Delft University of Technology
Parallel and Distributed Systems Report Series

Human Mobility in Virtual and Real Worlds

Siqi Shen, Niels Brouwers, and Alexandru Iosup
{S.Shen,N.Brouwers,A.Iosup}@tudelft.nl

Draft version.

report number PDS-2013-xxx

ISSN 1387-2109

Published and produced by:
Parallel and Distributed Systems Section
Faculty of Information Technology and Systems Department of Technical Mathematics and Informatics
Delft University of Technology
Zuidplantsoen 4
2628 BZ Delft
The Netherlands

Information about Parallel and Distributed Systems Report Series:
`reports@pds.twi.tudelft.nl`

Information about Parallel and Distributed Systems Section:
`http://pds.twi.tudelft.nl/`

Abstract

The design and tuning of networked virtual environments (NVEs), such as World of Warcraft (WoW), require understanding the in-NVE mobility characteristics of their citizens. Although many mobility-aware NVE systems already exist, their validation and further development have been hampered by the lack of public datasets and of comparison studies based on multiple datasets. To address these two issues, in this work we collect from WoW mobility traces for over 30,000 virtual citizens, and compare these traces with traces collected from Second Life (SL) where the environment is designed and changed significantly by the citizens themselves. Furthermore, motivated by the existence of numerous studies and models of networked real-world environments (NRE), we systematically compare the characteristics of two NVE and two NRE mobility traces. Our comparative study reveals that long-tail distributions characterize well various mobility characteristics, that the invisible boundary of human movement also appears for NVEs, and that area-visitation shows personal preferences. We also find several differences between NVE and NRE mobility characteristics.

Contents

1 Introduction **4**

2 Background and Related work **5**
2.1 Terminology and mobility characteristics 5
2.2 Related work 5

3 Datasets **6**
3.1 Dataset Description 6
3.2 Data Collection and Sanitation 7
3.3 Data Sanitation 7

4 Characterization of Human Mobility **8**
4.1 Method for Distribution Fitting 8
4.2 Flight Length (C1) 8
4.3 Pause Duration (C2) 11
4.4 Area Popularity (C3) 12
4.5 Invisible Movement Boundary (C4) 14
4.6 Personal Preference in Area Visitation (C5) 15

5 Conclusion and Ongoing Work **17**

6 Appendix **22**

List of Figures

1 Flight length distribution of (left) the WoW dataset, and (right) the SL dataset. 9
2 Flight length distribution of (left) the GPS dataset, and (right) the GPS-2 dataset. 9
3 Distribution fitting of (left) WoW dataset, and (right) the GPS dataset. 10
4 Distribution fitting of (left) SL dataset, and (right) the GPS-2 dataset. 10
5 Pause duration distribution of (left) the WoW dataset, and (right) the SL dataset. 11
6 Pause duration distribution of (left) the GPS dataset, and (right) the GPS-2 dataset. 11
7 Distribution fitting of (left) WoW dataset, and (right) the GPS dataset. 12
8 Number of area visits of (left) the WoW dataset, and (right) the SL dataset. 13
9 Number of area visitors of (left) the WoW dataset, and (right) the SL dataset. 13
10 14
11 15
12 16
13 Normalized number of distinct visited areas of (left) the WoW dataset, and (right) the SL dataset. 16
14 Fitting results for (left) the WoW dataset, and (right) the GPS dataset. 17
15 Gini coefficient of personal preference weight (left) the WoW dataset, and (right) the SL dataset and GPS dataset. 17

List of Tables

1 Dataset overview. 6
2 Flight Lengths. 22
3 Pause Durations. 22
4 Number of vistors per area. 22
5 Number of visits per area. 22
6 Percentage of visited areas. 23
7 Personal preference weight Gini. 23
8 GPS visit times. 23

1 Introduction

Networked virtual environments (NVEs), including Massively Multiplayer Online Games (MMOGs) such as World of Warcraft (WoW), already serve tens of millions of users world-wide. Making the current and future NVEs more appealing to their citizens, more scalable to unexpected surges in temporal and spatial popularity, and more efficient in their resource use, depends on understanding user behavioral patterns (the input workload of any NVE system). Complementing much previous research in the design and tuning of NVE systems, and in particular in characterizing and modeling NVE workloads [1, 2], we focus in this work on the mobility of NVE citizens. To facilitate the design, validation, and comparison of mobility models and mobility-aware systems, and further motivated by the scarcity of public mobility datasets, we collect for this work a large-scale dataset from WoW and share it through the Game Trace Archive [3][1]. We also conduct a comprehensive, comparative characterization of the mobility of citizens in WoW and other, conceptually different NVE. Besides its findings, our characterization stresses the need for more mobility datasets to be shared among NVE researchers. Furthermore, we also do a high-risk, high-return investigation: motivated by the existence of datasets from networked *real-world* environments (NREs) and by the similarity between some NVEs (e.g., WoW) and NREs, we conduct a comparative analysis of mobility in NVEs and NREs.

Understanding in-NVE mobility can be useful to tune existing designs of NVEs and to innovate in the design of future NVEs. For example, recent advances in server cluster architectures [4] and peer-to-peer overlays [5] need to be validated against mobility workloads and, perhaps, tuned further to specific characteristics, e.g., their structure may need to be tuned to the area visitation characteristics, etc. For cloud-based architectures supporting NVEs such as [6], the load of various servers is strongly correlated with player mobility, due to player interaction [7, 8], cell visitation [4], etc. As has been shown in preliminary work on this topic [9], cloud-based workloads can be much more efficiently supported if the leasing of resources is in-tune with the workload.

NVE mobility is difficult to understand not only because public datasets are scarce, but also because NVEs cover a broad spectrum of applications. Among the most popular NVEs are MMOGs such as World of Warcraft and user-created NVEs such as Second Life (SL). For WoW, the game developer designs the virtual world to resemble a medieval, albeit fantasy-based, real-world environment. The citizens of WoW need to be highly mobile, to be able to finish quests of the storyline, trade goods, and socialize with the other players. Different from WoW, the virtual world of SL is created by the users themselves; this user-generated content should primarily foster socialization, collaboration, and even supervised learning. We pose and investigate the following research question: *How similar are WoW and SL avatar mobility patterns?* To answer it, we collect a new dataset of WoW mobility traces, and conduct a comprehensive and comparative study across multiple NVE datasets.

The scarcity of NVE mobility datasets is not paralleled by the existence of NRE mobility data. Although few NRE datasets are public, large-scale studies of millions of real-world citizens have appeared in the last decade [10, 11]. A high-risk, high-return idea would be to use these traces in NVE scenarios or even create NVE mobility models based on real world models, for example, when the NVE is by design similar to an NRE for which mobility is well understood, either spatially or w.r.t. the activities that users mostly engage in. WoW and many other NVEs have been designed starting from real-world cities (e.g., medieval cities), and equipped with traditional city-center functions such as meeting and trading. To immerse users, the movement of users in virtual worlds is designed to be as similar as possible to movement in the real world, albeit faster. The high-risk with using NRE traces in NVE scenarios is that the characteristics of NVE and NRE mobility may never match, in spite of the intents of the NVE designers. For example, real-world users do feel the physical effects of movement, including tiredness, legal restrictions, sometimes even cost, etc. The high-return is that the known NRE mobility traces are orders-of-magnitude larger than any of the NVE mobility traces previously reported, and there are many NRE mobility models already developed [12]. Thus, in this work we also set to answer the research question *How similar are the characteristics of mobility in NVEs and NREs?* In this work, we compare two NVE and two NRE mobility traces, and show evidence that their characteristics share many

[1] We will release the data before the conference.

common patterns. We also point out their main differences, which need future research before NRE mobility related research can be used in NVE studies.

In summary, our main contribution is twofold:

1. We collect a detailed and large-scale mobility dataset from the NVE World of Warcraft (Section 3). We plan to share the dataset via the Game Trace Archive [3].

2. We conduct a comprehensive study of human mobility characteristics in both virtual- and real-world environments (Section 4). The analysis in this work can help NVEs designers better planning resources and provide a base for building a mobility model for simulation.

2 Background and Related work

In this section, we introduce the terminology and compare previous mobility studies focusing on virtual environments with our work.

2.1 Terminology and mobility characteristics

We consider in this work the mobility of a population of individuals. In the rest part of this work, we use *mobility pattern* and *mobility characteristic* interchangeably. Following traditional mobility terminology [11], we define: *Citizens* (avatars, persons) are the moving entities. *Flight* is a straight-line trip without pause or significant directional change. The "angle model" of Rhee et al. [13] allows several consecutive straight line trips to be connected into a single flight if the angle between consecutive trips does not change the general direction of the flight. *Waypoints* (or locations) are the endpoints of a flight. *Pause duration* is the time spent by an individual in a waypoint.

In this work, we focus on five mobility characteristics which have been investigated in the past and shown to significantly affect the performance and reliability of NREs. Some of the characteristics have been also shown to have an impact of the performance of NVEs too. These characteristics are:

- **(C1) Long-tail distribution of flight lengths** [10]: Gonzalez et al. [10] find that the distribution of flight lengths for mobile phone users is long-tail. This means that human usually travel short distances and occasionally travel long distance.

- **(C2) Long-tail distribution of pause durations** [13], similarly to (C1).

- **(C3) Skewed popularity of areas** [14, 11]; for example, certain areas of cities are very popular, while others are rarely visited.

- **(C4) Invisible boundary of human movement:** By analyzing the trajectories of mobile phone users, Gonzalez et al. [10] found that human trajectories are bounded by a characteristic distance. Most of the time, people only travel between and around a few preferred locations.

- **(C5) Different personal preferences for areas** [15].

2.2 Related work

Much of the prior work [1, 16, 17, 2, 18] has focused on network measurement, online population, and session behavior. Kinicki and Claypool [2] find that walking in SL induces more network traffic than standing still. Chambers et al. [1] find that the online population of players has regular daily cycles. Since the late-2000s, several studies [19, 20, 21] collect and analyze mobility traces of NVEs. We compare our work with these, in the following.

Dataset	World Type	Citizens	Locations	Duration	Sampling
WoW	Virtual	31,290	4 cities	weeks	1/s
SL [19]	Virtual	26,714	5 regions	3 days	1/10s
GPS [23]	Real	1,366	3 cities	1 week	1/6s
GPS-2 [13]	Real	52	2 campuses	days	1/10s

Table 1: Dataset overview.

Closest to our work, Liang et al. [19] collect trace from SL, analyze the session behavior, contact patterns, and mobility patterns, in SL. For mobility patterns, they find that the number of visits to different cells of a region is skewed (C3), accumulated pause duration of avatars stay inside a cell is skewed (C2), and total travel distance of avatars is skewed (similar to C1). Our studies complement each other. We study the five characteristics (C1)–(C5) of two virtual world and two real world mobility traces, to find the similarity and difference of mobility between virtual- and real- world.

Pittman and GauthierDickey [14] analyze traces of two NVEs: WoW and Warhammer; they find that the popularity of different areas is skewed (C3), which they model using the Weibull distribution, and the durations each player stays in different zones vary. Varvello et al. [21] find that in SL, the popularity of zones is skewed (C3), and about half of the players form groups of good friends who meet frequently at the same locations (similar to C5). In contrast to [21], Miller and Crowcroft [22] find that, in their observation scenario of WoW battleground, most movement is individual rather than group-based. La and Michiardi [20] investigate (C3) and mobile communication related metrics such as inter-contact time, using traces collected from SL, and use the traces to evaluate the performance of wireless network protocols.

3 Datasets

In this section we introduce the four datasets used in this work. We describe the datasets, and the data collection and sanitation processes. The limited infrastructure and development effort we have used for collecting our dataset give evidence that *collecting mobility traces from virtual worlds can lead to datasets of large size and fine detail, while making the collection process feasible for many researchers.*

3.1 Dataset Description

To understand human mobility, we have collected a very large and detailed dataset from a popular virtual world, World of Warcraft, and used four other *public* datasets that were collected by others from virtual and real worlds. The characteristics of these six datasets are summarized in Table 1.

`WoW` *dataset, technical details:* Similarly to most large MMOGs, WoW is not built as a single, seamless virtual environment. Instead, it operates concurrently a few hundreds of independent realms of identical design. Each world is serviced by a server. We have collected mobility traces from 3 capital cities (`Ironforge`, `Stormwind`, and `Orgrimmar`) of the popular Silvermoon realm (server located in Europe), accounting for about 60% of the users, and from 1 capital city (Stormwind, we call it `Stormwind-2` from now on) of the popular Argent-Dawn realm (server located in Europe).

The `SL` dataset was collected by Liang et al. [19] from Second Life, a popular user-created virtual world, by using a custom SL client. We use in this work the publicly available dataset of about 27,000 virtual citizens distributed across 5 non-connected regions.

The `GPS` observe the mobility of volunteers. For the `GPS` dataset, Bohte and Maat [23] have distributed GPS devices to over 1,000 volunteers in 3 Dutch cities, and record per-citizen locations every 6 seconds for 1 week.

The `GPS-2` dataset, Rhee et al. [13] have recorded 125 GPS tracks from 52 citizens; this public dataset does not allow us to map trajectories to individuals. Rhee et al also collected GPS trajectories from Disney world and State Fair, we do not analysis them in these paper.

Overall, the two datasets collected from virtual worlds follow each several tens of thousands of citizens, that is, one or two orders of magnitude more than the datasets collected from real worlds. Our dataset is large-scale (over 30,000 citizens) and multi-location (4 cities); it was also collected using fine-grained sampling (1 sample every second) over a significant period (several weeks). Thus, ours is one of the largest and most detailed datasets available to researchers, to date[2].

3.2 Data Collection and Sanitation

We have collected the `WoW` dataset from the virtual environment provided by World of Warcraft (WoW), a popular massively multiplayer online game (MMOG) with over 2,000,000 daily active players. WoW is designed to emulate a medieval-like environment with fantastic elements, and contains cities and wildland. WoW citizens travel in this world in real time, with the movement speed limited to 8–10 in-game distance units per second. In WoW, each virtual citizen can observe the presence and activity of any other virtual citizen within a radius of about 100 in-game distance units (10–15 seconds of movement away); unlike the real world, the observation range in WoW is not affected by interposing objects such as buildings or other citizens. Similarly to most large MMOGs, WoW is not built as a single, seamless virtual environment. Instead, it operates concurrently a few hundreds of independent worlds of identical design. Each world is serviced by a server cluster; we will investigate this architecture in Section **??**. We have collected mobility traces from 3 capital cities (Stormwind City, Ironforge, and Orgrimmar) of the popular Silvermoon world, accounting for about 60% of the users, and from 1 capital city (Stormwind City) of the popular Argent-Dawn world.

To collect the `WoW` dataset, we have developed a customized WoW client and used it to observe a selected number of cities. To observe mobility in a city, we deploy virtual citizens such that the areas they observe cumulatively overlap the surface of the city. Our client logs-in several WoW clients and coordinates them to observe a complete city. The client uses 6 machines per measurement, each running several WoW clients and collecting their observed data. Due to the availability of machines, which are regular PCs used for coursework at our university, during week days we can only collect data during the night. In total, we have obtained data for 3 complete week-ends and about 20 week-day evenings during April and September, 2011, resulting in mobility information for over 30,000 virtual citizens (see Table 1).

3.3 Data Sanitation

The datasets used in this work may include incorrect data. We describe in the following the steps we followed to remove incorrect data. `SL, GPS-2, iMote` are sanitated by their owners.

For GPS-based datasets, the locations recorded by GPS devices may be incorrect: we have observed in these datasets human speeds faster than the speed of sound. We have removed all locations that could only result from non-human speed.

For our `WoW` dataset, data are collected simultaneously from multiple clients (sensors) operating from multiple machines. Although the machines are synchronized, the locations recorded by the sensors whose circular areas of observation overlap may be different, due to network delays between the local sensor and the remote WoW server. We aggregate the recorded locations in overlapped areas and use the average value of locations recorded within the same second. Less than 1% of these aggregates involve ranges of values exceeding 10 in-game units (1 second of movement).

For GPS-based datasets, the locations recorded by GPS devices may be incorrect: we have observed in these datasets human speeds faster than the speed of sound. We have removed all locations that could only result from non-human speed.

[2] We will make our dataset publicly available prior to the conference.

4 Characterization of Human Mobility

In this section, we answer the question *How similar are WoW and SL avatar mobility traces?* and *How similar are virtual and real-world human mobility traces?* To answer this question, we investigate the characteristics (C1)–(C5) (see Section 2.2) for `WoW`, `SL` and `GPS` (Section 3). We only investigate (C1)–(C2) for `GPS-2` due to the relative small sample sizes and the mobility of citizens are limited to campus scenarios.

Where the datasets comprise multiple locations, we analyze both the entire dataset and each location, in turn. Unless otherwise noted, we have obtained similar results for each investigated dataset. To study characteristics among different traces, we look at the basic descriptive statistics, and then use the distribution fitting method (described in Section 4.1) to look at the trend and distribution of data. We present here only a selection of representative results.

Our main finding is that the mobility characteristics for the two virtual world (`WoW`, `SL`) traces have many similarities. The flight length (C1), pause duration (C2), and area popularity (C3) follow long-tail distributions; avatars only visit a small portion of virtual cities (C4); and preference to visit only a few, preferred areas does exists (C5). In comparison, for `GPS`, the flight length is longer; and the personal preference to some areas is higher.

4.1 Method for Distribution Fitting

We now describe the method used to create piecemeal statistical models for each characteristic, per trace.

For each trace considered in this work, we attempt to fit the empirical data corresponding to each characteristic (C1)–(C5) with a set of well-known probability distributions that are available in most simulation and experimental toolboxes, namely the power-law, the exponential, the Weibull, the log-normal, the gamma, the normal, the general Pareto, and the truncated power-law (or truncated Pareto) distributions(the upper cutoff of data is set to be 99.9% quartile value). The fitting is performed using *maximum likelihood estimation* (MLE) [24], which determines for a distribution the parameters that lead to the best fit with given empirical data. For number of visitors(C3), we first obtained the best fitted continuous distribution, and then fit the data with discrete version of best fitted distribution to improve accuracy. (see following paragraphs for the definition of "best fitted").

Then, we use a method for assessing the *goodness-of-fit* (GoF) that has been shown to have good results for large datasets in distributed systems studies [25, 26]. In this method, the results of MLE fitting are tested using a *goodness-of-fit* (GoF) procedure that combines the the *Kolmogorov-Smirnov* (KS) and the *Anderson-Darling* (AD) GoF tests. Using both of these tests provides a more robust GoF test than using any of the KS and AD tests individually, since the KS test is more sensitive to the center of distributions and the AD test is more sensitive to the tail. The method uses 0.05 as the significance level for the *p-value*, below which the null hypothesis that the fitted distribution represents the empirical data is rejected. The p-value used by this method is the average of 1,000 p-values, each of which is calculated by randomly selecting 30 samples from the empirical data and applying the GOF tests to the selected data.

Last, in our distribution fitting method we we consider a probability distribution as a *good fit* for the studied properties only if that distribution passes both the KS and AD tests for all of our three datasets. The *best fit* of a property is, from a set of distributions that are good fits, the distribution that has the smallest Akaike information criterion with correction (AICc) [27], which takes into account the number of parameters and the likelihood of fit.

4.2 Flight Length (C1)

Figure 1 shows the cumulative distribution function (CDF) of the flight lengths of `WoW` (left) and `SL` (right). The flight lengths of `WoW` traces are long-tail distributed and the flight lengths for all four cities are similar. The mean values of flight lengths in the four virtual cities are around 20 to 25 meters. Most (about 85% to 90%) of

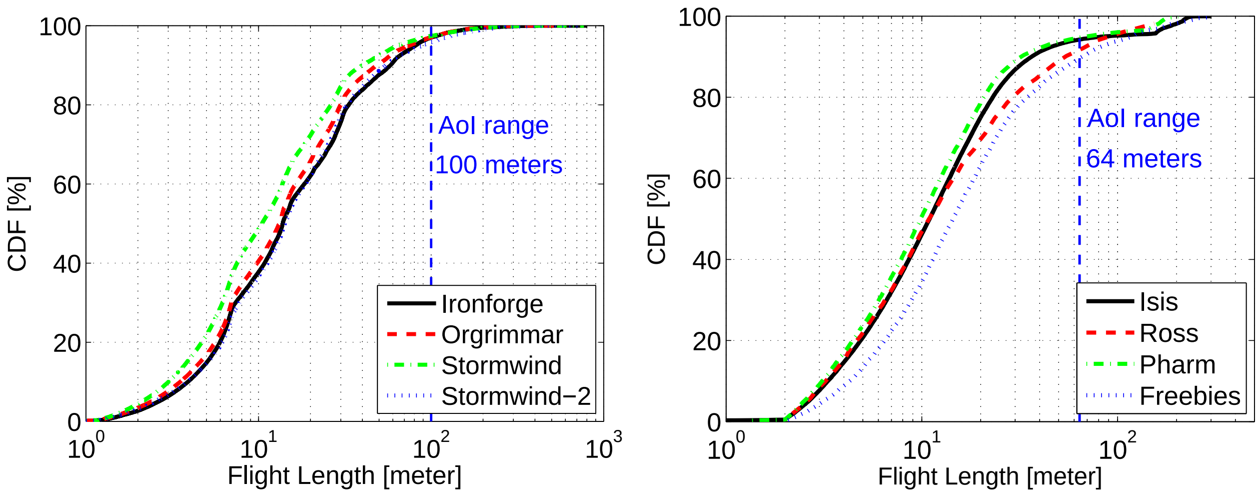

Figure 1: Flight length distribution of (left) the `WoW` dataset, and (right) the `SL` dataset. (Logarithmic scale on horizontal axis.)

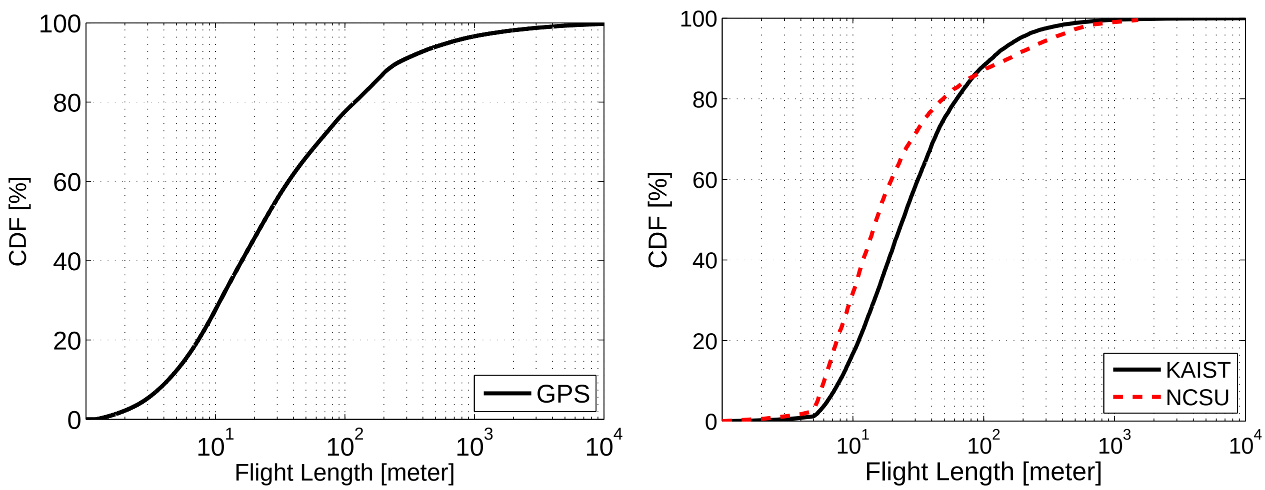

Figure 2: Flight length distribution of (left) the `GPS` dataset, and (right) the `GPS-2` dataset. (Logarithmic scale on horizontal axis.)

the flights are shorter than the area of interest (AoI) range (100 meters) of WoW. For the `SL` traces, the mean values of flight lengths in the four zones are around 19 to 29 meters. Most (80% to 90%) of the flight lengths are shorter than the AoI range (64 meters). This may suggest that when avatars travel in virtual worlds, most of them travel within the boundary of AoI, and occasionally avatars travel long distances.

As Figure 2 shows, for the two real world datasets: the mean value of flight lengths of `GPS` is 215 m, while the mean values of flight lengths for `KAIST` and `NCSU` are 61 m and 71 m, respectively. The flight lengths of the two real world datasets have longer tail than the two virtual world datasets: the 99% percentiles for the two virtual world datasets are about 150 m to 230 m, while the 99% percentiles for the two real world datasets are about 600 m to 4,000 m.

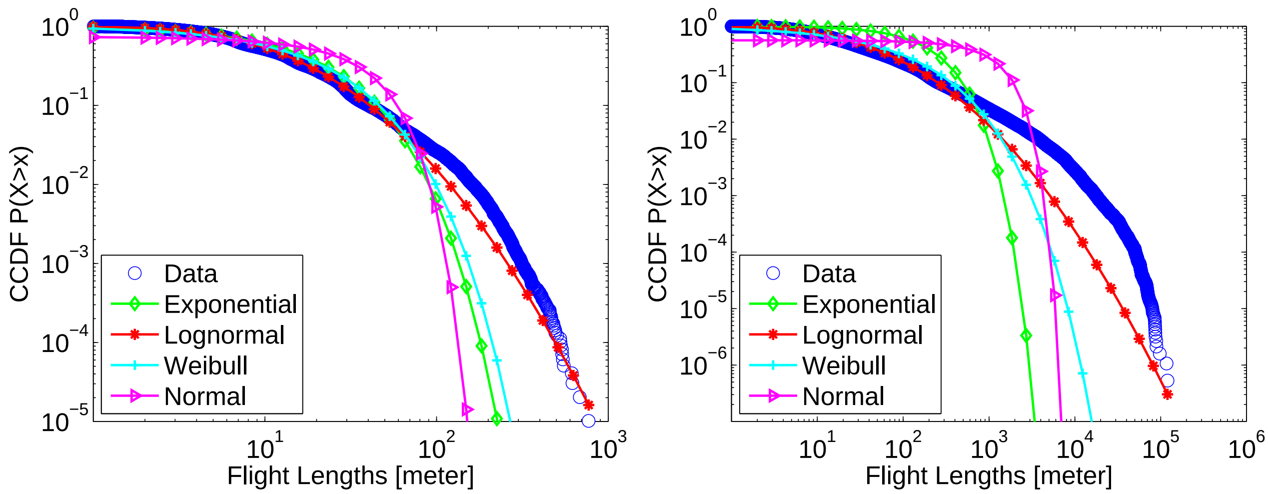

Figure 3: Distribution fitting of (left) `WoW` dataset, and (right) the `GPS` dataset. (All axes use logarithmic scales.)

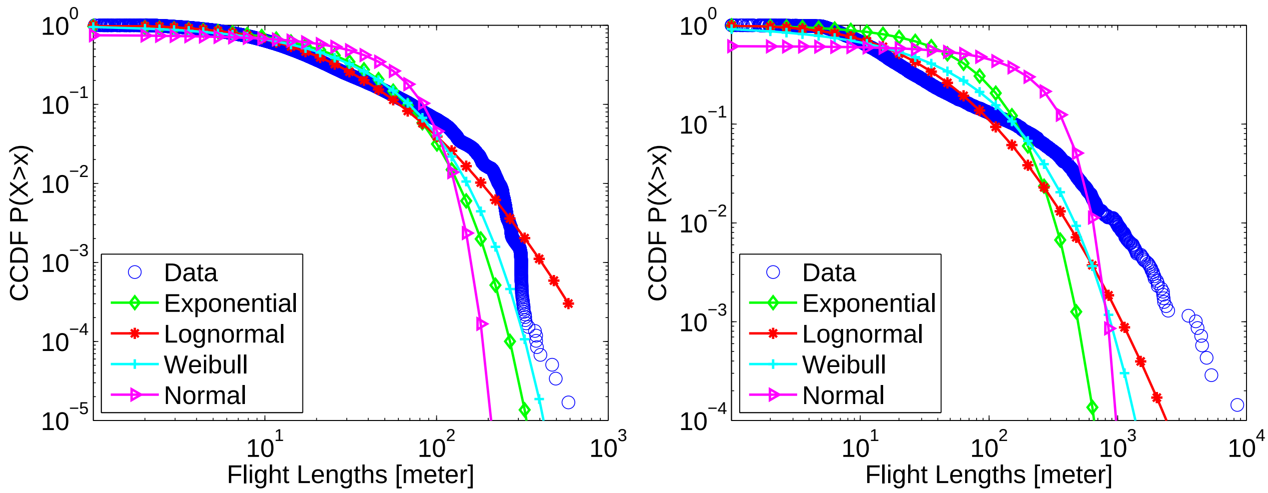

Figure 4: Distribution fitting of (left) `SL` dataset, and (right) the `GPS-2` dataset. (All axes use logarithmic scales.)

Figure 3 depicts the results of fitting for `Stormwind`, `WoW` and `GPS`, while Figure 4 shows the fitting results for `Freebies`, `SL` and `NCSU`, `GPS-2`. The vertical axis shows the complementary cumulative distribution function (CCDF) of the flight lengths, in logarithmic scale (Note that the scales of the two figures are different). We find that the best fit for `Stormwind` is LogNormal distribution (mean $\mu = 2.4$ deviation $\sigma = 1$). For the `GPS` data, the best fit is a LogNormal distribution ($\mu = 3.4$ $\sigma = 1.65$) (the distribution fitting diverge a bit when the flight lengths are higher than $1,000\,m$). The flight lengths distributions for the two virtual world datasets (`WoW` and `SL`) and the two real world datasets (`GPS` and `GPS-2`) follow long-tail distributions, and all of them can be best fitted using the LogNormal distribution.

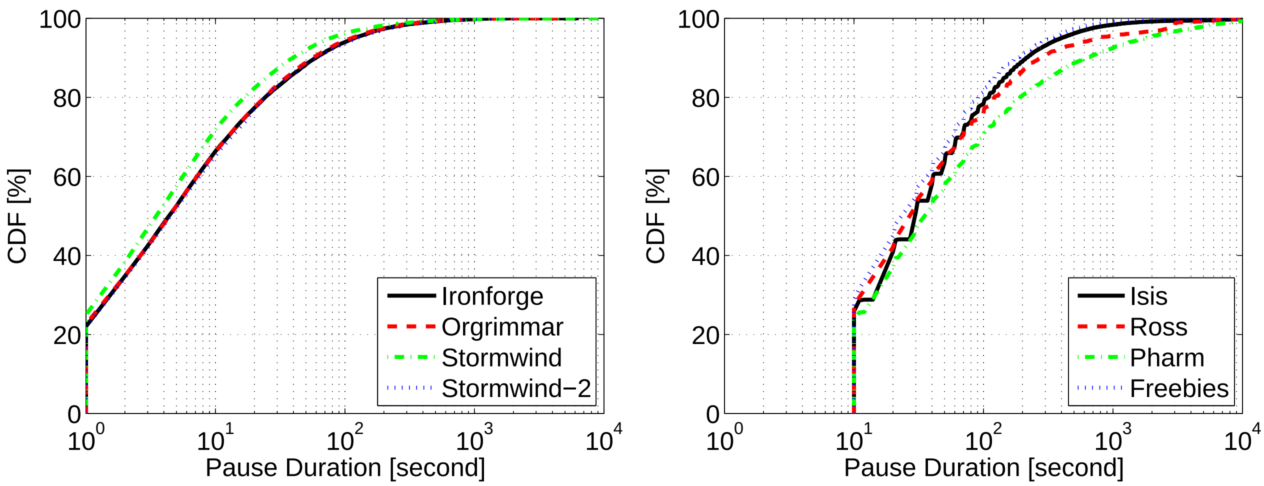

Figure 5: Pause duration distribution of (left) the `WoW` dataset, and (right) the `SL` dataset. (Logarithmic scale on horizontal axis.)

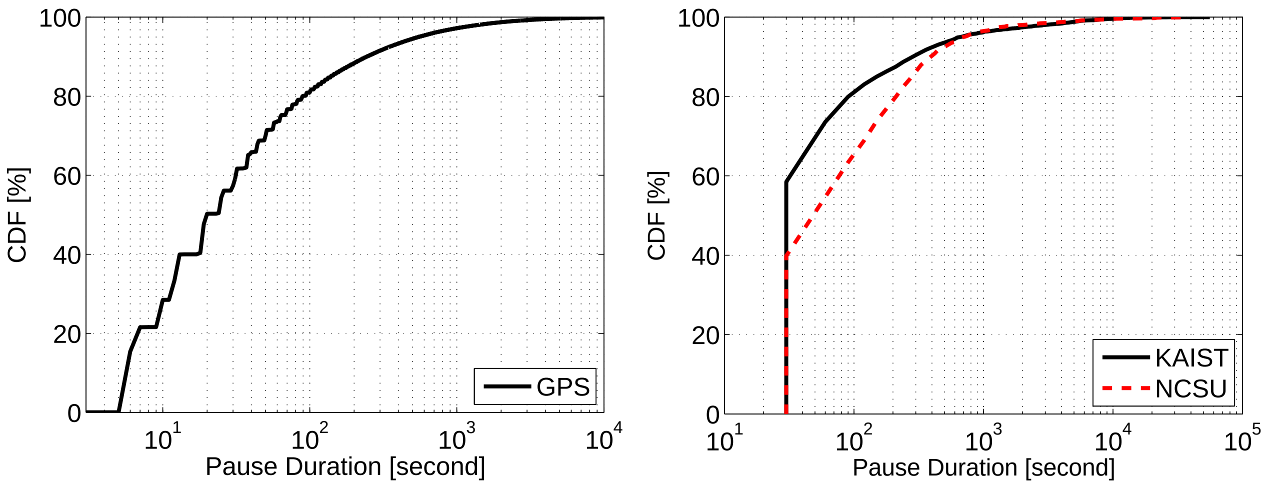

Figure 6: Pause duration distribution of (left) the `GPS` dataset, and (right) the `GPS-2` dataset. (Logarithmic scale on horizontal axis.)

4.3 Pause Duration (C2)

Figure 5 shows the pause durations distribution of the `WoW` and the `SL` datasets. Overall, the pause durations of both datasets are long-tail, about 80% of the pause durations of `WoW` is shorter than 30 seconds. The pause duration of `Stormwind` is slightly lower than the other three cities, while the other three have very similar distributions. For the `SL` dataset, about 70% to 80% of the pause durations is shorter than 100 seconds. The pause durations for the `Pharm` zone is higher than the other because the main activities of that zone is camping (staying in the same location). The pause durations of the `WoW` datasets are significantly shorter than the `SL` datasets. The difference of the pause durations for the two datasets may be caused by the design of the two NVEs: `SL` focuses more on social aspects, while `WoW` is more task-oriented and the interactivity between

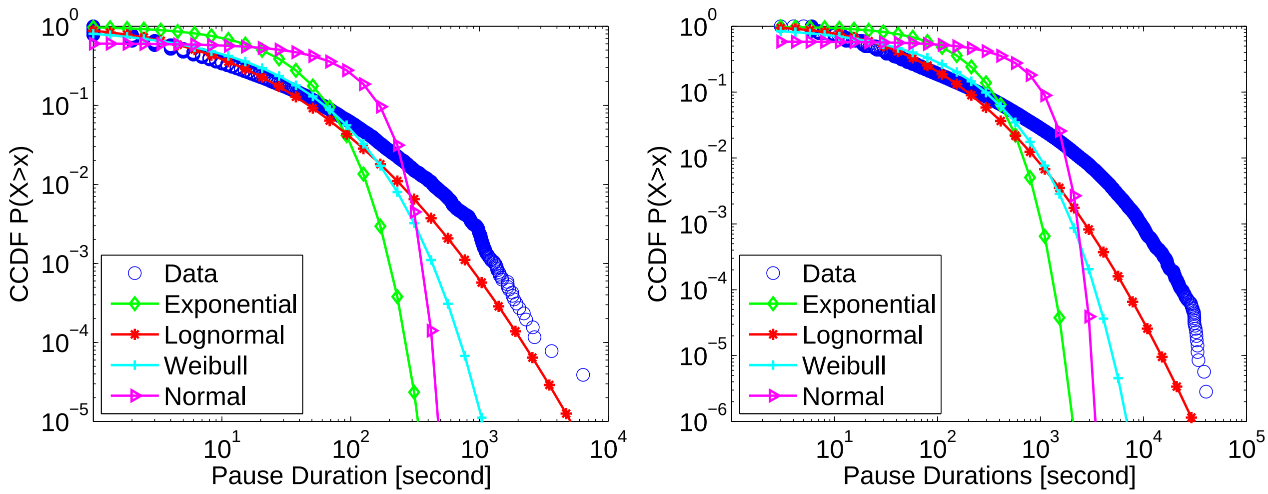

Figure 7: Distribution fitting of (left) `WoW` dataset, and (right) the `GPS` dataset. (All axes use logarithmic scales.)

players is more frequent.

As Figure 6 shows, for the real-world datasets, the pause durations for those datasets are long-tail too. The average pause duration of the `GPS` dataset is about 2.5 minutes, and the 99% percentile of pause durations is about 40 minutes. For the `GPS-2` dataset, the mean values range from 5.5 to 6 minutes, and the 99% percentiles are around 1.5 hours.

Figure 7 depicts the results of distribution fitting for the `Stormwind`, `WoW` and `GPS`. The pause durations observed in `Stormwind` can be best modeled using the LogNormal distribution ($\mu = 1.63$ $\sigma = 1.45$). The fitting result for the `GPS` dataset is the LogNormal distribution ($\mu = 3.41$ $\sigma = 1.44$). In summary, the pause durations distributions for the two virtual world datasets: `WoW` and `SL` and the two real world datasets: `GPS` and `GPS-2` follow long-tail distributions, and all of them can be best fitted using the LogNormal distribution.

4.4 Area Popularity (C3)

To investigate area popularity, we first split the environments into rectangular grids, where each cell is an area. Rectangular grids are convenient for setting up simulation scenarios and may enable fair comparison between different city scenarios. We explore different values for the size of each area, which is the parameter of the splitting procedure; we split maps into areas of $10\,m \times 10\,m$ up to $50\,m \times 50\,m$. For each area size, we quantify the popularity of the resulting areas using two main indicators: the *number of area visits*, defined for each area as the number of pauses observed in that area; and the *number of area visitors*, defined for each area as the number of unique visitors paused in that area. Intuitively, the former indicator quantifies the total traffic through an area well, whereas the latter does not account for returning visitors. Although these two indicators may also depend on the period over which they are observed, for our datasets we have found that a period of 1 day, which is a typical period for human activity, is sufficient for contouring the distributional shape.

Number of area visits: Figure 8 (left) shows the number of area visits for `Ironforge`, by splitting the map into areas of $10\,m \times 10\,m$, $20\,m \times 20\,m$, and $50\,m \times 50\,m$. The visitation count increases with the increasing size of the areas. Large portions of the map are not visited at all, about 75% of the $10\,m \times 10\,m$ areas are not visited once, and about 40% of the $50\,m \times 50\,m$ areas are not visited. The visitation count is long-tail; for $10\,m \times 10\,m$ areas, the 85% percentile is 10 while the maximal value is about 1,921. Figure 8 (right) shows the results for `SL`, when splitting the map into areas of $10\,m \times 10\,m$ size. Similarly to `WoW`, large parts of the map are not visited, 3 out of 4 zones have 80% unvisited areas; and the distribution of the number of area visits of `SL` is long-tail.

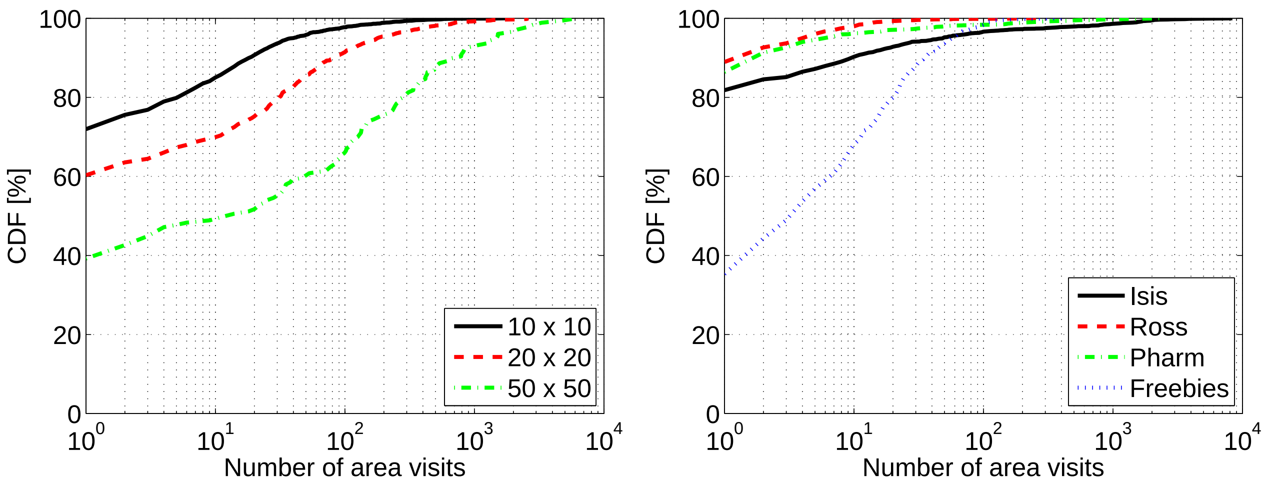

Figure 8: Number of area visits of (left) the WoW dataset, and (right) the SL dataset. (Logarithmic scale on horizontal axis.)

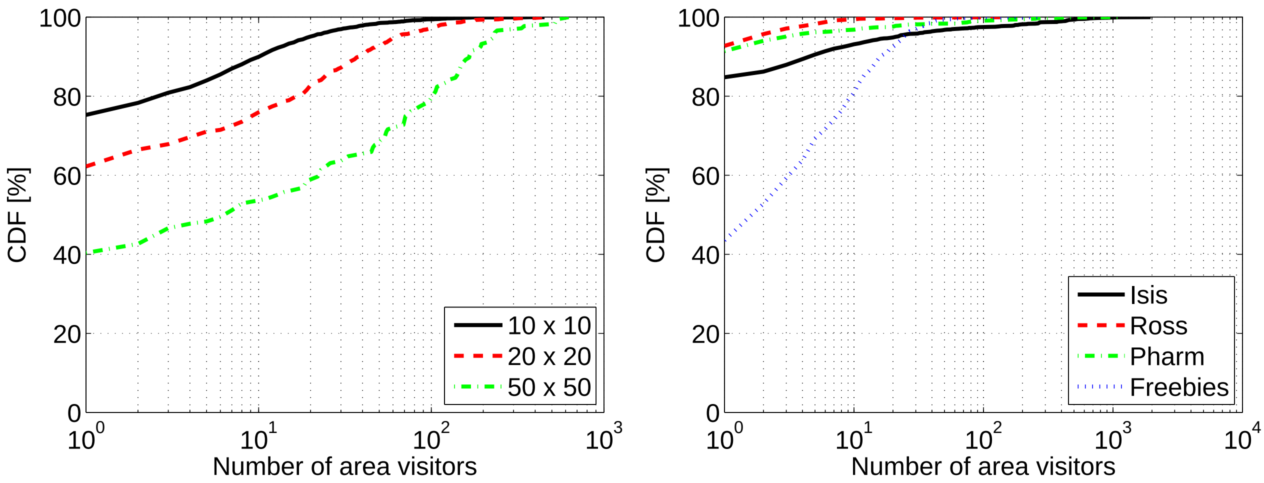

Figure 9: Number of area visitors of (left) the WoW dataset, and (right) the SL dataset.

The number of area visits for GPS is long-tail too, when it is partitioned into areas of $10\,m \times 10\,m$, over 99% of the areas is not visited at all, while the most popular area is visited about 900 times.

Number of area visitors: Figure 9 shows the number of area visitors for Ironforge and SL. The number of area visitors is smaller than the number of area visits, but it is long-tail too. Figure 9 (left) shows the number of visitors for Ironforge, by splitting the map into areas of $10\,m \times 10\,m$, $20\,m \times 20\,m$, and $50\,m \times 50\,m$. For $10\,m \times 10\,m$ areas, the 85% percentile is 6 while the maximal value is about 453. Figure 9 (right) shows the results for SL, when splitting the map into areas of $10\,m \times 10\,m$ size. Similar to WoW, large parts of the maps are not visited, and the distributions of the number of area visitors for SL are long-tail. For the GPS dataset, when it is partitioned into areas of $10\,m \times 10\,m$, the most popular areas is visited by 80 persons, and when it is partitioned into areas of $50\,m \times 50\,m$, the most popular area is visited by 173 persons. The distributions of the number of area visitors for the WoW, SL, and GPS datasets are long-tail.

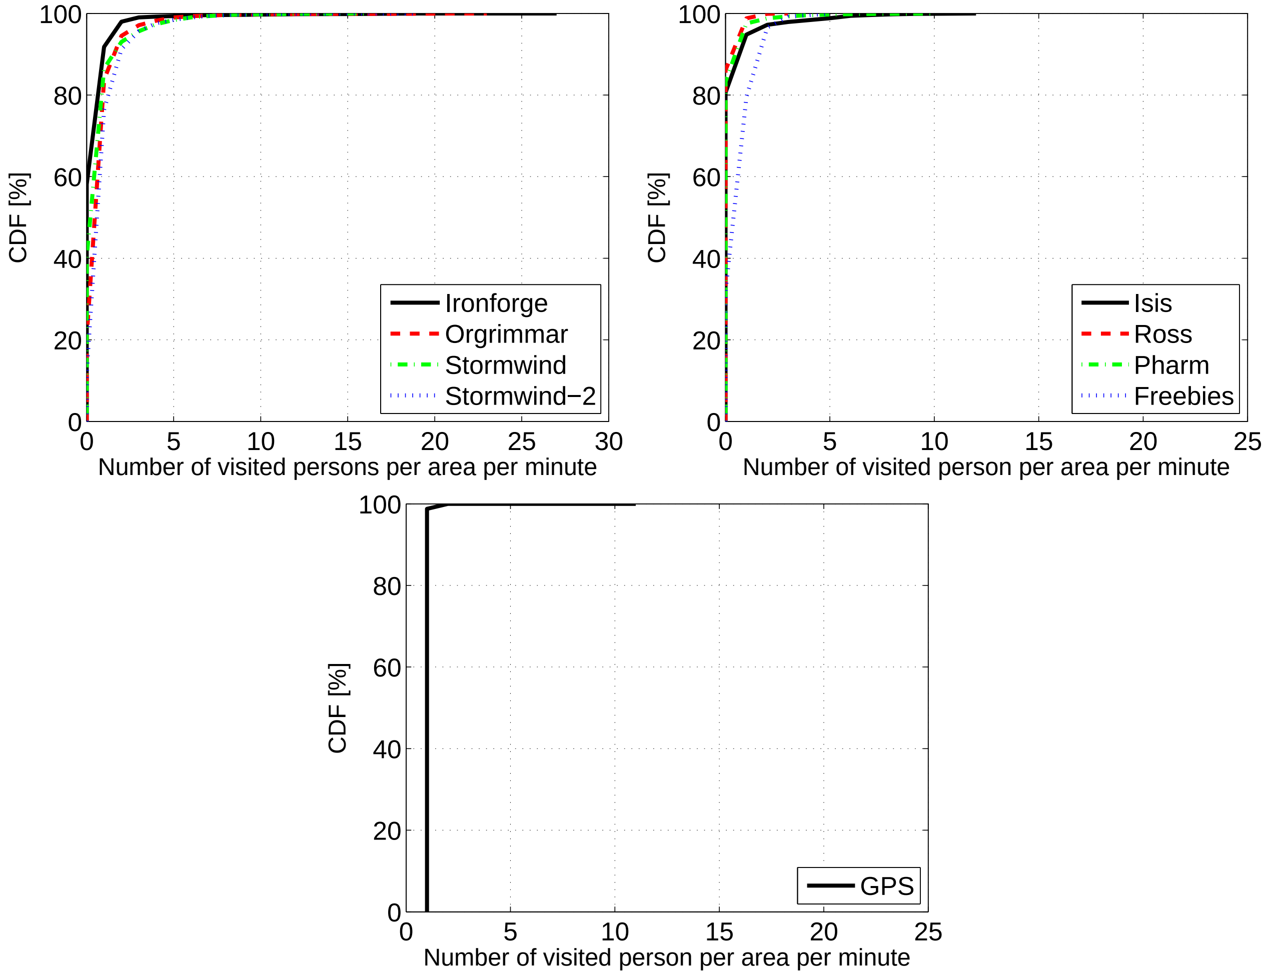

Figure 10: Maximal number of visited person per minute per area of (left) the `WoW` dataset, and (right) the `SL` dataset.

Maximal number of visitors per minute:

4.5 Invisible Movement Boundary (C4)

We now look at the invisible movement boundary, that is, the phenomenon that humans tend to travel mostly within a fixed and reduced set of locations around home and office (see Section 2.2). We find that *the invisible movement boundary is present in both real and virtual worlds*. To quantify the boundary, we use the proxy metric normalized number of distinct visited areas, measured per person. Figure 13 shows the number of distinct areas, normalized by the total number of visited areas per map. The higher this value is, the higher the probability of avatars meeting each other. This metric can be useful for modeling mobility: when generating waypoints on maps, the model can limit the avatar to visit only a small subset of waypoints. As Figure 13 shows, for `WoW` and `SL`, the normalized number of distinctive areas is low. Most (about 95%) of the avatars visit less than 5% of the visited areas; only a few persons visit more than 10% of the visited areas. In average, each avatar visits about 0.4 to 1% of the areas in the `WoW` dataset; and in `SL`, each avatar visits about 1.2 to 2% of the areas.

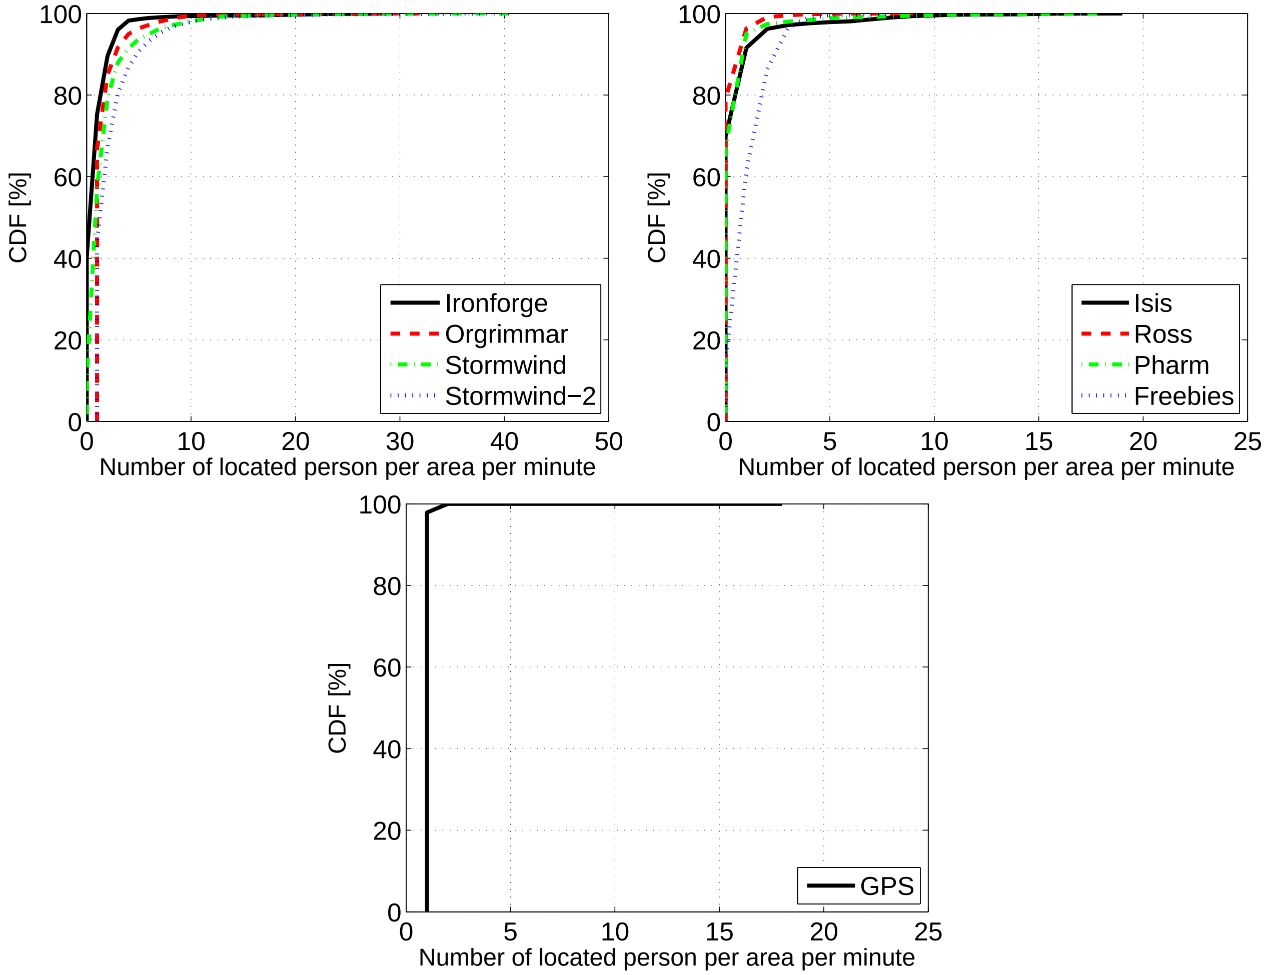

Figure 11: Maximal number of located person per minute per area of (left) the `WoW` dataset, and (right) the `SL` dataset.

For the `GPS` dataset, most (about 95%) of the avatars visit less than 0.5% percent of the visited areas. We attribute the significantly lower values for the `GPS` data to the fact that the real world cities are much bigger than the virtual world cities: the `GPS` dataset cover a map about $30\,km \times 30\,km$, while the largest virtual cities in the `WoW` and `SL` is smaller than $2\,km \times 2\,km$. For the empirical distributions: the normalized number of distinct visited areas for `Ironforge` can be fitted best by the Weibull distribution (scale $a = 2.34$, shape $b = 1.99$), and the best fit for the `GPS` dataset is the Weibull distribution ($a = 0.12$, $b = 2$). However, the `SL` traces can be better modeled using the LogNormal distribution.

4.6 Personal Preference in Area Visitation (C5)

In SL, some avatars like to visit the same group of persons [21]; and real-world citizens have strong preferences for different areas [23]. We study the personal preferences of virtual and real- world in this section. For each of the area the avatars visited, we count the number of time the avatar visited that area as *personal preference weight*. Then for each person, we calculate the Gini coefficient (also called Gini index) of the personal preference

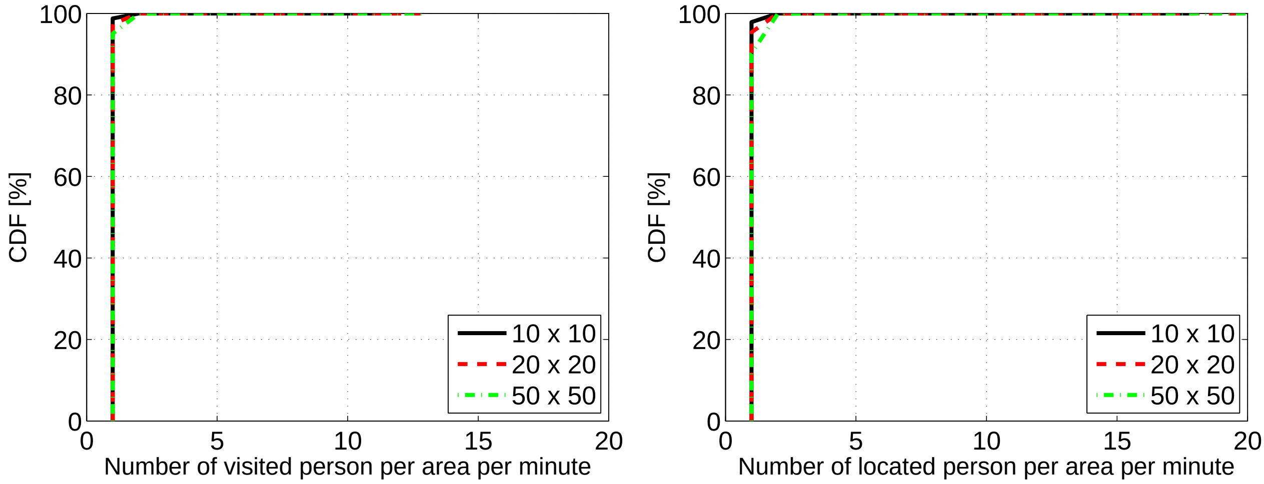

Figure 12: Popularity GPS.

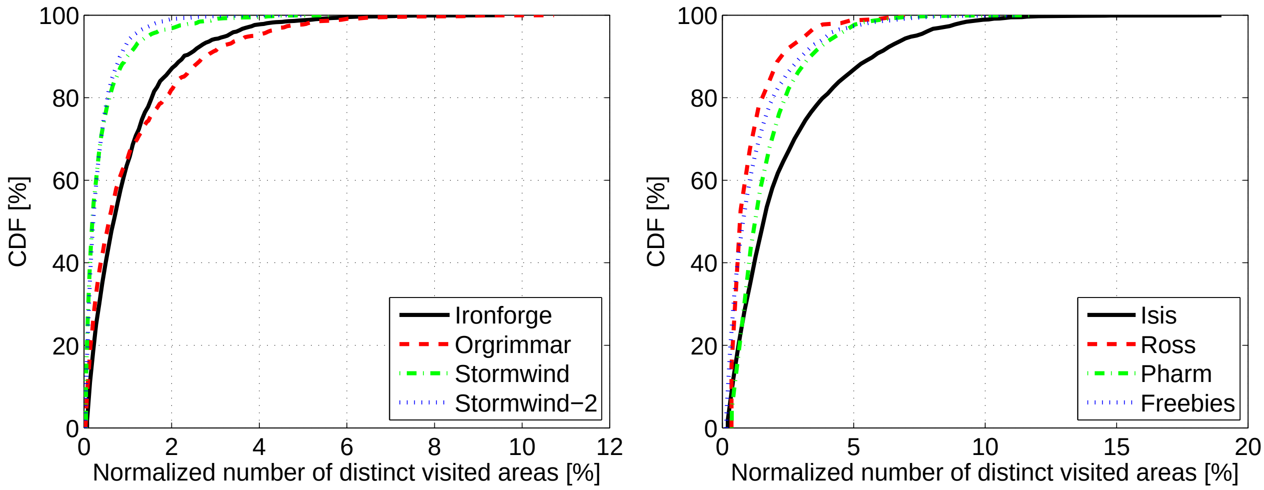

Figure 13: Normalized number of distinct visited areas of (left) the WoW dataset, and (right) the SL dataset.

weight. The Gini coefficient is used to quantify the inequality of personal preference (a value of 1 means very unequal, whereas 0 means perfectly equal).

Figure 15 shows the Gini coefficient distribution of each person for WoW, SL, and GPS. For this figure, we remove the persons that visit less than 5 areas (the result is similar without removal). In general, the two virtual world datasets have similar Gini coefficient distributions: most (80% to 95%) of the Gini coefficients are lower than 0.4. For the GPS dataset, about 40% are higher than 0.4. The probability distribution functions of the Gini coefficients for all datasets are bell-shape curves, can be modeled using the Weibull distribution. For reference, we also generate personal preference weights using the power-law distribution with exponent α range from 2.5 to 4; the Gini coefficients of these weights are depicted as three vertical lines labelled as $\alpha = 2.5, 3, 4$ in Figure 15. For mobility modeling purpose, we find that if each individual assigns the personal preference weights based on a power-law distribution, then most of the exponents of the power-law distribution are between 2.5 to 4.

The Gini coefficients of the personal weights in GPS dataset is higher than in the two virtual world datasets.

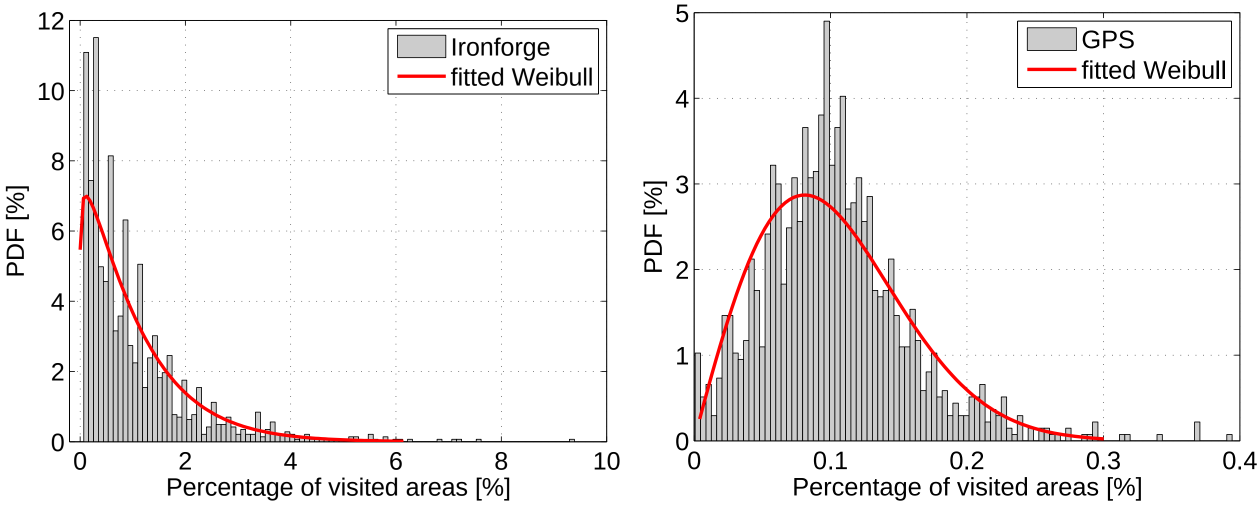

Figure 14: Fitting results for (left) the WoW dataset, and (right) the GPS dataset.

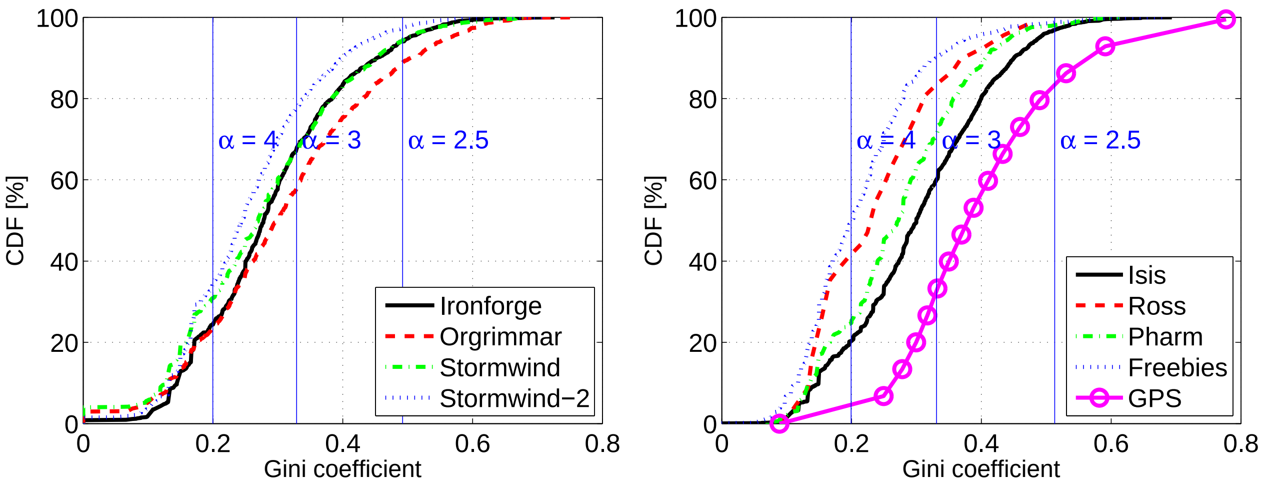

Figure 15: Gini coefficient of personal preference weight (left) the WoW dataset, and (right) the SL dataset and GPS dataset.

This may suggest that the personal preference for areas is stronger in real-world environments than in virtual worlds, and has higher predictability in real-world human mobility than in virtual-world avatar mobility. As possible explanations, we point to the higher rate of movement, to the less restrictive of movement, and to other lower penalties for movement (legal restrictions, cost, etc.) in virtual vs real-world mobility.

5 Conclusion and Ongoing Work

Understanding the characteristics of and modeling human mobility is important for the design and tuning of modern networked multimedia systems. Our main contribution is the collection, characterization, and modeling of mobility traces. We have collected detailed position information of virtual world mobility traces from a popular networked virtual environment (NVE), World of Warcraft. We have shown that the human mobility traces collected from virtual- and real-world environments have many similar characteristics and analyze the main differences. We have developed a new human mobility model, SAMOVAR, which can generate synthetic yet realistic traces for virtual and real world environments. We have validated our model and compared it through simulations with several other mobility models. Our simulation study includes an in-depth study of the impact of human mobility characteristics on the performance of NVE. Last, we have shown evidence that SAMOVAR leads to useful insights into the performance of NVEs and wireless network protocols. We are currently using and extending SAMOVAR for the design and tuning of NVE algorithms.

Acknowledgments

The datasets used in this work were gracefully provided by the authors of [28, 29, 13, 23]. We would like to thank Kees Maat, Hugo Ledoux and Liu Liu to allow us to use their GPS dataset. The authors would further like to thank Otto Visser for his help regarding World of Warcraft, Dick Epema for his comments and Koen Langendoen for supporting this work.

References

[1] C. Chambers, W.-C. Feng, S. Sahu, D. Saha, and D. Brandt, "Characterizing online games," *ToN*, vol. 18, no. 3, 2010. 4, 5

[2] J. Kinicki and M. Claypool, "Traffic analysis of avatars in Second Life," in *NOSSDAV*, 2008, pp. 69–74. [Online]. Available: http://dl.acm.org/citation.cfm?id=1496063 4, 5

[3] Y. Guo and A. Iosup, "The game trace archive," in *NetGames*, 2012. 4, 5

[4] Y. Deng and R. W. H. Lau, "On delay adjustment for dynamic load balancing in distributed virtual environments," *TVCG*, vol. 18, no. 4, 2012. 4

[5] S.-Y. Hu, J.-F. Chen, and T.-H. Chen, "VON: a scalable peer-to-peer network for virtual environments," *IEEE Netw*, vol. 20, no. 4, 2006. 4

[6] C.-Y. Huang, C.-H. Hsu, Y.-C. Chang, and K.-T. Chen, "GamingAnywhere: an open cloud gaming system," in *MMSys*, 2013. 4

[7] D. T. Ahmed and S. Shirmohammadi, "Improving online gaming experience using location awareness and interaction details," *Multimedia Tools Appl.*, vol. 61, no. 1, 2012. 4

[8] V. Nae, A. Iosup, and R. Prodan, "Dynamic resource provisioning in massively multiplayer online games," *IEEE TPDS*, no. 3, pp. 380–395, 2010. 4

[9] Y.-T. Lee and K.-T. Chen, "Is server consolidation beneficial to mmorpg? a case study of world of warcraft," in *CLOUD*, 2010. 4

[10] M. C. González, C. Hidalgo, and A.-L. Barabási, "Understanding individual human mobility patterns." *Nature*, vol. 453, no. 7196, 2008. 4, 5

[11] R. Becker et al, "Human mobility characterization from cellular network data," *Commun. ACM*, vol. 56, no. 1, 2013. [Online]. Available: http://doi.acm.org/10.1145/2398356.2398375 4, 5

[12] A. Keränen, T. Kärkkäinen, and J. Ott, "Simulating Mobility and DTNs with the ONE," *JCM*, vol. 5, no. 2, 2010. 4

[13] I. Rhee et al., "On the Levy-Walk Nature of Human Mobility," *Infocom*, 2008. 5, 6, 19

[14] D. Pittman and C. GauthierDickey, "Characterizing virtual populations in massively multiplayer online role-playing games," *AMM*, 2010. [Online]. Available: http://link.springer.com/chapter/10.1007/978-3-642-11301-7_12 5, 6

[15] C. Song et al., "Modelling the scaling properties of human mobility," *Nature Phys.*, vol. 6, no. 10, 2010. 5

[16] S. Choy, B. Wong, G. Simon, and C. Rosenberg, "The Brewing Storm in Cloud Gaming: A Measurement Study on Cloud to End-User Latency," in *NetGames*, 2012. 5

[17] W.-C. Feng, F. Chang, W.-C. Feng, and J. Walpole, "A traffic characterization of popular on-line games," *ToN*, vol. 13, no. 3, pp. 488–500, Jun. 2005. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1458759 5

[18] A. Petlund, P. Halvorsen, P. F. Hansen, T. Lindgren, R. Casais, and C. Griwodz, "Network traffic from anarchy online: analysis, statistics and applications: a server-side traffic trace," in *MMSys*, 2012. 5

[19] H. Liang et al., "Avatar mobility in user-created networked virtual worlds: measurements, analysis, and implications," *Multimedia Tools and Applications*, vol. 45, no. 1-3, 2009. 5, 6

[20] C. La and P. Michiardi, "Characterizing user mobility in Second Life," in *Online social networks*, 2008. 5, 6

[21] M. Varvello, S. Ferrari, E. Biersack, and C. Diot, "Exploring Second Life," *ToN*, vol. 19, no. 1, 2011. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5545464 5, 6, 15

[22] J. L. Miller and J. Crowcroft, "Avatar movement in World of Warcraft battlegrounds," in *NetGames*, 2009. 6

[23] W. Bohte and K. Maat, "Deriving and validating trip purposes and travel modes for multi-day GPS-based travel surveys," *Transportation Research Part C: Emerging Technologies*, vol. 17, no. 3, 2009. 6, 15, 19

[24] A. Clauset, C. Shalizi, and M. Newman, "Power-law distributions in empirical data," *SIAM Review*, vol. 51, no. 4, p. 661, 2009. 8

[25] D. Nurmi, J. Brevik, and R. Wolski, "Modeling machine availability in enterprise and wide-area distributed computing environments," in *Euro-Par*, 2005. 8

[26] D. Kondo et al., "The failure trace archive: Enabling comparative analysis of failures in diverse distributed systems," *CCGrid*, 2010. 8

[27] K. P. Burnham, "Multimodel inference: Understanding AIC and BIC," *Sociol. Meth. & Res.*, vol. 33, no. 2, 2004. 8

[28] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms," *Infocom*, 2006. 19

[29] K. Lee, S. Hong, S. J. Kim, I. Rhee, and S. Chong, "SLAW: A New Mobility Model for Human Walks," in *Infocom*, 2009. 19

6 Appendix

Flight lengths	Mean	Q1	Median	Q3	99%
Ironforge	24.6	6.6	13.8	29.6	157.6
Orgrimmar	22.3	6.4	13.2	26.5	153.9
Stormwind	19.6	5.5	10.5	21.9	161.7
Stormwind-2	25.9	6.8	14.2	28.6	202.2
Isis	22.2	5.7	11.0	19.9	217.0
Ross	23.4	5.5	10.9	23.8	186.7
Pharm	19.0	5.2	9.9	18.1	170.6
Freebies	29.1	7.4	14.1	29.0	230.2
GPS	214.6	9.0	23.8	84.8	3833.8
KAIST	61.3	12.8	24.1	49.7	573.1
NCSU	71.4	8.5	15.1	35.6	973.3

Table 2: Flight Lengths.

Pause durations	Mean	Q1	Median	Q3	99%
Ironforge	29.2	2.0	5.0	18.0	434.0
Orgrimmar	27.9	2.0	5.0	17.0	398.0
Stormwind	27.0	2.0	4.0	13.0	350.0
Stormwind-2	29.2	2.0	5.0	19.0	434.2
Isis	169.1	10.0	30.0	82.0	1571.2
Ross	230.4	10.0	30.0	92.0	3857.2
Pharm	441.0	11.0	40.0	133.0	8448.6
Freebies	123.3	10.0	30.0	70.0	1471.2
GPS	149.8	10.0	20.0	64.0	2449.0
KAIST	269.5	30.0	30.0	90.0	5520.0
NCSU	313.7	30.0	60.0	180.0	5349.0

Table 3: Pause Durations.

Popularity (visitor)	Mean	Q1	Median	Q3	99%	Max
Ironforge	5.0	0.0	0.0	2.0	78.1	505
Orgrimmar	2.4	0.0	0.0	0.0	50.0	486
Stormwind	1.2	0.0	0.0	0.0	28.0	322
Stormwind-2	2.4	0.0	0.0	0.0	45.0	668
Isis	12.2	0.0	0.0	0.0	428.5	1955
Ross	0.5	0.0	0.0	0.0	7.0	159
Pharm	4.1	0.0	0.0	0.0	93.3	1180
Freebies	35.9	4.0	13.5	42.5	250.5	1477
GPS	1.4	1.0	1.0	1.0	7.0	80

Table 4: Number of vistors per area.

Popularity (visitor)	Mean	Q1	Median	Q3	99%	Max
Ironforge-visits	13.4	0.0	0.0	3.0	252.2	2129
Orgrimmar-visits	7.4	0.0	0.0	0.0	115.7	4146
Stormwind-visits	3.3	0.0	0.0	0.0	67.0	960
Stormwind-2-visits	5.1	0.0	0.0	0.0	97.0	2430
Isis-visits	45.2	0.0	0.0	0.0	1537.4	8357
Ross-visits	1.1	0.0	0.0	0.0	15.0	284
Pharm-visits	10.9	0.0	0.0	0.0	239.8	2275
Freebies-visits	81.5	8.0	26.0	89.0	743.5	3235
GPS-visits	3.8	1.0	2.0	3.0	34.0	884

Table 5: Number of visits per area.

GPS visits	Mean	Q1	Median	Q3	99%	Max
Ironforge	1.0	0.3	0.7	1.4	9.4	0.4900
Orgrimmar	1.1	0.2	0.6	1.5	10.7	0.5684
Stormwind	0.4	0.1	0.2	0.5	5.7	0.5722
Stormwind-2	0.4	0.1	0.2	0.5	5.0	0.5088
Isis	2.5	0.8	1.7	3.4	19.0	0.4651
Ross	1.2	0.7	0.7	1.4	6.9	0.3856
Pharm	1.8	1.1	1.4	2.2	11.9	0.3561
Freebies	1.3	0.5	0.9	1.7	11.3	0.4773
GPS	0.1	0.1	0.1	0.1	0.4	0.2782

Table 6: Percentage of visited areas.

Personal preference Gini	Mean	Q1	Median	Q3	Max	Gini
Ironforge	0.2880	0.2023	0.2781	0.3582	0.7262	0.2278
Orgrimmar	0.3074	0.2083	0.2978	0.3986	0.7614	0.2625
Stormwind	0.2749	0.1714	0.2697	0.3619	0.6780	0.2687
Stormwind-2	0.2542	0.1667	0.2448	0.3176	0.7084	0.2346
Isis	0.3022	0.2245	0.3000	0.3839	0.6935	0.2125
Ross	0.2427	0.1523	0.2333	0.3007	0.6335	0.2373
Pharm	0.2730	0.2000	0.2727	0.3407	0.6131	0.2125
Freebies	0.2140	0.1500	0.2015	0.2648	0.6825	0.2345
GPS	0.3976	0.3122	0.3780	0.4671	0.8815	0.1632

Table 7: Personal preference weight Gini.

GPS visits times	Mean	Q1	Median	Q3	99%	Max
gps10-times	3.8	1.0	2.0	3.0	34.0	884
gps20-times	6.1	1.0	2.0	4.0	72.0	2354
gps50-times	11.8	1.0	2.0	6.0	180.0	4789

Table 8: GPS visit times.