Competition for Shortest Paths on Sparse Graphs

Chi Ho Yeung and David Saad

The Nonlinearity and Complexity Research Group, Aston University, Birmingham B4 7ET, United Kingdom

(Received 1 February 2012; published 14 May 2012)

Optimal paths connecting randomly selected network nodes and fixed routers are studied analytically in the presence of a nonlinear overlap cost that penalizes congestion. Routing becomes more difficult as the number of selected nodes increases and exhibits ergodicity breaking in the case of multiple routers. The ground state of such systems reveals nonmonotonic complex behaviors in average path length and algorithmic convergence, depending on the network topology, and densities of communicating nodes and routers. A distributed linearly scalable routing algorithm is also devised.

DOI: 10.1103/PhysRevLett.108.208701 PACS numbers: 89.75.Hc, 02.50.−r, 05.20.−y, 89.20.−a

Routing and path selection are at the heart of many communication and logistics applications. For instance, instant messengers, Internet telephony, and payment security verification require packets to be delivered instantly or otherwise lose their functionality [1,2], while the efficiency of transportation networks depends crucially on effective path selection [3,4]. Existing routing algorithms are mostly based on minimizing path lengths. Some use routing tables that register the shortest distance to various destinations but are insensitive to traffic congestion [5,6]; others control congestion by monitoring queue length or latency heuristically [7] or merely optimize routing selfishly [8]. Both analytical understanding of routing and the development of efficient distributive principled routing algorithms, which minimize route length while restricting congestion, remain a challenge.

Path optimality and congestion control have been extensively studied within the physics community in other contexts, such as the research of spanning [9,10] and Stenier trees [11] with quenched link weights, to mimic broadcast or multicast systems. However, these studies ignore interaction terms (overlap costs) that depend on the specific choice of paths. Other approaches such as preferential random walk and diffusion methods are used to reduce traffic congestion but result in heuristic protocols which route packets through suboptimal paths in a probabilistic manner [12–14]. Active suppression of overlaps is essential for realistic analysis of routing as they give rise to congestion; for instance, the road transit time is described by a quartic function of traffic volume [15].

Mapped onto a statistical physics framework, routing poses both theoretical and numerical challenges due to the multiplicity of possible routes between communicating nodes and the nonlinear costs induced by the interaction between overlapping routes, akin to a nonlocal repulsion force between polymers. While overlap costs have been partially addressed by assigning quenched link weights [9], these do not fully reflect the complex interaction between dynamically assigned routes. Techniques used to analyze polymers [16], for instance, in self-avoiding walks [17] and the traveling salesman problem [18], are prime candidates for the analysis of routing but do not consider the cost of interaction between paths.

In this Letter, we study a scenario whereby numerous senders seek the shortest possible route to a few receivers while minimizing traffic congestion. The problem is relevant to node pairs on a network that communicate via designated routers, or nodes, possibly sensors, that communicate via an outlet router or base stations. It is also relevant to transportation networks where traffic gravitates towards one or several centers, such as city center or hub airports [3,19]. Using the cavity approach [20], we examine analytically the dependence of the typical path length on the system topology, the density of communicating node pairs and routers, and their location. We identify the conditions for ergodicity breaking in solution space and observe oscillations in typical path lengths and algorithmic convergence in regular graphs. We show that allocating routers on hubs, which seems to be the natural choice, is indeed optimal in many respects. An applicable linearly scalable local message passing algorithm is also derived which optimizes, in a principled manner, individual routes subject to the mitigation of global congestion.

Model.—We consider a sparse network of N nodes, where each node $i = 1, \ldots, N$ is randomly connected to k_i neighbors denoted by $\mathcal{N}_i$, where $k_i \ll N$ is randomly drawn from a distribution $\rho(k)$ representing the graph topology. We randomly select one node as a receiver (router) denoted by r [21] and a fraction $0 \leq f_s \leq 1$ of the other nodes as senders that communicate with r through a single path each. We assign $\sigma_{ij}^{s\to r} = 1$ if the communication between sender s and r passes from i to j and $\sigma_{ij}^{s\to r} = 0$ otherwise. To minimize path length subject to a nonlinear cost for overlapping paths, we introduce the Hamiltonian $\mathcal{H} = \sum_{(ij)}[\sum_s(\sigma_{ij}^{s\to r} + \sigma_{ji}^{s\to r})]^\alpha$, counting the number of paths passing through the edge (ij), with $\alpha > 1$ to nonlinearly penalize overlaps; both analysis and algorithm are generic for any $\alpha \geq 1$.

To facilitate the analysis, each node i is assigned with a communication load Λ_i, reflecting its role: The receiver

node r being a global sink is assigned with $\Lambda_r = -\infty$, senders are assigned with $\Lambda = +1$ as they originate communication, and all other nodes are assigned with $\Lambda = 0$. We further denote the net directed integer flow $I_{ij} \equiv \sum_s [\sigma_{ij}^{s\to r} - \sigma_{ji}^{s\to r}]$ from i to j, where $I_{ji} = -I_{ij}$, and use it to ensure that all arriving or originated communications leave node i so that the excess communication load $R_i = \Lambda_i + \sum_{j\in\mathcal{N}_i} I_{ji}$ vanishes, i.e., $R_i = 0$, $\forall\, i \neq r$.

To minimize path length subject to a nonlinear cost for overlapping paths, one should minimize $\mathcal{H}$, which is computationally difficult. Instead, we minimize the Hamiltonian $\mathcal{H}' = \sum_{(ij)} |I_{ij}|^\alpha$ subject to $R_i = 0$, $\forall\, i \neq r$, as it can be shown [22] that both Hamiltonians share an identical ground state. An example of a small system attaining the ground state of $\mathcal{H}'$ with $\alpha = 2$ is shown in Fig. 1(a).

This model can be adapted to accommodate various routing scenarios. For instance, senders with integer $\Lambda > 1$ may be introduced to model data segmentation and routing through multiple paths; multiple receivers may be considered as is the case in wireless sensor networks where any reachable base station will do [7]. The Hamiltonian can also be modified to accommodate different forms of cost and objectives on nodes and links.

To analyze the system's ground state behavior, we employ the cavity approach [20] and assume that only large loops exist, so that neighbors of node i become statistically independent if it is being removed. At zero temperature, this allows one to define the energy at each vertex $E_i(I_{il})$, representing the sum of net flows on a tree terminated at the edge from vertex i to parent l, with all flows in the tree optimized given I_{il} [23]. For any $i \neq r$, one obtains a recursion relating the energy at vertex i, $E_i(I_{il})$, to $E_j(I_{ji})$ of its neighbors (descendents) j other than (parent) l, as sketched in Fig. 1(b) and given by

$$E_i(I_{il}) = \min_{\{\{I_{ji}\}|R_i=0\}} \Big[|I_{il}|^\alpha + \sum_{j\in\mathcal{N}_i\backslash l} E_j(I_{ji}) \Big]. \qquad (1)$$

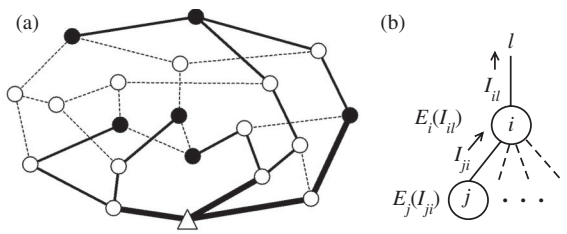

FIG. 1. (a) An example of the ground state of $\mathcal{H}'$ with $\alpha = 2$ in a simple network with 6 senders (●) and a single receiver (△) among 20 nodes. Dashed lines correspond to idle links, while thin and thick solid lines correspond to communication loads of 1 and 2, respectively. To minimize $\mathcal{H}'$, communication from the top left sender is routed through a long path. (b) A schematic diagram representing the derivation of Eq. (1).

Intuitively, for each node i one minimizes the nonlinear cost and descendent energies given the constraint of $R_i = 0$. For receiver nodes, being universal sinks, $E_r = |I_{rj}|^\alpha$, $\forall\, j \in \mathcal{N}_r$. For brevity we denote the right-hand side of Eq. (1) by a functional $\mathcal{M}[\underline{E}_{il}; \Lambda_i, I_{il}]$, where $\underline{E}_{il} \equiv \{E_j | j \in \mathcal{N}_i\backslash l\}$. The function $E(I)$ is extensive and difficult to iterate; it is replaced by the intensive quantity $E^V(I) \equiv E(I) - E(0)$ [24] to obtain the recursive relation

$$E_i^V(I_{il}) = \mathcal{M}[\underline{E}_{il}^V; \Lambda_i, I_{il}] - \mathcal{M}[\underline{E}_{il}^V; \Lambda_i, 0]. \qquad (2)$$

To evaluate various physical quantities, one employs population dynamics [22,25] to iterate Eq. (2) and obtain a stable distribution $P[E^V(I)]$. The ground state energy is evaluated by the average energy of an additional node and link; by denoting $\underline{E}_i^V \equiv \{E_j^V | j \in \mathcal{N}_i\}$, this is given by

$$\langle E_{\text{node}} \rangle = \langle \mathcal{M}[\underline{E}_i^V; \Lambda, 0] \rangle_{\Lambda, k, \underline{E}_i^V}, \qquad (3)$$

$$\langle E_{\text{link}} \rangle = \langle \min_I [E_{j_1}^V(I) + E_{j_2}^V(-I) - |I|^\alpha] \rangle_{E_{j_1}^V, E_{j_2}^V}. \qquad (4)$$

Angled brackets denote averages over sender and degree distributions $p(\Lambda)$, $\rho(k)$, and $P[E^V(I)]$. Equation (3) corresponds to the energy change due to contributions of all k trees per node, while Eq. (4) is obtained by considering two trees with a common edge [22]. The energy per transmission is given by $\langle E \rangle = (\langle E_{\text{node}} \rangle - \frac{\langle k \rangle}{2} \langle E_{\text{link}} \rangle)/f_s$; the numerator corresponds to the average energy per node [22].

To obtain the average path length, we calculate a linear-cost function L^V parallel to E^V:

$$L_i^V(I_{il}) = |I_{il}| + \sum_{j\in\mathcal{N}_i\backslash l} (L_j^V[I_{ji}^*(I_{il})] - L_j^V[I_{ji}^*(0)]), \qquad (5)$$

where the optimal flow $I_{ji}^*(I_{il})$, given I_{il}, is determined from Eq. (2). While E^V provides the optimal energy, L^V evaluates typical traffic in the optimal state. The average path length per communication is obtained by using the joint distribution $P(E^V, L^V)$ in the form $\langle L \rangle = (\langle L_{\text{node}} \rangle - \frac{\langle k \rangle}{2} \langle L_{\text{link}} \rangle)/f_s$, where averages $\langle L_{\text{node}} \rangle$ and $\langle L_{\text{link}} \rangle$ are evaluated similarly to Eq. (5) via the corresponding flow I^* obtained from Eqs. (3) and (4).

Solutions.—Once we obtain numerically a stable distribution $P[E^V(I)]$ by iterating Eq. (2), both $\langle E \rangle$ and $\langle L \rangle$ can be evaluated through Eqs. (3)–(5). Results shown in Fig. 2, for $\alpha = 2$, $f_r = 0.01$ [21], and different connectivity distributions representing regular, Erdös-Rényi (ER), and scale-free (SF) networks, show a linear increase in $\langle E \rangle$ for most of the range of f_s regardless of network topology. This indicates that the potential quadratic increase in cost due to overlaps is mitigated through rerouting, balancing path length, and overlap costs. The typical path length $\langle L \rangle$ (inset) increases initially with f_s for all networks (SF being shorter) due to overlap costs but saturates for ER networks while showing a decrease for regular and SF graphs at high f_s. This indicates that longer alternative routes are also

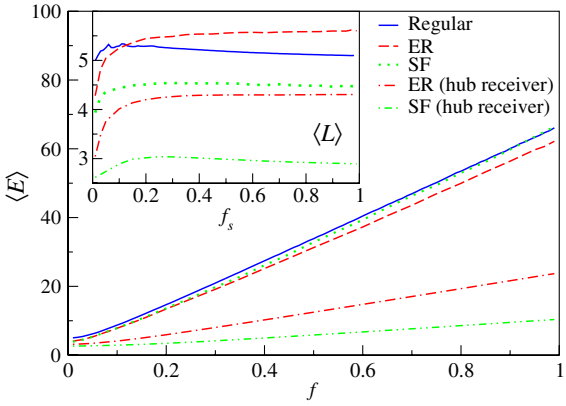

FIG. 2 (color online). Analytical results of $\langle E\rangle$ and $\langle L\rangle$ per communication as a function of f_s on regular, ER, and SF networks and ER and SF networks with hub receivers, all with $\langle k\rangle = 3$ and $f_r = 0.01$.

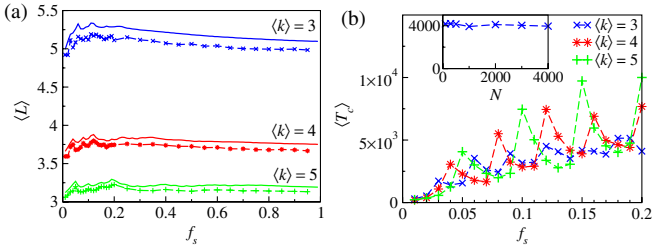

FIG. 3 (color online). (a) Analytic $\langle L\rangle$ values (solid lines) compared to simulation results (dashed lines, symbols) for regular networks with $\langle k\rangle = 3, 4, 5$, a single receiver, and $N = 100$ nodes, averaged over 1000 realizations. (b) Algorithmic convergence time $\langle T_c\rangle$ averaged over instances ($\approx 98\%$) which converge within 5×10^4 updates per node. Inset: $\langle T_c\rangle$ as a function of N for $\langle k\rangle = 3$, $f_s = 0.2$, and a single receiver.

congested, making shorter congested ones more cost-effective. It suggests that simple shortest path routing schemes may achieve close-to-optimal performance for either low or high traffic densities in regular and SF networks but are less effective for intermediate densities and ER networks.

To examine the effect of receiver connectivity, we study the case where receivers are deployed to the largest hubs in ER and SF networks. Figure 2 shows a more moderate increase in energy and path length compared to random deployment. In addition, communication costs in SF networks become lower with respect to ER networks, reversing their random case positions. This suggests that receiver local connectivity dominates network performance and potentially explains why most communication networks are SF with hub receivers.

Careful examination of average path length $\langle L\rangle$ in regular networks as a function of f_s reveals a fine structure of small peaks, in addition to the main peak, at multiples of $\langle k\rangle/N$ as shown in Fig. 3(a). These occur when traffic is balanced predominantly around the receiver, which we term the balanced receiver phenomenon. For example, the long route from the top left sender in Fig. 1(a) is chosen to ensure that links connecting the receiver are occupied exactly by 2 communications, as a high cost is incurred by imbalanced traffic. This behavior is also reflected in the expected algorithmic convergence time (number of update steps per node) $\langle T_c\rangle$ shown in Fig. 3(b). The balanced receiver phenomenon may appear in real-life networks; for instance, one may attempt to balance car traffic heading towards a city center by varying toll or speed limits. This phenomenon is masked in the results for ER and SF networks due to the degree variability of receivers and since the balanced receiver phenomenon is more pronounced at low degree receivers.

Computation.—One computational challenge in Eq. (2) is the extremization on the integer domain subject to an equality constraint. We use the convexity of $|I|^\alpha$ ($\alpha \geq 1$) to show that $E^V(I)$ is convex and denote the energy change when I_{il} increases or decreases by 1 as $\Delta_i^{\pm}(I_{il}) = E_i^V(I_{il} \pm 1) - E_i^V(I_{il})$. The optimality condition $\Delta_{j_1}^{+}[I^*_{j_1 i}(I_{il})] + \Delta_{j_2}^{-}[I^*_{j_2 i}(I_{il})] \geq 0$, $\forall\, j_1, j_2 \in \mathcal{N}_i \backslash l$, and the convexity of $E_i^V(I_{il})$ yield

$$\Delta_i^{\pm}(I_{il}) = |I_{il} \pm 1|^\alpha - |I_{il}|^\alpha + \min_{j \in \mathcal{N}_i \backslash l} \{\Delta_j^{\pm}[I^*_{ji}(I_{il})]\},$$

$$I^*_{ji}(I_{il} \pm 1) = \begin{cases} I^*_{ji}(I_{il}) \pm 1, & j = \text{argmin}_{j \in \mathcal{N}_i \backslash l}\{\Delta_j^{\pm}[I^*_{ji}(I_{il})]\}, \\ I^*_{ji}(I_{il}), & \text{otherwise.} \end{cases} \tag{6}$$

This simplifies the computation [22] and facilitates a study of the model's ergodicity breaking properties. Equation (6) implicitly assumes localized interdependence of flows that corresponds to the replica symmetry (RS) [20] property. Small variations $\delta\Delta_i^{\pm}(I_{il})$ in the cavity field of Eq. (6) lead to expected perturbations of the form $\langle(\delta\Delta_i^{\pm})^2\rangle = f_s\langle(\delta\Delta_j^{\pm})^2\rangle$ for all system nodes. Perturbations decay for $f_s < 1$ indicating an RS [26] behavior and a unique state. This is supported by simulations, as almost all cases converge in a small number of updates.

Algorithm.—Interestingly, the analysis also gives rise to an algorithm for optimizing individual instances; to do that, one iterates Eq. (2) until convergence of the E^V's and computes the flows by using either Eq. (3) or (4) to identify individual paths. The situation becomes more complicated when solutions of Eq. (3) or (4) exhibit degeneracy, which is typical for integer α. One can then break the degeneracy by assigning randomly to each link (ij) a small quenched bias $\epsilon_{(ij)}$ and incorporate in the iteration of Eq. (2) the modified cost of $|I_{ij}|^\alpha + |I_{ij}|\epsilon_{(ij)}$. The iterations then converge to a particular solution with no degeneracy. The algorithm is linearly scalable as shown in the inset in Fig. 3(b), but the bias increases convergence time with respect to the zero bias case.

Breaking of ergodicity.—As in other combinatorial optimization problems, a replica symmetry breaking (RSB) [20] may emerge under some conditions, resulting in numerous low-lying states that hinder algorithmic convergence. Clearly, this prevents one from obtaining a complete analytical solution of the problem; algorithmically, it may lead to frequent switching of variable states, routing instabilities, and a loss of information.

The framework studied so far is replica symmetric. However, we will now introduce two similar routing scenarios that exhibit RSB behavior under some conditions. One scenario is that of multiple receiver classes, where users (sensors) transmit information for specific routers (receivers). In the case of two receiver types, say, A and B with $(\Lambda^A, \Lambda^B) = (-\infty, 0)$ and $(0, -\infty)$, respectively, senders are initialized with either $(\Lambda^A, \Lambda^B) = (+1, 0)$ or $(0, +1)$. The corresponding Hamiltonian is $\mathcal{H}' = \sum_{(ij)}(|I_{ij}^A| + |I_{ij}^B|)^\alpha$, and vertex energies $E^V(I_{ij}^A, I_{ij}^B)$ become two-dimensional. As E^V's are no longer convex for the entire range, only an approximate solution can be obtained by a recursion relation of cavity fields similar to Eq. (6). To identify the RS-RSB transition, we introduce perturbations to the cavity fields, which may depend on those induced by the descendants [22].

Figure 4(a) shows the RS-RSB phase transition as a function of the fractions f_s and f_r of sender and router nodes with respect to the system size, respectively, for regular ER networks with random and hub routers (f_r refers equally to each of the receiver types). Generally, RSB emerges for large f_s and small f_r as a large number of communications from distant senders are optimized simultaneously, leading to an extensive frustration of a finite fraction of all communications. While ER networks with a random receiver have a larger RSB phase than regular networks, it shrinks significantly when receivers are deployed on the hubs. This shows again the advantage of hub receivers, since the RS phase, where simple algorithms suffice, extends to systems with denser communications. Since the allocation of specified receivers to senders induces RSB, we expect a more extensive RSB phase with the increase in receiver types.

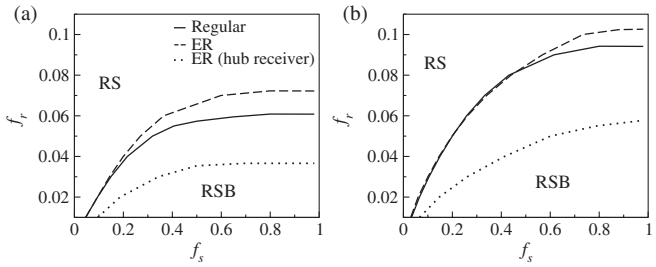

FIG. 4. Phase diagram for fractions f_s and f_r of sender and receiver (router) nodes, respectively, and two types of receivers on regular and ER networks with random and hub receivers. (a) Senders with a fixed router allocation and (b) sender pairs with a flexible router choice.

In another scenario, senders pair up to establish a communication line via a receiver, which can be of either type A or B, by setting their respective communication loads (Λ^A, Λ^B) to either $(1, 0)$ or $(0, 1)$. The RS-RSB transition can be similarly identified, but one has to account also for the perturbations when a partner sender switches its receiver type, resulting in the qualitatively similar but more extensive RSB phase shown in Fig. 4(b).

These results provide insights into the complexity of various routing problems and, in particular, to peer-to-peer communications. The case of a single receiver type can be interpreted as routing communications via a central router, while cases with two receiver types can be considered as user pairs communicating through a preset or flexible routers, where RSB emerges and becomes more prominent with the increasing number of receiver types or routing flexibility. This suggests that, instead of building a peer-to-peer network by setting each node as a receiver of a different type (with direct routing to individuals), one can use a limited number of receiver types by setting up a small number of central routers and benefit from simpler algorithms and smoother convergence.

Optimal routing in the presence of overlap costs is a difficult problem mainly due to the multiplicity of possible routes and the interaction between them; it is highly relevant to communication networks and exhibits a complex and interesting behavior as a function of the system topology and sender and receiver densities and types. Analytically obtained results provide insights into the suitability of various network topologies as communication networks and the preferred location of network servers. Routing scenarios whereby senders communicate individually or jointly through specific receivers exhibit RSB behavior with direct impact on the algorithmic hardness of the routing problem and implications for the study of constrained polymers on graphs. The application of statistical mechanics to communication networks is highly promising due to both the insight gained and the potential for better principled routing algorithms.

We thank K. Y. M. Wong, R. Urbanke, C. Fraguoli, N. Ruozzi, and A. Mozeika for fruitful discussions. This work is supported by EU FET project STAMINA (FP7-265496) and Royal Society Exchange Grant No. IE110151.

[1] C. Huitema, *Routing in the Internet* (Prentice-Hall, Englewood Cliffs, NJ, 1995).
[2] J. T. Moy, *OSPF: Anatomy of an Internet Routing Protocol* (Addison-Wesley, Reading, MA, 1998).
[3] Z. Wu, L. A. Braunstein, S. Havlin, and H. E. Stanley, Phys. Rev. Lett. **96**, 148702 (2006).
[4] B. J. Kim, C. N. Yoon, S. K. Han, and H. Jeong, Phys. Rev. E **65**, 027103 (2002).
[5] R. Bellman, Q. Appl. Math. **16**, 87 (1958).

[6] E. W. Dijkstra, Numer. Math. **1**, 269 (1959).
[7] S. Rangwala *et al.*, in *Proceedings of SIGCOMM 2006* (Association for Computing Machinery, New York, 2006), p. 63.
[8] T. Roughgarden and É. Tardos, J. ACM **49**, 236 (2002).
[9] J. D. Noh and H. Rieger, Phys. Rev. E **66**, 066127 (2002).
[10] R. Dobrin and P. M. Duxbury, Phys. Rev. Lett. **86**, 5076 (2001).
[11] M. Bayati, C. Borgs, A. Braunstein, J. Chayes, A. Ramezanpour, and R. Zecchina, Phys. Rev. Lett. **101**, 037208 (2008).
[12] L. A. Adamic, R. M. Lukose, A. R. Puniyani, and B. A. Huberman, Phys. Rev. E **64**, 046135 (2001).
[13] W. X. Wang, B.-H. Wang, C.-Y. Yin, Y.-B. Xie, and T. Zhou, Phys. Rev. E **73**, 026111 (2006).
[14] B. Danila, Y. Yu, J. A. Marsh, and K. E. Bassler, Phys. Rev. E **74**, 046106 (2006).
[15] E. Köhler and M. Skutella, SIAM J. Optim. **15**, 1185 (2005).
[16] M. Daoud, J. P. Cotton, B. Farnoux, G. Jannink, G. Sarma, H. Benoit, C. Duplessix, C. Picot, and P. G. de Gennes, Macromolecules **8**, 804 (1975).
[17] J. F. Stilck, K. D. Machado, and P. Serra, Phys. Rev. Lett. **76**, 2734 (1996).
[18] M. Mézard and G. Parisi, J. Phys. (Paris) **47**, 1285 (1986).
[19] D. R. Wuellner, S. Roy, and R. M. D'Souza, Phys. Rev. E **82**, 056101 (2010).
[20] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
[21] Alternatively, in the large system limit, one may select a small fraction of the system nodes f_r to act as receivers rather than selecting a single receiver.
[22] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.108.208701 for proofs for the equivalence of the ground states of $\mathcal{H}$ and $\mathcal{H}'$ and the convexity of $E^V(I)$, as well as complementary equations and discussions related to the material presented in this Letter.
[23] K. Y. M. Wong and D. Saad, Phys. Rev. E **74**, 010104(R) (2006).
[24] C. H. Yeung and K. Y. M. Wong, J. Stat. Mech. (2010) P04017.
[25] M. Mézard and R. Zecchina, Phys. Rev. E **66**, 056126 (2002).
[26] O. Rivoire, G. Biroli, O. C. Martin, and M. Mézard, Eur. Phys. J. B **37**, 55 (2003).