Biased Selection for Building Small-World Networks ⋆

Andrés SEVILLA[1], Alberto MOZO[2], M. Araceli LORENZO[2], Jose Luis LÓPEZ-PRESA[3], Pilar MANZANO[1], and Antonio FERNÁNDEZ ANTA[4]

[1] Dpto Informática Aplicada, U. Politécnica de Madrid, Madrid, Spain
asevilla@eui.upm.es, pmanzano@eui.upm.es
[2] Dpto Arquitectura y Tecnología de Computadores, U. Politécnica de Madrid, Madrid, Spain
amozo@eui.upm.es, alorenzo@eui.upm.es
[3] DIATEL, U. Politécnica de Madrid, Madrid, Spain
jllopez@diatel.upm.es
[4] Institute IMDEA Networks, Leganés, Spain
antonio.fernandez@imdea.org

Abstract. Small-world networks are currently present in many distributed applications and can be built augmenting a base network with long-range links using a probability distribution. Currently available distributed algorithms to select these long-range neighbors are designed ad hoc for specific probability distributions. In this paper we propose a new algorithm called Biased Selection (BS) that, using a uniform sampling service (that could be implemented with, for instance, a gossip-based protocol), allows to select long-range neighbors with *any* arbitrary distribution in a distributed way. This algorithm is of iterative nature and has a parameter r that gives its number of iterations. We prove that the obtained sampling distribution converges to the desired distribution as r grows. Additionally, we obtain analytical bounds on the maximum relative error for a given value of this parameter r. Although the BS algorithm is proposed in this paper as a tool to sample nodes in a network, it can be used in any context in which sampling with an arbitrary distribution is required, and only uniform sampling is available.
The BS algorithm has been used to choose long-range neighbors in complete and incomplete tori, in order to build Kleinberg's small-world networks. We observe that using a very small number of iterations (1) BS has similar error as a simulation of the Kleinberg's harmonic distribution and (2) the average number of hops with greedy routing is no larger with BS than in a Kleinberg network. Furthermore, we have observed that before converging to the performance of a Kleinberg network, the average number of hops with BS is significantly smaller (up to 14% smaller in a 1000×1000 network).

1 Introduction

Overlay networks are currently present in many distributed global applications and services. Overlay networks based on a small-world topology are an efficient and flexible alternative to structured overlays. Small-world networks can be built augmenting a base

⋆ This research was supported in part by Comunidad de Madrid grant S2009TIC-1692 and Spanish MICINN grant TIN2008–06735-C02-01.

network with long-range links, where the long-range neighbors are chosen using a probability distribution [10,7]. In order to obtain small-world networks by augmentation, a few distributed protocols have been proposed [5]. Among them, the simplest are epidemic protocols based on gossiping. However, these protocols only implement specific probability distributions to select long-range neighbors, i.e., each protocol is designed in an ad-hoc way for a given distribution. Gossip-based protocols have been designed for the uniform distribution and an approximation of Kleinberg's harmonic distribution [9,3,2]. In this paper we propose a local algorithm that, using a uniform sampling service (that could be implemented with, for instance, a gossip-based protocol), allows to select long-range neighbors with *any* arbitrary distribution. This algorithm is of iterative nature and we have found experimentally that in a small number of iterations converges to the desired distribution.

1.1 Related Work

Small world networks have been introduced in an attempt to explain the properties of social networks, and in particular the surprisingly small diameter and short routing in these networks [11]. One line of work on small-world networks has to do with synthesizing networks that have these properties. This can be done by starting from a base network (representing acquaintances geographically close) and adding long-range links (representing distant acquaintances). This process is called *augmentation* of the base network. Watts and Strogatz [14] considered a network augmentation in which the long-range neighbors are chosen uniformly at random. However, Kleinberg [10] has shown that a polylogarithmic greedy routing is achieved only if the long-range neighbors are chosen with specific distributions. A network built this way is commonly called a Kleinberg network. This seminal result has led to a large amount of subsequent work on construction of small-world networks [1,5,6,7]. Most of the algorithms that have been proposed are centralized. As far as we know, the first distributed algorithm to build a small world network is due to Duchon et al. [5]. Bonnet et al. [3] have proposed two gossip-based protocols to select long-range links, which are modified versions of Cyclon [13]. One selects the long-range neighbors with uniform probability, while the other selects them with an approximation to the Kleinberg distribution. This latter protocol has been improved in [2]. There have been other gossip-based protocols that sample the network nodes with uniform probability [4,9,13].

1.2 Contributions

In this paper we provide an algorithm, called *Biased Selection* (BS), that implements a sampling service in a set S with any probability distribution. The probability distribution is proportional to probabilities (represented as weights) assigned to the element in S. The algorithm BS is very simple and completely local. It only needs access to a uniform sampling service (that could be implemented with a gossip-based protocol) and to the weight assigned to each element returned by this service. This algorithm has a parameter r that determines the number of times the uniform sampling service is used (number of rounds) before returning a sample. (In fact, the times the uniform sampling

service is used is exactly $r + 1$.) We prove that the obtained sampling distribution converges to the desired distribution as r grows. Additionally, we obtain analytical bounds on the maximum relative error for a given value of this parameter r. Although the BS algorithm is proposed in this paper as a tool to sample nodes in a network, it can be used in any context in which sampling with an arbitrary distribution is required, and only uniform sampling is available.

To evaluate by simulation the performance of the algorithm, we use it to choose long range neighbors in a torus, in order to build small-world networks similar to Kleinberg's [10]. In this network, every node i in the torus chooses another node j as its long-range neighbor with a probability proportional to $1/d(i,j)^2$, where $d(i,j)$ is the Euclidean distance[5] from i to j. For simplicity, we will call this the *Kleinberg distribution*. The obvious way to choose the long range neighbor of a node i implies to know all the nodes, the distance from i to each of them, and to compute the associated probabilities, which requires $\Omega(n)$ operations in a network of n nodes. To do this for all nodes requires $\Omega(n^2)$ operations. On the other hand, BS does not need to know all the nodes, and only requires the distance from i to the nodes returned by the uniform sampling service. If BS uses r rounds, the overall number of operations needed for each node i is $O(r)$, and $O(rn)$ operations for the whole network. We have observed experimentally that the value of r required is much smaller than n.

Comparing the samples obtained simulating BS and samples from a simulation of the Kleinberg distribution, we observe that in a very small number of rounds (10 in a 100×100 torus), both simulations have matching average relative error and very similar maximum relative error. Then, we build networks by adding to each node in the torus a long range neighbor. We evaluate the performance of greedy routing in networks in which the long range neighbor is chosen with the BS algorithm (BS-network) when different number of rounds is used. We compare these results with networks that use the Kleinberg distribution (K-network) and the uniform distribution (U-network). We observe that the average number of hops of greedy routes in the BS network converges to the values obtained in the K-network as the number of rounds increases. Furthermore, with only a few rounds, the average number of hops in the BS network is sensibly smaller than in the U-network (especially for large networks). Surprisingly, we have observed that before converging to the K-network performance, the average number of hops of the BS-network is significantly smaller (up to 14% smaller in a 1000×1000 network). In fact, the best performance of BS-networks is achieved with a small number of rounds. The origin of this behavior is left for future study. Finally, we have done similar experiments adding long range neighbors to incomplete tori. These are obtained by deciding whether to remove each node with a fixed probability. The experiments have been done with two probability values, namely 0.8 and 0.3, obtaining dense and sparse networks. The results observed are consistent with those obtained in the complete torus.

[5] Observe that Kleinberg used, instead, Manhattan distance on a grid.

1.3 Structure of the Rest of the Paper

In Section 2 we introduce concepts and notation that will be used in the rest of the paper, along with the description of the experimental environment that will be used. In Section 3 the BS algorithm is presented, its correctness is proven, and a bound on its convergence rate is derived. Finally, in Section 4 simulation and experimental results are presented.

2 Definitions and Experimental Setup

2.1 Definitions

Although the BS algorithm is proposed in this paper as a tool to sample nodes in a network, it will be presented in a more general form, to emphasize the fact that it could possibly be used in other contexts. Let S be a set of n elements such that each element $i \in S$ has an associated weight $w(i) > 0$. The problem to solve is to sample the set S with a probability distribution p such that the probability of choosing i is proportional to $w(i)$. Let us denote $\eta = \sum_{j \in S} w(j)$. Then, the sampling probability of $i \in S$ has to be $p(i) = w(i)/\eta$. The challenges of sampling S are the following:

1. We assume that the whole set is not available.
2. The weight values can only be consulted for individual previously known elements.

These restrictions prevent, for instance, from even computing the value η. However, in order to be able to solve the problem, we assume the availabilty of a sampling primitive $USel_S$ that returns an element i of S chosen with uniform probability. Once an element $i \in S$ is obtained using $USel_S$, its weight $w(i)$ can also be obtained.

2.2 Experimental Setup

In the experiments conducted in this paper, we consider a 2-dimensional torus topology. A pair of integer values (x, y) is used to locate each node into the 2-dimensional space. The former node coordinates range from 0 to $m - 1$, and so the number of nodes in the network is m^2. In this topology, the distance between two nodes located at positions (x_1, y_1) and (x_2, y_2) is the Euclidean distance in the torus, computed as:

$$d_e = \sqrt{\left(\min\left(|x_1 - x_2|, m - |x_1 - x_2|\right)\right)^2 + \left(\min\left(|y_1 - y_2|, m - |y_1 - y_2|\right)\right)^2}$$

To test the BS algorithm, we design two different types of experiments. The first experiment (Section 4.1) shows the average and maximum relative error values of the BS algorithm with respect to the Kleinberg probability distribution. These values are also compared with the relative error values obtained with a real simulation of the ideal Kleinberg distribution. The second group of experiments (Sections 4.2 and 4.3) compares BS with the Kleinberg simulator with respect to the average number of hops when these algorithms are used to greedily route packets in a network. This group of experiments is executed in two different scenarios. Firstly, a complete torus with m^2 nodes

is used. In this torus, each node has four local neighbors and one long range neighbor. Secondly, we use an incomplete torus, where nodes are eliminated using a random uniform probability. In this network, the expected number of present nodes is m^2q, being q the probability of node presence. Note that q is equal for every node, and the presence of each node is independent of the presence of other nodes. In this case, to allow for greedy routing, each node i has links to seven neighbors. These include one long range neighbor, and six local neighbors, that are the closest node in each of the six 60° wedge in a circle centered on node i [15].

Additionally, when using the BS algorithm, the number of rounds r will determine the accuracy of the values obtained. Hence, experiments are executed several times. We start with $r = 0$ rounds (uniform distribution) and we gradually increase r until the BS algorithm converges to the Kleinberg distribution. Each experiment is repeated 10 times with different seeds, and we present the average of these executions.

3 Biased Selection

3.1 The Biased Selection Algorithm

We present here the algorithm that can be used to sample the set S as defined in Section 2 with the desired probability distribution. The algorithm is called *Biased Selection* and presented in Figure 1. The input of the algorithm is a value r that specifies the number of rounds the algorithm must execute before returning the sample. As will be shown, the larger the number of rounds r, the closer the output of $BSel_S(r)$ gets to the desired probability distribution.

```
1  function BSel_S(r)
2  x ← USel_S
3  for i ← 1 to r do
4      y ← USel_S
5      set x ← y with probability w(y)/(w(x)+w(y))
6  end for
7  return x
```

Fig. 1. Biased Selection Algorithm for set S.

3.2 Correctness

We first show that, as r goes to infinity, the probability distribution of the output values of $BSel_S(r)$ converges to the desired probability distribution p. Let x_i be the value stored in variable x after i iterations of the for loop, being x_0 the value assigned to x in Line 2. Let us consider the infinite run of $BSel_S(\infty)$. The infinite sequence of values

$x_0, x_1, ...$ can be seen as a Markov chain[6] M on the finite state space S. We will first show that M has a unique stationary distribution π such that $\pi = A\pi$, where A is the transition matrix of M. Finally, we show that π matches the probability distribution p.

The transition matrix $A = [a_{ij}]$ of the Markov chain M can be obtained from the algorithm of Figure 1 in the following way. For each $i, j \in S$ and $i \neq j$,

$$a_{ij} = \Pr[x_{t+1} = j | x_t = i] = \frac{1}{n} \frac{w(j)}{w(j) + w(i)} = \frac{1}{n} \frac{p(j)}{p(j) + p(i)}.$$

Additionally, $a_{ii} = 1 - \sum_{j \neq i} a_{ij}$. Observe that for all $j \neq i$, $a_{ij} < 1/n$ (recall that $n = |S|$), and hence $a_{ii} > 0$. We show now that M is *ergodic*. Let $a_{ij}^{(s)}$ be the probability of reaching state j in s steps starting from state i. A Markov chain is ergodic if it is finite, *irreducible* ($\forall i, j \in S, \exists s : a_{ij}^{(s)} > 0$), and *aperiodic* ($\forall i, j \in S, \gcd\{s : a_{ij}^{(s)} > 0\} = 1$) [12]. Ergodicity implies that the stationary distribution is unique.

Lemma 1. *The Markov chain M is ergodic, and hence has a unique stationary distribution π that satisfies $\pi = \pi A$.*

Proof. The fact that the Markov chain M is irreducible follows directly from the fact that $a_{ij} > 0$ for all $i, j \in S$. Additionally, since $a_{ii} > 0$ for all $i \in S$, M is aperiodic. Hence, it is ergodic, and has a unique stationary distribution π that satisfies $\pi = \pi A$ [12].

Let us now prove that the stationary distribution is in fact the desired distribution p.

Theorem 1. *The output of $BSel_S(r)$ converges to the probability distribution p as r tends to infinity.*

Proof. From the above lemma, the distribution of values x_r output by the algorithm converge to the stationary distribution π of Markov chain M as r goes to infinity. All that has to be proven is that $p = \pi$, i.e., that $p = pA$. Consider any $i \in S$, we need to prove that $p(i) = \sum_{j \in S} p(j) a_{ji}$. Replacing,

$$\begin{aligned}
\frac{w(i)}{\eta} &= \sum_{j \in S} \frac{w(j)}{\eta} a_{ji} \\
&= \sum_{j \neq i} \frac{w(j)}{\eta} \frac{1}{n} \frac{w(i)}{w(j) + w(i)} + \frac{w(i)}{\eta} (1 - \sum_{j \neq i} a_{ij}) \\
&= \sum_{j \neq i} \frac{w(i)}{\eta} \frac{1}{n} \frac{w(j)}{w(j) + w(i)} + \frac{w(i)}{\eta} (1 - \sum_{j \neq i} a_{ij}) \\
&= \frac{w(i)}{\eta} \left(\sum_{j \neq i} a_{ij} + 1 - \sum_{j \neq i} a_{ij} \right) \\
&= \frac{w(i)}{\eta}
\end{aligned}$$

Since this holds for all $i \in S$, the proof is complete.

[6] Some familiarity of the reader with Markov chains is assumed.

3.3 Convergence Rate

We study now the number of rounds r that are needed for the distribution of the output values of $BSel_S(r)$ to be *almost* the same as the probability distribution p. To measure the distance between both distributions, we will use the *relative pointwise distance* as defined in [12]. This parameter measures the largest relative error between the distributions, for all possible final and initial values. Observe that $a_{ij}^{(r)}$ is the probability that $BSel_S(r)$ outputs j if the initial value of x is i, i.e. $a_{ij}^{(r)} = \Pr[x_r = j | x_0 = i]$. Then, the *maximum relative error* is defined as

$$\Delta(r) = \max_{i,j \in S} \frac{|a_{ij}^{(r)} - p(j)|}{p(j)}.$$

In order to bound $\Delta(r)$ we first prove that the Markov chain M is *time-reversible*, which holds if $a_{ij}p(i) = a_{ji}p(j)$ [12].

Lemma 2. *The Markov chain M is time-reversible.*

Proof. Replacing in $a_{ij}p(i) = a_{ji}p(j)$, we get

$$a_{ij}p(i) = \frac{1}{n}\frac{w(j)}{w(j)+w(i)}\frac{w(i)}{\eta} = \frac{1}{n}\frac{w(i)}{w(j)+w(i)}\frac{w(j)}{\eta} = a_{ji}p(j)$$

Lemmas 1 and 2 are useful to bound $\Delta(r)$ because of the following result, derived from Proposition 3.1 in [12].

Lemma 3 ([12]). *Let A be the transition matrix of an ergodic time-reversible Markov chain, p its stationary distribution, and $1 = \lambda_0 > \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{n-1}$ its (real) eigenvalues. Then, for all $r \geq 1$ the maximum relative error satisfies*

$$\Delta(r) \leq \frac{\lambda^r}{p_{\min}},$$

where $\lambda = \max_{k \geq 1} |\lambda_k|$ and $p_{\min} = \min_{i \in S} p(i)$.

Clearly, $\lambda = \max(\lambda_1, |\lambda_{n-1}|)$, where $\lambda_1 < |\lambda_{n-1}|$ only if $\lambda_{n-1} < 0$. Let us define $a_{\min} = \min_{i \in S} a_{ii}$. From the Gershgorin Circle Theorem [8], we have that $\lambda_{n-1} \geq 2a_{\min} - 1$. Then, $\lambda \leq \max\{\lambda_1, 1 - 2a_{\min}\}$. To bound λ_1 we use the conductance of M.

Definition 1. *Consider the Markov chain M. For any set $B \subset S$, denote $C(B) = \sum_{i \in B} p(i)$ and $F(B) = \sum_{i \in B, j \notin B} a_{ij}p(i)$. The* conductance *of M is defined as*

$$\Phi = \min_{\emptyset \subset B \subset S : C(B) \leq 1/2} \frac{F(B)}{C(B)}.$$

Lemma 3.3 in [12] shows that $\lambda_1 \leq 1 - \frac{\Phi^2}{2}$. Then, we can bound λ_1 as follows.

Lemma 4. *The eigenvalue λ_1 of the ergodic time-reversible Markov chain M satisfies*

$$\lambda_1 \leq 1 - \frac{1}{2}\left(\max\{(1-1/n)p_{\min}, 1/(4np_{\max})\}\right)^2,$$

where $p_{\min} = \min_{i\in S} p(i)$ and $p_{\max} = \max_{i\in S} p(i)$.

Proof. Consider any set B such that $\emptyset \subset B \subset S$ and $C(B) \leq 1/2$. Let us denote $\Phi(B) = \frac{F(B)}{C(B)}$. We will obtain two lower bounds for $\Phi(B)$.

First, since $C(B) \leq 1/2$, then $\Phi(B) \geq 2F(B)$. In order to bound $F(B)$, we observe that $\frac{p(i)p(j)}{p(j)+p(i)}$ decreases with the values of $p(i)$ and $p(j)$, which implies that $\frac{p(i)p(j)}{p(j)+p(i)} \geq \frac{p_{\min}^2}{2p_{\min}} = \frac{p_{\min}}{2}$. On the other hand, $|B| \cdot |S \setminus B| \geq n-1$. Then,

$$F(B) = \sum_{i\in B, j\notin B} a_{ij}p(i) = \frac{1}{n}\sum_{i\in B, j\notin B}\frac{p(i)p(j)}{p(j)+p(i)} \geq \frac{1}{n}\sum_{i\in B, j\notin B}\frac{p_{\min}}{2} \geq \frac{n-1}{n}\frac{p_{\min}}{2}.$$

This implies that $\Phi(B) \geq (1-1/n)p_{\min}$. For the second bound, observe that

$$F(B) = \frac{1}{n}\sum_{i\in B, j\notin B}\frac{p(i)p(j)}{p(j)+p(i)} \geq \frac{1}{2np_{\max}}\sum_{i\in B, j\notin B} p(i)p(j) = \frac{\left(\sum_{i\in B} p(i)\right)\left(\sum_{j\notin B} p(j)\right)}{2np_{\max}}.$$

Then, since $C(B) = \sum_{i\in B} p(i)$,

$$\Phi(B) = \frac{F(B)}{C(B)} \geq \frac{\sum_{j\notin B} p(j)}{2np_{\max}} \geq \frac{1}{4np_{\max}},$$

where the second inequality follows from $\sum_{j\notin B} p(j) = 1 - C(B)$ and $C(B) \leq 1/2$.

Since both bounds hold for any B, we obtain that $\Phi \geq \max\{(1-1/n)p_{\min}, 1/(4np_{\max})\}$. Combining this bound with Lemma 3.3 in [12], the claim follows.

From the above results we can bound the relative pointwise distance as follows.

Theorem 2. *The maximum relative error of the Markov chain M satisfies*

$$\Delta(r) \leq \left(1 - \min\{\frac{1}{2}\left(\max\{(1-1/n)p_{\min}, 1/(4np_{\max})\}\right)^2, 2a_{\min}\}\right)^r / p_{\min},$$

where $p_{\min} = \min_{i\in S} p(i)$, $p_{\max} = \max_{i\in S} p(i)$, and $a_{\min} = \min_{i\in S} a_{ii}$.

4 Experimental Results

4.1 Accuracy of Biased Selection

To be able to evaluate the goodness of BS, we compare its relative error with the one measured in a simulation of the Kleinberg distribution. We define the relative error e_i for node i in a collection C of s samples as $e_i = \frac{|fsim_i - fkl_i|}{fkl_i}$, where $fsim_i$ is the

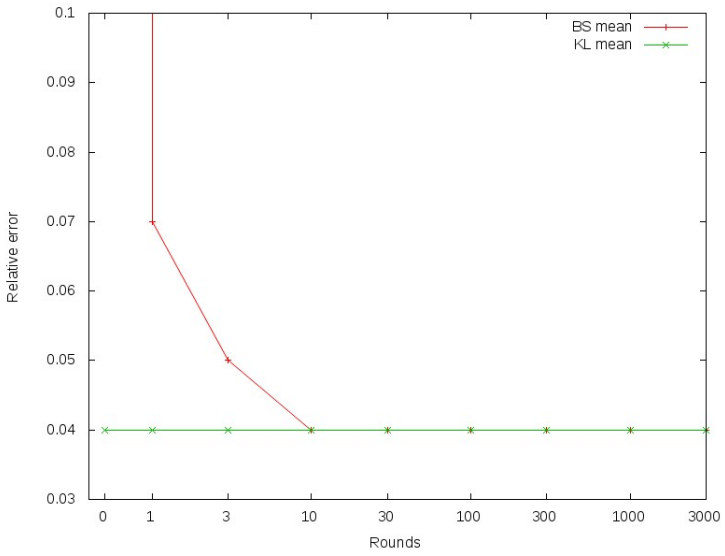

Fig. 2. Average relative error distribution of the Kleinberg simulator and BS (100×100 torus).

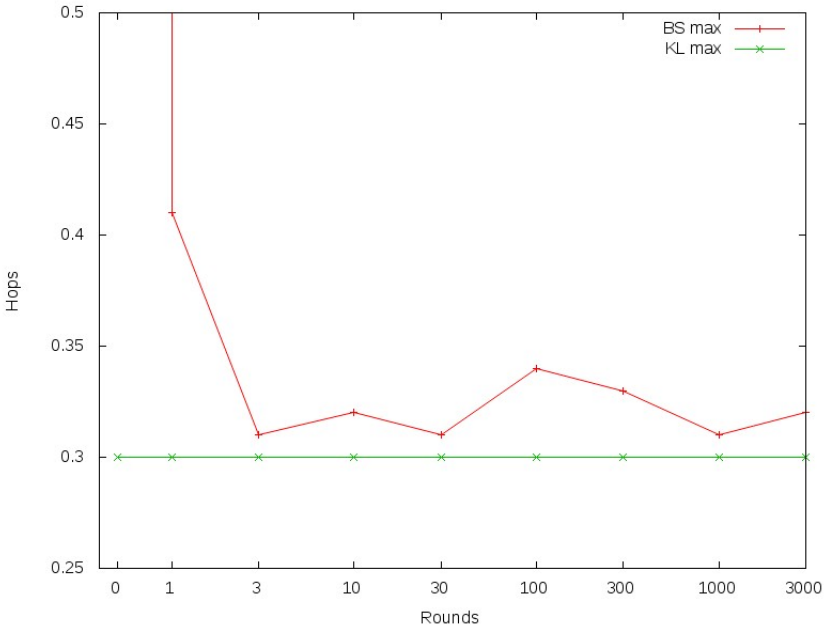

Fig. 3. Maximum relative error distribution of the Kleinberg simulator and BS (100×100 torus).

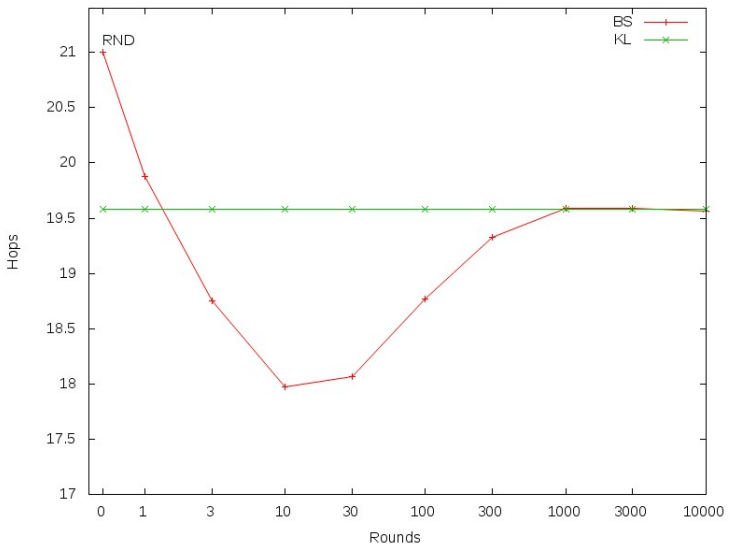

Fig. 4. Comparative of average routing hops (100×100 torus).

number of instances of i in collection C, and $fkl_i = p(i) \cdot s$ is the expected number of instances of i with the ideal Kleinberg distribution. The experiments have been done in a 100×100 torus. For each experiment, a collection of around 13.5 million samples has been used in order to guarantee that every node appears on average at least 100 times. Additionally, in the BS algorithm, experiments have been performed using an increasing number of rounds, trying to reach a behavior similar to that of the Kleinberg simulator. In the 10,000-node torus used this happens for relatively small number of rounds ($r = 10$). We have not performed this experiment on larger networks (more than 100×100 nodes), due to limitations in the execution time needed to handle experiments above that size.

In Figures 2 and 3 we show, respectively, the average and maximum relative error values obtained using the BS algorithm against a Kleinberg simulator. Round numbers approximately follow an exponential sequence of $r = 0$, 1, 3, 10 , 30, 100, 300, 1000, and 3000. It must be noted that for $r = 0$, our BS algorithm is equivalent to a uniform random distribution. We can observe that, for $r \geq 10$ the mean values of the relative errors in the BS algorithm match those obtained using the Kleinberg simulator.

4.2 Building Small-World Networks with Biased Selection

As previously commented, in this scenario we built a complete torus with m^2 nodes. In the experiments we perform 500,000 search operations, choosing source and destination uniformly at random for each search. We measure the average number of hops needed to reach the destination with greedy routing. Let h_K be the average number of

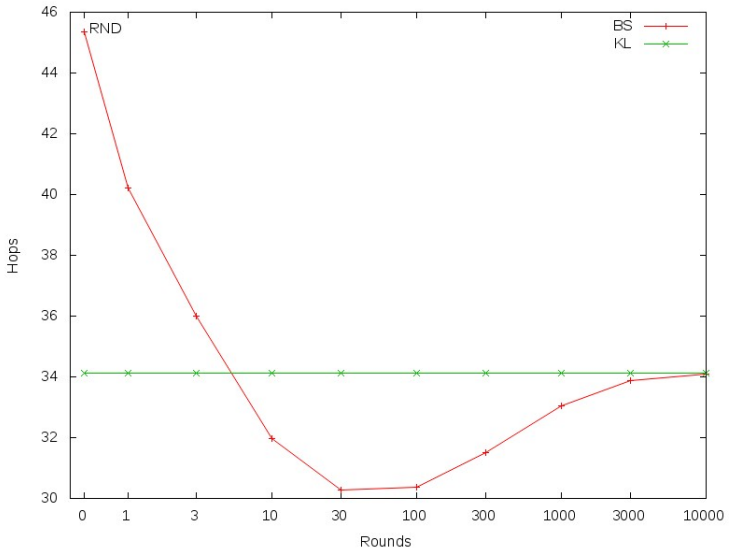

Fig. 5. Comparative of average routing hops (300 $\times$ 300 torus).

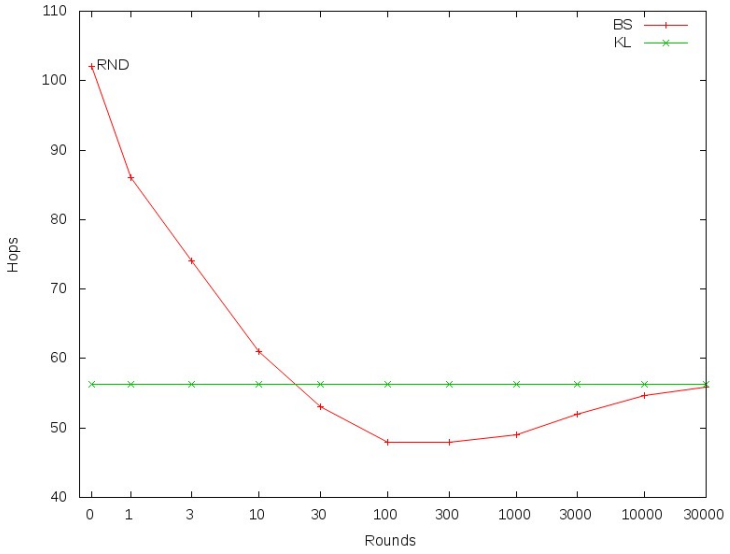

Fig. 6. Comparative of average routing hops (1000 $\times$ 1000 torus).

hops measured with the Kleinberg simulator, and $h_{BS}(r)$ the average number of hops measured when using BS with r rounds. We consider that BS and Kleinberg have converged for r rounds when, for all $r' \geq r$, it holds that $|h_{BS}(r') - h_K|/h_K < 0.05$ (they are off by less than 5%). First, we run the Kleinberg simulator, and then we run the BS algorithm, starting from $r = 0$ and gradually increasing the number of rounds until convergence is reached. We use round numbers that approximately follow an exponential sequence, $r = 0$, 1, 3, 10, 30, 100, 300, 1000, 3000, 10,000, and 30,000. We denote the smallest of these values of r that satisfy convergence as r_{conv}. It must be noted that using 0 rounds in BS is equivalent to using a uniform distribution to choose the long range neighbor. The experiment was run using three different torus sizes, 100×100 (Figure 4), 300×300 (Figure 5) and 1000×1000 (Figure 6).

From these experiments, we can conclude that:

- As analytical results showed in section 3.3, the execution of the BS algorithm converges to the Kleinberg simulator results when using a sufficient number of rounds. The approximate number of rounds needed is 100 in the 100×100 torus, 1,000 in the 300×300 torus, and 10,000 in the 1000×1000 torus. Note that the number of rounds needed grows when the torus size increases.
- Using a uniform distribution to choose the long range link (equivalent to BS with 0 rounds) produces worse results than the Kleinberg and BS simulators with $r \geq 1$. The difference increases as the network size grows.
- With a relatively small number of rounds, BS outperforms the Kleinberg simulator. This singular behavior appears in all the experiment executions. In a 100×100 torus, with 10 rounds, BS results are 8% better than the Kleinberg results and 10% better than the uniform distribution. In a 300×300 torus, with 30 rounds, we get improvements of 12% and 25% when comparing to the Kleinberg simulation and the uniform distribution, respectively. Finally, in a 1000×1000 torus, with 100 rounds, improvements are of 14% and 45% when comparing to the Kleinberg simulation and the uniform distribution, respectively. It can be seen that the number of rounds required by BS to obtain the minimum average number of hops grows as the network size increases. We denote this number of rounds as r_{min}.

In Table 1, we present r_{min} and r_{conv} as a function of the size of the network.

4.3 Building Incomplete Small-World Networks with Biased Selection

In this scenario we built an incomplete torus using two different probability values (0.8 and 0.3) to determine node presence in 100×100 and 300×300 topologies. As before, in this experiment, we perform 500,000 search operations, using alternatively the Kleinberg and the BS simulators. The number of rounds in the BS algorithm started with 0 and continued until reaching convergence with the Kleinberg simulator.

Figure 7 (100×100 torus) and Figure 8 (300×300 torus) compare the performance of the BS algorithm with that of the Kleinberg simulator in a topology with a probability value $q = 0.8$ of node presence in the network. Additionally, Figure 9 (100×100 torus) and Figure 10 (300×300 torus) compare the former simulators using a topology with a probability value $q = 0.3$ for node presence in the network. Finally, in Table 1, we show r_{min} and r_{conv} for each network considered.

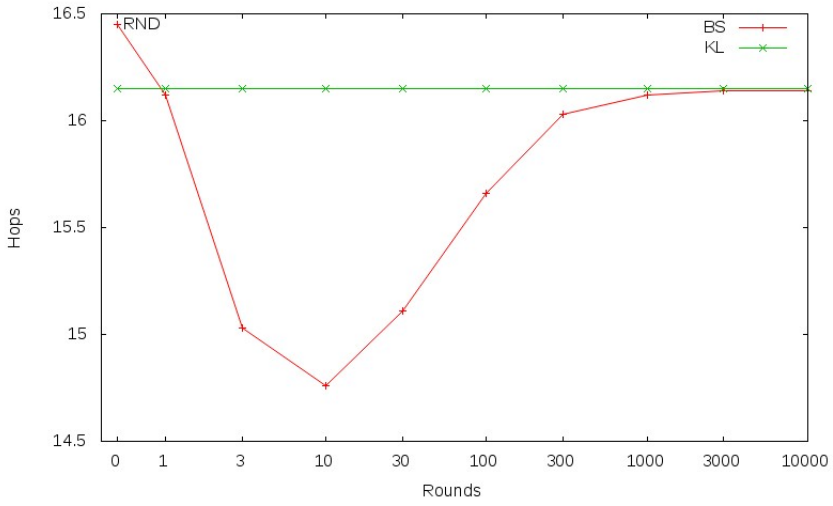

Fig. 7. Comparative of average routing hops (100×100 Incomplete torus $q = 0.8$)

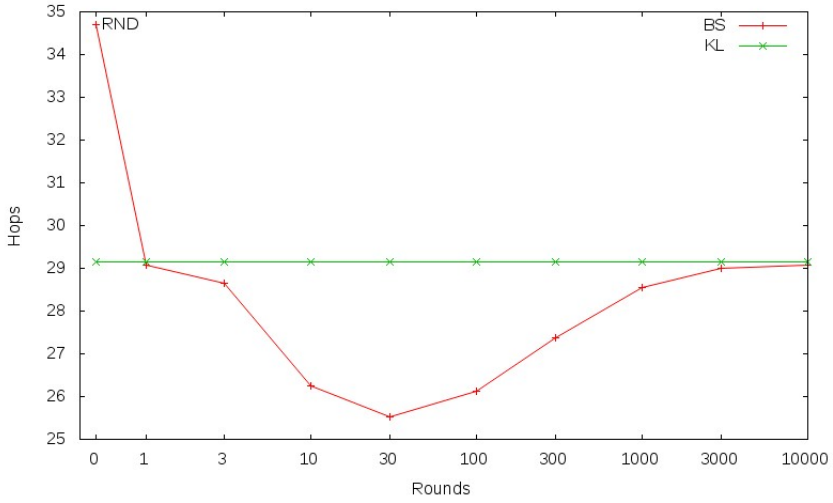

Fig. 8. Comparative of average routing hops (300×300 Incomplete torus $q = 0.8$)

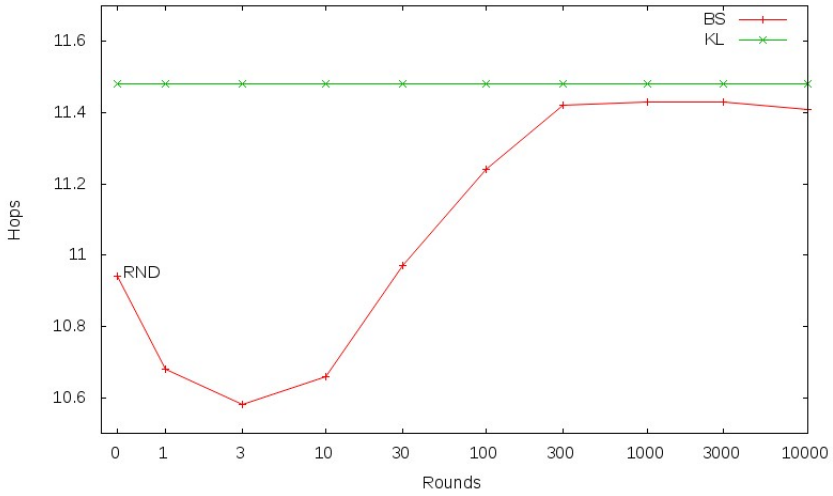

Fig. 9. Comparative of average routing hops (100×100 Incomplete torus $q = 0.3$)

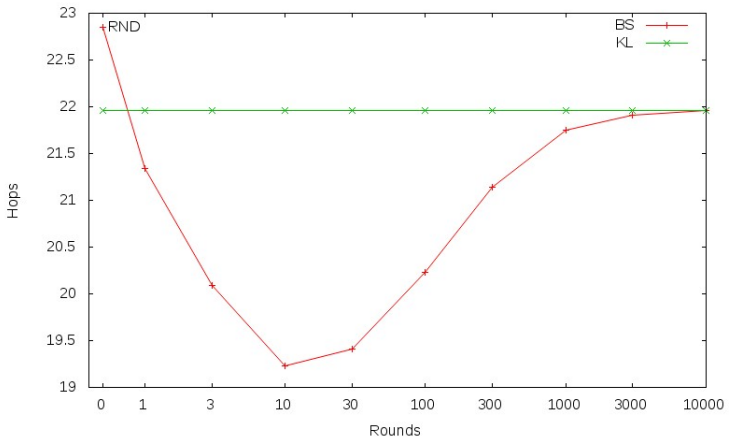

Fig. 10. Comparative of average routing hops (300×300 Incomplete torus $q = 0.3$)

Topology	Torus size	Nodes (avg)	r_{min}	r_{conv}
Complete	100×100	$10,000$	10	100
Complete	300×300	$90,000$	30	$1,000$
Complete	1000×1000	$1,000,000$	300	$10,000$
Incomplete ($q = 0.8$)	100×100	$8,000$	10	100
Incomplete ($q = 0.8$)	300×300	$72,000$	30	$1,000$
Incomplete ($q = 0.3$)	100×100	$3,000$	3	30
Incomplete ($q = 0.3$)	300×300	$27,000$	10	300

Table 1. Number of BS rounds for minimal number of routing hops r_{min} and for convergence r_{conv} in the different experiments.

From these experiments, we can conclude that:

- The results are similar to those obtained in the previous section when using a complete torus. The values of r_{min} and r_{conv} for networks with roughly the same number of nodes are similar. These results are especially interesting because an incomplete torus with $q = 0.3$ is almost a random network, so it seems that the results obtained are not associated only to a torus topology. In an incomplete network with $q = 0.8$ the obtained improvements are: (a) in the 100×100 torus, with 10 rounds, BS results are 8% better than the Kleinberg results and 11% better than the uniform distribution; and (b) in the 300×300 torus, with 30 rounds, BS results are 12% better than the Kleinberg results and 26% better than the uniform distribution. Using $q = 0.3$: (a) in the 100×100 torus, with 3 rounds, BS results are 8% better than the Kleinberg results and 4% better than the uniform distribution; and (b) in the 300×300 torus, with 10 rounds, BS results are 12% better than the Kleinberg results and 15% better than the uniform distribution.
- In these experiments, the average number of routing hops and the number of rounds needed to converge seem to be somewhat smaller than the values obtained in the complete torus experiments. We presume that this result may be due to two factors: (a) the number of local neighbors is greater than in the previous experiment (6 neighbors versus 4); and (b) each routing hop in this network generates a larger advance than in a complete network.

5 Conclusions and Future Work

In this paper we proposed a simple, iterative and local algorithm (BS) that allows us to select long-range neighbors with any arbitrary distribution to build small-world networks. BS uses a uniform sampling service and only needs one parameter determining the number of rounds needed by the algorithm to converge to the desired distribution. In this work, we use the Kleinberg distribution as the target. We also proved the algorithm convergence and obtained analytical bounds on the maximum relative error for a given value of the algorithm parameter. We evaluated the algorithm by simulation in different scenarios, obtaining convergence with the Kleinberg simulator results. We

also observed that, before converging to the Kleinberg distribution, the BS algorithm provided a smaller average number of hops, up to 14 % smaller in a 1000×1000 network.

Future work will provide an analytical description of this behavior and evaluate the algorithm with topologies different from the torus. We will also study the algorithm in a dynamic network scenario. Finally, we would like to compare the properties of BS with aggregation protocols.

Acknowledgement

The authors would like to thank Juan A. Carrasco for useful discussions on Markov chains and for referring them to the Gershgorin Circle Theorem.

References

1. L. Barrière, P. Fraigniaud, E. Kranakis, and D. Krizanc. Efficient routing in networks with long range contacts. In J. L. Welch, editor, *Proc. DISC*, volume 2180 of *Lecture Notes in Computer Science*, pages 270–284. Springer, 2001.
2. Marin Bertier, François Bonnet, Anne-Marie Kermarrec, Vincent Leroy, Sathya Peri, and Michel Raynal. D2ht: The best of both worlds, integrating rps and dht. In *EDCC*, pages 135–144. IEEE Computer Society, 2010.
3. F. Bonnet, A.-M. Kermarrec, and M. Raynal. Small-world networks: From theoretical bounds to practical systems. In *Proc. OPODIS*, pages 372–385, 2007.
4. F. Bonnet, F. Tronel, and S. Voulgaris. Brief announcement: Performance analysis of cyclon, an inexpensive membership management for unstructured P2P overlays. In S. Dolev, editor, *Proc. DISC*, volume 4167 of *Lecture Notes in Computer Science*, pages 560–562. Springer, 2006.
5. Philippe Duchon, Nicolas Hanusse, Emmanuelle Lebhar, and Nicolas Schabanel. Towards small world emergence. In Phillip B. Gibbons and Uzi Vishkin, editors, *SPAA*, pages 225–232. ACM, 2006.
6. P. Fraigniaud, C. Gavoille, and C. Paul. Eclecticism shrinks even small worlds. *Distributed Computing*, 18(4):279–291, 2006.
7. Pierre Fraigniaud and George Giakkoupis. On the searchability of small-world networks with arbitrary underlying structure. In Leonard J. Schulman, editor, *STOC*, pages 389–398. ACM, 2010.
8. S. Gerschgorin. Über die abgrenzung der eigenwerte einer matrix. *Izv. Akad. Nauk. USSR Otd. Fiz.-Mat. Nauk*, 7:749–754, 1931.
9. M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, and M. van Steen. Gossip-based peer sampling. *ACM Trans. Comput. Syst.*, 25(3), 2007.
10. J. M. Kleinberg. Navigation in a small world. *Nature*, 406(6798), August 2000.
11. Stanley Milgram. The small world problem. *Psychology Today*, 2:60–67, 1967.
12. Alistair Sinclair and Mark Jerrum. Approximate counting, uniform generation and rapidly mixing markov chains. *Inf. Comput.*, 82(1):93–133, 1989.
13. S. Voulgaris, D. Gavidia, and M. van Steen. Cyclon: Inexpensive membership management for unstructured P2P overlays. *J. Network Syst. Manage.*, 13(2), 2005.
14. D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.
15. A.C.C. Yao. On constructing minimum spanning trees in k-dimensional space and related problems. *SIAM Journal of Computing*, 1982.