

Characterizing Virtual Populations in Massively Multiplayer Online Role-Playing Games*

Daniel Pittman and Chris GauthierDickey

Department of Computer Science
University of Denver
{dpittman | chrisg}@cs.du.edu

Abstract. Understanding player distributions, sessions, and movements in a Massively Multiplayer Online Role-Playing Game (MMORPG) is essential for research in scalable architectures for these systems. We present the first detailed measurement study and the first models of the virtual populations in two popular MMORPGs, World of Warcraft™ and Warhammer Online™. Our results show that while these two types of MMORPGs are significantly different in play style, the features of their virtual populations can be modeled similarly, allowing future researchers to accurately simulate these types of games.

1 Introduction

Measuring and modeling player distributions, session lengths, and player movements in a virtual world are essential to research in architectures for massively multiplayer online games (MMOs). Accurate models based on empirical evidence significantly strengthens researchers' arguments that one particular architecture is better than another. While prior measurement research has been done on MMOs, it has primarily focused on traffic modeling and characterization. Though traffic measurement is important and useful for research in MMOs, an understanding of how players move around and populate the virtual world allows one to explore architectures that take advantage of this information and analyze those that do not.

Our research addresses this issue. We provide the first set of measurement-based models for population distributions and movements *within* two different massively multiplayer online role-playing games (MMORPGs) that can be used for simulation and analysis of new architectures. MMORPGs are a subset of all possible MMOs, and while other types of MMOs are possible (for example, one based on real-time strategy rules), to date MMORPGs have been the most commercially successful.

The ultimate goal of our research is to design a realistic and empirically-based simulation model from measurements taken from current, commercial MMORPGs. To achieve this, we have measured overall populations, session lengths, player distributions, and player movements over several months on two MMORPGs: World of Warcraft™ and Warhammer Online™, both which are classified as MMORPGs but have significantly different play styles. We measured these two games in order to test our

* Published in the *Proceedings of The 16th ACM International Multimedia Modeling Conference (MMM 2010)*. Original publication can be found at <http://www.springerlink.com>.

hypothesis that regardless of the style of play in the game, players' behavior could be modeled using a unified set of functions. We find our hypothesis to be correct. While both games have significantly different play styles, the resulting models are quite similar and would allow a researcher or game designer to simply change a few parameters in order to simulate various scenarios. While the models we generated may not translate to all types of MMOs, these two representative games have provided excellent models for future research and simulations.

2 Related Work

Over the last several years, a significant number of measurement studies have been performed on MMOs, though most of these have focused on traffic patterns and network characteristics. Chambers et al. studied network patterns related to players and the server of small networked multiplayer games [1]. Similar to our data, their measurements also show diurnal patterns of game populations. Kim et al. measured network patterns on Lineage II, a popular MMO in Korea [5]. Their work focused on network packet sizes, RTTs, session times and inter-session arrival times. The data they recorded shows a similar power-law distribution for session times as the times that we have observed. Ye et al. devised a set of performance models for MMORPG servers and networks based on concurrent player population [11]. Chen et al. profiled packet interarrival times, packet load distribution, and bandwidth utilization of *ShenZhou Online*, a popular MMORPG [2]. Svoboda et al. modeled traffic patterns and sessions lengths for players of WoW using both wireless and wired Internet connections [9].

Beyond network traffic measurement, some research has looked at traffic patterns, session lengths, and latency when measured with respect to players and user behaviors. Tarng et al. performed a long term study of WoW in order to see if it was possible to model subscription lengths of players based on how much a user plays an MMO [10]. They showed that while it is possible to predict short term behavior, long term prediction is much more difficult. Claypool and Claypool characterized latency requirements of various online games in terms of the deadline in which a user command must be processed and the precision of the commands the user is issuing [3]. Traffic patterns and session lengths of WoW were profiled with respect to different player action categories by Suznjevic et al. [8]. They hypothesized that mobile devices could be used for some of the less traffic intensive player activities. Fernandes et al. characterized traffic patterns in Second Life during different player activities in the world [4]. Kinicki et al. expanded on the work of Fernandes et al. by considering object and avatar interactions of the player in the virtual world when modeling traffic characteristics [6]. Finally, Szabó et al. provided a model from which you can detect the activity of a user within a MMORPG by correlating the traffic patterns observed through passive monitoring and packet level introspection.

While all of the related work provides important contributions towards modeling MMOs, especially in terms of traffic behavior, our work is the first to provide details of the virtual world, its population distributions, and player movements. Note that our preliminary results were published in [7], but these results only show initial measurements of World of Warcraft over a smaller data set without any modeling. In this paper,

we examine a much larger and complete data set of two MMORPGS, we create models for simulation and analysis, and we show the similarity between both MMORPGs even though their play style differs significantly.

3 Methodology

Two methods can be used for measuring virtual populations and behaviors of players in MMORPGs. The first method is to analyze logs generated directly from an MMORPG or from customized clients which log the behavior for you. This method has the advantage of being accurate, though few companies are willing to share logs from their games or allow you to modify their clients and further the logs may not contain the needed information. The second method is to use probing-based measurements to try to infer properties of the system. We use the second method for our research.

In order to measure population information, we designed a set of scripts that run from the game clients using the Lua¹ scripting interface provided by both WoW and WAR². For WoW, we modified the Census+ add-on to collect broad information about all players currently online³. WAR's add-on was custom written, but was based on functionality of Census+. We also wrote an additional add-on for both MMORPGs to record continuous detailed information about a randomly selected subset of players.

To measure the games, we performed server queries from the clients using the *who* service, which allows a player to search for another player in the game, and the *friends list*, which is updated by the server whenever the client queries the friends list. Using the *who* service, we performed back-to-back snapshots of the server population. However, since an entire snapshot takes several minutes, we populated our friends list with a random set of players which were not seen in the previous snapshot. The friends list allows us to track a small subset of players including when they log on and off and where they are at during each query. Using our techniques, we observed over 115,000 individual players and tracked player movements on over 75,000 sessions. Note that while all MMOs do not use Lua as a scripting interface, the *who* service and friends list tends to be universal and therefore similar techniques could be applied to other MMOs.

From both games we measured the total populations over time, the lengths of each session observed, the zones each player visited (including the order visited), and the time spent in each zone. WoW was measured over a 4 month period on the Aerie Peak server while WAR was measured over 2 weeks on the Volkmar server. We examined data from other servers and it was similar to the results presented here, thus these two servers are sufficiently general for both games.

4 Measurements and Models

When analyzing the data we recorded, our goal is two-fold. First, we want to verify previous work regarding arrival rates and session lengths. Prior research has done this

¹ <http://www.lua.org>

² Source code available from <http://www.cs.du.edu/~chrisg/measurements>.

³ <http://www.warcraftrealms.com/censusplus.php>




Fig. 1: *Average daily population*: This figure shows the average daily population of the WoW server we measured. Players typically play more in the evenings and both earlier and later on the weekdays. Tuesdays are “patch days”, when server maintenance occurs, explaining the empty server at that time.

through traffic and log analysis; our measurements were taken from within the games themselves. Second, and more importantly, we provide measurements and models for player distributions and movements within the virtual space of both games.

4.1 Daily Populations and Arrival/Departure Rates

In our first set of measurements, we examined population fluctuations over time and arrival and departure rates. Due to space limitations, we only show the measurements on WoW, however we observed similar measurements in WAR. Note that WoW measurements include both factions within the game.

Population Over Time: We measured the total number of players in the game every 15 minutes during our measurement period and averaged the results by hour each day of the week. Our hypothesis was that more players were online during evening and weekend hours, due to weekly obligations such as work and school and therefore architectures would need to address these cycles. Figure 1 shows the 24 hour daily cycle with each line representing one of the days of the week.

We note three important aspects of our graphs. First, populations have an average peak at almost 3600 players. Given the imprecision incurred by the measurement method, we estimate that a typical World of Warcraft server will support up to 4000 players. Second, we see that weekend play stands out from weekday play in that the realm experiences a significantly higher average population earlier in the day. This implies that servers must be provisioned for weekend play. Last, we see an almost 5-fold increase in the number of players from the lowest point (at 4AM) to the highest point (7PM) of the realm population. This implies that servers must also be over-provisioned to handle peak loads during the evenings and are only partially loaded during the early mornings.

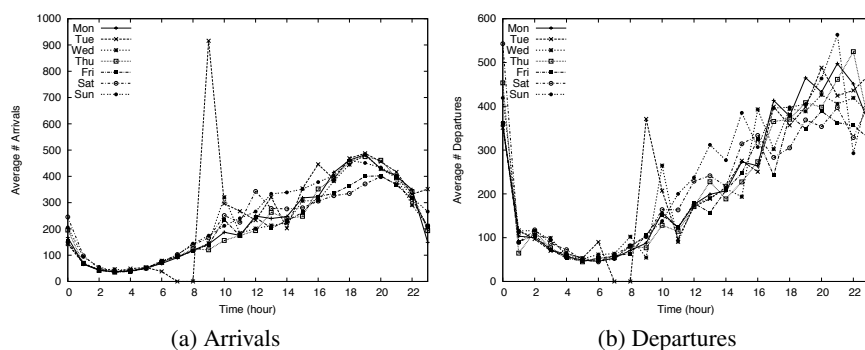


Fig. 2: *Arrival and Departure Rates*: Figure (a) shows the arrival rate in terms of the number of new players seen this hour. Figure (b) shows the departure rate in terms of the number of players seen in the prior hour that are no longer online.

Arrival and Departure Rates: To further understand the population fluctuations and to help understand the amount of churn that occurs in an MMORPG, we measured the number of arrivals and number of departures per hour and averaged this again on each day of the week. Figures 2 show these results.

In these two figures, we see that the amount of *churn*, or the number of players joining and leaving the game, is high during peak playing times. Figure 2 (a) shows similar trends of arrivals during the weekdays, but has an increased arrival rate on the weekends during earlier hours of the day. Figure 2 (b) shows that the number of departures increases towards the end of the day. Together, *we see that during peak playing times, over 1,000 players join and leave the game per hour.*

In terms of the magnitude of the difference between minimum and maximum arrival and departure rates, these results show that arrival rates and departure rates differ by a factor of 10 while. In terms of MMO architectures, 1,000 players joining and leaving per hour may not appear to be a huge burden. However, given that WoW claims to have over 10 million subscriptions, a theoretical maximum of 4,000 players per realm indicates that approximately 1 million players log on and off *per hour* of the WoW servers. This is a significant amount of churn that an MMORPG architecture would need to handle.

4.2 Session Lengths

Session lengths were measured by adding a random subset of players seen in the most recent snapshot to the *friends list*, allowing us to track how long a character is played in the game. Our measurements in Figure 3 show that contrary to anecdotal stories, most sessions were short lived.

Figure 3 (a) shows the CDF calculated from all observed session times in WoW. From this figure, we see that only a small percentage of players we observed played for longer than 400 minutes (8 hours), while most players played for less than 200 minutes (3 hours). We calculated the mean session time to be 80 minutes, with a maximum

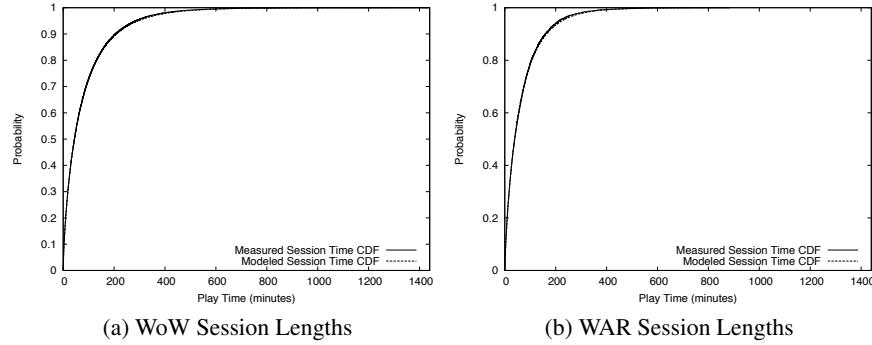


Fig. 3: *Session Times Observed*: Figure (a) shows the CDF of all the sessions we recorded in WoW. In addition, the data was fit to a Weibull distribution, and plotted on the same graph. Figure (b) shows the CDF of all the sessions recorded in WAR along with the model created by fitting the data to a Weibull distribution. Both lines in both figures are barely distinguishable due to the close fit of the models.

observed session time of 1440 minutes (24 hours) and a minimum session time of 1 second. While we did not track players for longer than 24 hours, we will consider how many players were online for extensive periods of time in future work.

Figure 3 (b) shows the CDF calculated from all sessions observed in WAR. In WAR, almost 90% of all sessions were less than 200 minutes with a mean session time of 89 minutes, a maximum time of 882 minutes (14 hours), and a minimum session time of 1 second.

For both models, we used the least-squares method to find a fit for a model of the measured data. Given the trend of the data, we determined that a Weibull distribution would fit well. The models are plotted in Figure 3 (a) and (b). Next, we validated our models by plotting the residuals between the predicted and measured values (not shown due to space limitations), and found that the residuals for both models had a standard deviation of 0.004, indicating an extremely close fit. Thus, given a uniformly distributed random variable $0 \leq x \leq 1$, the session lengths for WoW and WAR can be modeled as follows:

$$Session_{WoW}[x] = 1 - e^{-(x/69.75)^{0.7522}}, Session_{WAR}[x] = 1 - e^{-(x/59.81)^{0.8322}}$$

This result verifies that MMORPGs experience considerable churn. A large fraction of sessions are short lived while only a small fraction are stable. We believe that what may be happening here is that players may be logging on to check to see if friends or guild members are currently online, checking in-game mail, or checking auctions at the auction house. If this is true, then the implication is that load on an architecture could be reduced by providing an external interface to these services that does not require logging into the game. *Given the predictability of player session time, one may conclude that game developers should target playing experiences for session times that reach the*

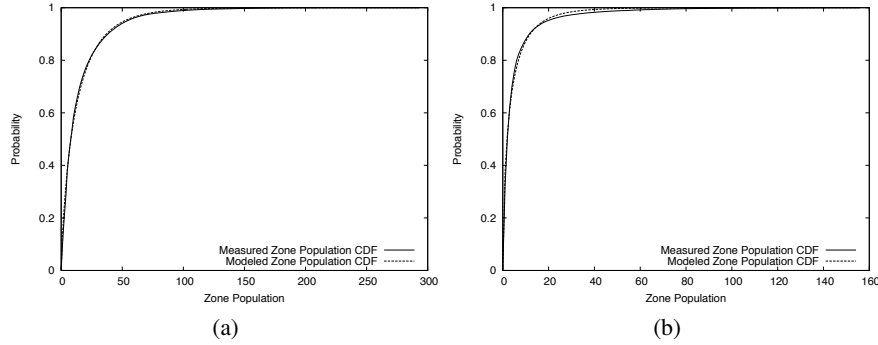


Fig. 4: *Distribution of players per zone*: Figure (a) shows the distribution of players per zone in WoW while Figure(b) shows the distribution of players per zone in WAR. Note that both CDFs do not include the zones without any players: WoW had 36% of the zones empty while WAR had 78% of the zones empty.

majority of players. Researchers, on the other hand, can use session times to predict how long players will connect to a given architecture.

4.3 Player Distributions

We next measured the distribution of players in the virtual world of both games. Throughout the world regions are statically divided into *zones*. We measure how many players are in each zone over the measurement period. After examining the data, we realized that a large percentage of the zones had 0 players in them. To model this correctly, we calculated the quantity of 0 population zones we examined (36.44% in WoW and 78.56% in WAR) and removed these from the data set for the purpose of modeling the remaining data. We then stretched the remaining points to cover the probability from 0 to 1. Using the least-squared method, we fit the data using a Weibull distribution.

Figure 4 (a) and (b) show the measured data and the fitted Weibull distributions of the zone populations. Note that in both games, only a few zones have more than 50 players, while the majority of zones have fewer than 10 players. For WoW, we saw an average of 121 players in a zone, with a minimum of 0 and a maximum of 293 players. On WAR, we saw an average of 74 players in a zone with a maximum of 156 players and a minimum of 0.

As with session times, player distributions were modeled very closely using a Weibull distribution with a standard deviation of 0.007 for WoW and 0.008 for WAR of the residuals from the measured data and models. Thus, given a uniformly distributed random number $0 \leq p \leq 1$, we can model WoW and WAR population distributions as follows:

$$Population_{WoW}[p] = \begin{cases} 0 & \text{if } p \leq .3644 \\ \lfloor 1 - e^{-(((p-.3644)/.6356)/12.744)^{0.7822}} \rfloor & \text{otherwise} \end{cases}$$

$$Population_{WAR}[p] = \begin{cases} 0 & \text{if } p \leq .7856 \\ \lfloor 1 - e^{-((p-.7856)/.2144)/3.256} \rfloor^{0.6417} & \text{otherwise} \end{cases}$$

The measurements of player distributions are important because they show that players are *not* uniformly distributed in the virtual world as much of the prior research in scalable game architectures has assumed. Clearly, given a uniform distribution of players, almost any architecture can be reasonably well-balanced so that it scales well. However, a Weibull distribution indicates that players tend to group in large numbers in only a few zones, causing stress on any architecture as it has to handle the increased number of interactions between players. *Therefore, game designers and researchers must consider this Weibull distribution of players in which a few zones contain a large number of players while many zones only have a few (or no) players when characterizing the potential load on a MMO architecture.*

4.4 Player Movements

To model player movements, we measured the number of zones visited in a session, how long they remained in a zone, and what zones they chose to travel to from their current zone.

Number of Zones Visited: We hypothesized that a linear relationship exists between the number of zones visited during a session and the session length. To test this hypothesis, we measured how many zones the players travelled to each session in both WoW and WAR and plotted the results in Figure 5 (a) and (b). The number of zones visited are not unique zones, but the total number of times a player moved from one zone to another. From these figures, we see that our hypothesis held for the 80% of session times in both games, i.e., those 200 minutes and below in WoW and those 100 minutes and below in WAR. On both graphs, the hypothesis no longer seemed to hold for the highest 20% of the sessions. One explanation may be that players who are on for long periods of time behave differently in the game than those on for shorter periods. Note that the game will disconnect players who remain idle for longer than 10 minutes. Thus, even these long sessions consist of active players or bots.

In both cases, we model this behavior using a simple linear equation. For WoW, we found that the equation $y = 0.070x + 0.831$ works well while for WAR we found that the line at $y = 0.014x + 1.20$ works well. For future work, we plan on exploring how the longer sessions can be modeled more accurately.

Time in Zones: We then observed the distribution of time that players spend in any given zone. This information is important because it helps us understand whether players spend an even amount of time in each zone or perhaps spend only a small amount of time in a majority of zones but a large amount of time in one or two zones. Figure 6 (a) and (b) shows the measured and modeled CDFs of the time players spend in a zone in both WoW and WAR. As with session times, the time in each zone also followed a Weibull distribution, indicating that players did in fact spend most of there time in

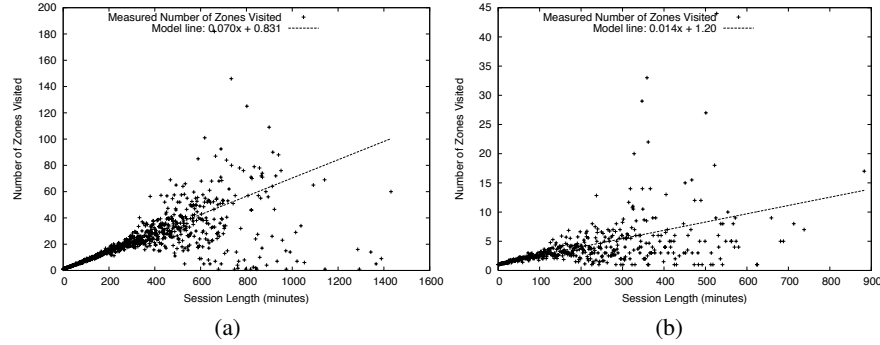


Fig. 5: *Zones Visited vs. Session Lengths*: These figure shows the number of zones visited plotted against the session time in minutes. (a) shows the results from WoW while (b) shows the results from WAR. Simple linear equations are used in both figures to model the data.

a few zones, and a small amount of time in the rest of the zones they visited. These distributions are as follows:

$$ZoneTime_{WoW}[p] = 1 - e^{-(x/8.189)^{0.5674}}, ZoneTime_{WAR}[p] = 1 - e^{-(x/24.42)^{0.6669}}$$

As with the previous models, we measured the residuals between the measured and modeled data and found that the standard deviation of the WoW function was 0.009 and for WAR it was 0.002.

Player Movement: The final aspect we measured with regards to player movement was *how* players moved between zones. Our hypothesis was that the random waypoint model of player movement is not accurate for MMORPGs. Our results indicate that instead a log-normal distribution of waypoint choices is more accurate.

To model this type of player movement, we examined the zone locations of all of the players we tracked and recorded where they were from one measurement to the next. We created a Markov chain of zone to zone player movement from this data. A Markov chain is represented by a two dimensional matrix where each row represents the probability of a transition from the row header (or zone in this case) to a given column header (a destination zone). To visualize the matrix, we created a square image where each pixel represents a cell in the matrix and is colored gray according to its probability in the table, with black being a probability of 1 and white a probability of 0. Figure 7 (a) and (b) shows the result of this visualization.

In order to create a player movement model, we also wanted to be able to randomly generate Markov chains that had the same properties as the measured Markov chains. We found that a log-normal CDF worked well in modeling these probabilities. The results are seen in Figure 7 (c) and (d). The modeled CDFs are defined as follows, where given a uniformly distributed number $0 \leq p \leq 1$ and where *erf* is the Gauss error




Fig. 6: *Time Spent in a Zone*: Figure (a) shows the measured and modeled CDFs of the time spent in a zone in WoW. Figure(b) shows the measured and modeled CDF time spent per zone in WAR.

function:

$$Choice_{WoW}[p] = \frac{1}{2} \left(1 + \operatorname{erf} \left(\frac{15.45 + \log(p)}{6.399\sqrt{2}} \right) \right)$$

$$Choice_{WAR}[p] = \frac{1}{2} \left(1 + \operatorname{erf} \left(\frac{12.23 + \log(p)}{5.235\sqrt{2}} \right) \right)$$

Note that in using this function to create the Markov table would require each line to be adjusted so that it summed to 1, though this should be trivial in practice.

5 Conclusion and Future Work

In this paper we have provided measurement results from and empirically-based models for two popular MMORPGs. Further, by measuring two MMORPGs with significantly different play styles, we show that the models are not unique to a single game, but can be modified slightly to apply to MMORPGs even with different play styles. Given that each distribution used has one or more parameters, these parameters can be altered to create a variety of scenarios to test an MMORPG architecture under. For example, by changing the parameters of the Weibull distributions, one can stretch the function to cover more values or cause its initial rise to cover a larger percentage of the probability space. However, the values provided here are excellent starting points as they model two of the most popular MMORPGs today. As future work, we plan to continue measuring other MMOs to compare the results with our models.

References

1. Chambers, C., Feng, W., Sahu, S., Saha, D.: Measurement-based characterization of a collection of online games. In: Proceedings of the Internet Measurement Conference (2005)

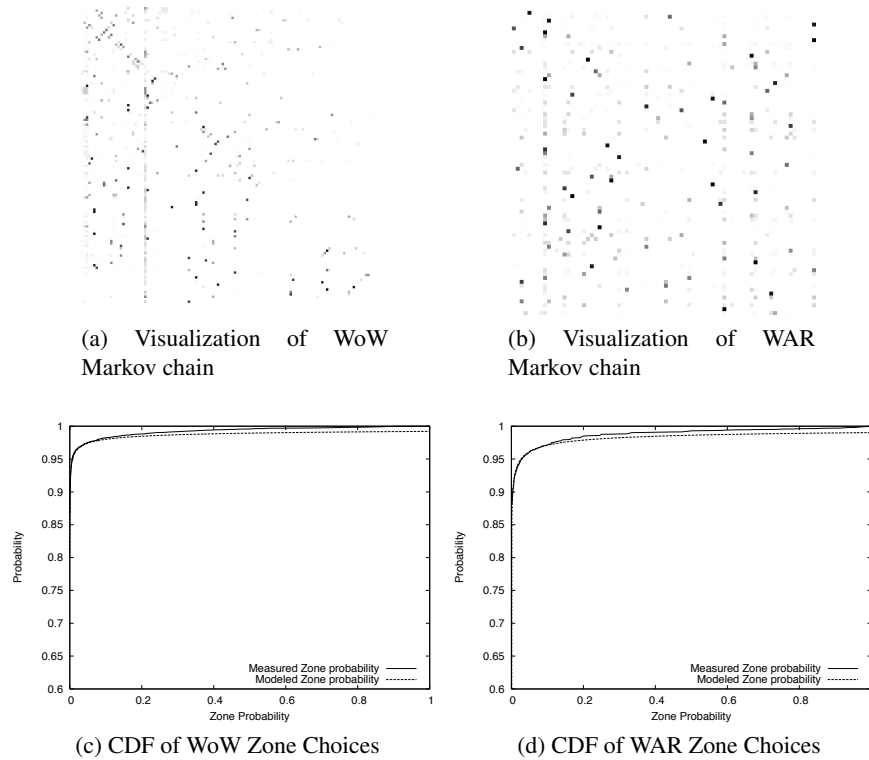


Fig. 7: *Visualization and CDF of Markov Chain Probabilities*: Figures (a) and (b) are visualizations of the Markov chains generated from the measured player movements between zones in WoW and WAR. Figures (c) and (d) shows the modeled log-normal CDFs of the probabilities from the measured Markov chains.

2. Chen, K.T., Huang, P., Huang, C.Y., Lei, C.L.: Game traffic analysis: an mmorpg perspective. In: NOSSDAV '05: Proceedings of the international workshop on Network and operating systems support for digital audio and video. pp. 19–24. ACM, New York, NY, USA (June 2005)
3. Claypool, M., Claypool, K.: Latency and player actions in online games. *Communications of the ACM* 49(11), 40–45 (2006)
4. Fernandes, S., Kamienski, C., Sadokn, D., Moreira, J., Antonello, R.: Traffic analysis beyond this world: the case of second life. In: ACM NOSSDAV (June 2007)
5. Kim, J., Choi, J., Chang, D., Kwon, T., Choi, Y., Yuk, E.: Traffic Characteristics of a Massively Multiplayer Online Role Playing Game. In: Proceedings of ACM NetGames (2005)
6. Kinicki, J., Claypool, M.: Traffic analysis of avatars in second life. In: NOSSDAV '08: Proceedings of the 18th International Workshop on Network and Operating Systems Support for Digital Audio and Video. pp. 69–74. ACM, New York, NY, USA (May 2008)
7. Pittman, D., GauthierDickey, C.: A measurement study of virtual populations in massively multiplayer online games. In: Proceedings of ACM NetGames (September 2007)

8. Suznjevic, M., Dobrijevic, O., Matijasevic, M.: Mmorpg player actions: Network performance, session patterns and latency requirements analysis. *Multimedia Tools and Applications* (May 2009), <http://dx.doi.org/10.1007/s11042-009-0300-1>
9. Svoboda, P., Karner, W., Rupp, M.: Traffic analysis and modeling for world of warcraft. In: *Communications, 2007. ICC '07. IEEE International Conference on Communications*. pp. 1612–1617 (June 2007)
10. Tarng, P.Y., Chen, K.T., Huang, P.: An analysis of wow players' game hours. In: *NetGames '08: Proceedings of the 7th ACM SIGCOMM Workshop on Network and System Support for Games*. pp. 47–52. ACM, New York, NY, USA (October 2008)
11. Ye, M., Cheng, L.: System-performance modeling for massively multiplayer online role-playing games. *IBM Systems Journal* 45(1), 45–58 (January 2006)