McGill University
School of Computer Science
Game Research at McGill

Measuring the Impact of Cooperation in Halo: Reach

GR@M Technical Report No. 2012-1

Martin Ashton Clark Verbrugge
School of Computer Science, McGill University
Montréal, Québec, Canada, H3A 2A7
martin.ashton@mail.mcgill.ca clump@cs.mcgill.ca

August 21, 2012

Abstract

Social aspects of multiplayer games are well known as contributors to game success, with online friendships and socialization expected to expand and strengthen a player-base. Understanding the nature of social behavior and determining the impact of cooperation on gameplay is thus important to game design. Here we make use of the data exposed through the public interface of a contemporary first-person shooter, HALO: REACH, originally intended to foster player community. We use this to investigate the extent of cooperation among player groups and the effect on individual player behavior. Our results demonstrate that players who enter as a group into the multiplayer matchmaking system have, on average, a significantly higher win-to-loss ratio than players who enter the matchmaking system alone. Principal component analysis of individual behaviors reveals a set of novel player types, adapted to the multiplayer context and quite distinct from player types found in other contexts. These data illustrate how player characteristics are strongly related to the actual game context, and show the impact of cooperative behavior on game design.

1 Introduction

Multiplayer games benefit from strong social engagement; a large and cohesive player community is essential for providing players with an abundance of other players to act as opponents or team-mates, and is also seen as mechanism to encourage player retention. Understanding the nature and impact of player behaviors in multiplayer game contexts is thus important to game growth as well as how game environments may be structured or tailored to encourage community.

In this work, we examine the impact of cooperation among players and attempt to model individual player behavior in the context of a contemporary multiplayer first person shooter (FPS). We focus our attention on HALO: REACH, a FPS whose relevance in the genre is reflected by its classically-rooted game mechanics, and whose popularity is stated by its 1.3 billion games played within its first four months of release [6, 1]. Our approach makes use of the growing trend by game studios to offer web-portals for users to display and compare personalized gameplay data; examples of such portals include Battlefield 3's Battlelog [10], the World of Warcraft home page [5], the Bungie.net Stats page [7] and the Steam Community page [22]. Although intended to foster game community [19], these systems include a wealth of user data, such as achievements, rank, total kills, total deaths, wins and score, and so provide interesting opportunities to analyze large volumes of real-world gameplay data [15, 11].

Our investigation builds on the HALO: REACH Stats API, which allows us to crawl a representative subset of the HALO: REACH player base and investigate player-player interactions. The edge-weights of the resulting graph provide insight into the level of coordination among players. We discover that in the average case, the win-to-loss ratio of players who enter as a cohesive group into the multiplayer matchmaking system scales in proportion to the size of the group.

We further use the data to examine the nature of individual players in this multiplayer context. We apply an exploratory principal component analysis (PCA) to the aggregate set of player gameplay performances, isolating specific sets of player behaviors. This step reveals a set of five stable components which can be used to describe a player's behavior, and demonstrate a relatively unique set of player types, distinct from the player types found in more general game contexts [4, 2, 23].

Specific contributions of this work include:

- The development of a software system which we use to sample the body of HALO: REACH gameplay data via the HALO: REACH Stats API.

- The inference of coordination among player groups by isolating tight-knit components in the player subgraph. This data shows the size and extent of game "friendships," as well as the impact on game balance of cohesive groups.
- The use of principal component analysis (PCA) on FPS multiplayer data in an attempt to categorize individual player behavior in this genre. Our results reveal interesting and novel player types strongly adapted to the game genre.

The next section of this paper provides background information on HALO: REACH. Section 3 describes the system we developed to collect data from the HALO: REACH API. Section 4 presents our data analysis of the player subgraph, and the PCA results of the gameplay data. Section 5 presents related work in the domain of game metrics, and we conclude with Section 6.

2 Background

2.1 Halo: Reach

HALO: REACH is a recent installment of the HALO franchise, which has established itself as a popular series in the console FPS genre [18]. The game features a single-player campaign, as well as an array of multiplayer modes including campaign coop, free-for-all and team games. The measurements performed in our study only pertain to competitive team-based game types; in this context a match-making system allows players to form teams, with results from the ensuing gameplay published to the post-game lobby.

Matchmaking. The HALO: REACH matchmaking system allows players to be matched with one another according to a desired *playlist.* A playlist in this matchmaking system can be understood as a set of possible game modes. For example, the "Team Objective" playlist organizes players into two teams of size 4 and places them into an objective-based game type such as capture-the-flag or king-of-the-hill. Larger team sizes are also permitted in other playlists, such as "Big Team Battle," which pits two teams of up to 8 players against each other.

The matchmaking system is provided through the Xbox Live network service associated with the console. This service allows players to organize and maintain lists of friends, which can then be used to join the HALO: REACH matchmaking system as a cohesive group, maximizing the ability of a set of friends to play together. The service also provides persistent player identities through a unique *gamertag* assigned to each service account. It is possible to play with up to three guests on the same console, in which case these players are identified as guests of the host gamertag. Guests are restricted to playing in non-competitive playlists, and as such, our metrics do not encompass guest gameplay data.

Outside of group entry, the matchmaking system uses the Xbox Live "TrueSkill" ranking system to match players of similar skill together [16]. In contrast to the well-known ELO rating system used in Chess, the TrueSkill system makes use of a Bayesian update mechanism to estimate a lower bound on a player's skill, iteratively narrowing the associated uncertainty of this skill estimate from repeated gameplays [12]. The TrueSkill ranking system updates its results based on the final outcome of the game and operates under the assumption that this outcome reflects the skill value of each participant. The TrueSkill system then matches teams according to the summed skill-rating calculated for each team.

Gameplay. The most popular and competitive playlists involve two teams of four players, for a total of eight players in a game. Each player then begins by selecting a preset *loadout,* which dictates his or her two starting weapons. Loadouts also define the player's special ability such as: invisibility and radar-jamming, sprinting, jump-jetting, temporary invincibility, and dodging. Loadout presets can also be changed while the player is respawning.

Highly competitive playlists usually provide players with a default medium-range rifle capable of precise headshots regardless of the selected loadout. Players can only carry up to two weapons at a time, which can be swapped for more powerful weapons found on the game-map, such as sniper rifles and rocket launchers. Players also have shields which limit damage taken, but must recharge a few moments after taking damage. If a player's shields are depleted, damage will be taken by enemy fire and must be regenerated by consuming medkits found on the map; shieldless, several weapons also become capable of dispatching him or her with a single headshot. A player who is killed is temporarily removed from gameplay, and must wait before respawning. The game ends if a team obtains the required number of kills, or if the objective has been captured a set amount of times, or if time runs out.

Post-game lobby. Once a game finishes, the players are placed in the post-game lobby. A summary of each player's performance is presented here, which includes his or her number of kills, deaths, assists, betrayals, headshots, as well as a detailed account of the player's earned medals, derived from achieved headshots as well as many other possible player feats, such as double-kills, assisted kills, and killing sprees. Medals do not have a direct in-game effect, and are mainly used to encourage competitive social display and provide players with more difficult and long-term goals. As we will show later, they can also serve as a means of measuring and understanding behavior. Medals are broken down by the game system into four basic categories:

- **Multi Medals** - Awarded for a series of kills within four seconds of each other.
- **Spree Medals** - Awarded for a number of kills in a row without dying.
- **Style Medals** - Awarded for feats, such as headshots, assists, and assassinations.
- **Other Medals** - Awarded for additional feats, such as sniper-rifle headshots, melee kills, or hitting an opponent with a vehicle.

In addition to the Xbox Live friends list, players who enjoyed playing together can opt to form a party from the post-game lobby. This mechanism allows for new, online friends to be discovered, and these newly-grouped players can continue playing on the same team in subsequent matches formed by the matchmaking system.

2.2 Stats API

Individual HALO: REACH game statistics are hosted on the Bungie.net Stats page. For each game, users can browse the information provided in the post-game lobby and can view a temporal progression of events on the map.

The majority of this data can also be accessed via calls to the HALO: REACH Stats API, allowing development-oriented players to aggregate or explore the data. To use this API, a valid Xbox Live account is required, along with a subscription to the Bungie Pro service. The main purpose of Bungie Pro is to provide users with the functionality to record, render and download high-quality, in-engine videos. Bungie Pro also supplies developers with API keys to pipe HALO: REACH game statistics through their web servers or applications. The endpoints of the HALO: REACH Stats API are compatible with PHP and .NET. Client applications written with the .NET framework must reference a Windows Communication Foundation (WCF) service to obtain the API class definitions; in either case a valid API key must be included with each method call, and a rate limit of 300 requests per minute is enforced.

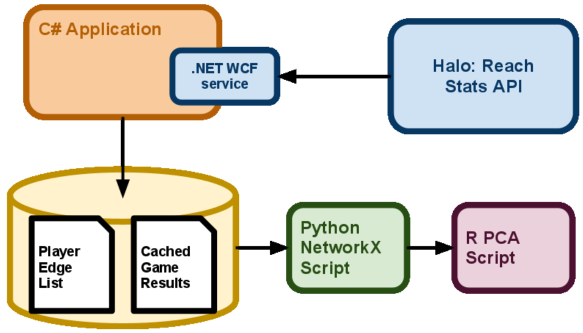

Figure 1: Halo: Reach multiplayer data analysis system.

3 Data Gathering

In this section, we outline the framework we designed and used to gather and process player data. The overall design of our data gathering framework is illustrated in Figure 1. A program written in C# uses the .NET endpoint of the HALO: REACH Stats API to gather post-game player data, which is then cached for processing by network and statistical analysis libraries.

Discovering random player names is difficult, and so our approach is based on crawling through players connected by common gameplay. A round of data-collection begins from a chosen seed player, whose gamertag is added to a *visited players* list. This introduces a potential bias from the choice of seed, which we address below by evaluating multiple data sets. We inspect a player's 75 most recent competitive team games, assigning unique game-IDs to avoid duplicate game inspection. We enqueue newly-encountered gamertags found in each game and update an undirected, weighted edge list to identify the players who have played together in the same game. We also inspect each player's performance per game, defined by a row vector containing his or her score, team standing, individual standing, kills, deaths, assists, suicides, betrayals, number of medals earned per category, unique medals earned per category, average death distance, average kill distance, and number of headshots.

The next iteration pops a gamertag from the queue and adds it to the list of *visited players*. We repeat the process of game inspection, edge-list maintenance and game performance caching for each player added to the *visited players* list. HALO: REACH is a popular game, and so continuing this process until a closed set is found is not feasible, at least not given the rate-limits imposed by the API. We thus terminate rounds of data-collection manually, and consider the impact of data-set size in our analysis below.

The Python NetworkX module is used to construct a player graph from the weighted edge list [17]. Network metrics are calculated on this graph to plot node-degree distribution, edge-weight distribution and average component size. We also aggregate the cached gameplay performances of each player to compute his or her average performance per game. This data is processed by R's principal component analysis (PCA) algorithm, in particular using singular value decomposition on the centered and uniformly scaled column values [13]. The resulting output is a list of orthogonal vectors (principal components) whose linear combination can be used to describe the individual performance of each player according to a new basis. Each principal component accounts for a specific proportion of the variance in the input data. The use of PCA

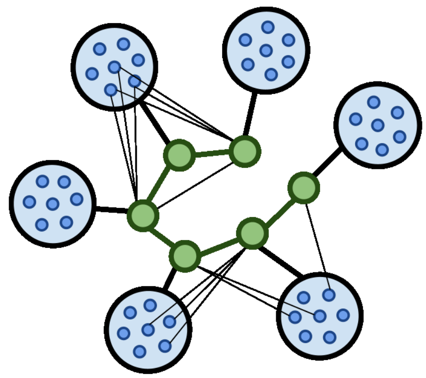

Figure 2: A *conceptual simplification* of our network graph; the green nodes represent the visited players, while the blue nodes represent the satellite players.

in this experiment aims at reducing the number of initial variables used to qualify a player's performance (score, team standing, individual standing, kills, etc) into a smaller set of correlated variables.

4 Data Analysis

The results analyzed in this section are based on four rounds of data collection, yielding a total of 6 700 unique visited players and 384 124 uniquely inspected games. The first three rounds of data collection follow a breadth-first search from three different seed gamertags. The fourth round of data collection uses the same seed as the first round of data collection, but a random walk is used instead of a breadth-first sampling approach. Our concern is that a breadth-first crawl would yield a biased set of visited players alike in skill and behavior. However, the similarities between network topologies and PCA results support the equivalence of both sampling methods.

4.1 Player Graph

We construct an undirected graph from the player edge list. The set of nodes in this graph corresponds to the set of all players encountered in the inspected games. The edges in this graph represent the fact that two players have played in the same game, with the weight of the edge denoting the number of times a pair of connected players have played games together.

Our data crawl is artificially truncated to maintain feasibility of the crawling process, and so not all players we encounter are fully explored. As illustrated in Figure 2, a node in the graph can thus belong to one of two sets:

1. **Visited Players**: the set of players whose recent games have been inspected and used to further crawl the player base.
2. **Satellite Players**: the set of players who have appeared in at least one inspected game, but who have not been included in the set of visited players.

Figure 3 presents the node degree distribution of the player graph constructed from the third round of data collection, which encompasses 2 443 visited players and 3 812 330 edges. Note that the total of 411 877

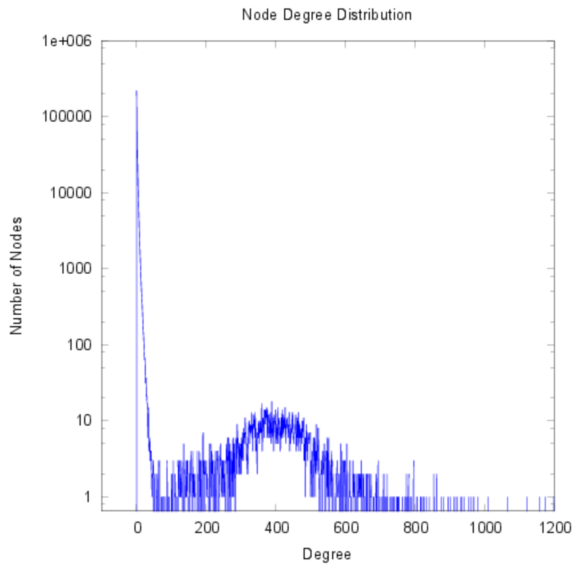

Figure 3: The *node degree distribution* of players from our third (largest) round of data collection, plotting number of nodes versus node degree. The numerous nodes with low degree correspond to the satellite players, while the relatively few nodes with high degree correspond to the visited players.

nodes in this graph far exceeds the number of visited players—the sharp peak on the left corresponds to the satellite players, while the hump near the middle corresponds to the set of visited players. Each visited player is necessarily connected to each player appearing in his or her inspected games, while each satellite player is connected to at least one visited player in addition to the other players appearing in the same inspected game(s). A satellite player can be connected to more than one visited player, in particular if he or she is friends with several visited players, but not to other satellite players.

By construction, the player graph is initially defined as one large component. That is to say, there exists at least one path from one node to every other node in the graph. To evaluate the impact of cooperation among the players in this data set, we begin by identifying the players who play together often. We simplify the terminology by labeling a pair of players as "friends" if the two have played in at least eight games together. This threshold is justified by an analysis of component size built by progressively pruning edges of increasing weight. The results are shown in Figure 4, and strongly indicate a stabilizing point of around eight. Given the size of the HALO: REACH player-base, it is unlikely that two strangers will be matched into the same eight games regardless of their similarity in skill rating. Since players can opt to continue playing together by forming a party in the post-game lobby, our use of the word "friend" does not necessarily denote a friendship in the Xbox Live service. As evidenced by the capacity to reliably enter the same game, friendship implies a level of coordination between two players. We assume that such a friendship also implies a willingness to cooperate.

Our interest in identifying friendship groups is based on a hypothesis that friendship has an impact on team success in competitive situations. Since friendships can be formed based on successful, randomly-formed competitions, the causal direction of such a relation is not trivial to disentangle, but is interesting from either perspective as a potential factor in game design and balancing. Friends can enter the matchmaking system as a party of up to eight players, although the majority of competitive playlists only allow a maximum party-size of four. For each visited player in our merged data set, we compound his or her wins and losses with respect to the number of his or her friends in the game. Figure 5 shows the increasing progression

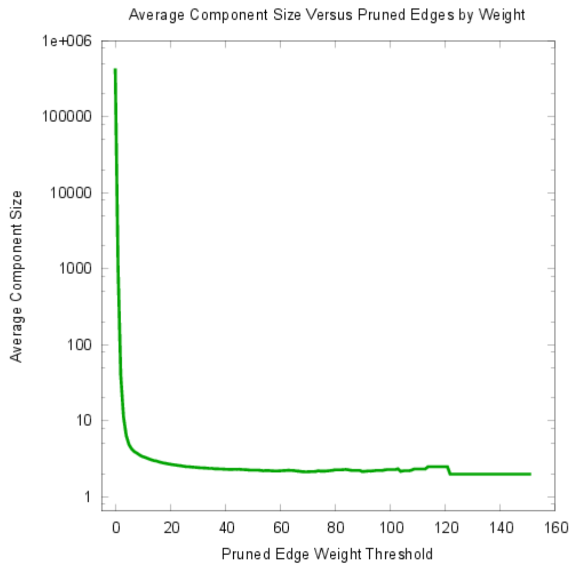

Figure 4: *Average Component Size.* The graph begins as one large component. As edges are pruned, the average component size quickly stabilizes to a value less than 4. Based on these results, we choose a weight threshold of 8 to qualify player pairs as friends.

of the average win/loss ratio in relation to the number of friends in the game. The relation is overall very clear and relatively linear, with the dip at $x = 4$ and $x = 5$ explained by the jump from playlists only allowing a maximum of four players per party, to those accepting a maximum of eight players per party. With four or five friends involved, a team must be composed of eight players, but is not yet fully dominated by friendships, and so still requires significant and necessarily suboptimal coordination with strangers.

4.2 PCA

The second part of our experiment aims to isolate and categorize individual player behaviors found in our sample. Competitive, team-based play encourages aspects of cooperation, but also provides the potential for different play styles through the variety of available loadouts and weapons, such as stealthy guerilla tactics, long-range sniping or close-range assault. It is also worth noting that the real-time and dynamic nature of multiplayer gameplay does not necessarily allow players to distinctively exhibit their preferential play styles—a failed assassination cannot be retried by reverting the game to a previous state, as is the case with most single player games. Thus, a strict adherence to a specific play style may be discouraged in favor of coordinated team play.

Our analysis is based on the average performance per game of each visited player. We considered as much of the individual data available per player as possible, applying a principal component analysis as detailed in Section 3. This analysis yields a set of five major principal components whose cumulative proportion of variance is valued at 79%. In other words, the weighted linear combination of these five vectors accounts for 79% of the variance among the 6 700 player performances. We present a simplification of these principal components in Table 1 to illustrate the correlation between performance metrics.

Players who have a large PC1 weight show an overall aptitude for and an interest in obtaining medals, attesting to their comfort in varied gameplay scenarios.

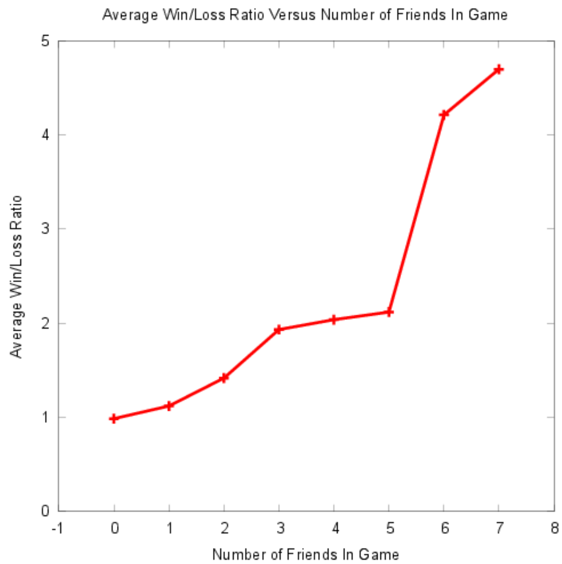

Figure 5: *Average Win/Loss Ratio Versus Number of Friends In Game.* The average win/loss ratio is 0.99 for players whose games contained 0 of their friends, indicating a balanced number of wins and losses. This ratio rises to 1.97 for players whose games contained 3 of their friends, and to nearly 5 for players playing with 7 friends, illustrating the strong benefit of repeated cooperation.

A large positive weight along PC2 corresponds to a player's ability to perform headshots. It is worth noting that this variable negatively correlates with assists. Intuitively, players who achieve many headshots are awarded with the killing blow and not the assist. The positive correlation between headshots and kill/death distance is explained by the longer effective range of weapons capable of headshots.

Players with a positive weight for PC3 are better suited to obtain assists, while a positive weight along PC4 indicates a disposition towards long-range engagements. In this later component, the average kill/death distance appears to be correlated with the spree medals, but oppositely correlated to score. Note that objective-based games increase a player's score according to each objective captured and not according to his or her number of kills. As such, players who are far removed from the action are also less likely to score objective points.

The Standing metric is a binary ranking (lower values meaning higher ranks), and indicates whether the player's team has won (Standing = 0) or lost (Standing = 1). PC5 is focused on this metric. Thus, a player with a large negative value along PC5 indicates that the player's team loses more often (Standing closer to 1). By contrast, a player with a more positive weight for PC5 indicates that the player's Standing is closer to 0, and so implying that his or her team wins more often. In PC5, Standing is also correlated with Individual Standing, a more fine-grain ranking value that measures the player's overall performance in a game compared to the other player scores, regardless of team. An Individual Standing of 0 indicates the first-place position, while a value of 7 denotes the last place position in a game of 8 players. Observe that in PC5 these two measures are both correlated, and so players whose teams lose more often (Standing closer to 1) would tend to have larger Individual Standings (ie worse, or larger individual rankings).

The correlation of the variables within the principal components reveals that player types roughly correspond to the use of the available game mechanics of HALO: REACH. This is in contrast to broader player types found in other game contexts [4, 2, 23], and in partial contradiction to our initial expectations that both assistive "heart" types and traditional "griefing" or "club" behaviours [4] would be the most strongly

Performance Metric	PC1 (38%)	PC2 (14%)	PC3 (12%)	PC4 (9%)	PC5 (6%)
Games Played					
Score	+			-	
Standing					(-)
Individual Standing	-				-
Kills	+				
Dying		-	-	-	-
Assists		-	(+)		
Suicides					
Betrayals		-	-		
Multi Medals	+		-		
Other Medals	+	-		+	
Spree Medals	+			+	
Style Medals	+	+		-	
Total Medals	+				
Unique Multi Medals	+		-		
Unique Other Medals		-	+	+	
Unique Spree Medals	+			+	
Unique Style Medals	+		+	-	
Total Unique Medals	+		+		
Avg Death Distance		+		(+)	
Avg Kill Distance		+		(+)	
Headshots	+	(+)			

Table 1: *Simplified Principal Components.* We conserve the sign of the coefficients whose absolute value is greater than half of the absolute value of the largest coefficient. Opposite signs indicate opposite variable correlation. The percentage beneath each principal component (PC) represents its proportion of variance. Bracketted signs indicate an exceptionally large absolute value for a coefficient and thus a dominant behavior.

represented. Game types are instead dominated by skill in general (PC1), followed by those more focused on relatively distinct, but comparable play styles—headshots (PC2), wounding (PC3, which implies some amount of cooperation), distance kills and sprees (PC4), and a facility towards teamplay (PC5).

5 Related Work

A variety of other studies and efforts have been developed that aim to evaluate players or player behaviors, both during game design and as part of post-facto analysis. In prior work related to human computer interaction, for instance, the authors of [14] have developed the TRUE (Tracking Real-Time User Experience) system to collect streams of timestamped *user initiated events*, ranging in context from spreadsheet applications to video games. The TRUE system was used in the development of Halo 2 to isolate problem areas of the single-player campaign. A similar metric-oriented approach has been used in the development of Halo 3, namely to determine the fairness of multiplayer maps [20].

The process of collecting gameplay metrics in order to complement game design is furthermore illustrated in [8], where the authors describe the use of their own data collection system, *Tracktivity,* during the testing phases of the racing game *Split/Second*. Events triggered by players were monitored by Tracktivity to determine the difficulty of a race-track and to help designers tailor the game's overall learning curve. The authors of [9] additionally attest to the benefit of game metrics in conjunction with the game design of *Kane*

& Lynch and *Fragile Alliance*, two modern first person shooters.

The alternative direction has also been explored. The effect of level design on player behavior is investigated in [21], in which the authors use game metrics gathered from the stealth shooter game, *Hitman: Blood Money*. Measurements including character position, rotation angle, actions performed and narrative choices are analyzed to qualify a player's overarching behavior, or *persona*. The game can thus label a player as either a *mass murderer*, a *silent assassin*, a *mad butcher* or *the cleaner*, depending on the prevailing nature of his or her actions.

Quantitative assessments of player behavior have also been performed in the context of massively multi-player online games, in particular using World of Warcraft (WoW) [3, 15]. Recent work detailed in [24], for example, made use of survey data and behavioral metrics to identify the expression of psychology-based personality traits among players in WoW. Participants of this study filled a questionnaire to determine their alignment along the "Big-5" personality traits which include Extraversion, Agreeableness, Conscientious-ness, Emotional Stability and Openness to Experience. To measure the behavior of each participant, a web crawler was developed to aggregate per-character achievement data in the WoW Armory over a period of 4 months. Personality traits were then shown to have significant correlations with various aspects of gameplay, suggesting that virtual worlds are an appropriate medium in which to infer personality traits. The sequential evolution of achievement-data mined from the WoW Armory has additionally been used to predict player behavior and character progression in [11].

6 Conclusions and Future Work

Although not directly aimed at facilitating user analysis, vendor-sponsored game community sites provide a wealth of user data. Our use of the HALO: REACH Stats API has allowed us to quantitatively evaluate the cooperative behavior of FPS players in a competitive, team-based environment. The network structure of the sampled player set reveals an interesting correlation between winning and the number of online relations. It is not suprising that friends exhibit better team coordination, but our findings suggest that multiplayer game balance could be improved by matching competitive teams based on cohesion or some other group ranking metric. A further inspection of individual player behaviors in this multiplayer FPS context has yielded a set of five relatively stable descriptors tied to the underlying game mechanics of HALO: REACH; the intensely competitive setting tends to focus player types on success, which can be achieved either in general or through specializations.

A deep understanding of actual player demographics, as well as the interplay of mechanics and player be-haviours within a given game environment has an abundance of potential uses. Our own interests are aimed at developing techniques for dynamically measuring and thus adapting gameplay as a means of improving the overall player experience. A specific direction evolving from this work is to develop a suitable dy-namic, online system that can usefully measure team cohesion. Such a metric would make an interesting and potentially useful adjunct to the skill estimation systems currently in use by popular team-based games.

7 Acknowledgements

This research was supported by the Fonds de Recherche du Quebec - Nature et Technologies. Special thanks are also extended to the talented people at Bungie Studios.

References

[1] L. Alexander. Gamers Hit 1.3 Billion Halo: Reach Matches. http://www.gamasutra.com/view/news/32111/Gamers_Hit_13_Billion_Halo_Reach_Matches.php, Dec. 2010.

[2] A. Alix. Beyond P-1: Who plays online? In *Proceedings of DiGRA 2005 Conference: Changing Views–Worlds in Play*, 2005.

[3] M. Ashton and C. Verbrugge. Measuring cooperative gameplay pacing in world of warcraft. In *FDG'11: Foundations of Digital Games Proceedings*, FDG '11, New York, NY, USA, 2011. ACM.

[4] R. Bartle. *Designing Virtual Worlds*. New Riders, 2003.

[5] Blizzard Entertainment. World of Warcraft. http://us.battle.net/wow/en//, Dec. 2011.

[6] Bungie, Inc. bungiestats.jpg. http://www.bungie.net/images/News/Inline10/121710/bungiestats.jpg, Dec. 2010.

[7] Bungie, Inc. Halo Reach Online. http://www.bungie.net/stats/reach/online.aspx, Dec. 2011.

[8] E. J. Chapresto, K. Mitchell, and F. J. Seron. Capture and Analysis of Racing-Gameplay Metrics. *IEE Software*, 28(5), 2011.

[9] A. Drachen and A. Canossa. Towards gameplay analysis via gameplay metrics. In *Proceedings of the 13th International MindTrek Conference: Everyday Life in the Ubiquitous Era*, MindTrek '09, pages 202–209, New York, NY, USA, 2009. ACM.

[10] Electronic Arts. Battlefield 3 Battlelog. http://battlelog.battlefield.com/bf3/servers/, Dec. 2011.

[11] B. Harrison and D. L. Roberts. Using sequential observations to model and predict player behavior. In *FDG'11: Foundations of Digital Games Proceedings*, FDG '11, New York, NY, USA, 2011. ACM.

[12] R. Herbrich, T. Minka, and T. Graepel. Trueskill(tm): A bayesian skill rating system. In *Advances in Neural Information Processing Systems 20*, pages 569–576. MIT Press, 2007.

[13] I. Jolliffe. *Principal Component Analysis*. Springer-Verlag, 2nd edition, 2002.

[14] J. H. Kim, D. V. Gunn, E. Schuh, B. Phillips, R. J. Pagulayan, and D. Wixon. Tracking real-time user experience (true): a comprehensive instrumentation solution for complex systems. In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, CHI '08, pages 443–452, New York, NY, USA, 2008. ACM.

[15] 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.

[16] Microsoft Research. Trueskill(tm) Ranking System: Details. http://research.microsoft.com/en-us/projects/trueskill/details.aspx, Dec. 2011.

[17] NetworkX Developers. NetworkX. http://networkx.lanl.gov/, Nov. 2011.

[18] C. Nutt. A decade on, halo charts its course. `http://www.gamasutra.com/view/feature/6365/a_decade_on_halo_charts_its_course.php`, May 2011.

[19] S. Shen and A. Iosup. The XFire Online Meta-Gaming Network: Observations and High-Level Analysis. In *IEEE HAVE 2011 Proceedings*, HAVE 2011, 2011.

[20] C. Thompson. Halo 3: How Microsoft Labs Invented a New Science of Play. *Wired Magazine*, 09(15), 2007.

[21] A. Tychsen and A. Canossa. Defining personas in games using metrics. In *Proceedings of the 2008 Conference on Future Play: Research, Play, Share*, Future Play '08, pages 73–80, New York, NY, USA, 2008. ACM.

[22] Valve Corporation. Steam Community. `https://steamcommunity.com/`, Dec. 2011.

[23] N. Yee. Motivations for play in online games. *CyberPsychology and Behavior*, (9):772–775, 2007.

[24] N. Yee, N. Ducheneaut, L. Nelson, and P. Likarish. Introverted elves & conscientious gnomes: the expression of personality in world of warcraft. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 753–762, New York, NY, USA, 2011. ACM.