

Network dismantling

Alfredo Braunstein^{a,b,c}, Luca Dall'Asta^{a,c}, Guilhem Semerjian^d, and Lenka Zdeborová^{e,1}

^aDepartment of Applied Science and Technology, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Turin, Italy; ^bHuman Genetics Foundation, Via Nizza 52, 10126 Turin, Italy; ^cCollegio Carlo Alberto, Via Real Collegio 30, 10024 Moncalieri, Italy; ^dLaboratoire de Physique Théorique, Ecole Normale Supérieure, Paris Sciences et Lettres (PSL) Research University, Sorbonne Universités, Université Pierre et Marie Curie (UPMC), CNRS, 75005 Paris, France; and ^eInstitut de Physique Théorique, CNRS, Commissariat à l'Energie Atomique (CEA) and Université Paris Saclay, 91191 Gif-sur-Yvette, France

Edited by Giorgio Parisi, University of Rome, Rome, Italy, and approved September 16, 2016 (received for review March 29, 2016)

We study the network dismantling problem, which consists of determining a minimal set of vertices in which removal leaves the network broken into connected components of subextensive size. For a large class of random graphs, this problem is tightly connected to the decycling problem (the removal of vertices, leaving the graph acyclic). Exploiting this connection and recent works on epidemic spreading, we present precise predictions for the minimal size of a dismantling set in a large random graph with a prescribed (light-tailed) degree distribution. Building on the statistical mechanics perspective, we propose a three-stage Min-Sum algorithm for efficiently dismantling networks, including heavy-tailed ones for which the dismantling and decycling problems are not equivalent. We also provide additional insights into the dismantling problem, concluding that it is an intrinsically collective problem and that optimal dismantling sets cannot be viewed as a collection of individually well-performing nodes.

graph fragmentation | message passing | percolation | random graphs | influence maximization

A network (a graph G in the discrete mathematics language) is a set V of N entities called nodes (or vertices), along with a set E of edges connecting some pairs of nodes. In a simplified way, networks are used to describe numerous systems in very diverse fields, ranging from social sciences to information technology or biological systems (reviews are in refs. 1 and 2). Several crucial questions in the context of network studies concern the modifications of the properties of a graph when a subset S of its nodes is selected and treated in a specific way. For instance, how much does the size of the largest connected component of the graph decrease if the vertices in S (along with their adjacent edges) are removed? Do the cycles survive this removal? What is the outcome of the epidemic spreading if the vertices in S are initially contaminated, constituting the seed of the epidemic? On the contrary, what is the influence of a vaccination of nodes in S preventing them from transmitting the epidemic? It is relatively easy to answer these questions when the set S is chosen randomly, with each vertex being selected with some probability independently. Classical percolation theory is nothing but the study of the connected components of a graph in which some vertices have been removed in this way.

A much more interesting case is when the set S can be chosen in some optimal way. Indeed, in all applications sketched above, it is reasonable to assign some cost to the inclusion of a vertex in S : vaccination has a socioeconomic price, incentives must be paid to customers to convince them to adopt a new product in a viral marketing campaign, and incapacitating a computer during a cyber attack requires resources. Thus, one faces a combinatorial optimization problem: the minimization of the cost of S under a constraint on its effect on the graph. These problems thus exhibit both static and dynamic features, the former referring to the combinatorial optimization aspect and the latter referring to the definition of the cost function itself through a dynamical process.

In this paper, we focus on the existence of a giant component in a network: that is, the largest component containing a positive fraction of the vertices (in the $N \rightarrow \infty$ limit). On the one hand, the existence of a giant component is often necessary for the network to fulfill its function (e.g., to deliver electricity or

information bits or ensure possibility of transportation). An adversary might be able to destroy a set of nodes with the goal of destroying this functionality. It is thus important to understand what an optimal attack strategy is, possibly as a first step in the design of optimal defense strategies. On the other hand, a giant component can propagate an epidemic to a large fraction of a population of nodes. Interpreting the removal of nodes as the vaccination of individuals who cannot transmit the epidemic anymore, destroying the giant component can be seen as an extreme way of organizing a vaccination campaign (3, 4) by confining the contagion to small connected components [less drastic strategies can be devised using specific information about the epidemic propagation model (5, 6)]. Another related application is influence maximization as studied in many previous works (7–9). In particular, optimal destruction of the giant component is equivalent to selection of the smallest set of initially informed nodes needed to spread the information into the whole network under a special case of the commonly considered model for information spreading (7–9).

To define the main subject of this paper more formally, following ref. 10, we call S a C -dismantling set if its removal yields a graph with the largest component that has size (in terms of its number of nodes) at most C . The C -dismantling number of a graph is the minimal size of such a set. When the value of C is either clear from the context or not important for the given claim, we will simply talk about dismantling. Typically, the size of the largest component is a finite fraction of the total number of nodes N . To formalize the notion of destroying the giant component, we will consider the bound C on the size of the connected components of the dismantled network to be such that $1 \ll C \ll N$. It should be noted that we defined dismantling in terms of node removal; it could be rephrased in terms of edge removal (11), which turns out to be a much easier problem. The dismantling problem is also referred to as fragmentability of graphs in graph theory literature (12–14) and optimal percolation in ref. 15.

Significance

Many systems of interest can be represented by a network of nodes connected by edges. In many circumstances, the existence of a giant component is necessary for the network to fulfill its function. Motivated by the need to understand optimal attack strategies, optimal spread of information, or immunization policies, we study the network dismantling problem (i.e., the search for a minimal set of nodes in which removal leaves the network broken into components of subextensive size). We give the size of the optimal dismantling set for random networks, propose an efficient dismantling algorithm for general networks that outperforms by a large margin existing strategies, and provide various insights about the problem.

Author contributions: A.B., L.D., G.S., and L.Z. designed research, performed research, contributed new reagents/analytic tools, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. Email: lenka.zdeborova@cea.fr.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1605083113/-DCSupplemental.

Table 1. The (1RSB) cavity predictions for the decycling number of ER random graphs of average degree d and the decycling number reached by the Min-Sum algorithm on graphs of size $N = 10^7$ nodes

d	$\theta_{\text{dec}}(d)$	$\theta_{\text{dec}}^{\text{MS}}(d)$
1.5	0.0125	0.0135
2.5	0.0912	0.0936
3.5	0.1753	0.1782
5	0.2789	0.2823

shall call q light tailed when this is the case), then there is actually an equality between these two parameters, $\theta_{\text{dis}}(q) = \theta_{\text{dec}}(q)$.

The first claim follows directly from the above observation on the decycling number of forests. After a decycling set S of G has been found, one can add to S additional vertices to turn it into a C -dismantling set, the additional cost being bounded as $\theta_{\text{dis}}(G, C) \leq \theta_{\text{dec}}(G) + 1/(C+1)$. Taking averages of this bound and the limit $C \rightarrow \infty$ after $N \rightarrow \infty$ yields directly i .

To justify our second claim, we consider a C -dismantling set S of a graph G . To turn S into a decycling set, we need to add additional vertices to break the cycles that might exist in $G \setminus S$. The lengths of these cycles are certainly smaller than C , and removing at most one vertex per cycle is enough to break them. We can thus write $\theta_{\text{dec}}(G) \leq \theta_{\text{dis}}(G, C) + n_C(G)/N$, with $n_C(G)$ denoting the number of cycles of G of length at most C . We recall that the existence of a second moment of q implies that $n_C(G)$ remains bounded when $N \rightarrow \infty$ with C fixed. Considering the limit $N \rightarrow \infty$ and property i , property ii follows.

Network Decycling

In this section, we shall explain the results on the decycling number of random graphs that we obtained via statistical mechanics methods and how they can be exploited to build an efficient heuristic algorithm for decycling arbitrary graphs.

Testing the Presence of Cycles in a Graph. The 2-core of a graph G is its largest subgraph of minimal degree 2; it can be constructed by iteratively removing isolated nodes and leaves (vertices of degree 1) until either all vertices have been removed or all remaining vertices have degree at least 2. It is easy to see that a graph contains cycles if and only if its 2-core is nonempty. To decide if a subset S is decycling, we remove the nodes in S and perform this leaf removal on the reduced graph. To formalize this procedure, we introduce binary variables $x_i^t(S) \in \{0, 1\}$ on each vertex $i \in V$ of the graph, t being a discrete time index. At the starting time $t=0$, one marks the initially removed vertices by setting $x_i^0(S) = 1$ if $i \in S$ and 0 otherwise, and let the x variables evolve in time according to

$$x_i^{t+1}(S) = \begin{cases} 1 & \text{if } x_i^t(S) = 1, \\ \mathbb{I} \left[\sum_{j \in \partial i} (1 - x_j^t(S)) \leq 1 \right] & \text{if } x_i^t(S) = 0, \end{cases} \quad [3]$$

where $\partial i = \{j : (ij) \in E\}$ denotes the local neighborhood of vertex i , and \mathbb{I} denotes the indicator function (that is, one if its argument is true and zero otherwise). One can check that the x_i s are monotonous in time (they can only switch from zero to one); hence, they admit a limit $x_i^*(S)$ when $t \rightarrow \infty$. At this fixed point, $x_i^*(S) = 0$ if and only if i is in the 2-core of $G \setminus S$; hence, the sufficient and necessary condition for S to be a decycling set of G is $x_i^*(S) = 1$ for all vertices i .

Note that the leaf removal procedure can be equivalently viewed as a particular case of the linear threshold model of epidemic propagation or information spreading. By calling a removed vertex infected (or informed), one sees that the infection (or information) of node i occurs whenever the number of its infected (or informed) neighbors reaches its degree minus one. This equivalence, which was already exploited in refs. 15 and 23,

allows us to build on previous works on minimal contagious sets (20, 21, 23) and influence maximization (7–9).

Optimizing the Size of Decycling Sets. From the point of view of statistical mechanics, it is natural to introduce the following probability distribution over the subsets S to find the optimal decycling sets of a given graph

$$\hat{\eta}(S) = \frac{1}{Z(\mu)} e^{\mu|S|} \prod_{i \in V} \mathbb{I}[x_i^*(S) = 1], \quad [4]$$

where $|S|$ denotes the number of vertices in S , μ is a real parameter to be interpreted as a chemical potential (or an inverse temperature), and the partition function $Z(\mu)$ normalizes this probability distribution. From the preceding discussion, this measure gives a positive probability only to decycling sets, and their minimal size can be obtained as the ground-state energy in the zero-temperature limit:

$$\theta_{\text{dec}}(G) = \frac{1}{N} \lim_{\mu \rightarrow -\infty} \frac{1}{\mu} \ln Z(\mu). \quad [5]$$

The computation of this partition function remains at this point a difficult problem; in particular, the variables x_i^* depend on the choice of S in a nonlocal way. One can get around this difficulty in the following way: because the evolution of x_i^t is monotonous in time, it can be completely described by a single integer, $t_i(S) = \min\{t : x_i^t(S) = 1\}$, the time at which i is removed in the parallel evolution described above. Note that $t_i(S) = 0$ if and only if $i \in S$ and $t_i(S) > 0$ otherwise. We use the natural convention $\min \emptyset = \infty$; hence, the nodes i in the 2-core of $G \setminus S$ are precisely those with an infinite removal time $t_i(S) = \infty$. The crucial advantage of this equivalent representation in terms of the activation times is its locality along the graph. Indeed, the dynamical evolution rule (Eq. 3) can be rephrased as static equations linking the times t_i on neighboring vertices:

$$t_i(S) = \begin{cases} 0 & \text{if } i \in S, \\ \phi_i(\{t_j\}_{j \in \partial i}) & \text{if } i \in V \setminus S, \end{cases} \quad [6]$$

$$\text{with } \phi_i(\{t_j\}_{j \in \partial i}) = 1 + \max_2(\{t_j(S)\}_{j \in \partial i}), \quad [7]$$

where we denote \max_2 the second largest of the arguments [reordering them as $t_1 \geq t_2 \geq \dots \geq t_n$, one defines $\max_2(t_1, \dots, t_n) = t_2$]. In the leaf removal procedure, one vertex is removed in the first step after the time at which all but one of its neighbors has been




Fig. 1. Fraction of nodes in the largest component as a function of the fraction of removed nodes for an ER random graph of average degree $d=3.5$ and size $N=78,125$. We compare the result of our Min-Sum algorithm (MS) with random (RND), adaptive largest degree (DEG), adaptive EC, adaptive CI centrality, and SA.

time t_i of site i . From the solution of this combinatorial optimization problem, we construct one of the minimal decycling sets S by including vertex i in S if and only if $t_i^* = 0$. It remains now to find a good approximation for h_i ; we compute it by the Min-Sum algorithm, which corresponds to the $\mu \rightarrow -\infty$ limit of BP and is similarly based on the exchange of messages $h_{ij}(t_i, t_j)$ between neighboring vertices, an analog of $\eta_{ij}(t_i, t_j)$, but interpreted as a minimal cost instead of a probability. We defer to [SI Appendix](#) for a full derivation and implementation details, stating here the final equations.

$$h_i(t_i) = \psi_i(t_i) + \sum_{k \in \partial i} L_{ki}(t_i) + M_i(t_i), \quad [10a]$$

$$h_i(0) = \psi_i(0) + \sum_{k \in \partial i} R_{ki}(0), \quad [10b]$$

$$M_i(t_i) = \min \left\{ 0, \min_{k \in \partial i} \{R_{ki}(t_i) - L_{ki}(t_i)\} \right\}, \quad [11]$$

for $T \geq t_i > 0$, where L_{ij} , R_{ij} , M_{ij} , h_{ij}^0 , and h_{ij}^1 form a solution of the following system of fixed point equations for messages defined on each directed edge $i \rightarrow j$ of the graph:

$$L_{ki}(t_i) = \min_{t_k < t_i} h_{ki}^0(t_k), \quad [12a]$$

$$R_{ki}(t_i) = \min \left\{ h_{ki}^0(t_i), \min_{t_k > t_i} h_{ki}^1(t_k) \right\}, \quad [12b]$$

$$M_{ij}(t_i) = \min \left\{ 0, \min_{k \in \partial \setminus j} \{R_{ki}(t_i) - L_{ki}(t_i)\} \right\}, \quad [12c]$$

$$h_{ij}^0(t_i) \propto \psi_i(t_i) + \sum_{k \in \partial \setminus j} L_{ki}(t_i), \quad [12d]$$

$$h_{ij}^1(t_i) \propto \psi_i(t_i) + \sum_{k \in \partial \setminus j} L_{ki}(t_i) + M_{ij}(t_i), \quad [12e]$$

$$h_{ij}^0(0) \propto \psi_i(0) + \sum_{k \in \partial \setminus j} R_{ki}(0), \quad [12f]$$

where \propto includes now an additive normalization constant. An intuitive interpretation of all of these quantities and equations is provided in [SI Appendix](#); let us only mention at this point that the message $h_{ij}^0(t_i)$ [respectively $h_{ij}^1(t_i)$] is the minimum feasible cost on the connected component of i in $G \setminus j$ under the condition that i is removed at time t_i in the original graph, assuming that j is not removed yet (respectively assuming that j is already removed from G).

This system can be solved efficiently by iteration. The computation of one iteration takes $O(|E|T)$ elementary (+, −, ×, min) operations, where $|E|$ denotes the number of edges of the graph, and a relatively small number of iterations is usually sufficient to reach convergence. In principle, one should take the cutoff T on the removal times to be greater than N to solve the decycling problem. We found, however, that using large but finite values of T (i.e., constraining the diameter of the tree components after the node removal) did not increase extensively the size of the decycling set; in the simulations presented below, we used $T = 35$. Note that our algorithm is very flexible, and many variations can be implemented by appropriate modifications of the cost function. For example, we exploited the possibility to forbid the removal of certain marked nodes i by setting $\psi_i(t_i) = \infty$ for them.

Results for Dismantling

Results on Random Graphs. The outcome of our algorithm applied to an ER random graph of average degree 3.5 is presented in Fig. 1.

Here, the red circle corresponds to the output of its first stage (decycling with Min-Sum), which yields, after the removal of a fraction 0.1781 of the nodes, an acyclic graph in which the largest components contain a fraction 0.032 of the vertices. The red line corresponds to the second stage, which further reduces the size of the largest component by greedily breaking the remaining trees. We compare with simulated annealing (SA; black circle) as well as several incremental algorithms that successively remove the nodes with the highest scores, where the score of a vertex is a measure of its centrality. Other than a trivial function that gives the same score to all vertices [hence removing the vertices in random order (RND)] and the score of a vertex equal to its degree, we used the eigenvector centrality (EC) measure and the recently proposed CI measure (15). We used all of these heuristics in an adaptive way, recomputing the scores after each removal. Additional details on all of these algorithms can be found in [SI Appendix](#).

We see from Fig. 1 that the Min-Sum algorithm outperforms the others by a considerable margin: it dismantles the graph using 13% fewer nodes than the CI method. The Monte Carlo-based SA algorithm performs rather well but is considerably slower than all of the others.

In Fig. 2, we zoom in on the results of the second stage of our algorithm and perform a finite size scaling analysis, increasing the size of the dismantled graphs up to $N = 10^6$. In this way, we identify a threshold for decycling (and thus, for dismantling) by the Min-Sum algorithm that converges toward the value $\theta_{\text{dec}}^{\text{MS}} \approx 0.1782$, which is close but not equal to the theoretical prediction of the 1RSB calculation $\theta_{\text{dec}}^{\text{1RSB}} \approx 0.1753$ (vertical arrow in Fig. 2). Fig. 2, *Inset* shows a remarkable scaling that indicates that the size of the largest component after dismantling by removing a given fraction of nodes does not depend on the graph size.

Combinatorial optimization problems typically exhibit a very large degeneracy of their (quasi)optimal solutions. We performed a detailed statistical analysis of the quasioptimal dismantling sets constructed by our algorithm, exploiting the fact that the Min-Sum algorithm finds different decycling sets for different realizations of the random tie-breaking noise $\varepsilon_i(t_i)$.

For a given ER random graph of average degree 3.5 and size $N = 78,125$, we ran the algorithm for 1,000 different realizations of the tie-breaking noise $\varepsilon_i(t_i)$ and obtained 1,000 different decycling sets, all of which had sizes within 40 nodes of one another. Randomly chosen pairs among these 1,000 decycling sets coincided, on average, on 82% of their nodes. For each node, we computed its frequency of appearance among the 1,000 decycling sets that we obtained. We then ordered nodes by this frequency and plotted the frequency as a function of this




Fig. 4. Fraction C/N taken by the largest component in the Twitter network achieved after removing a fraction S/N of nodes using the Min-Sum (MS) algorithm and the adaptive versions of CI and EC measures. The red circle marks the result obtained by decycling using MS (followed by the curve from the optimal tree-breaking process). The branches at lower values of S/N are obtained after the application of the RG strategy from the graph obtained when the largest component has $C = 100$ nodes. The black circle denotes the dismantling fraction obtained by SA.

ordering in Fig. 3. We see that some nodes appear in almost all found sets, that about 60% of nodes do not appear in any sets, and that a large portion of nodes appears only in a fraction of the decycling sets. We compare the frequencies of nodes belonging to one typical set found by Min-Sum and the CI heuristics.

An important question to ask about dismantling sets is whether they can be thought of as a collection of nodes that are in some sense good spreaders or whether they are a result of highly correlated optimization. We use the result of the previous experiment and remove the nodes that appeared most often (i.e., have the highest frequencies in Fig. 3). If the nature of dismantling was additive rather than collective, then such a choice should further decrease the size of the discovered dismantling set. This scenario is not what happens; with this strategy, we need to remove 20.1% of nodes to dismantle the graph compared with the 17.8% of nodes found systematically by the Min-Sum algorithm. From this observation, we conclude that dismantling is an intrinsically collective phenomenon, and one should always speak of the full set rather than a collection of influential spreaders.

We also studied the degree histogram of nodes that the Min-Sum algorithm includes in the dismantling sets and saw that, as expected, most of the high-degree nodes belong to most of the dismantling sets. Each of the dismantling sets also included some nodes of relatively low degrees; for instance, for an ER random graph of average degree $d=6$ and size 5^7 , a typical decycling set found by the Min-Sum algorithm has around 460 (i.e., around 17% of the decycling set) nodes of degree 4 or lower. To assess the importance of low-degree nodes for dismantling, we ran the Min-Sum algorithm under the constraint that only nodes of degree at least 5 can be removed, and we find decycling sets almost as small (only about 50 nodes; i.e., 0.2% larger) as without this constraint. From this observation, we conclude that none of the low-degree nodes (even those with high CI centrality) are indispensable for dismantling, going against a highlight claim of ref. 15.

More General Graphs. Up to this point, our study of dismantling relies crucially on the relation to decycling. For light-tailed random graphs, these two problems are essentially asymptotically equivalent. However, for arbitrary graphs that contain many small cycles, the decycling number can be much larger than the dismantling one. We argue that, from the algorithmic point of view, decycling still provides a very good basis for dismantling. For instance, consider a portion of $N=532,000$ nodes of the Twitter network already analyzed in ref. 15. The decycling solution found by Min-Sum improves considerably the results obtained with the CI and EC heuristics (Fig. 4).

In a network that contains many short cycles, decycling removes a large proportion of nodes expressly to destroy these short cycles.

Many of these nodes can be put back without increasing the size of the largest component. For this reason, we introduce a reverse greedy (RG) procedure, in which starting from a dismantled graph with dismantling set S , maximum component size C , and a chosen target value $C' > C$ for the maximum allowed component size, removed nodes are iteratively reinserted. At each step, among all removed nodes, the one that ends up in the smallest connected component is chosen for reinsertion (details are in *SI Appendix*). The computational cost of this operation is bounded by $k_{\max} C' \log(k_{\max} C')$, where k_{\max} is the maximal degree of the graph; the update cost is thus typically sublinear in N .

In graphs where decycling is an optimal strategy for dismantling, such as the random graphs, a vanishing fraction of nodes can be reinserted by the RG procedure before the size of the largest component starts to grow steeply. For real world networks, the RG procedure reinserts a considerable number of nodes, negligibly altering the size of the largest component. For the Twitter network in Fig. 4, the improvement obtained by applying the RG procedure is impressive: 32% fewer nodes for the CI method and 20% fewer nodes for the Min-Sum algorithm, which ends up being the best solution that we found, removing only 3.4% of nodes to dismantle into components smaller than $C=1,000$ nodes. RG makes it possible to reach, and even improve, the best result obtained with SA that solves the dismantling problem directly and is not affected by the presence of short loops (*SI Appendix* has details on SA). Qualitatively similar results are achieved on other real networks [e.g., on the YouTube network with 1.13 million nodes (27), the best dismantling set that we found with Min-Sum + RG included 4.0% of nodes; this result is a 22% improvement with respect to the CI heuristics].

The RG procedure is introduced as a heuristic that provides a considerable improvement for the examples that we treated. The theoretical results of this paper are valid only for classes of graphs that do not contain many small cycles, and hence, our theory does not provide a principled derivation or analysis of the RG procedure. This point is an interesting open direction for future work. More detailed study (both theoretical and algorithmic) of dismantling of networks for which decycling is not a reasonable starting point is an important direction of future work.

ACKNOWLEDGMENTS. A.B. and L.D. acknowledge support by Fondazione Cassa di Risparmio di Torino (CRT) under the initiative La Ricerca dei Talenti. L.D. acknowledges the European Research Council Grant 267915.

- Boccaletti S, Latora V, Moreno Y, Chavez M, Hwang DU (2006) Complex networks: Structure and dynamics. *Phys Rep* 424(4):175–308.
- Barrat A, Barthélemy M, Vespignani A (2008) *Dynamical Processes on Complex Networks* (Cambridge Univ Press, Cambridge, UK).
- Pastor-Satorras R, Vespignani A (2002) Immunization of complex networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 65(3 Pt 2A):036104.
- Cohen R, Havlin S, Ben-Avraham D (2003) Efficient immunization strategies for computer networks and populations. *Phys Rev Lett* 91(24):247901.
- Britton T, Janson S, Martin-Löf A (2007) Graphs with specified degree distributions, simple epidemics, and local vaccination strategies. *Adv Appl Probab* 39(4):922–948.
- Altarelli F, Braunstein A, Dall'Asta L, Wakeling JR, Zecchina R (2014) Containing epidemic outbreaks by message-passing techniques. *Phys Rev X* 4(2):021024.
- Kempe D, Kleinberg J, Tardos E (2003) Maximizing the spread of influence through a social network. *KDD* 03:137–146.
- Chen N (2008) On the approximability of influence in social networks. *SODA* 08: 1029–1037.
- Dreyer PA, Jr, Roberts FS (2009) Irreversible k -threshold processes: Graph-theoretical threshold models of the spread of disease and of opinion. *Discrete Appl Math* 157(7): 1615–1627.
- Janson S, Thomason A (2008) Dismantling sparse random graphs. *Comb Probab Comput* 17(02):259–264.
- Bradonjic M, Molloy M, Yan G (2013) Containing viral spread on sparse random graphs: Bounds, algorithms, and experiments. *Internet Math* 9(4):406–433.
- Edwards K, McDiarmid C (1994) New upper bounds on harmonious colorings. *J Graph Theory* 18(3):257–267.
- Edwards K, Farr G (2001) Fragmentability of graphs. *J Comb Theory Ser B* 82(1):30–37.
- Edwards K, Farr G (2008) Planarization and fragmentability of some classes of graphs. *Discrete Math* 308(12):2396–2406.
- Morone F, Makse HA (2015) Influence maximization in complex networks through optimal percolation. *Nature* 524(7563):65–68.
- Albert R, Jeong H, Barabási AL (2000) Error and attack tolerance of complex networks. *Nature* 406(6794):378–382.
- Callaway DS, Newman MEJ, Strogatz SH, Watts DJ (2000) Network robustness and fragility: Percolation on random graphs. *Phys Rev Lett* 85(25):5468–5471.
- Cohen R, Erez K, ben-Avraham D, Havlin S (2001) Breakdown of the internet under intentional attack. *Phys Rev Lett* 86(16):3682–3685.
- Karp RM (1972) Reducibility among combinatorial problems. *Complexity of Computer Computations*, eds Miller RE, Thatcher JW, Bohlinger JD (Springer, Boston), pp 85–103.
- Altarelli F, Braunstein A, Dall'Asta L, Zecchina R (2013) Large deviations of cascade processes on graphs. *Phys Rev E Stat Nonlin Soft Matter Phys* 87(6):062115.
- Altarelli F, Braunstein A, Dall'Asta L, Zecchina R (2013) Optimizing spread dynamics on graphs by message passing. *J Stat Mech* 09:P09011.
- Zhou HJ (2013) Spin glass approach to the feedback vertex set problem. *Eur Phys J B* 86(11):1–9.
- Gugliola A, Semerjian G (2015) Minimal contagious sets in random regular graphs. *J Stat Phys* 158(2):300–358.
- Mugisha S, Zhou HJ (2016) Identifying optimal targets of network attack by belief propagation. *Phys Rev E Stat Nonlin Soft Matter Phys* 94(1-1):012305.
- Mézard M, Parisi G (2001) The Bethe lattice spin glass revisited. *Eur Phys J B* 20:217–233.
- Mézard M, Montanari A (2009) *Information, Physics and Computation* (Oxford Univ Press, London).
- Leskovec J, Krevl A (2014) *SNAP Datasets: Stanford Large Network Dataset Collection*. Available at snap.stanford.edu/data. Accessed September 30, 2015.