


PUBLICATION
INTERNE
N° 1849



**SMALL-WORLD NETWORKS:
IS THERE A MISMATCH BETWEEN THEORY AND PRACTICE?**

F. BONNET A.-M. KERMARREC M. RAYNAL

Small-World Networks: Is there a mismatch between theory and practice?

F. Bonnet* A.-M. Kermarrec** M. Raynal***

Systèmes communicants
Projet ASAP

Publication interne n° 1849 — Mai 2007 — 15 pages

Abstract: In small-world networks, each peer is connected to its closest neighbors in the network topology, as well as to additional long-range contact(s), also called shortcut(s). In 2000, Kleinberg showed that greedy routing in a n peer small-world network, performs in $O(n^{\frac{1}{3}})$ steps when the distance to shortcuts is chosen uniformly at random, and in $O(\log^2 n)$ when the distance to shortcuts is chosen according to a harmonic distribution in a d -dimensional mesh. Yet, we observe through experimental results that peer to peer gossip-based protocols achieving small-world topologies where shortcuts are randomly chosen, perform well in practice.

The motivation of this paper is to explore this mismatch and attempt to reconcile theory and practice in the context of small-world overlay networks. More precisely, based on the observation that, despite the fact that the routing complexity of gossip-based small-world overlay networks is not polylogarithmic (as proved by Kleinberg), this type of networks ultimately provide reasonable results in practice. This leads us to think that the asymptotic big $O()$ complexity alone might not always be sufficient to assess the practicality of a system. The paper consequently proposes a refined routing complexity measure for small-world networks. Simulation results confirm that random selection of shortcuts can achieve “practical” systems. Then, given that Kleinberg proved that the distribution of shortcuts has a strong impact on the routing complexity, arises the question of leveraging this result to improve upon current gossip-based protocols. We show that it is possible to design gossip-based protocols providing a good approximation of Kleinberg-like small-world topologies. Along, are presented simulation results that demonstrate the relevance of the proposed approach.

Key-words: Clustering protocol, Complexity analysis, Epidemic protocol, Grid, Kleinberg’s shortcut selection, Large scale system, Overlay network, Peer to peer system, Random-based sampling, Random topology, Routing, Shortcut, Simulation experiment, Small-world.

(Résumé : tsvp)

* IRISA, Campus de Beaulieu, 35042 Rennes Cedex, France Francois.Bonnet@irisa.fr

** IRISA, Campus de Beaulieu, 35042 Rennes Cedex, France Anne-Marie.Kermarrec@irisa.fr

*** IRISA, Campus de Beaulieu, 35042 Rennes Cedex, France raynal@irisa.fr

Petits mondes : Y a-t-il désaccord entre pratique et théorie ?

Résumé : Ce rapport présente une étude des petits mondes fondés sur des protocoles épidémiques.

Mots clés : Analyse de complexité, protocole épidémique, Grille, Sélection des liens longs à la Kleinberg, Système à grande échelle, Réseau couvrant, Système pair à pair, Echantillonnage aléatoire, Topologie aléatoire, Routage, Lien long, Expérimentation, Petit monde.

1 Introduction

Distributed systems have experienced a dramatic scale shift over the past decade. Peer to peer (P2P) overlay networks have been at the center of distributed systems research both in the theoretical and practical communities, often in a fully de-correlated manner though. The research yields rather different expectations whether theory or practice is considered. On one hand, practical implementations target effectiveness for the most frequent cases, potentially at the price of lack of theoretical “worst-case” guarantees. On the other hand, theoretical analysis provides lower bound guarantees without always leading to solutions viable in practice. Routing is one of the main issues encountered in these systems. Focusing on the cost of routing in 2-dimensional torus topologies, this paper is an attempt to *(i)* reconcile theory and practice in the context of small-world overlay networks, where each peer is connected to its closest neighbors in the topology and additional long-range contact(s) (shortcut(s)) and *(ii)* leverage both areas to provide provably-efficient systems. Its main contributions are the following.

- The paper first confronts lower bound results of routing in small-world networks with the efficiency achieved in practice using epidemic (also called gossip-based) protocols. We observe through simulations that the expected gap in the routing performance between the two approaches to select shortcuts in a small-world network, is not entirely reflected when it comes to practice. Based on this observation, we argue that the asymptotic complexity analysis alone is not sufficient to assess the practicality of a small-world topology. The paper refines this analysis and characterizes the cost of routing, in terms of the average number of hops in both the grid and uniform topologies.
- Indirectly, the paper provides a fresh look at epidemic-based overlay networks and argue that they can achieve small-world topologies. We investigate the improvement of epidemic-based small-world networks to fully leverage Kleinberg’s results in practical settings. More specifically, we provide the design and preliminary results of a gossip-based protocol, biasing the peer sampling component used to create shortcuts, so that it provides a good approximation of Kleinberg’s harmonic distribution.

Small-world networks Small-world networks have been introduced as an analytical way of understanding and exploiting the *six degrees of separation* stating that two random individuals are separated by small chains of acquaintances [14]. When applied to computing networks, this can be achieved by each node in a mesh, knowing its closest neighbors and having additional shortcuts in the graph. While Watts and Strogatz [19] considered shortcuts as picked up uniformly at random, Kleinberg refined this result, demonstrating that meshes augmented with shortcuts provide a polylogarithmic routing and navigation under a greedy routing protocol, as long as the distances from the peers to their shortcuts follow a specific distribution (*d*-harmonic) [12, 13]. One of the main results of Kleinberg’s work is the determination of the magnitude order of the routing complexity in such networks (this model is further detailed in Section 2). This result has been of the up-most importance in the community, leading to a full range of works improving upon the routing complexity based on an increase knowledge of the system or a slightly different greedy algorithm (e.g., [2, 7, 15]).

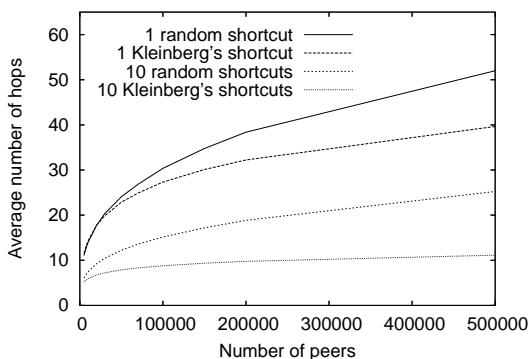
Epidemic-based overlay networks Epidemic-based (or gossip-based) protocols were first introduced to reliably disseminate data in large-scale networks [3, 4]. In the practical world, epidemic-based protocols have received an increasing attention as a scalable and reliable solution to build and maintain P2P overlay networks of arbitrary structure [9, 10]. Their convergence properties, reliability and simplicity make them however attractive for much more than data dissemination [6, 11]. More specifically, they have been applied in a wide variety of settings and are now turned into a generic tool to build and maintain large-scale overlay networks. It turns out that depending on the peer locally chosen for the interaction and the information exchanged, gossip-based protocols can be used to build overlay networks ranging from fully random-like unstructured networks to fully (DHT-like) structured networks (e.g., [8]). In this paper, we take a fresh look at overlay networks based on epidemic protocols and consider them with respect to small-world networks.

Epidemic protocols may be used to construct P2P overlay networks achieving graph properties very close to those of random graphs [5, 10, 16]. Typically, a gossip-based peer sampling service provides each peer with a set of long range contacts in a large-scale overlay network [10]. Resulting graphs are extremely robust and remain connected even in the presence of a large number of failures. In the context of this paper, we consider such a peer sampling service to be a way to implement randomly chosen shortcuts of small-world networks. Gossip-based protocols have also been used to create overlays optimized with respect to application-specific metric (e.g., clustering peers according to a proximity metric). It is actually relatively straightforward to use such gossip-based clustering protocols [17, 18] to choose the local neighbors in a small-world network.

Motivation The paper focuses on systems provided with *(i)* an underlying peer sampling gossip-based protocol that provides each peer with a random sample of the system (i.e., each peer is provided with shortcuts randomly chosen), and *(ii)* a gossip-based clustering protocol that provides each peer with a set of close neighbors (according to the considered

underlying topology). Such a combination creates therefore a small-world topology according to the Watts and Strogatz model [19].

While, as proved in [13], the routing expectation (expected number of hops) is $O(n^{\frac{1}{3}})$ (n being the total number of peers), it appears that performance results in practical settings turn out to be reasonable and follow the exponential convergence time of epidemic-based protocols, thus qualifying such networks for efficient routing. On another side, would the peer sampling service provide a sample following the distribution as defined in the Kleinberg model [13], the routing cost should greatly improve to $O(\log^2 n)$.



The previous observations constitute the starting point of our work. Simulating a system whose size ranges from 5000 to 500,000 peers, using six close neighbors and 1 or 10 shortcut(s) in a uniform topology, we first compared the random selection of shortcuts against Kleinberg's selection, and computed the average number of hops between any pair of peers in the system (this is depicted on the figure on the left). It appears that (as expected) the average number of hops between two peers is significantly improved when using the latter choice. The discrepancy increases with the size of the system and the number of shortcuts. These results show however that a random selection of neighbors keeps the average number of hops within reasonable bounds. This gives the motivation of our work, namely, understand the mismatch between practice and theory and leveraging the theory to influence and improve practical systems, thus bridging a gap between theory and practice in small-world networks.

Contributions Kleinberg's results were obtained on a d -dimensional grid topology. As such a topology is not always encountered in practice, we consider in this paper an additional topology where peers are randomly and uniformly distributed. Accordingly, the paper refers to the *grid topology* and the *uniform topology*, respectively.

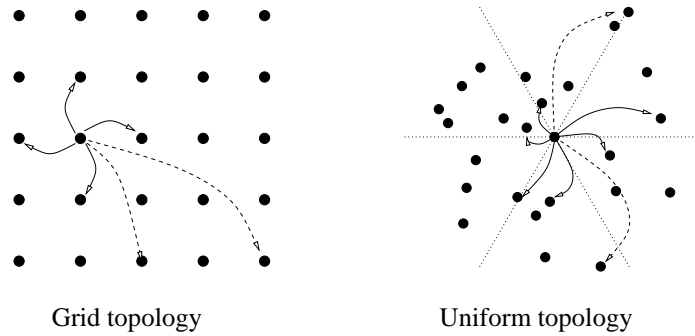
Based on the observed mismatch, we argue that relying on the magnitude order of the complexity analysis is not sufficient to draw conclusions on the practicality of an approach. The results obtained in practical systems tend to show that the system performs well regardless of the $O(n^{\frac{1}{3}})$ expected complexity of the greedy routing strategy. Using random shortcuts, the first contribution provides a refinement of the analysis of the routing cost in both the grid and uniform topologies. Simulations results demonstrate the good accuracy of the number of hops provided by the proposed analysis. As Kleinberg's selection of shortcuts yields an improved routing, we apply our approach to that model too. While we have not been able to provide results for the grid topology, we give an analysis for the uniform topology. We have also conducted simulations, the results of which are extremely encouraging. The table on the right summarizes the paper contributions. Finally, we propose the design and preliminary evaluations of a gossip-based protocol leveraging the theory to achieve an approximation of a Kleinberg-like small-world network.

Topology	Routing cost according to the way shortcuts are selected	
	Random	Kleinberg's model
Grid	$O(n^{\frac{1}{3}}) + \text{Section 3.2 (Eq 3)}$	$O(\log^2 n)$
Uniform	$O(n^{\frac{1}{3}}) + \text{Section 3.3 (Eq 5)}$	$O(\log^2 n) + \text{Section 4 (Eq 8)}$

Roadmap The paper is organized as follows. System models are defined in Section 2. Section 3 refines the routing analysis of the random selection of shortcuts in both the grid and the random topology. Section 4 applies this analysis to a small-world network where the selection of shortcuts follows the distribution proposed by Kleinberg, in the context of a uniform topology. Section 5 proposes to improve gossip-based sampling protocols by providing them with shortcuts defined from an approximation of Kleinberg's shortcuts selection. Finally, Section 6 concludes the paper.

2 System models and simulation setup

This section describes the system models considered in the paper. Dealing with failures is out of the scope of this paper (the reader interested in fault-tolerant routing can consult [1]). Network dynamic is also left out for this study and left for future work (robustness in face of high dynamics is one of the main strengths of epidemic-based protocols and we are confident that resulting protocols should sustain high dynamic).


 Figure 1: The grid topology and the uniform topology ($n = \ell^2 = 25$)

As announced in the Introduction, we consider the grid and the uniform topologies. When compared to the grid topology, the uniform topology is more practical in the sense that it is easier to achieve in a practical setting and applicable to a wider variety of applications. Theory has shown that Kleinberg's selection of shortcuts achieves polylogarithmic routing complexity. Yet, we observe that a random selection of shortcuts achieves lower but close enough performance. Consequently, these two shortcut selection processes are considered, resulting in four system models.

2.1 Base topologies

Both the grid and the uniform topologies considered here are based on a 2-dimensional torus. They differ in the way the peers are positioned on that torus (see Figure 1). In the following, a peer A is denoted by its name or its coordinates in the corresponding topology.

Grid topology That topology consists of a network made up of $n = \ell \times \ell$ peers corresponding to the points of a $\ell \times \ell$ square. The location of each peer in the grid is defined by coordinates (i, j) with $i, j \in \{0, 1, \dots, \ell - 1\}$. As we consider a torus, it should be noted that the peer located at $(0, 0)$ is in between the peers located at $(0, 1)$ and $(0, \ell - 1)$. More generally, in this topology, the distance between two peers A and B located at (i_1, j_1) and (i_2, j_2) is defined as their Manhattan distance, i.e.,

$$d_m = \min(|i_2 - i_1|, \ell - |i_2 - i_1|) + \min(|j_2 - j_1|, \ell - |j_2 - j_1|).$$

Uniform topology In this topology, the positions of the n peers are chosen uniformly at random in a 2-dimensional torus $[0 : 1] \times [0 : 1]$. More precisely, the pair of coordinates (A_x, A_y) associated with a peer A are chosen from the set $[0 : 1]$ following a uniform random distribution. In this topology, the distance between two peers located at (i_1, j_1) and (i_2, j_2) is the classical Euclidean distance, i.e.,

$$d_e = \sqrt{\min(|i_2 - i_1|, 1 - |i_2 - i_1|)^2 + \min(|j_2 - j_1|, 1 - |j_2 - j_1|)^2}.$$

2.2 Neighbor selection

In a small-world network, each peer, fully characterized by its location in the torus, maintains a *view* of the system. That view is made up of two sets of neighbors: a set of *local neighbors* (or *local contacts*), which are close neighbors in the graph and a set of long-range neighbors, called *shortcuts*, chosen according to a selection distribution.

We consider a *greedy routing* algorithm to navigate a small-world network. This means that, at each hop, a message is routed to a peer, the position of which is *closer* to the destination (according to the metric relevant to the considered topology), thus ensuring that the distance to the destination always decreases as the routing process progresses. At each routing peer, the neighbors from these two sets are considered to select the peer to which a message has to be routed.

2.2.1 Local contact selection

To ensure the effectiveness of a greedy routing, a peer needs to know about its closest neighbors according to the distance measured in the corresponding topology. Figure 1 represents local contacts with plain lines. The local contacts differ between the grid and the uniform topology. More explicitly, we have the following.

- **Grid topology:** To allow for a greedy routing, each peer must know at least its four closest neighbors which all are at distance 1. The local contacts of a peer (i, j) are peers $(i, j + 1 \bmod \ell)$, $(i, j - 1 \bmod \ell)$, $(i + 1 \bmod \ell, j)$, and $(i - 1 \bmod \ell, j)$.

- **Uniform topology:** To allow for a greedy routing, each peer must know at least six of its closest neighbors, one in each wedge of the space as shown in the right part of Figure 1. (If a node does not have a contact belonging to one of these wedges, it can easily be shown that the greedy routing may fail [20]¹.) Partitioning evenly the space around each peer into 60° wedges, and assigning a local contact belonging to each of these wedges ensures that the greedy routing can be implemented using only local contacts.

2.2.2 Shortcut selection

Shortcuts are added to the view of each peer to speed up the routing process, providing them with candidates to perform “large” routing steps. The complexity achieved by a greedy routing is highly sensitive to the way such shortcuts are chosen. We consider two selection algorithms in this paper, providing each peer with q shortcuts. In Figure 1, we have $q = 2$, and the two shortcuts are depicted with dashed lines.

- **Random selection:** As introduced in the Watts and Strogatz model [19] and implemented using the peer sampling protocol in the context of epidemic-based algorithms, the selection is done uniformly at random. Each peer A is provided with q shortcuts by choosing q peers uniformly at random from the set of all the peers of the network that are not local contacts of A .
- **Kleinberg’s selection:** As proposed in [12], shortcuts can be added according to a non-uniform distribution. Selecting shortcuts this way has proven to significantly reduce the cost of a greedy routing, achieving polylogarithmic complexity. In Kleinberg’s model, a peer A selects a peer B as a shortcut with a probability proportional to the value $\delta(B) = \frac{1}{d(A,B)^2}$ ($d()$ denotes the Manhattan distance or the Euclidean distance, according to the network model). More precisely, a peer B is chosen by a peer A with probability $\frac{\delta(B)}{\sum_{B' \in S} \delta(B')}$ where S denotes the set of all the peers that are not a local contact of A . In the following, this selection mode will be referred as *Kleinberg’s selection*.

2.3 Simulation setup

Due to page limitation, the simulation setup is not described in detail. The results are illustrated by comparing analytical results against simulation results obtained using the simulator PeerSim [21]. This simulator allows us to choose the topology, the number of nodes, the number of local contacts p , and the number of shortcuts q . For each generated network, a high number (500,000 if not specified otherwise) of pairs of peers have been randomly selected to evaluate the cost of the routing. The simulator also implements the gossip-based protocols evaluated in Section 5.

3 Small-worlds with randomly selected shortcuts

3.1 Preliminaries

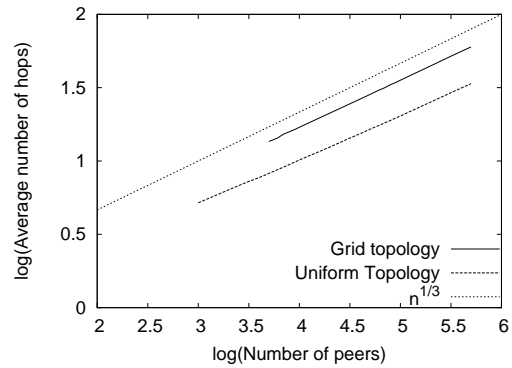
As already noticed, studying the complexity of greedy routing in a grid topology in [12, 13], Kleinberg proved that a random selection of shortcuts in a small-world network gives rise to a routing cost with an expected number of hops that is at least $\alpha n^{\frac{1}{3}}$, where n denotes the number of peers and α is a coefficient -not explicitly determined- that is independent of n .

The good performance achieved in small-world networks by epidemic protocols (and more specifically by the random selection of shortcuts achieved by peer sampling protocols), led us to think that knowing the value of α is actually interesting. The idea here is that knowing more precisely that value enables us to analyze efficient implementations, and allows consequently for a better understanding of why gossip-based protocols are practically efficient.

Although Kleinberg’s study is only on the grid topology, we are confident that the same kind of results can be extended to the uniform topology. As an attempt to check this experimentally, we conducted simulations to compare the routing cost using shortcuts defined from a random selection in both the grid and the uniform topology.

¹In practice, there is no need of this condition if the algorithm takes enough local contacts. It is possible to show that the probability of these special cases is then close to 0.

The corresponding results are presented in the figure on the right. We observe that, in both cases, the simulation results match Kleinberg’s analysis for the grid topology. As we can see, the $\frac{1}{3}$ coefficient of the complexity appears in both topologies (as expected by the theoretical analysis). When considering the grid topology, these simulations are in perfect agreement with the corresponding theoretical results. So, we are pretty confident that Kleinberg’s analysis in the context of a grid topology can certainly be applied to the random topology as well.



3.2 Routing analysis in the grid topology

In order to understand the unexpected good performance of practical implementations of the grid topology, we evaluate the routing cost in this network model more accurately. To that end, rather than dealing only with the magnitude order of the complexity, the routing cost has been analyzed by determining the number of hops needed to route from any peer to any other peer. As previously described, the $n = \ell^2$ peers are positioned on a $\ell \times \ell$ square over a torus. Let r denotes the radius of the local contacts: each peer knows the other peers that are within distance r of it. Let us recall that q denotes the number of shortcuts each peer randomly chooses in the network. For the purpose of the analysis, we first introduce the notion of k -neighborhood of a peer.

Evaluating the size of the k -neighborhood The k -neighborhood corresponds to the set of peers that are at distance k of a given peer. The size of k -neighborhood is required to compute probabilities related to the location of shortcuts. Due to the regularity of the grid topology, let us observe that, for any peer, there is one peer at distance 0 (itself, obviously), there are 4 peers at distance 1, 8 peers at distance 2, and more generally $4d$ peers at distance d for $d < \frac{\ell}{2}$.

For $d \geq \frac{\ell}{2}$ the number of peers at distance d depends on the parity of ℓ . If ℓ is odd, there are $4(\ell - d)$ peers at distance d for $d \geq \frac{\ell}{2}$. If ℓ is even, there are $2\ell - 2$ peers at distance $\frac{\ell}{2}$ and $4(\ell - d)$ peers at distance d for $\frac{\ell}{2} < d < \ell$ and 1 node at distance ℓ . These values are summarized in the table at the right. For the sake of simplicity, we assume that ℓ is odd. Thus, the previous numbers can be easily written as follows: each peer has $4 \min(\ell - d, d)$ peers at distance d , $1 \leq d < \ell$.

	odd ℓ	even ℓ
$d = 0$	1	1
$0 < d < \frac{\ell}{2}$	$4d$	$4d$
$d = \frac{\ell}{2}$		$2\ell - 2$
$\frac{\ell}{2} < d < \ell$	$4(\ell - d)$	$4(\ell - d)$
$d = \ell$	0	1

Expected number of hops for $r = q = 1$ The average number of hops needed to route from a peer to another peer depends solely on the distance between these peers. Intuitively, if the destination is near to the source, the routing process will be faster than if the peers are far from each other. So, let us define $f(d)$, a function that gives the average number of routing hops between any two peers that are at distance d . We have the following when $r = q = 1$.

- $d = 0$: The destination is the source. Obviously the number of required hop is $f(0) = 0$.
- $d = 1$: The destination is at distance 1 from the source. The destination belongs to the local contacts of the source peer. Then the number of required hop is consequently $f(1) = 1$.
- $d > 1$: The destination peer does not belong to the local contacts of the source peer (since $d > r = 1$). As, for the next hop, a greedy algorithm takes either a shortcut or a local contact, two cases need to be studied. If a local contact is chosen, the distance to the destination decreases only by one at each hop. If a shortcut is chosen, the distance can be greatly reduced if the shortcut is close to the destination. Interestingly, it is possible to evaluate the benefit of the shortcut (i.e., the distance between the shortcut and the destination). Let $d(i)$ be the probability that the shortcut is at distance i to the destination. We obtain the following recursive expression for $f(d)$:

$$\forall d > 1 : \quad f(d) = 1 + \left(\sum_{i=0}^{d-2} d(i)f(i) \right) + \left(1 - \sum_{i=0}^{d-2} d(i) \right) f(d-1). \quad (1)$$

The first term (value 1) corresponds to the hop the algorithm has to perform to progress towards the destination. The second term ($\sum d(i)f(i)$) corresponds to the use of the shortcut by the algorithm: after using the shortcut, a distance i remains to the destination with a probability $d(i)$. The last term corresponds to the case when a local contact is used; the distance decreases then only by one as explained before.

From $f(0)$, $f(1)$, and the recursive definition of $f(d)$, it is possible to compute $f(d)$ for any value of d . It remains however to determine $d(i)$. In fact, $d(i)$ is already known, as we show it. The number of peers that are at distance i of a given peer has been determined previously: it is the i -neighborhood of that peer. The value of $d(i)$ is related to this number: since shortcuts are randomly chosen, the probability that the shortcut is at distance i of the destination is equal to the number of peers at distance i divided by the total number of peers. Consequently, we have $d(i) = \frac{4 \min(i, \ell-i)}{n}$ ⁽²⁾ for $i > 0$ and $d(0) = 1$. Despite the fact that we do not know a closed formula for $f(d)$, that function can be efficiently computed using a mathematical software. Figure 2(a) shows that the average value of $f(d)$ over all possible distances, i.e., $\sum_{i=0}^{\ell-1} d(i)f(i)$, matches almost exactly the results obtained from simulations.

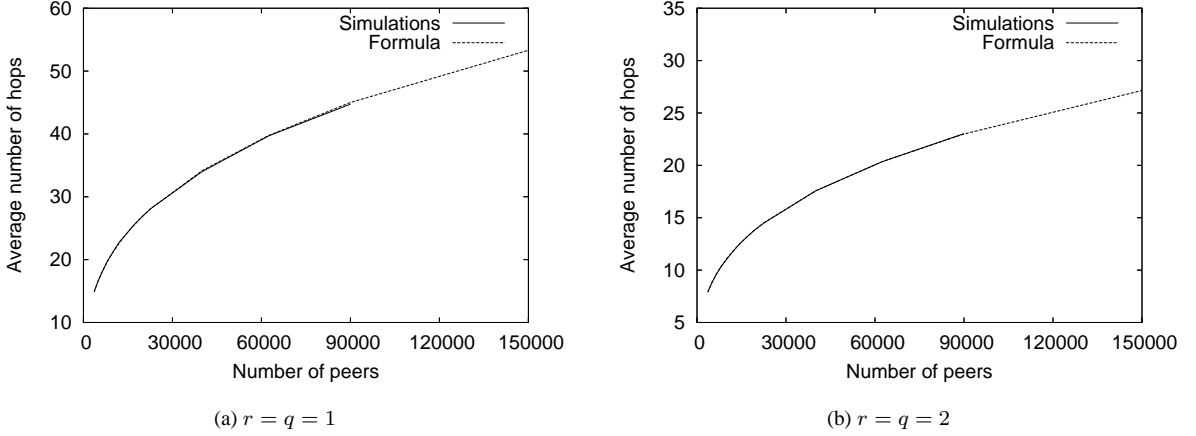


Figure 2: Comparison between simulations results and the formula, grid topology

Expected number of hops for any r, q The previous analysis can be generalized by taking a larger set of local contacts or more than one shortcut. To that end, appropriate adaptations in the recurrence formula have to be considered.

- $r > 1$: All the peers at distance r (or less) of a peer are its local contacts. It follows that for a distance i to the destination such that $i \leq r$, only one hop is required and consequently $f(i) = 1$. Moreover, in any other case where the routing uses a local contact, the distance to the destination decreases by at least r in one hop.
- $q > 1$: Each peer chooses q shortcuts. It follows that the probability to have a shortcut close to the destination increases. Let $d_q(i)$ denote the probability that the closest shortcut to the destination is at distance i . The function $d_q()$ can easily be recursively computed as follows:

$$d_q(i) = d_{q-1}(i) \left(1 - \sum_{k=0}^{i-1} d_1(k) \right) + \left(1 - \sum_{k=0}^i d_{q-1}(k) \right) d_1(i). \quad (2)$$

The first term corresponds to the case when the best shortcut amongst $q - 1$ shortcuts is at distance i and the last shortcut is at least at distance i . The second term corresponds to the case when $q - 1$ shortcuts are at distance greater than i and the last shortcut is at distance i . These two cases describe all the possibilities.

We consequently obtain the following generalized recurrence formula for the grid topology:

$$\forall d > r : \quad f(d) = 1 + \left(\sum_{i=0}^{d-r-1} d_q(i) f(i) \right) + \left(1 - \sum_{i=0}^{d-r-1} d_q(i) \right) f(d-r). \quad (3)$$

Figure 2(b) represents the routing cost for a system with $r = 2$ and $q = 2$. Again, we observe an almost perfect match between simulations and computations.

²The denominator should be $n - 5$ because a shortcut can not be chosen amongst the local contacts of a peer. However, in a large scale network, n is a reasonable approximation of $n - 5$.

3.3 Routing analysis in the uniform topology

The analysis for the uniform topology is a little bit more involved. The main difference with the grid is related to the distance between peers. Instead of being integers, distances are now real numbers in $[0, \frac{\sqrt{2}}{2}]$ (let us recall that the nodes are on a $[0 : 1] \times [0 : 1]$ torus). Thus, the analysis has to be adapted to take in account this new constraint. The function $f()$ has to accept real numbers instead of integers as input.

Let us recall that p denotes the number of local contacts and q denotes the number of (randomly chosen) shortcuts. In the grid topology, the local contacts are defined from a radius; we need now to estimate the corresponding radius r_p . As the torus $[0 : 1] \times [0 : 1]$ represents an area of 1 unit, the average surface is $\frac{p}{n}$ when considering p peers. If a disk is used to approximate that area, its radius is $r_p = \sqrt{\frac{p}{n\pi}}$.

In the first previous analysis, there were probabilities on the position of the best shortcut: $d_q(i)$ corresponded to the probability that the best shortcut (among q possible shortcuts) is at distance i of the destination. Here, as $f(d)$, the function $d()$ is a function with real inputs: $d_q(i)$ denotes the density of probability that the best shortcut is at distance i of the destination. As proved in appendix A the function $d_1()$ is as follow:

$$\begin{aligned} 0 \leq i \leq \frac{1}{2} : & \quad d_1(i) = 2\pi i, \\ \frac{1}{2} < i \leq \frac{\sqrt{2}}{2} : & \quad d_1(i) = -2\pi i + 8i \arcsin\left(\frac{1}{2i}\right). \end{aligned}$$

The other functions $d_q()$ can be computed recursively from $d_1()$ using the same method as in the grid topology (see Equation 2), namely (the derivative symbol dk is omitted in order not to overload the formulas):

$$d_q(i) = d_{q-1}(i) \left(1 - \int_{k=0}^i d_1(k)\right) + \left(1 - \int_{k=0}^i d_{q-1}(k)\right) d_1(i). \tag{4}$$

Taking into account these modifications, we obtain the following recurrence formula for the uniform topology:

$$\begin{aligned} 0 < d \leq r_p : & \quad f(d) = 1, \\ \forall d > r_p : & \quad f(d) = 1 + \left(\int_{i=0}^{d-r_p} d_q(i)f(i)\right) + \left(1 - \int_{i=0}^{d-r_p} d_q(i)\right) f(d - r_p). \end{aligned} \tag{5}$$

Figure 3 compares simulations wrt the formula for networks of size ranging from 4,000 up to one million. Each peer knows $p = 20$ local contacts and $q = 2$ random shortcuts are chosen. We observe a slight discrepancy between our formula and the simulation results. This comes from the fact we always consider that a local contact is located on the circle of radius r_p , and consequently the distance to the destination is reduced by r_p each time a local contact is used. In a real setting, as a local contact may be within the disk of radius r_p , the gain may be smaller (it would actually be possible to take this fact into account at the price of a much more complicated formula).

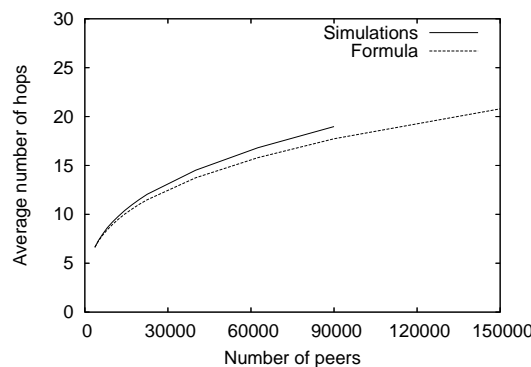


Figure 3: Comparison between simulations results and the formula (for $p = 20$ and $q = 2$), uniform topology

4 Small-worlds with shortcuts according to Kleinberg's distribution

Section 3 has investigated the routing cost of greedy routing in small-world networks where shortcuts are randomly chosen. Here, the performance of routing algorithms based on Kleinberg's shortcut selection is analyzed. The analysis is presented only for the uniform topology. This is due to the fact that, when one wants to determine the location of the best shortcuts in the grid topology, the absence of symmetry generates an extremely large number of cases, which makes such a computation "unfeasible". As opposed to the cases studied so far, the distribution of these locations does not depend only on the distance between two peers. (This problem does not appear in the uniform distribution, rendering the analysis feasible.)

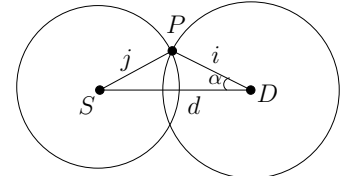
Local contacts analysis The effect of the p local contacts on the routing performance is the same, be the q shortcuts selected randomly or according to Kleinberg's distribution. So, in our analysis, the study of locals contacts remains the same, namely, there is an estimated radius $r_p = \sqrt{\frac{p}{n\pi}}$ which corresponds to the area approximately covered by the local contacts.

Distribution of shortcut locations Since the shortcuts are no longer chosen following the uniform distribution, that distribution becomes more complex. As described in Section 2, a shortcut B is selected by a peer A following a probability proportional to the inverse of the square of the distance between A and B . We already know the distribution of the distance between two peers, that is expressed by the function $d_1(\cdot)$. We consequently obtain the density of probability $dist(\cdot)$ that the shortcut B is at distance i from A from the following formula (defined only for $i > r_p$ since shortcuts can not be taken amongst local contacts):

$$\forall r_p < i \leq \frac{\sqrt{2}}{2} : \quad dist(i) = \frac{\frac{d_1(i)}{i^2}}{\int_{k=r_p}^{\frac{\sqrt{2}}{2}} \frac{d_1(k)}{k^2}}. \quad (6)$$

Expected number of hops As in the previous analysis, we need to compute the probability for a peer to use one of its shortcuts in the routing process. The previous recurrence formulas remain correct if the function $d(\cdot)$ is appropriately adapted. Let us start with only one shortcut ($q = 1$). We are looking for the density of probability $d'_1(\cdot)$ that the shortcut is at distance i to the destination.

The figure on the right depicts the following situation: S is the source node, D the destination, and P the shortcut; d , i , and j denotes the distances SD , DP , and SP , respectively. From a geometrical analysis we conclude that $j = \sqrt{d^2 + i^2 - 2di \cos(\alpha)}$. Summing all the possible positions of P over the circle gives:



$$d'_1(d, i) = i \int_{\alpha=0}^{2\pi} dist\left(\sqrt{d^2 + i^2 - 2di \cos(\alpha)}\right).$$

Let us notice that, as opposed to the previous analysis, where the function $d'_1(\cdot)$ depends on a single parameter (i), this function now depends also on a second parameter measuring the Euclidean distance d between the source and the destination. More generally, the function $d'_q(\cdot)$, for more shortcuts $q > 1$, can be computed (similarly to Relation 4) from the value of $d'_1(\cdot)$. We then obtain the following value:

$$d'_q(d, i) = d'_{q-1}(d, i) * \left(1 - \int_{k=0}^i d'_1(d, k)\right) + \left(1 - \int_{k=0}^i d'_{q-1}(d, k)\right) d_1(d, i). \quad (7)$$

Finally, similarly to Equation 5, the routing cost can be computed with the following recurrence:

$$\forall d > r_p : \quad f(d) = 1 + \left(\int_{i=0}^{d-r_p} d'_q(d, i) f(i)\right) + \left(1 - \int_{i=0}^{d-r_p} d'_q(d, i)\right) f(d - r_p). \quad (8)$$

Figure 4 shows the impact of the number of shortcuts on the average number of hops in a 200,000 peer system and compares the random shortcut selection against the Kleinberg's one. As expected, the routing performance improves with the number of shortcuts regardless of the shortcuts selection. However, we observe that the routing performance decreases faster with Kleinberg's selection: with one shortcut only, Kleinberg's selection improves the number of hops of 16% with one shortcuts and 48% with 10 shortcuts. Although the average number of hops is less than 20 in a 200,000 peer system using randomly selected shortcuts, and therefore qualify for a practical system, there is still room for improvement by leveraging the Kleinberg's selection method in practice.

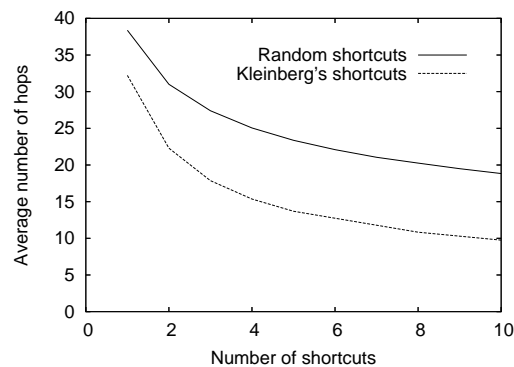


Figure 4: Impact of the number of shortcuts on the routing performance in a 200,000 peer system

5 Kleinberg-like epidemic-based small-world networks

Gossip-based protocols have been recognized as a sensible and efficient paradigm for building peer to peer overlay networks of arbitrary structure. Current gossip-based protocols can achieve already small-world topologies with random shortcuts. This section presents the design of a gossip-based protocol implementing a small-world overlay network where shortcuts are selected according to an approximation of Kleinberg's selection.

A generic gossip-based protocol Let us consider a system made up of n peers uniquely defined by their coordinates³. Each peer maintains a set of neighbors (IP address of other peers in the system) called its *view*, reflecting its knowledge of the membership of the system⁴. This creates a connection graph, where an edge between two peers A and B means that each of them belongs to the view of the other one. Each peer executes an active thread and a passive thread. The size of a view is c (c being a parameter of the system). Periodically each peer A runs the active thread: (i) it selects from its view a peer B to gossip with, (ii) sends a message to B containing a subset of its view, and (iii) merges its own view with the information received from B , truncating its view back to c . The passive thread on A (i) sends to B a subset of its view upon receiving a gossip message from B and (ii) merges its own view with the information received from B , truncating its view back to c . (Many details of the protocol are omitted due to space limitations, but full details can be found in [10].) It turns out that the resulting connection graph strongly depends on the peer selection, the state exchanged during the gossip and the processing of the state to compute the resulting view.

Local contact selection Using the generic protocol described above, a clustering gossip-based protocol may be used to create the local contacts in a small-world. Let us consider the uniform topology case. As shown in Figure 1, each peer needs to maintain a peer in each of the six wedges attached to each peer (recall that each wedge covers 60°). In such a context, a gossip-based algorithm may easily be implemented as follows. *Peer selection*: the closest peer, according to the Euclidean distance, in one of the wedge of the circle, is chosen to gossip with (at random if several candidates). *State exchanged*: in this preliminary version, the whole views are exchanged. *State processing*: the closest peers, according to the Euclidean distance and optimizing along all directions, are kept. (In a dynamic system, such a clustering protocol might be run in parallel with a random peer sampling protocol.)

Random shortcuts It has been shown in [10] that such a gossip-based protocol can be used to provide a random sample to each peer. (Running the Cyclon [16] protocol for example results in a graph, the properties of which are close to those of random graphs with respect to the average path length, clustering coefficient and diameter.) Therefore, this gossip-based protocol can be run together with the clustering protocol mentioned above in a straightforward manner, thereby implementing the random shortcuts of a small-world random graph. For the purpose of comparison, we use the Cyclon protocol in our simulations of gossip-based random small-world overlays (Figure 5).

Creating Kleinberg's shortcuts In order to leverage the potential of the d -harmonic distribution as defined in [12], we bias the peer sampling service in order to approximate the distribution advocated by Kleinberg. To that end, we propose to change the nature of the state exchanged between peers upon gossip, in order to match as closely as possible

³Peers could join and leave the system dynamically although we do not consider dynamics in our experiments.

⁴As, in this paper, we consider the use of gossip-protocols for overlay maintenance only, the state of a peer is fully represented by its view.

this distribution. The algorithm on peer A is implemented as follows. *Peer selection*: a peer B is chosen uniformly at random from A 's view (of size c). *State exchanged*: $\frac{c}{2}$ peers are chosen in A 's view to reflect the Kleinberg's distribution. More specifically, a peer A keeps a peer C from its view with a probability proportional to $\delta(C)$ (as defined in 2.2.2, i.e., $\delta(C) = \frac{1}{d(A,C)^2}$ where d is the distance separating A and C). The remaining $\frac{c}{2}$ peers are sent to B during the gossip operation. This enables to minimize the loss of information during the gossip operation. *State processing*: the view is purged from the $\frac{c}{2}$ non chosen entries from the view, sent over during the gossip operation. (This is described in more details in Appendix B.)

Figure 5 compares the routing performance of a gossip-based protocol implementing an approximation of Kleinberg distribution against a gossip-based protocol implementing the random selection (peer sampling service). In addition we compared those simulations with the hypothetical ideal simulation mode⁵. Those preliminary results confirm that gossip-based protocols can be used to achieve in a fully decentralized way a close approximation of Kleinberg-like small-world overlay networks. We are currently investigating a broader exploration of the parameter space.

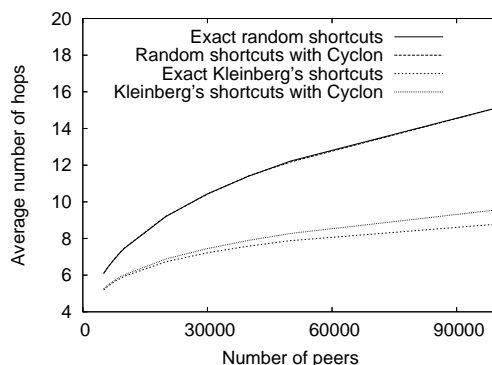


Figure 5: Ideal versus gossip-based selection of random and Kleinberg shortcuts in small-world overlay networks

6 Conclusion

This work has been motivated by the observation that, despite theoretical evidence that a random selection of shortcuts in small-world networks should lead to poor routing performances, such systems are reasonably efficient in practice. To have a better understanding of this gap between theory and practice, we (i) precisely analyzed the average number of routing hops (required in a greedy routing strategy), both in the grid topology and the uniform topology, and (ii) compared the random selection and Kleinberg's selection of shortcuts. Not surprisingly, this analysis confirms the superiority of Kleinberg's selection, but nonetheless demonstrates that the random selection of shortcuts can be considered in practice. Simulation results show that there is an almost perfect match between the results observed in practice and our analysis. Small-world topologies can already be implemented using a two layer gossip-based protocol. We then propose the design and preliminary evaluation of a gossip-based protocol to implement a Kleinberg small-world topology in a fully decentralized way, biasing the peer sampling service to approximate a Kleinberg selection of shortcuts. To conclude, we believe that, contrarily to what could have been a priori suspected, there is no mismatch between theory and practice. This study emphasizes the need to go beyond the big O complexity analysis to evaluate the practicality of small-world overlay networks. Further investigation is now needed to refine the biased gossip-based protocol, explore arbitrary topologies, apply such techniques to other routing strategies and consider dynamic settings.

Acknowledgment

We would like to thank Rachid Guerraoui and Etienne Rivière for their constructive comments on a draft of this paper.

References

- [1] Aspnes J., Diamadi Z. and Shah G. Fault-tolerant Routing in Peer-to-Peer Systems. *Proc. 21th ACM Symposium on Principles of Distributed Computing (PODC'01)*, ACM Press, pp. 223-232, 2001.

⁵In this mode, all nodes are considered in the simulator to select the shortcuts.

- [2] Barrière L., Fraigniaud P., Kranakis E. and Krizanc D., Efficient Routing in Networks with Long Range Contacts. *Proc. 15th Int'l Symp. on Distributed Computing (DIS'01)*, Springer-Verlag LNCS #2180, pp. 270–284, 2001.
- [3] Birman K.P., Hayden M., Ozkasap O., Xiao Z., Budiu M. and Minsky V., Bimodal Multicast. *ACM Transactions on Computer Systems*, 17(2):41-88, 1999.
- [4] Demers A.J., Greene D.H., Hauser C., Irish W. and Larson J., Epidemic Algorithms for Replicated Database Maintenance. *Proc. 6th ACM Symposium on Principles of Distributed Computing (PODC'87)*, pp. 1-12, 1987.
- [5] Eugster P.Th., Guerraoui G., Handurukande B., Kermarrec A.-M. and Kouznetsov P., Lightweight Probabilistic Broadcast, *ACM Transactions on Computer Systems*, 21(1):341-374, 2003.
- [6] Eugster P.Th., Guerraoui G., Kermarrec A.-M. and Massoulié L., Epidemic Information Dissemination in Distributed Systems. *IEEE Computer*, 37(5):60-67, 2004.
- [7] Fraigniaud P., Gavoille C. and Paul C., Eclecticism Shrinks even Small Worlds. *Proc. 23th ACM Symposium on Principles of Distributed Computing (PODC'04)*, ACM Press, pp. 169-178, 2004.
- [8] Fraigniaud P., Gauron Ph. and Latapy M., Combining the Use of Clustering and Scale-free Nature of User Exchanges into a Simple and Efficient P2P System. *Proc. European Conf. on Parallelism (EUROPAR'05)*, 2005.
- [9] Jelasity M. and Babaoglu O., *T-Man*: Gossip-based Overlay Topology Management. *Proc. Engineering Self-Organising Applications (ESOA'05)*, 2005.
- [10] Jelasity M., Guerraoui G., Kermarrec A.-M. and van Steen M., The Peer Sampling Service: Experimental Evaluation of Unstructured Gossip-based Implementations. *Proc. 5th ACM/IFIP/USENIX Int'l Conference on Middleware*, Springer-Verlag LNCS, pp. 79-98, 2004.
- [11] Jelasity M., Montresor A. and Babaoglu O., Gossip-Based Aggregation in Large Dynamic Networks. *ACM Transactions on Computer Systems*, 23(3): 219-252, 2005.
- [12] Kleinberg J., Navigation in a Small World. *Nature*, 845(406), 2000.
- [13] Kleinberg J., The Small-World Phenomenon: an Algorithmic Perspective. *Proc. 32nd ACM Symposium on Theory of Computing*, ACM Press, pp. 163-170, 2000.
- [14] Milgram S., The Small-World Problem. *Psychology Today*, 61(2):60-67, 1967.
- [15] Naor M. and Wieder U., Know Thy Neighbor's Neighbor: Better Routing for Skip-Graphs and Small Worlds. *Proc. Int'l Workshop on Peer-to-Peer Systems (IPTPS'04)*, Springer-Verlag LNCS #3279, pp. 269-277, 2004.
- [16] Voulgaris S., Gavidia D. and van Steen M., CYCLON: Inexpensive Membership Management for Unstructured P2P Overlays. *Journal of Network and Systems Management*, 13(2):197-217, 1005.
- [17] Voulgaris S., Rivière E., Kermarrec A.-M. and van Steen M., Sub-2-Sub: Self-Organizing Content-Based Publish and Subscribe for Dynamic and Large Scale Collaborative Networks. *Proc. 5th Int'l Workshop on Peer-to-Peer Systems (IPTPS'06)*, 2006.
- [18] Voulgaris S. and van Steen M., Epidemic-style Management of Semantic Overlays for Content-Based Searching. *Proc. Int'l Conf. on Parallel and Distributed Computing (Euro-Par'05)*, Springer-Verlag LNCS, 2005.
- [19] Watts D.J. and Strogatz S.H., Collective Dynamics of Small-World Networks. *Nature*, 393: 440-442, 1998.
- [20] Yao A.C.C., On Constructing Minimum Spanning Trees in k -dimensional Space and Related Problems. *SIAM Journal of Computing*, 11:721-736, 1982.
- [21] <http://peersim.sourceforge.net/>

Notation

- n is the total number of nodes.
- d and $d()$ denote distances.
- p denotes the number of local contacts of a peer.
- r denotes the radius within which are located the local contacts.
- q denotes the number of shortcuts of a peer.

A Determination of $d_1()$

Let us recall that $d_1(i)$ denote the density of probability that the random shortcut is at distance i from the considered node. Since the selection is done by random, we just have to evaluate the function that gives the proportion of nodes at distance less than i from a source. d_1 is simply the differentiate of this function.

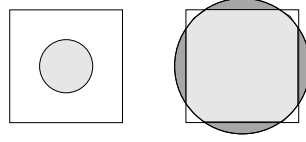


Figure 6: Determination of the probability density $d_1(i)$

For $0 \leq i \leq \frac{1}{2}$ the proportion of nodes at distance less than i is easy to calculate. Indeed it corresponds to the average number of peers included in the disk of radius i divided by the total number of peers. As represented on the left side of Figure 6, we have to determine the average number of peers in the grey zone. It is easy to see that this number of peers is $\frac{\pi i^2}{n}$. Consequently, for $0 \leq i \leq \frac{1}{2}$, we have $d_1(i) = 2\pi i$.

For $\frac{1}{2} < i \leq \frac{\sqrt{2}}{2}$, the calculation is a little bit more involved. This is because, as shown in the right part of Figure 6, the disk goes out of the torus. We have consequently to evaluate the average number of peers included only in the light grey zone. We determine that number by subtracting the area of the dark grey zones from the area of the disk. More precisely, the light grey area $S(i)$ is equal to:

$$\begin{aligned}
 S(i) &= \pi i^2 - 4 \int_{x=\frac{1}{2}}^i \int_{y=-\sqrt{i^2-x^2}}^{\sqrt{i^2-x^2}} 1, \\
 &= \pi i^2 - 8 \int_{x=\frac{1}{2}}^i \int_{y=0}^{\sqrt{i^2-x^2}} 1, \\
 &= \pi i^2 - 8 \int_{x=\frac{1}{2}}^i \sqrt{i^2-x^2}, \\
 &= \pi i^2 - 2\pi i^2 + \sqrt{4i^2-1} + 4i^2 \arcsin \frac{1}{2i}, \\
 &= -\pi i^2 + \sqrt{4i^2-1} + 4i^2 \arcsin \frac{1}{2i}.
 \end{aligned}$$

With a differentiation, we obtain the result, namely, $d_1(i) = -2\pi i + 8i \arcsin(\frac{1}{2i})$, for $\frac{1}{2} < i \leq \frac{\sqrt{2}}{2}$.

B Creating Kleinberg's shortcuts in a gossip-based protocol

Both Cyclon and the proposed protocol follow the same pattern. They differ only in the way a peer selects the set of peers it sends to another peer during a round (steps denoted (2) and (3) in the following). While these peers are randomly selected in Cyclon, they are selected according to their distance in the proposed protocol.

The behavior of each protocol is explained through the following example. Let A and B be two peers whose views (of size $c = 6$) are $v_A = \{B, C, D, E, F, G\}$ and $v_B = \{U, V, W, X, Y, Z\}$, respectively.

A round in Cyclon

- 1 A randomly selects a peer from its view (say B).
- 2 A randomly selects $\frac{c}{2} - 1$ other peers from its view, say C and E , and sends to B the set $A \text{ to } B = \{B, C, E\}$.
- 3 When B receives the set $A \text{ to } B$ from A , it randomly selects $\frac{c}{2}$ peers from its view (e.g., the set $B \text{ to } A = \{U, W, Z\}$), and sends it to A . It executes $v_B \leftarrow (v_B \setminus B \text{ to } A) \cup A \text{ to } B$ to obtain its new view. So, the view of B is now $v_B = \{A, V, C, X, Y, E\}$.
- 4 Finally, when A receives the set $B \text{ to } A$ from B , it executes $v_A \leftarrow (v_A \setminus A \text{ to } B) \cup B \text{ to } A$ to obtain its new view, that becomes $v_A = \{U, W, D, Z, F, G\}$.

A round in the proposed peer sampling protocol

- 1 A randomly selects a peer from its view (say B).
- 2 Among the other peers in A 's view ($v'_A = v_A \setminus \{B\} = \{C, D, E, F, G\}$), A chooses $\frac{c}{2}$ peers to be kept in its view and sends the other peers to B . The peer selection is done according to Kleinberg's distribution. A computes the value $\frac{1}{d(A,x)^2}$ for each peer $x \in v'_A$. Then A keeps a peer in its view according to a probability proportional to the previous value. More precisely, a peer x is chosen with probability

$$\frac{\frac{1}{d(A,x)^2}}{\sum_{y \in v'_A} \frac{1}{d(A,y)^2}}.$$

For example A keeps the peers $\{D, E, F\}$ and then sends to B the set $A_to_B = \{B, C, G\}$.

- 3 When B received the set A_to_B from A , it keeps $\frac{c}{2}$ peers in its view (according to Kleinberg's distribution), e.g., the set $\{U, V, Z\}$, and sends the remaining peers to A , i.e., the set $B_to_A = \{W, X, Y\}$ to A . It finally executes $v_B \leftarrow (v_B \setminus B_to_A) \cup A_to_B$ to update its view, that becomes $v_B = \{U, V, B, C, G, Z\}$.
- 4 Finally, when A receives the set B_to_A from B , it executes $v_A \leftarrow (v_A \setminus A_to_B) \cup B_to_A$ to update its view, that becomes $v_A = \{W, X, D, E, F, Y\}$.