

Dynamic Resource Provisioning in Massively Multiplayer Online Games



Journal:	<i>Transactions on Parallel and Distributed Systems</i>
Manuscript ID:	TPDS-2008-08-0321.R2
Manuscript Type:	Regular
Keywords:	C.2.4.a Client/server < C.2.4 Distributed Systems < C.2 Communication/Networking and Information Technology < C Computer Systems Organization, C.3.d Real-time and embedded systems < C.3 Special-Purpose and Application-Based Systems < C Computer Systems Organization, C.4.g Measurement, evaluation, modeling, simulation of multiple-processor systems < C.4 Performance of Systems < C Computer Systems Organization, I.6.8.f Gaming < I.6.8 Types of Simulation < I.6 Simulation, Modeling, and Visualization < I Computing Methodologies, J.7.g Real time < J.7 Computers in Other Systems < J Computer Applications



Manuscript Only

Dynamic Resource Provisioning in Massively Multiplayer Online Games

Vlad Nae, Alexandru Iosup, Radu Prodan

Abstract—Today’s Massively Multiplayer Online Games (MMOGs) can include millions of concurrent players spread across the world and interacting with each other within a single session. Faced with high resource demand variability and with misfit resource renting policies, the current industry practice is to over-provision for each game tens of self-owned data centres, making the market entry affordable only for big companies. Focusing on the reduction of entry and operational costs, we investigate a new dynamic resource provisioning method for MMOG operation using external data centres as low-cost resource providers. First, we identify in the various types of player interaction a source of short-term load variability, which complements the long-term load variability due to the size of the player population. Then, we introduce a combined MMOG processor, network, and memory load model that takes into account both the player interaction type and the population size. Our model is best used for estimating the MMOG resource demand dynamically, and thus for dynamic resource provisioning based on the game world entity distribution. We evaluate several classes of online predictors for MMOG entity distribution and we propose and tune a neural network-based predictor to deliver good accuracy consistently under real-time performance constraints. We assess using trace-based simulation the impact of the data centre policies on the quality of resource provisioning. We find that the dynamic resource provisioning can be much more efficient than its static alternative even when the external data centres are busy, and that data centres with policies unsuitable for MMOGs are penalised by our dynamic resource provisioning method. Last, we present experimental results showing the real-time parallelization and load balancing of a real game prototype using data centre resources provisioned using our method and show its advantage against a rudimentary client threshold approach.

I. INTRODUCTION

Massively Multiplayer Online Games (MMOGs) have emerged in the past decade as a new type of large-scale distributed application characterized by a huge real-time virtual world entertaining millions of players spread across the world, e.g., World of Warcraft counts over 6 million unique players daily. Invented and promoted by industry, MMOGs have started to attract the interest of scientific researchers as well, and challenges such as scalability, trust, and data consistency have been identified by the distributed systems [1] and database [2] communities. In this paper, we draw the attention to the resource provisioning problem for MMOGs as a new direction of research.

Today’s MMOG operate as client/server architectures, in which the game server simulates a world via computing and

database operations, receives and processes commands from the clients, and inter-operates with a billing and accounting system [3], [4]. Based on the actions submitted by the players, the game servers compute the global state of the game world represented by the position and interactions of the entities, and send appropriate real-time responses to the players containing the new relevant state information. Depending on the game, typical response times to ensure fluent play must be between 100 milliseconds in online *First Person Shooter (FPS)* action games [5] and 1-2 seconds for *Role-Playing Games (RPG)*. A good game experience is critical in keeping the players engaged, and has an immediate consequence on the income of the MMOG operators. Failing to deliver timely simulation updates leads to a degraded game experience and triggers player departure and account closing [3], [4].

Today, a single computer is limited at around 500 simultaneous and persistent network connections, and databases can manage the update of around 500 objects per second [2]. To support at the same time millions of *active concurrent players* and many more other game entities, MMOG operators need to install and operate a large dedicated multi-server infrastructure [4], with hundreds to thousands of computers hosting the distributed load of each game [3]. However, due to the dynamic character of MMOGs, both on the short and on the long term, the game providers have to over-provision their infrastructure, which leads to a low and inefficient resource utilisation and new providers finding it difficult to join the market. For example, the operating infrastructure of the *Massively Multiplayer Online Role Playing Game (MMORPG)* World of Warcraft has over 10,000 computers [6].

In this paper, we propose a dynamic resource provisioning solution that addresses the over-provisioning and the low-cost market joining problems. We first present in Section II a realistic model for an MMOG ecosystem that describes the resource providers, the application, and their integration. Extending previous work where resources were provided by a single data centre [7], [8], in our model the MMOG resources are provided by multiple geographically distributed data centres with different lease policies. In contrast to previous models of client/server Internet applications [7]–[10], which focus on non-interacting user requests, our model is designed around the notion of player interactions within the same application instance. In Section III we demonstrate through the analysis of traces from the popular MMOG RuneScape [11] that MMOGs are more dynamic than previously believed [12] and that the interaction between players is another major component of the MMOG workload, complementing the player population size.

Motivated by the dynamic MMOG workloads, we propose

V. Nae and R. Prodan are with the University of Innsbruck; A. Iosup is with the Delft University of Technology. Contact: Vlad@dps.uibk.ac.at, Radu@dps.uibk.ac.at, and A.Iosup@tudelft.nl.

1 a dynamic provisioning method in which the amount of re-
2 sources is first predicted, and then obtained dynamically from
3 data centres external to the MMOG operator. We devise in
4 Section IV an analytical MMOG load model for three types of
5 resources: processor, memory, and network. The model takes
6 into account both the player interaction type and the population
7 size. By feeding real-time measurements or online predictions
8 into the model, the game operators can dynamically estimate
9 the amount of resources needed for running their MMOG.
10 However, measuring the values of our proposed MMOG load
11 model's parameters in real time may not be cost-efficient or
12 even feasible due to the system scale and geographical spread.
13 Moreover, player interactivity types are not uniform across the
14 whole game, making real-time prediction difficult to adapt to
15 MMOGs. We address these two problems in Section V by
16 proposing a novel strategy for applying predictions to MMOG
17 player interaction. We partition the simulated world into small
18 areas where real-time prediction proves to be accurate, and
19 aggregate the estimates to get an overall prediction. We evalu-
20 ate several classes of online predictors from moving averages
21 to simple neural networks, and find that the simple neural
22 networks have better accuracy than the studied alternatives
23 for a variety of MMOG workloads.

24 In Section VI we evaluate the efficiency and the Quality of
25 Service (QoS) of our dynamic resource provisioning. We show
26 through simulations that dynamic resource provisioning which
27 favours data centres that lease fewer resources at a time and
28 for shorter periods of time considerably reduces the MMOG
29 operation costs with a reasonable loss of QoS, even when
30 the centres are busy. We also find that when game operators
31 use our dynamic resource provisioning the data centres are
32 incentivised to offer MMOG-friendly hosting policies when
33 competition exists on the data centre market.

34 We designed, implemented, and deployed our methods
35 in the edutain@grid environment [13] targeting a platform
36 for scalability, QoS provisioning, and user-friendly business
37 models for real-time online interactive applications, with string
38 focus on MMOG. In Section VII we present a prototype
39 experiment of applying our integrated methods on a real online
40 game in a distributed infrastructure.

41 Finally, in comparison with the related work surveyed in
42 Section VIII, our study is the first to investigate the resource
43 provisioning for a multi-MMOG, multi-data centre ecosystem.

44 II. MMOG ECOSYSTEM MODEL

45 In this section we introduce a model and platform for an
46 MMOG ecosystem in which a global network of data centres
47 host services that execute many MMOGs at the same time. Our
48 multi-MMOG and multi-data centre model extends previous
49 work, which focuses on either a single application (usually a
50 Web service) and/or a single data centre [9], [14], [15].

51 A. Application Model

52 MMOGs are large-scale simulations of persistent *game*
53 *worlds* comprising various objects or *entities* that we classify
54 in four categories: (1) *avatars* are in-game representation of
55 the players; (2) *bots* or *non-player characters* (NPC) are

56 mobile entities that have the ability to act independently; (3)
57 *movable objects* (such as boxes or guns) are passive entities
58 which can be manipulated but do not initiate interactions; and
59 (4) *immutable entities* or decor.

60 The mostly employed architectural model for MMOGs is
client/server [3], with *game operators* maintaining the servers
that simulate a distributed game world. The simulation consists
for every MMOG of an identical set of steps to be executed
each discrete game time unit (*game tick*), described in Sec-
tion IV-A. The clients dynamically connect to a joint *game*
session and interact with each other by sending play actions
such as movements, shootings, operations on game objects, or
chat. To ensure scalability and real-time response, an MMOG
session is distributed on multiple game servers, and each
player is mapped to an avatar on one of the servers, usually
to one in its closest proximity to minimise latencies. We call
the entities hosted by game server in a distributed session as
active entities (from the point of view of that server). The vast
majority of game servers follow a similar computational model
implementing an infinite loop, where in each loop iteration
(also called tick) there are certain steps to be performed: (i)
processing events coming from the connected clients and other
servers; (ii) processing the states of the active entities; and (iii)
broadcasting state update to the connected clients.

61 B. Game Session Parallelization

62 Today's MMOGs use three main session parallelization
63 techniques to serve at the same time hundreds of thousands of
64 players: zoning, replication, and instancing. These techniques
65 use the combined capacity of a group of servers working in
66 parallel; in this sense, we use in this work the terms *server* and
67 *server group* interchangeably. We describe in the following
68 each of these techniques, in turn. In Section VII we present
69 experiments with a real game that uses all three techniques.

70 Spatial scaling of a game session is achieved through
71 a conventional parallelization technique called *zoning* [16],
72 based on similar data locality concepts as in scientific parallel
73 processing. Zoning partitions the game world into geograph-
74 ical areas to be handled independently by separate machines
75 (see Figure 1). Zones are not necessarily of same shape and
76 size, but should have an even load distribution that satisfies
77 the QoS requirements. Today, zoning is successfully employed
78 in slower-paced (compared to fast-paced FPS action games)
79 MMORPG, where the transition between zones can only hap-
80 pen through certain portals (e.g. special doors, teleportation,
81 accompanied on the screen by a load clock or some standard
82 animation video) and requires an important amount of time.
83 Typically, zones are started manually by the game operators
84 based on the current load, player demand, or new game world
85 and scenario developments.

86 The second technique called *replication* [17] targets par-
87 allelization of game sessions with a large density of players
88 located and interacting within each other's area of interest (see
89 Figure 1). Such situations are typical to fast-paced FPS action
90 games in which players typically gather in certain hot-spot
91 action areas that overload the game servers that are no longer
92 capable of delivering state updates at the required rate. To

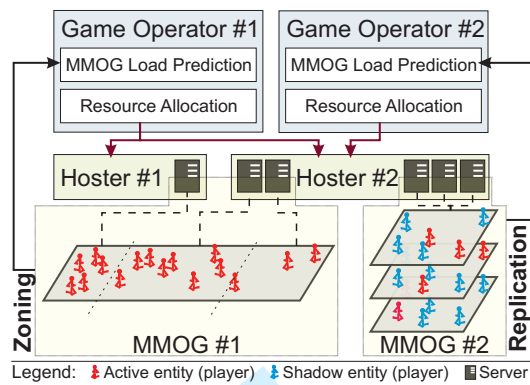


Fig. 1. The MMOG ecosystem architecture.

address this problem, replication defines a novel method of distributing the load by replicating the same game zone on several CPUs. Each replicated server computes the state for a subset of entities called *active entities*, while the remaining ones, called *shadow entities* (which are active in the other participating servers), are synchronised across servers.

The third technique called *instancing* is a simplification of replication which distributes the session load by starting multiple parallel instances of the highly populated zones. The instances are completely independent of each other, meaning that two avatars from different instances will not see each other, even if they are located at nearby coordinates.

C. Load Complexity Model

The load of a MMOG depends not only on the number of active concurrent players, but also on the number and type of their interactions (see Section III-D for empirical evidence). The interaction type and count span a wide range depending on the game design. The number of interactions between the entities may be very low (e.g. for puzzle games where a player interacts with the system after long periods of thinking), to low (e.g. for MMORPGs where small groups of people interact with a sparse environment), and to very high (e.g. for FPS action games where many players test their reaction time in a confined area).

Assuming the number of entities is n , the interaction complexity may range from $O(n)$ for games in which players are mostly solitary or the game does not need to make many state changes or compute complex interactions, to $O(n^2)$ for games in which many players acting individually are interacting, and to $O(n^3)$ for games in which groups of many players are interacting. To reduce the computational load, most MMOGs simulate and send updates only for world regions representing the *area of interest* of each avatar [18]. When using such techniques, the interaction complexity may become $O(n \cdot \log n)$ from $O(n^2)$, and $O(n^2 \cdot \log n)$ from $O(n^3)$.

The lack of responsiveness of an MMOG may also be caused by high latency, independently from the game operator's server and bandwidth capacity. However, depending on the game design, MMOGs have different *latency tolerance*. The latency tolerance has been investigated by previous research [5], [19]; for example, for FPS action games laten-

cies above 100ms severely disrupt the gameplay, while for MMORPGs any latency below 1.5s is tolerable.

D. Hosting Model

We consider the hosting platform as consisting of data centres scattered around the world, where each centre pools together resources that may serve several game operators simultaneously (see Figure 1). For simplicity, we assume that each data centre consists of a single computing resource, which can be a shared or distributed memory parallel machine owned by a *hoster*. The game operators submit requests to the data centre by specifying the type, number, and duration for which the resources are desired. We consider four resource types: *CPU*, *memory*, input from the external network (*ExtNet[in]*), and output to the external network of a data centre (*ExtNet[out]*); here, the external network connects the data centre with the Internet.

Depending on the data centre's service model (either best-effort or advance reservation-based), resource requests are queued or immediately fitted in the schedule, respectively. Once the available resources are matched against the requests, these resources are *allocated* to the game operators. From the game operator's point of view, we say that the resources have been *provisioned*. We use from here on the terms resource allocation and resource provisioning interchangeably. The allocated resources are reserved for executing the MMOG servers for the entire duration of the game operator's request (task preemption or migration are not supported).

Our hosting model considers the size and the duration of the minimal resource allocation, which may be not only for a resource as a whole (e.g. a server in Web data centres [14] or a multi-processor node in a Grid system [20]), but also for a fraction of that resource (e.g. a virtual machine running on a physical node [10], or a channel of an optical network). Similarly, the minimal duration for which a resource may be allocated may be between a few seconds (servicing one user request by a Web service) to several months (a typical value for Web server hosting). We define the *resource bulk* as the minimum number of resources that can be allocated for one request, expressed as the multiple of a minimal resource size. Similarly, we define the *time bulk* as the minimum duration for which a resource allocation can be performed expressed as multiple of a minimal time period. A data centre may choose to allocate resources for MMOGs only in bulks under a space-time *hosting policy* imposed by the hosters.

E. Complete Ecosystem Model

The MMOG ecosystem comprises multiple hosters and multiple game operators, where each hoster may set a different space-time policy for its resources. The game operators handle simultaneously MMOGs of different genres and designs with different interactivity types and counts, and with different latency tolerance. The relative success of each game is characterised by the number of registered players. Figure 2 displays the number of MMORPG players over time for the USA and European markets based on the survey of Woodcock [21] for dates until June 2006, and on our own research afterwards. The




Fig. 2. Number of MMORPG players over time.

chart shows that there are currently six games with more than 500 thousand players each. The total number of MMORPG players is well approximated by the exponential trend $\alpha \cdot e^{\beta x}$, where $\alpha = 7 \cdot 10^{-9}$ and $\beta = 0.028$ give a Pearson's coefficient of determination $R^2 = 0.974$. Assuming the same rate of growth, we can estimate over 60 million players by 2011 in the US and EU markets alone. The large number of MMOG players, for the whole ecosystem and for each game in part, is an important motivation for our work.

The game operators make requests based on the load of their game servers and the data centres respond with offers based on their local time-space renting policy. The resource allocation is realised by a request-offer matchmaking mechanism according to three criteria that favour the game operator. First, the number and the type of resources requested must match with the offer, and when they do not match an offer that includes at least the requested amounts is selected. Second, depending on the game latency tolerance, the resources closest to the request are preferred. Third, to deal with data centre hosting policies, the finer grained resources with the shorter period of reservation time are preferred.

III. MMOG WORKLOAD ANALYSIS

Previous work on MMOG workload characterisation focused on highly interactive games with few users playing together [22], traced small to mid-sized MMOGs [12], [23], or collected data from only one server from a large distributed MMOG [24]. In contrast, our analysis focuses on all server groups of one of the largest commercial MMOGs, for which we analyse the workload at both server and network level based on server location and user interactions.

A. RuneScape Traces

RuneScape [11], ranked second by number of players in the US and European markets (see Figure 2), is not a traditional MMORPG, but combines elements of RPG and FPS (and other genres) in specific parts of the game world called *minigames*, where player interaction follows different rules. Thus, various levels of player interactivity coexist and the game load cannot

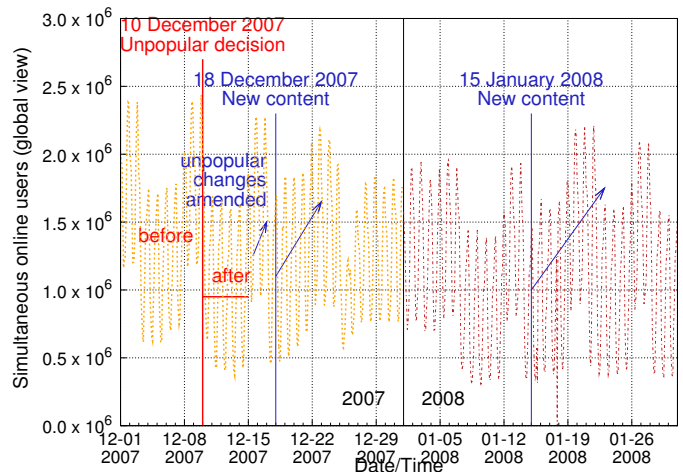


Fig. 3. The global RuneScape active concurrent players.

be trivially computed, for example using the linear models employed in [15]). We started monitoring and collecting traces from the official RuneScape Web page [11] in August 2007. The traces are sampled every two minutes and contain the number of players over time for each server group used by the game operators. In this work, we analyse the traces for a period of over six months until March 2008.

B. Global Number of Players

The number of RuneScape players has surged over the past two years, starting with the introduction of the minigames. From 180,000 active players (i.e. players that played at least once in the last month) at the beginning of 2005, the game is estimated to have now over 5,000,000 active players from over 8,000,000 open accounts in 2007 [25], [26]. Our study of the official list of top RuneScape players counted over 3,000,000 active players in September 2007. Since a player needs to be efficiently active for about a month to become a top player, we conclude that RuneScape converts into *dedicated* players between 30% and 60% of the starting players.

Our study shows that the maximum global number of active concurrent players for RuneScape is around 250,000. However, this number is strongly driven by the mood of the player base. Figure 3 depicts the number of active concurrent users for RuneScape over a period of two months. A highly unpopular decision issued on December 10, 2007 resulted in massive account cancellations which dropped the number of active concurrent players by over 30,000 units (a quarter of its value) in less than one day. Under intense pressure, the game operators agreed to amend the changes and the number of active concurrent players raised again to 95% of the previous value. On December 18, 2007 and January 15, 2008 the game operators released new content, which caused an over 50% surge of the number of active concurrent players after about one week in each case.

This high MMOG variability in the number of active and concurrent players determines a very dynamic resource requirement, which means that static resource provisioning would lead to significant over-provisioning that is a strong motivation for our work.

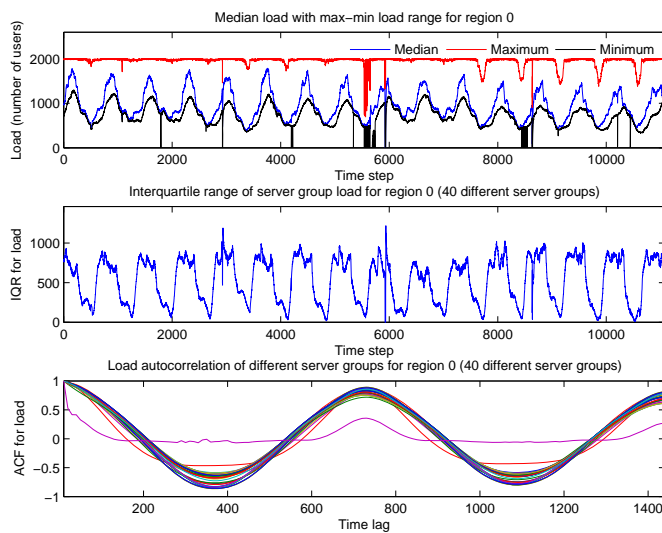


Fig. 4. RuneScape workload for region zero (Europe).

C. Patterns in the Regional Number of Players

Our RuneScape traces characterise a global population spread across five geographical regions including Europe and the US East and West Coasts. We present in this section the analysis of the number of players that use the server groups of region zero representing Europe. The input for this analysis is a subset of two weeks (middle to end of August 2007) from the original RuneScape traces aggregating over 11,000 data samples, where each sample contains one load value for each server in the region. The analysis of the other four regions, as well as of the rest of the traces, yielded very similar results. The top subplot of Figure 4 shows the minimum, the median, and the maximum load measured in number of online users in any server group in the region at each time step. The median load shows a diurnal pattern and a strong load variation during the peak hours when the median is about 50% higher than the minimum. There are also some heavy fluctuations caused by few sporadic and short-lived server group outages which fall outside the scope of this analysis.

To characterise the load variability between different server groups, the middle subplot depicts the load interquartile range (IQR) over time, defined as the difference between the 75th and 25th percentiles of a data set. Similarly to the median load, the load variability has a diurnal cycle. Unlike the load of e-business and Web servers [27], the median load shows no weekend effects, e.g. the load does not differ significantly between weekend and normal work days.

To establish the duration of the cycles observed in the top and middle subplots, the bottom subplot displays the autocorrelation function for each of the European server groups. We see a significant peak at around 720 (720 samples * 2 minutes per sample = 24 hours) and a strong negative peak at around 360 (12 hours), which shows again the diurnal pattern of the load of most servers. However, the same subplot shows that the load of 2 – 5% of the servers is always 95% except for outages, which does not follow a diurnal pattern.

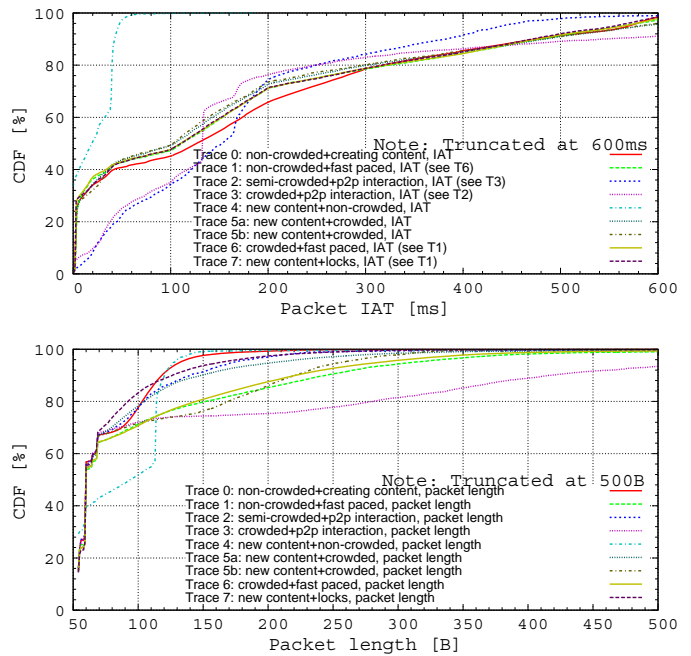


Fig. 5. The CDF of packet length (*top*) and the CDF of packet inter-arrival time (*bottom*) for different levels of player interaction.

D. Player Interaction Influence on Server Load

A fundamental premise of our work is that MMOG workloads depend on the player interaction. We demonstrate in this section that this is indeed the case. Using the `tcpdump` tool, we collected eight RuneScape traces labelled T_1, \dots, T_8 , and analysed the packet size distribution and the inter-arrival time (IAT) between consecutive packets. Each trace is collected from a session of at least five minutes and at most one hour. To ensure the independence of the measurements, the traces were collected at different dates over a six month period, while the traces T_{5a} and T_{5b} were collected from the same environment at consecutive periods of time.

Figure 5 illustrates that the network load depends on the number and type of player interaction. For traces T_1 and T_6 that involve a fast-paced game, the level of interaction (i.e. crowded or non-crowded) does not increase the server load as the players are very sensitive to delays. Thus, for fast-paced games the server needs to send packets as often as possible and including as much information as possible. For traces T_2 (market) and T_7 (new content) that involve direct player-to-player interaction, the packet sizes are similar but their IAT is very different. The statistical moments of the IAT of T_7 are lower than those of T_2 , as for T_2 the need for updates is conditioned by players starting and agreeing to trades, with more thinking time than for the player actions in T_7 . For the trace T_4 that involves group interaction, the packets need to arrive more often (lower IAT than for other traces) and to include information about more objects (higher packet size).

IV. MMOG LOAD MODELLING

We propose in this section an analytical model for the load of MMOGs. Our model covers the three main types of resources used by MMOGs: CPU, memory, and network.

Let us consider N clients connected to a distributed game session aggregating a total of H (parallel, cluster) machines from different hosters. Let us further consider that the game world is populated by BE moving bots or NPCs (entity category (ii) in Section II-A). On each the machine, there are only AE active entities and C clients connected.

A. CPU Load Model

For modelling the load of one machine in a distributed session, we distinguish three basic time consuming activities within one game tick: (1) the computation of the interaction between pairs of entities (consumed time for each computation: t_i); (2) the reception of event messages from each client (t_m); and (iii) the update of entity states received from/sent to another machine (t_u). In order to keep the complexity of the model acceptable, we assume that NPC entities do not interact, which is true for many MMOGs. We model the CPU time t_M spent for sending and receiving messages from a server to each client (active avatars) as:

$$t_M = C \cdot t_m.$$

The CPU time t_U spent by the server for processing state updates from the other machines is:

$$t_U = (N - C) \cdot t_u + (BE - AE) \cdot t_u + AE \cdot t_u,$$

and the time t_I spent by the server for computing the interactions between the active entities is:

$$t_I = I \cdot t_i,$$

where I is the total number of interactions involving the active entities. Obviously, the computation of interactions that do not involve active entities is allotted to other machines.

For quantifying the interactions between entities, we use a generic function $f(e_1, e_2)$ which has to be instantiated for each interaction type introduced in Section II-C:

$$f(e_1, e_2) = \begin{cases} e_1 + e_2, & \text{for } O(n) \text{ interaction;} \\ e_1 \cdot \log(e_2), & \text{for } O(n \cdot \log(n)) \text{ interaction;} \\ e_1 \cdot e_2, & \text{for } O(n^2) \text{ interaction;} \\ e_1^2 \cdot \log(e_2), & \text{for } O(n^2 \cdot \log(n)) \text{ interaction;} \\ e_1^2 \cdot e_2, & \text{for } O(n^3) \text{ interaction.} \end{cases}$$

where e_1 and e_2 are two classes of interacting entities.

Let IC denote the number of avatars interacting with any other entities (either avatars or NPCs). Furthermore, we define p_{ci} as the average number of interactions involving active avatars entities expressed as a percentage of IC . Analogously, we define p_{ei} as the average number of interactions involving active NPCs expressed as a percentage of BE . The total number of interactions is composed of the number of interactions between active avatars and the number of interactions between active avatars and NPC entities:

$$I = p_{ci} \cdot f(IC, IC) + p_{ei} \cdot f(IC, BE).$$

Consequently, the CPU time t_I for processing the interactions involving all active entities can be calculated as follows:

$$t_I = (p_{ci} \cdot f(IC, IC) + p_{ei} \cdot f(IC, BE)) \cdot t_i.$$

Approximating the time for sending/receiving an event as equal to the time for updating the state of one entity ($t_m = t_u$), the total CPU time consumed in one tick becomes:

$$t_C = (N + BE) \cdot t_u + (p_{ci} \cdot f(IC, IC) + p_{ei} \cdot f(IC, BE)) \cdot t_i.$$

Furthermore, quantifying t_i with regard to t_u as $t_i = p_{ui} \cdot t_u$, where p_{ui} is the ratio between the time necessary for one entity update and the time for computing one interaction (in percentage), the CPU time consumed in one tick becomes:

$$t_C = (N + BE + p_{ui} \cdot p_{ci} \cdot f(IC, IC) + p_{ui} \cdot p_{ei} \cdot f(IC, BE)) \cdot t_i.$$

Finally, considering t_{SAT} as the tick saturation threshold, we can define the CPU load function:

$$L_{CPU} = \frac{t_C}{t_{SAT}} = \frac{N + BE + p_{ui} \cdot p_{ci} \cdot f(IC, IC) + p_{ui} \cdot p_{ei} \cdot f(IC, BE)}{v},$$

where v is the CPU speed expressed as an integer representing the number of t_u -long tasks the CPU is able to perform in a t_{SAT} -long time interval.

B. Memory Load Model

The memory model is less complex than the processor load model, since all machines keep the entity state records for all entities participating in the game session. First, we take into account the game-dependent constants such as the amount of memory m_{game} needed to run the actual game engine with no game world loaded and no clients connected. Next, we define m_{world} as the amount of memory used for the game world being played. As for entity-related memory constants, let m_{cs} denote the amount of memory needed to store the state of one avatar, and m_{es} the amount of memory needed to store the state of an NPC entity. The interaction between entities does not have a significant impact on the memory load and we ignore it. Aggregating all this data, the memory consumption M on a machine taking part in a distributed session is:

$$M = N \cdot m_{cs} + BE \cdot m_{es} + m_{game} + m_{world}.$$

As a consequence, the final memory load function is:

$$L_{MEM} = \frac{M}{M_{machine}} = \frac{N \cdot m_{cs} + BE \cdot m_{es} + m_{game} + m_{world}}{M_{machine}},$$

where $M_{machine}$ represents the amount of memory available on a machine.

C. Network Load Model

In terms of network consumption, we define the outgoing network bandwidth usage for a machine running a server of a distributed game session as follows:

$$D_{out} = C \cdot d_{cout} + (H - 1) \cdot (C + AE) \cdot d_{updt},$$

where d_{cout} represents the amount of data sent to a client and d_{updt} the amount of data exchanged between machines for updating a single entity state.

The incoming network bandwidth usage for a machine in a distributed session is defined as:

$$D_{in} = C \cdot d_{cin} + (N - C + BE - AE) \cdot d_{updt},$$

where d_{cin} is the amount of data received from a client.

Finally, we define an overall network load function as the maximum between the incoming and outgoing loads, since the network is congested once one of the two maxima is reached:

$$L_{NET} = \max \left(\frac{C \cdot d_{cout} + (H - 1) \cdot (C + AE) \cdot d_{updt}}{B_{out}}, \frac{C \cdot d_{cin} + (N - C + BE - AE) \cdot d_{updt}}{B_{in}} \right),$$

where B_{in} and B_{out} denote the input, respectively output network bandwidths.

D. Complete Load Model

We integrate the presented models into a complete resource load model for MMOGs, where the load of the entire system is imposed by the maximum load of the individual resources:

$$L = \max(L_{CPU}, L_{MEM}, L_{NET}).$$

To stress the generality of our approach, MMOG classes can be defined using the set of constants involved in all the models previously described:

$$MMOG_{class} = \{(BE), (m_{cs}, m_{es}, m_{game}, m_{world}), (d_{cout}, d_{cin}, d_{updt})\}.$$

Note that BE is not MMOG-dependent, but rather game world and play style (e.g. single versus team play)-dependent. Nevertheless, we included it among the constants because we can consider games running different game worlds and play styles as belonging to different MMOG classes.

V. MMOG LOAD PREDICTION

In this section we introduce a MMOG load prediction solution comprising a load prediction strategy and a tuned neural network-based predictor that offers good real-time prediction accuracy for a variety of MMOG workloads. We also present the evaluation of several classes of predictors leading to the selection of the neural network-based predictor.

A. MMOG Load Prediction Strategy

We have shown in Section III that the load of MMOGs is more dynamic than previously believed, mostly because of the player interactions. Therefore, fast and accurate load prediction algorithms are required to dynamically allocate resources for MMOGs which, besides the entity count, also consider the entity interaction. However, due to the different play styles of the players logged in at different moments of time, the number of interactions can vary greatly. Thus, predicting the number of interactions for the whole game in real-time leads to low accuracy. To address this problem, our prediction strategy for MMOGs is to first partition the game world into subareas, where the size of a subarea is small enough to be characterised by the entity count and interaction type (see Section II-C). Then, the entity count for each subarea is used as input for the model described in Section IV. We define the overall entity distribution as the set

of entity counts for each subarea. The overall MMOG load prediction is obtained by using the entity distribution as input to many instances of the load model.

The main problem to be solved for this strategy is finding a predictor with high accuracy for a variety of MMOG workloads and good performance. As the overall entity distribution is required for one overall MMOG load prediction, the predictor needs to be able to deliver tens of thousands of predictions per second. In the remainder of this section we address the problem of finding such a predictor.

B. Predictor Families

Two options are available for quantitative predictions in MMOGs: explanatory models and time series prediction. While explanatory models can deliver good accuracy with little computation, they are difficult to obtain for complex applications such as MMOGs, and are tightly-coupled to the application instance and sometimes to platform for which they have been constructed. With MMOGs relying on frequent and large updates to maintain interest among the players (in Figure 3 we see a rate of one update per month), the explanatory models quickly become unmaintainable. Thus, our work is based on prediction algorithms that use historical values to discover patterns in the historical data series, and extrapolate these patterns into the future.

Many such prediction algorithms have already been proposed [28]. Simple prediction algorithms like exponential smoothing and variants thereof are computationally inexpensive and can be applied in parallel on several data sets, but their predictive power is limited. More elaborated prediction algorithms like autoregressive (AR), integrated (I), moving average (MA) models, and combinations thereof like ARMA or ARIMA try to find the best prediction model for the given data set. Although their predictive power is higher, such methods are also more time consuming and resource intensive, thus being ill suited for highly dynamic MMOGs.

C. Neural Network-based Prediction

As an alternative to other prediction algorithms, we investigated the use of neural networks [29] that provide a robust approach to approximating real- or discrete-valued target functions. Low complexity neural networks are capable of approximating complex and noisy functions, provided that good input preprocessing is performed. Our predictor uses one separate neural network for each game subarea which receives as input the entity count at equidistant past time intervals and delivers as output the entity count at the next time step.

A downside of neural networks is that they require a long offline training phase consisting of several sub-phases. First, the *data set collection* phase is a long process in which the game is observed by gathering entity count samples for all subareas at equidistant time steps. The second *training* phase uses most of the collected samples as training sets, and the remaining ones as test sets. The training phase runs for a number of eras, until a convergence criterion is fulfilled. A training era consists of three steps: (1) presenting all the training sets in sequence to the network; (2) adjusting the network's weights to

1 better fit the expected output (the real entity count for the next
2 time step); and (3) testing the network's prediction capability
3 with the different test sets. Separating the training from the
4 test sets is crucial for avoiding memorisation and ensures that
5 the network has enough generalisation potential for delivering
6 good results on new data sets.
7
8

9 D. Traces for Testing MMOG Predictors

10 To experiment and validate the neural network prediction,
11 we developed a distributed game simulator which realistically
12 emulates the behaviour of game players. The motivation for
13 using an emulator is twofold: (1) we do not have available
14 the exact coordinates of entities in the RunScape game world
15 (and we do not have access to the code or the documentation
16 of the RuneScape server either); and (2) through this emulator
17 we are able to give further evidence that the player interaction
18 determines the server load (see also Section III-D).

19 1) *Emulator*: We have designed a MMOG emulator which
20 simulated a virtual environment populated by avatars. The
21 environment consists of real maps from the popular FPS game
22 Counter Strike [30]. The avatars are driven by several Artificial
23 Intelligence profiles that match the four behavioural patterns
24 most encountered in MMOGs [3]: the *achiever* determines
25 the avatar to seek and interact with opponents; the *socialiser*
26 causes the avatar to act in a group together with its teammates;
27 the *explorer* leads the avatar for discovering uncharted zones
28 of the game world (not guaranteeing any interaction); and the
29 *killer* simulates a well-known tactic in FPS games to hide
30 and wait for the opponent (thus gaining a tactical advantage
31 through the element of surprise). To also account for the
32 mixed behaviour encountered in deployed MMOGs [3], each
33 entity has its own preferred profile, but can change the profiles
34 dynamically during the emulation.
35

36 2) *Generated Traces*: We used this emulator to generate
37 eight different data traces for a duration of one day each with
38 a sampling rate of two minutes, modelling four parameters
39 (see Table I): peak hours, peak load, overall dynamics, and
40 instantaneous dynamics. The peak hours correspond to the
41 periods with high player count in online gaming such as late
42 afternoons (see Section III). The peak load represents the
43 highest load observed in a MMOG which is a good measure
44 for its relative popularity. The overall dynamic represents the
45 variability of the entity interaction over a period of one day,
46 while the instantaneous dynamic indicates the same variability
47 over a period of two minutes. The eight data sets exhibit
48 three major types of signals: *type I* with high instantaneous
49 dynamics and medium overall dynamics (sets 2, 3, and 4);
50 *type II* with low instantaneous dynamics (sets 6, 7 and 8); and
51 *type III* with medium instantaneous dynamics (sets 1 and 5).
52

53 E. Prediction Results

54 In this section, we compare the neural network prediction
55 against other well-known fast prediction methods such as last
56 value, average, moving average, sliding window median, and
57 three levels of exponential smoothing. We used a well-tuned
58 Multilayer Perceptron with a [6, 3, 1] structure representing
59 the number of neurons on each layer, which delivered the
60

TABLE I
CONFIGURATION PARAMETERS OF THE TRACE DATA SIMULATION SETS.

Data set	Player behaviour [%]				Peak hours modelling	Peak load	Overall dynamics (17 h.)	Instantaneous dynamics (2 min.)
	Aggressive	Scout	Team player	Camper				
Set 1	80%	10%	0%	10%	No	+++	+++++	+++++
Set 2	60%	10%	0%	20%	No	++++	+++++	+++++
Set 3	70%	20%	0%	10%	No	+++++	+++++	+++++
Set 4	70%	30%	0%	0%	No	++++	+++++	+++++
Set 5	30%	40%	30%	0%	Yes	+++++	+++++	+++++
Set 6	10%	80%	10%	0%	Yes	+++++	+++++	+++++
Set 7	20%	40%	40%	0%	Yes	++++	+++++	+++++
Set 8	20%	80%	0%	0%	Yes	++++	+++++	+++++




Fig. 6. Accuracy of seven prediction algorithms applied to MMOG data.

most accurate results in a series of offline network tuning experiments [31]. The goal of our experiment is to minimise the *prediction error*:

$$PE = \frac{\sum_{i=1}^N |n_i^{(real)} - n_i^{(pred)}|}{\sum_{i=1}^N n_i^{(real)}}$$

where N is the total number of samples in the trace data set, and $n_i^{(real)}$ and $n_i^{(pred)}$ are the real, respectively predicted entity counts at time step i .

Each prediction method receives as input every trace data set described in Table I, and outputs for each input set a sample prediction. The results in Figure 6 show that, apart from having lower prediction errors, the quality of our neural network-based predictor is its ability to adapt to various types of input signals. In contrast, other algorithms exhibit poor performance for some types of signals (e.g. the average method is the second most accurate for Type I signals, but performs poorly for all Type II and for some Type III signals). Notably, our neural network predictor was significantly better than the others for the sets with high instantaneous dynamics (Types I and III signals).

Figure 7 depicts the duration of one prediction on a Intel Core Duo E6700 (2.66GHz) processor. Although the neural network predictor is the slowest with an average prediction duration of approximately seven microseconds, it is nevertheless fast enough and suitable to MMOGs. To justify this statement, let us consider a RuneScape setup consisting of a hoster managing 200 game servers, each one with an upper client limit of 2,000. Moreover, let us consider the worse case scenario in which our predictor would use a number of subareas equal to the number of players, and that all game worlds are handling the maximum number of players (i.e. 250,000 subareas in total, see Section III-B). The approximate total prediction time using our proposed predictor would be 14 milliseconds per server, and 2.8 seconds per hoster consid-




Fig. 7. Statistical properties of the duration of one prediction for four prediction algorithms applied to MMOG data.

TABLE II
THE EXPERIMENTAL SPACE COVERAGE.

Section	Resource provisioning	Prediction algorithm	Update models	Hosting policies	Latency tolerance	Number of MMOGs
Section VI-B	static, dynamic	all	$O(n^2)$	optimal	none	one
Section VI-C	dynamic	Neural	all	optimal	none	one
Section VI-D	dynamic	Neural	$O(n^2)$	all	none	one
Section VI-E	dynamic	Neural	$O(n^2)$	optimal	all	one
Section VI-F	dynamic	Neural	$O(n^2)$	optimal	none	several

ering a single threaded execution. From a realistic prediction interval of two minutes, the prediction time would amount approximately 2%, leaving the remaining 98% for resource allocation and load balancing actions.

VI. MMOG RESOURCE PROVISIONING

In this section we present an evaluation of our MMOG resource provisioning model described in Section II-E. As outlined in Table II, we cover an experimental space with six axes: the provisioning mechanisms, the prediction algorithms, the player interaction, the hosting policies, the latency tolerance, and the multi-MMOG workloads.

A. Experimental Setup

1) *Metrics*: We evaluate each experiment by using three metrics: resource over-allocation, resource under-allocation, and number of significant under-allocation events.

The *resource over-allocation* characterises the percentage of a resource (i.e. CPU, memory, network) that has been allocated from the amount used for the seamless execution of a MMOG session. We define the resource over-allocation at time instance t as the cumulated over-allocation for all M resources participating in the game session, where $\alpha_m(t)$ is the total allocated resource and $\lambda_m(t)$ is the resource usage (the generated load):

$$\Omega(t) = \frac{\sum_{m=1}^M \alpha_m(t)}{\sum_{m=1}^M \lambda_m(t)} \cdot 100[\%].$$

The *resource under-allocation* characterises the percentage of resources that have not been allocated from the amount necessary for the seamless execution of the MMOG, but considering that missing resources on one machine can be hidden by over-allocating the resource on other machines. We define resource under-allocation at time instance t as the average difference between the allocated and used resource. The min function limits the maximum under-allocation value to at most zero and thus, an over-allocation at a certain time

does not reduce impact of an under-allocation at a different time instance (that is, $\Omega(t)$ and $\Upsilon(t)$ are not correlated):

$$\Upsilon(t) = \frac{\sum_{m=1}^M \min(\alpha_m(t) - \lambda_m(t), 0)}{M} \cdot 100[\%].$$

The *number of significant under-allocation events* indicates the number of times the under-allocation causes game play disruption over a long period of time. In this work, we consider an under-allocation as being disruptive (and frustrate players that may quit the game) if its absolute value is over 1% for a period of time of at least two minutes.

2) Environment: We

performed experiments in a simulated RuneScape-like environment. The input workload consisted of the first two weeks from the RuneScape trace data analysed in Section III which, with the metrics being evaluated every two minutes, gives over 10,000 metric samples for each simulation, ensuring statistical soundness. The data centres are located on four continents and seven countries, as depicted in Table III. Each machine in the specified setup is capable of handling at least one game server at full load (e.g., 2,000 simultaneous clients for RuneScape). The data centres were configured to use different hosting policies, where each policy describes one time and one resource bulk for each type of resource (see Section II-D). The measurement unit for the policy resources is a generic “unit” which represents the requirement for the respective resource in a fully loaded RuneScape game server (e.g., one external outward network unit is equivalent to a real bandwidth value of 3 MB/s – see also Figure 5).

In our simulation, the game operators perform a prediction of the game load (that is, of the number of players and of the interactions per zone) every two minutes and request an appropriate amount of resources to the data centres. The protocol how resources are matched based on their number, type, location, and time in this ecosystem was presented in Section II-E. We assume zero overhead in resource allocation, provisioning, and setup from data centres to game operators.

B. Prediction Impact

In this experiment, we evaluate the impact of the prediction method on the provisioning process. We assign the *HP-1* and

Continent	Location		Data Machines	
	Country	Centres (total)		
Europe	Finland	2	8	
	Sweden	2	8	
	U.K.	2	20	
	Netherlands	2	15	
North America	U.S. (West)	2	35	
	Canada (West)	1	15	
	U.S. (Central)	1	15	
	U.S. (East)	2	32	
Australia	Canada (East)	1	10	
	Australia	2	8	

TABLE III
DATA CENTRE PHYSICAL CHARACTERISTICS.

Hosting policy	CPU	Memory	External net.		Time [minutes]
			in	out	
HP-1	0.25	n/a	6	0.33	360
HP-2	0.25	n/a	4	0.5	360
HP-3	0.22	2	n/a	n/a	180
HP-4	0.28	2	n/a	n/a	180
HP-5	0.37	2	n/a	n/a	180
HP-6	0.56	2	n/a	n/a	180
HP-7	1.11	2	n/a	n/a	180
HP-8	0.37	2	n/a	n/a	360
HP-9	0.37	2	n/a	n/a	720
HP-10	0.37	2	n/a	n/a	1440
HP-11	0.37	2	n/a	n/a	2880

TABLE IV
HOSTING POLICIES.

TABLE V
DYNAMIC RESOURCE ALLOCATION RESULTS.

Predictor type	Avg. Over-allocation [%]		Avg. Under-allocation [%]			Υ > 1% events	
	CPU	ExtNet	CPU	ExtNet	ExtNet		
Neural network	25.90	995.27	66.04	-0.09	0	317	
Average	32.41	1023.43	69.29	-12.84	0	-2.46	8123
Last value	25.11	989.10	65.36	-0.16	0	0	608
Moving average	24.92	992.06	65.69	-0.33	0	-0.03	1142
Sliding window	24.97	992.73	65.76	-0.41	0	-0.03	1423
Exponential smoothing	24.76	977.85	64.11	-0.42	0	-0.03	1429

HP-2 hosting policies described in Table IV to the data centres presented in Table III in a round-robin fashion. When two data centres have the same location (column “Country” in Table III), their hosting policies are set differently to HP-1, respectively HP-2, and their machine size to half the number of resources at that location (column “Machines” in Table III).




Fig. 8. Cumulative under-allocation events for five predictors.

Table V. For over-allocation, we observe two performance classes: the poor performance class with one member (the average predictor), and the normal performance class with the other five predictors as members. The reason for the big over-allocations for the external network bandwidth is that the two utilised policies were not well fitted to the input workload (i.e., the policies in Table IV included too much external network bandwidth relative to the CPU). We will show in Section VI-E that the resources from data centres with such policies are not used when other suitable alternatives exist.




Fig. 9. Static versus dynamic resource over-allocation.

We further rank the five predictors from the normal performance class using the under-allocation metric. First, only the neural network and the last value predictors lead to no under-allocation of the external network bandwidth, in addition to leading to the lowest CPU under-allocation. Figure 8 depicts the cumulative number of significant under-allocation events over time for the five predictors with a normal over-allocation value, which shows that the neural network predictor exhibits the lowest number and the most stable evolution. We conclude that our novel neural network predictor enables the best resource provisioning, followed by the last value predictor which confirms our results from Section V-E.

Figure 9 comparatively displays the resource over-allocation resulted from the use of static, respectively dynamic provision-

TABLE VI
STATIC VERSUS DYNAMIC PROVISIONING FOR VARIOUS INTERACTIONS.

Interaction type	Static provisioning		Dynamic provisioning		Υ > 1% events
	Over-allocation [%]	Over-allocation [%]	Under-allocation [%]	Under-allocation [%]	
$O(n)$	55.71	8.47	0	1	
$O(n \cdot \log(n))$	71.86	16.07	-0.024	22	
$O(n^2)$	146.03	27.77	-0.066	103	
$O(n^2 \cdot \log(n))$	180.60	36.26	-0.096	191	
$O(n^3)$	242.04	54.62	-0.130	304	

ing strategy (using the neural network predictor) for the same workload. As expected, the dynamic provisioning of resources achieves better results, its average over-allocation being around 25%, compared to 250% for the static allocation. The over-allocation for dynamic provisioning is not the outcome of unreliable predictions and can be lowered if the data centres policies are more favourable. In this experiment, the deallocation of resources was only allowed at least six hours after the start of the allocation (column “Time” in Table IV).

C. Player Interaction Impact

In this section we study the impact of player interaction on the dynamic provisioning using the neural network predictor.

Figure 10 shows the resource over- and under-allocation over time for the $O(n)$, $O(n \cdot \log(n))$, $O(n^2)$, $O(n^2 \cdot \log(n))$, and $O(n^3)$ MMOG update models described in Section II-A, for dynamic resource provisioning. We observe that the higher the complexity of the update model is, the greater the fluctuations in resource over-allocation are. At the same time, the significant under-allocation events become more frequent as the complexity of the update model increases. This is further confirmed by Figure 11 which depicts the cumulative number of significant under-allocation events over time, at the end of the two simulation weeks this number being significantly higher for $O(n^3)$ than for $O(n)$.

Table VI

compares the average static and dynamic resource provisioning mechanisms for various interaction types. The static provisioning has five to seven times




Fig. 11. Cumulative under-allocation events.

higher resource over-allocation than the dynamic one, but no under-allocation events. However, the number of significant under-allocation events over the whole simulated period remains below 3% for dynamic provisioning (at most 30 samples from over 10,000 samples used in the simulation). When even this low occurrence cannot be tolerated, a mechanism that increases the over-allocation shall be used.

D. Data Centre Hosting Policy Impact

We further evaluate the influence of the hosting policies on the dynamic resource provisioning, where each hosting policy expresses the sizes of the resource and of the time




Fig. 10. Over- and under-allocation for three update models.

bulks (see also Section II). When more resource types are involved in the matchmaking process, there is a separate bulk size for each resource type. Because the resource and the time bulks have a combined influence on the provisioning of resources, we conduct three separate experiments that use a different resource hosting policy setup each. We first study in Section VI-D1 the individual impact of the resource and time bulk parameters on resource provisioning by varying one of them and keeping the other constant. Then, we study the combined effect of the resource and time bulks by using a heterogeneous hosting policy setup in Section VI-D2.

1) *The independent impact of the resource and time bulks:* We first estimate the impact of the CPU resource bulk variation on the resource allocation performance by using five hosting policies from *HP-3* to *HP-7* in Table IV. The resource bulks for other resource types and the time bulk are kept constant. The values selected for the CPU resource bulk, i.e. from 0.22 to 1.11, are not evenly distributed in the selected interval, which reflects the real-life policies of data centres that try to maximise their own resource usage and do not willingly adapt to a specific MMOG's resource requirements. We will show in Section VI-E that in our MMOG ecosystem, the game operators are not always forced to accept such conditions and can penalise the data centres with unsuitable hosting policies by not using their resources.

Experiment batch *A* in Table VII illustrates the influence of the CPU resource bulk on the dynamic resource provisioning. There is a visible tendency of higher over-allocation values for bigger CPU resource bulks. Conversely, we can observe an increase in significant under-allocation events as the CPUs are offered with finer grained quantities. In conclusion, the finer grained policies have the potential to increase the resource provisioning efficiency, but also the risk of increasing the number of significant under-allocation events. An optimal value for the resource bulk granularity can be determined with respect to the type of game serviced and its tolerance to resource shortages.

To observe the impact of the time bulk, we vary it from 0.1 to 2.0 days while keeping the resource bulks constant. To this end, we use the policies *HP-5*, and *HP-8* to *HP-11* described in Table IV. The results of experiment batch *B* presented in Table VII show that the efficiency of resource provisioning can be significantly improved by using resources from the data centres whose policies specify the shortest time bulks. The increase of the average under-allocation is low if the time bulks

are set to realistic values, e.g. above one hour.

2) *The combined impact of the resource and time bulks:* To study the combined impact of the resource and time bulks, we aggregate the policies used in the previous two experiments by taking all possible CPU-time bulk combinations (the other resource bulks are identical for all involved policies). The results from the experiment batch *C* depicted in Table VII follow the same pattern as those obtained when varying the time bulk in Section VI-D1 with higher under-allocation values. We conclude that changing the time bulk has a higher impact on the dynamic resource provisioning than changing the resource bulk.

E. MMOG Latency Tolerance Impact

We now investigate the impact of the MMOG latency tolerance on the quality of the dynamic resource provisioning. We consider an ideal network behaviour, where the latency between players and data centres is exclusively determined by the physical distance between them. The higher the latency tolerance of a MMOG, the further away the servers can be located from the users, and the longer the list of data centres from which resources can be dynamically provisioned is. We show in this experiment that higher latency tolerance leads to resources of the data centres with unsuitable hosting policies being unused when other suitable alternatives exist.

We define five classes of maximal physical distance between the players and the server locations, where in practice the distance values would depend on the design of each MMOG: (1) *same location*, when users must be handled by resources at the same location; (2) *very close*, when resources can be allocated within a radius of 1,000 km from their users; (3) *close*, within a radius of 2,000 km; (4) *far*, within a radius of 4,000 km; and (5) *very far*, when any server can serve any user. We consider from the setup described in Table III only the data centres located in the North American region, and select from the workload only the requests that arrive at these data centres. The hosting policies are coarse grained (i.e. with large resource and time bulks) for the data centres located on the US East Coast and become gradually finer grained for the data centres located at the Central and West Coast locations.

We first consider a restricted workload in which only the data centres from the US East Coast location receive allocation requests from game operators. Figures 12 shows the distribution of the allocated resources for various latency

TABLE VII
EXPERIMENTAL RESULTS SHOWING THE IMPACT OF DIFFERENT RESOURCE AND TIME BULKS ON RESOURCE PROVISIONING.

Experiment batch	Allocation time [hours]	CPU resource [units]	Over-allocation [%]	Under-allocation [%]	Under-allocation events
A: Time bulk variation	3	0.56	57.1	-0.028	78
	6	0.56	60.78	-0.022	62
	12	0.56	68.42	-0.020	54
	24	0.56	77.79	-0.014	46
	48	0.56	78.75	-0.008	23
B: Resource bulk variation	3	1.11	112.79	0	0
	3	0.56	51.03	-0.048	90
	3	0.37	57.1	-0.028	78
	3	0.28	46.62	-0.064	137
	3	0.22	38.95	-0.061	151
C: Time bulk variation [average over all resource bulk sizes]	3	0.22 – 1.11	56.83	-0.047	127
	6	0.22 – 1.11	59.86	-0.040	108
	12	0.22 – 1.11	66.36	-0.029	91
	24	0.22 – 1.11	79.93	-0.018	70
	48	0.22 – 1.11	81.16	-0.011	41

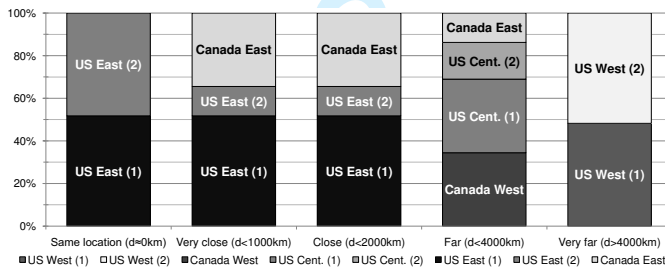


Fig. 12. Resource allocation distribution for US East Coast resource requests only and various latency tolerance values.

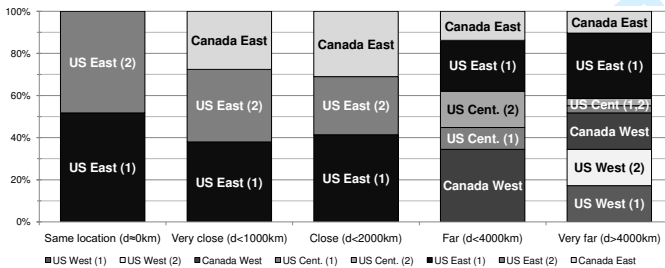


Fig. 13. Resource allocation distribution for all North American resource requests and various latency tolerance values.

tolerance values. As expected, because of the hosting policies setup, the desirability for a data centre increases with its distance from the US East Coast. For the maximal latency tolerance, all requests are served on resources located at the maximal distance from the US East Coast, i.e. on the data centres from the US West Coast.

We now investigate the system behaviour in a realistic situation under the combined workload of all North American game operators. Figure 13 shows for this setup the distribution of the allocated resources for various latency tolerance values. Due to resource contention, the resource allocation follows different patterns than in the optimal case. Figure 14 depicts the resource allocation for all North American data centres, which demonstrates that the US East Coast data centres with the most unsuitable hosting policies are penalised by having more unused resources, especially when the latency tolerance admits far and very far maximal service distances. In addition, the US East Coast requests are served under the best policies even in a busy system.

TABLE VIII
AVERAGE OVER- AND UNDER-ALLOCATION FOR CONCURRENT MMOGS.

MMOG workload structure	A [%]	B [%]	C [%]	Over-allocation [%]	Under-allocation [%]	$ \Upsilon _{\text{events}} > 1\%$
0	0	100		35.98	-0.12	240
5	5	90		36.89	-0.12	231
10	10	80		36.50	-0.12	220
25	25	50		35.80	-0.10	213
33	33	33		33.84	-0.09	200
0	100	0		33.24	-0.09	216
100	0	0		19.79	-0.03	75

F. Impact of Servicing Multiple MMOGs

In this last experiment, we evaluate the dynamic provisioning when servicing multiple types of MMOGs. We select three MMOG types with different interaction complexities, as defined in Section VI-C: *MMOG A* uses an $O(n \cdot \log(n))$, *MMOG B* an $O(n^2)$, and *MMOG C* an $O(n^2 \cdot \log(n))$ interaction complexity. We run seven scenarios in which the system has to handle resource requests for the three selected MMOG types in different proportions. The workload structure for each scenario is depicted in the “Workload structure” group of columns from Table VIII.

When the workload is dominated by the more compute-intensive *B* or *C* MMOGs, i.e. the first six rows in Table VIII, the behaviour of dynamic provisioning is stable. The over-allocation under a workload of only *C* MMOGs is less than 3% higher than under a workload of only *B* MMOGs. In the seventh scenario, when the workload comprises only the (less compute-intensive) *A* MMOGs, the dynamic resource provisioning is significantly better than in the other scenarios. We conclude that the quality of the dynamic provisioning is determined by its biggest consumer. Under these circumstances, game operators of a different MMOG type (e.g. type *A*) may find it more convenient to install their own infrastructure. As an alternative, we plan to investigate in future work the impact of prioritizing resource requests according to the MMOG interaction complexity.

VII. REAL-WORLD EXPERIMENT

We designed and implemented our MMOG ecosystem within the edutain@grid project [13], aiming to provide a platform for scalability, QoS provisioning, and user-friendly business models for real-time online interactive applications, with strong focus on MMOGs. edutain@grid is using as pilot




Fig. 14. Far maximal allocation distance (left) and Very far maximal allocation distance (right) for all North American data centres and resource requests.

a commercial FPS game application called Hunter developed by the Darkworks game development company with the headquarters in Paris, France. Parallelization of a game session according to the zoning, mirroring, and instancing techniques outlined in Section II-B is achieved by means of a generic portable library called Real-Time Framework (RTF) [32].

We used in our experiments an FPS demonstrator application that uses the open-source three-dimensional graphics engine (OGRE – <http://www.ogre3d.org/>) for graphics and sound, and the RTF library for game session parallelization. The resource testbed consists of six servers provided by the Amis telecommunication company located in Maribor, Slovenia and managed by one resource allocation service. The game world consists of two adjacent zones, handled in the beginning by two idle game servers. Clients connect to a game operator located at the University of Innsbruck, which processes their requests and arranges for the real-time connection the game servers. To stress the servers, we generated waves of incoming clients, implemented as non-player characters (or bots) and modelled using the profiles described in Section V-D.

We ran three distributed sessions using different resource provisioning methods. One session used our dynamic provisioning method and the other two were managed using a fixed threshold on the number of clients to create replication servers and migrate clients. We set the replication thresholds to 40 and 50 clients, respectively. The goal of each session was to accommodate a total of 190 clients which connect in four waves (one minute apart from each other) of 80, 30, 30 and 50 clients with as little resource utilisation as possible, and at the same time providing good QoS for the clients (i.e. the least under-allocation).

Figure 15 shows three histograms with the total number of clients, their zone distribution, the resource provisioning, and load balancing measures taken for the distributed session using our provisioning method. Figure 16 presents the under-allocation events for the three distributed sessions. The average under-allocation for the session managed by our dynamic resource provisioning method is 0.66%, and for the two sessions utilising the client threshold is 0.86% and 8.69%, respectively. Although the under-allocation for the client threshold method with a threshold of 40 is relatively close to the one exhibited by our method, there is a difference in the resource utilisation. In this case, the client threshold method utilises a total of six machines, whereas our provisioning method only needed five machines to accommodate the same amount of clients. It is possible to reduce the resource utilisation using a higher client threshold, but the resource under-allocation events increase drastically, as seen in the session with the client threshold of 50 which utilised only five machines and had an average under-




Fig. 15. Number of clients connected over time in a distributed session.




Fig. 16. Under-allocation event comparison.

allocation of 8.69%. We conclude that our dynamic resource provisioning method performs better than the rudimentary client threshold-based method by providing a better service to the clients while offering a good resource utilisation.

VIII. RELATED WORK

We have already reviewed in this article existing works related to our MMOG ecosystem, workload, and prediction models. We now turn our attention to the related work in the area of resource provisioning and identify three main directions from the resource provider's perspective.

A. Data Centres

The case when resources from one data centre are shared between multiple applications with statistical performance guarantees has received much attention [9], [10], [14], [33], [34]. In all these approaches the variables characterizing requests (such as "service time") can be expressed independently of the system state, for example with a random variable whose behaviour is well characterized by a well-known statistical distribution. However, an important component of the resource

1 demands in the MMOG ecosystem is the interaction between
2 players, which makes the resource demands of MMOGs very
3 different from traditional web applications. To express the
4 interaction between concurrent system users, in our model a
5 request is dependent on transient system state parameters, such
6 as the number of active players and the location of players (see
7 Section IV-A).

8 From this body of related work, much attention has been
9 given to modelling single-tier web applications; with the no-
10 table difference in expressing interaction between concurrent
11 users, these models are similar to ours. The MUSE system [14]
12 allocates periodically a percentage of the resource capacity for
13 each service class such that a target utilisation is achieved at
14 the data centre level. An approach based on virtual machines,
15 as opposed to physical resources, has also been explored [33].
16 Resource demand profiles for business applications are con-
17 structed for various durations (e.g. hour, day, week) and
18 resources are allocated in advance to provide statistical perfor-
19 mance guarantees [9]. A similar approach based on application
20 profiles was proposed in [10]. Closest to our work, the benefit
21 of provisioning resources from single data centres has been
22 evaluated for databases and Web services [8]. Our work differs
23 from this approach in two significant aspects. First, MMOGs
24 have a different load model and their load in particular depends
25 also on the interaction between users. Second, we consider
26 multiple data centres to handle the different load patterns in
27 different geographical locations specific to MMOGs.

28 Recently, research has focused on multi-tier models of web
29 applications [8], [34]. They use probability distributions to
30 characterise resource demands, queueing theory to analyse
31 the system, and proactive and reactive resource provisioning;
32 the target environment is a single data centre. In contrast,
33 our approach seems less sophisticated, but it already fosters
34 emerging behaviour in a multi data centre MMOG ecosystem,
35 that is, the emergence of a complex system from the large
36 number of simple choices and interactions; for example,
37 we have shown in Section VI-E evidence that hosters have
38 incentives to offer MMOG-friendly hosting policies when the
39 market is competitive.

40 B. Grid Computing

41 The problem of dynamically allocating geographically dis-
42 tributed resources to applications has been a popular topic
43 in Grid computing research. Recent work investigates mecha-
44 nisms for resource allocation across single- and multi-cluster
45 Grids [20], [35]. They assess the performance of various
46 resource allocation mechanisms for typical Grid workloads
47 comprising batches of scientific and engineering jobs [36].
48 Unlike MMOGs, scientific Grid applications do not change
49 their resource requirements at runtime. Moreover, the Grid
50 resource allocation policies only allow for whole resources be
51 allocated at a time, while our work also considers the sub-
52 unitary allocation sizes specific to business data centres.

53 Closest to our work, the industrial game hosting platform
54 Butterfly.net Grid (now renamed the Emergent Platform) [37]
55 uses Grid technology to provide on-demand access to cluster
56 resources. Their hosting policy only considers multi-unitary

resource bulks and long time bulks; as such, this platform fits
well into our MMOG ecosystem as a typical large hoster.

57 C. Peer-to-Peer Computing

58 Peer-to-peer computing has emerged as a scalable and low-
59 cost technology, and as a potential alternative to traditional
on-demand resource provisioning. When employing peer-to-
peer technology, the game operators make use of the resources
of their clients instead of renting them from hosters. The
NPSNET project [38] uses a peer-to-peer approach in which
all the game computation is performed on client resources. The
SimMud [39] project uses a similar approach to NPSNET,
but also balances and optimises the use of resources. How-
ever, three problems have prevented so far the adoption of
peer-to-peer technology for MMOGs: the lack of appropriate
business models, the wide-spread attempts of cheating, and
the low availability of peers observed for other peer-to-peer
systems (such as the Gnutella and the BitTorrent file sharing
networks [40], [41]).

60 IX. CONCLUSION AND FUTURE WORK

We focused in this work on MMOGs as a new type of
large-scale distributed simulation with a growing user base of
tens to millions of players. To ensure that the user demand is
satisfied at all times, game operators resort to static resource
provisioning by building and maintaining computing platforms
of up to 10,000 machines located on several continents for a
single MMOG. In this paper we proposed a more efficient
alternative based on the dynamic resource provisioning and
management of data centre resources. Ours is the first thorough
investigation of an MMOG ecosystem, that is, of a multi-
MMOG, multi-data centre environment.

We showed in this work that the number and the type of
interactions between players, and between players and the
environment, are an important contributor to the game load.
To address it, we have introduced a new MMOG model
that focuses on the interaction count and type between game
entities, shown that interaction leads to much more dynamic
resource demands than previously believed, and proposed a
novel prediction algorithm based on neural networks that is
fast yet accurate; our algorithm performed significantly better
than the six time predictors also investigated in this work.
We have further investigated the performance of the resource
provisioning and management of data centre resources with a
large variety of scenarios that focus both on MMOG-specific
properties and on data centre hosting policies. Most impor-
tantly, we have shown that the static resource provisioning can
be on average from five up to ten times more inefficient than
dynamic allocation under the same conditions, and that the
game operators can penalise the data centres with unsuitable
hosting policies, by not using their resources. Last, we have
designed and implemented our methods on top of the platform
offered by the EU project edutain@grid, and presented an
experiment showing the real-time resource provisioning for
a real game prototype.

REFERENCES

- [1] C. Neumann, N. Prigent, M. Varvello, and K. Suh, "Challenges in peer-to-peer gaming," *Computer Comm. Rev.*, vol. 37, no. 1, pp. 79–82, 2007.
- [2] W. M. White, C. Koch, N. G. 0003, J. Gehrke, and A. J. Demers, "Database research opportunities in computer games," *SIGMOD Record*, vol. 36, no. 3, pp. 7–13, 2007.
- [3] R. Bartle, *Designing Virtual Worlds*. New Riders Games, 2003.
- [4] A. Shaikh, S. Sahu, M.-C. Rosu, M. Shea, and D. Saha, "On demand platform for online games," *IBM Systems Journal*, vol. 45, no. 1, pp. 7–20, 2006.
- [5] T. Beigbeder, R. Coughlan, C. Lusher, J. Plunkett, E. Agu, and M. Claypool, "The effects of loss and latency on user performance in unreal tournament 2003," in *NETGAMES*, 2004, pp. 144–151.
- [6] B. Hack, M. Morhaime, J.-F. Grollemund, and N. Bradford, "Introduction to vivendi games," Presentation. [Online] Available: <http://www.vivendi.com/>, Jun 2006.
- [7] R. P. Doyle, J. S. Chase, O. M. Asad, W. Jin, and A. Vahdat, "Model-based resource provisioning in a web service utility," in *USENIX Symposium on Internet Technologies and Systems*, 2003.
- [8] B. Urgaonkar, P. J. Shenoy, A. Chandra, and P. Goyal, "Dynamic provisioning of multi-tier internet applications," in *International Conference on Autonomic Computing*. IEEE CS Press, 2005, pp. 217–228.
- [9] J. Rolia, X. Zhu, M. F. Arlitt, and A. Andrzejak, "Statistical service assurances for applications in utility grid environments," *Perform. Eval.*, vol. 58, no. 2-3, pp. 319–339, 2004.
- [10] A. Karve, T. Kimbrel, G. Pacifici, M. Spreitzer, M. Steinder, M. Svridenko, and A. N. Tantawi, "Dynamic placement for clustered web applications," in *WWW*. ACM Press, 2006, pp. 595–604.
- [11] Jagex, Ltd., "Runescape," MMOG. [Online] Available: <http://www.runescape.com/>, Nov 2007.
- [12] K.-T. Chen, P. Huang, and C.-L. Lei, "Game traffic analysis: An mmorg perspective," *Computer Networks*, vol. 50, no. 16, pp. 3002–3023, 2006.
- [13] T. Fahringer, C. Anthes, A. Arragon, A. Lipaj, J. Müller-Iden, C. Rawlings, R. Prodan, and M. Surridge, "The edutain@grid project," in *GECON*, ser. LNCS, vol. 4685. Springer, August 2007, pp. 182–187.
- [14] M. Aron, P. Druschel, and W. Zwaenepoel, "Cluster reserves: a mechanism for resource management in cluster-based network servers," in *SIGMETRICS*, 2000, pp. 90–101.
- [15] M. Ye and L. Cheng, "System-performance modeling for massively multiplayer online role-playing games," *IBM Systems Journal*, vol. 45, no. 1, pp. 45–58, 2006.
- [16] W. Cai, P. Xavier, S. J. Turner, and B.-S. Lee, "A scalable architecture for supporting interactive games on the internet," in *16th Workshop on Parallel and Distributed Simulation*, 2002, pp. 60–67.
- [17] J. Miller-Iden and S. Gorlatch, "Rokkatan: scaling an RTS game design to the massively multiplayer realm," *Computers in Entertainment*, vol. 4, no. 3, p. 11, 2006.
- [18] A. R. Bharambe, J. R. Douceur, J. R. Lorch, T. Moscibroda, J. Pang, S. Seshan, and X. Zhuang, "Donnybrook: enabling large-scale, high-speed, peer-to-peer games," in *SIGCOMM*. ACM Press, 2008, pp. 389–400.
- [19] M. Claypool, "The effect of latency on user performance in real-time strategy games," *Computer Networks*, vol. 49, no. 1, pp. 52–70, 2005.
- [20] A. Iosup, D. Epema, T. Tannenbaum, M. Farrellee, and M. Livny, "Interoperating grids through delegated matchmaking," in *SC*. ACM Press, 2007.
- [21] B. S. Woodcock, "An analysis of mmog subscription growth," Report, 21 Edition. [Online] Available: <http://www.mmogchart.com>, Jun 2006. [Online]. Available: <http://www.mmogchart.com>
- [22] W.-C. Feng, F. Chang, W.-C. Feng, and J. Walpole, "A traffic characterization of popular on-line games," *IEEE/ACM Trans. Netw.*, vol. 13, no. 3, pp. 488–500, 2005.
- [23] W.-C. Feng, D. Brandt, and D. Saha, "A long-term study of a popular MMORPG," in *NETGAMES*. ACM Press, 2007, pp. 6–11.
- [24] D. Pittman and C. Gauthier, "A measurement study of virtual populations in massively multiplayer online games," in *NETGAMES*. ACM Press, 2007, pp. 25–30.
- [25] BBC News, Technology, "British gaming firm takes on the world," News Item. [Online] Available: <http://news.bbc.co.uk/2/hi/technology/7090490.stm>, Nov 2007.
- [26] The Screen Digest Group, "Western world MMOG market: 2006 review and forecasts to 2011," Research Report. [Online] Available: <http://www.screen Digest.com/Reports>, Mar 2007.
- [27] D. A. Menascé, V. Almeida, R. H. Riedi, F. Ribeiro, R. C. Fonseca, and W. M. Jr., "A hierarchical and multiscale approach to analyze e-business workloads," *Perform. Eval.*, vol. 54, no. 1, pp. 33–57, 2003.
- [28] G. E. P. Box, G. M. Jenkins, and G. C. Reinsel, *Time Series Analysis, Forecasting and Control*. Prentice Hall, 1994.
- [29] C. M. Bishop, *Neural Networks for Pattern Recognition*. Oxford University Press, 1996.
- [30] I. GameData, "Counter strike," <http://www.counter-strike.com>.
- [31] V. Nae, R. Prodan, and T. Fahringer, "Neural network-based load prediction for highly dynamic distributed online games," in *Euro-Par*. Springer Verlag, 2008.
- [32] F. Glinka, A. Ploss, J. Miller-Iden, and S. Gorlatch, "RTF: A real-time framework for developing scalable multiplayer online games," in *NetGames*. ACM Press., 2007.
- [33] P. Padala, K. G. Shin, X. Zhu, M. Uysal, Z. Wang, S. Singhal, A. Merchant, and K. Salem, "Adaptive control of virtualized resources in utility computing environments," in *EuroSys*. ACM Press, 2007, pp. 289–302.
- [34] B. Urgaonkar, P. J. Shenoy, A. Chandra, P. Goyal, and T. Wood, "Agile dynamic provisioning of multi-tier internet applications," *TAAS*, vol. 3, no. 1, 2008.
- [35] M. Siddiqui, A. Villazón, and T. Fahringer, "Grid allocation and reservation - grid capacity planning with negotiation-based advance reservation for optimized qos," in *ACM/IEEE Conference on High Performance Networking and Computing (SuperComputing)*, 2006, p. 103.
- [36] A. Iosup, C. Dumitrescu, D. H. Epema, H. Li, and L. Wolters, "How are real grids used? The analysis of four grid traces and its implications," in *GRID*. IEEE CS Press, 2006, pp. 262–270.
- [37] Gamebryo, "Butterfly Grid/Emergent Platform," [Online] Available: <http://www.emergent.net/>, Aug 2008.
- [38] M. R. Macedonia, D. P. Brutzman, M. J. Zyda, D. R. Pratt, P. T. Barham, J. Falby, and J. Locke, "NSPNET: A multiplayer 3D virtual environment over the internet," in *SI3D'95*, 1995, pp. 93–94.
- [39] B. Knutsson, H. Lu, W. Xu, and B. Hopkins, "Peer-to-peer support for massively multiplayer games," in *INFOCOM*, 2004, pp. 96–107.
- [40] M. Ripeanu, A. Iamnitchi, and I. T. Foster, "Mapping the gnutella network," *IEEE Internet Computing*, vol. 6, no. 1, pp. 50–57, 2002.
- [41] J. Pouwelse, P. garbacki, D. Epema, and H. Sips, "The Bittorrent P2P file-sharing system: Measurements and analysis," in *IPTPS*, 2005.