

# Peer-to-Peer-based Infrastructure Support for Massively Multiplayer Online Games

Simon Rieche, Klaus Wehrle  
Distributed Systems Group  
RWTH Aachen University  
{rieche,wehrle}@cs.rwth-aachen.de

Marc Fouquet, Heiko Niedermayer, Leo Petrak, Georg Carle  
Computer Networks and Internet  
University of Tübingen  
{fouquet,niedermayer,petrak,carle}@informatik.uni-tuebingen.de

**Abstract**—Online games are an interesting challenge and chance for the future development of the Peer-to-Peer paradigm. Massively multiplayer online games (MMOGs) are becoming increasingly popular today. However, even high-budget titles like World of Warcraft that have gone through extensive beta-testing suffer from downtimes because of hard- and software problems. Our approach is to use structured P2P technology for the server infrastructure of MMOGs to improve their reliability and scalability. Such P2P networks are also able to adapt to the current state of the game and handle uneven distributions of the players in the game world. Another feature of our approach is being able to add supplementary servers at runtime. Our system allows using off-the-shelf PCs as infrastructure peers for participation in different game worlds as needed. Due to the nature of the Economy of Scale the same number of hosts will provide a better service than dedicated servers for each game world.

## I. INTRODUCTION

Multiplayer games played over the Internet have become very popular in the last few years. An interesting subcategory are the so-called massively multiplayer online games (MMOGs) that allow thousands of player characters to share a single game world. Such a world is usually run on a high-performance and high-availability server cluster. However, even with games that have been extensively beta-tested, downtimes of several hours because of hard- or software failures are not uncommon. Downtimes, especially in the first few weeks after the release, can negatively affect the image of the game and the company that created it.

Traditionally, a cluster of servers contains one virtual world of a MMOG. Such infrastructure is inflexible and error-prone. One would rather like to have a system that allows disconnecting a server at runtime while others take over its tasks. Server-based MMOGs can have performance problems if players are concentrated in certain parts of the game world or some worlds are overpopulated. Thus, there is also a need for load balancing mechanisms. Peer-to-Peer (P2P) systems quite naturally support the use of load balancing.

In this paper we use a structured P2P technology for the organization of the infrastructure and thus for the reduction of downtimes in MMOGs. By using a CAN-based approach we split the game world in disjunctive zones and distribute them on different nodes of the P2P network.

Online games are an interesting challenge and chance for the future development of the P2P paradigm. A wide variety

of aspects of only theoretically solved and especially yet completely unsolved problems are covered by this application. Security and trust problems appear as well as the need to prevent cheating. The application is not as tolerant to faults as instant messaging or file sharing. Consistent data storage is a problem, decisions and transactions have to be performed in a decentralized way. Moreover, the P2P network is not used as pure lookup service, but more as a communication and application-specific social structure.

The rest of this paper is organized as follows: First we discuss related work in Section II and give a brief introduction to MMOGs and their challenges in Section III. Section IV shows our approach to use structured P2P Systems for MMOGs and section V the evaluation with player traces from a real MMOG. Finally, Section VI provides conclusions.

## II. RELATED WORK

Some efforts have been undertaken to design a MMOG on a P2P basis, with the server tasks being shared among the player's PCs. In [1] and [2] the game world is divided into zones, some peers become zone owners that take responsibility for computing the server tasks for such a zone. While this approach is fascinating on one hand, it suffers from a number of practical problems. Players constantly connect to and disconnect from the game, often without warning if the PC crashes or suddenly gets disconnected from the network, so one always needs backup machines to replace disconnected zone owners. Allowing player's computers to calculate parts of the game mechanics makes it harder to avoid cheating. Another problem is that persistent player data, for example the progress of the player's characters in a role-playing game, need to be saved in a way that makes sure that no information is lost, as players may have invested a lot of work into the game. At the same time one has to prevent players from cheating by modifying this data—which would be possible if it was stored on the player's local hard drives. This suggests that some kind of infrastructure provided by the game manufacturer would still be needed, a full P2P approach does not appear practical. Another challenge is the low upstream bandwidth of most current Internet connections, probably only peers with a good connectivity could be considered for becoming zone owners.

Solipsis [3] uses a different approach, as it tries to build a P2P network-based on the neighborhood relations of the

player’s avatars. Each peer has direct connections to all other peers that are in visible range of the player’s avatar. There is a real implementation of Solipsis, however currently it is little more than a distributed chat client. If one wanted to make it a “real” MMOG, one would be confronted with the same problems as described above. It is especially difficult to make such an approach cheat-proof [4].

Many algorithms exist for load balancing in structured P2P systems [5]–[8]. However, most of these systems are not applicable for online games. The virtual server approach [5] is based on the idea of managing multiple partitions of a Distributed Hash Table’s (DHT) address space in one node. Thus, one physical node may act as several independent logical nodes. Each virtual server will be considered by the underlying DHT as an independent node.

### III. MASSIVELY MULTIPLAYER ONLINE GAMES

MMOGs differ from other games mainly in the number of simultaneous players. They allow thousands of players to share a single game world. This type of games has its origin in the text-based Multi User Dungeons (MUDs) that first appeared in the 1970s. Blizzard launched *World of Warcraft* [9] in 2004, which became a worldwide best-seller.

Most of today’s MMOGs titles are role playing games, the term *Massively Multiplayer Online Role Playing Games* (MMORPGs) is well-known in the gaming industry. In this game type, the player’s character earns experience while adventuring through the game world. In role playing games, combat itself is more a matter of luck and strategy than of quick reflexes. The player has to choose the abilities that he wants to use. Using skills, magic spells or weapons also requires luck, because whether an opponent is hit and how bad he is hit is computed randomly. This is why MMORPGs are relatively tolerant with respect to delay; it is by far not as important to have a “low ping” as with an ego-shooter.

#### A. Server Tasks

The central server cluster in a traditional MMOG has to perform various different tasks, including:

- **Management of player positions in the world:** The server has to know the position of each player and distribute this information to other nearby players.
- **Management of the world:** Controlling monsters, non-player characters, treasures, the weather and all other dynamic aspects.
- **Combat:** To avoid cheating, the effects of actions taken by the players in combat should be calculated on a server.
- **Chat:** Usually a global in-game-chat is implemented on the server.
- **Accounting for the players:** The player’s characters must be surely saved. This is a task that makes the creation of a real P2P multiplayer game—with all clients acting as servers and without a central infrastructure—very hard.




Fig. 1. The architecture for the P2P-based infrastructure for MMOGs. Each infrastructure host is responsible for different virtual zones in the world.

#### B. Challenges

The first and possibly greatest challenge that the server infrastructure of a MMOG faces are the first few days after the game has been launched. An eagerly expected game like *World of Warcraft* can have a huge number of players on the servers after a few days. For such games it is quite common to have different areas for players with different levels of experience. Thus, an additional challenge here is that all those new players will concentrate in the beginner areas of the game and the load on the servers will not be balanced equally. This is a general problem that one also has to face when the first rush on the game is over. Players do not evenly distribute over the world, they tend to form large groups to raid villages of enemy guilds or realms. Sophisticated mechanisms to balance the load are required here. It would be especially advantageous if one could simply connect another server to the cluster in a high-load condition and the load would be split automatically.

The leading MMORPG in the market—Blizzard’s *World of Warcraft* with over 6 Million players worldwide—is constantly suffering from overloaded servers. Players are complaining about lag, downtimes of the game worlds and long queues when trying to connect to the game. As a consequence, players in China even went on strike in March 2006 [10].

Another challenge that needs to be addressed is the reliability of the game. Games occasionally experience server instabilities caused by hard- or software failures. In this situation one wishes to have the possibility to simply disconnect a machine that shows a problem, the load should automatically be split up among the remaining servers. Other challenges that are out of our scope in this paper are data management, various hand-overs, ways to guarantee a certain level of consistency, and many more.

### IV. P2P-BASED INFRASTRUCTURE SUPPORT FOR MMOGS

This section shows our approach to use structured P2P technology for the organization of the infrastructure for MMOGs. The game world is split in disjunctive zones and distributed on different nodes of the P2P network.

#### A. Architecture

As [2] and [1], we use the fact that only a limited region of the game world is interesting for the player, as his character

can only see a limited area and also has a limited movement speed. This allows splitting the game world into regions, whereas the current region and a small number of neighboring regions are in the player's visible range.

The game world is defined by a map, which is split into disjunctive zones in  $d$  dimensions (mostly two or three) and distributed on different nodes of the P2P network. As CAN is a structured P2P System based on an  $d$ -dimensional coordinate space, it is best-suited for map-based scenarios. Our approach does not use the DHT functionality of CAN, only the functionality of a structured P2P system is needed, like adding nodes or maintenance of the coordinate space. Figure 1 shows a distribution of servers over the map, which is mapped into a zone-based structured P2P system. Since the zones at the edge of the identifier space are connected with the opposite zones for routing and load balancing, also game worlds simulating a terrestrial globe are possible. The players are assigned to the servers according to the position in the game world.

The game world is distributed on a server infrastructure. This infrastructure is not necessarily located in a single data center. However in most cases it makes sense to locate the peers that are responsible for adjacent regions in the same network to minimize delay and costs for internal traffic. But our approach could also be used for a super-peer network in a more distributed setting for more delay-tolerant games.


Besides using the CAN address space as a map to store location-specific information, one can also still use it as a traditional DHT to store data that is not directly related to a location in the game world. For example, to realize the global chat functionality that is common with MMOGs, one needs to look up player characters by name and not by their position in the game world. This can be achieved by simply storing a key-value-pair in the CAN DHT.

### B. Accounting

Accounting can be done via a standard Authentication, Authorization, and Accounting (AAA) [11], [12] server infrastructure. This can be enforced since the game control is not directed to untrusted peers, but to infrastructure nodes that organize as P2P network. Also it is possible to store the player's characters permanently on the game hosts or a dedicated machine and to make backups of this data regularly.

### C. Availability

Structured P2P systems include mechanisms, that allow neighboring peers to take over the area of a server that has disconnected [13]. However the state information that was stored about this area would be lost, if it has not been replicated somewhere in the network. As the ratio of disconnects is far lower in our case than with a player P2P network—the game servers should be more reliable than users peers—the replication problem is not as serious as in [2] though. A server that owns area  $A$  can simply propagate changes to the game state to one or more other servers  $B - E$  that own a neighboring area (cf. Figure 2(a)). It would be appropriate to choose nodes with low load to store the replicas. Note that parts of this information are also useful in  $B - E$ , as some



(a) Replication mechanism: Nodes propagate their state to neighbors. If a host suddenly disconnects, a neighbor can take over immediately.

(b) View of a player character: A node has a view across zones borders. The messages are send trough the host, responsible for each zone.

Fig. 2. Replication mechanism and view of player characters in the structured P2P infrastructure.

of players in these areas may be close to the borders and therefore able to see what happens in area  $A$  (cf. Figure 2(b)). Other neighbors receive information that might be interesting for their players, e.g. information about objects and characters that are within visible range of the zone border.

If a failure occurs, the host with the lowest load is chosen among the machines that replicated the information, e.g. the host that owns area  $D$ . Since the game state for area  $A$  is already known at this node, no further copy is needed to run the game. Then, the load-balancing algorithms presented below can be used to re-distribute the load on the physical servers.

### D. Pooling

Any P2P approach has to support joins and leaves of nodes. As a consequence, new hosts can be added to the P2P infrastructure when needed. Suitable hosts are not only high-end servers; the game could also run on standard PCs. On the basis of the virtual server concept hosts may have virtual servers in different game worlds. This has two major implications. First, the downtime of one or only few servers will probably not be able to stop the game. The availability of the game will be much higher than today. Only when a new version has to be installed on all servers and clients that can not interoperate with the old version there will be a definite stop. Second, since hosts can be pooled and used for the game or instance that currently needs them most the so-called Economy of Scale<sup>1</sup> comes into play. The consequence is a better service with the same number of resources or a potential to reduce costs.

### E. Security

The main security functionality can be realized using a Authentication, Authorization, and Accounting (AAA) [11],

<sup>1</sup>The Economy of Scale in this context is also called Multiplexing Gain in the theory of communication networks. It states that an increase in demand does not necessarily need the same increase in capacity to achieve the same service quality.

[12] infrastructure. These dedicated AAA servers grant peers access to the world and registered players access to the game. Since the infrastructure peers are controlled by the game company (or a community of trusted peers) all nodes in the network are trusted. The peers have to be authenticated since an outside peer could try to enter the network and attack it. Players have to be authenticated, but this, too, can be solved by the AAA infrastructure.

#### F. Load Balancing

Load balancing is an important issue, as one can not expect the players to distribute uniformly in the game world. Some places in the game world, cities for example, can be expected to be more populated than areas in the countryside. This knowledge could be used to statically configure the servers of an MMOG in a way that distributes the expected load evenly. But in many games, it is common to form large groups for raiding enemy territories for example. In such a case, dynamic load balancing is required.

The virtual server approach [5] is based on the idea of managing multiple partitions of a structured P2P address space in one node. Thus, one physical node may act as several independent logical nodes. Each virtual server will be considered as an independent node by the underlying structured P2P System. Within a CAN system, one virtual server is responsible for a zone of the address space, whereas the corresponding physical node may be responsible for several different and independent zones. The basic advantage of this approach is the simplicity to place and transfer virtual servers among arbitrary nodes. This operation is similar to the standard join or leave procedure in a structured P2P system. Every participating node manages many virtual servers so load can be moved between nodes by moving a whole virtual server to another node. Additionally, zones with too many players inside can be split and one part can be sent to another server. Also, adjacent zones with low numbers of players can be merged for a lower number of internal messages.

When an area has to be split because it is overpopulated, one of five algorithms is used to decide how the split should be performed.

- **SplitCenter** splits areas along the center horizontally or vertically in exchange. If area  $A$  was split into areas  $B$  and  $B'$  horizontally, then area  $B$  will be split into  $C$  and  $C'$  vertically.
- **MaxDistToBorders** also splits along the center. This algorithm decides whether to split horizontally or vertically by maximizing the average distance of the players to the newly created border. The idea behind this is, to minimize internal traffic as looking and walking over borders always creates overhead.
- **IntelliDistance** again splits along the center and decides whether to split horizontally or vertically, such that as few players as possible can see the other side of the new border. The rationale behind this is the same as with **MaxDistToBorders**.
- **EqualNumbers** splits along the center and decides

whether to split horizontally or vertically in a way that makes the number of players in the new regions is as equal as possible. The expectation is that this algorithm will perform better than **SplitCenter** in terms of load-balancing.

- **VarAreas** splits horizontally or vertically as **SplitCenter** does. However it places the border not along the center but at the centroid of the players. As with **EqualNumbers** the expected result is better load-balancing.

## V. EVALUATION

In this section we evaluate our approach with respect to the load balancing issues. First, we briefly introduce the simulation methodology. Secondly, we discuss how players move in the game world and finally present our simulation results.

The two major criteria for the evaluation are:

- **The variance of the physical server load** is the variance of “computational” load we determined for the various physical servers. Ideal would be 0, i.e. same load for all.
- **The number of internal messages** is a way to measure the overhead introduced by distributing the game world on multiple servers. Internal messages are exchanged between physical servers, they are caused by players looking or walking over area borders and when players are assigned to another server because a virtual server has been moved.

#### A. Player Behavior

We evaluated our approach with a simulation using artificially generated as well real traces of user-movement:

- **Random walk trace data** is a collection of generic data represents movement of users according to the Random Walk Mobility Model. It covers up to 300 users with four hours of movement time.
- **Random waypoint trace data** has the same properties by using the Random Waypoint Mobility Model.
- **Freewar trace data** consists of real player-movements traced in Freewar [14] over a period of up to five hours. Approximately 400 players were online in this period of time, the number of concurrent players varies over time since users join or leave the game. So there is a dynamic number of players over time. Also the Freewar traces show an extremely uneven distribution of the players over the world with hot spots in cities.

#### B. Results

Figure 3 shows the coefficient of variation (CV) of the physical server load, with all five algorithms and with real player-traces from Freewar. One can see that *IntelliDistance* performs best, *SplitCenter* is the second-best algorithm. *EqualNumbers* and *VarAreas* have a very high CV. The reason for this is that those algorithms already achieve a relatively good load-balancing with few split operations, so the game-world is not split into as many areas as with the other algorithms. Figure 3 shows also the CV for artificial “Random Walk” and “Random Waypoint” traces. One can see that the CV is generally much lower here, as the players are distributed more evenly. Again




Fig. 3. The effectiveness of the load-balancing algorithms with real trace data of the MMOG Freewar and artificial trace data.

*EqualNumbers* and *VarAreas* have the highest CV, this time *SplitCenter* was the best algorithm.

In Figure 4 one can see the number of internal messages caused by distributing the game world on multiple peers. Internal messages are triggered by players that are close to an area-border so they can see the other side, by players crossing an area border (handover) and by virtual servers that are moved from one physical server to the next.

The algorithms *MaxDistToBorder* and *IntelliDistance* perform approximately as good as *SplitCenter* with the real player-traces from Freewar. Splitting areas in a way that minimizes the number of players who can see beyond the border in the moment of the split is ineffective. The algorithms *EqualNumbers* and *VarAreas* perform especially well here. At the first glance it appears strange that the algorithms designed for good load-balancing are good with regard to the number of internal messages while they appeared bad regarding the coefficient of variation above. The explanation is that these algorithms split the world more effectively, so fewer splits are required to prevent servers from being overloaded. Therefore the game world is split into fewer areas which reduces the number of internal messages.

Figure 4 also shows the number of internal messages with artificial “Random Walk” and “Random Waypoint” traces. Surprisingly *EqualNumbers* and *VarAreas*, the algorithms with the best results with real player data, perform specially bad here. In the artificial traces, the players are distributed on the map evenly, therefore those two algorithms gain little advantage from their power to split areas. However they have disadvantages when merging regions as - specially with variable areas - the chances of having a common border with the neighboring area are reduced.

This shows also that artificial trace data is of little use for evaluation of MMOGs. Real players concentrate at certain hot spots on the map, they move in groups and sometimes even show a “flash crowd”-like behavior. Therefore “Random Walk” and “Random Waypoint” do not generate results that would be applicable to real games. In a real MMOG, one may want to start splitting the world using the *SplitCenter* algorithm, as it produces good results and appears to be robust against changes of the player distribution. *VarAreas* is even more effective and may be the algorithm of choice if the player distribution is rather uneven and shows hot-spots. But switching between algorithms is also possible at run-time.




Fig. 4. The number of internal messages with real trace data of the MMOG Freewar and artificial trace data.

## VI. CONCLUSIONS AND FUTURE WORK

In this paper we use a structured P2P technology for the organization of the infrastructure and thus for the reduction of downtimes in MMOGs. By using a CAN-based approach we split the game world in disjunctive zones and distribute them on different nodes of the P2P network. Thus, we get the possibility to dynamically connect and disconnect machines to and from the peer-cluster and to load-balance the game accordingly to the actions of the players.

The P2P technology makes it possible to run multiple game worlds on a pool of servers. The physical location of these servers is of minor importance, they do not have to run in the same data center. Placing the servers of a single game world at different locations introduces additional overhead; however this may be justified by enhanced reliability or maybe by an improved locality. Even if one whole location should be disconnected from the network, the servers at other locations could take over seamlessly and without loss of data.

We will continue working on P2P support for MMOGs. Future work includes more a detailed investigation of handovers and state consistency between the servers and even better load-balancing algorithms.

## REFERENCES

- [1] T. Iimura, H. Hazeyama, *et al.*, “Zoned federation of game servers: a Peer-to-Peer approach to scalable multi-player online games,” in *Proc. of NETGAMES*, Portland, OR, 2004.
- [2] H. Lu, “Peer-to-Peer Support for Massively Multiplayer Games,” in *Proc. of IEEE INFOCOM*, Hong Kong, China, 2004.
- [3] J. Keller and G. Simon, “Solipsis: A Massively Multi-Participant Virtual World,” in *Proc. of PDPTA*, Las Vegas, NV, 2003.
- [4] C. GauthierDickey, D. Zappala, *et al.*, “Low Latency and Cheat-Proof Event Ordering for P2P Games,” in *Proc. of NOSSDAV*, Ireland, 2004.
- [5] A. Rao, K. Lakshminarayanan, *et al.*, “Load Balancing in Structured P2P Systems,” in *Proc. of IPTPS*, Berkeley, CA, 2003.
- [6] S. Rieche, K. Wehrle, *et al.*, “A Thermal-Dissipation-based Approach for Balancing Data Load in DHTs,” in *Proc. of LCN*, Tampa, FL, 2004.
- [7] D. Karger and M. Ruhl, “Simple Efficient Load Balancing Algorithms for Peer-to-Peer Systems,” in *Proc. of IPTPS*, San Diego, CA, 2004.
- [8] J. Byers, J. Considine, *et al.*, “Simple Load Balancing for DHTs,” in *Proc. of IPTPS*, Berkeley, CA, 2003.
- [9] Blizzard Entertainment, [www.worldofwarcraft.com](http://www.worldofwarcraft.com).
- [10] Interfax China, “Gamers threaten mass protest against The9’s operation of WoW in China,” 2006.
- [11] C. de Laat, G. Gross, *et al.*, “Generic AAA Architecture,” IETF, RFC 2903, 2000.
- [12] S. Farrell, J. Vollbrecht, *et al.*, “AAA Authorization Requirements,” IETF, RFC 2906, 2000.
- [13] S. Ratnasamy, P. Francis, *et al.*, “A Scalable Content-Addressable Network,” in *Proc. of the ACM SIGCOMM*, San Diego, CA, 2001.
- [14] J. Cernik, “Freewar - MMORPG Browsergame,” [www.freewar.de](http://www.freewar.de).