora, M. Marchiori, and A. Rapisarda, Physica A **320**, 622 (2003).
[12] M. Girvan and M.E.J. Newman, Proc. Natl. Acad. Sci. U.S.A. **99**, 8271 (2002).
[13] A.E. Motter, T. Nishikawa, and Y. Lai, Phys. Rev. E **66**, 065103 (2002).
[14] A.E. Motter and Y. Lai, Phys. Rev. E **66**, 065102(R) (2002).
[15] Y. Moreno, R. Pastor-Satorras, A. Vázquez, and A. Vespignani, Europhys. Lett. **62**, 292 (2003).
[16] Y. Moreno, J.B. Gomez, and A.F. Pacheco, Europhys. Lett. **58**, 630 (2002).
[17] V. Latora and M. Marchiori, Phys. Rev. Lett. **87**, 198701 (2001).
[18] S. Wasserman and K. Faust, *Social Networks Analysis* (Cambridge University Press, Cambridge, 1994).
[19] The model can be easily generalized to directed graphs.
[20] In the case of an unweighted graph e_{ij} is 1 if there is an arc joining node i to node j, and 0 otherwise.
[21] J. Smith, Commun. ACM **31**, 1202 (1988).
[22] R. Jain, *The Art of Computer Systems Performance Analysis* (Wiley, New York, 1991).
[23] The harmonic composition of N numbers $x_1, x_2, \ldots, x_N$ is defined as $[\Sigma_{i=1}^{N}(1/x_N)]^{-1}$.
[24] K.-I. Goh, B. Kahng, and D. Kim, Phys. Rev. Lett. **87**, 278701 (2001).
[25] Another possibility which gives the same results is to set $C_i = \alpha \max_j L_j(0) \forall i = 1,2,\ldots,N$.
[26] As an initial damage to the network here we consider the removal of a node, although the dynamics of the model can be started by the removal of one arc if one wants to simulate, for instance, the breakdown of a line in an electrical power grid.
[27] P. Erdös and A. Rényi, Publ. Math. (Debrecen) **6**, 290 (1959).
[28] A.-L. Barabási and R. Albert, Science **286**, 509 (1999).
[29] For load-based removals in BA scale-free graphs we have found that α_c is independent of the size of the system. In fact, we have obtained $\alpha_c = 1.28 \pm 0.05$, 1.3 ± 0.05, 1.28 ± 0.05, respectively, for $N = 1500$, 2000, 2500. Such a result is similar to that which is obtained in a simpler model, the fiber-bundle model, studied in Ref. [16].
[30] R. Pastor-Satorras, A. Vázquez, and A. Vespignani, Phys. Rev. Lett. **87**, 258701 (2001).
[31] http://moat.nlanr.net/AS/Data/ASconnlist.20000102.946809601.
[32] D.J. Watts and S.H. Strogatz, Nature (London) **393**, 440 (1998).