

Arrakis

CSE 451: Operating Systems Autumn 2013

Irene Zhang
iyzhang@cs.washington.edu

Arrakis

An **OS** that provides **low-latency, high-throughput** network I/O for **datacenter applications** using **hardware virtualization**.

The OS is a bottleneck for many datacenter applications.

Many applications are I/O-intensive,

- memcached, key-value stores
- load balancer and other middleboxes

And datacenters have high-speed networks.

- 10-100 Gb/s
- Infiniband, fiber and other technologies

Some numbers ...

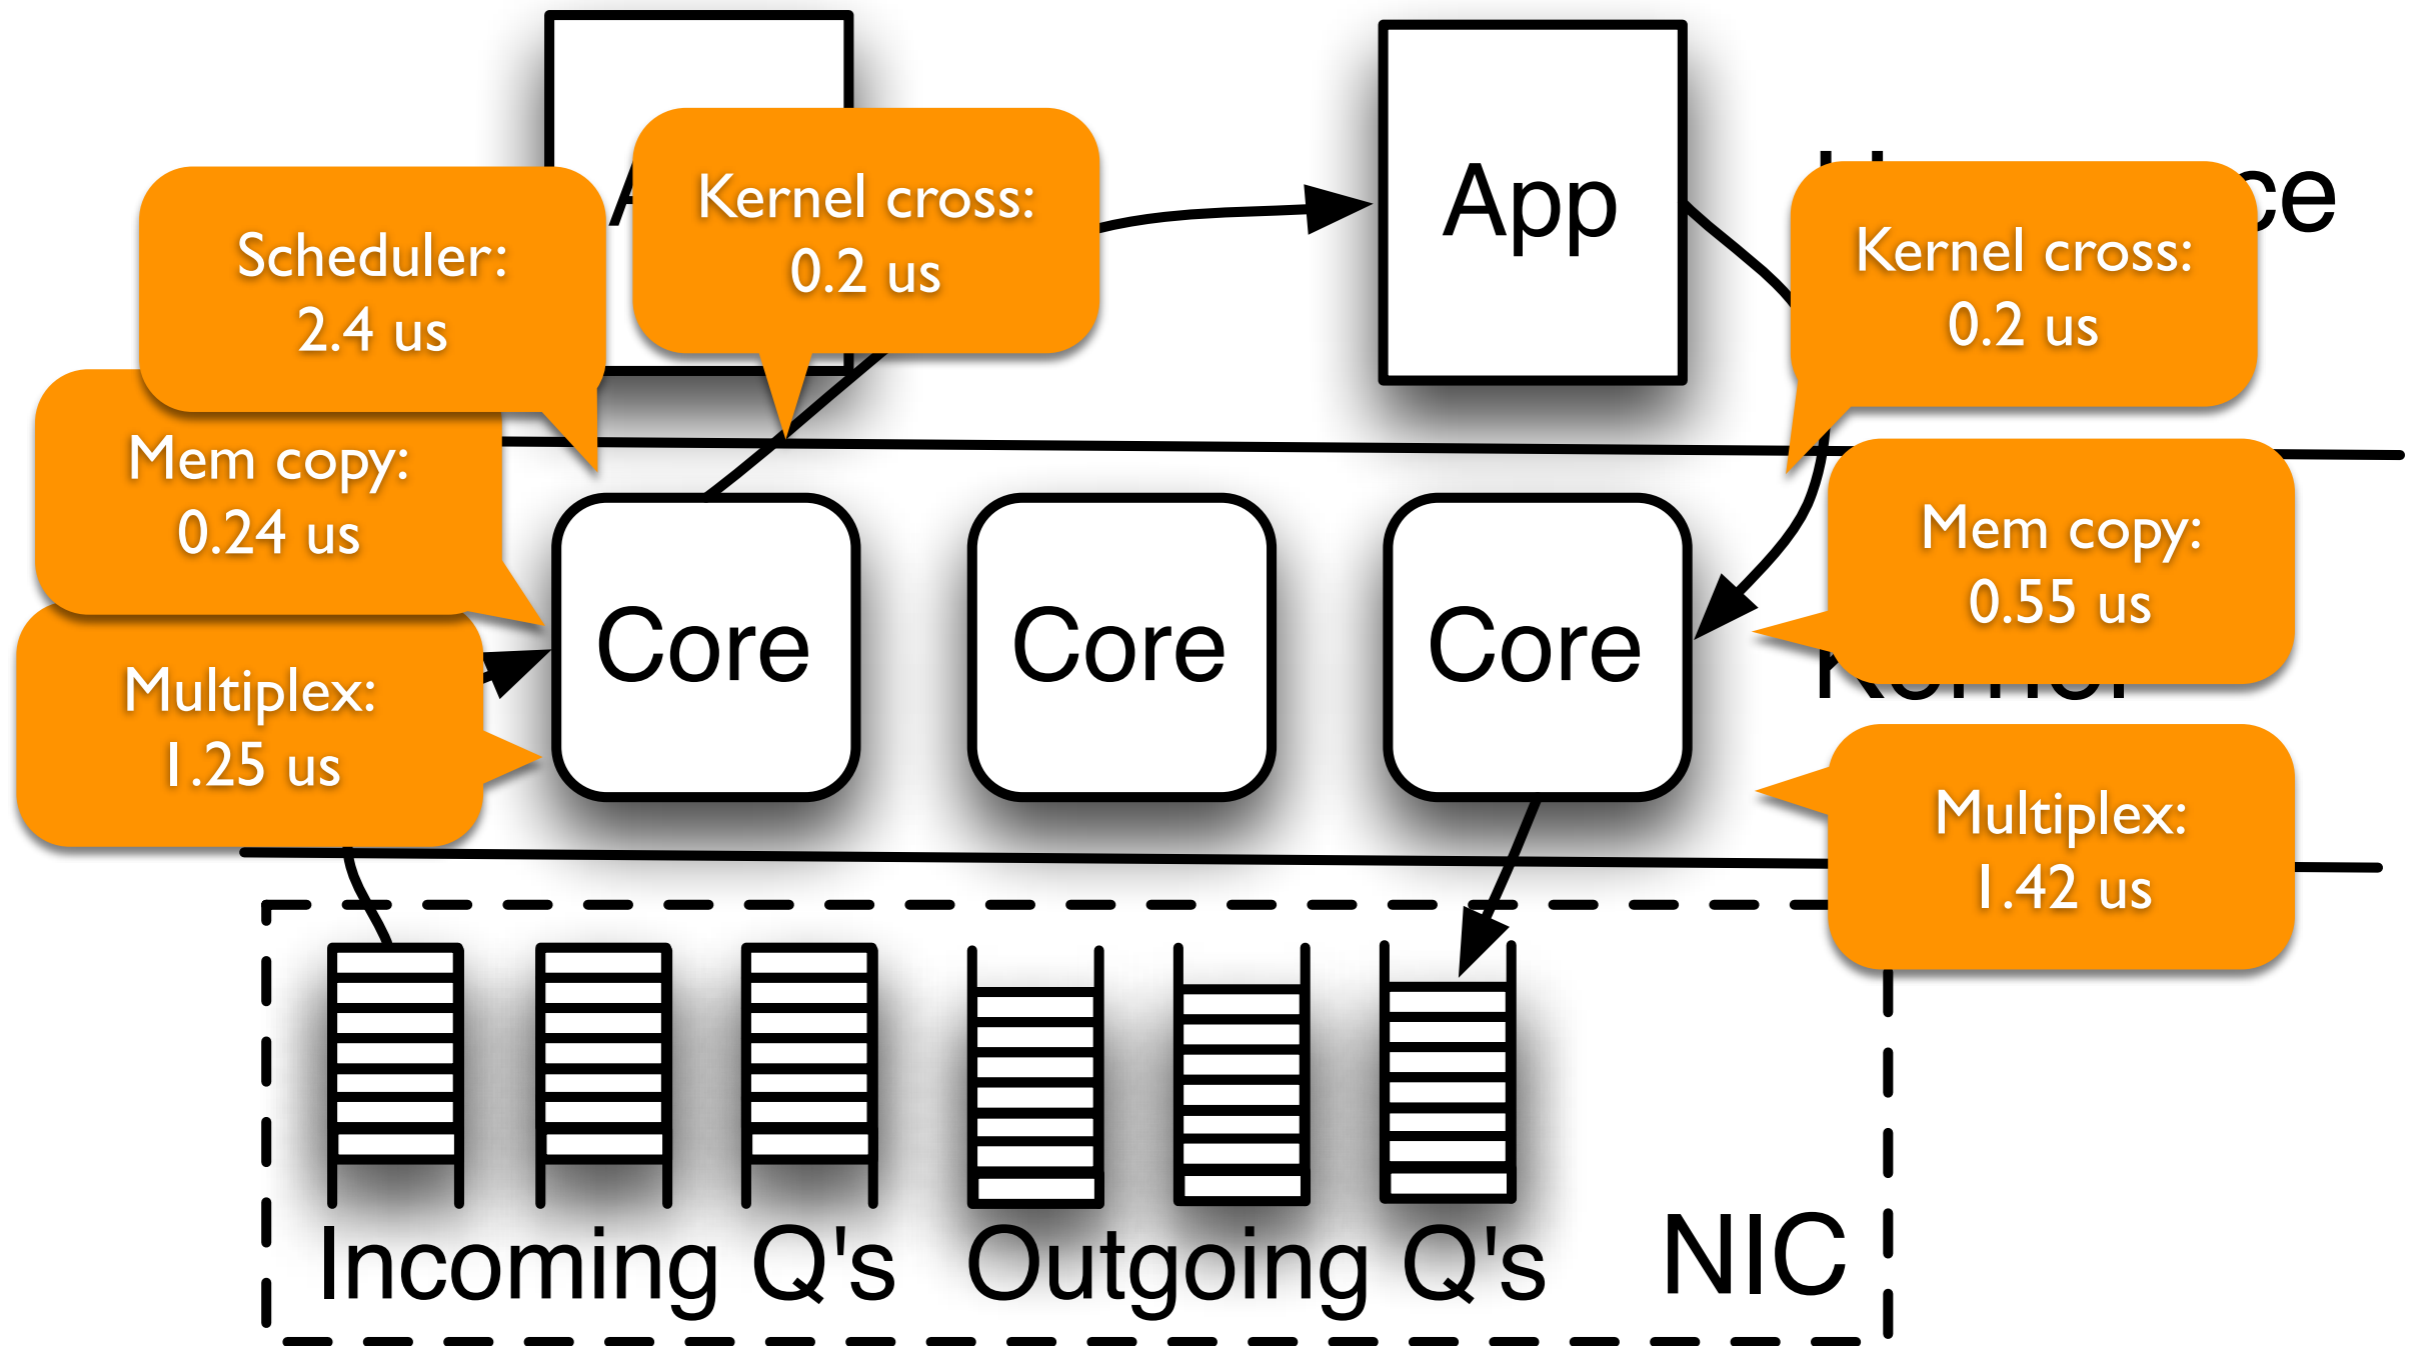
Time through 100 ft of cable: **100 ns**

Time through a 10G switch: **350 ns**

Time for memcached to respond: **1 us**

Time for OS to process a packet: **6 us**

OS overheads



Why is the OS so slow?

The OS provides

- multiplexing
- isolation
- higher level abstractions (sockets)

which have become *relatively* slow.

Key Idea

Hardware virtualization provides fast multiplexing and isolation in the networking card

allowing us to *eliminate* the OS from most network I/O.

Virtualization

- Technology that allows more than one OS to share one physical machine.
- Each OS runs in its own virtual machine.
- Virtual machine monitor multiplexes the hardware among virtual machines.

Network Virtualization

- VMM multiplexes one network card among several virtual machines.
- Each virtual machine gets one virtual network card (vNIC).
- Guest OS still runs a network driver and TCP stack.

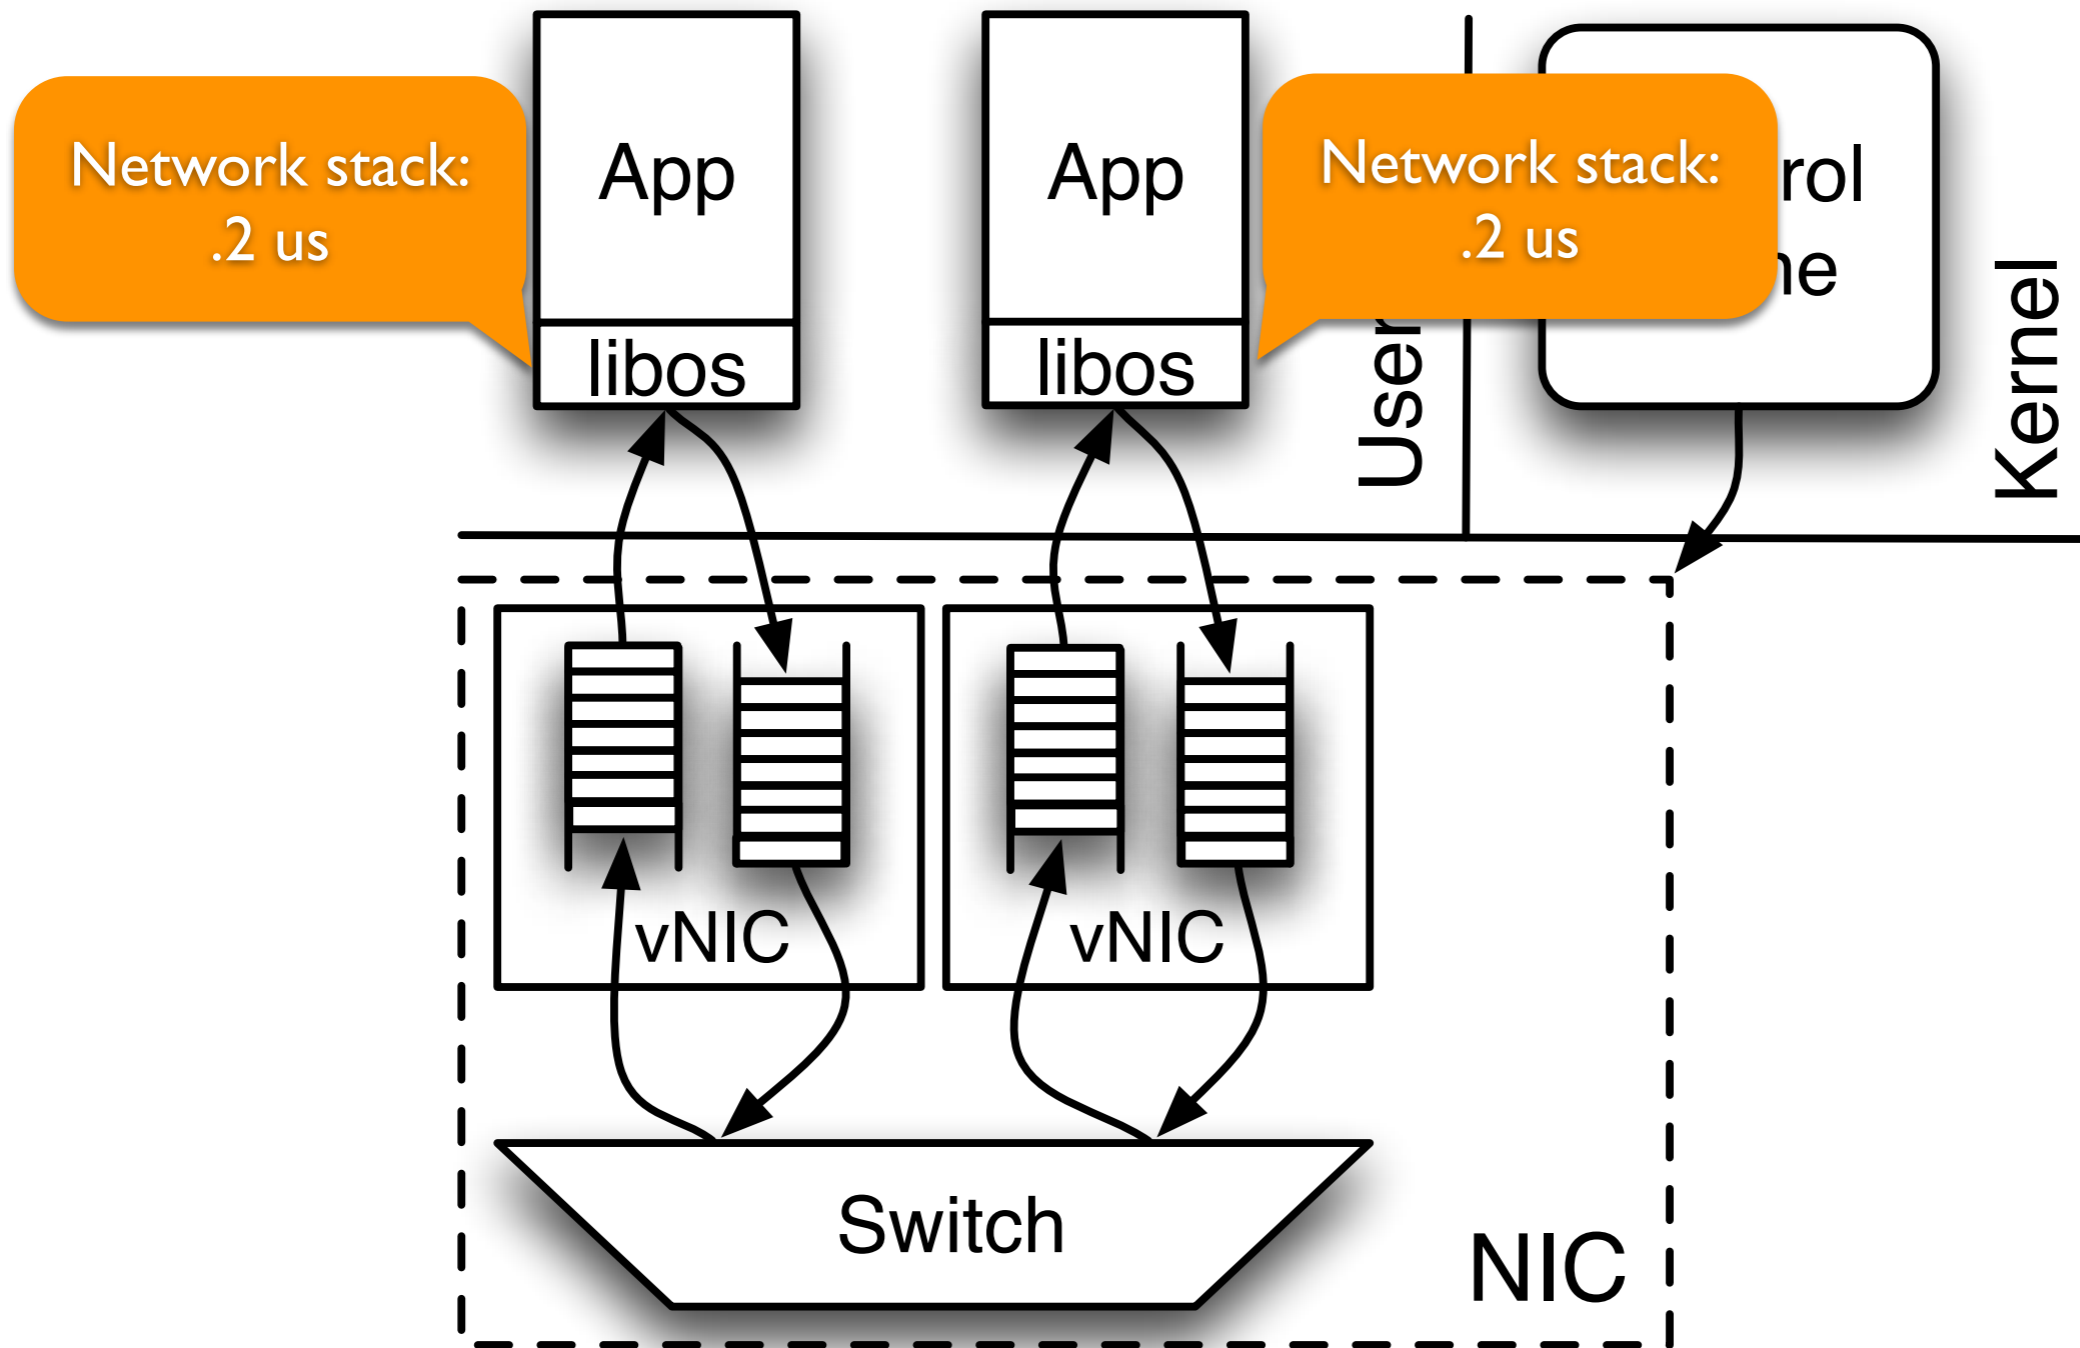
Hardware Virtualization

- Network card directly exposes a number of vNICs.
- VMM allocates one vNIC to each VM.
- Guest OS directly talks to networking card without VMM interposition on each packet.

Arrakis OS

- OS directly exposes vNICs to applications.
- Each application runs a tiny user-level TCP stack with no multiplexing.
- No intervention from the OS (kernel crossings/copying) for each network packet.

Arrakis



Some more numbers ...

Time through 100 ft of cable: **100 ns**

Time through a 10G switch: **350 ns**

Time for memcached to respond: **1 us**

Time for OS to process a packet: **6 us**

Time for Arrakis to process a packet: **200 ns**

Interested?

Come work with us!

Or learn more:

<http://arrakis.cs.washington.edu>