A Long-Term Study of a Popular MMORPG

Wu-chang Feng
Portland State University

David Brandt
CCP Games

Debanjan Saha
IBM Research

ABSTRACT

Over the last decade, Massively Multi-player On-line Role Playing Games (MMORPGs), have become big business. In typical MMORPGs, players pay a monthly subscription to the game publisher who hosts the game and provides periodic content updates. To be successful, game publishers must characterize their player population so that they can provision sufficient resources to support the game and so that they can update the game in a timely manner. To this end, this paper provides the first, long-term study of a popular MMORPG from its launch. Our dataset encompasses the entire lifetime of "EVE Online" [1, 2], a science-fiction based MMORPG that has supported nearly 1 million unique players, 67 million player sessions, and 17,000 player years of gameplay since its launch in May 2003.

1. INTRODUCTION

On-line games, and particularly massively multi-player on-line role playing games (MMORPGs), have seen significant growth over the last decade. Although a large number of MMORPGs currently exist, the closed nature of the application has limited measurement studies to a handful of short-term traffic analyses at the client and server [3, 4, 5]. For a MMORPG to be successful, it is important for the game publisher that is providing the content and hosting infrastructure to understand long-term player workloads and behavior. To address this problem, this paper examines the complete session history of a popular, science-fiction based MMORPG called "EVE Online".

The study of MMORPG usage is typically limited due to the proprietary nature of the industry and the need to ensure player privacy. With the help from CCP, the publisher of EVE Online, the dataset obtained includes the complete session log for EVE Online from its launch. Specifically, the logs contain start and end times of each session along with an anonymized user number. Table 1 shows the characteristics of the dataset. Using this dataset, our study targets two key areas important to game publishers:

`EVE Online` trace	
Start time	Tue May 6 2003
End time	Sun Mar 12 2006
Total sessions	67,060,901
Total unique players	925,928
Total player time	17,204 years

Table 1: Data set

- *How predictable is the MMORPG workload?* For game publishers hosting a game, it is important to be able to predict a game's workload in order to provision sufficient server and network resources. Previous work has shown that the workload for a number of on-line games is highly periodic and predictable in the short-term while unpredictable over longer periods of time [6]. Unfortunately, the dataset used in [6] includes mostly games of the first-person shooter (FPS) genre and does not include a highly successful subscription MMORPG such as EVE Online.
- *How predictable are individual players?*: A game publisher must characterize its players to understand their playing behavior. By understanding factors that contribute to players joining the game, players continuing their subscriptions, and players leaving the game, the publisher can provide in-game incentives and game updates to keep the player active.

2. WORKLOAD PREDICTABILITY

Game publishers must provide players with an appropriate amount of resources to play the game. In performing this task, it is ideal if the workload is fairly predictable. Previous work studied the problem using "public-server" first-person shooter games [6] where the infrastructure for playing the game is actually provided by the players themselves. Using our dataset, we examine this property in the context of a MMORPG hosted by the game publisher.

2.1 Predictable daily and weekly changes

Figure 1(a) shows the player load over four consecutive weeks in March 2004 for EVE Online. As the figure shows,

NetGames '07, September 19–20, 2007, Melbourne, Australia.

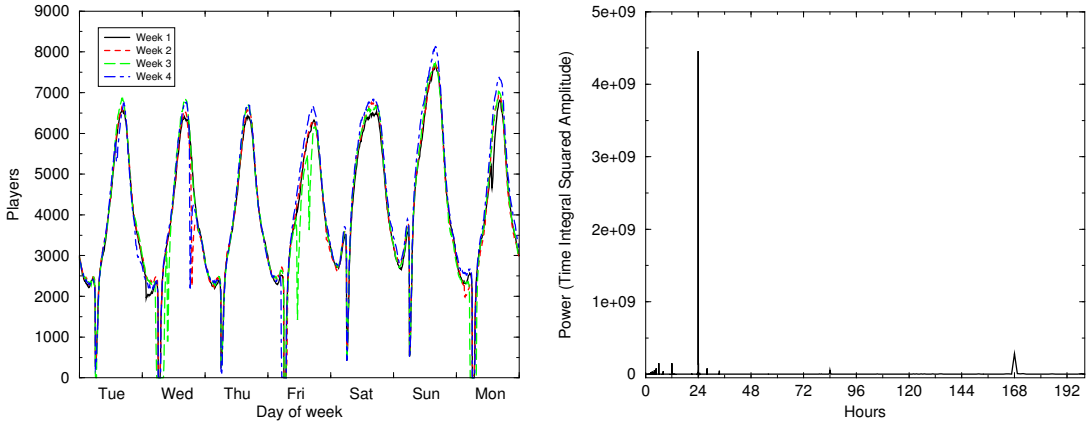

(a) Weekly player load (March 2004)

(b) FFT of player load

Figure 1: Player load predictability and periodicity.

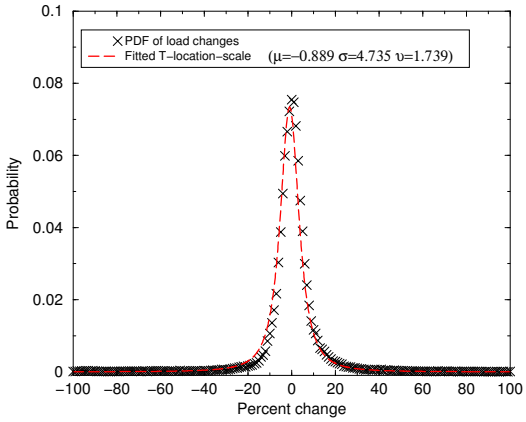

Figure 2: Weekly load variation.

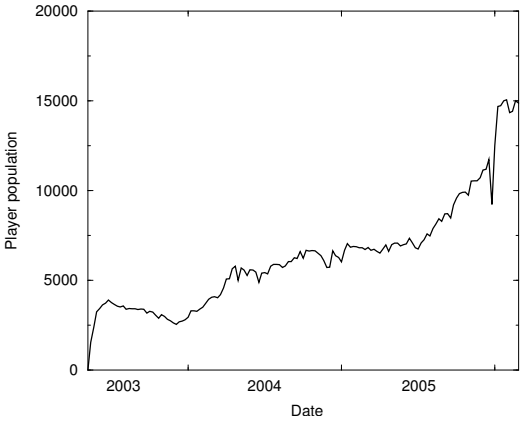

Figure 3: Average weekly player population.

the workload has strong daily fluctuations that are similar to those seen in other on-line applications. To further quantify this, Figure 1(b) plots the Fast Fourier Transform of player load over the entire trace in order to show the frequency components that drive the EVE Online workload. The figure shows that along with a strong daily peak, the workload contains a smaller weekly one. This peak corresponds to the increase in player load on the weekends.

In addition to periodic load changes, Figure 1(a) also shows that there is very little variation in player load on a week-to-week basis. Figure 2 shows that this holds across the entire trace by plotting the PDF of percent load changes from week-to-week over the duration of the trace. As the figure shows, most of the week-to-week variation is under 10%, making the short-term provisioning problem simple. The shape of the distribution follows a 't' location-scale distribution with almost identical parameters as that fitted for the GameSpy collection of on-line games [6].

2.2 Unpredictable long-term fluctuations

By virtue of having the entire lifetime of EVE Online, we can examine the predictability of player load across long time periods. Figure 3 plots the average population of active players for each week the game has been on-line. A player is considered "active" if he or she has played the game at least once in the last month. As the figure shows, the active player population has increased irregularly throughout the game's lifetime in an unpredictable pattern.

3. PLAYER POPULATION GROWTH

While predicting the overall workload is important for provisioning resources, the ultimate goal for a game publisher is to increase the overall number of players subscribed to the game. This means ensuring that the rate at which players "join" the game always is larger than the rate at which players "quit" for good.

3.1 Joining and quitting are correlated

In order to study player growth, we measured the number of players playing for their first time and the number of players playing for their final time across the trace. At the end of the trace, any player that has played in the last month

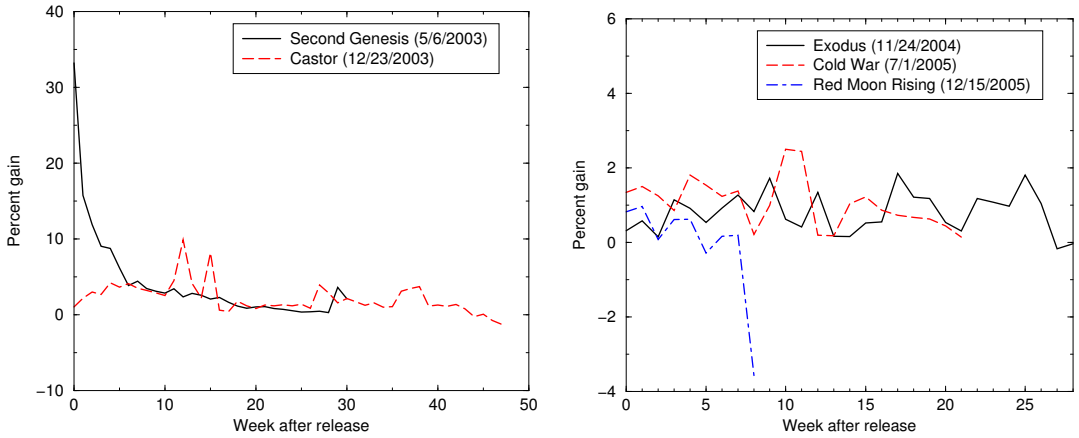

Figure 5: Weekly percent gain in players across 5 software updates.

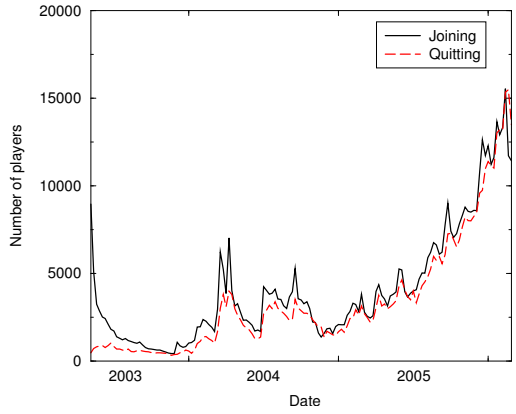

Figure 4: Weekly rates of joining and quitting.

of the trace is not included in the count for players who have "quit" the game. Thus, a player who played only in the first month and the last month of the trace has "joined" the game, but has not "quit" yet. Figure 4 plots the weekly rate of players joining and quitting. As the figure shows, aside from the game's launch, the two rates follow each other closely. This indicates that most players try the game out for a short time before quitting because the game does not interest them. The low quitting rate at the beginning of the trace can be attributed to the fact that the majority of those players played the beta version of the game and already knew they liked the game. The figure also shows several large growth spurts early in 2004 that are a result of a promotion targeting players of a competing sci-fi MMORPG "Earth and Beyond" [7] which was being closed [1].

3.2 Updates slightly impact player growth

One of the major controls a game publisher has on subscription growth is the ability to provide fresh and compelling content. To this end, the game is often updated to keep players interested. Ideally, such updates would have direct impact on the rates of joining and quitting the game. Figure 5 examines this by plotting the percentage gain in active players for the first five software updates for EVE Online each week after the update was released. As the figure shows, aside from the release of the initial game (Second Genesis), new software updates appear to have a slight impact on overall subscriber growth. The most significant spikes in subscriber growth occur 10 weeks after the second software update (Castor) and are the result of the Earth and Beyond promotions.

3.3 Updates slightly impact player minutes

Although software updates may not have a large impact on overall player subscriptions, they may impact the amount of time an individual player spends on the game. This is important when considering that some games are examining "in-game" advertising as a means of revenue. By increasing "eye-ball" time for such advertisements, game updates might drive such revenue. Figure 6 shows the average minutes played per player on a weekly basis for each of the first five updates to the game. As the figure shows, the new content slightly increases the average playing time per player. The limited impact of the new content on play time can potentially be attributed to the fact that EVE Online players already spend more than 20 hours a week playing the game and that lack of new content may not be the reason they spend less time with the game.

3.4 Player churn increases over time

One of the problems that MMORPGs have is the ability to attract and retain new players. There are several potential reasons for this. Due to the persistent nature of such games, one reason is that new players often come in at a severe disadvantage to those who have played from the game's launch. Since old players have acquired significant virtual wealth and power, the disparity in abilities often discourages new players from continuing to play. Another reason is

[1] On 3/17/2004, EA announced the closure of Earth and Beyond effective 9/2004. CCP's promotion began on 3/29/2004.

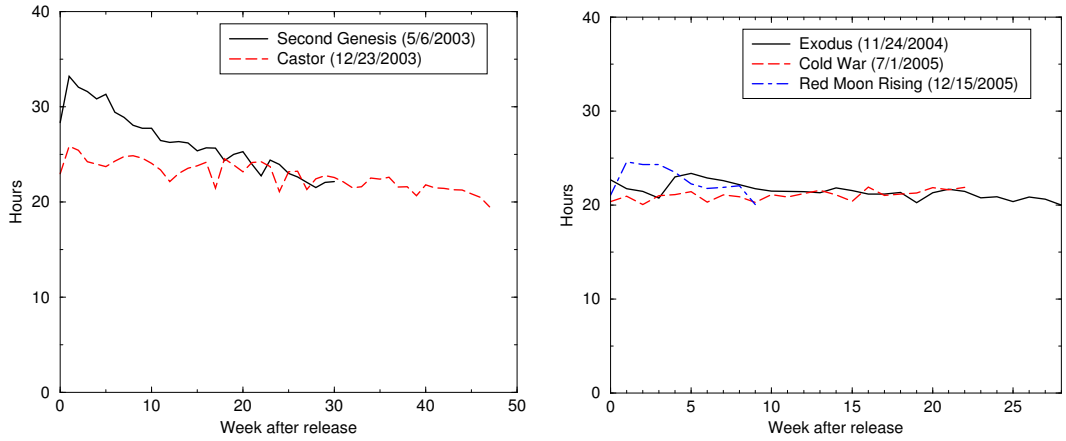

Figure 6: Weekly minutes played per player across 5 software updates.

that serious fans of the particular game genre join as soon as the game is released leaving a population of less enthusiastic players to join later. In order to quantify this, Figure 7 plots the retention rate of players for 4 different months across the first two years of EVE Online. For each month, the figure plots the percentage of players who joined in that month who are still playing the game a certain time into the future. For example, the top curve shows that more than 70% of the players who joined during the month of 6/2003 were still playing the game the following month and that towards the end of 2005, nearly 30% of the players who joined during 6/2003 were still playing the game. The figure shows some startling trends in that as the game has matured, it has become increasingly more difficult to keep new players. After starting with extremely high retention rates, the one-month retention rate for EVE Online decreases steadily to around 25% after two years. While the initial high retention rates at the beginning of the game can be attributed to the fact that the beta version of the game has "weeded" out many uninterested players, the steady decrease in retention rates over time appears to indicate that many new players become disinterested due to the imbalance of power. Figure 7 also shows that after the initial months, player retention rates decay similarly regardless of when players joined. This indicates that after the first several months, the "quitting" rate for players is roughly the same regardless of when they started playing.

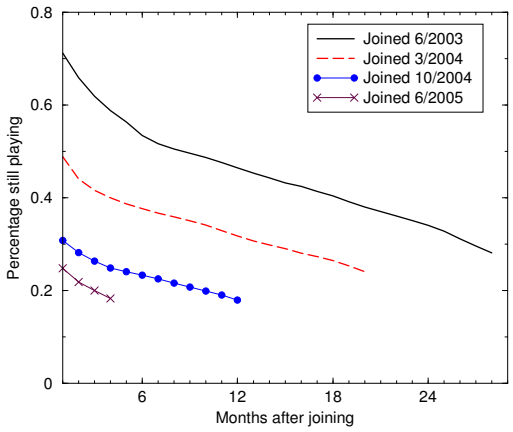

Figure 7: Player retention over time for 4 different months.

4. INDIVIDUAL CHARACTERISTICS

While the previous sections examined aggregate player characteristics, it is also useful to examine individual player behavior as well.

4.1 Player sessions reveal player disinterest

Figure 8 shows the PDF of the number of sessions that players play before quitting the EVE Online. As the figure shows, a large number of players only play the game a handful of times before quitting. The distribution fits a Weibull distribution closely.

Figure 9(a) shows the PDF of player session times through-

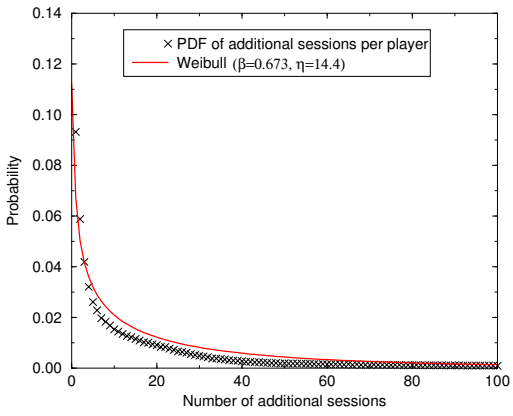

Figure 8: Distribution of number of sessions per player.

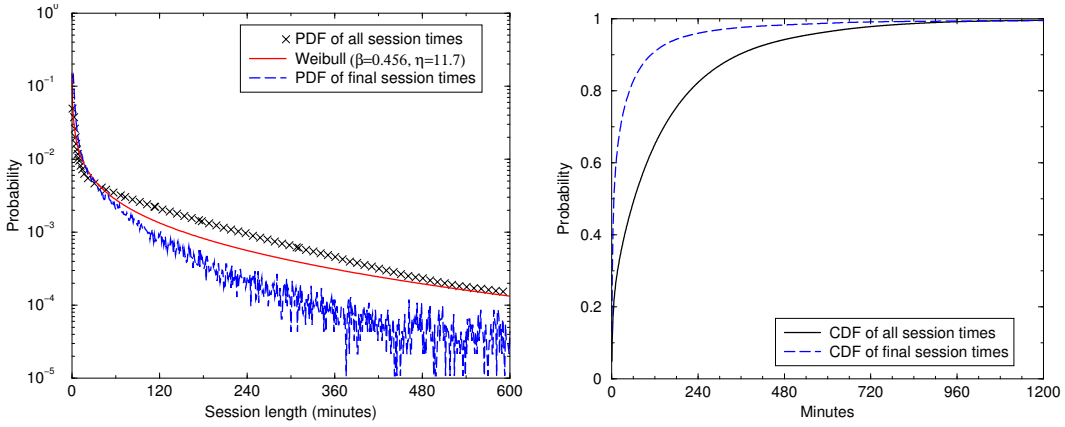

(a) PDF (b) CDF

Figure 9: Player session time and final session time distributions.

out the trace. Overall, many players don't stay for long. However, the distribution does have a long-tail with a small fraction of players that play for an extremely long time. The distribution itself can be fit with a Weibull distribution similar to that used in previous studies [6]. It is ideal if a game publisher can detect players that are about to quit by studying their session times. To this end, the figure also shows the PDF of "final" session times across the entire trace. Final session times refer to the length of time a player spent on the game during his/her last session before quitting the game for good. As the figure shows, the final session for players quitting the game is much shorter than normal and thus represents a useful signal for game publishers to catch players about to quit. Figure 9(b) shows the difference more clearly by plotting the CDF of both.

Similar to session times, intersession times can be a good predictor of player behavior. Figure 10(a) shows PDF of intersession arrivals across the entire trace. As the figure shows, a large majority of players return to play within several days of their previous session. The distribution shows periodic daily spikes indicating playing times are scheduled based on the time of day. The figure also shows the PDF of "final" intersession arrival times for players that have quit the game for good. As with final session times, final intersession times are typically much larger than normal. Figure 10(b) shows this more clearly by plotting the CDF of intersession times overall and of final intersession times.

4.2 Individual history reveals disinterest

The above distributions indicate that aggregate session times decrease and aggregate intersession times increase for players that are quitting the game altogether. Ideally, a game publisher would be able to detect such characteristics on an individual basis and deliver incentives or software updates in order to keep the player interested in playing. To this end, one metric to examine is the percentile of a player's final session and final intersession times with respect to his/her prior times. That is, what percentile does a player's final session and intersession times fall into when compared to the player's previous session and intersession

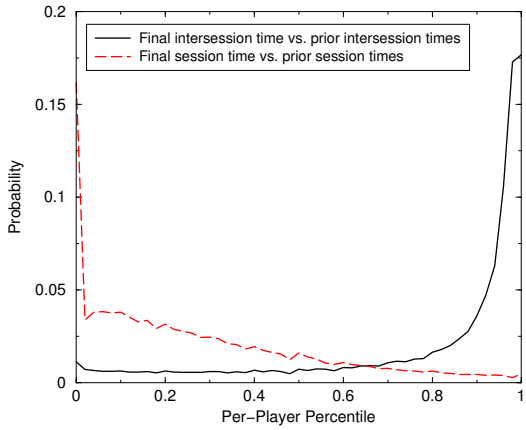

Figure 11: "Final" session and intersession time percentiles.

times? Figure 11 plots the percentile statistics of the final session and intersession times of all players that have quit the game. As the figure shows, most of the final session times are well below the 50th percentile of the player's previous session times while most of the final intersession times are well above the 50th percentile of the player's previous intersession times. From the graph, it appears that final intersession times are heavily weighted towards one end, making it a better indicator for detecting players that are quitting.

Related to session and intersession times is the amount of time a player spends on the game over time. Figure 12 shows the average time individual players spend during each week as a function of their "play history" [6]. Using this measure, a player playing the game for the first time corresponds to the 0% point and a player playing the game for the final time corresponds to the 100% point of their play history. As the figure shows, players steadily decrease their usage before eventually quitting the game, giving game publishers another useful metric for determining when to deliver incentives to individual players.

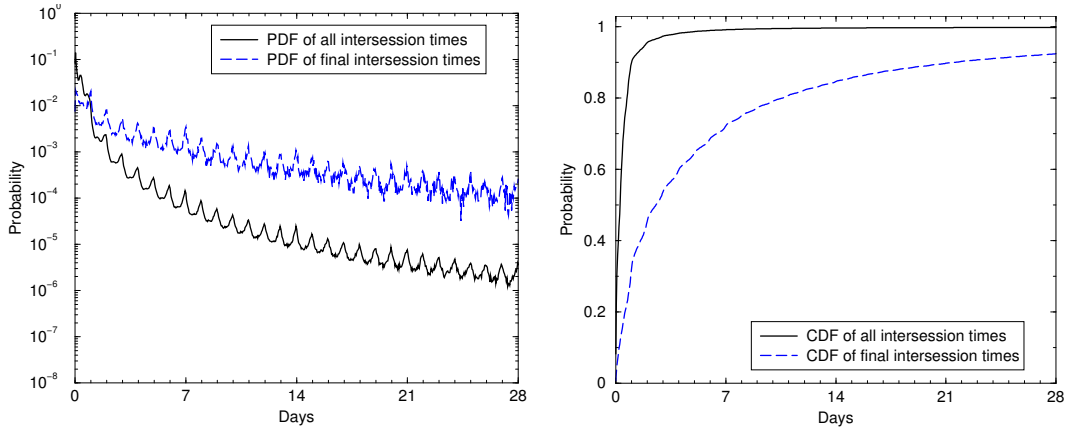

Figure 10: Player intersession time and final intersession time distributions.

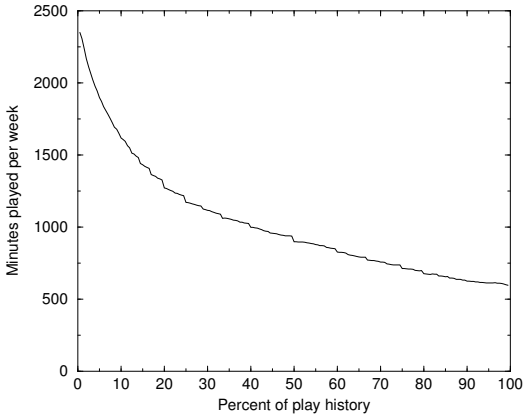

Figure 12: Minutes per player as a function of lifetime.

5. CONCLUSION

This paper has provided the first, long-term analysis of a popular MMORPG. Our results show that (1) its workload is highly predictable in the short-term, (2) content updates have only a slight impact on subscription growth and player usage, (3) player churn increases as the game matures, and (4) intersession times provide a reasonable metric for identifying players that are about to quit.

6. ACKNOWLEDGMENTS

This material is supported by an IBM Faculty Partnership Award as well as the generous donations of Intel Corporation and CCP Games, Inc.

7. REFERENCES

[1] "EVE Online," `http://eve-online.com/`.

[2] CCP, "CCP Games," `http://ccpgames.com/`.

[3] K. Chen, P. Huang, C. Huang, and C. Lei, "Game Traffic Analysis: An MMORPG Perspective," in *NOSSDAV*, June 2005.

[4] J. Kim, J. Choi, D. Chang, T. Kwon, Y. Choi, and E. Yuk, "Traffic Characteristics of a Massively Multiplayer Online Role Playing Game and Its Implications," in *NetGames*, October 2005.

[5] T. Fritsch, H. Ritter, and J. Schiller, "The Effect of Latency and Network Limitations on MMORPGs - (A Field Study of EverQuest II)," in *NetGames*, October 2005.

[6] C. Chambers, W. Feng, S. Sahu, and D. Saha, "Measurement-based Characterization of a Collection of On-line Games," in *Internet Measurement Conference*, October 2005.

[7] "Earth and Beyond Trial Program," `http://eve-online.com/enb`.