

The Life and Death of Online Gaming Communities: A Look at Guilds in World of Warcraft

Nicolas Ducheneaut¹, Nicholas Yee², Eric Nickell¹, Robert J. Moore¹

¹Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304 USA

{nicolas, nickell, bobmoore}@parc.com

²Stanford University
Department of Communication
Stanford, CA 94305 USA
nyee@stanford.edu

ABSTRACT

Massively multiplayer online games (MMOGs) can be fascinating laboratories to observe group dynamics online. In particular, players must form persistent associations or “guilds” to coordinate their actions and accomplish the games’ toughest objectives. Managing a guild, however, is notoriously difficult and many do not survive very long. In this paper, we examine some of the factors that could explain the success or failure of a game guild based on more than a year of data collected from five World of Warcraft servers. Our focus is on structural properties of these groups, as represented by their social networks and other variables. We use this data to discuss what games can teach us about group dynamics online and, in particular, what tools and techniques could be used to better support gaming communities.

Author Keywords

Online communities, Massively Multiplayer Online Games, social networks, group dynamics, data analysis tools

ACM Classification Keywords

H.5.3. [Collaborative Computing]: online games

INTRODUCTION

Massively Multiplayer Online Games (MMOGs) are now hosting millions of players in their rich 3D virtual worlds. These games are collaborative by design [23]: players often have to band together to accomplish the game’s objectives, and trading items and information is essential to a player’s advancement [17]. This need for repeated collaboration translates into formal, persistent groups that are supported out-of-the box by nearly all MMOGs: guilds.

Guilds are essential elements in the social life of online

gaming communities. Guild members have access to simple tools to coordinate with each other. Most commonly these include an in-game roster showing who is currently logged on and a private chat channel to broadcast messages to them. Guilds frame a player’s experience [20] by providing a stable social backdrop to many game activities, and their members tend to group with others more often and play longer than non-affiliated players [9]. At the “high-end” of a game, guilds can even become indispensable: “raids” requiring coordination among up to 40 players are essential to advancement and it is almost impossible to assemble a pick-up group of this size – some formal coordination mechanisms are required, and the guilds provide such an environment. Being a member of an “elite” or “uber” guild, renowned for its ability to tackle the hardest challenges, is therefore a badge of honor. Admission to these prestigious social groups often requires going through a “trial period”, as well as being sponsored by one of the members [23].

But overall, guilds are incredibly diverse. Some are small groups with pre-existing ties in the physical world and no interest in complex collaborative activities. Others are very large, made up mostly of strangers governed by a command-and-control structure reminiscent of the military. In previous work, we have explored the range of possibilities between these two extremes and documented the motivations that lead players to guilds of one type or the other [26]. Across all types, one trend was particularly clear: guilds are fragile social groups, and many do not survive very long (see also [9]).

This fragility is almost certainly due to a broad combination of factors. Leadership style, for instance, is often cited by players [26]. Game design is another contributor: players “burn out” due to the intense “grind” required to advance in MMOGs [29] and leave the game, abandoning their guild at the same time. “Drama” (public conflict between two or more guild members) and internal politics (e.g., arguments over who gets access to the most powerful “loot” dropped by monsters) have also been the demise of many guilds. All these factors and many others have been documented in the aforementioned previous works.

One set of factors, however, remains unexplored: the structural properties of these groups. Indeed, many

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2007, April 28-May 3, 2007, San Jose, California, USA.
Copyright 2007 ACM 978-1-59593-593-9/07/0004...\$5.00.

variables influence a guild's composition. Is it mostly made up of high-level players or beginners? Are many classes (e.g. warriors, priests) represented or do the members favor a particular one? Is there any formal organization or are the members partnering in an essentially ad-hoc fashion? The list could go on much longer, but it is reasonable to hypothesize that some aspects of the structure of a guild contribute to its eventual success, just like the structure of any organization plays a role in its efficiency [16].

To explore this aspect of the social life of guilds in more detail we therefore decided to use data from our ongoing study [9] of World of Warcraft (WoW), the most popular US-based MMOG so far with more than 8 million subscribers [6]. Our quantitative observations allow us to compute "social accounting metrics" [5] that reflect the structural properties of guilds and their possible impact on a group's survival in the long term. Our approach is inspired by a well-established line of research at CHI that seeks to measure the structural properties of online communities, with the hope of eventually increasing their navigability and the enjoyment of their members [e.g. 21, 25]. We use our results to discuss what gaming communities can teach us about the social dynamics of online groups, as well as the potential for creating new tools to help understand and manage these unique online social spaces.

METHODS

The use of quantitative data for social science research, at CHI or elsewhere, is often criticized for ignoring the rich, qualitative context that the metrics emerge from. Before presenting our analyses it is therefore worth mentioning that, as serious gamers and researchers, we have been observing social interactions in MMOGs "from the inside" for several years. For this paper, all the authors have accumulated hundreds of hours of play time in World of Warcraft, getting exposed in the process to a very broad palette of social experiences. We have all joined guilds, big and small, successful and doomed to failure, since the launch of the game in November 2004. This deep, personal experience with the game's environment frames our analyses and allows us to make sense of our numbers in a contextualized manner.

Our current project and its approach was influenced in great part by an interesting design choice made by Blizzard Entertainment, producers of WoW. Indeed, WoW was built such that its client-side user interface is open to extension and modification by the user community. Thanks to this open interface, we have been able to develop custom applications to collect data directly from the game. For this study we rely on WoW's "/who" command, which lists the characters currently being played on a given server. Our software periodically issues "/who" requests and takes a census of the entire game world every 5 to 15 minutes, depending on server load. Each time a character is observed our software stores an entry of the form:

Alpha,2005/03/24,Crandall,56,Ni,id,y,Felwood,Ant Killers.

The above represents a level 56 Night Elf Druid on the server Alpha, currently in the Felwood zone, grouped ("y"), and part of the Ant Killers guild. Using this application we have been collecting data continuously since June 2005 on five different servers: PvE(High) and PvE(Low), respectively high- and low-load player-versus-environment servers; PvP(High) and PvP(Low), their player-versus-player equivalents; and finally RP, a role-playing server. Overall we observed more than 300,000 unique characters to date. We then used the accumulated data to compute a variety of metrics reflecting these characters' activities [9] and, in particular, the structure of their guilds.

For instance, we can easily measure the observed size of guilds (by counting the number of characters with a given tag) and track some aspects of their membership (for instance, by counting the number of characters of a given level and class). We can also get a sense of the organization of each guild by looking at their social networks. To do so, we rely on three variables: the "zone" information, the "grouped" flag, and finally the "guild" data. We assume that characters from the same guild who are grouped in the same zone are highly likely to be playing together. If so, we create a tie between them, where the strength of the tie is proportional to the cumulative time these characters have spent together. We then use the accumulated data to compute a variety of social network analysis metrics for each character and each guild, such as their centrality and density [24]. We also rely on visualization tools we developed to observe the evolution of these networks and other metrics over time (these tools are described later in the paper).

Before going any further it is important to mention some inherent limitations of our data. First, note that we are collecting information about characters, not players. Players often create several characters or "alts" (some actively played, some acting as "mules" for storage and trading). We believe however that this does not affect the validity of our analyses for two reasons: 1) our observations show that all the "alts" of a player are generally members of the same guild; 2) except for a few "alcoholics," players tend to focus on developing one character exclusively for a reasonably long stretch of time instead of constantly switching between many, simply because WoW's design makes the latter very unproductive – players cannot keep up with the "grind" required to advance and fall behind the rest of their guild. Considering that our sample periods are quite short (one month or less, see next section), it is therefore highly probable that each sample contains on average data limited to a player's current "main", their mule, and perhaps an additional "alt" leveled at the same time. Since we are looking at aggregate, guild-level structural measures, not individual patterns of behavior, this relatively uniform spread of the number of characters played at any give time should therefore not skew our analyses too much.

We also rely heavily on a character's location to construct our social networks, which is not immune to distortion. For

instance, characters are often left “AFK” (Away From Keyboard) in the game’s main cities before or at the end of a play session – their physical proximity there does not necessarily reflect any kind of joint activity. We therefore exclude cities from our sample when computing social networks. It is also entirely possible for characters from the same guild to be in the same zone and not playing together – they could each be grouped with strangers. While this can be a common occurrence in the “entry level” zones of the games that are densely populated, our experience shows this clearly tapers off as characters gain in level. We therefore believe that, while our social networking data might be a bit noisy and possibly creates more (or stronger) ties between guild members than really exist, this effect is not overwhelming.

With this in mind, we now turn to the analysis of our data.

THE DEMOGRAPHICS OF GUILDS IN WOW

Before looking at the impact of a guild’s structure on its survival, it is worth describing some high-level properties of these social groups – in particular, how their membership evolves over time. This will help us characterize some of the difficulties they face over their life cycle.

Briefly restating some data from earlier research [9], guilds in WoW tend to be quite small: the average size is 16.8, with a median of 9. The largest observed guild had 257 members. The 90th percentile of the distribution is 35, and the distribution of guild sizes over our entire sample follows a power law (see Figure 1) – a property shared by many other online phenomena [13].




Figure 1 - Distribution of guild sizes

As mentioned earlier, we also know that guilds are relatively fragile – almost a quarter of the guilds we observe at any point in time have disappeared after a month [9]. For guilds that survive, membership tends to be fluid. Starting with the 6,188 guilds in our December 2005 sample, we tabulated two rosters: a “full” roster for each guild at the beginning of the month and a “current” roster one month later. We repeated the procedure up to the July 2006 sample. Note that a character who is in the full roster but not the current one is not simply a character who was not observed towards the end of the month. For this difference to occur, they must have “deguilded” (that is, they are not bearing any guild tag) or joined another guild (they are bearing a different guild tag).

Thus for each guild, the difference between those two roster sizes is the member churn - the number of characters who were at one point in the guild but are not there any longer. Table 1 lists the average churn for guilds of different sizes. The churn percentage is around 25% and fairly stable across guilds of all sizes. In other words, if we see a guild that currently has 20 members, then over the past month, there were 5 members who have left the guild.

sizeBands	N	Minimum	Maximum	Mean	Std. Deviation
2-10	3230	0	103	2.21	4.926
11-35	1221	0	135	6.28	10.848
36-75	412	0	140	14.16	17.458
76-150	261	0	132	22.50	22.467
> 150	161	1	140	28.94	26.938

Table 1 - Mean monthly churn, by guild size

We wanted to get a sense of the pattern of migration from guilds to one another. Also, we were interested in how often people left guilds and whether this changed over the level spread. For each character over a one week sample period in August 2006 (131,984 characters), we calculated the following variables:

- 1) Unguild Event - for each time a character is observed in a guild in snapshot X but not observed to be in a guild in snapshot X+1, we increment their unguild event score by 1.
- 2) Guild Switch Event - for each time a character is observed in a guild in snapshot X and then observed in a different guild in snapshot X+1, we increment their guild switch event by 1.
- 3) Guided - whether a character is guilded or not at the end of the sampling period.




Figure 2 – Unguild, switch, and guilded events across levels

We found that unguild events were far more frequent than guild switch events and this effect magnified over the level spread (Figure 2). Between levels 21-40, unguild events are 3 times more frequent than guild switch events (4% vs. 13%); between levels 41-60, unguild events are 7 times more likely than guild switch events (3% vs. 21%). When characters leave a guild, it takes them some time to find a new home – the more so as they increase in level.

This seems to fit well with some of the guild difficulties we mentioned earlier. After an episode of “drama”, leaders will often forcefully remove the offending member(s) without

notice. Conversely, players can get so frustrated or unsatisfied with a guild that they would rather leave and be alone. This, combined with the fact that many guilds require a “trial period” before accepting new members, explains why some players can find themselves in a prolonged interim without guild affiliation. As admission criteria also become more stringent with rising levels, it seems logical that unguild events would far outnumber guild switches over the life of a character.

Of course, WoW is a dynamic world and as servers mature, we would expect guild size and stability to change. To assess the evolution of guilds we focused on a 6-month long period in our data (July 2005 to January 2006), looking at guild membership every 2 weeks (yielding about 100,000 observed characters in each 2-week sample). First, we looked at the percentage of characters who were in guilds. There was a mild positive increase over time. This increase in percentage of gilded characters could mean one of two things: there may be more guilds that spring up, or characters are joining existing guilds. Figure 3 suggests the latter is the case: over time, established guilds attract more and more characters and increase in size.

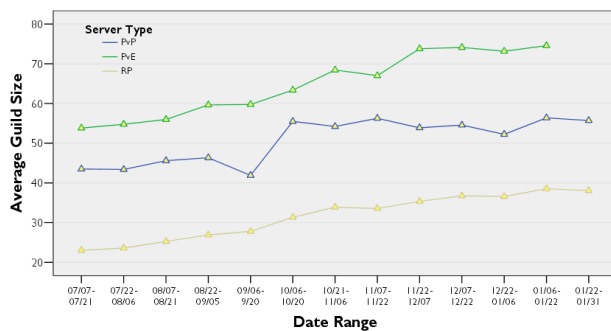


Figure 3 - Average guild size over time, by server type

Over time, guilds also stabilize. As Figure 4 shows, members are less likely to quit a guild as a server matures. Overall, these analyses suggest that over time, characters on a server are more and more likely to be in a guild; the guilds they join tend to be established guilds; and over time, guild turn-over decreases.

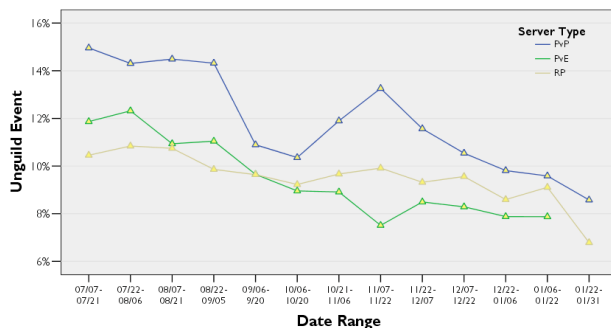


Figure 4 - Guild churn over time, by server type

We also looked at whether churn was different across the server types. The data showed that member churn was

significantly and consistently higher on PvP servers than RP or PvE servers, by about 75% to 100% (Figure 4). Again, this seems to confirm broad trends reported elsewhere [9, 26]. The PvP worlds are more dangerous places, and guilds may be serving a more utilitarian function than on other servers: if the guild fails to deliver the required amount of protection and reward, players start looking elsewhere. This also fits with reports that PvP players tend to be more achievement-oriented [28] and instrumental in their approach to group selection, as opposed to role-players who value group life more.

Summary: consolidation and specialization

The above data reveals interesting population dynamics within and across guilds. Overall it looks like guilds are often in flux, but there seems to be a trend towards consolidation where “the rich get richer”: some guilds survive longer than others, grow in size, and attract most of the churn from other guilds. In parallel there might also be a trend towards specialization, with the most established guilds focusing on specific aspects of the game (e.g. PvP, raids) and filtering new members accordingly, which increases the time required for players to find a new guild when leaving another – the more so at higher levels.

This leads us to the central question of this paper: What causes the rich to get richer in WoW? Can we explain the survival and growth of guilds using structural variables?

THE IMPACT OF GUILD STRUCTURE

Since our software collects data from the client-side of the game, we cannot measure the structural properties of a guild exhaustively. Still, the “/who” command we rely on covers a broad range of variables, and many of these could potentially have significant impacts. We had access to the following indicators:

- **Size:** number of characters bearing a given guild tag during the sampling period. As we saw earlier bigger guilds tend to attract more members over time. It is therefore reasonable to hypothesize that size has positive impact on a guild’s evolution.
- **Density:** connections between guild members can be mapped out as a matrix. The density of a guild is the percentage of matrix cells that are filled in. In previous work we saw that guild social networks in WoW tend to be very sparse [9]. We wanted to explore whether or not guilds benefit from higher social connectivity.
- **Centrality:** for each guild member, their degree centrality is the number of connections they have divided by the total number of connections they can have (i.e., the guild size - 1). The guild’s centrality is the average of all of its character’s centrality scores.
- **Maximum subgraph size:** largest interconnected cluster of members in a guild’s social network. This measure gives a rough sense of how large subgroups can get within a guild. Larger groups often experience more

coordination issues and overhead, which could impact survivability and performance.

- **Mass count:** the number of subgraphs larger than three in a guild's social network, that is, how many independent subunits there are. Fragmentation of the membership might create more manageable and more successful groups within a guild, or it could impede information sharing and be detrimental.
- **Level (average, median, and standard deviation) and number of level 60 characters:** indicators of the level of player experience in a guild. A large number of level 60 players knowing a lot about WoW could presumably help a guild in the long run. And overall guilds of higher level might fare better than lower ones.
- **Average time spent together:** a measure of schedule compatibility – the higher the value, the more members are online at the same time (we normalize this value using each guild's size to be able to compare them). Schedule incompatibilities are often mentioned by players as an important reason for leaving a guild [26].
- **Average time spent in instances:** an indicator of the importance of planned activities in a guild, as opposed to ad-hoc quest parties.
- **Class balance:** a good play group in WoW often has representatives of different classes, since they are highly complementary by design. We use a chi-square score to measure overall balance or imbalance. The chi-square score calculates the deviation of each class count from the expected count for a given size (e.g, there being 8 classes for each faction, a perfectly balanced guild of 80 members would have 10 members of each class). Bigger scores mean bigger imbalances (we normalize the result using each guild's size).

Having computed the above for each guild in our sample, we then tried to assess their impact on two success indicators for a guild: its survival, and the rate of advancement of its members.

Guild Organization and Survival

To study guild survival, we took two month-long samples, one from July 2005 and the other from December 2005, and extracted all unique guilds in both. If a guild seen in the early sample was not observed in the later one, we marked it as "dead". Otherwise, we marked it as "survived". Using this method, we had 3,537 unique guilds in our July sample. Of those, 1,917 (or 54%) were not seen again in December and marked as "dead".

We then ran a logistic regression with survival as the dependent variable and all the metrics mentioned earlier as predictors. The Cox & Snell R-Square for the resulting model was .200 (Table 2) – a number that may initially seem low but is in fact well within the accepted norms for

similar social science research [8]¹. And again, we openly acknowledge that our model cannot be entirely accurate since we can only collect a limited number of variables.

Step	-2 Log likelihood	Cox & Snell R Square	Nagelkerke R Square
1	4088.392	.200	.268

Table 2 - Guild survival model summary

Observed	12_Survival	Predicted		Percentage Correct
		.00	1.00	
Step 1	12_Survival	.00	1.00	76.5
		1467	450	64.8
	Overall Percentage	570	1050	71.2

Table 3 - Classification table for the survival model

	B	S.E.	Wald	df	Sig.	Exp(B)
@07_guildsize	.054	.011	24.481	1	.000	1.055
@07_max_subgraph_size	-.048	.012	15.845	1	.000	.953
@07_mass_count	.077	.109	.502	1	.479	1.081
@07_total_count	-.080	.072	1.237	1	.266	.923
@07_density	3.586	.916	15.339	1	.000	36.081
@07_centrality	-1.527	.889	2.948	1	.086	.217
@07_LvlAvg	.021	.006	13.935	1	.000	1.021
@07_LvlMedian	-.002	.003	.369	1	.533	.998
@07_LvlSD	.046	.010	23.283	1	.000	1.048
@07_Count60	-.022	.008	7.143	1	.008	.978
@07_AvgTogether	.090	.074	1.466	1	.226	1.094
@07TogetherRatio	-1.243	.377	10.887	1	.001	.288
@07_InstanceTimeRatio	7.011	1.782	15.481	1	.000	1109.076
@07_ClassBalance	.003	.005	.423	1	.515	1.003
@07_BalanceRatio	-.288	.054	28.135	1	.000	.750
Constant	-1.541	.195	62.413	1	.000	.214

Table 4 - Regression coefficients for the survival model

Using a strict cut-off, the model provided by the logistic regression was accurate in 76.5% of the "death" cases and 64.8% of the "survival" cases (Table 3) - better than chance alone. The model identified six significant predictors of survival (Table 4) we can rank using the Wald test. In order of importance, we find:

- **Class balance ratio (28.135):** unsurprisingly, more balanced guilds survive better than others. More importantly, this can also explain why churn is so high across guilds, and why some get bigger while others disappear entirely. Indeed, we know from previous research that the distribution of classes over the entire population is very imbalanced [9, 10] – priests (a crucial healing class), for instance, are in notoriously short supply. And therefore, their presence in one balanced guild means class imbalance in another. The quest for a well-balanced roster leads to churn, as players from the needed classes are recruited away from one guild to another (this could be especially prevalent for guilds focusing on “endgame” content).

¹ Cohen states that an R of .37 would be considered “large” (with a corresponding R-Square value of .14), for data collected during highly-controlled experimental conditions. Considering that our analysis was conducted on a large naturalistic sample with a great deal of extraneous noise, a R-Square of .200 is therefore quite high.

- Guild size (24.481): as expected, bigger guilds are more likely to survive.
- Level standard deviation (23.283): a wider level spread contributes positively to survival. Our hypothesis that a concentration of high level characters would increase the guild's knowledge pool, and therefore its survival, does not seem to hold here. But an alternative explanation could be that a wide level spread is indicative of fresh recruits joining the ranks, replacing natural attrition through burn-out and transfers to competing guilds.
- Maximum subgraph size (15.845): controlling for guild size, guilds with smaller subgroups are more likely to survive – perhaps because they avoid coordination issues, as we hypothesized.
- Time in instances (15.481): interestingly, guilds that focus on the most complex game areas survive better. Since these dungeons usually require more planning and coordination than simply “roaming the world”, it could be a reflection of a more organized guild (as opposed to one limited to ad-hoc quest groups).
- Density (15.339): better connected guilds apparently survive more often than others. Anthropologists like Dunbar [11] have proposed that a certain amount of “social grooming” is necessary to hold a group together. A larger number of ties might be indicative of higher cohesion and more peer pressure to participate in guild activities, increasing its odds of success.

While far from providing a definitive answer, these analyses show that simple structural indicators can enrich our understanding of group dynamics online and help predict their long-term survival. In the context of online game guilds, attracting a large number of members is key but the composition and organization of this membership is equally important. In particular, guild leaders need to make sure that class and level spread are as broad as possible. It is especially important to prevent the guild from becoming “top heavy” with too many level 60 characters. As this would be hard to achieve through chance alone, a pro-active recruitment strategy is probably needed. Organizing “instance runs”, as opposed to purely ad-hoc groups, also seems to contribute positively to survival

Moreover, while guilds benefit from a dense internal social network, the size of their largest subgroup can become a problem. This indicates that large group activities in WoW (e.g. 40-man raids) require significant coordination efforts that few guilds can manage successfully. We discuss the implications of these findings later in this paper.

Guild Organization and Player Advancement

While MMOG players join guilds for many reasons [26], the primary motive is often game-related. A guild provides access to shared resources, knowledge, and game partners that can all facilitate progress through the game. The extent

to which this actually works, however, can be limited: grouping can be an inefficient way of advancing in an online game [9]. We explored the relationship between guild structure and the progress of its members.

For a measure of player advancement, we computed a standardized character advancement score. A character's raw advancement is simply the number of levels the character has advanced over one month (for the analyses below, from July to August 2006). In this case, we subtracted the starting level from the ending level. Because a 10 level advancement by a level 1 character is much less significant than a 10 level advancement by a level 50 character (the later stages of the game require much more time and effort to progress), we standardized character advancement by calculating the average (and standard deviation) of advancement for every starting level. In other words, we compared each character only with others who also started at the same level at the same time. This was done by calculating the z-score of advancement for every character. Characters who were already level 60 at the beginning of the sampling period were excluded.

We then computed a standardized guild advancement score – simply the average of the standardized advancement scores of every member in that guild. This guild score was thus a reflection of how much the guild as a whole advanced during the sampling period. Again, characters who were already level 60 at the beginning of the sampling period were excluded.

Using the same predictors as in the previous section on guild survival, we ran a multiple regression with guild advancement as the dependent variable. The R-Square for the resulting model was .098 (Table 5) – smaller than before but still within acceptable limits. The model identified five significant predictors of character advancement (Table 6).

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.313 ^a	.098	.094	1.31596

Table 5 - Guild advancement model summary

Model	Unstandardized Coefficients		Standardized Coefficients	t	Sig.	
	B	Std. Error	Beta			
1	(Constant)	-.574	.101		-5.681	.000
	@08_guildsize	-.016	.005	-.505	-3.296	.001
	@08_max_subgraph_size	.018	.006	.470	3.043	.002
	@08_mass_count	.250	.065	.107	3.876	.000
	@08_total_count	.033	.039	.021	.839	.402
	@08_density	-.134	.431	-.010	-.312	.755
	@08_centrality	.898	.424	.078	2.117	.034
	@08_LvIAvg	.001	.003	.008	.328	.743
	@08_LvIMedian	.002	.002	.025	1.244	.213
	@08_LvISD	.014	.005	.056	2.711	.007
	@08_AvgTogether	.139	.034	.226	4.044	.000
	@08TogetherRatio	.574	.186	.071	3.078	.002
	@08_InstanceTimeRatio	-2.092	1.069	-.035	-1.957	.050
	@08_ClassBalance	.007	.002	.056	2.890	.004
	@08_BalanceRatio	-.018	.024	-.019	-.762	.446

Table 6 - Regression coefficients for the advancement model

In order of importance based on the standardized coefficients we find:

- Guild size (-.505): players progress faster in smaller guilds – an interesting contrast to the earlier model.
- Maximum subgraph size (0.470): the larger the subgroups in a guild, the faster players advance. This is again opposite to the survival model. But since “fast” guilds are also smaller (see above), it is still probable that these groups are not too large.
- Schedule compatibility (“Together ratio”) (0.186): perhaps unsurprisingly, guilds with members whose time online overlaps significantly have a positive impact on advancement – they make finding partners for joint play sessions easier.
- Mass count (.107): a guild fragmented into many cohesive subunits is more beneficial to its members’ advancement. This fits well with WoW’s design: most “quests” are designed to be challenging enough for small groups of up to 5 players. Guilds where players can repeatedly team with up to 4 other members of approximately the same level (see below) should therefore facilitate advancement.
- Class balance (0.056): here again, a well balanced guild has a positive effect on its members’ progress – presumably because forming balanced and efficient leveling groups is easier.
- Levels standard deviation (0.056): the broader the range of levels in a guild, the faster players progress. This is most probably because such a spread does not constrain players to a fixed rate of advancement. For example, if the bulk of a guild progressed from level 25 to 30 in a given month, characters below 20 and above 35 would have trouble finding partners of the appropriate level. A large level spread ensures that there will always be someone in the guild with a level close enough to play with – and this whether each player advances faster or slower than the guild’s norm.

While some of the predictors differ from our earlier analysis of guild survival, similar trends can also be seen. In order to benefit their members’ progress, guilds apparently need to be broken down into separate subgroups that cater to different level bands, thus facilitating teaming and leveling. Unsurprisingly, schedule compatibility is also important: the more members’ playtime overlaps, the easier it is to form a group and progress more quickly. But interestingly, size does not help. On the other hand, playing with a broader subsection of the guild (the “max subgraph” variable) is useful, most probably because it corresponds to having a more diverse choice of partners. There is therefore an interesting tension between advancement and survival: growing and partitioning a guild into small subunits increases the group’s chances of survival, but it is less beneficial to each individual member.

Many of the predictors we identified above and in the previous section might sound “obvious” to long-term WoW players – and indeed, they fit our own intuition about

successful strategies in the game fairly well. But our data allows us to substantiate such intuitions and highlight trends that could prove important for the design of future online gaming communities. We now discuss the implications of our findings in more depth.

DISCUSSION

Small Is Beautiful: Designing For Successful Gaming Communities

Anthropologist Robin Dunbar proposed that “there is a cognitive limit to the number of individuals with whom any one person can maintain stable relationships” [11]. Based on studies of the group size of a variety of primates, Dunbar predicts that 150 is the “mean group size” for humans. This, in turn, matches census data obtained from villages and tribes in many cultures. But Allen argued that, online, group size will usually plateau at a number lower than “Dunbar’s number” of 150 [1]. Citing evidence from several online communities (in particular another MMOG, Ultima Online), Allen hypothesizes that the optimal size for creative and technical groups (as opposed to exclusively survival-oriented groups such as villages) is around 45 to 50. The data we obtained from WoW gave us the opportunity to further test this hypothesis in the context of gaming communities. Interestingly, our numbers are very close to Allen’s hypothesis: most guilds in WoW have 35 members or fewer.

WoW therefore confirms that, in games as in other online social spaces, mass collective action can be difficult to achieve. Returning to Dunbar, this difficulty could be due to limited “social grooming” [11], that is, repeated interactions between the members of a guild. As we saw above, a number of simple game design factors conspire against the formation of cohesive subgroups in guilds - schedule incompatibilities, level gaps, class imbalances, etc. As a result social networks in guilds tend to be sparse, and it is well known that when the likelihood of two individuals working together again is low, people tend to behave selfishly [2] – and leave. Such trends can be exacerbated where individuals self-select for achievement and an instrumental orientation to online play: as we saw, churn is highest on PvP servers.

It has been argued before that online communities (Usenet newsgroups for instance) can favor the emergence of very large groups [15], because the medium itself reduces the costs of communication and coordination, but online games like WoW are almost the antithesis of these pioneering online social groups in this respect. In particular, WoW exacerbates the challenge of finding people with similar interests: no information is readily available about the makeup of a guild, its collective interests, its needs for new members of particular levels and classes, etc. Most of this information is traded out-of-game (if at all) on forums that are not visited by all players. Yet our analyses show that simple variables could be used to better match players to

guilds: for instance, making each guild's roster publicly visible in-game could go a long way.

Some online game designers seem to have taken notice. Subscribers to Sony Online Entertainment's (SOE) Everquest II can get access to dedicated tools to publish information about their guilds on SOE's web site – provided they pay an additional fee for the “premium” service. Considering the importance of such information for the long-term health of guilds, we would argue that online games would benefit from providing such a service in-game and for free.

Guilds in WoW are also susceptible to a form of “tragedy of the commons” against which previous online communities had developed rules and institutions [15]. In particular, leaving a guild has no cost to the player: typing “/gquit” is enough to remove oneself from the group. As such, nothing prevents players from leaving a guild as soon as their personal objectives are accomplished. To be sure, high-level players who behave selfishly will tarnish their reputation and news travels fast on a WoW server, decreasing their chances of finding a new group. Still, no mechanisms are in place to build up a player's attachment to his/her guild, which probably encourages churn. But here again we see signs of interesting design changes: in City of Heroes, another MMOG produced by NCSoft, guild (“supergroup”) members are expected to play in “SG mode,” which means that they receive fewer “influence” (in-game currency) points for their actions because part of the influence is converted to “prestige,” the guild currency, for the guild's use. Here membership is actually exacting a definite cost, which should make the boundaries of a guild less porous and potentially reduce free-riding.

Another worrying trend emerging from our data is that guilds seem to have a tendency towards entropy over the long run. Groups get larger and larger, monopolizing the most-needed players and concentrating the game's most coveted rewards in the hands of a few. This has the potential to negatively impact playability over time, in two opposite ways illustrated by our data: large guilds can become “top-heavy” and susceptible to burn-out; new players can have a harder time progressing since few groups are available to cater to their needs. The difficult issue seems to be to encourage “healthy” levels of churn that prevent guild stagnation yet do not threaten their survival and growth.

Overall, WoW is a fascinating example of group dynamics in an online environment with little to no support to group formation and coordination. It is interesting to note that WoW's designers may have overestimated the size that a group can reach organically under these conditions: the 90th percentile for guild size, 35, falls just short of what is required to access the game's toughest (and most rewarding) content: 40-player raids. As such, a very large number of players cannot enjoy a substantial portion of the game, simply because they cannot grow a group to the

necessary size (a problem we explore in more depth in [10]). When designing group activities in online games, short of providing an extensive set of tools to support large social units, the best principle might therefore be that “small is beautiful” [19] – a somewhat ironic conclusion for massively multiplayer environments with millions of subscribers. Blizzard seems to have adopted a similar view: the majority of new high-end dungeons they recently added require only 10- or 20-player groups, well within the reach of a 35-members guild.

A Social Dashboard for Managing Gaming Communities

As we mentioned above, games like WoW provide few tools out-of-the-box to facilitate the large-scale, collaborative activities MMOGs are famous for. Yet monitoring simple variables, like the ones we used in our models, could help identify some important problems in groups. Both players and game managers could benefit from tools to track group-survival metrics: the former could adapt their guild's recruitment strategy to increase their chance of success in game, and the latter could monitor the health of guilds across an entire server to assess the impact of their game's design on collaboration.

Inspired by similar efforts focused on other online communities [e.g. 18, 22], we developed a prototype Social Dashboard to visualize and explore the guild survival metrics we described earlier. We have used this tool internally in our research, and hope to release it to players and game designers alike in the near future. We present it below as a simple example of what could be done when mining social interaction data from online games.

The Social Dashboard's main screen presents an overview of some key guild survival metrics (only three are shown in Figure 5: guild size, density, and number of subgroups) for an entire game server. Each gauge clearly indicates “dangerous” and “critical” thresholds for each variable, based on the models we described earlier. This gives the user (here most probably a community manager) a sense of the most important areas to address – on this particular server for example, guilds are too small.




Figure 5 - The Social Dashboard's main screen

To understand the problem in more detail, the user can click one of the gauges to access a report on the evolution of the metric over a given time period (in Figure 6, over a month). Aggregate values for the entire server are available (the

blue lines – light blue are daily values, dark blue is a mobile mean over seven days), as well as specific data for any given guild on the server (the dark lines). Guilds can be selected from the complete list to the right or from pre-computed short lists of groups that have passed the “dangerous” or “critical” thresholds for this metric (in the example above, the selected guild’s size has been collapsing over the past month and just dipped below the survival threshold, represented by the dotted green line).

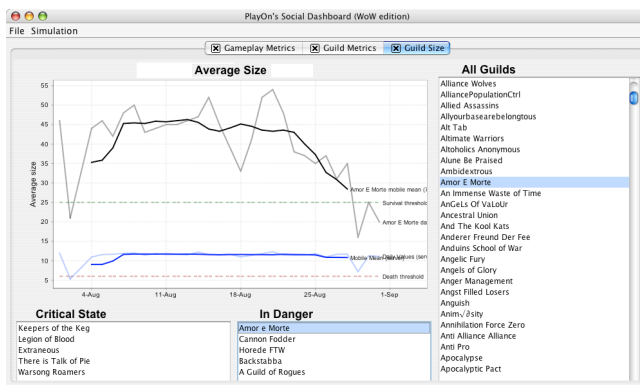


Figure 6 - Evolution of a survival metric over time

Finally, the user can explore which factors in the guild’s composition and organization might have contributed to the problem identified earlier. The Social Dashboard can display the evolution of a guild’s social network over time (Figure 7), allowing the user to observe the changing roles of veteran guild members and newcomers alike, as well as the impact of members leaving. The network displays additional information relevant to guild survival, such as a player’s class and level. Various components of the network can be isolated using standard simplification techniques (e.g., eliminate nodes based on degree or strength of ties)².




Figure 7 - Social network for a guild early in the month

² We use a deterministic layout algorithm to ensure the position of each player in the network remains the same from one analysis session to the next. Our dynamic network visualization package was implemented on top of the Prefuse toolkit [12]; the algorithm itself was inspired by the Kamada-Kamui layout [14] used in the SoNIA project [4].

Learning Teamwork From Games

As we saw earlier, successful guilds in WoW are both big and divided into multiple, small subgroups (around 6 players per subgroup for most, see [9]). From the perspective of organization theory, successful guilds are therefore organic, team-based organizations [7]. This fits the game environment well: most tasks require small groups (5 participants for most quests) with complementary skills and similar levels (if the gaps between levels in a group exceeds 10, the higher-level participants do not earn “experience points”). Guilds provide the opportunity for forming such cohorts that will progress through the game at the same pace. But in parallel, the overall size of the guild provides access to resources that could not be obtained otherwise. In a large guild, players can specialize in crafting special items for other players, getting other items in returns. The larger the guild, the more this specialization makes sense – in other words, guilds reduce transaction costs [27]. Getting information and help from guildmates is also generally easier than asking random strangers. As such, the exchange of information and resources provides an incentive for joining a large guild, while the structure of in-game activities encourages small teams.

These findings are particularly interesting in light of the recent debate about the educational value of games that are not originally designed with the teaching of specific skills in mind. For instance, it has been argued that the “video game generation” is acquiring valuable knowledge from games that will help them transform the workplace [3]. Our observations indicate that MMOGs like WoW certainly familiarize their players with organizational forms that are prevalent in today’s work environment. Players are also given clear roles (their class) that naturally steer them into specific positions in their guild’s social network. This may later affect the way these players behave in the workplace (for instance, WoW players might prefer working in small teams with clearly-defined individual responsibilities). The relationship between online games and “real world” behavior in organizations is clearly an opportunity for future research.

CONCLUSION

Online games can be fascinating laboratories to observe the dynamics of groups online. In games as in other online social spaces, growing and sustaining large communities can be quite difficult. Our findings reinforce earlier research showing that there might be a hard limit on the size of a viable organic group online, possibly set at around 35 group members or less. This has important implications for the design of current and future games, since most require players to form substantially larger social units that might be unsustainable without additional support. Somewhat surprisingly, games like WoW do not offer much collaboration infrastructure to their player associations, despite years of research on cooperation and conflict online. If players had access to simple data to evaluate a guild’s profile, a great deal of churn could

possibly be avoided. We presented one tool we designed to address this problem, the Social Dashboard – but much more could be done.

Still, some guilds manage to optimize aspects of their organization to increase their chances of growth and survival. While our data is inherently limited and we believe more factors are at play, our analyses show that simple models can help isolate some beneficial structural properties for a guild. In WoW, this means simultaneously growing a guild while partitioning the members into small, balanced subgroups (in terms of class and levels) that are best suited to doing quests and other activities. The guild itself serves as a broader social environment where resources and services can be exchanged. This “optimal” organization is a direct consequence of WoW’s design and might not sound surprising to veteran players. Still, we have been able to show that there is apparently little room to deviate from these built-in constraints. This, in turn, steers the players towards certain forms of teamwork that might transfer to group activities outside of games. Such data is particularly relevant in light of current debates about the educational value of MMOGs and their possible impact on the workplace.

ACKNOWLEDGEMENTS

The authors would like to thank Cabell Gathman for her insightful comments on early drafts of this paper.

REFERENCES

- [1] Allen, C.: The Dunbar number as a limit to group sizes. <http://tinyurl.com/e3t8w>.
- [2] Axelrod, R.: The evolution of cooperation. Basic Books, New York (1984)
- [3] Beck, J.C., Wade, M.: Got game: How the gamer generation is reshaping business forever. Harvard Business School Press (2004)
- [4] Bender-deMoll, S., MacFarland, D.A.: The art and science of dynamic network visualization. *Journal of Social Structure*, 7 (2). (2006)
- [5] Bernheim Brush, A.J., Wang, X., Combs Turner, T., Smith, M.A.: Assessing differential usage of Usenet social accounting meta-data. In: Proceedings of CHI 2005, ACM, New York, (2005), 889-898
- [6] Blizzard: January 11, 2007 press release. <http://www.blizzard.com/press/070111.shtml>.
- [7] Burns, T., Stalker, G.M.: The management of innovation. Tavistock Publications, London (1961)
- [8] Cohen, J.: Statistical power analysis for the behavioral sciences (2nd ed.). Lawrence Erlbaum Associates, Hillsdale, NJ (1988)
- [9] Ducheneaut, N., Yee, N., Nickell, E., Moore, R.J.: "Alone Together?" Exploring the social dynamics of massively multiplayer online games. In: Proceedings of CHI 2006, ACM, New York, (2006), 407-416
- [10] Ducheneaut, N., Yee, N., Nickell, E., Moore, R.J.: Building a MMO with mass appeal: a look at gameplay in World of Warcraft. *Games and Culture*, 1 (4). (2006) 1-38
- [11] Dunbar, R.I.M.: Coevolution of neocortical size, group size and language in humans. *Behavioral and brain sciences*, 16 (4). (1993) 681-735
- [12] Heer, J., Card, S.K., Landay, J.A.: Prefuse: a toolkit for interactive information visualization. In: Proceedings of CHI 2005, ACM, New York, (2005), 421-430
- [13] Huberman, B.A.: The laws of the Web: patterns in the ecology of information. MIT Press, Cambridge, MA (2001)
- [14] Kamada, K., Kawai, S.: An algorithm for drawing general undirected graphs. *Information processing letters*, 31. (1989) 7-15
- [15] Kollock, P., Smith, M.A.: Managing the virtual commons: cooperation and conflict in computer communities. In: Herring, S. (ed). Computer mediated communication, John Benjamins Publishing Company, Philadelphia, PA, (1996)
- [16] Mintzberg, H.: The structuring of organizations. Prentice Hall (1978)
- [17] Nardi, B., Harris, J.: Strangers and Friends: Collaborative Play in World of Warcraft. In: Proceedings of CSCW 2006, ACM, New York, (2006)
- [18] Sack, W.: Conversation Map: An interface for very large-scale conversations. *Journal of Management Information Systems*, 17 (3). (2001) 73-92
- [19] Schumacher, E.F.: Small Is Beautiful: Economics As If People Mattered. Blond & Briggs Ltd., London (1973)
- [20] Seay, A.F., Jerome, W.J., Lee, K.S., Kraut, R.E.: Project Massive: A study of online gaming communities. In: Proceedings of CHI 2004, ACM, New York, (2004), 1421-1424
- [21] Smith, M.A.: Measures and maps of Usenet. In: Lueg, C., Fisher, D. (eds.) From Usenet to Cowebs: Interacting with Social Information Spaces, Springer, Berlin, (2003), 47-78
- [22] Smith, M.A., Fiore, A.T.: Visualization components for persistent conversations. In: Proceedings of CHI 2001, ACM Press, NY, Seattle, WA, (2001), 136-143
- [23] Taylor, T.L.: Play between worlds. The MIT Press, Cambridge, MA (2006)
- [24] Wasserman, S., Faust, K.: Social network analysis: methods and applications. Cambridge University Press, Cambridge, UK (1994)
- [25] Whittaker, S., Terveen, L., Hill, W., Cherny, L.: The dynamics of mass interaction. In: Proceedings of CSCW 1998, November 14 - 18, 1998, Seattle, WA USA, (1998), 257-264
- [26] Williams, D., Ducheneaut, N., Xiong, L., Zhang, Y., Yee, N., Nickell, E.: From treehouse to barracks: The social life of guilds in World of Warcraft. *Games and Culture*, 1 (4). (2006) 338-361
- [27] Williamson, O.E., Masten, S.E.: The Economics of Transaction Costs. Edward Elgar (1999)
- [28] Yee, N.: The Daedalus Gateway. <http://www.nickyee.com/daedalus>.
- [29] Yee, N.: The labor of fun: How video games blur the boundaries of work and play. *Games and Culture*, 1. (2006) 68-71