On the Capacity of Small-World Networks

Rui A. Costa João Barros

***Abstract*— Recent results from statistical physics show that large classes of complex networks, both man-made and of natural origin, are characterized by high clustering properties yet strikingly short path lengths between pairs of nodes. Breaking with the traditional approach to these so called *small worlds* that relies mainly on graph parameters directly related to connectivity, we investigate the *capacity* of these networks from the perspective of network information flow. Our contribution includes upper and lower bounds for the capacity of standard and navigable small-world models based on added shortcuts, and the somewhat surprising result, that, with high probability, random rewiring does not alter the capacity of a small-world network.**

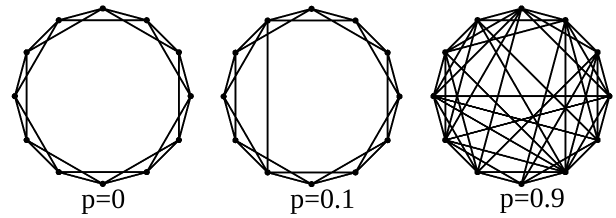

Fig. 1. Small-World model with added shortcuts for different values of the adding probability p.

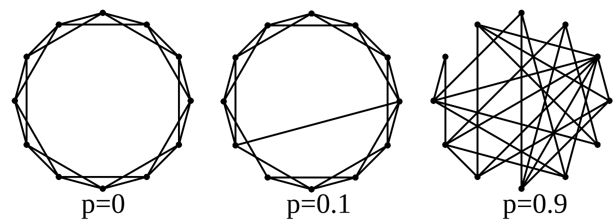

Fig. 2. Small-World model with rewiring for different values of the rewiring probability p.

I. Introduction

Small-World graphs, i.e. graphs with high clustering coefficients and small average path length, have sparked a fair amount of interest from the scientific community, mainly due to their ability to capture fundamental properties of relevant phenomena and structures in sociology, biology, statistical physics and man-made networks. Beyond well-known examples such as Milgram's "six degrees of separation" [1] between any two people in the United States (over which some doubt has recently been casted [2]) and the Hollywood graph with links between actors, small-world structures appear in such diverse networks as the U.S. electric power grid, the nervous system of the nematode worm *Caenorhabditis elegans* [3], food webs [4], telephone call graphs [5], citation networks of scientists [6], and, most strikingly, the World Wide Web [7].

The term small-world graph itself was coined by Watts and Strogatz, who in their seminal paper [8] defined a class of models which interpolate between regular lattices and random Erdös-Rényi graphs by adding shortcuts or rewiring edges with a certain probability p (see Figures 1 and 2). The most striking feature of these models is that for increasing values of p the average shortest-path length diminishes sharply, whereas the clustering coefficient remains practically constant during this transition.

Since then, most contributions in the area of complex networks focus essentially on connectivity parameters such as the degree distribution of the nodes, the clustering coefficient of the graph, the shortest path length between two nodes, or the *betweenness* of a node (i.e. the total number of shortest paths that pass trough it). In spite of its arguable relevance — particularly where communication networks are concerned — the *capacity* of small-world networks has, to the best of our knowledge, not yet been studied in any depth by the scientific community.

The authors are with the Laboratory of Artificial Intelligence and Computer Science (LIACC/UP) and the Department of Computer Science at the University of Porto, Porto, Portugal. URL: `http://www.dcc.fc.up.pt/˜barros/`.

The main goal of this paper is thus to provide a preliminary characterization of the capacity of small-world networks from the point of view of network information flow. Our main contributions are as follows:

- *Capacity Bounds on Small-World Networks with Added Shortcuts:* We prove a high concentration result which gives upper and lower bounds on the capacity of a Small-World with shortcuts of probability p, thus describing the capacity growth due to the addition of random edges.
- *Capacity bounds for Navigable Small-World Networks:* We define a navigable small world network inspired by [9], which allows for efficient distributed routing, and prove a high concentration result for its capacity, as well.
- *Rewiring does not alter the Capacity:* We construct tight upper and lower bounds for the capacity of small worlds with rewiring and prove that, with high probability, capacity will not change when the edges are altered in a random fashion.

The rest of the paper is organized as follows. Sec. II gives an overview of related work pertaining the capacity of communication networks and properties of small-world models. Then, in Sec. III, we provide precise definitions for the three small-world models of interest in this work, so that the main results can be stated and proved in Sec. IV. Finally, Sec. V offers some concluding remarks.

II. Related Work

Although the capacity of networks (described by general graphs with or without edge capacities) supporting multiple communicating parties is largely unknown, progress has recently been reported in several relevant instances of this problem. In the case where the network has one or more independent sources of information but only one sink, it is known that routing offers an optimal solution for transporting messages [10] — in this case the transmitted information behaves like *water in pipes* and the capacity can be obtained by classical network flow methods. Specifically, the capacity of the network follows from the well-known Ford-Fulkerson *max-flow min-cut* theorem [11], which asserts that the maximal amount of a flow (provided by the network) is equal to the capacity of a minimal cut, i.e. a nontrivial partition of the graph vertex set V into two parts such that the sum of the capacities of the edges connecting the two parts (the cut capacity) is minimum. In [12] it was shown that network flow methods also yield the capacity for networks with multiple *correlated* sources and one sink.

The case of general multicast networks, in which a single source broadcasts a number of messages to a set of sinks, is considered in [13], where it is shown that applying coding operations at intermediate nodes (i.e. *network coding*) is necessary to achieve the max-flow/min-cut bound of the network. In other words, if k messages are to be sent then the minimum cut between the source and each sink must be of size at least k. A converse proof for this problem, known as the *network information flow problem*, was provided by [14], whereas linear network codes were proposed and discussed in [15] and [16]. Max-flow min-cut capacity bounds for Erdös-Rényi graphs and random geometric graphs were presented in [17].

Another problem in which network flow techniques have been found useful is that of finding the maximum stable throughput in certain networks. In this problem, posed by Gupta and Kumar in [18], it is sought to determine the maximum rate at which nodes can inject bits into a network, while keeping the system stable. This problem was reformulated in [19] as a multicommodity flow problem, for which tight bounds were obtained using elementary counting techniques.

Since the seminal work of [8], key properties of small-world networks, such as clustering coefficient, characteristic path length, and vertex degree distribution, have been studied by several authors (see e.g. [20] and references therein). The combination of strong local connectivity and long-range shortcut links renders small-world topologies potentially attractive in the context of communication networks, either to increase their capacity or simplify certain tasks. Recent examples include resource discovery in wireless networks [21], design of heterogeneous networks [22], [23], and peer-to-peer communications [24].

When applying small-world principles to communication networks, we would like not only that short paths exist between any pairs of nodes, but also that such paths can easily be found using merely local information. In [9] it was shown that this *navigability* property, which is key to the existence of effective distributed routing algorithms, is lacking in the small-world models of [8] and [25]. The alternative navigable model presented in [9] consists of a grid to which shortcuts are added not uniformly but according to a harmonic distribution, such that the number of outgoing links per node is fixed and the link probability depends on the distance between the nodes. For this class of small-world networks a *greedy* routing algorithm, in which a message is sent through the outgoing link that takes it closest to the destination, was shown to be effective, thus opening the door towards a capacity-attaining solution.

III. Small-World Models

We start by presenting rigorous definitions for the three small-world models used in the rest of the paper. For convenience, all of these models are constructed based on a ring lattice, but it is worth pointing out that the presented methodology can be equally applied to other classes of base lattices. In the following, we also assume that all edges have unitary weight. Before proceeding with the model descriptions, we require a precise notion of distance in a ring.

Definition 1: Consider a set of n nodes connected by edges that form a ring (see Fig. 3, left plot). The *ring distance* between two nodes is defined as the minimum number of *hops* from one node to the other. If we number the nodes in clockwise direction, starting from any node, then the ring distance between nodes i and j is given by $d(i,j) = min\{|i-j|, n+i-j, n-|i-j|\}$.

For simplicity, we refer to $d(i,j)$ as the *distance* between i and j. Next, we define a k-connected ring lattice.

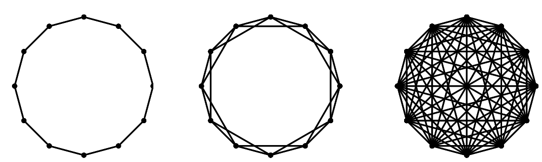

Fig. 3. Illustration of a k-connected ring lattice: from left to right $k = 2, 4, 12$.

Definition 2: A *k-connected ring lattice* (see Fig. 3) is a graph $L = (V_L, E_L)$ with nodes V_L and edges E_L, in which all nodes in V_L are placed on a ring and are connected to all the nodes within distance $\frac{k}{2}$.

Notice that in the definition of a k-connected ring lattice, all the nodes have degree k. Based on this topology, we can construct the following models (see Figs. 1 and 2).

Definition 3 (Small-World Network with Shortcuts [25]): Consider a k-connected ring lattice $L = (V_L, E_L)$ and let E_C be the set of all possible edges between nodes in V_L (i.e. (V_L, E_C) is a fully connected graph). To obtain a *small-world network with shortcuts*, we add to the ring lattice L each edge $e \in E_C \backslash E_L$ with probability p.

Definition 4 (Navigable Small-World Network): Starting with a k-connected ring lattice, add one edge to each node i randomly according to the probability distribution $p(i,j) = d(i,j)^{-r}$, where $d(i,j)$ denotes the distance between nodes i and j and $r > 0$ is a fixed parameter.

Definition 5 (Small-World Network with Rewiring [8]): Consider a k-connected ring lattice and choose a vertex and the edge that connects it to its nearest neighbor in a clockwise sense. With probability p, reconnect this edge to a vertex chosen uniformly at random over the entire ring. Repeat this process by moving around the ring in clockwise direction, considering each vertex in turn until one lap is completed. In each step ensure that none of the edges is duplicated and that no edge that was removed is placed again in the graph. Next, consider the edges that connect vertices to their second-nearest neighbors clockwise. As before, randomly rewire each of these edges with probability p, and continue this process, circulating around the ring and proceeding outward to more distant neighbors after each lap, until each edge in the original lattice has been considered once.

IV. Capacity of Small-World Networks

In Sec. II, we argued that the max-flow min-cut capacity provides the fundamental limit of communication for various relevant network scenarios. Motivated by this observation, we will now use network flow methods and random sampling techniques in graphs to derive a set of bounds for the capacity of the small-world network models presented in the previous section.

A. Preliminaries

We start by introducing some notation. Let G be an undirected and unweighted graph and let G_s be the graph obtained by sampling on G, such that each edge e has sampling probability p_e. From G and G_s, we obtain G_w by assigning to each edge e the weight p_e, i.e. $w(e) = p_e, \forall e$. We denote the global minimum cuts of G_s and G_w by c_s and c_w, respectively. It is helpful to view a cut in G_s as a sum of Bernoulli experiences, whose outcome determines if and edge e connecting the two sides of the cut belongs to G_s or not. It is not difficult to see that the value of a cut in G_w is the expected value of the same cut in G_s.

The next theorem gives a characterization of how close a cut in G_s will be with respect to its expected value.

Theorem 1 (From [26]): Let $\epsilon = \sqrt{2(d+2)\ln(n)/c_w}$. Then, with probability $1 - O(1/n^d)$, every cut in G_s has value between $(1-\epsilon)$ and $(1+\epsilon)$ times its expected value.

Notice that although d is a free parameter, there is a strict relationship between the value of d and the value of ϵ. In other words, the proximity to the expected value of the cut is intertwined with how close the probability is to one. *Theorem 1* yields also the following useful property.

Corollary 1: Let $\epsilon = \sqrt{2(d+2)\ln(n)/c_w}$. Then, with high probability, the value of c_s lies between $(1-\epsilon)c_w$ and $(1+\epsilon)c_w$.

Before using the previous random sampling results to determine bounds for the capacities of small-world models, we prove another useful lemma.

Lemma 1: Let $L = (V_L, E_L)$ be a k-connected ring lattice and let $G = (V_L, E)$ be a fully connected graph, in which edges $e \in E_L$ have weight $w_1 \geq 0$ and edges $f \notin E_L$ have weight $w_2 \geq 0$. Then, the global minimum cut in G is $k \cdot w_1 + (n-1-k) \cdot w_2$.

Proof: We start by splitting G into two subgraphs: a k-connected ring lattice L with weights w_1 and a graph F with nodes V_L and all remaining edges of weight w_2. Clearly, the value of a cut in G is the sum of the values of the same cut in L and in F. Moreover, both in L and in F, the global minimum cut is a cut in which one of the partitions consists of one node (any other partition increases the number of outgoing edges). Since each node in L has k edges of weight w_1 and each node in F has the remaining $n-1-k$ edges of weight w_2, the result follows. ■

B. Capacity of Small-World Networks with Added Shortcuts

With this set of tools, we are to state and prove our first main result.

Theorem 2: With high probability, the value of the capacity of a small-world network with added shortcuts lies between $(1-\epsilon)c_w$ and $(1+\epsilon)c_w$, with $\epsilon = \sqrt{2(d+2)\ln(n)/c_w}$ and $c_w = k + (n-1-k)p$.

Proof: Let G_w be a fully connected graph with n nodes and with the edge weights (or equivalently, the sampling probabilities) defined as follows:

- The weight of the edges in the initial lattice of a small-world network with added shortcuts is one (because they are not removed);
- The weight of the remaining edges is p, (i.e. the probability that an edge is added).

Notice that G_w is a graph in the conditions of Lemma 1, with $w_1 = 1$ and $w_2 = p$. Therefore, the global minimum cut in G_w is $c_w = k + (n-1-k)p$, where k is the initial number of neighbors in the lattice. Using Corollary 1, the result follows. ■

The obtained bounds are illustrated in Fig. 4.

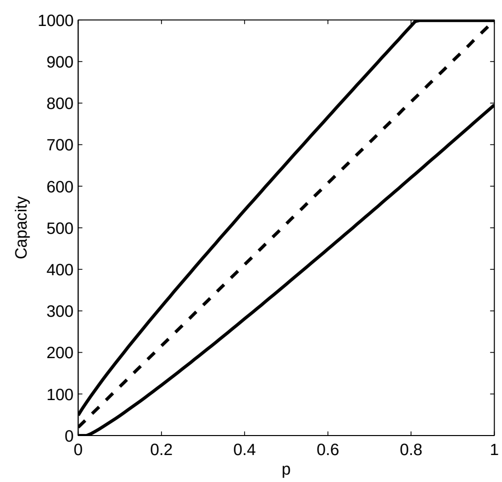

Fig. 4. Bounds on the capacity of a small-world network with added shortcuts, for $n = 1000$, $k = 20$, and $d = 1$. The dashed line represents the expected value of the capacity, and the solid lines represent the bounds. Naturally, the capacity increases with p, as the number of added links become larger.

C. Capacity of a Navigable Small-World Network

For the class of navigable small-world networks defined in Sec. III, in which the probability of a given edge being added depends on the distance between the nodes, we obtain the following result.

Theorem 3: With high probability, the capacity of the navigable small-world network has a value in the interval

$[(1-\epsilon)c_w),(1+\epsilon)c_w)]$, with $\epsilon=\sqrt{2(d+2)\ln(n)/c_w}$ and

$$c_w = k + (1+a_n)\cdot(\frac{n-a_n}{2})^{-r} + 2\cdot\sum_{i=k+1}^{\frac{n-a_n}{2}-1} i^{-r},$$

where $a_n=\frac{1-(-1)^n}{2}$.

Proof: Consider the fully connected graph $G_w=(V_L,E)$ with weights defined as follows: the weights of edges $(i,j)\in E_L$ is set to one and those of $(i,j)\notin E_L$ are equal to $w(i,j)=d(i,j)^{-r}$, i.e. the probability of adding edge (i,j).

Notice that the ring distance between two nodes does not depend on which node is numbered first. It is therefore correct to state that all the nodes have the same number of nodes at distance d. We also have that all the edges in the ring lattice unitary weight. Based on these two observations and the fact that G_w is a fully connected graph, it is clear that the global minimum cut in G_w, denoted c_w, is a cut in which one of the partitions consists of a single node, say node 1. Thus, we may write

$$c_w = k + \sum_{i\in A} d(1,i)^{-r},$$

with

$$A=\{i:(1,i)\notin E_L\}=\{i:d(1,i)>k\}.$$

Now, we must distinguish between two different situations: even n and odd n. If n is even, it is not difficult to see that the single node that maximizes the distance to node 1 is node $\frac{n}{2}+1$, with $d(1,\frac{n}{2}+1)=\frac{n}{2}$. Notice that, for distances d inferior to $\frac{n}{2}$, there are two nodes at a distance d to node 1. Therefore, if n is even, we have

$$c_w = k + (\frac{n}{2})^{-r} + 2\cdot\sum_{i=k+1}^{\frac{n}{2}-1} i^{-r}$$

. When n is odd, it is also easy to see that there are two nodes that maximize the distance to node 1, nodes $\frac{n+1}{2}$ and $\frac{n+3}{2}$, with the maximum distance being $\frac{n-1}{2}$. Therefore, if n is odd,

$$c_w = k + 2\cdot\sum_{i=k+1}^{\frac{n-1}{2}} i^{-r}$$

Using Corollary 1 and observing that $a_n=\frac{1-(-1)^n}{2}$ is equal to 1 if n is odd and equal to 0 if n is even, we obtain the desired bounds. ∎

The result is illustrated in Fig. 5.

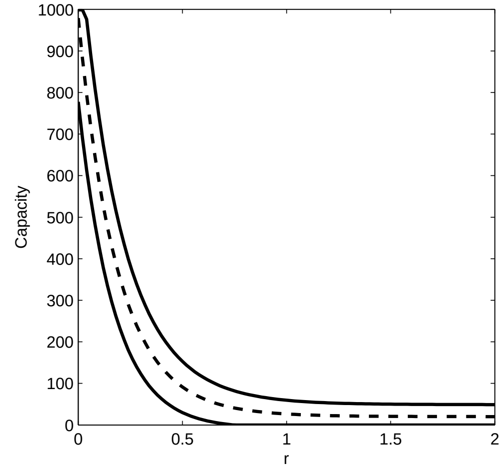

Fig. 5. Bounds for the capacity of a navigable small-world network for $n=1000$, $k=20$, $d=1$, and different values of parameter r. The dashed line represents the expected value of the capacity and the solid lines represent the bounds. As expected, the capacity decreases sharply with higher r, because increasing r decreases the probability of adding new edges.

D. Capacity of Small-World Networks with Rewiring

In the previous classes of small-world networks, edges were added to a k-connected ring lattice (with minimum cut k) and clearly the capacity could only grow with p. The next natural step is to ask what happens when edges are not added but rewired with probability p, as described in Sec. III. Before presenting a theorem that answers this question, we will prove the following lemma.

Lemma 2: Let G_w be a weighted, fully connected graph, whose weights correspond to the edge probabilities of a small-world network with rewiring, and let c_w be the global minimum cut in G_w. Then, $c_w\geq k$.

Proof: We start with the initial lattice edges $(l,m)\in E_L$, and assign the weight $1-p$ to their counterparts in G_w. In order to determine the weight of the non-initial edges that result from rewiring, consider the following events:

- $R(i,j)$: "Rewire the edge $(i,j)\in E_L$";
- $C_i(j,l)$: "Rewire $(i,j)\in E_L$ to $(i,l)\notin E_L$".

Notice that $\mathcal{P}(R(i,j))=p,\forall i,j$.

Let i and j be two non-initially connected nodes. The notation $i\leftrightarrow j$ denotes the event that the nodes i and j are connected.

$$\begin{aligned}
\mathcal{P}(i\leftrightarrow j) =& \mathcal{P}([\cup_{x=1}^{k/2}(R(i,i+x)\cap C_i(i+x,j))] \\
&\cup[\cup_{x=1}^{k/2}(R(j,j+x)\cap C_j(j+x,i))]) \\
=& \sum_{x=1}^{k/2}(\mathcal{P}(R(i,i+x)\cap C_i(i+x,j)) \\
&+\mathcal{P}(R(j,j+x)\cap C_j(j+x,i))) \\
=& \sum_{x=1}^{k/2}(\mathcal{P}(C_i(i+x,j)|R(i,i+x))\mathcal{P}(R(i,i+x)) \\
&+\mathcal{P}(C_j(j+x,i)|R(j,j+x))\mathcal{P}(R(j,j+x))) \\
=& p\cdot(\sum_{x=1}^{k/2}(\mathcal{P}(C_i(i+x,j)|R(i,i+x)) \\
&+\mathcal{P}(C_j(j+x,i)|R(j,j+x))))
\end{aligned}$$

We have $\mathcal{P}(C_i(i+x,j)|R(i,i+x))=\frac{1}{m}$, where m is the number of possible new connections from node i when we rewired the edge $(i,i+x)$. It is possible that, occurring some rewiring or not, none of the choices to a new link is the node i. In this case, $m=n-k-1$. Notice that this is the highest it can get, therefore $m\leq n-k-1$. Then

$$\mathcal{P}(C_i(i+x,j)|R(i,i+x))\geq\frac{1}{n-k-1}.$$

Analogously, $\mathcal{P}(C_j(j+x,i)|R(j,j+x)))\geq\frac{1}{n-k-1}$. Therefore,

$$\mathcal{P}(i\leftrightarrow j)\geq p\cdot(\sum_{x=1}^{k/2}\frac{2}{n-k-1})=\frac{pk}{n-k-1}.$$

There are k initial edges and $n-k-1$ non-initial edges in each node.

Consider a fully connected, weighted graph F with the weights defined as follows: all the edges $(i,j) \notin E_L$ have the weight $\frac{pk}{n-k-1}$; and all the others edges $(i,j) \in E_L$ have the weight $1-p$. Notice that F is a graph in the conditions of Lemma 1, with $w_1 = 1-p$ and $w_2 = \frac{pk}{n-k-1}$. Therefore, $c_F = k(1-p) + (n-k-1)\frac{pk}{n-k-1} = k$. Notice that, in this situation, all the weights in F are a lower bound of the weights in G_w. Therefore, a cut in F is a lower bound for the corresponding cut in G_w. Then, the global minimum cut in F is a lower bound for all the cuts in G_w, in particulary, for c_w: $c_w \geq c_F = k$. ■

With this lemma, we are now ready to state and prove our last result.

Theorem 4 (Rewiring does not alter capacity.): With high probability, the capacity of a small-world network with rewiring, c_s, verifies $c_s \geq (1-\epsilon)k$, with $\epsilon = \sqrt{2(d+2)ln(n)/k}$.

Proof: Based on Lemma 2 and Corollary 1, we have that, with high probability, $c_s \geq (1-\epsilon_w)k$, with $\epsilon_w = \sqrt{2(d+2)ln(n)/c_w}$. Now, from the fact that $c_w \geq k$, we have that $\epsilon = \sqrt{2(d+2)ln(n)/k} \geq \epsilon_w$. Then, $(1-\epsilon_w)k \geq (1-\epsilon)k$, and the first part of the result follows.

Next, we prove by contradiction that $c_s \leq k$. Suppose that the proposition $c_s > k$ is true. Let c_i be the cut in which one of the partitions consists of node i, $i = 1, ..., n$. Because c_s is the global minimum cut in G_s, we have that $c_i > k$, $\forall i = 1, ..., n$. Notice that c_i is the degree of node i. Then, because in the k-connected ring lattice all nodes have degree k and all nodes in G_s have degree greater than k (because $c_i > k, \forall i$), we have that the number of edges in G_s must be greater than the number of edges in the k-connected ring lattice. But this is clearly absurd, because in the construction of G_s, we do not add new edges to the k-connected ring lattice, we just rewire some of the existent edges. ■

V. Concluding Remarks

We studied the max-flow min-cut capacity of three fundamental classes of small world networks. Using classical network flow arguments and concentration results from random sampling in graphs, we provided bounds for both standard and navigable small-world networks with added shortcuts. In addition, we presented a tight result for small-world networks with rewiring, which permits the following interpretation: *With high probability, rewiring does not alter the capacity of the network*. This observation is not obvious, because we can easily find ways to rewire the ring lattice in order to obtain, for instance, a *bottleneck*. But, according to the previous results, such instances occur with very low probability.

Possible directions for future work include tighter capacity results, extensions to other classes of small-world networks (e.g. weighted models and those used in peer-to-peer networks [24]), and understanding if and how small-world topologies can be exploited in the design of capacity-attaining network codes. At a more conceptual level, we are intrigued by the possibility that the notion of capacity may help us answer a very central question: *why* small-world topologies are ubiquitous in real-world networks.

References

[1] S. Milgram, "Phychology today," *Physical Review Letters*, vol. 2, pp. 60–67, 1967.

[2] J. S. Kleinfeld, "Could it be a big world after all?," Society, 2002.

[3] T.B. Achacoso and W.S. Yamamoto, *AY's Neuroanatomy of C. elegans for Computation*, CRC Press, 1992.

[4] R.J. Williams and N.D. Martinez, "Simple rules yield complex food webs," *Nature*, vol. 404, pp. 180–183, 2000.

[5] James Abello, Adam L. Buchsbaum, and Jeffery Westbrook, "A functional approach to external graph algorithms," in *ESA '98: Proceedings of the 6th Annual European Symposium on Algorithms*, London, UK, 1998, pp. 332–343, Springer-Verlag.

[6] M.E.J. Newman, "The structure of scientific collaboration networks," in *Proc. Natl Acad. Sci.*, 2001, vol. 98, pp. 404–409.

[7] A. Broder, "Graph structure in the web," *Comput. Netw.*, vol. 33, pp. 309–320, 2000.

[8] Duncan J. Watts and Steven H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, no. 6684, June 4 1998.

[9] Jon Kleinberg, "The small-world phenomenon: an algorithm perspective," in *STOC '00: Proceedings of the thirty-second annual ACM symposium on Theory of computing*, New York, NY, USA, 2000.

[10] April Rasala Lehman and Eric Lehman, "Complexity classification of network information flow problems," in *SODA '04: Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, Philadelphia, PA, USA, 2004, pp. 142–150, Society for Industrial and Applied Mathematics.

[11] L.R. Ford and D.R. Fulkerson, *Flows in Networks*, Princeton University Press, Princeton, NJ, 1962.

[12] J. Barros and S. D. Servetto, "Network Information Flow with Correlated Sources," Accepted for publication in the IEEE Transactions on Information Theory, October 2005. Available at `http://www.dcc.fc.up.pt/ barros0`.

[13] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network Information Flow," *IEEE Trans. Inform. Theory*, vol. 46, no. 4, pp. 1204–1216, 2000.

[14] S. Borade, "Network Information Flow: Limits and Achievability," in *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, Lausanne, Switzerland, 2002.

[15] S.-Y. R. Li, R. W. Yeung, and N. Cai, "Linear Network Coding," *IEEE Trans. Inform. Theory*, vol. 49, no. 2, pp. 371–381, 2003.

[16] R. Koetter and M. Médard, "An Algebraic Approach to Network Coding," *IEEE/ACM Trans. Networking*, vol. 11, no. 5, pp. 782–795, 2003.

[17] Aditya Ramamoorthy, Jun Shi, and Richard D. Wesel, "On the capacity of network coding for random networks," in *Proc. of 41st Allerton Conference on Communication, Control and Computing*, Allerton, Illinois, October 2003.

[18] P. Gupta and P. R. Kumar, "The Capacity of Wireless Networks," *IEEE Trans. Inform. Theory*, vol. 46, no. 2, pp. 388–404, 2000.

[19] C. Peraki and S. D. Servetto, "On Multiflows in Random Unit-Disk Graphs, and the Capacity of Some Wireless Networks," Submitted to the IEEE Trans. Inform. Theory, March 2005. Available from `http://cn.ece.cornell.edu/`.

[20] S.N. Dorogovtsev and J.F.F. Mendes, *Evolution of Networks: From Biological Nets to the Internet and WWW*, Oxford University Press, 2003.

[21] A. Helmy, "Small worlds in wireless networks," *IEEE Communications Letters*, vol. 7, no. 10, pp. 490–492, October 2003.

[22] Alex Reznik, Sanjeev R. Kulkarni, and Sergio Verdu, "A small world approach to heterogeneous networks," *Communication in Information and Systems*, vol. 3, no. 4, pp. 325–348, 2004.

[23] S. Dixit, E. Yanmaz, and O. K. Tonguz, "On the design of self-organized cellular wireless networks," *IEEE Communications Magazine*, vol. 43, no. 7, pp. 86–93, 2005.

[24] G. S. Manku, M. Naor, and U. Wieder, "Know thy neighbor's neighbor: The power of lookahead in randomized p2p networks," in *Proceedings of STOC*, 2004.

[25] Mark E. J. Newman and Duncan J. Watts, "Scaling and percolation in the small-world network model," *Physical Review E*, vol. 60, no. 6, pp. 7332–7342, December 1999.

[26] David R. Karger, "Random sampling in cut, flow, and network design problems," in *STOC '94: Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, New York, NY, USA, 1994.