Does Mobility Matter?

Ólafur Helgason, Sylvia T. Kouyoumdjieva and Gunnar Karlsson
ACCESS Linnaeus Centre
KTH, Electrical Engineering
Stockholm, Sweden
{olafur.helgason, stkou, gk}@ee.kth.se

***Abstract*—In modern society, wireless devices are commonly carried by humans. The wireless communication is therefore affected by pedestrian mobility in urban outdoor and indoor spaces which is the scenario we consider in this work. Many of the mobility models currently used for evaluating wireless communication systems have poor resemblance to reality. Although advances have recently been made, there is still a lack of understanding on which elements of mobility affect system performance. In the civil-engineering field of transport and urban planning there exist advanced pedestrian mobility models, used for designing and dimensioning public spaces for pedestrian crowds and emergency evacuation. These models capture micro-mobility of pedestrians better than most mobility models used in mobile networking since the application domain requires that they realistically capture node interactions with its physical environment as well as other nodes. In this work we use Legion Studio, a commercial simulator, to explore which elements of pedestrian mobility are important with respect to system performance and how sensitive the connectivity metrics of nodes are to input mobility parameters. These studies give insight into whether relatively simple mobility models suffice for evaluating wireless systems. Furthermore, they contribute to our understanding of which parameters are important for modelling mobility and the accuracy in which these parameters need to be estimated to give dependable results.**

I. Introduction

It is well known that mobility significantly affects the performance of wireless communication systems [1], [2]. On the one hand, mobility causes disruptions when connected ad-hoc networks are partitioned as nodes move out of communication range. On the other hand, mobility can be exploited to forward messages across network partitions by having nodes physically store, carry and forward messages [3]. Mobility is therefore a crucial component of these systems and it is essential to realistically capture mobility when evaluating their performance.

In many wireless systems the mobile terminals are carried by humans and this is the scenario that we target for; pedestrians in urban areas. Much of the performance evaluation of these types of systems uses synthetic mobility models where nodes randomly move in a closed area. These random waypoint/direction models however reflect reality poorly and therefore a series of models have been proposed that address some of the shortcomings [4], [5], [6], [7], [8]. Although advances have recently been made towards more realistic mobility models, there are still many elements of human mobility that are ignored or inadequately captured in current models. As examples of this, they commonly assume a free flow of nodes and ignore node-to-node interactions. The structure of space in which mobility occurs is either not considered or only in a very limited way. They do not consider cultural aspects of mobility such as differences in personal space requirements and walking speed, and in contrast to real-life they usually assume a closed-system and thus do not consider the effects of arrival processes and sojourn times in the area under inspection. We feel that there is still a lack of understanding of which elements of mobility are important when it comes to evaluating wireless communication systems, in particular with respect to the sensitivity of system performance to individual mobility parameters.

In the area of urban planning and transportation research, advanced mobility models for flows of pedestrians as well as for individual pedestrians in a flow have been devised [9]. These models are primarily used for designing and dimensioning large-scale public spaces in order to optimize for large crowds, emergency and evacuation strategies. It is thus of primary importance to realistically capture the structure of the physical space, how nodes interact with it and how they interact with each other. These models therefore better capture micro-mobility of pedestrians than most mobility models used in mobile networking. We believe that these models can be of great use for evaluating the performance of mobile communication and we are convinced that they can give insight into which elements of mobility affect performance.

In this paper we study how mobility affects the performance of wireless communication systems. Our approach is based on a detailed and realistic micro-level mobility model implemented in Legion Studio[1]; a commercial pedestrian simulator package that is commonly used by architects and civil engineers to design and dimension large-scale public spaces. The mobility model in Legion Studio is based on advanced analytical models that have been calibrated by measurement studies [10]. To the best of our knowledge, Legion Studio is the most advanced and realistic simulation model available for micro-level pedestrian mobility and it has, as far as we know, not been used before in the context of evaluating mobile communication systems.

The focus of our current work is on how micro-level mobility parameters (such as speed, arrival process and personal space requirements) affect the connectivity parameters (such as contact duration, contact rate and number of neighbours) of

[1] http://www.legion.com

mobile wireless systems. In particular, the goal of our study is to explore the performance sensitivity of mobile communication with respect to the micro-level mobility metrics and we want to address the following:

- Is connectivity highly sensitive to even small changes in input mobility parameters?
- Are some input parameters more important than others? If so, which?

We believe that these are fundamental questions to ask when it comes to modelling mobility for the sake of evaluating wireless communication systems. They contribute to understanding whether mobility can be captured by relatively simple analytic or simulation models to give meaningful performance results, or if other more advanced measures are needed. Also, they help us understand whether a rough estimate of input parameters is sufficient or if accurate estimates of the empirical distributions are needed.

The rest of this paper is organized as follows. In section II we further motivate and position our work with respect to the work of others. Section III gives an overview description of Legion Studio, the mobility simulator used in our studies, and section IV describes our simulation settings and results. In section V we conclude and discuss future work.

II. Background & Related Work

Human mobility can be classified as consisting of three levels: strategic, tactical and operational mobility [9]. At the strategic level humans decide their daily movement patterns and the activities they would like to perform, such as go to work or take a walk in the park. The tactical level, considers the implementation of a strategic decision, such as choosing a way of travel or a shortest path, taking into consideration environmental factors like obstacles and congestion. At the operational level, the physical process of human movement is considered, including walking speed, physical size of nodes and interaction with other traffic due to queuing or to avoid collisions. The performance of mobile wireless communication systems is likely to be affected in a different way by each level. Decisions at the strategic and tactical levels determine the regularity and routines in moving patterns which in turn affect the inter-meeting time with other individual nodes. Some routing protocols for delay-tolerant networks try to take advantage of such non-randomness in node mobility patterns to efficiently route messages to a given destination node [11], [12]. Mobility at the operational level presumably affects node connectivity and the duration of individual contacts while on the move. This determines the amount of data that can be transferred over each contact which is of great importance when forwarding packets between nodes according to a routing protocol or when disseminating data for content distribution applications [13], [14], [15].

A series of works have identified that mobility can significantly affect the performance of wireless communication systems [16], [1], [2]. By simulating various ad-hoc routing protocols under different mobility models, the work in [1] confirms that the choice of mobility model matters and that performance ranking of protocols may vary with the mobility model used. Our work differs from [1] in that we focus strictly on a certain scenario, namely pedestrian mobility. Also, currently we do not focus on particular applications or (ad-hoc) routing protocols but are more interested in how sensitive the node connectivity metrics are to variations in mobility input parameters at the operational level.

In the recent past, random mobility models (such as random waypoint [17] or random direction [18]) have commonly been used to evaluate mobile wireless systems. These models have a poor resemblance with real life mobility for various reasons:

- At the operational level, they assume a free flow of nodes and therefore do not capture interactions between nodes, queuing effects or node interactions with the physical neighborhood and with obstacles.
- At the strategic level, they neither capture the social clustering of humans nor regularities in human mobility such as due to working day cycles or social networks.
- They assume a closed system and therefore they do not allow for assessing the effect of arrival processes and sojourn times in an area.

Recently a number of mobility models have been proposed to address some of these shortcomings. The Manhattan mobility model in [1] captures obstacles by assuming that nodes are confined to streets and it models pedestrian interactions by assuming that fast nodes slow down when blocked by a slower node. It is however 1-dimensional in the sense that it does not assume that nodes have physical size nor that streets have width. In contrast, the Legion mobility model we use is a much more detailed model of pedestrian behaviour than the Manhattan model. In Legion Studio, as in real-life, fast nodes may slow momentarily down due to slow nodes, but if space allows, they will overtake when there is opportunity. The obstacle mobility model [5] also captures movements of nodes where obstacles block or restrict their freedom to roam or their signal propagation. It does however not capture detailed operational-level pedestrian mobility. Another mobility model that focuses on pedestrians in urban areas is proposed in [19]. It assumes that pedestrians can be modelled by a finite number of flows where each node belongs to one particular flow. They formulate the flows as a linear programming optimization problem, and by measuring the densities of pedestrians at different observation points in the graph topology they can solve the optimization problem to obtain the rate of each flow. This model is thus not a detailed operational level model, but rather models tactical mobility. Some mobility models do not explicitly model the physical movements of nodes but instead model the temporal clustering due to regularity in their movements. This clustering and non-randomness in human mobility stems from mobility at the strategic level and it strongly affects the node inter-meeting times which is particularly important for routing in delay-tolerant networks, as previously discussed. These models can thus be classified as purely strategic or tactical level models and they can be based on social networks [6], [8] or diurnal regularity in human

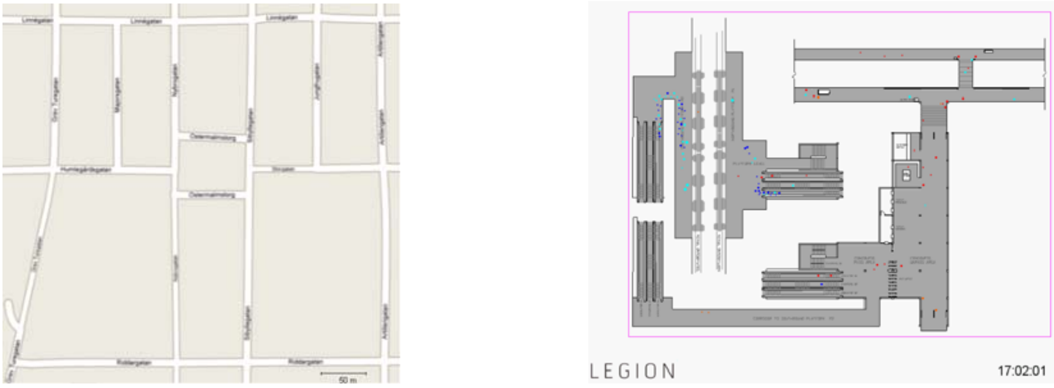

(a) (b)

Fig. 1. The simulation scenarios: A part of a downtown Stockholm (left) and a two-level subway station (right).

movements [7], [20]. Although advances have recently been made with more realistic human mobility models there is still a lack of understanding which parameters affect the performance of mobile communications, not least at the operational level.

Recently there has been significant interest in measuring the connectivity of mobile devices [2], [21], [22]. Connectivity traces are collected from mobile devices that are distributed to a set of users in a controlled group (e.g. conference participants, a sample of students on campus). These studies are mostly analyzing the frequency of transfer opportunities by studying the distribution of inter-meeting times between devices. Less attention has been paid to studying the individual contacts and what affects their duration and frequency. Although these experiments have shown interesting results they are severely limited in many ways. The traces inherently suffer from limitations with the respective radio technology employed which can make it hard to draw general conclusions about the effects of mobility. As an example of technical limitations, the Bluetooth devices frequently used in measurements perform neighbour discovery approximately every 2 minutes. This is likely to filter out many of the short contact opportunities as we address in the current work. Also, performing large-scale measurements is expensive and resource demanding and still each experiment only captures a particular scenario. Reproducing the environment between individual experiments to single out the effect of one factor is difficult because of issues such as variations in the radio channel, mobility and randomness in the MAC layer contention. We therefore conclude that although collecting traces is important, simulations are, and will remain, an invaluable tool for evaluating the performance of mobile networks.

III. Legion Studio mobility simulator

Legion Studio is a commercial simulation software package used by architects and civil engineers for designing large-scale public spaces via simulation of pedestrian behavior. Legion Studio can import AutoCAD drawings of real life structures and buildings in urban areas such as subway stations or shopping malls. The multi-agent pedestrian model is based on advanced analytical and empirical models [23] which have been calibrated by measurement studies [10]. Each pedestrian is represented as a two-dimensional entity with a circular form and a size that approximates the size of an actual pedestrian. Most of the navigation decisions can be made by the entities themselves based on the individual interaction with other nodes and with the surrounding environment, although overriding certain choices is possible. The movement patterns follow the *least effort* principle where, just like in reality, each entity tries to minimize the dissatisfaction before choosing its next move. It should be noted that dissatisfaction in Legion Studio can have different aspects, including inconvenience, frustration and spatial discomfort. The three levels at which an entity can take its navigation decision, strategic, tactical (macro-navigation) and operational (micro-navigation), allow correct capturing of issues such as queuing behind slower nodes or at bottlenecks, as well as personal space requirements. Dynamics increase because Legion Studio allows the use of open systems, where entities can enter or leave the system according to a certain pattern. The ability of registering the spatial and temporal position of each node provides extensive information that can be used for examining user behaviors in the context of evaluating mobile communication systems.

In summary, Legion Studio offers a more sophisticated way of simulating operational-level pedestrian mobility than most current mobility models in the field of communication networks. For a more detailed description of the simulation model we refer to [23].

IV. Evaluation

Evaluation scenarios

Our evaluation considers two scenarios with different characteristics. The first scenario is an outdoor urban scenario that models the Östermalm area of central Stockholm, shown in Fig. 1(a). It consists of a grid of interconnected streets where each street is 2 m wide and lengths vary between 20 m and 200 m. There are 12 passages that connect the area to the outside world and we assume that all streets have equal node arrival rates denoted by λ. Upon arriving at an intersection, nodes continue to move on the same street (if possible) with

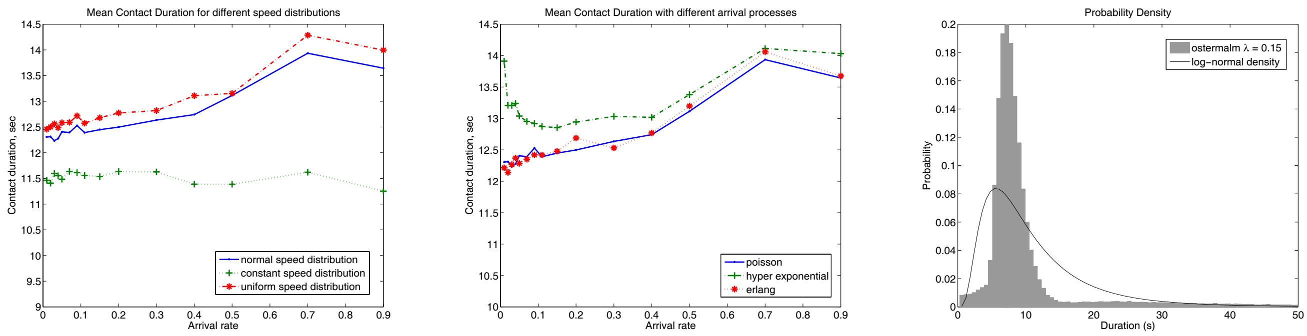

(a) (b) (c)

Fig. 2. The impact of speed distribution (2(a)) on and arrival process (2(b)) on mean contact duration. 2(c) shows the histogram of the contact duration fitted log-normal pdf.

probability 0.5 or turn to other adjoining streets with equal probabilities. The scenario is relatively large with an active area of approximately 5872 m^2 that the nodes can move in. Nodes in the scenario move constantly so it is characterized by high mobility. We note that we have experimented with different routing probabilities and replaced the center streets with a wide square and found that this does not significantly affect our result.

The second scenario is a two-level indoor subway station, shown in Fig. 1(b). The train platform is connected to the entry-level with escalators and nodes can arrive either in batches from an incoming train or on foot through one of the doors on the entry-level. The train arrivals/departures thus contribute to burstiness in the node arrival and departure processes. Nodes will congregate at the platforms when waiting for the next train arrival or in the store or restrooms at the entry level. Since the station is relatively densely populated with people and since its structure severely constrains mobility, the physical interaction of nodes is high, which in turn can lead to queuing. The escalators are the main bottlenecks where queuing effects are observed. The active area of the scenario is about 1921 m^2.

The mobile nodes may communicate over short-range radio, such as Bluetooth or 802.11. We assume a simple model of the physical layer in which nodes connect if they are within a transmission range Δ of one another. Currently we do not consider radio or data-link issues such as interference, shadowing, fading or MAC-layer contention. In dense pedestrian scenarios these issues are certainly important if a large fraction of the nodes have a communication device. Many of the scenarios considered here however have a relatively low density. Also, if the penetration of active devices in a crowd is low, the mobility of the device-carrying nodes is still constrained by the crowd. This will affect some of the connectivity parameters studied here, in particular the contact duration.

Simulation setup

Each simulation run conducted in Legion Studio results in a mobility trace file, containing a snapshot of the positions of all nodes in the system every 0.6 s. We assume a node transmission range of $\Delta = 10$ m. For obtaining the connectivity metrics from the mobility traces we parse them with custom scripts or feed them into a system simulator [24]. In processing the simulation runs we sample only values after steady state has been reached to avoid any bias due to the initial transient.

Performance metrics

We are interested in evaluating how the connectivity of the nodes is affected by mobility and how sensitive the connectivity metrics are to changes in mobility input parameters. We study both the mean and full empirical distributions of the following connectivity metrics:

- *Contact duration* - This is the time that two nodes are physically within direct communication range. In this work we do not consider a particular radio technology but we note that our definition of contact duration includes the node discovery and contact setup time which can be different for different radios. The amount of data that can be transferred over a contact thus depends the channel bitrate and the remaining contact duration after connection setup.
- *Contact rate* - We calculate the contact rate as the number of contacts of non-zero duration that a node makes, divided by the lifetime of the node in the simulation. Since we study open systems, the contact rate is a more suitable metric than the number of contacts since the rate is (by definition) normalized by the simulation sojourn time of a node.
- *Inter-contact time* - In this work we define the inter-contact time as the elapsed time from the beginning of one contact to the beginning of the next (in some works this is referred to as the inter-any-contact time). We do not consider inter-meeting times (i.e. elapsed time between contacts of the same pair of nodes) in this study since it is mainly determined by mobility at the strategic and tactical levels. The current study focuses on operational-level mobility.

A. Östermalm scenario results

With the Östermalm scenario we consider how the input speed distribution and arrival process affects the connectivity metrics and their distributions. We point out that the input speed distribution only gives the *target speed* of nodes but in the simulator, the actual empirical speed distribution is different since node speed is affected by inter-node dynamics and the physical space (i.e. nodes may slow down due to congestion and they can momentarily speed up when overtaking other slower nodes). Legion Studio also allows one to configure different cultural profiles that modify node parameters such as the physical size, dissatisfaction levels due to queuing or constraints, and desired inter-node distance. We have explored the effect of varying these parameters and found that the connectivity metrics show high insensitivity to changes. Therefore we do not discuss these further.

In Fig. 2 we study the effect of the speed distribution and arrival process on the mean contact duration. In 2(a) we consider three different speed distributions: constant with speed 1.3 m/s, uniform(0.6; 2.0), and truncated normal(0.6; 2.0) with a mean of 1.3 m/s. The reason for choosing 1.3 m/s as a mean speed value is based on the study in [25] and we have found that changing the absolute value only shifts the performance curves by a constant factor but does otherwise not affect performance (graphs are omitted due to space constraints). In order to facilitate the comparison, in all cases a Poisson arrival process is considered. Due to different factors (for example, cultural profiles or aging) people do however not move at constant speed. Therefore the normal and uniform speed distributions are closer to reality and these give similar results, with the uniform distribution giving slightly longer contact durations. The reason is that with the uniform distribution there is a larger amount of slow nodes in the scenario which increases the average contact durations.

Fig. 2(b) studies the effect of different arrival processes on connectivity. All the arrival processes have the same mean λ: a Poisson process with rate λ, 4-stage Erlang where each stage has rate $4\,\lambda$ and a two-phase hyperexponential inter-arrival time distribution with arrival rates 0.35λ and 5.7λ in the first, respectively, second phase, and selection probabilities 0.31 and 0.69. The Erlang distribution has a coefficient of variation $1/2$ and the hyper-exponential distribution has a coefficient of variation 2. The speed distribution in all cases is normal with mean 1.3 m/s. At low arrival rates the hyper exponential arrivals result in longer contact durations than the other processes. This is due to burstiness of the process: bursty periods of short inter-arrival times are followed by longer inter-arrivals. Nodes arriving during a bursty period have a high probability of being connected upon arrival and thus stay connected for a relatively long time as they move in a connected group. We believe that arrival processes in real life are in many cases bursty since common phenomena in urban life contribute to this, such as traffic lights, elevators and arrivals of trains and buses. At higher arrival rates the distinction between arrival processes diminishes and the average contact duration increases. This is due to the clustering of nodes in the scenario which leads to queuing that forces faster nodes to slow down, thus increasing the contact duration with the surrounding nodes.

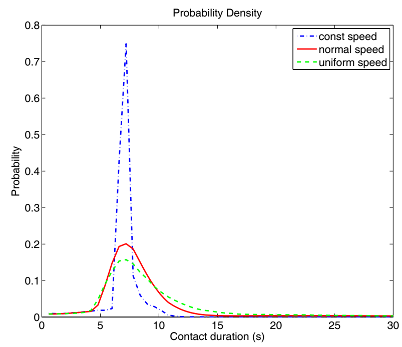

Fig. 3. Empirical pdf of contact duration for three speed distributions.

In addition to studying the mean values we have also looked at how the distribution of the contact duration is affected by the speed and arrivals. Fig. 2(c) shows the histogram of contact durations under a Poisson arrival process with a per-entry arrival rate of $\lambda = 0.15\ s^{-1}$. The histogram shows a sharp peak in the contact durations at approximately 8 s. This is due to contacts with nodes moving in opposite directions. The mean duration of this type of contacts is given by $d = 2\Delta/|v_a - v_b|$ where Δ is the communication range and $|v_a - v_b|$ is the relative speed of the nodes. For $\Delta = 10$ m and a mean speed of 1.3 m/s the average duration of these contacts is 7.7 s. The histograms for the Erlang and Hyper-exponential have high resemblance and are therefore omitted from the plot. In other words, we find that the arrival process does not alter the shape of the contact duration and this is further verified by a statistical comparison. We have compared the empirical contact distribution with five standard probability distributions: exponential, power-law, gamma, weibull and log-normal. Although all distributions are rejected by a Kolmogorov-Smirnov (KS) test, a log-normal distribution with parameters with $\mu = 2.2$ and $\sigma = 0.7$ gives the closest fit (i.e. lowest KS statistics) for all three arrival processes. As seen from Fig. 2(c) the log-normal distributions is however only a rough estimate and it does not well capture the sharp spike at 8 s.

Fig. 3 illustrates how the speed distribution affects the contact duration. The speed affects both the height and the width of the spike due to contacts in opposite directions and this is well reflected in the mean value analysis in Fig. 2(a). The basic shape of the empirical contact distribution however remains the same.

In summary, we see that the speed distribution and arrival processes have a small effect on the mean contact duration. The shape of the underlying empirical distributions for the contact durations however remains the same and they are all roughly approximated by the same log-normal distribution.

Fig. 4(a) shows that up to a certain point ($\lambda = 0.7\ s^{-1}$) the contact rate increases approximately linearly with the arrival rate, and that the absolute value is highly insensitive to the

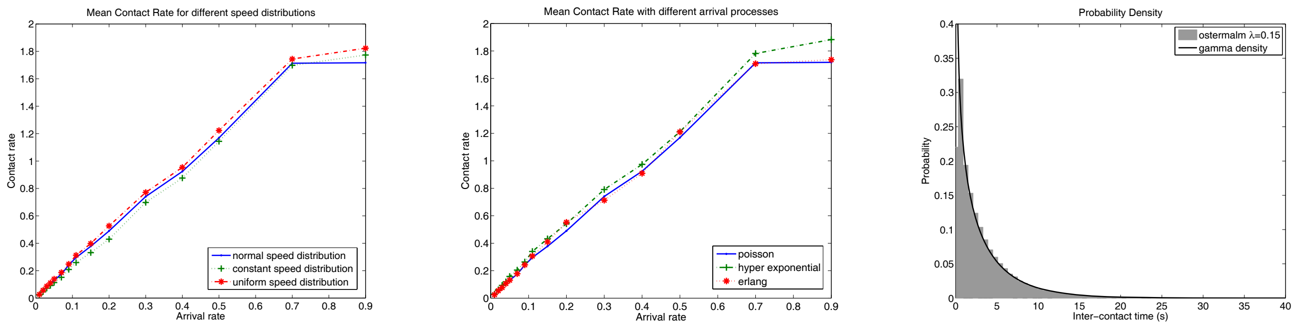

(a) (b) (c)

Fig. 4. The impact of speed distribution (4(a)) and arrival process (4(b)) on mean contact rate. 4(c) shows the histogram of the inter-contact time and a fitted gamma pdf.

shape of the speed distribution. The violation of the linearity, observed in the end of the graphs marks the beginning of a saturation phase. Due to the high density, nodes are severely constrained in their mobility and injecting more nodes into the scenario does barely increase the opportunities for establishing new contacts. The same tendency can be observed in Fig. 4(b) and the graph demonstrates that changing the shape of the arrival processes result in minor differences in the mean value of the contact rate.

We have found that the empirical distribution of the inter-contact times is only slightly affected by the different arrival processes. The inter-contact times are approximated by a gamma distribution for all arrival processes and the empirical inter-contact means and gamma parameters are listed in table I. Furthermore, Fig. 4(c) shows the histogram of the empirical inter-contact time distribution for a Poisson arrival rate of $\lambda = 0.15$ s along with the fitted gamma probability density (the gamma densities and histograms for Erlang and Hyper-exponential are omitted since they cannot be separated from the Poisson case).

Similar to the contact duration, the mean contact rate and the inter contact time distribution show high insensitivity with respect to changes both in the walking speed distribution and the arrival process. We find that the inter-contact distribution is well approximated by the gamma distribution while the log-normal fit for the contact durations should only be considered a rough estimate.

B. Subway scenario results

As previously mentioned, mobility in the subway scenario is different from the Östermalm scenario in that nodes do not constantly move but can pause at the store and the train platforms. Nodes arrive from the entrances and from arriving trains that follow a given schedule in Legion Studio, which we cannot modify. Therefore we do not consider different arrival processes for this scenario but focus on the one given by the scenario. We have also found that the empirical distributions for both contact duration and contact rate are insensitive to the speed distribution (detailed results are omitted due to space constraints).

	Poisson	Erlang	Hyper-exponential
Mean (s)	3.2	2.8	3.0
Gamma params. a (shape) and b (scale)	$a = 0.66$ $b = 4.78$	$a = 0.54$ $b = 5.24$	$a = 0.66$ $b = 4.62$

TABLE I
PARAMETERS FOR THE EMPIRICAL INTER-CONTACT TIME FROM THE ÖSTERMALM SCENARIO WITH POISSON ARRIVALS AND $\lambda = 0.15\ s^{-1}$.

To compare the empirical distributions from the subway and Östermalm scenarios it seems reasonable to select data from an arrival rate for the Östermalm scenario that gives comparable average node density ρ. We have that $\rho = N/A$ where N is the mean number of nodes in the scenario and A is the size of the active area. N can be obtained from Little's law as $N = \lambda_{tot}T$ where T is the mean sojourn time of nodes in the scenario and λ_{tot} is the mean total arrival rate. We can use $\hat{T}$, the measured mean sojourn time from our simulations, to obtain N according to Little's law and therefore the node density. We have found that the mean node density in the subway scenario is approximately 0.09 $nodes/m^2$ and a per-entry arrival rate of $\lambda = 0.15\ s^{-1}$ gives approximately the same node density in the Östermalm scenario. Therefore we compare Östermalm statistics for $\lambda = 0.15\ s^{-1}$ with the subway scenario.

Fig. 5(a) shows the histogram of contact durations. The mean observed contact duration is 18.1 s and we find that a log-normal distribution with parameters $\mu = 2.3$ and $\sigma = 1.1$ gives the best fit among the five standard probability distributions previously mentioned. The histogram indicates that the contact duration distribution exhibits multimodality as we see a small peak at around 1 s, another one at about 8 s and a third peak at approximately 20 s. Compared to the contact duration in the Östermalm scenario (Fig. 2(c)) the histogram for the subway scenario is more spread out with a less intense peak due to nodes moving in opposite directions (at about 8 s) and a thicker tail most likely due to waiting and/or queuing effects. This is further reflected in the mean contact durations: 12.3 s for Östermalm and 18.1 s for the subway. The multimodality of the distribution suggests that the contact duration cannot be accurately modelled by a single standard distribution but should rather be approximated with

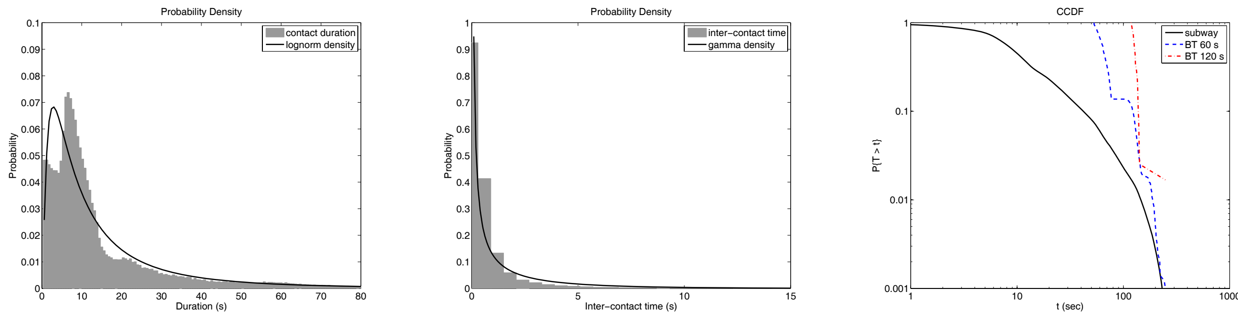

(a) (b) (c)

Fig. 5. 5(a): Histogram of empirical contact duration for the subway scenario and a fitted log-normal pdf. 5(b): Histogram of inter-contact time for subway scenario and a fitted gamma pdf. 5(c): Empirical contact duration tail distribution compared with a Bluetooth radio model.

a mixture of distributions. At present we however leave this modelling work for the future.

The histogram for the inter-contact times is shown in Fig. 5(b). The mean inter-contact time is 1.1 s, compared to a mean of around 3.0 s for the Östermalm scenario as seen in table I. As in the Östermalm scenario, a gamma distribution is found to give the best but with different parameters ($a = 0.2$ (shape) and $b = 5.6$ (scale)). It is noticeable that for both scenarios, the inter-contact times can be reasonably well fitted by a standard gamma distribution while the empirical contact durations are more complex and are therefore not as well fitted by approximations with the standard distributions we use for reference.

C. Comparison with a Bluetooth radio model

As discussed in Section II, a number of experiments have recently been conducted to explore the connectivity of mobile devices carried by humans. Most of these experiments are conducted by distributing Bluetooth devices to a group of users. The devices periodically scan their neighbourhood for other devices and the contact statistics are logged.

Bluetooth uses a frequency hopping spread spectrum scheme and due to this, scanning for neighbours is an operation that consumes significant power and prevents normal data flow. A Bluetooth node is discoverable when it is in the *inquiry* state and when it performs a neighbour scan it enters the *inquiry scan* state in which it is itself not discoverable by other scanning nodes. For these reasons neighbour scanning is usually only performed periodically (sometimes with a randomized period to avoid synchronization between nodes). In contrast, our simulation model assumes a radio that is always on and detects all nodes within communication range. To compare with measurement traces and to understand how Bluetooth affects the node connectivity of our model we have therefore simulated the neighbour discovery of Bluetooth and calibrated it based on parameters used in measurement studies. We simulate the Bluetooth node discovery as follows: Each node alternates between a *inquiry scan* state and an *inquiry* state. During inquiry scan, the node discovers its current direct neighbours and neighbours that are also in the inquiry scan state are not discovered. The identities of all nodes seen during the inquiry scan are stored and those nodes that are also neighbours in the next inquiry scan are classified as contacts and the contact duration is updated accordingly. A contact with a neighbour is completed if the neighbour is not seen in two consecutive scans; this definition of the contact duration is consistent with the measurements previously mentioned [22].

In Fig. 5(c) we plot the empirical tail distribution of the contact durations for the subway scenario with our previous radio model and compare with results using a Bluetooth radio. The duration of the inquiry scan is 10 s and the time interval between inquiry scans is 60 s and 120 s respectively, the latter being representative for the setup in measurement studies. The time between scans is further randomized by adding or subtracting (with equal probability) 0 to 12 seconds uniformly selected at random.

The effect of superimposing the Bluetooth neighbour discovery on our simple radio model is significant. In the original subway scenario we measured 534592 contacts in total. Using Bluetooth the total number of contacts is reduced to 2932 for a sleep period of 60 ± 12 s and only 69 contacts when the sleep period is 120 ± 12 s. This indicates clearly that Bluetooth is very inefficient at exploiting the short contacts that arise when people are on the move. For opportunistic networking, it must instead rely on longer contacts such as when people are stationary at work, sitting on a bus or in a lecture hall.

Most studies on experimental Bluetooth traces have focused on analyzing the inter-meeting time of nodes but [22] and [26] also studied the contact duration and inter-contact time as we do in this work. A comparison with these works shows some significant difference in results. The contact durations seen in our Bluetooth simulations are considerably shorter than those in the conference environments studied in the experiments and our results do not confirm with the power-law decay experienced there (see figure 4 in [26] and figure 3 in [22]). The main reason is that the scenario in our simulations is quite different from the relatively low-mobility conference environments where nodes may spend long times in proximity. In the subway station, nodes are highly mobile and only briefly pause, such as when waiting for a train and standing in the coffee shop. Also, as previously mentioned, our simulations focus on operational-level mobility and do therefore not capture effects of diurnal regularities and social patterns in node movements.

V. Conclusions

In this work we have studied the effect of mobility on node connectivity in two pedestrian scenarios; a city section and a subway station. We have statistically analyzed the empirical distributions for the connectivity metrics and how sensitive they are to changes in mobility input parameters, such as cultural profile of nodes, walking speed distribution and arrival processes. For each scenario we find that the empirical distributions of the connectivity metrics are relatively insensitive to changes in input parameters, showing only a modest change in mean values and retaining their basic shape. This is positive since if the statistics were highly sensitive to modest changes in input parameters, capturing them by a model would be hard and it would be difficult to identify valid parameters for the model. Comparing connectivity metrics across scenarios however shows more difference as can be seen when comparing figures 4(c) and 5(a). This suggests that when modelling mobility, accurately capturing the scenario and its structure is more important than a detailed estimation of input mobility parameters.

Recent experiments using Bluetooth devices to log contact statistics of humans have suggested that both contact durations and inter-contact times (often referred to as inter-any-contact times in those works) are approximated by power-law distributions [26]. This contradicts our results; in the scenarios we consider the inter-contact times are well approximated by gamma distributions and the contact durations are approximated by log-normal distributions (although the log-normal approximation for the contact duration does not fully capture the shape of the empirical distribution). There are at least three reasons for this discrepancy. First, we study different scenarios with higher mobility than in the conference-like environments often considered in experiments. Second, our current simulation models do not consider strategic-level mobility and therefore effects due to diurnal regularities and social effects are not captured. Third, the neighbour discovery of Bluetooth returns a very coarse sample of the underlying contact process, as we show in Fig. 5(c). By simulating the Bluetooth radio model we have shown that the number of discovered contacts is reduced by several orders of magnitude and that Bluetooth essentially filters out most of the short contact opportunities and much of the effects of operational level mobility that we consider in this work. It is possible that this crude sampling results in the power-law behaviour observed in measurements. A low power radio with a faster neighbour discovery and shorter contact setup times than Bluetooth is desired to unleash the potential of opportunistic networking.

Acknowledgments

We thank Vladimir Vukadinović for his earlier contribution to the work presented here.

References

[1] F. Bai, N. Sadagopan, and A. Helmy, "The important framework for analyzing the impact of mobility on performance of routing for ad hoc networks," *Ad Hoc Networks Journal - Elsevier*, vol. 1, pp. 383–403, 2003.

[2] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on the design of opportunistic forwarding algorithms," in *Proc. IEEE Infocom*, Barcelona, Spain, 2006.

[3] M. Grossglauser and D. N. C. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Trans. Netw.*, vol. 10, no. 4, pp. 477–486, 2002.

[4] V. Vukadinovic, O. Helgason, and G. Karlsson, "A mobility model for pedestrian content distribution," in *Proc. Simutools, SCENES Workshop*, Rome, Italy, March 2009.

[5] A. Jardosh, E. M. Belding-Royer, K. C. Almeroth, and S. Suri, "Towards realistic mobility models for mobile ad hoc networks," in *Proc. Mobicom*, San Diego, CA, USA, 2003, pp. 217–229.

[6] M. Musolesi and C. Mascolo, "A community based mobility model for ad hoc network research," in *Proc. REALMAN*, Florence, Italy, 2006.

[7] F. Ekman, A. Keränen, J. Karvo, and J. Ott, "Working day movement model," in *Proc. MobilityModels*, Hong Kong, China, 2008, pp. 33–40.

[8] K. Lee, S. Hong, S. J. Kim, I. Rhee, and S. Chong, "Slaw: A mobility model for human walks," in *Proc. Infocom*, Rio de Janeiro, Brazil, 2009.

[9] S. P. Hoogendoorn and P. H. L. Bovy, "Pedestrian route-choice and activity scheduling theory and models," *Transportation Research Part B: Methodological*, vol. 38, no. 2, pp. 169 – 190, 2004.

[10] J. L. Berrou, J. Beecham, P. Quaglia, M. A. Kagarlis, and A. Gerodimos, "Calibration and validation of the legion simulation model using empirical data," in *Pedestrian and Evacuation Dynamics*. Germany: Springer, 2007, ch. 3, pp. 167–181.

[11] A. Lindgren, A. Doria, and O. Schelén, "Probabilistic routing in intermittently connected networks," in *Proc. SAPIR 2004*, August 2004.

[12] A. Balasubramanian, B. Levine, and A. Venkataramani, "Dtn routing as a resource allocation problem," in *Proc. SIGCOMM*, Kyoto, Japan, 2007.

[13] G. Karlsson, V. Lenders, and M. May, "Delay-Tolerant Broadcasting," *IEEE Transactions on Broadcasting*, vol. 51, no. 1, March 2007.

[14] M. Papadopouli and H. Schulzrinne, "Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices," in *Proc. MobiHoc*, Long Beach, CA, USA, 2001, pp. 117–127.

[15] S. Jung, U. Lee, A. Chang, D.-K. Cho, and M. Gerla, "Bluetorrent: Cooperative content sharing for bluetooth users," in *Proc. PerCom*, White Plains,USA, 2007.

[16] T. Camp, J. Boleng, and V. Davies, "A survey of mobility models for ad hoc network research," *Wireless Communications & Mobile Computing (WCMC)*, vol. 2, no. 5, pp. 483–502, 2002.

[17] D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile Computing*. Kluwer Academic Publishers, 1996, pp. 153–181.

[18] E. Royer, P. Melliar-Smith, and L. Moser, "An analysis of the optimum node density for ad hoc mobile networks," in *Proc. IEEE ICC*, Helsinki, Finland, 2001.

[19] K. Maeda, K. Sato, K. Konishi, A. Yamasaki, A. Uchiyama, H. Yamaguchi, K. Yasumoto, and T. Higashino, "Getting urban pedestrian flow from simple observation: realistic mobility generation in wireless network simulation," in *Proc. MSWiM*, Montréal, Canada, 2005.

[20] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi, "Understanding individual human mobility patterns," *Nature*, vol. 453, pp. 779–782, 2008.

[21] N. Eagle and A. Pentland., "Reality mining: Sensing complex social systems," *Personal and Ubiquitous Computing*, vol. 10, no. 4, pp. 255–268, 2005.

[22] A.-K. Pietiläinen and C. Diot, "Experimenting with opportunistic networking," in *Proc. MobiArch*, Krakow, Poland, 2009.

[23] M. Kagarlis, "Method and apparatus of simulating movement of an autonomous entity through an environment," United States Patent No. US 7,188,056 B2, March 2007.

[24] O. R. Helgason and K. V. Jónsson, "Opportunistic networking in omnet++," in *Proc. Simutools, OMNeT++ workshop*, Marseille, France, 2008.

[25] "Pace Of Life," http://www.paceoflife.co.uk/.

[26] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Pocket switched networks and human mobility in conference environments," in *Proc WDTN*, Philadelphia, Pennsylvania, USA, 2005.