



Working Day Movement Model

Frans Ekman, Ari Keränen, Jouni Karvo and Jörg Ott
Helsinki University of Technology TKK, Dept. of Communications and Networking
{frans.ekman|ari.keranen|jouni.karvo|joerg.ott}@tkk.fi

ABSTRACT

Abstract movement models, such as Random Waypoint, do not capture reliably the properties of movement in the real life scenarios. We present and analyse a movement model for delay-tolerant network simulations that is able to produce inter-contact time and contact time distributions that follow closely the ones found in the traces from the real-world measurement experiments. We validate the movement model using the ONE simulator.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Store and forward networks

General Terms

Design, Experimentation, Measurement, Verification

Keywords

Movement Model, Simulation, DTN, Delay-Tolerant Networking, Mobility Models, Routing

1. INTRODUCTION

Movement of the network nodes is essential for the performance of delay-tolerant networks (DTN). A movement model that captures the behaviour of the nodes in the real usage scenarios is thus needed for a reliable assessment of a new protocol.

There are two types of movement models that have been proposed for these analyses — generic high level models that aim to produce movement accurate enough with statistical measures, and models that describe incidental scenarios, hoping for a more accurate depiction of single devices.

While efficient to use in simulations, the high level models, such as Random Waypoint (RWP) [9], often imply that the scenarios for which the protocols are simulated have huge numbers of nodes, so that the relevant protocol features are given statistically realistic distributions of events.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobilityModels '08, May 26, 2008, Hong Kong SAR, China.
Copyright 2008 ACM 978-1-60558-111-8/08/05 ...\$5.00.

For scenarios with few nodes, the differences between different usage scenarios become more significant. Thus, movement models that depict more precisely some specific types of movement are needed.

We present a new movement model to be used in DTN simulations, called *Working Day Movement Model*. The model presents the everyday life of average people that go to work in the morning, spend their day at work, and commute back to their homes at evenings. The model intuitively depicts the movement pattern of people, but we also verify the model by simulating it and compare the statistical features of the model to real-world traces.

This paper is organised as follows. Section 2 reviews related work on mobility models based upon which our model is presented in Section 3. The simulation tool we use for evaluation is introduced in Section 4. Section 5 shows the evaluation results, and finally, Section 6 concludes the paper.

2. RELATED WORK

Mobility models have been under active research recently, see [2]. *Inter-contact times* and *contact durations* are typical metrics for characterising mobility in sparsely populated DTNs. An inter-contact time, or sometimes referred to as an inter-meeting time, is the time interval between contacts for a node pair. It is defined as the time interval a node pair is not in contact with each other. The contact time or contact duration is the time a contact between two mobile nodes lasts. Inter-contact times correspond to how often nodes will have an opportunity to send packets to each other, while the contact durations limit the amount of data that can be sent. Usually, inter-contact time distributions and contact time distributions are used in comparisons.

Musolesi et al. [14] show that simple mobility models have very different properties in terms of inter-contact time and contact durations compared to real user traces. The distribution of inter-contact times does have practical implications. Chaintreau et al. [4] show that if the distribution of inter-contact times is power-law distributed with exponent less than one, any possible routing algorithm in a delay tolerant network will produce an infinite average delay for packet delivery. They also analysed four different traces of real people moving, and concluded that the inter-contact times are power-law distributed with the power-law exponent less than one.

Karagiannis et al. [10] show that the inter-contact times are only power-law distributed up to 12 hours, and have an exponential cut-off after that. A possible cause for the phenomenon is the daily routines people have.

Han Cai et al. [1] show that simple models on a boundless area can produce a power-law distribution of inter-contact times. Additionally, they show that the exponential cut-off is in many cases a side-effect of the bounded area. The motivation behind this is that nodes that would move over the edges on an infinite area are forced to stay within the area, thereby meeting other nodes sooner than they otherwise would and the number of long inter-contact times will be smaller.

Kim et al. [12] extracted various parameters such as speed and pause time distributions from real user traces. These parameters were then used in a synthetic mobility model. Their model is validated based on the number of nodes within different regions on the map at different hours. This might not be a sufficient criterion alone to determine the suitability of a model, since the performance of protocols and applications in DTN is highly related to the nature of the contacts.

Rhee et al. [15] uses Levy Walks to generate movement traces. The model is very similar to random walk, except that the flight lengths and pause times are drawn from a power law distribution. They manage to produce similar inter-contact time distributions as many real world traces, but the model does not capture characteristics as heterogeneity among nodes, repetitiveness, group mobility or any relationships between nodes.

The community based mobility model [14] is based on the idea that nodes favour squares with higher social attractiveness. The social attractiveness is based on how many friends are in the same square. Changing friends depending on the time of day results in periodic patterns like for example people meeting their work colleagues in the day and their family in the evening. This model lacks group movement and the movement is relatively homogeneous. The paper does not show the inter-contact time distribution behaviour for more than up to roughly one third of a day.

The time-variant mobility model [8] is somewhat similar. In this model, nodes move to different squares at different times of day in a periodic manner, thereby creating some heterogeneity in both time and space. Nodes do not move in groups and the movement is homogeneous in the sense that every node follows the same instructions.

Little work has been made on indoor movement, especially combined with outdoor movement. Some models and thoughts of office scenarios can be found in [13] and [5], but in a context not really practical for delay tolerant network simulations.

Hsu and Helmy [7] show by studying real user traces that nodes are very often turned on/off and only visit a small portion of the WLAN access points in campus areas. Moreover, they find that node mobility while using network is very low and one node only meets a small portion of all other nodes in the area. Furthermore, they reveal repetitive patterns with a period of one day and heterogeneity among nodes. According to them, the biggest issue with most synthetic models is that they are not capturing such characteristics as heterogeneous behaviour, switching devices on/off or relationships between users.

Various group mobility models exist and analysis of the impact of group mobility [6] has been made. However, to the best of our knowledge, group mobility has never been a component of a larger model covering many other aspects like community, daily routines, heterogeneity, etc. The same ap-

plies to most of the ideas presented above; they only model one aspect of mobility. Our approach is to combine these different elements to create a new movement model.

3. WORKING DAY MOVEMENT MODEL

We have developed a new mobility model by combining different movement model elements together. These models are called *submodels*. The model consists of three different major activities that the nodes can be doing. They are being at home, working and some evening activity with friends. These activities are the most common and capture most of a working day for the majority of people. More subtle variations and many other activities exist but, for now, we assume that they are reasonably well captured by the activities we have modeled or their overall impact is small and leave introducing further diversity (by means of more submodels) to future work.

On a more detailed level, the activities differ from each other. These submodels repeat every day, resulting in periodic repetitive movement. Their parametrisation and adding further submodels as needed allows fine-tuning the model to meet the needs of the target scenarios.

Communities and social relationships are formed when a set of nodes are doing the same activity in the same location. For example, nodes with the same home are family members, while nodes with the same office location are colleagues from work.

Nodes are doing the activities on a daily basis starting from home in the morning. Each node is assigned a wakeup time, which determines when the node should start from home. This value is drawn from a normal distribution with mean 0 and configurable standard deviation. The node uses the same wakeup time every morning during the whole simulation. The variance in the wakeup time models the differences in rhythms in real life.

At the wakeup time, nodes leave their homes, and use different transport methods to travel to work. Nodes travel between activities either by car or by bus, which are both different submodels. The working time is configurable. After the working hours, the nodes decide, by drawing, whether they go out for the evening activity, or return home. Again, different submodels are used for transitions between the locations. Different user groups have different locations where the activities take place.

3.1 Home Activity Submodel

The home activity submodel is used for the evenings and nights. Each node is initially assigned a map point as its home location. Having reached this location, the node walks a short distance away and stays still until the wakeup time.

We do not model any movement inside homes. Node activities at home can consist of the device lying on some table until the next day, people watching TV, cooking, sleeping etc., where the movements within the house are not relevant.

3.2 Office Activity Submodel

The office activity submodel is a 2-dimensional model for movement inside an office where the employee has a desk and sometimes needs to walk to other places for meetings or just to quickly talk to someone. Minder et al. [13] present a model for meetings where organisation structure is taken into account. We do not use such a model because we are

actually interested in the contacts of nodes, due to the application to delay tolerant networking.

Habetha et al. [5] used a more detailed office movement model, where employees are moving in rooms and corridors. The walls will have a significant effect on the path-loss. For a simpler modelling, we do not model the signal attenuation on walls.

The model adopted is as follows. The office is entered from a specific map point, called a *door*. The office is a square where the upper left hand corner is the door. Each node is assigned a coordinate inside the building where the node's desk is located.

The movement inside the office starts immediately when the node reaches the door; the node starts walking towards the desk with the walking speed defined in the settings. When it reaches its desk, it stops for an amount time, drawn from a Pareto distribution. When the node wakes up from the pause, it selects a new random coordinate inside the office, walks there and waits for an amount of time also drawn from the same Pareto distribution. The movement between the desk and randomly selected coordinates repeats until the work day is over. The purpose of nodes moving between their desk and random coordinates is that nodes having their desks close to each other will meet each other more frequently and nodes with their desks located next to each other will be in contact most of the time.

Earlier research suggests that the length of meetings at an office follow a log-normal distribution [13]. However, the study covers only team meetings, which does not necessarily correlate with pause times in movement. A truncated Pareto distribution is suggested in [15] for general movement inside buildings. We choose the Pareto distribution for our pause times inside the office. We also added parameters to turn off the pausing completely and to have an infinite pause time, in which case nodes stay at their desk for the whole workday.

Obviously, there are a variety of different jobs and buildings where people move according to different patterns. Our model is a first level abstraction, but we are working on further parameterizations and extensions of the submodel to introduce broader diversity (and determine how this diversity impacts the mobility metrics).

3.3 Evening Activity Submodel

The evening activity submodel models the activities that nodes can do in the evening, i.e. after work. This activity is done in groups. The evening activity model can be interpreted as shopping, walking around the streets or going to a restaurant or a bar. Each node is in the beginning of the simulation assigned a favourite meeting spot. Immediately when a node ends its working day, it is assigned to a group based on its favourite meeting spot. If all groups for a given favourite meeting spot are full, a new one is created with a randomly selected and uniformly distributed size with minimum and maximum values defined in settings. The node then uses the transport submodel to move to the meeting spot. The node waits at the meeting spot until all the nodes of the group are present. Then they start moving according to the map based movement model, which is actually a random walk on streets. They all walk in a group along roads a certain distance defined in settings, and then they pause for a longer time defined in settings, and finally split up and walk back to their homes.

3.4 Transport Submodel

Nodes move between home, office and evening activity using the transport submodel. During the initialisation, a configurable percentage of nodes in each group are set to use a car for transportation between activities. Nodes not moving by car will use the bus or walking submodel. Nodes moving by car only use the car submodel for all transportations. Supporting different types of transport models adds additional heterogeneity and has impact on the performance of routing protocols, since quicker nodes, like cars for instance, can transfer packets longer distances quickly.

- **Walking submodel**

Nodes that walk use streets to advance with a constant speed towards the destination. Dijkstra's algorithm is used for finding the shortest path to the destination.

- **Car submodel**

Nodes owning a car can travel at a higher speed between different locations. Otherwise it does not differ from walking. Within an activity submodel, car owners behave as the other nodes.

- **Bus submodel**

Nodes without a car can use buses for travelling faster. There are pre-defined bus routes on the city map. The buses run these routes according to a schedule. Buses can carry more than one node at a time.

Each node that does not own a car knows one bus route. It can use any bus driving that route. The nodes make the decision of taking the bus if the Euclidean distance from the node's location to the nearest bus stop summed with the Euclidean distance from the destination to the nearest bus stop is shorter than the Euclidean distance between the node's location and the destination. Otherwise, it walks the whole distance. If the node decides to take the bus, it uses the walking submodel to the closest bus stop and waits for the bus. When the bus arrives, the node enters it and travels until the bus comes to the bus stop nearest the destination. Then it switches back to the walking submodel to reach the destination.

3.5 The Map

All nodes move on a map. The map defines the space and routes in which the nodes can move; it contains all the information of the locations of the houses, offices and meeting spots, as well as the bus routes with bus stops. The design of the map is an important part of the mobility model. Since all the movement of the nodes is determined by activities with specific locations, the placement of these locations define how nodes are moving on a larger scale, i.e., in which areas of the map nodes will be doing different activities. The positions of these locations can be node group specific, which makes it possible to create small districts within the map. Therefore, the map can be used to limit node movement to small areas, which we refer to as increasing the *locality*. On one hand, houses, offices and meeting spots can be spread randomly on the map, thereby, having very little locality and nodes meeting easily. On the other hand, it is possible to restrict node movement to very small areas by creating lots of small districts, thereby increasing the locality. These may also be combined: different sized district

where some overlap others allows to have high locality but also some movement between districts, which corresponds to nodes coming to some district to work or meet friends, while others are leaving their district for similar reasons. Nodes moving between districts, not located next to each other, will have to pass through other districts, thereby appearing as drive through traffic in the intermediary districts. Figure 1 shows an example setup. The most suitable configuration of districts is environment specific.



Figure 1: A map of Helsinki city’s central areas divided into 4 artificial districts

4. EXPERIMENTAL SETUP

We implemented our movement model as an extension to the Opportunistic Network Environment (ONE) simulator [11]. The ONE is a highly customisable communication network simulator for delay tolerant networking that has several movement models implemented, from simple Random Waypoint to more realistic Map Based Movement models that can import map data and constrain node movement to the streets and roads of the imported data. ONE can also visualise the imported data and node movement using a GUI which helps on validating the model in an intuitive way. Based on the node movement and nodes’ radio device’s range, ONE generates contact information reports that can be used for detailed analysis. This contact information also feeds the simulation engine embedded in ONE which supports multiple DTN routing protocols [11].

4.1 Mobility Modelling

ONE can import mobility data from real-world traces or other mobility generators. Movement models and report modules are loaded dynamically based on the given configuration so that the simulator can be easily extended with new modules and the modules used in different scenarios can be changed as needed. In a simulation setting, any number of types of mobile nodes—referred to as a *node group*—may be defined. A node group shares a common set of simulation parameters like speed and pause time distributions, message buffer size, and radio range, among others. Different node groups can also use different movement model modules.

The basic version of ONE supports the Random Waypoint [9] mobility model, arbitrary mobility models by using externally generated movement data, and different map-based movement models. All map-based movement models obtain their configuration data using files formatted with a subset of the Well Known Text (WKT) format. WKT files can be edited and generated from real world map data using Geographic Information System (GIS) programs such as OpenJUMP¹. With map-based movement models, the nodes move using roads and walkways from the map data. In addition, different node groups can be set to use only certain parts of the map, thus allowing to distinguish between cars and pedestrians so that the former do not drive on pedestrian paths or inside buildings.

The simple random Map-Based Movement model (MBM) is a derivative of the Random Walk model, where nodes move to randomly determined directions on the map following the roads as defined by the map data. The Shortest Path Map-Based Movement model (SPMBM) is a derivative of the Random Waypoint model, where nodes use Dijkstra’s shortest path algorithm to calculate shortest paths from the current location to a randomly selected destination, by using the roads or paths.

Finally, some nodes may have pre-determined routes in the map that they follow. This Route-Based Movement model (RBM) uses the same map data but nodes always select the next destination on the route they are currently travelling. This mode of movement is useful for modelling e.g., bus and tram routes.

4.2 Implementation

The Working Day Movement model was added to the ONE as a combination of many mobility models. One movement model implements the main model controlling the movement of the nodes going to work, to their homes and meeting their friends. The main model passes the responsibility to lower level models handling different activities and transportation. Additional information about the destination is passed to the transportation models by the main model so that the nodes can find the way to the right place. The main movement model also decides whether to travel by bus or by walking between activities, and whether to do some evening activity or not.

Buses are an extension of the Route-Based Movement model, using bus routes defined in WKT files. Buses interact with passengers through a bus control system. The bus control system acts as a mediator between buses and passengers, informing passengers when the bus stops and buses when passengers enter them. Each bus control system has a unique ID, which is used to link bus node groups together with normal node groups in the settings file.

The evening activity makes use of a similar control system, defined in the settings for each node group, to facilitate the group movement. The locations of offices, meeting spots and homes are listed in WKT files defined in settings for each group, or randomly selected by ONE if no WKT files are provided.

4.3 Measurements in ONE

Creation of reports for various events with the ONE is implemented with the help of different event listeners. The default package of the ONE contains some commonly used

¹<http://openjump.org/>

reporting tools and some new ones were developed. In the case of inter-contact times and contact times, the ONE gives as output a list with event lengths and observation counts, from which it is easy to calculate a CCDF. We have implemented a contacts per hour report generator, which counts the number of contacts happening each hour. Finally, we have implemented two different reporting tools to measure encounters. The first one counts the number of other nodes a node has encountered (unique encounters), and provides a distribution for the fraction of user population nodes have encountered. The second one counts the number of total encounters and the number of unique encounters for each node separately, and provides a list with all the nodes and their measurements.

5. SIMULATION

We validated our model by comparing it to real user traces in terms of inter-contact times, contact durations and contacts per hour. The inter-contact times and contact durations are commonly used in the literature for characterizing connectivity in DTNs, while the contacts per hour metric has been used to measure activity at different hours of the day. Data from real world measurements is available for all three metrics, thus allowing us to validate our model.

We had over 1000 nodes moving on a map of the Helsinki centre area with the surrounding districts with the size of roughly $7000 \times 8500 \text{m}^2$. The area was divided into 4 main districts, see Figure 1. Additionally, 3 overlapping districts were created to simulate movements between the centre and other districts, and one district to cover the whole simulation area. See Table 1 for details about the assignments of nodes, offices and meeting spots. Every district, except the one covering the whole map, was assigned its own bus route and 2 buses. The district covering the whole map has 4 buses driving on one route.

Table 1: The assignment of nodes, offices and meeting spots to the different districts

District	Nodes	Offices	Meeting spots
A	150	30	4
B	50	10	1
C	100	20	2
D	100	20	2
E (A and B)	100	20	2
F (A and C)	150	30	4
G (A and D)	150	30	4
H (Whole map)	200	40	5

Half of all the nodes were set to travel by car. The walking speed for nodes was set to 0.8–1.4m/s and for buses 7–10m/s with a 10–30s waiting at each stop. The probability to do some evening activity after work was set to 0.5 with the group size 1–3. The working day length was 28800s and the pause times inside the office were drawn from a Pareto distribution with coefficient 0.5 and minimum value 10s. The office size was set to a $100\text{m} \times 100\text{m}$ square. The size of the office was chosen so that it would compensate for the lack of floors, walls, etc. The differences in schedules of nodes were drawn from a normal distribution with a standard deviation of 7200s. We also added 10 nodes moving according to the SPMBM model in the background to simulate taxis, delivery of goods, etc.

The transmit range of all nodes was set to 10m. The nodes were considered to be in contact when they were closer to each other than the transmit range. In other words, immediately when two nodes were in reach of each other, a connection was established. In the real world, there is usually a connection setup delay and the frequency of scanning for other devices is usually limited to keep energy consumption low (leading to a detection delay). It is worth noting that this phenomenon has probably affected real world contact traces used for comparison. A study about optimal probing of contacts can be found in [18]. Considering scanning intervals is subject to ongoing work.

We used a warmup period of half a day, which is sufficient due to the periodic nature of the mobility model.

For comparison, we simulated a RWP scenario on a same sized simulation area with 1000 nodes, moving with speed 0.5–5m/s and pause time 1–3600s, both uniformly distributed.

We estimated the inter-contact and contact time distributions by sampling them from simulation runs of length $T = 5 \cdot 10^5 \text{s}$. Due to the finite simulation time, the longer events are less likely to get observed. This is because a larger fraction of them has the beginning or end outside the simulation time. This leads to a systematic error so that it is not easy to say whether there is an exponential decay in an empirical distribution or just this systematic error. To avoid this uncertainty, we adjust the experimental distributions as follows.

We assume that the events are uniformly distributed over a longer period of time. Then, consider the probability of an event of length x , $p(x)$. Only events that begin during the time interval $[0, T - x]$ will get recorded. To compensate this, the estimated probability density function $\hat{p}(x)$ is

$$\hat{p}(x) = \frac{T}{T - x} p'(x),$$

where $p'(x)$ denotes the measured density. We use $\hat{p}(x)$ for creating the Complementary Cumulative Density Functions (CCDF, $P[X > x]$). The same method was used also on the experimental data used as a reference. This removes the effect caused by different length of the measurement period between different measurements.

We ran the simulation with the settings described above and measured inter-contact times, contact durations, contacts per hour and unique encounters compared to total encounters for each node plotted on a scatter diagram.

Figure 2 shows that the distribution of inter-contact times is similar to the one from the iMote traces [16]. We have a power-law distribution up to roughly half a day, after which an exponential decay follows.

Our model has not been configured to produce a power-law with the exact same exponent as the one in the iMote traces, since the coefficient can vary between environments. Power-law exponents were calculated for different traces in [4], and the exponents ranged from 0.4 to 0.9 when no correction procedures like we used had been applied to the results. In our model, it is possible to vary the power-law exponent by altering different parameters to better match a specific environment.

From Figure 3 we see that the contact durations follow a similar curve as the one from the iMote trace [3] experiments. Our contacts per hour graph (Figure 4) is similar to what Song et al. [17] obtained from the Dartmouth traces. Additionally we see from the scatter diagram (Figure 5),

how much heterogeneity the model has compared to RWP. A dot indicates how many total encounters and unique encounters a node has. If nodes move according to different patterns and speeds within different sized areas, nodes will meet different amounts of other nodes. Additionally, a node restricted to a small area will keep meeting the same nodes over and over again, while a scout node exploring the whole simulation area will mostly meet new nodes. Therefore, the scatter diagram can be used to measure heterogeneity.




Figure 2: Inter-contact times for RWP (---), Working Day Movement (—) and iMote trace (—)




Figure 3: Contact durations for RWP (---), Working Day Movement (—) and iMote trace (—)

We experimented with three scenarios with different numbers of background SPMBM nodes, to better understand how mixing different movement models affects the inter-contact time distribution. Figure 6 shows that more SPMBM nodes smoothen the inter-contact times distribution and the exponential cutoff gets less sharp. This is due to the exponential nature of the SPMBM movement.

To investigate the effect of the evening activity, we conducted two experiments with different values for the evening activity probability; one where nodes always do evening activity and one where nodes never do. The first scenario





Figure 4: Contacts per hour (normalized) for RWP (---), Working Day Movement (—) and Dartmouth trace (—)



(a) Default scenario

(b) RWP

Figure 5: Total- vs. unique encounters for each node plotted as a dot

created a peak in the contacts per hour graph at the time when most nodes were doing their evening activity, while the second scenario lowered the activity graph to almost the same state as during night. We concluded that the probability of evening activity parameter can be used to adjust the model to different environments where there is more or less night-life.

To determine whether the way nodes move between different activities makes a difference, we experimented with two scenarios: one where all the nodes move by car and another where all use the bus. A high percentage of nodes travelling by car seems to reduce significantly the fraction of the population a node has encountered. Cars move fast and always along the shortest path, therefore the probability of meeting new nodes gets smaller. The bus travels longer routes and nodes have much more opportunities to meet others at bus stops or in the bus. We also get slightly longer contact durations with buses, but it seems as if most of the contact durations originate from other things.

To understand the impact of the map, we created a Manhattan like map where one block is a square with the side 180m. We compared the results to the results from the Helsinki map, and surprisingly, there were not many differences. To make the comparison easier we decided to remove the districts and randomly select the locations of the homes,

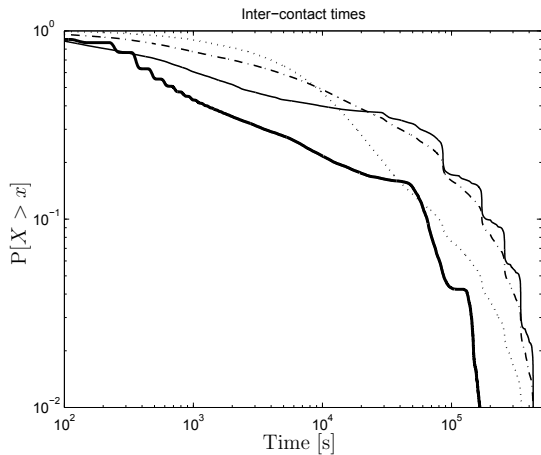


Figure 6: Inter-contact times for 0 (—), 100 (---) and 500 (· · ·) SPMBM nodes compared to the iMote trace (—)

offices and meeting spots. We found that there was no impact on the inter-contact times. The contact durations and fraction of nodes encountered were slightly affected. That is probably due to the fact that the bus routes were different, which affects the time travelled inside the bus and the amount of nodes encountered.

We also ran a simulation where no districts were used. The homes, offices and meeting spots were uniformly distributed on the whole map. Figure 7 shows that districts reduce the fraction of the user population encountered by nodes, and are therefore a good mechanism to further increase locality.

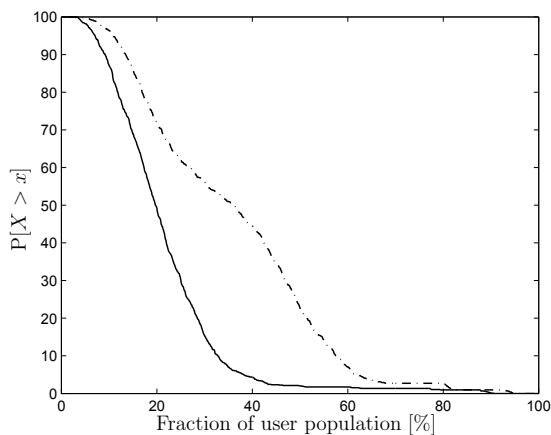


Figure 7: Fraction of user population encountered with districts (---) and without (—)

We also varied the movement inside offices. In the first scenario we let the nodes constantly move inside the office without pausing at all and in the second scenario nodes were not moving at all (infinite pause time). Figure 8 illustrates how this changed the power-law exponent of the inter-contact time distribution.




Figure 8: Inter-contact times when nodes constantly move (---), do not move (· · ·) inside offices compared to default settings (---) and iMote trace (—)

5.1 Addressing Modelling Complexity

A model capturing movement of nodes at level of detail like ours comes with a certain degree of complexity in terms of configuration. Specifying where thousands of nodes live and work is not feasible to do manually. Even more challenging is the process of obtaining reliable data for all the different parameters like people’s schedules, where they work and which buses they use. Therefore, we advise users of the model to create settings and other data files with the help of scripts, taking as input data that is crucial for the experiments and available to the user.

A model capturing movement of nodes at a detailed level must make a tradeoff between configurability and ease of use. Scripts taking as input only a few parameters from which the parameters of the real model can be derived from are easier to work with and thus more practical. Some statistical data, like residential density within different areas of a city, might be available to researchers; hence, a script can place homes and offices accordingly while randomizing the rest. Currently, if homes, offices, or meeting spots are not configured by the user, they are all uniformly distributed across the map. We are planning to provide a configuration template for Helsinki based upon real-world data as a starting point; similar ones can be developed for other cities and shared as a basis for reproducible results.

5.2 Limitations of the Model

Like every model, we need to provide an abstraction of reality to maintain the complexity at an acceptable level. All of the activities are modelled at a high level with simple abstractions on a two dimensional plane. Floors, walls and different obstacles affect mobility and especially the contacts between nodes, but are not covered in our model. A real working day exhibits more subtleties than we currently cover, including, e.g., lunch breaks and shopping activities. Furthermore, in our experiments, devices are always turned on. Even though modelling of attenuation and on/off times is not a responsibility of the mobility model, it is worth noting that measurements from real user traces, typically used for validation of mobility models, have been affected by

these details. While not included in our current model and implementation, further submodels can be added for more fine-grained capturing daily activities and parameterisation allows further fine-tuning. Similarly, on/off periods can be added by revising the corresponding mechanisms available in the ONE.

Finally, we currently do not model traffic in the sense of queues, traffic lights, or speed limits. Neither do the roads have any width so all nodes are moving along the same path and car/bus frames do not attenuate signals.

6. CONCLUSION

We have developed a new mobility model capturing several different mobility characteristics at a lower level of abstraction than many other models have. We have shown that the model is heterogeneous in both time and space and produces similar distributions of inter-contact times and contact durations as real user traces. Additionally, we have explored how different parameters affect mobility and revealed that the geometry of the map has a small impact on the outcome compared to many other configurable properties. Therefore, when modelling a specific environment with our model, researchers should concentrate more on other characteristics of the environment than the road net.

Next steps will include implementing additional submodels and developing the old ones further. Modelling of traffic at a more detailed level is also on the roadmap. Furthermore, we are developing scripts and tools to better handle the complexity while we are also conducting experiments to explore the parameter space in hope of eliminating redundant parameters. Finally, we have plans to improve contact scanning and to develop trace recording mechanisms using the same procedures as the real world experiments to improve validation of mobility models.

Acknowledgments

This research was funded by Nokia Research Centre in the SINDTN project and by the Academy of Finland in the DISTANCE project (grant no. 117429).

7. REFERENCES

- [1] CAI, H., AND EUN, D. Y. Crossing over the bounded domain: from exponential to power-law inter-meeting time in manet. In *MOBICOM (2007)*, E. Kranakis, J. C. Hou, and R. Ramanathan, Eds., ACM, pp. 159–170.
- [2] CAMP, T., BOLENG, J., AND DAVIES, V. A survey of mobility models for ad hoc network research. *Wireless Communications and Mobile Computing* 2, 5 (2002), 483–502.
- [3] CHAINTREAU, A., HUI, P., CROWCROFT, J., DIOT, C., GASS, R., AND SCOTT, J. Pocket switched networks: Real-world mobility and its consequences for opportunistic forwarding. Tech. Rep. UCAM-CL-TR-617, University of Cambridge, Computer Laboratory, Feb. 2005.
- [4] CHAINTREAU, A., HUI, P., CROWCROFT, J., DIOT, C., GASS, R., AND SCOTT, J. Impact of human mobility on the design of opportunistic forwarding algorithms. In *Proc. INFOCOM'06 (2006)*.
- [5] HABETHA, J., NADLER, M., AND CALVO, D. Dynamic clustering with quality of service guarantees and forwarder selection in wireless ad hoc networks. In *Proc. Asia Pacific Conference on Communications (Seoul, Korea, Nov. 2000)*.
- [6] HOG, X., GERLA, M., PEI, G., AND CHIANG, C.-C. A group mobility model for ad hoc wireless networks. In *Proc. ACM/IEEE MSWiM '99 (Seattle, WA, Aug. 1999)*.
- [7] JEN HSU, W., AND HELMY, A. IMPACT: Investigation of mobile-user patterns across university campuses using WLAN trace analysis. *CoRR abs/cs/0508009 (2005)*. informal publication.
- [8] JEN HSU, W., SPYROPOULOS, T., PSOUNIS, K., AND HELMY, A. Modeling time-variant user mobility in wireless mobile networks. In *Proc. INFOCOM'07 (2007)*, pp. 758–766.
- [9] JOHNSON, D. B., AND MALTZ, D. A. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing*, Imielinski and Korth, Eds., vol. 353. Kluwer Academic Publishers, 1996.
- [10] KARAGIANNIS, T., BOUDEEC, J.-Y. L., AND VOJNOVIC, M. Power law and exponential decay of inter contact times between mobile devices. In *Proc. MOBICOM'07 (2007)*, E. Kranakis, J. C. Hou, and R. Ramanathan, Eds., pp. 183–194.
- [11] KERÄNEN, A., AND OTT, J. Increasing reality for DTN protocol simulations. Technical Report, Helsinki University of Technology, Networking Laboratory, July 2007.
- [12] KIM, M., KOTZ, D., AND KIM, S. Extracting a mobility model from real user traces. In *Proc. INFOCOM'06 (2006)*.
- [13] MINDER, D., MARRÓN, P. J., LACHENMANN, A., AND ROTHERMEL, K. Experimental construction of a meeting model for smart office environments. In *Proc. First Workshop on Real-World Wireless Sensor Networks (REALWSN 2005)*, SICS Technical Report T2005:09 (June 2005).
- [14] MUSOLESI, M., AND MASCOLO, C. A community based mobility model for ad hoc network research. In *Proc. 2nd ACM/SIGMOBILE International Workshop on Multi-hop Ad Hoc Networks: from theory to reality (REALMAN'06) (May 2006)*.
- [15] RHEE, I., SHIN, M., HONG, S., LEE, K., AND CHONG, S. On the Levy-walk nature of human mobility: Do humans walk like monkeys? Technical Report, North Carolina State University, 2007.
- [16] SCOTT, J., GASS, R., CROWCROFT, J., HUI, P., DIOT, C., AND CHAINTREAU, A. CRAWDAD trace cambridge/haggle/imote/infocom (v. 2006-01-31). Downloaded from <http://crawdad.cs.dartmouth.edu/cambridge/haggle/imote/infocom>, Jan. 2006.
- [17] SONG, L., AND KOTZ, D. F. Evaluating opportunistic routing protocols with large realistic contact traces. In *Proc. CHANTS'07 (New York, NY, USA, 2007)*, pp. 35–42.
- [18] WANG, W., SRINIVASAN, V., AND MOTANI, M. Adaptive contact probing mechanisms for delay tolerant applications. In *MOBICOM (2007)*, E. Kranakis, J. C. Hou, and R. Ramanathan, Eds., ACM, pp. 230–241.