

---

# Soroban: Attributing latency in virtualized environments

---

*Lucian Carata*  
*lucian.carata@cl.cam.ac.uk*

*James Snee, Oliver R.A. Chick, Ripduman Sohan,  
Ramsey Faragher, Andrew Rice, Andy Hopper*

---

Tail latency increases when you virtualize software.

---

*But who's responsible?*

---

of individual requests

Tail latency increases when you virtualize software.

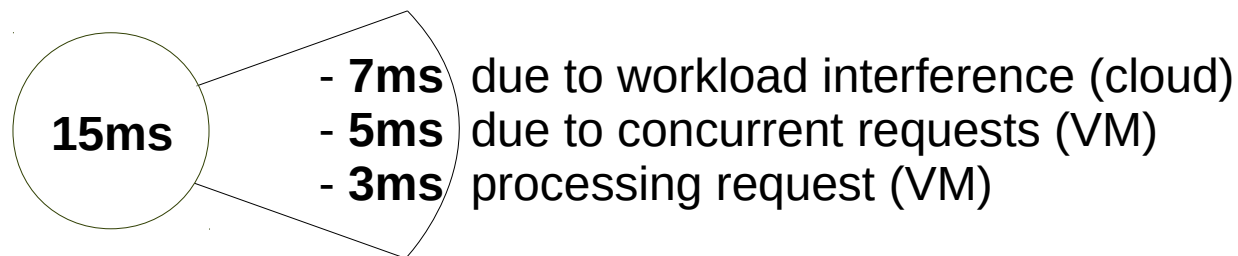
---

*But who's responsible?*

# Attribution as a cloud problem

---

- Typical issues:
  - tail latency
  - perf variability & isolation
  - workload colocation side-effects
- Measurement only tells half of the story



*Attribution gives an in-depth view of time spent doing an action*

# Attribution as a cloud problem

---

- Typical issues:
  - tail latency
  - perf variability & isolation
  - workload colocation side-effects
- Measurement only tells half of the story
  - Precise measurements
  - Context (relating different measurements)
  - Inference

*Attribution gives an in-depth view of time spent doing an action*


# Attribution based on measurement

---


- Setup:
  - Lighttpd (modified with Soroban API),
  - 16 concurrent VMs,  
Low contention (periodic CPU load, low network load)
- **Goal:**
  - determine how much is the concurrent workload slowing down each request
- Start from a simple idea (soroban-enabled measurement):
  - look at the time the lighttpd VM was scheduled out
  - during each request!



# Measurement




# Measurement







# Measurement



# Measurement



# Measurement



# Soroban

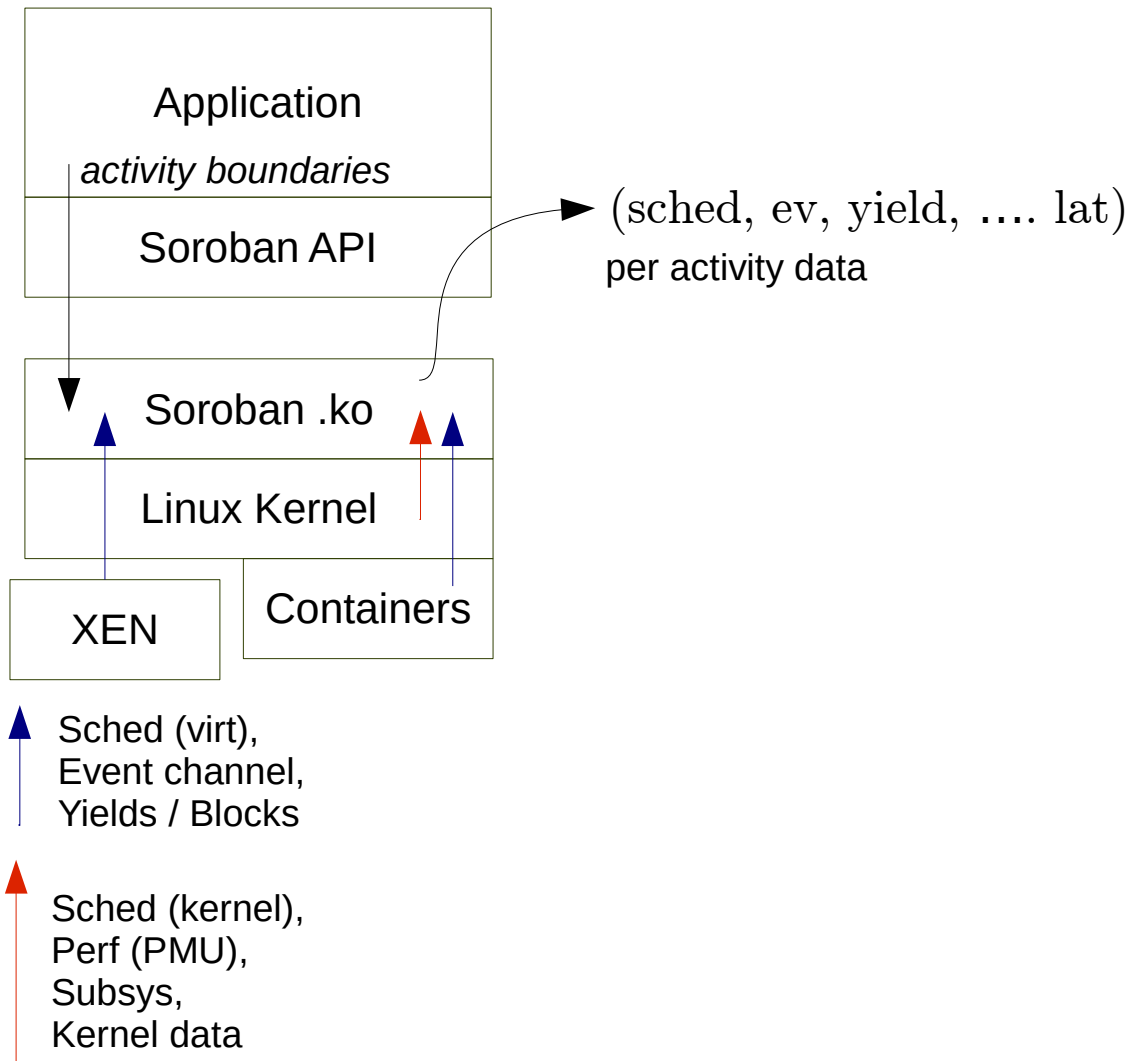
---

- Precise measurement API
  - Activity boundaries
- Kernel module
  - Provides measurement context
  - Aggregation
- Inference
  - Machine learning to determine attribution model

# Soroban


## Measurement

## Training



# Soroban

## Measurement



▲ Sched (virt),  
Event channel,  
Yields / Blocks


▲ Sched (kernel),  
Perf (PMU),  
Subsys,  
Kernel data

(sched, ev, yield, .... lat)  
per activity data

## Training

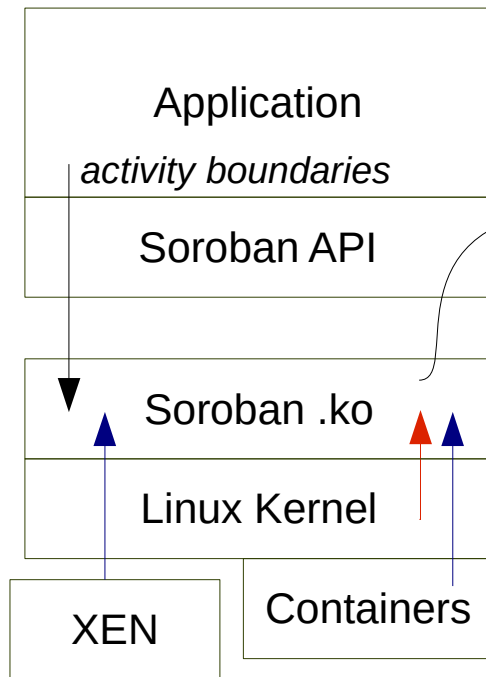
Measurements in virtualized  
environment

Reference latency profile  
(bare metal)



# Soroban

## Measurement



▲ Sched (virt),  
Event channel,  
Yields / Blocks


▲ Sched (kernel),  
Perf (PMU),  
Subsys,  
Kernel data

(sched, ev, yield, .... lat)  
per activity data


## Training

Measurements in virtualized  
environment

Reference latency profile  
(bare metal)

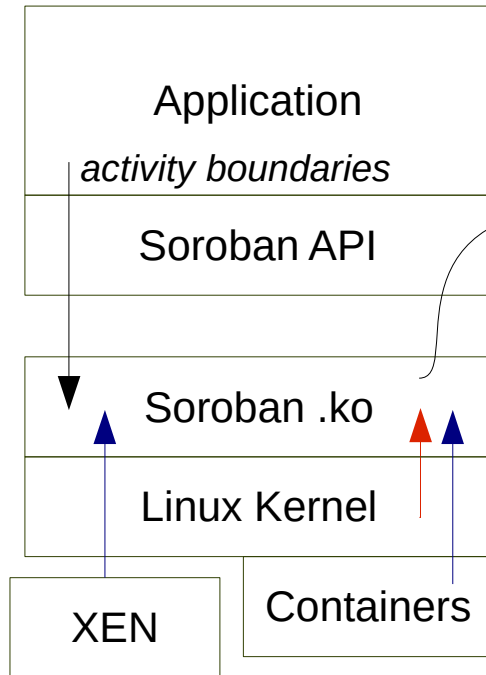


quantile-quantile difference



# Soroban

## Measurement



▲ Sched (virt),  
Event channel,  
Yields / Blocks


▲ Sched (kernel),  
Perf (PMU),  
Subsys,  
Kernel data

(sched, ev, yield, .... lat)  
per activity data

## Training

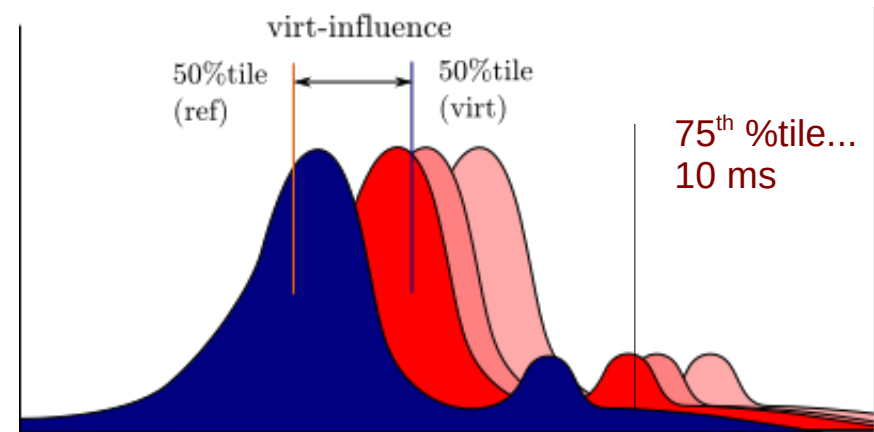
Measurements in virtualized  
environment

Reference latency profile  
(bare metal)



23 ms ... 75<sup>th</sup> %tile


quantile-quantile difference





# Soroban

## Measurement



▲ Sched (virt),  
Event channel,  
Yields / Blocks

▲ Sched (kernel),  
Perf (PMU),  
Subsys,  
Kernel data

## Training


Measurements in virtualized environment

(sched, ev, yield, .... lat)  
per activity data

23 ms ... 75<sup>th</sup> %tile

quantile-quantile difference

Reference latency profile  
(bare metal)




75<sup>th</sup> %tile...  
10 ms

13 ms

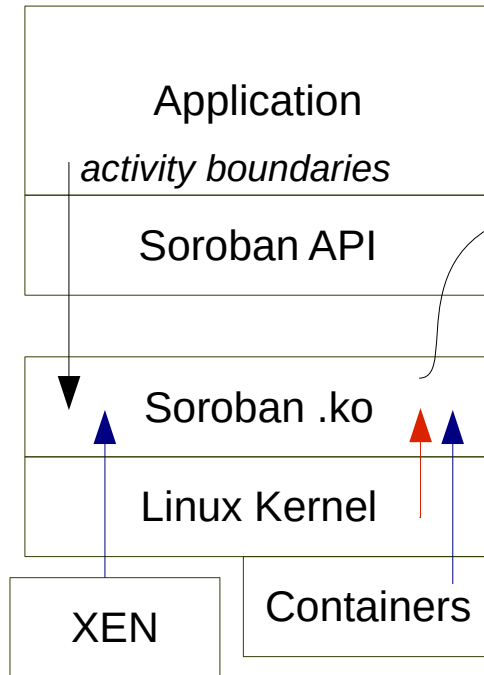
(sched, ev, yield, .... lat, **virt-influence**)

Gaussian Process  
regression model



# Soroban

## Measurement



▲ Sched (virt),  
Event channel,  
Yields / Blocks


▲ Sched (kernel),  
Perf (PMU),  
Subsys,  
Kernel data

## Measurements in virtualized environment

(sched, ev, yield, .... lat)  
per activity data

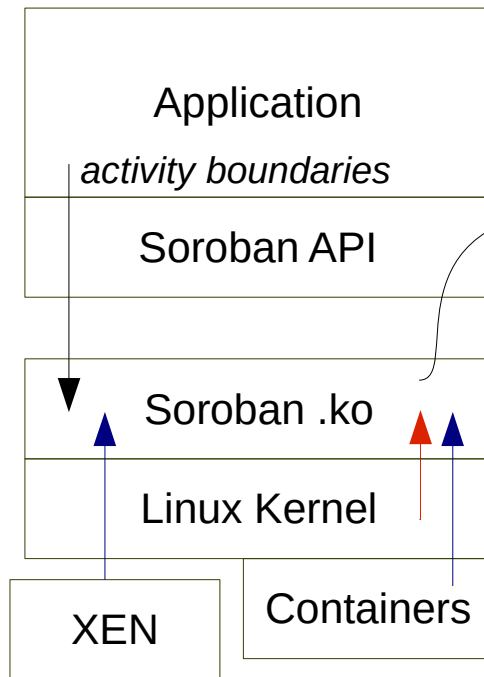
## Attribution

## Gaussian Process regression model



# Soroban

## Measurement



▲ Sched (virt),  
Event channel,  
Yields / Blocks


▲ Sched (kernel),  
Perf (PMU),  
Subsys,  
Kernel data

(sched, ev, yield, .... lat)  
per activity data

## Attribution

Measurements in virtualized  
environment

Gaussian Process  
regression model



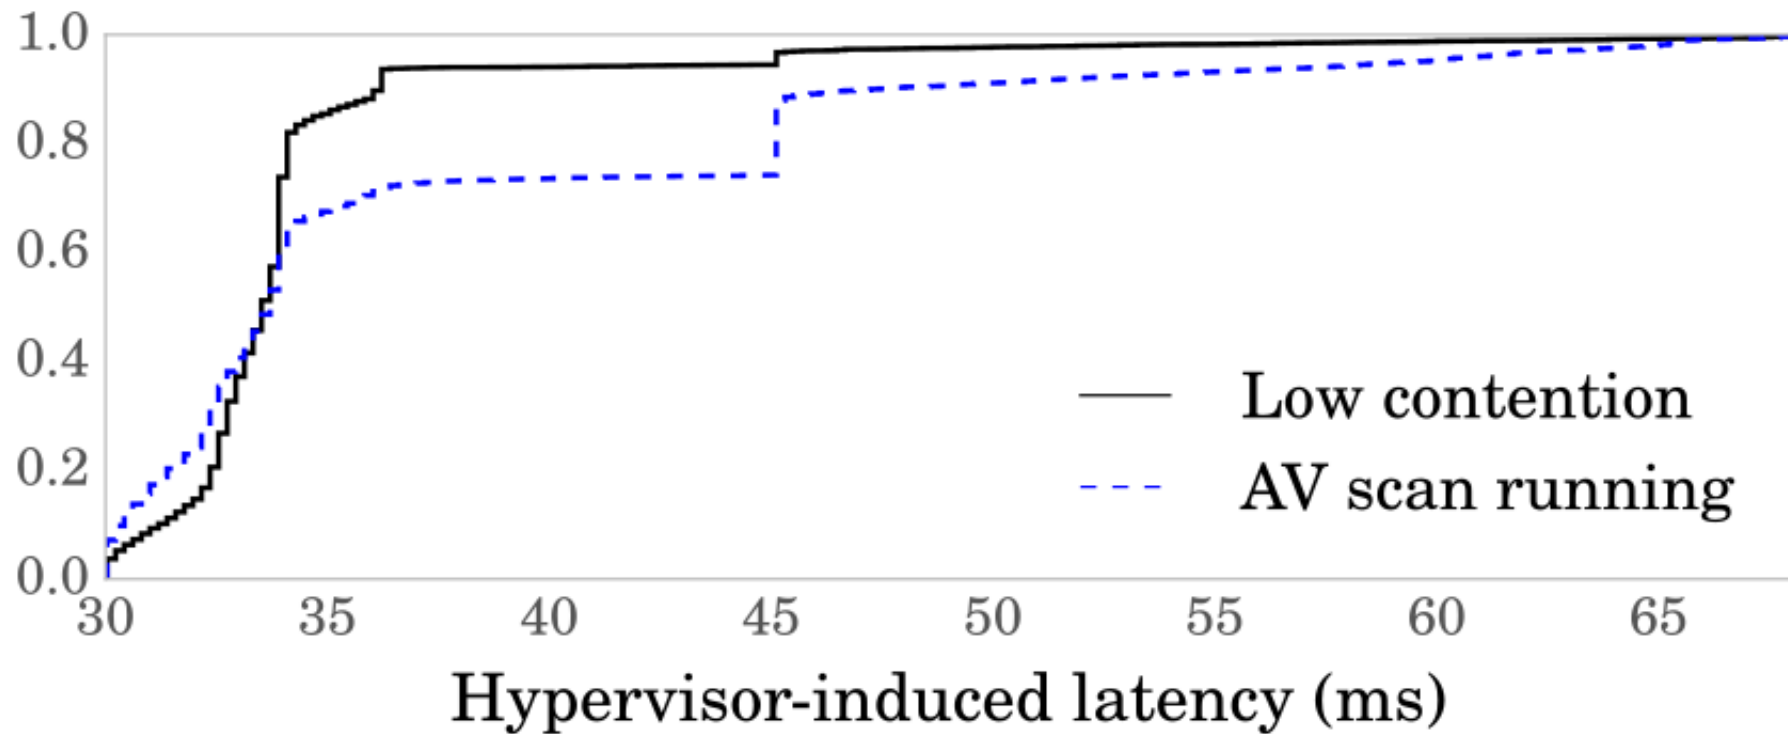
(**virt-influence**, confidence)  
prediction

Hypervisor-induced latency  
Containerisation-induced latency

# Results

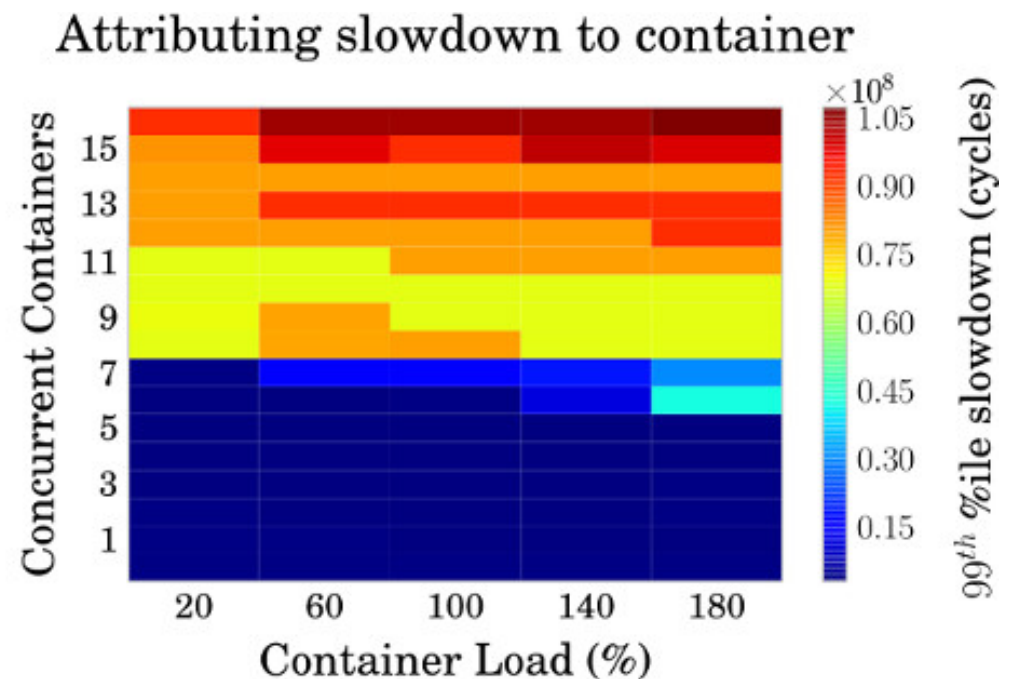
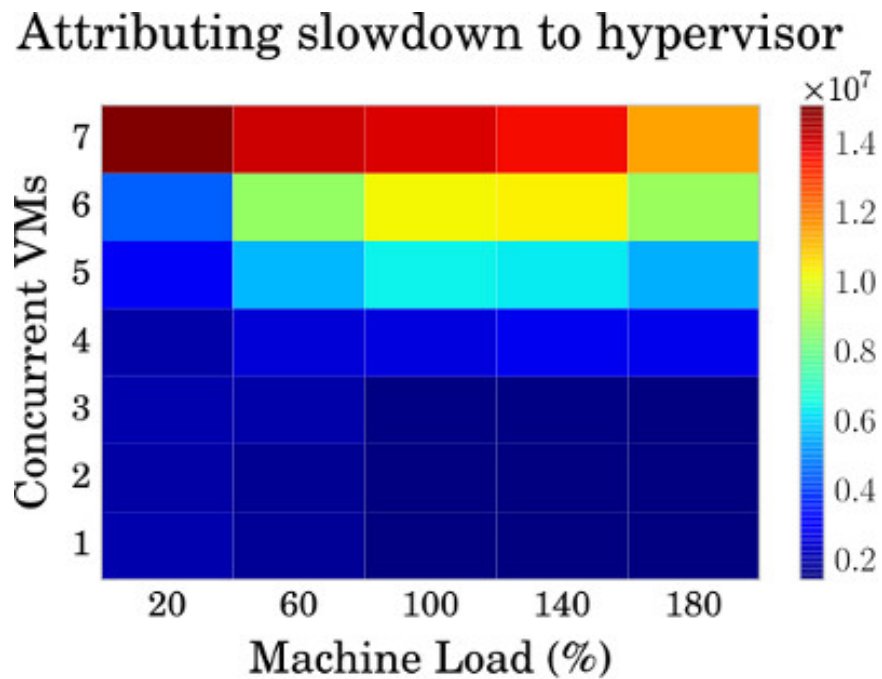
- Anti-Virus Scan performed in DOM0

CDFs of hypervisor-induced latency



# Results

- Systematically varying load and contention



-





# Discussion

- Moving to multi-hop requests / actions
- Automating app instrumentation
- Cloud provider transparency
- Finer-grained charging?



[www.cl.cam.ac.uk/rscfl](http://www.cl.cam.ac.uk/rscfl)

*lucian.carata@cl.cam.ac.uk*  
*dtg-resourceful@cl.cam.ac.uk*



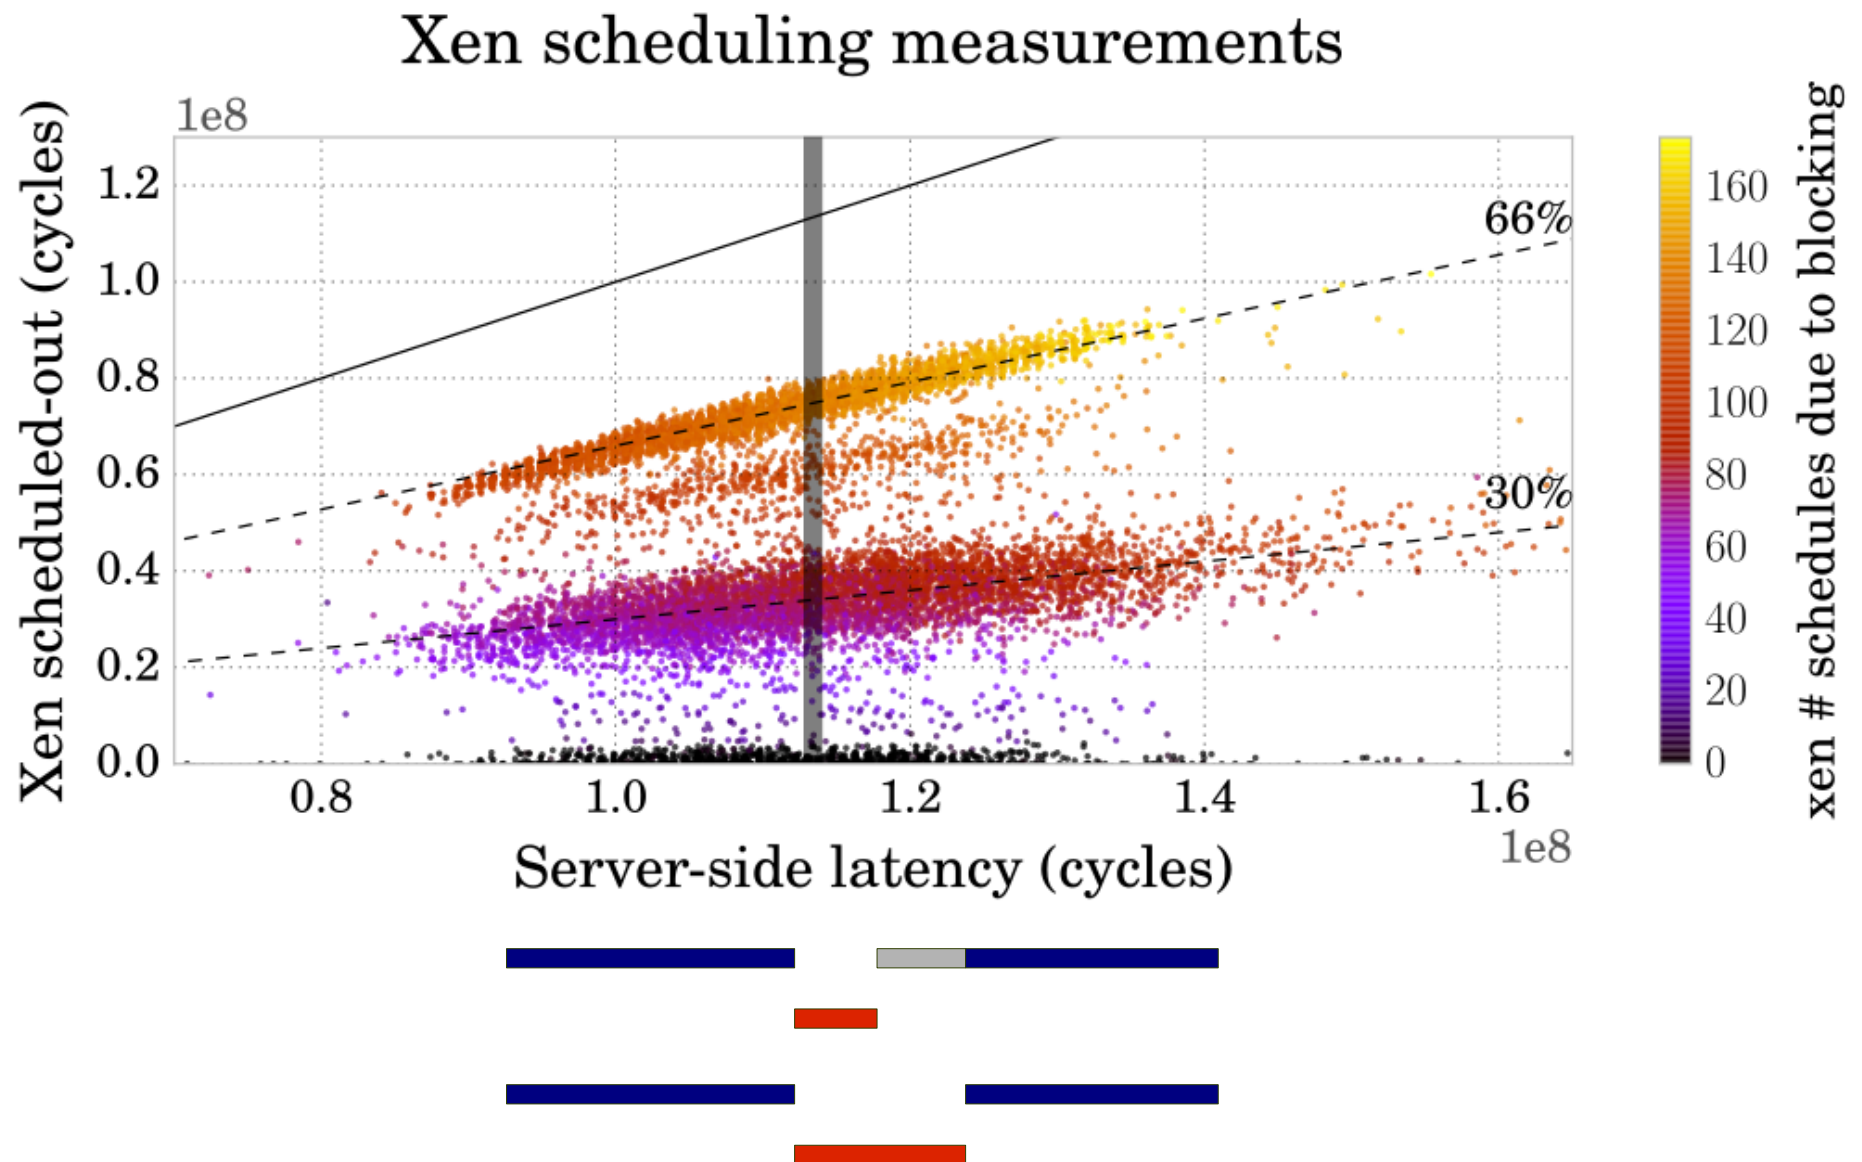
---

# Discussion Slides

---



# Measurement



# Measurement

