

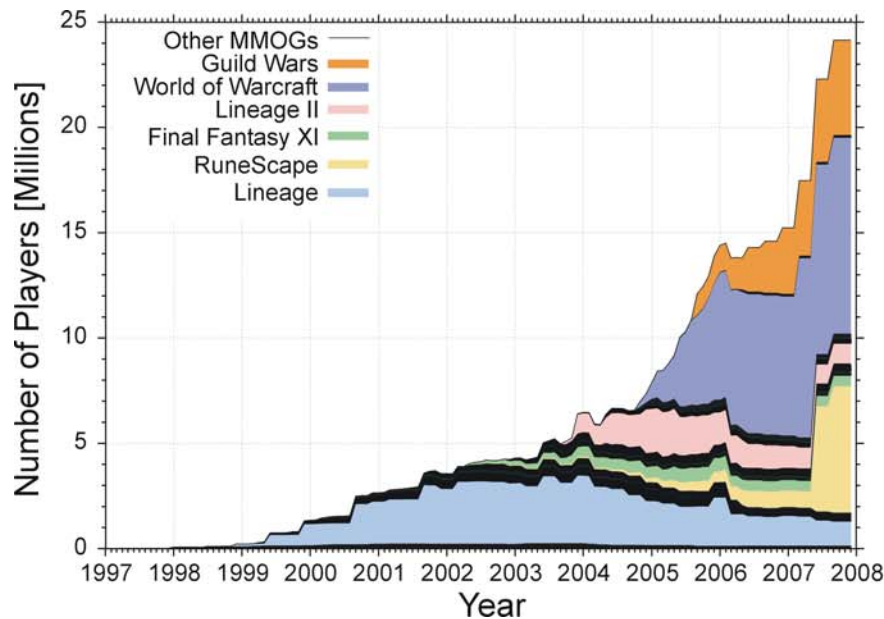
CAMEO: Continuous Analytics for Massively Multiplayer Online Games

Alexandru Iosup


Parallel and Distributed Systems Group
Delft University of Technology

MMOGs are a Popular, Growing Market

- 25,000,000 active players (est. 60,000,000 by 2012)
- Over 150 MMOGs in operation
- Market size 7,500,000,000\$/year



Sources: MMOGChart, own research.



Sources: ESA, MPAA, RIAA.

What is an MMOG?

1. Content

graphics, maps,
puzzles, quests +

2. Virtual world simulation

explore, do, learn,
socialize, compete




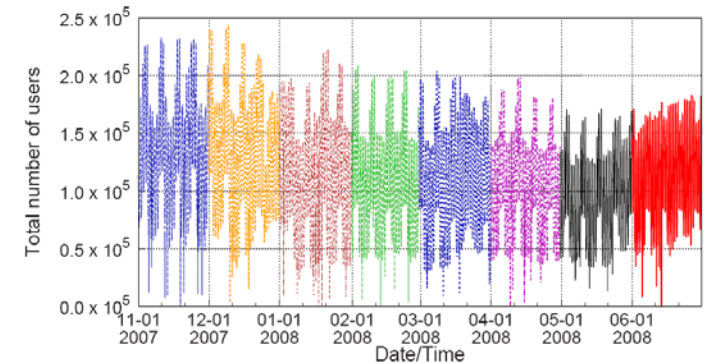
ively Multiplayer Online Games

Continuous Analytics for MMOGs

MMOG Data =
raw and derivative information
about the virtual world

Continuous Analytics for MMOGs =
Analysis of MMOG data s.t.
important events are not lost

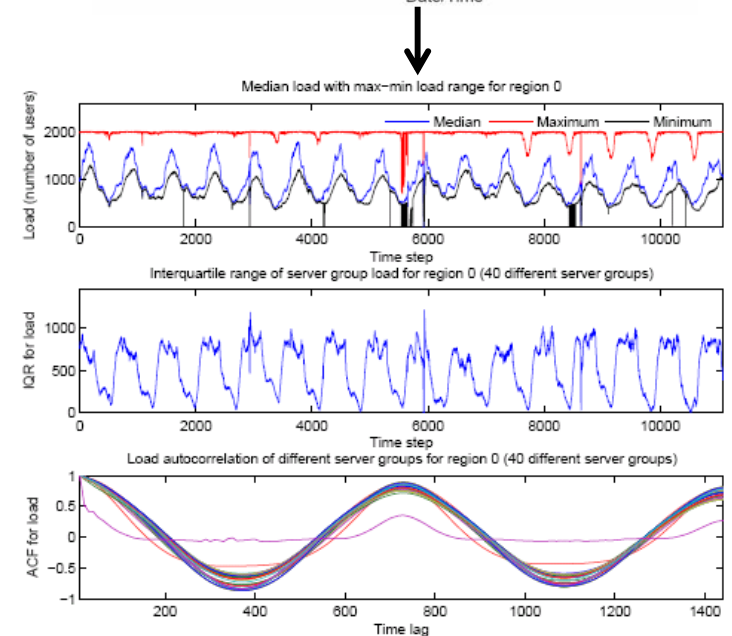

- Data collection
- Data storage
- Data analysis
- Data presentation
- ... at MMOG rate and scale



Continuous Analysis for MMOGs

Main uses by the gaming industry

1. Understand play patterns (decide future investments)
2. Prevent and detect cheating or disastrous game exploits (think MMOG economy reset)
3. Support player communities
4. Broadcasting of gaming events
5. Data for advertisement companies (new revenue stream for MMOGs)



Other Uses for MMOG Data

Social Sciences

- The emergence and performance of ad hoc groups in contemporary society
- Emergent behavior in complex systems

Economy

- Contemporary economic behavior

Psychology

- Games as coping mechanism (minorities)
- Games as cure (agoraphobia)

Biology

- Disease spread models



Last Updated: Tuesday, 21 August 2007, 00:04 GMT 01:04 UK

<http://news.bbc.co.uk/2/hi/health/6951918.stm> BBC

Virtual game is a 'disease model'

An outbreak of a deadly disease in a virtual world can offer insights into real life epidemics, scientists suggest.



Scientists believe the game error could offer a valuable insight

The "corrupted blood" disease spread rapidly within the popular online World of Warcraft game, killing off thousands of players in an uncontrolled plague.

The infection raged, wreaking social chaos, despite quarantine measures.

[Listen](#) [How it reflects reality](#)

The experience provides essential clues to how people behave in such crises, Lancet Infectious Diseases reports.

In the game, there was a real diversity of response from the players to the threat of infection, similar to those seen in real life.

Some acted selflessly, rushing to the aid of other characters even though that meant they risked infection themselves.

“ The players seemed to really feel they were at risk and took the threat of infection seriously ”

Professor Nina Fefferman, from Tufts University School of Medicine

Others fled infected cities in an attempt to save themselves.

Challenges for MMOG Continuous Analytics

Data characteristics

- Large-scale datasets (TB/year in MMOG database, but 10Ks users tracked by largest research studies)
- Dynamic size of relevant data (user activity peaks)

Data producer vs. data user

- Most users cannot get access to main MMOG database
- Users have different cost policies (accuracy vs. cost)

How to support continuous analytics for MMOGs?

Outline

1. Motivation and Problem Statement
- 2. The CAMEO Framework**
3. Experimental Results
4. Conclusion

The CAMEO Framework

Main Targets

Meet the two challenges

- Data characteristics and Data producer vs. data user

Control the accuracy-cost trade-off

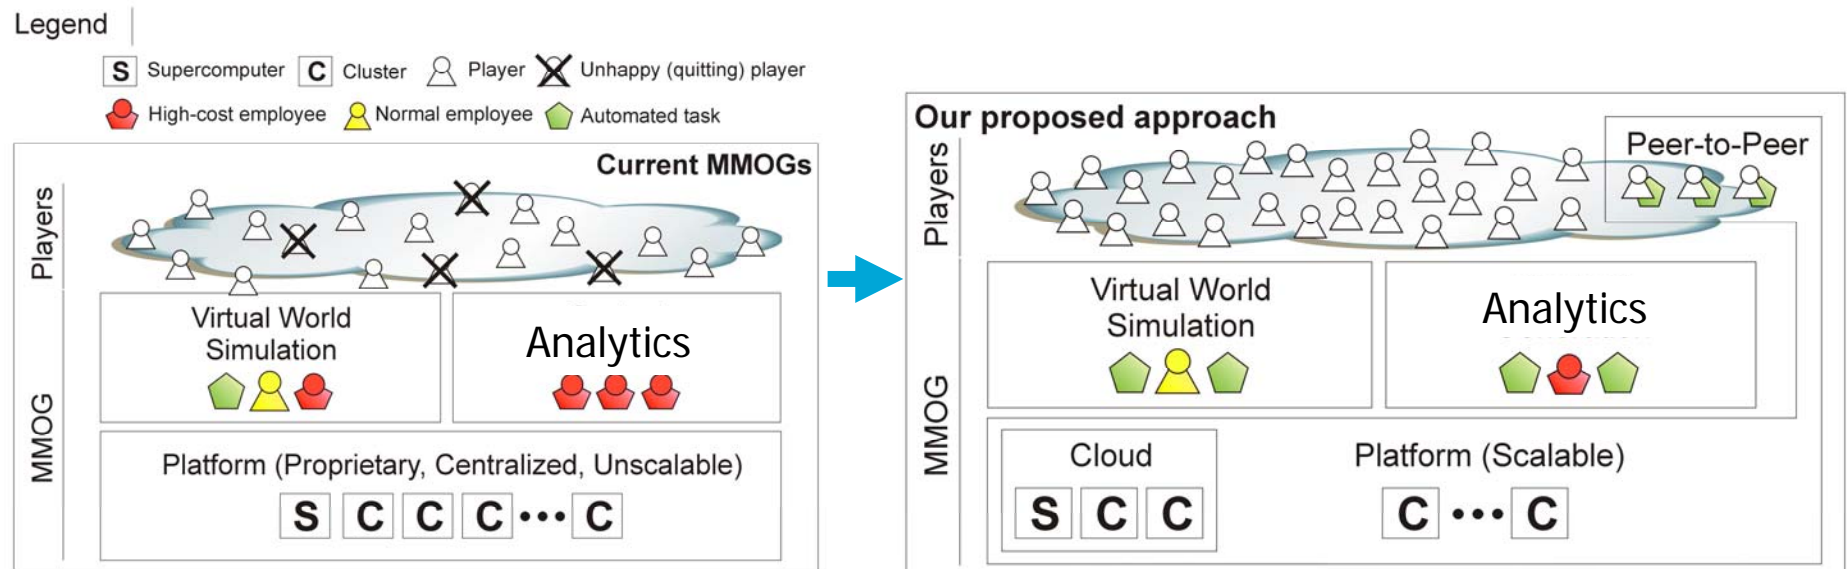
- Get more data = use and pay for more resources
- How much data for results to be relevant?
- How to efficiently manage resources to get enough data?

Auto-tuning to measured system dynamics

- Number of players and related metrics (similar to traditional system provisioning and operation)
- Player activity (very useful for player communities and ads)

The CAMEO Framework

Low upfront costs, efficient and scalable capacity



- Use own resources for predicted load
- **Use cloud (on-demand, paid, guaranteed) resources for excess load**
- Users (peers) may also provide service (future)

Outline

1. Motivation and Problem Statement
2. The CAMEO Framework
- 3. Experimental Results**
4. Conclusion

Experimental Setup

Goal: continuous analytics for RuneScape, the second-most popular MMOG today (3M active players, over 100M accounts)

Technical goal: use Amazon EC2, the largest commercial cloud provider and proponent of open cloud API

- 1 EC2 Compute Unit (ECU) = CPU power of a 1.0-1.2 GHz 2007 Opteron or Xeon proc.
- Pay only used ECUs and bandwidth
- CAMEO currently uses m1.small resources

Resource Type	Cores (ECUs)	RAM [GB]	Architecture [bit]	I/O Performance	Disk [GB]	Cost [\$/h]
m1.small	1 (1)	1.7	32	Med	160	0.1
m1.large	2 (4)	7.5	64	High	850	0.4
m1.xlarge	4 (8)	15.0	64	High	1,690	0.8
c1.medium	2 (5)	1.7	32	Med	350	0.2
c1.xlarge	8 (20)	7.0	64	High	1,690	0.8

Sample MMOG Analytics Results


Skill Level Distribution in RuneScape

- Dataset: **2,899,407 players (largest MMOG msmt.)**
 - 1,817,211 over level 100
 - Max skill 2,280

- Number of mid- and high-level players is significant**




Content generation challenge for MMOGs*



* A. Iosup, POGGI: Puzzle-Based Online Games on Grid Infrastructures EuroPar 2009 (Session Distinguished Papers – Part 1, Thu@10:30am)


Is the Infrastructure Service Guaranteed ?



EC2 Instance acquisition time grows slowly (2' by 2010) Steady delivered performance^{*}

* S. Ostermann, A. Iosup, N.M. Yigitbasi, R. Prodan, T. Fahringer, and D.H.J. Epema An Early Performance Analysis of Cloud Computing Services for Scientific Computing, PDS-2008-006, December 2008

Resource Consumption Control



- Control how many resources are consumed: bursts
- Control how resources are consumed: dynamic vs. steady

The Cost of MMOG Continuous Analytics

Billing Statement: April 1, 2009

Billing Cycle for this Report: March 1 - March 31, 2009

[Expand All](#) | [Collapse All](#)

Rate	Usage	Totals
Amazon Elastic Compute Cloud View/Edit Service		
Amazon EC2 running Linux/UNIX		
\$0.10 per Small Instance (m1.small) instance-hour (or partial hour)	2,097 Hrs	209.70
Amazon EC2 Bandwidth		
\$0.100 per GB Internet Data Transfer - all data transfer into Amazon EC2	611,005 GB	61.10
\$0.170 per GB Internet Data Transfer - first 10 TB / month data transfer out of Amazon EC2	507,121 GB	86.21
Taxes		67.83
Charges due on April 1, 2009+		424.85

- Put a price on MMOG analytics (here, **\$425/month**)
- Trade-off accuracy vs. cost, runtime is constant

Outline

1. Motivation and Problem Statement
2. The CAMEO Framework
3. Experimental Results
- 4. Conclusion**

Conclusion

MMOGs and Cont.Analytics

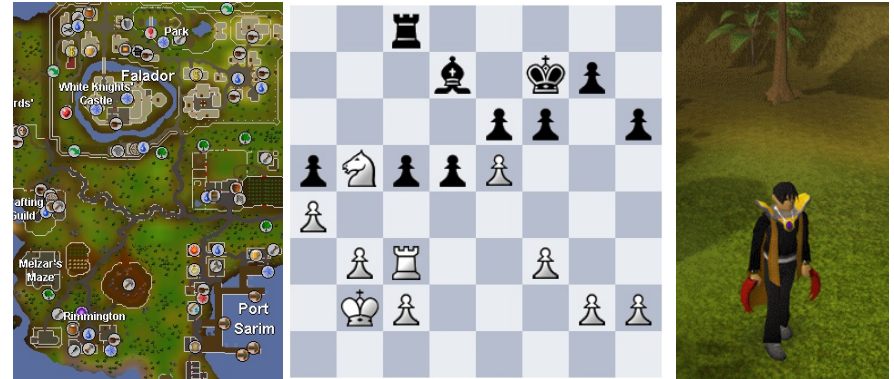
- Million-users, multi-bn. market
- Need for continuous analytics

Current Technology

- Upfront payment
- Cost and scalability problems

Our Approach

- Use clouds as on-demand, paid, guaranteed infrastructure
- Automate most analytics tasks



The CAMEO Framework

- Control accuracy-cost trade-off
- Auto-tune process
- Opens new avenues for research: system and data

The Future of CAMEO

- Full automation
- More clouds and MMOGs
- **Help building an MMOG Workloads Archive**

Thank you for your attention!
Questions? Suggestions? Observations?

Alexandru Iosup

A.Iosup@tudelft.nl

<http://www.pds.ewi.tudelft.nl/~iosup/> (or google "iosup")

Parallel and Distributed Systems Group

Delft University of Technology

