A Large-Scale, Longitudinal Study of User Profiles in World of Warcraft

Jonathan Bell, Swapneel Sheth, Gail Kaiser
Department of Computer Science
Columbia University, New York, NY, USA
{jbell, swapneel, kaiser}@cs.columbia.edu

ABSTRACT

We present a survey of usage of the popular Massively Multiplayer Online Role Playing Game, World of Warcraft. Players within this game often self-organize into communities with similar interests and/or styles of play. By mining publicly available data, we collected a dataset consisting of the complete player history for approximately six million characters, with partial data for another six million characters. The paper provides a thorough description of the distributed approach used to collect this massive community data set, and then focuses on an analysis of player achievement data in particular, exposing trends in play from this highly successful game. From this data, we present several findings regarding player profiles. We correlate achievements with motivations based upon a previously-defined motivation model, and then classify players based on the categories of achievements that they pursued. Experiments show players who fall within each of these buckets can play differently, and that as players progress through game content, their play style evolves as well.

Categories and Subject Descriptors

H.1.2 [**Information Systems**]: User/Machine Systems—*Software Psychology*; H.5.1 [**Information Interfaces and Presentation**]: Multimedia Information Systems—*Artificial, augmented, and virtual realities*; H.2.8 [**Information Systems**]: Database Applications—*Data mining*

Keywords

virtual worlds; user profiles; video games; web information mining; world of warcraft

1. INTRODUCTION

It is vital to study how games succeed and fail in order to improve future titles. By studying player's behavior, we can understand users' tendencies towards different sorts of game content at a macro level — their *player profile*. To this end, work has been done to study players of these games both qualitatively [8, 19, 21, 22] and quantitatively [10, 13]. Qualitative studies typically involve conducting user studies, collecting information through online surveys. Quantitative studies, on the other hand, are typically based on more longitudinal data, typically game *metrics* recorded by

game studios or researchers. In either case, the information gathered can be used in many ways within game development, allowing developers to customize users' interactions with the game. For instance, work has been done to identify player's motivations for playing [22], and to correlate those motivations with player retention [8], allowing developers to use this data to improve user retention.

We recognize that the quantitative study of games is a developing field, with many data sources that have not been fully exploited, and that there are still questions unasked. We seek to combine findings from existing qualitative game studies with quantitative gameplay data in order to better understand how people play participate in online games. We propose to investigate the following three research questions, in the context of the Massively Multiplayer Online Role Playing Game (MMORPG), World of Warcraft (WoW) [6]:

1. Can we cluster characters based on the type of goals (e.g., what sort of motivations could exist for that goal) that they complete in game?

2. Do characters' play profile evolve as they level?

3. Do characters with different play profiles play through the game content at different rates?

To begin to answer these questions we performed a large-scale data crawl of Blizzard's WoW Armory [6], building and improving on prior techniques to massively increase our data sample from hundreds of thousands of characters [13] to tens of millions. The WoW Armory contains a complete listing of player data for all currently active players. While previous work [10,13] has made use of data from The Armory (described further in Section 2), we do not believe that prior researchers have fully investigated the available data.

Among other records, the Armory includes details of players' *achievements* — records of in-game accomplishments. Achievements can be simple, for example, "Stable Keeper," which is awarded upon obtaining ten riding mounts (which are easily purchased in-game). Other achievements can be incredibly complex, for example, "A Tribute to Immortality," which requires players to form a group with 24 others, and then survive five "heroic" (very difficult) encounters without allowing any member of the group to die. We gathered complete data for over six million characters (including their achievement records) as well as partial data for another six million characters and analyzed it. These 12 million characters were sampled across 578 different servers in the Americas, Europe, and Asia.

The remainder of this paper is structured as follows: We describe the related work in the next section. Section 3

describes in detail our approach to retrieving character usage data from Blizzard. Section 4 presents the findings from our study, and Section 5 concludes our paper with a discussion of our results.

2. RELATED WORK

This sort of study has been proposed for other sorts of virtual communities before, for example Orgaz and colleagues showed a general technique to extract data from a VW and cluster it, but used different attributes from non-game data sets [14]. Games such as World of Warcraft provide a rich set of user data, which has been the focus of many interesting experiments and games research. We describe several recent related papers (within the games community) next.

Hullett *et al.* [11] analyzed data from *Project Gotham Racing 4*, an XBOX 360 car racing game. They used data from thousands of users to provide feedback regarding the most popular cars, game modes, and event types to the game development team. The game development team would presumably use this data to make the next version of the game better and more appealing to their users.

Lewis and Wardrip-Fruin [13], one of the first papers that attempted a large scale survey of WoW using publicly available data, used a web crawler and screen scraper to collect information on 136,047 characters. They used the collected data to analyze game characteristics such as classifying players based on what items they were holding, time it takes to reach a certain level based on player class, and number of deaths based on player class. They showed that game data that was previously only available to internal developers at the game companies was now available publicly to the world. Moreover, they presented a tool to easily collect the data, to allow researchers to gain insight into these games and lead to interesting qualitative studies.

Harrison and Roberts [10] used the above-mentioned WoW crawler to create player models to predict a player's behavior in a game, over a sample of approximately 15,000 players. They validated that their predictions were accurate by using cross-validation and measuring precision and recall for their models. They showed that their model is statistically significantly better than a baseline algorithm. Ashton and Verbrugge [1] create a monitoring plugin for WoW to measure the level of difficulty of game play.

The papers mentioned so far have relied upon a quantitative, "data-driven" approach for research. On the other hand, there have been many papers by Yee and colleagues (*e.g.,* [8, 19, 21, 22]) that have primarily used user studies and online surveys to explore different aspects of WoW such as player motivations, personality, and demographics. These surveys typically involve several thousand respondents, found through online message boards. For example, Debeauvais *et al.* [8] used online user surveys to analyze player commitment and retention in WoW. 2,865 WoW players completed the survey and Debeuavais *et al.* used the answers to analyze topics such as number of hours played per week, number of years the respondents had been playing WoW, and the ratio of respondents who stopped playing the game and returned to it at a later point. In a "data-driven" paper, they performed a census of approximately 220,000 characters and collected additional data using automated bots [9]. This is the largest dataset of characters used in any paper so far. This data was used to address such game metrics as playing time by character level, in-game demographics (such as

character races, classes, and genders), and character abandonment rate by class.

T.L. Taylor has several relevant works as well. In [17], Taylor explores the notions of how language, age, nationality, etc. affect game play and game culture. In [12], Jakobsson and Taylor explore the role of social networks and online communities in Everquest [15] (another MMORPG). Taylor has also written a book [16] that explores multiplayer culture and ethnography in Everquest.

While our work on WoW is inspired by these papers mentioned above, our work differs in many aspects. First, as we use the WoW API, we were able to gather data for about 12 million characters. This is at least one to two orders of magnitude more than any of the papers mentioned so far. Thus, we believe, that our work is more easily generalizable and limits selection and sampling bias. Second, we require no extra effort on behalf of the players of WoW or participant recruitment for our study. We use the publicly available data provided by Blizzard and their WoW API for all our findings. This makes our study very easily replicable and extendable, which would be hard to do with some of the other work mentioned here.

Third, we also address the lack of many large scale longitudinal studies for MMORPGs. WoW was released in November 2004; our data dates back to December 2006 (note that while achievements didn't exist this long ago, there were many that were granted after-the-fact when the achievement system was created). We thus have user data for the last five years (out of seven that the game has been in existence). This helps us to answer our research questions by looking at five years of data, rather than over a few months (or less) as reported by most of the other work. Finally, our research questions (described in Sections 1 and 4) are different from those already explored by the research community so far.

3. METHODOLOGY

To build our experiments, we gathered data for approximately 12 million characters over a 16 day timespan, with data dating back as far as December 2006. While we are unable to determine the oldest character collected (due to the limitation that achievements didn't begin until 2008), we have identified that our dataset includes characters who began playing before December 2006 — before Patch 2.0.1 was released (this patch made it impossible from that point on to earn a set of titles — and we have players who have those titles, who must have earned them before this patch).

For half (approximately 6 million) of the characters, we collected only basic biographical information (class, in-game gender, in-game race, and level). For the remaining half (also approximately 6 million characters), we captured as much information as was feasible to store, which included: achievements, professions, quests, raids, reputation, titles, mounts, and pets in addition to the same basic biographical information. For achievements we captured the date and time that the character completed the achievement. For raids we captured the number of times that the character successfully completed each raid and each raid boss, broken down by difficulty level. For professions we captured only the chosen professions and levels.

After collecting the data, we categorized each achievement by associating it with a "motivation" — social, immersion, and achievement — using the terms described by

Yee [22]. Finally, we ran a clustering algorithm over the coded achievement-player data.

3.1 Data Acquisition

As previous work [10, 13] has found, Blizzard's World of Warcraft "Armory" [6] provides an excellent resource for data. The Armory provides a complete listing of character data, including what items the character has equipped, what achievements they have completed, what quests they have completed, which companions and mounts they have obtained, their professions, and their reputation. In the summer of 2011, Blizzard released a RESTful web service [7], creating a new and more efficient means to access Armory data that was unavailable to previous work [10, 13].

The web service allows developers to query for information on a specific character, or to list the characters in a guild. It does not allow developers to list available guilds or all characters on a server. The service also is restricted to 3,000 requests per day per IP address. Therefore, we created an optimized three-phase character retrieval system. In the first phase, we discover names of possible guilds. In the second phase, we discover the names of members of those guilds (we discuss the validity of including only characters in guilds in Section 5). Finally, we query the Blizzard web service to obtain complete information for each discovered character.

3.1.1 Guild Discovery

To discover the names of guilds, we turned to the website WarcraftRealms.com [18]. WarcraftRealms.com provides an interface to a repository of basic census information — we found no other comparably large data set available. WarcraftRealms.com collects the names and guilds of all characters "seen" in game by users of an in-game UI modification. We retrieved the list of 126,317 guilds discovered by the WarcraftRealms.com census, and used that as a seed for discovering characters in the next step. These guilds were distributed across 578 servers — with 64,423 of the guilds on European servers, 55,403 on American servers, and 6,491 on Asian servers.

3.1.2 Character Discovery

With guild names in hand, Blizzard's API provides a simple means to discover character names. By providing the API with a guild and server name, we can easily retrieve a JSON formatted list of characters in that guild. Along with each character's name, we also were able to retrieve their level, class, and in-game race and gender. For example, to retrieve the list of all characters in the guild "Foundation" on server "Lightnings Blade," we need only make one request, to `http://us.battle.net/api/wow/guild/lightnings-blade/Foundation?fields=members`. This list includes all characters in the guild — including inactive members. In WoW, players who cancel their subscription become listed as "inactive" after 2 weeks, at which point their detailed information is archived and removed from all publicly accessible interfaces, but remain listed in their guild (unless they are specifically removed from it).

From this process, we discovered approximately 12 million characters. These characters are roughly evenly distributed across all of the World of Warcraft servers.


Figure 1: Crawler Design

3.1.3 Character Retrieval

After retrieving a list of characters, we moved on to the final phase of our crawling process: character retrieval. The Blizzard API for characters provides developers with access to all dimensions about a character in a single request. For example, this sample request retrieves the basic profile for the character "Wizlock" on the server "Borean Tundra," along with data for achievements, quests, professions, raid progress, companions, mounts, titles, and reputation: `http://us.battle.net/api/wow/character/Borean%20Tundra/Wizlock?fields=achievements,quests,professions,progression,companions,mounts,titles,reputation`.

Due to the request limit, we retrieved all data that we considered to be remotely relevant, so as to avoid the need to make more requests in the future. The only data provided that we did not access was: the list of items equipped by the character, the list of talents learned by the character, and PvP team membership.

We parallelized this process across approximately 60 machines, with each machine retrieving 30 characters at a time, then updating the database, retrieving another 30, and so on. We used this chunking strategy to decrease update-lock contention on the database server. We used a "depth-first" retrieval of characters. That is, we retrieved all characters that we could for a particular server before moving on to the next one. This was done quite intentionally to be sure that we had as many complete server datasets as we could (we anticipated being unable to complete our entire in-depth crawl, see section 3.1.5).

We ran this process from September 29th, 2011 through October 14, 2011, during which time we visited 8,878,429 different characters. Of these, 2,112,940 appeared to be inactive characters, and 6,765,489 were active. For the approximately four million characters that we did not visit and the two million who we visited and found inactive, we still had basic data (level, class, and in-game gender and race) — leaving us with approximately 6 million characters that had complete information, and 6 million with only partial.

We anticipate that with 8-10 more days we could have completed exploring all characters that we had discovered, but stopped our crawl at Blizzard's request.

3.1.4 Metadata Retrieval

Data returned from the character API is relatively compact: it contains references to achievement IDs and quest IDs but lacks metadata about these quests and achievements. Achievement information was retrieved from an undocumented feed in the Blizzard API, at `http://us.battle.net/api/wow/data/character/achievements`. Quest information was retrieved from the locale files included in the installation of WoW. We used a tool called MPQEditor [23] to unpack the data archives included in WoW.

3.1.5 Crawler Design and Scalability

We created a distributed system, running in the background on a lab of 60 machines, with each slave controlled by a central master. The master node contained a MySQL server which stored all "to-do" characters to retrieve, and all resulting character information (this high level design is shown in figure 1).

We found that while the published limit was 3,000 requests per day, some nodes were able to achieve over 12,000 requests per day. Further investigation showed that the actual limit of requests per day was determined by system load - hosts making over 3,000 requests per day were allowed to continue if load was low on the server. Nonetheless, each host was designed to completely stop making any requests for 24 hours upon receiving a single message from the server that it was at its limit, so as to not put a burden on Blizzard's servers. Each request was also signed with the first author of this paper's name and email address as the user agent HTTP header, to facilitate communication should an issue arise.

By focusing the time that the process ran to late at night, we were able to increase the number of requests before being throttled to approximately 8,500 complete characters per host per day. At peak grid performance, we captured 1,171,516 characters in a single day, with each node taking approximately 1.5 seconds to process each character. After two weeks of retrieving characters at an average of 554,901 characters per day, we were asked by Blizzard to scale down our crawl to fewer nodes, but stopped outright, satisfied with the total number of characters for which we now had data.

3.2 Data Cleaning

We had to manually remove data from some of our analyses in several cases where it was clearly incorrect. This arose when an achievement was retroactively awarded, as it was timestamped to the first time that the character was logged in after it was awarded. This misinformation would distort our temporal analyses (although not our overall analyses), so we censored the following data from our time-scale analyses:

1. Achievements awarded before January 2009: Achievements were not introduced until mid October, 2008 [4], so in the three months following this rollout we ignore all achievement timestamps, suspecting that in any of these cases the user may have just logged in, and received an achievement retroactively.

2. Upon the release of the "Cataclysm" system patch in November, 2010 [5], the world map was revamped. This led to many characters simultaneously receiving achievements for "exploring" entire areas which they hadn't already fully explored. Similarly, a new

	Our Sample		Previous Results [18]	
Class	Size	Percent	Size	Percent
Warrior	1,241,040	10%	341,966	10%
Paladin	1,465,169	12%	444,556	13%
Hunter	1,409,296	11%	410,360	12%
Rogue	1,095,414	9%	273,573	8%
Priest	1,243,601	10%	341,966	10%
Death Knight	1,304,117	10%	341,966	10%
Shaman	1,105,101	9%	273,573	8%
Mage	1,331,620	11%	376,163	11%
Warlock	990,729	8%	205,180	6%
Druid	1,446,516	11%	410,360	12%

Table 1: Character class distributions

	Our Sample		Previous Results [18]	
Race	Size	Percent	Size	Percent
Human	2,188,208	18%	581,342.88	17%
Orc	942,570	8%	205,179.84	6%
Night Elf	1,592,896	13%	410,359.68	12%
Undead	1,094,839	9%	205,179.84	6%
Gnome	737,307	6%	170,983.20	5%
Troll	913,159	7%	205,179.84	6%
Tauren	1,074,708	9%	273,573.12	8%

Table 2: Character race distributions

achievement was created called "Surveying the Damage," awarded upon visiting the areas of the map that had been changed in the world revamp. However, we found that many characters received this achievement simultaneously: our suspicion is that characters were awarded this achievement immediately if they had *already* visited the changed areas (before they had changed).

In both of these cases, we continued to use these achievements when investigating characters' overall activities, but ignored them when tracking characters' progress over time.

3.2.1 Data Profile

As a sanity check, we compared the general statistics of our sampled characters with those from the website WarcraftRealms.com [18]. While our dataset was seeded from their list of guilds, our dataset contains both significantly more characters — we gathered (at least) partial data for over 12 million characters compared with their 3 million — and significantly more data — our data set includes complete information on achievement and quest completion, while theirs does not. We compared statistics on character class and character race from our sample to theirs, and were pleasantly surprised to see very similar (nearly identical) figures, which we believe validates our sampling method. We considered comparing these results against other academic publications, such as [9], however, these data sets pre-date the most recent expansions to the game, and therefore do not reflect the same environment. Note that WarcraftRealms.com does

Title	Description	Achievement	Social	Immersion
1000 Conquest Points	Earn 1000 Conquest Points	X		
Arrested Development	Allow all three of Corla's zealots to evolve, then defeat Corla after slaying the evolved zealots in Blackrock Caverns on Heroic Difficulty.	X	X	X
Archavon the Stone Watcher (10 player)	Defeat Archavon the Stone Watcher in 10 player mode.		X	
It's Happy Hour Somewhere	Drink 25 different types of beverages.			X
Silence is Golden	Defeat Atramedes in Blackwing Descent without any raid member's sound bar going over 50%.	X	X	
The Immortal	Within one raid lockout period, defeat every boss in Naxxramas without allowing any raid member to die during any of the boss encounters in 25-player mode.	X	X	
The Harder they Fall	Discover how orc Chieftan Hargal was killed by collecting the following artifacts.			X

Table 3: Sample listing of WoW achievements with coding

not have statistics for the races Goblin, Blood Elf, Draenei, and Worgen displayed, so we omitted those races from our comparison, and percentages in that table will not add to 100. We believe that our approach to using a seed guild list from a third party (WarcraftRealms.com) combined with authoritative guild membership and character information (from Blizzard) provides the most representative sampling given available resources. Table 1 presents a comparison of class distribution, while Table 2 presents the comparison of race distributions.

3.3 Data Analysis

Key to our analysis was a coding of achievements into the categories identified by Yee — social, achievement, and immersion [22]. As defined by Yee, social oriented play focuses on socializing, relationship building and teamwork, while immersion oriented play focuses on discovery (of lore, for example), role-playing, customization or escapism. Per his categorization, achievement oriented play focuses on advancement (of progress, power or status), game mechanics (optimization, tempting or analysis), or competition (challenging others).

Using the subcomponents of each category, we labeled each achievement as being in at least one of those categories using our domain-knowledge of World of Warcraft. For example: all "raid" achievements (a raid is a difficult set of encounters with bosses and their minions, typically requiring either 10, 25 or 40 players to work together to achieve) were categorized as "social." Achievements that were based on competition or character advancement in-game were categorized as "achievement." For instance, achievements of the variety "Level 10" (granted for reaching that level, similar achievements exist for each 10 levels from 20-80 and 85) or "1500 Quests Completed" (again, similar achievements exist at different progressions) were categorized as "achievement." Achievements where characters were completing tasks related to lore or discovery in game were categorized as "immersion," and included (among others) "holiday" achievements that were related to in-game holidays.

Some achievements were labeled with more than one category, as we found appropriate, and several examples of these cases can be found in table 3. While we have not included a complete listing of achievement codings in this paper (there


Figure 2: Characters clustered by the percentage of their achievements that were social, immersion, and achievement oriented.

are 1,700 achievements), we have included this information as an appendix in a technical report for reference [3].

4. RESULTS

4.1 Character Motivation Clusters

Our first research question was to validate our technique of clustering characters by the motivations that we tagged achievements. To answer this question, we performed a Scalable EM clustering using Microsoft's SQLServer Analysis Services 2008. Figure 2 shows a graphical view of the clustered characters. In this figure, the x-axis presents that percentage of a character's achievements that were socially-

	A	**B**	**C**	**D**
Population	1,652,063	1,429,169	1,022,530	974,554
Social	12% ± 6	7% ± 8	1% ± 3	28% ± 6
Immersion	38% ± 6	20% ± 7	38% ± 29	31% ± 5
Achievement	66% ± 4	87% ± 6	35% ± 24	61% ± 4

Table 4: Characters clustered by the percentage of their quests that were social, immersion, and achievement oriented. The standard deviation is also shown

Figure 3: Achievement motivation by level

Level	Immersion	Social	Achievement
0-9	3.23 ± 0.87	0.47 ± 0.87	0.86 ± 0.87
10-19	4.26 ± 0.77	0.83 ± 0.77	20.47 ± 0.77
20-29	6.28 ± 1.85	0.24 ± 1.85	25.26 ± 1.85
30-39	4.15 ± 0.67	0.24 ± 0.67	18.51 ± 0.67
40-49	7.12 ± 1.29	0.23 ± 1.29	22.27 ± 1.29
50-59	8.83 ± 1.91	0.47 ± 1.91	23.31 ± 1.91
60-69	12.09 ± 1.68	0.65 ± 1.68	41.47 ± 1.68
70-79	21.57 ± 2.73	2.23 ± 2.73	57.3 ± 2.73
80-84	108.36 ± 35.14	56.15 ± 35.14	157.55 ± 35.14
85	40.19 ± 4.81	68.59 ± 4.81	114.39 ± 4.81

Table 5: Achievements completed by type, by level. The standard error is also shown.

4.3 Time to Level

We investigated the time that each character took to reach new levels.

Since October 2008, Blizzard has captured when each character achieves a new "major" level (10, 20, 30, 40, 50, 60, 70, 80, and 85). As with all achievements, these are timestamped, allowing us to infer how much time passes between when a character reaches each of these milestones. We define the time to level between each milestone l_i as $l_i.\text{date} - l_{i-1}.\text{date}$. Note that the achievements are timestamped with "real-world" timestamps, and are not necessarily related to play time. Therefore, our definition of time to level includes time spent *not playing the game*, and captures the time between levels spent outside of the game, and captures not just how much play-time it takes to level a character, but how often a user plays.

We handle the following cases specially: (1) If a character reached level i at the time when i was the maximum, we calculate their time to reach level $i + 10$ as the difference between when they reached level $i + 10$ and when level $i + 10$ was available, rather than when the character reached i (this was the case for levels 60, 70, and 80, which all at one point were the maximum level attainable in the game). (2) If a character is recorded as reaching level i before December 2008, we ignore that data point, as it is "bogus." For characters active before December 2008, Blizzard retroactively granted all appropriate level achievements simultaneously, in November 2008.

4.3.1 By Character Classification

We classified each character as either achievement, social, or immersion oriented based on the overall percentage of their achievements that fell into each category — placing the character into whichever bucket the majority of their achievements were in. Table 6 presents these results (with standard error and population size), while Figure 4 presents them graphically. There are two interesting findings here: (1) using this classification technique, the majority of characters were categorized as "achievement driven," and (2) characters who focused on immersive achievements took far longer to reach level 85. We attribute the first finding to the overall distribution of achievements within our categorization. The second finding is interesting in that it suggests that we have been able to identify characters who "stop to smell the roses" — those who may go out of their way to explore game content and who do not focus on reaching end-game content sooner.

oriented, the y-axis presents the percentage of a character's achievements that were immersion-oriented, and the z-axis presents the percentage of a character's achievements that were socially-oriented. Larger circles represent larger clusters. The results are also presented for analysis in Table 4 as average values for each cluster with standard deviations. Note that because many achievements were categorized in multiple buckets (e.g., both social and achievement), the total percentages in each column may not sum to 100.

4.2 Character Motivations by Level

We also investigated character motivations over time, to see if characters received different sorts of achievements at different levels. We sampled all characters for whom we had a complete set of data — characters who started playing after December 2008 (the introduction of achievements) and are currently level 85 (the maximum level). Table 5 presents the average number of achievements (plus or minus the standard error) completed by characters broken down by category and by level. Figure 3 shows the same data in a graphical form. Note that there are two interesting trends here: (1) characters receive far more achievements later in the game than they do earlier (perhaps by design, perhaps not), and (2) social achievements don't really come into play until level 70. We believe that this is indicative that characters don't typically perform large group tasks until they achieve the maximum level — which was first 60, then 70, then 80, and most recently 85 (note that as Table 5 shows, characters *do* on occasion get some social achievements early in the game, but the average number received is less than 1). This may also be due to the structure of the game: large group play ("raids") are not available until level 60. Within our sampled characters, some of them were playing when level 70 was the maximum level reachable, and most were playing when 80 was the maximum attainable level.

	Total time to 85	**Population Size**
Immersion	$9,074 \pm 23.74$	56,745
Social	$4,911 \pm 9.99$	204,734
Achievement	$5,001 \pm 2.39$	3,499,764
Total	$5,050 \pm 2.32$	3,761,243

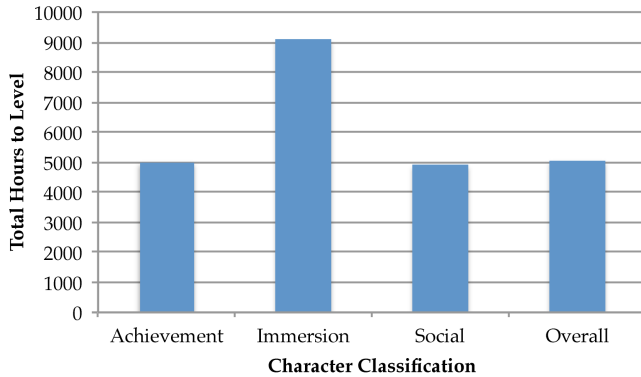**Table 6: Total time to level by play category**

Figure 4: Time to level, grouped by character type

4.3.2 By Start Date

We also segregated the character data to characters who began playing before the most recent expansion to WoW, "Cataclysm," released on November 23, 2011 [20], and those who began playing afterwards. In Figure 5, we present the time to level for 4,304,385 pre-cataclysm characters and 2,950,351 post-cataclysm characters. Note that to maintain the scale of the graph, we doubled the time that characters took to get from level 80 to 85, and marked this as "level 90." The finding is striking: characters who began playing after the release of "Cataclysm" leveled much more rapidly. There are several outside variables that we have identified that may influence this: (1) Blizzard has reportedly made early content easier. (2) There are new promotions to make it easier to level characters faster when you refer a friend to play. (3) After reaching end-game content on one character, it is possible to buy special items to allow the same player to level other characters faster. (4) Players may take longer to level their first character than later characters.

We are also interested in the spike of time to level in characters who took very long periods to level between 60 and 70, 70 and 80, and 80 and 85. Each of levels 60, 70, and 80 were at one point the maximum level achievable, and although we smoothed the data over these periods (see section 4.3, above), there are still large jumps in time to level. We attribute this to time spent *out of game*, rather than in game. In other words, these spikes represent players abandoning their characters upon reaching what was then the "end game" content, and not resuming until some time well after the next expansion was released. Since a single player may have multiple characters, it is also certainly possible that upon the release of an expansion, the player levels one of their characters first, plays the new end-game content, and then, only later levels other characters that they have.
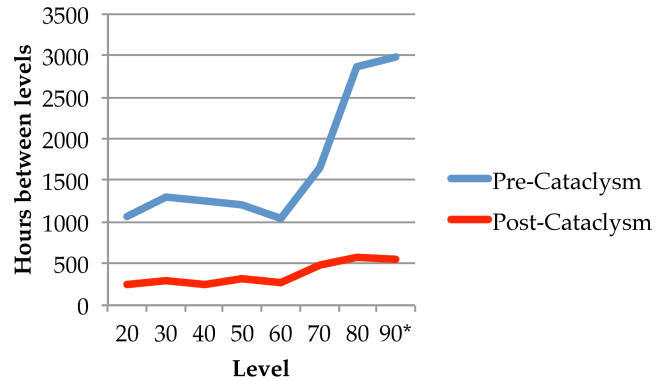

Figure 5: Time to level. Note that the time between 80 to 85 is doubled and presented as level "90" to maintain scale

5. DISCUSSION

5.1 Threats to Validity

5.1.1 Achievement Coding

Several of our analyses are grounded in the coding of achievements into the motivational buckets defined by Yee [22]. Therefore, the results that we see are dependent on the correct labels being applied. To create a consistent labeling, we conducted a two-pass algorithm with two different labelers. First, one student (with domain expertise in World of Warcraft; not an author on this paper) labeled every single achievement themselves. Then, the primary author (who also has domain expertise) reviewed the list for inconsistencies (e.g., two similar achievements being labeled differently), and then conferenced regarding any remaining concerns. While there may be room for discussion in what bucket some achievements belong in, we believe that we consistently labeled similar achievements. We have not included a complete listing of coded achievements in this paper, but have included them as an appendix in a technical report [3] — with 1,700 achievements, the listing is quite long.

5.1.2 Sampling Bias

As described in Section 3.1, we collected information for only characters in guilds. However, we argue that although we may not have sampled characters not in guilds, we have gathered a sample which is still representative of the entire population of characters, as represented by our comparison in Section 3.2.1. Nonetheless, it is possible that our data set is skewed towards "power-players" — with the hypothesis that the most casual of players do not join guilds. However, we believe that there are indeed a large number of "casual" guilds within the game, and that in comparison to previous attempts to sample players in WoW (e.g., posting surveys on WoW related websites) [8, 19, 22], our approach yielded data with at least no more bias.

Note that it would be possible to perform similar analyses on unguilded characters, but this would first require discovering them. To do so, one could create an add-on in the game (e.g. [9]) that logs the names of all characters seen. This technique would however still be biased to only sample characters who come into contact with the logger,

and would severely limit the amount of data that could be collected within a reasonable time period.

5.2 Ethical Implications

5.2.1 Privacy

While the data presented in this paper is highly aggregated, presenting only aggregate data over a population of millions of characters, there are underlying concerns regarding the use and distribution of the data. Each character is identified by a unique name, and although that name is not directly linkable to a real-world personality, in immersive games such as World of Warcraft, leakage of an in-game name may constitute a serious breach of privacy. However, at the same time, all data that we gathered is already publicly available through the Blizzard World of Warcraft website [6], or in game. We consulted with our institutional review board on this matter, and were notified that per our institution's policies, as long as our data does not tie back directly to real world individuals, it should not be considered human subjects research [2].

Nonetheless, as Computer Science researchers and participants in online communities, we recognize that online identities can be just as important as offline identities, and are protecting individual character and guild names. Even though the data that we used is available publicly already, the public interface is limited in analysis, and for example, does not support for the sort of aggregations that we performed in the research presented in this paper. To this end, we are more than happy to share our distributed data collector (online, at our lab's git repository: `http://code.psl.cs.columbia.edu`), and anonymized data (upon request).

5.2.2 Other Ethical Issues

Retrieving large amounts of data from APIs can potentially degrade the quality of service for other users of the API. As responsible researchers, we made sure to avoid doing so by: (1) scheduling our crawls to occur during late-night hours, (2) by immediately responding to any error message from the API servers with a simultaneous shutdown of all crawler nodes, and (3) by measuring the response time of the API service before and during our crawl, making sure that we did not impact the API's response time. Additionally, we notified Blizzard before we began this process, so as to give them a direct line of communication to us should we unintentionally cause service issues. Moreover, each API request was signed with the first author's email address and name in the "User Agent" HTTP header, making it further obvious how to contact us should the need arise.

5.3 Legal Implications

The data analyzed in this paper were collected from a public website owned and operated by Blizzard Entertainment. The API documentation website [7] provides a location where one would imagine to find a terms of service, titled "Chapter 3. API Policy," but at both the time of our crawl and the time of writing, this section simply contained the text "The final policy document is being reviewed and will be published at a later date." Since we were unable to find a terms of service or API policy document, we are led to believe that the data that we gathered was in the public domain, and free to be analyzed and published as we may desire. Nonetheless, we wish to avoid what may be perceived

to be irresponsible behavior on the part of Blizzard, and are *not* publicly posting the data set, but rather will only release anonymized portions of the data upon request.

6. FUTURE WORK

We have only begun to analyze the vast quantities of player data that we have collected from World of Warcraft. We have not began to study other potentially interesting dimensions that we constructed from Blizzard's data — quests, raid progress, professions, mounts, companions, and reputation. It would be worthwhile to investigate means to associate multiple characters who are owned by the same player. We would also like to work to make the data collected publicly available, as long as doing so remains ethical and legal.

7. CONCLUSION

In this paper, we presented what we believe to be the first truly large-scale and longitudinal study of World of Warcraft players. While previous work has been based on qualitative research over thousands of users, or quantitative research over hundreds of thousands of characters, we analyzed over six million characters. We presented a brief description of the tool that we used to crawl the data as well as analyses showing that we can cluster characters based on their in-game actions. We showed that each character's play profile can evolve during the course of play, and finally, we broke time to level down by character motivation, showing that players who play WoW differently level at different rates. Our data show that there is room for new and novel research in game studies by using massive data sets gathered through public APIs, and suggests topics for future academic study.

9. REFERENCES

[1] M. Ashton and C. Verbrugge. Measuring Cooperative Gameplay Pacing in World of Warcraft. In *FDG'11: Proceedings of the Sixth International Conference on Foundations of Digital Games*, June 2011.

[2] A. Barry. Personal communication, November 2011.

[3] J. Bell, S. Sheth, and G. Kaiser. A Large-Scale, Longitudinal Study of Player Motivations in World of Warcraft. Technical Report CUCS-046-11, Columbia University, 2011.

[4] Blizzard Entertainment. Patch 3.0.2 - Game - World of Warcraft. `http://us.battle.net/wow/en/game/patch-notes/3-0-2`.

[5] Blizzard Entertainment. Patch 4.0.3 - Game - World of Warcraft. `http://us.battle.net/wow/en/game/patch-notes/4-0-3`.

[6] Blizzard Entertainment. World of Warcraft. `http://us.battle.net/wow/en`.

[7] Blizzard Entertainment. World of Warcraft Armory API. `http://blizzard.github.com/api-wow-docs/`.

[8] T. Debeauvais, B. Nardi, D. J. Schiano, N. Ducheneaut, and N. Yee. If you build it they might stay: retention mechanisms in World of Warcraft. In *Proceedings of the Sixth International Conference on the Foundations of Digital Games*, New York, NY, USA, June 2011. ACM.

[9] N. Ducheneaut, N. Yee, E. Nickell, and R. J. Moore. Building an MMO with mass appeal: A look at gameplay in World of Warcraft. *Games and Culture*, 1:281–317, 2006.

[10] B. Harrison and D. L. Roberts. Using Sequential Observations to Model and Predict Player Behavior. In *Proceedings of the Sixth International Conference on the Foundations of Digital Games*, New York, NY, USA, June 2011. ACM.

[11] K. Hullett, N. Nagappan, E. Schuh, and J. Hopson. Data analytics for game development (NIER track). In *Proceeding of the 33rd International Conference on Software Engineering*, ICSE '11, pages 940–943, New York, NY, USA, 2011. ACM.

[12] M. Jakobsson and T. L. Taylor. The Sopranos meets EverQuest: Socialization processes in massively multi-user games. In *FineArt Forum*, volume 17, 2003.

[13] C. Lewis and N. Wardrip-Fruin. Mining game statistics from web services: a World of Warcraft armory case study. In *Proceedings of the Fifth International Conference on the Foundations of Digital Games*, FDG '10, pages 100–107, New York, NY, USA, 2010. ACM.

[14] G. B. Orgaz, M. D. R-Moreno, D. Camacho, and D. F. Barrero. Clustering avatars behaviours from virtual worlds interactions. In *Proceedings of the 4th International Workshop on Web Intelligence & Communities*, WI&C '12, pages 4:1–4:7, New York, NY, USA, 2012. ACM.

[15] Sony Online Entertainment. EverQuest. `http://www.everquest.com/`.

[16] T. L. Taylor. *Play between worlds: Exploring online game culture*. The MIT Press, 2006.

[17] T. L. Taylor. Does World of Warcraft Change Everything? How a PvP Server, Multinational Playerbase, and Surveillance Mod Scene Caused Me Pause. *Digital Culture, Play, and Identity: A*, pages 187–201, 2008.

[18] WarcraftRealms.com. Warcraft Realms Data Feeds. `http://www.warcraftrealms.com/dataexport.php`.

[19] D. Williams, N. Yee, and S. E. Caplan. Who plays, how much, and why? Debunking the stereotypical gamer profile. *Journal of Computer-Mediated Communication*, 13(4):993–1018, July 2008.

[20] WoWWiki. Patches - WoWWiki - Your guide to World of Warcraft. `http://www.wowwiki.com/Patches`.

[21] N. Yee. The demographics, motivations, and derived experiences of users of massively multi-user online graphical environments. *Presence: Teleoper. Virtual Environ.*, 15:309–329, June 2006.

[22] N. Yee. Motivations for Play in Online Games. *CyberPsychology & Behavior*, 9(6):772–775, Dec. 2006.

[23] L. Zezula. MPQ Archives. `http://www.zezula.net/en/mpq/download.html`.