

# Achieving High Throughput by Transparent Network Interface Virtualization on Multi-core Systems

Huiyong Zhang, Yuebin Bai, Zhi Li  
Niandong Du, Wentao Yang

Beihang University, China



# Outline

1. Why Transparent Network Interface Virtualization?
  2. Motivation
  3. Xen network design changes and optimizations
  4. Performance results
-

# Network I/O Virtualization

- **Software-based approaches**

- Paravirtualization I/O model(Xen)
- Device emulation(Hyper-V, KVM etc.)

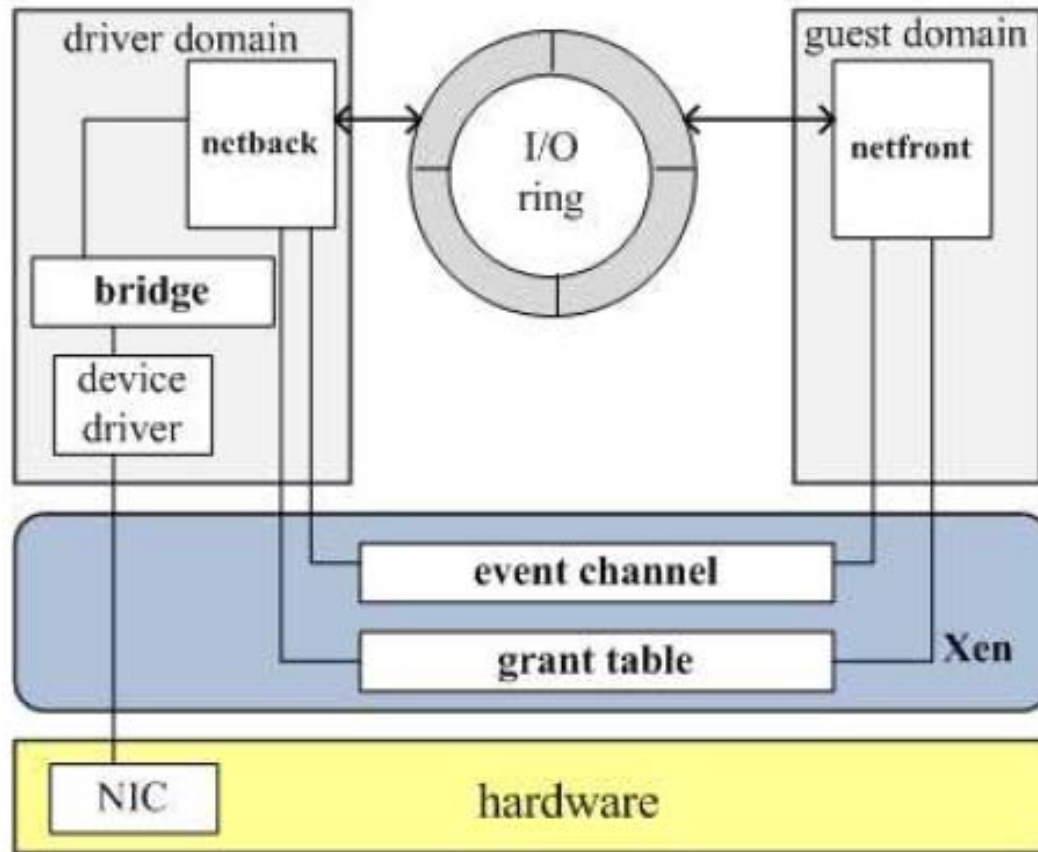
- **Hardware-based approaches**

- Direct I/O(VT-D, IOMMU)
- Self-virtualization devices(SR-IOV)

# Transparent Network Virtualization

- Live migration
- Fault isolation
- Minimize the VM
- ...

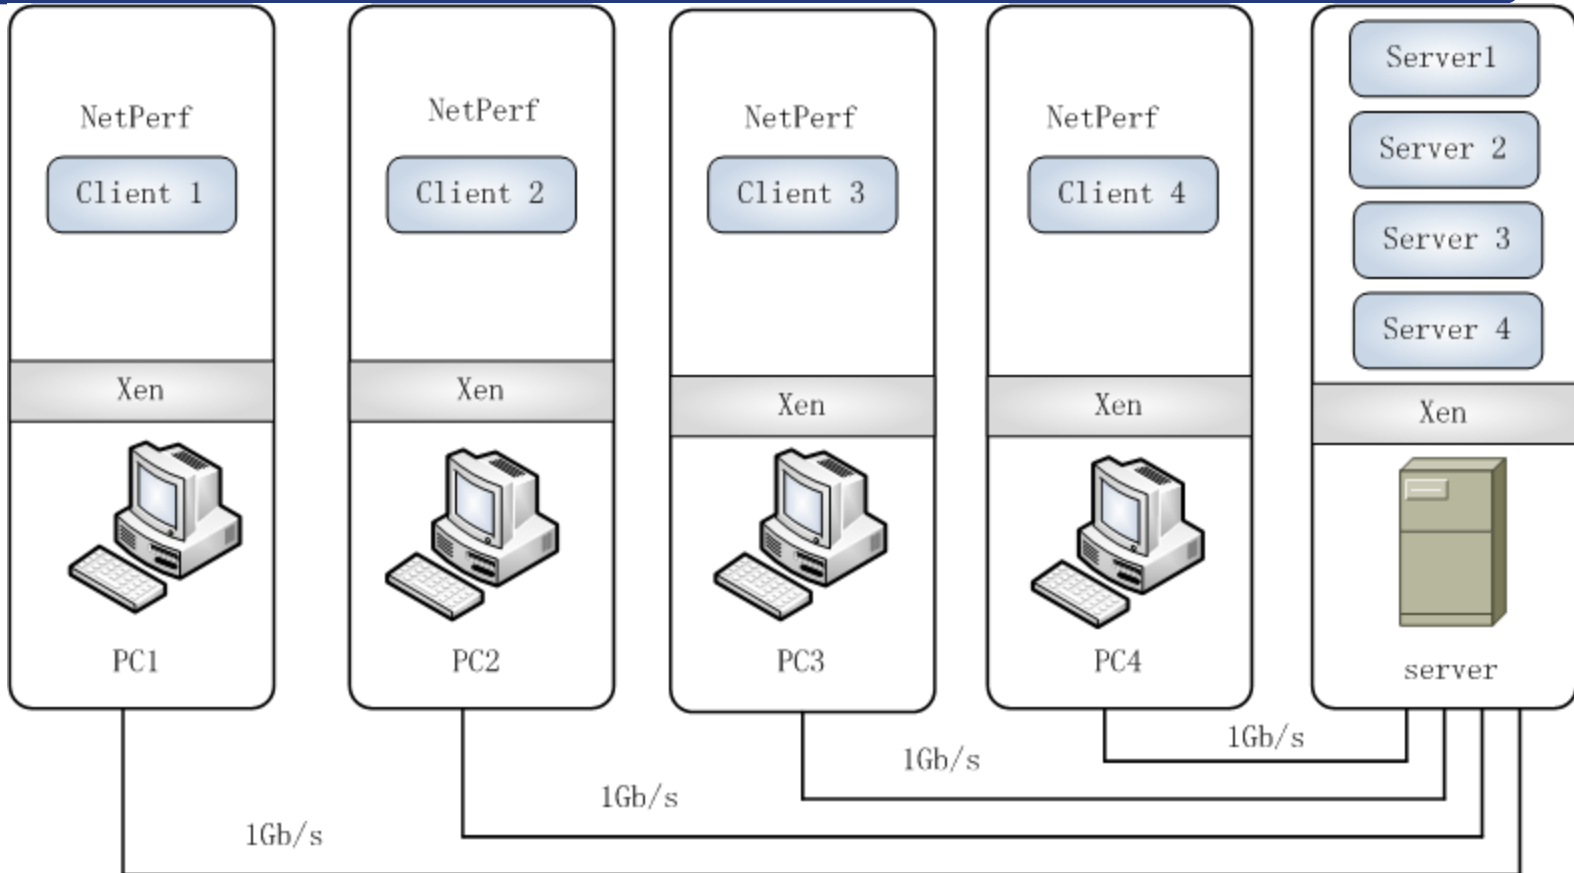
# Xen Network I/O



## I/O Channel

- I/O ring
- Event channel
- Grant table

# Experimental Setup

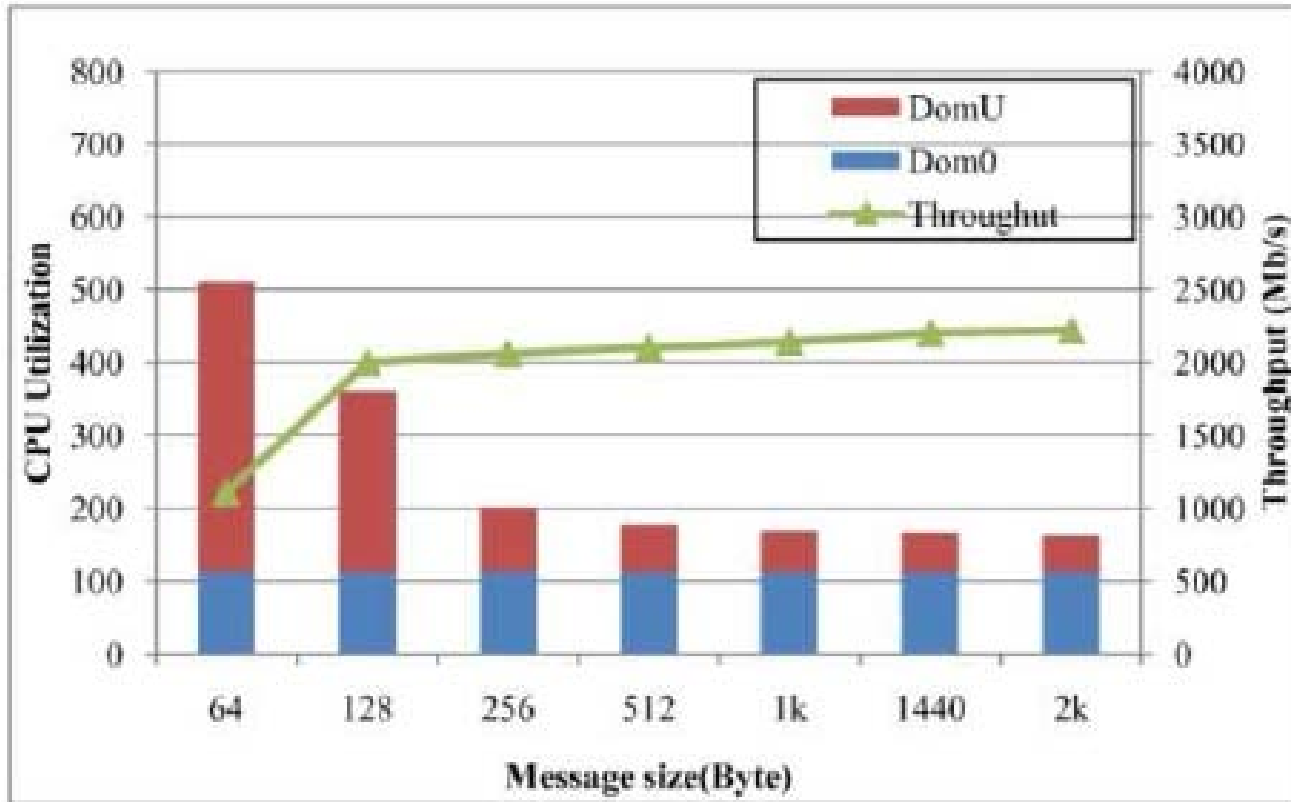


Physical server: Xeon 5405(x2) 4-port 1Gb/s NIC

Virtual Server: 1VCPU 512MB memory

Dom0: 4VCPU 2GB memory

# Xen I/O Performance



**Only achieve 2.3Gb/s throughput**

# Performance Analysis

## Dom0 CPU utilization

```
debian:/home/zhanghuiyong# xm vcpu-list
```

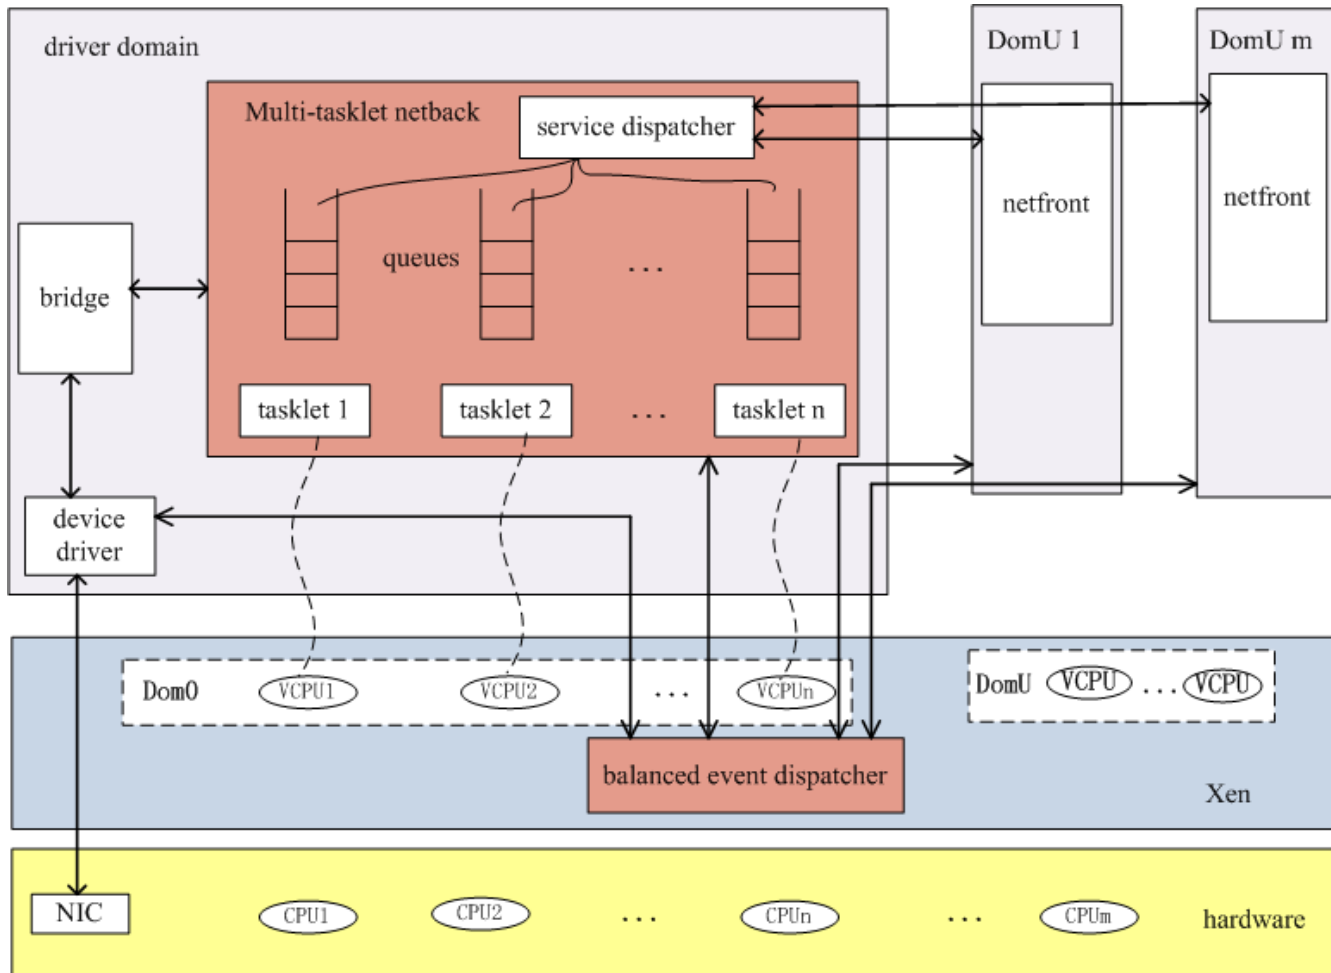
Name	ID	VCPU	CPU	State	Time(s)	CPU Affinity
Domain-0	0	0	1	-b-	820.7	any cpu
Domain-0	0	1	2	r--	21.9	any cpu
Domain-0	0	2	3	-b-	17.5	any cpu
Domain-0	0	3	0	-b-	25.6	any cpu

### ●Reason

- All interrupts are delivered to VCPU0
- Only one-pair tasklets in the netback



# Xen network design changes



- Balanced event dispatcher
- Multi-tasklet netback

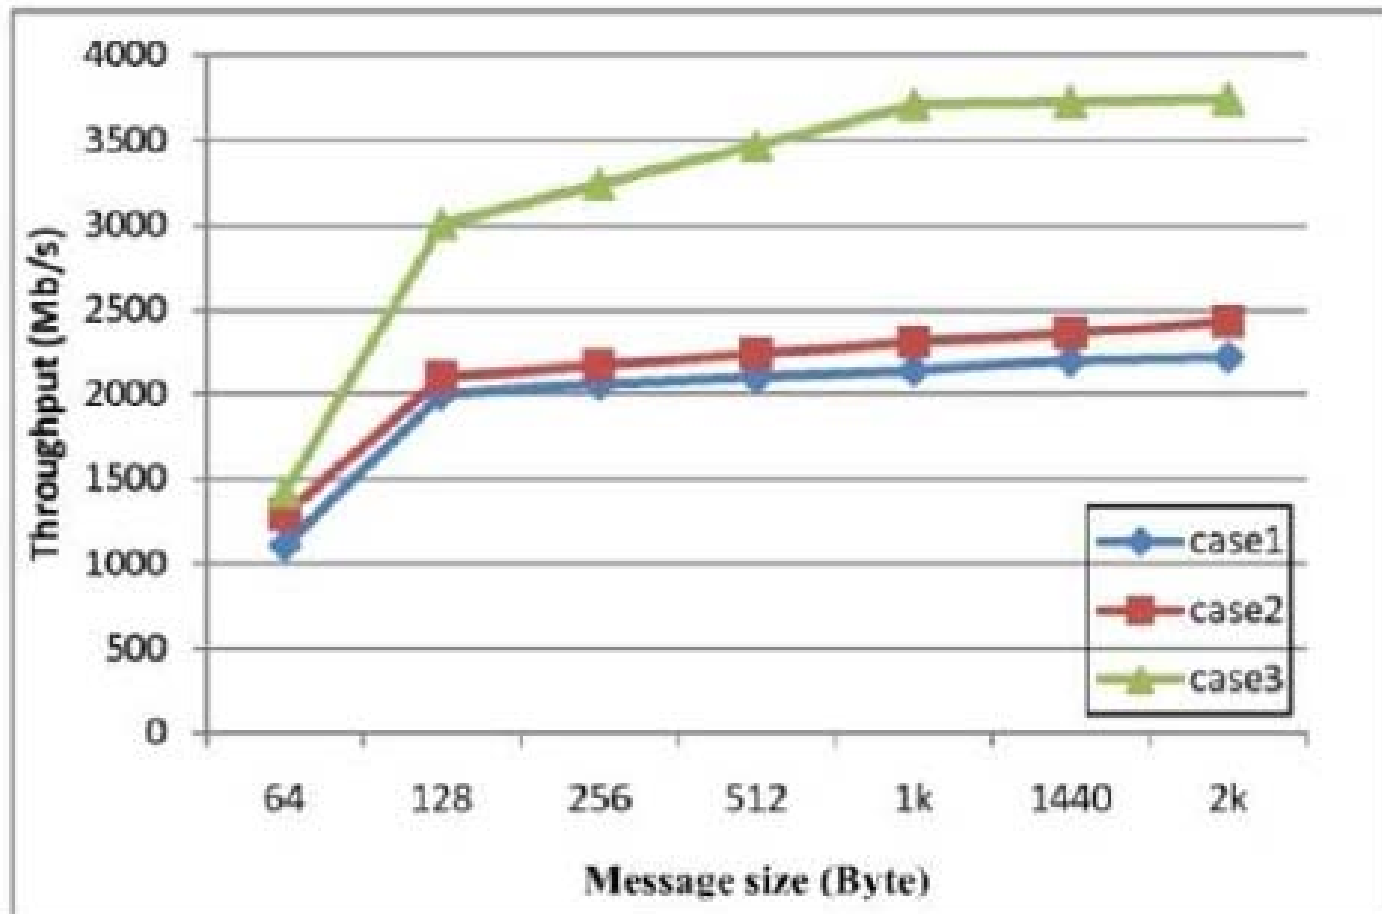
# Balanced Event Dispatcher

- When network traffic is low
  - deliver all events notification to VCPU0  
(better cache behavior)
- When network traffic is high
  - deliver all events notification to all VCPUs  
(to achieve high throughput)

# Multi-tasklet Netback

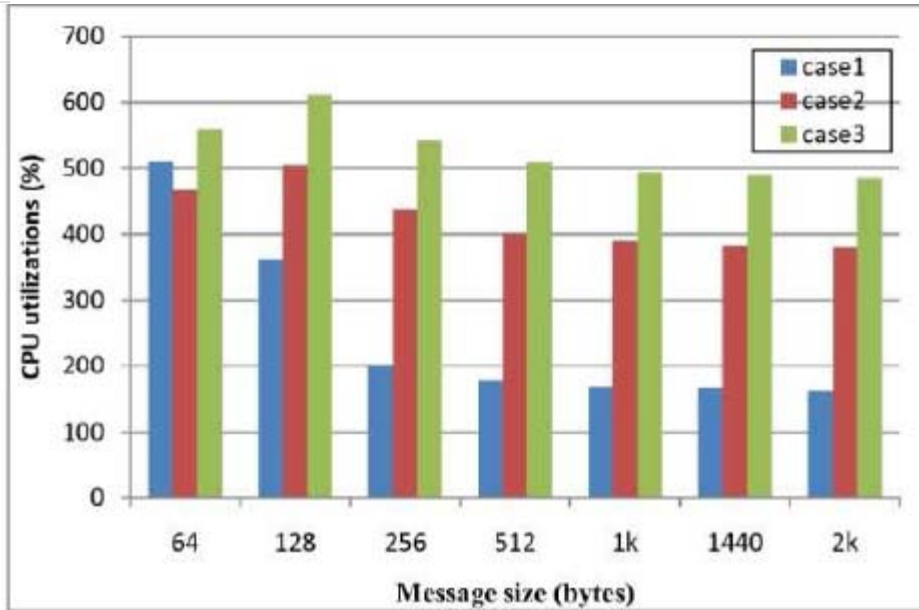
- Define N pairs tasklet, where  $N = \text{Dom0 VCPU \#}$
- Each tasklet has its own tx\_queue and rx\_queue
- Assign each virtual network interface to a tasklet, According to the event dispatcher

# Performance




Can achieve 3.7G/s, improved by 68%

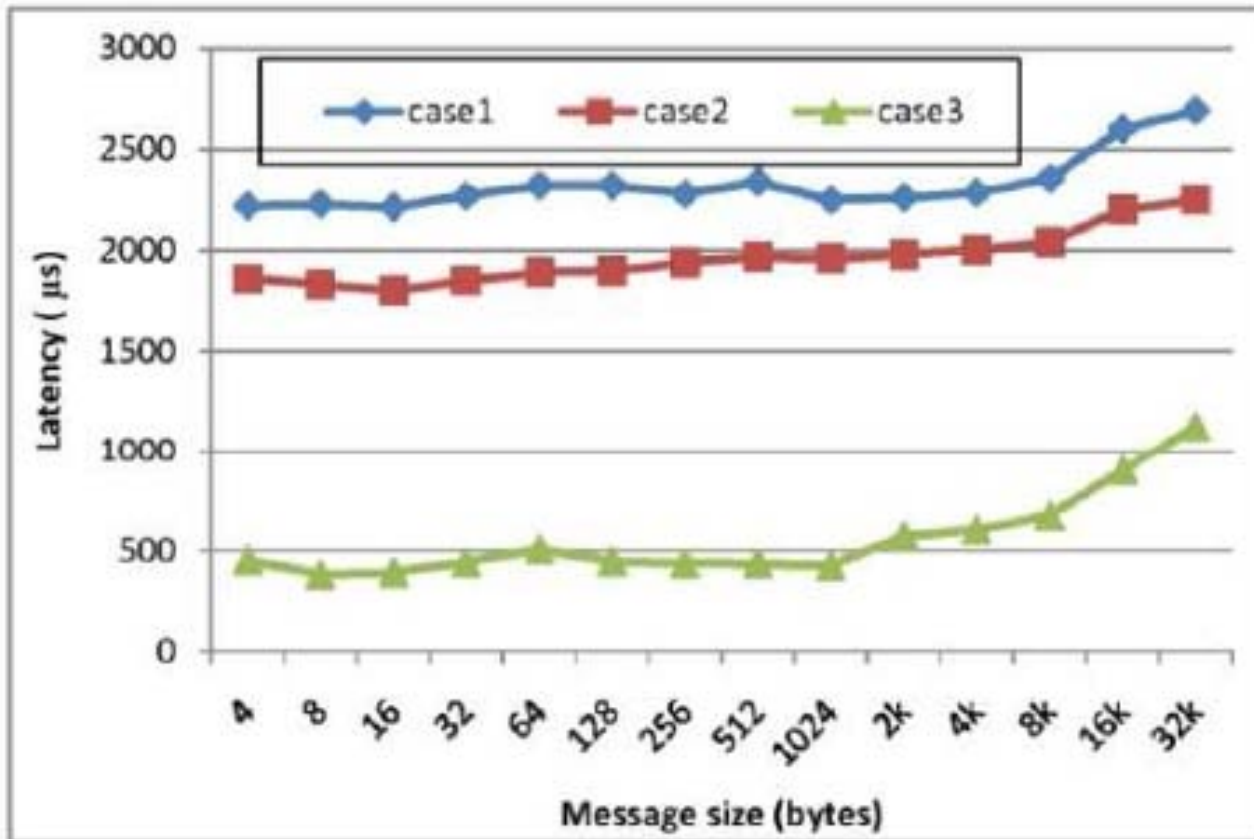
# Performance



CPU resources  
are utilized  
more effectively



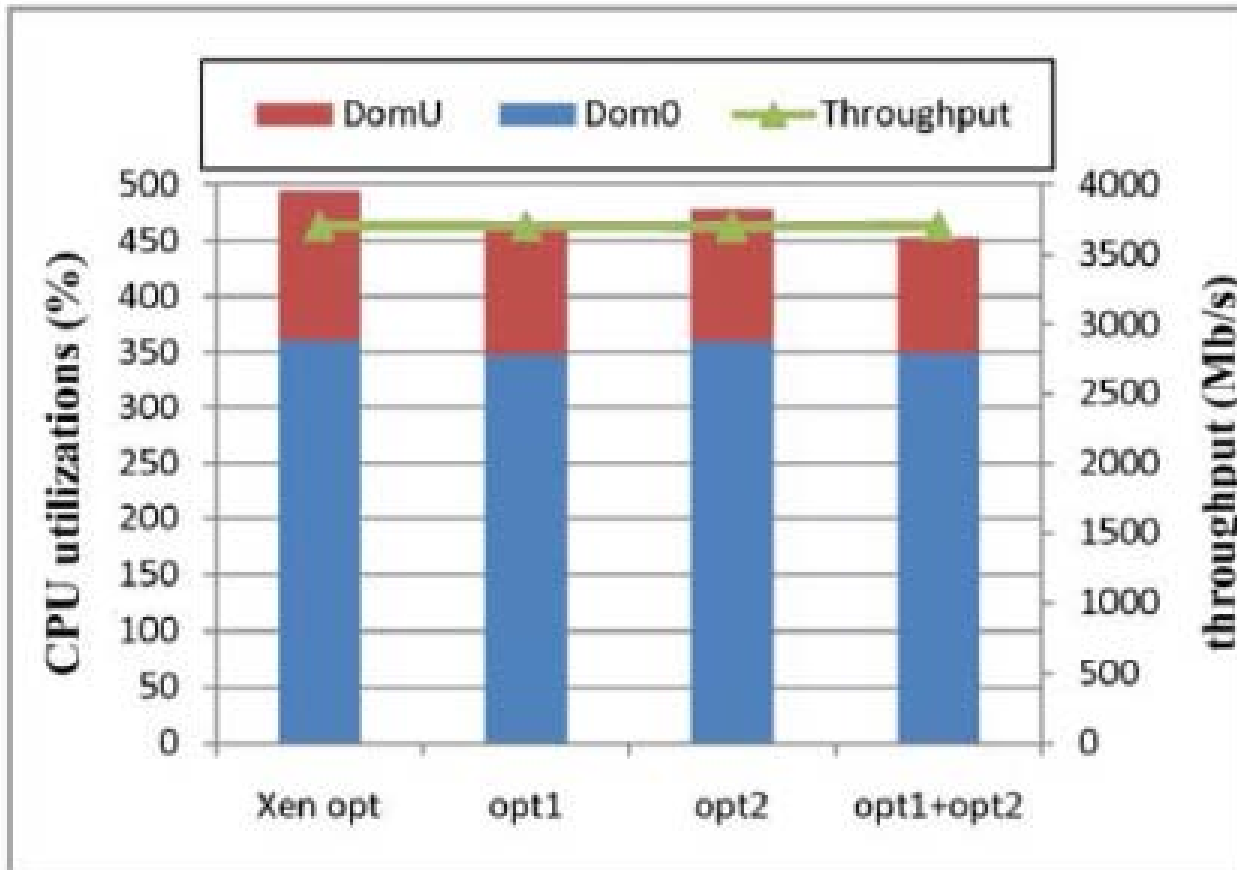
# Performance



Improve the TCP ping-pong Latency

# Optimization

- Event Notification Control in Netback
- LRO in Netfront



## **Achieving High Throughput by Transparent Network Interface Virtualization on Multi-core Systems**



**Thanks!**