The Cost of Simplicity: Understanding Datacenter Scheduler Programming Abstractions

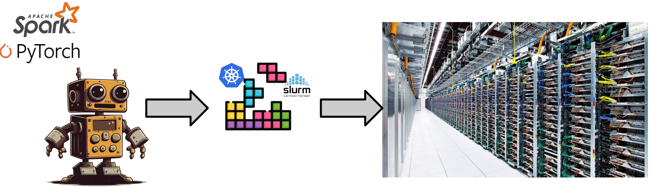

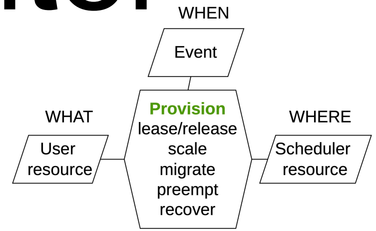

Aratz Manterola Lasa[1], **Sacheendra Talluri**[2], Tiziano De Matteis[2], Alexandru Iosup[2]

[1]WarpStream Labs
[2]Vrije Universiteit Amsterdam

@Large Research

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Why APIs/abstractions?

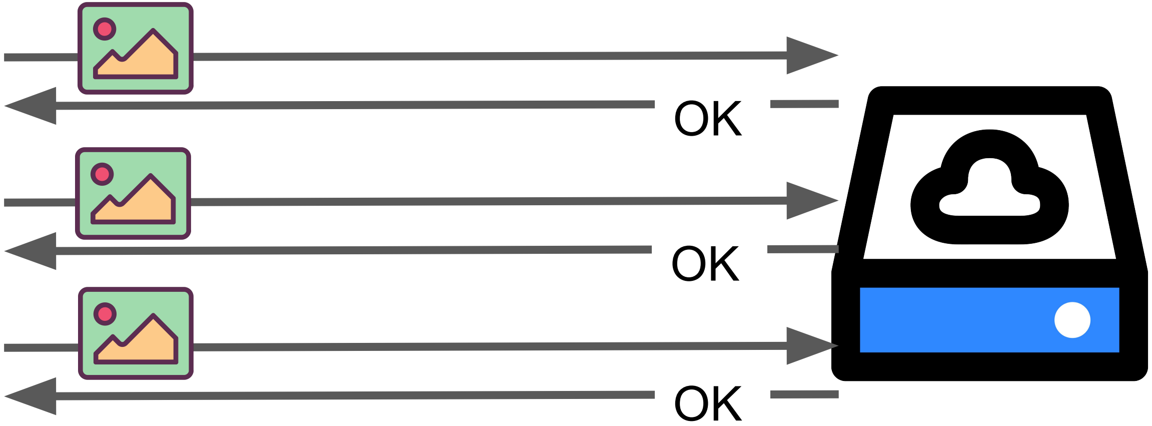

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Why APIs/abstractions?

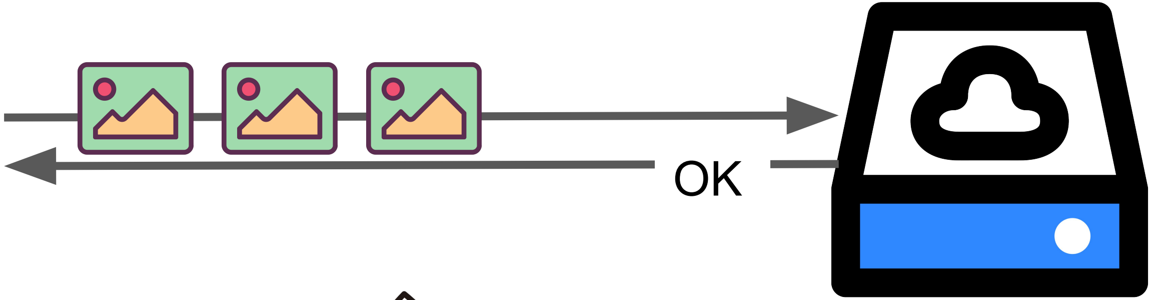

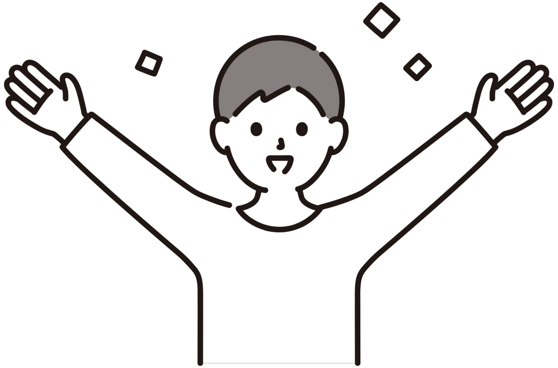

Less work for the user!

Better resource management by the provider

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Datacenter System Model

Workload

Bag of Tasks

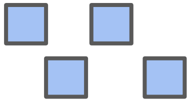

Workflow

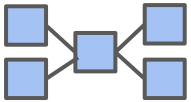

Resource Manager

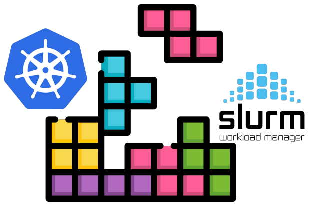

Datacenter / Cluster / Virtual Cluster

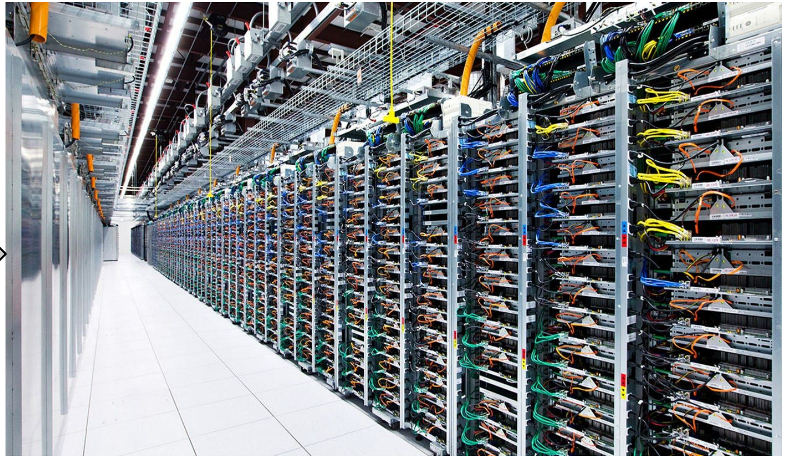

© Amazon

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Runtimes Can Use Complex Resource Management APIs

APACHE Spark

PyTorch

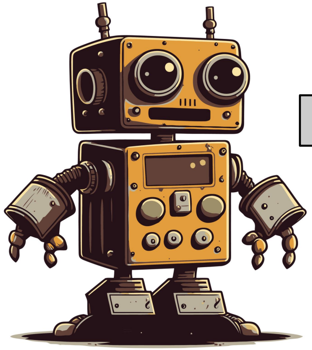

CC Mehedi Hasan

Resource Manager

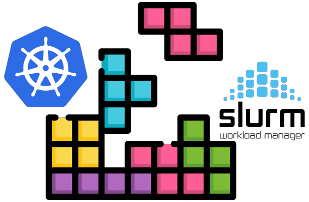

Datacenter / Cluster / Virtual Cluster

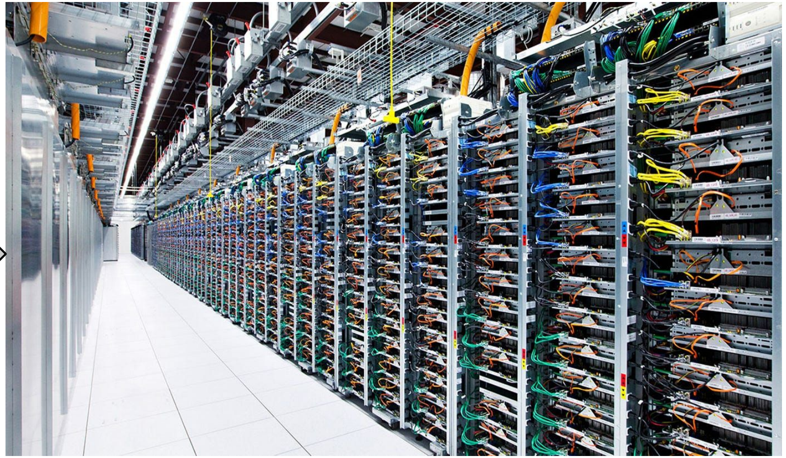

© Amazon

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Runtimes Can Use Complex Resource Management APIs

Apache Spark

PyTorch

CC Mehedi

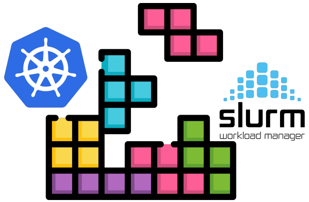

Datacenter /
Cluster /
Virtual Cluster

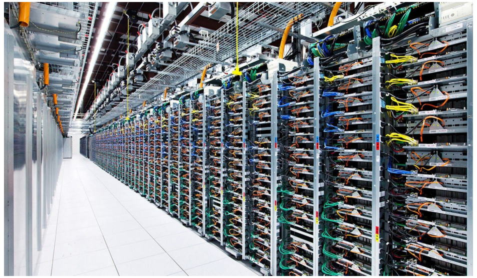

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Runtimes Can Use Complex Resource Management APIs

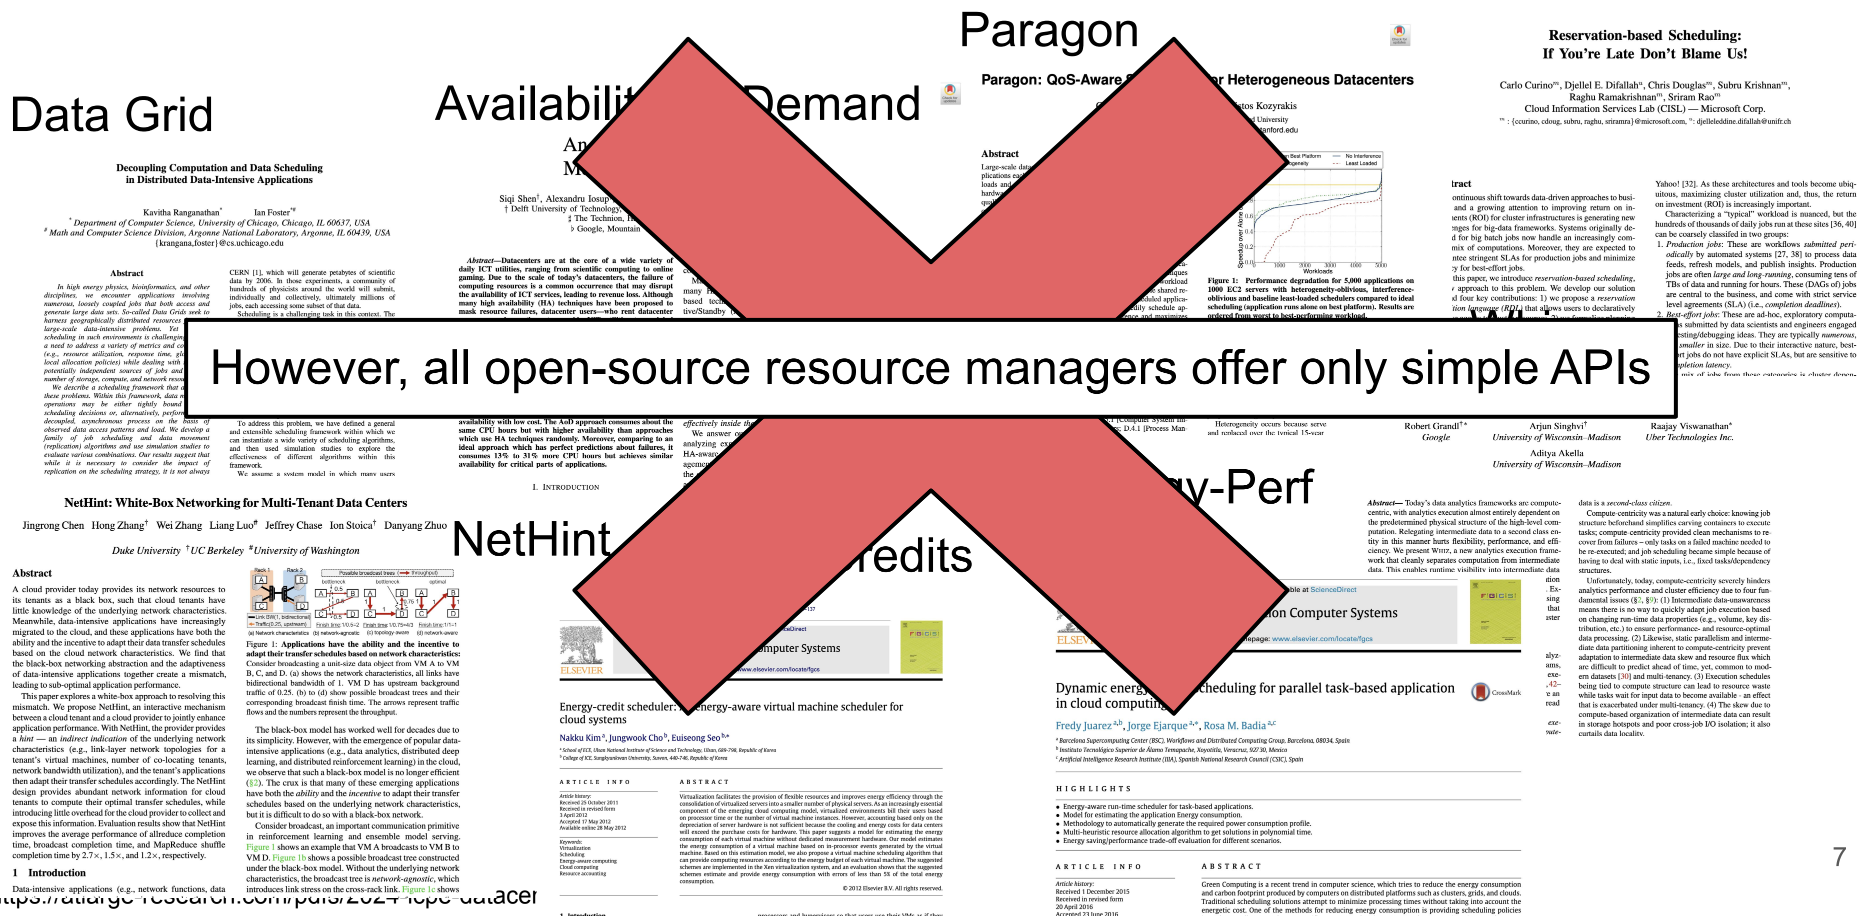

However, all open-source resource managers offer only simple APIs

Research Questions

RQ1. What API features are missing from commercial open-source resource managers?

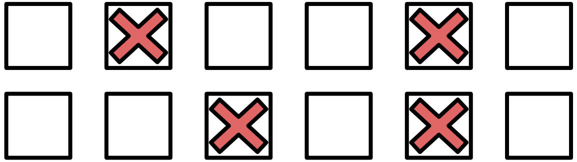

RQ2. What is the performance cost of these missing features?

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Research Questions

RQ0. What is the reference architecture for resource manager APIs?

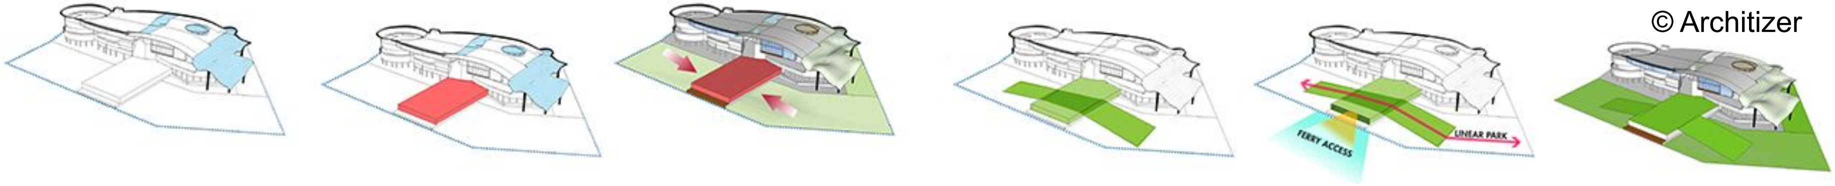

© Architizer

RQ1. What API features are missing from commercial open-source resource managers?

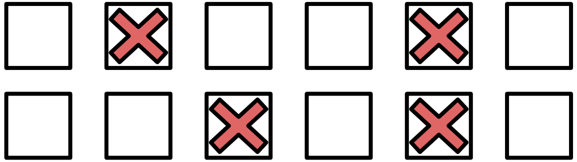

RQ2. What is the performance cost of these missing features?

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Resource Manager API Ref. Arch. Design Process

Systematic Literature Survey
Keywords eg.: scheduler
AND (datacenter OR API)
Sorted by: normalized citations

First 15 papers

Features from 5 industrial schedulers

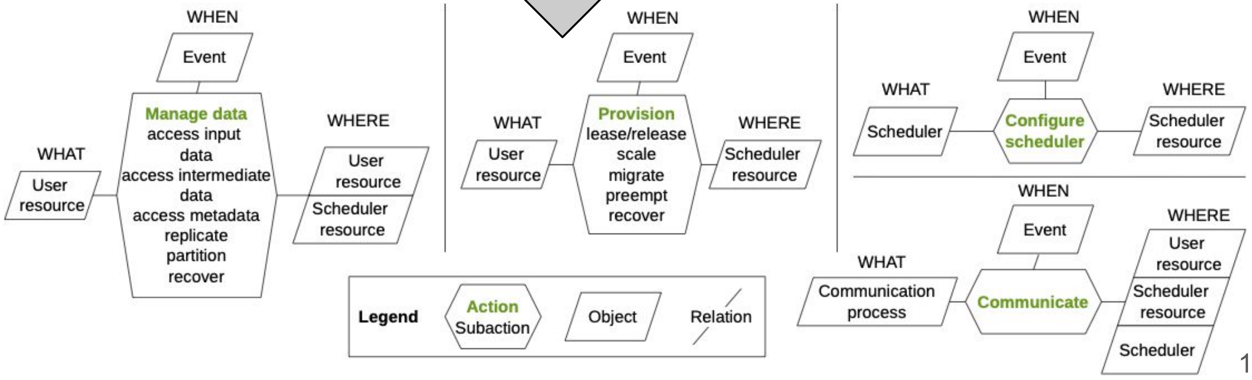

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Resource Manager API Reference Architecture

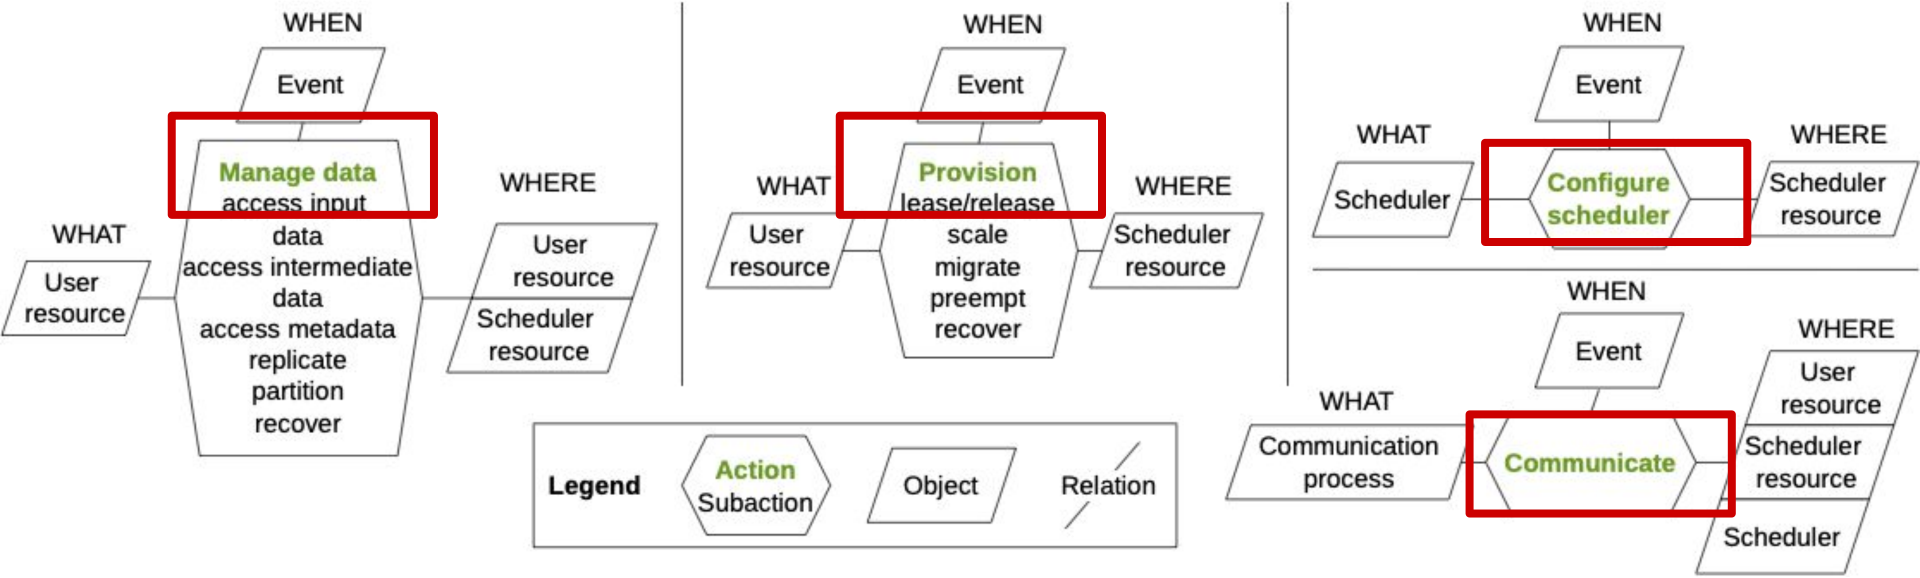

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Resource Manager API Reference Architecture - Provision

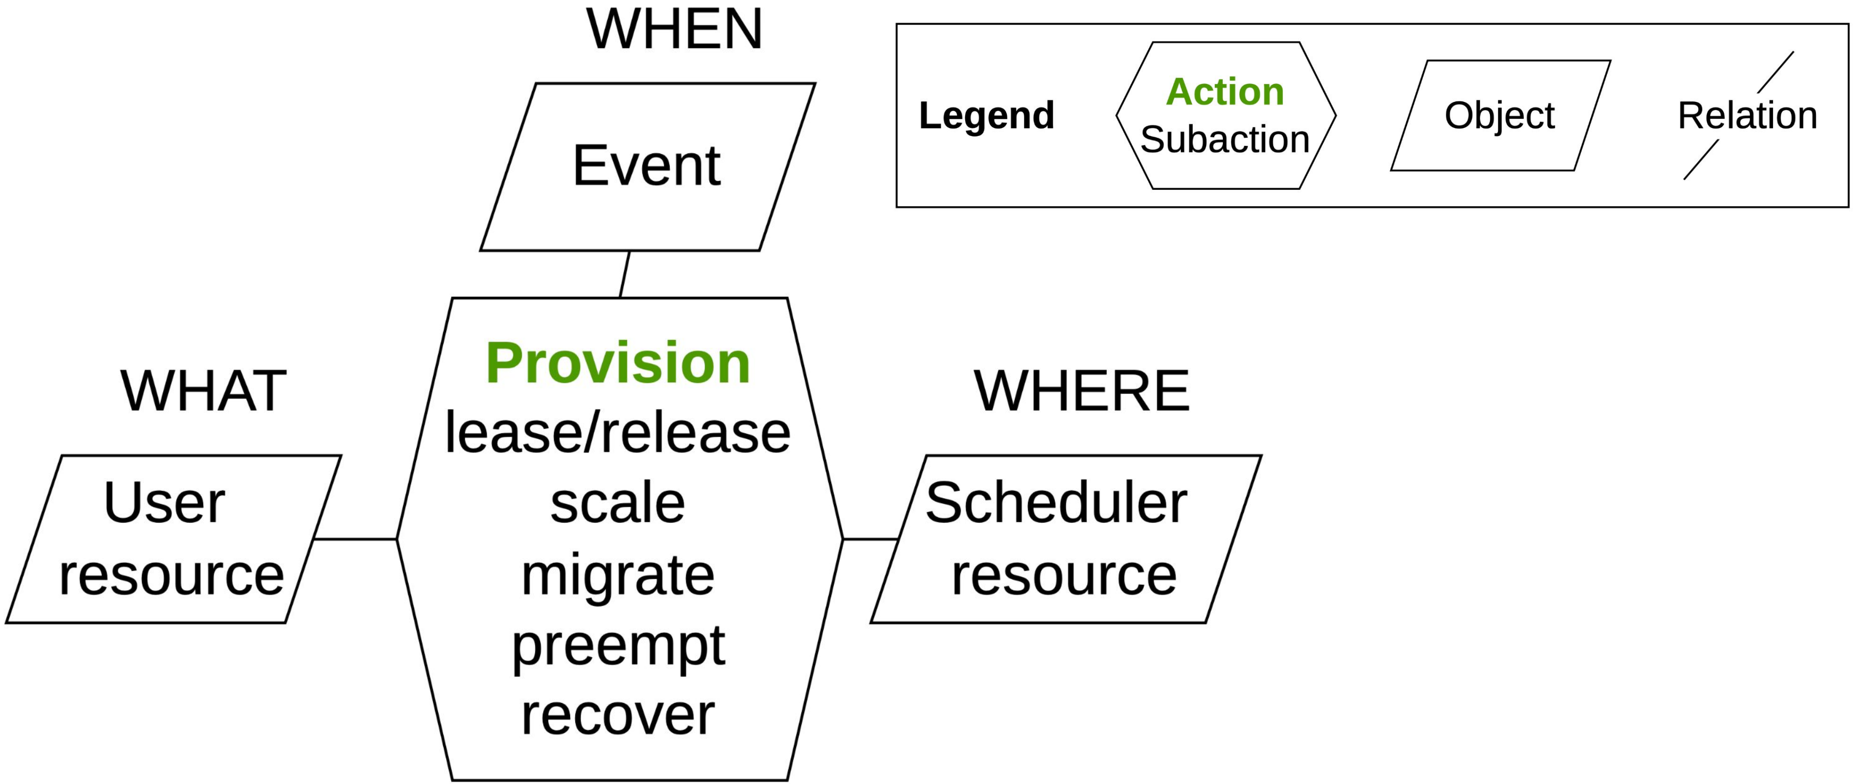

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Resource Manager API Reference Architecture - Syntax

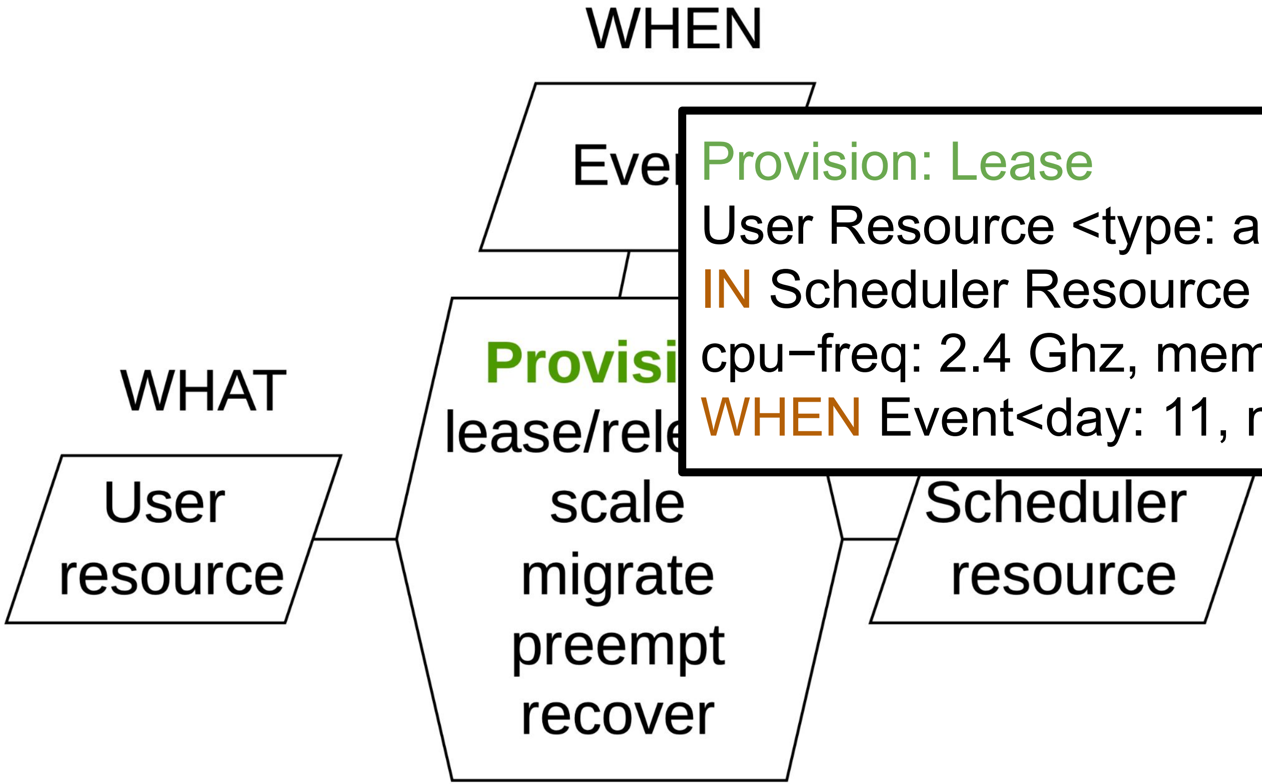

Provision: Lease
User Resource <type: app, id: 1, runtime: 1h>
IN Scheduler Resource <type: vm, cores: 8, cpu−freq: 2.4 Ghz, memory: 32 Gb>
WHEN Event<day: 11, month: 12, year: 2023>

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Resource Manager API Reference Architecture - Syntax

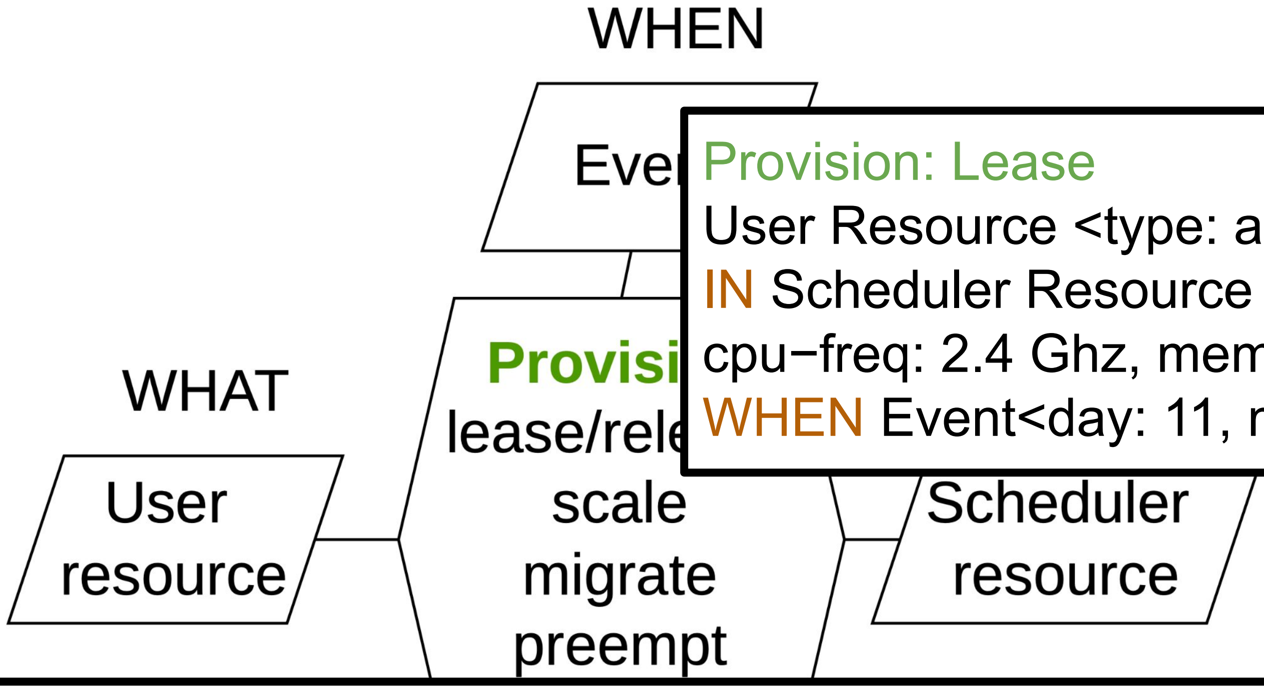

Provision: Lease
User Resource <type: app, id: 1, runtime: 1h>
IN Scheduler Resource <type: vm, cores: 8, cpu−freq: 2.4 Ghz, memory: 32 Gb>
WHEN Event<day: 11, month: 12, year: 2023>

RQ0. What is the reference architecture for resource manager APIs?

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Mapping Industrial Scheduler API Actions to Ref. Arch.

Action	Sub-Action	Schedulers				
		Ku	Sl	Sp	Co	Ai
Provision	Lease / release	●	●	●	●	●
	Scale	●	○	◗	○	○
	Migrate	○	○	○	○	○
	Preempt	◗	●	○	●	○
	Recover	◗	◗	●	◗	◗
Configure scheduler		●	◗	●	●	●
Manage data	Access input	●	◗	●	●	●
	Access interm.	○	○	◗	○	○
	Access metadata	○	○	○	○	○
	Replicate	○	○	●	○	○
	Partition	○	○	●	○	○
	Recover	◗	○	●	●	○
Communicate		◗	●	◗	◗	◗

Ku = Kubernetes
Sl = SLURM
Sp = Spark
Co = HTCondor
Ai = Airflow

● = Full support
◗ = Partial support
○ = No support

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

APIs Missing in Industrial Schedulers

Action	Sub-Action	Schedulers: Ku	Sl	Sp	Co	Ai
Provision	Lease / release	●	●	●	●	●
	Scale	●	○	◗	○	○
	Migrate	○	○	○	○	○
	Preempt	◗	●	○	●	○
	Recover	◗	◗	●	◗	◗
Configure scheduler						
Manage data	Access input	●	◗	●	●	●
	Access interm.	○	○	◗	○	○
	Access metadata	○	○	○	○	○
	Replicate	○	○	●	○	○
	Partition	○	○	●	○	○
	Recover	◗	○	●	●	○
Communicate		◗	●	◗	◗	◗

Little Support for Data Management

Ku = Kubernetes
Sl = SLURM
Sp = Spark
Co = HTCondor
Ai = Airflow

● = Full support
◗ = Partial support
○ = No support

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

APIs Missing in Industrial Schedulers

Action	Sub-Action	Schedulers: Ku	Sl	Sp	Co	Ai
Provision	Lease / release	●	●	●	●	●
	Scale	●	○	◗	○	○
	Migrate	○	○	○	○	○
	Preempt	◗	●	○	●	○
	Recover	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Configure scheduler		[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Manage data	Access input	●	◗	●	●	●
	Access interm.	○	○	◗	○	○
	Access metadata	○	○	○	○	○
	Replicate	○	○	●	○	○
	Partition	○	○	●	○	○
	Recover	◗	○	●	●	○
Communicate		◗	●	◗	◗	◗

No Support for Migration

Ku = Kubernetes
Sl = SLURM
Sp = Spark
Co = HTCondor
Ai = Airflow

● = Full support
◗ = Partial support
○ = No support

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

APIs Missing in Industrial Schedulers

Action	Sub-Action	Ku	Sl	Sp	Co	Ai
		Schedulers				
Provision	Lease / release	●	●	●	●	●
	Scale	●	○	◗	○	○
	Migrate	○	○	○	○	○
	Preempt	◗	●	○	●	○
	Recover					
Configure scheduler						
Manage data	Access input	●	◗	●	●	●
	Access interm.	○	○	◗	○	○
	Access metadata	○	○	○	○	○
	Replicate	○	○	●	○	○
Communicate		◗	●	◗	◗	◗

No Support for Migration

Ku = Kubernetes
Sl = SLURM
Sp = Spark
Co = HTCondor
Ai = Airflow

● = Full support
◗ = Partial support
○ = No support

RQ1. What API features are missing from commercial open-source resource managers?

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Cost of Missing APIs - Experiments

3 Experiments Using the OpenDC Simulator

1. Resource Reservation
2. Container Migration
3. Storage Metadata

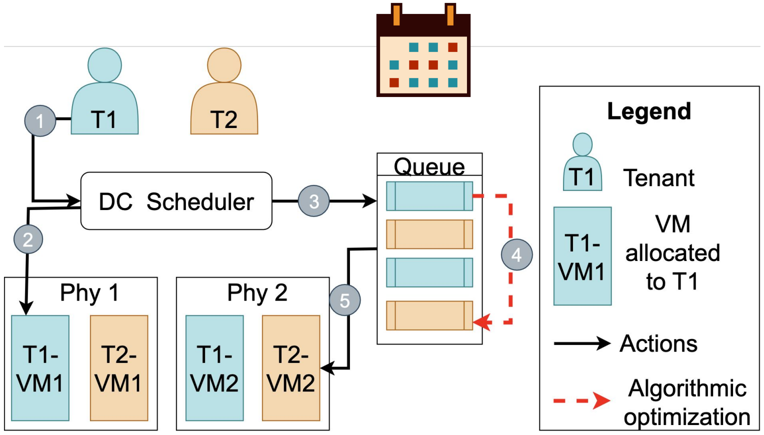

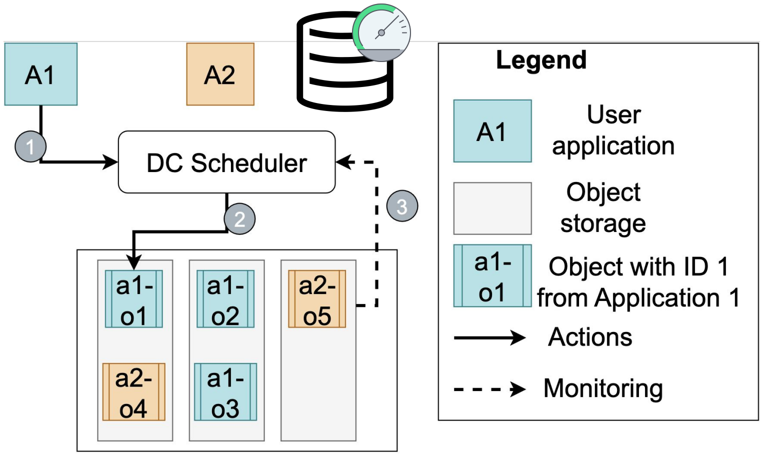

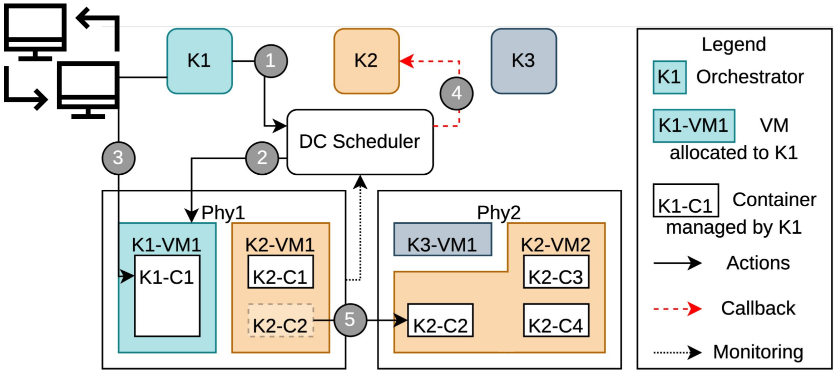

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Container Migration Experiment - Setup

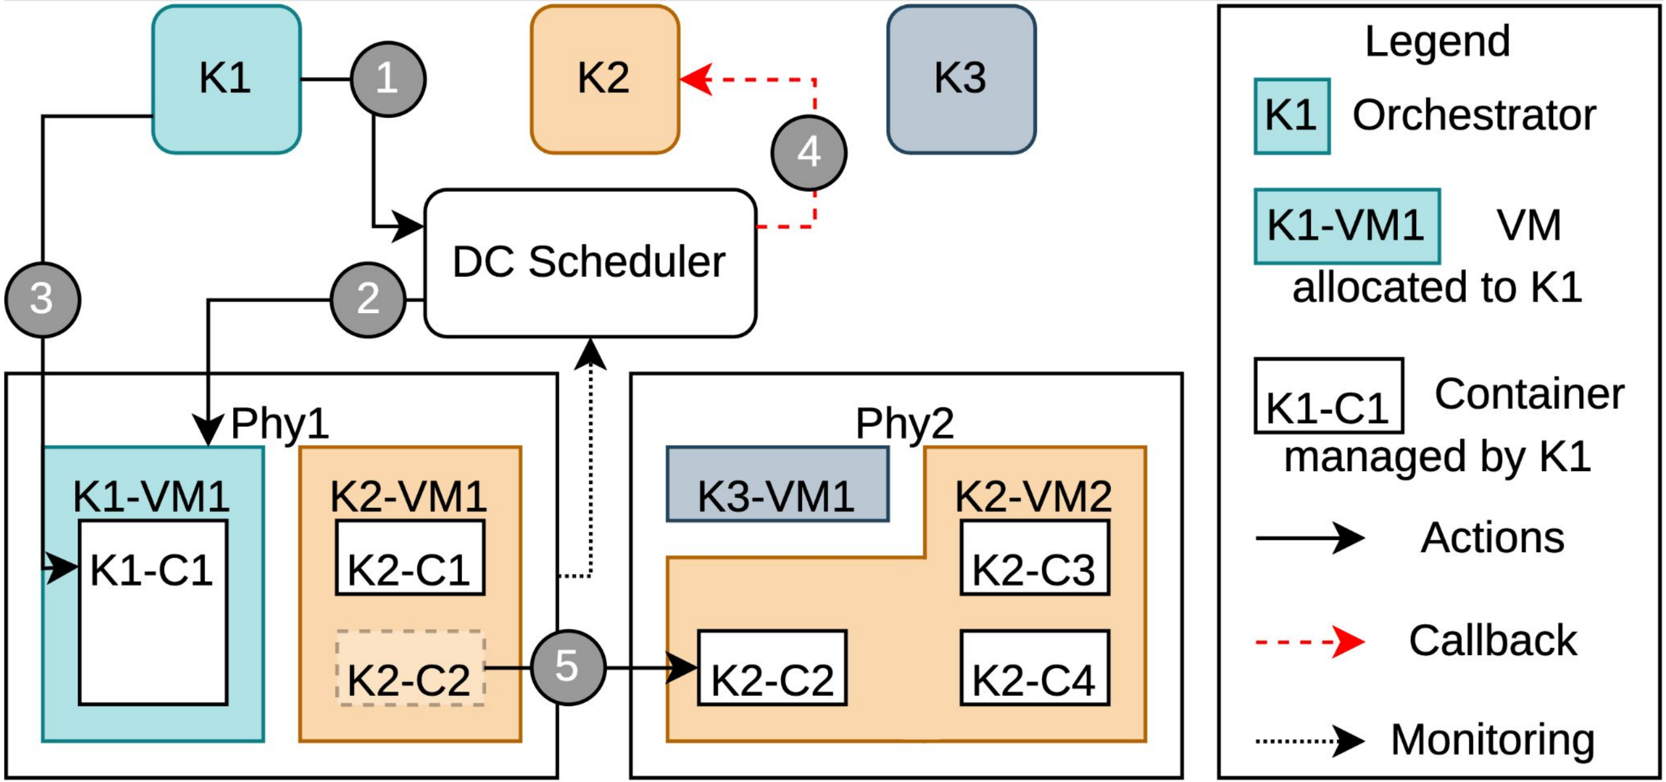

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Virtual Machines - Resource Underutilization

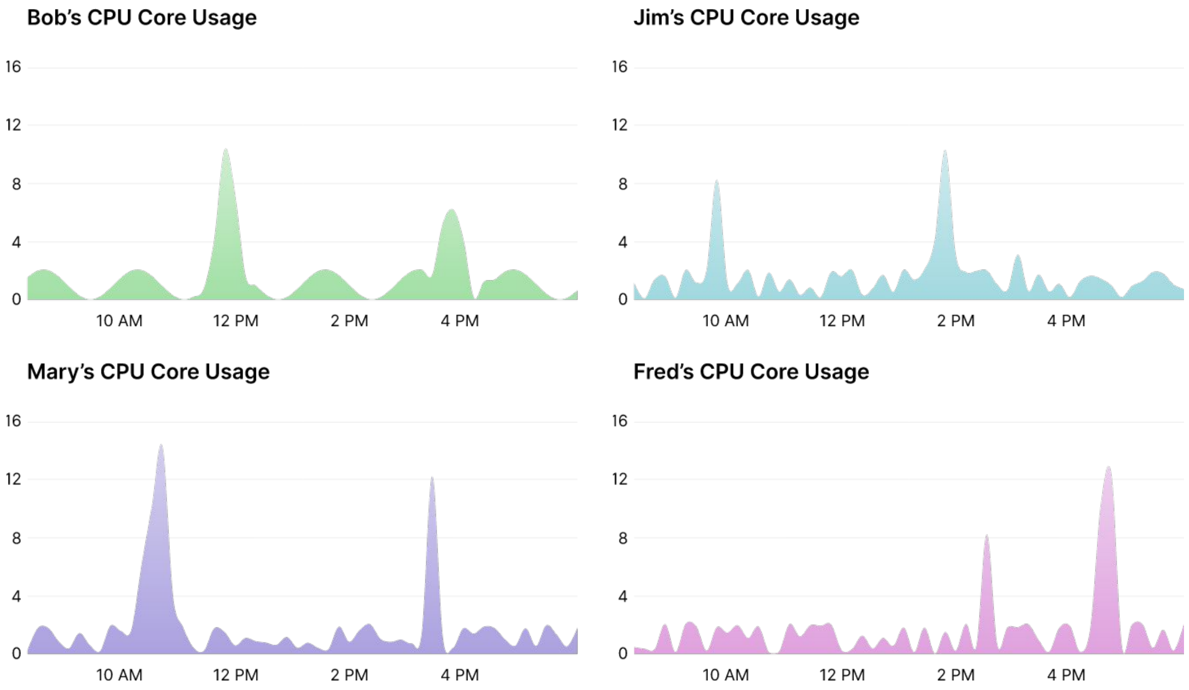

© https://coder.com/docs/v1/latest/guides/admin/resources

Oversubscription

1 CPU for 3 VMs

Oversubscription ratio = 3

Time sharing

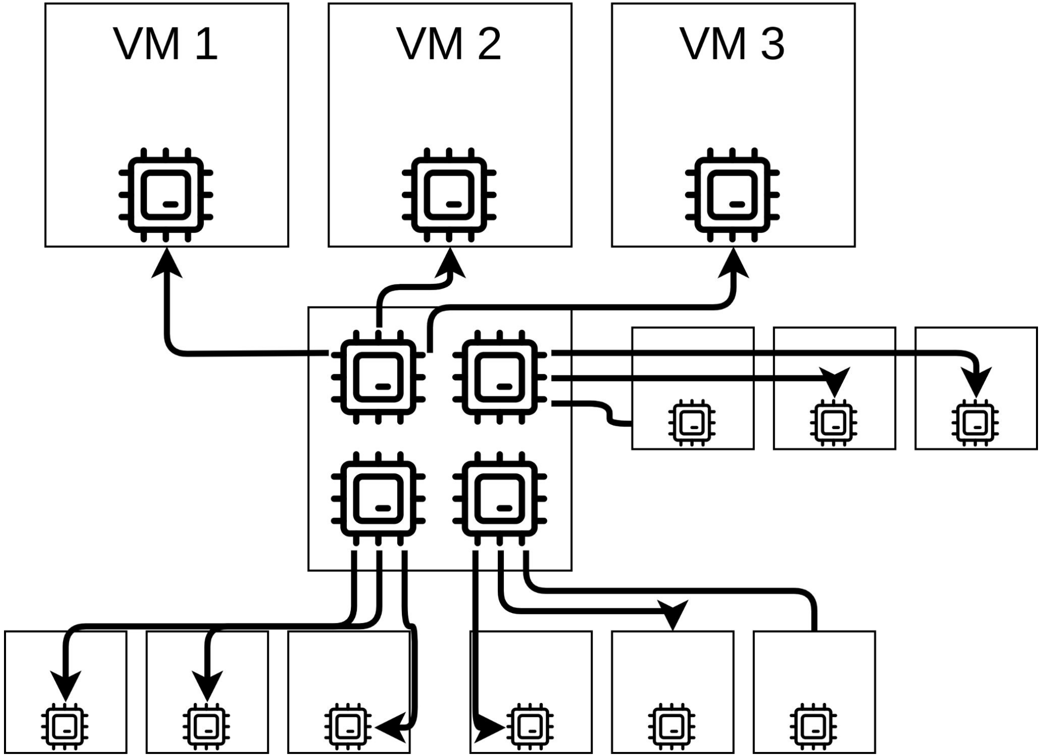

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Oversubscription - Resource Contention

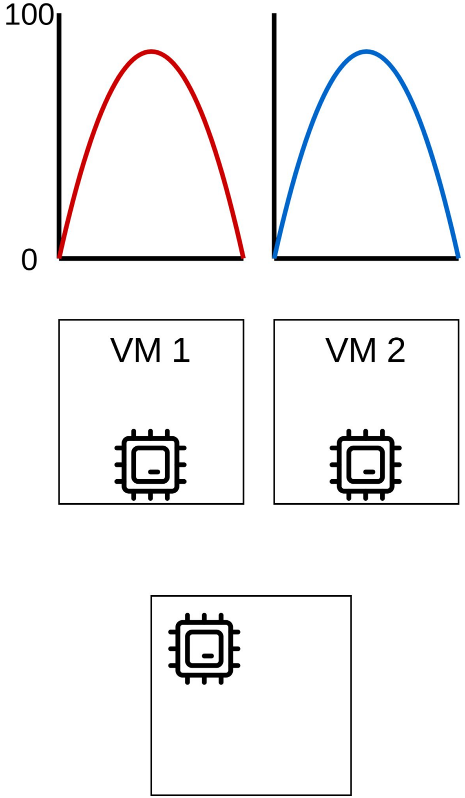

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Migration

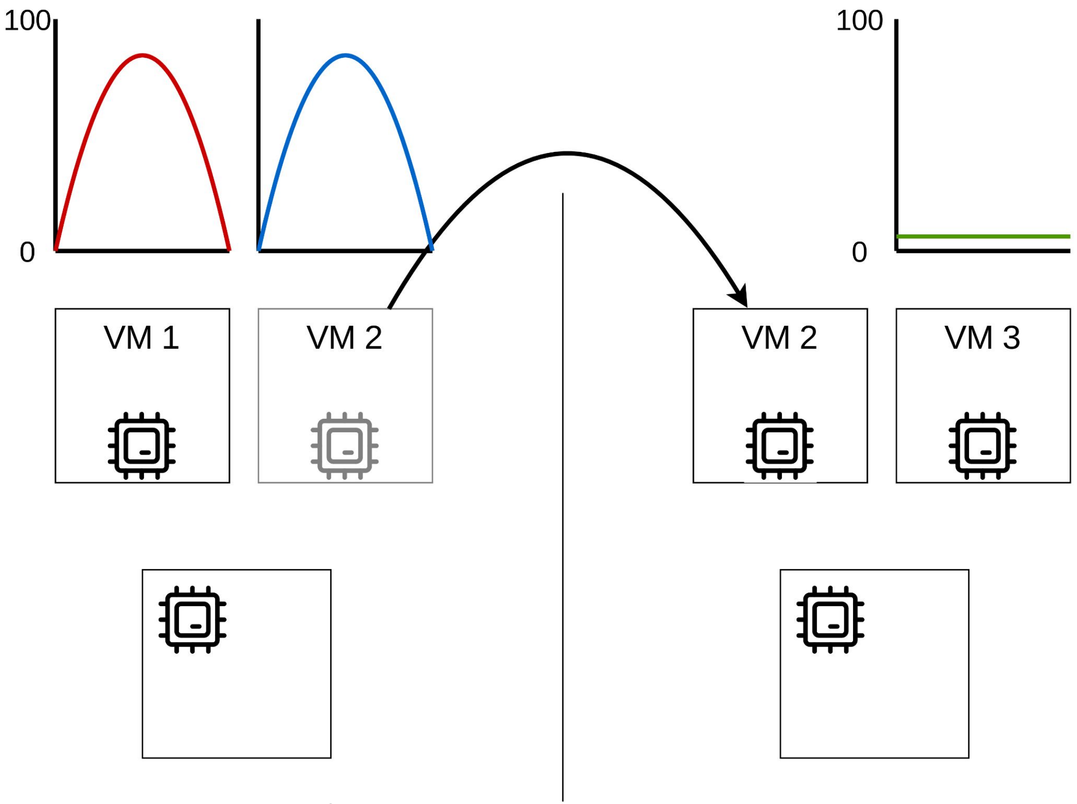

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Container Migration

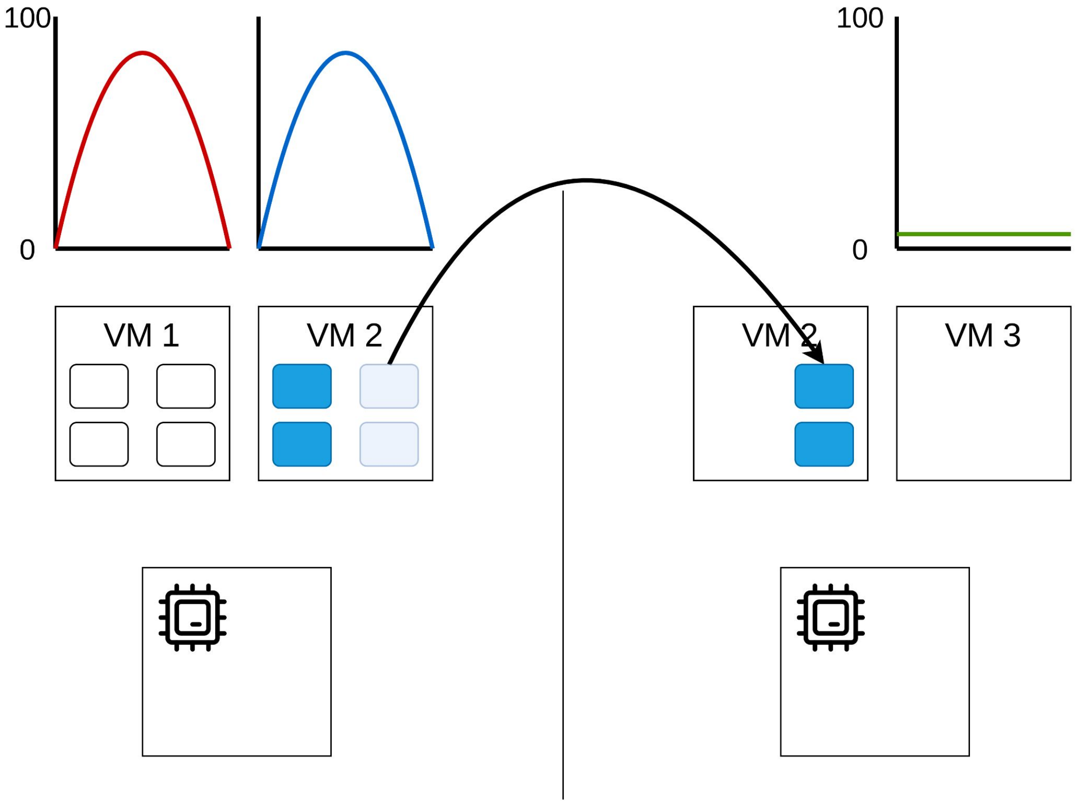

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Container Migration Experiment - Setup

Cluster Setup:
1 Physical Cluster
5 Kubernetes Virtual Clusters
512 Mbps migration speed
80% average CPU utilization target
    Physical cluster size calibrated for each trace

3 Traces:
Google 2011 (~25 machines)
Azure 2017 (~65 machines)
BitBrains 2015 (~100 machines)

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Container Migration - Results - Azure P90 App Runtime

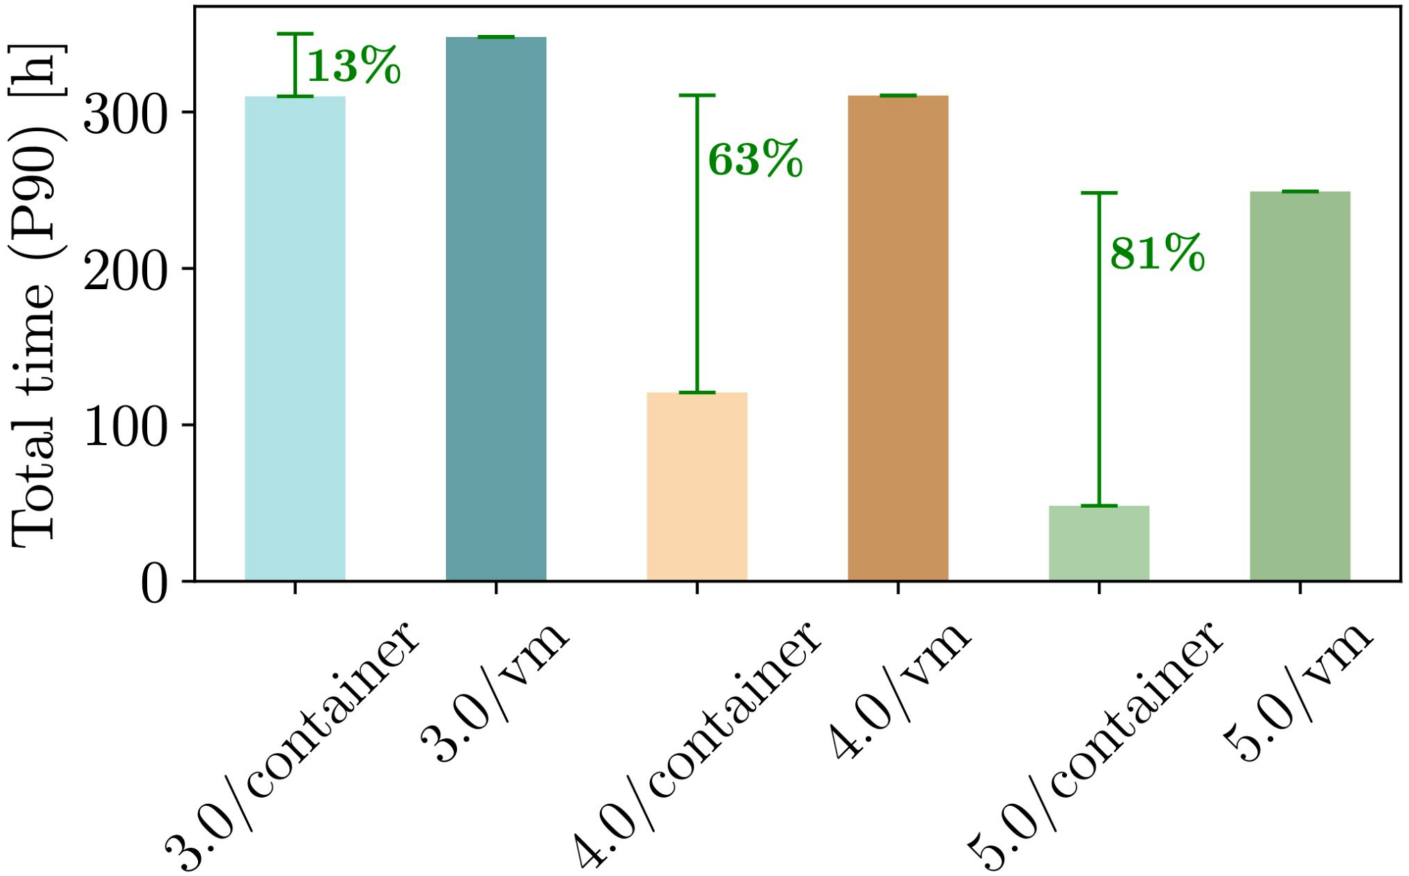

Oversubscription Ratio / Migrated object

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Container Migration - Results - Azure P90 App Runtime

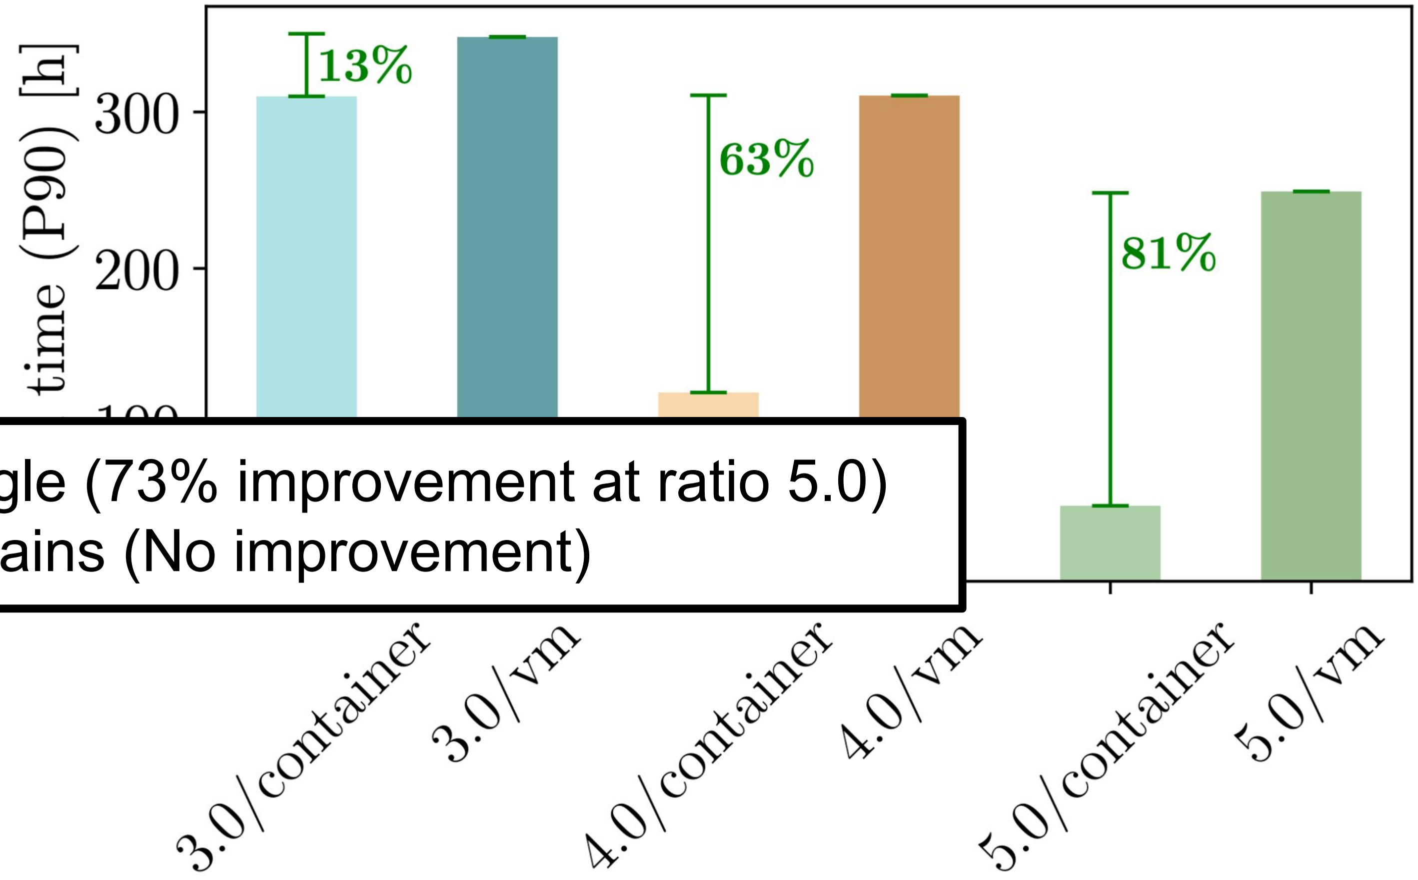

Google (73% improvement at ratio 5.0)
Bitbrains (No improvement)

Oversubscription Ratio / Migrated object

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Key Takeaways

- Distributed Systems Runtimes can take advantage of complex APIs for resource management
- Current open-source resource managers are missing complex resource management abstractions
  - Migrations, Data Management
- Complex abstractions offer performance benefits for some workloads, but not all
  - 17% to 81% improve in 90th percentile app runtime using container migration callbacks for Azure trace, but not the Bitbrains trace

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Why didn’t these APIs make it into industrial schedulers?

1. New APIs/scheduling techniques only work for specific workloads
   a. Widely-deployed schedulers need generality

2. Need to implement separately for different systems and system versions
   a. Can't use Spark 3.2 scheduler with Spark 2.4
   b. Cannot use a SLURM scheduler with Kubernetes

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Future Work - Kubernetes-based Embeddable Scheduler

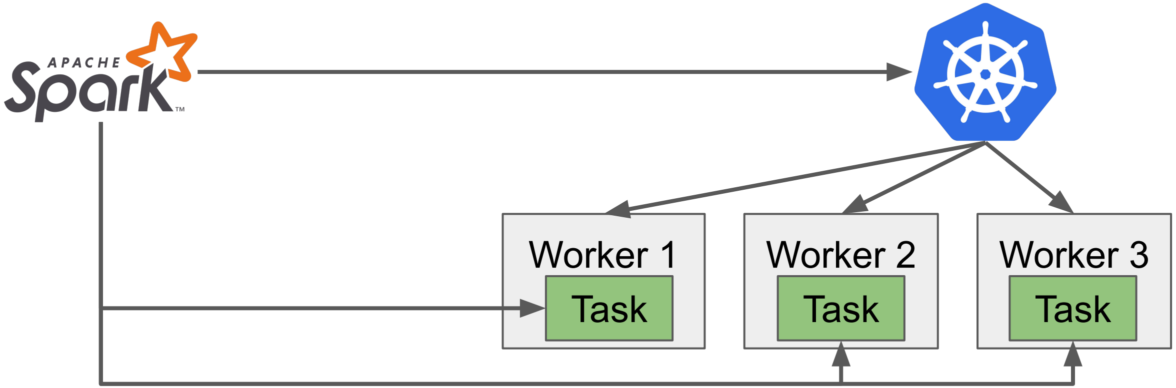

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Future Work - Kubernetes-based Embeddable Scheduler

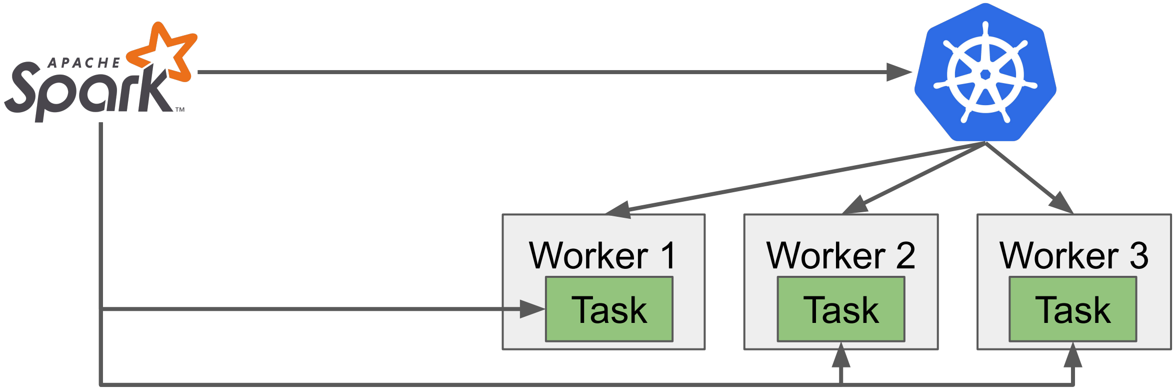

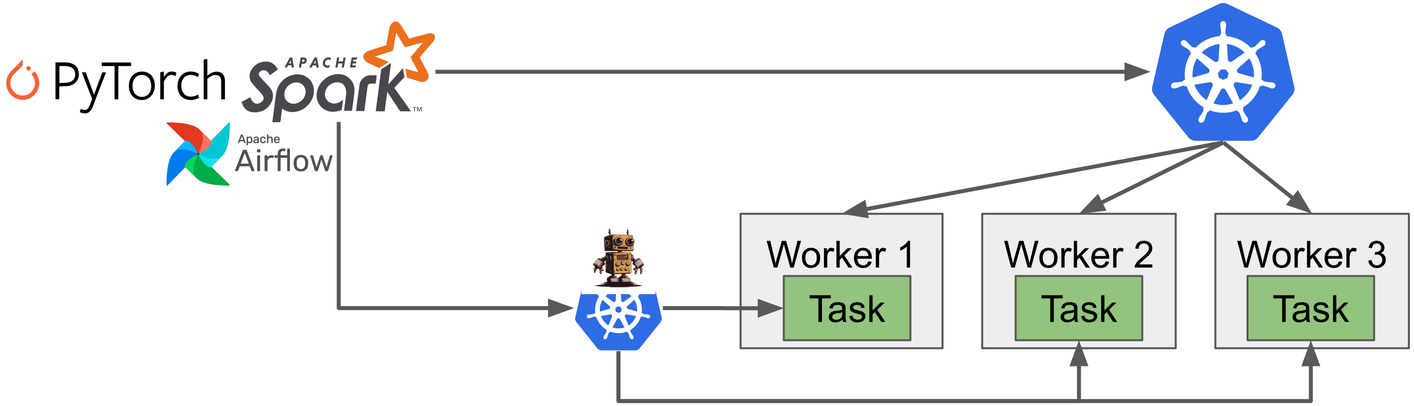

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Future Work - Kubernetes-based Embeddable Scheduler

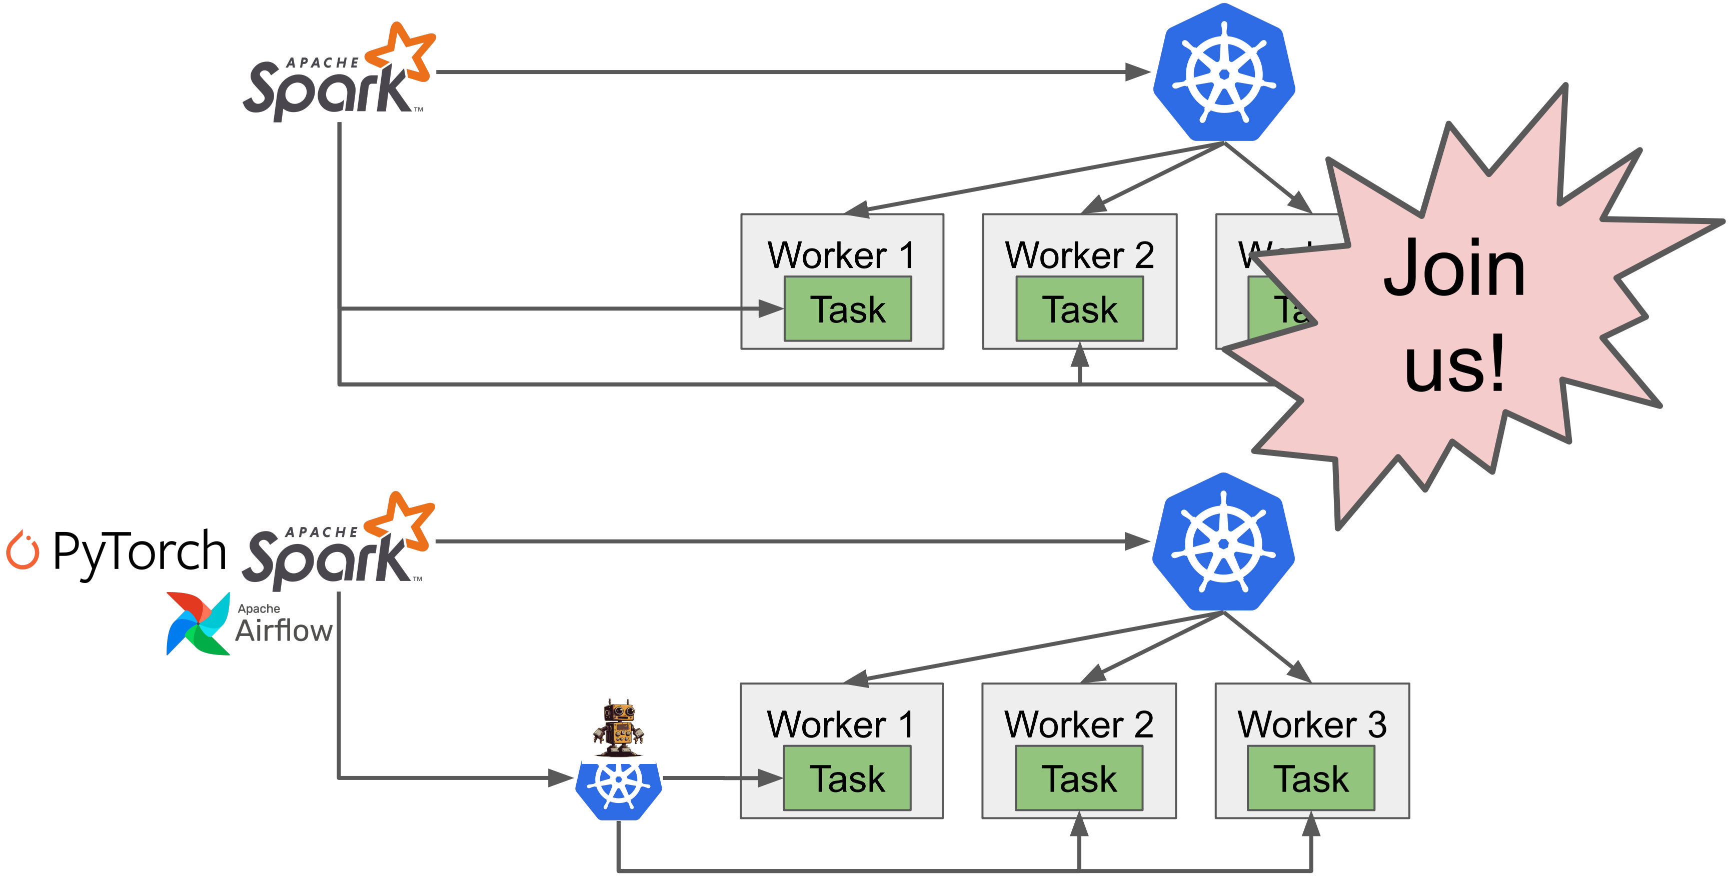

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Key Takeaways

- Distributed Systems Runtimes can take advantage of complex APIs for resource management
- Current open-source resource managers are missing complex resource management abstractions
  - Migrations, Data Management
- Complex abstractions offer performance benefits for some workloads, but not all
  - 17% to 81% improve in 90th percentile app runtime using container migration callbacks for Azure trace, but not the Bitbrains trace

https://atlarge-research.com/pdfs/2024-icpe-datacenter-scheduler.pdf

Further Reading

[this work] Talluri, S., Herbst, N., Abad, C., De Matteis, T., & Iosup, A. (2024). ExDe: Design space exploration of scheduler architectures and mechanisms for serverless data-processing. Future Generation Computer Systems, 153, 84-96.

[related work on scheduler APIs] Manterola Lasa, A., Talluri, S., De Matteis, T., & Iosup, A. (2024, May). The Cost of Simplicity: Understanding Datacenter Scheduler Programming Abstractions. In Proceedings of the 15th ACM/SPEC International Conference on Performance Engineering (pp. 166-177).

[using the tools] Talluri, S., Herbst, N., Abad, C., Trivedi, A., & Iosup, A. (2023, May). A Trace-driven Performance Evaluation of Hash-based Task Placement Algorithms for Cache-enabled Serverless Computing. In Proceedings of the 20th ACM International Conference on Computing Frontiers (pp. 164-175).

[reference architecture] Andreadis, G., Versluis, L., Mastenbroek, F., & Iosup, A. (2018, November). A reference architecture for datacenter scheduling: design, validation, and experiments. In SC18: International Conference for High Performance Computing, Networking, Storage and Analysis (pp. 478-492). IEEE.