

CAMEO: Enabling Social Networks for Massively Multiplayer Online Games through Continuous Analytics and Cloud Computing

Alexandru Iosup

Delft University of Technology,
The Netherlands

Adrian Lascateu, Nicolae Tapus

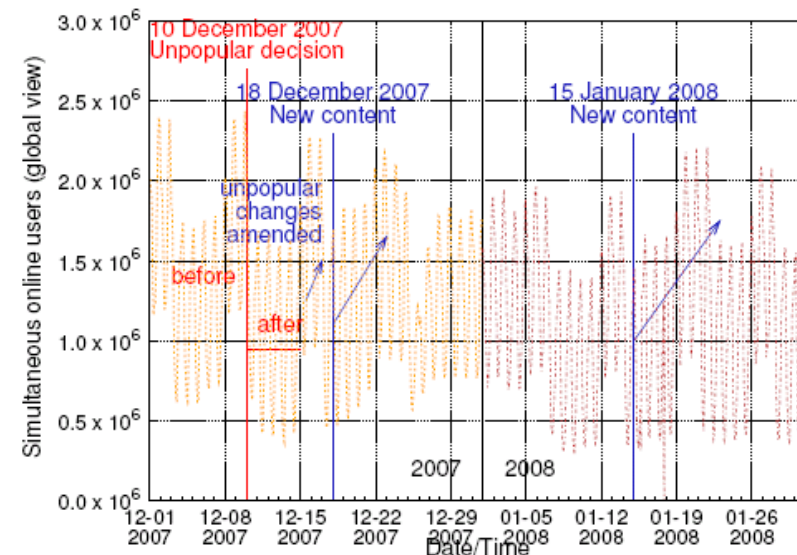

Politehnica U. of Bucharest,
Romania

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

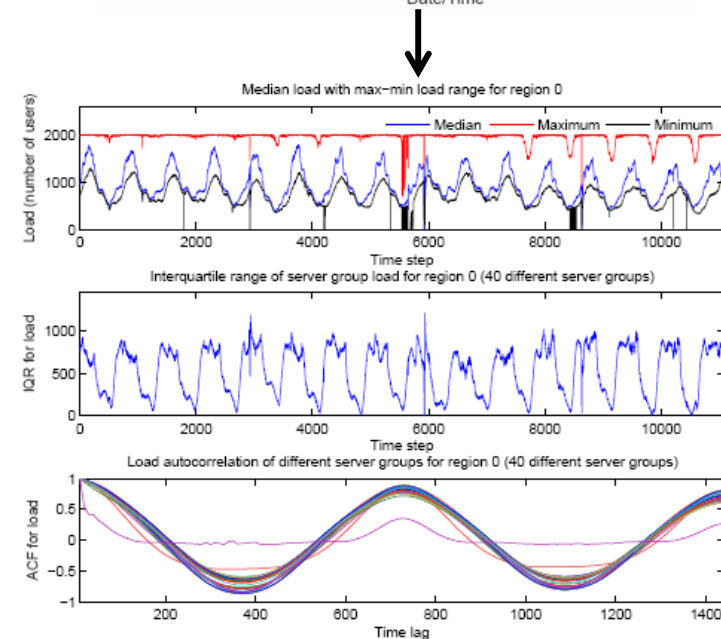

- Data collection
- Data storage
- Data analysis
- Data presentation
- ... at MMOG rate and scale



Continuous Analysis for MMOGs

Main Uses By and For Gamers

1. Support player communities
2. Understand play patterns
(decide future investments)
3. Prevent and detect cheating or disastrous game exploits
(think MMOG economy reset)
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



NU.nl GAMES
Vr. 8 mei 2009. Het laatste nieuws lees je het eerst op NU.nl

World of Warcraft gebruikt bij onderzoek Mexicaanse griep

Uitgegeven: 2 mei 2009 11:45
Laatst gewijzigd: 2 mei 2009 11:45

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.

Analytics for MMOG Social Networks



1. Address community needs

- Specific: casual vs hard-core gamers
- Dynamic: size over time, accuracy vs cost/time/size

2. Using on-demand technology

3. Data management and storage


- TB/year for large games (e.g., EverQuest II)
- Web/Web 2.0 interfaces (RuneScape, NCSoft Dungeon Runners)

4. Performance, scalability, robustness

Background on Cloud Computing



VS



Tropical Cyclone Nargis (NASA, ISSS, 04/29/08)

- “The path to abundance”
- On-demand capacity
- Pay what you use
- Great for web apps (EIP, web crawl, DB ops, I/O)
- “The killer cyclone”
- Not so great performance for compute- or data-intensive applications¹
- Long-term perf. variability²


1- Iosup et al., Performance Analysis of Cloud Computing Services for Many-Tasks Scientific Computing, IEEE Trans. on Par. Distrib. Sys. 2011.

2- Iosup et al., On the Performance Variability of Production Cloud Services, Technical Report PDS-2010-002, [Online] Available: <http://pds.twi.tudelft.nl/reports/2010/PDS-2010-002.pdf>

Outline

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

The CAMEO Framework



1. Address community needs

- Can analyze skill level, experience points, rank
- Can assess community size dynamically

2. Using on-demand technology: Cloud Computing

- Dynamic cloud resource allocation, Elastic IP

3. Data management and storage: Cloud Computing

- Crawl + Store data in the cloud (best performance)

4. Performance, scalability, robustness: Cloud Comp

CAMEO: Analytics Capabilities

1. Various pieces of information


- Skill level, experience points, rank

2. Single and Multi-snapshot analysis

3. Analysis functions already implemented

- Ranking by one or more pieces of information
- Community statistical properties for a piece of information
- Identification of Top-K players in single/multi-snapshot
- Evolution of (Top-)K players
- Evolution of average community skill
- Identification of players with special skill combos

CAMEO: Cloud Resource Management

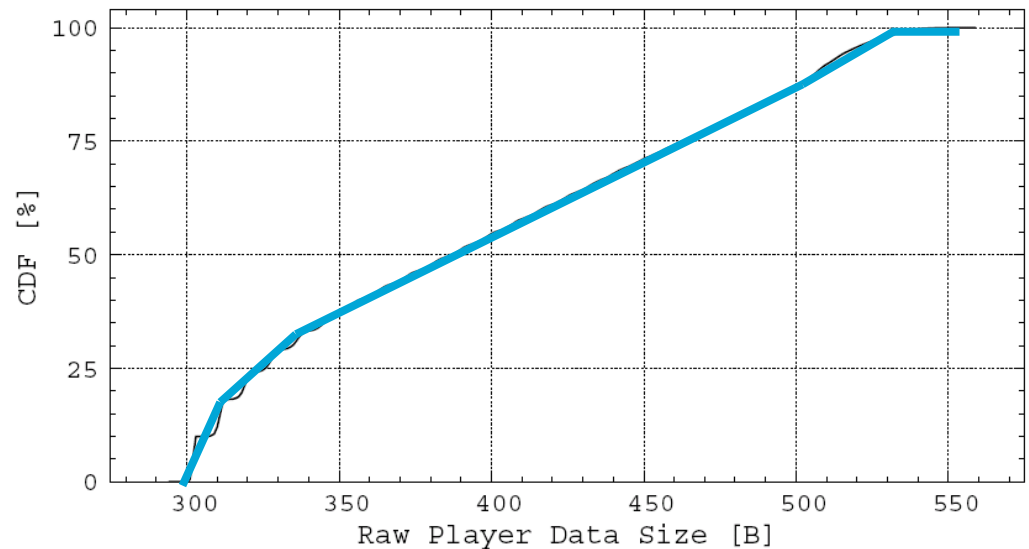


- Snapshot = dataset for a set of players
- More machines = more snapshots per time unit

CAMEO: Exploiting Cloud Features

- **Machines close(r) to server**

- Traffic dominated by small packets (latency)



- **Elastic IP to avoid traffic bans**
(legalese: acting on behalf of real people)

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 (7M active players, over 135M accounts*)

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

- 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

* G. Iddison, Our first look at RuneScape HD, talk at Leipzig GC, Aug 27, 2008.


Sample MMOG Analytics Results (1/2)

Skill Level Distribution in RuneScape

- Dataset 1: **2,899,407** players
 - 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 (Best Paper Award)

Sample MMOG Analytics Results (2/2)

Skill Level Distribution in RuneScape

- Dataset 2: **3,531,478 players (largest MMOG msmt.)**
 - 3,239,089 over level 100
 - Max skill 2,488
- **Distribution changed over time**



The Cost of MMOG Continuous Analytics

Billing Statement: April 1, 2009

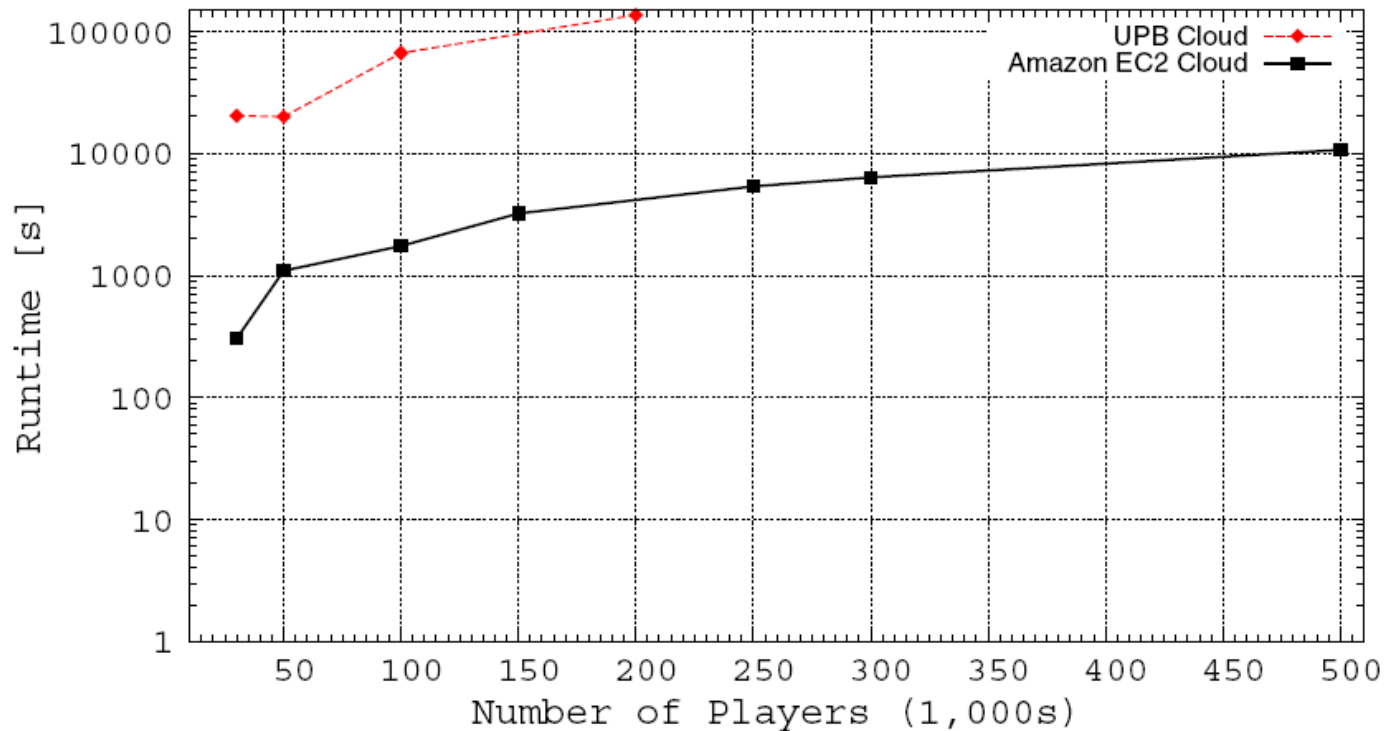
Billing Cycle for this Report: March 1 - March 31, 2009

[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

Performance Results: Why Choosing the Cloud Matters



- Location of machines influences MMOG analytics performance (data acquisition)

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

- Cloud Computing features
- 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

