MASSIVIZING COMPUTER SYSTEMS – THE ICT INFRASTRUCTURE MEMEX

VU ON OPERATIONAL DATA ANALYTICS IN THE 21ST CENTURY

Sponsored by: GRAPH MASSIVIZER, NWO

@Large Research
Massivizing Computer Systems

http://atlarge.science

bit.ly/AIScalPerf23

ODA = Process for data-driven analysis of ICT infrastructure operations → insight, optimization

@AIosup

VU VRIJE UNIVERSITEIT AMSTERDAM

Prof.dr.ir. Alexandru IOSUP

Contributions from the MCS/AtLarge teams. Many thanks!
Many thanks to our collaborators, international working groups, authors of all images included here. Also thanks to Gianfranco Bilardi and the ScalPerf'23 organizing team.

US IN 1 MINUTE

WE’RE MASSIVIZING COMPUTER SYSTEMS!

VU AMSTERDAM < SCHIPHOL < THE NETHERLANDS < EUROPE

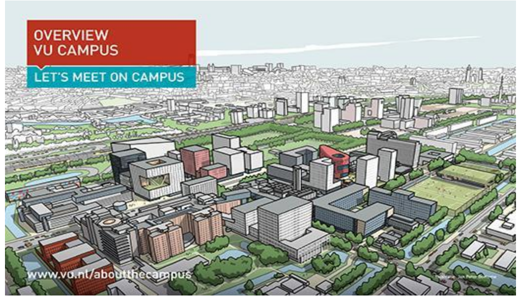

Amsterdam

founded 10th century

pop: 850,000

VU

founded 1880

pop: 24,500

(+ 4,600 staff)

The Netherlands

pop: 16.5 M

Europe

VU UNIVERSITY AMSTERDAM

http://atlarge.science

URRENT TEAM

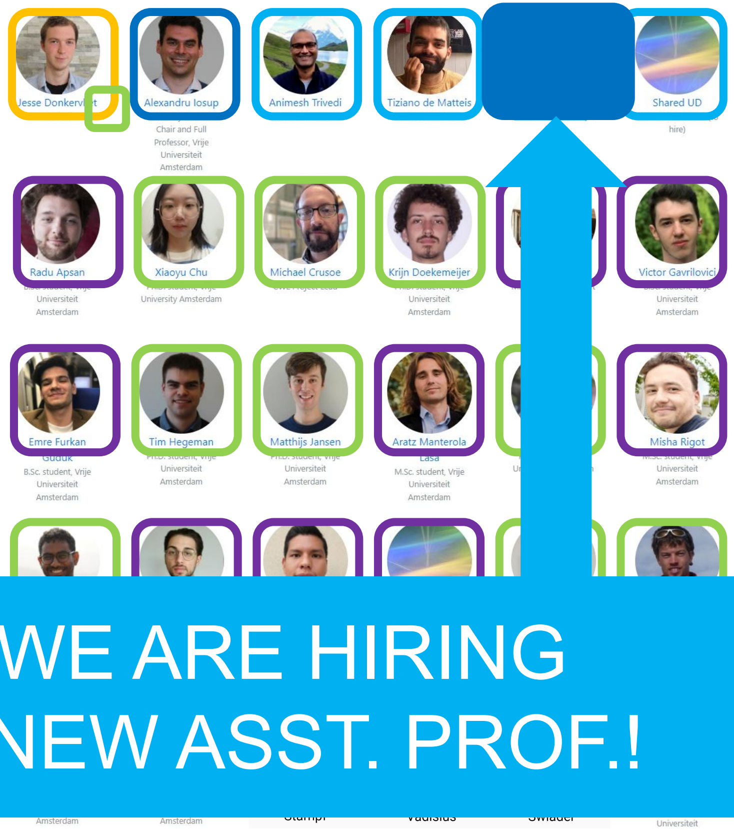

Professor

Assistant Prof.

Teacher

Visitor/P.-doc

Ph.D. student

Early Scientist MSc, BSc HP

ALUMNI

VISITORS

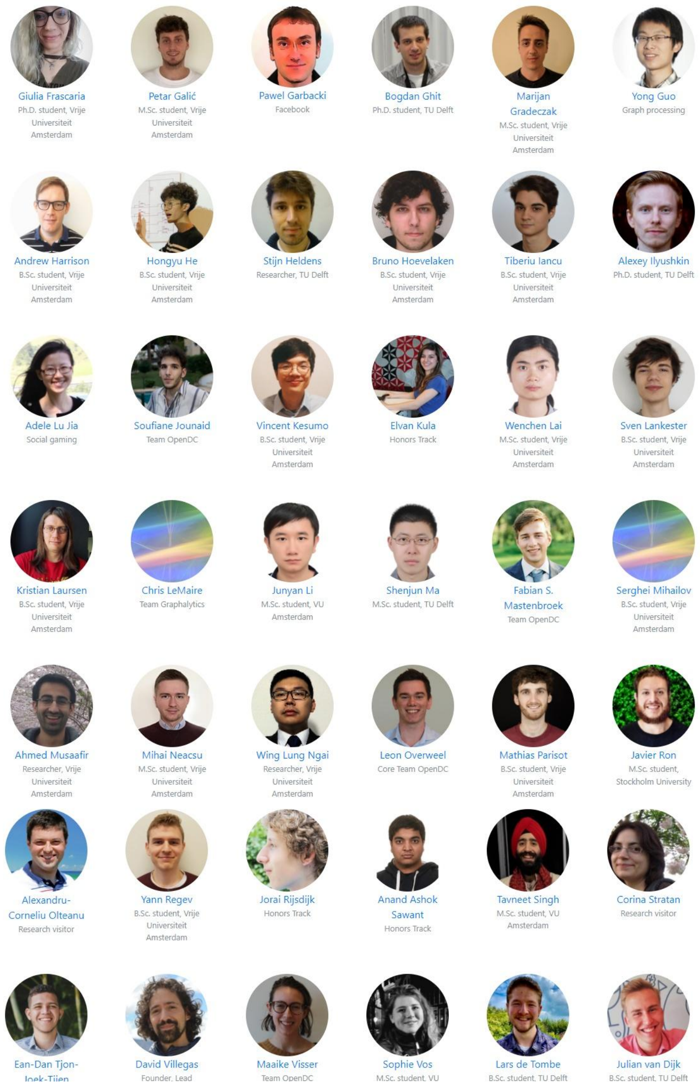

WE ARE HIRING
A NEW ASST. PROF.!

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

DIVERSITY: CHALLENGES AND OPPORTUNITIES

Team and collaborators last year (biased selection)

Prof.dr. Sanne Abeln

Prof.dr. Carole Goble

dr. Chrysa Papagianni

dr. Paola Grosso

Laura Stampf

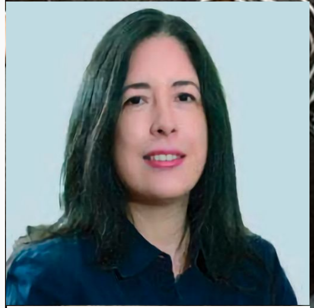

Prof.dr. Cristina Abad

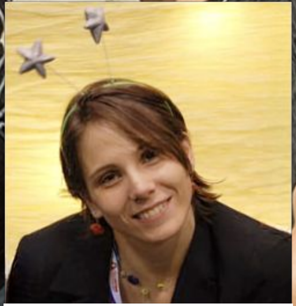

Prof.dr.ir. Ana Lucia Varbanescu

Renske de Wit, MSc

WHO AM I?
PROF. DR. IR. ALEXANDRU IOSUP

- Education, my courses:
  - Honours Programme, Computer Org. (BSc)
  - Distributed Systems, Cloud Computing (MSc)
- Research, 15 years in DistribSys:
  - Massivizing Computer Systems
  - About 30 young researchers in the team
- 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



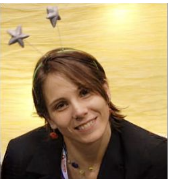


SINCE LAST YEAR – RE-BOOTED THE COMPSYS NL COMMUNITY…

ICT-RESEARCH PLATFORM NETHERLANDS

SIG FCSN + Manifesto on Computer Systems and Networking Research

Clear vision for the field in the NL, 2021-2035

Signed

50 PIs / Leads

07 universities

05 relevant societal stakeholders

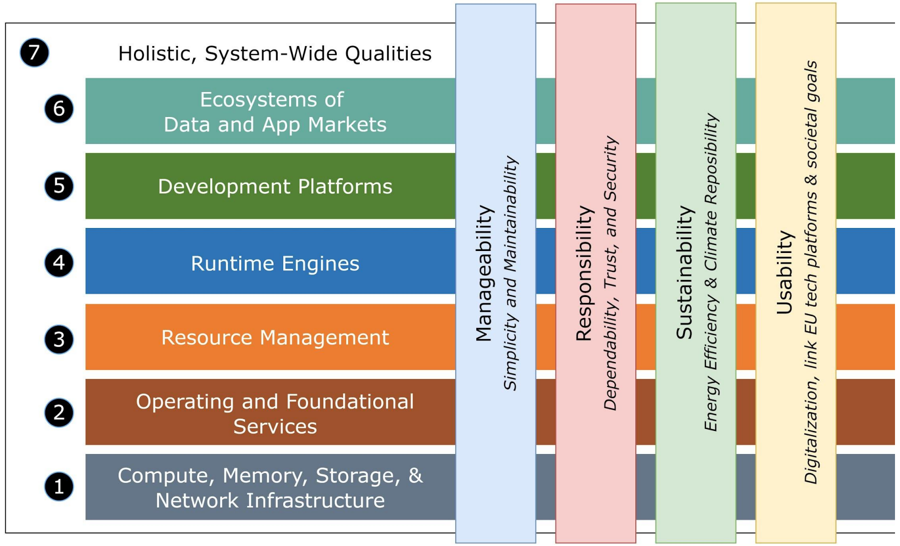

Available

Full version https://arxiv.org/pdf/2206.03259

Who's Who in CompSysNL? https://bit.ly/CompSysNLWhosWho

VU VRIJE UNIVERSITEIT AMSTERDAM

6G FUTURE NETWORK SERVICES

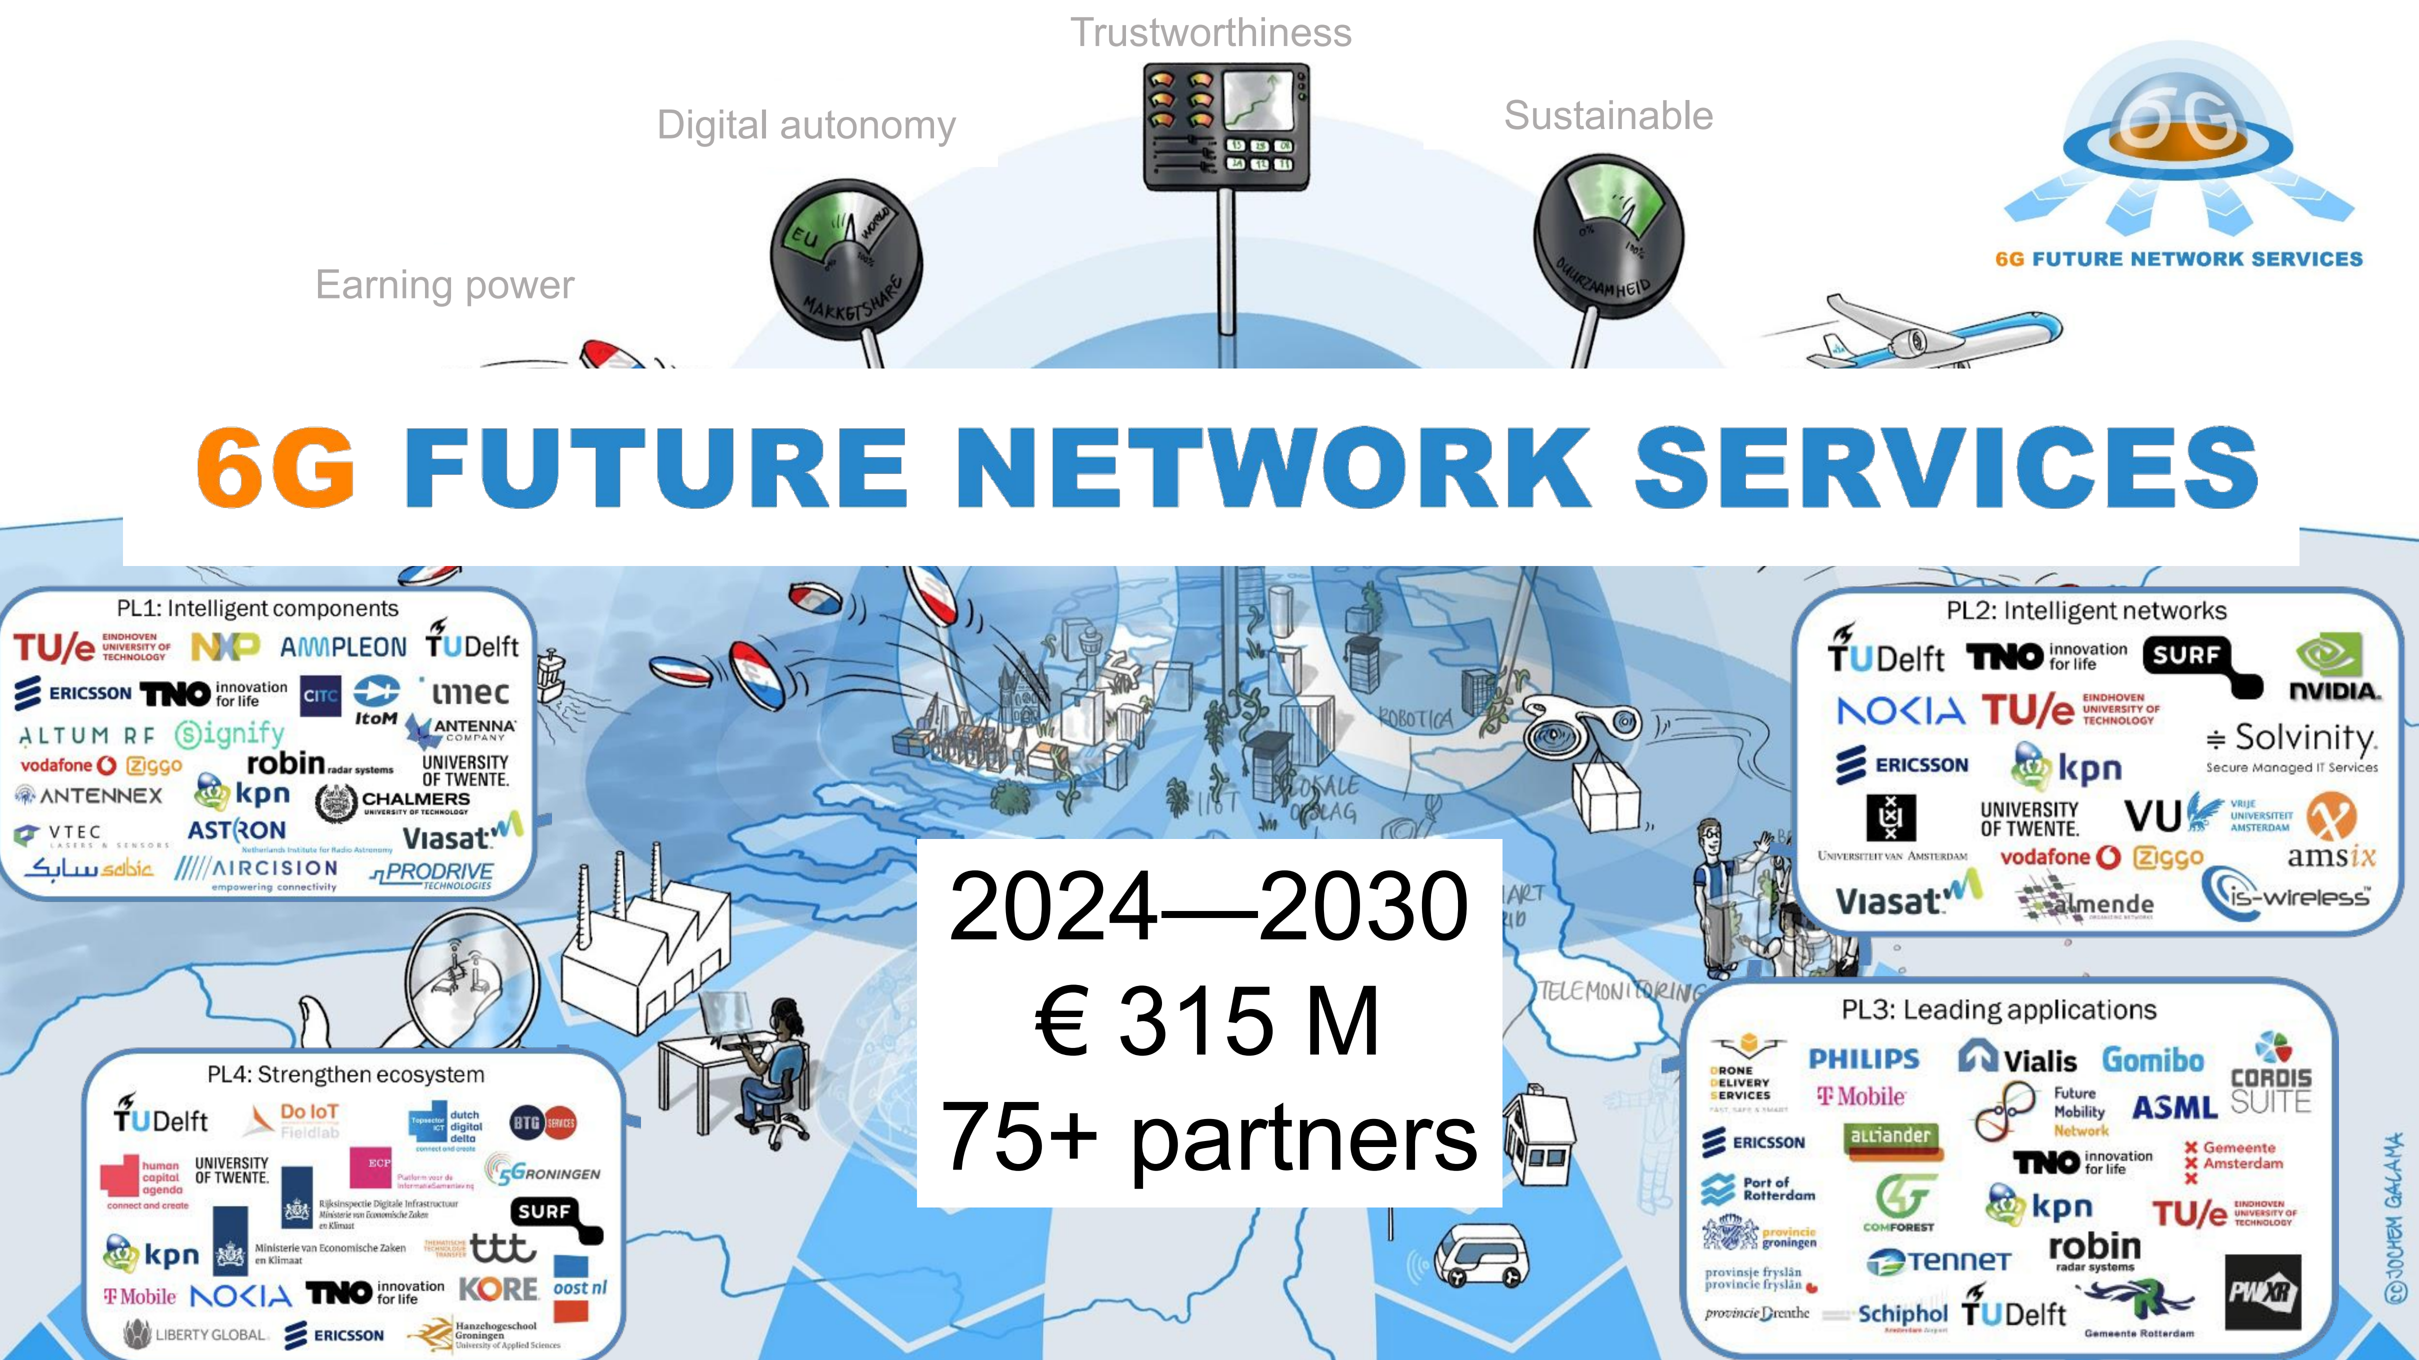

2024—2030
€ 315 M
75+ partners

ONE PROJECT TO MENTION…

GRAPH MASSIVIZER

Big Graph Processing: Used in AI/ML, FinTech, ICT Infra., Industry 4.0, Energy Mgmt.*, etc.

Vision: Massivizing computer systems approaches are key to enable big graph ecosystems

CACM Cover/Featured article, Sep 2021

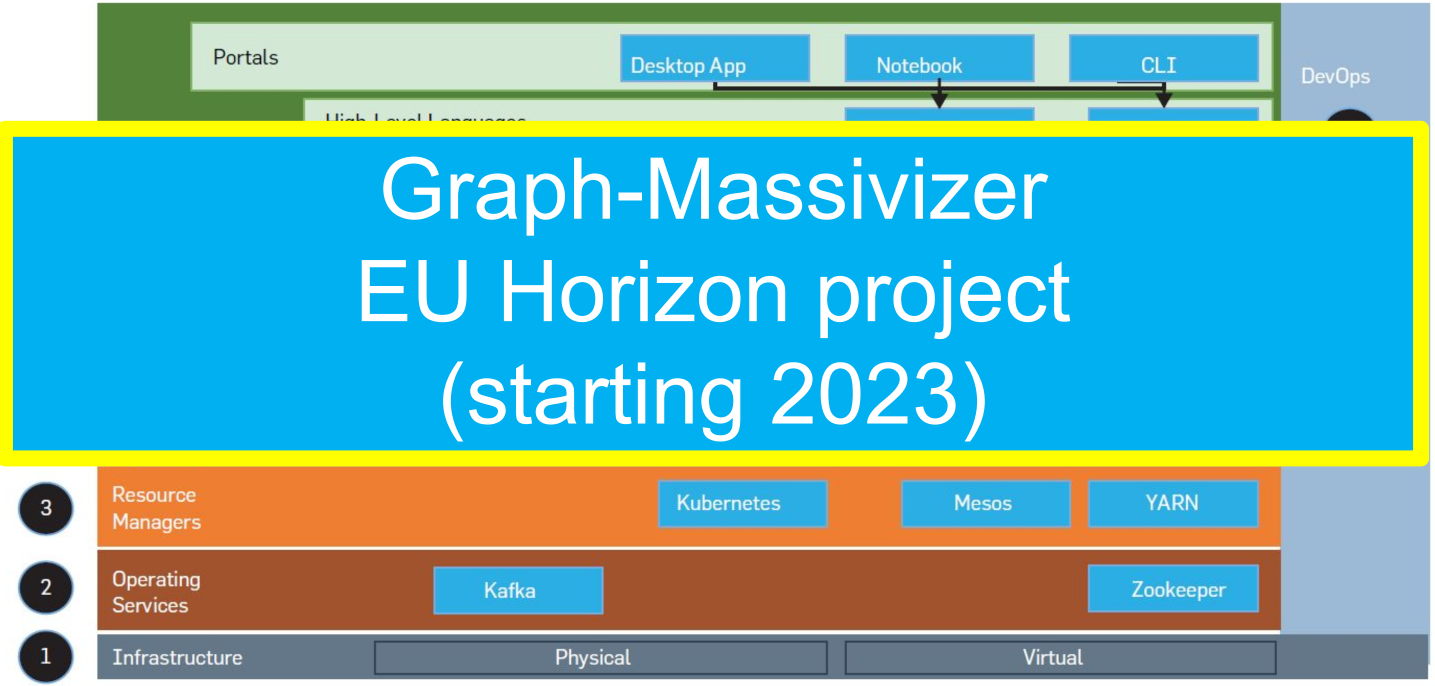

(*) Digital twin for datacenters, with partners CINECA, UniBo, etc.

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

Radu Prodan, Dragi Kimovski, Andrea Bartolini, Michael Cochez, Alexandru Iosup, Evgeny Kharlamov, Jože Rožanec, Laurenţiu Vasiliu, Ana Lucia Vărbănescu (2022) Towards Extreme and Sustainable Graph Processing for Urgent Societal Challenges in Europe. IEEE Cloud Summit.

ONE CONFERENCE TO MENTION…

https://2024.ccgrid-conference.org/

CCGRID 2024

Philadelphia, USA
May 5-9, 2024

The 24th IEEE/ACM international Symposium on Cluster, Cloud and Internet Computing

Program Chairs

Alexandru Iosup, Vrije Universiteit Amsterdam, Netherlands
Xubin He, Temple University, USA
Beth Plale, Indiana University, USA

1. Hardware Systems and Networking
2. Software Systems and Platforms
3. Machine Learning (ML) for Systems and Systems for ML
4. Future Compute Continuum and Seamless Ecosystems
5. Applications and Workflows
6. Performance Monitoring, Modeling, Analysis, and Benchmarking
7. Distributed and Parallel Storage Systems
8. Education about Cluster, Cloud and Internet Computing

THIS IS THE GOLDEN AGE OF COMPUTER ECOSYSTEMS

1

GENERALITY OF MASSIVE COMPUTER ECOSYSTEMS

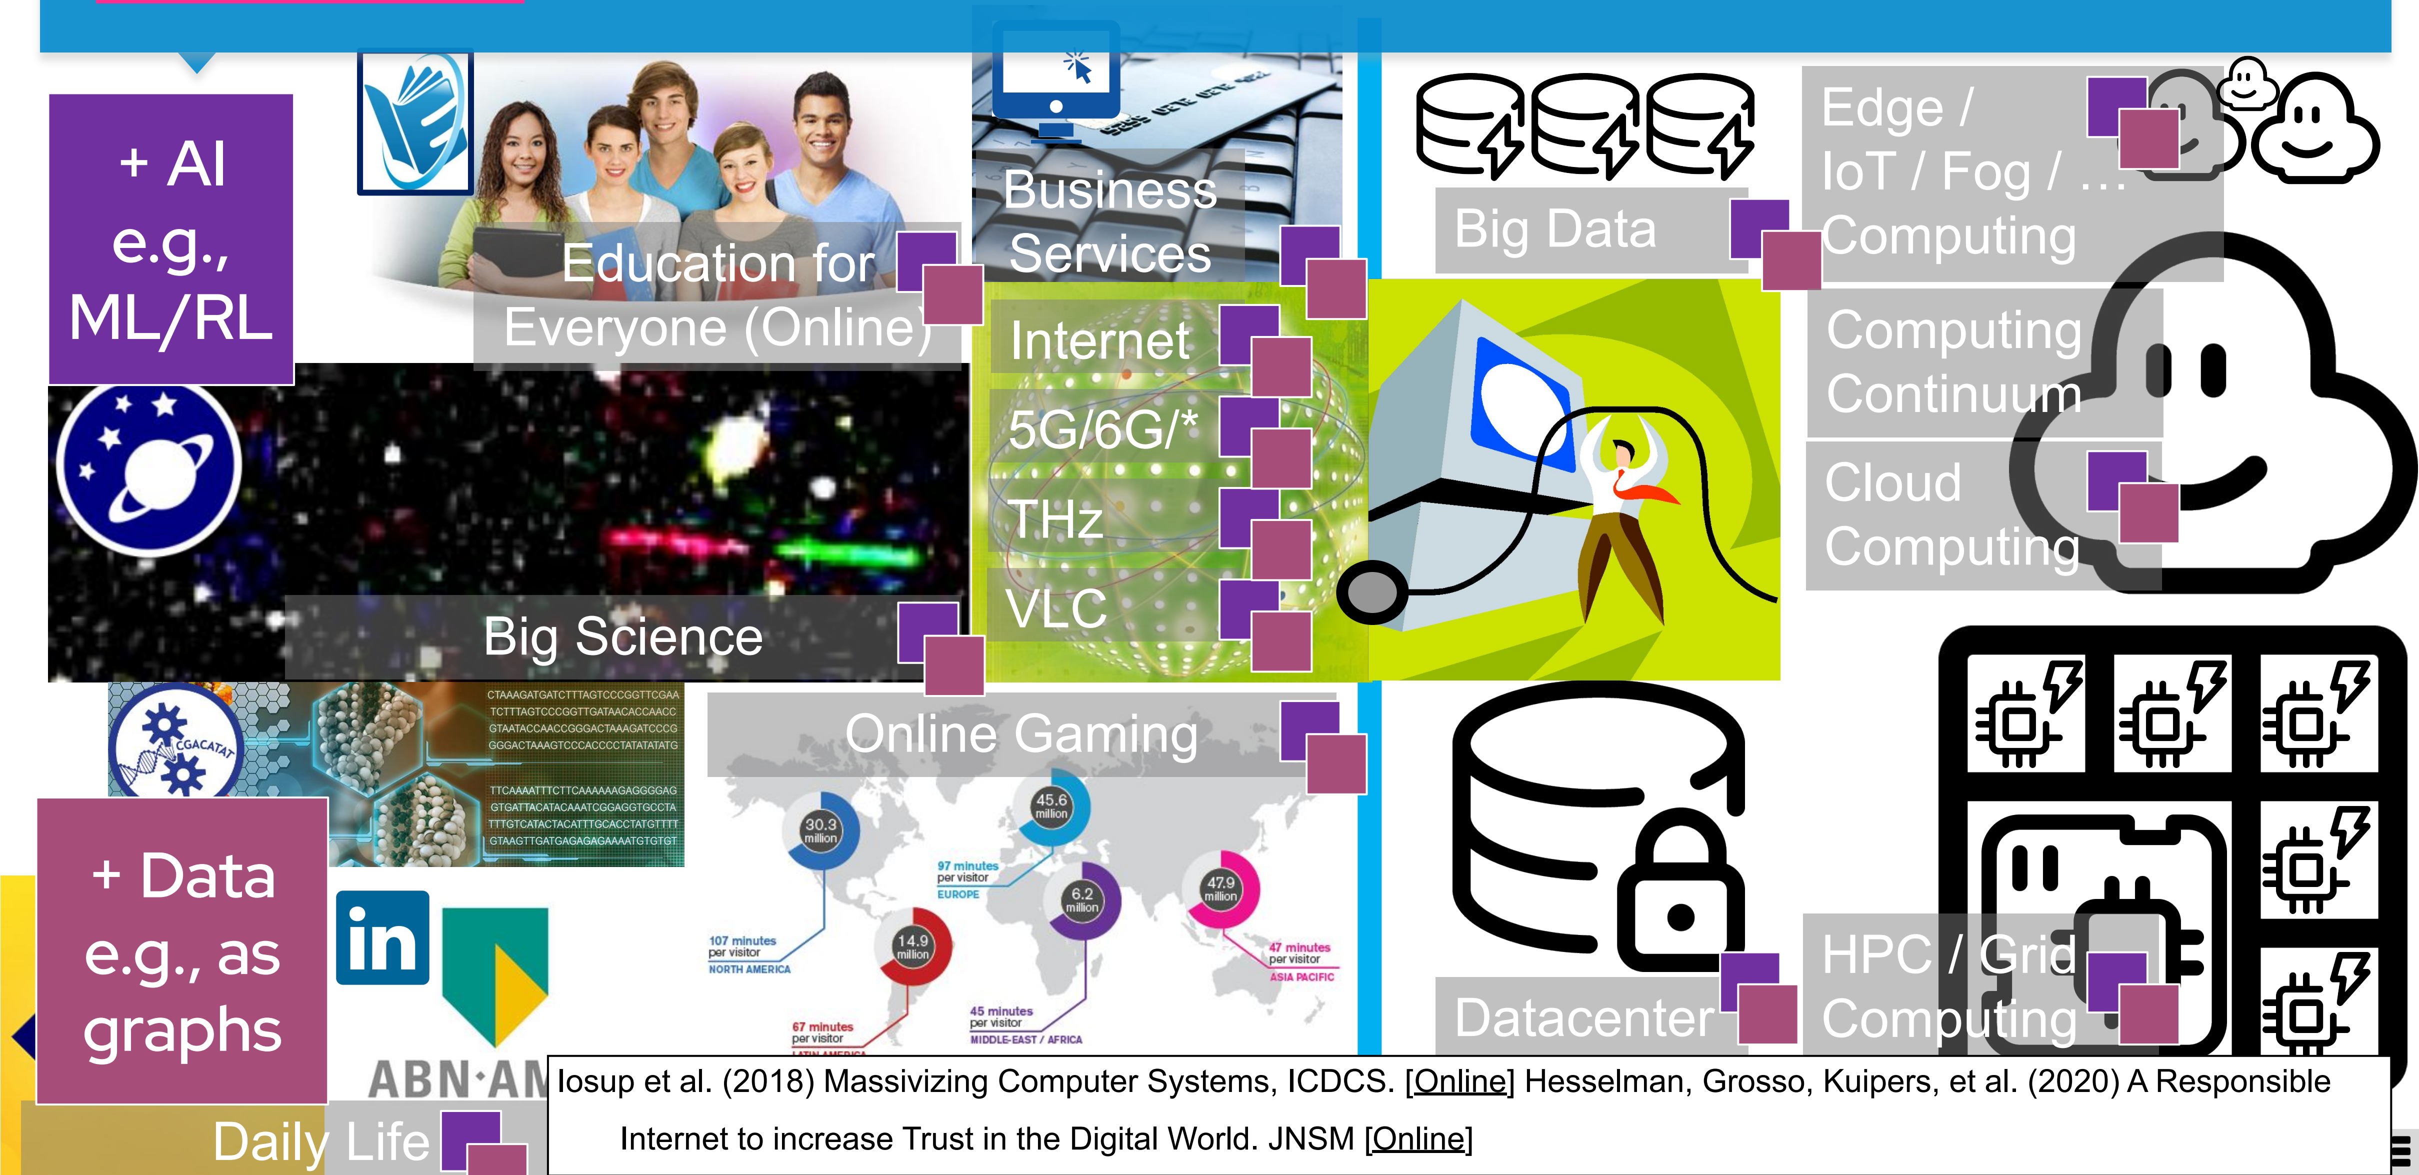

Iosup et al. (2018) Massivizing Computer Systems, ICDCS. [Online] Hesselman, Grosso, Kuipers, et al. (2020) A Responsible Internet to increase Trust in the Digital World. JNSM [Online]

USING REPRODUCIBLE, COMPLEX WORKFLOWS …

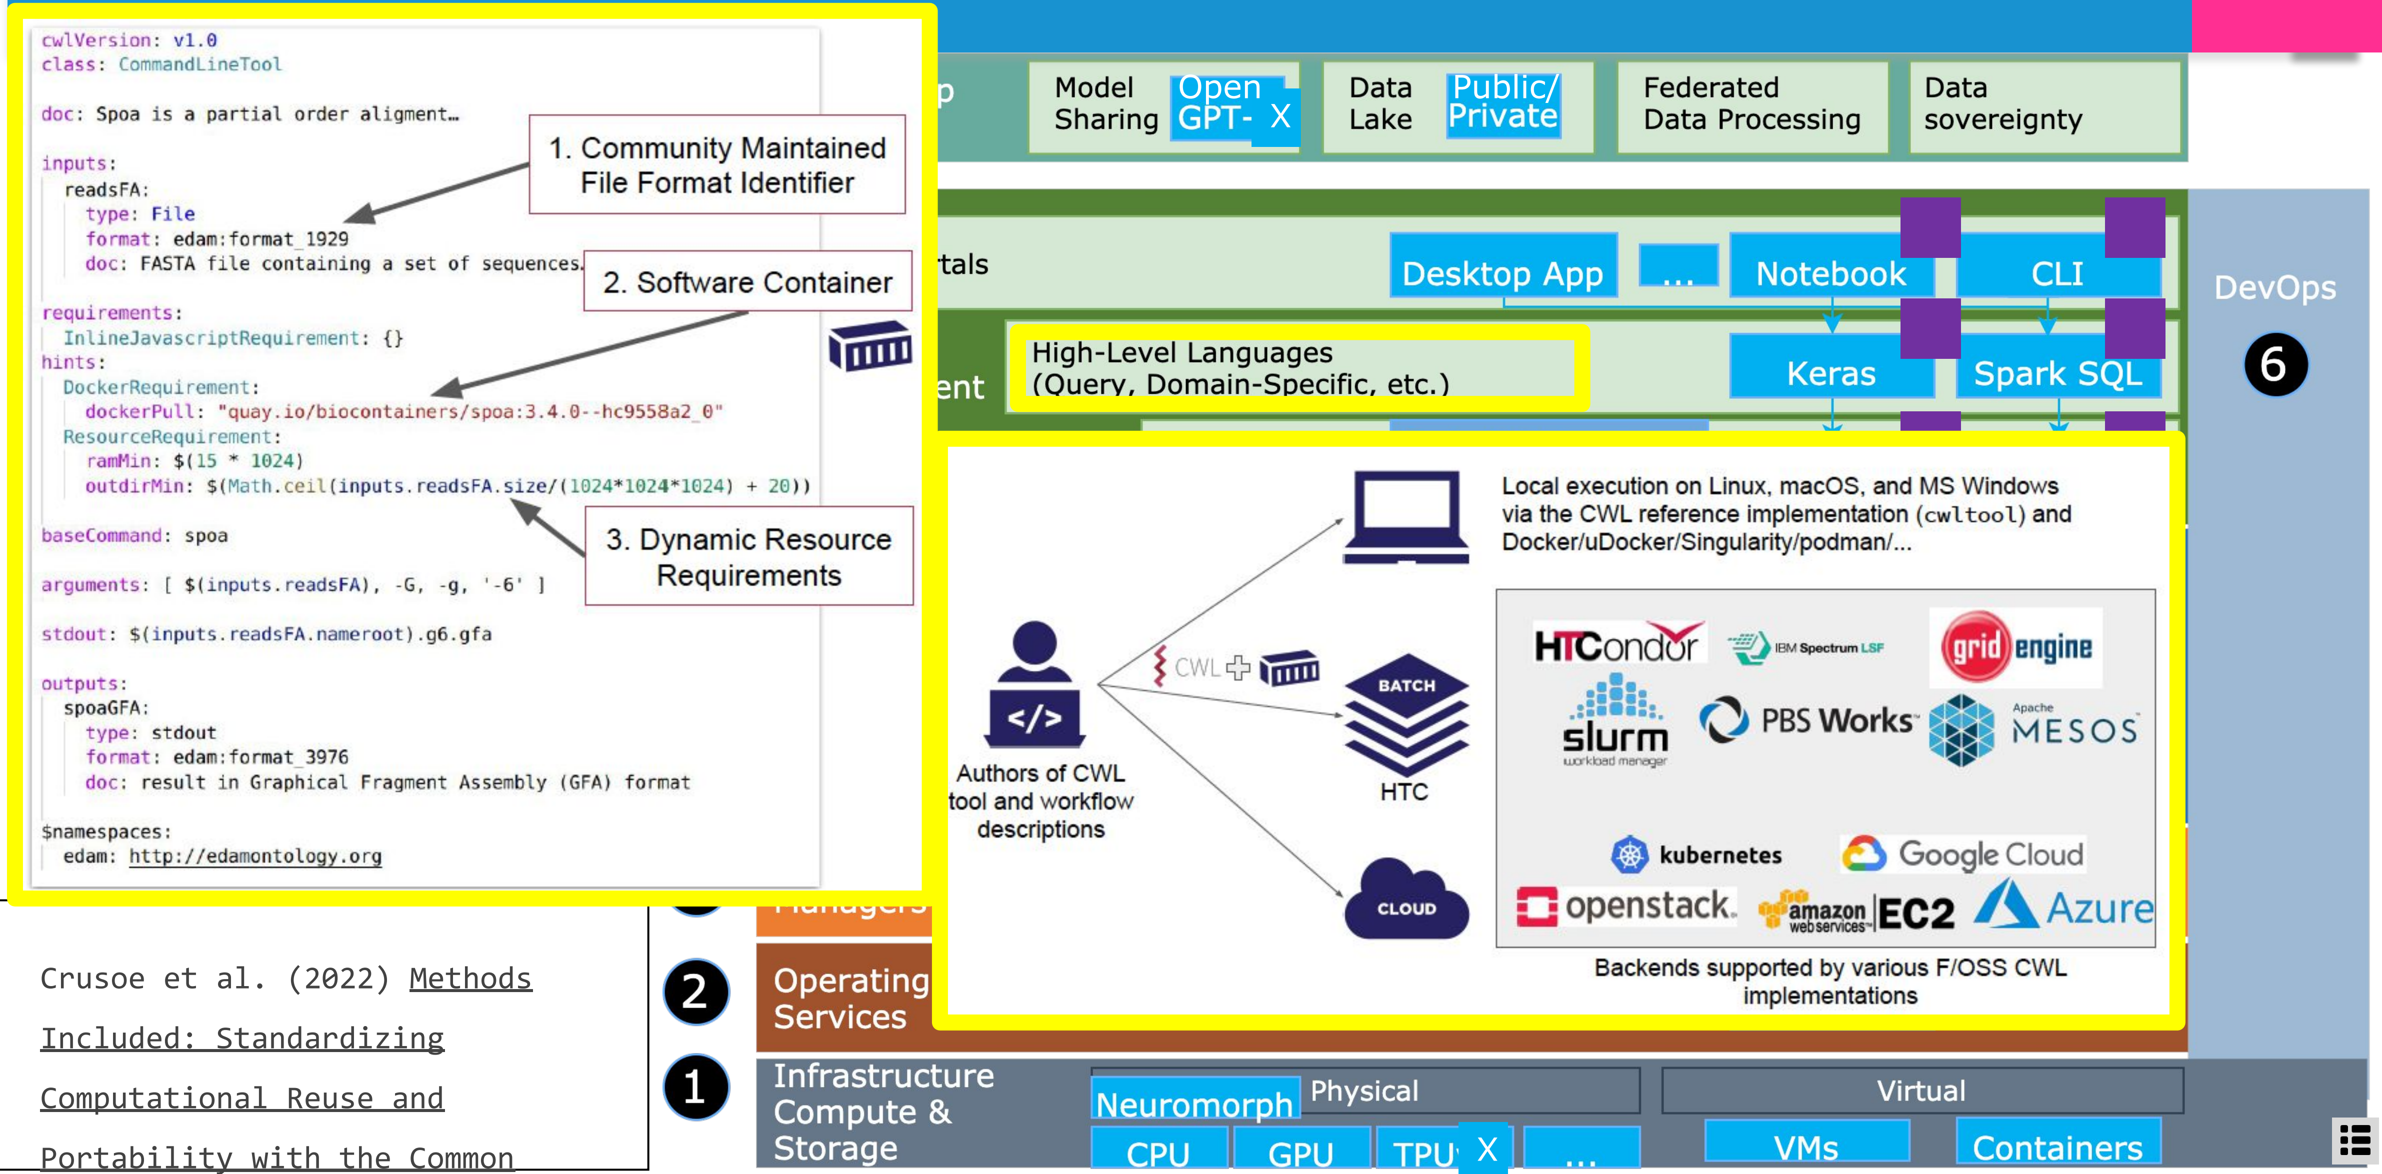

Crusoe et al. (2022) Methods Included: Standardizing Computational Reuse and Portability with the Common

… RUNNING ON CONTINUUM RESOURCES & SERVICES

ISSUES: COMPLEXITY, NON-TECHNICAL

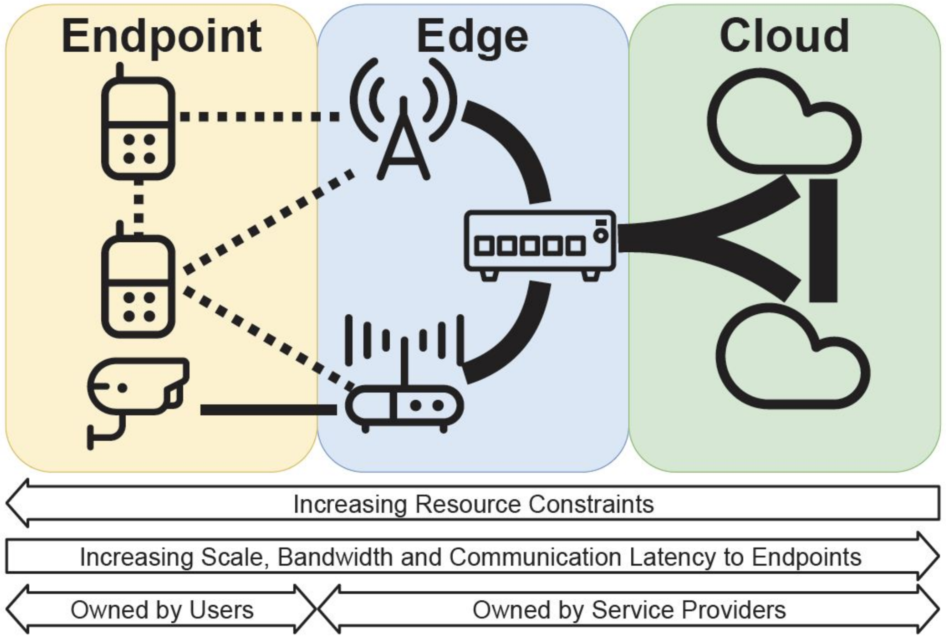

Trivedi, Wang, Bal, Iosup (2021) Sharing and Caring of Data at the Edge. HotEdge.

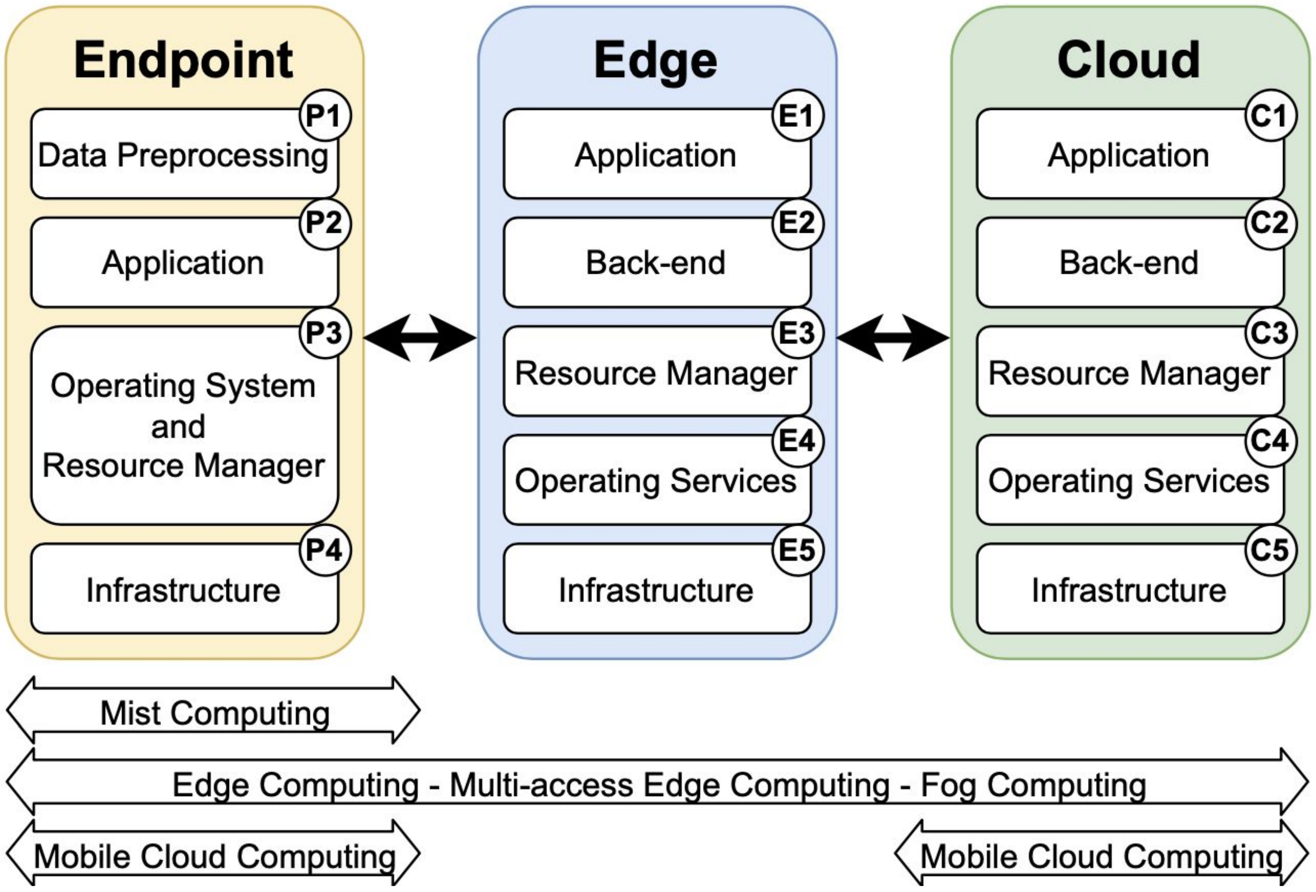

Jansen, Al-Dulaimy, Papadopoulos, Trivedi, Iosup (2023) The SPEC-RG Reference Architecture for the Edge Continuum. CCGRID. Open access:

… IN A SMARTLY ORCHESTRATED ICT ECOSYSTEM …

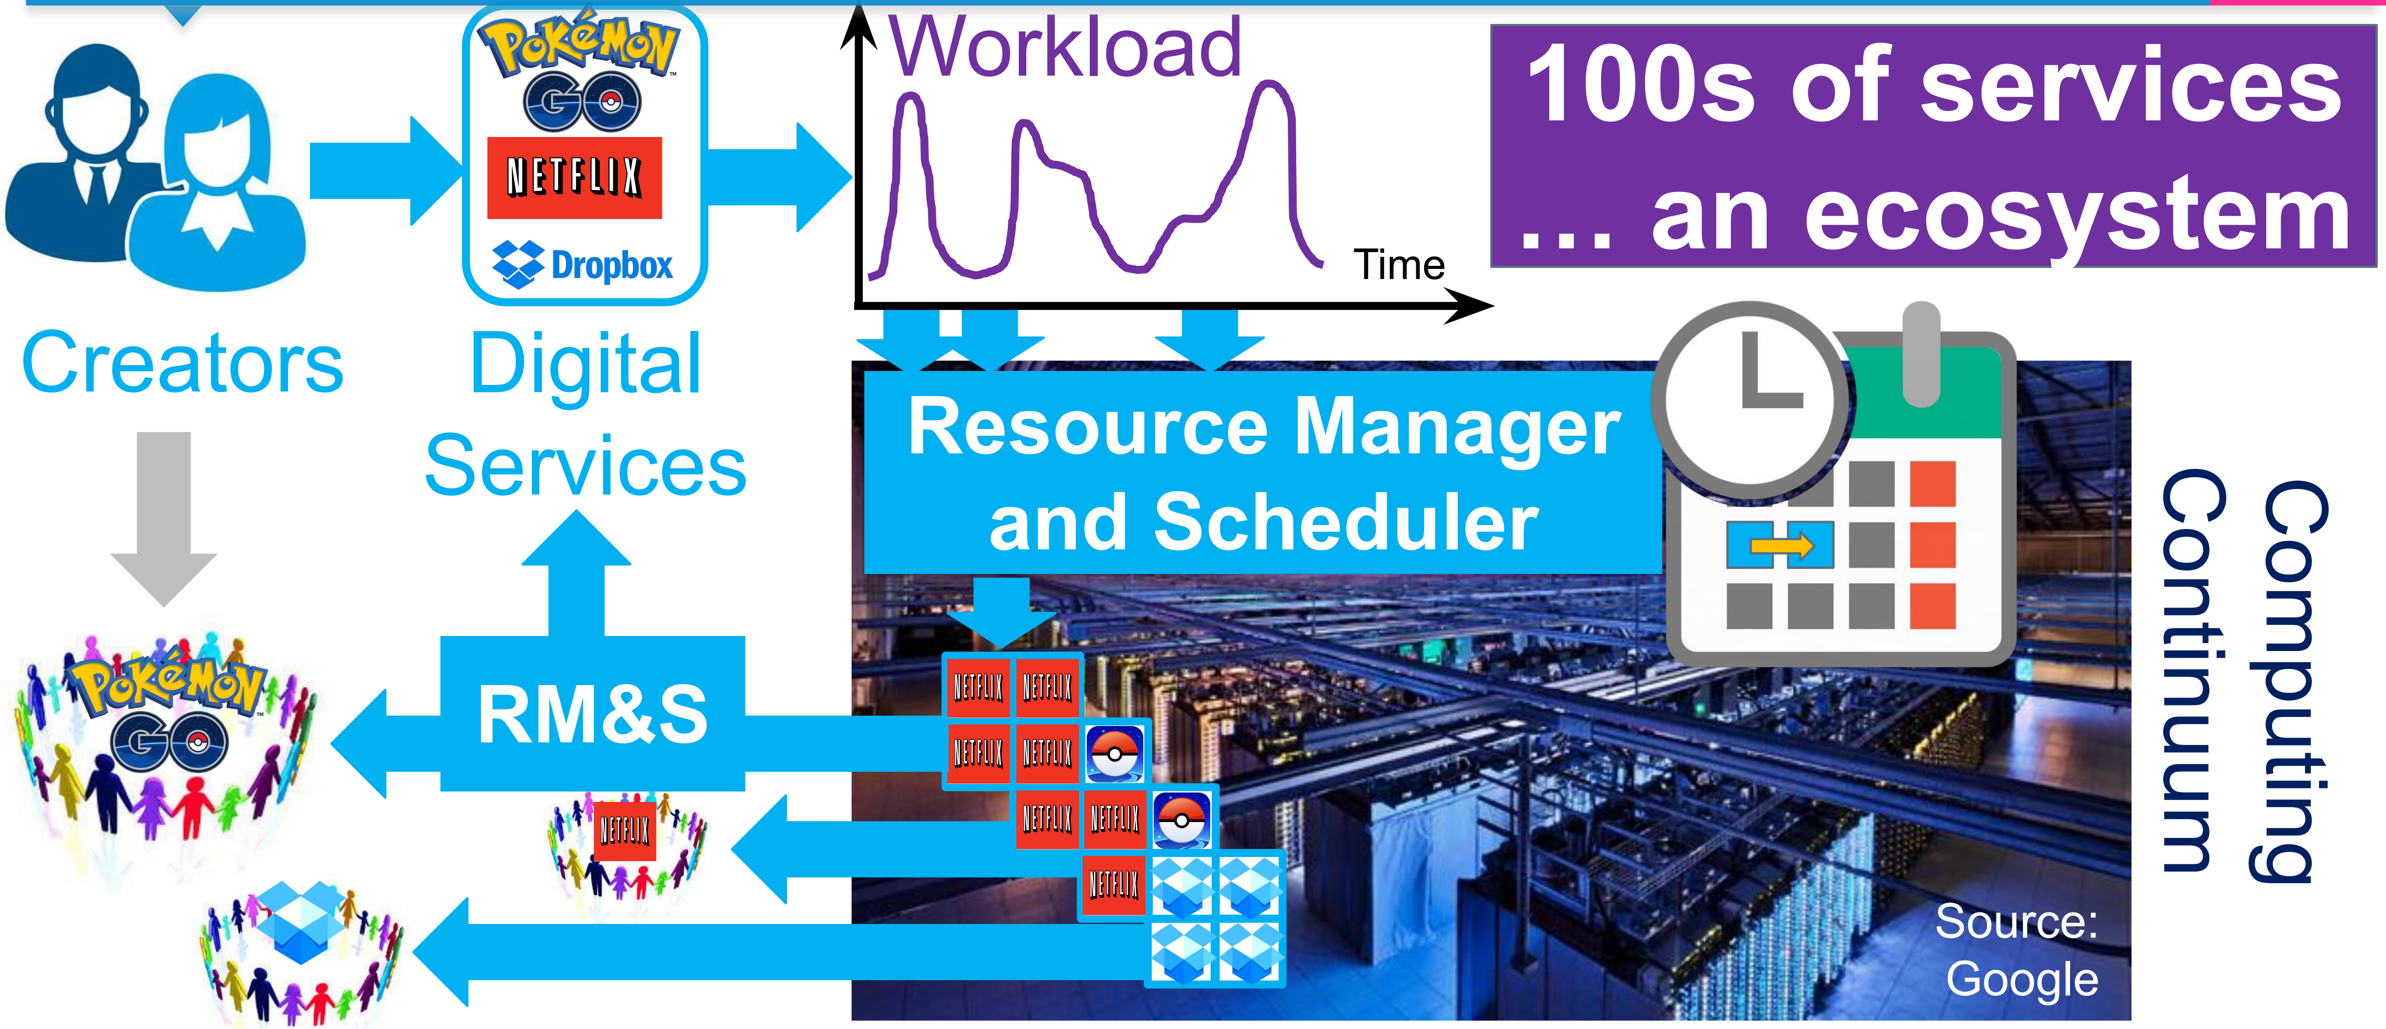

Extreme Automation, Performance, Dependability, Sustainability

…DELIVERING SERVERLESS COMPUTING PROPERTIES

^

bit.ly/MassivizingServerless22

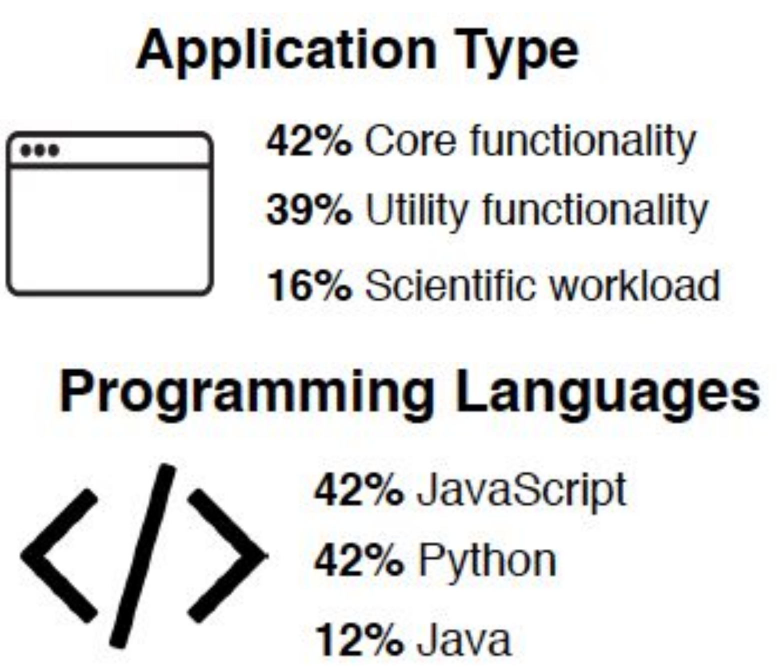

Simon Eismann, Joel Scheuner, Erwin Van Eyk, Maximilian Schwinger, Johannes Grohmann, Nikolas Herbst, Cristina L. Abad, Alexandru Iosup (2022) The State of Serverless Applications: Collection, Characterization, and Community Consensus. IEEE Trans. Software Eng. 48(10)

Serverless computing = extreme automation + fine-grained, utilization-based billing

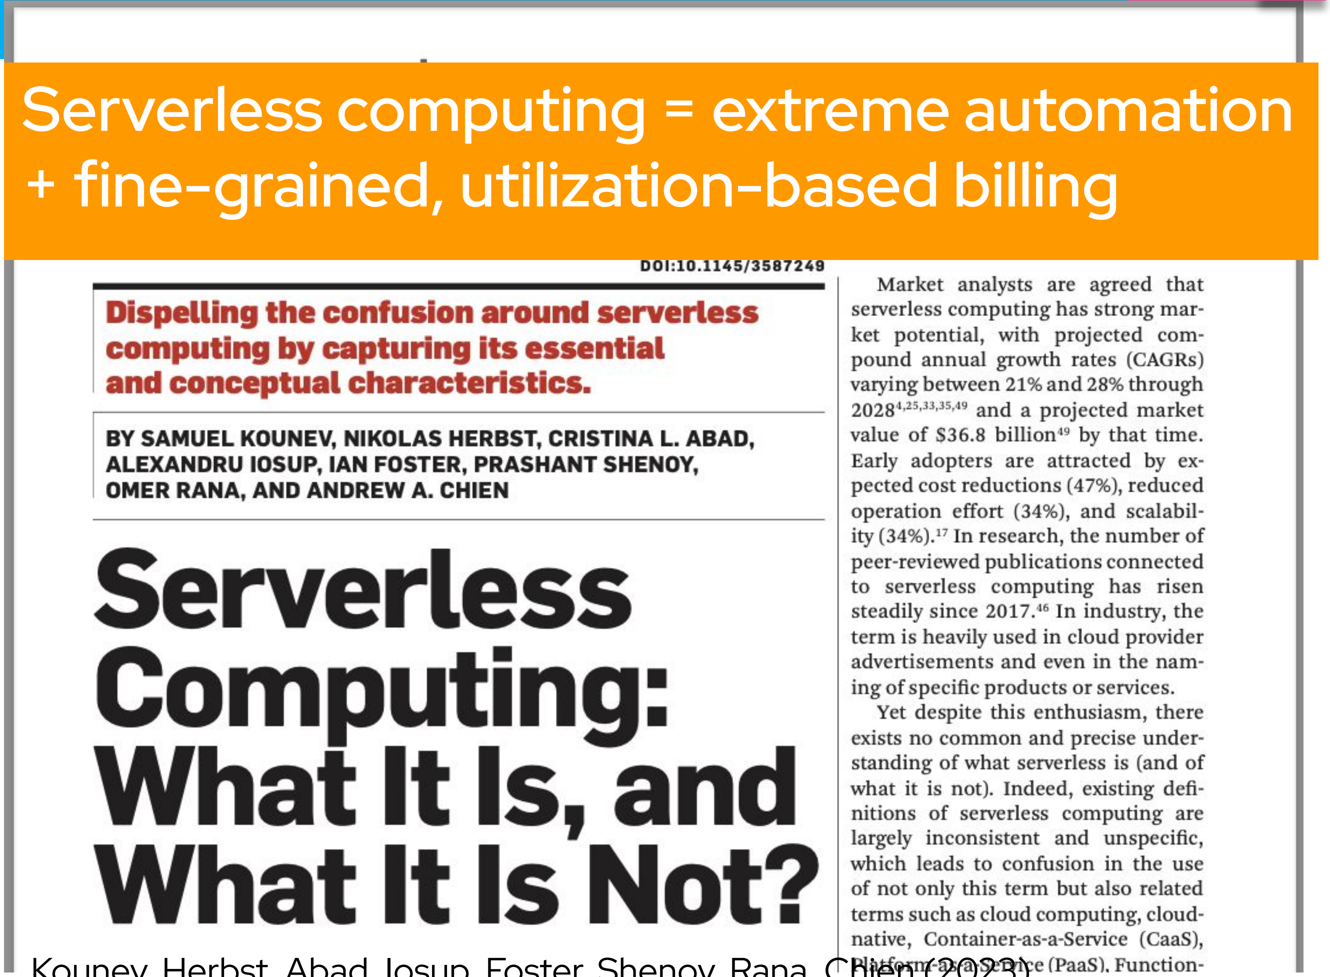

DOI:10.1145/3587249

Dispelling the confusion around serverless computing by capturing its essential and conceptual characteristics.

BY SAMUEL KOUNEV, NIKOLAS HERBST, CRISTINA L. ABAD, ALEXANDRU IOSUP, IAN FOSTER, PRASHANT SHENOY, OMER RANA, AND ANDREW A. CHIEN

Serverless Computing: What It Is, and What It Is Not?

Market analysts are agreed that serverless computing has strong market potential, with projected compound annual growth rates (CAGRs) varying between 21% and 28% through 2028[4,25,33,35,49] and a projected market value of $36.8 billion[49] by that time. Early adopters are attracted by expected cost reductions (47%), reduced operation effort (34%), and scalability (34%).[17] In research, the number of peer-reviewed publications connected to serverless computing has risen steadily since 2017.[46] In industry, the term is heavily used in cloud provider advertisements and even in the naming of specific products or services.

Yet despite this enthusiasm, there exists no common and precise understanding of what serverless is (and of what it is not). Indeed, existing definitions of serverless computing are largely inconsistent and unspecific, which leads to confusion in the use of not only this term but also related terms such as cloud computing, cloud-native, Container-as-a-Service (CaaS), Platform-as-a-Service (PaaS), Function-

Kounev, Herbst, Abad, Iosup, Foster, Shenoy, Rana, Chien (2023) Serverless Computing: What It Is, and What It Is Not? CACM. Sep 2023 issue.

BUT WE CANNOT TAKE THIS TECHNOLOGY FOR GRANTED

2

(We need science to tackle the issues)

RESPONSIBILITY OF MASSIVE COMPUTER ECOSYSTEMS

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

But availability not as high as believed

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.

&

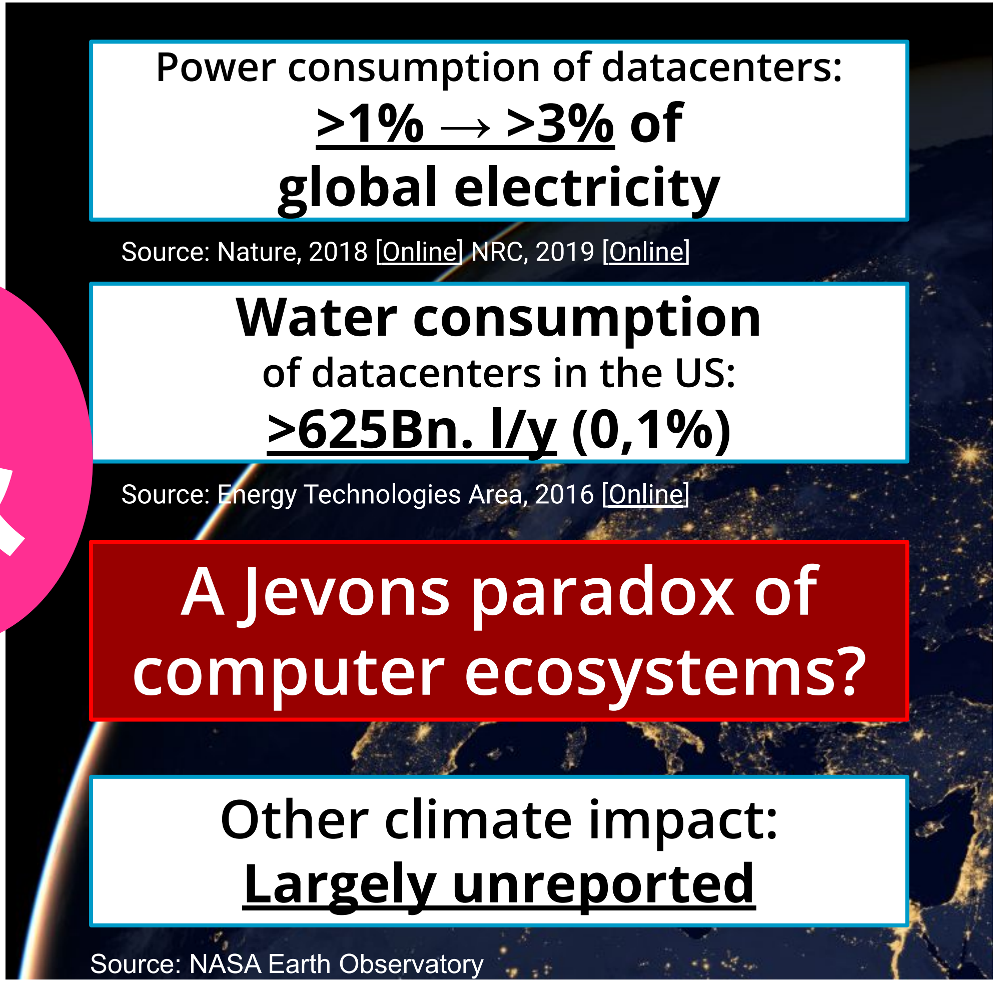

Power consumption of datacenters:
>1% → >3% of global electricity

Source: Nature, 2018 [Online] NRC, 2019 [Online]

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

Source: Energy Technologies Area, 2016 [Online]

A Jevons paradox of computer ecosystems?

Other climate impact:
Largely unreported

Source: NASA Earth Observatory

S. Talluri, L. Overweel, L. Versluis, A. Trivedi, A. Iosup (2021) Empirical Characterization of User Reports about Cloud Failures. ACSOS.

TECHNOLOGY EVOLUTION

END OF MOORE'S LAW/DENNARD SCALING → COMPLEX, DISTRIBUTED ECOSYSTEMS

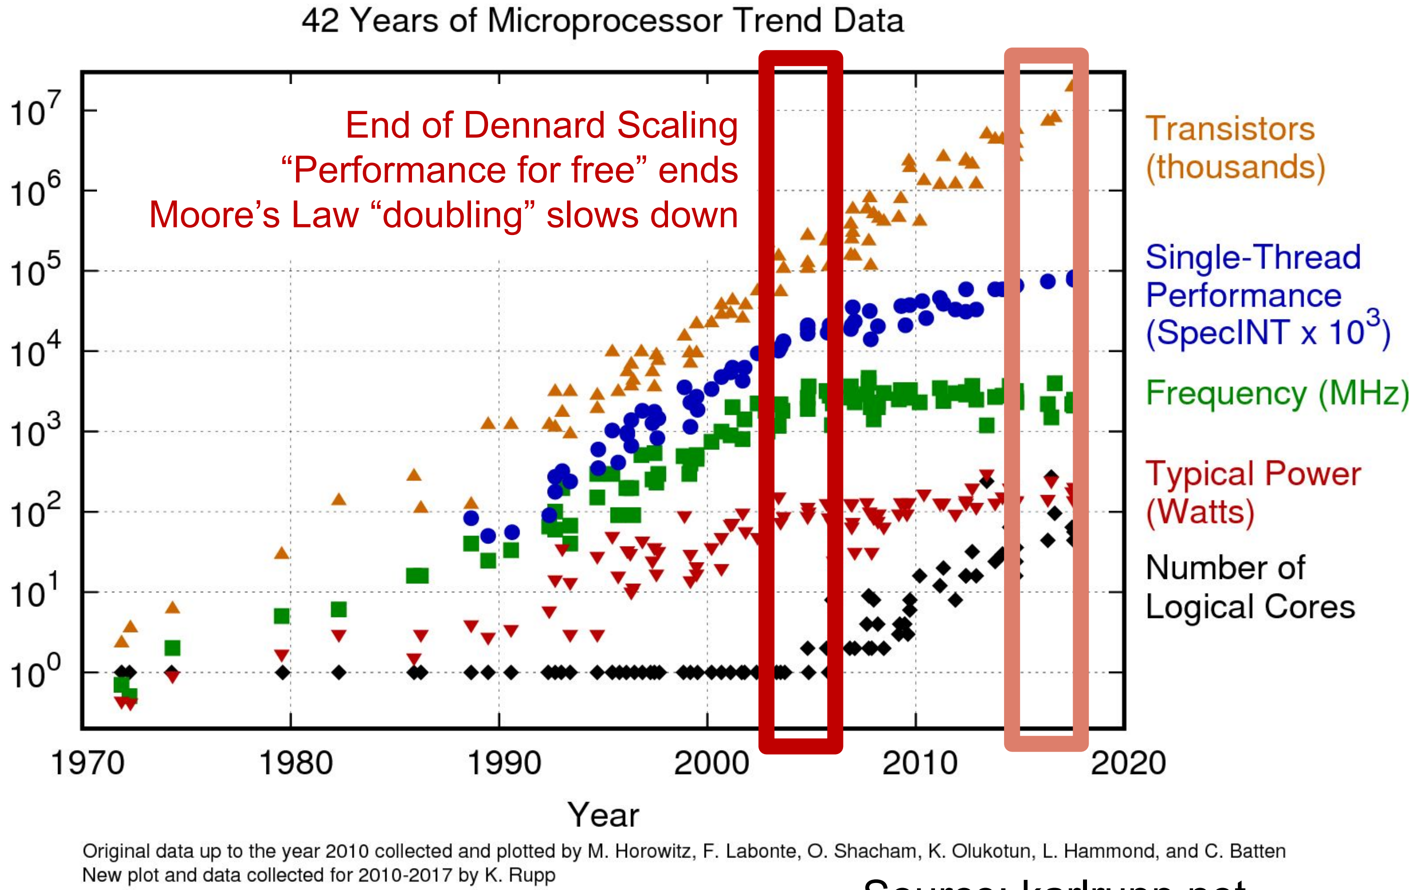

Source: karlrupp.net

Need for parallelism and distribution

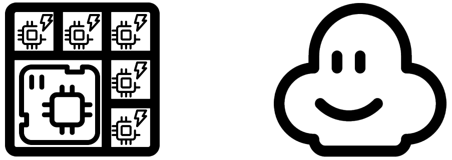

Cost of energy and data movement still high

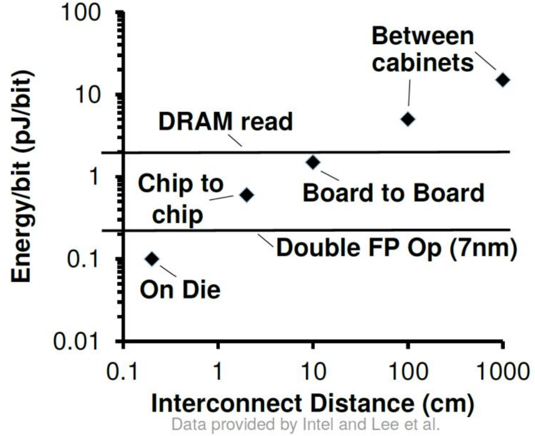

COMPLEXITY GROWS

COMPLEX, DISTRIBUTED ECOSYSTEMS DO NOT ACT LIKE REGULAR COMPUTER SYSTEMS

Ecosystems don't have easily…

- Simplicity
- Maintainability
- Responsibility
- Sustainability
- Usability

which includes:

- Synchronization
- Consistency, consensus
- Performance
- Scalability, elasticity
- Availability, reliability
- Energy-efficiency

Iosup, Kuipers, Trivedi, et al. (2022) Future Computer Systems and Networking Research in the Netherlands: A Manifesto. CORR

Operational goals are becoming more complex

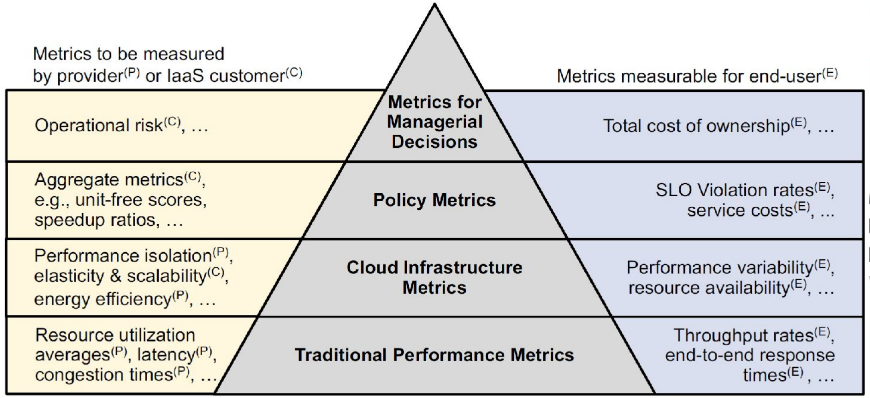

N. Herbst, E. Van Eyk, A. Iosup, et al. (2018) Quantifying Cloud Performance and Dependability: Taxonomy, Metric Design, and Emerging Challenges. TOMPECS 3(4).

Operational techniques are becoming more complex

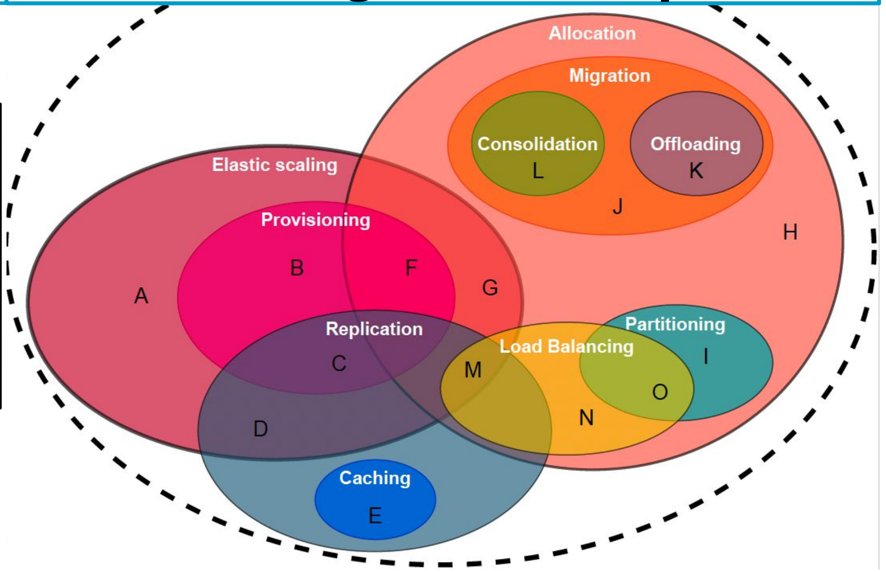

Stijn Meijerink, Erwin van Eyk, Alexandru Iosup (2021) Multivocal Survey of Operational Techniques for Serverless Computing. White Paper.

FEW CAN OPERATE COMPLEX IT ECOSYSTEMS

THE WORKFORCE GAP, IN THE NETHERLANDS & IN EUROPE

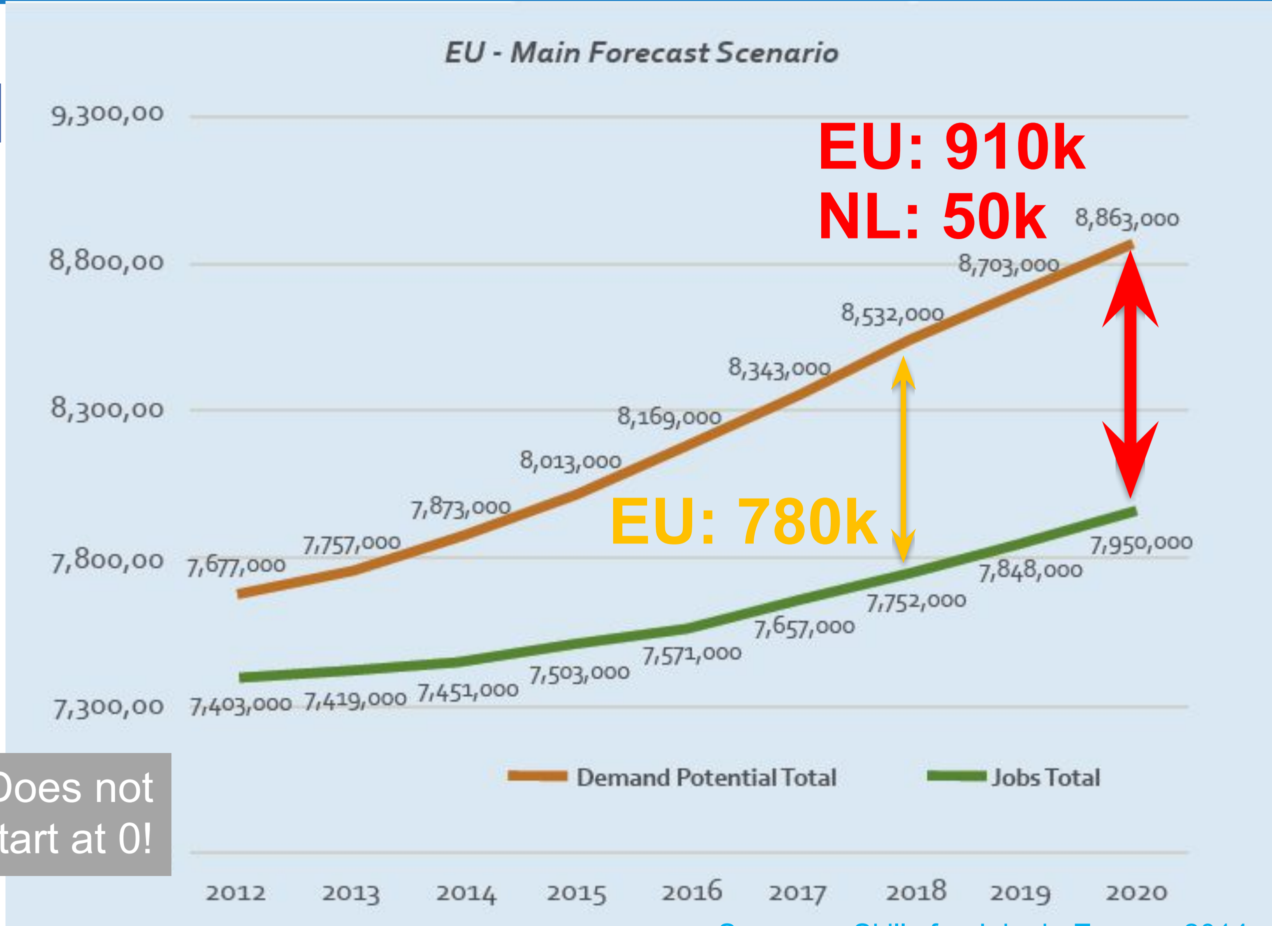

Does not start at 0!

Source: e-Skills for Jobs in Europe, 2014

IN THIS TALK:
BUILDING AN ICT INFRASTRUCTURE MEMEX TO ADDRESS LONG-TERM ODA NEEDS

3

(What we need in CompSys infrastructure)

Building the Infrastructure Memex: VU on Operational Data Analytics in the 21st Century

THIS TALK IN A NUTSHELL: A HOLISTIC VIEW, BASED ON ECOSYSTEM INTROSPECTION

Technology **not ready, many issues**

A Why does this happen?

B What to do about it?

In modern computer systems, issues are often linked

Source: Alexandru’s personal library.

DISCOVERY = LARGE-SCALE, LONG-TERM STUDY

+

UNCOVERING THE MYSTERIES OF OUR PHYSICAL UNIVERSE

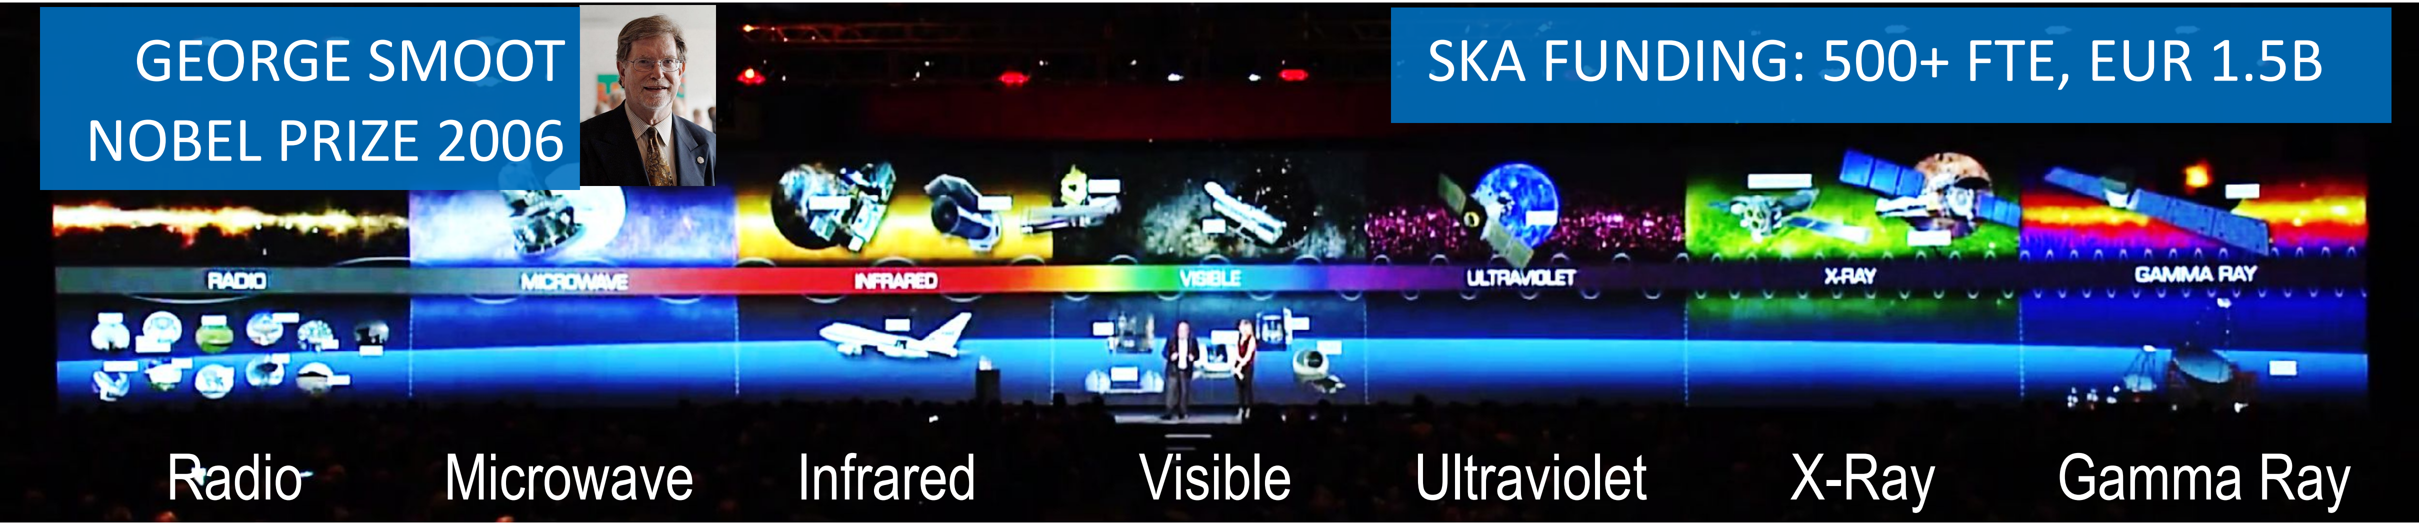

James Cordes, The Square Kilometer Array, Project Description, 2009 [Online]

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

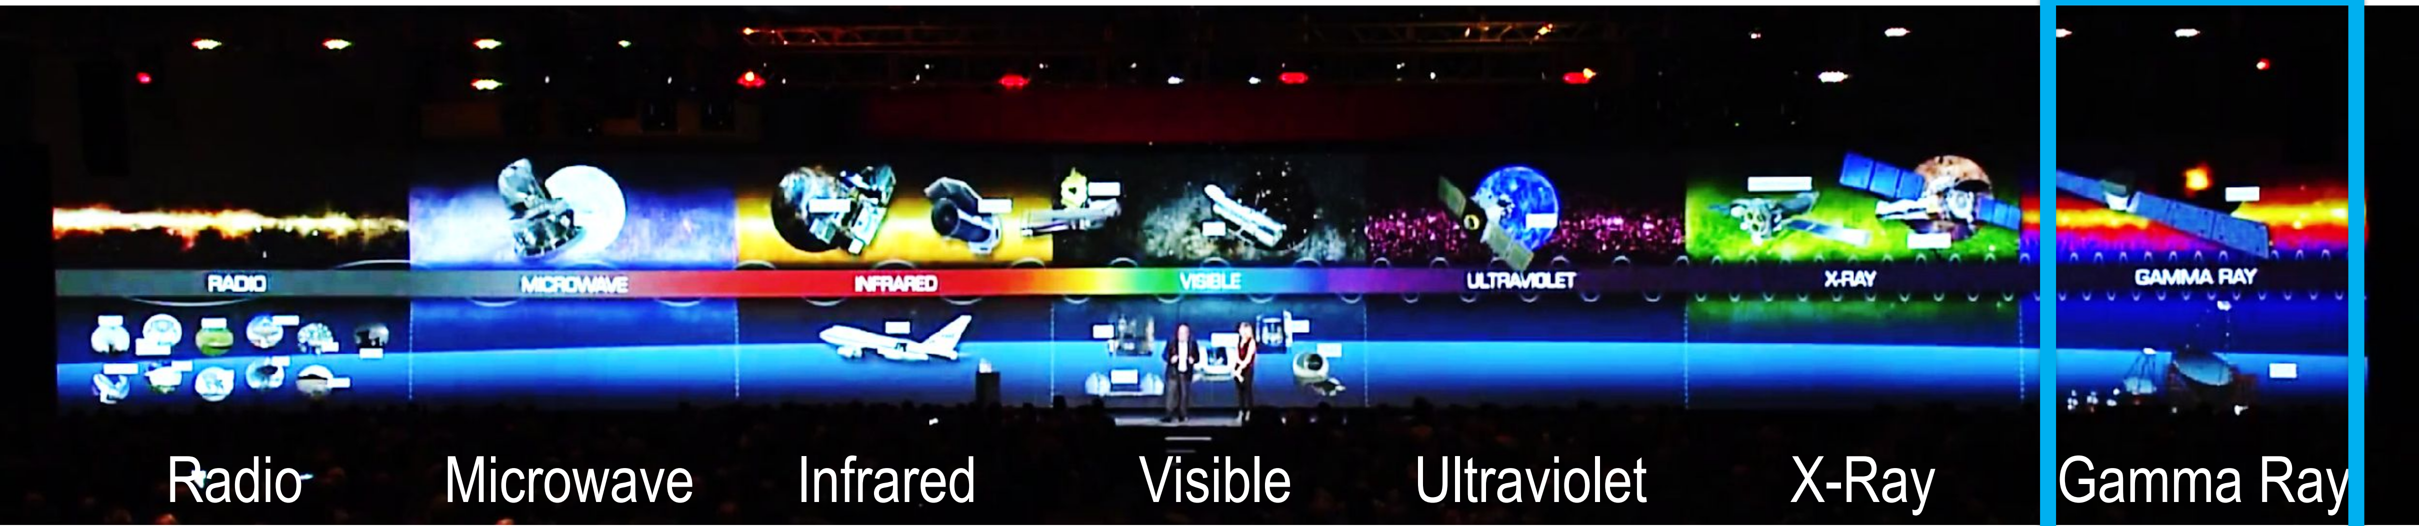

Radio	Microwave	Infrared	Visible	Ultraviolet	X-Ray	Gamma Ray
Cloud, Grid, Edge, Fog, etc.	One aspect: BigData, P2P	Sci.&Eng. Apps+Sys.	Consumer Apps+Sys.	Enterprise Sys.	Systems, Ecosystems	Performance Availability, etc.

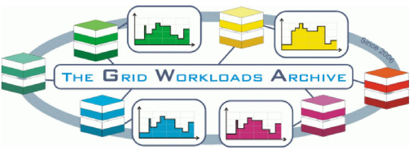

[Iosup et al.

P2P T.A.

[Zhang et al.

[Versluis et al. TPDS'21]

Groups, Workflows

[Iosup et al.

Game Trace Archive

[Guo et al.

Business -Critical

[Shen et al.

[Ghiț et al.

[Versluis et al. FGCS'22]

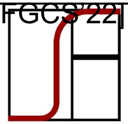

[Iosup et al.

OUR VISION: THE ICT INFRASTRUCTURE MEMEX

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

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

Find and eradicate performance issues

Get quantitative evidence

Enable new designs and automation

Consider culture and ethics

Understand how entire ecosystems behave and evolve

Cloud, Grid, Edge, Fog, etc.

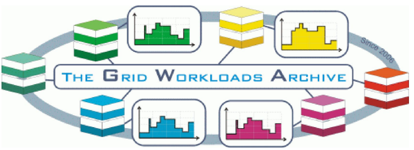

[Iosup et al.

One aspect: BigData, P2P

P2P T.A.

[Zhang et al.

Sci.&Eng. Apps+Sys.

[Versluis et al. TPDS'21]

Groups, Workflows

[Iosup et al.

Consumer Apps+Sys.

Game Trace Archive

[Guo et al.

Enterprise Sys.

Business -Critical

[Shen et al.

Systems, Ecosystems

[Ghiț et al.

Performance, Availability, etc.

[Versluis et al. FGCS'22]

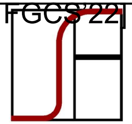

[Iosup et al.

TYPICAL RESULTS

4

DISCOVER PHENOMENA, EXPLAIN WHAT'S HAPPENING

A

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.

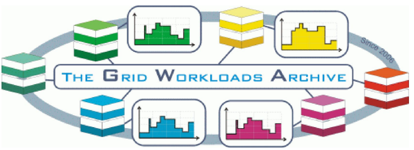

[Iosup et al. FGCS'08]

P2P T.A.

[Zhang et al. CoNext'10]

[Versluis et al. TPDS'21]

Groups, Workflows

[Iosup et al. IEEE IC'11]

Game Trace Archive

[Guo et al. NETGAMES'12]

Business -Critical

[Shen et al. CCGRID'15]

[Ghiț et al. CCGRID'14]

[Versluis et al. FGCS'22]

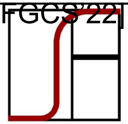

[Iosup et al. CCGRID'10]

One Basic Result

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

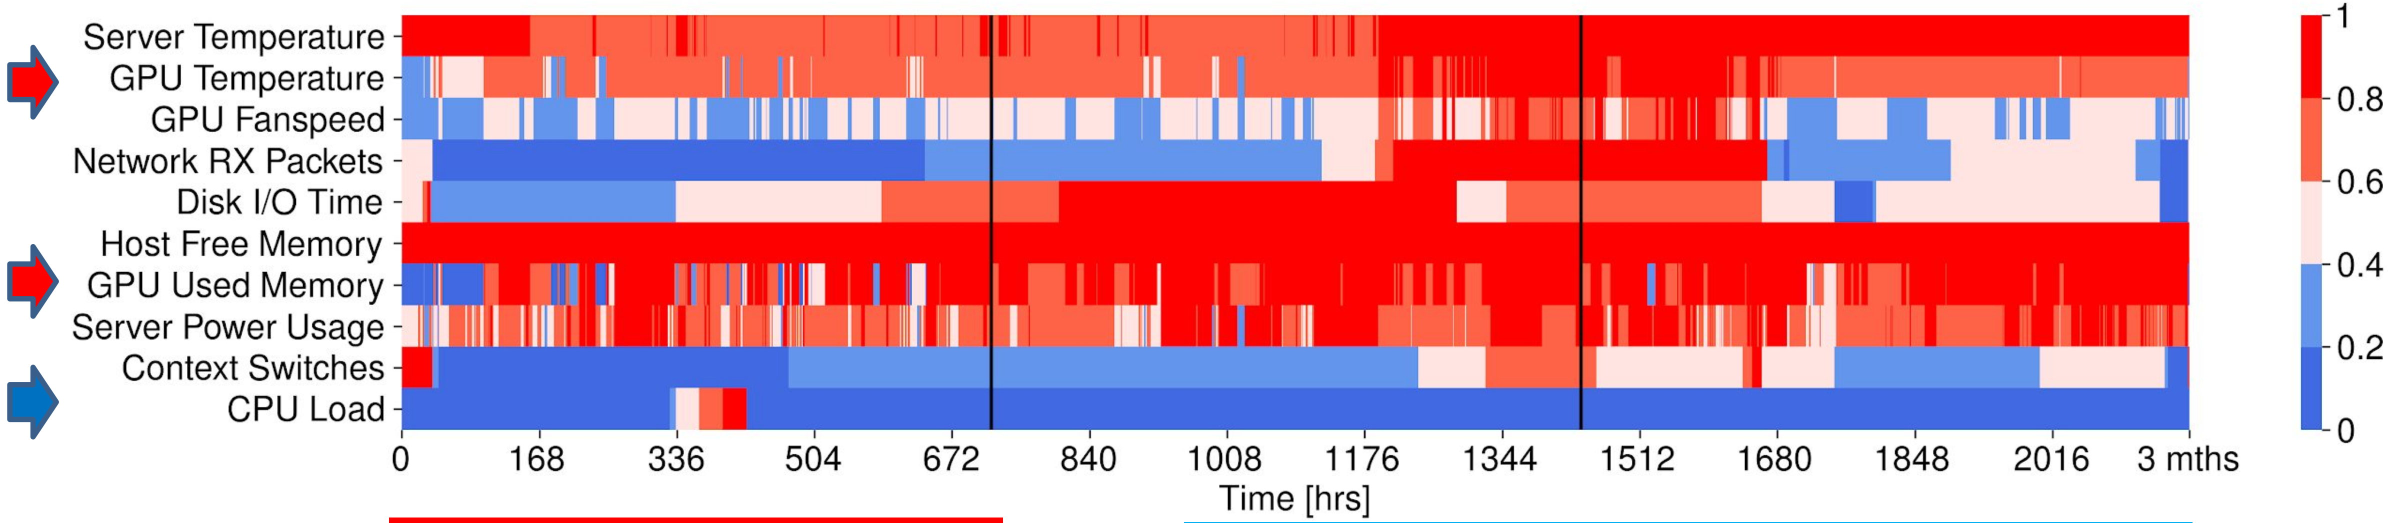

Hot GPUs

Cold CPUs*

* when GPUs

Could use less powerful, less expensive CPUs

One Phenomenon: BoTs = Dominant Programming Model for Grid Computing

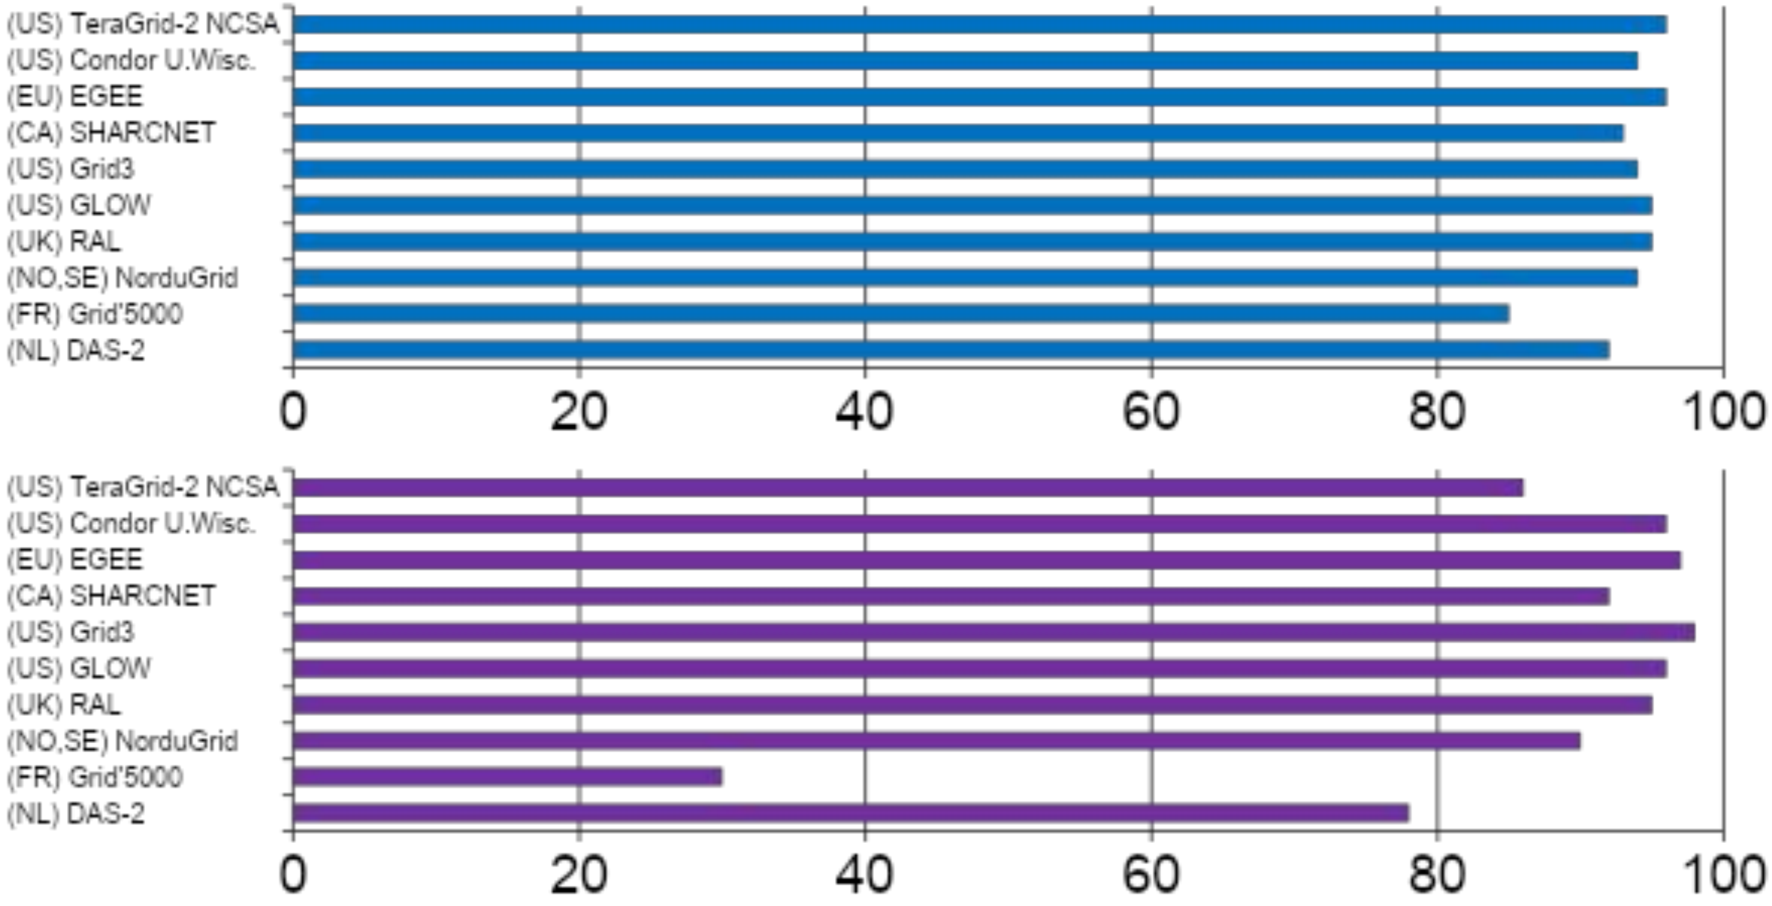

Each bar is for one grid

Each bar is a long-term workload trace

Iosup and Epema: Grid Computing Workloads. IEEE Internet Computing 15(2): 19-26 (2011)

THE WORKFLOW TRACE ARCHIVE

96 traces

Laurens Versluis

METADATA AND TRACES FOR YOUR WORKFLOW SYSTEMS

WORKFLOWS ARE COMMON IN MANY DOMAINS

EXCEPT IN SCI., DESIGN, & ENG.

THE WORKFLOW TRACE ARCHIVE
 WORKFLOWS ARE DIVERSE!

Protein BLAST
protein ▸ protein

Shell

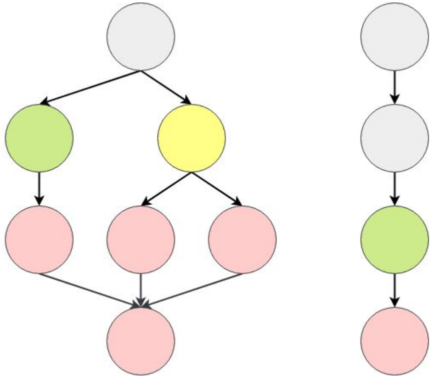

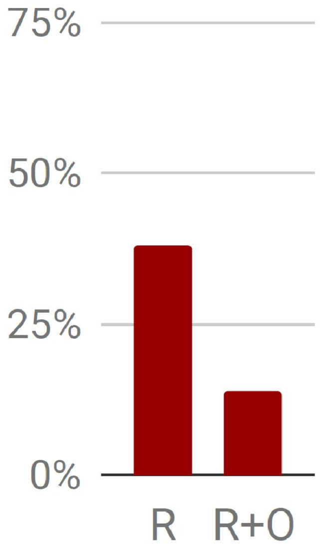

universität innsbruck

Alibaba.com® Global trade starts here.™

WorkflowHub

Shell

TWO SIGMA

Los Alamos NATIONAL LABORATORY

Google

TU Delft Delft University of Technology

Pegasus

VU UNIVERSITY AMSTERDAM

spec Research

[Versluis et al. The Workflow Trace Archive]
IEEE TPDS + http://arxiv.org/pdf/1906.07471

http:// wta.atlarge.science

USE A SIMULATOR TO ENABLE ICT DIGITAL TWINS

B

… CAN WE AFFORD A? WHAT IF B HAPPENS? HOW DOES C EVOLVE? X vs. Y … vs. Z?

[OpenDC](https://opendc.org) simulator

- 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

and more…

Learn more: opendc.org

Electricity expenses become High Risk

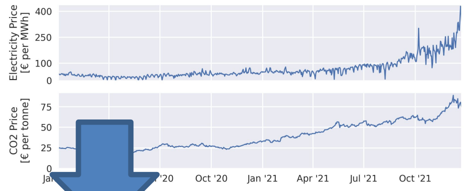

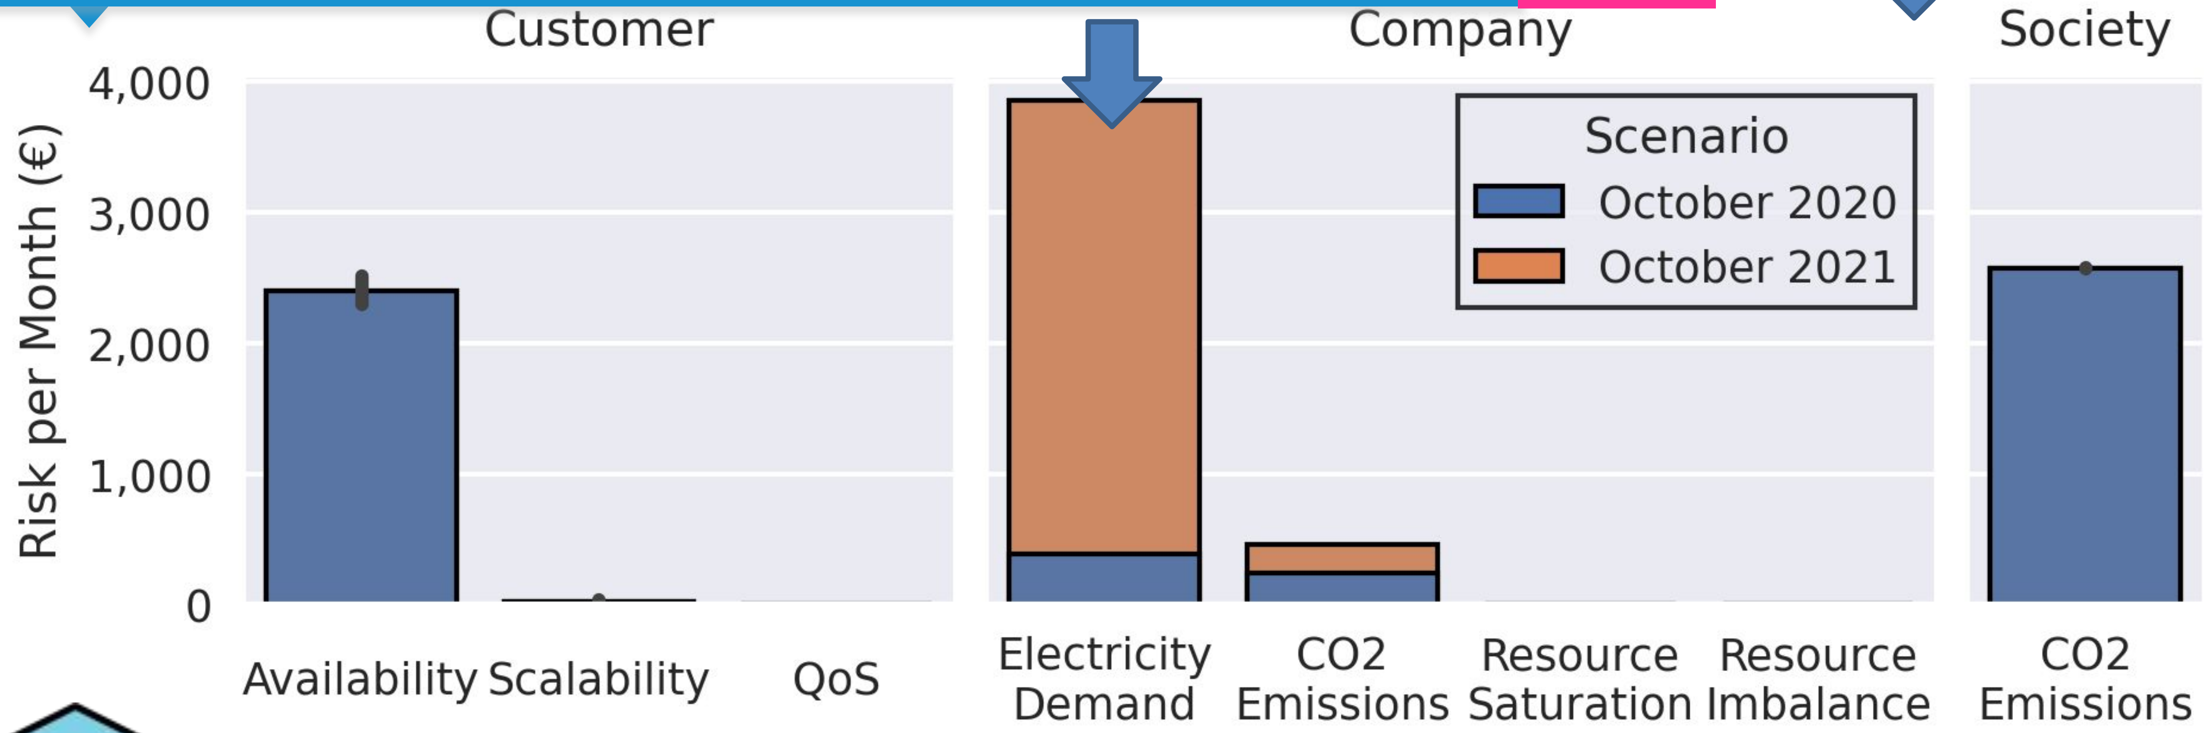

Electricity expenses could become **primary risk** in datacenters

Experiments are very expensive!
An Environmental Perspective

GradeML > Systems Memex >

112

Core-hours

2.1

kWh

Delft →

→ Rotterdam

In Simulation

8,600,000,000

Core-hours

116,000,000,000

kWh

x 55 billion!

In Reality

Exact numbers confidential, depend on topology

UNDER THE HOOD: WHAT'S IN THE ICT INFRASTRUCTURE MEMEX?

5

(How to get to Operational Data Analytics?)

SOME QUESTIONS

bit.ly/
AIScalPerf23

We're building the ICT Infrastructure Memex to enable Operational Data Analytics into the 21st century. We have many (developing) theories and practical results. We know others are working on this topic and seek discussion and an active collaboration. To discuss:

- Theory: What are the core components of ODA? How do they relate - what is a good reference architecture for ODA?
- Theory: Are performance and availability just parts of a conceptual continuum? What else is in there and how to include it in ODA?
- Theory: We say: Just touch it with your lower lip, briefly, and move away if too hot. How to reason about energy use, both short- and long-term?
- Practice: Ontology building for sharing traces collected in ICT infrastructure. Do we need ontologies? What is a good middle ground between implicit ontologies and exhaustive ontologies?
- Practice: What kinds of ODA benefit from more complex types of analysis, e.g., graph analytics and learning?
- Ethics: If a tree falls in front of you, do you have to see it? From plausible deniability to responsibility in ICT infrastructure management.

ONE SYSTEM MODEL: FITS AI/ML, BIG DATA, SCIENTIFIC, ENGINEERING, BUSINESS CRITICAL, ONLINE GAMING, OTHER APPS

ML app is a small part, but now we can do complex workflows

Rest is systems, HW+SW, including HPC

Adapted from:

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

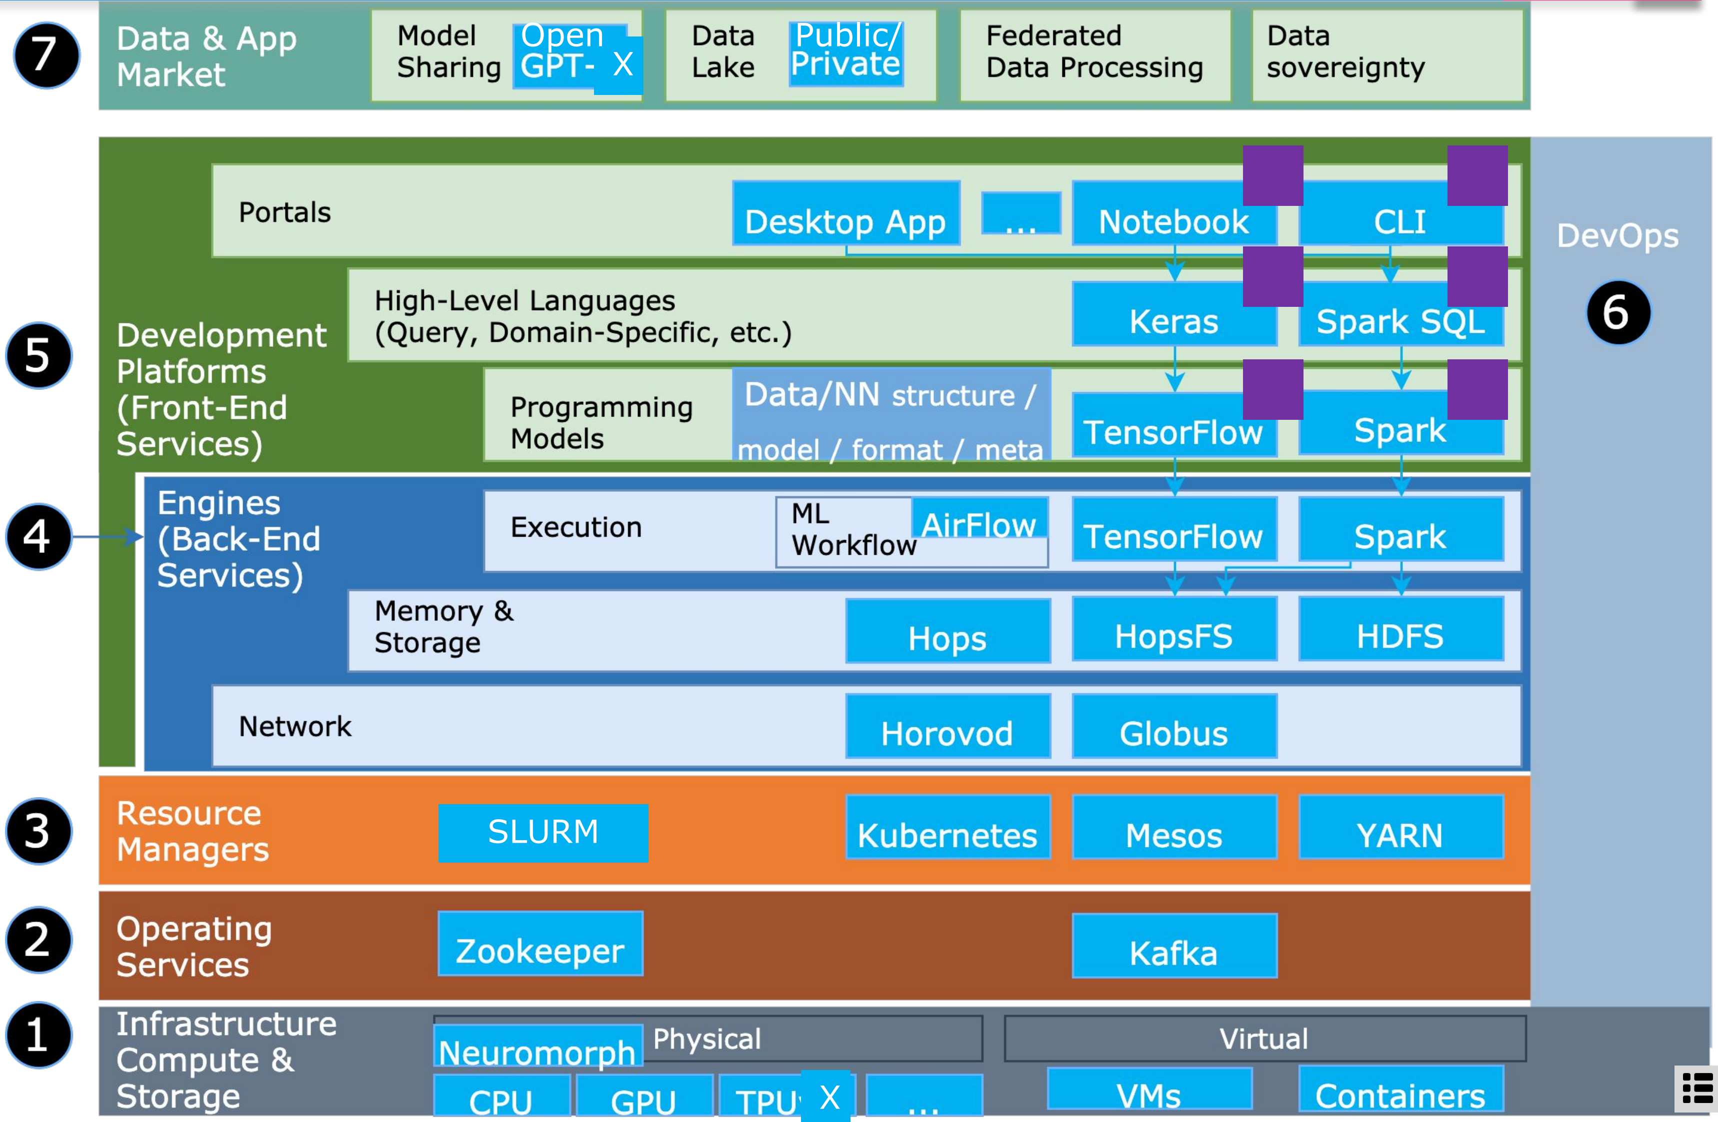

ODA REFERENCE ARCHITECTURE (OPERATIONALIZE)

Problem:
Data often seen as overhead, loose process

Solution:

1. Link data collection to per-layer capabilities
2. Enable data science process

Ongoing work

Shekhar Suman, Xiaoyu Chu, Martin Molan, Andrea Bartolini (UniBo), Iosup, et al. (2023) Ontology for

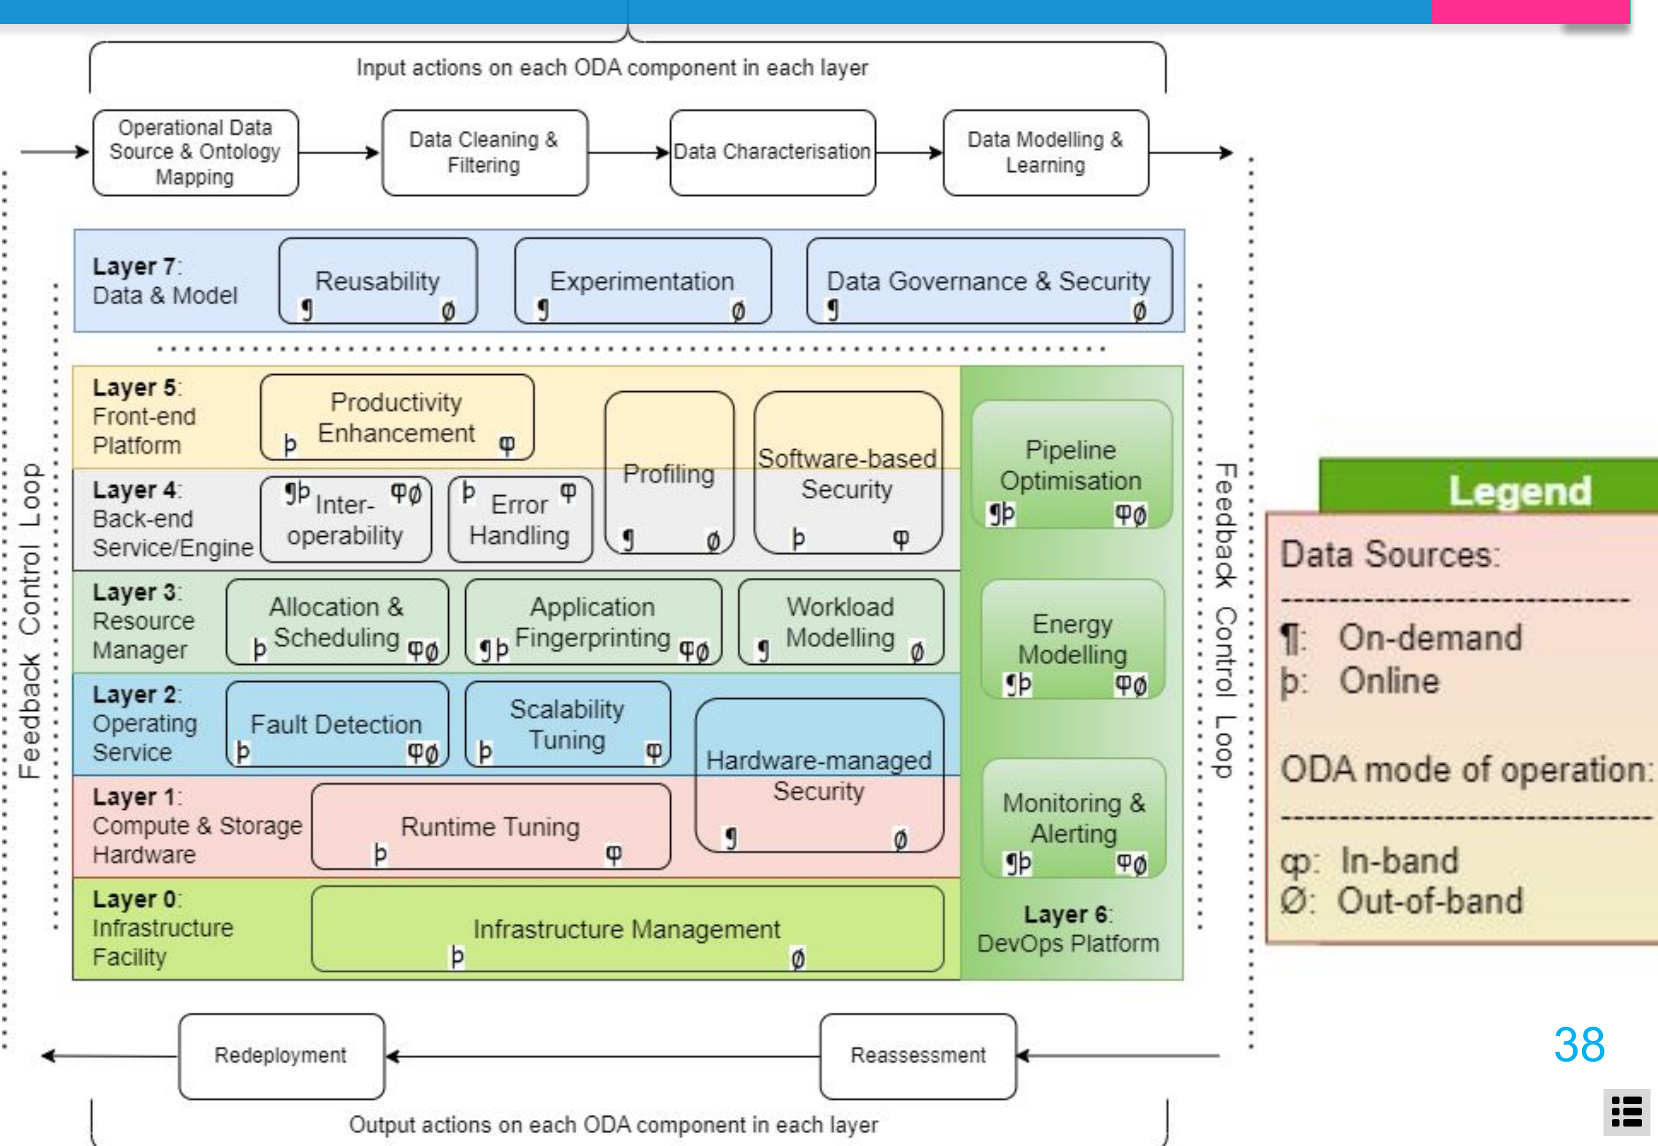

TUNABLE CONSISTENCY FOR ODA (OPERATIONALIZE)

Problem:
Too much ODA data.
Solution:

1. Define various levels of importance for data
2. Ensure various levels of consistency across levels
3. Intra-level, enable tunable consistency, e.g., Dyconits

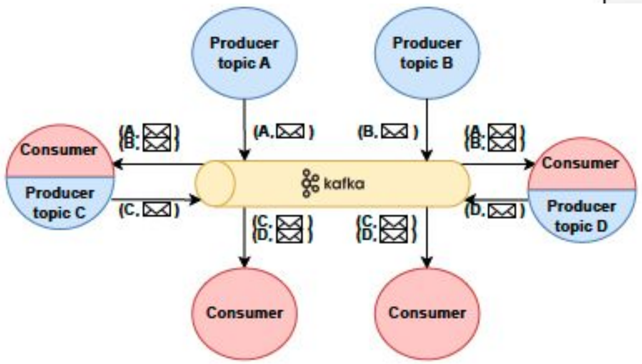

Multi-Hop Topology

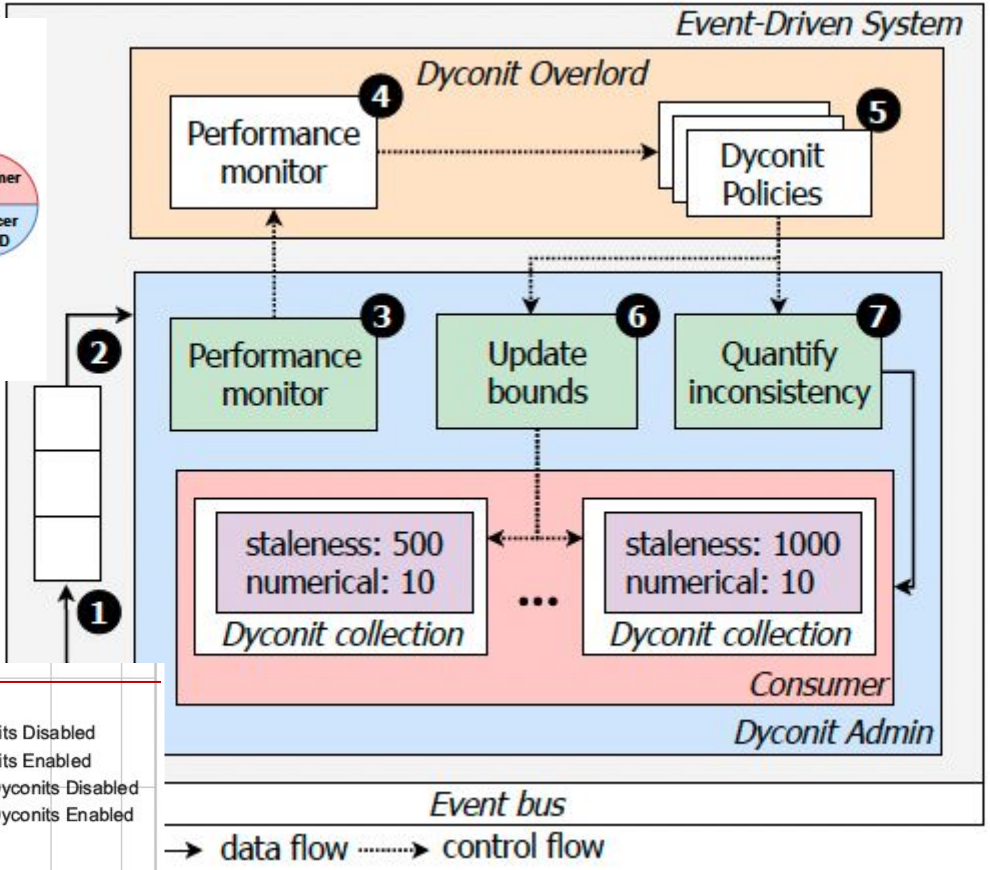

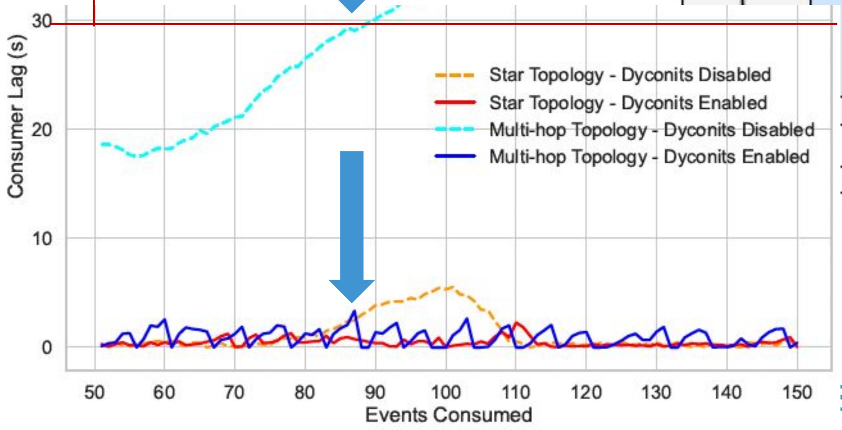

Ongoing work

Jurre Brandsen, Jesse Donkervliet, Iosup, et al. (2023) Hestia: General Dyconit Middleware + Donkervliet et al. (2021)

A NARROW ONTOLOGY FOR ODA (STANDARDIZE ‘D’)

Problem:
Different data formats, collection processes.

Solution:

1. Define a narrow ontology
2. Integrate into data science process
3. Implement for usability, e.g., graph, time-series, relational database

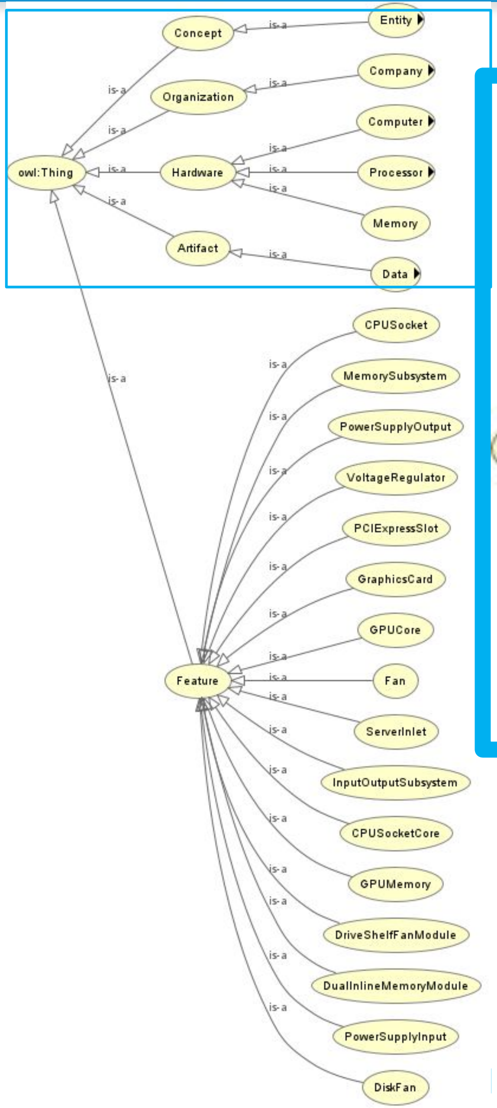

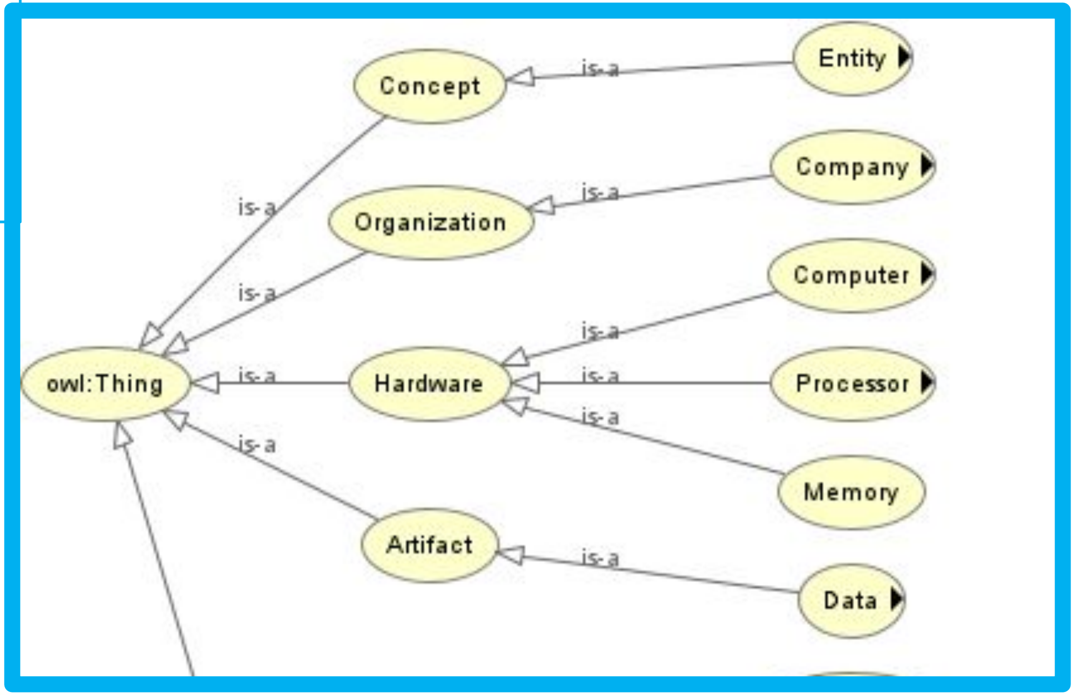

Ongoing work

Xiaoyu Chu, Shekhar Suman, Martin Molan, Andrea Bartolini (UniBo), Iosup, et al. (2023) Ontology for

FROM LOGS TO PROVENANCE (EXTEND 'D')

Problem:
Much happens
in-between and across
components.

Solution:

1. Drop-in, multi-granularity, multi-source provenance as core component
2. Implement for usability, e.g., debugging, tuning, auditing, reproducing

Ongoing work

Gilles Magalhaes, Tiziano De Matteis, Iosup, et al. (2023) Provenance in graph processing

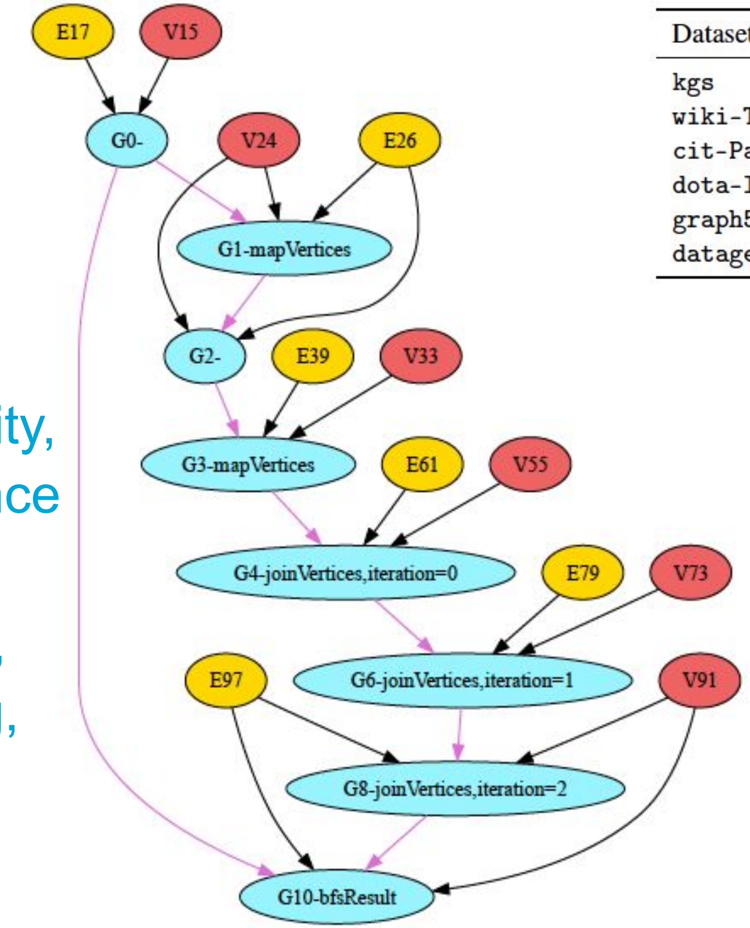

Dataset	Graphalytics scale	# nodes	# edges	Size (compressed)
kgs	XS	832,247	17,891,698	69 MB
wiki-Talk	XS	2,394,385	5,021,410	34.9 MB
cit-Patents	XS	3,774,768	16,518,947	119.1 MB
dota-league	S	61,170	50,870,313	114.3 MB
graph500-22	S	2,396,657	64,155,735	202.4 MB
datagen-7_8-zf	S	16,521,886	41,025,255	544.3 MB

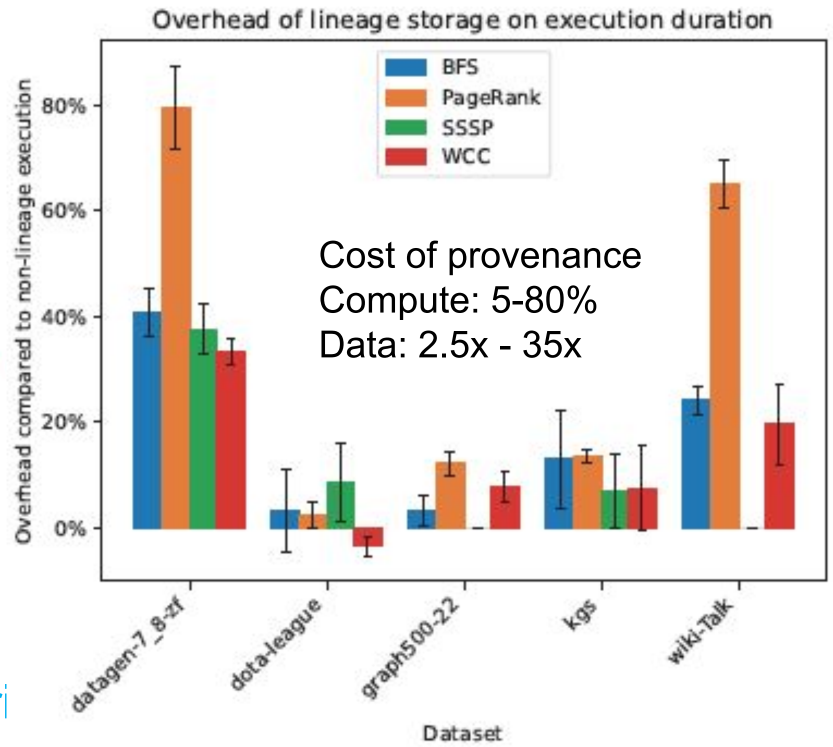

A DIGITAL TWIN FOR ODA (EXTEND 'A')

Problem:
Link between data and decision-making.

Solution:

1. Define most common types of analysis
2. Integrate into digital twinning process
3. Implement for usability and flexibility

Ongoing work

Iosup et al. (2017-2022) OpenDC 2.0, (2023) Graph-Greenifier

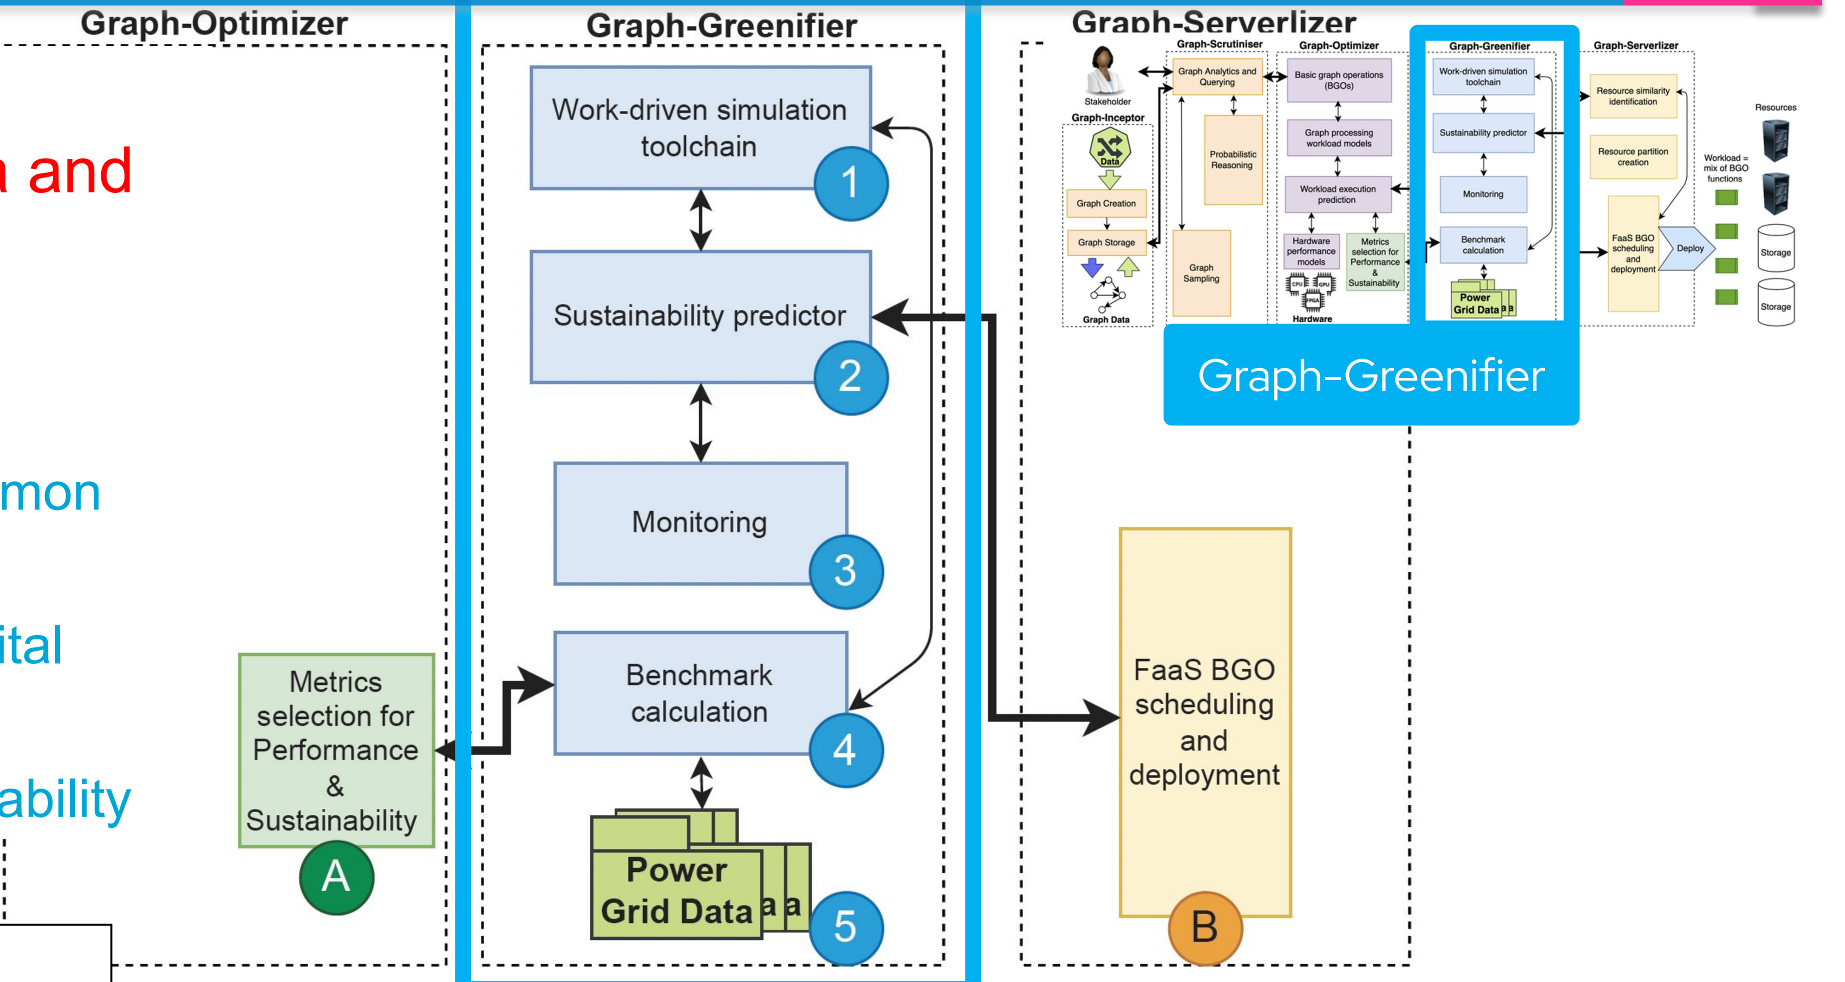

Graph-Greenifier → RM&S decisions, green labeling with data-driven explanation

MACHINE LEARNING OR LONG-TERM OBSERVATION (OR BOTH) WILL BE NEEDED (EXTEND 'A')

Problem: Complexity really means complexity.

Solution:

1. Express core components as processes
2. Observe and model continuously

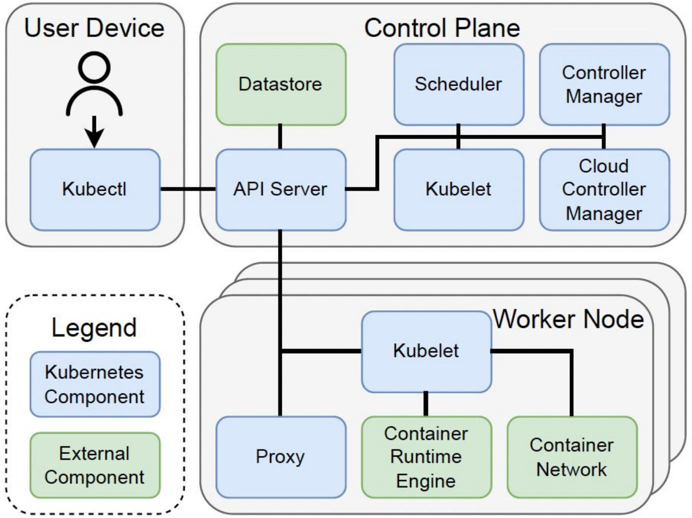

Component	File Resources	File Params	CLI Params
API server	39	149	152
Controller manager	37	147	132
Kubelet	31	206	133
Proxy	9	61	59
Scheduler	69	315	55
Other	50	189	0
Total	234	1,067	531

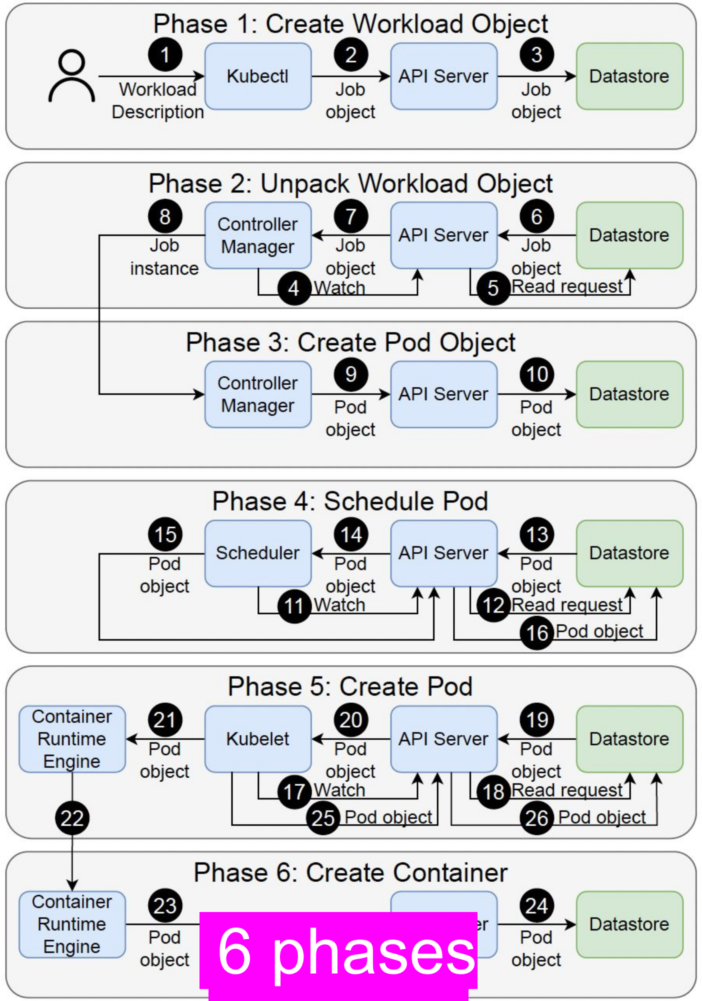

6 phases
26 steps

Ongoing work

Matthijs Jansen, Animesh Trivedi, Iosup et al. (2023) K8s exploration

EDUCATION

Integrate the ICT infrastructure memex, and more generally ODA, into our compsys coursework and education processes. It's fun, stimulates curiosity, leads to learning!

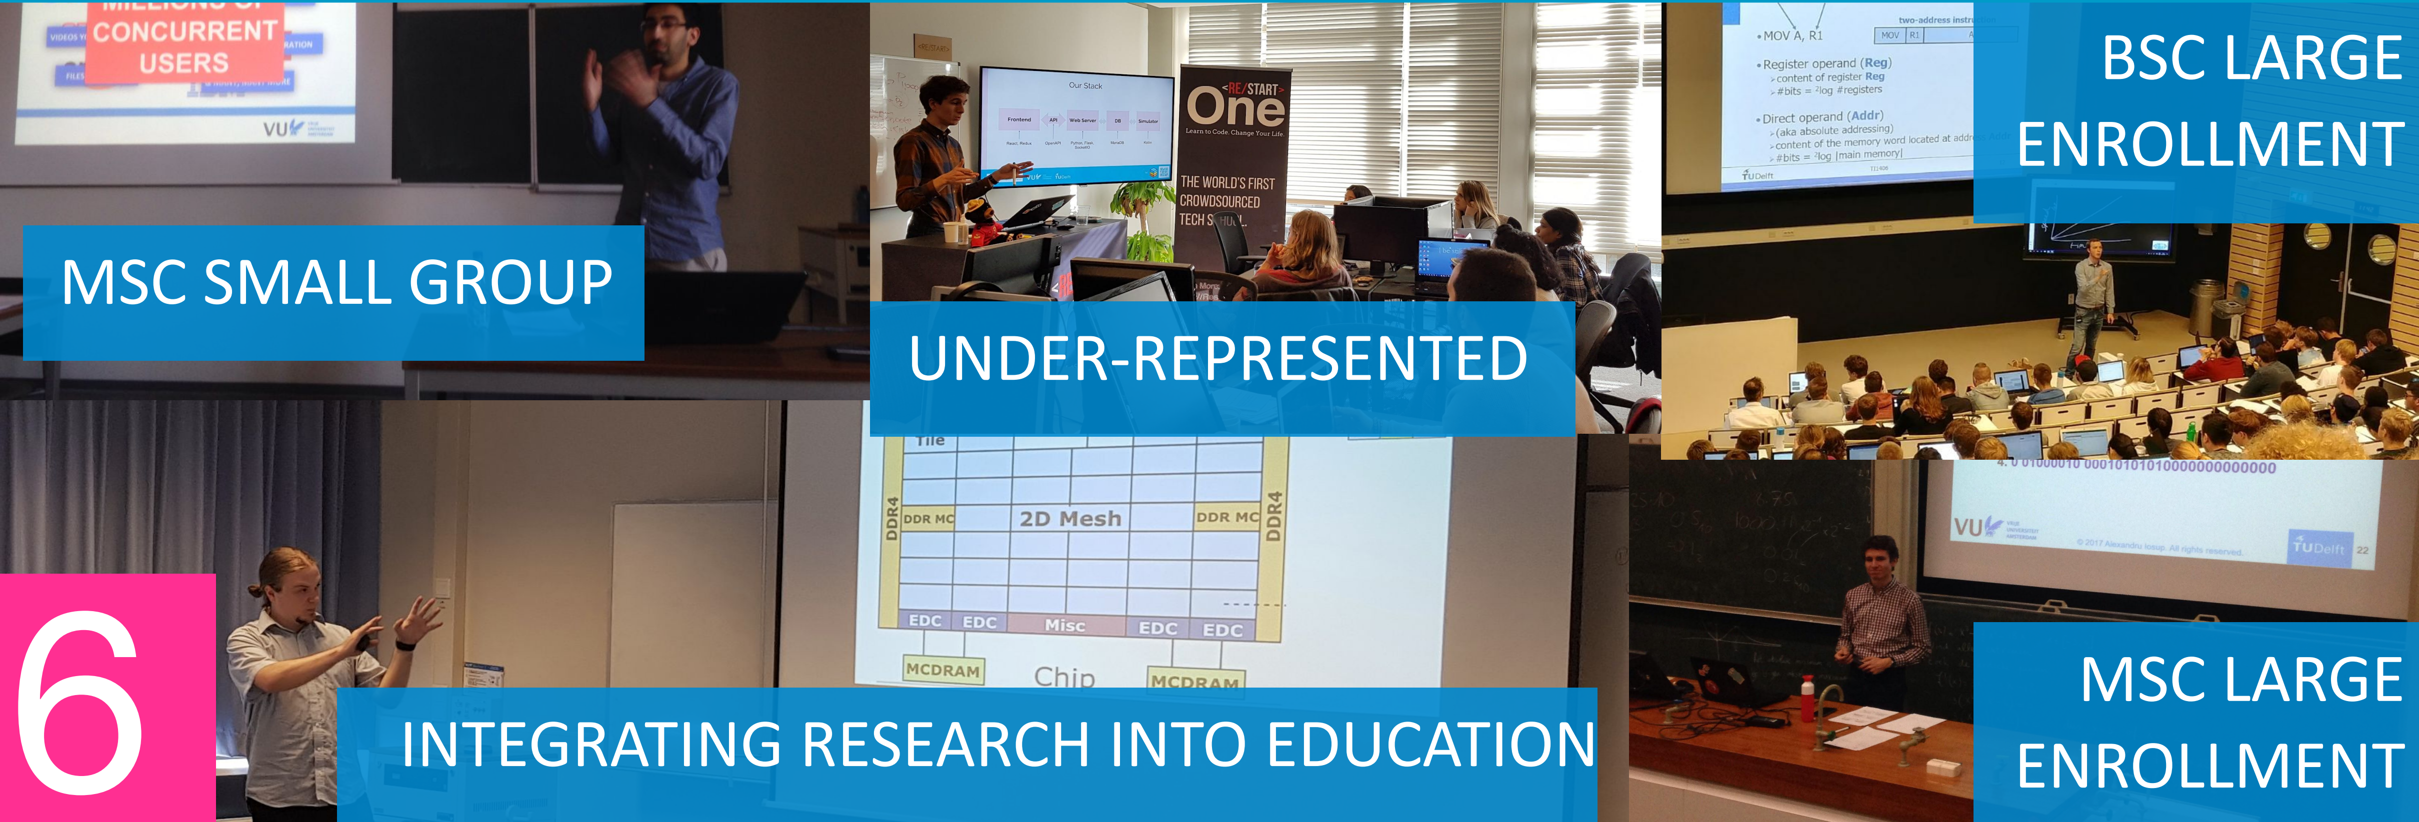

6

TAKE-HOME

We're building the ICT Infrastructure Memex to enable Operational Data Analytics into the 21st century. We have many (developing) theories and practical results. We know others are working on this topic and seek discussion and an active collaboration. To discuss:

- Theory: What are the core components of ODA? How do they relate - what is a good reference architecture for ODA?
- Theory: Are performance and availability just parts of a conceptual continuum? What else is in there and how to include it in ODA?
- Theory: We say: Just touch it with your lower lip, briefly, and move away if too hot. How to reason about energy use, both short- and long-term?
- Practice: Ontology building for sharing traces collected in ICT infrastructure. Do we need ontologies? What is a good middle ground between implicit ontologies and exhaustive ontologies?
- Practice: What kinds of ODA benefit from more complex types of analysis, e.g., graph analytics and learning?
- Ethics: If a tree falls in front of you, do you have to see it? From plausible deniability to responsibility in ICT infrastructure management.

bit.ly/AIScalPerf23

WANT TO READ MORE ON THE TOPIC?

MASSIVIZING COMPUTER SYSTEMS

FURTHER READING

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

1. Alexandru Iosup, Fernando Kuipers, Ana Lucia Varbanescu, Paola Grosso, Animesh Trivedi, Jan S. Rellermeyer, Lin Wang, Alexandru Uta, Francesco Regazzoni (2022) Future Computer Systems and Networking Research in the Netherlands: A Manifesto. CoRR abs/2206.03259.
2. Kounev, Herbst, Abad, Iosup, et al. (2023) Serverless Computing: What It Is, and What It Is Not? CACM 66(9).
3. Jansen et al. (2023) The SPEC-RG Reference Architecture for The Compute Continuum. CCGRID.
4. Versluis et al. (2023) Less is not more: We need rich datasets to explore. FGCS 142.
5. Crusoe et al. (2022) Methods included: standardizing computational reuse and portability with the Common Workflow Language. CACM 65(6).
6. Andreadis et al. (2022) Capelin: Data-Driven Capacity Procurement for Cloud Datacenters using Portfolios of Scenarios. IEEE TPDS 33(1).
7. Eismann et al. (2022) The State of Serverless Applications: Collection, Characterization, and Community Consensus. IEEE Trans. Software Eng. 48(10).
8. Radu Prodan, Dragi Kimovski, Andrea Bartolini, Michael Cochez, Alexandru Iosup, Evgeny Kharlamov, Jože Rožanec, Laurenţiu Vasiliu, Ana Lucia Vărbănescu (2022) Towards Extreme and Sustainable Graph Processing for Urgent Societal Challenges in Europe. IEEE Cloud Summit.
9. Sakr, Bonifati, Voigt, Iosup, et al. (2021) The future is big graphs: a community view on graph processing systems. Commun. ACM 64(9).
10. Mastenbroek et al. (2021) OpenDC 2.0: Convenient Modeling and Simulation of Emerging Technologies in Cloud Datacenters. CCGRID.
11. Talluri et al. (2021) Empirical Characterization of User Reports about Cloud Failures. ACSOS.
12. Versluis and Iosup (2021) A survey of domains in workflow scheduling in computing infrastructures: Community and keyword analysis, emerging trends, and taxonomies. FGCS.
13. Papadopoulos et al. (2021) Methodological Principles for Reproducible Performance Evaluation in Cloud Computing. IEEE Trans. Software Eng. 47(8).
14. Donkervliet et al. (2021) Dyconits: Scaling Minecraft-like Services through Dynamically Managed Inconsistency. ICDCS.
15. Versluis et al. (2020) The Workflow Trace Archive: Open-Access Data From Public and Private Computing Infrastructures. IEEE TPDS 31(9).