

MobIPLity: A trace-based mobility scenario generator for mobile applications

Nuno Cruz^{1,*}, Hugo Miranda²

¹Instituto Superior de Engenharia de Lisboa, Instituto Politécnico de Lisboa, LaSIGE, Portugal

²Universidade de Lisboa - Faculdade de Ciências, LaSIGE, Portugal

Abstract

The understanding of human mobility patterns is key for the development and evaluation of ubiquitous applications. To overcome the scarcity and difficulties in capturing mobility data, models have been devised. In general, each model replicates some of the observed metrics, while neglecting others. However, all tend to ignore diversity, in the roles and goals of the users but also in the devices that are used to access the WiFi network.

This paper presents the mobility traces from the access records of 49000 devices to the eduroam WiFi network of IPL for 7 years. Traces are made publicly available in the expectation that its large scale permits to support evaluations base on real mobility data, thus removing the uncertainty that emerges from the use of synthetic mobility models. Traces emphasise differences between device types, with impact on aspects like observed trace duration, speed, pause times, ICTs and availability, which can hardly be replicated on synthetic mobility models.

Received on 13 February 2015; accepted on 30 April 2015; published on 13 July 2015

Keywords: mobility patterns, mobile computing, traces

Copyright © 2015 N. Cruz and H. Miranda, licensed to ICST. This is an open access article distributed under the terms of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>), which permits unlimited use, distribution and reproduction in any medium so long as the original work is properly cited.

doi:10.4108/ue.2.5.e2

1. Introduction

Simulations play a fundamental role in the performance evaluation of mobile applications and protocols as they permit to circumvent the difficulties in deploying large scale and long term real experiments. Network simulators put to test an implementation of the application against abstractions of the environment, traffic and device movement. Therefore, the reliability of the experiments performed by network simulators is strongly influenced by their capability to reproduce observed settings in each of these abstractions.

The device mobility is dictated by the movement of their owners, which makes of human mobility a key factor influencing simulations. Research on human mobility has been pursuing two approaches: pure synthetic and trace-based synthetic mobility models. In pure synthetic mobility models nodes move according to some predefined statistical function. The

commitment of the rules defined for synthetic models to the replication of observed user movement patterns vary but this model has been criticised for its inability to reproduce human movement patterns, when evaluated by metrics like inhomogeneity [1]. As an example, consider random way-point [2], one of the most popular synthetic mobility models, where nodes unrealistically cycle between moving in straight lines to a random location and pause, both for random amounts of time.

Trace-based synthetic mobility models derive statistical distributions from observations of user movement, thus trying to mirror properties observed in real traces of human mobility (e.g. [3–5]). Traces are provided either by volunteers, which make their location available, or by a third party performing passive observation. Unfortunately, the number of trace-based samples made publicly available is scarce, present a small time span and/or number of users.

To circumvent the limitations of the pure and trace based synthetic mobility model classes, this paper proposes a new class, named *trace-set mobility*

*Corresponding author. Email: ncruz@deetc.isel.ipl.pt

models. The novelty of this class is that it creates mobility scenarios exclusively from data observed in real traces, disregarding any statistical approximation. As a contribution for the creation of this class, the paper presents a data-set composed of the observation of the 49000 wireless devices that connected to the eduroam network of the Polytechnic Institute of Lisbon, Portugal (IPL) between 2005 and 2012. This data-set is made publicly available using a web interface that exports data in bonnmotion [6] format, thus facilitating its use on a broad range of network simulators.

The contributions of this paper are two fold. First, it presents MobIPLity, a trace-set mobility model and scenario generator. MobIPLity combines the moments of association and disassociation of mobile devices to access points (APs) with the knowledge on APs geographic location to create realistic mobility scenarios. The scenarios can be filtered by many distinct parameters, including the number of devices, number of access points visited by each device, duration and device type. These parameters facilitate the generation of mobility scenarios that better purport the characteristics of the group of devices for which an application is developed.

Secondly, it characterises and compares the traces of MobIPLity with some trace-based models found in the literature. MobIPLity is characterised using temporal, spatial and social metrics frequently cited in the literature such as inhomogeneity, inter-contact times, jump-size and pause-times. Results show that trace-based models fail to emulate multiple device types and that they require a careful and difficult configuration to accurately model the behaviour that was observed in MobIPLity. The paper also shows that the different device types exhibit distinct mobility patterns, with impact on the metrics typically modelled by trace-based synthetic models.

The paper is organised as follows. The next section makes a brief overview of different mobility models, efforts to collect mobility data and metrics that have been used to characterise human mobility. Section 3 makes an overview of the data-set used by MobIPLity. Sections 4 and 5 respectively present the methodology used to convert the raw data-set on trace-based mobility scenarios and the different metrics associated with it. A discussion of the results and their comparison with other mobility models is presented on Sec. 6. Section 7 concludes the paper and outlines the plans for the continuation of this work.

2. Related work

2.1. Characterisation of Human Mobility

The human mobility has been characterised along spatial, temporal and social axis [7]. Spatial axis considers aspects like node density and distance,

portrayed by metrics like *jump size* (sometimes referred as *flight*) and *inhomogeneity*. Jump size characterises the average distance travelled by users and is affected by the characteristics of the area, for example, by the distance between buildings. Trace-based mobility models have been modelling jump sizes with either log-normal [8] or truncated power law distributions [4, 5]. The *inhomogeneity metric* aims at evaluating the dispersion of the devices on the physical space, in order to highlight hot-spots, something the proposers [9] consider to be a natural characteristic of human mobility. The variation of the Random Way-Point presented in [3] and the Disaster Area [10] are good examples of mobility models enforcing a heterogeneous node distribution. A lower inhomogeneity value is expected from random distributions, while a higher value shows that users are creating groups, formed by nodes placed in popular locations.

Time-varying properties of human mobility characterise patterns such as workday/weekend variations and *pause times*, i.e. the time spent on a specific place. Spatial properties are usually tied with temporal ones, associating the time to the distance travelled between two points, populating metrics such as *speed*.

The social axis characterises the meetings between participants. In combination with the temporal axis, they contribute to determine how long or how frequently two or more persons meet. Multiple models with a strong focus on the social relationships established between participants have been proposed. Metrics considered include attraction (found for example in [5]) but also repulsion. Both properties are explored in [11] by combining the modelling of relationships with individual walks and group trips. The *inter-contact time* (ICT), defined by the time interval between two consecutive contacts of two persons, is a frequently used metric to relate the temporal and social axis. Trace-based synthetic mobility models frequently model ICT using a truncated power law distribution [4, 12, 13].

2.2. Mobility Models

Research on human mobility has been pursuing two approaches: pure synthetic and trace-based synthetic mobility models. Pure synthetic mobility models use random distributions to simulate device movement. Classical examples are the Random Way-point Mobility Model [2] (RWP), the Disaster Area [10] and the Manhattan Grid. An advantage of pure synthetic mobility models is their simplicity in the generation of mobility scenarios, which facilitated cross comparison of mobile applications and protocols. Intuition suggests that random movement would be the most challenging for evaluating mobile applications' performance. However, it

has been shown that, in addition to their disparate modelling of human behaviour, synthetic mobility models typically bias node distribution and speed in a non-natural way [1, 14]. Limitations of the RWP have been addressed, for example in [3, 15–18].

Trace-based synthetic mobility models, on the other hand, attempt to mirror patterns observed in human movement by modelling nodes behaviour according to some probabilistic distribution functions. The mechanisms used for collecting data inspiring trace-based mobility models can be arranged in two categories. Intrusive approaches (for example [19, 20]) are those that obtain their data directly from the device carried by the user. These approaches benefit from the precision of the data, captured by dedicated software or hardware. Unfortunately, these studies are constrained by the considerable amount of resources involved, which limit their time scale and number of participants and may bias conclusions concerning the identification of patterns.

Non intrusive approaches, of which [21–23] are good examples, use logs collected by external devices (like access points or indoor-localisation devices) to produce traces with the user location at each instant. In spite of the privacy issues raised with the collection of the data, non intrusive approaches are those that present the capability to scale better in both number of users and time span. Unfortunately, surveys on mobility models [7, 24] indicate a scarcity of traces from mid-2008 onward, thus excluding the generalisation of mobile devices observed with the emergence of the last generation of smart phones and tablets. If available, more recent traces could evidence the emergence of new mobility and contact patterns among users resulting from the wave of mobile devices, such as the iPhone and Android OS based smart-phones, debuted respectively in 2007¹ and 2008².

The WiFi network of the Dartmouth College has been serving for collecting a considerable number of traces, for example during the 17 weeks of the 1999/2000 [25] and 2003/2004 winter semesters [22]. The method for collecting the traces is very similar to the one used in this paper and which is further described in Sec. 4. Authors used the logs to model real user tracks and defined a threshold walking speed, below which users were assumed to have stopped before moving to the destination. In comparison with the work presented in this paper, the study of 2003/2004 evaluates a larger number of access points, but a lower number of users and a shorter time frame. An interesting result of this study was the definition of a trace-based synthetic mobility model [8] inspired on the mobility patterns of

a limited (198) number of VoIP handsets. The model addressed social, spatial and temporal features and considered hot-spots, workday/weekend distinction, and mobile and stationary sets although it is affected by the particularities of the users hosting the devices.

Results on a two month study on the eduroam infrastructure of the universities of Minho and Vigo in 2010 can be found in [23]. The methodology followed is very close to the one used in our study. The access point association to physical spaces allowed to separate network traffic originating in residential from academic areas. Authors found that the APs with more users are not necessarily the ones with more network traffic. In addition, the paper confirms the expectations of a weekly usage pattern for this network, with the vast majority of users connecting only on weekdays. In terms of mobility, authors conclude that 90% of the users connect to more than one AP monthly, with about 35% visiting at least 5 APs. Unfortunately, the small analysis period of this study makes the notion of mobility disperse in time and of little relevance in the characterisation of real mobility.

Both SLAW [4] and HCMM [5] are trace-based synthetic mobility model which aggregate conclusions reported on other studies on human mobility with well-known characteristics like social and location attraction to popular places and the patterns of movement within confined areas. Metrics such as inter-contact time (ICTs), pause-times and jump-sizes are also inspired on previous research and modelled as truncated power-law distributions [12, 13]. SLAW and HCMM were validated by comparing its output with other models, either through statistical fitting of the generated results or by evaluating the performance of routing protocols for Delay-Tolerant Networks (DTN) using different mobility traces. In GeSoMo [11], the authors additionally extended social features by noticing the existence of a repulsion force in social relationships. GeSoMo models movement as a combination of individual walks and group trips, the latter representing users that walk in group to a predetermined popular destination. The authors also included the knowledge about statistical distributions from previous studies. GeSoMo and HCMM use as input a social network model, defined as a set of relations between users.

3. MobIPLite Data-Set

The MobIPLite data-set presented in this paper is composed by the log records produced, between January 1st, 2005 and December 31st, 2012, by all Access Points (APs) of the eduroam WiFi network of the Lisbon Polytechnic Institute (IPL). A total of 48699 devices and 30629 distinct users accessed the network

¹<http://en.wikipedia.org/wiki/IPhone>

²[http://en.wikipedia.org/wiki/Android_\(operating_system\)](http://en.wikipedia.org/wiki/Android_(operating_system))

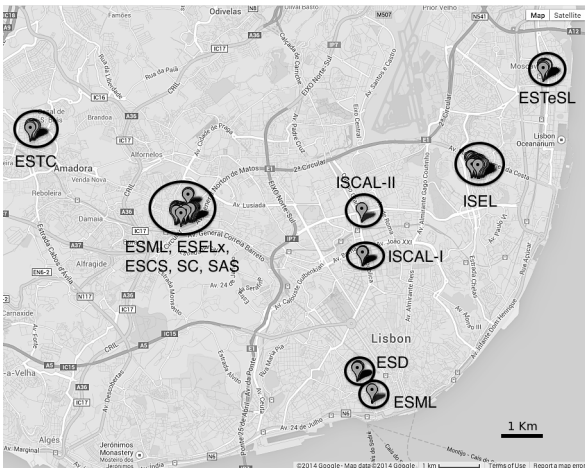


Figure 1. Location of IPL sites

during this time frame, producing about 31 million records.

IPL is the 7th largest teaching institution in Portugal with approximately 1300 teachers and 15000 students, distributed by 10 distinct sites on the Lisbon metropolitan area (see Fig. 1). The IPL's eduroam network is supported by approximately 200 Cisco Systems APs, covering a total of 26 buildings and inter-building areas. Records are originated from all the users accessing the network, thus also including visitors from other institutions.

The study partitioned the devices in two classes: *Small Mobile Devices* (SMD) are small, can be used on the move and are usually turned on uninterruptedly. Examples of SMDs are smart-phones, PDAs and tablets. The second class, *Laptops*, group the larger devices, usually executing a classical operating system (Linux, Windows or MacOS).

Devices were distributed by these classes using the information voluntarily provided in the DHCP [26] *vendor*, *parameter request list* and *hostname* message fields to learn their operating systems. It should be noted that information in these fields can be biased by: *i*) intentional changes to the information sent in these records by the owners of the devices; and *ii*) failure to comply with the DHCP protocol specifications from some OSs. No attempt to circumvent these limitations has been made. Furthermore, because DHCP logs were being discarded until the end of 2008, results for years between 2005 and 2008 include only the devices that connected at least once since 2009.

Overall, it was possible to identify the operating system running on 81.6% of the devices that connected to the network between 2005 and 2012. Of these, 7592 were associated to the *SMDs* class and 33054 are assumed to be *Laptops*. Analysis of the data-set show results for a third class, *All*, which aggregates *SMDs*, *Laptops* and the 8260 devices for which no




Figure 2. Operating systems observed (by ascending order)




Figure 3. Evolution of devices, users and access points




Figure 4. Laptops vs SMDs

classification was possible. The distribution of devices per operating system is depicted in Fig. 2. Devices that have presented themselves with several operating systems are represented once per operating system detected.

Figure 3 shows a continuous growth of the number of users and devices, although at distinct rates, specially since 2010. This is coincidental with an increase in the sales of smart-phones observed at the national level and suggests that the number of users accessing the network with more than one device has been increasing. Figure 4, which compares the proportion of SMDs and Laptops in each year, supports this claim.

The collection of mobility data is centred on the logs produced by the RADIUS [27] protocol. Log entries reproduce the RADIUS session concept thus considering the association of each device to a single AP. Records contain the device MAC address, AP id, user name, session start and stop times.

In ideal conditions, each RADIUS session should represent the association of a device to an access point.




Figure 5. Evolution of number of sessions with time

However, the number of sessions observed is slightly amplified due to: *i*) automatic handover between APs, triggered by variations in signal strength; *ii*) incompatibilities between client drivers and protocol versions running in the AP and; *iii*) operating system energy saving mechanisms that may turn off the radio interface when it is not in use. Interpretations of the results which rely on the number of sessions should therefore be made with some caution and take into account these factors. To mitigate some obvious anomalies, logs have been edited by:

- merging in a single record consecutive sessions between the same device and AP with an interval of less than 5 seconds. These sessions are attributed to network card or driver problems;
- removing concurrent sessions of the same device to distinct APs. This is an impossibility that can only be explained if the device did not disassociate correctly from one AP before associating to the next and the former artificially defined the session stop time upon a timeout. In this case, the session stop time of the earliest session was corrected to happen immediately before the start time of the latest;
- removing sessions with the stop time equal to the start time. Sessions with these characteristics are created when a user has some problem when connecting to the network, although the network considers the user authenticated (thus creating the RADIUS record).

The evolution of the total number of RADIUS sessions with time is presented in Fig. 5. The temporal evolution on the absolute number of sessions must consider the gradual capacity growth of the eduroam network (cf. Fig. 3), which in case of user mobility can increase the number of sessions established on the same path but taken in different years.

The increasing penetration of SMDs on the Eduroam network and its impact on the user mobility is further supported by Fig. 6. The figure depicts the yearly evolution of 2 interdependent metrics: the number of APs visited daily and the session duration, from both




Figure 6. Evolution of the average number of sessions and their duration with time

the users and SMDs perspective, and suggests that user mobility patterns between 2005 and 2012 can be partitioned in 3 distinct periods:

The period between 2005 and 2007 is characterised by the stability of both the average number of access points visited and the session length. Their small value indicates that devices tend to be fixed. These results are consistent with the ratio of 1 to 1 between users and devices (cf. Fig. 3) and with the small proportion (approximately 5%) of SMDs, observed in Fig. 4. SMDs visited on average a number of APs comparable with those of the remaining devices, a result that is attributed to the limitations of the first generation of SMDs (namely Personal Digital Assistants), where power efficiency of the wireless network interface was still a concern and motivated users to make a judicious use of their devices battery.

The period between 2007 and 2009 is characterised by an increase of nearly 100% of the average duration of sessions but with no change in the number of visited access points per user. This period is coincident with the increase of traffic, in contrast with the number of users, which continues to grow at an almost linear pace. The short distance between the number of devices and the number of users (cf. Fig. 3), and the stability of the number of visited APs, suggests that this period is uniquely characterised by an increase in the volume of IPL eduroam network use, attributed to the emergence of a considerable number of Internet based services, of which Peer-to-Peer networks are a remarkable example. To support our claim, it should be noted that this period precedes the flat rate business model, that became latter generalised for residential Internet access.

The year of 2010 marks the beginning of a new pattern where users connect to the network at a larger number of locations, although by shorter amounts of time. In this period, the average session duration falls progressively to values that, in 2012, become comparable to 2005. Simultaneously, the average number of visited APs increases by more than 50%. This result confirms our expectations that a significant change is taking place in the wireless

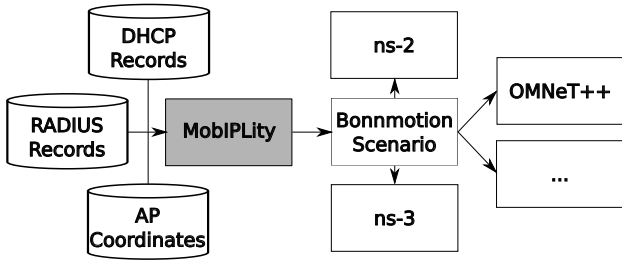


Figure 8. MobIPLity Work-flow

network usage pattern. This change is attributed to the wider deployment of SMDs, a claim supported by: *i*) an increase in the ratio between devices and users, depicted in Fig. 3; and *ii*) an increase of the ratio of SMDs over the total number of devices, evident in Fig. 4. It should also be noted that SMDs exhibit a pattern that differentiates from the average results, both when considering the number of visited APs (considerably more than the average) and the session duration (considerably shorter), thus permitting to confirm the increased mobility of SMDs.

Figure 7 depicts the number of distinct devices recorded by RADIUS logs per day. As expected, the plot exhibits an irregular pattern consistent with the different activity levels that can be found on workdays, weekends and summer and winter breaks in the campus.

4. MobIPLity

As depicted on Fig. 8, MobIPLity combines records produced by the DHCP and RADIUS services to create mobility scenarios that closely reflect observed user behaviour. In MobIPLity, DHCP records contribute with the identification of the device type and RADIUS identifies the participants and the moment of each association/dissociation event. The geographical coordinates of the access points contribute with the estimate of the location of each participant at each association/dissociation moment. Scenarios are produced in bonnmotion format, a popular mobility scenario generator capable of interacting with numerous network simulators.

This section begins with the presentation of the algorithm being used for extracting the traces from the information above, after what it briefly addresses the web interface used to make publicly available this information. The mechanisms that have been putted in place to guarantee user privacy are the focus of Sec. 4.3.

4.1. Trace Generation

The MobIPLity trace-set mobility model is created from a set $E \subseteq D \times A \times \{IN, OUT\} \times T$ where D is the set of wireless devices, A the set of access points of the

network annotated with their geographical coordinates and T are time stamps. The set is populated with 2 events $(d, a, IN, t_1), (d, a, OUT, t_2)$, for each RADIUS log record made available by the eduroam network of IPL. In these events t_1 and t_2 are time-stamps reflecting respectively d 's association/dissociation to AP a .

Let $E_d \subseteq E$ be the subset of E containing all the events recorded for device d . The set E_d is expected to respect two invariants: *i*) devices are always associated with an access point before being disassociated from it, i.e., $\forall (d, a, OUT, t') \in E_d, \exists (d, a, IN, t) \in E_d : t \leq t'$; and *ii*) in any point in time, a device is associated at most to one access point, i.e., $\forall (d, a, IN, t), (d, a', IN, t') \in E_d \wedge t \leq t', \exists (d, a, OUT, t'') \in E_d \wedge t \leq t'' \leq t'$. It should be noted that invariant *i*) is trivially assured by the access points software and invariant *ii*) by the corrections applied to the RADIUS logs that have been outlined in Sec. 3.

We define $E'_d = e_{d,0}, e_{d,1}, \dots, e_{d,n}, d \in D, e_{d,i} \in E_d, i > 0$ as the temporally ordered set of events for device d . It should be noted that according to invariants *i*) and *ii*), $e_{d,2j}, j \geq 0$ are events of type *IN* and, conversely, $e_{d,2j+1}, j \geq 0$ are all events of type *OUT*.

A trace $W_d = w_0, w_1, \dots, w_{2n-1}, n \geq 1$ for some device d , is defined as a sequence of way-points $w_i = F(e_{d,2j+i}), 0 \leq i \leq 2n-1, j \geq 0, e_{d,2j+i} \in E'_d$. The way-points are defined by a geographical coordinate and a time stamp, returned by a function F applied to consecutive events (not necessarily $e_{d,0}$) in E'_d . The output of function F depends of:

- the position of the way-point on the trace;
- the type (*IN, OUT*) of the event;
- the transmission radius estimated for the access point;
- the coordinates of the access point;

The general case is depicted in Fig. 9a. w_0 is set with the time stamp of $e_{d,2j}$ and the coordinates of the access point in this event. Subsequent transformations of pairs of events on pairs of way-points $w_{2i+1}, w_{2i+2}, i \geq 0$ will return coordinates overlapping a vector $\overrightarrow{AP_A AP_B}$, with AP_A, AP_B being the coordinates of the access points in the corresponding events $e_{d,2j+2i+1}, e_{d,2j+2i+2}$. The coordinates of the way-points are dictated by the transmission radius of the access points, as w_{2i+1} (resp. w_{2i+2}) will be placed at the intersection of the vector with the transmission radius of AP_A (resp. AP_B). Time stamps of w_1 and w_2 are copied from the corresponding events. Notice that, according to the definition of E'_d above, events $e_{d,2j+2i+1}, e_{d,2j+2i+2}$ are respectively an *OUT* and an *IN* record, thus signalling the moment at which d abandoned the area covered by AP_A and the moment at which d associated with AP_B . The algorithm is successively repeated for each pair of events and way-points.




Figure 7. WiFi Devices Connected Per Day

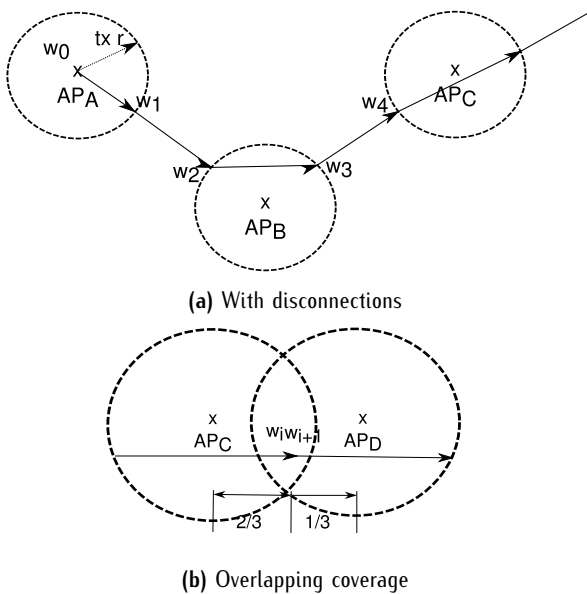


Figure 9. Trace extraction examples

The handling of the particular case occurring when the coverage area of two consecutive access points visited by the device intersects is depicted in Fig. 9b. The algorithm reflects the conservative approach of wireless interface drivers. The two way-points receive the time stamp of the *IN* record and are set at $2/3$ of the distance between the access points. This model reflects the expected conservative behaviour of the driver of performing one hand-off only when its benefits become evident.

Trace Termination Conditions The conditions for terminating a trace are motivated by the need to signal

the cases where the device abandoned the network, with the users moving to locations MobiPLity is unable to track and which would be otherwise represented by very slow movements that were not effectively observed. Traces are terminated at an *OUT* event by creating a way-point with the coordinates of the access point. MobiPLity identifies two conditions for interrupting a trace:

- when the speed for traversing the distance between two consecutive access points falls below some threshold;
- when two consecutive connections to the same AP exceed a time threshold;

Both conditions are triggered by thresholds that by default are set to $0.5ms^{-1}$ and $120s$ respectively, but can be configured according to the user preferences. Expectations are that these thresholds are sufficient for identifying the cases where the user abandoned the campus (for example to go home) in one location and re-entered it at a different (first condition) or at the same gate (second condition). In these cases, a new trace will be started from the next *IN* event for the same device.

4.2. Web Interface

To facilitate the dissemination of the traces, a web interface has been prepared and made publicly available at <http://edata.e.ip1.pt>. Traces are extracted on-demand from an *E* set stored at a local database and running the MobiPLity algorithm described in Sec. 4.1. The algorithm outputs traces in the format used by the bonnmotion mobility scenario generator and analysis tool [6]. The algorithm unifies the output, shifting the time stamps and creating an initial way-point for

Table 1. Trace extraction options

Parameter	Description
Start/End Date	Not before/after dates of the records
# Devices	Number of devices
Device' Type	Laptops/SMD/Any
Points	Minimum number of APs in a trace
Duration	Trace duration
Location	School for extraction
Axes	2D/3D representation
Warm Up	Duration of the warm up period
Cool Off	Duration of the cool off period
Enhanced trace	Include terminated traces
AP Range	Radius of the AP coverage
Speed	Min. speed for consecutive APs
Time	Max. time between connections

each trace with the location predicted for each device at scenario starting time. Finally, all access points are consistently positioned using a random factor.

Table 1 lists the parameters that can be configured for the generation of each trace. These parameters have been defined in order to facilitate the generation of a large number of instances, possibly, with distinct characteristics. Parameters can be arranged in four categories:

The *number of devices*, *device type*, *points*, *duration* and *location* parameters have a direct impact on the metrics that are more frequently referred in the characterisation of the mobility scenarios. Section 5 further addresses this aspect, by focusing on the role of the *location* and *device type* parameters in the generation of distinct mobility scenarios. The web interface permits to choose from 12 distinct locations, creating mobility scenarios whose areas range from small to medium campus sizes ($0.001km^2 - 0.063km^2$) and to a metropolitan area ($40km^2$) by considering the aggregation of the locations depicted in Fig. 1.

The *AP Range*, *Speed* and *Time* parameters influence the MobIPLity algorithm. As described in Sec. 4.1, *AP Range* has impact on the determination of the device location in traces while *Speed* and *Time* dictate the conditions for individual trace termination.

The *Warm Up*, *Cool Off*, *Axis* and *Enhanced trace* parameters address the more technical aspects related with the generation of the scenarios. The *Warm up* and *Cool off* parameters ensure that the devices remain active for the entire duration of the scenario. Therefore, MobIPLity exclusively selects for the scenario devices that have visited at least one access point in the *Warm Up* and *Cool Off* periods that respectively precede and succeed the time interval selected for the scenario. Alternatively, the *Enhanced trace* option permits the inclusion of devices that connect/disconnect to the eduroam network during the scenario. Unfortunately,

Enhanced trace conflicts with a number of network simulators which do not consider device disconnection in their mobility parameters.

Finally, the *start/end date* parameter permits to create multiple instances of scenarios with the same characteristics, knowing that MobIPLity will select the first moment after *start date* where all the remaining conditions can be satisfied for creating the scenario.

4.3. Enforcement of User Privacy

To protect the confidential nature of the data, original records are kept at a secure location and cannot be disclosed. In compliance with the bonnmotion file format, the algorithm exclusively outputs (time,coordinates) pairs for the distinct devices. Therefore, no identification that could be associated directly with a user or device is released. In addition, original data is obfuscated by: *i*) Positioning way-point coordinates in each new scenario generation while maintaining coherence; and *ii*) starting all scenarios at time 0, without disclosing the offset between the requested start date and the effective beginning of the scenario.

To somewhat limit any judicious analysis of the data that could be crossed with information made available from other sources, all requests of scenarios will be moderated. Boundaries on the duration of the scenarios may also be applied.

5. Mobility Analysis

This section presents and discusses the characteristics of the mobility patterns found on the 2012 subset of the MobIPLity trace-set. This year was chosen for being the most recent for which the data has been completely processed. Recency is a fundamental aspect as it better reflects the most up-to-date use of the technology, with an higher number and variety of device types.

Analysis proceeds in side-by-side comparison of the 2 alternative types of devices, permitting to confirm the existence of distinct mobility patterns for users carrying large (“Laptops”) and small (“SMD”) devices. To further increase diversity, the ISEL and IPL locations are considered. These are the contrasting extremes concerning node density. ISEL is the engineering school of IPL, located in a single site with an area of $0.063km^2$ and provides the largest number of devices from a single location. The IPL location considers records collected from access points at all schools (including ISEL) and presents a very small node density as campus are distributed over $40km^2$ of the Lisbon metropolitan area (Cf. Fig. 1).

Table 2 makes an overview of the dimension of the 2012 MobIPLity’s trace-set. It should be noted that columns for IPL consider all the institution, thus including ISEL. Still, ISEL accounts with approximately 40% of the devices and of the number of traces.

Table 2. Overview of the 2012 trace set

	All		Laptops		SMDs	
	IPL	ISEL	IPL	ISEL	IPL	ISEL
Devices	24141	10080	14947	7066	5403	2056
Traces	2075731	985816	1061686	602620	641394	250209

The first part of the evaluation provides an overview of the metrics observed in the traces produced from the complete data-set. Section 5.2 refines these results by performing an in-depth evaluation of 2 specific traces.

5.1. General Data-set Analysis

In the analysis of the complete data-set of 2012, discussion proceeds in two complementary perspectives. The “Per trace” perspective makes no association between the traces. I.e., traces are considered and averaged individually. In contrast, the “Per device” approach first aggregates the traces produced by each device and then proceeds by making an evaluation of the results on a device-by-device basis.

Trace duration. Figure 10 shows the complementary cumulative distribution function (CCDF) of the duration of the traces. As expected, the figure shows a small number of extremely long traces, exceeding 10 days, which are attributed to laptops connected on student dorms.

However, the figure also shows that less than 18% of the traces for Laptops exceed 2 hours and that for SMDs this value further decreases to about 10% of the traces in IPL and 7% in ISEL. These results are confirmed when observing the average session duration of each device on Figs. 10c and 10d. Such small proportion of “long traces” is surprising. One would expect that the usage pattern reflected the increasing use of mobile devices on the campus and, therefore, that trace durations were consistently higher.

The small duration of the traces, and the consistently lower average duration of SMDs when compared with Laptops, is attributed to the energy-saving mechanisms that can be found on mobile devices. These mechanisms automatically disable the wireless interface when not in use or when the screen is turned off. This is an aspect that has been consistently ignored in trace-based mobility models and is even hard to reproduce in network simulators. However, this feature has a non-negligible impact on the design and evaluation of many protocols and applications for ad hoc and delay-tolerant networks which assume “always on” connectivity of the devices. As a simple example, consider the impact of intermittent connectivity on the route discovery phase (that uses flooding) of many reactive routing protocols for MANETs, such as DSR [28] and AODV [29]. A more in depth investigation of the impact of the power saving

mechanisms is out of the scope of this paper and left as future work.

Speed. Speed plots (Fig. 11) exhibit some abnormal patterns of devices moving up to $1000ms^{-1}$. However, these are found on less than 0.01% of the traces and are attributed to the ping-pong effect that results from a combination of the fast roaming of the devices between overlapping APs and the trace generation algorithm used. This is a problem that has been observed in other models (e.g. [8]) and has a negligible impact as these fast speeds occur for very small amounts of time and distances. It should be noted that a portion of the 15% of the traces with an average speed above the average human movement speed on the IPL trace set can be attributed to users moving between sites, and that 18% of the IPL traces and 30% of the ISEL traces have a speed under 1m/s, which simply suggests users walking at low speed.

A comparison by device type shows that SMDs consistently present an average trace speed higher than Laptops. This result confirms the distinct utilisation pattern which can be easily observed in real life, with users operating their SMDs while walking. The “Per Device” perspective for IPL is still affected by the ping-pong effect. However, it is possible to observe a non-negligible number of devices (around 1%) roaming across distinct campus by exhibiting average speeds on the range of 30Km/h.

Distances Travelled. The distance travelled is evaluated using two metrics. Trace length is depicted on Fig. 12 and measures the length of each trace in meters. The geographical disposition of IPL sites and the roles of some of its members results in some traces obtaining surprising values of 100 Km. However, the “Per Device” averages are more predictable and only reach 11Km for IPL and 800 meters for ISEL.

As expected, the higher mobility of SMDs is confirmed by longer average traces in conjunction with shorter durations. However, looking at the complete trace-set of IPL we cannot differentiate between different device type distributions. This is expected, as users that carry a laptop are expected to equally carry a SMD and as such when they travel between IPL locations they carry both devices with them. Table 3 shows that about 27% of all devices on IPL and 16% on ISEL are static, contributing for the 70% of the traces without movement. However, the size of the MobIPLity trace set is sufficient to attenuate this large proportion

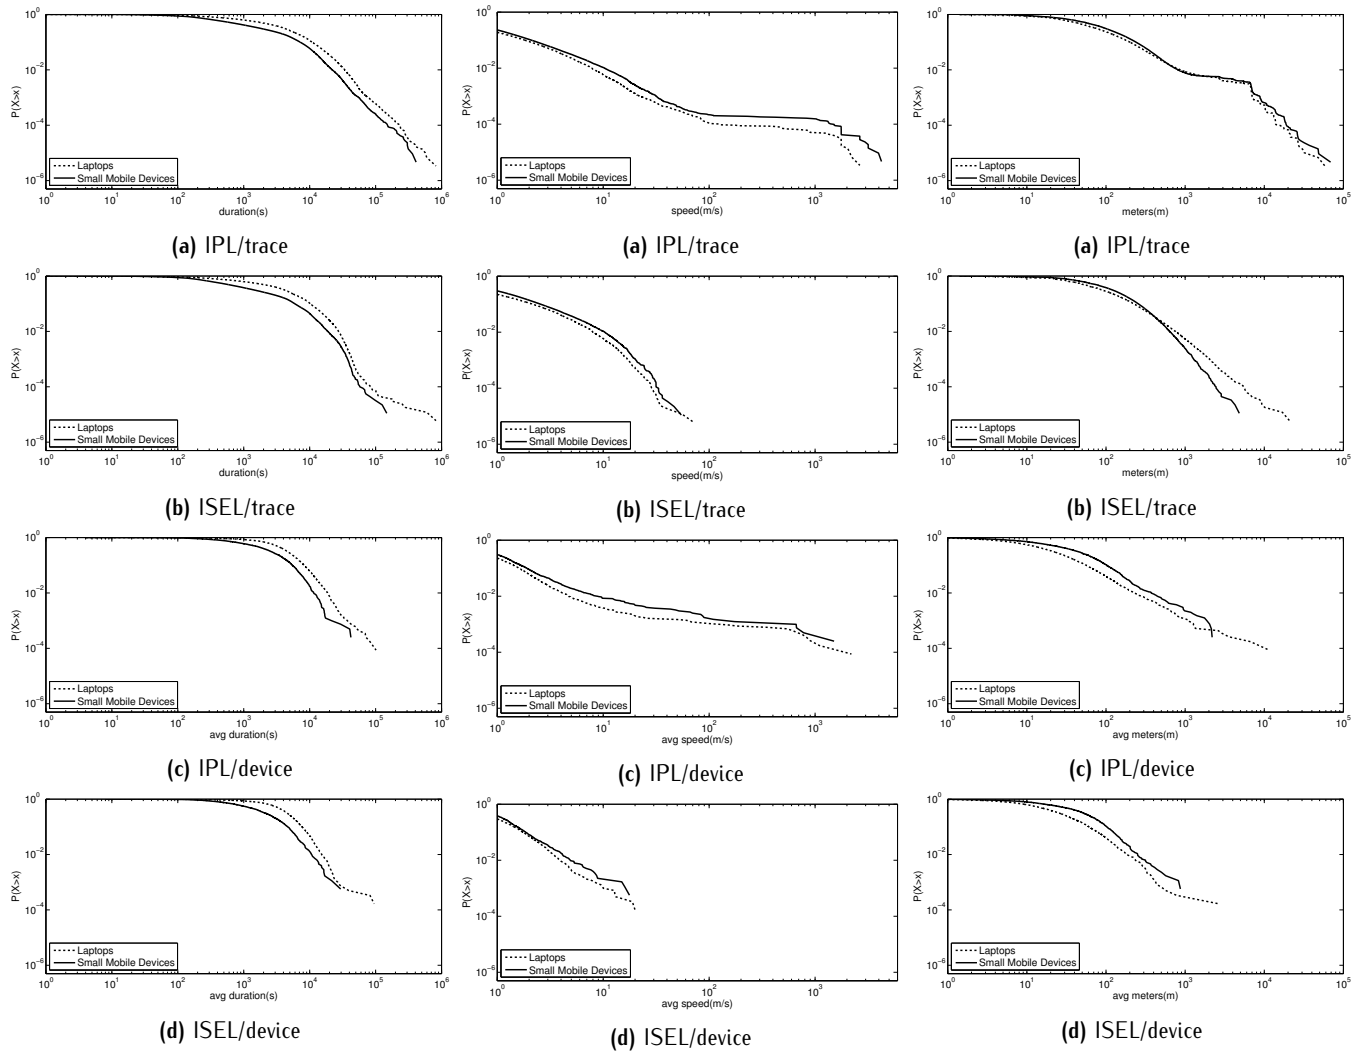


Figure 10. Trace duration (seconds)

Figure 11. Trace speed (m/s)

Figure 12. Trace length (meters)

as more than 600000 traces for IPL can be found exhibiting movement. Authors believe that such a large number should be considered sufficient to support a trace-set mobility model.

The distribution of jump sizes (i.e. the distance travelled between way-points) is depicted in Fig. 13. The irregular pattern observed at Fig. 13a, with knees at 100m, 7000m, 8000m and 10000m shows how the roaming of users among the multiple IPL campus impact the model. The smaller campus area of ISEL justifies the smaller travelled distances, which never exceed the 200m.

Pause times. Figure 14 shows consistently briefer pause times for SMDs on both IPL and ISEL trace sets. This supports common knowledge of SMDs showing a higher mobility, what contrasts with the expected large pause times for laptops, typically operated by steady users. The longer tail on the plot for Laptops can be caused by laptops that are kept at teachers offices, or

at students dorms. The logarithmic scale of the graph hides the large difference between the maximum pause time for laptops (almost 8 days) and a maximum of 2 days for SMDs. The difference between these values is consistent on IPL and ISEL.

Disconnection time. Figure 15 presents the CCDF for the time for which devices were disconnected, creating distinct traces. This metric was only obtained for devices that returned to the network after a disconnection. The figure clearly shows the impact of the academic environment where the data was collected. The plot knees evidence a considerable number of disconnections of 12 hours, 2 days, 12 days, 2 months and 6 months. These periods represent either weekend/weekday periods, vacations and semesters. We also found that laptops have a higher probability of being disconnected frequently for periods of 90 minutes, which is the duration of classes. In contrast, the figure indicates that SMDs have traditionally

Table 3. Number of samples with no distance travelled

	Location	All	Laptops	SMDs
Per Trace	IPL	2075731/1449669 (70%)	1061686/767041 (72%)	641394/426939 (66%)
	ISEL	985816/682205 (69%)	602620/435860 (72%)	250209/159873 (63%)
Per Device	IPL	24141/6395 (27%)	14947/3486 (23%)	5403/1417 (26%)
	ISEL	10079/1637 (16%)	7066/1092 (16%)	2056/337 (16%)

Samples with zero distance/Total number of samples (% of samples with zero distance)








(a) IPL/trace

(a) IPL/trace

(a) IPL/trace

(b) ISEL/trace

(b) ISEL/trace

(b) ISEL/trace
Figure 13. Jump size (meters)

Figure 14. Pause times (s)

Figure 15. Disconnection time (s)

smaller disconnection times, what is attributed to the power saving mechanisms available on these devices.

5.2. Scenario Analysis

This section uses some metrics discussed in the related work to evaluate the following 2 scenarios created using MobiPLity.

3 days This scenario puts together three traces of 2h each, extracted respectively from the 22nd of May, 18th of October and 6th of December 2012. The days and periods were individually selected to create a single 3600s scenario with the highest possible number of devices. This scenario was motivated by the objective of defining an environment as similar as possible to the one that can be more frequently found in the literature. Therefore, this scenario does not include devices whose trace terminated during the 2h period of the collection.

Disconnected The “disconnected” scenario further reflects the full potential of MobiPLity by eliminating any constraints aimed to reproduce the conditions usually found in the literature. It was extracted from the 18th of October 2012, a date chosen because it presents the largest number of devices in a single 3600s trace.

The disconnected scenario includes interrupted traces. Recall from Sec. 4.2 that interrupted traces include devices that turn off their radio during the period of the study.

Table 4 details the configuration parameters used in MobiPLity to produce both scenarios, according to the designation introduced in Table 1. The “disconnected” and the “3 day” scenarios share all the configuration parameters except for the option to include traces interrupted during the period. Expectations are that the differences observed between the two scenarios can give some hints on the impact of node disconnection.

Overall, this section considers 8 distinct data-sets, resulting from the combination of the two scenarios (resp. “3 days” and “disconnected”), two locations (ISEL and IPL) and two device types (Laptops and SMD). In the general case, all the data-sets consider 100 devices. The exception is for SMDs in the “3 days” scenarios where only respectively 15, 20 and 18 devices for ISEL and 28, 43 and 45 devices for IPL could be found.

ICT. Figure 16 depicts the CCDF for the Inter-contact times (ICT) of the scenarios. It should be noted that the ICT metric only considers pairs of devices that become in touch a second time, ignoring all the cases where devices are in contact at most once. This supports the irregularity of the plots for the “3 day” scenarios,

Table 4. Trace extraction options for the 3 days and disconnected scenarios

Parameter	3 days	Disconnected
Device' Type	Laptops/SMD	
Start/End Date	May 22, Oct 18, Dec 6, 2012	Oct 18, 2012
Points	2	
Duration	7200s	
Location	ISEL/IPL	
Axes	2D	
Warm Up	7200s	n.a.
Cool Off	7200s	n.a.
Enhanced trace	No	Yes
AP Range	50m	
Speed	0.5ms ⁻¹	
Time	120s	

attributed to the small number of devices that remained connected during the 2h period.

Results show longer inter-contact times for devices that are geographically bounded to ISEL, what should be expected, considering the highest density of the network (in comparison with IPL) which increases the probability of the nodes to become in proximity more frequently. In general, plots suggest frequent reconnections among pairs of devices, with only 1% of them interrupted by more than 1000s.

Jump size. The distinctive dimensions of the IPL and ISEL campus become evident in the CCDF of the jump sizes, depicted in Fig. 17. Results for IPL (Figs. 17a and 17b) exhibit a step pattern attributed to the distances between the different schools of the institution and to the need of some students and professors to commute between them.

A comparison between SMDs and Laptops shows that, in general, jump sizes of the former have a higher probability of being shorter than the latter. This can be attributed to the mobility of SMDs, which may connect to APs while being carried in the pocket of their users and which can be operated while the user is moving.

Figure 17a presents an interesting exception to the relation of the curves presented by laptops and SMDs given that laptops have a lower probabilities of moving throughout all the IPL. However, Fig. 17b contradicts the “3 day” results. Since the difference between both traces is restricted to the minimal speed of travel (which in Fig. 17a) must be above 0.5ms⁻¹, it is safe to assume that the abnormal behaviour is due to the speed at which the devices travelled such long distances. In general, jump size results tend to support the claim that SMDs have a higher mobility, which produces larger traces passing through multiple APs, while laptops are disconnected and reconnected at a new location.

Jump sizes are tightly associated with the physical dispersion of IPL, which creates groups of users that

Table 5. Inhomogeneity values

	All (σ)	Laptops (σ)	SMD (σ)
IPL-Oct 18th	0.79 (0.01)	0.9 (0.01)	0.77 (0)
ISEL-Oct 18th	0.59 (0.01)	0.62 (0.02)	0.68 (0.05)
IPL Disc.	0.76 (0)	0.87 (0.01)	0.75 (0)
ISEL Disc.	0.66 (0.03)	0.71 (0.01)	0.56 (0.03)
IPL RWP	0.4 (0.05)		
ISEL RWP	0.35 (0.06)		
IPL SLAW	0.63 (0.02)		
ISEL SLAW	0.52 (0.05)		

either remain close on one school or travel between several. This issue has been previously identified, for example, in HCMM [5] where scenario generation considers the possibility to set-up a number of groups and create bell shaped normal distributions.

Pause times. The CCDF of pause times is presented on Fig. 18. It should be noted that in MobIPLity, pause times are particularly small as the methodology followed for trace definition tend to maintain a node in movement even if at a very small speed (Cf. Sec. 4.1). Therefore, a change in the methodology was applied. The results presented on the figure consider a device to be stopped if the distance between way-points is less than 1m. Still, it is interesting to observe that SMDs have different pause time distributions with lower probabilities of having higher values, something that supports the mobility characteristics expected for SMDs.

Our results are in contrast with those presented in [8] where pause times were defined by detecting users walking at a low speed. Authors of [8] observed that pause times exhibited a log-normal distribution. In contrast, the application of the Akaike test to our datasets indicates that the results are closer to a power law distribution.

Inhomogeneity. Table 5 shows results for the inhomogeneity metric. Samples for this metric were obtained at 4 different times on the scenarios (at the beginning, 1/3, 2/3 and at the end of the scenario) of Oct 18th, the day with the most SMDs present on the network. Laptops show a higher inhomogeneity, indicating a larger concentration and irregular distribution for these devices. This is consistent with our expectations as it suggests that laptop users tend to be grouped, for example in classrooms or libraries. The lower value of inhomogeneity for SMDs confirms their pseudo-random deployment over the area.

6. Discussion

The evaluation above permitted to identify a number of metrics that are “scenario agnostic”. Trace duration, trace length, pause times, disconnection time and ICT

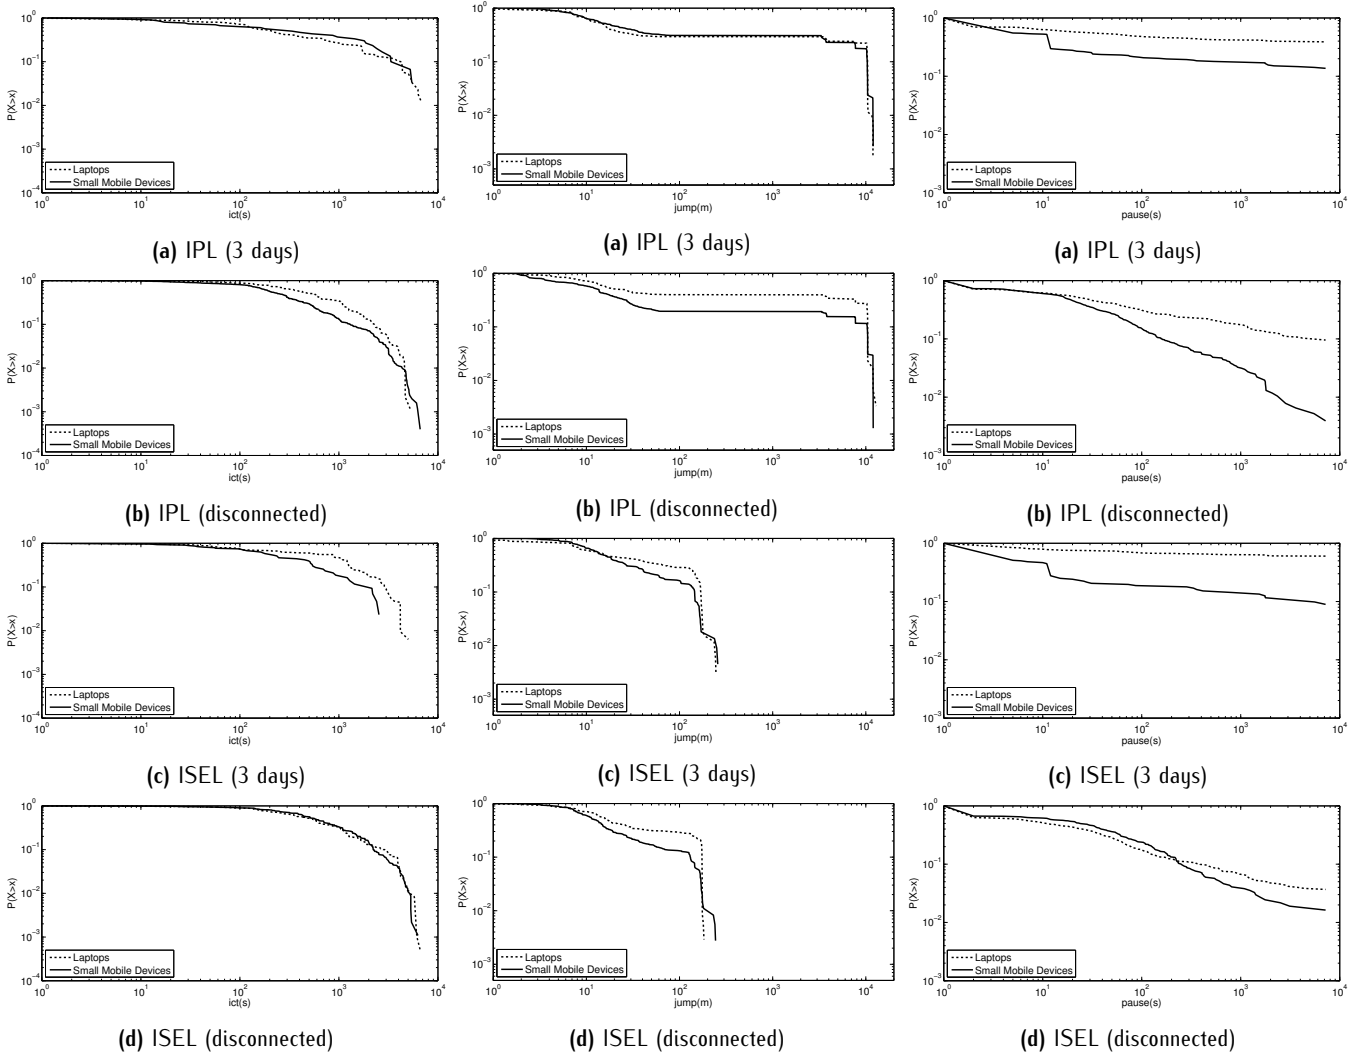


Figure 16. Inter-Contact Times (s)

Figure 17. Jump Size (meters)

Figure 18. Pause times (s)

are examples of metrics where differences observed between ISEL and IPL are minimal. On the opposite side, the evaluation also permitted to identify “scenario dependent metrics” of which trace speed, jump size and inhomogeneity are good examples. To facilitate the comparison with trace-based mobility models, we compare MobiPLity with SLAW [4] and RWP using ICT (a scenario agnostic metric) and Inhomogeneity (a scenario dependent metric).

To compare ICTs, the setup and data presented in Sec. 5.2 was used. Scenarios that geographically emulate ISEL and IPL and with a similar number of devices were arranged for SLAW and the Random Waypoint (RWP). To better replicate the real conditions, the location of Access Points was passed to SLAW as hot-spots, permitting the creation of a model as accurate as possible. Figure 19 depicts and compares the CCDF results of ICTs for MobiPLity, RWP and SLAW. The figure shows that for IPL, RWP distributes the nodes so

homogeneously that prevents generation of long ICTs, thus limiting the visibility of the data on the figure. Despite setting SLAW to emulate IPL on the number of access points, SLAW has limitations on representing the same ICTs as MobiPLity, which by itself presents similar ICTs for IPL and ISEL. Unfortunately, neither SLAW or RWP distinguish different device types although our results show that the device type plays a significant role on ICT.

The inhomogeneity metric, depicted in Table 5, was calculated for scenarios synthetically generated by SLAW and RWP that replicate the conditions found in IPL and ISEL (dimension, duration, devices and number of hot-spots/access points). Results of our mobility records are similar to the ones obtained by SLAW, and as expected, diverge from the randomness found in RWP where the metric value is low.

To better understand the differences between metrics, mobility models and device types, the Akaike test was




Figure 19. Comparison with ICT for SLAW and RWP

Table 6. Akaike test results

	SLAW	Dartmouth	MobiPLity		
	All	All	All	Laptops	SMDs
Duration	-	-	LN	W	LN
Dur./dev	-	-	G	G	P
Length	-	-	P	P	P
Len./dev	-	-	P	P	P
Speed	-	-	W	W	P
Speed/dev	-	E	P	P	P
Disconnection	-	-	GEV	LN	GEV
Pause	P	LN	P	P	P
Jump	P	-	P	P	P

(P: Pareto, LN: Log-Normal, G: Gamma, E: Exponential, W: Weibull, GEV: Extreme value)

used to compare the fitness of ISEL traces for the complete year of 2012 to well-known distributions. ISEL traces were chosen due to their containment in a single campus, approximating the configurations found in the literature. Results presented in Table 6 were obtained using Matlab automatic fitting using all supported distributions and sorted by Akaike criteria. It is interesting to notice the difference between fitting results for different device types, where for the vast majority of SMD metrics Pareto is chosen as the best fit. This is in contrast with Laptops, with metrics being approximated by a bigger variety of distributions.

Results confirm our suspicions on the existence of different mobility models for these devices. The table shows the distributions used by SLAW and Dartmouth [8]. The difference in chosen statistical models for pause times are evident. Pause time calculation methodology has impact on the fitting process. In [8], pause times were defined by detecting users walking at a low speed. The paper claims that pause times have a log-normal distribution, something that is not supported by our observations which show a power law distribution.

7. Conclusions and Future Work

The availability of real world mobility traces contributes to improve the predictability of new protocols,

applications and algorithms. This paper presented MobiPLity, a data set aggregating the access records to the eduroam network of the Polytechnic Institute of Lisbon between 2005 and 2012 and a methodology and website to extract mobility scenarios from this data-set. Expectations are that the effort to provide a public web interface and a set of indicative metrics will allow developers to rapidly create a scenario in a common, widely adopted format, that could be used in application or framework evaluations using simulations.

Having a source of real mobility data facilitates the modelling. In spite of the limitations dictated by the specific academic environment where the data was collected, MobiPLity permits to create distinct scenarios by considering the several schools of IPL and the distinct movement patterns that can be observed when considering Laptops, Small Mobile Devices and a combination of both. This paper claims that MobiPLity can be considered as a good starting point for the evaluation of a multitude of applications.

The recency of the data set permitted to observe the most recent pattern changes on mobility, that result from the increasing popularity of small dimension devices, effectively increasing user mobility. Interestingly, evaluation showed that power saving mechanisms that are standard on these devices reduce the possibility of spontaneous communication between devices and with the environment. This is an aspect that cannot be neglected by both mobility models and network simulators neither left for applications developers to address.

Analysis of the data continues. As future work, authors plan to analyse the data-sets of 2013 and 2014, in particular, to investigate if some changes in mobility patterns continue to evolve. In addition, authors plan to use this data-set to investigate users contact patterns, in particular, on the capability to predict future contacts.

Acknowledgement. Work described in this paper was partially supported by FCT through funding of LaSIGE Strategic Project, ref. PEst-OE/EEI/UI0408/2014.

References

- [1] BETTSTETTER, C., RESTA, G. and SANTI, P. (2003) The node distribution of the random waypoint mobility model for wireless ad hoc networks. *Mobile Computing, IEEE Transactions on* 2(3): 257–269. doi:10.1109/TMC.2003.1233531.
- [2] JOHNSON, D.B. and MALTZ, D.A. (1996) Dynamic source routing in ad hoc wireless networks. In IMIELINSKI, T. and KORTH, H. [eds.] *Mobile Computing* (Kluwer Academic Publishers), 153–181.
- [3] GYARMATI, M., SCHILCHER, U., BRANDNER, G., BETTSTETTER, C., CHUNG, Y.W. and KIM, Y.H. (2008) Impact of random mobility on the inhomogeneity of spatial distributions. In *IEEE Global Telecommunications Conf. (GLOBECOM)*: 1–5. doi:10.1109/GLOCOM.2008.ECP.85.

- [4] LEE, K., HONG, S., KIM, S.J., RHEE, I. and CHONG, S. (2009) Slaw: A new mobility model for human walks. In *INFOCOM 2009, IEEE*: 855–863. doi:10.1109/INFCOM.2009.5061995.
- [5] BOLDRINI, C. and PASSARELLA, A. (2010) Hcmm: Modelling spatial and temporal properties of human mobility driven by users' social relationships. *Computer Communications* 33(9): 1056–1074. doi:10.1016/j.comcom.2010.01.013, URL <http://www.sciencedirect.com/science/article/pii/S0140366410000514>.
- [6] ASCHENBRUCK, N., ERNST, R., GERHARDS-PADILLA, E. and SCHWAMBORN, M. (2010) Bonnmotion: A mobility scenario generation and analysis tool. In *Procs. of the 3rd Int'l ICST Conf. on Simulation Tools and Techniques, SIMUTools '10 (ICST, Brussels, Belgium, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering))*: 51:1–51:10. doi:10.4108/ICST.SIMUTOOLS2010.8684.
- [7] KARAMSHUK, D., BOLDRINI, C., CONTI, M. and PASSARELLA, A. (2011) Human mobility models for opportunistic networks. *Communications Magazine, IEEE* 49(12): 157–165. doi:10.1109/MCOM.2011.6094021.
- [8] KIM, M., KOTZ, D. and KIM, S. (2006) Extracting a mobility model from real user traces. In *INFOCOM 2006. 25th IEEE Int'l Conf. on Computer Communications. Procs.*: 1–13. doi:10.1109/INFCOM.2006.173.
- [9] SCHILCHER, U., GYARMATI, M., BETTSTETTER, C., CHUNG, Y.W. and KIM, Y.H. (2008) Measuring inhomogeneity in spatial distributions. In *Vehicular Technology Conf., 2008. VTC 2008. IEEE*: 2690–2694. doi:10.1109/VETECS.2008.589.
- [10] ASCHENBRUCK, N., GERHARDS-PADILLA, E., GERHARZ, M., FRANK, M. and MARTINI, P. (2007) Modelling mobility in disaster area scenarios. In *Procs. of the 10th ACM Symp. on Modeling, Analysis, and Simulation of Wireless and Mobile Systems, MSWiM '07 (New York, NY, USA: ACM)*: 4–12. doi:10.1145/1298126.1298131.
- [11] FISCHER, D., HERRMANN, K. and ROTHERMEL, K. (2010) Gesomo: A general social mobility model for delay tolerant networks. In *Mobile Adhoc and Sensor Systems (MASS), 2010 IEEE 7th Int'l Conf. on*: 99–108. doi:10.1109/MASS.2010.5663959.
- [12] KARAGIANNIS, T., LE BOUDEC, J.Y. and VOJNOVIC, M. (2010) Power law and exponential decay of intercontact times between mobile devices. *Mobile Computing, IEEE Transactions on* 9(10): 1377–1390. doi:10.1109/TMC.2010.99.
- [13] CHAINTREAU, A., HUI, P., CROWCROFT, J., DIOT, C., GASS, R. and SCOTT, J. (2007) Impact of human mobility on opportunistic forwarding algorithms. *Mobile Computing, IEEE Trans. on* 6(6): 606–620. doi:10.1109/TMC.2007.1060.
- [14] YOON, J. and LIU, M. AND NOBLE, B. (2003) Random waypoint considered harmful. In *Proceedings of the Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2003) (IEEE)*, 2: 1312–1321.
- [15] TSENG, Y.C., LI, Y.F. and CHANG, Y.C. (2003) On route lifetime in multihop mobile ad hoc networks. *IEEE Transactions on Mobile Computing* 2(4): 366–376.
- [16] NAVIDI, W. and CAMP, T. (2004) Stationary distributions for the random waypoint mobility model. *IEEE Transactions on Mobile Computing* 3(1): 99–108.
- [17] KRAAIER, J. and KILLAT, U. (2005) - the random waypoint city model -: user distribution in a street-based mobility model for wireless network simulations. In *WMASH '05: Proceedings of the 3rd ACM international workshop on Wireless mobile applications and services on WLAN hotspots (New York, NY, USA: ACM Press)*: 100–103. doi:10.1145/1080730.1080749.
- [18] ROJAS, A., BRANCH, P. and ARMITAGE, G. (2005) Experimental validation of the random waypoint mobility model through a real world mobility trace for large geographical areas. In *MSWiM '05: Proceedings of the 8th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems (New York, NY, USA: ACM Press)*: 174–177. doi:10.1145/1089444.1089474.
- [19] PIORKOWSKI, M., SARAFIJANOVIC-DJUKIC, N. and GROSSGLAUSER, M. (2009), CRAWDAD data set epfl/mobility (v. 2009-02-24), Downloaded from <http://crawdad.cs.dartmouth.edu>.
- [20] THIAGARAJAN, A., RAVINDRANATH, L., BALAKRISHNAN, H., MADDEN, S. and GIROD, L. (2011) Accurate, low-energy trajectory mapping for mobile devices. In *Procs. of the 8th USENIX Conf. on Networked Systems Design and Implementation, NSDI'11 (Berkeley, CA, USA: USENIX Association)*. URL <http://dl.acm.org/citation.cfm?id=1972457.1972485>.
- [21] JAIN, R., SHIVAPRASAD, A., LELESCU, D. and HE, X. (2004) Towards a model of user mobility and registration patterns. *SIGMOBILE Mob. Comput. Commun. Rev.* 8(4): 59–62. doi:10.1145/1052871.1052877.
- [22] HENDERSON, T., KOTZ, D. and ABYZOV, I. (2008) The changing usage of a mature campus-wide wireless network. *Computer Networks* 52(14): 2690–2712. URL <http://www.sciencedirect.com/science/article/pii/S1389128608001813>.
- [23] MULHANGA, M., LIMA, S. and CARVALHO, P. (2011) Characterising university w lans within eduroam context. In BALANDIN, S., KOUCHERYAVY, Y. and HU, H. [eds.] *Smart Spaces and Next Generation Wired/Wireless Networking (Springer), Lecture Notes in Computer Science* 6869, 382–394.
- [24] ASCHENBRUCK, N., MUNJAL, A. and CAMP, T. (2011) Trace-based mobility modeling for multi-hop wireless networks. *Computer Communications* 34(6): 704 – 714. doi:10.1016/j.comcom.2010.11.002, URL <http://www.sciencedirect.com/science/article/pii/S0140366410004755>.
- [25] TANG, D. and BAKER, M. (2000) Analysis of a local-area wireless network. In *Proc. of the 6th Annual Int'l Conf. on Mobile computing and networking (Mobicom'00) (ACM)*: 1–10.
- [26] DROMS, R. (1997) *Dynamic Host Configuration Protocol*. Request for Comments 2131, Internet Engineering Task Force.
- [27] RIGNEY, C. (2000), RADIUS Accounting, RFC 2866 (Informational). URL <http://www.ietf.org/rfc/rfc2866.txt>. Updated by RFCs 2867, 5080, 5997.

- [28] JOHNSON, D.B., MALTZ, D.A., BROCH, J. *et al.* (2001) DSR: the dynamic source routing protocol for multi-hop wireless ad hoc networks. *Ad hoc networking* 5: 139–172.
- [29] PERKINS, C., BELDING-ROYER, E. and DAS, S. (2003) *Ad hoc on-demand distance vector (AODV) routing*. Request For Comments 3561, IETF.