Sponsored By

Google intel STORDIS The Open Networking Expert

Towards an Open P4 Programmable Hardware Platform

Gordon Brebner
Fellow
Xilinx Labs

XILINX

Brief NetFPGA history

- Original Stanford boards
  - For teaching (CS344 class) and research (e.g., Ethane → OpenFlow)
  - NetFPGA-v1 (8x10M, 2001-2004)
  - NetFPGA-v2 (4x1G, 2004-2007), then upsized as -v2.1 to become NetFPGA 1G

- Sequence of NetFPGA community boards, carrying Xilinx FPGAs:
  - NetFPGA 1G: Virtex-II Pro (2007) - Stanford lead
  - NetFPGA 10G: Virtex-5 (2010) - Cambridge+Stanford lead
  - NetFPGA SUME: Virtex-7 (2013) - Cambridge lead
- Much onus on the leading university for hands-on work
  - Board design and infrastructure development to create usable platforms to enable the community

- Very successful in terms of world coverage, number of groups, publications

NetFPGA transition: 2010's state → 2020's state

	Today	Future: "NetFPGA 2020"
Open source community	Standalone NetFPGA.org	Alongside P4.org within Linux Foundation and ONF umbrella
P4 focus	Add-on via P4→NetFPGA flow	Integrated P4.org role: as *the* open hardware infrastructure reference
Platform architecture	Limited reference designs	Adaptable platforms supporting P4 standard architectures and APIs, alongside current tools
Boards	Custom NetFPGA boards with (dated) cheap FPGAs	Generic contemporary boards of minimal specification, plus allow for any custom partner-vendor offerings

NetFPGA.org and P4.org

- P4.org is the home of the P4 packet processing language and its ecosystem
- Goal: NetFPGA becomes the standard adaptable hardware node infrastructure for P4.org

- Retain the successful NetFPGA brand identity
- … Placed systematically into a broader context, and aligned with the P4 movement

- Benefits from broader open source community infrastructure

- Makes a step away from being something specialist in the crazy world of FPGAs
- … but hardware aficionados still have a role: platform building, and adding custom externs

P4→NetFPGA workflow

- Implemented by Steve Ibanez (Stanford), available since July 2017

- P4 compiled using Xilinx SDNet …
- … then block implanted into pipeline of NetFPGA SUME switch reference design

- High community interest – around 100 SDNet licenses donated by Xilinx already
  - Used in teaching at Cambridge and Stanford
  - Apparent increase in NetFPGA SUME boards sale driven by its availability

- First practical step towards a future P4-aligned vision for NetFPGA

P4.org trend: Two architecture models

- P4.org has Architecture Working Group

- First product is Portable Switch Architecture (PSA)
  - Currently v1.1 spec: https://p4.org/p4-spec/docs/PSA-v1.1.0.pdf

- New sub-group discussing Portable NIC Architecture (PNA)
  - Second standard architecture, modelling NIC rather than switch
  - Started after multiple expressions of interest within P4 community

- (No-one thinking of yet more models beyond these two canonical models)
  - But some Stanford/Xilinx research on Programmable Target Architecture (PTA)
  - … Modest extension of P4 to allow specification and implementation of a target architecture
  - … Open-ended way of creating one's own architecture, if underlying technology allows it

Common infrastructure for NetFPGA 2020

- Avoid divergent collection of reference designs for different boards
- One standard adaptable platform design each for NIC and switch, open sourced
  - Based on PNA and PSA architecture specs (PNA when available)
  - P4-programmable data plane
  - P4Runtime (and other standard runtime software)
- Minimize customized or specialized aspects in standard base platform
- … Open-ended hooks to allow addition of in-line and lookaside components
- Standard platform and its runtime software becomes the NetFPGA 2020 brand
  - Not specific underlying boards

Xilinx Labs supported base components

Xilinx Vivado IP

Xilinx Labs open IP

Network-side packet interface

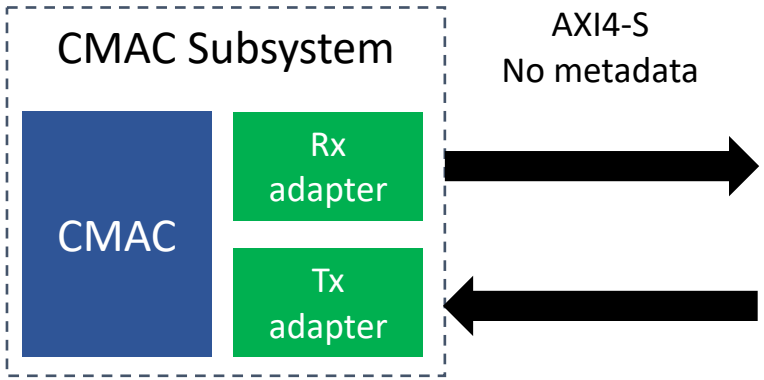

CPU-side packet and control interface

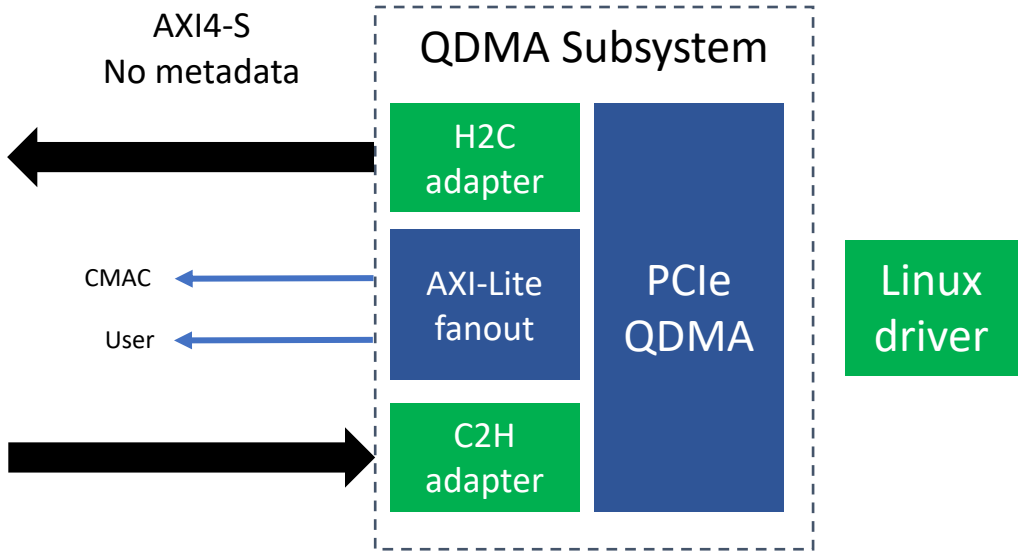

† Possible: Include option for Rx timestamp in metadata

† Future: Include one or more memory interface components

Xilinx Labs supported base shell: One-port NIC-style

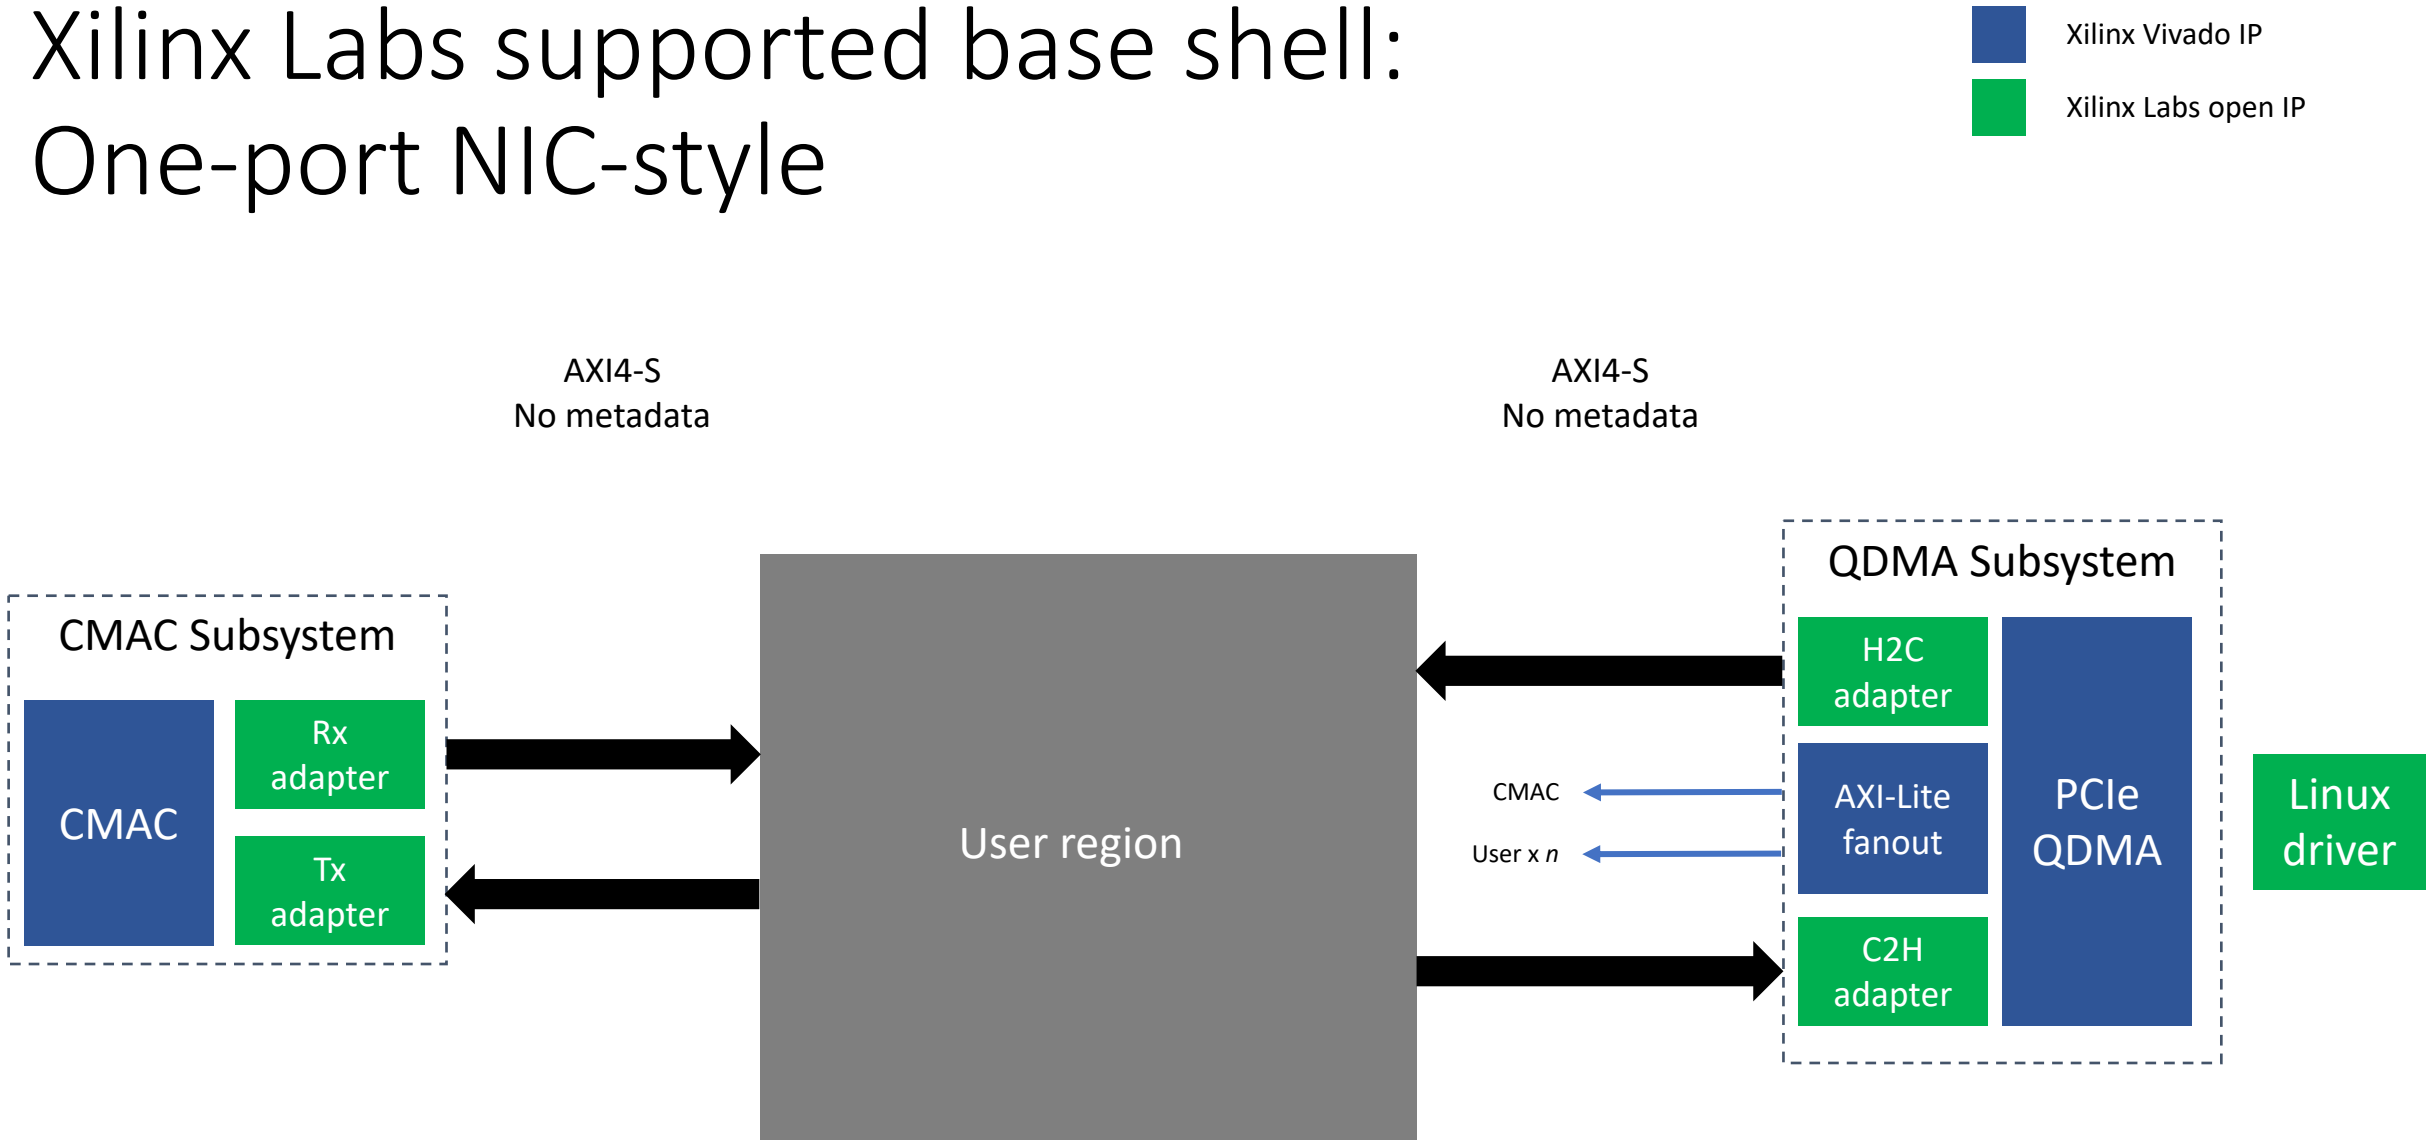

† Future: Include one or more memory interface components

NetFPGA supported base shell: Two-port Switch-style

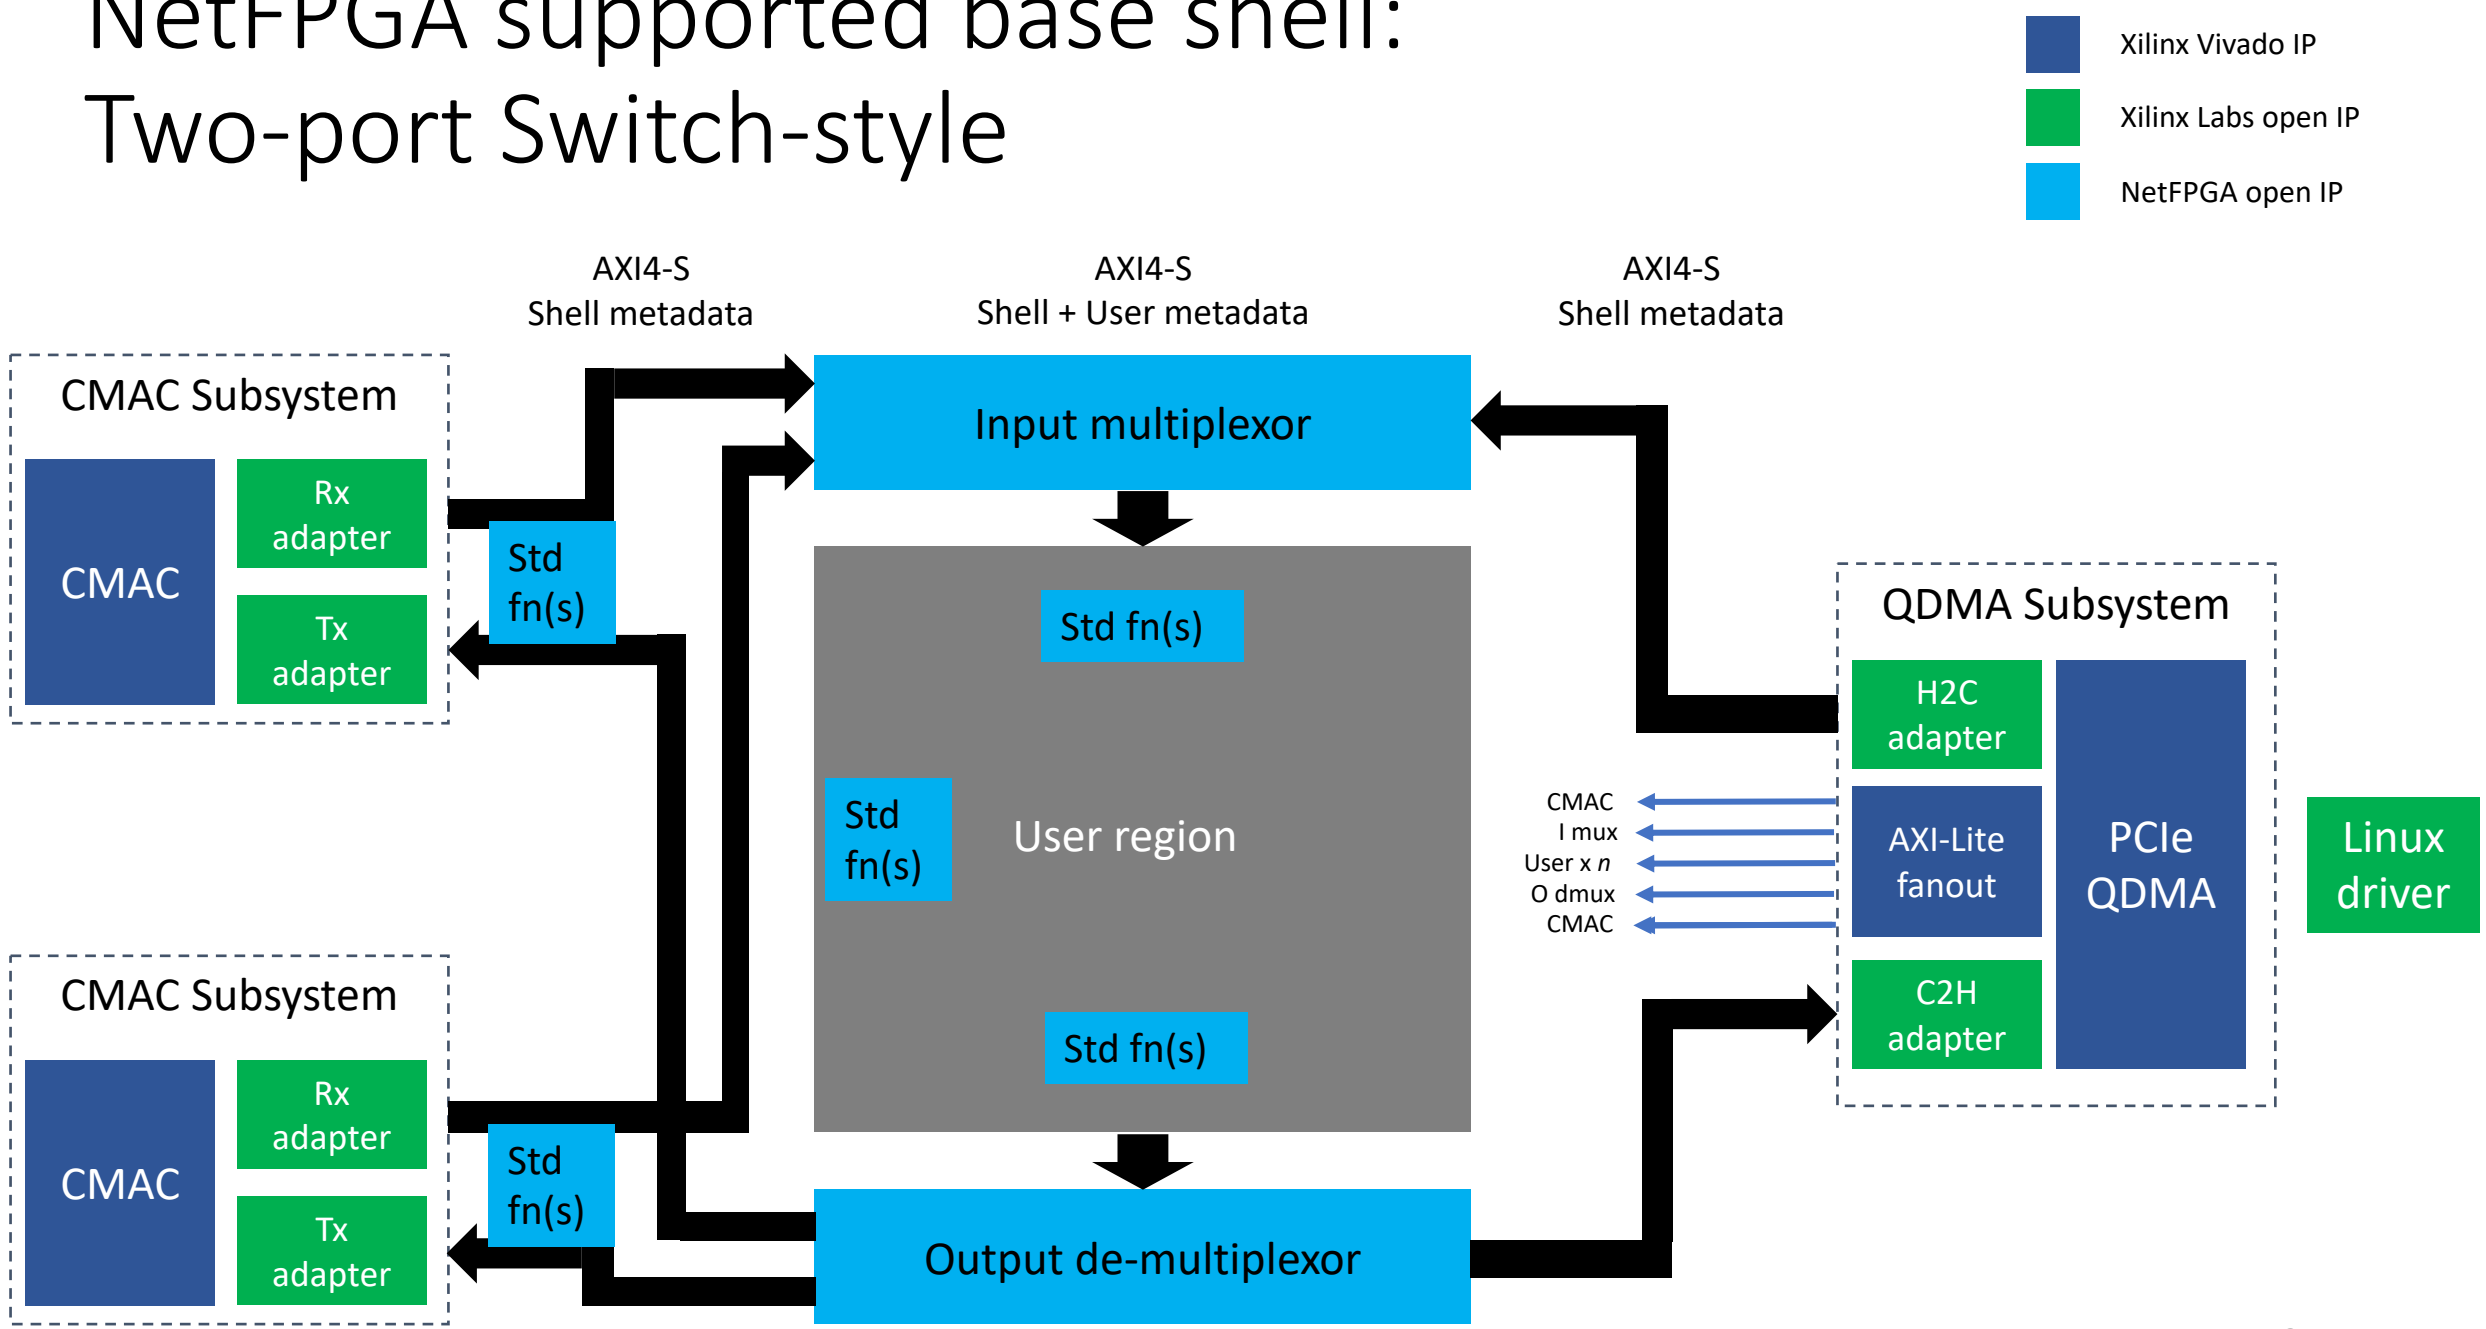

User region examples

- Reference switch (including packet length compute, any width conversions)
- P4-NetFPGA workflow → open P4 PSA-style hardware platform
- [Shell extension] 4x25G or 8x25G switch
- [Shell extension] Programmable traffic management

Current target boards

Platforms designed and parameterized to be easily ported

- Legacy/prototype [To be deprecated later in 2020]
  - Xilinx VCU1525
  - BittWare SOC-250 (MPSoC: FPGA + ARM)

- Mainstream current family of Xilinx Alveo boards
  - Alveo U200, U250, U280

- Future prospects
  - Alveo U50 [Soon]
  - Alveo Versal [Upward port to next Xilinx generation from Alveo Ultrascale+]

Status

- Platform architecture
  - Initial NIC and Switch versions completed during Oct-Dec 2019
  - Stable versions became available Jan-Mar 2020
- Boards (see earlier slide for more detail)
  - Alveo Ultrascale+ portfolio supported, plus some legacy, and some future plans
- P4 focus
  - Updated P4→NetFPGA tool flow completed during Jan-Mar 2020
  - Including brand-new Xilinx P4 compiler, and targeting 100G throughput initially
- Open source community
  - Small set of NetFPGA 2020 early adopters during Jan-Aug 2020 [COVID-19 impact …]
  - Open up general user community from Sep 2020: workshop planned
  - Joint community plan with P4.org to be discussed

Xilinx Labs base shell: Early Access Program (Jan-Aug 2020)

Organization	Lead	Potential open source contributions
Lawrence Berkeley National Lab	Yatish Kumar	• BX NPU array • RISC-V array
Cambridge	Andrew Moore	• NetFPGA switch • HBM • Open tester • Applications
Oxford	Noa Zilberman	
QMU, London	Gianni Antichi	
Stanford	Nick McKeown	• Multi-port NIC • P4 drop-in • Applications
Cornell	Nate Foster	
Yale	Robert Soulé	
Princeton	Jen Rexford	• P4 network
Brown	Theo Benson	• Virtualization
Harvard	Minlan Yu	• Traffic control
Warwick	Suhaib Fahmy	• HLS compute drop-in • Applications
KAUST	Marco Canini	
NUS, Singapore	Bingsheng He	• Alveo cluster
Virginia Tech	Edson Horta	• FPGA+ARM target

Xilinx support for NetFPGA 2020

- Alveo board academic donations or subsidies (or strategic long-term loans)
- *Ad hoc* extra support to community leaders (= Cambridge and Stanford to date)
- Seeding of the standard adaptable platform designs through open sourced components/shells/drivers
- SDNet tool license donations
- IP license donations (for any non-free IP)
- Active participation in P4.org community
- Xilinx Labs research collaborations

Summary of P4-related goals for NetFPGA 2020

- Associate NetFPGA.org more closely with P4.org
  - Position NetFPGA as providing the programmable hardware platform for P4
- Broaden the church from hardware-then-networking users to include networking-then-hardware users
  - P4 as a focus, alongside Verilog as the foundation
- Reduce historic focus on custom boards first, then infrastructure, then community building
  - Support P4 standard NIC/Switch architectures and runtime APIs, implementable on range of boards
- Use latest FPGAs, with simplified future migration to next generations
  - Keep up to date: more logic resource, faster interfaces, use embedded ARM cores

Sponsored By

Google intel STORDIS The Open Networking Expert

Thank You

Contact info for further questions:
Gordon.Brebner@Xilinx.com

XILINX