PARD

Programmable Architecture for Resourcing-on-Demand

Outline

- Background
- Overview
- Summary

Background

- Contention
  - trade-off: high utilization v.s. low latency
- Differentiated Service
  - QoS (Quality of Service)
- SDN (Software Defined Network)
  - Seperate data plane and control plane

A computer is inherently a network

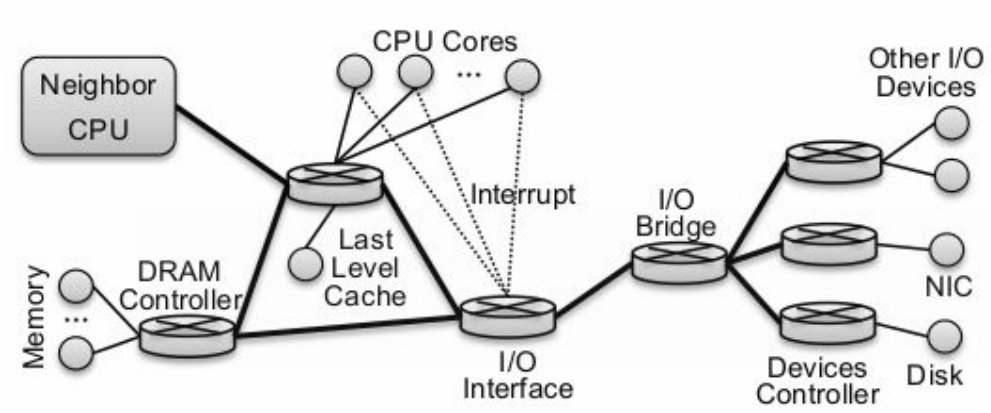

Overview

- Tagging Mechanism
- Programmable Control Plane
- Platform Resource Manager (PRM)
- Programming Methodology

Tagging Mechanism

- DS-id
- DS-id register
- life-time ?
  - transfer through cache hierarchy
- who owns ?
  - CPU
  - I/O device

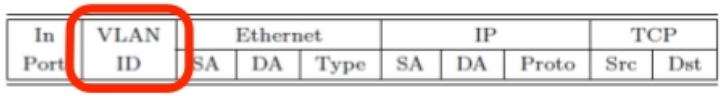

packet format of OpenFlow proxy

Programmable Control Plane

- Fast, inflexible firmware
- Slow, flexible program
- Three DS-id indexed tables
  - parameter table: resource allocation policy
  - statistic table: resource usage information
  - trigger table: event triggers (conditions)

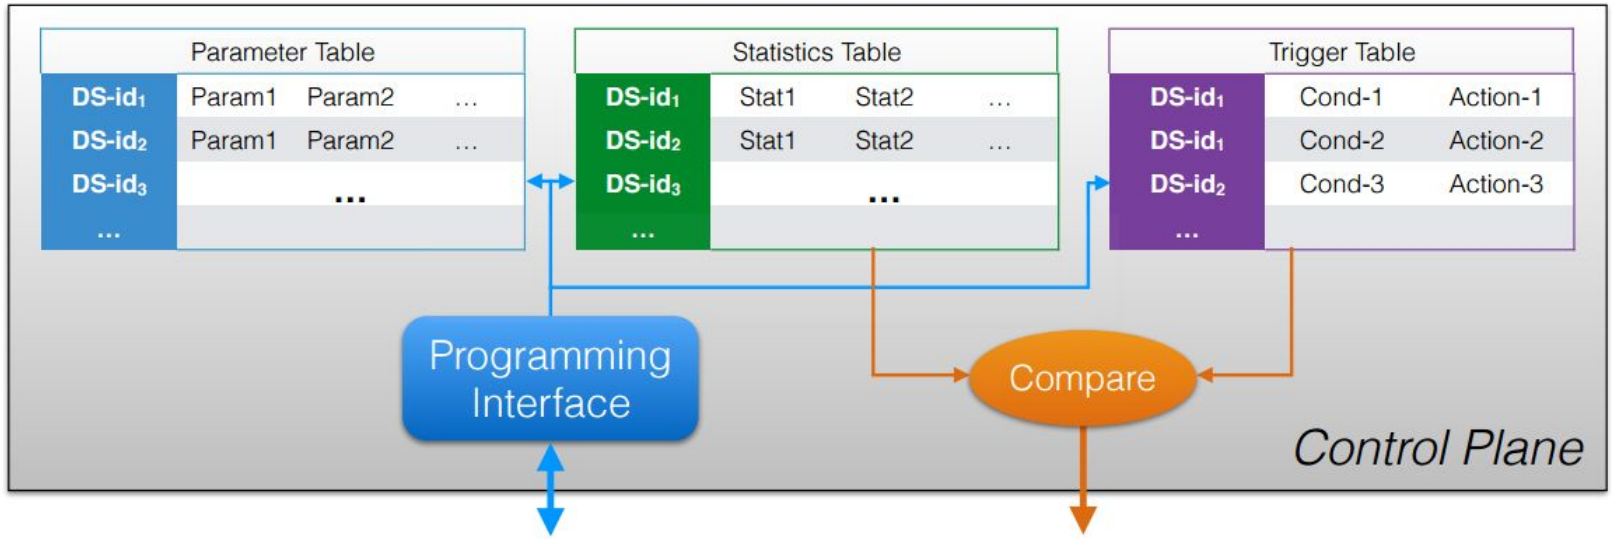
Parameter Table
DS-id1 Param1 Param2 …
DS-id2 Param1 Param2 …
DS-id3 …
…
Statistics Table
DS-id1 Stat1 Stat2 …
DS-id2 Stat1 Stat2 …
DS-id3 …
…
Trigger Table
DS-id1 Cond-1 Action-1
DS-id1 Cond-2 Action-2
DS-id2 Cond-3 Action-3
…
Programming Interface
Compare
Control Plane

Platform Resource Manager

- Centralized
  - connects all control planes and tag registers
- SoC
  - CPU, RAM, flash, bus, Ethernet adaptor, ...
  - Linux-based firmware

Programming Methodology

- Trigger　action
- Action
  - program / scrips
  - file primitives
- File abstraction

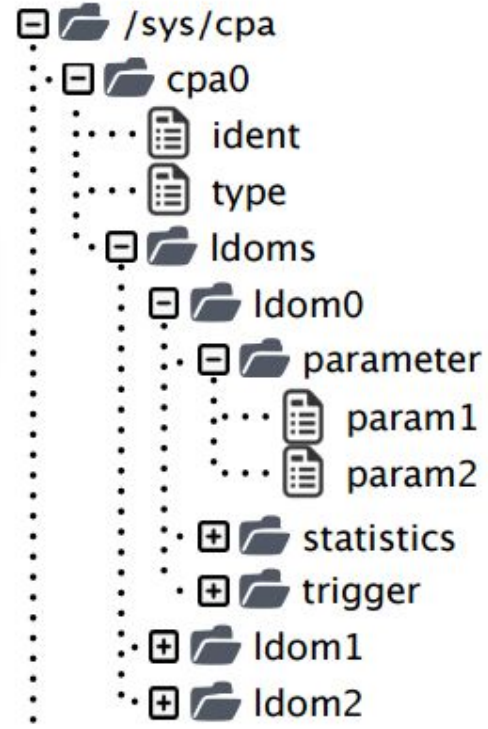

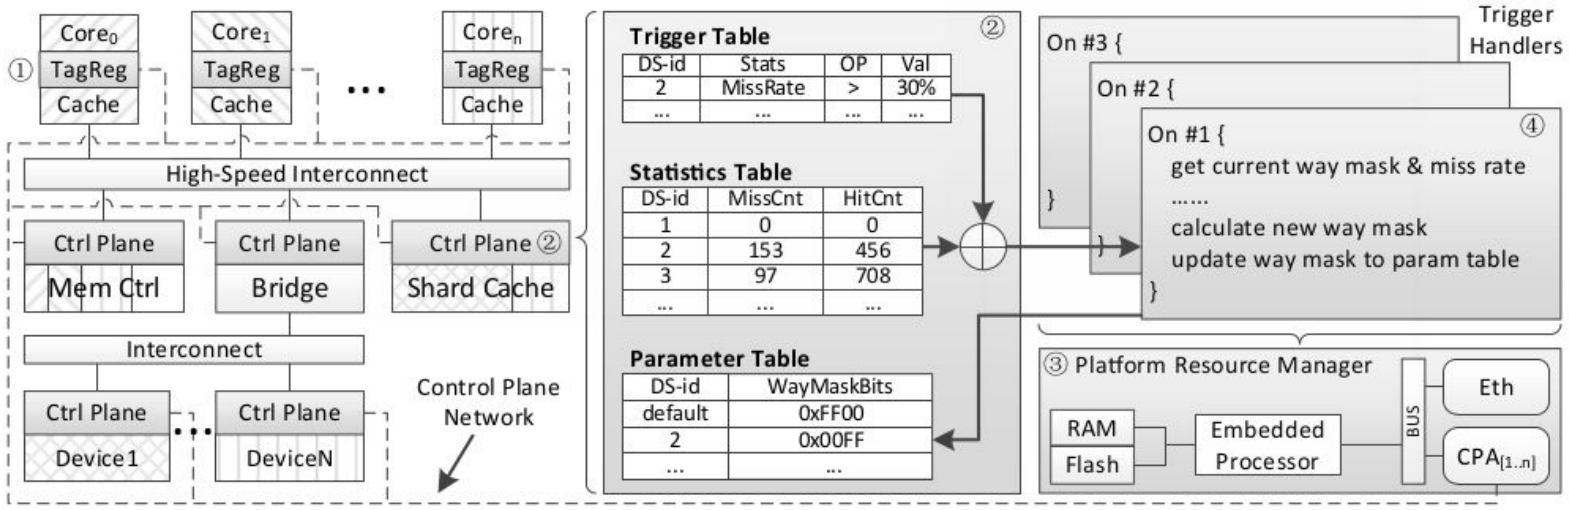
Core0
Core1
Coren
TagReg
Cache
①
High-Speed Interconnect
Ctrl Plane
Mem Ctrl
Bridge
Shard Cache
②
Interconnect
Device1
DeviceN
Control Plane
Network
Trigger Table
DS-id Stats OP Val
2 MissRate > 30%
Statistics Table
DS-id MissCnt HitCnt
1 0 0
2 153 456
3 97 708
Parameter Table
DS-id WayMaskBits
default 0xFF00
2 0x00FF
Trigger
Handlers
On #3 {
On #2 {
On #1 {
④
get current way mask & miss rate
……
calculate new way mask
update way mask to param table
③ Platform Resource Manager
RAM
Flash
Embedded
Processor
BUS
Eth
CPA[1..n]

Summary

- Convey QoS requirements
  - DS-id
  - DS-id indexed table
- Uniformed managing interface
  - can manage multipe devices
  - cache partition: only one type of resources

References

- Supporting Differentiated Services in Computers via Programmable Architecture for Resourcing-on-Demand (PARD) (ASPLOS 2015, keynote, paper)