How Much Independent Should Individual Contacts be to Form a Small–World? ⋆
(Extended Abstract)

Gennaro Cordasco and Luisa Gargano

Dipartimento di Informatica ed Applicazioni, Università di Salerno, Italy.
E-mail: {cordasco,lg}@dia.unisa.it

Abstract. We study Small–World graphs in the perspective of their use in the development of efficient as well as easy to implement network infrastructures. Our analysis starts from the Small–World model proposed by Kleinberg: a grid network augmented with directed long–range random links. The choices of the long–range links are independent from one node to another. In this setting greedy routing and some of its variants

have been analyzed and shown to produce paths of polylogarithmic expected length. We start from asking whether all the independence assumed in the Kleinberg's model among long–range contacts of different nodes is indeed necessary to assure the existence of short paths. In order to deal with the above question, we impose (stringent) restrictions on the choice of long–range links and

we show that such restrictions do not increase the average path length of greedy routing and of its variations. Diminishing the randomness in the choice of random links has several benefits; in particular, it implies an increase in the clustering of the graph, thus increasing the resilience of the network.

1 Introduction

In this paper we consider Small-World (SW) networks based on Kleinberg's model [5]. We investigate the possibility of diminishing the amount of randomness that nodes need in the choice of their long–range links, while keeping the short routes of the original model. Our proposal has the advantage that the obtained networks show high clustering and, hence, they are particularly resilient to failures.

1.1 Small–World (SW) Networks

The study of many large–scale real world networks shows that such networks exhibit a set of properties that cannot be totally captured by the traditional models: regular graphs and random graphs. Indeed, many biological and social networks occupy a position which is intermediate between completely regular and random graphs. Such networks, commonly called *Small–World* networks, are characterized by the following main properties:

- they tend to be sparse;
- they tend to have short paths (as random graphs);
- they tend to be clustered (unlike sparse random graphs).

⋆ This work was partially supported by the Italian FIRB project "WEB-MINDS", http://web-minds.consorzio-cini.it/.

The study of SW graph was pioneered by Milgram [10] in the 1960s. In his famous experiment, people were asked to send letters to unknown targets only through acquaintances. The result confirmed the believe that random pairs of individuals are connected by short chains of acquaintances.

The work by Milgram has been followed by several studies finalized to a better understanding and modeling of the SW phenomenon. In particular, Watts and Strogatz [16] noticed that, unlikely random graphs, real data networks tend to be clustered. Thus, Watts and Strogatz proposed thinking about SW networks as combining an underlying regular (high–diameter) graph with a few random links. They showed that several real networks fall into this category, e.g. the network of actors cooperations or the neural wiring diagram of the worm *C. elegans*.

Recently, Kleinberg [5] reconsidered an important algorithmic aspect of Milgram's experiment: not only short path exist but individuals are able to deliver messages to unknown targets using short routes. Kleinberg proposed a basic model that uses a two-dimensional grid as underlying interconnection, the grid is then augmented with random links: Each node has an undirected local link to each of its grid neighbors and one directed long-range random link; this last link connects to a node at lattice distance d with probability proportional to d^{-2}. Kleinberg proved that such graphs are indeed *navigable,* that is, a simple *greedy* algorithm finds routes between any source and target using only $O(\log^2 n)$ expected hops in a network on n nodes [5]. Notice that navigability is an interesting property for a graph. Such graphs, in fact, can be easily used in the development of efficient network infrastructures, such as for Peer–to–peer systems, where neither flooding nor complex routing protocols are to be used.

Subsequently, routing strategies that make use of an augmented topological awareness of the nodes have been investigated. In particular it was shown that if each node knows the long-range contacts of its closest $O(\log n)$ neighbors on the grid (*Indirect greedy routing* [8, 13]) then $O(\log^{1+1/s} n)$ expected hops suffice in an s–dimensional lattice, for any fixed s. Papers [11] and [12] consider the improvements obtainable over greedy routing in the case in which the topological awareness of a node is augmented by the knowledge of the long–range contacts of all neighbors of the node (*Neighbor–of–neighbors* routing).

Augmenting an overlay network with random links is also at the base of randomized Peer–to–Peer networks. Two examples are randomized Chord and Symphony [12]. Such networks are both obtained by adding $O(\log n)$ random long range links to each node of a ring. They allow to obtain optimal routing [12].

1.2 Low Randomness Small–World Networks

In a SW network, the additional long–range random links represent the chance, that play a large role in creating short paths through the network as a whole. In this paper we consider the following question:

> *Do the long–range contacts really need to be completely random, or some "long–range clustering" could instead be envisaged in such a "navigable" network?*

In other words, we investigate the problem of whether all the independence assumed in the Kleinberg's model among different long–range contacts is indeed necessary to assure the existence of short paths. We show that, up to a certain extend, the answer to the above question is that the same (greedy) routing performances can be maintained if we limit the amount of randomness nodes use in the choice of long–range contacts.

Why to reduce randomization? In the perspective of using SW graphs as a base for the development of network infrastructures, we notice that besides diameter and degree, a very important property for a such an infrastructure is the resilience to simultaneous node failures.

The resilience of a network grows with the clustering of the graph which provides the overlay network. High clustering provides several alternative pathways throw which flow can pass, thus avoiding the failed component [14]. Clustering is a very interesting concept that is found in many natural phenomena and, roughly speaking, determines how tightly neighbors of a given node link to each other [9]. In a graph the clustering is related to the level of randomization of the graph itself. In particular, it was shown in [16] that the smaller is the randomization the higher is the clustering of the considered system. It is also worth noticing that the use of randomization increases the difficulties in the implementation and testing of applications.

Therefore, it is worth investigating if, in analogy with real SW graphs [16], a SW interconnection can be obtained by using a limited amount of randomness. Such a small value would allow both SW requirements (small average path length and high clustering) to be obtained together with easy routing strategies.

1.3 Our Results.

We start from Kleinberg's model (that is an s–dimensional lattice augmented with long–range links) and proceed in two steps. In a first step, we limit the choice of long–ranges of a node u to be done only among those other nodes that differ from u itself in exactly one coordinate; namely, a node $u = \langle u_1, \ldots, u_s \rangle$ has its long–range contacts chosen among nodes $v = \langle v_1, \ldots, v_s \rangle$ for which there exists i ($1 \leq i \leq s$) such that $v_j = u_j$ for each $j \neq i$. We show that all the routing properties immediately translate to this restricted model, some times with easier proofs.

In a second part, we enforce even more the restriction nodes have in establishing their long–range contacts. We introduce the notion of *communities*: keeping the restriction that long–range links can only connect two nodes differing in exactly one coordinate, nodes are partitioned into (random) groups and all nodes belonging to the same group are subject to additional restrictions in the choice of their long range contacts depending on the group they belong to.

We show that a logarithmic (in the number of nodes) number of different communities is sufficient to assure the SW property for the resulting graph. Namely, we analyze the routing performances of greedy, indirect and neighbor–of–neighbor routing strategies in dependence of the number of communities; in particular, if the number of communities is at least logarithmic all the routing strategies attain the same performances as in Kleinberg's original model.

The main results presented in this paper, together with those known for the original Kleinberg model, are summarized in Table 1.

Road Map: In Section 2 we give the basic definitions and present the notation used in the paper. In Section 3 we define Restricted–Small–World networks and analyze their routing properties. Section 4 is devoted to the study of SW networks with communities. Section 5 concludes the paper.

1.4 Related Work

The Small–World phenomenon was first demonstrated by Milgram's famous experiment [10]. Such an experiment showed that not only pairs of people were connected by short chains of acquaintances but also that people were able to route letters in a few hops by forwarding them to one of their acquaintances. Since then there has been quite a lot of work in the literature concerning the SW phenomenon and models for SW networks. A first model was proposed by Watts and Strogatz [16]. Following the notion that SW graph interpolate between completely regular graphs and purely random graphs, such model is based on randomly rewiring each edge of a regular graph with a given probability p. For suitable values of p, this gives rise to networks having the short path lengths and the high level of clustering observed in studied real networks (e.g. the network of actors cooperations or the neural wiring diagram of the worm *C. elegans*).

The model we consider in this paper (namely, an s-dimensional grid augmented with long–range links according to a given probability distribution – cfr. Definition 1) was introduced by Kleinberg in [5]. Kleinberg's proposal gave rise to a vast subsequent literature on routing in SW like networks. The $O(\log^2 n)$ expected number of hops shown by Kleinberg for greedy routing was proved tight in [1]. Moreover [13] and [8] prove that if the knowledge of a node is augmented with the knowledge of the long–range contacts of some of its neighbors then $O(\log^{1+1/s} n)$ can be achieved by greedy routing in a s–dimensional lattice, for any fixed s. A review of the vast literature on the subject can be found in [7].

SW like networks have been considered in the context of routing in peer–to–peer systems. In particular [12] analyzes the use of neighbor–of–neighbor greedy routing strategy, were the local knowledge of a node is augmented with the knowledge of the long–range links of all its contacts. They show that such routing reaches the optimal $O(\log n/\log\log n)$ expected number of hops in SW percolation networks and in the randomized version of the Chord network [15] (both having $O(\log n)$ degree). Paper [2] introduces, in order to speed-up bootstrap in the Chord ring, the concept of classes; such concept has some resemblance to that of communities.

Paper	**Results**	**Avg #steps**	**Networks**	**Observations**
[5]	Greedy	$O((\log^2 n)/q)$	$\mathcal{K}(n,s,q)$	
[1, 13]	Greedy	$\Omega((\log^2 n)/q)$	$\mathcal{K}(n,s,q)$	
[11, 12]	NoN	$O((\log^2 n)/(q\log q))$	$\mathcal{K}(n,s,q)$	$s=1$
[13]	IR	$O(\log^{1+1/s} n)$	$\mathcal{K}(n,s,q)$	$s=O(1)$
[8]	IR	$O((\log^{1+1/s} n)/(q^{1/s}))$	$\mathcal{K}(n,s,q)$	$s=O(1)$
[8]	IR	$\Omega(\log^{1+1/s} n)$	$\mathcal{K}(n,s,q)$	$s=O(1)$
This paper	Greedy	$O((\log^2 n)/q)$	$\mathcal{R}(n,s,q), \mathcal{R}_c(n,s,q)$	
This paper	IR	$O((\log^{1+1/s} n)/(q^{1/s}))$	$\mathcal{R}(n,s,q), \mathcal{R}_c(n,s,q)$	
This paper	NoN	$O((\log^2 n)/(q\log q))$	$\mathcal{R}(n,s,q), \mathcal{R}_c(n,s,q)$	

Table 1. Performance of variants of greedy routing (see Section 2.1 for more details). Routing strategies: Greedy (Greedy routing), IR (Indirect greedy routing), NoN (Neighbor–of–neighbor greedy routing). Grid networks having n nodes, s-dimensions and q long-range contacts are considered: $\mathcal{K}(n,s,q)$ (Kleinberg–Small–World network $\mathcal{K}(n,s,q,p)$ with probability $p(d)$ proportional to d^{-s}, cfr. Definition 1), $\mathcal{R}(n,s,q)$ (Restricted–Small–World network, cfr. Definition 2), $\mathcal{R}_c(n,s,q)$ (Small–World network with communities, cfr. Definition 3).

2 Preliminary Notation and Definitions

In the following we denote by d the distance between two points $v = \langle v_1, \ldots, v_s \rangle$ and $u = \langle u_1, \ldots, u_s \rangle$ on an s-dimensional toroidal lattice having $n = m^s$ nodes. The metric distance is defined as $d(v,u) = \sum_{i=1}^{s} d_i$ where $d_i = (u_i - v_i) \bmod m$.

Definition 1. ***Kleinberg–Small–World network (*** $\mathcal{K}(n,s,q,p)$ ***):*** *Consider n nodes lying on a toroidal s-dimensional grid $\{0, \ldots, m-1\}^s$ where each node maintains two types of connections:*

> **short–range contacts** *($2s$ connections): Each node has a direct connection to every other node within distance 1 on the grid;*
> **long–range contacts** *(q connections): Each node v establishes q directed links independently according to the probability distribution p (on the integers): each link has endpoint u with probability $p(d(u,v))$.*

All reported results on Kleinberg's model assume $p(d)$ proportional to d^{-s} with normalization factor $\sum_{u,v} d(u,v)^{-s}$. We notice that in Definition 1, we assume the interconnection be formed by a torus while the original Kleinberg's model was based on a grid. For sake of simplicity we will show our results on the torus, however they could be also obtained in case of a grid (e.g. with no wrap–around).

In the following we will denote by:

- $N(v)$ the neighborhood of v, that is, set of $2s$ neighbors of v on the s-dimensional torus;
- $L(v)$ the set of the q long-ranges contacts of node v;
- $N_r(v) = \{v \mid d(u,v) \leq r\}$ the ball of radius r and center v.

We recall that,

$$|N_r(v)| = \frac{2^s}{s!} \cdot r^s + \nu(r), \qquad (1)$$

where $\nu(x)$ is a positive polynomial of degree $s - 1$[3].

2.1 Routings Strategies

Consider a SW network $\mathcal{K}(n,s,q,p)$. We shortly review the routing strategies adopted in the Small–World related literature and that will be subsequently used in this paper. We denote by v be the node currently holding the message and by t the target key.

Greedy Routing uses only the local knowledge of v: a message is forwarded along the link (either short or long) that takes it closest to the target t.

Indirect greedy routing and *neighbor–of–neighbor greedy routing* are obtained through an additional *topological awareness* given to the nodes: Each node v is in fact aware of the long-range contacts of some other nodes.

Indirect Greedy Routing (IR) assumes that node v is aware of the long-range contacts of its $|N_r(v)|$ closest neighbors, for some $r > 0$. Formally, IR entails the following decision:

1. Let $L_r(v) = \bigcup_{u \in N_r(v)} L(u)$ denote the set of all long-ranges contacts of nodes in $N_r(v)$;

2. Among the nodes in $L_r(v) \cup N(v)$, assume that z is the closest to the target (with respect to the distance $d()$): If $z \in L(v) \cup N(v)$ then route the message from v **to** z directly; otherwise, let $z \in L(u)$ ($u \in N_r(v)$ for some $u \neq v$) and route the message from v **to** z **via** u.

Neighbor–of–Neighbor Greedy Routing (NoN) assumes that each node knows its long-range contacts, and on top of that it holds the long-range contacts of its neighbors. Here we restrict ourself to consider only the long-range contacts of nodes in $L(v)$, this will be sufficient to get the improvements over greedy assured by NoN. Formally, a NoN greedy step entails the following decision:

1. Consider the set $L^2(v) = \cup_{u \in L(v) \cup \{v\}} L(u)$ of long–range contacts of the long–ranges of v;

2. Among the nodes in $L^2(v) \cup N(v)$, let z be the closest to the target (with respect to the distance $d()$): If $z \in L(v) \cup N(v)$ then route the message from v **to** z directly, otherwise let $z \in L(u)$ ($u \in L(v)$ for some $u \neq v$) and route the message from v **to** z **via** u.

We recall that in the *NoN routing*, u may not be the long-range neighbor of v which is the closest to the target; indeed the algorithm could be viewed as a greedy algorithm on the square of the graph induced by the long-range contacts.

3 Restricted–Small–World networks

In a restricted–Small–World network we allow each node to make long–range connections only with nodes that differ from it in exactly one coordinate. A connection between two nodes u and v is created with probability proportional to $d^{-1}(u,v)$. In particular this probability is $p(d(u,v)) = \frac{1}{\lambda \cdot d(u,v)}$, where λ is the inverse normalized coefficient, $\lambda = s \cdot \sum_{j=1}^{m} \frac{1}{j} \approx \ln n$. Different connections are established by independent trials.

Definition 2. *A* ***Restricted–Small–World network (***$\mathcal{R}(n,s,q)$***):*** *is a network* $\mathcal{K}(n,s,q,p)$ *with probability distribution* p *s.t. for any* u *and* v*, the probability of having a long–range link from* u *to* v *is*

$$p(d(u,v)) = \begin{cases} \frac{1}{d(u,v)\ln n} & \text{if } v \text{ and } u \text{ differ in exactly one dimension} \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to observe that any outgoing link goes along a generic dimension i with probability:

$$\sum_{j=1}^{m} p(j) \approx \frac{\ln n^{1/s}}{\ln n} = 1/s.$$

We will show that all the results obtained on Kleinberg's SW networks [5, 13, 8, 12] can be easily proved to remain valid in spite of the restrictions we impose on the long-range connections.

3.1 Greedy Routing

Theorem 1. *The average path length is* $O\left(\frac{\log^2 n}{q}\right)$ *for the greedy routing on* $\mathcal{R}(n,s,q)$ *when* $1 \leq q \leq \log n$*.*

Proof. Consider a generic node $v = \langle v_1, \ldots, v_s \rangle \in \{0, \ldots, m-1\}^s$ that holds a message destined for node $t = \langle t_1, \ldots, t_s \rangle$ at distance d. For each $i = 1, \ldots, s$, let d_i the distance between v and t on dimension i that is $d_i = (t_i - v_i) \bmod m$.

The routing proceeds by greedily diminishing the distance on dimension 1 first, until $d_1 = 0$, then the distance on dimension 2 is considered, and so on until the target is reached.

Consider now a fixed dimension i. Clearly the analysis is equivalent to that of greedy routing on a ring with long-range connections added according to Definition 1.
Let ϕ denote the event that the current node is able to diminish the remaining distance, from d_i to at most d_i/k in one hop. The expected number of nodes encountered before a successful event ϕ occurs is $O\left(\frac{k \log n}{q}\right)$. Since the maximum number of times the remaining distance could possibly be diminished is $\log_k d_i \leq \log_k m$, $d_i \leq m$, it follows that the average number of hops we need on each dimension is $O\left(\frac{k \log n}{q} \cdot \log_k m\right)$. By repeating for each dimension (that is , up to s times),

$$O\left(s \cdot \frac{k}{\log k} \frac{\log m \log n}{q}\right) = O\left(\frac{k}{\log k} \frac{\log^2 n}{q}\right). \tag{2}$$

By choosing $k = 2$, the result follows. □

3.2 Indirect Greedy Routing

In the following we analyze *indirect greedy routing*. Step 2 in the definition of IR routing can be specialized to $\mathcal{R}(n, s, q)$ as follows:

> *2'.* Among the nodes in $L_r(v) \cup N(v)$, assume that z is the closest to the target (with respect to the distance $d()$): If $z \in L(v) \cup N(v)$ then route the message from v **to** z directly, otherwise first route the message from v **to** $u \in N_r(v)$ (at most r hops), then use the u's long-range to z (1 hop), and finally go from z **to** a node w using a inverse path with respect to the path from v to u, in such a way that w differs from v in exactly one dimension (at most r hops). Formally, if $u = \langle v_1 + d_1, v_2 + d_2, \ldots, v_s + d_s \rangle$ then $w = \langle z_1 - d_1, z_2 - d_2, \ldots, z_s - d_s \rangle$.

Let $N = |N_r(v)|$. We can repeat the proof of Theorem 1 by noticing that Nq long-ranges contacts are available at each greedy step and choosing the parameter k in equation (2) so that $qN = k \ln n$. Hence, we get that the number of indirect steps to reach the destination is $O\left(\frac{k}{\log k} \frac{\log^2 n}{qN}\right) = O(\log_k n)$. Each step requires at most $2r + 1$ hops where, by (1), $r = O\left(s \left(\frac{k \ln n}{q}\right)^{1/s}\right)$. Hence, the average path length is

$$O\left(r \log_k n\right) = O\left(\frac{s \ln n}{\ln k} \left(\frac{k \ln n}{q}\right)^{1/s}\right).$$

Choosing $k = e^s$, we obtain the following result.

Corollary 1 *The average path length is* $O\left(\frac{\log^{1+1/s} n}{q^{1/s}}\right)$ *for the indirect routing on* $\mathcal{R}(n, s, q)$ *when each node is aware of the long-range contacts of its* $\frac{e^s \ln n}{q}$ *closest neighbors.*

Remark 1. Unlikely in [8, 13], where the same result holds only for value of the parameter s independent of n (multiplicative factors in s are discarded in the asymptotic notation), our results are expressed also for a non-constant number s of dimensions. In particular the results in [8, 13] are obtained using an awareness of $O(\log n/q)$ neighbors; this awareness allows to obtain an average path length $O\left(s \cdot \frac{\log^{1+1/s} n}{q^{1/s}}\right)$.

Corollary 2 *The average path length is $O(\log n)$ for the indirect routing on $\mathcal{R}(n,s,q)$ when $s \geq \ln\ln n$ and each node is aware of the long-range contacts of its $\ln^2 n/q$ closest neighbors.*

3.3 NoN Greedy Routing

In the case of neighbor–of–Neighbor greedy routing, we obtain the following result.

Theorem 2. *The average path length is $O\left(\frac{\log^2 n}{q\log q}\right)$ for NoN routing on $\mathcal{R}(n,s,q)$ when $1 < q \leq \log n$.*

4 Small–World networks with communities

In this section we impose more strict restrictions on the choice of long–range contacts by the nodes in the network. Namely, we assume that nodes are partitioned into c groups, called *communities*. Each node randomly choose one of the communities to belong to; each node in community i, for $0 \leq i < c$, can choose its long–range contacts only among a subset of nodes depending on the parameter i.

Definition 3. *A **Small–World network with communities** ($\mathcal{R}_c(n,s,q)$): is a network $\mathcal{K}(n,s,q,p)$ with probability distribution p such that for any u and v the probability of having a long–range from v to u is obtained as follows:*

i) Node v chooses the community its belongs to by uniformly at random selecting an integer c_v in the interval $[0,c)$.

ii) Node v chooses uniformly at random an integer σ in the set $[1,s]$.
Let $t = q \bmod s$, $T = \{\sigma, \sigma+1 \bmod s \ldots, \sigma+t-1 \bmod s\}$, and, for $i = 1\ldots s$,

$$q_i = \begin{cases} \lceil q/s \rceil & \text{if } i \in T \\ \lfloor q/s \rfloor & \text{otherwise.} \end{cases} \quad (3)$$

For $i = 1,\ldots,q$, the i^{th} long–range link from v has endpoint u at distance $d(v,u)$ with probability

$$p(d(v,u)) = \begin{cases} \frac{1}{q} & \text{if } v \text{ and } u \text{ differ in exactly one dimension and } u \text{ is a feasible} \\ & \text{endpoint (i.e. } d(v,u) = \lfloor \gamma_i^{\ell+\frac{c_v}{c}} \rfloor \text{ for some } \ell = 0,\ldots,q_i-1) \\ 0 & \text{otherwise.} \end{cases}$$

where γ_i denotes a real number satisfying $\ln\gamma_i = (\ln m)/q_i$.

Observation 1 *For each node v there are exactly q feasible endpoints each with probability $\frac{1}{q}$. Hence a feasible endpoint of a node v is a long-range of v with probability $1-\left(1-\frac{1}{q}\right)^q \geq 1-e^{-1} > 1/2$.*

4.1 Routing in Small–World networks with communities

In this section we show that by introducing communities we reduce the amount of randomness with no harm to the efficiency of the system.

The following preliminary result will be a tool in the analysis of the performances of the various routing strategies.

Lemma 1. *Fix a dimension i. Let $k > 1$ an integer and let Φ denote the probability that a node of a $\mathcal{R}_c(n,s,q)$ network (where $1 \leq q \leq \log n$) is able to diminish with one hop the distance on a fixed dimension i, from d_i to at most d_i/k, then $\Phi = \Omega\left(\frac{q}{k\log n}\right)$ if $c \geq \frac{2k\ln n}{q}$.*

Greedy Routing

Theorem 3. *The average path length is* $O\left(\frac{\log^2 n}{q}\right)$ *for the greedy routing on* $\mathcal{R}_c(n,s,q)$ *when* $1 \leq q \leq \log n$ *and* $c \geq \frac{4\ln n}{q}$.

Proof. By following the same arguments as in proof of Theorem 1 we only need to show that $\Phi = \Omega\left(\frac{q}{\log n}\right)$, where Φ denotes the probability that the current node is able to halve the distance on a fixed dimension i. Using Lemma 1 with $k = 2$ we obtain the desired value. □

Indirect Greedy Routing Consider an indirect greedy routing step as described in Section 3.2. Let v be the node currently holding the message, consider $N_r(v)$ such that it contains a set $N_c \subseteq N_r(v)$ of nodes which belong to different communities with $|N_c| = O\left(\frac{k\ln n}{q}\right)$. Fix any dimension i for which the distance from v to the target on dimension i is $d_i > r$ (otherwise we route on this dimension using short–range links). For each node in N_c, consider the event that one of its long–ranges allows to to diminish in one indirect step the distance on dimension i from d_i to at most d_i/k; since nodes in N_c belong to different communities, such events are independent. Then applying Lemma 1, the probability that at least one node in N_c is able to diminish, with one indirect step, the distance to the target on dimension i, from d_i to at most d_i/k is $\Omega(1)$.

If $c \geq \frac{2k\ln n}{q}$ and $|N_r(v)| \geq \frac{k\ln n}{q}$ we have $|N_c| = O\left(\frac{k\ln n}{q}\right)$ and we can reach the destination using $O(\log_k n)$ indirect greedy routing step. Therefore, since $r = O\left(s\left(\frac{k\ln n}{q}\right)^{1/s}\right)$ the average path length is as for Restricted–Small–World.

Corollary 3 *The average path length is* $O\left(\frac{\log^{1+1/s} n}{q^{1/s}}\right)$ *for indirect routing on* $\mathcal{R}_c(n,s,q)$ *when* $1 \leq q \leq \log n$, $c \geq \frac{2e^s\ln n}{q}$ *and each node knows the long-range contacts of its* $\frac{e^s\ln n}{q}$ *closest neighbors.*

NoN Greedy Routing By analyzing NoN Greedy Routing in SW network with communities, we can prove the following Theorem.

Theorem 4. *The average path length is* $O\left(\frac{\log^2 n}{q\log q}\right)$ *for the NON routing on* $\mathcal{R}_c(n,s,q)$ *when* $1 < q \leq \log n$ *and* $c > \log n$.

5 Conclusions

Our Theorems 3 and 4 answer in a positive way our initial question: *Do the long–range contacts really need to be completely random, or some "long–range clustering" could instead be envisaged in such a "navigable" network?* In a sense, we show that it is not necessary to use a completely eclectic network in order to obtain a SW. Indeed, such result can be obtained using a limited amount of heterogeneity, namely only a logarithmic number of communities.

A part of their theoretical interest, such networks can be used toward the design of efficient as well as easy to implement network infrastructures based on the SW approach. Diminishing the amount of randomness used for random links increases the clustering of the network. Hence, one can get interconnected networks which, in addition to convenient graph properties (such as low average path length and degree) and beside providing the efficient and easy routing algorithms (as offered by Kleinberg's model) offer an increased resilience (due to a higher clustering).

References

1. Lali Barrière, Pierre Fraigniaud, Evangelos Kranakis and Danny Krizanc. "Efficient Routing in Networks with Long Range Contacts". Proc. *DISC 01*, LNCS 2180, pp. 270–284, Springer, 2001.
2. Giovanni Chiola, Gennaro Cordasco, Luisa Gargano, Alberto Negro, and Vittorio Scarano. "Overlay networks with class". *Proc. I-SPAN 2005*, IEEE Press, Dec. 2005.
3. John H. Conway and Neil J.A. Sloane. "Low Dimensional Lattices VII: Coordination Sequences". *Proc. Royal Soc. A453*, pages 2369–2389, 1997.
4. Peter Sheridan Dodds, Roby Muhamad, and Duncan J. Watts. "An Experimental study of Search in Global Social Networks". In *Science Volume 301*, pages 827–829, August 2003.
5. Jon M. Kleinberg. "The Small–World Phenomenon: An Algorithm Perspective". In *Proc. of 32th ACM Symp. on Theory of computing (STOC '00)*, pages 163–170, May 2000.
6. Jon M. Kleinberg. "The Small–World Phenomena and the Dynamics of informations". In *Advances in Neural Information Processing Systems (NIPS)*, December 2001.
7. Jon M. Kleinberg. "Complex Networks and Decentralized Search Algorithms". In *International Congress of Mathematicians (ICM)*, 2006.
8. P. Fraigniaud, C. Gavoille, C. Paul. "Eclecticism Shrinks Even Small Worlds" *PODC 2004.*
9. D. Loguinov and A. Kumar and V. Rai and S. Ganesh Graph-Theoretic, "Analysis of Structured Peer-to-Peer Systems: Routing Distances and Fault Resilience". *In Proc. of the ACM SIGCOMM '03 Conference*, 2003.
10. Stanley Milgram. "The Small World Problem". In *Psychology Today*, pp. 60–67, May 1967.
11. Gurmeet Singh Manku, Mayank Bawa, and Prabhakar Raghavan. "Symphony: Distributed hashing in a Small World". In *Proc. of USENIX Symp.on Internet Technologies and Systems (USITS)*, March 2003.
12. Gurmeet Singh Manku, Moni Naor, and Udi Wieder. "Know thy Neighbor's Neighbor: The Power of Lookahead in Randomized P2P Networks". *Proceedings of 36th ACM Symp. on Theory of Computing (STOC '04)*, pages 54–63, June 2004.
13. Chip Martel and Van Nguyen. "Analyzing Kleinberg's (and other) Small–world models". *In the 23rd ACM Symposium on Principles of Distributed Computing, pages 179–188, 2004.*
14. David Newth and Jeff Ash. "Evolving cascading failure resilience in complex networks". In *Proc. of* 8*th Asia Pacific Symp. on Intelligent and Evolutionary Systems*, December 2004.
15. Ion Stoica, Robert Morris, David Liben-Nowell, David R. Karger, M. Frans Kaashoek, Frank Dabek, and Hari Balakrishnan. "Chord: A Scalable Peer-to-Peer Lookup Protocol for Internet Applications". In *IEEE/ACM Trans. Networking, Vol. 11, No. 1*, pp. 17–32, 2003.
16. Duncan J. Watts and Steven H. Strogatz. "Collective dynamics of 'smallworld' networks". *In Nature, volume 393, Macmillan Publishers Ltd*, pages 440–442, June 1998.