SUNH & SUPERp

***S**cale **U**p **N**etwork **H**eader*
***S**cale **UP** **E**the**R**net **p**rotocol*

Tom Herbert
LPC 2025

XDPnet

SUNH and SUPERp

Network protocols to meet the demands of scale-up networks for AI

- High tput low latency
- HW implementation
- Don’t boil the ocean!

XDPnet

Scale Up Network Header (SUNH)

- Compressed network layer header
- 16 bit addresses (low order bits of host IP address)
- Traffic class, Next hdr, Hop Limit, Flow label

XDPnet

SUNH compared to IPv4

SUNh compared to IPv4, twelve bytes saved

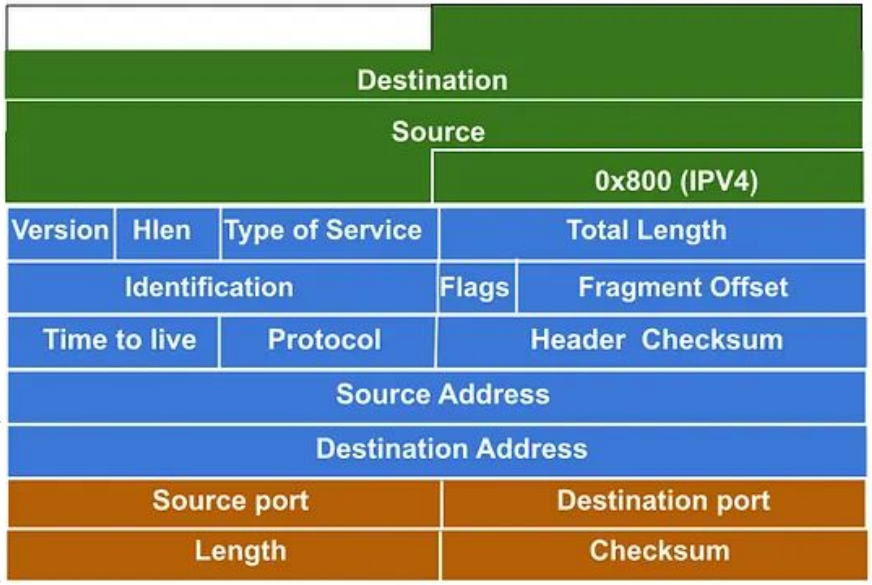

XDPnet

Scale UP Ethernet protocol (SUPERp)

- Transport/ULP protocol for scale-up
- Support memory operations, other HPC ops
- Transaction oriented with initiator/target
- Reliable, flow control, ordered or unordered
- Amenable to SW and HW implementation

XDPnet

Packet Delivery Layer Header (PDL)

- Cnx oriented, reliable communication TCP)
- PSNs with ACKS, SACK bitmap
- RX flow control, order or unordered delivery

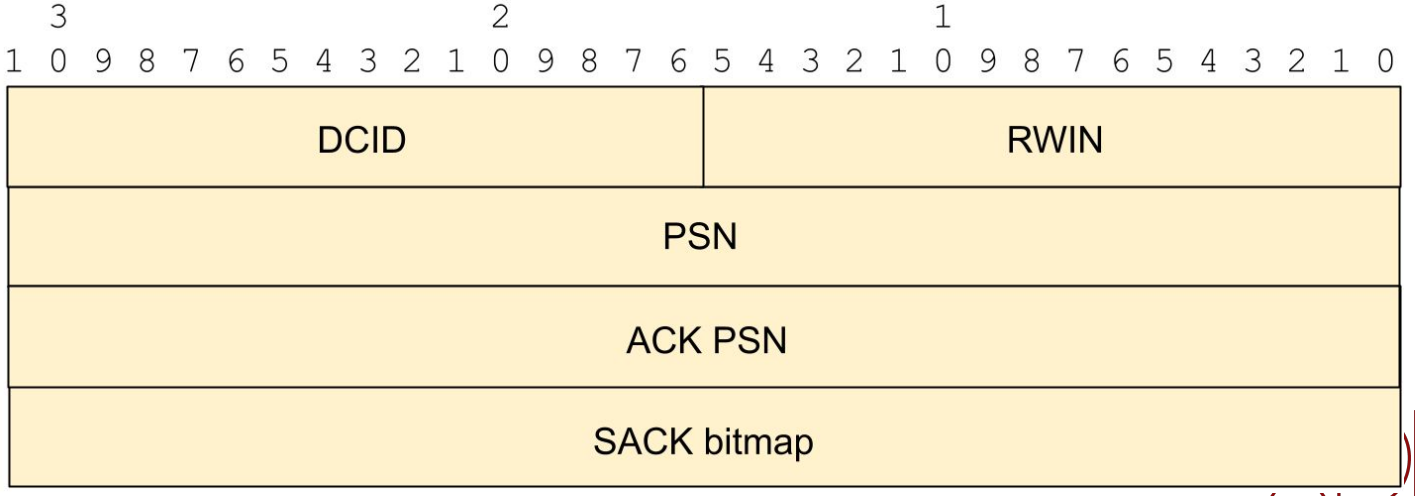

Transaction Layer (TAL)

- Transactions have one or more operations
- XIDs, ACKs, and flow control
- Transaction may span multiple messages

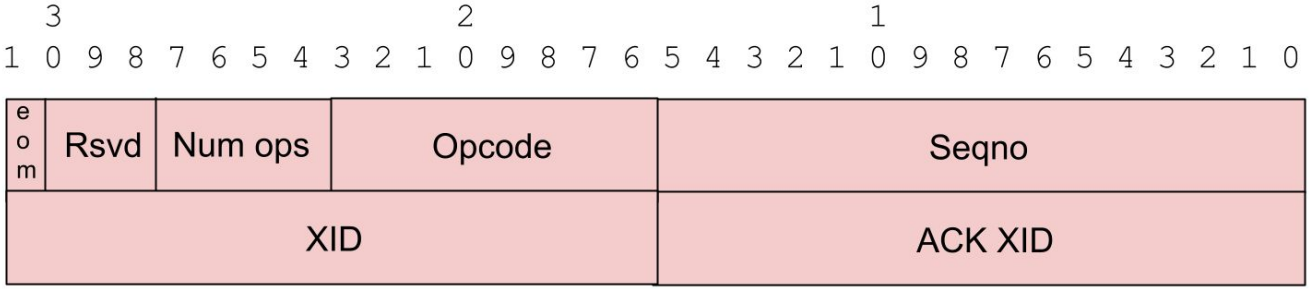

XDPnet

Operations

Read

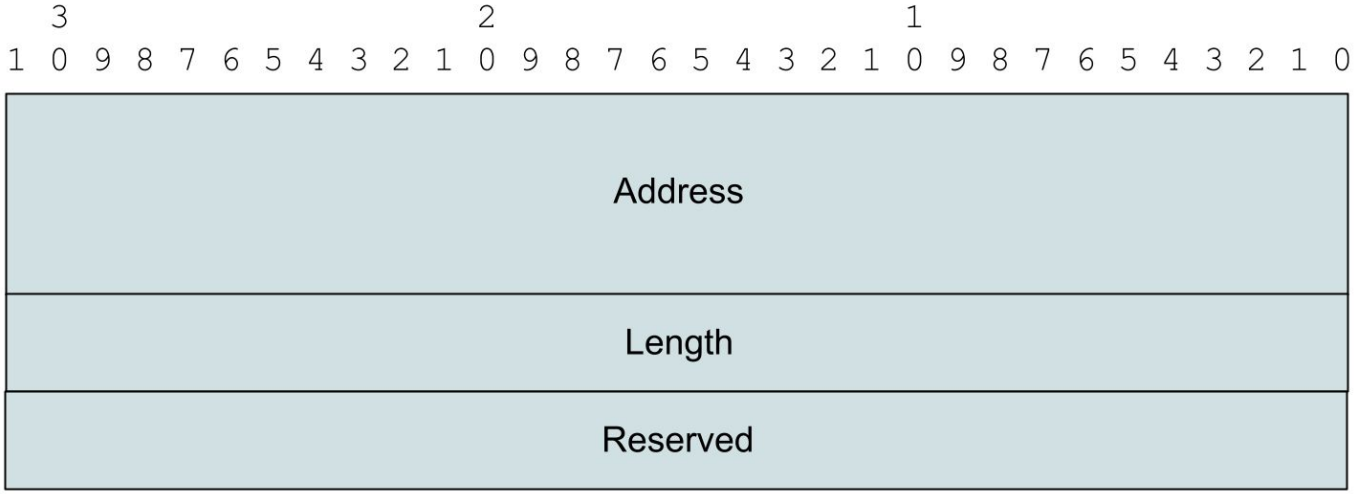

Read response

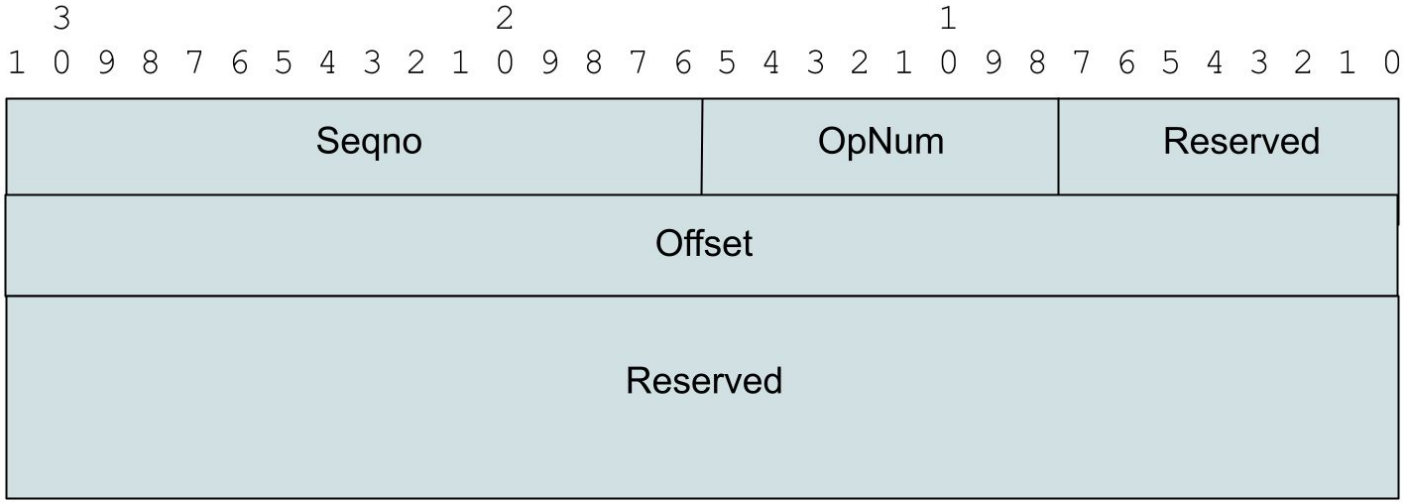

XDPnet

SUPERp packet

Packet operations

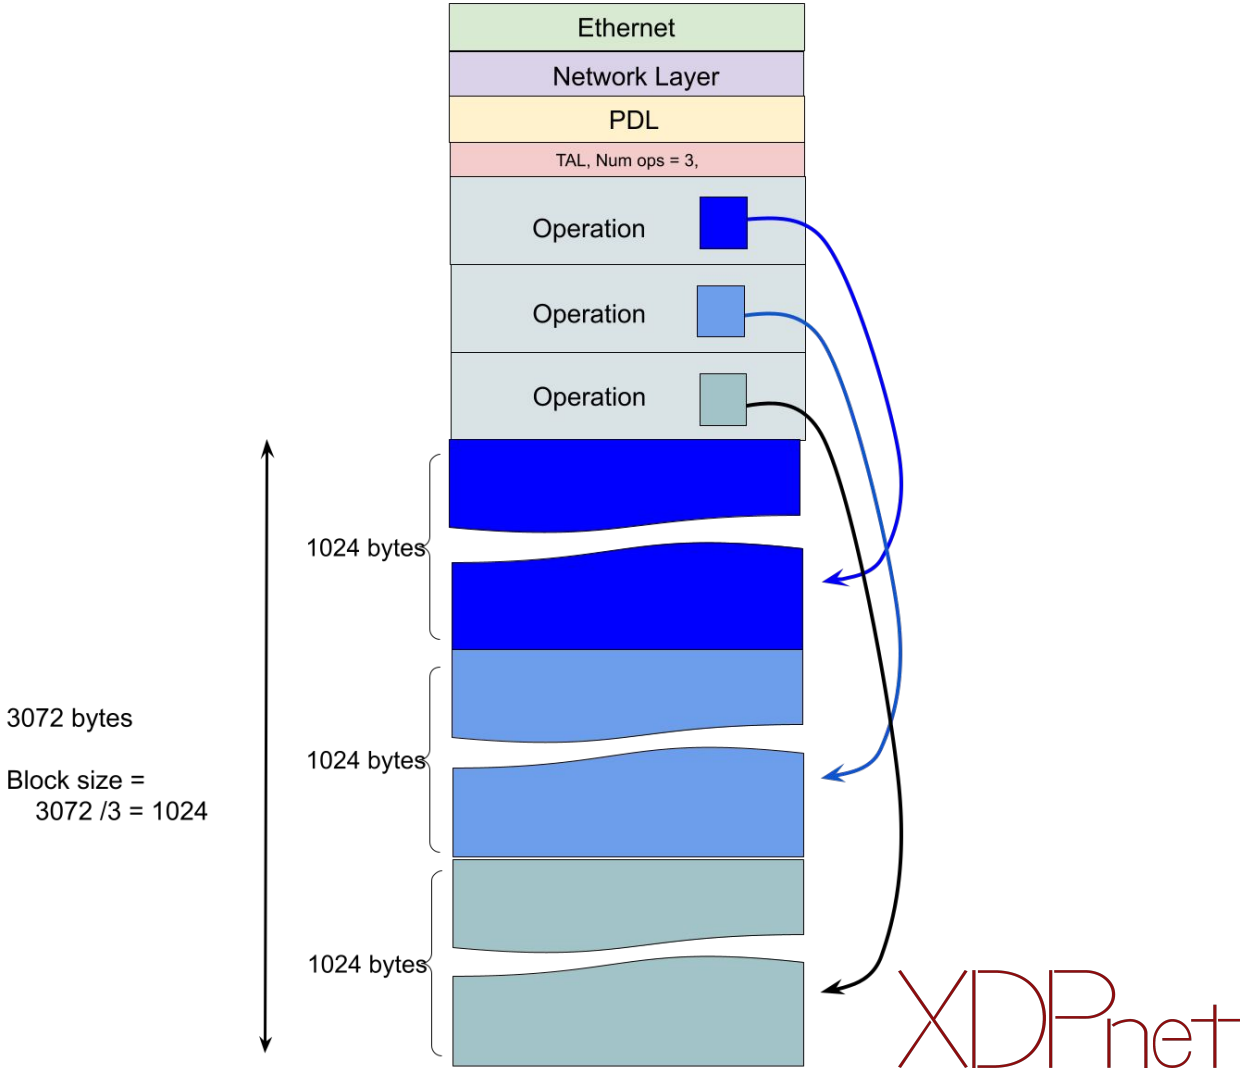

SUPERp protocol layers

Transaction Layer with operations

Packet Delivery Layer

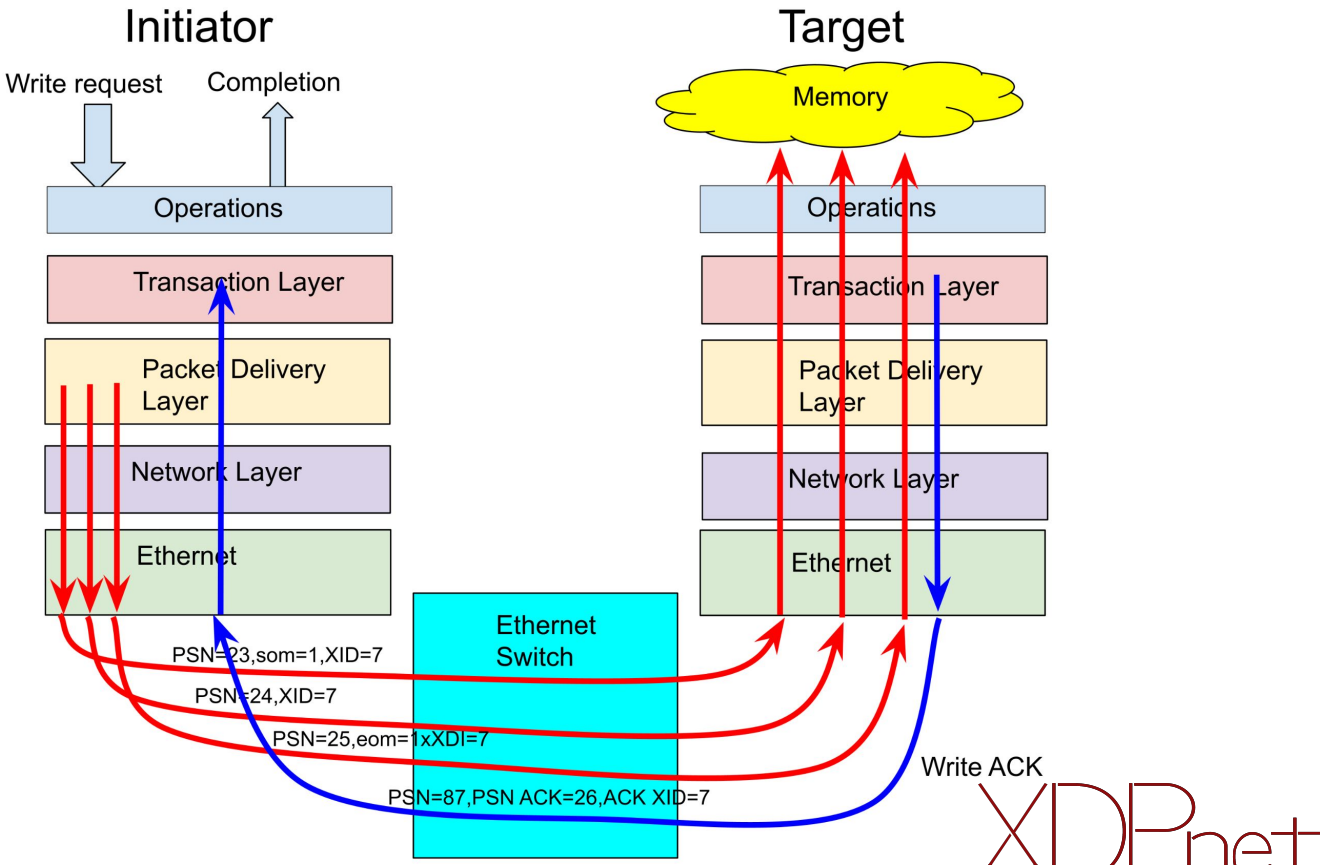

XDPnet

SUNH implementation

- Essentially just another IP header
- SUNH as compressed IP protocol
  - SUNH address is low order bits of IPv4 or IPv6 address
  - netns/interface config gives IPv4 or IPv6 SUNH prefix
  - Transparent to application
- Something like net/sunh

SUPERp

- New transport protocols should be eBPF
- XDP2 allows clean offload to hardware
- PDL an TAL offer modularity
- Hooks for operations
- Offloading SUPERp

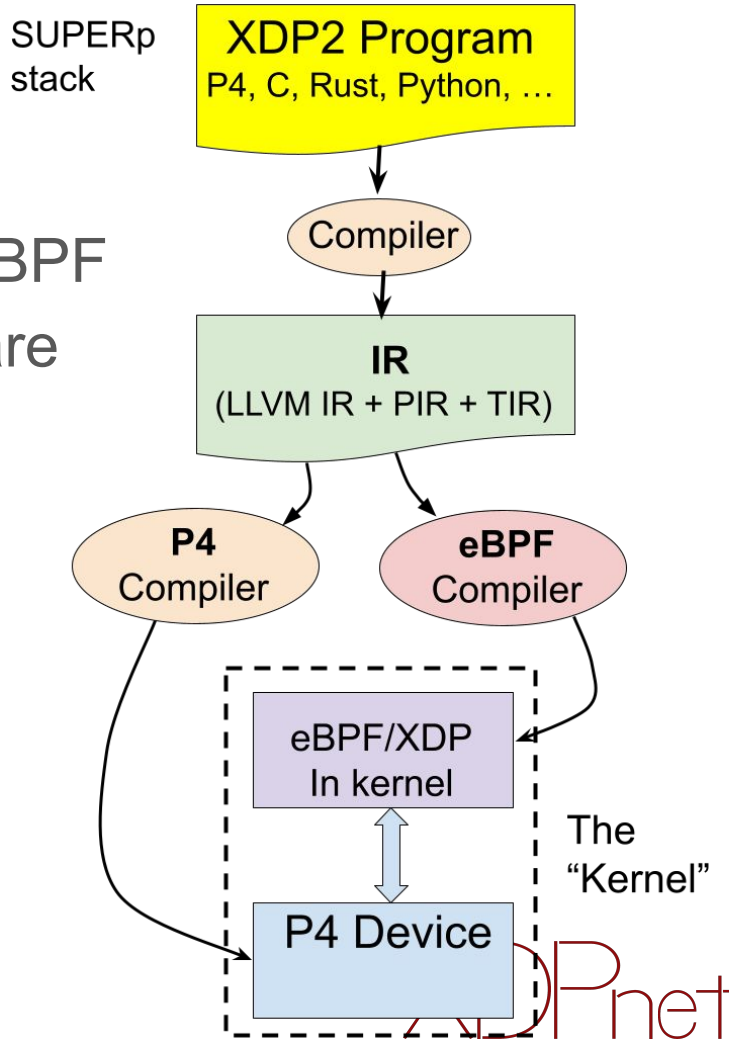

Development

- Specification
  - https://www.ietf.org/archive/id/draft-herbert-sunh-00.txt
  - Will propose to ESUN in OCP
- Protocol formats, packet examples, parser
  - xdp-dev/xdp2 on GitHub
  - e.g. include/sunh and include/superp
- Linux kernel
  - Pending implementation in eBPF (xdp2->eBPF)

XDPnet

Thank you!

XDPnet