# Multitenancy for Fast and Programmable Networks in the Cloud

Tao Wang\*, Hang Zhu\*, Fabian Ruffy, Xin Jin, Anirudh Sivaraman, Dan Ports, and Aurojit Panda

NEW YORK UNIVERSITY, JOHNS HOPKINS UNIVERSITY, Microsoft Research

by: Group 5 Jiawei Yao, Jiayu Li, Xiaowan Guo





# Emergence of programmable network devices

- Pipeline-based programmable devices
  - $\rightarrow$ In-network switches
  - $\rightarrow$ At-host SmartNICs
- Enable wide-range innovations for classical networked systems →Consensus: NOPaxos, NetPaxos
  - →Caching: NetCache, IncBricks
  - →Storage: NetChain, SwitchKV









# A hybrid compile-time and run-time solution

### ➤ Compile-time program linker

- Target generic resources (e.g., SRAMs/TCAMs, action units, etc.)
- But static

### ➤ Run-time memory allocator

- Target stateful memory
- But limited



# Goals of compile-time linker

- Restrict resource usage
- Provide isolation
  - Ensure tenant program does not inference with others'
  - Ensure no infinite packet resubmitting
  - Ensure no loop forwarding configuration















# **Run-time memory allocator efficiency**

### ≻Experimental Setting

≻64 tenants submit 1-min heavy hitter detection task against source IP address within its /6 subnets

≻10-min CAIDA trace replay

### ≻Evaluation metric

- ≻Utility: memory hit ratio
- ≻Satisfaction: time fraction w/ utility > 0.9
- ≻We show the mean and 5th percentile



# Conclusion

- ≻A hybrid solution for multi-tenancy support
- ≻Compile-time linker: general but static
- ≻Run-time memory allocator: dynamic but limited

UNIVERSITY of WASHINGTON 21

# **Critique: Strengths**

- Resource efficiency
  - ➤ Overhead is negligible.
- Isolation
  - > Apply in the context where each tenant runs its own program and is distrustful of others tenants
- Targets pipeline-based ASICs and multitenant
- ♦ A hybrid solution
  - > Consider dynamics and developed run-time memory reallocation mechanism

### **Critique: Weaknesses**

- Only considered a model where consumers submit entire programs to run on programmable devices.
  - higher-level API? simpler, less device-specific programming? lightweight isolation, e.g. paravirtualization?
- As a device OS, failed to provide other common services for applications
  - > such as, abstractions for scalable and fault-tolerant storage
- In the evaluation part, only 3 testcases performance were displayed.
  > Also, lacked comparisons with other systems and architectures



# <section-header><list-item><list-item><list-item><list-item><list-item><list-item><list-item>

## **Identify GAPS**

### ★ Future Work

- $\circ~$  Seek new hardware design, both  $\ensuremath{\textbf{general}}$  and  $\ensuremath{\textbf{dynamic.}}$ 
  - For example, HW supports for partial reconfiguration, or dedicated hardware for IP & MAC address translation.
- In the program linker, consolidate among tenants programs during merging to conserve hardware resources.

### ★ Open Problems

 Diversity of resource types makes it hard to determine an efficient and fair policy for partitioning among tenants.
 Tenants may lie about requirements to get higher allocations.



### Reference

- 1. <u>https://www.usenix.org/system/files/hotcloud20\_paper\_wang.pdf</u>
- 2. https://www.usenix.org/system/files/hotcloud20-paper74-slideswang.pdf

