### QuartzV: Bringing Quality of Time to Virtual Machines

### <u>Sandeep D'souza</u> and Raj Rajkumar Carnegie Mellon University





IEEE RTAS @ CPS Week 2018

## A Shared Notion of Time

- Coordinated Actions
- Ordering of Events



### A Shared Notion of Time is useful

 $\rightarrow$  Replace Communication with Local Computation\*

**Carnegie Mellon University** 

\*Liskov, *Distributed Computing* '93

### **Coordination in Space and Time\***



```
<sup>3</sup> D'souza et al., HotCloud '17
```

## Cyber-Physical Edge-Cloud

- The cloud is *key* to enable geographical scaling
  - data storage
  - host the intelligence behind CPS
- Low-latency requirements of CPS
  - Safety-critical + real-time performance
  - A hierarchy of edge and cloud deployments <sup>#</sup>



### Virtualization is *key* for utilizing cloud and edge resources

## Virtualization

- OS-level (Containers)
- Hypervisor-based (VMs)
  - CPU Instructions
    - Emulation
    - Hardware-Acceleration
  - Peripheral Access
    - Full Virtualization
    - Para-Virtualization



Hardware-Accelerated CPU Instructions + Para-Virtual Peripheral Access  $\rightarrow$  near-native performance for VMs

## Outline

- Motivation
- Background
  - Quality of Time (QoT)
  - Timelines
  - QoT Stack for Linux
- QuartzV
- Experimental Evaluation
- Conclusion

# Quality of Time (QoT)\*

- Quantified
  - using clock parameters:
    - accuracy, precision, drift....
  - w.r.t a *reference clock* (time)
- Each timestamp has high-probability bounds
  - Timestamp  $\epsilon \{t-\varepsilon_{\mu}, t+\varepsilon_{h}\}$



# The *end-to-end* uncertainty in the notion of time delivered to an application by the system

## QoT and Fault Detection

- Failure Scenario:
  - Clock Synchronization degrades
  - **Reported QoT** must reflect degradation
- Application-specific *failover* mechanisms
  - Physical and Analytical Redundancy



QoT can enable fault detection

 $\rightarrow$  *fault-tolerant* coordination in CPS

## **Enabling Coordination at Scale**

- Timeline\*: Virtual time reference
  - time-based coordination service
  - Platform-independent API
- Coordinated actions, distributed components
  - o all components *bind* to a timeline
  - each *specifying* its required QoT



### Timelines *abstract* away clock synchronization

### **QoT Stack for Linux\***



Support for ARM and x86 platforms

open source, modular implementation, no change to the Linux kernel

**Carnegie Mellon University** 

10

## Outline

- Motivation
- Background
- QuartzV
  - Challenges
  - Para-Virtual Design & Implementation
  - Full-Virtualization Support
- Experimental Evaluation
- Conclusion

### Virtualization and Time

- Higher clock-read and interrupt latencies\*
  - Overhead of additional abstraction layers
- Delivering application-specific requirements
  - different **QoT requirements**
  - different time scales



### Providing *near-native* timing performance is a challenge

### **Quartz-V: Challenges**

- Low-Latency Clock Reads + QoT estimates
- Supporting multiple VMs  $\rightarrow$  multiple timelines
- Maintaining VM Isolation
  - prevent *malicious VMs* from affecting correct timing
- Portability
  - o avoiding modifications to the kernel, hypervisor & application source

### Objective: *near-native* performance while *maintaining* isolation and portability



### Quartz-V: Key Ideas

- Extends the QoT Stack for Linux
  - **QEMU-KVM** hypervisor extensions for Virtual Machines
- Para-Virtual (PV) approach: "Time-as-a-service"
  - **PV guests specify** Timeline and QoT requests
  - Host performs clock sync + QoT Estimation
    - transfers the sync parameters to the guest
  - **PV guests read** a **para-virtual** host clock (KVM-Clock)
    - compute the "timeline" time using host-provided parameters

$$\begin{array}{c} tl_{now} = tl_{last} + tl_{skew} * (core_{now} - core_{last}) \\ \hline current & offset & clock & time passed since the clock \\ time & skew & was last synchronized \end{array}$$



### qot\_virtd: QoT Virtualization Daemon

- qot\_virtd is the interface between:
  - QoT-aware *applications* in the guest VM, and
  - QoT Stack *system-services/kernel-module* on the host



qot\_virtd acts as a server for client QoT-aware applications in a VM

## Specifying Application Requests

- *not* on the application critical path
  - reliable, need not be low latency
- virtio\_serial
  - *bi-directional* serial communication
  - scalable across multiple VMs



### **Provide Clock + QoT Parameters**

- Clock parameters required
  - to *compute* the current time
- Clock reads are on the critical path
  - reliable and low latency
- Inter-VM Shared Memory (ivshmem)
  - *Host clock-sync service* writes parameters
  - Guest Applications can read parameters
    - reading memory is *low latency*
    - VMs have read-only access, maintains isolation

QuartzV provides *near-native* performance

while *maintaining* isolation and portability



## Fully-Virtual QoT Stack

- Some VMs/Hypervisors
  - do not support paravirtualization
- The entire QoT Stack inside a VM
  - Fully-Virtual deployment
  - clock sync achieves *lower accuracy*
  - clock reads have higher latency



The Fully-Virtual QoT Stack can provide full functionality,

but lower performance w.r.t. QuartzV

## Outline

- Motivation
- Background
- QuartzV
- Experimental Evaluation
  - Performance Evaluation
  - Scalability Evaluation
  - Prototype Industrial Robotics Application
- Conclusion

## **Clock-Synchronization Accuracy**

- Comparison using PTP (qot\_ptp)
  - 8-core i7 CPU, with 16GB RAM
  - Clock Evaluation Test-Bed
- QuartzV
  - yields near-native accuracy
  - clock sync on the host
- Fully-Virtual QoT Stack
  - yields *lower* accuracy
  - clock sync in the VM
  - overhead of virtualized network stack



### QuartzV can provide *near-native* clock-sync accuracy

### **CPU-stress Scalability Results**



QuartzV is more *resilient* to CPU-intensive applications

### **Network-stress Scalability Results**



QuartzV is more resilient to malicious network-intensive VMs

### **Network-Regulation Scalability Results**



Bandwidth Regulation can *prevent* malicious network-intensive VMs from degrading clock-synchronization accuracy

## QuartzV: Clock-Read Latency Scalability

- Concurrent VMs performing clock reads
- QoT Clock v/s CLOCK\_MONOTONIC
  - higher latency,
  - due to *applying* projection parameters,
  - and QoT calculations
- Slight increase in latency with new VMs
  - contention in *reading* the hardware timer
  - same trend for both clocks



Shared Memory supports *simultaneous* read-only access

 $\rightarrow$  **Bottleneck-free** clock-read implementation

### **Coordinated Industrial Robotics**



### **Conclusion and Future Work**

- Geo-Distributed CPS: "Coordination at scale using time"
- QoT-awareness enables Fault Detection
- Quartz-V: Adds QoT-awareness to VMs
  - uses the *para-virtual* approach
    - *delivers near-native* timing performance
    - maintains isolation and scales to multiple VMs
    - resilient to malicious resource-intensive VMs
- Fully-Virtual QoT Stack for hypervisors/VMs with no para-virtual support
- Future Work
  - Building a Geo-Scale QoT-based CPS Coordination Framework
- **Open source:** https://bitbucket.org/rose-line/qot-stack/src

QuartzV can enable hosts of new geo-distributed coordinated Cyber-Physical Systems



# Thank You ! Questions ?

Sandeep D'souza sandeepd@andrew.cmu.edu





### **QoT Estimation Results**



Estimated QoT (99.999% confidence)

*bounds* the measured clock-synchronization error