

Network Information Flow in *Navigable* Small-World Networks

Rui A. Costa

João Barros

Abstract—Small-world graphs, exhibiting high clustering coefficients and small average path length, have been shown to capture fundamental properties of a large number of natural and man-made networks. In the context of communication networks, *navigable* small-world topologies, i.e. those which admit efficient distributed routing algorithms, are deemed particularly effective, for example in resource discovery tasks and peer-to-peer applications. Intrigued by the fundamental limits of communication in networks that exploit this type of topology, we study two classes of navigable small-world networks from the point of view of network information flow and provide inner and outer bounds for their max-flow min-cut capacity. Our contribution is in contrast with the standard approach to small world networks which privileges parameters pertaining to connectivity.

I. INTRODUCTION

Random graphs play an important role as mathematically tractable models for complex, large-scale networks. The most recent addition to this set of tools is a class of objects generally designated by *small world* graphs, which exhibit high clustering coefficients (i.e. neighboring nodes are likely to be connected) and small average path length — the diameter of a graph with n nodes is in fact bounded by a polynomial in $\log n$. The term small-world graph itself was coined by Watts and Strogatz, who in their seminal paper [1] defined a class of models which interpolate between regular lattices and random Erdős-Rényi graphs by adding long-range shortcuts with a certain probability p , as illustrated in Fig. 1. The most salient 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.




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

Since their discovery, small-world graphs have been shown to capture fundamental properties of relevant phenomena and structures in sociology, biology, statistical

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/>.




Fig. 2. Kleinberg small-world model. Each node is directly connected to all neighbors within h hops, and also to q more distant nodes through so called shortcuts. In the shown example, where $h = q = 1$, lightly shaded circles represent the nodes that are directly connected to node (i, j) , i.e. the four direct neighbors of (i, j) and one additional node (k, l) connected by a shortcut.

physics and man-made networks, with examples ranging from Milgram’s “six degrees of separation” [2] between any two people in the United States to 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]. It is therefore not surprising that key attributes of small-world networks, 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 through it), have become the focus of intense research (see e.g. [8] 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 capacity or simplify certain tasks. Recent examples include resource discovery in wireless networks [9] and design of heterogeneous networks [10], [11]. Another relevant application is related to overlay networks for peer-to-peer communications, for which small world properties are deemed to be particularly useful [12].

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 [13] Kleinberg showed that this *navigability* property, which is key to the existence of effective distributed routing algorithms, is lacking in the small-world models of [1] and [14].

The alternative navigable model presented in [13] 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, as illustrated in Fig. 2. 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 information flow in a distributed fashion.

Motivated by their potential to improve the transfer of data in networks with multiple parties, we set out to investigate the fundamental limits of network information flow in small world networks. In [15], we focused on the original (non-navigable) models of [1] and [14], and proved a high concentration result that gives upper and lower bounds on the max-flow min-cut capacity of said networks. The main goal of this paper is to provide a preliminary characterization of the capacity of *navigable* small-world networks, for which highly efficient distributed routing algorithms are known to exist and distributed network coding strategies (allowing processing at intermediate nodes) are likely to be found. Our main contributions are as follows:

- *Capacity Bounds for Kleinberg Small-World Networks:* We construct upper and lower bounds for the max-flow min-cut capacity of Kleinberg graphs derived from a square lattice and illustrate how the choice of connectivity parameters affects communication.
- *Capacity Bounds for Navigable Small-World Networks on Ring Lattices:* Arguing that the corners present in Kleinberg’s models introduced undesirable artefacts in the computation of the capacity, we define a navigable small world network based on a ring lattice and derive a high-concentration result for the capacity of this instance, as well.

The rest of the paper is organized as follows. Sec. II gives an overview of related work pertaining the capacity of communication networks. Then, Sec. III provides precise definitions for the two small-world models of interest in this work, so that the main results can be stated and proved in Sec. IV. The paper concludes with Sec. V.

II. OTHER 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 [16] — 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 [17], 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 [18] 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 [19], 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 [20], whereas linear network codes were proposed and discussed in [21] and [22]. Max-flow min-cut capacity bounds for Erdős-Rényi graphs and random geometric graphs were presented in [23].

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 [24], 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 [25] as a multicommodity flow problem, for which tight bounds were obtained using elementary counting techniques.

Since small world graphs were proposed as models for complex networks [1] and [14], most contributions in the area of complex networks focus essentially on connectivity parameters such as the degree distribution, the clustering coefficient or the shortest path length between two nodes (see e.g. [26]). 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.

III. NAVIGABLE SMALL-WORLD NETWORKS

We start by presenting rigorous definitions for the two small-world models used in the rest of the paper. In the following, we also assume that all edges have unitary weight.

Definition 1 (Kleinberg Small-World graph, see Fig. 2): We begin from a two-dimensional grid and a set of nodes that are identified with the set of lattice points in an $n \times n$ square, $\{(x, y) : x \in \{1, 2, \dots, n\}, y \in \{1, 2, \dots, n\}\}$, and we define the *lattice distance* between two nodes (x, y) and (w, z) to be the number of *lattice steps* (or hops) separating them: $d((x, y), (w, z)) = |w - x| + |z - y|$. For a constant $h \geq 1$, the node (u_1, u_2) is connected to every other node within lattice distance h (we denote the set of this initial edges as E_1). For universal constants $q \geq 0$ and $r \geq 0$, we also construct edges from (u_1, u_2) to q other nodes using independent random trials; the i^{th} edge from (u_1, u_2) has endpoint (v_1, v_2) with probability proportional to $[d((u_1, u_2), (v_1, v_2))]^{-r}$. To ensure a valid

probability distribution, we consider the set of nodes that are not connected with (u_1, u_2) in the initial lattice, i.e.

$$N_{(u_1, u_2)} = \{(v_1, v_2) : d((u_1, u_2), (v_1, v_2)) > h\}$$

and divide $[d((u_1, u_2), (v_1, v_2))]^{-r}$ by the appropriate normalizing constant

$$s(u_1, u_2) = \sum_{(v_1, v_2) \in N_{(u_1, u_2)}} [d((u_1, u_2), (v_1, v_2))]^{-r}.$$

In the next section, we will see that this model exhibits unexpected effects related to the corners of the chosen base lattice. Motivated by this observation, we construct a somewhat different model, which uses a ring lattice but still keeps the key relationship between shortcut probability and node distance that assures the navigability of the model. Before proceeding with its definition, we require a precise notion of distance in a ring.

Definition 2: 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 $d(i, j) = \min\{|i - j|, n + i - j, n - |i - j|\}$.

Since the type of distance will always be clear from the context, using the same notation for ring distance and lattice distance does not cause confusion. 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 3: 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 now present a precise definition of an alternative model.

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.

IV. MAX-FLOW MIN-CUT BOUNDS IN NAVIGABLE 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 an 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 [27]): 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 $(1 + \epsilon)c_w$.

B. Capacity Bounds for Navigable Small-World Networks based on Ring Lattices

We start with the somewhat simpler class of navigable small-world networks based on ring lattices and prove the following result.

Theorem 2: 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

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

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 h . 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

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




Fig. 4. 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.

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 h inferior to $\frac{n}{2}$, there are two nodes at a distance h to node 1. Therefore, if n is even, we have

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

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. 4.

C. Capacity of Kleinberg Small World Networks

Before proceeding with the bounds for the capacity of Kleinberg small-world networks, we require an algorithm to calculate the normalizing constants

$$s(x, y) = \sum_{(i,j) \in N(x,y)} [d((x, y), (i, j))]^{-r},$$

for $x, y \in \{1, \dots, n\}$. For this purpose, we note that the previous sum can be written as

$$\begin{aligned} s(x, y) &= \sum_{(i,j) \neq (x,y)} (|i-x| + |j-y|)^{-r} \\ &\quad - \sum_{(i,j) \notin N(x,y)} [d((i, j), (x, y))]^{-r}. \end{aligned}$$

Clearly, the first term can be easily calculated. Thus, the challenging task is to present an algorithm that deals with the calculation of $\sum_{(i,j) \notin N(x,y)} [d((i, j), (x, y))]^{-r}$. The nodes $(i, j) \notin N(x, y)$ are the nodes initially connected to node (x, y) , i.e., the nodes at a distance $t \leq h$ from node (x, y) . It is not difficult to see that the nodes at a distance t from node (x, y) are the nodes in the square line formed by the nodes $(x-t, y)$, $(x+t, y)$, $(x, y+t)$ and $(x, y-t)$. Then, we could just look at nodes in the square formed by the nodes $(x-h, y)$, $(x+h, y)$, $(x, y+h)$ and $(x, y-h)$ and sum all the corresponding distances to node (x, y) . A corner effect occurs when this square lies outside the base lattice. Assume that we start by calculating the distances to the nodes in line $y+i$, with $i \geq 0$.

To avoid calculating extra distances (i.e., distances of nodes that are out of the grid), we must make sure that this line verifies $y+i \leq n$ and also $y+i \leq h$. For this reason, i must vary according to $i \in \{0 \dots \min\{h, n-y\}\}$. Now, in each line $y+i$, we first look at the nodes in the right side of (x, y) , i.e., we calculate the distances of the nodes $(x+j, y+i)$, with $j \geq 0$. Now, notice that in the line y , we have h points on the right side of (x, y) that are in the square (regardless of whether they are in the grid). Because the distance is the minimum number of steps in the grid, we have that in line $y+i$ there are $h-i$ points at the right side of (x, y) that are inside the square. This way, j must vary according to $j \in \{0 \dots \min\{h-i, n-x\}\}$. Now, when looking at the nodes at the left side (i.e., the nodes $(x-j, y+i)$, with $i \geq 1$), the idea is the same. The only difference is that, in this case, the variation for j is $j \in \{1 \dots \min\{h-i, x-1\}\}$. Then, we proceed analogously for the lines below (x, y) , i.e., the lines $y-i$, with $i \in \{1 \dots \min\{h, y-1\}\}$. This algorithm is summarized in Table I. The matrix z is a buffer for the distances, i.e., $z(u_1, u_2) = d((x, y), (u_1, u_2))$. We impose $z(x, y) = 0$, because $d((x, y), (x, y))^{-r}$ is also calculated in this procedure.

The following quantities will be instrumental towards

TABLE I

ALGORITHM FOR COMPUTING NORMALIZING CONSTANTS

Algorithm 1:

```

 $z = [0]_{n \times n}$ 
for  $i = 0 : \min\{h, n-y\}$ 
  for  $j = 0 : \min\{h-i, n-x\}$ 
     $z(x+j, y+i) = (i+j)^{-r}$ 
  for  $j = 1 : \min\{h-i, x-1\}$ 
     $z(x-j, y+i) = (i+j)^{-r}$ 
for  $i = 1 : \min\{h, y-1\}$ 
  for  $j = 0 : \min\{h-i, n-x\}$ 
     $z(x+j, y-i) = (i+j)^{-r}$ 
  for  $j = 1 : \min\{h-i, h-(m_1-i), x-1\}$ 
     $z(x-j, y-i) = (i+j)^{-r}$ 
 $z(x, y) = 0$ 
 $z = \sum_{i=1}^n \sum_{j=1}^n z(i, j)$ 
 $s(x, y) = \sum_{(i,j) \neq (x,y)} (|i-x| + |j-y|)^{-r} - z$ 

```

characterizing the capacity:

$$\begin{aligned}
M &= \max \left\{ \frac{h(h+3)}{2} + q, (1-\epsilon)c_w \right\} \\
\epsilon &= \sqrt{2(d+2) \ln(n^2)/c_w} \\
c_w &= \frac{h(h+3)}{2} + \sum_{x=1}^{h+1} \sum_{y=h+2-x}^n f(x,y) \\
&\quad + \sum_{x=h+2}^n \sum_{y=1}^n f(x,y) \\
f(x,y) &= q \cdot (g_{(x,y)}(1,1) + g_{(x,y)}(x,y)) \\
g_{(x,y)}(a,b) &= \left(1 - \frac{(x+y-2)^{-r}}{s(a,b)} \right)^{q-1} \cdot \frac{(x+y-2)^{-r}}{s(a,b)} \\
s(1,1) &= \sum_{i=h+1}^{n-1} (i+1) \cdot i^{-r} \\
&\quad + \sum_{i=0}^{n-2} (n-1-i) \cdot (n+i)^{-r}.
\end{aligned}$$

Recall that $s(x,y)$ can be calculated using *Algorithm 1*. The proof of the capacity will rely heavily on the following lemma.

Lemma 1: Let G_w be the weighted graph associated with Kleinberg's model of a small-world graph, and c_w be the global minimum cut in G_w . Then, for $h < n-1$, c_w is given by (1).

Proof: All the edges $e \in E_L$ have weight 1 (because they are never removed), all nodes in G_w have degree n^2-1 , and the weights of these edges depend only on the distance between the nodes they connect. Therefore, the global minimum cut in G_w must be a cut in which one of the partitions consists of one node. Because the weight of an edge in G_w decreases with the distance between the nodes that it connects, the global minimum cut in G_w must be a cut in which one of the partitions consists of a single node that maximizes the distance to other nodes. Therefore, c_w must be a cut in which one of the partitions consists of a corner node: $(1,1)$, $(1,n)$, $(n,1)$ or (n,n) .

Now, let $w((u_1, u_2), (v_1, v_2))$ be the weight of the edge connecting the nodes (u_1, u_2) and (v_1, v_2) . Assume, without loss of generalization, that c_w is the cut in which one of the partitions consists of node $(1,1)$. This way, $c_w = \sum_{(u_1, u_2) \neq (1,1)} w((1,1), (u_1, u_2))$. Now, we must count how many edges connecting node $(1,1)$ are in E_L , therefore, having weight 1. For this, we define an auxiliary way to numerate diagonals: $\{(1,1)\}$ is the diagonal 0, $\{(1,2), (2,1)\}$ is diagonal 1, and so on.

It is not difficult to see that the nodes in the i^{th} diagonal have a distance i to node $(1,1)$ ($i = 1, \dots, 2(n-1)$). Now, for $i \leq n-1$, there are $i+1$ nodes in the i^{th} diagonal and, for $i = n+j$ ($j = 0, \dots, n-2$), there are $n-1-j$ nodes in the i^{th} diagonal. Then, there are $\sum_{i=1}^h i+1 = h(h+3)/2$ nodes initially connected to node $(1,1)$ (again, with $h < n-1$), then there are $h(h+3)/2$ edges with weight 1. Therefore,

we have that:

$$\begin{aligned}
c_w &= \frac{h(h+3)}{2} + \sum_{x=1}^{h+1} \sum_{y=h+2-x}^n w((1,1), (x,y)) \\
&\quad + \sum_{x=h+2}^n \sum_{y=1}^n w((1,1), (x,y)).
\end{aligned}$$

We can calculate $s(1,1)$ as

$$(1) \quad s(1,1) = \sum_{i=h+1}^{n-1} (i+1) \cdot i^{-r} + \sum_{i=0}^{n-2} (n-1-i) \cdot (n+i)^{-r}.$$

Next, we determine the weights, $w((u_1, u_2), (v_1, v_2))$. Consider two nodes that are not connected initially, (u_1, u_2) and (v_1, v_2) , and the edge $((u_1, u_2), (v_1, v_2))$. This edge can be added in two different trials: one for node (u_1, u_2) and another one for node (v_1, v_2) . Because we do not consider multiple edges, these can be viewed as two mutually exclusive trials. Therefore, the weight of this edge is the sum of the probabilities of adding this edge when looking at node (u_1, u_2) and when looking at node (v_1, v_2) . Let us focus on node (u_1, u_2) . The trial "add edge $((u_1, u_2), (v_1, v_2))$ " follows a Binomial distribution, with q Bernoulli experiences, with success probability

$$\begin{aligned}
a_{(u_1, u_2)}(v_1, v_2) &= \frac{[d((u_1, u_2), (v_1, v_2))]^{-r}}{s(u_1, u_2)} \\
&= \frac{(|u_1 - v_1| + |u_2 - v_2|)^{-r}}{s(u_1, u_2)}.
\end{aligned}$$

Therefore, the weight of the edge $((u_1, u_2), (v_1, v_2))$ is

$$\begin{aligned}
w((u_1, u_2), (v_1, v_2)) &= q \cdot ((1 - a_{(u_1, u_2)}(v_1, v_2))^{q-1} \cdot a_{(u_1, u_2)}(v_1, v_2) \\
&\quad + (1 - a_{(v_1, v_2)}(u_1, u_2))^{q-1} \cdot a_{(v_1, v_2)}(u_1, u_2)).
\end{aligned}$$

Now, observing that $a_{(1,1)}(x,y) = \frac{x+y-2}{s(1,1)}$ and $a_{(x,y)}(1,1) = \frac{x+y-2}{s(x,y)}$, and using expression (1) for c_w , the result follows. ■

We are now ready to state our main result.

Theorem 3: For $h < n-1$ the capacity of a Kleinberg small-world network graph lies, with high probability, in the interval $[M, (1+\epsilon)c_w]$.

Proof: Using Lemma 1 and Corollary 1, we have that, with high probability,

$$c_s \in [(1-\epsilon)c_w, (1+\epsilon)c_w].$$

A tighter lower bound can be obtained for c_s as follows. Each node has a number of initial edges, determined by h , and q additional shortcut edges. The nodes with less initial edges are obviously the corner nodes, which exhibit $\frac{h(h+3)}{2}$ initial connections. Therefore, we have that

$$c_s \geq \frac{h(h+3)}{2} + q,$$

and the result follows. ■

The bounds of *Theorem 3* are illustrated in Fig. 5.




Fig. 5. Bounds for the capacity of a navigable small-world network for $n = 80$ (i.e. 1600 nodes), $h = 2$, $r = 2$ e $d = 1$, and different values of the shortcut parameter q . The white dots represent the expected value of the capacity and the dark dots represent the bounds computed according to Theorem 3.

V. CONCLUDING REMARKS

We studied the max-flow min-cut capacity of two classes of navigable small world networks. In both cases, we derived upper and lower bounds and illustrated their dependency on the parameters of the chosen topology.

In [13], Kleinberg explains that, in order to obtain a probability distribution, $d((u_1, u_2), (v_1, v_2))^{-r}$ should be divided by $\sum_{(v_1, v_2) \neq (u_1, u_2)} [d((u_1, u_2), (v_1, v_2))]^{-r}$. As we have shown, the previous expression is not an accurate normalizing constant, because the candidates for new connections from node (u_1, u_2) are *not all* the nodes of the base lattice, but *only* those nodes that are initially not connected to node (u_1, u_2) . For the goals of [13], this aspect is not very important, because in the cases of interest for that particular work (very small values of h relatively to the total number of nodes n^2) the difference between the two quantities is irrelevant in the construction of the model. However, as our work shows, using the correct normalizing factor is crucial towards bounding the capacity. The main reason is that this normalizing constant differs from node to node. In order to calculate the weights of the edges connecting a single node, we need to compute this normalizing constant for every node in the base lattice. Thus, the accumulation of errors affects the calculation of c_w , often leading to erroneous bounds.

Possible directions for future work include tighter capacity results, extensions to other classes of small-world networks (e.g. weighted models and other navigable topologies used in peer-to-peer networks [12]), and understanding if and how small-world topologies can be exploited in the design of capacity-attaining network codes and distributed network coding algorithms.

REFERENCES

- [1] Duncan J. Watts and Steven H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, no. 6684, June 4 1998.
- [2] S. Milgram, "Psychology today," *Physical Review Letters*, vol. 2, pp. 60–67, 1967.
- [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] S.N. Dorogovtsev and J.F.F. Mendes, *Evolution of Networks: From Biological Nets to the Internet and WWW*, Oxford University Press, 2003.
- [9] A. Helmy, "Small worlds in wireless networks," *IEEE Communications Letters*, vol. 7, no. 10, pp. 490–492, October 2003.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] Rui A. Costa and J. Barros, "On the capacity of small world networks," Submitted for publication, November 2005.
- [16] 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.
- [17] L.R. Ford and D.R. Fulkerson, *Flows in Networks*, Princeton University Press, Princeton, NJ, 1962.
- [18] 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/~barros>.
- [19] 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.
- [20] S. Borade, "Network Information Flow: Limits and Achievability," in *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, Lausanne, Switzerland, 2002.
- [21] 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.
- [22] R. Koetter and M. Médard, "An Algebraic Approach to Network Coding," *IEEE/ACM Trans. Networking*, vol. 11, no. 5, pp. 782–795, 2003.
- [23] 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.
- [24] P. Gupta and P. R. Kumar, "The Capacity of Wireless Networks," *IEEE Trans. Inform. Theory*, vol. 46, no. 2, pp. 388–404, 2000.
- [25] 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/>.
- [26] Steven H. Strogatz, "Exploring complex networks," *Nature*, vol. 410, pp. 268–276, 2001.
- [27] 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.