Explosive Percolation in Random Networks

Dimitris Achlioptas,[1] Raissa M. D'Souza,[2*] Joel Spencer[3]

Networks in which the formation of connections is governed by a random process often undergo a percolation transition, wherein around some critical point, the addition of a small number of connections causes a sizable fraction of the network to suddenly become linked together. Typically such transitions are continuous, so that the percentage of the network linked together tends to zero right above the transition point. Whether percolation transitions could be discontinuous has been an open question. Here, we show that incorporating a limited amount of choice in the classic Erdös-Rényi network formation model causes its percolation transition to become discontinuous.

A large system is said to undergo a phase transition when one or more of its properties change abruptly after a slight change in a controlling variable. Besides water turning into ice or steam, other prototypical phase transitions are the spontaneous emergence of magnetization and superconductivity in metals, the epidemic spread of disease, and the dramatic change in connectivity of networks and lattices known as percolation. Perhaps the most fundamental characteristic of a phase transition is its order, i.e., whether the macroscopic quantity it affects changes continuously or discontinuously at the transition. Continuous (smooth) transitions are known as second-order and include many magnetization phenomena, whereas discontinuous (abrupt) transitions are known as first-order, a familiar example being the discontinuous drop in entropy when liquid water turns into solid ice at 0°C.

We consider percolation phase transitions in models of random network formation. In the classic Erdös-Rényi (ER) model (*1*), we start with n isolated vertices (points) and add edges (connections) one by one, each edge formed by picking two vertices uniformly at random and connecting them (Fig. 1A). At any given moment, the (connected) component of a vertex v is the set of vertices that can be reached from v by traversing edges. Components merge under ER as if attracted by gravitation. This is because every time an edge is added, the probability two given components will be merged is proportional to the number of possible edges between them which, in turn, is equal to the product of their respective sizes (number of vertices).

One of the most studied phenomena in probability theory is the percolation transition of ER random networks, also known as the emergence of a giant component. When rn edges have been added, if $r < ½$, the largest component remains miniscule, its number of vertices C scaling as $\log n$; in contrast, if $r > ½$, there is a component of size linear in n. Specifically, $C \approx (4r - 2)n$ for r slightly greater than ½ and, thus, the fraction of vertices in the largest component undergoes a continuous phase transition at $r = ½$ (Fig. 1C). Such continuity has been considered a basic characteristic of percolation transitions, occurring in models ranging from classic percolation in the two-dimensional grid to random graph models of social networks (*2*).

Here, we show that percolation transitions in random networks can be discontinuous. We demonstrate this result for models similar to ER, thus also establishing that altering a network-formation process slightly can affect it dramatically, changing the order of its percolation transition. Concretely, we consider models that, like ER, start with n isolated vertices and add edges one by one. The difference, as illustrated in Fig. 1B, is that to add a single edge we now pick two random edges $\{e_1,e_2\}$, rather than one, each edge picked exactly as in ER and independently of the other. Of these, with no knowledge of future edge-pairs, we are to select one and insert it in the graph and discard the other. Clearly, if we always resort to randomness for selecting among the two edges, we recover the ER model. Whether nonrandom selection rules can delay (or accelerate) percolation in such models has received much attention in recent years (*3*–*6*).

A selection rule is classified as "bounded-size" if its decision depends only on the sizes of the components containing the four end points of $\{e_1,e_2\}$ and, moreover, it treats all sizes greater than some (rule-specific) constant K identically. For example, a bounded-size rule with $K = 1$ due to Bohman and Frieze (BF) (*3*), the first selection rule to be analyzed, proceeds as follows: If e_1 connects two components of size 1, it is selected; otherwise, e_2 is selected. So, in Fig. 1B, e_2 would be selected. Bounded-size rules, in general, are amenable to rigorous mathematical analysis, and in (*3*, *4*) it was proven that such rules are capable both of delaying and of accelerating percolation. In contrast, unbounded-size rules seem beyond the reach of current mathematical techniques. A crucial point is that the percolation transition is strongly conjectured to be continuous for all bounded-size rules (*4*). This conjecture is supported both by numerical evidence and mathematical considerations, though a fully rigorous argument has remained elusive.

[1]University of California at Santa Cruz, Santa Cruz, CA xxxxx, USA. [2]University of California at Davis, Davis, CA 95616, USA. [3]New York University, New York, NY xxxxx, USA.

*To whom correspondence should be addressed. E-mail: raissa@cse.ucdavis.edu

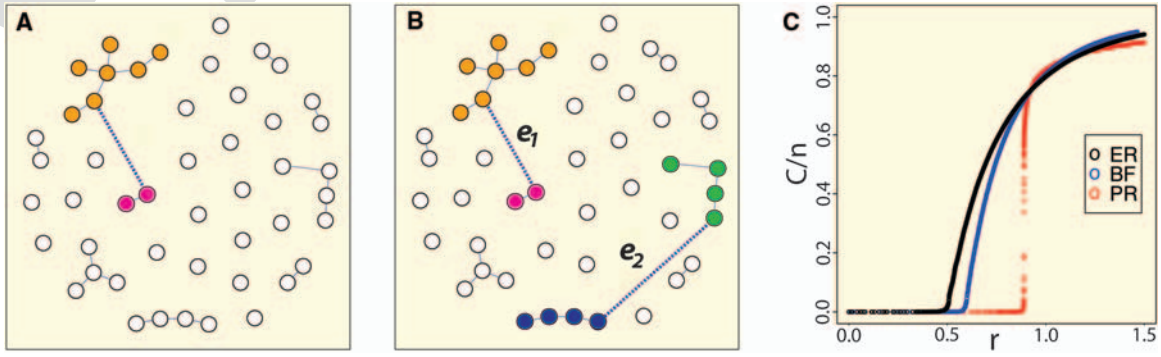

Fig. 1. Network evolution. (**A**) Under the Erdös-Rényi (ER) model, in each step two vertices are chosen at random and connected by an edge (shown as the dashed line). In this example, two components of size 7 and 2 get merged. (**B**) In models with choice, two random edges $\{e_1,e_2\}$ are picked in each step yet only one is added to the network based on some selection rule, whereas the other is discarded. Under the product rule (PR), the edge selected is the one minimizing the product of the sizes of the components it merges. In this example, e_1 (with product $2 \times 7 = 14$) would be chosen and e_2 discarded (because $4 \times 4 = 16$). In contrast, the rule selecting the edge minimizing the sum of the component sizes instead of the product would select e_2 rather than e_1. (**C**) Typical evolution of C/n for ER, BF (a bounded size rule with $K = 1$), and PR, shown for $n = 512{,}000$.

MS no: RE1167782/HJ/COMP/MATH

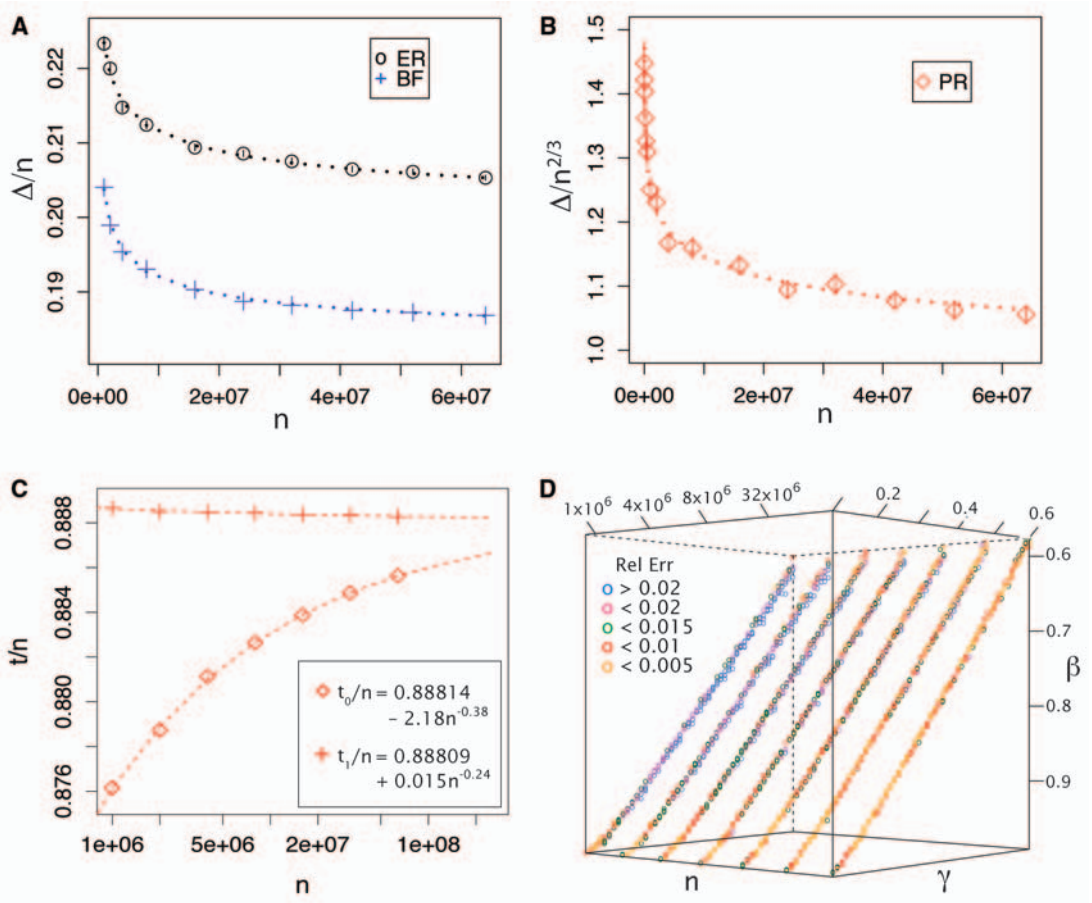

Fig. 2. (**A**) The ratio Δ/n for ER and BF for increasing system sizes. (**B**) The ratio $\Delta/n^{2/3}$ for PR for increasing system sizes. (**C**) Convergence to $r_c = 0.888\ldots$ from above and below (the two curves fitted independently). (**D**) A linear scaling relation is obeyed in the range $\gamma \in [0.2,0.6]$, shown here for $A = 0.5$. Color shows convergence with increasing system size n to the relation $\gamma + 1.2\beta = 1.3$. Our numerical experiments establish this scaling relation for $A \in [0.1,0.6]$ and we expect that in larger system sizes this range of A would broaden, particularly the lower end.

Here, we provide conclusive numerical evidence that, in contrast, unbounded-size rules can give rise to discontinuous percolation transitions. For concreteness, we present evidence for the so-called product-rule (PR): Always retain the edge that minimizes the product of the sizes of the components it joins, breaking ties arbitrarily (Fig. 1B). Thus, the PR selection criterion attempts to reduce the aforementioned gravitational attraction between components. We note that other unbounded-size rules also yield first-order transitions. For example, results similar to those for PR hold when "product" is replaced by "sum." It is also worth noting that the criterion employed by PR can also be used to accelerate percolation by always selecting the edge that maximizes rather than minimizes the product of the size of the components it merges (and similarly for sum). Nevertheless, in that case, the percolation transition remains continuous, reflecting the completely different evolution of the component-size distribution in the maximizing versus the minimizing case.

Let C denote the size of the largest component, t_0 denote the last step for which $C < n^{1/2}$, and t_1 the first step for which $C > 0.5n$. In continuous transitions, the interval $\Delta = t_1 - t_0$ is always extensive, i.e., linear in n. For example, $\Delta > 0.193n$ in ER. In contrast, as we show in Fig. 2B, Δ is not extensive for the product rule; indeed, $\Delta < 2n^{2/3}$ and it appears that $\Delta/n^{2/3} \to 1$, so that the fraction of vertices in the largest component jumps from being a vanishing fraction of all vertices to a majority of them "instantaneously." Although t_0/n and t_1/n converge to $r_c = 0.888\ldots$ (Fig. 2C), the variance in the value of t_0 and t_1 is enough to prevent the direct observation of a first-order transition. That is, measuring the size of the largest component as a function of the number of steps and averaging it over different realizations smears out the transition point, motivating our introduction of Δ and its measurement along different realizations. Specifically, each data point in Fig. 2, A to C, represents an average over an ensemble of 50 independent identically distributed realizations, and the dashed lines are the statistical best fits to the data (for details, see the supporting online material). Our computer implementation makes use of efficient procedures (*7*) for tracking how components merge as edges are added.

Our choice of $n^{1/2}$ and $0.5n$ above for defining Δ was simply illustrative. To demonstrate the discontinuity of PR's percolation transition, it suffices to find constants $A > 0$ and $\beta,\gamma < 1$ such that the number of steps between $C < n^{\gamma}$ and $C > An$ is smaller than n^{β}. Indeed, we have discovered a general scaling law associated with PR's percolation. For a range of values for A, we find that the same simple linear scaling relation governs the boundary of valid parameter choices, namely $\gamma + \lambda\beta = \mu$, where to the best of our numerical estimates, $\lambda \approx 1.2$ and $\mu \approx 1.3$. Convergence to this behavior for $A = 0.5$ is shown in Fig. 2D. Here, each data point depicts an individual realization, and color is used to show the relative error between the empirical value and that predicted by the scaling relation (see supporting online material for details).

We have demonstrated that small changes in edge formation have the ability to fundamentally alter the nature of percolation transitions. Our findings call for the comprehensive study of this phenomenon, and of its potential use in bringing phase transitions under control.

References and Notes

1. P. Erdős, A. Rényi, *Publ. Math. Inst. Hungar. Acad. Sci.* **5**, 17 (1960).
2. M. E. J. Newman, D. J. Watts, S. H. Strogatz, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 2566 (2002).
3. T. Bohman, A. Frieze, *Random Structures Algorithms* **19**, 75 (2001).
4. J. Spencer, N. Wormald, *COMBINATORICA* **27**, 587 (2007).
5. A. Beveridge, T. Bohman, A. Frieze, O. Pikhurko, *Proc. Am. Math. Soc.* **135**, 3061 (2007).
6. M. Krivelevich, E. Lubetzky, B. Sudakov, "Hamiltonicity thresholds in Achlioptas processes"; available at http://arxiv.org/abs/0804.4707 (2008).
7. M. E. J. Newman, R. M. Ziff, *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **64**, 016706 (2001).

8. We thank Microsoft Research, where our collaboration initiated, for its support. D.A. is supported in part by NSF CAREER award CCF-0546900, an Alfred P. Sloan Fellowship, and IDEAS grant 210743 from the European Research Council.

Supporting Online Material

www.sciencemag.org/cgi/content/full/VOL/ISSUE/PAGE/DC1
SOM Text

28 October 2008; accepted 16 January 2009
10.1126/science.1167782