Container Start Times: Empirical Analysis and Predictability

12th Symposium on Software Performance

Martin Straesser, Samuel Kounev
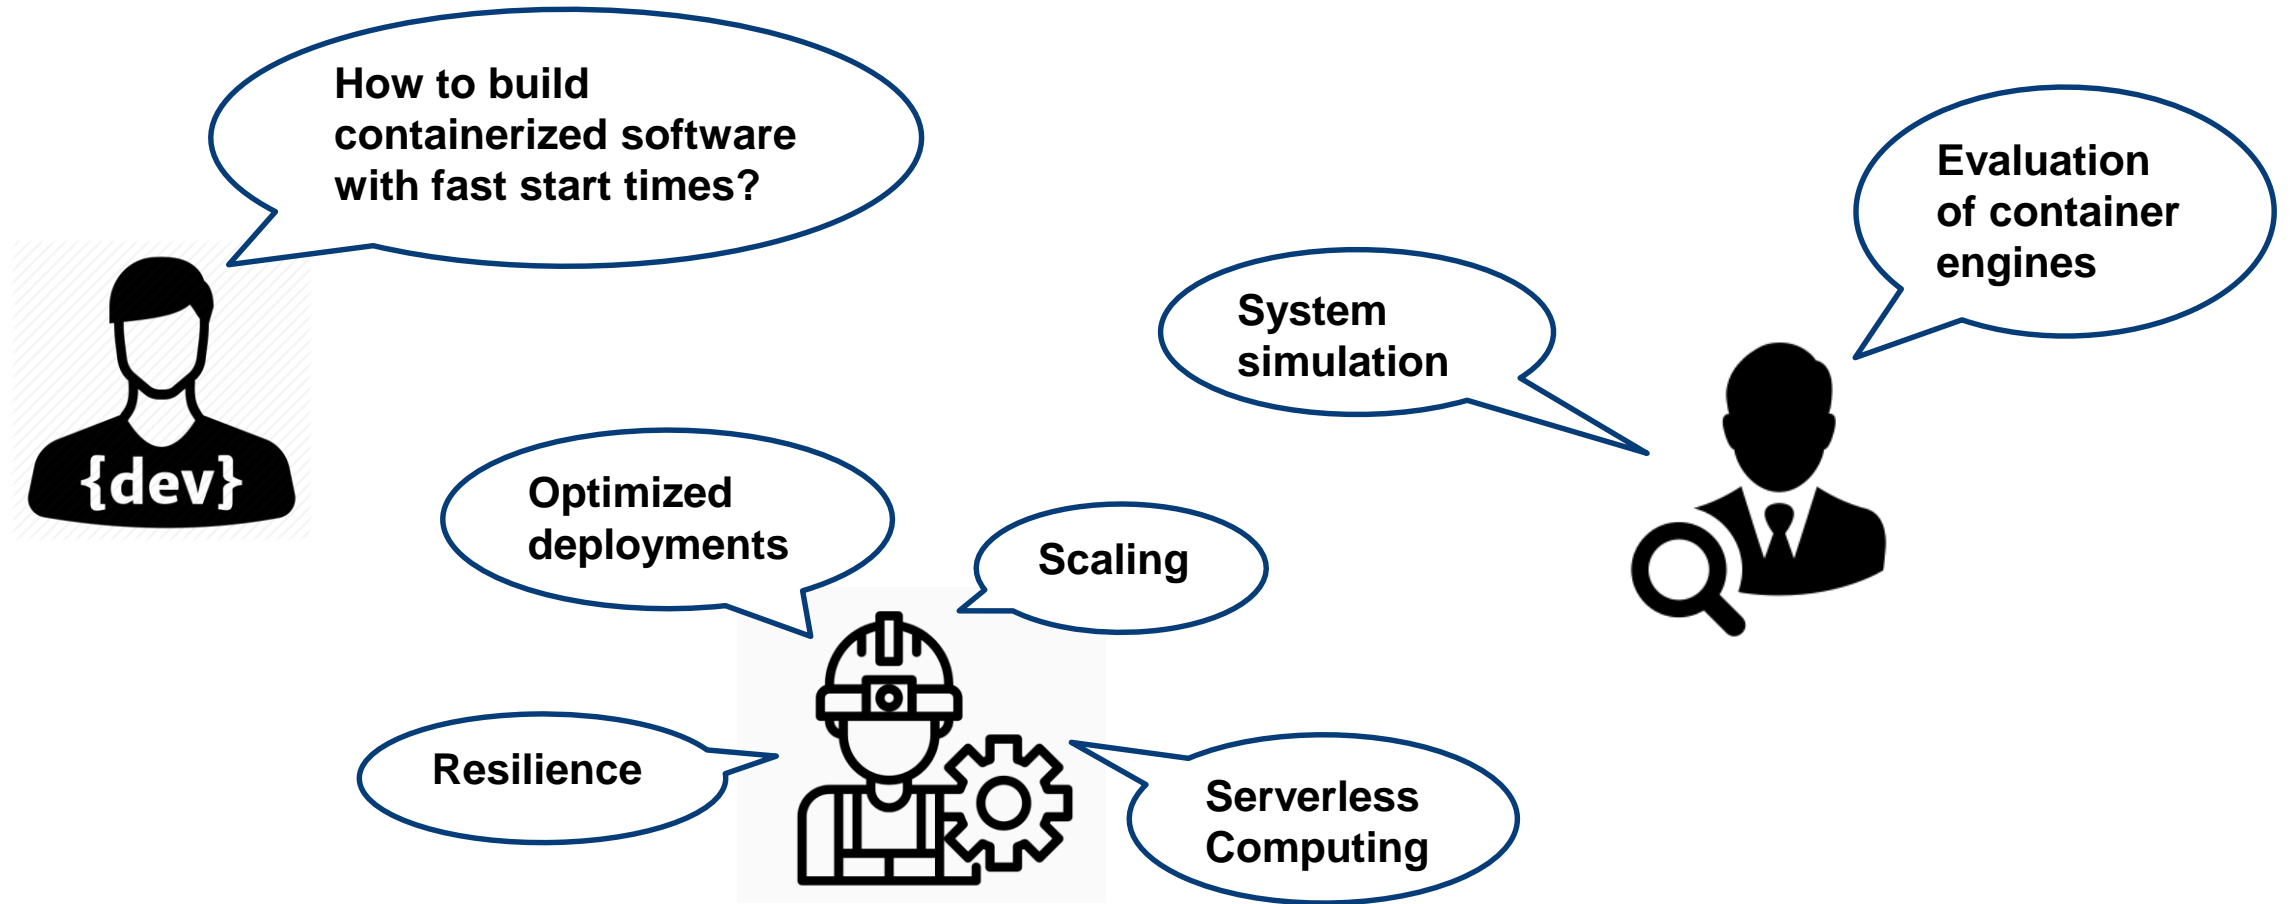
10.11.2021

Leipzig, Germany

https://se.informatik.uni-wuerzburg.de

Motivation

➢ Container start times are interesting for various domains

How to build containerized software with fast start times?

Evaluation of container engines

System simulation

Optimized deployments

Scaling

Resilience

Serverless Computing

Problem Statement

➢ What is the start time of a container?

- The time difference between the following two events

- Event 1: The container engine (e.g. Docker) receives the command to start a new container

- Event 2: The container starts the execution of the entrypoint command

- Note: This time difference can be evaluated without deeper knowledge about the container and without dependencies of the container (in contrast to readiness time)

➢ Central research questions

- What are influencing factors for container start times?

- Can we predict the start time of an arbitrary container?

➢ Idea: Test a large set of containers from the public repository Docker Hub

Influencing Factors For Container Start Times

Image

- Image size
- Number of layers
- Size of layers
- Platform architecture
- Platform OS
- Number of exposed ports
- Volumes
- Start command
- ...

Deployment

- (Image in local repo?)
- Node architecture
- Node OS
- Node kernel
- Container runtime
- Storage driver
- Number of containers deployed
- ...

UNI
WÜ

Study Phases

Design and Development

- Setting the project scope
- Study and setup design
- Software development

Measurements in Private Cloud

- Process and software testing
- Determine settings for public cloud (e.g., measurement repetitions, multi-threading, …)

Measurements in Public Cloud

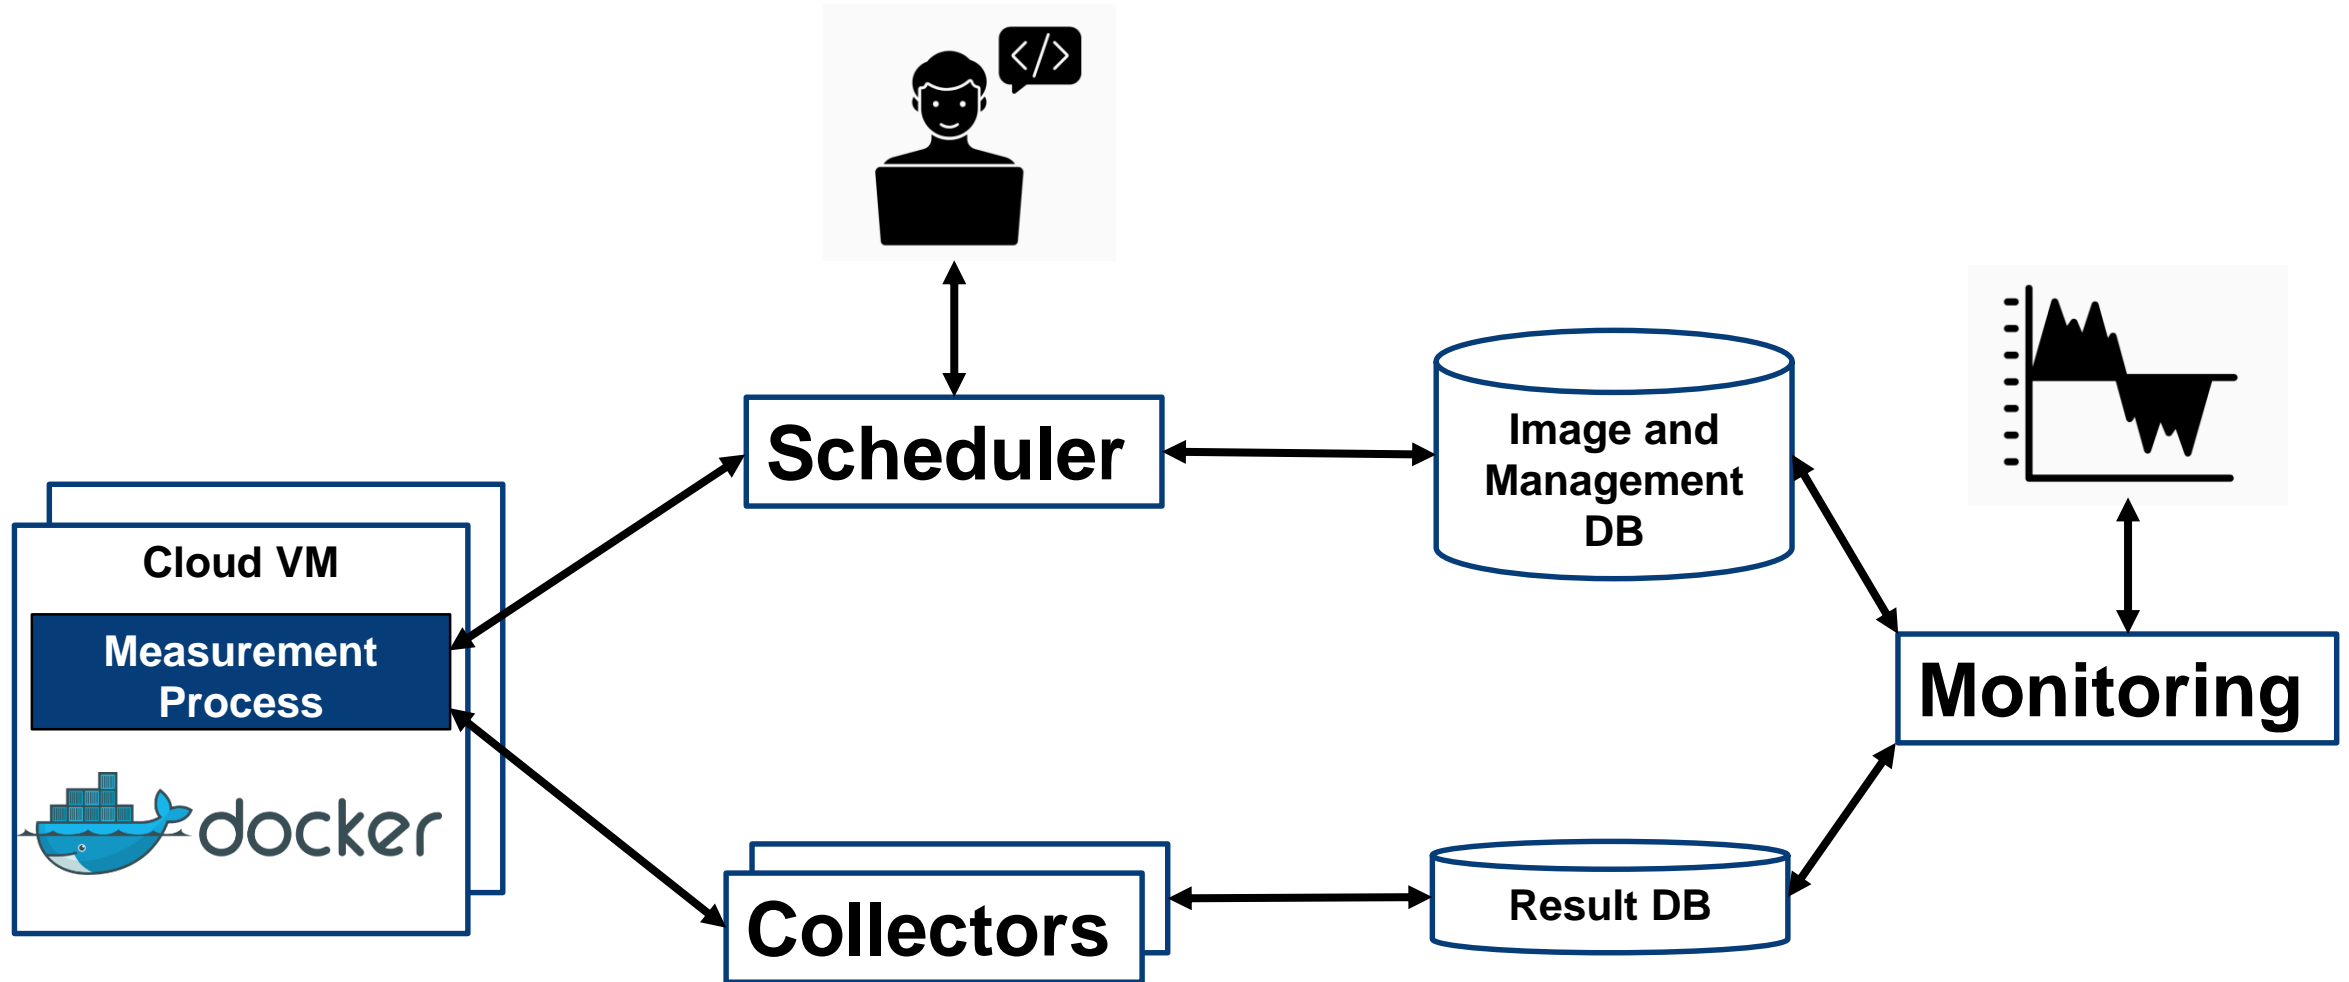- Container start time measurements in a representative environment

Google Cloud

Result Analysis

- Data analysis and visualization
- Prediction model
- Identify open questions and possible future work

Experiment Setup

Result Data and Contributions

Image Metadata

ImageID	Image Size	Volumes	Ports	Layers	…
x	12	5	2	32	…
y	45	1	1	13	…
…					
z	34	0	1	8	…

Measurement Results

ImageID	Start Time	…
x	1231	…
x	1564	…
…		
y	2344	…
y	2654	…
…		

➢ Contributions

- Dataset

- Statistical measures starting one specific image

- Regression data (manifest parameters → start time)

- Prediction model

- Others (e.g. evolution of start times over different versions)

UNI WÜ

Extension Points and Recent Challenges

➢ Extension points

- Validating results for start times in environments with container orchestration (e.g. Kubernetes)

- Extend evaluation with parameters from deployment (e.g. host OS, storage driver, …)

- Evaluate other container engines

- …

➢ Recent challenges

- Getting enough pullable images from Docker Hub (no full list available and search results are always limited to 400 images)

- Developing the software for measurement execution and result collection (databases, scheduling, live monitoring, failure prevention and handling, …)

Methodology

➢ We explicitly consider best practices from latest ACM SIGSOFT Empirical Standards [1]

➢ We maintain a question catalogue which answers fundamental questions for all stages of the study

General Standard	Action Research	Benchmarking	Case Study	Case Survey
Data Science	Engineering Research	Experiments	Grounded Theory	Longitudinal
Meta Science	Mixed Methods	Optimization Studies	Qualitative Surveys	Quantitative Simulation
Questionnaire Surveys	Systematic Reviews			

[1] ACM Empirical Standards. 2021. https://acmsigsoft.github.io/EmpiricalStandards/docs/

Methodology

1. **Preliminaries**
   1. Motivation
   2. Problems, objectives, research questions
   3. Related key concepts
   4. Relationship to other studies

2. **Study Execution**
   1. Overview of methodology
   2. Data description
   3. Environment of data acquisition
   4. Temporal aspects of data acquisition
   5. Data collection and processing

3. **Investigation of Results**
   1. Data analysis
   2. Result presentation
   3. Answers to research questions

4. **Wrapping up**
   1. Assumptions, limitations and threats to validity
   2. Replicability
   3. Open questions and future work

5. **Supplementary Materials**
   1. Datasets and source code
   2. Technical documentation

Conclusion

Problem

- No systematic study on container start times with a large number of images in a representative environment
- Influencing factors are only partially known

Idea

- Measure the start time of a large set of containers from a public repository
- Analyze image and deployment factors

Benefits

- Deeper knowledge about container start times is beneficial in various domains, e.g., development of containerized software, serverless computing, …