Easily Repairable Networks: Reconnecting Nodes after Damage

Robert S. Farr,[1,2,*] John L. Harer,[3,†] and Thomas M. A. Fink[1,4,‡]
[1]London Institute for Mathematical Sciences, 35a South Street, Mayfair, London W1K 2XF, United Kingdom
[2]Unilever R&D, Colworth Science Park, MK44 1LQ Bedford, United Kingdom
[3]Department of Mathematics, Duke University, Rayleigh-Durham, North Carolina 27708-0320, USA
[4]Centre National de la Recherche Scientifique, Paris 75248, France

(Received 8 May 2014; revised manuscript received 25 July 2014; published 26 September 2014)

We introduce a simple class of distribution networks that withstand damage by being repairable instead of redundant. Instead of asking how hard it is to disconnect nodes through damage, we ask how easy it is to reconnect nodes after damage. We prove that optimal networks on regular lattices have an expected cost of reconnection proportional to the lattice length, and that such networks have exactly three levels of structural hierarchy. We extend our results to networks subject to repeated attacks, in which the repairs themselves must be repairable. We find that, in exchange for a modest increase in repair cost, such networks are able to withstand any number of attacks.

DOI: 10.1103/PhysRevLett.113.138701 PACS numbers: 89.75.Hc, 89.75.Da, 89.75.Fb, 89.75.Kd

Infrastructure networks, such as distribution systems (power and water) [1,2], communication grids (cellular and World Wide Web), and transport networks (road and rail) [3], underlie much of modern society but are vulnerable to both natural phenomena (tsunamis, hurricanes, and earthquakes) and man-made threats (accidents, terrorism, and depletion). The loss is exacerbated by a tendency to collapse even though the fraction of each network damaged may be small [2,4,5]. Increasing their resilience is, therefore, a topic of practical concern [6].

To safeguard against the threat of disasters, many researchers have focused on robustness, whereby damage is absorbed due to internal redundancy. Robustness tends to be the strategy adopted by biological networks, such as the circulatory and nervous systems and leaf venation [7], which also must function reliably under environmental insults [8]. Robustness, however, is not the only strategy for increasing resilience. Recently, the European Union science agency appealed for "resilience concepts [that] take into account the necessity to … implement a substitution process in a crisis or disaster, aiming to deal with a lack of … capacities necessary to assume the continuity of basic functions and services, until recovery from negative effects and return to the normal situation" [9]. This "substitution process," or workaround, involves finding a short-term fix until the damaged part itself can be repaired [10]. When the typical cost of a workaround (averaged over all possible failure modes) is low, we say that the system is repairable. Repairability in this sense is the subject of this Letter.

For concreteness, we define the two resilience strategies as follows: A network is *robust* if, after an error in part of it, it is able (or more likely) to function normally on account of internal redundancy. A network is *repairable* if, after an error in part of it, it is able (or more likely) to function normally on account of intervention in other parts of it.

Before considering network resilience, we briefly outline optimal infrastructure networks. The simplest models take connectivity to be the sole determinant of function. Such models are appropriate for certain networks under light load, such as roads, electricity supply [11], and communication networks [12]. Here the network cost typically grows with the total length of the edges; if the cost equals the total length, optimal solutions are minimal spanning trees [11]. As a network becomes more heavily used, connectivity alone is no longer sufficient, and the capacity of the edges must also be considered. For models derived from resistor networks, efficiency translates to minimum power dissipation. If one associates a cost $R^{-\gamma}$ with each resistance R and specifies a total cost, then for planar networks with $\gamma < 1$, loopless geometries are known to be optimal in an unchanging environment [13,14]. More generally, this approach provides an explanation for fractal branching networks in biology, and ultimately for allometric growth laws [15].

In the presence of unexpected events, rather than a static environment, the traditional approach to guaranteeing function is building in redundancy (localized, such as extra paths from source to sink, or distributed, such as checksums in digital data). For models of the Internet, where simple connectivity suffices, the exponent of the distribution of node degrees [16] is the key parameter, for both random [17] and directed [18] attacks, and where failures may cascade [4,5,19]. For planar resistor networks with $\gamma < 1$, if there is local damage, or the loads or sources fluctuate, the designs that emerge from numerical optimization have redundancy with a hierarchical loop structure over many length scales [7,20,21].

The above work has focused on network robustness, whereas we want to understand network repairability. In doing so, we wish to capture the constraints on real infrastructure networks, in particular, the cost of capacity rather than just connectivity. Resistor networks account for this in a natural way, but analytic results often depend on large matrix computations [22–24]. Models of intermediate complexity, such as packet congestion [25,26], have therefore been proposed. Here we consider a related, but time-continuous, model, where the required capacity of each edge is equal to the number of downstream nodes it has to serve. An example is shown in Fig. 1(b), where houses are supplied with water from a central tower.

In this Letter, we do three things, each of which corresponds to a separate part. (i) We introduce a model of repairable networks in which a break at edge i can be mitigated by adding an edge j, and show that the cost of repair is the flux at i times the length, less 1, of the loop through i and j. (ii) We introduce the concept of an easily repairable network (ERN), which minimizes the expected cost of repair $\langle c \rangle$ after a single attack. We prove that ERNs have exactly three levels of structural hierarchy. (iii) When attacks are sufficiently numerous to strike the same place repeatedly, the repairs themselves must be easily repairable. To address this, we describe steady-state ERNs, able to withstand any number of attacks in exchange for a modest increase in $\langle c \rangle$.

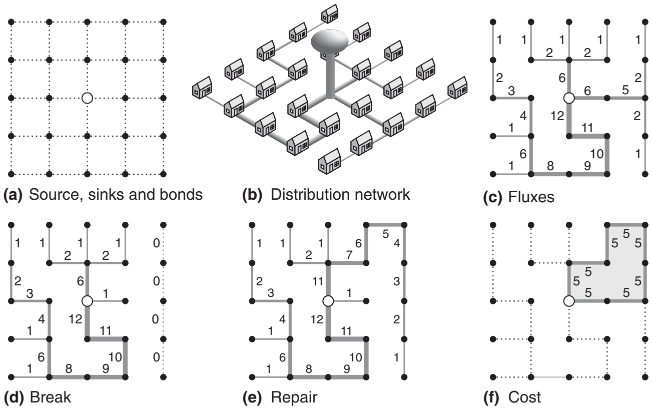

FIG. 1. A distribution network and the process of damage and repair. (b) Each house requires one unit of water per day from a tower. The flux into each house is equal to the number of houses downstream from it, plus one. Each day a random pipe breaks, and a new pipe must be built elsewhere to reconnect the cut off houses. Moreover, some pipes upstream and downstream from the new pipe must be resized to accommodate the change in demand. What network minimizes the expected total change in flux, or pipe altered, per break? (a) The source (open circle), nodes (closed circles), and bonds, which are possible edges (dashed lines). (b) The tree of edges connecting the nodes to the source. (c) The fluxes $\mathbf{J}_i$ at each edge. (d) An edge is broken, splitting the tree in two. (e) A new bond is made an edge, reconnecting the subtrees, and the updated fluxes $\mathbf{J}'_i$. (f) The changes in fluxes $|\mathbf{J}'_i - \mathbf{J}_i|$. They vanish everywhere except on the loop through the broken and repairing edges. The cost of repair is the flux through the broken edge times the remaining loop length: $c = 5 \times (8 - 1) = 35$.

Model of repairable networks.—For some systems, mistakes are reversible: the unintended change can be reverted or the broken part can be replaced like for like [27] (an on-screen typo, a dropped pen, a flat tire). More often than not, however, the error is irreversible or the repair time is unacceptably long (a printed typo, a missing ingredient, a jet engine failure). In these cases, it is necessary to find a workaround, that is, to restore the broken part by intervening in other parts.

As a model of the latter, suppose we need to continuously transport a commodity or information from a single source node (hereafter, source) to a collection of N downstream nodes (hereafter, nodes), each of which consumes the substance at a unit rate. We imagine that the source and N nodes form the vertices of an underlying lattice L, which we take to be either a square with side length l ($L_\square$, Fig. 1) or triangular with a hexagonal boundary with side length $l/2$ ($L_\triangle$). Note that l counts nodes not edges, and is odd. The total number of vertices is l^2 ($L_\square$) or $\frac{3}{4}(l^2 - 1) + 1$ ($L_\triangle$), so that $N_\square = l^2 - 1$ and $N_\triangle = \frac{3}{4}(l^2 - 1)$. The bonds of L are the possible conduits for transport, but not all of these bonds will be used in practice; the ones that are, we call edges. Each edge i carries a signed flux $\mathbf{J}_i$, such that the flux leaving the source is N, and the net flux into each of the N nodes is 1. The direction of positive flux is defined for each edge according to the flow in the original network. We assume (without loss of generality, as we shall see later) that the network of edges is a tree T [Fig. 1(b)]. This loop-free property allows us to assign fluxes on T unambiguously [Fig. 1(c)]. It also means there must be N edges, because into each node flows exactly one edge.

We now consider what happens when the network is broken and then repaired. In our model, a break consists of disabling a single edge i (i.e., edge i becomes just a bond; bonds have zero flux), which disconnects the tree T into two disjoint trees [Fig. 1(d)]. We then proceed to repair the network by adding an edge $j \neq i$ such that the new network is once more a tree, with a new set of fluxes $\{\mathbf{J}'_i\}$ [Fig. 2(e)]. We define the cost (intervention required) c_{ij} from this break-repair operation as the sum of the absolute changes in flux, but omitting the flux in the broken edge (we do not pay for the attack):

$$c_{ij} = -|\mathbf{J}_i| + \sum_{k\text{ an edge}} |\mathbf{J}'_k - \mathbf{J}_k|. \qquad (1)$$

This makes sense: if we imagine a fluid to flow along small unit flux pipes in parallel, or cars to travel along unit flux lanes of a highway, c is the number of pipes or lanes to add or remove (where reversing involves adding and removing). Note in particular that we pay if capacity is reduced; this is a valid strategy if there is an ongoing maintenance cost attached to the capacity of each edge (e.g., metabolism in biological tissues), so that it is rational to pay up front to eliminate this. This is the natural choice from a mathematical perspective, but may not always be physically realistic. We later consider the alternative model where spare capacity is free; here we find the costs are proportional to those in our original model.

In order to evaluate c_{ij}, suppose we have any two valid networks T' and T'' with the same source location. If we subtract the fluxes in T'' from those in T', the resultant pattern of fluxes must have no sources or sinks. Therefore, it must either vanish or be a sum of closed flux loops. Now consider our original network T. If we make a new network containing the original edges of T and the added edge k, this network is no longer a tree, but contains exactly one loop, of length d_{ij}. When we take the difference of the fluxes in the original and repaired networks, this is the only path that can have a nonzero flux [Fig. 1(f)], which must be the original flux at i (since after repair, the flux in this edge is zero). Consequently, the cost of repair is equal to the length of the loop through i and j less 1 (we omit the broken edge) times the flux at i: $c_{ij} = J_i(d_{ij} - 1)$, where $J_i = |\mathbf{J}_i|$. Now a broken network can be mended in a number of ways, corresponding to different choices of the bond j. We desire the *cheapest* repair, which is the one with the least cost:

$$c_i = J_i \min_j (d_{ij} - 1). \tag{2}$$

Hereafter, we take "repair" to mean "cheapest repair" and call the loop that arises from considering both the broken edge and the repair a dormant loop. Because the cost of repair c_i depends on the particular edge i that is broken, which is unknown, we want the expected cost of repair:

$$\langle c \rangle = \frac{1}{N} \sum_{i \text{an edge}} J_i \min_j (d_{ij} - 1). \tag{3}$$

What network minimizes $\langle c \rangle$? Figure 2 shows some examples of networks, with associated values for $\langle c \rangle$.

Easily repairable after one break.—The expected cost of repair $\langle c \rangle$ will be a minimum if two conditions are met: the expected absolute flux $\langle J \rangle \equiv N^{-1} \sum J_i$ in the original network is a minimum, and all the individual loop lengths d_{ij} are minimal. Without knowing to what extent the two are independent, we first ask, What networks minimize $\langle J \rangle$? To answer this, we note that every node must be connected to the source by some path, and along this path flows a unit of flux. There may also be confluent fluxes to other nodes flowing in the same edges, but we can conceptually treat this flux separately, even if in practice we do not keep track of all the individual streams. Accordingly, the expected flux $\langle J \rangle$ is $1/N$ times the sum of the lengths of all these streams. A minimum of $\langle J \rangle$ is achieved when each path is a geodesic of L between the node and the source. Such geodesics in general will not be unique, but at least one network composed of them must be a tree, because any loop can be broken by removing the distal edge and diverting the incoming streams into one side of this loop. If the source is at the lattice center, we find $\langle J \rangle_{\min} = l/2$ for a square lattice ($L_\square$) and $l/3$ for a triangular lattice ($L_\triangle$). Since the minimum loop length d_{ij} on $L_\square$ and $L_\triangle$ is 4 and 3, we find

$$\langle c \rangle_\square \geq 3l/2 \quad \text{and} \quad \langle c \rangle_\triangle \geq 2l/3. \tag{4}$$

Because both of these criteria can be met simultaneously, as Fig. 3 demonstrates, the above bounds can be achieved. We call the solutions, which are not unique, easily repairable networks.

The structure of ERNs on regular lattices is remarkable in that they have exactly three levels of hierarchy: connected to the source are primary arms (1-arms), from which branch secondary arms (2-arms), from which branch terminal hairs of length 1 (3-arms). This is in contrast to robust resistor networks, which have a hierarchical branching structure over many generations [7,28]. The steps to our proof are as follows. (i) A 1-arm must lie along a coordinate axis, or else the path to at least one of the nodes on that axis

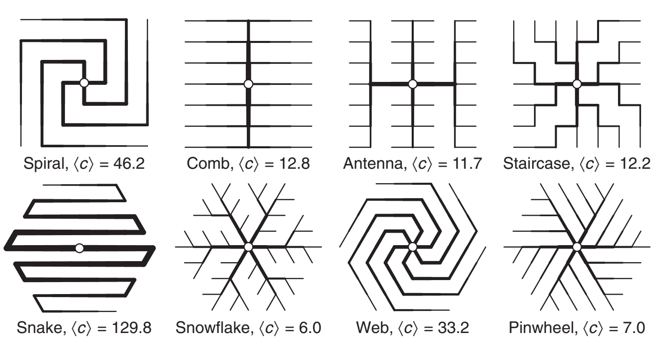

FIG. 2. Different strategies for building a distribution network and their expected cost of repair $\langle c \rangle$.

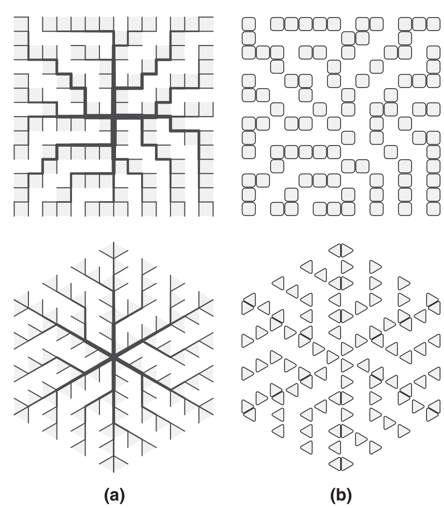

FIG. 3. Networks that are easy to repair after one break. The original networks (a) and the overlapping dormant loops (b).

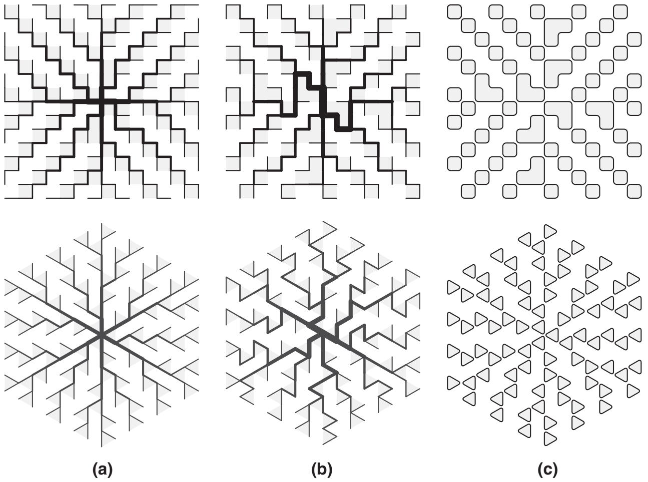

FIG. 4. Networks that are easy to repair after many breaks. The original networks (a), the networks after many break-repair cycles (b), and the independent dormant loops (c).

would include a bend. (ii) When a 2-arm splits, there can be only two daughter branches (including the 2-arm); there are only two directions that are geodesic: the two away from the origin. (iii) In any 2-arm split, the two daughter nodes cannot both split; if they did, a closed loop would be formed. (iv) The node joining two consecutive edges must have a split, or else the distal edge could not be part of a minimal dormant loop.

Easily repairable after many breaks.—Thus far we have considered networks subject to a single break. What happens when there is a series of breaks? As indicated earlier, we suppose that a break cannot be repaired immediately but is open to intervention later on (e.g., if there is another attack in the neighborhood). Clearly, repairing the first break (at a cost of $3l/2$ or $2l/3$) leaves the network fully functional. However, there is a more subtle effect: the repairs themselves may not be optimally repairable. Consequently, as the number of breaks increases, the network degrades, and the expected cost of repair $\langle c\rangle$ goes up.

How do we fix this? The structure of an ERN is such that (i) it is geodesic, thereby minimizing $\langle J\rangle$, and (ii) each of its edges is part of a minimal dormant loop, thereby minimizing d. Network degradation is due to (i') the failure of the repaired network to be geodesic, and (ii') the increase in dormant loop length. Now (i') is because a minimal latent loop can have 4 ($L_\square$) or 3 ($L_\triangle$) orientations, or spins, and not all of these are geodesic; since repair is performed on the only absent edge, eventually the spins will occur with equal frequency. And (ii') is the result of dormant loops that share an edge: when an edge belonging to two such loops is broken, repairing it coalesces them into one larger dormant loop. Therefore, to optimize networks under multiple breaks, we must ensure that the dormant loops are independent, having no edges in common. Figure 4 shows

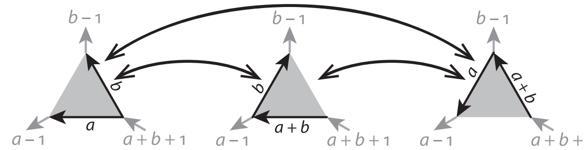

FIG. 5. The three spins of a dormant loop. Only the left spin is geodesic. Each break and repair sends the loop from one spin to another. The total flux along the geodesic loop is $a+b$. After many break-repair cycles, the total flux, averaged over all spins, is $\frac{4}{3}(a+b)$. Thus, at steady state, the expected flux $\langle J(\infty)\rangle$ is $4/3$ of that before any breaks $\langle J(0)\rangle$.

that for $L_\triangle$, this is achievable; for $L_\square$, there will always be at least a small fraction of dormant loops sharing an edge, and in fact, these dominate the change in $\langle c\rangle$ after many breaks. Figure 5 shows the calculation for the average repair cost per break after many breaks. The final result in the many-break limit is

$$\langle c\rangle_\square \to 7l/4 + O(1) \quad \text{and} \quad \langle c\rangle_\triangle \to 8l/9. \qquad (5)$$

In both cases, the networks are able to withstand any number of breaks in exchange for a modest increase in the expected repair cost. We call the solutions steady-state ERNs. Their design aims to achieve three properties (all being simultaneously possible for the triangular lattice): (i) the initial network is geodesic, (ii) the dormant loops are minimal, (iii) the dormant loops do not overlap.

We briefly consider the case where we do not have to pay to immediately reduce capacity. For an ERN on $L_\triangle$, the expected cost of repair is identical in both cases, because none of the fluxes change direction, and we do not pay for the break itself. For a steady-state ERN on $L_\triangle$, analysis similar to that in Fig. 5 (but now involving six states) reveals that asymptotically $\langle c\rangle_\triangle \to 5l/9$.

Conclusion.—Some aspects of repairability we have not addressed here, but briefly mention: First, if the edges are of unequal length, we can generalize our model by letting d_{ij} be the physical length of the loop rather than the number of edges in it. For small perturbations from the regular lattice L there will be no new solutions (since moving away from a solution optimal on L incurs a cost of order unity), but the perturbations will break degeneracy of the optimal solutions. Second, it is well known that in scale-free networks directed attacks against high-degree nodes can cause major damage. We have assumed attacks are undirected. To inflict the most damage, an attack would target the highest flux edges, namely, those adjacent to the source. In a triangular ERN, the cost of repair for such a directed attack is $3l/4$ times more than the undirected case. Third, while we only considered a single source, we expect multiple sources to merely partition L into Voronoi regions, without introducing added congestion [29]. Fourth, applying our model to other topologies, such as scale-free networks [16], may be possible, but would require their redesign to include pervasive loops.

We believe our model of distribution networks captures essential features of their real-world analogs in a form simple enough to be analytically tractable. The structure of ERNs and steady-state ERNs embodies useful design directions for engineering applications, such as resource distribution, smart electricity grids [30], and communication networks. More generally, it helps quantify the concept of repairability, and is a framework for understanding it as an alternative to robustness in achieving resilience.

This project was sponsored by the Department of Defense (Defense Threat Reduction Agency), and received support from the EU Framework 7 project GROWTHCOM (FP7-ICT-2013-10) and Boston Consulting Group.

*robert.farr@unilever.com
†harer@math.duke.edu
‡tf@lims.ac.uk

[1] R. Carvalho, L. Buzna, F. Bono, E. Gutiérrez, W. Just, and D. Arrowsmith, Phys. Rev. E **80**, 016106 (2009).
[2] R. Albert, I. Albert, and G. L. Nakarado, Phys. Rev. E **69**, 025103(R) (2004).
[3] A. Tero, S. Takagi, T. Saigusa, K. Ito, D. P. Bebber, M. D. Fricker, K. Yumiki, R. Kobayashi, and T. Nakagaki, Science **327**, 439 (2010).
[4] A. E. Motter and Y.-C. Lai, Phys. Rev. E **66**, 065102(R) (2002).
[5] A. E. Motter, Phys. Rev. Lett. **93**, 098701 (2004).
[6] See www.unisdr.org/we/inform/publications/26462.
[7] E. Katifori, G. J. Szöllösi, and M. O. Magnasco, Phys. Rev. Lett. **104**, 048704 (2010).
[8] A. S. Iyer-Pascuzzi, T. Jackson, H. Cui, J. J. Petricka, W. Busch, H. Tsukagoshi, and P. N. Benfey, Dev. Cell **21**, 770 (2011).
[9] European Commission, Horizon 2020: Disaster resilience: safeguarding and securing society (2014).
[10] For an example involving regulatory networks, see C. Campbell and R. Albert, BMC Syst. Biol. **8**, 53 (2014).
[11] O. Borůvka, Elektrotech. Obz. **15**, 153 (1926).
[12] B. M. Leiner, V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. Postel, L. G. Roberts, and S. S. Wolff, Commun. ACM **40**, 102 (1997).
[13] S. Bohn and M. O. Magnasco, Phys. Rev. Lett. **98**, 088702 (2007).
[14] M. Durand, Phys. Rev. Lett. **98**, 088701 (2007).
[15] G. West, J. H. Brown, and B. J. Enquist, Science **276**, 122 (1997).
[16] A.-L. Barabási and R. Albert, Science **286**, 509 (1999).
[17] R. Cohen, K. Erez, D. ben-Avraham, and S. Havlin, Phys. Rev. Lett. **85**, 4626 (2000).
[18] R. Cohen, K. Erez, D. ben-Avraham, and S. Havlin, Phys. Rev. Lett. **86**, 3682 (2001).
[19] M. L. Sachtjen, B. A. Carreras, and V. E. Lynch, Phys. Rev. E **61**, 4877 (2000).
[20] D. Hu and D. Cai, Phys. Rev. Lett. **111**, 138701 (2013).
[21] F. Corson, Phys. Rev. Lett. **104**, 048703 (2010).
[22] N. Rubido, C. Grebogi, and M. S. Baptista, Europhys. Lett. **101**, 68001 (2013).
[23] N. Rubido, C. Grebogi, and M. S. Baptista, Phys. Rev. E **89**, 012801 (2014).
[24] E. Lopez, S. Carmi, S. Havlin, S. V. Buldryrev, and H. E. Stanley, Physica (Amsterdam) **224D**, 69 (2006).
[25] G. Yan, T. Zhou, B. Hu, Z.-Q. Fu, and B.-H. Wang, Phys. Rev. E **73**, 046108 (2006).
[26] Y.-H. Chen, B.-H. Wang, L.-C. Zhao, C. Zhou, and T. Zhou, Phys. Rev. E **81**, 066105 (2010).
[27] B. Hu, F. Li, and H.-S. Zhou, Chin. Phys. Lett. **26**, 128901 (2009).
[28] G. Caldarelli, *Scale-Free Networks: Complex Webs in Nature and Technology* (Oxford University Press, Oxford, 2007).
[29] S. Carmi, Z. Wu, S. Havlin, and H. E. Stanley, Europhys. Lett. **84**, 28005 (2008).
[30] M. Amin, Nature (London) **499**, 145 (2013).