

Player Dynamics in Massively Multiplayer Online Games

Xinyu Zhuang* **Ashwin Bharambe***
Jeffrey Pang* **Srinivasan Seshan***

October 2007
CMU-CS-07-158

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

*Computer Science Department, School of Computer Science, Carnegie Mellon University

Keywords: games, massively multiplayer, measurement, player dynamics, churn

Abstract

Millions of people now participate in Massively-Multiplayer Online Games (MMOGs), placing tremendous and often unpredictable maintenance burdens on their operators. Thus, understanding the dynamic nature of MMOG players is critical for the designers and implementers of the systems and networks that host MMOGs. This paper improves our understanding of player dynamics in MMOGs by collecting and analyzing a 5-month long measurement study of World of Warcraft, a leading commercial MMOG. Our novel findings include the following: First, the distribution of player session lengths is similar to that of peer-to-peer file sharing sessions, despite a game's interactive nature, but has a shorter tail. Second, we find several good predictors of session length, such as in-game character level or age. Third, despite the social nature of MMOGs, interactivity between players accounts for little of the dependence between sessions of different players, and most sessions are not correlated at long time scales. Fourth, changes to a game's virtual world can cause dramatic shifts in the population densities of in-game locations, which are otherwise relatively stable.

1 Introduction

The popularity of Massively-Multiplayer Online Games (MMOGs) has grown astronomically in the past few years, now with more than 12 million participants [17] and market value surpassing \$1 billion [20]. The most popular MMOG is Blizzard Entertainment’s World of Warcraft [25], which has 8.5 million subscribers alone [26] and is the focus of this study.

MMOGs are computer games where thousands of players interact with each other online in a persistent virtual world. Unlike other games where players can win or lose, players in an MMOG can never reach a state where they are considered to have won the game and, thus, usually subscribe to play for months or years. Players not only interact with each other and the virtual world, but sometimes also participate in building the virtual world itself (*e.g.*, in Second Life [22]). The large number of persistent connections, computation burden of maintaining a real-time virtual world, and bandwidth requirements of online content distribution each place tremendous demands on the systems and networks that host MMOGs. Therefore, operators must understand the dynamic nature of players in order to better manage their resources.

In order to reduce some of these burdens, there has also been interest in building MMOGs as distributed systems, for example, as peer-to-peer applications where game clients act as peers (*e.g.*, [2], [9], [16]). In such systems, player behavior is even more critical to the reliability of the game. In particular, distributed game designs must take into account the churn inherent in MMOG player participation.

This paper presents the results of an extensive study of player dynamics in World of Warcraft, the leading commercial MMOG. By making innovative use of the World of Warcraft client along with an infrastructure to automate the client, we monitored the dynamic behavior of over 1000 players within the game world for 5 months. In addition to these measurements, we contribute detailed analysis of the variation in player participation over time, the characteristics of player session lengths, downtimes, inter-arrival times, and availability, the aggregate player churn rate, the degree of player independence, and characteristics of stay time in different locations in the game. Our key novel findings include the following:

- The distribution of player session lengths is very similar to that of peer-to-peer file sharing sessions, despite a game’s interactive nature, but has a shorter tail.
- We find several good predictors of session length, such as in-game character level or age.
- Despite the social nature of MMOGs, interactivity between players accounts for little of the dependence between sessions of different players, and most sessions are not correlated at long time scales.
- Changes to a game’s virtual world can cause dramatic shifts in the population densities of in-game locations, which are otherwise relatively stable.

The remainder of this paper is organized as follows. Section 2 presents related work. Section 3 describes our measurement methodology and data sets. Section 4 presents detailed analysis of our measurements. Section 5 summarizes key implications of our results for operators.

2 Related Work

Previous work similar to ours falls into three categories: MMOG studies, studies of other game genres, and characterization of user behavior in other peer-to-peer systems.

User Behavior in MMOGs. Chen *et. al.*'s study of player interaction in the MMOG ShenZhou Online [6, 7] is the most similar work to ours. There are two key differences that distinguish our study: First, their measurement technique is markedly different than ours, making use of network-level packet traces collected from the local network hosting the MMOG. Our techniques do not rely on access to an MMOG operator's network and thus are more likely to be applicable to other games. Second, their study only examined player activity over periods of less than one day, whereas we measure and analyze player behavior over several months. Therefore our study reveals long-term trends that have not yet been studied.

The PlayOn [18] project is an on-going census and measurement study of World of Warcraft, with emphasis on social behavior and interaction [10]. Similar to our study, the PlayOn project also uses client modifications to gather their measurements. However, they focus on sociological aspects of gaming, whereas our focus is on the effects of player behavior on systems and networks provisioned for MMOGs.

User Behavior in Other Games. Chambers *et. al.* [5] conducted a measurement study of several online games, focusing on the first-person-shooter (FPS) game Counter-Strike. While some of the games that were measured can loosely be considered MMOGs, the majority were games that consist of separate game rounds lasting at most a few hours. For example, in FPS games like Counter-Strike, Quake, or Halo, rounds rarely last more than 30 minutes. Thus, player sessions in such games are heavily influenced by the length of game rounds. As expected, we observe that player behavior is markedly different in an MMOG. Feng *et. al.* [11] and Henderson and Bhatti [14] also present measurements of FPS games.

User Behavior in P2P Systems. As mentioned earlier, there have been several proposals for building multiplayer games as distributed systems [2, 9, 16]. Thus, it is useful to compare player dynamics in an MMOG to the peer dynamics in common peer-to-peer applications for which there has been considerable study.

Stutzbach and Rejaie [24] conducted an analysis of churn rates in P2P file sharing systems, and provided a set of guidelines to follow when attempting to measure churn. We use similar metrics for measuring player sessions in MMOGs. Sen and Wang [23] provide an analysis of peer-to-peer file sharing traffic, with some information on session lengths. Gummadi *et. al.* [13] also provide a characterization of session lengths in the P2P application Kazaa. In contrast to file sharing applications, MMOGs are interactive and players only remain online while actively participating in the game. Surprisingly, we find that some characteristics, such as session length, are not dramatically different in MMOGs despite this distinction. Nonetheless, we also find that there is not a stable core of long-lived players as there are long-lived peers in file sharing networks.

Guha *et. al.* [12] present a study of peer behavior in the Skype voice-over-IP P2P network and Yu *et. al.* [27] present a study of user behavior in a video-on-demand (VoD) system. As with MMOGs, there has been interest in building these types of systems in a P2P fashion. Although these applications also require some user attention during participation, they do not require as much

| Experiment | Length (days) | Sample Size | Mechanism | Polling Interval (minutes) |
|------------|---------------|-------------|-----------|----------------------------|
| Long-term | 71 | 1,100 | polling | 5 |
| Location | 21 | NA | polling | 2 |
| Detailed | 28 | 50 | callback | NA |

Table 1: Summary of the experiments conducted.

interaction as games and are not typically social in nature. Moreover, the average duration of phone calls is much shorter than the average MMOG session time, while sessions of participants in VoD systems will be heavily influenced by the length of videos, whereas MMOGs are open-ended.

3 Methodology

3.1 World of Warcraft

We first describe the object of this study, World of Warcraft (WoW), and aspects of the game that are pertinent to understanding our results. All of these characteristics have analogs in other Massively Multiplayer Role Playing Games (MMORPGs), the most common MMOG sub-genre. Thus, we believe our study is representative of most MMORPGs.

Realms. To accommodate 8.5 million subscribers [26], World of Warcraft uses *realms*, which are independent copies of the game world, each hosting several thousand players. Players choose a realm in which they wish to play in, and their character is then bound to that realm. The typical realm chosen for the measurement study was Frostmane.¹ While a realm might actually be served by a cluster of machines, in this paper we refer to the host of the realm as the *server*.

Most other popular MMORPGs, such as Everquest and Final Fantasy Online, use identical mechanisms to manage large subscriber bases (realms are also called shards).

Characters and Levels. A *character* corresponds to an in-game avatar controlled by a player, which possesses a set of attributes as well as an inventory of equipment. The key attributes that concern us are the character’s *name* and *level*. A character’s *name* is a unique identifier, and a character’s *level* is an in-game measure of the character’s experience. A higher level character is more powerful and has access to more content in the game. Characters gain levels by performing actions such as killing monsters or completing quests. A character’s level is strictly increasing from 1 to 70; characters cannot lose levels. Nonetheless, the average player takes months to advance a character beyond level 60. In this paper, we use the terms player and character interchangeably.

To our knowledge, all MMORPGs have some measure of character ability such as player level. Killing monsters and completing quests to advance levels are also common traits.

Factions. In each WoW realm, there are two opposing game *factions*, the Alliance and the Horde.

¹Frostmane is a typical Player-vs-Player realm. The other two realm types are Player-vs-Environment, in which the only difference is that players can not attack each other, and Role-Playing, which are advertised as hosting only players that act “in-character” but are not technically any different. We do not believe these differences in realm types dramatically impact our results.

Characters are assigned a faction upon creation. The game prevent members of one faction from discovering information about members of the other. Thus we were only able to monitor a single faction throughout our experiments. Since factions are purely in-game constructs, we do not believe that there are major differences in player behavior with respect to the metrics we are concerned with. The faction chosen for measurement was the Alliance.

Other MMORPGs that use factions include City of Heroes and Guild Wars.

Zones. The WoW game world is partitioned into *zones* or different locations in the game. For example, each town, city, and dungeon in the game is often a different zone. In general, movement from one zone to another is seamless.

Although we do not know whether zones in WoW are partitioned across servers, some other MMOGs, such as Second Life and ShenZhou Online, are known to host different zones on different servers. MMOG middleware, such as BigWorld [3], Project Darkstar [19], and Zona [28], all partition zones across different servers, while allowing for seamless travel between zones.

3.2 Measurement Infrastructure

The goal of our measurement study was to characterize player dynamics. Since such information is not made available by commercial MMOG operators and their infrastructure is closed to outside parties, we made use of the standard World of Warcraft client to obtain such measurements. Making use of the client as a measurement platform is also advantageous because it has minimal impact on the game being measured. Nonetheless, since the game client was not designed to be used as a measurement tool, we had to overcome a few challenges in order to make use of it this way.

Collecting Player Information. The first challenge was obtaining player information from the server. Fortunately, WoW implements a client-side user interface scripting environment where users can write scripts to manipulate parts of the client. Scripts are written in Lua [15], an embedded scripting language, and a collection of Lua scripts is usually referred to as an addon. This allowed us to use the client to gather information about players. We developed a set of measurement addons for each experiment that monitored player status over time (*i.e.* a player's online status, current level and location in the game).

Controlling the Client. We developed a manager program to control the client, so that it would automatically connect to the game server and collect measurements without human intervention.

3.3 Sampling the Player Population

Due to limitations in the features provided by the client, it was not possible to monitor the entire player population. Instead, we took samples of the entire population and measured players in the sample set. We developed an addon to collect as many player names as possible, and we created our samples by taking random subsets of the collected names.

The CensusPlus [4] addon uses a similar technique in order to publish a census about WoW characters. CensusPlus makes use of a WoW addon as well, but also gathers its measurements from

| Location | Description |
|--------------------|-------------------------------|
| Stormwind City | Busy central city. |
| Stranglethorn Vale | Quest area. |
| Burning Steppes | Along a popular travel route. |

Table 2: List of game locations that were chosen for measurement.

results submitted by a large number of volunteers. Thus their census reports a more accurate count of the number of players in the game: a little over 10,000 Alliance players. Our addon gathered 9,530 names, which makes the set of names we collected and sampled close to the entire player population.

3.4 Experiments

We now describe the three different experiments that we conducted. Table 1 summarizes the key features of each experiment.

Long-Term Measurement Experiment. This first experiment monitored the behavior of a random set of 1,100 unique players over 71 days, or nearly 3 months. The main goal of this experiment was to obtain an understanding of the behavior of a large number of players over a long period of time.

Players were monitored by polling the WoW server every 5 minutes to determine whether a player was online, and if so, the player’s level and current location in the game. Due to client limitations, polling the status of a single character took 5 seconds. Thus an entire poll took about 110 seconds. Polls were also spaced 5 minutes apart in order to keep the amount of traffic being sent to the game server from being too large.

Location Measurement Experiment. The location measurement experiment was designed to characterize the nature of player movement within game areas. Six locations (referred to as *zones* in the game) were chosen and continually monitored for players entering and leaving each location.

Table 2 lists three of the six locations we chose to measure, along with brief descriptions. Stormwind City is one of the capital cities for Alliance players. Stranglethorn Vale contains a large number of quests. Burning Steppes is along the only land and air route between two major game areas and thus acts as a transit area. We present the results for these three because they roughly correspond to the three main in-game activities: socializing/trading in cities, fighting in quest areas, and traveling.

A poll-based monitoring system was used for the location experiment (though using a different polling technique than the long-term experiment). Polls were conducted every 2 minutes, and attempted to obtain the details of every player found in the location. The exact time taken by each poll varied, depending on how many players there were in each location.

Note that game zones in WoW are large areas, and some may encompass several sub-zones (villages, caves, *etc.*). Thus it is not entirely accurate to characterize Stranglethorn Vale as just a questing area because it contains several important towns as well. Also, all zones have quests available to players. Nonetheless, our characterizations are relatively accurate when compared to

the other zones in the game.

Detailed Measurement Experiment. The 5 minute polling interval used in the long-term experiment implies that it did not capture changes at shorter time scales (in the order of seconds or minutes). Thus, we conducted another similar experiment with a much finer-grained level of measurement.

Unlike the previous two experiments, the detailed experiment did not use a polling mechanism to determine player status. Instead, a callback-based technique was used that allowed our addon to be immediately notified by the server when a player’s status changed. However, in order to use this more accurate measurement technique, we had to reduce the number of monitored players to 50.

3.5 Measurement Limitations

Player Aliasing. Due to limitations in the information exposed to clients by the game server, there is no straightforward way to correlate a character with a particular player (*i.e.*, no unique identifiers such as IP address or account number are revealed) or vice versa. Players are able to create multiple characters so it is not possible to tell whether two characters actually belong to the same player or not, introducing possible aliasing. However, players cannot log on with multiple characters at the same time. Potential aliasing implies that individual players may log in more frequently than our measurements suggest (with different characters).


Missing Players. Occasionally during polling, the measurement addon would receive an error from the WoW stating that the player being polled was not found. The most likely reason for this was that the player’s account was deleted. Another cause could simply be a case of a server error. Whenever this happened, we recorded the player’s status as being *missing*. Missing players were pruned out of our measurements. About 90% of the players we monitored were active at some point during our measurement period. In addition to this limitation, our experiments did not monitor any players that were not subscribers at the beginning of our measurement period (*e.g.*, players that joined later). Nonetheless, subscriptions generally last for months, so we do not believe that the behavior of these players would change our results substantially.

Measurement Outages. There were several sources of outages throughout the measurement period. The WoW server occasionally offline for a few hours, usually for maintenance purposes. There were also outages due to server-side errors.

The other major outages were due to patches to the game client. At two points during our measurement study, patches were released that were incompatible with the Lua measurement addons. In one case, the patch updated the in-game Lua interpreter to version 5.1, from version 5.0. This caused problems with the addons because they were not written for the newer version of Lua. In another case, an expansion set (The Burning Crusade) was released, which caused problems with parts of the addons that had to be fixed.

As a result of the outages, some of our datasets were partitioned. Figure 1 provides a visualization of our measurement periods and summarizes each continuous measurement period.

Timezone of Players. There was no way to determine the local timezones of the players being



| Experiment | Period | Days |
|------------|------------------------|------|
| Long-term | 11/15/2006 - 12/4/2006 | 19 |
| Long-term | 12/10/2006 - 1/3/2007 | 24 |
| Long-term | 2/5/2007 - 3/5/2007 | 28 |
| Location | 3/6/2007 - 3/27/2007 | 21 |
| Detailed | 3/27/2007 - 4/24/2007 | 28 |

Figure 1: Summary of measurement period and periods of continuous measurement.

monitored, but we do know that the server was using Central Standard Time (CST). Analysis of the data also yields a very obvious diurnal pattern, as can be seen in Figure 2, represented using Eastern Standard Time (EST). If we assume that most players enter the game in the evening, then EST or CST are the most likely timezones for the majority of players in the game. We use EST when discussing the results.

4 Analysis

The results of the measurement study are shown here. We first give a brief description of each characteristic we discuss.

Player Count. The number of players online in the game at some point in time.

Session Length. The time between a player’s login to the game and logout from the game.

Downtime. The time a player spends offline between sessions.

Availability. The fraction of time a player is in the game.

Player Independence. A measure of the correlation and dependence of different player’s play-times.

Playtime Concentration. A measure of how much individual players concentrate their playtime on certain times of the day or days of the week.

Inter-arrival time. The time in between two player arrivals, *i.e.* between one player entering the game and the next.

Staytime. The amount of time a player spends in a particular game location.

Location density. The number of players in a given zone at a particular time.

4.1 Player Count

Understanding the variation in player count allows system properties to be adjusted based on the load that will be experienced. To determine the player count in the long-term experiment, the total number of players found online at each polling period was determined. The resulting distribution has a median of 63 players (5.7%). There were never more than 200 players online at the same time, which given our sample size of 1,100 players is less than 18%.

Interestingly, only 993 players *ever* appeared in the game, about 90% of the total population being monitored. Many of the remainder had accounts deleted.

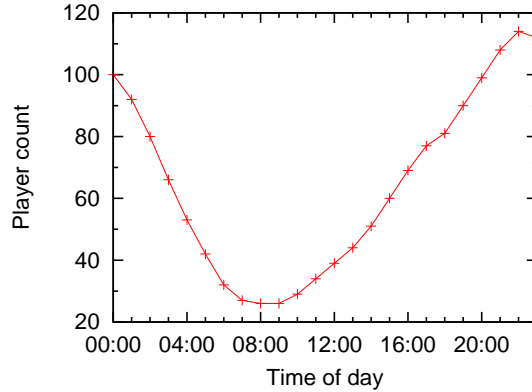


Figure 2: Median number of players per poll versus time of day the poll took place.

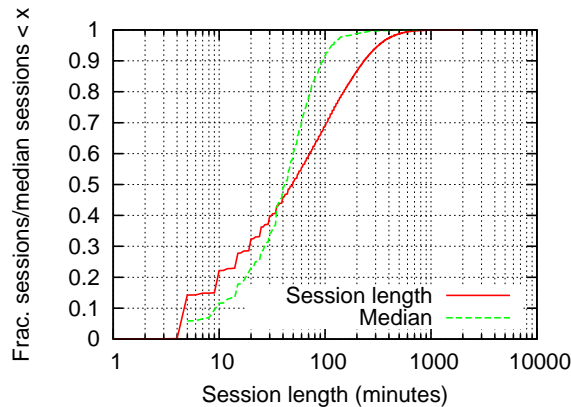


Figure 3: CDF of all session lengths, and median session lengths of each player. Note the logarithmic scale on the x-axis.

Player count is expected to vary with a diurnal pattern. Chambers *et. al.* [5] report such a pattern in player population of three popular games, as well as a weekly pattern. A plot of median player count versus the time of day the poll took place over all polls yielded the graph in Figure 2, which clearly demonstrates time-of-day effects. Possible day-of-week effects were also investigated, but no substantial patterns were found, *i.e.* people log on no more during the weekends than they do during the weekdays, in contrast to what was reported in [5].

4.2 Session Length

We now turn our attention to session lengths. Session lengths reflect the amount of churn a system faces, and also determine how long a peer can be used to provide service in a peer-to-peer system.

Figure 3 shows the CDF of the session lengths recorded during the long-term experiment, plotted on a log-linear scale. The median session length is 50 minutes. There were only 4 sessions

that lasted longer than a day.²

Another useful statistic to look at is the distribution of median session lengths of each player. This distribution characterizes the average behaviors of players, and is also shown on the same figure. The median of this distribution is 42 minutes. The longest median session length was around 12 hours, and the next longest was 7 hours.

An initial comparison of the two distributions shows that the median distribution has fewer extreme values. For example, the 95th percentile of the session length distribution is 1,976 minutes, as compared to that of the median distribution which is 117 minutes. Since we are looking at a median, this is expected, but the difference between the two also suggests that many of the extreme values in the session length distribution are actually distributed amongst all players. In other words, instead of there being a group of players that always have exceptionally long session lengths (and thus have a very high median), every player will occasionally have very long sessions along with shorter ones (yielding a more average median).

To investigate this hypothesis further, we determined the coefficient of variation (ratio of standard deviation to mean) of session length for each player. We obtained a median value over all players of slightly more than 1. This indicates large dispersion in the session lengths of *each player*, which agrees with our earlier speculation.

Comparison to Other Studies. The median session lengths of players in the ShenZhou Online MMOG reported in [6] for each of their traces was 122 minutes and 86 minutes, with a combined median of 100 minutes. The discrepancy from our measurements may be due to their much smaller measurement window; their median values were obtained by statistical methods in survival analysis, as a large number of their observed sessions were truncated due to the small measurement window.

Different measurement studies of P2P filesharing systems have produced differing results for median session lengths. These range from a few minutes [13] to close to an hour [21]. An average session length of 42 minutes falls within this range and is thus probably not that different from a P2P system. This is an interesting result, because the nature of filesharing applications is very different from that of an online game; a game player is almost always an active participant throughout the session, whereas a person downloading a file is likely to leave the filesharing application running while she does something else.

Results from Stutzbach and Rejaie [24] suggest that session lengths in P2P datasets (specifically BitTorrent) are not heavy-tailed, but best described by Weibull distributions. Analysis of the session length data collected by the long-term experiment gives similar results. A Weibull distribution with shape 0.84 and scale 4914.4 yields the best fit to our data. One important difference in our results and the results obtained by [24] is that they recorded sessions far longer than we did; some of their sessions lasted for weeks, whereas the longest session we found was slightly longer than a single day.

Predictability of Sessions. We now turn our attention to the predictability of session lengths. First, we investigate if there is a relationship between the lengths of the sessions a player has over the first week of the long-term experiment, and the lengths of all sessions a player has in total.

²Out of the 4, 2 were from the same player, and occurred slightly more than 2 weeks apart. In fact the top 10 longest sessions originated from a set of 5 players.

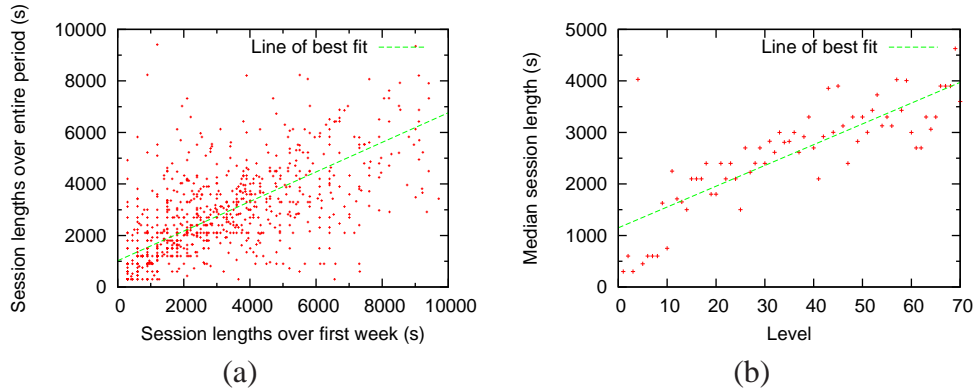


Figure 4: (a) Median session lengths over first week versus median session length over entire measurement period. (b) Level versus median length of sessions started at that level.

Figure 4(a) shows a plot of the median session lengths of players over the first week versus the overall median session lengths. Pearson’s correlation coefficient of these variables is 0.76, which means we can generally expect players who have had long sessions to continue doing so.

Another parameter that appears to influence session length is the level of the player at the beginning of the session. Figure 4(b) shows a plot of player level versus the median length of all sessions initiated by a player of this level. Pearson’s correlation coefficient for this dataset is 0.81, corresponding to a high degree of correlation. Thus, a higher level player is likely to have longer sessions.

Due to the strictly increasing nature of player levels, and the fact that most players’ goal in the game is to increase their level, we can generally assume that player level reflects the amount of time elapsed since the player created her character. Our results suggest that the longer a player has been involved with the game, the longer he is likely to play per session.

Finally, we see the same diurnal pattern with respect to both session lengths and session counts as in Figure 2. Longer sessions all begin in the evening and at night. Sessions that begin in the early morning tend to be short, particularly from the period from 5-7am. Thus, the length of a session could be predicted based on the time the session is started. In addition, since there are more sessions at night than in the day, this suggests that during the day there are both shorter sessions *and* fewer sessions.

Stutzbach and Rejaie [24] describe a pitfall in characterizing churn, where long sessions are not properly accounted for because of the limited size of the measurement period. We believe this is not a problem in our experiments, because sessions are rarely longer than a day, as compared to our measurement period of several weeks/months. We also investigated the effect of biases during the end of the measurement period by filtering out all sessions that began less than 2 days before the end of our measurements. Comparing the filtered dataset to the unfiltered one yielded negligible differences.

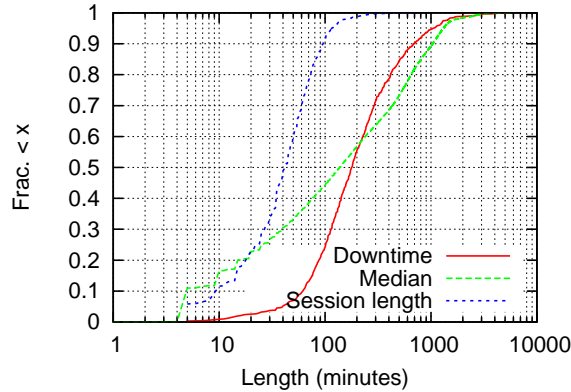


Figure 5: CDFs of median downtimes and median session lengths for each player.

4.3 Downtime

Downtime provides an indication of the time it will take an offline node to return online. It is of particular relevance to distributed games where nodes provide storage, because the data stored on a node will be unavailable during that node’s downtime.

The distribution of downtimes and median downtimes of each player is shown in Figure 5, together with the distribution of median session lengths from Figure 3 for comparison. Downtimes are longer than session lengths, as is expected since the average player is likely to spend less time in the game than out of it.

Surprisingly, the median of the median downtimes is 179.38 minutes, which is only about 3 hours. One would imagine this value be much higher, at the very least to accommodate for 7-8 hours of sleep, not to mention around 8 hours of work.

Comparing the downtime distribution to the median distribution, we find that there tend to be more short downtimes than players with shorter median downtimes. The 5th percentile of the downtime distribution is 300 seconds, while the 5th percentile of the median distribution is 2,400 seconds. This is similar to what we found with respect to session lengths and median session lengths.

Determining the coefficient of variance of each player’s downtime yielded a median coefficient of 1.22, indicating large dispersion in the downtimes of each player.

Thus, a possible explanation for the low median downtime is that players tend to have many sessions held close to one another, resulting in many short downtimes, and also a few sessions held very far apart, resulting in the much larger mean value. The period of downtime when the player is asleep or at work would then correspond to the fewer longer downtimes, which also suggests that a player has multiple sessions every night, held close together.

The longest observed downtime was about 4 days.

The relationship between downtime and the time of day as shown in Figure 6 is the inverse of that in 2, as expected. Nodes that leave in the early morning will not return for several hours, and nodes that leave in the evening will be back in a few hours or less.

It is possible to categorize the many short downtimes that players experience as *transient fail-*

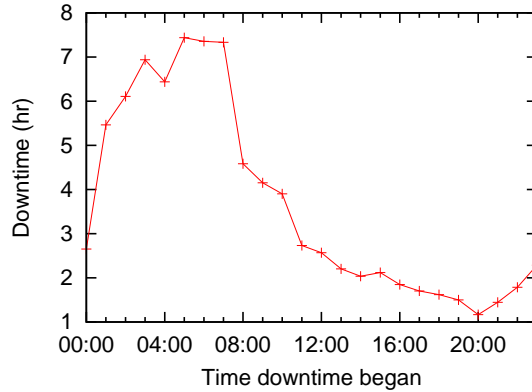


Figure 6: Relationship between downtime and time of day.

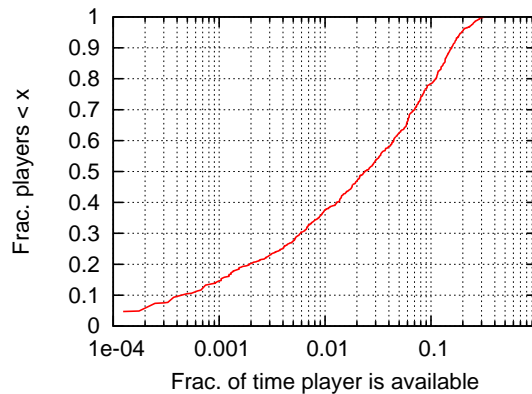


Figure 7: CDF of player availability over February.

ures, *i.e.* temporary disconnections from the system. In particular, they may not indicate that the player desires to stop playing the game; rather, they may be caused by network or machine failures, or the user needing to leave the game for a short time. In particular, there exists a mechanism in WoW that causes players to be logged out of the game after a few minutes of inactivity. Thus, if a player needs to leave his computer temporarily, she would be disconnected from the game. In games that do not have such a feature, the player would probably remain connected to the game (but not be active) for the short period of downtime we observed.

4.4 Availability

In distributed game designs where nodes are used to store game state, the level of replication required in the system is directly affected by how often the nodes are online. It is thus important to understand player availability. The availabilities of all players in the long-term experiment over the month of February (specifically from the 3rd period of measurement, Feb 5 2007 to Mar 5 2007) were determined by finding the ratio between the total amount of time spent online and the total amount of time elapsed (28 days). The distribution is shown in Figure 7.

The median availability is 0.024, *i.e.* around 2% or about 30 minutes of play time per day, on average. Note that 20% of players have availability over 10%, over 2 hours per day, on average. The highest recorded availability was 31.5%.

Comparison to Other Studies. Gummadi *et. al.* [13] determined the median availability for peer-to-peer clients to be 0.2%. This is only a lower bound, due to the methods used to characterize availability, and is far lower than our recorded 2%. Bhagwan *et. al.* [1] in contrast reports the availability of peers in the Overnet filesharing application to have a median of 7%. Thus, the availability characteristics of game players are different from that of file-sharing networks, but the distribution of peer availabilities are within an order of magnitude.

An additional complication with regards to availability is whether or not we can consider node availabilities to be independent. Peer-to-peer storage systems are very heavily influenced by dependencies between nodes; assumptions are typically made about node independence in order to decide on a suitable level of replication. For example, CFS [8] creates k replicas of each object, and assumes that the k nodes the replicas are on will fail independently. Thus, player independence (or lack thereof) should be taken into account when reasoning about availability. We discuss this issue in more detail in the next section.

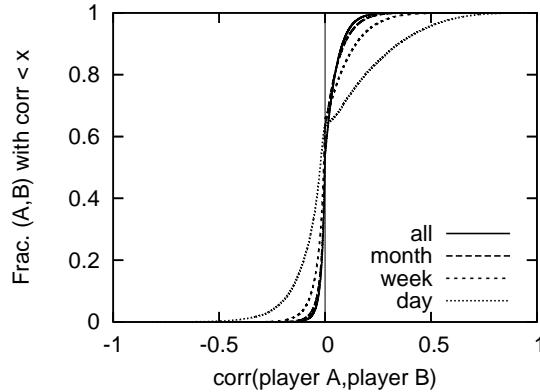


Figure 8: CDF of Pearson's correlation coefficient of each pair of player's online times.

4.5 Player Independence

Next we consider how independent player online times are. That is, we test whether we can consider the online times of players independent at various time scales and examine sources of dependence. As previously discussed, the degree of player independence determines the level of state replication needed in a distributed game. System could also exploit player dependence by finding sets of players that play during the same time. Then, if a player from such a set logs on or off, the system can then predict the other players will do so soon and take preemptive measures.

Figure 8 shows a CDF of Pearson's correlation coefficient $\text{corr}(\text{player A}, \text{player B})$ for each pair of players' online times at various time scales. Each player's time series is represented with a 1 when online and 0 otherwise (using 5 minute buckets, ignoring times when we are missing measurements). We saw earlier that there is a strong diurnal pattern in online times so we expect strong daily positive dependence in when players play, as shown by the "day" line. In fact, there is also daily negative dependence between some pairs of players because players focus their play times at slightly different times of day. Nonetheless, we see that the correlation between online times diminishes when we examine longer time scales. Indeed, at the timescale of 1 month, over 75% of player pairs have correlation coefficients between -0.05 and 0.05 which suggests that their online times are close to independent. This supports our previous finding that there did not appear to be any longer term periodicity in online times.

Nonetheless, Figure 8 also demonstrates that a small percentage of player pairs have higher correlation in their online times even at long time scales. We investigate the cause of this correlation by examining these player pairs. We define player A and player B to be *partners* if the correlation of their online times $\text{corr}(A,B) > t$.

Figure 9(a) shows a CDF of the number of partners per player when looking at a timescale of 1 month for several moderate correlation thresholds t . The majority of players have no partners and most that do have partners only have 1 or 2. Indeed, only 2% of players have partners with correlation > 0.5 , so most partner relationships are fairly weak. Figure 9(b) shows a graph of partner relationships for $t = 0.3$, where each vertex is a player and an edge is present if two players are partners. This graph demonstrates that the majority of partner relationships are pairwise. However,




Figure 9: (a) Number of partners per player, defined by high correlation. (b) A graph of players (vertices) and partners (edges) for $t = 0.3$. (c) CDF of the % time spent in the same zone for partners vs non-partners. All figures consider a monthly timescale.

there are 20 to 30 players that are partners with 10 to 30 players in the same cluster. These players tend to play more than others (all are online at least 14% of the time), but there are more players that play just as much and have few or no partners.

Overall, these results appear to be consistent with the findings of Chen et al. [7], which suggest that, at session time scales, most players play solo and the majority of groups are “duos.” Our results suggest that these relationships persist to longer timescales. To test this hypothesis, we attempt to determine whether this correlation in online times is actually due to “players playing together” or whether the correlation is due to other external factors. We conjecture that players that are playing together will exhibit locality in the virtual world. That is, they will tend to play in the same zones at the same time so that they can interact.

Figure 9(c) shows a CDF of the percentage of time, when online simultaneously, that partners and non-partners are in the same zone, for $t = 0.3$. We only consider player pairs that were simultaneously online for at least 1 hour so that we have a sufficient number of samples. If correlation in online times is due to players playing together, we expect partners to play more often in the same zone than non-partners. This is what we observe in the figure. However, partners are only 2 times more likely to play in the same zone than non-partners, on average (6% vs. 3%). Moreover, Kendall’s τ correlation coefficient of $\text{corr}(A,B)$ and the percentage of time spent in the same zone is only 0.05, indicating that the relationship between the two, while positive, is weak. Therefore, there must be external factors beyond players playing together that cause correlated online times. This conclusion implies that interactivity between two players in a game is unlikely to be sufficient to predict how strongly their online times will be dependent at long time scales.

In summary, we find that the majority of players’ online times are close to pairwise independent at timescales longer than 1 week. Of those players that have online times that show moderate positive correlation, most are only correlated with 1 or 2 other players. Some of this correlation is explained by players playing together, but other external factors also appear to play a part in dependence of online times.

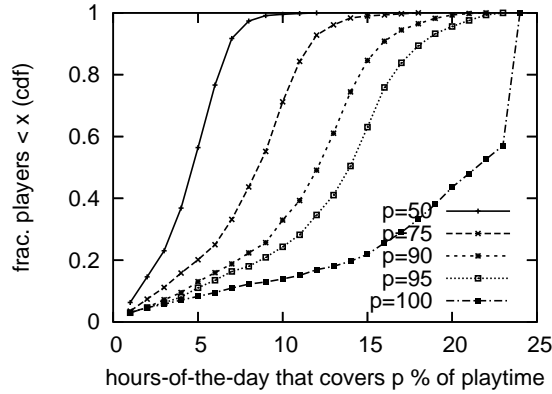


Figure 10: CDF of per-player concentration of playtime on different hours of the day.

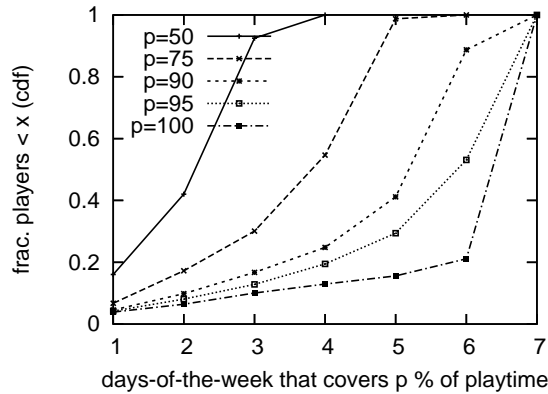


Figure 11: CDF of per-player concentration of playtime on different days of the week.

4.6 Playtime Concentration

So far, we’ve only looked at player availability in aggregate. Another question is whether individual players concentrate their play time during particular times of the day or days of the week. If a player’s playtime is concentrated during particular periods, the game can better predict when particular players will be online and offline using their history.

Figure 10 shows a CDF of the per-player concentration of play time on different hours of the day. Each line p corresponds to the minimum number of hours of the day (12AM, 1PM, 2PM, etc.) that cover at least $p\%$ of total playtime. For example, the $p = 50$ line shows that about 55% of players spend at least 50% of their playtime during at most 5 different hours of the day. Note that if a player play time was uniformly spread over X hours, it would require $X/2$ hours of the day to cover the majority of it. The $p = 95$ line shows that 14 different hours to account for 95% of most players’ playtime. Therefore, it appears that the majority of players spend almost all of their time playing on one half of the day, as expected (e.g., on the half during which they are awake).

However, most players do not appear to have more specific “designated” times of day during which they play frequently. For example, the $p = 50$ line also shows that less than 40% of players




Figure 12: CDF of inter-arrival times.

spend the majority of their playtime during 4 or fewer hours of the day. Note that even a player that was equally likely to play during any hour in one half of the day would spend the majority of their playtime on only 6 hours of the day. Moreover, the $p = 95$ line shows that 14 different hours are required to account for 95% of most players’ playtime. This implies that, beyond the side of the day on which a player is awake, players do not concentrate their playtime during a small number of hours.

Figure 11 shows a CDF of the per-player concentration of play time on different days of the week (Sunday, Monday, Tuesday, etc.). The $p = 50$ line shows that only 42% of players play the majority of their time on at most 2 days of the week. Note that even if a player played equal amounts of all days of the week, he or she would spend the majority of their time on only 3.5 days. The $p = 95$ line shows that most players spread 95% of their playtime over at least 5 days of the week. Therefore, most players do not appear to concentrate their play time during a small number of days of the week.

We conclude that although most players spend the vast majority of their play time on one side of the day, beyond that, most players do not strongly concentrate their play time during a few hours of the day or days of the week. This suggests that most players do not have “set schedules” for when they decide to play.

4.7 Inter-arrival Times

We now discuss node inter-arrival time. The inter-arrival time is a useful metric for characterizing how frequently players enter the game, and affects performance based on the overhead incurred when adding a new node to the system.

An important aspect about inter-arrival times is that they are likely to become shorter when the population increases. Thus, when measuring inter-arrival times over a large population such as with the long-term experiment, we often see inter-arrival times on the order of seconds. Since the long-term experiment has a polling interval of 5 minutes, it is not suited for gathering statistics at such a fine granularity. We thus focus on the results obtained by the detailed experiment instead.

Figure 12 shows the distribution of inter-arrival times as gathered from the detailed experiment.

The median of the distribution is 1189s, around 20 minutes.

While Stutzbach and Rejaie [24] suggest that inter-arrival times are likely to deviate from an exponential distribution due to time-of-day effects, fitting our data to an exponential distribution yielded a decent fit with the rate parameter $\lambda = 0.00049$, although there was deviation at the tail. We decided to work with the assumption that the inter-arrival time distribution we measured was exponential, and used that to determine a relationship between sample size and inter-arrival time.

If we consider the arrival of a player to be a Poisson process, then the inter-arrival time distribution is the result of 50 Poisson processes, each having a rate of around 0.00001 per second. Thus, the mean inter-arrival time of a sample of size n is $1/(0.00001n) = 100000/n$.

There were an estimated 10,000 players on the server we monitored. This yields a mean inter-arrival time of 10 seconds, which is quite short.

Churn Rate. We also determined the *churn rate* of players entering and leaving the game, *i.e.* the rate of both arrivals and departures. Our results show that an average of 1.1 events per minute. In addition, each event roughly corresponded to a single player; over an interval of 30 minutes, each player who entered or left the game during the period contributed an average of 1.2 events. This agrees with the finding that inter-arrival times are inversely proportional to the population size.

If we assuming that each event was caused by a unique player, this is equal to 0.1% of the population either entering or leaving the game every minute. Given that the median number of players online at any time is close to 6% of the total population, that amounts to 1.6% of the online population.

4.8 Stay Time

The staytime of a zone has similar impacts on location-based schemes as session lengths do on systems in general. For example, if a system partitions the game world such that each location is handled by a different machine, short staytimes imply frequent migration of players between those machines.

We expect staytime to vary according to characteristics of each zone. Staytime was determined from the location measurement experiment.

Figure 13 shows the distribution of stay time for 3 of the zones monitored in the location measurement experiment. Note that all 3 curves begin at the 2 minute mark, because the minimum staytime we were capable of recording was 2 minutes. It also does not make sense to discuss a staytime of 0, since that would correspond to a player *not* being in the location.

As is expected, the staytime for Burning Steppes (along a transport route) is very short, with the majority of players leaving the location within 2 minutes. Stormwind City in contrast has a median staytime of about 7 minutes, and Stranglethorn Vale (a questing area) has a median staytime of 18 minutes. The distributions of staytime for Stranglethorn Vale and Burning Steppes were very positively skewed, less so for Stormwind City. We were able to fit Stormwind City's staytime distribution to an exponential distribution, whereas the other locations were better fit by Weibull distributions.

Level and Staytime. The relationship between player level and staytime was also investigated.

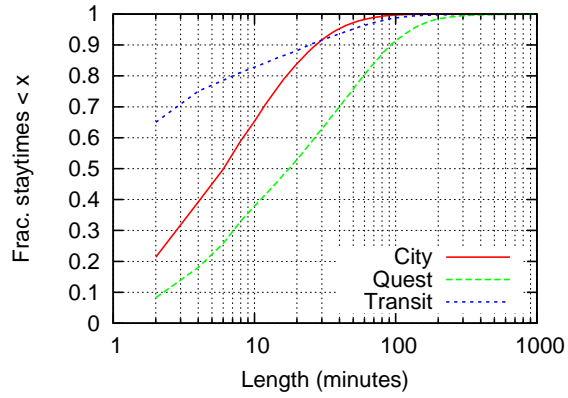


Figure 13: CDFs of the staytimes of city (Stormwind City), quest (Stranglethorn Vale) and transit (Burning Steppes) locations.

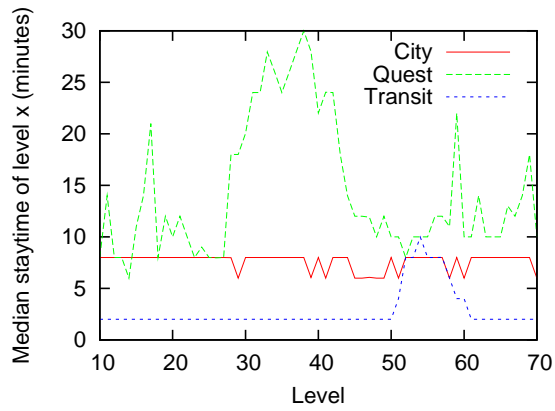


Figure 14: Median staytimes versus player level for 3 different locations.

Figure 14 shows the variation in the median staytime (of all players of a certain level) over player levels. We leave out the staytimes for levels 1 through 10 because there were too few players of those levels found in Stranglethorn Vale and Burning Steppes.

This graph has several interesting features. First, the median staytime of players in Stormwind City seems to be independent of level. In contrast, there are noticeable peaks in the staytimes of the other two locations, particularly at the 30-45 range for Stranglethorn Vale and the 50-60 range for Burning Steppes. These peaks exactly correspond to the *level range* of those locations. The level range refers to the recommended range of levels that players should be if they wish to enter a particular location. Players who have a level below the range are unable to defeat most of the monsters in the location, while players whose level is too high will gain little benefit from doing so. In addition, most of the quests in a location can only be performed by players within the level range or higher.

This suggests that the players who had long staytimes in Stranglethorn Vale and Burning Steppes were mainly there to perform quests, which take some time. Players who visited Stormwind City however were probably there for trade purposes (purchasing items, selling items) and left the area once their business was done. Finally, the players who visited Burning Steppes not to perform quests were on their way to someplace else and passed through very quickly.

One would expect that the session lengths of players in a location of appropriate level would be longer than that of other players, at the very least by virtue of the fact that their staytimes are long. To find a relationship between location, level and session length, we used the data from the long-term experiment and determined the remaining amount of time a player spent in the game (as opposed to just the location) after the player entered various locations. Despite our expectation, there was no discernible relationship between a player's level and the amount of remaining time she spent in the game upon entering a fixed location, nor was there a relationship between the current location and the remaining game time. We did however find a relationship between player level and staytime similar to that between player level and session length in the long-term dataset.

We speculate the lack of relationships are due to the fact that players do not spend all their online time in a single location. A low-level player might spend a short time in a high-level area, before proceeding to a low-level area where she spend a lot of time. Calculations of the amount of locations a player passes through in a session, which we term *mobility*, yielded a median mobility of 3 locations per session.

There are likely to be various other location properties that affect a player's staytime, for example the presence of more treasure or special quests. However, it is difficult to accurately characterize game locations in order to make these relationships evident, so these characterizations remain as future work.

4.9 Location Density

Another location-related metric we studied is the location density. Dense areas with many players generally require more processing and bandwidth usage, because there is likely to be more interaction between players. It is particularly interesting to observe the *stability* of location density in the game world, *i.e.* how the densities of various locations change over time. This can impact




Figure 15: Differences in mean location density over December and over February.

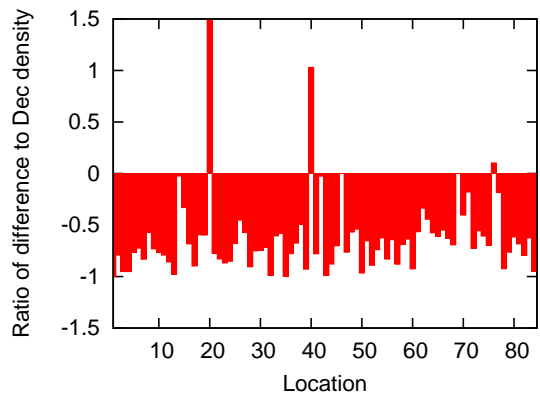


Figure 16: Ratios of difference in mean location density and mean location density over December.

load balancing decisions, for example how often systems should re-evaluate resource assignment to locations.

We make use of data from the long-term experiment to determine location density, as opposed to from the location experiment, because we wish to characterize a large number of locations. In the course of the long-term experiment, players were found to be in a total of 117 zones. Location density was determined by counting the number of players that were in a particular zone at each poll. Note that this is the number of players out of the sample set of 1,100 players, not the entire population.

Out of the 117 locations found, 33 were locations that were only accessible to players after the release of the expansion set. Thus, there were no players found in these locations before the February dataset. In order to differentiate between these new locations and the old ones, they were grouped together as locations 84-117 in our results.

Figure 15 shows the differences in mean location densities over December and over February, for each of the 117 locations. The vertical line delineates the locations present in the original game and the new locations only present in the expansion. A negative value indicates that there were less players in this location overall during February as compared to during December.

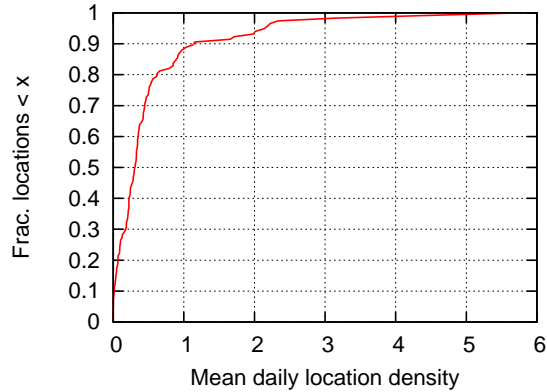


Figure 17: CDF of mean location density for each location.

As can be seen from figure, the change for most locations was negative, which meant that there were more players in those locations in December than there were in February. In addition, we found that a large majority of players were spending their time in the expansion locations. Thus, there was a definite shift in density before and after the release of the expansion.

In Figure 16, the differences in density were divided by the mean density over December. We assigned a ratio of 0 for locations where the mean density over December was zero, and we omitted the expansion locations since they did not exist in December. The location with the largest density ratio is actually truncated; its density ratio is close to 8^3 . Again, we can see that the densities of many locations dropped by at least 50%.

We also found that there were in general a few very dense areas, and many more less dense areas. This was particularly true for December, where we observed one location with very high density. This location turned out to be Ironforge, a capital city which also holds the auction-house, where players may bid for and offer items.

The distribution of the mean location density of each location is shown in Figure 17. This clearly confirms the earlier observation of there being a few exceptionally dense areas, as the tail of the graph is very long.

Finally, we observed the variation in density over the month of February for various locations. We found that density was stable over the month; although the location density varied depending on the number of players who were online, in general locations with high location densities had values that remained high throughout the month.

These results indicate that over a moderate time-period (*e.g.*, a month), location densities are likely to be stable, but popularity will shift dramatically due to in-game changes. There are also locations in the game which will almost always be crowded, and others that are usually nearly empty. A similar claim was presented in [7]; our results also demonstrate the stability of this property over time.

³A new dungeon was added to this location (Deadwind Pass) in the expansion, which caused the popularity of the location to increase sharply.

5 Summary and Implications

This section summarizes our results and discusses their implications for the design of massively multiplayer games.

5.1 Summary

Although most session lengths in MMOGs are not long, they are still within acceptable levels for many P2P support systems, such as DHTs. We also find several predictors of session lengths, such as a player's level and a player's previous median session length.

Nonetheless, session lengths are generally distributed evenly amongst all players, as opposed to there being particular players with consistently long sessions. In addition, the longest game sessions found were far shorter than the longest peer-to-peer filesharing sessions observed by other studies, *i.e.* game players are less long-lived than filesharing peers in the limit.

In terms of player availability, players spend less time online than they do offline, as expected. Nonetheless, most are available 2% of the time and 20% fraction are available more than 10% of the time, well within an order of magnitude of most filesharing peers, which is surprising given a game's interactive nature. This corresponds to most peers being online for at least 29 minutes every day, on average, and 20% of players being online for more than 2 hours every day, on average.

The distribution of inter-arrival times we have found suggests an inverse relationship between inter-arrival times and population size, where the interarrival-time of a population of size n is $100000/n$.

Players appear to play at independent times over longer timescales (more than a week), and while some dependence is due to players playing together, other external factors also appear to contribute to correlation. At shorter timescales, the diurnal effect causes to dependence in play times.

Players spend long periods of time only in areas where they stand to gain benefits and advance in the game (*i.e.* areas with suitable quests). In this regard, players appear to act rationally with respect to game rules. The distribution of location densities is very skewed, with some locations much more popular than others. Density is also stable over shorter periods of time, although major in-game changes, such as the addition of new locations, can cause density to shift.

5.2 Recommendations

Overall, our results suggest that multiplayer game designs should take into account the high churn rate present in large games, and ensure that any overhead incurred when players join and leave the game is sufficiently low. This also applies to systems where there is overhead in a player entering and leaving a game location, such as games where different zones are hosted on different servers. In addition, games should allocate more resources to popular areas as opposed to uniformly across all locations. Dynamic load balancing is not as necessary since density tends to be stable, and drastic changes in density only appear to be caused by major changes in the game that can be anticipated.

Predictors of location staytimes can be used to determine how long players will stay in a location and thus use the resources associated with that location.

Systems should also take into account daily fluctuations in parameters such as population count, for example arranging schedules such that a maximal amount of resources are available during the daily peaks.

Distributed games face challenges with regard to making use of game clients to maintain state. Node availability is low and churn rates are high, thus any design where state is stored on the clients must use heavy replication. Games should also use the predictors of session length and staytime we have found to identify suitable nodes for hosting game state. However, identifying long-lived players in order to form a backbone like what was suggested in [24] has far less benefit in the context of distributed games than it does in filesharing applications, because game clients are shorter-lived.

For state that does not last long and is frequently accessed, designs must take into account the daily dependence in node availability due to the diurnal effect. State which is longer-lived however can benefit from player independence at longer timescales. Our results also suggest that distributed system simulations can safely model most peers as independent in a distributed MMOG.

The session characteristics of players suggest that Distributed Hash Tables (DHTs) that are resilient to churn are suitable for use in distributed games. Rhea *et. al.* [21] showed that most DHT systems can cope relatively well with median session lengths of around 40 minutes, although any less than this quickly leads to very poor performance. Thus, making use of a DHT in the construction of a distributed game, as was done in [2, 16], is feasible, as far as coping with churn is concerned.

5.3 Future Work

Our measurement study is a first step towards a better understanding of player dynamics in multiplayer games. We plan to make measurements of similar properties in other MMOGs such as Second Life, in order to examine what effects, if any, different game mechanics and rules have on player behavior. We also hope to perform trace-driven evaluation of system designs using our observations of player behavior. In particular, we wish to examine how such player dynamics influence the performance of recently proposed distributed game architectures as well as distributed object stores.

References

- [1] BHAGWAN, R., SAVAGE, S., AND VOELKER, G. Understanding availability. In *Proceedings of the 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03)* (Feb. 2003).
- [2] BHARAMBE, A., PANG, J., AND SESHAN, S. A Distributed Architecture for Multiplayer Games. In *Proceedings of the Symposium on Networked Systems Design and Implementation* (San Jose, CA, May 2006).
- [3] Big World. <http://www.microforte.com/>.

- [4] Censusplus world of warcraft census. <http://warcraftrealms.com/census.php>.
- [5] CHAMBERS, C., FENG, W.-C., SAHU, S., AND SAHA, D. Measurement-based characterization of a collection of on-line games. In *Internet Measurement Conference* (2005).
- [6] CHEN, K.-T., HUANG, P., AND LEI, C.-L. Game traffic analysis: An MMORPG perspective. *Computer Networks* 51, 3 (2007). Article In Press.
- [7] CHEN, K.-T., AND LEI, C.-L. Network game design: Hints and implications of player interaction. In *Proceedings of ACM NetGames 2006* (Singapore, Oct 2006).
- [8] DABEK, F., KAASHOEK, M. F., KARGER, D., MORRIS, R., AND STOICA, I. Wide-area cooperative storage with CFS. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP '01)* (Chateau Lake Louise, Banff, Canada, October 2001).
- [9] DICKEY, C. G., ZAPPALA, D., AND LO, V. A fully distributed architecture for massively multiplayer online games. In *NetGames '04: Proceedings of 3rd ACM SIGCOMM workshop on Network and system support for games* (New York, NY, USA, 2004), ACM Press, pp. 171–171.
- [10] DUCHENEAUT, N., YEE, N., NICKELL, E., AND MOORE, R. J. “alone together?”: exploring the social dynamics of massively multiplayer online games. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems* (New York, NY, USA, 2006), ACM Press, pp. 407–416.
- [11] FENG, W., CHANG, F., FENG, W., AND WALPOLE, J. Provisioning on-line games: A traffic analysis of a busy counter-strike server. In *Proceedings of the Internet Measurement Workshop* (Nov 2002).
- [12] GUHA, S., DASWANI, N., AND JAIN, R. An Experimental Study of the Skype Peer-to-Peer VoIP System. In *Proceedings of The 5th International Workshop on Peer-to-Peer Systems (IPTPS '06)* (Santa Barbara, CA, February 2006), pp. 1 – 6.
- [13] GUMMADI, K. P., DUNN, R. J., SAROIU, S., GRIBBLE, S. D., LEVY, H. M., AND ZAHORJAN, J. Measurement, modeling, and analysis of a peer-to-peer file-sharing workload. In *SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles* (New York, NY, USA, 2003), ACM Press, pp. 314–329.
- [14] HENDERSON, T., AND BHATTI, S. Modelling user behaviour in networked games. In *MULTIMEDIA '01: Proceedings of the ninth ACM international conference on Multimedia* (New York, NY, USA, 2001), ACM Press, pp. 212–220.
- [15] IERUSALIMSCHY, R., DE FIGUEIREDO, L. H., AND FILHO, W. C. Lua – an extensible extension language. *Softw. Pract. Exper.* 26, 6 (1996), 635–652. Also see <http://www.lua.org/>.
- [16] KNUTSSON, B., LU, H., XU, W., AND HOPKINS, B. Peer-to-peer support for massively multiplayer games, 2004.
- [17] MMOG chart. <http://www.mmogchart.com/>.
- [18] PARC, X. Playon project.
- [19] Project Darkstar. <http://darkstar.thresher.com/>.

- [20] Research: Western mmo market passes \$1bn in 2006. Gamasutra, Mar. 2007. http://www.gamasutra.com/php-bin/news_index.php?story=13233.
- [21] SEAN RHEA, DENNIS GEELS, T. R., AND KUBIATOWICZ, J. Handling churn in a dht. Tech. Rep. UCB/CSD-03-1299, EECS Department, University of California, Berkeley, 2003.
- [22] Second life. <http://secondlife.com/>.
- [23] SEN, S., AND WANG, J. Analyzing peer-to-peer traffic across large networks. In *IMW '02: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement* (New York, NY, USA, 2002), ACM Press, pp. 137–150.
- [24] STUTZBACH, D., AND REJAIE, R. Understanding churn in peer-to-peer networks. In *IMC '06: Proceedings of the 6th ACM SIGCOMM on Internet measurement* (New York, NY, USA, 2006), ACM Press, pp. 189–202.
- [25] World of warcraft. <http://www.worldofwarcraft.com/>.
- [26] WoW hits 8.5m subscribers. CVG, Mar. 2007. <http://www.computerandvideogames.com/article.php?id=159422>.
- [27] YU, H., ZHENG, D., ZHAO, B. Y., AND ZHENG, W. Understanding user behavior in large-scale video-on-demand systems. *SIGOPS Oper. Syst. Rev.* 40, 4 (2006), 333–344.
- [28] Zona. <http://www.zona.net/>.