

The Heisenberg Measuring Uncertainty in Lightweight Virtualization Testbeds

Quan Jia, Zhaohui Wang, Angelos Stavrou

Computer Science Department, George Mason University



Outline

- Background
- Architecture
- Methodology
- Experiment Set Up
- Experiment Result
- Conclusion
- Future Work

The Heisenberg Uncertainty Principle

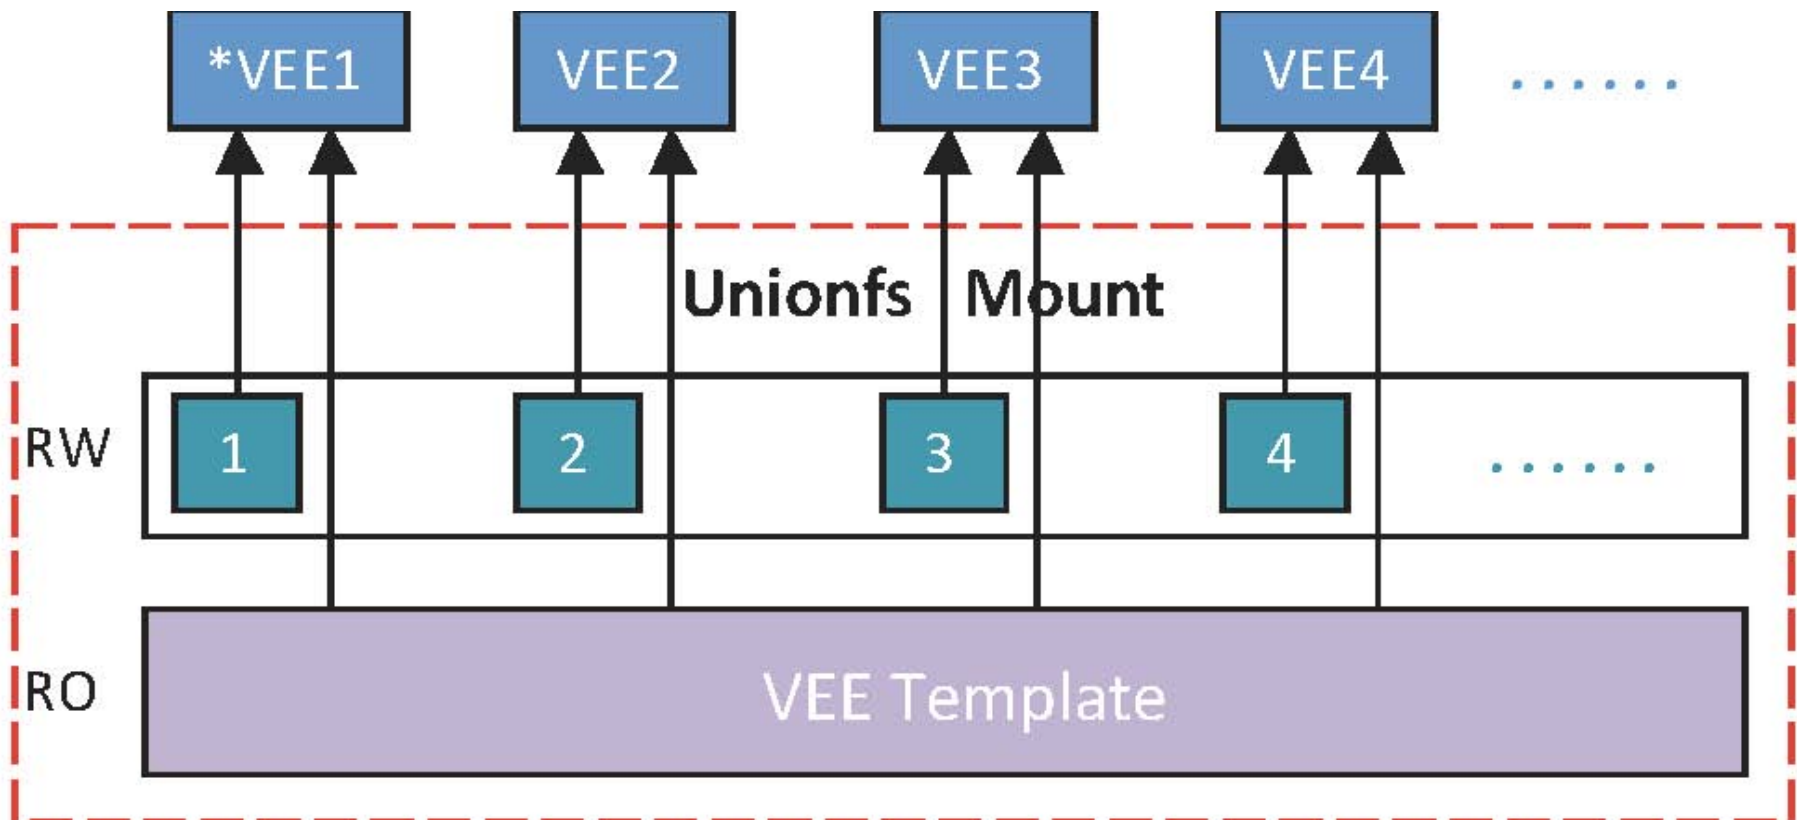
$$\Delta X \Delta P \geq \frac{\hbar}{2}$$

In Computer Science?

Increasing the frequency of measurement can decrease the number of concurrent containers we can utilize without interfering with the performance of the experiment itself.

- When testbeds scale up to hundreds of lightweight containers
 - ▶ how system resources managed
 - ▶ how the containers interact with each other
 - ▶ Is the fidelity maintained

The architecture



*VEE: Virtual Execution Environment, also known as container

Technical Approach

- Analysis of the memory consumption by container
 - ▶ Shared and non-shared memory pages
 - ▶ Disk consumptions
- Run-time analysis
 - ▶ Number of containers running on the same host
 - ▶ Different sampling interval
 - ▶ Statistics from /proc
 - ▶ Measure CPU, memory, network load and observe system behavior

Static Analysis of Memory Consumption

Processes/kB	RSS	shared	non-shared
init	860	604	256
syslogd	640	508	132
dbus-daemon	684	508	176
sshd	992	644	348
Sum		2264	912

Table:ubuntu-8.04-i386-minimal container template process memory consumption in kB

Experimental Setup


□ Client side

- Dell PE 1950 2.66GHz
- OpenVZ containers + Unionfs stackable filesystem
- Run experiment with 100-1400 containers
- 1 wget process per container.
- 1-10s random sleeping time between requests
- 4 sampleing intervals (0.1,0.01,0.005,0.001)

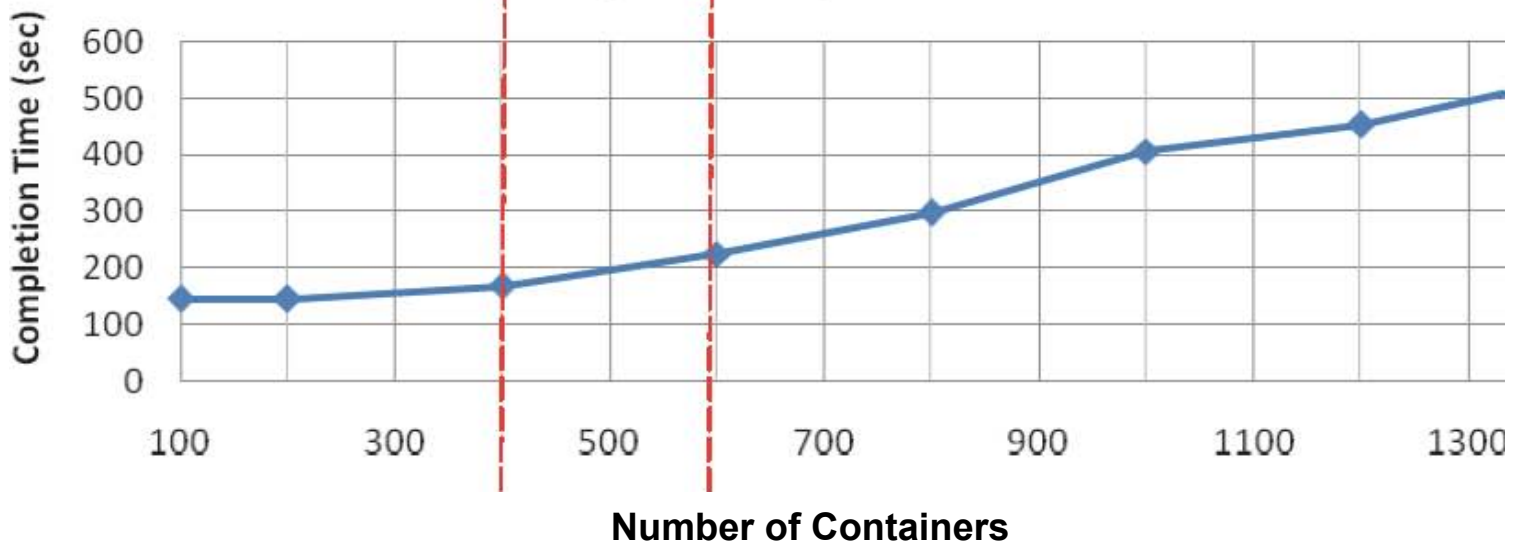
□ Server side

- Dell PE 1950 2.66GHz
- Apache 2.2.9 on Fedora 9
MaxClient : 4000
Gigabit LAN connection


Available System Memory



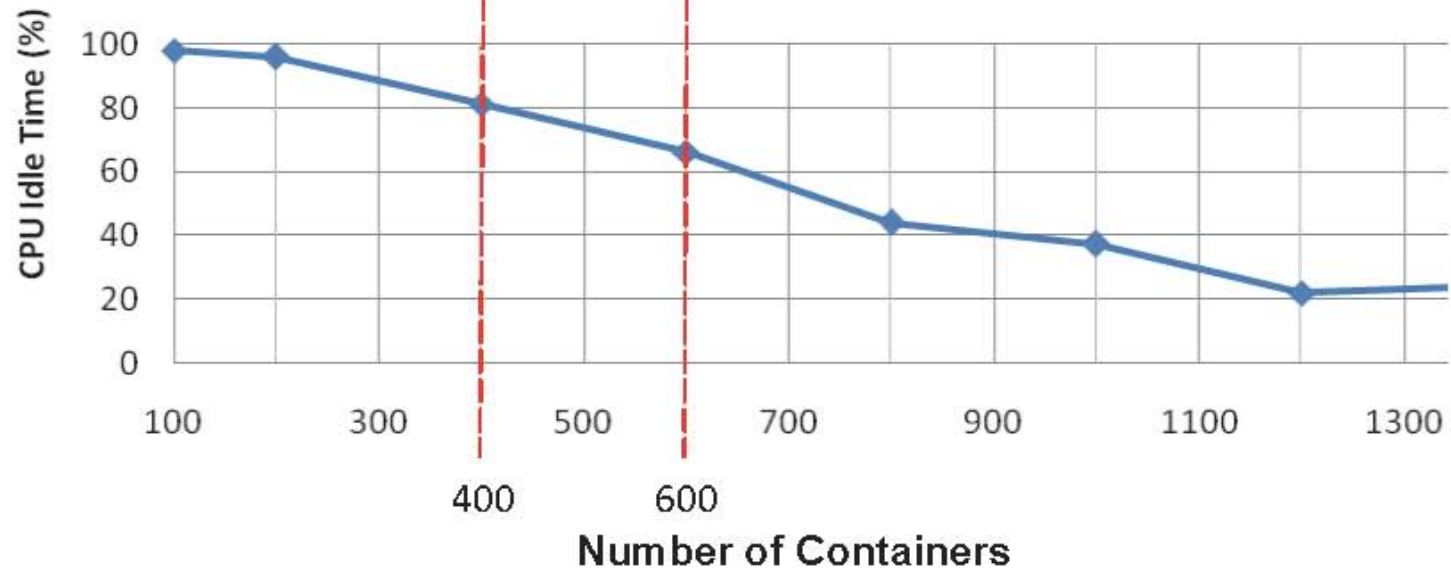
Average Completion Time




Average Network Throughput




Average CPU Idle Percentage



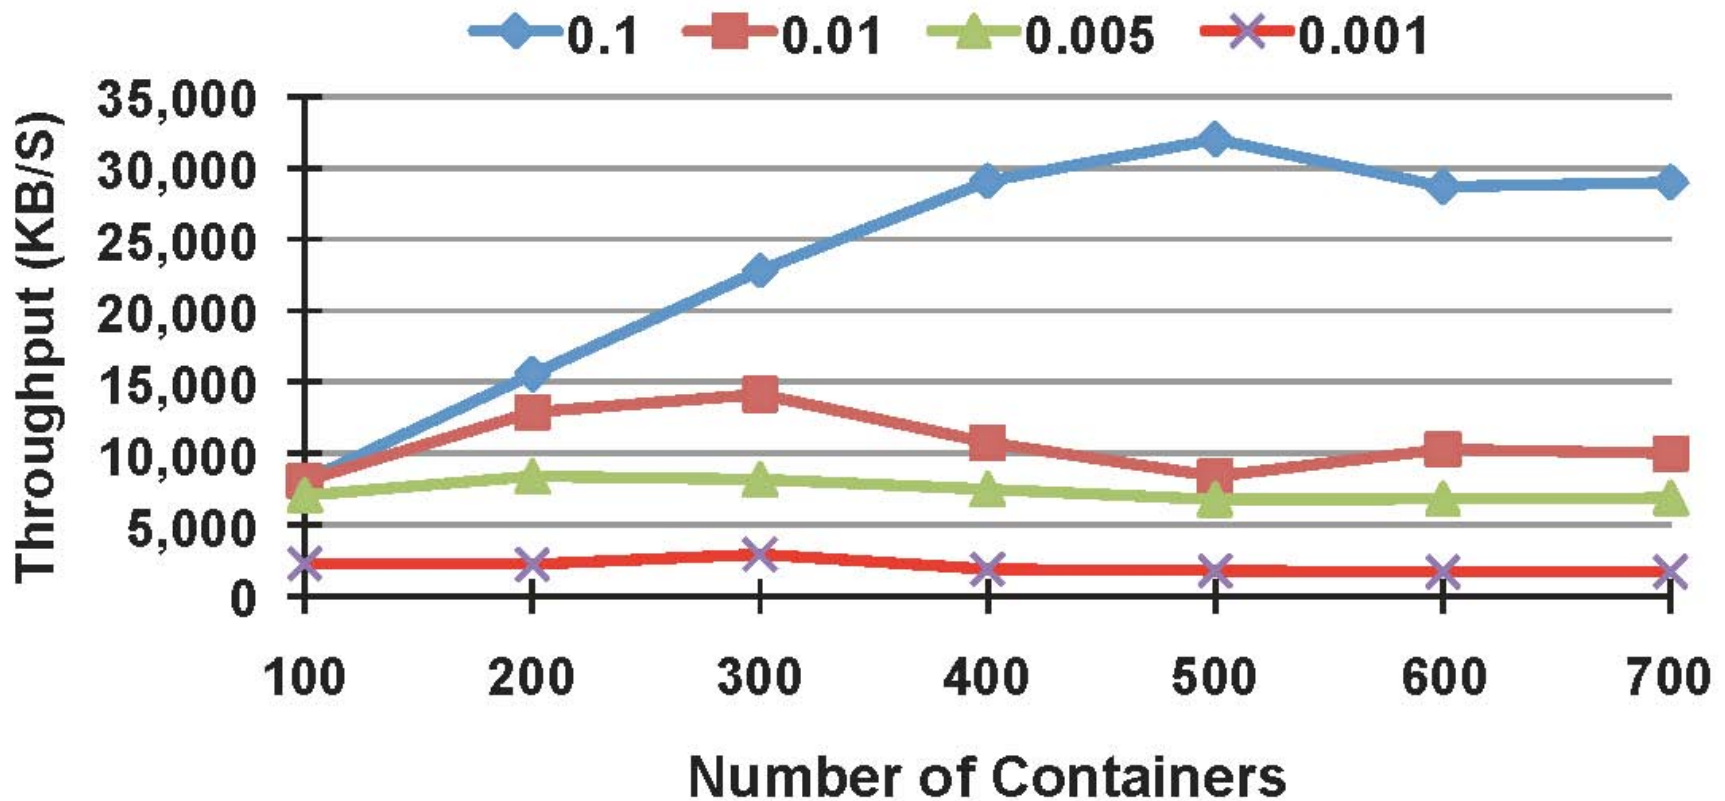
Experiment Completion Time w Load



Container Average CPU Utilization



Container Aggregate Network Throughput



Contribution and Limitation

- First time unveiled uncertainty problem due to system resource contention in a lightweight virtualization environment
- The desirable accuracy of measurement is largely dependent on the high sampling frequency, which potentially deprives the containers of available resources and adversely interferes with the experiment
- No quantitative analysis or formula yet

Conclusion

- Virtualization has some limitations due the sharing of host resources (CPU, network, memory and disk) among same host VEEs
- The Heisenberg uncertainty principle for host resource measurements: increasing the precision and fidelity of the resource measurements can interfere with the behavior of the experiment
- It is not a trivial task to determine the maximum number of VEEs that can be run concurrently in a physical machine without perturbing the experimental outcome

Future Vision

- More experiments to delve deeper and investigate the kernel data structures to achieve fine-grained resource management

- Modeling the host environment, guide the implementation of virtualization