

From the physics of interacting polymers to optimizing routes on the London Underground

Chi Ho Yeung^{a,1}, David Saad^a, and K. Y. Michael Wong^b

^aNonlinearity and Complexity Research Group, School of Engineering and Applied Science, Aston University, Birmingham B4 7ET, United Kingdom; and ^bDepartment of Physics, The Hong Kong University of Science and Technology, Hong Kong

Edited by Susan N. Coppersmith, University of Wisconsin–Madison, Madison, WI, and approved July 2, 2013 (received for review January 22, 2013)

Optimizing paths on networks is crucial for many applications, ranging from subway traffic to Internet communication. Because global path optimization that takes account of all path choices simultaneously is computationally hard, most existing routing algorithms optimize paths individually, thus providing suboptimal solutions. We use the physics of interacting polymers and disordered systems to analyze macroscopic properties of generic path optimization problems and derive a simple, principled, generic, and distributed routing algorithm capable of considering all individual path choices simultaneously. We demonstrate the efficacy of the algorithm by applying it to: (i) random graphs resembling Internet overlay networks, (ii) travel on the London Underground network based on Oyster card data, and (iii) the global airport network. Analytically derived macroscopic properties give rise to insightful new routing phenomena, including phase transitions and scaling laws, that facilitate better understanding of the appropriate operational regimes and their limitations, which are difficult to obtain otherwise.

Path optimization affects many of our daily activities. Although much attention has been dedicated to routing algorithms for Internet applications, such as instant messengers and peer-to-peer systems (1, 2), many other essential routing applications have attracted less attention, ranging from water distribution networks (3) to sensor networks (4), military convoy movements (5), and journey planners (6, 7). In many applications, enormous costs are incurred due to traffic congestion or non-essential and redundant capacity. Due to the computational costs involved, most existing routing algorithms are static and based on selfish decisions, with nonadaptive routing tables indicating the shortest path to destinations regardless of local traffic (8, 9). Dynamic routing protocols do exist, but they are either heuristic, probabilistic, or insensitive to other individual path decisions that dynamically constitute the traffic (10, 11). A more global approach that takes into account all individual path decisions is crucial for efficient use of overstretched infrastructure. For instance, one may suppress congestion by minimizing overlaps with other routes or decrease the number of active nodes by consolidating paths to reduce infrastructure demands or energy consumption. The latter is particularly important in the context of the Internet because it can consume up to 4% of the electricity generated (12). Future applications include individualized routing and optimal resource management of prebooked air and road traffic.

The difficulty in deriving a globally optimal algorithm, in contrast to greedy local ones, lies in the simultaneous assignment of multiple interacting paths to minimize a global cost, because the optimal path between any particular source–destination pair depends on all other path choices. Such interaction is highly nonlocal, because paths between different source–destination pairs may partially overlap. Existing algorithms either ignore these interactions (8, 9) or use heuristics to approximate them (10, 11); both approaches result in suboptimal solutions. A substantial effort has been devoted to the development of highly efficient routing methods: for instance, multicommodity flow algorithms (13–19). However, most methods are based on weighted linear objective functions and real variables, and they aim specifically at satisfying capacity constraints; they have limited flexibility in addressing the variety of nonlinear cost functions

one may want to optimize in different scenarios, especially concave costs and integer variables. A more detailed discussion is provided in *SI Appendix, section S4*.

Here, we utilize statistical physics-based methods used in the study of interacting polymers (20) and spin glasses (21, 22) to obtain both a macroscopic description of the routing problem and microscopic solutions for given instances; the latter leads to a simple, generic, and distributed routing optimization algorithm. The algorithm resembles message-passing techniques that have been developed independently in a number of disciplines (21, 23, 24) and have been successfully applied to a variety of problems, ranging from prototyping (25) to solving hard computational problems (26) and control of complex systems (27). Here, we demonstrate the potential and efficacy of our routing algorithm by applying it to random networks, individualized routing on the London subway network, and the global airport network. Together with other benchmark tests described in *SI Appendix*, we demonstrate that our algorithm achieves better optimization compared with existing heuristics and state-of-the-art approximation algorithms in various routing scenarios; moreover, it is distributed and principled and does not require fine-tuning of free parameters.

In addition to the significant algorithmic advances, several macroscopic phenomena, including a phase transition, scaling rules as a function of network size, and nonmonotonic growth in mean path length as a function of traffic volume, are revealed; these cannot be obtained by numerical studies and provide unique insights and understanding of optimal routing on sparse networks.

Model

Consider a system of M polymers interacting on a network of N nodes. Each node $i = 1, \dots, N$ is connected to k_i neighbors denoted by the set \mathcal{L}_i and the connectivity matrix $A_{ij} = A_{ji} = 1$ when i and j are connected and is zero otherwise. Each polymer $\nu = 1, \dots, M$ has two fixed ends and occupies a path described by a self-avoiding walk on the network (i.e., consecutive segments occupy topological neighbors and each polymer ν goes through a node at most once). We denote the variable $\sigma_i^\nu = 1$ when polymer ν occupies node i and $\sigma_i^\nu = 0$ otherwise, and the number of polymers occupying i as $I_i = \sum_\nu \sigma_i^\nu$. To penalize or encourage polymer overlap, we define the Hamiltonian \mathcal{H} to be a nonlinear function of the normalized flow $\lambda_i = I_i/M$, namely,

$$\mathcal{H} = M \sum_i \phi(\lambda_i). \quad [1]$$

The analytical solution and derived algorithm are generic for any ϕ . Although the current framework focuses on undirected polymers and costs that incur at vertices, it is clear that costs may incur at the edges and edges may be directed and weighted in

Author contributions: C.H.Y. and D.S. designed research; C.H.Y. and D.S. performed research; C.H.Y., D.S., and K.Y.M.W. contributed new reagents/analytic tools; C.H.Y. analyzed data; C.H.Y., D.S., and K.Y.M.W. interpreted the results; and C.H.Y. and D.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: chyeung@gmail.com.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1301111110/-DCSupplemental.

some applications. Our framework, derivation, and algorithms accommodate costs on edges (using a factor graph representation), as well as directed and weighted polymers, making them suitable for most routing scenarios. The derivation and corresponding algorithms are given in *SI Appendix, section S3*. We would like to point out that the algorithm presented below already accommodates directed traffic.

This model is equivalent to a setting of M source–destination pairs, which we term “communications,” each of which occupies a path on a network with N nodes. The variable λ_i is thus the normalized traffic on node i , and \mathcal{H} is the corresponding cost function. In the physical framework and the zero-temperature limit, we minimize \mathcal{H} to obtain the ground state of the system or the optimal path configuration of the corresponding routing system. Some simple forms of \mathcal{H} are already meaningful; for instance, $\phi(x) = x^\gamma$, where the cases with $\gamma > 1$ penalize overlaps to suppress congestion, whereas $\gamma < 1$ encourages overlaps to aggregate traffic (28–30). The case of $\gamma = 1$ reduces to $\mathcal{H} \propto \sum_\nu (\sum_i \sigma_i^\nu)$, whose ground state corresponds to shortest path routing.

Methods

Theoretical Approach. The main obstacle in accounting for the interaction between paths is in keeping track of the cost at local nodes or edges while maintaining path integrity between the two end points and avoiding redundant loops. Therefore, in addition to the cost at the various nodes, given by Eq. 1, we have introduced a technique used in polymer physics (20) in the study of self-avoiding walks (31, 32) to enforce the appropriate path constraints.

The method is based on representing each node as an n -component vector \vec{S} of length \sqrt{n} . Denoting the angular integration over \vec{S} as $\int_{\mathcal{S}} d\vec{S}$, it has been shown (20) that all positive moments of S_a vanish in the limit $n \rightarrow 0$, except the second moment $\frac{1}{C_n} \int_{\mathcal{S}} d\vec{S} S_a^2 = 1$, for any component a in \vec{S} , where $C_n = \int_{\mathcal{S}} d\vec{S}$ is a normalization constant. It is then implied that when $n \rightarrow 0$, all non-vanishing terms that contribute in

$$\prod_{i=1}^N \left(\frac{1}{C_n} \int_{\mathcal{S}_i} d\vec{S}_i \right) S_{x,a} S_{y,a} \prod_{(kl)} \left(1 + A_{kl} \vec{S}_k \cdot \vec{S}_l \right) \quad [2]$$

are of the form $A_{xk_1} A_{k_1 k_2} \dots A_{k_{l-1} k_l} S_{x,a}^2 S_{k_1,a}^2 S_{k_2,a}^2 \dots S_{k_{l-1},a}^2 S_{y,a}^2$, where k_i represents the i th node index of the corresponding path/polymer segment; these sequences represent self-avoiding paths over nodes $(x, k_1, k_2, \dots, k_l, y)$, joining the end nodes x and y (20). Each node that is part of these paths incurs a cost as in Eq. 1; a sum over all possible paths of all communications provides the partition function \mathcal{Z} , as detailed in *SI Appendix, section S1.1*. To obtain typical macroscopic properties, one needs to average \mathcal{Z} over topologies (given a degree distribution) and node pair choices, termed “quenched disorders” in statistical physics. This requires the use of the replica or cavity method of spin glass theory (21, 22), as presented in *SI Appendix, section S1*.

The aim of the analysis is twofold.

- i) At the macroscopic level, we derive the stable traffic distribution $P(l)$ in the limit of very large systems to obtain the average cost (energy) $\langle E \rangle = \langle \phi(l/M) \rangle$; the average path length, given by the total occupancy divided by M (i.e., $\langle L \rangle = \frac{1}{M} \langle l \rangle$); and the average fraction of idle nodes, given by $f_{\text{idle}} = \langle \delta(l) \rangle$, as




Fig. 1. Optimized path configurations on a regular random network. The network comprises 50 nodes (each with $k=3$) and 10 source–destination pairs. The corresponding costs are $\mathcal{H} \propto \sum_i l_i^2$ (A) and $\mathcal{H} \propto \sum_i l_i^{2.5}$ (B). The path of each communication is illustrated by nodes and edges of a specific color, whereas black nodes are shared by more than one path. The size of a node is proportional to the amount of traffic through it, and square nodes represent the source or destination of each communication. Networks are visualized by the NetDraw software (40).

detailed in *SI Appendix, section S1.4*. Angled brackets denote an average over $P(l)$, which includes averages over all variable states for a given network and over choices of network and end-point instances.

- ii) At the microscopic level, the cavity-based analysis (33) translates to an algorithm that optimizes path configuration in a principled, distributed, and computationally efficient manner.

Optimization Algorithm. The analytical solutions for infinite systems translate into an optimization algorithm valid for finite systems, as detailed in *SI Appendix, section S2*. The derived algorithm is based on sending a couple of messages $a_{j \rightarrow i}^\nu$ and $b_{j \rightarrow i}^\nu$ at the zero temperature limit, from node j to node i for each index ν ; these characterize the energy contributions of communication ν at edge $j \rightarrow i$, originated from the source and destination directions, respectively. The messages take the form:

$$a_{j \rightarrow i}^\nu = \begin{cases} \min_{l \in \mathcal{L}_j \setminus \{i\}} [a_{l \rightarrow j}^\nu] - \min \left[-\phi^*(\lambda_j^{*\nu}), \min_{l, r \in \mathcal{L}_j \setminus \{i\}} [a_{l \rightarrow j}^\nu + b_{r \rightarrow j}^\nu] \right], & \Lambda_j^\nu = 0 \\ - \min_{l \in \mathcal{L}_j \setminus \{i\}} [b_{l \rightarrow j}^\nu], & \Lambda_j^\nu = 1 \\ \infty, & \Lambda_j^\nu = -1 \end{cases} \quad [3]$$

$$b_{j \rightarrow i}^\nu = \begin{cases} \min_{l \in \mathcal{L}_j \setminus \{i\}} [b_{l \rightarrow j}^\nu] - \min \left[-\phi^*(\lambda_j^{*\nu}), \min_{l, r \in \mathcal{L}_j \setminus \{i\}} [a_{l \rightarrow j}^\nu + b_{r \rightarrow j}^\nu] \right], & \Lambda_j^\nu = 0 \\ \infty, & \Lambda_j^\nu = 1 \\ 0, & \Lambda_j^\nu = -1 \end{cases} \quad [4]$$

where $\Lambda_j^\nu = +1, -1$ for source and destination, respectively, and is zero otherwise; the general cost function ϕ and the set of nodes in the neighborhood of node j are denoted as \mathcal{L}_j . The value of $\lambda_j^{*\nu}$ is given by the solution of λ_j^ν in

$$\lambda_j^\nu = \frac{1}{M} + \frac{1}{M} \sum_{\mu \neq \nu} \left\{ |\Lambda_\mu| + (1 - |\Lambda_\mu|) \Theta \left(-\phi^*(\lambda_j^\nu) - \min_{l, r \in \mathcal{L}_j} [a_{l \rightarrow j}^\mu + b_{r \rightarrow j}^\mu] \right) \right\}. \quad [5]$$

The step function $\Theta(x)$ takes values $\Theta(x) = 0, 0.5, 1$ for $x < 0, x = 0$, and $x > 0$, respectively. Solutions of Eq. 5 are obtained by setting $\lambda_j^\nu = l/M$ and a test integer l starting from $l=0$ until a self-consistent λ_j^ν is found. Finally, after the set of messages in Eqs. 3 and 4 converges to nonfluctuating values, the optimal configuration of path ν on each node j is given by

$$\sigma_j^\nu = |\Lambda_\nu| + (1 - |\Lambda_\nu|) \Theta \left(-\phi^*(\lambda_j^{*\nu}) - \min_{l, r \in \mathcal{L}_j} [a_{l \rightarrow j}^\nu + b_{r \rightarrow j}^\nu] \right), \quad [6]$$

where $\lambda_j^{*\nu}$ is the solution of Eq. 5 after convergence and $\sigma_j^\nu = 1$ if the communication ν passes through node j and zero otherwise. The generalized algorithms that accommodate weighted and directed communications, generic costs on nodes and edges, and separate costs defined on directed edges are given in *SI Appendix, section S3*. The computational complexities of these algorithms are discussed in *SI Appendix, section S2.2*.

In some instances, the iterative equations fail to converge; this suggests that solution space in the infinite system case is fragmented and nonergodic; this corresponds to “replica symmetry breaking” (21, 22), a complicated energy landscape with numerous local minima that typically hinder algorithmic convergence (details are provided in *SI Appendix, section S5*). This is typical in the case of hard computational problems. Convergence is improved by assigning a random bias ϵ_i to each node (34), akin to an external field, guiding the system to one of the local minima. These biases can be easily incorporated in the present formalism by replacing $\phi(x)$ with $\phi_i(x)$ for each node i , such that $\phi_i(x) = \phi(x) + x\epsilon_i$. In cases where a large number of source–destination pairs are identical, we further replace ϵ_i by ϵ_i^ν for each communication ν to break the degeneracy brought about by Eq. 5. Details can be found in *SI Appendix, section S2.1*.




Fig. 3. Optimized traffic at individual airports and London subway stations (*inset*). The airports and stations are plotted in descending order of traffic in the optimized state of $\mathcal{H} \propto \sum_i I_i^{0.5}$ (red lines). Symbols (\times) in blue correspond to the optimized traffic with $\mathcal{H} \propto \sum_i I_i^2$. Squared symbols refer to airports and stations mentioned in the main text that have much higher traffic than that indicated by the red lines. The optimized airport traffic is obtained from the single instance shown in Fig. 5, and the optimized subway traffic is obtained by averaging over the 30 passenger sets as in Table 1.

algorithm optimizes a given generic cost at a price of modest increase in the average path length.

To evaluate the performance of the suggested algorithm (with $\gamma = 2$ only), we compared our results against those obtained using

a representative state-of-the-art congestion-aware routing algorithm, which we call the “min-cap” (MC) algorithm (13), based on multicommodity flow. Because the latter aims to optimize a linear cost, we have introduced a tunable parameter α such that the quadratic cost is optimized by an extensive search for an optimal α^* (*SI Appendix, Fig. S8*). Details are provided in *SI Appendix, section S4*. We emphasize that this comparison is limited to congestion-aware algorithms ($\gamma \geq 1$) because we have not identified existing efficient route optimization algorithms for concave costs that facilitate route consolidation (e.g., the results shown in Figs. 1*B*, 2*B*, and 4*B*).

Table 2 shows a modest gain in cost over the optimized MC results at individual α^* for each run, which is far less than the gain obtained with respect to the Dijkstra algorithm. Nevertheless, our algorithm provides a lower energy for all α values, which is unachievable by the MC algorithm even after fine-tuning (*SI Appendix, Fig. S8*). Our algorithm also results in a shorter average path length L in addition to a lower cost E in random regular graphs (*SI Appendix, Table S1*) used as a controlled benchmark problem. Moreover, it is distributed and principled; does not require fine-tuning of free parameters; and, most importantly, has the flexibility to accommodate any (nonpathological) cost function designed to address specific needs.

Path Adaptivity. Fig. 5 illustrates the adaptivity of our algorithm after removing the London subway station “Bank” (black node). Nodes and edges in Fig. 5 that show an increase (decrease) in optimized traffic are colored red (blue), respectively, with their size and thickness proportional to the magnitude of increase (decrease). Nodes and edges in Fig. 5 with no traffic changes are shown in white and black, respectively. In the case of optimization using $\gamma = 2$, the original traffic through Bank is rerouted either via “Embankment” or via “Old Street.” This redistribution

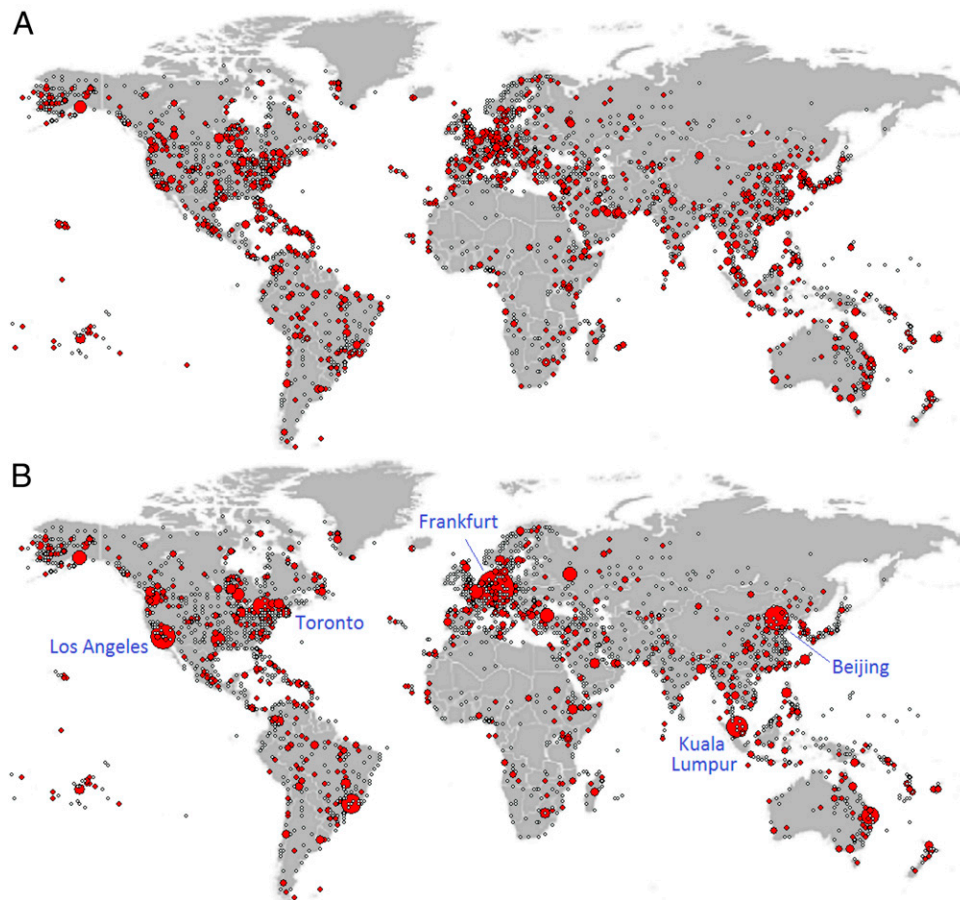


Fig. 4. Optimized traffic at individual airports of the global air network. A total of 2,480 airports constitute nodes, whereas the edges represent the existence of direct flights between airport pairs (36). Because the real demand in terms of source–destination pairs is unavailable, it was artificially generated by selecting a set of 300 randomly drawn source–destination pairs. Red nodes correspond to airports with nonzero traffic; the size of nodes indicates the air traffic through particular airports, and edges are omitted for clarity. (A) For $\mathcal{H} \propto \sum_i I_i^2$, traffic is routed to be almost uniformly distributed to reduce congestion. (B) For $\mathcal{H} \propto \sum_i I_i^{0.5}$, air traffic consolidates at the main hubs.

M value compared with $\gamma = 0.5$, implying that traffic homogeneity is achieved at smaller M in the case of $\gamma = 2$.

Although similar behaviors are observed for ER graphs (SI Appendix, section S6), SF networks show a much slower decrease of $\langle L \rangle$ after attaining its maximum, possibly due to the intrinsic node degree inhomogeneity that leads to traffic inhomogeneity even at large M . This suggests that shortest path routing is effective when M is large and topology is homogeneous but not in networks with a high degree of variability.

The scaling property of path lengths is shown in Fig. 6A. Rescaled path lengths $(\langle L \rangle - L_1)/(N/\log N)$ with $\gamma = 2$ at system sizes $n = 100, 200, 500$, and $1,000$, plotted as a function of the rescaled number of communications, $M/(N/\log N)$, fall on top of each other almost identically. A similar data collapse is also observed in ER graphs shown in SI Appendix, section S4. This finding implies that the nonmonotonic behavior observed for path lengths, and thus the network sensitivity to congestion, depends on M and N only through $M/(N/\log N)$. The latter is roughly proportional to the average traffic on a node because $\log N$ is proportional to the average shortest distance between any two nodes in random regular networks (37, 38) and ER graphs (39). In other words, the optimal behavior of routing on these graphs depends only on the average node traffic, regardless of system size and number of communications. The rescaling also appears in the reduced cost obtained by our algorithm, as shown in Fig. 6B. Note that theoretical results have been obtained in the infinite system limit; finite N values for theoretical results have been introduced merely to determine the scaling properties of M .

We have also examined the fraction of idle nodes as a function of γ . This revealed a phase transition, an abrupt change in the

fraction of idle nodes around the $\gamma = 1$ value (SI Appendix, section S7 and Fig. S11). The implication is that even a small change in the power γ is sufficient to power down unnecessary routers or close redundant subway stations effectively, with little impact on the cost or average route length.

Discussion

Optimal routing is one of many hard problems on networks that one should tackle to use limited and usually overstretched resources efficiently. The common characteristic of these problems is their global nature, and thus the difficulty in solving them at both macroscopic and microscopic levels with limited computational resources. By applying methods from the physics of interacting polymers and disordered systems, we obtained typical properties of routing problems and derive a readily applicable, principled, generic, distributed, and adaptive routing algorithm. Improvements over state-of-the-art algorithms in the intermediate traffic regime where $M \sim N \log N$ are considerable but are modest in the very sparse and dense traffic regimes. These findings will have a direct impact on a number of different research areas of practical and societal relevance, ranging from traffic to communication and logistics; however, more importantly, they may open the way for solving many other crucial and nonlocalized problems on networks.

ACKNOWLEDGMENTS. This work is supported by the European Union Future and Emerging Technologies project Statistical Mechanics Inspired Methods for Green Autonomous Networking (FP7-265496), Royal Society Exchange Grant IE110151, and the Research Grants Council of Hong Kong (Grant 605010).

- Huitema C (1995) *Routing in the Internet* (Prentice Hall, Englewood Cliffs, NJ).
- Moy JT (1998) *OSPF: Anatomy of an Internet Routing Protocol* (Addison-Wesley, Reading, MA).
- Vasan A, Simonovic SP (2010) Optimization of water distribution network design using differential evolution. *Journal of Water Resources Planning and Management* 136:279–287.
- Rangwala S, Gummadi R, Govindan R, Psounis K (2006) Interference-aware fair rate control in wireless sensor networks. *Computer Communication Review* 36:63–74.
- Chardaire P, McKeown GP, Verity-Harrison SA, Richardson SB (2005) Solving a time-space network formulation for the convoy movement problem. *Operations Research* 53(2):219–230.
- Beckmann M, McGuire CB, Winsten CB (1956) *Studies in the Economics of Transportation* (Yale Univ Press, New Haven).
- Wardrop JG (1952) Some theoretical aspects of road traffic research. *Proceedings of the Institute of Civil Engineers* 1(Part II):325–378.
- Bellman R (1958) On a routing problem. *Quarterly of Applied Mathematics* 16(1):87–90.
- Dijkstra EW (1959) A note on two problems in connexion with graphs. *Numerische Mathematik* 1:269–271.
- Xie H, Qiu L, Yang YR, Zhang Y (2004) On self adaptive routing in dynamic environments. *Proceedings of the 12th IEEE International Conference on Network Protocols* (IEEE, Washington, DC), pp 12–23.
- Zhu Y, Dovrolis C, Ammar M (2006) Dynamic overlay routing based on available bandwidth estimation: A simulation study. *Computer Networks Journal* 50:739–876.
- Baliga J, Hinton K, Tucker RS (2007) Energy consumption of the Internet. *Proceedings of the Conference on Optical Internet and the 32nd Australian Conference on Optical Fibre Technology* (IEEE, Washington, DC), pp 1–3.
- Shahrokhi F, Matula DW (1990) The maximum concurrent flow problem. *Journal of the Association for Computing Machinery* 37:318–334.
- Leighton T, Rao S (1988) An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. *Proceedings of Foundations of Computer Science 29th Annual Symposium* (IEEE, Washington, DC), pp 422–431.
- Awerbuch B, Azar Y, Plotkin S (1993) Throughput-competitive on-line routing. *Proceedings of Foundations of Computer Science 34th Annual Symposium* (IEEE, Washington, DC), pp 32–40.
- Garg N, Könemann J (1998) Faster and simpler algorithms for multicommodity flow and other fractional packing problems. *Proceedings of Foundations of Computer Science 39th Annual Symposium* (IEEE, Washington, DC), pp 300–309.
- Awerbuch B, Khandekar R (2007) Greedy distributed optimization of multi-commodity flows. *Proceedings of the Twenty-Sixth Annual ACM Symposium on Principles of Distributed Computing* (ACM, New York), pp 274–283.
- Barnhart C, Hane CA, Vance PH (2000) Using branch-and-price-and-cut to solve origin destination integer multicommodity flow problems. *Oper Res* 48:318–326.
- Castro J, Nabona N (1996) An implementation of linear and nonlinear multicommodity network flows. *Eur J Oper Res* 92:37–53.
- Daoud M, et al. (1975) Solutions of flexible polymers. Neutron experiments and interpretation. *Macromolecules* 8(6):804–818.
- Mézard M, Parisi G, Virasoro MA (1987) *Spin Glass Theory and Beyond* (World Scientific, Singapore).
- Nishimori H (2001) *Statistical Physics of Spin Glasses and Information Processing* (Oxford Univ Press, Oxford, UK).
- Pearl J (1982) Reverend Bayes on inferenced engines: A distributed hierarchical approach. *Proceedings of the Second National Conference on Artificial Intelligence* (AAAI, San Francisco), pp 133–136.
- Gallager RG (1998) *Information Theory and Reliable Communication* (Wiley, NY).
- Barash Y, et al. (2010) Deciphering the splicing code. *Nature* 465(7294):53–59.
- Mézard M, Parisi G, Zecchina R (2002) Analytic and algorithmic solution of random satisfiability problems. *Science* 297(5582):812–815.
- Liu YY, Slotine JJ, Barabási AL (2011) Controllability of complex networks. *Nature* 473(7346):167–173.
- Bohn S, Magnasco MO (2007) Structure, scaling, and phase transition in the optimal transport network. *Phys Rev Lett* 98(8):088702.
- Banavar JR, Colaiori F, Flammini A, Maritan A, Rinaldo A (2000) Topology of the fittest transportation network. *Phys Rev Lett* 84(20):4745–4748.
- Shao Z, Zhou H (2007) Optimal transportation network with concave cost functions: Loop analysis and algorithms. *Phys Rev E Stat Nonlin Soft Matter Phys* 75(6 Pt 2):066112.
- Batchelor MT, Nienhuis B, Warnaar SO (1989) Bethe-Ansatz results for a solvable O(n) model on the square lattice. *Phys Rev Lett* 62(21):2425–2428.
- Stilck JF, Machado KD, Serra P (1996) Nature of the collapse transition for polymers. *Phys Rev Lett* 76(15):2734–2737.
- Mézard M, Montanari A (2009) *Information, Physics, and Computation* (Oxford Univ Press, Oxford, UK).
- Yeung CH, Saad D (2012) Competition for shortest paths on sparse graphs. *Phys Rev Lett* 108(20):208701.
- Transport for London, Oyster card data. Available at www.tfl.gov.uk. Accessed April 16, 2012.
- Openflights.org, Airport network data. Available at <http://openflights.org/data.html>. Accessed April 12, 2012.
- Newman MEJ, Strogatz SH, Watts DJ (2001) Random graphs with arbitrary degree distributions and their applications. *Phys Rev E Stat Nonlin Soft Matter Phys* 64(2 Pt 2):026118.
- Kim M, Medard M (2004) Robustness in large-scale random networks. *Proceedings of the IEEE Infocom Conference. 23rd Annual Joint Conference of the IEEE Computer and Communications Societies* (IEEE, Washington, DC), pp 2364–2373.
- Bollobás B (1985) *Random Graphs* (Academic, London).
- Borgatti SP (2002) NetDraw Software for Network Visualization. Analytic Technologies: Lexington, KY.