


# Cloud Computing Support for Massively Social Gaming (Rain for the Thirsty)



**Alexandru Iosup**

**Parallel and Distributed Systems Group  
Delft University of Technology**


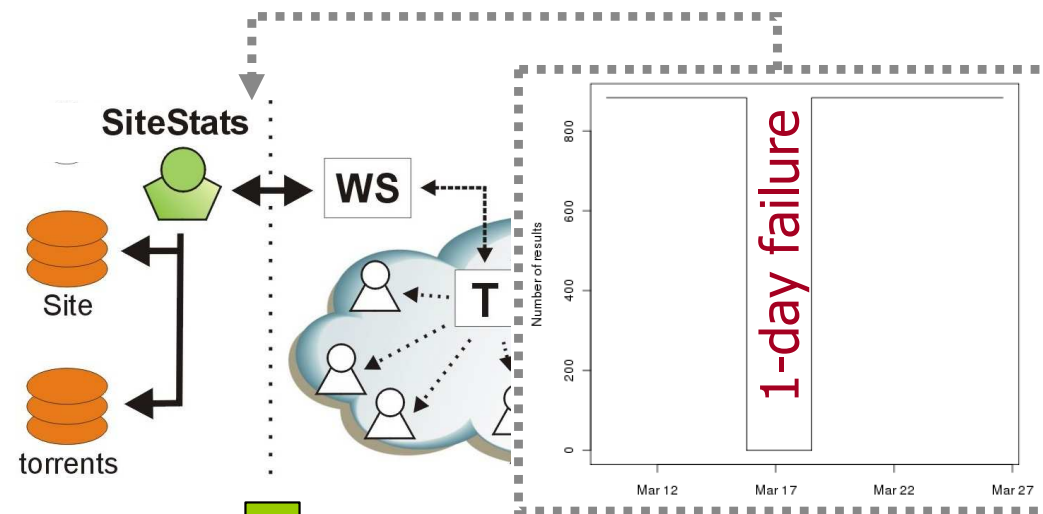
**Our team:** **Undergrad** Adrian Lascateu, Alexandru Dimitriu (UPB, Romania), Saleem Anwar (Vrije Universiteit, the Netherlands), ..., **Grad** Vlad Nae (U. Innsbruck, Austria), Nezh Yigitbasi (TU Delft, the Netherlands), **Staff** Dick Epema, Henk Sips (TU Delft), Thomas Fahringer, Radu Prodan (U. Innsbruck), Nicolae Tapus, Mihaela Balint, Vlad Posea (UPB), Guillaume Pierre (Vrije U.).

**Intermezzo:  
Tips on how clouds can  
help computer science now!**

# How Can Clouds Help? [1/2]

## Heart-Beat-Based Resource Allocation

- BTWorld [LSAP10]
  - Observe *global* status
  - Global = 300+M users
- Data collection failure
  - 1-day, City quarter, campus-wide power-down
- Cloud help
  - Heart-beat monitor
  - No update triggers "ALLOCATE!"
  - (Can pay more)




[LSAP10] wojciechowski, Capota, Pouwelse, and Iosup, BTWorld: Towards Observing the Global BitTorrent File-Sharing Network, ACM LSAP 2010, (accepted).

# How Can Clouds Help? [2/2]

## Dataset Storage for Computer Science

- Critical datasets in computer science
  - Grid Workloads Archive
  - Failure Trace Archive
  - Peer-to-Peer and Game Trace Archives (soon)
  - ... ITA, CRAWDAD, ...
- 1,000s of scientists
  - From theory to practice
- Cloud help
  - Free storage
  - (Can pay for processing)




# Agenda

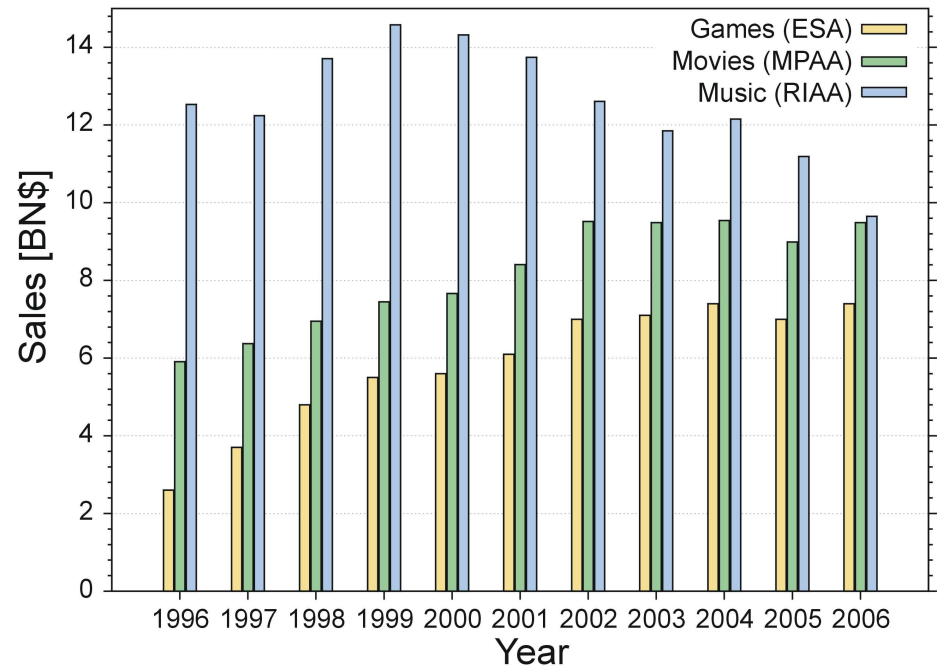
- 1. Background on Massively Social Games**
2. The Mission Slide... Actually, Slides
3. Challenges for Massively Social Games  
(Opportunities for Cloud Computing)
4. The CAMEO Framework for Game Analytics
5. Lessons Learned About Cloud Computing
6. Conclusion

# MSGs are a Popular, Growing Market

- **25,000,000** subscribed players (from 150,000,000+ active)
- **Over 10,000** MSGs in operation
- **Market size 7,500,000,000\$**/year



Sources: MMOGChart, own research.



Sources: ESA, MPAA, RIAA.

# What's in a name?

## Massively Social Gaming

(online) games with massive numbers of players (100K+), for which social interaction helps the gaming experience



**Romeo and Juliet**

### 1. Virtual world

Explore, do, learn, socialize, compete

+

### 2. Content

Graphics, maps, puzzles, quests, culture

+

### 3. Game analytics

Player stats and relationships

# Agenda

1. Background on Massively Social Games
- 2. The Mission Slide... Actually, Slides**
3. Challenges for Massively Social Games  
(Opportunities for Cloud Computing)
4. The CAMEO Framework for Game Analytics
5. Lessons Learned About Cloud Computing
6. Conclusion



# TU Delft and Massively Social Gaming

## Mission

Enable the development, deployment, and operation of Massively Social Gaming for small businesses and amateur game developers

## Strategy

- Be the first to identify MSG opportunities
- **Design and build fully functional, cloud-based MSGs;** uncover fundamental operational laws in the process
- Broaden impact through multi-disciplinary, international team
- Educate academics and academic education

**Launched Sep 2009, various funding sources**

# TU Delft and Cloud Computing

## Mission

Explore the capabilities of cloud computing to support real applications with massive social impact, such as massively social gaming

## Strategy

- **Do not run out of hyperbole** when describing cloud prospects; Apply decade-long grid and peer-to-peer systems expertise
- **Understand the capabilities of the cloud paradigm by designing and building fully functional applications**
- Broaden impact through multi-disciplinary, international team
- Educate academics and academic education

**Launched early-2008, various funding sources**

# Agenda

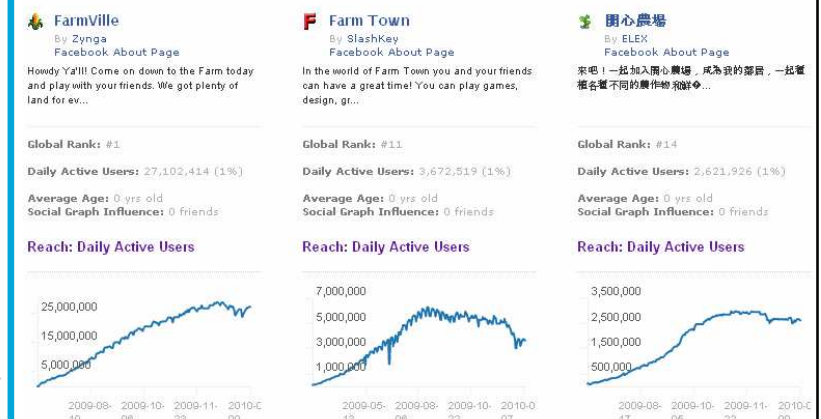

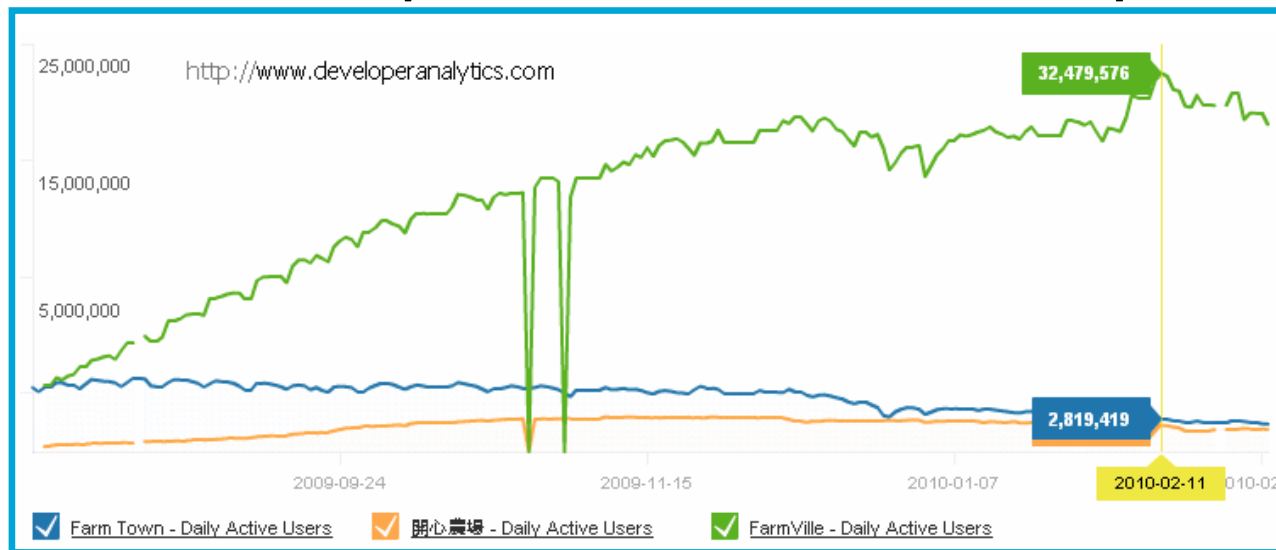
1. Background on Massively Social Games
2. The Mission Slide... Actually, Slides
- 3. Challenges for Massively Social Games  
(Opportunities for Cloud Computing)**
  - 1. Platform Challenge**
  - 2. Content Generation Challenge**
  - 3. Game Analytics Challenge**
4. The CAMEO Framework for Game Analytics
5. Lessons Learned About Cloud Computing
6. Conclusion

# Research Challenge: Solve the Platform Problem of MMOGs

## The Platform Problem of MMOGs

Scaling quickly to millions of players, efficient hosting

- 1M in 4 days, 10M in 2 months
- Up-front and operational costs
- Response time & Scalability



ACM SuperComputing 2008 + IEEE TPDS 2010 (in print),  
IEEE CCGrid 2009, ACM NetGames 2009 + IJMCA (invited)

# Research Challenge: Solve the Content Problem of MMOGs


## The Content Problem of MMOGs

Generating content on time for millions of players

- Player-customized: Balanced, Diverse, Fresh
- Up-front and operational costs
- Response time & Scalability



A. Iosup, POGGI: Puzzle-Based Online Games on Grid Infrastructures EuroPar 2009 **Best Paper Award** + Elsevier CCPE 2010 (accepted)





# Research Challenge: Solve the Analytics Problem of MMOGs

## The Analytics Problem of MMOGs


Analyzing the behavior of millions of players, on-time

- Data mining, data access rights, cost v. accuracy, ...
- Reduce upfront costs
- Low response time & Scalable

ROIA 2009



● worldofwarcraft.com ● mmo-champion.com ● thottbot.com ● wowwiki.com

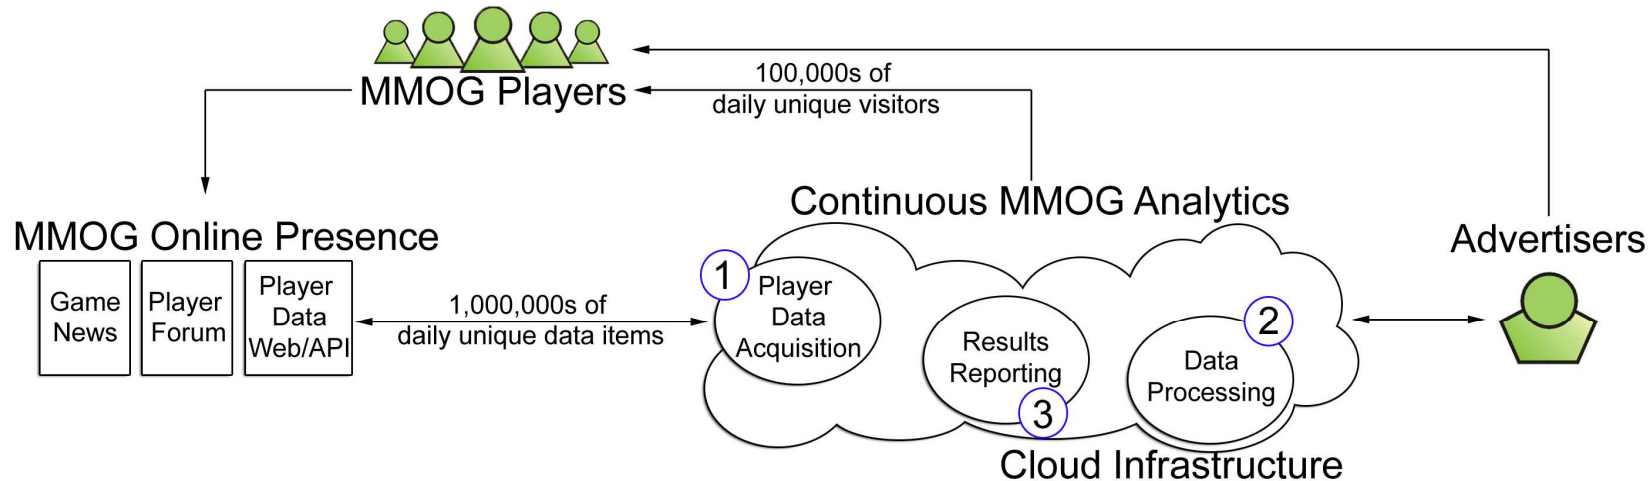


Cloud F

# Agenda

1. Background on Massively Social Games
2. The Mission Slide... Actually, Slides
3. Challenges for Massively Social Games  
(Opportunities for Cloud Computing)
- 4. The CAMEO Framework for Game Analytics**
5. Lessons Learned About Cloud Computing
6. Conclusion

# Third-Party Game Analytics



## Continuous Analytics for MSGs

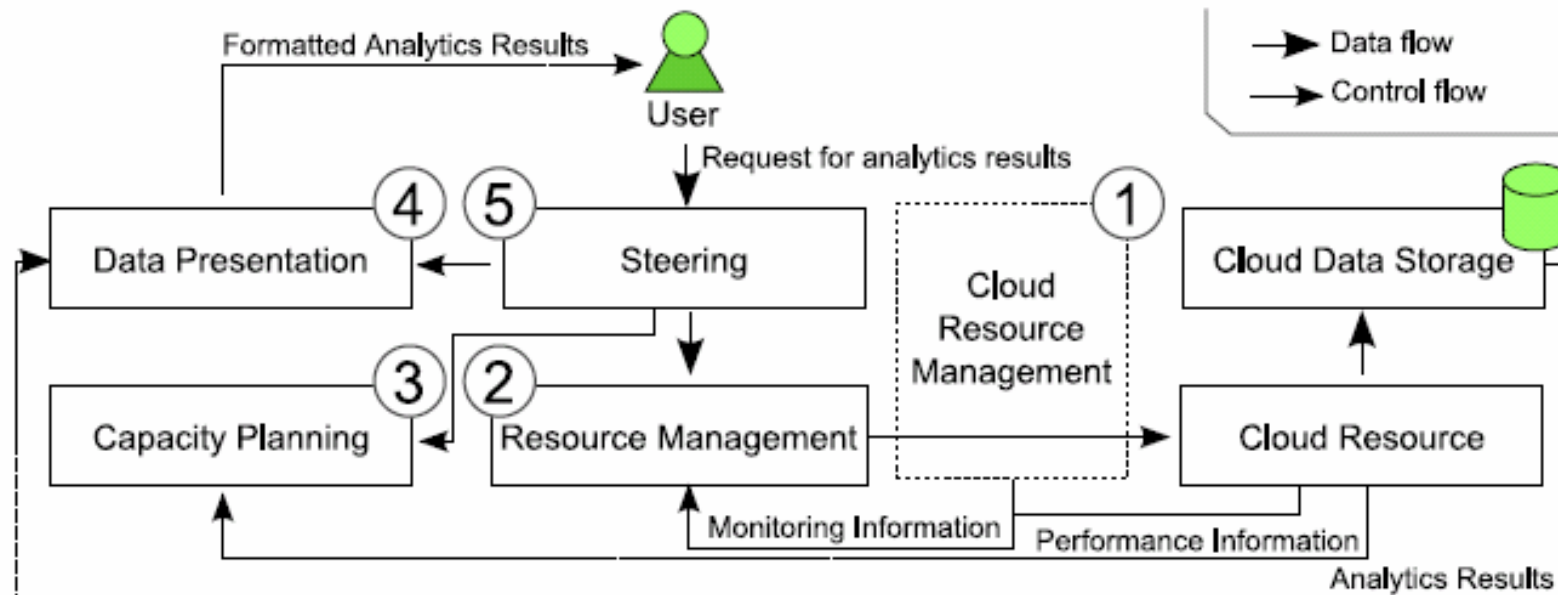
Analysis of raw and derivative MSG data s.t. important events are not lost

- Millions of users for most popular games (WoW, RuneScape, \*Ville)
  - State-of-the-Art third-parties: analytics for 10,000s players
- Dynamic size of relevant data (user activity peaks)
- Users have different requirements (accuracy vs. cost)



# The CAMEO Framework [ROIA09]

## Continuous MSG Analytics on the Cloud



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

[ROIA09] Iosup, CAMEO: Continuous Analytics for Massively Multiplayer Online Games on Cloud Resources. ROIA, Euro-Par 2009 workshops, LNCS 6043, pp. 289--299. Springer, Heidelberg (2010)

# Sample Game Analytics Results


## Skill Level Distribution in RuneScape

- **RuneScape**: 135M+ open accounts (world record)
- Dataset: **3M players (largest measurement, to date)**
  - 1,817,211 over level 100
  - Max skill 2,280


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



**New Content  
Generation Challenge**



# Cost of Continuous RuneScape Analytics

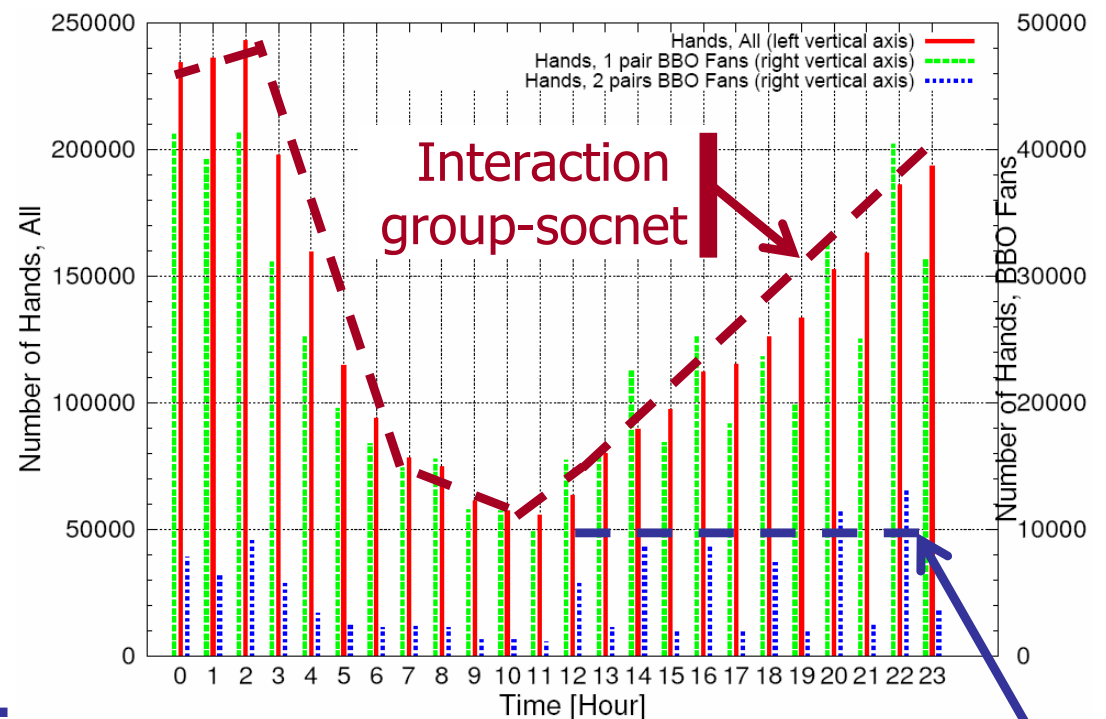
Billing Statement: April 1, 2009			
Billing Cycle for this Report: March 1 - March 31, 2009			
			<a href="#">Expand All</a>   <a href="#">Collapse All</a>
Rate	Usage	Totals	
 <b>Amazon Elastic Compute Cloud</b> <a href="#">View/Edit Service</a>			
<b>Amazon EC2 running Linux/UNIX</b>			
\$0.10 per Small Instance (m1.small) instance-hour (or partial hour)	2,097 Hrs	209.70	
<b>Amazon EC2 Bandwidth</b>			
\$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	
<b>Taxes</b>			<b>67.83</b>
<b>Charges due on April 1, 2009+</b>			<b>424.85</b>

- Put a price on MMOG analytics (here, **\$425/month**, or less than **\$0.00015/user/month**)
- Trade-off accuracy vs. cost, runtime is constant

# Sample Game Analytics Results

## BBO Activity and Social Network

- **Bridge Base Online (BBO):** 1M+ players, top free site
- **Dataset: 100K players**
  - 9K group
  - Social relationships from bridge pairing
- **Large (~10K) online social groups can coordinate**
- **Identified player behavior**  
community builder,  
community member, random player, faithful player



# Agenda


1. Background on Massively Social Games
2. The Mission Slide... Actually, Slides
3. Challenges for Massively Social Games  
(Opportunities for Cloud Computing)
4. The CAMEO Framework for Game Analytics
- 5. Lessons Learned About Cloud Computing**
6. Conclusion

# The Real Cloud



<http://www.flickr.com/photos/dimitrisotiropoulos/4204766418/>

VS



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

- “The path to abundance”
- On-demand capacity
- Cheap for short-term tasks
- Great for web apps (EIP, web crawl, DB ops, I/O)
- “The killer cyclone”
- Not so great performance for scientific applications<sup>1</sup> (compute- or data-intensive)
- Long-term perf. variability<sup>2</sup>

1- Iosup et al., Performance Analysis of Cloud Computing Services for MTC-Based Scientific Computing, (under submission).

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>

# Low Performance for Sci.Comp.

- Evaluated the performance of resources from four production, commercial clouds.
  - GrenchMark for evaluating the performance of cloud resources
  - Four production, commercial IaaS clouds: Amazon Elastic Compute Cloud (EC2), Mosso, Elastic Hosts, and GoGrid.
- **Finding: cloud performance low for sci.comp.**


S. Ostermann, A. Iosup, N. Yigitbasi, R. Prodan, T. Fahringer, and D. Epema, A Performance Analysis of EC2 Cloud Computing Services for Scientific Computing, Cloudcomp 2009, LNICST 34, pp. 115–131, 2010.

N. Yigitbasi, A. Iosup, D. Epema, S. Ostermann: C-Meter: A Framework for Performance Analysis of Computing Clouds. Proc. of CCGRID 2009: 472–477

... more under submission ...

# Cloud Performance Variability

- Performance variability of production cloud services
  - Infrastructure:  
Amazon Web Services
  - Platform:  
Google App Engine



- Year-long performance information for nine services
- **Finding: about half of the cloud services investigated in this work exhibits yearly and daily patterns; impact of performance variability depends on application.**

A. Iosup, N. Yigitbasi, and D. Epema, On the Performance Variability of Production Cloud Services, (under submission).



## MSGs

- Million-user, multi-bn market
- Content, World Sim, Analytics

## Current Technology

- Upfront payment
- Cost and scalability problems
- Makes players unhappy

### Publications Gaming and Clouds

2008: ACM SC, TR Perf

2009: ROIA, CCGrid, NetGames,  
EuroPar (**Best Paper Award**),  
CloudComp, TR variability

2010: IEEE TPDS, Elsevier CCPE

2011: Book Chapter CAMEO

### Graduation Forecast

2010/2011: 1PhD, 2Msc, 4BSc

# Conclusion

## Our Vision

- Scalability & Automation
- Economy of scale with clouds

## Ongoing Work

- Content: POGGI Framework
- Platform: edutain@grid
- Analytics: CAMEO Framework

## The Future

- **Happy players**
- **Happy cloud operators**

# Thank you for your attention!

## Questions? Suggestions? Observations?

More Info:

- <http://www.st.ewi.tudelft.nl/~iosup/research.html>
- [http://www.st.ewi.tudelft.nl/~iosup/research\\_gaming.html](http://www.st.ewi.tudelft.nl/~iosup/research_gaming.html)
- [http://www.st.ewi.tudelft.nl/~iosup/research\\_cloud.html](http://www.st.ewi.tudelft.nl/~iosup/research_cloud.html)

## Alexandru Iosup

[A.Iosup@tudelft.nl](mailto:A.Iosup@tudelft.nl)

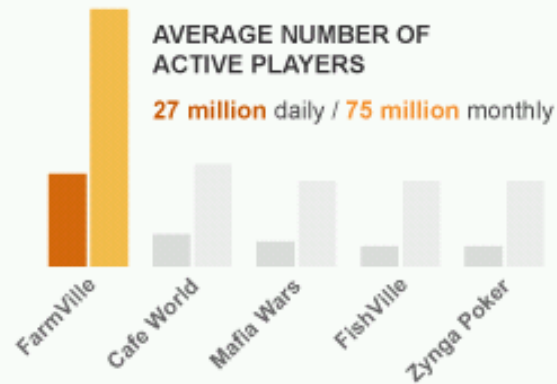
<http://www.pds.ewi.tudelft.nl/~iosup/> (or ~~google~~ <sup>bing</sup> "iosup")

Parallel and Distributed Systems Group  
Delft University of Technology

# Additional Slides

# FarmVille, a Massively Social Game

## ZYNGA GAME FarmVille



### PLAYER PROFILE

N/A average age  
60% female, 40% male




# 118

### THAT'S A LOT

FarmVille boasts 118 million total installs. It has more monthly active users than the population of France.


### MOST POPULAR TIME TO PLAY (EST)



Sources: CNN, Zynga.

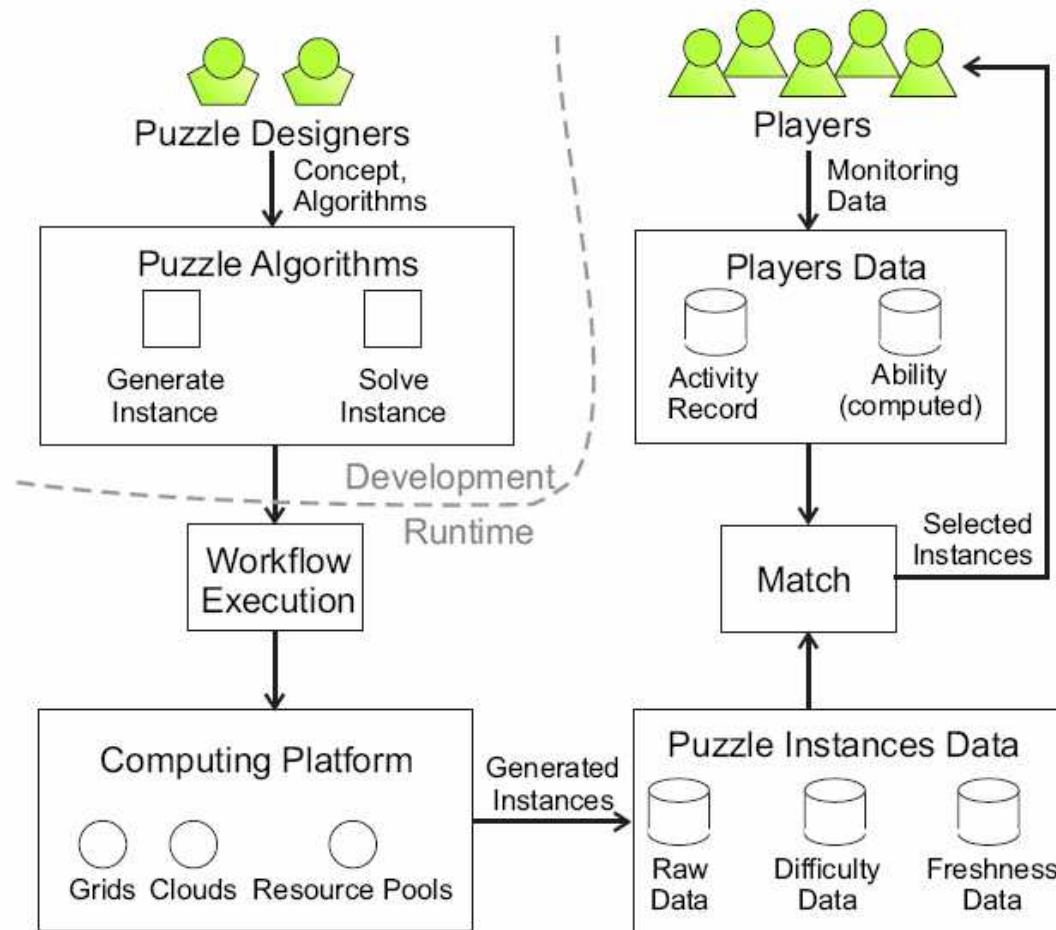
Source: InsideSocialGames.com

# Cloud Resource Consumption Control



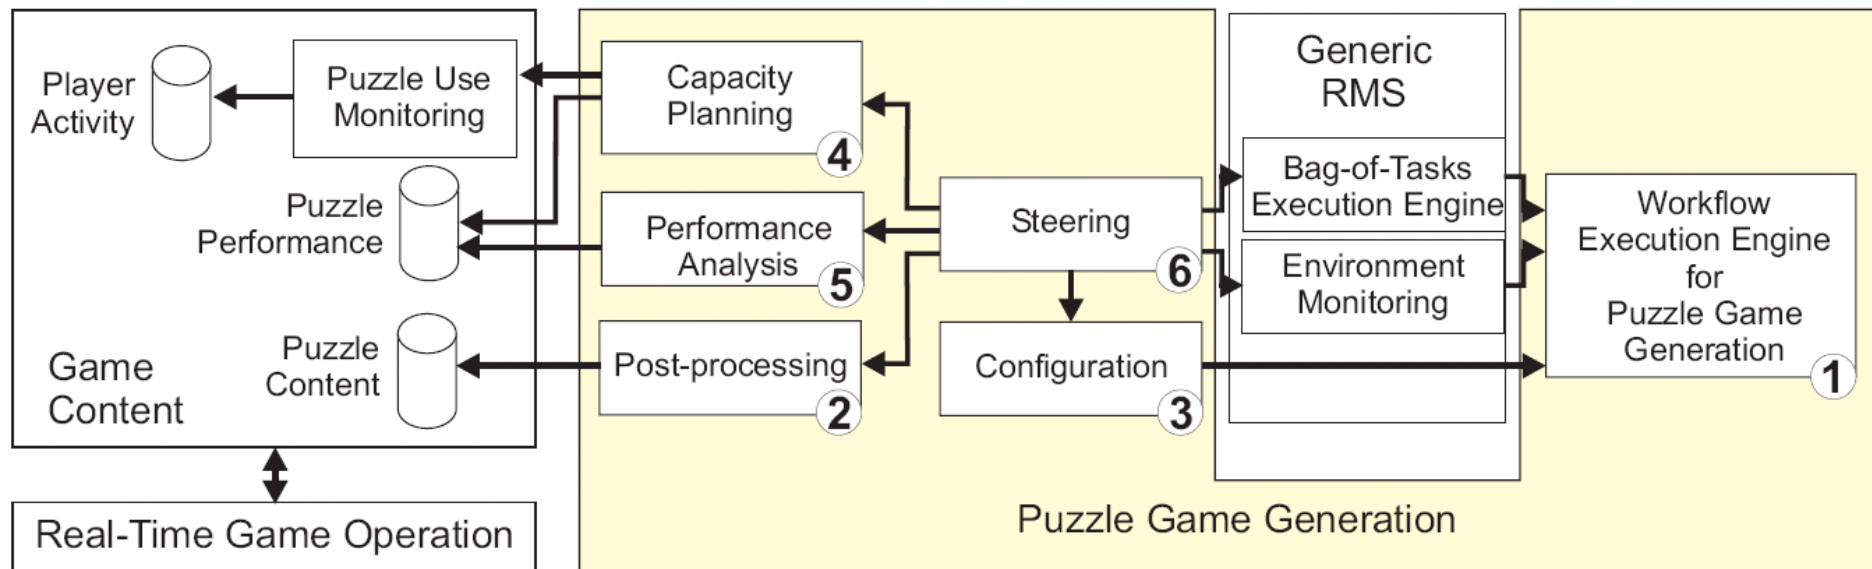
- Control how many resources are consumed: bursts
- Control how resources are consumed: dynamic vs. steady

# The New Content Generation Process



Only the puzzle concept, and the instance generation and solving algorithms, are produced at development time

# The POGGI Framework



## Focus on game content generation on grids

- Use existing middleware
- Control MMOG-specific workload demands and variability (soft guarantees for low response time by pre-generating content)

## ... but do not forget lessons on system design

- Add components for capacity planning and process monitoring

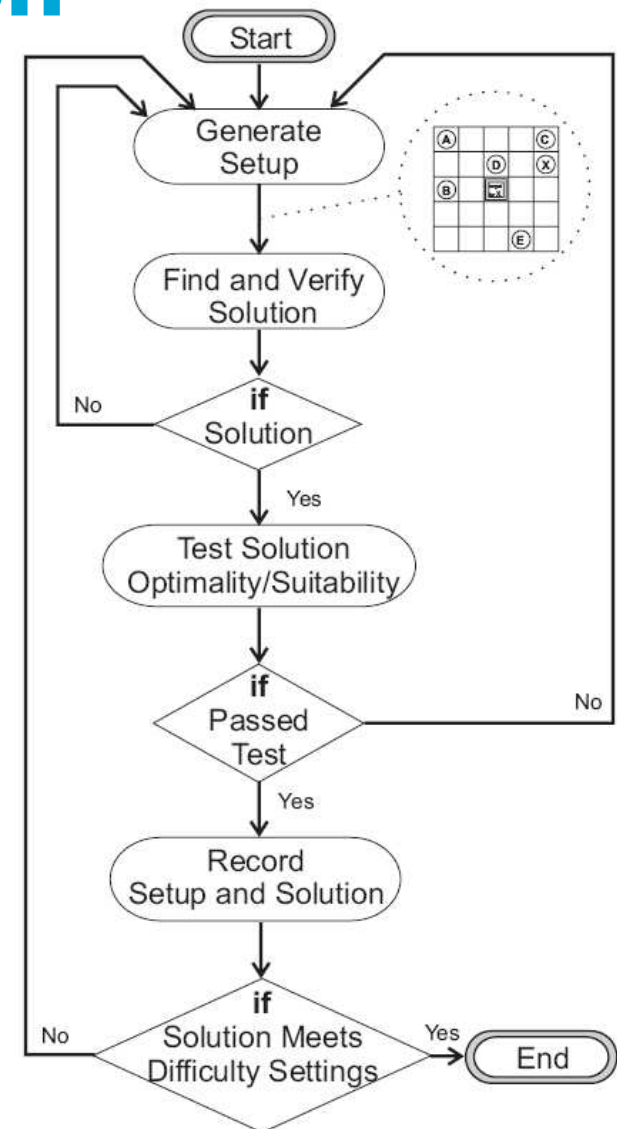
# Workflow Execution Engine for Puzzle Instance Generation

## Generic engine for puzzle generation

- Can plug-in different puzzles
- Can plug-in different solvers
- Can plug-in different policies for instance generation

## Reduce execution overheads

- By-pass RMS (similar to Condor glide-ins, Falcon/Swift, etc., but for WFs instead of tasks)
- Execute on single resource (current implementation, simplicity)





# Puzzle-Specific Considerations


## Generating Player-Customized Content

### Puzzle difficulty

- Solution size
- Solution alternatives
- Variation of moves
- Skill moves

### Player ability

- Keep population statistics and generate enough content for more likely cases
- Match player ability with puzzle difficulty
- Take into account puzzle freshness



X:Right A:Right B:Up X:Up  
(Best solution: 4 moves)

Target: Pins:



B:Up X:Up B:Left C:Down C:Left  
B:Down B:Right B:Down E:Right E:Down  
E:Right B:Up A:Up B:Left C:Down  
C:Right E:Down X:Left E:Left X:Down  
X:Left  
(Best solution: 21 moves)

# Why Not Let Players Generate Puzzles?

## How to control production pipeline?

After all, game developers sell content not technology.

## How to select content?

Ranking problems, diversity problems.

## How to avoid game exploits?

Virtual currency = Real currency

server	faction	100 Gold	Amount	Price	Quick link	Shop profile	Shop rating
Aegwynn	Alliance	\$0.57	5mn Gold	\$15.90	Buy now!	OgPal	3.0/10
	Horde	\$0.45					

Source: mmobux.com, Aug 2009

**User-generated content is clearly an interesting research area, but that's another story.**