Characterization of Human Mobility in Networked Virtual Environments

Siqi Shen, Niels Brouwers, Alexandru Iosup, Dick Epema
Delft University of Technology, Delft, The Netherlands
{S.Shen, N.Brouwers, A.Iosup, D.H.J.Epema}@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.

Categories and Subject Descriptors

H.5 [**Information Systems Applications**]: Multimedia Information Systems

General Terms

Measurement

Keywords

Network virtual environment, mobility characterization

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. Complementing much previous research in the design and tuning of NVE systems, and in particular in collecting, characterizing, and modeling NVE workloads [4, 9, 14], 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 [8]. We also conduct a comprehensive, comparative characterization of the mobility of citizens in WoW and other, conceptually different NVE. 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 [6] and peer-to-peer overlays [10] 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 NVEs such as [11], the load of various servers is strongly correlated with player mobility, due to player interaction [1, 19], cell visitation [6], etc. As has been shown in preliminary work on this topic [16], 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 [17, 22]. 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 [13]. 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 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), and share the dataset via the Game Trace Archive [8].

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 [22], 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. [12] 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** [17]: human usually travel short distances and occasionally travel long distance.

- **(C2) Long-tail distribution of pause durations** [12].

- **(C3) Skewed popularity of areas** [21, 22]; for example, certain areas of cities are very popular, while others are rarely visited.

- **(C4) Invisible boundary of human movement** [17]: most of the time, people only travel between and around a few preferred locations.

- **(C5) Different personal preferences for areas** [3].

2.2 Related work

Much of the prior work [4, 5, 7, 14, 20] has focused on network measurement, online population, and session behavior. Since the late-2000s, several studies [9, 15, 24] collect and analyze mobility traces of NVEs. We compare our work with these, in the following.

Closest to our work, Liang et al. [9] 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 [21] 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. [24] 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 [24], Miller and Crowcroft [18] find that, in their observation scenario of WoW battleground, most movement is individual rather than group-based. La and Michiardi [15] 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, and the data collection processes.

3.1 Dataset Description

We have collected a very large and detailed dataset from a popular virtual world, World of Warcraft, and used three other datasets that were collected by others from virtual and real worlds. The characteristics of the datasets are summarized in Table 1.

Dataset	World	Citizens	Space	Time	Granularity
WoW (ours)	Virt.	31,290	4 cites	2w	1s
SL [9]	Virt.	26,714	4 zones	days	10s
GPS [2]	Real	1,366	3 cities	1w	6s
GPS-2 [12]	Real	52	2 campuses	days	30s

Table 1: Dataset overview.

Our dataset, `WoW`, 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). The `SL` dataset is collected and released by Liang et al. [9] from Second Life, and includes about 25,000 citizens from 4 zones: `Isis`, `Ross`, `Pharm` and `Freebies`. The `GPS` dataset is collected by Bohte and Maat [2] by distributing GPS devices to over 1,300 persons live in 3 cities for transportation research. The `GPS-2` dataset is collected by Rhee et al. [12], and contains traces of about 50 persons from two campuses: `KAIST` and `NCSU`.

3.2 Data Collection

`WoW` adopts a sharding architecture with multiple independent realms with same starting scenario. Each realm may have different types of interaction styles: normal, role-playing and player versus player (PVP). We collected 17,583 users' trace from 3 capital cities (`Ironforge`, `Orgrimmar`, and `Stormwind`) of the popular Silvermoon server (normal realm in Europe) and 13,707 users from `Stormwind` (we call it `Stormwind-2` from now on) of the popular Argent-Dawn server (role-playing realm in Europe).

Each virtual citizen of WoW can observe the presence and activity of any other virtual citizen within a radius of about 100 in-game meters. Unlike the real world, the observation range in WoW is not affected by interposing objects such as buildings or other citizens. To collect the `WoW` dataset, we have developed a tracing script and used it to observe selected cities. Our script logs-in several regular WoW clients and coordinates them to observe a large part of a city. To observe mobility in a city, our client deploys virtual citizens such that their observation areas cumulatively cover the surface of the 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 weekends and about 20 week-day evenings during April and September 2011, resulting in mobility traces for over 30,000 anonymous virtual citizens. As described in our technical report [23], we have performed typical data sanitation on our datasets, including of non-pedestrian movement, overall removing less than 1% of the raw data.

4. CHARACTERIZATION

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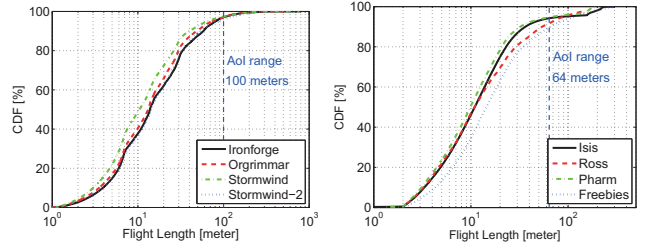
Figure 1: Flight length distribution of (left) the `WoW` dataset, and (right) the `SL` dataset (horizontal logscale).

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. The complete results, including per-trace detailed results and the detailed results of validation (fitting), are available as a technical report [23].

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

Because virtual worlds may distort the sizes of buildings and the speed of avatars, in comparison with real-world environments, we are interested in study the general trends and distributions of mobility characteristics, beside basic statistics values. 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 the Exponential, the Weibull, the LogNormal, the Gamma, the Normal, and the General Pareto distributions. The fitting is performed using *maximum likelihood estimation* (MLE), which determines for a distribution the parameters that lead to the best fit with given empirical data. Then, we use a method for assessing the *goodness-of-fit* (GoF) by using a combination of Kolmogorov-Smirnov (KS) and the Anderson-Darling (AD) GoF tests. We determine a distribution is the best fit for the empirical data if it passes the two GoF tests and has the smallest Akaike information criterion with correction(AICc) value. We call the probability distribution of data is "long-tail" if the tail of the probability distribution is longer than the fitted exponential distribution of data.

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 the flights are shorter
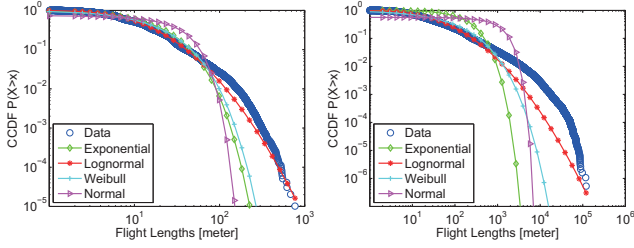
Figure 2: Distribution fitting of (left) `WoW` dataset, and (right) the `GPS` dataset (all axes logscale).
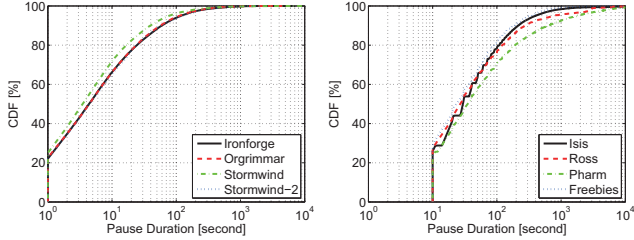

Figure 3: Pause duration distribution of (left) the `WoW` dataset, and (right) the `SL` dataset (horizontal logscale).

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.

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 2 depicts the results of fitting for `Stormwind`, `WoW` and `GPS`. 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.

4.3 Pause Duration (C2)

Figure 3 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
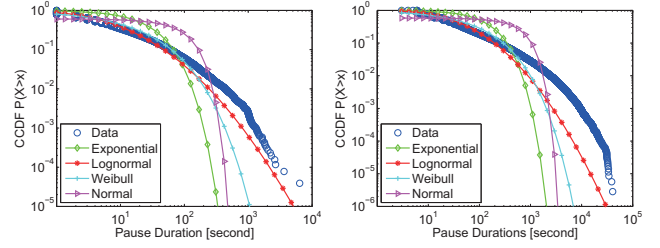

Figure 4: Distribution fitting of (left) `WoW` dataset, and (right) the `GPS` dataset. (all axes logscale.)

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 players is more frequent.

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 4 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.

Number of area visits: Figure 5 (left) shows the number of area visits of a 1 day trace from `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
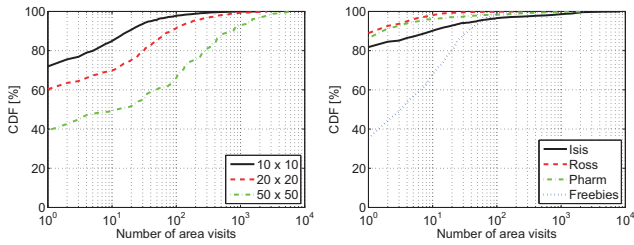
Figure 5: Number of area visits of (left) the WoW dataset, and (right) the SL dataset (horizontal logscale).
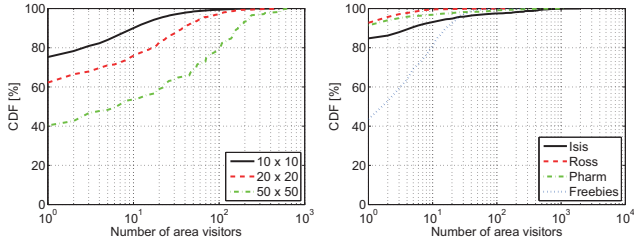

Figure 6: Number of area visitors of (left) the WoW dataset, and (right) the SL dataset.

10 while the maximal value is about 1,921. Figure 5 (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. 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 6 shows the number of area visitors of a 1 day trace from Ironforge and SL. The number of area visitors is smaller than the number of area visits, but it is long-tail too. Figure 6 (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 6 (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.

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 7 shows the number of distinct areas, normalized by the to-
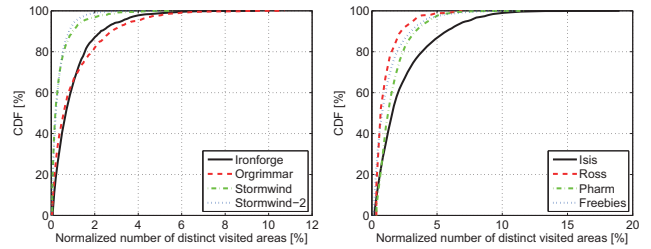
Figure 7: Normalized number of distinct visited areas of (left) the WoW dataset, and (right) the SL dataset.

tal 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 7 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.

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 [24]; and real-world citizens have strong preferences for different areas [2]. 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 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 8 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 $\alpha = 2.5, 3, 4$
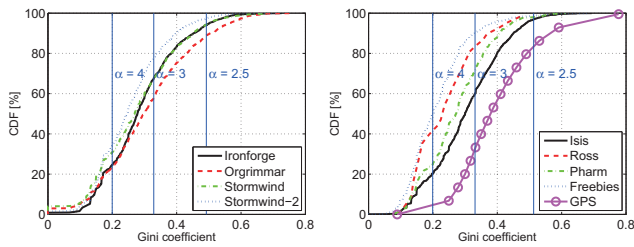
Figure 8: Gini coefficient of personal preference weight (left) the `WoW` dataset, and (right) the `SL` dataset and `GPS` dataset.

in Figure 8. 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. 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 mobility is important for the design and tuning of networked virtual environments (NVEs), but has been hampered until now by the lack of datasets and of comparative studies. Our main contribution in this work is the collection and comparative characterization of mobility traces in NVEs; furthermore, we have extended the comparison to also include networked real-world environment (NREs). We have collected detailed position information of virtual world mobility traces from World of Warcraft (WoW). Using our traces, the public Second Life (SL) traces (SL is of different genre and content-generation method than WoW), and two NRE traces collected by others and kindly shared with us, we have conducted a comprehensive, comparative study of mobility characteristics. Our study has shown evidence that long-tail distributions characterize well flight lengths, pause durations, and area popularity; that the invisible boundary of human movement also appears for NVEs, and that avatars do have preferences to different areas. We have also indicated several differences between NVE and NRE mobility characteristics: the flight lengths distributions have longer tail for NREs, the personal preference in area-visitation is more pronounced for NREs. We are developing system scaling techniques based on these observations.

6. ACKNOWLEDGEMENT

7. REFERENCES

[1] D. T. Ahmed and S. Shirmohammadi. Improving online gaming experience using location awareness and interaction details. *Multimedia Tools Appl.*, 61(1), 2012.

[2] 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*, 17(3), 2009.

[3] C. Song et al. Modelling the scaling properties of human mobility. *Nature Phys.*, 6(10), 2010.

[4] C. Chambers, W.-C. Feng, S. Sahu, D. Saha, and D. Brandt. Characterizing online games. *ToN*, 18(3), 2010.

[5] 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.

[6] Y. Deng and R. W. H. Lau. On delay adjustment for dynamic load balancing in distributed virtual environments. *TVCG*, 18(4), 2012.

[7] W.-C. Feng, F. Chang, W.-C. Feng, and J. Walpole. A traffic characterization of popular on-line games. *ToN*, 13(3):488–500, June 2005.

[8] Y. Guo and A. Iosup. The game trace archive. In *NetGames*, 2012.

[9] H. Liang et al. Avatar mobility in user-created networked virtual worlds: measurements, analysis, and implications. *Multimedia Tools and Applications*, 45(1-3), 2009.

[10] S.-Y. Hu, J.-F. Chen, and T.-H. Chen. VON: a scalable peer-to-peer network for virtual environments. *IEEE Netw*, 20(4), 2006.

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

[12] I. Rhee et al. On the Levy-Walk Nature of Human Mobility. *Infocom*, 2008.

[13] A. Keränen, T. Kärkkäinen, and J. Ott. Simulating Mobility and DTNs with the ONE. *JCM*, 5(2), 2010.

[14] J. Kinicki and M. Claypool. Traffic analysis of avatars in Second Life. In *NOSSDAV*, pages 69–74, 2008.

[15] C. La and P. Michiardi. Characterizing user mobility in Second Life. In *Online social networks*, 2008.

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

[17] M. González et al. Understanding individual human mobility patterns. *Nature*, 453(7196), 2008.

[18] J. L. Miller and J. Crowcroft. Avatar movement in World of Warcraft battlegrounds. In *NetGames*, 2009.

[19] V. Nae, A. Iosup, and R. Prodan. Dynamic resource provisioning in massively multiplayer online games. *TPDS*, 22(3), 2010.

[20] 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.

[21] D. Pittman and C. GauthierDickey. Characterizing virtual populations in massively multiplayer online role-playing games. *AMM*, 2010.

[22] R. Becker et al. Human mobility characterization from cellular network data. *Commun. ACM*, 56(1), 2013.

[23] S. Shen, N. Brouwers, and A. Iosup. Human mobility in virtual and real worlds: Characterization, modeling, and implications. Technical Report PDS-2011-007, TU Delft.

[24] M. Varvello, S. Ferrari, E. Biersack, and C. Diot. Exploring Second Life. *ToN*, 19(1), 2011.