Elasticity in Graph Analytics? A Benchmarking Framework for Elastic Graph Processing

Alexandru Uta, Sietse Au, Alexey Ilyushkin, Alexandru Iosup

a.uta@vu.nl

Massivizing Computer Systems

Vrije Universiteit Amsterdam

Elasticity in Computer Systems

- Mechanism to adjust the resources based on a utilization metric
- Supply (machines) vs. demand (workload)
- Increase or decrease no. of machines

Metrics:

- CPU utilization
- Number of connections
- Response time

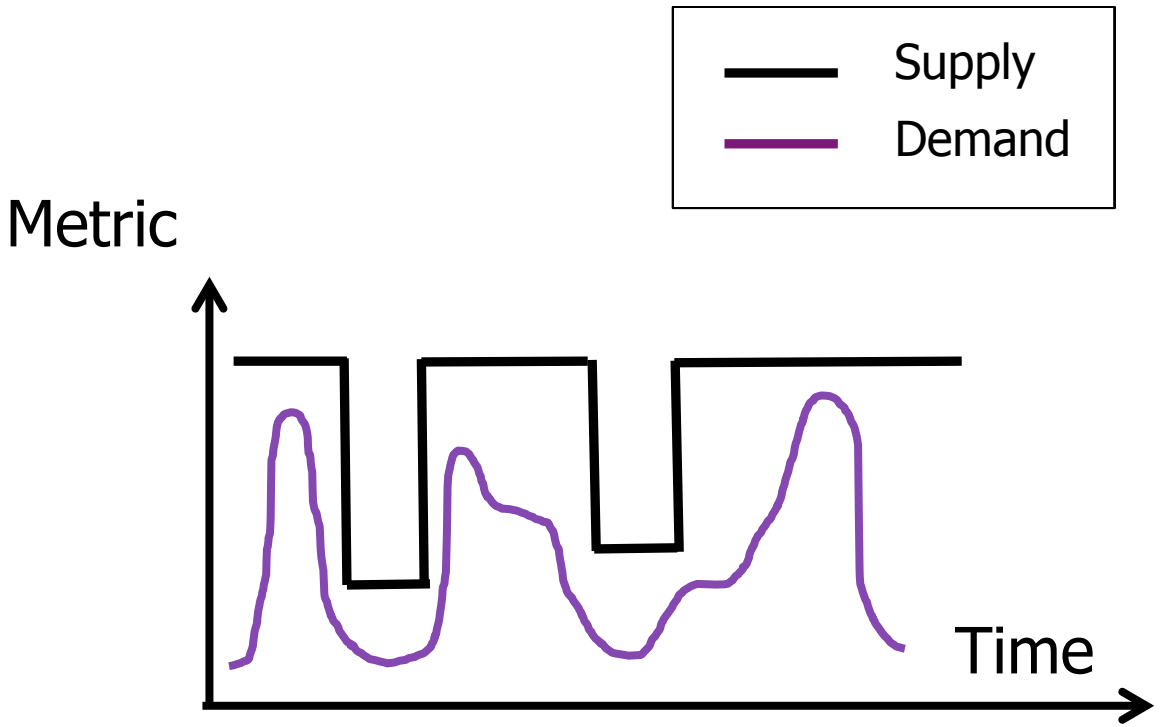

Distributed Graph Processing - Pregel

Bulk-synchronous parallel (BSP) model

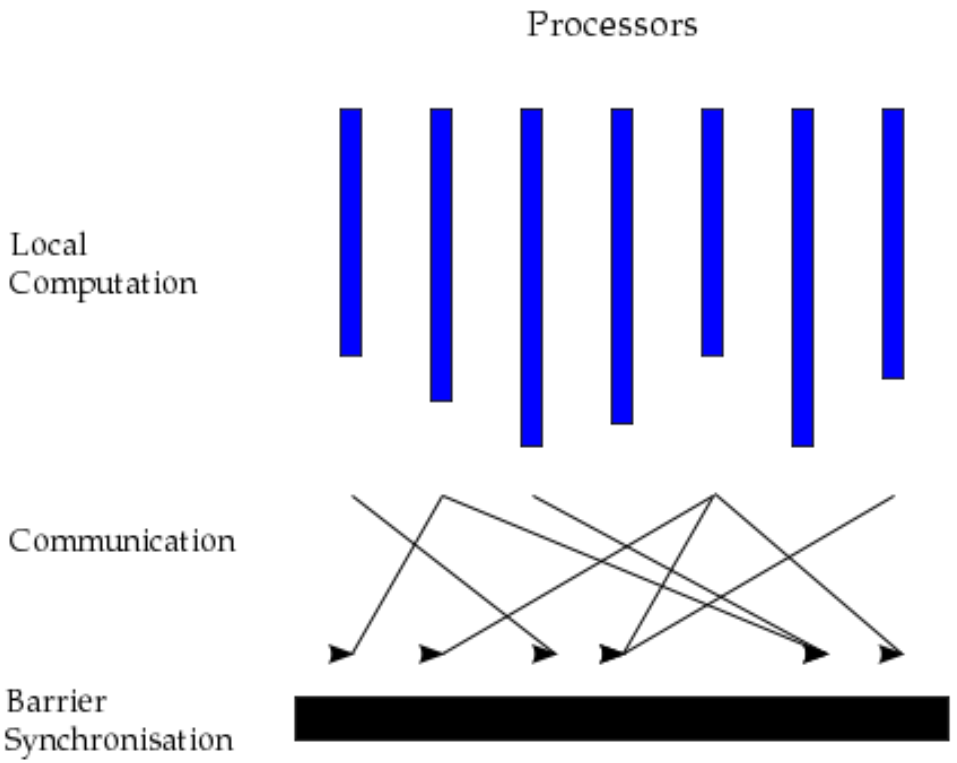

Pregel

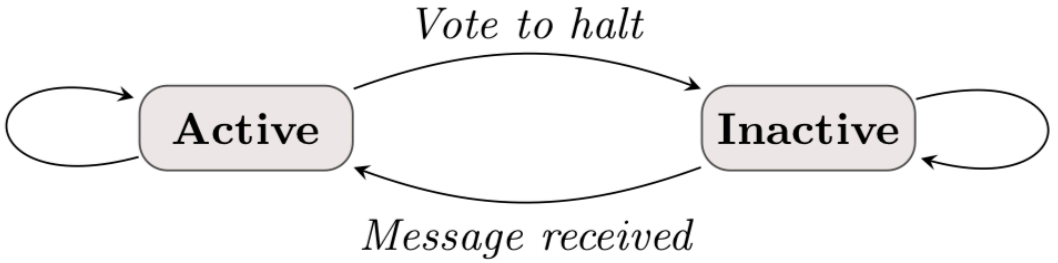

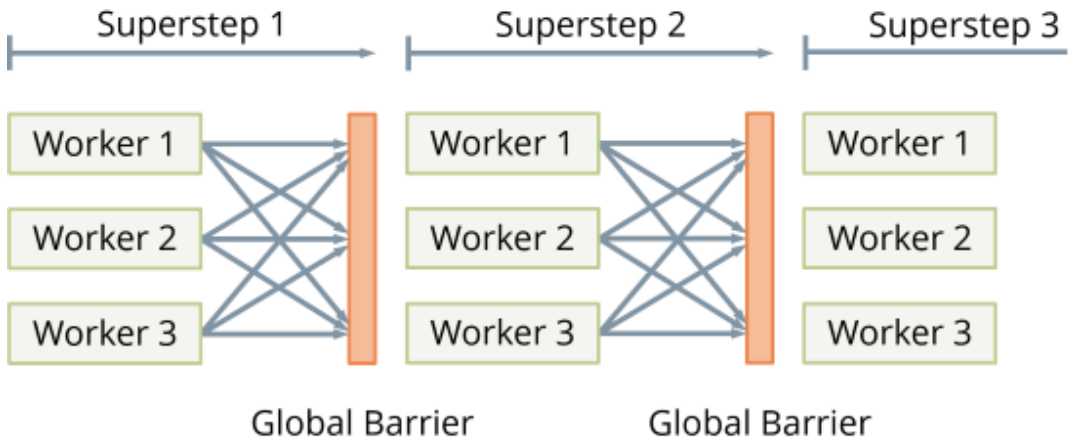

Elasticity in Graph Analytics

Metric

- How to define demand?

How to express elasticity in Graph Analytics?

Time

Variability in Graph Analytics

- Demand = active vertices
- Supply = no. of machines
- Algorithm BFS

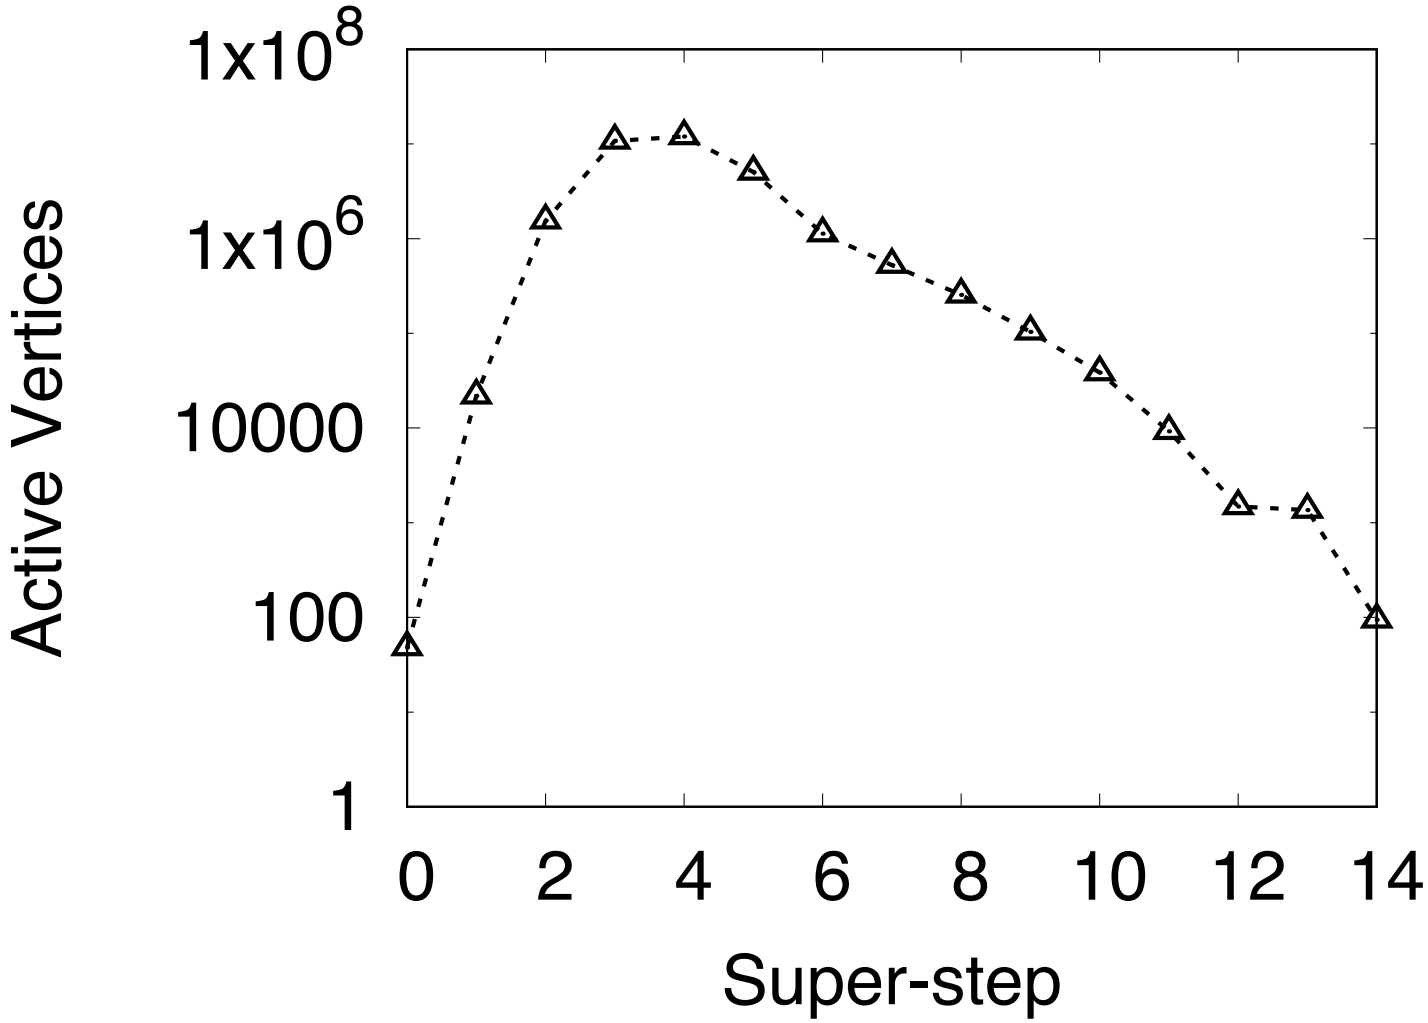

Active Vertices – Dependent on Algorithm

- Demand = active vertices

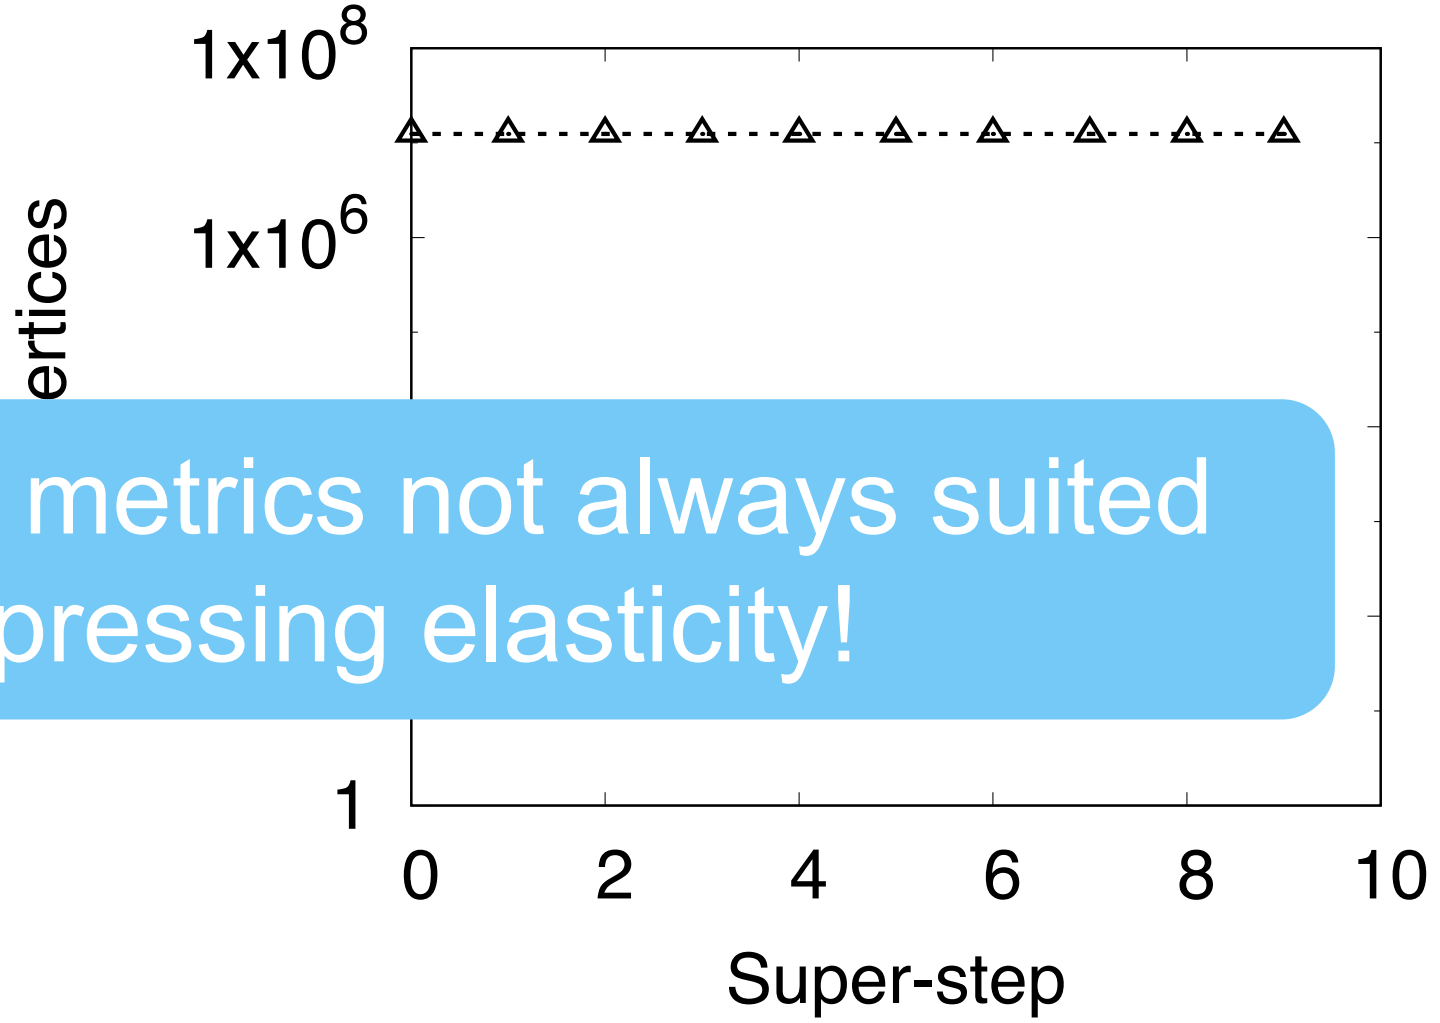

Graph-related metrics not always suited for expressing elasticity!

Elasticity using Systems-level Metrics – CPU Load

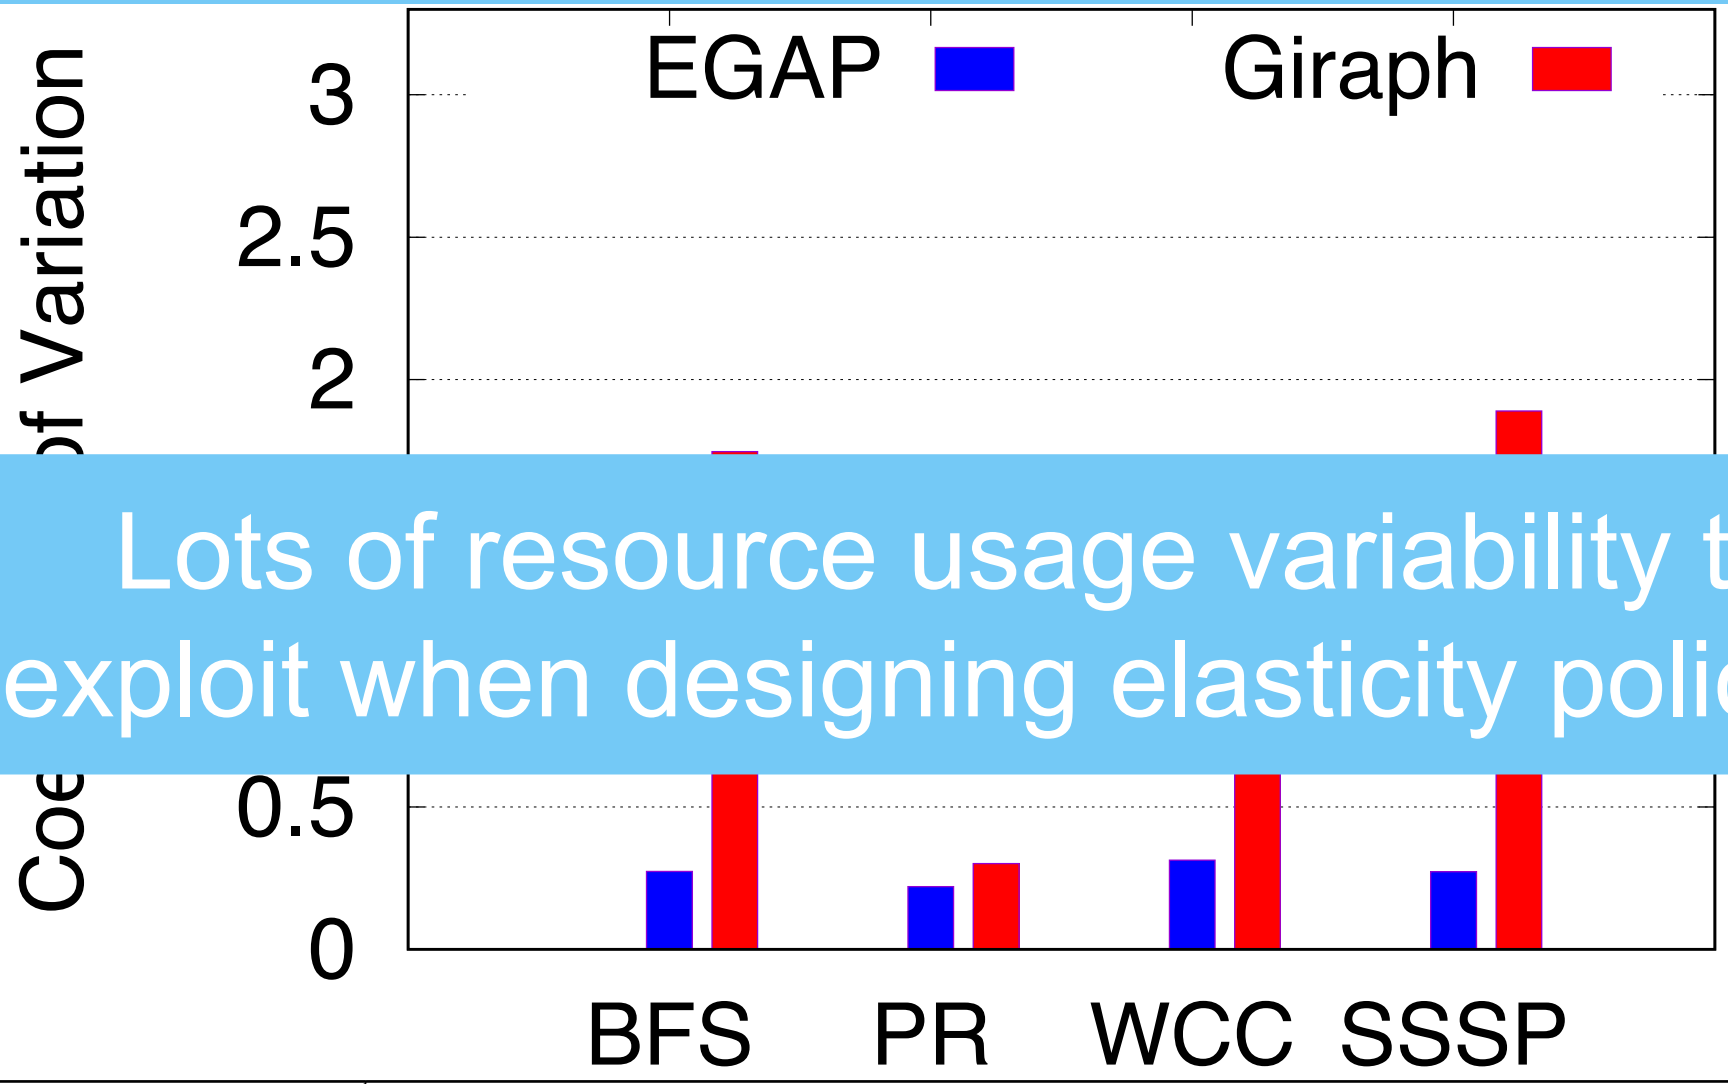

Lots of resource usage variability to exploit when designing elasticity policies!

- CoV computed per worker per superstep

Benchmarking Elasticity

Q: How to qualitatively assess elasticity in graph analytics?

Benchmarking Elasticity

- How to qualitatively assess elasticity in graph analytics?

(1) Workloads — LDBC Graphalytics

Benchmarking Elasticity

- How to qualitatively assess elasticity in graph analytics?

(1) Workloads — LDBC Graphalytics

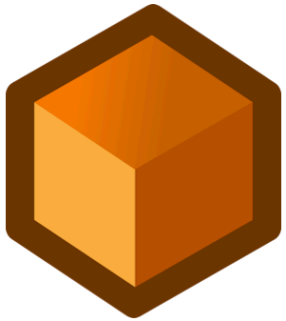

Graphalytics

Open-source Graph Processing Benchmark Suite

- State-of-the-art [1]
- Variety of algorithms and datasets
- Renewal process
- 6 Algorithms
  - BFS, PR, WCC, LCC, CDLP, SSSP
- Many datasets (small to large)
  - Scale-free vs. non-scale-free

[1] Iosup et al., VLDB '16

Benchmarking Elasticity

- How to qualitatively assess elasticity in graph analytics?

(1) Workloads	LDBC Graphalytics
(2) Metrics	Accuracy, Wrong Provisioning, Instability

- **Defined by SPEC Cloud RG**
- **How "good" is the autoscaling policy?**
- **Accuracy**
  - Over- and under-provisioning
- **Wrong-provisioning**
  - How much time spent over- or under-provisioned?
- **Instability**
  - How fast/slow does the system adapt to workload change?

Benchmarking Elasticity

- How to qualitatively assess elasticity in graph analytics?

(1) Workloads	LDBC Graphalytics
(2) Metrics	Accuracy, Wrong Provisioning, Instability
(3) Autoscalers	Generic, System-level

- **Autoscaling policies**
- **Generic policies**
  - SPEC RG Cloud
  - Proven to work for a wide variety of workloads
  - history-, regression-based etc.
- **System-level autoscaling policies**
  - CPU
  - Wallclock time
  - Network traffic

Benchmarking Elasticity

- How to qualitatively assess elasticity in graph analytics?

(1) Workloads — LDBC Graphalytics

(2) Metrics — Accuracy, Wrong Provisioning, Instability

(3) Autoscalers — Generic, System-level

---

(4) Elastic Graph Analytics Platform

- **EGAP = new operational model**
- JoyGraph reference implementation EGAP
- No other feasible alternative
- Pregel-based Scala implementation
- 11 kLoC
- 4 person-months

JoyGraph EGAP

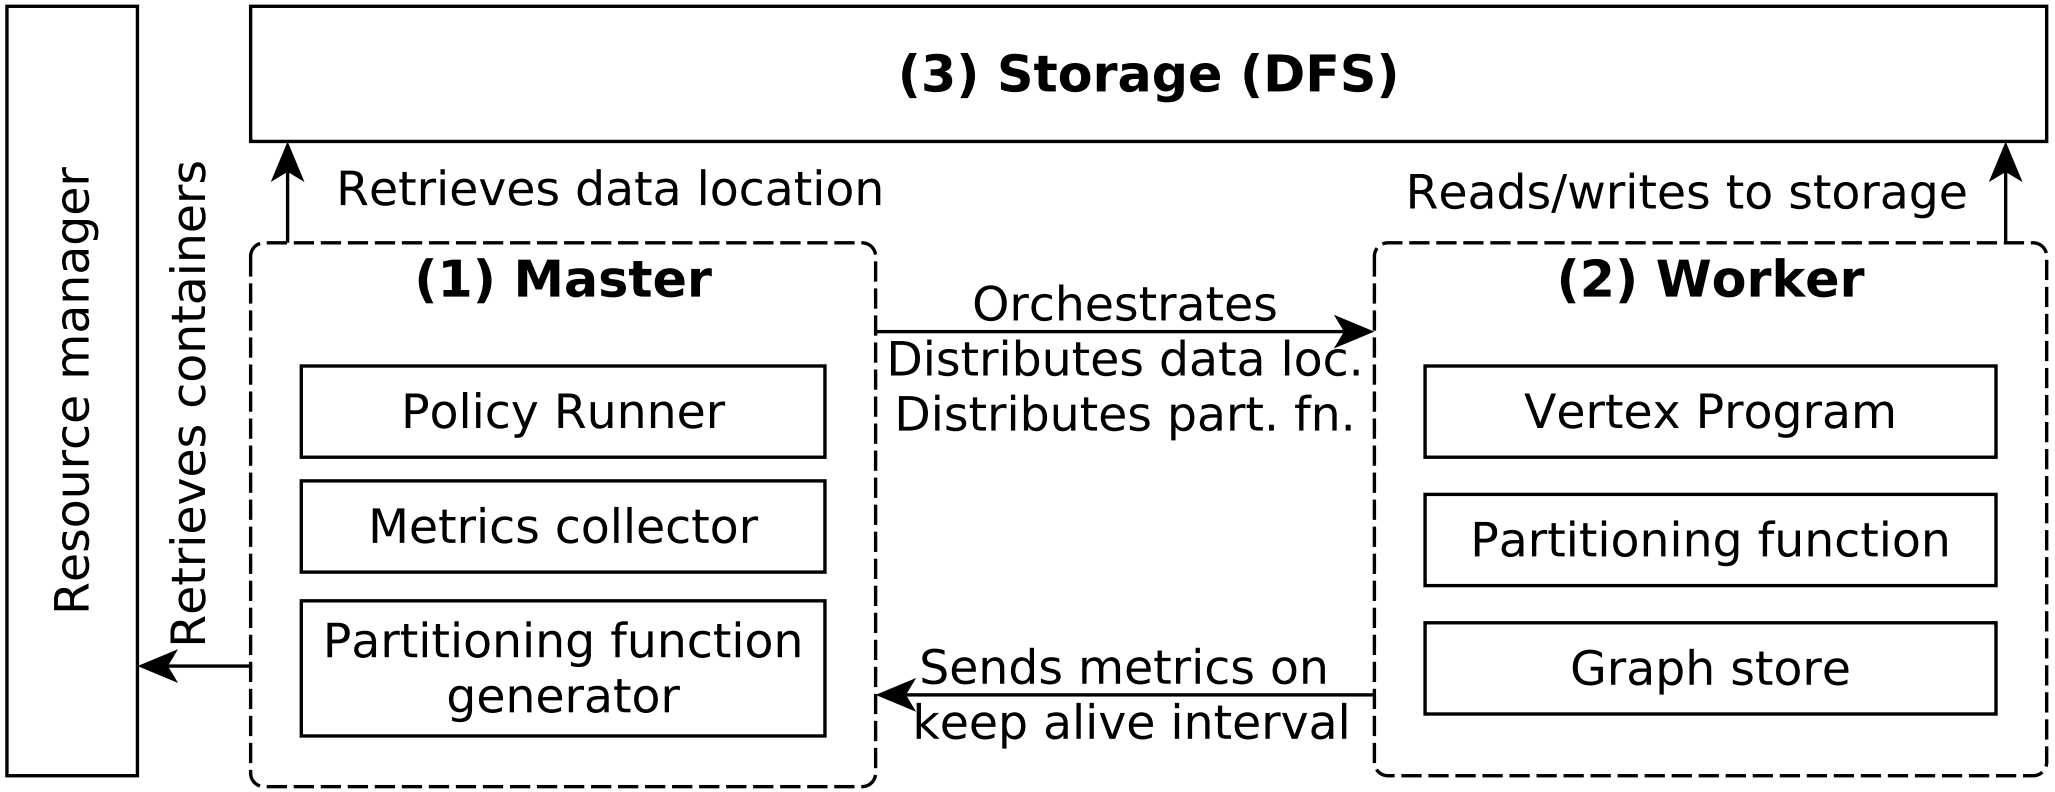

Experiment Setup

- **Q1: How does JoyGraph EGAP compare to state-of-the-art systems?**
- **Q2: What does elasticity mean for graph analytics?**
- Realistic big data setup:
  - Own cluster: 16-core, 64GB RAM, Infiniband network – machines
  - Proven workload: Graphalytics datasets and algorithms
  - Proven metrics: SPEC RG Cloud metrics
  - Proven and new system-level autoscalers

JoyGraph EGAP vs. State-of-Art

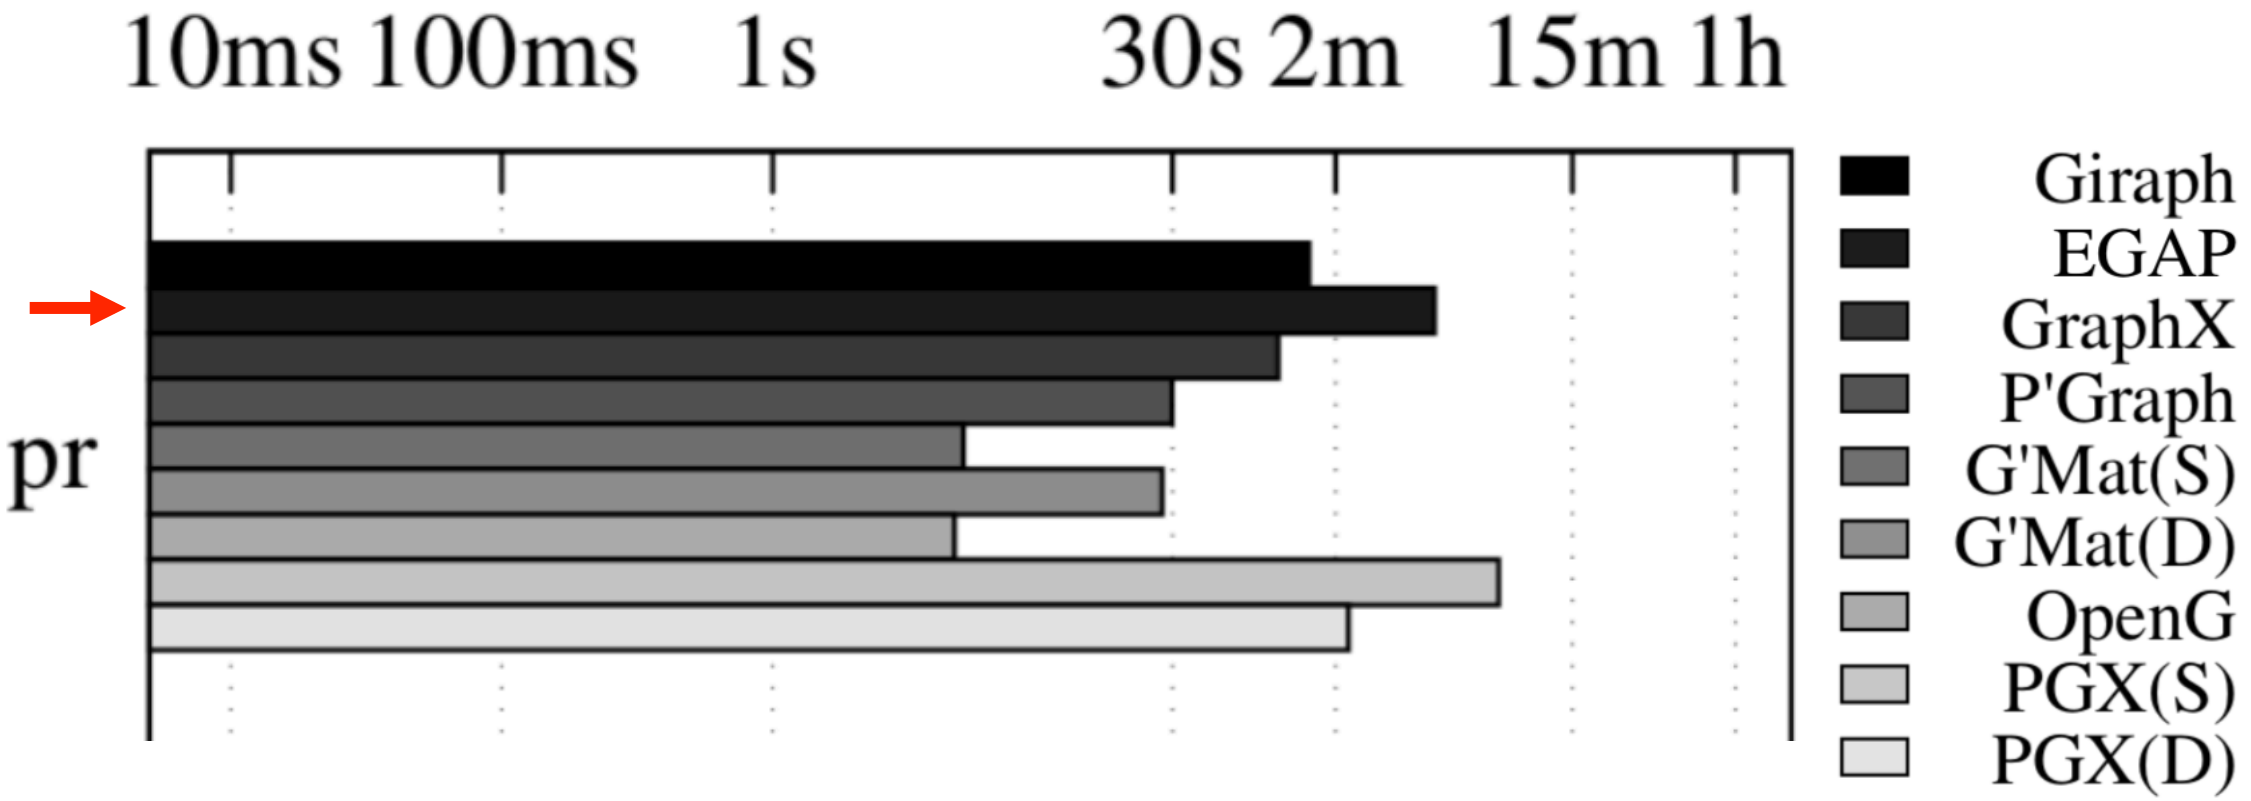

JoyGraph EGAP vs. State-of-Art

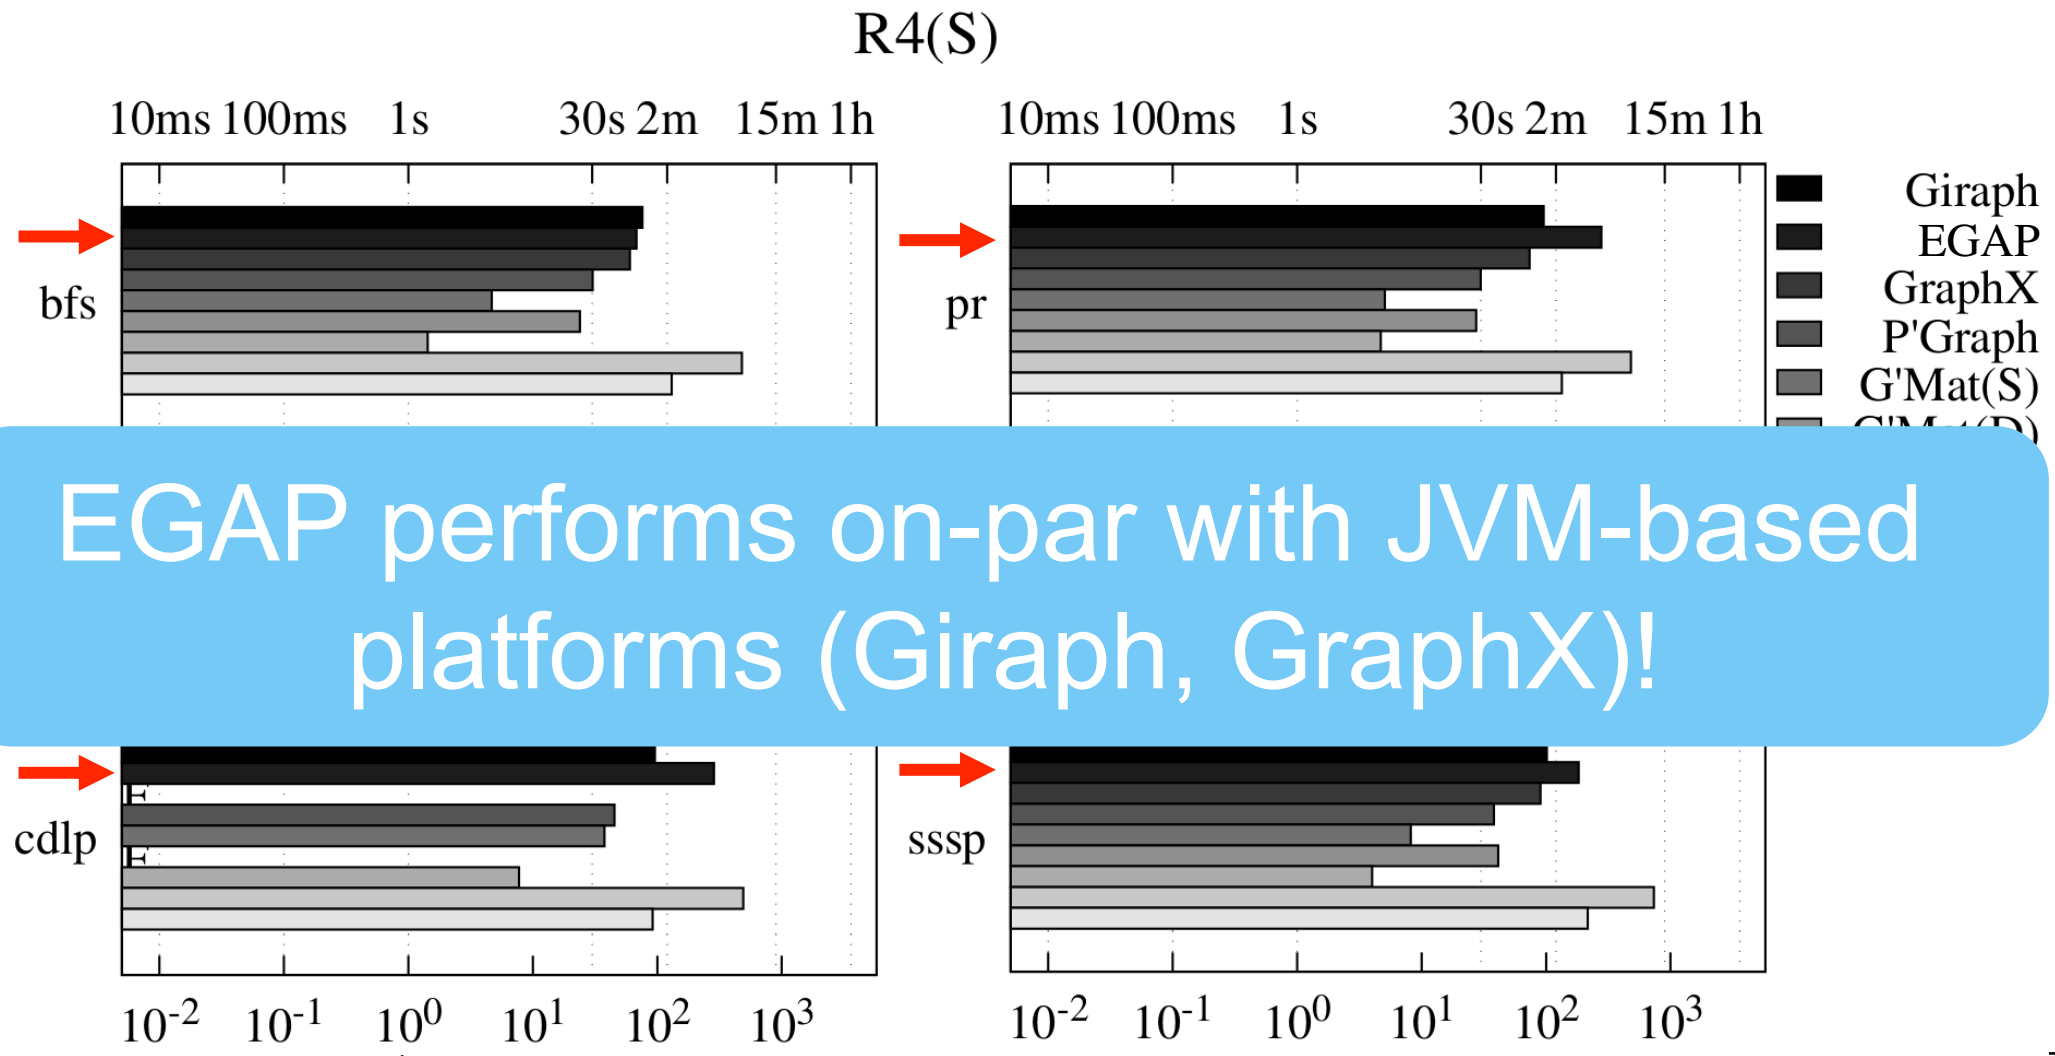

EGAP performs on-par with JVM-based platforms (Giraph, GraphX)!

Elasticity Metrics

AUTOSCALER METRICS FOR GRAPH500-25 AND ALGORITHM BFS.

	a_U	a_O	$\bar{a}_U$	$\bar{a}_O$	t_U	t_O	k	k'
React	0.11	1.27	0.00	0.82	0.12	21.19	36.98	58.95
AKTE	0.06	1.01	0.00	20.25	0.28	6.27	60.72	38.18
ConPaaS	0.07	2.02	0.00	38.52	0.07	12.79	34.94	62.59
Reg	0.12	2.71	0.00	54.21	0.13	16.93	44.82	51.18
Hist	0.13	2.91	0.00	58.18	0.14	19.39	52.36	43.61
NP	0.22	0.02	0.00	0.00	4.40	0.36	24.67	18.21
CPU	1.16	0.00	0.00	0.00	23.23	0.00	49.64	23.25
WCP	0.97	0.55	0.00	0.00	19.47	10.90	59.35	39.53

Elasticity Metrics – System vs. Generic Policies

AUTOSCALER METRICS FOR GRAPH500-25 AND ALGORITHM BFS.

	a_U	a_O	$\bar{a}_U$	$\bar{a}_O$	t_U	t_O	k	k'
React			0.00	0.82	0.12	21.19	36.98	58.95
AKTE			0.00	20.25	0.28	6.27	60.72	38.18
ConPaaS			0.00	38.52	0.07	12.79	34.94	62.59
Reg			0.00	54.21	0.13	16.93	44.82	51.18
Hist			0.00	58.18	0.14	19.39	52.36	43.61
NP			0.00	0.00	4.40	0.36	24.67	18.21
CPU			0.00	0.00	23.23	0.00	49.64	23.25
WCP			0.00	0.00	19.47	10.90	59.35	39.53

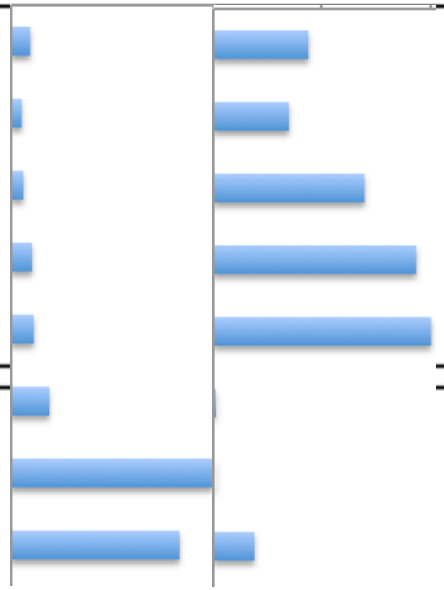

Generic = over-provision, system = under-provision

Elasticity Metrics – System vs. Generic Policies

AUTOSCALER METRICS FOR GRAPH500-25 AND ALGORITHM BFS.

	a_U	a_O	$\bar{a}_U$	$\bar{a}_O$	t_U	t_O	k	k'
React			0.00	0.82	0.12	21.19		
AKTE			0.00	20.25	0.28	6.27		
ConPaaS			0.00	38.52	0.07	12.79		
Reg			0.00	54.21	0.13	16.93		
Hist			0.00	58.18	0.14	19.39		
NP			0.00	0.00	4.40	0.36		
CPU			0.00	0.00	23.23	0.00		
WCP			0.00	0.00	19.47	10.90		

Both types of austoscalers are unstable.

Elastic Scaling Behavior - PageRank

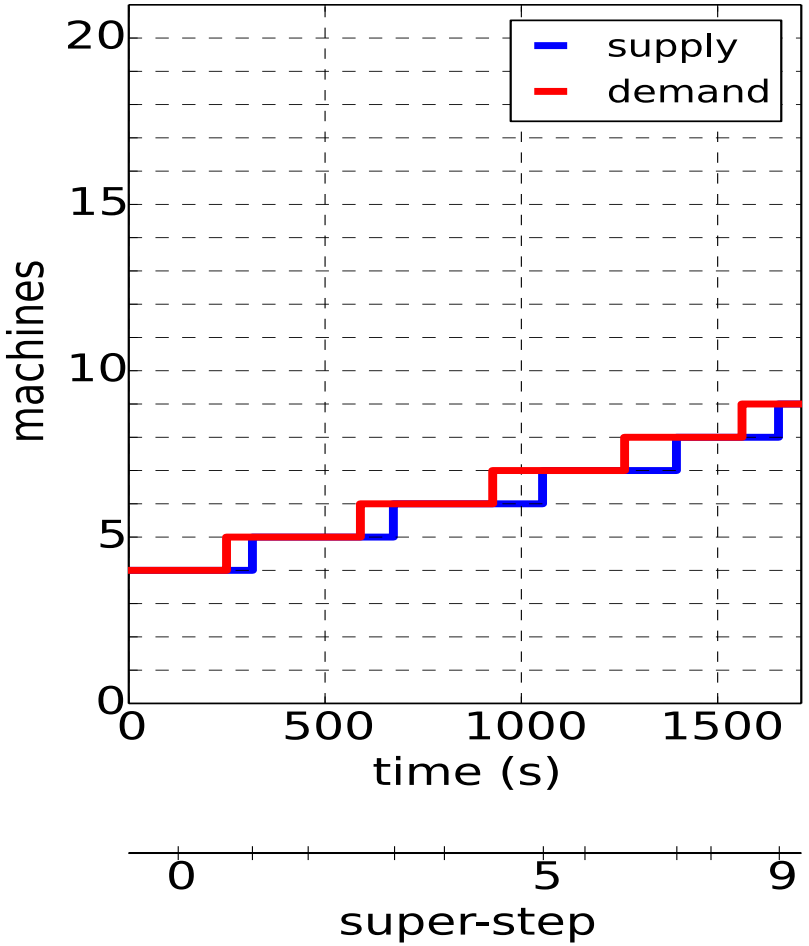

(a) WCP policy

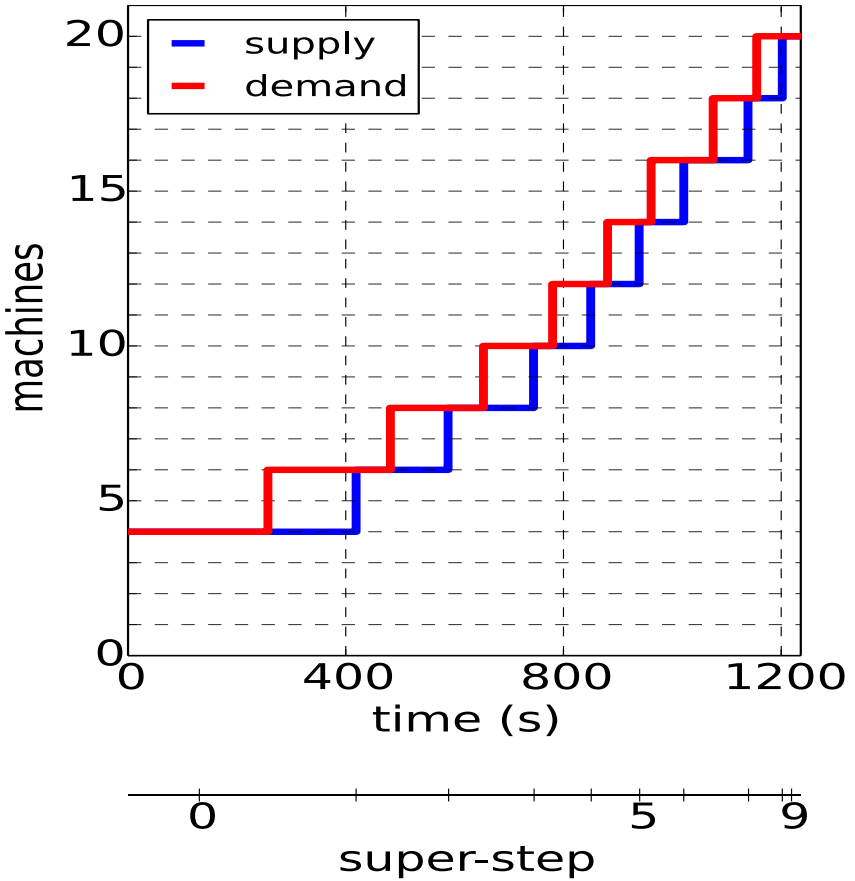

(b) AKTE Policy

Elasticity Overhead in Graph Analytics

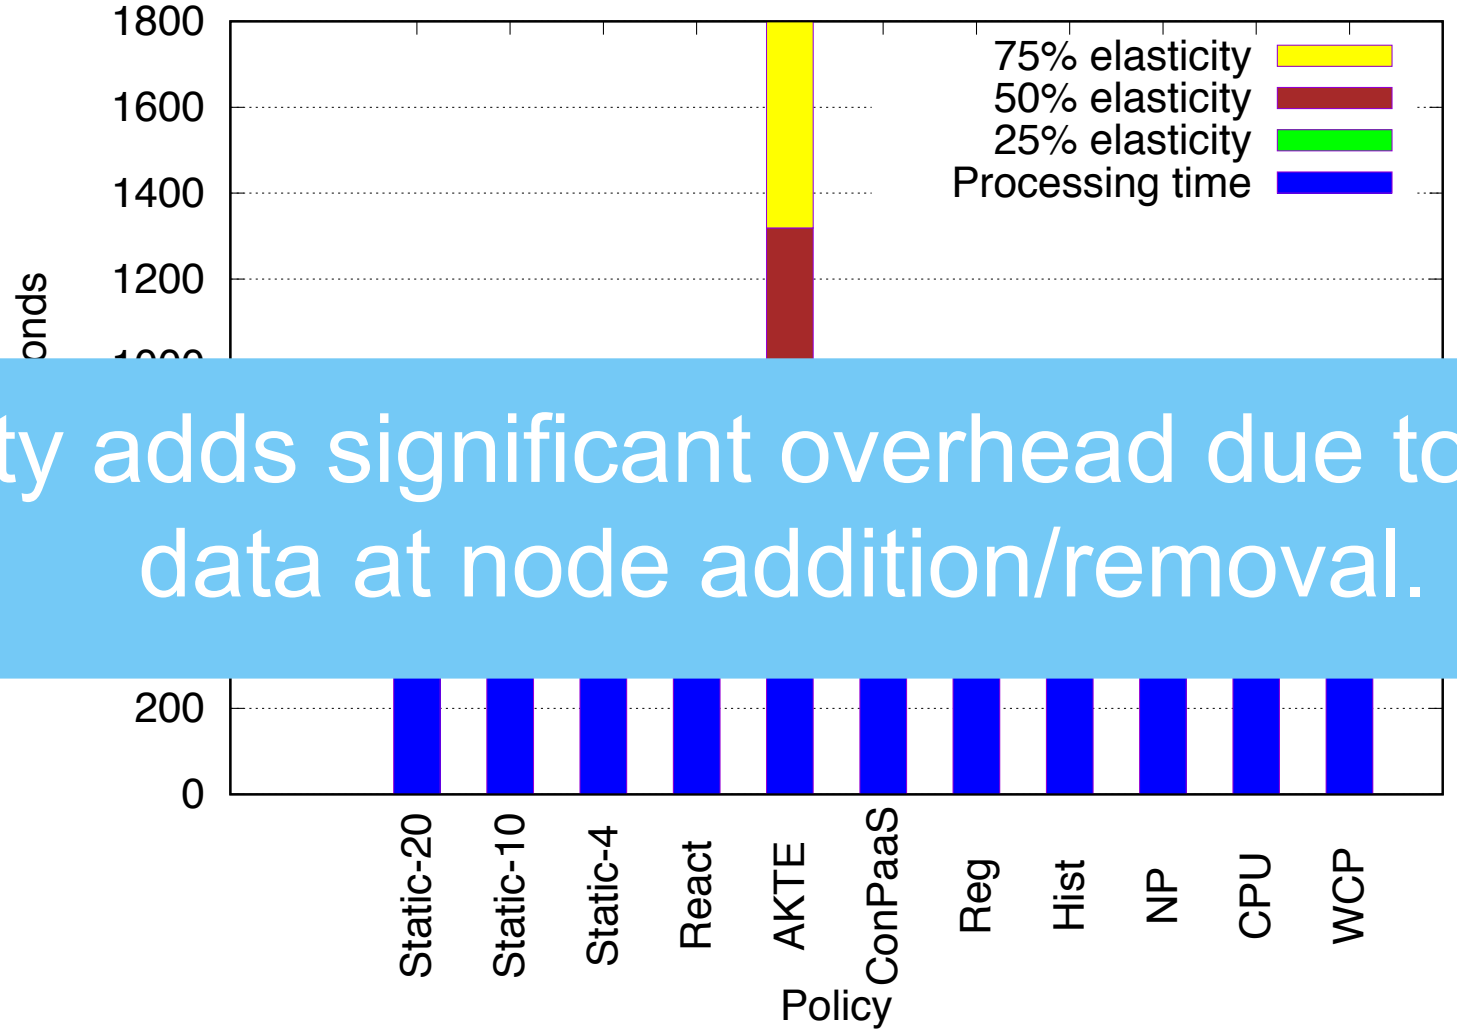

Elasticity adds significant overhead due to shuffling data at node addition/removal.

Take-home Message

- Graph analytics exhibits high variability, of kinds unexplored until now
- EGAP + 1st benchmark for assessing elasticity in graph analytics
- No "best" elastic scaling policy: performance ~ algorithm, dataset
- Elasticity adds significant communication overhead
- Future work: investigate better communication schemes, new models of elasticity in clouds (e.g., serverless/FaaS)

Alexandru Uta
a.uta@vu.nl