

Networks Become Navigable as Nodes Move and Forget

Augustin Chaintreau^{1*}, Pierre Fraigniaud^{2**}, and Emmanuelle Lebhar^{2***}

¹ Thomson, Paris, France

² CNRS and University Paris Diderot, France

Abstract. We propose a dynamic process for network evolution, aiming at explaining the emergence of the small world phenomenon, i.e., the statistical observation that any pair of individuals are linked by a short chain of acquaintances computable by a simple decentralized routing algorithm, known as greedy routing. Our model is based on the combination of two dynamics: a random walk (spatial) process, and an harmonic forgetting (temporal) process. Both processes reflect natural behaviors of the individuals, viewed as nodes in the network of inter-individual acquaintances. We prove that, in k -dimensional lattices, the combination of these two processes generates long-range links mutually independently distributed as a k -harmonic distribution. We analyze the performances of greedy routing at the stationary regime of our process, and prove that the expected number of steps for routing from any source to any target in any multidimensional lattice is a polylogarithmic function of the distance between the two nodes in the lattice. Up to our knowledge, these results are the first formal proof that navigability in small worlds can emerge from a dynamic process for network evolution. Our dynamica process can find practical applications to the design of spatial gossip and resource location protocols.

Keywords: Small world phenomenon, dynamic process, routing, spatial gossip, resource location, random walks.

1 Introduction

Models relating geography and social-network friendship enable a good understanding of the small world phenomenon, a.k.a., six degrees of separation between individuals [13, 31]. In these models, the probability of befriending a particular person is assumed to be inversely proportional to the number of closer people, fitting with what was observed experimentally (cf. [30]). Under this assumption, it was proved that, using ad hoc probability distributions, many classes of graphs are navigable, that is, a simple decentralized routing procedure enables

* Email: Augustin.Chaintreau@thomson.net.

** Email: Pierre.Fraigniaud@liafa.jussieu.fr. Additional supports from the ANR projects ALADDIN and ALPAGE, and from the COST Action 295 DYNAMO.

*** Email: Emmanuelle.Lebhar@liafa.jussieu.fr. Additional supports from the ANR project ALADDIN, and from the COST Action 295 DYNAMO.

efficient routing from any source to any target. (By efficient, we mean, as it is standard in this framework, that routing from any source s to any target t takes a polylogarithmic expected number of steps). For instance, such a navigability property is satisfied in multi-dimensional meshes [26], in graphs of bounded ball growth [15], and more generally in graphs of bounded doubling dimension [35]. In all these cases, a graph G , that may not only represent geography but also other proximity measures like professional activities, religious beliefs, etc., is enhanced with additional links chosen at random. More precisely, every node is given some *long-range links* pointing at other nodes in the graph. For each long-range link added at a node u , the probability that the head of this link is v is inversely proportional to the size of the ball of radius $\text{dist}_G(u, v)$ centered at u in G , hence depending on the density of G around u . This setting applies to weighted graphs too [28], and to infinite graphs as well [15]. For instance, in the k -dimensional lattice \mathbb{Z}^k , the probability that u has a long-range link pointing at v is essentially proportional to $1/d^k$ where d is the distance between u and v in the lattice. This choice of the long-range links enables greedy routing³ to perform in polylogarithmic expected number of steps (as a function of the distance in the lattice between the source and the target), and it was shown to be necessary.

1.1 Navigability as an emerging property

In [27] (Problem 7), Jon Kleinberg asks about "what kinds of growth processes or selective pressures might exist to cause networks to become more efficiently searchable". Many attempts have been made to explain how the density-based distribution of the long-range links can emerge with time from the evolution of a network. Inspired by the world wide web or by P2P file-sharing systems, all the models we are aware of have considered the augmentation process (or rewiring) of a static graph used by its nodes for searching information. Our work uses a different approach, starting from the following observations. On the one hand, individuals usually communicate with people they have met in the past. We thus assume that long-range connections are between remote people who have met once in the past. In other words, long-range links are emerging from nodes mobility, that we model by random walks in this paper. On the other hand people tend to forget some of their former acquaintance with time. This forgetting mechanism represents the well understood fact that one cannot maintain close relationships with an explosive number of people. Thus we couple the random walk process with a forgetting process, and prove that this ideal setting is sufficient to insure polylogarithmic navigability with simply one long-range connection per node.

³ Greedy routing [26] aims at modeling the routing strategy performed by the individuals in Milgram experiment. In a graph G enhanced with long-range links, a node u handling a message of destination t selects among all its neighbors, including its long-range contact(s), the one that is the closest to the target t according to the distance in the base graph G , and forwards the message to that node.

The main contribution of this paper is therefore to show that navigability emerges from the combination of two separate processes: a spatial process and a temporal process. Our model also provides a setting for the analysis of dynamic social networks (see e.g. [7] for an example of application to spatial gossip).

1.2 Rewiring processes

Clauset and Moore [11] proposed the following rewiring process for the multidimensional lattice, inspired by the actions of surfers on the web. While routing from a source s to a target t , if the target is not reached after τ steps, then the long-range link of s is rewired to point at the current node x . The threshold τ is set based on the distance (in the lattice) between s and t , and on the expected time of greedy routing from s to t when the k -dimensional lattice is augmented using the k -harmonic distribution [26]. The simulation results presented in [11] show that the distribution f of the link lengths converges to the k -harmonic distribution. Sandberg and Clarke [33] proposed a different rewiring process, based on Freenet feedback mechanisms [10]. This iterative process selects, at each phase, two nodes s and t uniformly at random, and constructs the greedy path $s = x_0, x_1, \dots, x_{k-1}, x_k = t$ from s to t . For every $i \in \{0, 1, \dots, k\}$, the long-range link of x_i is rewired with probability p , to point at t . The $k + 1$ decisions (rewiring or not) are taken mutually independently. This process is analyzed in [32]. It is proved that, under some hypotheses, the process converges. Moreover, the stationary distribution f of the link lengths can be fully characterized. In the k -dimensional lattice, it is close to the k -harmonic distribution when p is set appropriately, and simulations show that greedy routing in rings and meshes enhanced using the stationary distribution f performs as efficiently as when these networks are enhanced using the 1- and 2-harmonic distributions, respectively.

For both [11] and [33], the complete formal analysis of the process remains open (even the formal characterization of the stationary distribution of the processes described in [11] remains open). The difficulty of the analysis is due to the dependencies between the long-range links generated by the processes. In particular, the computation of the greedy routing performances is a challenge when the long-range links are not mutually independent. So, building further theory upon these two models looks quite difficult.

In this paper, we propose a dynamic network model based on the combination of two simple processes: a random walk process, and a harmonic forgetting process. We prove that the combination of these two processes generates long-range links mutually independently distributed as a distribution that resembles the density-based distribution, and from which navigability provably emerges.

1.3 Sketch of our network evolution process

In our network evolution process, called move-and-forget, or M&F for short, individuals are modeled by tokens moving from node to node in the k -dimensional lattice \mathbb{Z}^k , for some fixed integer $k \geq 1$ (the dimension of the lattice may be

related to the number of proximity criteria used by the individuals for routing). Initially, each node is occupied by exactly one token. These tokens are moved mutually independently during the execution of the dynamic process, according to a random walk.

Tokens are attached to the heads of the long-range links, whose tails are the nodes from where the tokens initially started their random walks. Using the analogy of individuals moving in the geographical world, each long-range link indicates an acquaintance between an individual located at a fixed geographical point (where the token initially stood) and some individual located at some geographical coordinates (where the token currently stands).

The random walk process is coupled with another dynamic: nodes may forget their contacts through their long-range links. The motivation for our forgetting process is that individuals may loose touch with former friends, but they meet new people among which some may become friends. Since older acquaintances indicate stronger relationships, they have less probability to be forgotten than recent ones. We assume that whenever an acquaintance has been known for twice longer, she or he is twice less likely to be forgotten now. Therefore a long-range link of age a , that is a long-range link that survived a steps of the forgetting process, is forgotten with probability $\phi(a) \propto 1/a$. When a long-range link is forgotten by a node, it is rewired to point at this node (hence creating a self-loop). The token at the head of the forgotten link is removed, and a new token is launched at the node. (A new local relationship replaces an old remote relationship).

Note that M&F is defined independently from the dimension k of the lattice: tokens execute random walks, and they are forgotten with a probability that depends only of their ages.

1.4 Our results

We prove that, for any fixed integer $k \geq 1$, the M&F rewiring process sketched above converges in the k -dimensional lattice to a stationary regime where the distribution f of the link lengths that resembles the k -harmonic distribution. Precisely, we prove that there exists $d_0 \geq 0$ and two positive constants c and c' , such that, for any $\mathbf{d} = (d_1, \dots, d_k) \in \mathbb{Z}^k$ with $|d_i| \geq d_0$ for all $i \in \{1, \dots, k\}$, we have

$$\frac{c}{\|\mathbf{d}\|^k \cdot \ln^{1+\varepsilon} \|\mathbf{d}\|} \leq f(\mathbf{d}) \leq \frac{c' \ln^{k/2} \|\mathbf{d}\|}{\|\mathbf{d}\|^k \cdot \ln^{1+\varepsilon} \|\mathbf{d}\|}$$

where $\varepsilon > 0$ is a fixed (arbitrary small) parameter of M&F, and $\|\cdot\|$ denotes the ℓ_∞ norm.

Moreover, M&F guarantees the mutual independence of the long-range links. As a consequence, the performances of greedy routing in the lattice enhanced using the distribution f can be analyzed formally. We prove that the expected number of steps of greedy routing from any source s to any target t at distance d in the k -dimensional lattice satisfies

$$\mathbb{E}[X_{s,t}] \leq O(\ln^{2+\varepsilon} d).$$

	Convergence	Navigability
A. Clauset and C. Moore (2003)	Simulations	Simulations
O. Sandberg and I. Clarke (2007)	Proof	Simulations
Move-and-forget (M&F)	Proof	Proof

Table 1. Properties of known network evolution processes compared to M&F

Therefore, greedy routing performs polylogarithmically as a function of the distance between the source and the target. In particular, the performances of greedy routing are essentially the same as the ones obtained by Kleinberg [26] using the ad hoc k -harmonic distribution [26].

Up to our knowledge, these results are the first formal proof that navigability in small worlds can emerge from a dynamic process for network evolution (see Table 1). Moreover, M&F is simple (by just coupling two simple dynamics), naturally distributed (each node takes care of just its token), robust (the loss of one token simply requires to launch a new token), and scalable (by direct adaptations of the infinite lattice setting to square toroidal meshes of arbitrary sizes).

Last but not least, M&F can find practical applications, including the design of distributed spatial gossip and resource location protocols.

The omitted proofs of the results in this extended abstract can be found in [8].

1.5 Related works

The search for a network evolution process that could explain the emergence of the small world phenomenon in social networks started with the pioneering work of Watts and Strogatz [36] who proposed a rewiring process in the cycle, generating networks possessing small diameter and large clustering coefficient, simultaneously. Adding random matchings to cycles, as in [5], yields graphs with small diameter, but non necessarily with large clustering coefficient. As far as navigability is concerned, these networks do not support efficient decentralized routing mechanisms [26]. Preferential attachment model [2, 34] enables the design of efficient search procedures under specific circumstances (see [18] and the references therein), the recent lower bounds in [14] show that polylogarithmic routing cannot be achieved in general in such networks. As far as we know, the only network evolution models from which polylogarithmic routing emerges are the aforementioned ones [11, 33].

Following up the seminal work of Kleinberg [26], a large literature has been dedicated to the analysis of greedy routing in graphs enhanced by long-range links set according to various kinds of probability distributions (see, e.g., [1, 15–17, 35]). These papers proved that several large classes of graphs can be enhanced by long-range links so that greedy routing performs in polylogarithmic expected number of steps. A lower bound of $\Omega(n^{1/\sqrt{\log n}})$ expected number of steps for greedy routing in arbitrary graphs has been proved in [20], and an upper bound of $O(n^{1/3})$ has been proved in [19]. Lower bounds for the cycle can be found in [3, 4, 21].




Fig. 1. Dynamic of the long-range links in M&F.

2 The Move-and-Forget (M&F) Rewiring Process

2.1 Process description

Random walks Let $k \geq 1$ be an integer. The rewiring process move-and-forget (M&F for short) assumes that each node in the k -dimensional lattice \mathbb{Z}^k is initially occupied by exactly one token. These tokens move mutually independently according to random walks. That is, each token is given a set of k fair coins c_i , $i = 1, \dots, k$. At each step of its walk, each token flips its k coins, and moves in the i th dimension of the lattice in the positive direction if c_i is head, and in the negative direction if it is tail. More precisely, let $X(t) \in \mathbb{Z}^k$ denotes the position of a token in the lattice after t steps of M&F, assuming that the token initially started at node $(u_1, \dots, u_k) \in \mathbb{Z}^k$. We have $X(0) = (u_1, \dots, u_k)$, and, for $t \geq 1$, $X(t) = (X_1(t), \dots, X_k(t))$ satisfies

$$X_i(t) = \begin{cases} X_i(t-1) + 1 & \text{with probability } 1/2; \\ X_i(t-1) - 1 & \text{with probability } 1/2. \end{cases} \quad (1)$$

Setting of the long-range links Tokens are attached to the heads of the long-range links, whose tails are the nodes from where the tokens initially started their random walks (see Figure 1(a)). The head of a long-range link is called the long-range *contact* of the tail of this link. Hence the long-range contact of a node u is the node v currently occupied by the token launched by node u .

Forgetting process Nodes may forget their contacts through their long-range links. More precisely, a long-range link of age $a \geq 0$, that is a long-range link that survived a steps of the forgetting process, is forgotten with probability $\phi(a)$. When a long-range link is forgotten by a node, it is rewired to point at this node (see Figure 1(b)). The token at the head of the forgotten link is removed, and a new token is launched at the node. This new token starts another random walk in \mathbb{Z}^k . Hence, if $A(t) \in \mathbb{N}$ denotes the age of the long-range link of some node u , that is the number of steps between time t and the last time this link was rewired during the execution of M&F, and if $C(t)$ denotes the long-range contact of node u at step t , then we have $C(t) = X(A(t))$.

The forgetting function ϕ has a huge impact on the distribution of the long-range link lengths. In this paper, we will consider $\phi(a) \propto 1/a$. The precise setting of ϕ will appear more complex for technical reasons only⁴ (series convergence for infinite lattices, normalization, etc.). In fact, its behavior essentially reflects a decreasing of the forgetting probability that is inversely proportional to the age of the relationships. The precise setting of ϕ is described in the next section which explains the connections between the random walk X , the forgetting function ϕ , and the distribution f of the long-range link lengths.

2.2 Setting of the forgetting function

We first prove that the age of the long-range link resulting from the execution of M&F at a node has a stationary distribution.

Lemma 1. *For any function ϕ in $[0, 1)$ such that the series of general term $\prod_{i=1}^j (1 - \phi(i))$ is finite, $(A(t))_{t \geq 0}$ is a Markov chain which is irreducible, aperiodic, and positive recurrent, with stationary probability distribution π where*

$$\pi(a) = \frac{\prod_{i=1}^a (1 - \phi(i))}{\sum_{j \geq 0} \prod_{i=1}^j (1 - \phi(i))},$$

for all $a \geq 0$.

Definition 1. *We define the forgetting probability ϕ as the following function:*

$$\phi(a) = \begin{cases} 0 & \text{if } a = 0, 1, \text{ or } 2; \\ 1 - \frac{a-1}{a} \left(\frac{\ln(a-1)}{\ln a} \right)^{1+\epsilon} & \text{if } a \geq 3; \end{cases} \quad (2)$$

where $\epsilon > 0$ is arbitrary small.

Note that $\phi(a) = \frac{1}{a} + o\left(\frac{1}{a}\right)$. Indeed,

$$\left(\frac{\ln(a-1)}{\ln a} \right)^{1+\epsilon} = \left(1 + \frac{\ln(1-1/a)}{\ln a} \right)^{1+\epsilon} = 1 - \frac{1+\epsilon}{a \ln a} + o\left(\frac{1}{a \ln a}\right)$$

If ϕ is defined according to Eq. (2), then Lemma 1 enables to give a close formula for π .

Lemma 2. *If ϕ is defined according to Eq. 2, then there exists a constant $c > 0$ such that $\pi(0) = \pi(1) = \pi(2) = c$ and for any $a \geq 3$,*

$$\pi(a) = \frac{c}{a \ln^{1+\epsilon} a}.$$

⁴ For instance, one needs $\sum_{a \geq 0} \phi(a)$ to diverge since otherwise the Markov chain $A(t)$ would be transient, and links could survive infinitely with positive probability. However, on the one hand, just setting $\phi(a) = 1/a$ would make $A(t)$ recurrent null (and thus for any a we would have $\Pr\{A(t) = a\}$ converging to 0 as t goes to infinity), but, on the other hand, setting $\phi(a) = 1/a^\alpha$ with $\alpha < 1$ would not yield navigability.

Finally, the relationship between the stationary distribution of the long-range link ages and the stationary distribution of the long-range link lengths is made explicit in the following lemma.

Lemma 3. *The distribution of the long-range links converges to the distribution f satisfying, for any $\mathbf{d} \in \mathbb{Z}^k$,*

$$f(\mathbf{d}) = \sum_{a \geq 0} \pi(a) \cdot \Pr\{X(a) = \mathbf{d}\}.$$

3 Analysis of the dynamic process M&F

In this section, we analyze the stationary distribution of the long-range link lengths in the k -dimensional lattice, and prove that this distribution resembles the k -harmonic distribution.

Theorem 1. *There exist $d_0 \geq 0$ and two positive constants c and c' such that, for any $\mathbf{d} = (d_1, \dots, d_k) \in \mathbb{Z}^k$ with $|d_i| \geq d_0$ for all $i \in \{1, \dots, k\}$, we have*

$$\frac{c}{\|\mathbf{d}\|^k \cdot \ln^{1+\epsilon} \|\mathbf{d}\|} \leq f(\mathbf{d}) \leq \frac{c' \ln^{k/2} \|\mathbf{d}\|}{\|\mathbf{d}\|^k \cdot \ln^{1+\epsilon} \|\mathbf{d}\|}$$

where $\epsilon > 0$ is the fixed parameter of M&F, and $\|\cdot\|$ denotes the ℓ_∞ norm.

To prove the theorem, we first prove that, for large distances d , a random walk of age a cannot be of length d unless $a \geq \Omega(d^2)$. More precisely, we establish an exponentially small upper bound for the probability for a long-range link to be of length d at age $a = o(d^2)$. Second, we prove that if the age a is sufficiently large, then the chance for a random walk to reach a given distance d at age a is proportional to $\frac{1}{\sqrt{a}}$. Summing this probability over all values of a larger than d^2 allows us to conclude that the transform of the age distribution π described in Lemma 3 is approaching the k -harmonic distribution.

We compute an estimation of $\Pr\{X(a) = d\}$ when a is sufficiently large. We will use the following asymptotic equivalent of the binomial coefficient, that can be derived by application of the Stirling formula. Let n_i and m_i be two sequences of positive integers such that $n_i \rightarrow \infty$, $m_i \rightarrow \infty$, and $n_i - m_i \rightarrow \infty$ when i grows to infinity. Then, as i grows to infinity

$$\binom{n_i}{m_i} \sim \frac{1}{\sqrt{2\pi}} \cdot \sqrt{\frac{n_i}{m_i \cdot (n_i - m_i)}} \cdot \frac{n_i^{n_i}}{m_i^{m_i} \cdot (n_i - m_i)^{n_i - m_i}}. \quad (3)$$

Lemma 4. *Let X be a random walk in \mathbb{Z} . For any $C > 0$, $\zeta > 0$, there exists $d_0 > 1$ such that, for any $|d| \geq d_0$ and $a \geq \frac{d^2}{C \cdot \ln |d|}$, we have*

$$(1 - \zeta) \cdot \sqrt{\frac{2}{\pi \cdot a}} \exp\left(-\frac{3d^2}{4a}\right) \leq \Pr\{X(a) = d\} \leq (1 + \zeta) \cdot \sqrt{\frac{2}{\pi \cdot a}} \exp\left(-\frac{d^2}{4a}\right).$$

We are now ready to prove of the lower bound of Theorem 1. For the sake of simplicity, let us first assume that the dimension of the lattice is 1. In this case, one can apply the results from the previous section directly. For any $a \geq \frac{3}{4}d^2$ we have

$$\exp\left(-\frac{3d^2}{4a}\right) \geq 1/e.$$

Therefore, for any $\zeta > 0$, there exists d_0 large enough and $a \geq \frac{3}{4}d^2$ such that Lemma 4 yields:

$$\Pr\{X(a) = d\} \geq \frac{1-\zeta}{e} \sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{a}}.$$

Thus :

$$f(d) = \sum_{a \geq 0} \Pr\{X(a) = d\} \pi(a) \geq \frac{1-\zeta}{e} \sqrt{\frac{2}{\pi}} \sum_{a \geq \frac{3}{4}d^2} \frac{1}{a^{3/2} \cdot \ln^{1+\epsilon}(a)}.$$

More generally, in the k -dimensional lattice, let us denote the position of the random walk by $X(a) = (X_1(a), \dots, X_k(a))$. From the setting of M&F, each X_i is an unbiased random walk in dimension 1, and the X_i s are mutually independent. We can thus apply all the results from the previous section independently for each coordinate of $\mathbf{d} = (d_1, \dots, d_k)$. Assuming that

$$|d_i| \geq d_0 \text{ for all } i \in \{1, \dots, k\},$$

we can apply Lemma 4 to every dimension. We get:

$$a \geq \frac{3}{4} \|\mathbf{d}\|^2 \implies \forall i \in \{1, \dots, k\}, \Pr\{X_i(a) = d_i\} \geq \frac{1-\zeta}{e} \sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{a}} \exp\left(-\frac{3d_i^2}{4a}\right).$$

For $a \geq \frac{3}{4} \|\mathbf{d}\|^2$, we have,

$$\frac{3d_i^2}{4a} \leq \frac{3\|\mathbf{d}\|^2}{4a} \leq 1 \text{ and thus } \exp\left(-\frac{3d_i^2}{4a}\right) \geq 1/e.$$

As a consequence, by Lemma 4,

$$\Pr\{X(a) = \mathbf{d}\} = \Pr\{X_1(a) = d_1, \dots, X_k(a) = d_k\} \geq \left(\frac{1-\zeta}{e} \sqrt{\frac{2}{\pi}} \frac{1}{\sqrt{a}}\right)^k$$

hence

$$f(\mathbf{d}) = \sum_{a \geq 0} \Pr\{X(a) = \mathbf{d}\} \pi_A(a) \geq \left(\frac{1-\zeta}{e} \sqrt{\frac{2}{\pi}}\right)^k \sum_{a \geq \frac{3}{4}\|\mathbf{d}\|^2} \frac{c}{a^{1+(k/2)} \cdot \ln^{1+\epsilon}(a)}.$$

The lower bound is then a direct consequence of the following result with $N = \frac{3\|\mathbf{d}\|^2}{4}$.

Lemma 5. *For any $\epsilon > 0$, and any $N \geq e^{2(1+\epsilon)}$, we have*

$$\frac{2/(k+1)}{N^{k/2} \ln^{1+\epsilon} N} \leq \sum_{a \geq N} \frac{1}{a^{1+(k/2)} \ln^{1+\epsilon} a} \leq \frac{2/k}{(N-1)^{k/2} \ln^{1+\epsilon}(N-1)}. \quad (4)$$

Due to lack of space, the proof of the upper bound of Theorem 1 is omitted. It uses similar arguments for large values of a and uses a Chernoff bound for small values of a , given by the following lemma.

Lemma 6. *Let X be a random walk in \mathbb{Z} . Then, for any age $a > 0$ and any distance $d \in \mathbb{Z}$, we have $\Pr\{X(a) = d\} \leq 2 \cdot \exp\left(-\frac{d^2}{32 \cdot a}\right)$.*

4 Applications

In the previous section, we have shown that the distribution f of the long-range link lengths is provably converging to a distribution that resembles the k -harmonic distribution. In this section, we show that greedy routing can be formally analyzed at the stationary state of this distribution. Greedy routing can be formally analyzed for two reasons: (1) The distribution f of the long-range links constructed by M&F can be bounded formally (cf. Theorem 1); (2) The long-range links resulting from M&F are mutually independent. Based on these two facts, we can establish the theorem below.

Theorem 2. *In the k -dimensional lattice augmented with the long-range links at the stationary distribution of the dynamic process M&F, the expected number of steps of greedy routing from any source node s to any target node t at distance d is $O(\ln^{2+\epsilon} d)$.*

In the rest of the section, we discuss how M&F can find practical applications to the design of spatial gossip and resource location protocols.

Gossip-based protocols, a.k.a., epidemic algorithms [12], have been introduced as a methodology for designing robust and scalable communication schemes in distributed systems. Roughly, in each step, each node u chooses some other node v , and sends a message to it. By applying such scheme at each node, an information originated at some source s will eventually reach its target(s). This methodology can be adapted to various problems, including information spreading, resource location, etc. In [24], Kempe et al. introduced *spatial gossip*, which allowed them to derive efficient solutions for many communication problems. In spatial gossip, nodes are arranged with uniform density in the k -dimensional Euclidean space, and, at each step of the gossip protocol, node u chooses node v with probability $\propto 1/d^{\rho k}$ where $\rho > 0$ is a fixed parameter, and d is the distance between u and v . In particular, it is shown that, for $\rho \in (1, 2)$, spatial gossip enables to propagate information at distance d in time polylogarithmic in d . In [25], Kempe and Kleinberg showed that spatial gossip enables to solve larger classes of problems, including MST construction and permutation routing. In

particular, they prove that permutation routing using spatial gossip with $\varrho = 1$ performs in polylogarithmic expected number of steps.

In [7], we show how the M&F process could facilitate the implementation of the protocols in [24, 25] for networks that take advantage of node mobility, as in, e.g., [9, 22, 23].

References

1. I. Abraham and C. Gavoille. Object Location Using Path Separators. 25th ACM Symp. on Principles of Distributed Computing (PODC), pages 188-197, 2006.
2. R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Review of modern physics* 74:47-97, 2002.
3. J. Aspnes, Z. Diamadi, and G. Shah. Fault-tolerant routing in peer-to-peer systems. In 21st ACM Symp. on Principles of Distributed Computing (PODC), pages 223-232, 2002.
4. L. Barrière, P. Fraigniaud, E. Kranakis, and D. Krizanc. Efficient Routing in Networks with Long Range Contacts. In 15th International Symposium on Distributed Computing (DISC), LNCS 2180, pages 270-284, Springer, 2001.
5. B. Bollobás and F. Chung. The diameter of a cycle plus a random matching. *SIAM J. Discrete Math.* 1(3):328-333, 1988.
6. P. Bremaud. *Markov Chains, Gibbs Field, Monte Carlo Simulation and Queues.* Springer-Verlag, 1999.
7. A. Chaintreau, P. Fraigniaud and E. Lebhar. Opportunistic spatial gossip over mobile social networks. In 1st ACM SIGCOMM Workshop on Online Social Net. (WOSN), to appear, 2008.
8. A. Chaintreau, P. Fraigniaud and E. Lebhar. Networks Become Navigable as Nodes Move and Forget. Technical Report, arXiv:0803.0248v1, 2008.
9. A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, J. Scott and R. Gass. Impact of Human Mobility on Opportunistic Forwarding Algorithms. *IEEE Trans. Mob. Comp.* 6(6):606-620, 2007.
10. I. Clarke, O. Sandberg, B. Wiley and T. W. Hong. Freenet: A Distributed Anonymous Information Storage and Retrieval System. In Workshop on Design Issues in Anonymity and Unobservability, LNCS 2009, pages 46-66, 2001.
11. A. Clauset and C. Moore. How Do Networks Become Navigable? Technical Report, arXiv:0309.415v2, 2003.
12. A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic Algorithms for Replicated Database Maintenance. *Operating Systems Review* 22(1):8-32, 1988.
13. P. Dodds, R. Muhamad, and D. Watts. An experimental study of search in global social networks. *Science* 301, no5634, pages 827-829, 2003.
14. P. Duchon, N. Eggeman, and N. Hanusse. Non-Searchability of Random Power Law Graphs. 11th Int. Conference on Principles of Distributed Systems (OPODIS), Springer, LNCS 4878, pages, 2007.
15. P. Duchon, N. Hanusse, E. Lebhar, and N. Schabanel. Could any graph be turned into a small-world? *Theoretical Computer Science* 355(1):96-103 (2006).
16. M. Flammini, L. Moscardelli, A. Navarra, and S. Perennes. Asymptotically optimal solutions for small world graphs. In 19th International Symposium on Distributed Computing (DISC), LNCS 3724, pages 414-428, 2005.

17. P. Fraigniaud. Greedy routing in tree-decomposed graphs: a new perspective on the small-world phenomenon. In 13th Annual European Symposium on Algorithms (ESA), pages 791–802, 2005.
18. P. Fraigniaud, P. Gauron, and M. Latapy. Combining the use of clustering and scale-free nature of user exchanges into a simple and efficient P2P system. In 11th European Conference on Parallel Computing (EuroPar), Springer, LNCS 3648, pages 1163–1172, 2005.
19. P. Fraigniaud, C. Gavoille, A. Kosowski, E. Lebhar, and Z. Lotker. Universal Augmentation Schemes for Network Navigability: Overcoming the \sqrt{n} -Barrier. In 19th ACM Symp. on Parallelism in Algorithms and Architectures (SPAA), pages 1-9, 2007.
20. P. Fraigniaud, E. Lebhar, and Z. Lotker. A doubling dimension threshold $\Theta(\log \log n)$ for augmented graph navigability. In 14th European Symp. on Algorithm (ESA), LNCS 4168, 376-386, 2006.
21. G. Giakkoupis and V. Hadzilacos. On the complexity of greedy routing in ring-based peer-to-peer networks. In 26th ACM Symp. on Princ. of Dist. Comp. (PODC), 2007.
22. M. Grossglauser and D. Tse. Mobility Increases the Capacity of Ad Hoc Wireless Networks. *IEEE/ACM Trans. on Net.* 10(4):477-486, 2002.
23. S. Jain, K. Fall and R. Patra. Routing in a delay tolerant network. In Proc. ACM SIGCOMM, 2004.
24. D. Kempe, J. Kleinberg, and A. Demers. Spatial gossip and resource location protocols”, In 33rd ACM Symposium on Theory of Computing, 163-172, 2001.
25. D. Kempe and J. Kleinberg. Protocols and impossibility results for gossip-based communication mechanisms. In 43st IEEE Symp. on Foundations of Computer Science, 471-480, 2002.
26. J. Kleinberg. The small-world phenomenon: an algorithmic perspective. In 32nd ACM Symp. on Theory of Computing (STOC), pages 163-170, 2000.
27. J. Kleinberg. Complex networks and decentralized search algorithm. Nevanlinna prize presentation at the International Congress of Mathematicians (ICM), Madrid, 2006.
28. R. Kumar, D. Liben-Nowell, and A. Tomkins. Navigating Low-Dimensional and Hierarchical Population Networks. In 14th European Symposium on Algorithm (ESA), LNCS 4168, 2006.
29. D. Liben-Nowell and J. Kleinberg. The link prediction problem for social networks. *Journal of the American society for information science and technology*, 58(7):1019-1031, 2007.
30. D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, and A. Tomkins Geographic routing in social networks *Proc. of the Natl. Academy of Sciences of the USA*, vol. 102/3, pages 11623–11628, 2005.
31. S. Milgram. The Small-World Problem. *Psychology Today*, pages 60-67, 1967.
32. O. Sandberg. Neighbor Selection and Hitting Probability in Small-World Graphs. *Annals of Applied Probability* (to appear), 2008.
33. O. Sandberg and I. Clarke. The evolution of navigable small-world networks. *Tech. Report 2007:14*, Chalmers University of Technology, 2007.
34. H. Simon. On a class of skew distribution functions. *Biometrika* 42(3/4):425-440, 1955.
35. A. Slivkins. Distance estimation and object location via rings of neighbors. In 24th Annual ACM Symp. on Princ. of Dist. Comp. (PODC), pages 41–50, 2005.
36. D. Watts, and S. Strogatz. Collective Dynamics of Small-World Networks. *Nature* 393:440-442, 1998.