MASSIVIZING COMPUTER SYSTEMS

MASSIVIZING ONLINE GAMES -OR- HOW I STOPPED WORRYING AND LEARNED TO LOVE DOING RESEARCH IN GAMING SYSTEMS

@Large Research
Massivizing Computer Systems

http://atlarge.science

bit.ly/MassivizingGames22

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

Sponsored by:
NWO

VU VRIJE UNIVERSITEIT AMSTERDAM

Contributions from the MCS team. Many thanks!
Many thanks to our collaborators, international working groups, authors of all images included here.
Also to Annalise Fishell, Study Association STORM for the invitation!

Prof.dr.ir. Alexandru IOSUP

US IN 1 MINUTE

WE'RE MASSIVIZING COMPUTER SYSTEMS!

VU AMSTERDAM < SCHIPHOL < THE NETHERLANDS < EUROPE

Amsterdam
founded 10th century
pop: 850,000

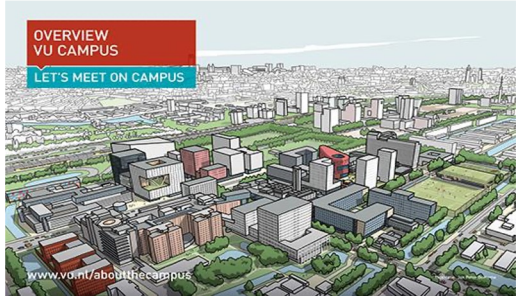

VU
founded 1880
pop: 23,500

The Netherlands

Europe

VU UNIVERSITY AMSTERDAM

http://atlarge.science

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft

CURRENT TEAM

WE ARE LOOKING FOR A NEW ASST. PROF.!

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

ALUMNI

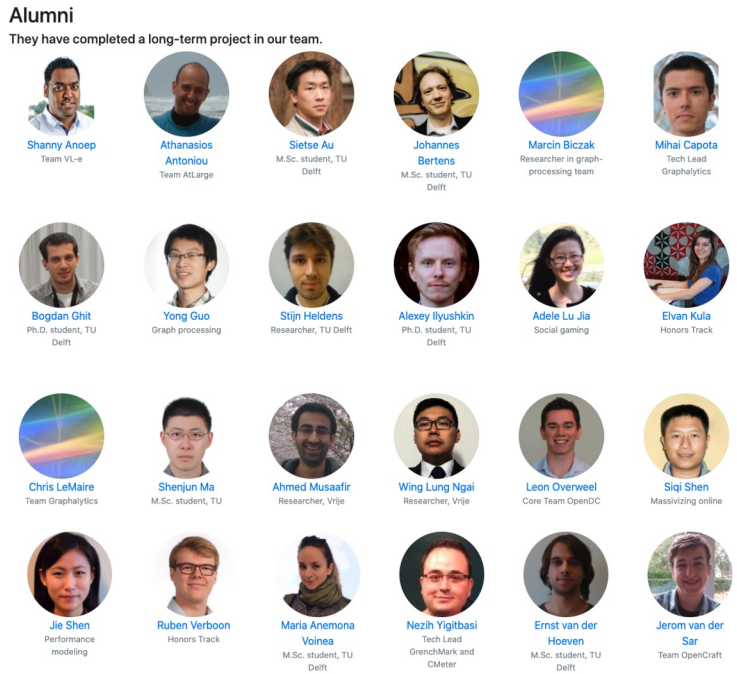

VISITORS

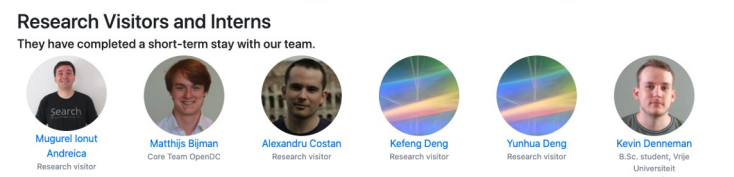

WE ARE A FRIENDLY, DIVERSE GROUP, OF DIFFERENT RACES AND ETHNICITIES, GENDERS AND SEXUAL PREFERENCES, VIEWS OF CULTURE, POLITICS, AND RELIGION. YOU ARE WELCOME TO JOIN!

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.

VU VRIJE UNIVERSITEIT AMSTERDAM

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

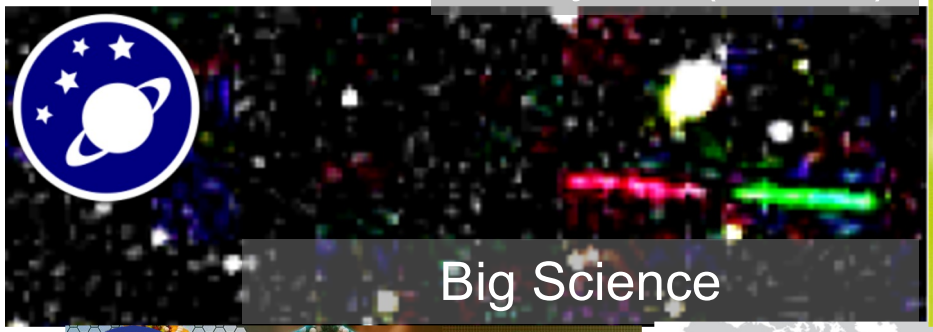

Big Science

Internet

5G/6G/*

THz

VLC

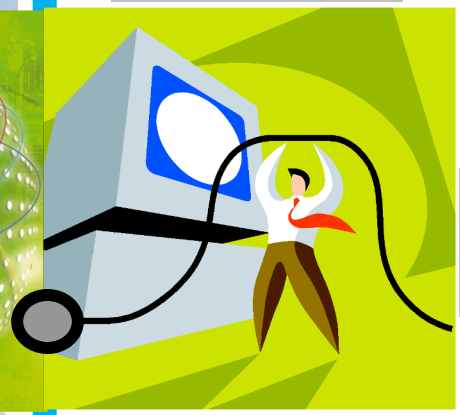

Cloud Computing

Online Gaming

Datacenter

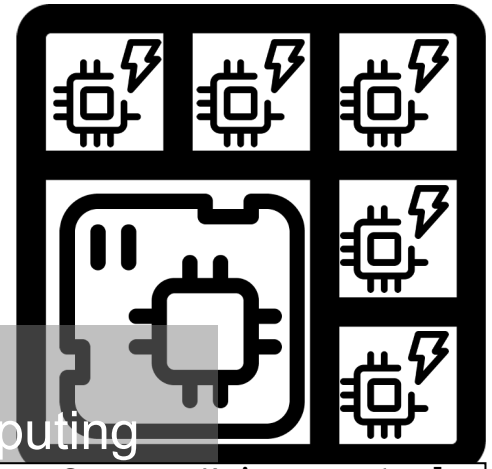

Grid Computing

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

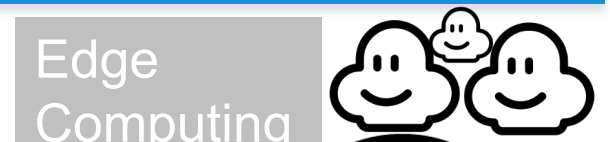

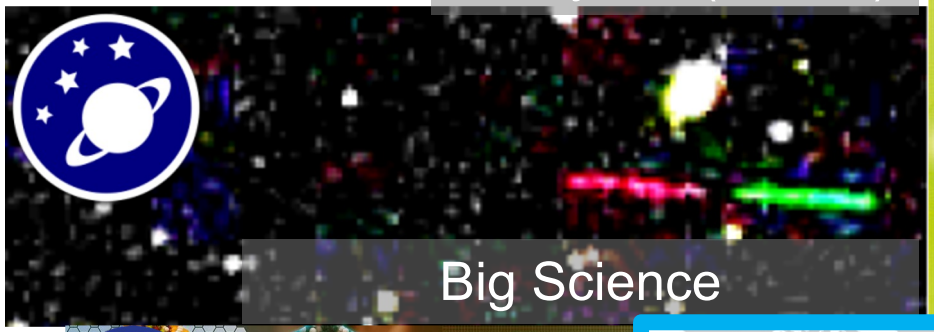

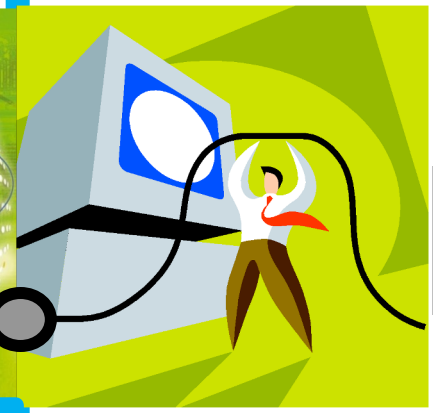

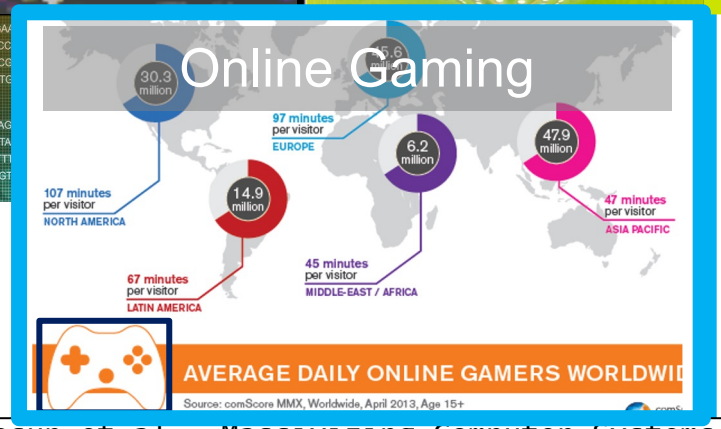

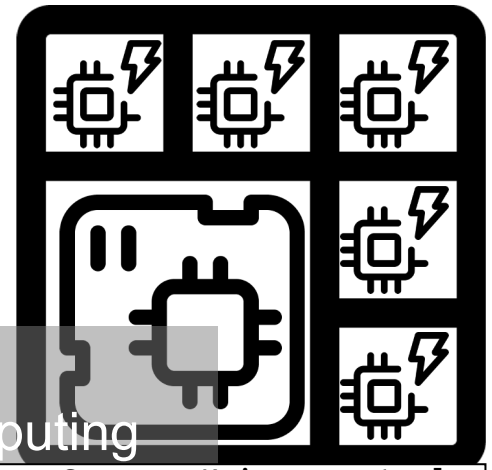

Daily Life

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 GAMES, LOTS OF CULTURE, LOTS OF REVENUE

65%
Puzzle
(e.g., *Tetris, Candy Crush Saga*)

57%
Arcade & Other
(e.g., *Pac-Man, Super Mario Party, Just Dance*)

46%
Skill & Chance
(e.g., *Solitaire, Bingo*)

43%
Action
(e.g., *Legend of Zelda, Uncharted*)

42%
Shooter
(e.g., *Call of Duty, Halo*)

41%
Simulation
(e.g., *The Sims, Animal Crossing*)

41%
RPG & Narrative
(e.g., *The Witcher, Skyrim*)

40%
Strategy
(e.g., *XCOM, Clash Royale*)

38%
Racing & Vehicle Simulation
(e.g., *Forza, Microsoft Flight Simulator*)

30%
Fighting
(e.g., *Street Fighter, Super Smash Bros.*)

28%
Sports
(e.g., *Madden NFL, NBA 2K*)

$177.8 BN
2020

2022 Essential Facts About the Video Game Industry. Source: The Entertainment Software Association.

LOTS OF GAMERS, LOTS OF CULTURE, LOTS OF REVENUE

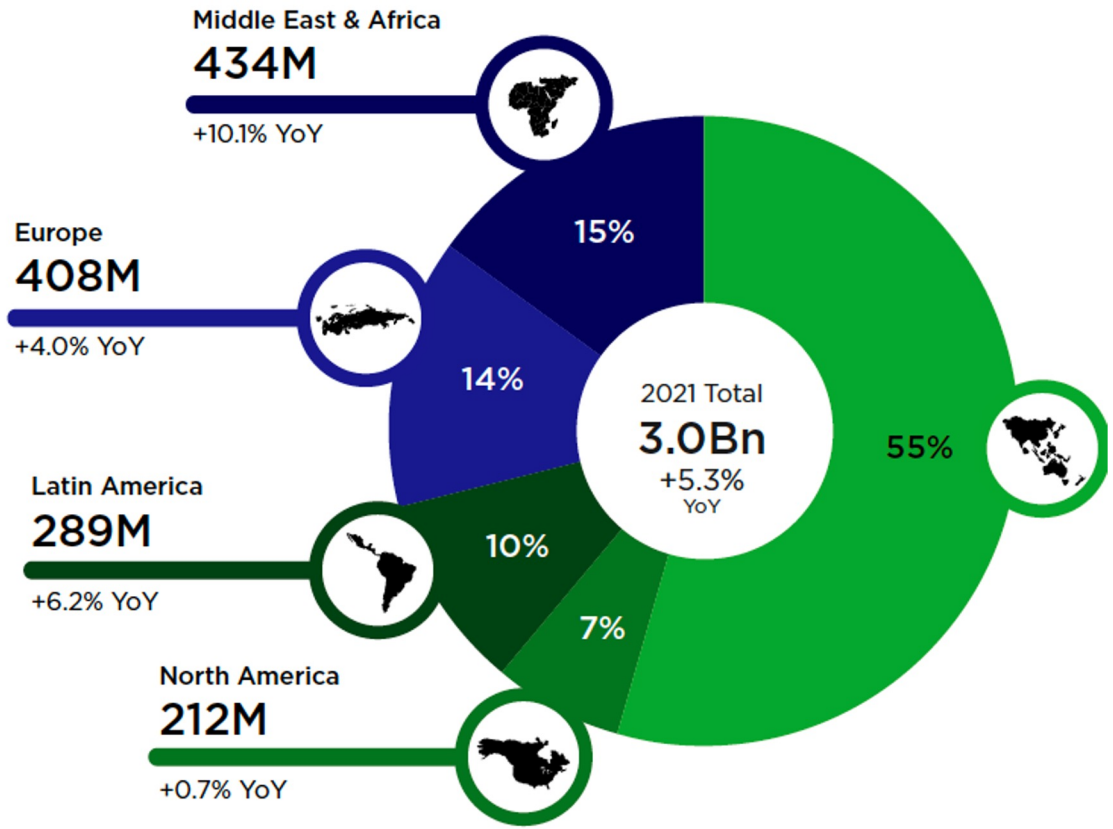

2021 Global Players. Source: Newzoo

2022 Essential Facts About the Video Game Industry. Source: The Entertainment Software Association.

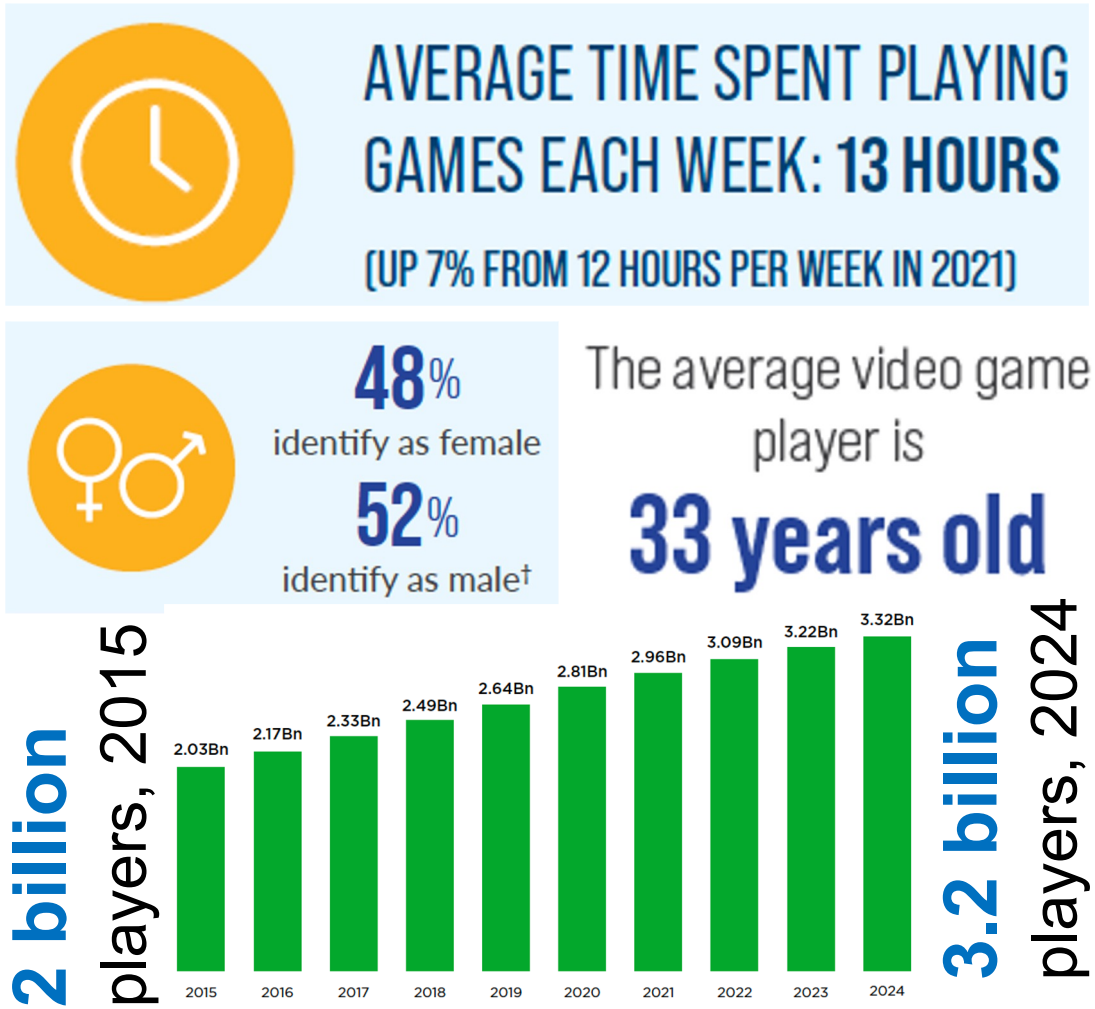

GAMES ARE INCREASINGLY MORE SOCIAL, BUILD COMMUNITIES

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.

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

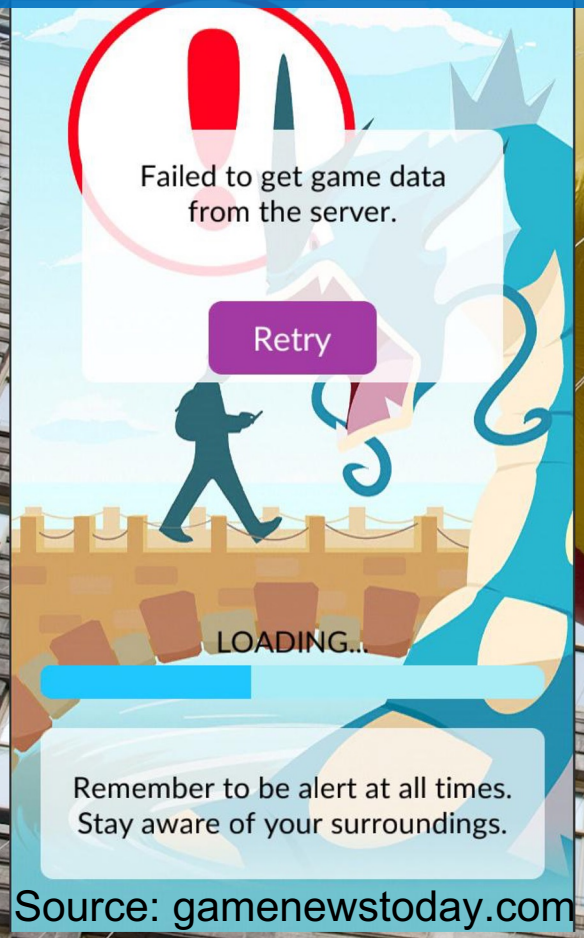

Source: gamenewstoday.com

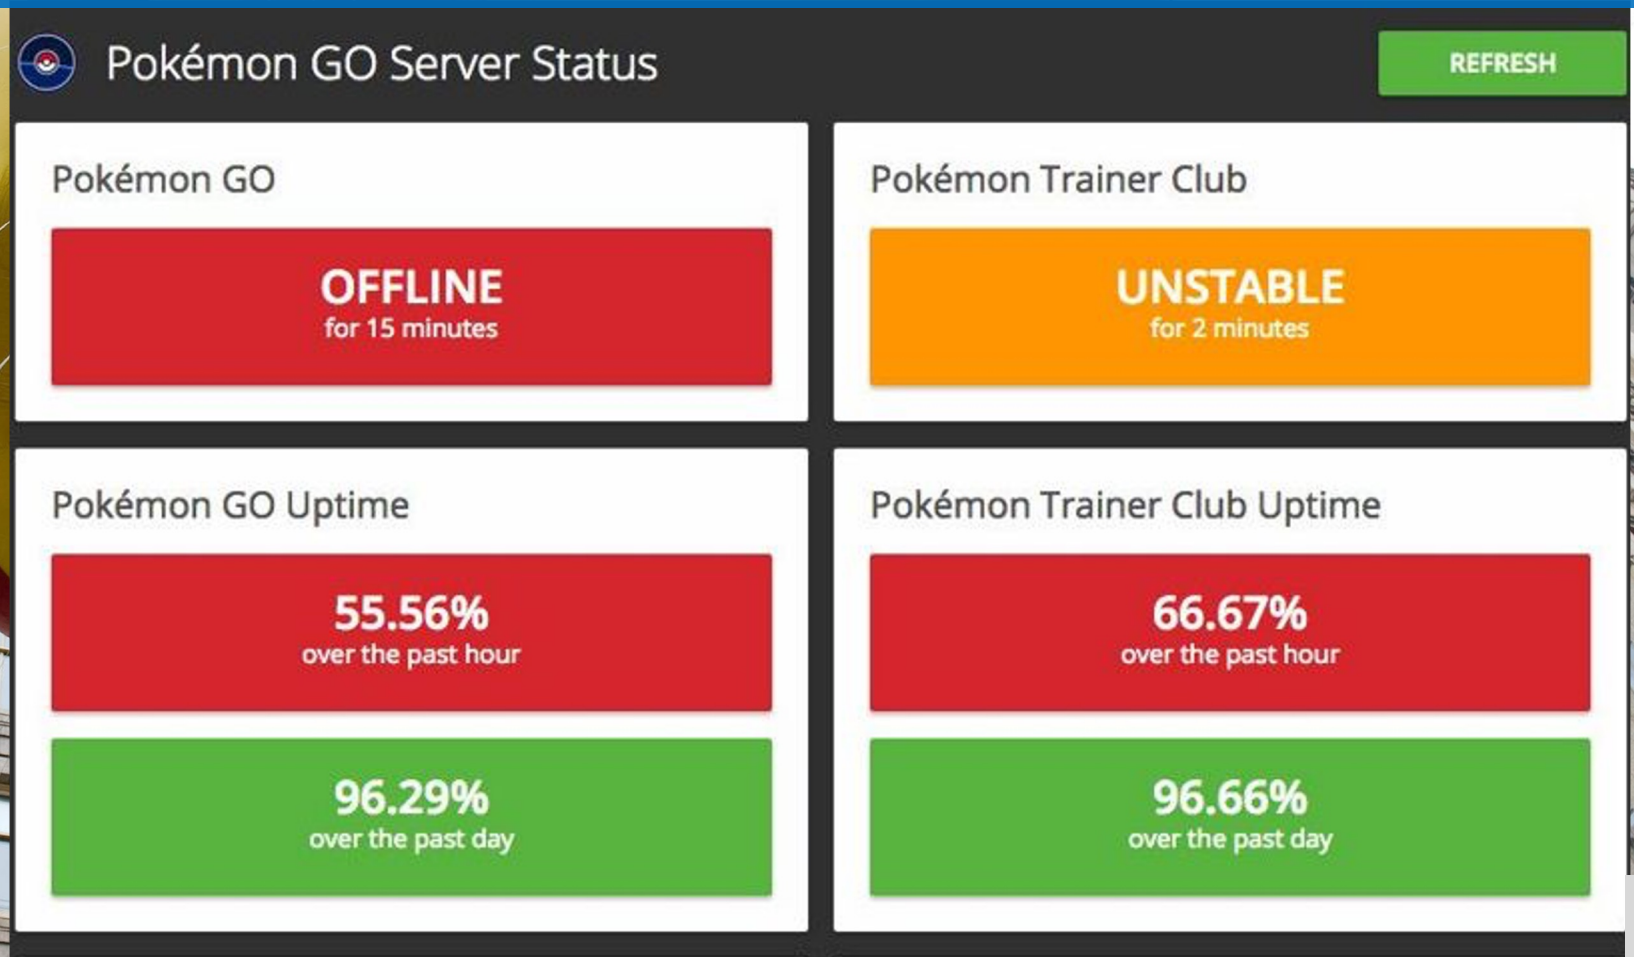

PHENOMENON: PERFORMANCE DROPS IN GAMES

UNCOVERING THE PRESENCE OF PERFORMANCE ISSUES, EVEN LEADING TO CRASHES

Polygon

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

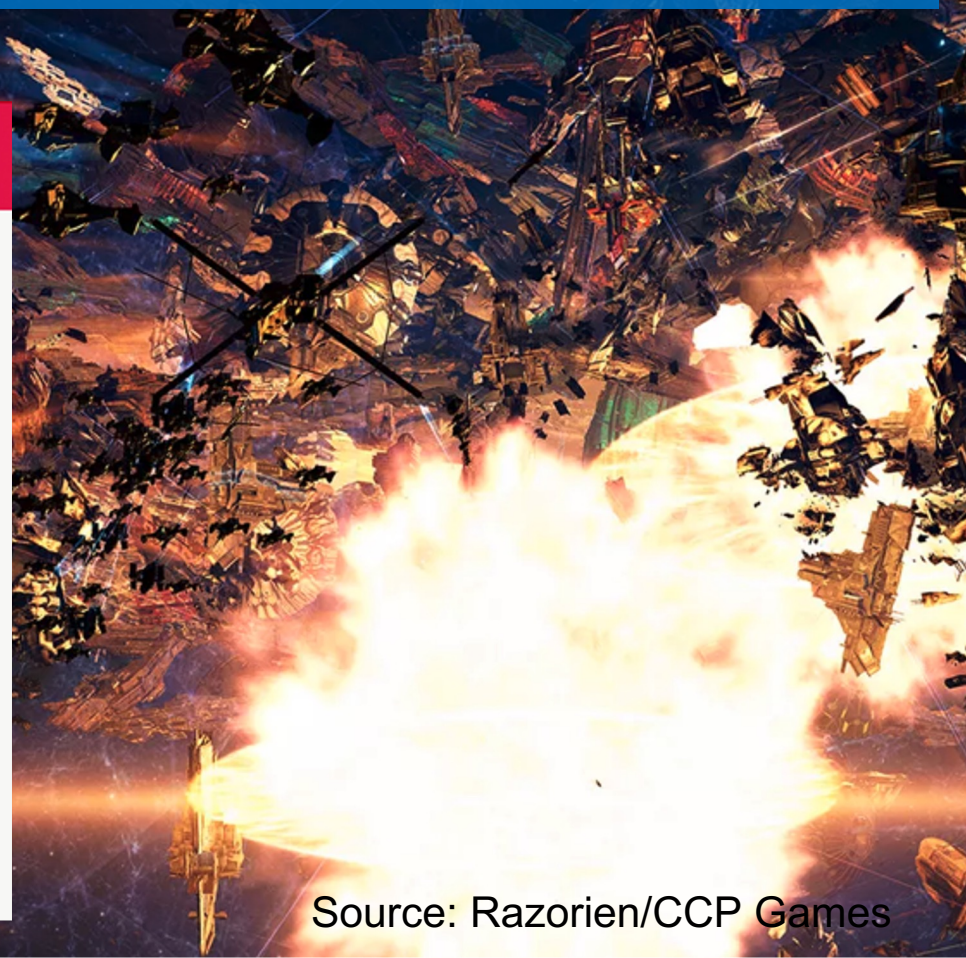

Source: Razorien/CCP Games

PHENOMENON: GAME SERVICE SUSTAINABILITY

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

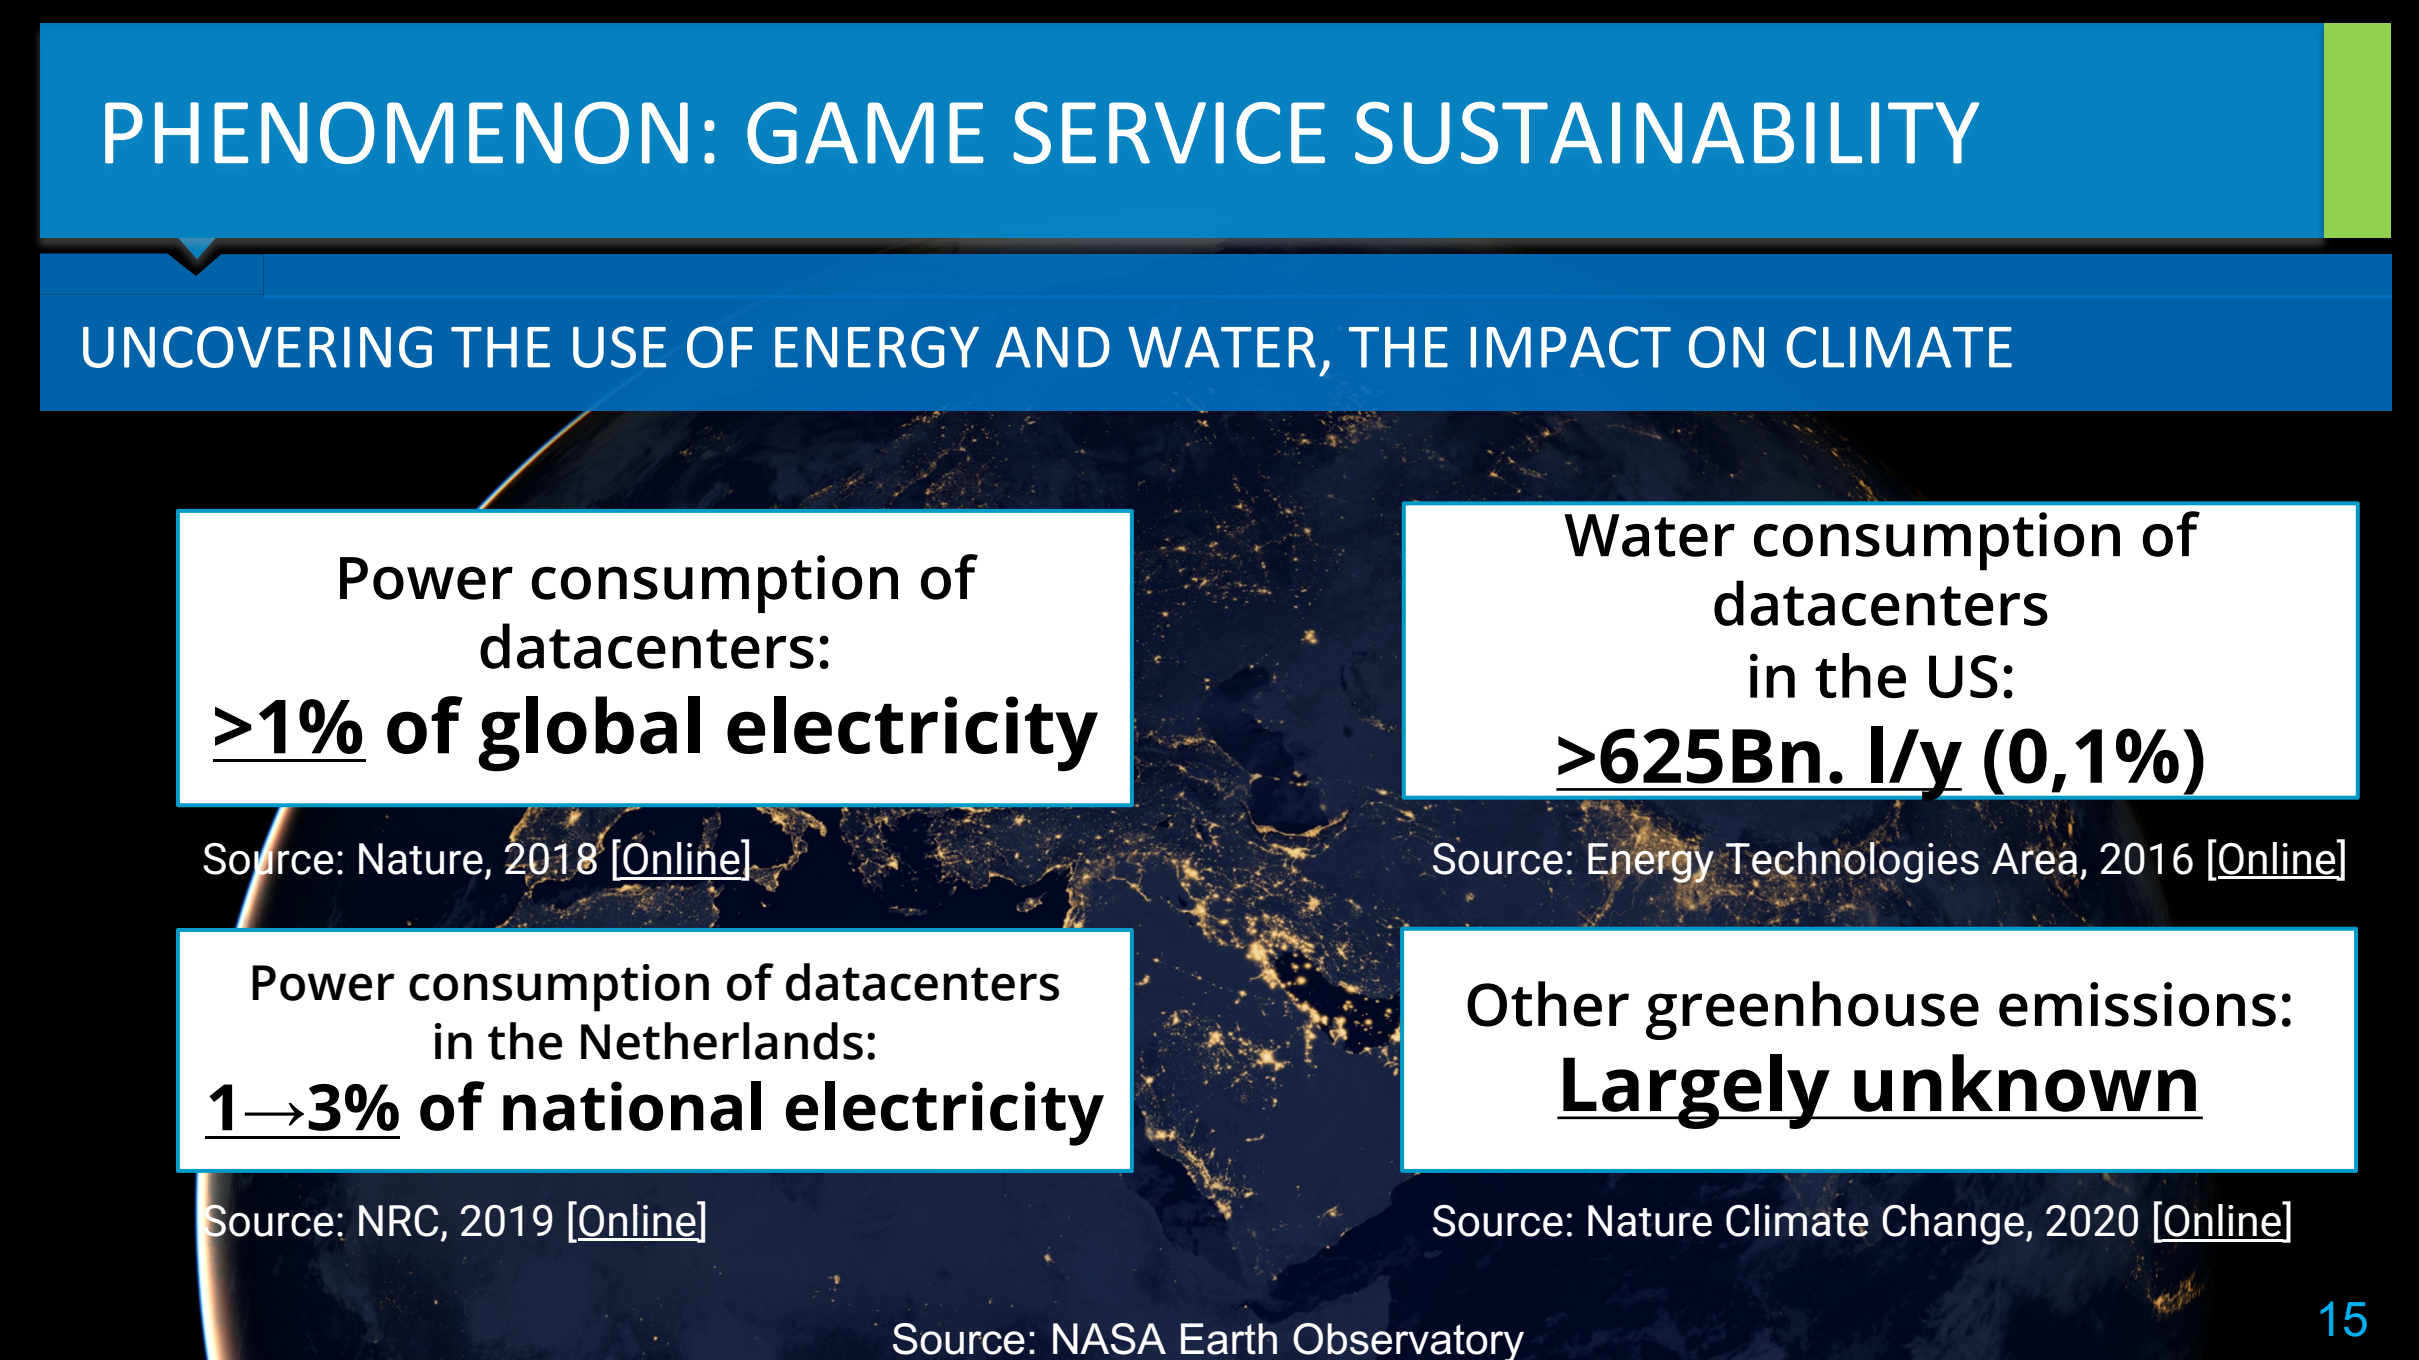

Power consumption of datacenters:
>1% of global electricity

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]

Source: NASA Earth Observatory

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

SHARE

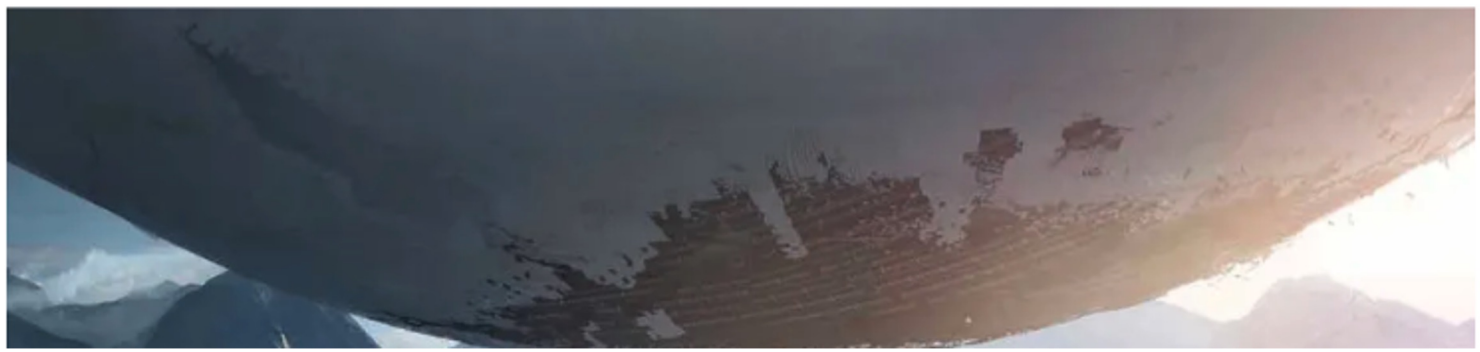

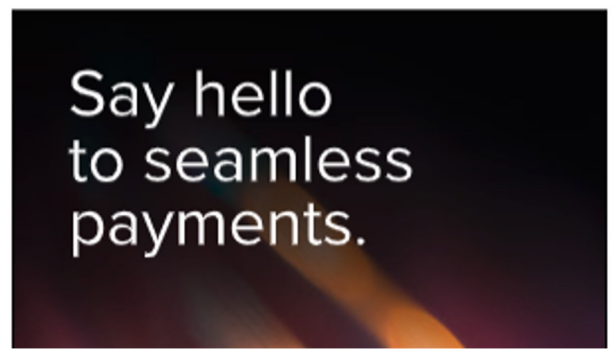

THIS TALK, IN A NUTSHELL:

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 computing

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

!

A new science, of complex, smart computer ecosystems

3

(operational simplicity for the user)

AN ANALOGY: MASSIVIZING CLIMATE SCIENCE

TAKE A HOLISTIC VIEW, BASED ON COUPLED NATURAL SYSTEMS

Can be understood only with coupled models

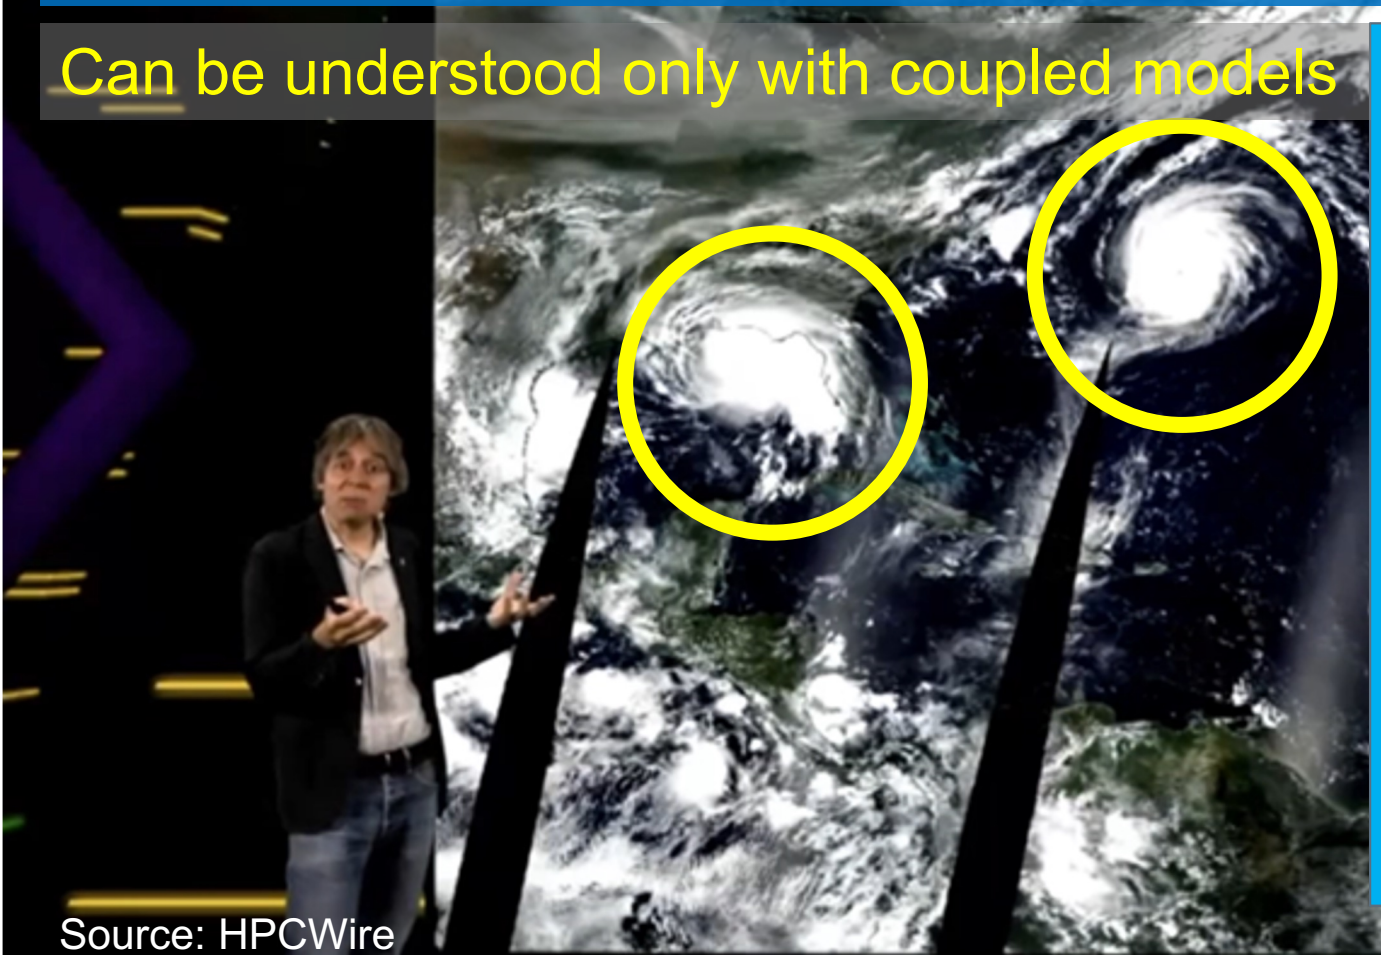

Source: HPCWire

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

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-TRIVIALLY FROM ITS SYSTEMS AND USERS (RECURSIVELY)
WHICH APPROACH?	MODERN DISTRIBUTED SYSTEMS AND PROBLEM-SOLVING

Iosup et al., Massivizing Computer Systems, ICDCS 2018. [Online]

(ONLINE) GAMES OF MASSIVE COMPUTING ECOSYSTEMS

Do you recognize this App?

Here is how it operates…

A TYPICAL ECOSYSTEM: SERVICE, DATACENTER, SCHEDULER

Creators

Digital Services

Workload

Time

100s of services … an ecosystem

Resource Manager and Scheduler

(full-/micro-/nano-) Datacenter

Performance, Dependability, Sustainability

LET'S UNDERSTAND HOW THINGS WORK!

4

Message: 4_A

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

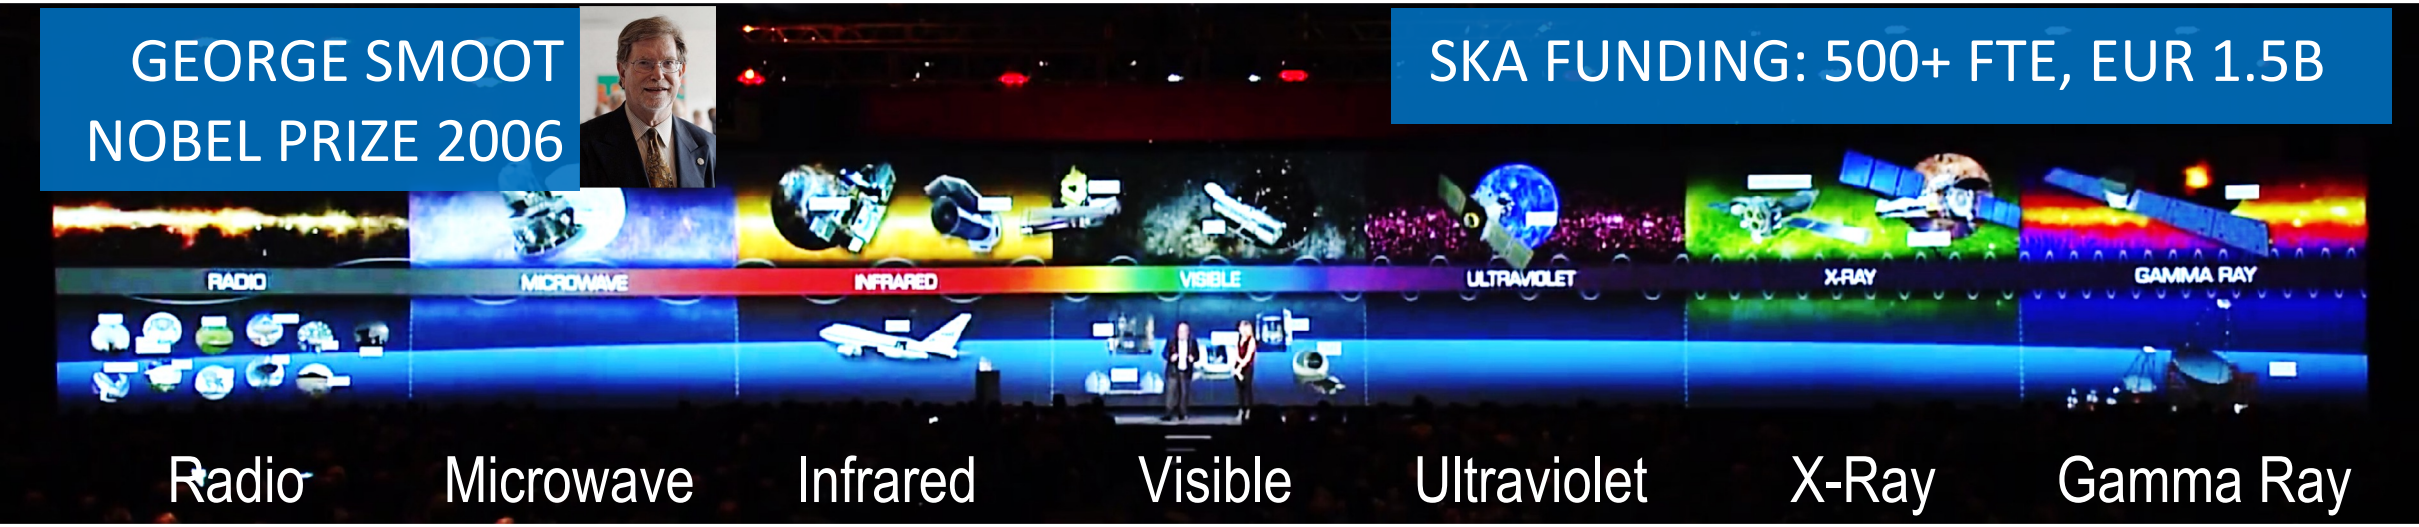

James Cordes, The Square Kilometer Array, Project Description, 2009 [Online]

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

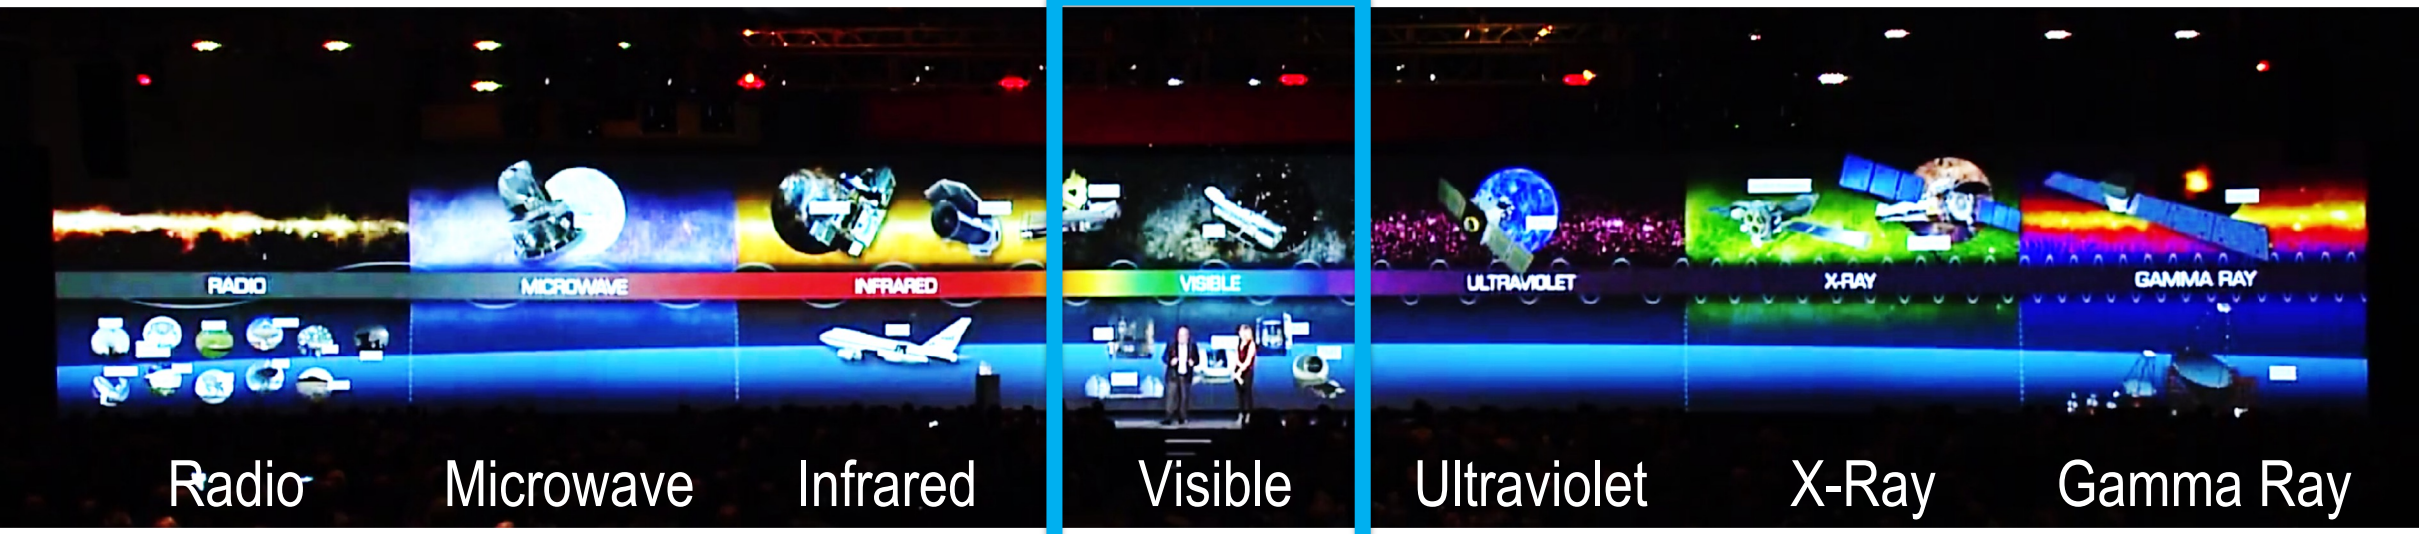

Radio	Microwave	Infrared	Visible	Ultraviolet	X-Ray	Gamma Ray
Cloud, Grid, Edge, Fog, etc.	One aspect: BigData, P2P	Sci.&Eng. Apps+Sys.	Consumer Apps+Sys.	Enterprise Sys.	Systems, Ecosystems	Performance, Availability,
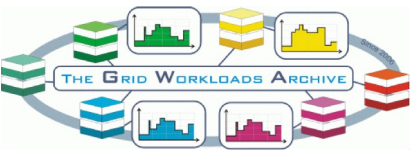	P2P T.A.	[Versluis et al. TPDS'21] Groups, Workflows	Game Trace Archive	Business - Critical		
[Iosup et al. EGCS'08]	[Zhang et al. CoNext'10]	[Iosup et al. IEEE IC'11]	Guo et al. NETGAMES'12]	[Shen et al. CCGRID'15]	[Ghit et al. CCGRID'14]	[Iosup et al. CCGRID'10]

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

Cloud, Grid, Edge, Fog, etc.	One aspect: BigData, P2P	Sci.&Eng. Apps+Sys.	Consumer Apps+Sys.	Enterprise Sys.	Systems, Ecosystems	Performance, Availability,
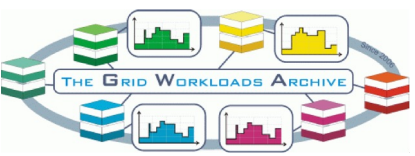				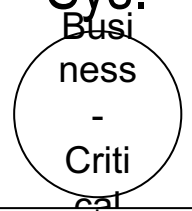		
[Iosup et al. EGCS'08]	[Zhang et al. CoNext'10]	[Versluis et al.'19] [Iosup et al. IEEE IC'11]	Guo et al. NETGAMES'12]	[Shen et al. CCGRID'15]	[Ghit et al. CCGRID'14]	[Iosup et al. CCGRID'10]

Message: 4_B Synthesize the ecosystem reference architecture

MEANINGFUL DISCOVERY

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

FUNDING: > 3B USD

- 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

International Human Genome Sequencing Consortium, Initial sequencing and analysis of the human genome, Nature 409, Feb 2011. [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

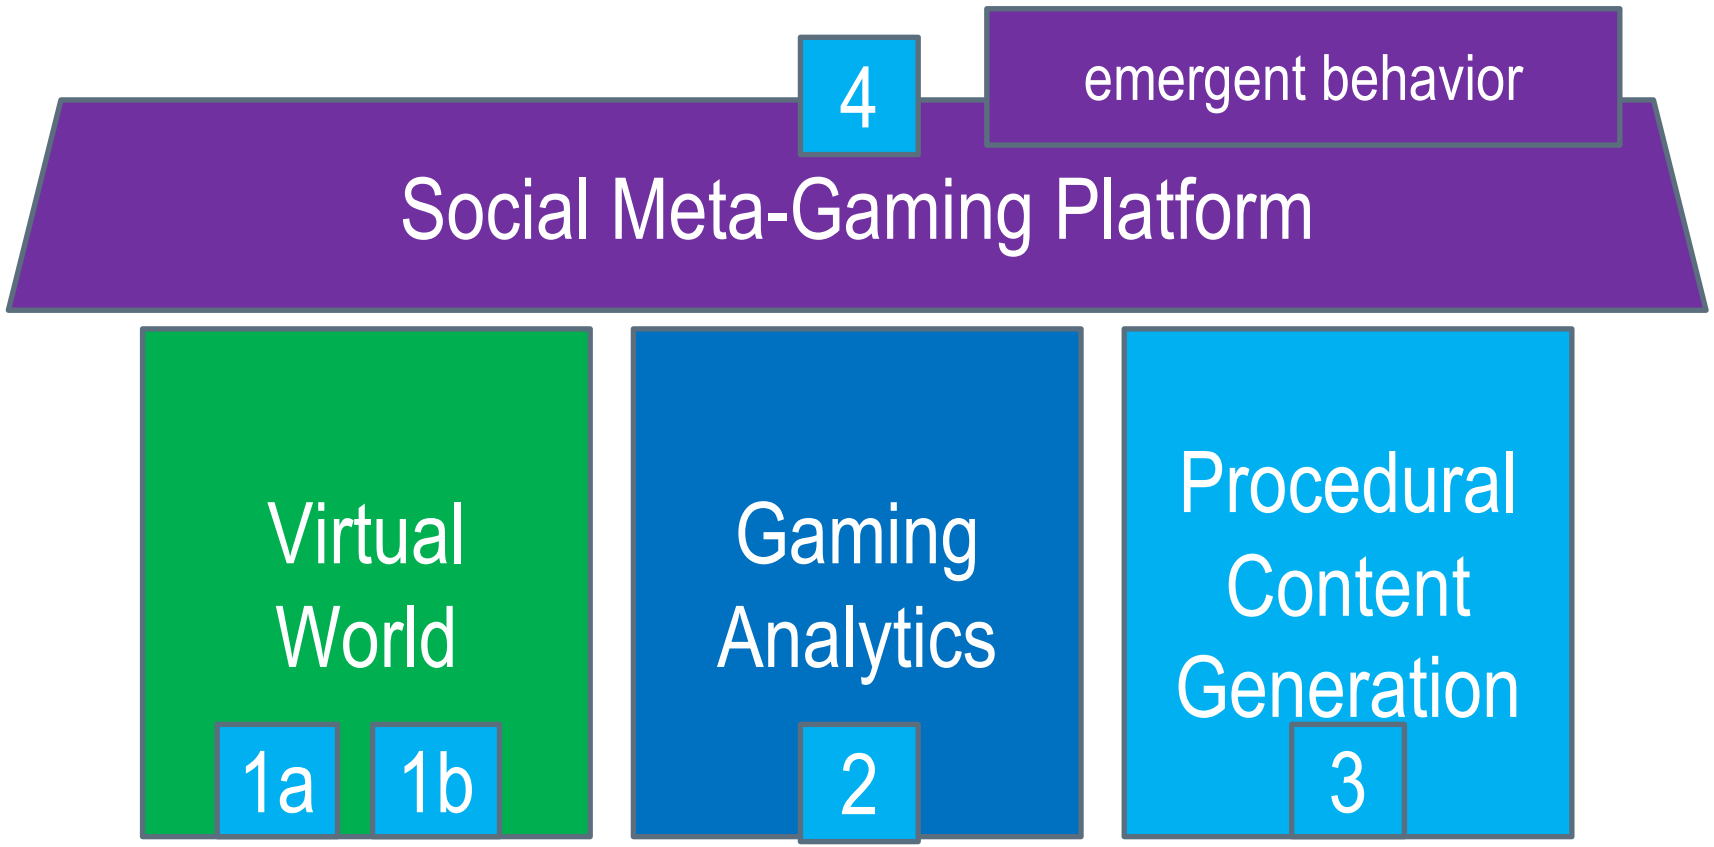

Iosup, Shen, Guo, Hugtenburg, Donkervliet, Prodan (2014) Massivizing online games using cloud computing: A vision. ICME Workshops [Online]

HOW TO MANAGE SYSTEM COMPLEXITY?

THE COMPLEXITY CHALLENGE

Focus on DevOps + Applications,
5 Core Layers:

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

IOSUP ET AL. REFERENCE ARCHITECTURE FOR DCS, 2016

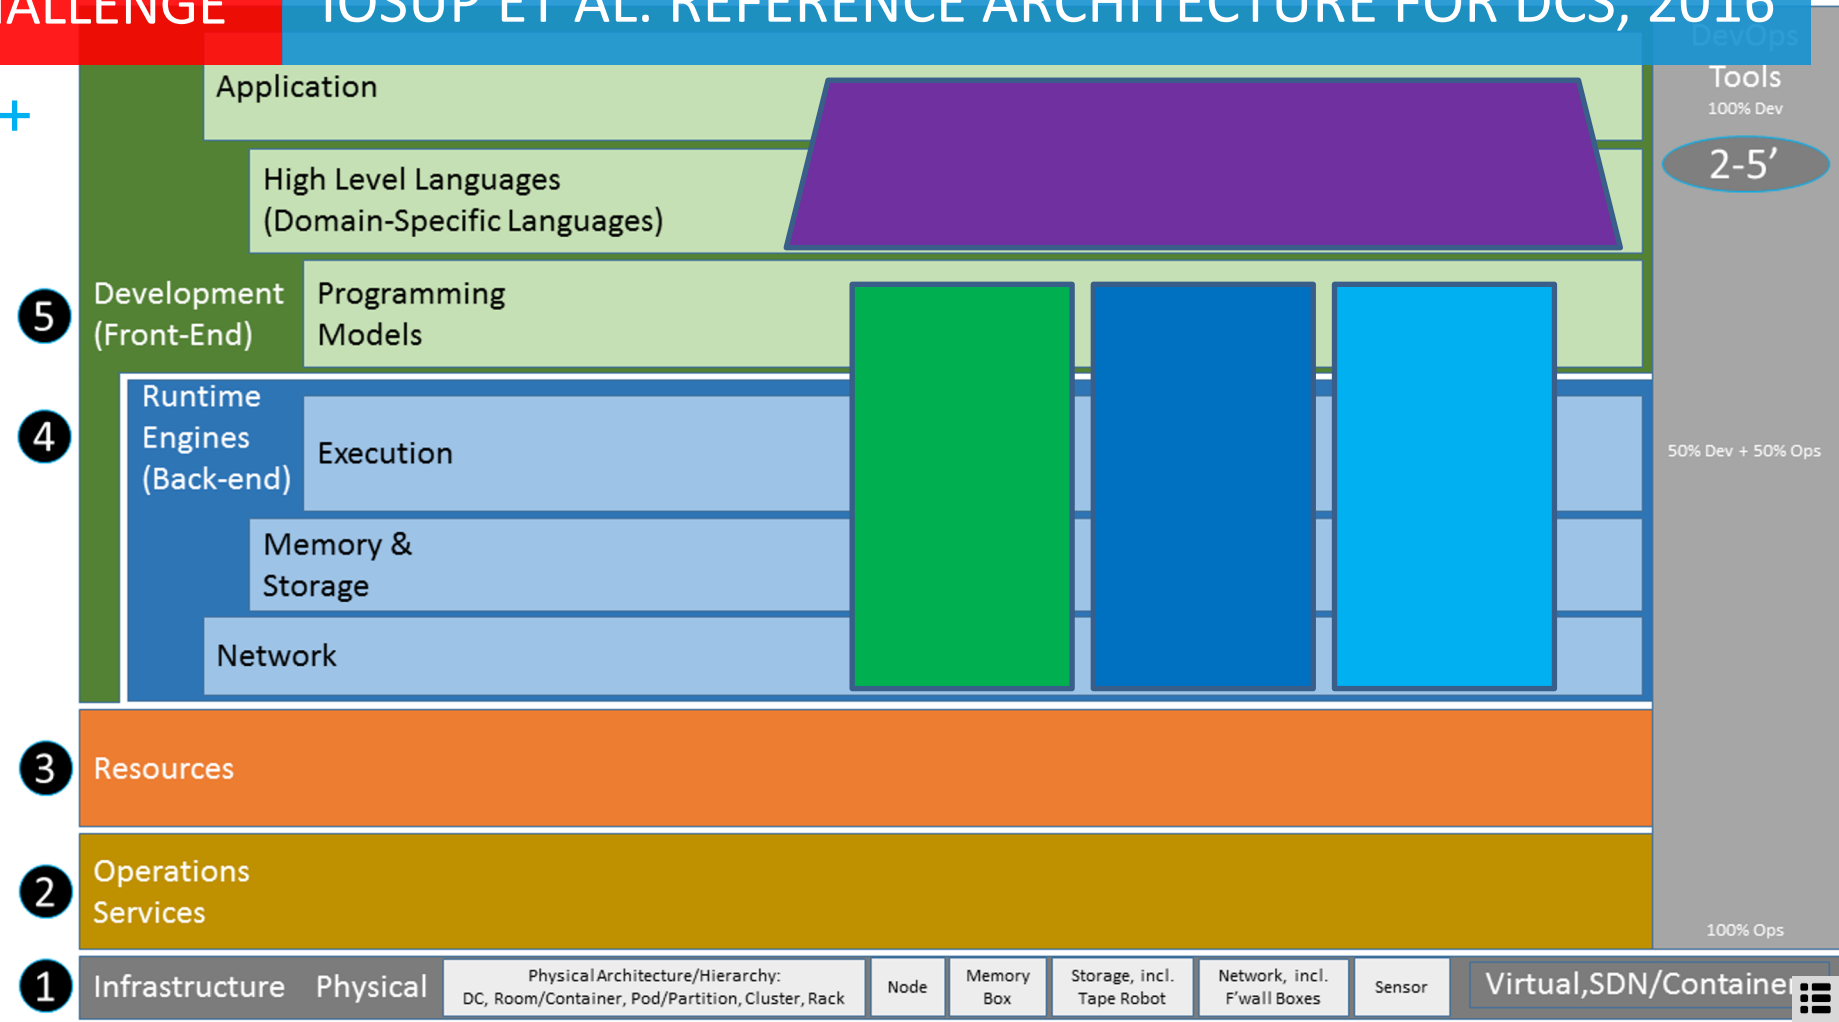

[Iosup et al. ICDCS'18]

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.

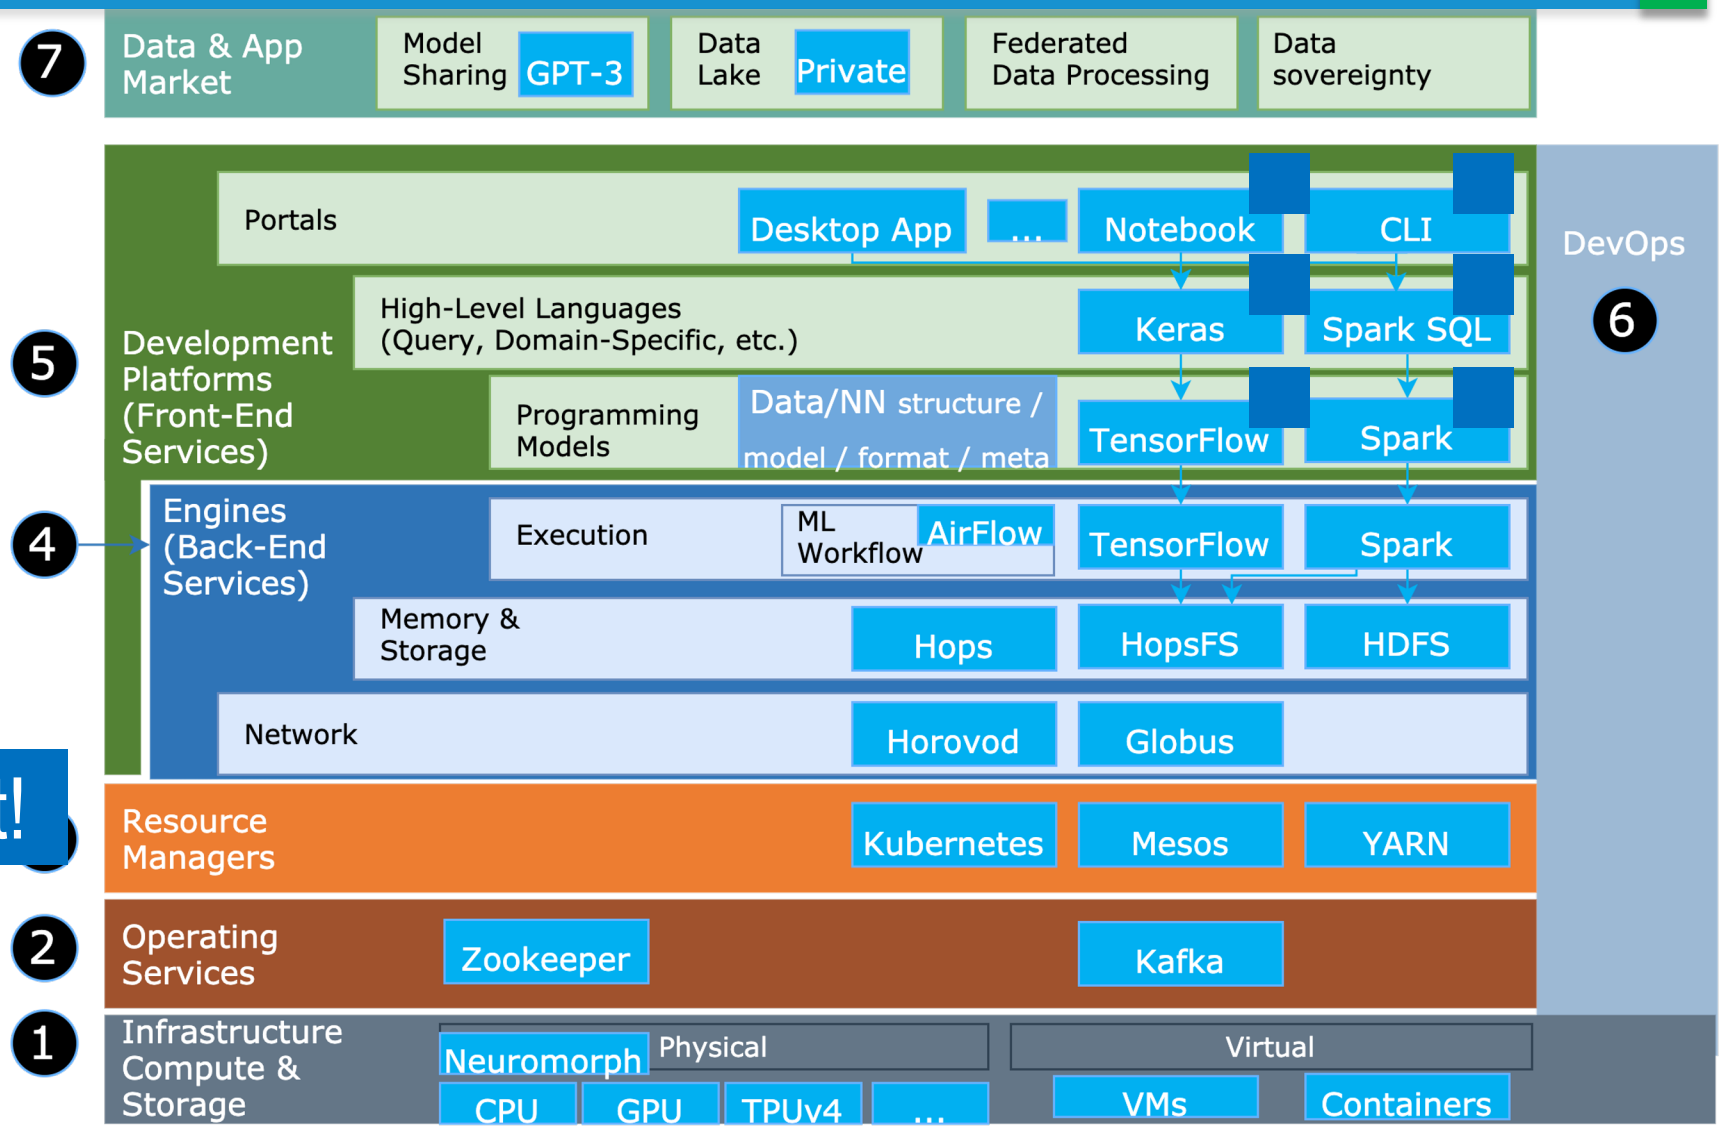

THE SUPER-DISTRIBUTION PRINCIPLE

THE COMPLEXITY CHALLENGE

RECURSIVE ECOSYSTEMS, 2018

Georgios Andreadis

ANDREADIS ET AL. REFERENCE ARCHITECTURE FOR SCHEDULERS IN DCS

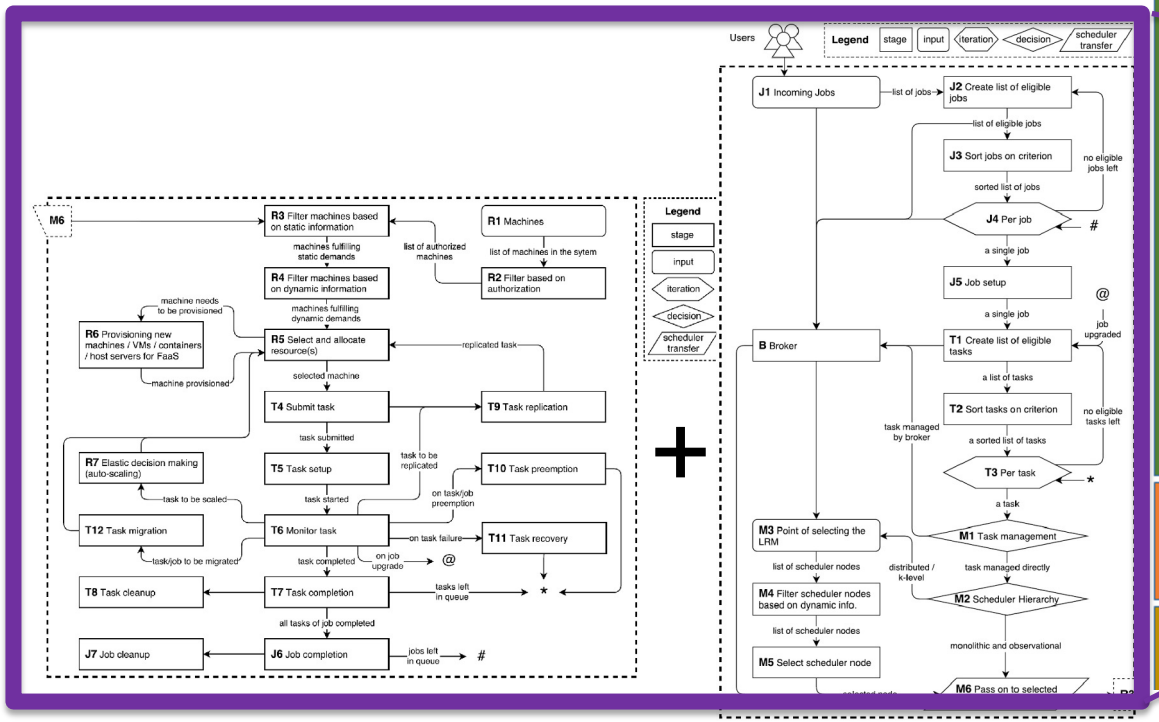

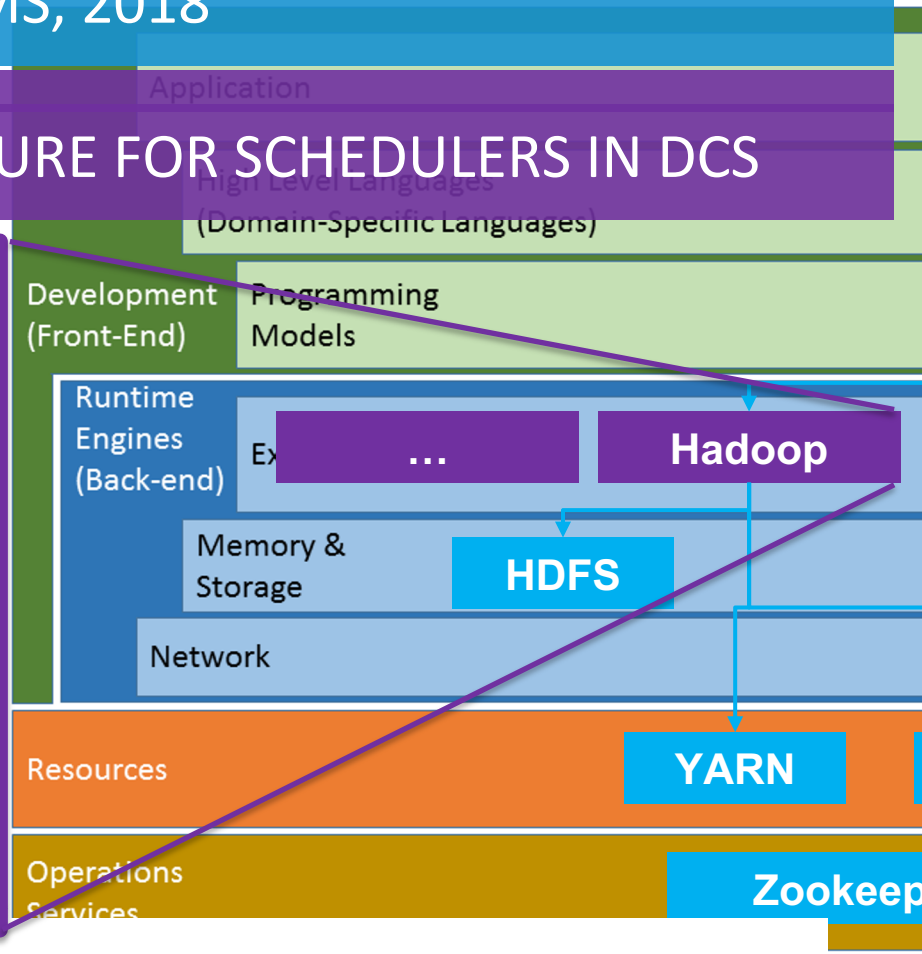

[Andreadis et al. SC'18]

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 phyisico/in vitro*)

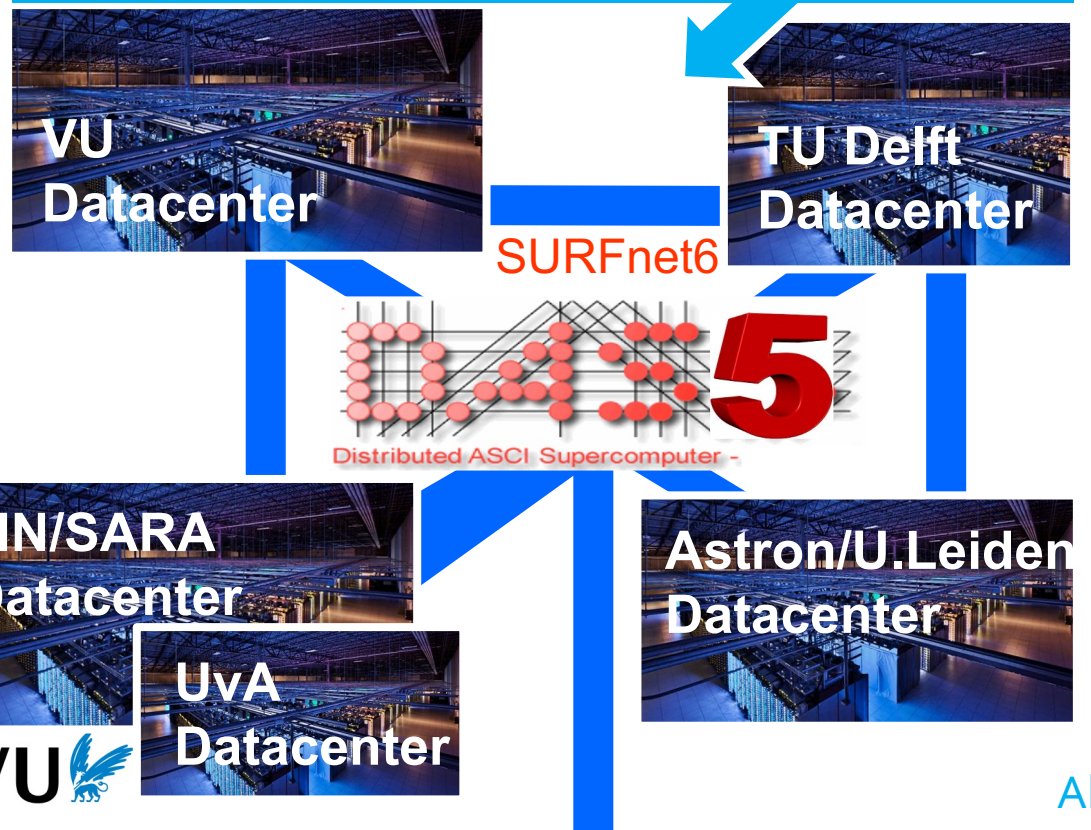

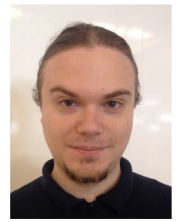

Alex Uta

Georgios Andreadis

Fabian Mastebroek

acheendra

Maria

Laurens Versluis

+

amazon web services™

We also use clouds

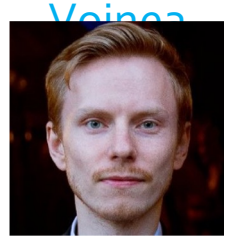

Alexey Ilyushkin

+

OpenDC

And simulators (*in silico*)

… 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

- Short-term resource management
- Long-term capacity planning
- Sophisticated model
- Support for many kinds of workloads and resources
- Validated for various scenarios
- Work with major NL hoster
- Used in training

Fabian Mastenbroek

and more...

Yardstick

A Benchmark for Minecraft-like Services

Jerom van der Sar
Jesse Donkervliet
Alexandru Iosup

Contact: opencraft@atlarge-research.com

www.atlarge

TU Delft

VU VRIJE UNIVERSITEIT AMSTERDAM

Minecraft only scales to hundreds of players

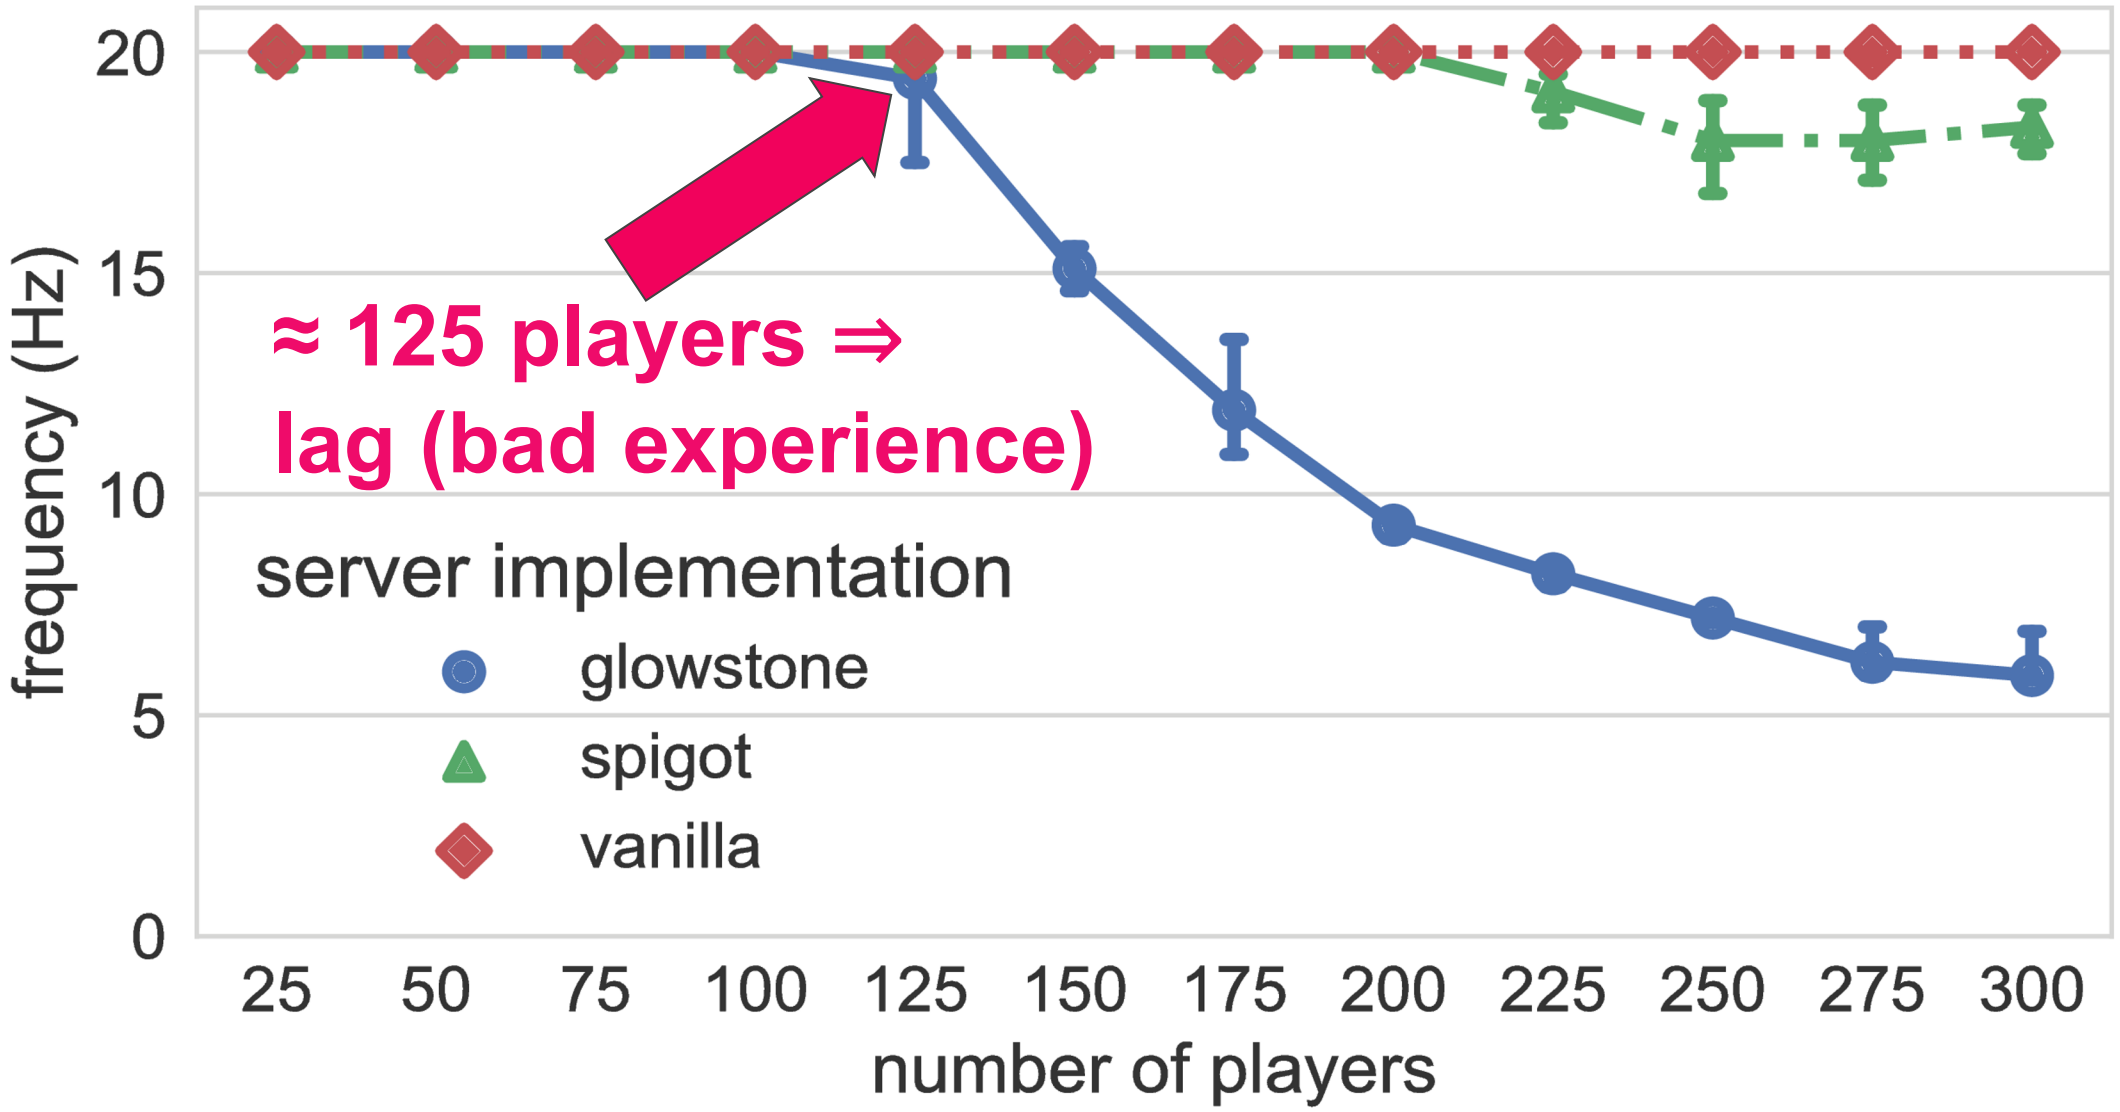

LET'S DESIGN & BUILD IN DISTRIBUTED 5 ECOSYSTEMS!

Message: 5.1 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

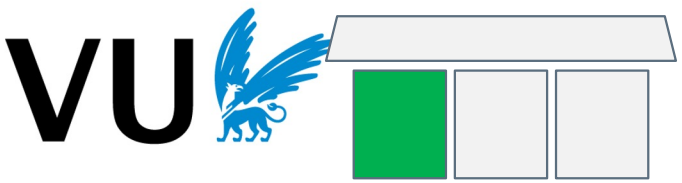

Done

Open challenge

HOW TO SCALE GAMES ACROSS THE ECOSYSTEM? 5.1A

IT'S SCALING TECHNIQUES!

REFERENCE VIEW ON OPERATIONAL TECHNIQUES

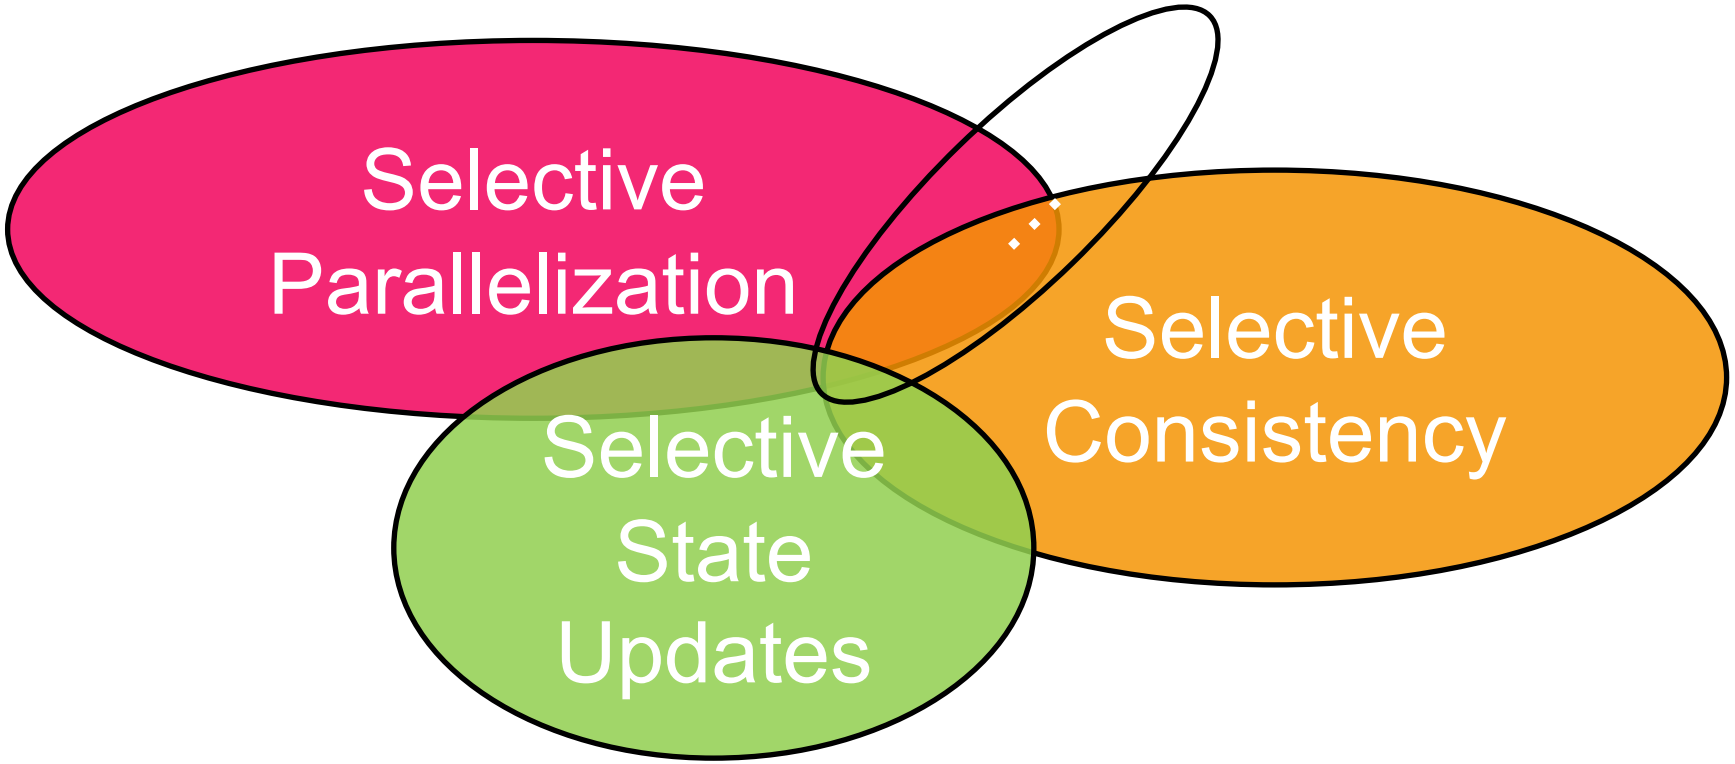

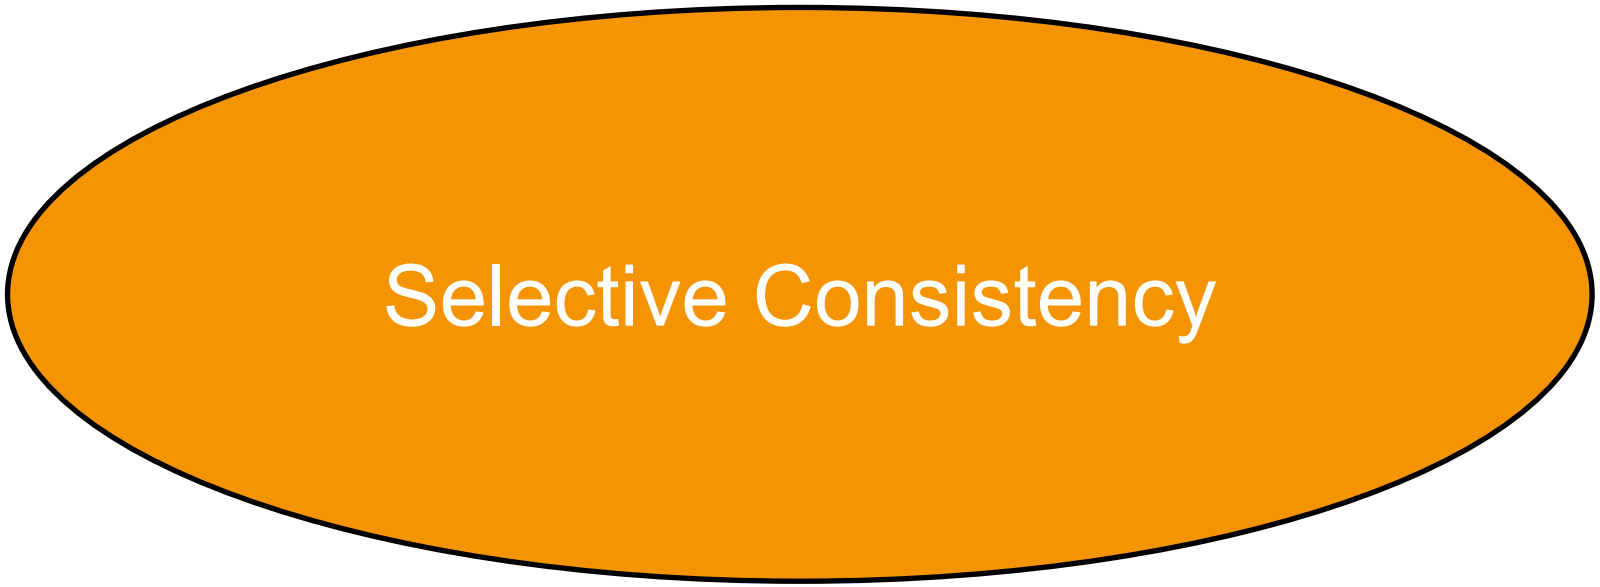
Selective Consistency

CHALLENGE: CONSISTENCY IN HIGH-STAKES APPS

Consistency Example: Current Approach

Tick: 1

Server

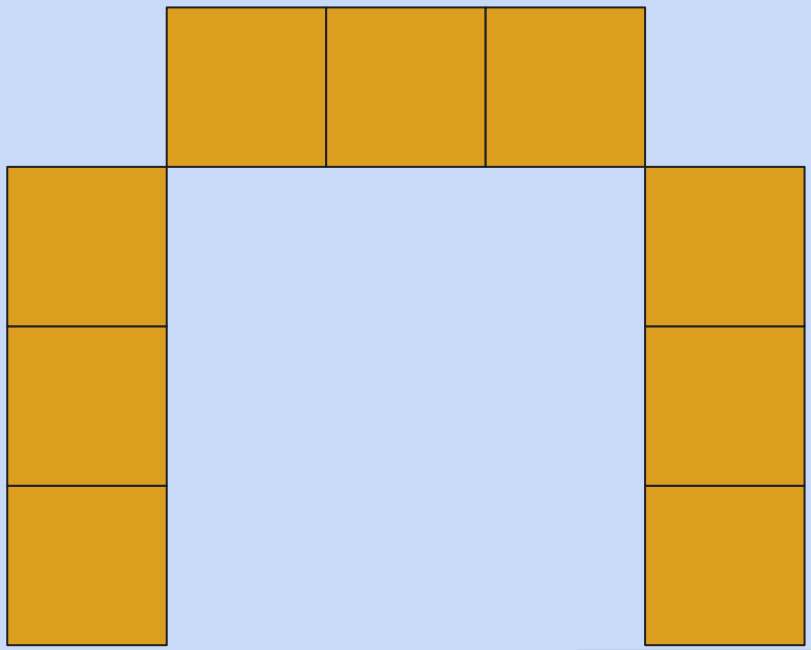

Client

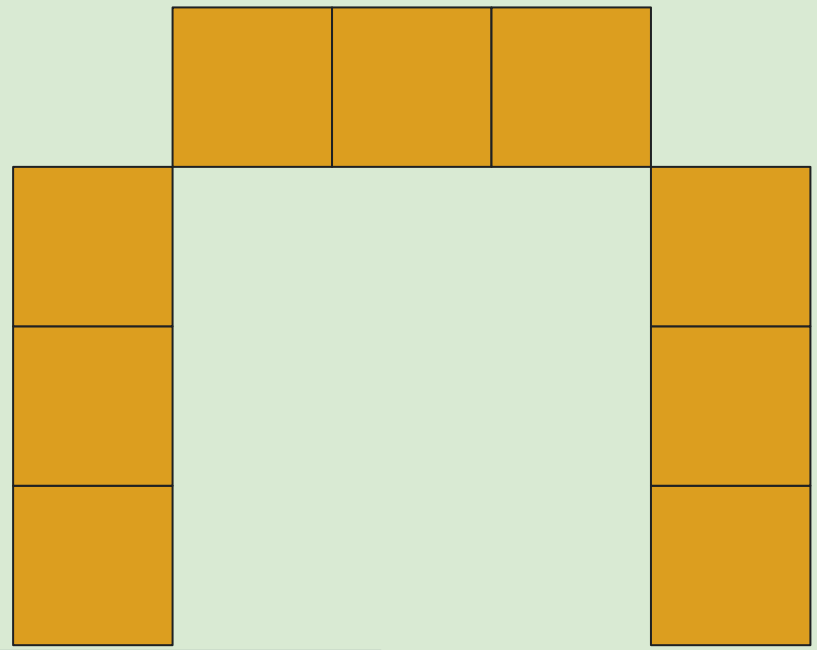

Slides by Jesse Donkervliet

End

CHALLENGE: CONSISTENCY IN HIGH-STAKES APPS

Consistency Example: Current Approach

Tick: 2

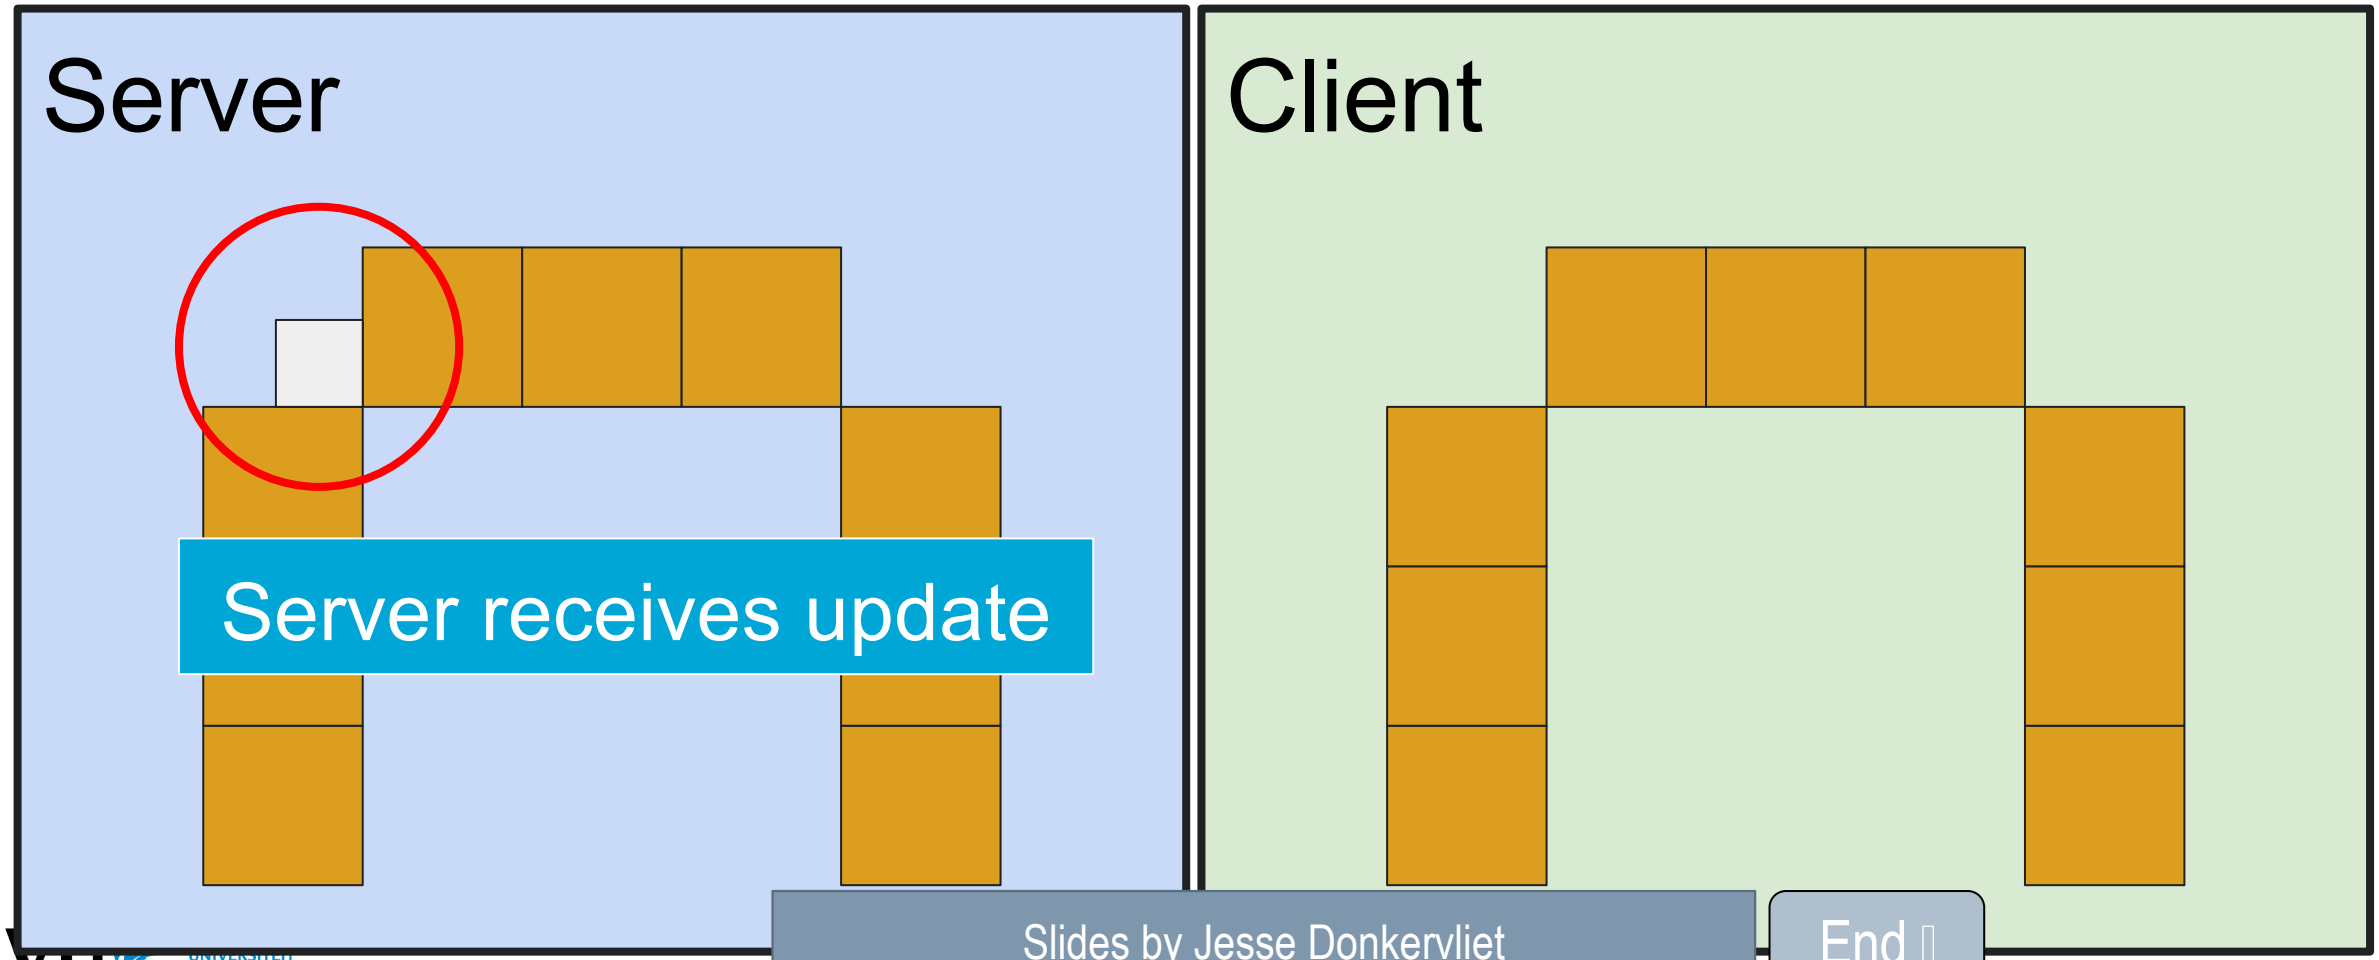

Slides by Jesse Donkervliet

End

CHALLENGE: CONSISTENCY IN HIGH-STAKES APPS

Consistency Example: Current Approach

Tick: 3

Server

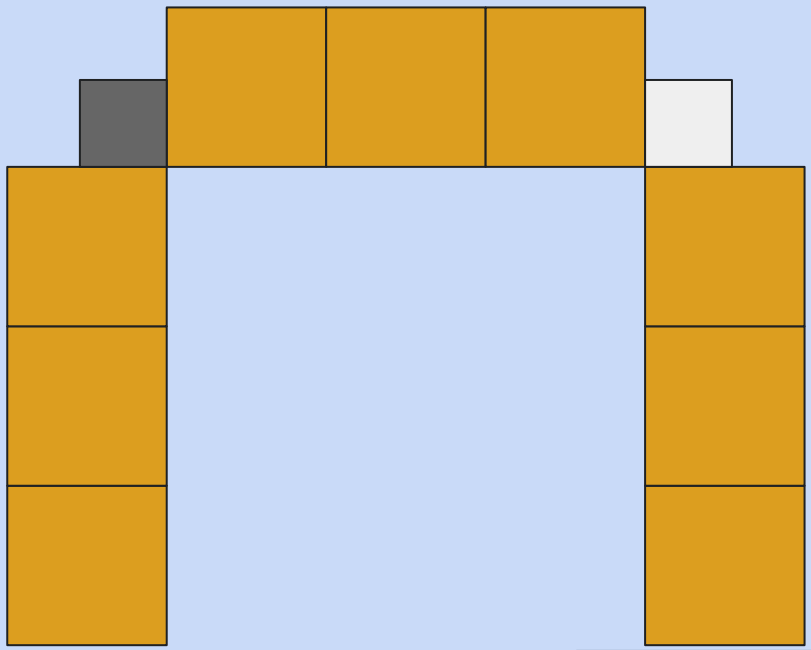

Client

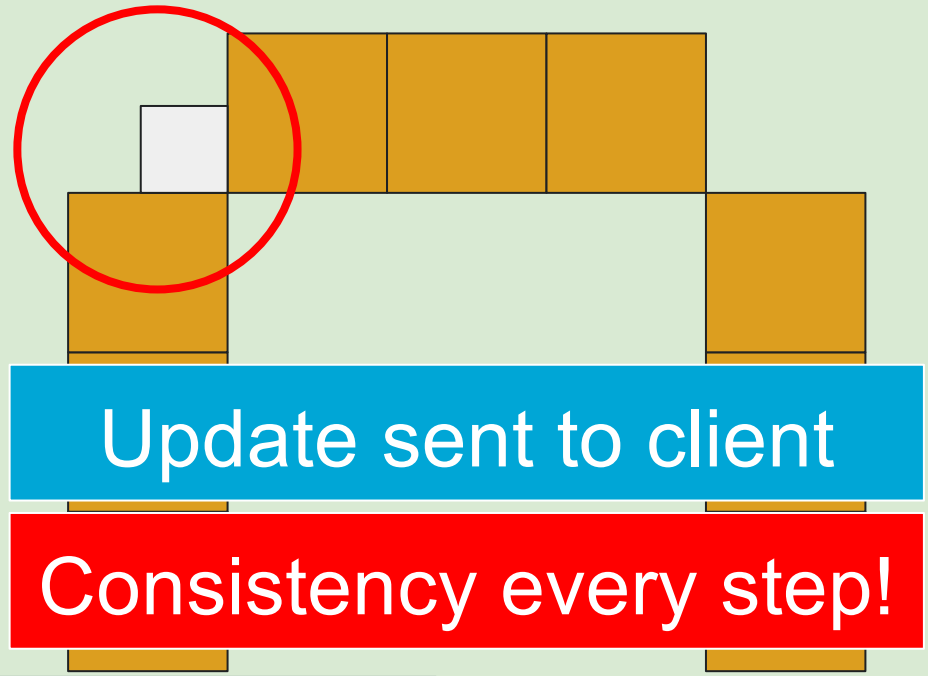

Update sent to client

Consistency every step!

Slides by Jesse Donkervliet

End

Dyconits: Scaling Minecraft-like Services through Dynamically Managed Inconsistency

Jesse Donkervliet, Jim Cuijpers, Alexandru Iosup

jesse.donkervliet@vu.nl
@jdonkervliet
https://atlarge-research.com/opencraft/

Dyconit Policies

Here be algorithms!

Use Feedback to Improve Operation

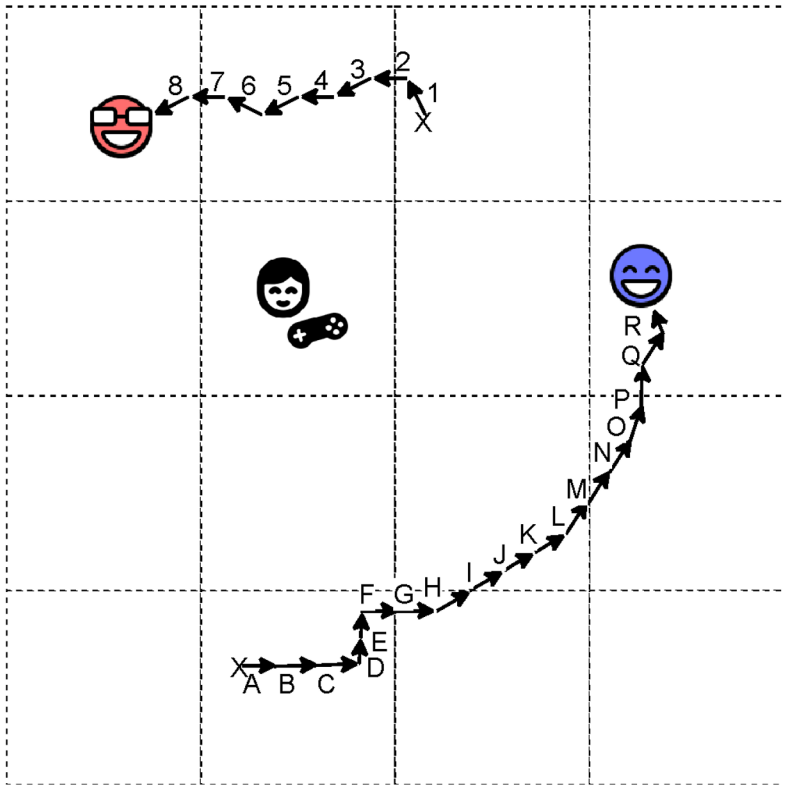

(a) No Dyconits

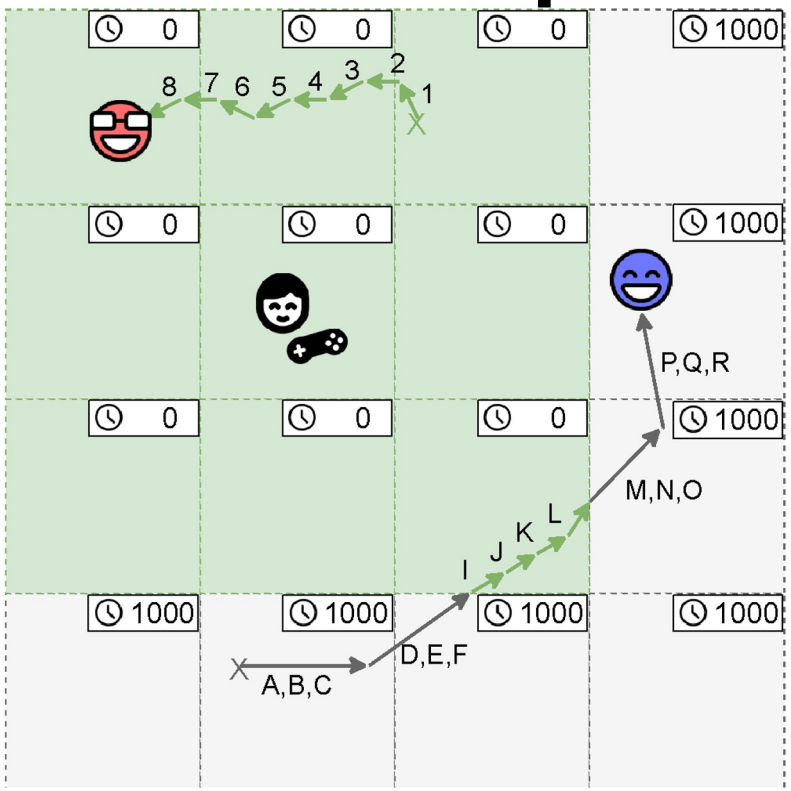

(b) Area Of Interest policy

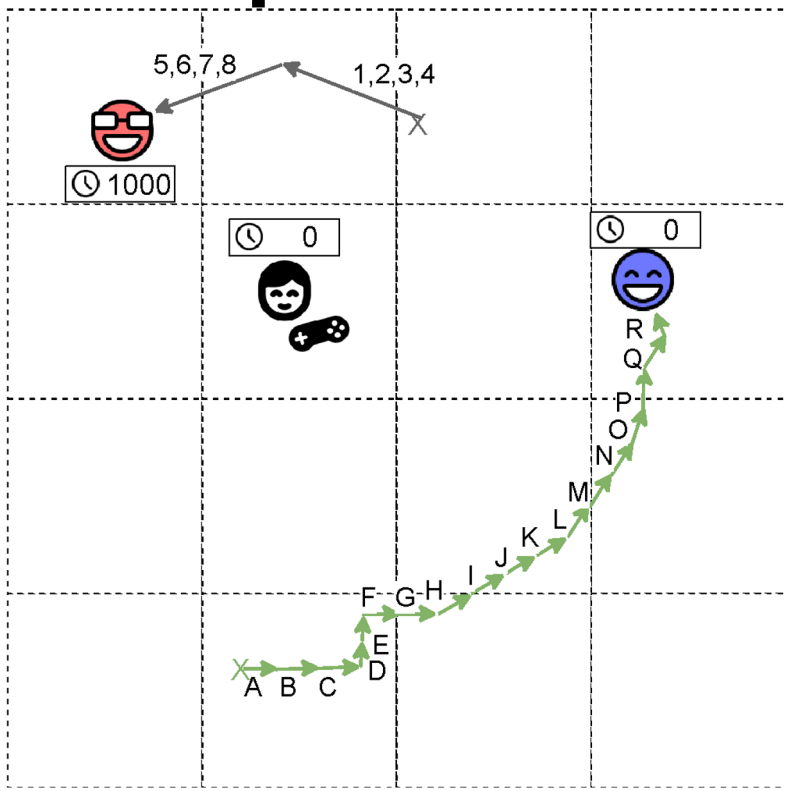

(b) Interest Set policy

Dyconit Policies Improve Scalability

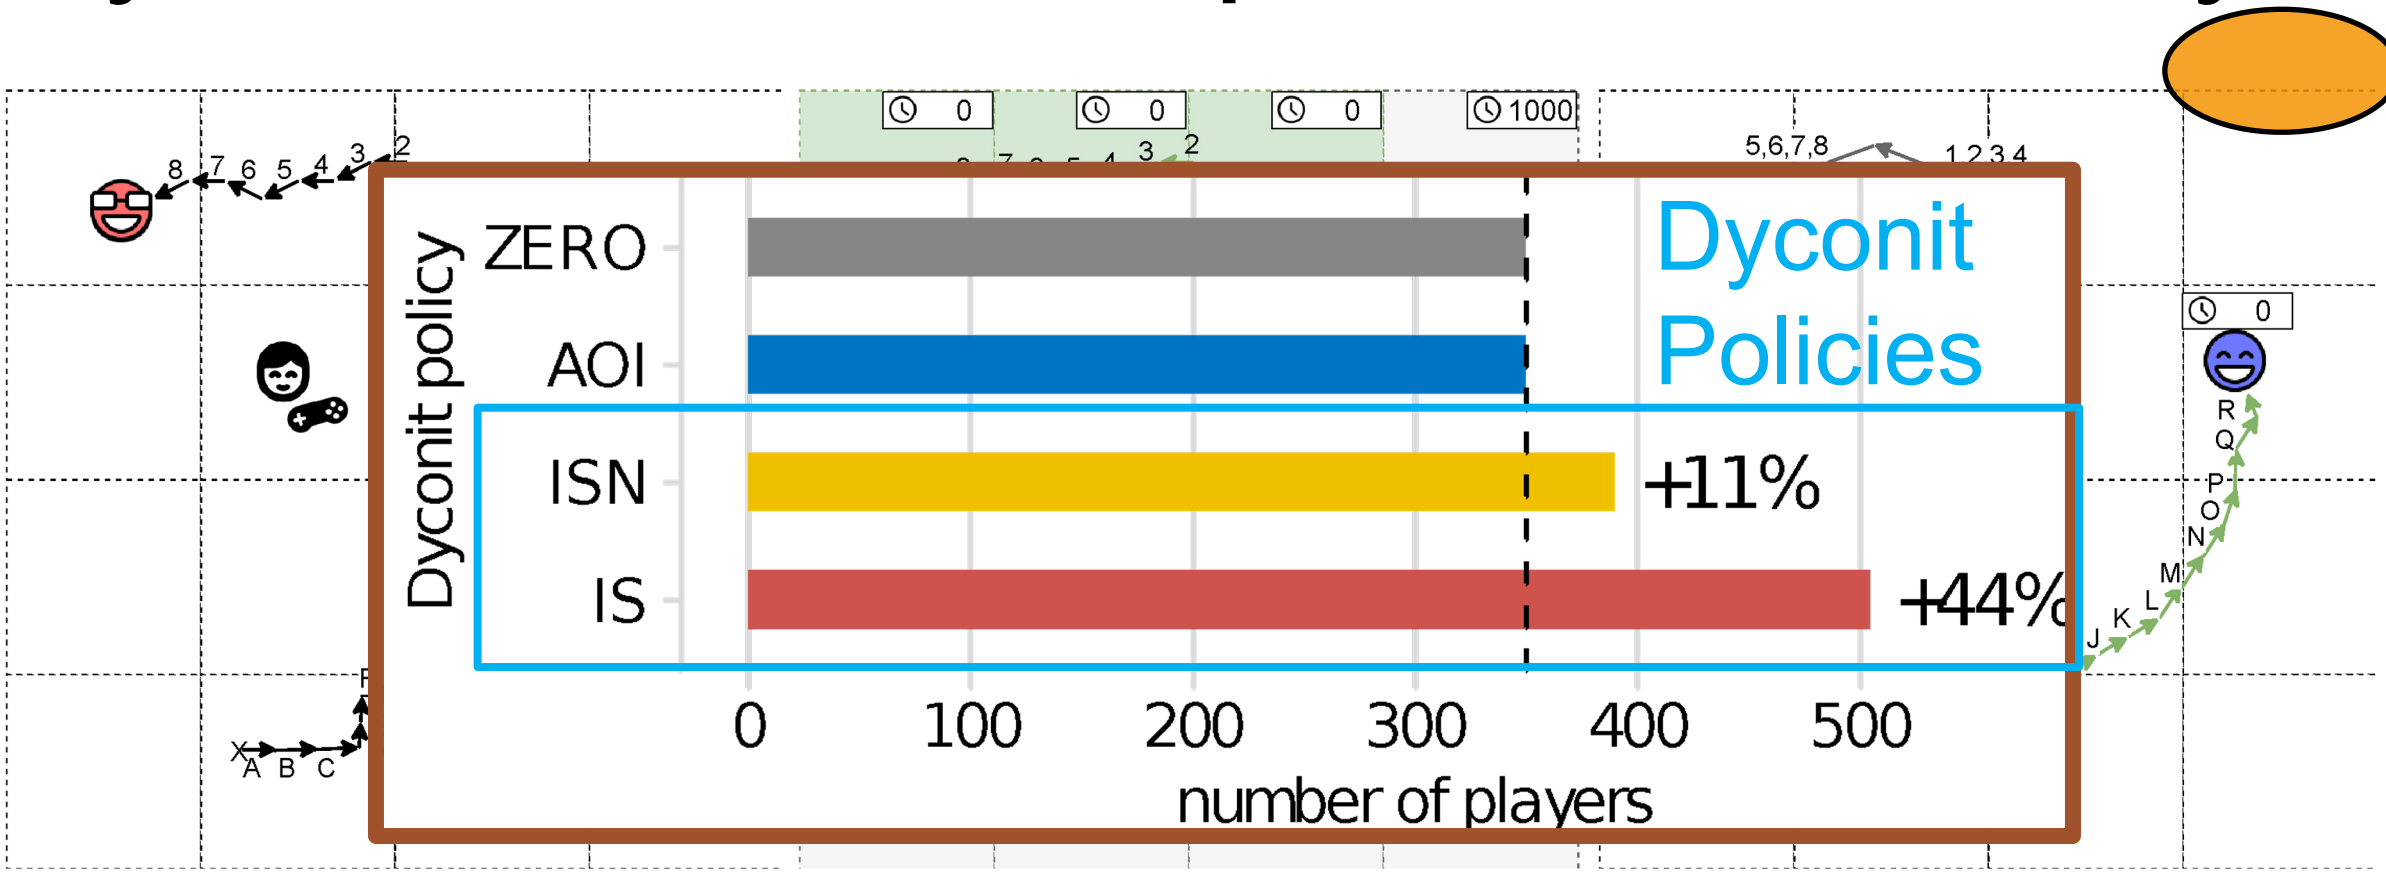

(a) No Dyconits (b) Area Of Interest policy (b) Interest Set policy

HOW TO AUTOMATE OPS ACROSS THE ECOSYSTEM 5.1B

IT'S OPERATIONS!

REFERENCE VIEW ON OPERATIONAL TECHNIQUES

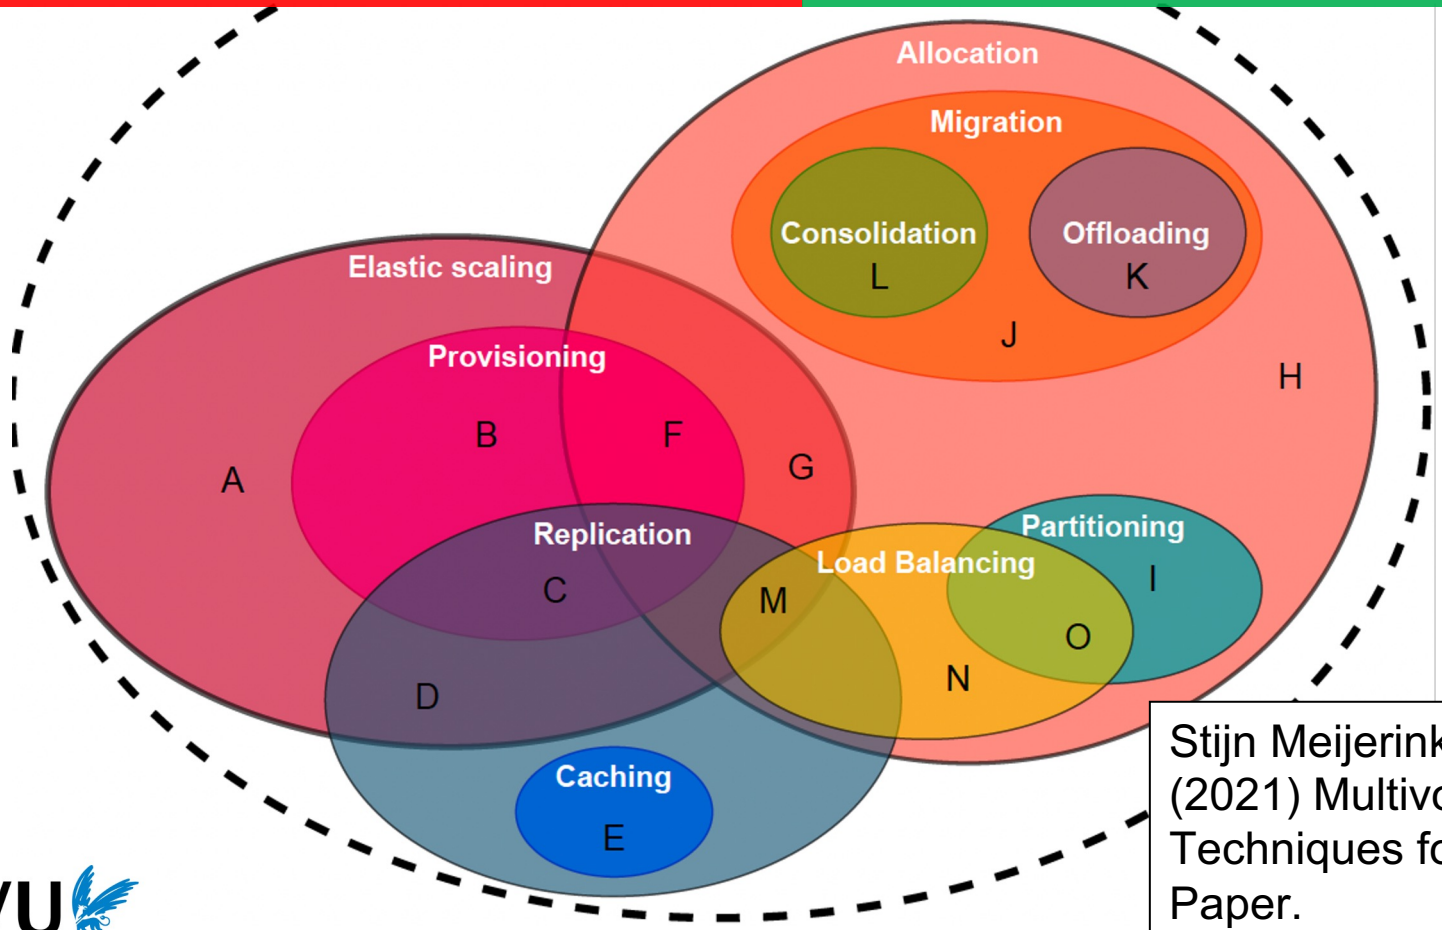

Provisioning (B)

Offloading (K)

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

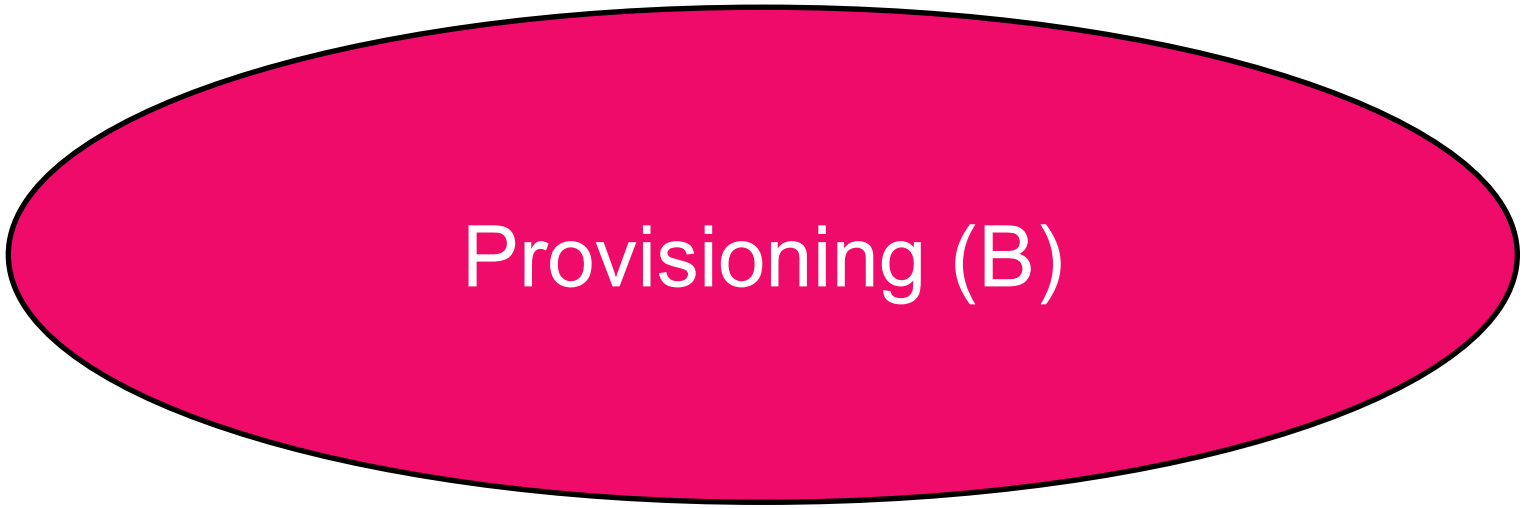
Provisioning (B)

World of Warcraft, a Traditional HPC User?!

(since 2003)

Ethics

- 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

Inter-operable

Scalability

Autonomous

Reusable

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

Resource Provisioning and Allocation
Static vs. Dynamic Provisioning

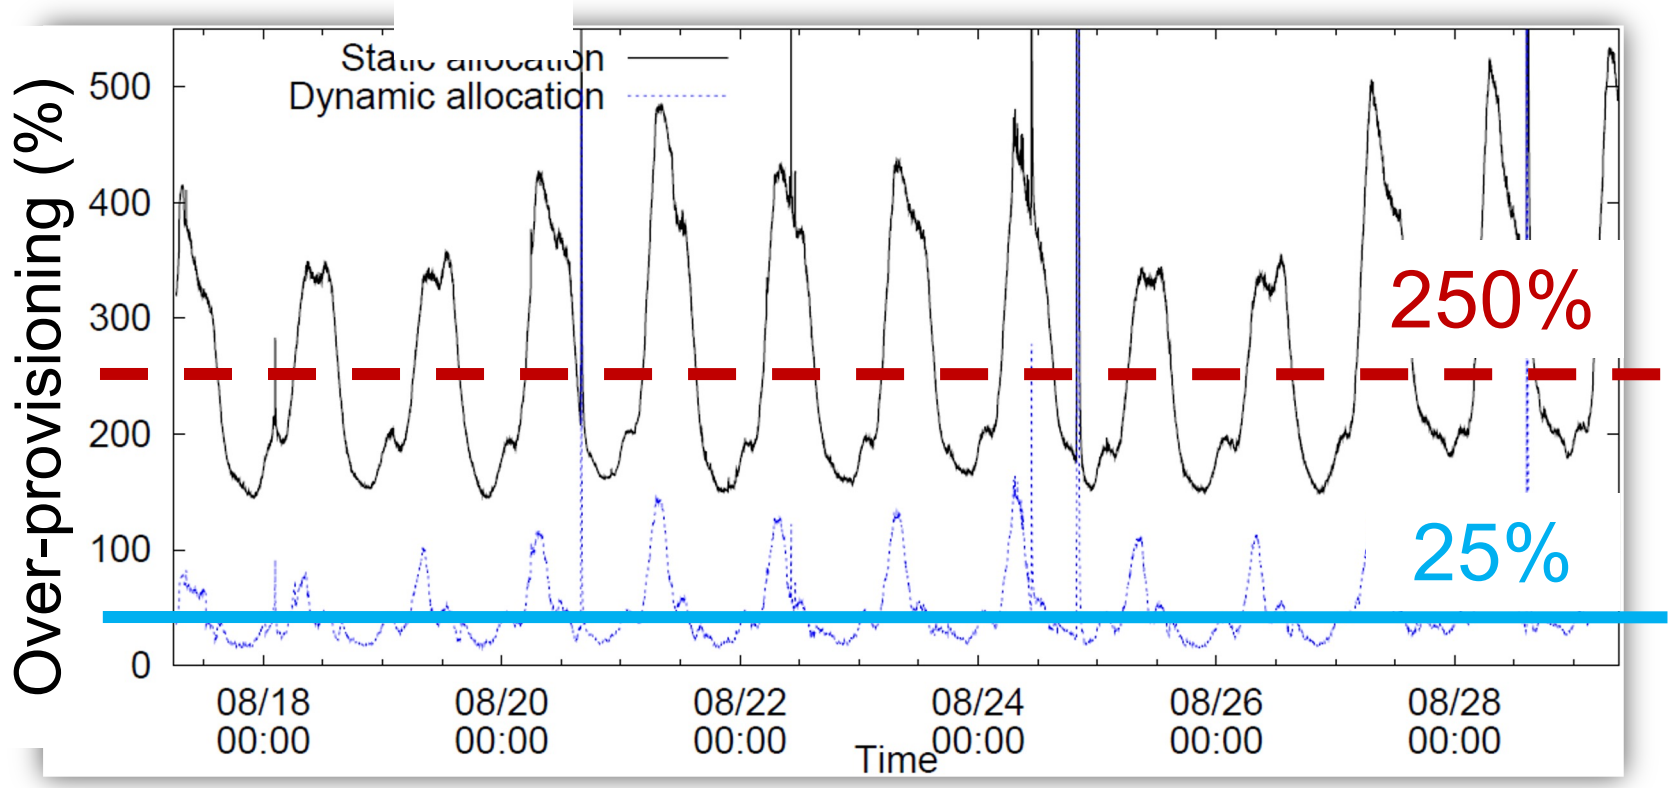

[Source: Nae, Iosup, and Prodan, ACM SC 2008]

Message: 5.2 Analytics ~ Only beginning to scratch the surface

2 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

2.3 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 limited topics
  - Analyzed vocabulary using toxicity-detection rules

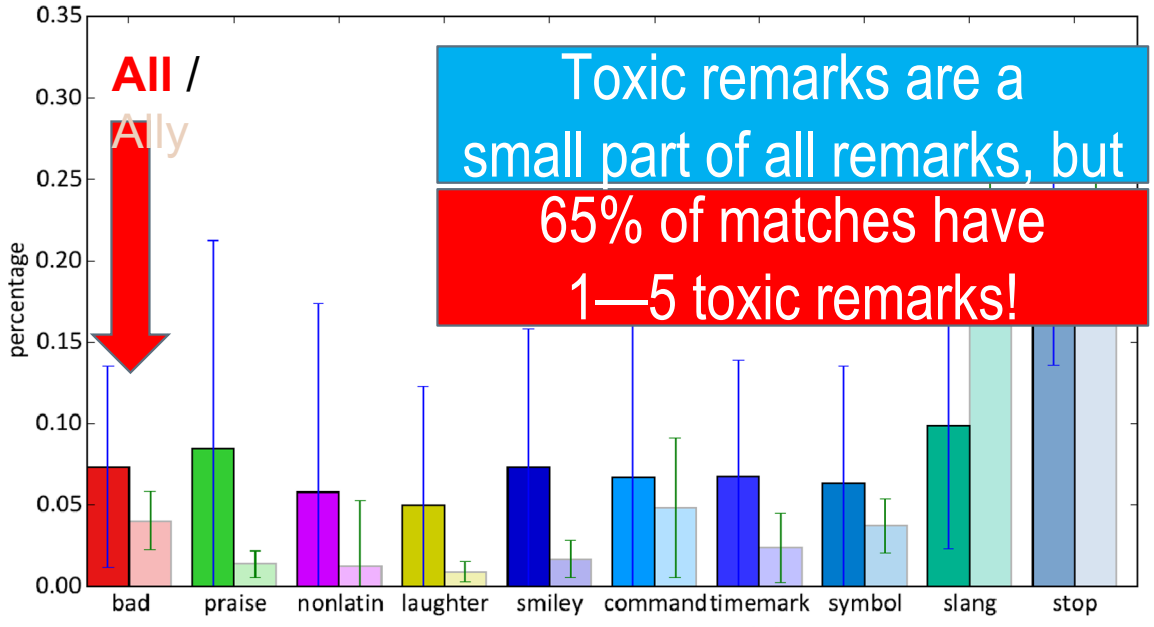

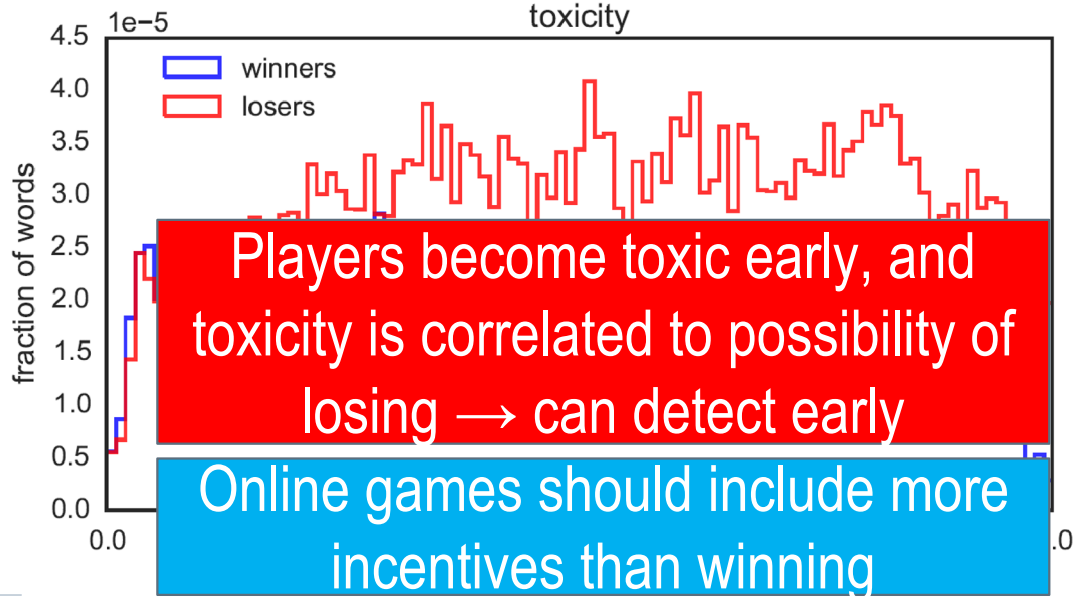

Maertens, Shen, Iosup, Kuipers. Toxicity detection in multiplayer online games. NETGAMES 2015: 1-6 (Best Paper Award)

Message: 5.3

Content ~ Computational power → emerging creativity, freshness

3 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-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: newsmap research team

Done

Open challenge

(Procedural) Game Content (Generation)

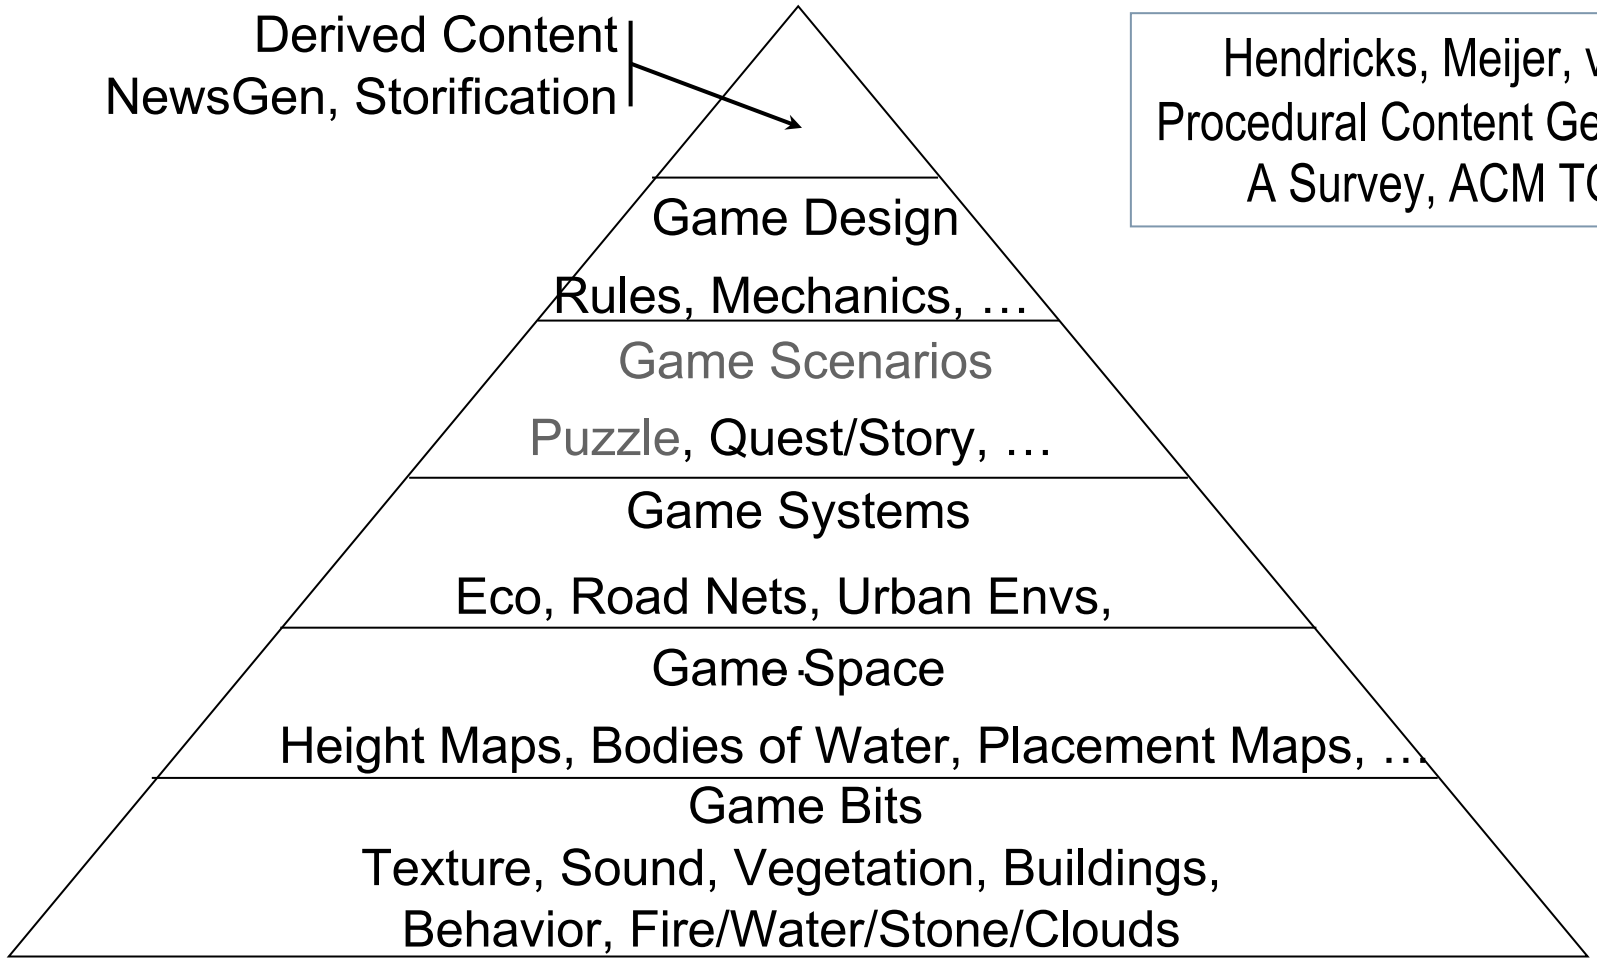

Hendricks, Meijer, vd Velden, Iosup, Procedural Content Generation for Games: A Survey, ACM TOMCCAP, 2012

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

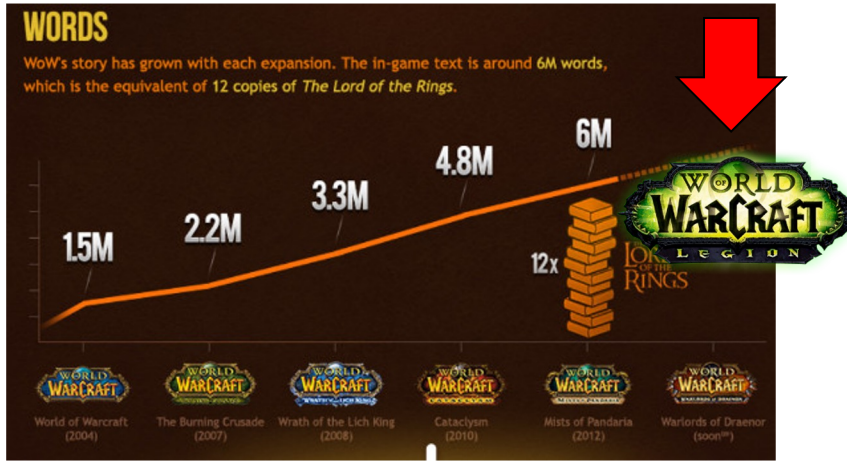

- **Frequent updates, fast-pace DLC**
  - Budget: $, upfront + $ / update
  - Breaks industry models: often based on technology, etc.
  - Players are happy

CITIES SKYLINES

RUNESCAPE

Source: Sergey Galyonkin, SteamSpy: Some things you should know about Steam, Jun 2015.

The POGGI Content Generation Framework

Scalability

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

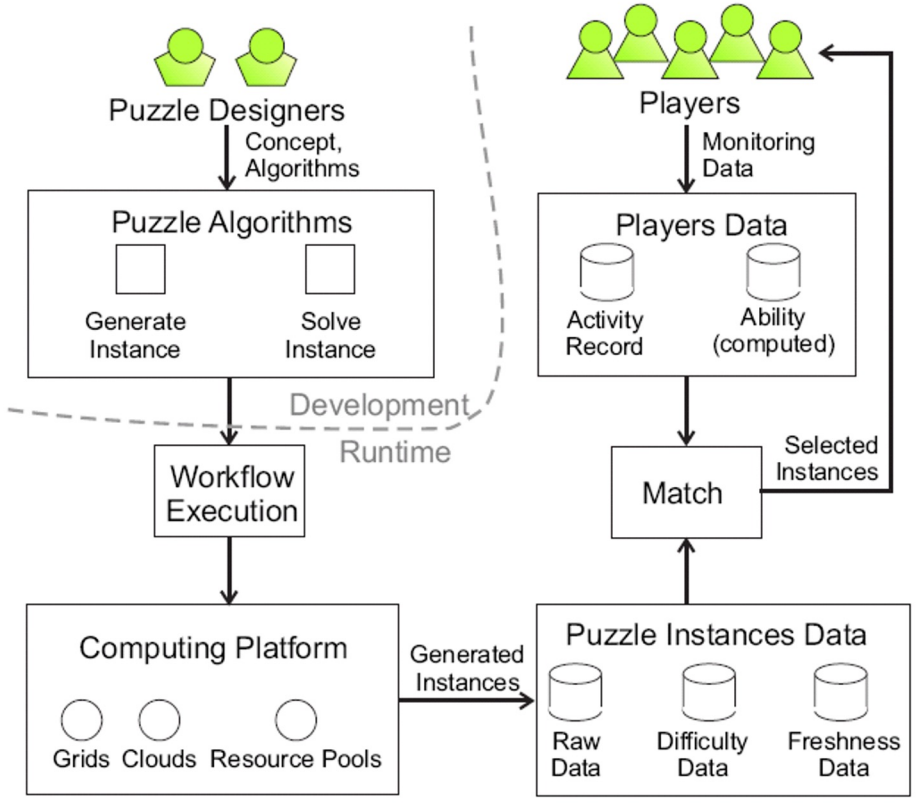

Smart system to recommend instances to players

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

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

Puzzle-Specific Considerations

Generating Player-Customized Content

Scalability (of content)

Puzzle difficulty

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

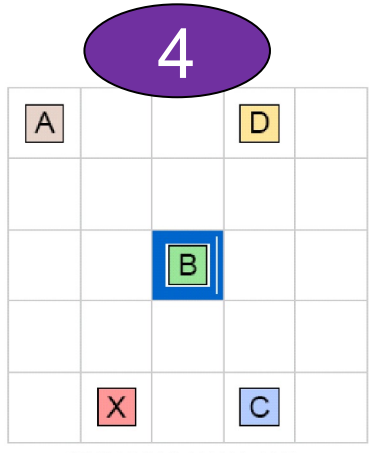

X:Right A:Right B:Up X:Up
(Best solution: 4 moves)

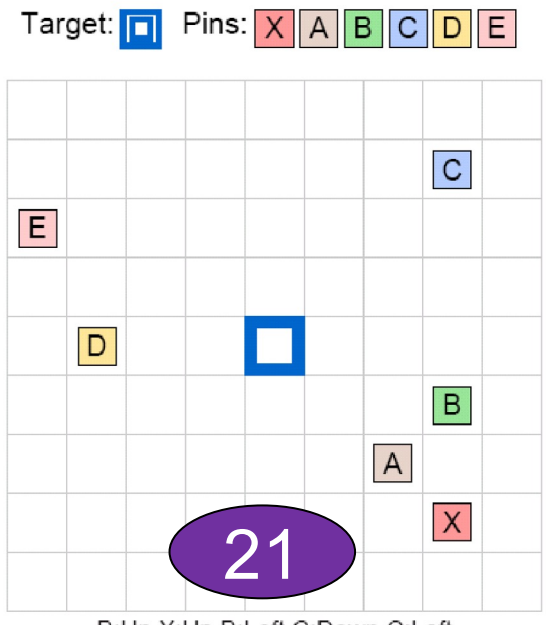

B:Up X:Up B:Left C:Down C:Left
B:Down B:Right B:Down E:Right E:Down
E:Right B:Up A:Up B:Left C:Down
C:Right E:Down X:Left E:Left X:Down
X:Left
(Best solution: 21 moves)

Player ability

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

Message: 5.4 Meta-gaming ~ Community techniques emerging

4 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

R. Garfield. Metagames. Horsemen of the Apocalypse: Essays on Roleplaying, 2000. Reproduced as Lost in the Shuffle: Games Within Games, http://www.wizards.com/Magic/magazine/Article.aspx?x=mtg/daily/feature/96

THE CAMEO FRAMEWORK

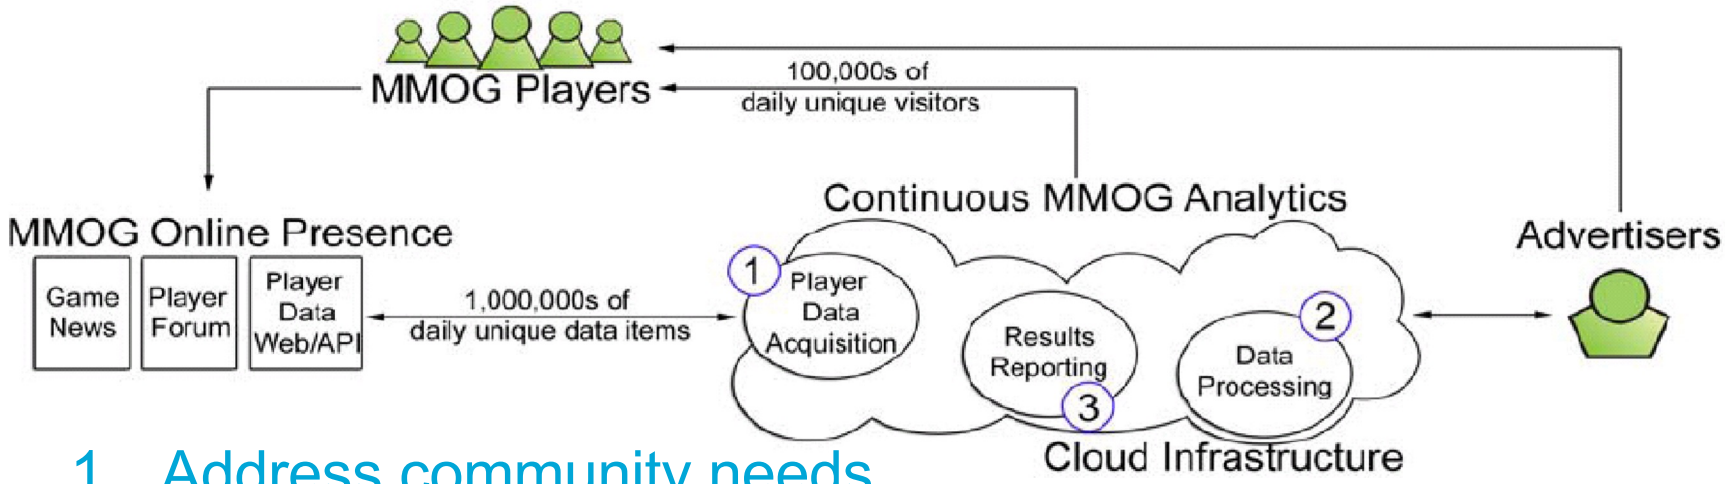

1. Address community needs
   - Can analyze skill level, experience points, rank
   - Can assess community size dynamically
2. Using on-demand technology: Cloud Computing
   - Dynamic cloud resource allocation, Elastic IP
3. Data management and storage: Cloud Computing
   - Crawl + Store data in the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

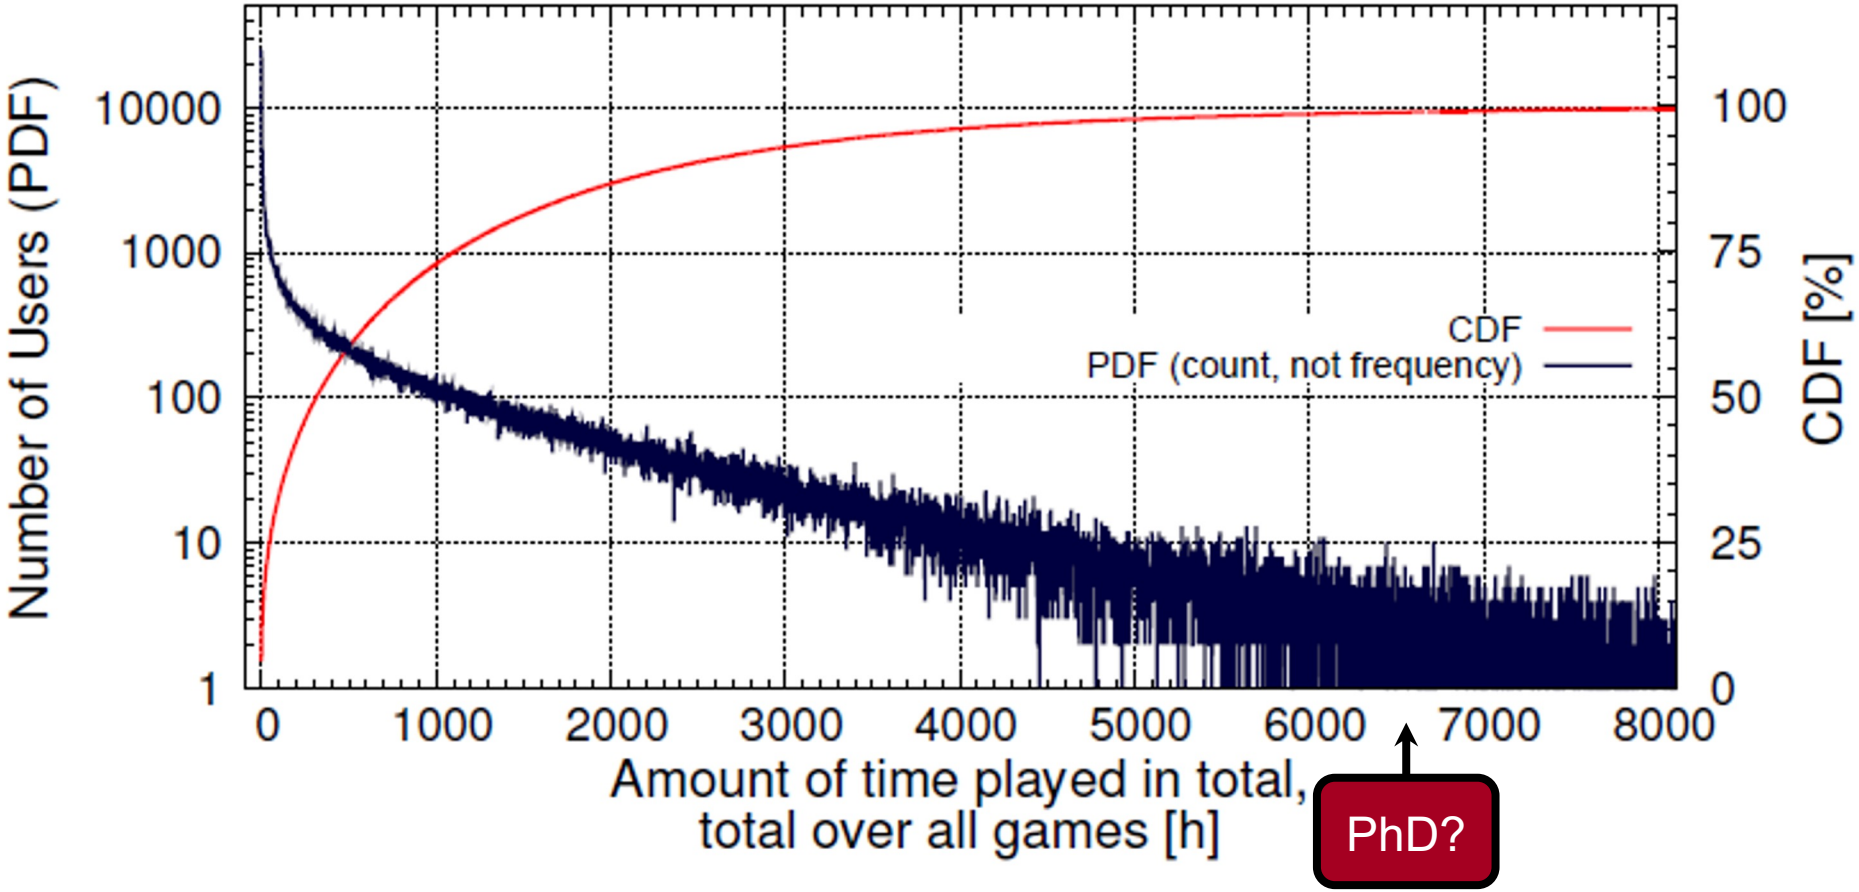

S. Shen, and A. Iosup, The XFire Online Meta-Gaming Network: Observation and High-Level Analysis, MMVE 2011

Message: 6

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

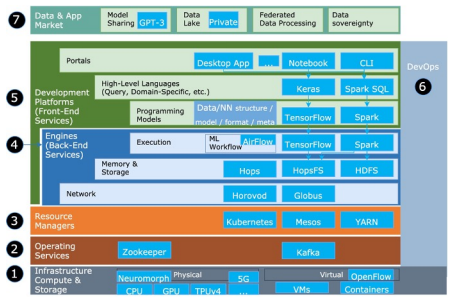

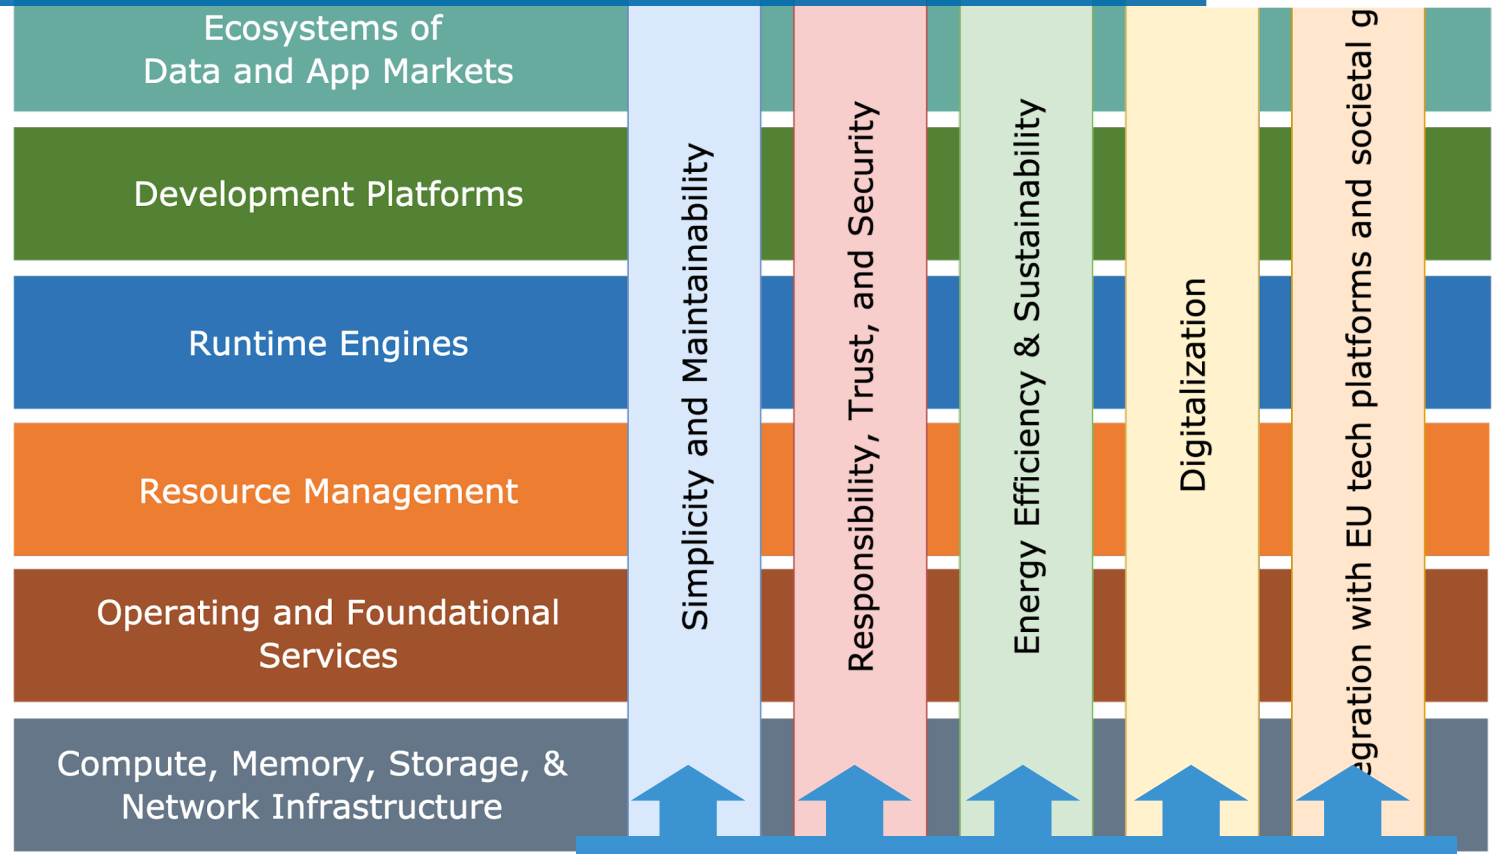

A.Iosup@vu.nl
http://atlarge.science

TAKE-HOME MESSAGE:

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

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

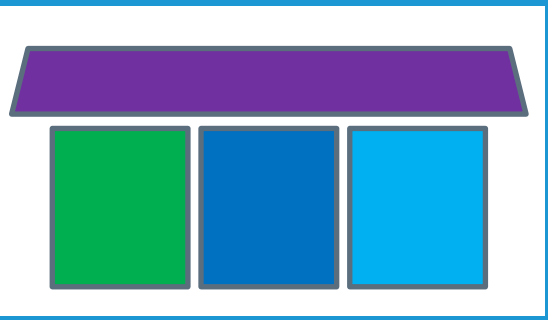

Online Gaming used to be networking, is now <u>all</u> computing

Online Gaming used to be game worlds, is now <u>all</u> kinds of apps

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!

WANT TO READ MORE ON THE TOPIC / JOIN OUR TEAM?

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. ISPDC 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

EXTRAS

A TYPICAL ECOSYSTEM: SERVICE, DATACENTER, SCHEDULER

Creators

Digital Services

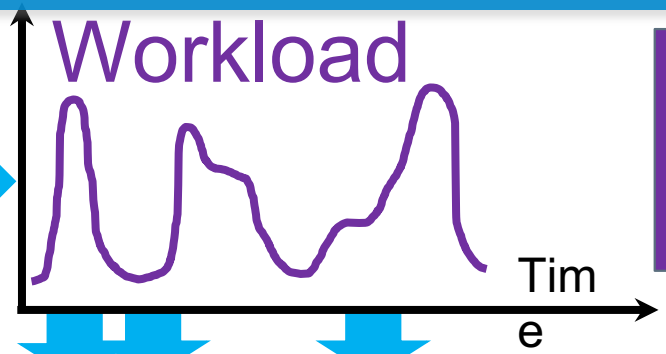

100s of services … an ecosystem

Resource Manager and Scheduler

(full-/micro-/nano-) Datacenter

Performance, Dependability, Sustainability

50+ PLATFORMS … EMERGENT FEATURES

3A

THE COMPLEXITY CHALLENGE

REFERENCE ARCHITECTURE OF FAAS PLATFORMS

Workflow Composition Layer

- W1 Workflow Registry
- W2 Workflow Engine
- W3 Workflow Scheduler
- W4 Workflow Execution Store

Function Management Layer

- F1 Function Registry
- F2 Function Builder
- F3 Function Deployer
- F4 Function Instance
- F5 Function Router
- F6 Function Autoscaler

Resource Orchestration Layer

- R1 Naming Service
- R2 Resource Manager
- R3 Resource Scheduler
- R4 Node Agent

Business Concerns

Operational Concerns

Erwin van Eyk

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

BEYOND THE DATACENTER: THE COMPUTING CONTINUUM

THE COMPLEXITY CHALLENGE

REFERENCE ARCHITECTURE OF FAAS PLATFORMS

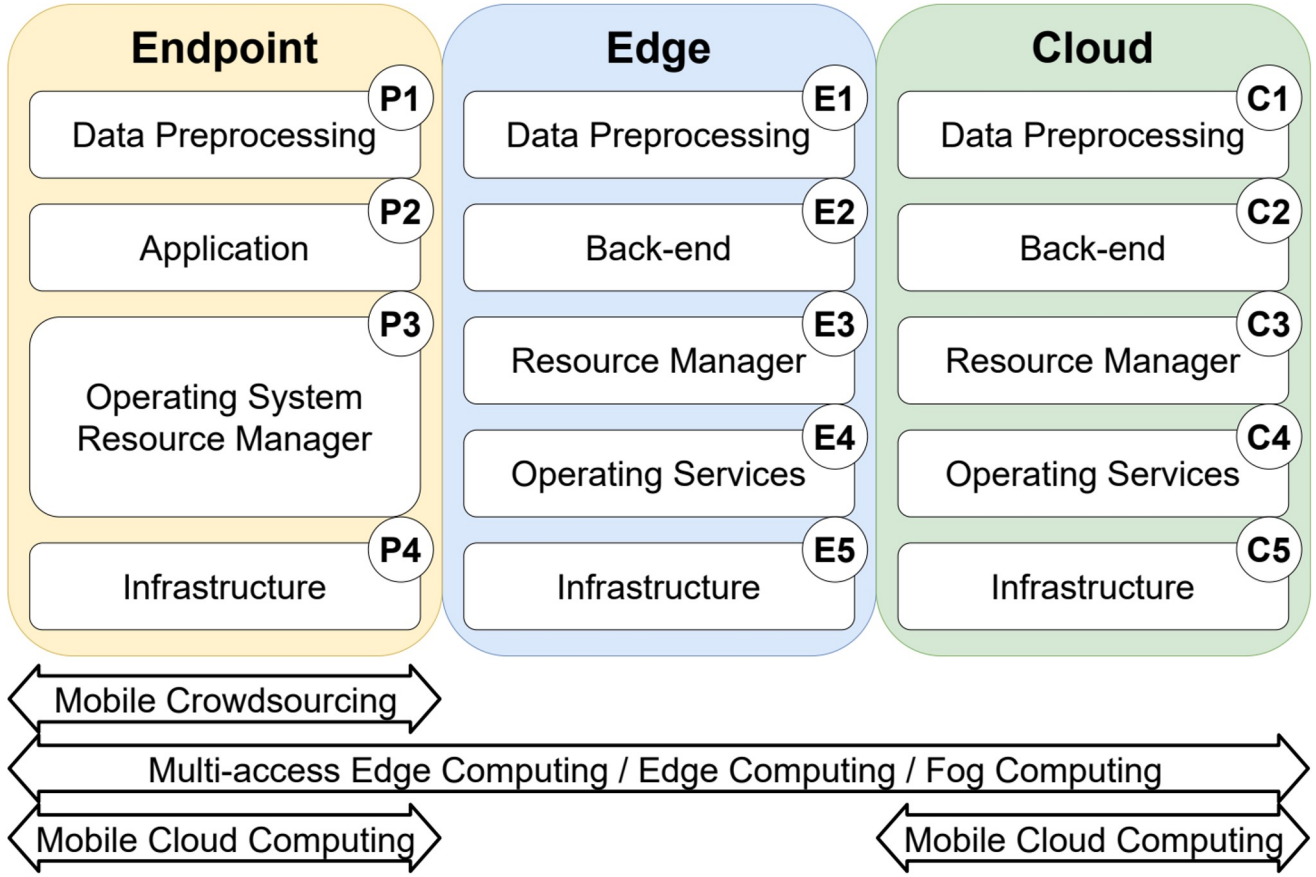

[Jansen et al. (2022) Ongoing Work

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

spec Research Serverless

4A

THE COMPLEXITY CHALLENGE

REFERENCE VIEW ON OPERATIONAL TECHNIQUES

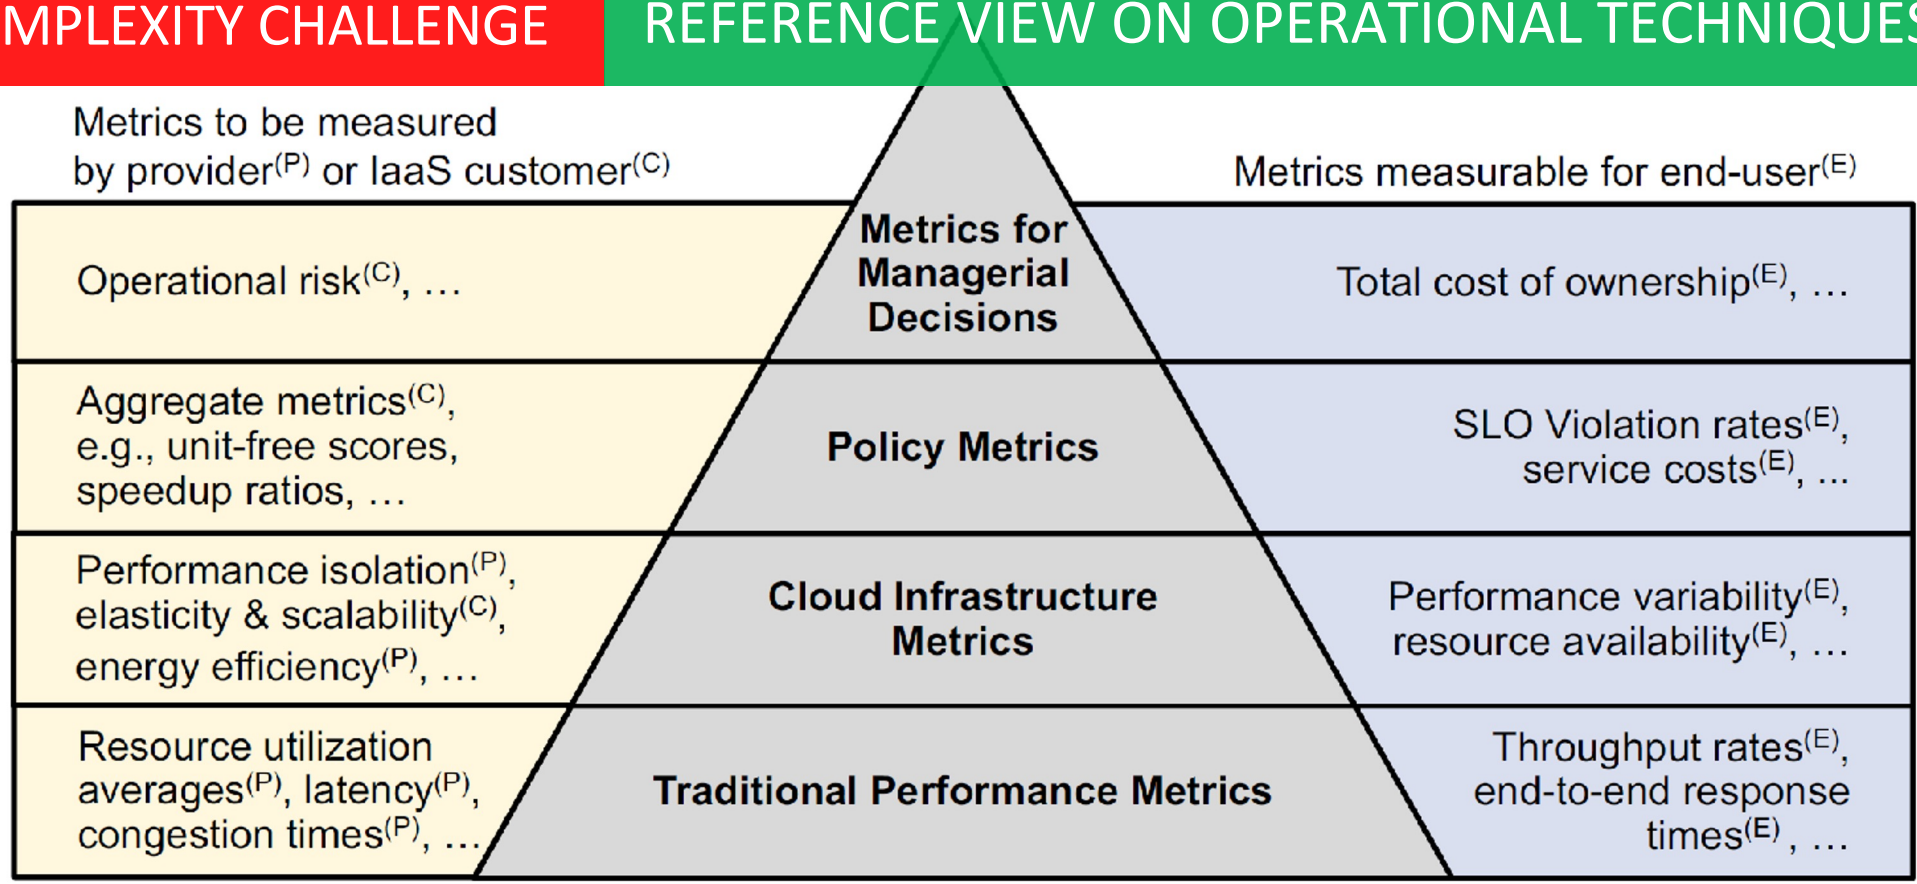

N. Herbst, E. Van Eyk, C. L. Abad, R. Iosup, et al. (2018) Quantifying Cloud Performance and Dependability: Taxonomy, Metric Design, and

VU

HOW TO AUTOMATE X ACROSS THE ECOSYSTEM?

4B

IT'S OPERATIONS!

REFERENCE VIEW ON OPERATIONAL TECHNIQUES

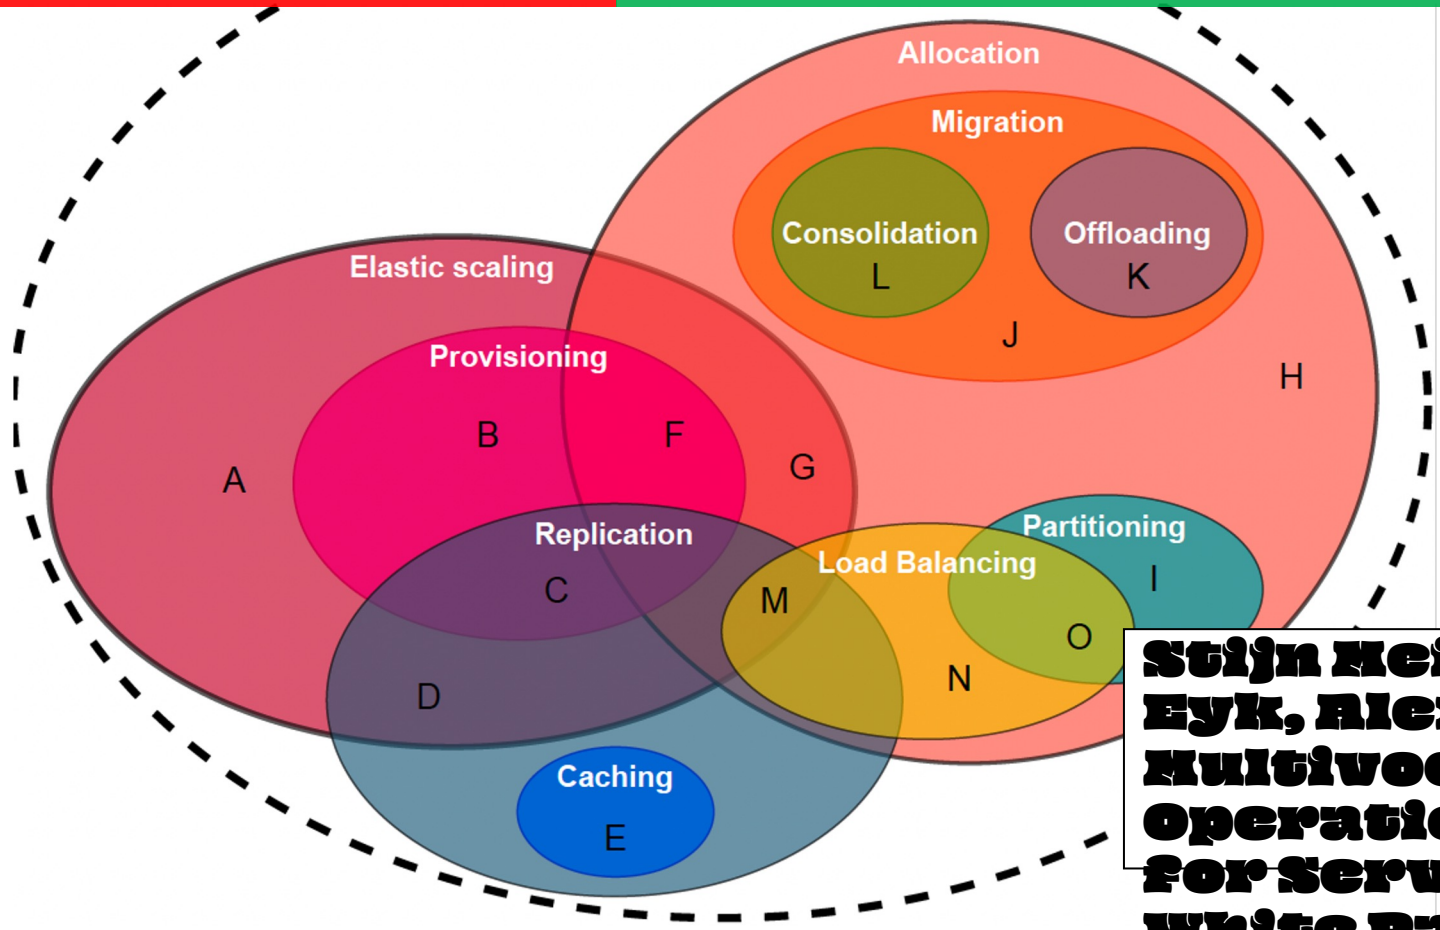

Stijn Heijerink, Erwin van Eyk, Alexandru Iosup (2021) Multivocal Survey of Operational Techniques for Serverless Computing. White Paper.

THE ECONOMIC IMPACT OF MASSIVE COMPUTER ECOSYSTEMS

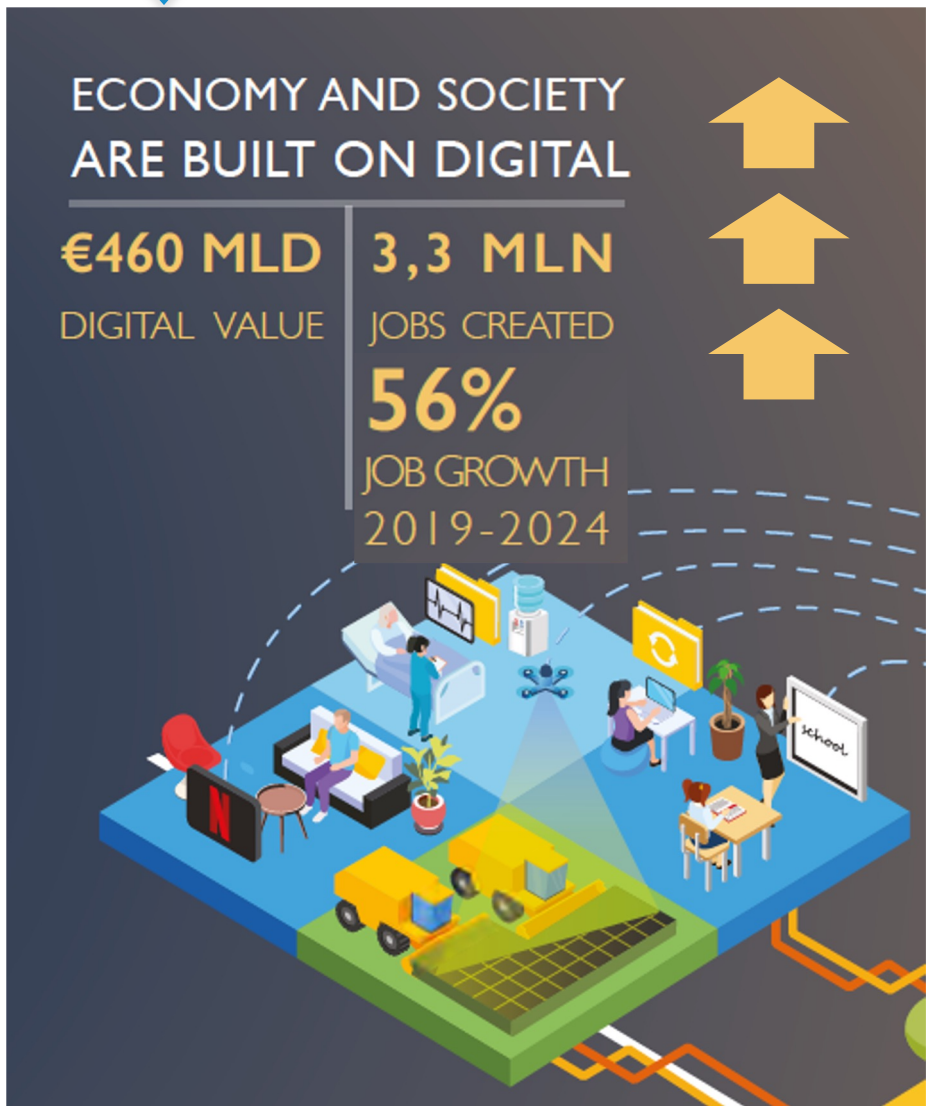

DIVERSE SERVICES FOR ALL

EVERY €1 → €15 ADDED VALUE

Impacting >60% of the NL GDP (1 trillion EUR/y)

Attracting >20% of all foreign direct investments in NL

Sources: Iosup et al., Massivizing Computer Systems, ICDCS 2018 [Online] / Dutch Data Center Association, 2020 [Online] / Growth: NL Gov't, Flexera, Binx 2020. Gartner 2019. IA 2017.

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

Iosup et al., Lecture Notes in Distributed Computer Systems, ICDCS