The Impact of Virtualization on the Performance of Massively Multiplayer Online Games

Vlad Nae, Radu Prodan, Thomas Fahringer
Institute for Computer Science
University of Innsbruck
Technikerstrasse 21a, 6020 Innsbruck, Austria
Email: {vlad,radu,tf}@dps.uibk.ac.at

Alexandru Iosup
Parallel and Distributed Systems
Delft University of Technology
Mekelweg 4, 2628CD Delft, Netherlands
Email: A.Iosup@tudelft.nl

***Abstract*—Today's highly successful Massively Multiplayer Online Games (MMOGs) have millions of registered users and hundreds of thousands of active concurrent users. As a result of the highly dynamic MMOG usage patterns, the MMOG operators pre-provision and then maintain throughout the lifetime of the game tens of thousands of compute resources in data centers located across the world. Until recently, the difficulty of porting the MMOG software services to different platforms made it impractical to dynamically provision resources external to the MMOG operators' data centers. However, virtualization is a new technology that promises to alleviate this problem by providing a uniform computing platform with minimal overhead. To investigate the potential of this new technology, in this paper we propose a new hybrid resource provisioning model that uses a smaller and less expensive set of self-owned data centers, complemented by virtualized cloud computing resources during peak hours. Using real traces from RuneScape, one of the most successful contemporary MMOGs, we evaluate with simulations the effectiveness of the on-demand cloud resource provisioning strategy for MMOGs. We assess the impact of provisioning of virtualized cloud resources, analyze the components of virtualization overhead, and compare provisioning of virtualized resources with direct provisioning of data center resources.**

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 real-time virtual world entertaining millions of players spread across the globe [1]. To support the variable and latency-aware resource demands of the players, the MMOG operators maintain expensive, world-wide multi-server infrastructures. In contrast, the new cloud computing technology based on the virtualization of resources promises to provide homogeneous resources for MMOG hosting. However, the virtualization overheads may cancel out the benefits of provisioning resources when they are actually needed. To understand the trade-offs of virtualization, in this work we investigate the impact of virtualization on the performance of MMOGs.

The current industry approach based on resource ownership guarantees that resources are available when needed, is subject to over-provisioning. The operating infrastructures of the leading MMOGs such as World of Warcraft [2] comprise each thousands of computers in tens of physical locations; resource ownership can take up to 40% of the game revenue [3]. However, the resource demand of MMOGs is highly dynamic [4]; to guarantee good quality of service, a large portion of the pre-provisioned resources are unnecessary. We have proposed in previous work [4] a new MMOG ecosystem consisting of game operators and multiple data centers, and studied the impact of various resource lease policies on resource allocation and provisioning for MMOGs. We have shown through simulations that dynamic resource provisioning has the potential to considerably reduce the MMOG operation costs with a reasonable performance loss (expressed as quality of service breaches). However, in our previous work we have not taken into account the performance penalties incurred by the virtualization overheads, or the competition for resources from multiple MMOG operators. Towards this end, our contribution is twofold:

1) We extend our MMOG ecosystem to accommodate both non-virtualized and virtualized resources, and multiple MMOG operators (Section III);
2) We evaluate through trace-based simulation the use of virtualized and non-virtualized resources, either separately or in conjunction, in scenarios involving one or more MMOG providers (Section IV).

II. Background

In this section we summarize our previous work on the MMOG ecosystem and the virtualization concept.

A. The MMOG Ecosystem

We have previously modeled [4] the MMOG ecosystem with two components: the MMOG operators and the MMOG platform providers (*hosters*). Each MMOG is managed by one *game operator* that is responsible for the real-time experience of the connected players and which negotiates with existing data center and cloud providers the necessary resources in order to achieve this goal. The hosters act as a global network of conventional *data centers* complemented by a set of virtualized *cloud computing providers* that host in cooperation multiple MMOG sessions at the same time.

Hosters The hosters operate two major services. A *load prediction* service, presented in detail in [5], is in charge of projecting the future distribution of entities in the game world that is demonstrated to have the highest impact on the server load. We devised accurate analytical models for translating the entity distribution prediction and possible interactions into estimating the game server load. Based on the projected load,

978-1-4244-5605-5/09/$26.00 ©2009 IEEE

a *resource allocation* service [4] provisions additional local servers to the game session (through the zoning, replication, or instancing parallelization techniques) that accommodate the player load while guaranteeing the real-time quality of service constraints. For example, by timely foreseeing critical hot-spots (i.e. excessively populated area of interest generating a large number of interactions), one can dynamically provision additional servers on some new resources and take timely load balancing actions that transparently redistribute the game load before the servers become overloaded.

Our hosting model considers the *size* and *duration* of the minimal resource allocation which may be not only for a resource as a whole (e.g. a server in Web data centers [6] or a multi-processor node in a Grid system [7]), but also for a fraction of that resource (e.g. a virtual machine running on a physical node [8], 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.

Operators The game operators handle simultaneously MMOG sessions of different types. The game operators issue resource requests based on the predicted load of the games they operate (either statically or dynamically computed), and the hosters respond with offers based on their local time-space renting policy. Depending on the hoster's service model (either best-effort or advance reservation-based), resource requests are queued or immediately fitted in the schedule, respectively. Using one or several important metrics (e.g. virtualization overheads, geographical proximity, data locality, resource proportionality), the game operator applies a *resource selection policy* using one or several of the following four operations:

1) *classifying* groups several resources into classes based on metric value ranges;
2) *sorting* orders the resources based on metric values;
3) *filtering* eliminates resources based on metric values;
4) *prioritising* gives higher allocation priority to resources with important metric values.

The operators submit resource requests to the hoster by specifying the type, number, and duration for which the resources are desired. As with our previous work [4], we currently consider four resource types that are relevant for MMOG hosting: *CPU*, *memory*, input from the external network (*ExtNet[in]*), and output to the external network of a data center (*ExtNet[out]*). Once the available resources are selected, they 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

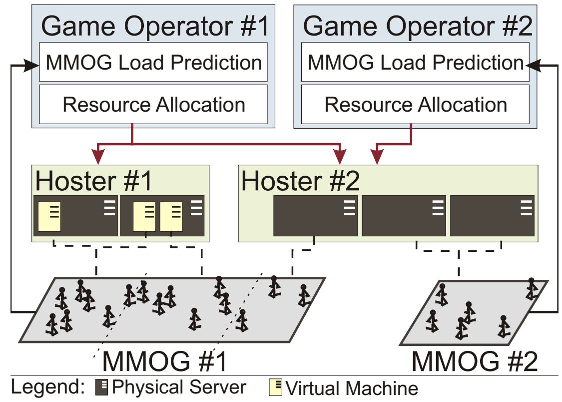

Fig. 1. The MMOG ecosystem architecture.

for the entire duration of the game operator's request (task preemption or migration are not supported).

B. Virtualization and Clouds

Virtualization is a key technology for hiding from the users the low-level physical characteristics of a computing platform by showing another abstract, higher-level emulated platform instead represented by a so called *Virtual Machine*(VM). The most common virtualization environments today are VMWare [9] and Xen [10], while there exist other solutions used by smaller communities such as Oceano, VMPlants, Kadeploy, XGE and KVM. Using virtualization, a computing cloud may address with the same shared set of physical resources a large user base with different needs, cheaply.

Similar to scientific computing software and other complex software, MMOGs have many library dependencies that make the deployment process complex and requiring manual intervention. In most cases, special requirements have to be fulfilled as discussed in [11] (for scientific applications), which makes the automatic deployment and provisioning of resources impractical without additional platform support such as virtualization. Although the virtualization layer adds performance overheads (that we will model in Section III-B), correctly configured VM images can be used to automatically deploy and further operate MMOGs on many hardware platforms. Thus, the use of virtualization promises to ensure deployment and provisioning scalability; we characterize the trade-offs raised by MMOG virtualization in the remainder of this work.

III. The Extended MMOG Model

In this section we extend the MMOG ecosystem introduced in [4] from multiple MMOG sessions operated by a single game operator on non-virtualized resources to multiple MMOG sessions operated by several competing game operators on non-virtualized and virtualized resources.

A. MMOG Hosting and MMOG Operation

Hosters Our enhanced hosting model consists of *both* conventional data centers that operate pre-installed game servers [4] *and* additional general purpose providers of virtualized cloud resources (see Figure 1). Thus, the extended hosting platform consists of both data centers and cloud providers scattered around the world, where each hoster pools together resources that may serve several games simultaneously. We add to the model a virtualization policy, which describes the characteristics of the virtualization process employed by a hoster; these characteristics are detailed in Section III-B.

Operators In our previous work [4] the MMOG operators employed four operations to select resources: classifying, sorting, filtering, and prioritizing (see Section II-A). The actual implementation and use four operations form a versatile policy-based resource selection mechanism. In this work we consider three such policies:

1) *"efficient"* targets a low over-allocation by selecting the offers with the values closest to the current request (sorting) and shortest allocation time (classifying);
2) *"good service"* targets a low under-allocation by first selecting the offers proportional to the request (sorting), then the offers geographically closest to the request (classifying), and finally the offers with the lowest virtualization overhead (prioritising);
3) *"balanced"* is a compromise between the "efficient" and the "good service" mechanisms trying to achieve a low under-allocation while keeping reasonable over-allocation thresholds. This policy first selects the proportional offers to the request (sorting) and then classifies on virtualization overheads.

B. Virtualization Overhead

We consider in this work two aspects of the virtualization overhead, corresponding to the VM instantiation and to the VM execution. In the following we detail each of these two components, in turn.

Instantiation VM instantiation is the process by which a VM is started on a selected resource. Our model for VM instantiation that considers four performance aspects expressed by the corresponding *virtualization overheads*: VM image preparation (characterized by the time t_c), VM transfer (t_x), VM start (t_s), and VM removal (t_r). Performance-wise, the total time needed to instantiate a VM can be expressed as $T = t_c + t_x + t_s + t_r$. Conventional data centers do not exhibit this overhead, but are restricted to pre-deployed software and thus lack the flexibility of dynamic provisioning and on-the-fly deployment of MMOG servers. Table I presents realistic values for the modeled parameters, and summarizes them into a theoretical average VM (the "Avg. VM" row). We set the VM preparation time to zero after assuming conservatively that MMOG images can only be pre-created, the VM start time to 80 seconds as determined in our evaluation of real commercial clouds [12], and the VM removal time to zero as images are not saved to avoid data corruption across sessions.

Execution Performance studies using general purpose benchmarks have shown that the overhead incurred by virtualization can be below 5% for computation and 15% for networking [13], and similar for other resources that are relevant for MMOG hosting; we call an aggregate of this performance overhead the *virtualization penalty*. We expect that future specialized MMOG hosters that will employ cloud computing/virtualization technology will achieve a similarly low virtualization penalty.

IV. EXPERIMENTAL RESULTS

Based on the extended model introduced in the previous section, we have evaluated through trace-based simulation the use of virtualized and non-virtualized resources, either separately or in conjunction, in scenarios involving one or more MMOG providers. We summarize the results of the evaluation in this section. We first show in Section IV-B that the use of virtualized resources without adequate policies can incur severe degradation of the gameplay experience. Then, in Section IV-C we analyze the impact of the virtualization parameters on the gameplay experience. Finally, in Section IV-D we evaluate a complex scenario involving multiple MMOG providers competing for resources from the same pool.

A. Experimental Setup

Input Workload We performed experiments using traces taken in our previous work [4] from RuneScape [14], a real MMOG ranked second after World of Warcraft by number of active paying customers in the US and European markets [4]. RuneScape 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 be trivially computed, for example using the linear models employed in [15].

Environment We performed experiments in a simulated RuneScape-like environment with the input workload consisting of the first two weeks from the trace data. The traces are sampled every two minutes (called *simulation steps*) and contain the number of players over time for each server group used by the game operators. This gives over 10,000 metric samples for each simulation, ensuring statistical soundness. The data centers are located on three continents and seven countries; their characteristics are summarized in Table II. For virtualized resources, we used the VM overhead model introduced in Section III-B.

Performance Metrics We evaluate the quality of the game experience using a *resource under-allocation* metric that characterizes the percentage of resources that have not been allocated from the amount necessary for the seamless execution of the MMOG. We define resource under-allocation $U(t)$ (as percentage) within one simulation step T as:

Location		Data	Machines
Continent	Country	Centers	(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
	Canada (East)	1	10
Australia	Australia	2	8

TABLE II
DATA CENTER PHYSICAL CHARACTERISTICS.

Real VM middleware	Image Creation	Transfer	Start	Removal
Oceano	0 (pre-created)	0 (present)	$130s$	n/a
VMPlants	$27s + 100s \times S$	$90s + 80s \times S$		n/a
Kadeploy	0 (pre-created)	$200s + 0.33s \times N$		$0.2s \times N$
Shirako	$100s$	$20s + 2s \times N$		n/a
VW	0 (pre-created)	$110s + 54s \times N$		n/a
VD caches	$24.8s + 0.09s \times N$		n/a	n/a
Avg. VM	$t_c = 0s$	$t_x = Xfer(S) + 0.09s \times N$	$t_s = 80s$	$t_r = 0s$

TABLE I

REALISTIC PARAMETER VALUES FOR THE VM INSTANTIATION MODEL, WHERE S IS THE SIZE OF THE DATA TO TRANSFER IN GB, N IS THE NUMBER OF VMS TO INSTANTIATE, $Xfer(S)$ IS THE TRANSFER TIME FOR DATA OF SIZE S FROM THE DATA SOURCE TO THE RESOURCE THAT INSTANTIATES THE VM (E.G. FOR A 1GBPS TRANSFER BANDWIDTH IN IDEAL CONDITIONS $Xfer^{ideal}_{1Gbps}(S) = 10s \times S$).

Policy	CPU [units]	Mem. [units]	Time [min.]	t_s [s]	t_x [s]	Virt. penalty
Ideal	0.4	0.25	360	0	0	0
Policy 1	0.4	0.25	360	30	5	1%
Policy 2	0.4	0.25	360	60	10	3%
Policy 3	0.4	0.25	360	90	20	5%
Policy 4	0.4	0.25	360	120	40	10%

TABLE III

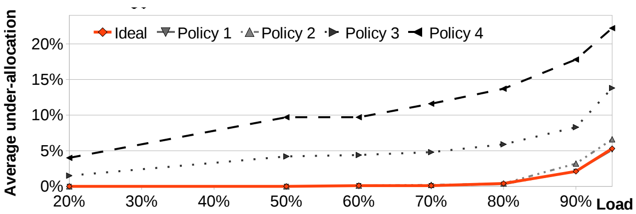

Fig. 2. The impact of the virtualization policy on the resource under-allocation under increasing load.

$U(t) = \frac{(A-L)\cdot t}{T\cdot L_{\max}} \cdot 100$, where t is the duration of the under-allocation event, A is the amount of allocated resources, L is the amount of needed resources (measured from the traces), and $L_{\max}$ is the maximum load determined by game design. Resource under-allocation has an important impact on the gameplay experience: not having enough resources quickly increases the game response time and degrades the game experience. For example, in a FPS game, during an under-allocation event, the minimum required client update frequency will not be met (e.g. updates at less than 25Hz). As a result, players may quit triggering mass-departure events [4]. The exact relationship between resource under-allocation and game experience is genre and even game-dependent, but can be determined through traditional gameplay testing [16].

B. Verifying the Virtualization Premise

An important premise of our work is that the virtualization policy has an important impact on the resource under-allocation, and thus on the quality of the game experience. We now show that this is indeed the case. Towards this end, in this experiment we assess the resource under-allocation for five virtualization policies that can be used by hosters in practice. The five policies are summarized in Table III. The "Ideal" policy has no virtualization overheads, while the policies "Policy 1" through "Policy 4" have increasingly larger virtualization overheads and thus become increasingly unfavorable for MMOG hosting.

We use the five virtualization policies, in turn, with the experimental setup presented in Section IV-A. Figure 2 shows the impact of the five virtualization policies on the resource under-allocation, under increasing load. The vertical distance between the curves corresponding to each policy indicates that different virtualization policies have an important impact on the average resource under-allocation. In particular, Policies 3 and 4 incur large resource under-allocations, which would affect severely the quality of the game experience. The small difference between Policies 1 and 2, and the Ideal virtualization policy, is due to the use of a load prediction service (see Section II-A) that effectively cancels the overheads of virtualization as resources are allocated before they are actually required. We conclude that the virtualization policy has an important impact on resource under-allocation.

Experiment Focus	Value Ranges for Characteristic t_s [seconds]	S size [GB]	Xfer BW [Mbit]	Virt. penalty
VM start	$20-170$	0.5	100	5%
VM size	80	$0.25-1.5$	100	5%
VM x'fer. bw.	80	0.5	$100-1000$	5%
Virt. penalty	80	0.5	100	$2-12\%$

TABLE IV

THE VALUE RANGES FOR THE CHARACTERISTICS INVESTIGATED IN SECTION IV-C. EACH ROW PRESENTS THE VALUE OR VALUE RANGE USED IN THE EXPERIMENT THAT INVESTIGATES ONE FOCUS CHARACTERISTIC.

C. Analysis of Virtualization Policies

To explain the effects of the virtualization policy on the game experience we explore the impact of each of the components of the virtualization policy on the resource under-allocation, when a single MMOG is hosted.

We used the experimental setup presented in Section IV-A, and varied in turn each of the four characteristics affecting the virtualization overheads introduced in Section III-B: the VM start time (t_s), the VM size (S), the bandwidth available for transferring the VM image, and the virtualization penalty. The ranges of values explored in the four separate experiments are summarized in Table IV. In each experiment, the single values are selected based on realistic data published by major providers of virtualized resources and other benchmark reports. We set the VM size set to 0.5 GB (which is average for base images with pre-deployed software), the transfer bandwidth to 100 Mbps, and the performance penalty of the game servers running inside virtualized resources to 5% [13]. We also study in each experiment the impact of the system load, for which we look at four values that are common in practice: 70%, 80%, 90%, and 95%.

Figures 3 and 4 show that average resource under-allocation increases linearly with the VM size and VM start time. The same growing pattern holds for different degrees of relative

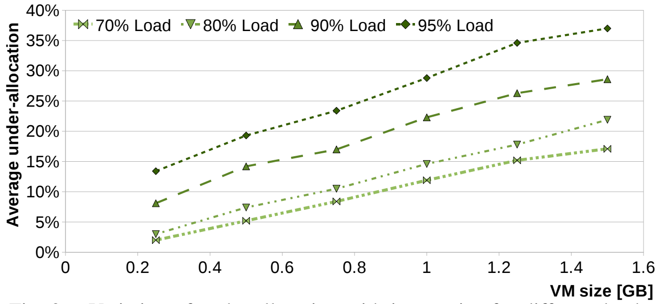

Fig. 3. Variation of under-allocation with image size for different loads.

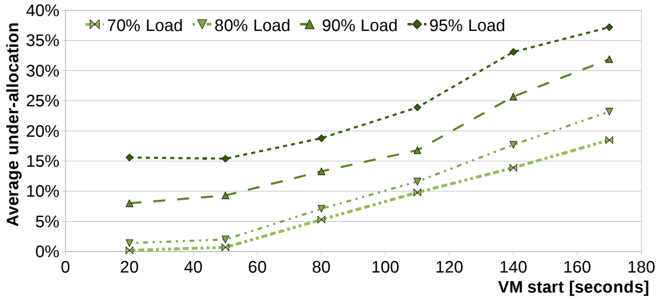

Fig. 4. Variation of under-allocation with image boot time for different loads.

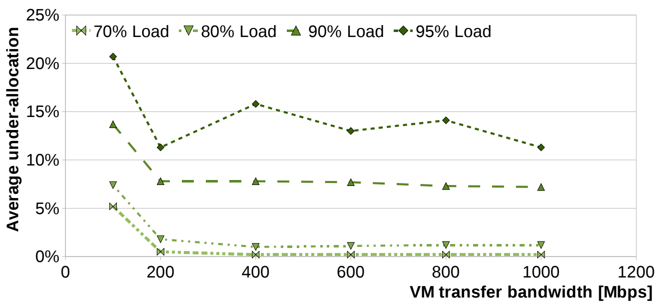

Fig. 5. Variation of under-allocation with image transfer bandwidth.

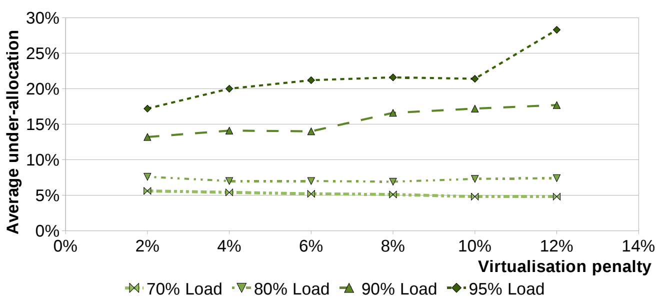

Fig. 6. Variation of under-allocation with virtualization penalty.

load with no artifacts. We conclude that the impact of these parameters on the quality of gameplay is predictable, the resource under-allocation roughly growing by 10% per GB of VM image size and for each minute of VM start time.

Figures 5 and 6 show that at loads lower than 90% the VM transfer bandwidth and virtualization penalty have little or no impact on the resource under-allocation, since their effect is negligible compared to the other virtualization parameters and can be easily hidden by the proactive (prediction-based) resource allocation. However, at relative load values higher than 90%, the under-allocation exhibits irregular behaviour.

Policy	VM start [s]	Image size [GB]	Xfer BW [Mbps]	Virt. penalty	Time [min.]
No Virt	0	0	0	0%	60
Virt 1	60	0.25	400	2%	120
Virt 2	80	0.50	300	5%	180
Virt 3	120	0.75	200	10%	240
Virt 4	180	1.00	100	12%	300

TABLE V
VIRTUALIZATION POLICIES FOR THE MULTI-MMOG HOSTING SCENARIO.

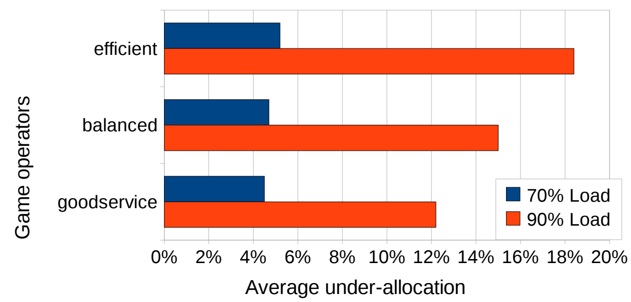

Fig. 7. Performance comparison of the three resource selection policies

D. Resource Competition for Multi-MMOGs

In this last experiment we investigate the impact of the policy for resource allocation employed by MMOG operators (see Section III-A) when operators with different policies compete for resources. We used the experimental setup presented in Section IV-A and considered several MMOG operators that use the multi-hoster environment. The three resource allocation policies for MMOG operators are assigned in equal proportion (33%) to the operators. Additionally, we designed the five different policies summarized in Table V, which we assigned to hosters in equal proportion (20%). Policy "No Virt" represents the allocation policy for non-virtualized resources (ideal from a virtualization standpoint), while the policies "Virt 1" through "Virt 4" are defined with increasingly unfavorable virtualization parameter values (overheads). The competition for resources is ensured by setting the average load high; we experiment with two values, 70% and 90%.

Figure 7 depicts the average under-allocation for the three resource allocation policies, for 70% and 90% load. The similar performance of the policies at 70% load is due to the lack of competition for resources. In contrast, for the 90% load the "good service" game operator provides the best under-allocation, followed by the "balanced" and the "efficient" operators. We conclude that the selection mechanism is crucial for achieving good performance when the resources are in short supply and game operators compete for offers.

To better understand the behavior of the resource allocation policies, Figure 8 shows the distribution of allocated resources for each policy. The "good service" exhibits a strong preference towards the resources with low virtualization overheads and is using the least amount of resources provisioned from Virt 4 (unfavorable) hosting. The "efficient" game operator favors Virt 1 resources, which have the shortest time bulk but high virtualization overheads.

V. RELATED WORK

We have reviewed throughout this article work related to our MMOG ecosystem. 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: data centers, Grid computing, and cloud computing.

The case when resources from one data center are shared between multiple applications with statistical performance guarantees has received much attention [6], [8], [17]. In all these approaches the variables characterising requests (such as "service time") can be expressed independently of the system state, for example with a random variable whose behaviour is well characterised by a well-known statistical distribution. In contrast, an important component of the resource demands in the MMOG ecosystem is the interaction between players [4], which makes the resource demands of MMOGs very different from traditional web applications.

The problem of dynamically allocating geographically distributed resources to applications has been a popular topic in Grid computing [7], [18] and, more recently, in cloud computing research [19], [20]. Work in these areas investigates mechanisms for resource allocation across single- and multi-cluster grids and clouds, but only for scientific workloads and for unitary resource allocations. Closest to our work, the industrial game hosting platform Butterfly.net Grid (now renamed the Emergent Platform) [21] uses Grid technology to provide on-demand access to cluster 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.

VI. Conclusion and Ongoing Work

The MMOG operators who can afford it currently operate dedicated multi-server infrastructures spread across the world to support the highly variable resource demand and to cope with the players' sensitivity to latency. The recent introduction of virtualization in data centers promises much lower costs through uniform-behaving resources that can be provisioned on-demand for MMOG hosting from specialized commercial clouds. However, the expected cost reduction also comes with a loss of performance due to virtualization overheads. In this work we investigate the impact of virtualization on the performance of MMOGs. Towards this end, we have first proposed an enhanced MMOG ecosystem that extends on previous work with two new aspects: (1) multiple competing game operators managing several MMOG sessions; and (2) provisioning of virtualized cloud computing resources for MMOG hosting together with conventional data center machines. Then, we have investigated through trace-based simulation the effect of using virtualized resources on the quality of gameplay. We have learned that using virtualized resources can negatively affect the quality of gameplay when the system is heavily loaded, which is common in practice. We have shown that both the virtualization policies employed by MMOG hosters and the resource allocation policies used by MMOG operators are important for achieving good performance. Notably, we find that MMOG hosting on virtualized resources is practical if the virtualization overheads are taken into account, although non-virtualized resources are still preferable for good performance.

We are currently conducting more experiments that study other interesting aspects of multi-MMOG operation in competition for resources such as the virtualization impact on cost and resource over-allocation, and different player interaction and complexity models.

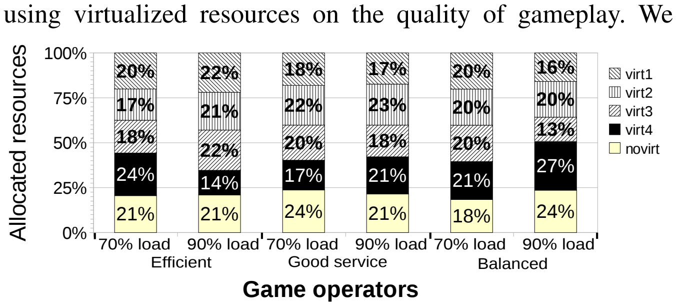

Fig. 8. Distribution of virtualized resources among the game operators at 70% and 90% load.

References

[1] R. Bartle, *Designing Virtual Worlds*. New Riders, 2003.
[2] I. Blizzard Entertainment, "World of warcraft," http://www.worldofwarcraft.com/.
[3] DFC Intelligence Team, "Online games report," 2004, [Online] Available: http://www.dfcint.com. [Online]. Available: http://www.dfcint.com
[4] V. Nae, A. Iosup, R. Prodan, D. Epema, and T. Fahringer, "Efficient management of data center resources for massively multiplayer online games," in *SC*, 2008.
[5] V. Nae, R. Prodan, and T. Fahringer, "Neural network-based load prediction for highly dynamic distributed online games," in *Euro-Par*. Springer Verlag, 2008.
[6] 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.
[7] A. Iosup, T. Tannenbaum, M. Farrellee, D. H. J. Epema, and M. Livny, "Inter-operating grids through delegated matchmaking," *Sci. Prog.*, vol. 16, no. 2-3, pp. 233–253, 2008.
[8] A. Karve, T. Kimbrel, G. Pacifici, M. Spreitzer, M. Steinder, M. Sviridenko, and A. N. Tantawi, "Dynamic placement for clustered web applications," in *WWW*, 2006, pp. 595–604.
[9] B. Walters, "Vmware virtual platform," *Linux Journal*, p. 6, 1999.
[10] D. Chisnall, *The Definitive Guide to the Xen Hypervisor*. Prentice Hall International, 2007.
[11] R. Bradshaw, N. Desai, T. Freeman, and K. Keahey, "A scalable approach to deploying and managing appliances," in *TeraGrid*, Madison, WI, 2007.
[12] S. Ostermann, A. Iosup, N. Yigitbasi, R. Prodan, T. Fahringer, and D. Epema, "An early performance analysis of cloud computing services for scientific computing," Delft University of Technology, PDS-2008-006, 2008.
[13] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. L. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, "Xen and the art of virtualization," in *SOSP*, 2003, pp. 164–177.
[14] Jagex, Ltd., "Runescape," http://www.runescape.com.
[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] T. Fullerton, C. Swain, and S. Hoffman, *Game Design Workshop: Designing, Prototyping, and Playtesting Games*. CMP Books, 2005.
[17] 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.
[18] M. Siddiqui, A. Villazón, and T. Fahringer, "Grid capacity planning with negotiation-based advance reservation for optimized QoS," in *SC*, 2006, p. 103.
[19] M. R. Palankar, A. Iamnitchi, M. Ripeanu, and S. Garfinkel, "Amazon S3 for science grids: a viable solution?" in *DADC*. ACM, 2008, pp. 55–64.
[20] E. Walker, "Benchmarking Amazon EC2 for HP Scientific Computing," *Login*, vol. 33, no. 5, pp. 18–23, Nov 2008.
[21] Gamebryo, "Butterfly Grid/Emergent Platform," [Online] Available: http://www.emergent.net/, Aug 2008.