Characterizing On-line Games

Chris Chambers, Wu-chang Feng, *Member, IEEE,* Sambit Sahu, *Member, IEEE,* Debanjan Saha, *Member, IEEE,* David Brandt

***Abstract*—On-line games are a rapidly growing Internet application. In order to run a successful on-line game, game companies and game infrastructure providers must properly manage game workloads and content so that they can maximize player satisfaction while minimizing their own costs. Towards this end, this paper provides a comprehensive, long-term analysis of several popular on-line games and their players using one of the richest data sets available for on-line games.**

I. Introduction

On-line gaming is an increasingly popular form of entertainment on the Internet. While a popular on-line game can have millions of subscribers paying a monthly fee, such successes are highly unpredictable and are balanced by the substantial costs in developing and hosting on-line games. As a result, it is important for game and hosting companies to maximize revenue while minimizing costs. Towards this end, this paper presents a first-of-a-kind measurement-based study of game players and game workloads. Unlike previous game measurement studies [1], [2], [3], [4], our data set includes a wide range of on-line game genres including first-person shooters (FPS), massively multiplayer on-line role-playing games (MMORPG), and casual games over a long duration of time. Using this data set, this paper examines the aggregate player populations as well as the individual player behavior of a number of popular on-line games. Our study targets two key areas that are important to game and hosting providers including:

- *Provisioning resources* One of the biggest costs in supporting an on-line game is the hardware and network resources that must be purchased to maintain a reasonable playing experience. Overprovisioning resources can negatively impact the company's bottom line growth by forcing it to sit on idle hardware while underprovisioning resources can negatively impact its top line growth by causing frustrated players to end their subscriptions. To address this problem, we examine issues in provisioning resources for launching and maintaining a game. In addition, we examine the ability for games to share infrastructure with other network applications to reduce hardware costs. Using the collected traces, our results show that the popularity of these games follows a distinct power law distribution making the provisioning of resources at launch-time extremely difficult. However, as games mature, their aggregate populations do become predictable, allowing game and hosting companies to more easily allocate resources to meet demand. Our results also show that usage behavior of interactive applications follows strict, geographically-determined, time-of-day patterns with limited opportunities for resource sharing.
- *Satisfying players*: A successful game must be able to keep its paying players satisfied. To maintain a healthy top line, game publishers must provide timely content updates, reliable uptime, and balanced gameplay to attract new players while keeping old ones. To address this problem, we examine what influences player behavior and satisfaction throughout the lifetime of a game. Using the collected traces, our results show game players can be difficult to satisfy, but that they show measurable disinterest in a game just before quitting it altogether. Such predictability allows game publishers to tailor incentives in order to maintain player subscriptions.

This material is based upon work supported by an IBM Faculty Partnership Award. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of IBM.

GameSpy trace	
Start time	Fri Nov 1 2002
End time	Fri Dec 31 2004
Total games	550
Total player time	337,765 years

Casual games trace	
Start time	Wed Jun 1 2005
End time	Wed Jun 28 2006
Total games	110
Total player time	128,331 years

`cs.mshmro.com` trace	
Start time	Tue Apr 1 2003
End time	Mon May 31 2004
Total connections	2,886,992
Total unique players	493,889
Median session time	27 minutes

`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
Median session time	64 minutes

TABLE I
Data sets

II. Methodology

The study of on-line game usage is typically limited due to the proprietary nature of the industry. To overcome this, we have collected several unique data sets that allow us to analyze properties that have not been possible previously. The data sets collected are shown in Table I.

A. GameSpy

One problem with measuring on-line game usage is the limited access to game server hosting data. Game companies typically keep the access and usage behavior of their players confidential. There are two factors that enable the measurement of aggregate game player populations, however: (1) on-line games use a centralized authentication server to keep track of the players that are playing and (2) information on overall player numbers per game is usually exported publicly. Several game portal services collect such player numbers over a large number of games and report the information in real-time. Among these services is the GameSpy network, which provides real-time player population data on individual games in a structured format that can be readily collected and analyzed [5]. Currently, there are over 550 on-line games that are being tracked across various genres including first-person shooter games (FPS), massively multi-player on-line role-playing games (MMORPG), real-time strategy games (RTS), card and board games, and sports games. The most popular games tracked by the Gamespy network are from the FPS genre however, and therefore when we refer to gamers we are predominately referring to FPS gamers. To study on-line game population behavior, we have collected a data feed from GameSpy for more than two years since November 2002. Our redundant collection facility periodically samples the GameSpy data every 10 minutes. Note that the availability of the data is sensitive to many factors, including service outages at the portal and our own outages. These outages have been manually removed from the data analysis. The trace represents more than 300,000 years of player time spent on games over the course of two years.

B. Casual games

Similar to GameSpy, there are a number of on-line game portals that track casual games. Casual games offer much simpler gameplay and are downloaded from portals directly. To study the behavior of such games, we collected a trace of casual game player population data as exported by MSN Games using the same methodology as the GameSpy dataset. The trace represents more than 128,000 years of player time spent on casual games over an 18-month time period [6].

C. cs.mshmro.com Counter-Strike server

To study the behavior of individual players, we examined two long-term traces of popular on-line games. The first trace is of one of the busiest and longest running Counter-Strike servers in the country located at `cs.mshmro.com` [7], [4]. Counter-Strike, a first-person shooter (FPS), is a Half-Life modification and has one of largest service footprints with 35,000 servers and over 4.5 billion player minutes per month [8]. Of all of the active Counter-Strike servers, `cs.mshmro.com` was among the busiest 20 servers in May 2004 as ranked by ServerSpy [9] with more than 40,000 connections per week and more than 60 player years in activity since its launch in August 2001. The server has 20 available slots for players and was continuously full throughout the duration of the trace. The server has hosted almost half a million unique players within a single year as identified by Steam IDs, unique identifiers bound to each player.

D. EVE Online

We also study player behavior using a trace that consists of the complete session history of "EVE Online", a popular science-fiction MMORPG. As of June 2006, EVE Online was estimated to be the 10th most popular MMORPG [10]. Although a handful of short-term traffic analyses of MMORPGs have been done [2], [1], [11], this is the first long-term study of player behavior throughout the lifetime of a popular MMORPG. The anonymized trace contains the start and end times of each player session and represents over 17,000 player years in activity since its launch in May 2003. The game has hosted more than 900,000 unique players as identified by EVE Online's User ID number.

III. Provisioning Resources

Provisioning resources for games involves supplying just enough resources to support the demands of the current player population. This is difficult due to the dynamic nature of game popularity and game players over time. Accurately predicting game workloads allows game hosting providers to allocate the appropriate amount of resources for a game. In order to determine whether this is feasible, we analyze the aggregate population data from the above traces. Specifically, we investigate whether any simple trends or patterns can be used to accurately predict the game workload, whether the workload is stable and if so, over what time scale.

A. Game popularity follows a power-law

Figure 1 shows the average weekly population of a representative group of games from the dataset.[1] The number is calculated by measuring the instantaneous player population for each game every 10 minutes and averaging them over each week. Figure 1(a) and (b) show the top games from the GameSpy trace, Figure 1(c) shows the top games from the casual games trace, and Figure 1(d) shows EVE Online. As the figure shows, games have population sizes and variations that are diverse, making it difficult to predict their long-term behavior. Such a pattern makes it difficult to provision resources at launch-time of a game and over long periods of time. To determine the distribution of on-line game popularity, the average player population measured from the casual games and GameSpy data sets were used. Figure 2 shows the popularity data for each on a log-log scale. As the figure shows, this distribution is heavily skewed in favor of the most popular games, with the first ranked game having substantially more players than the next most popular. This distribution of popularity is similar to a power-law distribution. Power-law distributions are of the form $f(x) = ax^{\lambda}$ where a and λ are constants. Applied to game popularity, $f(x)$ is then the average player population and x is the rank of the game. Power law distributions occur in a large number of places including

[1] Many games experience a drop during the *Sobig* virus in late 2003 [12].

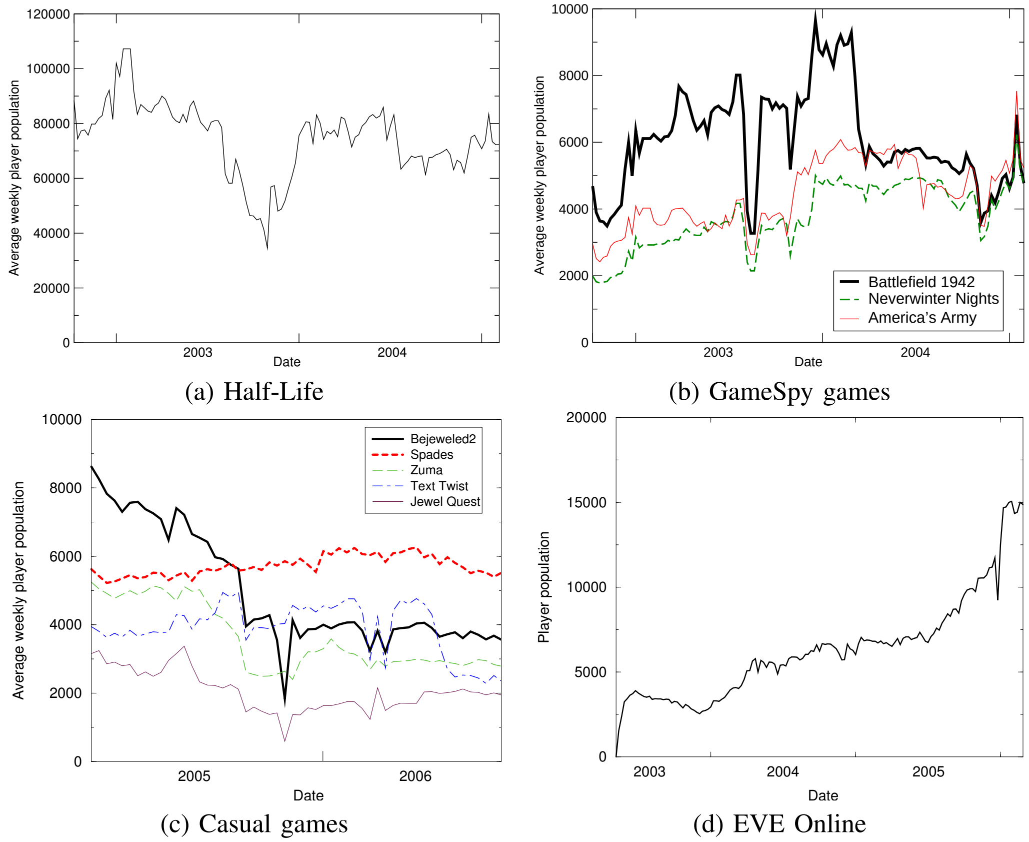

(a) Half-Life
(b) GameSpy games
(c) Casual games
(d) EVE Online

Fig. 1. Average weekly population across several popular games

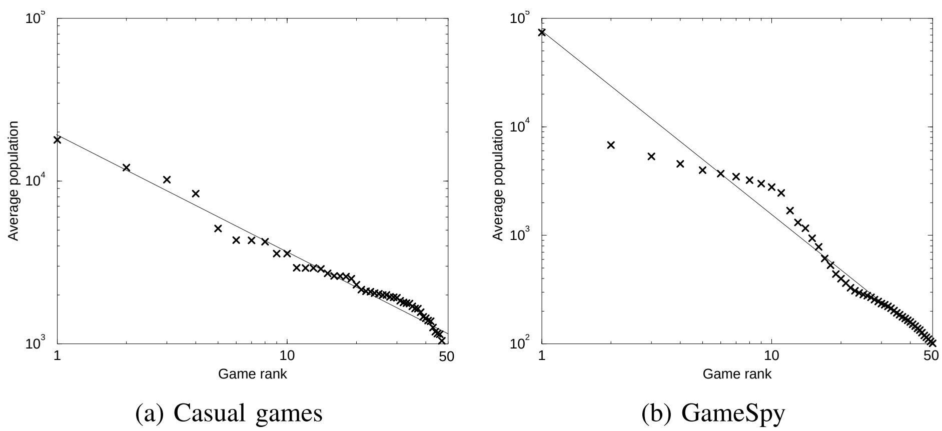

(a) Casual games
(b) GameSpy

Fig. 2. Game popularity for each data set (log scale)

word usage in the English language, web page popularity, and the population of cities. An intuition for these distributions is that whenever choices are made between many options, and each choice affects other choices, the choices tend to pile up on a few popular selections. Games and servers create communities: in selecting one, each player's choice affects and is affected by the popularity and reputation of that game or server. A perfect power-law distribution would graph as a straight line on a logarithmic scale in both the x and y axis. The straight line fits in both plots show qualitatively that the data does follow a power law distribution. This distribution has an unfortunate implication for provisioning server resources for on-line games: the host must plan for several orders of magnitude of change in popularity (and therefore resources) in either direction. As a result, this indicates that on-demand infrastructure such as those proposed by a number of commercial vendors [13], [14], [15], [16], [17] can significantly reduce the costs and risks of provisioning on-line games.

B. *Game workloads are predictable in the short-term*

Intuitively, it is reasonable to assume that game usage is strongly tied to the daily and weekly activities of players. Figure 3 shows the global player population of four consecutive weeks for two popular games from the GameSpy trace

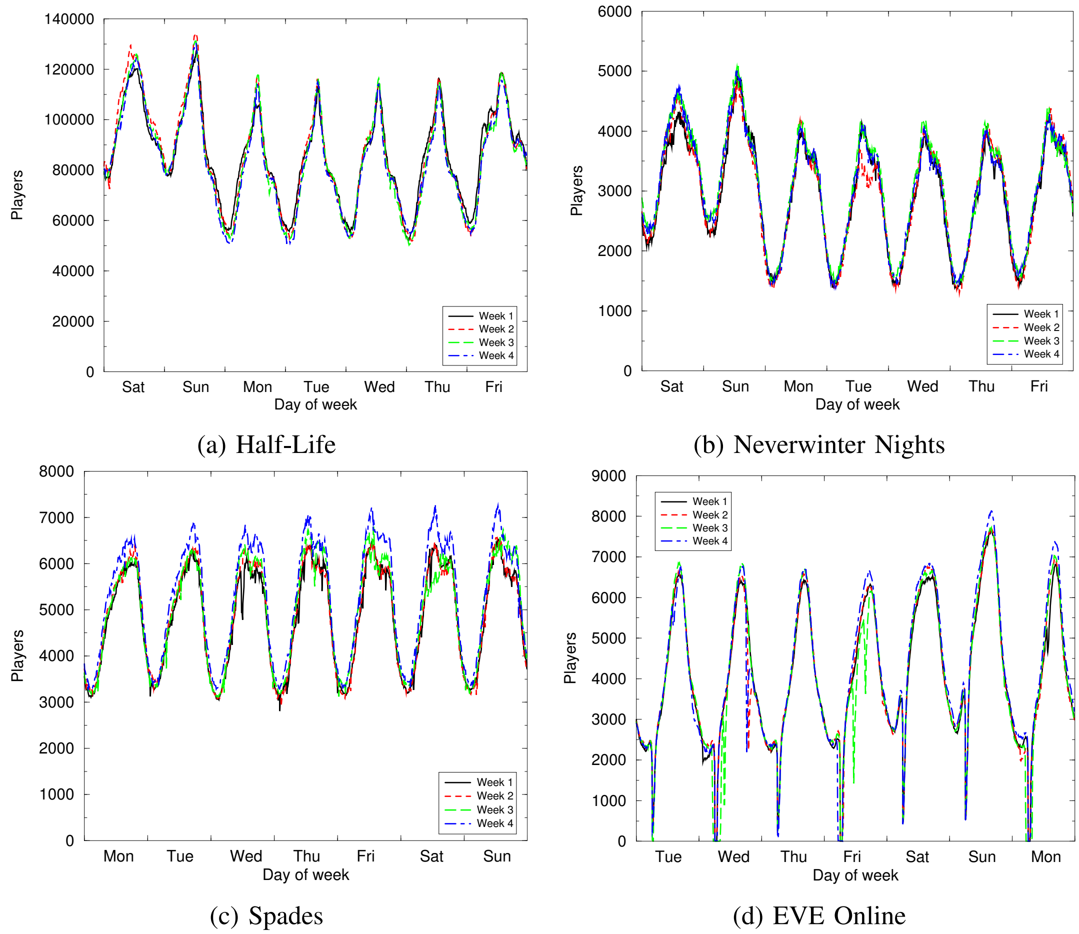

(a) Half-Life

(b) Neverwinter Nights

(c) Spades

(d) EVE Online

Fig. 3. Player load for four popular games over a 4-week period

(Half-Life and Neverwinter Nights), one popular game from the casual game trace (Spades), and EVE Online. Note that the data for the GameSpy games were taken from March 2003, while the data for the casual games and EVE Online were taken from March 2006. As expected, the figure shows that the workload has regular daily cycles and that over this one month period the workload does not vary significantly from week-to-week. In fact, for all games, the trends as well as the maximum and minimum points match up at identical points in time during the week. We observe similar results over other parts of the year with the only anomalies caused by service outages and by holidays. Note that one anomalous characteristic is the daily outages for EVE Online. This is a result of daily scheduled maintenance on the EVE Online cluster. The results show that aggregate short-term game workloads have similar characteristics regardless of game genre (i.e. casual games, FPS games, and MMORPG games).

To further demonstrate the cyclical nature of gaming workloads, we take the a year's worth of game server load samples across the same games and plot the Fast Fourier Transform (FFT) of the data. The FFTs have been scaled so that they can be plotted together. As Figure 4 shows, the FFT contains strong peaks at the 24-hour cycle for each of the games. There is also a smaller peak at the 168-hour (one week) cycle. This corresponds to an increase in player usage on the weekends

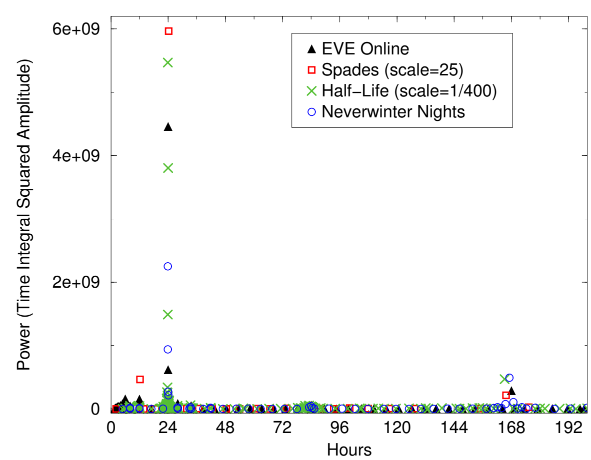

Fig. 4. FFT of player load across four games.

during some parts of the year.

In order to quantify the week-to-week variation of game workloads, Figure 5 shows distribution of week-to-week load changes of the top 5 GameSpy games of 2004 (Half-Life, Battlefield 1942, Medal of Honor: Allied Assault, America's Army, and Neverwinter Nights), the top 5 casual games of 2005 (Bejeweled2, Spades, Zuma, Text Twist, and Jewel

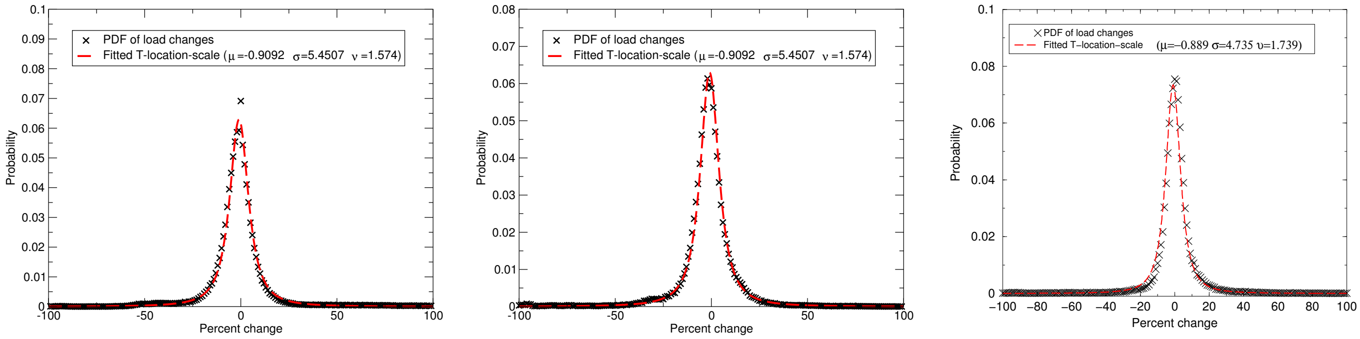

(a) Top 5 GameSpy games (b) Top 5 casual games (c) EVE Online

Fig. 5. Week-to-week PDF of percent load changes

Quest), and EVE Online. Specifically, the figure plots the distribution of instantaneous load changes between identical points in time of consecutive weeks for the games. The distributions closely fit a 't' location-scale distribution with three parameters, a scale parameter $\sigma > 0$, a location parameter μ, and a shape parameter $\nu > 0$. The density function for this distribution is:

$$f(x) = \frac{\Gamma(\frac{\nu+1}{2})}{\sigma\sqrt{\nu\pi}\Gamma(\frac{\nu}{2})}\left(\frac{\nu+(\frac{x-\mu}{\sigma})^2}{\nu}\right)^{-\frac{v+1}{2}}$$

Note that if x is 't' location-scale distributed, $\frac{x-\mu}{\sigma}$ is Student's 't' distributed with ν degrees of freedom. As illustrated in Figure 5, we find a very good fit for all the three plots. Based on this observation, we draw two main conclusions with regard to resource usage:

- As the figures show, most week-to-week load variations are under 15% of the previous week's workload. Specifically, for weekly variations, 86% of the GameSpy trace, 85% of the Casual game trace and 90% of the EVE Online trace are within 15% of the previous week. Such behavior makes it relatively easy for game and infrastructure providers to provision and predict resource usage on a weekly basis.
- The above distribution fitting of load variations indicates that it is feasible to model the week-to-week load variations using standard distributions.

C. Game workloads are synchronized

With the movement toward hosted game services [18], [19] as well as on-demand computing infrastructure for games [13], [14], [15], [16], [17], there has been a great deal of interest in reducing the cost of running game servers by sharing server resources dynamically across multiple games and applications. There are two ways games can be multiplexed with each other. One way would be to coarsely and statically assign physical servers to particular games based on the popularity of the game. As shown earlier, such an approach can provide a lot of benefit for game companies at launch time. Another way to multiplex games with each other would be to dynamically reallocate servers based on instantaneous demand for a particular game. An implicit assumption that gives value to the latter method is that different games have usage patterns that are substantially different. Thus, rather than have each game provision server resources based on their individual peak usage, resources would be provisioned for the global peak.

In order to investigate the extent to which different games

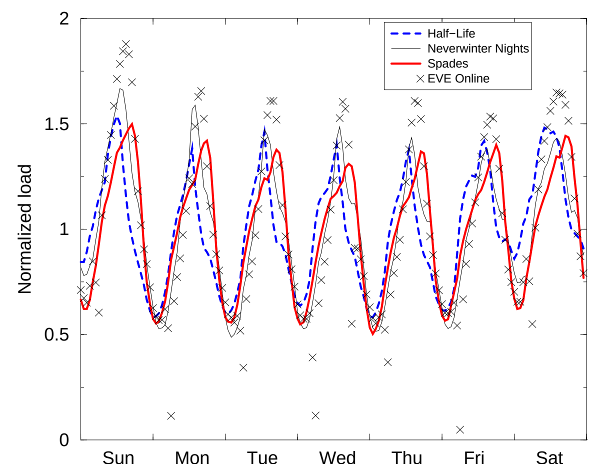

Fig. 6. Normalized load of four popular games over a representative week

North American cereal manufacturer	
Start time	Mon Aug 13 2001
End time	Sun Aug 19 2001
Total requests	10,368,896
Average requests per hour	61,720
Content transferred	59.6 GB

North American credit card company	
Start time	Tue Aug 14 2001
End time	Mon Aug 20 2001
Total requests	112,590,195
Average requests per hour	670,180
Content transferred	366.4 GB

International beverage manufacturer	
Start time	Tue Aug 14 2001
End time	Sat Aug 18 2001
Total requests	11,932,946
Average requests per hour	99,441
Geographically resolvable	11,829,429
Average requests per hour (North America)	39,312
Average requests per hour (Europe)	18,072
Content transferred	51.1 GB

TABLE II
WEB SITE LOGS FOR WEEK OF AUGUST 13, 2001

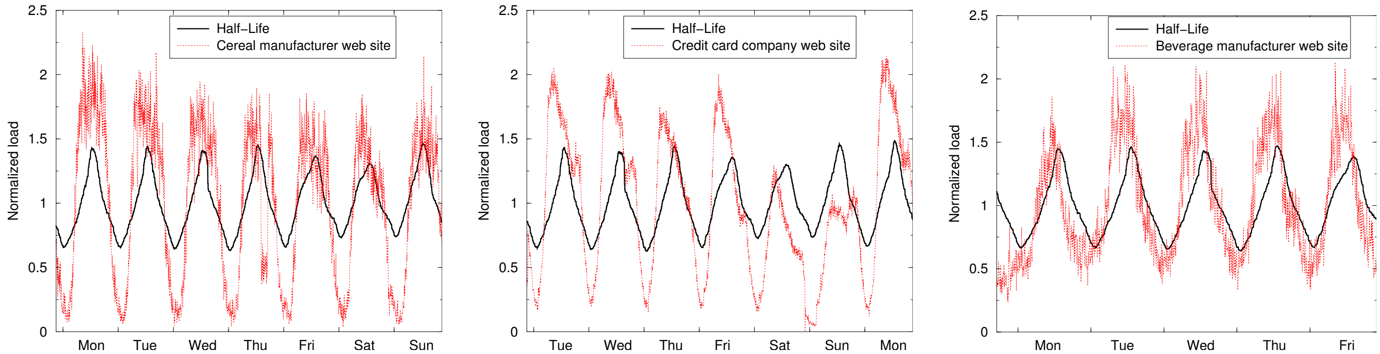

(a) North American cereal manufacturer

(b) North American credit card company

(c) International beverage manufacturer

Fig. 7. Aggregate normalized load between Half-Life and commercial web sites

can be multiplexed with each other, we examined the aggregate player populations of four popular games. The games examined include an FPS game (Half-Life), two MMORPGs (Neverwinter Nights, EVE Online), and a casual game (Spades). Player populations of these games were collected over a representative one week period. In order to compare the games directly, independent of their popularity, each game's population data was normalized by the mean population for that particular game during the week: 85,209 for Half-Life, 2984 for Neverwinter Nights, 4842 for Spades and 4096 for EVE Online. Figure 6 plots the normalized player loads for the four games over the one week period. As the figure shows, player populations fluctuate significantly based on the time of day from lows close to half of the mean to peaks close to twice the mean. The figure also shows that populations across games have peaks in close proximity to each other, making it difficult to achieve significant statistical multiplexing gain between different games. Unfortunately, this result indicates that the opportunity for on-demand computing to reduce peak resource usage across its game applications is quite limited.

D. *Games and interactive workloads are synchronized*

While it is difficult to obtain statistical multiplexing gain between different games, on-demand computing infrastructure could still be useful for multiplexing between other applications such as web servers. In order to examine this, we obtained web server logs over a week for three commercial sites. The sites included those for a North American cereal manufacturer, a North American credit card company, and an international beverage manufacturer. Table II describes the traces of the web servers, all from the week of August 13, 2001. The servers themselves were located in geographically distributed data centers and the individual logs from each site were aggregated and sorted into a single log file. Using these traces, we plotted the normalized load for the web server against the normalized global aggregate load of Half-Life during the same week in August 2004. Note that, similar to the game workload, the web site workloads were also normalized based on the average request rate throughout the trace as given in Table II in order to directly compare the workloads.

Total connections	71,253
Geographically resolvable	30,226
From North America	9,414
Average connection rate per hour (North America)	56.0
From Asia	9,814
Average connection rate per hour (Asia)	58.4
From Europe	8,788
Average connection rate per hour (Europe)	52.3
From other continents	2,210

TABLE III
CONNECTION DATA FOR `cs.mshmro.com` FOR WEEK OF MAY 23, 2004

As Figure 7 shows, workloads for web and on-line games share similar daily periodic peaks. This particular week of game traffic does not have a strong weekend rise (perhaps due to being from the summer), but the web traffic does slump during the weekends as the figure shows [2]. Interestingly, Half-Life shows considerably less variance than the North American websites, but similar variance to the international beverage manufacturer website. Intuitively, it makes sense that applications and web sites with global usage patterns are more consistently busy and have less daily variance. Due to the international popularity of Half-Life, its usage pattern is quite similar to that of the international beverage company's web site. Overall, these results indicate that infrastructure sharing between applications during the week will have a somewhat limited benefit with some potential for multiplexing gain during the weekends and during the "off hours" for geocentric applications. That is, workloads are driven by the time of day users are typically awake, which for games and interactive applications are identical.

E. *Games exhibit strong, diurnal geographic patterns*

One of the salient features of globally distributed, on-demand computing infrastructure is that it can easily shift resources geographically close to where the demand is coming from. Intuitively, it makes sense that a predictable, diurnal pattern drives global resource consumption and hence, the provisioning of server resources. This is especially the case

[2] Note that for the international beverage, we were only able to obtain a 5 day trace of their web site that did not include weekend days.

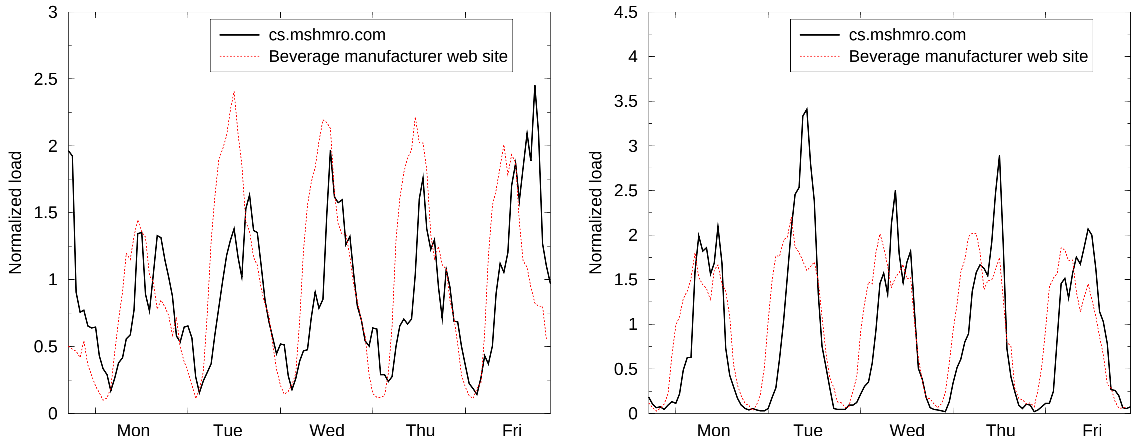

(a) North America (b) Europe

Fig. 9. Normalized load for `cs.mshmro.com` and the international beverage company website

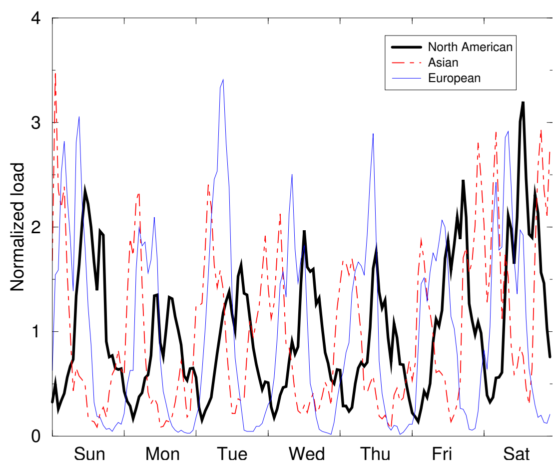

Fig. 8. Aggregate normalized load per-continent for `cs.mshmro.com`

for applications that require human participants such as games. To study this phenomenon, we examined a one-week period of `cs.mshmro.com` (Sunday May 23, 2004 to Saturday May 29, 2004). Note that due to the server's popularity, it was almost always completely full with 22 players during the entire period and thus, did not exhibit the load variation as the aggregate load shown in Figure 3(a). While the server had an overall constant load, we examined the geographic distribution of the load as it changed based on the time of day. Using a commercial geographic IP address mapping tool [20], the continent that each player connected from was resolved and aggregated. As Table III shows, players connect from geographically diverse regions. Using the resolved connections, the per-continent load normalized by the mean connection arrival rate listed in the table was plotted. As Figure 8 shows, each continent shows a predictable, diurnal pattern of activity with the only difference being a time-zone shift. This geographic pattern is consistent with the qualitative results shown previously for this server [21]. The diurnal pattern indicates that people have regular daily activity patterns that are tied to working hours. The lack of overlap especially between players in Asia and North America also indicates that, contrary to common perception, there aren't a significant number of players that stay up all night playing with people residing on different continents.

It is interesting to note that in contrast to the Half-Life aggregate load and international beverage company web site load (Figure 7(c)), the per-continent load of `cs.mshmro.com` exhibits a large variance similar to the North American web site loads shown in Figures 7(a) and 7(b). We hypothesize that when the usage patterns of international services are broken out into individual regions, the resulting load variances are similar to those of regional servers such as the cereal manufacturer and the credit card company. To test this hypothesis, we compared the per-continent load between `cs.mshmro.com` and the international beverage company web server trace. Figure 9 shows the per-continent, normalized load of the game and web server for North America and Europe. Note that for the web server data, the loads are normalized by the average request rate per-continent as shown in Table II. As expected, the per-continent load fluctuations and variance of the game server are similar to those found for the web site. The figure also shows that usage of both applications are highly synchronized when broken down into geographic regions. The degree of synchronization thus limits the benefits that geographically distributed, on-demand computing infrastructure has on interactive applications such as games and web.

IV. Player Characterization

A successful game must keep its players satisfied. Thus, it is important for game providers to understand the behavior of its players so that they can adequately address their needs. In order to study player characteristics, we analyze the trace of `cs.mshmro.com` and of EVE Online to track individual players throughout their play cycle. Note that we do not have a similar trace for casual games. Thus, our results on player characteristics are specific to FPS and MMORPG games.

A. *Gamers are impatient*

Quantifying the patience of on-line gamers is important for adequate server provisioning. For some Internet applications

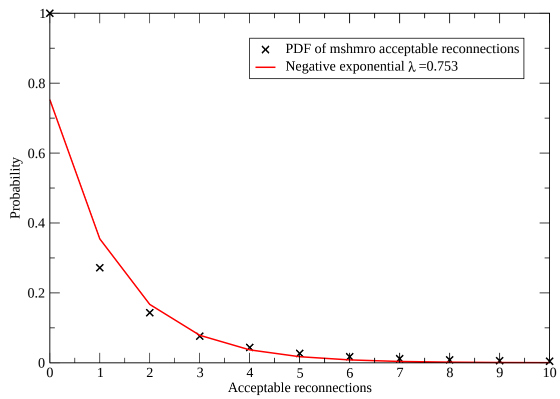

Fig. 10. PDF of player impatience based on number of acceptable reconnects

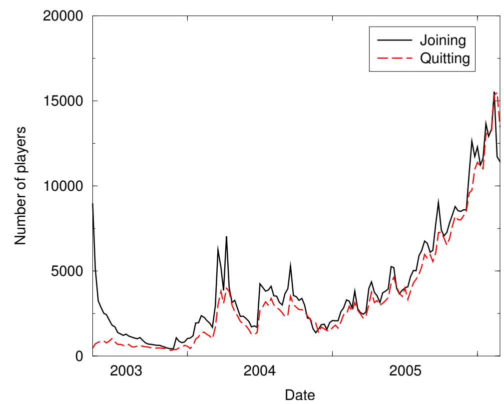

Fig. 11. Weekly rates of joining and quitting.

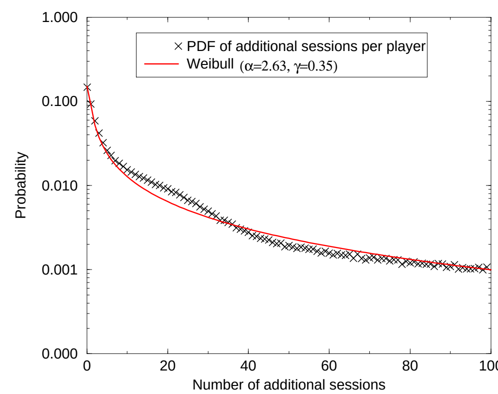

Fig. 12. Distribution of number of sessions per player for EVE Online.

such as web-browsing, users are impatient [22], while for others such as peer-to-peer file sharing, users are very patient [23]. Although our session trace of EVE Online only records successful connections and can not be used to study impatience, our trace of `cs.mshmro.com` records successful connections as well as connection attempts that are refused by the server. Due to its popularity, the server turns away thousands of players daily with several players reconnecting several times in a row, waiting for a spot on the server to free up. We assume that a player's willingness to reconnect to the same busy server repeatedly is an indication of their patience.

To quantify player patience we group each player's connection history into sessions, and consider a session of length N evidence of that player's willingness to reconnect after $N-1$ connections. Figure 10 shows the probability distribution of acceptable reconnections per player. As the figure shows 73% of the players are unwilling to reconnect to the server enough to play even once. One of the reasons players do not reconnect is that game clients have a "Quick Start" mechanism that many players use. The mechanism works by downloading a list of candidate servers from the master server and cycling through them one by one until a successful session is established. Thus, such clients may not lack patience, but rather are automatically redirected elsewhere. Aside from the first data point, the rest of the graph represents a client's patience in connecting to our busy server. When considering a client that does not subsequently try to reconnect as a 'failure', the trend of the graph is similar to that seen in reliability models for components whose lifetimes are exponentially distributed [24]. As shown in Figure 10, a negative exponential distribution with parameter $\lambda = 0.753$ [3] roughly fits the data with an R-square goodness fit of 0.978. Players, therefore, exhibit a remarkable degree of impatience with busy game servers.

B. Player churn is substantial

The ultimate goal for a game publisher is to increase the number of players subscribed to the game. To study player growth and churn, we examine the EVE Online trace and measure the number of players playing for their first time and the number of players playing for their final time across the trace. Our definition of players who have "quit" is that the player has not played for the last month of the trace. For such players, the time of the last recorded playing session by the player is considered the time he/she has "quit" the game.

Figure 11 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 were playing the beta version of the game and thus already knew they wanted to continue playing the game. To validate the observation that many players only stay for a short time, Figure 12 shows the PDF of the number of sessions that players play before quitting 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 having parameters α=2.63 and γ=0.35 [4] with an R-square goodness fit of 0.997.

One of the problems that MMORPGs have is that player churn increases over time. That is, the game's ability to attract and retain new players decreases over time. 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

[3] Estimated using the Maximum Likelihood Estimation methodology [24].

[4] Estimated using the probability plotting methodology [24].

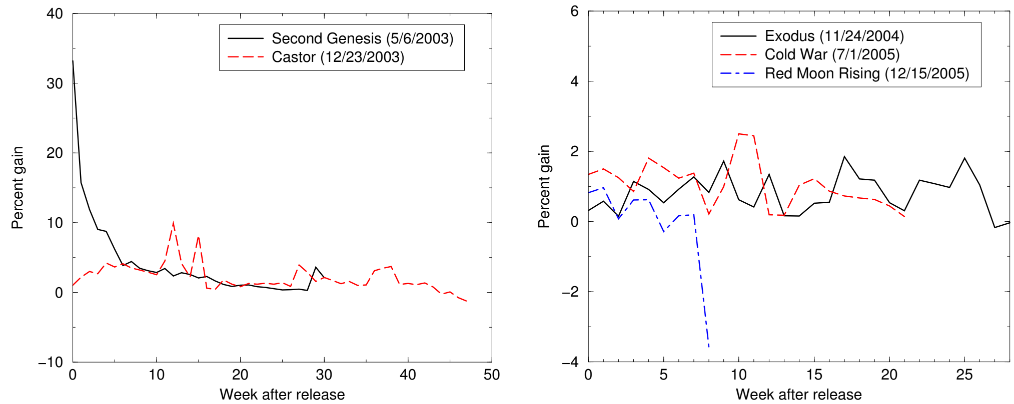

Fig. 14. Weekly percent gain in players after EVE Online updates.

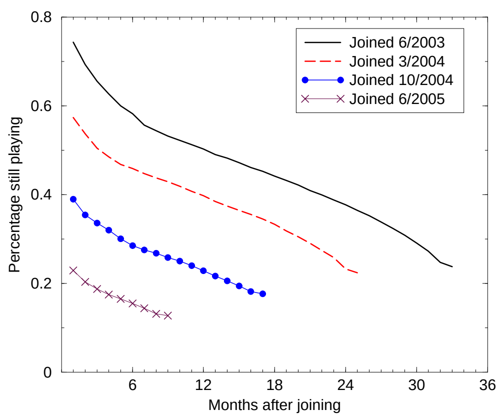

Fig. 13. Player retention over time for 4 different months.

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 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. To quantify this, Figure 13 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, but only around 30% are still playing by the end of 2005. The figure shows 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. Note that the curves shown in Figure 13 have similar slopes. That is, roughly 10% of the original player population quits the game every 6 months. Thus, while newer player populations have a much smaller initial retention rate, they subsequently decay at a rate similar to older player populations. Finally, in comparing retention rates to the overall player population of EVE Online shown in Figure 1(d), the continual growth in popularity of the game coupled with the gradual decline in retention rates indicates a much larger number of players are joining and subsequently quitting the game over time.

C. Updates slightly impact player growth and gameplay

One way for a game publisher to increase subscriptions is to provide new content. Thus, on-line games are often updated to keep players interested. Ideally, such updates would have direct impact on player retention. Figure 14 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 only 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 were the result of two marketing campaigns. The first involved a promotion targeting players of a competing sci-fi MMORPG "Earth and Beyond" [25] which was being shuttered [5]. The second involved a promotion at a popular game conference that involved giving away a large number of free trial accounts.

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 15 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

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

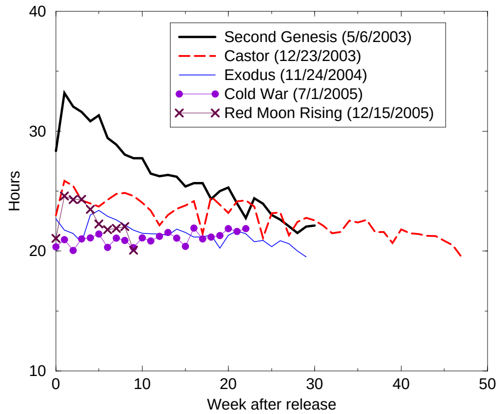

Fig. 15. Weekly minutes played per player after EVE Online updates.

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.

D. Gamers reveal when they lose interest

Since acquiring players is very difficult, one strategy game publishers might employ is to identify active players that might be losing interest in the game. By detecting waning interest, publishers may be able to deliver new content or customized incentives on a per-player basis to keep the player engaged with the game. Key to our analysis of player interest is the notion of a player's *play history*. The play history of a player consists of the sequence of play sessions for each player sessions from their first to their final recorded session. In order to ensure that a player has indeed "finished" playing the game at the end of their play history, players who have played the game during the last month of the trace are considered to be still active and are not included in our measurements. Since each player may progress through his or her game interest at a different rate, we normalize each of these measurements based on the duration each player is active on the server. Thus, a player playing the game for the first time corresponds to the 0% point of their play history while a player playing the game for the final time corresponds to the 100% point of their play history.

To study each player's play history, we examine the EVE Online trace which includes every game session from every player that has played the game. The first metric we examine is the amount of time a player spends playing the game per week. The intuition is that the longer a player spends on a game, the less likely the player is to quit the game. Figure 16(a) shows the average weekly playtime individual players spend based on their play history. As the figure shows, players steadily decrease their usage before eventually quitting the game. Because the figure averages in all users together regardless of how much they play, we also plot a normalized version that is based on each player's average weekly playtime. Specifically, the normalized playtime is the ratio of a players weekly playtime and their average weekly playtime over their play history. Figure 16(b) plots the average normalized playtime for players throughout their play history. As expected, players play substantially less than their average towards the end of their play history, thus giving game publishers a useful metric for determining when to deliver incentives to individual players.

Since the end of a player's play history is clearly important, we focus on this part of a player's play history using our other metrics: session times and intersession times. Figure 17(a) shows the PDF of player session times throughout 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 fitted with a Weibull distribution. To detect players that are about to quit, 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 another useful signal for game publishers to catch players about to quit. Figure 17(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 18(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 18(b) shows this more clearly by plotting the CDF of intersession times overall and of final intersession times.

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 times? Figure 19 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.

V. Conclusions

To understand the increasingly popular on-line game application, this paper presents a comprehensive analysis of a collection of on-line game player and game workload data from a number of unique sources across a wide range of game genres. Our key findings are that:

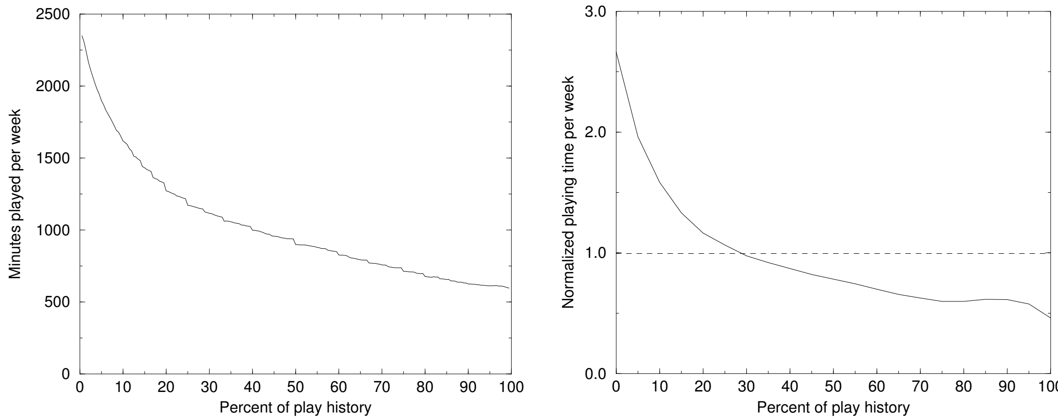

(a) Average playtime across players (b) Normalized per-player playtime

Fig. 16. Weekly playtime over play history for EVE Online.

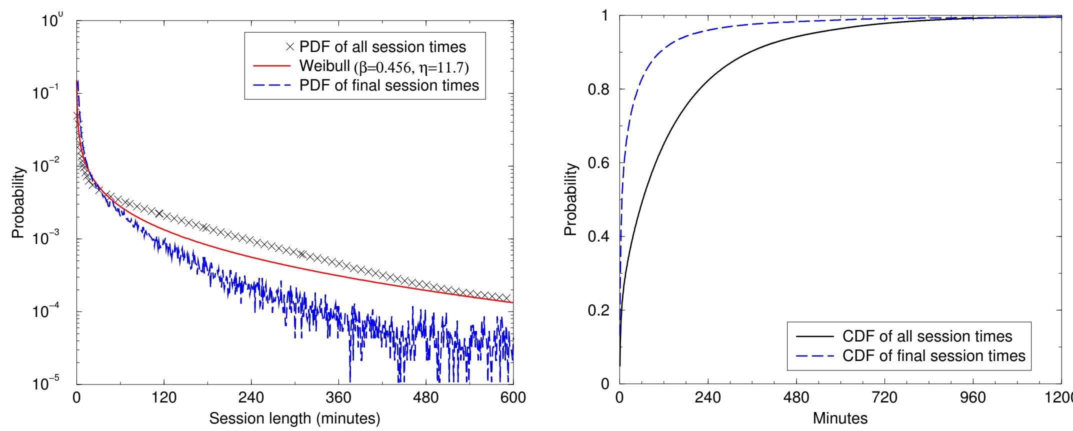

(a) PDF (b) CDF

Fig. 17. Player session time and final session time distributions for EVE Online.

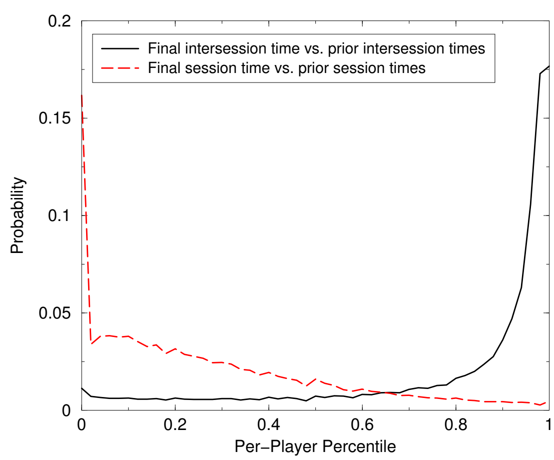

Fig. 19. “Final” session and intersession time percentiles for EVE Online.

- *Game popularity follows a power-law*
- *Aggregate game workloads across all genres are similar and predictable over the short-term*
- *Game workloads are synchronized with each other as well as other interactive applications such as the web*
- *Gamers have no tolerance for busy servers*
- *Player churn is substantial and increases over time*
- *Gamers change their play behavior when they are about to quit altogether*

References

[1] 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.

[2] K. Chen, P. Huang, C. Huang, and C. Lei, “Game Traffic Analysis: An MMORPG Perspective,” in *Proceedings of NOSSDAV*, June 2005.

[3] M. Claypool, D. LaPoint, and J. Winslow, “Network Analysis of Counter-Strike and Starcraft,” in *IEEE Internation Performance, Computing, and Communications Conference*, April 2003.

[4] W. Feng, F. Chang, W. Feng, and J. Walpole, “Provisioning On-line Games: A Traffic Analysis of a Busy Counter-Strike Server,” in *Proc. of the Internet Measurement Workshop*, November 2002.

[5] GameSpy Industries, “GameSpy: Gaming’s Home Page,” 2002, http://www.gamespy.com/.

[6] MSN, “MSN Games: Free Online Games,” 2005, http://zone.msn.com/.

[7] mshmro.com, “Counter-strike Server,” http://www.mshmro.com/.

[8] AMDZone, “Valve Releases Hammer Port of Counter-Strike Server,” 2003, http://www.amdzone.com/releaseview.cfm?ReleaseID=1050.

[9] ServerSpy, “ServerSpy.Net: World Server Ranks,” 2004, http://www.serverspy.net/site/serverranks/.

[10] B. Woodcock, “MMOGCHART.COM,” 2006, http://www.mmogchart.com/.

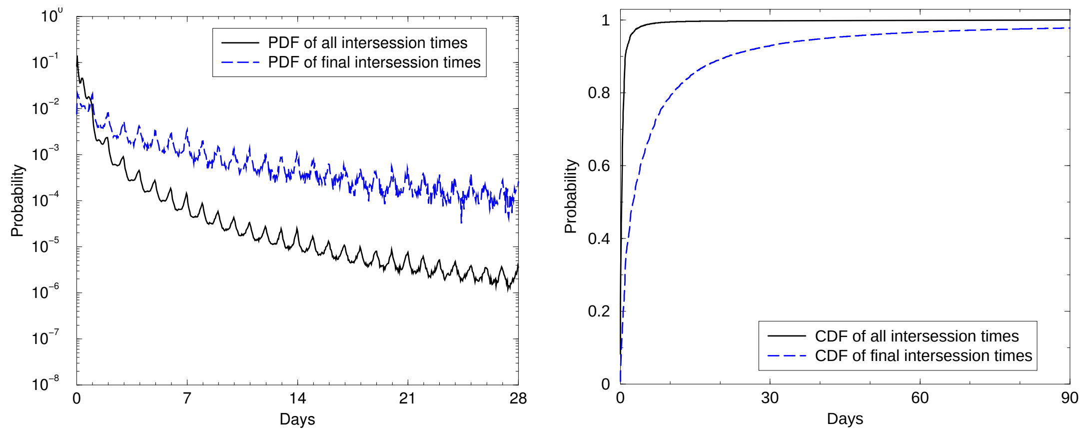

Fig. 18. Player intersession time and final intersession time distributions for EVE Online.

[11] 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.
[12] CNN, "SoBig.F Breaks Virus Speed Records," 2003, http://www.cnn.com/2003/TECH/internet/08/21/sobig.virus.
[13] IBM Corp., "On demand business," http://www.ibm.com/ondemand.
[14] E. Manoel *et al.*, *Provisioning On Demand: Introducing IBM Tivoli Intelligent ThinkDynamic Orchestrator*. IBM International Technical Support Organization, December 2003, http://www.redbooks.ibm.com.
[15] D. Saha, S. Sahu, and A. Shaikh, "A Service Platform for On-Line Games," in *NetGames*, Redwood City, CA 2003.
[16] A. Shaikh, S. Sahu, M. Rosu, M. Shea, and D. Saha, "Implementation of a Service Platform for Online Games," in *NetGames*, August 2004.
[17] Butterfly.net, Inc., "Butterfly Grid Solution for Online Games," 2003, http://www.butterfly.net/.
[18] Microsoft Corporation, "Xbox Live," 2003, http://www.xbox.com/live.
[19] Electronic Arts, Inc., "EA.com," 2003, http://www.ea.com/.
[20] Geobytes, Inc., "Geobytes Home Page," 2003, http://www.geobytes.com/.
[21] W. Feng and W. Feng, "On the Geographic Distribution of On-line Game Servers and Players," in *NetGames 2003*, May 2003.
[22] L. Catledge and J. Pitkow, "Characterizing Browsing Strategies in the World-Wide Web," *Computer Networks and ISDN Systems*, vol. 27, no. 6, pp. 1065–1073, 1995. [Online]. Available: citeseer.ist.psu.edu/catledge95characterizing.html
[23] K. Gummadi, R. Dunn, S. Saroiu, S. Gribble, H. Levy, and J. Zahorjan, "Measurement, Modeling, and Analysis of a Peer-to-Peer Workload," in *Proceedings of ACM SOSP*, October 2003.
[24] ReliaSoft Corporation, "Life Data Analysis and Reliability Engineering Theory and Principles Reference from ReliaSoft," 2003, http://www.weibull.com/lifedatawebcontents.htm.
[25] "Earth and Beyond Trial Program," http://eve-online.com/enb.

PLACE PHOTO HERE

C hris Chambers is a game developer at Zipper Interactive. He holds a Ph.D. degree in Computer Science from Portland State University.

PLACE PHOTO HERE

W u-chang Feng is an associate professor of Computer Science at Portland State University. He holds a B.S. degree in Computer Engineering from the Pennsylvania State University and an MSE and Ph.D. degree in Computer Science Engineering from the University of Michigan.

PLACE PHOTO HERE

S ambit Sahu is a research staff member in the Networking Software and Service group at the IBM T.J. Watson Research Center. He received a Ph.D. degree from the University of Massachusetts at Amherst.

PLACE PHOTO HERE

D ebanjan Saha is a senior manager of the Service Enablement Infrastructure group at the IBM T.J. Watson Research Center. He holds a B. Tech degree from IIT, India and M.S. and Ph.D. degrees in Computer Science from the University of Maryland at College Park.

PLACE PHOTO HERE

D avid Brandt is currently the Chief Technical Officer at Networked Alternate Reality Creations (NARC). Prior to NARC, he was a senior network programer at CCP Games, the publisher of the EVE Online MMORPG.