

# UDP Interface for hls4ml

# Inhalt

## Part I: Introduction

- Setup
- Things to Consider
- Goal

## Part II: HLS - High-Level Synthesis

- Overview
- Comparison to VHDL



## Part III: Ethernet + UDP/IP

- Ethernet PHY
- Networking Workflow with UDP/IP
- Ethernet Frame
- IPv4 Packet
- UDP Packet
- Implementation

# Part I: Introduction

# Setup

Working with



Source: <https://www.realdigital.org/img/16ebb6e013abdd79d4a33733c8ff0fc9.png>

[https://cdn-reichert.de/bilder/web/xxl\\_ws/A300/DIGIL\\_410-292\\_1.png](https://cdn-reichert.de/bilder/web/xxl_ws/A300/DIGIL_410-292_1.png)

Xilinx Z7000 series (PYNQ Z2)  
SoC (System on Chip)

Nexys A7  
„Pure“ FPGA

# Setup

- hls4ml
  - Developed at CERN
  - Mostly uses TensorFlow
- FINN
  - Developed at Xilinx / AMD
  - Uses PyTorch

# Things to Consider

- Prerequisite for Neural Networks on FPGAs: Getting data in and out of the board
- A simple standard interface might be useful: UDP
- Complex logic can be developed faster with HLS – high-level synthesis
- Note: Xilinx SoC chips add capability to load data into the programmable logic via CPU/RAM

# Goal

- Being able to send and receive data via ethernet to neural network applications on FPGAs without relying on SoC-capabilities
- Computer-to-board or board-to-board communication

# Part II: HLS – High-Level Synthesis



# Overview

- Abstraction level above common hardware description languages (VHDL, Verilog)
- Uses C++, C, or SystemC compiled down to VHDL or Verilog
- Enables faster development of complex hardware modules
- Enables automatic implementations of industry-standard module interfaces (AXI4 and other handshake protocols)

# Comparison to VHDL

```
1 library IEEE;
2 use IEEE.STD_LOGIC_1164.ALL;
3 use IEEE.NUMERIC_STD.ALL;
4
5 entity toggler is
6     Generic (
7         FREQ_IN: integer := 50000000
8     );
9     Port (
10        CLK: in std_logic;
11        O: out std_logic
12    );
13 end toggler;
14
15 architecture Behavioral of toggler is
16     signal cnt: unsigned(31 downto 0) := (others => '0');
17     signal toggle: std_logic := '0';
18 begin
19     counter_proc: process(CLK)
20     begin
21         if rising_edge(CLK) then
22             if cnt = FREQ_IN then
23                 cnt <= (others => '0');
24                 toggle <= not toggle;
25             else
26                 cnt <= cnt + 1;
27             end if;
28         end if;
29     end process;
30
31     O <= toggle;
32 end architecture;
```



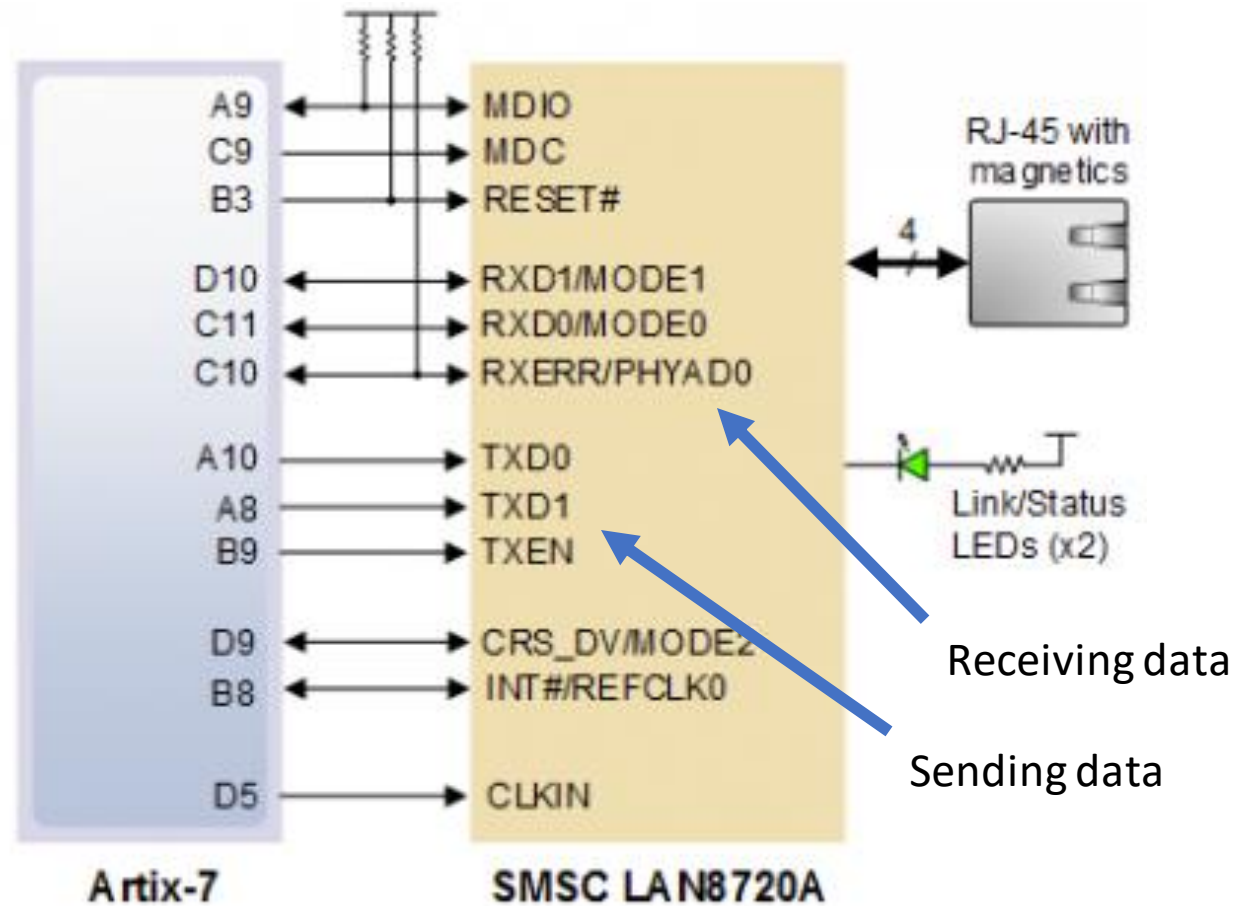
```
1 // toggler.hpp
2 #include <ap_int.h>
3
4 void toggler(ap_uint<1> &o);
5
6
7 // toggler.cpp
8 #include "toggler.hpp"
9
10 void toggler(ap_uint<1> &o) {
11     static ap_uint<32> cnt = 0;
12
13     if (cnt == 50000000) {
14         cnt = 0;
15         o = !o;
16     } else {
17         cnt++;
18     }
19 }
```

# Part III: Ethernet + UDP/IP

# Ethernet PHY – MAC Interface


Ethernet PHY is responsible for

- Encoding/Decoding
- Synchronization of clock signal
- Negotiation of speed and other settings with the communication partner




# Networking Workflow with UDP/IP

## Sending



## Receiving



# Ethernet Frame

Offset	Byte	0								1								2								3							
Byte	Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	Destination MAC Address																															
4	32	Destination MAC Address																Source MAC Address															
8	64	Source MAC Address																															
12	96	Ethertype																Payload (46 – 1500 bytes)															
16	128	Payload (46 – 1500 bytes)																															
?	?	Frame Check Sequence																															

# Ethernet Frame

- Uses 48-bit MAC addresses
- Ethertype: Protocol used within payload data (IPv4, IPv6, ARP, ... )
- Frame Check Sequence: Checksum appended to the end of the frame; CRC32 algorithm (polynomial division of data by predefined polynomial)
- Additionally, each Ethernet Frame transmission is encapsulated with a preamble (0xAA 0xAA 0xAA 0xAA 0xAA 0xAA 0xAA) and the start of frame delimiter (0xAB) at the front and the interpacket gap (idle time between frames) at the back.

# IPv4 Packet

Offset	Byte	0								1								2								3							
Byte	Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	Version				IHL				DSCP				ECN				Total Length															
4	32	Identification																Flags				Fragment Offset											
8	64	Time to Live								Protocol								Header Checksum															
12	96	Source IP Address																															
16	128	Destination IP Address																															
20	160	Options (Used if IHL > 5)																															
24	192																																
28	224																																
32	256																																
36	288																																



# IPv4 Packet

- Uses 32-bit IP addresses
- Protocol: Protocol used within payload data (UDP, TCP, ICMP, ... )
- Header Checksum: Checksum of the header data; simple ones complement addition of 16-bit words

# UDP Packet

Offset	Byte	0								1								2								3							
Byte	Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	Source Port																Destination Port															
4	32	Length																Checksum															
8	64	Payload (18 – 1472 bytes)																															

# UDP Packet

- Uses 16-bit ports
- Checksum: Checksum of the packet data; simple ones complement addition of 16-bit words
- May contain packets of other subprotocols in its payload (DHCP, DNS, ... )

# Implementation

- Simple UDP echo server
- Connected directly via ethernet to a computer
- Code available at <https://github.com/lp247/FPGA-Networking>
- Limitations:
  - General network connections not possible, more protocol implementations like ARP necessary

Thank you for your attention!