PACKET PROCESSING SCALING

A Flexible Scalable Approach

Puskar Upadhye
Senior Principal Engineer

1. INTRODUCTION / MOTIVATION

Overview of technology

- Compute nodes today comprise of multi-core CPUs
- Cluster of such nodes are popular for solving previously unsolvable problems
- Host Bus Adapter (HBA) generically refer to hardware used to connect nodes to the network

Motivation and problem-solving strategy

- HBAs typically parse/classify/translate address for incoming packets.
- Packets land in resource like queues which are provisioned as a group (context) exclusive to a process on a CPU
- Packet Process Scaling (PPS) enables spreading of packets across context

Connection with Marvell products / cross BUs

- QLogic® FC-HBA products
- Marvell® FastLinQ® series Ethernet controller
- Marvell DPU

2. FRAMEWORK, ARCHITECTURE, TECHNICAL FOUNDATION

- PPS spreads incoming packets across different context
- Does not use hardwired knowledge of the header fields
- A programmable match and indexing is performed over the first N bytes of the packet
- Context ID is generated using modular arithmetic
- Potentially a context can be mapped to one core on a multi-core node

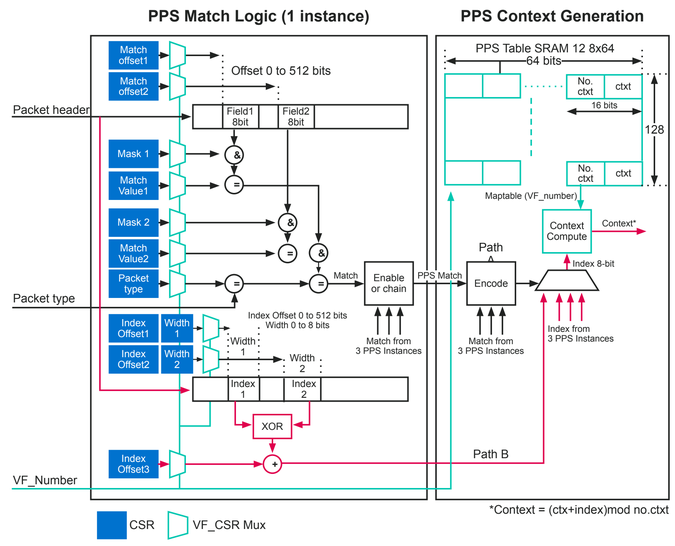

3. DATA FLOW

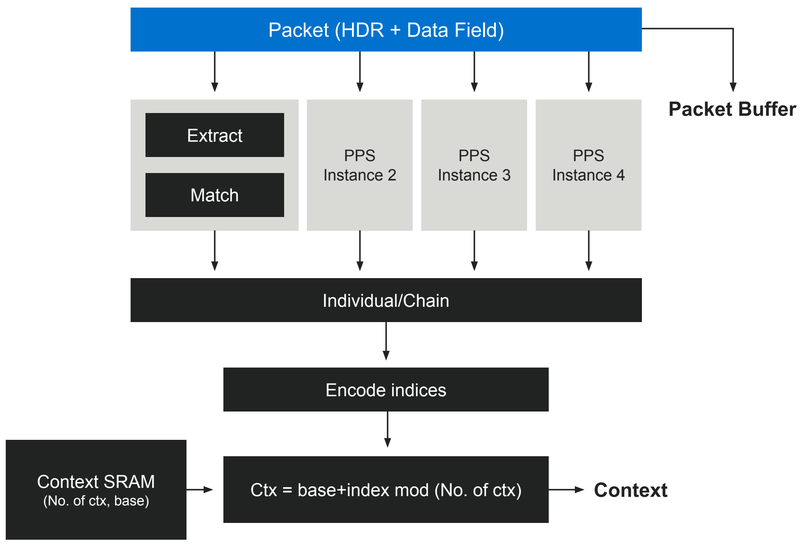

4. SCALABILITY

- Multiple PPS match instances can be added to support more packet types or exhaustive match criteria
- The PPT Table depth can be increased to support multiple virtual functions
- The Configurable State Register (CSR) scale linearly with the virtual functions
- A 16-bit PPS table entry can support up-to 256 contexts (2 socket Sapphire Rapids could go up-to 224)

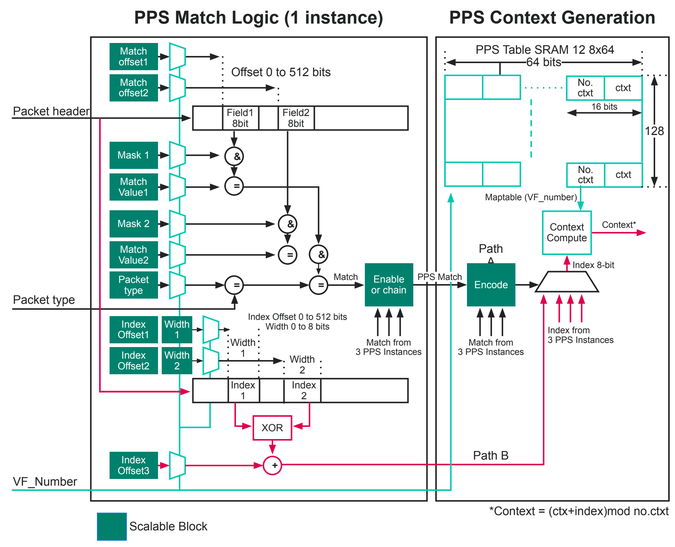

5. USE CASE

How can it be used in product or software application

- PPS in the hardware spreads the packets easily across larger number of context
- As CPU core/thread counts grow, packet distribution scales accordingly
- Applications can thus leverage higher degree of parallelism

Current challenges

- Device software/firmware would need to be enhanced to support the configuring packet filtering rules

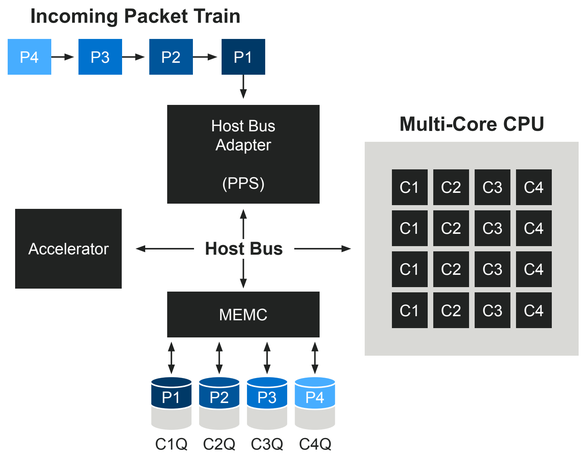

6. CONCLUSIONS, RECOMMENDATIONS, NEXT STEPS

Technology summary

The PPS is a feature designed to spread incoming packets over contexts with no hard-wired knowledge of header fields of the packet. This enables higher utilization of the CPU cores and directly impacts the performance (context per core/thread).

Key advantages

- Silicon area is optimized by using modular arithmetic to generate context value.
- Providing flexibility via CSRs allows for future proofing of the design.

1. PPS block can be plugged into existing design before packets are distributed into queues
2. PPS block is designed to be configurable and scalable for implementing future design with complex packet filtering need