### Modern Network Hardware Akshay Narayan, MIT CSAIL

FreeFlow: Software-based Virtual RDMA Networking for **Containerized Clouds** 

Direct Universal Access: Making Data Center Resources Available to FPGA

Stardust: Divide and Conquer in the Data Center Network

Blink: Fast Connectivity Recovery Entirely in the Data Plane

### Papers

## Where is this hardware?

FreeFlow: Software-based Virtual RDMA Networking for **Containerized Clouds** 

Direct Universal Access: Making Data Center Resources Available to FPGA

- Stardust: Divide and Conquer in the Data Center Network
- Blink: Fast Connectivity Recovery Entirely in the Data Plane

## What do datacenters look like?

### **Highlights for emphasis**

Sleeping servers with blankets Mood Lighting



## What do datacenters look like?



### Datacenter Host Hardware FreeFlow: Software-based Virtual RDMA Networking for Containerized Clouds Server 1 **Server N** Μ Μ Μ IOH IOH PCle PCle NIC NIC CORE NETWORK **DATACENTER NETWORK** / INTERNET



FreeFlow: Software-based Virtual RDMA Networking for Containerized Clouds

# What is RDMA?

### "Remote, Direct Memory Access"





NIC Image from @graziaprato via Justine Sherry

FreeFlow: Software-based Virtual RDMA Networking for Containerized Clouds

## Ok, so "Virtual RDMA for Containerized Clouds"?

"Container"

How can we let containers use RDMA?

"Virtual RDMA Networking"

compute isolation + filesystem isolation + network interface isolation for processes



Software pretending to be a RDMA NIC which talks to the real RDMA NIC

## Datacenter Host Hardware

Direct Universal Access: Making Data Center Resources Available to FPGA



Direct Universal Access: Making Data Center Resources Available to FPGA

# What is an FPGA?



"Temporal Computing"

Small amounts of data at a time

Switch to doing different things quickly

### "Field Programmable Gate Array"



"Spatial Computing"

Large amounts of data at a time

Switch operations slowly

Direct Universal Access: Making Data Center Resources Available to FPGA





Stardust: Divide and Conquer in the Data Center Network

## Datacenter Network Hardware



Stardust: Divide and Conquer in the Data Center Network

# The Dream: One Big Switch



### Conventional design: fat-tree network

Image from M. Al-Fares et al., SIGCOMM 2008



What users want: "One Big Switch"

Image from S. Agarwal et al., SIGCOMM 2018



### Stardust: Divide and Conquer in the Data Center Network

# Big Switches





## Internet Network Hardware



### Blink: Fast Connectivity Recovery Entirely in the Data Plane



Image adapted from S. Ratnasamy, UC Berkeley CS 168



### Blink: Fast Connectivity Recovery Entirely in the Data Plane



Image adapted from S. Ratnasamy, UC Berkeley CS 168



Blink: Fast Connectivity Recovery Entirely in the Data Plane





