

Using Crowdsourcing for Local Topology Discovery in Wireless Networks

Alessandro Checco¹, Carlo Lancia², and Douglas J. Leith¹

¹*Hamilton Institute, National University of Ireland Maynooth*

²*University of Rome Tor Vergata*

Abstract

In this paper we introduce the idea of estimating local topology in wireless networks by means of crowdsourced user reports. In this approach each user periodically reports to the serving basestation information about the set of neighbouring basestations observed by the user. We show that, by mapping the local topological structure of the network onto states of increasing knowledge, a crisp mathematical framework can be obtained, which allows in turn for the use of a variety of user mobility models. Using a simplified mobility model we show how obtain useful upper bounds on the expected time for a basestation to gain full knowledge of its local neighbourhood, answering the fundamental question about which classes of network deployments can effectively benefit from a crowdsourcing approach.

I. INTRODUCTION

In modern wireless networks there is an increasing need for Topology Discovery (TD), particularly in networks that are deployed in an unplanned, decentralised manner like WiFi hotspots and 3G/LTE Femtocells. Appropriate knowledge of the network topology allows the design of more efficient routing and interference-avoidance algorithms, improved allocation of limited network resources, etc. In this paper we focus on Local Topology Discovery (LTD), a class of problems where each wireless basestation only aims at the knowledge of its first-hop neighbour basestations.

The first-hop neighbour information provided by LTD is sufficient to allow implementation of distributed algorithms for solving a number of fundamental resource allocation tasks within wireless networks. One example is channel allocation in WiFi networks to minimise interference [1];

another is the allocation of scrambling codes in Femtocell networks, where these codes are used as basestation identifiers and therefore code confusion (neighbouring basestations using the same code) can jeopardise the functionality of the whole network, see [2] and references therein. More generally, first-hop neighbour information is sufficient for both distributed graph colouring and routing, and also for problems of joint optimisation of power and channel allocation.

Importantly, the combination of LTD with distributed algorithms also offers a number of useful benefits:

Incomplete information. The geographic location of small basestations deployed by users (in homes, hotspots etc) is typically not known, nor easily discovered. The network is therefore constrained to work with partial information.

Scalability. Distributed algorithms requiring only local topology information are inherently much more scalable than centralised algorithms or algorithms that require greater topology information.

Adaptability. Wireless networks are time-varying e.g. due to the rollout of new basestations, changes in the radio-propagation environment, changes in the distribution of traffic load, etc. Use of local topology information potentially allows rapid adaptation in response to local changes without a global reconfiguration of the network.

In the present paper we focus on the process of LTD via *crowdsourcing*, meaning that the task of detecting and reporting the existence of conflicting neighbours is delegated to users. Each user periodically reports to the serving basestation information about the set of neighbouring basestations observed. For example, in 3G-femtocell networks users would report both ID and scrambling code of the basestations that lie within the handover range, as described in [3].

We address the fundamental question about which classes of network deployments can effectively benefit from a crowdsourcing approach, estimating the expected time for a basestation to gain knowledge of its local neighbourhood.

The requirements on this value can be more or less stringent depending on the application; for example, in joint power and channel allocation problems, a very quick adaptability is required, because the basestation needs to re-evaluate its neighbourhood every time a change of power (and thus of coverage area) occurs. On the other hand, the local-neighbourhood estimation-time requirement is less tight in the case of channel allocation because the network is already deployed and only few node change in a relatively large time scale.

The use of crowdsourcing for LTD is appealing because it is easy to implement and virtually cost-free. For example, in 3G-femtocell networks it is already the case that users periodically report the list of basestations within handover range, but this information is currently disregarded by the serving basestation [4, §7.4.1]. However, to the best of our knowledge, no previous work has addressed the use of crowdsourcing for LTD in wireless networks.

Our main contributions are the following: (i) the problem of user-reports-based LTD is stated for the first time through a crisp mathematical formulation; (ii) the local topological structure of the network is mapped onto states of increasing knowledge, and in this way the use of a general user mobility model is naturally allowed for; (iii) we concentrate on a specific, simplified mobility model that is useful for gaining insight into those situations where crowdsourcing via user reports is likely to yield the greatest benefit; (iv) under certain conditions, that simplified model is shown to provide an upper bound on the time of topology discovery, thus it can be used as a design tool (see Section III-E).

A. Related Work

TD has been investigated because of its applications to geographical position discovery [5], routing protocols problems [6], and ad-hoc networks configuration in general [7].

Although some papers on wireless optimisation (e. g. [1]) try to relax, at the expense of performance, the hypothesis that each basestation has local knowledge of the network topology, many decentralised schemes do assume it [8–11]. Our work investigates when that assumption holds if a crowdsourcing approach is taken on, while the subsequent step of devising an optimisation scheme falls out of the scope of this project.

Crowdsourcing approaches have been investigated for different applications, e. g. for estimating both density and number of attendees of large events [12]. However, to the best of our knowledge, the effects of crowdsourcing on neighbours discovery capability have never been studied in the literature so far. This function is already available on commercial femtocells [3] and its implementation for interference reduction is recommended in [2] and [13].

II. LOCAL TOPOLOGY DISCOVERY MODEL

Given a set of wireless basestations $\mathcal{A} = \{a_0, \dots, a_N\}$, let $A(a_i) \subset \mathbb{R}^2$ denote the coverage area of access point a_i . Please note that $A(a_i)$ generally depends on the transmission power of

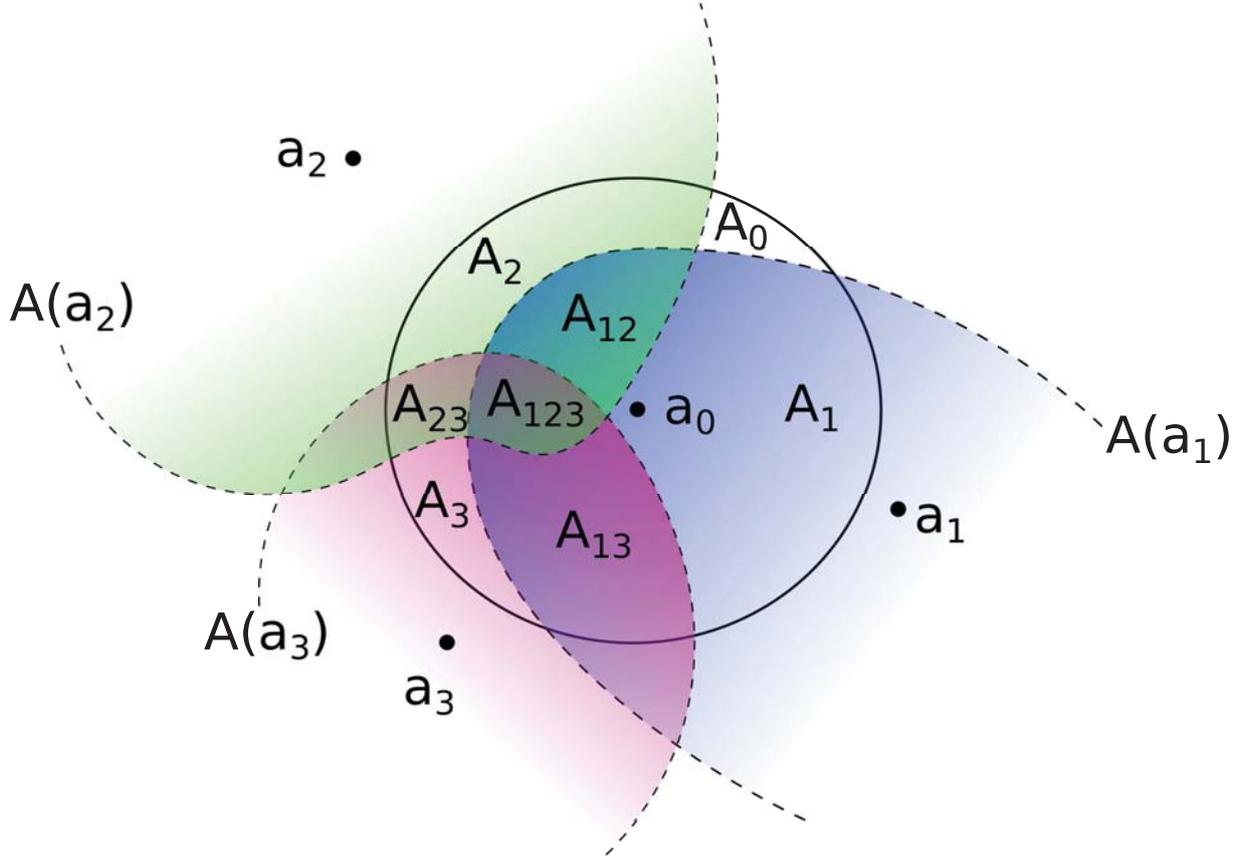


Figure 1: Example of a scenario in which the access point a_0 has three interfering neighbours: a_1, a_2, a_3 . The coverage area $A(a_0)$ can be tessellated with the sets $A_1, A_2, A_3, A_{12}, A_{13}, A_{23}, A_{123}$.

a_i and on the radio propagation properties of the medium. We focus on serving access point a_0 and let \mathcal{B} denote the neighbouring basestations that have non-void intersection with $A(a_0)$, i. e.

$$\mathcal{B} = \{a_i \in \mathcal{A}, i > 0: A(a_0) \cap A(a_i) \neq \emptyset\}.$$

We will hereafter use the symbol N to denote the cardinality of \mathcal{B} , i. e. $N = |\mathcal{B}|$.

Let $\mathcal{P}(\mathcal{B})$ denote the powerset of \mathcal{B} . A *tessellation* of the area $A(a_0)$ is the collection of tiles $\{A_i\}_{i \in \mathcal{P}(\mathcal{B})}$ such that

$$A(a_0) = \bigcup_{i \in \mathcal{P}(\mathcal{B})} A_i, \quad (1)$$

where

$$\begin{aligned} A_{\mathbf{i}} &= \bigcap_{j \in \mathbf{i}} A(a_j) \cap A(a_0), \quad \mathbf{i} \neq \emptyset, \\ A_{\emptyset} &= A(a_0) \setminus \bigcup_{\mathbf{i} \in \mathcal{P}(\mathcal{B}) \setminus \emptyset} A_{\mathbf{i}}. \end{aligned} \quad (2)$$

In what follows each element A_j composing the tessellation is referred to as a *tile*, and we will use the vector notation \mathbf{j} to represent a set of neighbouring basestations. Let us consider for example $\mathbf{i} = \{a_1, a_2\}$; then, the tile A_j is the portion of $A(a_0)$ that is covered by a_1 and a_2 only, see Figure 1. For simplicity of notation, we will write $A_0 := A_{\emptyset}$.

Whenever a user is in A_j , it will *report* \mathbf{j} to access point a_0 . In other words, a_0 will be aware of the existence of those neighbouring basestations $a_i \in \mathbf{j}$. The rate of these reports depends on the mobility model assumed (see Section III).

To keep the model as conservative as possible, and to encompass the frequent case of half-duplex basestations, we assume a_0 cannot detect the existence of any neighbour even though a_0 lies in one of the neighbours coverage area.

Let \mathcal{K}_t denote the *knowledge set* of access point a_0 at time t , i. e. the set of neighbours that a_0 is aware of. Given a sequence of reports $\{\mathbf{j}^1, \dots, \mathbf{j}^t\}$, we have that $\mathcal{K}_t = \bigcup_{i=1}^t \mathbf{j}^i$. \mathcal{K}_t is a growing set, i. e. $|\mathcal{K}_t|$ is non-decreasing in t . Clearly, the knowledge state at time t , \mathcal{K}_t , take values in $\mathcal{P}(\mathcal{B})$.

Definition 1 (Full Knowledge). Given an integer T and a finite sequence of reports $\{\mathbf{j}^1, \dots, \mathbf{j}^T\}$, the basestation a_0 is said to have *Full Knowledge* (FK) of its neighbours at time T if

$$\mathcal{K}_T = \bigcup_{s=1}^T \mathbf{j}^s = \mathcal{B}.$$

Remark 1. If a_0 has Full Knowledge (FK) of its neighbours at time T , so it has at all times $T + t$ for $t \geq 0$. In other words, once a_0 has reached FK, it cannot lose it.

Definition 2 (First time of FK). Given a sequence of reports $\{\mathbf{j}^1, \mathbf{j}^2, \dots\}$, the *first time of FK* τ for the basestation a_0 is the first time the latter reaches FK of its neighbours, i. e.

$$\tau := \min\{T \geq 0 \text{ such that } \mathcal{K}_T = \mathcal{B}\}. \quad (3)$$

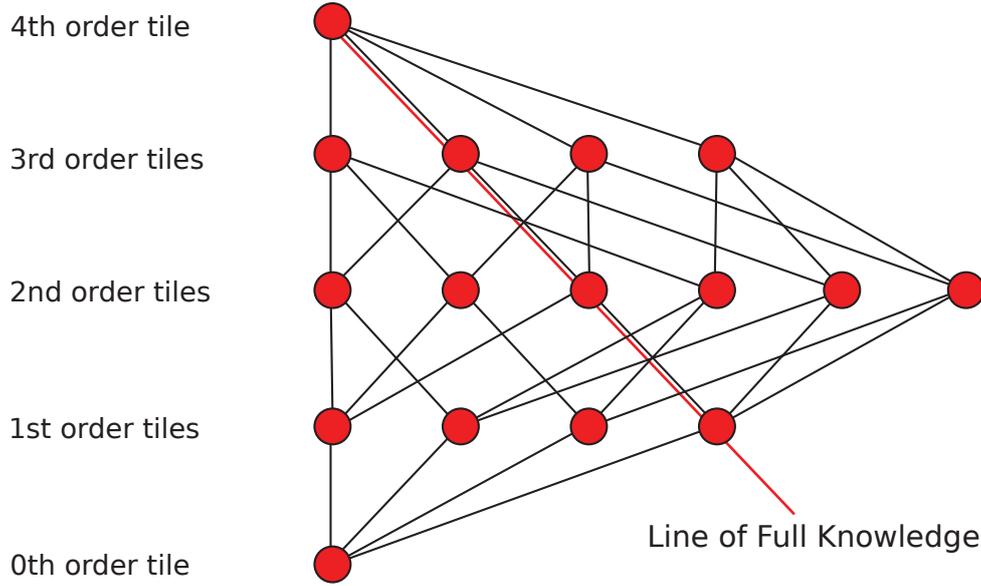


Figure 2: Hypercube representation of the tessellation for $N = 4$. There is one zeroth order tile, namely A^C , four first order tiles, A_1, A_2, A_3 and A_4 , six second order tiles, $A_{12}, A_{13}, A_{14}, A_{23}, A_{24}, A_{34}$, four third order tiles, $A_{123}, A_{124}, A_{134}, A_{234}$ and a fourth order tile, i.e. A_{1234} .

Remark 2. The characterisation of the first time of FK generally depend on the realisation of a sequence of user reports; this means that τ is a random variable. More precisely, by (3), τ is a *stopping time*, see e. g. [14].

We end the section with a note on the tessellation:

Remark 3. A generic tessellation of \mathcal{B} can be represented as an hypercube $H = \{0, 1\}^N$ by identifying the vertices of H with the tiles A_j that the tessellation is composed of. The number of tiles of a generic tessellation of \mathcal{B} is 2^N as well as the vertices of an hypercube, represented as vectors of size N . The tiles of the tessellation can be mapped onto the vertices of the hypercube by identifying the i -th component of the vertices $x \in H$ with $a_i \in \mathcal{B}$. In other words,

$$A_j \leftrightarrow x \quad \Leftrightarrow \quad x_i = \mathbb{1}_{\{a_i \in j\}}, \quad i = 1, \dots, N,$$

where $\mathbb{1}$ is the indicator function. We define the *order* of a tile as the number of neighbours a report from that tile would give knowledge of; the number of k -th order tiles is $\binom{N}{k}$. A report from a k -th order tile is equivalent to k first-order reports. In particular, FK is attained with a

report from the N -th order tile, or at least two reports from two distinct $(N - 1)$ -th tiles, etc. This property can be graphically represented by what we call the *Line of Full Knowledge*, see Figure 2. The line of FK is clearly not unique¹; the aim of Figure 2 is only to illustrate that a sequence of T reports $\{\mathbf{j}^1, \dots, \mathbf{j}^T\}$ is a path on the hypercube H , and that FK is attained whenever a line of FK is reached at a time smaller than T .

Since \mathcal{K}_t , the knowledge state at time t , takes on values in the same set $\mathcal{P}(\mathcal{B})$, we can map also the knowledge states on the hypercube H . In other words, a sequence of reports $\{\mathbf{j}^1, \mathbf{j}^2, \dots, \mathbf{j}^t\}$ is equivalent to a single report from tile $\bigcup_{s=1}^t \mathbf{j}^s = \mathcal{K}_t$.

We can now define the main problems of this work.

Problem 1 (Expected first time of Full Knowledge). *Given an access point a_0 , a set \mathcal{B} of neighbours with given position and coverage area, and a sequence of user reports, we want to characterise the expectation of the first time of FK, i. e.*

$$\mathbb{E}(\tau) = \sum_{t \geq 1} t \mathbb{P}(\tau = t).$$

Obviously, the way the user(s) moves inside the coverage area $A(a_0)$ heavily affects the difficulty of the problem and its answer. However, the formulation of Problem 1 has the great advantage of decoupling the notion of FK from the user mobility model; addressing the mean value of the first time of FK is also an enabler to the estimate of the tail of the distribution of τ – through Markov’s inequality, for example. Further, from a numerical point of view, the expected time of FK may be achieved via a Monte Carlo simulation once the set \mathcal{B} and the mobility model in use are fixed.

There may exist situations where we are content to characterise the first time in which only partial knowledge of the local topology is attained. For example, we may be interested in the first moment when the neighbouring basestations that have been already discovered, i. e. the elements of the knowledge set \mathcal{K}_t , are enough to describe a given fraction of the local topology. This idea motivates the following

Problem 2 (Expected first time of δ -knowledge). *Let ϱ be a measure over $\mathcal{P}(\mathcal{B})$ and fixed $\delta \in (0, 1]$. Given an access point a_0 , a set \mathcal{B} of neighbours with given position and coverage*

¹For example, there are N tiles of order $N - 1$, but only 2 are part of a given line of FK.

area, and a sequence of user reports, we want to characterise the expectation of the first time of δ -knowledge $\mathbb{E}(\tau_\delta)$, where

$$\tau_\delta = \min \left\{ T \geq 0 \text{ such that } \frac{\sum_{k \in \mathcal{P}(\mathcal{K}_T)} \varrho(A_k)}{\sum_{j \in \mathcal{P}(\mathcal{B})} \varrho(A_j)} \geq \delta \right\}.$$

When $\delta = 1$ and $\varrho(A_j) > 0$ for each $j \in \mathcal{P}(\mathcal{B})$, Problem 2 is equivalent to Problem 1. Indeed, $\frac{\sum_{k \in \mathcal{P}(\mathcal{K}_T)} \varrho(A_k)}{\sum_{j \in \mathcal{P}(\mathcal{B})} \varrho(A_j)} \geq 1$ if and only if $\mathcal{K}_T \equiv \mathcal{B}$.

We will hereafter consider the measure $\varrho(A_k) = |A_k|$. This leads to the following interpretation: δ -knowledge is attained when the knowledge set \mathcal{K}_t defines for the first time a tessellation that covers a fraction of $A(a_0)$ larger or equal than δ . Equivalently, τ_δ is the first time when the tiles that would give new information² cover a fraction of $A(a_0)$ that is smaller than $1 - \delta$.

Remark 4. The concept of δ -knowledge is fundamental in the simulation phase, when we want to know whether user reports can effectively be used to give knowledge of the local topology. Indeed, it is likely that the neighbours a_i whose coverage area do not overlap with $A(a_0)$ save for a nearly negligible portion, will be discovered after a very long time; in other words, the leading contribution to $\mathbb{E}(\tau)$ will be represented by the mean first visit time of the user(s) to $A(a_i)$. Discarding a_i from the picture, the concept of δ -knowledge let us focus on the quantitative analysis of LTD, see Section IV.

III. TELEPORT MOBILITY

The characterisation of τ , the first time of FK, depends on the users mobility model that is assumed. The latter describes how users enter, exit, and move within $A(a_0)$.

The users evolution can then be represented as a pair $U_t = (n_t, X_t)$ where n_t is the number of users that lie in $A(a_0)$ at time t , and $X_t = (x_t^1, x_t^2, \dots, x_t^{n_t})$ is a vector with the position of the n_t users. We assume the evolution of U_t to be driven by a discrete-time Markov chain (MC) throughout the paper.

The realisation of $\{U_t\}_{0 \leq t \leq T}$ completely determines the sequence of user reports $\{\mathbf{j}^1, \dots, \mathbf{j}^T\}$ to the access point a_0 , cf. Remark 2. Since \mathcal{K}_t only depends on \mathcal{K}_{t-1} and U_t , then the bivariate process (U_t, \mathcal{K}_t) is a MC.

²In the sense that the cardinality of the knowledge set \mathcal{K}_T would increase.

It will prove useful to consider a simplified mobility model in which a single user can instantaneously teleport to any tile:

Model 1. (Teleport Mobility) A single user moves within $A(a_0)$ according to a discrete-time MC taking on values in $\mathcal{P}(\mathcal{B})$. The user cannot exit $A(a_0)$ and no other user can enter it. At each step, the user instantaneously teleports with a probability that is proportional to the measure of the destination tile. In formulas, the transition probabilities are

$$\mathbb{P}(i \rightarrow j) = \frac{|A_j|}{A(a_0)}. \quad (4)$$

Remark 5. Model 1 greatly simplifies the characterisation of τ , the first time of FK. Indeed, with this mobility model \mathcal{K}_t is independent of U_t , and the sole process \mathcal{K}_t is hence sufficient to describe the process of gathering knowledge from the user reports. We will hereafter refer to \mathcal{K}_t as the *knowledge chain*.

Assuming Model 1, we can easily describe the process of gathering knowledge from user reports as a discrete-time random walk on the hypercube $H = \{0, 1\}^N$ (which we have introduced in Remark 3); having knowledge of n neighbouring basestations is in fact equivalent to receiving a report from the n -th order tile that give information about all of them.

Let $P(\cdot, \cdot)$ the transition kernel of the knowledge chain. If $\mathbf{k} \not\subseteq \mathbf{l}$, then $P(\mathbf{k}, \mathbf{l}) = 0$ because such transition would mean a loss of knowledge; in other words, $|\mathcal{K}_t|$ is non-decreasing as a function of t . Conversely, when $\mathbf{k} \subseteq \mathbf{l}$, a transition from \mathbf{k} to \mathbf{l} happens if the user moves to a tile that contains the missing information ($\mathbf{l} \setminus \mathbf{k}$) and do not add more than that information. Therefore,

$$P(\mathbf{k}, \mathbf{l}) = \begin{cases} \sum_{m \in \mathcal{P}(\mathbf{k})} \frac{|A_{\{m \cup (\mathbf{l} \setminus \mathbf{k})\}}|}{A(a_0)} & \text{if } \mathbf{k} \subseteq \mathbf{l}, \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

The following result holds:

Lemma 1. The matrix P is upper triangular.

Proof: Let us consider the following partial ordering relation among the states:

$$\mathbf{k} \preceq \mathbf{l} \quad \Leftrightarrow \quad \mathbf{k} \subseteq \mathbf{l}.$$

By (5), $P(\mathbf{k}, \mathbf{l}) \neq 0$ only if $\mathbf{k} \preceq \mathbf{l}$. Therefore, any mapping

$$\mathcal{P}(\mathcal{B}) \ni \mathbf{l} \quad \Leftrightarrow \quad \mathbf{l} \in \{1, 2, \dots, 2^N\}$$

such that

$$\mathbf{k} \preceq \mathbf{l} \quad \Leftrightarrow \quad k \leq l$$

will put the matrix P into an upper triangular form. In particular, we can order the states by increasing cardinality and in lexicographic order³. ■

The explicit computation of the whole matrix P using (5) is expensive in general – P is a $2^N \times 2^N$ matrix! However, as stated above P is upper triangular, while in Section III-C we show that it is possible to explicitly characterise its spectrum. For the reader's reference, Table I shows the matrix P for $N = 3$.

$$P = \frac{1}{A(a_0)} \begin{pmatrix} |A_0| & |A_1| & |A_2| & |A_3| & |A_{12}| & |A_{13}| & |A_{23}| & |A_{123}| \\ 0 & |A_0| + |A_1| & 0 & 0 & |A_2| + |A_{12}| & |A_3| + |A_{13}| & 0 & |A_{23}| + |A_{123}| \\ 0 & 0 & |A_0| + |A_2| & 0 & |A_1| + |A_{12}| & 0 & |A_3| + |A_{23}| & |A_{13}| + |A_{123}| \\ 0 & 0 & 0 & |A_0| + |A_3| & 0 & |A_{13}| + |A_1| & |A_2| + |A_{23}| & |A_{12}| + |A_{123}| \\ 0 & 0 & 0 & 0 & |A_0| + |A_1| + |A_2| + |A_{12}| & 0 & 0 & |A_3| + |A_{13}| + |A_{23}| + |A_{123}| \\ 0 & 0 & 0 & 0 & 0 & |A_0| + |A_1| + |A_3| + |A_{13}| & 0 & |A_2| + |A_{12}| + |A_{23}| + |A_{123}| \\ 0 & 0 & 0 & 0 & 0 & 0 & |A_0| + |A_2| + |A_3| + |A_{23}| & |A_1| + |A_{12}| + |A_{13}| + |A_{123}| \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

Table I: Example of transition matrix P for $N = 3$.

A. Expected Time of Full Knowledge

Let $\mathbf{k}^* = \{1, 2, \dots, N\}$ be the state of FK. By formula (5), $P(\mathbf{k}^*, \mathbf{k}^*) = 1$. This means that the chain has an absorbing state, and the hitting time of this state is just τ , the first time of FK. Hence, we can compute the expected time of FK simply by

$$\mathbb{E}[\tau] = (I - Q)^{-1} \mathbf{1}, \quad (6)$$

³For N neighbouring basestations, i.e. with 2^N different tiles, this would mean the sequence $\{1\}, \{2\}, \dots, \{N\}, \{1, 2\}, \dots, \{N-1, N\}, \dots, \{1, 2, \dots, N\}$.

where Q is obtained from P by removing the row and the column relative to state \mathbf{k}^* and $\mathbf{1}$ is the column vector of ones [15]. In a similar way it is possible to compute the other moments of τ .

Even if $I - Q$ is upper triangular and can be block decomposed, the computation of its inverse may not be affordable when the cardinality of \mathcal{B} grows. In Section III-D we will bound the probability of the event $\{\tau > t\}$.

B. Expected Time of δ -knowledge

Regarding Problem 2, we can easily modify matrix P to obtain the expected time of δ -knowledge. Every state $\mathbf{k} \in \mathcal{P}(\mathcal{B})$ such that

$$\sum_{\mathbf{l} \in \mathcal{P}(\mathbf{k})} \frac{|A_{\mathbf{l}}|}{A(a_0)} \geq \delta$$

can be aggregated in the absorbing state, summing the corresponding column of P in the last column, and then eliminating the column and row corresponding to state \mathbf{k} . In this way it is possible to compute $\mathbb{E}[\tau_\delta]$ using (6).

C. Eigenvalues

The following result fully characterises the spectrum of the matrix P :

Theorem 2. For $\mathbf{k} \in \mathcal{P}(\mathcal{B})$, the eigenvalues of P have the form

$$\lambda_{\mathbf{k}} = \frac{1}{A(a_0)} \left(|A_0| + \sum_{\substack{\mathbf{l} \subset \mathbf{k} \\ |\mathbf{l}|=1}} |A_{\mathbf{l}}| + \cdots + \sum_{\substack{\mathbf{l} \subset \mathbf{k} \\ |\mathbf{l}|=m}} |A_{\mathbf{l}}| + \cdots + \sum_{\substack{\mathbf{l} \subset \mathbf{k} \\ |\mathbf{l}|=|\mathbf{k}|}} |A_{\mathbf{l}}| \right).$$

Proof: The matrix P being upper triangular by Lemma 1, the entries $P(\mathbf{k}, \mathbf{k})$ are the eigenvalues of the matrix. Let us then imagine to have the knowledge chain in state \mathbf{k} . The only way for the chain to undergo a self-transition ($\mathbf{k} \rightarrow \mathbf{k}$) is that the user reports any combination of neighbouring basestations that have already been discovered. In other words, the knowledge chain undergoes a self-transition if and only if the user reports an element of $\mathcal{P}(\mathbf{k})$. Therefore,

$$\lambda_{\mathbf{k}} = \frac{1}{A(a_0)} \sum_{\mathbf{l} \in \mathcal{P}(\mathbf{k})} |A_{\mathbf{l}}|.$$

Last formula is equivalent to the thesis. ■

Since each eigenvalue is a sum of positive elements, the second-largest eigenvalue $\tilde{\lambda}$ can be obtained by maximising over the tiles of order $N - 1$:

$$\tilde{\lambda} = \max_{\mathbf{k}: |\mathbf{k}|=N-1} \lambda_{\mathbf{k}}. \quad (7)$$

D. Convergence Properties, Bounds

Using (7), it is possible to obtain the following result:

Lemma 3. Given $\varepsilon > 0$, let

$$S(1 - \varepsilon) = \frac{\log \varepsilon}{\log \tilde{\lambda}}. \quad (8)$$

Then, $S(1 - \varepsilon)$ reports are sufficient to achieve FK with probability greater or equal than $(1 - \varepsilon)$.

Proof: Using Theorem 1 and Equation (7) on P ,

$$\mathbb{P}(\tau > t) \leq \tilde{\lambda}^t. \quad (9)$$

For a small target tolerance ε of not achieving FK,

$$\text{if } t \geq \frac{\log \varepsilon}{\log \tilde{\lambda}} \quad \Rightarrow \quad \mathbb{P}(\tau > t) \leq \varepsilon. \quad (10)$$

■

δ -knowledge Convergence Bounds: Using the same manipulation of the matrix P described in Section III-B, Corollary 3 can be applied to the modified matrix to obtain a bound for the number of steps to have δ -knowledge with high probability.

E. Numerical Estimation of the Time of δ -knowledge

Model 1 is equivalent to a single user teleporting instantaneously to a random point within the coverage area of the basestation; time is discrete. Thus, at each time the basestation a_0 receives a user report from a point drawn according to the uniform probability distribution over the coverage area $A(a_0)$. Having in mind the numerical characterisation of the first time of δ -knowledge, the teleport model is particularly convenient. This task could be in fact carried out within the Monte Carlo paradigm by simply throwing sufficiently many points at random inside the coverage area $A(a_0)$. In other words, it is possible to numerically study the process through which δ -knowledge is achieved by sampling sufficiently many times a probability density function that is uniform over the coverage area $A(a_0)$.

Model 1 may prove itself unsatisfactory in a real life scenario. The main problem is that if we generate a sequence of user reports according to it, any two elements of the sequence are independent, whereas in general they are not. In each mobility model where the trajectory taken by the user is physically feasible, the user positions communicated by two successive reports are in fact correlated due to the motion constraints.

Let us imagine that a single user travels inside the coverage area $A(a_0)$ according to an unknown mobility model, and let $T(t)$ be the trajectory taken by the user. Sampling the trajectory at equally-spaced discrete times, we obtain an embedded sequence of user location, which correspond to an embedded sequence of user reports. Next, we can analyse the sequence and understand after how many steps δ -knowledge has been reached. By multiplying this number of steps by the time lapse between two consecutive reports (inter-report time), the time of δ -knowledge can be obtained for that particular realisation of the user-reports sequence. Finally, the procedure above can be repeated sufficiently many times to estimate with a Monte Carlo method the expected time necessary for a_0 to reach δ -knowledge.

As mentioned above, in a general mobility model it is likely that two successive user reports are correlated. These correlations may decay as the inter-report time grows larger and larger. As an example, let us imagine that a single user travels inside the coverage area $A(a_0)$ according to a MC. Let π be the equilibrium probability measure of the chain and let t_{mix} be the mixing time of the chain, i. e. the time needed for the chain to reach equilibrium. If the inter-report time is chosen comparable to t_{mix} then the time lapse between two successive reports will be sufficient for the MC to forget the past trajectory; in other words, the correlations between consecutive reports will be negligible. As a consequence, the user locations will be independently drawn from the probability measure π , and the matrix P describing the knowledge evolution will become

$$P(\mathbf{k}, \mathbf{l}) = \begin{cases} \sum_{m \in \mathcal{P}(\mathbf{k})} \pi(A_{\{m \cup (\mathbf{l} \setminus \mathbf{k})\}}) & \text{if } \mathbf{k} \subseteq \mathbf{l}, \\ 0 & \text{otherwise.} \end{cases} \quad (5')$$

Therefore, the formulation and the results developed in Sections III-A–III-D are still valid if we consider a single-user mobility model based on a MC, provided that the time lapse between two consecutive reports is of the order of the mixing time of the chain. Under the assumption that user reports are sent at a frequency comparable with the inverse mixing time of the mobility MC, we can compute an upper bound on the time of δ -knowledge. Any reporting rate higher

than $\frac{1}{t_{\text{mix}}}$ will in fact still guarantee that a_0 achieves δ -knowledge of its neighbourhood in at most $\mathbb{E}[\tau_\delta] \cdot t_{\text{mix}}$ seconds on average⁴.

We end this section by briefly mentioning a straightforward application of Model 1 in a multi-user scenario. Let us imagine that n users may enter, move within, and exit $A(a_0)$ according to a hidden mobility model. We assume that n is a very large number and that it is possible to statistically characterise the stationary user-density by means of a probability measure π over $A(a_0)$. At each time every user may independently send a report with a very small probability p . Then, the number of report received by a_0 in a given time interval is approximately Poissonian and the time lapse between two successive report is exponential with parameter $\lambda = np$. Next, let $m = \mathbb{E}[\tau]$ be the expected time of FK, expressed in number of reports, returned by (5') and (6); the expected time to achieve FK is the expectation of the first time for a Poisson process of parameter λ to hit the state m . We will come back to this scenario in Section V.

IV. SIMULATIONS

A. Teleport Model on Random Positioned basestations

In this section we offer a preliminary assessment of the possibility to use the machinery developed so far in real applications. To this purpose, we developed a simulation framework in MATLAB and studied a scenario where 8 basestations are positioned in a plane uniformly at random. Each basestation has a circular coverage area of the same size. We considered 350 different configurations, with the constraint that the coverage area of a_0 has non-void intersection with the coverage area of the remaining basestations, meaning that FK is achieved as soon as all 7 neighbours are reported to a_0 . We compute the tessellation of each configuration using a classical Monte Carlo sampler. For each of these 350 configurations, we computed the expected time of 0.9-knowledge $\mathbb{E}[\tau]$ together with the number of steps sufficient to guarantee 0.9-knowledge with 90% confidence, i. e. $S(0.9)$. The inter-report time being fixed during this first experiment, the amount of time in seconds to achieve 0.9-knowledge is directly proportional to the number of steps just evaluated.

Figure 3 displays the empirical probability mass function of these two quantities. $\mathbb{E}[\tau]$ is centered around 10 steps, while $S(0.9)$ is shifted on higher values, as expected being an upper

⁴Recall that $\mathbb{E}[\tau_\delta]$ is measured in number of reports.

bound. Figure 4 shows the empirical cumulative distribution function of $\mathbb{E}[\tau]$ and $S(0.9)$. We see that 16 steps are sufficient to achieve 0.9-knowledge for nearly all scenarios (95 %), while we need 22 steps using $S(0.9)$. We also notice that the bound obtained from (8) is a conservative estimation, because it uses only the second-largest eigenvalue $\tilde{\lambda}$. Indeed, it takes into account only the slowest way to reach the desired knowledge, while the problem has a rich combinatorial structure that cannot be completely captured by (9).

B. Random Walk on a Grid

In order to investigate and confirm the ideas of Section III-E, we simulated the reports sent with different inter-report times by a random walker that moves within $A(a_0)$ under the condition of reflective boundary, and compared this mobility model with Model 1 (see Section III) for a set of 8 basestations positioned as described at the beginning of this section.

In Figure 5 we let the inter-report time increase and compare the average time to achieve 0.9-knowledge according to both the random walk (green line) and the teleport model (blue). We see that that, if the inter-report time is sufficiently large, the empirical mean time to achieve 0.9-knowledge for the random walk model is well approximated by that of Model 1.

We assume typical femtocell parameters, i. e. that coverage radius is 50 m and that the user do a step in a grid of 2.5 m every 5 s. Figure 5 also shows that when reports are sent each 6 min or less, the time of 0.9-knowledge is smaller than 1 h, but at such high frequency the bound $S(0.9)$ (red line) is not valid anymore. The reason why more reports than Model 1 are needed in the case of high-frequency reports is the following: since the inter-report time is short, it is likely that many reports will be sent from the same tile, i. e. the knowledge chain will undergo many self-transitions.

It is important to notice that the inter-report time used in Figure 5 are far from the theoretical order of magnitude of the random walk mixing time. Yet Figure 5 suggests that, for a family of scenarios, it should be possible to determine the value of the inter-report time such that the average time to achieve 0.9-knowledge may be well predicted by Model 1. Once that value of the inter-report time is found, the value of $\mathbb{E}[\tau_{0.9}]$ returned by Model 1 may serve as an upper bound to the actual time to achieve 0.9-knowledge when smaller inter-report times are implemented.

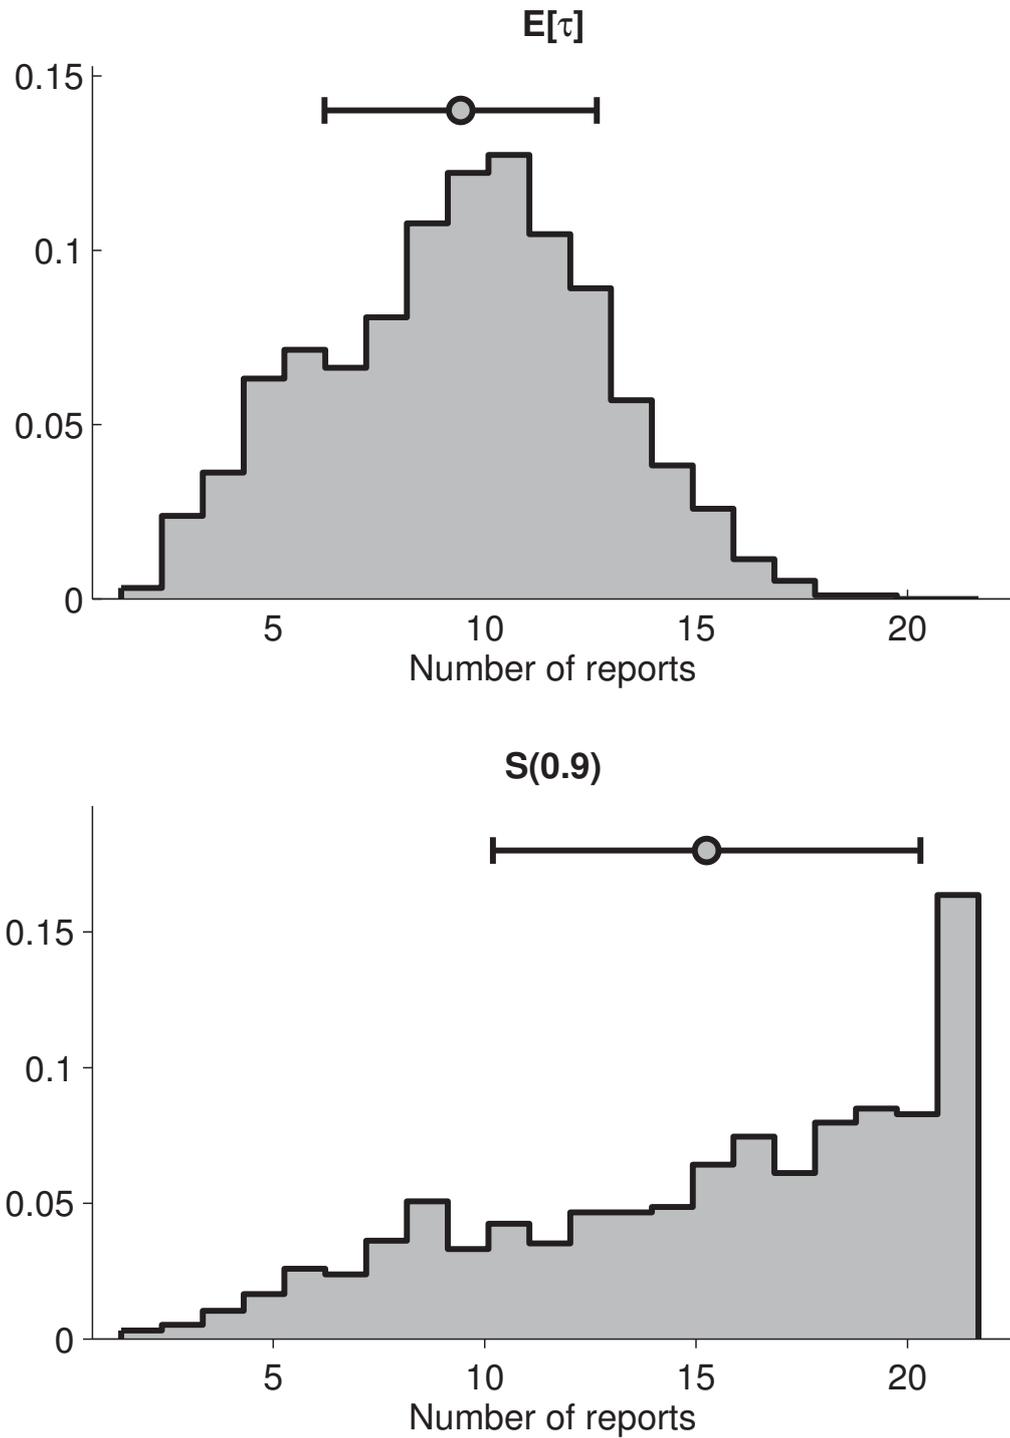


Figure 3: Empirical probability mass function of the expected time of 0.9-knowledge, and the number of steps to have 0.9-knowledge with 90% confidence, for the teleport model on random positioned basestations. Since the inter-report time is fixed, the simulation time is directly proportional to the number of reports.

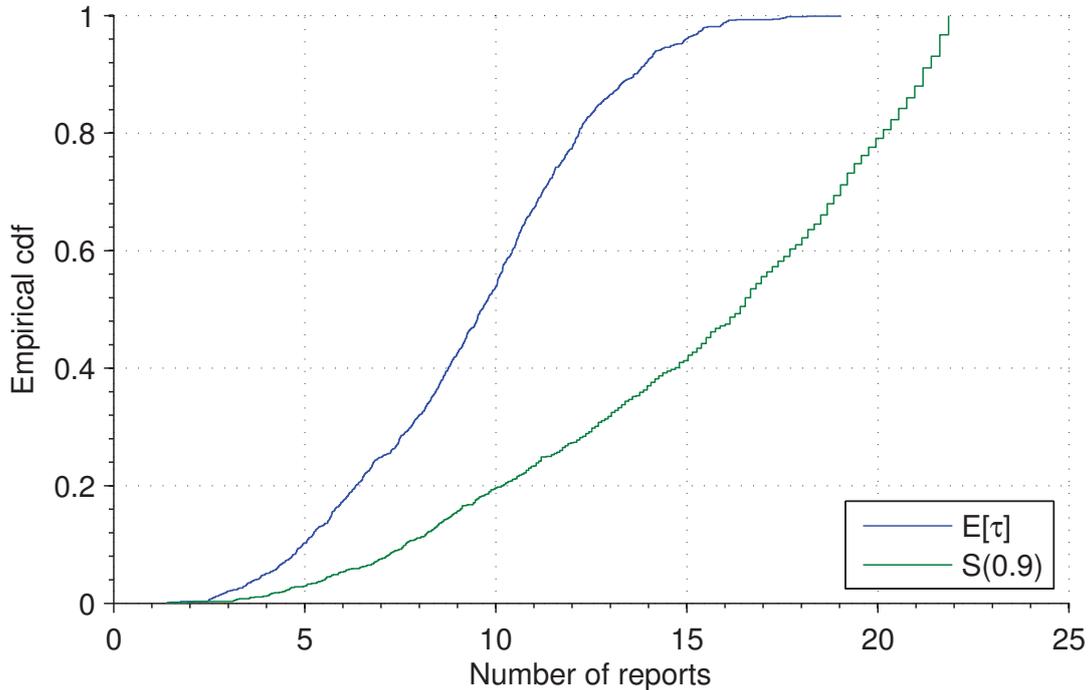


Figure 4: Empirical cumulative distribution function of the expected time of 0.9-knowledge, and the number of steps to have 0.9-knowledge with 90% confidence, for the teleport model on random positioned basestations. Since the inter-report time is fixed, the simulation time is directly proportional to the number of reports.

C. A Realistic Scenario

A received power map for 4 basestations in the Hynes convention centre have been generated using the Wireless System Engineering (WiSE) [16] software, a comprehensive 3D ray tracing based simulation package developed by Bell Laboratories. basestations are assumed transmitting at a frequency of 2.1 GHz with a power of 34 mW. We assume there is a macrocell that covers the whole building, and we estimate its time of full knowledge. As before, a Monte Carlo simulation has been made to estimate the tessellation, and then the expected time of δ -knowledge has been computed using a teleport mobility model (Model 1), as explained in Section III-B.

Figure 6 shows the corresponding coverage areas when powers detection threshold is -70 dBm. Although the shape of the coverage areas and their intersection is much more complex than the simple scenario depicted in Section IV-A, it is still possible to construct the tessellation by considering which coverage areas each spatial point lies in. For example, point a lies in the

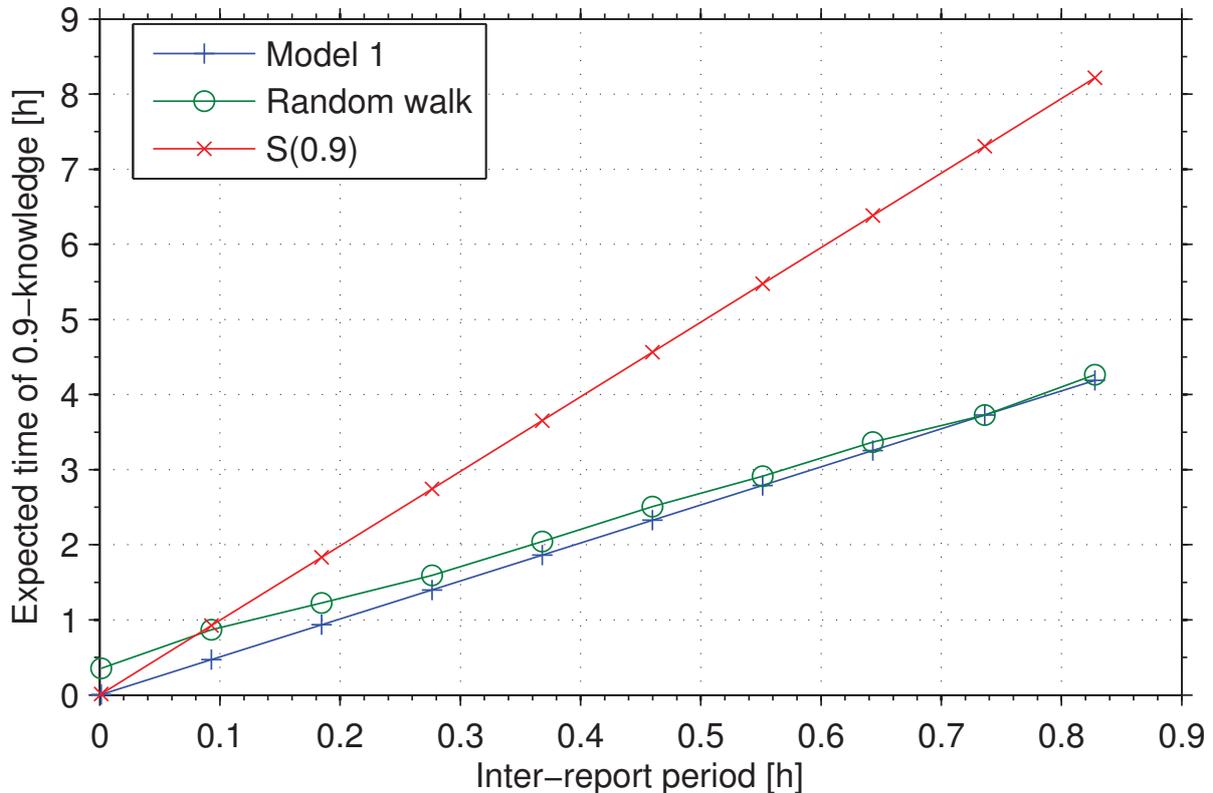


Figure 5: Empirical mean 0.9-knowledge time (in hours) of a random walk vs. inter-report period, compared with Model 1 and its bound $S(0.9)$ in a femtocell grid of 8 basestations.

coverage area of basestations 1, 2, and 3, so it belongs to the tile A_{123} .

Figure 7 displays the expected time of δ -knowledge, $\mathbb{E}[\tau_\delta]$ when δ is varied. We notice a step-function-like behaviour, with a new step that is added every time a new state become absorbing, as explained in Section III-B.

Figure 8 shows the behaviour of $\mathbb{E}[\tau]$, the expected time of FK, when the user detection threshold varies from a very conservative value of -60 dBm to a more realistic one of -100 dBm. When the users are more sensitive, the coverage areas, and the higher order tiles in particular, are bigger, leading to better performances. In particular, we see that an average of 14 steps are enough to achieve FK.

These results confirm that the values obtained placing random basestation with circular coverage in Section IV-A are compatible with real world scenarios.

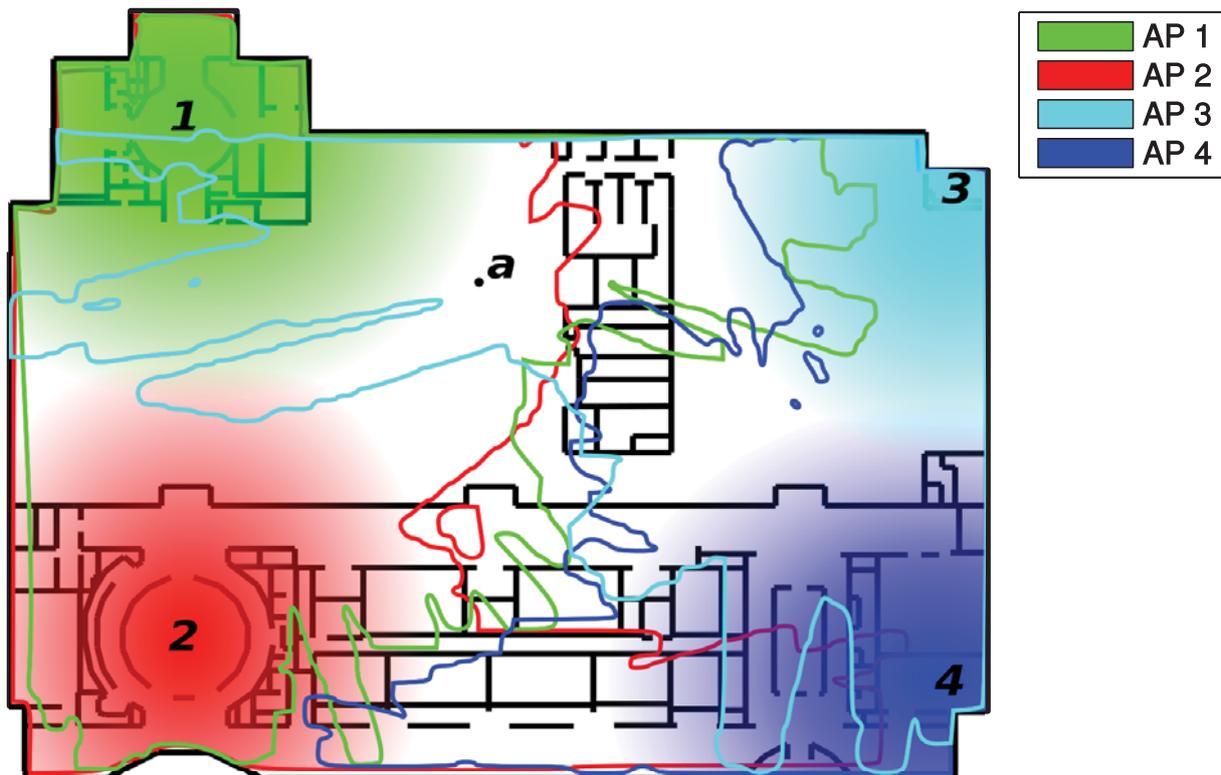


Figure 6: Coverage areas at Hynes convention centre. The coloured lines delimit the extension of the coverage areas. Basestations are transmitting at 2.1 GHz with a power of 34 mW. Point a lies in the tile A_{123} .

V. USE CASES

In this section we present some use cases that are representative of the possible practical applications of the results we have presented so far. Given a particular wireless application, the focus is on whether user reports can effectively be used to efficiently achieve knowledge of the network local topology.

In the case of femtocell deployment for residential use, each basestation typically serves a very small number of devices. Using data of typical residential densities and coverage areas, a statistic of the tessellation can be devised. If it is possible to establish a time \tilde{T} after which the user position can be considered as drawn from a uniform distribution, then $S(\delta)$ is an upper bound of the time of δ -knowledge for all the inter-report times smaller than or equal to \tilde{T} .

Opposite to the previous example, femtocells deployed in congested places like a mall have an extremely large basin of potential users. However, in situations where users main interest is other

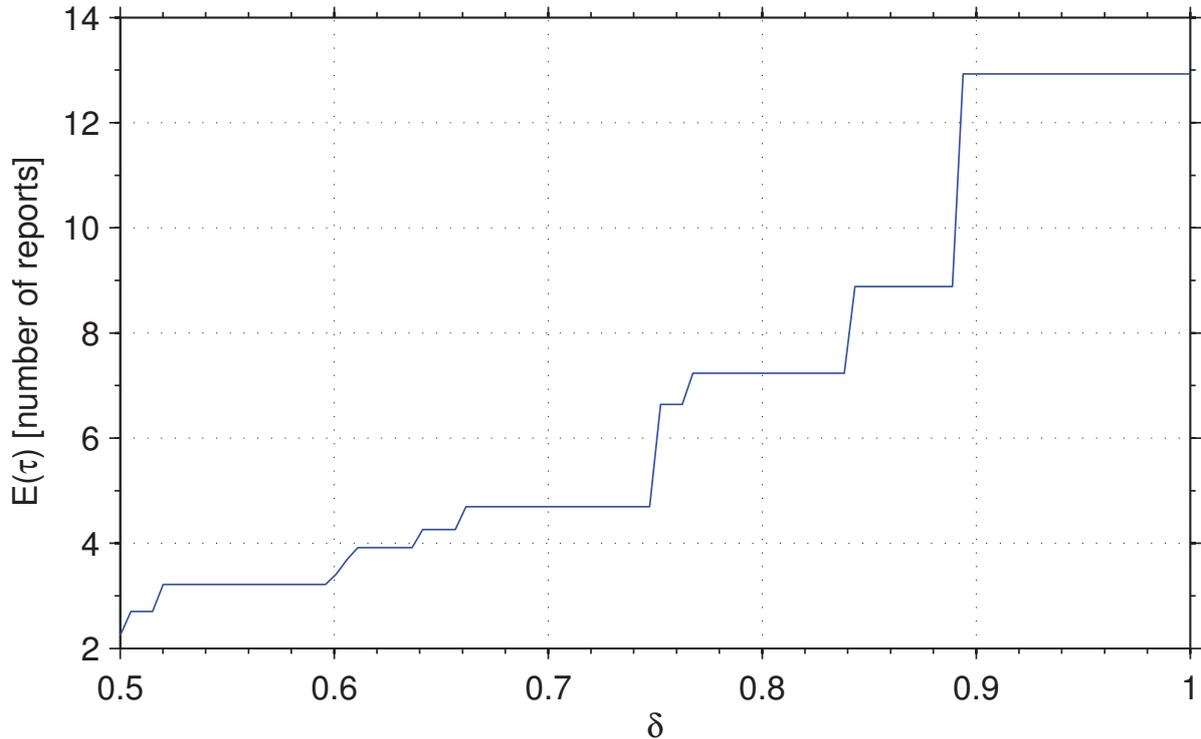


Figure 7: Expected time of δ -knowledge, $\mathbb{E}[\tau_\delta]$ using Model 1 in Hynes convention centre for different values of the parameter δ . The inter-report time is set to the very same value used for Figures 3 and 4.

than connecting to the internet, it is reasonable expecting the single-user reporting-activity to be rather sporadic. Therefore, the Poissonian approximation that we have mentioned at the end of Section III-E may be applicable. In these case, characterising the time to achieve δ -knowledge is possible through a statistic of the typical tessellations.

VI. CONCLUSIONS

We introduced the problem of user-reports-based Local Topology Discovery, providing a crisp mathematical formulation of it in the case of Model 1.

We showed that Model 1 can effectively be used as an upper bound for a wide range of mobility models when the user reports frequency is lower than the inverse mixing time of the MC of the actual mobility model (in practice it can be used if reports are sent every hour).

In Section III-D we provide an useful method to estimate the time of δ -knowledge when the problem is too big to solve exactly using Equation (6).

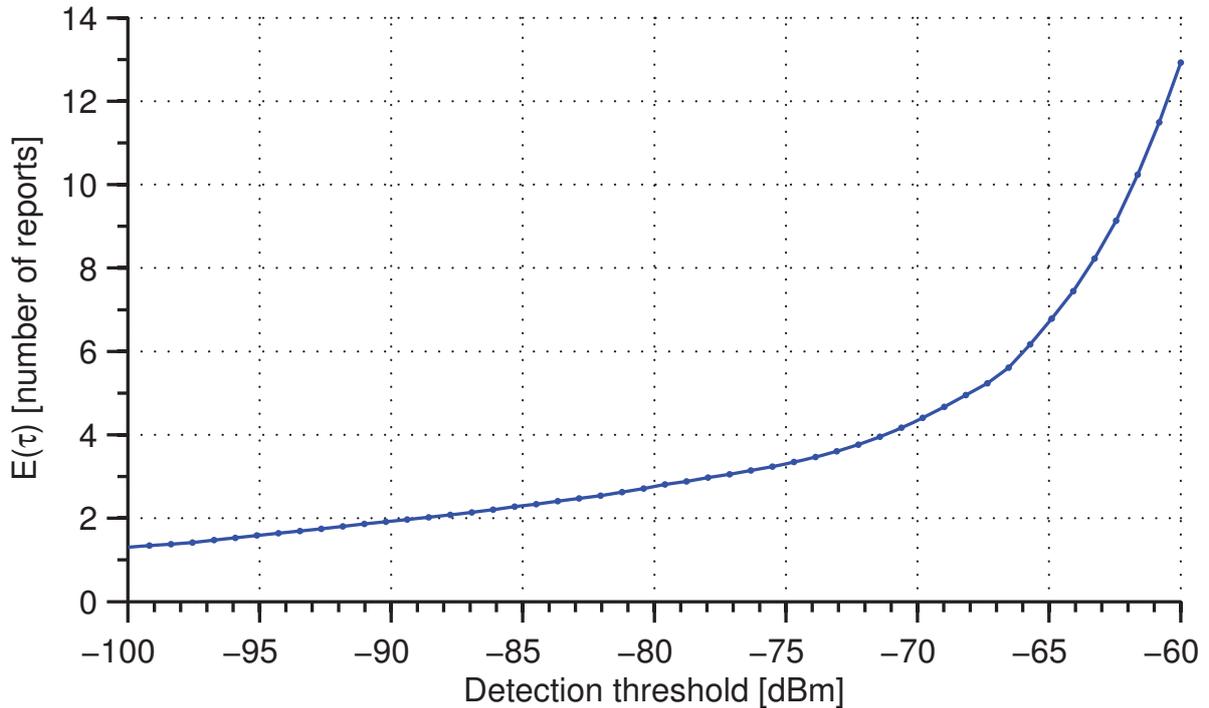


Figure 8: Expected time of FK, $\mathbb{E}[\tau]$, using Model 1 in Hynes convention centre, for different values of the user-detection threshold. The inter-report time is set to the very same value used for Figures 3 and 4.

Simulations on random scenarios with typical femtocell parameters show that the expected number of reports in order to have an high degree of knowledge of the local topology is very small. Roughly speaking, a user moving at 0.5m/s according to a random walk model, and providing a report every hour, can guarantee the basestation will have 0.9-knowledge with high probability in less than 5 h, and in less than two hours if reports are sent every 15 minutes. Since the local topology is not typically expected to change every day, this are acceptable times.

The simulations on more realistic scenarios (Section IV-C) give very similar results in term of time of 0.9-knowledge.

These preliminary results encourage to implement the user reports function, corroborating the heuristic recommendations in [13] and [2]. In the case of femtocells, such implementation should be easy, because the hardware and the firmware are already capable of managing user reports.

A. Future Work

A more extensive study on more realistic scenarios is required, where the typical topological properties of a urban area are taken in account (see Section V).

Similarly, an analysis of more realistic mobility models is desirable: our work encompass the simple case of Model 1, that can be used to estimate any other Markovian model (i. e. any mobility model that can be described with a Markov process) with unique stationary measure only when the report frequency is sufficiently low, namely slower than the mixing time of the Markov Chain. An analysis of the behaviour of the mobility models during their transient behaviour is left for future work.

REFERENCES

- [1] A. Checco and D. J. Leith, "Learning-based constraint satisfaction with sensing restrictions," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, pp. 811–820, 2013.
- [2] A. Checco, R. Razavi, D. J. Leith, and H. Claussen, "Self-configuration of scrambling codes for WCDMA small cell networks," in *IEEE 23rd International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, Sydney, Australia, September 2012.
- [3] P. Sapiano, "Discovering neighbouring femto cells," US Patent EP2 214 434, August, 2010. [Online]. Available: <http://www.freepatentsonline.com/EP2214434A1.html>
- [4] Technical Specification Group Radio Access Network, "Home Node B (HNB) Radio Frequency (RF) requirements (FDD)," 3rd Generation Partnership Project, TR 25.967, 2011.
- [5] R. Choudhury, S. Bandyopadhyay, and K. Paul, "A mobile agent based mechanism to discover geographical positions of nodes in ad hoc wireless networks," in *6th Asia-Pacific Conference on Communications (APCC2000)*, Seoul, Korea, vol. 30, 2000.
- [6] S. Marwaha, C. Tham, and D. Srinivasan, "A novel routing protocol using mobile agents and reactive route discovery for ad hoc wireless networks," in *Networks, 2002. ICON 2002. 10th IEEE International Conference on*. IEEE, 2002, pp. 311–316.
- [7] R. Choudhury, S. Bandyopadhyay, and K. Paul, "A distributed mechanism for topology discovery in ad hoc wireless networks using mobile agents," in *Proceedings of the 1st ACM international symposium on Mobile ad hoc networking & computing*. IEEE Press, 2000, pp. 145–146.

- [8] F. Kuhn and R. Wattenhofer, “On the complexity of distributed graph coloring,” in *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*. ACM, 2006, pp. 7–15.
- [9] Ö. Johansson, “Simple distributed $\Delta + 1$ -coloring of graphs,” *Information Processing Letters*, vol. 70, no. 5, pp. 229–232, 1999.
- [10] M. Luby, “Removing randomness in parallel computation without a processor penalty,” in *Foundations of Computer Science, 1988., 29th Annual Symposium on*. IEEE, 1988, pp. 162–173.
- [11] M. Szegedy and S. Vishwanathan, “Locality based graph coloring,” in *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*. ACM, 1993, pp. 201–207.
- [12] F. Movahedi Naini, O. Dousse, P. Thiran, and M. Vetterli, “Opportunistic Sampling for Joint Population Size and Density Estimation,” *IEEE Trans. Mobile Comp.*, 2013.
- [13] J. Edwards, “Implementation of network listen modem for WCDMA femtocell,” in *Cognitive Radio and Software Defined Radios: Technologies and Techniques, 2008 IET Seminar on*. IET, 2008, pp. 1–4.
- [14] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*. AMS, 2009.
- [15] F. Gehring and P. Halmos, *Finite Markov Chains*. Springer-Verlag, 1976, ch. 3: Absorbing Markov Chains, p. 224.
- [16] S. Fortune, D. Gay, B. Kernighan, O. Landron, R. Valenzuela, and M. Wright, “WiSE design of indoor wireless systems: practical computation and optimization,” *IEEE Comput. Sci. Eng.*, vol. 2, no. 1, pp. 58–68, 1995.