

A Parsimonious Model of Mobile Partitioned Networks with Clustering

Michał Piórkowski
School of Computer and
Communication Sciences, EPFL
1015 Lausanne, Switzerland

Natasa Sarafijanovic-Djukic
School of Computer and
Communication Sciences, EPFL
1015 Lausanne, Switzerland

Matthias Grossglauser
Internet Laboratory
Nokia Research Center
00180 Helsinki, Finland

Abstract—Mobile wireless networks frequently possess, at the same time, both dense and sparse regions of connectivity; for example, due to a heterogeneous node distribution or radio propagation environment. This paper is about modeling both the mobility and the formation of clusters in such networks, where nodes are concentrated in clusters of dense connectivity, interspersed with sparse connectivity. Uniformly dense and sparse networks have been extensively studied in the past, but not much attention has been devoted to clustered networks.

We present a new mobility model for clustered networks, which is important for the design and evaluation of routing protocols. We refer to our model as Heterogeneous Random Walk (HRW). This model is simple, mathematically tractable, and it captures the phenomenon of emerging clusters, observed in real partitioned networks. We provide a closed-form expression for the stationary distribution of node position and we give a method for “perfect simulation”.

Moreover, we provide evidence, based on mobility traces, for the main macroscopic characteristics of clustered networks captured by the proposed mobility model. In particular, we show that in some scenarios, nodes have statistically very similar mobility patterns. Also, we discuss cluster dynamics and the relationship between node speed and node density.

I. INTRODUCTION

Most routing protocols for MANETs (mobile ad hoc networks) are designed for dense networks with permanent end-to-end connectivity. Recently, there has been significant interest in protocols designed for partitioned networks, where end-to-end connectivity exists rarely or never. Many research groups have studied partitioned networks in the last few years [1], [5], [12], [13], [17], [21], [24], [26], [27], [6].

The authors of [6] distinguish between two types of mobile partitioned networks: *sparse* and *clustered*. What distinguishes a clustered network from a sparse network is the number and size of its (connected) components. As a result of this qualitative difference in component size, it is typically sufficient in sparse networks to restrict communication between nodes to single hops, as this does not sacrifice many relay opportunities. For this reason, many existing DTN routing algorithms exploit only one-hop communication [1], [12], [15], [17], [28], [27], [29], [32], with the exception of [19], [21], [24], [6]. In contrast, in clustered networks, restricting communication to single hop relaying would be very limiting, given the larger clusters and the presence of longer paths.

In this paper, we propose a new mobility model for clustered networks, which we refer to as Heterogeneous Random Walk

(HRW). This model captures some of the qualitative properties observed in real mobility scenarios. Under this mobility model, every node performs an independent random walk on a plane, which is divided into regions characterized by different average node speeds. Within a low speed region the density of nodes is high, giving rise to rich connectivity; within high speed regions the connectivity is sparse.

The HRW mobility model is interesting for four major reasons. First, the HRW model is parsimonious, requiring only a small number of parameters to control the macroscopic properties of the model. Nevertheless, the model captures clustering and spatial heterogeneity, a ubiquitous feature of real mobility processes.

Second, although other models [2], [10], [16], [26] rely on similar mobility processes, e.g., random waypoint or random direction, they assume that each node has a set of so-called *preferred locations*. Such locations are visited by a node more often than other locations. In other words, nodes are not statistically equivalent, and their individual mobility can be well predicted. This is a strong assumption, with important implications for the design of robust routing protocols. This is not the case in the HRW mobility model, where nodes move independently following the same law. We provide evidence in this paper that in some scenarios, the mobility patterns of different nodes can be almost indistinguishable. This may be problematic for DTN routing protocols that assume a-priori that the movement patterns of nodes are very different and predictable, e.g., [1], [12], [17].

Third, another interesting macroscopic property of the HRW model is the stability of cluster locations. In reality, cluster formation can often be attributed to features of the natural or constructed environment (e.g., railway station, warehouse, parking lot, watering hole). For DTN routing, this is an important feature that can be exploited to design efficient schemes. This becomes clear in light of the approaches presented in [15], [24], [30]. Specifically, with our Island Hopping (IH) routing algorithm [24], we show that the existence of stable clusters can be exploited in routing protocols *even if individual nodes are statistically indistinguishable*.

Fourth, the HRW model captures an interesting relationship between speed and density, for which we provide some evidence in this paper. In the model, differences in node density (and hence connectivity) in different regions, can be viewed

as a consequence of different average speeds in these regions. We do observe such an inverse relationship between speed and density in at least one realistic mobility scenario, but suspect that it may be quite universal.

The remainder of the paper is structured as follows. In Section II we describe the model, calculate its stationary distribution, and give a method for “perfect simulation”. In Section III, we analyze real mobility traces to validate the key properties of the model. In Section IV, we compare several mobility models and their ability to predict the performance of an epidemic routing protocol. In Section V, we discuss related work. Finally, in Section VI we conclude the paper with a brief summary of our contributions.

II. THE HETEROGENEOUS RANDOM WALK MODEL

Here, we define the HRW model, calculate the stationary node distribution and conditions for the emergence of clusters, and we discuss how to perform a simulation that starts directly from the stationary distribution (i.e., “perfect simulation” [3].) We first define our model as a diffusion process, then show how to correct for boundary effects in a discrete-time approximation.

A. Definition

The simulation area is a torus of unit side length. We have n nodes that move independently of each other. We divide the torus into two regions A_l and A_h , where each node performs a two-dimensional Brownian motion with heterogeneous speeds (variances) in these two regions, with a slower speed in A_l than in A_h . The region A_l can be generated, for example, as the union of m random disks of radius r_l , and A_h its complement (cf. Figure 1). The difference in variance in the two regions gives rise to different node densities.

More formally ([14], Chapter 15), each node moves according to a process $Z = \{Z(t), t \geq 0\}$ taking values in $[0, 1]^2$ with the coordinate representation $Z(t) = (Z_x(t), Z_y(t))$, where :

- $Z(0)$ is an initial position in the torus;
- Z_x and Z_y satisfies the following conditions:

$$a_x(z, t) = \lim_{h \rightarrow 0} \frac{1}{h} E[Z_x(t+h) - Z_x(t) \mid Z(t) = z] = 0 \quad (1)$$

$$\begin{aligned} b_x(z, t) &= \lim_{h \rightarrow 0} \frac{1}{h} E[(Z_x(t+h) - Z_x(t))^2 \mid Z(t) = z] \\ &= b(z) := \begin{cases} \sigma_l^2 & \text{if } z \in A_l \\ \sigma_h^2 & \text{if } z \in A_h \end{cases}, \quad \sigma_l < \sigma_h, \end{aligned} \quad (2)$$

and similarly for $a_y(z, t)$ and $b_y(z, t)$;

$$\begin{aligned} a_{xy}(z, t) &= \lim_{h \rightarrow 0} \frac{1}{h} E[(Z_x(t+h) - Z_x(t)) \times \\ &\quad \times (Z_y(t+h) - Z_y(t)) \mid Z(t) = z] = 0. \end{aligned} \quad (3)$$

Thus, *infinitesimal means* $a_x(z, t)$ and $a_y(z, t)$ of the changes in Z_x and Z_y , respectively, are equal to 0 (1).

Infinitesimal variances $b_x(z, t)$ and $b_y(z, t)$ of the changes in Z_x and Z_y , respectively, are equal to σ_l^2 if a node is in A_l or σ_h^2 if a node is in A_h (2), where $\sigma_h^2 = \gamma \sigma_l^2$ with $\gamma > 1$. Note that the variance reflects the speed of the node. Moreover, we constrain ourselves to a process where the infinitesimal changes in Z_x and Z_y are uncorrelated (3).

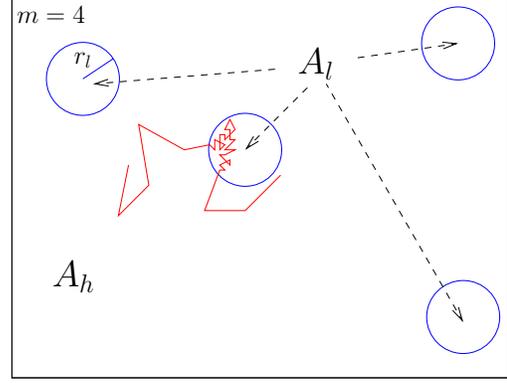


Fig. 1: Heterogeneous Random Walk - nodes move more slowly in region A_l than in A_h .

B. Stationary Distribution and Connectivity

A node’s position in the stationary regime has pdf

$$f(z) = \beta / b(z), \quad \beta = \frac{1}{|A_l|/\sigma_l^2 + |A_h|/\sigma_h^2}, \quad (4)$$

where $|A_l|$ and $|A_h|$ are the surface areas of the regions A_l and A_h , respectively (cf. Appendix A). The node density $\lambda(z) = n f(z)$, and hence, node densities λ_l and λ_h in regions A_l and A_h , respectively, relate as $\lambda_l = \gamma \lambda_h$ ($\gamma = \sigma_h^2 / \sigma_l^2 > 1$), i.e., the density of nodes is γ times larger inside the disks than outside of them. Thus, *the heterogeneous speed gives rise to the heterogeneous node density that is inversely proportional to the speed.*

The stationary probabilities that a node is located inside or outside a cluster are $\pi_l = \beta |A_l| / \sigma_l^2$ and $\pi_h = \beta |A_h| / \sigma_h^2$, respectively. We use this in Section IV to fit the parameters of our model to a mobility trace.

Let us now discuss the relationship between the connectivity of nodes and the parameters of our model. The regime of interest for our model is when node density is high enough for clusters to form inside the disks, and sparse outside. We rely on results from continuous percolation theory [18]. Assume a Poisson point process with density λ and assume that two nodes are connected iff they are within distance r . Then, if $\lambda r^2 > (\lambda r^2)_{cr} \approx 1.43$ (where $(\lambda r^2)_{cr}$ is called the percolation threshold), an infinite cluster appears with a positive probability. This is called the supercritical regime. In the subcritical regime $\lambda r^2 < (\lambda r^2)_{cr}$, the clusters are almost surely finite. Although these results hold for the infinite plane, they are good approximations as long as the node density is sufficiently high ([22], Proposition 2).

Therefore, we choose the parameters of our model such that we are in the supercritical regime in A_l , i.e., $\lambda_l r^2 > (\lambda r^2)_{cr}$,

and such that we are in the subcritical regime in A_h , i.e., $\lambda_h r^2 < (\lambda r^2)_{cr}$.

The following table summarizes the parameters of the model.

n	# of nodes
m	# of disks (clusters)
r_l	disk radius
$\gamma = \sigma_h^2 / \sigma_l^2$	ratio between variances in A_l and A_h
r	connectivity range

C. Perfect Simulation and Discrete-Time Approximation

In this subsection, we provide an algorithm allowing for “perfect simulation”. The algorithm initializes node positions to start the simulations immediately in a stationary regime, thus obviating the need for a transient warmup phase.

We begin the simulations as follows. We need to place n nodes according to the stationary distribution given by (4)¹. First, we find the initial number N_l of nodes in A_l , which is $\sim \text{Bin}(n, \pi_l)$ (and $N_h = n - N_l$). We then need to place N_l nodes uniformly at random in the union of disks. One straightforward way to achieve this is by dropping nodes randomly on the torus, and retaining only those falling into a disk, until we have enough nodes. We proceed in analogous fashion for the N_h external nodes.

After placing the nodes, we run our simulations as follows. At every time step each node moves for $\Delta = (\Delta_x, \Delta_y)$, where Δ_x and Δ_y are independent variables chosen independently for every node and for every time step. Each of these variables has a Gaussian distribution with mean 0 and variance σ_l^2 if a node is in A_l or σ_h^2 if it is in A_h .

As the HRW model is a diffusion process with spatially dependent coefficient $b(z)$, our simulations will lead to a systematic error, as shown in [8]. The authors in [8] propose an elegant solution to this problem in a general case when $b(z)$ is a differentiable function. This solution consists in correcting every step length Δ , depending on the gradient of $b(z)$ during this step. As in our case $b(z)$ is a step function their solution is not straightforwardly applicable. Nevertheless, following their approach we find the correction for Δ for our model (see Appendix B).

D. Implications for Routing

One motivation for our model is to shed light on the design of routing and forwarding algorithms for partitioned mobile networks, and more specifically, on the use of control information in these algorithms. The design and analysis of routing algorithms for such networks is more challenging and complex than for connected networks, because (i) a route is only realized over time, and (ii) the exchange of control information itself has to rely on opportunistic connectivity, and is therefore subject to the dynamics of the network topology.

The design of routing protocols usually relies on a set of assumptions on the network topology and dynamics. For example, a particular mobility model (random waypoint, random walk, etc.) or a class of network topologies is postulated. If such assumptions are too strong, there is a danger of tailoring the algorithm to these assumptions, instead of designing a robust algorithm which is effective over a wide range of settings. Therefore, we need to carefully check these assumptions, and test candidates against different models to assess their robustness.

We introduce the HRW model because we believe that it represents a challenging but realistic corner case that DTN routing algorithms should be able to handle. Specifically, the key distinguishing features of the model are that (i) nodes are statistically equivalent, and (ii) clustering arises as a property of the environment, rather than as a property of the nodes.

It should be clear that these two features would hamper any attempt to base routing decisions on *nodal statistics* (e.g., preferred locations for particular nodes, or encounter frequencies between specific pairs of nodes), which appears to be an implicit or explicit assumption built into many routing protocols [15], [31].

This also allows us to shed some light on an important consideration in protocol design, namely: what control information should nodes collect and exchange to drive routing decisions? In the corner case represented by our model, we can shed some light on this question, which we now briefly sketch without giving formal proofs.

Consider a modified version of the HRW model, where nodes move on a square grid rather than in Euclidean space, and where two nodes are connected iff they are located at the same grid position. This model is interesting to study routing protocols, because forwarding a message to another node has no benefit in terms of distance (because both nodes must be located at the same grid position); the only benefit may lie in the future evolution of a node which receives the copy.

Different routing protocols proposed in the literature collect and exchange different statistics about the network (e.g., encounter times with other nodes). We can bound the performance of such routing protocols by an oracle that reveals the entire past and current state to every node, and ask how a routing protocol should be designed within the HRW model.

First, consider so-called single-copy routing strategies, where only a single copy of a message exists at any given time in the network [28]. In the (grid-based) HRW model, the optimal routing strategy would never forward a message. This is because conditional on the current state of the network, the future evolution of any node is i.i.d. As a consequence, there is no value in statistics about the network state.

Second, let us focus on multi-copy strategies. One approach is the so-called spray-and-wait approach [27], where copies of a message are initially sent to a set of relay nodes; each relay node then waits to meet the destination to hand off the message, but no further copies are made. In the HRW model, it can be shown that the choice of the set of nodes carrying the message is irrelevant; therefore, again, past information has no

¹The node positions are “almost” a Poisson process with heterogeneous intensity $\lambda(z)$, except for the condition that the total number of nodes is exactly n instead of $Po(n)$.

value under HRW mobility ².

Third, and more subtly, if copy decisions are allowed at later points in time, then there *is* value in statistics about the past, despite the Markovian structure of the system. This is because information about the past can help to generate additional copies of a message at opportune times. For example, when a node with a copy of a message encounters other nodes that have recently met the destination, then this is an indication that the destination is near; it might then be beneficial to generate additional copies of the message to “catch” the destination. It is an open question what an optimal multi-copy strategy based on an oracle would look like, what the best possible performance penalty of a scheme without oracle is, and what statistics such a scheme needs to collect.

In summary, we can (very loosely) say that in a system where nodal mobility is independent and not predictable conditional on the current location (random walk), multi-copy schemes can nevertheless be designed so that the set of copies of a message *does evolve towards the message destination*. However, this cannot be achieved through either a single-copy or a spray-and-wait approach, where copy decisions are only made once, but requires that copy decisions are made repeatedly along the way. Our work [24] is such an example, where the nodes collectively learn the topology of clusters, and continuously create and destroy message copies, ensuring that the “bundle” of messages as a whole makes progress towards the destination, even though the underlying nodes perform independent random walks.

III. VALIDATION

In this section, we show that certain macroscopic characteristics specific to clustered networks are prevalent in real mobility traces and that they can be captured by the HRW mobility model.

For the purpose of our study we use two large GPS-based mobility traces. The main data set contains mobility traces of taxi cabs in San Francisco, USA³. It contains GPS coordinates of approximately 500 taxis collected over 30 days in the San Francisco Bay Area. Each taxi is equipped with a GPS receiver and sends a *location-update* (timestamp, identifier, geo-coordinates) to a central server. The location-updates are quite fine-grained - the average time interval between two consecutive location updates is less than 10 sec, allowing us to accurately interpolate node positions between location-updates.

The second data set is less detailed, and we use it only for one experiment below. It contains GPS traces collected by mobile phone users subscribed to the Nokia Sportstracker service⁴. This data set is larger than the first in terms of number of mobile nodes, but also covers a much larger geographical region and time period (more than one year). Mobile phone users can upload workout and activity traces to the Nokia Sportstracker web site to store and share with others. For the

purpose of this study we focused on subsets from Helsinki, Finland (3757 distinct GPS traces), Stockholm, Sweden (1056 distinct GPS traces) and London, UK (2488 distinct GPS traces). The data set combines traces from different activities (running, walking, cycling, etc.), and as such only provides occasional snapshots of a person’s long-term mobility. Therefore, it is less suited for the purposes of this study, where we are interested in long-term mobility patterns of a set of nodes.

A. Statistical Equivalence of Nodes

In the following experiment, we test to what extent the nodes in the San Francisco taxi trace follow distinguishable mobility patterns, e.g., because drivers might have location preferences. This is an important consideration, because it measures the potential for a DTN routing algorithm to make judicious decisions about *which node* a packet to forward to.

We assume a null hypothesis that all the node mobility patterns are drawn from the same underlying distribution, and use a Pearson’s χ^2 test to decide whether we need to reject the null hypothesis, separately for each trace. More specifically, we look at the distribution of counts of visits to a grid of square cells of equal size for a time period of one day. This time period is a conservative upper bound of the time-scale of interest in a routing protocol. We simply take the true distribution as the aggregate over all nodes.

For this analysis we first divide the whole area of San Francisco into a set of square cells of equal size (1000×1000 m). For each day, we compute the χ^2 statistic for each node to measure the difference between its cell counts and that of the entire population. See Figure 2 for four snapshots of daily counts; the plots suggest that the distributions are indeed quite similar.

Call O_z a count for cell z , and E_z the count under the reference distribution. To prevent overestimation of statistical significance for small counts, we only include cells with sufficient probability mass, i.e., the set $Q := \{z : E_z \geq 5\}$. The test statistic is $X^2 = \sum_{z \in Q} \frac{(O_z - E_z)^2}{E_z}$. We find the p -value of the test as $p = \mathbf{P}\{X^2 \leq \chi^2(|Q| - 1)\}$, where $\chi^2(|Q| - 1)$ is a χ^2 random variable with $|Q| - 1$ degrees of freedom. The test rejects the null hypothesis if the p -value is smaller than the significance level (e.g., $\alpha = 0.05$).

In Figure 3 we plot the CDF of the p -values found for each node on each day for the whole period of 30 days. We find that most of the traces (more than 60%) are statistically indistinguishable (p -value $> \alpha = 0.05$) from the overall population. We note that over longer time scales, we would reject the null hypothesis more frequently, because small differences in the underlying distributions would be amplified relative to the sampling noise. However, in most applications of interest (such as forwarding in DTN routing protocols), the time-scale of interest tends to be of the order of minutes to hours; over these time-scales, for the mobility trace at hand, we have to assume that a majority of nodes follow similar mobility patterns.

²Except possibly through the *number* of copies that are generated.

³<http://cabspotting.org>

⁴<http://sportstracker.nokia.com/>

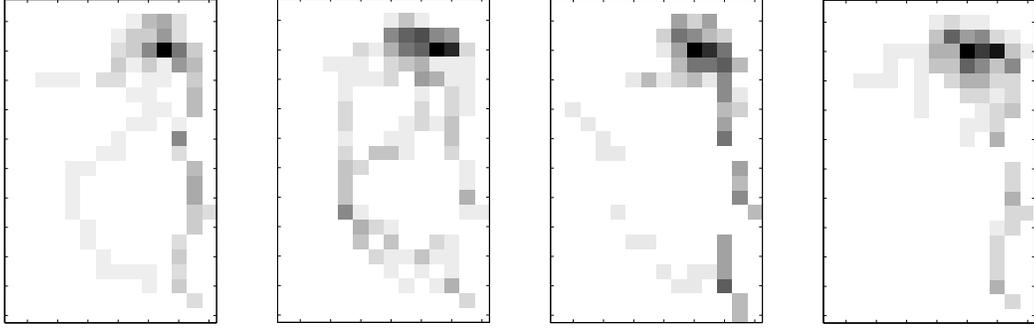


Fig. 2: Spatial histogram of cell counts for four different taxis during one day. A darker pixel indicates higher count.

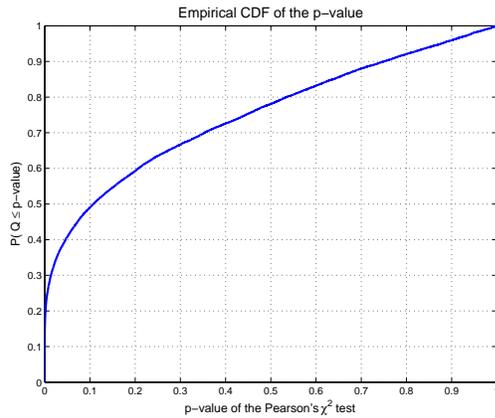


Fig. 3: The CDF of the p -value of the Pearson's χ^2 test that verifies the statistical equivalence of mobile nodes from the San Francisco data set.

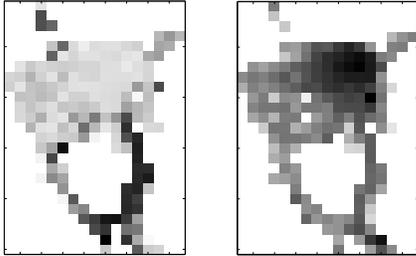


Fig. 4: Comparison of the speed map (left) and density map (right), based on a 5 hour period (from 16h00 until 21h00 local time on a working day) from the San Francisco trace, with cell size 1000×1000 m. Darker pixels indicate higher speed/density.

B. Speed and Density Maps

Here we focus on the spatial distribution of vehicle speed and density. To find the average speed at a location, we again divide the whole area of San Francisco into a set of square cells of equal size. Then, for each cell z we record the vehicle's speed at which it enters cell z to obtain an average *cell speed* v_z .

Next, we focus on the density of nodes over the same grid of square cells. At every time instant, we record the number

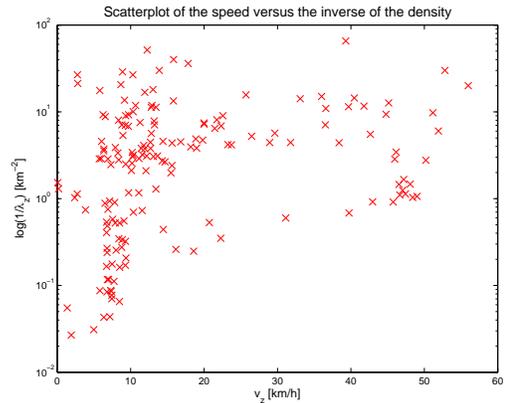


Fig. 5: Correlation between the speed and the inverse of density, for cells with at least 20 updates.

of vehicles visiting a particular cell z . Next, we find the *node density* λ_z for each cell z . As a result, we obtain a so-called *density map*. To generate an example density map, we use the same 5 hour snapshot of the trace as for the speed map above. As the spatial distribution of node density is heavy tailed, for the sake of visualization, we present an example map of $\log(\lambda_z)$ in Figure 4b.

The figure does suggest that speed and density are negatively correlated, i.e., that locations of high speed see low node density, and vice versa. In Figure 5, we show a scatterplot of speed and density for all cells z at which there were at least 20 location-updates (to average out noise). This plot suggests that the negative correlation is present, but amid a lot of noise.

The HRW model controls node density and average speed with a single parameter $\sigma_{l,h}^2$. High node density inside A_l arises because nodes move more slowly, and therefore spend relatively more time than in A_h . We conjecture that such an inverse relationship may be quite universal, because the mobility of nodes may be more “constrained” by other nodes in high-density areas, limiting their speed. This is quite easy to see in vehicular settings (highway vs. downtown traffic jam).

C. Connectivity

Let us first define the *connectivity graph*. Here again, as in Section II, we assume a (scaled) unit disk model for connectivity. The mobile nodes and the corresponding wireless links

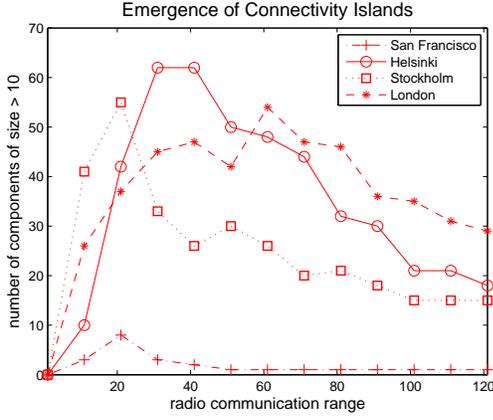


Fig. 6: Emergence of connectivity islands in the connectivity graph - each data point represents the number of components of size larger than 10 in the connectivity graph $G(V, E)$ for different values of r .

define the *connectivity graph* $G(V, E)$, where $V(G)$ is the set of mobile nodes and $E(G)$ is the set of radio links between mobile nodes, i.e., $E(G) = \{e = (i, j) | d_{ij} \leq r\}$. We define H_k as a component of G , with $C(G) = \{H_1, H_2, \dots, H_K\}$ the set of all components, i.e., $G = \bigcup_{k=0}^{K=K} H_k$.

To show that islands of connectivity may emerge we study how the structure of the connectivity graph changes with increasing r . More precisely, we focus on the size of components $|H_k|$ present in the connectivity graph $C(G)$ generated at a random time instant from both the taxi cab and Nokia Sportstracker data sets.

For this experiment, we *densify* our trace in the following way. We assume that the sampled nodes are a representative subset of the overall traffic, and that their movements are stationary and ergodic. Under this assumption, we can generate a denser sample of instantaneous node locations (in our case, 5000 nodes) by sampling uniformly at random from the entire dataset. This denser sample brings out more detail in the connectivity graph.

In Figure 6 we show the results for four cities: San Francisco, Helsinki, Stockholm and London. We observe a clear trend in all the cases: the number of components rises quickly with the communication range r towards a maximum; past this critical value, the number of components starts to decrease, because smaller components start coalescing into larger ones. However, this decrease tends to be slow; even when r becomes a multiple of the critical value, we still have many components left. This is because node locations are distributed very non-uniformly in space, which prevents percolation into a single giant cluster when r grows. Thus, we conclude that disconnected network topologies with a large number of components seem to be a robust phenomenon which persists over a wide range of radio ranges.

D. Cluster Dynamics

In the previous subsection, we studied components in instantaneous snapshots of the connectivity graph. We now turn

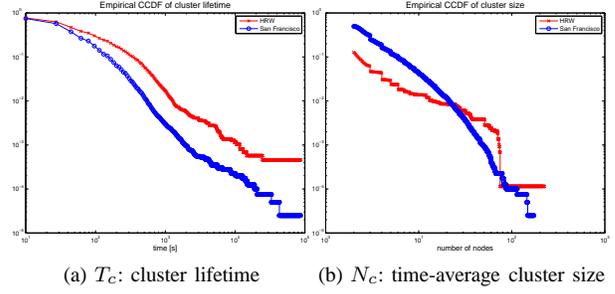


Fig. 7: Comparison of cluster lifetime and cluster size in both the trace and the model.

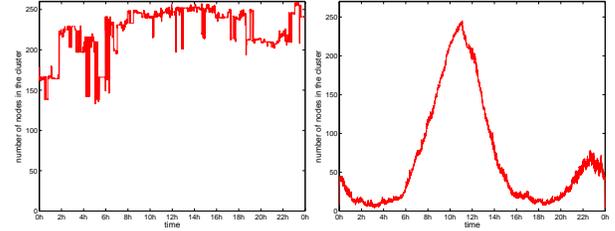


Fig. 8: Comparison of the size evolution of the largest long lasting clusters in the model and San Francisco trace.

to the dynamics of the connectivity graph over time, to verify to what extent the HRW model approximates reality.

For this, we focus on clusters, which can be thought of as sequences of “related” components over time. Note that from one time step to the next, the node set of a component may change slightly because of mobility, and components can merge or split. In order to meaningfully analyze cluster dynamics, we need to establish a correspondence between components at subsequent time steps. To identify such clusters, we rely on the CCL algorithm proposed in our earlier work [25].

We would like to validate how well our model captures the cluster dynamics over real mobility scenario. Therefore, we analyze the cluster dynamics of both real clusters observed in the San Francisco trace, as well as clusters produced in the model whose parameters were inferred from the San Francisco trace. We explain in detail the method for extracting the HRW model parameters from a trace in the Section IV.

We analyze the following two metrics:

- T_c : *cluster lifetime* - the time a cluster is present in the network;
- N_c : *cluster size* - the average number of nodes of a cluster over its lifetime.

In Figure 7, we show the empirical complementary CDF for the cluster lifetime T_c and the average cluster size N_c . In the case of the mobility trace, we observe that the cluster lifetime spans from 10 seconds up to one day; similarly for the cluster size - it varies from 2 to 200. We observe that in both the real-life mobility and in the synthetic trace the distributions of the lifetime and the cluster size.

We can make the following observations. First, looking at the evolution of the size of the largest cluster in both the model

and the trace (cf. Figure 8), we see - not surprisingly - strong diurnal fluctuations in one but not the other.

Second, the lifetime of clusters both as produced by the model and inferred from the trace is heavy-tailed. However, the model produces more stable clusters, which in part results from the absence of diurnal fluctuations in the model, and from the sharp difference between high and low-density regions. In a real setting, the boundaries are of course more fuzzy.

Third, the cluster sizes for the model and the trace are distributed rather differently. The model exhibits clear modes (the largest for approx. 250 nodes, the size of the largest cluster), which correspond to the individual supercritical disks fitted to the trace (cf. Figure 9). Given the homogeneity of the mobility processes, the size of each cluster tends to fluctuate in a narrow range. In the trace, there is more of a continuous spectrum of cluster sizes, again, in part, because of fluctuations during the day. Small clusters arise more frequently in the trace than in the model, which is a result of the subcritical region not being uniform as stipulated in the model.

IV. APPLICATION TO EPIDEMIC DISSEMINATION

In this section, we evaluate the predictive power of the HRW model in the context of epidemic dissemination, one of the most prominent paradigms for routing in DTNs. The HRW model has very few parameters, similar to other prominent mobility models. We argue that the HRW model appears to predict the performance of an epidemic dissemination protocol more accurately than other models, which suggests that clustering and spatial heterogeneity are salient features of real mobility scenarios that should not be abstracted away. However, we stress that this evaluation is not exhaustive enough to conclude that the HRW model is appropriate for all scenarios of interest, and should be viewed as preliminary.

We first start by explaining how we fit the models under consideration to the mobility trace. First, the average speed $\bar{\sigma}$ of a node under HRW is given by:

$$\bar{\sigma}^2 = \pi_h \sigma_h^2 + \pi_l \sigma_l^2, \quad (5)$$

where π_l (π_h) is the probability that a random node is located inside (outside) a cluster. Hence, from (5) and $\gamma = \sigma_h^2/\sigma_l^2$, we can find the speeds inside and outside clusters.

For the HRW model, we first estimate the average speed $\bar{\sigma}^2$ of all the taxis operating within the entire period of 30 days. Next, as in the previous section, we rely on the CCL algorithm to identify the clusters in the trace. Then, we estimate γ as the ratio of node density inside and outside clusters. Finally, we find the empirical probability of a random node being associated with a cluster, i.e., π_l - which is given by the ratio between the total time nodes associate with clusters to the total time they operate in the network. Then, using (5) we find σ_l and σ_h . In addition, we infer also the cluster locations and their radii by identifying the center of mass of long-term clusters and the average distance of a node from each cluster to the cluster's center of mass. For the San Francisco trace, this yields the following parameters of the HRW mobility model:

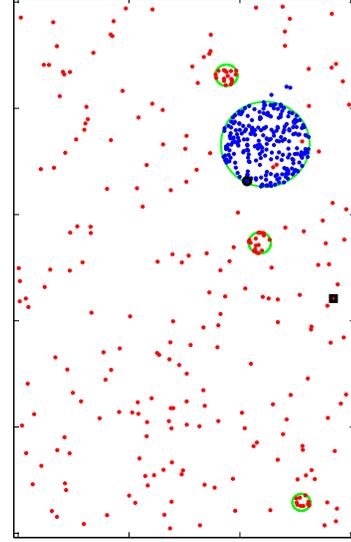


Fig. 9: A snapshot of the execution of the epidemic dissemination on the HRW mobility model fitted to the San Francisco trace. Dots represent nodes and bright green circles show the perimeter of the supercritical disks. The black disk represents the source of the message, whereas a black square represents the destination. Blue (red) dots are nodes which do (not) hold a copy of the message.

n	500
m	4
r	300 [m]
r_l	400, 500, 500, 2000 [m]
σ_l^2, σ_h^2	0.28, 9.11
$\lambda_l r^2, \lambda_h r^2$	1.75, 0.05
π_l, π_h	0.57, 0.43

To fit the RW and RWP mobility models, we set the node speed equal to $\bar{\sigma}$, and we set the RWP pause time to zero. For all the models we simulate 500 mobile nodes on a rectangle of size 15x25 km (the same size as for the San Francisco trace). Finally, we set the beaconing period of the epidemic protocol to 10 sec.

For the purpose of our study, we use the following performance metrics:

- *message delay* - the time interval between sending the message by the source and receiving its copy at the destination;
- *message copies* - total number of copies made for a message in the network during the message delay time interval;
- *minimum delay path* - the hop-length of the minimum delay path between the source and destination discovered by the protocol.

In Figure 9 we show an execution snapshot of dissemination in the HRW model. We perform 300 independent simulations for each mobility model, where both the source and destination are chosen independently at random. For the trace-based

simulation, for each day during the period of 30 days, we ran 10 simulations, varying the destination and the message sending time.

In Figure 10 we give the performance evaluation results of the epidemic protocol using a box-plot for the three metrics. First, we see that a (homogeneous) random walk gives a message delay which is much larger than the real delay; the HRW model is still conservative, but much closer to the actual performance. This suggests that the speed at which the message spreads benefits greatly from clusters (which allow a rapid spread of the message to other nodes), and from the low-density complement (where messages can travel quickly). The RWP model underestimates the message delay, but is closest to reality. Second, the three models predict the number of message copies fairly well. Third, the RWP model underestimate the minimum delay path. This is because the RWP model has a very small mixing time, which makes it likely that the message reaches its destination after only a small number of hops.

V. RELATED WORK

Exhaustive surveys of mobility models can be found in [4], [7]. Here we only focus on macroscopic mobility models designed for partitioned networks. To the best of our knowledge there are only a few such mobility models [2], [10], [16], [20], [26].

In [2] Bittner et al. propose the *Area Graph Based Mobility Model* that considers the inhomogeneity of spatial node distribution. It is based on a directed graph with areas of different densities. In [10] Hsu et al. propose a similar approach called the *Weighted Waypoint Mobility Model* that captures node preferences in choosing destinations. Unlike in [2], Hsu et al. design their model based on mobility data (mobility survey on a campus). Although their model is mathematically tractable, they do not provide a closed-form expression for time-stationary distribution. Hence, they do not have a method for “perfect simulation”. The same holds for the model called *Clustered Mobility Model* proposed in [16] by Lim et. al. All three models above exhibit a skewed distribution of location visits. This could result in forming islands of connectivity at popular locations, but it is not clear if they will remain stable. Moreover, their occurrence is due to a “preferential attachment” mechanism. In contrast, in the HRW model, clusters emerge in because of low node speed in popular areas.

Musolesi et al. in [20] present a mobility model that incorporates sociological relationships between mobile nodes, but because of this they pay the price of high model complexity. This is why their model is not easily applicable in the simulation of partitioned networks. The *Community Based Mobility Model* proposed by Spyropoulos et al. in [26], and further extended in [11], is designed to recreate the characteristics of mobility in sparse networks. It captures well the skewed location visiting preferences. A careful adjustment of its parameters may lead to the appearance of connectivity islands, e.g. by assigning the same community to groups of nodes and synchronizing their movement, such that a sufficient

number of nodes will be present in the community at the same time. Our model achieves the same through a simpler, more parsimonious model.

VI. CONCLUSION

Based on evidence from mobility traces, we propose a parsimonious mobility model called Heterogeneous Random Walk. It captures some of the important mobility characteristics observed at the macroscopic level in real mobility data, yet is mathematically simple. We give a closed-form expression of the time-stationary node distribution and the recipe for “perfect simulation”.

A key feature of the model is that nodes follow independent and statistically equivalent mobility patterns, despite the presence of long-term clusters. The validity of this assumption very much depends on the scenario at hand. For example, while it is likely that the mobility of private cars would exhibit strong differences from one car to the next, we have shown one scenario in this paper where different nodes follow very similar patterns. In summary, we believe that our model represents a useful corner case to consider in the design of routing protocols for partitioned networks, in order to guard against unrealistic assumptions on the predictability of node movements.

ACKNOWLEDGMENTS

We would like to thank George Theodorakopoulos for valuable feedback and discussions on this paper. We also thank Holly Cogliati for help in improving the manuscript. We are indebted to the Exploratorium - the museum of science, art and human perception for making the San Francisco taxi traces publicly available through the cabspotting project.

The work presented in this paper has been supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5006-67322.

REFERENCES

- [1] A. Balasubramanian, B. Levine, and A. Venkataramani. DTN Routing as a Resource Allocation Problem. *SIGCOMM Comput. Commun. Rev.*, 37(4):373–384, 2007.
- [2] S. Bittner, W.-U. Raffel, and M. Scholz. The Area Graph-based Mobility Model and its Impact on Data Dissemination. In *PERCOMW '05: Proceedings of the Third IEEE International Conference on Pervasive Computing and Communications Workshops*, pages 268–272, Washington, DC, USA, 2005.
- [3] J.-Y. L. Boudec. Understanding the Simulation of Mobility Models with Palm Calculus. *Performance Evaluation*, 64(2):126–147, 2007.
- [4] T. Camp, J. Boleng, and V. Davies. A Survey of Mobility Models for Ad Hoc Network Research. *Wireless Communications & Mobile Computing (WCMC): Special Issue on Mobile Ad Hoc Networking: Research, Trends and Applications*, 2(5):483–502, 2002.
- [5] A. Chaintreau, P. Hui, C. Diot, R. Gass, and J. Scott. Impact of Human Mobility on Opportunistic Forwarding Algorithms. *IEEE Transactions on Mobile Computing*, 6(6):606–620, 2007.
- [6] Y. Chen, W. Zhao, M. Ammar, and E. Zegura. Hybrid Routing in Clustered DTNs with Message Ferrying. In *MobiOpp '07: Proceedings of the 1st International MobiSys Workshop on Mobile Opportunistic Networking*, pages 75–82, San Juan, Puerto Rico, 2007.

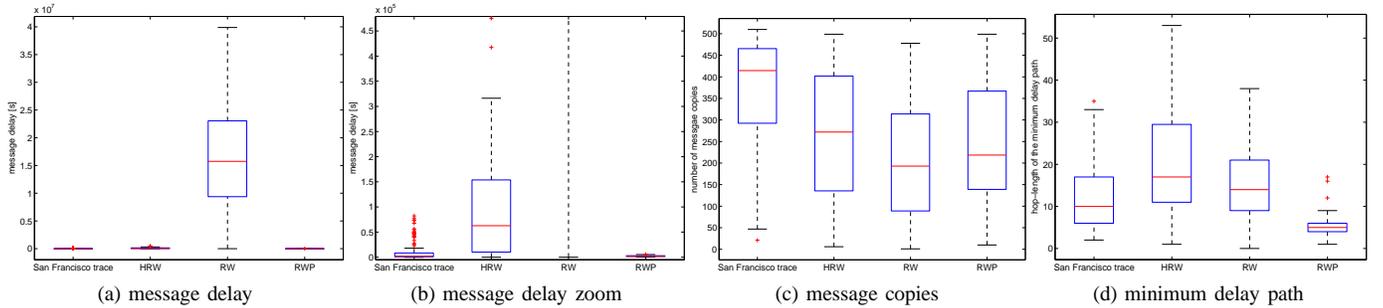


Fig. 10: Performance evaluation of the simple epidemic dissemination protocol under different mobility processes (San Francisco trace, HRW, RW and RWP).

- [7] B. Fan and A. Helmy. *Wireless Ad Hoc and Sensor Networks*, chapter A Survey of Mobility Models in Wireless Adhoc Networks, pages 1–29. Kluwer Academic Publishers, 2004.
- [8] L. Farnell and W. G. Gibson. Monte Carlo Simulation of Diffusion in a Spatially Nonhomogeneous Medium: Correction to the Gaussian Steplength. *J. Comput. Phys.*, 198(1):65–79, 2004.
- [9] C. W. Gardiner. *Handbook of stochastic methods : for physics, chemistry and the natural sciences*. Springer, 2004.
- [10] W. Hsu, K. Merchant, H. Shu, C. Hsu, and A. Helmy. Weighted Waypoint Mobility Model and its Impact on Ad Hoc Networks. *SIGMOBILE Mob. Comput. Commun. Rev.*, 9(1):59–63, 2005.
- [11] W. J. Hsu, T. Spyropoulos, K. Psounis, and A. Helmy. Modeling Time-variant User Mobility in Wireless Mobile Networks. In *Proceedings of the 26th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, Anchorage, Alaska, March 2007.
- [12] S. Jain, K. Fall, and R. Patra. Routing in a Delay Tolerant Network. In *SIGCOMM '04: Proceedings of the 2004 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 145–158, Portland, Oregon, USA, 2004.
- [13] T. Karagiannis, J.-Y. L. Boudec, and M. Vojnović. Power Law and Exponential Decay of Inter Contact Times between Mobile Devices. In *MobiCom '07: Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking*, pages 183–194, Montréal, Québec, Canada, September 2007.
- [14] S. Karlin and H. M. Taylor. *A second course in stochastic processes*. Academic Press, 1981.
- [15] J. Leguay, T. Friedman, and V. Conan. Evaluating Mobility Pattern Space Routing for DTNs. In *Proceedings of the 25th IEEE International Conference on Computer Communications (INFOCOM)*, Barcelona, Spain, April 2006.
- [16] S. Lim, C. Yu, and C. Das. Clustered Mobility Model for Scale-Free Wireless Networks. In *Proceedings of the 31st IEEE Conference on Local Computer Networks*, pages 231–238, Tampa, Florida, USA, Nov. 2006.
- [17] A. Lindgren, A. Doria, and O. Schelén. Probabilistic Routing in Intermittently Connected Networks. *SIGMOBILE Mob. Comput. Commun. Rev.*, 7(3):19–20, 2003.
- [18] R. Meester and R. Roy. *Continuum percolation*. Cambridge University Press, 1996.
- [19] M. Musolesi, S. Hailes, and C. Mascolo. Adaptive Routing for Intermittently Connected Mobile Ad Hoc Networks. In *WoWMoM '05: Sixth IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks*, pages 183–189, June 2005.
- [20] M. Musolesi and C. Mascolo. Designing Mobility Models based on Social Network Theory. *SIGMOBILE Mob. Comput. Commun. Rev.*, 11(3):59–70, 2007.
- [21] J. Ott, D. Kutscher, and C. Dwertmann. Integrating DTN and MANET Routing. In *CHANTS '06: Proceedings of the 2006 SIGCOMM Workshop on Challenged Networks*, pages 221–228, Pisa, Italy, 2006.
- [22] M. Penrose and A. Pisztor. Large Deviations for Discrete and Continuous Percolation. *Advances in Applied Probability*, 28(1):29–52, 1996.
- [23] H. Risken. *The Fokker-Planck equation : methods of solution and applications*. Springer, 1989.
- [24] N. Sarafijanovic-Djukic, M. Piórkowski, and M. Grossglauser. Island Hopping: Efficient Mobility-assisted Forwarding in Partitioned Networks. In *SECON '06: 3rd Annual IEEE Communications Society on Sensor and Ad Hoc Communications and Networks*, pages 226–235, Reston, VA, USA, Sept. 2006.
- [25] N. Sarafijanovic-Djukic, M. Piórkowski, and M. Grossglauser. Island Hopping: Efficient Mobility-assisted Forwarding in Partitioned Networks. Technical report, EPFL, 2007.
- [26] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Performance Analysis of Mobility-assisted Routing. In *MobiHoc '06: Proceedings of the 7th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 49–60, Florence, Italy, 2006.
- [27] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Efficient Routing in Intermittently Connected Mobile Networks: the Multiple-copy Case. *IEEE/ACM Trans. on Networking*, 16(1):77–90, 2008.
- [28] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Efficient Routing in Intermittently Connected Mobile Networks: the Single-copy Case. *IEEE/ACM Trans. on Networking*, 16(1):63–76, 2008.
- [29] A. Vahdat and D. Becker. Epidemic Routing for Partially Connected Ad Hoc Networks. Technical report, Duke University, 2000.
- [30] E. Yoneki, P. Hui, S. Chan, and J. Crowcroft. A Socio-aware Overlay for Publish/Subscribe Communication in Delay Tolerant Networks. In *MSWIM 07: Proceedings of the 10th ACM Symposium on Modeling, Analysis, and Simulation of Wireless and Mobile Systems*, pages 225–234, Chania, Greece, Oct. 2007.
- [31] X. Zhang, J. Kurose, B. N. Levine, D. Towsley, and H. Zhang. Study of a Bus-based Disruption-tolerant Network: Mobility Modeling and Impact on Routing. In *MobiCom '07: Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking*, pages 195–206, Montréal, Québec, Canada, 2007.
- [32] W. Zhao, M. Ammar, and E. Zegura. A Message Ferrying Approach for Data Delivery in Sparse Mobile Ad Hoc Networks. In *MobiHoc '04: Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 187–198, Tokyo, Japan, 2004.

APPENDIX

A. Stationary Distribution

Proof: The diffusion process is described by a Fokker-Planck equation (or Forward Kolmogorov equation), which in the two-dimensional case is given by ([9], p.145):

$$\begin{aligned}
 \frac{\partial f(z,t)}{\partial t} &= -\frac{\partial}{\partial x}(a_x(z,t)f(z,t)) - \frac{\partial}{\partial y}(a_y(z,t)f(z,t)) + \\
 &+ \frac{1}{2} \frac{\partial^2}{\partial x \partial y}(b_{xy}(z,t)f(z,t)) + \\
 &+ \frac{1}{2} \frac{\partial^2}{\partial x \partial y}(b_{yx}(z,t)f(z,t)) + \\
 &+ \frac{1}{2} \frac{\partial^2}{\partial x^2}(b_x(z,t)f(z,t)) + \\
 &+ \frac{1}{2} \frac{\partial^2}{\partial y^2}(b_y(z,t)f(z,t)), \tag{6}
 \end{aligned}$$

where $f(z,t) = P(Z(t) = z | Z(t_0) = z_0)$ with $Z(t) = (Z_x(t), Z_y(t))$ denoting the two dimensional diffusion, and with the values of the coefficients given by (1), (2), and (3).

Our goal is to find a stationary distribution given by $f(z) = \lim_{t \rightarrow +\infty} f(z, t)$. Hence, $f(z)$ is a unique solution of (6) when $t \rightarrow +\infty$, in which case it holds $\frac{\partial f(z, t)}{\partial t} = 0$.

To fully describe a diffusion process we also need to know the boundary conditions. Here we have periodic boundary conditions (see [9], p. 119 and p. 121) given by:

$$f(x, 1) = f(x, 0), f(1, y) = f(0, y); J_x(x, 1) = J_x(x, 0), J_y(1, y) = J_y(0, y), \quad (7)$$

where

$$J_x(x, y) = -0.5 \frac{\partial}{\partial x} (b(x, y) f(x, y)), J_y(x, y) = -0.5 \frac{\partial}{\partial y} (b(x, y) f(x, y)) \quad (8)$$

The unique solution of (6) with the conditions (7) is given by (4). ■

B. Steplength Correction for Simulation

As explained in Section II-C, due to the spatially heterogeneous coefficient $b(z)$, our discrete-time simulation have a systematic error. Therefore, a step length Δ should be corrected whenever $b(z)$ has different values during this step. Figures 11 and 12 show how to make the correction. We consider only cases when $b(z)$ changes its value once during one step, because we assume Δ is small enough so that the cases when $b(z)$ changes its value more than once during one step happen rarely.

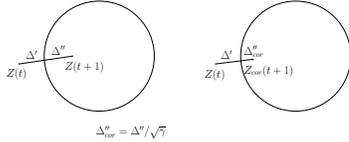


Fig. 11: Case I - a node enters a disk.

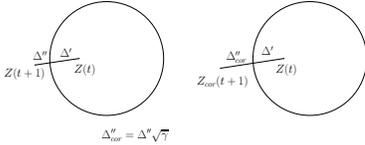


Fig. 12: Case II - a node leaves a disk.

We look at the projections of the step length Δ into x and y coordinates and hence we show how to calculate the correct step length $\Delta_{corr} = \Delta + \epsilon$ in one-dimensional case (cf. Figure 13).

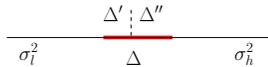


Fig. 13: One-dimensional case.

To calculate a correction term ϵ we start from the following equation [8]:

$$\epsilon = \frac{\int_{-\infty}^{\Delta} g(x) dx - \int_{-\infty}^{\Delta} g_{corr}(w) dw}{g_{corr}(\Delta)}, \quad (9)$$

where Δ is the step length for the case when $b(x) = \sigma_l^2$ (i.e., $b(x)$ is a constant for all x), and $g(x)$ is the distribution of the step length Δ , and $g_{corr}(w)$ is the distribution of the exact step length Δ_{corr} (when $b(x)$ is a step function).

By the normalization technique given in [23], Chapter 5, we obtain:

$$g_{corr}(w) = \begin{cases} g(w) & \text{if } b(w) = \sigma_l^2 \\ \frac{\sigma_l}{\sigma_h} g(w) = \frac{1}{\sqrt{\gamma}} g(w) & \text{if } b(w) = \sigma_h^2 \end{cases} \quad (10)$$

We divide Δ in two parts as shown in Figure 13, Δ' where $b(x) = \sigma_l^2$ and Δ'' where $b(x) = \sigma_h^2$. Then, using (9) and (10) we obtain $\epsilon = -(1 - 1/\sqrt{\gamma})\Delta''$ which means that:

$$\Delta_{corr} = \Delta' + \frac{1}{\sqrt{\gamma}} \Delta'' \quad (11)$$