

Group Movement in World of Warcraft Battlegrounds

John L. Miller^{1,2}

¹Microsoft Research, Cambridge
Cambridge, United Kingdom
+44 1223 479813

johnmil@microsoft.com

Jon Crowcroft²

²University of Cambridge Computer Laboratory
Cambridge, United Kingdom
+44 1223 763633

jac22@cl.cam.ac.uk

Keywords: Avatar Movement; Massively Multiplayer Online Games; World of Warcraft Battlegrounds; Waypoint Navigation; Hotspots;

Abstract

Distributed Virtual Environment (DVE) topology management and message propagation schemes have been proposed for many years. Evaluating DVE message propagation schemes requires a variety of assumptions whose verity significantly affects results, such as details about avatar movement characteristics. We implemented two schemes for waypoint and hotspot detection, and examined their applicability for characterizing avatar movement. We confirmed that waypoint detection doesn't yield good results for characterizing human avatar movement, and gained new insight into why by rendering avatar movement as point clouds. We implemented an existing hotspot detection model, and proposed an enhancement to help overcome one limitation of cell-based hotspot detection. We were able to immediately apply this hotspot detection technique to help analyze group movement. We discovered that although a third of movement time in the battlegrounds is spent in inter-node journeys, less than a quarter of these journeys are made in groups.

1. Introduction

Distributed Virtual Environments come in all shapes and sizes, from simple turn-based games to the more prevalent real-time three dimensional simulations. Evaluating performance and scalability of these systems requires assumptions about game and player avatar behavior. Often times these assumptions are based upon anecdotal experience, or the traffic patterns of a game purpose-built by researchers to evaluate their platform.

The choices made for evaluating systems can vary system impact by orders of magnitude. For example, does an avatar move once per second or continuously, up to the render frame rate? Do avatars move independently, or in groups? Do they navigate using fixed features, or is their movement too complex and varied to characterize in this fashion? Building a DVE and showing that it *can* be used is very different from building a DVE which *will* be used. We believe that existing broadly deployed DVE's are the best source of models for evaluating DVE research results intended for DVE adoption.

World of Warcraft (WoW) is by far the world's most popular DVE, and fertile ground for gathering data about actual player avatar behaviors. We instrumented a relatively small part of the WoW experience – player-versus-player battlegrounds - to focus our findings. Battlegrounds were chosen because of a combination of tractability and applicability to DVE performance modeling. We examined three common assumptions people make about DVE performance when evaluating their DVE research, all related to the way avatars move. We investigated whether a waypoint model is a good fit to describe avatar movement, whether movement patterns result in significant hotspots, and whether player avatars organize into coherent groups for inter-hotspot journeys.

We found that waypoints cannot be easily applied to player avatar movement, and verified this by implementing and applying an appropriate existing waypoint detection model described by (Thawonmas, Kurashige, & Chen, 2007). More positively, we found significant evidence of hotspots, with 5% of visited territory accounting for 30% of all time spent in a typical battleground. Finally, we found that player avatars usually make significant journeys alone rather than in groups, despite clear game incentives to group.

The remainder of this paper provides supporting evidence for these findings. We start by describing World of Warcraft for those not familiar with the game, and provide particulars about the battleground we evaluated, Arathi Basin. We detail our methods for gathering data on player avatar movement within battlegrounds, and accuracy and completeness of that data. Finally, we provide data relevant to waypoint, hotspot, and group movement models for DVE player avatar movement.

2. Background

2.1 World of Warcraft

World of Warcraft (WoW) is the most popular DVE in history. With more than 11 million subscribers worldwide (Blizzard Entertainment, 2008), World of Warcraft has the majority of market share for all massively multiplayer online games (62% as of 2008 (Woodcock, 2008)). Because of its ubiquity, WoW is especially relevant as a DVE user study test bed.

2.1.1 Types of Experiences

WoW experiences can be divided into four main categories: world PvE / PvP, capital cities, instances, and battlegrounds. This paper focuses on battlegrounds. This section provides a basic overview of the other categories to support this decision.

World PvE ('player vs. environment') and *World PvP* ('player vs. player') are activities where player avatars wander the game world individually or in small groups. The game world is large and detailed – requiring more than an hour of real time for an avatar to walk across one of its four continents - and player avatars are relatively sparse. Indeed, player avatars are usually out of mutual interaction range unless they explicitly seek each other out. The world requires significant resources to simulate and to communicate with DVE clients. The game scales by running many simultaneous world copies called *shards*, with each avatar belonging to exactly one shard. A typical shard has between a few hundred and a few thousand active players online at any given time (Pittman & GauthierDickey, 2007), out of a population of tens of thousands assigned to that shard.

The game world contains ten large cities called *capital cities*. These offer a plethora of facilities, and are densely populated relative to the rest of the world. For example Dalaran, the current end-game capital in World of Warcraft, can have a quarter of the active population on a given shard concentrated in less than 0.1% of the world's geography. While densely populated, capital city interactions tend to be infrequent and lightweight, with most avatars sitting still and performing social or character maintenance activities.

Instances are small, self-contained adventures which groups of players play together. Just as a shard is one of many copies of the game world, an instance is one of many copies of that adventure. Instances are shared by either self-selected or randomly matched set of avatars – typically five, though up to 40 are allowed in certain 'raid' instances. Together they solve puzzles and fight particularly difficult AI-controlled avatars. An instance can be thought of as a private PvE experience requiring a group of players to complete. Unlike other PvE experiences, the players tend to stay together and move in a tight group from place to place.

Battlegrounds are a special type of instance, with a PvP focus. Like PvE instances, there can be many identical battlegrounds active and reachable from a given shard. A single battleground instance can be

populated by player avatars from multiple shards, known collectively as a battle group. Battlegrounds are characterized by scenarios which reward PvP, usually to achieve an objective or dominate a resource. Battlegrounds have intense continuous activity with between 20 and 240 mutually interacting participants split into two opposing sides called factions. For comparison, we measured 40kbps of incoming TCP traffic in a capital city with more than 250 people. A similar traffic measurement in the Wintergrasp battleground with 200 people often reaches 250 kbps of inbound TCP packets per client, and can jump over 500 kbps in times of heavy activity.

We've done preliminary analysis of game traffic for each of the categories above, evaluating avatar mobility, traffic magnitude, and completeness of the data we were able to capture. That preliminary analysis is not yet formalized enough to present here, but guided us to select PVP battlegrounds as the environment for the first leg of our avatar research. Our tools are able to gather a complete picture of activity in select battlegrounds, which we cannot easily do for world PvE. Battlegrounds have significant and varied avatar mobility and interaction, unlike instances, where avatars tend to all stay within interaction range of each other. Battlegrounds also have more inter-avatar activity than capital cities, where most traffic is generated by avatar movement, and joiners, and AI-controlled avatars.

Battlegrounds with their relatively high traffic requirements, inter-avatar interaction, and movement characteristics make a good test environment for DVE research related to player avatar movement. We chose the Arathi Basin battleground to measure these behaviors, as it is big enough to be interesting, and small enough to be tractable for measurement and analysis.

2.2 The Arathi Basin Battleground

There are six different battlegrounds in World of Warcraft. Battlegrounds are organized around inter-faction – Alliance and Horde – competition. Arathi Basin is a 30-person battleground where teams compete for control of five stationary flags. Gaining control of a flag requires a team member to use the flag without interruption for ten seconds, and to prevent any enemy faction team members from using the flag for an additional minute. Each team receives points every few seconds based on the number of flags they control. The first team to reach 1600 points wins the battleground. Both teams are rewarded, with winners receiving better rewards, incenting each team to win.

The battleground is approximately 600 yards by 600 yards in size, with flags evenly spaced around the center of the map as shown by the circled huts in Figure 1. The circled houses at either end of the battleground are the starting point for each faction, alliance at the northwest, and horde at the southeast. In terms of movement, some terrain slows down avatars, or is impassable. For example, water slows most avatars down to approximately four yard per second, a quarter of normal mounted movement speed. Most cliffs and steep hills cannot be traversed, and falling off them can injure or kill an avatar unless they have special, relatively rare talents. The view from the lumber mill plateau is shown in Figure 2.

Traversing the map requires about one minute mounted, or two minutes on foot, assuming no enemy engagements. Avatars typically must be within either melee range (5 yards) or ranged combat range (usually 30 yards) to interact. We call the latter distance an *interaction interval*. The nearest flags are between 170 and 250 yards from each other in a straight line, several interaction intervals apart. In other words, players halfway between two flags can interact with each other, but not with players near the flags.

Players are rewarded for controlling flags and for killing avatars from the enemy faction. When an avatar is killed, it is turned to a ghost and teleported to the graveyard near the closest controlled flag, or to the faction base if no flags are controlled. Every 30 seconds all ghosts at a graveyard are resurrected, and granted full health and mana.

Battles are usually less than a half hour long, with some turnover in participants. Real life or network problems force some players to drop out, and they are replaced by others waiting to battle. As a result, a given player avatar may be in a battleground for as long as the entire match, or as little as a few seconds.

2.3 Avatar Behavior and Traffic Classification

Little research has been done on avatar movement patterns.

(Pittman & GauthierDickey, 2007) provides an examination of World of Warcraft shard populations using WoW's built-in extensibility. Pittman and GauthierDickey found that the workloads used in simulations to evaluate DVE infrastructure were unrealistic. User sessions in WoW are on average less than half an hour, but can reach 24 hours. Peak populations on a shard are typically five times their minimum population.

(Mitterhofer, 2009) describes a mechanism for detecting automated cheats called 'bots' in World of Warcraft by using waypoints to characterize avatar movement. They found that many cheats use a scripted form of automation where the avatar follows the same path repeatedly. They provide a waypoint extraction algorithm and verify that it works to discriminate between player and automated avatars.

Their waypoint extraction algorithm acts upon a series of avatar movement traces. They apply the Douglas-Peucker (Douglas & Peucker, 1973) line simplification algorithm to reduce the traces to simpler lines with fewer vertices. They search for clusters of vertices, and label these as waypoints, as they are endpoints for many paths. Mitterhofer found that scripted bot movement tended to replay the same movement paths, and even with jitter resulted in movement that closely followed detected waypoints.

This approach is useful for simplifying player avatar movement as well as bot movements. Later we describe our efforts to apply the waypoint detection algorithm for waypoints to characterize player avatar movement.

(Miller & Crowcroft, 2009) analyzes avatar movement in the Arathi Basin battleground. It makes several assertions about the existence of waypoints and hotspots, and suggests that group movement is not prevalent. This article builds upon the authors' earlier research, quantifying and qualifying issues with waypoint detection, and providing precise definitions and metrics for analyzing group movement.

Some avatar movement models use a hotspot model to describe places avatars are likely to move towards and congregate. (Thawonmas, Kurashige, & Chen, 2007) describes an algorithm for automatically detecting landmarks, and a method for predicting movement between landmarks. They divide the virtual space into a regular grid. They count the number of visits each avatar makes to each cell in the grid, and compute the weighted entropy for the distribution of player visits. Cells are designated as landmarks based on their entropy, prioritized from highest to lowest entropy. Once a cell has been chosen as a landmark, its eight neighbor cells in the grid are omitted from landmark consideration, even if one or more of them have the next highest entropy value.

Evaluation of proposed DVE systems often uses a synthetic workload based on previous research, or on a model generated by the evaluators. For example, (Krause, 2008) compares three different categories of DVE infrastructure using a synthetic workload based upon an average session time of 100 minutes.

Avatars in his evaluation are simulated using a combination group and waypoint model, where groups of simulated avatars agree on a next point to visit, and move there together. Several other frameworks (Lui & Chan, 2002) (Matsumoto, Kawahara, Morikawa, & Aoyama, 2004) (Morillo, Orduna, Fernandez, & Duato, 2005) (Rueda, Morillo, & Orduna, 2007) assume movement and arrival / departure properties of participants without any obvious experimental basis. The assumptions used in all these cases are logical,

but without firm experimental grounding. We believe using a model based on actual DVE participant behavior can provide valuable insight into actual system performance under load.

Player avatar migration patterns are of critical importance for evaluating geometric routing schemes. For example, VAST (Hu & Liao, 2004) (Backhaus & Krause, 2007) and Delaunay triangulation (Matteo, 2007) organize themselves based upon the position of player avatars, and their overhead and efficacy depend heavily upon the density, distribution, and dynamics of those avatars. Likewise, region-based DVE architectures (Fan, Taylor, & Trinder, 2007) (Jardine & Zappala, 2004) (Yamamoto, Murata, Yasumoto, & Ito, 2005) organize region clients by their avatar location, and are affected by the frequency of avatar transitions between zones, and any tendency of avatars to cluster in hotspots.

Understanding the movement of avatars is important for correctly evaluating region-based and geometric-mesh based DVE schemes. The remainder of this paper provides information about avatar behavior in a real-world DVE which can be used to inform future DVE framework evaluations.

3. Methodology

Our goal is to capture all movement events for a set of battleground sessions to ground our evaluation of waypoint and hotspot presence, and group movement. To do this, we need an exhaustive trace of all movement during the battleground session. Unfortunately, WoW clients only receive avatar movement data which is immediately relevant to them. In practical terms, this means movement data for avatars which are within avatar visual range – approximately 250 yards – and which are not blocked by large obstructions such as cliffs.

Each WoW client transmits its own avatar movement updates to the server, which the server redistributes to other clients. Updates are absolute positions, consisting of a client identifier, three-dimensional Cartesian position, facing information, and additional information we don't decode. The server does not send position information for avatars the client can't perceive, such as very distant or stealthed (invisible) avatars.

Experimentation confirmed two well-placed observers receive movement updates for most of the Arathi Basin map. Moving observer avatars into these positions takes between 1 and 2 minutes from the start of a battle, depending upon enemy activity. Battles in our sample set ranged from 4 to 23 minutes in duration.

We used Microsoft Network Monitor 3.3 to capture network traffic, and FRAPS to capture video from a game client's rendered view. FRAPS videos enabled the game client view to be replayed to answer questions about activity in the game associated with specific times in the network traces.

The movement data we captured is correct, but incomplete. The two leading causes of missing data are an observer being out of position, and stealthed non-observer avatars. Observers started out of position, missing some data from the start of the game. Once in position, observers were sometimes attacked and killed, moving them out of position for a minute or more. Avatar death results in the slain avatar's ghost being teleported to the nearest faction-owned graveyard. Resurrection introduces on average a 15 second delay, and returning to post takes another 0.5 to 2 minutes. We used observers with stealth capabilities to reduce the chance of detection, and therefore of being targeted by the enemy. Our observers' positions in the map are marked with white X's in Figure 3. One was at the north edge of the lumber mill plateau, the other on top of a waterfall at the south end of the mine valley to the East.

Our observers avoided combat, and so effectively filled two of the 15 slots in the Alliance team with non-contributors. This biased the results of the battles, but not significantly. Our sample set has a good mix of battle results, with Alliance winning nearly half the observed games, in one case by a score of

1600-0. We were able to capture battles with scores ranging from 1600-1590 (the closest a battle can be) to 1600-0, the most disparate final score.

4. Analysis

We captured a few dozen traces of Arathi Basin battles, and retained 13 where our observers were mostly at their assigned posts. The rendered movement traces for battle 980 is shown in Figure 6. We analyzed our data to verify its correctness, and to provide information for others to evaluate suitability of avatar movement models used for evaluating DVE's. The three main phenomena we wanted to investigate were: appropriateness of waypoint models for guiding movement, existence of hotspots for hotspot-based movement models, and grouping/flocking.

Waypoints. We define waypoints as fixed navigation markers used through all battleground instances. We expected flags and graveyards to be strong candidates for waypoints for movement models, along with geographical choke-points.

Hotspots. Hotspots are situational gathering points in a map, where a disproportionate number of avatars spend time during a given battle. The map has natural hotspots in the form of avatar starting locations and flags. We were curious if other hotspots would show up, and if hotspots were consistent across battles.

Flocking and grouping. Logic dictates there should be significant grouping in movement. All avatars for each team begin at the same point, their faction base, and are released simultaneously. For avatars who die – typically every avatar several times per battle - resurrection is synchronized, with all waiting ghosts at each graveyard resurrected every thirty seconds. Battle dynamics incent avatars to group as well, to maintain numeric superiority.

Before describing our findings, it is worth discussing overall battleground and avatar participation characteristics.

4.1 Avatar Participation Characteristics

We had a series of qualitative questions. First, we wanted to provide an estimate of turnover in the battleground population. In other words, were there joiners and leavers? If so, how long was a typical session? Also, we knew we were missing movement for some avatars some of the time, and wanted to quantify the missing data.

Table 1 summarizes this information for each of the thirteen battles we analyzed, excluding our two observers. 'Lost by' shows the score difference between the winning and losing teams. 'Avatars' shows the number of unique avatars recorded during the battle. The battleground allows in a maximum of 30 simultaneously present avatars, 28 factoring out our observers, but departures can be replaced. 'Duration' gives the total time of each battle in seconds, from when avatars are released from their base to when one team wins and the battle concludes. 'Average play' gives the percentage of the total battleground duration an average participant played. Since many enemy avatars were not observed until the observers were in position – up to two minutes after the match started - this number is biased downwards. 'Average recorded' shows the percentage of avatar participation time successfully recorded for that battle. Average recorded was calculated by summing the total seconds played by all avatars, and subtracting out gaps in the traces for each avatar.

Our data set includes a good sampling of battle scores, ranging from the largest to smallest possible difference, with an average difference of 910. The average battle had 35 unique participants, each present for an average of 69% of the battle. Participant turnover was on average 25% during the course of a battle. We recorded a total of 392 unique avatars. We successfully recorded movement and position for avatars 73% of their participation time. As mentioned earlier, gaps were caused primarily by avatars

becoming invisible and therefore undetectable, and observers being killed and temporarily out of range of some avatars.

The remainder of this section describes relevance of waypoint, hotspot, and grouping models to DVE player avatar movement.

4.2 Waypoints

Waypoints are fixed points in the environment used as intermediate or final destinations for linear navigation. Drawn graphically, the path for an avatar following waypoints would resemble a series of straight line segments connecting the waypoints visited in the sequence they are visited. Each waypoint would typically be a location where the avatar changes direction.

Intuitively, strong waypoint candidates for the Arathi Basin battleground are graveyards, flags, and points on the optimal (non-water, non-cliff) routes between graveyards and flags. While we can manually identify such waypoints, there is no guarantee our identification would be correct. Instead, we took the more general analytical approach outlined in (Mitterhofer, 2009).

Using this algorithm, waypoints are extracted from movement traces by a combination of two strategies: *k-means⁺⁺ cluster analysis* and *path simplification*.

Avatar movement traces consist of a series of points which can be joined together to form a sequence of lines. If avatars are using waypoints for navigation, there should be clusters of line endpoints at the waypoints where some avatars change direction. If no avatars change direction at a given waypoint, it is not actually a waypoint.

Assuming waypoints are used for navigation, their presence can be obscured by human error in movement. Small diversions as avatar controllers delay turning or move in a non-optimal path can confound automatic detection of waypoints.

To mitigate this variation - and as suggested by Mitterhofer et. al. - we perform line simplification using Douglas-Peucker line simplification algorithm. This recursive algorithm reduces line complexity by approximating complex polylines with simpler, albeit less accurate lines with fewer points. The algorithm was invented nearly forty years ago, and is still considered one of the best general line simplification algorithms.

Lines are simplified as follows: Given an input path $P_{1,k}$ of points $\{p_1, \dots, p_k\}$ and an error tolerance E expressed in the same scalar system as the points:

1. Find the point p_j in $\{p_1, \dots, p_k\}$ which lies furthest from the line (p_1, p_k) , and call its distance from the line $d_{j,1,k}$.
2. If $(d_{j,1,k} \leq E)$ the line is simplified to two points, $\{p_1, p_k\}$ and processing is complete. Otherwise:
3. The line is simplified to the union of the results of applying the Douglas-Peucker algorithm to $P_{1,j}$, and $P_{j,k}$. This line will consist of at LEAST the points $\{p_1, p_j, p_k\}$.

Line simplification significantly reduces the number of points required to represent an avatar's movement, making it easier to find clusters of points. Indeed, we successfully used this technique with an error tolerance of 30 yards to simplify the movement trace shown in Figure 4 to that in Figure 5. We then applied *k-means⁺⁺ cluster analysis* to try to cluster the majority of remaining points into waypoints.

In our case, the *k-means⁺⁺* algorithm chooses clusters as follows:

First, choose a set of k seed cluster centers. We do this by first choosing a random point from the data set, and then for $(k - 1)$ iterations:

1. For each data point not in the cluster center set, calculate a probability of being selected equal to the square of the distance from the candidate to the closest cluster center set member.
2. Choose a random point to add to the cluster center set from the candidates, weighted by the probability of selection.

Next, execute normal k -means cluster analysis using this set of k seed cluster centers as the starting point.

Even with dramatic simplification of the avatar movement paths via the Douglas-Peucker algorithm, we found that k -means⁺⁺ cluster analysis was unable to identify a consistent set of waypoints to describe avatar movement. Successive runs identified potential waypoints, but many of these diverged widely between different runs, depending upon the initial waypoints selected according to algorithm. Reviewing the point cloud for both simplified and original movement traces battleground revealed the cause. While there are clear asymmetries in point density for describing movement traces, no reasonable number of clusters can encapsulate the majority of simplified points, even when dramatic simplification tolerances are used (e.g. allowing errors greater than an interaction interval).

We conclude that waypoints are not appropriate for characterizing general player avatar movement in World of Warcraft battlegrounds. As an aside, this conclusion supports Mitterhofer et. al.'s implication that waypoints are not a good fit for characterizing typical player-controlled avatar movement.

4.3 Hotspots

This section describes our efforts to characterize avatar movement patterns using hotspots. Hotspots are portions of the battleground where avatars spend the most time. Mathematically, hotspots are determined by dividing the map into equal sized cells, summing the number of seconds spent by each avatar in each cell, and designating the k cells with the highest totals as the k hotspots. Cells in the 8-neighbors of an existing hotspot are precluded from being designated as hotspots, to prevent runs of adjacent hotspots. In essence we implemented the technique for hotspot detection proposed and applied in (Thawonmas, Kurashige, & Chen, 2007), described earlier.

We expected hotspots at each of the five flags because they are game objectives with most activity occurring at them, and at the seven graveyards because participants die many times each battle and await resurrection in them. We did indeed find that most of the first 12 hotspots contained either a flag or graveyard. However, the order of 'hotness' of these points of interest varied from battle to battle. In some battles particular graveyards and flags were never included in the top hotspots. Non-graveyard, non-objective hotspots *were* encountered, reflecting battles where significant concentrations of avatar activity occurred away from flags and graveyards. Although such hotspots were seen in most battles, we were unable to determine a common characteristic for them. While specific non-objective hotspots were present in more than one battle, none were present in the same location for a majority of battles.

Example hotspots are shown in dark grey or black in the player time-density maps in Figure 7, with the five most active hotspots labeled 1 through 5. The trace on the left shows hotspots at the stables, mine, and blacksmith flags. The fourth hotspot is on the path from the alliance base to the stables, and the fifth at the farm flag. The trace on the right shows hotspots at the farm flag and graveyard, the route between the farm and blacksmith, and then the lumber mill flag and blacksmith graveyard.

As we forecast, hotspots were usually located where there was heavy contention over a flag, or a battle which migrated from flags towards arriving combatants. Based on this, we believed an adaptive hotspot-based model, taking into account current populations at hotspots, would be useful for gaining insight into avatar movement within battlegrounds, possibly allowing a generative movement model.

One weakness we noticed in this hotspot model was the tendency for cells which would otherwise have been hotspots to be excluded because of their abutment against a higher-ranked cell. This biased calculations, as it prevented hotspots from being correctly located.

To compensate for this, we extend the cell-based hotspot model using a centroid which need not be aligned on a hotspot boundary. Centroids are circles centered on a hotspot, but without the restriction to align at the granularity of cell boundaries. An example showing centroids overlaid on traditional cell-based hotspots is shown in Figure 8.

To calculate centroids with a radius of one interaction interval (a diameter of 60 yards), we perform normal hotspot calculation with cell width one third of the desired diameter, e.g. 20 yards. For each 9-square centered on a given hotspot, we calculate the center of mass for avatar movement in that 9-square, weighted by avatar dwell time in each cell. The resulting 30-yard diameter hotspot does a better job of covering the actual hotspot than a simple cell-based designation. As you can see in Figure 8, the centroids are often not centered on the hotspots, encapsulating a different set of territory than a simple cell-based hotspot mechanism allows.

Using these revised hotspot definitions, and other points of interest, we were able to perform deeper analysis of avatar movement. In particular, we evaluated the tendency of avatars to move between points of interest in groups.

4.4 Group Movement

We define group movement as the coordinated movement multiple avatars between points of interest. Points of interest in our model include avatar spawn points – as many trips start at spawn points – and at hotspots, typically flags and the sites of heavy battles.

There is strong incentive for players to group within a battleground: a lone combatant has very little chance of defeating multiple enemies. A fight between members of two factions almost always goes to the force with greater numbers. Success in combat implies success at controlling flags, which in turn leads to battleground victory and greater in-game rewards.

Regardless of the benefits of group movement, two factors provide a disincentive for forming and maintaining such groups. First, the difficulty of coordinating group formation and maintenance using default communication channels such as text chat. Second, the conflict between group and individual goals: without an accepted group leader, these often diverge.

Even when a group is well coordinated (for example via a voice over IP solution such as Ventrilo) and has an acknowledged leader, maintaining group coherence is difficult. If a group member is slain, they become a ghost, and must resurrect and travel back to the body of the group. Barring enemy interference, this can take up to two minutes, half the battle length in some cases. In the meantime, the group typically continues towards its next objective, with subsequent deaths splintering the group further.

The question we seek to answer: is group movement a reasonable model to apply to describe the majority of travel between points of interest? (Miller & Crowcroft, 2009) proposes a general grouping metric called ‘affinity.’ We extend this metric with a more precise definition.

Leveraging the work on hotspot identification above, we define the concept of a *journey*. A journey is a trip between points of interest which are at least two interaction intervals (60 yards) apart. Taking a circuitous route between two nearby hotspots does not typically constitute a journey, but travelling between two hotspots whose closest edges are 60 yards apart – without passing through a third point of interest – does. Two avatars which move between the same two points of interest maintaining a distance of no more than one interaction interval are considered as having affinity for that journey.

For our analysis we chose the seven graveyards / spawn points as fixed points of interest, and another seven centroids (including hotspots which happen to overlap spawn points) chosen by the algorithm described in the previous section as per-battleground instance points of interest. The centroids included at least four of the five flags in every battle, though with the centroid at slightly different locations. Most battles had at least one non-flag, non-graveyard centroid.

As mentioned previously, we have good trace data for 73% of overall avatar sessions. We divided movement segments for which we had data into four categories:

1. Inter-centroid movement, journeys between different centroids. Some are degenerate cases (return to the same centroid) or too short to be eligible for our calculations, but still fall in this category.
2. Centroid-anchored movement. These segments typically departed from a centroid, then terminated before reaching another centroid. An avatar which spawns at a graveyard and is killed before it can reach another centroid would be in this category.
3. Intra-centroid movement, movement segments within a graveyard or centroid which do not leave that centroid.
4. Extra-centroid movement. These segments are usually artifacts of missing data. The battleground start is a special case. If the starting bases are not picked as centroids by automatic hotspot detection, then the initial path of any avatars killed or departing the battle before reaching a centroid will be in this category.

Table 2 provides a summary of the percentage of recorded time which falls into each of the four movement categories. Of the more than 60 avatar-hours of traces obtained, approximately 32% is inter-centroid journey data. This gross figure includes inter-centroid journeys which are too short to qualify for further analysis: the actual percentage of eligible inter-centroid journey traces is 2/5 of the inter-centroid journey traces, or 14% of overall captured traces. This resulted in 951 candidate journeys for analysis.

We ran affinity analysis across all of our battleground traces, with the results summarized in Table 3. Note that partial affinity figures are excluded: either an entire journey is considered as having affinity between two or more avatars, or none of it is.

Our results confirm and quantify earlier findings. Despite incentives, the majority of journeys in battlegrounds are made alone, rather than in a group. Only 17% of journeys by time are made in a group. The number looks slightly better when viewed in terms of journey counts rather than journey durations, but is still less than 25% affinity. Interestingly, the disparity between affinity journey seconds and journey count percentages indicates that longer journeys are less likely to be made in groups than shorter ones.

5. Conclusions and Future Work

This paper provides an analysis of real DVE movement traces using existing techniques. It also proposes some refinements of those techniques.

We present movement data taken from battleground traces of World of Warcraft, the world's most popular DVE. We analyzed this data in terms usually used for generating DVE avatar movement patterns: waypoints, hotspots, and group movement.

We showed that player avatar movement cannot be characterized using existing waypoint generation models. We implemented an existing hotspot detection model, and augmented it with a centroid detection algorithm to provide more appropriate hotspot descriptions.

We proposed four categories of avatar journeys within battlegrounds, and analyzed inter-centroid journeys. We found that more than three quarters of avatar journeys between points of interest in battlegrounds are made alone, despite strong game incentives to move in groups.

While this work answers some questions, it raises others.

We've captured and parsed captures of real avatar movement in WoW, and shown that waypoint models are not sufficient to characterize them. However, we haven't provided a full solution to replace the waypoint model. Is there a mathematical model or algorithm which accurately characterizes avatar movement, and can such a model allow generation of avatar movement and activity traces consistent with our findings? Such a model would ease the task of finding a representative set of traces for evaluation. Based on our results, hotspots seem to provide a good foundation to start from. We do not intend to research in this direction, but the research community could benefit from the efforts of anyone who does.

We briefly mentioned preliminary work in characterizing traffic and movement patterns in four categories of WoW avatar activity: PvE, capital cities, PvP, and Instances. We intend to continue this work, providing a more comprehensive comparison of these categories, with formal analysis beyond network traffic byte counts to confirm that our categories – or those proposed in earlier literature – are correct. If successful, this characterization can provide guidance for design and evaluation of DVE message propagation algorithms.

Our avatar movement traces allow realistic evaluation of existing area of interest-based message propagation. We plan to use simulations to comparatively evaluate archetypical message propagation models, such as client-server and publish-subscribe. Realistic message attribution combined with movement traces from an actual DVE should allow high-fidelity examination of message propagation both in 'normal' cases, and in times of peak activity and flux, for example with many avatars moving in intersecting paths, unusually high traffic, and so on.

We believe an understanding avatar movement patterns for real-world DVE's is critical for designing solutions to the challenges of large-scale DVE's. We hope this paper provides some insights in this domain.

6. Acknowledgements

We would like to thank the anonymous reviewers for both NetGames 2009 and this special issue of IJAMC whose comments led to significant improvements in the paper and guidance in the research. We would also like to thank Richard J. Black for allowing us to use his NetMon trace parsing code, which significantly sped our work.

7. References

- Backhaus, H., & Krause, S. (2007). Gain and Loss of a Voronoi-based Peer-to-Peer Approach for MMOG. *NetGames '07: Proc. ACM Press*.
- Blizzard Entertainment. (2008, 11 21). *World of Warcraft Subscriber Base Reaches 11.5 Million Worldwide*. Retrieved 04 15, 2010, from Blizzard.Com: <http://us.blizzard.com/en-us/company/press/pressreleases.html?081121>
- Douglas, D. H., & Peucker, T. K. (1973). Algorithms for the Reduction of the Number of Points Required to Represent a Line or its Caricature. *The Canadian Cartographer*, 10(2), pp. 112-122.
- Fan, L., Taylor, H., & Trinder, P. (2007). Mediator: A Design Framework for P2P MMOGs. *NetGames '07: Proc. ACM Press*.

- Hu, S.-Y., & Liao, G.-M. (2004). Scalable peer-to-peer networked virtual environment. *NetGames '04: Proc.* (pp. 129-133). ACM Press.
- Jardine, J., & Zappala, D. (2004). A hybrid architecture for massively multiplayer online games. *Netgames '04 Proc.* (pp. 60-65). ACM.
- Krause, S. (2008). A Case for Mutual Notification: A survey of P2P protocols for Massively Multiplayer Online Games. *NetGames '08 Proc.* ACM Press.
- Lui, J. C., & Chan, M. F. (2002). An efficient partitioning algorithm for distributed virtual environment systems. *Parallel and Dist. Systems, IEEE Trans. on*, 13(3), 193-211.
- Matsumoto, N., Kawahara, Y., Morikawa, H., & Aoyama, T. (2004). A scalable and low delay communication scheme for networked virtual environments. *IEEE GlobCom Workshop 2004*, (pp. 529-535).
- Matteo, C. D. (2007). Dynamic Clustering in Delaunay-Based P2P Networked Virtual Environments. *Netgames '07 Proc.* (pp. 105-110). ACM Press.
- Miller, J. L., & Crowcroft, J. (2009). Avatar Movement in World of Warcraft Battlegrounds. *Network and System Support for Games*. IEEE.
- Mitterhofer, S. a. (2009). Server-Side Bot Detection in Massively Multiplayer Online Games. *IEEE Security and Privacy*, 29 - 36.
- Morillo, P., Orduna, J. M., Fernandez, M., & Duato, J. (2005). Improving the Performance of Distributed Virtual Environment Systems. *Parallel and Dist. Systems, IEEE Trans. on*, 16(7), 637-649.
- Pittman, D., & GauthierDickey, C. (2007). A Measurement Study of Virtual Populations in Massively Multiplayer Online Games. (pp. 25-30). ACM Press.
- Rueda, S., Morillo, P., & Orduna, J. M. (2007). A Saturation Avoidance Technique for Peer-to-Peer Distributed Virtual Environments. *Cyberworlds '07 Proc. 0*, pp. 171-178. IEEE Computer Society.
- Thawonmas, R., Kurashige, M., & Chen, K.-T. (2007). Detection of Landmarks for Clustering of Online-Game Players. *The International Journal of Virtual Reality*, 11-16.
- Woodcock, B. S. (2008, 04). MMOG Subscriptions Market Share - April 2008. *MMOGChart.com*(<http://www.mmogchart.com/Chart7.html>).
- Yamamoto, S., Murata, Y., Yasumoto, K., & Ito, M. (2005). A distributed event delivery method with load balancing for MMORPG. *NetGames '05 Proc.* (pp. 1-8). ACM.

Biographical Notes

Jon Crowcroft is the Marconi Professor of Networked Systems in the Computer Laboratory, of the University of Cambridge. Prior to that he was professor of networked systems at UCL in the Computer Science Department.

He is a Fellow of the ACM, a Fellow of the British Computer Society and a Fellow of the IEE and a Fellow of the Royal Academy of Engineering, as well as a Fellow of the IEEE. He was a member of the IAB 96-02, and went to the first 50 IETF meetings; was general chair for the ACM SIGCOMM 95-99; is recipient of Sigcomm Award in 2009.

John L. Miller is a software architect for Microsoft Research, Cambridge, specializing in distributed systems and distributed system security. He helped develop and test several versions of Windows, usually working on peer-to-peer and multimedia infrastructure and services. He is also a PhD student at the University of Cambridge working in distributed virtual environment security and scalability.



Figure 1 - Arathi Basin map



Figure 2 - Arathi Basin from the Lumber mill Plateau



Figure 3 - Placement of observers in Arathi Basin

Lost By (points)	Avatars	Dur (s)	Avg Play	Avg Rec.
10	36	1423	72%	81%
300	38	1296	67%	75%
420	46	1208	52%	62%
720	36	1015	69%	75%
870	36	957	63%	62%
950	33	671	71%	69%
960	37	951	69%	76%
980	36	891	70%	78%
1050	33	885	79%	78%
1180	37	658	61%	60%
1370	36	765	66%	65%
1490	32	583	78%	83%
1600	20	266	76%	82%
AVERAGE	35	890	69%	73*%

Table 1 – Avatar Participation Summary

Lost By (points)	Recorded (s)	Inter Centroid	Centroid Anchored	Intra Centroid	Extra Centroid
10	30,786	9,922	4,499	15,764	601
300	25,870	7,299	6,367	10,577	1,627
420	19,470	6,169	3,751	9,024	526
720	20,056	6,596	3,214	9,397	849
870	14,314	4,432	3,166	5,935	781
950	11,100	3,304	3,021	4,333	442
960	19,531	5,343	3,582	8,641	1,965
980	18,245	6,642	3,037	7,599	967
1050	19,000	6,743	4,227	6,890	1,140
1180	10,409	2,498	1,861	5,045	1,005
1330	12,285	3,717	3,265	4,221	1,082
1390	12,431	4,760	2,866	4,584	221
1600	3,442	1,278	719	1,415	30
AVERAGE	16,688	5,285	3,352	7,187	864
		32%	20%	43%	5%

Table 2 – Avatar Movement Trace Categories

Lost By (points)	Journey (s)	Affinity (s)	Journey (count)	Affinity (count)	% Affinity (s)	% Affinity (count)
1590	5,769	735	174	41	12.7%	23.6%
1300	3,296	754	96	24	22.9%	25.0%
1180	2,940	405	95	19	13.8%	20.0%
880	2,250	331	86	20	14.7%	23.3%
730	1,205	52	34	4	4.3%	11.8%
650	1,154	64	44	7	5.5%	15.9%
640	2,279	519	76	24	22.8%	31.6%
620	3,250	679	107	30	20.9%	28.0%
550	2,663	748	76	26	28.1%	34.2%
420	821	239	28	10	29.0%	35.7%
270	1,594	274	57	13	17.2%	22.8%
210	1,882	304	57	13	16.1%	22.8%
0	703	94	21	3	13.4%	14.3%
AVERAGE	2,293	400	73	18	17.4%	24.6%

Table 3 – Avatar Journey Affinity




Figure 4 – Point Cloud for Battle 1590



Figure 5 – Simplified Point Cloud (+/- 30 yard tolerance)




Figure 6 - Battle 980 movement paths




Figure 7 - Activity density maps from two different battles




Figure 8 - Density Map with Centroids