Second Life: a Social Network of Humans and Bots

Matteo Varvello[*]
Alcatel-Lucent, Holmdel, USA
matteo.varvello@alcatel-lucent.com

Geoffrey M. Voelker
University of California, San Diego, USA
voelker@cs.ucsd.edu

ABSTRACT

Second Life (SL) is a virtual world where people interact and socialize through virtual *avatars*. Avatars behave similarly to their human counterparts in real life and naturally define a *social network*. However, not only human-controlled avatars participate in the social network. Automated avatars called *bots* are common, difficult to identify and, when malicious, can severely detract from the user experience of SL. In this paper we study the SL social network and the role of bots within it. Using traces of avatars in a popular SL region, we analyze the social graph formed by avatar interactions. We find that it resembles natural networks more than other online social networks, and that bots have a fundamental impact on the SL social network. Finally, we propose a bot detection strategy based on the importance of the social connections of avatars in the social graph.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Measurement techniques; H.5.1 [**Multimedia Information Systems**]: Artificial, augmented, and virtual realities

General Terms

Measurement, Design

Keywords

Second Life, Social Networking, Bot Detection

1. INTRODUCTION

Second Life (SL) is a *virtual world* accessible by multiple participants through the Internet [16]. With more than 15 million residents, SL is one of the most popular virtual worlds. The virtual land is composed of *regions* that users access using human-controlled characters called *avatars*. Avatars live a parallel life in SL: they explore the virtual world, meet other users, communicate, play, trade, etc.

---

[*]Work done while a visiting scholar at UCSD

Researchers have shown that avatars tend to interact similarly to human beings in real life [12][17]. They meet, spend time together and make friends. This behavior suggests that avatars construct an *online social network*, an Internet-based network that represents the social relationships existing among human beings. Popular examples of online social networks are the ones created on social sites (e.g., Facebook [9] and Cyworld [7]), content sites (e.g., Flickr [10]), game sites (e.g., World of Warcraft [22]), etc.

Avatars can also be controlled via automated scripts, or *bots*. Bot interactions can impact the virtual world in a variety of ways. Automated avatars can be useful for welcoming human avatars to a region, and measurement bots can benignly record world state. Some bots can be offensive, however: they can spy on user behavior or, as with the CopyBot, even clone avatar identities [20].

Since offensive bots can significantly detract from the user experience, the SL provider disconnects avatars that have not moved after 15 minutes. Not surprisingly, bots perform repetitive actions, such as short walks, to circumvent this heuristic [17]. Furthermore, users can also report directly to the SL provider when they suspect the presence of bots or misbehaving avatars. The SL provider decides whether or not to expel the reported avatar, but this procedure is manual and can require several days. Detecting bots therefore remains a compelling problem for virtual worlds like SL.

In this paper we study the social network formed by avatars interacting in SL, and compare it with other social networks. We then examine the impact bots have on the SL social graph, and, using avatar relationships defined by the social graph, propose an automated approach for identifying suspicious avatars as bots. We use measurements of avatar behavior as a basis for constructing the SL social network. We enhance a *crawler* developed in previous work [17] to collect traces of avatar behavior in SL for longer time periods. We crawl a very popular SL region for 10 days, tracing the behavior of 3,291 unique avatars.

Initially, we find that the complete SL social network is a small-world network. Further, it is much more similar to natural networks [18][19] than popular online social networks [13][21]. The creation of social relationships in SL requires an *active* interaction among avatars. Conversely, a social link between two users in online social networks frequently represents only the acceptance of a friendship request and does not require ongoing interaction.

Many features point to the presence of bots in the social network and that these bots have a fundamental impact on its structure. Bots interact with a very large number of avatars, but they construct very fragile relationships. Indeed we find that, after removing suspected bots from the SL social network, it is no longer a small-world network. As a result, we conclude that analyses of user behavior in SL should distinguish between human and bot avatars.

Finally, inspired by the insight into bots provided by the social

graph analysis, we propose a bot detection strategy that identifies avatars as likely bots based on the importance of their social connections in the social graph. Applied to the trace, the strategy easily detects the bot used by our own crawler and overall estimates that 4–7% of the avatars are suspected bots. This estimation requires further confirmation, however, and defines our future work.

2. RELATED WORK

Despite its early success, only recently have measurement projects characterized user and system dynamics in SL. La and Michiardi [12] collect traces of avatar behavior in three SL regions for a few hours. They use these traces to compare avatar to human behaviors. Interestingly, they show that the distribution of avatar contact times measured in SL is very similar to the distribution of human contact times measured in real-world experiments [4][14]. In previous work, Varvello et al. [17] collect public data across the entire SL virtual world. They show that only few regions are quite popular, and nearly 30% of the regions attract no visitors. Although we rely on measurement methodologies similar to those used in [12][17], our analysis of the SL social network extends and complements these studies.

Online social networks have received comparatively more attention. Mislove et al. [13] study the structural properties of current popular online social networks, e.g., Flickr [10], Orkut [15], etc. They show that online social networks have structural properties very different from natural networks (e.g., they exhibit a high level of local clustering). Subsequently, Chun et al. [5] analyze Cyworld [7], a large South Korean social networking service. They compare the *friend relationships network* — the network defined by friendships among Cyworld users — with the *activity network* — the network defined by user activities. They show that the two networks have a similar structure, suggesting that interactions between users in a social network tend to follow the declaration of friend relationships. Recently, Wilson et al. [21] conduct a similar study to the one in [5]. They focus on Facebook [9], currently the most popular social networking service. In contrast with the study conducted in [5], they show that the *interaction network* derived from user activities in Facebook exhibits significantly lower levels of the small-world properties shown in the Facebook social network. We adopt methodology from these efforts to analyze the SL social network, which in contrast arises more naturally as a consequence of human-based avatar interactions.

3. DATA AND METHODOLOGY

In this section we describe our methodology for constructing a social graph among SL avatars based upon a trace of SL avatar interactions.

3.1 Contact and Social Graph

We define contacts in SL based upon proximity in the virtual world. An intuitive result in human communication is that "closer together" means "more likely to converse" [8]. Recent work has shown that avatars share behavior similar to human beings [12][17], e.g., they gather in popular places to meet friends. Hence, we assume that the distance between avatars plays an important role in avatar communication.

We assume that two avatars are interacting, i.e., there exists a *contact* between them, when their Euclidean distance is less than an *interaction range* R. We recognize that this assumption may identify contacts where avatars are not directly interacting (e.g., avatar "strangers" passing each other in a street), but it still intuitively captures avatar contacts and the possibility of interaction. We de-

Region	Japan Resort
Trace Length	10 days
Crawling Frequency	90 secs
Unique Avatars	3,291

Table 1: Traces Summary

fine *contact time* as the time interval in which two avatars have an Euclidean distance smaller than R. Finally, we define *session time* as the continuous online time of an avatar.

We now introduce the *contact graph* similar to that previously described in [14]. The contact graph $G_t = (V_t, E_t)$ is the collection of avatars connected to a SL region (V_t) as well as the edges connecting the avatars (E_t) at a time t. G_t is an *unweighted* graph, i.e., edges are not assigned any a-priori weight. An edge $< i, j >_t$ connecting nodes i and j in $V(G_t)$ equals 0 when the Euclidean distance d between the coordinates of avatars i and j at time t is greater than R, whereas it equals 1 when $d < R$. G_t is also an *undirected* graph, i.e., $< i, j >_t = < j, i >_t$ for any i and j.

Similar to the activity network studied in [5] and the interaction network studied in [21], we now introduce the *social graph* as the network of friendships that users construct with their behaviors. Formally, the social graph $G = (V, E)$ is the collection of avatars visiting SL (V) as well as the edges connecting these avatars (E). G is a *weighted* and *directed* graph, i.e., each edge $< i, j >$ connecting two avatars i and j in $V(G)$ has two associated weights $w_{i,j}$ and $w_{j,i}$. We compute $w_{i,j}$ and $w_{j,i}$ as the ratio of the sum of all contact times between avatars i and j and the sum of the session times for i and j, respectively. Intuitively, $w_{i,j}$ is the fraction of "virtual" time avatar i spends being close to j. Therefore, $w_{i,j}$ and $w_{j,i}$ captures the "importance" of the social connection between avatars i and j (acquaintances, friends or relatives).

3.2 Data Collection and Limitations

In order to construct an instance of the SL social graph, we collect traces of avatar behavior with the crawler we developed in previous work [17]. The crawler consists of a modified SL client that exploits standard avatar capabilities to monitor the virtual world. For the specific task of monitoring avatar behaviors on a region, the crawler extracts from a map of the region the high level coordinates of the avatars, then it reaches each of these coordinates and scans the surrounding area. All these operations require some time to be accomplished, implying a "crawling frequency".

In that work, the crawler only monitored avatars for 3 days due to IP blacklisting from the SL provider. Such a short time period, however, is not sufficient to capture social relationships among avatars. To overcome this limitation, we exploit DHCP reassignment in a private ADSL connection provided by a French ISP to enable the crawler to use a range of IP addresses to connect to SL. We changed the crawler to toggle its Internet connection by interacting with the home gateway router. This procedure triggers the ISP to assign a new public IP address to the connection, circumventing the IP blacklisting mechanism adopted by the SL provider.

Ideally, we would want to capture the interactions of avatars across all of the 18,000 regions comprising SL. However, monitoring SL at this scale is not feasible using our current methodology (e.g., it requires public IP addresses linear with the number of regions). Instead, as an initial study we focus the crawler on a single, highly-popular region [17]. The social graph we analyze therefore corresponds to the portion of the social graph originated by the avatars just interacting in this region. Given that avatars ex-

Traces	α (in)	α (out)	$\frac{C}{C_{er}}$	$\frac{L}{L_{er}}$
Flickr [13]	1.74	1.78	47,200	-
Orkut[13]	1.5	1.5	7,240	-
Facebook [21]	1.5	1.5	21,866	-
Film Actors [19][3]	3	3	2,925	1.22
Power grid [19][2]	-	-	16	1.50
C.elegans [19][11]	2.2	2.2	5.6	1.17
Second Life	2.2	2.2	70	1.1

Table 2: Comparison of the SL social network with several natural networks and online social networks.

hibit similar behaviors in different SL regions [12][17], studying one popular region can still shed some lights on the entire SL.

We use the crawler on the popular "Japan Resort" region for 10 days between July 22 and August 2, 2008, and we monitor 3,291 unique avatars. According to available SL statistics [16], these avatars account for about 1% of the unique avatars logged in during this time. Table 1 summarizes the main characteristics of the traces collected.

The traces contain a large gap of about 20 hours due to a major region outage, and other minor gaps of a few minutes likely due to server updates. These holes do not represent a loss of information since no avatar activity was possible in the region.

To choose a value for the interaction range R, we compute the minimum distance observed between any pair of avatars at each crawling snapshot. We found that 99% of the minimum distance values are less than 5 meters. Therefore, we set R to 5 meters.

4. SOCIAL GRAPH CHARACTERISTICS

In this section we analyze the social graph characteristics of Second Life avatars, compare them to measured characteristics of other social networks, and show that bots have a fundamental impact on the SL social graph.

4.1 A "real" online social network

We first characterize the features of the SL social graph defined by the traces of avatar contacts in the "Japan Resort" region. When discussing these features, we also compare and contrast them with the characteristics of both online social networks and natural networks. Table 2 summarizes the SL social graph features and how they compare to the characteristics of other social networks as determined by previous studies. We focus on the analysis of the "complete" graph G, i.e., two nodes i and j in G are connected if $w_{i,j}$ and consequently $w_{j,i}$ are non-zero.

4.1.1 Degree

The *degree* of a node $v \in V(G)$ is the number of edges incident to v, and represents the number of unique avatars encountered by avatar v during the crawl period. We start with a quantitative analysis of the node degree distribution computed over G as defined at the end of the crawl period. Although not shown due to space limitations, 90% of the values in the node degree distribution are between 1–30. This range implies that avatars are either very isolated, e.g., they explore the region without meeting any other avatar, or they interact with a restricted set of avatars, e.g., they meet their friends. The highest percentiles of the node degree distribution, however, approach 300. Intuitively, avatars with such high node degrees are either extremely social or controlled by automated scripts. A further analysis of the traces shows that these avatars are the most *persistent* as well, i.e., they are consecutively connected for about 90% of the crawl duration. This persistence

suggests that their node degree represents the number of avatars they *meet*, rather than the number of avatars they *interact* with. We attribute this (un)social behavior to bots, and further analysis below increasingly corroborates this conjecture.

Given that social networks are often characterized by a power law distribution of their node degrees [13][18], we examine how well a power law fits the node degree distribution of G. As in other studies, we use the maximum likelihood method to calculate the best power law coefficient α as well as its lower bound x_{min}, and the Kolmogorow-Smirnow goodness-of-fit to evaluate the fit quality (D) [6].

Figure 1(a) plots the estimation of α, x_{min} and D for the node degree distribution computed on G every 15 minutes for the duration of the trace. During the first 24 hours, the estimations of α, x_{min} and D change rapidly. This phenomenon is due to the fact that we incrementally construct the social graph using the interactions that occur among avatars over time. Therefore, the initial hours simply represent a transient phase in the definition of the social graph. Subsequently, α slowly decreases to a value of 2.2 and D varies around 0.05, indicating that the power law well approximates the node degree distribution. The estimation of x_{min} oscillates around a degree value of 20–25, i.e., the power law fit is verified for about 25% of the avatars.[1] For comparison, observations of most natural networks indicate a value of α between 2 and 3 [2][3][11], whereas a value of α smaller than 2 is observed in most online social networks [21][13].

Figure 1(a) shows two high peaks in the estimation of both α and x_{min} at $t = 96$ hrs and $t = 108$ hrs, respectively. These times correspond with two short interruptions of the region service likely due to a server update. At those times the tail of the node degree distribution becomes more skewed, causing a shift of the power law fit, i.e., higher x_{min} and α. This shift indicates that the avatars responsible for the very high values in the node degree distribution are the most responsive in re-connecting to the region when the service becomes available. The same phenomenon is even more visible at $t = 192$ hrs when the SL service returns after an outage of 20 hrs (Section 3.2). In this case, the node degree distribution is significantly impacted and α and x_{min} return to the values measured before the outage only after 48 hrs. We conjecture that while real users delay before returning to a region in the presence of server failures, these high-degree nodes correspond to bots that reconnect to the region as soon as the service is available using automated probing.

4.1.2 Clustering Coefficient

The *clustering coefficient* is often used to characterize the extent to which nodes in social graphs form a small-world network. The clustering coefficient for a node $v \in V(G)$ is the ratio of the number of edges between the nodes within v's neighborhood and the total number of edges that could possibly exist between them [18]. Figure 1(b) shows a scatter plot of a node's degree and its clustering coefficient for all nodes in the SL social graph at steady state (i.e., the graph constructed at time $t = 256$ hrs). We observe no clear relationship between a node's degree and its clustering coefficient for nodes with a degree less than 30 (about 90% of the nodes in G). Conversely, the clustering coefficient of nodes with a degree greater than 30 quickly decreases since the node degree value grows so high. Looking at the neighborhoods of these nodes in more detail reveals that, for 99% of them, the median value of the weight of their edges is smaller than 0.1, i.e., they spend less than 1% of their virtual time in close proximity to the avatars they con-

---

[1]Note that [5][13] verify the power law fit for about 10% of users.

(a) Power law analysis of the node degree over time for α, x_{min}, D.

(b) Node degree (log scale) vs. clustering coefficient; $t = 256$ hrs.

(c) Percentiles of the betweenness centrality distribution over time.

Figure 1: Second Life social graph analysis.

tact. This result suggests that avatars associated with nodes with very high degree establish very fragile social relationships, again suggesting that these avatars are likely bots.

4.1.3 Betweenness Centrality

The *betweenness centrality* of a node reflects its relative importance in a graph, e.g., the popularity of a person within a social network. Formally, the betweenness centrality for a node v in $V(G)$ is the number of times node v occurs in a shortest path between any two other nodes in $V(G)$ divided by all existing shortest paths in G [18]; the *shortest path length* for a node v in $V(G)$ is the minimum number of edges connecting v to all other nodes in $V(G)$.

Figure 1(c) shows several percentiles of the distribution of the betweenness centrality over time (the gaps correspond to SL service outages). Figure 1(c) shows that nearly 90% of the nodes appear in less than 1% of all shortest paths. Interestingly, this result is similar to what Mtibaa et al. [14] observed in the social network formed via a mobile social application. However, in the SL social network we also observe that about 5% of the nodes are very central and intersect up to 10% of all the existing shortest paths. These very central nodes, though, also have high degrees and again most likely correspond to bots. This results indicates that bots may have a significant impact on the social network.

4.2 Impact of bots on the social network

In this section, we explore the impact of suspected bots on the structure of the SL social network. We compare the structure of the complete social network with embedded social networks formed by filtering weak social connections among avatars. Since suspected bots introduce many of these weak social connections, removing them enables us to compare the SL social network with and without bots. More formally, we analyze the set of graphs G' formed by filtering from G all edges with a weight smaller than a threshold W. Then we compare G' with the Erdos-Renyi graph, i.e., a purely random graph commonly used to identify *small-world* graph properties [18]. In our analysis, we always consider the largest connected component of the graph [18].

The high-level result of the following analysis is that the complete SL social network can be classified as a small world-network. However, after removing the edges suspected to be associated with bots, the resulting network is no longer a small-world network.

4.2.1 Clustering Coefficient

We denote by C the *average clustering coefficient* computed among all nodes in $V(G)$. We compare C with the clustering coefficient C_{er} computed in a Erdos-Renyi graph with the same number of nodes and edges as G [18]. We recall that $\frac{C}{C_{er}} >> 1$ indicates a possible small-world graph.

Figure 2(a) shows the evolution over time of the ratio $\frac{C}{C_{er}}$ as a function of the threshold W on edge weights $w_{i,j}$ between any two nodes i and j. For a graph formed using a weight threshold W, an edge only appears in the graph if $w_{i,j} > W$. We start by focusing on the analysis of the complete graph G, i.e., the curve obtained with $W = 0\%$ (all edges included). After the initial transient phase, the ratio $\frac{C}{C_{er}}$ slowly grows until reaching 70. Interestingly, the ratio $\frac{C}{C_{er}}$ measured for SL is several orders of magnitude smaller than the values measured for several online social networks (e.g., 21,866 in Facebook [21] and 47,200 in Flickr [13]), whereas it is much closer to the ratio measured in many natural networks (e.g., 16 in the electrical power grid of the western United States, and 5.6 in the neural network of the nematode worm C. elegans [1][19]).

We now increasingly remove edges corresponding to weak social connections. With $W = 1\%$, the ratio $\frac{C}{C_{er}}$ is reduced by $\frac{1}{3}$ when removing from G all edges with weight smaller than 1%, i.e., the presence of highly clustered portions of the graph are reduced significantly. This change indicates that weak social connections (mostly originated by bots) play an important role in the definition of the social graph structure. Interestingly, when increasing the value of W from 1% to 50% the curves approach the curve obtained with $W = 0\%$. This result indicates that the subgraph G' composed only using edges with a large edge weight (i.e., strong social connections mostly associated to human-controlled avatars) is again a possible small-world graph.

4.2.2 Shortest Path

We denote by L the *characteristic path length* for a graph G, i.e., the median of the means of the shortest path lengths connecting each node $v \in V(G)$ to all other nodes [18]. We now compare the characteristic path length L in G with the characteristic path length measured in a Erdos-Renyi graph (named L_{er}) with the same number of nodes and edges as G [18]. Again, analyses of social networks compute the ratio $\frac{L}{L_{er}}$ since values close to one indicate small-world networks.

Figure 2(b) shows the evolution over time of $\frac{L}{L_{er}}$ for a range of

(a) $\frac{C}{C_{er}}$ evolution over time.

(b) $\frac{L}{L_{er}}$ evolution over time.

(c) CDF of the percentage of human-to-bot (H2B) interactions.

Figure 2: Sensitivity of small-world graph metrics and human-to-bot interactions on social connection threshold W.

values of W. The curve for the complete graph with $W=0$ shows that $\frac{L}{L_{er}}$ is nearly one. Together with the high clustering coefficient of G measured above, these results indicate that the complete SL social network is a small-world network [18]. However, for $W=1\%$ the ratio $\frac{L}{L_{er}}$ assumes a small value. Most edges with a small weight are associated with nodes with a very high degree — bots — that are central in the social network. When we remove these edges, the largest connected component of the graph is reduced, resulting in smaller characteristic path lengths. However, increasing the value of W to 10% — effectively filtering out most of the edges due to bots — the ratio $\frac{L}{L_{er}}$ increases again. In other words, although we have further reduced the number of edges, the size of the largest connected component remains stable. This result indicates that human-controlled avatars tend to share at least one edge with a large weight.

5. SOCIAL-BASED BOT DETECTION

Second Life currently relies upon a simple idleness heuristic and user reports to disconnect suspected bots from the system. Based on our analysis of the SL social graph, we found a variety of features of the graph that correlate with bot behavior. As a result, we propose that SL can further rely upon characteristics of the SL social network to differentiate between human avatars and bots.

A social-based bot detection strategy could use the social graph features directly. A provider may identify as bots all avatars with a node degree larger than 100 and a clustering coefficient smaller than 0.2 (Figure 1(b)) or a value of centrality larger than 5% (Figure 1(c)). However, this approach requires tuning several independent thresholds, computing the social metrics for all avatars, and a centralized authority.

Instead, we propose a bot detection strategy that classifies avatar interactions as *human-to-human* and *human-to-bot* based on the importance of their social connections in the social graph. Human avatars will have strong social connections as they interact in a community, while bot avatars will have weak connections. Using social connections naturally combines the specific graph features into a single abstraction without having to set thresholds for each feature independently. For simplicity, we consider the case of a single server hosting a region. Moreover, we assume that the server computes and maintains over time the social graph.

5.1 Description

We say that whenever two avatars A and B meet (i.e., they enter

their respective interaction range R) the server "classifies" their interaction. To do so, it searches for a path connecting nodes A and B in the social graph G entirely composed of edges $< i, j >$ that have both weights $w_{i,j}$ and $w_{j,i}$ greater than a threshold W, with $0 < W \leq 1$. We require both weights for all edges $< i, j >$ in the path A-B to be larger than W to capture the effect that both avatar i and j "agree" in the importance of their social connection. If no path A-B exists, e.g., B is a newcomer not yet known to the community, the server delays making a decision since it does not have enough information to classify the interaction. If at least one path exists and respects the weight condition, it labels the interaction A-B as *human-to-human*. Otherwise, it labels the interaction A-B as *human-to-bot*. At this time, the server does not know whether A or B is the bot, it simply suspects that one of the two avatars involved in the interaction is computer-controlled.

The bot detection strategy as described so far generates a set of classifications for the avatar interactions. Based on these classifications, we then suspect avatars to be bots if most of their interactions are labeled as human-to-bot. Having detected suspicious bot avatars, the SL provider can then decide what kind of policy to apply to the bot avatars. For example, it can ignore useful and benign bots while banning offensive bots (e.g., the CopyBot). To further reduce false positives, SL can also use CAPTCHAs on suspected bots or the "bot reports" collected from users for additional evidence about bots considered offensive by the users community.

The server can lazily compute the volume of human-to-bot interactions per avatar, e.g., during time periods with low load. Moreover, we anticipate classifying avatar interactions in a distributed fashion by having avatars perform local computations and collecting them at servers. Correlating classifications among neighbors can guard against falsification.

5.2 How many bots are out there?

We now use the bot detection strategy to estimate the total number of bots observed in the 10 day traces. We implement a Matlab simulation that takes as input the avatar traces (Section 3.2), computes and maintains the social network (Section 3.1), and labels the avatar interactions as *human-to-human* or *human-to-bot* (Section 5.1). In the simulations we vary W, the minimum level of reciprocal acquaintances required among two avatars to trust their social connection.

In Figure 2(c), we compute for each avatar the percentage of its interactions with other avatars labeled as human-to-bot (H2B) and

plot the CDF across all avatars as a function of W. When the percentage of interactions labeled as H2B is relatively low ($\leq 20\%$), we find that the results are sensitive to the minimum weight W required to trust a social connection. For example, the percentage of avatars that have less than one percent of their interactions labeled as H2B is equal to 86% when $W=1\%$, while it reduces to 62% for $W=50\%$. This trend occurs because increasing the threshold W increases the probability of generating false positives (labeling some human-to-human interactions as H2B), and decreasing W increases the probability of generating false negatives (labeling some bot interactions as human-to-human).

Figure 2(c) shows, however, that the curves nearly overlap in the right portion of the graph. When we consider avatars that have 20–100% of their interactions labeled as H2B, the results relatively are insensitive to the minimum social connection threshold W. We conjecture that these avatars are very likely bots. If so, then if we set to 30% the percentage of avatar interactions that need to be labeled as H2B to trigger a bot detection, the bot detection strategy indicates that 4–7% of the avatars in SL are bots.

We do not yet have ground truth data to quantify the accuracy of the detection strategy, and a comprehensive evaluation remains future work. Anecdotally, though, we can apply the strategy to our crawler by injecting its movement patterns into the avatar traces. Using $W=3\%$, we observe that about 90% of the crawler interactions are labeled as H2B, suggesting that the strategy could rely upon high interaction thresholds to further reduce false positives. Moreover, we find that most of the false positives are generated at the beginning of the crawl, i.e., when the social network is not yet well-defined (Section 4.1). For example, in the first six hours the number of interactions between the crawler bot and the other avatars that are labeled as H2B increases from 10% to 90%. A bot detection strategy could further take into account these transient behaviors to improve accuracy.

6. CONCLUSIONS AND FUTURE WORK

This paper studies the Second Life social network and the role of bots within it. We find that the SL social network is more similar to natural networks than to popular online social networks. We also find evidence of substantial bot activity, and that bot interactions fundamentally impact the structure of the social network. We then propose a bot detection strategy based upon avatar connections in the social graph, and estimate a surprisingly large bot presence. In future work, we plan to systematically evaluate the accuracy of this bot detection strategy, and explore a design and deployment in practice.

Acknowledgement

7. REFERENCES

[1] T. B. Achacoso and W. S. Yamamoto. *AYs Neuroanatomy of C Elegans for Computation*. November 1991.

[2] R. Albert, I. Albert, and G. L. Nakarado. *Phys. Rev.*, (69), 2004.

[3] A. L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, October 1999.

[4] A. Chaintreau, A. Mtibaa, L. Massoulie, and C. Diot. The Diameter of Opportunistic Mobile Networks. In *Conext'07*, NY, USA, Dec. 2007.

[5] H. Chun, H. Kwak, Y.-H. Eom, Y.-Y. Ahn, S. Moon, and H. Jeong. Comparison of online social relations in volume vs interaction: a case study of cyworld. In *IMC'08*, Vouliagmeni, Greece, November 2008.

[6] A. Clauset, C. R. Shalizi, and M. E. J. Newman. Power-law distributions in empirical data, June 2007.

[7] Cyworld. http://www.cyworld.com.

[8] H. Edward. *The Hidden Dimension*. Doubleday, Garden City, 1966.

[9] Facebook. http://www.facebook.com.

[10] Flickr. http://www.flickr.com.

[11] H. Jeong, B. Tombor, R. Albert, Z. N. Oltvai, and A. L. Barabási. The large-scale organization of metabolic networks. *Nature*, 407(6804):651–654, October 2000.

[12] C.-A. La and P. Michiardi. Characterizing User Mobility in Second Life. In *WOSN'08*, Seattle, USA, August 2008.

[13] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and Analysis of Online Social Networks. In *IMC'07*, San Diego, USA, October 2007.

[14] A. Mtibaa, A. Chaintreau, J. LeBrun, E. Oliver, A.-K. Pietilainen, and C. Diot. Are you moved by your Social Network Application? In *WOSP'08*, New York, NY, USA, August 2008.

[15] Orkut. http://www.orkut.com.

[16] Second Life. http://www.secondlife.com.

[17] M. Varvello, F. Picconi, C. Diot, and E. Biersack. Is There Life in Second Life? In *Conext'08*, Madrid, Spain, Dec. 2008.

[18] D. J. Watts. *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press, 1999.

[19] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.

[20] Wikipedia. CopyBot. http://en.wikipedia.org/wiki/CopyBot.

[21] C. Wilson, B. Boe, A. Sala, K. P. N. Puttaswamy, and B. Y. Zhao. User Interactions in Social Networks and their Implications. In *EuroSys'09*, Nuremberg, Germany, October 2009.

[22] World of Warcraft. http://www.worldofwarcraft.com/.