

# MASSIVIZING COMPUTER SYSTEMS

MASSIVIZING ONLINE GAMES: THE SCIENCE, DESIGN, AND  
ENGINEERING OF DISTRIBUTED ECOSYSTEMS (FOR ONLINE GAMING)

@Large Research  
Massivizing Computer Systems



<http://atlarge.science>

[bit.ly/MassivizingGames22](http://bit.ly/MassivizingGames22)

Massivizing Online Games =  
Rich challenge of computer  
systems, with societal impact!



Contributions from the MCS team. Many thanks!  
Many thanks to our collaborators, international working groups,  
authors of all images included here.  
Also many thanks for the invitation, to all involved!

Sponsored by:



Prof.dr.ir. Alexandru

IOSUP

# US IN 1 MINUTE




WE'RE  
MASSIVIZING  
COMPUTER  
SYSTEMS!



# VU AMSTERDAM < SCHIPHOL < THE NETHERLANDS < EUROPE



Amsterdam  
founded 10<sup>th</sup> century  
pop: 850,000



VU  
founded 1880  
pop: 23,500



# WHO AM I?

## PROF. DR. IR. ALEXANDRU IOSUP

- Education, my courses:
  - > Systems Architecture (BSc)
  - > Distributed Systems, Cloud Computing (MSc)
- Research, 15 years in DistribSys:
  - > Massivizing Computer Systems
- About me:
  - > Worked in 7 countries, NL since 2004
  - > I like to help... I train people in need
  - > VU University Research Chair + Group Chair
  - > NL ICT Researcher of the Year
  - > NL Higher-Education Teacher of the Year
  - > NL Young Royal Academy of Arts & Sciences
  - > Knighted in 2020











# <http://atlarge.science>

## CURRENT TEAM

This is us, now.

-  Professor
-  Assistant Prof.
-  Teacher
-  Visitor/P.-doc
-  Ph.D. student
-  Early Scientist

## WE ARE HIRING A NEW ASST. PROF.!

# WE ARE A FRIENDLY, DIVERSE GROUP, OF DIFFERENT RACES AND ETHNICITIES, GENDERS AND SEXUAL PREFERENCES, AND VIEWS OF CULTURE, POLITICS, AND RELIGION. YOU ARE WELCOME TO JOIN!

## Alumni

They have completed a long-term project in our team.

Shanny Aneep Team VL-e	Athanasios Antoniou Team AtLarge	Sietsje Au M.Sc. student, TU Delft	Johannes Bertens M.Sc. student, TU Delft	Marcin Biczak Researcher in graph-processing team	Mihai Capota Tech Lead Graphalytics
Bogdan Ghit Ph.D. student, TU Delft	Yong Guo Graph processing	Stijn Heldens Researcher, TU Delft	Alexey Ilyushkin Ph.D. student, TU Delft	Adele Lu Jia Social gaming	Elvan Kula Honors Track
Chris LeMaire Team Graphalytics	Shenjun Ma M.Sc. student, TU Delft	Ahmed MUSAfir Researcher, Vrije	Wing Lung Ngai Researcher, Vrije	Leon Overweel Core Team OpenDC	Siqi Shen Massivizing online
Jie Shen Performance modeling	Ruben Verboon Honors Track	Maria Anemona Voinea M.Sc. student, TU Delft	Nozhi Yigitbasi Tech Lead GrenchMark and CMeter	Ernst van der Hoeven M.Sc. student, TU Delft	Jerom van der Sar Team OpenCraft

## Research Visitors and Interns

They have completed a short-term stay with our team.

Mugurel Ionut Andreia Research visitor	Matthijs Bijman Core Team OpenDC	Alexandru Costan Research visitor	Kefeng Deng Research visitor	Yunhua Deng Research visitor	Kevin Dehnenman B.Sc. student, Vrije Universiteit
-------------------------------------------	-------------------------------------	--------------------------------------	---------------------------------	---------------------------------	------------------------------------------------------

## ALUMNI

## RS

# MASSIVIZING COMPUTER SYSTEMS: OUR MISSION

<http://atlarge.science/about.html>



1. Improve the lives of millions through impactful research.



2. Educate the new generation of top-quality, socially responsible professionals.



3. Make innovation available to society and industry.





# THIS IS THE GOLDEN AGE OF COMPUTER ECOSYSTEMS

1

# THIS IS THE GOLDEN AGE OF MASSIVE COMPUTER ECOSYSTEMS



Education for Everyone (Online)




Business Services




Big Data

Edge Computing



Big Science



Internet

5G/6G/\*

THz

VLC

Cloud Computing




CTAAAGATGATCTTAGTCCCGGTTTCGAA  
TCTTTAGTCCCGTTGATAACCAACC  
GTAATACCAACCGGGACTAAAGATCCCG  
GGGACTAAAGTCCACCCCTATATATG  
  
TTCAAAATTTCTCAAAAAAGAGGGGAG  
GTGATACATACAAATGGGAGGTGCCTA  
TTTGTACACTACATTTGCACCTATGTTT  
GTAAGTTGATGAGAGAAAATGTGTG



Datacenter

Grid Computing



Daily Life



Iosup et al., Massivizing Computer Systems, ICDCS 2018. [Online] Hesselman, Grosso, Kuipers, et al. (2020) A Responsible Internet to increase Trust in the Digital World. JNSM [Online]



# THIS IS THE GOLDEN AGE OF MASSIVE COMPUTER ECOSYSTEMS



Education for Everyone (Online)




Business Services




Big Data

Edge Computing



Big Science




Internet


5G/6G/\*

THz


VLC



Cloud Computing




CGACATAT  
CTAAAGATGATCTTTAGTCCCGGTTTGGAA  
TCTTTAGTCCCGGTTTATAACCAACCAAC  
GTAATACCAACCGGGACTAAAGATCCCG  
GGGACTAAAGTCCCAACCCCTATATATATG  
TTCAAAATTTCTTCAAAAAAGAGGGGAG  
GTGATTACATACAAATGGGAGGTGCCTA  
TTTGTGCATACATTTTGCACCTATGTTT  
GTAAGTTGATGAGAGAAAATGTGTG



Datacenter

Grid Computing

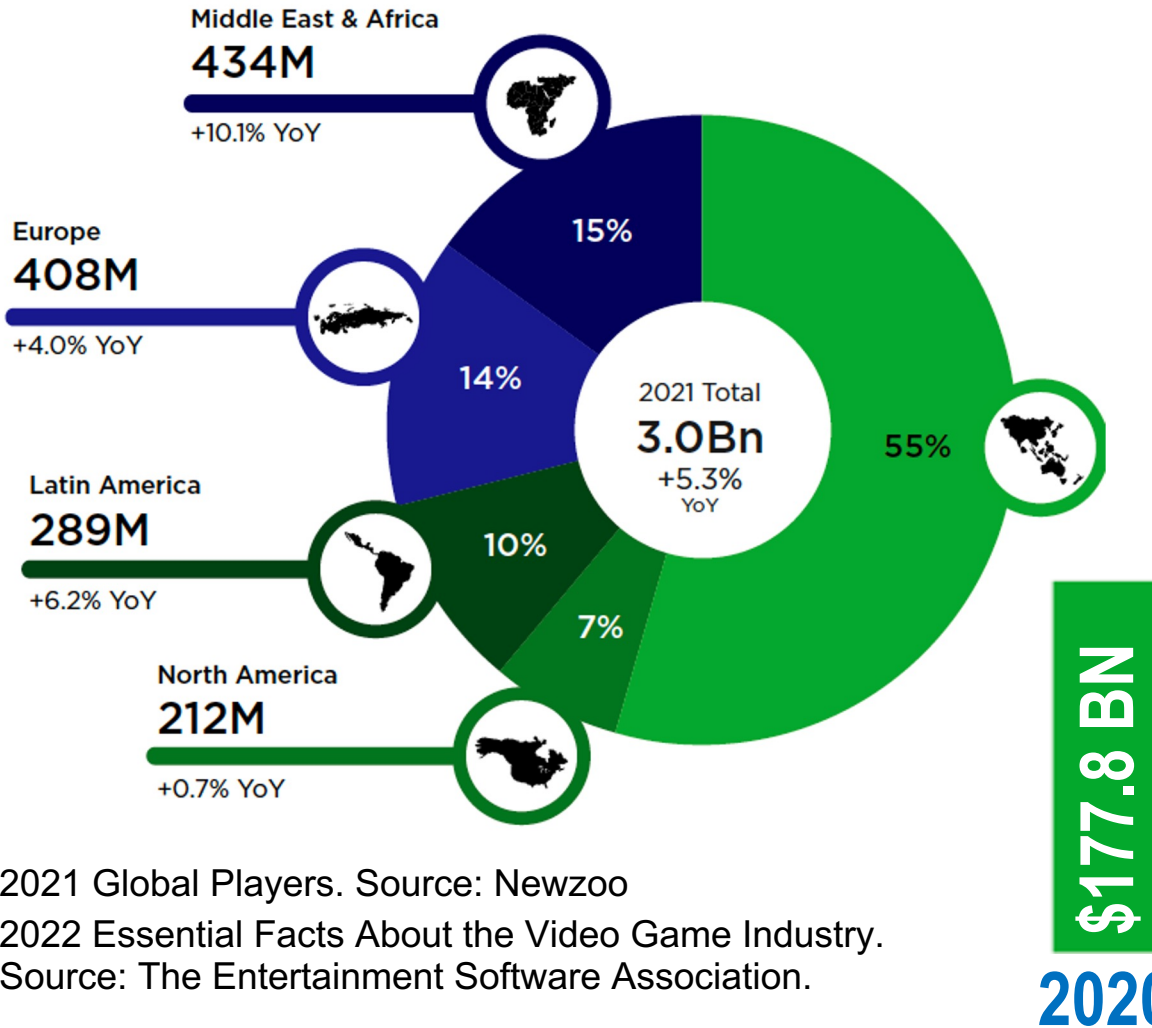


Daily Life



Iosup et al., Massivizing Computer Systems, ICDCS 2018. [Online] Hesselman, Grosso, Kuipers, et al. (2020) A Responsible Internet to increase Trust in the Digital World. JNSM [Online]

# LOTS OF GAMERS, LOTS OF CULTURE, LOTS OF REVENUE



**AVERAGE TIME SPENT PLAYING GAMES EACH WEEK: 13 HOURS**

(UP 7% FROM 12 HOURS PER WEEK IN 2021)




**48%** identify as female

**52%** identify as male†

The average video game player is

**33 years old**



2021 Global Players. Source: Newzoo


2022 Essential Facts About the Video Game Industry.

Source: The Entertainment Software Association.



# ME AS A GAMER: MULTI-PLATFORM/-GENRE + COMPLETIONIST

Level (61)





# GAMES ARE INCREASINGLY MORE SOCIAL, BUILD COMMUNITIES



**83%**

of players say games introduce people to new friendships and relationships (up from 78%)



**46%**

of players have met a good friend, spouse, or significant other through video games (up from 42%)



**61%**

of players have met people through video games they otherwise would not have met (up from 54%)



**61%**

of players agree video games have helped them stay connected to friends/family (up from 53%)



**55%**

of players say video games have helped them develop deeper relationships with others



**78%**

of players agree video games promote social interaction



**72%**

of players see benefits of games for existing relationships



**89%**

of players see benefits of games for new relationships



**67%**

of parents agree video games helped their child connect with friends and family during the pandemic

2013 Game Masters museum exhibition. Personal collection.  
2022 Essential Facts About the Video Game Industry.  
Source: The Entertainment Software Association.



# Games and Beyond

## Take Minecraft

- Over **125 million people** play Minecraft every month
- **40,000+ mods**
- **100+ games “like M’craft.”**

## Generally Beneficial Features

- **Entertainment**
- **Education**
- **Activism**
- **Social Interaction**



Minecraft: Connecting More Players Than Ever Before

by Helen Chiang, Studio Head, Mojang Studios • May 18, 2020 @ 6:00am



## HISTORY BLOCKS

8-10 yrs old 11-13 yrs old 14-18 yrs old 18+ yrs old Geography

Reading and Writing Service Learning & Social Good

The purpose of this activity is to guide teachers through an activity in which students reconstruct Unesco world heritage sites Minecraft.

## How Greenpeace Used Minecraft to Stop Illegal Logging in Europe’s Last Lowland Primeval Forest

Good game, everybody

By Angela Natividad | January 22, 2018

## Pussy Riot, Idles to play Minecraft virtual festival ‘Block by Blockwest’

It starts this Saturday

By Makena Kelly | @kellymakena | Apr 22, 2020, 1:33pm EDT



BUT WE CANNOT  
TAKE THIS  
TECHNOLOGY  
FOR GRANTED

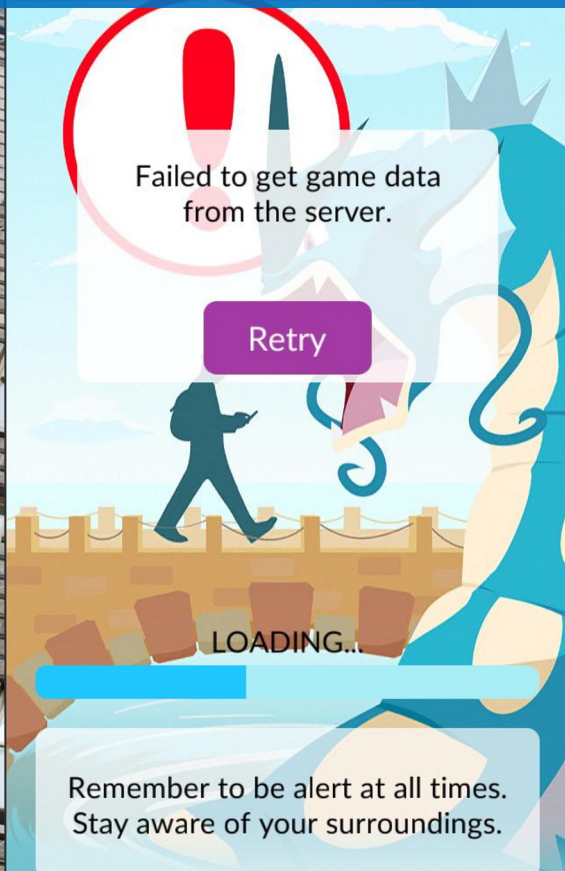
2

(So, this is why I am giving this talk)



# PHENOMENON: UNAVAILABILITY OF GAMING SERVICES

## UNCOVERING THE PRESENCE OF FAILURES



Source: gamenewstoday.com

### Pokémon GO Server Status

REFRESH

#### Pokémon GO

**OFFLINE**  
for 15 minutes

#### Pokémon Trainer Club

**UNSTABLE**  
for 2 minutes

#### Pokémon GO Uptime

**55.56%**  
over the past hour

**96.29%**  
over the past day

#### Pokémon Trainer Club Uptime

**66.67%**  
over the past hour

**96.66%**  
over the past day





# PHENOMENON: PERFORMANCE DROPS IN GAMES

UNCOVERING THE PRESENCE OF PERFORMANCE ISSUES, EVEN LEADING TO CRASHES



Source: <http://bit.ly/EveOnline21Crash>

NEWS

## Players in Eve Online broke a world record — and then the game itself

*Developers said they're not 'able to predict the server performance in these kinds of situations'*

By [Charlie Hall](#) | [@Charlie\\_L\\_Hall](#) | Jan 5, 2021, 2:54pm EST



Source: Razorien/CCP Games

# PHENOMENON: GAME SERVICE SUSTAINABILITY

UNCOVERING THE USE OF ENERGY AND WATER, THE IMPACT ON CLIMATE

Power consumption of  
datacenters:  
**>1% of global electricity**

Source: Nature, 2018 [\[Online\]](#)

Water consumption of  
datacenters  
in the US:  
**>625Bn. l/y (0,1%)**

Source: Energy Technologies Area, 2016 [\[Online\]](#)

Power consumption of datacenters  
in the Netherlands:  
**1→3% of national electricity**

Source: NRC, 2019 [\[Online\]](#)

Other greenhouse emissions:  
**Largely unknown**

Source: Nature Climate Change, 2020 [\[Online\]](#)



# PHENOMENON: CHEATING, OTHER TOXICITY IN GAMES

UNCOVERING THE PRESENCE OF TOXICITY AND FINDING WAYS TO ERADICATE IT

GAMING ENTERTAINMENT TECH

Source: [The Verge](#).

## Destiny 2 cheat creators come to a \$13.5 million settlement with Bungie

*A traveler-sized sum*

By [Alice Newcome-Beill](#) | Jun 17, 2022, 3:09pm EDT

f   SHARE



Valorant will start monitoring your voice chats starting July 13th

*It's part of Riot's efforts to crack down on disruptive behavior*

By [Emma Roth](#) | Jun 26, 2022, 5:05pm EDT

f   SHARE





# THIS TALK, IN A NUTSHELL:



Massivizing Online Games =  
Rich challenge of computer  
systems, with societal impact!

Online Gaming used to be art,  
is now also massive computing

Online Gaming used to be  
networking, is now all of computing


Online Gaming used to be game  
worlds, is now all kinds of apps

What to do?

New science

Lots of work





# Massivizing Online Games = Rich challenge of computer systems, with societal impact!



1. This is the Golden Age of massive computer ecosystems
2. But we cannot take this technology for granted
3. A new science to help massivizing computer ecosystems
4. Understand how things work: observe, synthesize, experiment
5. Design and build better systems:  
virtual worlds, analytics, content generation, meta-gaming
6. What's next? PLATO in the 21<sup>st</sup> century
7. Beyond games: Massivizing computer systems is a national priority in the NL!

Lots to do  
next!



A new science, of  
complex, smart  
computer ecosystems

3

(operational simplicity  
for the user)

# DISTRIBUTED ECOSYSTEMS, OUR DEFINITION

1. Set of 2+ **constituents**, often **heterogeneous**
2. Each constituent is a system or an ecosystem (**recursively**)
3. Constituents are **autonomous**, cooperative or in competition
4. Ecosystem **structure** and **organization** ensure responsibility
  1. Completing functions and providing services
  2. Providing desirable non-functional properties
  3. Fulfill agreements with both operators and clients, clients in the loop
5. Long and short-term **dynamics** occur in the ecosystem

losup et al., Lecture Notes in Distributed Systems, Section 1.1.1

losup et al., Massivizing Computer Systems, ICDCS 2018. [\[Online\]](#)

# AN ANALOGY: MASSIVIZING CLIMATE SCIENCE

TAKE A HOLISTIC VIEW, BASED ON COUPLED NATURAL SYSTEMS

Can be understood only with coupled models

\* In climate science, issues are often linked. The same occurs in massive computer (eco)systems.




(ONLINE) GAMES = MASSIVE COMPUTING ECOSYSTEMS

Do you recognize this App?



Here is how it operates...

ECOSYSTEM = SERVICES + COMPUTING + SMARTS + GOALS



Computing  
Continuum

Extreme Automation, Performance, Dependability, Sustainability

# THIS IS THE MODERN SCIENCE OF COMPUTER ECOSYSTEMS

## IN A NUTSHELL

### WHO?

 SCIENTISTS,  ENGINEERS,  DESIGNERS,  MANAGERS, ETC.

### WHAT? MAIN GOAL

UNDERSTAND AND CONTROL DISTRIBUTED ECOSYSTEMS, TO  
TURN THEM INTO EFFICIENT, AUTOMATED UTILITIES

### HOW? CENTRAL PARADIGM


ECOSYSTEM OPERATION AND CHARACTERISTICS DERIVE  
NON-TRIVIALY FROM ITS SYSTEMS AND USERS (RECURSIVELY)

### WHICH APPROACH?

MODERN DISTRIBUTED SYSTEMS AND PROBLEM-SOLVING







# Massivizing Online Games = Rich challenge of computer systems, with societal impact!



1. This is the Golden Age of massive computer ecosystems
2. But we cannot take this technology for granted
3. A new science to help massivizing computer ecosystems
4. Understand how things work: observe, synthesize, experiment
5. Design and build better systems:  
virtual worlds, analytics, content generation, meta-gaming
6. What's next? PLATO in the 21<sup>st</sup> century
7. Beyond games: Massivizing computer systems is a national priority in the NL!

Lots to do  
next!



# LET'S UNDERSTAND HOW THINGS WORK!

4



Message: **4<sub>A</sub>**

Synthesize the  
ecosystem  
reference  
architecture

# MEANINGFUL DISCOVERY

BUT ... IS THERE A SYSTEMATIC WAY TO APPROACH THESE PHENOMENA?



- The Human Genome Project:
  - > Physical map covering >90% human genome
  - > Sequence data made available open-access
- Big Science:
  - > Took >10 years to complete
  - > Led by US, work by 20 groups in CN, DE, FR, JP, UK, US
- Big impact:
  - > Decrease cost of sequencing
  - > Facilitate biomedical research

FUNDING: > 3B USD

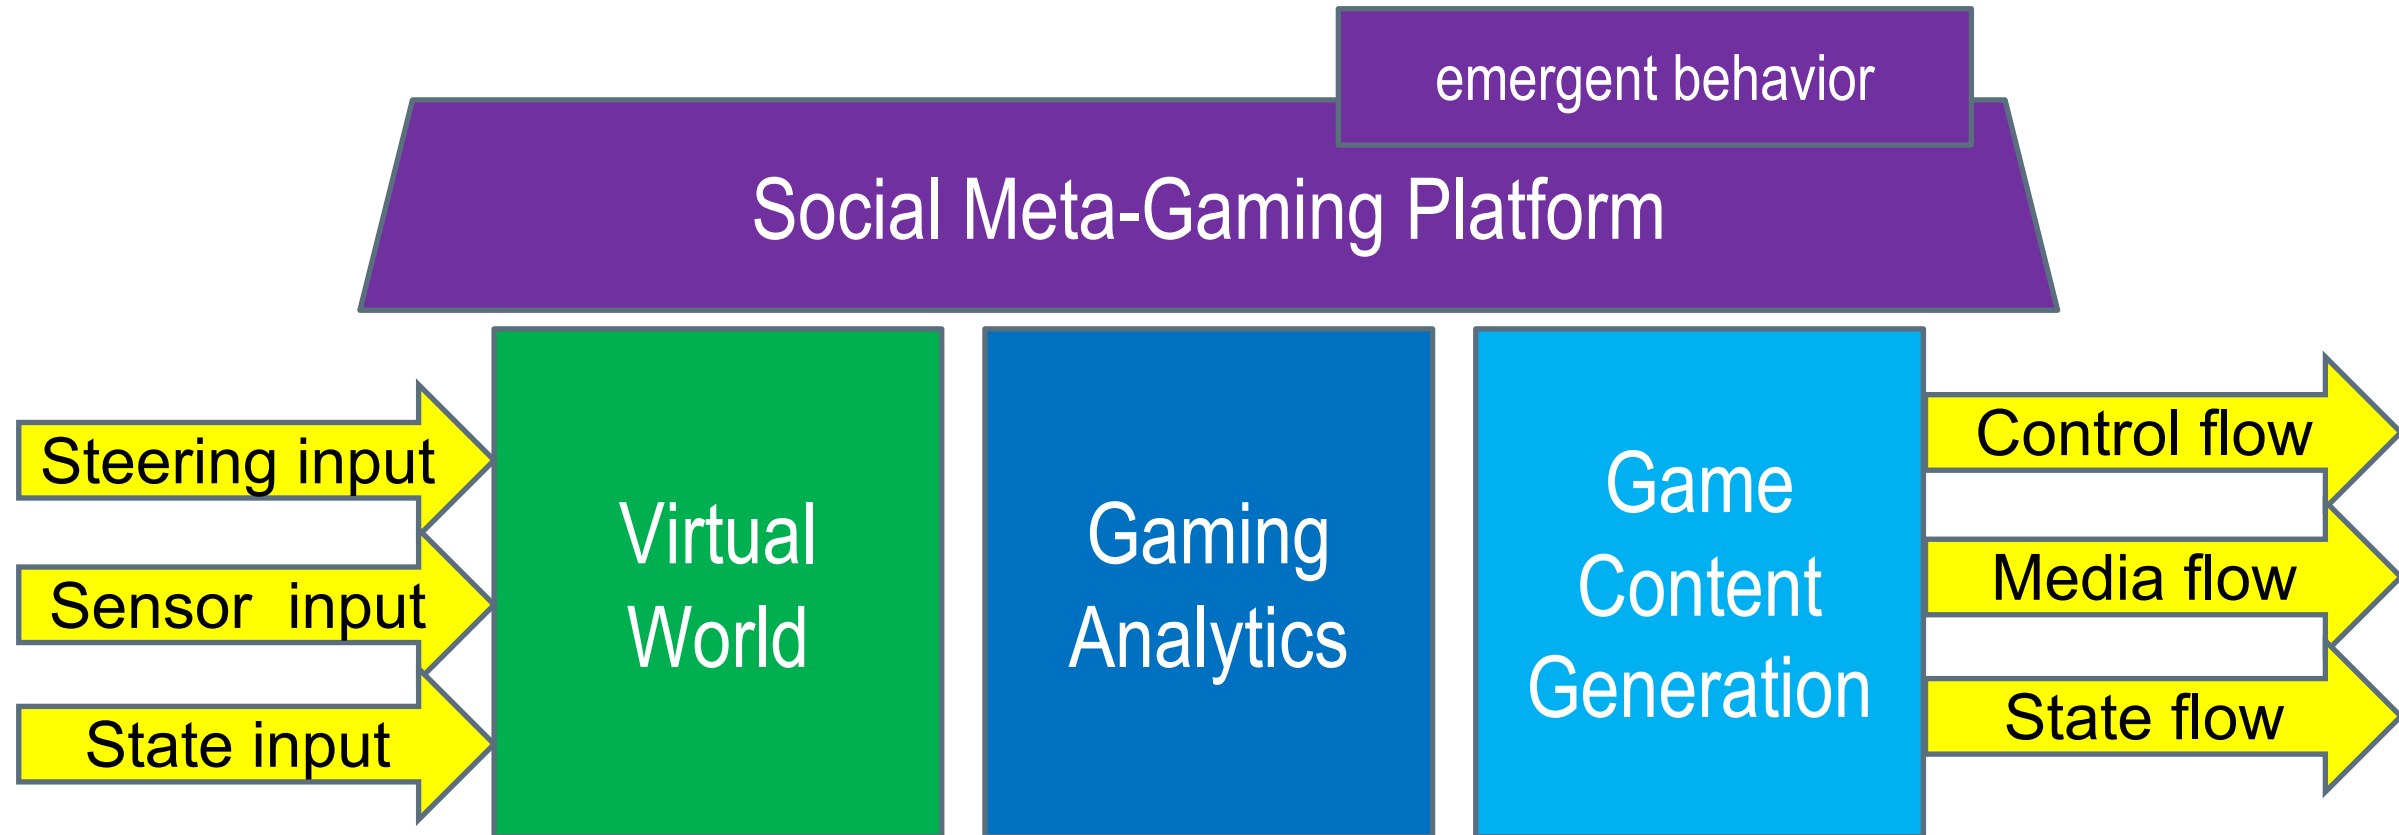
International Human Genome Sequencing Consortium, Initial sequencing and analysis of the human genome, Nature 409, Feb 2001. [\[Online\]](#)

Julie Gould, The Impact of the Human Genome Project, Naturejobs blog, 2015. [\[Online\]](#)

# ONLINE GAMING ECOSYSTEMS, HIGH-LEVEL VIEW

THE COMPLEXITY CHALLENGE

IOSUP ET AL. MASSIVIZING ONLINE GAMES



Iosup, Shen, Guo, Hugtenburg, Donkervliet, Prodan (2014) Massivizing online games using cloud computing: A vision. ICME Workshops [Online]

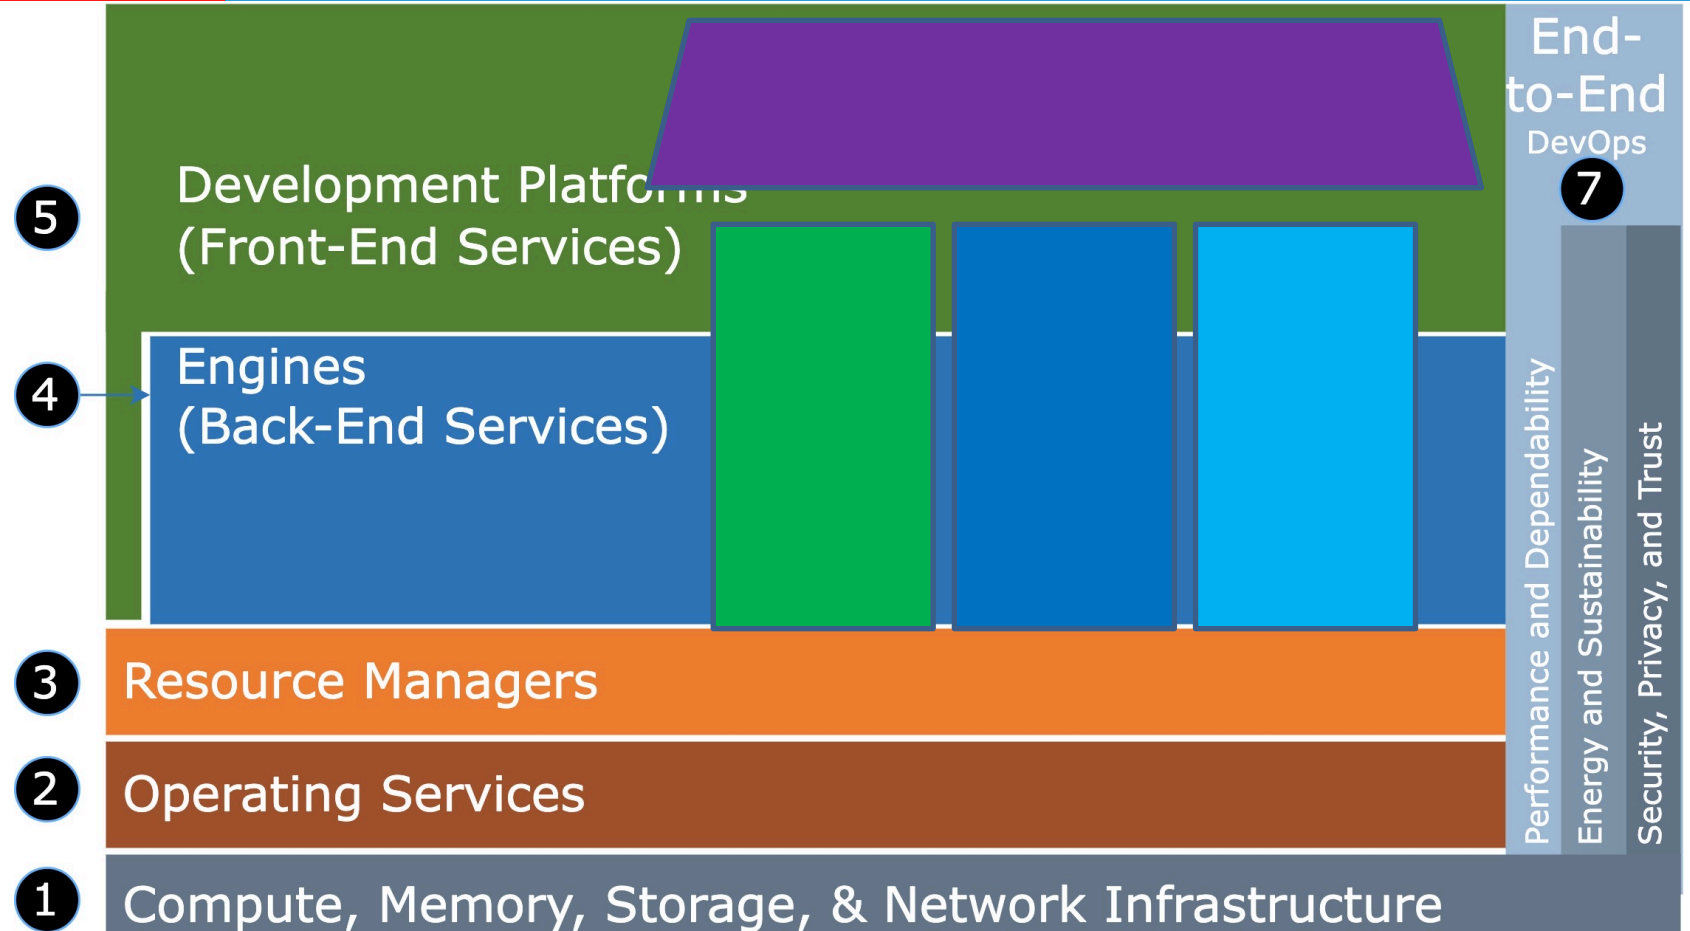
# HOW TO MANAGE SYSTEM COMPLEXITY?

## THE COMPLEXITY CHALLENGE

IOSUP ET AL. REFERENCE ARCHITECTURE FOR DCS, 2016


Focus on DevOps + Applications,  
5 Core Layers:

5. Development (Front-end)
4. Runtime Engines (Back-end)
3. Resources
2. Operations Services
1. Infrastructure





# AI/ML/DL OPERATIONS FOR GAME SERVICES



ISSUES: COMPLEXITY,  
NON-TECHNICAL

IOSUP ET AL. 2021

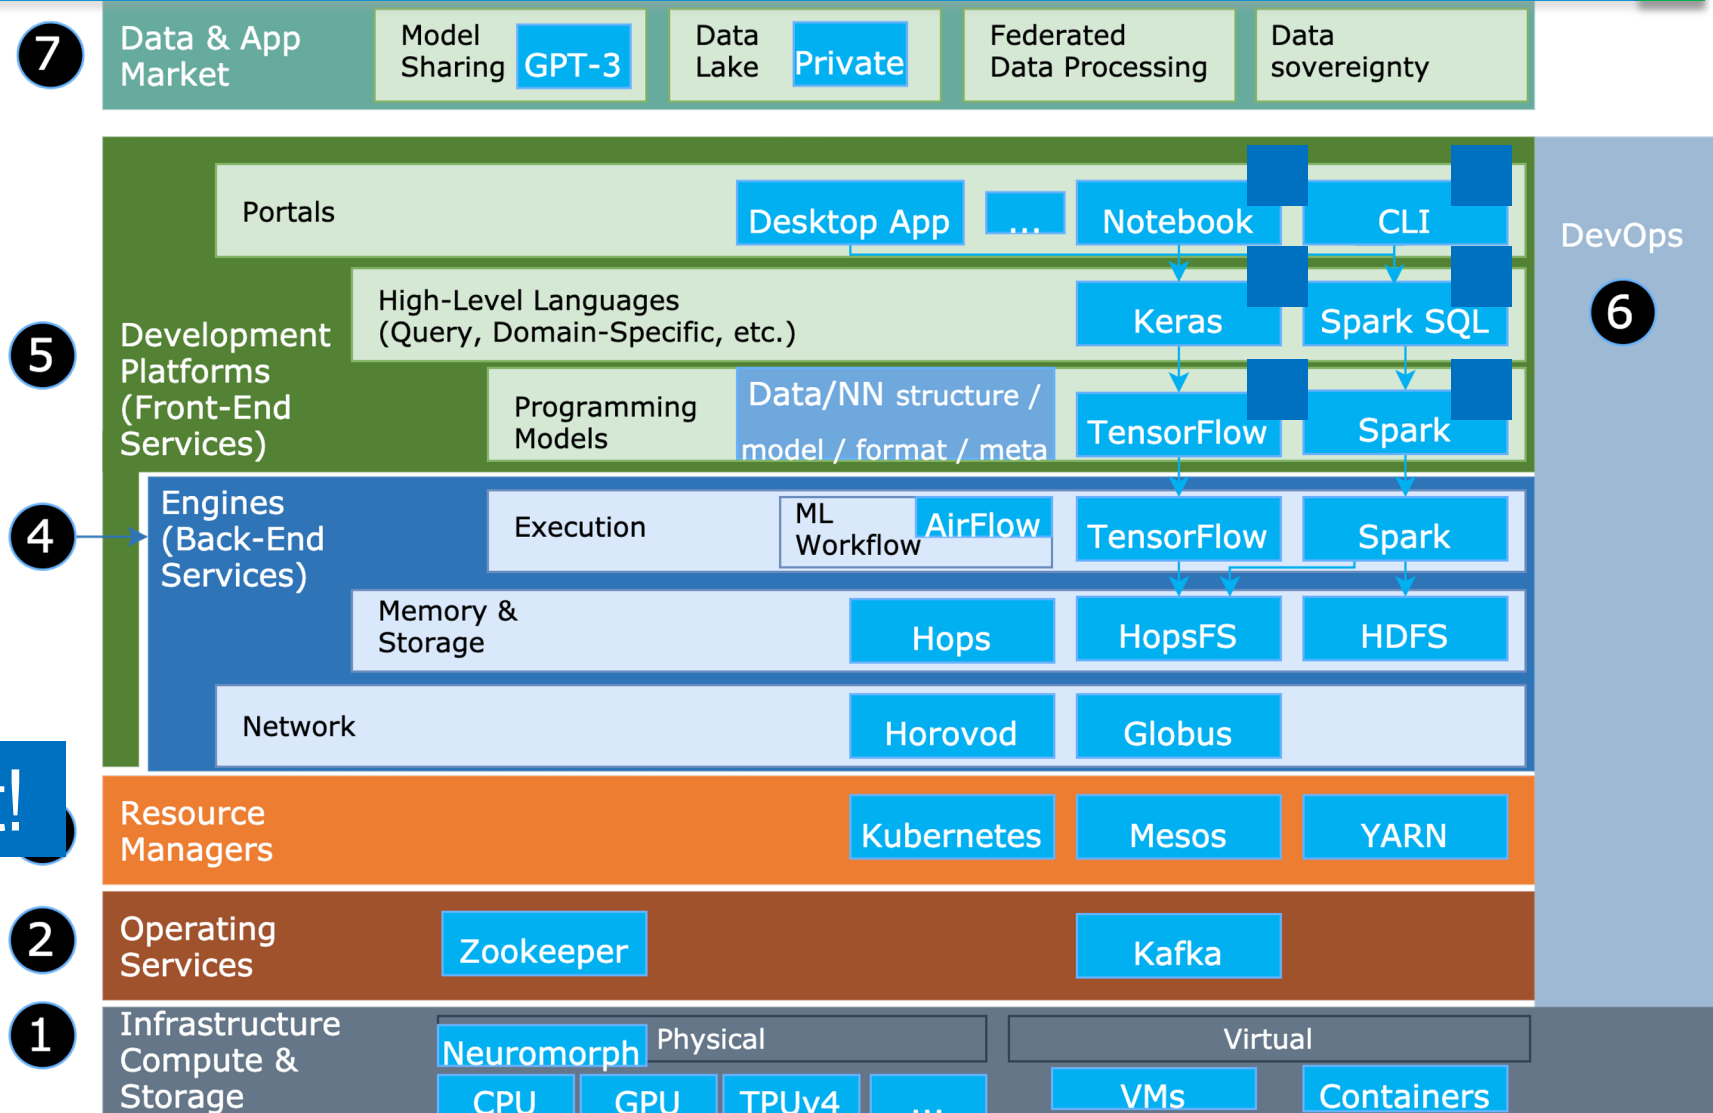
Gaming  
Analytics

2

ML app = small part!

Adapted from:

Sakr, Bonifati, Voigt, Iosup, et al. (2021) The Future Is Big Graphs! CACM.





Message: **4<sub>B</sub>**

Observe, measure  
ecosystem  
operation  
(mostly dynamics)



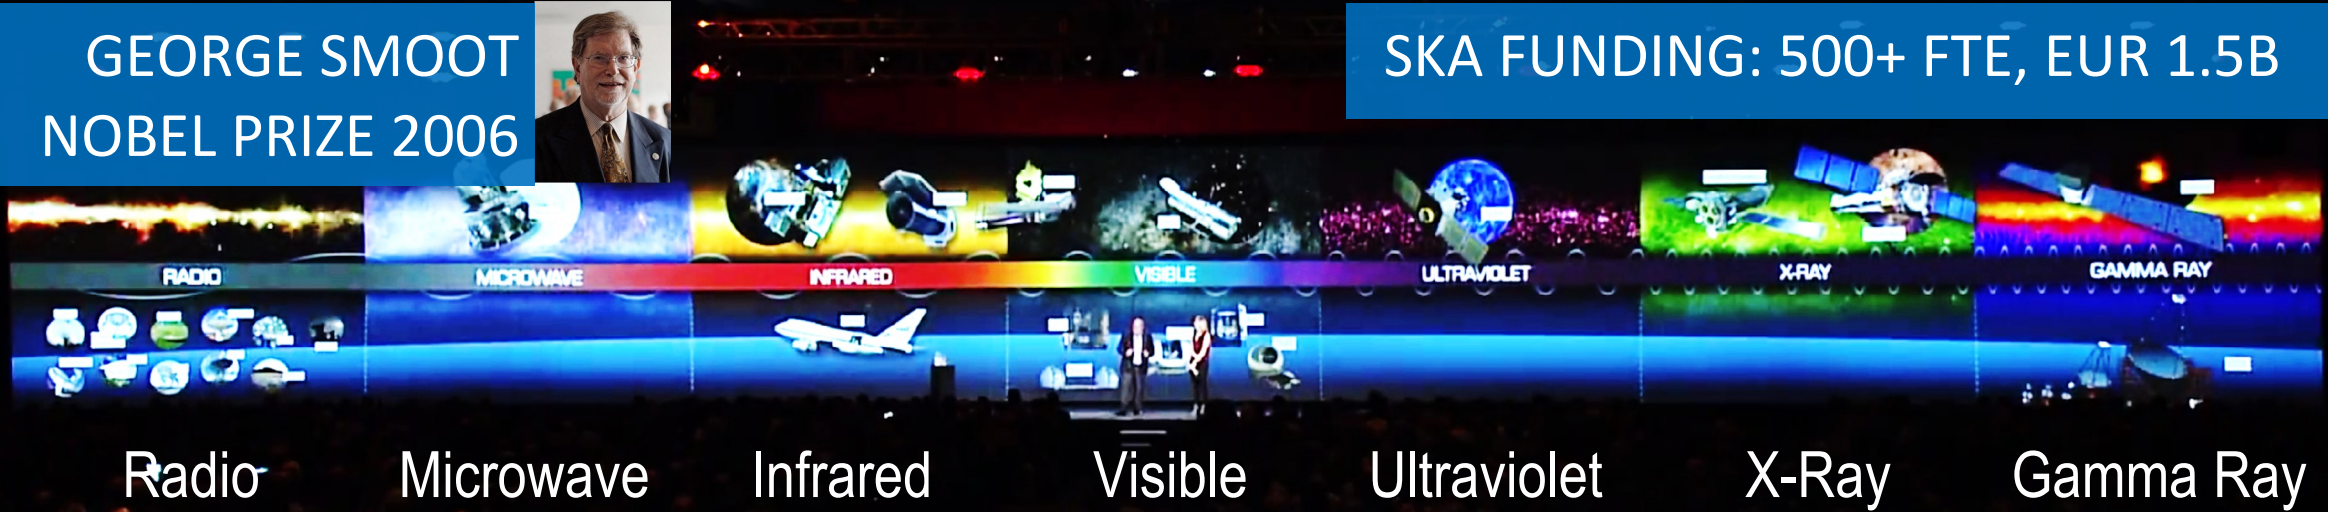
# DISCOVERY = LARGE-SCALE, LONG-TERM STUDY

UNCOVERING THE MYSTERIES OF OUR PHYSICAL UNIVERSE

GEORGE SMOOT  
NOBEL PRIZE 2006



SKA FUNDING: 500+ FTE, EUR 1.5B



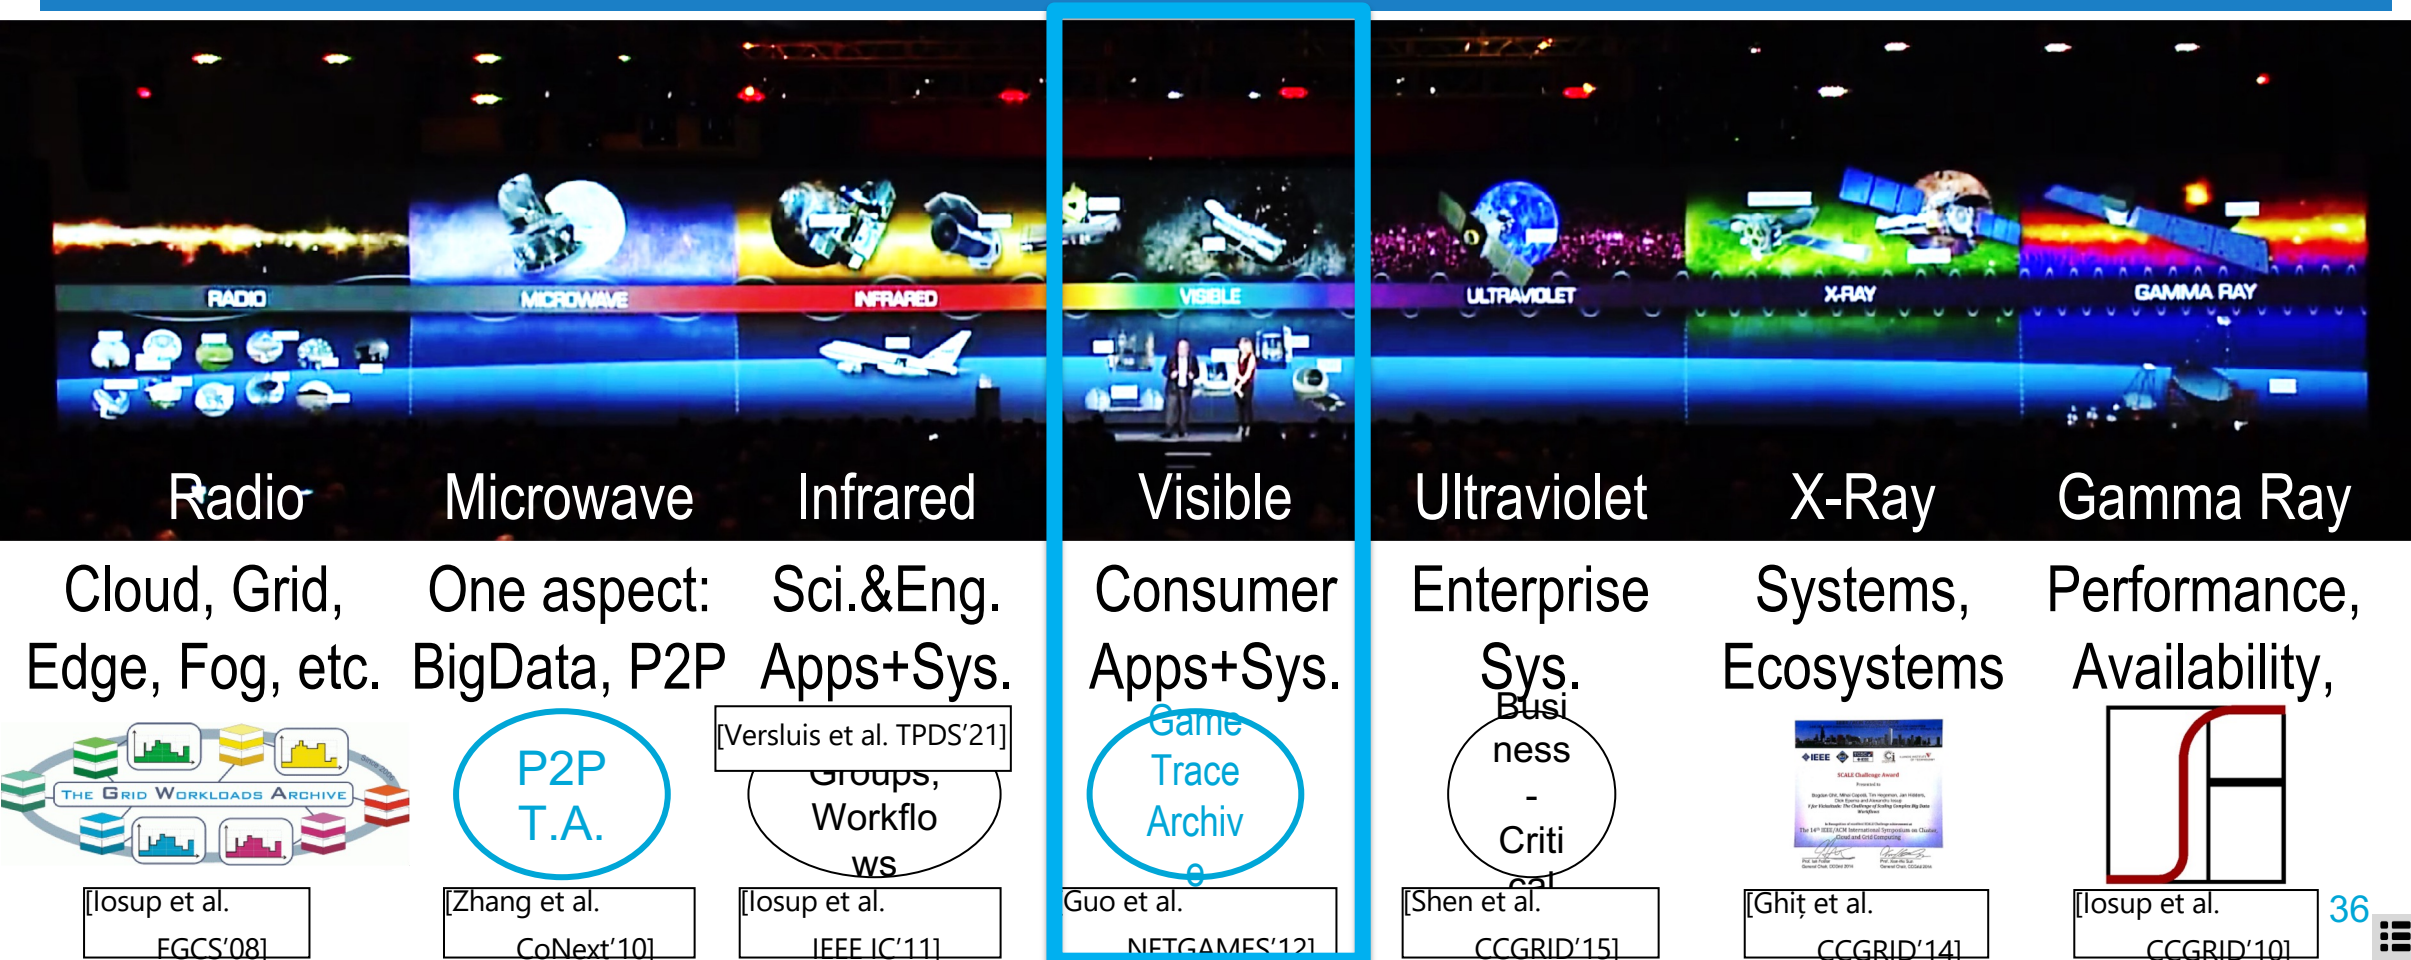
James Cordes, The Square Kilometer Array, Project Description, 2009 [Online]

The Square Kilometer Array Factsheet, How much will it cost?, 2012 [Online]

Phil Diamond and Rosie Bolton, Life, the Universe & Computing: The story of the SKA Telescope, SC17 Keynote. [Online]

# DISCOVERY = LARGE-SCALE, LONG-TERM STUDY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL



# OUR VISION: THE DISTRIBUTED SYSTEMS MEMEX

Bush (1945) As we may think. The Atlantic, Jul 1945.

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

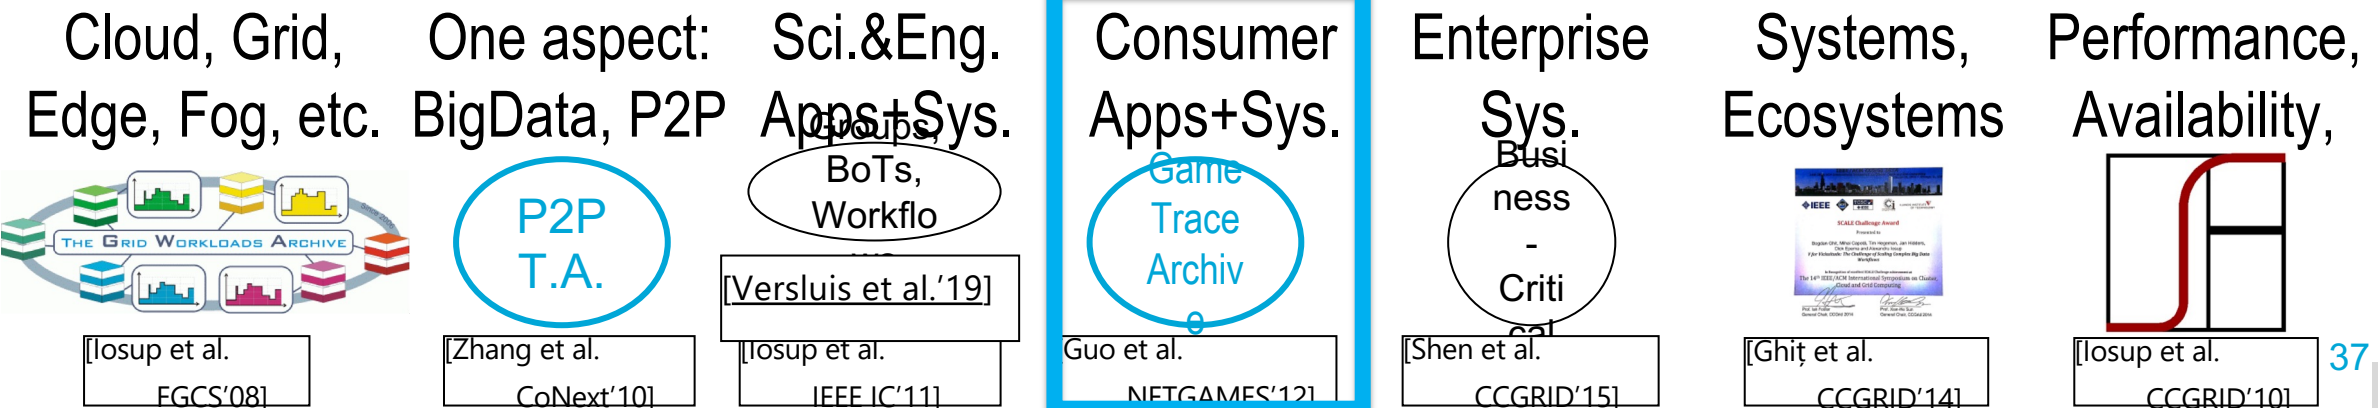
Find and eradicate performance issues

Quantitative evidence

Enable new designs and automation

Cultural and ethical concerns



Understand how entire ecosystems behave and evolve



# UNKNOWN PHENOMENA: INTER-, ADAPT-, EXAPTATION

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

SOME OF OUR DISCOVERIES



BOTS, NOT  
PARALLEL JOBS

GROUPS NOT  
RARE, DOMINANT

COMMUNITY  
DYNAMICS

SYSTEMIC  
VARIABILITY

CORRELATED,  
NOT IID FAILURES

Cloud, Grid,  
Edge, Fog, etc.

One aspect:  
BigData, P2P

Sci.&Eng.  
Apps+Sys.

Consumer  
Apps+Sys.

Enterprise  
Sys.

Systems,  
Ecosystems

Performance,  
Availability, etc.



[Iosup et al.  
FGCS'08]

[Zhang et al.  
CoNext'10]

[Iosup et al.  
IEEE IC'11]

[Guo et al.  
NETGAMES'12]

[Shen et al.  
CCGRID'15]

[Ghit et al.  
CCGRID'14]

[Iosup et al.  
CCGRID'10]






# One Basic Result




Uta et al., Beneath the SURFace: An MRI-like View into the Life of a 21st-Century Datacenter. USENIX; login 2020.



Hot GPUs

Cold CPUs\*

\* when GPUs



Could use less powerful,  
less expensive CPUs

Message: 4c



Experiment with the  
ecosystem, real-  
world + simulation




# EXPERIMENTAL METHODS OF DISCOVERY

UNIQUE OPPORTUNITY: WE DRINK OUR OWN CHAMPAGNE (*IN VIVO*)!

Our Prototypes (*in physico/in vitro*)




SURFnet6



Alex Uta



Georgios  
Andreadis




Fabian  
Mastebroek



Vishal  
Kachheendra



Maria  
Vein




Laurens  
Versluis



We also use clouds




And simulators (*in silico*)



Alexey Ilyushkin



# Minecraft only scales to hundreds of players



# ... CAN WE AFFORD X? WHAT IF Y? A vs. B ... vs. Z?

TOO COSTLY TO CONDUCT REAL-WORLD EXPERIMENTS, SO WE BUILT A SIMULATOR



OpenDC  
simulator



Learn more:  
[opendc.org](https://opendc.org)



- Short-term resource management
- Long-term capacity planning
- Sophisticated model → many Qs, goals
- Supports many kinds of workloads
- Supports many kinds of resources
- Validated for various scenarios
- Work with major NL hoster
- Used in training, education, research

Fabian Mastenbroek



*and more...*

# Electricity Expenses become High Risk



Electricity expenses become **primary risk** in datacenters



# Experiments are also Expensive!

## An **Environmental** Perspective

112

Core-hours

2.1

kWh



**In Simulation**

8,600,000,000

Core-hours


116,000,000,000

kWh



**In Reality**

*Exact numbers confidential, depend on topology*



# Massivizing Online Games = Rich challenge of computer systems, with societal impact!



1. This is the Golden Age of massive computer ecosystems
2. But we cannot take this technology for granted
3. A new science to help massivizing computer ecosystems
4. Understand how things work: observe, synthesize, experiment
5. Design and build better systems:  
virtual worlds, analytics, content generation, meta-gaming

6. What's next? PLATO in the 21<sup>st</sup> century

7. Beyond games: Massivizing computer systems is a national priority in the NL!

Lots to do  
next!



LET'S DESIGN

& BUILD IN

DISTRIBUTED



5

ECOSYSTEMS!





Message:

5<sub>A</sub>



Scaling ~ creative

palette of


techniques, design

and use wisely

# 5a SCALING IN AND OUT THE VIRTUAL WORLD OF A GAME

Goal: Build an efficient platform for massive scalability

1. Close to players
2. No upfront costs, no maintenance, pay for what is actually used
3. Compute platforms: multi-cores, GPUs, clusters, all-in-one!
4. Auto-scaling mechanisms and policies
5. Performance guarantees
6. Hybrid deployment model
7. Geo-distributed scheduling
8. Code for various compute platforms—platform profiling
9. Load prediction miscalculation costs real money
10. What are the services?
11. Vendor lock-in?
12. My data



Done

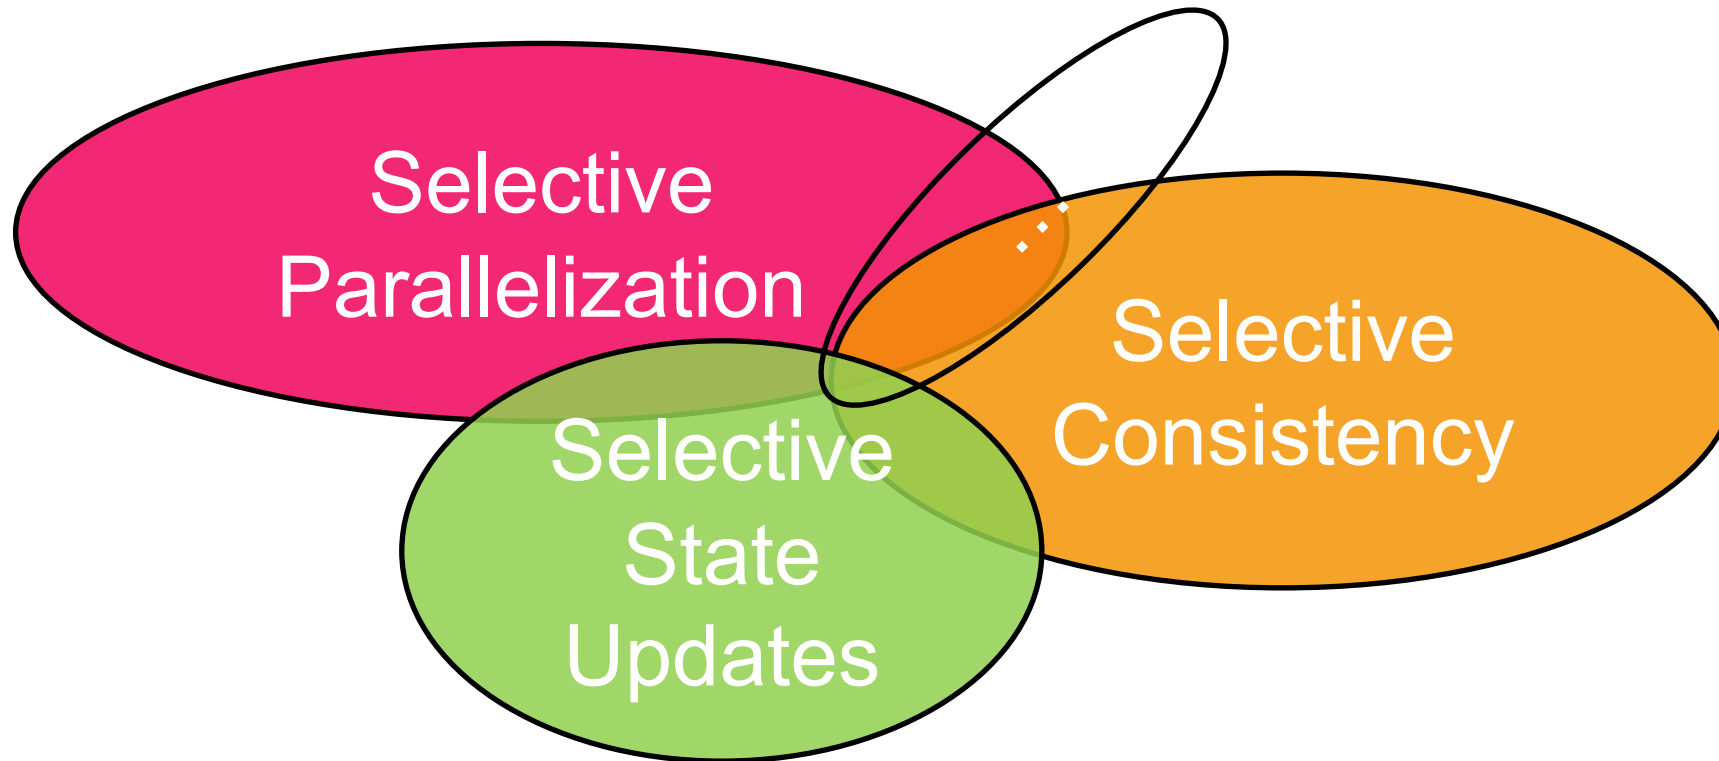
Open challenge

# HOW TO SCALE GAMES ACROSS THE ECOSYSTEM?

5A1


IT'S SCALING TECHNIQUES!

REFERENCE VIEW ON OPERATIONAL TECHNIQUES




# Selective Consistency

# CHALLENGE: CONSISTENCY IN HIGH-STAKES APPS




Tick: 1






# CHALLENGE: CONSISTENCY IN HIGH-STAKES APPS




Tick: 2


## Server



## Client

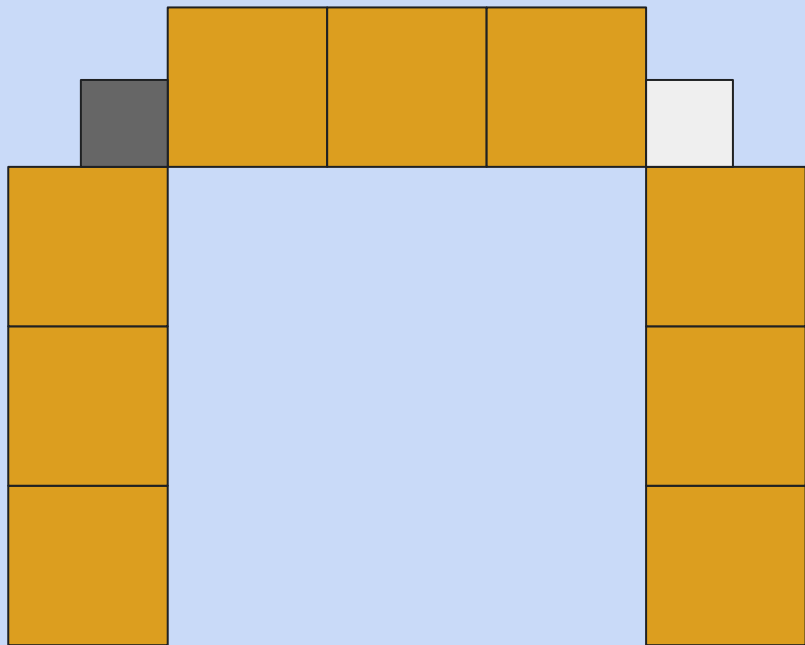


# CHALLENGE: CONSISTENCY IN HIGH-STAKES APPS




Tick: 3

## Server



## Client



# Dyconits: Scaling Minecraft-like Services through Dynamically Managed Inconsistency

Jesse Donkervliet, Jim Cuijpers, Alexandru Iosup



[jesse.donkervliet@vu.nl](mailto:jesse.donkervliet@vu.nl)

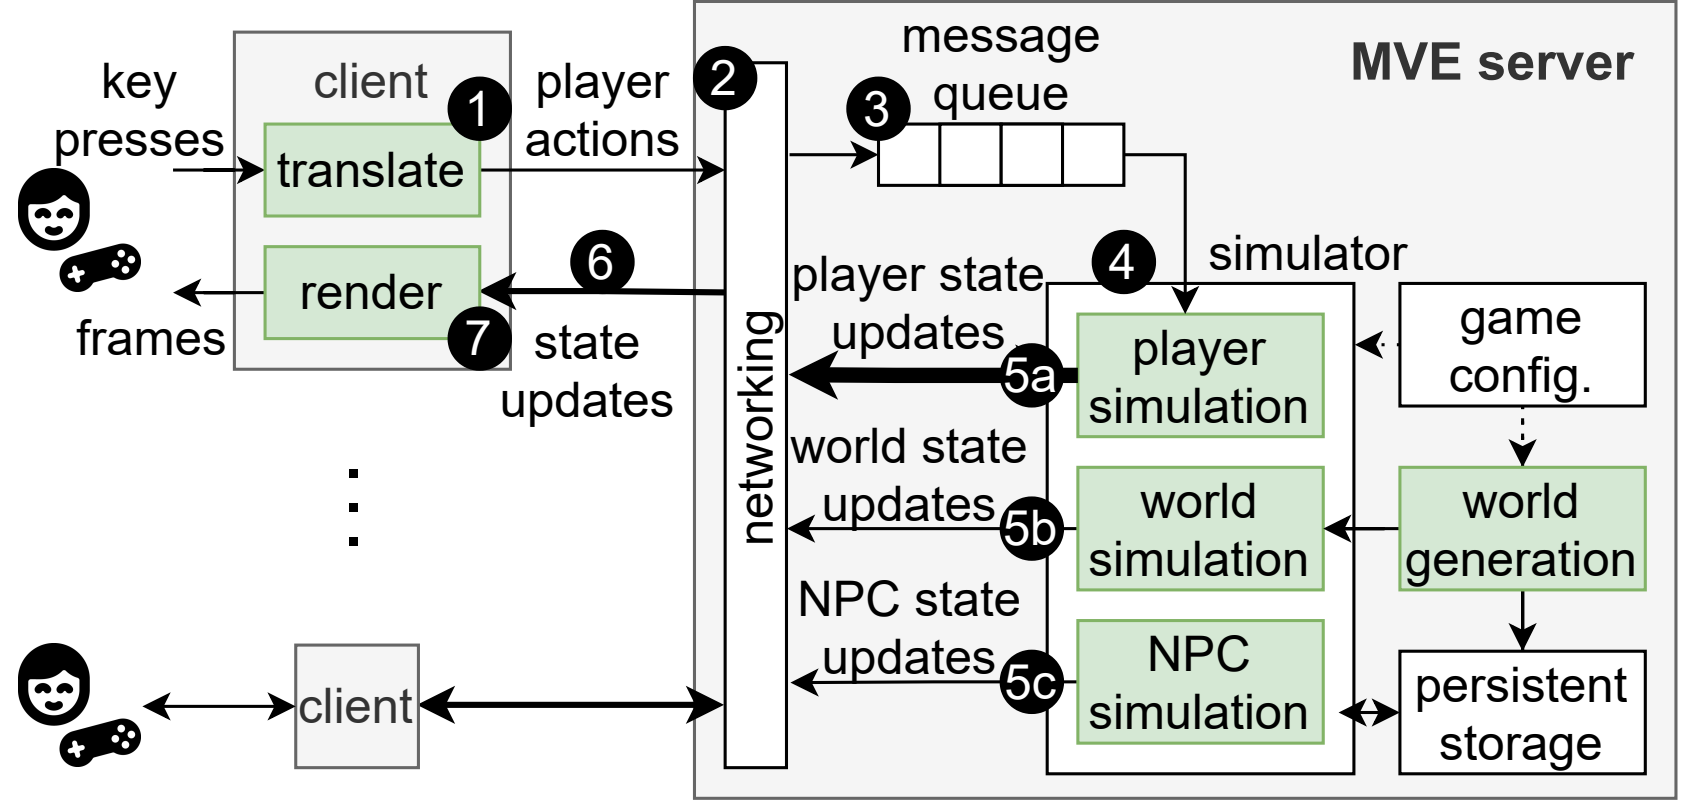


[@jdonkervliet](https://twitter.com/jdonkervliet)



<https://atlarge-research.com/opencraft/>

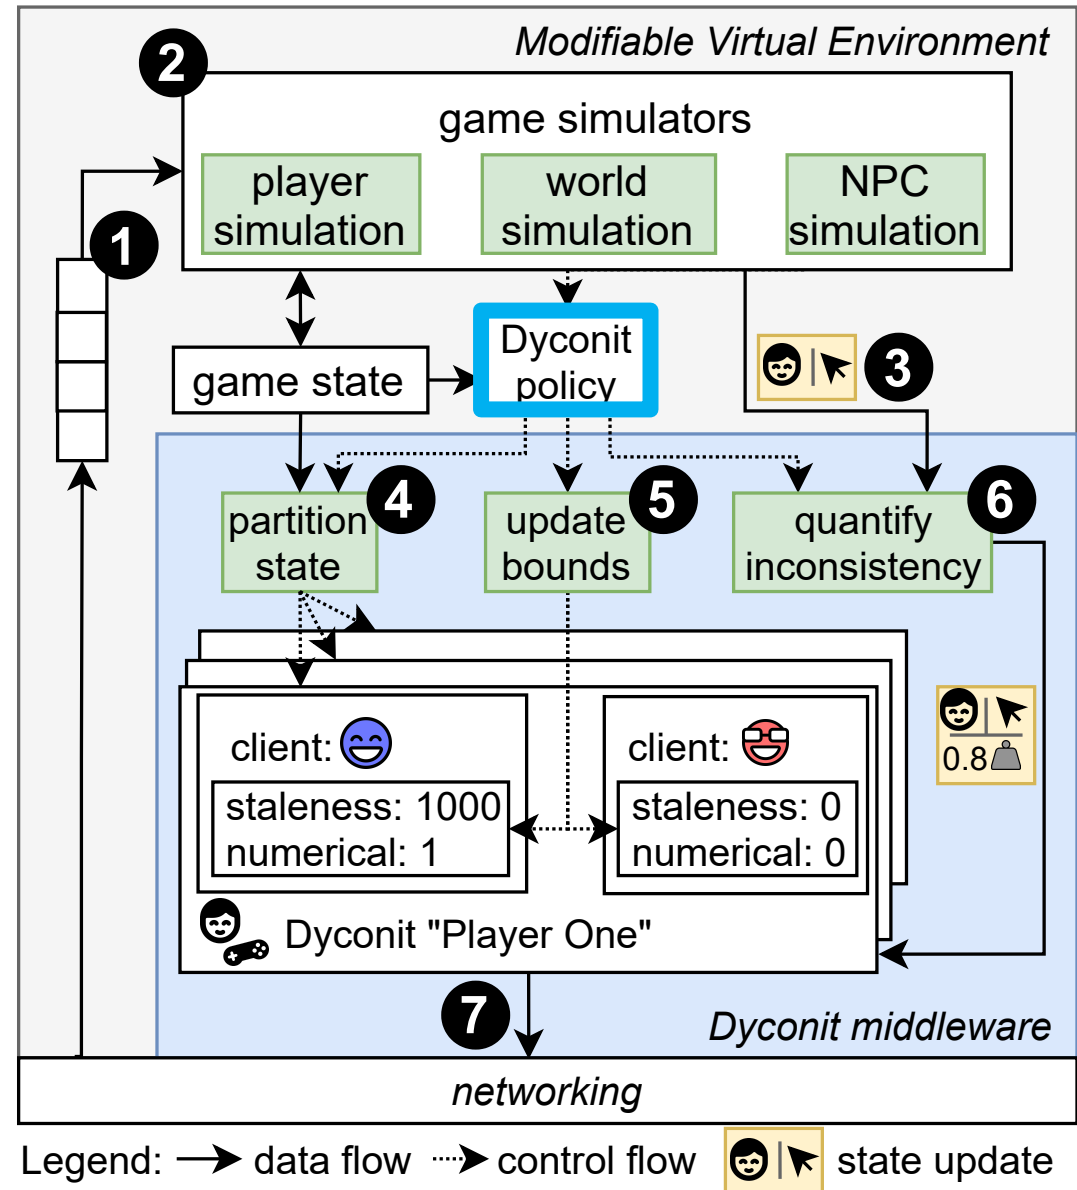
# MVE System Model





# Design of the Dyconit Middleware

1. Fundamental DS concept  
**Conits**
2. Mechanism to deliver function/SLA  
**Dyconits**
3. Technique for short-term dynamics  
**Adaptive dyconit management, with policies**



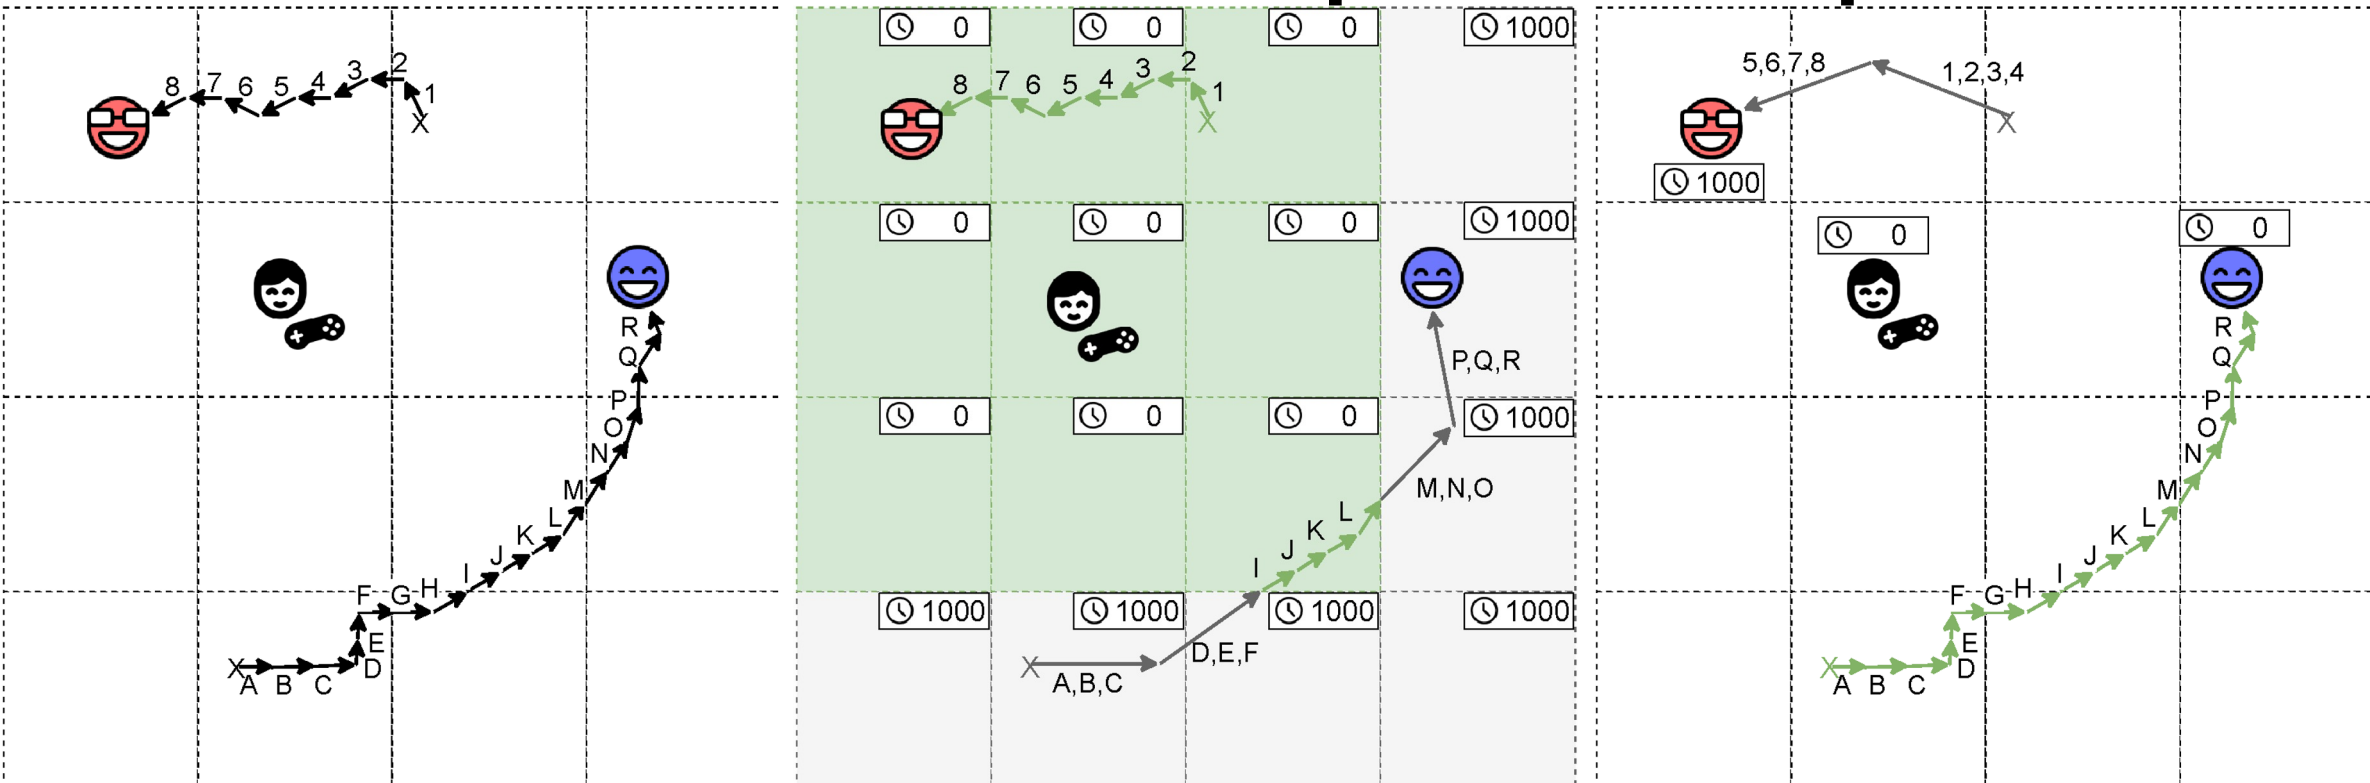

# Dyconits: Using the conit consistency model

Quantify inconsistency in the system using three metrics

1. **Staleness** (how old is this update?)
  2. **Numerical error**  
(how large is the impact of this update?)
  3. **Order error**  
(how does this update relate to other updates?)
- + **DyConit = conit + dynamic conit management**  
(change bounds at runtime, based on actual gameplay)

# Dyconit Policies Here be algorithms!

## Use Feedback to Improve Operation




(a) No Dyconits

(b) Area Of Interest policy

(b) Interest Set policy

# Dyconit Policies Improve Scalability



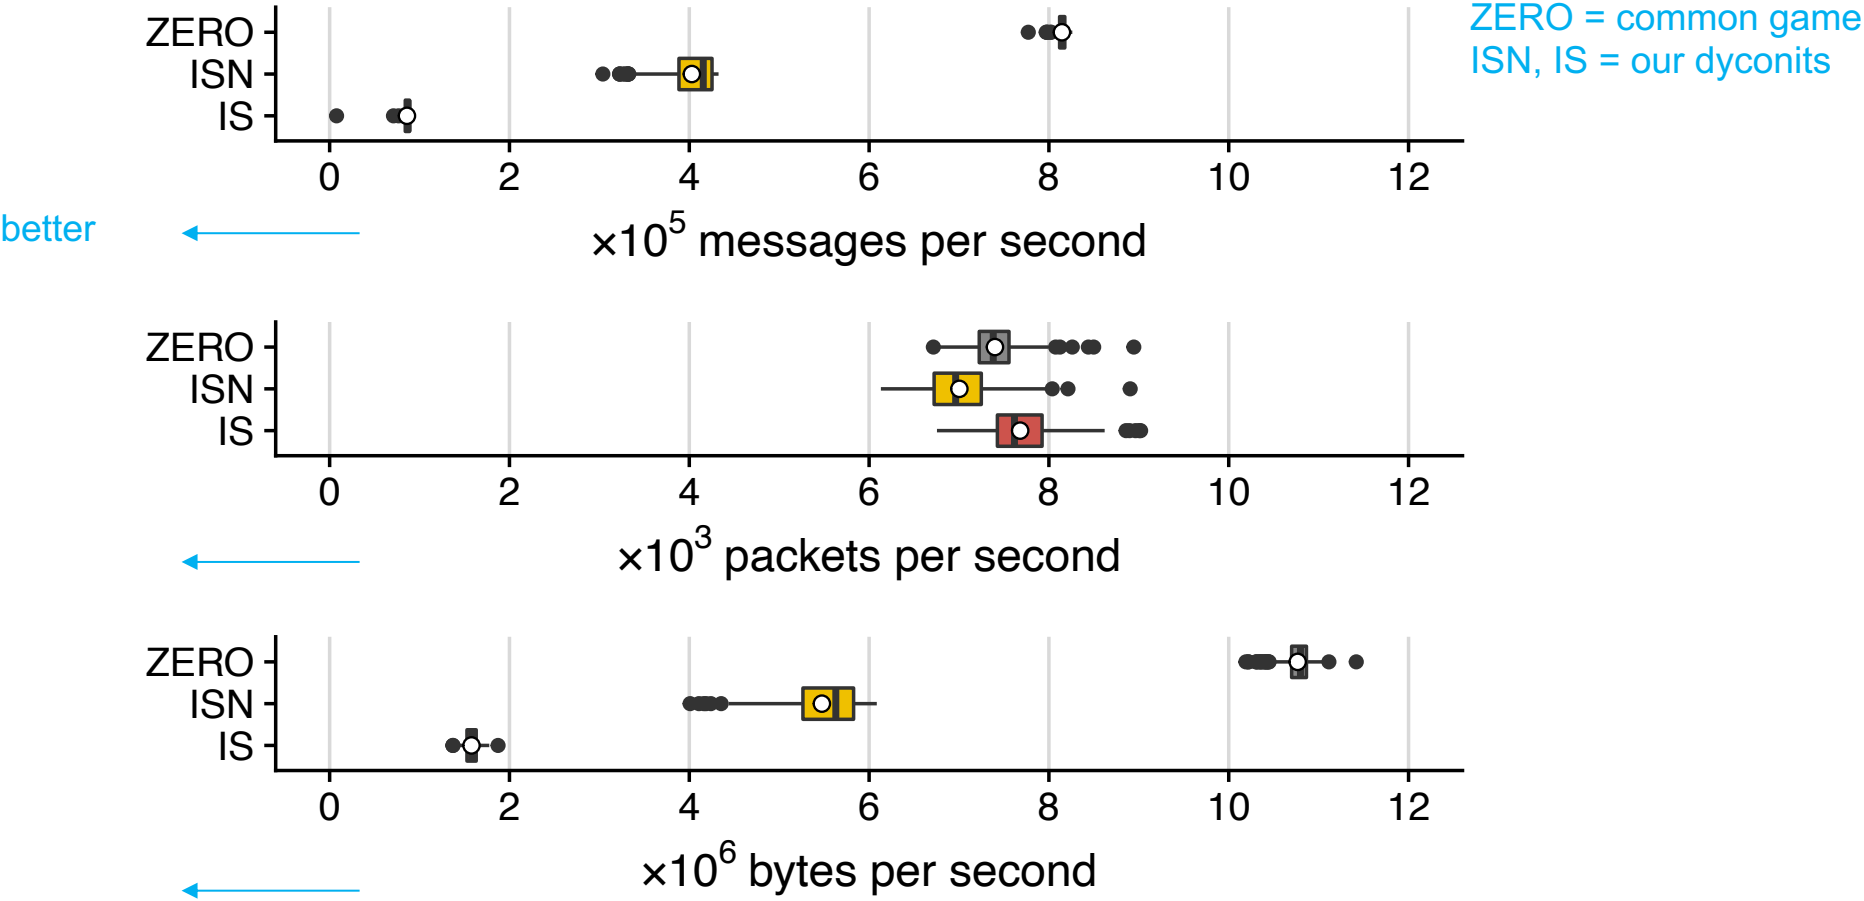
(a) No Dyconits

(b) Area Of Interest policy




(b) Interest Set policy



# Dyconits Reduce Bandwidth Usage




# Dyconits Bound Inconsistency



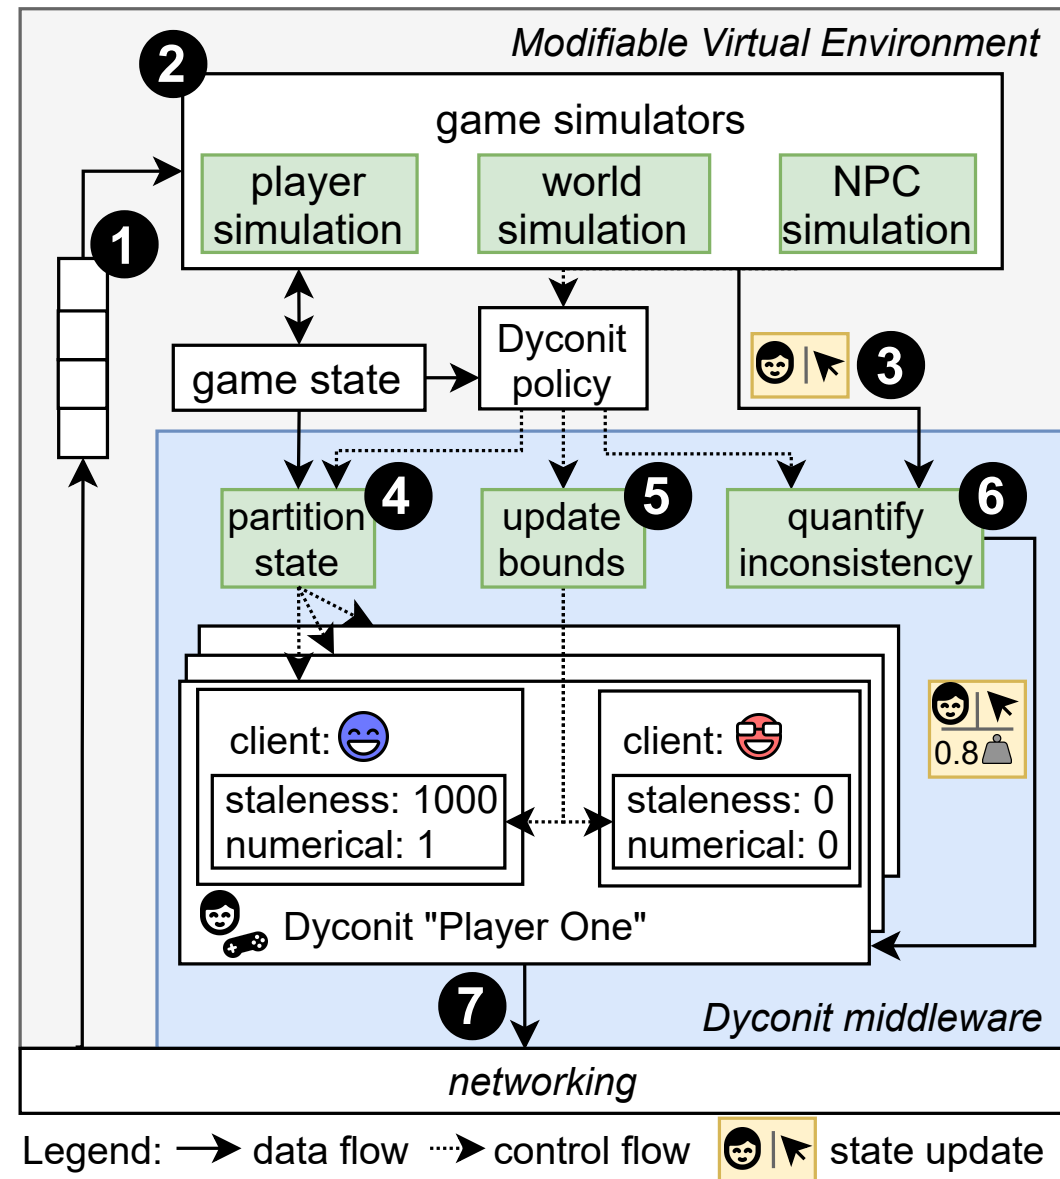
ZERO = common game  
ISN, IS = our dyconits

# Dyconits Can Dynamically Trade-off Consistency for Performance



# Main Contribution



1. Design of *Dyconits* to address scalability issues
2. Realization of a Minecraft-like game using Dyconits; Game and Dyconits code publicly available
3. Real-world experiments to evaluate scalability






# Selective Parallelization

# Game parallelization models




- **Zoning**: huge game-world division into geographical sub-zones – each zone is handled by different machines
- **Mirroring**: the same game-world handled by different machines, each one handling a subset of all entities
- **Instancing/sharding**: multiple instances of the same zone with independent states



# Selective State Updates

# Remember RTS Games?



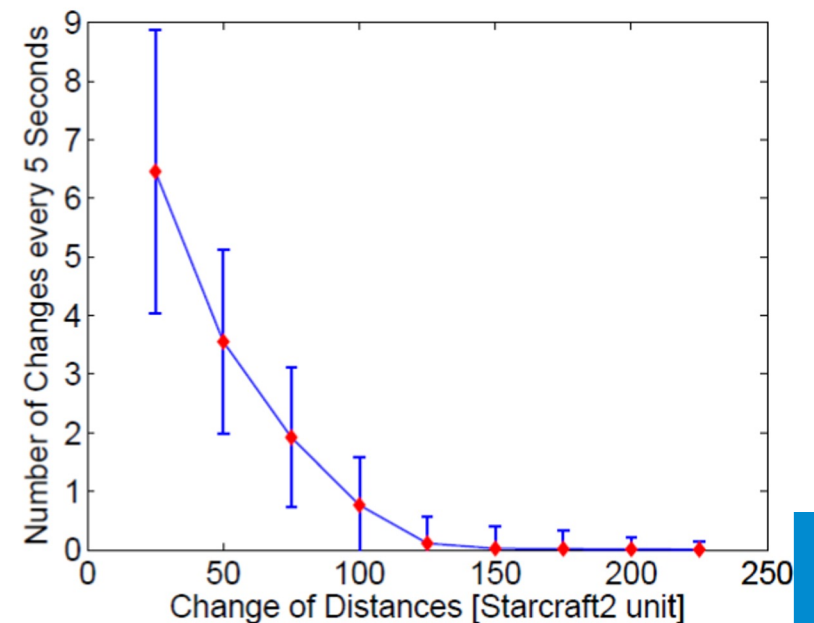

- Players control tens up to hundreds of units.
- Players need to take decisions in real-time.






# Traditional Aol does **not** work

- Area of Interest (Aol) = traditional mechanism for RPG: only receive 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**




# 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, player receives different types of information (**commands** or **state**) about game world
- Each player can have multiple Areas of Simulation (AoS) and Areas of Update (AoU), and also multiple No-Update Areas (NUA)

# Experimental results







- Simulator and prototype RTS game
- Evaluate in two Cloud platforms: EC2 and Azure
- Prototype about 20k lines of C++ code
  - Based on an open source game (~6k lines)
- Up to 200 players and **10,000 battle units**
  - **State-of-the-art unplayable at 1-2,000**
  - **Crashes not uncommon due to CPU and Network bottlenecks**
- Using our AoS-based method can lead to  
**lower CPU** consumption than pure event-based method (RTS) and  
**lower network consumption** than pure update-based method (RPG)

# Experimental results

AoS can lead to lower resource use per player



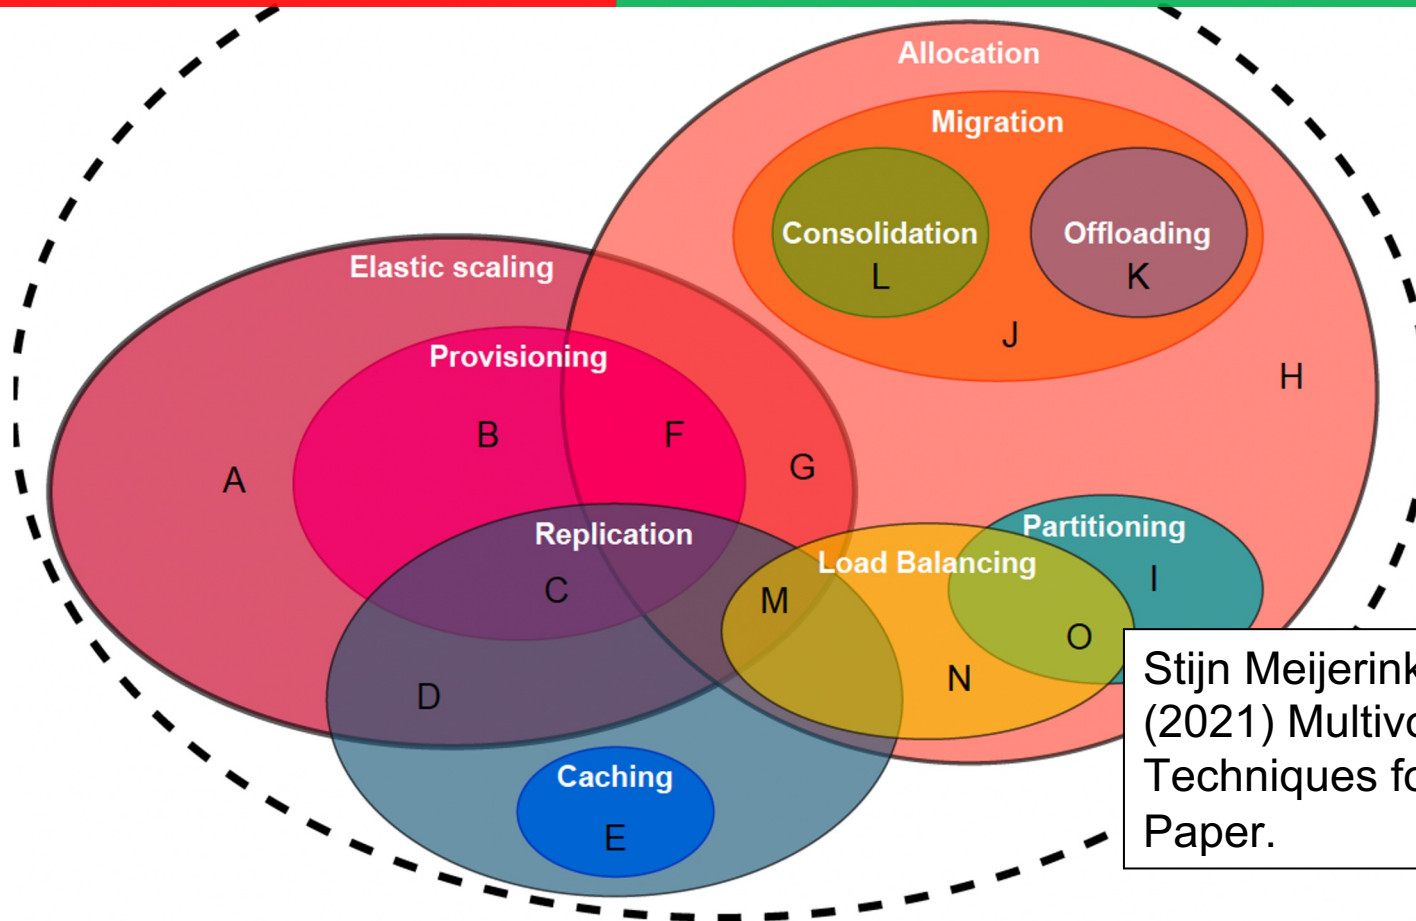
Allows **better scalability** than current RTS and RPG games



# HOW TO AUTOMATE X ACROSS THE ECOSYSTEM?

IT'S OPERATIONS!

REFERENCE VIEW ON OPERATIONAL TECHNIQUES



Stijn Meijerink, Erwin van Eyk, Alexandru Iosup (2021) Multivocal Survey of Operational Techniques for Serverless Computing. White Paper.

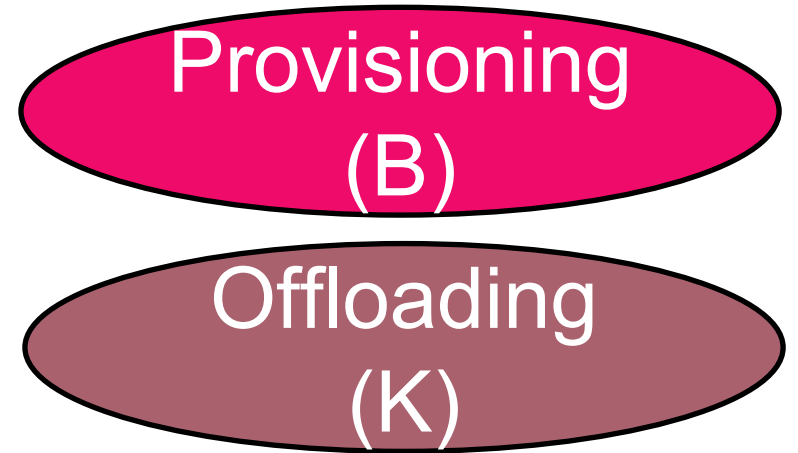
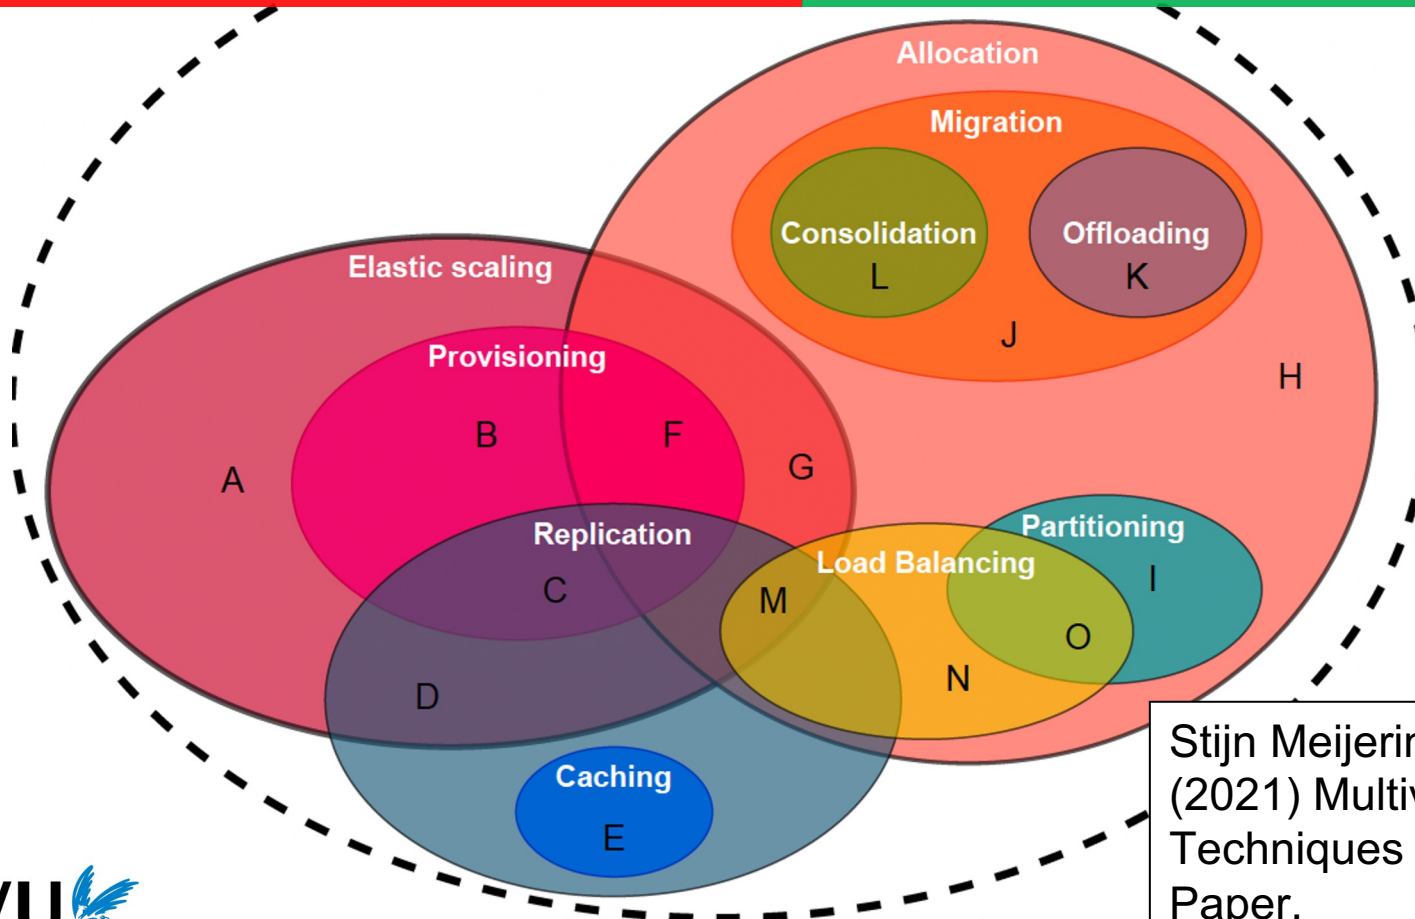


# HOW TO AUTOMATE OPS ACROSS THE ECOSYSTEM

5A2

IT'S OPERATIONS!


REFERENCE VIEW ON OPERATIONAL TECHNIQUES



Stijn Meijerink, Erwin van Eyk, Alexandru Iosup (2021) Multivocal Survey of Operational Techniques for Serverless Computing. White Paper.

# Provisioning (B)

# World of Warcraft, a Traditional HPC User?!



(since 2003)

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



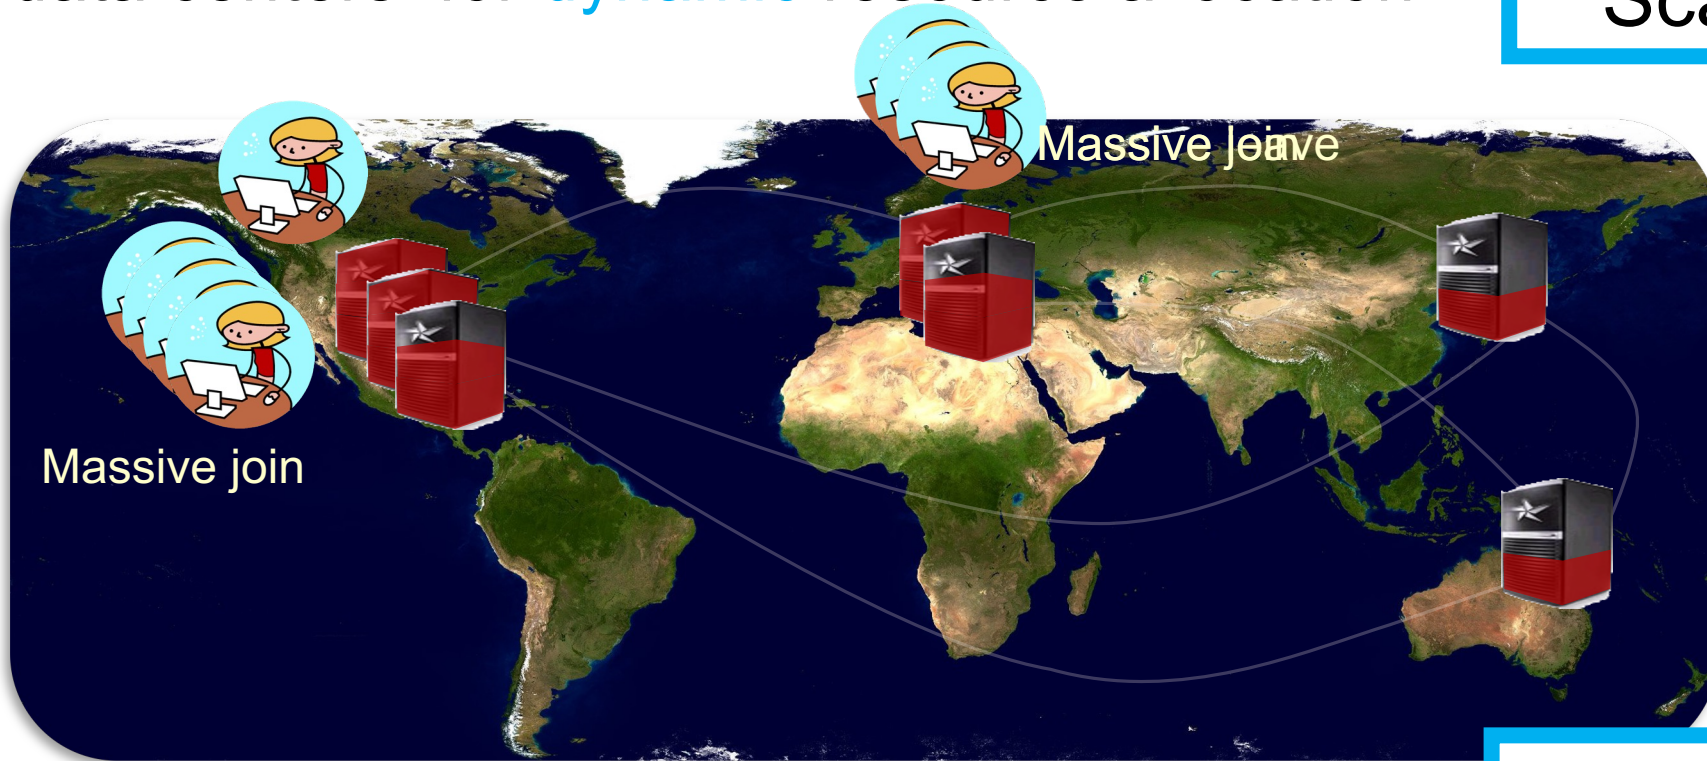


# Cloud-based hosting model

- Using data centers for **dynamic** resource allocation

Inter-operable

Scalability



PubNub



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

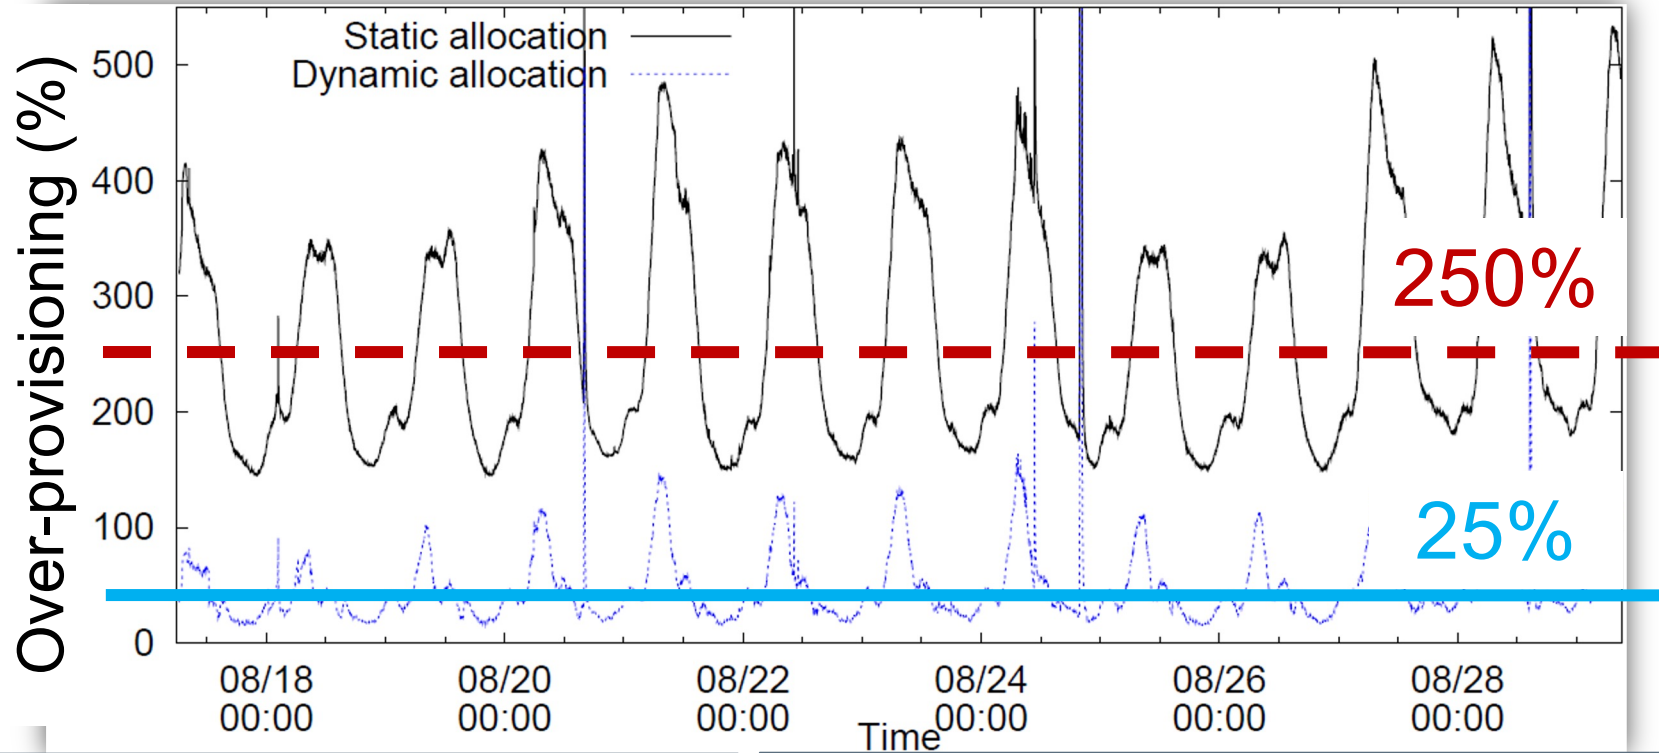
Autonomous

Reusable

# Resource Provisioning and Allocation

## Static vs. Dynamic Provisioning

[Source: Nae, Iosup, and Prodan, ACM SC 2008 / TPDS 2011]



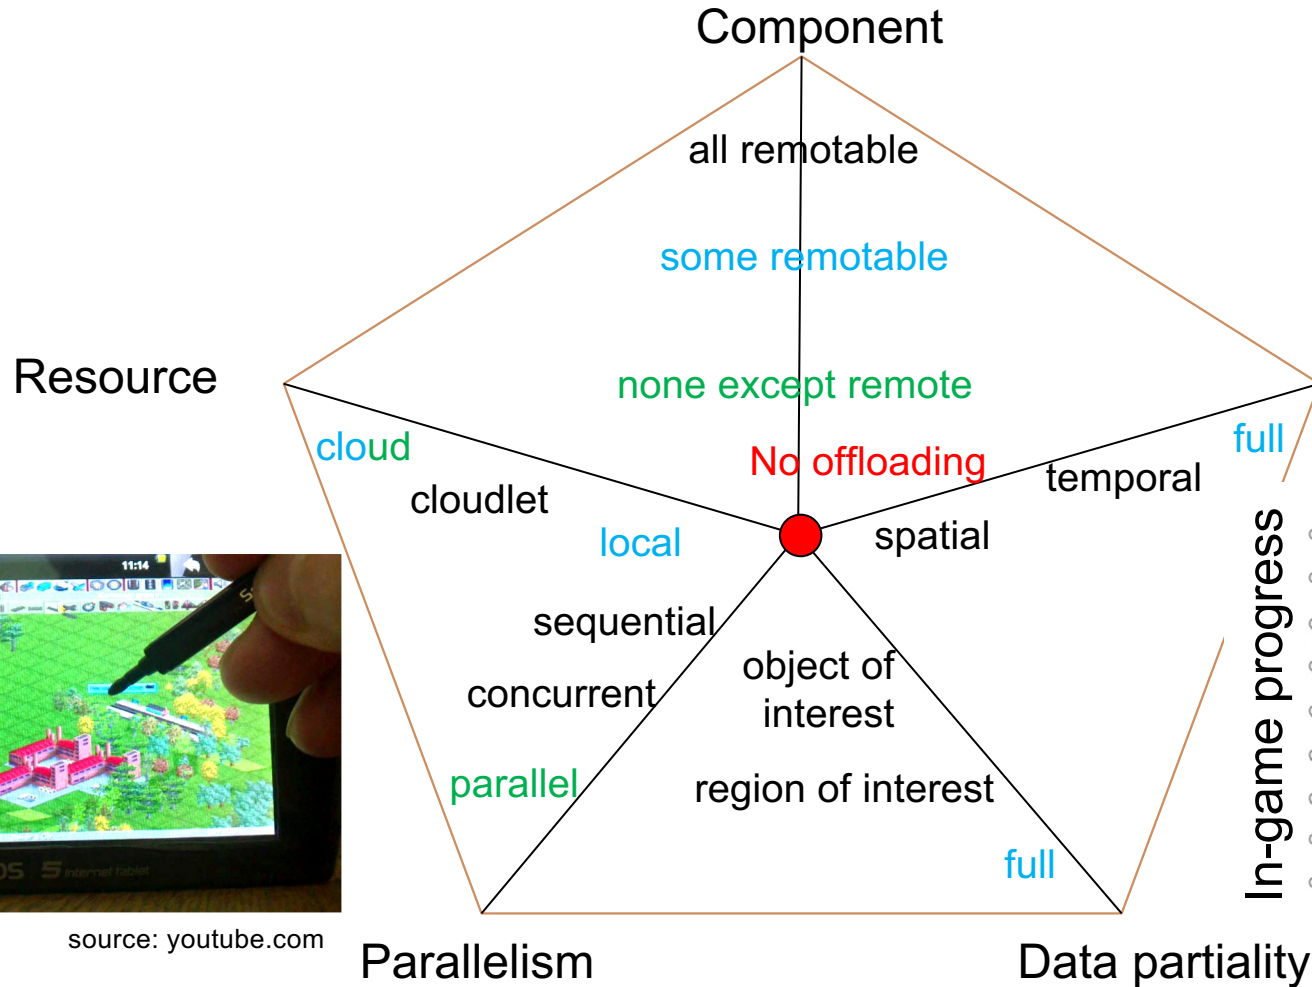
Consolidation is important!

... but need to rethink metrics, prediction, sched., OS, etc.



Offloading (K)

## Scaling through Decomposition □ True “Cloud Gaming” Exploratory Space for Offloading

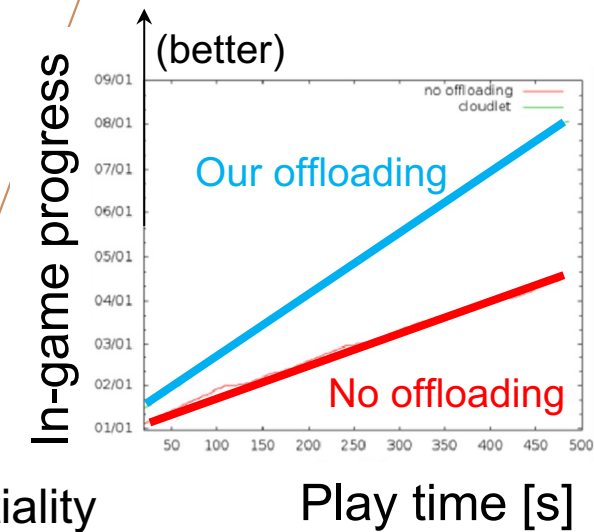


Cloud gaming / SuperServer:  
OnLive,  
Gaikai,  
Sony Online,  
Microsoft Live,  
etc.

Process partiality



source: youtube.com

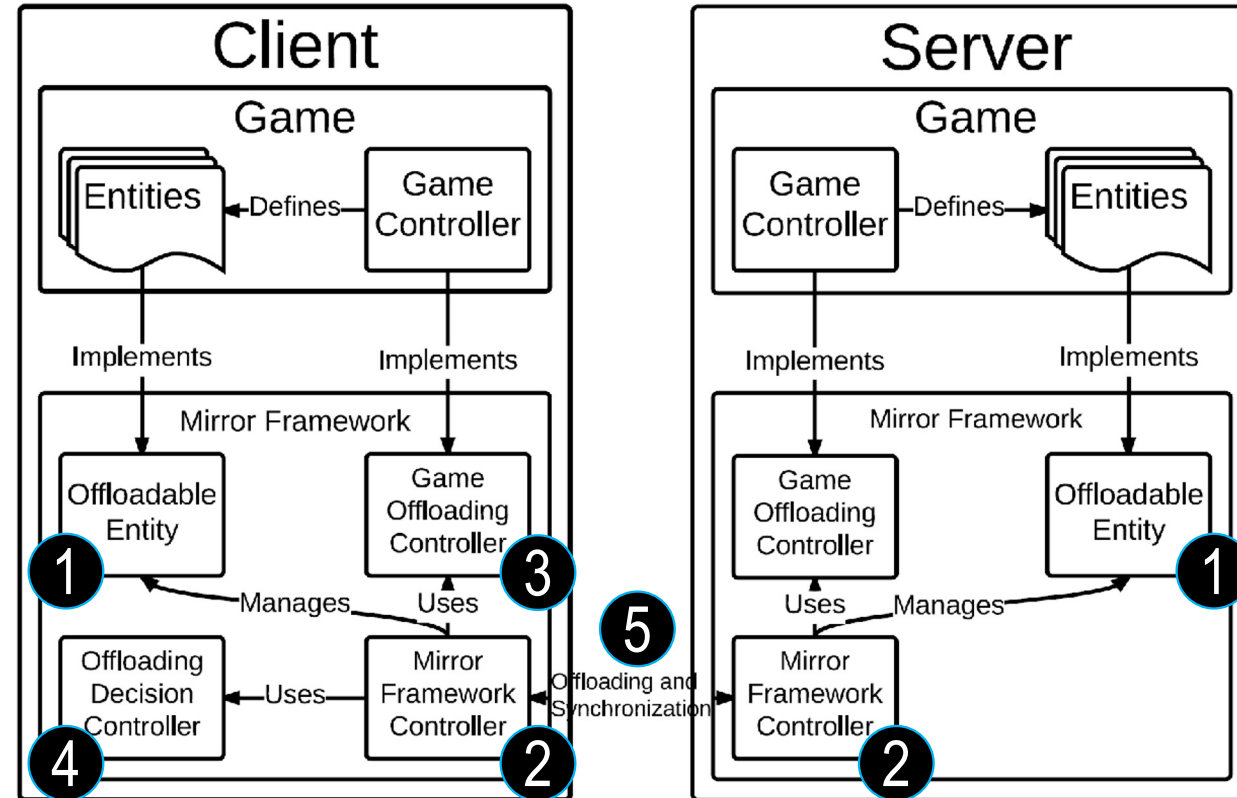


A.-C. Olteanu, N. Tapus, A. Iosup: Extending the Capabilities of Mobile Devices for Online Social Applications through Cloud Offloading. CCGRID 2013: 160-163.

# The Mirror Architecture for Game Offloading

- 1 Offloadable Entity**
  - (A part of) a game-object class
  - Mirrored on client and server
- 2 Mirror Framework Controller**
  - Manages offloading sessions, all
- 3 Game Offloading Controller**
  - Exposes main game-loop as API
- 4 Offloading Decision Controller**
  - Only on client
- 5 Offloading + Synchronization Protocols**

## Scalability






M. H. Jiang, O. W. Visser, I. S. W. B. Prasetya, Alexandru Iosup:

Mirror: A computation-offloading framework for sophisticated mobile games. CCPE 2018.

# Performance: use OpenTTD+Mirror

## Scalability



OpenTTD+Mirror+Offload All delivers significantly better performance than OpenTTD without offloading

# Portfolio Scheduling: Allocation (H)



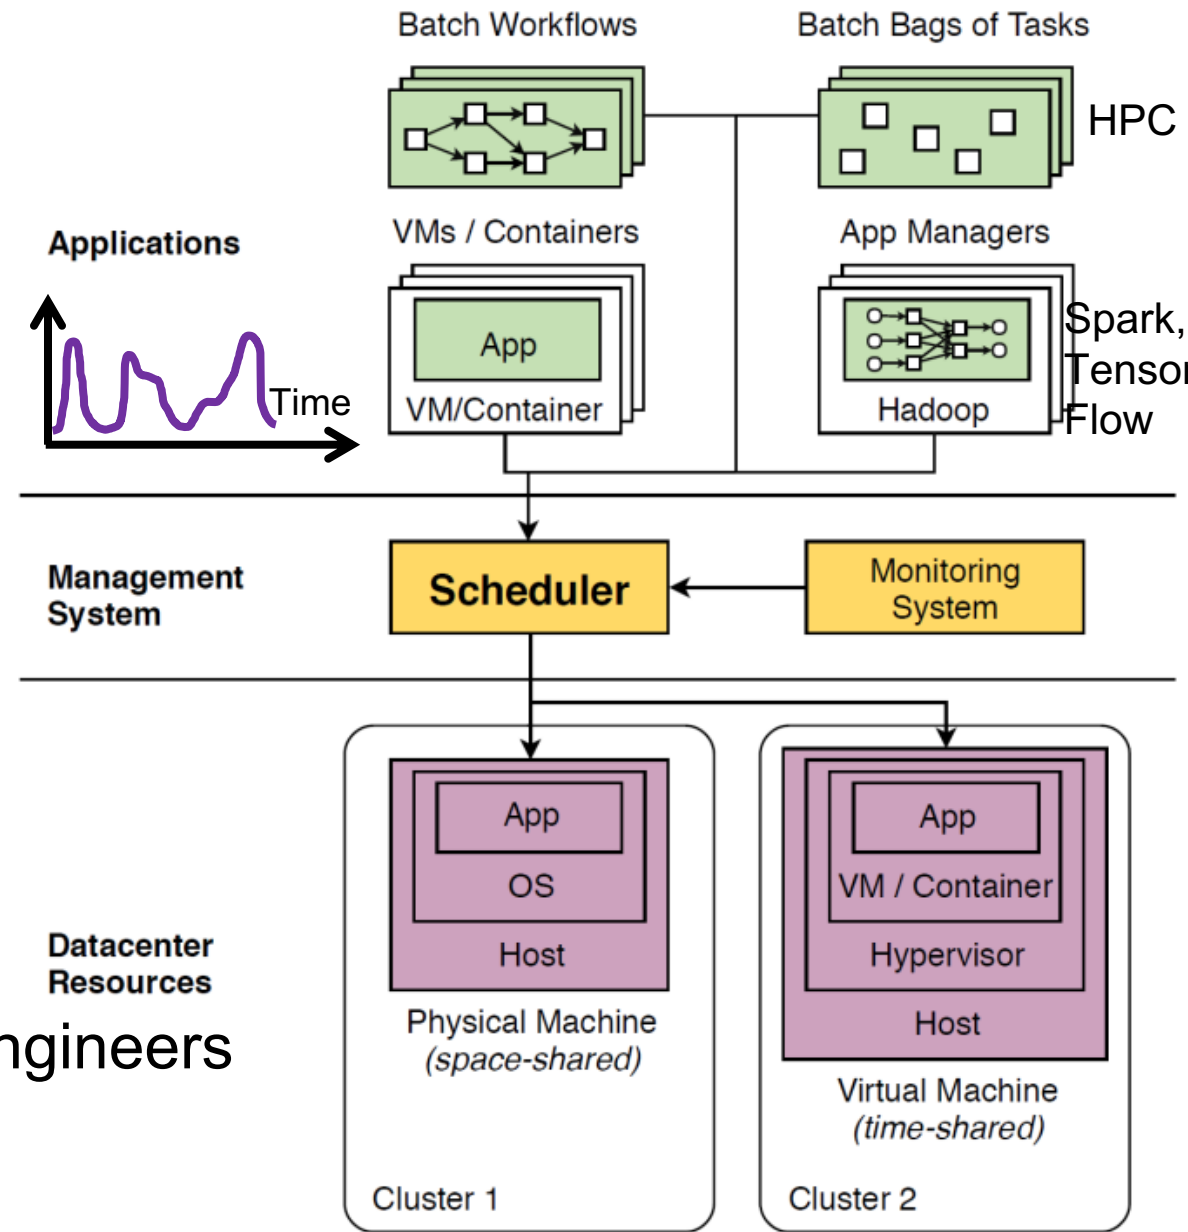
# The Technical Version

Scheduler needs to

- Make and take decisions
- (classic theory helps here)


but also to

- Use system feedback (Monitoring, possibly also eng.)
- Address workload and resource heterogeneity, dynamicity, structure / topology, etc.
- (2010s+) Explain decisions to engineers



# Portfolio Scheduling, In A Nutsh

- Datacenters cannot work without one or even several schedulers
- Instead of ephemeral, risky schedulers, we propose to

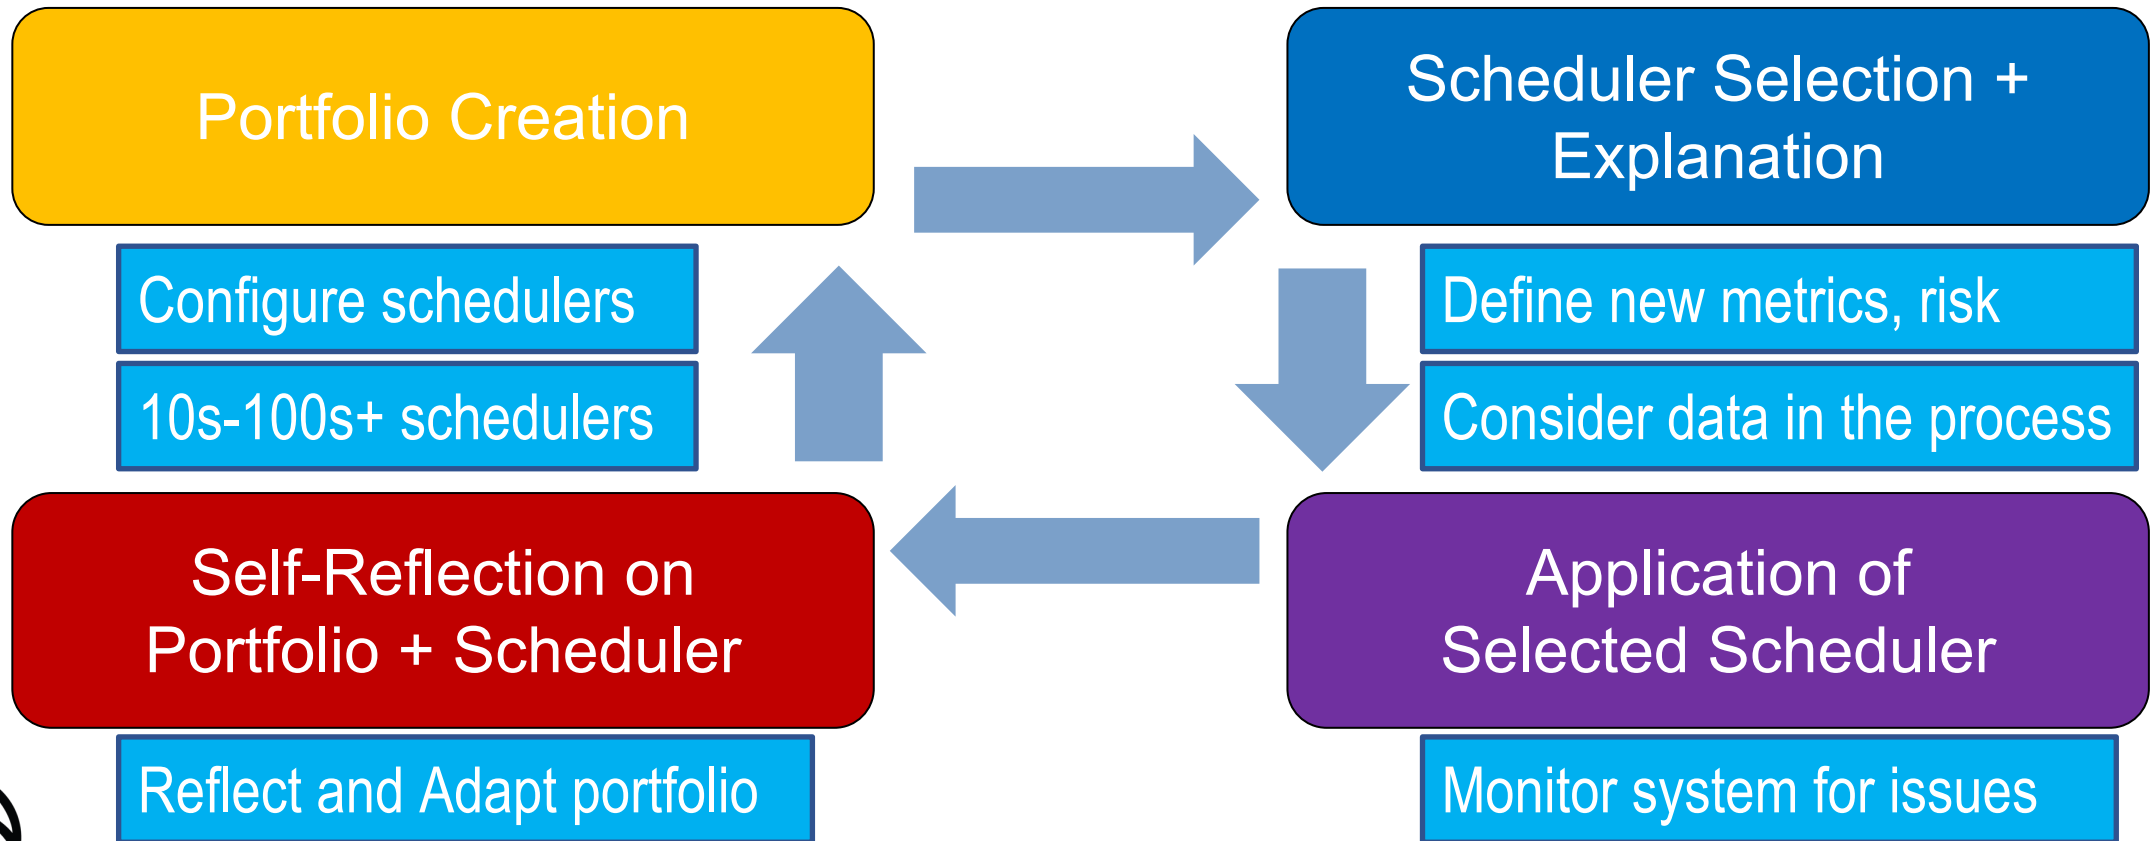


1. Create a set of schedulers (resource provisioning and allocation policies)
  2. Select active scheduler online, apply for the next period, analyze results
- (Repeat)

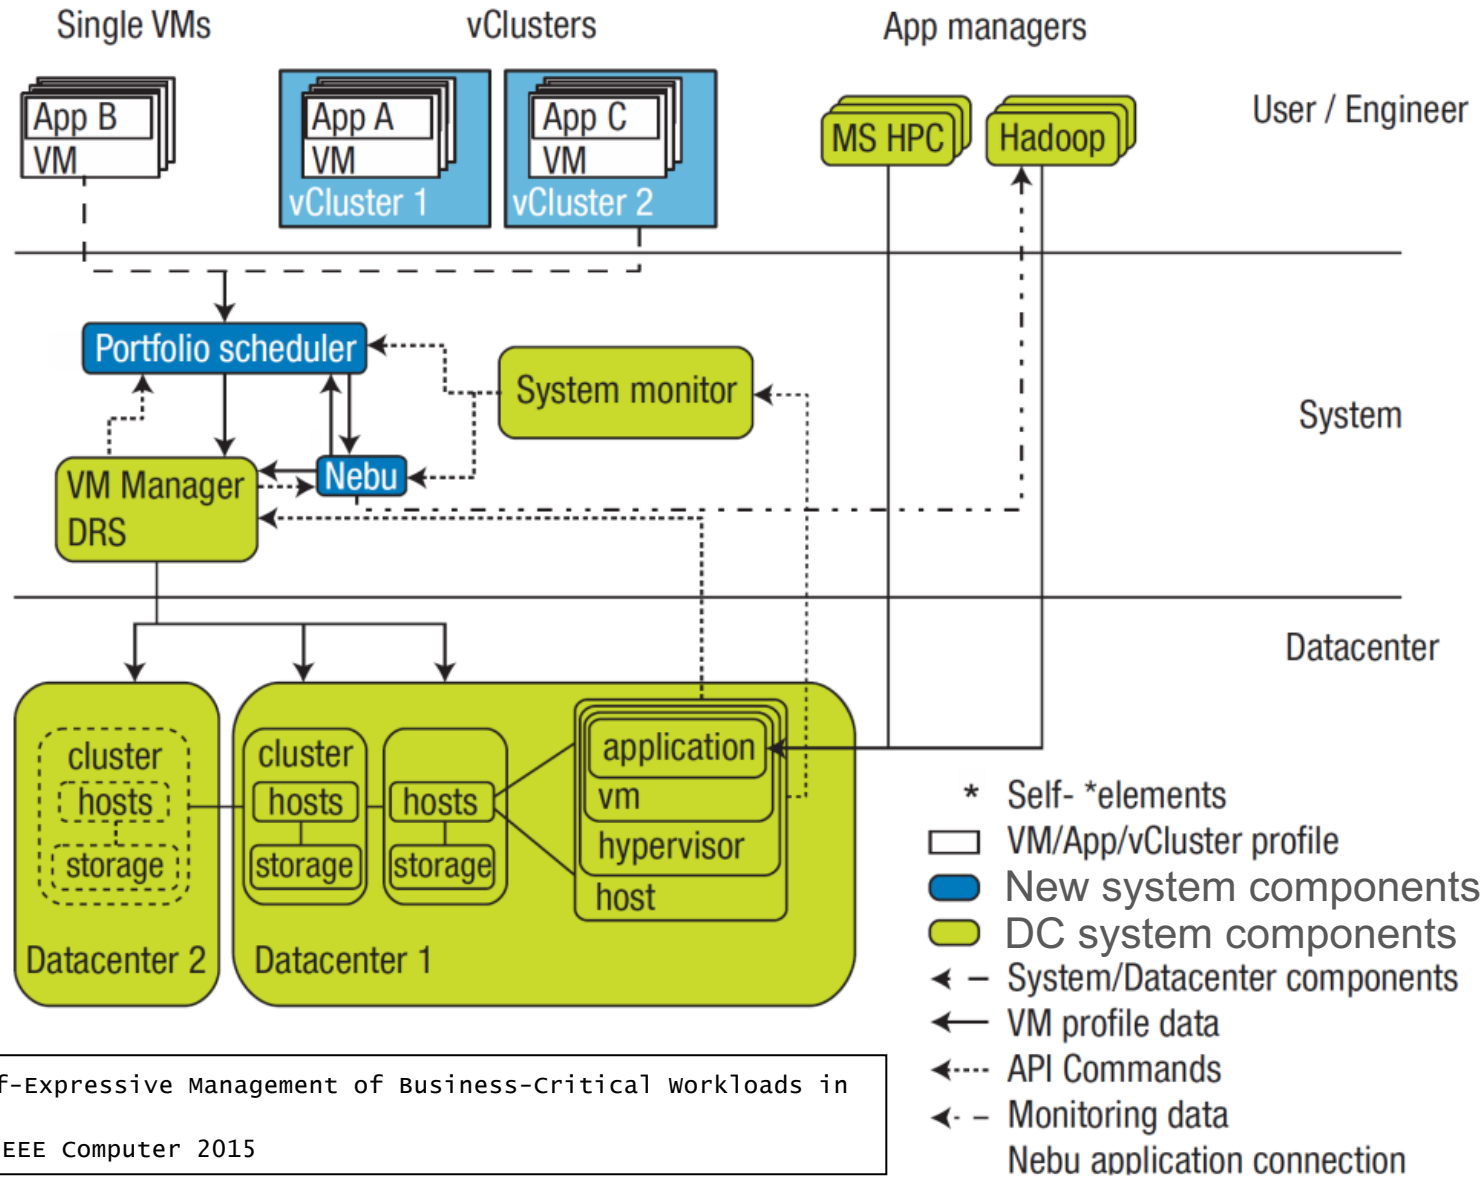


# Generalized Portfolio Scheduling for Computer Systems

## Portfolio Scheduling



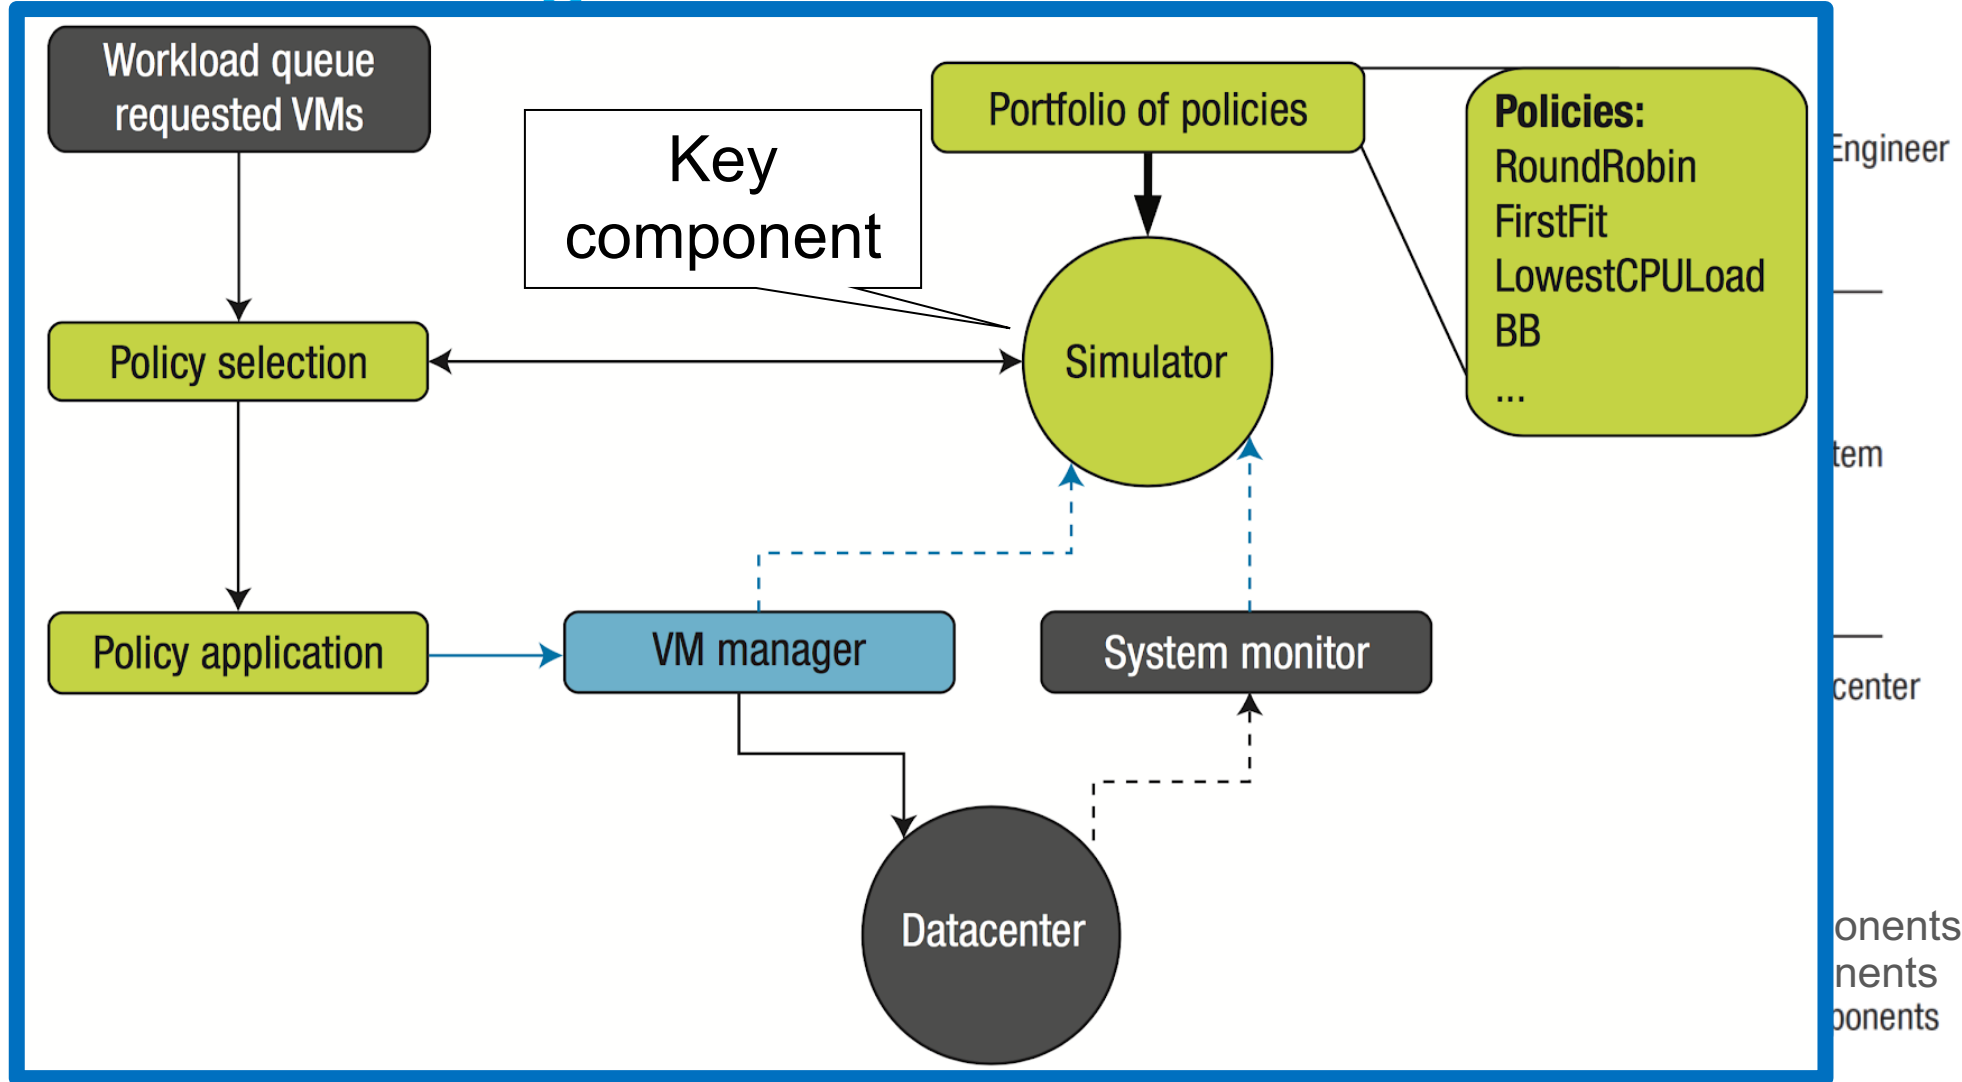
# Portfolio Scheduling in Datacenters



v. van Beek et al. Mnemos: Self-Expressive Management of Business-Critical workloads in Virtualized Datacenters. IEEE Computer 2015



# Portfolio Scheduling in Datacenters



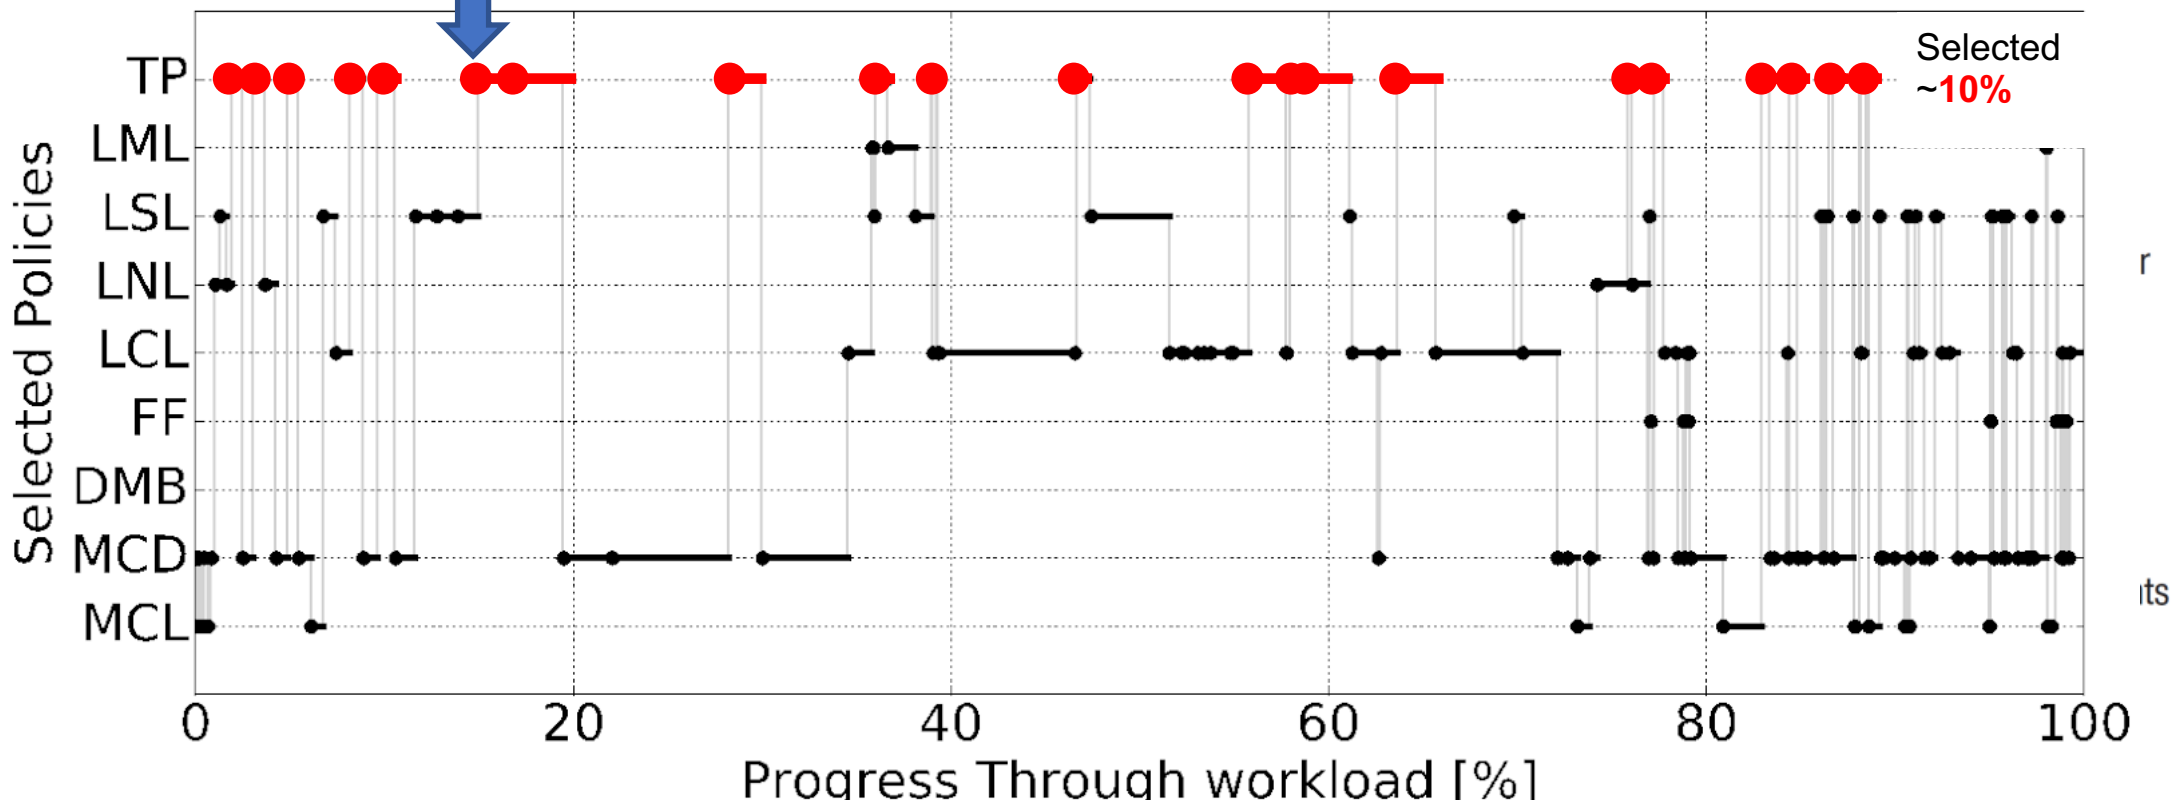
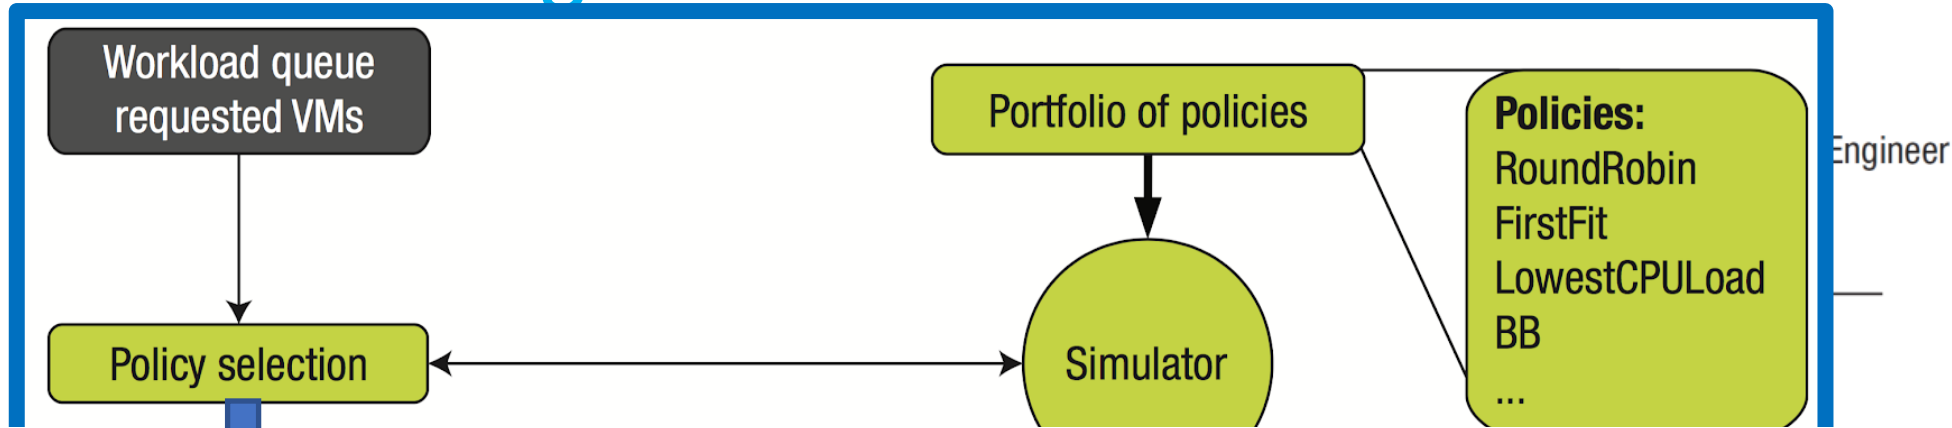
v. van Beek et al. Mnemos: Self-Expressive Management of Business-Critical workloads in Virtualized Datacenters. IEEE Computer 2015

←--- API Commands  
 ←- Monitoring data  
 ←- Nebu application connection






# Portfolio Scheduling in Datacenters

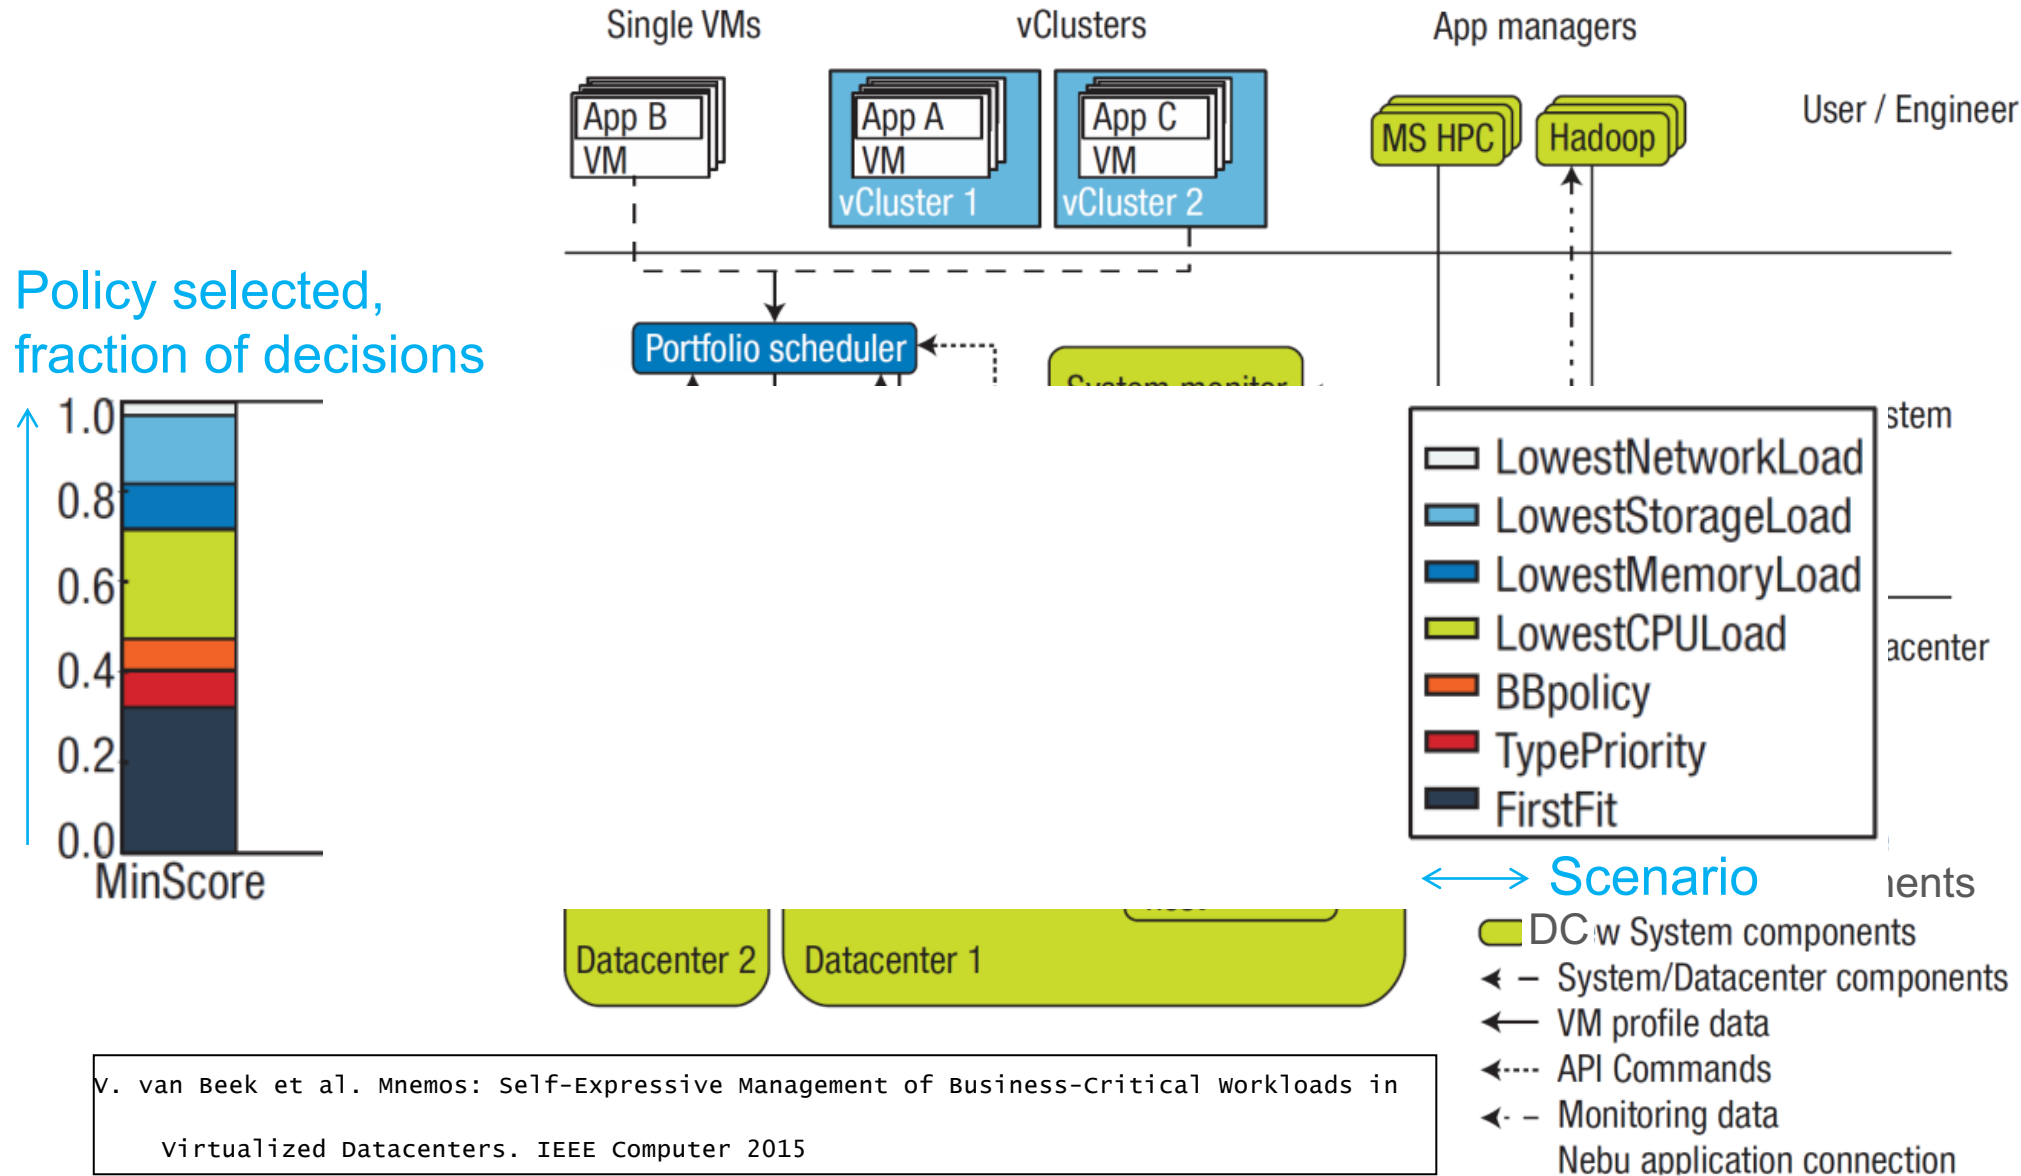


# Portfolio Scheduling in Datacenters



1. Portfolio scheduler achieves the lowest risk of all scenarios.
2. Portfolio scheduler achieves at least **35%** lower DRR compared to individual policies.
3. **40%** lower DRR than commercial production system **REPLAY**.

# Portfolio Scheduling in Practice: Risk Management



v. van Beek et al. Mnemos: Self-Expressive Management of Business-Critical workloads in


Virtualized Datacenters. IEEE Computer 2015





Message:


5<sub>B</sub>



Analytics ~ Only

beginning to scratch  
the surface


# GAMING ANALYTICS: SOCIAL EVERYTHING!



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

## Goal: Improve gaming experience

1. **Ranking / Rating**
2. **Matchmaking / Recommendations**
3. **Detect and combat toxicity**
4. **Exporting data**
5. **Play Style/Tutoring**



Done

Open challenge



# Blizzard: World of Warcraft's Quest for Big Data

Source: Amanda Gerdes, talk at Terradata Partners Conference, NL, Dec 2016.

- New Stakeholders
  - Game Development Team
  - Engagement Manager + Team
  - Data Scientists
  - External Partners + crowdsourcing
- New Requirements
  - Balanced decentralized data warehouse
  - 90% of decision-making done by algorithms
  - From transactional to data-driven and model-driven behavioral data
- New Applications
  - Matchmaking
  - Event prediction
  - Recommendations
  - Fraud detection and classification
  - Micro-segmentation for campaigns
  - Third-party / crowdsourced analysis


# TOXICITY DETECTION IN ONLINE GAMES

- Data collection and cleansing

- Representative MOBA game
- DotAlicious ~13k games, Feb 2012
- Identified ~10k games w/ victory/loss

- Analysis

- Used chat logs, both **ally-chat** and **all-chat**
- Natural Language Processing  $\square$  limited topics
- Analyzed vocabulary using toxicity-detection rules



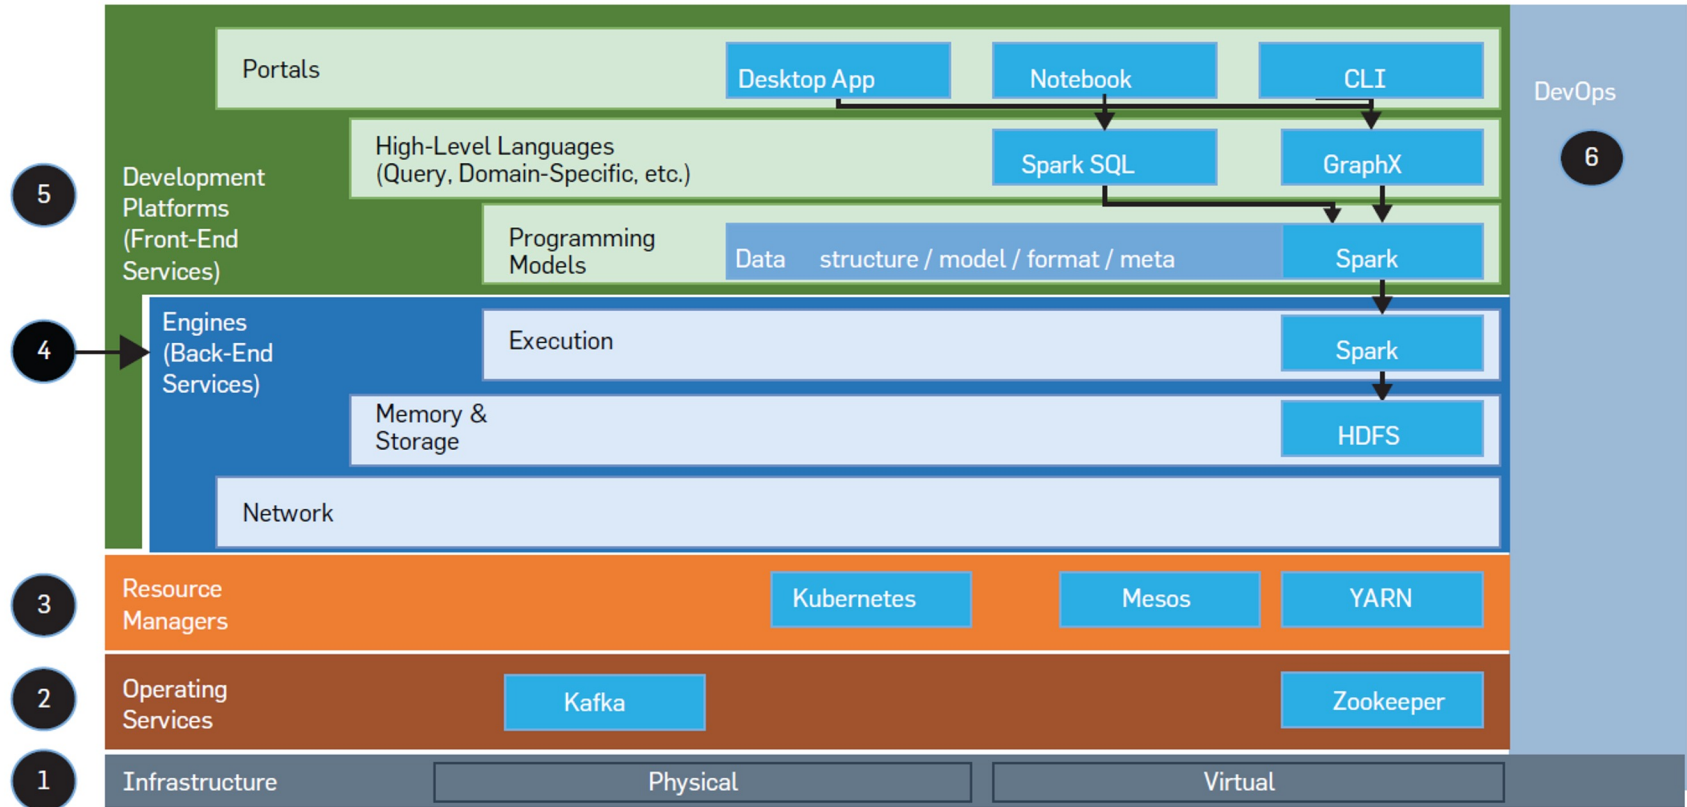
# GRAPH PROCESSING FOR GAME SERVICES

CACM Cover/Featured article,  
Sep 2021

for the next decade and beyond.

## The Future Is Big Graphs: A Community View on Graph Processing Systems

GRAPHS ARE, BY nature, 'unifying abstractions' that can leverage interconnectedness to represent, explore, predict, and explain real- and digital-world phenomena. Although real users and consumers of graph instances and graph workloads understand these abstractions, future problems will require new abstractions and systems. What needs to happen in the next decade for big graph processing to continue to succeed?



# INTERACTION GRAPHS: FROM GAME INSTANCES TO SOCIAL TIES

- How to map the relationships in matches to graphs?
  - 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. Link if percent matches played together.


# 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		Team 2	
Player	Cluster	Player	Cluster
A	1	F	2
B	2	G	5
C	1	H	3
D	3	I	6
E	4	J	3

Bonding score: 0 3 5 7

Matchmaking based on interaction graphs much better than random







Message:

5c



Content ~

Computational

power → emerging

creativity, freshness

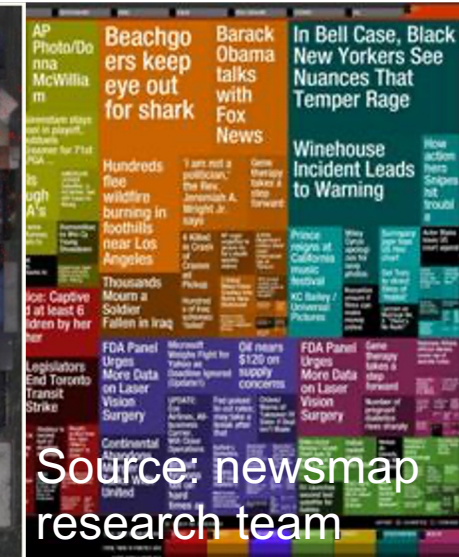
# CONTENT, CONTENT, CONTENT!

Goal: Produce and distribute content for 1BN players


1. Game Analytics → Game statistics
2. Content distribution/ Streaming content
3. Crowdsourcing
4. Storification
5. Auto- / user-generated game content
6. Adaptive game content at scale
7. Make procedural content as appealing as designer content



Source: America's Army  
3rd party heat map



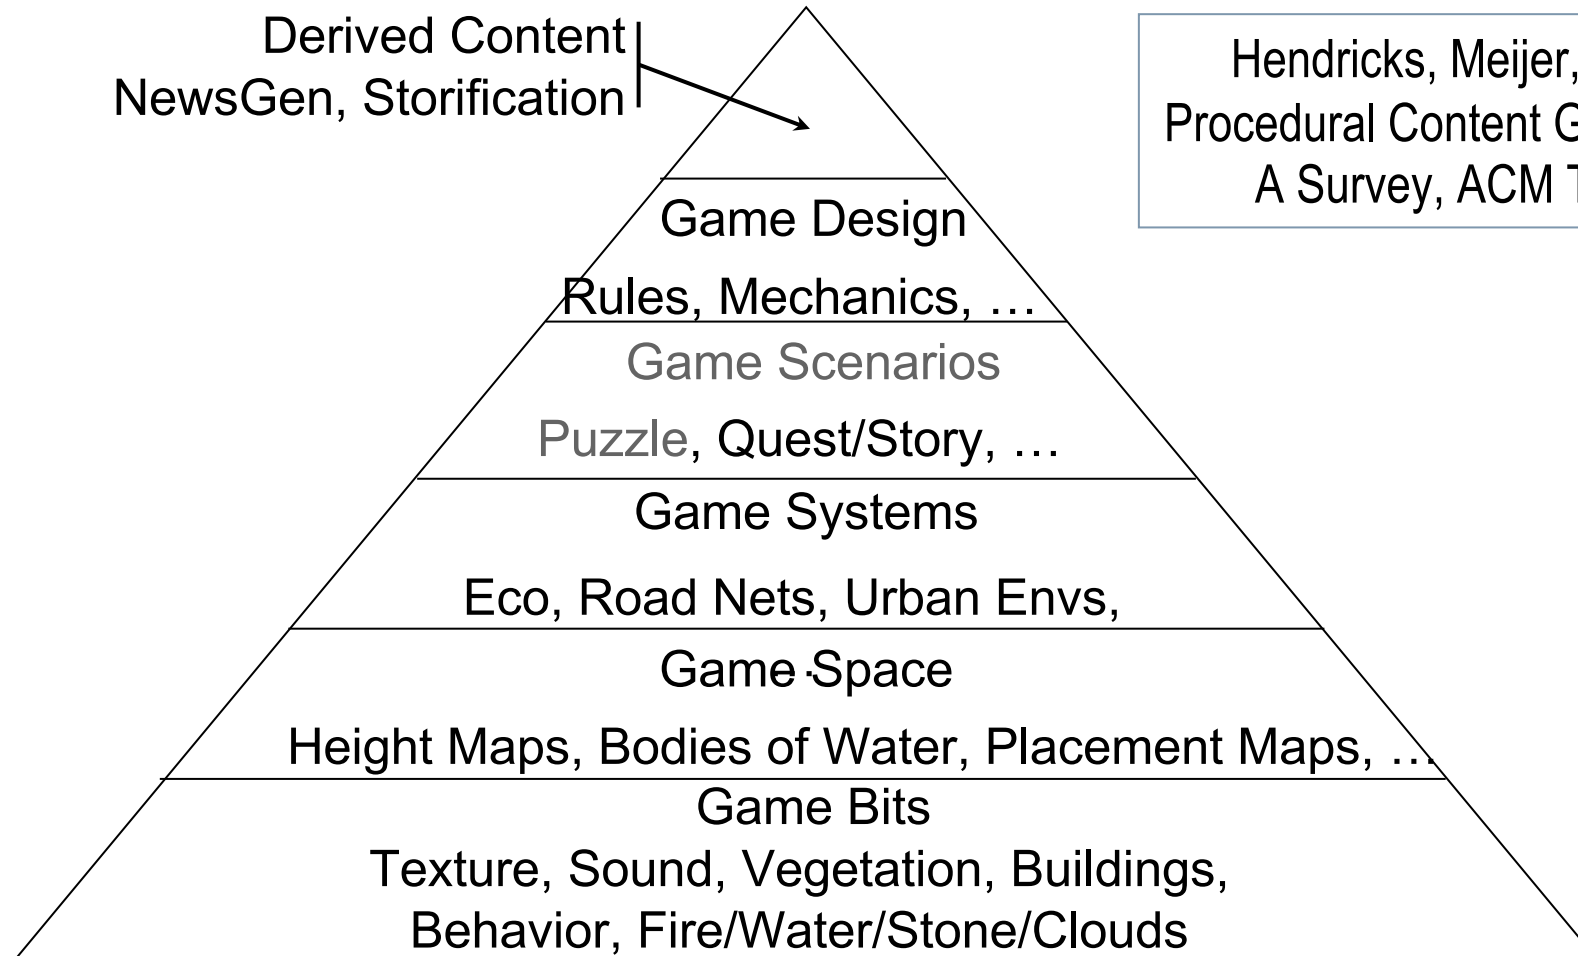
Source: newspaper  
research team



Done

Open challenge

# (Procedural) Game Content (Generation)



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

# STRATEGIES FOR CONTENT GENERATION


- **Clone**

- Budget: \$
- Zero recognition from game designers
- Amazing career in sales, does not work on Steam, players are unhappy



- **Franchise and 18-/24-month DLC**

- Budget: \$\$\$, upfront + \$-\$\$ / update
- Little recognition from game designers
- Works fine for sales, works on Steam, players unhappy



- **Frequent updates, fast-pace DLC**

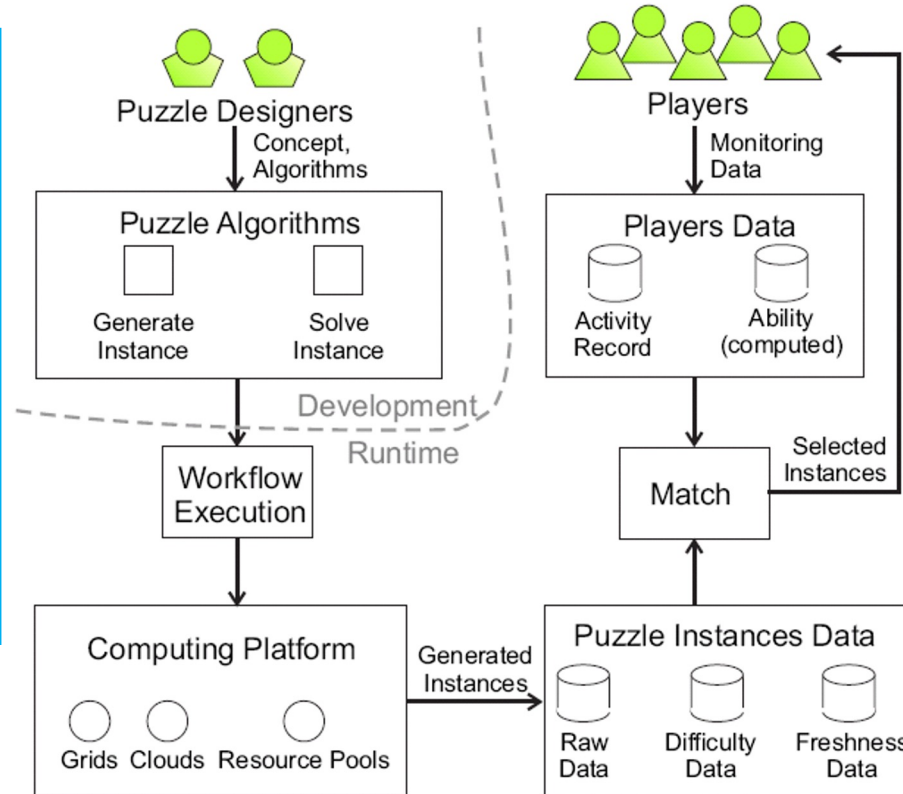
- Budget: \$, upfront + \$ / update
- Breaks industry models: often based on technology, etc.
- Players are happy



# The POGGI Content Generation Framework

## Scalability

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



Smart system to  
recommend instances  
to players

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




# Puzzle-Specific Considerations

## Generating Player-Customized Content

### Puzzle difficulty








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

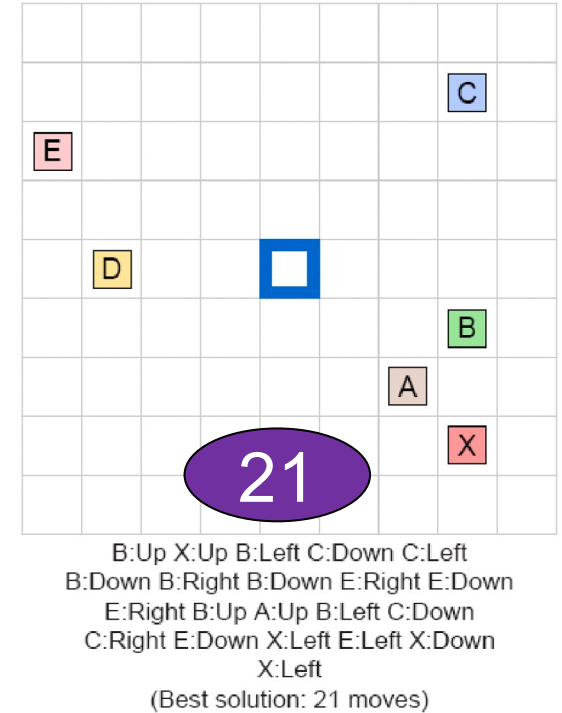


### Player ability

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

## Scalability (of **content**)


Target:  Pins:      





Message:

5<sub>D</sub>



Meta-gaming ~

Community

techniques

emerging

# META-GAMING: BEYOND THE DESIGNED INTERACTION

**Meta-gaming**="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)

## Goal: Self-Organizing Gaming Communities of Millions of Players

1. Recording player behavior
2. Understanding player behavior
3. Ranking / Rating
4. Play style analysis Done
5. Exporting data Open challenge



XLink Kai: Evo VII  
global network gaming  
teamxlink.co.uk

# TWITCH.TV (JUSTIN.TV)

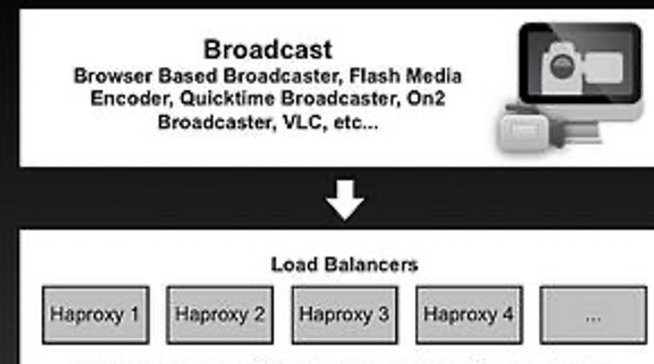


- Functionality
  - Video streaming: concurrent streams avg. 30k / max. 2m
  - Chat: 10b messages/day
  - API-based access to functionality + data: avg. 50k reqs./s
  - Data science infrastructure
  - DevOps: distributed everything
- Main lesson:
  - Distributed ecosystem

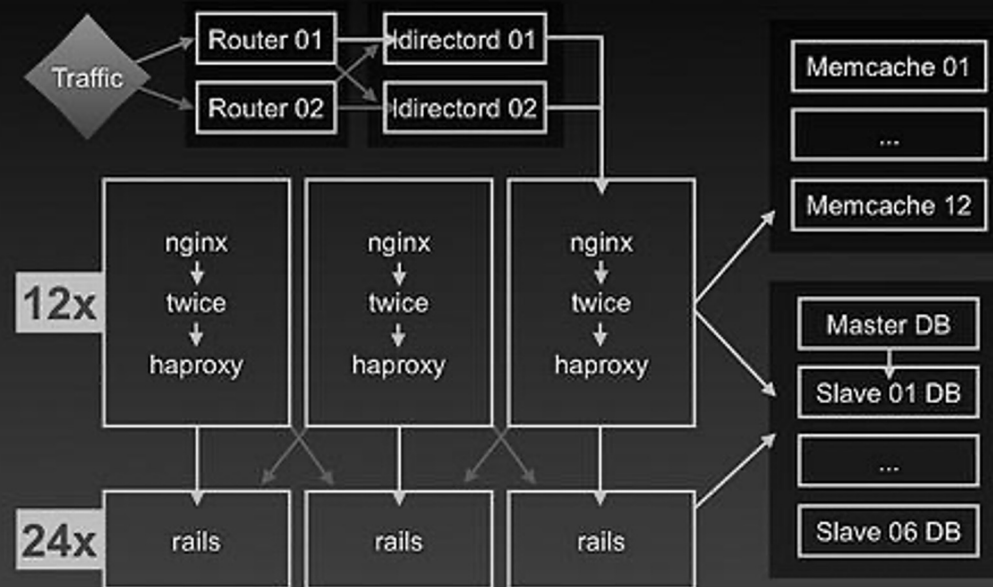
<https://blog.twitch.tv/twitch-engineering-an-introduction-and-overview-a23917b71a25>  
<http://highscalability.com/blog/2010/3/16/justintv-vs-live-video-broadcasting-architecture.html>



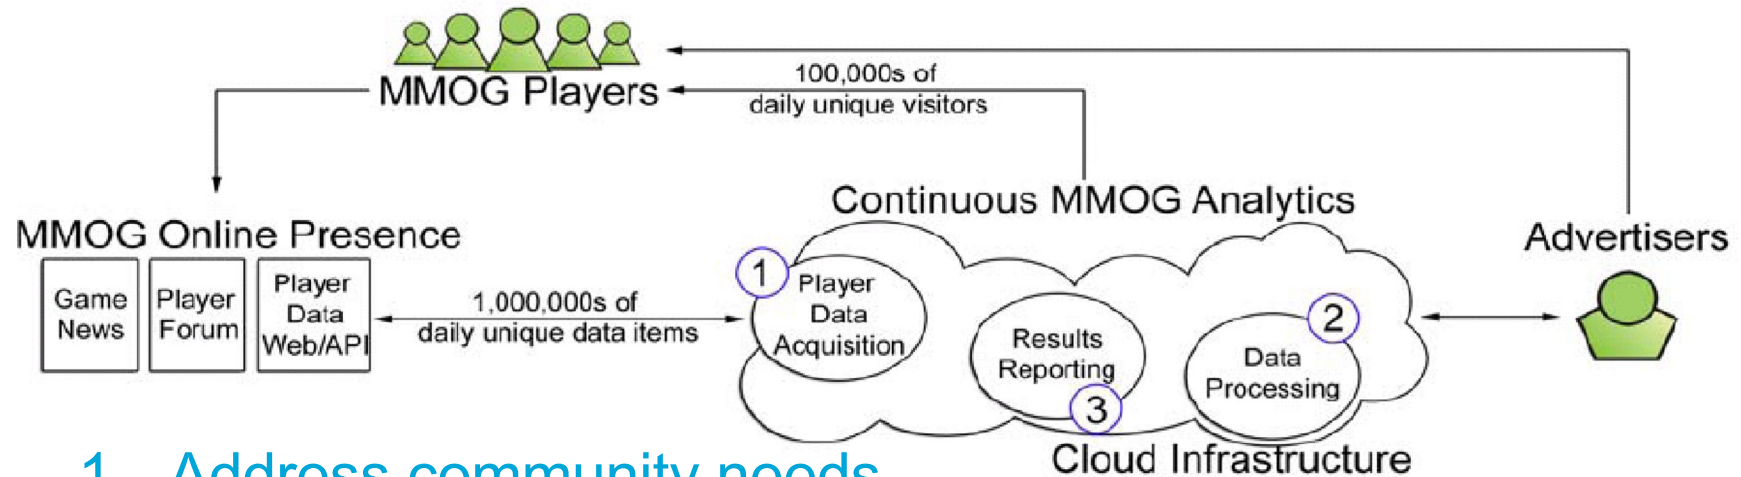
## Live Video Architecture: Broadcasting



## Web Architecture: Backend



# 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 Computing


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

Observed XFire: 2008—2011 (3+ years), 500K of 20M players



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



# Massivizing Online Games = Rich challenge of computer systems, with societal impact!



1. This is the Golden Age of massive computer ecosystems
2. But we cannot take this technology for granted
3. A new science to help massivizing computer ecosystems
4. Understand how things work: observe, synthesize, experiment
5. Design and build better systems:  
virtual worlds, analytics, content generation, meta-gaming

6. What's next? PLATO in the 21<sup>st</sup> century

7. Beyond games: Massivizing computer systems is a national priority in the NL!

Lots to do  
next!

Message: 6



What's next?



--OR--

PLATO in the

21<sup>st</sup> century

# PLATO: Early Games, Early Social Apps (1970s)

- PLATO
  - Early learning management system
  - Distributed, thousands of deployments
  - Basics of messaging, communication
  - Basics of many online game genres
  - Basics of social networking

**“Control Data PLATO® increased**

**“PLATO computer-based education will save Merrell-National over 12,000 classroom hours on just one drug product training program alone.”**

Gary J. Wilson, Sales Training Manager  
Merrell-National Laboratories  
Division of Richardson-Merrell Inc.



“Today’s modern ethical drugs demand sales people who are better informed, more technically oriented and better able to relate vital product information to help doctors do a better job. To meet the challenge, Merrell-National Laboratories has initiated a program to improve the quality of pharmaceutical representative training.

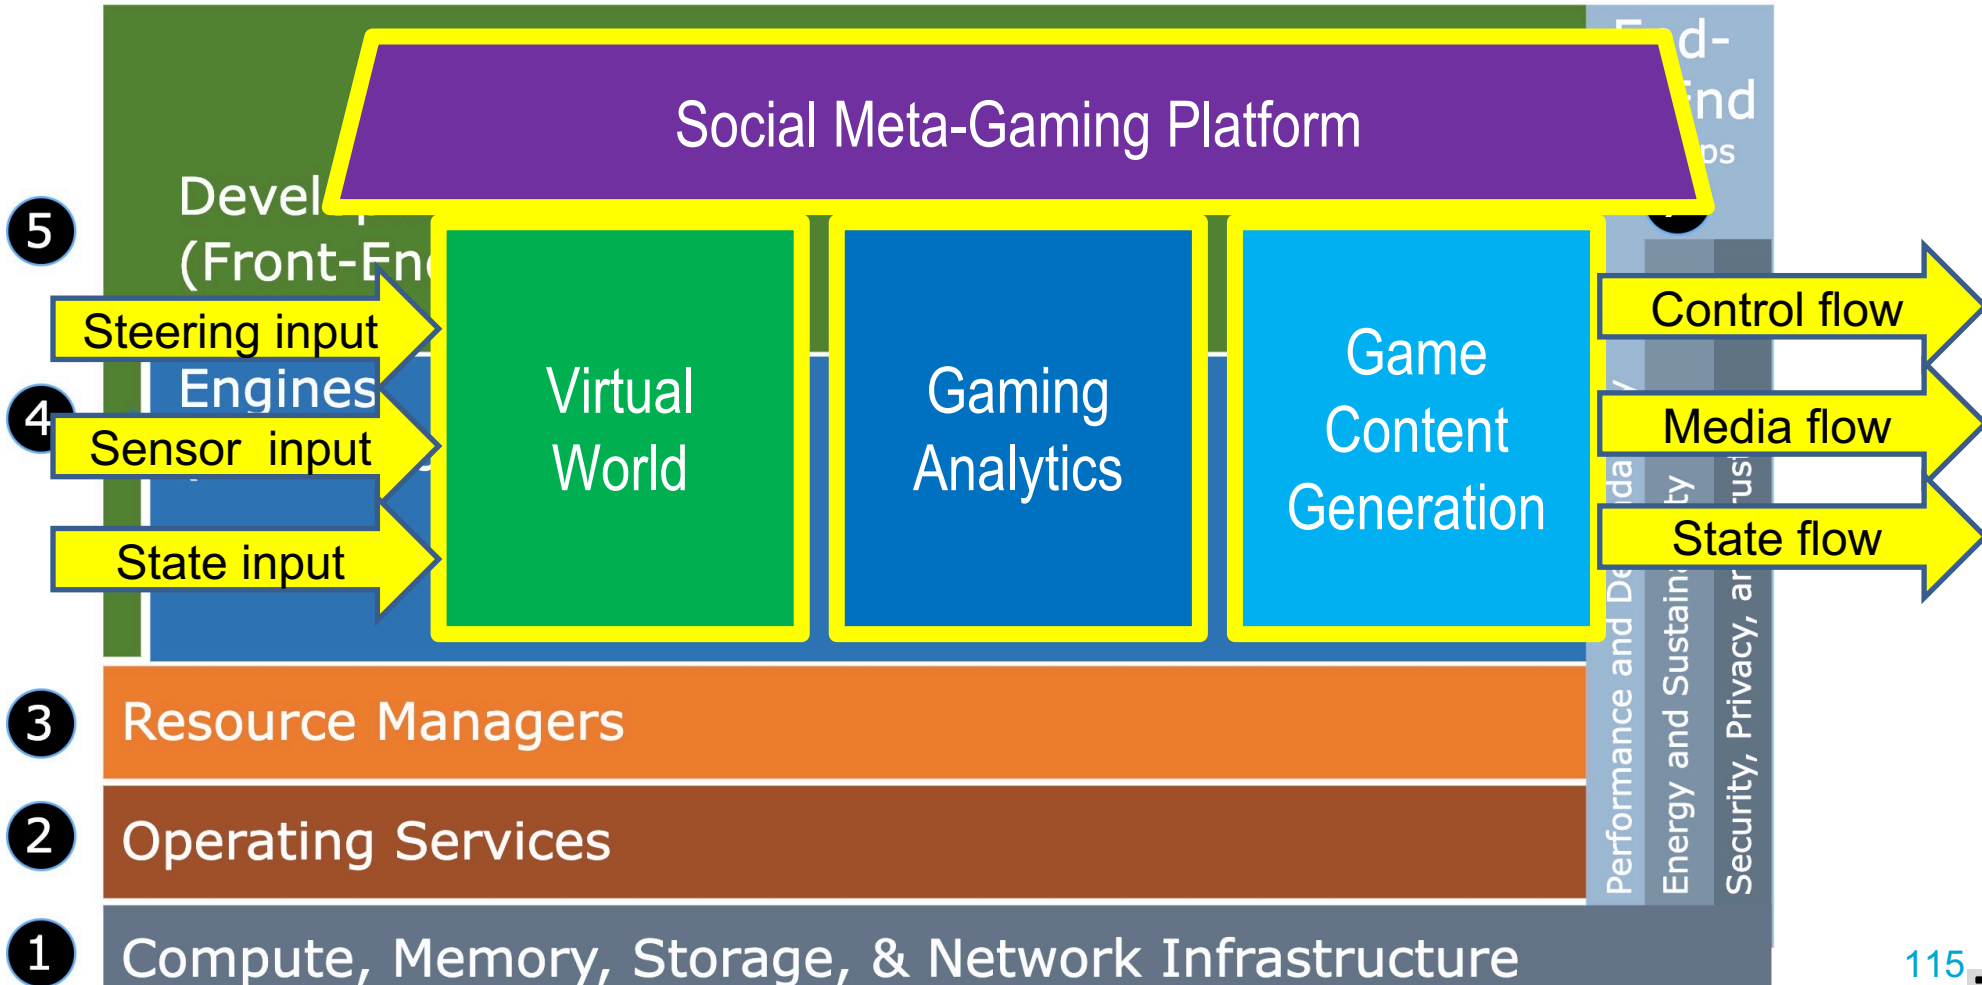
“At the heart of this program is the Control Data PLATO system, a truly versatile and cost-effective approach to computer-based



*For further information, write on your business letterhead to Control Data Education Company, HQN111, P.O. Box 0, Minneapolis, MN 55440. Or call 612/853-7600.*

# HOW TO MANAGE SYSTEM COMPLEXITY?

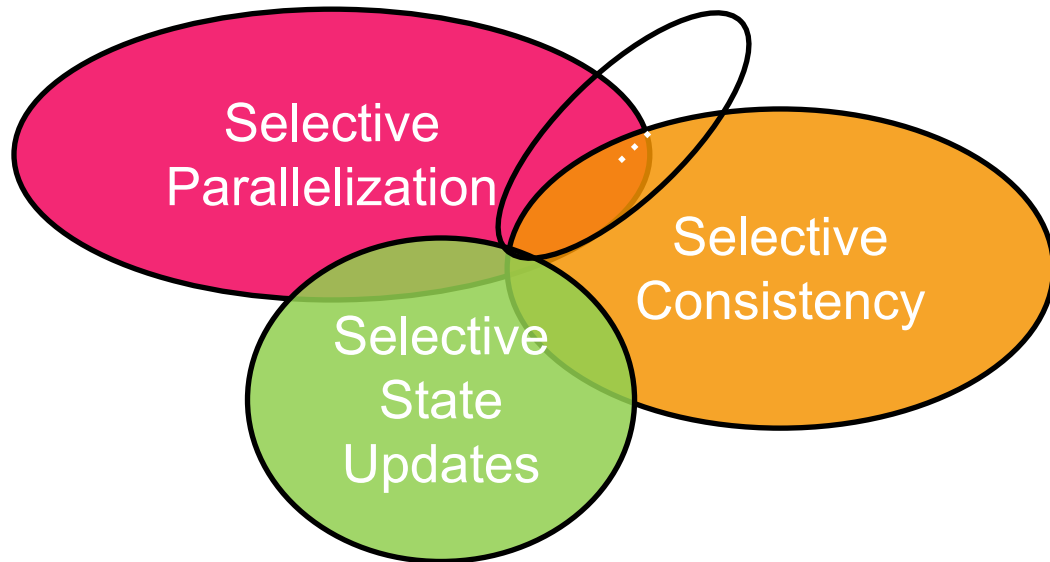
6 Data & App Market



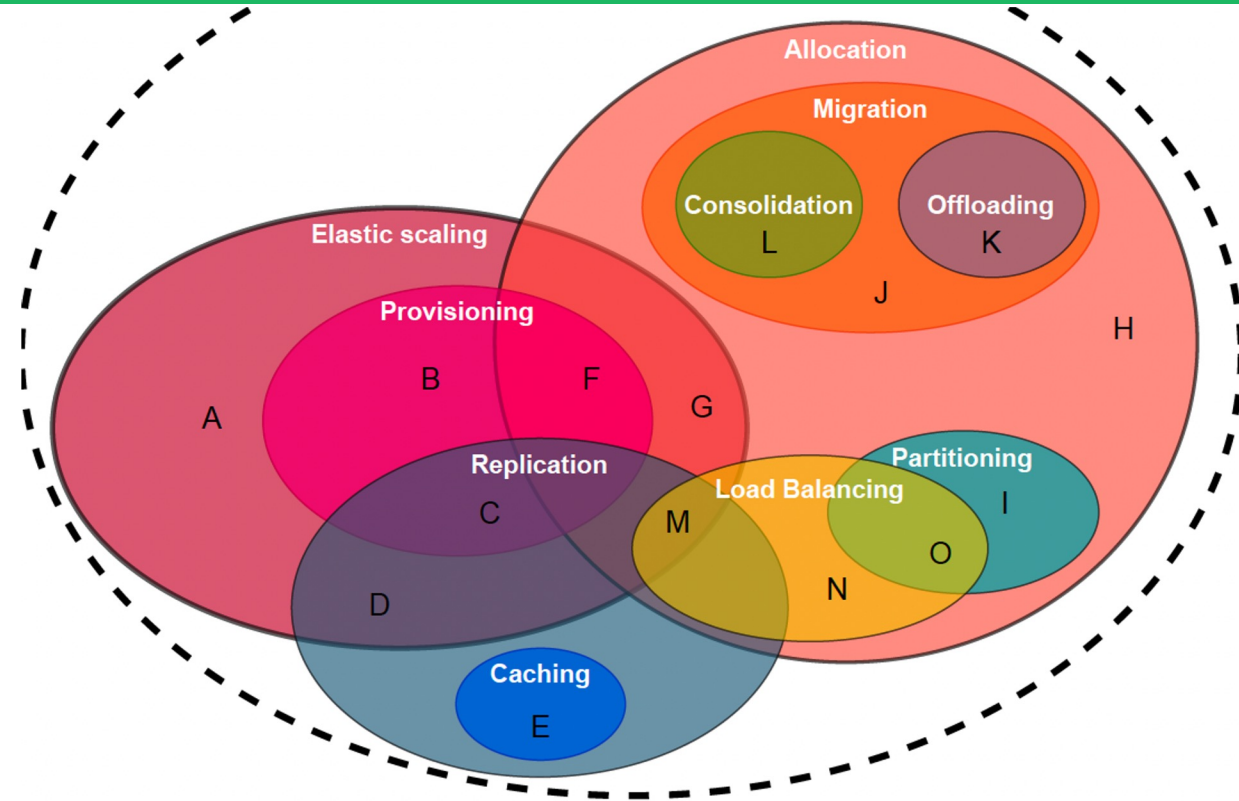


# HOW TO EXPLORE SYSTEMATICALLY...

## EXTREME SCALING TECHNIQUES



## OTHER OPERATIONAL TECHNIQUES



# HOW TO MANAGE SERVERLESS OPERATIONS?

EXTREME AUTOMATION

REFERENCE ARCHITECTURE OF FAAS PLATFORMS

## Workflow Composition Layer



## Function Management Layer



## Resource Orchestration Layer



Business Concerns

Operational Concerns

# HOW TO ACHIEVE FINE-GRAINED BILLING AND UTILIZATION-BASED BILLING?



Erwin van Eyk

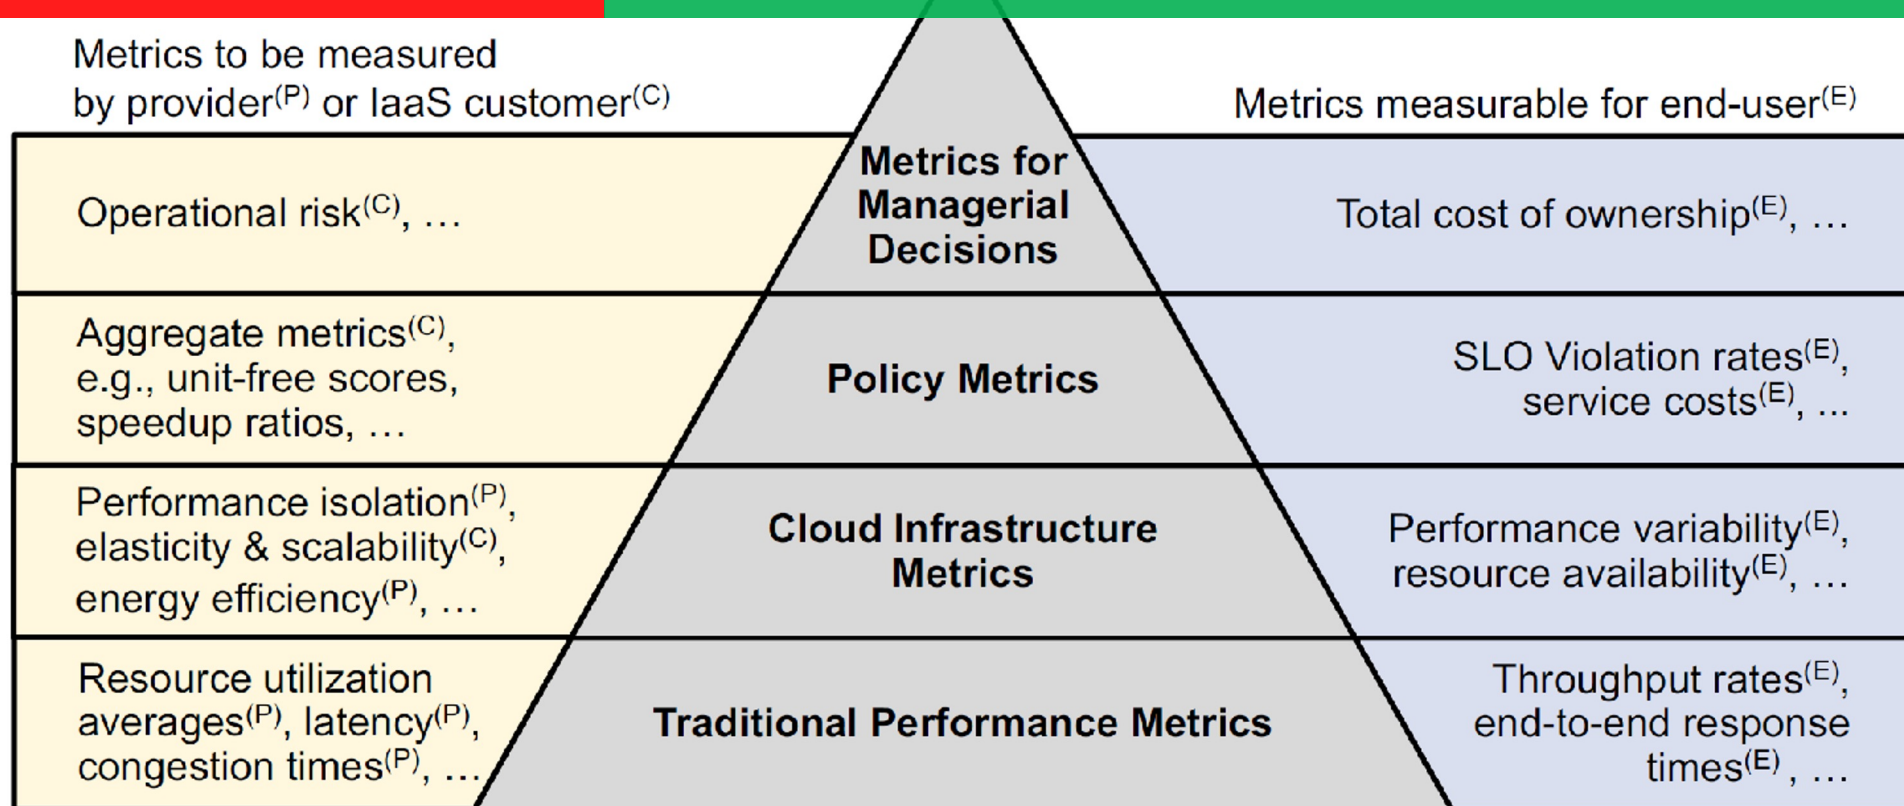
[van Eyk et al. (2019) Serverless is More: From PaaS to Present Cloud Computing, IEEE Internet Computing] [\[Online\]](#)

# WHAT TO MEASURE? HOW TO INFER/PREDICT?



## MANAGEABILITY CHALLENGE

## REFERENCE VIEW ON METRICS



N. Herbst, E. Van Eyk, C. L. Abad, A. Iosup, et al. (2018) Quantifying Cloud Performance and Dependability: Taxonomy, Metric Design, and Emerging Challenges. TOMPECS 3(4): 19:1-19:36

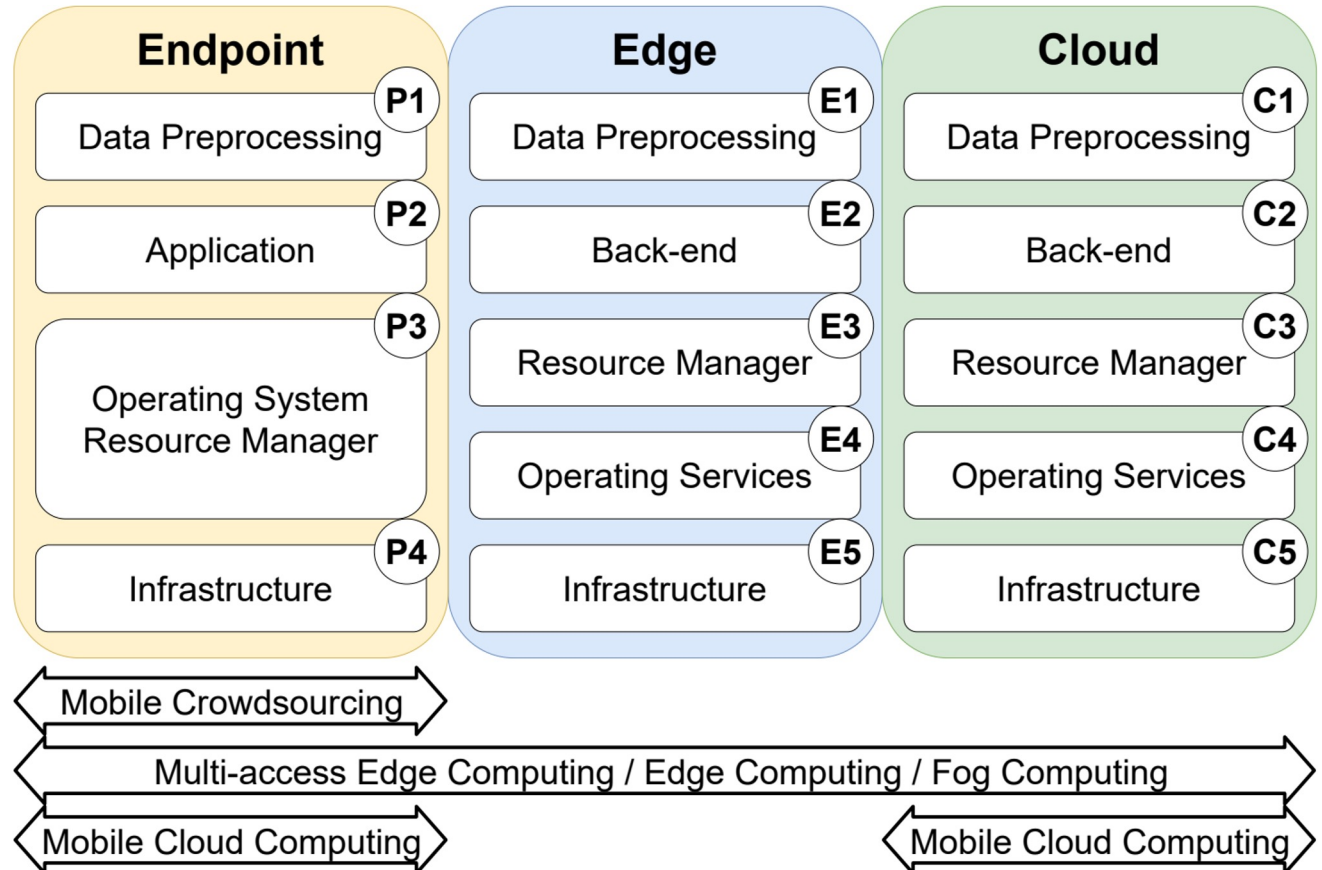
# HOW TO USE THE COMPUTING CONTINUUM




## THE COMPLEXITY CHALLENGE

## REFERENCE ARCHITECTURE ACROSS PLATFORMS

HOW TO  
CONTAINERIZE  
EVERYTHING?



[Jansen, Trivedi, Iosup, et al. (2022) Ongoing Work



# Massivizing Online Games = Rich challenge of computer systems, with societal impact!



1. This is the Golden Age of massive computer ecosystems
2. But we cannot take this technology for granted
3. A new science to help massivizing computer ecosystems
4. Understand how things work: observe, synthesize, experiment
5. Design and build better systems:  
virtual worlds, analytics, content generation, meta-gaming
6. What's next? PLATO in the 21<sup>st</sup> century

7. Beyond games: Massivizing computer systems is a national priority in the NL!

Lots to do  
next!





Message:

7

Beyond games:

Massivizing

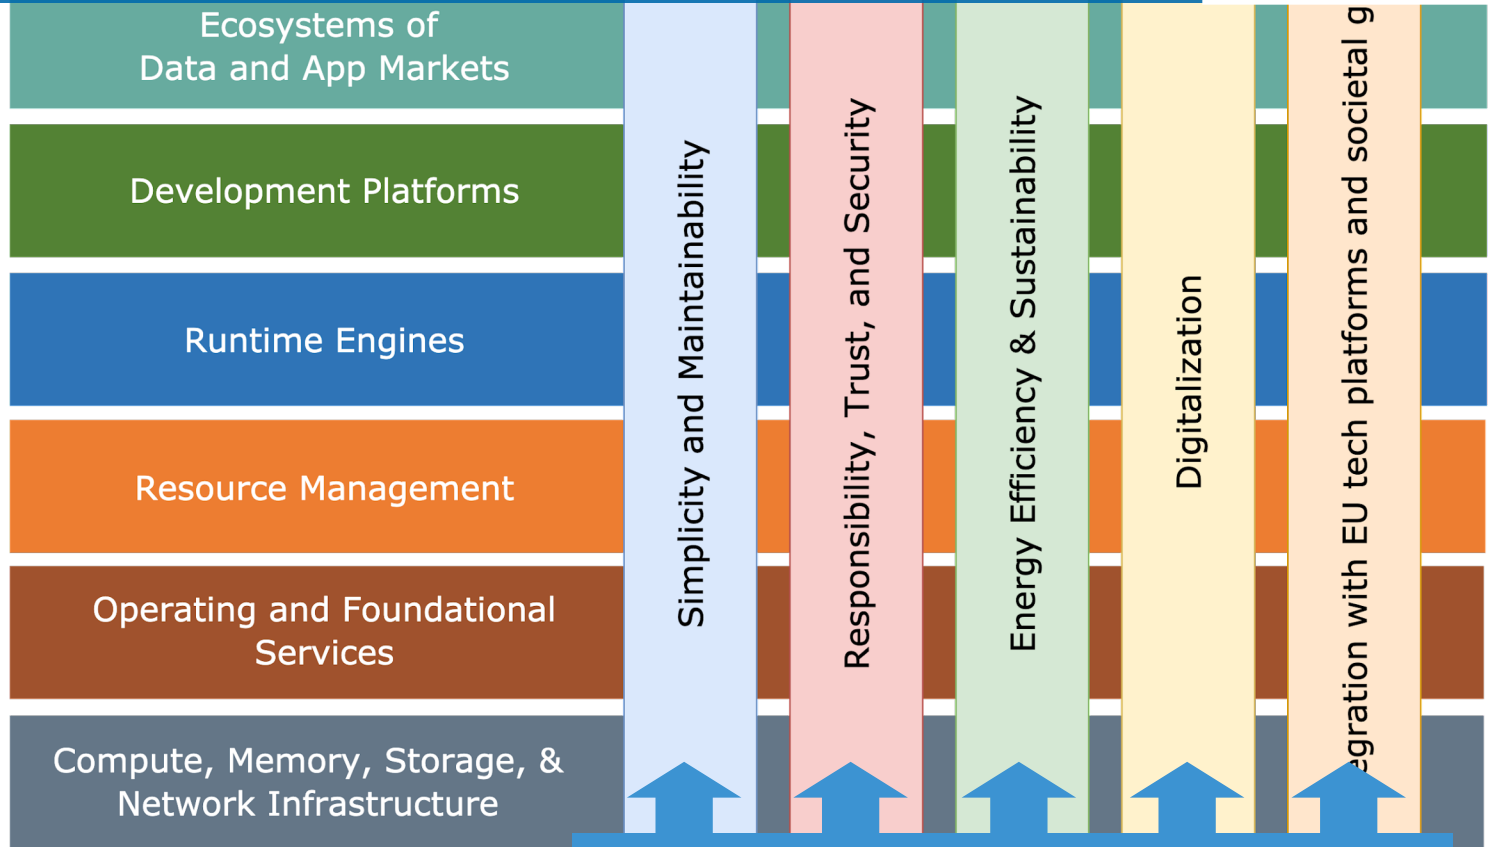
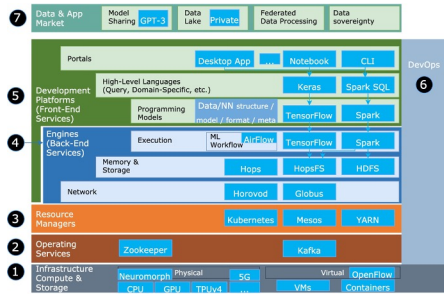
computer systems =

NL priority!


<https://bit.ly/ManifestoCompSysNL> ▶  
<https://arxiv.org/abs/2206.03259> ▶



# A LARGER VISION OF HOW COMPUTING WILL HELP OUR SOCIETY



A.iosup@vu.nl  
<http://atlarge.science>



# Massivizing Online Games = Rich challenge of computer systems, with societal impact!



1. This is the Golden Age of massive computer ecosystems
2. But we cannot take this technology for granted
3. A new science to help massivizing computer ecosystems
4. Understand how things work: model, simulate, experiment
5. Design and build: virtual worlds, social networks, meta-gaming
6. Workforce for the 21<sup>st</sup> century
7. Beyond games: Massivizing computer systems is a national priority in the NL!

Take-Home Message >

Lots to do  
next!


# TAKE-HOME MESSAGE:




Massivizing Online Games =  
Rich challenge of computer  
systems, with societal impact!

Understand how things work:  
observe, synthesize, experiment

Design and build better systems:  
virtual world, analytics, content  
generation, and meta-gaming  
techniques, individual and all put  
together into gaming ecosystem



Lots to do  
next!



# Massivizing Online Games = Rich challenge of computer systems, with societal impact!



1. This is the Golden Age of massive computer ecosystems
2. But we cannot take this technology for granted
3. A new science to help massivizing computer ecosystems
4. Understand how things work: observe, synthesize, experiment
5. Design and build better systems:  
virtual worlds, analytics, content generation, meta-gaming
6. What's next? PLATO in the 21<sup>st</sup> century
7. Beyond games: Massivizing computer systems is a national priority in the NL!

Lots to do  
next!





WANT TO READ  
MORE ON THE  
TOPIC?

# MASSIVIZING COMPUTER SYSTEMS



## FURTHER READING

<https://atlarge-research.com/publications.html>

1. Crusoe et al. (2022) Methods Included: Standardizing Computational Reuse and Portability with the Common Workflow Language. CACM. (accepted, in print)
2. Andreadis et al. (2022) Capelin: Data-Driven Capacity Procurement for Cloud Datacenters using Portfolios of Scenarios. TPDS.
3. Sakr, Bonifati, Voigt, Iosup, et al. (2021) The Future Is Big Graphs! CACM.
4. Eismann et al. (2021) A Review of Serverless Use Cases and their Characteristics. TSE.
5. Eismann et al. (2021) Serverless Applications: Why, When, and How? IEEE Softw. 38(1): 32-39 (2021)
6. Mastenbroek et al. (2021) OpenDC 2.0: Convenient Modeling and Simulation of Emerging Technologies in Cloud Datacenters. CCGRID.
7. Versluis and Iosup (2021) A survey of domains in workflow scheduling in computing infrastructures: Community and keyword analysis, emerging trends, and taxonomies. FGCS.
8. Uta et al. (2020) Is Big Data Performance Reproducible In Modern Cloud Networks? NSDI.
9. Donkevliet et al. (2021) Dyconits: Scaling Minecraft-like Services through Dynamically Managed Inconsistency. ICDCS.
10. Versluis et al. (2020) The Workflow Trace Archive. TPDS.
11. Hegeman et al. (2021) GradeML. HotCloudPerf.
12. Uta et al. (2020) Beneath the SURFace: An MRI-like View into the Life of a 21st-Century Datacenter. login USENIX
13. Iosup, Hegeman, et al. (2020) The LDBC Graphalytics Benchmark. CoRR.
14. Abad, Iosup, et al. An Analysis of Distributed Systems Syllabi With a Focus on Performance-Related Topics.

# MASSIVIZING COMPUTER SYSTEMS



## FURTHER READING

<https://atlarge-research.com/publications.html>

1. Iosup et al. The AtLarge Vision on the Design of Distributed Systems and Ecosystems. ICDCS 2019 ← Start here
2. Uta et al. Is big data performance reproducible in modern cloud networks? NSDI 2020
3. Van Eyk et al. The SPEC-RG Reference Architecture for FaaS: From Microservices and Containers to Serverless Platforms, IEEE IC 2019
4. Papdopoulos et al. Methodological Principles for Reproducible Performance Evaluation in Cloud Computing. TSE 2019 and (journal-first) ICSE 2020
5. van Beek et al. Portfolio Scheduling for Managing Operational and Disaster-Recovery Risks in Virtualized Datacenters Hosting Business-Critical Workloads. ISPD 2019
6. van Beek et al. A CPU Contention Predictor for Business-Critical Workloads in Cloud Datacenters. HotCloudPerf19

+ Iyushkin et al. Performance-Feedback Autoscaling with Budget Constraints for Cloud-based Workloads of Workflows. Under submission

Etc.

# MASSIVIZING COMPUTER SYSTEMS



## FURTHER READING

<https://atlarge-research.com/publications.html>

1. Iosup et al. Massivizing Computer Systems. ICDCS 2018 ← start here
  2. Andreadis et al. A Reference Architecture for Datacenter Scheduling, SC18
  3. Van Eyk et al. Serverless is More: From PaaS to Present Cloud Computing, IEEE IC Sep/Oct 2018
  4. Uta et al. Exploring HPC and Big Data Convergence: A Graph Processing Study on Intel Knights Landing, IEEE Cluster 2018
  5. Talluri et al. Big Data Storage Workload in the Cloud. ACM/SPEC ICPE 2019.
  6. Toader et al. Graphless. IEEE ISPDC'19.
  7. Jiang et al. Mirror. CCPE 2018.
  8. Ilyushkin et al. Autoscalers. TOMPECS 2018.
  9. Versluis et al. Autoscaling Workflows. CCGRID'18.
  10. Uta et al. Elasticity in Graph Analytics? IEEE Cluster 2018.
  11. Herbst et al. Ready for rain? TOMPECS 2018.
  12. Guo et al. Streaming Graph-partitioning. JPDC'18.
  13. Iosup et al. The OpenDC Vision. ISPDC'17.
  14. Iosup et al. Self-Aware Computing Systems book.
  15. Iosup et al. LDBC Graphalytics. PVLDB 2016.
- Etc.

# MASSIVIZING COMPUTER SYSTEMS



## FURTHER READING ON GAMING

1. Jerom van der Sar, Jesse Donkervliet, Alexandru Iosup. Yardstick. ICPE (2019)
2. M. H. Jiang, Otto W. Visser, I. S. W. B. Prasetya, Alexandru Iosup: A mirroring architecture for sophisticated mobile games using computation-offloading. Concurrency and Computation: Practice and Experience 30(17) (2018)
3. Adele Lu Jia et al. (2016) When Game Becomes Life: The Creators and Spectators of Online Game Replays and Live Streaming. TOMCCAP 12(4): 47:1-24.
4. Marcus Märtens, Siqi Shen, Alexandru Iosup, Fernando A. Kuipers: Toxicity detection in multiplayer online games. NETGAMES 2015: 1-6
5. Adele Lu Jia et al. (2015) Socializing by Gaming: Revealing Social Relationships in Multiplayer Online Games. TKDD 10(2): 11:1-29.
6. Iosup et al. (2014) : Analyzing Implicit Social Networks in Multiplayer Online Games. IEEE Internet Computing 18(3): 36-44 (2014).
7. Alexandru Iosup: POGGI: generating puzzle instances for online games on grid infrastructures. Concurrency and Computation: Practice and Experience 23(2): 158-171 (2011)
8. Siqi Shen, Otto W. Visser, Alexandru Iosup: RTSenv: An experimental environment for real-time strategy games. NETGAMES 2011: 1-6
9. Alexandru Iosup, Adrian Lascateu, Nicolae Tapus: CAMEO: Enabling social networks for Massively Multiplayer Online Games through Continuous Analytics and cloud computing. NETGAMES 2010: 1-6
10. Shen, Deng, Iosup, and Epema: Scheduling Jobs in the Cloud Using On-Demand and Reserved Instances. Euro-Par 2013: 242-254.
11. Vlad Nae, Alexandru Iosup, Radu Prodan: Dynamic Resource Provisioning in Massively Multiplayer Online Games. IEEE Trans. Parallel Distrib. Syst. 22(3): 380-395 (2011)
12. Vlad Nae, Radu Prodan, Alexandru Iosup, Thomas Fahringer: A new business model for massively multiplayer online games. ICPE 2011: 271-282



