Massivizing Online Gaming
Distributed Computing Challenges and High Quality Time

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

* Representing the whole @large team.

What's In a Game?
Defence of the Ancients (DotA)

What's In a Game? MSG, MMO(G), …

Over 250,000,000 active players

Massively Social Gaming =

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

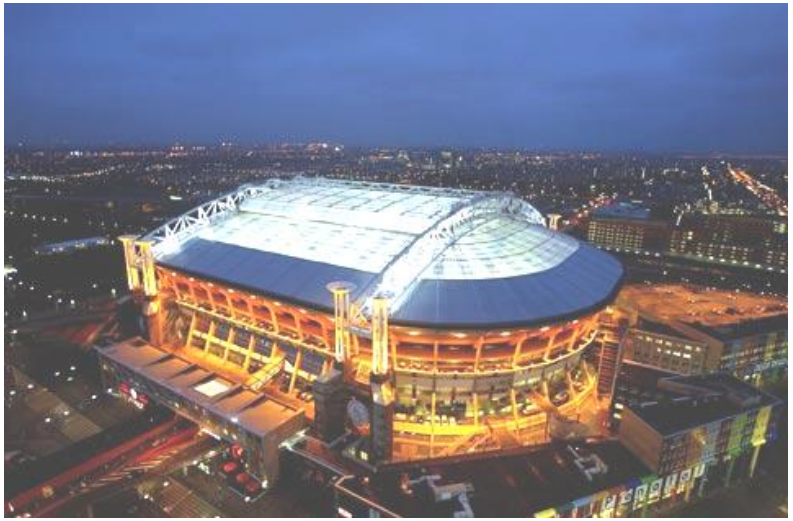

1. **Virtual World Sim**
   Explore, do, learn, socialize, compete
   +
2. **Game Data**
   Player stats and relationships, others
   +
3. **Game Content**
   Graphics, maps, puzzles, quests, culture

Online Game Types

- MMO Role Playing Games (MMORPG)
  - Adventure, role-play: Runscape, World of Warcraft
  - Thousands of players sharing one persistent game session in a huge game world, play sessions ~ hours
  - Latency <1s acceptable
- First Person Shooter (FPS)
  - Action games: Counter Strike, Halo, …
  - Max. 64 players in one session, 5-30’
  - Latency <100ms needed
- Real-Time Strategy (RTS)
  - Economic and battle strategy games: DotA, Starcraft, Age of Empires series
  - Currently few players, hundreds of objects, 30-45’
  - Latency <350ms needed

Why Social Gaming?

Massivizing Social Gaming = Rich Challenge (of Distributed Computing)

Online Gaming used to be art, may now be computing

Online Gaming used to be multimedia, is now DC

Online Gaming used to be communication, is now *all* DC

Online Gaming used to be v-worlds, is now *many apps*

Agenda

1. What's in a Name?
2. **Three Current Challenges**
   1. **Platform Scalability Challenge**
   2. Gaming Analytics Challenge
   3. Content Generation Challenge
3. Conclusion

The Grand Challenge: Massivizing Online Games!

Scaling In and Out for Online Games

(Platform Scalability Challenge)
Build efficient platform for massive scalability

- Close to players
- No upfront costs, no maintenance—PaaS + middleware for Gaming
- Compute platforms: multi-cores, GPUs, clusters, all-in-one!
- Scaling in and scaling out mechanisms and scheduling policies
- Performance guarantees
- Hybrid deployment model
- Code for various compute platforms—platform profiling
- Load prediction miscalculation costs real money
- What are the services?
- Vendor lock-in?
- *My* data

Challenges:World of Warcraft, a Traditional HPC User?!

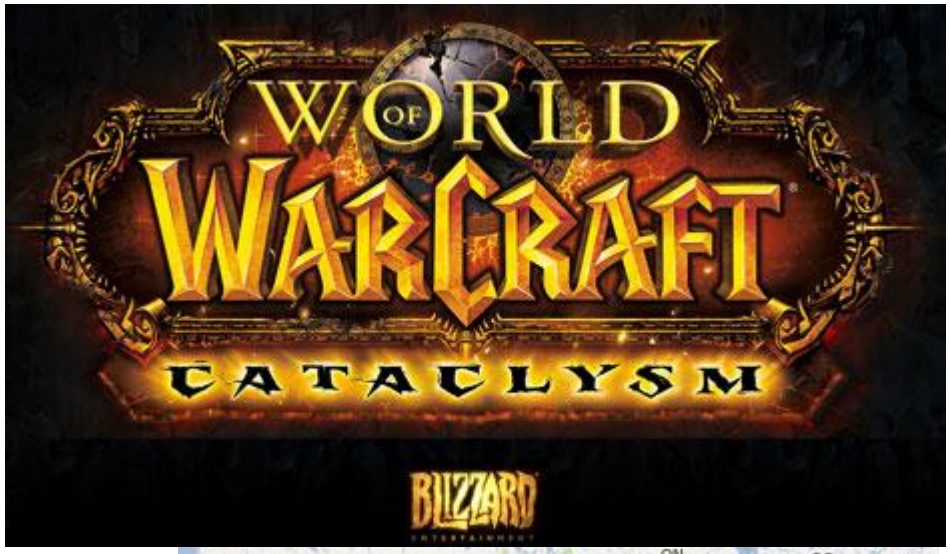

- 10 data centers
- 13,250 server blades, 75,000+ cores
- 1.3PB storage
- 68 sysadmins (1/1,000 cores)

http://www.datacenterknowledge.com/archives/2009/11/25/wows-back-end-10-data-centers-75000-cores/

Cloud-based hosting model

- Using data centers for dynamic resource allocation

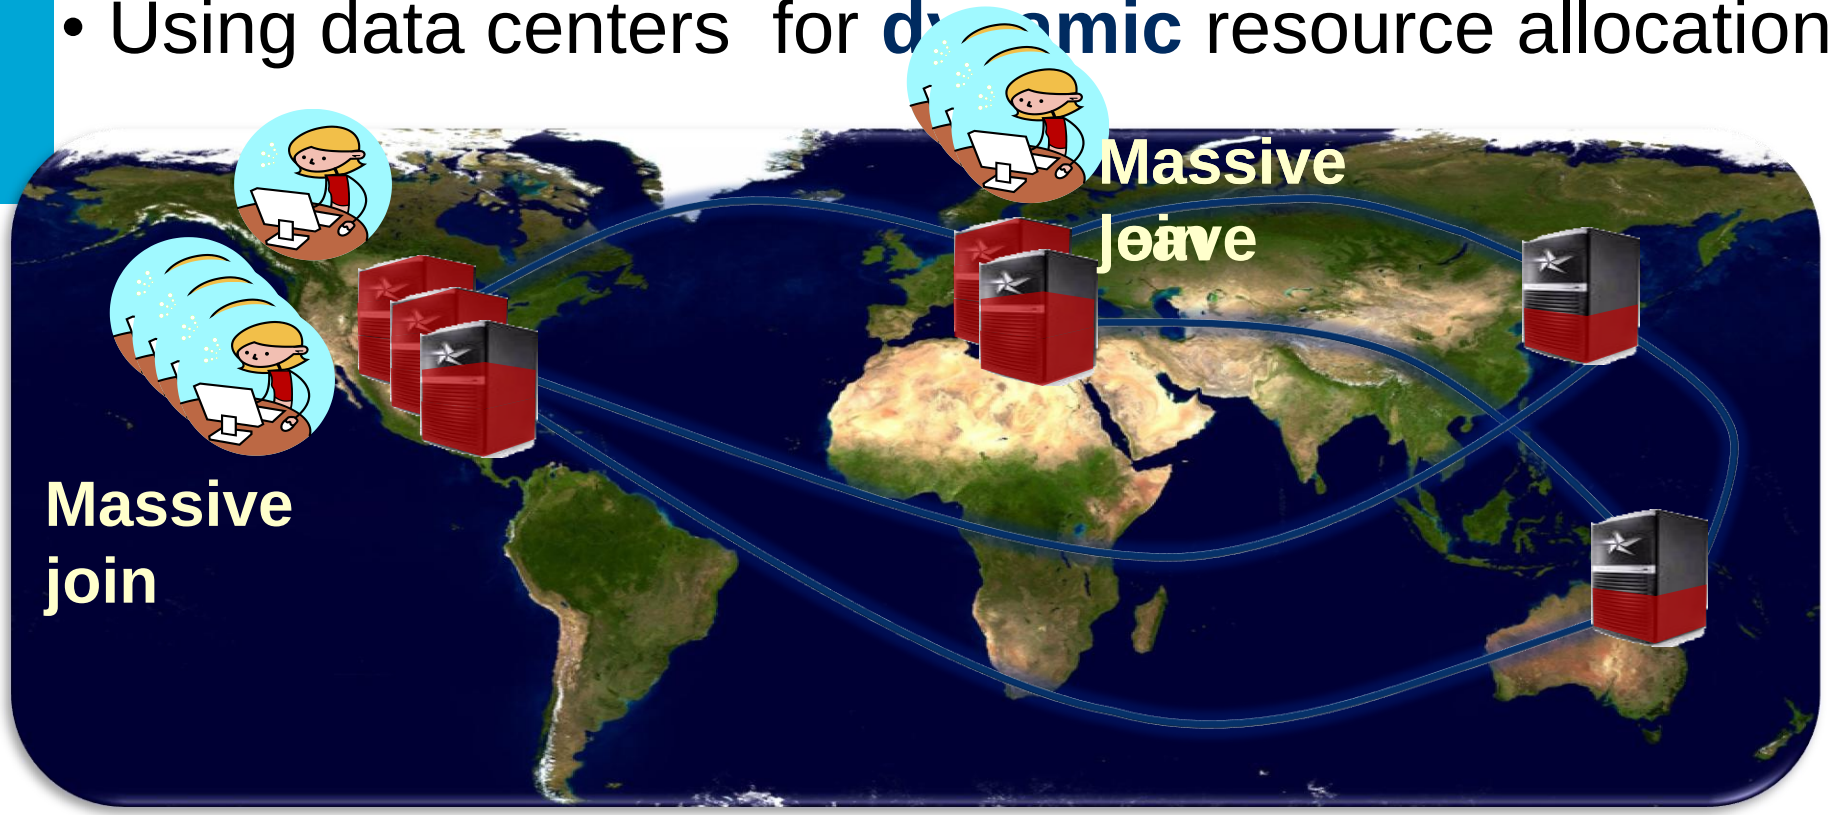

- Main advantages:
  1. Significantly lower over-provisioning
  2. Efficient coverage of the world is possible

[Source: Nae et al., SC 2008]

Vlad Nae, Alexandru Iosup, Radu Prodan: Dynamic Resource Provisioning in Massively Multiplayer Online Games. IEEE Trans. Parallel Distrib. Syst. 22(3): 380-395 (2011)

Resource Provisioning and Allocation
Static vs. Dynamic Provisioning

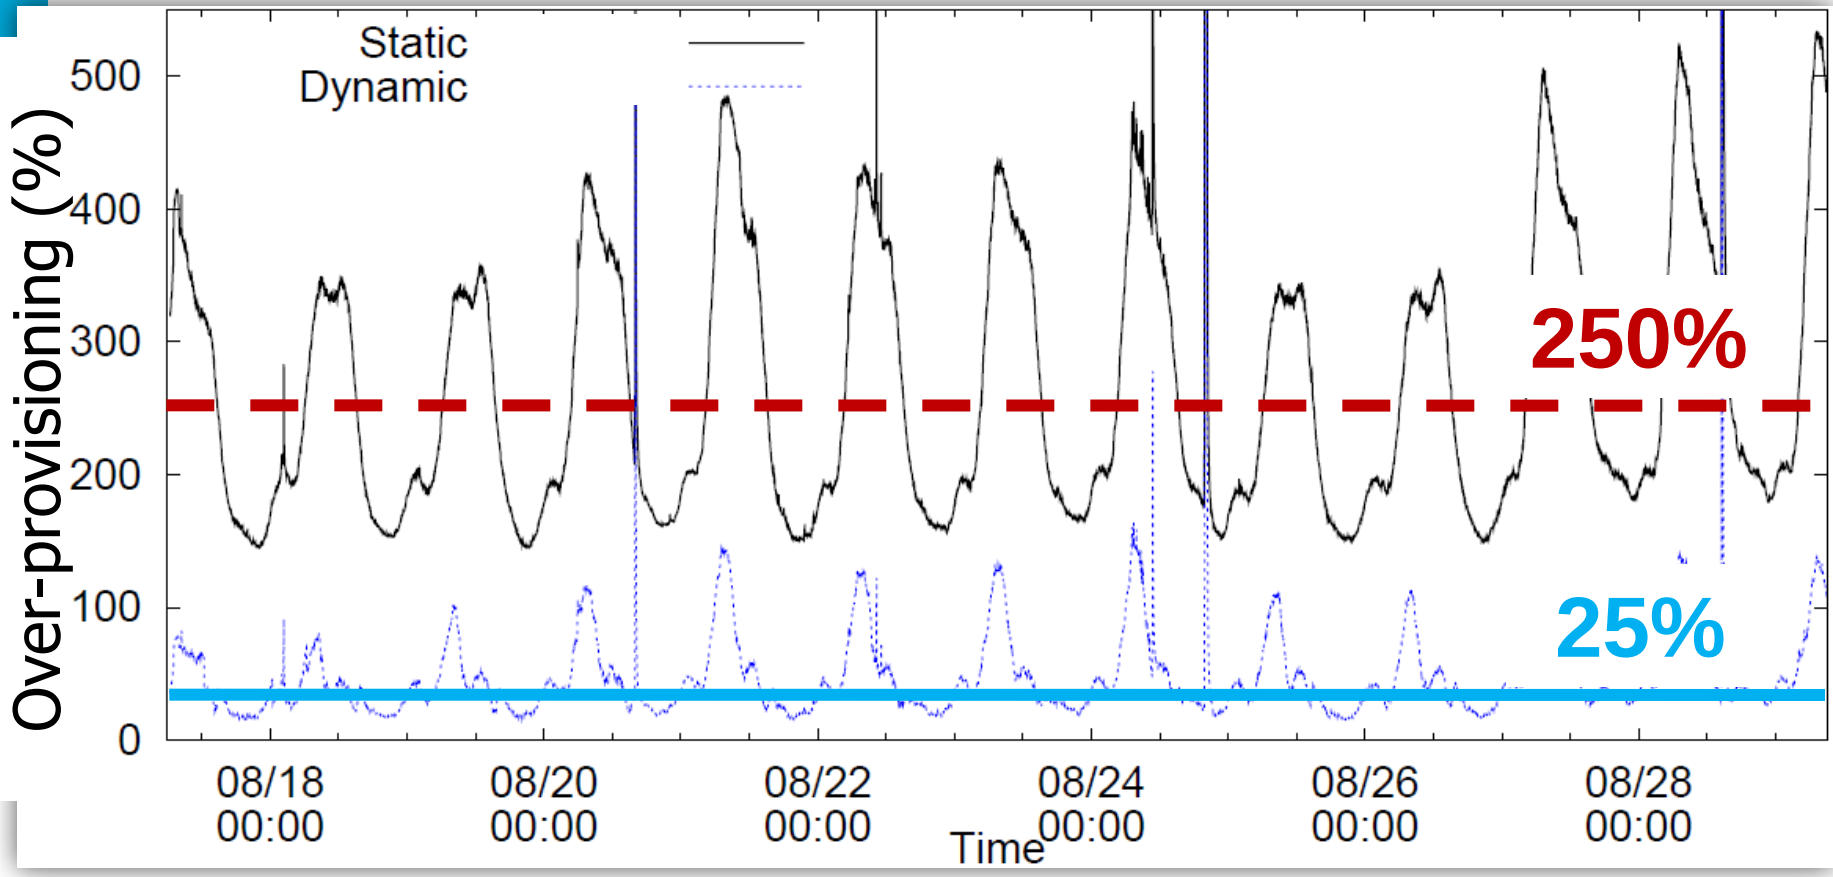

Vlad Nae, Alexandru Iosup, Radu Prodan: Dynamic Resource Provisioning in Massively Multiplayer Online Games. IEEE Trans. Parallel Distrib. Syst. 22(3): 380-395 (2011)

Remember RTS Games?

- Players control tens up to hundreds of units.
- Players need to take decisions in real-time.

> Our approach for consistency: Area of Simulation

Area of Simulation: Core Idea

- Partition the game into multiple areas (rectangular)
- Each player pays attention to different areas + attention level
- Depending on attention level and machine performance, the player will receive different types of information (**commands** or **state**) about the game world
  - AoS: Area of Simulation = high-attention area, full simulation based on input commands (CPU-intensive)
  - AoU: Area of Update = low-attention area, receives state (Net)
  - NUA: No update area
- Each player can have multiple AoS and AoU

S. Shen, A. Iosup, D. H. J. Epema, and S.-Y. Hu. Area of Simulation: Mechanism and Architecture for Multi-Avatar Virtual Environments, ACM Trans. Multimedia Comput. Comm. App. 2015.

Experimental results

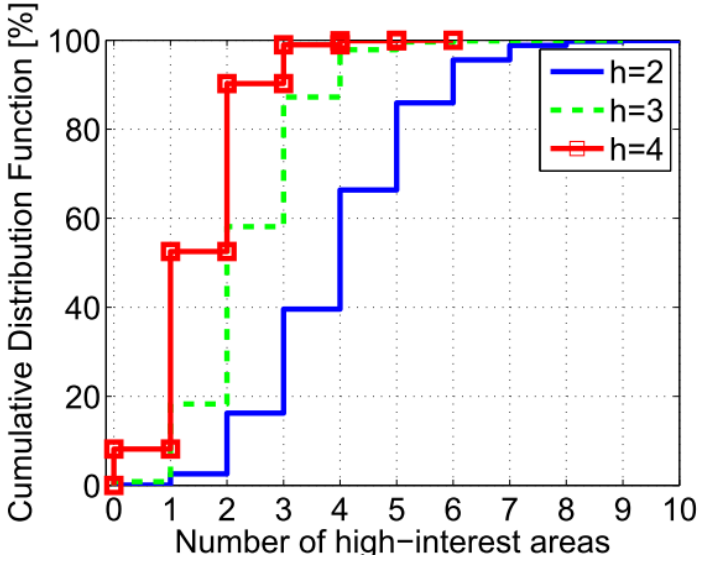

h=high:normal interest

- **Area of Simulation is needed**
  - N vs 1 Area of Interest (traditional)
- Simulator and real-world prototype RTS game
  - Prototype about 20k lines of C++ code
    - Based on an open source game (~6k lines)
  - Evaluated in two Cloud platforms
- Our AoS-based method leads under most circumstances to
  - Higher scalability Up to 200 players and **10,000 battle units**

    vs. state-of-the-art: unplayable at 1,000-2,000 battle units + crashes
  - **Lower CPU** consumption than pure event-based method (RTS) and **lower network** consumption than pure update-based method (RPG)

S. Shen, A. Iosup, D. H. J. Epema, and S.-Y. Hu. Area of Simulation: Mechanism and Architecture for Multi-Avatar Virtual Environments, ACM Trans. Multimedia Comput. Comm. App. 2015.

Mobile/Remote Gaming and the SuperServer ("Cloud" Gaming?!)

(Platform Challenge, concerning thin clients)
Support MSGs on mobile/remote devices

- Mobiles everywhere (2bn+ users)
- Gaming industry for mobiles is the new Growing Market
- SuperServer to stream game video
- SuperServer to generate content for low-capability devices?
- Battery for 3D/ Networked games?
- Where is my server? (Ad-hoc mobile gaming networks?)
- Security, anti-cheating

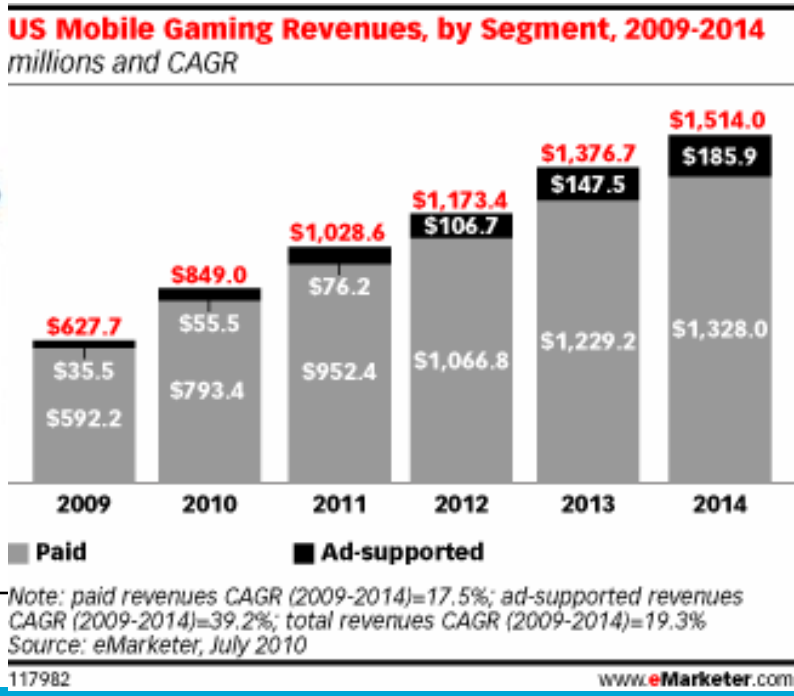

Computational Model for the Server

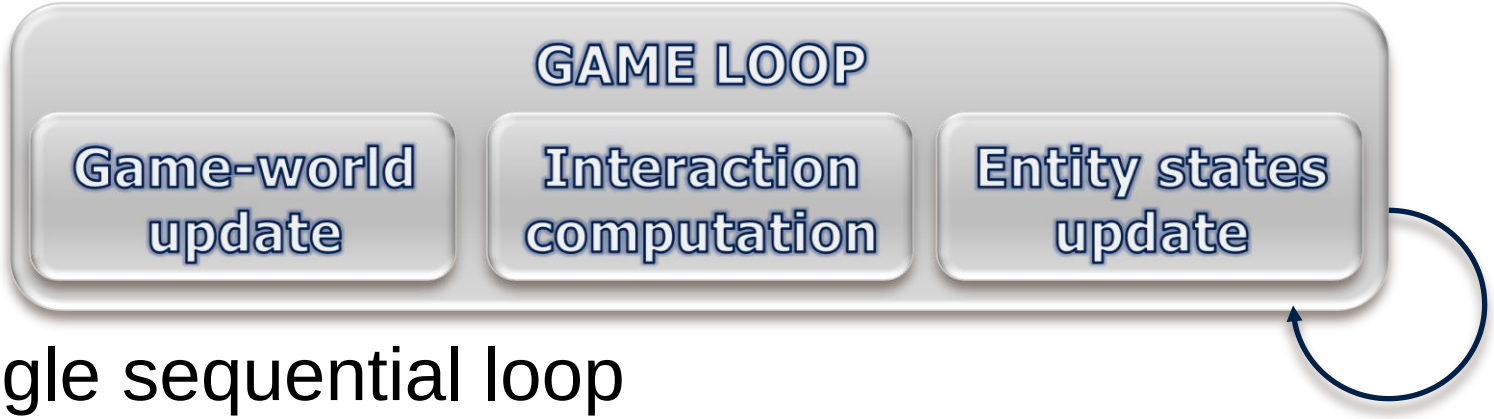

- Single sequential loop
- 3 steps in each loop:
  1. Game-world state update
  2. Entity interaction computation (dominant for MSGs)
  3. Entity state updates (expensive part: rendering)
     - offload if possible and needed

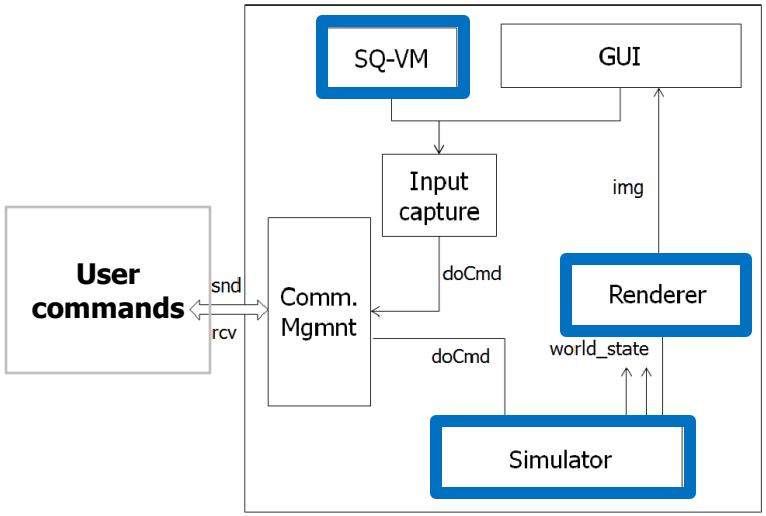

Exploratory Space for Offloading

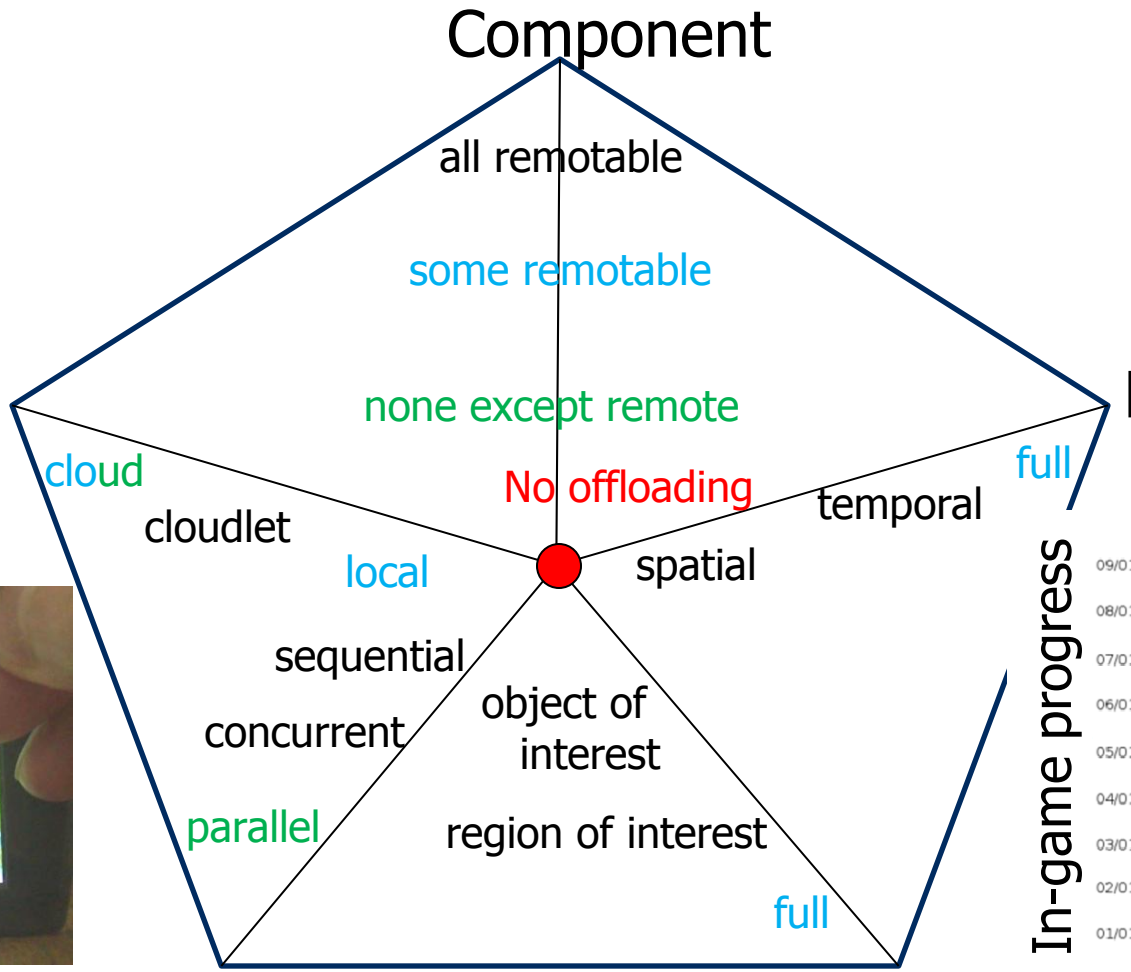

Cloud gaming / SuperServer: OnLive, Gaikai, Sony Online, Microsoft Live, etc.

source: youtube.com

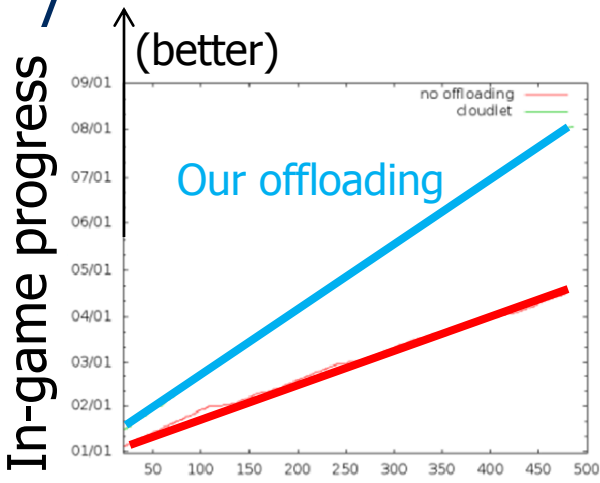

A.-C. Olteanu, N. Tapus, A. Iosup: Extending the Capabilities of Mobile Devices for Online Social Applications through Cloud Offloading. CCGRID 2013: 160-163.

Agenda

1. What's in a Name?
2. **Three Current Challenges**
   1. Platform Scalability Challenge
   2. **Gaming Analytics Challenge**
   3. Content Generation Challenge
3. Conclusion

The Grand Challenge: Massivizing Online Games!

Social Everything! So Analytics

- **Social Network=undirected graph, relationship=edge**
- **Community=sub-graph, density of edges between its nodes higher than density of edges outside sub-graph**

(Gaming Analytics Challenge)
Analytics at Massive Scale

Improve gaming experience

- Ranking / Rating
- Matchmaking / Recommendations
- Play Style/Tutoring

Self-Organizing Gaming Communities

- Player Behavior

Challenges: Bungie, An Analytics Company Serving Petabytes?!

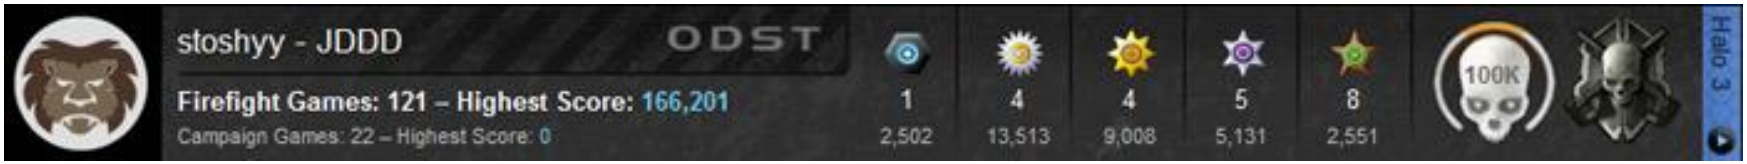

- Halo is one of the many successful game franchises
- Halo 3 players got, in 1.4PB/year
  - Detailed player profiles
  - Detailed usage stats
  - Ranking

- CERN produces ~15PB/year (10x larger)
  - (Not) faster than the speed of light, the Higgs boson (?)

The CAMEO Framework

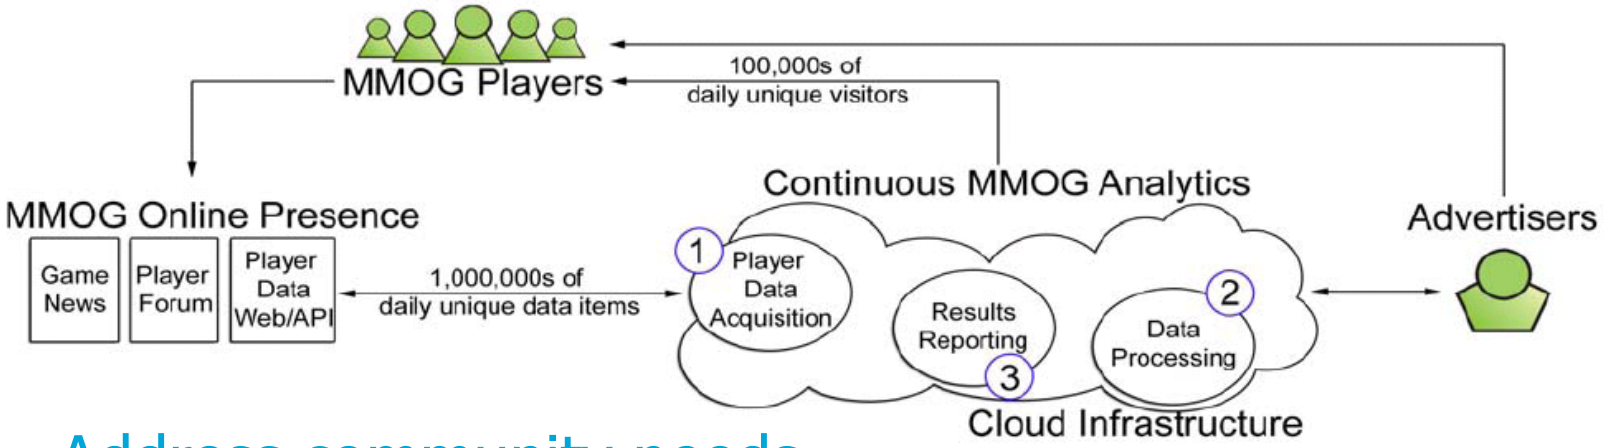

1. Address community needs
   - Can analyze skill level, experience points, rank
   - Can assess community size dynamically
2. Using on-demand technology: Cloud Comp.
   - Dynamic cloud resource allocation, Elastic IP
3. Data management and storage: Cloud Comp.
   - Crawl + Store data in the cloud (best performance)
4. Performance, scalability, robustness: Cloud Comp.

A. Iosup, CAMEO: Continuous Analytics for Massively Multiplayer Online Games on Cloud Resources. ROIA, Euro-Par 2009 Workshops, LNCS 6043, (2010)

@large: Sample Analytics Results

Analysis of Meta-Gaming Network

"When you play a number of games, not as ends unto themselves but as parts of a larger game, you are participating in a metagame." (Dr. Richard Garfield, 2000)

- XFire: since 2008 (3+ years), 500K of 20M players

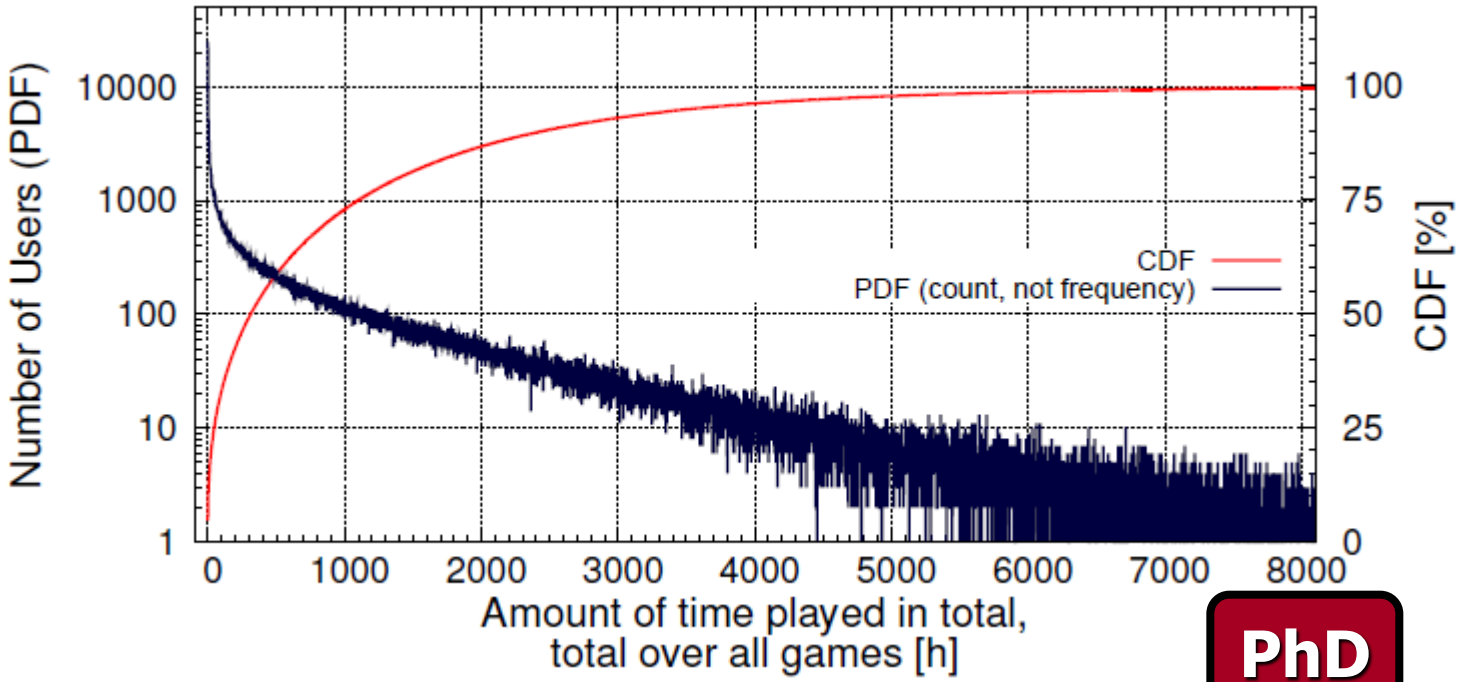

PhD

S. Shen, and A. Iosup, The XFire Online Meta-Gaming Network: Observation and High-Level Analysis, MMVE 2011

Interaction Graphs: From Game Instances to Social Ties

- How to map the relationships in matches to graphs?
  - A dataset D is mapped to a graph G:
  - Players are mapped to nodes
  - Relationships between players are mapped to edges

- We use six different mappings and various thresholds:
  - **SM:** two players occur more than n times in the same match
  - **SS:** two players occur more than n times on the same side
  - **OS:** two players occur more than n times on opposing sides
  - **ML:** two players have lost more than n matches together
  - **MW:** two players have won more than n matches together
  - **PP:** a directed version of the mappings above. A link exists if a player has played more than n percent of his matches together

L. Jia, S. Shen, R van de Bovenkamp, A. Iosup, F.A. Kuipers, D.H.J. Epema. Socializing by Gaming: Revealing Social Relationships in Multiplayer Online Games, ACM Trans. Knowl. Disc. Data. 2015.

From Players to Graph Clusters and Back (to Matchmaking)

- Players→Graph→Clusters (thresholding)
- Bonding score—enjoyment increases when playing together
  - +1/player if clustered players in same game

Team 1

Player	Cluster
A	1
B	2
C	1
D	3
E	4

Team 2

Player	Cluster
F	2
G	5
H	3
I	6
J	3

Bonding score: 0 3 5 7

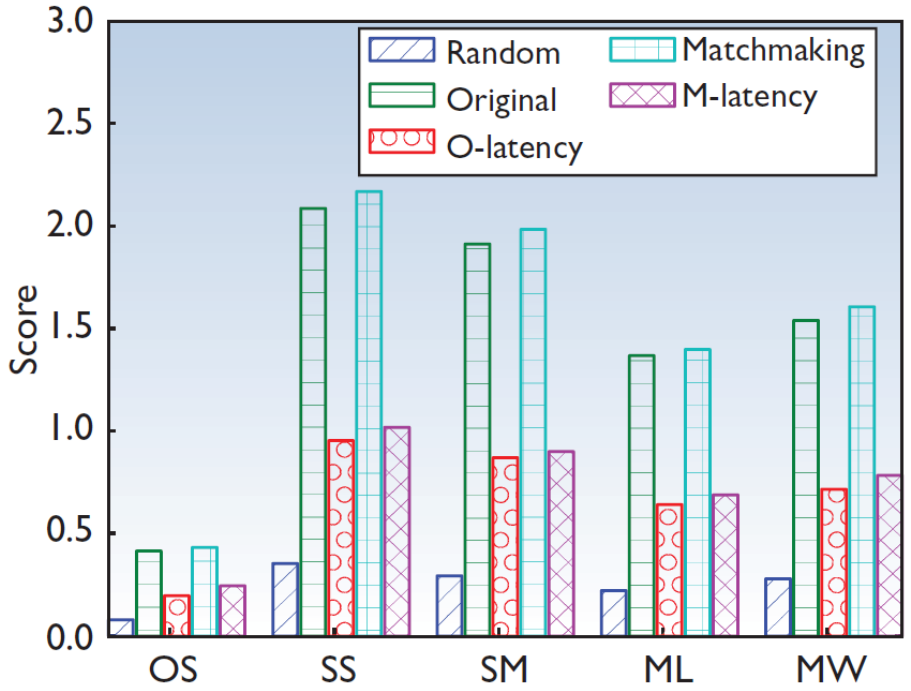

L. Jia, S. Shen, R van de Bovenkamp, A. Iosup, F.A. Kuipers, D.H.J. Epema. Socializing by Gaming: Revealing Social Relationships in Multiplayer Online Games, ACM Trans. Knowl. Disc. Data. 2015.

Agenda

1. What's in a Name?
2. **Three Current Challenges**
   1. Platform Scalability Challenge
   2. Gaming Analytics Challenge
   3. **Content Generation Challenge**
3. Conclusion

The Grand Challenge: Massivizing Online Games!

Content, Content, Content

(Content Challenge)
Produce and distribute content for 1BN players

- Game Analytics → Game statistics
- Content distribution/ Streaming content
- Crowdsourcing
- Storification
- Auto-generated game content
- Adaptive game content at scale

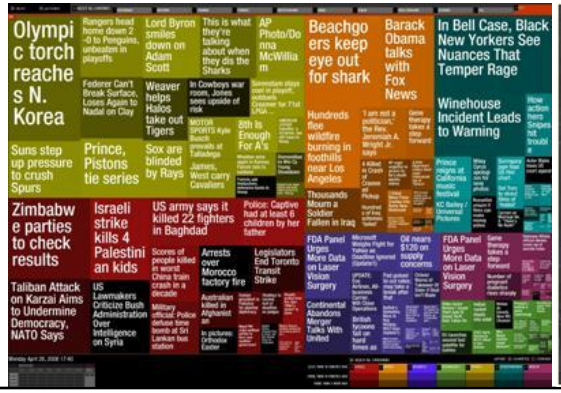

(Procedural) Game Content (Generation)

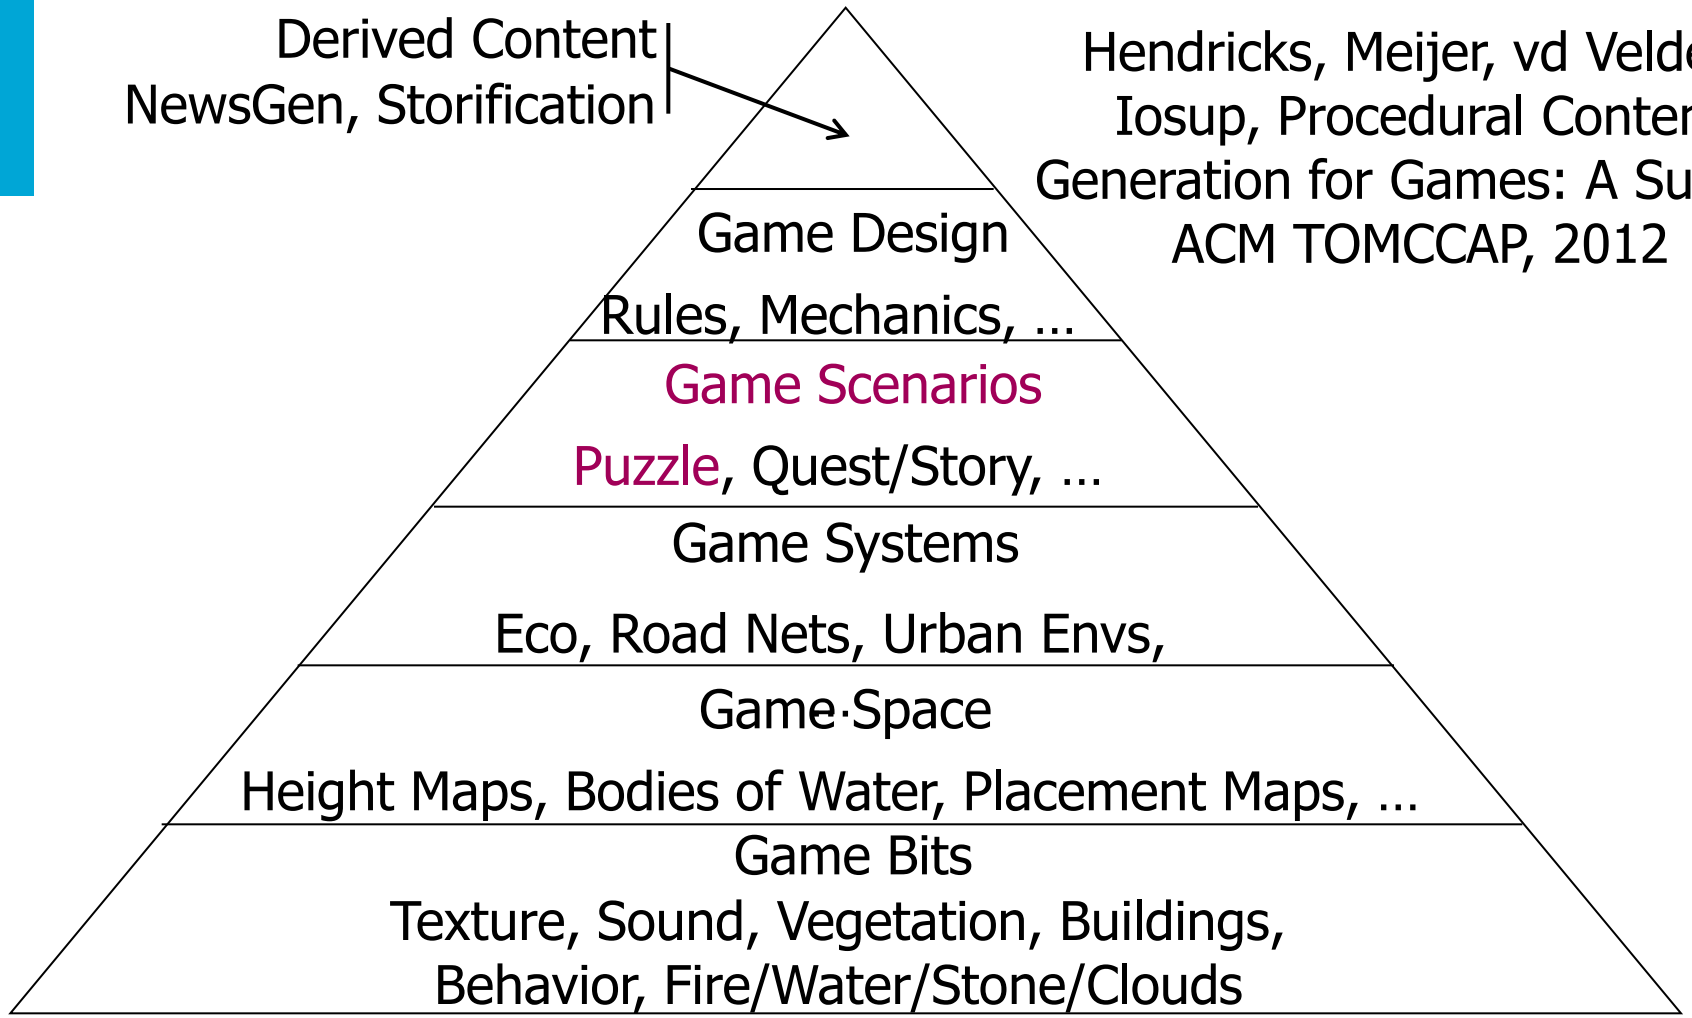

Hendricks, Meijer, vd Velden, Iosup, Procedural Content Generation for Games: A Survey, ACM TOMCCAP, 2012

The POGGI Content Generation Framework

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

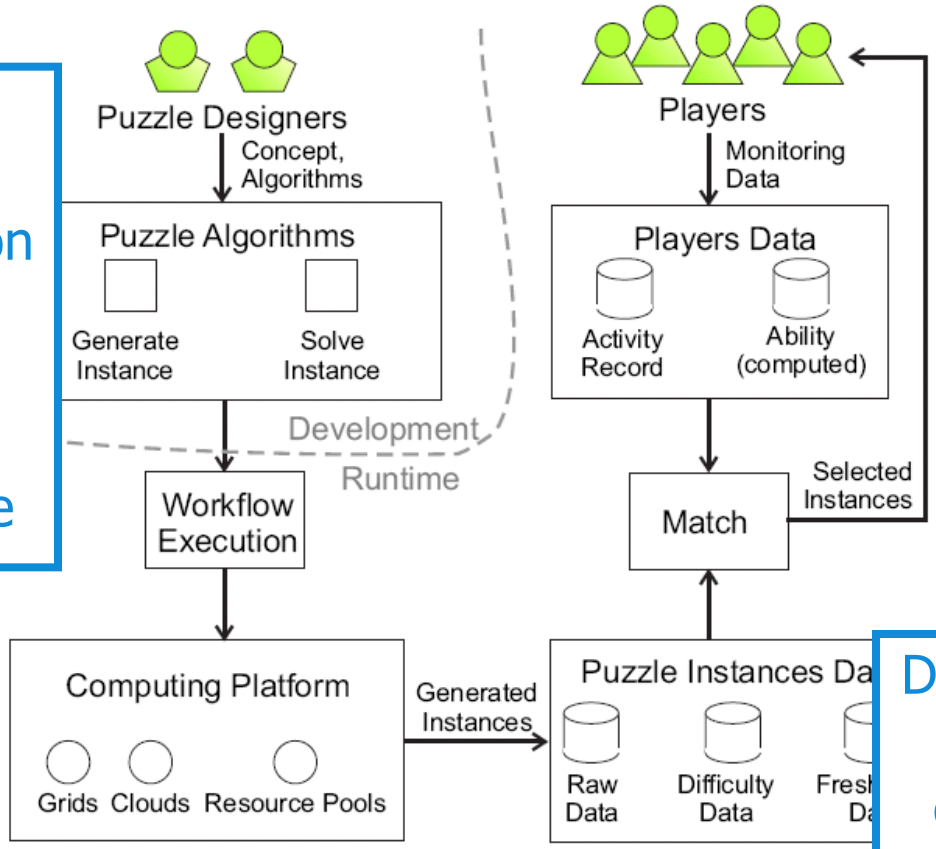

Distributed system to generate instances on-demand, reliably, efficiently, and with performance guarantees

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

Puzzle-Specific Considerations
Generating Player-Customized Content

Puzzle difficulty

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

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

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)

Player ability

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

Summary

Massivizing Online Gaming

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

Current Technology

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

Many Approaches

- Naming
- Messaging
- Consistency
- Replication
- Persistence of V-Worlds
- Etc.

@large: Our Vision

- Distributed Systems to help
- Economy of scale with clouds

@large: Ongoing Work

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

@large: The Future

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

Other Distributed Systems Issues

- Operation of replicated servers: performance guarantees
- Operation with slow user clients/networks
- Persistent worlds
- Content distribution
- The whole CAP spectrum
  - Consistency
  - Availability
  - Partition-tolerance

http://www.popscreen.com/v/6wEHS/Minecraft-Epic-Fail-Creeper

Thank you!
Suggestions? Questions?

- http://www.st.ewi.tudelft.nl/~iosup/research.html
- http://www.st.ewi.tudelft.nl/~iosup/research_cloud.html
- http://www.st.ewi.tudelft.nl/~iosup/research_gaming.html
- http://www.pds.ewi.tudelft.nl/

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

MSG Ecosystem: Model

- **Game operators**
  - Past player activity/business model → **Predicted load** → **requests**
- **Data centers**
  - **Local** time-space **renting policy** → **offers**
  - Time-Space renting policy, e.g., 1 node-hour
- Resource allocation: central request-offer matching
- Scheduling = rules for ranking request-offer match:
  1. The offer size and type vs. the request
  2. The geographical proximity offer-request
  3. The finer grained resources (quantity & time)

[Source: Nae, Iosup, and Prodan, ACM SC 2008]

Experimental Setup [1/3] Discrete-Event Simulator

- Input
  - Traces from RuneScape, a real top-5 MMOG
    - 7 countries, 3 continents
    - More than 130 game worlds
  - Consisting of
    - Geographical location
    - Number of clients
    - Over 10,000 samples at 2 min. interval, 2 weeks
- Output (for every time-step)
  - Resource allocation decisions
  - Resource allocation
  - Performance metrics

[Source: Nae, Iosup, and Prodan, ACM SC 2008]

Experimental Setup [2/3]
Environment

- 1 game operator
- 17 data centers
- 11 data center time-space renting policies

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

[Source: Nae, Iosup, and Prodan, ACM SC 2008]

Experimental Setup [3/3] Performance Metrics

- Resource over-provisioning [%]
  - The wasted resources vs. optimal provisioning at each simulation time step for all utilized machines (cumulative)
- Resource under-provisioning [%]
  - The amount of resources needed but not allocated, for all machines (computed individually)
- Significant under-provisioning events (count)
  - Count of events: resource under-provisioning is >1%, for a period of 2 minutes → **people leave**

[Source: Nae, Iosup, and Prodan, ACM SC 2008]

Impact of Load Prediction Accuracy

Q: How does the prediction accuracy impact resource provisioning? A: Good prediction matters.

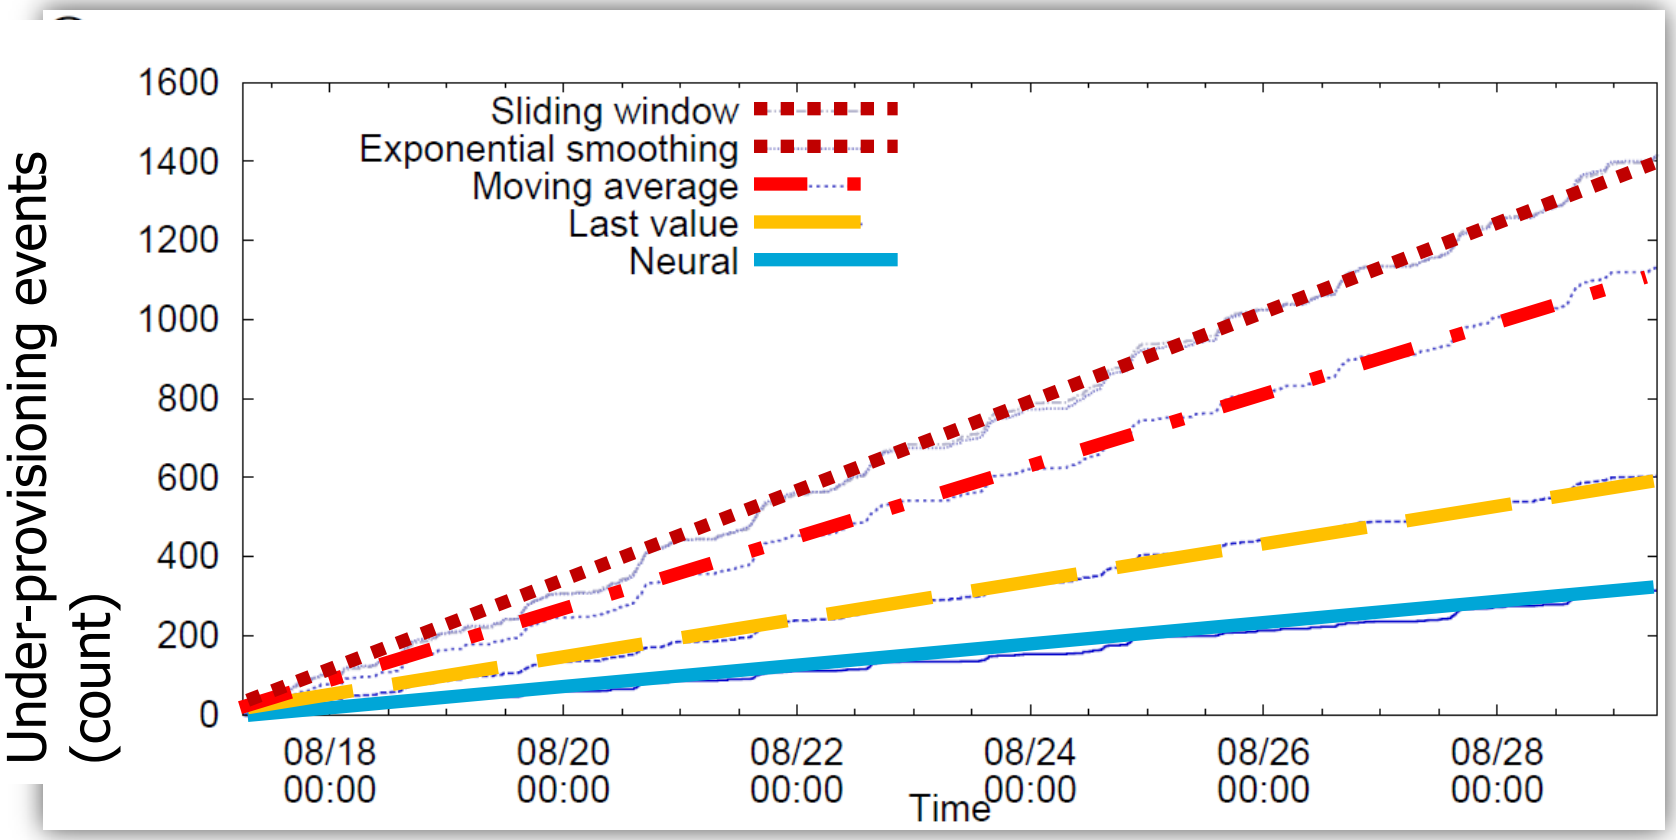

[Source: Nae, Iosup, and Prodan, ACM SC 2008]

Latency Tolerance: From None to High

Q: What is the impact of latency tolerance on hosting?

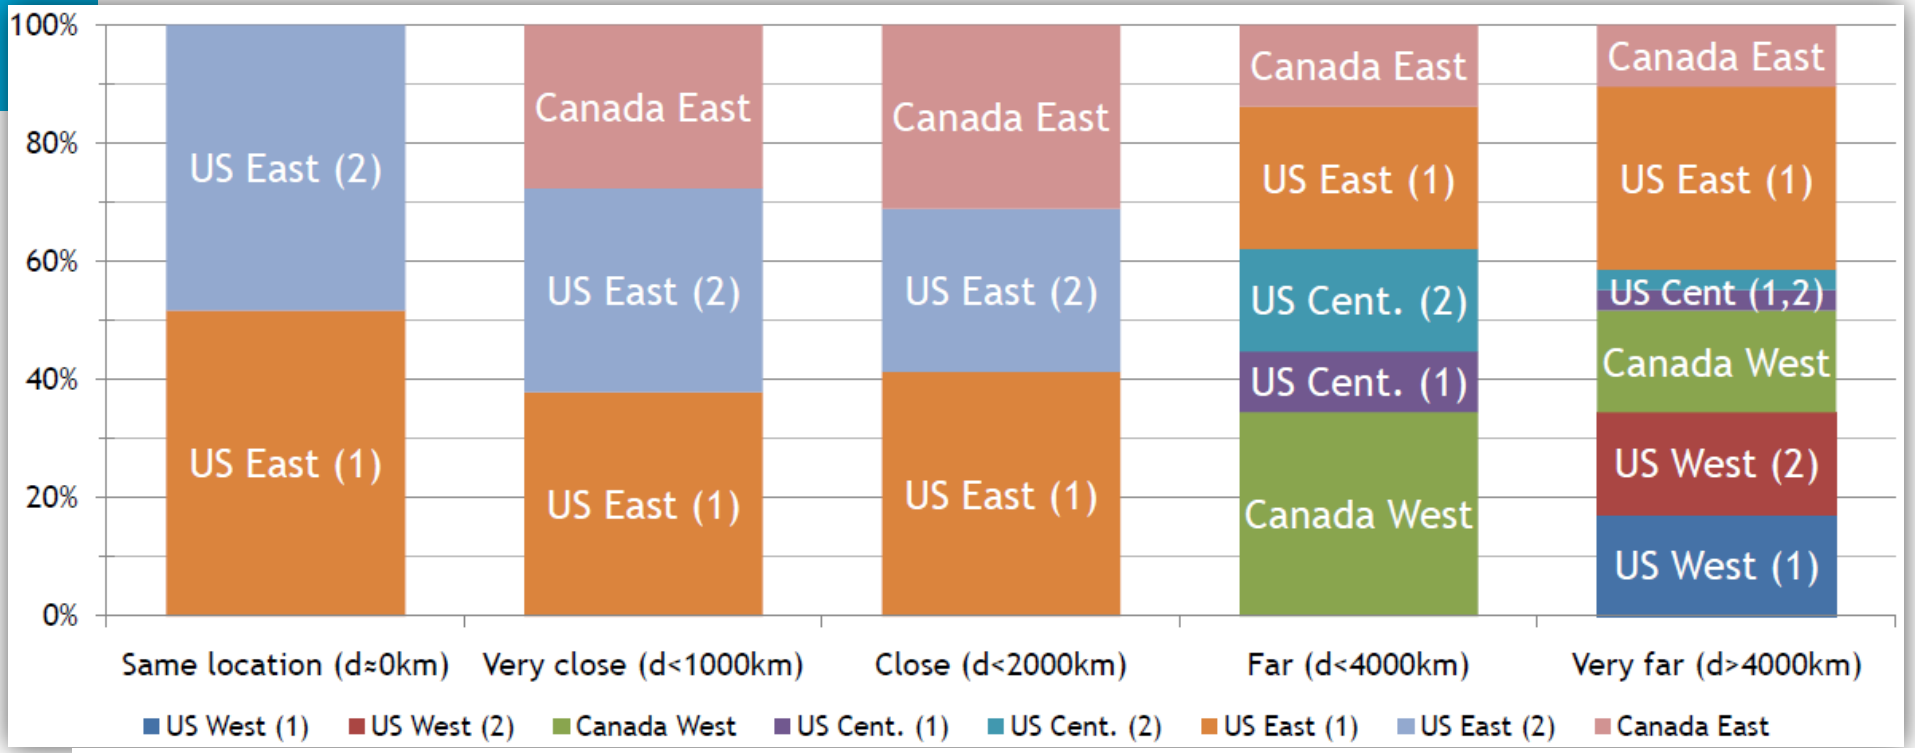

A: (left) very sensitive very costly

(mid) sensitive costly

(right) non-sensitive cheap

Traditional AoI does not work

- Area of Interest (AoI) = traditional mechanism for RPG:
  only receive/process information around avatar, but…
- …In RTS, each player has tens of units under control, so much more data to be transferred
- … In RTS, players change focus (interest) more often than in RPG and FPS, so higher management overhead

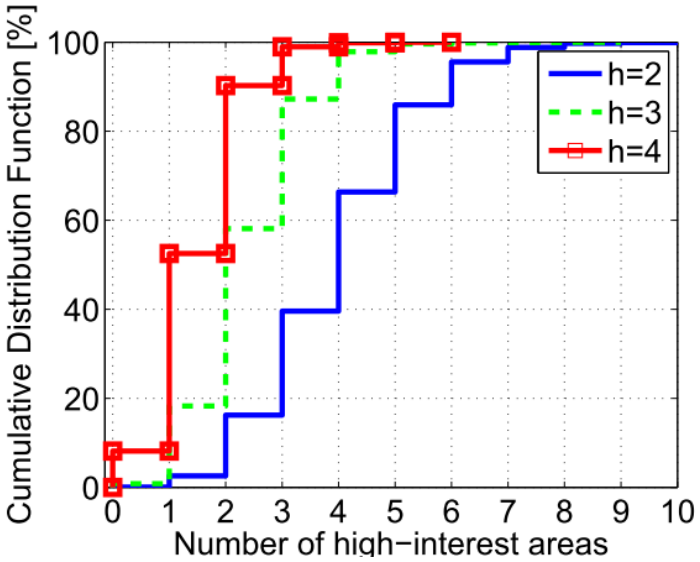

h=high:normal interest

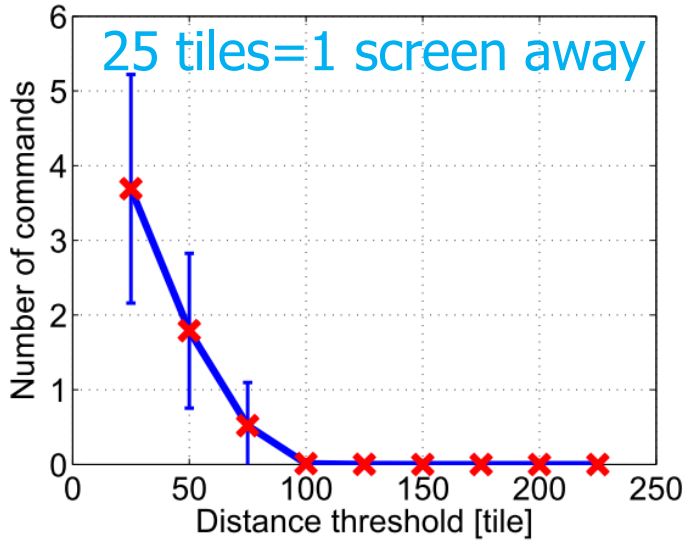

10s period, distance from first command

S. Shen, A. Iosup, D. H. J. Epema, and S.-Y. Hu. Area of Simulation: Mechanism and Architecture for Multi-Avatar Virtual Environments, ACM Trans. Multimedia Comput. Comm. App. 2015.

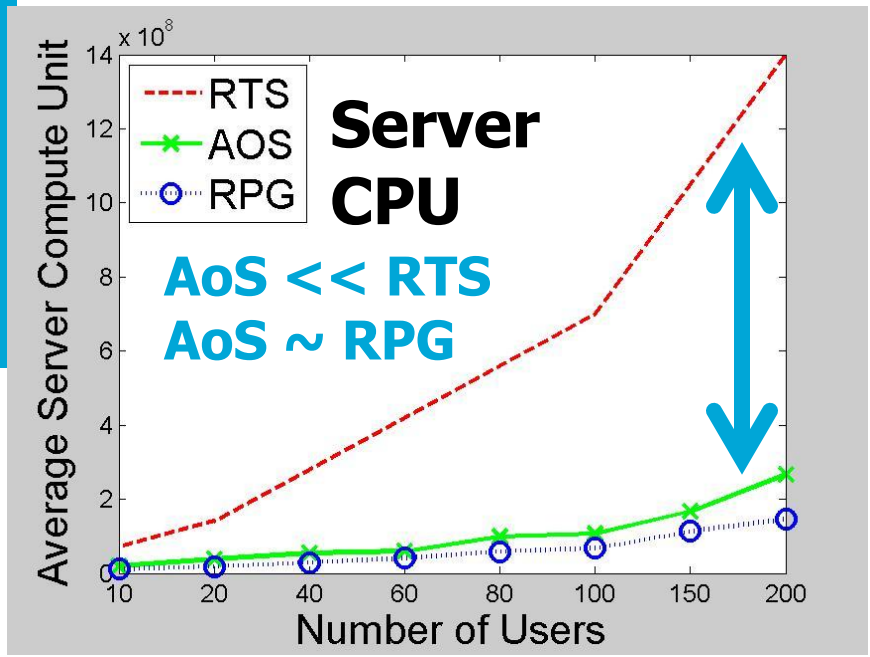

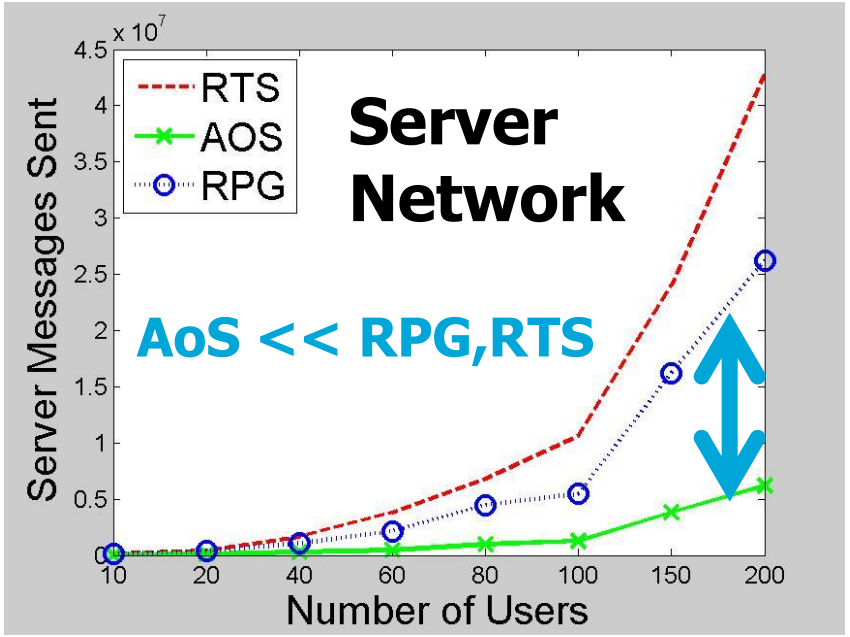

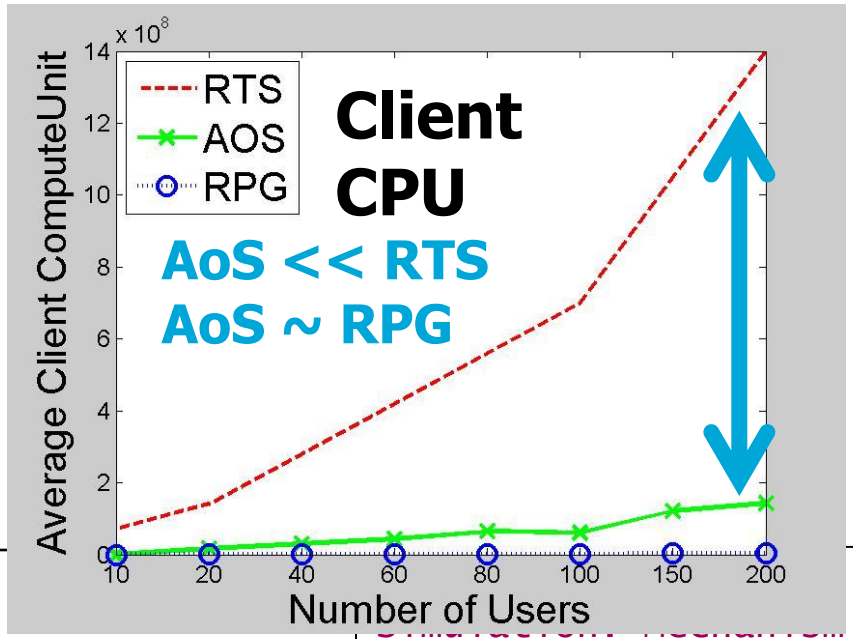

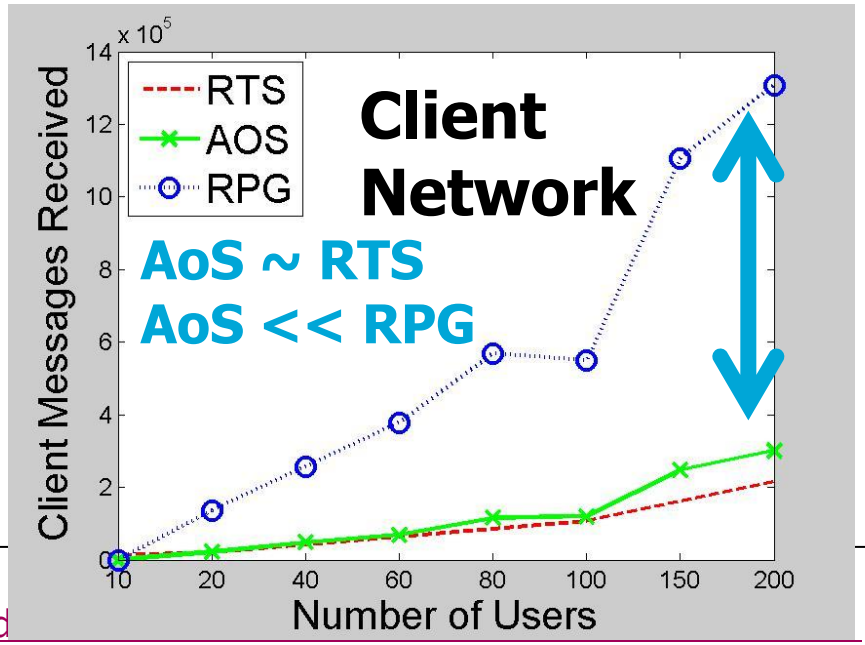

H.
and
Environments, ACM Trans. Multimedia Comput. Comm. App. 2015.

General Offloading Model

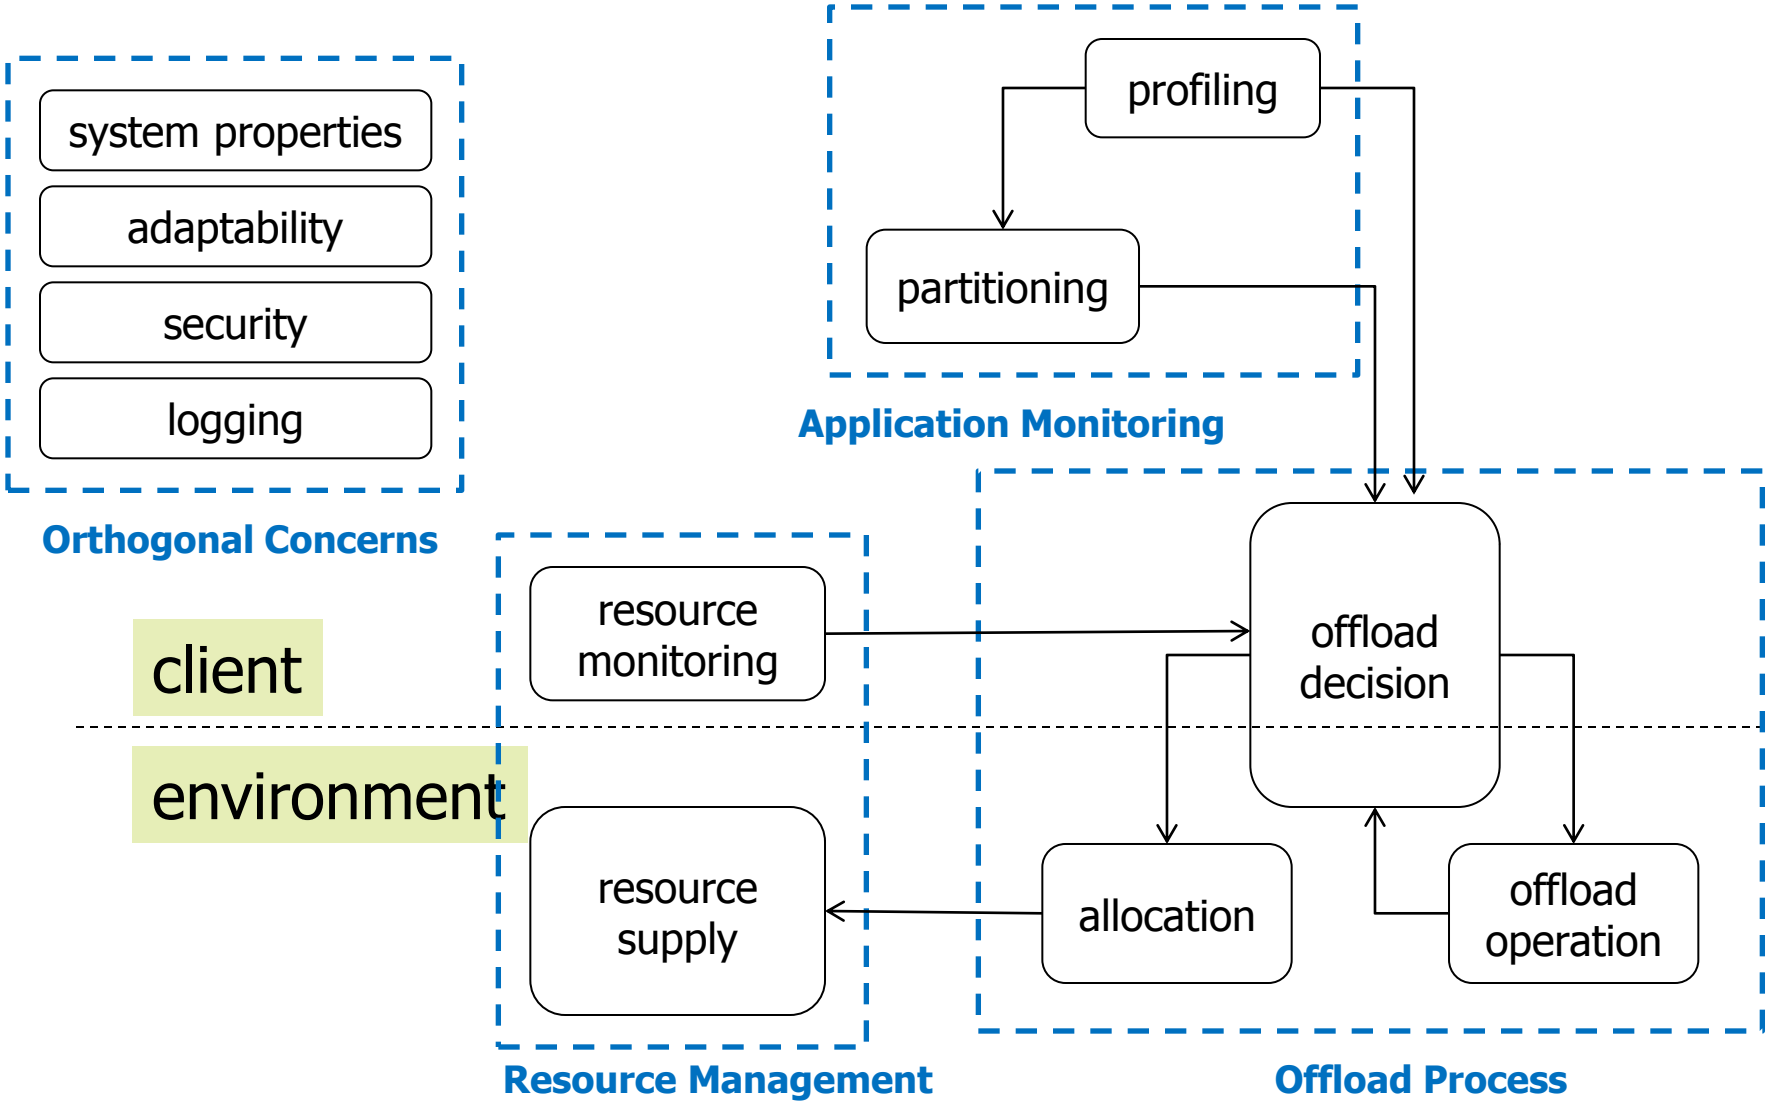

A.-C. Olteanu, N. Tapus, A. Iosup: Extending the Capabilities of Mobile Devices for Online Social Applications through Cloud Offloading. CCGRID 2013: 160-163.

Player-Customized Content

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

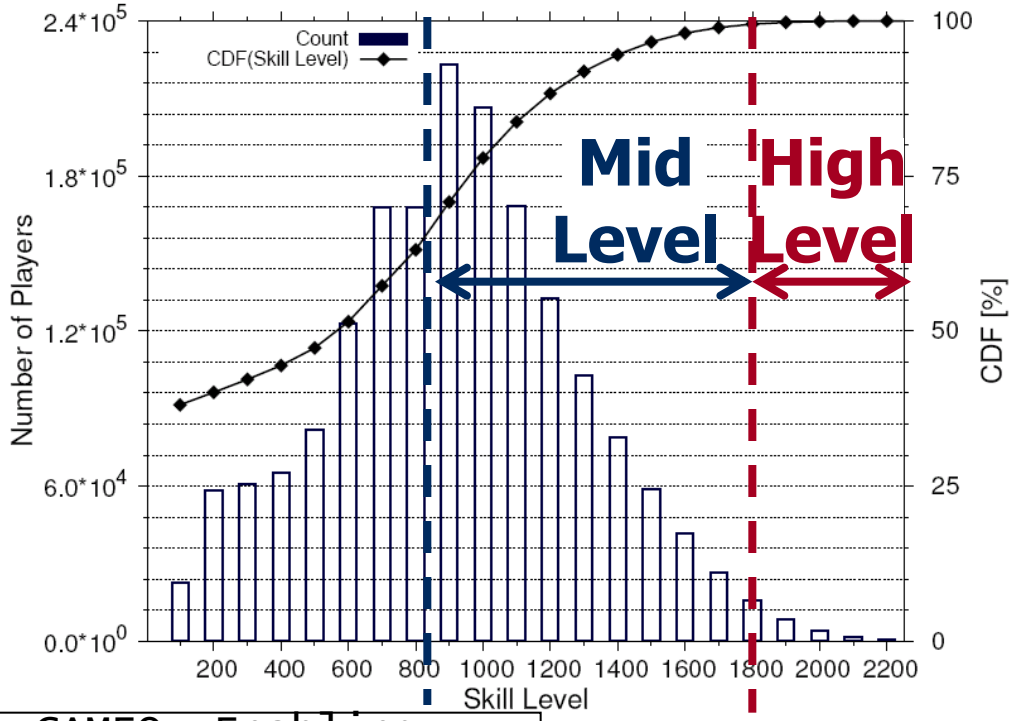

A. Iosup, A. Lascateu, N. Tapus, CAMEO: Enabling Social Networks for Massively Multiplayer Online Games through Continuous Analytics and Cloud Computing, ACM NetGames 2010.

Workload Model

How does the number of users evolve over time?

Macro Model

What workloads do user's actions trigger?

Micro Model

Workload Model

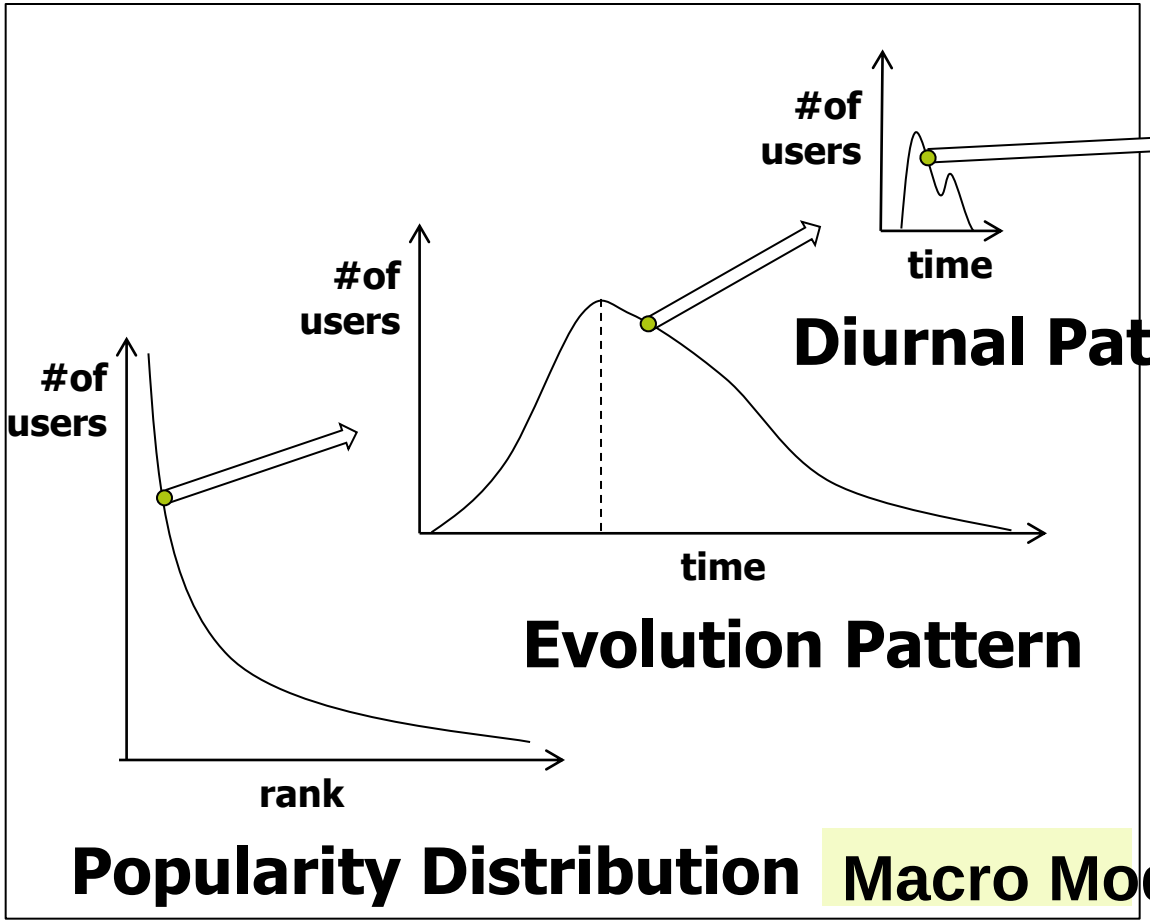

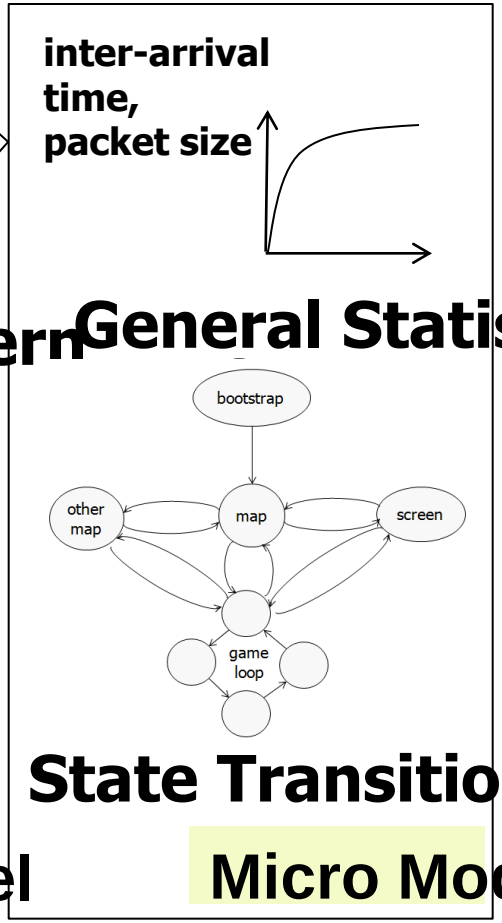

(w/ A.Iosup)

MSGs are a Popular, Growing Market

- **25,000,000+ subscribed players (from 250,000,000+ active)**
- **Over 10,000 MSGs in operation**
- **Subscription market size $7.5B+/year, Zynga $600M+/year**

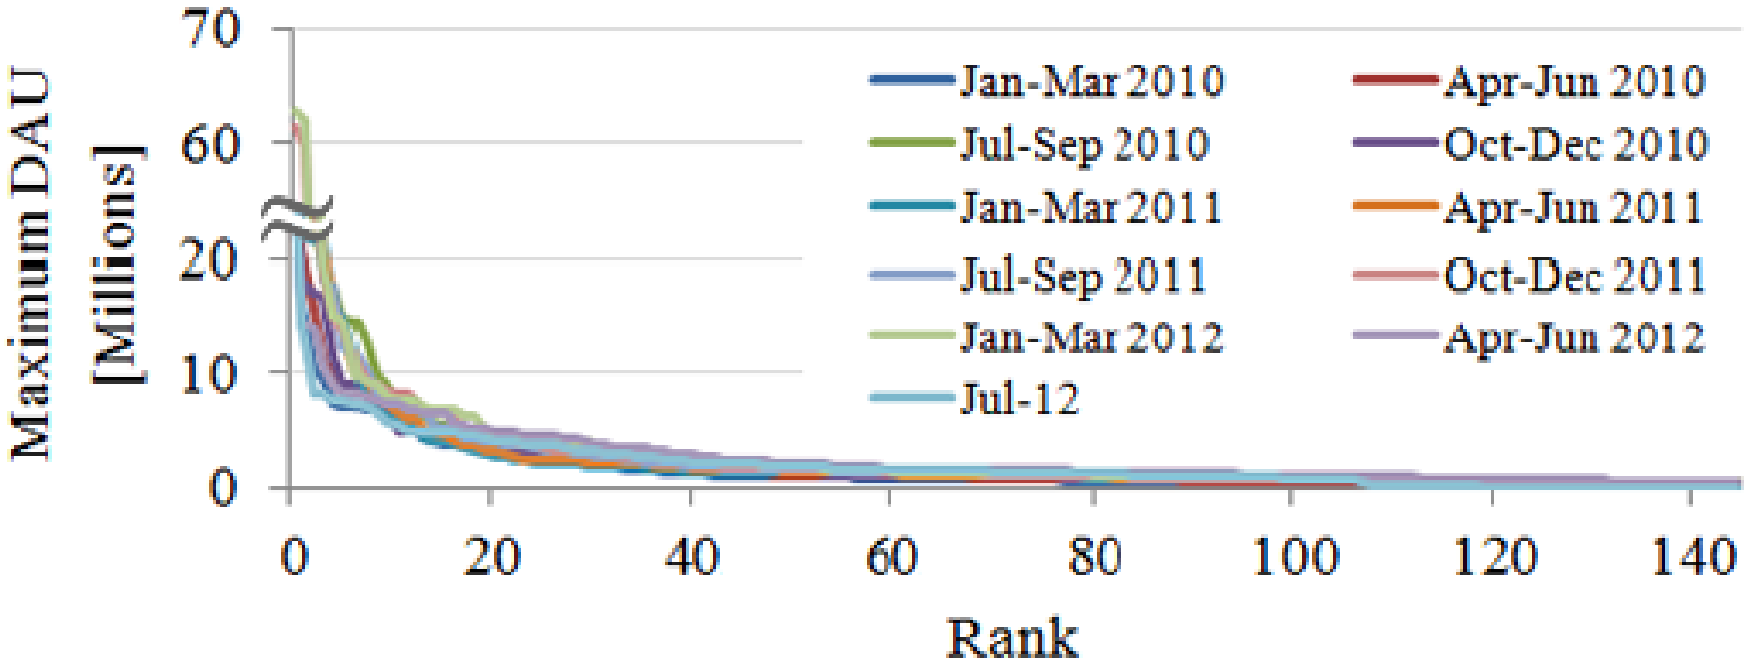

A.C. Olteanu, A. Iosup, and N. Tapus, Towards a workload model for online social applications. Proc. ICPE 2013 http://bit.ly/VzTei2

Simulation

Design & Implementation:

- generate evolution pattern and general network statistics (workload model)
- various cases: session duration, graphics, offloading mechanism

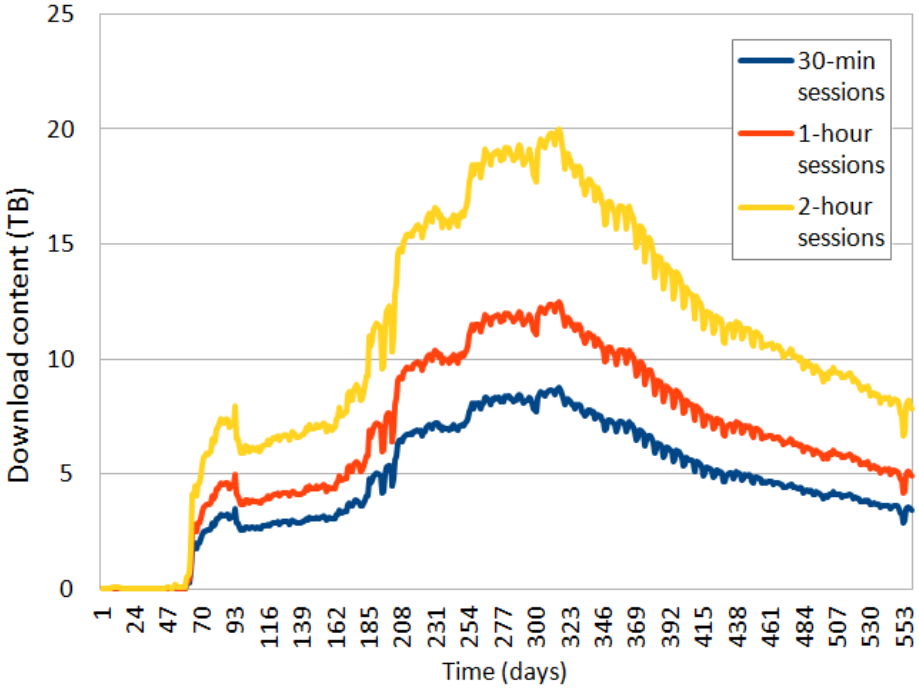

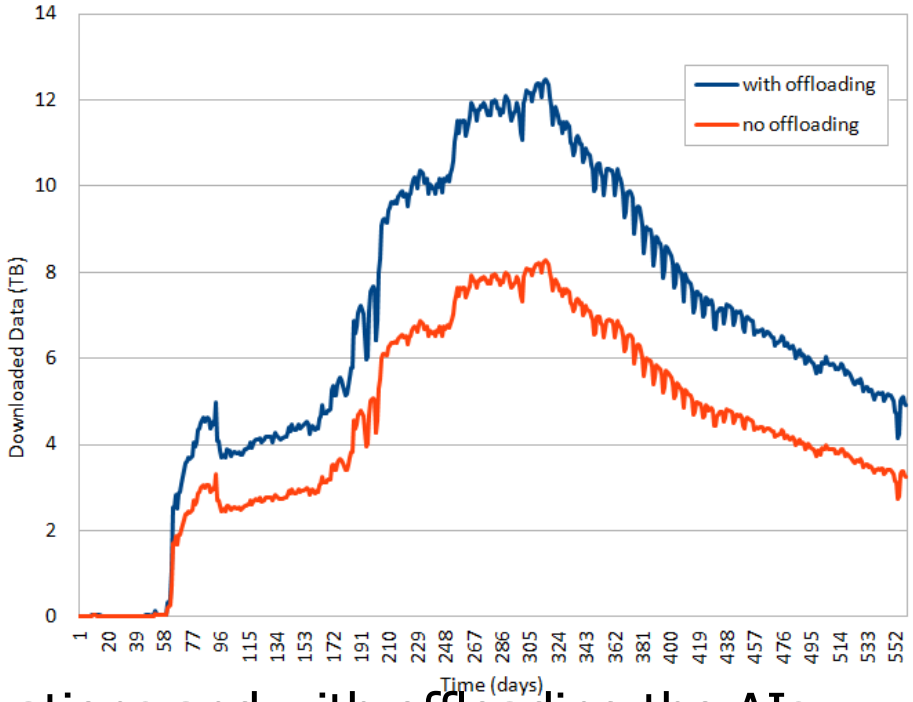

Simulations with various session durations and with offloading the AIs using a synthetic trace: peak_DAU = 1 mil., lifetime = 18 months